UNIVERSITA' DEGLI STUDI DI PADOVA

DIPARTIMENTO DI SCIENZE ECONOMICHE ED AZIENDALI "M. FANNO"

CORSO DI LAUREA MAGISTRALE IN ECONOMICS AND FINANCE

TESI DI LAUREA

"INFLATION AT RISK IN THE EURO AREA"

RELATORE:

CH.MO PROF. EFREM CASTELNUOVO

LAUREANDO: MATTEO GIRELLI

MATRICOLA N. 2038506

ANNO ACCADEMICO 2023 – 2024

Dichiaro di aver preso visione del "Regolamento antiplagio" approvato dal Consiglio del Dipartimento di Scienze Economiche e Aziendali e, consapevole delle conseguenze derivanti da dichiarazioni mendaci, dichiaro che il presente lavoro non è già stato sottoposto, in tutto o in parte, per il conseguimento di un titolo accademico in altre Università italiane o straniere. Dichiaro inoltre che tutte le fonti utilizzate per la realizzazione del presente lavoro, inclusi i materiali digitali, sono state correttamente citate nel corpo del testo e nella sezione 'Riferimenti bibliografici'.

I hereby declare that I have read and understood the "Anti-plagiarism rules and regulations" approved by the Council of the Department of Economics and Management and I am aware of the consequences of making false statements. I declare that this piece of work has not been previously submitted – either fully or partially – for fulfilling the requirements of an academic degree, whether in Italy or abroad. Furthermore, I declare that the references used for this work – including the digital materials – have been appropriately cited and acknowledged in the text and in the section 'References'.

Firma (signature) ……Matteo……Girelli…

Università degli studi di Padova

Department of Economics and Management "M. Fanno"

Master's degree in Economics and Finance

Dissertation:

Inflation at Risk in the Euro Area

Supervisor
Professor Efrem Castelnuovo
Università degli studi di Padova

Master Candidate
Matteo Girelli

Student Number
2038506

Academic Year

2023-2024

I hereby declare that I have read and understood the "Anti-plagiarism rules and regulations" approved by the Council of the Department of Economics and Management and I am aware of the consequences of making false statements. I declare that this piece of work has not been previously submitted – either fully or partially – for fulfilling the requirements of an academic degree, whether in Italy or abroad. Furthermore, I declare that the references used for this work – including the digital materials – have been appropriately cited and acknowledged in the text and in the section 'References'.

Matteo Girelli

Signature

Dedication:

I dedicate this work to my family, who have supported me in every way throughout my studies.
I dedicate this work to the friends I made during this Master's degree, with whom I have shared years of hardship, disappointments and successes, that we faced together. I could never ask for better friends.
I dedicate this to my girlfriend, with whom I lived during the preparation of our master's theses, for having always believed in me, as well as for having been an inspiration and a fan in equal measure.
Lastly, I dedicate this work to myself, for proving what I can do, and for letting it be enough.

Abstract

I study Inflation at Risk for four euro area countries: Italy, Germany, France and Spain.

I model a simple Phillips' curve relation and estimate a series of two-steps quantile regressions, at different forecasting horizons. I create two measures of upside risk: expected longrise and upside relative entropy, that correlate across countries to different degrees. The fitting of a theoretical distribution allows me to study the moments of the distribution of the Phillips curve; this yields interesting findings for both inflation at risk and uncertainty literatures.

I run a SVAR to assess the effects of central bank information (CBI) shocks on these risk measures, as well as other economic variables, comparing the IRFs across countries. CBI shocks cause an increase in interbank lending rates. In all countries except Germany, inflation responds positively, unemployment negatively and expected longrise increases. Germany is insulated from the effects of CBI shocks, in the sense that its IRFs are not statistically significant. Upside relative entropy responds heterogeneously across countries.

An extension that accounts for financial conditions shows that this heterogeneity is present also for downside risks, that also matter in terms of reaching pricing stability.

The results of this dissertation provide an inflation at risk perspective on questions of optimality of the single currency area, putting into question aspects such as feasibility of the mandate and benefits of a single monetary policy, in an area characterised by latent structural differences that give rise to fundamental heterogeneities in inflation risk.

Contents

Listing of figures

Listing of tables

Listing of acronyms

EA .... Euro Area

G4 .... Italy, Germany, France, Spain

IaR .... Inflation at Risk

QR .... Quantile Regression

VAR .... Vector Autoregression

SVAR .... Structural Vector Autoregression

IRF .... Impulse Response Function

CB .... Central Bank

ECB .... European Central Bank

Fed .... Federal Reserve Bank

VaR .... Value at Risk

CVaR .... Conditional Value at Risk

IMF .... International Monetary Fund

BIS .... Bank of International Settlements

TS .... Time series

CBI .... Central Bank Information

1 Introduction

1.1 Motivation

Setting the stage The central banks' mandates are the formalisation of the role they play in the economy. They are mainly tasked to maintain financial stability and price stability, although some central banks have also got additional objectives — for instance, the Fed's maximum employment mandate. In my dissertation, I focus on the ECB and the set of national central banks that together form the Eurosystem.

The main goal of the Eurosystem is summarised well by citing the ECB's website:

*"The European Central Bank and the national central banks together constitute the Eurosystem, the central banking system of the euro area. The main objective of the Eurosystem is to maintain **price stability**: safeguarding the value of the euro."*

The task of maintaining price stability, that is, to control inflation, is far from a trivial one; indeed, one could argue that the objective of managing the business cycle's fluctuations, directly linked to inflation dynamics, in order to achieve the target (of 2% for the Euro Area) has been systematically failed. The difficulty is in part caused by the fact that economic variables often have non-standard distributions, displaying skewness, fat-tails or even multimodality. Because of their mandate, central banks tend to focus much of their work on forecasting key economic indicators. Traditionally, economic forecasts rely on linear models providing estimates of the conditional mean (or mode) of such indicators; but given the point made above about the complexity of dealing with such variables in real scenarios, more information about the dynamics of economic variables is needed.

Recently, the focus has been shifting to studying and forecasting the entire conditional distribution of such variables. This new approach is briefly mentioned by Mario Draghi's speech in Sintra (2019)[1]:

"Monetary policy responded first in the summer of 2012 by acting to defuse the sovereign debt crisis, which had evolved from a ***tail risk*** *for inflation into a material threat to price stability"*

where "tail risk" refers to the risk of extreme and unexpected realizations of inflation, represented by the right tail of the predicted distribution. Traditional methods cannot provide us with such information, since they are mostly limited to modal forecasts.

I cite two more statements about the need for a new approach in monetary policy. The first is given by Yellen (2017)[2]

"The outlook is subject to considerable uncertainty from multiple sources, and dealing with these uncertainties is an important feature of policymaking."

while the second by from Greenspan (2004). [3]

"Uncertainty is not just an important feature of the monetary policy landscape; it is the defining characteristic of that landscape. The conduct of monetary policy in the United States at its core involves crucial elements of risk management, a process that requires an understanding of the many sources of risk and uncertainty that policymakers face and the quantifying of those risks when possible."

Risk or Uncertainty? At this point, given the previous citations' reference to risk and uncertainty, it is dutiful to define and distinguish generally the two terms; in doing so, I set the scope of my work.

Risk refers to situations where the outcome of a random variable is unknown but governed by a known probability distribution. Outcomes are associated to probabilities and risk can be quantified.

[1] At the time, Mario Draghi was still President of the ECB.
[2] At the time, Yellen was the Chair of the Federal Reserve.
[3] Given at a symposium sponsored by the Federal Reserve Bank of Kansas City in 2003.

For example, the risk of death for people with particular age and other characteristics can be assessed for life insurance; we could say that risk is "what we know that we don't know".

Uncertainty, on the other hand, is in layman terms "what we don't know that we don't know". It characterises cases in which one is not only unsure about the realisation of the outcome, but cannot easily associate a reliable confidence interval to the potential outcomes. Such situations are, for instance, Black Swan events like stock market crashes and the "Global Financial Crisis". The following table is a summary of the differences between these two concepts:

	Risk	**Uncertainty**
Interpretation	Quantifiable possibility that the realisation is different from the expected outcome (plus associated benefits/costs)	Unpredictability or unreliability of the prediction of an event; inaccurate confidence estimates and unknown (benefit or) damage by such an event.
Outcome	Can lead to both favourable or unfavourable outcomes.	
Computability	Can be quantified or measured accurately.	Difficult to measure reliably.
Information	Risk is based on known information and data.	Uncertainty arises from lack of information or incomplete knowledge.
Controllability	Manageable or mitigable.	Out of direct control.
Probability	Can be assigned a probability or likelihood.	Often involves unknown probabilities.

Table 1.1: Comparison of Risk and Uncertainty

At the time of Greenspan (2004), macroeconometrics practitioners mostly talked about uncertainty as a sort of continuum, ranging from well defined risks to the truly unknown. This interpretation encompasses both "Knightian uncertainty," in which the probability distribution of outcomes is unknown, and risk, in which uncertainty is delimited by a known probability distribution.

Nowadays, the vast literature on uncertainty has made leaps forward in its study and understanding since Greenspan (2004). Its definition too has altered; a more modern and precise one is "the change in the second moment of the distribution given a certain fixed value of the first moment, therefore a mean-preserving shock" (Castelnuovo, 2023).

In this project, however, I will purposely deal with this "continuum" described in Greenspan (2004). This is because I believe that the approach used here allows me to obtain results that can potentially contribute to both literatures, without needing to compartmentalise one result as "strictly about risk" or as "strictly about uncertainty".

Moreover, while both macroeconomic risks and the study of uncertainty are very interesting and hot research alleys in macroeconomics, tackling uncertainty in a specialised way is a much more complex effort, requiring advanced training and techniques.

Paradigm shifts Between the end of the 20th and the start of the 21st century, globalisation increased in pace and became more widespread, the world became more interconnected and supply chains longer (as well as, arguably, more fragile). National economies became more open, the job markets more mobile (in favour of lower income countries) and both financial and real investments caused larger amounts of money than ever before to start flowing continuously across countries.

Clearly, such "developments" caused big structural changes to the economies they took place in. For instance, the primary and secondary sectors in advanced countries suffered enormously from delocalisation to and competition from cheaper, less regulated developing countries. Advanced countries' trade balances suffered and they started to become more and more indebted.

In addition, the financial sector growth led its role in the economy to become increasingly important. The exponential surge of financial innovation, driven by deregulation and technological innovations, enabled firms and investors to share risks and to lower informational frictions, as well as transaction costs, spurring economic growth through efficiency gains. However, this also enabled speculators to profit from market moves that were detrimental to financial stability. Flash financial crises became more and more frequent and their effects on the real economy increasingly more deleterious. Because of these paradigm shifts, financial risk, economic risk and their interplay became a focus point of central banks. More flexible, robust and informative approaches to the forecast and study of economic variables became necessary.

1.2 Literature

The foundations of the field of "economic risk management" were laid down in the early 2000s by Greenspan (2004) and Kilian and Manganelli (2007). As the name suggests, the field adopts many ideas and methods coming from quantitative finance, computational finance and risk management, with the aim of applying these to the study of economic variables; the field also relies heavily on data analysis and Bayesian statistics for inference and superior forecasting performance.

Much of the effort is spent on improving the forecasts of potential outcomes, as well as the estimates of risk and uncertainty, through the use of linear and non-linear models and advanced techniques enabling the study of the whole distribution of random variables.

Recently, there has been a growing interest in this field, in particular by central banks and institutions — to name a few: ECB, IMF, BIS and Fed.

Here are some of the most relevant papers on macroeconomic risks, from these institutions or associated researchers:

- European Central Bank. (2019)
- Korobilis et al. (2021)
- Adams et al. (2021)
- Ciccarelli et al. (2024)
- Banerjee et al. (2024)

The most widespread applications of risk analysis in macroeconometrics and macro-finance are on GDP growth and inflation, but the approach can potentially be applied to other variables, such as oil price.

In the next paragraphs, I will talk about some of the most important papers for each application mentioned above; as for the literature that deals with the more technical aspects of my dissertation, I will cite the relevant papers within the related sections of the dissertation's body, or in Appendix B.

GROWTH AT RISK The economic indicator whose risk has been studied the most is GDP growth. The seminal paper of Adrian et al. (2019) "Vulnerable growth" has been a reference point on the topic. They study the American GDP growth (over a forecasting horizon) conditional on economic conditions, represented by the current period's GDP growth, and on financial conditions, captured by the NFCI index. They take a sample period from the first quarter of 1975 to the last quarter of 2015, focusing on downside risks. The sample thus includes periods of marked downside volatility, like during the Global Financial Crisis of 2008 (GFC). They devise the two-steps quantile regression, a novel econometric methodology: first they run quantile regression estimating a series of quantiles, then they fit a known distribution by matching the empirical quantiles to the theoretical ones; this allows them to estimate the moments of the distribution and to calculate risk measures. They find that deteriorating financial conditions lead to an increase in the conditional volatility and a decline in the conditional mean of GDP growth, affecting the lower quantiles more than the upper ones. The amplification mechanisms in the financial sector contribute to the observed dynamics of growth vulnerability, where downside risks to GDP growth increase with tighter financial conditions. They emphasize the link between financial stability and macroeconomic performance and suggest that DSGE models should account for nonlinear equilibrium relationships due to financial conditions. For these reasons, their findings have implications for both the macroeconometrics and macro-finance literatures.

Based on this method, Figueres and Jarociński (2020) study growth at risk in the whole Euro Area, from the first quarter of 1999 to the second quarter of 2018. Their regression equation is the same as in the previous paper, but their economic indicators are GDP growth for the EA and CISS, or "Composite Indicator of Systemic Stress" (in place of NFCI). This indicator, developed by Holló et al. (2012), is the best indicator of financial conditions among the ones considered, according to the authors.

They find that, as expected, financial conditions are negatively related to GDP growth, while current GDP growth is positively related to future GDP growth; also, the lower quantiles of the distribution of GDP growth are more sensitive to financial conditions than the upper ones. As CISS increases, output growth becomes more negatively skewed, and the relationship is nonlinear at the lower quantiles. Financial conditions explain well the volatility in the lower tail of the conditional distribution of GDP growth. Moreover, the conditional mean and variance of output growth correlate negatively. They infer this by following the two-steps quantile regression approach of Adrian et al. (2019).

Lastly, they reproduce concisely the analysis for inflation and do not find nonlinearities in the relationship of CISS and Inflation; additionally, the conditional mean and variance of inflation correlate positively (hence, inflation is more volatile when higher).

Inflation at Risk An example of a research paper thoroughly studying inflation at risk is the one by (López-Salido and Loria, 2024) for the USA.

Their paper examines the tails of the inflation distribution over a forecasting horizon, using quantile regressions in a panel of OECD countries. Their regression model is an Augmented New Keynesian Phillips curve, with average inflation over the previous period, long term inflation expectations, unemployment gap, oil price and financial conditions. They explore the variability in inflation risks during periods like the Global Financial Crisis and the Covid-19 pandemic, highlighting the nonlinear impact of financial conditions on inflation predictions. They also compare their model with alternative ones, using financial market quotes and survey data, finding that financial conditions significantly influence downside inflation risks. Additionally, they tackle the changing role of economic drivers of inflation risks over time and show the increasing importance of inflation expectations compared to other determinants.

Overall, despite stable average inflation, significant variability in inflation risks was observed, particularly during crises. Tight financial conditions were found to create substantial downside inflation risks, highlighting their importance in understanding inflation dynamics. Upside inflation risks emerged from fiscal stimulus and supply constraints during the pandemic, persisting even as inflation went down. Finally, the authors remark the need to consider tail risks and the role of financial conditions when assessing inflation dynamics, especially during periods of economic uncertainty.

CENTRAL BANK INFORMATION SHOCKS My dissertation also includes a SVAR identified through the use of a pure central bank information proxy, based on data from the EA-MPD (European Central Bank., 2019) and recalculated as in Jarociński and Karadi (2020) to fit my sample.

In Jarociński and Karadi (2020), the authors study monetary policy and central bank information shocks in the USA and Euro Area. For Europe, their dataset includes 280 ECB policy announcements from 1999 to 2016. Surprises are measured in narrow windows around press statements and conferences. The identification method combines standard High-Frequency Identification (HFI) and sign restrictions with Bayesian methods.

While HFI assumes that announcement surprises are affected only by announcement shocks, sign restrictions allow to differentiate between negative and positive co-movement shocks, based on interest rate and stock price movements. For the EA, they run SVARs with the proxies, German one-year government bond yield, observations from the STOXX 50 index, interpolated real GDP and GDP deflator, and BBB bond spread for financial conditions, comparing the results with the "entangled" case. The authors' main contributions are the proposal of a way to disentangle the monetary policy shocks and the information shocks, as well as the calculation of pure MP and CBI proxies for the euro area. They show that, in a structural VAR, the dynamic responses of economic variables to such shocks are different. Not purifying the monetary policy shock from the information component of CB's announcements causes to attenuate the estimated effects of monetary policy. Their approach finds stronger monetary transmission; moreover, the economy responds significantly and positively to this CBI shocks.

1.3 MY DISSERTATION

In this dissertation I study Inflation at risk, analysing the distribution of the average inflation rate over a forecasting horizon, conditional on present time inflation and unemployment rates, with a focus on upside risks. I employ the two-steps quantile regression devised by Adrian et al. (2019). The choice of these variables is done in order to emulate a very simple Phillips curve.

Next, I develop a measure of upside inflation risk, and I analyse the correlation in these measures across the Euro Area countries I study: Italy, Germany, France and Spain. This provides a different perspective on the optimality of the single currency area, in relation to the mandate of the central bank of managing inflation. Lastly, I study the impact of central banks' information shocks on inflation risks and other economic variables in a SVAR.

ORGANISATION The dissertation is roughly divided in the following parts:

- Overview and data description.
- Quantile Regression (QR) of the conditional distribution of the inflation rate over forecasting horizons
- Construction of the inflation risk measures.
  - matching quantiles and fitting a theoretical CQF
  - using the distribution to calculate the measures
- Comparison and correlation analysis of risk measures across countries
- SVAR analysis of response of economic indicators and risk measures to monetary policy shocks
  - Making use of a pure central bank information proxy.

In addition, there are two appendices: Appendix A extends the model to consider financial conditions and headline inflation rate Appendix B is a treatment of some technical elements.

2

Overview and data description

The following is the reference forecasting regression equation:

$$y_h = \alpha + \mathbf{x_t}'\boldsymbol{\beta} + \varepsilon_h \tag{2.1}$$

where:

$$y_h = \frac{1}{h}\sum_{i=1}^{h} y_{t+i}$$

y_h: discrete moving average of the inflation rate over the forecasting horizon[1].

and

$$\mathbf{x_t} = \begin{pmatrix} y_t \\ u_t \end{pmatrix}$$

$$y_t \text{ : Inflation rate at time } t$$

$$u_t \text{ : Unemployment at time } t$$

The use of a moving average is common in time-series analysis, because it allows to easily smooth out short-term fluctuations and to highlight longer-term trends or cycles; moreover, it is often more

[1] Adrian et al. (2019) use the notation y_{t+h} to denote the moving average of GDP growth over the forecasting horizon. I opted to change the notation in order to clarify that I am not forecasting the exact inflation rate realisation at time $t+h$, but the average over the horizon.

interesting to forecast the average value of a variable instead of the exact realisation.

I take the G4 countries' respective core inflation rates (less food and energy; growth rate from the same period in the previous year, not seasonally adjusted) and unemployment rates from FRED.

The sample goes from January 1991 to December 2019 (data are monthly). I purposely exclude the COVID-19 period from the analysis; this is because the outliers and structural breaks caused by the epidemic invalidate most of the time-series analyses, unless the extreme volatility and breaks are carefully modelled by making use of stochastic volatility models and other techniques, which ultimately lead to similar results as analyses that directly exclude the observations (European Central Bank., 2020).

The forecast horizons used are 6 months and 12 months.

Figure 2.1 shows the time series of inflation rate for the four countries.

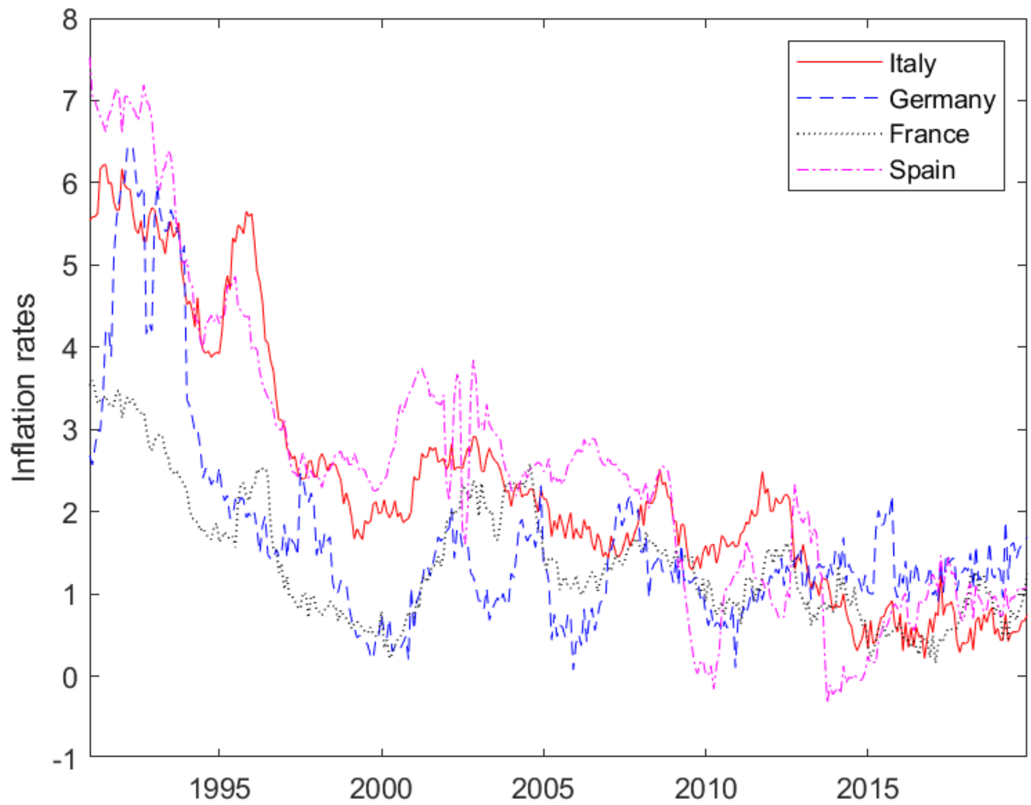

Figure 2.1: Inflation rate series, G4

Figure 2.2 compares inflation and unemployment rate series across countries. As we can see, since

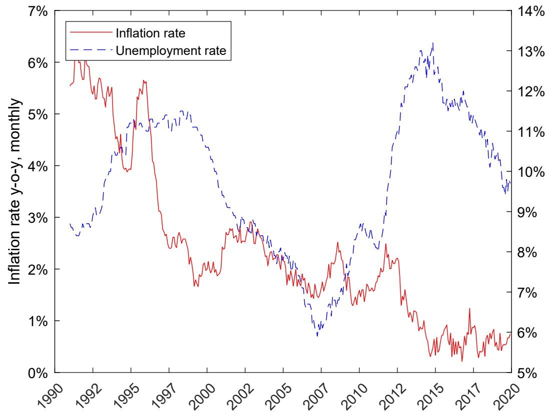

Italy

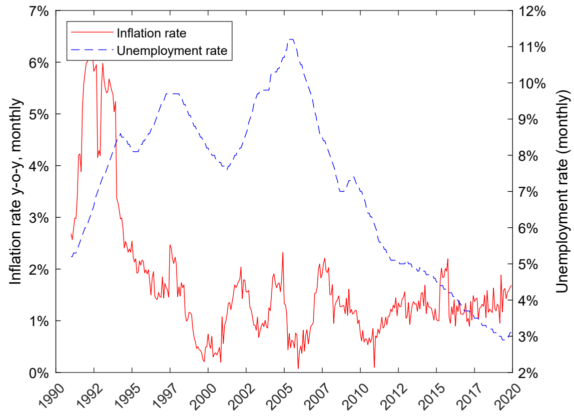

Germany

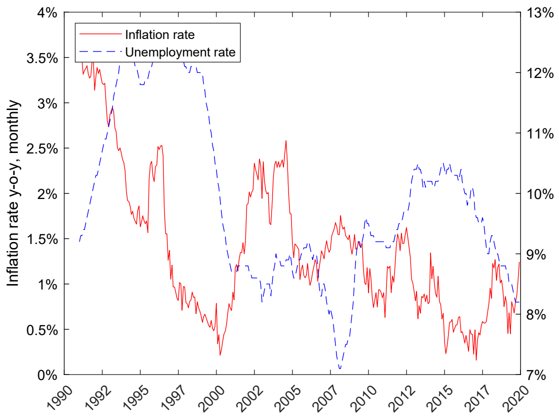

France

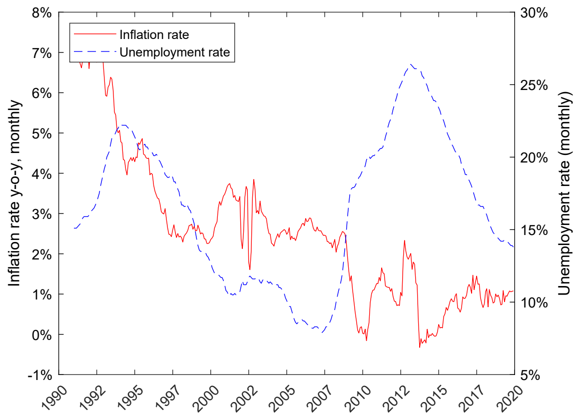

Spain

Figure 2.2: Inflation and Unemployment series

the introduction of the euro, the inflation rates have been co-moving in an increasing fashion. In order to estimate the parameters of the forecasting equation, I use quantile regression. In the next chapter, I present the QR estimation and explain the related results.

3

Quantile Regressions

3.1 The Quantile Regression estimator

I want to obtain an approximation of the distribution of the moving average of the inflation rate (over different forecasting horizons), conditional on the inflation and unemployment rates at the present time. An OLS regression would only provide an estimate of the expectation value of the conditional distribution of inflation; in terms of distributions, this is a modal estimate. This is clearly insufficient to characterise a statistical distribution.

One approach to remedy this is quantile regression. Developed by Koenker and Bassett (1978), QR allows to estimate the quantiles (one for each regression) of the conditional quantile function (CQF) of the dependent variable, y_h, conditionally on the regressors, x_t.

The quantile function is the inverse of the cumulative distribution function (CDF):

$$Q(p) = F^{-1}(p) \tag{3.1}$$

for $0 < p < 1$, where F is the CDF of a random variable X.

It maps a probability p to the value x of the random variable, so that the probability of the variable being less than or equal to x is p. The median is a special case where $p = 0.5$.

This concept is used, for instance, in finance, to calculate the value at risk (VaR) for returns, by setting p to certain level (a threshold) like 0.05. This allows to find the expected loss for extreme outcomes (fifth percentile and below) of the distribution of returns.

The conditional quantile function extends the concept of quantiles to conditional distributions. The CQF of Y at quantile τ, conditionally on the regressors $\mathbf{X}$, is given by:

$$Q_\tau(Y|\mathbf{X}) = F_y^{-1}(\tau|\mathbf{X}) \tag{3.2}$$

where $F_y(Y|\mathbf{X})$ is the cumulative distribution function of Y at y conditional on $\mathbf{X}$.

In other words, the CQF identifies the τth quantile of the distribution of Y, given the value x of the conditioning variables $\mathbf{X}$.

In reference to the previous example, this concept is used to calculate the Conditional Value at Risk for the conditional distribution of the returns.

To go back to the problem, I want to estimate some values (quantiles) of the CQF of y_h conditionally on $\mathbf{x_t}$, as in Equation 2.1.

Koenker and Bassett (1978) show that the quantile regression estimator $\hat{\beta}_\tau$ is an unbiased and consistent *linear* estimator of CQF [1].

The following is the linear programming problem of the QR estimator, for regressor *i*:

$$\begin{aligned}\hat{\beta}_{i\tau} = \arg \min_{\beta_{i\tau}\in\mathbb{R}^k} \sum_{t=1}^{T-h} \big(\tau \cdot \mathbf{1}_{(y_h \geq x_{it}\beta_i)} |y_h - x_{it}\beta_\tau| \\ +(1-\tau) \cdot \mathbf{1}_{(y_h < x_{it}\beta_i)} |y_h - x_t\beta_{i\tau}|\big)\end{aligned} \tag{3.3}$$

where $\mathbf{1}$ denotes the indicator function[2].

The predicted value from the quantile regression above is the quantile of y_h conditional on $\mathbf{x_t}$:

$$\hat{Q}_{y_h|\mathbf{x_t}}(\tau|\mathbf{x_t}) = \mathbf{x_t}\hat{\boldsymbol{\beta}}_\tau \tag{3.4}$$

Below I summarise the advantages of QR over OLS:

- Objective Function:
  - OLS estimates the conditional mean through the minimisation of the sum of squared residuals.
  - QR minimises the weighted sum of absolute residuals and thus is able to estimate the (conditional) median or any other quantile.

[1] Instead, OLS solves a linear programming problem for exactly the conditional expectation function (CEF).
[2] See Appendix B for a review.

- Assumptions:
  - OLS is parametric: assumes a particular distribution of the error term (the normal distribution).
  - QR is semi-parametric: no assumption on the parametric distribution of the error term. However, QR is still linear: the parameters are linear functions of the quantiles.
- Robustness:
  - OLS is inefficient when errors are non-normal and suffers heavily outliers.
  - QR is robust to non-normality of errors and outliers. It is particularly accurate at estimating the largest and smallest quantiles.
- Flexibility:
  - OLS is not invariant to monotonic transformations
  - QR is invariant to monotonic transformations (can take the log of quantiles and translate the results back to y).

In the full model I estimate the quantile regressions for 19 quantiles: from the 5th to the 95th in steps of 5. These 19 quantiles are calculated for each observation and for each regressor; the process is repeated for every forecasting horizon and again for every country. As it can be imagined, the computational load is large.

In the following sections of this chapter, I focus on Italy, for reasons of synthesis and exposition[3].

3.2 Univariate regressions

The first results from my analysis are shown in figure 3.1. The figure shows the univariate quantile regressions of the moving average of inflation rate over the forecasting horizons (6 and 12 months) on current inflation rate only, versus the same regression on unemployment rate only.

For the first two figures we can see a positive and linear relationship, as expected. For the last two figures, representing the QR of inflation rate on unemployment rate, it is clear how different the slopes of the OLS and median (Q50) lines are. This shows that the distribution of inflation rate (or its moving average over the forecasting horizon) is not symmetric and there could be outliers.

We can see that the relationship between inflation and unemployment is negative, as expected given that the model is effectively a very simple Phillips curve; however the relationship is weak (the slope is

[3] In my analysis I have produced many results. I needed to select appropriately what to show in this dissertation, hence I take Italy as reference; I show the results for the other countries in chapter 5.

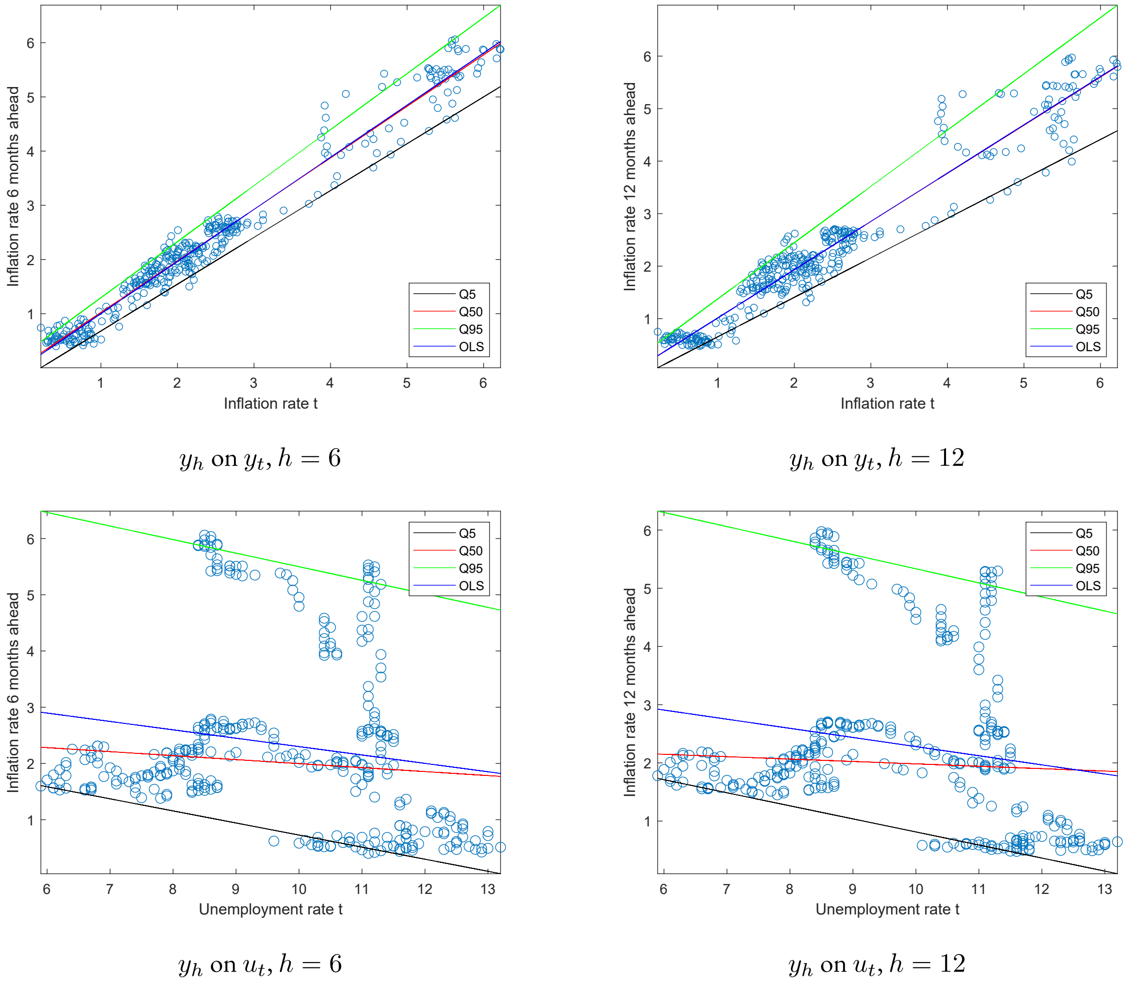

Figure 3.1: Univariate QRs

small in absolute value) and the distribution is characterised by high variability (the distance between the 5th and 95th quantiles is large).

3.3 Multivariate regressions

In Figure 3.2 I compute the multivariate quantile regression of the fully specified model (see Equation 2.1) and show separately the estimated quantile regression coefficients of inflation at time *t* and unemployment at time *t*, for each computed quantile and for the two forecasting horizons.

The red line with marks is the in-sample fit of the regression, that is, the value of the parameters $\beta(\tau)$ for each quantile τ. In addition, I trace the blue and black dashed lines representing the OLS estimate and the median, respectively. The bands, calculated for different confidence intervals, act as confidence bounds for the null hypothesis that the true data-generating process is a general, flexible *linear* model for the inflation and unemployment rates. These bounds are calculated by estimating a simulating VAR with 6 lags, Gaussian innovations and a constant using the full-sample evolution of inflation and unemployment rate. Following this, I bootstrap 1,000 samples to compute the bounds at different confidence level for the OLS relationship. The quantile coefficient estimates that fall outside this confidence bounds indicate that the relation between the dependent variable and the regressor is nonlinear in those quantiles.

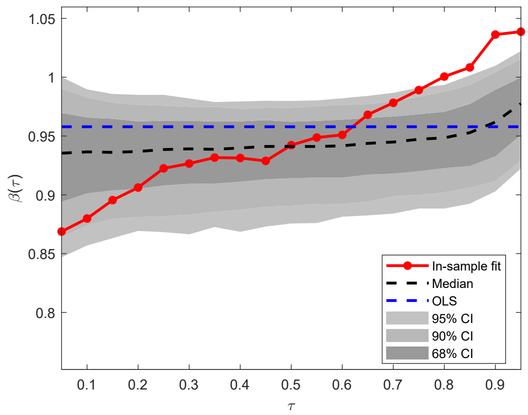

(a) Coefficients of y_t for quantiles τ, h = 6

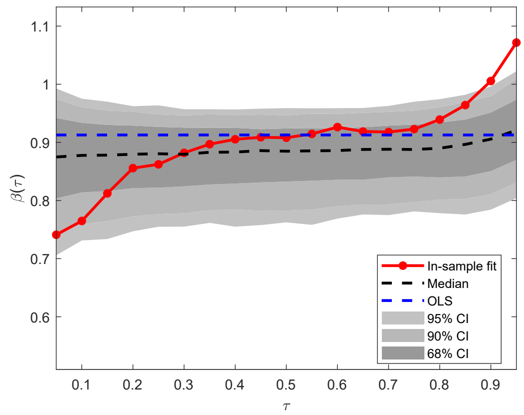

(b) Coefficients of y_t for quantiles τ, h = 12

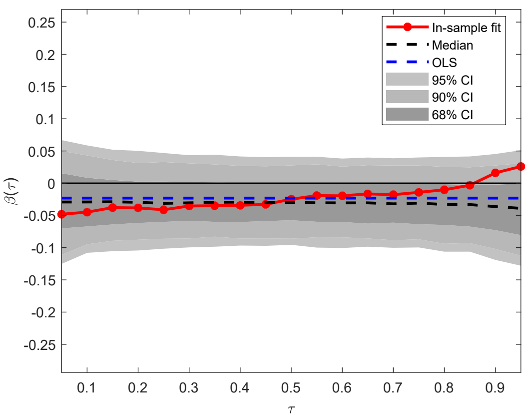

(c) Coefficients of u_t for quantiles τ, h = 6

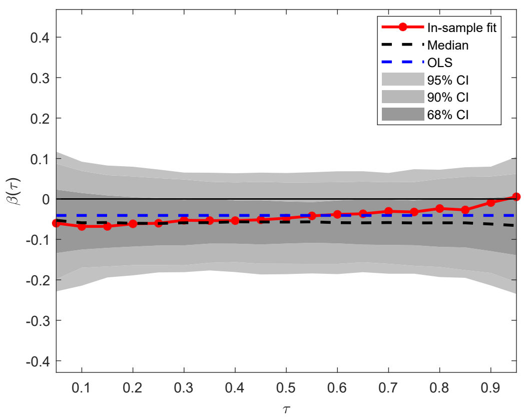

(d) Coefficients of u_t for quantiles τ, h = 12

Figure 3.2: Full QRs

The figure thus shows how the unemployment rate has a linear effect on the MA of the inflation rate over the forecasting horizons, while the present-time inflation rate has a nonlinear effect on its forecasted MA value: linearity is rejected at the 10% significance level for the lower quantiles and is strongly rejected at the higher quantiles (at least 5% significance). The regression slope changes across quantiles for the inflation rate predictor, while for the unemployment rate this is not the case.

3.4 Predicted Distributions

Figure 3.3 shows the time series evolution of the predicted distribution of the MA of the inflation rate over the forecasting horizons, conditionally on the full set of regressors. The continuous black line is the median forecast, while the dashed blue line is the actual value of inflation in that period. The different colours for the bands represent the 5th, 25th, 75th and 95th percentiles.

The variability naturally is higher for the longest forecasting horizon. We can also notice that the forecasts of the inflation rate are more accurate, even with this simple model, compared to the forecasts of other economic variables, such as GDP growth; this is because the inflation rate is a highly persistent process. Nevertheless, the variability has decreased substantially from the 90s to the present time. The figure shows how the tails of the predicted distribution change significantly over time: periods in which the tails were more left-skewed (negative skewness), such as during the GFC, give place to small but positive skewness regimes, such as around 2018. However, this figure is not the best at showing the structure of the predicted distribution; indeed I am only plotting some quantiles. I still do not have enough information to characterise it and because of this I cannot quantify risk rigorously. The next step is presented in the following chapter.

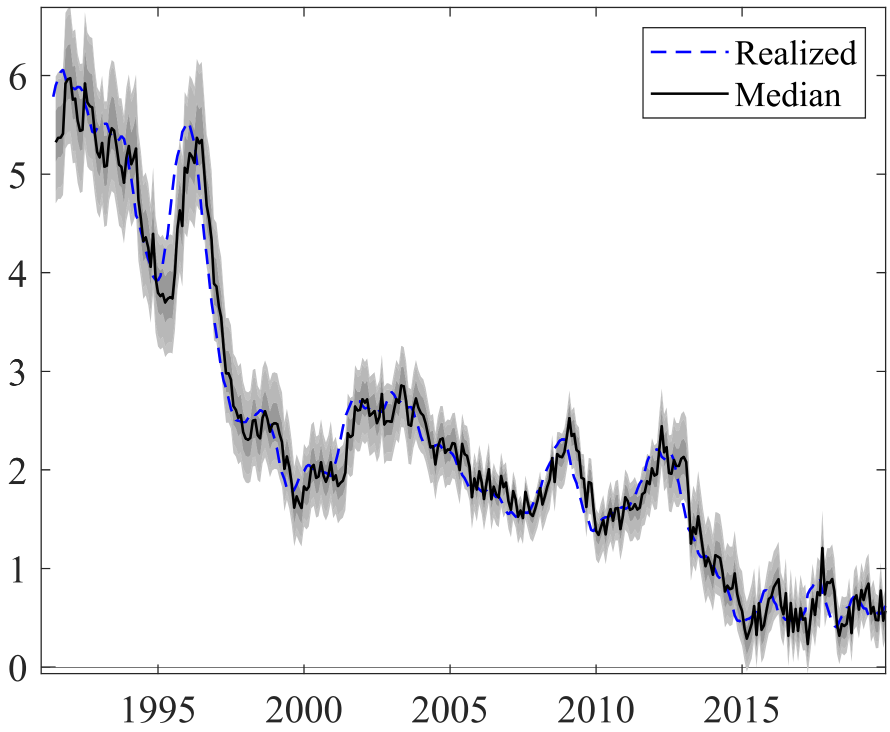

h6

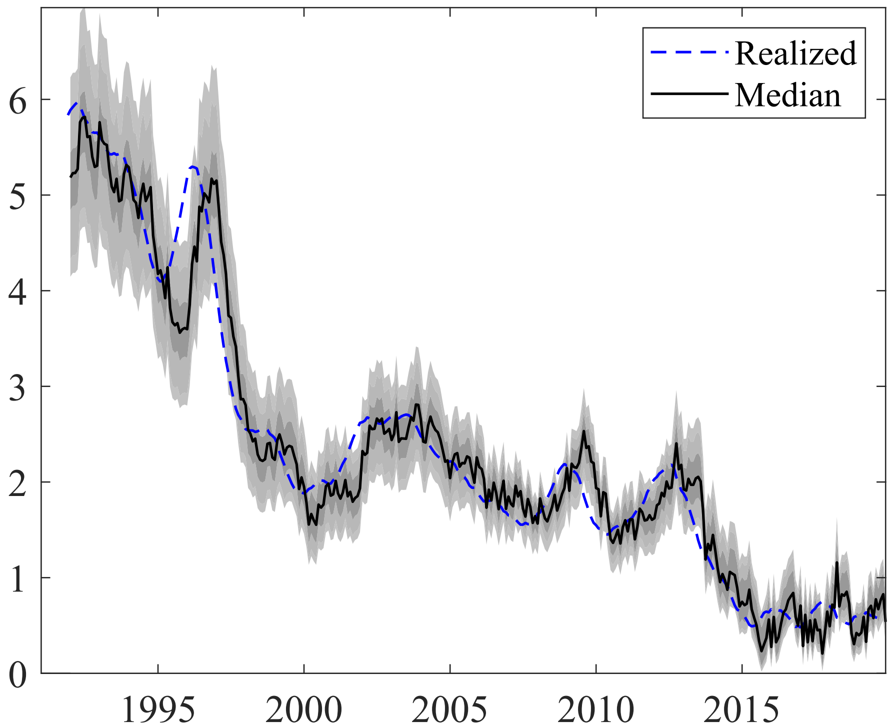

h12

Figure 3.3: Predicted distribution of inflation rate (TS)

4

Constructing the risk measures

At this point I have calculated a series of quantiles (see section 3.1) and produced some preliminary results. Now I proceed with the main focus of this dissertation, that is to study the risk of upside inflation, as well as its conditional distribution.

In this chapter, I explain the matching and fitting procedures of the two-steps quantile regression by Adrian et al. (2019), and present the main findings. I still pertain to the case-study of Italy, before comparing it to the rest of the G4 in chapter 5.

4.1 Matching quantiles and fitting the distribution

The task of estimating the full CQF can be accomplished by different means. One way is to use the spline interpolation to smooth the curve of the quantiles from the previous quantile regressions. However, this does not guarantee that the function obtained this way is going to be monotonic and injective (one-to-one). Thus, the result of this approach would likely not satisfy the requirements of a quantile and distribution functions, in addition to the fact that the function obtained will likely be non-invertible.

Schmidt (2013) proposes a parametric method to estimate multiple conditional quantiles, in which monotonicity and injective function requirements are satisfied by construction[1].

[1] There are other methods, but my aim here is not to make a comprehensive list of them.

My approach is to match the empirical quantiles calculated in chapter 3 to the theoretical ones of a known statistical distribution, so to smooth the conditional quantile function and to be able to recover the information I am interested in. For instance, by inverting the QF I can obtain the CDF, which can be differentiated into the PDF; in addition, I can calculate the moments, etc.

In order to do so, the researcher needs to choose a suitable distribution and optimisation methods. Ideally, the chosen distribution would be the most flexible (many parameters describing the shape of the distribution); however, since the fitting task ultimately reduces to a series of optimisation problems, the researcher has to weigh in computational load, complexity and sensibility to the initial conditions (the initial values of the parameters that have to be decided as input, as well as the bounds for the parameters). All of these elements can affect the accuracy of the fit; thus it may seem as the choice of this method is suboptimal, but as I will show later, the fit is good and the distribution chosen allows to capture well the shape of the CQF without being too burdensome or hard to estimate.

The choice falls on the *Skew-t* distribution by Azzalini and Capitanio (2003). Its PDF, for a random variable y, is:

$$f(y;\mu,\sigma,\alpha,\nu) = \frac{2}{\sigma} t\left(\frac{y-\mu}{\sigma};\nu\right) T\left(\alpha\frac{y-\mu}{\sigma}\sqrt{\frac{\nu+1}{\nu+(\frac{y-\mu}{\sigma})^2}};\nu+1\right) \tag{4.1}$$

Where t(·) and T(·) denote respectively the PDF and CDF of the Student t-distribution. The parameters[2] are location μ, scale σ, shape α, fatness ν. The shape parameter α controls the skewness. The skewed t-distribution is part of a general class of mixed distributions proposed by Azzalini (1985) and further developed by Azzalini and Dalla Valle (1996)

For certain values that they can take on, the parameters assume standard interpretations such as mean, variance, skewness, and degrees of freedom. Here are some notable special cases:

- $\alpha = 0$ yields the traditional Student's t.
- $\alpha = 0$ and $\nu = \infty$ yield the Gaussian with mean μ and standard deviation σ.
- $\nu = \infty$ and $\alpha \neq 0$ yields the skewed normal distribution

[2]The parameters are functions of the regressors. I don't report the dependence in my notation for simplicity.

For all observations, I choose the optimal parameters to minimise the squared distance between the estimated quantile function:

$$Q_{y_h|\mathbf{x_t}}(\tau|\mathbf{x_t}) \tag{4.2}$$

and the quantile function of the skewed-t distribution:

$$F^{-1}(\tau; \mu, \sigma, \alpha, \nu) \tag{4.3}$$

The quantiles matched are the 5th, 25th, 75th and 95th[3].

The following is the optimisation problem solved[4]:

$$\{\hat{\mu}_{t+h}, \hat{\sigma}_{t+h}, \hat{\alpha}_{t+h}, \hat{\nu}_{t+h}\} = \underset{\mu,\sigma,\alpha,\nu}{\operatorname{argmin}} \sum_{\tau} \left(\hat{Q}_{y_h|\mathbf{x_t}}(\tau|\mathbf{x_t}) - F^{-1}(\tau; \mu, \sigma, \alpha, \nu)\right)^2 \tag{4.4}$$

To cite Adrian et al. (2019), the first proposers of this method: "This can be viewed as an exactly identified nonlinear cross-sectional regression of the predicted quantiles on the quantiles of the skewed t-distribution".

4.2 Predicted densities and moments

Figure 4.1 shows the fully conditional predicted densities across time. This figure shows clearly how the densities change substantially across time. The evolution of the second, third and fourth moments is particularly interesting for risk analysis purposes.

Below, in Figure 4.2 I report the first and second moments, as well as the standardised third and fourth central moments, across time. They are, respectively, the mean, variance, skewness and kurtosis. The moments help me to characterise the predicted distributions.

Mean and variance obviously provide measures of location and dispersion associated to the estimated probability density function.

Skewness is a measure of asymmetry about the mean. Negative skewness implies left-tailed density;

[3]Alternatively I could have matched all the quantiles calculated in chapter 3 to calculate the optimal parameters, allowing them to be over-identified. Here I follow a more parsimonious and exactly-identified approach.

[4]The algorithm used solves a nonlinear OLS.

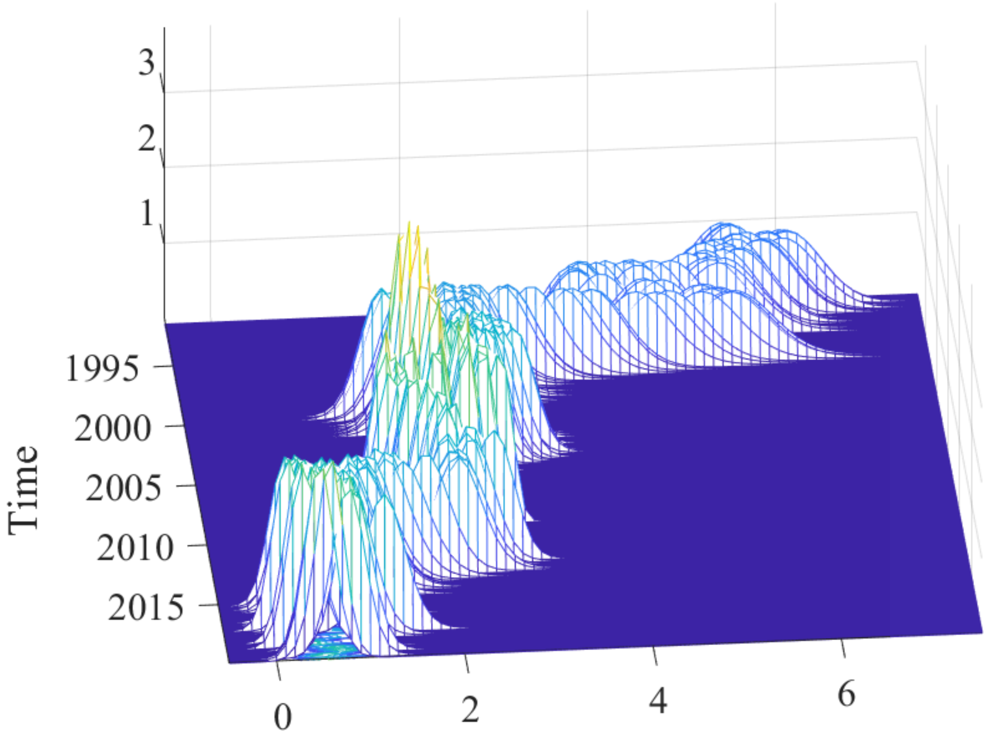

h6

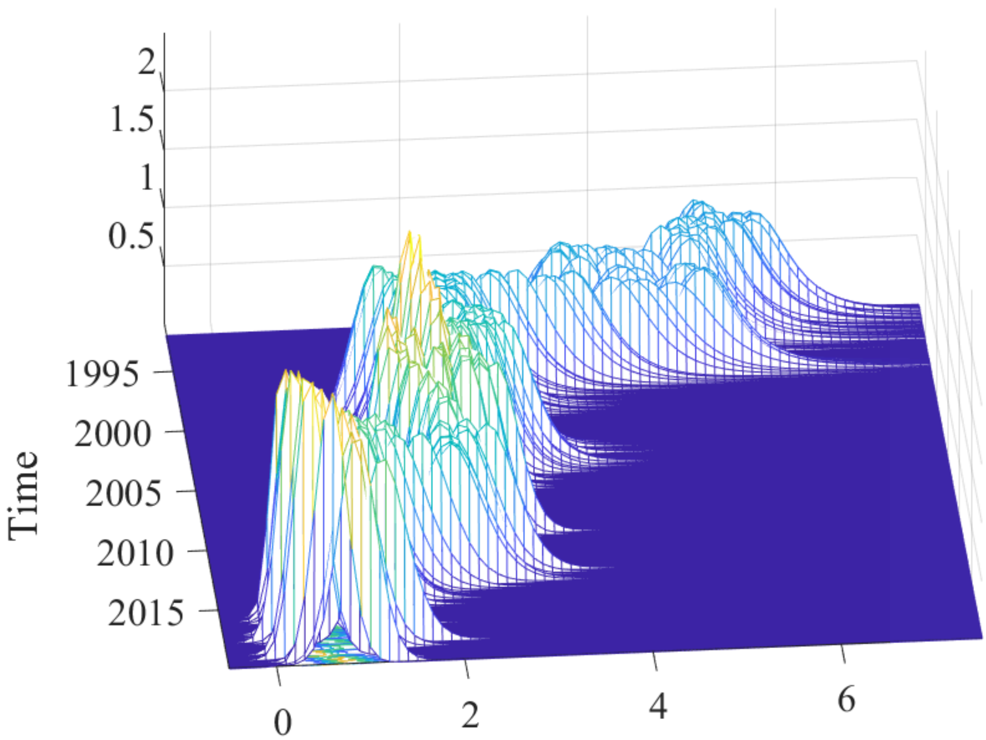

h12

Figure 4.1: Predicted distributions of inflation rate

conversely, positive skewness implies right-tailed density[5]. For instance, positive skewness means that the density is affected by some relatively rare, high-value occurrences in the data — in other words, positive outliers. Thus, right-skewness implies that we are at risk of unexpectedly high realisations of the random variable, and vice versa for the left-skewness

In order to quantify the amount of outliers, the kurtosis is calculated. Kurtosis is a measure of the "tailedness" of the probability distribution. It is important to remark that this number is related to the tails of the distribution, not its peak; hence, the characterisation of kurtosis as "peakedness" is erroneous. Higher kurtosis corresponds to greater extremity of deviations (or outliers); thus, we should refrain from attempting to infer visually the shape characteristics of the PDF simply by reading the values of the individual moments.

From Figure 4.2 we can see that, in recent years, the mean of the inflation rate goes down from the high levels of the 1990s to around target, since the adoption of the euro. Variance has decreased substantially. The skewness revolves around zero, with strong volatility in coincidence of the Global Financial Crisis. From around 2013 onwards, the predicted densities are slightly positively skewed, meaning that there is a potential for unexpectedly high realisations of inflation. When it comes to kurtosis, the predicted distributions are leptokurtic (positive excess kurtosis), meaning that the tails are "fatter": there is more probability mass in the tails compared to the Normal distribution, hence there is a higher risk of outliers. The kurtosis is very large in the case of predicted densities before 1998 in Italy, for the 12-months forecasting horizon. Moreover, in the 6-months forecasting horizon, we can see how the kurtosis becomes very volatile around the Global Financial Crisis (something that we see also in the skewness for the same horizon).

It is worthwhile to highlight how the moments change from a 6-months to a 12-months forecasting horizon. A similar level of volatility is found in Adrian et al. (2019)[6].

[5]Assuming uni-modality.
[6]Online Appendix.

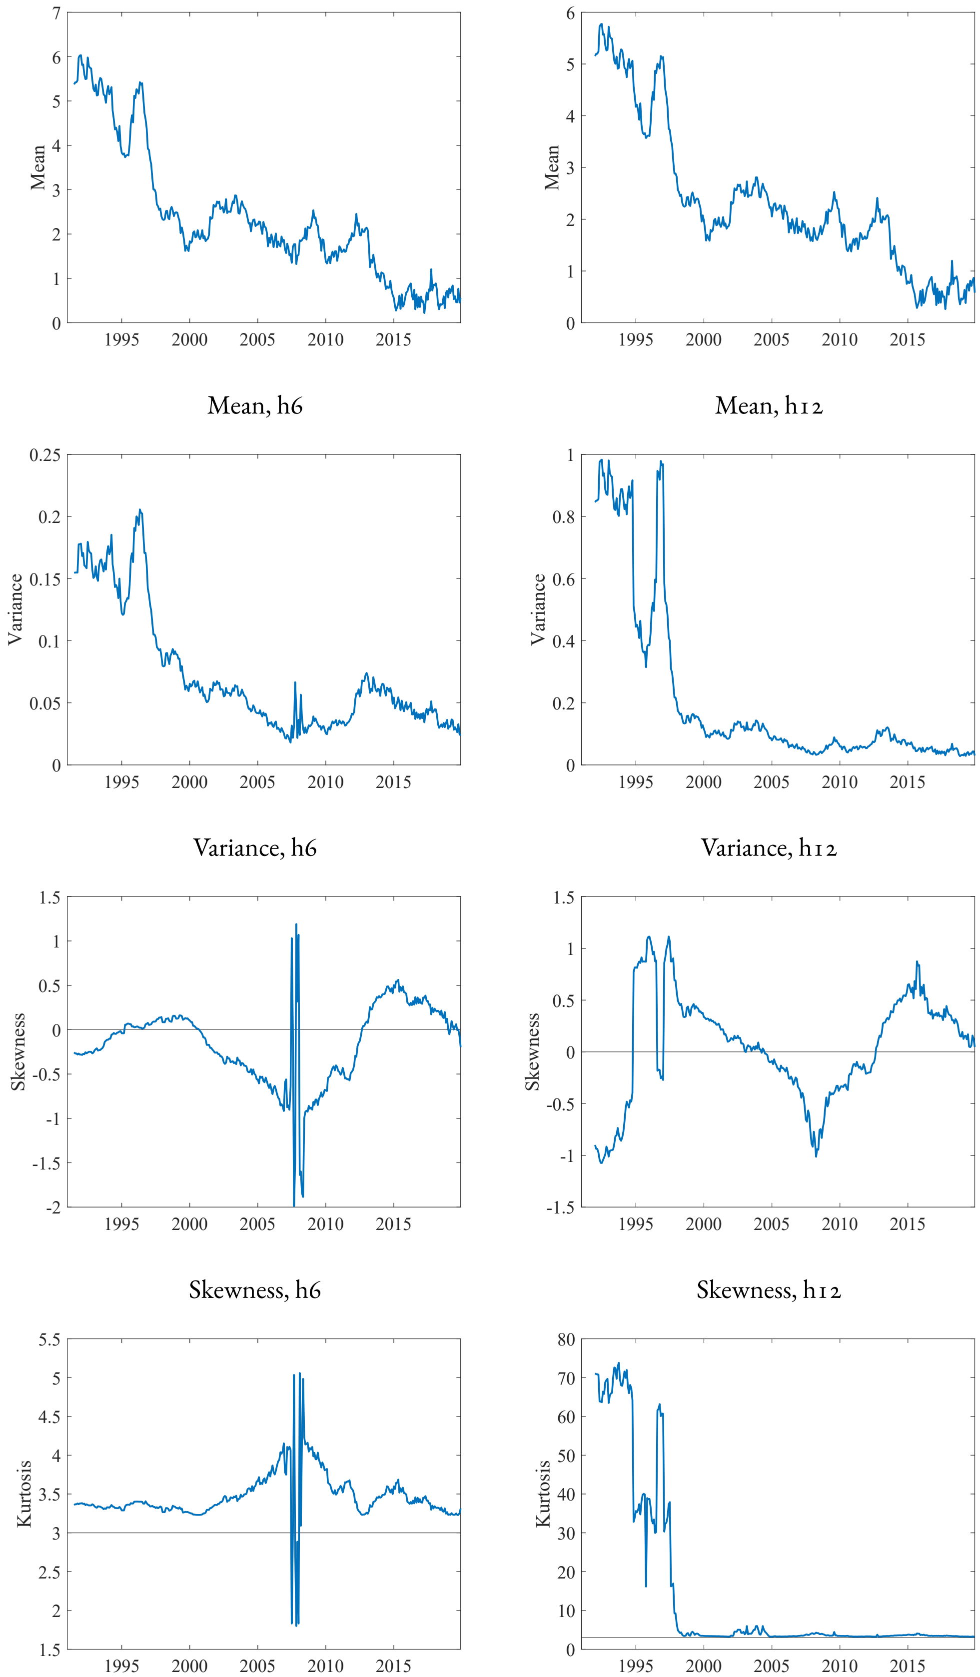

Mean, h6 | Mean, h12
Variance, h6 | Variance, h12
Skewness, h6 | Skewness, h12
Kurtosis, h6 | Kurtosis, h12

Figure 4.2: Moments of the predicted densities

4.3 Assessing the fitting process

Figure 4.3 shows the fitted conditional quantile functions

$$\hat{Q}_{y_h|\mathbf{x_t}}(\tau|\mathbf{x_t}) \tag{4.5}$$

superimposed on the empirical quantiles; this allows to have a visual assessment of the fitting process.

I pick three observations and corresponding CQFs: January 2013, January 2015 and December 2018. The first month is characterised by relatively stable inflation around target, the second month represents a period of very low inflation (deflation for the headline measure), while the third is again "normal times" towards the end of the sample and one year before COVID-19.

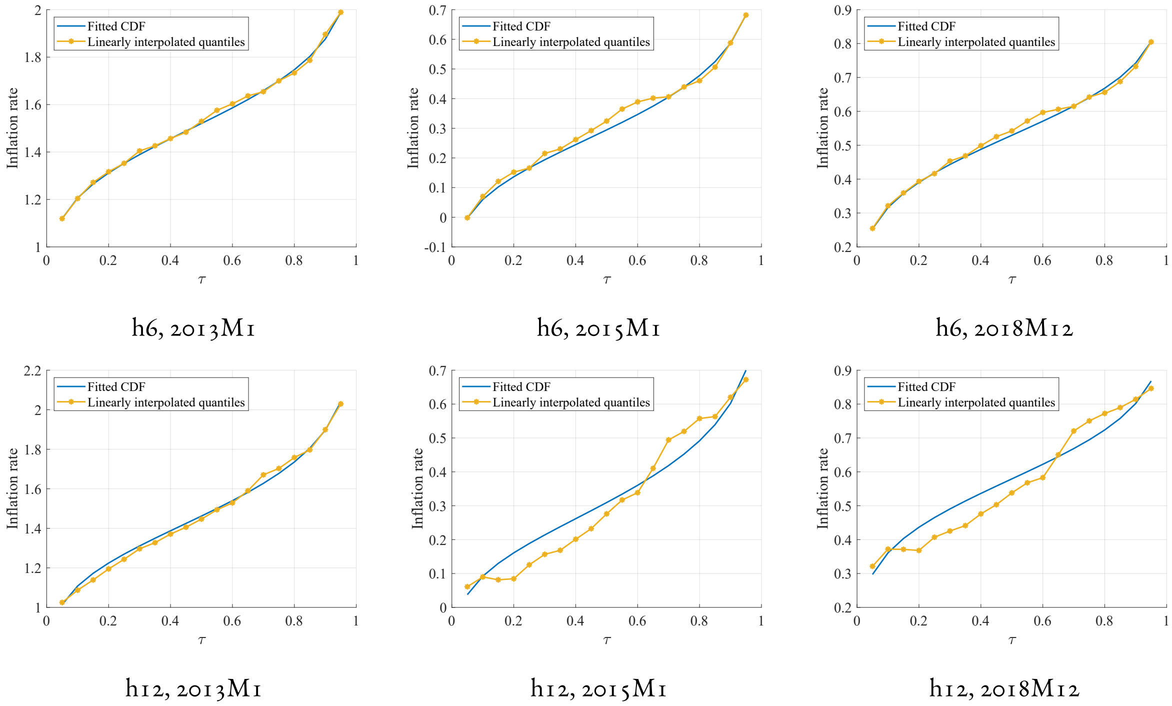

Figure 4.3: CQFs

The fitting process is good overall, especially for normal times (2015, 2018). However, we can see that for periods of low inflation (2015), the fit is considerably worse. It is also noticeable how the empirical quantiles (yellow marks) do not follow the shape of a proper quantile function (the linear interpolation is not monotonic and injective); hence the issue might not be in the choice of the distribution itself or its computation, but rather with the estimation of these quantiles in the quantile regression. The QR struggles in times of very low, zero or negative inflation.

A modification that could help to solve this issue might be the addition of an appropriate indicator accounting for financial conditions. López-Salido and Loria (2024) show that financial conditions capture well and significantly the lower-regime dynamics of inflation. I have estimated such a model and presented the full analysis for the new specification in Appendix A.

A similar exercise is repeated for Figure 4.4, where I show the estimated probability density functions of the full model, compared to the ones of a model without unemployment as regressor.

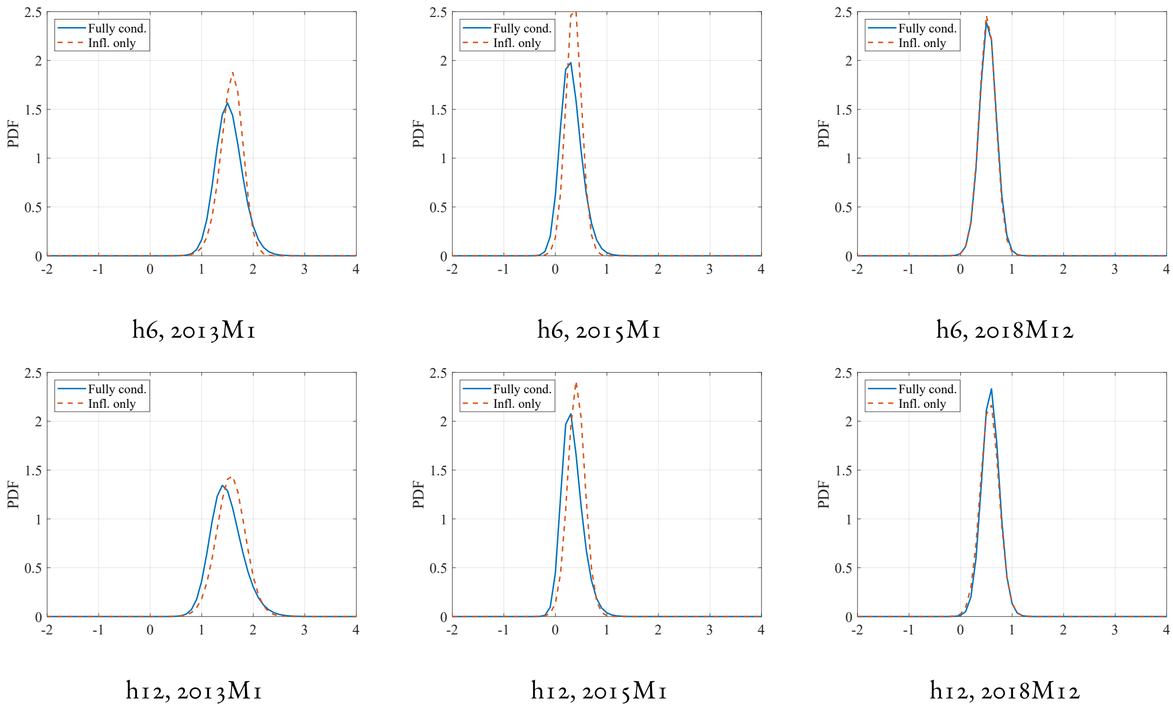

Figure 4.4: PDFs

This choice allows to see the effect of controlling for unemployment in the distribution of inflation: such an effect is not large on the median or the mode, but it affects also the other moments of the distribution mainly. The fully conditional densities are somewhat less "spiky".

4.4 The risk measures

The mean of the estimated density function is a modal forecast of the MA of inflation rate over the forecasting horizons, conditionally on the regressors. Now, I shall introduce the risk measures that allow me to capture the vulnerability of the inflation rate to extreme realisations, as well as to quantify what is the expected value of inflation for such unlikely but risky outcomes.

I am interested in upside risk of inflation, hence I focus on the following measures: expected longrise and upside relative entropy [7].

Expected Longrise In finance, the Conditional Value at Risk (CVaR) is the expected value of losses beyond the VaR probability cut-off. It considers the average of losses that occur beyond the worst-case threshold defined by VaR. It is also called expected shortfall. These concepts can be applied in the study of other objects, such as economic indicators. Let us go back to inflation: since I want to study upside risk, I calculate the CVaR in the opposite direction, obtaining the expected longrise:

$$LR_h = \frac{1}{\pi} \int_{1-\pi}^{1} \hat{F}^{-1}_{y_h|\mathbf{x_t}}(\tau|\mathbf{x_t})\, d\tau \tag{4.6}$$

The expected longrise is the expected value of the conditional distribution of inflation beyond the $1 - \pi = 90\%$ probability. It captures the tail behaviour of the conditional distribution in absolute terms. Figure 4.5 Reports the estimated expected longrise (as well as shortfall, shown for reference) over time. We can see how the series resembles the one for inflation.

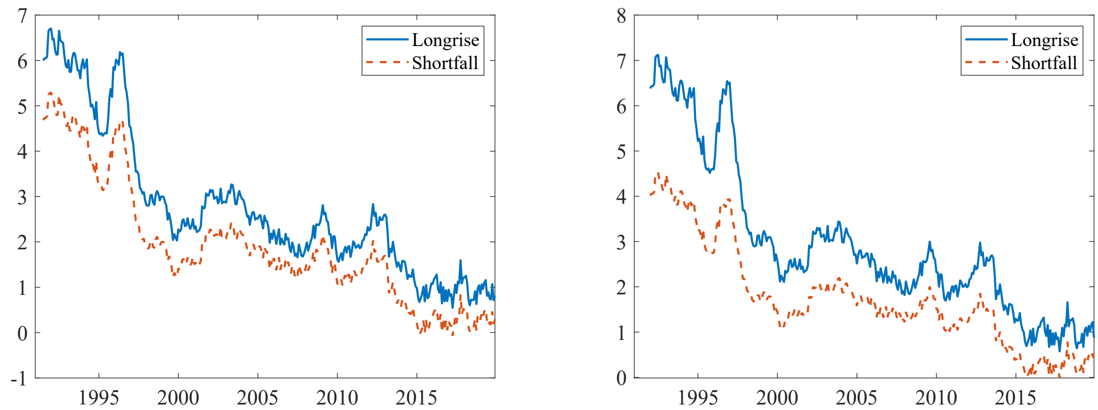

h6 h12

Figure 4.5: Expected longrise and shortfall

[7] Adrian et al. (2019) instead focus on downside vulnerability of GDP growth, hence studying downside relative entropy and expected shortfall.

Upside relative entropy Define the unconditional density[8] as:

$$\hat{g}_{y_h} \tag{4.7}$$

and the estimated (conditional) skew-t density as:

$$\hat{f}_{y_h|\mathbf{x_t}}(y \mid \mathbf{x_t}) = f(y; \hat{\mu}_h, \hat{\sigma}_h, \hat{\alpha}_h, \hat{\nu}_h) \tag{4.8}$$

Then, the Upside Relative Entropy of $\hat{g}_{y_h}$, relative to $\hat{f}_{y_h|\mathbf{x_t}}(y \mid \mathbf{x_t})$, is:

$$\mathcal{L}_t^U \left(\hat{f}_{y_h|\mathbf{x_t}}; \hat{g}_{y_h} \right) = -\int_{\hat{F}^{-1}_{y_h|\mathbf{x_t}}(0.5|\mathbf{x_t})}^{\infty} \left(\log \hat{g}_{y_h}(y) - \log \hat{f}_{y_h|\mathbf{x_t}}(y|\mathbf{x_t}) \right) \hat{f}_{y_h|\mathbf{x_t}}(y|\mathbf{x_t})\, dy \tag{4.9}$$

Entropy is a measure of the uncertainty or randomness of a dataset. It quantifies the expected value of the information content of the distribution, where "information content" is in the Shannon information sense [9].

Relative Entropy is the statistical distance between two distributions, over the full sample space. In this case, I take the two distributions to be the unconditional PDF and the conditional PDF, but integrating above the median. This yields the formula in Equation 4.9 [10].

Upside relative entropy measures, intuitively, the divergence between conditional and unconditional distributions, from the median above. This divergence is represented by extra probability mass in the right side of the conditional density compared to the unconditional one. A higher upside relative entropy means that the conditional density assigns positive probability to more extreme right tail outcomes than the unconditional density. This implies greater vulnerability, or risk, around the modal forecast, compared to the unconditional distribution. In essence, the extra probability mass in the right-tail, as represented by upside relative entropy, is "material" risk that is not captured in the unconditional distribution; hence, this measure could also give an idea of how important the set of regressors is in estimating risk.

[8] Calculated by first estimating a quantile regression of the moving average of inflation over the forecasting horizon on the constant alone, then by fitting the quantiles obtained. It is time-invariant.

[9] Entropy and the concepts related to it are topics that deserves at least an attempt at a proper treatment. However, here I want to focus on the more straightforward and useful interpretations in order to present the results effectively. That is why I have an appendix going over the more technical aspects (Appendix B).

[10] Equation 4.9 is, formally, the *differential conditional relative entropy* from the median above.

Figure 4.6 Shows the evolution of upside relative entropy over time (I report downside relative entropy too, for reference).

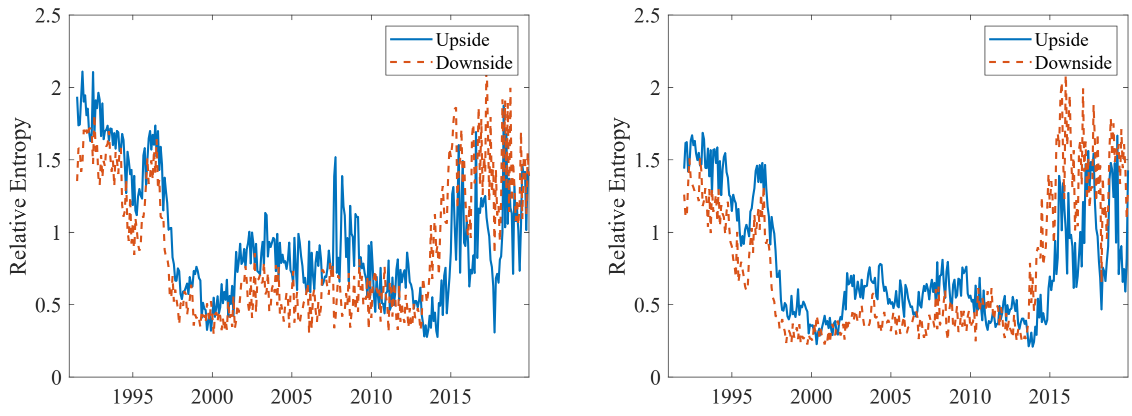

h6 h12

Figure 4.6: Upside and Downside relative entropies

It is striking how the evolution of the dynamics of the inflation rate (or the conditional distribution of its MA forecast) over the sample is such that, in recent years, inflation risks have decreased substantially in absolute terms, as shown by expected longrise and shortfall, while at the same time there has been a considerable increase in the upside and downside relative entropy. The simple Phillips curve relation has become more at risk of unforeseen extreme realisations both to the upside and downside, roughly from 2014. Moreover, URE was high but not volatile before 1997, while from 2014 URE started to increase as well as to become very volatile.

Upside relative entropy and expected longrise share an interesting connection with skewness: if both the conditional and unconditional distributions are positively skewed, then the upside relative entropy is low while the expected longrise is high.

4.5 Tests

I conclude this part by assessing the out-of-sample performance of the model, following the test procedures by Adrian et al. (2019).

In the following paragraphs, the out-of-sample expanding window forecasts are calculated as follows: first, using the subsample from January 1991 to July 2010 I estimate the predictive distribution for January 2011, 6 months later, and July 2011, 12 months later. Afterwards, the procedure is repeated expanding the sample one month at a time, until the sample ends (December 2019). The estimation steps in the previous sections are calculated for each iteration.

Remark Part of the poorer performance compared to Adrian et al. (2019) (notwithstanding the different economic indicator forecasted) is due to the fact that my sample is much shorter (1973 vs 1991). Furthermore, I want to underline the fact that in my dissertation the aim is not to make a good forecasting model to predict accurately inflation. The focus, up to here, has been the to study the dynamics of inflation rate and its relationship with unemployment, in terms of vulnerability. Had I wanted to forecast inflation accurately, I would have used a different methodology and specification of the Phillips curve. Still, it can be relatively useful to evaluate the out-of-sample forecast performance of this simple Phillips curve estimated by two-steps quantile regression.

Robustness Figure 4.7 compares, in the first row, the out-of-sample and in-sample quantiles for the forecasted inflation rate over 6 and 12 months. The gray dotted lines are the in-sample estimated 5th, 50th and 95th quantiles, while the blue solid lines are the out-of-sample ones. In the second row I show the out-of-sample upside relative entropy for the two forecasting horizons, in blue, compared to the in-sample ones, in dotted black.

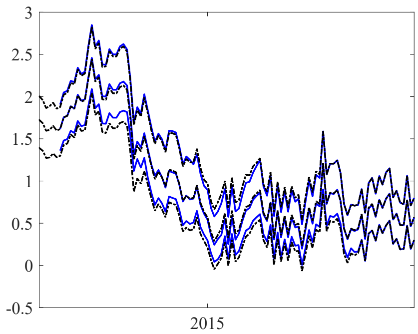

Quantiles, h6

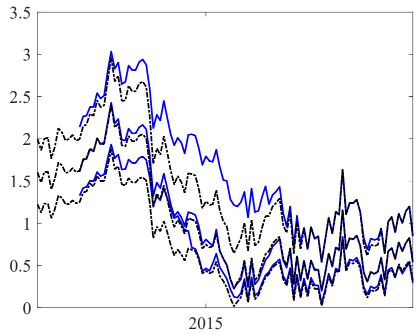

Quantiles, h12

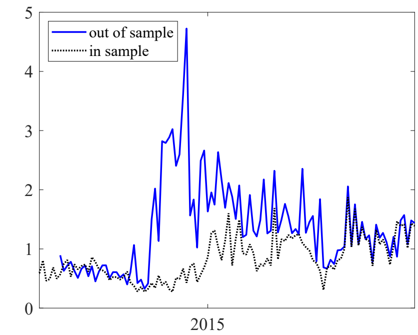

Upside R.E., h6

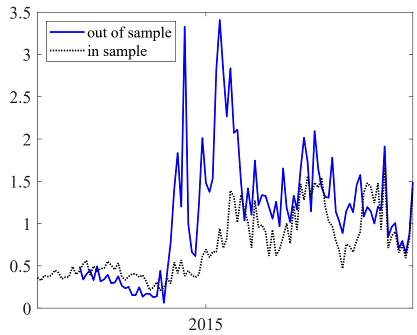

Upside R.E., h12

Figure 4.7: Out-of-sample prediction of quantiles and Upside relative entropy

We can see that the quantiles are forecasted accurately in the 6 months horizon; less so in the 12 months horizon. Real-time prediction of upside relative entropy is not very accurate in the first half of the sample, in fact it is overestimated. This can be attributed to the short window sample, but the theme of the out-of-sample performance being poor is similar for the quantiles around 2015. It also resonated with the poorer fit of the CQF around the same period. The simple Phillips curve relation clearly does not work well when the conditions of periods such as 2014-15 are in place.

Reliability and accuracy I compute the predictive score, that measures the accuracy of the density forecast. More precisely, I evaluate the model's predictive distribution at the realised value of the series. A higher predictive score implies a more accurate prediction. Figure 4.8 shows these measures, calculated for both forecasting horizons and for both the fully specified model and the model without unemployment rate.

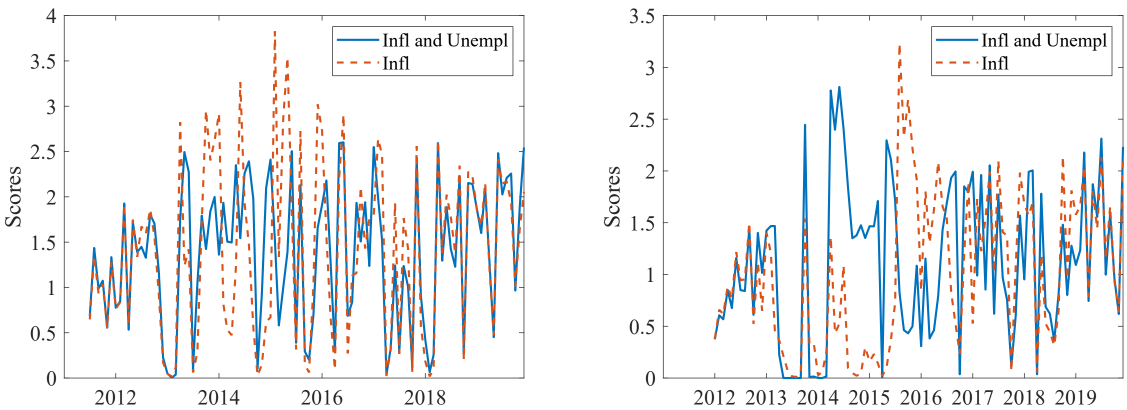

h6 h12

Figure 4.8: Predictive scores

The predictive scores of the fully specified model and the reduced model are comparable in the 6-months horizon, while the full model is better in the year-long horizon. This is further confirmation that the Phillips curve, in its simplest form, is not very powerful in predicting inflation.

Calibration To conclude the testing exercise, I evaluate the calibration of the predictive distribution. I do this by computing the empirical cumulative distribution of the probability integral transforms (PITs). This allows to see the percentage of observations that are below any given quantile. The closer the empirical cumulative distribution of the PITs is to the 45-degree line, the better the calibration. The confidence bands around the 45-degree line, that account for sample uncertainty, are calculated using the critical values obtained following the procedure by (Rossi and Sekhposyan, 2019).

The calibration tests show that the fully specified model is better calibrated than the one without unemployment, while also performing better than, for instance, the GDP growth and NFCI model

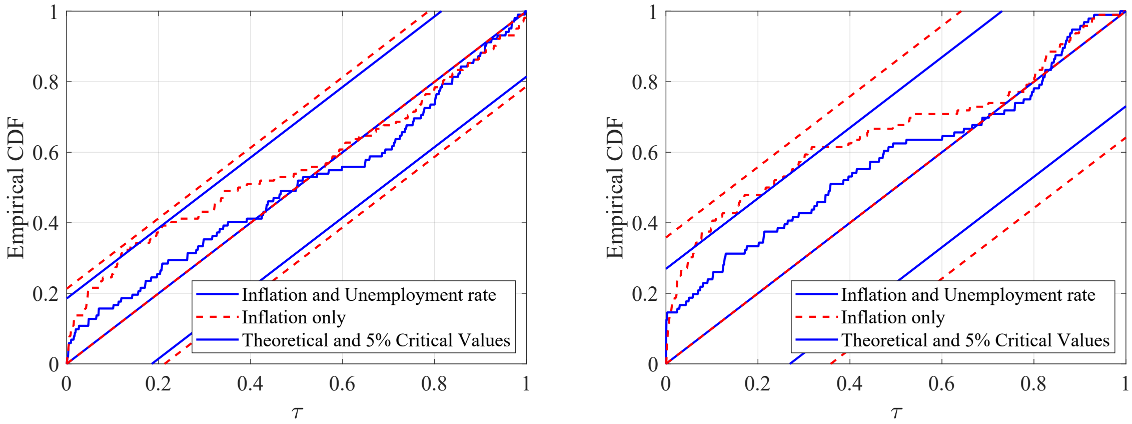

h6 h12

Figure 4.9: Empirical CDF of the PITs

by Adrian et al. (2019). This means that, overall, the quantiles are estimated quite well by the quantile regression; hence the inference made on the previous results gains further credibility.

5

Comparison across countries

In the previous chapter I fitted the skew-t distribution, calculated the predicted distributions and sample moments across time and developed the main inflation risk measures — expected longrise and upside relative entropy. In this chapter, I compare some important results obtained in chapter 4 for Italy with the analogous ones for Germany, France, Spain.

5.1 Two-steps QR comparison

Figure 5.1 and Figure 5.2 compare the estimates of the coefficients of inflation and unemployment for the multivariate quantile regression, akin to chapter 3 for both horizons.

We can see that the results are similar, for a given country, across horizons. The nonlinear effect of current inflation on the forecasted MA of inflation in the upper quantiles is present for Italy, Germany and, more weakly, for Spain; for France we cannot reject linearity in the upper quantiles. As for the lower quantiles, the effect of inflation on its forecasted average value is negative and nonlinear for Italy, Germany and France, with a significance of 5% or smaller; for Spain I find also a nonlinearity but with a positive effect of inflation on the lower quantiles of the forecast. Finally, the unemployment rate's effect on the forecasted inflation rate is linear in all countries.

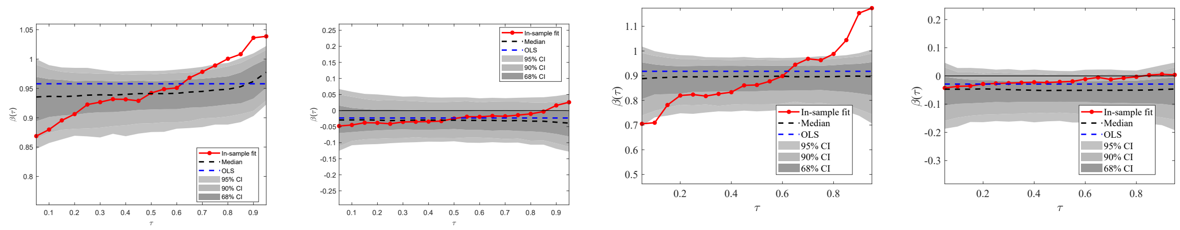

Coefficients of y_t, Italy, h6 Coefficients of u_t, Italy, h6 Coefficients of y_t, Germany, h6 Coefficients of u_t, Germany, h6

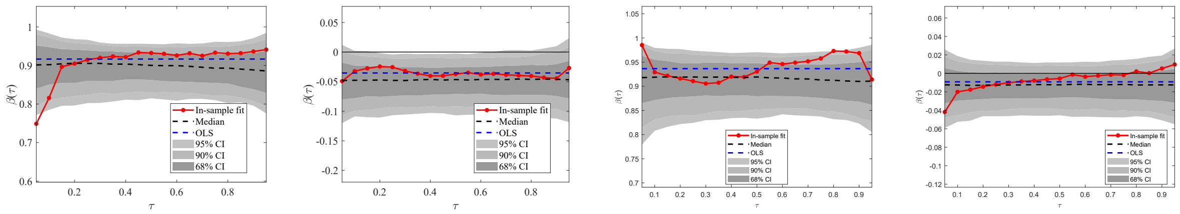

Coefficients of y_t, France, h6 Coefficients of u_t, France, h6 Coefficients of y_t, Spain, h6 Coefficients of u_t, Spain, h6

Figure 5.1: Coefficients from multivariate QR, comparison, h6

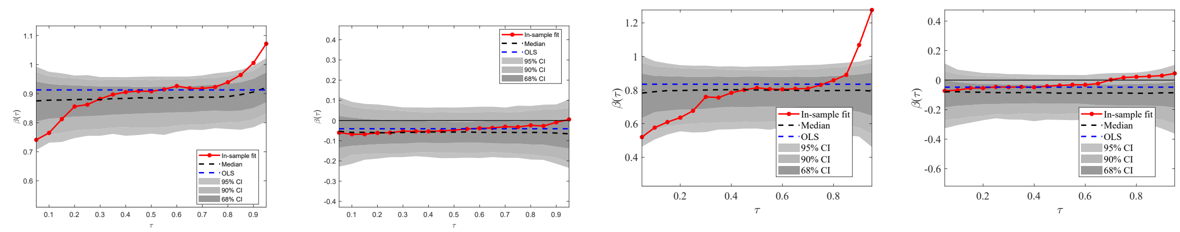

Coefficients of y_t, Italy, h12 Coefficients of u_t, Italy, h12 Coefficients of y_t, Germany, h12 Coefficients of u_t, Germany, h12

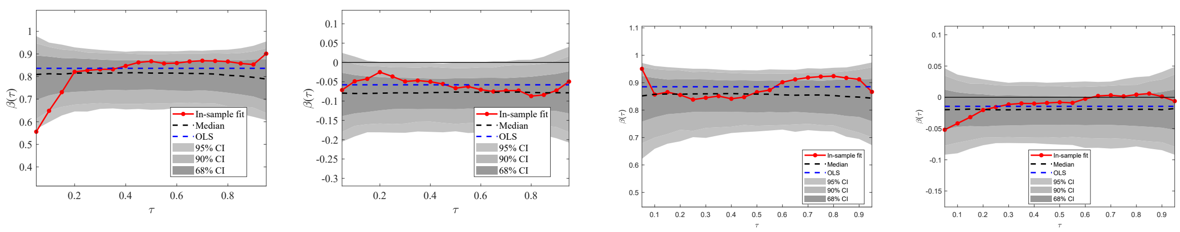

Coefficients of y_t, France, h12 Coefficients of u_t, France, h12 Coefficients of y_t, Spain, h12 Coefficients of u_t, Spain, h12

Figure 5.2: Coefficients from multivariate QR, comparison, h12

Figure 5.3 and Figure 5.4 display the time series of (selected) predicted quantiles, for the conditional forecast of the MA of inflation rate for both horizons and all countries.

The quantiles show a degree of heterogeneity in the skewness of the distributions across countries; for instance, Spain was in deflation risk in the last 5 years of the series while the other countries were not. France manifests left-skewness in the early part of the series, while Germany, in the same part, is right-skewed.

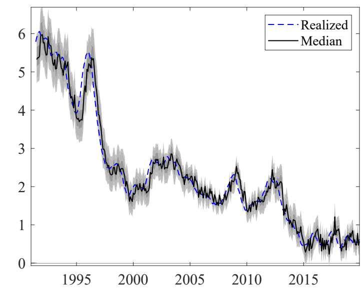

Italy, h6

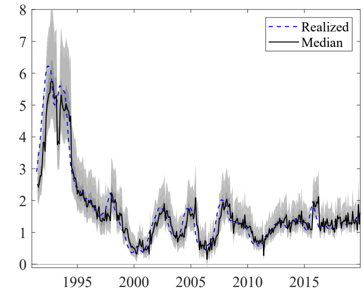

Germany, h6

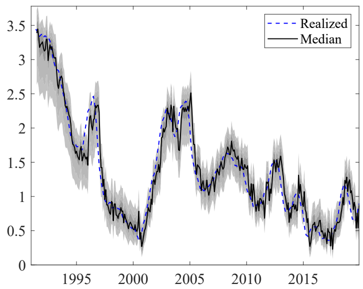

France, h6

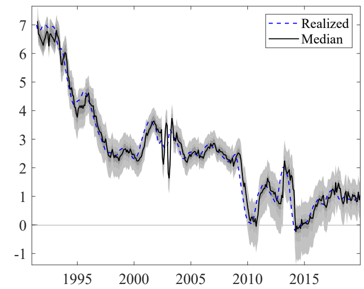

Spain, h6

Figure 5.3: Predicted distribution of inflation rate (TS), comparison, h6

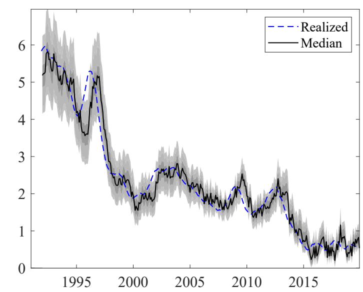

Italy, h12

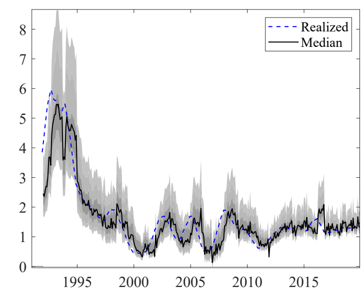

Germany, h12

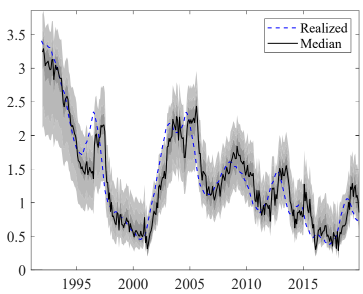

France, h12

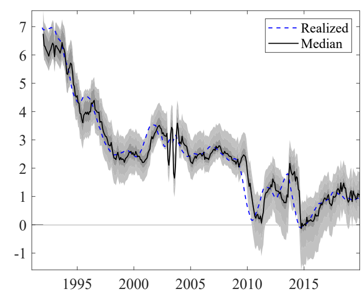

Spain, h12

Figure 5.4: Predicted distribution of inflation rate (TS), comparison, h12

To have a better idea of the shape and moments of the distribution of the forecasted average inflation rate, I show below the predicted densities across time of the G4 countries (after the fitting).

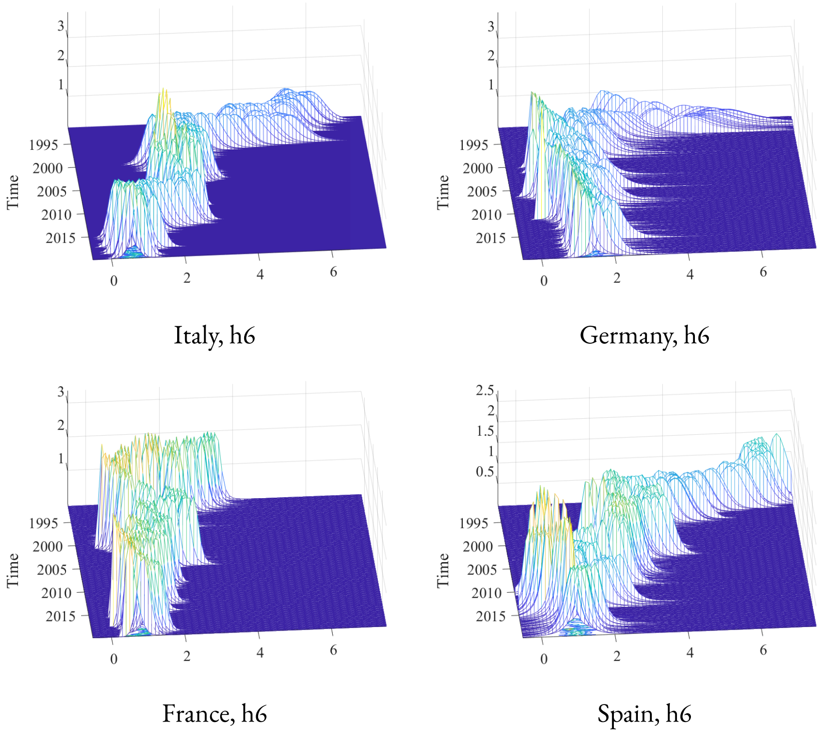

Figure 5.5: Predicted densities, G4 comparison, h6

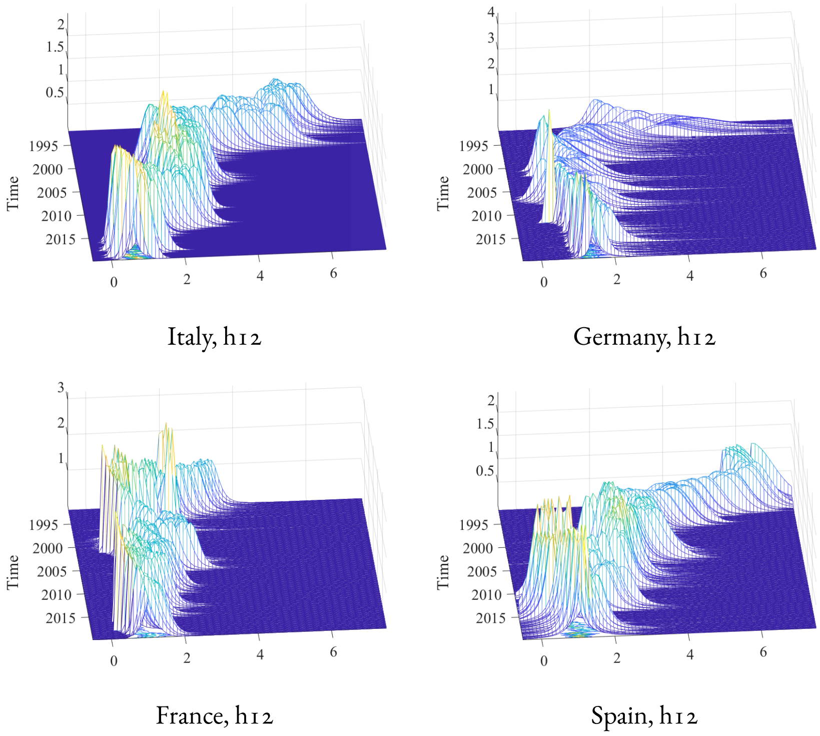

Figure 5.6: Predicted densities, G4 comparison, h12

The densities change dramatically from a country to another. The two extremes are Germany and France: the former presents the largest amount of right-skewness, especially at the start of the series, where the densities are characterised by a very large amount of variance. In all countries, variance starts to decrease substantially going towards the end of the sample; similarly, skewness seems to become more contained. Below are the moments of the fitted distributions across countries.

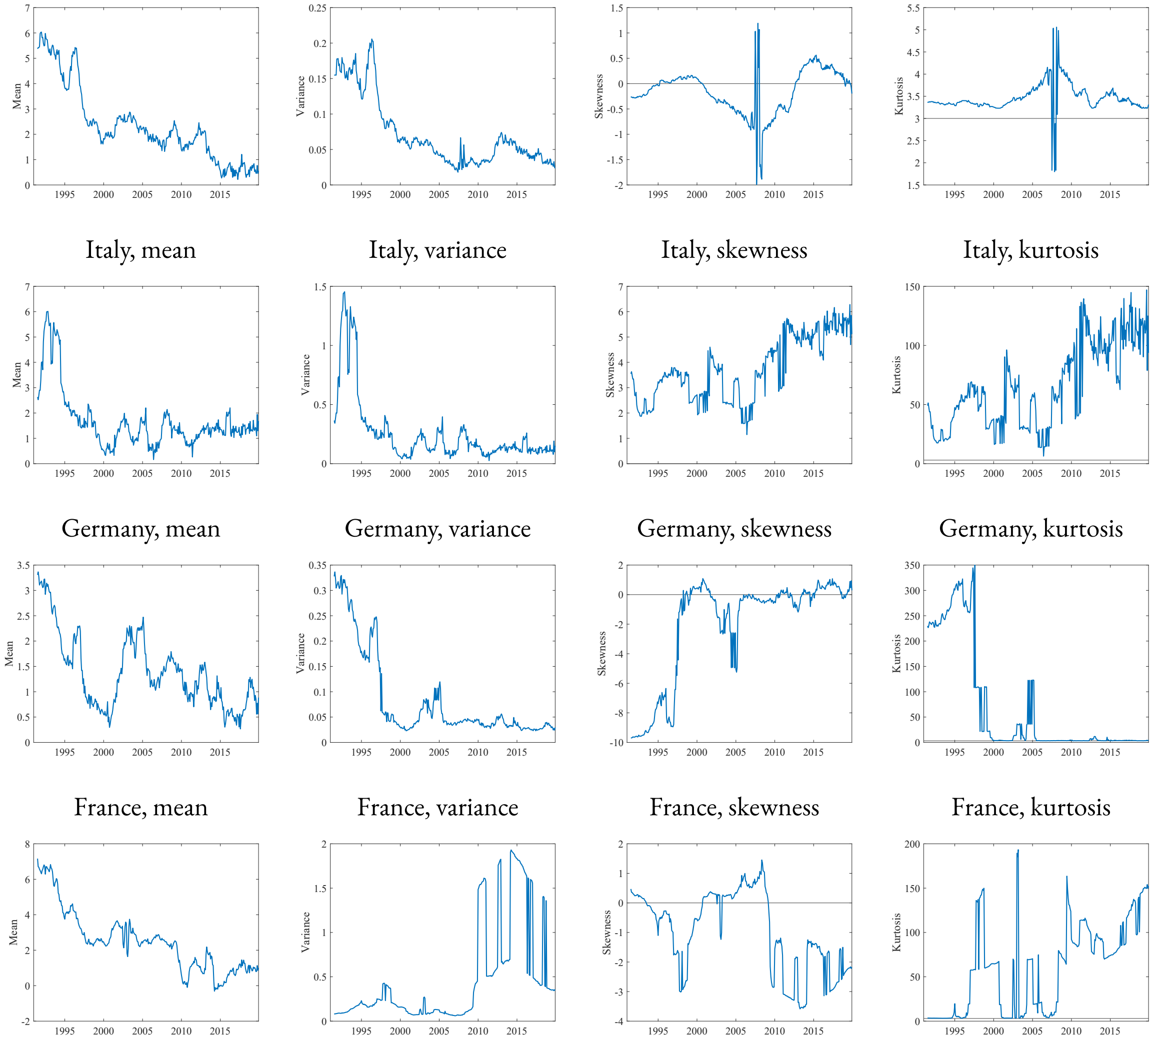

Italy, mean — Italy, variance — Italy, skewness — Italy, kurtosis

Germany, mean — Germany, variance — Germany, skewness — Germany, kurtosis

France, mean — France, variance — France, skewness — France, kurtosis

Spain, mean — Spain, variance — Spain, skewness — Spain, kurtosis

Figure 5.7: Moments, all countries, h6

All countries' distributions, except Spain, are characterised by low variance; Spain suffers from inflation risk even in the traditional "standard deviation" sense. Italy and France's predicted densities manifest barely any excess skewness and in the more recent years of the sample they are low in kurtosis, thus they are not significantly at risk of surprisingly high or low values. The opposite is the case for Germany and Spain. Keeping a focus on the more recent years of the sample, Spain's predicted density of the forecasted average inflation rate is left-skewed, thus there are significant outliers in the low values of inflation. Germany is instead right-skewed. Both countries present high levels of kurtosis, indicating a substantial degree of outliers.

Taking stock 1 Based on these measures, we can say that the picture does not look exceedingly promising in terms of the European Central Bank being able to continue the process of convergence of the inflation rates in the euro area to the target, at least when it comes to evaluating the chances of that happening based on inflation at risk and on the predicted Phillips curve relation. The distributions are very heterogeneous and some countries seem to be substantially more at risk than others from this first analysis. The degree of outliers is high in some countries, low in others.

5.2 Risk measures' comparison

5.2.1 Correlation of expected longrise

Figure 5.8 shows the expected longrise series for all countries.

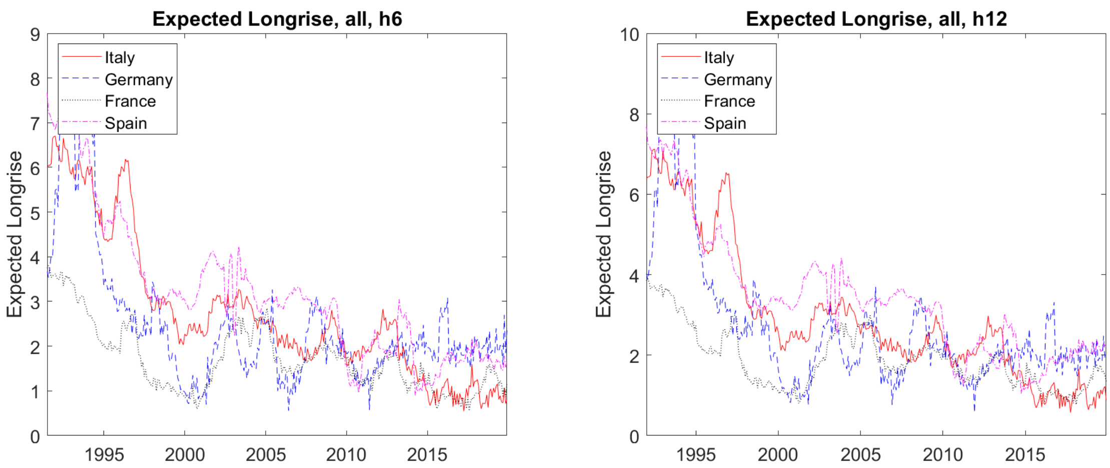

h6 h12

Figure 5.8: Expected Longrise series, comparison

Table 5.1 reports the linear (Pearson) correlation among these countries' inflation risks, as measured in absolute terms by the expected longrise. This is done for both horizons.

	Ita	Ger	Fra	Spa
Ita	1	0.68	0.78	0.88
Ger	0.68	1	0.63	0.61
Fra	0.78	0.63	1	0.72
Spa	0.88	0.61	0.72	1

h6

	Ita	Ger	Fra	Spa
Ita	1	0.61	0.70	0.88
Ger	0.61	1	0.53	0.52
Fra	0.70	0.53	1	0.66
Spa	0.88	0.52	0.66	1

h12

Table 5.1: Pearson correlation matrices of Expected Longrise across countries

Table 5.2 reports the Spearman correlation, which assesses how well the relationship between two variables can be described using a monotonic function, whether linear or nonlinear. The indices are for expected longrise and for both horizons.

	Ita	Ger	Fra	Spa
Ita	1	0.39	0.72	0.83
Ger	0.39	1	0.44	0.31
Fra	0.72	0.44	1	0.62
Spa	0.83	0.31	0.62	1

h6

	Ita	Ger	Fra	Spa
Ita	1	0.25	0.66	0.84
Ger	0.25	1	0.28	0.18
Fra	0.66	0.28	1	0.57
Spa	0.83	0.18	0.57	1

h12

Table 5.2: Spearman correlation matrices of Expected Longrise across countries

Expected longrise measures co-move significantly in the group of four euro area countries, with similar values, except for some outliers around 2015, as in the case of Spain, where the expected longrise is strongly negative. This is a period of very low inflation. The linear correlation matrix reveals that expected longrise measures are highly correlated among Italy, Spain and France; less so for Germany and the other countries. However, surprisingly, Germany is the most correlated to Italy out of all the countries. Spearman correlation sees even lower values of correlation between Germany and the other countries, while confirming a strong relationship among Italy, France and Spain.

My interpretation of this section is that monetary policy decisions and central bank communications, could potentially have relatively similar effects across the four countries, at least in terms of how these policies and communications impact the expected, worst-case scenarios' inflation rate values.

5.2.2 Correlations of Upside relative entropy

Figure 5.9 shows the upside relative entropy series for all countries.

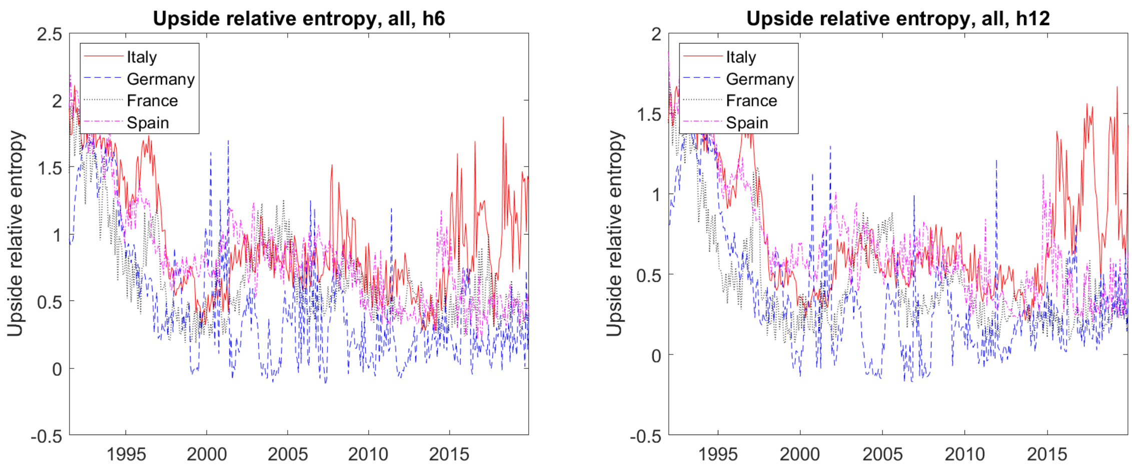

h6 h12

Figure 5.9: Upside relative entropy, comparison

Table 5.3 and Table 5.4 reports the Pearson and Spearman correlations, as before.

	Ita	Ger	Fra	Spa
Ita	1	0.56	0.67	0.64
Ger	0.56	1	0.5527	0.71
Fra	0.67	0.56	1	0.72
Spa	0.64	0.71	0.72	1

h6

	Ita	Ger	Fra	Spa
Ita	1	0.59	0.60	0.63
Ger	0.59	1	0.59	0.70
Fra	0.60	0.59	1	0.73
Spa	0.63	0.70	0.73	1

h12

Table 5.3: Pearson correlation matrices of Expected Longrise across countries

	Ita	Ger	Fra	Spa
Ita	1	0.34	0.57	0.43
Ger	0.34	1	0.31	0.50
Fra	0.57	0.31	1	0.49
Spa	0.43	0.50	0.49	1

h6

	Ita	Ger	Fra	Spa
Ita	1	0.42	0.48	0.42
Ger	0.42	1	0.35	0.42
Fra	0.48	0.35	1	0.5
Spa	0.42	0.42	0.5	1

h12

Table 5.4: Spearman correlation matrices of Upside relative entropy across countries

Also in the case of upside relative entropy the measures co-move across countries, but the relationship is noisier (given that the measure itself is more volatile compared to expected longrise). We can see that upside relative entropy is relatively stable at lower values for Germany, compared to Italy, for instance. The increase in relative entropy in the last years of the series that I mentioned in chapter 4 for Italy is not present in the case of the other countries; hence, we can deduce that the Phillips curve relation in the last observations of the series is still weak in the rest of the euro-area countries examined, compared to Italy; this is because there is not much divergence between the density conditional on the full set of regressors (Phillips curve) and the unconditional density.

When it comes to correlation among countries in this measure, the matrices report relatively low indices. Italy is most correlated with France and least with Germany, while Germany is most correlated with Spain.

Taking stock II The picture drawn by the correlations of the upside relative entropy is somewhat more pessimistic compared to expected longrise, in terms of symmetry and convergence, especially when considering some specific country pairs. In the next chapter, I study how the ECB can affect these risk measures through its communications associated to policy announcements, in a dynamic structural approximation of the real economy.

6

CBI shocks and inflation risk: a SVAR analysis

Following chapter 5, I continue my investigation of inflation at risk with a focus on the euro area. In this chapter I estimate a series of Structural Vector Autoregressions (SVAR) where I identify the effects of a central bank information shock, disentangled from the monetary policy shock in Jarociński and Karadi (2020)[1], on key economic variables and my inflation risk measures.

6.1 SVAR setup and identification

The specification of the SVAR is presented below:

$$\mathbf{y_t} = \begin{pmatrix} CBIproxy \\ URE \\ ExpLong \\ Infl \\ Eonia \\ Unemp \end{pmatrix}$$

[1] I re-estimate the VAR with sign-restrictions and internal instrument in Jarociński and Karadi (2020) to re-calculate the proxies for my sample length, based on the dataset EA-MPD (European Central Bank., 2019).

$$\mathbf{y_t} = \boldsymbol{\alpha} + \sum_{j=1}^{\rho} \mathbf{A_j y_{t-j}} + \mathbf{u}_t \tag{6.1}$$

$$\mathbf{u_t} = \mathbf{B}\epsilon_\mathbf{t} \tag{6.2}$$

where α represents the intercept (drift), ρ represents the lag-order of the VAR ($\rho = 3$) and u_t represents the vector of residuals from the OLS equation-by-equation regressions.

Formally, I estimate a recursive VAR with internal instrument (Plagborg-Møller and Wolf, 2021). This method imposes an arbitrary restriction in order to recover the structural shocks from the residuals of the reduced-form VAR and the inverted matrix of coefficients of the right-hand variables, thanks to the Cholesky decomposition of the variance-covariance matrix of the reduced-form residuals; at the same time, the internal instrument approach takes advantage of this restriction to ensure exogeneity of the proxy. Here I show more in detail how the method works:

$$\begin{pmatrix} CBIproxy_t \\ \dots \end{pmatrix} = \boldsymbol{\alpha} + \sum_{j=1}^{p} \mathbf{A_j} \begin{pmatrix} CBIproxy_{t-j} \\ \dots \end{pmatrix} + \boldsymbol{u_t} \tag{6.3}$$

$$\boldsymbol{u_t} = \mathbf{B}\epsilon_\mathbf{t} = \mathbf{B} \begin{bmatrix} \epsilon_{CBI,t} \\ \dots \end{bmatrix} \tag{6.4}$$

I impose that there is no contemporaneous effect of the CBI proxy on the other variables by setting to zero the parameter in the impact matrix $\mathbf{B}$ that multiplies $\epsilon_{CBI,t}$, effectively imposing that $\mathbf{B}$ is lower triangular. Following this, I calculate the Cholesky decomposition of the reduced-form residuals, that returns the product of a lower triangular matrix, the lower Cholesky factor, and its transpose. But since the impact matrix $\mathbf{B}$ is lower triangular, $\mathbf{B}$ is the lower Cholesky factor. This allows to estimate the impact effects of shocks, that is, the parameters of the structural shocks.

Plagborg-Møller and Wolf (2021) proved, among other important contributions, that internal-instrument SVARs are equivalent to Proxy SVARs.

The choice of a three-lags SVAR is based on the fact that the sample is short (1999-2019) and the parameters to be estimated are numerous (6 equations), hence I decided to be parsimonious. A lag-length test with 6 lags chooses a lag-length of 3[2]. I calculate the SVARs and the objects described in the next section for all four countries: Italy, Germany, France and Spain[3].

6.2 Analysis of the dynamic effects

After having estimated the SVAR, I proceed to calculate the Impulse Response Functions (IRFs), the main object of interest of this chapter. IRFs enable the quantification of the causal dynamic effect of a shock on the system. In other words, they allow to see the response over time of the SVAR's endogenous variables to an innovation (usually of one standard deviation) in one structural shock, assuming that the other structural shocks are kept to zero. Such responses should not be fundamentally at odds with economic theory and should go back to zero in the very long run. Figure 6.1 to Figure 6.4 report the IRFs.

Next, I calculate the Forecast Error Variance Decompositions (FEVD) and Historical Decompositions (HDs). The FEVDs show the portion of the variance of the SVAR's forecast errors (at a given horizon) due to each structural shock. In other words, it provides information about the relative importance of each structural shock in affecting the forecast error variance of the SVAR's endogenous variables. Figure 6.5 to Figure 6.8 show the FEVDs. The Historical Decompositions show the historical contribution of each structural shock in driving deviations of the SVAR's endogenous variables away from their equilibrium. It allows to track, at each point in time, the role of structural shocks in driving the SVAR's endogenous variables away from their steady state. Figure 6.9 to Figure 6.12 report the HDs.

The table below avoids possible confusion due to the different variable names in the pictures shown in this section, compared to the previous one.

CBI_MEDIAN	Central Bank Information proxy
EONIA_RATE	Eonia interest rate
CPGRLE01ITM659N (and similar)	Core inflation rate
LRHUTTTTITM156S (and similar)	Unemployment rate
EXPLONG_H6	Expected longrise, horizon 6 months
UPSRELENTR_H6	Upside relative entropy, horizon 6 months

[2] I do not report this test because of conciseness.

[3] I choose to report only the six-months forecasting horizon case because of the large amount of results. The implications are basically the same for the case of the year-long horizon.

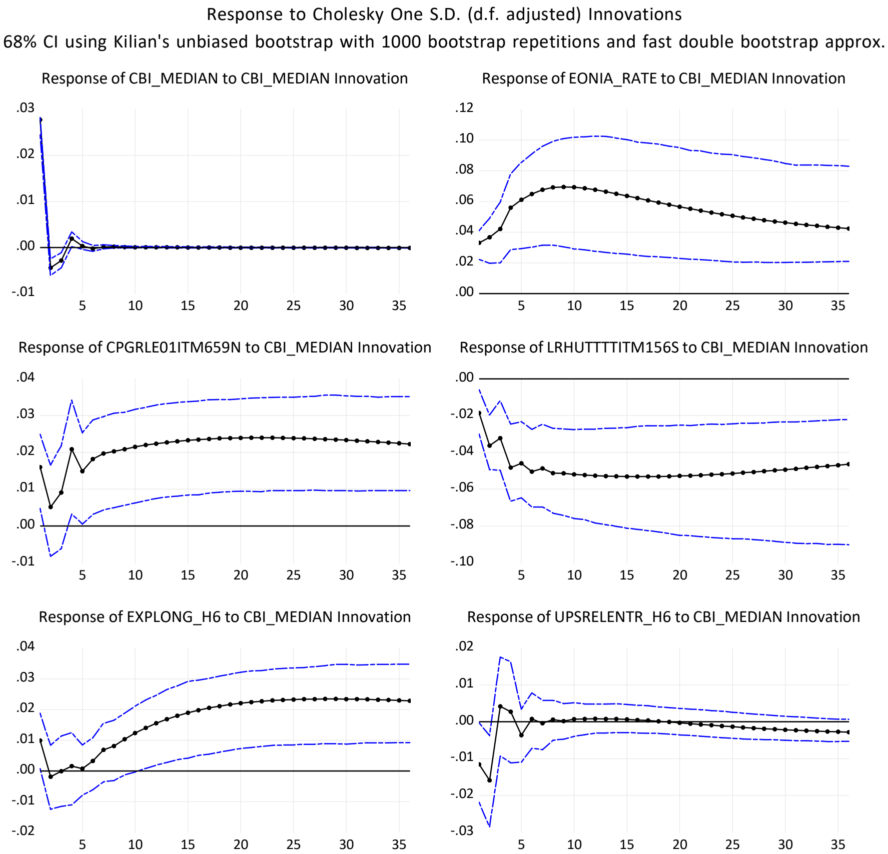

Figure 6.1: IRFs Italy h6

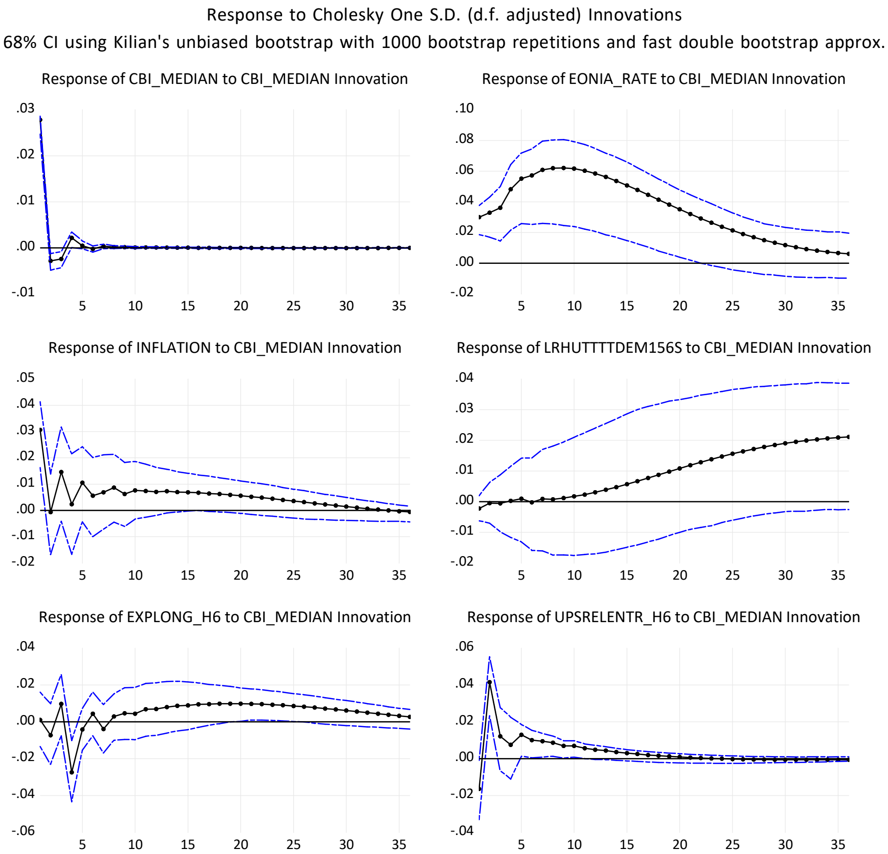

Figure 6.2: IRFs Germany h6

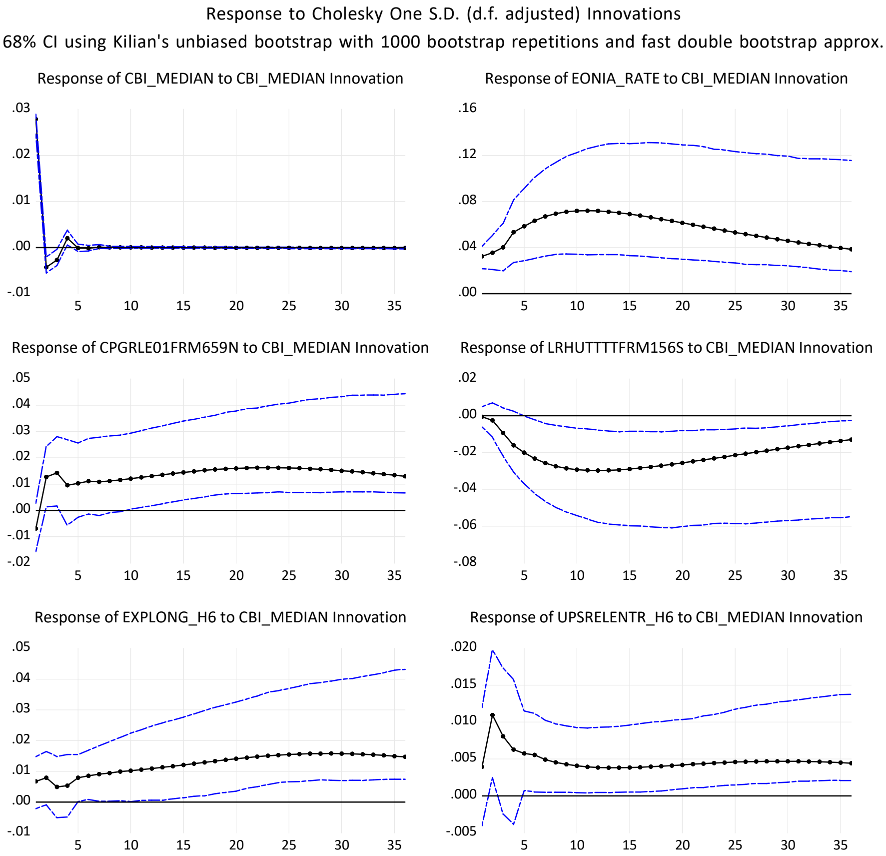

Figure 6.3: IRFs France h6

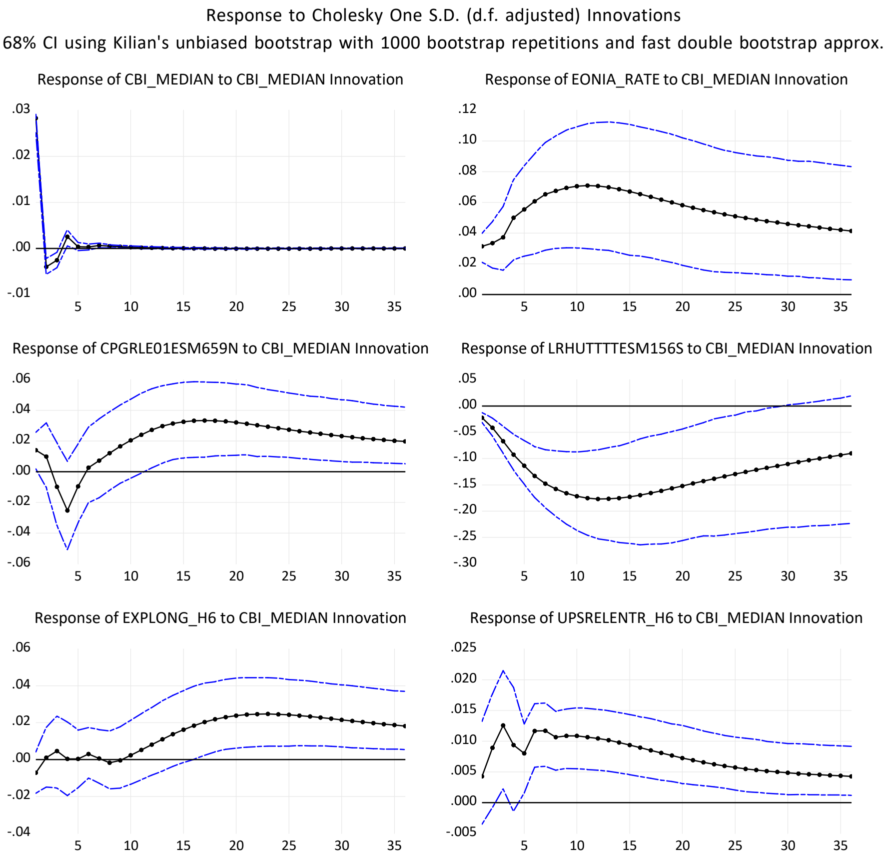

Figure 6.4: IRFs Spain h6

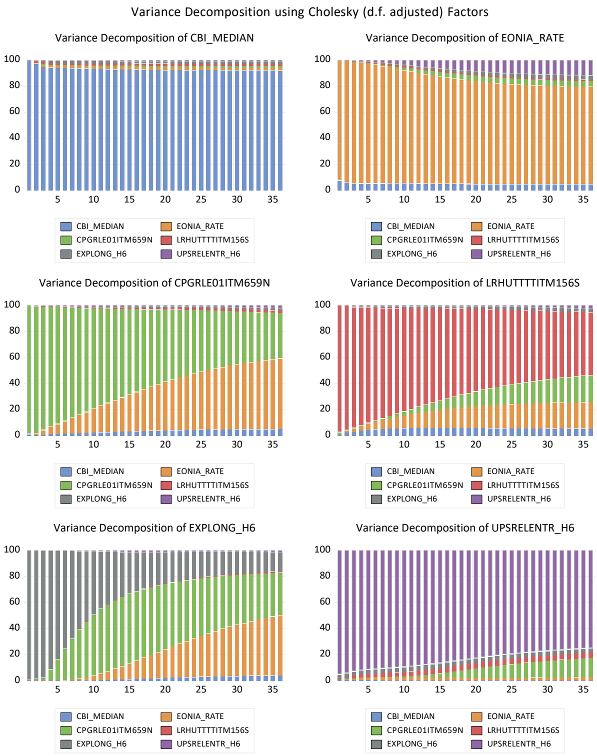

Figure 6.5: FEVDs Italy h6

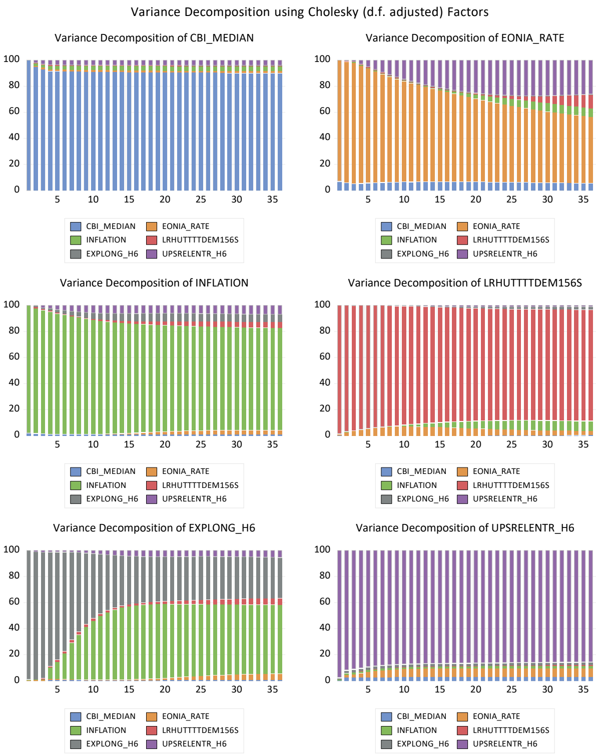

Figure 6.6: FEVDs Germany h6

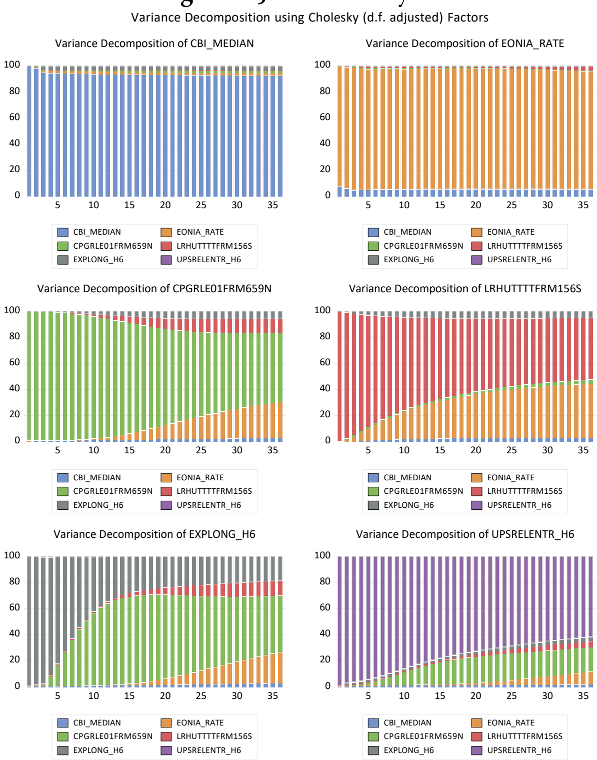

Figure 6.7: FEVDs France h6

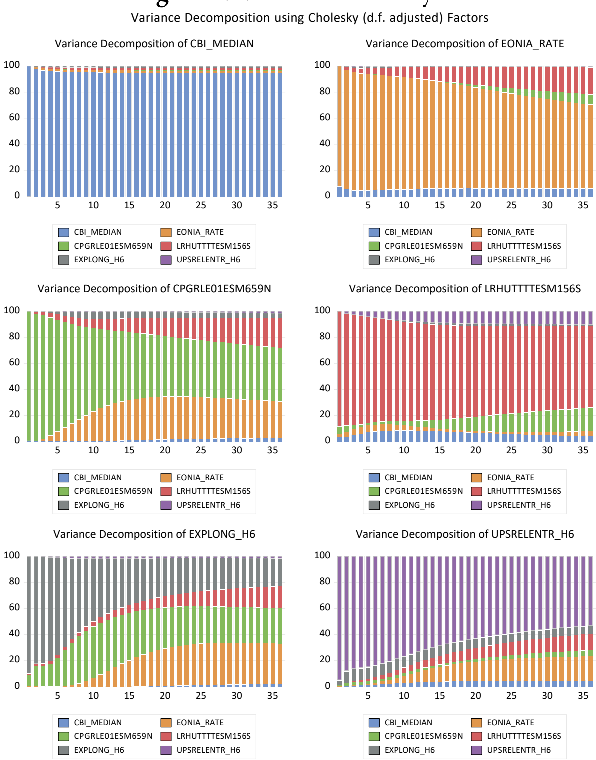

Figure 6.8: FEVDs Spain h6

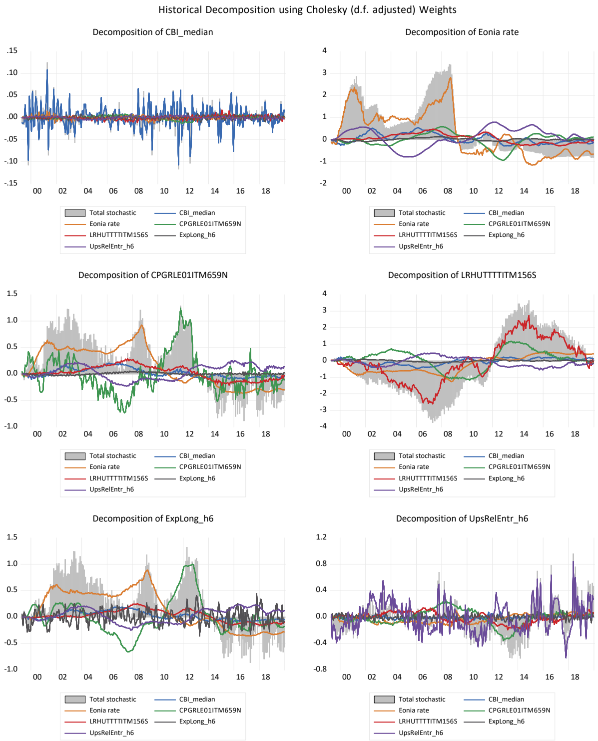

Figure 6.9: HDs Italy h6

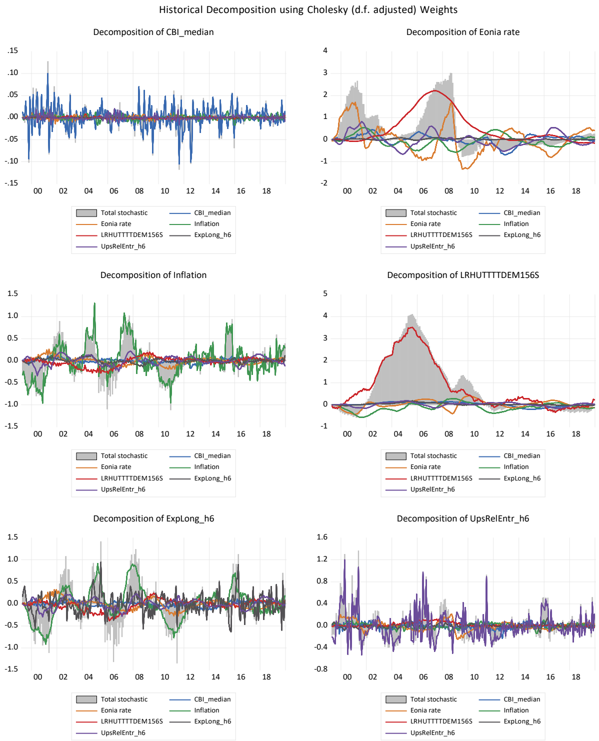

Figure 6.10: HDs Germany h6

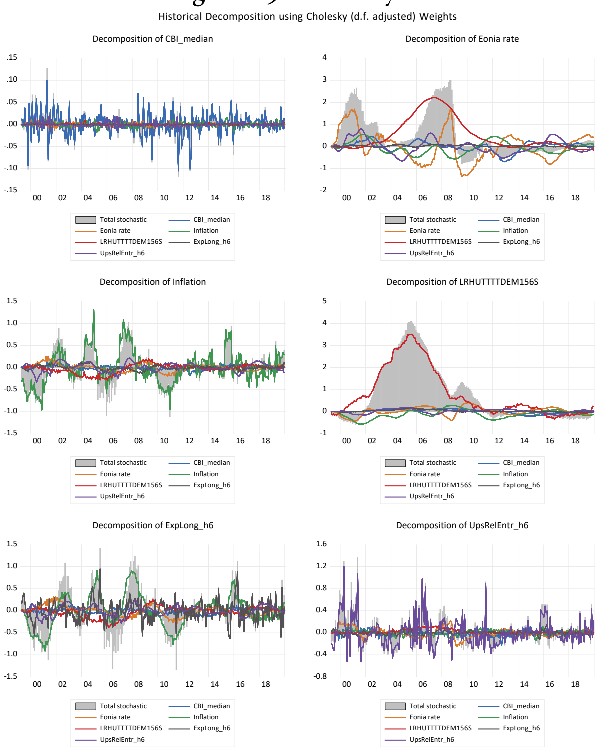

Figure 6.11: HDs, France h6

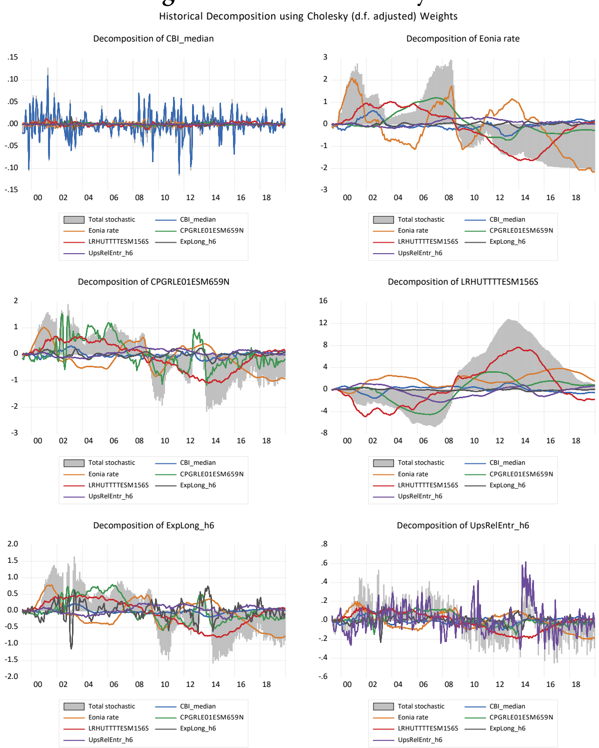

Figure 6.12: HDs, Spain h6

6.3 Results

Main findings For all four countries, a one standard deviation shock in the central bank information proxy leads to an increase in the EONIA rate, as expected, given that CBI shocks occur around policy announcements.

The inflation rate's immediate reaction to the shock is not significant for Italy, France and Spain, but becomes significant after 10 to 15 months; for Germany instead it is significant on impact, but the response quickly becomes non-significant.

In Italy, Spain and France the response of the unemployment rate is negative, coherently with the increase in inflation; in Germany this response is never significant, as we might expect given the behaviour of the inflation rate's impulse response (the point estimate of the unemployment rate in Germany goes back to zero over a very long horizon, not shown here).

Ultimately, we get to the IRFs of the risk measures. The expected longrise for Italy, France and Spain is non-significant in the first periods, but becomes significant after 15 to 20 periods, similarly to the inflation rate, although with delay. This means that the central bank communication leads not only to an increase in the modal forecast of the inflation rate, but also in the expected value of the conditional density of the inflation rate beyond the 90th percentile. Upside inflation risks increase, as the worst-case scenario inflation rate is higher in expectation. In Germany, this effect is not statistically significant, just like for inflation and unemployment.

Now we come to the impulse response functions of upside relative entropy[4].

For Italy, the upside relative entropy decreases significantly in the second to third period, before quickly becoming non-significant and going back to zero; hence the divergence between conditional and unconditional distributions, from the median above, decreases slightly. There is less probability mass in the right tail of the conditional density compared to the unconditional one. The conditional density attributes less probability to extreme right-tail outcomes compared to the unconditional density; this implies that vulnerability, or risk, around the modal forecast decreases.

Germany suffers a stronger and more significant increase in upside relative entropy around the same period, before its response too goes quickly back to zero.

France's upside relative entropy response is small and non-significant on impact, but it becomes significant for many periods before converging to zero (beyond the horizon shown here).

Spain's upside relative entropy response is similar to France's, although somewhat stronger.

[4]Go back to section 4.3 for reference on the interpretation of URE.

The FEVDs show that the relevance of the unemployment rate in explaining the dynamics of the expected longrise and the upside relative entropy is very low.

It is worth noting that, when both expected longrise and upside relative entropy increase, the conditional distribution of inflation becomes more right-skewed compared to the unconditional distribution.

Taking stock Overall, I find that European Central Bank's communication associated to policy announcements releases implicit information that is basically "sterile" for Germany, since the impulse responses of its inflation and unemployment rate, as well as expected longrise over 6 months, are small and not statistically significant[5]. However, the country suffers from a slight increase in vulnerability to inflation risk, in terms of the probability of extreme outcomes, as represented by upside relative entropy.

The same cannot be said for the other three countries. Italy, France and Spain see their inflation and unemployment rate respond significantly, although with a lag, and their risk in absolute terms increase, as represented by expected longrise. France and Spain become more vulnerable over a longer horizon with respect to Germany. Italy instead becomes slightly less vulnerable.

[5] If not for inflation rate in the very first period.

7

Discussion

Having presented the econometric methodology and main results of my dissertation, in this chapter I want to discuss potential extension, alternatives and caveats, some of which I have explored.

7.1 A modern Phillips Curve

Obviously, the model specification that I have adopted — a Phillips curve in levels of inflation and unemployment, not accounting for expectations, mark-ups and other potentially relevant determinants — does not allow me to make reliable and accurate forecasts of inflation — this was never the aim, as explained previously. The objective is to study the relationship between inflation rate itself and the main covariate prescribed by economic theory, the unemployment rate.

An immediate extension of my model would be the implementation of a more modern and advanced Phillips curve, such as the Augmented Phillips curve model used in the study of inflation at risk in the USA by López-Salido and Loria (2024)[1][2]. The authors' augmented Phillips curve is New Keynesian, since it accounts for inflation expectations; as we know from the literature, expectations are a major component affecting inflation dynamics. In addition to other relevant covariates, they add a regressor that accounts for financial conditions. Financial conditions explain well the downside

[1]Their specification extends the Phillips-curve by Blanchard et al. (2015)

[2]My methodological contribution, compared to López-Salido and Loria (2024), is the use of the two-steps quantile regression in the study of inflation.

risks of the inflation rate and I believe that an implementation of this variable with my econometric approach would improve further not only the explanatory power of the Phillips curve, but also the fit of the conditional quantile function, improving reliability and accuracy of the estimations and out-of-sample forecasts.

For these reasons, I have reproduced the core of my analysis for a model with the headline inflation rate, unemployment rate and financial conditions — the main results are shown in Appendix A.

Other covariates used by López-Salido and Loria (2024), that are relevant in the pandemic era and in even more recent times, are oil price shocks, supply shocks and, most importantly, fiscal policy. All these predictors are variables that I would like to extend my model for the euro area with.

7.2 A more flexible distribution

In chapter 1 I talked about how the field of economic risk management spawned from the need of studying the risks associated to economic variables, and how the tools and techniques adopted by researchers who venture in the field are often interdisciplinary, taking especially from quantitative finance.

In fields such as quantitative finance, financial engineering and computational finance, for various reasons, often the need for a flexible distribution arises. Adrian et al. (2019) in the two-steps QR use the Skew-t by Azzalini and Capitanio (2003). Azzalini published some of his most important works on this subject in *Biometrika*, an important biostatistics journal. Applications of his and his co-authors' Skew-t, to the best of my knowledge, are mainly in biological sciences.

Although there surely are applications of this distribution to finance as well, arguably the most adopted and preferred distribution based on the Student's t, in finance, is the Skewed Generalised T (SGT) distribution by Theodossiou (1998).

The probability density function of the SGT, as defined in Davis (2015) and similarly[3] in BenSaïda and Slim (2016) is as follows:

$$f_{SGT}(x;\mu,\sigma,\lambda,p,q) = \frac{p}{2\upsilon\sigma q^{\frac{1}{p}}\left(\frac{1}{p},q\right)\left(\frac{|x-\mu+m|^p}{q(\upsilon\sigma)^p(\lambda sign(x-\mu+m)+1)^p}+1\right)^{\frac{1}{p}+q}}$$

where B is the beta function, μ is the location parameter, controlling the mean (which becomes the mode in the presence of skewness), σ is the scale parameter controlling the variance, λ is the skewness parameter, p and q are parameters controlling the kurtosis, υ and m are functions of the parameters

[3] Up to a change in notation.

and the beta function and they can take themselves on different values depending on how we want to fit the distribution.

Although I believe that the Skew-t distribution is able to reproduce sufficiently well the real kurtosis of the data, for models such as the Phillips curve in my dissertation, where the kurtosis seems to take a prominent role compared to, for instance, Adrian et al. (2019) and their model of GDP growth at risk, I believe that using the 5-parameters SGT distribution would improve substantially the fit, in addition to assuring robustness when fitting a more complex and sophisticated Phillips curve, such as the one I talked about in section 7.1.

The SGT distribution can be used with the "sgt" package in R (Davis, 2015) or with the "flexible distributions" toolbox in Matlab (BenSaïda and Slim, 2016) [4].

7.3 A possible caveat

Finally, I believe it is appropriate to mention the reason for not studying the effects of a pure monetary policy shock on the economic indicators and risk measures in the G4.

The reason is that the only pure monetary policy proxy (that is, disentangled from the CB's implicit information shock) available for the euro area, the one by Jarociński and Karadi (2020), is proving to be problematic, in the sense that in SVARs it produces impulse response functions that are not compatible with economic theory and with similar results obtained for analogous models in other countries. The CBI proxy, instead, seems not to be affected by these issues, while still being able to provide interesting results and policy implications.

[4] The SGT distribution is usually fitted using MLE.

8

Conclusion

As shown in chapter 5, the modal forecast of the average conditional inflation rate in the four euro-area countries at study has undergone a convergence process in the sample considered. Variance has also substantially decreased. These are positive stylised facts attributable to the integration resulting from these countries entering a single currency area — the euro area — and relinquishing control of their sovereign monetary policy in favour of a communitarian one, controlled by the ECB and the Eurosystem.

However, I do not find that the same convergence process has taken place when it comes to the degree of risk and uncertainty in the distributions.

At a given period, one country can be at risk of lower inflation than expected while another of higher inflation than expected, as in the case of Spain and Germany around 2015 (Figure 5.3).

While some countries' Phillips curve relations seem to be robust and largely unaffected by outliers, like for Italy and France in recent times, in the same period other countries' relations seem to be vulnerable to them, while being skewed in opposite directions — as in the case of Germany and Spain (Figure 5.7).

A thought-stimulating result is the exchange of variance for kurtosis, something that we see clearly from the 3D plot of the predicted densities across time. It is a sort of switch from risk in the traditional sense, to uncertainty, as captured by the degree of outliers.

The correlations in expected longrise are roughly from 60 to 90 percent for expected longrise (Table 5.1) and from 40 to 85 percent for upside relative entropy (Table 5.3), depending on the country-pairs. I would argue that these figures and their variability are not alarming, but do not look very promising either, in terms of the degree of similarity in inflation risks — a condition that I believe necessary in order to obtain stability around the target inflation rate in the medium to long run. These correlations are calculated on the whole sample. A reduced sample would probably find stronger correlations for expected longrise in the very recent years (Figure 5.8); however, the same cannot be said for the risk of extreme realisations taking place unexpectedly, as represented by upside relative entropy (Figure 5.9).

The SVAR analysis in chapter 6 suggests that European Central Bank's communication are either coincidentally suitable or deliberately calibrated not to shock significantly the German economy. The analysis also shows evidence for structural heterogeneity among the G4 countries. Germany's inflation rate, unemployment rate and expected longrise do not respond to CBI shocks, while the upside relative entropy responds positively but is significant only in the short term. The other countries suffer from an increase in inflation and respond with a decrease in unemployment, both with a lag. Their expected longrise also responds to the upside, with a lag. Upside relative entropy responds heterogeneously, with varying sign and different degrees of significance. For the purposes of sticking to the target inflation rate, I believe that this heterogeneity is overall detrimental and makes the task of the central bank substantially mode difficult.

The policy implications of my dissertation are the following: First, the central bank should better calibrate and manage the information released by their communications associated to policy announcement, with the aim of increasing the symmetry in the responses of economic indicators and risk measures across countries. This could help in maintaining price stability.

Second, fiscal policy coordination across countries should be strengthened. Fiscal reforms should also be aimed at smoothing out the structural differences across countries. The ECB, when assessing monetary policy options, should tackle the issues with the Phillips curve highlighted in this dissertation, such as strong heterogeneity in vulnerability to inflation risks. New tools and regulations could be necessary in order to strengthen the symmetry across countries and protect integration.

From the point of view of the individual countries, and from the narrow perspective of price stability, it is not clear that the single currency area is beneficial, for instance, to southern European countries if the German economy is taken as reference for monetary policy decisions, and vice versa.

A

Appendix: IaR with headline inflation and financial conditions

Introduction This appendix extends the model introduced in chapter 2, by accounting for financial conditions, while implementing a slight modification: the use of the headline inflation rate instead of the core one. This last choice is taken because headline inflation, being generally more volatile due to the inclusion of items such as food and energy prices, can now be better studied by an expanded model.

The headline inflation rate for each country is taken by FRED. I keep the same unemployment rate indicator mentioned in chapter 2. The financial conditions indicator is the Sovereign Composite Indicator of Systemic Stress, or SovCISS — an improved, monthly version of the CISS (Holló et al., 2012), developed by European Central Bank. (2018).

SovCISS is a composite indicator measuring the multidimensional sovereign bond market stress in the euro area as a whole as well as in individual euro area member states. It integrates measures of credit risk, volatility and liquidity at different maturities. The statistical method is inspired by that of the Composite Indicator of Systemic Stress (CISS).

The use of two-steps quantile regression in the study of inflation at risk in the euro area, conditional on unemployment and financial conditions, represents a further innovation.

As López-Salido and Loria (2024) find, in the case of the USA, financial conditions are good explanatory variables of downside inflation dynamics. For this reason, in this appendix I take a different approach and study both expected longrise and shortfall, as well as the full relative entropy instead of just the upside one. In the SVARs I only include the full relative entropy as risk measure, given the addition of SovCISS as an additional equation and the shorter sample.

The sample goes from September 2000 to December 2019 [1].

The appendix structure encompasses a shorter version of the comparison among countries, akin to chapter 5, and the IRFs results, similarly to chapter 6.

Model The following is the extended model

$$y_h = \alpha + \mathbf{x_t}'\boldsymbol{\beta} + \varepsilon_h \tag{A.1}$$

where:

$$y_h = \frac{1}{h}\sum_{i=1}^{h} y_{t+i}$$

y_h: discrete moving average of the inflation rate over the forecasting horizon.

and

$$\mathbf{x_t} = \begin{pmatrix} y_t \\ u_t \\ f_t \end{pmatrix}$$

y_t : Inflation rate at time t

u_t : Unemployment at time t

f_t : Financial conditions at time t

Results Below I show the results from an analogous analysis to the one made in chapter 5 and chapter 6, to which the reader can refer for a detailed explanation.

We can see in Figure A.1 and Figure A.2 that the estimated parameters for inflation and unemployment at each quantile are different compared to the figures in chapter 5.

There is some degree of rejection of the linearity null for inflation in the lower quantiles, as well as in the upper quantiles for Spain. The unemployment rate is still linear, although some parameters

[1]The sample has been shortened because SovCISS is a shorter series.

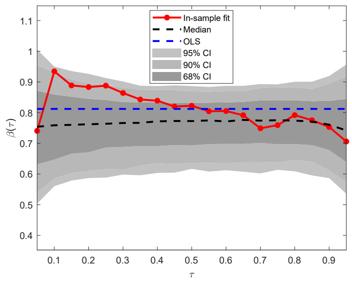

Coefficients of y_t, Italy

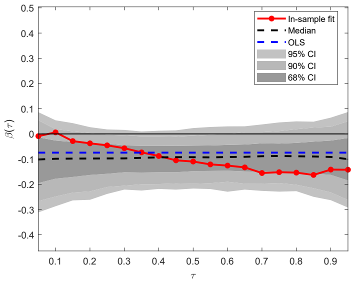

Coefficients of u_t, Italy

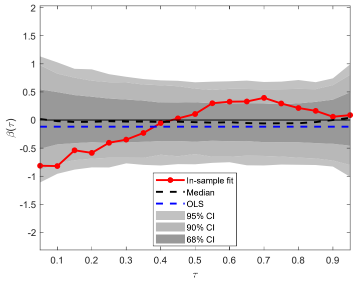

Coefficients of f_t, Italy

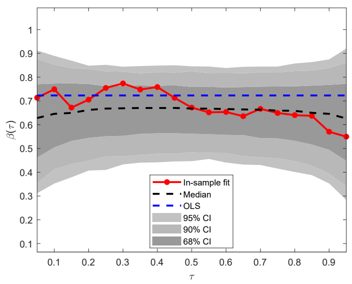

Coefficients of y_t, Germany

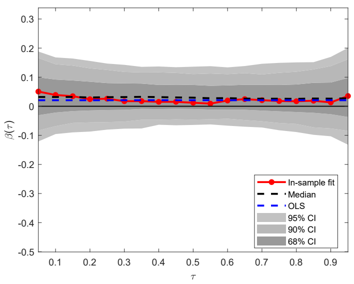

Coefficients of u_t, Germany

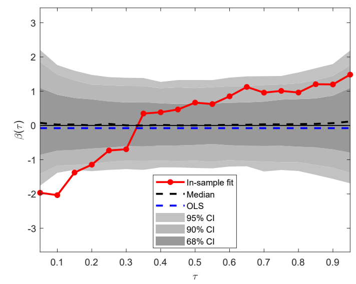

Coefficients of f_t, Germany

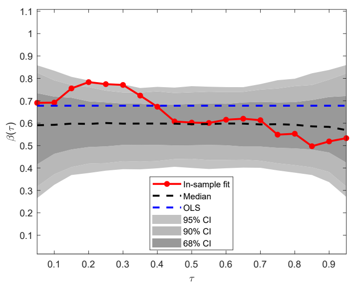

Coefficients of y_t, France

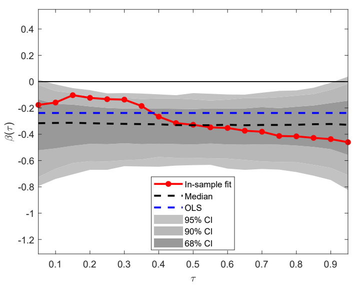

Coefficients of u_t, France

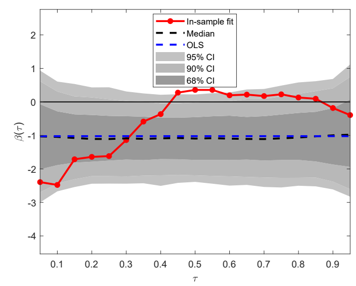

Coefficients of f_t, France

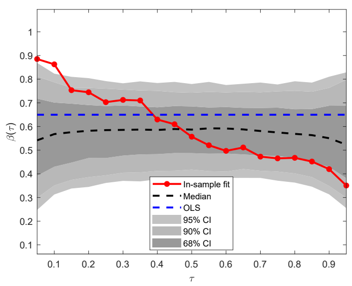

Coefficients of y_t, Spain

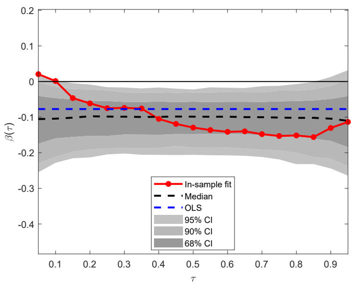

Coefficients of u_t, Spain

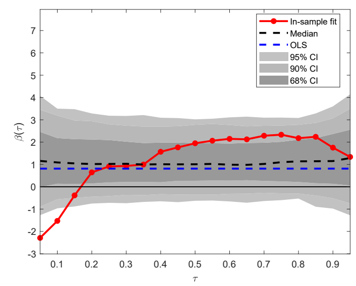

Coefficients of f_t, Spain

Figure A.1: Comparison, coefficients from multivariate QR, h6

reject the null at 10 percent in the lower quantiles; here, Spain strongly rejects linearity for the 10th quantile and below. Finally, the estimated parameters for the financial conditions change substantially across the quantiles' spectrum, similarly among countries. Italy, perhaps surprisingly, is the country suffering less from nonlinear effects of financial conditions on inflation. Germany and Spain manifest a strong nonlinearity in the lower quantiles, while France in the central quantiles (but only for the

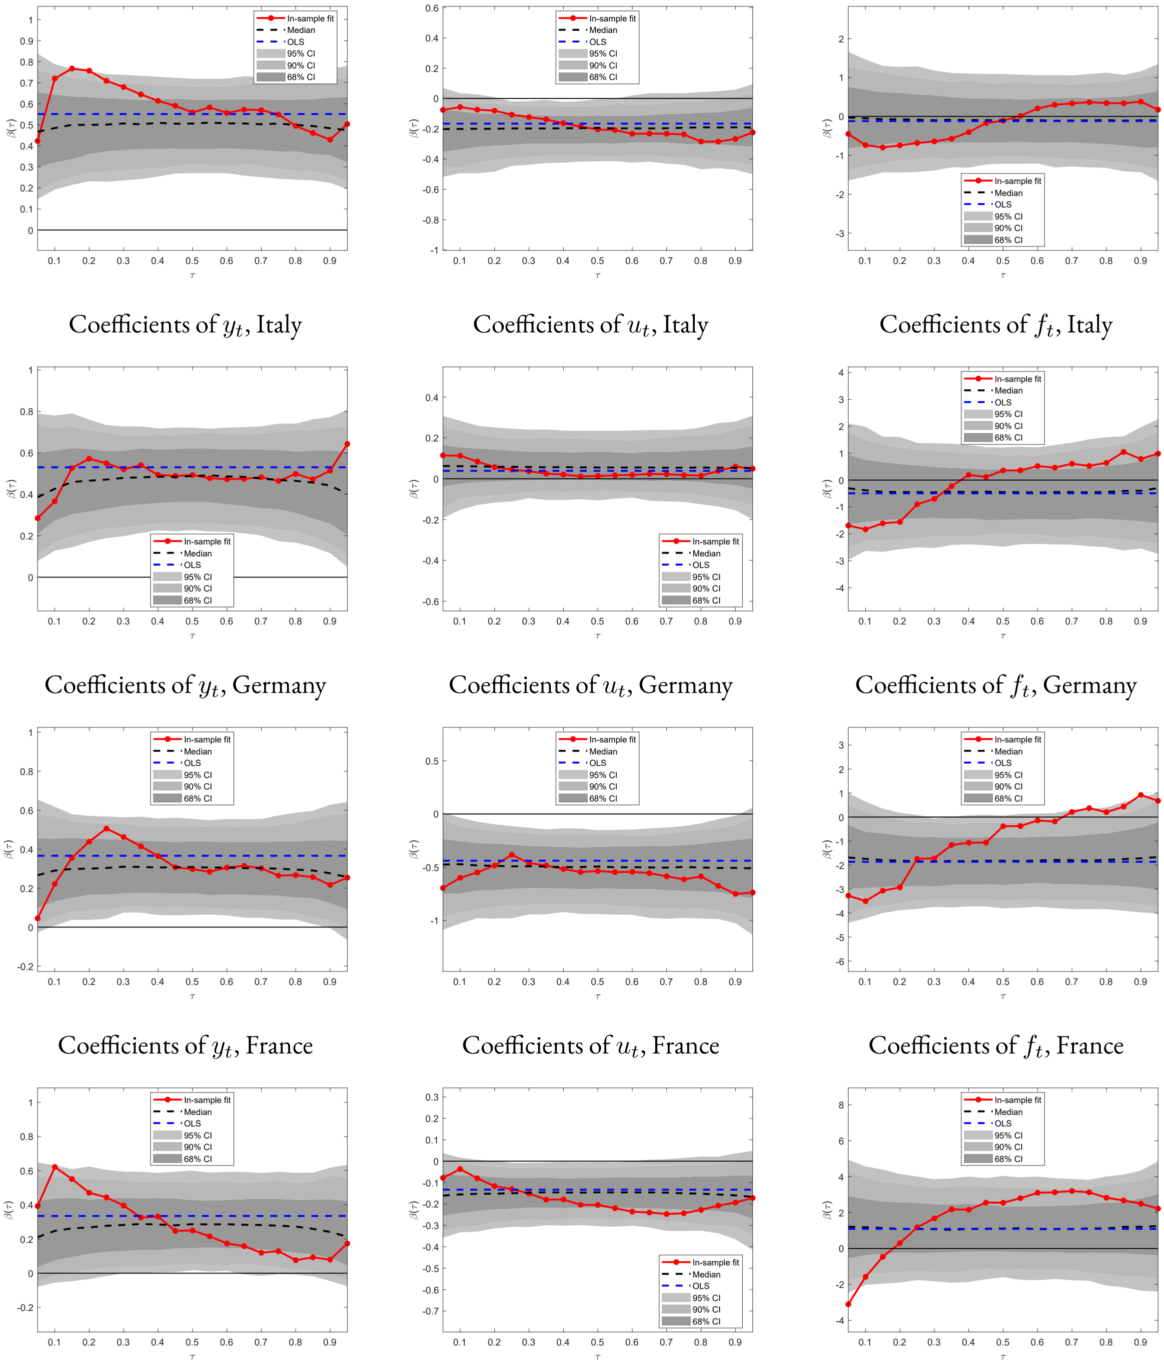

Coefficients of y_t, Italy — Coefficients of u_t, Italy — Coefficients of f_t, Italy

Coefficients of y_t, Germany — Coefficients of u_t, Germany — Coefficients of f_t, Germany

Coefficients of y_t, France — Coefficients of u_t, France — Coefficients of f_t, France

Coefficients of y_t, Spain — Coefficients of u_t, Spain — Coefficients of f_t, Spain

Figure A.2: Comparison, coefficients from multivariate QR, h12

6-months horizon).

As a general note, these nonlinearities seem to be stronger in the shorter forecasting horizon (h6) rather than in the year-long one (h12).

Figure A.3 and Figure A.4 display the time series of (selected) predicted quantiles, for the conditional forecast of the MA of inflation rate for both horizons and all countries.

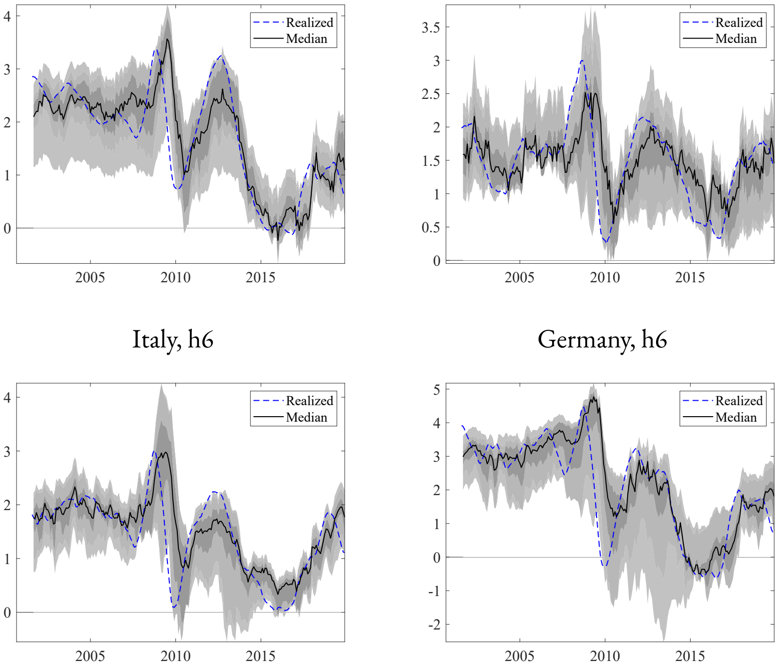

Italy, h6 Germany, h6

France, h6 Spain, h6

Figure A.3: Predicted distribution of inflation rate (TS), comparison, h6

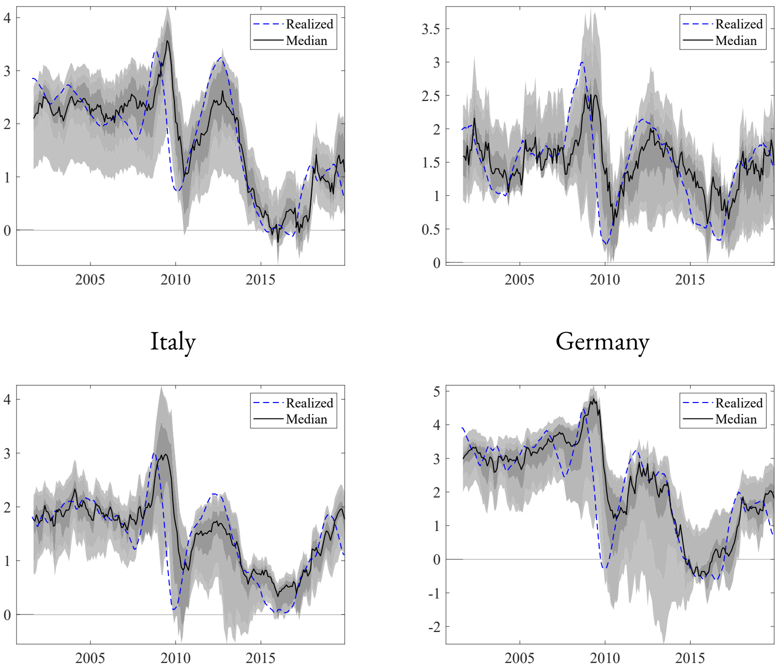

Italy Germany

France Spain

Figure A.4: Predicted distribution of inflation rate (TS), comparison, h12

The striking result is the overwhelming evidence of additional risk, both to the upside and to the downside, that is accounted for by extending the model with the financial conditions. Part of this additional risk is due to the use of headline inflation rate, but in a different version (not included in this dissertation) where I estimate a model with headline inflation rate and unemployment rate only, I do not see this level of variation in the quantiles, nor do I find evidence for the level of heterogeneity across countries that is present in this case.

While Italy and France's predicted distributions look similar, Germany and Spain's ones are very different, not in the modal forecast (and realised values) but in their risks. Spain suffered from substantial downside inflation risk roughly after 2007. Also, from a visual inspection, the forecasts are the least accurate for Germany and Spain out of the four countries.

Figure A.5 shows the selected CQFs for Italy, to verify if the fitting process has improved by accounting for financial conditions.

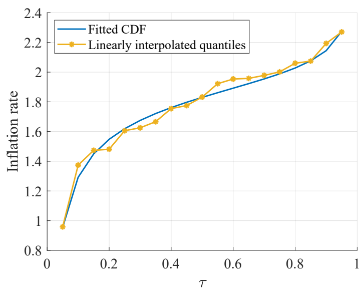

h6, 2013M1

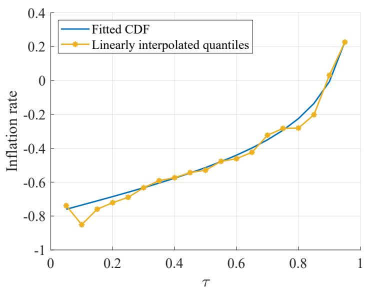

h6, 2015M1

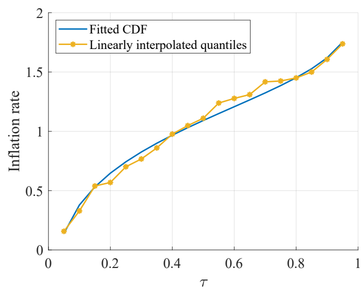

h6, 2018M12

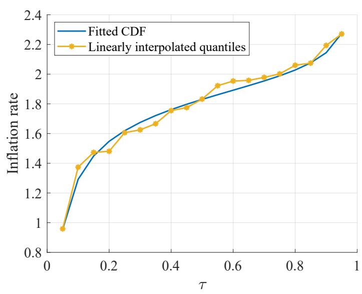

h12, 2013M1

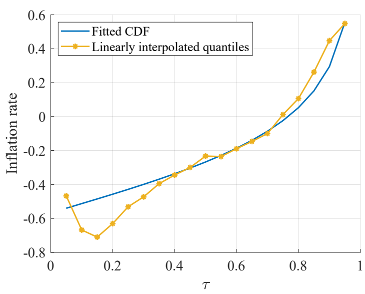

h12, 2015M1

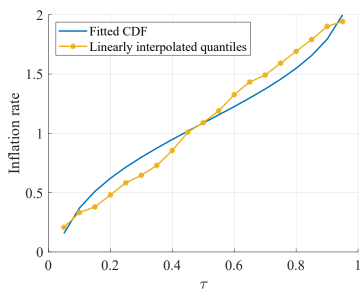

h12, 2018M12

Figure A.5: CQFs

The fitting process, at a visual inspection, has marginally improved, although there are still some issues for 2015 in its lower quantiles. Potential solution might be the addition of more explanatory variables, the adoption of an over-identification of the parameters, the alteration of the initial conditions or modification of the optimisation process, or the use of a different distribution altogether, such as the SGT (Theodossiou, 1998).

In Figure A.6 and Figure A.7 I show the predicted densities.

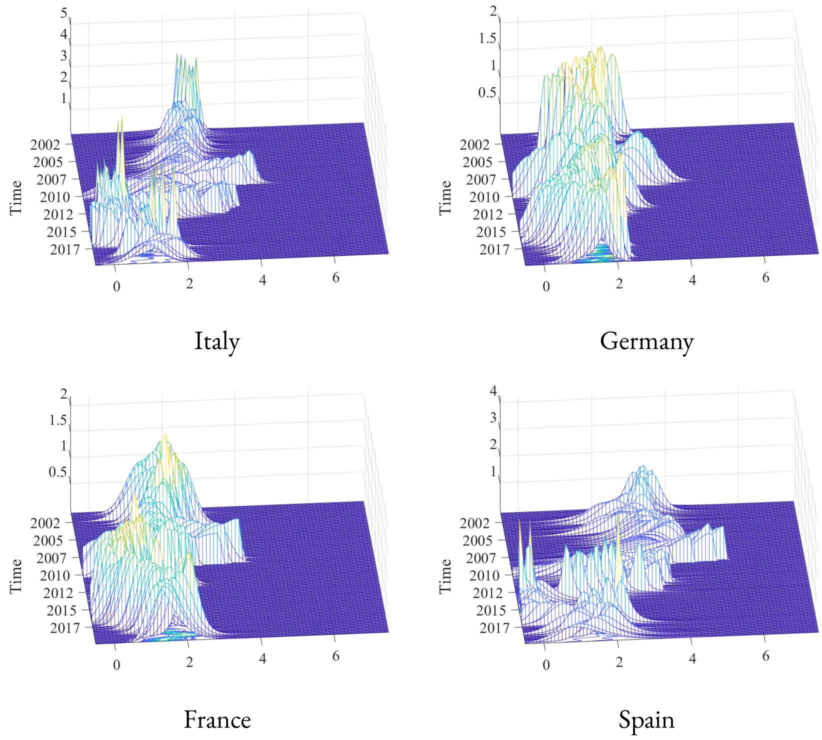

Figure A.6: Predicted densities, G4 comparison, h6

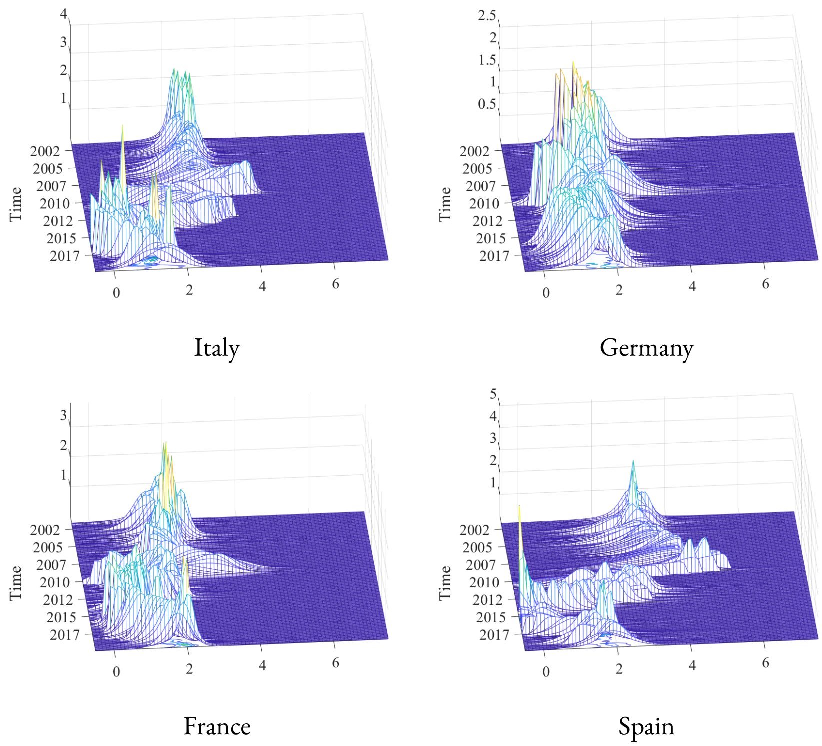

Figure A.7: Predicted densities, G4 comparison, h12

These vary emphatically over time, with very different levels of skewness and kurtosis, but relatively contained variance for most periods and countries (the worst-off, in this aspect, is again Spain).

The moments are reported in Figure A.7, for the 6-months forecasting horizon.

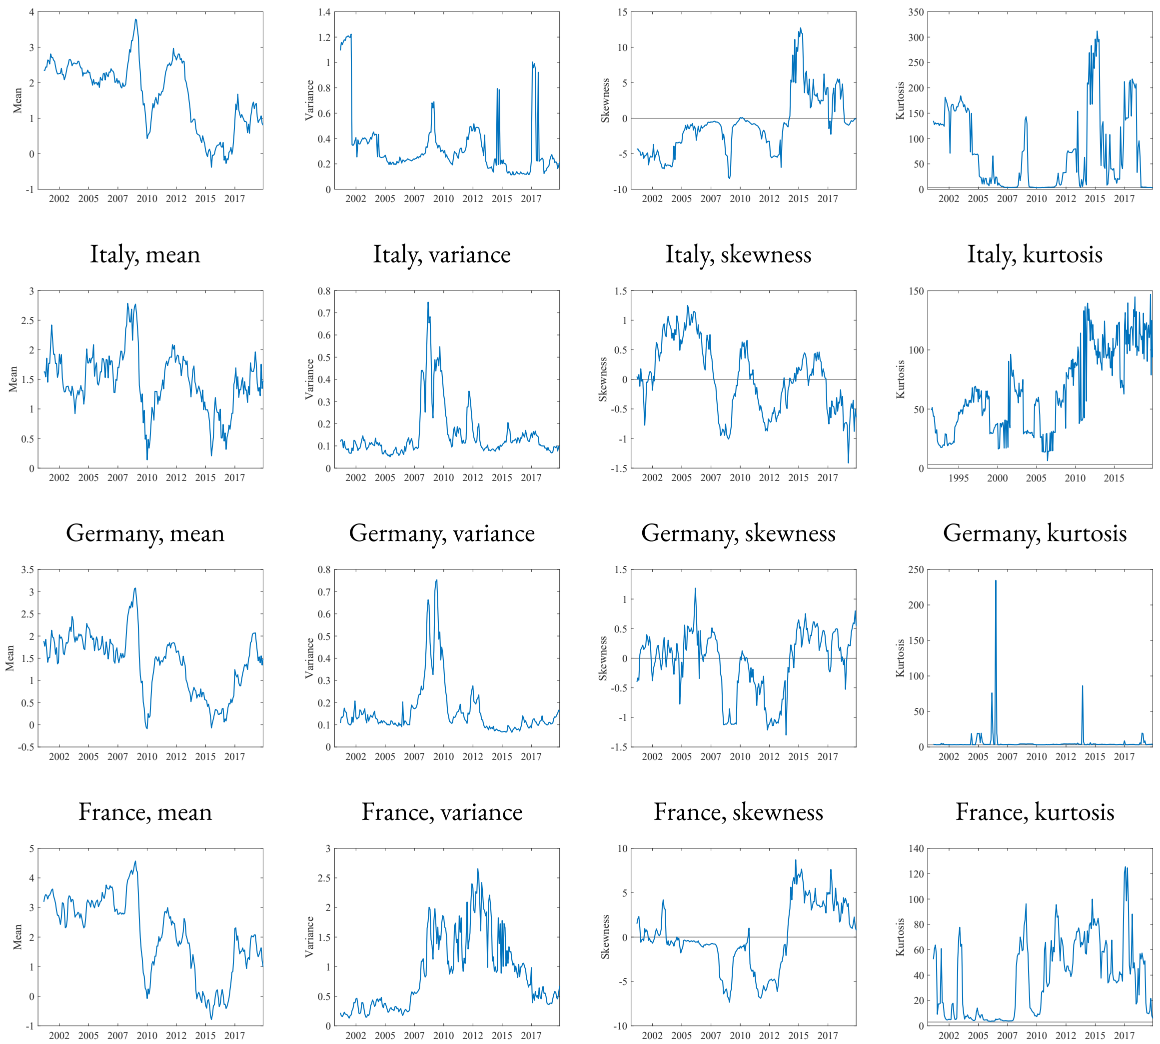

Italy, mean — Italy, variance — Italy, skewness — Italy, kurtosis

Germany, mean — Germany, variance — Germany, skewness — Germany, kurtosis

France, mean — France, variance — France, skewness — France, kurtosis

Spain, mean — Spain, variance — Spain, skewness — Spain, kurtosis

Figure A.8: Moments, all countries, h6

For recent years, I find a large degree of heterogeneity across countries in skewness and especially kurtosis. Italy, Germany and Spain show a large level of outliers and hence are vulnerable to unexpectedly high or low realisations.

As for the tests, I do not provide them here given the already extensive amount of results, but I summarise the findings: the predictive scores improve substantially, the empirical CDF of the PITs are still within the confidence bands and the out-of-sample prediction of quantiles and Upside relative entropy do not signal issues.

I compare the risk measures of full relative entropy and both expected shortfall and longrise in the next figures [2].

Figure A.9, Figure A.10 and Figure A.11 show the full relative entropy, expected longrise and expected shortfall comparisons across countries, for both horizons.

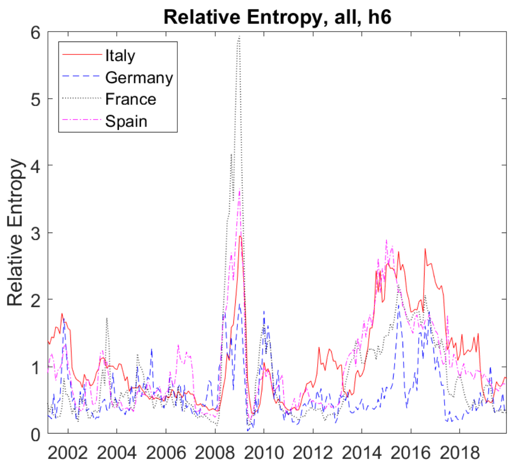

h6

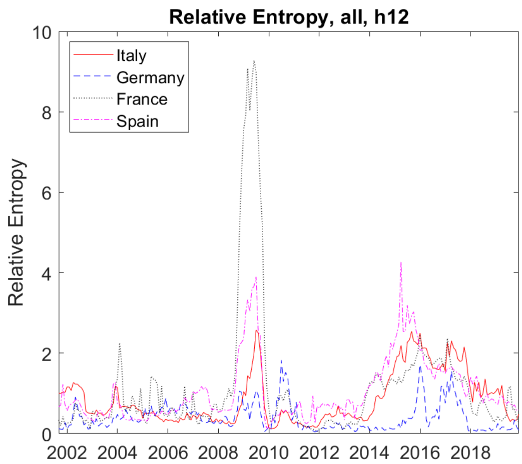

h12

Figure A.9: Full relative entropy, comparison

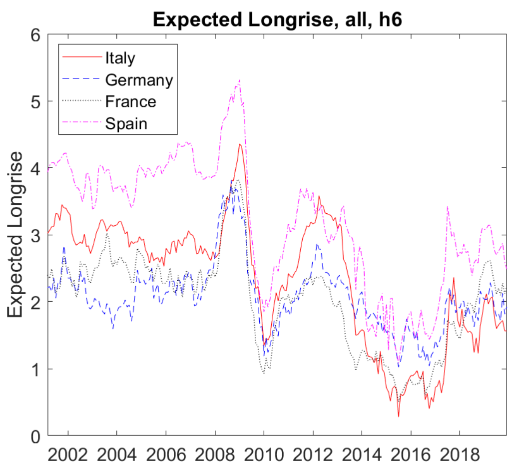

h6

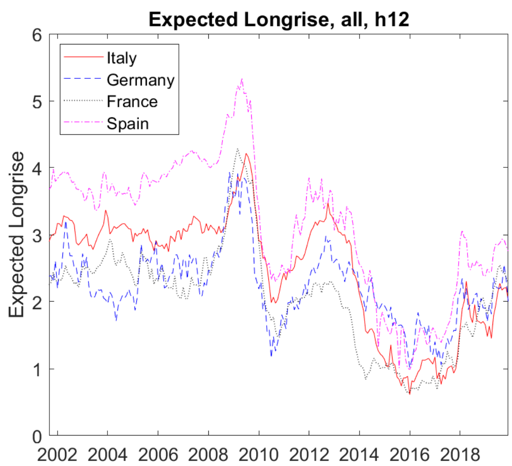

h12

Figure A.10: Expected longrise, comparison

(Full) relative entropy increased markedly around the 2008 GFC, particularly in France. I find also that from roughly 2015 to late 2018 relative entropy was high. This means that the risk of extreme values both to the upside and downside is larger for the conditional distribution than for the unconditional distribution.

[2] I redirect the reader to Adrian et al. (2019) for an explicit formulation of expected shortfall, which is just the 10% CVaR. Full relative entropy is analogous to URE, but with the integral taken over from $-\infty$ to ∞.

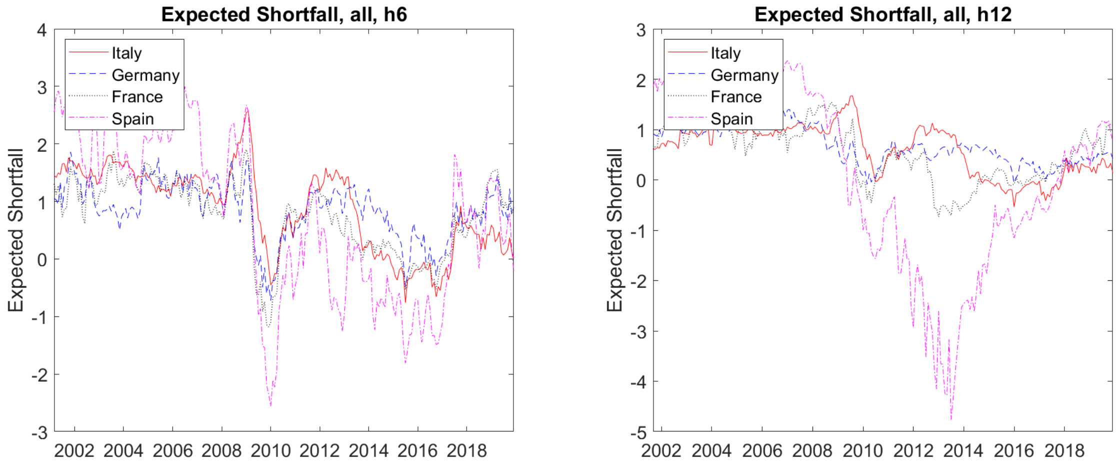

h6 h12

Figure A.11: Expected shortfall, comparison

While upside inflation risks in absolute terms, represented by expected longrise, seem to be quite correlated and symmetric across countries, the co-movement of downside inflation risks is affected by the extreme dynamics of Spain's expected shortfall, although the behaviour normalises towards the end of the series.

Ultimately, I show the impulse response functions and FEVDs of a SVAR model whose structural shocks are identified analogously to the exercise in chapter 6. The specification of the SVAR is the following:

$$\mathbf{y_t} = \begin{pmatrix} CBIshock \\ Eonia \\ Infl \\ Unemp \\ FinCond \\ FRE \end{pmatrix}$$

where "FinCond" stands for the SovCISS indicator of financial conditions and "FRE" is the acronym of full relative entropy.

The sample is 2000M9—2019M12, lag order is 3 and the SVAR is specified with a constant.

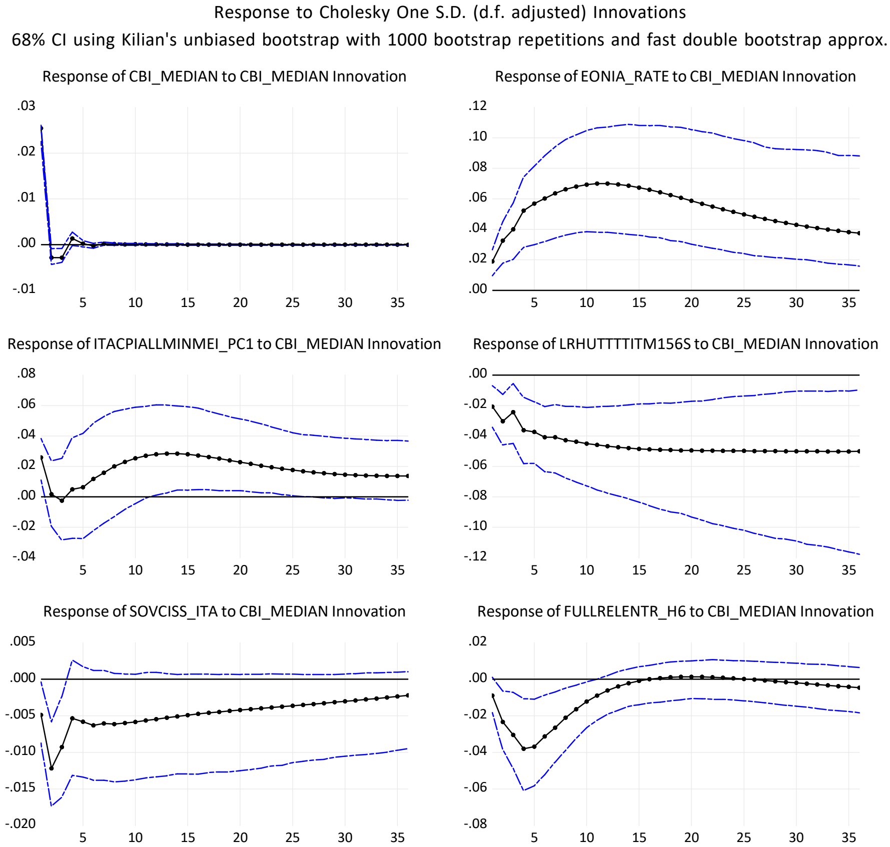

Figure A.12: IRFs Italy h6

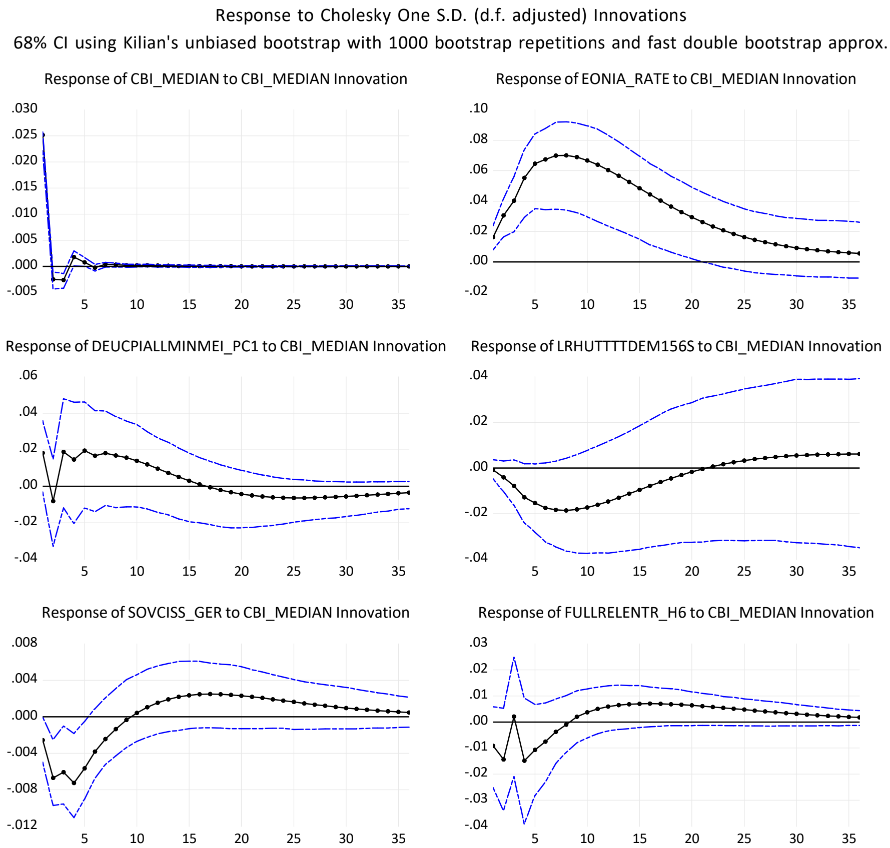

Figure A.13: IRFs Germany h6

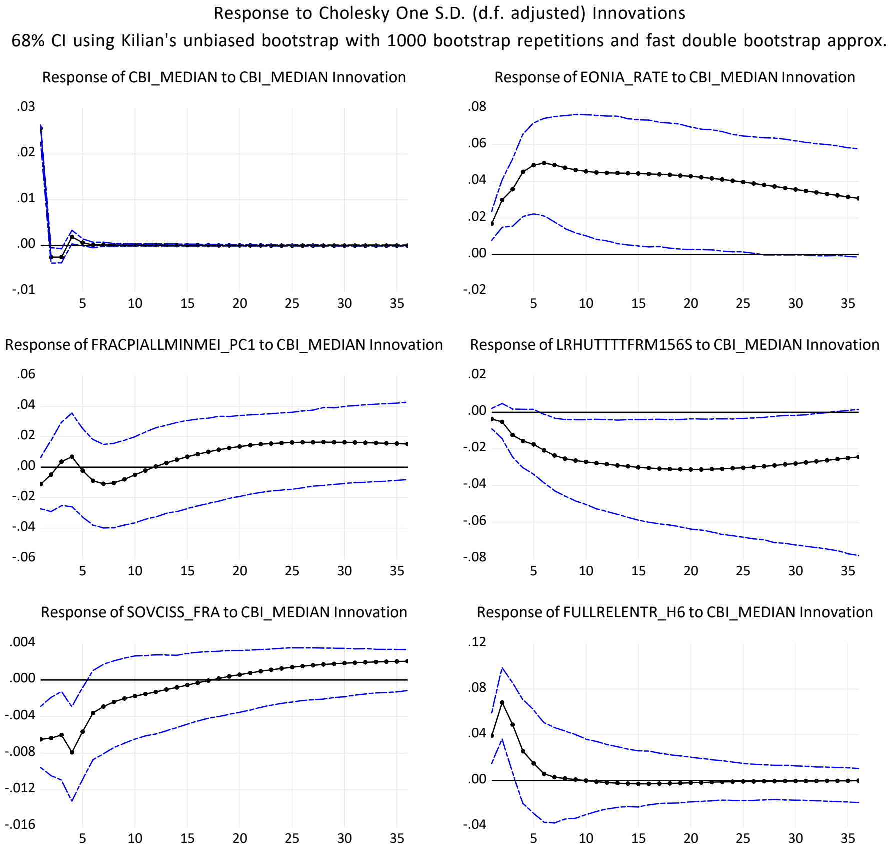

Figure A.14: IRFs France h6

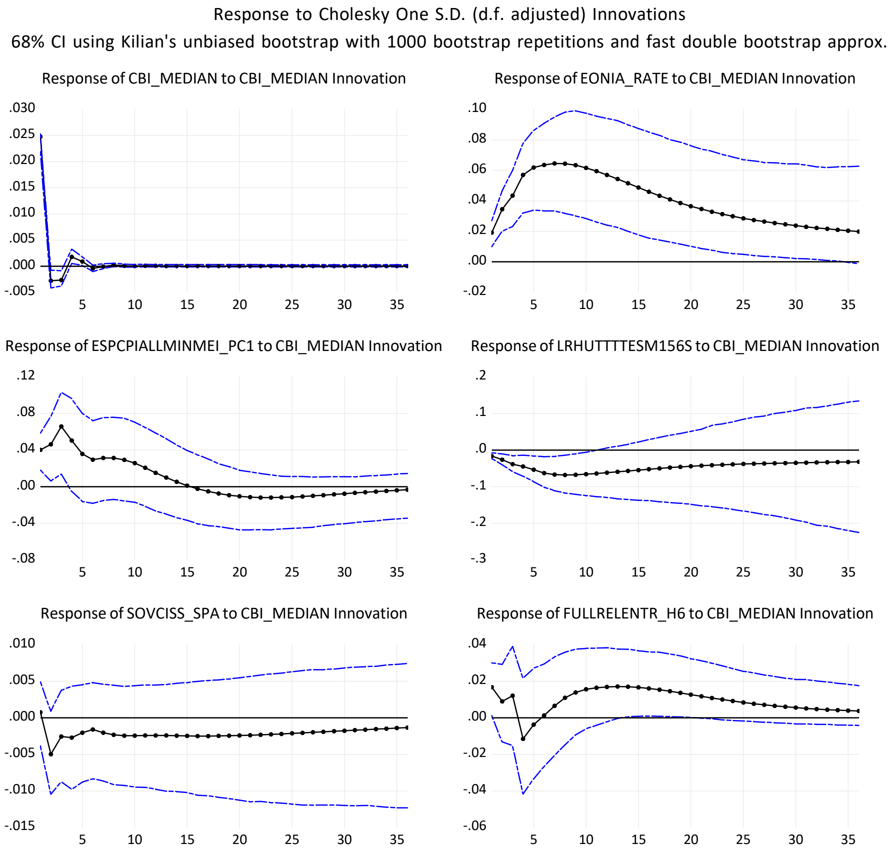

Figure A.15: IRFs Spain h6

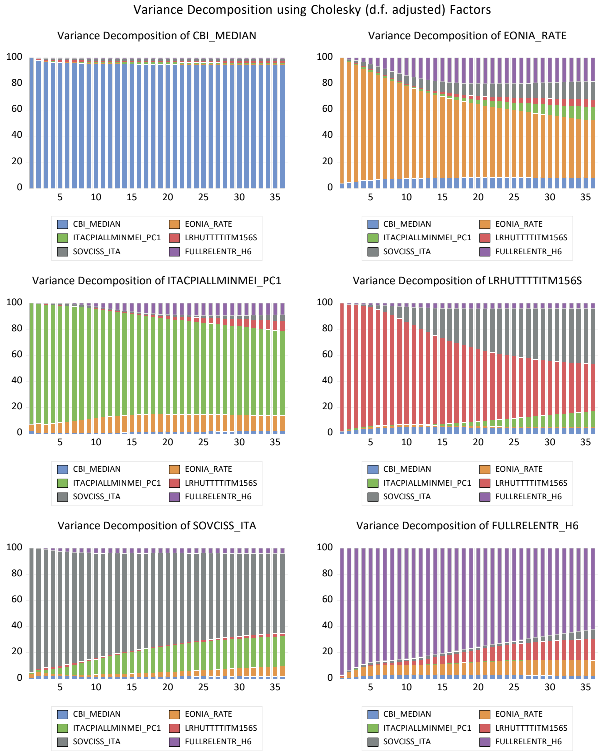

Figure A.16: FEVDs Italy h6

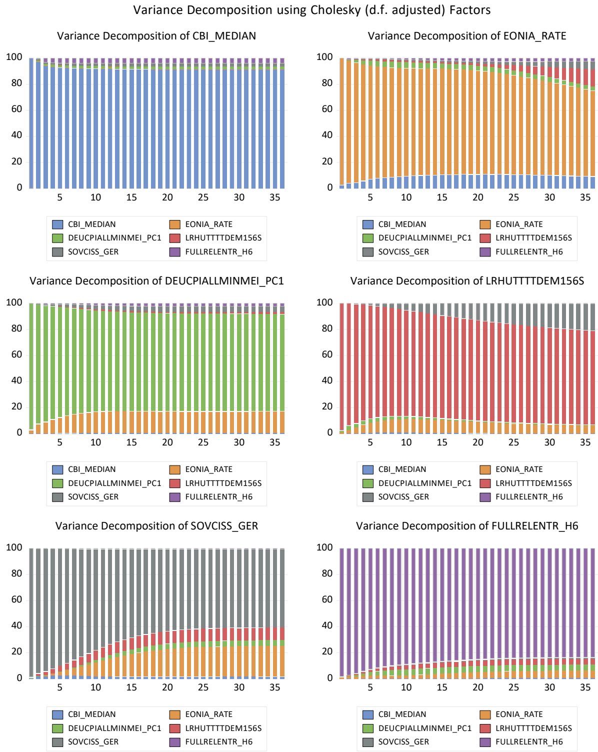

Figure A.17: FEVDs Germany h6

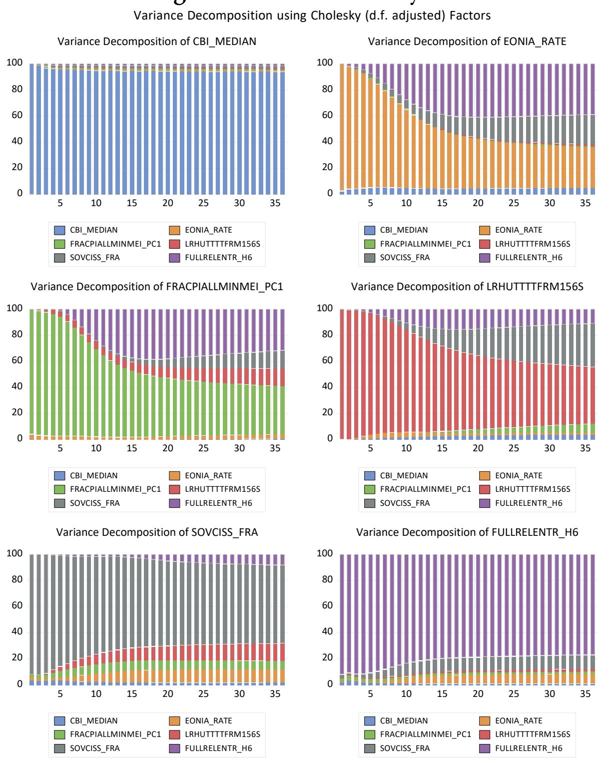

Figure A.18: FEVDs France h6

Figure A.19: FEVDs Spain h6

The effects of a Central Bank Information shock on the EONIA, inflation and unemployment rates are, qualitatively, unambiguously the same as in chapter 6 for all countries. Germany's inflation and unemployment rates are still "insulated" from shocks.

Most countries financial conditions' index respond beneficially to a CBI shock, with a decrease in the SovCISS index, although unremarkably in quantitative terms; Spain's financial conditions do not respond significantly.

When it comes to full relative entropy, the responses are, once again, heterogeneous: Spain and Germany do not respond significantly, France's FRE increases in the first periods while Italy's decreases.

The relevance of financial conditions in explaining relative entropy change across countries, while overall being not too powerful. Instead, the SovCISS indicator is similarly relevant as the unemployment rate in explaining the dynamics of inflation, while affecting importantly unemployment itself; hence I consider the financial conditions indicator a relevant predictor of inflation.

Conclusion This is further corroboration of the arguments presented in the body of the dissertation. The ECB faces a hard task in terms of achieving the necessary level of inflation risks' convergence, and the chances of reaching long-term price stability in the absence of important structural reforms aimed at achieving a higher degree of homogeneity in the fiscal policy, labour and financial markets, are probably low.

B

Appendix: Technicalities

B.1 Review of some concepts

B.1.1 Indicator function

The indicator function, denoted by $1_A(x)$, is a function defined on a set X that indicates membership of an element in a subset A of X. Formally, it is defined as:

$$1_A(x) = \begin{cases} 1 & \text{if } x \in A, \\ 0 & \text{if } x \notin A. \end{cases}$$

This means that if the element x is in the subset A, the function returns 1; otherwise, it returns 0.

B.2 Elements of Information Theory

Information theory is a field at the intersection of mathematics probability, statistics, computer science, statistical mechanics and electrical engineering, that studies the quantification, storage, and communication of information; in addition, information theory deals with randomness and uncertainty as well.

Its foundations were laid in the 1920's, but the foremost contribution to the field is arguably the one Shannon made in 1948, with the introduction of the concepts of self-information and entropy in 1948, in (Shannon, 1948).

B.2.1 Self-information

Self-information is a quantity derived directly from the probability; indeed it can be thought of as an alternative way of expressing it, that provides some mathematical advantages in some applications. The self-information (or Shannon information) can be interpreted as the quantification of the ”surprise” in the realisation of an outcome.

Shannon (1948)’s definition of self-information lies on three fundamental axioms:

- An event with probability 100 is perfectly unsurprising and yields no information.
- The less probable an event is, the more surprising it is and the more information it yields.
- If two independent events are measured separately, the total amount of information is the sum of the self-informations of the individual events.

Formally, the self-information of measuring X as its outcome x is defined as:

$$I_X(x) := -\log f_X(x) = \log \left(\frac{1}{p_X(x)} \right)$$

where f_X is the probability density (or mass) function of the random variable X and x is an outcome of X .

This concept is closely linked to entropy.

B.2.2 Entropy

The entropy of a random variable is the average level of information that we can infer from the realisation of a random trial, that is, a random sampling from the space of all the possible outcomes of the random variable.

The core idea of information theory is that the “informational value” of a communicated message depends on the degree to which the content of the message is surprising. If a highly likely event occurs, the message carries very little information. On the other hand, if a highly unlikely event occurs, the message is much more informative For instance, the fact that a die toss yields a 5, an event with 1/6 probability, conveys more “information” than tossing a coin and getting a 1, with 1/2 probability, because the first event is less likely than the second. Hence, the die toss has got higher entropy.

The entropy of a discrete random variable X, defined on the set χ and is distributed according to

$p : \chi \to [0, 1]$ such that $p(x) := \mathbb{P}[X = x]$, is:

$$H(X) = \mathbb{E}[I(X)] = \mathbb{E}[-\log p(X)]$$

where $\mathbb{E}$ is the expected value operator and I is information content of X. The self-information $I(X)$ is itself a random variable.

Entropy is made explicit, in discrete terms, as:

$$H(X) = -\sum_{x \in X} p(x) \log p(x)$$

One may also define the conditional entropy of two random variables X and Y as:

$$H(X|Y) = -\sum_{x,y \in \chi \times \mathcal{Y}} p_{X,Y}(x, y) \log \frac{p_{X,Y}(x, y)}{p_Y(y)}$$

where $p_{X,Y}(x, y) := \mathbb{P}[X = x, Y = y]$. This quantity should be understood as the remaining randomness in the random variable X given the random variable Y.

The entropy Shannon developed is restricted to discrete random variables. The corresponding formula for a continuous random variable with probability density function $f(x)$ with finite or infinite support $X \in \mathbb{R}$ is defined as differential entropy:

$$H(X) = E[-\log f(X)] = -\int_X f(x) \log f(x) dx$$

Additional notes Entropy can be formally defined as a measure. Among its many applications, it is ubiquitously used in communications and data compression, to calculate the smallest amount of information required to convey a message.

B.2.3 Relative Entropy

The relative entropy, of Kullback-Leibler divergence (Kullback and Leibler, 1951), is a type of statistical distance. More precisely, it is a divergence: a measure of how one probability distribution P is different from a second, reference probability distribution Q. The relative entropy of distribution P from Q can be interpreted as the expected excess surprise from using Q as a model instead of P, when the actual distribution is P.

For discrete random variables it is defined as:

$$D_{KL}(P \parallel Q) = -\sum_{x \in X} P(x) \log \left(\frac{Q(x)}{P(x)} \right)$$

while for continuous random variable it is defined as:

$$D_{KL}(P \parallel Q) = \int_{-\infty}^{\infty} p(x) \log \left(\frac{p(x)}{q(x)} \right) dx$$

where p and q denote the probability densities of P and Q

In other words, it is the expectation of the logarithmic difference between the probabilities P and Q, where the expectation is taken using the probabilities P.

Additional notes In applications, P typically represents the "true" distribution of data, observations, or a precisely calculated theoretical distribution, while Q typically represents a theory, model, description, or approximation of P. In order to find a distribution Q that is closest to P, we can minimize the KL divergence and compute an information projection.

Although it quantifies the divergence between two distributions and can be thought of as a "distance," the KL divergence is not a true metric. Unlike metrics, it is not symmetric with respect to the two distributions, and it does not satisfy the triangle inequality. While metrics are symmetric and generalize linear distance, satisfying the triangle inequality, divergences are asymmetric and generalize squared distance, in some cases satisfying a generalized Pythagorean theorem. In general, $D_{\mathrm{KL}}(P \parallel Q)$ does not equal $D_{\mathrm{KL}}(Q \parallel P)$.

References

Adams, P. A., Adrian, T., Boyarchenko, N., and Giannone, D. (2021). Forecasting macroeconomic risks. *International Journal of Forecasting*, 37(3):1173–1191.

Adrian, T., Boyarchenko, N., and Giannone, D. (2019). Vulnerable Growth. *American Economic Review*, 109(4):1263–1289.

Azzalini, A. (1985). A Class of Distributions Which Includes the Normal Ones. *Scandinavian Journal of Statistics*, 12(2):171–178. Publisher: [Board of the Foundation of the Scandinavian Journal of Statistics, Wiley].

Azzalini, A. and 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. arXiv:0911.2342 [stat].

Azzalini, A. and Dalla Valle, A. (1996). The multivariate skew-normal distribution. *Biometrika*, 83(4):715–726.

Banerjee, R., Contreras, J., Mehrotra, A., and Zampolli, F. (2024). Inflation at risk in advanced and emerging market economies. *Journal of International Money and Finance*, 142:103025.

BenSaïda, A. and Slim, S. (2016). Highly flexible distributions to fit multiple frequency financial returns. *Physica A: Statistical Mechanics and its Applications*, 442:203–213.

Blanchard, O., Cerutti, E., and Summers, L. (2015). Inflation and Activity – Two Explorations and their Monetary Policy Implications. Technical Report w21726, National Bureau of Economic Research, Cambridge, MA.

Castelnuovo, E. (2023). Uncertainty before and during COVID-19: A survey. *Journal of Economic Surveys*, 37(3):821–864.

Ciccarelli, M., Darracq Paries, M., Priftis, R., Angelini, E., Banbura, M., Bokan, N., Rigato, R. D., Fagan, G., Gumiel, J. E., Kase, H., Kornprobst, A., Ganoulis, I., Lalik, M., Montes-Galdón, C., Müller, G., Paredes, J., Santoro, S., Warne, A., Zimic, S., Brunotte, S., Cocchi, S., Magi, A., Invernizzi, M., and Von-Pine, E. (2024). ECB Macroeconometric Models for Forecasting and Policy Analysis. *SSRN Electronic Journal*.

Davis, C. (2015). The Skewed Generalized T Distribution Tree Package Vignette. *Open Access*.

European Central Bank. (2018). *Beyond spreads: measuring sovereign market stress in the euro area.* Publications Office, LU.

European Central Bank. (2019). *Measuring euro area monetary policy.* Publications Office, LU.

European Central Bank. (2020). *How to estimate a VAR after March 2020.* Publications Office, LU.

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

Greenspan, A. (2004). Risk and Uncertainty in Monetary Policy. *American Economic Review*, 94(2):33–40.

Holló, D., Kremer, M., and Duca, M. L. (2012). CISS - a composite indicator of systemic stress in the financial system. *ECB Working Papers Series*.

Jarociński, M. and Karadi, P. (2020). Deconstructing Monetary Policy Surprises— The Role of Information Shocks. *American Economic Journal: Macroeconomics*, 12(2):1–43.

Kilian, L. and Manganelli, S. (2007). Quantifying the Risk of Deflation. *Journal of Money, Credit and Banking*, 39(2-3):561–590.

Koenker, R. and Bassett, G. (1978). Regression Quantiles. *Econometrica*, 46(1):33.

Korobilis, D., Landau, B., Musso, A., and Phella, A. (2021). The Time-Varying Evolution of Inflation Risks. *SSRN Electronic Journal*.

Kullback, S. and Leibler, R. A. (1951). On Information and Sufficiency. *The Annals of Mathematical Statistics*, 22(1):79–86.

López-Salido, D. and Loria, F. (2024). Inflation at risk. *Journal of Monetary Economics*, page 103570.

Plagborg-Møller, M. and Wolf, C. K. (2021). Local Projections and VARs Estimate the Same Impulse Responses. *Econometrica*, 89(2):955–980.

Rossi, B. and Sekhposyan, T. (2019). Alternative tests for correct specification of conditional predictive densities. *Journal of Econometrics*, 208(2):638–657.

Schmidt, L. D. W. (2013). Quantile Spacings: A Simple Method for the Joint Estimation of Multiple Quantiles. *SSRN Electronic Journal*.

Shannon, C. E. (1948). A Mathematical Theory of Communication. *The Bell System Technical Journal*.

Theodossiou, P. (1998). Financial Data and the Skewed Generalized T Distribution. *Management Science*, 44(12):1650–1661. Publisher: INFORMS.

Acknowledgments

I want to thank my supervisor, Professor Efrem Castelnuovo.

Professor Castelnuovo has provided me with excellent supervision and timely, sensible feedback on my work. He has been very available in offering his assistance, even at short notice, when at times I was in a rush.

When I proposed him my idea to work on this project, that uses some techniques that go above and beyond a master's training, instead of rejecting my idea as too ambitious or "up in the air", he chose to accommodate my curiosity, to believe in my potential and to provide the expert guidance that enabled me to complete this project.

I am lucky and grateful to have been supervised by Professor Castelnuovo.