ALMA MATER STUDIORUM
UNIVERSITÀ DI BOLOGNA

DOTTORATO DI RICERCA IN

ECONOMICS

Ciclo 36

Gruppo Scientifico Disciplinare: 13/ECON-01 - ECONOMIA POLITICA

Settore Scientifico Disciplinare: ECON-01/A - ECONOMIA POLITICA

THREE ESSAYS IN EMPIRICAL MACROECONOMICS

Presentata da: Matteo Santi

Coordinatore Dottorato

Andrea Mattozzi

Supervisore

Matteo Barigozzi

Esame finale anno 2025

Abstract

This thesis presents three empirical contributions on relevant topics for monetary policy in the Euro Area. In the first chapter, I study macroeconomic tail risks in the EA by constructing conditional distributions of expected GDP growth and inflation and estimating the GDP-at-risk and the Inflation-at-risk. The time variance of the tails differs considerably, and asymmetries in the balance of risks are accompanied by a weak synchrony of GDP-at-risk estimates between member states, especially during crises; upside risks to inflation appear more synchronous. Results have mixed implications for monetary policy, as the weak synchrony of business cycles during crises indicates a possible failure of a criterion for optimal currency areas; on the contrary, the coordination of upside risks to prices is an enhancing factor for a common monetary policy. In the second chapter, I investigate the evolution of the natural rate (r^*) in the EA in the last twenty-five years. The estimation of r^* is performed by extracting the trend component of real rates from the risk-free yield curve, finding a marked drop around the Great Financial Crisis and a partial recovery after 2021. The r^* trend is decomposed into two components, capturing "quantity" imbalances between savings and investments and "quality" effects concerning the availability of safe assets, finding the former to be main driver of r^*. In the third chapter, I study the heterogeneity of monetary policy transmission to financial markets, focusing on the reaction of high-ESG stocks in the EA. The analysis combines a study at stock level and one on portfolios constructed by assuming different degrees of ESG-awareness. Results hint at a stronger sensitivity of high-ESG stocks and portfolios to MP shocks, particularly for securities issued by financial companies, and suggest that the price of these assets can be particularly sensitive to unexpected events in the financial markets.

Contents

Chapter 1

A High-Dimensional GDP-at-Risk and Inflation-at-Risk for the Euro Area

Adapted from tools originally developed in the financial risk management literature, the GDP-at-risk and the Inflation-at-risk are standard measures of tail risk in modern macroeconometrics. In this paper, these indicators are estimated for the Euro Area and its member states by leveraging a high-dimensional dataset in the construction of time-varying conditional distributions of GDP growth and inflation. The distributions obtained at country level are used to assess how the synchrony of the EA countries' business cycles has evolved since the introduction of the Euro. Results indicate significant asymmetries in the balance of upside and downside risks for both GDP and inflation, and a persistently weak synchrony of the left tails of the GDP growth distributions during episodes of crisis.

1.1 Introduction

Expectations on the future state of the economy carry with them a sizeable degree of uncertainty, stemming from the variety of factors driving economic variables and the complexity of their interactions. A common approach to deal with this uncertainty is to complement point estimates of future realizations of economic outcomes with probability distributions of the variables of interest. Forecasts of GDP growth and inflation are often presented in the form of probability distributions, conditional on a series of financial and economic indicators. The tails of these distributions, capturing outcomes in worst-case scenarios, can be especially informative and under scrutiny in turning points of the business cycle. A vast literature (Cecchetti (2006), Adrian, Boyarchenko, and Giannone (2019), Aikman et al. (2019) among others) and the recurrence of periods of macroeconomic stress have drawn policymakers' attention on these rare but harmful events, adapting modelling approaches borrowed from the financial risk management literature as the value-at-risk. In this paper, I estimate conditional distributions of GDP growth and inflation and construct measures of their tail risks for the Euro Area and its member countries. Such distributions are based on factors extracted from a high-dimensional dataset through a partial quantile regression approach (PQR, Giglio et al. (2016)). The advantage of this approach lies in the possibility to exploit a large information set, while avoiding at the same time making strong assumptions on the choice of the indicators in the conditioning set of the distribution. Once the conditional distributions for the EA member countries are estimated, I study the evolution of their synchrony since the introduction of the single currency, focusing on the tails. The synchrony of business cycles is a relevant feature in the optimal currency areas theory (OCAs, Mundell (1961), Krugman (2013)), particularly for the reduction of losses incurring from one-size-fits-all policies. In fact, the infrequency of asymmetric shocks, together with a uniform and ordered transmission to the real economy throughout the whole currency area is conducive to synchronized business cycles and, ultimately, to an easier conduct of monetary policy. Although the original OCAs theory assigns equal importance to the synchronization of growth and crisis periods, the history of the Economic and Monetary Union has shown that the main progress in the integration process was made during the latter. Given the relevance of these episodes for policy advancements, this paper studies explicitly their synchrony in the EA by considering tail risk indicators across its member states.

Results hint at significant asymmetries in the variance of the tails of the GDP distribution both in the EA and in its member states: while the left tail - measuring downside risks to growth - varies significantly in the analysed sample, the right tail remains quite stable over time. Consistently with the literature, the resulting distribution is skewed to the left, as the probability of episodes of severe recessions is higher than those of strong expansions. Moreover, the left tails of the GDP distributions of EA countries tend to become more disperse during crises, indicating that downside risks are often accompanied by risks of divergence among countries. An opposite - but smaller in extent - asymmetry is obtained when considering the inflation distribution, which displays a more volatile right tail, in contrast with a quite stable left part. In contrast to growth downside risks, upside risks to inflation appear more coincident across countries, indicating a good level of synchrony of price increases also during periods of stress in the analysed sample. In principle, this synchrony might be a positive factor for the effectiveness of a common monetary policy, as coordinated (in time and size) shocks to inflation might facilitate a broader consensus in the definition of monetary policies at the EA level. Finally, as a by-product of this analysis, I study the relationship between GDP growth and inflation at different quantiles of their distribution: this relation is estimated to be quite flat in low-growth/low-inflation scenarios, with a gradually increasing slope moving towards the right tail of the GDP conditional distribution.

The paper is structured as follows: Section 1.2 presents the contributions to the literature on the estimation of measures of GDP-at-risk and inflation-at-risk. The dataset and the empirical methodology applied in this work are introduced in Section 1.3. Results for the Euro area as a whole are in Section 1.4, while those for the single countries are in 1.5. Section 1.6 concludes.

1.2 Literature Review

This paper is closely related to the literature on macroeconomic tail risks. The notion of GDP/growth-at-risk (GaR) dates back to the seminal paper of Cecchetti (2006). The approach consists in modelling macroeconomic tail risks by adapting tools developed in asset pricing and financial econometrics with the original objective of quantifying the uncertain losses caused by the randomness of the price of assets in a portfolio. Given the distribution of GDP growth rates, the $\alpha\%$ growth-at-risk

is defined - along the lines of the value-at-risk - as the worst possible outcome (*i.e.* the lowest growth rate) over a specific time horizon, once excluded the worst $\alpha\%$ occurrences. Similarly, the $\alpha\%$ inflation-at-risk represents the highest inflation rate, once the $\alpha\%$ realizations on its right in the distribution are excluded. The paper notably pointed at the non-Gaussianity of GDP growth realizations in a panel of seventeen advanced economies: fat tails and a negative skewness highlight the relevance of rare events, whose probability would be underestimated assuming Normal distributions for GDP growth rates. In a subsequent work (Cecchetti (2008)), the modelling framework is extended by means of a QVAR model and applied to study the non-linear relation between equity booms and macroeconomic variables. The non-linear relationship between financial conditions and GDP growth - theorized also in the literature on intermediary sector's capital and leverage cycles (Adrian and Boyarchenko (2012), Brunnermeier and Sannikov (2014) among others) - is elaborated upon in the analysis of Adrian, Boyarchenko, and Giannone (2019). In this paper, the conditional distribution of growth-at-risk is modelled semi-parametrically, combining quantile regressions (QR) with parametric interpolations. They find for the US a strong relationship between the left tail of the growth rate distribution and the NFCI and highlight a strong asymmetry between the behaviour of conditional right tail, that remains quite stable across time, and the left tail, whose level varies significantly. A series of papers (Aikman et al. (2019) using QR for the UK, Busetti et al. (2021) by expectile regressions for Italy) extend this methodology to obtain distributions conditional on multiple variables. Other works (Hartwig et al. (2021), Giglio et al. (2016)) introduce optimized selection methodologies for the selection of stress indicators to be employed in the analysis of downside risks. The present paper is closely related to a work in this strand of research (Giglio et al. (2016)), with which it shares the methodology for the extraction of latent factors on the quantiles (PQR). Among alternative approaches to the modelling of non-linear relationships between financial vulnerabilities and macroeconomic outcomes topic, Galvão and Owyang (2018) propose factor-augmented models with regime transitions. Growth quantiles are modelled as GARCH processes in Brownlees and Souza (2021), or in a stochastic volatility framework in Gächter et al. (2025). These approaches can allow for predictive power gains, but at the cost of losing some significant insights of the linkages between financial conditions and macroeconomic outcomes. Alternative strategies to gauge this macro-financial nexus include the use of copula methods (Coe and Vahey (2020), on US data), conjectures on data generating processes that can match the observed non-linearity, as in Loria

et al. (2022), who propose a threshold model with state-dependent elasticities of real outcomes to financial conditions, or score driven models for the conditional GDP growth distributions (Delle Monache et al. (2023))

A line of research that is closely related to this work is represented by the studies on downside risks in the Euro area. Trying to close the gap in the advancement of research on this topic between the US and Europe, Figueres and Jarociński (2020) apply the framework of Adrian, Boyarchenko, and Giannone (2019) to European data, taking the region as a whole. Their findings support the hypothesis of a non-linear linkage between GDP growth and financial stress indicators, and select the CISS (Hollo et al. (2012)) as the most informative on the GDP growth left tail. The results are reviewed to a limited extent by an extension with specific focus on the choice of regressors for European GaR (Szendrei and Varga (2023)), carried out by LASSO methodologies. In a similar vein, Alessandri et al. (2019) apply a quantile regression approach to study the interactions between real variables and financial indicators to Italian data, finding quite volatile predictive abilities for downside risks to growth. Other works (Chavleishvili and Kremer (2023), Chavleishvili and Manganelli (2024)) employ QVAR models, or extend the framework to a mixed data sampling setting (Ferrara et al. (2022)). Finally, the paper by Lang et al. (2025) find that spreads and volatility indicators seem as relevant as financial vulnerability when evaluated on the short term, while only the latter seem to display some potential over the medium term. This paper contributes to this literature by estimating measures of tail risk to growth and inflation based on a high-dimensional dataset for the EA, and extending the estimation to member states considered separately, thus allowing for the analysis of business cycle synchrony based on the whole distribution of the variables of interest.

The paper is also connected to the literature on the assessment of OCAs requirements in the Euro Area, with a particular focus on the synchrony of business cycles: as stated by the original theory, the conduct of a common monetary policy for a group of countries is easier in the presence of shocks affecting member states in a symmetric fashion (Mundell (1961)). The concerns raised on the existence of a core-periphery gap in the Euro Area in a set of seminal papers (Bayoumi and Eichengreen (1992), Bayoumi and Eichengreen (1997)) were brought back in the midst

of the Sovereign Debt Crisis (SDC) by Krugman (2013), who pointed at the potentially disruptive effects of sudden stops of capital flows to peripheral countries in the absence of both a Europe-wide backing of banks and of a system of fiscal transfers, paired with unfeasible macroeconomic adjustments through currency devaluation. However, a series of more recent papers (Campos and Macchiarelli (2016), Campos and Macchiarelli (2021) among others) update the original estimates of Bayoumi and Eichengreen (1997), finding evidence of an increased synchrony of business cycles, driven by the adoption of the common currency and the higher competition induced by lower trade barriers in the Euro Area. Other works cast doubts on this convergence: estimating a SVAR identified with a combination of zero and sign restrictions, Kunovac et al. (2022) find a very slow progress in OCAs conditions since the introduction of the Euro, especially on account of periods of asymmetric stress as the SDC. Similarly, de Haan et al. (2024) find some evidence of improved synchrony among member states' output gaps, occurring with a non-monotonic process, susceptible to sudden stops during financial crises. The present paper contributes to this literature by assessing empirically the convergence of economic cycles considered at different quantiles, with a particular focus on tail outcomes. The synchrony of risk drivers and crisis episodes is particularly relevant both for monetary policy and for the possibility to build consensus on institutional reforms at common level, a feature that is connected to another OCAs criterion, the perception of a "common destiny" shared by member states (Baldwin and Wyplosz (2019)).

1.3 Data and Empirical Methodology

The approach followed in this paper involves a two-step estimation of the GDP growth and inflation distributions, conditional on the information extracted from the large dataset of Barigozzi, Lissona, and Tonni (2024). The dataset is described in 1.3.1; Section 1.3.2 presents the estimation procedure.

1.3.1 Data

The data source employed for the extraction of the quantile factors is the EA-MD-QD (Barigozzi, Lissona, and Tonni (2024)), a comprehensive dataset on a large list of economic and financial variables for the Euro Area. The dataset contains a total of 118 series, providing information on national account variables (GDP and its subcomponents), labour market indicators (employment and unem-

ployment in different economic sectors and alternative age groups), financial and credit aggregates (assets and liabilities of public and private entities), labour costs, real exchange rates, interest rates on different maturities, production indices, price indices (PPI and HICP on different categories of goods and services), as well as monetary aggregates and business confidence indicators[1]. In this application, both the dataset at the Euro area level and the national sub-datasets are employed (for Austria, Belgium, Germany, France, Greece, Italy, Netherlands, Spain and Portugal[2]). The extraction of the factors is conducted on observations at quarterly frequency on the 2001Q1-2023Q4 sample, and the construction of conditional distributions is carried out on a four quarters ahead horizon ($h = 4$).

The series included in the dataset are transformed in order to obtain stationarity, and an additional processing is performed on outliers and for real variables during the Covid period, which are treated as missing data and imputed using the EM algorithm, as in McCracken and Ng (2016)[3]. Figures A1.1 and A1.2 report the yearly GDP growth rates for the EA and for the countries analysed in this paper, comparing the imputed series with the original ones. Using this approach, the intensity of the Covid recession and of the following recovery is dampened, allowing for a more robust estimation of the conditional distributions of GDP growth.

1.3.2 Principal Quantile Regression (PQR)

In the PQR model (Dodge and Whittaker (2009), Giglio et al. (2016)), the quantiles of the target variable - GDP growth or inflation h-quarter ahead, in this work - are assumed to be linear combinations of a series of unobserved factors $f_t(\tau)$. Such factors, which condense the most relevant information for the behaviour of the target variable at each quantile τ, are in turn linear combinations of the series contained in the large dataset X_t, whose weights measure the co-variation

[1] A complete list of the series in the dataset is reported in the Appendix.

[2] Country-specific datasets are do not include all the series available at aggregate level, but are equally quite large, as each of them contains around one hundred series.

[3] Outliers are defined according to two criteria: a first group of observations is considered as an outlier on a statistical basis (namely, a data point is treated as missing if its difference from the sample median is larger than ten times the interquartile range), and a second group is treated as such in order to handle the Covid period: real variables - GDP and its sub-components, employment indicators, industrial production indices - are treated as missing values in the 2020-2021 period and replaced by their EM-counterparts, obtained using loadings estimated on pre-Covid data and factors extracted from financial variables in the considered period.

between the single series x_{it} and the quantiles of the target variable y_{t+h}.

$$X_t = \lambda_i'(\tau) f_t(\tau) + u_{it}(\tau) \tag{1.1}$$

In this non-linear factor model, the loadings $\lambda_t(\tau)$, the factors $f_t(\tau)$ and the errors $u_{it}(\tau)$ are all quantile-dependent. The specification allows to combine the possibility - provided by factor models - to summarize the relevant information into a small-dimensional matrix $f_t(\tau)$ with the advantage of capturing the information which is relevant at different quantiles for the target variable $y_{\tau,t+h}$, as in the quantile regression literature (Koenker and Bassett Jr (1978) among others). The extraction of the factors takes place in two steps: in the first one, the target series $y_{\tau,t+h}$ is regressed on each series x_{it} contained in the original dataset through a sequence of univariate quantile regressions which also include a constant term, as in Equation 1.2. The outcome of this step is a vector of quantile-specific weights $\phi_i(\tau)$, whose elements measure the relevance of each series for the quantiles of $y_{\tau,t+h}$.

$$y_{\tau,t+h} = c_\tau + \phi_i(\tau) x_{it} + \varepsilon_{\tau,it} \tag{1.2}$$

Subsequently, given the weights $\hat{\phi}_i$, the realized factors $f_t(\tau)$ are estimated by computing the cross-sectional covariance of $\hat{\phi}_i$ and x_{it}. As their respective weights, also these estimated factors are quantile-specific, and are used in the following stage to model the distribution of the target variable. Factors are therefore weighted averages of the variables in X_t, where weights are given by the slope coefficients $\hat{\phi}_i$. This model differs from the similar Quantile Factor Model (QFM) specification (Chen et al. (2021)) in the choice of the target for the extraction of quantile factors. While factors, in the QFM case, condense the information at the quantiles of the independent variables set, here they are estimated in order to extract the most relevant information for an external target variable y_{t+h} on each quantile[4].

In the second step of the algorithm, the quantiles of the target variable $y_{\tau,t+h}$ are modelled as linear combinations of the factors estimated in the previous phase, as in Equation 1.3. The constant $\alpha(\tau)$ and the $\gamma'(\tau)$ parameters - which capture the relationship between the factors and the

[4]The employed PQR specification and estimation procedure allows for the estimation of single-factor models (with one factor for each quantile). However, applying to X_t a rank minimization approach to select the number of factors in a QFM framework indicates the presence of a single factor on tail quantiles.

target at the different quantiles - are estimated through a time series quantile regression.

$$\mathrm{Q}_\tau(y_{t+h}) = \alpha(\tau) + \gamma(\tau)\hat{f}_t(\tau) + \eta_t(\tau) \tag{1.3}$$

Once estimated both the factors and the parameters, it is possible to construct a distribution for y_{t+h},, conditional on the information extracted from X_t . Conditional quantiles are therefore defined as[5]:

$$\hat{\mathrm{Q}}_\tau(y_{t+h}|X_t) = \hat{\alpha}(\tau) + \hat{\gamma}(\tau)\hat{f}_t(\tau) \tag{1.4}$$

A parametric extension based on the skewed-t distribution The quantiles of Equation 1.4 implicitly characterize a distribution function, as the quantile function corresponds to the inverse cumulative distribution function. This distribution has the advantage of being free of parametric assumptions, but it is defined only for a discrete set of percentiles. However, it is possible to obtain a full conditional distribution by fitting a skewed-t distribution (Azzalini and Capitanio (2003)) on the known quantiles: this flexible specification allows for location and scale shifts (controlled by the parameters μ and σ), as well as for asymmetries and fat tails through the fatness and shape parameters ν and α. The probability density function of the distribution is:

$$f(y;\mu,\sigma,\alpha,\nu) = \frac{2}{\sigma} t\left(\frac{y-\mu}{\sigma};\nu\right) T\left(\alpha\frac{y-\mu}{\sigma}\sqrt{\frac{\nu+1}{\nu+\left(\frac{y-\mu}{\sigma}\right)^2}};\nu+1\right) \tag{1.5}$$

Where $t(\cdot)$ and $T(\cdot)$ are respectively the PDF and the CDF of a Student t-distribution. As this distribution is fit for each quarter in the sample, it is possible to obtain time series of its conditional parameters and its moments[6]. In addition, the parametric distribution allows to estimate measures of tail risk based on the expected value-at-risk conditional on the outcome variable falling in the set of the q% most extreme scenarios. For what concerns GDP, these quantities are denoted as

[5] In this specification, the conditional quantiles of the h-quarter ahead target variable are functions of factors at time t. However, the inclusion of lagged factors does not impact significantly the results of the analysis.

[6] As in Adrian, Boyarchenko, and Giannone (2019), the parameters of the distribution are estimated for each time period by minimizing the squared distance between a series of selected empirical quantiles (the 5^{th}, 25^{th}, 75^{th} and 95^{th} percentiles) of the non-parametric distribution and the corresponding quantiles of the skewed-t distribution.

$$\{\hat{\mu}_{t+h}, \hat{\sigma}_{t+h}, \hat{\alpha}_{t+h}, \hat{\nu}_{t+h}\} = \underset{\mu,\sigma,\alpha,\nu}{\arg\min} \sum_\tau \left(\hat{Q}_{y_{t+h}|X_t}(\tau|X_t) - F^{-1}(\tau;\mu,\sigma,\alpha,\nu)\right)^2$$

Parameters measuring the first (μ) and the third (α) moment are left free to assume any value in $\mathbb{R}$, while the dispersion parameter σ is constrained to lie in the $\mathbb{R}^+$ set and the shape parameter capturing the degrees of freedom ν is a positive real number in $\mathbb{Z}^+$

Expected Shortfall (for the left tail) and Expected Longrise (for the right one):

$$ES_{t+h} = \frac{1}{q}\int_0^q \hat{F}^{-1}_{y_{t+h}|X_t}(\tau|X_t)d\tau \tag{1.6}$$

$$EL_{t+h} = \frac{1}{q}\int_{1-q}^1 \hat{F}^{-1}_{y_{t+h}|X_t}(\tau|X_t)d\tau \tag{1.7}$$

Analogous objects are constructed for the conditional distribution of inflation. In this second case, the left tail represents a measure of expected deflation risks, conditional on price growth lying in the lowest $q\%$ scenarios, while the right one measures risks of high inflation.

1.4 Conditional distributions for the EA

This section presents the results obtained by constructing conditional distributions at the Euro area level: the findings for GDP are in 1.4.1, while those on inflation follow in 1.4.2. Finally, some considerations on the relationship between these distributions are presented in 1.4.3.

1.4.1 GDP-at-risk

Figure 1.1 reports the time-varying conditional distribution of the four-quarter-ahead Euro area GDP growth in the analysed sample. The dark blue line is the median of the distribution, while the lighter blue shaded areas represent a set of relevant percentiles (corresponding to a α set to 5, 10 and 25 percent on both tails of the distribution), and the black line is the actual GDP growth four quarters ahead. The lowest blue line representing the 5^{th} percentile, therefore, is a measure of the GDP-at-risk. The actual growth rate follows quite closely the (lagged) median during before the Great Financial Crisis and in the period between 2014 and 2020, when the yearly growth rate settles around 1-2 percent. On the contrary, the gap between the median and the actual growth rate becomes quite large during episodes of crisis, when the contraction of the conditional median is far less pronounced than the fall of the actual growth rate: this indicates how considering and modelling specifically tail risks is particularly relevant during recessions, when point estimates of future outcomes are particularly unreliable given the higher conditional variance. The last four quarters displayed in the graph are an out-of-sample exercise for future realizations of GDP based on the latest observations available in the sample. The median of the distribution hovers around 0,

suggesting a stagnating baseline scenario, but it is surrounded by a large degree of uncertainty, as the conditional variance of the distribution is quite large especially in the left tail, hinting at the presence of significant downside risks to GDP growth[7].

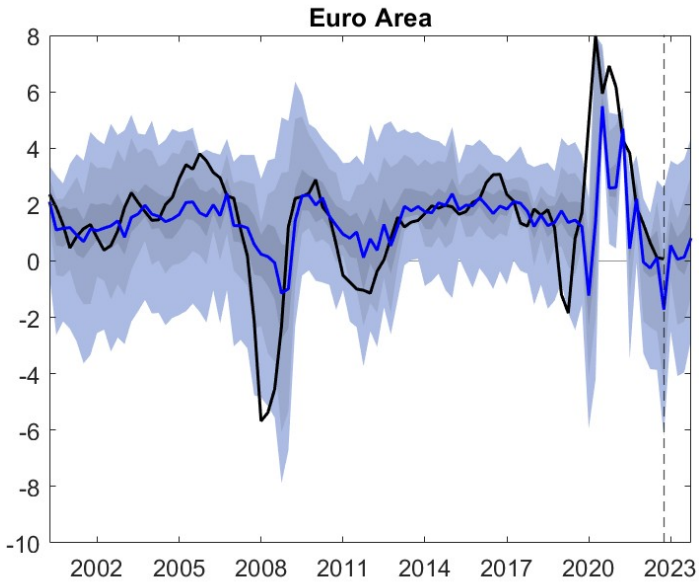

Figure 1.1: EA GDP growth conditional distribution (percentage points)
Black line: actual four-quarter-ahead GDP growth rate; Blue line: median; Shaded areas: 5^{th}, 10^{th}, 25^{th}, 75^{th}, 90^{th} and 95^{th} percentiles of the estimated conditional distribution.

An interesting feature of the obtained distribution is the asymmetric time-series variance on the two tails, a common result in the GaR literature: while the right tail of the distribution - representing unexpected economic booms - displays very little variation in the data sample, hovering around 3-4 percent, the converse applies to the left tail - measuring downside risks -, that varies much more significantly, with large drops during crisis episodes. The variance of the 5^{th} and the 10^{th} percentiles is around four times higher than that of the median or of the right tail of the distribution (Figure A1.3). Moreover, the dispersion of the distribution tends to increase when the median shrinks, suggesting the presence of a negative relationship between the first and the second moment. For what concerns the main drivers of GDP-at-risk in the high-dimensional dataset, the four-quarter-ahead left tail of the distribution co-moves strongly with indicators of real activity, as

[7]Given the limited time span covered by the data, predictive properties of the obtained distributions cannot be verified, as estimation is in-sample. However, the methodology allows for some improvement in terms of in-sample fit when compared to a standard quantile regression of GDP growth on its lags and the CISS indicator (as in Figueres and Jarociński (2020). The pseudo-R^2 for quantile regressions (Koenker and Machado (1999)) is reported in Table A1.2 in the Appendix.

the industrial production index for non-durable goods and for manufacturing firms. These variables are closely related also to the median of the distribution, confirming their leading properties for future economic activity. Other variables appear to be more specific to the left tail: the investment share of non-financial corporations, differently from production indices, does not appear among the most relevant series for the median of the distribution; price indices as the overall HICP and the services HICP, which hint at a relationship between the tails of the distributions of GDP growth and inflation; the residential property prices, a common measure of financial imbalances in the literature on GDP-at-risk and macroprudential policy (Aikman et al. (2019)). The relevance of GDP growth for both the left tail and the median is sensible, given the strong time dependence of macroeconomic outcomes[8].

Given the obtained conditional quantiles, it is possible to fit a parametric skewed-t distribution and to estimate its time-varying conditional moments, reported in Figure 1.2 with their moving average on a eight quarters window. The conditional mean of the GDP growth rate hovers between one and two percent in ordinary times, but displays quite significant downturns during periods of crisis. Moreover, during recessions, conditional variance tends to increase, so that a fall is the expected GDP growth is accompanied by an increase in the uncertainty around the conditional point estimates, particularly sizeable during the GFC. For what concerns higher moments, in most of the sample the conditional distribution is negatively skewed: this result is consistent with the asymmetry between tails observed in the conditional distribution, as the probability mass of the recessions episodes is larger than that representing unexpected periods of strong growth[9]. As far as kurtosis is concerned, it is noticeable that the drop in the probability of extreme data points caused by the reduction of variance in ordinary times is partially offset by the fatness of tails, as there is a moderate negative relationship between the second and the fourth moment.

[8]Figure A1.4 reports the conditional distribution for the Euro Area GDP growth obtained without the adjustment for the Covid period. The inclusion of the Covid outlier affects the results by blurring the relation between the macro-financial indicators in the X_t matrix and future GDP realizations. The effect is visible especially on the left tail of the distribution immediately before the Covid crisis, as the model tries to fit the future collapse in economic activity by interpreting it as a fall in conditional skewness. However, the asymmetry of the two tails is maintained, so that the main conclusions of the analysis remain valid to the robustness check.

[9]The estimated skewness decreases before crisis episodes (particularly, before the GFC and the SDC), suggesting the presence of some informative content of this indicator for economic outcomes in times of distress. However, this result is to be taken with caution as the estimation is in-sample and the indicator tends to increase during recessions (in contrast to the procyclical behaviour of the third moment observed in Delle Monache et al. (2023), Iseringhausen et al. (2023) among others).

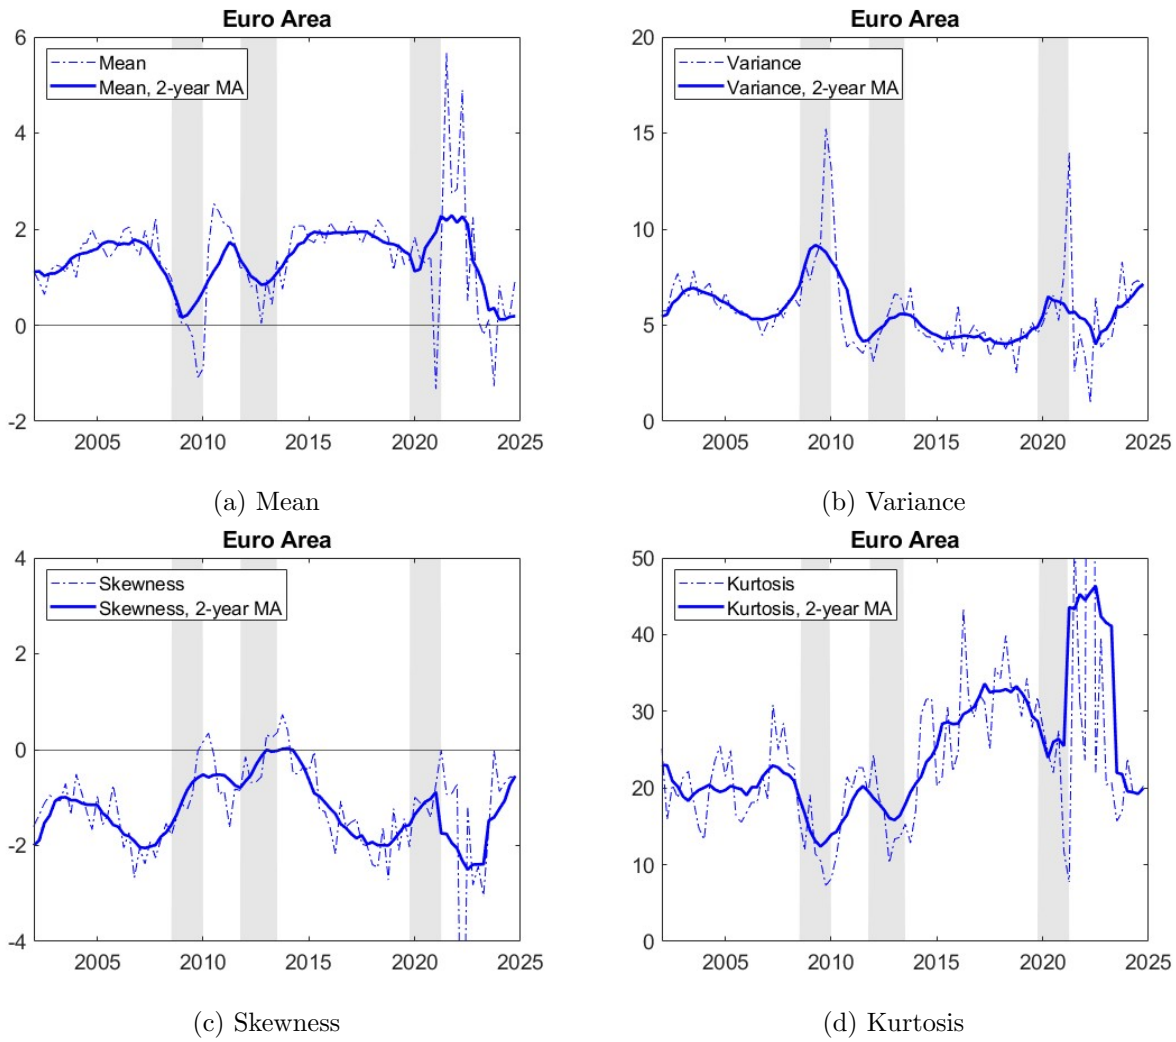

(a) Mean

(b) Variance

(c) Skewness

(d) Kurtosis

Figure 1.2: EA GDP growth conditional moments (percentage points)
Dashed lines: conditional moments; Solid lines: 2-year moving average of moments; Shaded areas: recessions.

Figure 1.3 reports the Expected Longrise (EL) and Expected Shortfall (ES) obtained from the parametric distribution. The information provided by these indicators partly overlaps with that provided by the non-parametric GDP-at-risk at the corresponding percentiles, but is enriched by incorporating the most extreme values of the distribution. Specifically, these indicators measure the expected change in GDP, conditional on GDP growth falling in the tails of its distribution. The asymmetry between the tails observed on the quantiles is confirmed also by these measures, as the time series variance of the EL is smaller than that of the ES. In particular, while the 95%-EL lingers around 5 percent throughout the whole sample, the 5-%ES displays major falls during crisis episodes. The potential for an unexpected growth leap of the EA economy, conditional on GDP

lying in the right tail on its distribution, appears therefore to be quite steady, while conditional downside risks are much larger, with possibly disastrous realizations during crisis episodes.

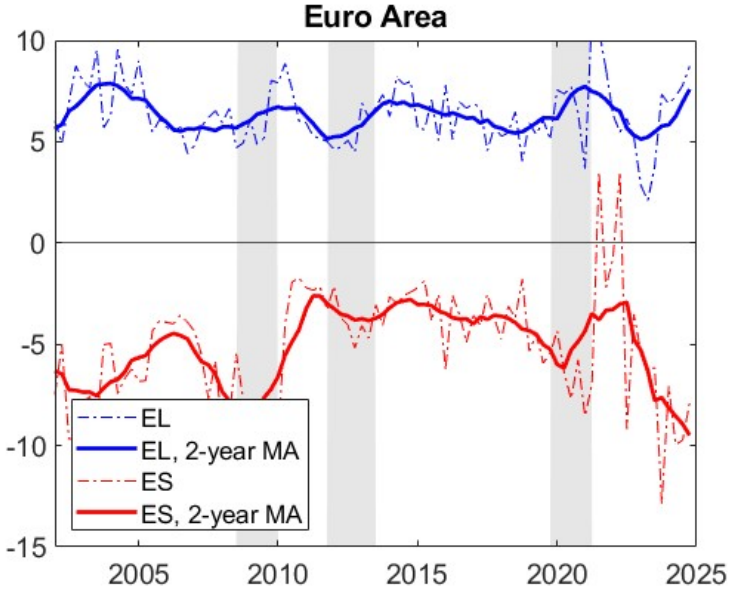

Figure 1.3: Expected Shortfall and Longrise, EA GDP growth (percentage points)
Dashed lines: ES and EL; Solid lines: 2-year moving average of EL/ES; Shaded areas: recessions.

1.4.2 Inflation-at-risk

The conditional distribution of four-quarter-ahead inflation for the EA is shown in Figure 1.4. Differently from GDP growth, the actual inflation level is quite aligned to the distribution median along the whole sample, at a level between 0 and 2 percent. The only exception is the large price increase of 2022-2023, during which actual inflation outpaced the conditional median by almost five percentage points, close to the right tail of the distribution. Notably, the asymmetry in tail variance observed for GDP growth is reversed for prices: lowest percentiles are stable throughout the sample around 0 percent - with the exception of the deflation risk observed during the GFC, the Sovereign Debt Crisis (SDC) and in the Covid period -, while the right tail of the distribution - a measure of inflation-at-risk - seems more volatile, with a jump around 2023. However, most of the right tail variance seems to be driven by the very extreme percentiles, while the 90% percentile remain quite constant across time, suggesting a lower volatility compared to the GDP-at-risk. The out-of-sample projections are compatible with a fast decline of inflation to the two percent target of the ECB, with a relatively low dispersion of the distribution if compared with its historical levels. Figure

A1.5 reports the time variance of the conditional quantiles, confirming the asymmetric behaviour of the two tails, even if asymmetries are smaller than those of GDP growth. In particular, only the most extreme percentiles display exceptionally volatile (or stable) trends, while those closer to the median are more similar to each other. Among the series which co-move more strongly with the right tail of the inflation conditional distribution, there are the monetary aggregates M1 and M2, which on the contrary do not seem to be much related to the median of the distribution. This result implies a link between monetary aggregates and future prices which seems relevant only when considering risks of very high inflation, while these indicators do not seem to be much informative of future inflation trends in baseline scenarios. Upside risks for prices stemming from credit are hinted at by the relevance of households' long-term liabilities for the right tail: these liabilities do not seem to be predictive of inflation in ordinary times, but become relevant for upside risks, given their ability to measure the build-up of financial imbalances. The correlation of such liabilities with confidence indicators in the construction sector explains the relevance of this latter for the right tail of the inflation distribution. On the other hand, producers' price indices for non-durable goods and the HICP for consumer goods are relevant for both the right tail and the median of the distribution.

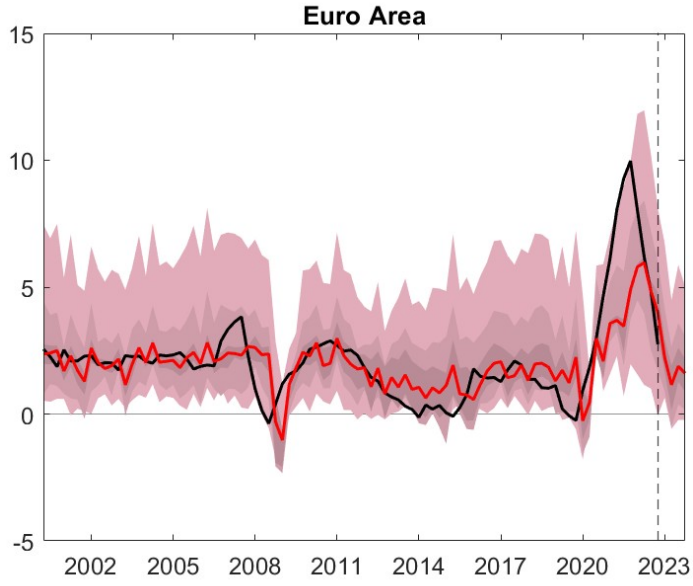

Figure 1.4: EA inflation conditional distribution (percentage points)
Black line: actual four-quarter-ahead inflation rate; Red line: median; Shaded areas: 5^{th}, 10^{th}, 25^{th}, 75^{th}, 90^{th} and 95^{th} percentiles of the estimated conditional distribution.

Figure 1.5 reports the conditional moments of the parametric inflation distribution. The conditional

mean stays around 2 percent, the ECB medium-term target, in most of the sample, with the only exceptions of the GFC and the latest inflation spike. The conditional variance of the distribution displays limited variations across time, with the only exception of the 2021-2022 period. The first and the second moment display a moderately positive correlation, suggesting that in periods of high inflation, also inflation uncertainty tends to increase. Conditional skewness is positive in most of the sample, as the tail right is fatter, especially in the 2014-2021 period; the distribution becomes more symmetric as its mean increases during the high inflation period. Finally, kurtosis remains quite constant in the sample, with a weakly negative correlation with conditional variance: changes in the probability of observing outliers caused by increases in conditional variance are partially offset by opposite movements in the fourth moment.

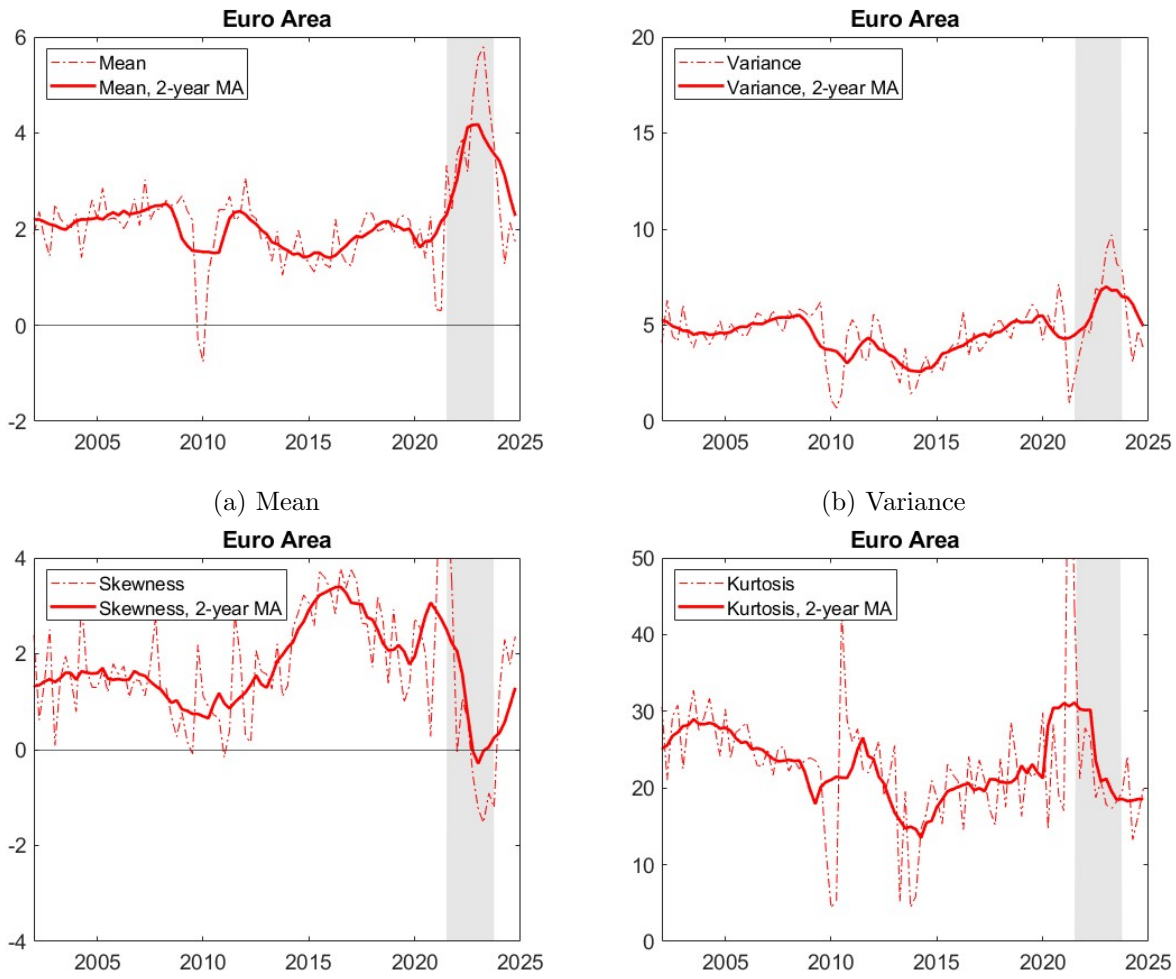

(a) Mean

(b) Variance

(c) Skewness

(d) Kurtosis

Figure 1.5: EA inflation conditional moments (percentage points)
Dashed lines: conditional moments; Solid lines: 2-year moving average of moments;
Shaded areas: quarters of high inflation (yearly HICP growth above 4 percent).

Figure 1.6 reports the counterparts of the EL and the ES for the inflation distribution, denoted here as "Expected High Inflation" (EHI) and "Expected Deflation" (ED), respectively. Results are opposite to what observed on the GDP distribution: as the expected outcome in the 95% percentile varies quite significantly across time, displaying a large fall around the GFC - when the recession muted upside risks to inflation - and a sudden increase in 2022, the expected deflation is quite steady, suggesting that the possible reduction of prices does not exceed -1 or -2 percent even during crises. As this measure is a conditional (on inflation lying in the lowest 5% scenario) VaR, this result does not indicate a stable probability of deflation, but rather a quite steady rate of expected price drops in the left tail scenarios. The only exception is the recent surge in inflation

and the subsequent monetary policy response, which induced an increase in conditional variance and temporarily stretched the tails of the distribution.

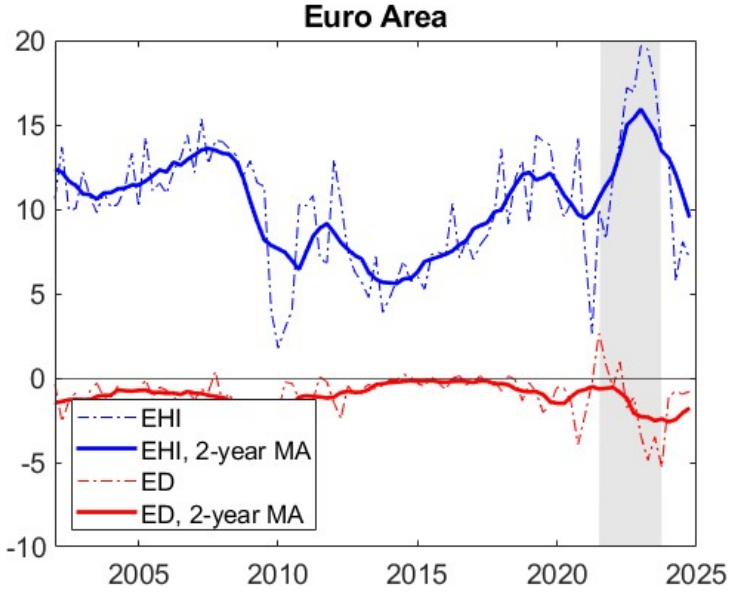

Figure 1.6: Expected High Inflation and Deflation, EA (percentage points)
Dashed lines: EHI and ED; Solid lines: 2-year moving average of EHI/ED;
Shaded areas: quarters of high inflation (yearly HICP growth above 4 percent).

1.4.3 Quantile AD/AS curves and probabilities of recession and high inflation

The conditional distributions of GDP growth and inflation constructed separately are now considered jointly, so to evaluate their relationship in different parts of their density functions. In the scatter plots of Figure 1.7, each observation represents a pair of (y, π) conditional quantiles in a given quarter, with the panels on the left representing y quantiles on the left tail (5^{th} percentile), and each row corresponding to a different π percentile (the 5^{th} and the 95^{th})[10][11]. The top left and bottom right plots report pairs of GDP growth and inflation considered at the same tail, and can therefore be interpreted as a conditional aggregate supply (AS) curve seen at different points of the distributions[12]. The curve outlined by conditional quantiles appears to be flat when this relationship is evaluated on the left tail, suggesting that positive demand shocks are unlikely to

[10]The conditional medians of the two variables are positively correlated, similarly to their unconditional counterparts, as reported by Figure A1.6 in the Appendix.

[11]The results of a similar exercise, carried out considering the distributions of the single member states, are presented in the Appendix.

[12]Considering conditional GDP growth and inflation at the same quantile implies evaluating the two variables in a context of a shock that affects them in the same direction (*e.g.* a demand shock). For this reason, the obtained (y, π) pairs on the main diagonal delineate a supply curve at different quantiles.

affect prices during times of weak economic activity. Conversely, the curve becomes steeper on the right tail: when output growth is close to its potential, or in periods of high inflation, expansionary demand shocks would translate into price increases, with very limited effects on economic activity. Analogously, the plots in the top right and bottom left corners can be read as conditional demand curves at the quantiles. The resulting AD curve is virtually flat on the right tail of the GDP growth distribution, implying that a positive supply shock would cause output to expand; on the contrary, the reaction of prices would be quite muted, consistently with the presence of downward price rigidities. Conversely, the demand curve is steeper on the left tail: as a consequence, the reaction of both output and prices to a negative supply shock would be rather intense[13].

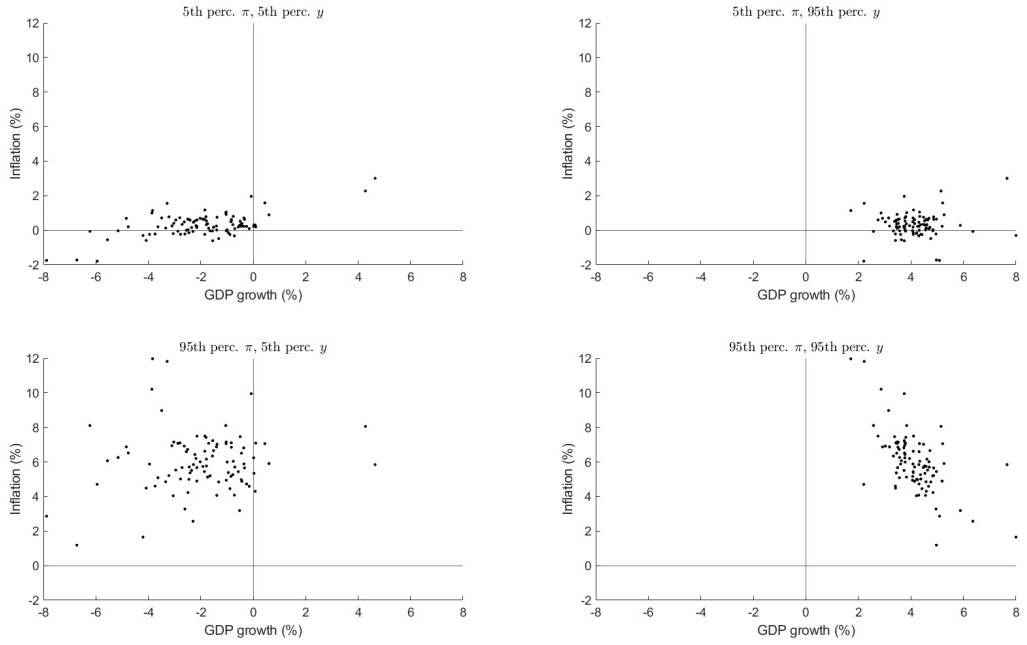

Figure 1.7: Conditional AD/AS curves at different quantiles of y and π, EA

Probabilities of recession and high inflation four quarters ahead based on the information provided by the high-dimensional dataset can be obtained as a by-product of the parametric distributions[14]. These are reported in Figure 1.8, where high inflation is defined as a yearly HICP increase of more

[13] A key limitation of this exercise is the absence of a model for the joint distribution of GDP growth and inflation; although the conditioning set is common, the two distributions are estimated separately, ignoring the correlation between the two marginals. However, some of its results can be suggestive of a methodology to model different parts of aggregate demand and supply curves using conditional quantiles.

[14] A time series of approximate probabilities of recession and high inflation can be obtained also non-parametrically by directly reversing the stepwise conditional quantile function. The obtained results - not reported for the sake of brevity - are broadly in line with the probabilities obtained from the parametric distribution.

than four percent. The model assigns the highest probabilities of a recession prior to the GFC, around the SDC, during the Covid crisis and after the inflation surge of 2022. High inflation risks remain very limited (around 10 percent) throughout the whole sample, and show a moderately negative relationship with recession risks until 2021. On the contrary, the sudden jump in inflation risk estimated in 2022 is accompanied by a just smaller rise of recession risks. This positive correlation, not present in the previous episodes of crisis, can be attributed to the supply-side origin of the shock, which exogenously pushed up prices and slowed down economic activity[15].

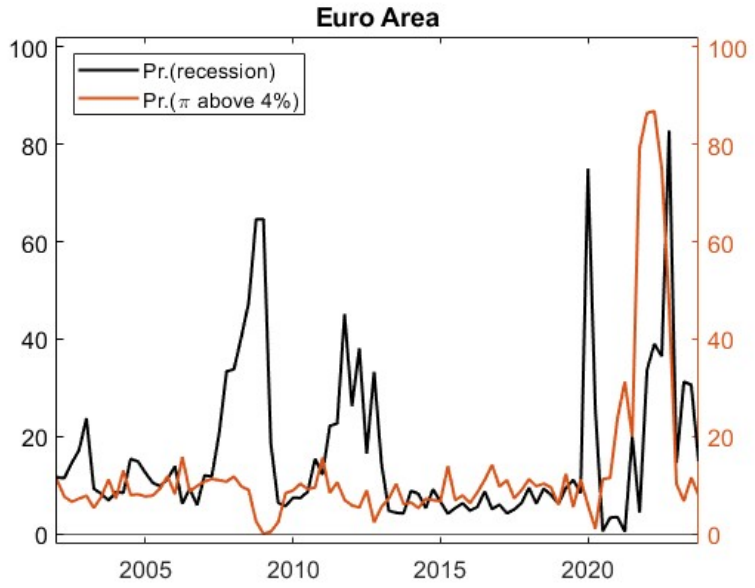

Figure 1.8: Conditional probabilities of recession and high inflation, EA (percentage points)

1.5 Assessing macroeconomic synchrony at the quantiles

This Section considers the main results obtained from conditional distributions at country level: the GDP growth conditional distributions for single member states are reported in 1.5.1, while the results on inflation distributions and their synchrony are in 1.5.2.

[15]Figure A1.17 reports the conditional probabilities of recessions and high inflation at country level. The timing of the increases in their probabilities show a good degree of synchrony among countries. However, while the sensitivity to price upside risks is very similar during the whole period, this conclusion does not apply to recessions. The estimated probability of a recession during the GFC ranged from 30-50 percent in France and Belgium to over 80 percent in Greece, Italy, Portugal and Spain; similar asymmetries arise during the SDC, hinting at a higher vulnerability of peripheral countries to downside risks.

1.5.1 GDP-at-risk

Figure 1.9 displays the GDP growth conditional distribution of the four main EA countries. In this exercise, factors are extracted from the large dataset at EA level using the single countries' GDP growth four quarters ahead as target: the time-varying distributions reported here are therefore conditional on macroeconomic and financial factors at EA level. The asymmetry between stable right tails (around 3 percent) and varying left tails is present in all the analysed countries[16]. Interestingly, while the timing of GDP-at-risk drops is quite coincident across countries, the value reached by the conditional 5^{th} percentile varies quite significantly during periods of stress. The left tail of the distribution of Germany and Spain appears to be particularly far from the median of the distribution, indicating a sizeable degree of downside risks, which look less pronounced in France. The whole conditional distribution of Italian GDP growth is shifted downwards with respect to that of the other countries, suggesting a persistently weaker growth potential also in baseline scenarios. Finally, the future perspectives outlined by the out-of-sample observations are consistent with weak growth prospects in all the considered countries[17].

The distributions obtained by conditioning on factors extracted from the EA dataset and from datasets at country level are compared in Figure A1.9. The subplots report the 5^{th}, the 50^{th} and the 95^{th} percentiles of the distribution conditioning on EA factors in blue, and their country-level counterparts in red. The median and the right tail are virtually overlapping in most of the sample and in all countries, suggesting that the information contained in the dataset at EA level - if treated by targeting GDP growth at national level when extracting quantile factors - differs only slightly from that provided by datasets at country level. On the other hand, also the left tails of the distributions are quite aligned, suggesting that also downside risks at country level are captured quite well by risk determinants extracted at EA level. However, there are some episodes of misalignment that suggest an additional degree of caution when dealing with low percentiles: common factors

[16] An exception to this asymmetry is represented by Greece - reported in Figure A1.7 -, where the dispersion of the distribution is quite high on both sides of the median point, in most of the sample. The time series variance of GDP growth conditional quantiles for the EA countries distribution - not reported for the sake of brevity - confirms the higher volatility of left tails in all the analysed countries.

[17] Figure A1.8 reports the distributions obtained without treating the original data to account for the Covid period. As for the distribution at EA level, the outlier obfuscates the relation between left tails and macro-financial factors to quite a limited extent, especially in the period before the pandemic. Nevertheless, the asymmetry between the two tails is confirmed.

suggested significant downside risks in Germany in the pre-GFC period, while the left tail of the distribution conditioned on country-level factors was very close to the conditional median, implying an additional degree of downside risks of external source that could not have been captured using only the German dataset. An opposite example is provided by the Italian and Spanish GDP-at-risk during the SDC, higher - in absolute terms - if conditioned on factors extracted at country level rather than from the EA dataset: this suggests the presence of country-specific downside risks - possibly related to interest rates paid on sovereign bonds and on debt imbalances, in this case - that were not observed through common factors.

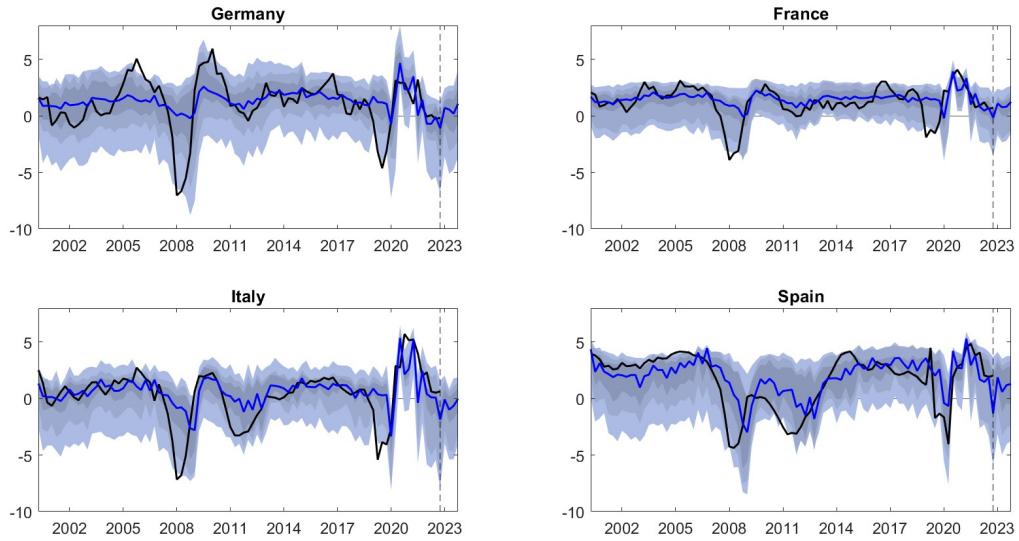

Figure 1.9: EA countries' GDP growth conditional distributions, common factors (percentage points)

Figure 1.10 reports the cross-sectional standard deviation of the conditional distributions in the nine countries, considered at different quantiles: this dispersion index can be used as an (inverse) measure of synchrony between GDP cycles, with the advantage - if compared to indices based on actual realizations - of capturing coordination on the whole distribution in each period. A first result is the low dispersion of median GDP growth rates across the whole sample, with small exceptions during episodes of crisis; the dispersion of actual growth rates follows quite closely the median dispersion during periods of growth. Similarly, the right tails of the distributions are quite synchronized throughout the period. On the contrary, the synchrony of left tails varies significantly across time, with very large spikes in periods of distress, when also the actual growth rates are

very disperse. Moreover, differently from other portions of the distribution, the standard deviation of GDP-at-risk tends to revert quite slowly to the lower levels observed before crises. Looking at the standard deviations from a medium-term perspective, the left tail is not only found to be systematically more disperse than the rest of the distribution, but there are also little to no signs of a long-run reduction in its cross-sectional dispersion, suggesting very little improvements in the synchronization of economic cycles in bad times across EA countries in the last twenty years. These findings corroborate some results from the recent literature on the empirical assessment of OCAs advancements in the Euro Area: as in Kunovac et al. (2022) and de Haan et al. (2024), progress towards a stronger synchrony is slow and non-monotonic, as crisis episodes hit countries asymmetrically inducing divergence, measured in this work by the cross-sectional standard deviation of left tails of the GDP growth conditional distributions[18].

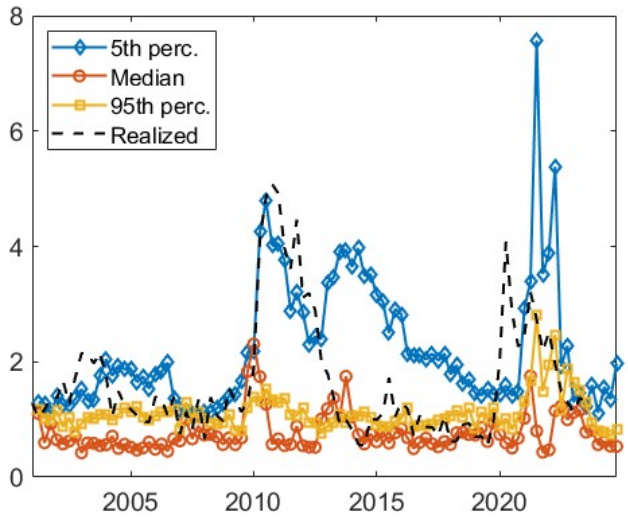

Figure 1.10: Conditional cross-sectional standard deviation for alternative τ, GDP growth

Figure 1.11 displays the conditional means of the GDP growth distributions of EA countries. Means tend to follow very similar trends, with coincident periods of growth and crisis in all countries; mean GDP growth lies around 2 percent in ordinary times for most countries, with the exception of Italy, where the central tendency of the distribution is close to 0 in most of the sample. The size of the impact of GFC is quite similar in all countries, with the exception of France and Belgium, while the conditional means of peripheral countries were more affected by the SDC. The trend of

[18]These results hold true when considering GDP growth distributions constructing by either conditioning the distributions on factors extracted from datasets at national level, or neglecting the adjustment for the Covid period.

conditional variances, reported in Figure 1.12, is quite similar to that of the EA distribution; the negative relationship between the first and the second moment can be observed both on the time series dimension (as variance tends to increase in periods of low conditional mean, for all countries) and on the cross section (as variance increases more during crises in countries where the fall of the first moment is stronger). The long-run level of variance is quite similar in most countries, and slightly lower only in Belgium and France; the large variance of the Greece distribution reflects the high dispersion already observed in the non-parametric distribution.

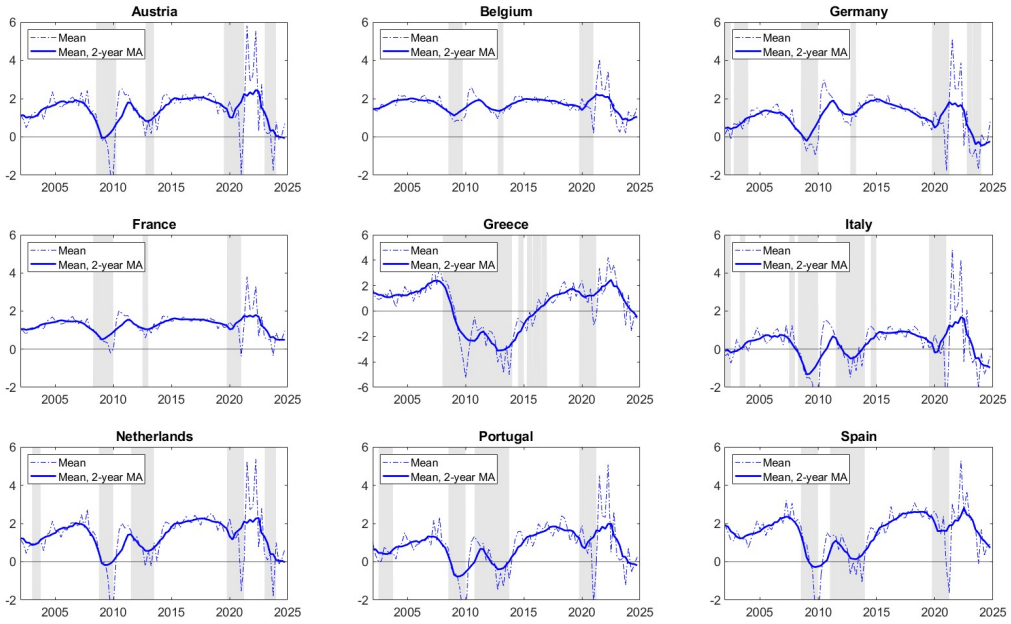

Figure 1.11: Conditional mean, EA countries GDP growth (percentage points)
Dashed lines: conditional mean; Solid lines: 2-year moving average of mean; Shaded areas: recessions.

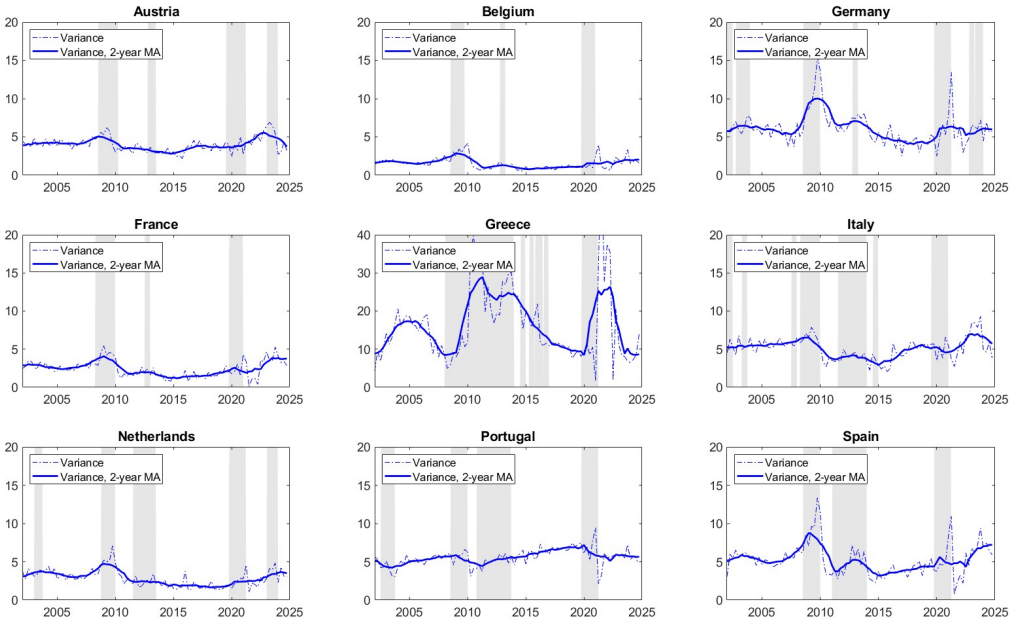

Figure 1.12: Conditional variance, EA countries GDP growth (percentage points)
Dashed lines: conditional variance; Solid lines: 2-year moving average of variance; Shaded areas: recessions.

1.5.2 Inflation-at-risk

The conditional distributions of inflation in the four main member states, obtained from factors extracted from the common dataset are reported in Figure 1.13. As for the EA, the only exception to the stability of the left tail of the distributions - that hovers around 0 percent during the whole sample - is the increase of deflation risk during crisis episodes. On the contrary, the percentiles on the right part of the distribution are quite volatile in all countries, and reach their highest point during the high inflation observed after 2022. During this episode, actual inflation rates do not follow the conditional median, as they instead do in most of the sample[19].

[19]The exposure of price leves to upside risks seems stronger in small open countries as Belgium, Greece and the Netherlands, whose distribution is more skewed to the right. Full results are reported by Figure A1.10 in the Appendix.

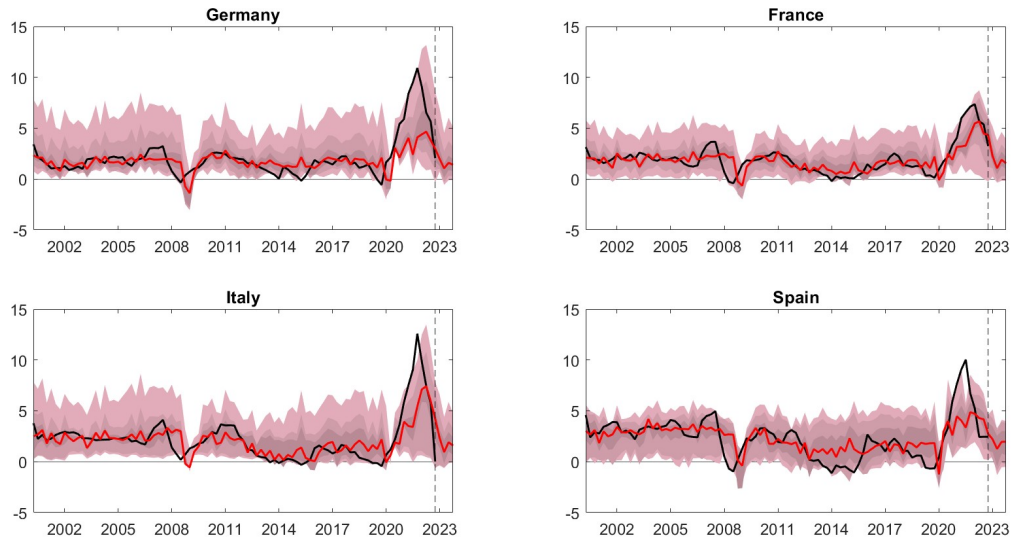

Figure 1.13: EA countries' inflation conditional distributions, common factors (percentage points)

Figure A1.11 compares the inflation distributions obtained conditioning on common (blue) and country-specific (red) factors. The dynamics of the left tails and of the conditional median are nearly equivalent in the whole sample, and also the right tails of the conditional distributions appear quite synchronous. These results indicate that the risks to price stability - both in periods of downward and upward pressures - are quite aligned throughout the Euro Area, to a wider extent than observed for GDP growth. The few episodes of departing right tails are observed during the pre-GFC period (in Germany, France and Italy), when upside risks to inflation were driven by common factors not fully captured by country-level indicators; moreover, these latter slightly underestimate the right tail during the Covid crisis in Italy and Portugal, suggesting that a relevant part of inflation risks was driven by international supply-side determinants, not entirely captured by internal factors.

The cross-sectional standard deviation of the conditional inflation distributions, evaluated at different quantiles, is displayed in Figure 1.14. In contrast with GDP growth, here the average dispersion of the 95^{th} percentile is higher in comparison to the median and the lower percentiles. The dispersion of the actual inflation rates tends to follow that of the median, with the exception of the high inflation period of 2022, when it was closer to the right tail. However, there are two departures

from the GDP results: first, for what concerns stress periods, the dispersion of the right tail of the inflation distribution varies in time much less than the indicator computed on the left tails of GDP growth, and tends to return quite quickly to previous levels, suggesting that the degree of divergence during times of upside risks to prices increases significantly less than what occurs in the case of downwards risks to economic activity; second, taking a longer term perspective, the dispersion of the distributions is more limited for inflation on all the percentiles, indicating a quite high and stable level of synchronization of price dynamics in the EA, reached already before the introduction of the common currency.

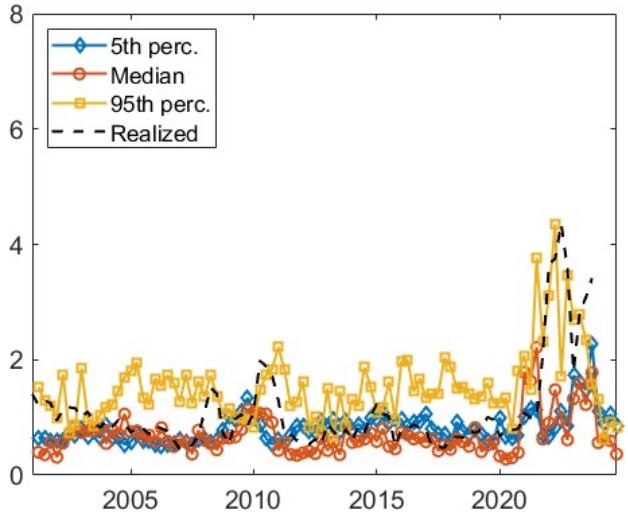

Figure 1.14: Conditional cross-sectional standard deviation for alternative τ, inflation

The conditional moments estimated from the inflation parametric distributions of inflation provide further evidence supporting the synchrony of inflation dynamics in the EA. Figure 1.15 reports the conditional means in the analysed EA countries: the most relevant result is the almost perfect co-movement of the time series, which decrease by similar proportions reflecting downward risks during the GFC and increase together in the high inflation period. Conditional means appear quite anchored to the ECB target of 2 percent in most of the sample and in all countries, with short-lived deviations during periods of crisis. The conditional variances of the inflation distributions at country level are reported in Figure 1.16. The second moment of the distributions is quite stable across time, displaying as observed on the EA moments a moderately positive correlation with the

conditional mean on the time series dimension and a particularly strong increase during the period of high inflation. On the cross section, is is interesting to notice the higher variance of inflation in small open economies as Belgium, Greece and the Netherlands, more exposed to exogenous price shocks than other larger economies.

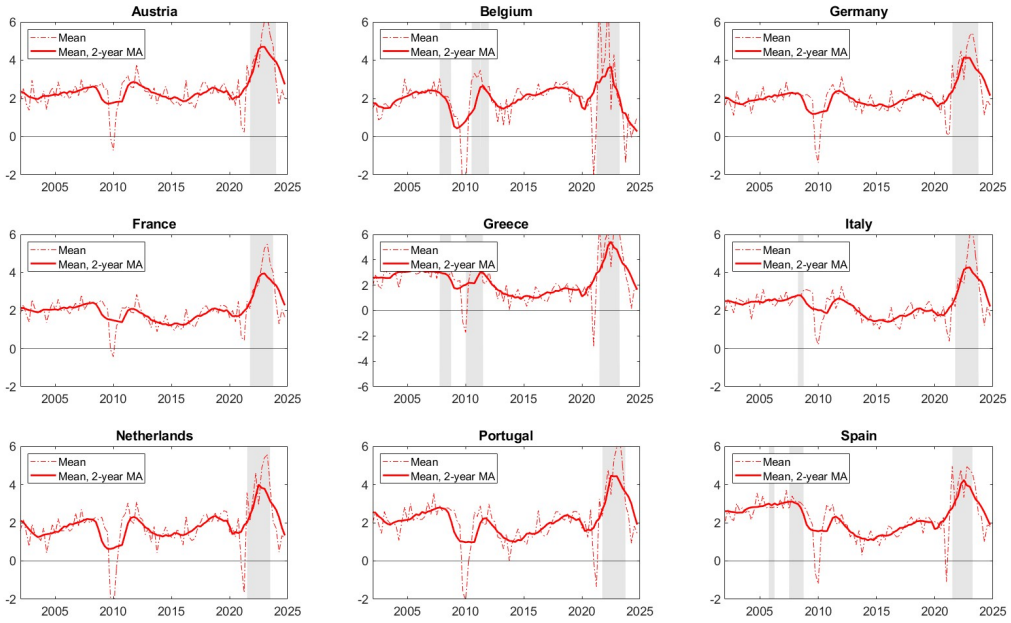

Figure 1.15: Conditional mean, EA countries inflation (percentage points)
Dashed lines: conditional mean; Solid lines: 2-year moving average of mean;
Shaded areas: quarters of high inflation (yearly HICP growth above 4 percent)

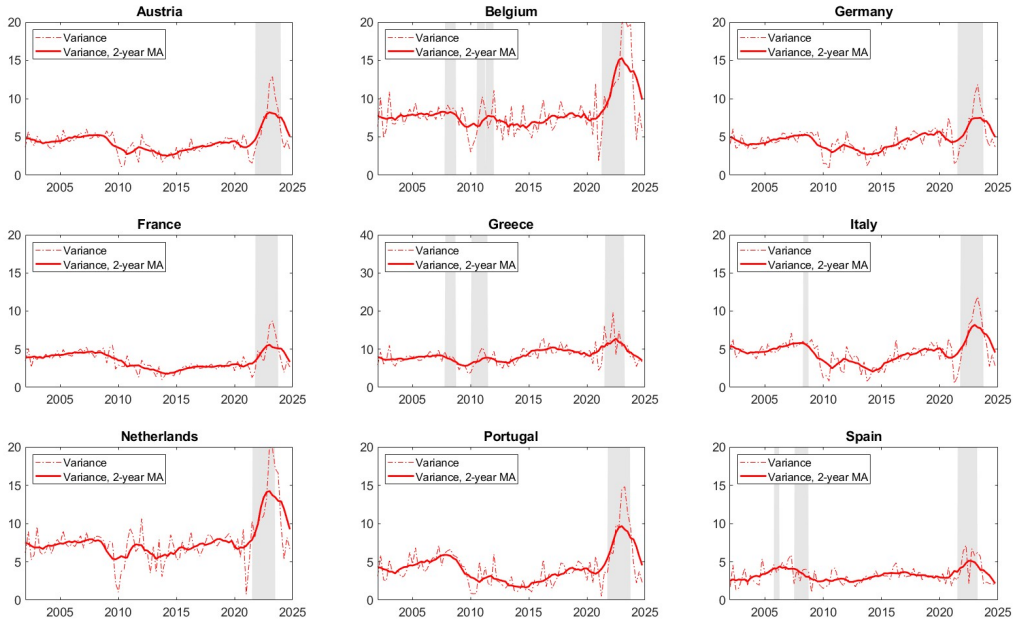

Figure 1.16: Conditional variance, EA countries inflation (percentage points)
Dashed lines: conditional variance; Solid lines: 2-year moving average of variance;
Shaded areas: quarters of high inflation (yearly HICP growth above 4 percent)

1.6 Conclusions

In this paper, the GDP-at-Risk and Inflation-at-risk for the Euro Area and its member states were estimated using the information provided by a large number of explanatory variables. For this purpose, a set of quantile-specific latent factors were extracted from a high-dimensional dataset, each one summarizing the relevant information at different points of the target variable distribution, and used them as independent variables in the construction of a conditional quantile function. In a subsequent step, a parametric distribution was fit on the interpolated quantiles, allowing for the estimation of time-varying conditional moments and measures of tail risk as the Expected Shortfall. Results indicate strong asymmetries in the GDP growth distribution, which displays fatter and volatile left tails in contrast with steadier right tails. Conditional mean and variance appear negatively related, suggesting that the uncertainty around outcomes is higher during recessions. Opposite - but minor in extent - asymmetries are found also in the inflation distribution, that displays fatter and more volatile right tails, while remaining fairly stable on its lower quantiles in the whole sample; the first moment of the inflation distribution remains substantially stable around

the ECB target of two percent in most of the sample. Leveraging on the distributions of GDP growth and inflation at country level, the analysis evaluated the synchrony of business cycles across EA member states, inversely measured by the cross-sectional dispersion of conditional quantiles of the distributions. The objective is therefore to complement the study of the synchrony between economic cycles by considering the whole conditional distributions of growth and inflation. Results hint at an imbalance between the synchrony between countries' inflation distributions - on the whole distribution, even if slightly lower on the right tail - and the dispersal observed between the left tails of the GDP distributions during crisis, as measured by the cross-sectional standard deviations between low percentiles and by the conditional probabilities of recession implied by the parametric distributions. On the other hand, the short deviations of the conditional distribution from the target levels of inflation indicate a good level of synchrony of upside risks to inflation across countries in the euro area, at least in the analysed sample, characterized by positive shocks to the price level that hit the member states' economies in quite a symmetric manner.

A1 Appendix

Table A1.1: Series in the EA-MD-QD dataset at EA level (Barigozzi, Lissona, and Tonni (2024))

Name	Source	Class	Name	Source	Class
Real GDP	EUR	R	HHs - Assets: ST Loans	EUR	F
Real Exports	EUR	R	HHs - Assets: LT Loans	EUR	F
Real Imports	EUR	R	HHs: Total Financial Liabilities	EUR	F
Real Gov. Final consumption	EUR	R	HHs - Liabilities: ST Loans	EUR	F
Real HHs consumption expenditure	EUR	R	HHs - Liabilities: LT Loans	EUR	F
Real HHs cons. expenditure: Durable Goods	EUR	R	NULCs: Industry	EUR	N
Real HHs cons. expenditure: ND Goods and Serv.	EUR	R	NULCs: Mining and Quarrying	EUR	N
Real Gross capital formation	EUR	R	NULCs: Manufacturing	EUR	N
Real GFCF	EUR	R	NULCs: Construction	EUR	N
Real GFCF: Construction	EUR	R	NULCs: Trade, Transport, Food, IT	EUR	N
Real GFCF: Machin. and Eq.	EUR	R	NULCs: Financial Activities	EUR	N
Adjusted HHs Real Disposable Income	EUR	R	NULCs: Real Estate	EUR	N
Actual Final Consumption Expenditure of HHs	EUR	R	NULCs: Prof., Scientific, Tech activities	EUR	N
Gross Profit Share of NFCs	EUR	R	Real Exchange Rate	EUR	F
Gross Investment Share of NFCs	EUR	R	Exchange Rate (US dollar)	EUR	F
Gross Investment Rate of HHs	EUR	R	3-Months Interest Rates	EUR	F
Gross HHs Savings Rate	EUR	R	6-Months Interest Rates	EUR	F
Total Employment (domestic concept)	EUR	R	LT Interest Rates	EUR	F
Employees (domestic concept)	EUR	R	IPI: Manufacturing	EUR	R
Self Employment (domestic concept)	EUR	R	IPI: Capital Goods	EUR	R
Hours Worked: Total	EUR	R	IPI: Consumer Goods	EUR	R
Employment: Agriculture, Forestry, Fishing	EUR	R	IPI: Durable Consumer Goods	EUR	R
Employment: Industry	EUR	R	IPI: Non Durable Consumer Goods	EUR	R
Employment: Manufacturing	EUR	R	IPI: Intermediate Goods	EUR	R
Employment: Construction	EUR	R	IPI: Energy	EUR	R
Employment: Trade, transport, food	EUR	R	TI: Manufacturing	EUR	R
Employment: Information and Communication	EUR	R	TI: Capital Goods	EUR	R
Employment: Financial and Insurance activities	EUR	R	TI: Consumer Goods	EUR	R
Employment: Real Estate	EUR	R	TI: Durable Consumer Goods	EUR	R
Employment: Prof., Scientific, Tech activities	EUR	R	TI: Non Durable Consumer	EUR	R
Employment: PA, education, health/social services	EUR	R	TI: Intermediate Goods	EUR	R
Employment: Arts and recreational activities	EUR	R	TI: Energy	EUR	R
Unemployment: Total	EUR	R	PPI: Capital Goods	EUR	N
Unemployment: Over 25 years	EUR	R	PPI: Consumer Goods	EUR	N
Unemployment: Under 25 years	EUR	R	PPI: Durable Consumer Goods	EUR	N
Real Labour Productivity	EUR	R	PPI: Non Durable Consumer Goods	EUR	N
Wages and salaries	EUR	N	PPI: Intermediate Goods	EUR	N
Employers' Social Contributions	EUR	N	PPI: Energy	EUR	N
Tot.Ec. - Assets: ST Debt Securities	EUR	F	HICP: Overall Index	ECB	N
Tot.Ec. - Assets: LT Debt Securities	EUR	F	HICP: All Items: no Energy & Food	ECB	N
Tot.Ec. - Assets: ST Loans	EUR	F	HICP: Goods	ECB	N
Tot.Ec. - Assets: LT Loans	EUR	F	HICP: Industrial Goods	ECB	N
Tot.Ec. - Liabilities: ST Debt Securities	EUR	F	HICP: Services	ECB	N
Tot.Ec. - Liabilities: LT Debt Securities	EUR	F	HICP: Energy	EUR	N
Tot.Ec. - Liabilities: ST Loans	EUR	F	Real GDP Deflator	EUR	N
Tot.Ec. - Liabilities: LT Loans	EUR	F	Residential Property Prices	FRED	N
NFCs: Total Financial Assets	EUR	F	Industrial Confidence Indicator	EUR	C
NFCs - Assets: ST Loans	EUR	F	Consumer Confidence Index	EUR	C
NFCs - Assets: LT Loans	EUR	F	Economic Sentiment Indicator	EUR	C
NFCs: Total Financial Liabilities	EUR	F	Construction Confidence Indicator	EUR	C
NFCs - Liabilities - ST Loans	EUR	F	Retail Confidence Indicator	EUR	C
NFCs - Liabilities - LT Loans	EUR	F	Services Confidence Indicator	EUR	C
GG: Total Financial Assets	EUR	F	Business Confidence Index	OECD	C
GG - Assets: ST Loans	EUR	F	Consumer Confidence Index	OECD	C
GG - Assets: ST Loans	EUR	F	Money Stock: Currency	ECB	F
GG: Total Financial Liabilities	EUR	F	Money Stock: M1	ECB	F
GG - Liabilities: ST Loans	EUR	F	Money Stock: M2	ECB	F
GG - Liabilities: LT Loans	EUR	F	Share Prices	OECD	F
HHs: Total Financial Assets	EUR	F	Passenger's Cars Registrations	ECB	R

HHs: Households; GFCF: Gross Fixed Capital Formation; NFCs: Non-Financial Corporations; ST/LT: Short/Long-Term;
GG: General Government; NULCs: Nominal Unit Labour Costs; IPI: Industrial Price Index;
PPI: Producer Price Index; TI: Turnover Index; HICP: Harmonized Index of Consumer Prices

A1.1 Imputation of GDP growth during the Covid period

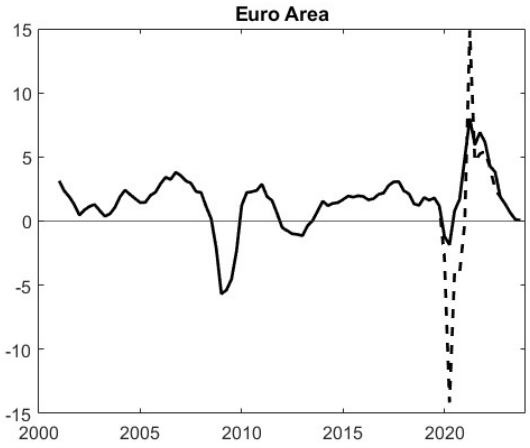

Figure A1.1: EA GDP growth, actual and imputed values (percentage points)

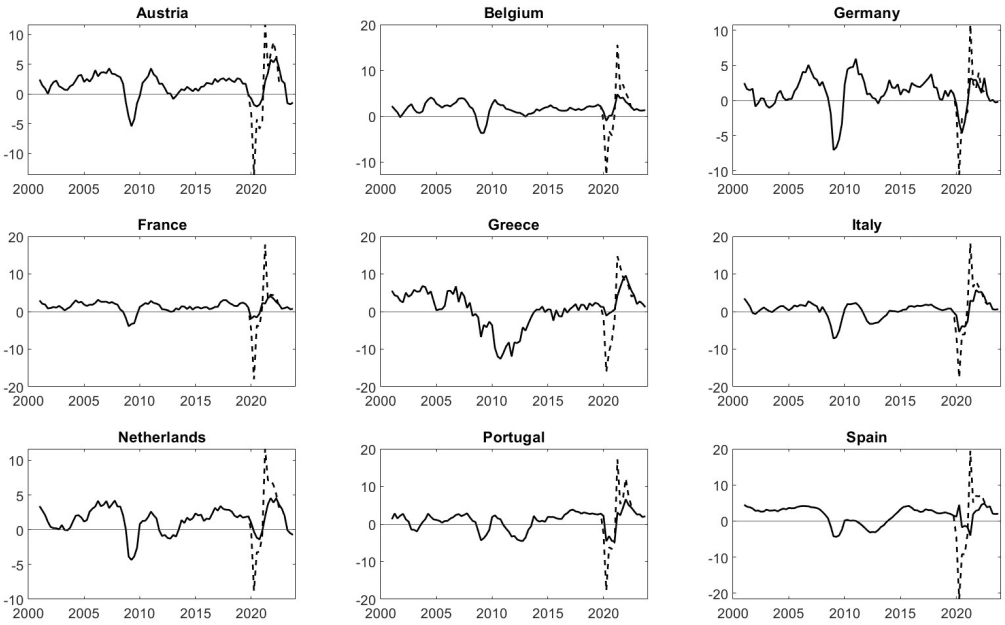

Figure A1.2: EA countries GDP growth, actual and imputed values (percentage points)

A1.2 Other results for the Euro Area

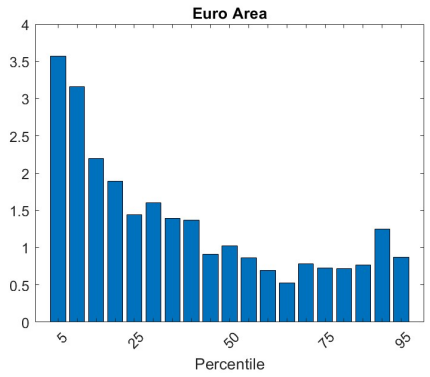

Figure A1.3: Variance of EA GDP conditional quantiles

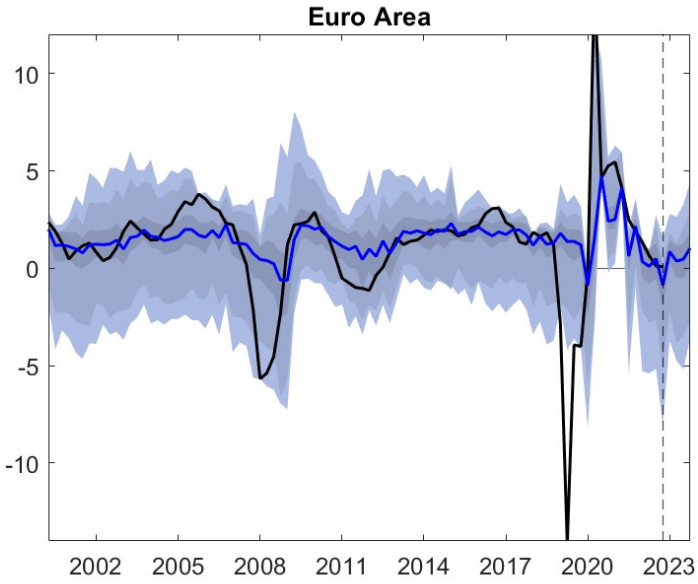

Figure A1.4: EA GDP growth conditional distribution, without Covid adjustment

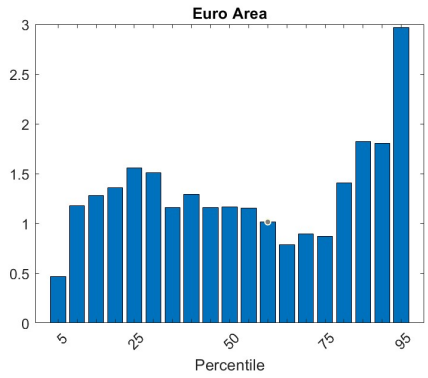

Figure A1.5: Variance of EA inflation conditional quantiles

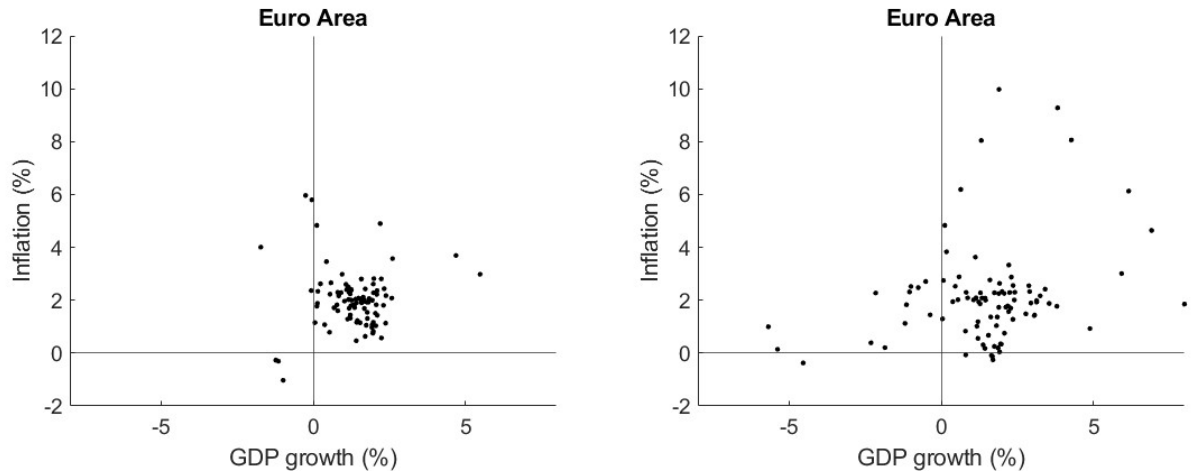

(a) Conditional medians

(b) Unconditional values

Figure A1.6: GDP growth and inflation, EA

Table A1.2: In-sample pseudo-R^2, EA

	GDP		Inflation	
Percentile	PQR	CISS	PQR	CISS
5	0.759	0.756	0.841	0.814
10	0.612	0.586	0.703	0.665
15	0.464	0.436	0.602	0.537
20	0.399	0.351	0.535	0.420
25	0.339	0.286	0.473	0.325
30	0.289	0.235	0.403	0.254
35	0.248	0.197	0.371	0.199
40	0.221	0.164	0.330	0.159
45	0.199	0.136	0.298	0.131
50	0.190	0.113	0.272	0.114
55	0.198	0.098	0.249	0.101
60	0.210	0.095	0.240	0.100
65	0.226	0.113	0.233	0.107
70	0.274	0.153	0.256	0.117
75	0.323	0.206	0.275	0.134
80	0.394	0.288	0.332	0.176
85	0.476	0.381	0.414	0.246
90	0.593	0.498	0.507	0.356
95	0.717	0.677	0.672	0.616

A1.3 Other results for the EA countries

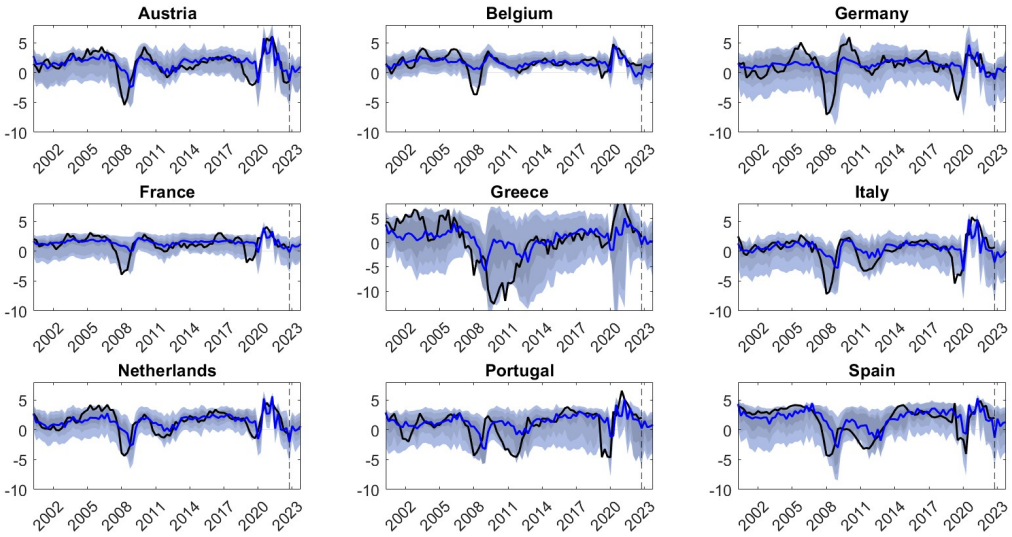

Figure A1.7: GDP growth conditional distributions, common factors (percentage points)

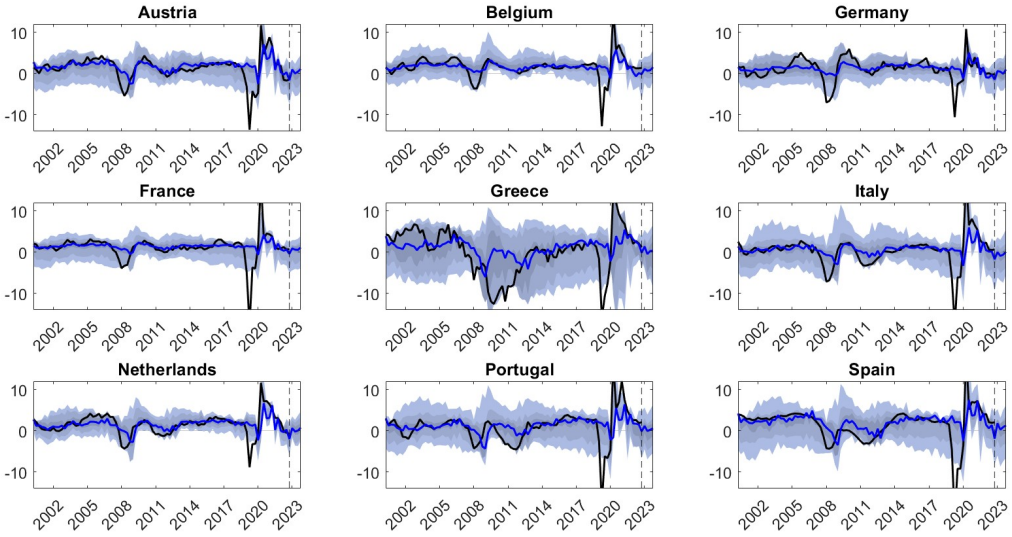

Figure A1.8: GDP growth conditional distributions, common factors (without Covid adjustment)

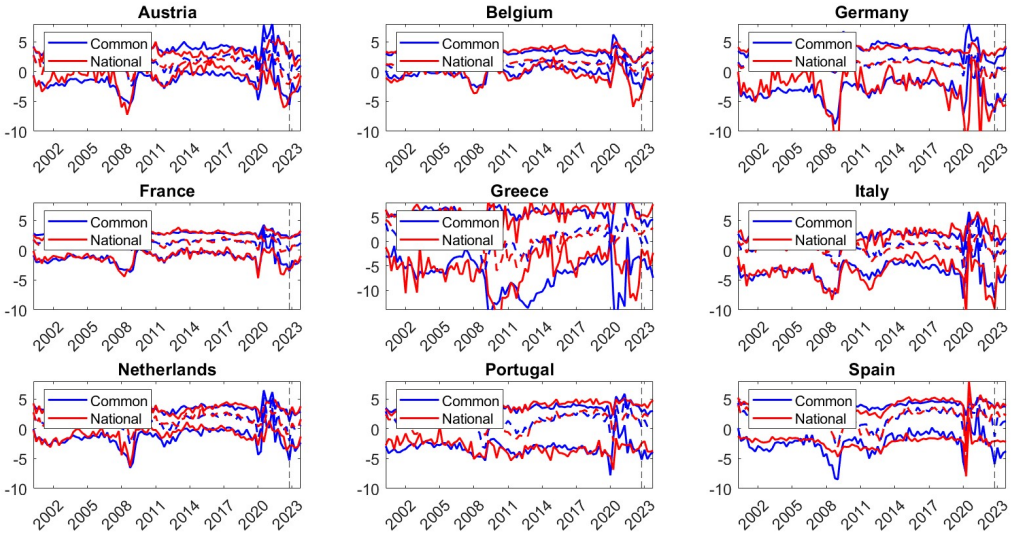

Figure A1.9: GDP growth conditional distributions, common and country-specific factors (percent)
Red lines: 5^{th}, 50^{th} and 95^{th} percentiles of the distributions conditional on EA factors (median in dashed lines).
Blue lines: 5^{th}, 50^{th} and 95^{th} percentiles of the distributions conditional on country-specific factors.

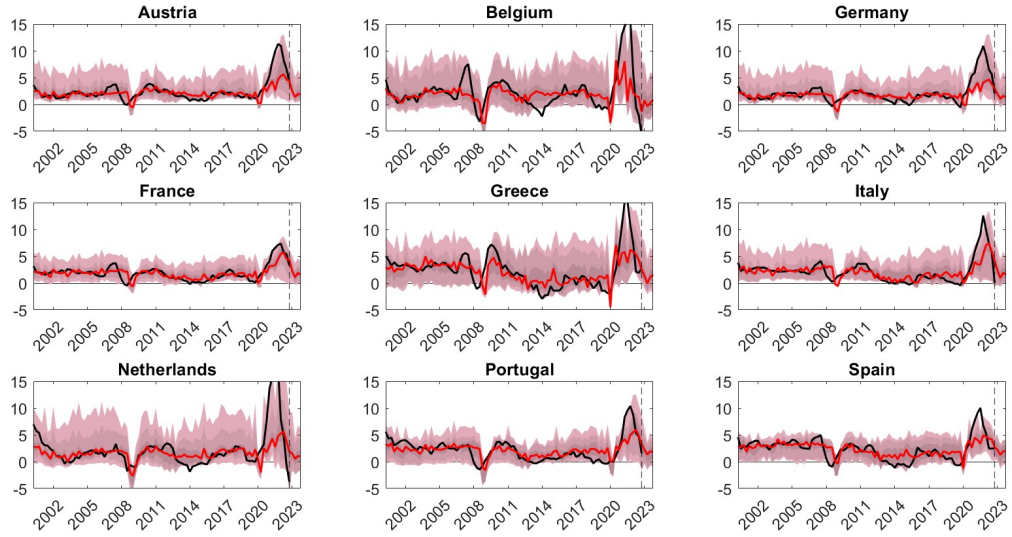

Figure A1.10: EA countries' inflation conditional distributions, common factors (percentage points)

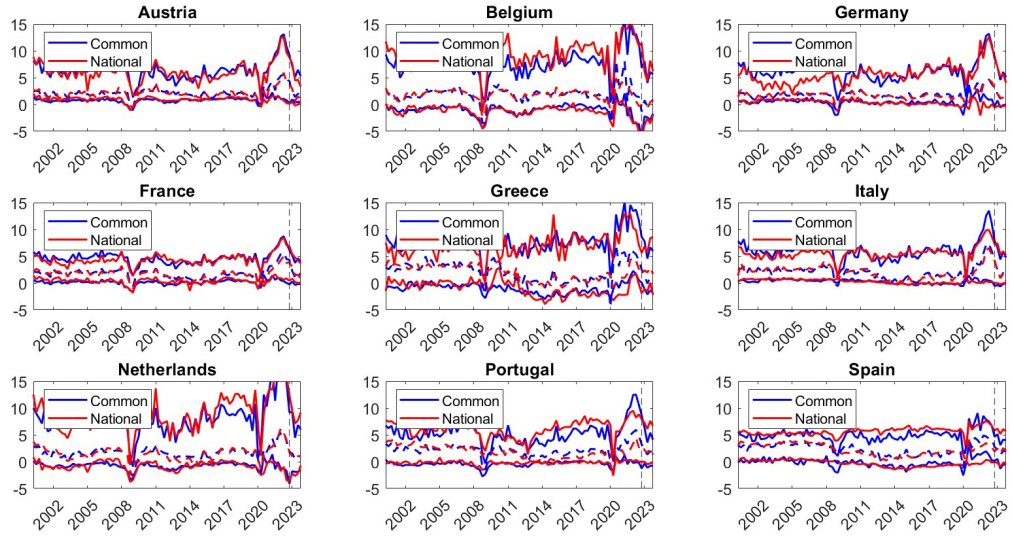

Figure A1.11: Inflation conditional distributions, common and country-specific factors (percent)
Red lines: 5^{th}, 50^{th} and 95^{th} percentiles of the distributions conditional on EA factors (median in dashed lines).
Blue lines: 5^{th}, 50^{th} and 95^{th} percentiles of the distributions conditional on country-specific factors.

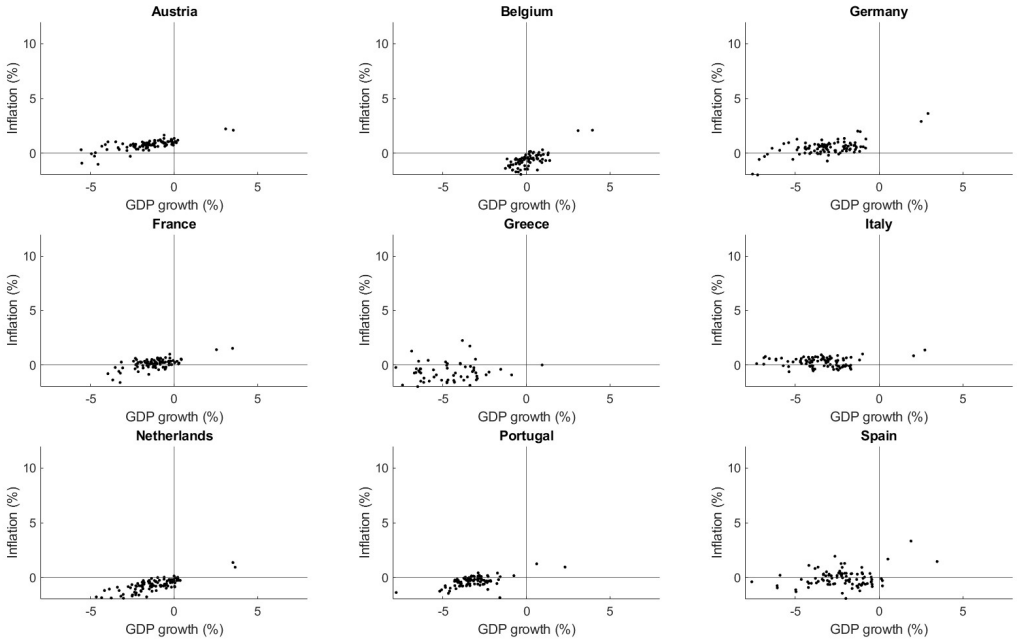

Figure A1.12: Conditional AS curves on the left tail, EA countries

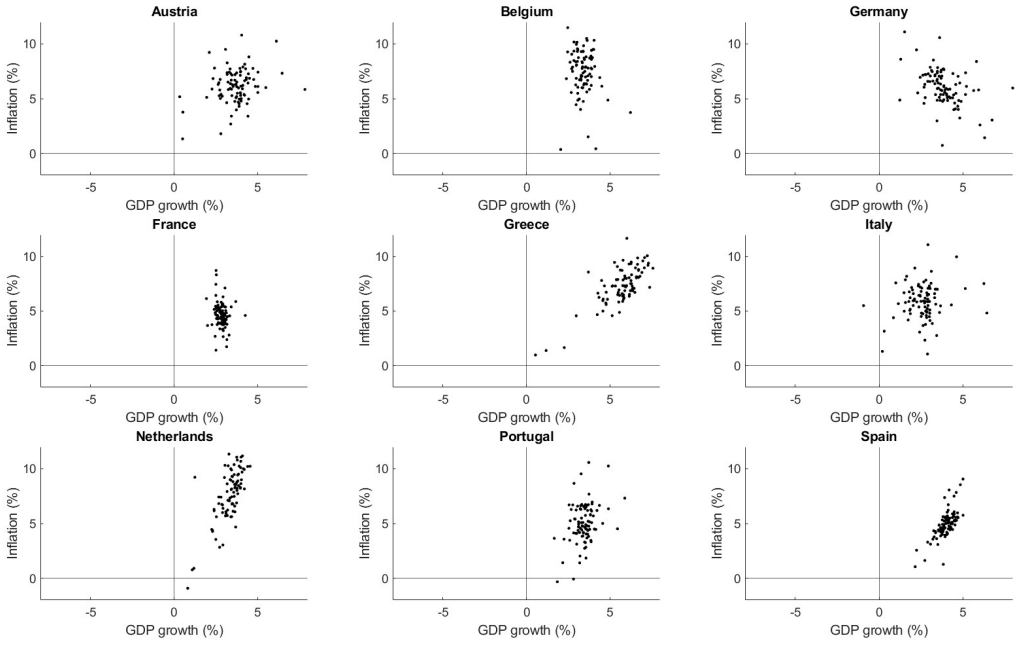

Figure A1.13: Conditional AS curves on the right tail, EA countries

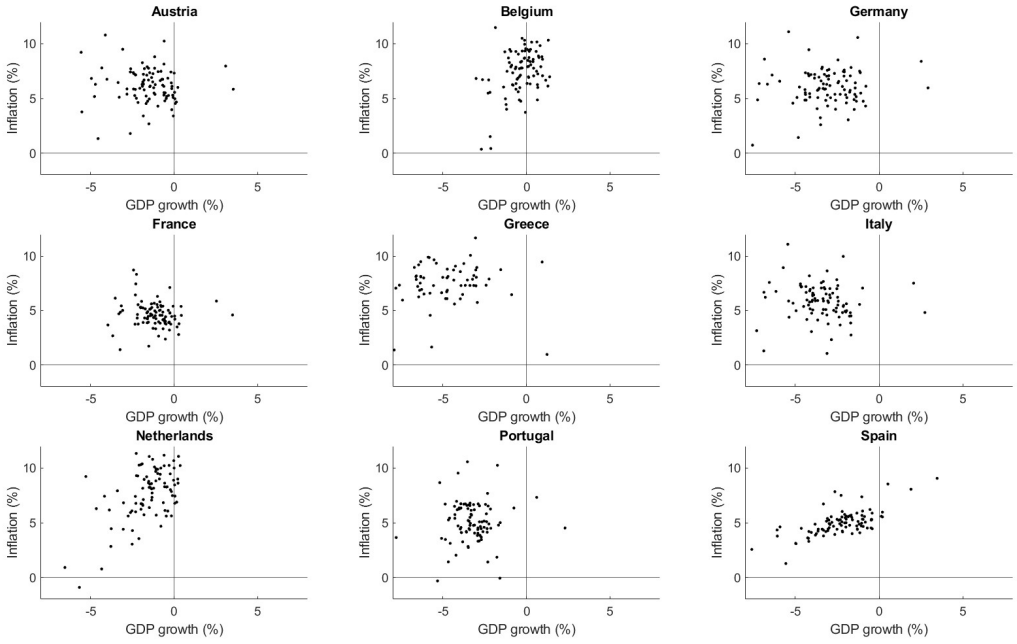

Figure A1.14: Conditional AD curves on the left tail, EA countries

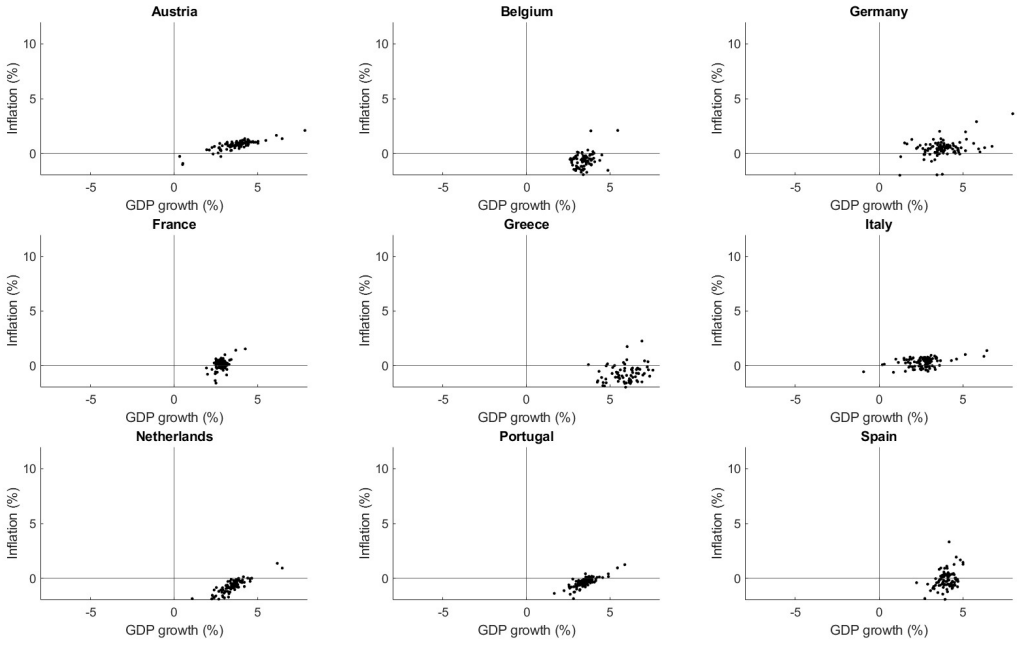

Figure A1.15: Conditional AD curves on the right tail, EA countries

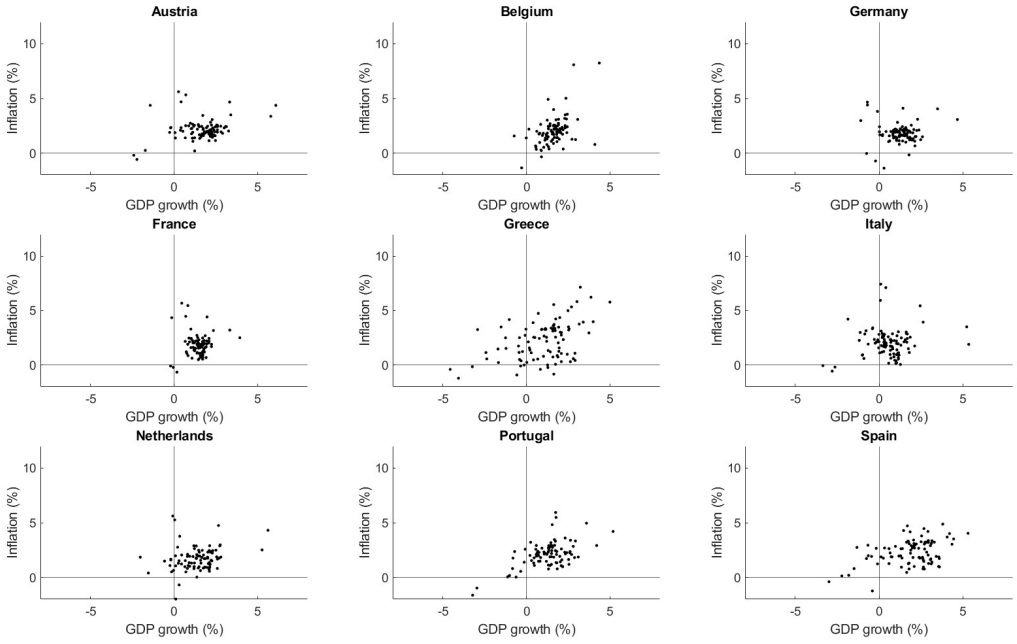

Figure A1.16: Conditional median of π and y, EA countries

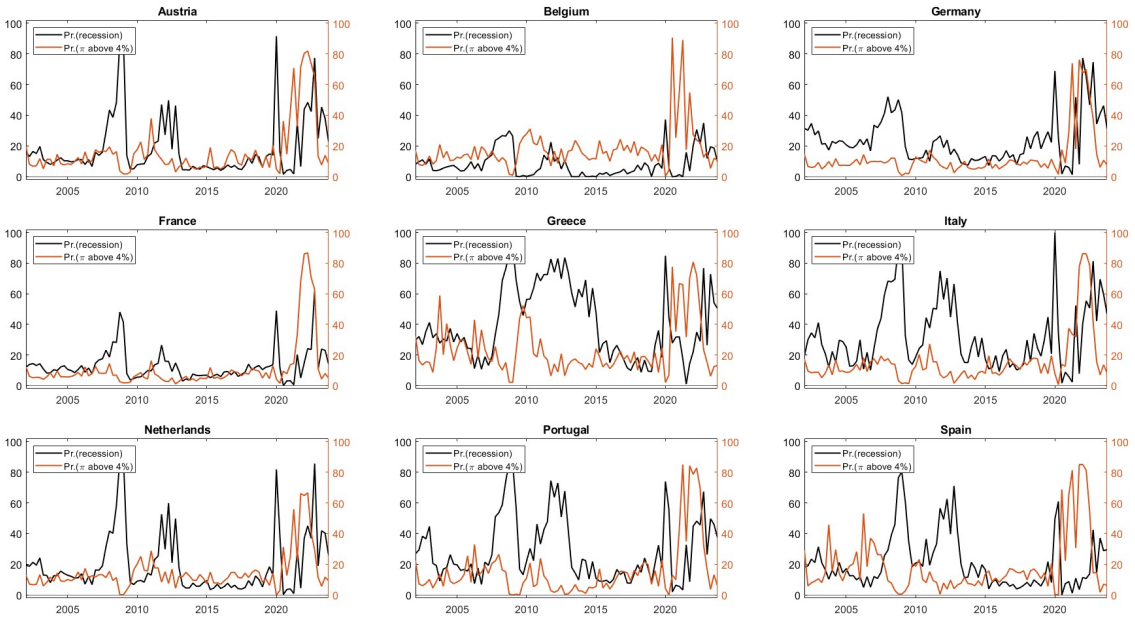

Figure A1.17: Conditional probabilities of recession and high inflation, EA countries (percent)

Chapter 2

Estimating the natural rate of interest in the Euro Area using the yield curve

The natural rate of interest r^* is an unobserved yet pivotal variable for the assessment of the stance of monetary policy. In this paper, I investigate its evolution in the Euro Area in the last twenty-five years by extracting a trend component from the term structure of risk-free interest rates. The employed approach allows to decompose the trend of the natural rate into two sub-trends, capturing stochastic discount factors and convenience yields. Results suggest that the decline of r^* observed in the analysed period is due to "quantity" imbalances between savings and investments, rather than to "quality" effects concerning the demand and supply of safe assets.

2.1 Introduction

The level of development in financial intermediation achieved in the present-day European economy has impressively changed the scope and the instruments of monetary policy. Quantities that were once treated as exogenous choice variables, and as such considered essential components of a central banker's toolbox, have been gradually put aside and replaced by others. Money supply, the chief instrument of monetary policy until the 1970s, has become obsolete in a context in which the role of private financial institutions in creating money has become more and more relevant. Economists and policymakers started to think of money supply as an endogenous variable and gave more weight to other instruments, deemed to be more efficient and handy: first and foremost, refinancing and deposit interest rates. In a setting in which policy rates play such an essential role, understanding how the real interest rate that is prevalent in the economy compares to the "natural" rate is crucial in assessing the stance of monetary policy. If the real rate is higher than the natural rate, the policy can be considered restrictive, while if the opposite is true the stance of monetary policy is expansionary. This assessment, however, is not as straightforward as it may seem: while the real rate that prevails on the market is virtually observed, the same cannot be said of the natural rate of interest r^* ($\bar{r}_t$ later in the paper). This theoretical construction (Wicksell (1898)) corresponds to the prevailing short term real interest rate in the absence of transitory perturbations, determined by the equilibrium of demand and supply of savings/investments, were output consistent with its natural level and inflation be constant. The level of the natural interest rate has severe implications for the effectiveness of monetary policy: when it is too low, and inflation expectations are low as well, as during the Great Recession, monetary policy is likely to be constrained by the zero lower bound (ZLB).

This paper aims to investigate the dynamics of the natural interest rate in the Euro Area in the last twenty-five years using the information provided by the term structure of interest rates, with flexible specifications that allow to overcome the non-linearities induced by the ZLB. The period under study is characterised by recessions stemming from causes of different nature (the subprime crisis, the sovereign debt crisis and the Covid crisis), addressed with diverse mixes of fiscal and monetary policy: a traditional response in the first case (anticyclical fiscal policy at national level, expansionary but conventional monetary policy), a mixed one in the second (restrictive fiscal pol-

icy, expansionary and unconventional monetary policy), and an unprecedented approach in the last episode (anticyclical fiscal policy with a moderate risk sharing at European level, unconventional monetary policy). The results presented in the paper confirm the long run fall of r^* already attested by a large body of literature.

The other subject of interest of this paper is the decomposition of the natural interest rate into two trend components: one that refers to stochastic discount factors, and therefore to the intertemporal choices of consumption/saving, and the other to the evolution of the premium for the safety and liquidity properties of safe assets (to which I will refer to as *convenience yields*, as in Krishnamurthy and Vissing-Jorgensen (2012) and Del Negro et al. (2017)). The dynamics of convenience yields are closely related to those in the demand and supply of safe assets: if there is an excess demand of safety, the premium increases, while the opposite occurs if there is abundance of safe assets. I find that most of the decline observed in the natural interest rate in Europe is due to trends in stochastic discount factors, and therefore to "quantity" effects, rather than to "quality" effects due to the scarcity of safe assets. These results, which are validated by alternative robustness checks, are particularly relevant at a time in which the European Union has started programs of fiscal policy at common fiscal level, whose financing requires the issuing of bonds for a total amount of around 800 billion euros between 2021 and 2026. A deeper understanding of the main drivers behind the dynamics of r^* is pivotal to gauge the channels through which this expansionary policy can affect it.

The paper is structured as follows: Section 2.2 presents the most relevant literature on the estimation of the natural rate of interest. The main empirical model used in the paper is introduced in Section 2.3, while Section 2.4 presents the extension that allows to disentangle the two subcomponents. The data and the estimation strategy are in 2.5; results are in 2.6. Section 2.7 concludes.

2.2 Literature Review

The seminal paper on the estimation of the natural rate of interest is that by Laubach and Williams (2003) (LW henceforth). In their work, the authors estimate a semi-structural model that includes a pseudo-IS curve and a Phillips curve in state space form. Output and inflation are

functions of the short term interest rate and of other unobserved states as the trend output level, the trend growth rate of productivity and the natural interest rate $\bar{r}_t$, the main object of interest. The extracted natural rate is found to vary significantly in the studied US sample, ranging from a maximum of 4 percent in the 1960s to a minimum of 2 percent in the mid 1990s. In a subsequent work, the authors analyse the effect of the Great Recession on $\bar{r}_t$ (Laubach and Williams (2016)), finding a sharp drop after 2007. Moreover, despite the quick economic recovery in the U.S., $\bar{r}_t$ showed very light signs of growth after 2015, suggesting that low rates were likely to prevail also in the future. Holston et al. (2017) found similar diminishing patterns after 2007 in Canada, in the United Kingdom and in the Euro Area using a similar methodology, a result confirmed by an analysis when taking into account the effect of the Covid pandemic (Holston et al. (2023)). Other authors refined the LW approach: Kiley (2015) employs Bayesian methods to a very similar model, and finds that that results are not robust to extensions of the model that include other potential demand shifters, such as corporate bond spreads or credit growth rates. Juselius et al. (2016) incorporate proxies for the aggregate leverage at market prices and for the private sector debt service burden: the resulting $\bar{r}_t$ is higher than the one obtained with a standard model, but the decreasing trend is confirmed. Lubik, Matthes, et al. (2015) relax some of the assumptions of the LW model by estimating a time-varying parameter VAR, and use the long-term conditional forecast of the real interest rate as a measure for $\bar{r}_t$, finding similar results. Finally, Fiorentini et al. (2018) note that the LW methodology delivers highly uncertain results when the IS curve or the Phillips curve are flat, however broadly confirming the results of Holston et al. (2017).

Alternatively, other authors have studied the evolution of $\bar{r}_t$ using financial-based approach, starting the strand of literature most closely related to this paper. Johannsen and Mertens (2016) estimate a trend-cycle model for interest rates and inflation that takes explicitly into account the presence of the effective lower bound: $\bar{r}_t$ is still found to be weakly decreasing after the 1990s, with considerable uncertainty on its level. Del Negro et al. (2017) specify a VAR with common trends on interest rates of different maturities and the inflation rate, estimating parameters with a Bayesian approach. A relevant contribution of this paper is the decomposition of the trend in $\bar{r}_t$ into a component that depends on stochastic discount factors and one related to convenience yields. They find that the drop of the natural rate observed after the 1990s in the U.S. is mostly due to an increase in con-

venience yields, while discount factors remain roughly constant. These results are confirmed by a DSGE analysis, in which $\bar{r}_t$ is the interest rate on a safe asset that would be observed in the absence of sticky prices and wages. Subsequent extensions of this class of models (Del Negro et al. (2019), Ferreira and Shousha (2023)) assess the existence of global trends in $\bar{r}_t$ and suggest that most of the decline in natural rates observed throughout the analysed advanced economies is driven by global trends in the convenience yields and by slower growth. A similar approach, which delivers analogous results, is followed by Christensen and Mouabbi (2024), who estimate $\bar{r}_t$ for the Euro Area from the term structure of the price of bonds with cash flows indexed to inflation.

Since the influential statement on the existence of a "saving glut" (Bernanke (2005)), a large literature has also investigated the possible explanations behind the drop in natural interest rates. The stubbornly low - or even negative - levels of the natural interest rate has led to suppose that the world economy had entered a long lasting period of *secular stagnation* (Summers (2014)), in which structural factors in the real economy (shrinking debt-financed investment, increasing savings and slower population growth) compress $\bar{r}_t$. Eggertsson et al. (2019) use a quantitative life-cycle to formalize the "new normal" hypothesis, in which the real interest rate needed to achieve full employment is permanently negative. The main drivers of this fall are related to aging and slower productivity growth, a view also supported by Brand et al. (2018), who also mention the role of stronger risk aversion and higher inequality. Rachel and Summers (2019) highlight the fact that natural rates decreased in advanced economies despite the presence of significant offsetting factors that absorbed private savings, such as higher government debt, pay-as-you-go pension systems and the insurance value of public health care programs. Del Negro et al. (2017) support a different explanation for the fall of $\bar{r}_t$, related to the shortage of safe assets and the subsequent increase of long-run levels of convenience yields. This paper contributes to this literature by both providing updated financial-based estimates of $\bar{r}_t$ in the Euro area and by taking a stance on this "quantity-quality" debate.

2.3 Model setup: extracting r^* from the yield curve

In the first part of the analysis, I exploit the information conveyed by the term structure of interest rates to estimate the level of the natural rate of interest $\bar{r}_t$. Elaborating on the approach of Del Negro et al. (2017), I specify a VAR with common trends in which the observed variables y_t are decomposed into a vector of trends $\bar{y}_t$ and a vector of cyclical components $\tilde{y}_t$:

$$y_t = \Lambda_t \bar{y}_t + \tilde{y}_t \tag{2.1}$$

The vector of observables y_t contains the time series of the nominal interest rates paid by safe and liquid bonds on five different maturities: 3 months (short-term), 2, 5, 7 and 10 years. I use this specification to extract the trend in the real rate of interest (the chosen measure of $\bar{r}_t$). In order to disentangle trends in real rates and time premia from trends in inflation, the vector contains a series for inflation π_t and one for inflation forecasts π_t^e: the full vector of observables is $y_t = (r_t^s, r_t^2, r_t^5, r_t^7, r_t^{10}, \pi_t, \pi_t^e)$. I assume that all the included nominal interest rate time series share the three trends in $\bar{y}_t = (\bar{r}_t, \bar{\pi}_t, \bar{tp}_t)$: one for the real rate, one for inflation and one for the time premium. The relationship between the observed time series y_t and the trends $\bar{y}_t$ is captured by the loadings Λ_t, a matrix with elements partly kept fixed due to the modelling approach - and partly left free for estimation, as explained more thoroughly in Section 2.5[1]. The first line of the model defines the nominal short term rate r_t^s as the sum of two trends (real rate and inflation), and of its cyclical component $\tilde{r}_t^s$:

$$r_t^s = \bar{r}_t + \bar{\pi}_t + \tilde{r}_t^s \tag{2.2}$$

In the row of Λ_t referring to equation 2.2, all parameters are fixed. The equation for the longest maturity interest rate (r_t^{10}) is modelled similarly as a sum of three trends plus a cyclical component.

$$r_t^{10} = \bar{r}_t + \bar{\pi}_t + \bar{tp}_t + \tilde{r}_t^{10} \tag{2.3}$$

Interest rates on medium maturities are analogously modelled as a sum of trends and cycles: the time-varying loadings μ_t^i that link observed rates and the trend on time premia are bounded between

[1]The model specification bears some similarity with dynamic term structure models with shifting endpoints (as in Bauer and Rudebusch (2020), among others).

zero and one. I impose this restriction by positing that in the long run the time premium of a shorter horizon bond cannot be larger than the one of a title with analogous safety and liquidity properties but a longer maturity[2].

$$r_t^2 = \bar{r}_t + \bar{\pi}_t + \mu_t^2 \ \bar{tp}_t + \tilde{r}_t^2 \tag{2.4}$$

$$r_t^5 = \bar{r}_t + \bar{\pi}_t + \mu_t^5 \ \bar{tp}_t + \tilde{r}_t^5 \tag{2.5}$$

$$r_t^7 = \bar{r}_t + \bar{\pi}_t + \mu_t^7 \ \bar{tp}_t + \tilde{r}_t^7 \tag{2.6}$$

Finally, inflation and inflation expectations are decomposed into a common trend $\bar{\pi}_t$ and a cyclical component specific to each series:

$$\pi_t = \bar{\pi}_t + \tilde{\pi}_t \tag{2.7}$$

$$\pi_t^e = \bar{\pi}_t + \tilde{\pi}_t^e \tag{2.8}$$

Trends are modelled as series of univariate indipendent random walks with a small variance Σ_t^e; cyclical components follow a stationary VAR process with time-varying autoregressive coefficients Φ_t and variance Σ_t^ϵ.

$$\bar{y}_t = \bar{y}_{t-1} + e_t \qquad e_t \ \sim \mathcal{N}(\mathbf{0}, \Sigma_t^e) \tag{2.9}$$

$$\tilde{y}_t = \Phi_t \ \tilde{y}_{t-1} + \epsilon_t \qquad \epsilon_t \ \sim \mathcal{N}(\mathbf{0}, \Sigma_t^\epsilon) \tag{2.10}$$

The described model setup allows to extract a trend component of the real interest rate, a common measure of the natural rate in the literature. The choice of the measures rests on an underlying assumption: as r^* is by definition the real rate consistent with output at its natural level, the long run trend of the real rate $\bar{r}_t$ cleaned of inflation trends and short-term fluctuations - captured by cyclical components - can be used as a measure of r^*. This approach is quite parsimonious in contrast to structural and semi-structural estimation frameworks of r^*, as very little assumptions on the relationship between real and financial variables are required[3].

[2]The estimated μ_t parameters are presented in Figure A2.6 in the Appendix.

[3]This advantage over macroeconomic-based approaches is particularly relevant if the Phillips Curve or the IS curve are flat, a cause of high uncertainty in the estimates of r^* (Fiorentini et al. (2018)). On the other hand, a drawback of this approach is that the estimated errors of some of the equations of the model are autocorrelated, potentially

2.4 Extension: stochastic discount factors and convenience yields

In the second part of the analysis I investigate empirically the drivers behind the long-term movements in the natural rate of interest, decomposing the trend component of the real rate $\bar{r}_t$ and disentangling two trends, one that refers to the stochastic discount factors and one capturing convenience yields. Section 2.4.1 presents two alternative explanations of the observed fall in $\bar{r}_t$, while the econometric specification employed to disentangle the two components is exposed in 2.4.2.

2.4.1 Two alternative explanations

The relationship between real interest rates, convenience yields and stochastic discount factors can be expressed using a Euler equation (as in Del Negro et al. (2017)) which considers risk premia, here comprised in the convenience yield component:

$$1 = \mathbb{E}_t[(1 + r_t)(1 + CY_t)M_{t+1}] \tag{2.11}$$

In Equation 2.11 r_t is the real rate paid by perfectly safe and liquid bond, CY_t is the convenience yield and M_{t+1} is the stochastic discount factor $\mathbb{E}_t\left[\beta\frac{u'(c_{t+1})}{u'(c_t)}\right]$. The real rate trend can be written in logarithmic terms as:

$$\bar{r}_t = \bar{m}_t - \bar{cy}_t \tag{2.12}$$

In Equation 2.12, $\bar{r}_t = log(1 + r_t), \bar{m}_t = -logM_t$ and $\bar{cy}_t = log(1 + CY_t)$. This relation clarifies the two explanations for the movements of $\bar{r}_t$ that are the object of this paper. The first one is related to stochastic discount factors, and therefore to intertemporal choices of saving and consumption. A fall in $\bar{r}_t$ can be due to an increase in savings supply which *ceteris paribus* causes the equilibrium real rate to fall, as in the left panel of Figure 2.1. An increase in the supply of savings can be due for instance to an aging population, increased inequality - as propensity to consume decreases as income rises, and therefore a more uneven distribution of income generates additional savings - or in general to any movement in the social preferences for patience that shifts savings supply to the right (Bernanke (2005)). A similar effect can be observed if this movement occurs on the demand side, as in the right panel of Figure 2.1. In this case, keeping the savings supply curve unchanged,

hinting at model misspecification. However, the conclusions on the estimated trend components of $\bar{r}_t$, $\bar{m}_t$ and $\bar{cy}_t$, which are robust to a series of robustness checks presented in the analysis, should not be severely affected.

the effect can be driven by a decrease in the demand for investment. The economic determinants can be related to demography (a slower population growth depresses investment), to technology (a slowdown in the TFP growth, as suggested by Summers (2014)), or to any fall in the will to invest (using Keynes' jargon, a crisis of the "animal spirits"). Both groups of explanations are related to changes in the *quantity* of desired savings and investment, regardless of the safety and liquidity properties of the traded assets.

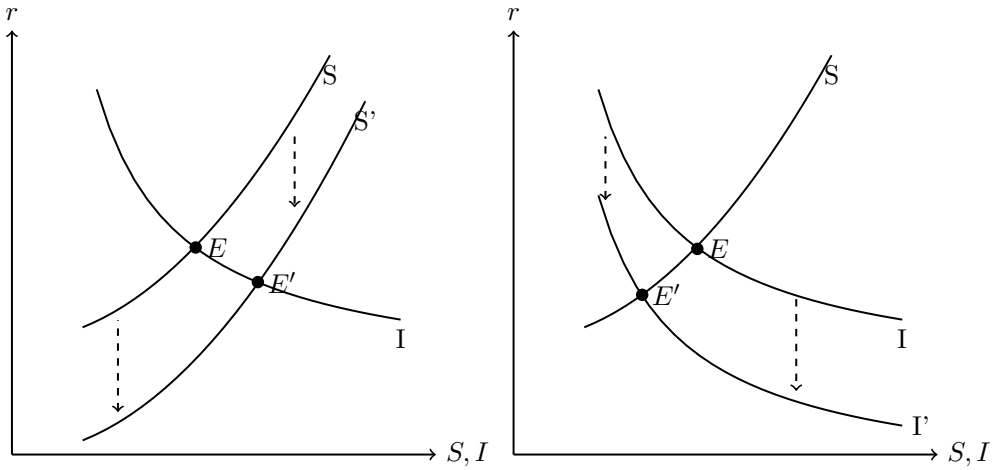

(a) Aging, inequality, patience.. (b) Slower pop. growth, weaker TFP..

Figure 2.1: Channels in the supply and demand for saving/investment

An alternative explanation behind the movements in the natural rate of interest is brought up by Del Negro et al. (2017), who put greater emphasis on the role of trends in convenience yields (cy), namely the premium paid by an investor for some specific benefits that holding a certain asset can guarantee. For instance, a safe asset can be used as collateral in borrowing activities, and for this reason its price can incorporate a convenience yield. The price of these desired characteristics of safety and liquidity is determined on the market, as the clearing price of the relative demand for safety, defined as the ratio between the demand for safe assets over the total demand for assets, and the relative supply of safety, defined analogously. The relative demand for safety is decreasing in the level of cy, while relative supply increases as the premium rises, as portrayed in Figure 2.2.

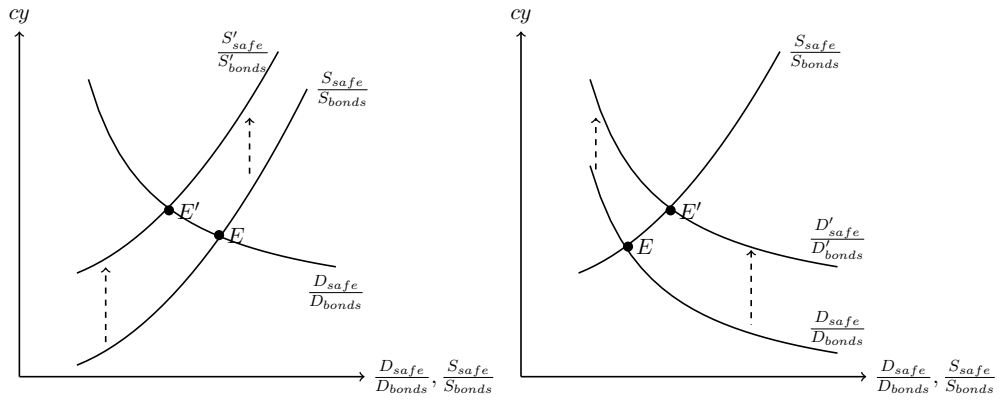

(a) A decrease in the supply of safety (b) An increase in the demand for safety

Figure 2.2: Channels in the supply and demand for safety/liquidity

The left panel of Figure 2.2 shows the effect of a deterioration of the average creditworthiness in the market, arising for instance in a context of financial distress. In the market for safety, if relative demand does not move in the meanwhile, relative supply shifts to the left. This causes the temporary excess demand to be cleared by an increase in the remuneration for the desired safety and liquidity properties, namely *cy*. As expressed by Equation 2.12, this causes the natural rate to diminish in the absence of counterbalancing movements on the "quantity" side. An equivalent conclusion can be reached going through the relative demand of safety, as shown in the right panel of Figure 2.12. It is common to observe *flight-to-quality* movements during periods of stress: large investors react to increases in volatility and uncertainty by selling assets deemed risky to buy safe bonds, commodities or currencies. This translates into an increase of the relative demand of safety, that moves to the right and causes convenience yields to rise if the relative supply does not shift to the right at the same time. This effect - that works through a channel related to the *quality* of traded assets rather than their *quantity* - eventually drives down the natural rate.

2.4.2 Identifying the two trends

In order to identify these two components, the model is modified as follows. Mirroring Equation 2.12, I replace the trend $\bar{r}_t$ with $\bar{m}_t - \bar{c}\bar{y}_t$ in Equations 2.2 - 2.6 (leaving Equations 2.7 and 2.8 unchanged) and include additional equations that allow to distinguish these components, by generalizing the approach of Krishnamurthy and Vissing-Jorgensen (2012) and Del Negro et al. (2017): in particular, I add to the y_t vector some series of interest rates paid by bonds without the safety and liquidity

features of the triple-A government bonds of the original model. I exploit the information carried by yields paid by private bonds with a lower rating. The difference between these series is captured by the trend in convenience yields $\bar{cy}_t$: while safe bonds incorporate them (Equation 2.13), the equation for risky ones (2.14) does not, so that the bond pays a higher premium. This part of the model reads as follows:

$$r_t^{10} = \bar{m}_t - \bar{cy}_t + \bar{\pi}_t + \bar{tp}_t + \tilde{r}_t^{10} \tag{2.13}$$

$$r_t^{Low} = \bar{m}_t + \bar{\pi}_t + \bar{tp}_t + \tilde{r}_t^{Low} \tag{2.14}$$

Moreover, I exploit the information provided by other series related to bonds with intermediate characteristics between the safe bond and the worst included in the model. Private bonds with average ratings have some of the properties of safe assets: for this reason they are modelled them assuming that the interest rate they pay is between the safe and the riskiest security.

$$r_t^{Mid} = \bar{m}_t + \lambda_t^{Mid}\bar{cy}_t + \bar{\pi}_t + \bar{tp}_t + \tilde{r}_t^{Mid} \tag{2.15}$$

The time-varying parameter λ_t, that links observed rates and the trend in convenience yields, is constrained to lie between zero (worst-rated bonds, no safety and liquidity) and -1 (safe assets, paying the lowest yield)[4]. Including this set of series allows to estimate more precisely the latent trends $\bar{m}_t$ and $\bar{cy}_t$, by expanding the information set at disposal. Section 2.6 and the Appendix present results obtained using different combinations of lower-than-AAA private bonds, as well as some robustness checks that make use of public bonds issued by peripheral European countries as Greece, Spain, Italy and Portugal. The resulting trends ($\bar{m}_t, \bar{cy}_t, \bar{\pi}_t, \bar{tp}_t$) delivered by the alternative specification are qualitatively consistent.

2.5 Data and Estimation

2.5.1 Data

The variable employed to measure the interest rate on riskless and liquid assets is the nominal yield paid by German government bonds at different points of the yield curve. These securities have

[4]The estimated λ_t parameters are presented in Figure A2.13 in the Appendix.

been constantly rated as triple-A by all agencies since the start of the analysed period and have the thickest market among the public bonds issued in the Euro Area, a feature that makes them particularly stable and liquid and justifies the presence of a convenience yield. The collected yields on German bonds refer to securities with 3-month, 2-, 5-, 7- and 10-year maturity[5].

For what concerns private bonds with lower ratings, I use the effective yield of the Euro High Yield Indices - provided by ICE Bank of America - rated as BB, B and CCC[6]. As an alternative, the interest paid by risky assets in the EA is measured by the yield of public bonds issued by peripheral European countries. For this analysis, I employ 10-year bonds issued by Spain, Italy and Greece, with additional robustness checks employing Portugal data. The data source for these time series is Eurostat. Finally, the employed measure of inflation is Eurostat HICP, while the measure of inflation forecasts is provided by the ECB's Survey of Professional Forecasters (SPF). All series are collected at quarterly frequency, and the model is estimated on the sample 2000Q1-2024Q4. Table A2.1 in the Appendix presents some descriptive statistics of the data employed in this analysis.

2.5.2 Estimation

The models described in Sections 2.3 and 2.4 can be conveniently rewritten in state space form. Using this representation, the observed variables y_t are modelled as functions of the vector α_t of unobserved states, composed in turn by the trend and cyclical components ($\alpha_t = (\bar{y}_t, \tilde{y}_t)'$). The vector α_t is a function of its past values through the time-varying matrix T_t; the model can be therefore expressed as a system of the measurement equation (2.16), where the matrix Z_t links y_t to the states α_t, and the transition equation (2.17):

$$y_t = Z_t \ \alpha_t + \varepsilon_t \qquad \varepsilon_t \sim \mathcal{N}(\mathbf{0}, \Sigma_t^{\varepsilon}) \tag{2.16}$$

$$\alpha_t = T_t \ \alpha_{t-1} + \eta_t \qquad \eta_t \sim \mathcal{N}(\mathbf{0}, \Sigma_t^{\eta}) \tag{2.17}$$

The error in Equation 2.17, η_t is composed by a trend and a cyclical component. Imposing that trends and cycles are independent components of the observed series, the covariance of the error Γ_t^{η}

[5]The source of data for the interest rates paid by German safe assets is Thomson Reuters.

[6]I carry out some robustness checks - which are presented in the Appendix - using private bonds with different ratings. Corporate bonds eligible to be included in these indices have at least one year of maturity and a minimum outstanding amount of Euro 100 million. The data on yields paid by low-rated private bonds are gathered from Thomson Reuters.

is by construction a block diagonal matrix:

$$\Sigma^{\eta} = \begin{bmatrix} \Gamma_t^e & \mathbf{0} \\ \mathbf{0} & \Gamma_t^\epsilon \end{bmatrix} \tag{2.18}$$

Trend components $\bar{y}_t$ are independent with respect to each other, so that the matrix Γ_t^e is diagonal; on the contrary, the variance matrix of the cyclical errors Γ_t^ϵ is full. The residual error matrix Γ^ε is kept fixed, so that the unknown elements of the system are the states α_t, the coefficients matrices Z_t and T_t and the block Γ_t^ϵ. The estimation of these unknown elements is performed by means of the Expectation-Maximisation (EM) algorithm (Sargent and Sims (1977), Watson and Engle (1983)). The algorithm works as if unobserved states α_t were observed, so to have a complete log-likelihood function $\ell_t^{(h)}$, and iterates between two steps: in the E-step, fixing the values of the parameters $\varphi_t^{(h)}$ which compose the matrices Z_t, T_t, Σ_t^e and Σ_t^η, the Kalman Smoother delivers an estimate of the expected value of the states $\mathbb{E}(\alpha_{t|T})$ that completes $\ell_t^{(h)}$; in the M-step, given $\mathbb{E}(\alpha_{t|T})$, the log-likelihood function is maximized to update the parameters $\varphi_t^{(h+1)}$. The procedure is repeated alternating these two steps until $\ell_t^{(h)}$ converges. This algorithm allows for constraints on the parameters φ_t of the kind described in Sections 2.3 and 2.4, and can also be adapted to time-varying parameters by applying the algorithm on a rolling window of observations[7].

Accounting for the effect of the ZLB

The time span that is object of the analysis includes approximately six years in which the ECB policy rate remained at zero. The main refinancing rate, the key object of conventional monetary policy, was lowered to 15 basis points in June 2014, then to 5 in September 2014 and to zero in March 2016. The impossibility to lower policy rates below zero can cause non-linearities to arise in the relationship between observed interest rates and long term drivers of the same. Although time-varying parameters are expected to accommodate the effect of non-linearities, I reinforce the approach in two additional ways. First, following the approach of Del Negro et al. (2017), I disregard the information provided by the short term interest rate r_t^s during the ZLB period, treating these observations as missing. In order to do so, I use the methodology for treating factor models with

[7]The bandwidth for the rolling window is of length $T^{1/2}$ (Giraitis et al. (2014)). I use a Gaussian kernel to give a greater weight to closer observations.

missing data developed by Bańbura and Modugno (2014). Their algorithm adapts the two-step EM algorithm: the E-step is modified to skip rows corresponding to missing observations when running the Kalman Filter; the M-step is adjusted in the updating equations of the matrices entering the measurement equation 2.16, Z_t and Σ_t^{ε}[8]. The second robustness check consists in modelling explicitly the non-linear relationship between r_t^s and $\bar{r}_t$ during the ZLB. In order to do so I specify a threshold factor model, as in Massacci (2017):

$$y_t = \mathbf{1}(x_{t-j} \leq \theta) Z_{1,t}\alpha_t + \mathbf{1}(x_{t-j} > \theta) Z_{2,t}\alpha_t + \varepsilon_t \tag{2.19}$$

$$\alpha_t = T_t \ \alpha_{t-1} + \eta_t \tag{2.20}$$

In the general form of the model, the loadings Z_t change depending on the value of an external variable x_{t-j} being greater or lower than the threshold θ. In the baseline specification, many parameters of the models are already time-varying, while the loading linking r_t^s and $\bar{r}_t$ is imposed by the model to be equal to 1 in Equation 2.2. This version of the model allows this loading to switch regime when the ECB policy rate hits a threshold θ. Equation 2.2 becomes:

$$r_t^s = \begin{cases} \bar{r}_t + \bar{\pi}_t + \tilde{r}_t^s & r_{t-j}^{ecb} \geq \theta \qquad \text{(ordinary times)} \\ \gamma_t \bar{r}_t + \bar{\pi}_t + \tilde{r}_t^s & r_{t-j}^{ecb} < \theta \qquad \text{(ZLB binds)} \end{cases} \tag{2.21}$$

The first line of Equation 2.21 corresponds to the "conventional" monetary policy condition, in which the policy rate can be lowered and the short term real rate is modelled as in the baseline model. When the policy rate r_t^{ecb} hits θ, the non-linear relationship between r_t^s and $\bar{r}_t$ is captured by the time-varying loading γ_t. This specification allows for an unknown threshold level and for the possibility of a lagged reaction of the system to the attainment of the ZLB, given by j. I estimate jointly the parameters of the state space model - including γ_t, - the threshold θ and the lag j using a combination of the EM algorithm and a gridpoint search method, as in Correal and Peña (2008)[9].

[8] More details on the algorithm of Bańbura and Modugno (2014) are provided in the Appendix.

[9] More details on the estimation algorithm of the threshold factor model (Correal and Peña (2008)) are provided in the Appendix.

2.6 Results

2.6.1 Model for the estimation of $\bar{r}_t$

This Section presents the estimates of $\bar{r}_t$ obtained using the models described in Section 2.3. Figure 2.3 reports the trend component of the real interest rate, estimated from a standard version of the model that includes five equations for safe interest rates - one for each maturity - and two for inflation (observed and expected)[10]. The effect of the ZLB is accounted for by means of the algorithm of Bańbura and Modugno (2014). At the beginning of the sample the estimated natural rate hovers around 1 percent, on a slowly declining trend; the fall accelerates at the beginning of the Great Financial Crisis, and hastens in 2011-2013, while $\bar{r}_t$ becomes negative. After reaching a minimum level of around -1.7 percent in 2018, the rate rebounds in the last years of the sample, reaching -0.2 percent in 2024. The estimated decline of $\bar{r}_t$ throughout the whole period amounts to about 1-1.5 percentage points[11].

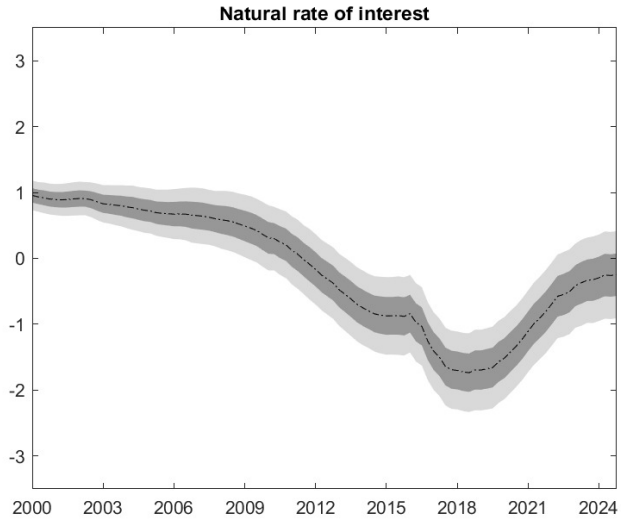

Figure 2.3: Natural rate - missing r_t^s (7 equations)

Figure A2.2 shows the results obtained by a model without the inflation forecasts π_t^e. Using this

[10] Confidence bands are obtained through a bootstrap algorithm described more in detail in the Appendix. Dark bands are 68% confidence bands, while the lighter ones are 95% bands.

[11] Figure A2.1 in the Appendix reports the results obtained with a standard version of the algorithm, thus neglecting the effect of the ZLB. The estimates corresponds to the baseline trend until 2015, before a slight divergence: the fall is somewhat slower in the first phase of the QE, and the model does not capture any recovery after the inflation surge and the series of policy rate hikes, as $\bar{r}_t$ stabilizes at a level close to -1.5 percent. The estimated fall over the whole period amounts to around 2-2.5 percentage points.

smaller model, the fall in $\bar{r}_t$ seems to be quite steady - without the slowdowns and the accelerations obtained using the full model - between 2002 and 2016. The qualitative conclusions on $\bar{r}_t$ are in line with those of the baseline model, but the estimated fall is a more intense (around -3 percent throughout the whole period)[12]. Figures A2.4 and A2.5 report the estimates obtained by explicitly modelling the non-linearity induced by the ZLB, by means of a threshold model (Correal and Peña (2008)). In Figure A2.4, the ZLB is assumed to affect structurally the relationship between the short rate r_t^s and the long run trend $\bar{r}_t$ with an immediate effect captured by the γ parameter at the (imposed) 0 basis points threshold for the ECB refinancing rate. The resulting r^* decreases slowly from 1.5 percent to around 1 percent in the pre-GFC periods, before a faster fall during the crisis. After 2015, the rate recovers by around 20 basis points and remains stable in the last quarters between -0.5 and -1 percent. The estimated fall in the whole time span amounts to about 2.5 percentage points. The extracted trend is very similar when allowing for an automatically selected threshold and a lagging effect of the ZLB on the α parameters: in this case (Figure A2.5), the estimated threshold γ remains at 0 basis points, while the effect of the non-linearity materializes after two quarters.

In Figure 2.4, the estimated natural rate is displayed with the *ex post* safe real rate of interest $(r_t^s - \pi_{t+1})$. In this exercise, the estimated $\bar{r}_t$ is employed of the assessment of the ECB monetary policy stance, by comparing the short term real rate to its natural counterpart. Periods where the real rate is above $\bar{r}_t$ can be labelled as "restrictive", as the central bank keeps rates above their long-run level consistent with output and inflation at natural level. On the contrary, when the real rate lies below $\bar{r}_t$, the monetary policy authority is pushing the economy beyond its natural level with "expansionary" policies, in order to overcome a recession or to limit deflation risks. The monetary policy stance cannot therefore be assessed by considering realized real rates separately from $\bar{r}_t$: a 0 percent real rate can either indicate a restriction or an expansion, depending on the level of the natural rate. The results show a moderately restrictive monetary policy in the very first years of the Euro. When the effects of the Dot-com bubble and of the Second Gulf War start

[12] Figure A2.3 shows the estimates obtained with a smaller model composed of six equations, where a smaller part of the yield curve is exploited. The qualitative conclusions that can be drawn from these models are quite in line with those delivered by the previous ones. The decline of $\bar{r}_t$ is particularly fast during the years of the GFC, and corresponds to a fall of around 2 percentage points.

to emerge in Europe, the policy turns expansive to support the economic cycle until 2006, when the ECB intervenes to slow down the first signs of inflation. The policy is restrictive until the beginning of 2009, the worst year of the GFC in Europe, at the start of the easing phase set to mitigate the effects of the crisis. The conventional policy stops being effective around 2013-2014, when real rates start to increase, driven by low demand and weak inflation dynamics. The ZLB impedes the possibility of further rate cuts to avoid deflation, in a context of fragmentation risks and weak economic activity. The implementation of unconventional monetary policies, together with a moderate economic growth, propels prices, lowering the real rate; the provided measure of stance indicates a moderately expansive monetary policy between 2017 and 2019. The policy is quite neutral in the first months of the Covid crisis, before becoming expansive at the beginning of the inflation surge in 2021. The rate hikes started in 2022, combined with the gradual disinflation process, lead the policy into restrictive territory since the second half of 2022.

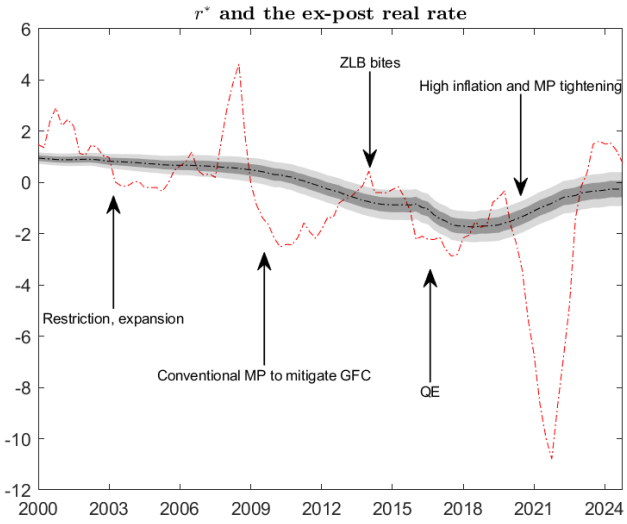

Figure 2.4: $\bar{r}_t$ and the real rate

2.6.2 Disentangling $\bar{m}_t$ and $\bar{cy}_t$

In this Section, I present the results of the extended model designed to separate the trends in stochastic discount factors $\bar{m}_t$ from those in the convenience yields $\bar{cy}_t$. Figure 2.5 shows the results obtained using a model that shares seven equations with the benchmark model for $\bar{r}_t$, extended with three equations for riskier private bonds, rated BB, B and CCC. According to these results,

obtaining by means of the algorithm of Bańbura and Modugno (2014), $\bar{m}_t$ decreased significantly during the years of the GFC, with a following a period of stability since 2018. The trend component $\bar{cy}_t$ seems to be quite stable in most of the sample. After 2018, $\bar{cy}_t$ started to decline, suggesting that the "quality" channel might even have slowed down the fall of $\bar{r}_t$ in part of the sample. Most of the fall of $\bar{r}_t$ identified in the previous Section seems to be explained by quantity factors - imbalances between the aggregate demand of investments and the supply of savings during the crisis -, rather than by safety factors.

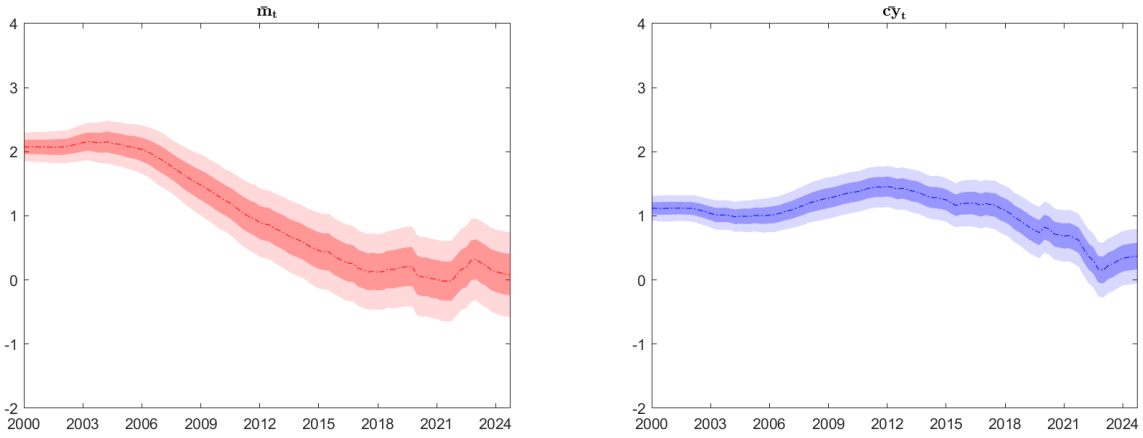

(a) Stochastic discount factors

(b) Convenience yields

Figure 2.5: $\bar{m}_t$ and $\bar{cy}_t$ - Model with private bonds rated BB, B and CCC

The conclusions on the prevalence of $\bar{m}_t$ are robust to a series of robustness checks[13]. Figure A2.8 reports the results obtained using private bonds with a slightly higher credit rating (AA, BBB and BB). Similarly to the benchmark model, $\bar{m}_t$ decreases significantly during the GFC, and rises slightly after 2020. Fluctuations in convenience yields are weaker for these safer bonds, with a small peak around the GFC and an overall steady behaviour. Results are similar also when using safer private bonds. Figure A2.9 shows the trends estimated with AAA, AA and BB bonds: the decreasing trend of $\bar{m}_t$ is evident since the beginning of the GFC, while $\bar{cy}_t$ is estimated to be substantially stable. In conclusion, the different model specifications point at $\bar{m}_t$ as the main driver of trends in real rates.

The prominence of $\bar{m}_t$ seems clear also when the model is specified with different series for dis-

[13]Results are virtually unaffected when using a standard EM algorithm that ignores the effect of the ZLB on the relationship between r_t^s and $\bar{r}_t$. Figure A2.7 shows the results obtained using this simpler procedure, that are almost equivalent to those of Figure 2.5.

entangling the two sub-trends, that is interest rates paid by bonds issued by peripheral countries of the Euro Area, perceived as riskier and less liquid especially in times of crisis. Figure 2.6 displays the estimates obtained with a model including the nominal interest rates on Spanish, Italian and Greek bonds, confirming the decreasing trend of $\bar{m}_t$ as the main driver of $\bar{r}_t$[14]. Using different combinations of public bonds, qualitative conclusions are not affected. Results are similar when excluding Greek bonds, whose risk - and yields - underwent large changes during the analysed period, as shown by Figure A2.11. Figure A2.12 displays the results obtained using Portuguese, Italian and Greek bonds: the size of the increase of $\bar{cy}_t$ is significantly smaller than the fall of $\bar{m}_t$.

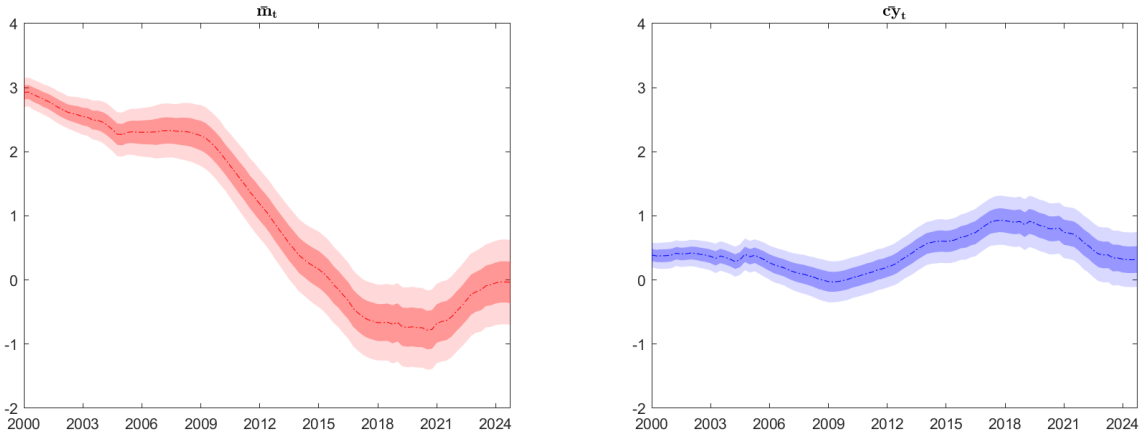

(a) Stochastic discount factors

(b) Convenience yields

Figure 2.6: $\bar{m}_t$ and $\bar{cy}_t$ - Model with ES, IT and GR public bonds

The evidence presented in this Section supports an explanation of the fall of $\bar{r}_t$ that highlights the role of "quantity" factors, rather than that of desired and supplied quality. This finding is in line with the note of Williams and Wu (2017), and more in general with the literature employing real variables approaches to estimate $\bar{r}_t$ (Laubach and Williams (2003), Holston et al. (2017) among others). This result is relevant for the design of fiscal policy at European level, as well as for the channels through which this latter can affect $\bar{r}_t$. Fiscal policy has a relevant role for the determination of the natural rate of interest, as argued by Rachel and Summers (2019), as it impacts $\bar{r}_t$ by absorbing private and public savings and affecting social preferences on intertemporal consumption. At the same time, a fiscal policy at centralized level can impact $\bar{r}_t$ also through a "quality" channel, by increasing the

[14]Using the standard estimation approach that ignores the ZLB delivers equivalent results, as shown by Figure A2.10 in the Appendix.

supply of safe assets, thus reducing $\bar{cy}_t$. The results of this Section, however, show that as most of the fall of $\bar{r}_t$ in the EA is not explained by safety issues, it is reasonable to surmise that there is little room for reducing $\bar{cy}_t$. On the contrary, there is room for absorbing excess savings through fiscal expansions, so to bring $\bar{r}_t$ back to its pre-crisis levels. The two channels are closely related, since a fiscal expansion financed issuing risky bonds can increase $\bar{m}_t$ and $\bar{cy}_t$ at the same time, but the main channel through which fiscal policy can increase $\bar{r}_t$ at the current state is the adoption of programs of public investment that absorb private savings.

2.6.3 By-products: estimates of trend inflation and time premia

This Section presents two by-products of the models used in the analysis: an estimate of the trend component of inflation, $\bar{\pi}_t$, common to all the series included in the y_t vector, and an estimate of the trend component of time premia $\bar{tp}_t$, linked to the observed nominal rates in y_t by the time-varying loadings μ_t^i. Figure 2.7 displays the trend inflation obtained from a standard version of the model estimated using the algorithm of Bańbura and Modugno (2014): the series hovers around one percent in the first ten years of the sample. Then, $\bar{\pi}_t$ displays some fluctuations between 2013 and 2017, before starting to increase after 2018, reaching 2 percent around 2021. The long run level of $\bar{tp}_t$, reported in Figure 2.8 seems quite stable in most of the sample, hovering around 0.5 percent, before increasing significantly in the last part of the sample[15].

[15]Figure A2.14 and A2.15 reports the trend estimated with the standard algorithm.

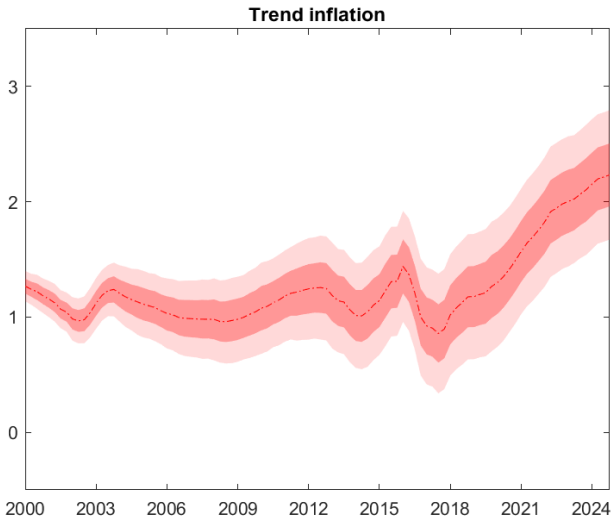

Figure 2.7: Trend inflation

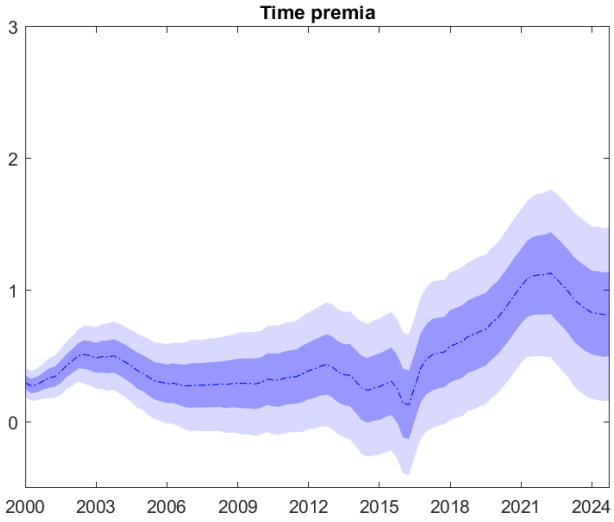

Figure 2.8: Time premia

2.7 Conclusions

This paper attempted to answer two questions: first, what can the yield curve reveal about the behaviour of the natural rate of interest in the Euro Area in the last twenty-five years?; second, are long run movements of r^* explained by "quantity" or "quality" dynamics?

For what regards the first question, the results show that r^* - net of some model uncertainty -

declined by about 1.5 percent in the EA since the early 2000s, with a pronounced acceleration during the Great Financial Crisis and a weak recovery after the pandemic. The time-varying level of r^* allows to interpret the stance of the ECB monetary policy in the analysed years: after a period of expansive policies in the first years of the sample, a restriction followed in 2006-2008; the ECB reacted to the GFC with an expansionary policy, that turned ineffective when the ZLB became binding; finally, the results show the effectiveness of QE in changing radically the monetary policy stance in recent years and the strength of the policy tightening cycle of 2022-23. As far as the second question is concerned, the results presented in the paper emphasize the role of imbalances in the supply and demand for savings/investment, significantly more relevant than the one played by safety issues. This finding has relevant implications in the understanding of the channels through which a common European fiscal policy can affect r^*, since most of the room for manoeuvre is on the "quantity" side, and namely in the possibility of absorbing excess savings. These conclusions are robust to various alternative model specifications and estimation approaches.

In conclusion, this paper confirmed the fall of the European natural rate of interest, already identified in the literature, and shed light on some of its possible determinants. Low levels of r^* cause issues for monetary policy and financial stability, but are not irreversible (Rachel and Summers (2019)) since policymakers can affect it with appropriate choices; this paper provides some insights on the channels through which fiscal policy can aspire to affect the natural rate.

A2 Appendix

A2.1 Additional results for $\bar{r}_t$

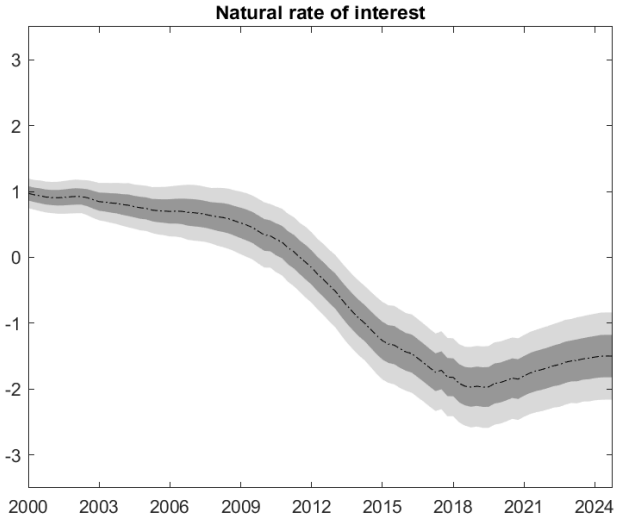

Figure A2.1: Natural rate - standard algorithm (7 equations)

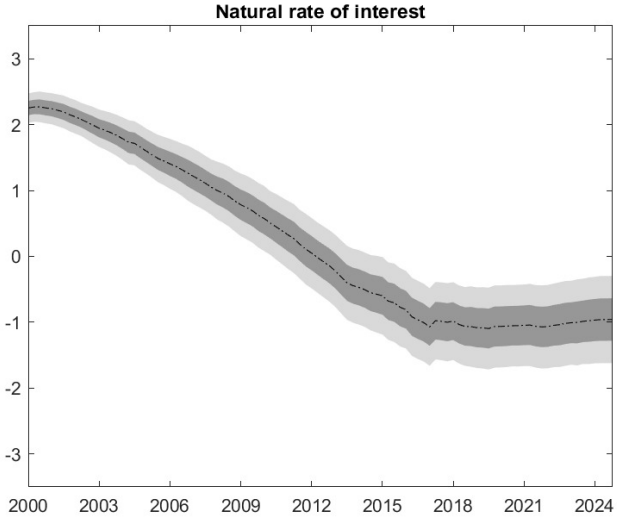

Figure A2.2: Natural rate - model without π_t^e

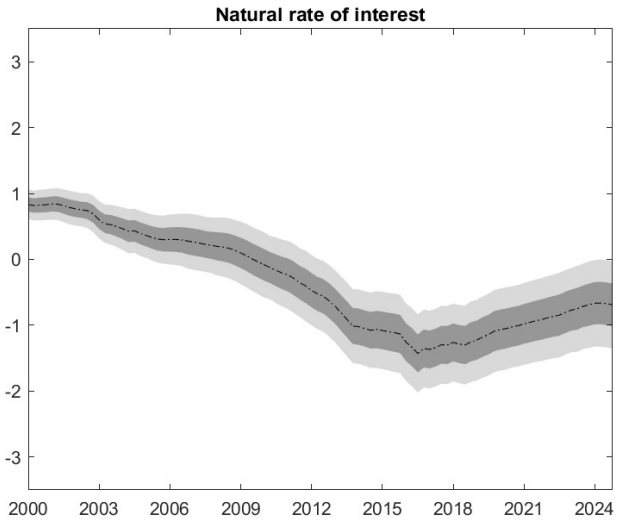

Figure A2.3: Natural rate - missing r_t^s (6 equations)

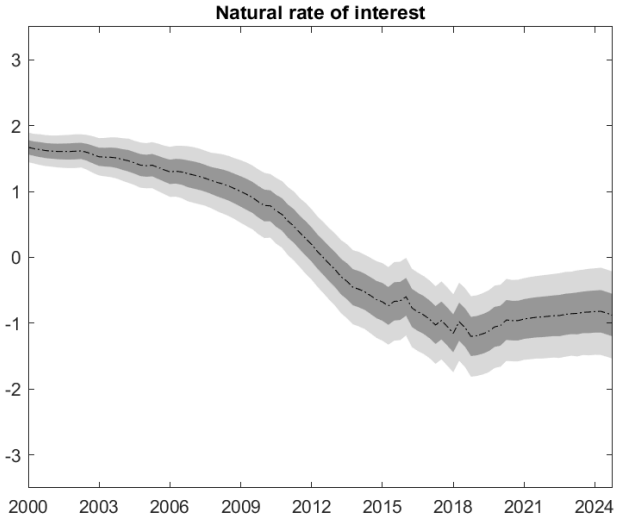

Figure A2.4: Natural rate - Threshold model

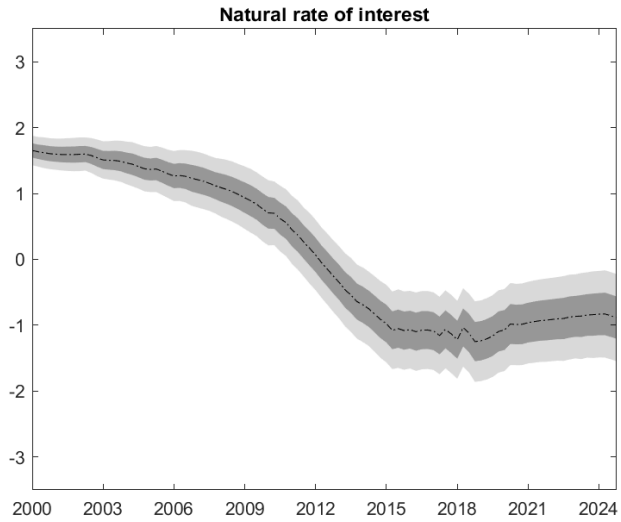

Figure A2.5: Natural rate - Threshold model with automatic selection of θ and number of lags

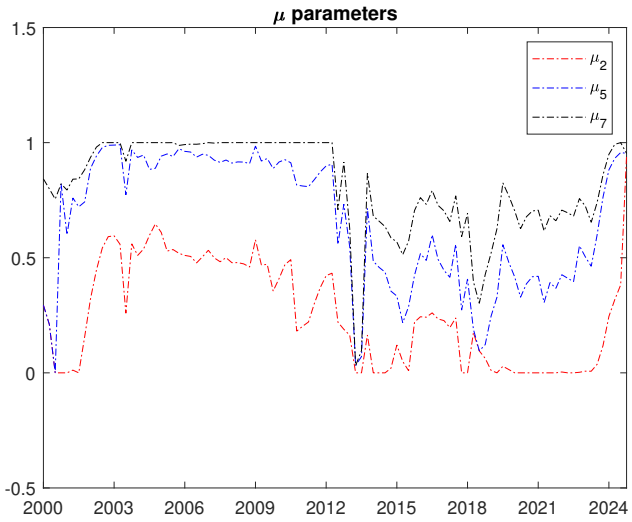

Figure A2.6: Estimated μ_t parameters (7 equations)

A2.2 Additional results for $\bar{m}_t$ and $\bar{cy}_t$

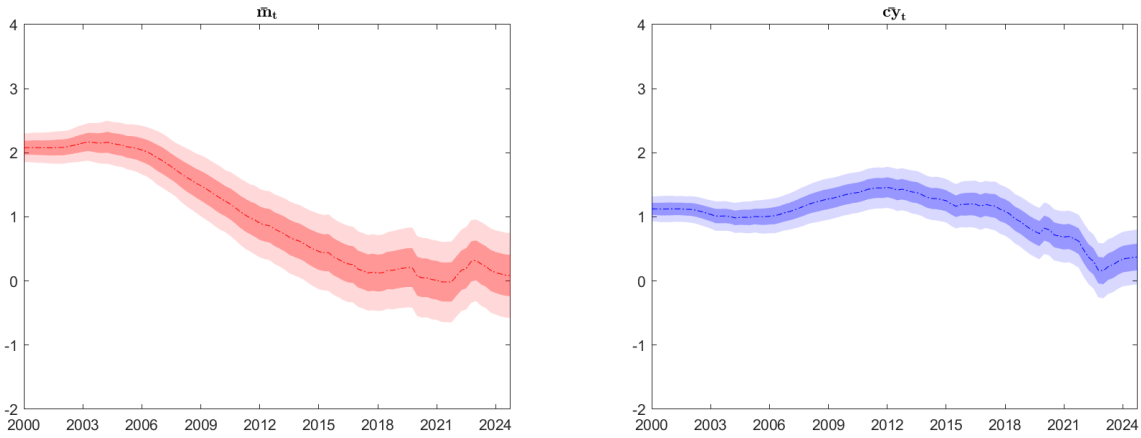

(a) Stochastic discount factors (b) Convenience yields

Figure A2.7: $\bar{m}_t$ and $\bar{cy}_t$ - Model with private bonds rated BB, B and CCC (non-missing r_t^s)

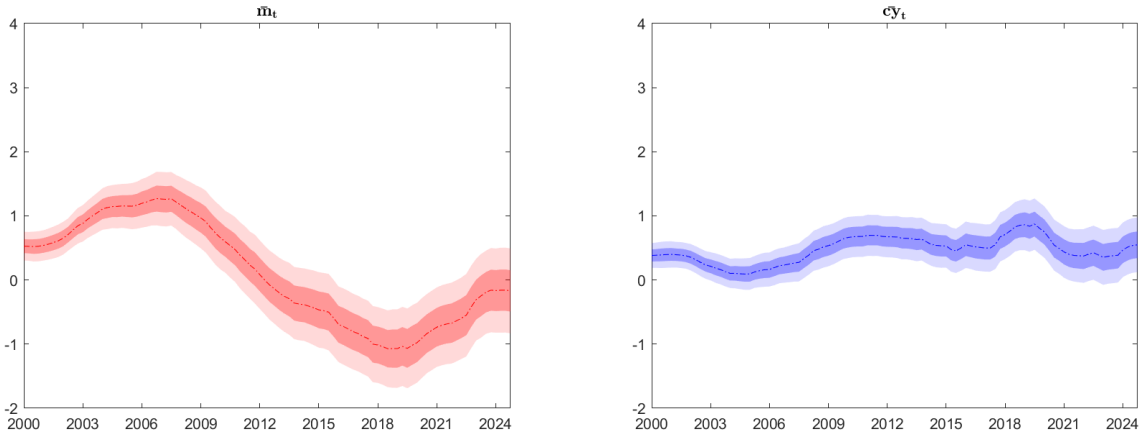

(a) Stochastic discount factors (b) Convenience yields

Figure A2.8: $\bar{m}_t$ and $\bar{cy}_t$ - Model with private bonds rated AA, BBB and BB

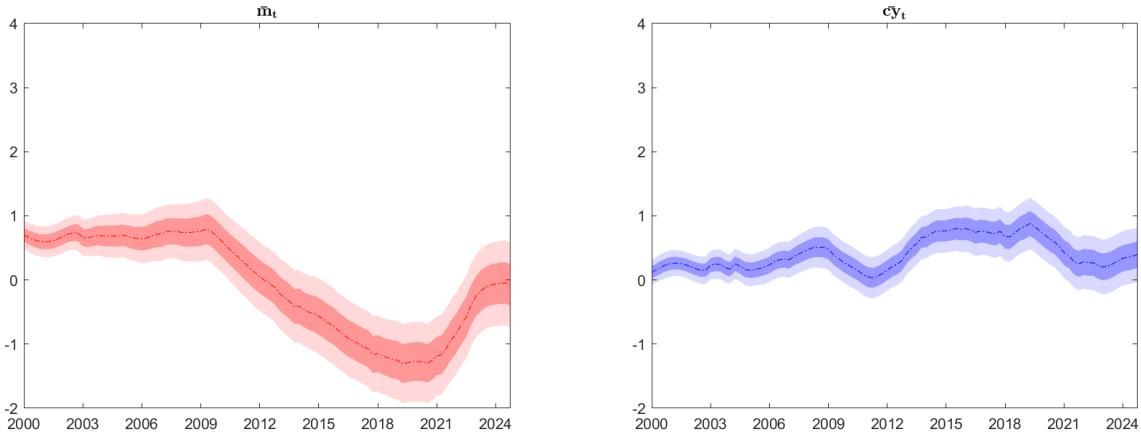

(a) Stochastic discount factors (b) Convenience yields

Figure A2.9: $\bar{m}_t$ and $\bar{cy}_t$ - Model with private bonds rated AAA, AA and BB

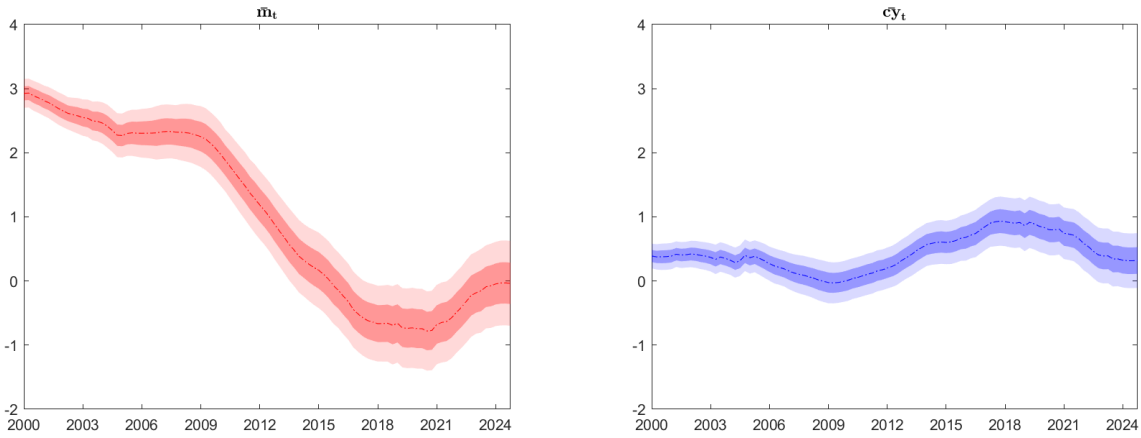

(a) Stochastic discount factors (b) Convenience yields

Figure A2.10: $\bar{m}_t$ and $\bar{cy}_t$ - Model with ES, IT and GR public bonds (non-missing r_t^s)

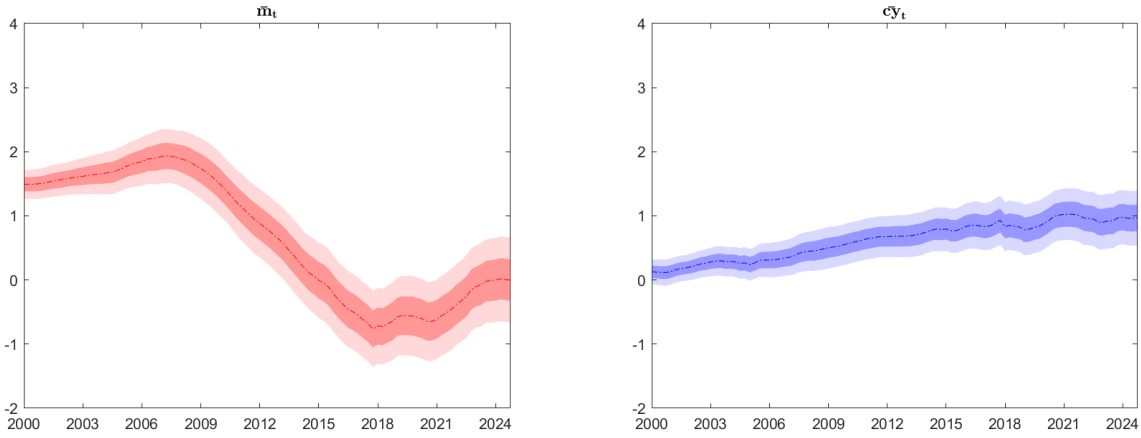

(a) Stochastic discount factors (b) Convenience yields

Figure A2.11: $\bar{m}_t$ and $\bar{cy}_t$ - Model with PT, ES and IT public bonds

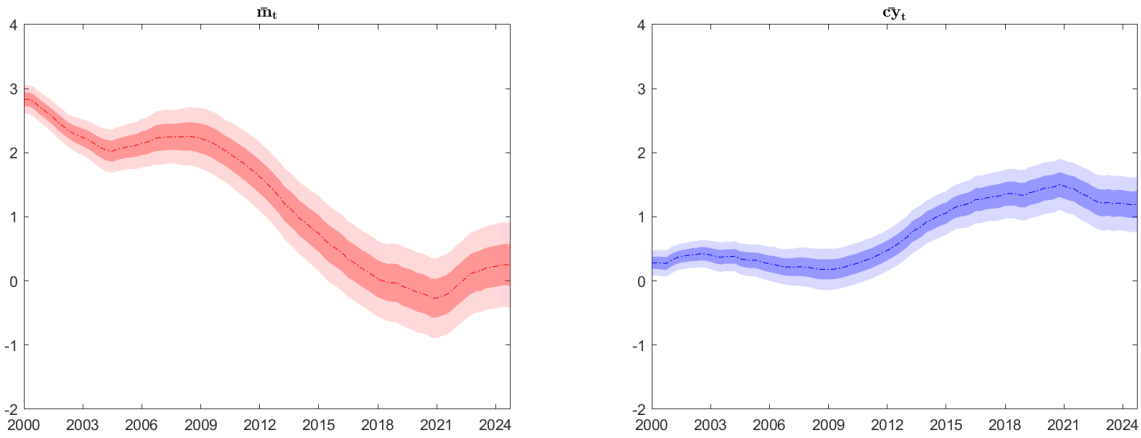

(a) Stochastic discount factors (b) Convenience yields

Figure A2.12: $\bar{m}_t$ and $\bar{cy}_t$ - Model with PT, IT and GR public bonds

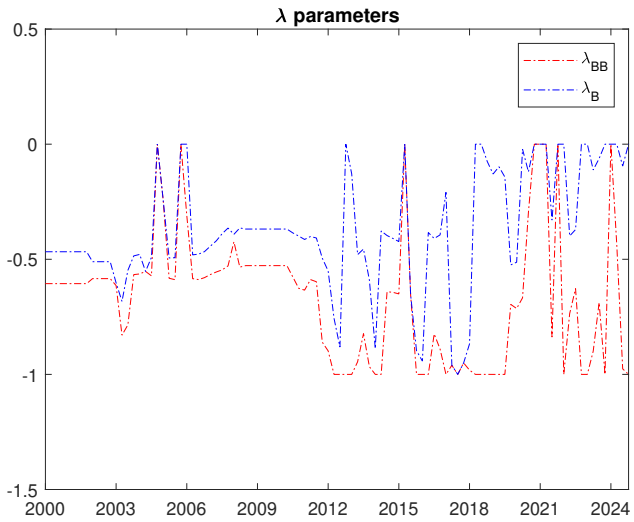

Figure A2.13: Estimated λ_t parameters - Model with private bonds rated BB, B and CCC

A2.3 Other results

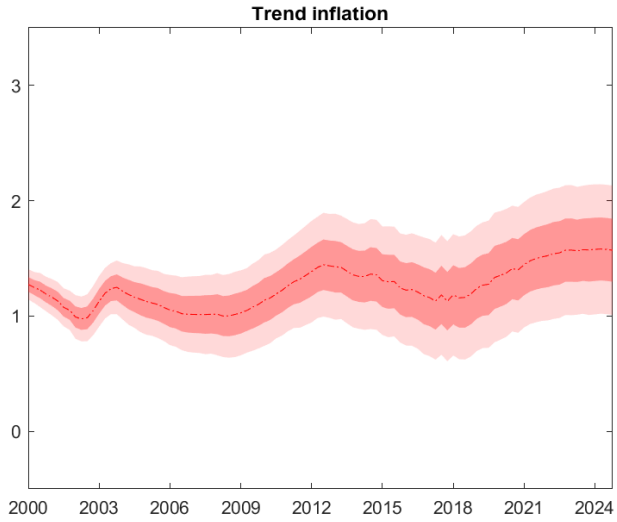

Figure A2.14: Trend inflation - standard algorithm (7 equations)

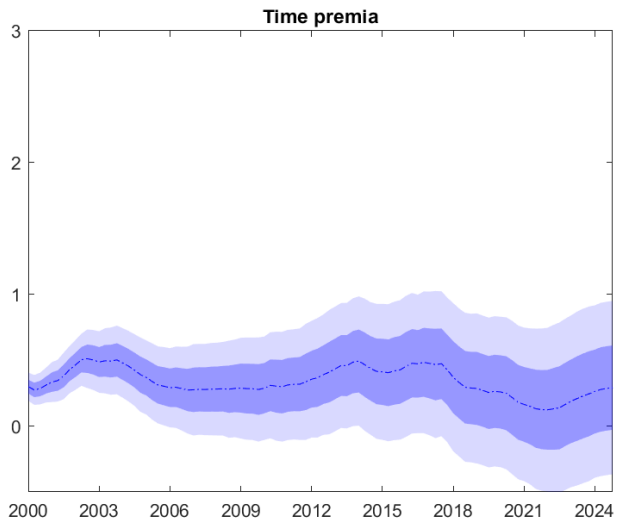

Figure A2.15: Time premia - standard algorithm (7 equations)

Table A2.1: Summary Statistics

	Mean	Std	Min	Max	Source
3-Month Bond Yield, Germany	1.31	1.85	-0.91	5.02	Thomson Reuters
2-Year Bond Yield, Germany	1.45	1.83	-0.80	5.13	Thomson Reuters
5-Year Bond Yield, Germany	1.80	1.85	-0.80	5.13	Thomson Reuters
7-Year Bond Yield, Germany	2.02	1.86	-0.73	5.36	Thomson Reuters
10-Year Bond Yield, Germany	2.30	1.80	-0.57	5.46	Thomson Reuters
10-Year Bond Yield, Greece	6.22	4.70	0.70	25.40	Eurostat
10-Year Bond Yield, Italy	3.62	1.45	0.64	6.61	Eurostat
10-Year Bond Yield, Portugal	4.04	2.51	0.09	13.22	Eurostat
10-Year Bond Yield, Spain	3.31	1.67	0.10	6.43	Eurostat
AAA Euro Corporate Bonds Index, BofA	2.51	1.64	-0.09	5.70	Thomson Reuters
AA Euro Corporate Bonds Index, BofA	2.68	1.86	-0.10	5.97	Thomson Reuters
A Euro Corporate Bonds Index, BofA	3.11	2.02	0.18	8.27	Thomson Reuters
BBB Euro Corporate Bonds Index, BofA	3.71	2.06	0.52	8.87	Thomson Reuters
BB Euro High Yield Bonds Index, BofA	6.02	3.39	2.13	18.26	Thomson Reuters
B Euro High Yield Bonds Index, BofA	8.52	3.82	3.89	25.77	Thomson Reuters
CCC Euro High Yield Bonds Index, BofA	16.05	9.80	6.81	48.72	Thomson Reuters
HICP Inflation Rate, Euro Area	2.14	1.81	-0.37	9.97	ECB
SPF Inflation Rate Forecast, Euro Area	1.73	0.57	0.80	4.80	ECB

A2.4 Models in SSF

Standard model for the estimation of $\bar{r}_t$

The standard model in state space form is the following:

$$y_t = \begin{bmatrix} r_t^s \\ r_t^2 \\ r_t^5 \\ r_t^7 \\ r_t^{10} \\ \pi_t \\ \pi_t^e \end{bmatrix} \qquad Z_t = \begin{bmatrix} 1 & 1 & 0 & \\ 1 & 1 & \mu_t^2 & \\ 1 & 1 & \mu_t^5 & \\ 1 & 1 & \mu_t^7 & \mathbf{I_{7\times 7}} \\ 1 & 1 & 1 & \\ 0 & 1 & 0 & \\ 0 & 1 & 0 & \end{bmatrix} \qquad \alpha_t = \begin{bmatrix} \bar{r}_t \\ \bar{\pi}_t \\ \bar{tp}_t \\ \tilde{r}_t^s \\ \tilde{r}_t^2 \\ \tilde{r}_t^5 \\ \tilde{r}_t^7 \\ \tilde{r}_t^{10} \\ \tilde{\pi}_t \\ \tilde{\pi}_t^e \end{bmatrix}$$

$$T_t = \begin{bmatrix} \mathbf{I_{3\times 3}} & \mathbf{0_{3\times 7}} \\ & \\ \mathbf{0_{7\times 3}} & \mathbf{\Phi_{t,7\times 7}} \end{bmatrix} \qquad \Sigma_t^{\eta} = \begin{bmatrix} \sigma^2_{\bar{r}_t} & 0 & 0 & \\ 0 & \sigma^2_{\bar{\pi}_t} & 0 & \mathbf{0_{3\times 7}} \\ 0 & 0 & \sigma^2_{\bar{tp}_t} & \\ & \mathbf{0_{7\times 3}} & & \mathbf{\Sigma_{t,7\times 3}} \end{bmatrix} \qquad \Sigma_t^{\varepsilon} = \sigma^2_{\varepsilon,t} \bullet \begin{bmatrix} \mathbf{I_{7\times 7}} \end{bmatrix}$$

Extended model for the estimation of $\bar{m}_t$ and $\bar{cy}_t$

The extended model in state space form is the following:

$$y_t = \begin{bmatrix} r_t^s \\ r_t^2 \\ r_t^5 \\ r_t^7 \\ r_t^{10} \\ r_t^{BB} \\ r_t^B \\ r_t^{CCC} \\ \pi_t \\ \pi_t^e \end{bmatrix} \qquad Z_t = \begin{bmatrix} 1 & -1 & 1 & 0 & \\ 1 & -1 & 1 & \mu_t^2 & \\ 1 & -1 & 1 & \mu_t^5 & \\ 1 & -1 & 1 & \mu_t^7 & \\ 1 & -1 & 1 & 1 & \mathbf{I_{10\times 10}} \\ 1 & \lambda_t^{BB} & 1 & 1 & \\ 1 & \lambda_t^B & 1 & 1 & \\ 1 & 0 & 1 & 1 & \\ 0 & 0 & 1 & 0 & \\ 0 & 0 & 1 & 0 & \end{bmatrix} \qquad \alpha_t = \begin{bmatrix} \bar{m}_t \\ \bar{cy}_t \\ \bar{\pi}_t \\ \bar{tp}_t \\ \tilde{r}_t^s \\ \tilde{r}_t^2 \\ \tilde{r}_t^5 \\ \tilde{r}_t^7 \\ \tilde{r}_t^{10} \\ \tilde{r}_t^{BB} \\ \tilde{r}_t^B \\ \tilde{r}_t^{CCC} \\ \tilde{\pi}_t \\ \tilde{\pi}_t^e \end{bmatrix}$$

$$T_t = \begin{bmatrix} \mathbf{I_{4\times 4}} & \mathbf{0_{4\times 10}} \\ & \\ \mathbf{0_{10\times 4}} & \mathbf{\Phi_{t,10\times 10}} \end{bmatrix} \qquad \Sigma_t^{\eta} = \begin{bmatrix} \sigma^2_{\bar{m}_t} & 0 & 0 & 0 & \\ 0 & \sigma^2_{\bar{cy}_t} & 0 & 0 & \mathbf{0_{4\times 10}} \\ 0 & 0 & \sigma^2_{\bar{\pi}_t} & 0 & \\ 0 & 0 & 0 & \sigma^2_{\bar{tp}_t} & \\ & \mathbf{0_{10\times 4}} & & & \mathbf{\Sigma_{t,10\times 4}} \end{bmatrix} \qquad \Sigma_t^{\varepsilon} = \sigma^2_{\varepsilon,t} \bullet \left[\mathbf{I_{10\times 10}}\right]$$

A2.5 The EM algorithm

This section presents the version of the EM algorithm used in this paper, which considers the presence of Gaussian weights ω^i in the updating step of the parameters φ_t (Z_t, H_t, Σ_t^{ε}) and Σ_t^{η}. The algorithm works by alternating two steps until convergence:

1. E-step: given the estimated parameters $\hat{\varphi}_t^{(h)}$, apply the Kalman Filter and the Kalman Smoother, that provide an estimate of the states $\mathbf{E}_{\hat{\varphi}^{(h)}}\left[\alpha_{t|T}\right]$

2. M-step: given the states $\mathbf{E}_{\hat{\varphi}^{(h)}}\left[\alpha_{t|T}\right]$, maximise the log-likelihood function and update the parameters $\hat{\varphi}_t^{(h+1)}$

While the E-step is not affected by the presence of the weights ω_i, the M-step is. The parameters $\hat{\varphi}_t^{(h+1)}$ are given by the following expressions:

$$\begin{aligned}
\hat{T}_t^{(h+1)} &= \left(\sum_{t=1}^{T}\mathbf{E}_{\hat{\varphi}^{(h)}}\left[\omega_t\alpha_t\alpha_{t-1}'\right]\right)\left(\sum_{t=1}^{T}\mathbf{E}_{\hat{\varphi}^{(h)}}\left[\omega_{t-1}\alpha_{t-1}\alpha_{t-1}'\right]\right)^{-1} \\
&= \left(\sum_{t=1}^{T}\omega_t\left[\alpha_{t|T}^{(h)}\alpha_{t-1|T}^{(h)'}+C_{t,t-1|T}^{(h)}\right]\right)\left(\sum_{t=1}^{T}\omega_t\left[\alpha_{t-1|T}^{(h)}\alpha_{t-1|T}^{(h)'}+P_{t-1|T}^{(h)}\right]\right)^{-1} \qquad (2.22) \\
\hat{\Sigma}_t^{\eta(h+1)} &= \frac{1}{T-1}\sum_{t=2}^{T}\mathbf{E}_{\hat{\varphi}^{(h)}}\left[\omega_t(\alpha_t-\hat{T}_t\alpha_{t-1})(\alpha_t-\hat{T}_t\alpha_{t-1})'\right] \\
&= \frac{1}{T-1}\sum_{t=2}^{T}\omega_t\left[\alpha_{t|T}^{(h)}\alpha_{t|T}^{(h)'}+P_{t|T}^{(h)}-\left(\alpha_{t|T}^{(h)}\alpha_{t-1|T}^{(h)'}+C_{t,t-1|T}^{(h)}\right)\hat{T}_t^{(h+1)}\right] \qquad (2.23) \\
\hat{Z}_t^{(h+1)} &= \left(\sum_{t=1}^{T}\mathbf{E}_{\hat{\varphi}^{(h)}}\left[\omega_t y_t\alpha_t'\right]\right)\left(\sum_{t=1}^{T}\mathbf{E}_{\hat{\varphi}^{(h)}}\left[\omega_t\alpha_t\alpha_t'\right]\right)^{-1} \\
&= \left(\sum_{t=1}^{T}\omega_t\left[y_t\alpha_{t|T}^{(h)'}\right]\right)\left(\sum_{t=1}^{T}\omega_t\left[\alpha_{t|T}^{(h)}\alpha_{t|T}^{(h)'}+P_{t|T}^{(h)}\right]\right)^{-1} \qquad (2.24) \\
diag(\hat{\Sigma}_t^{\varepsilon(h+1)}) &= \frac{1}{T}\sum_{t=1}^{T}\mathbf{E}_{\hat{\varphi}^{(h)}}\left[\omega_t(y_t-\hat{Z}_t\alpha_t)(y_t-\hat{Z}_t\alpha_t)'\right] \\
&= \frac{1}{T}\sum_{t=1}^{T}\omega_t\left[y_t y_t'+\hat{Z}_t^{(h+1)'}\left(\alpha_{t|T}^{(h)}\alpha_{t|T}^{(h)'}+P_{t|T}^{(h)}\right)\hat{Z}_t^{(h+1)}\right. \\
&\quad \left.-y_t\hat{Z}_t^{(h+1)'}\alpha_{t|T}^{(h)'}-\hat{Z}_t^{(h+1)}\alpha_{t|T}^{(h)}y_t\right] \qquad (2.25)
\end{aligned}$$

At the end of the last iteration H, the algorithm delivers an estimate of the states $\mathbf{E}_{\hat{\varphi}^{(H)}}\left[\alpha_t\right]$ and of the parameters $\hat{\varphi}_t^{(H)}$.

A2.6 EM algorithm with missing data

The EM algorithm can be adapted to the presence of patterns of missing data in the observed variables vector y_t, as in the case presented in Section 2.5.2. For what concerns the E-step, the Kalman Filter and Smoother can be conveniently adapted (Durbin and Koopman (2001)): when the filter finds a missing observation, it normally predicts the value of the vector $y^*_{t|t}$ but, in the absence of information on the forecast error $z_{jt} = y_{jt} - \mathbf{E}_{[}\alpha_{jt}|y_t]$ for the time series j, it skips the updating step only for this observation, treating the simulated $y^*_{t|t}$ as if it were an observed series in the following iteration of the loop over t. This procedure is standard in time series econometrics, and allows to overcome efficiently the absence of a subset of observations. As far as the second step of the algorithm is concerned, I use the version of the M-step developed by Bańbura and Modugno (2014) to use only the reliable information - that comes from the observed y_t, and not from the simulated series y^*_t - when updating of the parameters φ_t. Their methodology makes use of a selection matrix S_t, a diagonal matrix with size equal to the number of series in y_t, whose entries take a value of 1 if the observation corresponding to that row is present, and 0 if it is missing. Given that the observations y_t enter the M-step only in the relations for $\hat{Z}_t^{(h+1)}$ and $diag(\hat{\Sigma}_t^{\varepsilon(h+1)})$, only Equations 2.24 and 2.25 are modified, while the remaining part of the algorithm is unchanged. These equations become:

$$\hat{Z}_t^{(h+1)} = \text{vec}\left(\sum_{t=1}^{T} \omega_t \left[S_t y_t \alpha_{t|T}^{(h)\prime}\right]\right)\left(\sum_{t=1}^{T} \omega_t \left[\left(\alpha_{t|T}^{(h)}\alpha_{t|T}^{(h)\prime} + P_{t|T}^{(h)} \otimes\right) S_t\right]\right)^{-1} \tag{2.26}$$

$$\begin{aligned} diag(\hat{\Sigma}_t^{\varepsilon(h+1)}) = \frac{1}{T}\sum_{t=1}^{T} \omega_t \Big[& S_t y_t y_t' S_t' + S_t \hat{Z}_t^{(h+1)\prime}\left(\alpha_{t|T}^{(h)}\alpha_{t|T}^{(h)\prime} + P_{t|T}^{(h)}\right)\hat{Z}_t^{(h+1)} S_t \\ & - S_t y_t \hat{Z}_t^{(h+1)\prime}\alpha_{t|T}^{(h)\prime} S_t - S_t \hat{Z}_t^{(h+1)}\alpha_{t|T}^{(h)} y_t S_t + (\mathbf{I} - S_t) diag(\hat{\Sigma}_t^{\varepsilon(h)})(\mathbf{I} - S_t) \end{aligned} \tag{2.27}$$

As made clear by Equations 2.26 and 2.27, this algorithm exploits only the information coming from the actual observations collected in the $S_t y_t$ matrix, while the others receive a zero weight in the updating step.

A2.7 Estimation of the threshold factor model

The threshold model presented in Section 2.5.2 allows to formalize explicitly the presence of a non-linear relationship between the observed short-term safe rate r_t^s and the natural rate $\bar{r}_t$. However, it adds a layer of complexity in its estimation. The model in a general form is the following:

$$y_t = \mathbf{1}(x_{t-j} \leq \theta) Z_{1,t}\alpha_t + \mathbf{1}(x_{t-j} > \theta) Z_{2,t}\alpha_t + \varepsilon_t \tag{2.28}$$

$$\alpha_t = T_t\ \alpha_{t-1} + \eta_t \tag{2.29}$$

The elements to be estimated are three in this model: the unknown states α_t, the SSM parameters $\varphi_t = \{Z_{1.t}, Z_{2,t}, \Sigma_t^\varepsilon, T_t, \Sigma_t^\eta\}$ and the additional parameters introduced by the threshold structure of the model $\psi_t = \{\theta, j\}$. The parameter θ is the level which, when attained by the x_{t-j}, causes the loadings Z_t to switch regime, while j is the number of additional quarters needed for this effect (the "lagged reaction") to actually occur. The procedure I follow for the joint estimation of α_t, φ_t and ψ_t is the one introduced in Correal and Peña (2008), and combines the EM algorithm with a gridpoint search method. The set of admissible values for θ and j is first discretized to create a gridpoint: in my paper, the x_{t-j} variable is the ECB policy rate, and I allow the threshold θ to be equal to 0 (ZLB by definition) or to 0.05 (assuming the ZLB to be reached already when markets expect that the central bank cannot lower further its policy rate when it is at such a modest level). For what concerns the number of lags j, I assume it to be equal to 0 (immediate reaction within the quarter), one or two. The algorithm works as follows:

1. Apply the standard EM Algorithm as described in Section A2.5 for each possible combination of (θ, j), and for each of them collect the latent factors α_t, the SSM parameters φ_t and the log-likelihood function ℓ:

$$\ell(y_t; \alpha_t, \varphi_t) = -\frac{T}{2}\log\ \det(\Sigma_t^\varepsilon) - \frac{1}{2}\sum_{t=1}^{T}(y_t - Z_t\alpha_t)'(\Sigma_t^\varepsilon)^{-1}(y_t - Z_t\alpha_t) \tag{2.30}$$

2. Choose among all the combinations (θ, j) in the six-points grid the one that maximizes ℓ, and take the corresponding values of φ_t and α_t. The optimal number of lags and threshold level

are given by:

$$(\theta, j)^* = \arg\max_{\theta,\ j}\ \ell(y_t; \alpha_t, \varphi_t) \tag{2.31}$$

In the end, the algorithm provides estimates of the states α_t, the SSM parameters φ_t, the threshold parameters ψ_t.

A2.8 Construction of confidence bands

The EM algorithm delivers point estimates of the states $\alpha_t = (\bar{y}_t, \tilde{y}_t)$, as well as of the SSM parameters ϕ_t. Here I describe the procedure followed to construct confidence bands for the trend vector $\bar{y}_t$. The bootstrap-based algorithm is analogous to the one used in Barigozzi and Luciani (2021). For each bootstrap repetition $b = 1, .., B$, it follows these steps:

1. Simulate the states $\alpha_t^{(b)}$ (trend and cycles), starting from the EM point estimates α_t^{init}:

   (a) Trends

      i. Simulate trends at $t = 1$: $\bar{y}_1^{(b)} \sim \mathcal{N}(\bar{y}_1,\ \hat{P}_{1|T})$

      ii. Simulate errors: $e_t^{(b)} \sim \mathcal{N}(\mathbf{0},\ \hat{\Sigma}^e)$

      iii. Simulate trends for $t > 1$: $\bar{y}_t^{(b)} = \bar{y}_{t-1}^{(b)} + e_t^{(b)}$

   (b) Cycles

      i. Simulate cycles at $t = 1$: $\tilde{y}_1^{(b)} \sim \mathcal{N}(\tilde{y}_1,\ \hat{P}_{1|T})$

      ii. Simulate errors: $\epsilon_t^{(b)} \sim \mathcal{N}(\mathbf{0},\ \hat{\Sigma}_t^{\epsilon})$

      iii. Simulate cycles for $t > 1$: $\tilde{y}_t^{(b)} = \hat{\Phi}_t \tilde{y}_{t-1}^{(b)} + \epsilon_t^{(b)}$

2. Simulate bootstrap data for y_t: $y_t^{(b)} = \hat{\Lambda}_t \bar{y}_t^{(b)} + \tilde{y}_t^{(b)}$

3. Apply the EM algorithm on bootstrap data to obtain B new estimates of parameters and states $(\alpha_t^{(b),em})$

4. Normalize states: $\alpha_t^{(b)} = \alpha_t^{init} - \alpha_t^{(b)} + \alpha_t^{(b),em}$

In the end, compute standard deviations and confidence bands for the states using the simulated B time series: $[\alpha_t - z_{\beta/2}\ \sigma_t;\ \alpha_t + z_{\beta/2}\ \sigma_t]$

Chapter 3

High-ESG assets: Robust to the ECB monetary policy shocks?

High-ESG securities widely attracted both praise and criticism in recent years, both for their financial performances and for their capability of conveying investment towards sustainable activities. In this paper, I investigate the reaction of high-ESG stocks to ECB monetary policy shocks, by means of a SFAVAR model identified through high-frequency policy surprises. Subsequently, I construct optimal portfolios for alternative levels of investors' ESG-awareness, and study their robustness to monetary policy shocks. The results highlight a more volatile behaviour of high-ESG stocks at the occurrence of a shock, especially for securities issued by financial companies. As a consequence, portfolios that overweight this kind of securities are less robust to MP shocks than comparable portfolios built using a simple mean-variance approach.

3.1 Introduction

Portfolio choices regarding the optimal asset holding mix are influenced by a large variety of considerations. Typically, the properties of an asset investors are interested in are expected returns and risk, but many other consideration come into play in the selection of portfolios. Investors differ according to the quality of the information at their disposal, their ability to rationally process it and to their preferences; economic agents can derive utility not just from holding assets with certain financial properties, but also from other non-pecuniary features. In this paper, I focus on one of these features, that is the engagement of the issuing firms in sustainable activities related to the protection of the environment and - broadly speaking - to social-oriented projects. Sustainable investment is a major portion of total investment throughout financial markets: assets under management by sustainable investors - *i.e.* investors who take into account environment, social and governance (ESG) factors in their portfolio strategies - reached 30.3 trillion dollars in 2022 (GSIA (2023)), corresponding to 37.9% of total managed assets in the world. ESG investing has attracted public attention in the last decade, as institutional investors - particularly sensitive to signals coming from politics and society - started to increase their ESG holdings. European Union's directives as the CSRD (EU (2022)) and banking regulations as Basel III incorporated obligations to disclose information on sustainability-related topics. The new rules aim at creating a financial environment in which investors can benefit from a wider transparency on the impact of the issuers of the securities they choose to invest in. Ideally, this should also create incentives for businesses to change their production technologies, fostering innovation and accelerating the long needed transition to sustainability. A large literature has studied the specificities of sustainabilty-related securities: they attract a share of investors that hold onto them even in times of crisis, creating a segmented market with specific pricing dynamics (Pedersen et al. (2021)); the firms that issue often outperform the market during crises (Lins et al. (2017)); they can have insurance-like properties (Godfrey et al. (2009)) and are usually deemed safer (Ilhan et al. (2021), Hoepner et al. (2024)). The objective of this work is to investigate whether the specificity of ESG stocks in the European financial markets generates differences in the reaction of their price to monetary policy shocks. Monetary policy decisions are inspired by a principle of market neutrality, but often produce heterogeneous effects on different asset classes. The main question this work tries to answer

is whether high-ESG stocks react differently from others to monetary policy shocks in the Euro Area. A stronger reaction would suggest the presence of an additional layer of volatility, possibly connected with bubble-like behaviours, while a weaker reaction could be associated to a perception of safety of such assets. Safe assets are information-insensitive, liquid securities whose value is quite robust to shocks (Habib et al. (2020)), particularly attractive as a store of value or as collateral and preferred by risk-averse investors: banks' marketing divisions often push on the safety properties of high-ESG assets, advertised as a good diversification solution.

The empirical strategy I employ to investigate the differential impact of monetary policy shocks on ESG stocks is based on a structural factor-augmented VAR model, that yields impulse response functions specific to each stock in the sample. The effect of monetary policy shocks is identified using high-frequency external instruments: the reaction of stocks is then studied separately according to the listing market, the market capitalization and the economic sector of the issuers. Finally, I evaluate the impact of monetary policy shocks on portfolios constructed by the POET estimator (Fan et al. (2013)), assuming different degrees of preference for stocks with a high ESG scores. Results hint at a stronger sensitivity of high-ESG stocks and portfolios to monetary policy shocks, especially for issuers in the financial sector, suggesting that the price of these securities is not particularly robust to unexpected events. The subject has implications both from the point of view of an investor and from a policy perspective, since understanding how such a mass of securities held often by retail asset holders reacts to monetary shocks is crucial in modelling macro-financial interactions and choosing the optimal monetary policy strategy.

The paper is structured as follows: Section 3.2 presents the most relevant literature on the empirical assessment of the high-ESG securities peculiarities. Section 3.3 presents the employed data, as well as the identification strategy and the methodologies employed for model selection. The results on the differential reaction of stock prices are presented in 3.4, while the effects on portfolios constructed with alternative ESG constraints are in Section 3.5. Section 3.6 concludes.

3.2 Literature Review

This work is closely related to various strands of literature, first and foremost to the empirical papers on the relationship between ESG ratings and assets' financial performances. A relevant work in the field is Luo (2022), who employs a micro asset pricing model to show that low ESG quintile portfolios pay significantly higher than average returns, suggesting a risk premium on low ESG stocks. Their results hold when controlling for the Fama and French (1993) factors and are driven by low-liquidity firms, supporting the view of ESG assets as effective tools to reduce downturn market risks during "flight-to-liquidity" episodes. A related paper by Alessi, Hirschbühl, and Rossi (2022) studies the price dynamics of ESG assets using an extended version of the present-value model and finds a disconnect of the price of ESG assets from market fundamentals in recent times, suggesting an overall "overpricing" of these assets caused by a shift of investors' preferences towards sustainability-related features. Godfrey et al. (2009) study the insurance-like properties of assets issued by firms engaged in corporate social responsibility (CSR) activities: using a US database, they find that CSR-active firms experience milder declines in their market value, as reflected by their stock price. These results are rationalized by positing a non-financial link between a fraction of stakeholders and the firm, fostered by CSR activities. The seminal empirical paper on the relationship between CSR and financial performance is Waddock and Graves (1997), who find a positive association between socially-oriented activities and financial performance measures in the US. The work by Nofsinger and Varma (2014) studies the performance of high-ESG funds during the Great Financial Crisis (GFC), finding that ESG funds outperformed similar non-ESG funds during the crisis, at the cost of a systematic under-performance in good times. Lins et al. (2017) reach similar conclusions on US stocks, finding an outperformance during the GFC, and suggest that the effect could be explained by higher level of trust among investors, a scarce good notably during crises. Hoepner et al. (2024) analyze the sensitiveness of ESG assets to downside risks, and finds that ESG engagement significantly reduces the value-at-risk (VaR) and the second-order lower partial moment. A negative association between social performance and risk exposure is equally found by Sassen et al. (2016) on a European dataset, and corroborated by studies on other crisis episodes. Broadstock et al. (2021) finds an outperformance of high-ESG portfolios in China during the COVID-19 crisis, and Pisani and Russo (2021) comes to similar results with an event study

analysis of a set of European funds. A closely related paper is Patozi (2023), who investigates the financial response of green firms to FOMC policy shocks, and finds a weaker reaction to shocks of the price of stocks issued by green firms, with analogous insights provided by credit default swaps.

The paper is related to a second branch of studies, that analyzes empirically the difference between "green" and more carbon-intensive assets. Green assets are conceptually different from high ESG securities, but the two notions are closely related. The paper by Chava (2014) studies firms' cost of capital depending on their environmental profile, and finds that heavy polluters are asked higher dividends from investors and charged higher interest rates from banks, as a result of different exposure to environment-related risk. Matsumura et al. (2014) analyze the effect of carbon emissions intensity on firm value, highlighting a stock value penalization for firms with higher emissions, particularly in the case of missing information disclosure. These results are consistent with the findings of Hsu et al. (2023) on the outperforming returns of portfolios with high average emission intensity. This "pollution premium" is rationalized by presuming different exposures of firms to policy regime shift risks, already priced by markets. In a similar vein, Ilhan et al. (2021) investigate whether climate policy uncertainty is already priced by the option market, and find that protection from downside risk is estimated to be significantly higher in the US for highly polluting firms, especially before 2020. The presence of a positive carbon premium for direct and indirect emissions is supported by the analysis on US stock returns of Bolton and Kacperczyk (2021), who also note that divestment by institutional investors explains only in part the premium. Benchora et al. (2023) study the heterogeneous reaction of stock prices to FOMC policy shocks for different carbon emissions levels of the issuers, by developing a model to explain a lower sensitivity of these green securities and validating the hypotheses by means of a panel event-study on US firms data.

The second part of the work is also linked to the literature on the modelling of green assets' peculiarities. Zerbib (2022) constructs an asset pricing model (S-CAPM) with partial segmentation and heterogenous investors, that differ according to their attitudes towards ESG investment. Asset prices are distorted by exclusionary screening and ESG integration, leading to a taste premium and an exclusion premium explained by the reduction in the investors' base, as in the neglected stocks

model of Merton et al. (1987). Similarly, the paper by Pedersen et al. (2021) derives a solution to the investor's portfolio problem under heterogenous preferences for ESG features. In their model, investors can be ESG-*unaware*, ESG-*aware* or ESG-*motivated*. Using this adjusted version of the CAPM, they derive the highest attainable Sharpe ratio for each ESG level, and provide conditions under which ESG can affect returns that depend on the proportions of agents in the market. Finally, Pástor et al. (2021) develop a model of investing in which green assets present negative alphas, explained by the preference of some investors for green holdings and from the ability of these assets to hedge climate risk. However, the model still considers the possibility of positive alphas in specific cases, as in the case of a investors' tastes.

3.3 Data and identification approach

3.3.1 Data

The dependent variables employed in the analysis consist in financial data on European stock prices between 2014 and 2023. In order to exploit a large set of information on prices reaction to policy shocks, I collect the time series of the STOXX Europe 600 index constituents at daily frequency[1]. Among the stocks included in the index, I select those listed in Euro Area markets and for which a ESG score is available in the last ten years, restricting the sample to a total of 251 time series. ESG scores are collected the Refinitiv ESG Company Scores database. For what concerns the identification of the monetary policy shock, I use two sources: for the yield paid by safe and liquid (AAA) assets on different points of the yield curve (6 months, 1, 2, 3 and 5 years) at daily frequency, I resort to ECB; the reaction of money market to monetary shocks in a narrow time window around policy announcements, required for the high frequency identification, is provided by the EA-MPD dataset (Altavilla et al. (2019)). This comprehensive dataset includes information on the reaction of a large variety of safe rates at different maturities, yields paid by public bonds and other relevant financial metrics, isolating the variations occurred in the time window around monetary policy announcements.

[1]The STOXX Europe 600 index comprises the largest listed firms in Europe, accounting for around 90% of total market capitalization: for this reason, its components should provide a representative view of the dynamics of the European stock markets.

3.3.2 High-frequency identification of a monetary shock in a FAVAR

The behaviour of stock prices' series, processed to ensure stationarity, is modelled by means of a factor-augmented VAR (FAVAR): the large (n-dimensional) vector of observables is described as a linear combination of a small set of r factors F_t, which are in turn modelled as a VAR process. In addition to the factors, the VAR includes a safe rate time series, also used for the identification of the shock. Factors affect the vector of observables both instantaneously and dynamically through their AR structure. The FAVAR model can be written as:

$$Y_t = \Lambda F_t + e_t \tag{3.1}$$

$$\Phi(L)\alpha_t = \eta_t \tag{3.2}$$

Where Y_t is the vector of n observed variables, F_t are the r unobserved factors, Λ is the $n \times r$ loadings matrix, which defines the linear relationship between Y_t and F_t, while the Φ matrix contains the VAR coefficients of the factors in the transition equation. The vector $\alpha_t = (F_t, r_t^{2y})$ is composed by the extracted factors and by the yield of a two-year safe and liquid bond. I then identify a structural monetary shock and its effects on the Y_t vector (structural FAVAR, Belviso and Milani (2006) among others). Similarly to SFMs (Forni, Giannone, et al. (2009), Stock and Watson (2016)) SFAVARs combine the main advantages of factor models - ease in treating large datasets, good fit with macroeconomic data - with the possibility of obtaining stock-specific impulse response functions (IRFs) from orthogonal shocks, as in the SVAR literature (Sims (1980), among others) . The model is composed by the equations of the FAVAR (3.1-3.2) and the structural error equation, that links the VAR shocks η_t to the structural shocks ε_t through the $(r+1) \times 1$ matrix H:

$$\eta_t = H\varepsilon_t \tag{3.3}$$

The stock-specific IRF, defined as the effect of a unit increase in $\varepsilon_{j,t}$ on the i-th variable is:

$$IRF_{ij} = s_i \Lambda \Phi(L)^{-1} H s_j \tag{3.4}$$

Where s_i and s_j are selection matrices. The estimation of Λ and F_t is quite straightforward, as it can be carried out applying principal components, and so is also the ML estimation of the Φ coefficients. I provide below some additional details on the identification of the shock.

Extracting the instrument In the paper, the shock is identified by using an external instrument, following a methodology similar to in Gürkaynak et al. (2005) - on US data - or Altavilla et al. (2019), on EU data among others. The procedure consists of two steps: first, a common factor is extracted from the set of high-frequency reactions of safe rates in the EA-MPD dataset, adequately rotated to be interpreted as a monetary policy shock. In particular, I focus on the reaction of safe rates measured in the time window around the ECB press release: in Equation 3.5, the matrix Z_t^m contains the HF reactions in the EA-MPD, from which the common factor $\tilde{\zeta}_t$ is extracted.

$$Z_t^m = A\ \tilde{\zeta}_t + \tilde{Z}_t \tag{3.5}$$

As in Altavilla et al. (2019), the factor is rotated to maximize its loadings on the short-term part of the yield curve, and normalizing the loading to one at the shortest available maturity: this allows to treat the instrument as a proxy for a target shock, referred to as ζ_t from here onwards[2].

HF-IV identification In the second step of the procedure, ζ_t is used as an instrument for the VAR residual η_t^m in a IV setting, as in Stock and Watson (2018). Equation 3.2 can be rewritten as:

$$\alpha_t = \Phi(L)^{-1}\eta_t = \Phi(L)^{-1}H\varepsilon_t = \Theta(L)\varepsilon_t \tag{3.6}$$

Where the vector of VAR residuals η_t contains the residuals of the two-year rate equation (η_t^m) and those of the factors F_t (η_t^f). The estimate of the parameters in the H matrix is obtained by regressing F_t on η_t^m, using ζ_t as an external instrument:

$$\Theta_{m,f} = \frac{E(\eta_t^f\ \zeta_t)}{E(\eta_t^m\ \zeta_t)} \tag{3.7}$$

[2]I analyze the reaction during the press release, where the interest rate shock is identified by a single factor (defined as "Target" shocks in Altavilla et al. (2019)). In this part of the press announcement, the ECB announces with little details its policy decisions on the reference interest rates and on its open market operations. Forty-five minutes later, the policy event is completed by a press conference: the President reads a short description that explains some of the grounds behind the policy decision, gives other details, and answers questions from the audience.

The application of this identification strategy relies on the following assumptions:

$$E(\varepsilon_t^m \ \zeta) = \phi \neq 0 \qquad \text{(Relevance)}$$

$$E(\varepsilon_t^f \ \zeta) = 0 \qquad \text{(Contemporaneous exogeneity)}$$

The relevance assumption is quite reasonable, given that the disturbance of the two year rate ε_t^m is clearly correlated with the rates movements in the announcement window ζ. The requirement of contemporaneous exogeneity corresponds to presuming that the only effect ζ can have on the factors' innovations goes through ε_t^m. This is a realistic assumption, given that the factors are extracted from stock prices: it can be expected an immediate reaction to the policy shock only through interest rates movements[3]. Once obtained $\Theta(L)$, the impulse response function of factors can be projected on the single stock prices:

$$IRF_{ij} = s_i \Lambda \Theta(L) s_j \tag{3.8}$$

Using this specification, the stock-specific responses are proportional with respect to the others, with their linear dependence expressed by the loadings[4]. The monetary indicator chosen for identification is the daily series of the two-year triple-A interest rate, as in Altavilla et al. (2019)[5].

3.3.3 Selection of the number of factors

In this section, I delve into the selection of the number of factors r in the F_t vector. As the factors as not given an economic interpretation, I remain agnostic on this parameter. I employ the selection method of Alessi, Barigozzi, and Capasso (2010), which generalizes the criterion of Bai

[3] In this SFAVAR-IV setting, a lead-lag exogeneity condition is not strictly required (Stock and Watson (2018)). However, for lags of ζ, this condition is quite straightforward, as shocks are uncorrelated with the past values of the instrument. On the other hand, for leads of ζ the instrument should not be linearly dependent on the VAR error lags: this is sensible assuming a certain degree of uncertainty in the behaviour of the ECB and therefore in the immediate effects of its choices in the money market.

[4] This assumption is relaxed in the Appendix 3.4.3, where I replicate the main analysis by means of a CC-SVAR-like specifation (Forni, Gambetti, et al. (2020)). Both approaches capture the impact of the shock on the "systematic" component of prices/returns, seized by common factors. The effect of idyosincratic disturbances is not considered, focusing on the market impact of the policy shock.

[5] This maturity is a reasonable compromise between the need to identify an interest rate shock (more visible at short maturities) and the possibility to capture the effects of non-traditional monetary policy instruments (that are usually reflected more at medium-long maturities). However, as presented in Section 3.4.3, results are qualitatively equivalent using safe rates on different maturities.

and Ng (2002)[6]. In particular, I use the method based on the minimization of the loss function:

$$\hat{r}_{c,n}^{T} = \underset{0\leq k\leq r_{max}}{\arg\min}\ IC_{2,(c,n)}^{T}(k) = \underset{0\leq k\leq r_{max}}{\arg\min}\ \log[V(k)] + ck\left(\frac{n+T}{nT}\right)\log\left(\frac{n+T}{nT}\right) \tag{3.9}$$

The fit improvement obtained by adding factors, captured by the error variance $V(k)$, is weighted against a penalty function depending on the number of factors k, the size of Y_t (n and T), and the parameter c. The strength of the penalization increases with c, so that the optimal model becomes more parsimonious for higher values of c. In the selection algorithm, the minimization problem is carried out repeatedly on a grid of values of c, and the optimal r is selected at the first interval where the minimum of the $IC_{2,(c,n)}^{T}$ remains constant. Figure 3.1 presents the results of the selection criterion obtained allowing r to vary between zero and ten factors, with the red line showing the optimal number of factors as a function of c. Depending on the size of the window according to which we define r to be stable, the criterion suggests the presence of two or four factors[7].

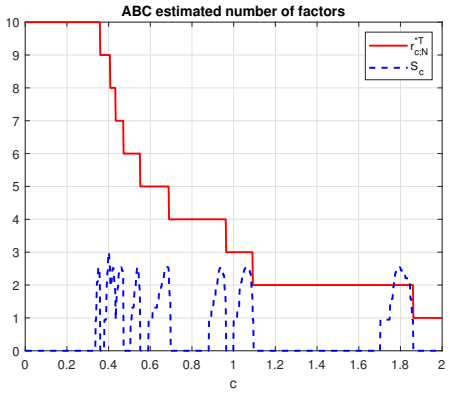

Figure 3.1: ABC criterion (Alessi, Barigozzi, and Capasso (2010))

As a further check, I apply alternative selection criteria based on the eigenvalues extracted from the Y_t matrix: these methods have the valuable properties of not requiring penalty functions or the estimation of threshold for eigenvalues. The standard Scree Test (Cattell (1966)) would indicate the presence of a single factor, as the difference between the first and the second eigenvalue is much larger than all the following differences, as shown by Figure 3.2a. This result is confirmed by the methods in Ahn and Horenstein (2013) (the ER test, based on the ratio between adjacent

[6]Specifically, the paper by Alessi, Barigozzi, and Capasso (2010) elaborate on the logarithmic version of the criteria of Bai and Ng (2002), IC_1 and IC_2.

[7]In order to ensure stationarity, the data employed for the analysis were preprocessed by taking first difference of log prices.

eigenvalues, and the GR test, based on the growth ratio of subsequent eigenvalues), all suggesting the presence of a single factor.

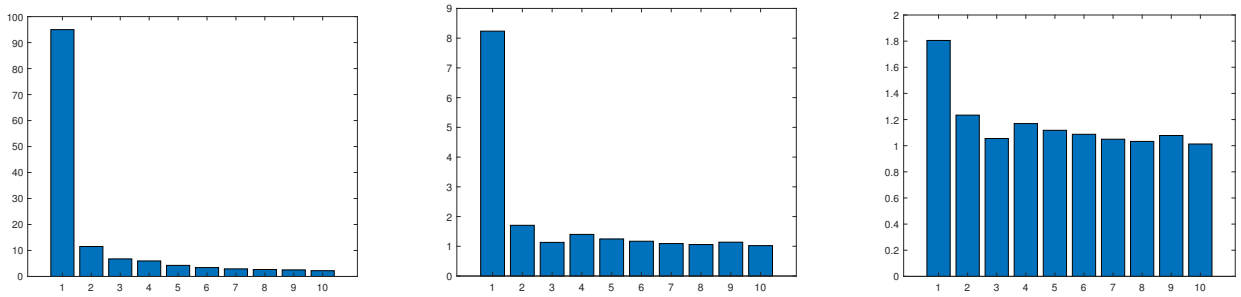

(a) First ten eigenvalues (b) ER (Ahn and Horenstein (2013)) (c) GR (Ahn and Horenstein (2013))

Figure 3.2: Eigenvalues-based criteria for the choice of r

The original criteria of Bai and Ng (2002) signal the presence of four factors, while the methodology of Onatski (2010) (based on the number of eigenvalues exceeding a divergence threshold for the divergence of eigenvalues) indicates five factors. Finally, the criterion of Trapani (2018), in which eigenvalues are randomized and drawn to test their divergence, suggests the presence of a single factor. Table 3.1 summarizes the number of factors obtained using the tested criteria.

Table 3.1: Criteria for the choice of r

Criterion	Factors
Bai and Ng (2002), IC_1	4
Bai and Ng (2002), IC_2	4
Alessi, Barigozzi, and Capasso (2010)	2-4
Onatski (2010)	5
Ahn and Horenstein (2013), ER	1
Ahn and Horenstein (2013), GR	1
Trapani (2018)	1

I start by estimating a model with four factors. Figure A3.1 in the Appendix reports the extracted factors, normalized to have unit variance; the corresponding loadings - listing stocks by their ESG score in ascending order - are in Figure A3.2 in the Appendix. The first factor explains around 34 percent of the overall variance of the returns in Y_t, and shows a particularly volatile behaviour during

the pandemic crisis; loadings all have the same sign, and are on average somewhat higher among high-ESG stocks. The second factor shows a more prolonged, but similar, pattern of volatility after 2020, and explains an additional 5 percent of variance: its loadings are on average close to zero and slightly decreasing at higher ESG scores. The volatility patterns of the third and the fourth factor are similar to the first two, but with loadings very close to zero on average. Given the similar factor patterns and the small loadings, the analysis - and the results presented in Sections 3.4 and 3.5 - is based on two-factor models[8].

3.4 Stock prices reactions

In this Section, I present the stock prices reaction to a monetary policy shock, obtained using a SFAVAR identified using the approach described in 3.3.2. I start by showing the results obtained on the full sample, and then move to present the estimated IRFs splitting the analyzed sample by country, company size and sector of economic activity[9].

3.4.1 Overall effects

The estimated reaction to a restrictive policy shock is negative for all the stocks included in the sample. This finding is broadly consistent with the results in Altavilla et al. (2019) and with an "Odyssean" interpretation of the ECB monetary policy shocks after 2014[10]. At the occurrence of a shock, stock prices fall on impact by 1.3 percent on average, given the chosen normalization; then, they stabilize within a few days. The model does not provide reliable information on the long-term reaction of stock prices to monetary policy shocks: as the research objective of this work is related mostly to the short-run effects, I limit the analysis to a two-week window after the occurrence of a shock. Figure 3.3 displays the on-impact effect of a 15 bp-equivalent restrictive shock for all the 251 stocks in the sample, ordered by the ESG score of their issuers. Examining the magnitude of

[8] A robustness check whose results are presented in the Appendix highlights that the conclusions of the work are robust to alternative choices of the number of factors.

[9] The model is estimated on a stationary version of the data (precisely, on the first difference of log prices); then, in order to have a measure of the impact on prices, I use a cumulated version of the impulse response functions on returns. The size of the shock is normalized to induce an increase of 15 basis points on the two-year safe rate. Confidence bands are constructed by bootstrapping the errors of the FAVAR and replicating the estimation of the IRFs.

[10] Positive "Odyssean" shocks, in contrast to "Delphic" ones, are characterized by reductions of stock prices (Andrade and Ferroni (2021)), as in these occasions the choices and words of the policymakers signal a strong commitment to transparent policy rules that depend on measurable and publicly known macroeconomic and financial variables.

the effect separately by ESG quartile, the average -1.3 percent reaction hides relevant differences. The average price reaction of stocks in the fourth quartile (around -1.6 percent) is stronger than the response of stocks in the the lower groups, and notably stronger than the first quartile stocks response. There are also small differences between the middle two quartiles and the first one, whose components' price seems more robust to the monetary shock (the average effect is -1.1 percent). The bootstrap confidence bands for quartile means displayed in Figure 3.3 (where red lines indicate 90%confidence bands are in red, while 68% bands are in blue) are based on the standard errors of on-impact effects, that are estimated to be smaller than the IRFs standard errors on longer time horizons. The main indication they provide is that stocks with a ESG score slightly above the sample median react very similarly to those slightly below. Moreover, looking at the 90 percent confidence bands, the difference between the first and the second quantile is way less distinct than the gap between the third and the fourth one. Therefore, the specificity seems to be related only to high-ESG stocks, and not necessarily - in the opposite direction - to low-ESG ones, whose reaction is not drastically different from the one of stocks in the quartiles closer to the median.

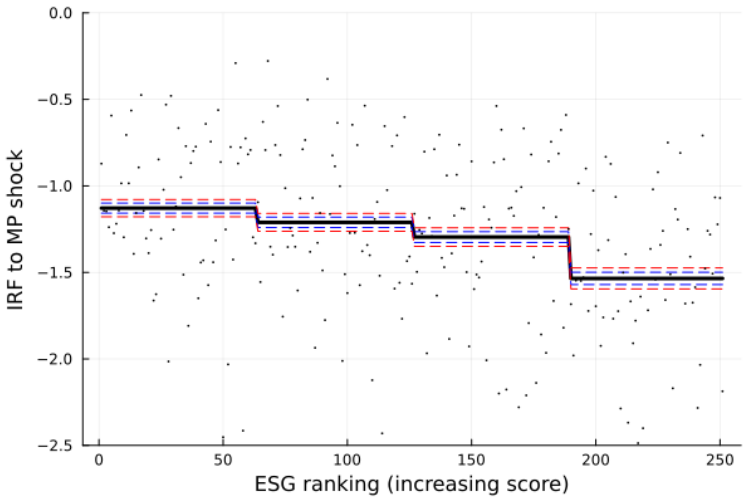

Figure 3.3: On-impact effects, with mean by ESG quartile (whole sample)

Other than by a within-mean comparison, the presence of asymmetries in the reaction of stock prices can be evaluated also by looking directly at the IRFs distribution as in the boxplot of Figure 3.4. Dividing stocks by ESG quartile as done previously, the median reaction in the fourth quartile (-1.5 percent) is always stronger than that of the lower quartiles (around -1.1 percent). There is some within-group variance, but the presence of an effect on the average reaction is confirmed: in this

respect, it is worth noting the 75th quartile IRF in the fourth group is roughly equal to the median of the other three groups. There is no identifiable difference between the two middle quartiles, and the lowest-ESG stocks' reaction is also quite close to the sample median as well.

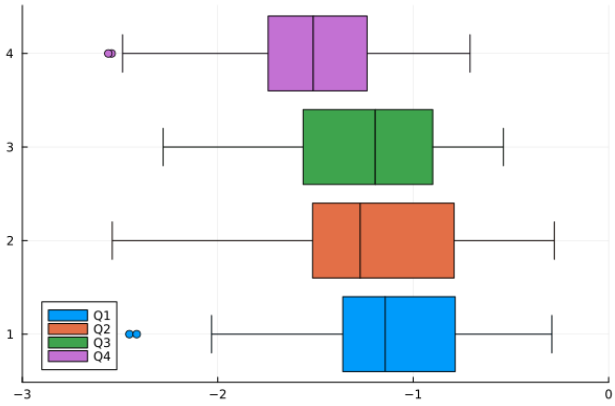

Figure 3.4: Boxplot of on-impact effects (whole sample)

I repeat the analysis looking at the minimum reached by IRF in the two-week window, rather than at on-impact effects. The mean effect by quartile computed on minima is displayed in Figure 3.5. The confidence bands of the means are now based on the bands of IRFs evaluated at their minimum point, that occurs on average after four or five days: uncertainty increases significantly after few days, as other disturbances affect the prices' dynamics. For this reason, also the confidence bands of quartile means are wider than those estimated on on-impact effects, and reaching conclusions on the lagged effect of the shock after some days is quite difficult. However, the average reaction of stocks in the fourth quartile is still stronger (-1.6 percent) than that of stocks in the other quartiles (around -1.2 percent, with no significant differences between the first three quartiles). The level of this minimum points is very similar to that of on-impact effects[11].

[11] Analyzing directly the minimum level of IRFs using a boxplot, as in Figure A3.3 in the Appendix, leads to very similar conclusions. As for on-impact effects, the median reaction of stocks in the fourth ESG quartile is stronger (-1.6 percent) than that of others; also here, the reaction of prices in the second quartile is very similar to that in the third one (around -1.2 percent).

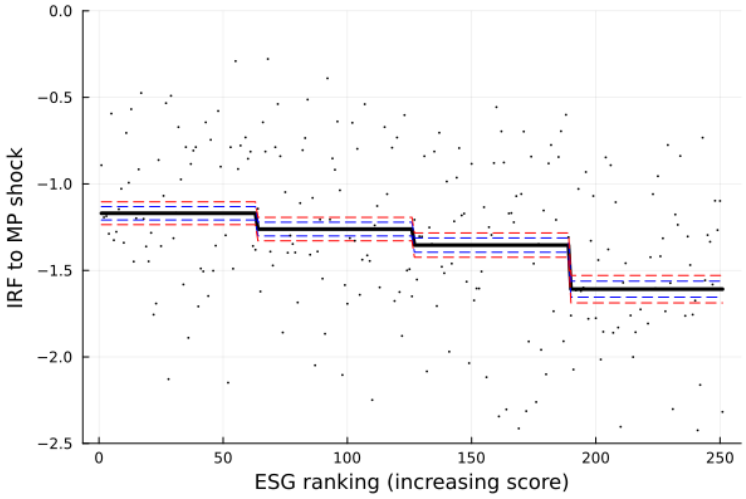

Figure 3.5: Minimum value of IRFs, with mean by ESG quartile (whole sample)

The results presented so far show that the price of high-ESG stocks is more reactive than the average to the ECB monetary policy shock. This finding indicates that these securities are not becoming information-insensitive, a quality commonly attributed to safe assets, and on the contrary that they are particularly volatile when a monetary event occurs. However, it is also consistent with the hypothesis of a slight overpricing of high-ESG securities, as their valuation incorporates elements that go beyond their fundamentals (namely, that exceeds their expected discounted cash flow). In turn, this overpricing of these assets would be caused by a preference shift towards sustainable securities (Alessi, Hirschbühl, and Rossi (2022)): when a shock occurs, therefore, the observed effect is stronger, as the pre-shock price was higher than what implied by the fundamentals. On the other hand, I find smaller evidence of peculiarities for stocks in the lowest ESG quartile: their reaction is marginally smaller than the average, suggesting that they also cannot be considered particularly information-insensitive.

3.4.2 Effects by country, size and sector of economic activity

The sample object of the analysis is representative of a very large and heterogeneous set of stocks. In this Section, I repeat the analysis dividing the stocks' sample according to the stock exchange in which the security is listed, the company size, as measured by its average market capitalization and the economic sector of the issuing firm.

Country Figure 3.6 reports the on-impact effects of the policy shock on stocks ordered by their ESG score, separately by the groups of countries in whose exchange the security is listed. The stock exchanges of Frankfurt and Paris are treated separately, while those in Southern European countries are presented together to have a more informative number of observations. Similarly, the last group is derived residually and includes the smallest stock exchanges in the Euro Area.

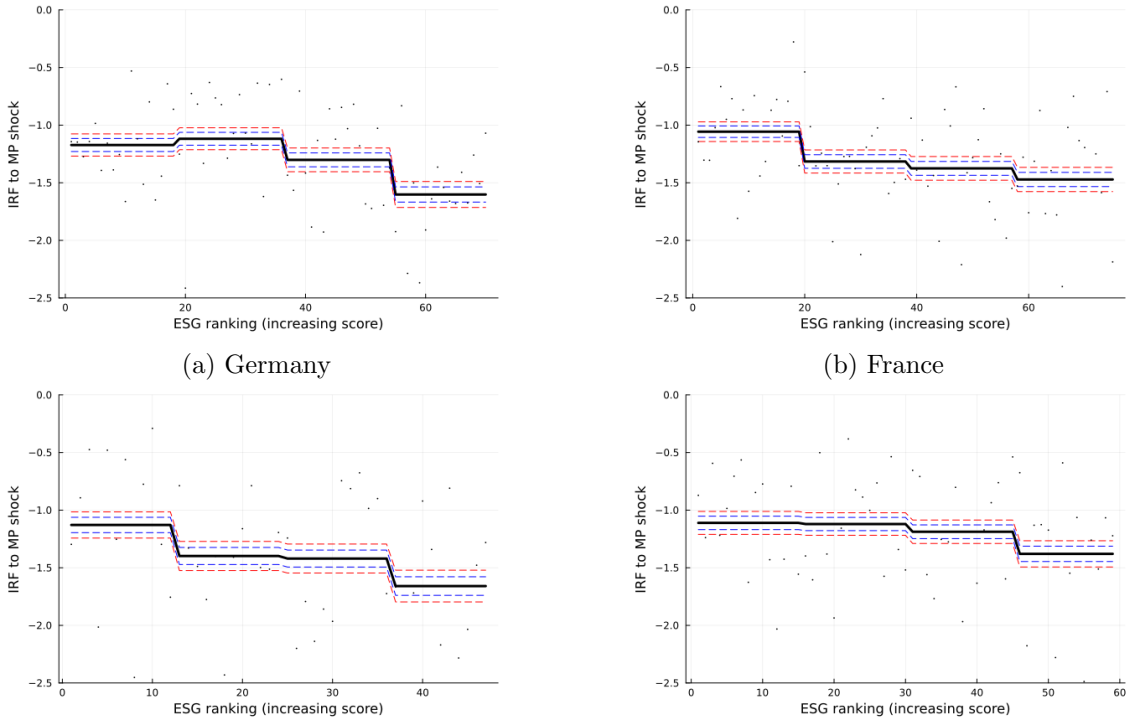

(a) Germany

(b) France

(c) Italy/Spain/Portugal

(d) Netherlands/Belgium/Finland/Austria

Figure 3.6: On-impact effects by country, with mean by ESG quartile

The impact on high-ESG stocks is stronger than the average in all the analyzed exchange subsets. There are some differences in the average reaction of stocks in the middle quartiles (that seem more robust to shocks in Germany and in the Netherlands, less in France), but the difference between the first and the fourth quartile is significant and broadly stable (around 0.4 percent, as summarized by Table 3.2). Therefore, the differential impact along the ESG dimension is not confined to specific exchanges; moreover, it is not caused by the combination of different ESG score levels and asymmetric effects of monetary policy across countries[12].

[12]Figure A3.4 in the Appendix reports the results by country obtained considering the minimum points reached by IRFs, rather than on-impact effects. As the sample size shrinks significantly dividing the sample into four groups, assessing whether the effects after a few days differ significantly is quite troublesome. However, point estimates are in line with those presented for on-impact effects; in addition, considering 68% confidence intervals the difference

Table 3.2: Q4 - Q1 difference

	Q1	Q4	Delta
Germany	-1.2	-1.6	0.4
France	-1.1	-1.5	0.4
Italy/Spain/Portugal	-1.2	-1.6	0.4
Neth./Bel./Aus./Fin.	-1.1	-1.4	0.3
Overall	-1.1	-1.6	0.5

Company size I present here the reactions to the policy shock analyzed on sub-samples of companies with similar size. As for the previous analysis on exchanges, this analysis aims to clarify whether the stronger reaction of high-ESG assets is concentrated in specific market segments. To this end, stocks are first divided into three equally sized groups according to their average market capitalization in the period of interest, and then sorted according to their ESG score. Figure 3.7 reports the on-impact effects for three groups of stocks with similar market capitalization. For different company sizes, patterns are very similar: on average, the reaction of stocks in the fourth quartile is around 0.4 percent stronger than the one of stocks in the first quartile, and the difference between these quartiles is significant. Average levels - for both the extreme quartiles - are also very resembling for different levels of market capitalization[13].

between the first and the fourth quartile is significant in all the four groups.

[13]The results obtained when considering IRFs minima rather that on-impact effects - presented in Figure A3.5 in the Appendix - are qualitatively equivalent. The Q4-Q1 gap is observable and significant across all the three sub-samples, confirming this strongest reaction in all the market segments defined according to capitalization.

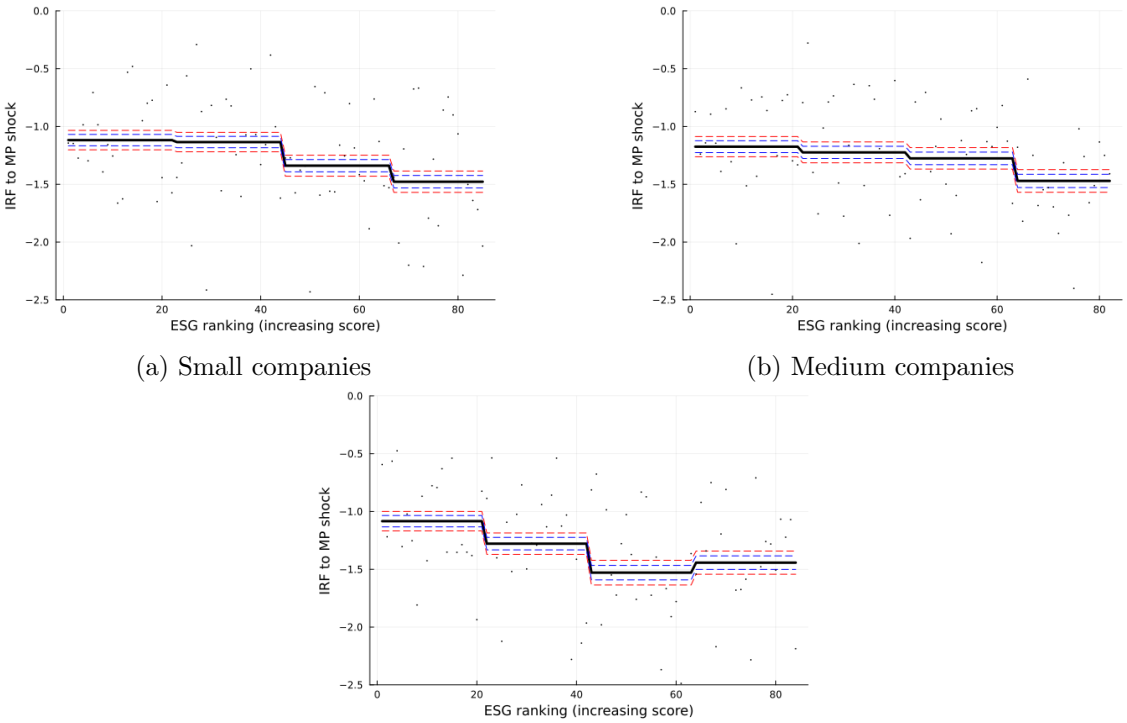

(a) Small companies

(b) Medium companies

(c) Large companies

Figure 3.7: On-impact effects by company size, with mean by ESG quartile

Table 3.3: Q4 - Q1 difference

	Q1	Q4	Delta
Small companies	-1.1	-1.5	0.4
Medium companies	-1.2	-1.5	0.3
Large companies	-1.0	-1.4	0.4
Overall	-1.1	-1.6	0.5

Economic sector The analysis is also repeated by first grouping stocks according to the economic sector of the issuing firm, so to evaluate whether the effect observed on the whole sample is actually driven by differences between sectors, or whereas there are significant differences in the within-sector ESG split. The sample is divided into four groups: Industry, Finance, ICT and Services.

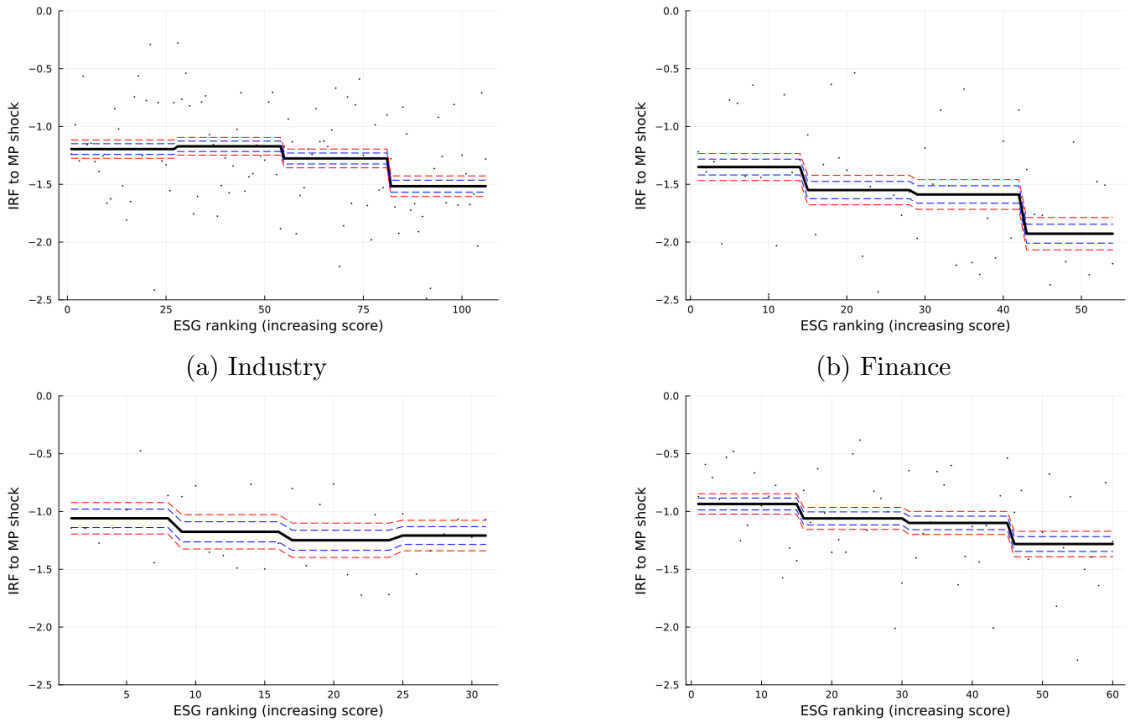

(a) Industry

(b) Finance

(c) ICT

(d) Services

Figure 3.8: On-impact effects by sector, with mean by ESG quartile

Figure 3.8 reports the on-impact effects for stocks divided by sector and ESG quartile. Results exhibit relevant differences among sectors: while for ICT firms the reactions of stocks in the first and fourth quantiles do not differ significantly, there is a much clearer pattern for stocks issued by financial companies. In between, among firms in the industry and services sectors, the Q4-Q1 gap is significant, but amounts only to a half the interquartile range among financial companies. As summarized by Table 3.4, the subset in which the clearest and strongest effect emerges is in fact the financial sector (with a 0.6 percent average gap and a decreasing trend across quartiles). These conclusions are confirmed by the analysis of IRFs minima (reported in Figure A3.6 in the Appendix): the difference is much more evident for financial companies than in the other sectors.

Table 3.4: Q4 - Q1 difference

	Q1	Q4	Delta
Industry	-1.2	-1.5	0.3
Finance	-1.3	-1.9	0.6
ICT	-1.1	-1.2	0.1
Services	-0.9	-1.2	0.3
Overall	-1.1	-1.6	0.5

These sectorial differences are of particular interest: the sector for which I estimate the strongest price reaction for high-ESG securities is one in which the link between ESG scores and environmental - or other clearly measurable - impacts is particularly blurred. Put differently, the most volatile stocks at the occurrence of a shock (probably as a result of a misalignment with respect to the fundamental value) are not the high-ESG stocks issued by industries actively engaged in sustainable activities in their core business, but those issued by banks and other financial companies whose high ESG score is explained by higher shares of loans to environmental-friendly economic activities (with all the ambiguities this evaluation brings with it), *ad hoc* consultancies to increase the S and G components of the score, and other elements whose link to actual environmental impacts is quite disputable. A work by Giannetti et al. (2023) shows that European banks with greater environmental disclosures provide more credit to high-emissions firms than to other comparable financial institutions, suggesting a certain degree of climate goals "overemphasizing" and a tendency to conceal the actual business model of such banks. Another possibly meaningful factor is that financial companies - and especially banks - are the only sector that issues shares (that can be labelled as "green" with ESG certifications often provided by associated companies) and at the same time offers investment and consultancy services for his customers. This creates an incentive to use this kind of certifications as a way to promote investment in these assets (on the basis of the positive impact this would have on the environment, or on other sensitive fields), increasing their demand and pushing their prices over their fundamental value. When a shock occurs, these

prices fall more as they tend back to their "fundamental" value. On the contrary, the reaction of stocks issued by companies in the other sectors varies little to nothing as a function of ESG scores. Firms in the other sectors do not have the possibility to convey investors towards their preferred securities as easily as financial companies; investment in these sectors - and the robustness of prices to monetary policy shocks - can be driven by other, more substantial factors.

3.4.3 Response to restrictive and expansionary shocks

The results of the previous sections presented the stock prices reaction to a monetary policy shock identified on a sample of mixed - both restrictive and expansionary - policy announcements. As shocks are identified up to their size and scale, the displayed results are based on a normalization: in this case, they measure the reactions to a 15-bp restrictive shock. According to the baseline results, therefore, the price of high-ESG stocks tends to fall more than the average at the occurrence of a restriction, and to grow more in the presence of an expansion. Here, I try to identify restrictions and expansions separately to investigate possible non-linearities.

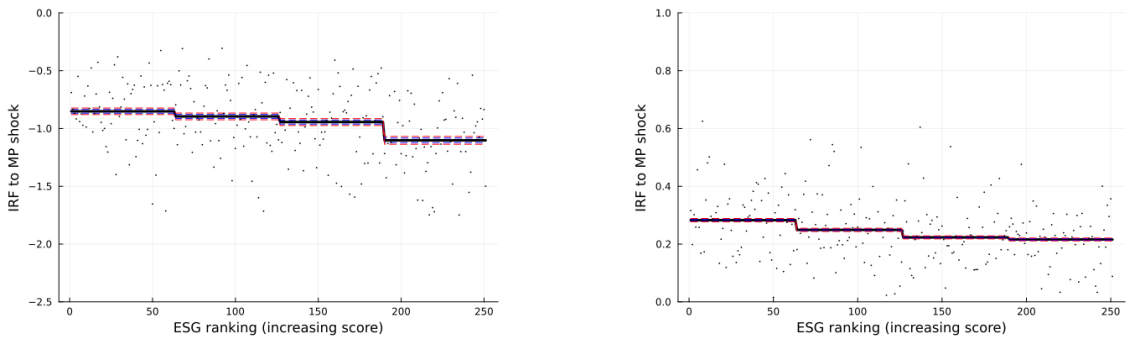

(a) Restrictive shocks (b) Expansionary shocks

Figure 3.9: On-impact effects of shocks of different sign

Figure 3.9 presents the results obtained by differentiating the two kinds of shocks[14]. The left panel shows how the reaction of restrictive shocks - on average - does not differ too significantly from what was observed on the full sample. The mean reaction to a 15-bp shock is a 1.0 percent fall, with the overreaction of high-ESG stocks remaining clear and significant. Analysing separately monetary

[14]This analysis is conducting by "muting" the instrument in the case of expansionary shocks, to capture the effect of policy restrictions, and by symmetrically muting it for restrictions of gauge the effect of a loosening.

expansions, I find a weaker average reaction of the stock prices in the sample - a 0.2 percent increase -, suggesting an asymmetric sensitivity of financial markets to monetary policy decisions. Moreover, the reactions of stocks in the different ESG quartiles are much more alike. Both low- and high-ESG assets tend to react very little to the monetary expansions of the analyzed period, with a very light over-reaction of stocks in the first two quartiles. Therefore, the value of high-ESG stocks seems to react more than average to restrictions (by falling more) and somewhat less to expansions (by increasing less than stocks with a lower score).

Robustness checks

The results presented in the previous Sections are generally robust also to other robustness checks, carried out by modifying some characteristics of the model with respect to its baseline version. Figure A3.8 in the Appendix presents the on-impact effects by ESG quartile obtained by means of a SFAVAR with a different number of factors. Estimating a single-factor model yields very similar impacts to the ones obtained with the baseline model: differences among ESG quartiles remain significant, and their size (0.5 percent) is in line with the estimates of the two-factor model. Models with a higher r capture a larger fraction of the price impact, but interquartile gaps are substantially unaffected. Despite the additional uncertainty these less parsimonious models bring with them, differences between quartiles remain significant.

In the baseline model, the time series used for the identification of the monetary shock is the average yield paid by a European triple-A bond on a two-year maturity. As shown in Figure A3.9, the main conclusions of the analysis are not dependent on this choice. Using rates on shorter maturities for identification produces similar price reactions to the one obtained with the two-year model on all ESG levels. Levels differ slightly using longer maturities: however, differences among ESG quartiles remain significant with all the tested maturities, a reassuring finding on the conclusions of the main analysis.

An additional robustness check involves considering the effect of Fama and French's (2015) five factors on stocks' returns over time, and the consequences this could have on the reaction to monetary policy shocks. In order to verify whether the observed effect on ESG scores was hiding the

influence of these factors, I estimate the SFAVAR model on a modified version of the data, obtained by taking the residuals of regressions of the original time series on the daily five factors (Mkt-Rf, SMB, HML, RMW and CMA). Figure A3.10 displays how, despite the average price reaction of these residuals is weaker than the one obtained in the baseline analysis - as a part of the impulse response is captured by the Mkr-Rf factor - it is still possible to observe the over-reaction of high-ESG stocks. Moreover, the difference between tail quartiles - measured as a proportion of the average effect - is quite in line with the results obtained on the original data. This suggests a specificity of the reaction of high-ESG stocks which is not captured by Fama and French's (2015) factors.

3.5 Reactions of ESG-aware and ESG-unaware portfolios

Up to this point, the analysis has focused on the reaction of independently picked stock prices to monetary policy shocks, then grouped by ESG quantile to obtain more general results. In this Section, I study the subject from another perspective: first, I construct optimal portfolios assuming different degrees of investors' desired ESG integration; second, I analyze the reaction of the value of the obtained portfolios to monetary policy shocks identified with a HF approach. This allows to mimic a more realistic setting, in which investors opt to diversify risks.

3.5.1 POET-based optimal portfolios with different degrees of ESG integration

The construction of optimal portfolios is carried out assuming, as in the previous analysis, that the market component of the behaviour of the 251 stock prices in the sample can be captured by a small number of factors. For this reason, I estimate the large-dimensional covariance matrix - needed for the portfolio optimization - using the POET estimator (Fan et al. (2013)). The technique works as follows: first, given the chosen number of factors r, PCA is applied. This allows to write the spectral decomposition of the POET covariance matrix $\hat{\mathbf{Q}}^{\mathbf{r}}$ as:

$$\hat{\mathbf{Q}}^{\mathbf{r}} = \sum_{i=1}^{r} \hat{\lambda}_{\mathbf{i}} \hat{\xi}_{\mathbf{i}} \hat{\xi}'_{\mathbf{i}} + \hat{\mathbf{S}}^{\tau}_{\mathbf{r}} \tag{3.10}$$

Where $\hat{\lambda}_i$ and $\hat{\xi}_{\mathbf{i}}$ are respectively the first r eigenvalues and eigenvalues of $\hat{\mathbf{Q}}^{\mathbf{r}}$. The POET estimator penalizes the off-diagonal elements of the principal orthogonal complement $\hat{\mathbf{S}}^{\tau}_{\mathbf{r}}$, but allows for some

cross-sectional correlation also after taking out the first r principal components. Thresholding is not applied on the diagonal entries, and works through a shrinkage function σ_{ij} applied on entries whose absolute value exceeds τ.

$$\hat{\mathbf{S}}_{\mathbf{r}}^{\tau} = (\hat{s}_{pxp}^{\tau}) \qquad \hat{s}_{ij}^{\tau} = \begin{cases} \hat{s}_{ij}, & i = j, \\ \sigma_{ij}(\hat{s}_i j)\mathbf{I}(|\hat{s}_i j| \geq \tau), & i \neq j \end{cases} \tag{3.11}$$

The final result is the POET covariance matrix $\hat{\mathbf{Q}}^{\mathbf{r}}$, given (Equation 3.10) by the sum of the covariance explained by the first r principal components and by its penalized complement. Figure A3.11 in the Appendix reports the heatmap of the sparse POET-based $\hat{\mathbf{Q}}^{\mathbf{r}}$ estimate. As expected, the highest values are on the diagonal entries (containing returns' standard deviations), but some cross-correlation is observable in the off-diagonal elements (captured both by the first PC and by the penalized complement $\hat{\mathbf{S}}_{\mathbf{r}}^{\tau}$)[15]. Given the estimated covariance matrix $\hat{\mathbf{Q}}^{\mathbf{r}}$, the optimization problem solved by an ESG-unaware investor can be modelled as the choice of the vector of weights $\mathbf{x}$ that, given the budget constraint defined by Equation 3.13, minimizes the portfolio variance needed to obtain a certain expected return s. I allow for a share of the endowment - normalized to one - to be stored as a liquid asset paying no interest (or equivalently as money), while I rule out short selling by constraining weights x_i to be non-negative[16].

$$\min_{x} \quad \mathbf{x}'\hat{\mathbf{Q}}^{\mathbf{r}}\mathbf{x} \tag{3.12}$$

$$\text{s.t.} \quad \sum_{i=1}^{N} x_i \leq 1 \tag{3.13}$$

$$\sum_{i=1}^{N} \bar{r}_i x_i \geq s \tag{3.14}$$

$$x_i \geq 0, \quad \forall i \in [0, N] \tag{3.15}$$

[15]The displayed covariance matrix is obtained assuming the presence of two factors ($r = 2$), as in the previous analysis, and applying a soft thresholding with $\tau = 0.5$. Results are robust to a series of checks - not presented here for the sake of brevity - carried out changing the number of factors r, the threshold τ or applying hard thresholding rules ($\sigma_{ij}(\hat{s}_i j) = \hat{s}_i j$).

[16]Despite the advantages of analyzing optimal portfolios over individual stock prices, this approach is still subject to relevant limitations. Portfolios are here constructed assuming a simple, static mean-variance approach, while long-term investor decisions might be more realistically modelled as a dynamic optimization problem (Campbell and Viceira (2002)). For simplicity, in this study I exclude the possibility to re-optimize the portfolio allocation.

In this framework, the optimization problem can be conveniently adapted to allow for investors' ESG integration in the selection of the portfolio. In particular, I add an additional constraint to the problem defined by Equations 3.12 - 3.15, that requires the overall ESG score of the portfolio - defined as a weighted average of its components - not to be smaller than a threshold ξ. The parameter ξ measures therefore the strength of ESG integration in the investor's preferences.

$$\sum_{i=1}^{N} ESG_i \; x_i \geq \xi \tag{3.16}$$

Figure 3.10 displays the optimal portfolio weights obtained assuming an ESG-unaware mean-variance approach and with the introduction of ESG integration at growing levels. Most of the low-ESG stocks in the left part of the graph, included in the portfolio chosen with the simple MVO (blue lines) are replaced by others - otherwise not selected - when the investor requires a mean ESG portfolio score of 65 (red lines). With more intense ESG preferences (mean score of 75), this effect becomes even stronger (green lines)[17].

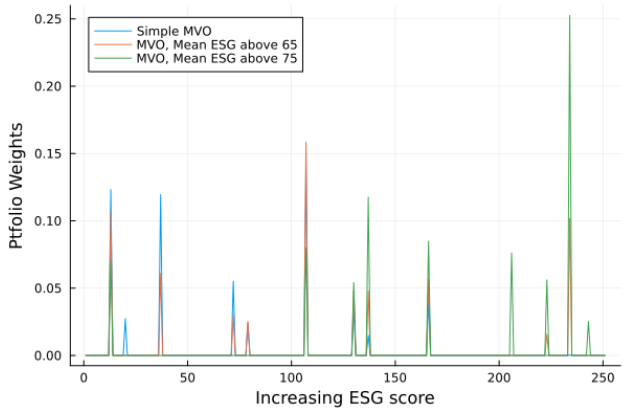

Figure 3.10: Optimal weights as a function of ESG integration, 1 percent return

Figure 3.11 reports the estimated efficient portfolio frontiers, together with a sample of randomly selected portfolios: for each level of desired expected returns, the frontier indicated the minimum possible level of portfolio volatility. The black line is the ESG-unaware frontier, obtained with a simple mean-variance approach, while the green lines are frontiers for investors with increasingly strong ESG integration considerations in their portfolio choices (respectively, mean portfolio scores

[17]The results presented here are obtained assuming the investor to require an average 1 percent monthly return; however, choosing different seeked returns does not impact significantly the results. Optimal weights in the case of 0.5 percent returns are reported in Figure A3.12, in the Appendix.

over 65, 75 and 85). The distance between the efficiency frontiers is considerable: as an example, the optimal portfolio for an ESG-unaware investor who requires an expected return of 1 percent exhibits a 2.6 percent standard deviation, while an investor who - in addition - exacts an portfolio ESG score of 75 needs to tolerate a 3.1 percent volatility. This difference becomes even more striking for investors with stronger ESG integration preferences, reaching 4.8 percent at a 85 level. This result can be explained by the reduced degree of diversification that an investor suffers from in order to satisfy her ESG preferences, and formally by the introduction of additional constraints in the optimization problem. On the other hand, high-ESG stocks do not seem to possess specific features as strong as to offset the impact on diversification possibilities of the additional constraint.

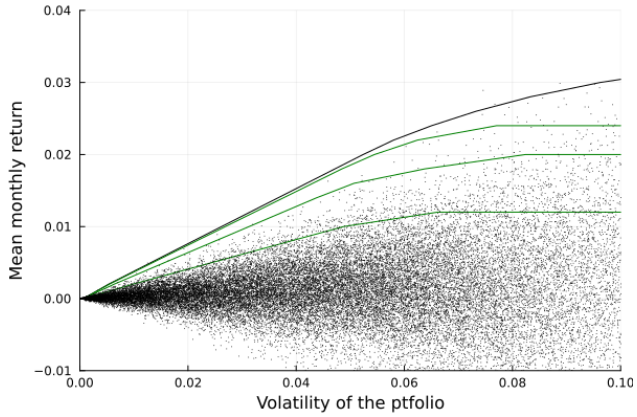

Figure 3.11: Optimal frontiers at alternative degrees of ESG integration

An alternative way to assess the interference of ESG integration in the optimal diversification problem is by estimating the optimal Sharpe Ratio (Sharpe (1966)). The ratio between the expected returns and the volatility of optimally selected portfolios is a standard measure of portfolio performance: here, I evaluate its relationship with a varying degree of desired ESG integration. Figure 3.12 reports the maximum attainable Sharpe Ratio as a function of the required portfolio ESG score (normalized to the market average score). It is possible to notice a substantial stability around 0.4 for investors who require an ESG score that is not greater than the market average. This implies that a mild integration is possible at a negligible cost in terms of efficiency. However, investors asking for a portfolio ESG score 10 or 20 percent higher than the market average already need to accept a much lower ratio, as diversification is severely impacted; at higher levels, the ratio deterioration becomes even faster, as indicated by the concavity of the curve.

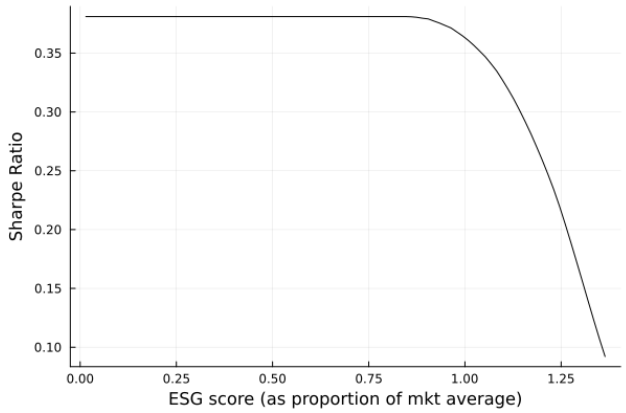

Figure 3.12: Maximum attainable Sharpe Ratio and ESG integration

The results presented so far expose how ESG integration hinders the extent to which diversification is feasible in the selection of a portfolios, and the effects of the additional constraints on the portfolio volatility. Here, instead, I return to the main subject of this work, focusing on the portfolios' reaction to monetary policy shocks. The approach I choose consists in the estimation of a SVAR model including four efficient portfolios with the same expected returns (1 percent) and an increasing level of ESG integration - constructed following the procedure described above - and the two-year safe interest rate[18]. Figure 3.13 reports the IRFs of the four portfolios to a restrictive 15 bp-equivalent monetary shock. Consistently with the findings of the analysis at stock level, the on-impact reaction of the portfolio constructed with the simple mean-variance approach is the smallest in the model (around -0.2 percent). The strength of the impact increases if we consider portfolios constructed with a more intense ESG integration, reaching -1.0 percent in the most extreme case. Considering the bootstrap-based confidence bands on impact, the difference between the ESG-unaware and the ESG-aware portfolios becomes significant - at a 68 percent level - only for strong preferences (over 75 on the portfolio average), while this is not the case for milder integration levels, in line with the evidence on the Sharpe Ratio.

[18]In order to preserve consistency with the results presented in Section 3.4, I employ the two-year rate for identification. However, results are qualitatively robust using alternative maturities. The identification of the shock being carried out by means of the high frequency methodology mentioned in Section 3.3.2.

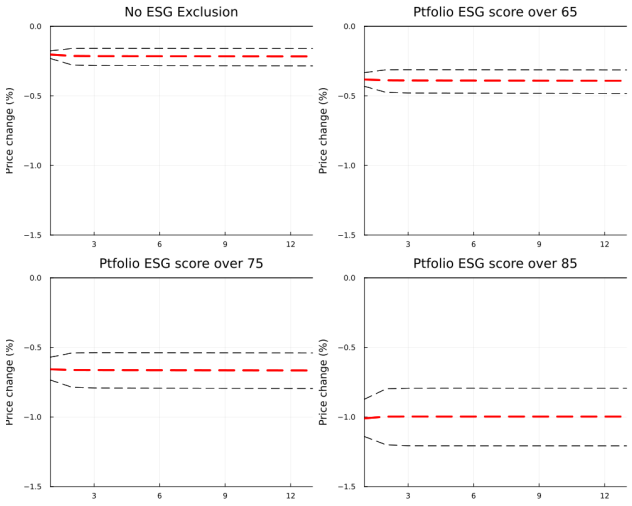

Figure 3.13: IRFs to MP shocks, optimal portfolios with different degrees of ESG integration

The higher sensitivity of ESG-aware portfolios to monetary policy shocks can be therefore explained as the sum of two effects: first, introducing constraints in the portfolio optimization problem reduces the diversification scope, coercing investors to accept some additional volatility for the same level of expected returns; secondly, the high-ESG assets, over-weighted in the aware portfolios, are more reactive to monetary shocks, as shown in Section 3.4, and therefore drive the portfolios response to be stronger. As a results of these two factors, the portfolios constructed with high degrees of ESG integration are particularly receptive to policy shocks, decisively distinguishing these securities from information-insensitive assets.

3.5.2 More radical preferences: complete exclusion of low-ESG stocks

In this Section, I examine the effects of the introduction of ESG considerations in portfolio selection with an alternative modelling approach: instead of assuming investors' to be interested in the average ESG score of their portfolios ("integration"), I restrict the sample of eligible stocks by setting a threshold on the minimum ESG score of the stocks' issuing firms ("exclusion"). Such selection criterion involves assuming more radical ESG preferences, with the objective of better approximating the behaviour of ESG-motivated retail investors. The optimization problem in this

case is:

$$\min_{x} \quad \mathbf{x}'\hat{\mathbf{Q}}^{\mathbf{r}}\mathbf{x} \tag{3.17}$$

$$\text{s.t.} \quad \sum_{i=1}^{N} x_i \leq 1 \tag{3.18}$$

$$\sum_{i=1}^{N} \bar{r}_i x_i \geq s \tag{3.19}$$

$$x_i \geq 0, \quad \forall i \in [0, N] \tag{3.20}$$

$$x_i = 0 \ \text{ if } \ ESG_i < \xi \tag{3.21}$$

The objective function 3.17, as well as the constraints 3.18-3.20, follow the formulation of the "Integrator's" problem, while 3.21 is modified: a stock cannot be included in the portfolio mix if its ESG score is lower than a threshold imposed by the investor based on the intensity of her preferences. Therefore, the constraint is set upstream - on the sample of acceptable stocks - rather than downstream on the portfolio average score, as it is for ESG integration.

Figure 3.14 reports the optimal portfolio frontiers constructed by solving the problem with alternative levels for the desided ESG minimum score (chosen consistently with the previous analysis on ESG integration). The effects of the introduction of ESG-based constraints on diversification are reflected by the higher volatility investors need to accept, keeping expected returns constant. As expected, this effect is even stronger for ESG-"excluders": given a 1 percent expected return, ESG-unaware portfolios display a 2.6 standard deviation, while this figure rises to 3.1 percent with a threshold at a 65 ESG level, and 4.8 with a 75 level.

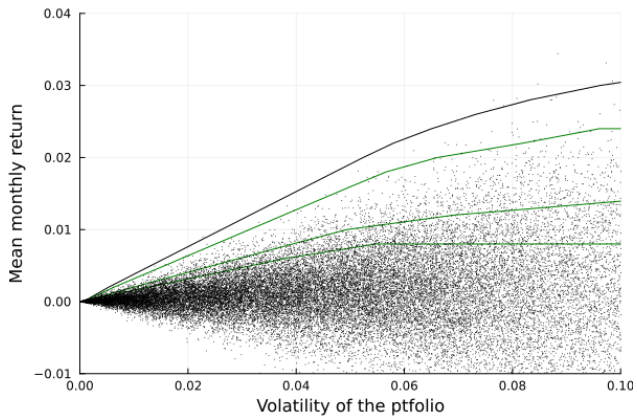

Figure 3.14: Optimal frontiers at alternative degrees of ESG exclusion

Figure 3.15 reports the maximum attainable Sharpe Ratio as a function of the intensity ESG exclusion, measured as a proportion of the average market score. Results are substantially in line with those provided by the analysis on ESG integration. The optimal Sharpe Ratio attainable by an investor who disregards ESG considerations is around 0.4, while some mild exclusion is feasible at the cost of quite a limited efficiency amount in the portfolio selection. On the contrary, stronger degrees of exclusion (10-20 percent above the market average) are feasible only accepting a much lower ratio. The main difference lies in the kinks of the curve: the decrease in the SR is not smooth as in the integration case, but it happens suddenly when a stock - present in the optimal portfolio mix for a given threshold - becomes inadmissible as ESG exclusion becomes more stringent.

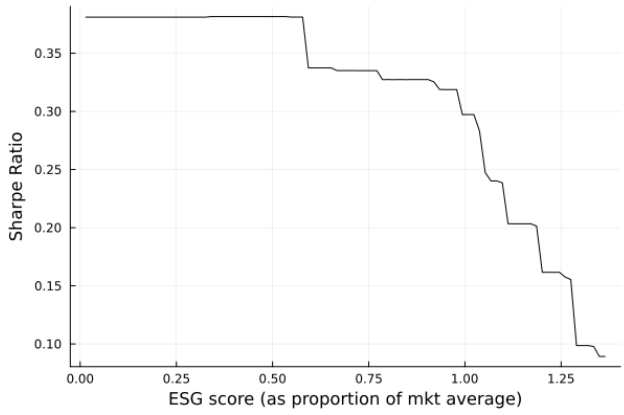

Figure 3.15: Maximum attainable Sharpe Ratio and ESG exclusion

Figure 3.16 reports the impulse response functions obtained by estimating a SVAR model comprising

portfolios with the same expected returns[19], constructed with alternative degrees of ESG exclusion, and the two-year safe rate required for the usual HF identification strategy. Consistently with the results provide by the analysis on ESG-integration, the on-impact reaction of ESG-aware portfolios is stronger than the reaction of the standard mean-variance portfolio. As expected, differences are larger than those estimated in the case of ESG integration, and are significant on impact - at a 68 percent level - already for a mild degree of exclusion preferences (65 score level). On the other hand, the wide confidence bands suggest not to take up outcomes on IRFs few days after the shock. Finally, also in this case this difference can be explained by the combined effect of the reduced diversification possibilities in the presence of ESG exclusion and by the higher sensitivity of high-ESG stocks to monetary policy shocks.

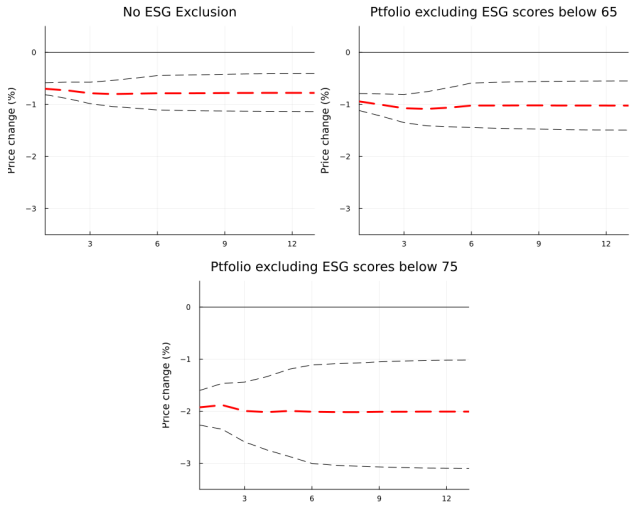

Figure 3.16: IRFs to MP shocks, optimal portfolios with different degrees of ESG exclusion

3.6 Conclusions

The purpose of this paper was to investigate how high-ESG stocks exchanged in the Euro Area react to monetary policy shocks occurring as a result of monetary policy announcements in the EA, both analyzing the behaviour of individually considered stock prices and studying the reaction of portfolios constructed assuming different degrees of investors' interest in ESG concerns. The

[19] As in the analysis on ESG integration, the chosen expected return is 1 percent. As it was not possible to construct a portfolio with such high returns and an ESG exclusion threshold at the 85 score level, the model is estimated only on the ESG-unaware portfolios and on portfolios constructed with thresholds at the 65 and 75 score level.

main result provided by the SFAVAR analysis is the over-reaction of the prices of stocks issued by firms with ESG scores in the fourth quartile: in the event of a restrictive 15-bp policy shock, the price of high-ESG stocks falls on average by around 0.5 percent more than that of stocks in the first two quartiles. Moreover, a further analysis proved how most of the effect is driven by stocks issued by financial companies, a group for which environmental impacts - and, more in general, - CSR activities are particularly hard to evaluate. These findings are consistent with those obtained with the portfolios investigation: with the same expected returns, portfolios constructed assuming higher degrees of ESG integration or exclusion are more volatile both overall - as shown by the analysis of efficient frontiers and by the Sharpe ratios - and in the event of monetary policy shocks. These results rest on a series of assumptions, which also represent the main limitations of this study. The employed identification approach in the SFAVAR assumes symmetric effects between monetary expansions and restrictions, and the produced IRFs are proportional with each other (some extensions of the model were introduced to account for these shortcomings). Furthermore, the portfolios of Section 3.5 are constructed assuming investors to hold only stocks or a safe and liquid asset paying no interest. Finally, portfolio selection is based on simple mean-variance approach corrected to account for ESG considerations: other modelling approaches were not considered for the sake of simplicity. In the context of European financial markets, the peculiarities of high-ESG securities are a topic of fundamental importance, both as results of shifts in investors' preferences and out of relevant legislative innovations which will compel a growing number of firms to disclose information on their CSR-related activities. This work has analyzed one of the features of these assets - the reactivity to monetary policy shocks -, with the objective of shedding light on a relevant aspect, that can in turn help providing a more complete view of the financial performances of this important asset class.

A3 Appendix

A3.1 Additional results from the SFAVAR model

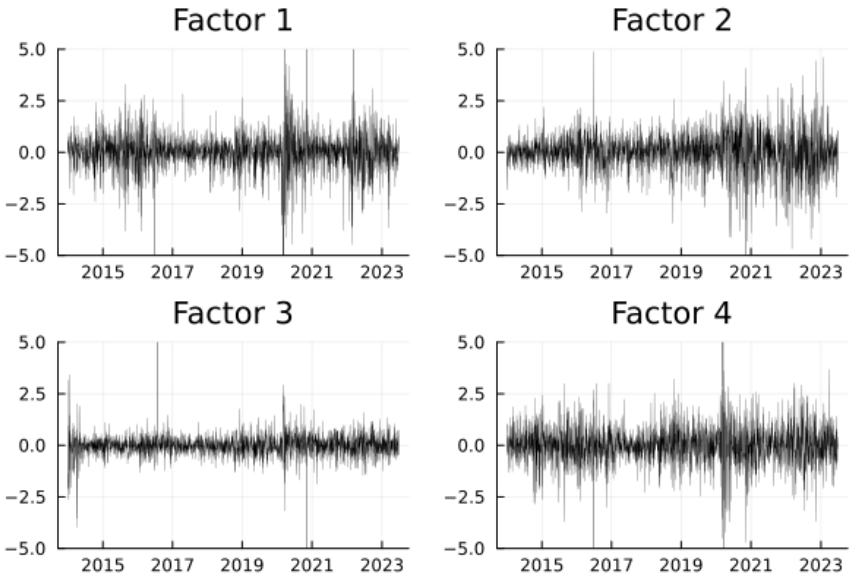

Figure A3.1: First four extracted factors, normalized

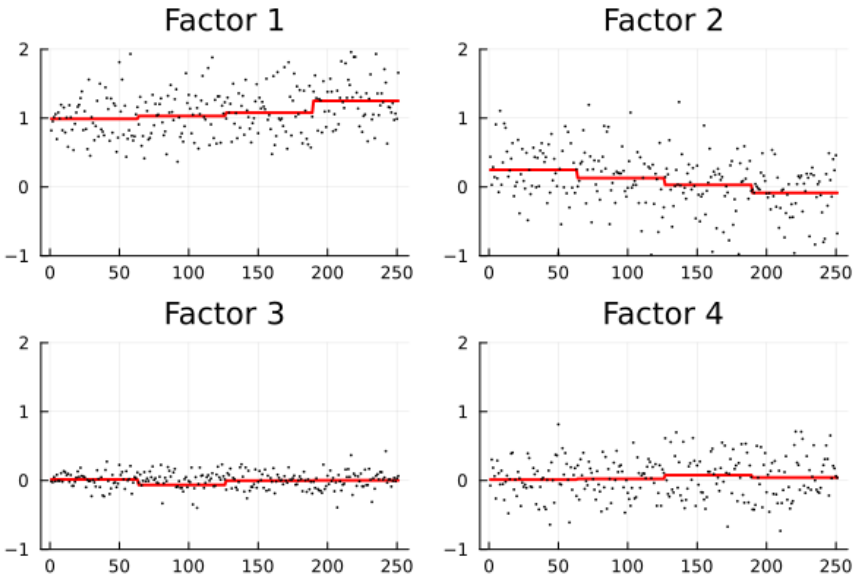

Figure A3.2: Loadings on the 251 stocks, with means by quartile of ESG ranking

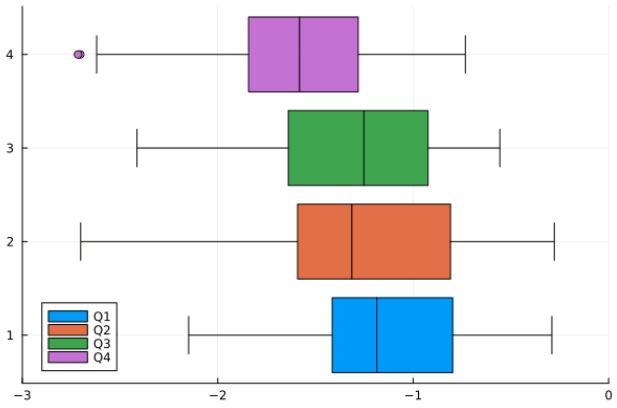

Figure A3.3: Boxplot of IRFs minima (whole sample)

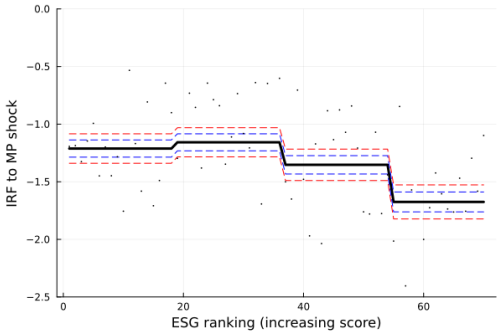

(a) Germany

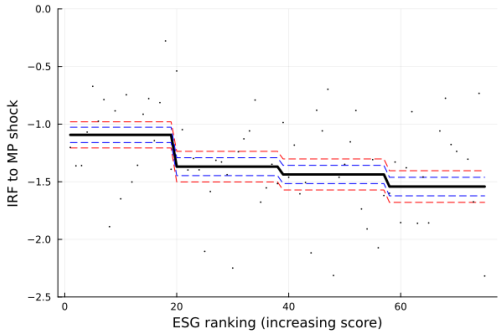

(b) France

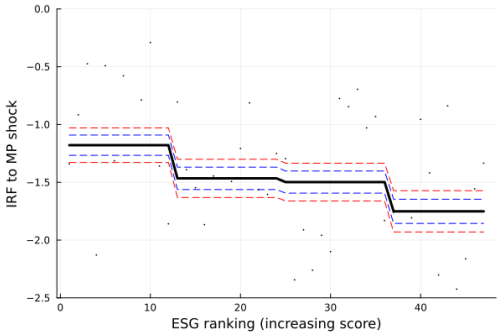

(c) Italy/Spain/Portugal

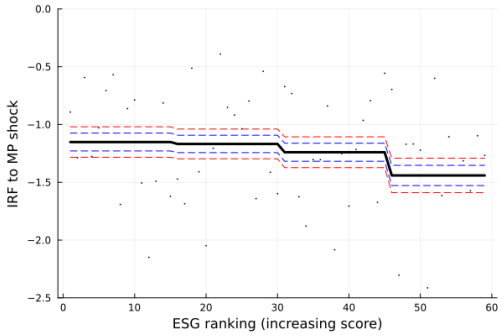

(d) Netherlands/Belgium/Finland/Austria

Figure A3.4: Minimum value of IRFs, with mean by ESG quartile (by country)

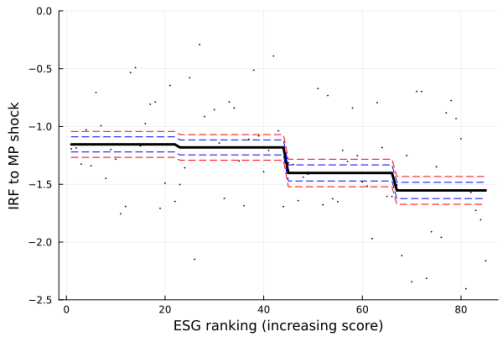

(a) Small companies

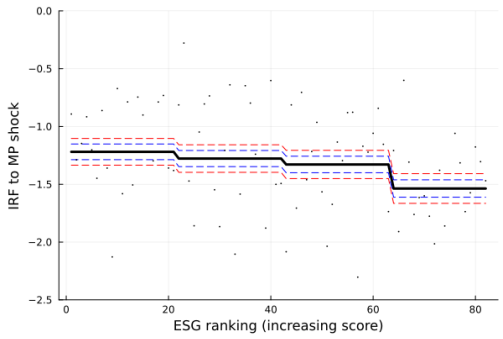

(b) Medium companies

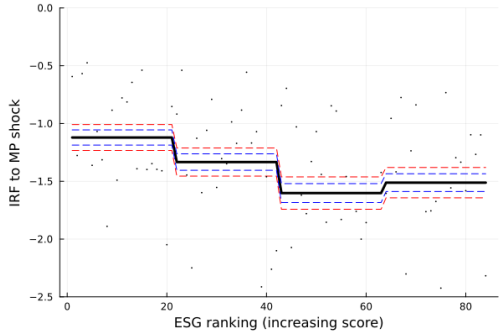

(c) Large companies

Figure A3.5: Minimum value of IRFs, with mean by ESG quartile (by company size)

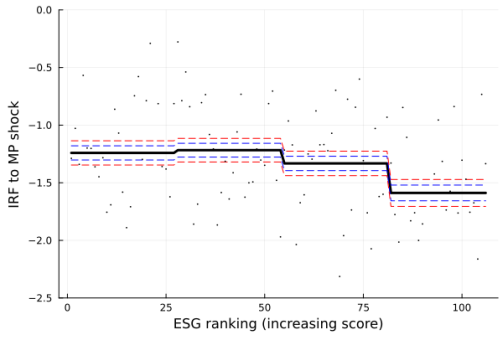

(a) Industry

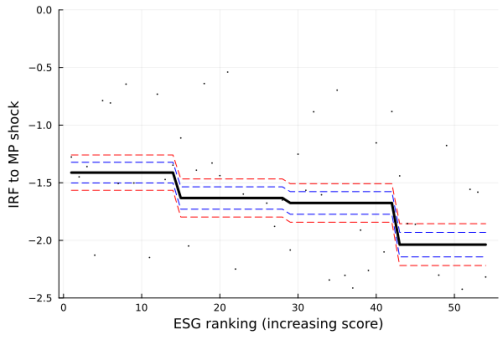

(b) Finance

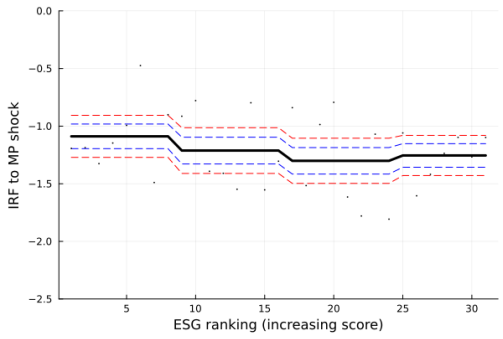

(c) ICT

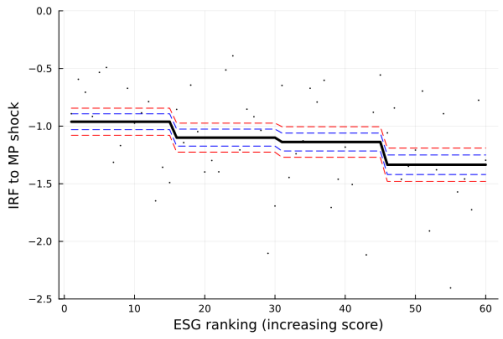

(d) Services

Figure A3.6: Minimum value of IRFs, with mean by ESG quartile (by sector)

A3.2 A CC-SVAR model for monetary policy shocks

A limitation of the SFAVAR approach followed in the main analysis is posed by a constraint on the stock-specific IRFs: since they are obtained indirectly through the dynamic effect of the structural shock on the factors, they are closely dependent on the estimated relationship between themselves and F_t, measured by the loadings Λ. For this reason, all the stock-specific stocks estimated through the SFAVAR are linearly dependent. In this robustness check, I relax this assumption by means of a CC-SVAR-like specification (Forni, Gambetti, et al. (2020)). The model is estimated as follows:

1. I apply PCA to estimate the r factors F_t. These are multiplied by loadings to obtain the common component χ_t for each stock, isolating the idiosyncratic part ξ_t

$$Y_t = \chi_t + \xi_t = \Lambda F_t + \xi_t \tag{3.22}$$

2. I estimate n bivariate VARs, each on a matrix containing a single common component in the χ_t matrix and the two-year rate series required for the identification of the shock

3. I apply the usual HF identification techniques on each of the bivariate VARs and obtain stock-specific IRFs computed isolating only the market component of the price reaction

The on-impact effects obtained from a CC-SVAR (where $r = 2$, as in the SFAVAR) are presented in Figure A3.7. Overall, the results confirm the conclusions reached with the the baseline model, with confidence bands being marginally narrower especially for medium-high ESG scores. The difference between the first and the middle quartiles is quite small (around 0.2 percent, in line with the SFAVAR), while the size of the Q4-Q1 gap is attested also by this model to be around 0.5 percent. Also for what concerns IRFs levels, results are in line with those presented with the baseline model.

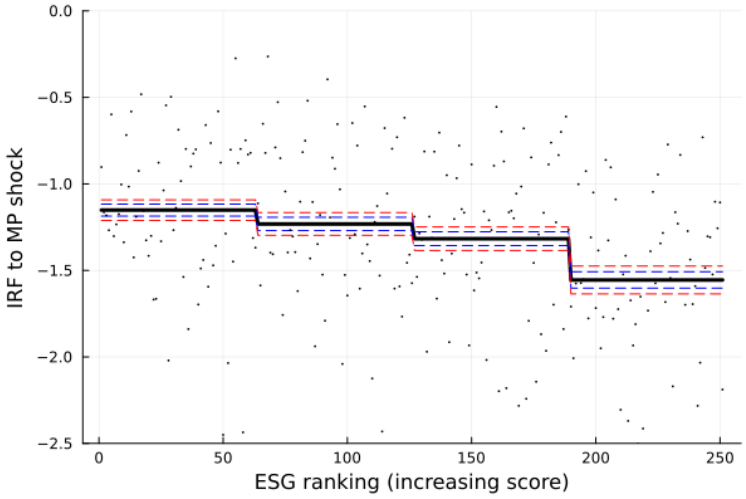

Figure A3.7: CC-SVAR on-impact effects, with mean by ESG quartile (whole sample)

A3.3 Other robustness checks

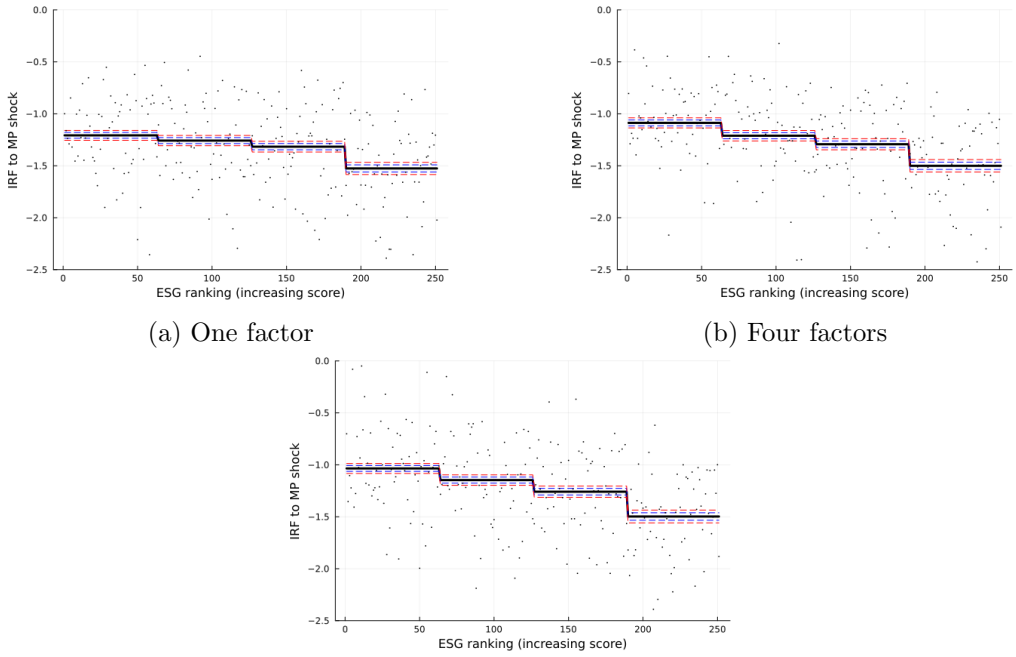

(a) One factor

(b) Four factors

(c) Five factors

Figure A3.8: On-impact effects with alternative number of factors

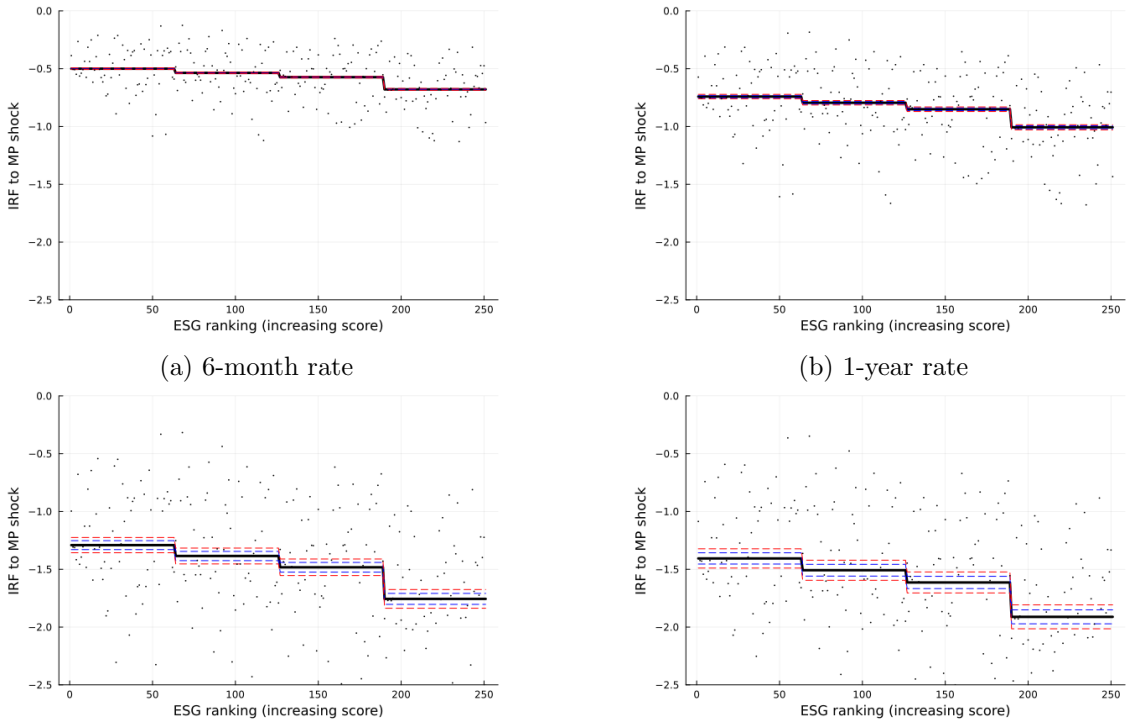

(a) 6-month rate

(b) 1-year rate

(c) 3-year rate

(d) 5-year rate

Figure A3.9: On-impact effects with alternative maturities for the identification TS

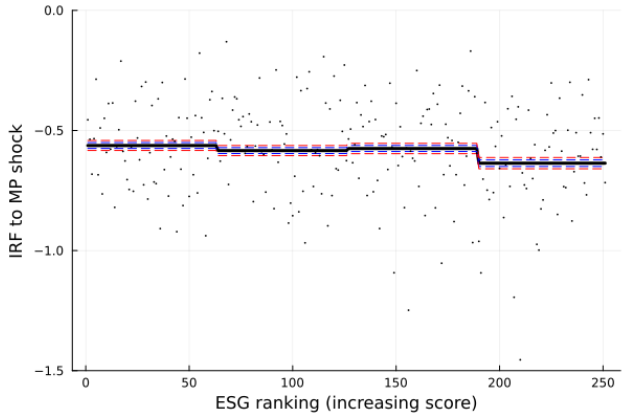

Figure A3.10: On-impact effects computed controlling for Fama and French's factors

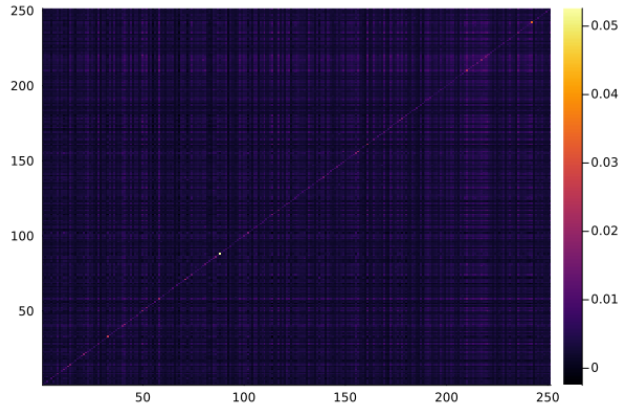

Figure A3.11: Heatmap of the POET covariance matrix

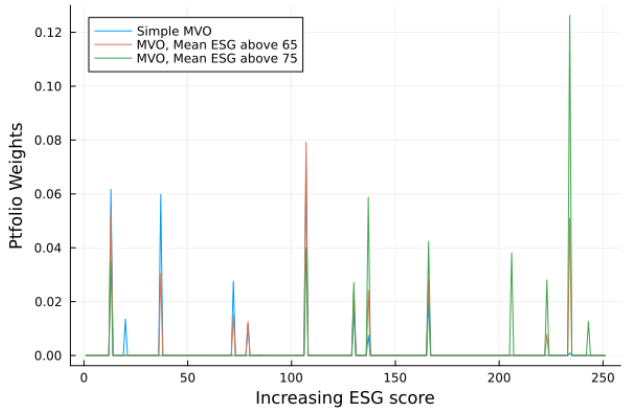

Figure A3.12: Optimal weights as a function of ESG integration, 0.5 percent return

Bibliography

Adrian, T., & Boyarchenko, N. (2012). Intermediary leverage cycles and financial stability. *Becker Friedman Institute for Research in Economics Working Paper*, (2012-010).

Adrian, T., Boyarchenko, N., & Giannone, D. (2019). Vulnerable growth. *American Economic Review*, *109*(4), 1263–89.

Ahn, S. C., & Horenstein, A. R. (2013). Eigenvalue ratio test for the number of factors. *Econometrica*, *81*(3), 1203–1227.

Aikman, D., Bridges, J., Hacioglu Hoke, S., O'Neill, C., & Raja, A. (2019). Credit, capital and crises: A gdp-at-risk approach.

Alessandri, P., Del Vecchio, L., Miglietta, A., et al. (2019). *Financial conditions and growth at risk in italy*. Banca d'Italia.

Alessi, L., Barigozzi, M., & Capasso, M. (2010). Improved penalization for determining the number of factors in approximate factor models. *Statistics & Probability Letters*, *80*(23-24), 1806–1813.

Alessi, L., Hirschbühl, D., & Rossi, A. (2022). *A sustainability transition on the move? evidence based on the disconnect from market fundamentals* (tech. rep.). JRC Working Papers in Economics and Finance.

Altavilla, C., Brugnolini, L., Gürkaynak, R. S., Motto, R., & Ragusa, G. (2019). Measuring euro area monetary policy. *Journal of Monetary Economics*, *108*, 162–179.

Andrade, P., & Ferroni, F. (2021). Delphic and odyssean monetary policy shocks: Evidence from the euro area. *Journal of Monetary Economics*, *117*, 816–832.

Azzalini, A., & Capitanio, A. (2003). Distributions generated by perturbation of symmetry with emphasis on a multivariate skew t-distribution. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, *65*(2), 367–389.

Bai, J., & Ng, S. (2002). Determining the number of factors in approximate factor models. *Econometrica*, *70*(1), 191–221.

Baldwin, R., & Wyplosz, C. (2019). *The economics of european integration*. McGraw Hill.

Bańbura, M., & Modugno, M. (2014). Maximum likelihood estimation of factor models on datasets with arbitrary pattern of missing data. *Journal of Applied Econometrics*, *29*(1), 133–160.

Barigozzi, M., Lissona, C., & Tonni, L. (2024). Ea-md-qd: Large euro area and euro member countries datasets for macroeconomic research. *Zenodo*.

Barigozzi, M., & Luciani, M. (2021). Measuring the output gap using large datasets. *The Review of Economics and Statistics*, 1–45.

Bauer, M. D., & Rudebusch, G. D. (2020). Interest rates under falling stars. *American Economic Review*, *110*(5), 1316–1354.

Bayoumi, T., & Eichengreen, B. (1992). Shocking aspects of european monetary unification.

Bayoumi, T., & Eichengreen, B. (1997). Ever closer to heaven? an optimum-currency-area index for european countries. *European economic review*, *41*(3-5), 761–770.

Belviso, F., & Milani, F. (2006). Structural factor-augmented vars (sfavars) and the effects of monetary policy. *Topics in Macroeconomics*, *6*(3).

Benchora, I., Leroy, A., & Raffestin, L. (2023). Is monetary policy transmission green?

Bernanke, B. S. (2005). Govenor bernanke speech on the global saving glut and the us current account deficit.

Bolton, P., & Kacperczyk, M. (2021). Do investors care about carbon risk? *Journal of Financial Economics*, *142*(2), 517–549.

Brand, C., Bielecki, M., & Penalver, A. (2018). The natural rate of interest: Estimates, drivers, and challenges to monetary policy. *ECB Occasional Paper*, (217).

Broadstock, D. C., Chan, K., Cheng, L. T., & Wang, X. (2021). The role of esg performance during times of financial crisis: Evidence from covid-19 in china. *Finance research letters*, *38*, 101716.

Brownlees, C., & Souza, A. B. (2021). Backtesting global growth-at-risk. *Journal of Monetary Economics*, *118*, 312–330.

Brunnermeier, M. K., & Sannikov, Y. (2014). A macroeconomic model with a financial sector. *American Economic Review*, *104*(2), 379–421.

Busetti, F., Caivano, M., Delle Monache, D., & Pacella, C. (2021). The time-varying risk of italian gdp. *Economic Modelling*, *101*, 105522.

Campbell, J. Y., & Viceira, L. M. (2002). *Strategic asset allocation: Portfolio choice for long-term investors*. Clarendon Lectures in Economic.

Campos, N. F., & Macchiarelli, C. (2016). Core and periphery in the european monetary union: Bayoumi and eichengreen 25 years later. *Economics Letters*, *147*, 127–130.

Campos, N. F., & Macchiarelli, C. (2021). The dynamics of core and periphery in the european monetary union: A new approach. *Journal of International Money and Finance*, *112*, 102325.

Cattell, R. B. (1966). The scree test for the number of factors. *Multivariate behavioral research*, *1*(2), 245–276.

Cecchetti, S. G. (2006). Gdp at risk: A framework for monetary policy responses to asset price movements. *NBER Working Paper*, (C0012).

Cecchetti, S. G. (2008). Measuring the macroeconomic risks posed by asset price booms. In *Asset prices and monetary policy* (pp. 9–43). University of Chicago Press.

Chava, S. (2014). Environmental externalities and cost of capital. *Management science*, *60*(9), 2223–2247.

Chavleishvili, S., & Kremer, M. (2023). Measuring systemic financial stress and its risks for growth.

Chavleishvili, S., & Manganelli, S. (2024). Forecasting and stress testing with quantile vector autoregression. *Journal of Applied Econometrics*, *39*(1), 66–85.

Chen, L., Dolado, J. J., & Gonzalo, J. (2021). Quantile factor models. *Econometrica*, *89*(2), 875–910.

Christensen, J. H., & Mouabbi, S. (2024). The natural rate of interest in the euro area: Evidence from inflation-indexed bonds.

Coe, P. J., & Vahey, S. P. (2020). Financial conditions and the risks to economic growth in the united states since 1875.

Correal, M. E., & Peña, D. (2008). Threshold dynamic factor model. *Revista Colombiana de Estadística*, *31*(2), 183–192.

de Haan, J., Jacobs, J. P., & Zijm, R. (2024). Coherence of output gaps in the euro area: The impact of the covid-19 shock. *European Journal of Political Economy*, *84*, 102369.

Del Negro, M., Giannone, D., Giannoni, M. P., & Tambalotti, A. (2017). Safety, liquidity, and the natural rate of interest. *Brookings Papers on Economic Activity*, *2017*(1), 235–316.

Del Negro, M., Giannone, D., Giannoni, M. P., & Tambalotti, A. (2019). Global trends in interest rates. *Journal of International Economics*, *118*, 248–262.

Delle Monache, D., De Polis, A., & Petrella, I. (2023). Modeling and forecasting macroeconomic downside risk. *Journal of Business & Economic Statistics*, 1–27.

Dodge, Y., & Whittaker, J. (2009). Partial quantile regression. *Metrika*, *70*, 35–57.

Durbin, J., & Koopman, S. J. (2001). Time series analysis by state space methods. *Time Series Analysis by State Space Methods*, 253.

Eggertsson, G. B., Mehrotra, N. R., & Robbins, J. A. (2019). A model of secular stagnation: Theory and quantitative evaluation. *American Economic Journal: Macroeconomics*, *11*(1), 1–48.

EU. (2022). Csrd: Directive (eu) 2022/2464 of the european parliament and of the council of 14 december 2022 amending regulation (eu) no 537/2014, directive 2004/109/ec, directive 2006/43/ec and directive 2013/34/eu, as regards corporate sustainability reporting (text with eea relevance).

Fama, E. F., & French, K. R. (1993). Common risk factors in the returns on stocks and bonds. *Journal of financial economics*, *33*(1), 3–56.

Fama, E. F., & French's, K. R. (2015). A five-factor asset pricing model. *Journal of financial economics*, *116*(1), 1–22.

Fan, J., Liao, Y., & Mincheva, M. (2013). Large covariance estimation by thresholding principal orthogonal complements. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, *75*(4), 603–680.

Ferrara, L., Mogliani, M., & Sahuc, J.-G. (2022). High-frequency monitoring of growth at risk. *International Journal of Forecasting*, *38*(2), 582–595.

Ferreira, T. R., & Shousha, S. (2023). Determinants of global neutral interest rates. *Journal of International Economics*, *145*, 103833.

Figueres, J. M., & Jarociński, M. (2020). Vulnerable growth in the euro area: Measuring the financial conditions. *Economics Letters*, *191*, 109126.

Fiorentini, G., Galesi, A., Pérez-Quirós, G., & Sentana, E. (2018). The rise and fall of the natural interest rate.

Forni, M., Gambetti, L., Marco, M., & Sala, L. (2020). Common component structural vars.

Forni, M., Giannone, D., Lippi, M., & Reichlin, L. (2009). Opening the black box: Structural factor models with large cross sections. *Econometric Theory*, *25*(5), 1319–1347.

Gächter, M., Hasler, E., & Huber, F. (2025). A tale of two tails: 130 years of growth at risk. *Macroeconomic Dynamics*, *29*, e46.

Galvão, A. B., & Owyang, M. T. (2018). Financial stress regimes and the macroeconomy. *Journal of Money, Credit and Banking*, *50*(7), 1479–1505.

Giannetti, M., Jasova, M., Loumioti, M., & Mendicino, C. (2023). "glossy green" banks: The disconnect between environmental disclosures and lending activities. *Banks: The Disconnect between Environmental Disclosures and Lending Activities (December, 2023). ECB Working Paper*, (2023/2882).

Giglio, S., Kelly, B., & Pruitt, S. (2016). Systemic risk and the macroeconomy: An empirical evaluation. *Journal of Financial Economics*, *119*(3), 457–471.

Giraitis, L., Kapetanios, G., & Yates, T. (2014). Inference on stochastic time-varying coefficient models. *Journal of Econometrics*, *179*(1), 46–65.

Godfrey, P. C., Merrill, C. B., & Hansen, J. M. (2009). The relationship between corporate social responsibility and shareholder value: An empirical test of the risk management hypothesis. *Strategic management journal*, *30*(4), 425–445.

GSIA. (2023). *2022 global sustainable investment review.* (tech. rep.).

Gürkaynak, R. S., Sack, B., & Swanson, E. (2005). The sensitivity of long-term interest rates to economic news: Evidence and implications for macroeconomic models. *American economic review*, *95*(1), 425–436.

Habib, M. M., Stracca, L., & Venditti, F. (2020). The fundamentals of safe assets. *Journal of International Money and Finance*, *102*, 102119.

Hartwig, B., Meinerding, C., & Schüler, Y. S. (2021). Identifying indicators of systemic risk. *Journal of International Economics*, *132*, 103512.

Hoepner, A. G., Oikonomou, I., Sautner, Z., Starks, L. T., & Zhou, X. Y. (2024). Esg shareholder engagement and downside risk. *Review of Finance*, *28*(2), 483–510.

Hollo, D., Kremer, M., & Lo Duca, M. (2012). Ciss-a composite indicator of systemic stress in the financial system.

Holston, K., Laubach, T., & Williams, J. C. (2017). Measuring the natural rate of interest: International trends and determinants. *Journal of International Economics*, *108*, S59–S75.

Holston, K., Laubach, T., & Williams, J. C. (2023). Measuring the natural rate of interest after covid-19. *FRB of New York Staff Report*, (1063).

Hsu, P.-H., Li, K., & Tsou, C.-Y. (2023). The pollution premium. *The Journal of Finance*, *78*(3), 1343–1392.

Ilhan, E., Sautner, Z., & Vilkov, G. (2021). Carbon tail risk. *The Review of Financial Studies*, *34*(3), 1540–1571.

Iseringhausen, M., Petrella, I., & Theodoridis, K. (2023). Aggregate skewness and the business cycle. *Review of Economics and Statistics*, 1–38.

Johannsen, B. K., & Mertens, E. (2016). A time-series model of interest rates with the effective lower bound. *Journal of Money, Credit and Banking*.

Juselius, M., Borio, C. E., Disyatat, P., & Drehmann, M. (2016). Monetary policy, the financial cycle and ultralow interest rates. *Bank of Finland Research Discussion Paper*, (24).

Kiley, M. T. (2015). What can the data tell us about the equilibrium real interest rate?

Koenker, R., & Bassett Jr, G. (1978). Regression quantiles. *Econometrica: journal of the Econometric Society*, 33–50.

Koenker, R., & Machado, J. A. (1999). Goodness of fit and related inference processes for quantile regression. *Journal of the american statistical association*, *94*(448), 1296–1310.

Krishnamurthy, A., & Vissing-Jorgensen, A. (2012). The aggregate demand for treasury debt. *Journal of Political Economy*, *120*(2), 233–267.

Krugman, P. (2013). Revenge of the optimum currency area. *NBER macroeconomics annual*, *27*(1), 439–448.

Kunovac, D., Palenzuela, D. R., & Sun, Y. (2022). A new optimum currency area index for the euro area.

Lang, J. H., Rusnák, M., & Greiwe, M. (2025). Medium-term growth-at-risk in the euro area. *IMF Economic Review*, 1–43.

Laubach, T., & Williams, J. C. (2003). Measuring the natural rate of interest. *Review of Economics and Statistics*, *85*(4), 1063–1070.

Laubach, T., & Williams, J. C. (2016). Measuring the natural rate of interest redux. *Business Economics*, *51*(2), 57–67.

Lins, K. V., Servaes, H., & Tamayo, A. (2017). Social capital, trust, and firm performance: The value of corporate social responsibility during the financial crisis. *The Journal of Finance*, *72*(4), 1785–1824.

Loria, F., Matthes, C., & Zhang, D. (2022). Assessing macroeconomic tail risk. *Available at SSRN 4002665*.

Lubik, T. A., Matthes, C., et al. (2015). Calculating the natural rate of interest: A comparison of two alternative approaches. *Richmond Fed Economic Brief*, (Oct).

Luo, D. (2022). Esg, liquidity, and stock returns. *Journal of International Financial Markets, Institutions and Money*, 101526.

Massacci, D. (2017). Least squares estimation of large dimensional threshold factor models. *Journal of Econometrics*, *197*(1), 101–129.

Matsumura, E. M., Prakash, R., & Vera-Munoz, S. C. (2014). Firm-value effects of carbon emissions and carbon disclosures. *The accounting review*, *89*(2), 695–724.

McCracken, M. W., & Ng, S. (2016). Fred-md: A monthly database for macroeconomic research. *Journal of Business & Economic Statistics*, *34*(4), 574–589.

Merton, R. C., et al. (1987). A simple model of capital market equilibrium with incomplete information.

Mundell, R. A. (1961). A theory of optimum currency areas. *The American economic review*, *51*(4), 657–665.

Nofsinger, J., & Varma, A. (2014). Socially responsible funds and market crises. *Journal of Banking & Finance*, *48*, 180–193.

Onatski, A. (2010). Determining the number of factors from empirical distribution of eigenvalues. *The Review of Economics and Statistics*, *92*(4), 1004–1016.

Pástor, L., Stambaugh, R. F., & Taylor, L. A. (2021). Sustainable investing in equilibrium. *Journal of Financial Economics*, *142*(2), 550–571.

Patozi, A. (2023). Green transmission: Monetary policy in the age of esg. *Available at SSRN 4327122*.

Pedersen, L. H., Fitzgibbons, S., & Pomorski, L. (2021). Responsible investing: The esg-efficient frontier. *Journal of Financial Economics*, *142*(2), 572–597.

Pisani, F., & Russo, G. (2021). Sustainable finance and covid-19: The reaction of esg funds to the 2020 crisis. *Sustainability*, *13*(23), 13253.

Rachel, Ł., & Summers, L. H. (2019). *On secular stagnation in the industrialized world* (tech. rep.). National Bureau of Economic Research.

Sargent, T. J., & Sims, C. A. (1977). Business cycle modeling without pretending to have too much a priori economic theory. *New methods in business cycle research*, *1*, 145–168.

Sassen, R., Hinze, A.-K., & Hardeck, I. (2016). Impact of esg factors on firm risk in europe. *Journal of business economics*, *86*(8), 867–904.

Sharpe, W. F. (1966). Mutual fund performance. *The Journal of business*, *39*(1), 119–138.

Sims, C. A. (1980). Macroeconomics and reality. *Econometrica: journal of the Econometric Society*, 1–48.

Stock, J. H., & Watson, M. W. (2016). Dynamic factor models, factor-augmented vector autoregressions, and structural vector autoregressions in macroeconomics. In *Handbook of macroeconomics* (pp. 415–525, Vol. 2). Elsevier.

Stock, J. H., & Watson, M. W. (2018). Identification and estimation of dynamic causal effects in macroeconomics using external instruments. *The Economic Journal*, *128*(610), 917–948.

Summers, L. H. (2014). Us economic prospects: Secular stagnation, hysteresis, and the zero lower bound. *Business economics*, *49*(2), 65–73.

Szendrei, T., & Varga, K. (2023). Revisiting vulnerable growth in the euro area: Identifying the role of financial conditions in the distribution. *Economics Letters*, 110990.

Trapani, L. (2018). A randomized sequential procedure to determine the number of factors. *Journal of the American Statistical Association*, *113*(523), 1341–1349.

Waddock, S. A., & Graves, S. B. (1997). Quality of management and quality of stakeholder relations: Are they synonymous? *Business & society*, *36*(3), 250–279.

Watson, M. W., & Engle, R. F. (1983). Alternative algorithms for the estimation of dynamic factor, mimic and varying coefficient regression models. *Journal of Econometrics*, *23*(3), 385–400.

Wicksell, K. (1898). *Interest and prices*. Ludwig von Mises Institute.

Williams, J. C., & Wu, J. C. (2017). Comments and discussion. *Brookings Papers on Economic Activity*, 295–316.

Zerbib, O. D. (2022). A sustainable capital asset pricing model (s-capm): Evidence from environmental integration and sin stock exclusion. *Review of Finance*, *26*(6), 1345–1388.