Essays on Volatility Networks and Uncertainty

Luca Rossi

TESI DOCTORAL UPF / ANY 2018

DIRECTORS DE LA TESI
Prof. Barbara Rossi i Prof. Christian Brownlees
Departament d'Economia i Empresa

A mi mujer
Inés

Acknowledgements

I would first like to especially thank my advisors Barbara Rossi and Christian Brownlees for their always precise and invaluable suggestions and comments they gave me. The dedication and passion about doing research is definitely one of the most important things I learned from them. I've always been very positively impressed by the time they always found to meet me and discuss my latest advances. I can but guess how busy they are with research, conferences, workshops, seminars, refereeing, advising other students' theses, and of course teaching. I never took your precious time for granted, so thank you. Thanks also to Majid Al-Sadoon and Geert Mesters for the many useful comments you gave me throughout the various stages of my research projects.

I dedicate this thesis to my wonderful wife, who unconditionally supported me throughout the hard task of pursuing a PhD since my very first year at UPF, where PhD courses and problem sets almost completely wiped out my spare time. To make matters worse, my less-than-A1 Spanish proficiency in the early stages of my stay in Barcelona didn't help our conversations, and hearing me uttering Spanish words was simply awkward. Still, if I can now say that I do speak Spanish is because of your patience in teaching me almost everything I know about it. Spanish besides, thanks for helping me becoming a better self, and for helping me every day in the truly fundamental things in life.

I thank my beautiful daughter Sofía. Thank you because your laughs and uncontainable joy are one of the most powerful and wondering things I've seen. Being interrupted by your laughter while I was working was the sweetest interruption, and smiling while working on a PhD thesis is priceless.

I owe too much to my parents, without your support I couldn't have enrolled in the BGSE Master. I still remember the day I had to decide whether to stay in Milan or going to Barcelona, and I quite had an intuition that I had to fly to Spain. You helped me make what turned out to be the best decision of my life.

Thank you, Inés and Fernando. I felt I was welcome in your family since the very first day we met. You helped Inés and me out in a great deal of situations, but I would especially thank you for taking care of Sofía while we had to work, so that a significant part of this PhD thesis has been made possible because of your support.

Israel and Tommi, I owe you many thanks. I really miss our coffees outside the library, and I wish we could see each other much more frequently. Yet, it was great to see you again after 2 and a half years. Chapter 1 of this thesis took this direction thanks to the endless stats discussions we've had during and after our year in Barcelona.

I've been very lucky to share office 20.142 with Christian, Yiru and Youel. It's been a true honor to be your office-mate and to share with you the good and bad moments of our PhDs. I wish you all the best and I'm sure you'll do great wherever you'll end up after grad school. Thanks also to all the members of the Econometrics seminars who gave me valuable pieces of advice.

Thanks to Alessio Anzuini and Pietro Tommasino for giving me the possibility to work with you on the paper which is Chapter 3 of this thesis. My summer internships at the Bank of Italy have been a precious opportunity to learn a lot and have a glimpse of how the central bank world works.

Thanks to UPF IT staff, particularly Marc Esteve and Raúl Cacho, who gave me prompt help about all the problems I had with remote computing. You set up the beautiful Marvin cluster right when I needed it. All the results in Chapter 1 have been made possible because of one big machine in that cluster. I am very grateful to Marta Araque and Laura Agustí too, who make the administrative requirements of the Economics PhD program much easier to understand and comply with.

A heartfelt thanks goes to Gianluca Femminis and Maurizio Motolese, who continuously provided me with good advice and answered the many questions I had regarding PhDs during my time at UCSC. Most importantly, I am enormously indebted with you because of the time you spent writing recommendation letters and giving me suggestions about how to write a good statement of purpose, both for PhD and scholarships applications.

Finally, very sincere gratitudes to Ferdinando Colombo, who taught me microeconomics and game theory from the very 101 course to the advanced ones. You've been able to go beyond the abstract formulas and show us the deep intuition hidden behind them, thus leading us to what I now believe is and should always be the true essence of microeconomics and economics in general, that is providing people with tools to reason about aspects of the real world in a smarter and non-trivial way. You first sparked my interest for economic theory, and time turned it into the firm decision to embark on doctoral studies. Thank you.

Abstract

This thesis empirically investigates different aspects of time-varying volatility. Chapter 1 estimates a large TVP-FAVAR and recovers a dynamic directed network of connections between European stock volatilities. We propose an ad-hoc estimation methodology that is shown to outperform both standard approaches and competing models. Chapter 2 focuses on tracking dynamic connectedness between US sectoral volatilities using Generalized Forecast Error Variance Decompositions with a Bayesian model. As opposed to estimates obtained with rolling windows, we allow parameters to vary in a more flexible way. We show that there exists a stable relationship between the network structure and the volatility regimes in place at a given time. Chapter 3 estimates the unexpected time-varying volatility component of fiscal budgets in Italy. We show that periods of higher unexpected fiscal volatility are likely to be recessionary. Expansionary policies are effective only when not accompanied by increases in uncertainty.

Resum

Aquesta tesi investiga empíricament diferents aspectes de la volatilitat variable. El Capítol 1 estima un TVP-FAVAR i recupera una xarxa de connexions dinàmiques entre les volatilitats de accions europees. Proposem una metodologia d'estimació ad-hoc que es demostri que supera els enfocaments estàndard i els models competidors. El Capítol 2 es centra en el seguiment de la connectivitat dinàmica entre les volatilitats sectorials dels Estats Units mitjançant descomposicions generalitzadas de variància d'errors de previsió amb un model Bayesià. A diferència de les estimacions obtingudes amb finestres enrotllables, permetem que els paràmetres variïn de manera més flexible. Mostrem que existeix una relació estable entre l'estructura de la xarxa i els règims de volatilitat vigents en un moment determinat. El Capítol 3 estima el component variable inesperat de la volatilitat dels pressupostos fiscals a Itàlia. Mostrem que els períodes de major volatilitat fiscal inesperada probablement són recessius. Les polítiques expansives només són efectives quan no s'acompanyen d'increments d'incertesa.

Preface

The common theme that links the chapters in this thesis is volatility. Interestingly, volatility has been the main ingredient in many recent theoretical and applied papers, and it has been so in a very diverse spectrum of economic and financial topics. Networks, granularity, connectedness, systemic risk, measures of uncertainty and its aggregate effects are (broadly speaking) among the strands of literature that have recently dealt with second moments as the main object of analysis. This thesis' chapters contribute to some of the aforementioned topics.

In Chapter 1 we estimate a time-varying directed network of volatilities using a large panel of European companies. We use a novel TVP-FAVAR estimation methodology that deals with the inherent complexity of the model and jointly improves both out-of-sample performance and the estimates of network connections. We propose a simple local validation algorithm to estimate tuning parameters in our penalized regressions. The algorithm consistently outperforms 10-fold Cross-Validation, and it runs up to 50 times faster. We replace uniform weighting of sample observations with forgetting factors, which maintain local flexibility of the parameters. Also, we show that shrinking the common factors coefficients within the TVP-FAVAR is key to improve out-of-sample performance. As opposed to standard TVP-FAVARs, our regressions outperform competing models and yield very large improvements. The resulting network structure is highly predictable at least in the very short term, and we find that the degree distributions strongly move together with the volatility regimes in place at a given time. Moreover, systematic volatility events activate additional connections on top of the ones that would be justified by the systematic episode itself. Individual nodes quickly move towards the middle of the network when hit by purely idiosyncratic events.

Chapter 2 studies sectoral volatility connectedness using US ETFs data. We estimate volatility spillovers as in Diebold and Yılmaz (2012), i.e. we use Generalized Forecast Error Variance Decompositions to construct the adjacency matrix of a weighted directed graph. We improve over their methodology by estimating the model with Koop and Korobilis (2013) TVP-VAR which also allows for Dynamic Model Switching in order to dynamically accommodate time-varying changes in fundamental model parameters. Our findings show that sectors are highly inter-connected and that two regimes split our sample between an earlier one of lower and volatile connectedness, and a latter one characterized by higher and more stable connections. Furthermore, highly volatile periods are associated with a change in network structure, which moves towards a star network, whereas quiet periods are characterized by a more asymmetric shape.

Chapter 3 is a joint paper with Alessio Anzuini and Pietro Tommasino. Our research departs from the previous chapters in my thesis and takes a structural approach by focusing on the macroeconomic effects of fiscal uncertainty. Indeed, economic uncertainty is an important factor behind macroeconomic fluctuations: in an uncertain environment, firms reduce hiring and investment, financial intermediaries are more reluctant to lend, and households increase their propensity to save. In our paper, we study the effects of the uncertainty which arises from fiscal policy decisions. To this end, we propose a new measure of fiscal policy uncertainty (FPU). In particular, we estimate a fiscal reaction function, allowing the volatility of the shocks to be time-varying. The time series of this volatility is our proxy for FPU. Looking at Italian data over the period 1991-2014, we find that an unexpected increase in our measure of FPU has a negative impact on the economy. One implication of this result is that the same change in the government budget can have different effects depending on whether it is associated with a reduction or an increase in FPU. Therefore, neglecting FPU may partly explain why the size (and sign) of fiscal multipliers differs so much across existing empirical studies.

Contents

List of Figures

List of Tables

Chapter 1

Evolving Networks and Volatility Forecasting: A Unified Approach

1.1 Introduction

In the last two decades, the literature on economic and financial networks burgeoned. One of the reasons is the recognition that aggregate fluctuations can often hardly be predicted by indistinctly aggregating individual ones, while they rather mainly depend on the complex network structure underlying the system under consideration, as well as its granularity.

In this paper, we introduce a new methodology to estimate time-varying directed networks from financial data, and we apply this method to stock volatilities of European companies. Crucially, our model is capable to improve at the same time *both* estimates of network connections and out-of-sample forecasting performance. We measure volatility as in Parkinson (1980), that is we exploit intra-day information with the daily high-low range. Compared to GARCH-like estimates, computing volatilities in this way leads to remarkably less persistent and more noisy estimates. This feature makes the dataset an interesting and challenging one to forecast, and a well suited field-test for our method.

We compute the dynamics of the network by estimating our model on rolling windows. Importantly, we add forgetting factors in the form of exponential weights to better pin down the local variation in parameters due to the most recent observations. In particular, we apply those weights to observations within a given window, thereby assigning increasingly higher importance to the most recent data. Standard rolling windows equally weight all the observations, thereby producing overly smoothed estimates, and this is especially true during periods where the network structure changes rapidly. Using forgetting factors addresses precisely this issue. Thanks to exponential weights and time-varying sparsity, the resulting model is elastic and flexible enough to better accommodate periods of sudden changes.

As recalled in Barigozzi and Brownlees (2017), previous literature shows that volatility has a factor structure, meaning that we have to control for systematic volatility components. We choose to augment each forecasting equation with common factors, thereby estimating a (TVP) FAVAR and correctly recovering parameters that are free from systematic components. Furthermore, we illustrate that shrinking the common factors coefficients together with the rest of predictors is necessary to improve forecasting performance, and corroborates the already established result that even though principal components maximize in-sample variation, they might not be equally helpful in out-of-sample forecasting.

Finally, we illustrate that standard 10-fold Cross-Validation procedures[1] fail when estimating the coefficients of our penalized regressions. We thus propose a local validation algorithm where the choice of the tuning parameter at each point in time depends on past forecasting performance. We show that locally validating the penalization parameter consistently outperforms 10-fold Cross Validation. Furthermore, thanks to the fact that the model has to be estimated only once (that is we do not need to repeatedly split the sample in multiple folds), results are available up to 50 times faster.

All our new features generate very large out-of-sample improvements. Indeed, estimating the TVP-FAVAR with uniform weights (i.e. the standard rolling window approach that does not use forgetting factors), using 10-fold Cross-Validation, and not shrinking common factors yields a particularly disappointing out-of-sample performance. Specifically, this model cannot beat standard competitors, and it is only applying our methodology that we are able to make our model the best performing one.

As for the network representation, we obtain interesting results. We first illustrate some facts about the resulting network structure, e.g. that connections are very predictable at least in the immediate one-day horizon, and that both the in-degree and the out-degree distributions are strongly regime-dependent. Indeed, we find that our model is able to change rapidly after key events occur, and that high systematic volatility episodes are often accompanied by additional non-systematic connections that might have not been activated had that event not occurred. More specifically, both degree distributions tend to spread out significantly during financial turmoils, even though they do so in a very different way. Indeed, whereas a fairly large portion of companies in our dataset are likely

[1]See Stone (1974).

to be heavily influenced by many others during highly volatile periods, only an extremely limited set of companies are the main responsible for the aforementioned firms to be significantly affected during volatility storms. This means that the network has a granular structure *especially* during stressed periods, and that close attention has to be paid to the few institutions that affect the highest amount of companies. Importantly, those very central companies are not the same throughout the sample period, i.e. their centrality rank vary significantly over time. For instance, we find that important episodes that are entirely idiosyncratic to individual nodes lead the corresponding companies to quickly become very central in the web of connections. Section 1.4.1 explains those results with more detail.

Apart from the methodological contribution[2] (which is a fundamental part of this paper), our research is mainly related to the recent literature that aims at empirically estimating financial networks. Diebold and Yılmaz (2009) measures returns and volatility spillovers in global equity markets with Generalized Forecast Error Variance Decompositions. The authors show that connectedness in returns exhibits an increasing trend over time (market integration) whereas volatility spillovers vary significantly more and display no trend. Later papers[3] apply a similar methodology to different economic and financial environments. Our setting borrows the rolling windows approach to compute the evolution of regression parameters, but we greatly improve their estimation strategy by adding forgetting factors as we explained above. Koop and Korobilis (2013) uses them in the Bayesian tradition, i.e. by updating forgetting factor estimates at each Kalman filter iteration. Our framework is purely frequentist, but this does not prevent one from using forgetting factors, which are easily implementable in the class of penalized regression models we estimate throughout our work. Section 1.3.2 explains the differences between Koop and Korobilis (2013) forgetting factors and ours in detail. Moreover, we also depart from Diebold and Yılmaz approach by validating the penalization parameter with our own algorithm, which (as highlighted above) is shown to be a superior alternative to k-fold Cross-Validation in time-varying-parameters settings like these. Billio, Getmansky, Lo and Pelizzon (2012) estimates Granger-causality networks on hedge funds, banks, broker/dealers, and insurance companies and, among other results, shows that the banking sector is crucial in transmitting shocks to other institutions. Brownlees, Nualart

[2]We refer to Barigozzi and Brownlees (2017) for a list of references in the research area concerning estimation of sparse VARs.

[3]See Diebold and Yılmaz (2012), Diebold and Yılmaz (2013), Diebold and Yılmaz (2014), Diebold and Yılmaz (2015), Demirer, Diebold, Liu and Yılmaz (2017).

and Sun (2015) uses Lasso algorithms to regularize inverse realized covariance estimators of log-prices, thus uncovering the partial correlation network structure of returns. Abbassi, Brownlees, Hans and Podlich (2016) uses both a market-based and a proprietary dataset to develop measures of connectedness, and show that market-based ones work comparatively well as a tool to monitor banks credit riskiness. Barigozzi and Brownlees (2017) introduces a new Lasso-based algorithm called NETS to analyze large panels of volatility measures and estimate their inter-connectedness, producing both a directed and an undirected static graph. Concerning the assignment of connections between companies, we strictly follow their approach and attach directed links whenever a certain company's volatility Granger-causes others. The methodology used in Barigozzi and Brownlees (2017) is shown to outperform several forecasting models. The authors control for common factors by first de-factorizing each series and estimating the model on the residuals thereafter. By controlling for systematic volatility within the model we are able to estimate it on the original non-defactorized dataset.

1.2 Data

We estimate our model on the panel of stock volatilities computed from the constituents of the FTSE Developed Europe Index. This index blends the 525 largest companies listed in 13 European countries, and it does so using market capitalization weights. Interestingly, some of the most important non Eurozone/non EU countries are also included. The complete list indeed comprises Belgium, Denmark, Finland, France, Germany, Italy, Netherlands, Norway, Portugal, Spain, Sweden, Switzerland, United Kingdom. Table 1.A.1 in the Appendix lists the companies included in the index as of May 20^{th}, 2016. As we will explain in the next paragraphs, we will have to drop some of them, so we only report the ones we use in our analysis.

Our method requires the panel to be balanced. However, the natural existence of disparate first trading days among companies lead to an inevitable trade-off between retaining more time periods but fewer companies, or the opposite. Nevertheless, we want our dataset to include a relevant portion of the most recent history. Finally, we want to avoid our results to be significantly influenced by missing observations, and we find a few companies with a relatively high amount of missing data. For reasons that will become clear shortly on, we treat any volatility that takes a value equal to zero as if it was a missing observation.

We prefer not to arbitrarily choose the number of companies and the sample length of the dataset we will use. Therefore, we do this automatically by searching for the panel that maximizes the number of data-points $N \cdot T$ subject to two constraints whose boundaries we deliberately choose. First, the final panel must start no later than January 1^{st}, 2003, and must finish no earlier than November 11^{th}, 2016. Second, every company in the panel must have no more than 5% of missing observations, where the percentage is computed over the given sample length (which changes throughout the panel search).

This optimization procedure produces the final dataset, which comprises $N = 322$ companies and $T = 3905$ time periods, with the sample starting on April 16^{th}, 2002, and ending on April 3^{rd}, 2017. The total number of missing observations is low and equal to 3.1% of all data-points, meaning that parameters estimates are not sistematically affected by our approximations which assign previous day volatility values to missing data.

In order to show in a succinct way the spectrum of companies in our panel, we use the Global Industry Classification Standard (GICS) to classify the sector and industry to which every company belongs to. Table 1.1 counts companies in every sector and industry. As we can see, three out of eleven sectors make up the majority of the companies in the sample. Those sectors are Financials, Industrials, and Consumer Discretionary. Table 1.2 shows that a similar reasoning applies for three (out of thirteen) countries, namely the United Kingdom, France, and Germany. The table also displays more detailed counts on companies belonging to a given sector *and* country. In general, having such a large sample ensures that we have a representative picture of the largest listed companies within the European countries we consider.

Table 1.1: Dataset summary: sectors and industries breakdown

Sector	Industry	#		%
Financials	Banks	29	**68**	**21.12%**
	Diversified Financials	17		
	Insurance	22		
Industrials	Capital Goods	43	**65**	**20.19%**
	Commercial&Professional Services	13		
	Transportation	9		
Consumer Discret	Automobiles&Components	12	**51**	**15.84%**
	Consumer Durables&Apparel	17		
	Consumer Services	5		
	Media	11		
	Retailing	6		
Consumer Staples	Food&Staples Retailing	9	**31**	**9.63%**
	Food Beverage&Tobacco	16		
	Household&Personal Products	6		
Materials	Materials	25	**25**	**7.76%**
Health Care	Health Care Equipment&Services	9	**22**	**6.83%**
	Pharm Biotech&Life Sciences	13		
Utilities	Utilities	15	**15**	**4.66%**
IT	Software&Services	6	**13**	**4.04%**
	Technology Hardware&Equipment	4		
	Semiconductors&Semic Equipment	3		
Energy	Energy	12	**12**	**3.73%**
Telecom Services	Telecommunication Services	12	**12**	**3.73%**
Real Estate	Real Estate	8	**8**	**2.48%**

Sectors are ordered by decreasing number of represented companies. Industry classification also follows GICS standards, and we highlight it here for exposition purposes only, whereas in the paper we will generally classify companies with the sector they belong to. The last column shows the percentage of companies which belong to a particular sector, as percentage of the total number of companies in the sample.

Table 1.2: Dataset summary: joint countries and sectors breakdown

	GBR	FRA	DEU	SWE	CHE	ESP	ITA	NLD	DNK	BEL	FIN	NOR	PRT	**All**	$\%_S$
Financials	15	10	6	7	6	6	8	2	1	4	1	1	1	**68**	21.12%
Industrials	17	12	8	9	6	4	0	4	4	0	1	0	0	**65**	20.19%
Consumer Discret.	17	16	7	3	2	1	3	0	0	0	1	1	0	**51**	15.84%
Consumer Staples	10	6	4	2	2	0	0	1	1	2	1	1	1	**31**	9.63%
Materials	5	2	6	1	2	1	0	2	1	2	2	1	0	**25**	7.76%
Health Care	4	4	3	1	3	0	1	0	5	1	0	0	0	**22**	6.83%
Utilities	4	2	2	0	0	4	1	0	0	0	1	0	1	**15**	4.66%
IT	1	5	3	2	0	0	0	1	0	0	1	0	0	**13**	4.04%
Energy	3	2	0	1	0	2	2	1	0	0	0	1	0	**12**	3.73%
Telecomm. Services	2	1	1	2	0	1	2	1	0	0	1	1	0	**12**	3.73%
Real Estate	4	2	0	0	1	0	0	1	0	0	0	0	0	**8**	2.48%
All	**82**	**62**	**40**	**28**	**22**	**19**	**17**	**13**	**12**	**9**	**9**	**6**	**3**	**322**	100%
$\%_C$	25.46%	19.25%	12.42%	8.70%	6.83%	5.90%	5.28%	4.04%	3.73%	2.80%	2.80%	1.87%	0.93%	100%	

This table lists both sectors and countries by decreasing numbers of companies which belong to the corresponding category. Approximately one third of the companies in the sample are British, French, and German companies in Financials, Industrials, and Consumer Discretionary sectors. The last row and the last column show the percentage of companies which belong to a particular country and sector respectively, as percentage of the total number of companies in the sample.

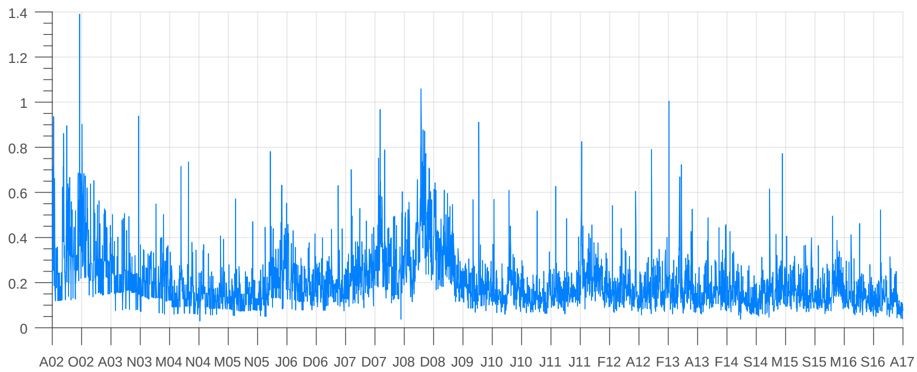

Figure 1.1: This picture clearly shows the noisy behavior of volatilities estimated with the high-low range. This illustrative example plots Swedish Match volatility, but it is representative of all other companies' series.

We forecast stock volatilities, here computed as in Parkinson (1980) exploiting intra-day information:

$$\sigma_{it}^2 = \frac{(h_{it} - l_{it})^2}{4 \log 2}, \tag{1.1}$$

where h_{it} (l_{it}) is the log of the high (low) price for company i at period t. Computing volatilities as in equation (1.1) ensures to detect high volatility episodes that resolve within the same day. For example, stock i could close at €100 at day $t-1$, spike up at €105 *during* day t, and close at €101 at day t. A GARCH model would simply use information stemming from the two closing days, i.e. it would use a €1 return, thus missing the large additional €4 spike, something that Pearson's method easily captures. Finally, the fact that this method is model-free makes the competition a fairer one. Indeed, it would not be surprising to find that a GARCH model would outperform competitors in forecasting volatilities estimated themselves with the same GARCH model. Figure 1.1 plots Swedish Match volatility as an example. Volatility series for the other companies in our sample are qualitatively very similar, but the reason why we choose to plot this particular company's volatility is because we will refer to it again in Section 1.4.1 when describing a peculiar idiosyncratic event that precisely occurred to this company.

Alizadeh, Brandt and Diebold (2002) shows that log-range volatility estimates are approximately normal. Even though we do not need normality for any of our results to hold, the dependent variable we use is $\log \sigma_{it}^2$. In this way we adapt the data to the linear model we use in order to better pin down regression parameters and consequently connections between companies. Using σ_{it}^2 as dependent variable could distort our results because of its highly skewed log-normal distribution.

1.3 Model

1.3.1 General Setting

We assume the non-common movements in (log) volatilities to be governed by a sparse VAR model. Barigozzi and Brownlees (2017) cites recent research which provides evidence for stock volatilities to have a factor structure[4]. However, our ultimate goal is to estimate a directed network of volatility connections that are free from influences of common components. If we were to estimate a simple VAR on this dataset, we would not be able to discern whether a given estimated link would be the result of common factors co-movements, pairwise relationships, or both. A possible venue that could overcome this issue would be to forecast the de-factorized series, i.e. to estimate the model on the residuals obtained from a static factor regression.

The problem with this approach is that de-factorization leads to a loss of information, and it implicitly assumes that factors coefficients do not suffer from over-fitting problems. However, we will later show that controlling for common factors with an out-of-sample perspective greatly improves forecasting performance as compared to merely controlling for them in-sample. This implies that by simply de-factorizing the series we are likely to overstate the importance of common factors. Moreover, forecast errors of a de-factorized series are hardly interpretable from an economic point of view.

We solve those issues by augmenting each forecasting equation (where the dependent variables are the non de-factorized series) with principal components. Doing this ensures that any estimated connection between companies would only stem from a pairwise relationship, because co-movements induced by common factors would be captured by their corresponding coefficients.

Our framework is related to Barigozzi and Brownlees (2017) which addresses the factor structure issue by de-factorizing the series with observed factors, whereas we maintain the advantages of forecasting the original dataset *and* obtaining a sparse network representation that is not influenced by common factors.

Similarly from Demirer et al. (2017) we recursively estimate elastic nets over moving windows, but we use an entirely different estimation strategy which is shown to perform best when compared to standard approaches. Elastic nets are a general class of penalized regression models first developed in Zou and Hastie (2005). The peculiar feature of those models is that they embed both Ridge and Lasso regressions. In particular, the researcher has to define an elastic net

[4]See references in the cited paper.

parameter $\alpha \in [0,1]$ (see equations (1.6)-(1.7) below) which governs the equilibrium between ℓ_2 (Ridge) and ℓ_1 (Lasso) regularization. While Ridge regression shrinks parameters towards zero but assigns a non-zero coefficient to all the predictors, Lasso shrinks parameters and potentially sets some of them to be exactly equal to zero, thus producing a more interpretable model which only retains the most relevant regressors.

Throughout our analysis we use two different window widths. The first, which we call W, is the one we use for estimating coefficients in our FAVAR. The second, which we denote by S, is a smoothing window which averages the most recent out-of-sample performances of each one of the different tuning parameters in our elastic net models. This averaging procedure is needed only when we use our own Local Validation (LV) scheme, which we will explain in detail in Section 1.3.4. Standard Cross-Validation (CV) procedures have been originally thought for cross-sectional settings and do not therefore average past periods performances. Therefore, let's define $P \equiv W + \mathbb{1}_{\text{LV}} \cdot S$ (where the indicator function $\mathbb{1}_{\text{LV}}$ takes value equal to one for models estimated with our LV scheme) and consider an N-variable TVP-FAVAR(p):

$$\mathbf{Y}_t = \boldsymbol{\mu}_t + \boldsymbol{\Gamma}_t \mathbf{F}_t + \sum_{m=1}^{p} \boldsymbol{\Phi}_{mt} \mathbf{Y}_{t-m} + \boldsymbol{\epsilon}_t, \quad t = P, \ldots, T \tag{1.2}$$

where $\mathbf{Y}_t$ is $N \times 1$, $\boldsymbol{\Gamma}_t$ is $N \times r$, $\mathbf{F}_t$ is $r \times 1$, $\boldsymbol{\Phi}_{mt}$ is $N \times N$, and we do not impose any distribution on $\boldsymbol{\epsilon}_t$. Equation (1.2) represents the system of equations we want to estimate. Static factors are retrieved *at each window* s from the following standard equation:

$$\begin{aligned} \bar{\mathbf{Y}}_d = \boldsymbol{\Lambda}_s \mathbf{F}_d + \boldsymbol{\nu}_d, \quad & d = s - (W - p) + 1, \ldots, s, \\ & s = W, \ldots, T, \end{aligned} \tag{1.3}$$

where $\bar{\mathbf{Y}}_d$ is equal to centered (but not standardized) $\mathbf{Y}_d$, $\boldsymbol{\Lambda}_s$ is the $N \times r$ matrix of window-specific (i.e. time-varying) factor loadings. We forecast the common factors with a VAR(1), using the same observation weights w_d as in the elastic net regression[5]. Therefore, we regress

$$\begin{aligned} \mathbf{F^*}_e = \mathbf{A}_s \mathbf{F^*}_{e-1} + \boldsymbol{\eta}_e, \quad & e = s - (W - p) + 2, \ldots, s, \\ & s = W, \ldots, T, \end{aligned} \tag{1.4}$$

where $\mathbf{F^*}_e = \sqrt{w_{e-s+W-p}}\, \mathbf{F}_e$, and $\mathbf{A}_s$ is $r \times r$. Now, define

[5]See equations (1.6)-(1.8) later on in this section.

- μ_{si} the i-th element of $\boldsymbol{\mu}_s$ scalar
- $\mathbf{B}_s = [\boldsymbol{\Gamma}_s \quad \boldsymbol{\Phi}_{1s} \dots \boldsymbol{\Phi}_{ps}]$ $N \times (r + p \cdot N)$
- $\boldsymbol{\beta}_{si}$ the i-th column of $\mathbf{B}'_s$ $(r + p \cdot N) \times 1$
- y_{si} the i-th element of $\mathbf{Y}_s$ scalar
- $\mathbf{X}_s = \left[\mathbf{F}'_s \quad \mathbf{Y}'_{s-1} \dots \mathbf{Y}'_{s-p}\right]'$ $(r + p \cdot N) \times 1$

For each window $s = P, \dots, T$ we then have

$$\mathbf{Y}_d = \boldsymbol{\mu}_s + \mathbf{B}_s \mathbf{X}_d + \boldsymbol{\epsilon}_d, \quad d = s - W + p + 1, \dots, s, \tag{1.5}$$

and we estimate each of the equations in the system (1.5) with elastic nets, allowing shrinkage to be company-specific. The reason for this last additional complexity is intuitive. Indeed, it is very likely that in a given time period some company is not experiencing any particular troubles, being therefore mainly affected by common factors and auto-regressive components. This means that its corresponding forecasting equation would need a relatively high parameters constraint. On the other hand, there might well exist other companies that are facing volatility storms, and their fragility might imply that their own volatility is not just influenced by common and own components, but also by a handful of other companies, which means those nodes would need a lower penalization parameter. As a consequence, a unique tuning parameter would likely lead to too high penalization for some companies and too low for others, undermining both forecasting ability and the actual reliability of the resulting network representation.

We define s and i to be time and company indexes respectively, whereas $l \in \{L, \dots, 1\}$ indexes decreasing values of the penalization parameter λ, and L is the grid size for λ[6]. Therefore, for each $s \in \{W, \dots, T\}$, for each $i \in \{1, \dots, N\}$, and for each $l \in \{L, \dots, 1\}$ we estimate

$$(\hat{\mu}_{sil}, \hat{\boldsymbol{\beta}}_{sil}) = \underset{(\mu_{sil}, \boldsymbol{\beta}_{sil})}{\operatorname{argmin}} \left[\frac{1}{2} \sum_{\substack{d=s-W\\+p+1}}^{s} w_d (y_{di} - \mu_{sil} - \boldsymbol{\beta}'_{sil} \mathbf{X}_d)^2 + \lambda_{sil} P_\alpha(\boldsymbol{\beta}_{sil}) \right], \tag{1.6}$$

where

$$P_\alpha(\boldsymbol{\beta}_{sil}) = (1-\alpha)\frac{1}{2}||\boldsymbol{\beta_{sil}}||^2_{\ell_2} + \alpha||\boldsymbol{\beta}_{sil}||_{\ell_1} = \sum_{j=1}^{r+p\cdot N} \left[\frac{1}{2}(1-\alpha)\beta^2_{silj} + \alpha|\beta_{silj}| \right], \tag{1.7}$$

[6]We follow Friedman, Hastie and Tibshirani (2010) in computing the grid of values.

$$w_d = \frac{\gamma^{s-d+1}}{\sum_{g=1}^{W-p} \gamma^g}, \tag{1.8}$$

and where β_{silj} is the j-th element of $\boldsymbol{\beta}_{sil}$. Importantly, $\gamma \in (0, 1]$ is the value we use for the forgetting factor, which results in the exponential weights computed as in equation (1.8). The exact functioning and choice of forgetting factors is explained in detail in the next section.

Our baseline model is a Lasso, which means that we should set $\alpha = 1$. However, Friedman et al. (2010) warns that in the presence of two highly correlated and relevant regressors the pure Lasso case tends to choose between one of the two, i.e. it sets (without loss of generality) a zero coefficient on the first predictor, while leaving the non-zero estimate only on the second one. This is undesirable, because we ultimately want to obtain reliable estimates of volatility connections between the companies we consider, and this behavior could potentially mute some important links. Luckily, the authors state that one can side-step this issue by estimating an elastic net with Lasso weight $\alpha = 1 - \varepsilon$, where ε is small, so we choose $\varepsilon = 0.01$. We then minimize the objective function in each iteration with the fast coordinate descent algorithm developed in Friedman et al. (2010). Finally, the network is computed by assigning a directed edge from variable (i.e. company) i to variable j whenever the first Granger-causes the second.

1.3.2 Forgetting Factors: an Introduction

As anticipated, we introduce a new feature, that is we weight observations with exponential decay as in equation (1.8). The reason why we do so is that equal weights lead to overly rigid estimates that react too slowly to news. Forgetting factors solve this problem by giving exponentially higher weight to the most recent observations.

For example, if window width $W = 400$ and $\gamma = 0.98$ (as in our baseline case), then observations 1-100, 101-200, 201-300, 301-350, 351-400 have a cumulative weight of 0.2%, 1.5%, 11.5%, 23.16% and 63.6% respectively. Similarly, approximately the last 100 observations receive a weight which is higher than the one they would obtain with equal weights (which would equal 0.25% in this case). Moreover, the most recent observation receives a weight of approximately 2%, which shows that thanks to forgetting factors we can use those relatively long windows (i.e. more information) without compromising the flexibility we need in our parameter estimates. Standard rolling windows would give the same weight to all time periods.

Weighting observations in a non-uniform way also implies that some information

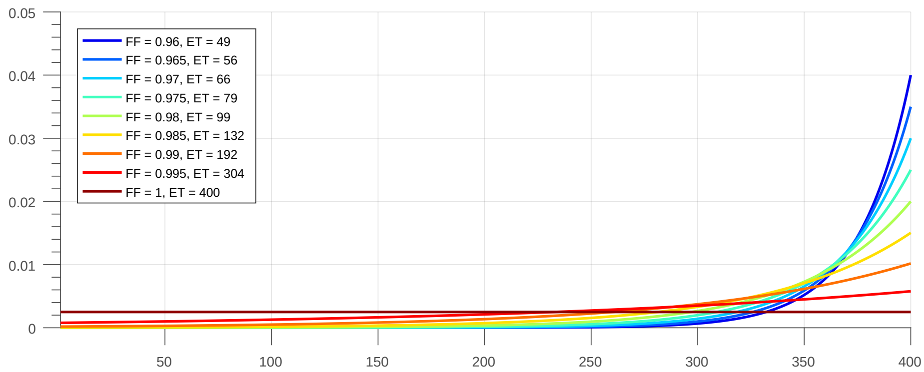

Figure 1.2: Different values for the forgetting factor. Note that seemingly small variations in the forgetting factor value lead to relevant changes in the amount (and type) of information used. In the extreme case we present, when setting the forgetting factor to 0.96, even if we are using a window length of 400 periods it is as if we were using only the informational content of 49 observations.

is neglected. More specifically, it is as if we were using less observations. Kish (1995) provides a formula to compute the effective sample size:

$$ET = \frac{\left(\sum_{t=1}^{T} w_t\right)^2}{\sum_{t=1}^{T} w_t^2}. \tag{1.9}$$

Figure 1.2 plots weights for different levels of the forgetting factor, and displays effective sample sizes for each value of γ. We can see that even if we are using a window width of 400 days, the effective number of observations we use is equal to 99. Note that when $\gamma = 1$ we obtain the standard uniform weights case.

The idea of using forgetting factors comes from Koop and Korobilis (2013), even though the way we use them is slightly different from their. Koop and Korobilis (2013) makes explicit modeling assumptions, that is they assume every parameter in their TVP-VAR to evolve as a random walk:

$$\boldsymbol{\beta}_{t+1} = \boldsymbol{\beta}_t + \boldsymbol{\eta}_t, \quad \boldsymbol{\eta}_t \overset{iid}{\sim} \mathcal{N}(\mathbf{0}, \boldsymbol{\Omega}_t), \tag{1.10}$$

where $\boldsymbol{\beta}_t$ is the vector containing all the coefficients in the TVP-VAR at time t, and it has dimension $N(1+pN) \times 1$. In their model, the forgetting factor equally shrinks each value in a pre-estimator of $\boldsymbol{\Omega}_t$ towards zero through the coefficient $\left(\frac{1}{\gamma} - 1\right)$, which is less than one for reasonable values of γ. The authors dynamically estimate γ using the grid $\{0.97, 0.98, 0.99, 1\}$. As a consequence, a low value of γ leads the model to forget faster, and the converse is true for high γs. In the limit, when $\gamma = 1$ the matrix $\boldsymbol{\Omega}_t$ has zeros everywhere, and all the coefficients become fixed. Therefore, the time-varyingness of their forgetting factor serves to increase/decrease the instantaneous flexibility of their parameters.

For practical purposes, our forgetting factor is very similar to the one used in Koop and Korobilis (2013), because we also have that past observations receive a non-linearly lower weight in the estimate of each parameter[7]. However, we do not make any assumption about parameters evolution, and our forgetting factor does not control the underlying covariance matrix $\mathbf{\Omega}_t$. Its task is simply to always weight less observations far away from current ones. Even when γ is fixed during the whole sample period (which is our baseline choice) data can decide to let coefficients change more during certain periods and less during others. Importantly, when $\gamma = 1$ (uniform weights) parameters do not stay fixed, even though data could estimate negligible changes anyway.

1.3.3 Adaptive Forgetting Factors

In principle, it is possible to estimate γ at each time period from a grid of values. We first do this and estimate the evolution of γ with the same local validation algorithm we describe in Section 1.3.4.

Table 1.3 displays in-sample and out-of-sample results for both the adaptive and the fixed forgetting factors cases. The first column is an average of the in-sample MSEs, where the average is taken over both companies and time periods. Similar considerations (with obviously different interpretations) are true for the remaining columns. Out-of-sample R^2 is the percentage improvement over the out-of-sample MSE of a benchmark model, which in our case is a 400-days-long time-varying mean. For ease of exposition, from here on we will refer to out-of-sample MSE simply as MSE, whereas the in-sample counterpart will still be fully specified. Notably, for all the sectors and individually for 84% of the companies in the sample, the best γ is between 0.975 and 0.985. Moreover, one can see that performance variation changes smoothly with γ. Those results illustrate that better out-of-sample performance of models estimated with $\gamma = 0.975$, $\gamma = 0.98$, or $\gamma = 0.985$ is not just a random average coincidence, but it is consistently present for the great majority of the companies. For long enough sample sizes which include significant periods of both calm and storm, we then suggest to use $\gamma = 0.98$ as the preferred *fixed* forgetting factor choice.

[7]Whereas in Koop and Korobilis (2013) observations j periods in the past *exactly* receive a weight equal to γ^j, in our framework this is not necessarily true. Indeed, given that we use rolling windows of length $W - p$, all the observations more than $W - p$ periods away get truncated and receive zero weight. Therefore, in our case we can say that observations j periods in the past receive a weight which is *proportional* to γ^j. This is clear by looking at equation (1.8).

Table 1.3: Adaptive versus fixed forgetting factors

	In-sample MSE	Out-of-sample MSE	In-sample R^2	Out-of-sample R^2
Mean	1.0729	1.0788	0%	0%
Lasso($\hat{\text{FF}}$)	**0.5258**	**0.7069**	**50.65%**	**33.96%**
Lasso(0.96)	0.5794	0.7519	45.28%	29.56%
Lasso(0.965)	0.5635	0.7363	46.87%	31.08%
Lasso(0.97)	0.5485	0.7256	48.36%	32.14%
Lasso(0.975)	0.5345	0.7186	49.77%	32.83%
Lasso(0.98)	**0.5232**	**0.7153**	**50.90%**	**33.17%**
Lasso(0.985)	0.5233	0.7167	50.93%	33.06%
Lasso(0.99)	0.5573	0.7283	47.77%	31.99%
Lasso(0.995)	0.6412	0.7533	39.78%	29.64%
Lasso(1)	0.7065	0.7878	33.63%	26.43%

The first two columns report unconditional results, i.e. in-sample and out-of-sample MSEs averaged over all companies and time periods. The third and fourth column list relative results, that is averaged percentage decrease in MSEs as compared to the benchmark model, which we label with "Mean".

Estimating γ (Lasso($\hat{\text{FF}}$)) increases out-of-sample performance by a statistically significant amount. When we compare MSEs resulting from a model estimated with the adaptive forgetting factor and from one with our preferred fixed choice (i.e. $\gamma = 0.98$) the Fluctuation Test proposed in Giacomini and Rossi (2010) rejects the null hypothesis of equal forecasting performance at each time period with a 95% confidence level. *However*, the economic significance of the additional forecasting power is certainly questionable, since estimating the forgetting factor increases out-of-sample R^2 by a mere 0.79%. Moreover (as is made clear in Section 1.3.6) allowing γ to vary over time requires a much higher computational cost due to the fact that the model has to be re-estimated for each value in the grid. Therefore, we decide to lose a little forecasting power in exchange for much lower computational costs, and stick to a model with fixed γ, which we set at 0.98.

For completeness of exposition, Figure 1.3 plots the estimates of the average forgetting factor. Interestingly, it generally seems that γ lowers during turbulent periods, and this is sensible since more flexibility is required in order to quickly pin down abrupt parameters change. Furthermore, as an illustrative example, Figure 1.4 plots the dynamic estimates of the forgetting factor for British Petroleum. The same reasoning we did above applies, and it does so in the idiosyncratic example we provide, that is γ drops to the minimum allowed value in the months after the Deepwater Horizon oil spill, which can safely be considered the worse shock the

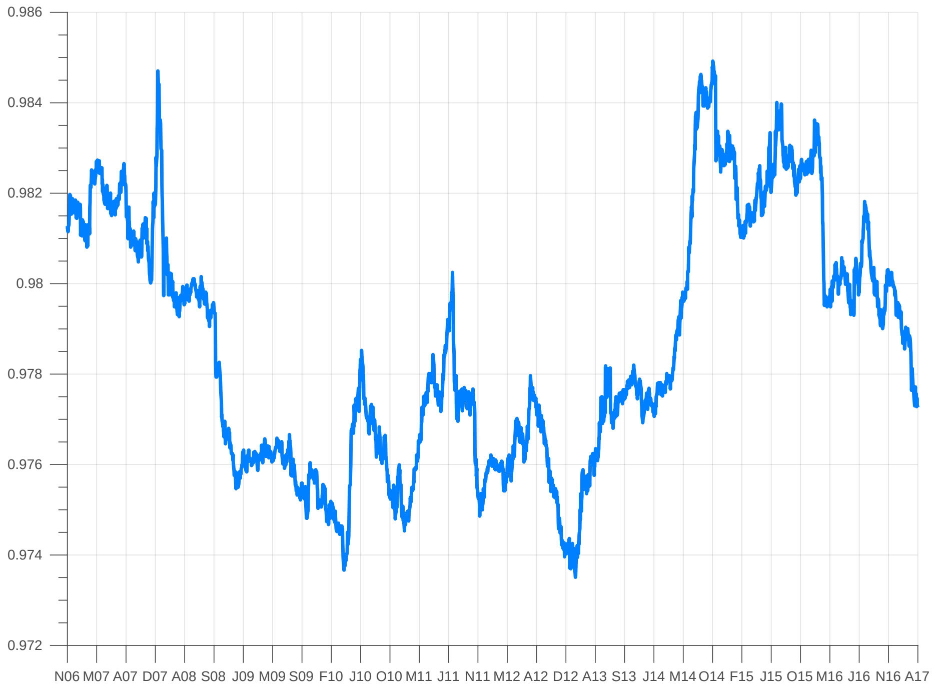

Figure 1.3: This figure averages forgetting factors estimates for all companies. The picture shows a general tendency to use less information during periods where systematic volatility is high. Moreover, the average γ almost always fluctuates inside the interval 0.975-0.985, meaning that apart from isolated, idiosyncratic, and temporary cases, this forgetting factor bandwidth contains optimal values for the great majority of the companies.

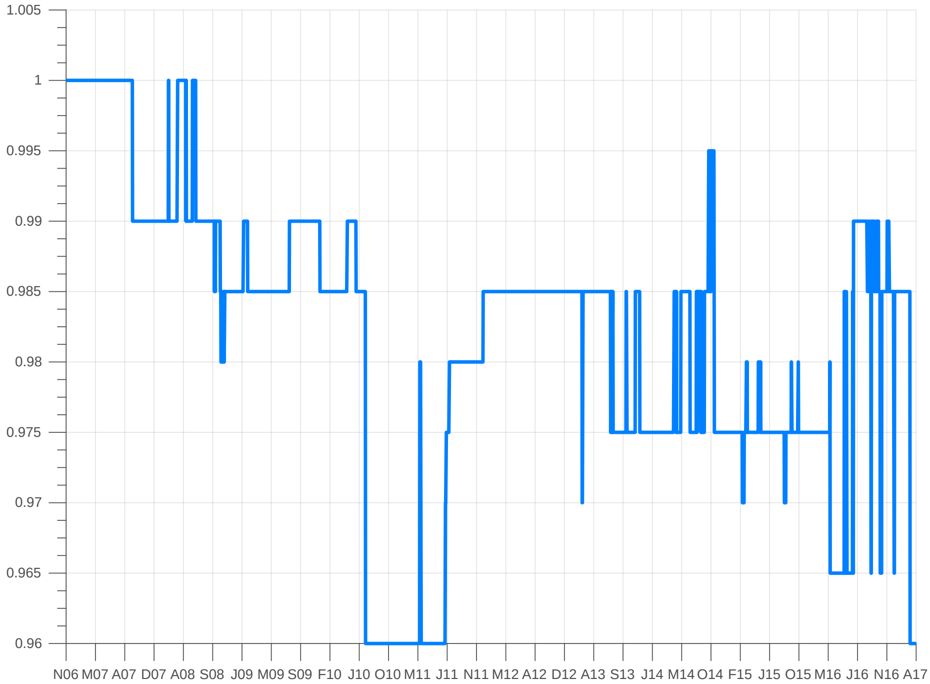

Figure 1.4: British Petroleum forgetting factor: the series fluctuates around 0.98, and reaches its minimum immediately after the disastrous oil spill whose company has been responsible for. Note that uniform weights (i.e. standard rolling windows) are optimal only for a small portion (approximately 10 months) of the sample length.

company had to bear. Finally, apart from a small time span, uniform weights are never the optimal choice.

A similar result is found for *all* the other companies in our dataset. The fact that using $\gamma < 1$ is always better than using $\gamma = 1$ is an important result because it means that always using more data is *never* the best thing to do. In this time series setting, this means that if we had (e.g.) 100 years of data we would not obtain better forecasting performance, and the opposite would be true. The reason is that parameters are changing, namely we do not want to average the values for the coefficients over too many observations, but we rather prefer to have a local perspective. In other words, $T \to \infty$ is not what we aim to.

1.3.4 Validation

As we said, the tuning parameter λ is chosen for each variable and at each point in time, allowing shrinkage to change over time. The most common way to estimate λ is through k-fold Cross-Validation. However, applying CV to time series data is known to be problematic because of time dependency between errors. Bergmeir, Hyndman and Koo (2015) shows that CV is still the optimal validation scheme only in the purely auto-regressive case. In the same paper, the authors cite the most relevant contributions where alternative CV methods have been developed for dependent data[8].

Far from being exclusively auto-regressive, our model is a good candidate for CV not to perform well, and when applying 10-fold CV on our TVP-FAVAR we indeed obtain unsatisfactory results: for some periods and for some companies we frequently observe over-fitting issues, where (in the Lasso specification) CV selects about 50 (out of 322) "relevant" regressors.

The over-fitting problem is very likely caused by time-varying data dependencies. For example, there might be a large number of cases where the test set contains observations that are best forecasted with low (high) shrinkage, whereas the training set might need a higher (lower) shrinkage. This is unlikely to happen systematically with independent data. Furthermore, CV almost surely selects training and test sets that contain non-consecutive observations, further decreasing the effectiveness of the penalization parameter choice. Those issues can contribute to strongly impair the overall ability to select the best tuning parameter at each window. One might then prefer to allow only the last observations to belong to the validation set, and all the previous ones to form the training set.

[8]We do not report those methods here and we refer to the references in the paper for a list of the relevant articles.

We then propose a local validation scheme, and later in the paper we show that it always outperforms CV. Crucially, LV is up to 50 times faster than 10-fold CV since the model has to be estimated only once.

More specifically, our methodology imposes the validation for λ to be local on a window of the S most recent observations. This can be achieved by the following steps[9]:

1. Set the final period of the first window to be $s = W$[10];

2. At time $t = s + (\mathbb{1}_{s>W} \cdot S)$, for a given decreasing sequence $\{\lambda_l\}_1^{l=L}$, estimate elastic nets with a window length $W - p$ and produce a forecast $y_{t+1|t,l}$ $\forall l$;

3. Move to $t = s + (\mathbb{1}_{s>W} \cdot S) + 1$ and repeat step 2 in a similar way. Compute one-step-ahead MSEs $\forall l$ as

$$\mathrm{MSE}_{t,l} = (y_t - y_{t|t-1,l})^2;$$

4. Repeat step 3 S times, until $t = s + S$ is reached;

5. Compute average MSE over the past S periods and $\forall l$ as

$$\overline{\mathrm{MSE}}_{t,l} = \frac{1}{S} \sum_{k=1}^{S} \mathrm{MSE}_{t-k+1,l};$$

6. Choose the tuning parameter[11] λ_l that minimizes average MSE:

$$\hat{\lambda}_t = \lambda_z | \overline{\mathrm{MSE}}_{t,z} = \min_{l \in \{1,\ldots,L\}} \overline{\mathrm{MSE}}_{t,l};$$

7. Set $s = s + 1$ and repeat steps 2,3,5 and 6;

8. Repeat step 7 until the final period $s = T$ is reached.

Intuitively, at every period the algorithm estimates L models on a wide spectrum of penalization parameters, and forecasts the relevant variable with the model that performed best on average in the previous S days. Each model is re-estimated

[9] For ease of exposition we omit the i subscript, but the following algorithm has to be thought of as being separately applied for each company $i \in \{1, \ldots, N\}$.

[10] This is the first period where coefficients estimates *for each value of the tuning parameter* can be retrieved. We will later see that the actual first usable estimates of the TVP-FAVAR parameters become available at time $t = W + S$.

[11] Note that in the formula to come the subscript in $\hat{\lambda}$ is a time subscript, whereas the one in λ refers to the grid position of the penalization parameter.

at every window on the same set of tuning parameters, and the optimal tuning parameter choice is updated accordingly.

We can then see that this algorithm turns out to be well suited for our model, because it lets λ to adapt based on past information, something that is desirable in a time-varying-parameters framework. Choosing a single tuning parameter for all time periods would have likely led to worse forecasting performance, and to a misleading dynamic representation of the network.

The reason why we use the average MSE to choose λ is because we want to smooth forecasting performances at each period. Given that we have to choose one parameter out of $L = 100$ possible ones (our baseline choice for L), we want to reduce as much as possible the probability of selecting one whose best very recent performance is simply the outcome of pure randomness. Therefore, we decide to average out-of-sample performance over $S = 400$ days. The (additional) reason why we can set S to be very high is because we have a lot of data, that is we can throw away the first $W + S = 800$ observations (our baseline choice) and still have a very interesting dataset to forecast. If we had had a smaller dataset (which would have been the case with macroeconomic applications) smoothing the estimates of the penalization parameters would have been more problematic and it very likely would not have led to the large performance improvements we observe in our setting.

1.3.5 Common Factors

Given that we use a fixed window width, Bai and Ng (2006) condition $\left(\frac{\sqrt{T}}{N} \to 0\right)$ approximately applies and factors can be treated as if they were observed. This ensures that the common factors we use are not contaminated by estimation error, something that would undermine our ability to correctly separate common movements from idiosyncratic ones. The forecasting performance of our baseline TVP-FAVAR as compared to the corresponding TVP-VAR is the same, and we interpret this result as evidence of the fact that estimated factors are as good as observed ones in our framework.

The first decision we have to make is how many factors to select. Andersen, Bollerslev, Diebold and Ebens (2001) and Luciani and Veredas (2011) decisively point towards a simple factor structure. Also, as Barigozzi, Brownlees, Gallo and Veredas (2014) suggests, overall volatility is likely to be the result of idiosyncratic volatility fluctuating around a systematic volatility component. The authors model the systematic component as a smoothly evolving trend, and find that it peaks during global turmoils such as the dot-com bubble or Lehman bankruptcy. In this

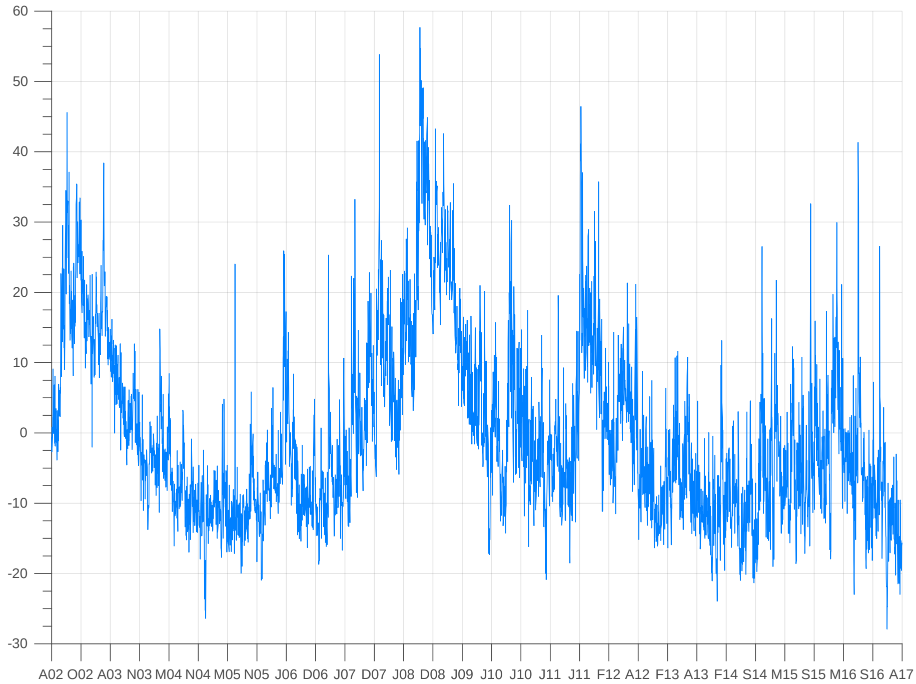

Figure 1.5: In this particular setting, the first principal component can be safely interpreted as the non-smoothed systematic volatility component. This figure plots the first common factor estimated on the full sample length, whereas in our analysis we re-estimate it on each specific window.

way, the single common factor underlying all volatility series is an approximate average volatility level that smooths out idiosyncrasies caused by individual companies. Therefore, it is easily understood that this trend is strongly related to the general underlying economic and financial forces around the periods one is analyzing. We then decide to follow this literature and we set the number of factors to be fixed at $r = 1$ throughout the whole sample period. Interestingly, we find that the standardized first principal component estimated over the whole sample is virtually the same as the standardized *simple* average of all the series (log-volatilities) we use to estimate the common factors. The correlation between the two series is indeed equal to 1 up to numerical error. Figure 1.5 illustratively plots the first principal component, and one can see significant peaks during (e.g.) Lehman collapse, but also during Greek, Italian, and Spanish sovereign debt crises.

The second fundamental reason why we decide to shrink coefficients on common factors is because computing principal components is an unsupervised procedure. Shrinking the coefficients is a simple way to supervise the common factors and possibly gain better out-of-sample performance, which we will show is our case. In some way, our approach is similar to the semi-supervised strategy proposed in

Bair, Hastie, Debashish and Tibshirani (2006), with a methodological difference. Whereas in Bair et al. (2006) principal components are supervised *before* they are computed, we first estimate them in the classic unsupervised fashion on the whole dataset and supervise the one we use by shrinking its coefficient *differently for each company* in each equation we estimate. Importantly, this supervision would not have been possible had we just got rid of common factors before running the forecasting regressions.

Controlling for common factors and using forgetting factors at the same time in our TVP-FAVAR raises a further issue. Indeed, recall that the only reason why we add principal components in our equations is because we want to get rid of the effect of common movements *at each time period*, and we want to do this in the same way for each observation in the window at hand. If we were to apply forgetting factors also to common factors, we would not be controlling for them in the uniform way we want. Intuitively, the most distant observations would control for common factors with exponentially lower strength. This would yield a non reliable network representation, because common and idiosyncratic effects would be confounded. Therefore, we need to apply uniform weights on each common factor observation, whereas we want to apply forgetting factors to the rest of predictors. Note that uniformly weighting common factors and applying exponential weights to the other regressors is equivalent to first de-factorizing the series (as Barigozzi and Brownlees (2017) does) and then applying forgetting factors to the residuals.

First, let's assume we want to do this on a simple OLS regression. Let's then suppose we want to estimate

$$\mathbf{Y} = \mathbf{X}\boldsymbol{\beta} + \boldsymbol{\epsilon} \tag{1.11}$$

with weighted OLS, where[12] $\mathbf{Y}$ is $T \times 1$, $\mathbf{X} = [\mathbf{F} \quad \mathbf{R}]$ is $T \times (r+K)$, $\mathbf{F}$ is $T \times r$, $\mathbf{R}$ is $T \times K$ and $\boldsymbol{\beta}$ is $(r+K) \times 1$. Now, define a $T \times T$ diagonal weighting matrix $\mathbf{W}$ with elements w_t on the diagonal ($t = 1, \ldots, T$), and compute $\mathbf{W}_c = \text{chol}(\mathbf{W})$ such that[13] $\mathbf{W}_c'\mathbf{W}_c = \mathbf{W}$. Weighted OLS is then computed by regressing $\mathbf{Y}^* = \mathbf{W}_c\mathbf{Y}$ on $\mathbf{X}^* = \mathbf{W}_c\mathbf{X}$. Therefore, we have:

$$\begin{aligned} \hat{\boldsymbol{\beta}} &= (\mathbf{X}^{*\prime}\mathbf{X}^*)^{-1}\mathbf{X}^{*\prime}\mathbf{Y}^* \\ &= (\mathbf{X}'\mathbf{W}_c'\mathbf{W}_c\mathbf{X})^{-1}\mathbf{X}'\mathbf{W}_c'\mathbf{W}_c\mathbf{Y} \\ &= (\mathbf{X}'\mathbf{W}\mathbf{X})^{-1}\mathbf{X}'\mathbf{W}\mathbf{Y} \end{aligned} \tag{1.12}$$

[12]In this particular example, notation has to be understood as being independent from the one we use throughout the rest of the paper, and is intended for illustrative purposes only.

[13]In a diagonal matrix, Cholesky factorization is tantamount to taking the square root of each diagonal element in the matrix.

Now, suppose that we want to apply weights only to the matrix $\mathbf{R}$ but not to the matrix $\mathbf{F}$. The solution is as simple as regressing $\mathbf{Y}^*$ on $\tilde{\mathbf{X}} = [\mathbf{F} \quad \mathbf{R}^*]$, where $\mathbf{R}^* = \mathbf{W}_c\mathbf{R}$. Therefore, we have that

$$\begin{aligned}
\hat{\boldsymbol{\beta}} &= \left(\tilde{\mathbf{X}}'\tilde{\mathbf{X}}\right)^{-1} \tilde{\mathbf{X}}'\mathbf{Y}^* \\
&= \left(\begin{bmatrix} \mathbf{F}' \\ \mathbf{R}'\mathbf{W}_c' \end{bmatrix} \begin{bmatrix} \mathbf{F} & \mathbf{W}_c\mathbf{R} \end{bmatrix} \right)^{-1} \begin{bmatrix} \mathbf{F}' \\ \mathbf{R}'\mathbf{W}_c' \end{bmatrix} \mathbf{W}_c\mathbf{Y} \\
&= \left(\underbrace{\begin{bmatrix} \mathbf{F}'\mathbf{F} & \mathbf{F}'\mathbf{W}_c\mathbf{R} \\ \mathbf{R}'\mathbf{W}_c'\mathbf{F} & \mathbf{R}'\mathbf{W}\mathbf{R} \end{bmatrix}}_{\equiv \mathbf{A}} \right)^{-1} \underbrace{\begin{bmatrix} \mathbf{F}'\mathbf{W}_c\mathbf{Y} \\ \mathbf{R}'\mathbf{W}\mathbf{Y} \end{bmatrix}}_{\equiv \mathbf{B}}
\end{aligned} \tag{1.13}$$

Intuitively, weighted OLS does two things: i) It gives relatively more importance to values of the *regressors* that have higher weight. This is apparent when looking at the bottom-right block of matrix $\mathbf{A}$ in equation (1.13), where the covariance matrix $\mathbf{R}'\mathbf{R}$ is weighted with the matrix $\mathbf{W}$. On the other hand, we can see that $\mathbf{W}$ does not appear in the top-left block of matrix $\mathbf{A}$, the reason being that we do not want to apply any weight on $\mathbf{F}$; ii) It gives relatively more importance to values of the *dependent variable* that have higher weight. This can be seen in both blocks of matrix $\mathbf{B}$ in equation (1.13), with a simple difference. Indeed, since we only want to weight $\mathbf{Y}$, the vector $\mathbf{F}'\mathbf{Y}$ is only partially weighted, that is we simply pre-multiply $\mathbf{Y}$ with $\mathbf{W}_c$. This is not the case in the bottom block, where the vector $\mathbf{R}'\mathbf{Y}$ is weighted with the whole matrix $\mathbf{W}$.

We reach a qualitatively similar result estimating the model with elastic nets, which (theoretically speaking) are simply a generalization of OLS regression. In particular, when the tuning parameter is equal to zero and $\mathbf{X}$ has full column rank, then the coefficients computed with elastic net are the same as the ones we find with least squares, which implies that the intuition provided by OLS is valid for elastic nets as well. Crucially, this parallelism holds true in the elastic net case also when the penalization parameter is positive and when $\mathbf{X}$ has not full column rank, as is (e.g.) the case in the high dimensional frameworks where elastic nets are frequently used. The only difference will lie in the actual estimates of the regression coefficients.

1.3.6 Computational Details and Speed Improvements

The fast coordinate descent algorithm we use exploits warm starts to obtain the optimal coefficients values much more efficiently and consequently faster than older methods. Nonetheless, recall that we have to run this model for each company and for each time period. Therefore, even though a single model is solved very quickly (hundredth of seconds) we have to solve $N \cdot (T - W) = 1,128,610$ models. This is the case with our (much faster) LV scheme. If we run the model with 10-fold CV we have to multiply the previous number by 10 (the number of folds) because the model is re-estimated at each fold, making the process even slower. Running various specifications of the model (required whenever one is interested in dynamically estimating the forgetting factor, or when we want to compare out-of-sample performances, as is our case) can then quickly lead to prohibitive computational times, and this is especially true for the specifications where we use 10-fold CV.

The only advantage of CV over our LV scheme is that CV is parallelizable. The way we program LV is such that we do not store coefficients at each point in time *and* for each of the 100 values of the tuning parameter. This option is feasible in theory and would allow to make LV parallelizable, but it is unfeasible in practice due to the then unreasonable memory requirements. The only feasible way to parallelize it would then be to use swap space, but this solution would surely wipe out the benefits of parallelization due to virtual memory I/O activity being extremely slow. Finally, running a single CV specification on one core takes between two and three *weeks*, making this not a feasible option.

Given all these considerations we parallelize only the CV specifications, and in order to obtain the results in reasonable times we use a large machine (128GB RAM) and we parallelize CV loops over 14 cores. This is the only way that allows to scale down CV estimation times to make them close to LV times, which nonetheless require much less computational power, being able to run on a single core with less than 8GB RAM.

In order to reduce the communication between workers and possibly gain further speed, we perform the parallelization over blocks of time periods, where we choose the size of each block to be equal to $\lfloor \frac{T-W}{15 \cdot nc} \rfloor = 16$ periods, where $nc = 14$ is the number of cores we use. Eventually, depending on the actual model, we need between 10 and 25 hours to have the results for only one specification.

In principle, we could use previous period parameter values as warm starts for next period coefficients. However, we would then have to abandon the standard warm starts computed with Friedman et al. (2010) algorithm, and it is not clear a priori that this alternative solution would be faster, especially in a rapidly changing environment like the one we are considering, where parameters can potentially change significantly from one day to the other. The authors themselves state that solving the model with their warm starts is as fast as running the single OLS counterpart for that model. Therefore, with a large enough grid of values for the tuning parameter (we use 100 values) using the standard algorithm should be enough to yield the solution in the fastest way.

1.4 Results

1.4.1 Network

Our network has $N(N-1) = 103,362$ potential links at each point in time, and the total number of estimated links ranges between 500 and 3,500, implying that the network density fluctuates between 0.5% and 3.4% approximately. Interestingly, the portion of links that change status (i.e. links that become active at time t when they were inactive at time $(t-1)$, and viceversa) is approximately equal to 30% of the total number of links *at each point in time*. This number is useful because it shows how stable connections are, and in turn how much one can look at network structure today to assess what are likely to be the connected links tomorrow. On average, 70% of the connections that are active/inactive at time t turn out to be in the same status also at time $t+1$, meaning that our network *structure* can be well predicted in the short term. As we anticipated, the forecasting performance of our TVP-FAVAR as opposed to the TVP-VAR counterpart are the same, but we find that the network density is significantly lower in the first case. In particular, we find almost 10% (on average) less connections than the baseline case, and this result corroborates the usefulness and effectiveness of controlling for common factors. Moreover, correlation between the two average degrees is equal to 0.985, meaning that the two series have (to all practical purposes) the same time profile, but at every period our model correctly gets rid of a roughly constant amount of spurious connections.

Figures 1.6-1.7 show how the network structure has evolved historically over time. Recall that our network is directed, thus we both have in-degree and out-degree distributions at each point in time. The metrics we use to summarize the structure are the average, standard deviation, and skewness of the degree distributions. Importantly, the average in-degree and the average out-degree are equivalent, so we will simply refer to them as the average (or mean) degree. As we can see, the mean degree spikes up especially during the 2008 financial crisis and the European sovereign debt crises. Note that those spikes did not necessarily need to be so pronounced, because we already are taking into account the systematic volatility by augmenting the forecasting equations with the common factor. In other words, if systematic events would affect every company through the common

Table 1.4: Degree distributions: correlation between main metrics

		AD	SDID	SDOD	SKID	SKOD
Average degree	AD	1	0.86	0.85	-0.72	-0.03
Std. dev. in-degree	SDID		1	0.58	-0.49	-0.19
Std. dev. out-degree	SDOD			1	-0.76	0.35
Skewness in-degree	SKID				1	-0.26
Skewness out-degree	SKOD					1

The table shows full-sample correlations between the first three moments of the in-degree and out-degree distributions. The only metric that seems to fluctuate independently is the skewness of the out-degree distribution, whereas all the other measures are strongly related between each others.

components alone, we would obtain that the average degree time path would entirely be independent of systematic events, exactly because we are controlling for them. The fact that we still see a positive correlation between the average degree and the occurrence of crises episodes is interesting because it means that crises periods tend to awake additional non-systematic connections that would have likely stayed dormant otherwise.

Furthermore, we find that the shape of the degree distributions tend to move with an approximately fixed scheme. Indeed, the standard deviation of both the in-degree and the out-degree distribution is very positively correlated with the average degree. The same is true (with negative correlation) for the skewness of the in-degree distribution, but not for the out-degree distribution one. Those patterns imply that during tranquil times the network structure tends to have very few connections in the great majority of nodes. On the other hand, turbulent times activate and spread out connections. For what concern in-degree distribution, those connections are scattered in a relatively regular way, meaning that we have many nodes with both a relatively high and low number of in and out connections. We can be sure of this because of the fact that skewness tends to approach zero during those times, thus ensuring that standard deviation is not increasing simply because of a few extremely highly in-connected nodes. This means that a relatively high number of companies are fragile to turmoils, because their volatilities really seem to be affected by a higher number of other companies' volatilities. On the other hand, connections in the out-degree distribution tend to be concentrated in a small (but varying) set of companies constantly during the whole sample period. Indeed, even though skewness is completely unrelated with the average degree, it is constantly high and positive (much more than the in-degree counterpart) implying that there always are an extremely low number of nodes that are very central and affect an extraordinarily high number of companies at each time period. Table 1.4 summarizes the existing correlations we cited between the metrics of the in and out-degree distributions.

Our TVP-FAVAR (as opposed to the TVP-VAR) can also answer questions regarding the nature of certain events. For instance, one could be interested in knowing whether a given key event has been simply systematic to the companies in the network. In our setting, a purely systematic event for company i is an event that does not increase company i's in-degree, the reason being that the systematic part is captured by the common factor coefficient. Crucially, note that a given event can happen to be systematic for company i but not for company j, and (as we will show and explain in the next paragraphs) this is indeed what occurs in our network during certain periods.

Table 1.5 ranks companies by their out-degree during three different days. The first is September 15^{th}, 2008, when Lehman Brothers went bankrupt. As we can see from the average degree plots (e.g. Figure 1.7), tensions were steadily mounting since at least the very beginning of 2007, although bad news from the US housing market already became a serious issue weeks before. During this stressed period we see a sharp increase in the average degree, which went from 1.5 to approximately 5 immediately before Lehman crash. At the onset of the financial panic (top panel of Table 1.5) we have a starting point of no significant higher out-degree centrality from financial sector companies. Indeed, even though we see two big insurance companies among the top three out-degreeing ones, the network structure was still relatively homogeneous as what regards sectoral distribution. However, starting from the very day after we observe an incredibly pronounced spike in *overall* average degree (again, see Figure 1.7), and this suggests that to *most* European companies the ongoing turbulence was not just systematic. After ten days from the announcement (middle panel in Table 1.5), when panic was widespread, the network structure had significantly changed: as we said, out-degree generally spiked up, and financial companies started to be more central, until some day later (October 10^{th}, 2008, lower panel of Table 1.5) out-degrees reached unprecedented levels and the five most central companies all belonged to the financial sector.

Looking at the reaction of the network structure after Lehman bankruptcy, it really seems that there were some companies for which that event was not purely systematic, and this can be inferred simply looking at the violent upward reaction of average degree, which went from 4 to more than 8 connections in very few days. If Lehman crash would have been totally systematic to the companies in

Table 1.5: Most important nodes, before and after Lehman crash

September 15^{th}*, 2008*	Out Rank	Out-degree	In Rank	In-degree	Sector
Imperial Brands	1	46	107	5	Cons Stapl
Ageas	2	42	107	5	Financials
Helvetia Holding	3	39	49	9	Financials
Endesa	4	34	148	3	Utilities
Orpea	5	32	220	0	Health Care
Saipem	6	29	107	5	Energy
DNB	7	23	148	3	Financials
Hexagon B	7	23	10	16	IT
Norsk Hydro	9	20	148	3	Materials
Repsol	9	20	86	6	Energy
September 25^{th}*, 2008*	Out Rank	Out-degree	In Rank	In-degree	Sector
Imperial Brands	1	75	115	7	Cons Stapl
Helvetia Holding	2	74	96	8	Financials
Ageas	3	71	220	2	Financials
DNB	4	54	157	5	Financials
Lundin Petroleum	5	46	220	2	Energy
Saipem	6	39	157	5	Energy
Norsk Hydro	7	38	175	4	Materials
Lloyds Banking Group	8	37	40	12	Financials
Endesa	9	36	40	12	Utilities
Pearson	9	36	255	0	Cons Discr
October 10^{th}*, 2008*	Out Rank	Out-degree	In Rank	In-degree	Sector
DNB	1	129	228	5	Financials
Helvetia Holding	2	123	107	10	Financials
Deutsche Bank	3	85	206	6	Financials
Baloise	4	80	206	6	Financials
Ageas	5	75	278	0	Financials
Saipem	6	63	91	11	Energy
SAP	7	59	248	4	IT
Sulzer	7	59	179	7	Industrials
Volkswagen	9	55	206	6	Cons Discr
Commerzbank	10	50	134	9	Financials

Apart from the standard network interpretation, in-degrees and out-degrees also have a statistical meaning in our model. For every company i, its in-degree represents the number of non-common/non-auto-regressive variables that forecast company i volatility, whereas company i out-degree is the number of variables Granger-caused by the company.

our network, we would have seen no increase in the average degree. Moreover, the additional in-connections particularly stem from the kind of companies that are most related to the systematic event we are considering, that is they mainly come from some of the companies in the financial sector, and this can be seen from the rapid increase in out-degree in the financial companies that appear in Table 1.5, which means that banks and insurance companies started to effectively forecast many other volatilities in the network. On the contrary, the very same event happened to be entirely systematic for those very central financial companies. Indeed, in a period where average degree more than doubled, Table 1.5 shows that the in-degree of the most central financial companies either stayed stable or went down in the aftermath of the financial panic, and the low in-degree rankings confirm that those companies were among the ones that were mainly affected by the common component. Importantly, those results are not general in the financial sector, but we found them only in the companies that were most central during the very turbulent period following Lehman bankruptcy. Indeed, Figure 1.10 shows (as an example) that Spanish bank BBVA started to suddenly be affected by a considerable amount of other companies' volatilities in the aftermath of Lehman crash.

As an individual and not necessarily representative example, let's look at how BBVA network position evolved during the sovereign debt crisis. Remember that Spain public debt and banking sector crisis exploded at the end of 2011 and reached its peak in July 2012, just before Draghi's "Whatever it takes" and the activation of 41 billion euros ESM bailout programme for Spanish banks. Again, our TVP-FAVAR reveals that the unfolding of the debt crisis was something more than simply systematic *to BBVA*, that is a few network connections played an additional role in determining BBVA volatility, as confirmed in Figure 1.10. Generally speaking, both the 2008 financial panic and the sovereign debt crises really seem not to be just systematic for the majority of the companies in the network.

We now want to provide two examples on completely idiosyncratic events that occurred to two of the companies in our sample, that is Swedish Match and British Petroleum (BP). Swedish Match is the only listed European manufacturer of smokeless tobacco products. However, snus and moisting snuff are banned in the EU since 1992, meaning that this law prevented Swedish Match to access the potentially enormous EU market during all our sample period. On July 14^{th}, 2011, the Swedish Government *directly* sent an official (public) letter to EU Commissioners, lobbying for lifting the ban. Figure 1.8 shows that immediately after this letter was sent Swedish Match betweenness spiked to unprecedented

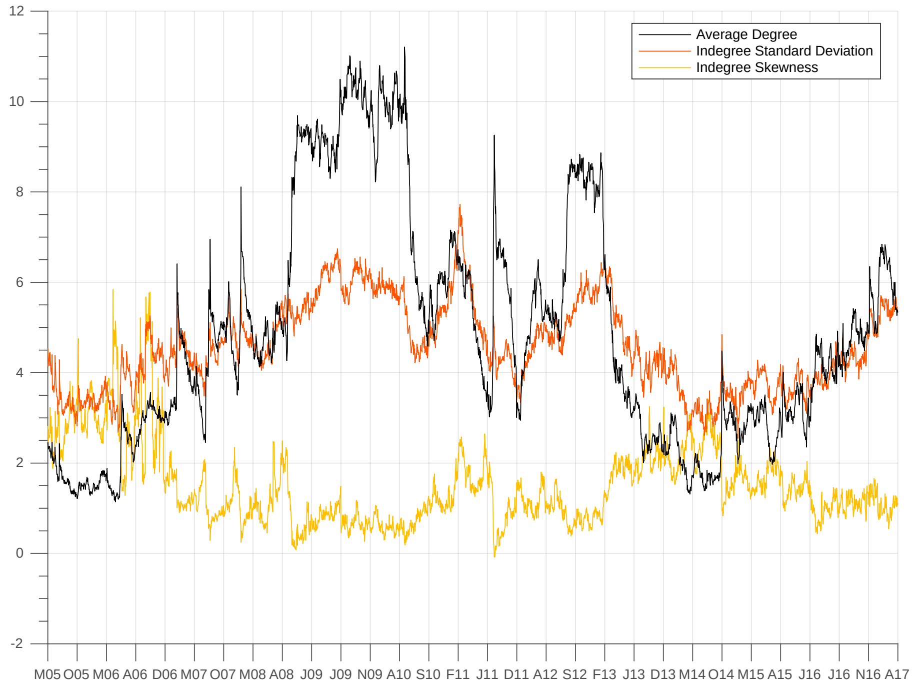

Figure 1.6: In-degree distribution. Note that not only the *number* of connections vary over time, but also their *composition* (as measured by second and third moments) does, and they all change together with volatility regimes.

levels and the company started to be the most central node in our network. We can now refer to Figure 1.1 again. It is noteworthy to see that Swedish Match volatility did not spike in the days surrounding the lobbying activity, that is nothing can be inferred by simply looking at volatility for this company. *However*, our model records Swedish Match to have the *overall* (i.e. considering all the companies in our sample) largest historical individual centrality metrics in the days following the publication of the letter. This result shows how powerful our model is in extracting key information from the very noisy large dataset we are considering.

A similar behavior is observed for British Petroleum (Figure 1.9) promptly after the Deepwater Horizon oil spill on April 20^{th}, 2010. After this catastrophic event BP starts to effectively forecast 55 companies' volatilities out of the average 5-6 it was forecasting beforehand. The examples reported above are meant to be illustrative only, and we could have added many more which are equally interesting. However, in the interest of brevity, we do not report them here but we leave them available on request.

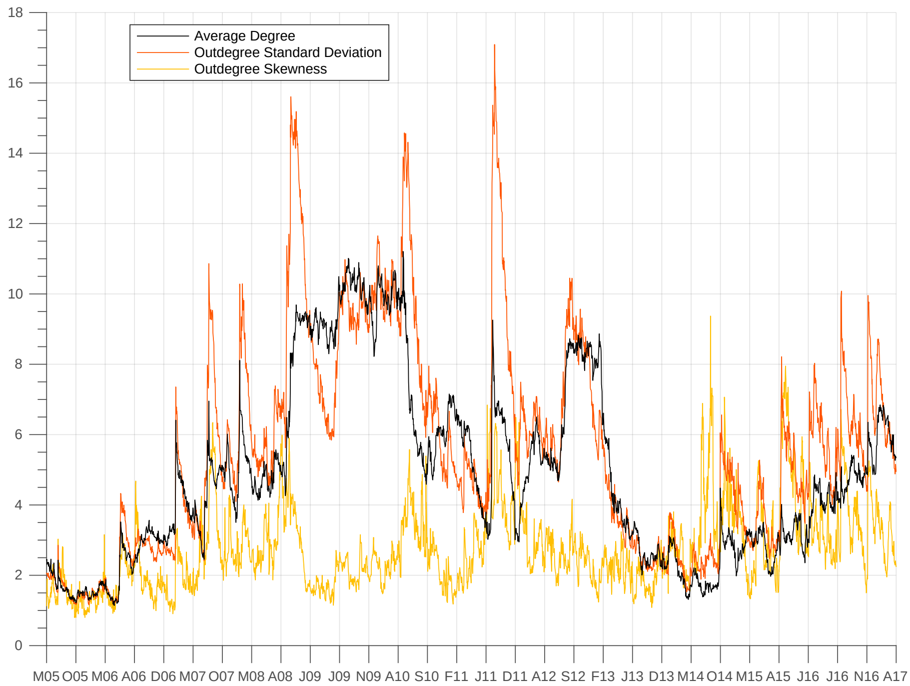

Figure 1.7: Out-degree distribution. The considerably high standard deviation and constantly positive skewness reveal that the out-degree distribution is mainly characterized by a few very central companies affecting the others.

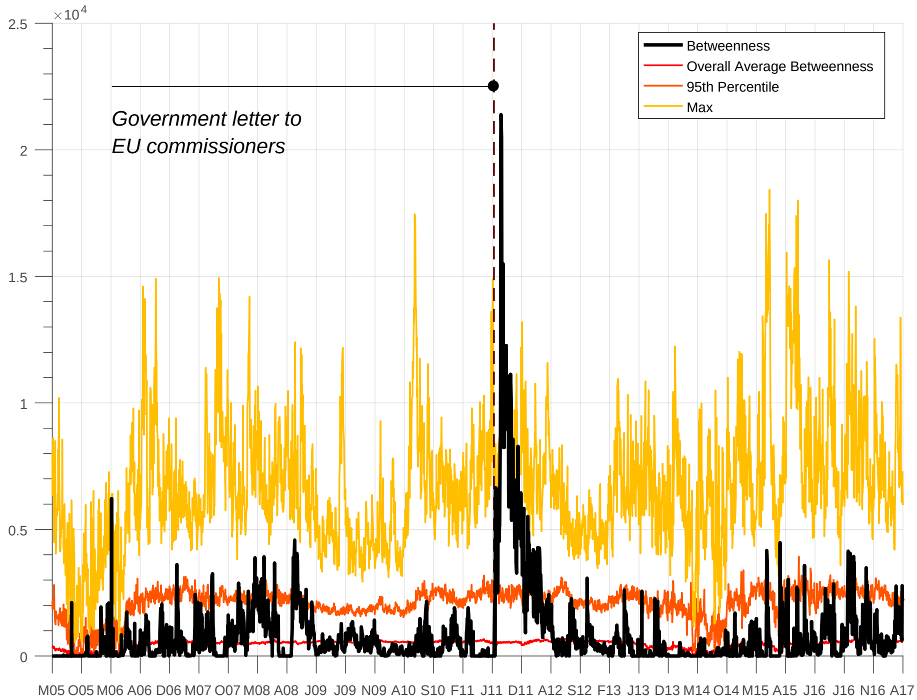

Figure 1.8: Swedish Match Betweenness. Betweenness is a network metrics that equals the overall number of shortest paths that pass through the company under consideration. A shortest path between two companies is the *smallest* set of edges that connects the two nodes.

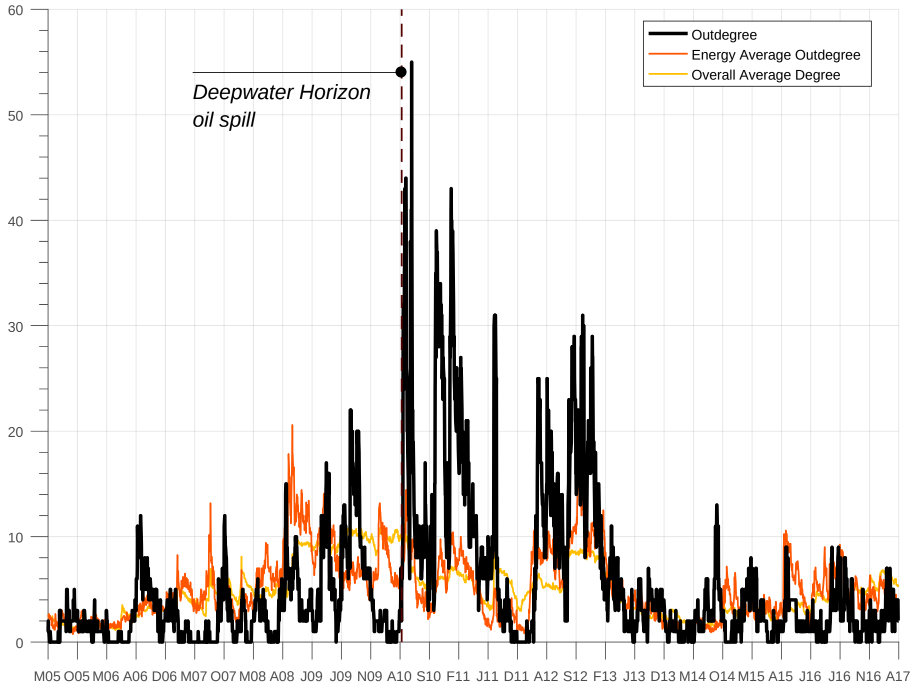

Figure 1.9: British Petroleum Out-degree. Note that after the highlighted event out-degree greatly exceeds both the overall average and the average out-degree of the sector to which the company belongs to, that is the Energy sector.

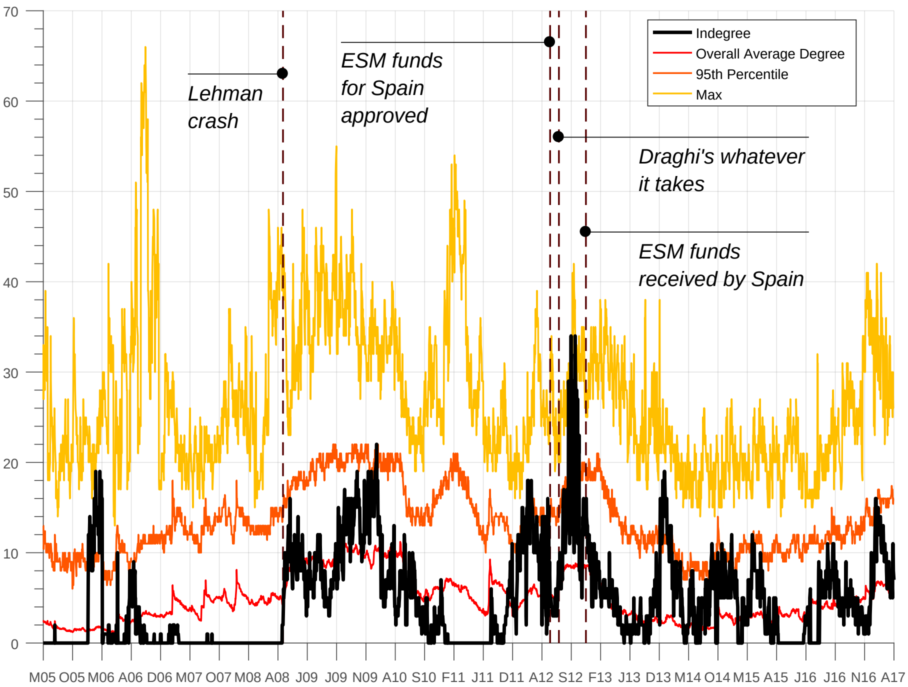

Figure 1.10: BBVA In-degree. Prior to the advent of the financial crisis, BBVA were mainly not influenced by other companies, but the situation drastically changes in correspondence with Lehman failure and the sovereign debt crises, which expose the bank to external idiosyncratic shocks.

Table 1.6: In-sample and out-of-sample results, TVP-FAVAR

	In-sample MSE	Out-of-sample MSE	In-sample R^2	Out-of-sample R^2
Mean	1.0792	1.0851	0%	0%
Lasso(LV,FF,SF)	**0.5591**	**0.7457**	**48.11%**	**31.00%**
Lasso(LV,FF,CF)	0.5431	0.7589	49.59%	29.75%
Lasso(CV,FF,CF)	0.4918	0.7887	54.22%	26.89%
Lasso(CV,FF,SF)	0.4581	0.7910	57.42%	26.63%
Lasso(LV,UW,SF)	0.7276	0.8124	32.17%	24.79%
Lasso(CV,UW,SF)	0.5644	0.8855	47.83%	17.87%
Lasso(LV,UW,CF)	0.6037	0.9035	44.12%	16.28%
Lasso(CV,UW,CF)	0.5695	0.9569	47.36%	11.22%
Ridge(LV,FF,SF)	**0.5077**	**0.7322**	**52.78%**	**32.26%**
PC	0.6764	0.8056	37.82%	25.95%
AR	0.8050	0.8150	25.08%	24.56%

Lasso models are ranked by out-of-sample R^2. Our baseline model is the best performing one, and the same is true when we estimate the same system of equations with Ridge regressions. For practical purposes, we only report Ridge results obtained when estimating the model with our novel features, whereas we omit the other specifications.

1.4.2 Forecasting

Table 1.6 displays our (average) forecasting results. The benchmark model is a time-varying average (denoted as Mean), which is the optimal forecast in case no additional information helps in better predicting the dependent variables. Our baseline model (denoted as Lasso(LV,FF,SF)) is a Lasso with window length $W = 400$, smoothing window $S = 400$, number of lags $p = 1$, forgetting factor $\gamma = 0.98$, and $r = 1$ principal component. The competing models are an AR(1) and principal components (PC) with one factor estimated at each window. Importantly, we also estimate every combination of Lasso models with our local validation scheme or 10-fold Cross-Validation, forgetting factors (FF) or uniform weights (UW), non-shrinked common factors (CF) or shrinked common factors (SF), with the purpose of better appreciating which feature improves the most out-of-sample performance in each situation. We rank the eight possible specifications with decreasing values of out-of-sample R^2.

Before commenting the main results we find, note that we also report results for Ridge regression estimated with the same specifications as our baseline Lasso model. One can see that out-of-sample performance is slightly better in Ridge as compared to Lasso. Nonetheless, we are confident that Lasso is a superior alternative to Ridge in computing the connections we are interested in. Indeed, the problem with Ridge is that it shrinks coefficients roughly by the same *proportion*. Therefore, explanatory power is distributed among all the coefficients in similar ways, and *no coefficient* will be exactly set to zero. On the other hand, Lasso shrinks coefficients approximately by the same amount κ, where the exact value of κ depends on the actual data and on the penalization parameter. All the coefficients that are less than κ in absolute value are set to zero (soft thresholding). This means that variables that are not helpful enough in predicting the outcome of interest are simply discarded. This is very useful in our setting, because it might well be that company i and company j are very weakly connected after controlling for all the other predictors, but this connection could well be economically insignificant. Lasso helps exactly in this respect, i.e. it gets rid of connections that are indeed insignificant, something that Ridge regression is not able to achieve *by construction*. Again, we therefore exchange very little forecasting power with a considerably more interpretable and insightful model.

Let's now go back to the general results we list in Table 1.6. In general, both baseline and competing models are able to explain and predict a significant additional portion of the data as compared to a simple time-varying mean.

Remarkably, and except for Ridge regression which performs very slightly better, our baseline (Lasso(LV,FF,SF)) is the best forecasting model, with an out-of-sample R^2 equal to 31.00%. Importantly, the model where we do not use any of our contributions (i.e. Lasso(CV,UW,CF)) is the worst performing one, with an out-of-sample R^2 equal to 11.22%. We want to stress how large this difference is, especially taking into account the fact that those numbers are averaged over a very large number of companies (322) and time periods (3105). This means that this result is *systematically* present and is not simply a one-time unlikely coincidence. The enormous difference in performances seems to be mainly due to the use of FF, although also LV and the fact that we are supervising the factors play a significant role in improving results. On average, applying forgetting factors increases forecasting performance by 11.03%, whereas validating the tuning parameter locally and shrinking common factors contribute 4.80% and 4.04% respectively (on average). Therefore, we can see that without LV, FF, and non-weighted factors supervision, Lasso would not be competitive as compared to AR and PC, and with a great gap in performance. It is only applying our features that we can significantly beat competing models.

Again, Table 1.6 demonstrates that our model not only forecasts better, but it also explains a significantly higher portion of the variation as compared

Table 1.7: In-sample and out-of-sample results, baseline model, sectoral disaggregation

	In-sample MSE	Out-of-sample MSE	In-sample R^2	Out-of-sample R^2
Financials	0.5498	0.7549	54.35%	**37.40**%
Real Estate	0.5505	0.7172	48.67%	33.50%
Materials	0.5369	0.7194	49.42%	32.46%
Energy	0.5314	0.7170	49.55%	32.25%
Consumer Discret.	0.5544	0.7351	47.89%	31.03%
<u>All sectors</u>	0.5591	0.7457	48.11%	31.00%
Utilities	0.5364	**0.7012**	46.24%	30.39%
Industrials	0.5566	0.7484	47.40%	29.44%
Telecom Services	0.5252	0.7034	46.85%	29.13%
IT	0.5425	0.7303	47.08%	29.08%
Consumer Staples	0.5716	0.7407	42.31%	25.28%
Health Care	0.6762	0.8598	39.74%	23.17%

The above results are disaggregated by sector, where each of them has different benchmark results. This explains why ranking sectors by out-of-sample R^2 (a relative measure) leads to a different ordering than the one obtained by ranking sectors by out-of-sample MSE (an absolute measure).

to competing PC and AR models. Moreover, the problems of estimating the model with CV are evident, because while in-sample performance is on average significantly higher as compared to the one of models estimated with LV, the reverse is true regarding out-of-sample results. This clearly suggests that the worries we raised in Section 1.3.4 are confirmed, and that CV suffers from an over-fitting problem in our framework.

Table 1.7 displays baseline results by sector, ranked by out-of-sample R^2. Financials and Real Estate are the sectors that perform significantly better than average, whereas Consumer Staples and Health Care are the ones where relative performance is worse. However, remember that our definition of out-of-sample R^2 is the percentage decrease in out-of-sample MSE as compared to a time-varying mean, and this mean is of course specific to each company. Therefore, in this case the ranking orders the sectors starting from the ones whose baseline improves the most over their respective benchmark models to the ones that improve the least. In absolute terms, the most predictable sectors are Utilities and Telecommunication Services, and this can be seen by looking at the out MSE column.

The results reported in Tables 1.6-1.7 are averaged over the whole sample length. However, our interest also focuses on possible time-variation in the relative forecasting performances of different models. Therefore, we apply the Fluctuation

Test as developed in Giacomini and Rossi (2010) to check whether there are clear occurrences that some model forecasts better than others during certain periods. In each of Figures 1.11-1.15, the top panel plots the (300-days) average rolling difference between model 1 and model 2 (in this order) MSEs *averaged* over all the N companies in our sample. The bottom panel is the same as the top one, but the MSEs differences are now standardized as in Giacomini and Rossi (2010), meaning that the resulting series is the day-by-day test statistic that we compare against the constant 95% two-sided critical values (the two black dashed lines). The actual figures for the critical values are not only determined by the confidence level, but also (using the authors' notation) by m, the number of time periods we use to smooth the loss differences, and by P, that is the number of periods where MSEs are available. Specifically, the critical values are an implicit function of $\mu = \frac{m}{P} = \frac{300}{3105} = 0.0967 \approx 0.1$ in our case. At any given period, if the test statistic is higher than the top critical value it means that out-of-sample performance of model 1 has been significantly worse than that of model 2 *during the last 300 days*. The converse is true when the test statistic is lower than the bottom critical value. Whenever the test statistic lies in the middle, it means that (on average) no statistically significant difference has been detected in that time span.

Figure 1.11 shows that Lasso is always performing better than AR (red bands are 95% confidence intervals), and this is true in particular during the financial panic and the European sovereign debt crisis. The reason for this might be that our multivariate models absorb spillovers faster than AR models which needs more time to adapt to quick and sudden breaks. Figure 1.12 highlights a similar result for PC regression. However, note that we do not see a spike in the relative performance during the crisis, and this could be a corroboration of our previous statement, i.e. that multivariate models capture spillovers in a more timely manner than AR models. Figures 1.13-1.15 demonstrate that each of the features we add individually improve out-of-sample performance consistently over time.

1.5 Concluding Remarks

In this paper we forecast the volatility panel of the FTSE Developed Europe Index constituents and we exploit the resulting forecasting relationships to compute a time-varying directed volatility network. We estimate volatilities in a model-free way, something that makes the forecast competition challenging, and we organize data as in a TVP-FAVAR, estimating it equation by equation for each time period with elastic nets. Controlling for common factors is crucial to be able to separate

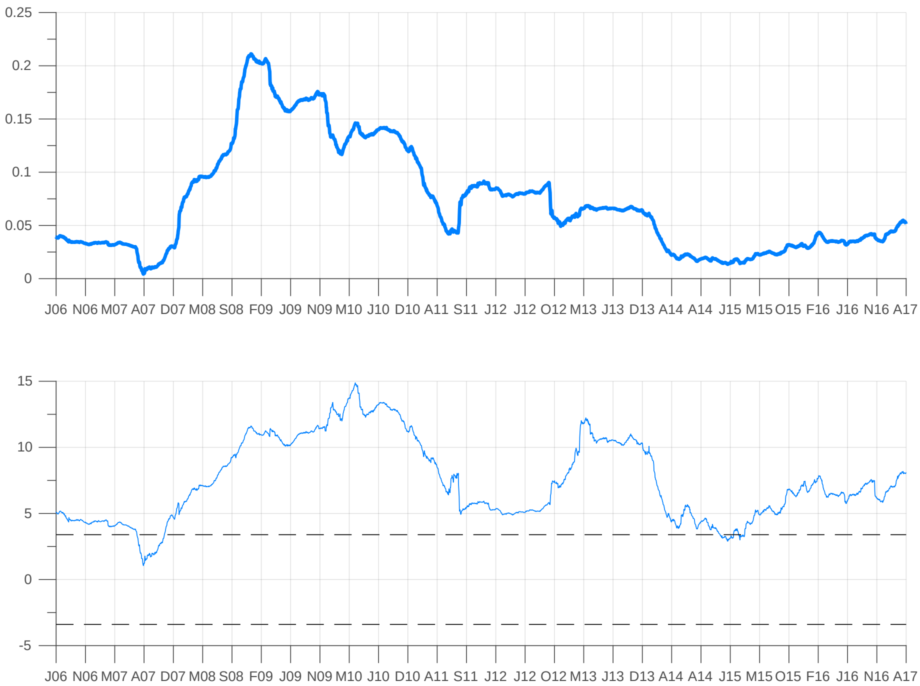

Figure 1.11: AR versus Lasso(LV,FF,SF): our baseline model always outperforms an AR(1), and reaches its best relative performance in the aftermath of the September 2008 financial panic.

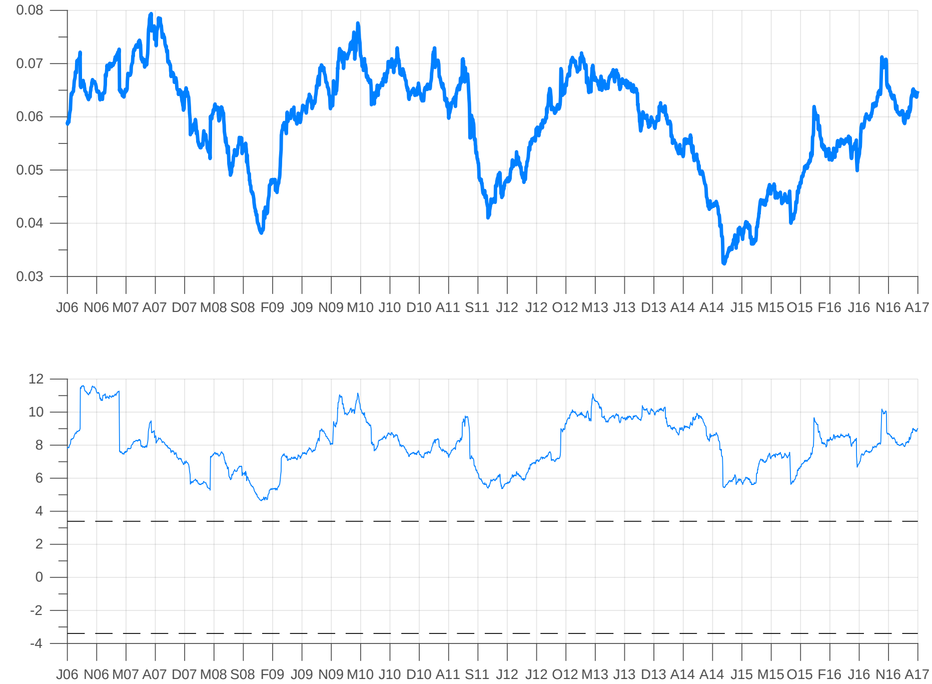

Figure 1.12: PC versus Lasso(LV,FF,SF): the difference in performances with respect to models estimated with principal components is less pronounced than the AR case, and the time profile seems to be independent of financial events.

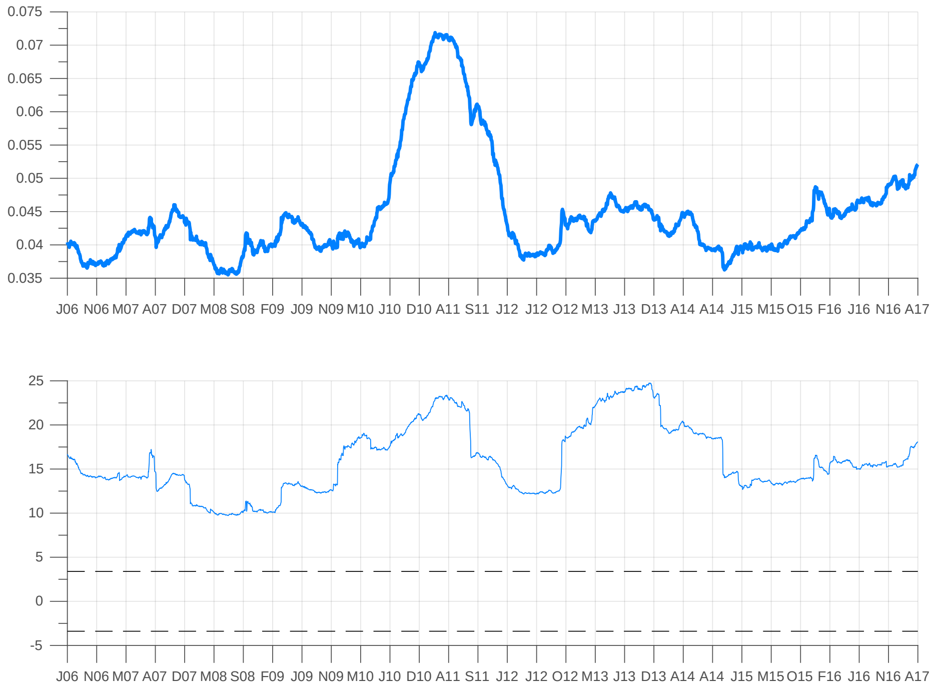

Figure 1.13: Lasso(CV,FF,SF) versus Lasso(LV,FF,SF): Our local validation algorithm improves over 10-fold Cross-Validation especially starting from the Sovereign debt crises.

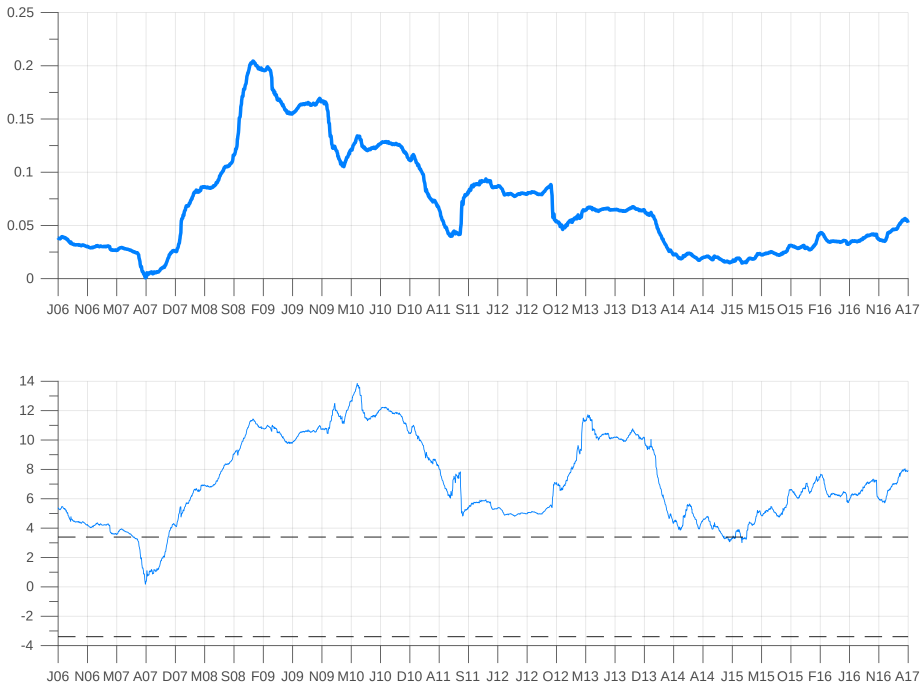

Figure 1.14: Lasso(LV,UW,SF) versus Lasso(LV,FF,SF): Again, our baseline performs better than a model estimated with standard rolling windows particularly during crises, where flexibility is crucial to locally identify quick and frequent changes in parameters.

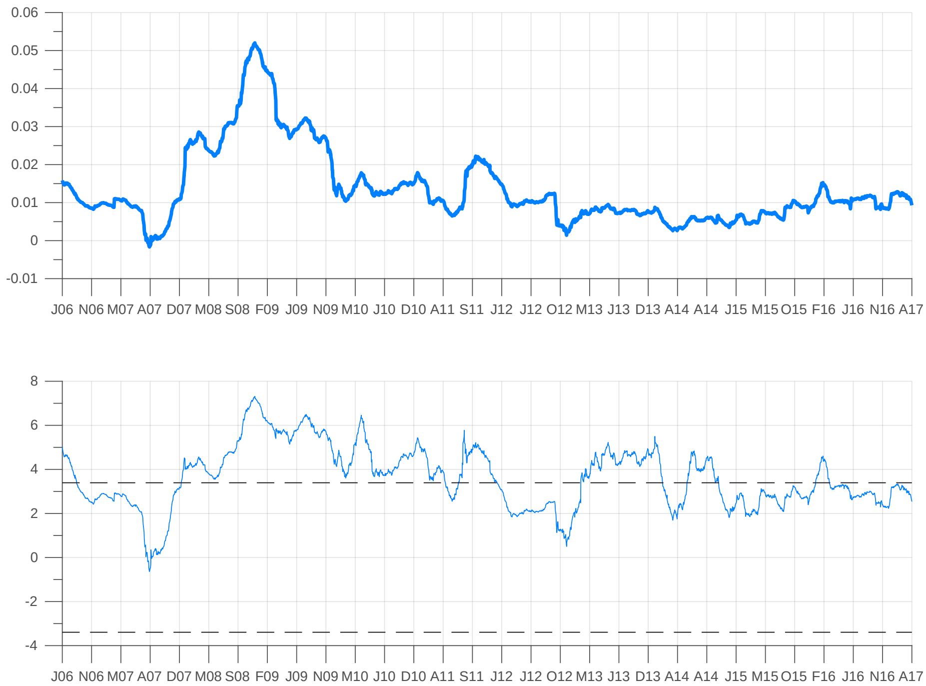

Figure 1.15: Lasso(LV,FF,CF) versus Lasso(LV,FF,SF): Shrinking common factors virtually always leads to better forecasts as compared to purely controlling for them in-sample, although improvements are not always statistically significant.

systematic effects from idiosyncratic ones. Moreover, doing it *within* the model (as opposed to de-factorizing the series) allows to forecast the original dataset.

We estimate the model with novel methods which are especially and exclusively suited for our time-varying-parameters setting. First, we propose a simple local validation algorithm to estimate tuning parameters in the elastic net. We show that our methodology always outperforms 10-fold Cross Validation, and it runs up to 50 times faster. Second, we maintain local flexibility of selected coefficients by applying exponentially decaying observation weights at each window, and we do so on all the predictors but the single (first) principal component, whose only function is to control for systematic volatility. Third, instead of simply controlling for common factors in-sample, we adopt a semi-supervised approach and shrink their coefficients whenever including them does not lower out-of-sample MSE. We show that the joint combination of the features we introduce yields very large out-of-sample gains. In general, our multivariate models outperform AR processes especially during periods of financial stress. This might be due to the fact that our model can accommodate spillovers earlier, whereas the same spillovers might need more time to update AR parameters properly.

We find that our network structure is on average very predictable in the very short term, and that the first three (two) moments of the in-degree (out-degree) distribution closely move together over the whole sample, thus revealing an interesting fact about volatility network structure, meaning that edges tend to be homogeneously distributed during quiet times, but they quickly concentrate in a few central nodes whenever very systematic events occur. Moreover, we find that those systematic events tend to trigger additional idiosyncratic connections *on top* of the ones that would be justified by common components, which we control for by augmenting the VAR with one common factor in each equation, as extensively explained throughout the paper. Finally, we provide completely idiosyncratic examples to show that whenever a node is individually affected by a particularly important event it promptly moves towards the middle of the network, thus suddenly becoming significantly more central.

Extensions of our model could head towards allowing for a time-varying set of nodes. This could be very useful in practice because it would allow the researcher to use companies that started or stopped trading in the middle of the sample, and it would mean to have both a larger and a more representative set of companies at each point in time. Moreover, the implementation would be relatively straightforward. On the other hand, one would have to find a way to represent the evolution of network metrics in a meaningful way, because we would need to take into account that some periods would have a more populated network than others. For instance, centrality measures such as in-degree and out-degrees could be computed in relative terms, e.g. as fractions of the time-varying number of nodes.

Finally, an obvious modification could be to estimate the model with adaptive elastic nets as in Zou and Zhang (2009) which enjoy the oracle property. As a consequence of the coefficient-specific shrinkage induced by those models and because of the particular set of predictors we use, we would surely obtain that the coefficient on the common factor would be the least shrunk regressor, and this would likely lead to an even sparser network representation.

1.A List of Companies

Table 1.A.1 lists all the companies we use in our analysis. Together with the name of the company, we add the corresponding GICS sector and industry, as well as the country where the company is listed.

Table 1.A.1: FTSE Developed Europe Constituents List

#	Company Name	Sector	Industry	Country
1	A.P. Moller-Maersk A	Industrials	Transportation	DNK
2	A.P. Moller-Maersk B	Industrials	Transportation	DNK
3	ABB	Industrials	Capital Goods	CHE
4	Aberdeen Asset Management	Financials	Diversified Financials	GBR
5	Abertis Infraestructuras	Industrials	Transportation	ESP
6	Acciona	Utilities	Utilities	ESP
7	Accor	Cons Discr	Consumer Services	FRA
8	Acerinox	Materials	Materials	ESP
9	Ackermans and Van Haaren	Financials	Diversified Financials	BEL
10	ACS	Industrials	Capital Goods	ESP
11	Adecco Group	Industrials	Comm&Prof Services	CHE
12	Adidas	Cons Discr	Cons Drbls&Apparel	DEU
13	Aegon	Financials	Insurance	NLD
14	Ageas	Financials	Insurance	BEL
15	Aggreko	Industrials	Comm&Prof Services	GBR
16	Air Liquide	Materials	Materials	FRA
17	Airbus	Industrials	Capital Goods	FRA
18	Akzo Nobel	Materials	Materials	NLD
19	Alba	Financials	Diversified Financials	ESP
20	Alfa Laval	Industrials	Capital Goods	SWE
21	Allianz	Financials	Insurance	DEU
22	Anheuser-Busch InBev	Cons Stapl	Food Bev&Tobacco	BEL

Continue in the next page

Continue from previous page

#	Company Name	Sector	Industry	Country
23	Ashtead Group	Industrials	Capital Goods	GBR
24	ASML Holding	IT	Semicond&SC Equip	NLD
25	Associated British Foods	Cons Stapl	Food Bev&Tobacco	GBR
26	AstraZeneca	Health Care	Pharm Biotech&Lf Sc	GBR
27	Atlas Copco A	Industrials	Capital Goods	SWE
28	Atlas Copco B	Industrials	Capital Goods	SWE
29	Atos	IT	Software&Services	FRA
30	Aviva	Financials	Insurance	GBR
31	AXA	Financials	Insurance	FRA
32	Babcock International Group	Industrials	Comm&Prof Services	GBR
33	BAE Systems	Industrials	Capital Goods	GBR
34	Baloise	Financials	Insurance	CHE
35	Banca Mediolanum	Financials	Diversified Financials	ITA
36	BBVA	Financials	Banks	ESP
37	Banco Comercial Portugues	Financials	Banks	PRT
38	Banco De Sabadell	Financials	Banks	ESP
39	Banco Popular	Financials	Banks	ESP
40	Banco Santander	Financials	Banks	ESP
41	Bankinter	Financials	Banks	ESP
42	Barclays	Financials	Banks	GBR
43	Barratt Developments	Cons Discr	Cons Drbls&Apparel	GBR
44	Barry Callebaut	Cons Stapl	Food Bev&Tobacco	CHE
45	BASF	Materials	Materials	DEU
46	Bayer	Health Care	Pharm Biotech&Lf Sc	DEU
47	Beiersdorf	Cons Stapl	Hhold&Pers Products	DEU
48	Bellway	Cons Discr	Cons Drbls&Apparel	GBR
49	BHP Billiton	Materials	Materials	GBR
50	BIC	Industrials	Comm&Prof Services	FRA
51	BMW	Cons Discr	Auto&Components	DEU
52	BNP Paribas	Financials	Banks	FRA
53	Boliden	Materials	Materials	SWE
54	Bollore	Industrials	Transportation	FRA
55	Boskalis Westminster	Industrials	Capital Goods	NLD
56	Bouygues	Industrials	Capital Goods	FRA
57	BP	Energy	Energy	GBR
58	British American Tobacco	Cons Stapl	Food Bev&Tobacco	GBR
59	British Group	Telecom	Telecom Services	GBR
60	Bunzl	Industrials	Capital Goods	GBR
61	Burberry Group	Cons Discr	Cons Drbls&Apparel	GBR
62	Cap Gemini	IT	Software&Services	FRA
63	Capita	Industrials	Comm&Prof Services	GBR
64	Carlsberg B	Cons Stapl	Food Bev&Tobacco	DNK

Continue in the next page

Continue from previous page

#	Company Name	Sector	Industry	Country
65	Carnival	Cons Discr	Consumer Services	GBR
66	Carrefour	Cons Stapl	Food&Staples Retail	FRA
67	Casino Guichard Perrachon	Cons Stapl	Food&Staples Retail	FRA
68	Centrica	Utilities	Utilities	GBR
69	Christian Dior	Cons Discr	Cons Drbls&Apparel	FRA
70	CNP Assurance	Financials	Insurance	FRA
71	Coloplast B	Health Care	Health Equip&Serv	DNK
72	Colruyt	Cons Stapl	Food&Staples Retail	BEL
73	Commerzbank	Financials	Banks	DEU
74	Compass Group	Cons Discr	Consumer Services	GBR
75	Continental	Cons Discr	Auto&Components	DEU
76	Credit Agricole	Financials	Banks	FRA
77	Credit Suisse	Financials	Diversified Financials	CHE
78	CRH	Materials	Materials	GBR
79	Croda International	Materials	Materials	GBR
80	Daimler	Cons Discr	Auto&Components	DEU
81	Danone	Cons Stapl	Food Bev&Tobacco	FRA
82	Danske Bank	Financials	Banks	DNK
83	Dassault Systemes	IT	Software&Services	FRA
84	Derwent	Real Estate	Real Estate	GBR
85	Deutsche Bank	Financials	Diversified Financials	DEU
86	Deutsche Boerse	Financials	Diversified Financials	DEU
87	Deutsche Lufthansa	Industrials	Transportation	DEU
88	Deutsche Post	Industrials	Transportation	DEU
89	Deutsche Telekom	Telecom	Telecom Services	DEU
90	Diageo	Cons Stapl	Food Bev&Tobacco	GBR
91	DNB	Financials	Banks	NOR
92	DSV B	Industrials	Transportation	DNK
93	Eon	Utilities	Utilities	DEU
94	Easyjet	Industrials	Transportation	GBR
95	EDP	Utilities	Utilities	PRT
96	Eiffage	Industrials	Capital Goods	FRA
97	Electrolux Ser B	Cons Discr	Cons Drbls&Apparel	SWE
98	Elisa	Telecom	Telecom Services	FIN
99	Enagas	Energy	Energy	ESP
100	Endesa	Utilities	Utilities	ESP
101	Enel	Utilities	Utilities	ITA
102	Engie	Utilities	Utilities	FRA
103	Eni	Energy	Energy	ITA
104	Ericsson B	IT	Tech Hardwr&Equip	SWE
105	Essilor International	Health Care	Health Equip&Serv	FRA
106	Euler Hermes Group	Financials	Insurance	FRA

Continue in the next page

Continue from previous page

#	Company Name	Sector	Industry	Country
107	Eurazeo	Financials	Diversified Financials	FRA
108	Eurofins Scientific	Health Care	Pharm Biotech&Lf Sc	FRA
109	Exor	Financials	Diversified Financials	ITA
110	Faurecia	Cons Discr	Auto&Components	FRA
111	FIAT Chrysler Automobiles	Cons Discr	Auto&Components	ITA
112	Fielmann	Cons Discr	Retailing	DEU
113	Fortum	Utilities	Utilities	FIN
114	Fraport AG Frankfurt	Industrials	Transportation	DEU
115	Fresenius Medical Care	Health Care	Health Equip&Serv	DEU
116	Gamesa	Industrials	Capital Goods	ESP
117	GEA Group	Industrials	Capital Goods	DEU
118	Gecina	Real Estate	Real Estate	FRA
119	Generali	Financials	Insurance	ITA
120	Genmab	Health Care	Pharm Biotech&Lf Sc	DNK
121	Getinge B	Health Care	Health Equip&Serv	SWE
122	GKN	Cons Discr	Auto&Components	GBR
123	GlaxoSmithKline	Health Care	Pharm Biotech&Lf Sc	GBR
124	Groupe Bruxelles Lambert	Financials	Diversified Financials	BEL
125	H Lundbeck	Health Care	Pharm Biotech&Lf Sc	DNK
126	Hammerson	Real Estate	Real Estate	GBR
127	Hannover Rueck	Financials	Insurance	DEU
128	HeidelbergCement	Materials	Materials	DEU
129	Heineken	Cons Stapl	Food Bev&Tobacco	NLD
130	Helvetia Holding	Financials	Insurance	CHE
131	Henkel Kgaa Ord	Cons Stapl	Hhold&Pers Products	DEU
132	Hennels and Mauritz B	Cons Discr	Retailing	SWE
133	Hermes International	Cons Discr	Cons Drbls&Apparel	FRA
134	Hexagon B	IT	Tech Hardwr&Equip	SWE
135	Hochtief	Industrials	Capital Goods	DEU
136	Howden Joinery Group	Industrials	Capital Goods	GBR
137	HSBC Holdings	Financials	Banks	GBR
138	Hugo Boss	Cons Discr	Cons Drbls&Apparel	DEU
139	Husqvarna AB B	Cons Discr	Cons Drbls&Apparel	SWE
140	Iberdrola	Utilities	Utilities	ESP
141	Imerys	Materials	Materials	FRA
142	IMI	Industrials	Capital Goods	GBR
143	Imperial Brands	Cons Stapl	Food Bev&Tobacco	GBR
144	Inditex	Cons Discr	Retailing	ESP
145	Industrivarden AB C Free	Financials	Diversified Financials	SWE
146	Infineon Technology	IT	Semicond&SC Equip	DEU
147	Informa	Cons Discr	Media	GBR
148	ING Group	Financials	Banks	NLD

Continue in the next page

Continue from previous page

#	Company Name	Sector	Industry	Country
149	Ingenico Group	IT	Tech Hardwr&Equip	FRA
150	Intertek Group	Industrials	Comm&Prof Services	GBR
151	Intesa Sanpaolo	Financials	Banks	ITA
152	Intesa Sanpaolo Rsp	Financials	Banks	ITA
153	Intu Properties	Real Estate	Real Estate	GBR
154	Investor B Free	Financials	Diversified Financials	SWE
155	ITV	Cons Discr	Media	GBR
156	JC Decaux	Cons Discr	Media	FRA
157	Jeronimo Martins	Cons Stapl	Food&Staples Retail	PRT
158	Johnson Matthey	Materials	Materials	GBR
159	K+S	Materials	Materials	DEU
160	KBC Group	Financials	Banks	BEL
161	Kering	Cons Discr	Cons Drbls&Apparel	FRA
162	Kesko B	Cons Stapl	Food&Staples Retail	FIN
163	Kingfisher	Cons Discr	Retailing	GBR
164	Klepierre	Real Estate	Real Estate	FRA
165	Koninklijke Philips	Industrials	Capital Goods	NLD
166	LafargeHolcim	Materials	Materials	CHE
167	Lagardere Groupe	Cons Discr	Media	FRA
168	Lanxess	Materials	Materials	DEU
169	Legal and General Group	Financials	Insurance	GBR
170	Linde	Materials	Materials	DEU
171	Lindt and Spruengli PC	Cons Stapl	Food Bev&Tobacco	CHE
172	Loyds Banking Group	Financials	Banks	GBR
173	London Stock Exchange Group	Financials	Diversified Financials	GBR
174	Lonza Grp Ag	Health Care	Pharm Biotech&Lf Sc	CHE
175	L'Oreal	Cons Stapl	Hhold&Pers Products	FRA
176	Lundin Petroleum	Energy	Energy	SWE
177	Luxottica	Cons Discr	Cons Drbls&Apparel	ITA
178	Louis Vuitton	Cons Discr	Cons Drbls&Apparel	FRA
179	MAN	Industrials	Capital Goods	DEU
180	Marks and Spencer Group	Cons Discr	Retailing	GBR
181	Mediaset	Cons Discr	Media	ITA
182	Mediobanca	Financials	Banks	ITA
183	Merck Kgaa	Health Care	Pharm Biotech&Lf Sc	DEU
184	Metro	Cons Stapl	Food&Staples Retail	DEU
185	Metso Corporation	Industrials	Capital Goods	FIN
186	Michelin B	Cons Discr	Auto&Components	FRA
187	Morrison (WM) Supermarkets	Cons Stapl	Food&Staples Retail	GBR
188	Muenchener Rueckversicherungs	Financials	Insurance	DEU
189	National Grid	Utilities	Utilities	GBR
190	Natixis	Financials	Banks	FRA

Continue in the next page

Continue from previous page

#	Company Name	Sector	Industry	Country
191	Next	Cons Discr	Retailing	GBR
192	Nokia	IT	Tech Hardwr&Equip	FIN
193	Nokian Renkaat	Cons Discr	Auto&Components	FIN
194	Nordea Bank AB	Financials	Banks	SWE
195	Norsk Hydro	Materials	Materials	NOR
196	Novartis (REGD)	Health Care	Pharm Biotech&Lf Sc	CHE
197	Novo-Nordisk B	Health Care	Pharm Biotech&Lf Sc	DNK
198	Novozymes AS	Materials	Materials	DNK
199	Old Mutual	Financials	Insurance	GBR
200	Orange	Telecom	Telecom Services	FRA
201	Orkla A	Cons Stapl	Food Bev&Tobacco	NOR
202	Orpea	Health Care	Health Equip&Serv	FRA
203	Pargesa Holding	Financials	Diversified Financials	CHE
204	Pearson	Cons Discr	Media	GBR
205	Pernod Ricard	Cons Stapl	Food Bev&Tobacco	FRA
206	Persimmon	Cons Discr	Cons Drbls&Apparel	GBR
207	Peugeot	Cons Discr	Auto&Components	FRA
208	Provident Financial	Financials	Diversified Financials	GBR
209	Prudential	Financials	Insurance	GBR
210	PSP Swiss Property	Real Estate	Real Estate	CHE
211	Publicis Groupe	Cons Discr	Media	FRA
212	Randstad Holdings	Industrials	Comm&Prof Services	NLD
213	Rational	Industrials	Capital Goods	DEU
214	Reckitt Benckiser Group	Cons Stapl	Hhold&Pers Products	GBR
215	Recordati	Health Care	Pharm Biotech&Lf Sc	ITA
216	Red Electrica Corp	Utilities	Utilities	ESP
217	RELX	Industrials	Comm&Prof Services	GBR
218	Relx NV	Industrials	Comm&Prof Services	NLD
219	Remy Cointreau	Cons Stapl	Food Bev&Tobacco	FRA
220	Renault	Cons Discr	Auto&Components	FRA
221	Rentokil Initial	Industrials	Comm&Prof Services	GBR
222	Repsol	Energy	Energy	ESP
223	Richemont A (BR)	Cons Discr	Cons Drbls&Apparel	CHE
224	Rolls-Royce Holdings	Industrials	Capital Goods	GBR
225	Royal Bank of Scotland Group	Financials	Banks	GBR
226	Royal DSM	Materials	Materials	NLD
227	Royal Dutch Shell B	Energy	Energy	GBR
228	Royal KPN	Telecom	Telecom Services	NLD
229	RSA Insurance Group	Financials	Insurance	GBR
230	RWE	Utilities	Utilities	DEU
231	Safran	Industrials	Capital Goods	FRA
232	Sage Group	IT	Software&Services	GBR

Continue in the next page

Continue from previous page

#	Company Name	Sector	Industry	Country
233	Sainsbury (J)	Cons Stapl	Food&Staples Retail	GBR
234	Saipem	Energy	Energy	ITA
235	Sampo Oyi	Financials	Insurance	FIN
236	Sandvik AB	Industrials	Capital Goods	SWE
237	SAP	IT	Software&Services	DEU
238	Sartorius Stedim Biotech	Health Care	Health Equip&Serv	FRA
239	SCA B	Cons Stapl	Hhold&Pers Products	SWE
240	Schibsted A	Cons Discr	Media	NOR
241	Schindler Holding (Cert)	Industrials	Capital Goods	CHE
242	Schindler Holding (Reg)	Industrials	Capital Goods	CHE
243	Schneider Electric	Industrials	Capital Goods	FRA
244	Schroders	Financials	Diversified Financials	GBR
245	Scor	Financials	Insurance	FRA
246	SEB	Cons Discr	Cons Drbls&Apparel	FRA
247	Securitas AB B	Industrials	Comm&Prof Services	SWE
248	Segro	Real Estate	Real Estate	GBR
249	SGS	Industrials	Comm&Prof Services	CHE
250	Shire	Health Care	Pharm Biotech&Lf Sc	GBR
251	Siemens	Industrials	Capital Goods	DEU
252	Skandinaviska Enskilda Banken A	Financials	Banks	SWE
253	Skanska B	Industrials	Capital Goods	SWE
254	SKF B	Industrials	Capital Goods	SWE
255	Sky	Cons Discr	Media	GBR
256	Smith and Nephew	Health Care	Health Equip&Serv	GBR
257	Smith (DS)	Materials	Materials	GBR
258	Smiths Group	Industrials	Capital Goods	GBR
259	SociÃ©tÃ© GÃ©nÃ©rale	Financials	Banks	FRA
260	Sodexo	Cons Discr	Consumer Services	FRA
261	Solvay A	Materials	Materials	BEL
262	SSE	Utilities	Utilities	GBR
263	St Gobain (Cie de)	Industrials	Capital Goods	FRA
264	St James Place	Financials	Insurance	GBR
265	Standard Chartered	Financials	Banks	GBR
266	Statoil ASA	Energy	Energy	NOR
267	STMicroelectronics	IT	Semicond&SC Equip	FRA
268	Stora Enso R	Materials	Materials	FIN
269	Straumann Hldg N	Health Care	Health Equip&Serv	CHE
270	Suedzucker	Cons Stapl	Food Bev&Tobacco	DEU
271	Sulzer	Industrials	Capital Goods	CHE
272	Svenska Handelsbnk B	Financials	Banks	SWE
273	Svenska Handelsbnk A	Financials	Banks	SWE
274	Swatch Group AG Reg	Cons Discr	Cons Drbls&Apparel	CHE

Continue in the next page

Continue from previous page

#	Company Name	Sector	Industry	Country
275	Swedbank AB Series A	Financials	Banks	SWE
276	Swedish Match	Cons Stapl	Food Bev&Tobacco	SWE
277	Swiss Re	Financials	Insurance	CHE
278	Syngenta	Materials	Materials	CHE
279	Tate and Lile	Cons Stapl	Food Bev&Tobacco	GBR
280	Taylor Wimpey	Cons Discr	Cons Drbls&Apparel	GBR
281	Technip	Energy	Energy	FRA
282	Tele2 AB	Telecom	Telecom Services	SWE
283	Telecom Italia	Telecom	Telecom Services	ITA
284	Telecom Italia Rsp	Telecom	Telecom Services	ITA
285	Telefonica	Telecom	Telecom Services	ESP
286	Telenor	Telecom	Telecom Services	NOR
287	Teleperformance	Industrials	Comm&Prof Services	FRA
288	TeliaSonera	Telecom	Telecom Services	SWE
289	Tesco	Cons Stapl	Food&Staples Retail	GBR
290	Thales	Industrials	Capital Goods	FRA
291	Thyssen Krupp	Materials	Materials	DEU
292	Total	Energy	Energy	FRA
293	Travis Perkins	Industrials	Capital Goods	GBR
294	Trelleborg Ab Ser B	Industrials	Capital Goods	SWE
295	UCB Cap	Health Care	Pharm Biotech&Lf Sc	BEL
296	Umicore	Materials	Materials	BEL
297	Unibail-Rodamco	Real Estate	Real Estate	NLD
298	Unicredit	Financials	Banks	ITA
299	Unilever	Cons Stapl	Hhold&Pers Products	GBR
300	UnipolSai	Financials	Insurance	ITA
301	United Inter Na	IT	Software&Services	DEU
302	United Utilities Group	Utilities	Utilities	GBR
303	UPM-Kymmene	Materials	Materials	FIN
304	Valeo	Cons Discr	Auto&Components	FRA
305	Veolia Environnement	Utilities	Utilities	FRA
306	Vestas Wind Systems	Industrials	Capital Goods	DNK
307	Vinci	Industrials	Capital Goods	FRA
308	Vivendi	Cons Discr	Media	FRA
309	Vodafone	Telecom	Telecom Services	GBR
310	Volkswagen	Cons Discr	Auto&Components	DEU
311	Volvo B	Industrials	Capital Goods	SWE
312	Vopak	Energy	Energy	NLD
313	Weir Group	Industrials	Capital Goods	GBR
314	Wendel	Financials	Diversified Financials	FRA
315	Williamd Demant Holding	Health Care	Health Equip&Serv	DNK
316	William Hill	Cons Discr	Consumer Services	GBR

Continue in the next page

Continue from previous page

#	Company Name	Sector	Industry	Country
317	Wolseley	Industrials	Capital Goods	GBR
318	Wood Group (John)	Energy	Energy	GBR
319	WPP	Cons Discr	Media	GBR
320	Zardoya Otis	Industrials	Capital Goods	ESP
321	Zodiac Aerospace	Industrials	Capital Goods	FRA
322	Zurich Insurance Group	Financials	Insurance	CHE

Concludes from previous page

Chapter 2

A Sectoral Analysis of Volatility Connectedness

2.1 Introduction

The 2008 financial breakdown demonstrated that seemingly isolated failures can soon lead to domino effects and generalized panic whenever the affected company is linked in critical ways to other companies in the network. Especially because of poor risk management (where too optimistic expectations and moral hazard likely induced excessive leverage) banks and insurance companies have been recognized as the main culprits. However, the intermediary function of financial companies and their subsequent key network position might have played an important role as well. In this work, we study the evolution of the US sectoral network structure, and particularly see whether it contributed to worsen the already dire situation during the financial crisis.

We estimate a network of sectoral volatilities using SPDR Exchange Traded Funds indexes which track 9 US broadly defined sectors. In order to estimate the network, we use Generalized Forecast Error Variance Decompositions in the spirit of Diebold and Yılmaz (2009) and their later contributions. However, our estimation strategy is completely different from their, because we directly model the evolution of VAR parameters and estimate the model in a single Kalman filter sweep. The standard rolling windows approach typically used in Diebold and Yılmaz articles attaches the same weight to all the observations in a given window, and we show that this leads to overly rigid estimates that do not promptly react to quick parameters change. We apply Koop and Korobilis (2013) model to address this issue. The authors introduce forgetting factors and built-in model switching features, which result in a flexible model that can better accommodate periods of sudden changes. Furthermore, we improve their methodology to estimate (instead of simply fixing) a key parameter of their model, and we show that this allows to appropriately pin down wide connectedness movements during the first part of our sample.

Our results show that this approach yields indeed superior estimates of connectedness, which presents two distinct regimes. Between 1999 and the 2007 financial crisis connectedness is on average lower but more volatile, whereas after the crisis it settles on permanently higher and less volatile levels. Also, sectoral spillovers are tightly concentrated around the mean, and this is especially true during highly stressed periods such as September 2008. This key result suggests that this particular network structure might have contributed to spill volatility between sectors at a faster rate, and this raises interesting questions regarding the underlying causal links between variations in network structure and fluctuations in volatility.

The literature on financial networks is relatively recent. One of the first theoretical contributions in economic and financial networks theory traces back to Allen and Gale (2000), where the authors build a simple model and show that a star (i.e. a fully connected) network in the banking system is less prone to create contagion episodes than a ring-like network where *each* company is connected to only two other companies. Gabaix (2011) provides evidence that when firms' size distribution is fat-tailed, idiosyncratic shocks are a source of aggregate fluctuations, thus highlighting the importance of large firms in the macroeconomic structure of a country. Acemoglu, Carvalho, Ozdaglar and Tahbaz-Salehi (2012) studies the relationships between the network structure of the economy and aggregate volatility, and show that in the presence of certain input-output linkages microeconomic idiosyncratic shocks can lead to aggregate fluctuations. Elliott, Golub and Jackson (2014) studies the effect of integration (greater dependence on counterparties) and diversification (more counterparties per organization) on the possibility of cascades. Increasing diversification allows to be better insured against another agent's failure. On the other hand, a low integration allows the financial organization to depend less on others, whereas a high one reduces the sensitivity to own investment outcomes. Acemoglu, Ozdaglar and Tahbaz-Salehi (2015) finds that the size of shocks matters: a densely inter-connected financial network promotes stability only when shocks are moderately small, whereas it spreads cascade failures whenever very large shocks occur.

Our results regarding the particular evolution of the network structure highlight the fact that the desirability of a given structure depends on the type of connections one is analyzing. Indeed, whereas a star network could be more robust when considering purely contractual relationships, it might well be a curse in settings like ours where high inter-connectedness likely leads to faster volatility spillovers, and the opposite is true for a ring network.

On the empirical side, Forbes and Rigobon (2002) proposes a measure of contagion that is free from the biases coming from time-varying volatility, and distinguishes contagion from inter-dependence. Lopes and Migon (2002) uses a factor stochastic volatility approach to study the transmission of financial crises across countries. Diebold and Yılmaz (2009) measures returns and volatility spillovers in global equity markets and show that connectedness in returns exhibits an increasing trend over time (market integration) whereas volatility spillovers display no trend but are more volatile. Later papers[1] apply the same methodology to different economic and financial environments. Billio, Getmansky, Lo and Pelizzon (2012) estimates Granger-causality networks on hedge funds, banks, broker/dealers, and insurance companies and, among other results, shows that the banking sector is crucial in transmitting shocks to other institutions. Brownlees, Nualart and Sun (2015) uses Lasso algorithms to regularize inverse realized covariance estimators of log-prices, thus uncovering the partial correlation network structure of returns. Abbassi, Brownlees, Hans and Podlich (2016) uses both a market-based and a proprietary dataset to develop measures of connectedness, and show that market-based ones work comparatively well as a tool to monitor banks credit riskiness. Barigozzi and Brownlees (2017) introduces a new Lasso-based algorithm called NETS to analyze large panels of volatility measures and estimate their inter-connectedness, producing both a directed and an undirected graph. Their methodology is shown to outperform several forecasting models.

Finally (on the pure networks literature) Newman (2003) is an introductory paper which describes the functioning of complex networks and provides details on some of the most important networks-related metrics. Newman (2001) generalizes the concepts of degree centrality, closeness centrality, and betweenness centrality to weighted networks. Jackson (2008) and Newman (2010) are introductory textbooks on the subject.

2.2 Data and Methodology

As we anticipated, we use data on the 9 sectoral Standard & Poor's Depositary Receipts (SPDR) Exchange Traded Funds for the US, which are available at a daily frequency starting from December 22^{nd}, 1998. The sectors are Consumer Discretionary, Consumer Staples, Energy, Financials, Health Care, Industrials, IT,

[1] See Diebold and Yılmaz (2012), Diebold and Yılmaz (2013), Diebold and Yılmaz (2014), Diebold and Yılmaz (2015), Demirer, Diebold, Liu and Yılmaz (2017).

Materials and Utilities. Figure 2.1 plots their respective indexes over time.

We compute volatility with the high-low range developed in Parkinson (1980), which (as opposed to standard GARCH-like measures) exploits intra-day information (high and low prices) and reveals volatility episodes that would otherwise not be visible if we would only use close prices as is done in GARCH models. Figure 2.2 plots our volatility measures together with a GARCH(1,1) estimate. As one can see, the two series largely co-move, although the one retrieved with the GARCH model is particularly more persistent.

As anticipated, we use the methodology developed in Diebold and Yılmaz (2009)[2]. Consider a covariance-stationary N-variable reduced-form TVP-VAR(p),

$$\mathbf{Y}_t = \sum_{i=1}^{p} \mathbf{\Phi}_{it}\mathbf{Y}_{t-i} + \boldsymbol{\epsilon}_t, \tag{2.1}$$

where $\boldsymbol{\epsilon}_t \sim \mathcal{N}(0, \mathbf{\Sigma}_t)$. The MA($\infty$) representation is

$$\mathbf{Y}_t = \sum_{i=0}^{\infty} \mathbf{A}_{it}\boldsymbol{\epsilon}_{t-i}. \tag{2.2}$$

The estimation of the network exploits the so-called Generalized Forecast Error Variance Decompositions (GFEVD from now on) developed in Koop, Pesaran and Potter (1996) for non-linear models, and applied to the linear case in Pesaran and Shin (1998). In a nutshell, the amount of fluctuations of variable i (j) explained by variable j (i) is used as a measure of the degree of (directional) connectivity from (to) variable j to (from) variable i. Diebold and Yılmaz (2014) shows that estimating the spillover effects with GFEVD is equivalent to estimating a weighted directed graph. Denoting the H-step-ahead GFEVD by $\theta_{ij,t}(H)$, for $H = 1, 2, \ldots,$ we have (dropping the subscript t for ease of notation)

$$\theta_{ij} = \frac{\frac{1}{\sigma_{jj}} \sum_{h=0}^{H-1} (\mathbf{e}_i' \mathbf{A}_h \mathbf{\Sigma} \mathbf{e}_j)^2}{\sum_{h=0}^{H-1} (\mathbf{e}_i' \mathbf{A}_h \mathbf{\Sigma} \mathbf{A}_h' \mathbf{e}_i)}, \tag{2.3}$$

where $\mathbf{e}_i$ is a selection vector that equals 1 at the i-th element and 0 otherwise.

Diebold and Yılmaz papers estimate both the static network (assuming fixed parameters throughout) and a dynamic one, where the VAR is recursively estimated using rolling windows. This approach has the advantage of being easy to estimate and model-free on both the parameters and residual covariance matrix evolution.

[2]Note that the original methodology relies on short term structural identification (Cholesky decompositions), whereas here we strictly follow Diebold and Yılmaz (2012) and do not impose any restriction on contemporaneous relationships.

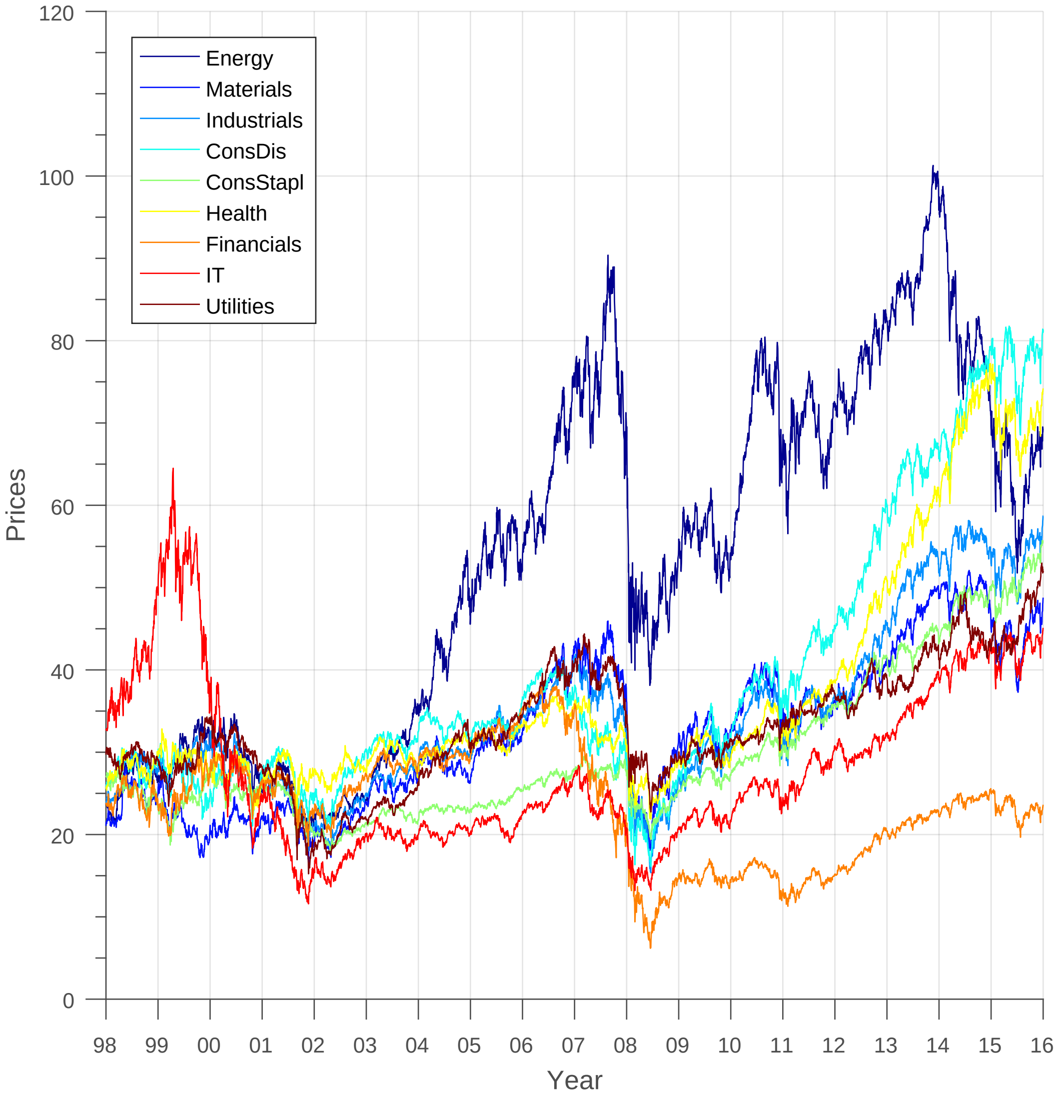

Figure 2.1: This figure plots the Standard & Poor's Depositary Receipts (SPDR) Exchange Traded Funds, which track 9 broadly defined US sectors. The series are highly correlated and most of them present a remarkably quick post-crisis re-bound, which lead the indexes to reach historically high values in the last part of the sample.

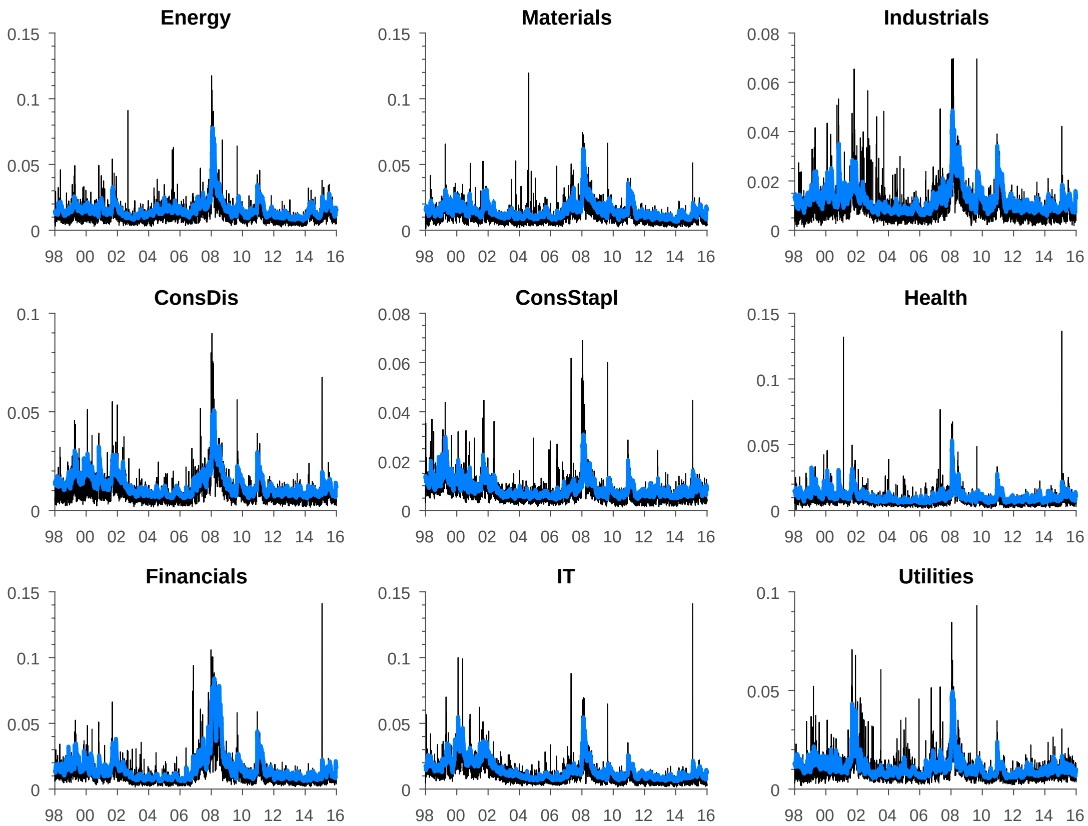

Figure 2.2: The figure compares range-based (black) and GARCH (blue) volatilities. Apparently, high-low range volatilities are considerably more noisy due to the fact that they incorporate intra-day information, and that they do not rely on any parametric assumption.

This comes at a cost however, namely that estimates obtained at time t are not a time-weighted moving average of the past w observations (where w is window width). This implies that time t innovation's weight is as low as $\frac{1}{w}$, and the same is true for time $t-w+1$ observation's weight. Since we are dealing with highly-parametrized objects, window width must be set to around $150-200$ observations in order not to incur in the curse of dimensionality problem. Therefore, the individual weight is approximately equal to 0.005-0.0067, whereas we would like to have a significantly higher one on current time t observation.

We address this issue by estimating a model that imposes some structure on the evolution of parameters and residual covariance matrix. The literature on time-varying-parameters VARs is now relatively large, and two choices would be Primiceri (2005) and Koop and Korobilis (2013). The first assumes both parameters and volatility to evolve as a random walk and estimate the model via MCMC methods. Because of the computational burden involved, this implies that the dimension of the VAR cannot be large. Instead, Koop and Korobilis (2013) proposes an estimation strategy that does not need to simulate the posterior distribution, thus generating results with incomparably lower estimation times. In particular, the authors estimate the time-varying covariance matrix with an Exponentially-Weighted Moving Average (EWMA) model, and the evolution of the parameters with the Kalman filter. Even more interestingly, they allow for model switching, which means that parameters are allowed to be relatively constant during normal periods, whereas they can change in a more abrupt manner during volatility storms. This is an important feature of the model that addresses the issue of estimating time-varying parameters when in fact they could be fixed over some period of time. From a practical standpoint, Koop and Korobilis (2013) VAR seems more appealing both methodologically and computationally, hence we opt for it.

Koop and Korobilis (2013) uses four high-level model parameters to be either chosen or estimated from the data. The first is (using the authors' notation) the parameters forgetting factor λ, which governs the instantaneous flexibility of the regression parameters. In other words, when λ is set to be equal to one parameters are fixed, whereas progressively lower values increase the probability for the regression parameters to vary by higher amounts from one period to the other. Instead of choosing a grid of values for λ, Koop and Korobilis (2013) estimates it following Park, Jun and Kim (1991):

$$\lambda_t = \lambda_{\min} + (1 - \lambda_{\min}) L^{-\mathrm{NINT}(\tilde{\epsilon}_t' \tilde{\epsilon}_t)}, \tag{2.4}$$

where $\tilde{\boldsymbol{\epsilon}}_t$ is the one-step ahead prediction error produced by the Kalman filter, and NINT$(\cdot)$ rounds to the nearest integer. Also, $\lambda_{\min} = 0.96$, and $L = 1.1$, which ensure λ to lie between $\lambda_{\min}$ and 1.

As we said, the covariance matrix is assumed to evolve with an EWMA law of motion:

$$\hat{\boldsymbol{\Sigma}}_t = \kappa \hat{\boldsymbol{\Sigma}}_{t-1} + (1 - \kappa)\hat{\boldsymbol{\epsilon}}_t \hat{\boldsymbol{\epsilon}}_t', \tag{2.5}$$

where κ is the covariance matrix forgetting factor κ, which is the second model parameter we have to select. Koop and Korobilis (2013) follows RiskMetrics (1996) which suggests values for κ in the interval $[0.94, 0.98]$. Since $\boldsymbol{\Sigma}_t$ is the covariance matrix of volatility, we expect it to present less pronounced swings as compared to the covariance matrix of returns. Therefore, we choose $\kappa = 0.98$ which yields reasonably time-varying estimates for $\boldsymbol{\Sigma}_t$. We then use the first 60 observations to compute the initial covariance matrix $\hat{\boldsymbol{\Sigma}}_0$.

Third, we use Koop and Korobilis (2013) Minnesota prior choice for our regression parameters. Therefore, we set $\mathbb{E}(\boldsymbol{\beta}_0) = 0$, thus shrinking parameters towards zero. Also, $\mathbb{V}(\boldsymbol{\beta}_0) = \underline{\mathbf{V}}$, and we define $\underline{\mathbf{V}}_i$ to be the diagonal elements of $\underline{\mathbf{V}}$. Then, we set

$$\underline{\mathbf{V}}_i = \begin{cases} \dfrac{\gamma}{r^2} & \text{for coefficients on lag } r, \quad \text{where } r = 1, \ldots, p, \\ \underline{a} & \text{for the intercepts,} \end{cases} \tag{2.6}$$

where p is lag length. In this way, parameters that associate observations that are relatively more distant between each other are shrinked more. The key parameter to be chosen is γ, which governs the general degree of parameters shrinkage. We estimate γ from a wide grid, that is $\gamma \in \{10^{-5}, 0.001, 0.005, 0.01, 0.05, 0.1\}$. On the other hand, we do not want to impose any informative prior to the intercepts, so we set $\underline{a} = 10^2$.

Finally, we have to choose the model switching forgetting factor α, which governs how past forecasting performances are weighted during the whole sample period. As an example, when $\alpha = 0.99$ forecast performance 20 periods in the past receives 80% as much weight as forecast performance last period, whereas when $\alpha = 0.95$ the same percentage decreases to about 35%. We extend the baseline setting in Koop and Korobilis (2013) to also estimate α, which we choose within the simple grid $\alpha \in \{0.95, 0.99\}$. The way we estimate α is the same as the method proposed by Koop and Korobilis (2013) to estimate κ and γ, that is we maximize geometrically weighted predictive densities over the entire set of different models. Recall from the paper that a single period forecasting performance is measured

Table 2.1: Full-sample connectedness

	Ener	Matrls	Ind	CD	CS	Hlth	Fincls	IT	Utils	**FROM**
Ener	.40	.11	.07	.08	.06	.06	.09	.08	.05	.60
Matrls	.09	.28	.11	.11	.07	.06	.13	.11	.04	.71
Ind	.05	.10	.28	.13	.07	.08	.12	.12	.05	.72
CD	.05	.09	.11	.26	.08	.09	.15	.13	.04	.74
CS	.05	.07	.08	.10	.31	.09	.11	.13	.06	.70
Hlth	.05	.06	.10	.11	.10	.31	.10	.12	.05	.69
Fincls	.05	.10	.10	.13	.08	.08	.30	.12	.04	.70
IT	.05	.08	.10	.12	.09	.08	.11	.33	.04	.67
Utils	.06	.06	.08	.08	.09	.08	.08	.08	.39	.61
TO	.46	.67	.75	.86	.63	.61	.90	.87	.37	.68
NET	-.14	-.03	.03	.12	-.06	-.08	.20	.20	-.24	

The table shows the GFEVD obtained by estimating a VAR over the whole sample. This representation can also be interpreted as a weighted and directed network of connections between US sectors. The weights are given by the portions of variations explained, whereas the in-degree (out-degree) spillovers have to be read from rows (columns).

with the predictive density $p_j(\mathbf{Y}_t|\mathbf{Y}_{t-1})$, where (in our case) j indexes models with all the possible combinations of κ, γ and α. The model that yields the maximum weighted predictive density is recursively chosen at each time period. As we will see, estimating α proves to be beneficial in our application.

2.3 Results

Section 2.3.1 presents average results for the whole sample, whereas the ensuing sections describe the results on the dynamic evolution of sectoral inter-connectedness. Our baseline model uses 4 lags and 10-days forecast horizon.

2.3.1 Average (Static) Volatility Connectedness

Table 2.1 is the connectedness table obtained from a fixed-parameters VAR, and it thus represents results that are averaged over the whole sample period. Each number in row i, column j is the portion of variation of variable i explained by variable j. The "**FROM**" column is the amount of variation of variable i explained by all the other variables, whereas the "**TO**" row is a sum of the portions of variation that variable i is able to explain in all the other variables. "**NET**" is the difference between "**TO**" and "**FROM**" and captures the extent to which a sector is a net sender of connections.

First of all, we see that overall connectedness during the whole period is remarkably high and equal to 68%. As a comparison, the only sample among the ones explored in Diebold and Yılmaz articles that exceeds this connectedness level is the one that considers U.S. financial institutions connectedness (78%).

From a network perspective, the level of connectivity is tantamount to the average degree centrality, or the mean of the degree distribution, i.e. a location measure. Nonetheless, the degree distribution as a whole (not only its mean) can give interesting insights, and our directed network is peculiar because it has two degree distributions (in-degree and out-degree) that share the same mean (which is equal to overall connectedness). A rapid inspection of the total "**FROM**" and "**TO**" measures shows that not only (as we already have commented) the sectors are highly inter-connected, but they are also similar in terms of spillovers received and given. The two degree distributions are indeed heavily concentrated around their mean, which allows to hypothesize that the network structure of the U.S. sectors (at least from a financial perspective) is likely to resemble a fully connected network. This first result (high connectedness and low dispersion of individual level of connectivity) is intuitive because we are considering sectors within the same economy, which means the input-output linkages are likely to be tight and permanent and the market is strongly integrated. Fears about the redditivity of one particular sector can easily transmit towards others because of the existence of contractual relationships.

We now analyze more in detail some of the measures we have obtained. Not surprisingly, the financial sector spills the most volatility over the other ones (90%). Being the key transmission channel of government policies and the sector that allows counterparties to settle their contracts, its centrality is easily understood. Moreover, the size of this sectors' constituents together with their high leverage can amplify any movement in their stock value and trigger even higher volatility phenomenons. Finally, it is important to note that the real estate market is included in this sector, which implies that the 2007 sub-prime crisis might have well increased the sectors' importance in spilling over volatility. Interestingly, both IT and Consumer Discretionary sectors also have a high out-degree centrality (87% and 86% respectively). In particular, the relevance of the Consumer Discretionary sector is in accordance with the fact that US GDP consumption component fluctuated between 65% and 68% during our sample, a considerably higher share as compared to the rest of the world. Consumer Discretionary includes items as media, retail, hotels, restaurants and leisure, textiles, apparel and luxury goods, household durables, automobiles, auto components, distributors, leisure equipment and products, diversified consumer services. It then potentially serves as a good benchmark to gauge the extent to which US consumers are spending and thus reviving the economy and the other sectors who depend a lot on consumer behavior. Fears about companies that directly serve consumers can quickly run onto the other sectors.

Energy and Utilities are the sectors whose idiosyncratic movements contribute the most to their overall volatility fluctuations (40% and 39% respectively) and at the same time they are the ones with the lowest out-degree spillovers. With net spillovers of -14% and -24% respectively, they are the sectors whose volatility fluctuations are mainly driven by the volatilities in the other sectors. Consumer Discretionary is the sector with the highest spillovers received, which points to the fact that there also exists a feedback effect whereby the Consumer Discretionary sector is also heavily affected by developments in the rest of the economy.

Finally, we see that Industrials, Consumer Discretionary, Financials, and IT form a sectoral cluster where pairwise inter-dependencies are higher as compared to those outside it. Nevertheless, these connections are not dramatically stronger than the ones outside the cluster, meaning that its economic significance is questionable.

2.3.2 Dynamic Volatility Connectedness

Both the Kalman filter approach (as explained in Section 2.2) and the standard rolling windows one used in (e.g.) Diebold and Yılmaz (2009) produce a set of parameters estimates and covariance matrices for each time period. We are then able to calculate daily connectedness metrics (as in Table 2.1) by computing GFEVD for all periods, thus obtaining dynamic estimates that might give further insights on whether and how the network structure has changed during our sample period.

Figure 2.3 plots the overall level of connectedness for both methods, and we find that the two series generally move together. However, a key difference exists whereby modeling the evolution of parameters with Koop and Korobilis (2013) method generates more volatile connectedness estimates. This result corroborates the fact that the rolling window approach is overly simplistic, and the problem with equally weighting all the periods within a given window is more evident during times of more volatile connectedness. Indeed, while we observe that our model promptly adapts, rolling windows (by assigning the same weight to the most distant and most recent observations) inevitably provide too sticky and consequently less precise connectedness estimates. Therefore, we strongly advocate replacing rolling windows with Kalman filtering techniques as in Koop and Korobilis (2013) to estimate this class of models. Overall, US sectoral connectedness features two distinct behaviors. On the one hand, the pre-2007 financial crisis period is characterized by (on average) lower but more volatile connectedness, whereas the opposite is true for the post-crisis period, where connectedness stabilizes at higher levels. Also, the last five years of the sample display a gently decreasing trend in connectedness.

Prior to the burst of the dot-com bubble we observe the most pronounced spike in overall connectedness in a short period of time, i.e. the period where the euphoria for tech companies was at its highest, and total connectedness increased from less than 30% at the end of 1999 to 70% in March 2000. Surprisingly, from there on every sector started to move more independently, with a consequent sharp decrease in overall connectedness. The lowest value is reached in October 2000 (25%) when a very volatile period that started in March 2000 ended, and from there on tech companies started their definitive collapse, with connectedness bouncing back to higher values. Importantly, the ensuing period has been characterized by a series of important events, with the March-November 2001 economic recession being an underlying threat to US economic stability, which together with the ongoing Enron scandal (culminated in December 2001 with the failure of the company) could be the main reason why we observe a volatile but positive trend in connectedness. The 9/11 terrorist attacks and the beginning of the war in Afghanistan one month later surprisingly led connectedness to decrease from 60% to 40%. On July 21^{th}, 2002 Worldcom (the second largest long distance telephone company in the US) files for the (at that time) largest bankruptcy in US history after being involved in a massive accounting fraud. This event led to a jump in connectedness up to a remarkable 75%. From there on connectedness decreased erratically until it reached 40% in October 2006.

The period between the end of 2006 and mid-2009 marked the transition between the two connectedness regimes identified earlier. Indeed, total connectedness increased rapidly at the end of 2006 when foreclosures rates started to apparently get higher than during normal times. When New Century (largest lender in the US) went under Chapter 11 (April 2007), overall connectedness trended further upwards to unprecedented levels. On February 7^{th}, 2008 the Senate passed the Economic Stimulus Act, a $170 billion package of tax rebates for low and middle income taxpayers, and corporate tax incentives. This could have contributed to push connectedness downwards. Indeed, we observe a marked drop in connectedness from 80% to 55% at the end of August 2008, and Bear Stearns bail out (March 14^{th}, 2008) did not lead to higher sectoral connectedness. Unfortunately however, September 2008 was one of the darkest moments in US financial and

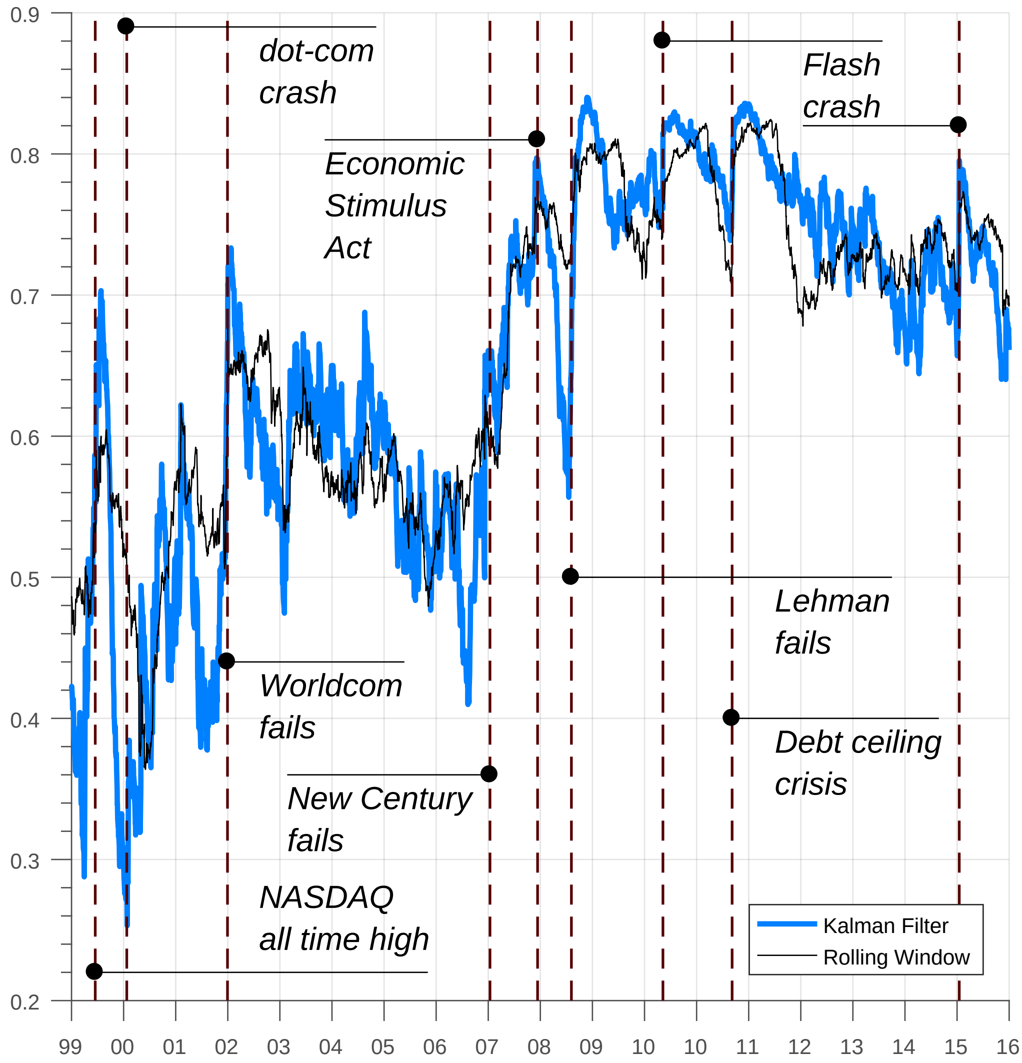

Figure 2.3: This figure clarifies the stark differences between the TVP-VAR estimated with the Kalman filter against the one obtained from standard rolling windows, and confronts the responsiveness of the two connectedess series with important episodes occurred during the sample. Apparently, our estimation strategy is able to pin down marked swings that rolling windows cannot detect.

economic history, and it led to the definitive change to the new regime of high and stable connectedness. On September 7^{th}, 2008 Fannie Mae and Freddie Mac are placed into conservatorship of the FHFA (Federal Housing Finance Agency). On September 15^{th} Lehman Brothers files for the largest bankruptcy in US history, leading markets to panic throughout the whole world. The day after (September 16^{th}) AIG is bailed out with an \$85 billion two-year loan. On September 25^{th} Washington Mutual (largest savings and loan bank, with \$307 billion in assets) goes bust and its remaining assets are sold to JP Morgan Chase for \$1.9 billion. After these catastrophic events, on October 3^{rd} President George W. Bush signs into law the TARP (Troubled Asset Relief Program), a \$700 billion program that allowed financial institutions to sell toxic assets to the US government. This step did not prevent connectedness to jump to its all-time high of almost 85%, a magnitude that shows how much the stock market was out of control during those days. This historical high level is reached in December 2008, when the Wall Street firm Bernard L. Madoff Investment Securities LLC goes bust in the largest accounting fraud in American history.

The beginning of 2009 is characterized by two important policy events. On the one hand the zero lower bound on the target interest rate is reached, and on the other Barack Obama signs the American Recovery and Reinvestment Act, a \$787 billion stimulus package mainly made of government spending measures. These measures lowered total connectedness by a modest 10%. In June 2009 the recession (which officially began in December 2007) ended, and the stock market stabilized. Nonetheless, as we said, total connectedness did not go back to normal levels and remained high throughout the rest of the sample, with a post-crisis minimum of 64% in June 2016.

The period going from July 2010 to June 2011 is characterized by a reduction in overall connectedness, and the most pronounced decline occurs after the Fed announces a second round (QE2) of \$600 billion bond purchases to encourage economic growth.

On July 31^{th}, 2011 the debt ceiling crisis is resolved with the Budget Control Act of 2011, whereby Republicans (who controlled the House of Representatives) agreed to raise the debt ceiling in exchange for spending cuts. The agreement came only two days before the ceiling would have been reached, and it was accompanied by the most volatile week after the 2008 financial crisis. During those days connectedness skyrocketed again to 82%, and increased up to almost 84% around the days of the first ever downgrade of US government debt from Standard & Poor's (from AAA to AA+). This level is indeed very close to the one reached during the days of financial failures and bail outs at the end of 2008.

Surprisingly enough, the approach of the fiscal cliff (which would have been reached on January 1^{st}, 2013) and the debates that preceded the enactment of the American Taxpayer Relief Act (ATRA) on the very same day as the one of the deadline did not lead to sector volatilities being more connected between each other. Similarly, no significant movement is observed after the announcement of the third round of bond purchases (QE3) on September 13^{th}, 2012, and after the resolution of the second debt ceiling crisis at the end of 2013.

The most significant largest spike in the final part of our sample is in Summer 2015, when concerns over Chinese economy spilled over to U.S. companies, and where uncertainty over U.S. monetary policy probably played a role as well. In particular, on August 24^{th} an enormous lack of liquidity coupled with herding behavior led the Dow Jones to lose more than 1,000 points in the first five minutes of trading. During this period volatility connectedness spiked up from 66% to almost 80% in a few days. Finally, on December 15^{th}, 2015 the Fed eventually increased target rates by 25 basis points, and after this decision a downward trend in connectedness started which led it to step down to around 65%.

The dynamics of overall connectedness are useful in providing a general picture of how connections have changed throughout the sample period. However, we also are interested in the variations of the degree distributions, and in analyzing whether certain sectors have been more central than others during given periods.

2.3.3 Directional Dynamic Volatility Connectedness

Figure 2.4 plots (as an illustrative example) the spillovers received by the Materials sector, although pictures for almost all the other sectors are fundamentally similar. We immediately notice that, when we look at totally disaggregated data, idiosyncratic components contribute the most to overall fluctuations. Also, we can observe a decreasing trend in idiosyncratic movements over time, which go from a (volatile) average of approximately 50% at the beginning of the sample to about 25% at the end of the sample. This result is entirely consistent with the one we found for total connectedness, and it reflects the decreased independence of every sector's volatility fluctuations from other sectors' movements. The only sectors that are an exception to this pattern are Energy and Utilities, which regained more independent fluctuations starting from 2011 approximately.

While results on total connectedness are clear, the same cannot be said for disaggregated ones, i.e. results regarding in-degree, out-degree, and net movements. Specifically, not only (as we saw in the previous section) connectedness is on average higher than in many cases analyzed in Diebold and Yılmaz applications,

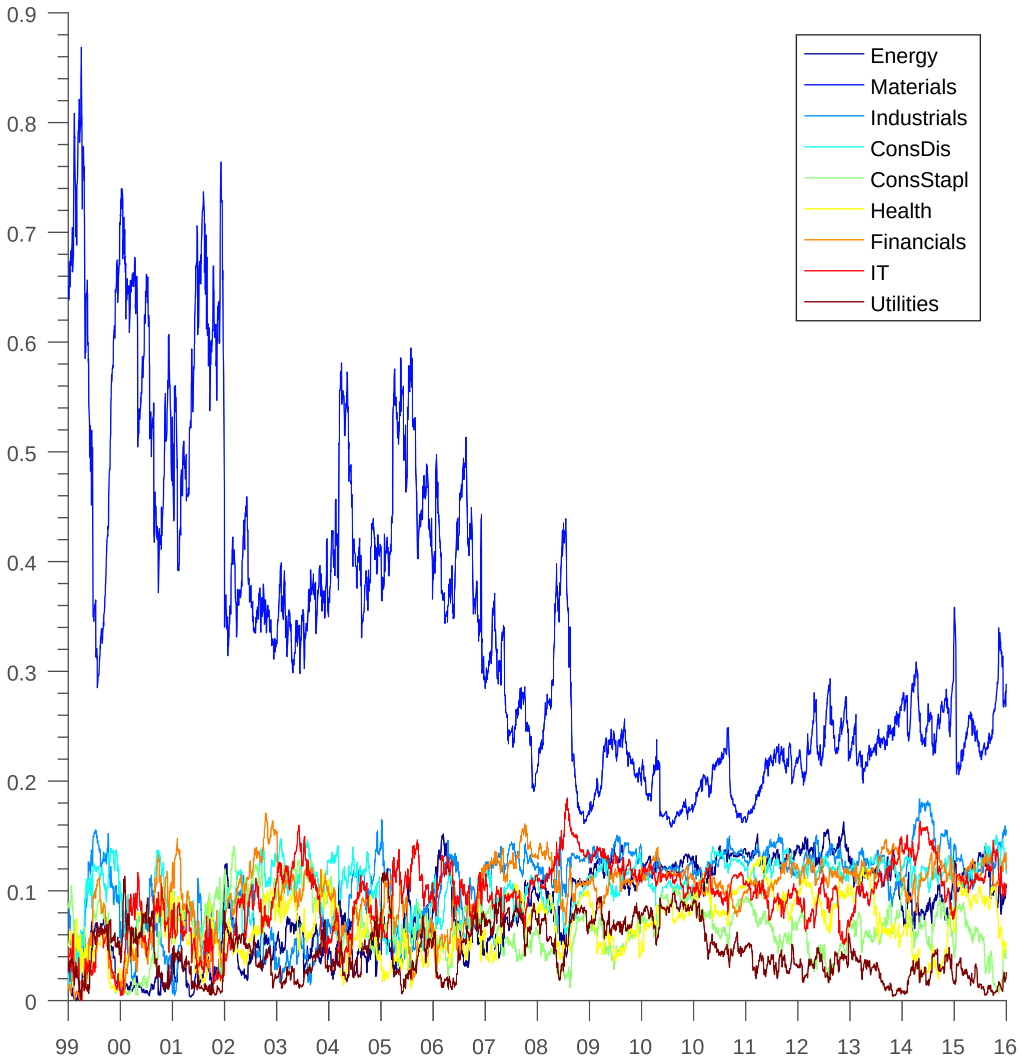

Figure 2.4: Materials, spillovers from other sectors. This picture shows the evolution over time of spillovers received by the Materials sector (which we choose for illustrative purposes only), as measured by the portion of variation in volatility explained by each sector. The idiosyncratic component dominates all the others, although with a decreasing importance especially after the 2008 financial crisis.

but it seems that sectors are so well connected between each other that any event that in principle could affect only one of them actually affects all the others contemporaneously (i.e. during that particular trading day), thus not leading to visible *distinct* change in spillovers, but rather to strong comovements throughout the whole sample. The result is an incredibly high correlation in spillovers. As an example, the burst of the dot-com bubble did not lead to a marked increase in the IT sector centrality, but it rather caused the generalized increase in connectedness that we already highlighted previously.

On the other hand, the housing burst was more systematic than the dot-com one since network components (as opposed to idiosyncratic ones) dominated more in determining all sectors' volatilities. With Lehman failure, our network showed an "alignment effect", whereby both in-degree and out-degree spillovers started to converge (see Figure 2.5 for an example, although all the other sectors have qualitatively very similar pictures, and the same is true for out-degree spillovers) and each sector started to give and receive spillovers to/from the other sectors in very similar magnitudes. This is the dynamic manifestation of the full inter-connectedness result we described in the static analysis, and its existence leads to a natural question: is the network structure causing higher volatility or is it the opposite? Do those effects reinforce each other? Also (as Allen and Gale (2000) first highlights) the finding that a star network is more robust than a ring one does not apply here, and the opposite might well be true. This seeming contradiction is due to the fact that the fundamental existence of a connection is profoundly different in the two cases. Indeed, in Allen and Gale (2000) an edge is placed whenever a contractual relationship exists, whereas in our case the strength of the edge is higher in cases where volatility is spilled over with larger flows. This (in our application) means that a star network is potentially more dangerous than a ring one, since a fully inter-connected structure allows volatility to propagate quickly throughout all the sectors, whereas a ring-like one could potentially slow down volatility transmissions.

Figure 2.6 plots the time-varying dispersion in spillovers received by every sector. Interestingly, the time-varying structure of our network is clear, and we also observe a contemporaneous -62.8% correlation with total connectedness. This means that not only there exists a close connection between the timing of important events and total connectedness, but a similar (opposite) relationship exists with the network structure. In particular, the network tends to have a relatively asymmetric structure during normal times, but it quickly converges towards a star network in periods when key events happen, just as total connectedness tends to increase during those same periods.

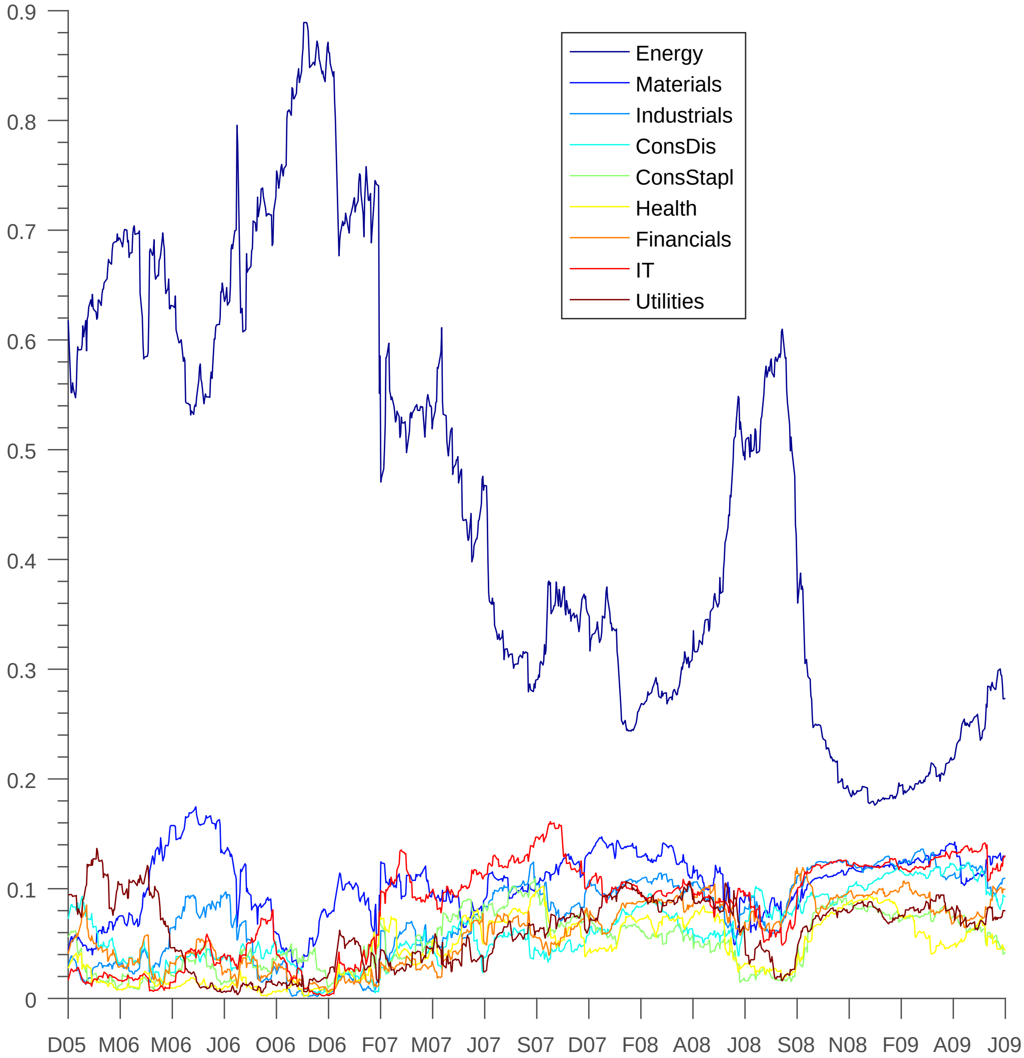

Figure 2.5: Energy, spillovers from other sectors during housing burst. This graph clearly illustrates that immediately after Lehman failure the network profoundly changed and started to have a star shape, where each sector became fully inter-connected with all the others. Notably, the idiosyncratic component of the Energy sector went down from 0.6 to less than 0.2 in very few days, whereas the other sectors' components evenly went up towards $\frac{1}{9} \approx 0.11$, where 9 is the number of sectors. This value, when reached by all the sectoral components, implies that the network is *exactly* fully inter-connected.

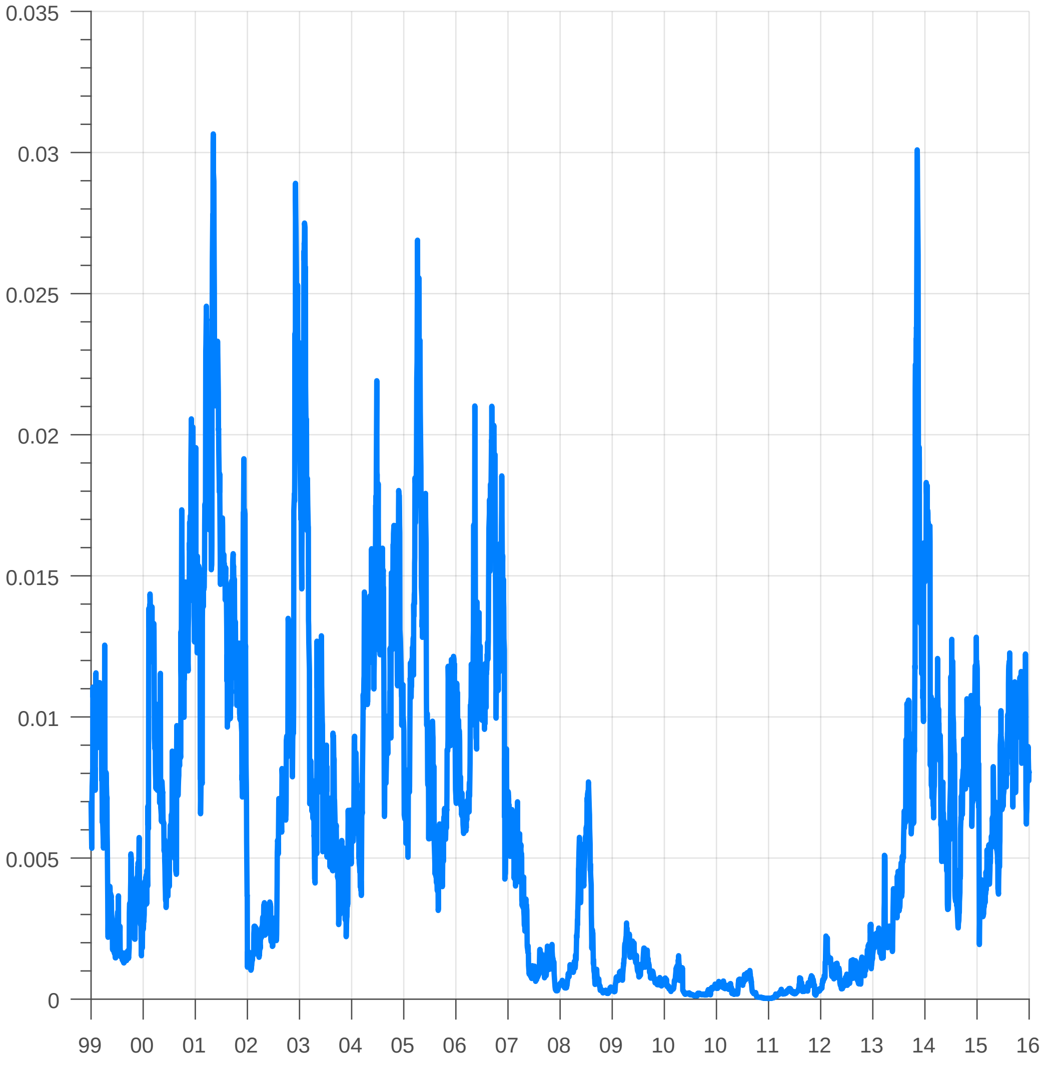

Figure 2.6: Overall network dispersion in spillovers received. Note that during the financial crisis and the sovereign debt crises the dispersion in spillovers reaches its minimum and remains constantly low, thus proving that the pattern we highlighted in Figure 2.5 is not a simple coincidence, but it is present for all sectors.

2.3.4 Filtering Results and Robustness Check

Parameters' forgetting factor estimates tend to have very high values, and many times (15.2% of the times) they hit the upper bound $\lambda = 1$, which corresponds to fixed parameters. Again, this result points to the need to have models that timely detect and correctly estimate time variation in parameters, but that are at the same time able to let the parameters be fixed during periods when they should. Note that Primiceri (2005) always assumes time-variation in parameters, and given the moderately large dimension of our TVP-VAR this would have caused important over-fitting problems. A further corroboration of this issue is that estimates of the shrinkage parameter γ are low during almost all the sample period, thus suggesting that the curse of dimensionality is indeed looming and has to be addressed. Remarkably, rolling windows would not have shrunk parameters either, and using them would thus have caused additional problems. Estimates of the model forgetting factor α are equal to 0.99 for the largest part of the sample, whereas the 0.95 value is selected a few times especially during the earlier observations where connectedness was more volatile, thus suggesting that estimating α together with the other model parameters led to a better fit.

As a robustness check, Figure 2.7 plots estimates of overall connectedness for different forecast horizons and number of lags in the TVP-VAR. The series are nearly identical in both cases, with slight changes only in the level of connectivity, but not in its timing. In the case of different horizons, connectedness is higher at higher horizons, as expected.

2.4 Concluding Remarks

In this paper we study the extent to which sectoral volatilities are inter-connected between them, and how they spill over other sectors. We find that there exist two distinct connectedness regimes whose dividing line lies during the 2007-2008 financial turmoil. The first regime presents a relatively low but volatile connectedness between sectors, whereas after the burst of the housing bubble connectedness stays at a seemingly permanently higher and more stable level. Also, we highlight the existence of a close and negative relationship between the level of total connectedness and the structure of the network as roughly measured by the dispersion in spillovers received. Those key results imply that during high volatility episodes spillovers become evenly distributed between all sectors, and in turn that sectoral volatilities positively influence each others in a stronger manner than what we would have obtained had sectoral connections resembled a ring

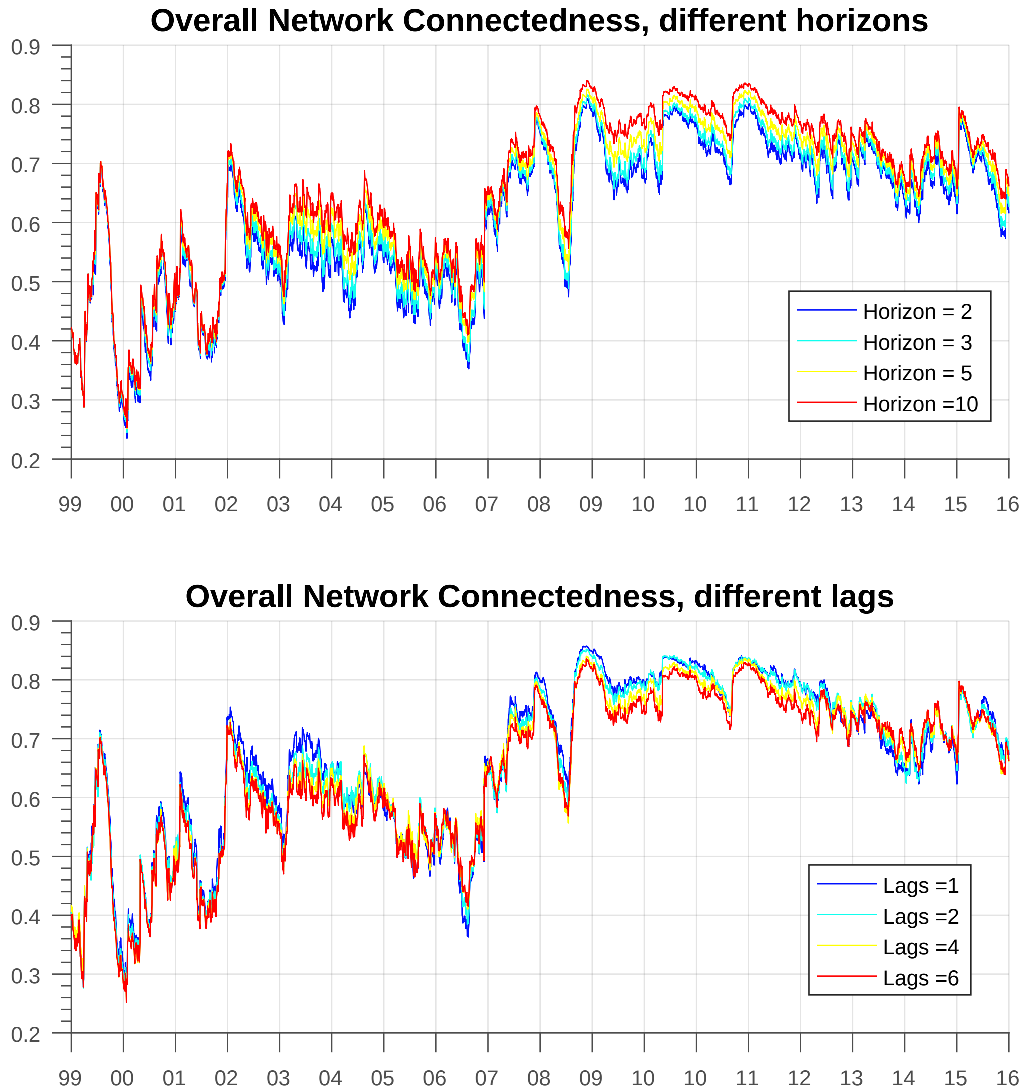

Figure 2.7: Overall connectedness, robustness check. The top panel illustrates estimates obtained by varying the forecast horizon and using 4 lags as in the baseline model. The lower panel fixes the horizon at 10 days and computes connectedness with different lags. In both cases results are virtually unchanged.

network.

While we are agnostic as to the true underlying structural links between volatility, connectedness, and the particular network structure, we believe that situations where the system is fully inter-connected could well facilitate further increase in volatilities. In future work it would be interesting to investigate this issue and try to uncover whether the network structure is mainly causing volatility, whether the opposite is true, or whether both cause each other.

On the technical side, estimating the TVP-VAR with Koop and Korobilis (2013) filtering method yields superior parameters estimates, and estimating the model switching parameter as well proved to be beneficial especially during the first part of the sample where connectedness was more volatile. On the other hand, the model is able to let parameters staying fixed during normal periods. This, together with Bayesian shrinkage, addresses the over-fitting issue in a satisfactory manner and remarkably produces more reliable connectedness estimates.

Anzuini, A., Rossi, L., and Tommasino, P. (2017), Fiscal Policy Uncertainty and the Business Cycle: Time Series Evidence from Italy. *Banca d'Italia, Temi di Discussione (Working Papers)*, 1151.

Chapter 3

Fiscal Policy Uncertainty and the Business Cycle: Time Series Evidence from Italy

Joint with Alessio Anzuini and Pietro Tommasino[1]

3.1 Introduction

After several years of recession or sub-par growth in many countries, several economists and policy makers have become convinced that widespread uncertainty might have concurred to the unsatisfactory pace of the recovery[2].

More generally, economic theory suggests that, under certain conditions, uncertainty shocks may be important in explaining economic fluctuations: firms may react to an increasingly uncertain environment by reducing hiring and investment, financial intermediaries may become more reluctant to lend, and households may increase their propensity to save, as supported by the evidence in the empirical literature. In this respect, Bloom (2014) is a comprehensive review of the stylized facts about uncertainty and its relationships on key economic variables, both from a theoretical and an applied point of view.

Economic uncertainty can take many forms, and may originate from several sources. To date, there are various indices of Economic Policy Uncertainty, but few indices of Fiscal Policy Uncertainty (FPU) which also mainly have focused on FPU stemming from US policies. In the present paper we propose a new measure of FPU, and study its effects on the macroeconomic situation in Italy. Indeed, for most of its recent history Italy has been characterized by fragile public finances and by a highly partisan and often fragmented political landscape. It is therefore

[1]The opinions expressed herein are of the authors and do not necessarily reflect those of Banca d'Italia.

[2]See e.g. IMF (2012).

an extremely appropriate laboratory to study FPU and its consequences.

Our proposed FPU index is constructed as follows: we first estimate a fiscal reaction function in order to capture how the fiscal stance reacts to economic developments. The key difference with respect to previous empirical exercises[3] is that the fiscal rule incorporates an innovation not only to the level, but also to the volatility of the fiscal stance (technically, we adopt a stochastic volatility model). As a second step, we feed a VAR model with the two series of innovations - i.e. innovations to the *level* and to the *volatility* of the fiscal variables of interest - and analyze how they impact the macro-economy.

Our FPU index peaks in correspondence with key historical events that greatly influenced Italian public budgets. Among the most significant spikes in FPU we find those recorded during the 1992 Exchange Rate Mechanism crisis and the introduction of the Euro in 1999. Moreover, our index has a positive and significant correlation with Baker, Bloom and Davis (2016) EPU index for Italy, while the differences have to be re-conducted to the fact that the two indexes are measuring different types of uncertainty. Finally, fiscal policy seems to react with fiscal consolidations both when debt figures increase and when the business cycle improves.

Our structural analysis finds that the effects of an expansionary fiscal shock are quite standard and in line with the previous VAR literature - i.e. we find positive multipliers. More interestingly, we also find that an increase in FPU has a negative impact on the economy. Additionally, when a positive fiscal shock is coupled with an unexpected increase in uncertainty the potentially positive effects become blurred, thus suggesting that expansionary policies can lose their effectiveness when implemented with high degrees of uncertainty. As a crucial policy implication, our results suggest that policy makers should be aware that closely and credibly targeting a pre-announced balance budget level could significantly increase the effectiveness of discretionary expansionary fiscal policies, and reduce the recessionary effects of contractionary ones. Therefore, attention should be paid not only to the first moment, but to the entire distribution of (conditional) future budget outcomes.

Our paper contributes to two different streams of the macroeconomic literature. First, the recent empirical research on the macroeconomics of uncertainty. As we already mentioned, uncertainty stems from several sources. Some papers have focused on stock-market-induced uncertainty, such as Bloom (2009) which uses peaks in stock market volatility (captured by a dummy variable) as a measure

[3]For a review, see Golinelli and Momigliano (2009).

of uncertainty[4]. Policy may be clearly another relevant source of macroeconomic uncertainty[5], and Baker, Bloom and Davis (2016) proposes a broad policy uncertainty index that is an average of three components: the frequency of references to economic policy uncertainty in the news, the amount of federal tax code provisions set to expire in future years and the extent of professional forecasters' disagreement over one-year-ahead inflation and government purchases of goods and services. Jurado, Ludvigson and Ng (2015) estimates macroeconomic uncertainty with a stochastic volatility model, but instead of using a particle filter (as we do) the authors adopt MCMC algorithms. More specific indicators are those related to trade policy and monetary policy, developed respectively by Handley (2014) and Creal and Wu (2016)[6].

The only paper that we are aware of that looks at fiscal uncertainty shocks is Fernández-Villaverde, Guerrón-Quintana, Kuester and Rubio-Ramírez (2015). We follow their econometric methodology, and proxy FPU with the time-varying volatility of the innovation of a fiscal reaction function[7]. However, differently from Fernández-Villaverde et al. (2015), we look at the overall (cyclically adjusted) primary deficit (CAPB) and not just to some of its components. This more encompassing variable is the most used indicator of the government's fiscal stance (incidentally, the CAPB also plays a relevant role in the context of the fiscal framework of the European Union).

Given our focus on the CAPB, our paper is directly relevant for a second stream of literature, namely that concerned with the macroeconomic effects of discretionary fiscal policy[8]. A review of that field is clearly outside the scope of this introduction, but it is well known that there is no consensus about the size - and even the sign - of fiscal multipliers. On one side, studies like Blanchard and Perotti (2002) and Romer and Romer (2010) find standard demand-driven

[4]See also the early paper by Romer (1990).

[5]Policy uncertainty (i.e. not knowing which policy will be implemented) may be in turn due to political uncertainty (i.e. not knowing who will be in power). The economic effects of this latter variable have been studied, for example, by Julio and Yook (2012) and Canes-Wrone and Park (2014).

[6]A related stream of literature neglects the real effects of policy uncertainty, focusing instead on its financial consequences. See e.g. Kelly et al. (2014) and Brogaard and Detzel (2015). Other papers, e.g. Gulen and Ion (2015), look at the microeconomic (firm-level) effects of changes in policy uncertainty. Incidentally, both Brogaard and Detzel (2015) and Gulen and Ion (2015) use the Baker et al. (2016) index.

[7]A similar methodology is adopted by Scotti (2013) and Jurado, Ludvigson and Ng (2015). Both papers aim at modeling macroeconomic volatility at large, not FPU.

[8]On the contrary, it is not easy to compare the results of Fernández-Villaverde et al. (2015), which looks at specific budgetary items, with those of papers which focus on more aggregated fiscal variables. The authors themselves acknowledge this limitation in the Appendix A of their paper.

Keynesian effects; on the other side, starting from Giavazzi and Pagano (1990), other authors have argued that the effects of a fiscal change can be non-Keynesian, with the possibility of expansionary fiscal consolidations and contractionary fiscal expansions, as is the case in Alesina and Ardagna (2013). Our main contribution to this debate is to show that fiscal policy-makers can influence the economy not only by changing the level of the budget deficit, but also by affecting its volatility. As a consequence, the same change in the government budget (say a budgetary expansion) can have different effects depending on whether it is associated with a reduction or an increase in the FPU. From an econometric viewpoint, this implies that a proper assessment of the impact of changes in the fiscal policy stance should correctly identify both the level and the uncertainty shock. From a policy perspective, our findings highlight the importance for policy makers of being credible, and avoid policies that are unsustainable in the long run.

The remainder of the paper is organized as follows. In Section 2 we outline the methodology we use to measure FPU, we describe our data, and present the results. In Section 3 we present our VAR estimates and we show the effects of the fiscal shocks on macroeconomic variables. Section 4 concludes.

3.2 Estimating Fiscal Policy Uncertainty

3.2.1 The Empirical Model: A Fiscal Rule with Time Varying Volatility

We estimate the following two-equation state space model:

$$\mathrm{def}_t = \beta_1 \mathrm{def}_{t-1} + \beta_2 \mathrm{debt}_{t-1} + \beta_3 \mathrm{gap}_{t-1} + e^{h_t} u_t, \qquad u_t \overset{iid}{\sim} \mathcal{N}(0,1), \tag{3.1}$$

$$h_t = \alpha_0 + \rho h_{t-1} + \gamma \varepsilon_t, \qquad \varepsilon_t \overset{iid}{\sim} \mathcal{N}(0,1), \tag{3.2}$$

where def_t is the cyclically adjusted ratio between the general government primary borrowing requirement and GDP at time t, debt_{t-1} is the debt ratio, gap_{t-1} is the output gap, and h_t is the log-volatility of the error term. Concerning the parameters, the βs have obvious interpretations, ρ is the persistence of the log-volatility and γ is the volatility of the shocks to log-volatility.

Equation (3.1) is a very standard fiscal reaction function, in the spirit of the ones analyzed in (e.g.) Galí and Perotti (2003) and in Golinelli and Momigliano (2009). Equation (3.2), instead, gives the law of motion for the volatility of the deficit, which in our model is not conditionally deterministic (as, for example, in a GARCH model) but includes a stochastic component[9]. Equation (3.2) captures our main idea: fiscal policy can in principle be affected by *two* kinds of innovations: level shocks (u_t) and FPU shocks (ε_t). Our FPU index is e^{h_t}, which (as opposed to h_t) intuitively represents the time-varying standard deviation of fiscal shocks.

The inclusion of a stochastic volatility element is important from an economic viewpoint to fully capture the nature of fiscal policy-making, but comes with some non-negligible computational costs. Indeed, it makes our model non-linear, precluding the use of standard econometric techniques, such as the Kalman filter, which requires instead linearity and Gaussianity. To estimate equations (3.1) and (3.2) we resort to particle filter estimation. This technique is similar in spirit to the Kalman filter: in both methods, non-data information (prior) and data information (likelihood) are combined to obtain an estimate of the variables of interest. Furthermore, as in the Kalman filter, the process h_t is not observable and has to be estimated together with all the relevant parameters.

However, differently from the Kalman filter, we do not have closed-form solutions for the posterior distributions, and the integrals involved in the computations of the posterior are approximated by using the discrete random samples obtained by drawing from the posterior[10].

We use the Liu and West (2001) version of the of the particle filter, which allows joint estimation of state and parameter vectors. To ensure convergence of the procedure, we follow Liu and West (2001) suggestion and introduce the following re-parameterization of our model:

$$\begin{aligned} \alpha_0 \equiv & (1-\rho)\omega, \\ \rho \equiv & \frac{\exp(\bar{\rho})}{\exp(\bar{\rho})+1}, \\ \gamma \equiv & (1-\rho^2)^{\frac{1}{2}} e^{\bar{\gamma}}, \end{aligned} \tag{3.3}$$

and we estimate ω, $\bar{\rho}$, and $\bar{\gamma}$ instead of α_0, ρ, and γ. Incidentally, the re-parameterization allows a relatively easy interpretation of the parameters we need to estimate[11]. Indeed, $\mathbb{E}(h_t) = \omega$, which means that ω is the log-modal volatility, and $\mathbb{V}(h_t) = e^{2\bar{\gamma}}$, i.e. $\mathrm{sd}(h_t) = e^{\bar{\gamma}}$.

---

[9]The advantages of a stochastic volatility model with respect to a GARCH are highlighted in Fernández-Villaverde and Rubio-Ramírez (2013).

[10]The algorithm for the basic version of the particle filter has been developed by Gordon, Salmond and Smith (1993). Other important contributions are included in Doucet, De Freitas and Gordon (2001).

[11]Without these transformation, suggested in Liu and West (2001), we could have had problems with estimating variances (which must be positive) and auto-regressive parameters (which must be inside the unit circle). Instead, ω, $\bar{\rho}$, and $\bar{\gamma}$ can assume any real value, as the logistic transform is constrained in the $[0, 1]$ interval, and the log transform ensures a positive parameter.

3.2.2 Data

Monthly data for general government borrowing requirements and public debt are taken from the official series published by the Bank of Italy (*Supplement to the Statistical Bullettin - The public finances, borrowing requirements and debt*). Even though the fiscal rule is estimated with data starting from 1981, the VAR we later estimate uses other variables that are only available from a later period. In particular, we estimate the VAR on the time span that goes from January 1991 to March 2014.

The borrowing requirement is computed on a cash basis, using changes in the stock of debt instruments, on which precise and almost complete information is available. As discussed in Levin (1993), it is controversial whether cash-basis or accrual-basis data (as in the national accounts) are the most appropriate when studying the impact of government operations on the economy. In our case, cash data are more reliable and sufficiently long. Moreover, deficit and debt data are built with the same methodology and criteria.

As it is customary in the literature[12], we exclude from the borrowing requirement debt settlements and privatization receipts, because the first refers to expenditures undertaken in past periods, while the latter cannot be considered as resources compulsorily subtracted from the private sector.

Since the GDP series has quarterly frequency, the fiscal reaction function is estimated on a quarterly basis[13]. Note that we could have tried to retrieve a monthly measure for GDP within a mixed frequency approach, therefore being able to estimate a monthly fiscal rule. However, in order to contain estimation errors we preferred to avoid estimating a further state variable (i.e. monthly GDP), hence we opted for aggregating the figures for borrowing requirements and working at a lower frequency. Indeed, note that although debt figures are available at a monthly frequency, they do not require temporal aggregation because they are stock variables. Nonetheless, we believe that a formal comparison of fiscal rules estimated at different frequencies (e.g. monthly versus quarterly) would alone deserve to be performed in future research.

[12] See e.g. Giordano, Momigliano, Neri and Perotti (2007).

[13] As public finances data are monthly, deficit figures are aggregated by sum, while for debt, we use the start-of-quarter figure. The output gap is obtained by HP-filtering the series for the log real GDP ($\lambda = 1600$). All the series are seasonally and calendar adjusted using X-ARIMA-12 RSA4c filtering.

3.2.3 Choosing the Priors

Economic theory does not offer any hint about the values of the parameters for debt and GDP in the fiscal reaction function, i.e. β_2 and β_3 in Equation (3.1), so we choose zero-mean uniform priors on a very wide support in both cases. Regarding the auto-regressive parameter β_1, one can reasonably expect it to lie between 0 and 1, since the CAPB series appears stationary and does not present negative auto-correlation. Therefore, we use a uniform prior on this support.

We do not have prior information about $\bar{\rho}$ (the logit of the persistence of the log-volatility), therefore we set $\bar{\rho}_{0|0} \sim \mathcal{N}(0, 1.5)$, which implies that ρ has an almost uniform density on the support $[0, 1]$, i.e. we expect log-volatility to be neither a negatively auto-correlated nor an explosive process. The parameter ω is the modal log-volatility of the fiscal shock. We expect e^{ω} to be lower than the unconditional (sample) volatility of def_t (Figure 3.1, lower left panel), and we consequently choose $\omega_{0|0} \sim \mathcal{N}(-3.94, 0.2)$. Choosing a prior for $\bar{\gamma}$ is particularly difficult so we choose it based on what are the likely effects on the standard deviation of the level shock once a one-standard-deviation volatility shock occurs (see Figure 3.2)[14]. Finally, recall that $\mathbb{E}(h_t) = \omega$, and $sd(h_t) = e^{\bar{\gamma}}$. Therefore, our prior for the log-volatility is $h_{0|0} \sim \mathcal{N}(\mathbb{E}(\omega_{0|0}), e^{\mathbb{E}(\bar{\gamma}_{0|0})})$.

3.2.4 Estimates of Fiscal Policy Uncertainty

Figure 3.3 plots the estimated series for the time-varying volatility $e^{\hat{h}_t}$ recovered with the particle filter[15]. Two of the three relative peaks of the index during the eighties (the one at the beginning of 1983 and the one in 1985) correspond to two well-known episodes of macroeconomic turbulence related to public finances. At the end of 1982 the Bank of Italy refused to buy government securities unsold on the primary market, creating uncertainty on sovereign bond markets and obliging the Parliament to pass a one-off one-year advance form the Central Bank. In 1985, the repayment of a dollar-denominated loan of a large public enterprise was associated with severe foreign exchange turbulence.

It is interesting to note that the two main peaks in the volatility series are in the nineties and coincide with critical moments in the recent history of Italian

[14] In the figure we show the effects on the standard deviation of the level shock when two distinct one-standard-deviation volatility shocks occur. The first is a negative shock (left density), whereas the second is a positive one (right density). The red line in the middle is the median of the prior for the standard deviation of the level shock.

[15] We run the particle filter using $M = 100,000$ and $R = 150,000$, where M is the number of particles that jointly approximate the posterior, and R is the number of draws in the auxiliary variable sampling step.

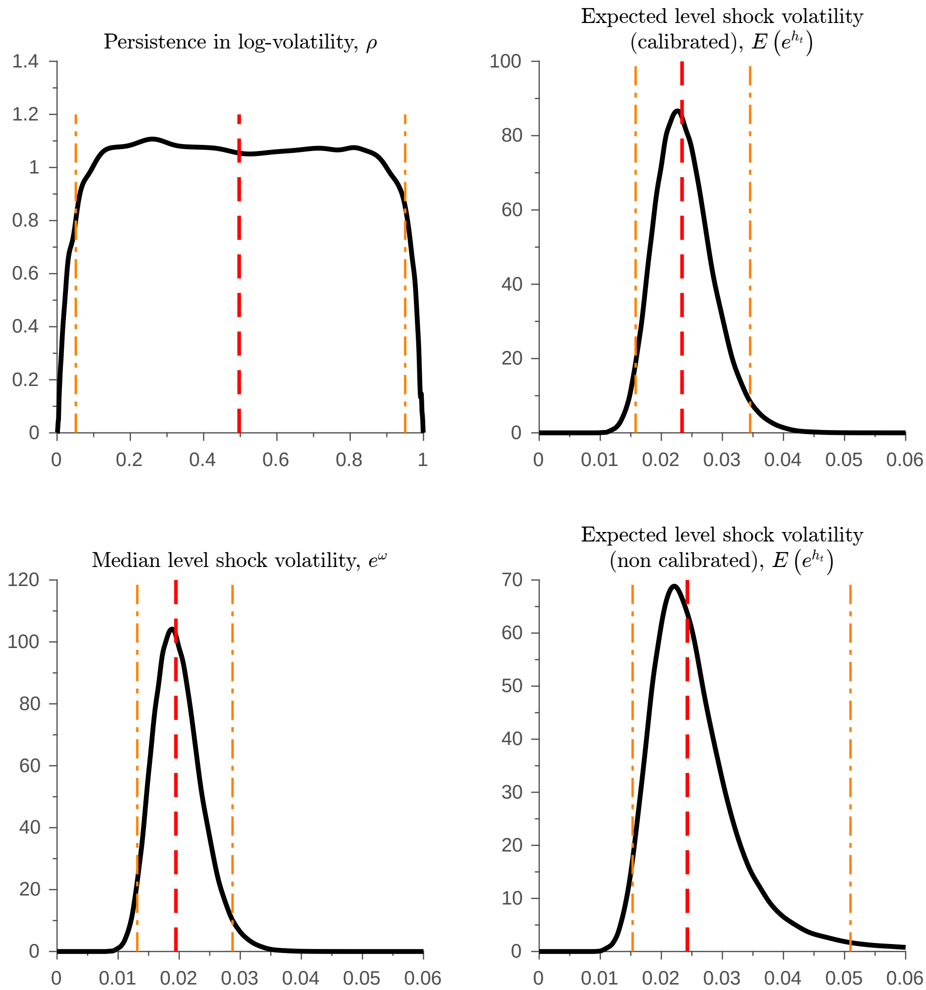

Figure 3.1: Priors for transformed parameters. Black lines are normal kernel density estimates, red lines are the median of the distribution, whereas orange bands are 95% prior probability intervals. The top left panel is the prior distribution for the auto-regressive stochastic log-volatility parameter, obtained as the logistic transform of the prior for the parameter $\bar{\rho}$. The bottom left panel is the prior distribution for the median volatility, where median log-volatility has been centered at the log-standard-deviation of residuals obtained from estimating the fiscal rule with OLS. The remaining two panels plot the prior distribution for the expected level shock volatility, where the upper one calibrates the parameters needed to compute it at their expected values, whereas the lower one accounts for their prior uncertainty. All resulting priors are very uninformative.

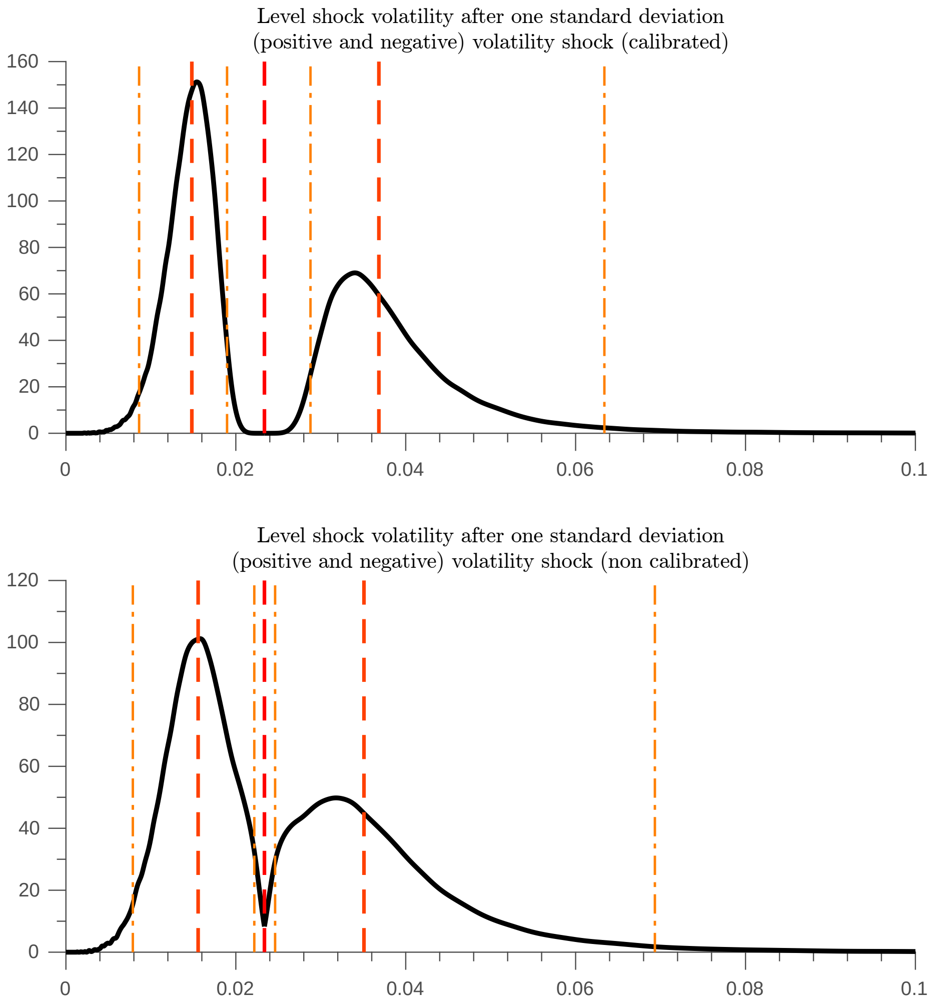

Figure 3.2: The figure shows the prior evaluation of $\bar{\gamma}$, which represents the log-volatility of log-volatility. The red line is the prior expected volatility, dark orange lines are prior median volatilities after one-standard-deviation positive and negative volatility shock, whereas light orange bands are their respective 95% prior probability intervals. The top panel calibrates the relevant parameters at their prior expected value, whereas the bottom one accounts for their prior uncertainty. As one can see, effects are large enough to be sure that the prior distribution we choose for $\bar{\gamma}$ is flat enough to be uninformative.

public finances. The first is in the second half of 1992, when a balance of payments crisis questioned the viability of the fixed exchange regime and the sustainability of Italian public finances. In this circumstance, the Government tightened budgetary policy (with emergency measures decided in July 1992), but ultimately (September 1992) the country was forced to abandon the European Exchange Rate Mechanism. The second peak is in the first half of 1999, i.e. the first months of the European Monetary Union. Starting from January 1^{st} 1999, the Euro area countries were subject to a single fiscal framework, the so called Stability and Growth Pact. Doubts about the implementation of the new rules can easily rationalize this spike in our measure of uncertainty. A further element of fiscal uncertainty was determined by the promise by the Government to give partly back to taxpayers – but only in case of EMU admission – a one-off tax which was levied in 1996. The fraction of the restitution was not specified ex-ante (it was decided only at the end of 1999 and turned out to be 60%). Also relevant might have been, in the same period, the introduction of two brand new taxes (the municipal and regional additions to the personal income tax) also meant to increase the degree of fiscal autonomy of local governments (which might in itself be considered something which increases fiscal uncertainty).

Finally, the local peak in 2001 can be rationalized as the effect of a significant turning point in fiscal policy, as the Parliament approved the first expansionary budget in years. With the benefit of the hindsight, fiscal out-turns also benefited by the windfall gains due to buoyant financial markets (which also influence our fiscal stance measure).

Our level shock series, although recovered in a completely different framework, correlates significantly with those recovered by Giordano et al. (2007). In particular, and as expected, it correlates positively with their tax shock series and negatively with their expenditure shocks series.

Figure 3.4 depicts (as a representative example) the output gap coefficient estimates obtained with our baseline particle filter against standard OLS estimates. First, recall that parameters are assumed to be fixed in our model, whereas the time-varying nature of the estimates obtained with the particle filter is due to the *posterior distribution* being time-varying, *not* the parameter itself. This behavior is natural in a filtering context, where estimates become more precise as observations arrive. Had we estimated parameters with OLS on an expanding window we would have obtained a very similar pattern. Interestingly, we see that at the beginning of the sample the posterior distribution has a very wide support, and that the point estimate is very unstable. This is sensible since very few observations are available in the first part of the sample. Nevertheless, as

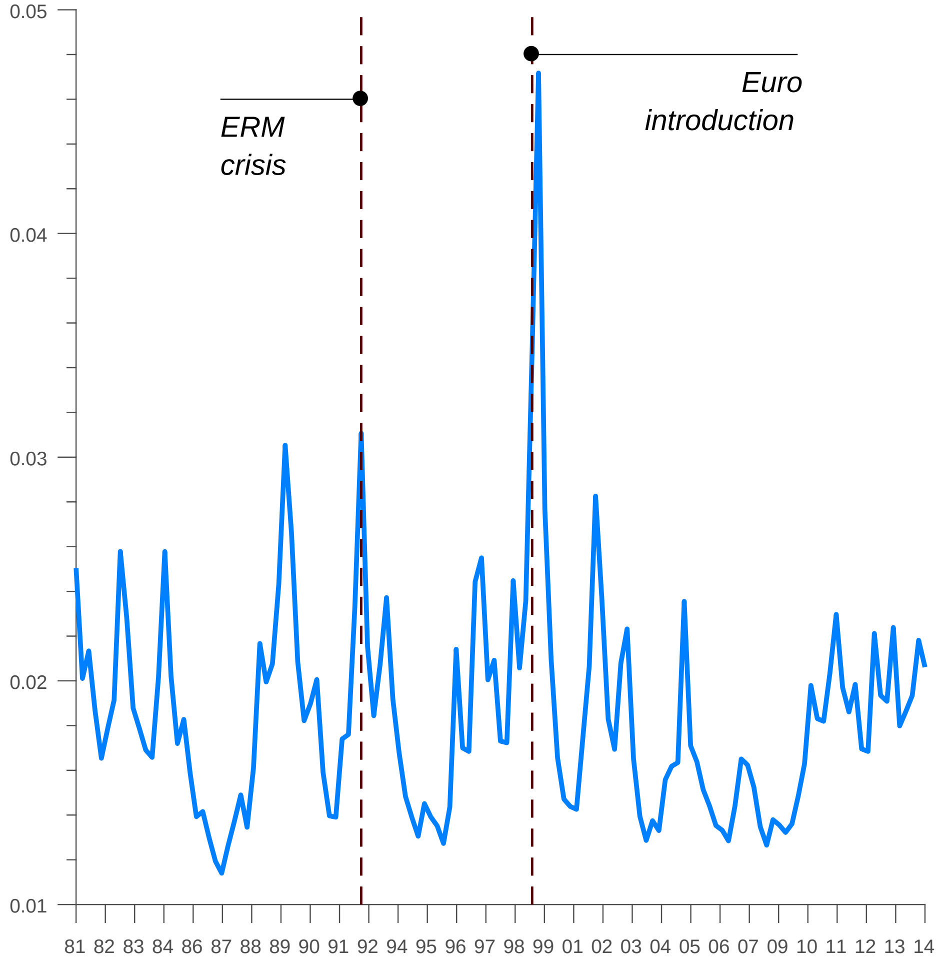

Figure 3.3: Fiscal Policy Uncertainty Index. This series is the average of the posterior distribution of the residual volatility recovered with the particle filter. Note that we could also have used the median as our centrality measure. However, the differences between the two series are negligible, thus we select the mean as our preferred central measure of fiscal volatility.

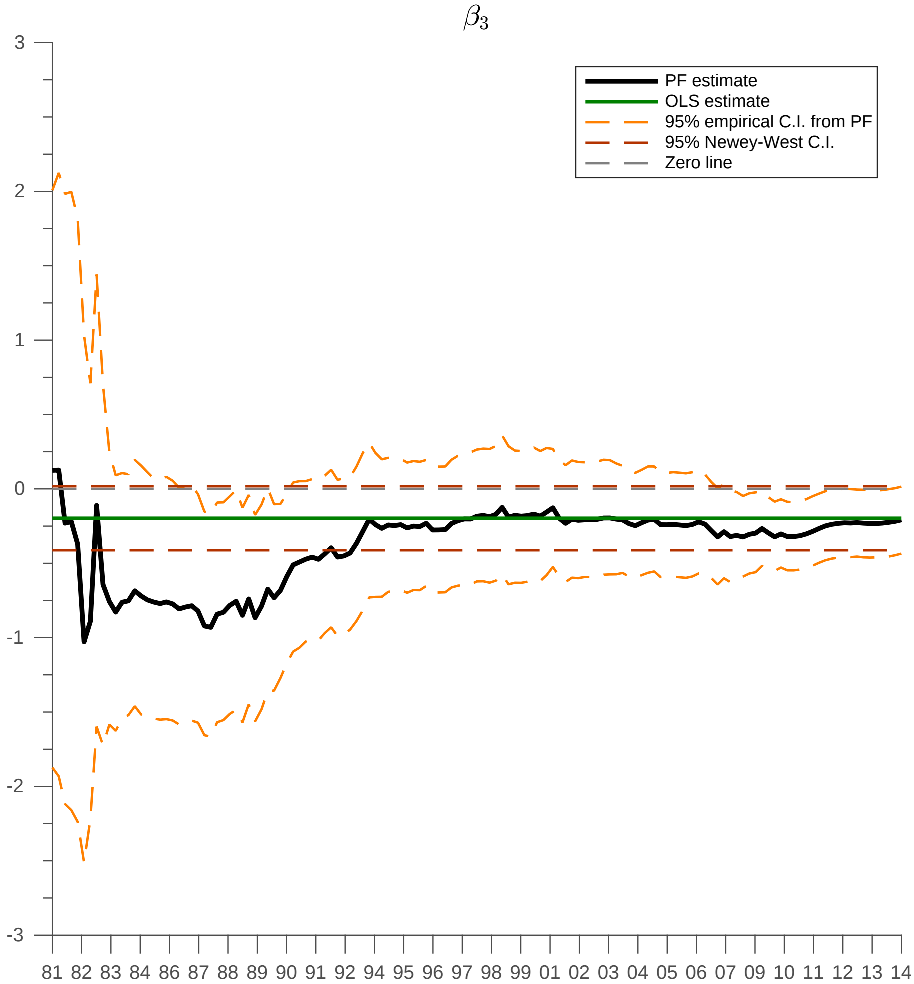

Figure 3.4: Particle Filter against OLS estimates. The graph shows that the filtering procedure produces very similar coefficients as those obtained with standard OLS. Using the particle filter allows to further disentangle the error component non-linearly with the underlying stochastic volatility model which we assume to govern residual volatility.

observations are added the support becomes narrower, until with the very last observation both the point estimate and the 95% probability intervals converge to their OLS counterparts. This is practical evidence of the fact that the prior we choose is flat enough to be uninformative. Results for the other coefficients are very similar. Moreover, residuals recovered with the particle filter and with OLS are the same to all practical purposes.

Figure 3.5 plots posterior estimates of fiscal rule parameters. Recall that we set uninformative priors, meaning that the shape of the posterior distribution is only driven by data information. Apart from the auto-regressive parameter β_1 (which is positive and significant, as expected) we find that the reaction to an increase

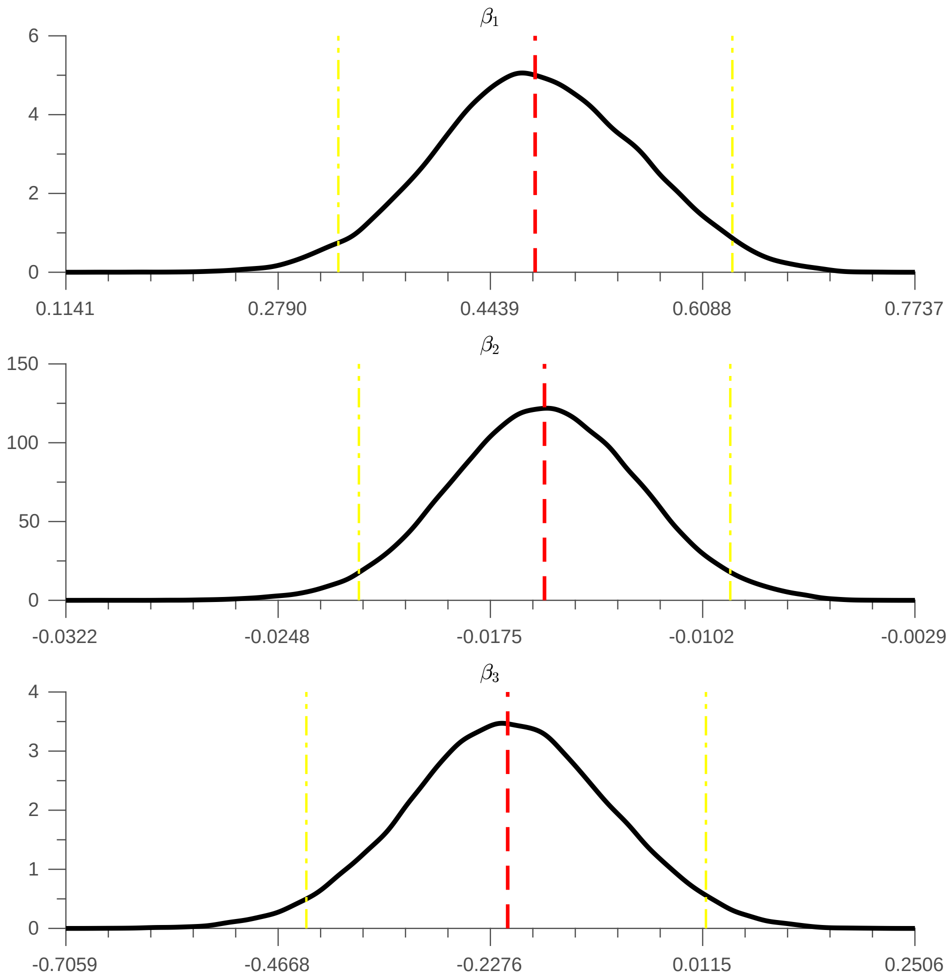

Figure 3.5: Fiscal rule coefficients: posterior estimates. Black curves are normal kernel density estimates, red lines are median estimates, whereas yellow bands are 95% posterior probability intervals. The empirical sample is based on $M = 100,000$ draws from the posterior distribution.

in the debt-GDP ratio is a significant fiscal consolidation (negative β_2). On the other hand, fiscal policy seems to be counter-cyclical (negative β_3) although the coefficient is marginally non-significant on a 95% probability interval.

3.3 Fiscal Policy Uncertainty and the Macroeconomy: a VAR Approach

3.3.1 Baseline Results

Having recovered the two series of the fiscal level shock and fiscal volatility shock we are now ready to analyze their impact on macroeconomic variables. In particular, we estimate a recursive auto-regressive model with conditioning exogenous variable corresponding to our measure of fiscal level shock and fiscal volatility shock (FPU). In the econometric literature this model is usually referred as a VARX model or as a rational distributed lag model[16]. Our system of equations is:

$$\mathbf{Y}_t = \boldsymbol{\delta}_0 + \boldsymbol{\delta}_1 \mathbf{t} + \boldsymbol{\delta}_2 \mathbf{t}^2 + \mathbf{A}(L)\mathbf{Y}_{t-1} + \mathbf{B}(L)\boldsymbol{\chi}_t + \mathbf{C}(L)\boldsymbol{\mu}_t + \boldsymbol{\upsilon}_t, \tag{3.4}$$

where the vector $\mathbf{Y}_t$ contains the log of real private GDP, the log of the private GDP deflator, log private employment and the 10 years Government bond yields. The variables $\boldsymbol{\chi}_t$ and $\boldsymbol{\mu}_t$ are respectively the fiscal level shock and the FPU determined outside the system of the equations. $\boldsymbol{\delta}_0$, $\boldsymbol{\delta}_1$ and $\boldsymbol{\delta}_2$ are vectors of coefficients, while $\mathbf{A}(L)$, $\mathbf{B}(L)$ and $\mathbf{C}(L)$ are finite-order polynomials in the lag operator L.[17] Finally, $\mathbf{t}$ is a time trend, and $\boldsymbol{\upsilon}_t$ a vector of white noise and mean-zero indipendently and identically distributed error terms.

Note that we could have included the fiscal shocks in the vector of endogenous variables, ordered first in the spirit of Romer and Romer (2010). As we will see later, doing this yields virtually the same results.

Our system is estimated using standard Bayesian techniques. In particular, we use a non-informative prior (Jeffrey's prior) distribution on parameter space and an inverse Wishart distribution as the conjugate prior for the covariance matrix. Antithetic acceleration is then used to improve convergence of the Monte Carlo draws.

We feed the estimated model with a one-standard-deviation shock on the unexpected variations in the cyclically adjusted primary balance (as a fraction of GDP) or, alternatively, a one-standard-deviation shock in unexpected FPU (i.e. the shocks to the log-volatility of the innovations to the balance budget). We show that not only the effects of the two shocks are quite different but also that not properly disentangling the two sources of the fiscal shocks and mixing them in a single shock would blur the effects of fiscal policy.

[16] See Lütkepohl (2005).

[17] Both AIC and BIC select 1 lag as preferred specification.

Figure 3.6 shows the conditional movements of the Italian macroeconomic variables after an unexpected expansionary fiscal shock. Standard Keynesian effects tend to dominate. GDP rises significantly and so does employment: the peak[18] response in GDP is reached after one quarter and, although positive, is rather weak (0.23%); the response in employment is strong and persistent reaching a peak (0.2%) after two quarters and remaining significantly positive for three years. The 10-year government bond yield tend to rise although not significantly.

When we feed the model with an unexpected shock to FPU, results are reversed. Private GDP persistently and significantly decreases reaching a negative trough after six quarters and employment persistently decreases reaching its lowest level after 10 quarters (results are shown in Figure 3.7). This result is consistent with the large theoretical literature on real options, whose key finding is that increases in the volatility of the profitability of investment opportunities induce firms to adopt a wait-and-see strategy, which is a rational behavior even when firms are *not* risk averse. The reason for this is that higher uncertainty increases the probability of receiving very low returns from investment, and the firm can *profitably* hedge from this undesirable outcome by waiting for uncertainty to resolve, and particularly by waiting until investment returns reach a given higher threshold. In this sense, an investment opportunity which has a considerable sunk cost that can be delayed over time can be seen as a call option, where the firm can decide whether and when to invest. In our framework, fiscal policy uncertainty is one of the many sources of uncertainty that a company faces, and (as we said above) if higher uncertainty leads firms to wait for new (capital and labor) investments, this leads *ceteris paribus* to a generalized depression, which is what we find to happen in our sample.

To sum up, the two fiscal shocks have an opposite impact on economic activity: GDP increases after a level shock (fiscal expansion) and decreases after a volatility shock (FPU increase). As we argued in the Introduction, ignoring the existence of the two dimensions of fiscal shock might be the underlying reason for the different effects recovered in the literature on fiscal multipliers.

Our results suggest that both the Keynesian and the non-Keynesian view of the effects of fiscal policy may be reconciled: a policy which increases the budget deficit tends to sustain growth if the way in which it is implemented decreases - or at least does not increase - FPU, but it can be contractionary otherwise. To illustrate this point, in Figure 3.8 we show the joint effect on the dynamic

[18]Importantly, recall that the figures to come are to be intended as *quarterly* (not annual) responses, which means that the corresponding annual effect would roughly be equal to four times the reported numbers.

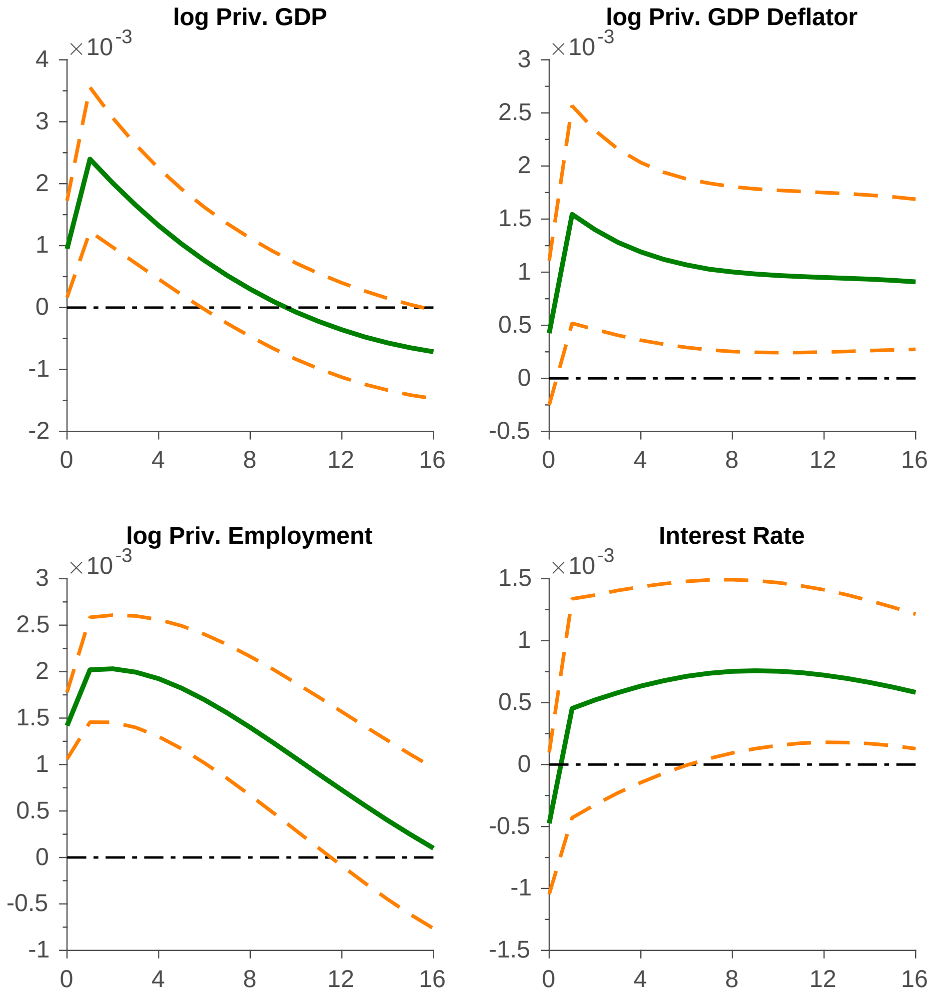

Figure 3.6: Impulse Response Functions - CAPB level shock. Orange bands are 68% bootstrapped confidence intervals, and responses are quarterly effects. A fiscal expansion drives private production and employment upwards, and puts pressure on prices and interest rates.

system of a one-standard-deviation expansionary fiscal shock and a simultaneous two-standard-deviations FPU shock[19]. In this case, the response of private GDP and private employment becomes largely insignificant. The fiscal expansion is in this case worthless as it induces a recessionary increase in FPU. The example corroborates our argument that governments should take into account, when assessing the effectiveness of a planned fiscal measure, the possible effects on uncertainty.

[19]In selecting a higher volatility shock we follow Fernández-Villaverde et al. (2015).

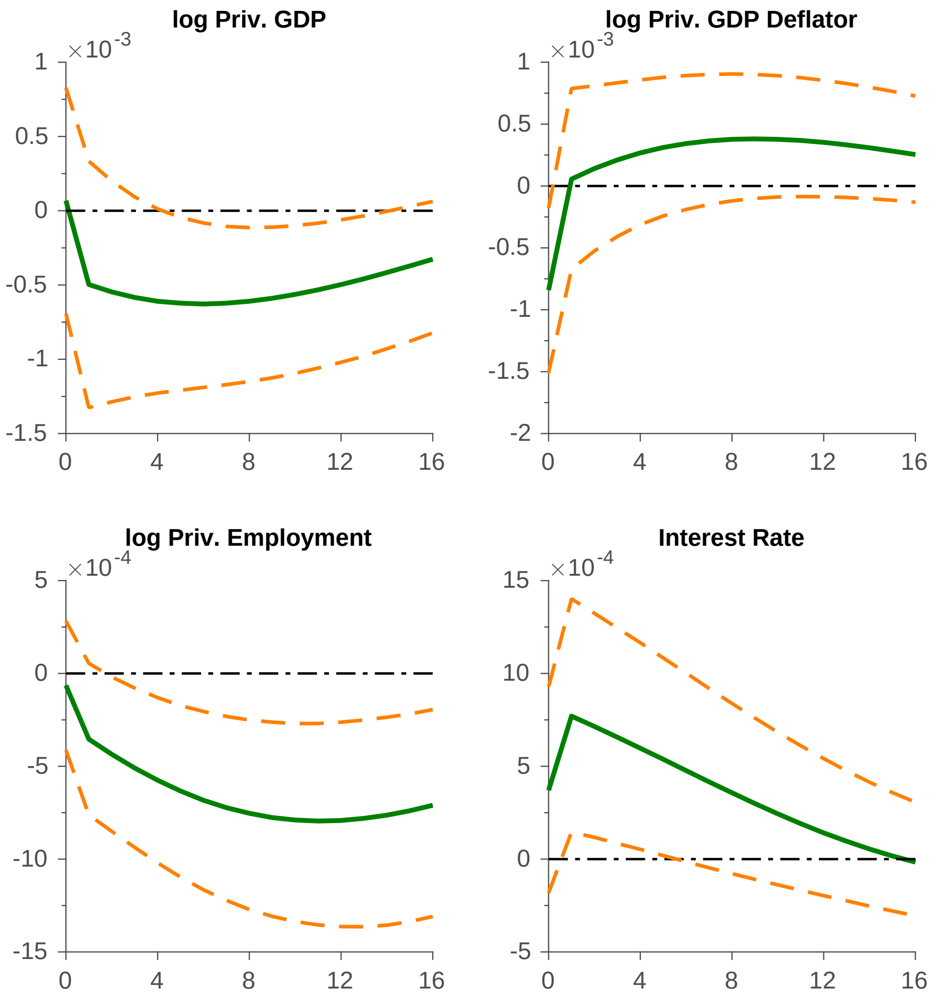

Figure 3.7: Impulse Response Functions - FPU shock. A one-standard-deviation shock in fiscal volatility leads economic activity to drop, and this is most likely due to the real options channel which leads private agents to a wait-and-see behavior which potentially delays investment and consumption decisions. Note that the interest rate response is still positive and significant, whereas prices increase although not significantly.

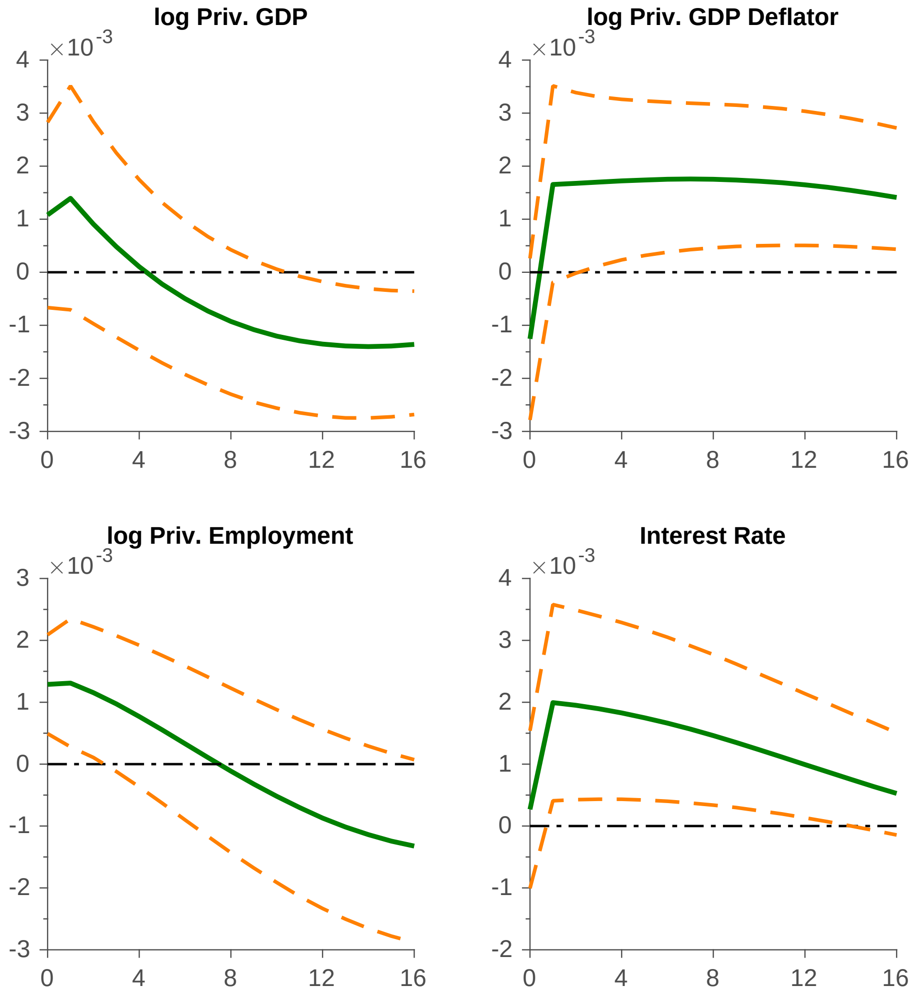

Figure 3.8: Impulse Response Functions - joint shocks to CAPB and FPU. In this experiment we simulate the effects of an expansionary fiscal policy that has been implemented with a substantial degree of uncertainty. Interestingly, the attempt to sustain the economy is completely offset by the adverse uncertainty effects that accompany the fiscal measure. Moreover, pressures on prices and interest rates sum up to reach higher levels than those obtained had the expansionary measure be enacted with no uncertainty.

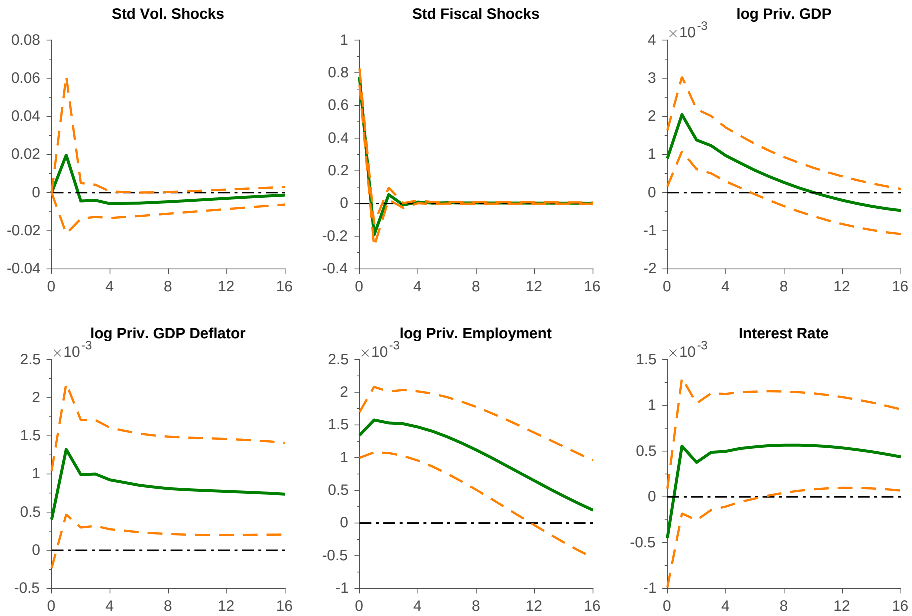

Figure 3.9: Impulse Response Functions - All variables endogenous - CAPB level shock. In this picture we plot robustness check results obtained by inserting the volatility and fiscal shocks series among the endogenous variables, and ordering them respectively first ad second in the VAR. We feed the model with a positive fiscal shock. Results are the same to all practical purposes.

3.3.2 Robustness Checks

In the current section, we report the results of several robustness checks.

Inclusion of the structural shocks among the endogenous variables

We perform these estimates as a robustness check even though we checked that, in our VAR, the fiscal structural shocks are statistically unrelated to the other endogenous variables, i.e. no variable Granger-causes our fiscal shocks. In particular the variables included in $\mathbf{Y}_t$ do not Granger-cause the structural shock ε_t. However, when u_t is considered results are less clear-cut and it might well be possible that there exist a particular lag structure configuration for which a single variable in the vector $\mathbf{Y}_t$ might Granger-cause u_t. In order to show that this possibility does not bias our results, we report also the estimates *à la* Romer and Romer (2010) including ε_t and u_t in the vector of endogenous $\mathbf{Y}_t$. As one can see, results (see Figures 3.9-3.10) are strongly comparable. Incidentally, we checked that changing the order of the variables in the VAR does not change the results.

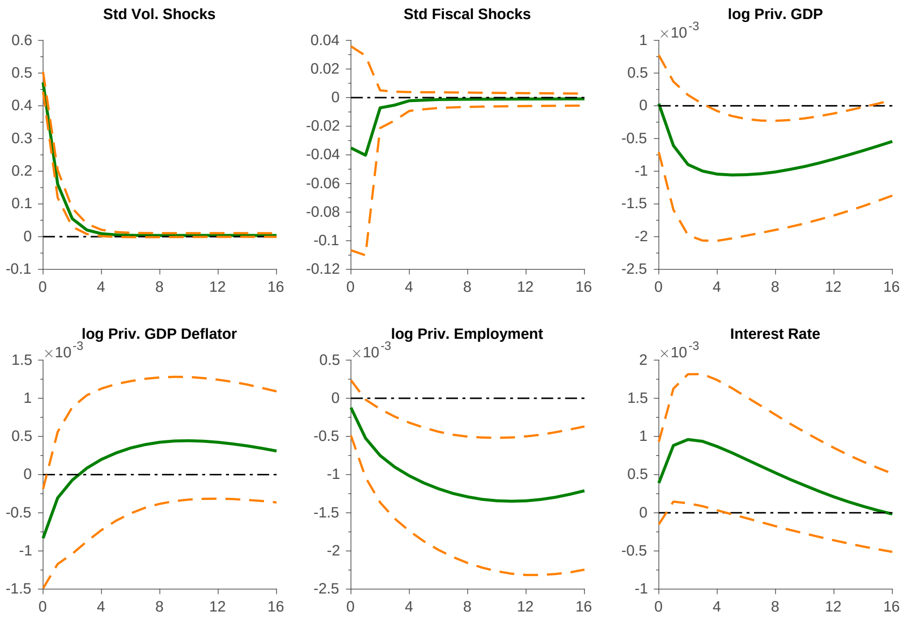

Figure 3.10: Impulse Response Functions - All variables endogenous - FPU shock. This picture shows that our baseline results concerning the effects of FPU on macroeconomic variables are robust to the inclusion of the two shocks series as endogenous variables in the VAR.

Different measures of fiscal balance

We estimated the fiscal rule with different measures of budget deficit. In particular, volatility estimates are robust to using the following dependent variables instead of the CAPB: total borrowing requirement (i.e. including interest outlays), change in the total borrowing requirement, change in the CAPB, cyclically *un*adjusted primary balance, change in the cyclically *un*adjusted borrowing requirement.

Indeed, running the particle filter with all the above measures yields similar filtered estimates for volatility. This is interesting because it means that our volatility estimates do not depend on the measures of budget balance we use in the fiscal rule, which instead is a hotly debated issue[20].

An alternative uncertainty index

Our result that an increase in uncertainty is contractionary is in line with Baker et al. (2016). In our case, though, the shock has a much cleaner interpretation. The uncertainty shock we identify is a pure FPU shock, while the one recovered in Baker et al. (2016) is a generic uncertainty shock which mixes uncertainty coming from fiscal policy with a generic economic policy uncertainty stemming

[20] See, e.g., Golinelli and Momigliano (2009).

from several other sources. Indeed, their EPU index is a weighted average of different series and does not isolate uncertainty coming from fiscal policy only, i.e. it is not a FPU index. Nevertheless, we replicated our VAR analysis using their index and found that the results are qualitatively the same. If one uses each sub-component instead of the general index, it appears that results are driven by the Google-counter news, whereas the series meant to proxy fiscal uncertainty, i.e. the dispersion of the consensus forecast on the budget, had no significant effect on the macroeconomic variables.

Figure 3.11 plots our volatility series together with Baker et al. (2016) index. While the latter is meant to capture a broad concept of uncertainty for Italy, our index is meant to isolate the fiscal policy uncertainty shock. Along this dimension our series may be considered a sub-index of Baker et al. (2016) and we should expect some positive correlation. Indeed, the correlation is equal to around 22%. We observe two periods (specifically the fourth quarter of 2011 and the first quarter of 2013) where Baker et al. (2016) index detects two large spikes in EPU, whereas our FPU series roughly stays constant. Those dates respectively coincide with the government change in November 2011 in the context of the Italian Sovereign debt crisis, and with the 2013 national elections which have seen great political fragmentation, internal divisions within parties, and the emergence of new political forces. The fact that our FPU index does not record significant increases in uncertainty is sensible given that it was *known* that further expansionary policies would not have been pursued, and that fiscal consolidation was about to start.

3.4 Concluding Remarks

The fact that economic uncertainty plays a role in shaping the business cycle should be by now relatively uncontroversial. As John Cochrane puts it, "*the question is: how much uncertainty is there? To what extent and by what mechanism does uncertainty influences GDP, investment and stock prices? The answer is certainly more than zero and less than infinity. As economists we need to look quantitatively at different causes of stagnation*"[21].

In this paper we take in this direction by isolating the uncertainty stemming from a specific source - namely governments decisions about the overall fiscal policy stance - and measuring its effects on the macro-economy. We find that when FPU - captured by a volatility shock in the government fiscal reaction

[21] Quoted in Mordfin (2014).

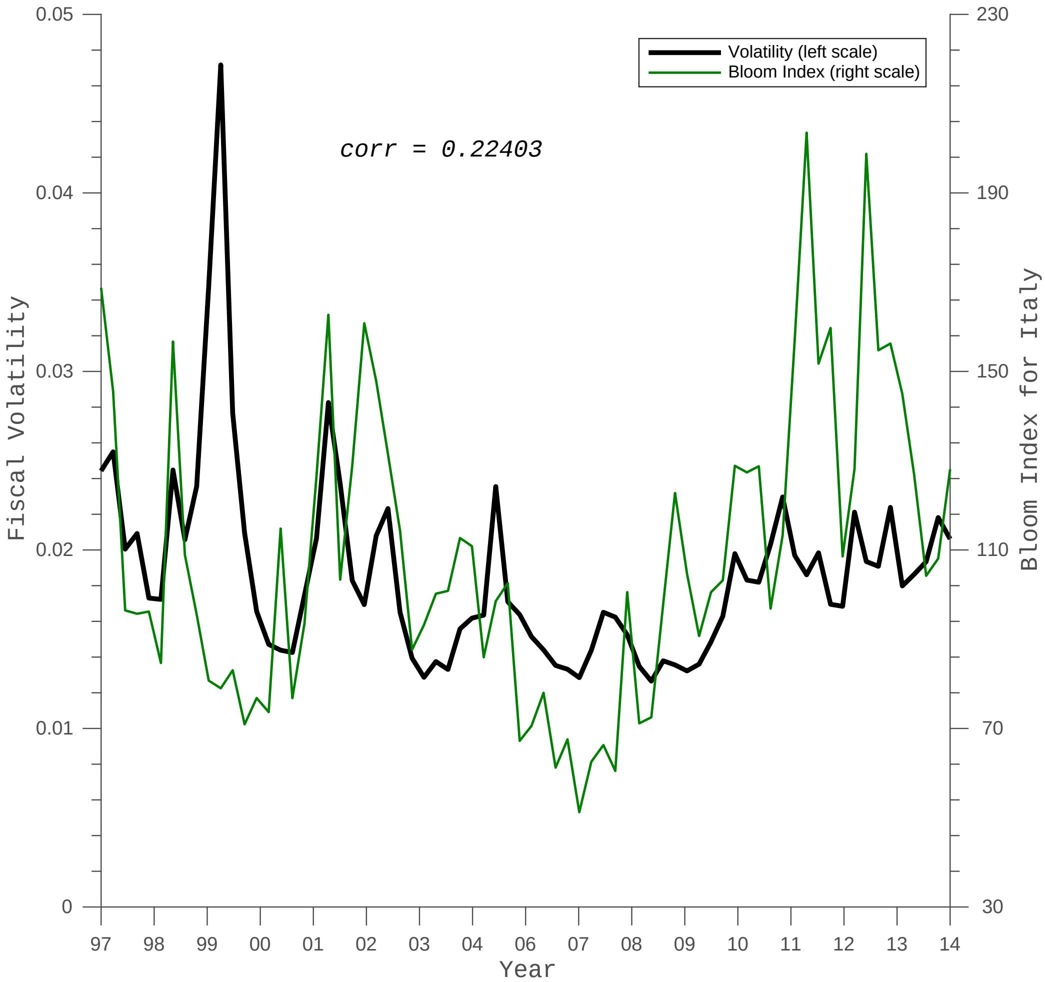

Figure 3.11: This picture compares our Fiscal volatility index against Baker et al. (2016) index. The two series co-move significantly, and the most relevant differences (at the beginning and at the end of the sample) are likely due to fundamental differences in the type of uncertainty measure they attempt to track.

function - unexpectedly increases, both GDP and its components decrease. This result highlights that fiscal policy does not simply have to care about choosing an appropriate deficit level, but it also has to anchor fiscal expectations. The same change in the public deficit may have very different macroeconomic consequences, depending on whether the choice of the government increases or decreases the uncertainty surrounding fiscal policy.

This should be taken into account by econometricians trying to measure the impact of budgetary consolidations and expansions and by fiscal authorities, which should rely on credible and well communicated medium term budgetary frameworks in order to avoid large and sudden policy adjustments.

Bibliography

Abbassi, P., Brownlees, C., Hans, C., and Podlich, N. (2016), Credit Risk Interconnectedness: What Does the Market Really Know? *Mimeo*.

Acemoglu, D., Carvalho, V.M., Ozdaglar, A., and Tahbaz-Salehi, A. (2012), The Network Origins of Aggregate Fluctuations. *Econometrica*, 80(5):1977–2016.

Acemoglu, D., Ozdaglar, A., and Tahbaz-Salehi, A. (2015), Systemic Risk and Stability in Financial Networks. *American Economic Review*, 105(2):564–608.

Alesina, A. and Ardagna, S. (2013), The Design of Fiscal Adjustments. *Tax Policy and the Economy*, 27:19–67.

Alizadeh, S., Brandt, M.W., and Diebold, F.X. (2002), Range-Based Estimation of Stochastic Volatility Models. *The Journal of Finance*, LVII(3):1047–1091.

Allen, F. and Gale, D. (2000), Financial Contagion. *Journal of Political Economy*, 108(1):1–33.

Andersen, T.G., Bollerslev, T., Diebold, F.X., and Ebens, H. (2001), The Distribution of Realized Stock Return Volatility. *Journal of Financial Economics*, 61: 43–76.

Bachmann, R., Elstner, S., and Sims, E.R. (2013), Uncertainty and Economic Activity: Evidence from Business Survey Data. *American Economic Journal: Macroeconomics*, 5(2):217–249.

Bai, J. and Ng, S. (2006), Confidence Intervals for Diffusion Index Forecasts and Inference for Factor-Augmented Regressions. *Econometrica*, 74(4):1133–1150.

Bair, E., Hastie, T., Debashish, P., and Tibshirani, R. (2006), Prediction by Supervised Principal Components. *Journal of the American Statistical Association*, 101(473):119–137.

Baker, S., Bloom, N., and Davis, S. J. (2016), Measuring Economic Policy Uncertainty. *Quarterly Journal of Economics*, 131(4):1593–1636.

Barigozzi, M. and Brownlees, C. (2017), NETS: Network Estimation for Time Series. *Mimeo.*

Barigozzi, M., Brownlees, C., Gallo, G.M., and Veredas, D. (2014), Disentangling Systematic and Idiosyncratic Dynamics in Panels of Volatility Measures. *Journal of Econometrics*, 182:364–384.

Bergmeir, C., Hyndman, R.J., and Koo, B. (2015), A Note on the Validity of Cross-Validation for Evaluating Time Series Prediction. *Monash University Working Paper.*

Billio, M., Getmansky, M., Lo, A.W., and Pelizzon, L. (2012), Econometric Measures of Connectedness and Systemic Risk in the Finance and Insurance Sectors. *Journal of Financial Economics*, 104:535–559.

Blanchard, O. and Perotti, R. (2002), An Empirical Characterization of the Dynamic Effects of Changes in Government Spending and Taxes on Output. *Quarterly Journal of Economics*, 117(4):1329–1368.

Bloom, N. (2009), The Impact of Uncertainty Shocks. *Econometrica*, 77(3): 623–685.

Bloom, N. (2014), Fluctuations in Uncertainty. *Journal of Economic Perspectives*, 28(2):153–176.

Brogaard, J. and Detzel, A. (2015), The Asset-Pricing Implications of Government Economic Policy Uncertainty. *Management Science*, 61(1):3–18.

Brownlees, C., Nualart, E., and Sun, Y. (2015), Realized Networks. *Mimeo.*

Canes-Wrone, B. and Park, J. W. (2014), Elections, Uncertainty, and Irreversible Investment. *British Journal of Political Science*, 44(1):83–106.

Creal, D.D. and Wu, J.C. (2016), Monetary Policy Uncertainty and Economic Fluctuations. *University of Chicago - Booth School of Business, Working Paper.*

Demirer, M., Diebold, F.X., Liu, L., and Yılmaz, K. (2017), Estimating Global Bank Network Connectedness. *Mimeo.*

Diebold, F.X. and Yılmaz, K. (2009), Measuring Financial Asset Returns and Volatility Spillovers, with Application to Global Equity Markets. *The Economic Journal*, 119:158–171.

Diebold, F.X. and Yılmaz, K. (2012), Better to Give Than to Receive: Predictive Directional Measurement of Volatility Spillovers. *International Journal of Forecasting*, 28:57–66.

Diebold, F.X. and Yılmaz, K. (2013), Measuring the Dynamics of Global Business Cycle Connectedness. *Mimeo.*

Diebold, F.X. and Yılmaz, K. (2014), On the Network Topology of Variance Decompositions: Measuring the Connectedness of Financial Firms. *Journal of Econometrics*, 182:119–134.

Diebold, F.X. and Yılmaz, K. (2015)a, Financial and Macroeconomic Connectedness: A Network Approach to Measurement and Monitoring. *Oxford University Press.*

Diebold, F.X. and Yılmaz, K. (2015)b, Trans-Atlantic Equity Volatility Connectedness: US and European Financial Institutions, 2004-2014. *Mimeo.*

Doucet, A., De Freitas, N., and Gordon, N.J. (2001), Sequential Monte Carlo Methods in Practice. *Springer-Verlag: New York.*

Elliott, M., Golub, B., and Jackson, M.O. (2014), Financial Networks and Contagion. *American Economic Review*, 104(10):3115–3153.

Fernández-Villaverde, J. and Rubio-Ramírez, J.F. (2013). Macroeconomics and volatility: Data, models, and estimation. In Acemoglu, D., Arellano, M., and Dekel, E., editors, *Advances in Economics and Econometrics: Theory and Applications, Tenth World Congress of the Econometric Society.* Cambridge University Press, Cambridge, UK.

Fernández-Villaverde, J., Guerrón-Quintana, P.A., Kuester, K., and Rubio-Ramírez, J.F. (2015), Fiscal Volatility Shocks and Economic Activity. *American Economic Review*, 105(11):3352–3384.

Forbes, K.J. and Rigobon, R. (2002), No Contagion, Only Interdependence: Measuring Stock Market Comovements. *The Journal of Finance*, LVII(105): 2223–2261.

Friedman, J., Hastie, T., and Tibshirani, R. (2010), Regularization Paths for Generalized Linear Models via Coordinate Descent. *Journal of Statistical Software*, 33(1).

Gabaix, X. (2011), The Granular Origins of Aggregate Fluctuations. *Econometrica*, 79(3):733–772.

Galí, J. and Perotti, R. (2003), Fiscal Policy And Monetary Integration In Europe. *Economic Policy*, 18(37):533–572.

Giacomini, R. and Rossi, B. (2010), Forecast Comparisons in Unstable Environments. *Journal of Applied Econometrics*, 25:595–620.

Giavazzi, F. and Pagano, M. (1990). Can severe fiscal contractions be expansionary? tales of two small european countries. In Blanchard, O. and Fischer, S., editors, *NBER Macroeconomic Annual*. MIT Press, Cambridge, MA.

Giordano, R., Momigliano, S., Neri, S., and Perotti, R. (2007), The Effects of Fiscal Policy in Italy: Evidence from a VAR Model. *European Journal of Political Economy*, 23:707–733.

Golinelli, R. and Momigliano, S. (2007), The Cyclical Reaction of Fiscal Policies in the Euro Area: The Role of Modelling Choices and Data Vintages. *Fiscal Studies*, 30(1):39–72.

Gomes Francisco, J., Kotlikoff Laurence, J., and Viceira, L.M. (2012), The Excess Burden of Government Indecision. *Tax Policy and the Economy*, 26(1):125–164.

Gordon, N.J., Salmond, D.J., and Smith, A.F.M. (1993), Novel Approach to Nonlinear/non-Gaussian Bayesian State Estimation. *Radar and Signal Processing, IEE Proceedings*, 140(2):107–113.

Grilli, V., Masciandaro, D., and Tabellini, G. (1991), Political and Monetary Institutions and Public Financial Policies in the Industrial Countries. *Economic Policy*, 6(13):341–392.

Handley, K. (2014), Exporting Under Trade Policy Uncertainty: Theory and Evidence. *Journal of International Economics*, 94(1):50–66.

IMF. (2012), How Does Uncertainty Affect Economic Performance? *BOX 1.3 World Economic Outlook*.

Jackson, M.O. (2008), Social and Economic Networks. *Princeton University Press*.

James, G., Witten, D., Hastie, T., and Tibshirani, R. (2013), An Introduction to Statistical Learning. *Springer, New York*.

Julio, B. and Yook, Y. (2012), Political Uncertainty and Corporate Investment Cycles. *Journal of Finance*, 67:45–83.

Jurado, K., Ludvigson, S.C., and Ng, S. (2015), Measuring Uncertainty. *American Economic Review*, 105(3):1177–1216.

Kish, L. (1995), Survey Sampling. *Wiley.*

Kontopoulos, Y. and Perotti, R. (2002), Fragmented Fiscal Policy. *Journal of Public Economics*, 86:191–222.

Koop, G. and Korobilis, D. (2013), Large Time-varying Parameter VARs. *Journal of Econometrics*, 177:185–198.

Koop, G., Pesaran, M.H., and Potter, S.M. (1996), Impulse Response Analysis in Nonlinear Multivariate Models. *Journal of Econometrics*, 74:119–147.

Levin, J. (1993). The cash deficit: Rationale and limitations. In M. Blejer, I. and Cheasty, A., editors, *How to Measure the Fiscal Deficit: Analytical and Methodological Issues*. IMF, Washington, DC.

Liu, J. and West, M. (2001). Combined parameter and state estimation in simulation-based filtering. In Doucet, A., de Freitas, N., and Gordon, N., editors, *Sequential Monte Carlo Methods in Practice*. Springer-Verlag, New York.

Lopes, H. F. and Migon, H.S. (2002), Comovements and Contagion in Emergent Markets: Stock Indexes Volatilities. *Case Studies in Bayesian Statistics*, VI: 285–300.

Luciani, M. and Veredas, D. (2011), A Simple Model for Vast Panels of Volatilities. *Technical Report, ECARES*, (28).

Lütkepohl, H. (2005), New Introduction to Multiple Time Series Analysis. *Springer-Verlag: Berlin.*

Mordfin, R. (2014), The Price of Policy Uncertainty. *Capital Ideas Magazine, Chicago Booth School of Business.*

Newman, M.E.J. (2001), Scientific Collaboration Networks. II. Shortest Paths, Weighted Networks, and Centrality. *Physical Review*, 64(016132):1–7.

Newman, M.E.J. (2003), The Structure and Function of Complex Networks. *Society for Industrial and Applied Mathematics Review*, 45(2):167–256.

Newman, M.E.J. (2010), Networks: An Introduction. *Oxford University Press.*

Park, D., Jun, B., and Kim, J. (1991), Fast Tracking RLS Algorithm Using Novel Variable Forgetting Factor with Unity Zone. *Electronics Letters*, 27:2150–2151.

Parkinson, M. (1980), The Extreme Value Method for Estimating the Variance of the Rate of Return. *The Journal of Business*, 53(1):61–65.

Pesaran, M.H. and Shin, Y. (1998), Generalized Impulse Response Analysis in Linear Multivariate Models. *Economics Letters*, 58:17–29.

Primiceri, G. (2005), Time Varying Structural Vector Autoregressions and Monetary Policy. *Review of Economic Studies*, 72:821–852.

RiskMetrics. (1996), Technical document, fourth edition. *J.P.Morgan/Reuters.*

Romer, C.D. (1990), The Great Crash and the Onset of the Great Depression. *Quarterly Journal of Economics*, 105:597–624.

Romer, C.D. and Romer, D.H. (2010), The Macroeconomic Effects of Tax Changes: Estimates Based on a New Measure of Fiscal Shocks. *American Economic Review*, 100(3):763–801.

Roubini, N. and Sachs, J. (1989), Political and Economic Determinants of Budget Deficits in the Industrial Democracies. *European Economic Review*, 33:903–938.

Scotti, C. (2013), Surprise and Uncertainty Indexes: Real-time Aggregation of Real-activity Macro Surprises. *International Finance Discussion Papers, Board of Governors of the Federal Reserve System*, (1093).

Stone, M. (1974), Cross-Validatory Choice and Assessment of Statistical Predictions. *Journal of the Royal Statistical Society. Series B*, 36(2):111–147.

Zou, H. and Hastie, T. (2005), Regularization and Variable Selection via the Elastic Net. *Journal of the Royal Statistical Society. Series B (Statistical Methodology)*, 67(2):301–320.

Zou, H. and Zhang, H.H. (2009), On the Adaptive Elastic Net with a Diverging Number of Parameters. *The Annals of Statistics*, 37(4):1733–1751.