Three essays on Bayesian econometrics

Georgios Papapanagiotou

University of Macedonia, Greece
October 2023

Three essays on Bayesian econometrics

doctoral dissertation submitted in fulfilment of the requirements for the degree of Doctor of Philosophy

Georgios Papapanagiotou
Department of Economics, University of Macedonia

Supervisor		
Theodore Panagiotidis	Professor	University of Macedonia
Committee members		
Theologs Dergiades	Associate Professor	University of Macedonia
Costas Milas	Professor	University of Liverpool
Defense committee		
Stylianos Fountas	Professor	University of Macedonia
Theologos Pantelidis	Professor	University of Macedonia
Kostas Mouratidis	Lecturer	University of Sheffield
Thanasis Stengos	Professor	University of Guelph

JEL classification codes: C11, C12, C52

Abstract

This thesis deals with the subject of Bayesian econometric methods in time-series analysis in the field of economics and finance. Each chapter constitutes an independent empirical application conducted in a Bayesian framework.

In the first chapter, we employ a Bayesian time-varying parameter Vector Autoregressive (TVP-VAR) model to examine the relation between the price of oil and investor sentiment. To measure investor sentiment, we construct a new proxy based on the search patterns of individuals on the Google engine. Using this new proxy, oil prices as well as benchmark macroeconomic and financial variables, we estimate a TVP-VAR that takes into account the changes in the transmission of investors sentiment shocks to oil prices over time. The results indicate that an unexpected increase in investor attention yields a long-lasting increase both in the price of oil and the stock market returns.

In the second chapter, we use alternative Bayesian Markov-switching Generalised Autoregressive Conditional Heteroscedasticity (MS-GARCH) models to analyse the behaviour of volatility of cryptocurrencies. In total, we consider 292 cryptocurrrencies for each of which we estimate the estimate 27 alternative MS-GARCH specifications. First, we evaluate the in-sample performance of each model using the information criteria. Next, we assess the ability of the models to perform one-day ahead conditional volatility and Value-at-Risk forecasts. The results indicate that for a wide range of cryptocurrencies (with different characteristics), Markov-switching models which two or more regimes outperform the ones with a sole regime. In addition, the findings suggest the presence of inverse leverage effect in the majority of the cryptocurrencies.

In the third chapter, we propose a new a methodology to test examine for Granger-causality in a time-varying framework. Specifically, we combine the estimates from a TVP-VAR with the null hypothesis of no Granger-causality that allows us to track changes in the causal relationship of variables in each period. This methodology offers several advantages compared the existing ones which mostly rely on rolling window algorithms. The performance (size and power) of the proposed methodology is evaluated through Monte Carlo simulations. As an empirical application, we exam-

ine the evolution of Granger-causal relationship between bitcoin returns and alternative variables. According to the results, other cryptocurrency returns, stock market returns and uncertainty Granger cause bitcoin returns during periods when bitcoin prices burst and bitcoin's trading volume Granger causes bitcoin returns during periods when bitcoin prices remain relatively steady.

Acknowledgements

The completion of this work could not have been possible without the contribution of many individuals to whom I express my deepest gratitude. First and foremost I would like to thank the supervisor of this work, professor Theodore Panagiotidis. His advice and guidance have been invaluable through the years, even before the beginning of this PhD. I would also like to thank the members of advisor board Theologos Dergiades and Costas Milas as well as Kostas Mouratidis, Thanasis Stengos, Georgios Bampinas and Jesus Otero for their help and cooperation through the years.

Next, I would like to thank my fellow PhD students for making working at the university's office more pleasant, both as colleagues and as friends.

The research work was supported by the Hellenic Foundation for Research and Innovation (HFRI) under the 3rd Call for HFRI PhD Fellowships (Fellowship Number: 81916).

Finally, I would like to thank my friends, my parents and my brother. Without their guidance and unconditional support and love, this goal would not have been accomplished.

Contents

List of Figures

Chapter 1

Introduction

Bayes' theorem was formulated by Thomas Bayes in 1763. Fifty years later, in 1814, Pierre-Simmon Laplace developed the Bayesian interpretation of probability. Since then, many authors have introduced many Bayesian methods for statistical inference. However, the computational cost of Bayesian methods made undesirable for most statisticians for more than one century.

The dawn of the 21st century witnessed the surge of powerful computers which facilitated the implementation of Bayesian methods. Concurrently, methodological approaches such as Markov Chain Monte Carlo (MCMC) algorithms which are used for Bayesian estimation have been further developed over the last decades.

From the methodological point of view Bayesian statistics and econometrics offer two major advantages compared to the classical frequentist inference in time series analysis. First, Bayesian models can be flexible enough to account for non-linearities in economic and financial data, since the assumption of constant parameters in a model is rather a poor one. For instance, it is documented that incidents like the global financial crisis in 2008 and the COVID-19 pandemic caused structural changes in both the macroeconomic and the financial variables.

Second, Bayesian methods a suitable for models with a large a number of variables. For example, Vector Autoregressive (VAR) models, a particularly popular class of models among economists, can often contain more than four or five endogenous variables. If we combine the large number of variables with the assumption of time-varying parameters, we end up with possibly hundreds of parameters to estimate. In this case, the challenge is to estimate such models without being affected by over-parameterisation issues. The Bayesian approach through the appropriate selection of shrinkage priors can identify the insignificant parameters and shrink them towards zero.

In this doctoral dissertation, we employ alternative Bayesian econometric methods

to conduct empirical applications. Each of the following chapters consists of independent essays analysing economic and financial time-series using Bayesian inference. In the first essay, we examine the relationship between the price of oil and investor attention. We begin the analysis by revisiting seminal VAR models which examine the relationship between oil prices, oil supply, oil demand and stock market returns. Next, we augment these models by adding a new index for investor attention. The index is of our construction and is based on Search Volume Index data for oil-related search terms in the Google engine. Our index offers a number of advantages compared to traditional indices and robustness checks reveal that this new index can explain better the evolution of the price of oil. All VAR models allow all parameters to gradually evolve over time. This assumption enables us to track changes in the transmission mechanisms of a shock from one variable to the other. Our findings reveal that a shock in investor attention causes a long-lasting increase both in the price of oil and the stock market returns. The results remain qualitatively when additional sentiment indices, such as proxies for uncertainty, are added to the model.

In the second essay, we evaluate the performance of models belonging in the family of Markov-switching Generalised Autoregressive Conditional Heteroscedasticity (MS-GARCH) models in modelling the volatility of 292 time-series of cryptocurrency returns. For each cryptocurrency we consider a total of 27 alternative MS-GARCH specifications. We focus on the number of regimes in the model and consider Markov-switching models with one, two and three states. In the first part of the analysis, we evaluate the goodness-of-fit of each model using the Deviance Information Criterion and the Bayesian Predictive Information Criterion. In the second step, we focus on the out-of-sample performance of the MS-GARCH models by analysing their ability to perform one-step-ahead and Value-at-Risk forecasts. The results indicate that while Markov-switching models with two states outperform the traditional ones (with a single state), further increasing the number of states does not improve the performance of the model. Furthermore, the results suggest the presence of inverse of leverage effect in the examined time-series.

In the last essay, we propose a new a methodology to test examine for Granger-causality in a time-varying framework. Specifically, we combine the estimates from a TVP-VAR with the null hypothesis of no Granger-causality that allows us to track changes in the causal relationship of variables in each period. This methodology offers several advantages compared the existing ones which mostly rely on rolling window algorithms. The performance (size and power) of the proposed methodology is evaluated through Monte Carlo simulations. As an empirical application, we examine the evolution of Granger-causal relationship between bitcoin returns and alterna-

tive variables. According to the results, other cryptocurrency returns, stock market returns and uncertainty Granger cause bitcoin returns during periods when bitcoin prices burst and bitcoin's trading volume Granger causes bitcoin returns during periods when bitcoin prices remain relatively steady.

Chapter 2

Oil shocks and investor attention

Abstract

We examine the existence of sentiment exposure in oil price returns. We augment Kilian's (*American Economic Review*, 2009, 99, 1053–1069) SVAR model including the effects of (1) stock market returns as in Kilian and Park (*International Economic Review*, 2009, 50, 1267–1287), (2) investors sentiment proxied by Google search volume index, (3) economic policy uncertainty (EPU), (4) time variation in both coefficients and the variance-covariance matrix and (5) an extended sample. Our empirical results show that changes in investor attention do exhibit a significant long-lasting impact on oil and stock market returns. Aggregate oil demand and supply shocks have a transitory effect on investor sentiment. We reveal that the impact of EPU is significant and short-lasting, while EPU responds strongly to shocks on oil prices and stock market returns. In all cases, the magnitude and sign of responses are affected by the timing of the shock. Our findings are robust to an alternative sentiment indicator and once the role of oil inventories is considered.

2.1 Introduction

Over the last two decades, energy commodities have become a popular asset class for financial institutions and retail investors, similar to equities and bonds. Meanwhile, oil price fluctuations stimulated the academic and policy debate regarding the impact of fundamental shocks such as supply and demand shocks, inventory shocks, or macroeconomic and monetary shocks on the price of oil (see for instance Kilian, 2008, 2009; Leduc and Sill, 2004; Kilian and Murphy, 2014; Baumeister and Kilian, 2016). The price of crude oil may also be related to non-fundamental shocks, such as expectations or herd behaviour. Moreover, oil price shocks have evolved over time and have effects on the real economy through consumer and firm behaviour that are not constant (Blan-

chard and Riggi, 2013; Hamilton, 2013; Kang et al., 2015). This essay offers new insights by examining the relationship between oil prices and investors sentiment using a time-varying structural vector autoregression (SVAR) model.

With regard to the direction of causality between the crude oil prices and investors sentiment, the literature follows two main branches. One the one hand, under the efficient market hypothesis, the traditional asset pricing models assert that investors are rational and that information is instantaneously incorporated into asset prices. Departures from efficiency have mainly been investigated through some form of investors' reaction to new information (Barberis et al., 1998; Hong and Stein, 1999). In this respect, investors' attention or sentiment must play a significant role in asset prices and returns. A low level of investor attention may lead to small fluctuations in oil prices, while information that has received broad attention from investors will instantly be comprised in prices, resulting in large fluctuations (Vozlyublennaia, 2014). On the other hand, oil price fluctuations may lead to pessimistic expectations about future economic conditions and to a subsequent reduction in consumer and investor sentiment, which induces households and firms to curb their consumption and investment expenditures, respectively (Guntner and Linsbauer, 2018).

One limitation of existing work relates to how investor's sentiment can be quantified. The empirical literature have taken several approaches to measuring investor sentiment. Some studies proxy for investor sentiment with market-based measures (such us trading volume, implied volatility or IPO volume and returns). Although these market-based measures have the advantage of being readily available at a relatively high frequency, they have the disadvantage of being the average outcome of many economic factors apparent from investor sentiment (Da et al., 2015). Others, use survey-based indices such as the University of Michigan Consumer Sentiment Index, the UBS /GALLUP Index for Investor Optimism, or investment newsletters. However, the survey measures are often available at a low frequency (monthly or quarterly), while in many cases the non-response rates are high and incentive for truth-telling is low, especially in personal or sensitive questions. More recently, Brandt and Gao (2019) use textual analysis to examine whether macroeconomic and geopolitical news sentiment have an affect on oil prices (for a comprehensive survey on textual analysis and sentiment see Algaba et al., 2020). Yet, there is a reasonable concern about the severity and impact of media bias. Sentiment in texts can stem from self-interest to generate a particular external outcome (Garz, 2014). Moreover, linguistic sentiment entails a degree of subjectivity either as a positive or as a negative opinion through language. In this respect, is quite complex to define the real sentiment disposition toward an entity.

We address both of these limitations by constructing the Google Search Volume

Index (GSVI) based on Search Volume Index (SVI) data by Google Trends and using this new index as a direct measure of investors attention on the crude oil market (for applications in financial markets see Da et al., 2011; Vlastakis and Markellos, 2012; Vozlyublennaia, 2014; Da et al., 2015). The advantage of using Google search data is that we can gauge people's active internet search in specific geographic areas at a high frequency (unlike surveys) and study their follow-up actions, as investors trade on their beliefs and move energy market prices accordingly. Moreover, as investors take the initiative to search online information about the oil market, the Google search index possesses their own attention behaviour, so that bias is eliminated (Da et al., 2011). Sockin and Xiong (2015) argue that key industrial commodities, such as crude oil, often serve as important price signals regarding the strength of the global economy for market participants. Indeed, in Figure 2.1 we observe a strong co-movement of GSVI with the St. Louis Financial Stress Index (STLFSI), especially during the two most recent recession periods.[1] Therefore, Google searches can be considered a good proxy not only for aggregate investors' attention to oil markets but also for the overall outlook of the economy.

Figure 2.1: Oil prices, Google Search Volume Index and the Financial Stress Index over the period 2004-2020. Oil prices, measured as refiner's acquisition cost of crude oil, are plotted against the right axis. The shaded areas denote the two most recent recession periods.

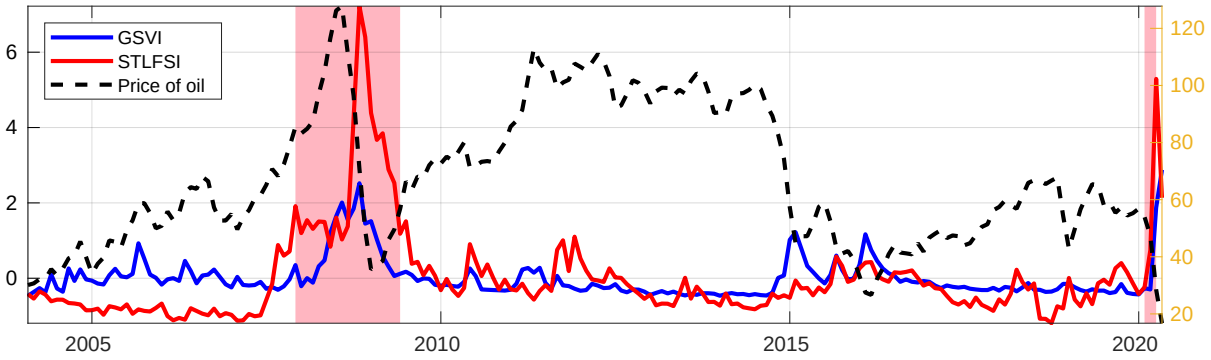

Apart from investors' sentiment, economic policy uncertainty may exert a great influence on real economic activity. A number of studies show that the supply-side oil shocks are relatively unimportant for the macro-economy compared to demand-side oil price shocks (Kilian, 2009; Hamilton, 2009a; Lippi and Nobili, 2012; Baumeister and Peersman, 2013); this finding is confirmed by Antonakakis et al. (2014) and Kang and Ratti (2013) who investigate the relationship between structural oil price shocks

[1]The St. Louis Financial Stress Index (STLFSI) quantifies financial stress in the U.S. economy using 18 key indicators of financial market conditions - 7 interest rates, 6 yield spreads, and 5 other indicators.

and economic policy uncertainty. A widely used measure of uncertainty is the Economic Policy Uncertainty (EPU) index provided by Baker et al. (2016). The index is a weighted average of four uncertainty components, namely the news-based policy uncertainty index, the tax expiration's index, the CPI forecast disagreement measure and the federal/state/local purchases disagreement measure. In this respect, EPU assesses macroeconomic uncertainty by combining economic uncertainty related to economic policy making and public views. Baker et al. (2016) argue that a rise in EPU is related with an increase in price volatility and a decrease in investment.

Our empirical analysis builds on the structural vector autoregressive (SVAR) approach and extents the Kilian and Park (2009) model by including the investor's sentiment index for oil (GSVI). The extended model consists of five endogenous variables, the change in oil production, the global economic activity, the real oil prices, the stock market returns and the GSVI index. We further allow for time-variation in both the parameters and the standard deviation by employing a Bayesian estimation of the SVAR model.[2] Our estimates show that the GSVI_{t-1} coefficient is negative in the stock returns equation for the entire sample period. Arguably, increases (decreases) in GSVI correspond with low (high) stock returns, a finding consistent with the sentiment induced mispricing hypothesis. Similar pattern is observed in the equation of real price of oil, but this relationship reverses after 2015 to positive. We also find that the standard deviation for all models increases considerably during the Global Financial Crisis in 2009, except the equation of stock market returns, where the standard deviation of the disturbance term reaches its highest value during the recent Covid-19 pandemic.

We further examine the importance of evolving parameters through the analysis of impulse response functions. We document that a positive shock in the GSVI variable yields a decrease in oil production which lasts only for two months. From this point on, oil production sharply increases and reaches pre-shock levels. Similar results are found in the responses of oil demand and oil prices. For the three aforementioned variables, the impulse responses are time invariant. Although, this is not the case with stock market: if the shock occurs before 2015, stock market returns decline after a couple of periods and if the shock occurs after 2015, stock market returns gradually increase. In both cases, the effect of the shock is both persistent and statistically significant.

Considering the role of economic policy uncertainty on firm level decisions and real economic activity, we subsequently estimate an alternative model by adding in the five-variable model the EPU index of Baker et al. (2016) as an indicator for economic

[2] Arampatzidis et al. (2021) show that the effect of the price of oil on the stock market depends both on the type and timing of the shock.

policy uncertainty. Higher volatility in oil prices has been linked with increased uncertainty at firms. Elder and Serletis (2010) and Rahman and Serletis (2011) show that uncertainty about changes in the real oil prices has a negative effect on real economic activity. Moreover, Bird and Yeung (2012) show that during periods of high (low) market uncertainty investors tend to ignore good (bad) news. In this regard, the aim of this exercise is to uncover the time-varying role of economic policy uncertainty in the crude oil - investor sentiment relationship.

Sensitivity analysis shows that our baseline findings are robust. First, we employ as an alternative sentiment measure the Index of Consumer Expectations (ICE), produced by the University of Michigan, which summarizes a monthly aggregation of around 500 survey respondents' expectations regarding the future outlook for their own personal finances and for the overall U.S. economy (see, e.g., Carroll et al., 1994; Matsusaka and Sbordone, 1995; Lagerborg et al., 2020). Second, we use specifications that consider the role of oil inventories to show that the results are not driven by an "identification problem" (see for example Hamilton, 2009a,b; Kilian and Murphy, 2014; Baumeister and Hamilton, 2019).

This essay contributes to the literature in several aspects. A great deal of literature has focused on verifying the impact of investor sentiment on crude oil returns by applying a limited number of search terms. This essay constructs a market specific investor sentiment index for crude oil markets based on a much wider set of attention terms and using the dynamic factor model (DFM) framework. This framework allows us to consider dynamics both in the factors and the idiosyncratic component, as well as heteroscedasticity in the idiosyncratic variance. Based on Bayesian estimation, this essay also studies the changing effects of the degree of investor sentiment in crude oil markets. Extending the model of Kilian and Park (2009) with the oil GSVI index, we find that information demand responds temporarily to oil supply and demand shocks and shocks in stock market returns. Also, by allowing for time-variation we find that these responses are time dependent (for example, a shock in oil prices after 2015 has a greater impact on GSVI). We further examine the importance of the changes of parameters of impulse response functions, where the sources of time variation are both the coefficients and the variance covariance matrix of the innovations. Doing so, we may identify changes in the impulse responses, mainly in magnitude but for some cases also in the sign of the response. Furthermore, to the best of our knowledge, this is the first paper that explores the role of economic policy uncertainty in the linkage between investor attention and crude oil market through a time-varying framework. Our findings show that the effect of policy uncertainty is time-varying and becomes significant and positive only during the period preceding the 2009 banking crises.

The rest of the essay is organised as follows: Section 2 describes the methodology. Section 3 describes the data. Section 4 presents the main findings of the study. Section 5 investigates the robustness of our results. The last section concludes.

2.2 Econometric methodology

Initially we construct the index for investors' sentiment based on internet searches for oil-related terms. The construction of the sentiment index is described in section 2.2.1. Once the new index is created, we estimate a TVP-VAR model with five variables: change in oil production, global real economic activity, real oil prices, stock market returns and investors' sentiment. We examine the effect of investors' sentiment on oil price fundamentals and stock market returns by examining the evolution of the estimated coefficients over time. In addition, we employ impulse response function analysis to examine the bilateral relationship between investors' sentiment and oil supply, oil demand, oil prices and stock market. The estimation of the TVP-VAR model is described in section 2.2.2. In the last part of the analysis we extend the main model by including a variable for economic policy uncertainty. In this case, we focus on the effect of a shock from and to the policy uncertainty variable.

2.2.1 Construction of the Google Search Volume Index

We start by constructing a list of sentiment-reveal search terms that are related to the crude oil market. We initiate the analysis with primary keywords such as "oil", "oil price", "crude oil","oil spot", "oil stock" etc.[3] More specifically, we have considered each search term in Google Trends search engine and we have kept the ten "top searches" related to each term. This generates approximately 300 related keywords for oil. Next, we remove duplicates, terms with no economic meaning and terms with inadequate data (query series with zero search volume). This leaves us with the final 186 oil-related search terms.

To construct the Google Search Volume Index (GSVI) we download from Google Trends the monthly SVI for 186 search terms related to oil, from January 2004 to May 2020 and employ the dynamic factor model (DFM) framework.[4] DFM is appropriate for handling large data sets without much loss of information.[5] In addition, DFM al-

[3]Table B1 reports all keywords used for the construction of the index.

[4]All time-series are stationary considering two unit root tests (the Augmented Dickey-Fuller test and the Phillips-Perron test) and are properly scaled to the highest value within the whole sampling period.

[5]For example, principal components analysis can lead to information loss if the number of components is not specified correctly.

lows for more flexible interpretation compared to principal components analysis and provides a better fit to the data (Sargent and Sims, 1977; Stock and Watson, 2016).

The DFM is written in the state space representation as:

$$x_t = \sum_{i=0}^{s} \Lambda_i f_{t-i} + e_t \tag{2.1}$$

$$f_t = \sum_{t=1}^{p} A_i f_{t-i} + u_t \tag{2.2}$$

where x_t is an $N \times 1$ vector of observable variables, Λ_i are $N \times q$ matrices of factor loadings, f_t are $q \times 1$ vectors of unobserved common dynamic factors which summarise the cross-covariance properties of x_t. The idiosyncratic component e_t, are $N \times 1$ stationary processes uncorrelated with f_t. The elements of e_t are weakly correlated cross-sectionally and/or serially. The VAR(p) model in equation 2.2, with $A_1 \ldots, A_p$ the matrices of autoregressive coefficients, approximates the dynamics of the common factors. The two error vectors are assumed to be independent.

The DFM described by equations 2.1 and 2.2 can be transformed into a static state space representation. We define the state vector as $F_t = (f_t', \ldots, f_{t-k+1}')'$ where $k = \max\{s+1, p\}$. If $k > p$ then $A_{p+1} = \cdots = A_k = 0$ in the companion matrix:

$$\mathbf{A} = \begin{bmatrix} A_1 & \ldots & A_p & \ldots & A_k \\ I_q & 0_q & . & . & 0_q \\ 0_q & \ddots & \ddots & . & \vdots \\ \vdots & \ddots & \ddots & . & 0_q \\ 0_q & . & 0_q & I_q & 0_q \end{bmatrix}$$

if $s+1 < k$, we set $\Lambda_{s+1} = \cdots = \Lambda_k = 0$ in the loadings matrix $\Lambda = (\Lambda_0\ \Lambda_1 \ldots \Lambda_k)$. The static factor form of the state-space model is given by the following equations:

$$x_t = \Lambda F_t + e_t$$
$$F_t = \mathbf{A} F_{t-1} + u_t^*$$

with $u_t^* = (u_t', 0_{1\times q}, \ldots, 0_{1\times q})'$. The number of estimated static factors $K = q(s+1)$ is determined using the information criterion methods provided by Bai and Ng (2002). The number of dynamic factors given a preselected number of static factors is computed using the four criteria Bai and Ng (2007).

The factors are estimated using Quasi ML - EM estimator developed by Doz et al. (2012) by iterating the two-step estimator of Doz et al. (2011). First an 'extended' state

space model is defined, with one more lag than the representation (2.1),

$$x_t = \sum_{i=0}^{s} \Lambda_i f_{t-i} + \mathbf{0}_{N\times q} f_{t-s-1} + e_t$$

where $\mathbf{0}_{N\times q}$ indicates an $N \times q$ null matrix. The static representation of this system then becomes:

$$\begin{aligned} x_t &= \grave{\Lambda} G_t + e_t \\ G_t &= \grave{\mathbf{A}} G_{t-1} + \grave{u}_t^* \end{aligned}$$

where $G_t = (F_t', f_{t-k}')' = (f_t', \ldots, f_{t-k+1}', f_{t-k}')'$ and the system matrices are redefined accordingly. This modification of the state space representation is necessary to make it possible to compute rapidly and efficiently all the quantities that are required in the M-step. Therefore, given a preliminary estimate of $\hat{\theta}$, the 'extended' system matrices are computed and the Kalman smoother is run, so as to get an unbiased predictor of the factors $\hat{G}_t$ with the associated covariance matrices. This is the E-step. Subsequently, the estimated factors are used to compute the sufficient statistics:

$$\begin{aligned} \hat{F}_t &= E(F_t|\mathbf{X}) \\ \hat{P}_{0,t} &= E(F_t F_t'|\mathbf{X}) \\ \hat{P}_{1,t} &= E(F_t F_{t-1}'|\mathbf{X}) \end{aligned}$$

where $\mathbf{X} = x_1, \ldots, x_T$, which are used in the M-step to compute (conditional) ML estimates of θ. Note that, given the covariance matrix of G_t, the calculation of $\hat{P}_{0,t}$ and $\hat{P}_{1,t}$ simply amounts to selecting appropriate submatrices. With the new estimate of θ in hand, the process is repeated until convergence. The first $q-$ element subvector $\hat{f}$ of $\hat{F}_t$, the ML estimate of the dynamic factor, is used as the GSVI.

To control convergence of the EM algorithm we employ a parameter distance criterion that adds the absolute deviation across all estimated parameters $\hat{\theta}$ between two consecutive steps. More formally, we calculate:

$$c_P^j = \frac{\sum_{i=1}^{h} |\hat{\theta}_I^{(j)} - \hat{\theta}_I^{(j-1)}|}{h} \tag{2.3}$$

where h is the number of elements in $\hat{\theta}$. We stop EM process when $c_M^P < 10^{-2}$.

Figure 2.1 plots the constructed index over the full sampling period, 2004M01-2020M04.

2.2.2 Estimation of the TVP-VAR model

Based on Primiceri (2005), Del Negro and Primiceri (2015) and Koop and Korobilis (2009) we consider the following time-varying parameter VAR model:[6]

$$y_t = c_t + \sum_{i=1}^{p} B_{i,t} y_{t-i} + u_t, \quad t = 1, \ldots, T \tag{2.4}$$

where y_t is an $n \times 1$ vector of endogenous variables, c_t is a $n \times 1$ vector of time-varying coefficients, $B_{i,t}$, $i = 1, \ldots, P$ are $n \times n$ matrices of time-varying coefficients and u_t are heteroscedastic shocks with variance-covariance matrix Ω_t. The lag order p is determined through the Schwarz information criterion (SIC). We can rewrite (2.4) as:

$$\begin{aligned} y_t &= X_t^{'} B_t + A_t^{-1} \Sigma_t \epsilon_t \\ X^{'} &= I_n \otimes [1, y_{t-1}, \ldots, y_{t-p}] \end{aligned} \tag{2.5}$$

where $B_t = [c_t, \text{vec}(B_{1,t}), \ldots, \text{vec}(B_{p,t})]$ and $\otimes$ denotes the Kronecker product and vec() the columnwise vectorisation of a matrix. A_t is a lower unitriangular matrix and Σ_t is a diagonal matrix such that, $A_t \Omega_t A_t^{'} = \Sigma_t \Sigma_t^{'}$. It follows that $\text{Var}(\epsilon_t) = I_n$. Cogley and Sargent (2002) consider similar decompositions of the variance-covariance matrix of a time-varying VAR model but with a time invariant matrix A_t. If α_t is the vector of unrestricted elements of A_t and σ_t the vector of diagonal elements of Σ_t, then the state equations are:

$$B_{t+1} = B_t + v_{t+1}, \quad v_t \sim N(0, Q) \tag{2.6}$$

$$\alpha_{t+1} = \alpha_t + \zeta_{t+1}, \quad \zeta_t \sim N(0, S) \tag{2.7}$$

$$\log \sigma_{t+1} = \log \sigma_t + \eta_{t+1}, \quad \eta_t \sim N(0, W) \tag{2.8}$$

The cardinalities of B_t, a_t and σ_t are $n_B = n + pn^2$, $n_A = \frac{n(n+1)}{2}$ and $n_\sigma = n$. All errors in the model are assumed to be jointly normally distributed with

[6] In this section we drop any notation used in the section 2.2.1.

$$\text{Var}\left(\begin{bmatrix} \epsilon_t \\ v_t \\ \zeta_t \\ \eta_t \end{bmatrix}\right) = \begin{bmatrix} I_n & 0 & 0 & 0 \\ 0 & Q & 0 & 0 \\ 0 & 0 & S & 0 \\ 0 & 0 & 0 & W \end{bmatrix}$$

As in Cogley and Sargent (2005) and Primiceri (2005), the prior distributions (initial conditions at $t = 0$) are specified and updated using a training sample. The training sample is used to estimate a time invariant VAR with Ordinary Least Squares (OLS). For the parameters B_0, A_0 and $\log \sigma_0$, the Normal distribution is assumed. The mean and variance of the parameters B_0 and A_0 are the mean and four times the variance of the respective estimated parameter of the OLS-VAR. In the case of $\log \sigma_0$ the mean of the prior distribution is chosen to be the logarithm of the OLS point estimates of the standard errors of the OLS-VAR model and the the variance covariance matrix is assumed to be the identity matrix. For the priors on error co-variances Q, S and W the hyperparameters are set to $k_Q = 0.01$, $k_S = 0.1$ and $k_W = 0.01$. Summarising, the priors take the form:

$$\begin{aligned} B_0 &\sim N(\hat{B}_{OLS}, 4\text{Var}(\hat{B}_{OLS})) \\ A_0 &\sim N(\hat{A}_{OLS}, 4\text{Var}(\hat{A}_{OLS})) \\ \log \sigma &\sim N(\log \hat{\sigma}_{OLS}, I_n) \\ Q^{-1} &\sim W(1 + n_B, ((k_q^2)(1 + n_B)\text{Var}(\hat{B}_{OLS}))^{-1}) \\ S^{-1} &\sim W(1 + n_A, ((k_q^2)(1 + n_A)\text{Var}(\hat{A}_{OLS}))^{-1}) \\ W^{-1} &\sim W(1 + n_\sigma, ((k_q^2)(1 + n_\sigma)I_n)^{-1}) \end{aligned}$$

Let M_t be a generic matrix. We define $M^\tau = [m_1', \ldots, m_\tau']'$, where $m_t = vec(M_t)$. Our aim is to estimate the posterior distributions of the unobservable states B^T, A^T and Σ^T and the hyperparameters of the variance covariance matrix V. The Bayesian estimation requires a Markov chain Monte Carlo (MCMC) algorithm. To exploit the blocking structure of the unknowns, a Gibbs sampling consisting of four steps is used. In each step B^T (the time-varying coefficients), A^T (the simultaneous relations), Σ^T (the stochastic volatility) and V (the hyperparameters) are drawn respectively, conditional on the observable data and the other parameters.

The state space model described in 2.5 and 2.6 is linear and Gaussian (conditional on A^T and Σ^T). This implies that B^T, A^T is a product of Gaussian densities and can be

drawn using a Gibbs sampler algorithm.[7] Here, we rely on the algorithm of Carter and Kohn (1994).[8] Since Σ^T is not a Gaussian product, we rely on the method of Kim et al. (1998) to transform it to a linear, Gaussian form. As a result, it can be drawn using the algorithm of Carter and Kohn (1994). Finally, we sample V by sampling Q, W and S independently from an inverse-Wishart distribution.

Once the reduced-form VAR model (2.4) is estimated, we construct the following structural VAR:

$$y_t = X_t^{'} B_t + \Xi \epsilon_t$$

where Ξ_t contains the necessary restrictions for structural identification and $\Xi_t \Xi_t' = \Omega_t$. Since the identification is based on a lower triangular scheme, it suffices to set $\Xi = A_t^{-1} \Sigma_t$.

The results are based on posterior sample of 50000 draws. The first 25000 draws are considered a burnin sample and are discarded. To assess the validity of our results we considered two robustness checks. In the first one, we increase the number of both burnin and posterior draws to 50000. The second check aims to take into account any possible autocorrelation among draws, (Korobilis, 2017). To this end, we build the posterior sample by keeping only every 10$^{\text{th}}$ draw. Both robustness checks provide similar results with the ones reported in the essay. These results are available upon request.

2.3 Data

The study utilises monthly time-series data over 2004M01-2020M04. Data for the oil production and prices of oil are obtained from the U.S Department of Energy. World crude oil production is measured in millions of barrels per day averaged monthly. Following Kilian (2009) we use the percent change in the oil supply by calculating the log differences and multiplying with 100 (denoted $\Delta prod_t$). To obtain the real price of oil (denoted rpo_t), we divide the refiner's acquisition cost of crude oil by the U.S. CPI (obtained from the Bureau of Labor Statistics).[9] The index of real aggregate demand, constructed by Kilian (2019) is used as an indicator of global real economic activity(denoted rea_t). The logarithm of closing prices of S&P500, deflated by the U.S. CPI

[7] A^T is drawn conditional on B^T and Σ^T

[8] Alternatively, one could employ the algorithm of Frühwirth-Schnatter (1994).

[9] Kuck and Schweikert (2017) and Galay (2019) argue that world oil markets maintain a long-run equilibrium.

[9] The index is an updated and corrected version of the Kilian (2009) index. See also Kilian and Zhou (2018) and Hamilton (2021).

is used as the stock market variable (denoted ret_t). As a proxy for economic policy uncertainty, we employ the index constructed by Baker et al. (2016) which is a weighted average of four uncertainty components (denoted epu_t). Finally, the sentiment index, GSVI, is constructed as described in subsection 2.2.1. SVI data used for the construction of the GSVI is available from 2004M01. The availability of SVI determines the sample size in our analysis.

In our analysis, we also perform robustness checks regarding the first main (five-variables) model. First, we replace GSVI with Michigan's University Index of Consumer Sentiment (ICS) (see section 2.5). ICS is constructed using the responses to five different questions that are part of a broader survey of consumer attitudes and expectations. Data for the ICS are available on monthly basis starting from January 1978 from the website of Michigan's University. Second, we examine how the results are affected when we take oil inventories into account. To construct the variable of crude oil inventories, we multiply the U.S. inventories by the ratio of OECD inventories (see also Hamilton, 2009a; Kilian and Murphy, 2014; Baumeister and Hamilton, 2019). We obtain data for the U.S. crude oil stocks (in millions of barrels) and for the OECD inventories of crude petroleum and petroleum products from the Energy Information Administration and the OECD databases, respectively. Table B2, in the Appendix, reports the summary statistics for all variables used in the analysis.

2.4 Results

2.4.1 A reexamination of the models of

Before we proceed with the main model, we revisit the seminal models of Kilian (2009) and Kilian and Park (2009) which are the base of the main models of this work. Kilian (2009) considers a trivariate model where the endogenous variables are the change in world oil production, the global real economic activity and the real price of oil. Since we do not use the GSVI in this model, we employ a longer sample, spanning from 1974M01 to 2020M04. In the case of constant parameters, we calculate the impulse response functions using SVAR and local projections (LP) estimators.[10] The LP are used as a less biased, robust to model specification estimator in contrast to the biased but efficient SVAR. The results are plotted in Figure 2.2. Both approaches produce qualitatively similar results. This finding is in line with the findings of Plagborg-Moller and Wolf (2021b) who showed that both methods estimate similar impulse responses.

A disruption in oil supply causes a decline in world oil production. After two months,

[10] Following the related literature, in all time invariant models we use 24 lags.

oil production returns to pre-shock levels. This is the only major difference between our findings and the findings of Kilian (2009) where the recovery in oil production was partial. The shock does not have a significant effect neither on global economic activity nor in oil price. A demand shock has a positive and persistent effect both on the global economic activity and the price of oil. When the OLS-VAR is considered, the responses of real oil prices are statistically insignificant after four quarters. Oil-specific demand shock does not affect oil production. However, it causes an increase both in real economic activity and in real price of oil. In general, the findings are in line with the results of Kilian (2009).

The impulse responses from the TVP-VAR model are presented in Figures 2.3 and B1. Figure 2.3 plots the impulse responses for every point in the sample over a horizon of 24 periods (months). The model contains four lags, based on the SIC. The selection of lag order is consistent with similar studies in the literature (Apergis and Miller, 2009; Degiannakis et al., 2013; Kang et al., 2015). According to the results, a negative supply shock results in a sharp decline in global oil production. Similar to the time invariant model, oil production is restored after a couple of periods. Time evolution does not affect the impulse responses of global oil production. On the contrary, the responses of the global economic activity and the price of oil to a supply shock depends on the moment the shock occurs. A positive shock in demand has a positive effect on real economic activity which lasts for a few periods followed by a gradual decline. The response of oil production is a momentary decline. The effect is fully reversed and a few periods after the shock, we observe a co-movement of oil production and oil prices. The magnitude of the increase is greater after 2010, two years after the Global Financial Crisis of 2008. In addition, a demand shock has a positive effect in the price of oil. An oil-specific shock has a positive effect on both the real price and the global oil production. A stronger increase in oil production is observed after 2015. This contradicts the findings from the time invariant model. Furthermore, an oil-specific shock has positive effect in real economic activity which gradually declines. The sign of long-run responses of real economic activity changes over time.[11]

[11]Figure B1, in Appendix 2.6, presents the evolution of impulse responses over the sampling period at 3, 12, 24 months after the shock and allows us to observe how the different effects of a shock based on the moment it happens.

Figure 2.2: Impulse responses using time invariant OLS and LP estimators. The SVAR model is based on Kilian (2009).

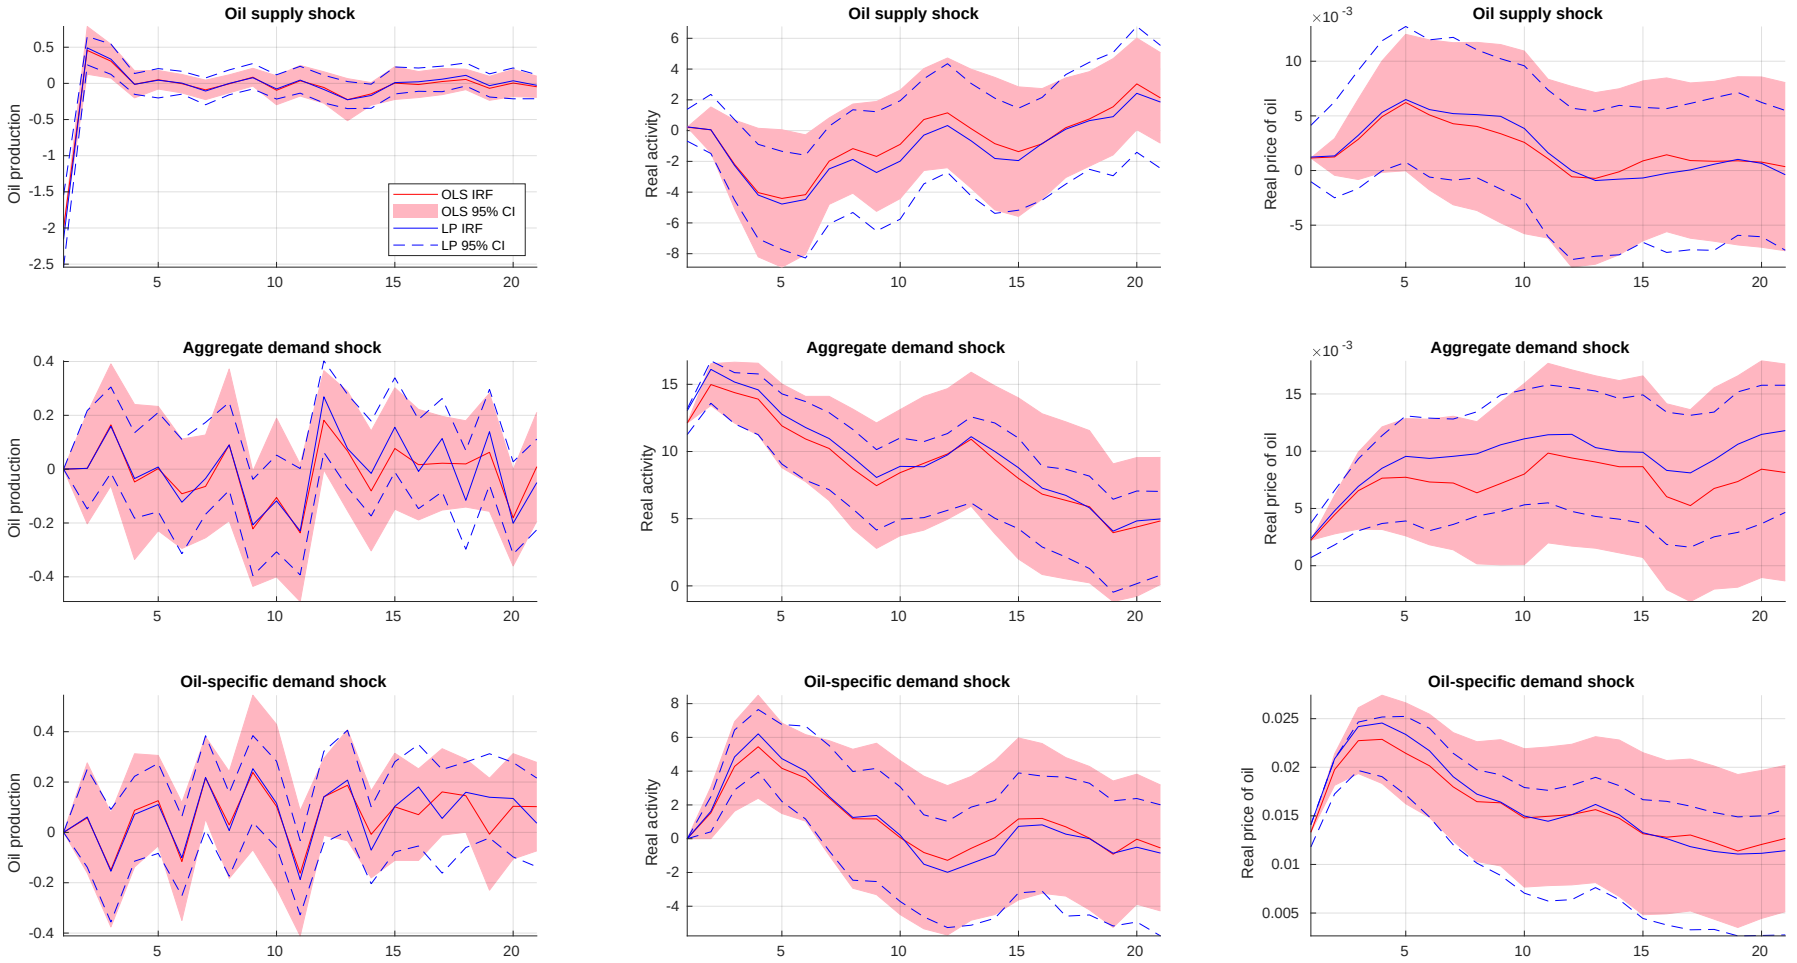

Figure 2.3: Time varying impulse responses. SVAR based on Kilian (2009).

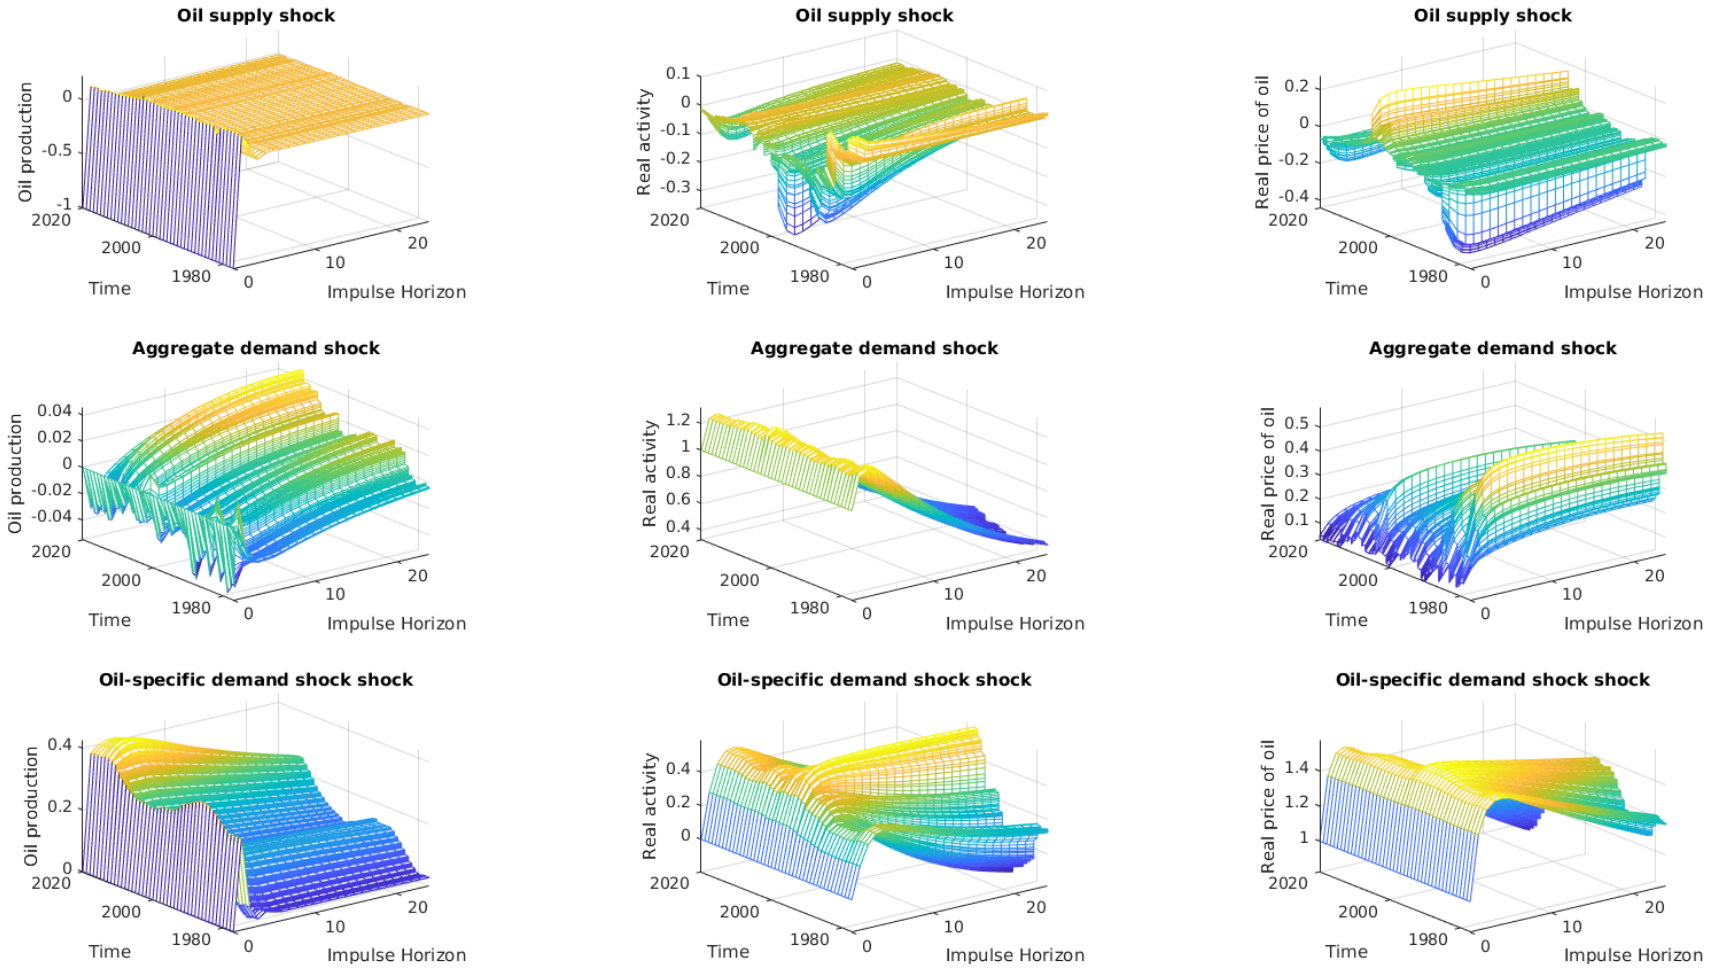

We augment the previous setup by including the stock market returns in the initial setup of Kilian and Park (2009). The aim of this exercise is to examine the relationship between stock prices and the price of oil. Similar to the trivariate model, we consider both time invariant and time-varying models with our sample starting from 1978M01.[12]

The impulse responses from the time invariant analysis are presented in Figure 2.4. The left column of Figure 2.4 plots the responses of stock market to oil supply and demand shocks. These findings are in line with the results of Kilian and Park (2009). First, we observe that a sudden disruption of in oil supply has no significant impact on the stock market returns. Second, an aggregate demand shock does not affect stock market returns on impact. However, after a period of six months after the shock, stock returns start increasing. This upward trend lasts until the end of the impulse horizon. According to the OLS estimates, the rise in stock market returns is not statistically significant (at the 95% significance level). However, LP estimates indicate that on the long-run (after 18 months), the increase in stock returns is substantial. Third, a shock in oil prices has an immediate positive impact on stock market returns which gradually increases for the a period of two years. The results from both estimators suggest that the response of stock market returns to the oil-specific demand shock is statistically significant over the whole impulse horizon.

The responses to a shock in the stock market returns are presented in the right column of Figure 2.4. The three examined variables (oil production, global real economic activity and oil prices) are not affected by a shock in the stock market. In all cases, the value of zero is included in the confidence interval (regardless of the estimator). In the case of an aggregate demand shock, stock returns drop in the long-run, but the effect is rather weak.

Figures 2.5 shows the impulses responses from the TVP-VAR model with four lags (a two dimension illustration of the 3-month, 12-month and 24-month impulse evolution is also presented in Figure B2, in Appendix 2.6).[13] The top rows of Figure 2.5 present the impulse responses of stock market to oil supply and demand shocks. We observe that these responses are strongly affected by the time the shock occurs. For example, while a sudden increase in the price of oil causes a decrease in stock returns before 2015, after that year, a shock in oil prices has a positive impact on the stock market returns. Furthermore, the shocks have a persistent impact on the stock market. The bottom rows of Figure 2.5 show the responses of the variables to a positive shock in the

[12]We report the findings regarding shocks only to and from stock market returns, since the results regarding the relationship of supply, demand and the price of oil are similar to the ones we obtained from the trivariate model.

[13]The responses of world oil production, global economic activity and oil price to supply and demand shocks do not differ from the responses in the model without stock returns and are not reported here, but are available from the authors upon request.

stock market returns. First, we observe a sharp rise in oil production which is adjusted to a pre-shock level after a few periods. Second, the shock has a long-run effect on the real economic activity and the real oil price. The sign of the magnitude of the effect varies, depending on the moment the shock takes place. For example, from 2007 to 2010, two years after the shock, economic activity has declined while after 2010 the shock has a positive effect on real economic activity. Finally, while oil prices increase shortly after the shock, the long-run response is also time dependent (negative over the period 2008-2014 and positive in the rest of the sampling period).

Figure 2.4: Impulse responses using time invariant OLS and LP estimators. The SVAR model is based on Kilian and Park (2009).

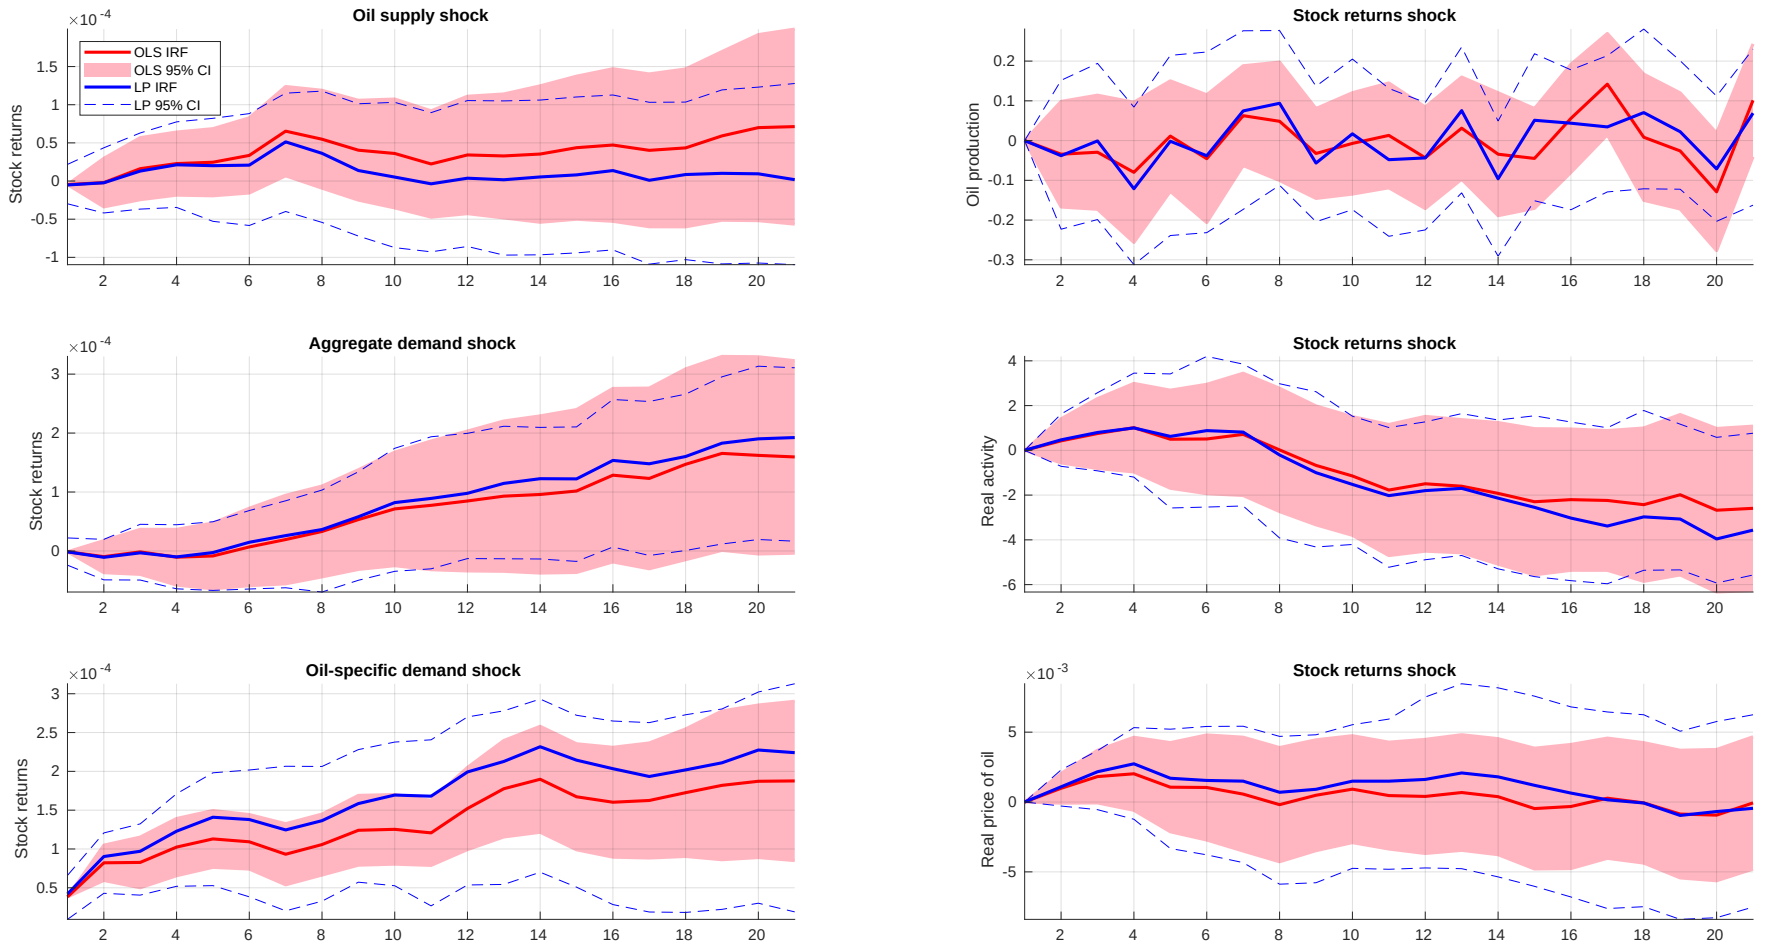

Figure 2.5: Time varying impulse responses to shocks to and from stock market returns. SVAR based on Kilian and Park (2009).

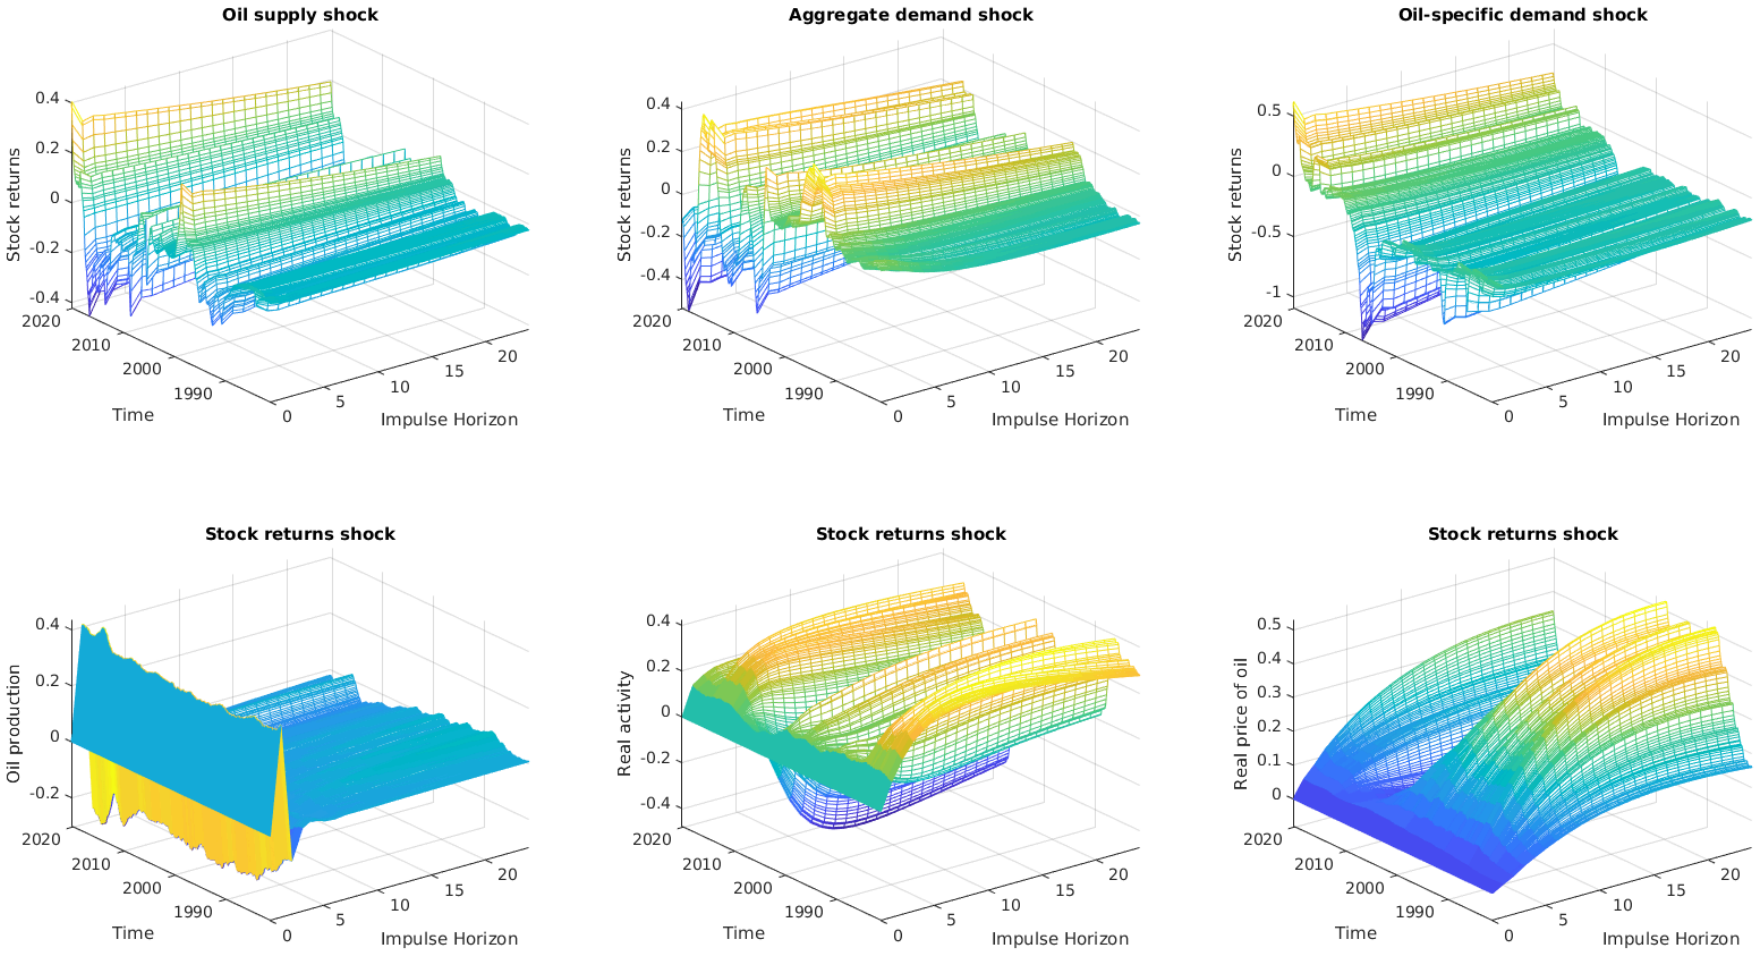

2.4.2 The effect of GSVI on the price of oil

Figure 2.6 shows the evolution of the estimated parameters over time (we consider two lags in the model based on the SIC).[14] The left column of Figure 2.6 presents the posterior means of the coefficients of the first lag of the GSVI variable (1 lag) in the five equations of the model: the oil production, oil demand, real price of oil, stock market returns and information demand equation. The shaded area in each graph reports the 68% credible interval. In the oil supply equation, the coefficients of lagged GSVI are negative, statistically significant and remain stable over the entire sample period. In the equation of oil demand, the coefficients of GSVI are positive and statistically insignificant for most of the sampling period. For a brief period, during 2016, the coefficient becomes statistically significant. There does seem to be a movement in the coefficients of the variable of information demand in the equation of real price of oil. The posterior coefficient increases over the first 5 years of the sample and then gradually declines until 2014. After 2014 it increases for almost two years and then remains relatively constant. In addition, the posterior mean of coefficients is negative until 2015 and it becomes positive after that period. We observe a similar pattern in the evolution over time of coefficients of GSVI in the stock market returns equation. However, the values of the coefficients are negative over the whole sample.

The effect of time is also evident on the standard deviation of the errors which justifies the selection of a model that accounts for stochastic volatility. The right column in Figure 2.6 presents the posterior means of the standard deviation of residuals in the five equations of the VAR model. In all cases, there is a steep increase in the standard deviation during 2009, caused by the Global Financial Crisis. Furthermore, there are cases in the sample where the standard deviations of the shocks increase for a short period. Finally, it is worth noticing the effect of the recent pandemic in the volatility of the error terms. Especially in the case of stock market returns where the standard deviation of the disturbance term is currently increasing and tends to exceed the maximum value which reached during 2009. The rate of increase is also higher than in 2009.

We now turn to the impulse response function analysis. The identification scheme builds on the work of Kilian (2009) and Kilian and Park (2009). We consider the following relationship between the reduced form errors u_t and structural errors ϵ_t:

[14] The selection of lag order is consistent with similar studies in the literature (Apergis and Miller, 2009; Degiannakis et al., 2013; Kang et al., 2015).

Figure 2.6: Posterior means of coefficients of lagged GSVI and posterior means of standard deviations of the residuals in the mean equations.

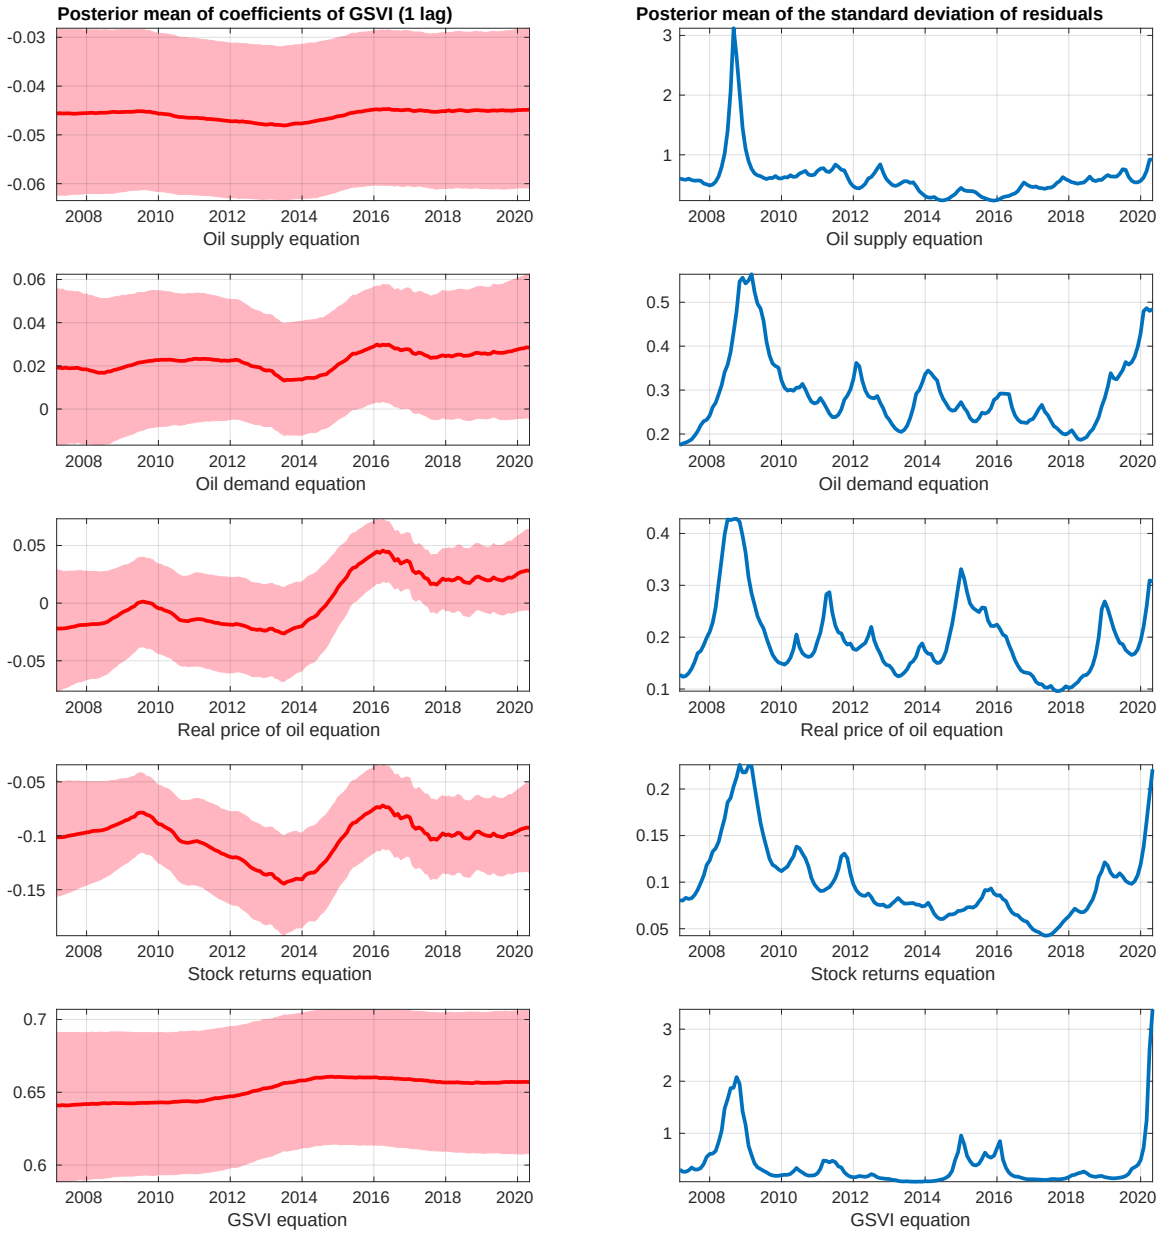

$$\begin{bmatrix} u_t^{\Delta prod} \\ u_t^{rea} \\ u_t^{rpo} \\ u_t^{ret} \\ u_t^{gsvi} \end{bmatrix} = \begin{bmatrix} \xi_{11} & 0 & 0 & 0 & 0 \\ \xi_{21} & \xi_{22} & 0 & 0 & 0 \\ \xi_{31} & \xi_{32} & \xi_{33} & 0 & 0 \\ \xi_{41} & \xi_{42} & \xi_{43} & \xi_{44} & 0 \\ \xi_{51} & \xi_{52} & \xi_{53} & \xi_{54} & \xi_{55} \end{bmatrix} \begin{bmatrix} \epsilon_t^{\Delta prod} \\ \epsilon_t^{rea} \\ \epsilon_t^{rpo} \\ \epsilon_t^{ret} \\ \epsilon_t^{gsvi} \end{bmatrix}$$

where we assume that stock market and information demand react contemporaneously to all supply and demand shocks and we treat shocks in oil prices as predetermined with the respect to the economy (Lee and Ni, 2002).

Our aim is to examine the bilateral relationship between GSVI and the price of oil. To this end, we consider the responses of the variables to shocks in the GSVI and the response of GSVI to shocks to the rest of the variables. The top row in Figure 2.7 presents the impulse responses of the variables to a shock in information demand. A positive shock in the GSVI variable yields a decrease in oil production which lasts only for two months. After that, oil production sharply increases and reaches pre-shock levels. The shock also causes a rise in aggregate demand which in turn yields a rise in the price of oil. Specifically, we observe a positive co-movement in aggregate demand and oil prices with the impact on the price of oil being greater than on aggregate demand. This result, combined with the delayed increase in oil production indicates that an unexpected increase in information demand acts as a mechanism for a shock in real economic activity, leading eventually in a positive co-movement of oil production and the price of oil. Finally, a sudden increase in information demand does not have an immediate effect on stock market. In addition, the effect is different, depending on the timing the shock in GSVI occurs: a shock in GSVI before 2015 causes stock market returns to start to fall four months after the shock. On the contrary, a similar shock after 2015 causes a rise in stock market returns after four months. In both cases, the effect of the shock is long-lasting and the upward/downward trends continues for more than two years.

The previous results suggest that unexpected changes in GSVI affect the behaviour of oil dynamics and stock market returns. However, oil demand and supply shocks may affect household disposable income available for other expenditures through energy prices. Since prior evidence suggests that oil price shocks are mainly transmitted through the demand side, we investigate the effect of structural oil supply and demand shocks on GSVI—a barometer of private households', retail investors and non-institutional traders perception of uncertainty and future economic conditions. The bottom row of Figure 2.7 shows the impact of different shocks on the GSVI index. All shocks affect investors' sentiment on impact, reflecting the rapid spread of news

worldwide. Furthermore, the sign of the response is affected by the time the shock occurs. For example, disruptions of the physical supply, which raise the real crude oil price, have a limited impact on the GSVI until 2010, but this effect is becoming stronger as we move to 2015 onward. The diminishing role for oil supply shocks suggests that households and retail investors expect that these shocks are short-lived, given that reduced production in one country is easily offset by other oil producers. On the demand side, both aggregate demand shocks associated with the global business cycle and oil-specific demand shocks significantly depress GSVI during the period prior to 2015. This impact turns to positive in subsequent years but is short-lived. In line with the results in Kilian (2009) and Guntner and Linsbauer (2018), positive aggregate demand shocks cause an increase of optimism among retail investors, followed by a significant reduction in the GSVI over the next 18 months. Other oil demand shocks such as an oil-specific demand shock, have a strong persistent positive effect on the GSVI over the recent years. The impact of a shock in stock returns on GSVI is positive and temporary over the last two years and during the global financial crisis period.

For illustrative purposes we also consider the impulses responses from a model with constant parameters. We estimate the impulse response functions using both an OLS-SVAR estimator and local projections (LP) by Jordá (2005). Following the related literature, in all time invariant models we use 24 lags.

The results are presented in Figure 2.8. In most cases, SVAR and LP produce qualitatively similar results.[15] LP produce narrower confidence intervals and as a result, suggest statistical significance of the responses more often compared to SVAR estimates.[16] The responses of aggregate demand and oil prices remain increasing, similar to the TVP-VAR model. The information demand shock leads to a decrease in stock market returns which is statistically significant only when we consider LP. Finally, the shock has no impact on crude oil production. The response of the GSVI to shocks in the rest of the variables is a temporary increase.

[15]This confirms the findings of Plagborg-Moller and Wolf (2021b)

[16]The LP are used as a less biased, robust to model specification estimator in contrast to the biased but efficient SVAR (Plagborg-Moller and Wolf, 2021a).

Figure 2.7: Impulse responses to an investors' sentiment shock (top) and impulses responses of GSVI to alternative shocks (bottom).

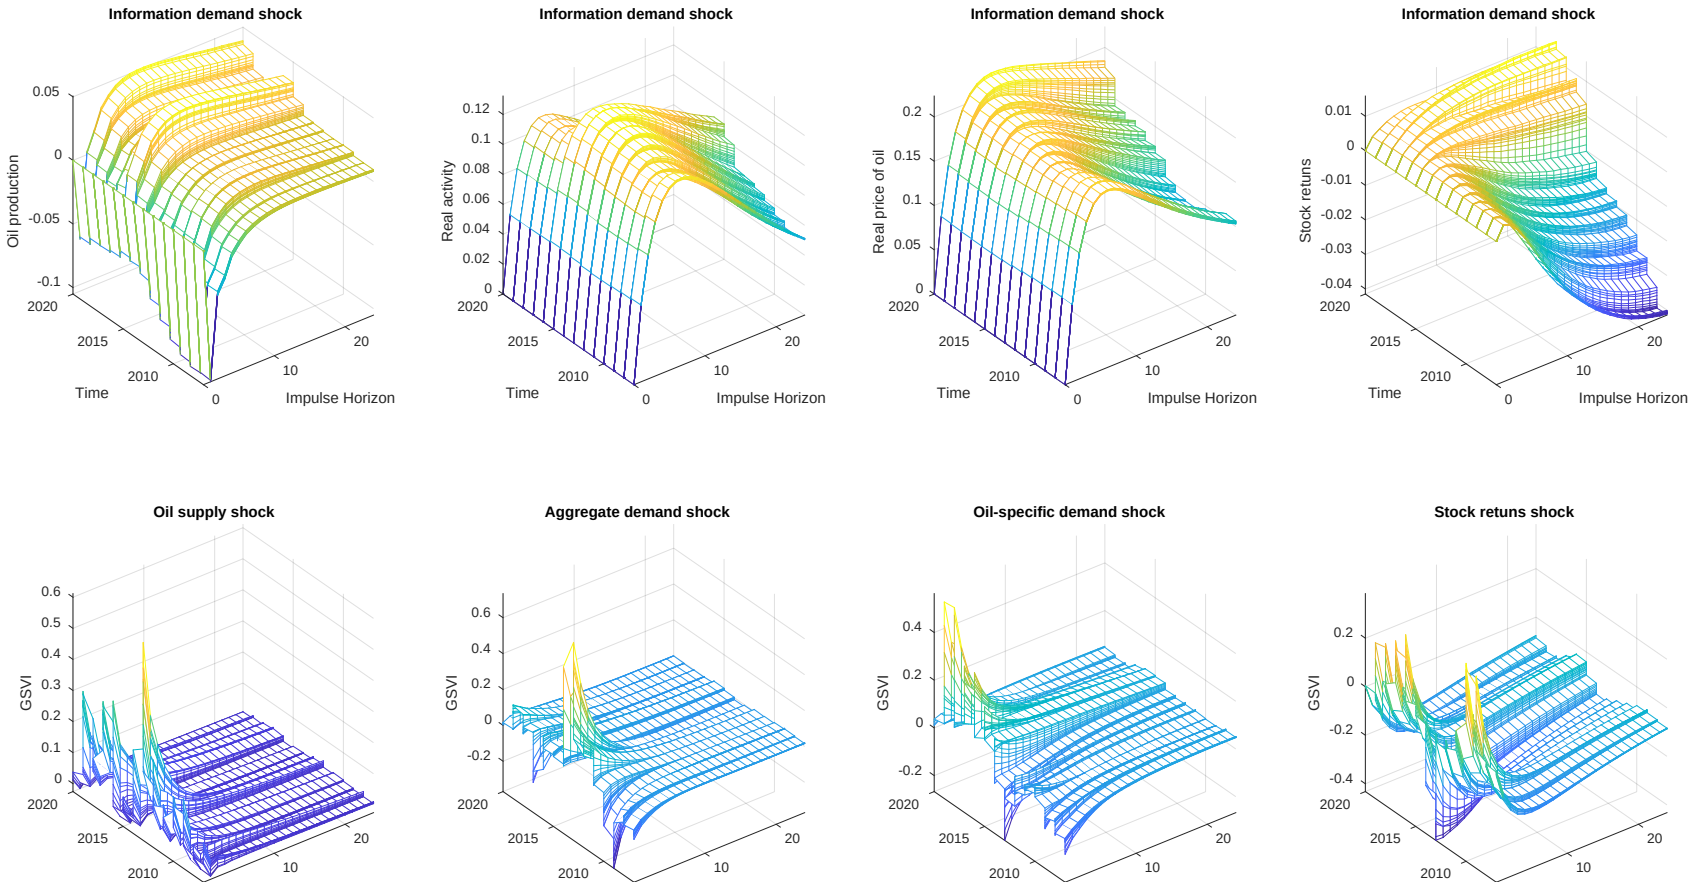

Figure 2.8: Impulse responses using time invariant SVAR and LP estimators. The first row plots the responses of the variables to an investors' sentiment shock. The second row plots the responses of GSVI to alternative shocks.

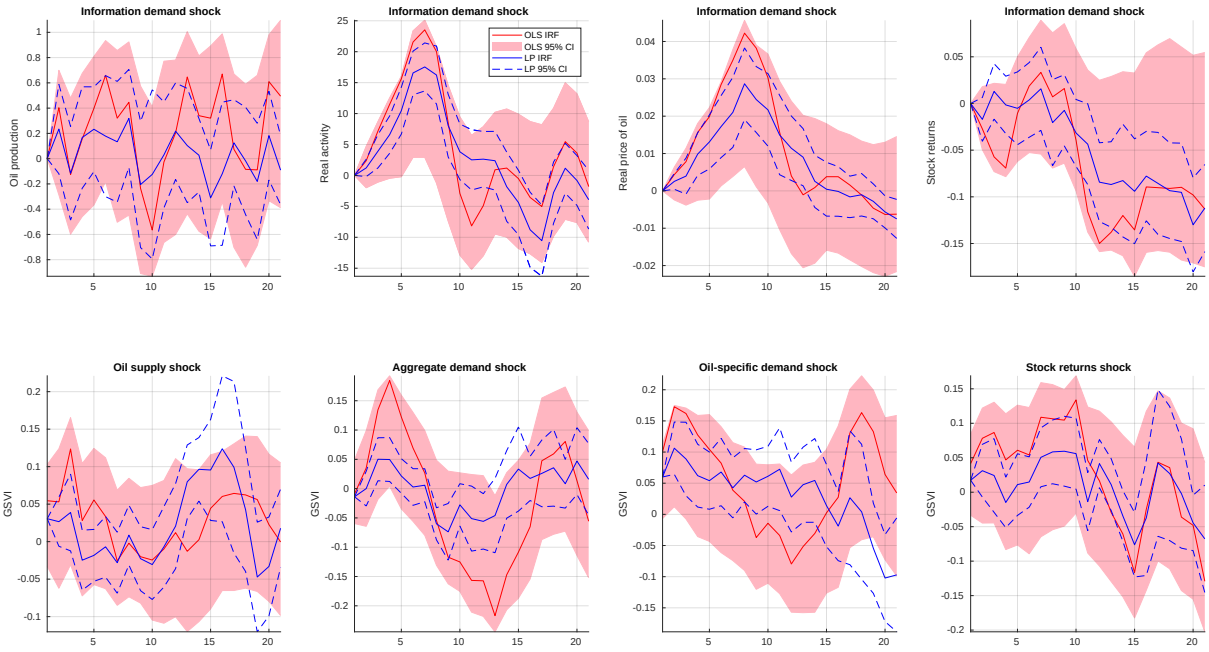

2.4.3 The impact of policy uncertainty on oil prices

In this section, we augment the baseline model by adding the EPU index. This model is based on the work of Kang and Ratti (2013) but differs in two respects. First, our model includes a measure of investors' sentiment and second, we consider a time-varying SVAR approach. By the same reasoning as in the previous model, the relationship between the reduced form and the structural errors is $u_t = \Xi\epsilon_t$ where $u_t^{'} = (u_t^{\Delta prod}, u_t^{rea}, u_t^{rpo}, u_t^{ret}, u_t^{epu}, u_t^{gsvi})$, $\epsilon_t^{'} = (\epsilon_t^{\Delta prod}, \epsilon_t^{rea}, \epsilon_t^{rpo}, \epsilon_t^{ret}, \epsilon_t^{epu}, \epsilon_t^{gsvi})$ and Ξ is 6×6 lower triangular matrix.

Figure 2.9 presents the impulse responses of the variables to a shock in the EPU index. The shock has an impact on oil production only in the short-run. Over the period 2007-2010 oil production increases after the shock while the after 2010, the shock causes a decrease. The shock has an negative impact on real economic activity if it occurs early in the sampling period. Similar response is observed for the oil prices. The response of stock market returns on the EPU index depends on time. From 2007 to 2008, an unanticipated increase in EPU affects stock market returns negatively in the short-run (the first ten periods). On the contrary, a shock after 2008 causes a slight

increase in stock returns. Finally, a shock in EPU has a marginal impact on GSVI which lasts only for a few periods.

Figure 2.10 plots the responses of the EPU index to a shock in the rest of the variables. An unexpected disruption in oil supply has no significant effect on the EPU index. On the contrary, a positive shock in aggregate demand leads to an increase in policy uncertainty. After the initial rise, EPU gradually decreases. However, over the recent years the effect of the shock persists until the end the impulse horizon. Shocks in oil prices and stock market returns have a similar effect on the EPU index. Both shocks yield a drop in policy uncertainty over the first two periods (months) which gradually returns to pre-shock levels. GSVI is affected from a shock in EPU only on impact. With the exception of a brief period from 2007 to 2008 when a shock causes an increase in information demand, a sudden rise in EPU has a weak negative effect on GSVI.

We now present the findings from a time invariant model. The top row of Figure 2.11 plots the impulse responses of the variables to a positive shock in EPU. Both the SVAR and the LP estimators produce similar results regarding the direction of each response, however, only the LP provide statistically significant estimates. An increase in EPU has a significant negative effect on up to fifteen months after the shock. Based on the SVAR estimates, the shock has a negative but insignificant impact on the price of oil. The shock also affects negatively stock market returns. When we use LP, we observe a substantial decrease in stock market returns that lasts up to thirteen months after the shock. Finally, the sentiment index decreases only for a couple of periods after the shock.

The second row of Figure 2.11 presents the responses of EPU. A positive aggregate demand shock has a negative effect on EPU that is statistically significant up to two months and from thirteenth to fifteenth month. Similarly, an oil-specific demand shock decreases EPU only on impact. A shock in stock market returns leads to an immediate decrease in EPU that lasts up to three periods after the shock. After that period, EPU rises to pre-shock levels for five months and then declines again for five more months. Finally, an increase in information demand does not affect EPU.

Figure 2.9: Impulse responses to an EPU shock.

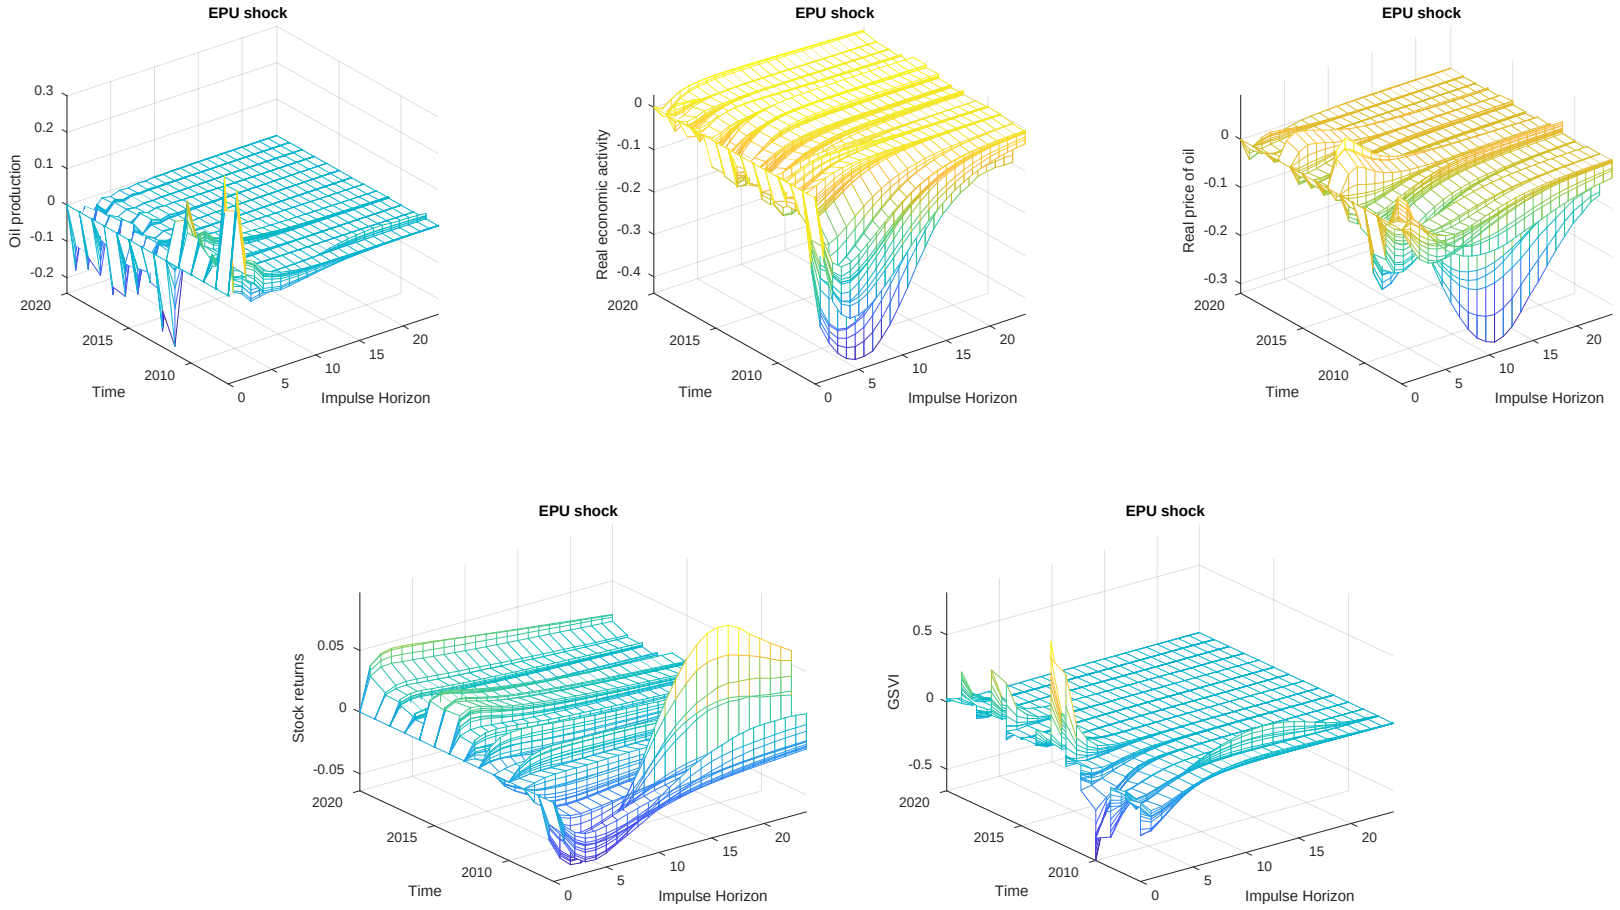

Figure 2.10: Impulse responses of the EPU index to alternative shocks.

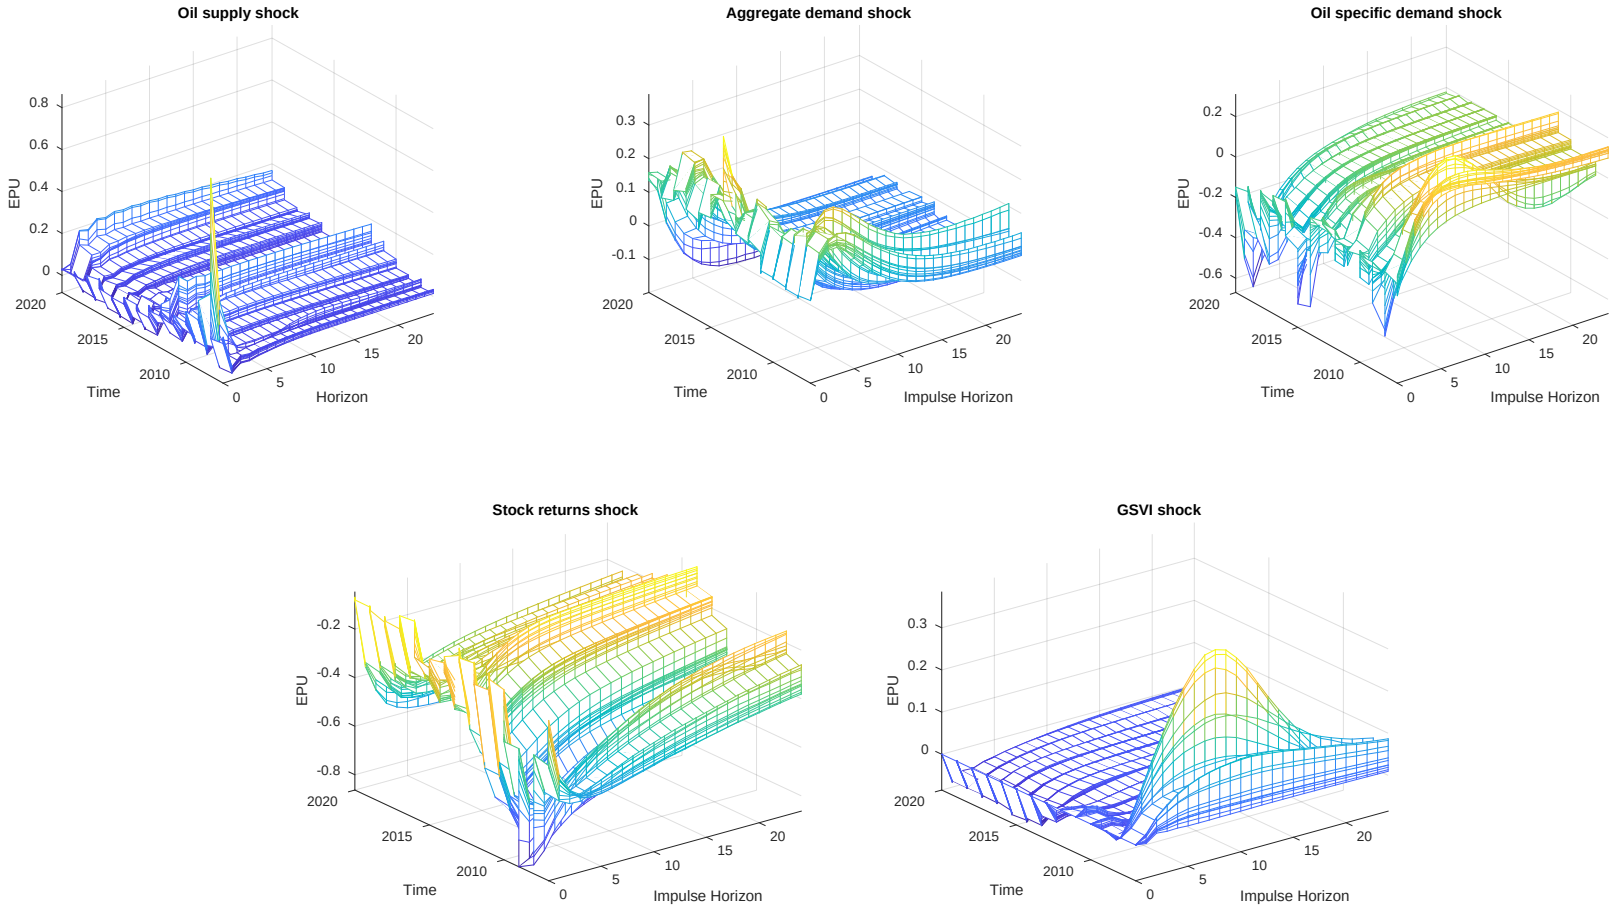

Figure 2.11: Impulse responses using time invariant SVAR and LP estimators. The first row plots the responses of the variables to an EPU shock. The second row plots the responses of EPU index to alternative shocks.

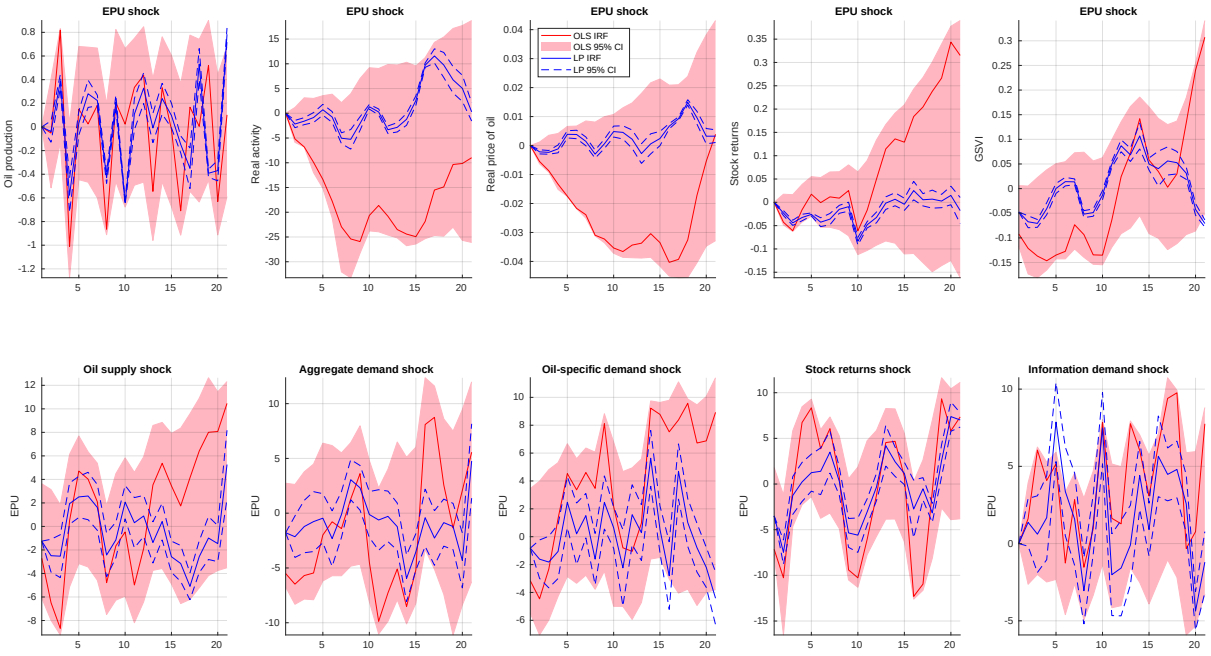

2.5 Robustness checks

2.5.1 Replacing GSVI with ICS

In this section we perform a robustness check regarding the model discussed in section 2.4.2 (the five-variable model). To examine the robustness of the results we use an alternative sentiment index which is often used in the literature, the Index of Consumer Sentiment (ICS) of Michigan's University (for applications, see Carroll et al., 1994; Matsusaka and Sbordone, 1995; Lagerborg et al., 2020). The index is available from 1978, however, for comparability purposes we estimate the model over the period 2004M01-2020M04.[17] Furthermore, we use the same number of lags, 2, as in the main model.

The first row in Figure 2.12 shows the impact of a positive shock in the ICS. Oil supply responds with an initial decrease which is reversed over the next month. After that initial response, oil production remains unresponsive. Real economic activity index,

[17]The results regarding the full sample are available from the authors upon request.

oil prices and stock returns are affected positively by the shock. The effect is more persistent in the case of oil prices and stock market returns. Overall, the responses of the first three variables are similar to the ones in the main model. However, the magnitude of the response is now smaller and independent of the time of the occurrence of the shock. Even in periods of known events (i.e. Global Financial Crisis), we observe no change in the response of the variables to alternative shocks. The latter lead us to conclude that the ICS can not capture the effect of time evolution. This explains the difference in the responses of stock market returns to shocks in GSVI and ICS over the first half of the sample.

The second row of Figure 2.12 presents the responses of ICS to shocks in the rest of the variables. Similar to GSVI, ICS is affected only for a couple of months from these shocks. While the response of ICS to each shock is similar over time, there is an increase in the magnitude of the response over the period 2018-2020 (oil production is an exception). Furthermore, while real economic activity increases after a shock from 2007 to 2018, after 2018 the shock causes a sharp decline in real economic activity.

Figure 2.12: Impulse responses using an alternative measure of investors' sentiment (GSVI is replaced by the ICS). The figure shows the impulse responses to shocks to and from the ICS.

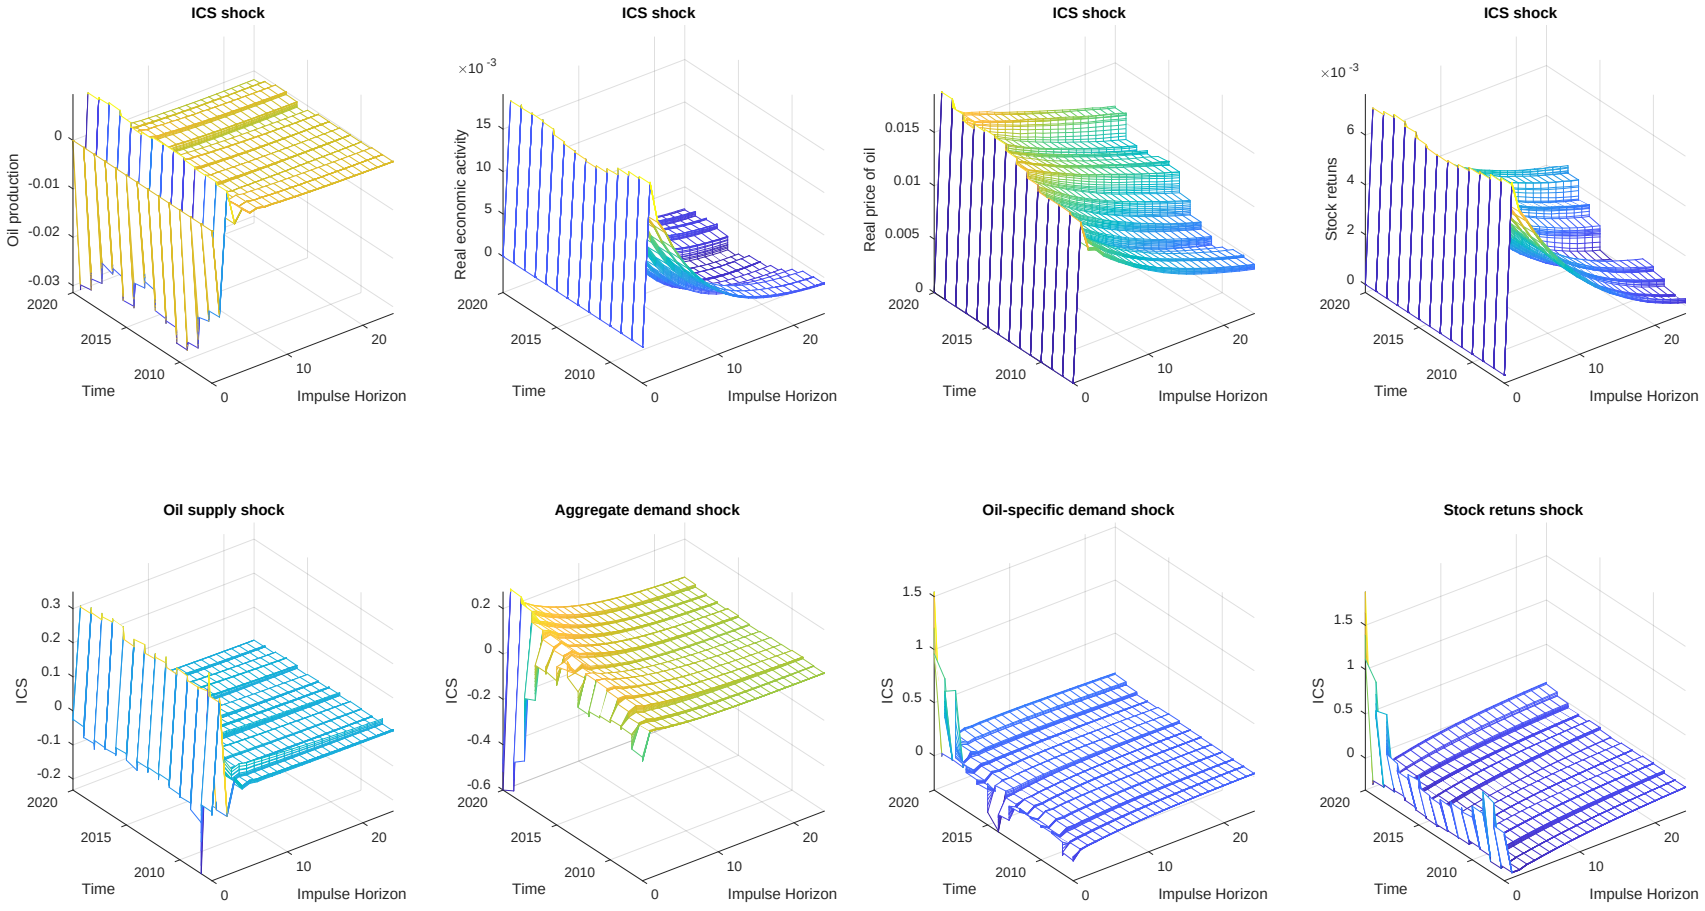

2.5.2 The role of oil inventories

The impact of oil inventories, through speculative trading, is well documented in the literature (see for instance Hamilton, 2009a,b; Kilian and Murphy, 2014; Baumeister and Hamilton, 2019; Känzig, 2020). However, as noted in Caldara et al. (2019) inventories could in principle quickly move to absorb differences between oil production and oil consumption, in turn affecting the dynamics of the oil market. An extension of our empirical analysis relates to the role of inventories, since our baseline model assumes that oil production is absorbed by consumption in every period. Thus, as an additional robustness check, we expand the five variable model by including a proxy for the oil inventories.

Our aim is to examine whether our findings are affected by the inclusion of oil inventories in the model. Doing so, we repeat the impulse response function analysis regarding shocks to and from the GSVI. These results are reported in Figure 2.13. In the first row, we plot the responses of the change in oil production, real economic activity, oil prices and stock market returns to a positive shock in information demand. The responses of the variables to a shock in the GSVI are similar to the responses in the main part of the analysis. In addition, the shock leads to an increase in oil demand which yields a positive shift in the price of oil. The increase both in oil demand and the price of oil is greater during the earlier period of the sample, something which we do not observe in the main results. The sign of the response of the stock market changes over time. With few exceptions, stock market returns decline if the shock occurs before 2015 and rise if the shock occurs after 2015. The second row of Figure 2.13 presents the impulse responses to supply and demand shocks. Similar to the main findings, the sign of the responses and the response of GSVI is affected mostly by the time of the shock.

Figure 2.13: Impulse responses to shocks from and to the GSVI. The model includes a proxy for oil inventories.

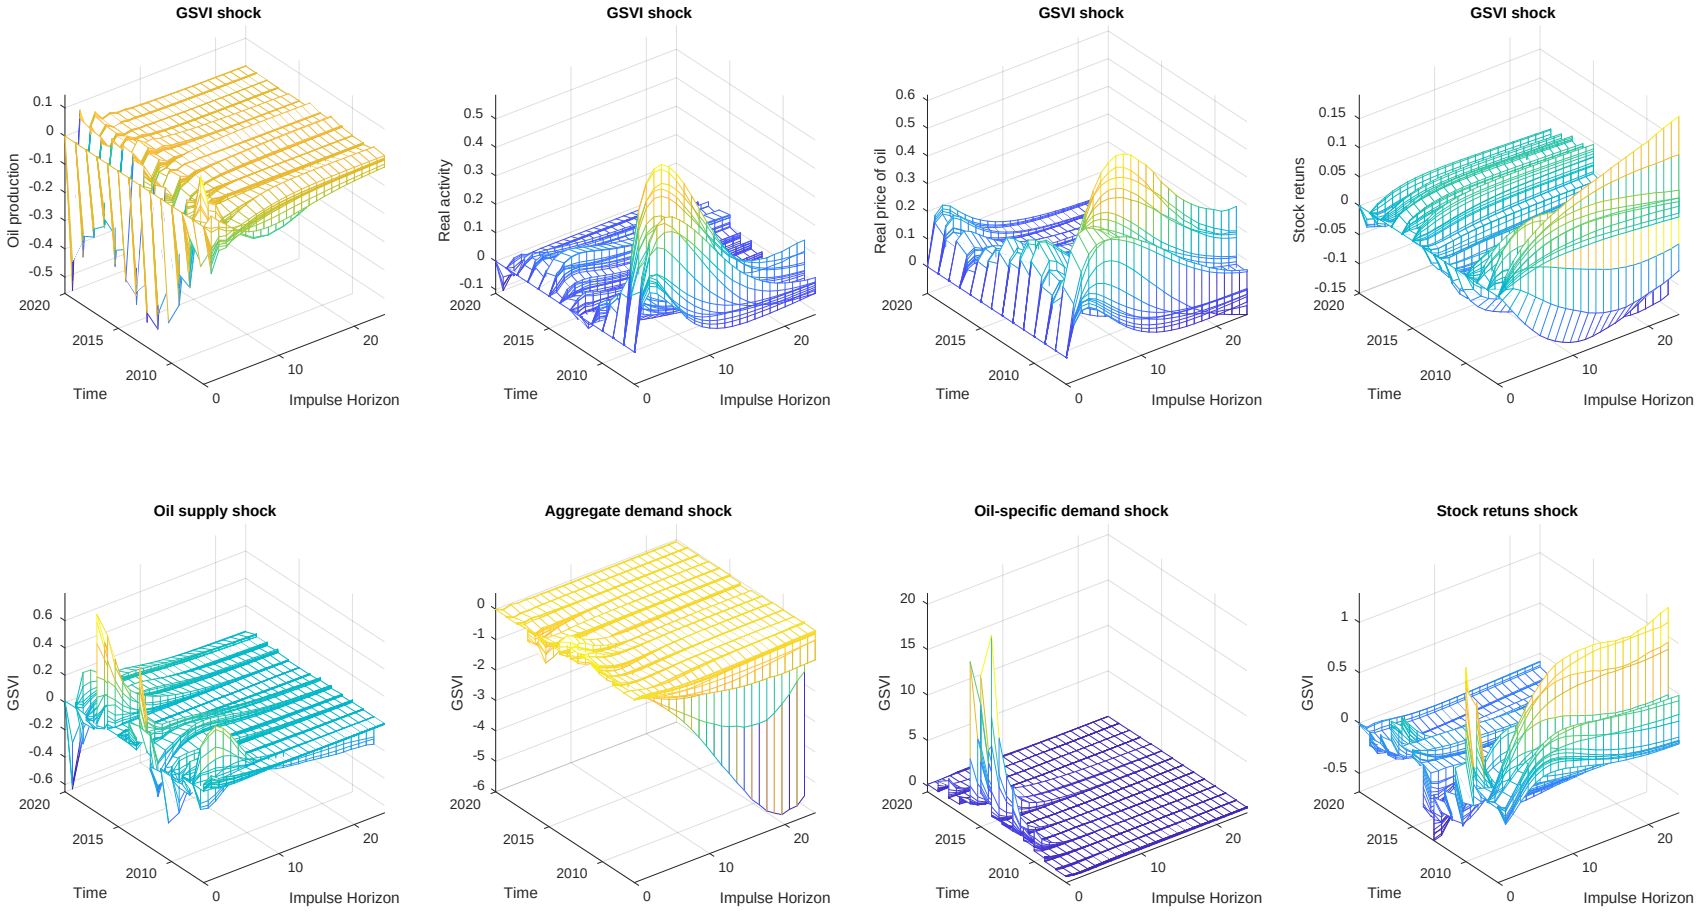

We subsequently consider the effect of oil inventories. In Figure 2.14 we plot the responses of all variables to a speculative shock. An unanticipated increase in oil inventory demand yields an increase in oil prices which results to a decrease in economic activity. Furthermore, a speculative shock leads to a rise in oil production which lasts only for one period. These findings are in line with the findings of Baumeister and Hamilton (2019). The speculative shock has a negative effect on stock market returns, a result also suggested by Ahmadi et al. (2016). The response of the GSVI is not statistically significant suggesting that a rise in inventory demand does not imply a rise in information demand. Figure 2.15 presents the impulse response of oil inventory to different shocks. Oil inventories do not respond to an increase in oil production but decline after a positive oil demand shock. This finding contradicts the results from Kilian and Murphy (2014) where the effect of a demand shock on impact increases oil inventories. However, this could be due to the sign restrictions, since in the next periods, the response of oil inventories becomes negative. An increase in prices reduces the oil stocks but the effect is rather weak. On the contrary, a positive shock in stock market has a strong impact on oil inventories. The effect is mostly negative, however for specific time periods, i.e. 2007-2009 and 2013-2014, the shock can lead to an increase in oil inventories. Finally, an unexpected increase in information demand leads to an decrease in oil inventories one month after the shock. The effect is reversed over the next few months. If the shock occurs from 2015 to 2017, the period after the sudden reduction of oil prices, the effect of an information demand shock on oil inventories is reversed.

Figure 2.14: Impulse responses of the variables to a speculative shock.

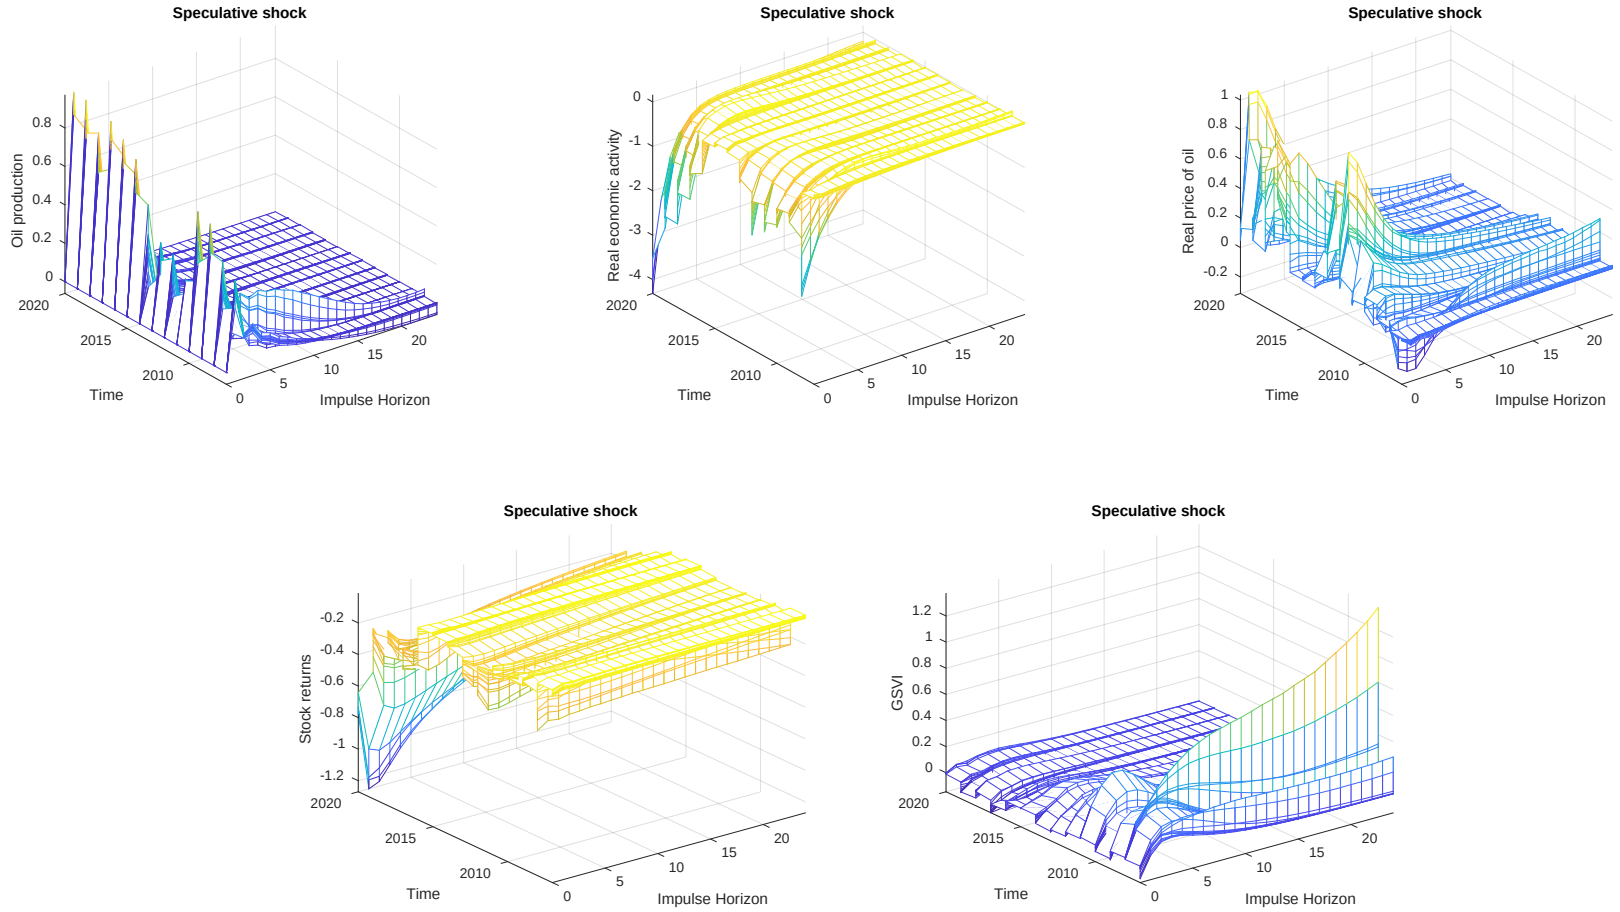

Figure 2.15: Impulse responses of oil inventories to alternative shocks.

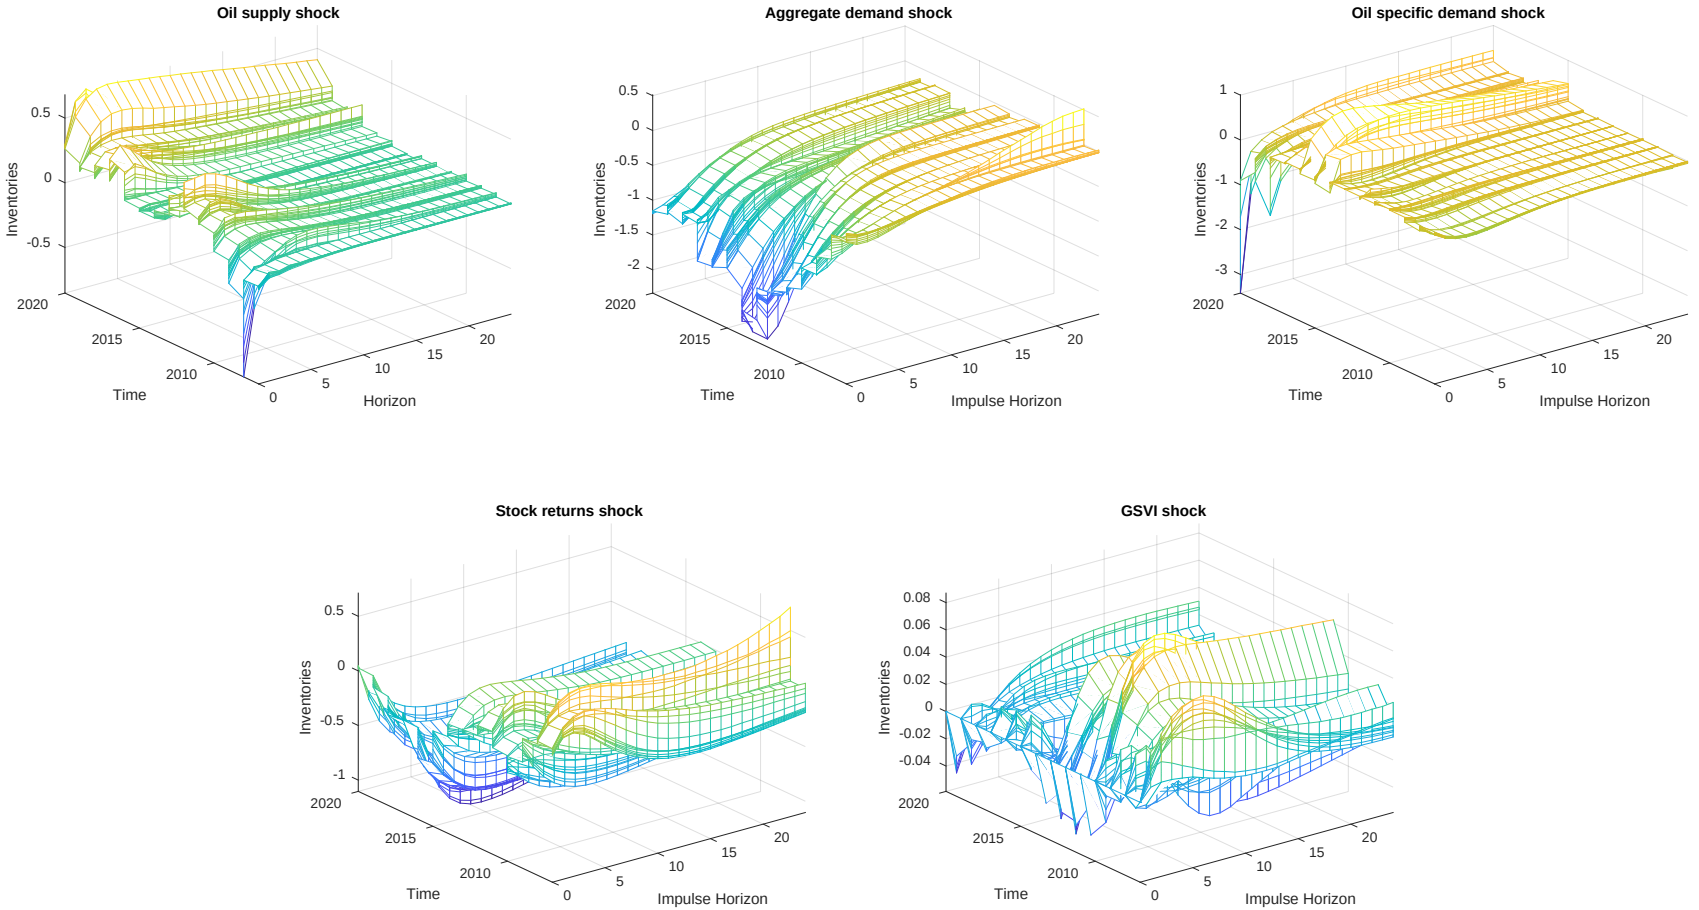

2.6 Conclusions

The empirical role of retail investors' sentiment shocks as a driver of oil price fluctuations remains debated in the literature, with findings hinging upon the identification assumptions used. This essay examines whether behavioural factors have an impact on oil prices. In order to get a direct measure of investor sentiment, we construct a Google search volume index (GSVI) by employing a dynamic factor analysis on 186 oil related search terms. The sentiment index is added to the model of Kilian and Park (2009) and the impulse response function analysis reveals that a sudden increase in information demand increases the price of oil through an increase in aggregate demand. Furthermore, the sign of the response of stock market returns to the shock depends on the timing of the shock. For example, a shock in the GSVI yields negative stock market returns before 2015 and positive stock market returns after 2015.

Motivated by previous evidence that oil price shocks are transmitted mainly through the demand side, we examine the effect of structural oil demand and supply shocks on GSVI —a barometer of households' and retail investors perception of uncertainty and current and future economic conditions. We show that the aggregate oil supply and demand shocks, oil-specific demand shocks and stock returns shocks have mostly a short-lived and significant impact on GSVI. The impact of aggregate oil supply and demand shocks is found similar in magnitude with the sign of the response to change over time only for the demand side.

Furthermore, economic policy uncertainty (EPU) plays also an important role in influencing financial and economic activities and may attract the attention of investors in the oil market, thus affecting the relationship between investor attention and oil market prices. By adding the EPU index in the model, we find that the effect of the EPU index on the variables is important but short-lasting. However, the EPU responds strongly to shocks in oil prices and stock market returns. In both examined models the results indicate that both the coefficients and transmission mechanisms of the shocks change over time. Finally, our findings are robust to an alternative measure of investor sentiment and after taking into account the role of oil inventories.

Bibliography

Ahmadi, M., M. Manera, and M. Sadeghzadeh (2016). Global oil market and the U.S. stock returns. *Energy 114*, 1277–1287.

Algaba, A., D. Ardia, K. Bluteau, S. Borms, and K. Boudt (2020). Econometrics meets sentiment: An overview of methodology and applications. *Journal of Economic Surveys 34*(3), 512–547.

Antonakakis, N., I. Chatziantoniou, and G. Filis (2014). Dynamic spillovers of oil price shocks and economic policy uncertainty. *Energy Economics 44*, 433–447.

Apergis, N. and S. Miller (2009). Do structural oil-market shocks affect stock prices? *Energy Economics 31*(4), 569–575.

Arampatzidis, I., T. Dergiades, R. Kaufmann, and T. Panagiotidis (2021). Oil and the U.S. stock market: Implications for low carbon policies. *Energy Economics 103*, 105588.

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

Bai, J. and S. Ng (2007). Determining the number of primitive shocks in factor models. *Journal of Business & Economic Statistics 25*(1), 52–60.

Baker, S., N. Bloom, and S. Davis (2016). Measuring economic policy uncertainty. *The Quarterly Journal of Economics 131*(4), 1593–1636.

Barberis, N., A. Shleifer, and R. Vishny (1998). A model of investor sentiment. *Journal of Financial Economics 49*(3), 307–343. Reprinted in Richard Thaler, ed., Advances in Behavioral Finance Vol. II, Princeton University Press and Russell Sage Foundation, 2005.

Baumeister, C. and J. Hamilton (2019, May). Structural interpretation of vector autoregressions with incomplete identification: Revisiting the role of oil supply and demand shocks. *American Economic Review 109*(5), 1873–1910.

Baumeister, C. and L. Kilian (2016, February). Forty years of oil price fluctuations: Why the price of oil may still surprise us. *Journal of Economic Perspectives 30*(1), 139–60.

Baumeister, C. and G. Peersman (2013). Time-varying effects of oil supply shocks on the us economy. *American Economic Journal: Macroeconomics 5*(4), 1–28.

Bird, R. and D. Yeung (2012). How do investors react under uncertainty? *Pacific-Basin Finance Journal 20*(2), 310–327.

Blanchard, O. J. and M. Riggi (2013). Why are the 2000s so different from the 1970s? a structural interpretation of changes in the macroeconomic effects of oil prices. *Journal of the European Economic Association 11*(5), 1032–1052.

Brandt, M. W. and L. Gao (2019). Macro fundamentals or geopolitical events? a textual analysis of news events for crude oil. *Journal of Empirical Finance 51*, 64–94.

Caldara, D., M. Cavallo, and M. Iacoviello (2019). Oil price elasticities and oil price fluctuations. *Journal of Monetary Economics 103*, 1–20.

Carroll, C. D., J. C. Fuhrer, and D. W. Wilcox (1994). Does consumer sentiment forecast household spending? if so, why? *The American Economic Review 84*(5), 1397–1408.

Carter, C. and R. Kohn (1994). On Gibbs sampling for state space models. *Biometrika 81*(3), 541–553.

Cogley, T. and T. Sargent (2002). Evolving post-world war ii u.s. inflation dynamics. In *NBER Macroeconomics Annual 2001, Volume 16*, pp. 331–388. National Bureau of Economic Research, Inc.

Cogley, T. and T. Sargent (2005). Drift and Volatilities: Monetary Policies and Outcomes in the Post WWII U.S. *Review of Economic Dynamics 8*(2), 262–302.

Da, Z., J. ENGELBERG, and P. GAO (2011). In search of attention. *The Journal of Finance 66*(5), 1461–1499.

Da, Z., J. Engelberg, and P. Gao (2015). The sum of all FEARS investor sentiment and asset prices. *Review of Financial Studies 28*(1), 1–32.

Degiannakis, S., G. Filis, and C. Floros (2013). Oil and stock returns: Evidence from european industrial sector indices in a time-varying environment. *Journal of International Financial Markets, Institutions and Money 26*(10), 175–191.

Del Negro, M. and G. Primiceri (2015). Time-varying structural vector autoregressions and monetary policy: A corrigendum. *The Review of Economic Studies 82*(4), 1342–1345.

Doz, C., D. Giannone, and L. Reichlin (2011). A two-step estimator for large approximate dynamic factor models based on Kalman filtering. *Journal of Econometrics 164*(1), 188–205.

Doz, C., D. Giannone, and L. Reichlin (2012). A quasi—maximum likelihood approach for large, approximate dynamic factor models. *The Review of Economics and Statistics 94*(4), 1014–1024.

Elder, J. and A. Serletis (2010). Oil price uncertainty. *Journal of Money, Credit and Banking 42*(6), 1137–1159.

Frühwirth-Schnatter, S. (1994). Data augmentation and dynamic linear models. *Journal of Time Series Analysis 15*(2), 183–202.

Galay, G. (2019). Are crude oil markets cointegrated? testing the co-movement of weekly crude oil spot prices. *Journal of Commodity Markets 16*(C), 100088.

Garz, M. (2014). Good news and bad news: evidence of media bias in unemployment reports. *Public Choice 161*(3), 499–515.

Guntner, J. H. F. and K. Linsbauer (2018, October). The Effects of Oil Supply and Demand Shocks on U.S. Consumer Sentiment. *Journal of Money, Credit and Banking 50*(7), 1617–1644.

Hamilton, J. (2009a). Causes and consequences of the oil shock of 2007-08. *Brookings Papers on Economic Activity 40*(1), 215–283.

Hamilton, J. (2009b). Understanding crude oil prices. *The Energy Journal 30*(2), 179–206.

Hamilton, J. (2013). Historical oil shocks. In R. Parker and R. Whaples (Eds.), *The Routledge Handbook of Major Events in Economic History*, pp. 239–265. New York: Routledge Taylor and Francis Group.

Hamilton, J. (2021). Measuring global economic activity. *Journal of Applied Econometrics 36*(3), 293–303.

Hong, H. and J. C. Stein (1999). A unified theory of underreaction, momentum trading, and overreaction in asset markets. *The Journal of Finance 54*(6), 2143–2184.

Jordá, O. (2005). Estimation and inference of impulse responses by local projections. *The American Economic Review 95*(1), 161–182.

Kang, W. and R. Ratti (2013). Oil shocks, policy uncertainty and stock market return. *Journal of International Financial Markets, Institutions and Money 26*(3), 305–318.

Kang, W., R. Ratti, and K. Yoon (2015). Time-varying effect of oil market shocks on the stock market. *Journal of Banking & Finance 61*, 150 – 163.

Kilian, L. (2008, December). The economic effects of energy price shocks. *Journal of Economic Literature 46*(4), 871–909.

Kilian, L. (2009). Not all oil price shocks are alike: Disentangling demand and supply shocks in the crude oil market. *American Economic Review 99*(3), 1053–1069.

Kilian, L. (2019). Measuring global real economic activity: Do recent critiques hold up to scrutiny? *Economics Letters 178*(3), 106–110.

Kilian, L. and D. P. Murphy (2014). The role of inventories and speculative trading in the global market for crude oil. *Journal of Applied Econometrics 29*(3), 454–478.

Kilian, L. and C. Park (2009). The impact of oil price shocks on the U.S. stock market*. *International Economic Review 50*(4), 1267–1287.

Kilian, L. and X. Zhou (2018). Modeling fluctuations in the global demand for commodities. *Journal of International Money and Finance 88*, 54–78.

Kim, S., N. Shephard, and S. Chib (1998). Stochastic volatility: Likelihood inference and comparison with arch models. *The Review of Economic Studies 65*(3), 361–393.

Koop, G. and D. Korobilis (2009). Bayesian multivariate time series methods for empirical macroeconomics. *Foundations and Trends in Econometrics 3*(4), 267–358.

Korobilis, D. (2017). Quantile regression forecasts of inflation under model uncertainty. *International Journal of Forecasting 33*(1), 11–20.

Kuck, K. and K. Schweikert (2017). A markov regime-switching model of crude oil market integration. *Journal of Commodity Markets 6*(C), 16–31.

Känzig, D. (2020). The macroeconomic effects of oil supply news: evidence from OPEC announcements. *American Economic Review 111*(4), 1092–1125.

Lagerborg, A., E. Pappa, and M. O. Ravn (2020). Sentimental business cycles. *Forthcoming in Review of Economic Studies 90*(3), 1358–1393.

Leduc, S. and K. Sill (2004). A quantitative analysis of oil-price shocks, systematic monetary policy, and economic downturns. *Journal of Monetary Economics 51*(4), 781–808.

Lee, K. and S. Ni (2002). On the dynamic effects of oil price shocks: a study using industry level data. *Journal of Monetary Economics 49*(4), 823–852.

Lippi, F. and A. Nobili (2012). Oil and the macroeconomy: a quantitative structural analysis. *Journal of the European Economic Association 10*(5), 1059–1083.

Matsusaka, J. G. and A. M. Sbordone (1995). Consumer confidence and economic fluctuations. *Economic Inquiry 33*(2), 296–318.

Plagborg-Moller, M. and C. Wolf (2021a). Local projection inference is simpler and more robust than you think. *Econometrica 89*(4), 1789–1823.

Plagborg-Moller, M. and C. Wolf (2021b). Local projections and VARs estimate the same impulse responses. *Econometrica 89*(2), 955–980.

Primiceri, G. (2005). Time varying structural vector autoregressions and monetary policy. *Review of Economic Studies 72*(3), 821–852.

Rahman, S. and A. Serletis (2011). The asymmetric effects of oil price shocks. *Macroeconomic Dynamics 15*(S3), 437–471.

Sargent, T. and C. Sims (1977). Business cycle modeling without pretending to have too much a priori economic theory. Working Papers 55, Federal Reserve Bank of Minneapolis.

Sockin, M. and W. Xiong (2015). Informational frictions and commodity markets. *The Journal of Finance 70*(5), 2063–2098.

Stock, J. and M. Watson (2016). Dynamic factor models, factor-augmented vector autoregressions, and structural vector autoregressions in macroeconomics. In J. Taylor and H. Uhlig (Eds.), *Handbook of Macroeconomics*, Volume 2, Chapter 8, pp. 415–525. Elsevier.

Vlastakis, N. and R. Markellos (2012). Information demand and stock market volatility. *Journal of Banking & Finance 36*(6), 1808–1821.

Vozlyublennaia, N. (2014). Investor attention, index performance, and return predictability. *Journal of Banking & Finance 41*(C), 17–35.

Appendix

Table B1: List of searched terms used for the construction of the GSVI.

a barrel of oil	current crude oil price(s)	oil barrel price(s)	oil symbol
barrel of crude oil	current crude oil	oil barrel	oil ticker symbol
barrel of crude	current crude	oil chart(s)	oil ticker
barrel of oil price(s)	current oil price(s)	oil close	oil today
barrel of oil	current oil	oil closing price	oil trading price
barrel oil price(s)	current price of a barrel of oil	oil commodities	oil trading
barrel oil	current price of crude oil	oil commodity	price barrel of oil
barrel price(s)	current price of crude	oil contract	price barrel oil
bloomberg crude	current price of oil	oil cost per barrel	price crude oil
bloomberg energy prices	daily oil price(s)	oil cost	price crude
bloomberg energy	energy fund(s)	oil crude	price for barrel of oil
bloomberg oil	energy index	oil current price	price of a barrel of oil
buy oil futures	energy mutual funds	oil demand	price of barrel of oil
closing oil price(s)	energy prices	oil future	price of barrel
closing price of oil	energy stocks	oil futures price(s)	price of crude oil per barrel
cost of a barrel of oil	future oil	oil futures	price of crude oil
cost of crude oil	futures price	oil graph	price of crude
cost of oil	gas barrel	oil index	price of oil barrel
crude future(s)	gas index	oil market(s)	price of oil
crude oil barrel price	gas per barrel	oil options	price oil barrel
crude oil barrel	gas price per barrel	oil per barrel	price oil
crude oil chart(s)	global oil	oil price barrel	price per barrel of oil
crude oil close	historic oil prices	oil price chart(s)	price per barrel today
crude oil future(s)	historical oil price(s)	oil price graph	price per barrel
crude oil index	history of oil prices	oil price history	real time oil
crude oil market	investing in oil	oil price index	spot crude
crude oil per barrel	latest oil price(s)	oil price per barrel	spot oil price(s)
crude oil price(s)	light crude oil	oil price quote	spot oil
crude oil pricing	light crude	oil price ticker	stock market oil
crude oil quote	light sweet crude oil	oil price today	sweet crude oil
crude oil spot	light sweet crude price	oil price trend	sweet crude price(s)
crude oil stock symbol	light sweet crude	oil price(s)	sweet crude
crude oil stock	live oil price	oil prices graph	symbol for crude oil
crude oil symbol	nymex crude future	oil pricing	the price of oil
crude oil ticker symbol	nymex crude oil	oil producers	time oil
crude oil ticker	nymex crude	oil quote(s)	today oil price
crude oil today	nymex oil	oil spot price	trade oil
crude oil	oil a barrel	oil spot	trading oil
crude price(s)	oil and gas prices	oil stock symbol	world oil prices
crude stock	oil barrel cost	oil stocks	
crude	oil barrel price today	oil supply	

Notes: All SVI data are downloaded from Google Trends.

Table B2: Summary statistics of the variables used in the analysis.

Variable / Statistic	Mean	Maximum	Minimum	St. dev.	Skewness	Kurtosis	Jarque-Bera stat.	ADF stat.
Oil production	7356	12859	3973	2425	0.721	2.265	21.36***	-2.283
REA	13.02	191.0	-160.0	78.47	0.360	2.289	8.355**	-3.948**
Oil price	67.04	127.7	16.74	24.38	0.395	2.105	11.65***	-2.923
Stock market	1712	3230	735.0	624.7	0.730	2.375	20.76***	-1.893
GSVI	-0.002	0.463	-3.164	0.548	-2.846	12.38	1066.51***	-2.207
EPU	119.6	283.1	57.20	40.42	0.955	4.250	42.53***	-1.714
ICS	83.98	103.8	55.30	12.18	-0.526	2.224	13.96***	-2.205

Notes: i) ***, ** and * denote rejection of the null hypothesis at the 1%, 5% and 10% significance level. ii) In the implementation of the ADF test we assume only constant in the test equation and for the selection of the lag-length we use the SIC.

Figure B1: Evolution of the 3-month, 12-month and 24-month impulse responses over time. SVAR model based on Kilian (2009).

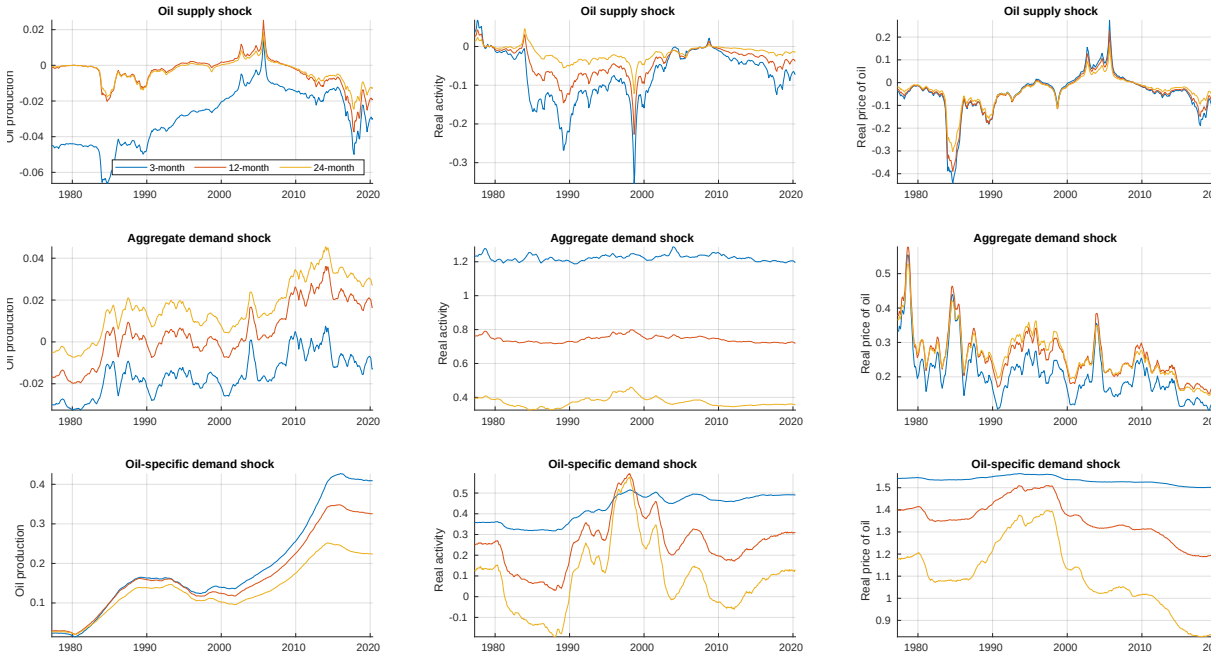

Figure B2: Evolution of the 3-month, 12-month and 24-month impulse responses over time. SVAR based on Kilian and Park (2009). Responses to shocks to and from stock market returns.

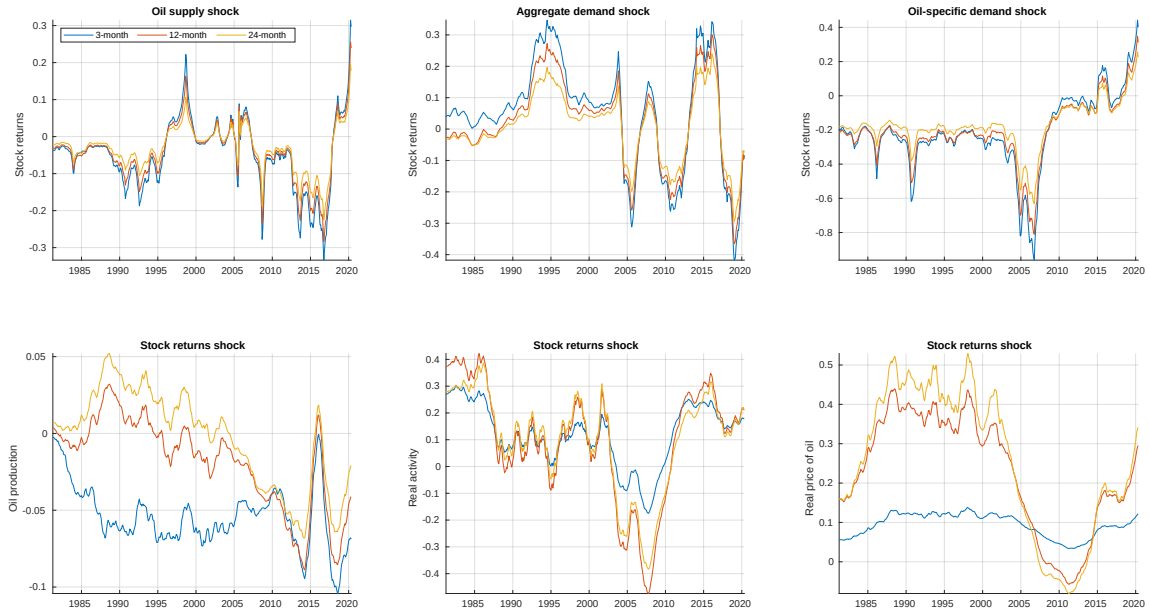

Chapter 3

On the volatility of cryptocurrencies

Abstract

We perform a large-scale analysis to evaluate the performance of traditional and Markov-switching GARCH models for the volatility of 292 cryptocurrencies. For each cryptocurrency, we estimate a total of 27 alternative GARCH specifications. We consider models that allow up to three different regimes. First, the models are compared in terms of goodness-of-fit using the Deviance Information Criterion and the Bayesian Predictive Information Criterion. Next, we evaluate the ability of the models in forecasting one-day ahead conditional volatility and Value-at-Risk. The results indicate that for a wide range of cryptocurrencies, time-varying models outperform traditional ones.

3.1 Introduction

The introduction of Bitcoin (Nakamoto (2009)) spurred the creation of new cryptocurrencies, with their current number exceeding 5000. Digital currencies facilitate electronic payments without the need of a bank (or other third party) intermediation.[1] According to Yermack (2013), Bitcoin cannot function as money, due to its nearly fixed supply, as Bitcoin mining is an energy consuming activity. However, there are similarities with exhaustible commodity resources (see the discussion in Gronwald (2019)). Furthermore, Bitcoin and digital currencies have additional features, such as hedging, diversifying and safe haven capabilities that make them appealing as an asset, (see for instance Dyhrberg (2016a,b); Bouri et al. (2017)).

Bitcoin like many financial assets exhibit volatility clustering and structural breaks in their volatility dynamics. Ignoring these features could negatively impact on the precision of volatility forecasts; see Lamoureux and Lastrapes (1990) and Bauwens

[1] Bitcoin also offers anonymity and low transactions fees, Panagiotidis et al. (2018).

et al. (2014). Since cryptocurrencies share similarities with other assets, one could expect that the volatility dynamics of cryptocurrencies also suffer from structural breaks. Recent studies examine the time-varying behaviour of Bitcoin returns and their volatility. Evidence from these reveals the presence of regime changes in the volatility of Bitcoin returns.[2] For example, Ardia et al. (2019) argue that Markov-switching GARCH (MSGARCH) models outperform single-regime specifications at predicting the one-day ahead Value at Risk (VaR).

The majority of the existing literature focuses on the analysis of Bitcoin and a limited number of major cryptocurrencies. This is, to some degree, justifiable since Bitcoin is the most traded digital currency in terms of market capitalisation. There are, however, other digital currencies that play an important role as speculative assets. The differences between cryptocurrencies in terms of both market capitalisation and investors' attention suggests that different modelling approaches should be employed in each case. In this essay, we perform an in-depth analysis of the volatility dynamics across a whole tranche of cryptocurrencies (292 in total).

We examine the presence of regime changes in the volatility of various cryptocurrencies using MSGARCH models. We consider a total of 292 cryptocurrencies including Bitcoin, Ethereum, Ripple and Tether which dominate the market of cryptocurrencies and capture the main bulge of the total market capitalisation.[3] We employ three traditional GARCH-type models, the GARCH, TGARCH (threshold GARCH) and EGARCH (exponential GARCH). For each model we consider three cases depending on the number of regimes (one, two and three regimes) and three cases depending on the specification of the conditional distribution (the Normal distribution, the Generalised Error Distribution (GED) and the Student's-t distribution). In total, for each currency, we estimate 27 alternative specifications of MSGARCH-type models. First, we compare the models in terms of goodness-of-fit using the Deviance Information Criterion (DIC) and the Bayesian Predictive Information Criterion (BPIC). Next, we asses the out-of-sample one-step-ahead volatility and density risk forecasting performance of the models. We evaluate the models performance using the mean squared (MSE), the mean absolute (MAE) errors (for volatility forecasts), the conditional coverage test (CC) of Christoffersen (1998) and the dynamic quantile test (DQ) of Engle and Manganelli (2004) (for risk forecasts). The models are estimated using Bayesian Markov chain Monte Carlo (MCMC) procedures. We choose the Bayesian approach since assessment of financial risk requires employing state-of-the-art econometric methodologies. We aim to compare the performance of MSGARCH models to their single-regime counter-

[2]The next section offers a review of the related literature.

[3]Throughout the essay the term 'major cryptocurrency' is used in terms of market capitalisation, at the time the study was conducted.

parts and to identify patterns for the cryptocurrencies based on the models that best describe them.

Our work expands and the work of Bouoiyour and Selmi (2015), Bouoiyour and Selmi (2015), Katsiampa (2017), Bouri et al. (2017) and Ardia et al. (2019) who assess the performance of GARCH models using information criteria. Furthermore, we complement the work of Stavroyiannis (2018) and Ardia et al. (2019) who examine the Value-at-Risk (VaR) of Bitcoin. Our contribution is twofold. First, we do not focus only on major cryptocurrencies but examine alternative digital currencies with different characteristics. Second, we employ MSGARCH models and we are able to detect structural changes in the volatility of cryptocurrencies without any a-priori specification on the time of the change.

It emerges from our analysis is that there is no one-model-fits-all solution in terms of modelling the volatility of cryptocurrencies; different GARCH-type models are found to be more suitable for different cryptocurrencies. More specifically, our analysis provide three main findings. First, MSGARCH models provide better results compared to their traditional counterparts. In-sample analysis suggests that MSGARCH models provide better goodness-of-fit results in more than 50% of the examined cryptocurrencies and out-of-sample analysis indicates that MSGARCH models provide better forecasts more than 60% (in some cases more than 70%) of times. Second, EGARCH models are selected less times by the employed criteria compared to the0 GARCH and TGARCH models. More formally, for a given number of regimes and a given conditional distribution, we obtain worse results (both in- and out-of-sample) from the EGARCH model compared to the non-exponential models. Finally, the analysis of two asymmetric models (the TGARCH and the EGARCH) indicates the presence of inverse leverage effect in most of the examined cryptocurrencies. That is, positive past returns affect the volatility of cryptocurrencies more than negative past returns.

The rest of the essay proceeds as follows: Section 2 discusses the related literature, the third one describes the econometric methodology, section 4 presents the main findings and the last one concludes.

3.2 Literature review

Despite the large number of active cryptocurrencies, the existing literature focuses on Bitcoin and some highly traded cryptocurrencies. One of the few exceptions is the work of Chu et al. (2017) where the seven most popular cryptocurrencies are considered. For each cryptocurrency, twelve GARCH models were fitted and compared based on five criteria.

Polasik et al. (2015) divide the Bitcoin analysis in four categories, one of which focuses on economic and financial issues from both theoretical and empirical perspective.[4] From a financial perspective, Urquhart (2016) studies the market efficiency of Bitcoin and finds that the results were affected by the sampling period. When examining the whole sample, the results indicate that the Bitcoin market is not weakly efficient. When the sample was split in two periods, the results reveal that informational inefficiency was caused by the first subsample. Bariviera (2017), Bouri et al. (2017) and Nadarajah and Chu (2017) also find Bitcoin to be in line with the efficient market hypothesis.

Bitcoin is considered an asset which can be used for speculative purposes, see for instance Baek and Elbeck (2015); Kristoufek (2014); Dyhrberg (2016b); Blau (2017), and Corbet et al. (2018). This can lead to extreme volatility and bubbles (see Fry and Cheah (2016)). These characteristics are closely related to speculation (Ahamed (2009), Reinhart and Rogoff (2009)). Williams (2014) shows that the volatility of Bitcoin price is seven times greater than that of gold. In addition, Baur et al. (2018) argue that returns, volatility and correlation characteristics of Bitcoin are distinctively different compared to gold and the US dollar.[5]

A large part of the related literature examines the volatility dynamics of Bitcoin. GARCH-type specifications provide the best in-sample performance according to Chu et al. (2017).[6] Glaser et al. (2014) employ the GARCH(1,1) model. Gronwald (2014) uses an autoregressive jump-intensity GARCH model and finds that Bitcoin prices are particularly marked by extreme price movements. Bouoiyour and Selmi (2015) examine the goodness-of-fit of GARCH model over two different periods, 2010 to 2015 and the first half of 2015. During the first interval, the threshold-GARCH estimates reveal the long duration of persistence. For the second period, the EGARCH is selected, displaying less volatility persistence. Klein et al. (2018) also focus on volatility persistence using the Fractionally Integrated APARCH (FIGARCH) to model volatility of Bitcoin. A number of papers examine the leverage effect on Bitcoin volatility. Bouri et al. (2017) find a negative relationship between the past shocks and volatility of Bitcoin before the first bubble burst in 2013 and no significant relation after. Katsiampa (2017) evaluates the performance of six different GARCH models using information criteria and selects the Asymmetric Component GARCH model as the most appropriate. Stavroyiannis (2018) implements a GJR-GARCH model to examine the VaR and related measures for the Bitcoin. The findings suggest that Bitcoin is a highly volatile currency which vio-

[4]For a detailed review on the Bitcoin literature see Panagiotidis et al. (2019).

[5]This essay is a replication and extension of Dyhrberg (2016a) who suggest that, in terms of volatility, the Bitcoin shares similarities with both gold and the dollar.

[6]HAR models are also used in Urquhart (2017).

lates the VaR measures more than the other assets (i.e. gold). Finally, Ardia et al. (2019) use MSGARCH models to investigate the presence of regime changes in the volatility dynamics of Bitcoin. They find that that asymmetric MSGARCH models perform better than their single-regime counterparts.

The work of Ardia et al. (2019) is not the only one to employ regime-switching models to model volatility of digital currencies. Caporale and Zekokh (2019) fit more than 1000 MSGARCH models to four cryptocurrencies and forecast one-day ahead VaR and expected shortfall. Mensi et al. (2019) support the existence of dual long memory in returns and volatility of Bitcoin and Ethereum. Cheikh et al. (2020) implement a Smooth Transition GARCH model to study the volatility dynamics of four major cryptocurrencies. Their findings support the existence of inverse leverage effect and safe haven hypothesis. Ma et al. (2020) propose a Markov regime-switching mixed-data sampling which improves the prediction accuracy of the realised variance of Bitcoin. Finally, Maciel (2021) argue that MSGARCH models provide better expected shortfall and VaR forecasts compared to their single-regime counterparts.

At the same time, Bayesian approaches have been employed in financial risk modelling. Geweke and Amisano (2010) compare Bayesian predictive distributions from five alternative models (including an ARCH model). Bauwens et al. (2010) proposed a MSGARCH model, estimated using a Bayesian MCMC algorithm. Bauwens et al. (2014) developed an estimation and forecasting method, based on a differential evolution MCMC method, for inference in GARCH models that allow for an unknown number of structural breaks at unknown dates. Balcombe and Fraser (2017) employ Bayesian Markov switching models to examine nine bubble containing series (including Bitcoin returns). Nearly all series appear to display strong regime switching. Ardia et al. (2017) compare the risk forecast performance of several GARCH-type models estimated via Maximum Likelihood (ML) and Bayesian techniques and find that Bayesian predictive densities improve the VaR backtest. Finally, Thies and Molnár (2018) utilise a Bayesian change point model to study structural breaks in average returns and volatility of Bitcoin. They find that regimes with higher volatility have higher average returns, however, the most volatile regime is the only regime with negative average returns.

3.3 Methodology

3.3.1 Estimation

Let y_t denote the returns of a cryptocurrency at day t. We assume $\mathbf{E}[y_t] = 0$ and $\mathbf{E}[y_t y_{t-l}] = 0$, that is the series has zero mean and are serially uncorrelated. We en-

sure that these assumptions are met by filtering the series with an AR(1) model. We follow Ardia et al. (2018) and express the general MSGARCH specification as:

$$y_t|(s_t = k, I_{t-1}) \sim D(0, h_{k,t}, \xi_k), \tag{3.1}$$

where $D(0, h_{k,t}, \xi_k)$ denotes a continuous distribution with zero mean, time-varying conditional variance $h_{k,t}$ and additional shape parameters gathered in the vector ξ_k and I_{t-i} is the information set observed up to time t. The stochastic variable s_t defined on the discrete space $\{1, \ldots, K\}$ evolves according to an unobserved first-order ergodic homogeneous Markov chain with $K \times K$ transition probability matrix $\mathbf{P} = (p_{i,j})$, where $p_{i,j}$ denotes the probability of a transition from state $s_{t-1} = i$ to state $s_t = j$. In the empirical analysis we allow up to $K = 3$ regimes. The standardised innovations are defined as $\eta_{k,t} = y_t / h_{k,t}^{1/2} \sim D(0, 1, \xi_k)$.

Conditionally on regime k, the conditional variance $h_{k,t} = h(y_{t-1}, h_{k,t-1}, \theta_k)$ follows a GARCH-type model, as in Haas et al. (2004).[7] We consider three specifications. The GARCH Bollerslev (1986), TGARCH Zakoian (1994) and EGARCH Nelson (1991). The models are described by the following equations:

$$\text{GARCH:} \quad h_{k,t} = \alpha_{0,k} + \alpha_{1,k} y_{t-1}^2 + \beta_k h_{k,t-1}, \tag{3.2}$$

$$\text{EGARCH:} \quad \ln(h_{k,t}) = \alpha_{0,k} + \alpha_{1,k}(|\eta_{k,t-1}| - \mathbf{E}[|\eta_{k,t-1}|]) + \alpha_{2,k}\eta_{t-1} + \beta_k \ln(h_{k,t-1}), \tag{3.3}$$

$$\text{TGARCH:} \quad h_{k,t}^{1/2} = \alpha_{0,k} + \alpha_{1,k} y_{t-1} \mathbf{I}\{y_{t-1} \geq 0\} + \alpha_{2,k} y_{t-1} \mathbf{I}\{y_{t-1} < 0\} + \beta_k h_{k,t-1}^{1/2}, \tag{3.4}$$

where $\mathbf{I}$ is the indicator function. For the conditional distribution we use the Normal, the GED and the Student's-t distribution.[8] Positivity and covariance stationarity conditions apply for all models. To ensure positivity in the GARCH and TGARCH models, we require that all coefficients are positive. No constraints are necessary in the EGARCH model. Covariance stationarity is obtained for the GARCH model by requiring that $\alpha_{1,k} + \beta_k < 1$. For the TGARCH model the constraint is $\alpha_{1,k}^2 + \beta_k^2 - 2\beta_k(\alpha_{1,k} + \alpha_{2,k}\mathbf{E}[\eta_{k,t}\mathbf{I}\{\eta_{k,t} < 0\}] - \alpha_{1,k}^2 - \alpha_{2,k}^2)\mathbf{E}[\eta_{k,t}^2\mathbf{I}\{\eta_{k,t} < 0\}] < 1$, see Francq and Zakoian (2010). Covariance stationarity for the EGARCH models requires that $\beta_k < 1$. The required constraints do not prevent the asymmetric models (TGARCH and EGARCH) from capturing the (inverse) leverage effect. In equation (3.3), ($\alpha_{2,k} > 0$) $\alpha_{2,k} < 0$ indicates the presence of (inverse) leverage effect. To examine the sign of asymmetry in equation (3.4), we examine whether the quantity $\alpha_{1,k} - \alpha_{2,k}$ is positive or not. In the first case,

[7] θ_k is the vector of additional regime-dependent parameters.

[8] The PDF of the GED is given by: $f(\eta; \nu) = \frac{\nu e^{-\frac{1}{2}|\eta/\lambda|^\nu)}}{\lambda 2^{(1+1/\nu)}\Gamma(1/\nu)}$, where $\lambda = \left(\frac{\Gamma(1/\nu}{4^{1/\nu}\Gamma(3/\nu)}\right)^{1/2}$, $\eta \in \mathbf{R}$ and the shape parameter ν is positive.

there is leverage effect while in the second case there is inverse leverage effect.

The models are estimated using a Bayesian MCMC approach. We denote $\mathbf{\Psi} = \{\xi_1, \theta_1, \ldots, \xi_k, \theta_k, \mathbf{P}\}$ the vector of model parameters, $f(y_t|I_{t-1})$ the density of y_t given the past observations, I_{t-1} . The likelihood function is then written as:

$$\mathcal{L}(\mathbf{\Psi}|I_T) = \prod_{t=1}^{T} f(y_t|\mathbf{\Psi}, I_{t-1}). \tag{3.5}$$

The conditional density of y_t is:

$$f(y_t|\mathbf{\Psi}, I_{t-1}) = \sum_{i=1}^{K}\sum_{j=1}^{K} p_{i,j} z_{i,t-1} f_D(y_t|s_t = j, \mathbf{\Psi}, I_{t-1}),$$

where $z_{i,t-1} = P[s_{t-1} = i|\mathbf{\Psi}, I_{t-1}]$ represents the filtered probability of state i at time $t-1$ obtained via the filter of Hamilton (1989) and f_D is the conditional density of y_t given $\mathbf{\Psi}$ and I_{t-1}.

Following Ardia (2008), the likelihood function is combined with a prior $f(\mathbf{\Psi})$ to build the kernel of posterior distribution $f(\mathbf{\Psi}|I_T)$. The prior is built from independent diffuse priors as follows:

$$\begin{aligned}
f(\mathbf{\Psi}) =& f(\boldsymbol{\theta}_1, \boldsymbol{\xi}_1) \cdots f(\boldsymbol{\theta}_k, \boldsymbol{\xi}_k) f(\boldsymbol{P}), \\
f(\boldsymbol{\theta}_k, \boldsymbol{\xi}_k) \propto& f(\boldsymbol{\theta}_k) f(\boldsymbol{\xi}_k) \mathbf{I}\{\boldsymbol{\theta}_k, \boldsymbol{\xi}_k \in \mathcal{CSC}_k\}, \\
f(\boldsymbol{\theta}_k) \propto& f_{\mathcal{N}}(\boldsymbol{\theta}_k; \boldsymbol{\mu}_{\boldsymbol{\theta}_k}, \text{diag}(\boldsymbol{\sigma}^2_{\boldsymbol{\theta}_k})) \mathbf{I}\{\boldsymbol{\theta}_k \in \mathcal{PC}_k\}, \\
f(\boldsymbol{\xi}_k) \propto& f_{\mathcal{N}}(\boldsymbol{\xi}_k; \boldsymbol{\mu}_{\boldsymbol{\xi}_k}, \text{diag}(\boldsymbol{\sigma}^2_{\boldsymbol{\xi}_k})) \mathbf{I}\{\xi_{k,1} > 0\}, \\
f(\boldsymbol{P}) \propto& \prod_{i=1}^{K} \left(\prod_{j=1}^{K} p_{i,j} \right) \mathbf{I}\{0 < p_{i,i} < 1\},
\end{aligned}$$

where $k = 1, \ldots, K$, $\mathcal{CSC}_k$ and $\mathcal{PC}_k$ denote the covariance stationarity condition and positivity condition in regime k (see Trottier and Ardia (2016)). $\xi_{k,1}$ is the asymmetry parameter in regime k. $f_{\mathcal{N}}(, \mu, \Sigma)$ denotes the multivariate Normal density with mean vector μ and covariance matrix Σ. The entries of prior means, $\mu_\bullet$ and variances, $\sigma^2_\bullet$ vectors are set to zero and 1000 respectively. The posterior must be approximated by simulation techniques, since is of unknown form. MCMC draws from the posterior are generated using the adaptive random walk Metropolis sampler of Vihola (2012). We set the number of the burn in draws to 5000. We then build a posterior sample of 5000 from 50000 draws by keeping every 10th draw.

3.3.2 Deviance and Bayesian predictive information criteria

The goodness-of-fit of each estimated model is evaluated using the Deviance (DIC) and Bayesian Predictive (BPIC) information criteria. DIC is a hierarchical modelling generalisation of the Akaike Information Criterion used in Bayesian model selection.[9] Deviance is defined as $D(\mathbf{\Psi}) = -2\ln\left(\mathcal{L}(\mathbf{\Psi}|I_T)\right) + C$, with C a constant that cancels out in all calculations that compare different models, and which therefore does not need to be known. The effective number of parameters of the models, p_D is calculated as in Spiegelhalter et al. (2002): $p_D = \overline{D(\mathbf{\Psi})} - D(\overline{\mathbf{\Psi}})$ where $\overline{\mathbf{\Psi}}$ is the expectation of Ψ and $\overline{D(\mathbf{\Psi})}$ is computed as the average of $D(\mathbf{\Psi})$ over the samples of $\mathbf{\Psi}$. DIC is then calculated as: DIC $= p_D + \overline{D(\mathbf{\Psi})}$. BPIC is an extension of the DIC, suggested by Ando (2008) to avoid over-fitting problems of DIC. BPIC is calculated as BPIC $= -2\mathbf{E}^{\mathbf{\Psi}}[\ln\left(\mathcal{L}(\mathbf{\Psi}|I_T)\right) + 2p_D]$

3.3.3 Forecasting volatility

For the one-step-ahead volatility forecast, we employ half of the total number of observations of y_t for estimation and test the performance of the models over the same number of out-of-sample observations. All models are estimated on a rolling window basis with the parameters being updated every 10 observations. Since volatility itself is unobservable, the comparison of volatility forecasts relies on an observable proxy for the latent volatility process. We assess the forecasting performance of the different GARCH types using employ two criteria, the Mean Square Error (MSE) and the Mean Absolute Error (MAE).

3.3.4 Density risk forecasting

The one-step-ahead Value-at-Risk forecast (VaR) is estimated as a quantile of the predictive density function (CDF), F, by numerically inverting the predictive CDF. Specifically, for a given risk level α,

$$\text{VaR}^{\alpha}_{T+1} = \inf\{y_{T+1} \in \mathbf{R} | F(y_{T+1}|I_T) = \alpha\}. \tag{3.6}$$

To compute the CDF, we first note that the one-step-ahead conditional probability density function (PDF) of y_{T+1} is a mixture of K regime-dependent distributions:

$$f(y_{T+1}|\mathbf{\Psi}, I_T) = \sum_{k=1}^{K} \pi_{k,T+1} \times f_{\mathcal{D}}(y_{T+1}|s_{T+1} = k, \mathbf{\Psi}, I_T),$$

[9] Brooks and Burke (2003) argue that unmodified traditional information criteria cannot be used for order determination of conditional heteroscedastic models.

with mixing weights $\pi_{k,T+1} = \sum_{i=1}^{K} p_{i,k} z_{i,T}$, where $z_{i,T} = \mathbf{P}[s_T = i|\mathbf{\Psi}, I_T]$, for $i = 1, \ldots, K$ are the filtered probabilities at time T. The CDF is then given by:

$$F(y_{T+1}|\mathbf{\Psi}, I_T) = \int_{-\inf}^{y_{T+1}} f(z|\mathbf{\Psi}, I_T)\mathrm{d}z.$$

Given a posterior sample $\{\mathbf{\Psi}^{[m]}, m = 1, \ldots, M\}$, the predictive PDF is obtained as:[10]

$$f(y_{T+1}|I_T) = \int_{\mathbf{\Psi}} f(y_{T+1}|\mathbf{\Psi}, I_T) f(\mathbf{\Psi}|I_T)\mathrm{d}\mathbf{\Psi} \approx \frac{1}{M}\sum_{1=m}^{M} f(y_{T+1}|\mathbf{\Psi}^{[\mathbf{m}]}, I_T),$$

and the predictive CDF is given by:

$$F(y_{T+1}|I_T) = \int_{-\inf}^{y_{T+1}} f(z|I_T)\mathrm{d}z$$

We test the accuracy of the VaR predictions using the conditional coverage (CC) and dynamic quantile (DQ) backtesting procedures. To investigate CC, Christoffersen (1998) proposed a series of VaR exceedance $d_t, t = S, \ldots, S + H$, where $d_t^{\alpha} = I\{r_t < \text{VaR}_t(a)\}$, usually referred to as the hitting series. Specifically, if correct CC is achieved by the model, VaR exceedances should be independently distributed over time.

The DQ test by Engle and Manganelli (2004) assesses the joint hypothesis that $\mathbf{E}[\delta_t^{\alpha}] = \alpha$ and that the hit variables are distributed independently. The implementation of the test involves the de-meaned process, $\text{Hit}_t^{\alpha} = \delta_t^{\alpha} - \alpha$. Under the correct model specification, unconditionally and conditionally, Hit_t^{α} has zero mean and is serially uncorrelated. The DQ test is then the traditional Wald test of the joint nullity of all coefficients in the following linear regression:

$$\text{Hit}_t^{\alpha} = \delta_0 + \sum_{l=1}^{L} \text{Hit}_{t-l}^{\alpha} + \delta_{L+1}\text{VaR}_{t-1}^{\alpha} + \eta_t.$$

Under the null hypothesis of correct unconditional and conditional coverage, we note that the Wald test statistic is asymptotically $\mathcal{X}^2$ distributed with $L + 2$ degrees of freedom.[11] Similar to volatility forecasts, we employ half of the total number of observations of the time-series for estimation and test the performance of the models over the same number of out-of-sample observations. The model parameters are estimated for every 10 observations. Ardia and Hoogerheide (2014) and Ardia et al. (2018) show, in the context of GARCH models, that the performance of VaR forecasts is not affected significantly when the updating frequency is low.

[10] By integrating out the parameter uncertainty.
[11] We follow the standard choice and set $L = 4$ lags.

The analysis is conducted in the R language using the MSGARCH and GAS packages by Ardia et al. (2019) and Ardia et al. (2019), respectively.[12] MSGARCH is used for model estimation and forecasting. The package is implemented such that positivity and covariance stationarity are ensured in the estimation. The package also provides estimates for the two information criteria. The GAS package is used to implement the CC and DQ tests and compute the p-values for the two tests.

3.4 Data

We download the closing prices for a total of 292 digital currencies from Yahoo Finance. Table B1 shows the symbols of the cryptocurrencies used in the empirical analysis. The downloaded time-series have different sample lengths based on the date of the release date of each cryptocurrency and the availability of the data. For all time-series the last observation is for 16/9/2020. The most extensive sample consists of 2583 observations (the first observation is on 1/10/2013). Data for Ripple (XRP) and Litecoin (LTC) are also available from 1/10/2013. In total, seventeen time-series have a sample length equal to 2583 observations.

We use the retrieved data and calculate the daily returns as the natural logarithmic differences of closing prices. Table 3.1 reports the summary statistics of the returns for thirty digital currencies with different size of market capitalisation (large, medium and small). We report the mean, the median, the standard deviation, the skewness and the kurtosis coefficients of the entire sample. For all three groups we observe similar results. Standard deviation ranges at similar levels and skewness is both positive and negative. Furthermore, in all groups we observe cases of extreme kurtosis (USDT, AE and IOC). In addition to the descriptive statistics with employ the Jarque-Bera and the ADF tests to examine the normality and stationarity of the returns. In all cases, the tests indicate that the series are stationary and do not follow the Normal distribution. These results are not reported and are available on request.

Figure 3.2 shows the closing prices for Bitcoin, Ethereum, Tether and Ripple scaled accordingly. Bitcoin and Ethereum are plotted against the left axis and Tether and Ripple against the right axis. We observe a burst in prices that starts in 2017 and lasts until the end of 2018. The next four subfigures plot the returns for the four cryptocurrencies where deviations from the mean are obvious. The means and standard deviations for the ten largest cryptocurrencies are presented in Figure 3.1 and respective values for all cryptocurrencies are presented in the Appendix, in Figures C1 and C2. The

[12]The MSGARCH package can be found in cran.r-project.org/MSGARCH and the GAS package in cran.r-project.org/GAS, see also Ardia et al. (2018).

Figure 3.1: Mean values for the returns (scaled to 10^3) and standard deviations for the ten major cryptocurrencies.

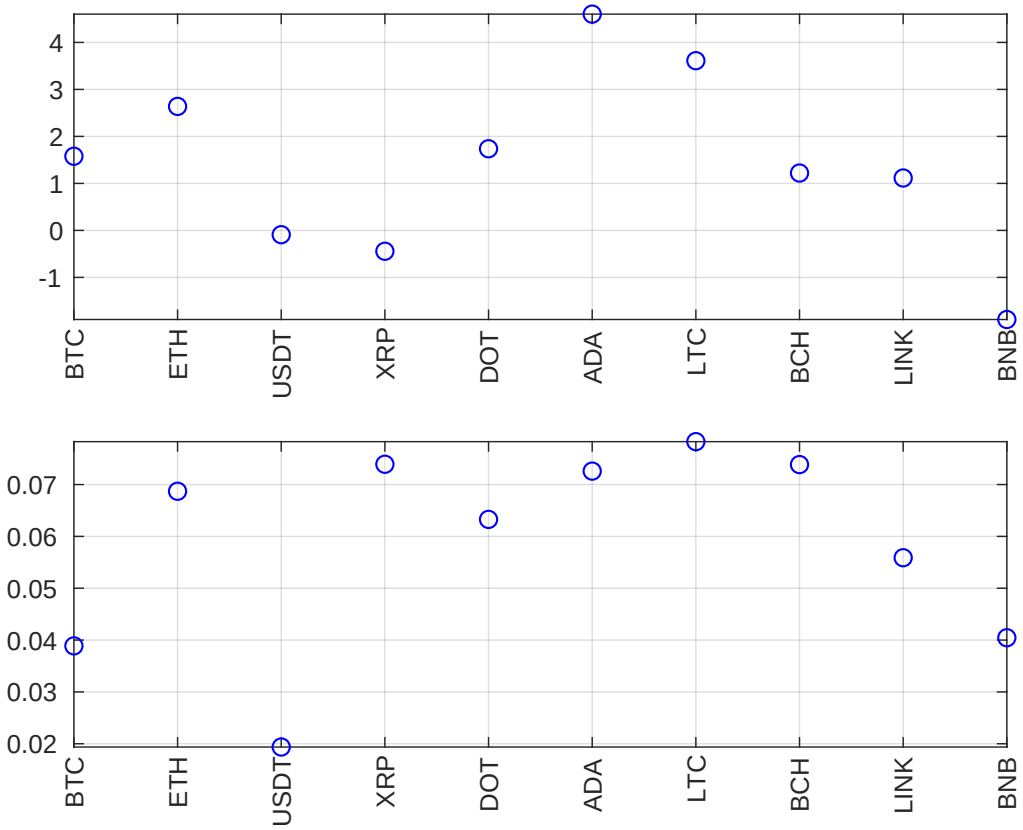

examination of the standard deviations of all cryptocurrencies suggests that digital currencies with smaller market capitalisation have a greater standard deviation. In addition, we observe cases with extreme standard deviation (compared to the rest of the cryptocurrencies) such as WICC, DEV and MOON.

Table 3.1: Descriptive statistics of return prices for 30 cryptocurrenies.

Symbol	Mean	Median	Std	Skewness	Kurtosis	Max	Min
Large market-cap							
BTC	1.57e-03	1.94e-03	0.03	-0.94	13.48	0.22	-0.46
ETH	2.63e-03	3.12e-05	0.06	-3.47	72.03	0.41	-1.30
USDT	-9.15e-05	0.00	0.01	-13.11	970.62	0.50	-0.68
XRP	-4.43e-04	-2.66e-03	0.07	0.23	10.26	0.43	-0.56
BSV	1.73e-03	-2.13e-03	0.06	2.80	45.66	1.02	-0.61
ADA	4.60e-03	7.89e-04	0.07	0.95	15.30	0.67	-0.54
LTC	3.61e-03	-1.05e-03	0.07	0.19	6.82	0.48	-0.61
BCH	1.22e-03	9.88e-05	0.07	2.31	26.89	0.86	-0.50
LINK	1.11e-03	-5.77e-04	0.05	0.41	14.71	0.51	-0.51
BNB	-1.89e-03	-4.17e-03	0.04	0.27	7.25	0.22	-0.18
Medium market-cap							
VSYS	-3.08e-03	-2.37e-03	0.05	-0.04	10.62	0.28	-0.25
SERO	6.26e-04	-3.52e-03	0.04	0.92	2.71	0.19	-0.12
MAID	8.44e-04	9.88e-04	0.06	-0.28	6.76	0.49	-0.57
TFUEL	9.09e-03	-8.19e-03	0.10	3.28	17.32	0.67	-0.25
AION	-1.43e-03	2.00e-03	0.05	-0.44	2.97	0.17	-0.27
AE	-1.56e-03	-5.49e-04	0.10	-2.39	34.73	0.62	-1.21
DEV	-2.70e-03	-4.16e-03	0.61	-0.23	31.49	4.42	-4.66
IOTX	5.96e-03	-3.43e-04	0.07	0.74	2.78	0.35	-0.18
XNC	3.76e-03	2.31e-03	0.02	0.85	6.87	0.13	-0.08
GXC	-1.39e-03	-2.29e-03	0.04	-1.12	6.87	0.11	-0.28
Large market-cap							
TUBE	-2.61e-03	-9.64e-03	0.08	0.76	2.10	0.36	-0.23
CURE	1.38e-04	-6.55e-03	0.05	0.97	3.32	0.24	-0.15
BST	-4.71e-03	-9.97e-03	0.24	-0.23	0.74	0.57	-0.83
NLC2	4.54e-04	-5.78e-03	0.16	1.12	37.65	2.03	-1.93
IOC	6.85e-04	-1.53e-03	0.17	0.07	276.73	4.03	-4.02
MGO	-5.02e-03	-1.68e-03	0.11	-0.25	4.55	0.62	-0.67
EDG	-1.57e-03	-6.09e-03	0.09	0.34	27.08	1.15	-1.02
SUB	-3.07e-03	-3.47e-03	0.09	-0.02	3.17	0.45	-0.56
ATB	-6.31e-03	-4.41e-03	0.12	0.07	8.42	0.71	-0.83
XUC	-3.94e-04	-1.28e-03	0.07	3.12	36.21	0.85	-0.40

Notes: The table reports the summary statistics for cryptocurrencies with different level of market capitalisation.

Figure 3.2: Daily closing prices and returns of the Bitcoin, Ethereum, Tether and Ripple.

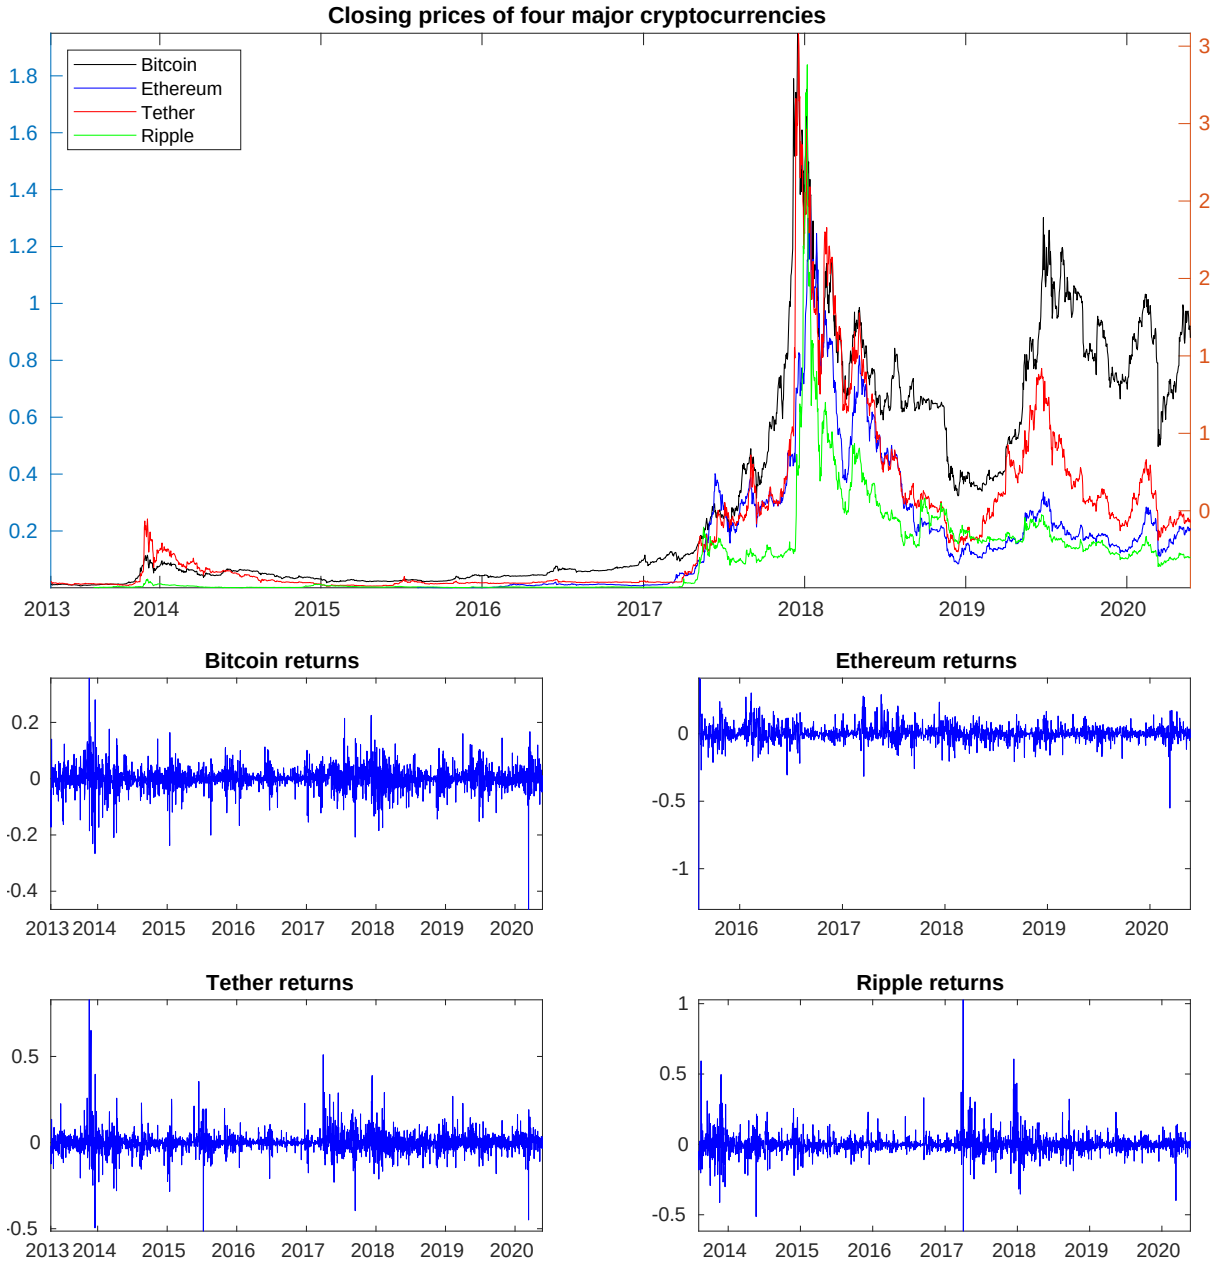

Notes: The first figure shows the closing prices of Bitcoin (black), Ethereum (blue), Tether (red) and Ripple (red) scaled appropriately. Bitcoin and Ethereum are plotted against the left axis and Tether and Ripple against the right axis. The next four figures plot the returns of the four cryptocurrencies.

3.5 Results

3.5.1 In-sample analysis

This section evaluates the performance of the models in terms of goodness-of-fit. Table 3.2 summarises the results regarding model selection through information criteria. Table 3.2 reports the absolute frequency at which each model is selected via DIC (Panel A) and BPIC (Panel B). The same results are also presented in the first subplot of Figure 3.3. According to the findings, DIC indicates single-regime models for 111 out 292 time-series, and regime-switching models for 181 time-series (119 two-regime and 62 three-regime models). On the contrary, BPIC selects single-regime models more often than regime-switching models. Specifically, BPIC indicates single-regime models for 162 digital currencies, two-regime for 91 and three-regime models for 39 models. Despite these differences, the two criteria select the same model in terms of goodness-of-fit for 192 out of 292 cryptocurrencies. If we confine the analysis to the set of cryptocurrencies for which both criteria indicate the same model, we observe that multi-regime models are preferred in more cases than single-regime models (100 over 92). In addition, we observe that the models with Student's-t conditional distribution outperform other models regardless of the GARCH-type and the number of regimes. These findings are consistent with the stylised empirical facts of financial time-series as discussed in Cont (2001) and Cont and Tankov (2004). Finally, we find that for a given number of regimes, the exponential model is outperformed by the other models.

For the three out of the four most traded cryptocurrencies, the two criteria indicate the same model as the most appropriate; for Bitcoin and Tether, the best model in terms of goodness-of-fit is the the two-regime TGARCH with Student's-t conditional distribution. In the case of Ethereum, the three-regime GARCH with Generelised distribution is selected. We also find evidence of regime changes in the volatility of Ripple, however the criteria indicate different models. In some cases, DIC and BPIC suggest models with different number of regimes for the same digital currency. Among these cryptocurrencies are Bitcoin Cash, Binance and Litecoin.

As discussed in the literature review, positive past returns have a greater effect on Bitcoin volatility. Here, we study the behaviour of leverage effect for all 292 cryptocurrencies, using the two asymmetric models (EGARCH and TGARCH). Table 3.3 reports the number of cryptocurrencies for which each model indicates the presence of inverse leverage effect in at least one regime. Clearly, for the majority of the time-series, both models support the hypothesis of inverse leverage effect. For example the traditional EGARCH model suggests positive asymmetry in the volatility of 177 out of 292 digital currencies. In total, there are 10 cryptocurrencies in which inverse leverage

Table 3.2: Model selection based on DIC and BPIC.

	$K=1$			$K=2$			$K=3$		
	N	Std	GED	N	Std	GED	N	Std	GED
	Panel A: DIC								
GARCH	1	29	21	5	28	7	6	17	6
EGARCH	0	14	7	10	13	6	0	7	4
TGARCH	0	30	9	10	34	6	3	15	4
	Panel B: BPIC								
GARCH	1	46	41	6	24	9	4	11	3
EGARCH	0	17	3	3	7	4	0	3	2
TGARCH	1	42	11	5	25	8	2	12	2
	Panel C: Counts of agreement between DIC and BPIC								
GARCH	1	27	21	3	18	5	4	9	3
EGARCH	0	8	3	3	6	4	0	3	2
TGARCH	0	26	6	4	22	2	1	9	2

Notes: The total number of cryptocurrencies is 292. The values denote the number of time-series for which the corresponding model is selected through DIC (panel A), BPIC (panel B) and both criteria (panel C). N, Std and GED denote the Normal, Student's-t and GED distributions, respectively. K refers to the number of regimes in the model.

effect is indicated by all models. These are LINK, XTZ, HIVE, XWC, PLC, ZNN, QBSR, QRK, SAPP, USNBT. On the contrary, there are 3 cases where according to the EGARCH, negative past returns affect the volatility more than positive past returns. These are CTC, DUN, XMC. For all cryptocurrencies, at least one specification of TGARCH indicates inverse leverage effect. Considering the major cryptocurrencies, we observe that most models indicate inverse leverage effect in at least one regime. For Bitcoin only the single-regime models with normal conditional distribution do not support the hypothesis of inverse leverage effect.

Table 3.3: Inverse leverage effect.

	$K=1$			$K=2$			$K=3$		
	N	Std	GED	N	Std	GED	N	Std	GED
EGARCH	177	196	202	245	240	241	265	261	262
TGARCH	187	218	208	249	255	257	268	268	267

Notes: The total number of cryptocurrencies is 292. The values denote the number of time-series for which the corresponding model indicates the existence of inverse leverage effect in at least one regime. N, Std and GED denote the Normal, Student's-t and GED distributions, respectively. K refers to the number of regimes in the model.

In the last part of the in-sample analysis, we calculate the mean and standard devi-

ation of conditional volatility of each estimated model. The results for the four major cryptocurrencies are presented in Figure C3. In all cases, the mean and the standard deviation values are close for all models. With the exception of Ripple, single-regime models exhibit a greater mean than MSGARCH models. For example, in the case of Ethereum, the two models with the higher mean values of the conditional volatility are the TGARCH and EGARCH with Student's-t distribution. Comparing the different digital currencies, we find that for most models the Bitcoin exhibits greater volatility. Similar plots for all examined cryptocurrencies are presented in the Appendix. From the examination of these figures it occurs that we can divide digital currencies in three groups. The first group contains the cryptocurrencies like Bitcoin and Ethereum for which all mean and standard deviation do not exhibit substantial differences for all models. The second group contains the cryptocurencies where the mean and standard deviation of few models differs from the rest. Tether and Ripple belong in the second group. The last group consists of the cryptocurrencies for which we cannot identify a patter for the calculated means and standard deviations, i.e. NMC and NPC. For most ctyptocurrencies and models we observe a mean higher than the standard deviation. There are however, a few exceptions such as Ethereum, NPCoim, TenX and Veritaseum.

3.5.2 Out-of-sample analysis

We now turn to the out-of-sample analysis. First, we compare the ability of the 27 models to correctly forecast the one-day ahead conditional volatility. We evaluate the performance of the the MSE and MAE loss functions. The main results are presented in Table 3.4 and the middle subfigure of Figure 3.3. Overall, two-regime models outperform both traditional and three-regime models. According to the MSE, two-regime models perform better than the rest of the models 79% of the time. That percentage is a little lower when MAE is used, approximately 73%. According to both criteria, TGARCH outperforms both GARCH and EGARCH models, regardless of the number of regimes or the conditional distribution. In addition, we obtain the most accurate forecasts when we fit TGARCH-type models to the major cryptocurrencies. For 218 out of 292 examined cryptocurrencies, both MSE and MAE agree on the model selection. Similar to the in-sample analysis, EGARCH models are outperformed by the other models. We obtain similar results if we focus only on these models (Panel C of 3.4).

We focus on the major digital currencies such as Bitcoin, Ehtereum, Ripple and Tether. For each one of these four cryptocurrencies both MSE and MAE indicate the same model as the best one. In the case of Bitcoin and Ethereum, single-regime mod-

Table 3.4: Predictive power of the models.

	$K=1$			$K=2$			$K=3$		
	N	Std	GED	N	Std	GED	N	Std	GED
	Panel A: MSE								
GARCH	9	5	7	22	20	27	2	2	2
EGARCH	3	1	1	9	10	5	0	1	0
TGARCH	11	6	3	50	33	55	4	1	3
	Panel A: MAE								
GARCH	11	5	9	22	17	23	4	2	3
EGARCH	2	0	3	10	16	2	3	1	0
TGARCH	10	9	5	39	31	54	5	2	4
	Panel C: Counts of agreement between MSE and MAE								
GARCH	7	4	6	15	12	19	2	2	1
EGARCH	1	0	1	6	9	1	0	1	0
TGARCH	9	5	3	35	27	45	4	1	2

Notes: The total number of cryptocurrencies is 292. The values denote the number of time-series for which the corresponding model is minimises the MSE (panel A), MAE (panel B) and both criteria (panel C). N, Std and GED denote the Normal, Student's-t and GED distributions, respectively. K refers to the number of regimes in the model.

els are chosen, the TGARCH with Normal distribution and the GARCH with Generalised distribution, respectively. In the case of Tether and Ripple, the two-regime TGARCH model is chosen (with Student's-t and generalised error distribution, respectively). Overall, MSGARCH models seem to perform better than their single-regime counterparts. In addition, asymmetric models are preferred to symmetric ones (for a given number or regimes). The latter also applies to the major cryptocurrency of the dataset such as Bitcoin, Tether and Bitcoin Cash where models that account for leverage effect appear as the most appropriate.

Next, we evaluate the performance of the models in forecasting one-day ahead VaR. Similar to conditional volatility predictions, all models are compared against each other in VaR forecasting based on the p-value of the CC and DQ tests. Table 3.5 presents the number of times each model outperforms all other models at forecasting the one-step-ahead VaR (these results are presented in the last panel of Figure 3.3). According to CC, the two-regime TGARCH model with Normal distribution is the best model in predicting VaR. Specifically, CC indicates this model as the best for 72 cryptocurrencies. Similar to volatility forecasts, MSGARCH models outperform single-regime models. The results are different when the DQ test is used. First, there is no model specification which clearly outperforms all others. Both the three-regime GARCH model with Student's t conditional distribution and the two-regime TGARCH with Normal dis-

Table 3.5: Model selection based on VaR forecasts.

	$K=1$			$K=2$			$K=3$		
	N	Std	GED	N	Std	GED	N	Std	GED
	Panel A: CC								
GARCH	15	7	2	17	17	4	9	8	3
EGARCH	2	2	0	19	16	3	1	2	1
TGARCH	15	13	7	72	30	11	3	10	3
	Panel B: DQ								
GARCH	18	16	4	24	13	13	17	27	8
EGARCH	3	3	3	11	8	7	0	1	0
TGARCH	12	17	6	26	19	10	8	14	4
	Panel C: Counts of agreement between CC and DQ								
GARCH	4	2	0	7	4	2	3	2	1
EGARCH	1	1	0	2	2	0	0	1	0
TGARCH	4	4	2	16	7	2	1	5	1

Notes: The total number of cryptocurrencies is 292. The values denote the number of time-series for which the corresponding model is selected through CC (panel A) and DQ (panel B). N, Std and GED denote the Normal, Student's-t and GED distributions, respectively. K refers to the number of regimes in the model.

tribution perform equally well (27 and 26 out of 292, respectively). Second, the GARCH models performs better than TGARCH models. Despite these differences, the EGARCH produces the worst results compared to the other models. CC and DQ agree on the selection of the model in 74 cases (out of 292). For these cases, the results are qualitatively the same with the ones we obtained using the CC test.

Considering the major cryptocurrencies, we observe that different GARCH models are chosen for different cryptocurrencies and as a result we can not arque that only one should be used in predicting the one-step-ahead VAR. In the case of Bitcoin, the three-regime GARCH model is selected. Regime-switching models are also suitable for Ripple. CC indicates a two-regime GARCH model with Normal distribution and DQ a three-regime GARCH model with Student's-t conditional distribution. The best model for forecasting the VaR of Ethereum and Tether is the traditional GARCH model. For the remaining popular digital currencies the asymmetric models (TGARCH) are preferred. Table 3.6 summarises the findings regarding ten major cryptocurrencies both for the in and the out-of-sample analysis.

Table 3.6: Summarised results for the ten major cryptocurrencies.

Symbol	DIC	BPIC	MSE	MAE	CC	DQ
BTC	TGARCH-T(2)	TGARCH-T(2)	TGARCH-N(1)	TGARCH-N(1)	GARCH-G(3)	GARCH-T(3)
ETH	GARCH-T(3)	GARCH-T(2)	GARCH-G(1)	GARCH-G(1)	GARCH-T(1)	GARCH-T(1)
USDT	TGARCH-T(2)	TGARCH-T(2)	TGARCH-T(2)	TGARCH-T(2)	GARCH-T(2)	GARCH-T(2)
XRP	TGARCH-T(3)	GARCH-T(2)	TGARCH-G(2)	TGARCH-G(2)	GARCH-T(2)	TGARCH-T(3)
BSV	GARCH-T(1)	GARCH-T(1)	GARCH-G(1)	TGARCH-T(1)	GARCH-T(2)	TGARCH-T(1)
ADA	GARCH-T(1)	GARCH-T(1)	GARCH-N(1)	GARCH-N(1)	GARCH-N(2)	TGARCH-T(2)
LTC	GARCH-T(3)	TGARCH-T(1)	GARCH-N(1)	GARCH-N(1)	GARCH-N(1)	TGARCH-G(3)
BCH	GARCH-T(3)	TGARCH-T(1)	EGARCH-N(1)	EGARCH-N(3)	EGARCH-T(1)	GARCH-T(1)
LINK	GARCH-T(1)	GARCH-T(1)	GARCH-N(1)	GARCH-N(1)	TGARCH-T(1)	TGARCH-T(1)
BNB	TGARCH-G(2)	GARCH-T(1)	GARCH-T(1)	GARCH-T(1)	EGARCH-T(1)	GARCH-G(1)

Notes: The table reports the results for the ten cryptocurrencies with the greatest capitalisation. -N, -T and -G denote the Normal, Student's-*t* and GED conditional distribution. The number in parentheses denote the number of the regimes in the selected MSGARCH model.

Furthermore, we report the percentages of cryptocurrencies for which the null hypothesis of correct unconditional and conditional coverage is rejected at the 5% significance level (Table 3.7). For all models the failure rate is less than 8%. Based on the CC test, the model with the highest failure rate is the traditional GARCH model with Normal conditional distribution and based on the DQ test, the two-regime EGARCH model with Student's t conditional distribution. For both tests, the model with the lowest rejection frequency is the TGARCH model with three regimes and Student's t conditional distribution. Overall, the rejection level decreases as the number of regimes increases. In addition, for a given number of regimes, the TGARCH models outperform the other models.

Table 3.7: Percentage of cryptocurrencies for which the validity of the VaR predictions is rejected.

	CC			DQ		
No. of regimes	1	2	3	1	2	3
GARCH-N	7.12	6.02	5.13	6.74	6.02	5.41
GARCH-Std	6.57	6.71	6.67	7.12	7.05	6.91
GARCH-GED	6.95	3.93	3.01	6.98	4.04	3.21
EGARCH-N	6.95	4.96	3.18	6.47	5.17	3.21
EGARCH-Std	6.71	6.74	6.57	6.95	7.15	6.95
EGARCH-GED	6.88	3.42	2.08	6.84	3.21	2.15
TGARCH-N	4.41	1.43	0.30	3.93	0.85	0.17
TGARCH-Std	4.86	1.16	0.58	4.34	0.75	0.30
TGARCH-GED	4.38	0.95	0.47	3.73	0.61	0.34

Notes: The values denote the percentage of cryptocurrencies for which the validity of the VaR forecasts is rejected for each model. N, Std and GED denote the Normal, Student's-t and GED distributions, respectively.

3.6 Conclusions

Cryptocurrencies have flourished over the last decade. Since the introduction of Bitcoin, hundreds of digital currencies have been created. Studies that examined Bitcoin and other cryptocurrencies have classified the cryptocurrencies as an asset used mostly for speculative purposes. This is an explanation for the extreme volatility of the returns of Bitcoin. Despite their similarities, there are significant differences between cryptocurrencies. For example, the level of market capitalisation between digital currencies and the attention each digital currency receives from investors are distinct. In this essay, we perform an extensive analysis of the volatility of the cryptocurrency

markets. We examine the volatility of the returns for a total of 292 different cryptocurrencies. For each one we estimate 27 GARCH-type models using Bayesian methods.

First, we perform an in-sample analysis. We evaluate the performance of each model in terms of goodness-of-fit using DIC and BPIC and investigate the presence of leverage effect. While DIC selects MSGARCH models more and BPIC selects single-regime models more, the two criteria agree on the model selection approximately 66% of cryptocurrencies. In this case, MSGARCH models are preferred in 52% of the cases. The examination of leverage effect suggests that most cryptocurrencies returns exhibit inverse leverage effect and respond more strongly to positive past returns than negative past returns.

Next, we perform the out-of-sample analysis and consider two forecast exercises. First, we evaluate the ability of the models to predict the one-step-ahead volatility. We find that two-regime models produce better forecasts than the rest of the models. Next, we assess the ability of the models to predict the one-step-ahead VaR based on conditional and unconditional convergence tests. For both tests the three-regime TGARCH models have the lowest rejection frequencies. Both from in-sample and out-of-sample analysis, we conclude that the EGARCH models are outperformed both by the GARCH and the TGARCH models.

Our study can be expanded in various ways. First, the analysis can be repeated, considering other digital currencies not included here. Second, it would be of interest to see whether including skewed versions of the conditional distributions would affect the results. Finally, we do not consider the GJR-GARCH model in our analysis. The positivity constraints, regarding the GJR-GARCH model, requires that the leverage coefficient is positive. However, this rules out the possibility of inverse leverage effect which is often observed in cryptocurrencies. To consider the GJR-GARCH model, we should ensure positivity as in Ardia et al. (2019).

Figure 3.3: Absolute frequency of model selection based on in-sample and out-of-sample ananlysis.

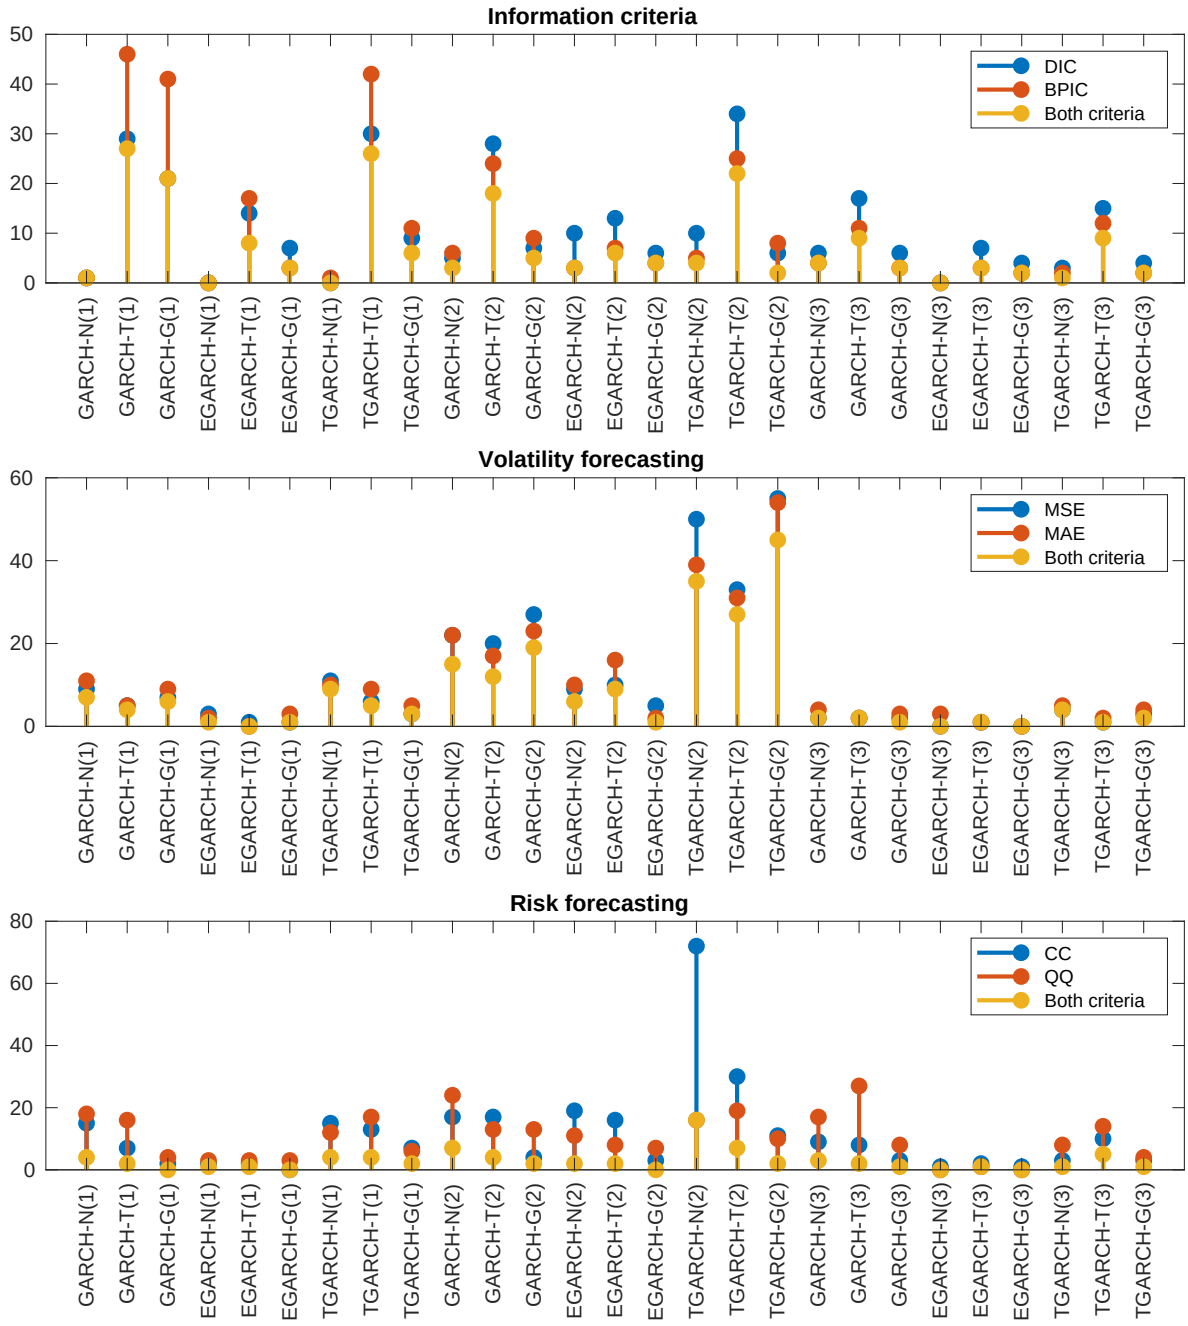

Notes: Each figure presents the absolute frequency at which each model is selected. The total number of examined cryptocurrencies is 292. N, T and G denote the Normal, Student's-t and Generalised distributions. The number in parentheses denote the number of regimes. For instance, according to Bayesian predictive information criterion, GARCH-T(1) is preferred in 46 out of the 292 cases.

Bibliography

Ahamed, L. (2009). *Lords of finance: The bankers who broke the world*. Penguin Press.

Ando, T. (2008). Bayesian predictive information criterion for the evaluation of hierarchical Bayesian and empirical Bayes models. *Biometrika 94*(2), 443–458.

Ardia, D. (2008). *Financial risk management with Bayesian estimation of GARCH Models: Theory and applications*, Volume 612, Lecture Notes in Economics and Mathematical Systems. Berlin Heidelberg: Springer-Verlag.

Ardia, D., K. Bluteau, K. Boudt, and L. Catania (2018). Forecasting risk with Markov-switching GARCH models: A large-scale performance study. *International Journal of Forecasting 34*(4), 733–747.

Ardia, D., K. Bluteau, K. Boudt, L. Catania, and D.-A. Trottier (2019). Markov-switching GARCH models in R: The **MSGARCH** package. *Journal of Statistical Software 91*(4), 1–38.

Ardia, D., K. Bluteau, and M. Ruede (2019). Regime changes in Bitcoin GARCH volatility dynamics. *Finance Research Letters 29*(3), 266–271.

Ardia, D., K. Boudt, and L. Catania (2018). Downside risk evaluation with the R package **GAS**. *The R Journal 10*(2), 410–421.

Ardia, D., K. Boudt, and L. Catania (2019). Generalized Autoregressive Score Models in R: The **GAS** Package. *Journal of Statistical Software 88*(6), 1–28.

Ardia, D. and L. Hoogerheide (2014). Garch models for daily stock returns: Impact of estimation frequency on value-at-risk and expected shortfall forecasts. *Economics Letters 123*(2), 187 – 190.

Ardia, D., J. Kolly, and D.-A. Trottier (2017). The impact of parameter and model uncertainty on market risk predictions from garch-type models. *Journal of Forecasting 36*(7), 808–823.

Baek, C. and M. Elbeck (2015). Bitcoins as an investment or speculative vehicle? A first look. *Applied Economics Letters 22*(1), 30–34.

Balcombe, K. and I. Fraser (2017). Do bubbles have an explosive signature in markov switching models? *Economic Modelling 66*, 81 – 100.

Bariviera, A. F. (2017). The inefficiency of bitcoin revisited: A dynamic approach. *Economics Letters 161*, 1–4.

Baur, D., T. Dimpfl, and K. Kuck (2018). Bitcoin, gold and the us dollar – a replication and extension. *Finance Research Letters 25*, 103 – 110.

Bauwens, L., B. De Backer, and A. Dufays (2014). A bayesian method of change-point estimation with recurrent regimes: Application to GARCH models. *Journal of Empirical Finance 29*, 207–229.

Bauwens, L., A. Dufays, and J. V. Rombouts (2014). Marginal likelihood for markov-switching and change-point garch models. *Journal of Econometrics 178*, 508 – 522.

Bauwens, L., A. Preminger, and J. Rombouts (2010). Theory and inference for a markov switching GARCH model. *The Econometrics Journal 13*(2), 218–244.

Blau, B. (2017). Price dynamics and speculative trading in bitcoin. *Research in International Business and Finance 41*(3), 493–499.

Bollerslev, T. (1986). Generalized autoregressive conditional heteroskedasticity. *Journal of Econometrics 31*(3), 307–327.

Bouoiyour, J. and R. Selmi (2015). What does bitcoin look like? *Annals of Economics and Finance 16*(2), 449–492.

Bouri, E., G. Azzi, and A. Dyhrberg (2017). On the return-volatility relationship in the bitcoin market around the price crash of 2013. *Economics E-Journal 11*, 1–16.

Bouri, E., R. Gupta, A. K. Tiwari, and D. Roubaud (2017). Does bitcoin hedge global uncertainty? evidence from wavelet-based quantile-in-quantile regressions. *Finance Research Letters 23*, 87 – 95.

Bouri, E., P. Molnár, G. Azzi, D. Roubaud, and L. I. Hagfors (2017). On the hedge and safe haven properties of bitcoin: Is it really more than a diversifier? *Finance Research Letters 20*, 192 – 198.

Brooks, C. and S. Burke (2003). Information criteria for GARCH model selection. *European Journal of Finance 9*(6), 557–580.

Caporale, G. M. and T. Zekokh (2019). Modelling volatility of cryptocurrencies using markov-switching garch models. *Research in International Business and Finance 48*, 143–155.

Cheikh, N. B., Y. B. Zaied, and J. Chevallier (2020). Asymmetric volatility in cryptocurrency markets: New evidence from smooth transition GARCH models. *Finance Research Letters 35*, 101293.

Christoffersen, P. (1998). Evaluating interval forecasts. *International Economic Review 39*(4), 841–862.

Chu, J., S. Chan, S. Nadarajah, and J. Osterrieder (2017). Garch modelling of cryptocurrencies. *Journal of Risk and Financial Management 10*, 17.

Cont, R. (2001). Empirical properties of asset returns: stylized facts and statistical issues. *Quantitative Finance 1*(2), 223–236.

Cont, R. and P. Tankov (2004). *Financial Modelling With Jump Processes*. Boca Raton London New York Washington D. C.: Chapman and Hall/CRC.

Corbet, S., B. Lucey, and L. Yarovya (2018). Datestamping the bitcoin and ethereum bubbles. *Finance Research Letters 26*, 81–88.

Dyhrberg, A. H. (2016a). Bitcoin, gold and the dollar - A GARCH volatility analysis. *International Review of Financial Analysis 16*, 85–92.

Dyhrberg, A. H. (2016b). Hedging capabilities of bitcoin. Is it the virtual gold? *Finance Research Letters 16*(C), 139–144.

Engle, R. and S. Manganelli (2004). Caviar: Conditional autoregressive value at risk by regression quantiles. *Journal of Business and Economic Statistics 22*(4), 367–381.

Francq, C. and J.-M. Zakoian (2010). *GARCH Models: zStructure, Statistical Inference and Financial Applications*. John Wiley & Sons.

Fry, J. and E.-T. Cheah (2016). Negative bubbles and shocks in cryptocurrency markets. *International Review of Financial Analysis 47*, 343–352.

Geweke, J. and G. Amisano (2010). Comparing and evaluating bayesian predictive distributions of asset returns. *International Journal of Forecasting 26*(2), 216 – 230. Special Issue: Bayesian Forecasting in Economics.

Glaser, F., K. Zimmermann, M. Haferkorn, M. Weber, and M. Siering (2014). Bitcoin - asset or currency? revealing users' hidden intentions. Technical report, ECIS 2014 Proceedings - 22nd European Conference on Information Systems.

Gronwald, M. (2014). The economics of bitcoins - market characteristics and price jumps. Technical report, CESifo Working Paper Series 5121, CESifo.

Gronwald, M. (2019). Is bitcoin a commodity? on price jumps, demand shocks, and certainty of supply. *Journal of International Money and Finance 97*, 86 – 92.

Haas, M., S. Mittnik, and M. Paolella (2004). A new approach to Markov-switching GARCH models. *Journal of Financial Econometrics 2*, 493–530.

Hamilton, J. (1989). A new approach to the economic analysis of nonstationary time series and the business cycle. *Econometrica 57*(2), 357–384.

Katsiampa, P. (2017). Volatility estimation for Bitcoin: A comparison of GARCH models. *Economics Letters 158*(3), 3–6.

Klein, T., H. P. Thu, and T. Walther (2018). Bitcoin is not the New Gold - A comparison of volatility, correlation, and portfolio performance. *International Review of Financial Analysis 59*, 105–116.

Kristoufek, L. (2014). What are the main drivers of the bitcoin price? evidence from wavelet coherence analysis. FinMaP-Working Papers 23, Collaborative EU Project FinMaP - Financial Distortions and Macroeconomic Performance: Expectations, Constraints and Interaction of Agents.

Lamoureux, C. G. and W. D. Lastrapes (1990). Persistence in variance, structural change, and the garch model. *Journal of Business & Economic Statistics 8*(2), 225–234.

Ma, F., C. Liang, Y. Ma, and M. Wahab (2020). Cryptocurrency volatility forecasting: A Markov regime-switching midas approach. *Journal of Forecasting 39*(8), 1277–1290.

Maciel, L. (2021). Cryptocurrencies value-at-risk and expected shortfall: Do regime-switching volatility models improve forecasting? *International Journal of Finance & Economics 26*(3), 4840–4855.

Mensi, W., K. H. Al-Yahyaee, and S. H. Kang (2019). Structural breaks and double long memory of cryptocurrency prices: A comparative analysis from bitcoin and ethereum. *Finance Research Letters 29*, 222–230.

Nadarajah, S. and J. Chu (2017). On the inefficiency of bitcoin. *Economics Letters 150*(3), 6–9.

Nakamoto, S. (2009). Bitcoin: A peer-to-peer electronic cash system. Available at: https://bitcoin.org/bitcoin.pdf.

Nelson, D. (1991). Conditional heteroskedasticity in asset returns: A new approach. *Econometrica 59*(2), 347–370.

Panagiotidis, T., T. Stengos, and O. Vravosinos (2018). On the determinants of bitcoin returns: A lasso approach. *Finance Research Letters 27*(C), 235–240.

Panagiotidis, T., T. Stengos, and O. Vravosinos (2019). The effects of markets, uncertainty and search intensity on Bitcoin returns. *International Review of Financial Analysis 63*, 220 – 242.

Polasik, M., A. Piotrowska, T. Wisniewski, R. Kotkowski, and G. Lightfoot (2015). Price fluctuations and the use of bitcoin: An empirical inquiry. *International Journal of Electronic Commerce 20*, 9–49.

Reinhart, C. and K. Rogoff (2009). The aftermath of financial crises. *American Economic Review 99*(2), 466–72.

Spiegelhalter, D., N. Best, B. Carlin, and A. Van Der Linde (2002). Bayesian measures of model complexity and fit. *Journal of the Royal Statistical Society: Series B (Statistical Methodology) 64*(4), 583–639.

Stavroyiannis, S. (2018). Value-at-risk and related measures for the bitcoin. *Journal of Risk Finance 19*(2), 127–136.

Thies, S. and P. Molnár (2018). Bayesian change point analysis of bitcoin returns. *Finance Research Letters 27*, 223 – 227.

Trottier, D.-A. and D. Ardia (2016). Moments of standardized Fernandez–Steel skewed distributions: Applications to the estimation of GARCH-type models. *Finance Research Letters 18*(8), 311 – 316.

Urquhart, A. (2016). The inefficiency of bitcoin. *Economics Letters 148*, 80 – 82.

Urquhart, A. (2017). The volatility of bitcoin. Technical report, SSRN Electronic Journal.

Vihola, M. (2012). Robust adaptive metropolis algorithm with coerced acceptance rate. *Statistics and Computating 22*(5), 997–1008.

Williams, M. (2014). Virtual currencies - Bitcoin risk. In *World Bank Conference*. Washington, DC. No. 21.

Yermack, D. (2013). Is bitcoin a real currency? an economic appraisal. NBER Working Papers 19747, National Bureau of Economic Research, Inc.

Zakoian, J.-M. (1994). Threshold heteroskedastic models. *Journal of Economic Dynamics and Control 18*(5), 931–955.

Appendix

Figure C1: Mean values for the returns of the cryptocurrencies.

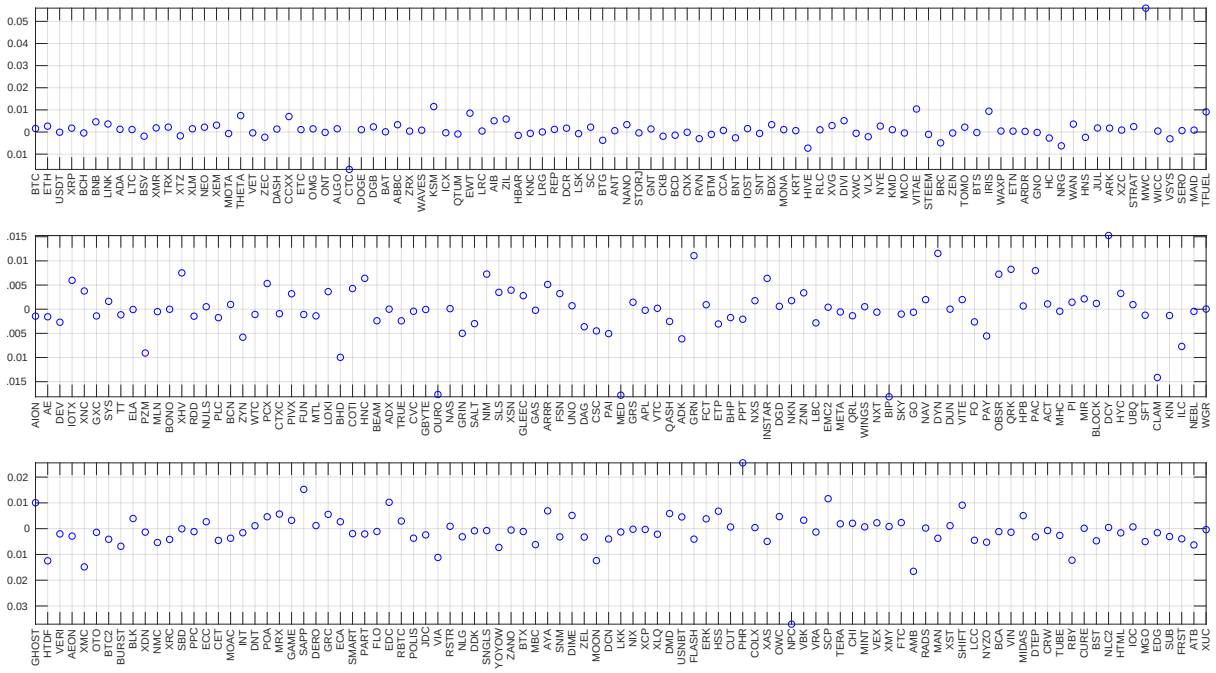

Figure C2: Standard deviations of the returns of the cryptocurrencies.

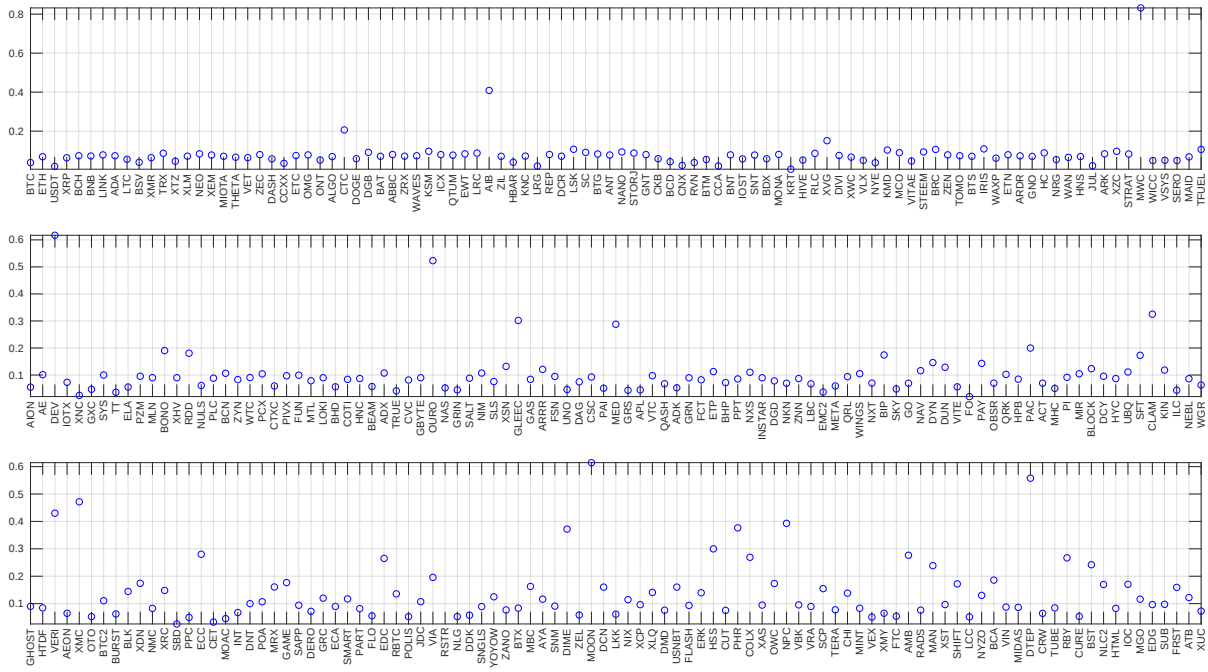

Figure C3: Mean and standard deviation of conditional volatility of cryprocurrencies.

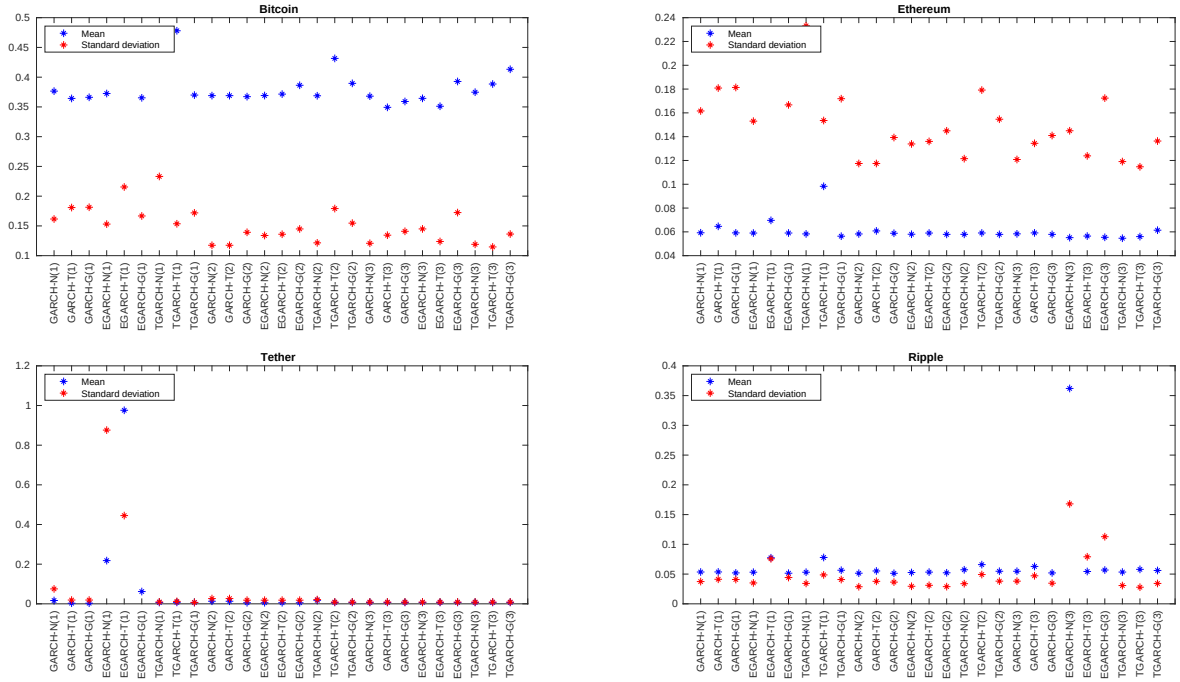

Notes: Each plot presents the mean and standard deviation of the conditional volatility for the different GARCH models. The total number of models is 27. N, T and G denote the Normal, Student's-t and Generalised distributions. The number in parentheses denote the number of regimes.

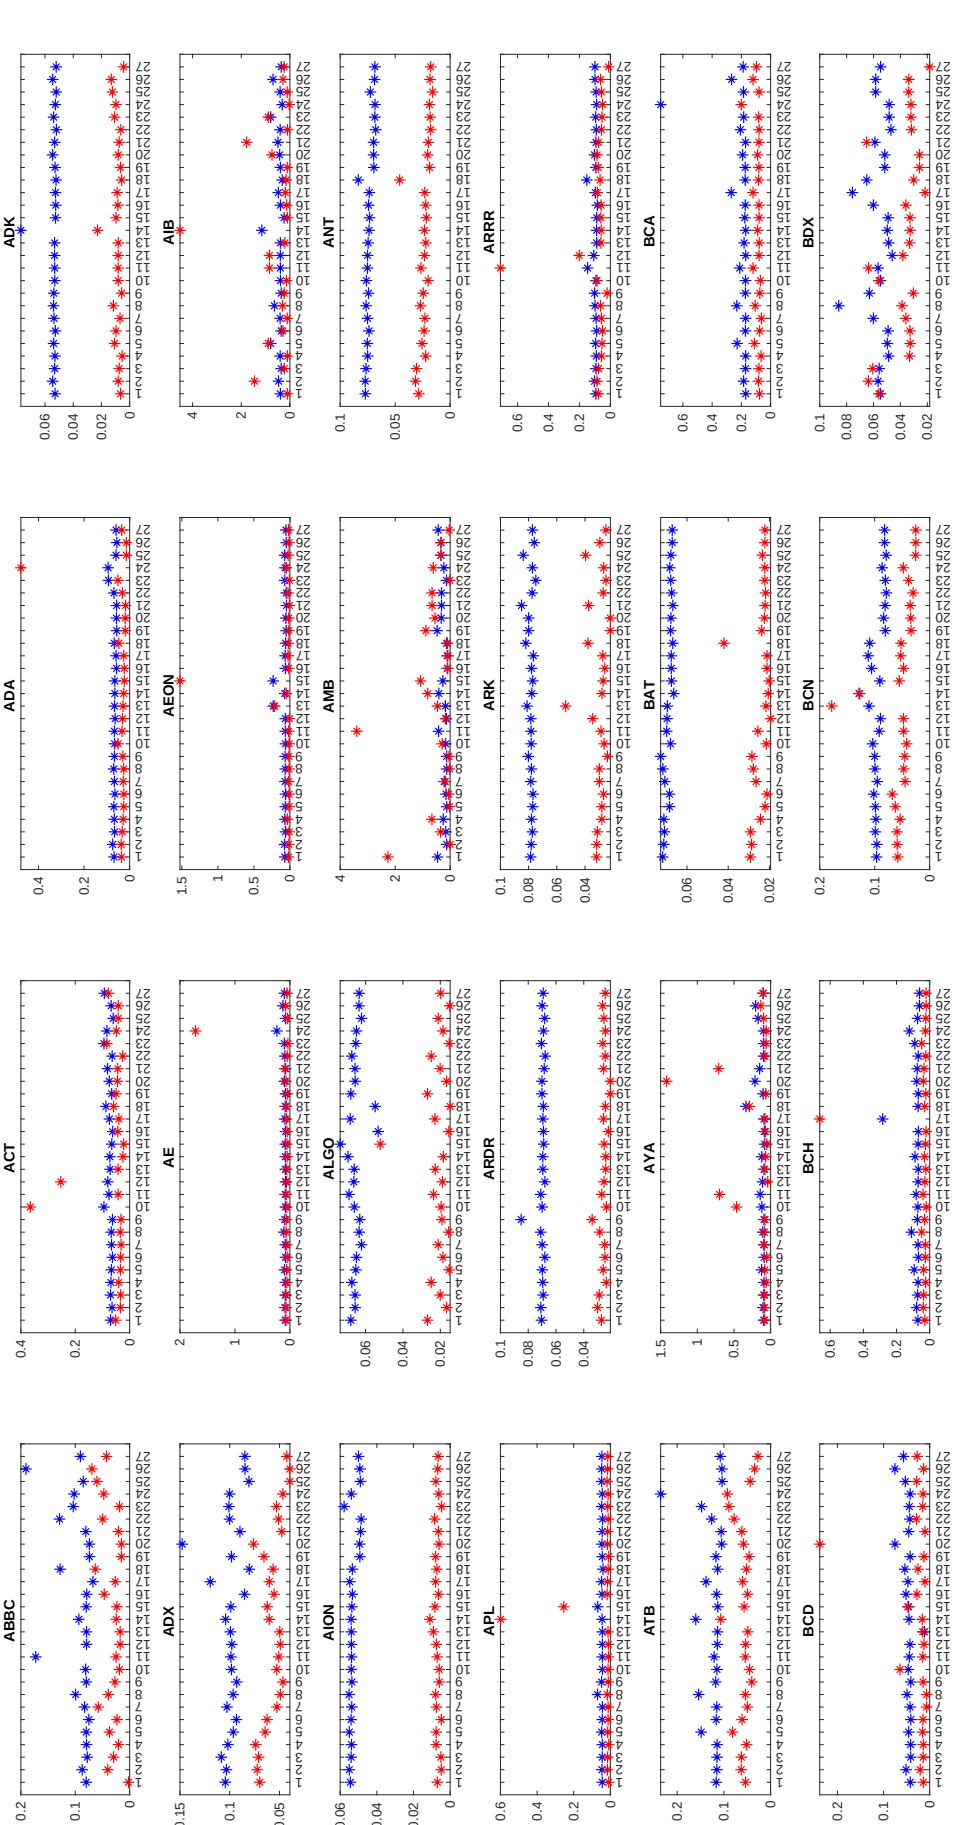
ABBC
ACT
ADA
ADK
ADX
AE
AEON
AIB
AION
ALGO
AMB
ANT
APL
ARDR
ARK
ARRR
ATB
AYA
BAT
BCA
BCD
BCH
BCN
BDX

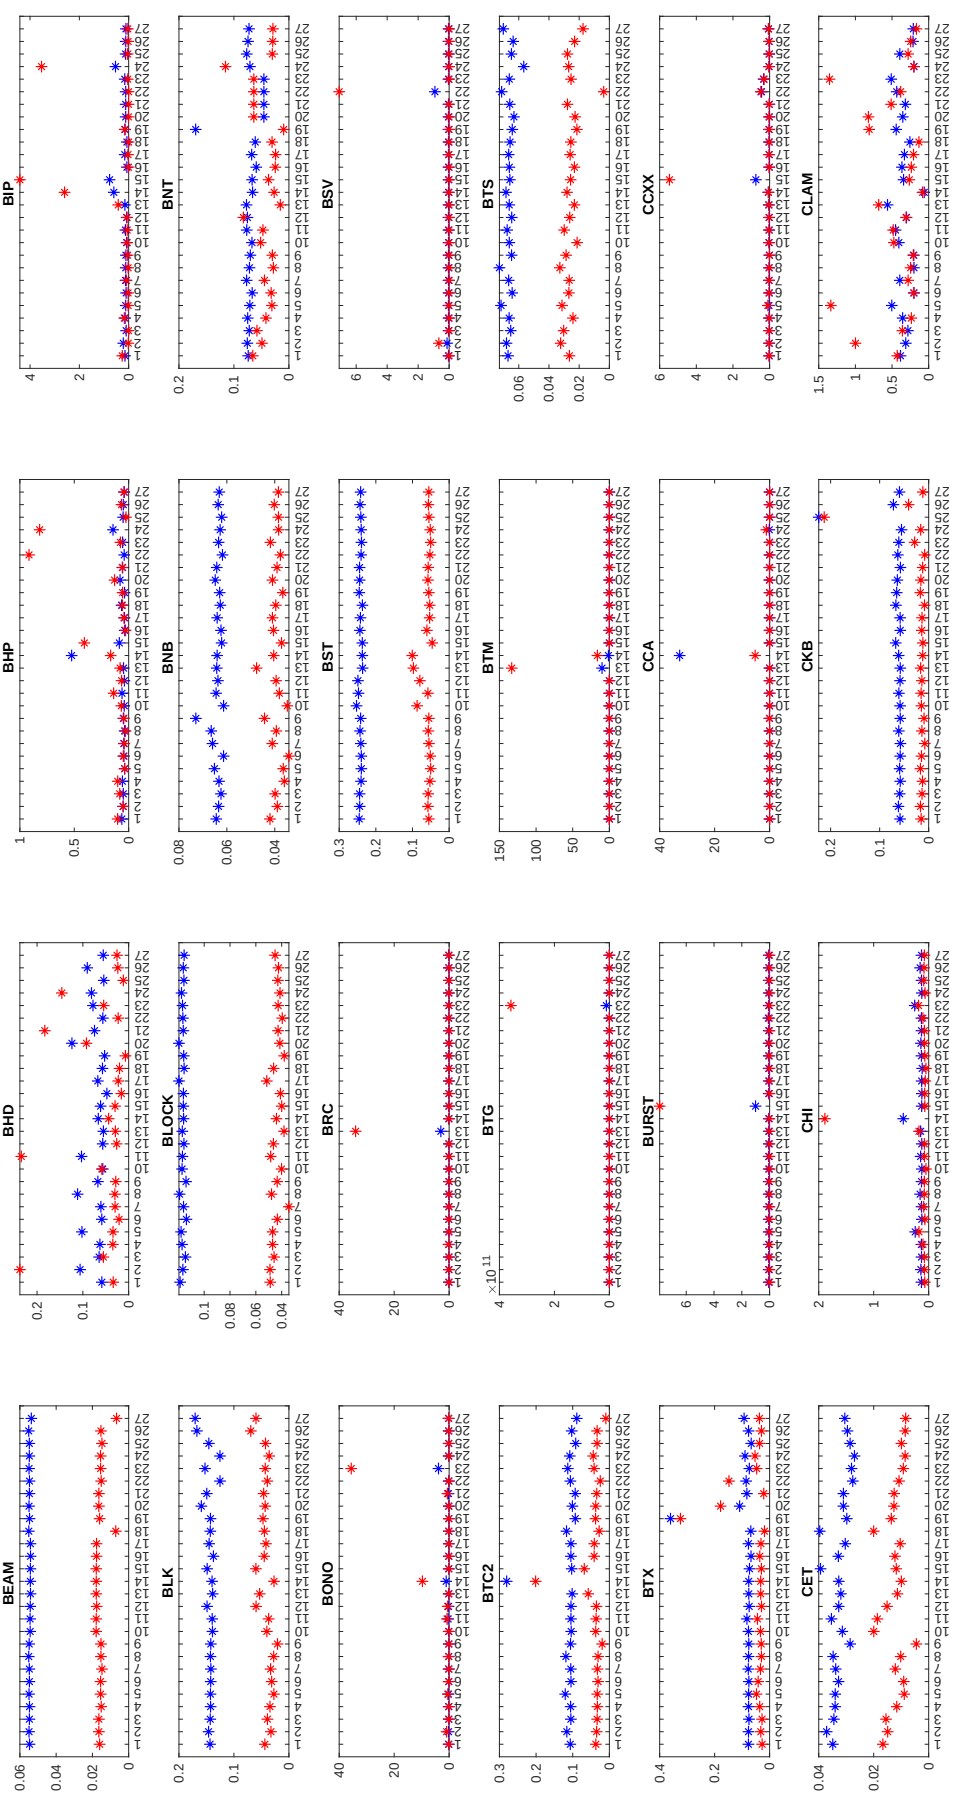
BEAM
BHD
BHP
BIP
BLK
BLOCK
BNB
BNT
BONO
BRC
BST
BSV
BTC2
BTG
BTM
BTS
BTX
BURST
CCA
CCXX
CET
CHI
CKB
CLAM

CNX
COLX
COTJ
CRW
CSC
CTC
CTXC
CURE
CUT
CVC
DAG
DASH
DCN
DCR
DCY
DDK
DERO
DEV
DGB
DGD
DIME
DIVI
DMD
DNT

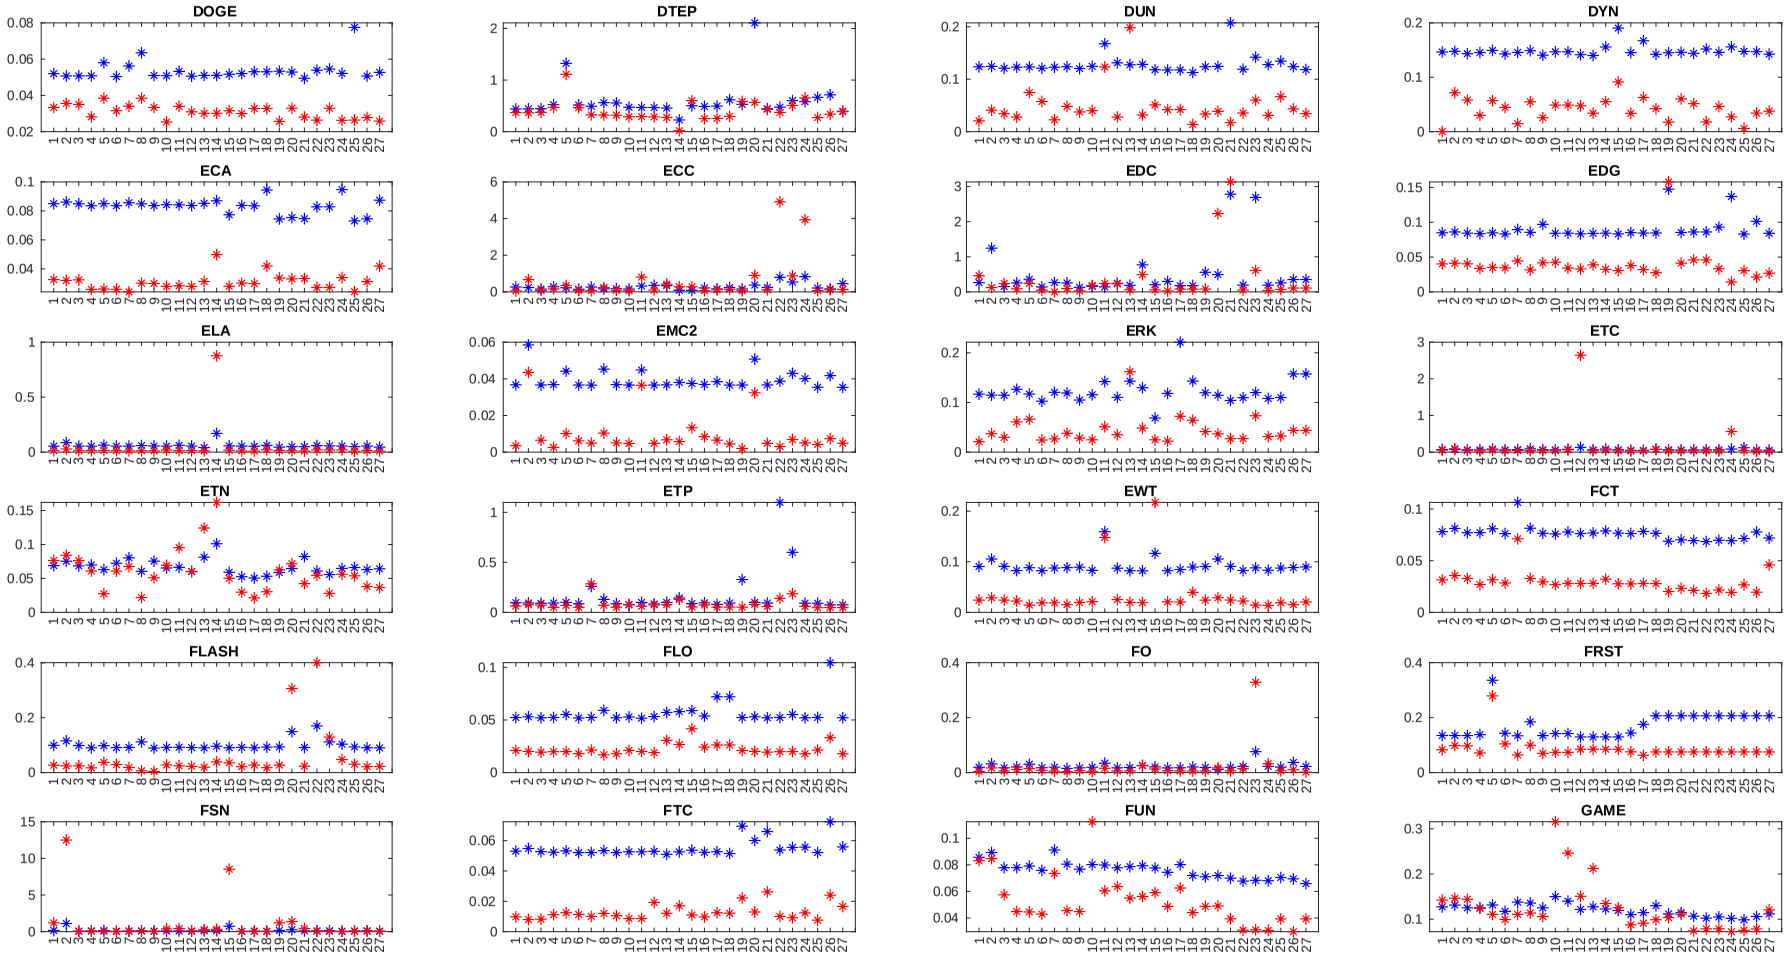
DOGE
DTEP
DUN
DYN
ECA
ECC
EDC
EDG
ELA
EMC2
ERK
ETC
ETN
ETP
EWT
FCT
FLASH
FLO
FO
FRST
FSN
FTC
FUN
GAME

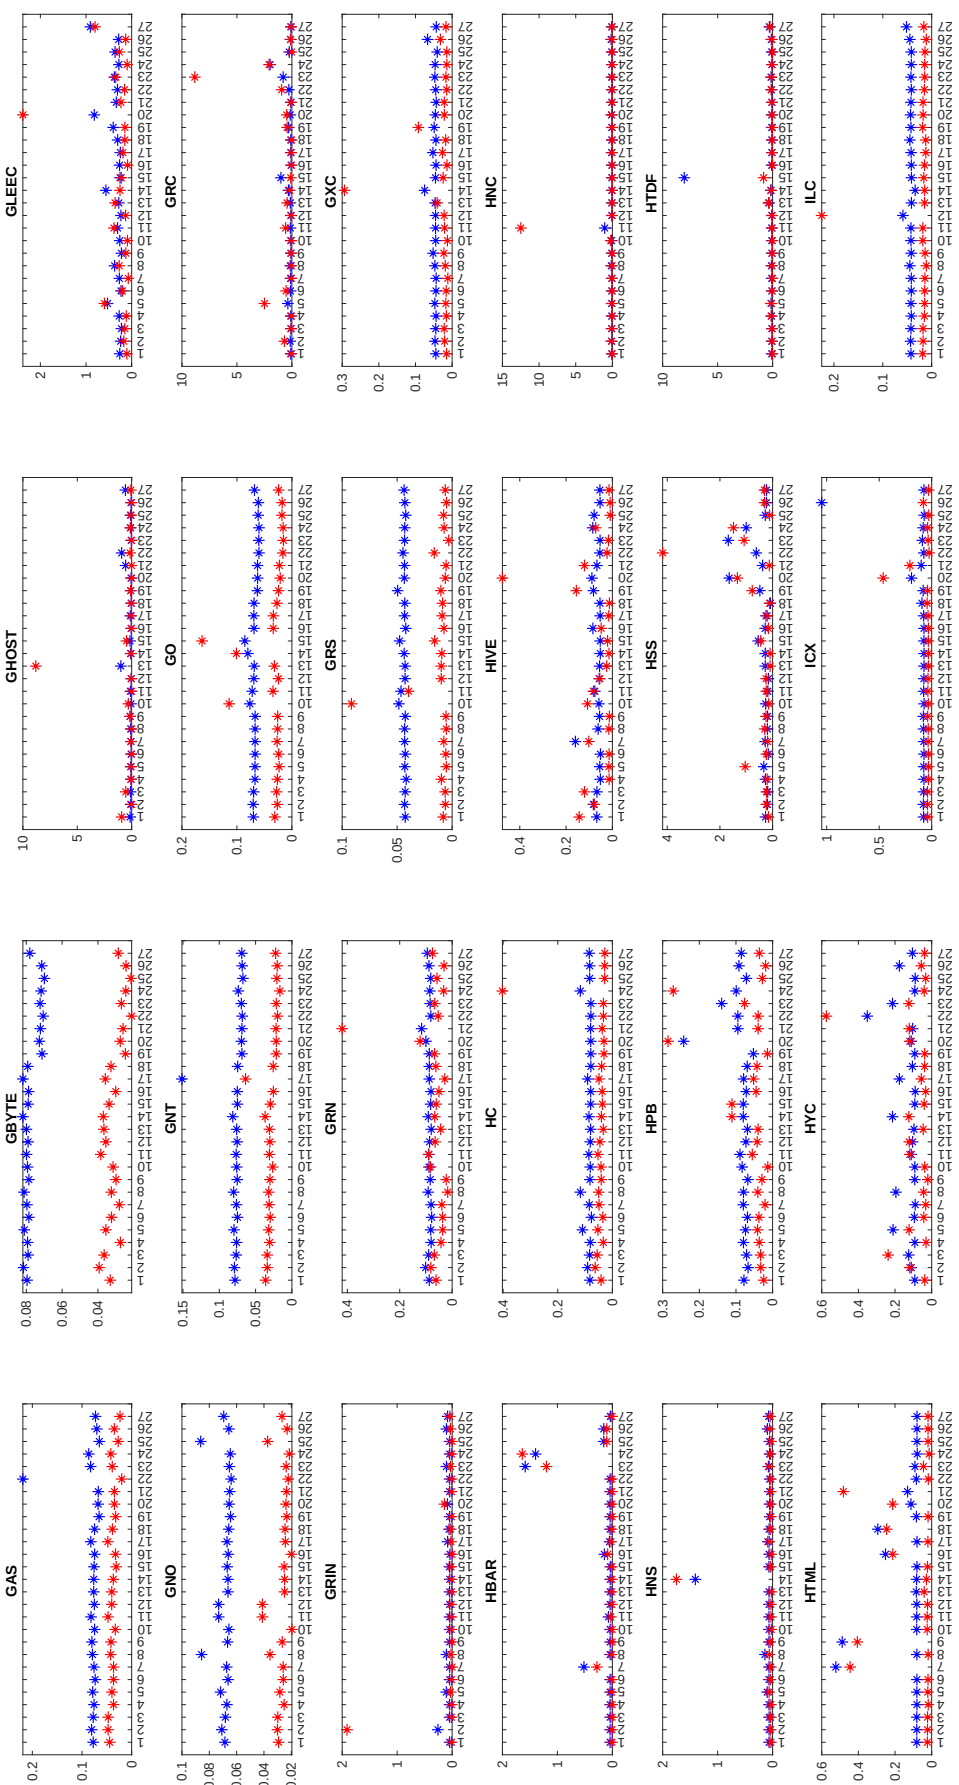
GAS
GBYTE
GHOST
GLEEC
GNO
GNT
GO
GRC
GRIN
GRN
GRS
GXC
HBAR
HC
HIVE
HNC
HNS
HPB
HSS
HTDF
HTML
HYC
ICX
ILC

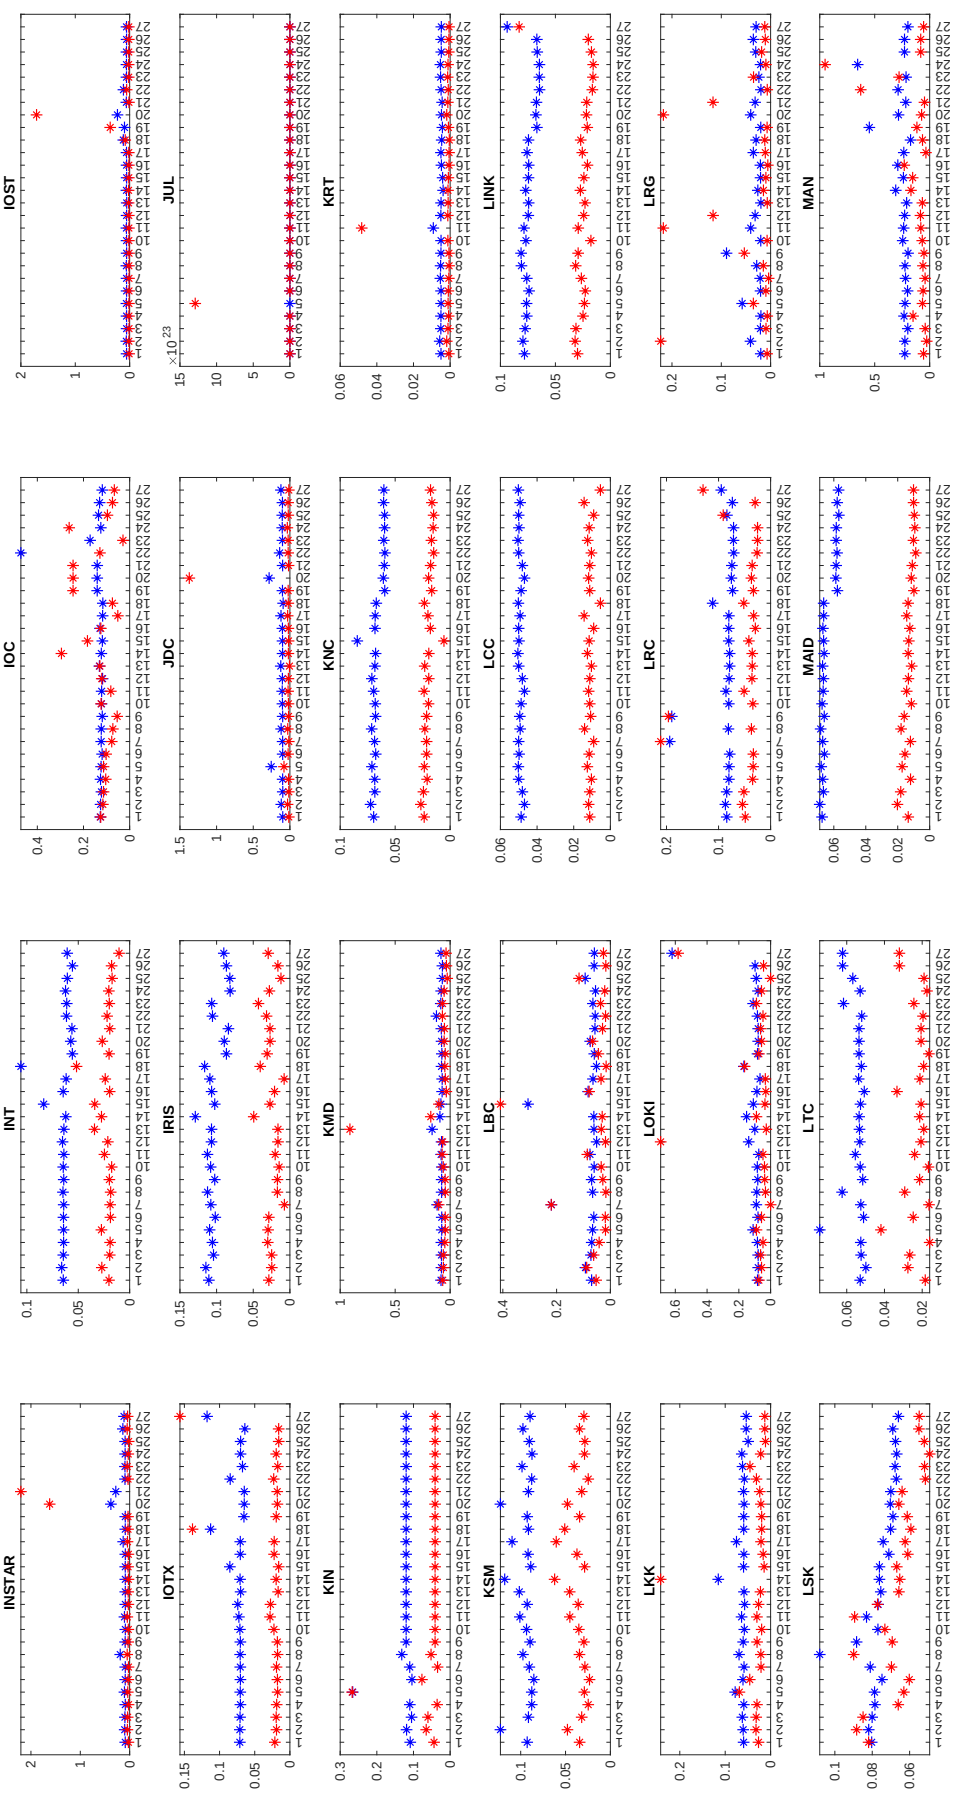
INSTAR
IOST
INT
IOC
IOTX
JUL
IRIS
JDC
KIN
KRT
KMD
KNC
KSM
LINK
LBC
LCC
LKK
LRG
LOKI
LRC
LSK
MAID
LTC
MAN

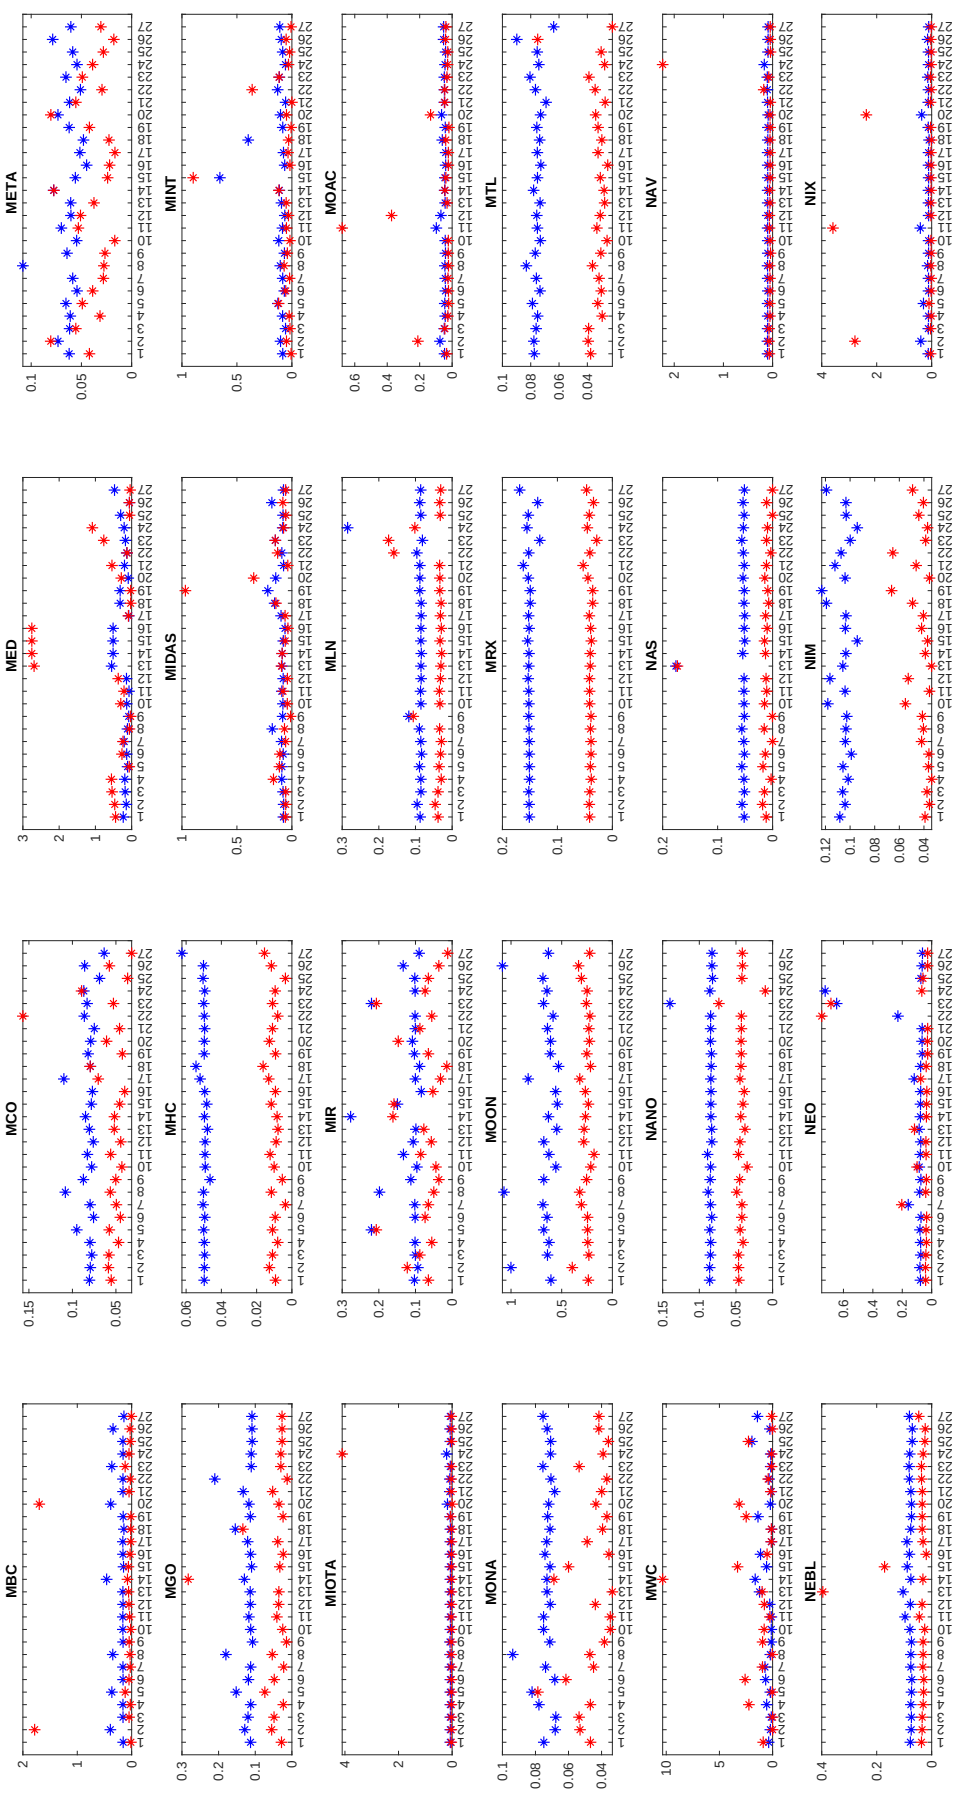
MBC
MCO
MED
META
MGO
MHC
MIDAS
MINT
MIOTA
MIR
MLN
MOAC
MONA
MOON
MRX
MTL
MWC
NANO
NAS
NAV
NEBL
NEO
NIM
NIX

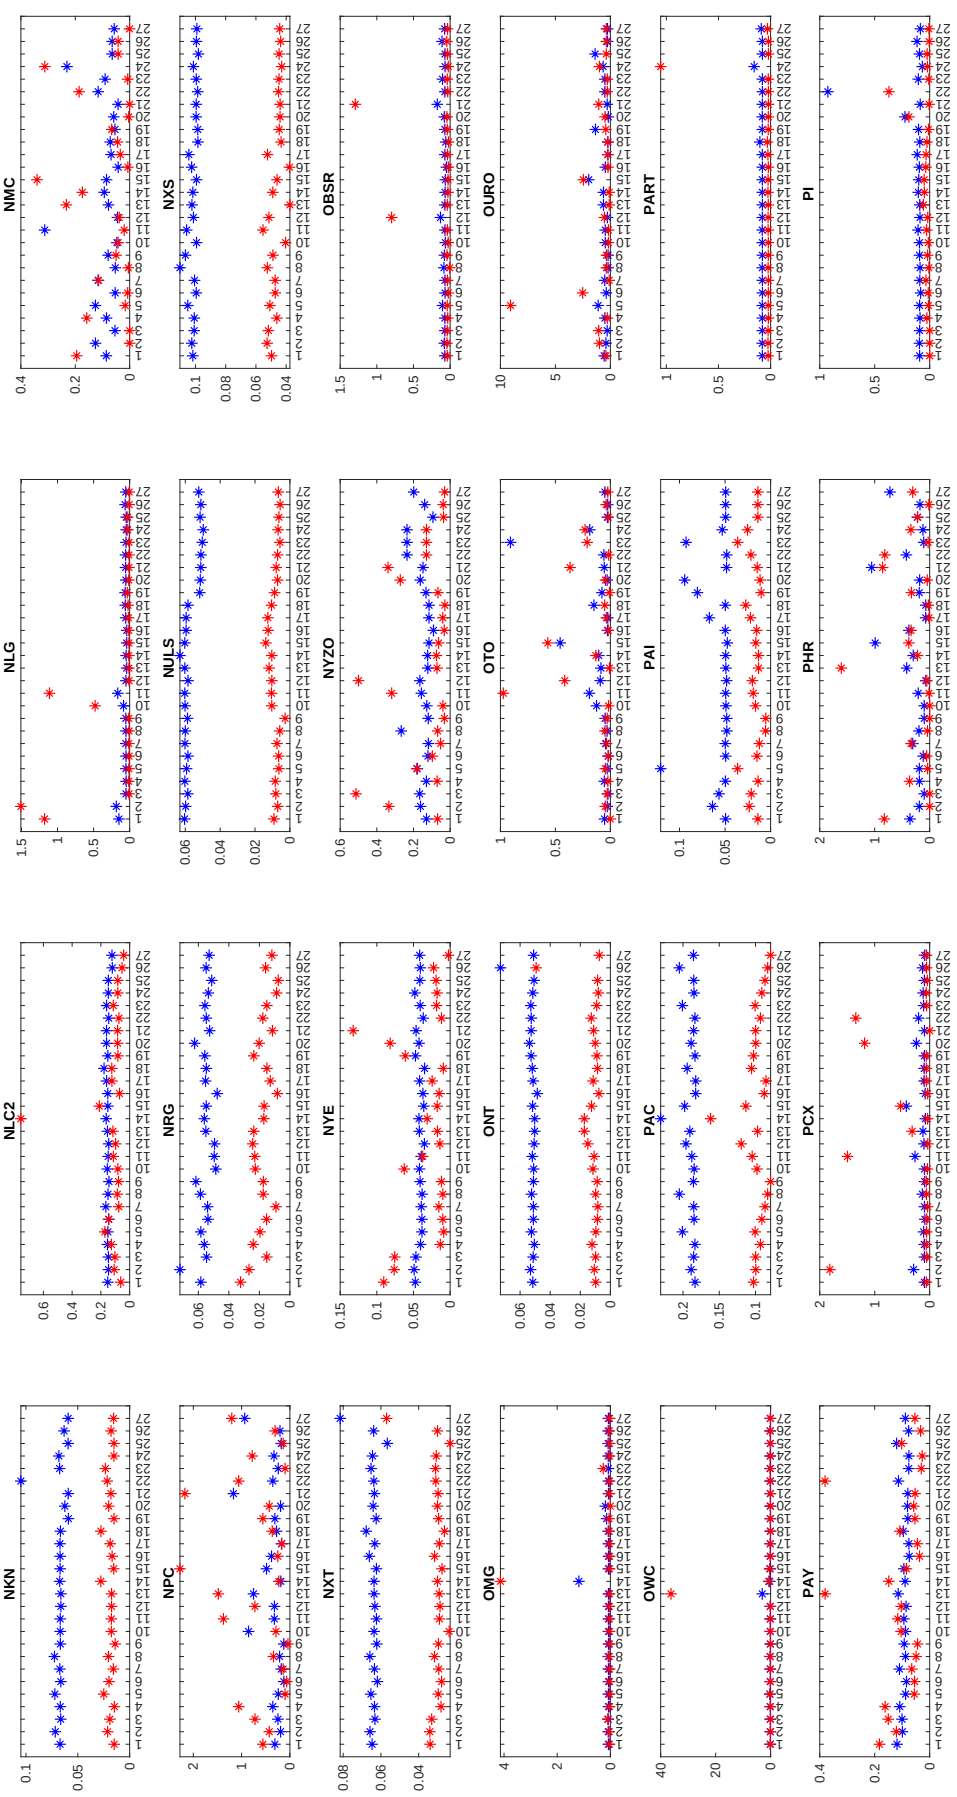
NKN
NLC2
NLG
NMC
NPC
NRG
NULS
NXS
NXT
NYE
NYZO
OBSR
OMG
ONT
OTO
OURO
OWC
PAC
PAI
PART
PAY
PCX
PHR
PI

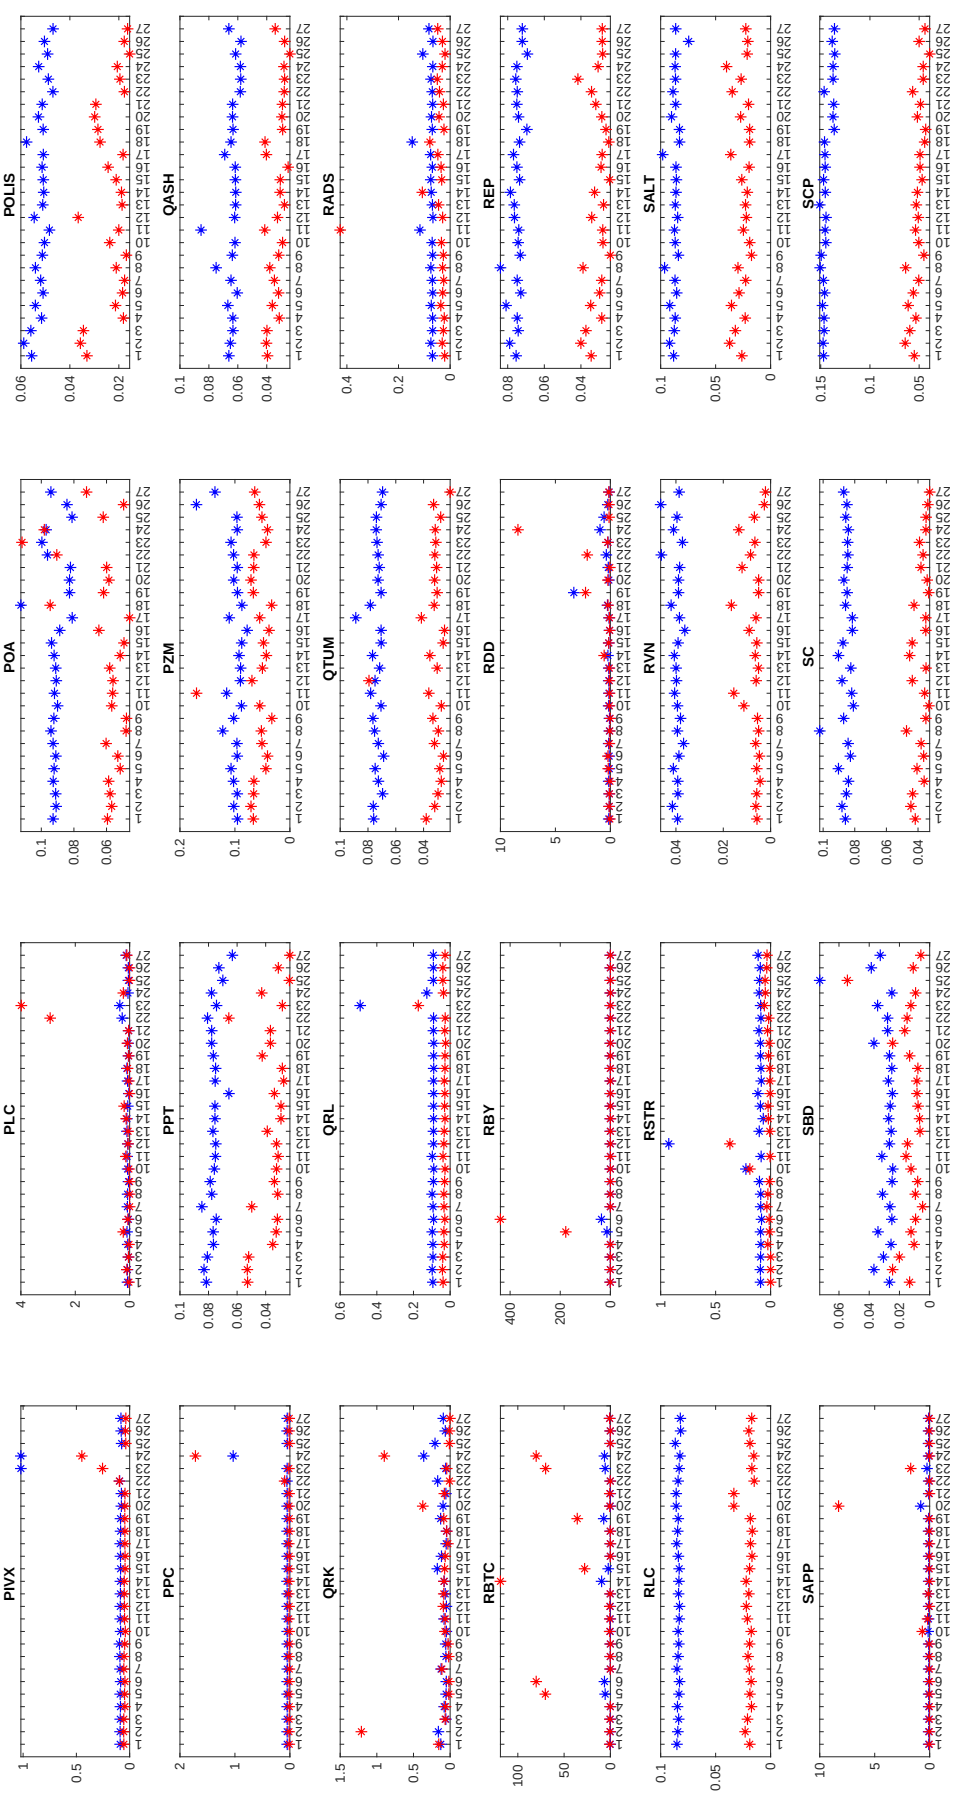
PIVX
PLC
POA
POLIS
PPC
PPT
PZM
QASH
QRK
QRL
QTUM
RADS
RBTC
RBY
RDD
REP
RLC
RSTR
RVN
SALT
SAPP
SBD
SC
SCP

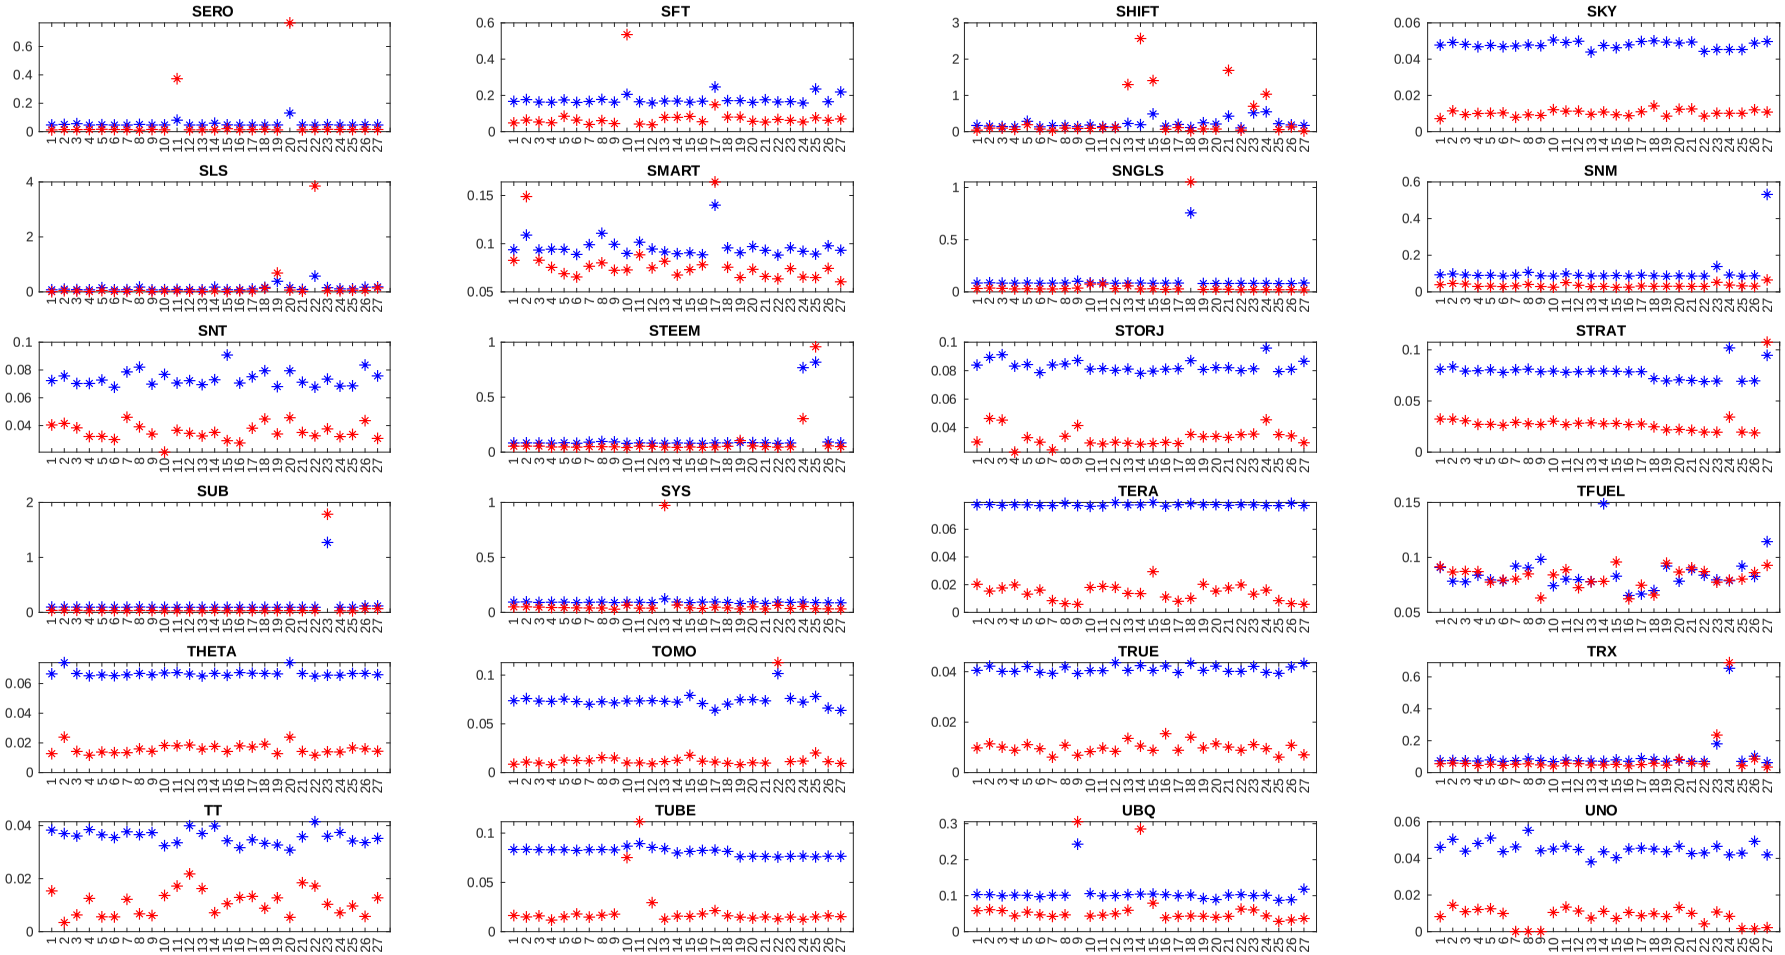
SERO
SFT
SHIFT
SKY
SLS
SMART
SNGLS
SNM
SNT
STEEM
STORJ
STRAT
SUB
SYS
TERA
TFUEL
THETA
TOMO
TRUE
TRX
TT
TUBE
UBQ
UNO

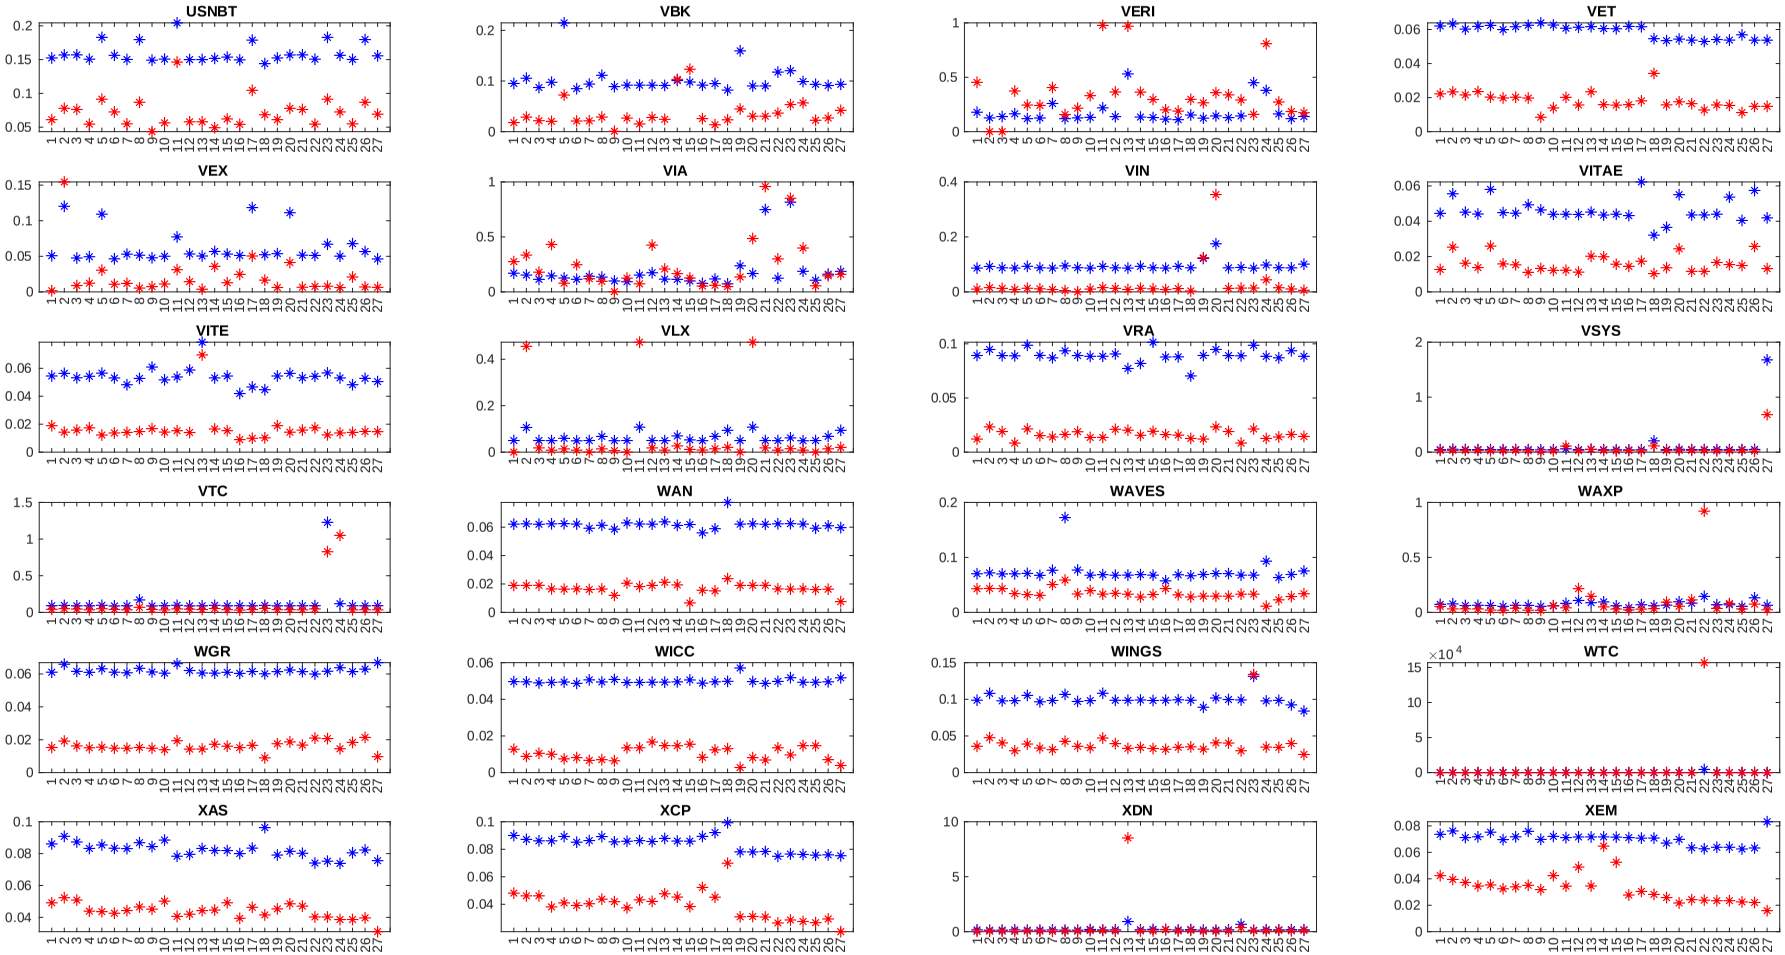
USNBT
VBK
VERI
VET
VEX
VIA
VIN
VITAE
VITE
VLX
VRA
VSYS
VTC
WAN
WAVES
WAXP
WGR
WICC
WINGS
WTC
XAS
XCP
XDN
XEM

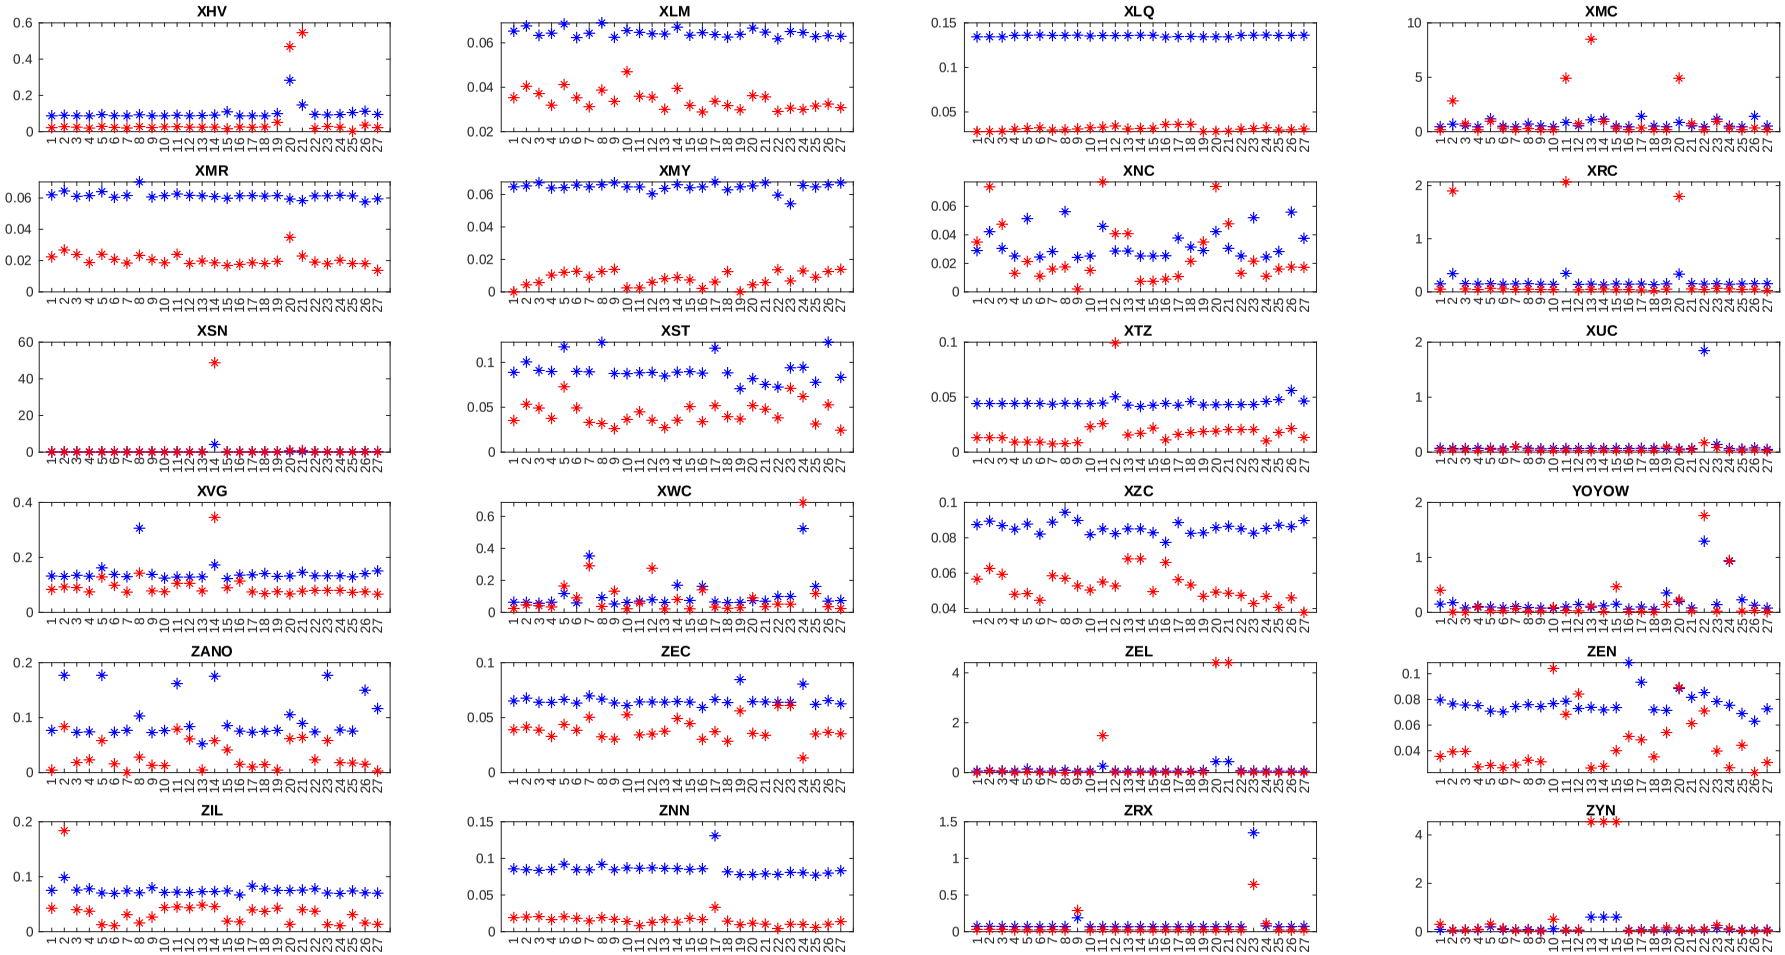
XHV
XLM
XLQ
XMC
XMR
XMY
XNC
XRC
XSN
XST
XTZ
XUC
XVG
XWC
XZC
YOYOW
ZANO
ZEC
ZEL
ZEN
ZIL
ZNN
ZRX
ZYN

Chapter 4

A time-varying Granger-causality analysis on bitcoin returns

Abstract

We examine the impact of alternative drivers of bitcoin returns. We employ a Bayesian time-varying parameter vector autoregressive model and Bayesian variable selection to detect variables with the highest posterior inclusion probability in the equation of bitcoin returns. Using these variables, we reestimate the TVP-VAR model and perform structural analysis, based on impulse response function and forecast error variance decomposition. To examine the causal relationship between bitcoin returns and their drivers, we employ a time-varying Granger-causality test based on the estimates of the TVP-VAR model and heteroscedastic-consistent Granger-causality hypothesis testing. We detect Granger-causality from information demand and cryptocurrency markets to bitcoin returns over the entire sample.

4.1 Introduction

More than ten years after its creation, bitcoin continues to attract the attention of both researchers and investors. Bitcoin has distinct differences from traditional currency as it is not used as a medium of exchange, despite its low transaction costs and peer-to-peer network, nor as a store of value due to its price fluctuations, (Blau, 2017). Furthermore, it can not function as money due to its almost fixed supply (Yermack, 2015). However, bitcoin has features that make it appealing as an asset. Guesmi et al. (2019) provide evidence that bitcoin can function as a risk diversifier and reduce liquidity risk since it is characterised by high average returns and low correlation with traditional financial assets. As an asset, bitcoin is used mainly for speculative purposes (among others, see Baur et al., 2018; Lee et al., 2020). This exploitation as a speculative asset

leads to bubbles (Cheah and Fry, 2015, 2016; Corbet et al., 2018) and extreme volatility (Bakas et al., 2022; Panagiotidis et al., 2022). Furthermore, Fry (2018) demonstrates that liquidity risks can result to boom-bust episodes even in the absence of a bubble and Xiang et al. (2022) argue that the bitcoin market capitalisation series exhibit explosive behaviour.

A strand of the bitcoin related literature focuses on the dynamics that drive bitcoin returns and the causal relationship between bitcoin returns and other factors. The literature identifies numerous factors as bitcoin returns determinants. Macroeconomic and financial variables such as stock market indices (Lahiani et al., 2021), exchange rates (Dyhrberg, 2016), exchange traded funds (Ciner et al., 2022) and economic policy uncertainty receive the most attention by researchers. The attractiveness of bitcoin as an asset is also documented as a driving force of its returns. As an investment asset, bitcoin draws the attention of retail investors who rely mostly on the internet as a source of information. Google searches (Li and Wang, 2017; Fry and Serbera, 2020; Liang et al., 2022; Raza et al., 2022), Wikipedia (Kristoufek, 2013), Twitter (Shen et al., 2019) and Reddit (Rothman, 2019) are shown to have a significant impact on bitcoin returns. Another set of variables that can affect its returns are the bitcoin supply and demand forces (de la Horra et al., 2019). However, supply forces have a smaller impact due to the nearly fixed amount of bitcoins. The nearly constant amount of bitcoins is caused by the difficulty in creating new bitcoins. New bitcoins are generated through a mining process (validating transactions on the network) which requires enormous amounts of electricity and computational power. Chen et al. (2021) examine the ability of several technological factors to forecast bitcoin returns.

Another strand of the literature employs a causality analysis. Lin (2021) employs a VAR model and Granger-causality analysis to investigate the relationship between cryptocurrency returns and Google searches. The causal relationship between bitcoin returns and Google searches is also examined by Dastgir et al. (2019) and Li et al. (2021). Dastgir et al. (2019) employ the Copula-based Granger-causality in the distribution test and find a bidirectional causal relationship between Google Trends search terms and bitcoin returns only in the tails of the distribution. Li et al. (2021) use a wavelet-based quantile Granger-causality test and find a symmetric causal relationship between Google searches and cryptocurrency returns in the short term. The authors do not consider only a single search term but proxy investor attention by constructing an index similar to Da et al. (2011). Shen et al. (2019) use a linear and nonlinear Granger-causality to examine the link between bitcoin returns and the number of tweets from Twitter. Bouri et al. (2018) study daily dependence between global financial stress and bitcoin using standard and quantile-based copula models. Their results

suggest that financial stress causes bitcoin returns at the left and right tail of the latter's conditional distribution. Financial stress, however, has limited directional predictability for Bitcoin. Finally, Ji et al. (2018) use the directed acyclic graph approach to study the interrelation between bitcoin and other financial asset classes. Despite the growing literature on the driving factors of bitcoin returns, the results are often diverse and contradicting and the discussion on the most important determinants continues.

The aim of this essay is to examine the effect of alternative factors on bitcoin returns. Contrary to previous studies, we do not focus only on a specific group of variables (e.g. stock markets or investor sentiment) but consider a wide range of factors from different areas. Specifically, we consider bitcoin volatility, investor sentiment indices, proxies for bitcoin supply and demand, stock market returns and volatility indices, commodities returns, exchange rates and interest rates. Overall, we consider 29 potential determinants of bitcoin returns.

The analysis is conducted on a time-varying framework using a Bayesian time-varying parameter vector autoregressive (TVP-VAR) model. In this framework, we model the evolution of the variance-covariance matrix of the errors using stochastic volatility, thus accounting for the volatility dynamics of bitcoin returns. The latter is often overlooked in the relevant literature. We perform the analysis using a two-step approach. In the first step, we estimate the TVP-VAR using all 29 potential factors (and bitcoin returns) and detect the most important variables of the model based on the Bayesian variable selection method by Korobilis (2013).[1] In the second step, we reestimate the TVP-VAR model using only the variables selected in the previous step and perform the analysis. We perform Granger-causality, impulse response function and forecast error variance decomposition analysis.

To take advantage of the time-varying framework in the Granger-causality analysis, we propose a time-varying Granger-causality approach that allows us to examine the evolution of causal relationship between two variables over time. The proposed methodology combines the estimates from the TVP model with the heteroscedastic-consistent Granger-causality hypothesis testing. Moreover, our approach does not rely on rolling windows or recursive algorithms (as in Shi et al., 2018) that require the estimation of the model on subsamples of the initial sample and thus does not suffer from inference issues caused by small sample sizes (Zapata and Rambaldi, 1997) and temporal aggregation (Otero et al., 2021). In addition, the causality analysis can be performed over the whole sample and does not require an initial window to initiate the algorithm.

[1]Instead of reducing the number of variables, one could employ the procedures of Song and Taamouti (2019) who propose alternative procedures to test for Granger-causality in a big data environment. However, these procedures assume a time invariant framework.

The contribution of this study is twofold. From a methodological perspective, we propose a new time-varying Granger-causality approach that does not depend on rolling-window or recursive evolving algorithms. We assess the performance of this new approach using Monte Carlo simulations. From an empirical perspective, we study the impact of alternative predictors on bitcoin returns. The results, based on the Bayesian variable selection algorithm, suggest that the most important factors that drive bitcoin returns are investor sentiment (proxied by Google searches and uncertainty cryptocurrency index), bitcoin trading volume, Nasdaq and ethereum returns. Using the newly proposed approach, we quantify the causal relationship between the five variables and bitcoin returns. We find that Google searches and ethereum returns Granger-cause bitcoin returns over the whole sample while the rest of the variables are only for specific periods. In addition, the causal relationship from trading volume to bitcoin returns exists only during periods when bitcoin prices remain relatively steady.

The rest of the essay proceeds as follows. Section 2 describes the methodology and the new time varying Granger-causality test. Section 3 discusses the Monte Carlo simulations. Section 4 presents the data used in the empirical analysis. Section 5 discusses the main findings. Section 6 compares the empirical results with the results from a recursive algorithm. Section 7 concludes and discusses the opportunities for further research.

4.2 Methodology

4.2.1 The TVP-VAR model

Let y_t be a vector of n time series, $t = 0, \ldots, T-1$. The VAR(p) model is written as:

$$y_t = b_t + \sum_{i=1}^{p} B_{i,t} y_{t-i} + \varepsilon_t, \tag{4.1}$$

where b_t is a $n \times 1$ vector of time-varying coefficients, $B_{i,t}$, $i = 1, \ldots, p$ are $n \times n$ matrices of time-varying coefficients and $\varepsilon_t \sim N(0, \Sigma_t)$. The TVP-VAR is a state space model, Koop and Korobilis (2010). Specifically, we define $X_t = I_n \otimes [1, y'_{t-1}, \ldots, y'_{t-p}]$ (where I_n is the $n \times n$ unit matrix) and $\beta_t = \text{vec}([c_t, (B_{1,t}), \ldots, (B_{p,t})]')$ (where vec() is the column-wise vectorisation of a matrix) and rewrite equation (4.1), the measurement equation, in block form as follows:

$$y_t = X_t \beta_t + \varepsilon_t. \tag{4.2}$$

Since the variance-covariance matrix of the error terms are allowed to smoothly

evolve over time, we model the evolution of Σ_t as in Cogley and Sargent (2005). Specifically, we consider the following decomposition of Σ_t:

$$\Sigma_t^{-1} = A' D_t^{-1} A,$$

where A_t is a $n \times n$ lower triangular matrix with diagonal elements equal to one and D_t is a $n \times n$ diagonal matrix with diagonal elements, $e^{h_{i,t}}$. Contrary to Primiceri (2005), we assume that A is a time invariant matrix. If we denote by α the vector of unrestricted elements of A, the elements below the main diagonal, and $h_t = (h_{1,t}, \ldots, h_{n,t})'$, we complete the state space model by obtaining the following state equations:

$$\begin{aligned} \beta_t &= \beta_{t-1} + u_t, \quad u_t \sim N(0, Q) \\ h_{i,t} &= h_{i,t-1} + u_{i,t}^h, \quad u_t \sim N(0, \sigma_{h,i}^2), \end{aligned} \tag{4.3}$$

where the initial conditions β_0 and $h_{i,0}$ are treated as unknown parameters and Q is assumed to be diagonal ($Q = \text{diag}(q_i)$, where $i = 1, \ldots, n + n^2 p$).

We complete the model by specifying the prior distributions for β_0, q_i, α, $h_{i,0}$. We consider the following independent priors:

$$\begin{aligned} \beta_0 &\sim N(a_0, B_0), \\ q_i &\sim IG(v_{0,q_i}, S_{0,q_i}), \\ \alpha &\sim N(\alpha_0, V_\alpha), \\ \sigma_{h,i}^2 &\sim IG(v_{0,h_i}, S_{0,h_i}), \\ h_0 &\sim N(s_0, S_0), \end{aligned}$$

where IG denotes the inverse gamma distribution. We consider alternative set of prior values to assess the validity of our findings. In all cases, the results remain similar to the ones presented here. The five model parameters and the two states are estimated using a Gibbs sampler described in Chan et al. (2019).

Once the model is estimated, we employ the Bayesian variable selection, proposed by Korobilis (2013), to select the most important variables in the model. We then reestimate the TVP-VAR model using using only the most important variables. In the new model, we employ the time-varying Granger-causality test as well as the impulse response function and forecast error variance decomposition analysis.

4.2.2 The time varying Granger-causality test

To describe the Granger-causality test, it is useful to define $Z_t = [1, y'_{t-1}, \ldots, y'_{t-p}]'$ and $R_{i,j}$ an appropriate coefficient restriction matrix for the null hypothesis that y_i does not Granger-cause y_j. Next, we can calculate

$$\Omega = T^{-1} \sum_{t=1}^{T-p} (\varepsilon_t \otimes Z_t)(\varepsilon_t \otimes Z_t)',$$
$$F = T^{-1} \sum_{t=1}^{T-p} Z_t Z'_t,$$
$$G = I_n \otimes F.$$

We estimate a series of heteroscedastic-consistent Wald statistics $W_{i,j,t}$, $t = p + 1, \ldots, T$ as follows:

$$W_{i,j,t} = T(R_{i,j}\beta'_t)'(R_{i,j}G^{-1}\Omega G^{-1}R'_{i,j})^{-1}(R_{i,j}\beta'_t). \tag{4.4}$$

For each t, $W_{i,j,t}$ follows a χ^2_q, where the degrees of freedom, q, is the number of restricted coefficients.

4.3 Monte Carlo simulations

4.3.1 Design of the Monte Carlo simulations

To assess the performance of the proposed time varying Granger-causality test, we perform several Monte Carlo experiments. We consider the following general data generating process (DGP):

$$\begin{aligned} y_{1,t} &= b_{10} + b_{11}y_{1,t-1} + b_{12}y_{2,t-1} + \varepsilon_{1,t} \\ y_{2,t} &= b_{20} + b_{21}y_{1,t-1} + b_{22}y_{2,t-1} + \varepsilon_{2,t} \end{aligned} \tag{4.5}$$

The DGP describes a bivariate VAR(1) model with a constant term. In all cases, $\varepsilon_{1,t}$ and $\varepsilon_{2,t}$ follow a normal distribution with zero mean and unit variance. Three alternative sample sizes are considered, $T \in \{100, 200, 400\}$. We examine the size and power properties of the test and the performance of the test in the presence of a structural break. To examine the size (incorrectly reject the null of no Granger-causality) properties of the test, we set $b_{21} = b_{12} = 0$, so that there is no Granger-causality between y_1 and y_2. The values of the rest of the parameters are all different from zero

and described in Table 4.1. To examine the power of the test, we use the same set of parameters as in the previous case but set $b_{12} = b_{21} = 0.4$. We consider a final case, where we there is a structural break in the parameters. Specifically, for the first half of the sample we set $b_{12} = 0$ and for the second half of the sample we set $b_{12} = 0.4$. We do so to validate the ability of the test to detect Granger-causality in specific subsamples. Figure 4.1 shows an example of two time series simulated using DGP3 and $T = 400$. Table 4.1 summarises the details for all DGPs.

Table 4.1: Parameter values for the Monte Carlo simulations.

DGP/Parameter	b_{10}	b_{11}	b_{12}	b_{20}	b_{21}	b_{22}	
DGP1	1	-0.3	0.0	1	0.0	0.6	
DGP2	1	-0.3	0.2	1	0.4	0.6	
DGP3	1	-0.3	0.2	1	0.0	0.6	t = 1,..., T/2
	1	-0.3	0.2	1	0.4	0.6	t = T/2+1,..., T

Notes: i) DGP1 and DGP2 are used to examine the size and power properties of the test, respectively. In DGP3 Granger-causality exists only in the second half of the sample. ii) We consider three sample sizes, $T = 100, 200, 400$. iii) Both error terms are independently and identically distributed following the normal distribution with zero mean and unit variance.

Figure 4.1: Example of simulated time series using DGP3. For $t = 1, \ldots, 200$, there is a unidirectional Granger-causal relationship from y_2 to y_1. For $t = 201, \ldots, 400$ there is a bidirectional causal relationship.

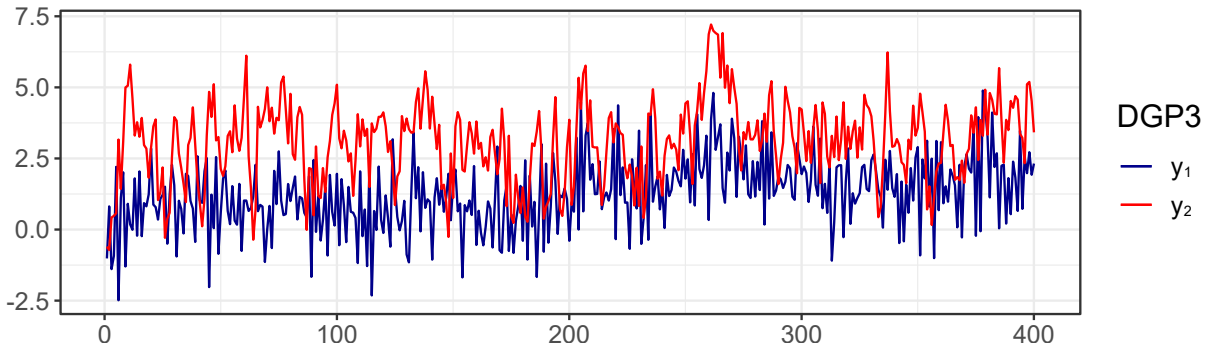

The TVP-VAR model is fitted to the generated time series. To estimate the posterior distributions, we use 20.000 iterations of the Gibbs sampler algorithm. The first half of the iterations are used as burn-in sample and the remaining are used to build posterior sample. The results of the Monte Carlo simulations are based on 10.000 iterations.

4.3.2 Results

For all DGPs, we consider the unidirectional Wald statistics, both from y_1 to y_2 and vice versa. First, we evaluate the size properties of the test using DGP1. For the test to be successful, we require that it accepts the null hypothesis of no Granger-causality, at the 5% significance level, for every observation in the sample. The left part of Table 4.2 presents the failure frequency of the test. We observe that even for smaller samples (i.e. $T = 100$), the test handles size reasonably well.

We use DGP2 to assess the power of the proposed test. We assume that the test successfully detects Granger-causality if it rejects the null hypothesis for the entire sample. The middle part of Table 4.2 reports the success rate of the test. Regardless of the sample size, the test is able to detect the causal relationship more than 90% of the time.

The last Monte Carlo experiment, DGP3, is used to examine the behaviour of the test in the presence of a structural break. Specifically, we investigate whether the test detects Granger-causality from y_1 to y_2 during a specific subsample. The fifth column of Table 4.2 shows the number of times that the test rejects the null hypothesis only during the latter half of the sample. However, we consider a period of $T/10$ observations before and after the break point, where the test is allowed to produce a false result. The results suggest that the new test successfully identifies the change in causal relationship in more than 80% of the cases. The success rate increases as the number of observations increases. Furthermore, we notice that change in b_{21} does not affect, the ability of the test to reject the null hypothesis of no Granger-causality from y_2 to y_1 for all observations in the sample. In particular, the last column of Table 4.2 indicates that the test indicates that y_2 Granger-causes y_1 for the entire sample, in more than 90% of the cases. This finding is, validates the results regarding the power of the test (obtained using DGP2).

4.4 Data

The empirical study utilises weekly data over the period 2015-2023. We download bitcoin closing prices from coinmarketcap.com and calculate the bitcoin returns as the logarithmic differences of the closing prices. In addition, we calculate bitcoin volatility as the square of the returns (Yousaf and Yarovaya, 2022).

We consider three indices regarding investor sentiment. The first is created using Search Volume Index (SVI) data, from Google Trends, for three search terms, 'bitcoin', 'bitcoin price' and 'BTC'. Alternative queries such as 'bitcoin returns', 'bitcoin

Table 4.2: Power and size properties of the proposed time-varying Granger-Causality methodology.

	DGP1		DGP2		DGP3	
T	$W_{1,2}$	$W_{2,1}$	$W_{1,2}$	$W_{2,1}$	$W_{1,2}$	$W_{2,1}$
100	0.143	0.164	0.915	0.940	0.802	0.903
200	0.065	0.089	0.953	0.982	0.816	0.947
400	0.030	0.034	0.995	1.000	0.853	0.961

Notes: i) DGP1 and DGP2 are used to examine the size and power properties of the test, respectively. In DGP3 Granger-causality exists only in the second half of the sample.

invest' or 'BTCUSD' are not used since the search volume for these queries is negligible compared to the search volume of the keywords that are used in the study. The three downloaded SVIs take values from 0 to 100 and are scaled accordingly (meaning that at each date, the SVIs are comparable). Furthermore, the three indices maximise on the same date (on 30/08/2019). Using the maximum as values as weights (rescaled so their sum equals 1), we define the Google Search Volume Index (GSVI) as the weighted average of the three indices.[2] We use the GSVI as a measure of attention for retail investors. We assume that Google searches quantify information demand from individual investors given that it is the most popular search engine. The other two indices of investor attention, are the Uncertainty Cryptocurrency (UCRY) policy and price indices, respectively. The two indices, proposed by Lucey et al. (2022), can be used to gauge, among other things, uncertainty around bitcoin stemming from policy makers and institutional investors. We opt not to include the Central Bank Digital Currency (CBDC) attention and uncertainty indices since CBCB is more related to stablecoins rather than bitcoin.

The next group of variables aims to capture supply and demand forces. We measure bitcoin supply using the number of total bitcoins in circulation. We proxy bitcoin demand using the trading volume and a number of network related measures, the number of confirmed transactions, the number of unique addresses, hash rate and network difficulty. The two latter variables are proxies of mining difficulty. Hash rate measures the computational power in the network. A higher hash rate indicates a more secure network and more users validating transactions. In addition, the higher the hash rate, the higher the chance that a user creates a new block in the blockchain and obtains bitcoin. Network difficulty measures the computational power required to

[2]Due to the small number of SVIs, other methods such as Bayesian averaging, principal components analysis or factor models.

verify a transaction and is adjusted, based on hash rate and the total number of users, to maintain the time required to verify a transaction of approximately ten minutes.

To analyse the relationship between stock markets and bitcoin returns, we employ the returns of five stock market indices, the S&P500, Dow Jones, Nasdaq, EURO STOXX 100, FTSE50, Nikkei 225 and the Shanghai composite index. We also include two stock market volatility indices, the CBOE market and the EURO STOXX 50 volatility indices. Regarding commodity markets, we consider two oil price indices, WTI and Brent and the price of gold. Finally, we examine the relationship between bitcoin returns and four exchange rates: euro, British sterling, Japanese yen and Chinese yuan. The exchange rates are expressed as the price of domestic currency to one US dollar, implying that an increase in the exchange rate denotes an appreciation.

All variables are obtained in weekly frequency from their respective data sources. Stationarity of the time series is examined through the ADF and the necessary transformations are performed. Figure D1 presents all variables used in the analysis in levels. Table 3.1 lists the variables and the data sources.

4.5 Empirical results

In what follows, for each TVP-VAR model, we consider a burn-in sample of 50,000 iterations and build the posterior distributions from the next 50,000 draws by keeping only each 10^th draw. This thinning process is used in order to mitigate auto-correlation among draws. The number of lags is based on the Schwarz information criterion and set equal to 1 in all cases. Finally, the results of the Granger-causality test are based on the 5% significance level

We estimate the TVP-VAR model using all variables presented in Section 4.4 and use the Bayesian variable selection algorithm to detect the most important regressors in the equation where the bitcoin returns are the dependent variable. Table 4.3 presents the posterior inclusion probability for each variable in the bitcoin equation. We find that apart from the lag of bitcoin returns, the variables with the greatest posterior inclusion probability are the cryptocurrency policy uncertainty index and the GSVI, two variables that proxy investor sentiment. The finding is supported by evidence from previous studies that conclude that information demand is the most important determinant of bitcoin returns, (see Panagiotidis et al., 2018). Out of the ten stock and commodity markets, only one, the Nasdaq index, has a sufficient high posterior inclusion probability. This result reinforces the argument of Liu and Tsyvinski (2020) and Liu et al. (2022) who argue that cryptocurrency returns have low exposure to traditional asset markets. Other variables with high posterior inclusion probability are the bit-

coin trading volume and ethereum returns, indicating that the behaviour of bitcoin returns is better explained by cryptocurrency factors (Liu and Tsyvinski, 2020).

Table 4.3: Posterior inclusion probability of the coefficients for the bitcoin equation in the TVP-VAR model.

BTC : 0.952	**TV : 0.672**	EURO: 0.436	SP500 : 0.590	SSEC : 0.296
VOL : 0.046	ADD : 0.376	GBP : 0.440	DJ : 0.466	VIX : 0.000
GSVI : 0.786	CTR : 0.000	YUAN: 0.236	**NASDAQ: 0.748**	VSTOXX: 0.000
UCRY policy: 0.949	HASH: 0.118	YEN : 0.010	STOXX : 0.402	WTI : 0.070
UCRY price : 0.553	DIF : 0.254	FFER: 0.288	FTSE : 0.448	BRENT : 0.146
ETH : 0.834	NBC : 0.196	ECB : 0.144	NIKKEI: 0.462	GOLD : 0.458

Notes: i) The values denote the posterior inclusion probability and are obtained using the Bayesian variable selection method proposed by Korobilis (2013). ii) The results refer to the bitcoin equation. iii) The five variables with the highest probability, highlight in bold, are considered for the main model. iv) See Table D1 for the definitions of the mnemonics.

Using these five variables and bitcoin returns (i.e. cryptocurrency policy uncertainty index, the GSVI, the bitcoin trading volume, the returns of Nasdaq index, the ethereum and bitcoin returns) we estimate the main TVP-VAR(1) model and perform the time-varying Granger-causality test, impulse response function analysis and forecast error variance decomposition.

We use the time-varying Granger-causality test, described in the section 4.2.2, to examine whether there is a unidirectional causal relationship between each variable and bitcoin returns. Figure 4.2 presents the sequences of Wald statistics for each variable, along with the critical value at the 5% significance level. The null hypothesis of no Granger-causality is rejected during the periods that the Wald statistic exceeds the critical value. We observe that the GSVI and ethereum returns Granger-cause bitcoin returns for the entire sample. The results indicate that the other three variables (bitcoin trading volume, cryptocurrency policy uncertainty and nasdaq returns) affect bitcoin returns only for specific periods of time. All three variables have a significant effect during the first part of the sampling period (approximately up to 2017). We observe that the policy uncertainty index becomes statistically significant during the third quarter of 2017, the period of the first major rise in the price of bitcoin. The index remains significant for a year, until 2018-10-26., long after bitcoin prices have returned to a steady level. After that period, the index becomes significant only during the last quarter of the sample. Regarding trading volume, we detect four periods where the null of Granger-causality is rejected. These are the first two years of the sample (up to 2017-09-15), during 2018, during the second quarter of 2019 and during the first three quarters of 2022. During these periods, bitcoin price was relatively constant, indicating that trading volume is better used to analyse the behaviour of bitcoin

during periods when no bursts occurs. Furthermore, if we consider a lighter confidence level (i.e. 10%), the null hypothesis is rejected for the majority of the sampling period, further supporting the results of Yarovaya and Zięba (2022) who show a causal relationship between trading volume and cryptocurrency returns. Finally, the Wald statistics for Nasdaq returns exceed the critical value from the start of the sample up to the third quarter of 2018.[3] Nasdaq returns also Granger-cause bitcoin returns from June 2021 until the end of the sample. That is, there is no causal relationship between Nasdaq and bitcoin returns during the second bubble burst at the end of 2020. However, there is a relationship during the most recent bubble burst which occurred at the second half of 2021.

Figure 4.3 presents the median impulse responses of bitcoin returns to a shock in each variable along with 95% credible set. Since the variance-covariance matrix of the residuals varies over time, we compute the impulse response functions using the median of all variance-covariance matrices. In all cases, the effect of each shock is short lasting, meaning that after four periods (weeks) bitcoin returns to pre-shock levels. The cryptocurrency policy uncertainty index is the only variable that does not have a significant effect on bitcoin returns. In addition, it is the only variable that causes a decrease in bitcoin returns. Considering the second sentiment index we observe that the positive shock in the GSVI yields a rise in bitcoin returns. Similarly, the results suggest that an unexpected increase in bitcoin trading volume leads to a rise in bitcoin returns. Furthermore, bitcoin returns respond positively to shocks both in cryptocurrency (i.e. ethereum) and stock (i.e. Nasdaq) market returns. In particular, a shock in the Nasdaq index yields to rise in bitcoin returns that lasts for two weeks. After that period, the response of bitcoin returns becomes statistically insignificant.

[3]There are a few exceptions during this period but only last for up to two consecutive observations.

Figure 4.2: Time varying Granger-causality results.

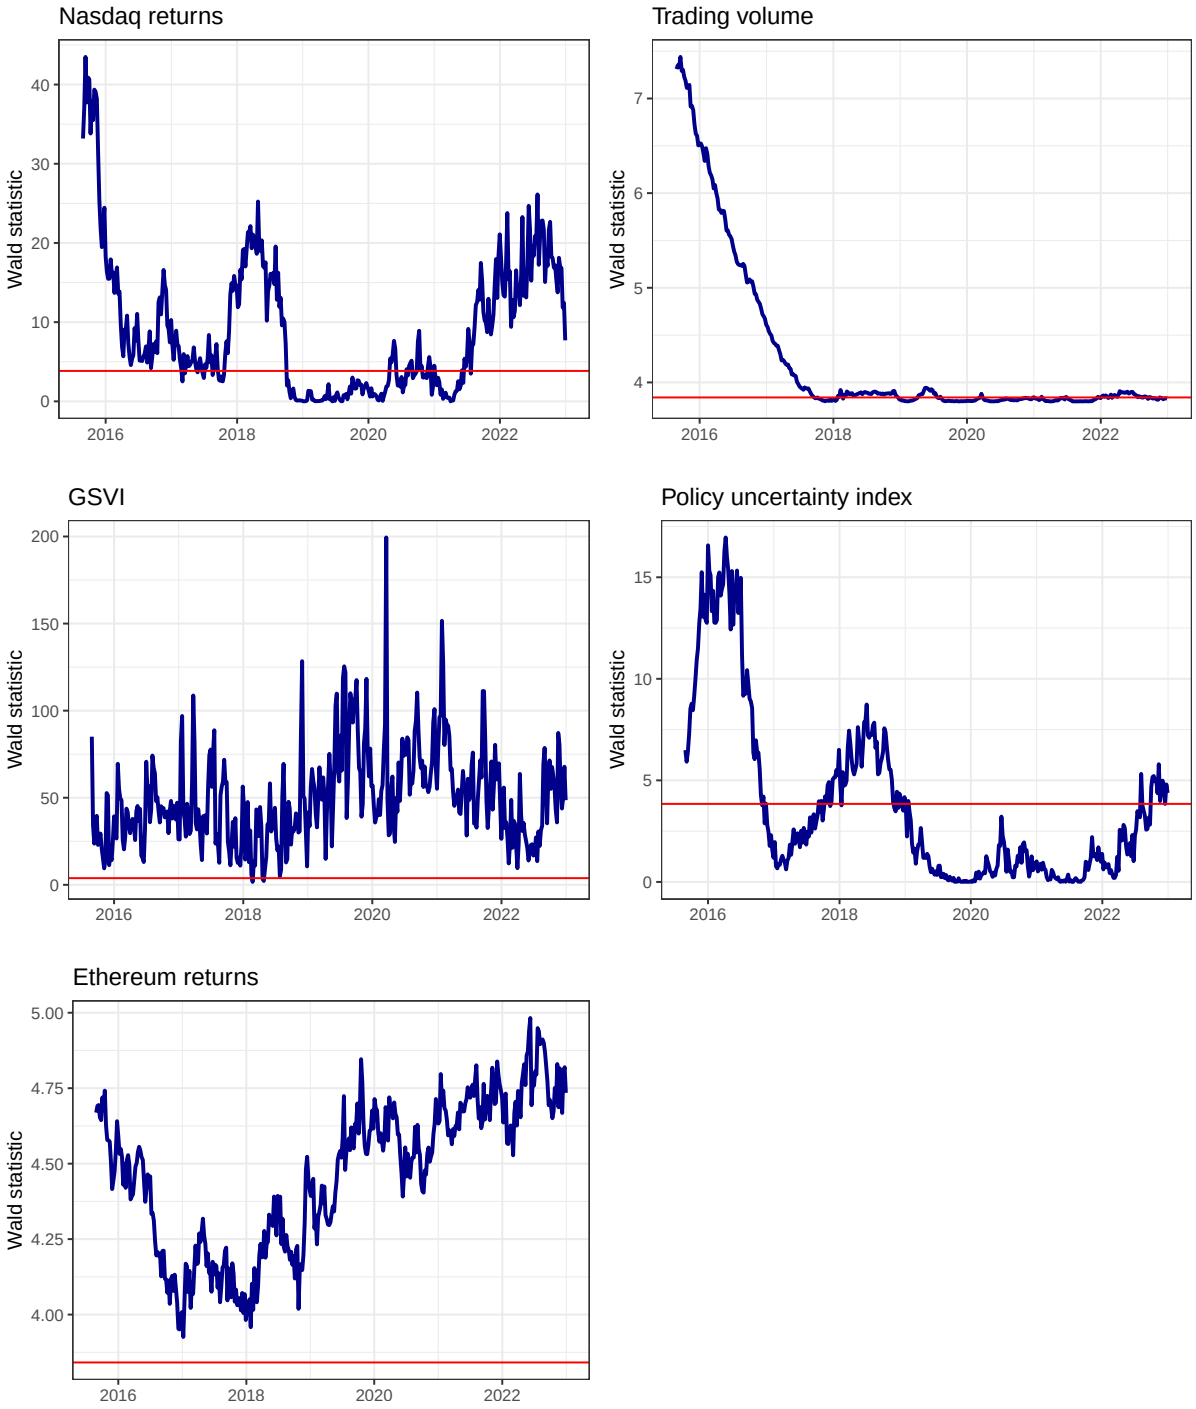

Notes:) i) The blue lines present the evolution of the Wald statistic during the sampling period. The red line presents the critical value at the 5% significance level. ii) The null hypothesis of Granger non-causality is rejected if the Wald statistic is greater than the critical value.

Figure 4.3: Impulse responses of bitcoin returns to a shock in each variable.

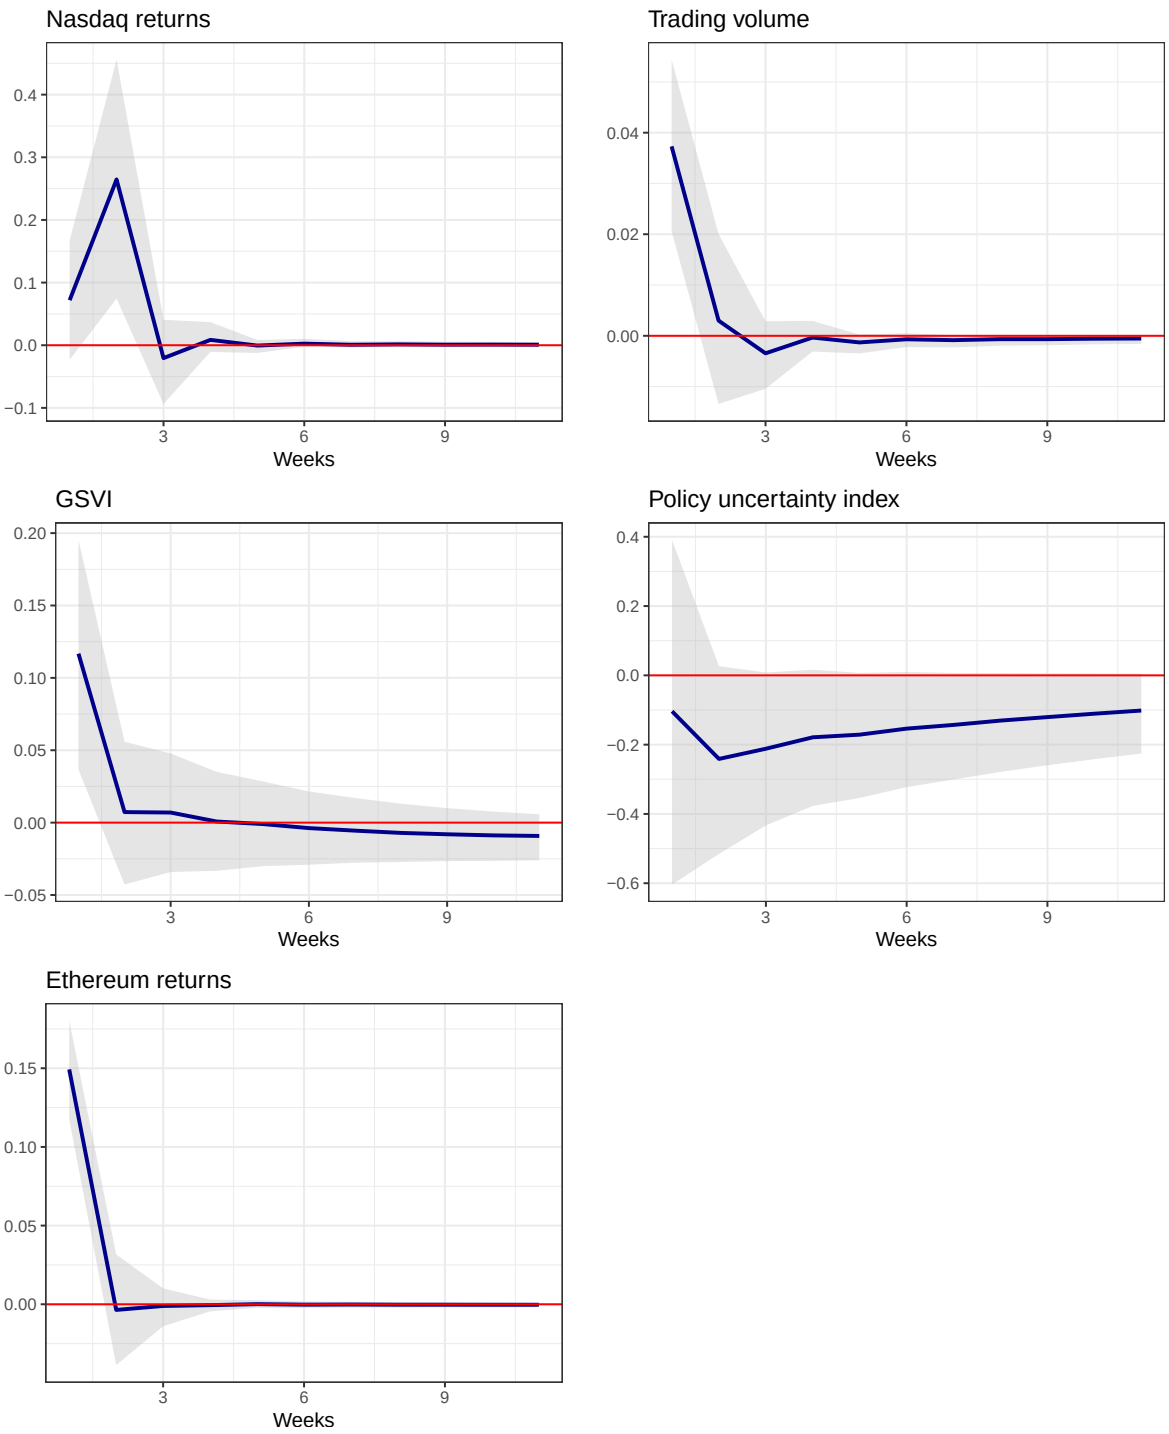

Notes: i) The blue lines presents the median impulse responses and the shaded area the 95% credible set. ii) Since stochastic volatility allows for the variance-covariance matrix to change over time, we estimate the impulse responses using the median matrix of the variance-covariance matrices.

Figure 4.4 plots the forecast error variance decomposition of bitcoin returns for a forecast horizon of 24 periods. In the initial period, 83% of the variance of bitcoin returns is explained by shocks in the bitcoin returns. Ethereum returns contribute to 10% of variation. The contribution of the remaining four variables is rather feeble, less than 7%. After one period, the percentage of variance attributed to bitcoin returns decreases by 5 percentage units. However, this difference is absorbed exclusively by Nasdaq returns, which explain approximately 6% of the forecast error variance of bitcoin returns. The contribution of the rest of the variables remains similar to that of the first period. Furthermore, as we move into later periods of the forecast horizon, we observe that shocks in each variable explain qualitatively the same percentage of the forecast error variance of bitcoin returns as in the second period.

To assess the validity of the results, we repeat the analysis and set the lag order equal to 2. The results remain qualitatively the same and are available upon request.

Figure 4.4: Forecast error variance decomposition of bitcoin returns.

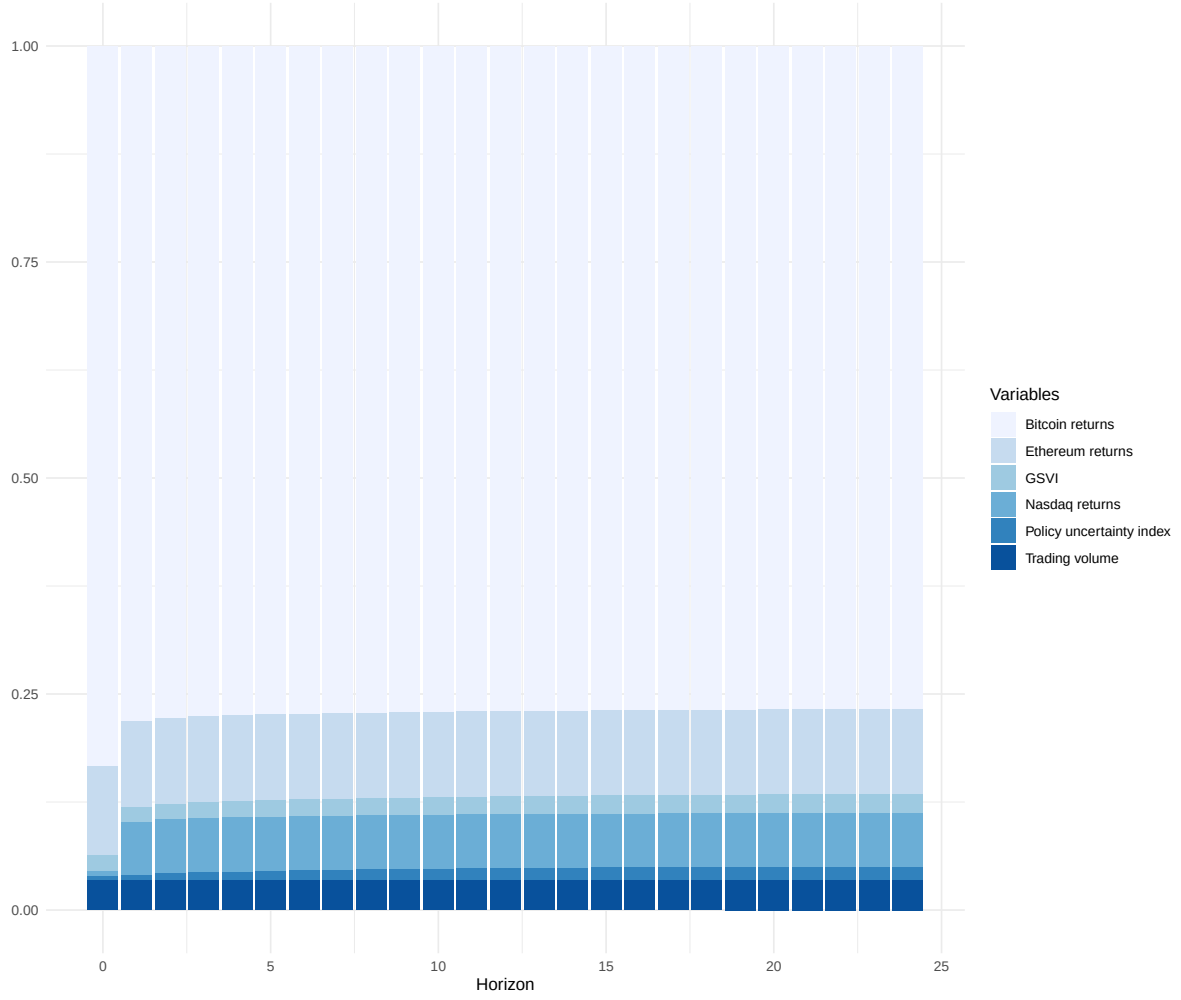

4.6 Comparison with a recursive algorithm

In this section we compare the proposed time-varying methodology with the recursive evolving window algorithm proposed by Shi et al. (2018) (for an additional emprical application see, Baum et al., 2021). For a given window of size r_0, the algorithm runs the Granger-causality test for all possible subsamples of length greater or equal than r_0. As a result, every point in the sample, except the ones in the initial window, is associated with a set of statistics. From each such set, we retrieve the supremum norm statistic and base statistical inference on the supremum norms. The recursive evolving algorithm since is a generalisation of both the rolling window and the forward expanding window algorithms.

Here, we consider a window size equal to 20% of the sample size (i.e. 77 observations). For all subsamples with length greater or equal than the selected window size, we estimate a time invariant VAR model with six endogenous variables (i.e. bitcoin returns, ethereum returns, Nasdaq returns, bitcoin trading volume, GSVI and the cryptocurrency policy uncertainty index) and then, the heteroskedastic-consistent Wald statistic for the null hypothesis that a variable Granger-causes bitcoin returns. Figure D2 presents the sequences of Wald statistics obtained from the recursive evolving algorithm. Note that there are no Wald statistics during the initial window size. The Figure also plots the series of Wald statistics obtained using the Bayesian time-varying methodology for ease of comparison between the two approaches. In the case of Nasdaq returns and the GSVI, we observe that the recursive evolving algorithm fails to reject the null hypothesis of no Granger-causality during specific periods. For example, according to the recursive algorithm, there is no Granger-causal effect from the GSVI to bitcoin returns during 2020 which contradicts the existing literature (for example, see Raza et al., 2022). In the case of policy uncertainty index and ethereum returns, the recursive algorithm indicates that there is a unidirectional Granger-causal relationship from the variables to bitcoin returns. Furthermore, the periods that the causal relationship is present, coincide with the periods indicated by the Bayesian time-varying approach. However, the critical value presented in Figure D2 is not size-adjusted. Shi et al. (2019) propose a bootstrapping algorithm that controls for size distortions (caused by the small sample size). If we consider a size-adjusted critical value, then the recursive algorithm detects a causal relationship for substantially shorter periods compared to the time-varying approach (note that the approach we propose in this essay does not require size-adjusted critical values). These results are available upon request.

4.7 Conclusions

In time series analysis, the assumption that the coefficients remain constant throughout the sample is not a realistic one. In this essay, we propose an approach to perform a time-varying Granger-causality test. The test combines the VAR estimates from a TVP-VAR model with heteroscedastic-consistent Granger-causality hypothesis testing. The proposed methodology does not rely on rolling window estimates. Consequently, this test is not subjected to small sample size distortions and does not require a subsample to initiate the recursive algorithm. The power and size properties of the test are evaluated through Monte Carlo simulations. We employ the new test to date-stamp the causal relationship between bitcoin returns and alternative variables. Specifically, in the empirical application, we consider a two-step approach. Initially, we estimate a Bayesian TVP-VAR with 30 variables (bitcoin returns and 29 potential determinants of bitcoin returns) and identify the most important variables of the model based on Bayesian variable selection algorithm. Next, we reestimate the TVP-VAR model using only the variables selected in the previous step and perform the analysis. Specifically, we employ the time-varying Granger-causality test which we discussed earlier, as well as impulse response function and forecast error variance decomposition analysis.

The results indicate that the most important variables in modelling bitcoin returns are information demand (proxied by the constructed GSVI), the cryptocurrency policy uncertainty index, bitcoin trading volume and ethereum and Nasdaq returns. Evidence from the time-varying Granger-causality test suggests that GSVI and ethereum returns Granger-cause bitcoin returns over the entire sampling period. The remaining three variables Granger-cause bitcoin returns only for specific periods in the sample, most profoundly during the first part of the sample over the period 2015 to 2017. Furthermore, the impulse response function analysis shows that with the exception of the policy uncertainty index, a shock in all variables yields a significant rise in bitcoin returns that lasts up to two weeks. On the contrary, an increase in the cryptocurrency policy uncertainty index leads to a decrease in bitcoin returns, however, the response is not significant. Finally, the forecast error variance decomposition indicates that regardless of the horizon, shocks in bitcoin returns contribute more than 78% of forecast error variance while shocks in Nasdaq and ethereum returns explain approximately 13% of the variance. The results reinforce the findings of previous studies and shed new light on the relationship between alternative variables and bitcoin returns. The empirical results also provide new insight into bitcoin market trading strategies.

It is worth emphasising that the time-varying Granger-causality approach we employ in this study is very general and goes well beyond the current bitcoin application.

It provides a rich and flexible environment for the detection of Granger-causality, that accounts for the disadvantages of similar methodologies. As a result, this methodology can be used in financial and economics applications where the data-set is characterised by structural breaks. In the cryptocurrency related literature, one could follow this approach, to examine for bidirectional causal relationship between a broad set of cryptocurrencies or examine the ability of traditional assets to Granger-cause alternative cryptocurrencies such as ethereum and litecoin.

Bibliography

Bakas, D., G. Magkonis, and E. Y. Oh (2022). What drives volatility in bitcoin market? *Finance Research Letters 50*, 103237.

Baum, C., S. Hurn, and J. Otero (2021). The dynamics of U.S. industrial production: A time-varying Granger causality perspective. *Econometrics and Statistics In Press*(Corrected Proof), 1–10.

Baur, D., K. Hong, and A. Lee (2018). Bitcoin: Medium of exchange or speculative assets? *Journal of International Financial Markets, Institutions and Money 54*, 177–189.

Blau, B. (2017). Price dynamics and speculative trading in bitcoin. *Research in International Business and Finance 41*(3), 493–499.

Bouri, E., R. Gupta, C. K. M. Lau, D. Roubaud, and S. Wang (2018). Bitcoin and global financial stress: A copula-based approach to dependence and causality in the quantiles. *The Quarterly Review of Economics and Finance 69*, 297–307.

Chan, J., G. Koop, D. Poirier, and J. Tobias (2019). *Bayesian Econometric Methods* (2 ed.), Volume 7. Cambridge University Press.

Cheah, E.-T. and J. Fry (2015). Speculative bubbles in Bitcoin markets? An empirical investigation into the fundamental value of Bitcoin. *Economics Letters 130*, 32–36.

Cheah, E.-T. and J. Fry (2016). Negative bubbles and shocks in cryptocurrency markets. *International Review of Financial Analysis 47*, 343–352.

Chen, W., H. Xu, L. Jia, and Y. Gao (2021). Machine learning model for bitcoin exchange rate prediction using economic and technology determinants. *International Journal of Forecasting 37*(1), 28–43.

Ciner, C., B. Lucey, and L. Yarovaya (2022). Determinants of cryptocurrency returns: A LASSO quantile regression approach. *Finance Research Letters 49*, 102990.

Cogley, T. and T. Sargent (2005). Drift and Volatilities: Monetary Policies and Outcomes in the Post WWII U.S. *Review of Economic Dynamics 8*(2), 262–302.

Corbet, S., B. Lucey, and L. Yarovaya (2018). Datestamping the bitcoin and ethereum bubbles. *Finance Research Letters 26*, 81–88.

Da, Z., J. Engelberg, and P. Gao (2011). In search of attention. *The Journal of Finance 66*(5), 1461–1499.

Dastgir, S., E. Demir, G. Downing, G. Gozgor, and C. K. M. Lau (2019). The causal relationship between bitcoin attention and bitcoin returns: Evidence from the copula-based granger causality test. *Finance Research Letters 28*(C), 160–164.

de la Horra, L., G. de la Fuente, and J. Perote (2019). The drivers of bitcoin demand: A short and long-run analysis. *International Review of Financial Analysis 62*, 21–34.

Dyhrberg, A. H. (2016). Bitcoin, gold and the dollar - A GARCH volatility analysis. *International Review of Financial Analysis 16*, 85–92.

Fry, J. (2018). Booms, busts and heavy-tails: The story of Bitcoin and cryptocurrency markets? *Economics Letters 171*, 225–229.

Fry, J. and J.-P. Serbera (2020). Quantifying the sustainability of bitcoin and blockchain. *Journal of Enterprise Information Management 33*(6), 1379 – 1394.

Guesmi, K., S. Saadi, I. Abid, and Z. Ftiti (2019). Portfolio diversification with virtual currency: Evidence from bitcoin. *International Review of Financial Analysis 63*, 431–437.

Ji, Q., E. Bouri, R. Gupta, and D. Roubaud (2018). Network causality structures among bitcoin and other financial assets: A directed acyclic graph approach. *The Quarterly Review of Economics and Finance 70*, 203–213.

Koop, G. and D. Korobilis (2010). Bayesian multivariate time series methods for empirical macroeconomics. *Foundations and Trends in Econometrics 3*(4), 267–358.

Korobilis, D. (2013). VAR forecasting using Bayesian variable selection. *Journal of Applied Econometrics 28*(2), 204–230.

Kristoufek, L. (2013). BitCoin meets Google Trends and Wikipedia: Quantifying the relationship between phenomena of the Internet era. *Scientific reports 3*(1), 1–7.

Lahiani, A., A. Jeribi, and N. B. Jlassi (2021). Nonlinear tail dependence in cryptocurrency-stock market returns: The role of bitcoin futures. *Research in International Business and Finance 56*, 101351.

Lee, A., M. Li, and H. Zheng (2020). Bitcoin: Speculative asset or innovative technology? *Journal of International Financial Markets, Institutions and Money 67*, 101209.

Li, R., S. Li, D. Yuan, and H. Zhu (2021). Investor attention and cryptocurrency: Evidence from wavelet-based quantile Granger causality analysis. *Research in International Business and Finance 56*, 101389.

Li, X. and C. A. Wang (2017). The technology and economic determinants of cryptocurrency exchange rates: The case of Bitcoin. *Decision Support Systems 95*, 49–60.

Liang, C., Y. Zhang, X. Li, and F. Ma (2022). Which predictor is more predictive for Bitcoin volatility? And why? *International Journal of Finance & Economics 27*(2), 1947–1961.

Lin, Z.-Y. (2021). Investor attention and cryptocurrency performance. *Finance Research Letters 40*, 101702.

Liu, Y. and A. Tsyvinski (2020, 09). Risks and returns of cryptocurrency. *The Review of Financial Studies 34*(6), 2689–2727.

Liu, Y., A. Tsyvinski, and X. Wu (2022). Common risk factors in cryptocurrency. *The Journal of Finance 77*(2), 1133–1177.

Lucey, B., S. Vigne, L. Yarovaya, and W. Yizhi (2022). The cryptocurrency uncertainty index. *Finance Research Letters 45*, 102147.

Otero, J., T. Panagiotidis, and G. Papapanagiotou (2021). Testing for exuberance in house prices using data sampled at different frequencies. *Studies in Nonlinear Dynamics & Econometrics 26*, 675–691.

Panagiotidis, T., G. Papapanagiotou, and T. Stengos (2022). On the volatility of cryptocurrencies. *Research in International Business and Finance 62*, 101724.

Panagiotidis, T., T. Stengos, and O. Vravosinos (2018). On the determinants of bitcoin returns: A lasso approach. *Finance Research Letters 27*(C), 235–240.

Primiceri, G. (2005). Time varying structural vector autoregressions and monetary policy. *Review of Economic Studies 72*(3), 821–852.

Raza, S. A., L. Yarovaya, and N. Shah (2022). Google Trends and cryptocurrencies: A nonparametric causality-in-quantiles analysis. *International Journal of Emerging Markets In Press*, 1–18.

Rothman, T. (2019). Trading the dream: Does social media affect investors' activity - the story of Twitter, Telegram and Reddit. *International Journal of Financial Research 10*(2), 147–152.

Shen, D., A. Urquhart, and P. Wang (2019). Does Twitter predict bitcoin? *Economics Letters 174*, 118–122.

Shi, S., S. Hurn, and P. C. Phillips (2019). Causal change detection in possibly integrated systems: revisiting the money–income relationship. *Journal of Financial Econometrics 18*(1), 158–180.

Shi, S., P. C. Phillips, and S. Hurn (2018). Change detection and the causal impact of the yield curve. *Journal of Time Series Analysis 39*(6), 966–987.

Song, X. and A. Taamouti (2019). A better understanding of Granger causality analysis: A big data environment. *Oxford Bulletin of Economics and Statistics 81*(4), 911–936.

Xiang, J., G. Guo, and Q. Zhao (2022). Testing for a moderately explosive process with structural change in drift. *Oxford Bulletin of Economics and Statistics 84*(2), 300–333.

Yarovaya, L. and D. Zięba (2022). Intraday volume-return nexus in cryptocurrency markets: Novel evidence from cryptocurrency classification. *Research in International Business and Finance 60*, 101592.

Yermack, D. (2015). Chapter 2 - Is Bitcoin a real currency? An economic appraisal. In D. Lee Kuo Chuen (Ed.), *Handbook of Digital Currency*, pp. 31–43. San Diego: Academic Press.

Yousaf, I. and L. Yarovaya (2022). The relationship between trading volume, volatility and returns of Non-Fungible Tokens: evidence from a quantile approach. *Finance Research Letters 50*, 103175.

Zapata, H. O. and A. N. Rambaldi (1997). Monte Carlo evidence on cointegration and causation. *Oxford Bulletin of Economics and Statistics 59*(2), 285–298.

Appendix

Table D1: Data sources.

	Variable	Source
	Bitcoin returns	Yahoo Finance
	Bitcoin returns' volatility (VOL)	authors' calculations
	Ethereum returns (ETH)	Yahoo Finance
Sentiment	SVI (GSVI)	Google Trends
	Uncertainty Cryptocurrency policy (UCRY policy)	Cryptocurrency indices
	Uncertainty Cryptocurrency returns (UCRY price)	Cryptocurrency indices
Supply	Total number of bitcoins in circulation (NBC)	blockchain.com
Demand	Trading volume (TV)	Yahoo Finance
	Confirmed transactions per day (CTR)	blockchain.com
	Unique addresses used (ADD)	blockchain.com
	Hash Rate (HASH)	blockchain.com
	Network Difficulty (DIF)	blockchain.com
Stock market	S&P500 (SP500)	Yahoo Finance
	Dow Jones NYSE index (DJI)	Yahoo Finance
	NASDAQ index (NASDAQ)	Yahoo Finance
	EURO STOXX 50 index (STOXX)	Yahoo Finance
	FTSE 100 index (FTSE)	Wall Street Journal
	Nikkei 225 index (NIKKEI)	Yahoo Finance
	Shanghai Composite Index (SSEC)	Yahoo Finance
	CBOE Market Volatility Index (VIX)	Yahoo Finance
	EURO STOXX 50 Volatility Index (VSTOXX)	Wall Street Journal
Commodities	Brent (BRENT)	FRED
	WTI (WTI)	FRED
	Gold (GOLD)	Yahoo Finance
Exchange rate	Euro (EURO)	FRED
	Sterling (GBP)	FRED
	Yuan (YUAN)	FRED
	Yen (YEN)	FRED
Financial variables	Fed Funds effective rate (FFER)	FRED
	ECB deposit facility rate (ECB)	FRED

Notes: i) We calculate the volatility of bitcoin returns as the square of bitcoin returns. ii) The mnemonics in the parentheses are used to denote the variable in the Tables and Figures.

Figure D1: Data in levels: Bitcoin and ethereum.

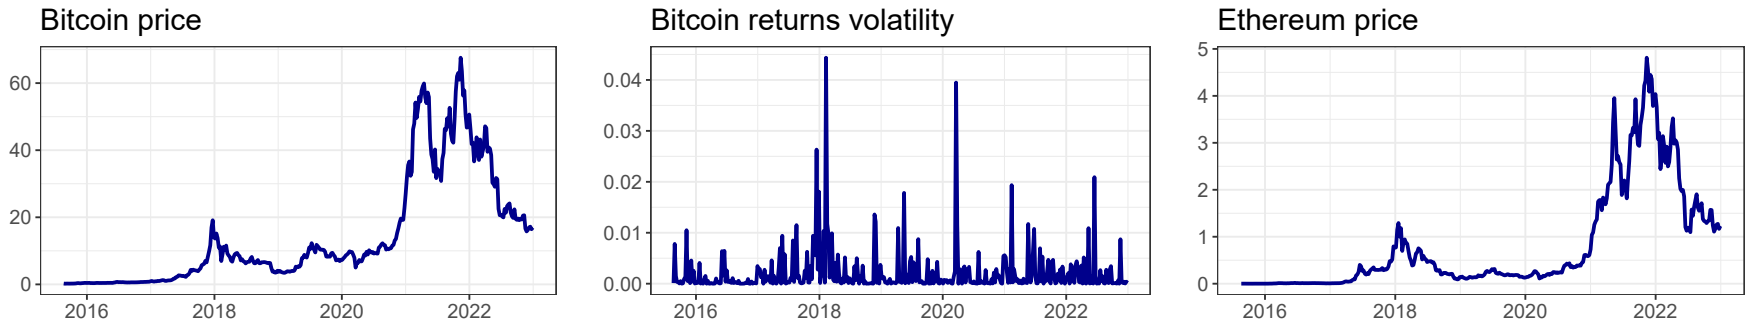

Notes: i) Bitcoin, ethereum and the stock market indices are in thousands of US dollars. Volume is in millions of US dollars. The number of unique addresses and daily confirmed transactions is in thousands. Hash rate is millions of TH/s. Difficulty is in trillions. The exchange rates are expressed as the price of domestic currency to one US dollar.

Figure D1: Data in levels (continued): Bitcoin related indices.

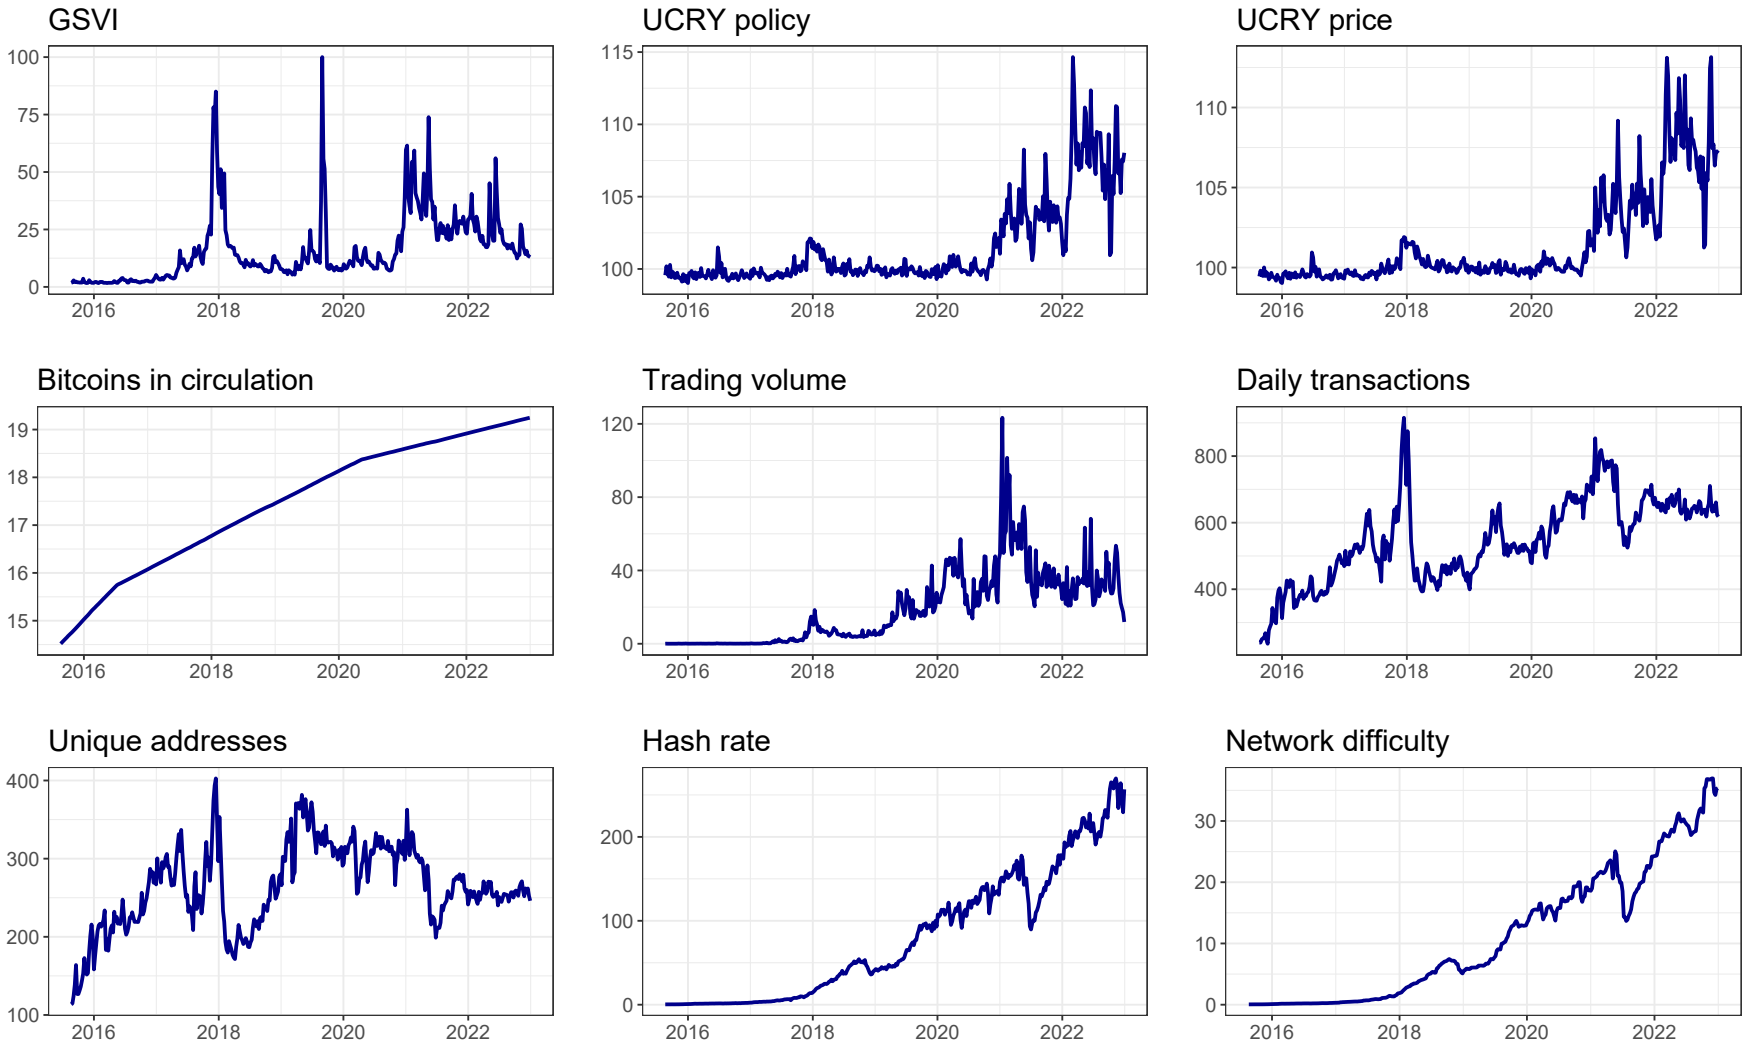

Figure D1: Data in levels (continued): Stock market prices and volatility.

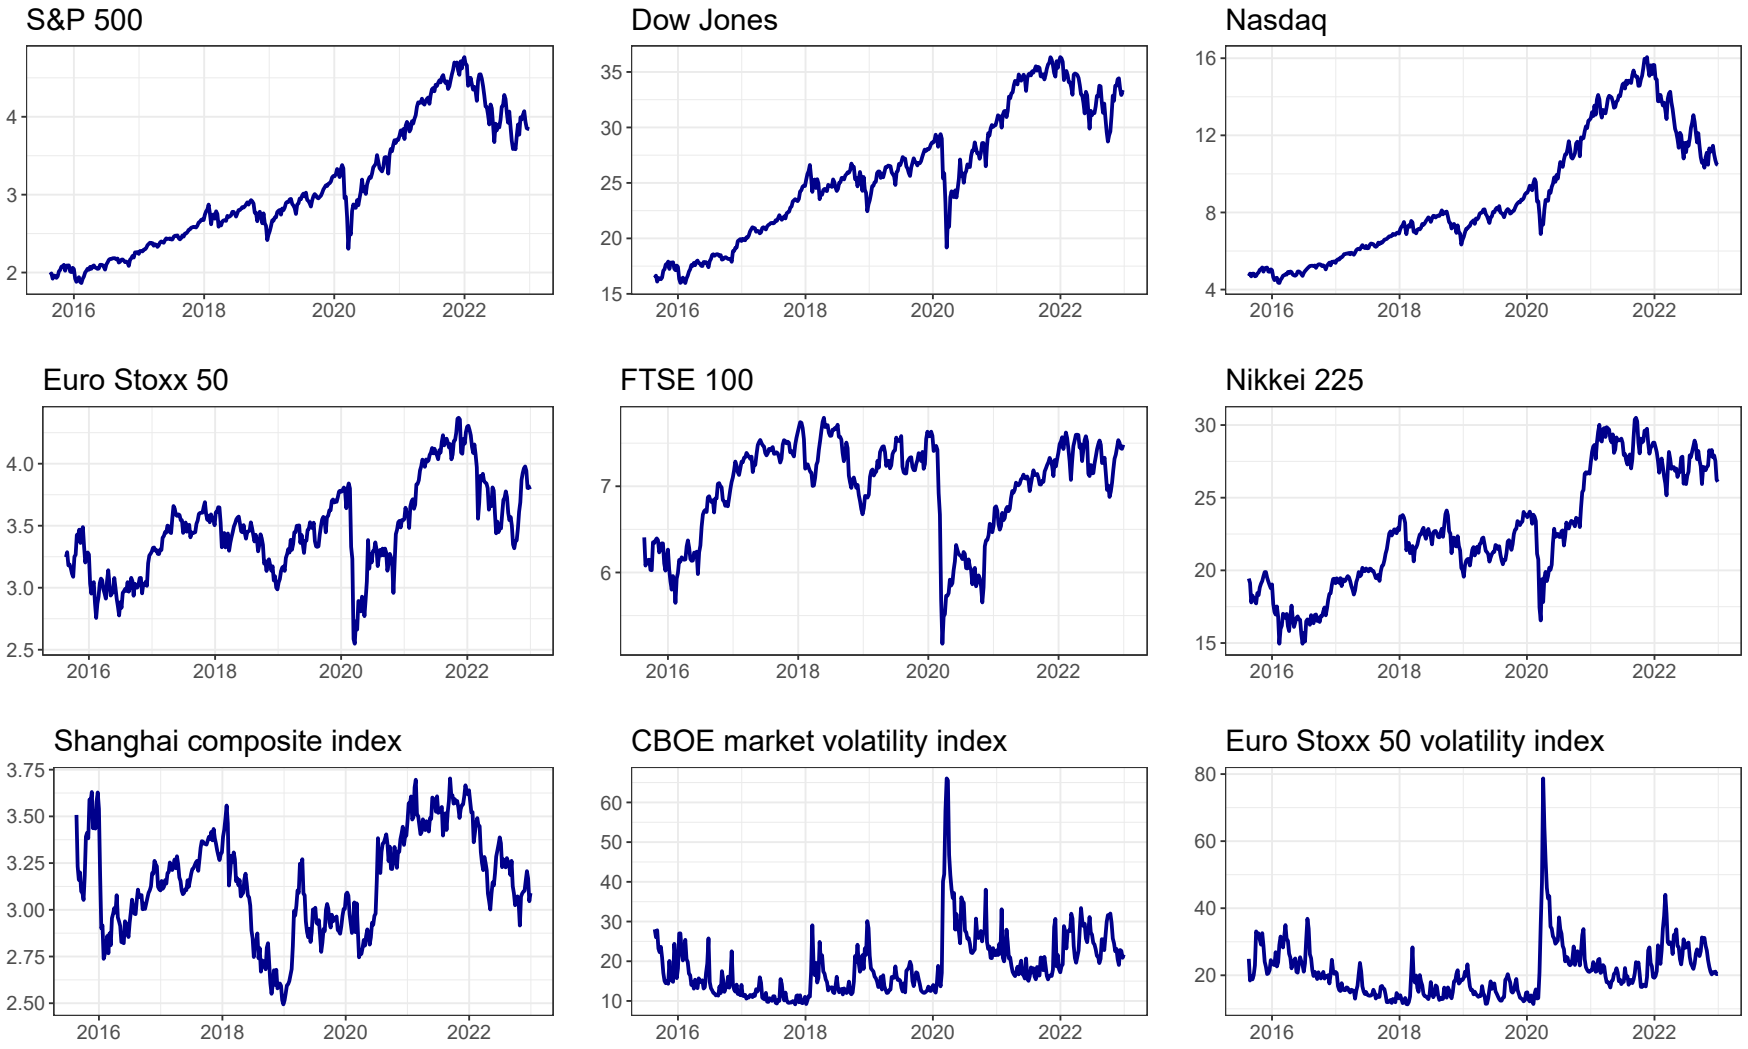

Figure D1: Data in levels (continued): Commodities, exchange rates and interest rates.

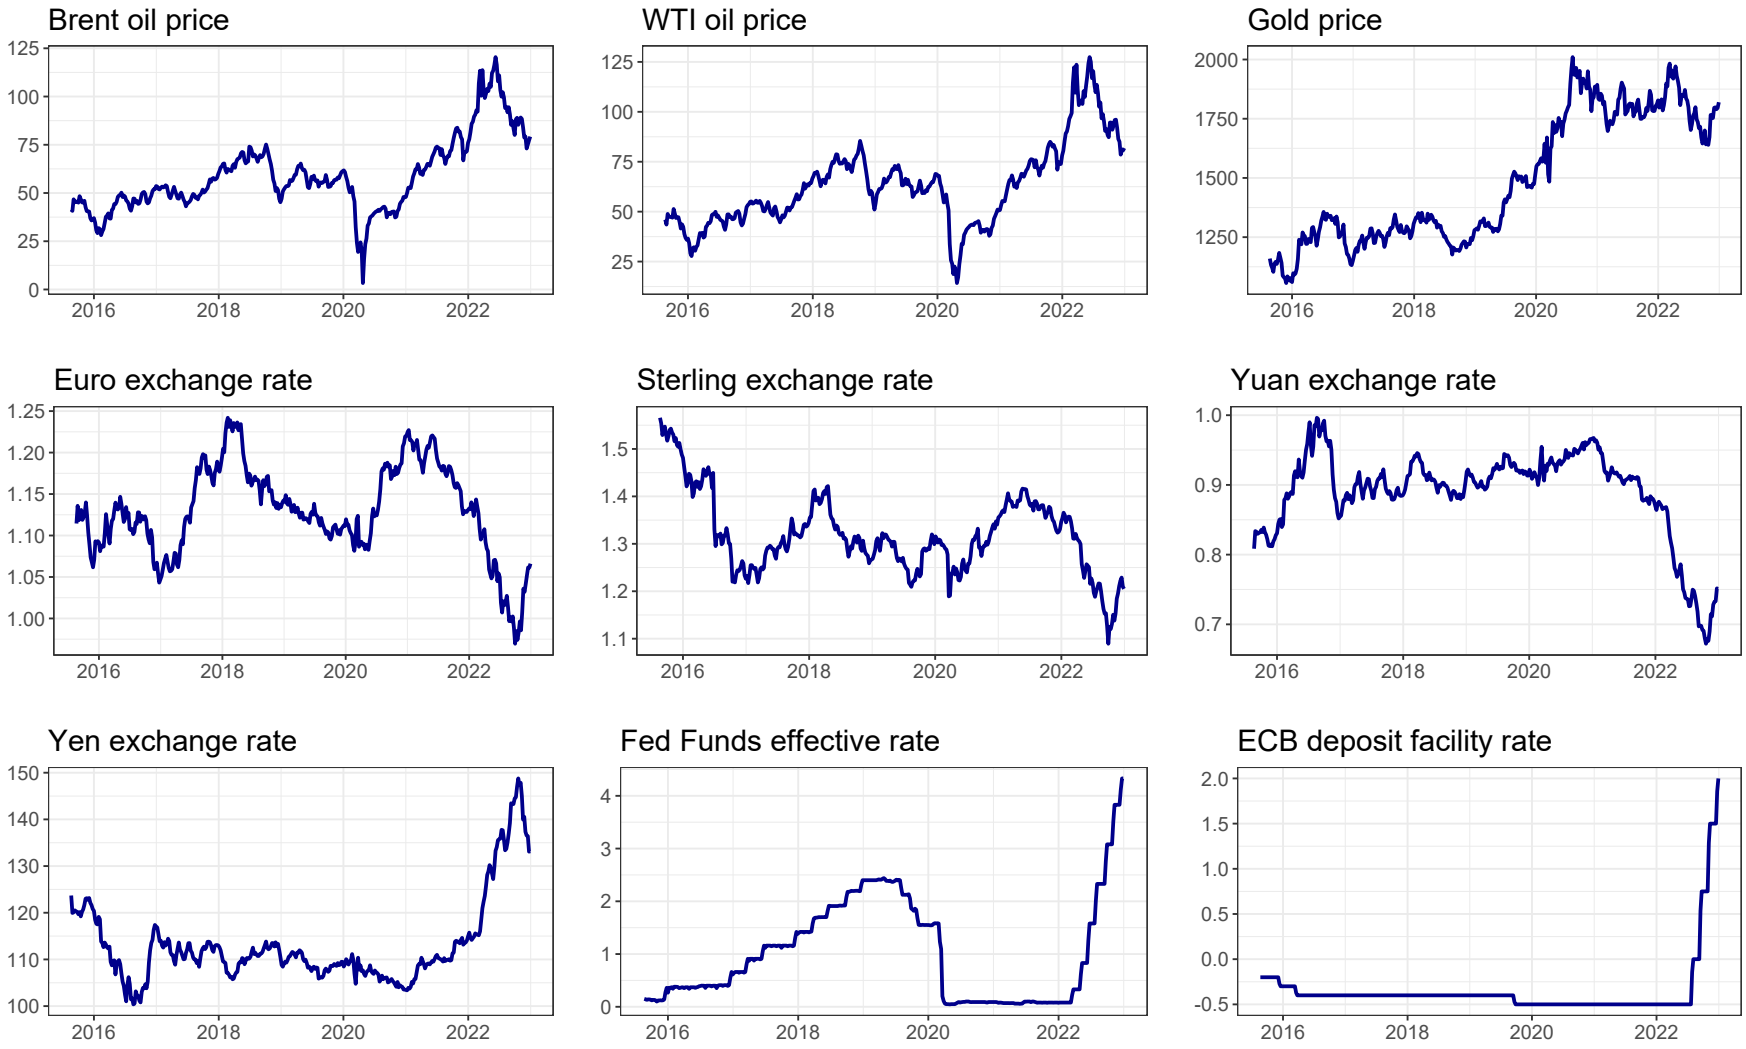

Figure D2: Granger-causality results from the recursive evolving window algorithm

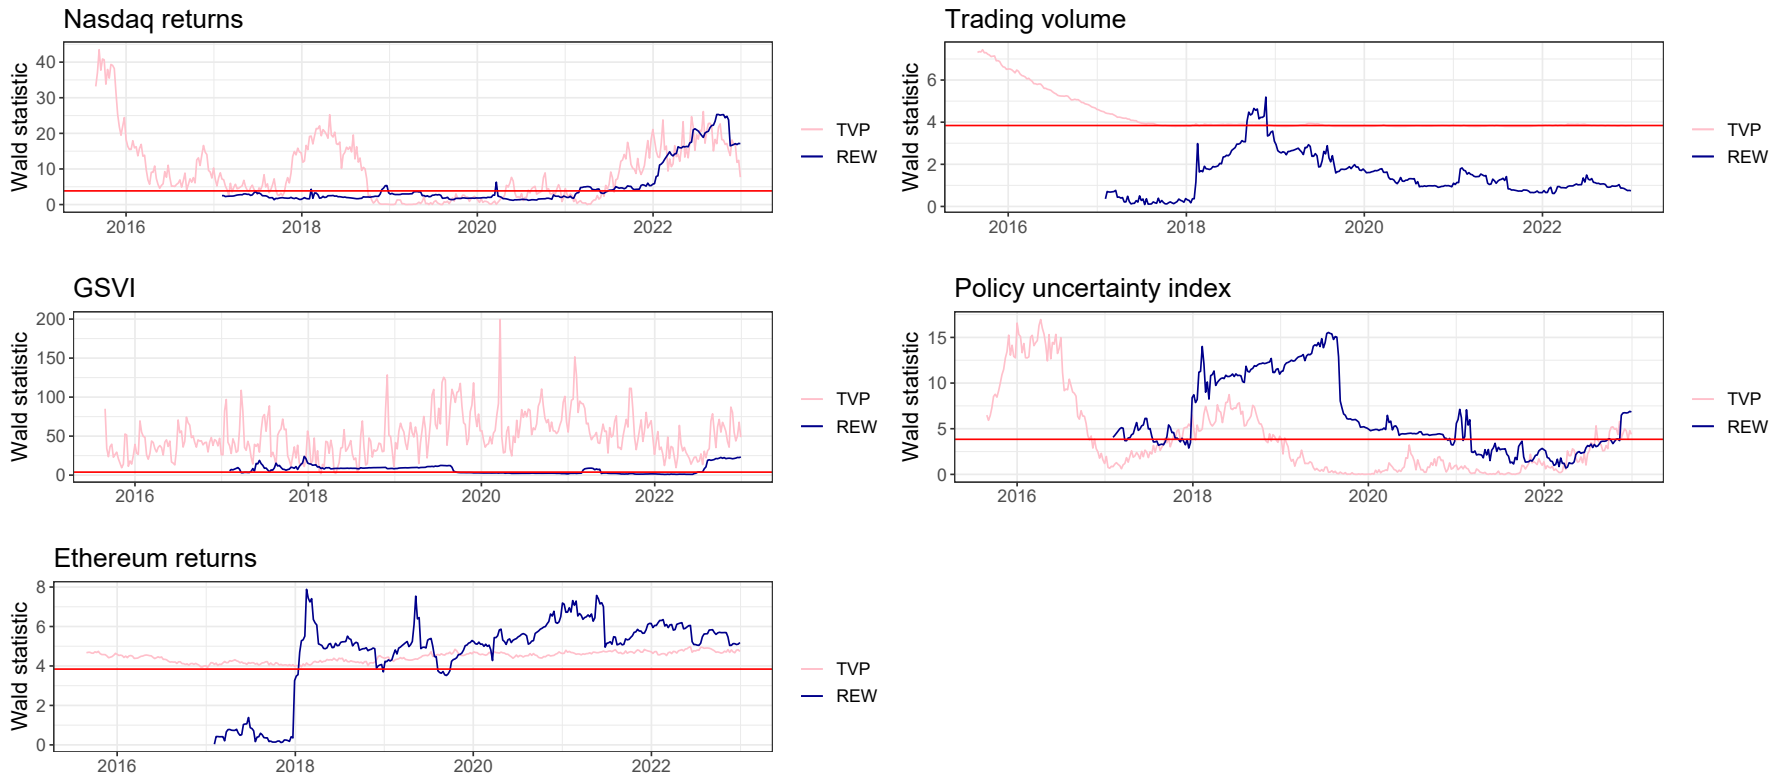

Notes: i) The red blue lines denote the sequence of Wald statistics obtained using the algorithm of Shi et al. (2019). The pink lines denote the sequence of Wald statistics obtained using the Bayesian TVP methodology. The red line represents the critical value at the 5% significance level. ii) The null hypothesis of Granger non-causality is rejected if the Wald statistic is greater than the critical value.