UNIVERSITA' CATTOLICA DEL SACRO CUORE

MILANO

Dottorato di ricerca in Economia e Finanza

Ciclo XXXII

S.S.D: SECS-P/01 - Economia Politica

THREE ESSAYS ON UNCERTAINTY: POLICY REACTIONS AND FINANCIAL CONSEQUENCES

Tesi di Dottorato di: Catalin Dragomirescu-Gaina

Matricola: 4612802

Anno Accademico 2018 / 2019

Dottorato di ricerca in Economia e Finanza

Ciclo XXXII

S.S.D: SECS-P/01 - Economia Politica

THREE ESSAYS ON UNCERTAINTY: POLICY REACTIONS AND FINANCIAL CONSEQUENCES

Coordinatore: Prof. Emanuele Bacchiocchi

Tesi di Dottorato di: Catalin Dragomirescu-Gaina

Matricola: 4612802

Anno Accademico 2018 / 2019

To Adina, Silvia and my parents

Synthesis

This thesis focuses on understanding broader or Knightian uncertainty and its relation with financial risk, given the overlaps existing between these two important concepts in the economics and finance literature. Each chapter tackles a different aspect of uncertainty, from a different angle and using a different methodology and data set. For this reason, the chapters are structured as standalone papers to assist readers and improve readability.

The first chapter of the thesis is titled: "Uncertainty spill-overs: when policy and financial realms overlap". It attempts to identify two different uncertainty shocks along with their policy effects and consequences, by modelling the complex intertwining between policy and financial realms that appears to be particularly relevant in the European context. The methodological differences in the construction and statistical properties of the two proxies, used for financial and policy uncertainty, facilitate the implementation of a recent structural identification approach based on magnitude restrictions. One of the main contributions of the chapter is to apply magnitude restrictions in a multi-country context with the aim of identifying two uncertainty shocks. This identification approach offers several advantages over other alternative structural identification methods, which the chapter discusses in details. After estimating the model, we recover the two structural uncertainty shocks, and find they match the dates and timing of some remarkable events that marked the recent history of the European project. Although there are significant cross-influences and overlaps between financial and policy uncertainty, the later reacts stronger to shocks in the former proxy; in other words, it is more likely that financial frictions and stress amplify uncertainty in the policy realm than vice-versa. The empirical results also point to ECB adopting a more pro-active stance towards policy uncertainty shocks in order to prevent further segmentation of the Euro Area financial market during periods of turmoil, but a more (passive or) accommodative stance towards financial uncertainty shocks.

The second chapter discusses the trade-off between prediction accuracy and reaction speed that allows hedge funds, some of the most astute investors today, to better time the market and profit during turmoil periods. The chapter is titled: "Trading off accuracy for speed: Hedge Funds' decision-making under uncertainty", and is co-authored with prof. D. Philippas and prof. M. Tsionas. A mathematical formulation of the trade-off casts the decision-making process in a Bayesian framework, while the empirical analysis employs different data-filtering techniques to distinguish between different prediction accuracy levels. According to the main results, less accurate predictions can speed up hedge funds' reactions to changes in the information set. For many hedge funds that claim to maintain a low *beta* and in the same time generate profits, market timing is essential, and therefore reaction speed becomes a means to achieve better timing. We justify our empirical findings in a simulation exercise, highlighting the importance of market timing abilities for active players like hedge funds.

The third chapter is titled: “On herding behaviour, ‘green’ energy and uncertainty” and is co-authored with prof. D. Philippas and prof. E. Galariotis. The chapter discusses challenges arising from the ongoing transition to a low-carbon economy and the portfolio choices that investors are facing during this process. The ‘green’ sector today looks as an exciting opportunity for investors in the energy sector, but uncertainty prevails in relation to the long-term economic viability of new ‘green’ technologies. In addition, the ‘green’ sector faces constant regulatory and policy challenges. With multiple uncertainty sources, investors should worry about price distortions driven by their own behavioural biases, which arise particularly in markets characterised by uncertainty and information frictions. This chapter aims at contributing to the discussion related to herding behaviour, and therefore learning in financial markets. In a context where investors can opt between investing in an old technology, like oil, and a new, ‘greener’ technology, we find that herding responds to oil returns and ‘green’ volatility shocks. Thus, investing into an old technology requires no more than information on returns, while opting for a new investment opportunity requires an entirely better information set.

Contents

Chapter 1

Uncertainty spill-overs: when policy and financial realms overlap[♯]

Abstract

This chapter aims at contributing to a new and growing empirical literature strand on uncertainty-related topics. No matter its source, financial- or policy-related, uncertainty feeds continuously onto itself, contaminating the real sector, and leading to identification challenges in empirical applications. We propose a new application of a recent identification approach to reveal and separate two different uncertainty sources. We model the complex intertwining between policy and financial realms, whose interactions create amplification mechanisms for country-specific uncertainty shocks, framing our empirical analysis within a multi-country model set in the European context. Stark methodological differences between our financial and policy uncertainty proxies allow us to use the structural identification approach based on magnitude restrictions proposed in De Santis and Zimic (2018) that offers several advantages over other alternative identification methods. Using impulse responses derived from a global VAR specification, we find persistent effects for both uncertainty shocks, including significant cross-border spill-overs. We reveal significant cross-influences between the two uncertainty proxies, with policy uncertainty reacting stronger to financial uncertainty shocks than vice-versa, in line with the existing evidence on the importance of financial frictions. Our identified structural shocks match the dates of some remarkable events that marked the recent history of the European project. With respect to ECB policy reactions, there are stronger but less persistent responses to financial uncertainty shocks compared to policy uncertainty shocks, pointing to ECB adopting a more pro-active stance towards the latter shocks, and a more (passive or) accommodative stance towards the former shocks. We suggest that a possible justification for such ECB actions might come from its attempt to tame policy uncertainty in order to prevent further segmentation of the Euro Area financial market.

JEL codes: C3, E58, E60, F36, F40

Keywords: policy uncertainty, financial integration, global VAR

[♯] *Acknowledgements*: I am grateful to my supervisor prof. Emanuele Bacchiocchi for his comments, suggestions and continuous feedback. I am thankful to prof. Gianluca Femminis for his constant encouragement and support while working on this chapter. I would also like to thank professors Elena Beccalli, Giulio Palomba and Eduardo Rossi for their many comments and suggestions, and to Alessandro Galesi for support on the Matlab codes used to run the empirical analysis in this chapter.

1. INTRODUCTION

For a few days every January in Davos, Switzerland, global financial elite mingles with political elite, central bankers and other policymakers. Most likely, policy and financial realms cannot be completely separated even if one dares considering centuries of history. From an analytical perspective, this leads to unexpected cross-influences that amplify each other, especially during uncertain times. Financial stress and market uncertainties can bring changes in policies or political contexts, as much as uncertainty stemming from policy changes creates anxieties for financial investors. No matter its source, financial or policy-related, uncertainty will feed onto itself, contaminating other areas and leading to identification challenges in empirical applications. Unfortunately, markets and investors are better equipped to evaluate and price risk rather than uncertainty, which is a broader concept encompassing risk and requiring proper analytical methods. We try to add to the existing stock of analytical methods able to disentangle among various sources of uncertainty in a multi-country context, where cross-border spill-overs and cross-influences are expected to pose additional identification challenges.

From this perspective, the European Union (EU), and the Euro Area (EA) in particular – with its rather incomplete institutional architecture –, make for an interesting case due to its high potential for uncertainty spill-overs. On the one hand, domestic policy uncertainty can reverberate at the European and global levels with serious financial consequences measured in terms of bond yields, financial stock prices or currency moves. In June 2015 the Greek government called a referendum over its bailout terms, generating chaos in European policy circles, but also among financial investors who feared a Euro Area (EA) breakdown; as market sentiment turned sour, Greek sovereign bond spreads reached unprecedented levels and the country was effectively cut off global financial markets, while domestic banks suffered and were forced to impose strict capital controls. On the other hand, it is the banking sector turmoil that echoes in the policy domain, as risks are transferred from the private to the public sector due to bank-rescue packages that increase sovereign and contagion risks (see Acharya et al. 2014; Attinasi et al. 2010; Bicu and Candelon, 2013; Stângă, 2014). Ireland perfectly illustrates this latter case, when the government introduced guarantees to address the weakness of the domestic banking sector in September 2008, right after the Lehman shock; as a result, banks' credit default swaps (CDS) came down but the Irish sovereign CDS spiked abruptly (Stângă, 2014; Leonello, 2018).

The present paper aims at exploring this complex intertwining, which is particularly prevalent in Europe, between policy and financial realms, whose interactions might create amplification mechanisms for country-specific uncertainty shocks. Whether such mechanisms work to amplify financial uncertainty, policy uncertainty, none or both is the most important research question we address in this chapter. We seek therefore to contribute to a new and rapidly expanding literature strand that deals with various uncertainty measures, their sources, effects, and cross-border spill-overs (see among many others; Bekaert, Hoerova, and Lo Duca, 2013; Caldara, Fuentes-Albero, Gilchrist, and

Zakrajšek, 2016; Bacchiocchi, 2017; Shin and Zhong, 2018; Ludvigson, Ma and Ng, 2019; Angelini, Bacchiocchi, Caggiano, Fanelli, 2019).

We also aim at understanding what specific role the European Central Bank (ECB) has played in counteracting various uncertainty sources and their spill-overs at the European level. Over the last decade, the ECB considerably expanded its policy toolkit, took greater supervisory and regulatory duties, and stepped in when there was no credible policy actor for global financial markets[1], up to the point of being called 'the only game in town'.[2] On the back of a rather complicated EA governance structure, the ECB provided an effective backstop to area-wide financial stress, while treating country-specific shocks with more flexibility. This is in spite of the fact that, in many instances, country-specific factors have penetrated the decision-making process in Brussels and Frankfurt.

Complex identification challenges arise within a multi-country settings, such as the EA, due to its significant financial integration, but incomplete political integration, where information frictions are important (see Freixas and Holthausen, 2004). During the European sovereign debt crisis, domestic banks in some EA periphery were given incentives to draw more central bank liquidity, largely against domestic sovereign bonds. Battistini, Pagano, and Simonelli (2014) and Acharya and Steffen (2015) provide empirical evidence on these mechanisms, where bailed-out periphery banks hold more periphery sovereign debt.[3] Recently, the theoretical work of Farhi and Tirole (2017), Leonello (2018), and Cooper and Nikolov (2018) sheds light on the feedback-loops between sovereigns and banks, but strong feedback-loops can blur the thin separation line between financial and policy realms.

Understanding the ECB role is an important topic because the EA suffers from a lack of institutional leadership to deal with several uncertainty sources. The existing literature on (monetary and fiscal) policy interactions within a common currency area does not provide us with sufficient clarifications in this regard (for a recent survey, see Foresti 2018). ECB faces numerous and delicate policy trade-offs in pursuing its price stability mandate, set according to the EU Treaties. A clearer distinction between policy and financial uncertainty shocks could improve ECB policy effectiveness, and even shield it from possible legal actions.[4] There have been many controversies surrounding ECB monetary policy conduct, especially with respect to its unconventional measures, like the various asset purchasing

[1] Mario Draghi's speech on 26th July 2012 has been considered a cornerstone moment for the EA sovereign debt crisis. See https://www.ecb.europa.eu/press/key/date/2012/html/sp120726.en.html.
[2] See Otmar Issing's comment at: https://www.centralbanking.com/central-banks/economics/2473842/otmar-issing-on-why-the-euro-house-of-cards-is-set-to-collapse.
[3] There are plenty of other empirically relevant studies on the moral hazard prevalent during the European sovereign debt crisis. Acharya, Drechsler, and Schnabl (2014) show that CDS for sovereigns and banks commove over the European crisis period, but not much before the crisis. Koijen, Koulischer, Nguyen, and Yogo (2017) document the home bias existing in vulnerable countries during the implementation of the ECB asset purchasing programmes.
[4] See the recent decision of the Court of Justice of the European Union in favour of the ECB's Public-Sector Purchase Programme (PSPP) at https://www.reuters.com/article/us-ecb-policy-court/ecb-wins-courts-backing-for-buying-government-debt-idUSKBN1OA0Q0.

programs implemented over the last decade.[5] In August 2011, for example, the Securities Markets Programme (SMP) made some sizeable bond purchases from the EA periphery, especially Italian and Spanish sovereigns, with some positive effects on spreads in unsettled market conditions. However, the program was soon suspended for Italian bonds as it became clear that the Berlusconi government was not delivering on its promised economic reforms; fast forward in November 2011, market confidence in the Italian government collapsed and a new prime-minister was appointed.

Given the large consequences stemming from the interaction of financial and policy realms within the EA, as discussed above, it is important to evaluate whether there are sizable spill-overs of country-specific uncertainty shocks, and whether ECB can play any specific role. Our main contribution is to approach these important research questions from an empirical perspective that is able to deal with the inherent identification challenges that arise in a multi-country setting. We use a global vector autoregressive (or GVAR) model specification (as in Dees et al., 2007; Georgiadis, 2015; Burriel and Galesi, 2018), and a new identification approach based on magnitude restrictions, recently proposed by De Santis and Zimic (2018). There are other few distinct but comparable approaches in a rapidly expanding empirical literature aiming at identifying (different types of) uncertainty shocks (e.g. Bacchiocchi, 2017; Shin and Zhong, 2018; Ludvigson, Ma and Ng, 2019; Angelini et al., 2019); as each methodological approach has its own merits, we regard them as largely complementary to ours. Inspired by event studies, the identification based on magnitude restrictions was proposed by De Santis and Zimic (2018) to expose spill-overs between U.S. and European sovereign bond yields. However, it is quite general and allows for the identification of shocks from within any strongly correlated variables, especially in cases of conceptual overlaps, like in the case of the two uncertainty proxies used in this study. An important aspect in our application is that the two proxies should focus on distinct data sources, and rely on different measurement approaches.

To capture financial uncertainty, we use the Composite Indicator for Systemic Stress (CISS), a highly relevant policy indicator for ECB, which also makes this indicator available on a weekly frequency, and for most EU Member States (see Hollo et al., 2010). Compared to other financial uncertainty measures that are probably more readily available (e.g. CDS, volatility, cross-sectional variation), composite indicators summarize a higher dimensional space and are more efficient in reflecting financial stress across several market segments.[6] Broader (or Knightian) uncertainty, instead, stemming from changes in the political landscape, rhetoric, opinions and policies is harder to measure (see Bekaert et al., 2013; Jurado, Ludvigson, and Ng, 2015; Baker, Bloom and Davis, 2016; Ferrara et

[5] In addition to PSPP, ECB conducted the Securities Markets Programme (SMP, May 2010-2012), and Outright Monetary Transactions (OMT, announced in September 2012) were targeted mostly at countries with severely impaired financial markets.

[6] Various studies, such as Fratzscher et al. (2016), Moder (2017), Burriel and Galesi (2018), Boeckx et al (2017) use the CISS index proposed in Hollo et al. (2010) to uncover transmission channels and consequences of financial stress across European markets.

al., 2018; Ludvigson et al., 2019). The recent literature is booming with different measures of this type of uncertainty, spanning different methodologies and data sources. However, some of the best known indicators rely heavily on media sources. In a highly influential paper, Baker et al. (2016) propose an *economic policy uncertainty* (EPU) measure based on the frequency of some relevant keywords in various newspapers (and other commonly available media sources); they further show their indicator is orthogonal to other common measures of risk and uncertainty, such as forecasts dispersion or financial volatility etc. We rely on EPU to measure policy uncertainty, mostly because of its wide availability for different EU and EA countries. Our selected CISS and EPU indexes, therefore, rely on different data sources and measurement methodologies. A closely related literature strand employs sovereign and banking risk measures derived from market instruments, like CDSs (see Bicu and Candelon, 2013; Stângă, 2014; Acharya, et al, 2014; Greenwood-Nimmo, Huang and Nguyen, 2019; Bettendorf, 2019). Our approach is broader, because in our case CISS reflects systemic rather than just bank-specific risks, while EPU reflects broader policy uncertainty rather than just sovereign risk.

Given the growing interest in uncertainty-related topics, we hope to contribute to this literature by investigating the dynamics of uncertainty arising from the interaction of financial and policy realms, where EA stands, unfortunately, as a fertile ground for research. The remaining of this chapter is organised as follows. Section 2 discusses the theoretical background relevant for our empirical analysis. Section 3 presents the data, along with its sources and limitations. Section 4 provides a detailed overview of the empirical approach, along with its main results and policy implications. Finally, section 5 concludes.

2. THEORETICAL BACKGROUND AND LITERATURE REVIEW

This section discusses the two main literature strands that directly relate to our empirical model. Firstly, we discuss the sovereign-bank nexus, and secondly, financial integration and the role of information frictions as uncertainty sources. The sovereign-bank nexus is important because it explains the interaction between financial and policy realms in a single-country setting. In multi-country settings, however, these theories cannot adequately explain the multiplicity of interactions that exist, for example, between, as well as among, EA sovereigns and EA banking sectors.

2.1. Sovereign-bank nexus

The sovereign-bank nexus, which is defined as the interaction between the financial and policy realms, is one of the main uncertainty sources in economics. What we are most interested in learning about is this very first stage of the uncertainty generating process, where policy and financial uncertainty usually combine and amplify each other, leading to identification challenges in empirical work. Then, once uncertainty arises, it propagates rapidly and inflicts the real sector affecting investment dynamics, asset prices, firms' balance sheets, credit spreads etc., amplified mainly by financial frictions (see among many others, Arellano, Bai, Kehoe, 2010; Christiano, Motto, Rostagno, 2014; Bloom, 2014; Gilchrist, Sim, Zakrajšek, 2014; Bloom et al., 2018).[7]

The main theoretical mechanisms underpinning the feedback loops between banks and sovereigns are best described in Farhi and Tirole (2017), Faia (2017), Leonello (2018), Allen, Carletti, Goldstein and Leonello (2018), Cooper and Nikolov (2018). We briefly summarize the two key mechanisms featuring in these models. On the one hand, as banks hold sovereign bonds in their books for liquidity and regulatory reasons, sovereign distress can contaminate the banking sector. On the other hand, the (implicit or explicit) guarantees provided by the government allow banking sector distress to inflict the public sector. Empirical evidence on these theoretical transmission mechanisms is provided, among many others, in Bicu and Candelon (2013), Stângă, (2014), Bettendorf (2019). While the evidence is clear, in reality there are some nuances one needs to consider. Government commitment to bailing out the banking sector depends on its fiscal capacity and debt dynamics, which explains why EA periphery banks had higher levels of domestic sovereign bonds in their books (Acharya, et al 2014; Koijen et al., 2017; Greenwood-Nimmo et al., 2019). Besides the fiscal costs of a bailout, the central bank can be involved along with the government, in which case there will be inflation and devaluation costs (Farhi and Tirole, 2017).

These theoretical models describing the sovereign-bank feedback-loops are all set within a single-country framework, and therefore cannot be easily extended to a multi-country settings, which is the main focus in this chapter. Difficulties arise from the lack of full political integration across the EA, and in particular the lack of a fully-fledged Banking Union. Recent institutional reforms at the EU level are welcome, although a lack political consensus is hindering further progress in this direction.[8]

2.2. Financial integration and the role of information frictions

For a multi-country perspective, a slight change in focus is in order. Without a fully operational Banking Union or further political integration, the theoretical mechanisms underlying the sovereign-bank nexus do not directly apply at the EA level. Therefore, balance sheet linkages that run through bond holdings

[7] Empirical evidence on these transmission mechanisms is provided in Stock and Watson (2012); Caldara et al., (2016), Alessandri and Mumtaz (2019). Related to this literature strand, Shin (2012), Cerutti, Claessens and Rose (2017) highlight the role of European banks in the transmission of cross-border financial risk spill-overs.

[8] A Single Resolution Mechanism working in conjunction with a Single Supervisory Mechanism (SSM) were recently established (second half of 2014), under the coordination of the ECB, together with competent supervisory authorities from EA Member States. These are two of the three pillars required for the Banking Union to function effectively. The third pillar, i.e. a common deposit guarantee across the entire EA, is still missing, despite ongoing technical discussions and negotiations. Therefore, there is no central authority at the EA level that can automatically provide full guarantees to banks and depositors from different countries. See https://ec.europa.eu/info/business-economy-euro/banking-and-finance/banking-union_en.

and government guarantees are no longer sufficient; instead, cross-border holdings that reflect capital flows across the EA, become part of the spill-overs transmission mechanism. Most importantly, uncertainty stemming from information frictions gathers a more prominent role than in a single-country setting.[9]

European cross-border banking has dramatically increased financial integration as a direct result of the two banking directives adopted in 1977 and 1989 aiming at eliminating restrictions, harmonizing regulation, and achieving better coordination in prudential supervision. Besides the benefits measured in terms of reduced costs and access to financial services, it was hoped that integration would increase the effectiveness of ECB monetary policy and improve its transmission mechanisms.[10] However, theory suggests that financial integration does not necessarily reduce information frictions and might even increase financial fragility.

Freixas and Holthausen (2004) show that integration of the EA interbank market can magnify the asymmetry of information in cross-border banking, creating a contagion channel and financial fragility. Depending on the amount of information frictions, their model allows for multiple equilibria. In particular, the model differentiates between financial *segmentation* and *integration*, where the former relates to a case where all interbank transactions occur within the national borders, liquidity distribution is inefficient and interest rates are higher, while the latter refers to the opposite case. The main theoretical insights from Freixas and Holthausen (2004) are that a segmented market equilibrium is always possible, but an integrated market equilibrium is not necessarily feasible at all times; sometimes, they find that the integrated market equilibrium is not even welfare improving due to increased financial fragility. In fact, more recently, Passari and Rey (2015) conclude that large welfare gains from financial integration, in general, are rather hard to find (in contrast to Allen et al., 2011). According to Freixas and Holthausen (2004), asymmetries leading to market segmentation arise when information remains locally bounded, like in the case of substantial differences in cultures and accounting practices (e.g. policy decisions to restrict risk modelling options for banks), or in local policy preferences with respect to prudential supervision (e.g. commitment to bail out a bank in financial distress). These few examples point to uncertainty sources that originate mainly in the policy rather than the financial realm, although anxieties are likely to arise in both policy and financial circles.[11] In a similar vein, more recently,

[9] Drawing on empirical work, De Grauwe and Ji (2013) advocate for a more active ECB role in counteracting self-fulfilling crises driven by investors' fears, not fundamentals, claiming that EA fragility stems from the lack of a "lender of last resort" for both banks and sovereigns. Their analysis underlines the importance of information frictions in a multi-country settings such as the EA, characterised by advanced financial and economic integration, but not enough political (including fiscal and other policies) integration.

[10] Legislative proposals to advance the integration of European capital markets, along with other segments of the financial market, are high on the policy agenda in Brussels and Frankfurt. Overall, financial integration had positive welfare effects over the first decade of the common currency, as summarized in Allen, Beck, Carletti, Lane, Schoenmaker and Wagner (2011).

[11] Obviously, differences in supervisory treatment should narrow under the newly established SSM framework, where the ECB is the direct supervisor for systemically important EA financial institutions. However, more recent data is needed to evaluate whether this is indeed the case.

Gârleanu, Panageas, and Yu (2015) present a theoretical model where access to financial markets is subject to information frictions, which lead to limited market integration in equilibrium. Moreover, because portfolio diversification (i.e. participation in distant markets) and leverage (i.e. taking more risks) are complements in their model, a symmetric equilibrium might fail to exist, just as in Freixas and Holthausen (2004).

Information frictions, along with asset commonalities, play a key role in other models as well (e.g. Acharya and Yorulmazer, 2008; Allen, Babus and Carletti, 2012). Allen, Babus and Carletti (2012) show that information contagion is more likely in clustered networks, where commonalities in banks' asset portfolios (and structures) are higher.[12] Information contagion refers to bad news about one bank that reveal (to depositors and investors) information about bad realisations of the common factor driving all banks' loan portfolios (and therefore systemic risk). They also claim that banks are 'informationally linked' as long as they use short-term financing, which allows their investors (who cannot clearly dissociate between banks due to opaqueness) to more easily reject rolling over the debt in case of adverse information (i.e. long-term financing would cancel this transmission channel). In Acharya and Yorulmazer (2008), banks undertake correlated investments in order to minimize the effect of information contagion on the expected cost of borrowing. Deep financial and economic integration across the EA make more likely a situation in which banks' loan portfolios share a common systematic factor that explains a higher share of the cross-sectional variation. For example, holding EA periphery versus EA core bonds brought substantial profits for European banks, an investment strategy that Acharya and Steffen (2015) have labelled as "the 'greatest' carry trade ever". These situations point instead to financial information as a potential source of uncertainty, with information frictions playing an amplifying role in this case.

In summary, while each of these theoretical mechanisms has its own merits, there is no clear consensus in the literature on the most relevant ones that can explain such complex, dynamic, double causality influences arising between financial and policy uncertainty within the EA. Starting from this reasoning, our empirical exercise can be seen as an attempt to shed light on these interactions that have important policy implications.

3. DATA

Our dataset focuses on the European region and consists of 24 individual countries and one aggregate, to which we add U.S., as summarised in Table 1 below.

[12] Notice that their *clustered* versus *unclustered* network structures resembles the *integrated* versus *segmented* interbank markets from Freixas and Holthausen (2004).

Table 1: Countries included in the empirical analysis

Euro Area	Other European Union	Others
Austria, AT	Czech Republic, CZ	United States, US
Belgium, BE	Hungary, HU	
Finland, FI	Poland, PL	
France, FR	Sweden, SE	
Germany, DE	Denmark, DK	
Italy, IT	United Kingdom, UK	
Ireland, IE		
the Netherlands, NL	**Other Europe**	
Spain, ES	Norway, NO	
Greece, EL	Switzerland, CH	
Portugal, PT	Turkey, TR	
Luxemburg, LU	Russia, RU	
Slovakia, SK		
Slovenia, SI		
Baltics, BA		

Note: Due to data limitations for specific indicators, we aggregate Latvia, Lithuania and Estonia into a single group, denoted as "Baltics". All indicators pertaining to Baltics are simple averages of available indicators.

The EA is represented here by 14 individual Member States and one aggregate, i.e. the Baltics. With respect to our country selection, some clarifications are in order at this point. Slovenia and Slovakia joined EA in 2007, and 2009 respectively, therefore, very early in the sample and before the European sovereign debt crisis. Although the Baltics joined the EA only recently (i.e. between 2011 and 2015), for the empirical analysis we consider them part of the EA given their small relative size, highly open economies, and the fact that all three have been in the European Exchange Rate Mechanism (ERM II) since mid-2000s – underlining the importance of ECB monetary policy for this aggregate. Regarding other EU member states that are not part of the EA, we include U.K., Denmark and Sweden, along with Czech Republic, Poland and Hungary as three of the most representative countries for Central and Eastern EU with the best data availability.[13] We also include Russia, Turkey, Norway and Switzerland, which are important commercial partners for EU; in addition, each of these countries has some particularities that justifies their inclusion in the sample: Russia is a source of policy uncertainty for Europe, especially during the 2014 annexation of Crimea; Turkey is an important global player in the

[13] Other EU members that are not part of the EA, e.g. Romania, Bulgaria and Croatia, suffer from severe limitations on data availability (i.e. shorter sample availability) for the main model's variables and, therefore, were not included in the analysis. Aggregating these countries is not a feasible option due to their larger heterogeneity than in the case of Baltics.

war against the terrorism that generated the massive immigration influx of 2015; Switzerland is an important financial hub; Norway is an important energy supplier for EU. Finally, we include U.S. as the main global financial centre, and an important source of policy and macroeconomic dynamics relevant for Europe and EA.

Our dataset consists in monthly time-series running from January 2003 to June 2018 (all data description and definitions are provided in Appendix 1 at the end of this chapter). Although CISS is available with a weekly frequency from the ECB data warehouse, EPU are available only with a monthly frequency. We believe that such a frequency is sufficient to uncover the most relevant spill-overs and cross-influences between the financial and policy uncertainty, due to the latter rather complex concept and measurement methodology. All country-specific EPU indexes have been calculated based on the same approach detailed in Baker et al. (2016), who propose searching the databases of major news publications in order to gauge the frequency of some relevant keywords pertaining to economic policy uncertainty domain. Obviously, speculations about un-announced policy changes, intentions or political declarations can be read almost daily in some economic and business publications, but time is of essence in order to observe sufficient political tension that eventually features prominently in the news (and gets captured in the EPU). Considering our country list, EPU time-series[14] are available for the following 11 countries: FR, DE, NL, ES, IT, EL, IE, SE, UK, RU and US. Most importantly, EU countries such as EL, IT, ES, FR, IE and UK, which have been the source of many peculiar events over the last two decades,[15] have both EPU and CISS available, allowing us to apply the identification from De Santis and Zimic (2018), which we discuss in the next section.

Besides the uncertainty proxies EPU and CISS, we include for each country the spread in 10-year sovereign yields against Germany, which is the analytical benchmark for the EA.[16] As a robustness check, we rebase all spreads against U.S., which represents instead the global benchmark. Including bond yields along with uncertainty proxies captures the inherent trade-off between risk and return.[17] Taking bond yield spreads against Germany should wipe out EA-aggregate uncertainty, which would be reflected in the dynamics of the German bond yields, ensuring therefore we indeed capture country-specific dynamics. To further reduce the risk that our results are influenced by aggregate dynamics, the GVAR rich specification allows us to include different measures of area-wide uncertainty computed as weighted averages of EPU and CISS indexes (see the definition of foreign variables in the next section). Besides weighted averages of country-specific EPU and CISS indexes, we include the volatility index

[14] We download all EPU data from www.policyuncertainty.com.
[15] Notice that only Portugal is missing from the list of so-called PIIGS countries.
[16] Data on 10-year sovereign spreads is available for all countries, except Turkey for which we use its 5-year sovereign yield.
[17] According to a recent Bloomberg article, financial investors still prefer high yields, despite high uncertainty stemming from a continuing political struggle between Italy and European Commission over fiscal plans. See https://www.bloomberg.com/opinion/articles/2019-06-12/italy-s-turmoil-means-nothing-to-bond-traders.

VIX[18] – which is a proxy for global risk appetite in the literature on global financial cycles (see Rey, 2015; Bruno and Shin, 2014; Miranda-Agrippino and Rey, 2015) as well as in the literature on global financial spill-overs (Chudik and Fratzscher, 2011; Bettendorf, 2019).

The original idea of the GVAR model specification is the complex re-weighting of country-specific vector autoregressive (VAR) models that reduces the parameters space and makes its estimation feasible (see Pesaran et al., 2004; Dees et al., 2007). To this end, we use a weighting scheme derived from data on bilateral portfolio exposures taken from the IMF's Coordinated Portfolio Investment Survey (CPIS), which includes cross-border investments in bonds and equities.[19] Weights based on portfolio flows, which are less volatile than other capital flows driven by changes in cross-border banking exposures, are more relevant for the model's main transmission mechanisms that reflect risk-return trade-offs across limited integrated markets (see discussion in Gârleanu et al., 2015).[20] Therefore, our specification only indirectly touches on the link between international capital flows and moves in sovereign spreads, i.e. the international portfolios rebalancing channel. According to this literature strand (see Rey, 2015; Bruno and Shin, 2014; Cerutti, Claessens and Ratnovski, 2017), global capital flows co-move with global risk factors and monetary policy changes in centre countries like U.S. and EA. In a similar vein, our empirical specification includes aggregate uncertainty and risk proxies (e.g. VIX, weighted averages of EPU and CISS), bond yield spreads against Germany (or US), and weights based on capital flows. In addition, by amplifying the effects of foreign shocks on the domestic economy, capital flows limit the policy options available to governments (Dragomirescu-Gaina and Philippas, 2015) and/or financial supervisory authorities (Allen et al., 2011), therefore further increasing policy uncertainty.

Due to data limitations for some countries, we use a fixed rather than a time-varying weighting matrix,[21] although the latter would probably only amplify the effects we uncover because of the time-varying profile of contagion among (as well as originating from) vulnerable EA countries. Table 2 below gives an overview on the stability of such portfolio exposures, displaying the average to standard deviation ratios computed over the 2001-2015 time period; lower values of the ratio correspond to more volatile flows, in general, while higher ratios stand for more stable flows. As expected, most EA countries (except EL, SK), together with UK and US have more stable (incoming and outgoing) portfolio flows compared to Eastern EU, Turkey and Russia.

---

[18] VIX is the implied volatility of the S&P500 stock index option prices (the Chicago Board of Options Exchange Market Volatility Index).

[19] Data source is http://data.imf.org/cpis. We average annual data over the 2000-2015 period (subject to availability; some countries, e.g. Baltics, had shorter time-series). The matrix is illustrated in Appendix 2.

[20] A similar weighting scheme based on CPIS data is employed, for example, in Hebous and Zimmermann (2013) and Greenwood-Nimmo et al. (2019), although most GVAR papers use weighting schemes based on bilateral trade flows. Eickmeier and Ng (2015) investigate several weighting schemes (e.g. based on bilateral trade, portfolio investment, foreign direct investment, banking exposures) and find that a combination between trade and financial weights works best to expose credit supply shocks in a GVAR including real and financial variables. See also Feldkircher and Huber (2016) for an analysis of different weighting schemes in GVARs.

[21] Large part of the GVAR literature simply employs fixed rather than time-varying weighting matrixes because the focus is on the interactions of the GVAR variables rather than the weights.

Table 2: Average-to-standard-deviation ratios for portfolio exposures over 2001-2015

Country	as a destination of flows	as a source of flows	Country	as a destination of flows	as a source of flows
EA countries			**Other EU countries**		
AT	2.4	2.4	PL	1.6	1.1
BE	2.1	2.2	HU	1.7	1.1
FI	2.4	2.1	CZ	1.5	1.6
FR	2.5	2.1	SE	2.0	2.0
DE	2.8	2.3	DK	1.7	2.0
EL	1.0	1.2	UK	2.6	2.3
IE	1.6	1.7	**Other Europe**		
IT	2.4	2.7	CH	2.2	2.3
LU	1.8	2.3	TR	1.4	0.6
NL	2.8	2.1	RU	1.6	0.7
PT	1.8	1.8	NO	1.6	1.5
SK	1.1	1.0	**Others**		
SI	1.0	3.7	US	2.5	1.9
ES	1.8	1.9			

Note: The table displays the mean value of these (average-to-standard-deviation) ratios computed over all country-pairs, where the indicated country is a destination or a source of portfolio flows (as mentioned on the first row), therefore, summarizing in a more efficient way a full matrix of statistics where each country pairs with all others. CPIS data is available for all countries from 2001 to 2015; exceptions are Lithuania (data available only for 2009-2015), Latvia (2006-2015), and Slovenia (2009 – 2015).

4. EMPIRICAL APPROACH

4.1. Preliminary data analysis

As the conceptual overlaps between policy and financial uncertainty were discussed in the previous sections, here we provide arguments for their empirical overlaps. As a preliminary analysis, Table 3 below displays the pair-wise correlations between country-specific EPU and CISS indexes, in logs, computed over the entire sample (for countries where EPU is available), at monthly frequencies.

As Table 3 illustrates, with the noticeable exception of U.K., almost all correlations are positive and statistically significant. For France and Ireland, correlations are slightly weaker when CISS lags EPU. The magnitude of the correlations is higher when EPU lags CISS in case of France, Germany, Netherlands, Spain and Ireland, but lower in case of Italy and Greece. We caution the readers not to make any causality inference from these correlations, which lack sufficient robustness and sometimes change with the sample size and period. This lack of robustness, instead, should be interpreted as an illustration of the dynamic nature of the interactions between policy (EPU) and financial (CISS) uncertainty, which might amplify or cancel each other, depending on the period, or the nature of the event that triggered the shock in a particular country. Once we identify the structural shocks from the reduced form residuals, we can investigate the overlapping of the structural shocks' time-series with some well-known episodes that marked the recent history of countries under consideration.

Table 3: Pair-wise correlations between country-specific EPU and CISS indexes

Country	EPU(t) x CISS(t-2)	EPU(t) x CISS(t-1)	EPU(t) x CISS(t)	EPU(t-1) x CISS(t)	EPU(t-2) x CISS(t)
France	0.104	0.1485**	0.194***	0.189***	0.147**
	(0.15)	(0.043)	(0.007)	(0.009)	(0.045)
Germany	0.1208	0.179**	0.2478***	0.218***	0.1637**
	(0.102)	(0.014)	(0.000)	(0.003)	(0.026)
Italy	0.4638***	0.4279***	0.4577***	0.4009***	0.3622***
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Netherlands	0.2355***	0.2392***	0.287***	0.2601***	0.2637***
	(0.001)	(0.001)	(0.000)	(0.000)	(0.000)
Spain	0.2685***	0.3177***	0.4008***	0.3735***	0.328***
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
UK	0.0309	0.0314	0.011	-0.0319	-0.0586
	(0.677)	(0.67)	(0.881)	(0.666)	(0.428)
Greece	0.3345***	0.3177***	0.3108***	0.3078***	0.2835***
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Ireland	0.125*	0.1118	0.1307*	0.1551**	0.1457**
	(0.091)	(0.129)	(0.075)	(0.035)	(0.048)

Note: The effective sample is: 2003:M01 – 2018:M06. The first rows display the lag/lead structure of the two time-series for which we compute the correlations, with t-1, t-2 and t+1, t+2 denoting 1 and 2 period lags, and leads respectively. Both EPU and CISS time series are in log terms. The p-values are provided in parentheses. The *, ** and *** denote statistical significance at 10%, 5% and 1% respectively.

4.2. The baseline GVAR specification with identification based on magnitude restrictions

The global VAR, or GVAR, was designed to simultaneously model cross-sectional dependence and time-series behaviour in macroeconomic data. This very flexible empirical framework was originally proposed by Pesaran et al., (2004), and extended by Dees et al., (2007). In essence, the GVAR is a collection of country-specific VARs, conveniently linked via a weighting matrix that makes the estimation feasible by reducing the parameter space. As discussed in section 3, we use financial weights derived based on IMF CPIS data, which reflect the importance of financial flows in explaining the dynamics of sovereign bond yield spreads, and the transmission of uncertainty spill-overs.

In principle, the GVAR model embeds three channels of cross-country interactions through: (i) foreign-specific variables, (ii) common factors and (iii) contemporaneous dependence of shocks. In this section, we allow for foreign-specific (or so-called star, i.e. *) variables to interact with domestic ones via the first channel, while in the next section, we introduce the second channel that works through common variables (i.e. the ECB monetary policy proxies). The third channel is implicitly accounted for through the estimated variance-covariance matrix in both this section and the next one. As long as the pairwise cross-country correlations left in the model residuals are low, most GVARs in the literature capture the cross-country interactions only through the first two channels, restricting[22] the variance-covariance matrix to be block-diagonal (e.g. Cesa-Bianchi 2013; Eickmeier and Ng, 2015; Feldkircher and Huber, 2016). However, since our focus is specifically on uncertainty spill-overs, we would like to capture the second-order moments of the data as well, and therefore leave the variance-covariance matrix unrestricted in the following analysis.

In the baseline specification, each country i is represented by a country-specific VAR model denoted as VARX (pi, qi), with pi and qi lags and $Y_{i,t}$ a vector of endogenous variables. Each country-specific model is specified as:

$$Y_{i,t} = a_i + \sum_{j=1}^{pi} B_{i,j} Y_{i,t-j} + \sum_{j=0}^{qi} C_{i,j} {Y^*}_{i,t-j} + v_{i,t} \qquad (1)$$

where a_i is a vector of intercepts; $B_{i,j}$ and $C_{i,j}$ are coefficient matrixes; and $v_{i,t}$ is a vector of idiosyncratic shocks, serially uncorrelated and with full variance-covariance matrix. The vector of endogenous variables $Y_{i,t}$ includes domestic variables, while foreign variables are denoted by ${Y^*}_{i,t} = \sum_{i \neq h} w_{i,h} Y_{h,t}$, which are specific to each country i and are constructed as weighted averages of country-specific endogenous variables using the CPIS weighting matrix, W, where for each i we have $\sum_{i \neq h} w_{i,h} = 1$.

For all EU countries, the domestic $Y_{i,t}$ vector includes three variables: EPU, CISS, and 10-year sovereign yield spread against Germany, denoted as $spread$. Obviously, three is the maximum size of the $Y_{i,t}$ vector for EU countries; this happens because for some countries there is no EPU available and, for Germany the sovereign spread is exactly zero, and so it is excluded as a variable. For non-EU countries, the vector $Y_{i,t}$ includes only EPU (for the sake of notation below, we assume there is an EPU available for all non-EU countries) and $spread$, but no CISS because ECB does not compute a CISS

[22] In technical terms, this assumption would amount to a lack of contemporaneous volatility spill-overs between the countries included in the sample, though it would still allow for indirect volatility spill-overs that work through the complex lag structure of the model.

index for these countries. For US, instead, we add VIX, which also serves as a global proxy for risk, in line with much of the existing literature on the determinants of sovereign spreads (see the discussion in section 3). The foreign country-specific vector ${Y^*}_{i,t}$ includes the foreign counterparts of domestic variables, so its size is set: to four for EU countries, to two for non-EU non-US countries, and to one for US. This symmetric (in terms of treating the two uncertainty proxies) but richer specification for EU countries captures the common European policy-making framework (i.e. through EPU^*), and the common financial regulatory framework (i.e. $CISS^*$). Except for US where it is endogenous[23], VIX features in the ${Y^*}_{i,t}$ vector of all countries, along with the foreign sovereign spreads denoted by $spread^*$. In summary, each VARX is specified as:

$$\text{EU countries}^{24}\text{: } Y_{i,t} = \begin{bmatrix} EPU \\ CISS \\ spread \end{bmatrix} \text{ and } {Y^*}_{i,t} = \begin{bmatrix} EPU^* \\ CISS^* \\ spread^* \\ VIX \end{bmatrix}$$

$$\text{Non-EU countries, except US: } Y_{i,t} = \begin{bmatrix} EPU \\ spread \end{bmatrix} \text{ and } {Y^*}_{i,t} = \begin{bmatrix} spread^* \\ VIX \end{bmatrix} \qquad (2)$$

$$\text{US: } Y_{i,t} = \begin{bmatrix} EPU \\ spread \\ VIX \end{bmatrix} \text{ and } {Y^*}_{i,t} = [spread^*]$$

Note that foreign variables are linear combinations of domestic ones, ${Y^*}_{i,t} = W_i Y_t$, with W_i being country-specific link matrices based on CPIS portfolio weights. We can therefore rewrite (1) as:

$$[I, -C_{i,0}] W_i Y_t = a_i + \sum_{j=1} [B_{i,j}, C_{i,j}] W_i Y_{t-j} + v_{i,t}$$

for each country, i. By staking all countries together, we obtain:

$$G_0 Y_t = g_0 + \sum_{j=1} G_j Y_{t-j} + v_t \qquad (3)$$

where $G_0 = \begin{pmatrix} [I, -C_{1,0}]W_1 \\ [I, -C_{2,0}]W_2 \\ \dots \end{pmatrix}$, $G_j = \begin{pmatrix} [B_{1,j}, C_{1,j}]W_1 \\ [B_{2,j}, C_{2,j}]W_2 \\ \dots \end{pmatrix}$, $g_0 = \begin{pmatrix} a_1 \\ a_2 \\ \dots \end{pmatrix}$ and $v_t = \begin{pmatrix} v_{1,t} \\ v_{2,t} \\ \dots \end{pmatrix}$. Provided that G_0 is invertible, we can write the GVAR in its reduced form as:

$$Y_t = h_0 + \sum_{j=1} H_j Y_{t-j} + u_t \qquad (4)$$

[23] Notice that there is no VIX^* because VIX is available only for US, and therefore the two would be identical. Moreover, the simplified specification of the foreign vector for US reflects is in line with much of the GVAR literature, reflecting the prominent (financial and economic) role of the US.

[24] For Germany, $Y_{i,t} = [EPU, CISS]'$, but the ${Y^*}_{i,t}$ is specified the same as for other EU countries.

where $h_0 = {G_0}^{-1} g_0$, $H_j = {G_0}^{-1} G_j$ are coefficients, and $u_t = {G_0}^{-1} v_t$ are reduced form residuals with variance-covariance matrix given by Ω_u.

With all variables expressed in logs (except for spreads), we estimate the model directly in levels, allowing an easy interpretation of impulse responses, which provide us with the main insights. Sims, Stock and Watson (1990) recommend against differencing even in the presence of unit roots, arguing that the goal of the analysis should be to determine the interactions between variables. They show that the VAR specified in levels delivers consistent estimates, even in the presence of stochastic trends and cointegration. Elliot (1998) further shows theoretically that imposing cointegration for near unit root variables can lead to large distortions. We do not estimate cointegrating relations, nor include time trends and error correction terms, also because our short sample and small set of variables would preclude a robust identification of these long-term relationships.[25]

Our sample includes more than 15 years of monthly observations. The main trade-off we are facing in the estimation is between model parsimony and its statistical properties (e.g. stability, residual tests). Kapetanios et al. (2007) notice that the quality of a VAR approximation to the true model depends on both the number of variables and the lag order; as the GVAR includes more variables than a normal VAR (i.e. both domestic and foreign variables in each country-specific model), small lag orders are regularly employed. We notice that setting a maximum lag length for domestic variables $pi = 3$ eliminates most residual autocorrelation (or serial dependence) and preserves a parsimonious specification (i.e. setting a higher lag order would further reduce autocorrelation). As for the maximum lag employed for foreign variables, qi, a smaller lag order is to be preferred because financial markets can react rapidly to foreign influences (e.g. media news, uncertainty boosts); in fact, accounting for the contemporaneous effects of foreign uncertainty proxies is key for estimating the cross-border spill-overs. Setting the maximum lag length $qi = 0$ for all country-specific models as in Burriel and Galesi (2018) does not guarantee model stability (i.e. all eigenvalues below unity) in all different specifications; therefore, allowing for $qi = 1$ in few specific VARX models, particularly for small countries (that are more likely to receive heavier influences from abroad, like for example AT, BE, IE, DK), appears the easiest fix to this stability problem and, in addition, maintains model parsimony and lowers autocorrelation further. Figure 1 depicts the estimated residual autocorrelation and the eigenvalues of the GVAR in the baseline specification. The large majority of residual autocorrelations lie within or close to the confidence bands (± 2 standard deviations), and all eigenvalues are less than one, despite some inherent persistency.

[25] Both theory and empirical studies provide evidence that European sovereign spreads are cointegrated with fundamentals (e.g. fiscal proxies, economic and financial proxies), which are omitted from our estimated GVAR (see De Santis, 2019).

Figure 1: Specification checks for baseline GVAR

Panel A: Residual autocorrelation

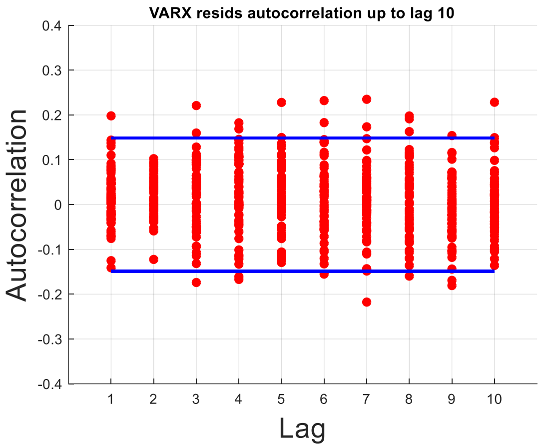

Panel B: Eigenvalues of GVAR

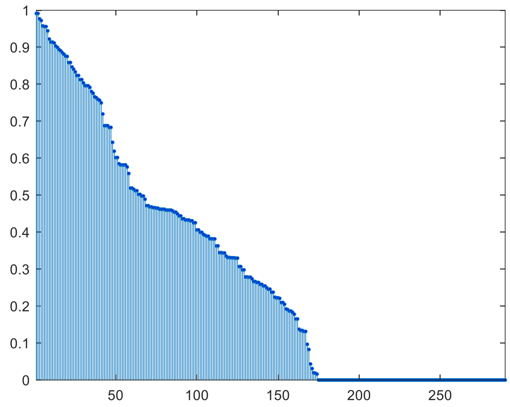

Note: Panel A plots the values of residual autocorrelation for all GVAR variables and all country-specific models, depending on the serial lag; the vast majority of them (96.5%) are lying within the indicated confidence bands ($\pm$2 standard deviations). Panel B plots the GVAR eigenvalues, all lying below unity.

As noted in Dees et al (2014), and also Dungey and Osborn (2014), dealing with multi-country models in general requires a different framework for conceptualizing the nature of shocks that one wishes to identify, particularly because of the strong cross-sectional dimension implied in such models. This is one of the contributions we bring to the uncertainty-related empirical literature, which deals largely with shock identification in single country models (noteworthy exceptions are Bicu and Candelon, 2013; Stângă, 2014; Acharya, Drechsler, and Schnabl, 2014; Bacchiocchi, 2017; Greenwood-Nimmo, Huang and Nguyen, 2019; Bettendorf, 2019). A GVAR specification can elegantly solve such challenges through the inclusion of country-specific foreign variables (and common variables) that can effectively reduce, and even eliminate, cross-sectional correlations in residuals. In our baseline specification, the average cross-sectional correlation is below 0.04 in absolute terms for CISS and EPU, and below 0.05 in absolute terms for spreads, with a maximum of 0.13 for some countries (e.g. FI, PL, ES, CZ). To better illustrate this point, not including the country-specific foreign vector ${Y^*}_{i,t}$ would rise all these cross-sectional correlations to within the 0.2 – 0.4 range.

In terms of identification, we follow De Santis and Zimic (2018) and implement structural identification through absolute magnitude restrictions. Any structural identification requires a mapping from reduced-form shocks, u, into structural ones, ε, say in the form: $u = S\varepsilon$, where S is a matrix that is the focus of any identification strategy. In practice, we are only interested in the identification of the two uncertainty shocks associated with the two uncertainty proxies and therefore with a partition of S that we denote as $S_{2\text{x}2}$. The GVAR estimated variance-covariance matrix associated with the first two equations in this case becomes $\Omega_{u,2\text{x}2} = S_{2\text{x}2}S_{2\text{x}2}{}'$, since the structural shocks ε are normalised and assumed to have unit variance. The main identification challenge is the lack of uniqueness for the matrix

$S_{2\text{x}2}$. In particular, for any orthonormal matrix, K, satisfying $KK' = I$, we can also write $\Omega_{u,2\text{x}2} = S_{2\text{x}2}K(S_{2\text{x}2}K)' = S_{2\text{x}2}KK'S_{2\text{x}2}{}' = S_{2\text{x}2}S_{2\text{x}2}{}'$ meaning that $S_{2\text{x}2}$ is not uniquely identified from the data without some additional assumptions.

Magnitude restrictions work by conveniently restricting the space where the coefficients of the $S_{2\text{x}2}$ matrix are required to lie, based on the simple assumption that the relative size of the contemporaneous response of uncertainty variable i to an uncertainty shock j, with $i \neq j$, must be smaller (in absolute terms)[26] than the contemporaneous response of uncertainty variable j to the same uncertainty shock j. In other words, when both variables i, j are scaled by their standard deviations, the indirect effect of a structural uncertainty shock ε_j on variable $i,\ i \neq j$, is lower than its direct effect on variable j. To some extent, these restrictions imply that any of our two uncertainty measures is better than the other one in capturing a structural shock that stems from its own domain – an implication that is not hard to accept given the obvious methodological differences between the two indicators. Indeed, despite the inherent statistical overlaps, CISS is a composite indicator designed and empirically tested (see Hollo et al., 2012) to reflect stress in different financial market segments rather than Knightian uncertainty, while EPU is designed to capture policy uncertainty as reflected in the media and related to government's initiatives, public proposals, or changes in rhetoric and opinions rather than financial stress.

De Santis and Zimic (2018) show that these simple inequality restrictions allow the unique identification of structural shocks in sovereign yields within a simple VAR focusing on EA. We rely on their proposed algorithm for small systems (see the Appendix from De Santis and Zimic, 2018), as we only require the identification of two structural shocks where convergence problems are not an issue. There is no particular difference between the working of the algorithm in a VAR settings compared to a GVAR one, apart from its different Matlab implementation and additional coding required into the GVAR toolbox, which is made available in Burriel and Galesi (2018). Appendix 3 at the end of this chapter details the main steps of the algorithm as implemented in our GVAR specification.

The advantages of using magnitude restrictions in our empirical setting are important to discuss in relation to other structural identification methods available in the VAR literature.[27] Firstly, the identification through magnitude restrictions does not impose any time precedence on the two uncertainty variables, like would be the case when applying a standard Choleski identification (which is just a special case of the identification based on magnitude restrictions as it imposes a zero contemporaneous response of some variables to some shocks).[28] In our case, imposing a time

[26] This means that the two uncertainty variables are allowed to move contemporaneously in any direction in response to a structural shock, as along as the relative (measured in terms of standard deviations) impact fulfils the respective inequality.
[27] It is important to mention that most of the GVAR literature uses generalised IRFs (or GIRFs) due to identification challenges in multi-country settings, as discussed in Dees et al. (2014). The GIRFs, however, have the main disadvantage that shocks cannot be given a structural interpretation.
[28] Bekaert et al. (2013) estimate a VAR specified in business cycle, monetary policy, risk aversion and expected market volatility, using a Choleski decomposition (with variables ordered as listed), and a combination of

precedence between two uncertainty proxies would be too strong of an assumption, given the complex, dynamic, double causality influences between policy and financial uncertainty (e.g. in the cases of Greece and Ireland the precedence of the shocks is reversed; see Farhi and Tirole, 2017). Secondly, an alternative identification method such as one based on sign restrictions would require not only more complex transmission mechanisms than the one implied here, but also strong theoretical predictions about these mechanisms. As long as the literature on measuring uncertainty and estimating its effects is still in its infancy, a perfect matching between theoretical concepts and empirical counterparts is challenging (see discussion in Jurado et al., 2015). Moreover, as noted in Caldara et al., (2016), different uncertainty shocks, despite differences in measurement, can have similar effects on other macroeconomic variables, complicating identification.[29] Thirdly, Bacchiocchi (2017) and Angelini et al., (2019) build on the original "identification through heteroskedasticity" idea proposed in Rigobon (2003) in order to identify uncertainty shocks in a VAR model. While their method is successful in dealing with endogeneity challenges that arise between uncertainty and real or financial variables, it requires that (at least some) structural parameters remain constant over time and across volatility regimes. Forth, Caldara et al., (2016) identify the effects of economic uncertainty and financial shocks by employing a penalty function approach, which shares some similarities with our identification approach. In their case, the structural shock should maximize the impulse response of its respective target variable over a pre-defined period. However, although they are able to identify the two structural shocks, they still use a sequential identification due to reverse causality fears.

To derive our main insights, we rely on the models' impulse response functions (IRFs), which are conveniently summarised in Table 4 below. To gauge statistical significance of the IRFs, we use bootstrapped 68% confidence intervals[30] based on 200 replications, with (a maximum of) 200 draws of the orthonormal matrix for each replication. In Appendix 4, Figures 4.1 – 4.6, we display detailed IRFs for all country-specific variables to uncertainty shocks originating in Italy, Spain, Greece, Ireland, France and UK, for which we have data on both EPU and CISS uncertainty proxies. Due to the large number of countries, and to save space, we only display the most representative ones in Appendix 4.

contemporaneous with long-run restrictions. They find that risk aversion decreases more strongly than volatility to a lax monetary shock, with both expected volatility and risk aversion extracted from VIX. Others like Baker et al. (2016) and Jurado et al. (2015) also employ Choleski decompositions, but use a single uncertainty proxy not two different ones.

[29] As the required inequality restrictions must be fulfilled only in absolute terms in our case, EPU and CISS are free to either co-move or move in opposite directions, and they might have similar effects on bond spreads.

[30] Burriel and Galesi (2018) and Anaya, Hachula and Offermanns (2017) also use 68% confidence intervals in GVAR applications, which are known to suffer from wider confidence bands due to over-parameterization.

Table 4: Summary findings based on IRFs to shocks in the baseline specification

Shock origin	Observed response	CISS responses to CISS shock	EPU responses to CISS shock	CISS responses to EPU shock	EPU responses to EPU shock
Italy	**domestic**	Significant, up to 9 months	Significant and quick, up to 6 months	Insignificant	Significant, up to 12 months
	cross-border spill-overs	Significant, up to 9-18 months	Significant and quick, up to 3-12 months	Insignificant	Significant, up to 9-12 months
Spain	**domestic**	Significant, up to 9 months	Insignificant	Insignificant	Significant, up to 6 months
	cross-border spill-overs	Significant, up to 9-24 months	Weakly significant in few countries	Mostly insignificant	Significant, up to 6-9 months
Greece	**domestic**	Significant, up to 12 months	Insignificant	Insignificant	Significant, up to 12 months
	cross-border spill-overs	Significant, up to 6-18 months	Short-lived, significant up to 9 months	Significant, up to 6-12 months	Significant, up to 6-12 months
Ireland	**domestic**	Significant, up to 18 months	Significant between 3-9 months	Insignificant	Insignificant beyond impact period
	cross-border spill-overs	Significant, up to 9-18 months	Significant and quick, between 3-12 months	Insignificant	Insignificant beyond impact period
France	**domestic**	Significant, up to 9 months	Insignificant	Insignificant	Significant, up to 9 months
	cross-border spill-overs	Significant, up to 9-18 months	Short-lived, significant up to 9 months	Insignificant	Significant, up to 9-12 months
UK	**domestic**	Significant, up to 9 months	Short-lived, significant up to 6 months	Insignificant	Significant, up to 9 months
	cross-border spill-overs	Significant, up to 9-24 months	Short-lived, significant up to 6-9 months for some countries	Significant, up to 3-9 months	Significant, up to 6-12 months

Note: The table displays a summary of the IRFs results derived for structural uncertainty shocks in the baseline GVAR specification. Statistical significance is based on bootstrapped 68% confidence intervals based on 200 replications, with a maximum of 200 draws of the orthonormal matrix for each replication.

Some main results from Table 4 stand out. Firstly, there are substantial and persistent cross-border spill-overs for both EPU and CISS shocks. Depending on the shock's country of origin, CISS or EPU might display more or less persistency in responding to own shocks. These observations are in line with the prevalence of different shocks in different countries and the narrative evidence available. For example in Spain and Ireland, CISS displays more persistent responses to its own shocks (i.e. mostly financial crises), while in Italy it is the EPU that displays more persistency to its own shocks (i.e. mostly political crises).

Secondly, there are substantial interactions (or cross-influences) between the two uncertainty proxies, revealing the important overlaps existing between policy and financial realms. In particular, most EPU responses to CISS shocks are significant and quick, lasting between impact and 12 months after the shock. On the contrary, CISS responses to EPU shocks, even when statistically significant, are only very weak and slow. In other words, policy uncertainty responds strongly and quickly to shocks in financial uncertainty, but financial uncertainty in general does not react to policy uncertainty shocks. This asymmetry is surprising given the perfectly symmetric treatment of the two uncertainty proxies in the specification of EU countries' VARX models (see equations 2). This is an important result in line with the existing evidence of the importance of financial frictions for (macro)economic policy stability (Allen et al., 2011).

Thirdly, there is more often the case that domestic interactions between EPU and CISS are insignificant, while their cross-border spill-overs are significant, like in the cases of Greece, France and Spain; this could mean that domestic events in these countries have been more relevant from a European rather than a domestic perspective.

A fourth result, not summarised in Table 4 to save space but easily revealed when inspecting Figures 4.1 – 4.6 panels C and F, in Appendix 4, refers to sovereign spreads' reaction to uncertainty shocks; these reactions are consistent with international portfolios rebalancing away from (mainly) EU periphery (though France, Netherlands, Baltics, and Eastern Europe can be also included here) and into US and UK bonds (along with other safe assets like Luxemburg, Norway).[31] Moreover, spreads' reactions to CISS shocks point to overshooting followed by undershooting their initial level within a two-year period, while reactions to EPU shocks are mostly positive with a slow return to the initial level over the same period of time. In other words, higher financial (though not policy) uncertainty is more likely to be associated with higher volatility in sovereign yields.

[31] Interestingly, and perhaps counter-intuitively, sovereign spreads of UK and Ireland decrease in reaction to domestic CISS shocks (and EPU shocks for UK), highlighting the European rather than local dimension of the stress events, and the two countries' role as financial hubs for European markets.

4.3. Extended GVAR specification: including ECB monetary policy

In this section we include ECB monetary policy, such that ECB can be interpreted as a synthetic country of the GVAR[32], as in Georgiadis, (2015), Burriel and Galesi (2018). Accordingly, we re-specify the reduced-form GVAR given in (1) as:

$$Y_{i,t} = a_i + \sum_{j=1}^{pi} B_{i,j} Y_{i,t-j} + \sum_{j=0}^{qi} C_{i,j} Y^{*}{}_{i,t-j} + \sum_{j=0}^{qi} D_{i,j} X_{t-j} + u_{i,t} \qquad (5)$$

where X_t includes the common variables that represent the ECB monetary policy, while $D_{i,j}$ is the associated coefficient matrix. We follow Boeckx et al. (2017) and Burriel and Galesi (2018) and define X_t such as to capture the main aspects of ECB monetary policy: (i) conventional monetary policy (CMP) and (ii) unconventional (UMP) policy tools, as well as (iii) a liquidity proxy; see data description in Appendix 1. We proxy CMP using the Main Refinancing Operations interest rate, which is the ECB main policy rate. As a liquidity proxy we use the spread between EONIA (i.e. the Euro Overnight Index average) and the Main Refinancing Operations rate. As UMP proxy we use the (log of) ECB balance sheet, which is the standard indicator in the literature. These same three indicators define ECB monetary policy in Boeckx et al., (2017) and Burriel and Galesi (2018), though here our focus is on evaluating ECB responses to uncertainty shocks. Accordingly, the extended GVAR specifies X_t as an autoregressive process with lag orders given by (px, qx) as:

$$X_t = m_x + \sum_{j=1}^{px} N_j X_{t-j} + \sum_{j=1}^{qx} P_j \tilde{Y}_{t-j} + u_{x,t} \qquad (6)$$

where N_j and M_j are (matrix) coefficients, while $\tilde{Y}_t$ is a vector of feedbacks from GVAR's endogenous domestic variables – capturing the response of ECB monetary policy to developments in the EA region (see Table 1), similar to Georgiadis (2015), Burriel and Galesi, (2018). As lag orders, we chose a parsimonious model matching the baseline specification and set $px = 3$ and $qx = 1$.

Central banks' balance sheets have recently attracted much research attention, with Gambacorta, Hofmann, and Peersman (2014) being the first to popularise the use of central bank balance sheet size as a proxy for unconventional monetary policy.[33] Starting with Fratzscher et al. (2016), there is a growing list of studies focusing on ECB's unconventional monetary policy effects and their cross-country spill-overs; see among many others Moder (2017), Burriel and Galesi (2018), Boeckx, Dossche, and Peersman (2017), Kucharčuková et al. (2016), Koijen et al., (2018). Moreover, as noted in Boeckx,

[32] In technical terms, equation (6) describes the dynamics of the dominant unit of the GVAR.

[33] ECB unconventional monetary policy tools include non-standard liquidity provision (LTROs) and several asset purchase programs. For example ECB conducted two Covered Bond Purchase Programs (CBPPs) between June 2009 and October 2012; between 2010 and 2012, ECB bought government bonds on the secondary market through its Securities Markets Program (SMP). After November 2014, ECB conducted an asset-backed securities purchase program (ABSPP) and a third CBPP.

et al., (2017), identification of unconventional policy shocks can be challenging if it does not sufficiently distinguish between policy- and demand-driven shocks. Given the fixed-rate tender with full allotment strategy of the ECB,[34] isolating the exogenous from the endogenous shifts in ECB's balance sheet becomes key for proper identification (see also Burriel and Galesi, 2018). We do not attempt to identify conventional nor unconventional monetary policy shocks in our simple model, as long as some relevant transmission mechanisms and indicators (e.g. real business cycle indicators) are missing. Expanding the GVAR specification to include a more detailed set of indicators and transmission mechanisms is beyond the scope of this paper and is left for future research.

The estimation of the extended GVAR proceeds as already described in the previous section. For consistency, we impose the same lag structure for the country-specific VARXs as in baseline, and observe this specification successfully passes stability and residual autocorrelations tests. In addition, Table 5 below presents, by country and equation, the F-statistics of the joint null that $D_{i,j} = 0$ in equation (5). Results are rather mixed, illustrating that ECB variables had jointly varying influences on each country; for Italy and Greece the impact of ECB proxies falls mostly on policy uncertainty, while for Portugal, Baltics and Central Europe it falls on financial uncertainty; sovereign spreads are instead affected by ECB in the cases of Greece, Spain, and some Central European countries.

Table 5: F-tests of the joint null that ECB policies had no influence on a EA country/variable

Country:	**Variable: Spread**	**Variable: CISS**	**Variable: EPU**	**F crit. (5%)**
Austria	3.6786*	0.4309		2.1539
Belgium	1.6759	1.1428		2.1539
Finland	1.6692	0.9005		2.6571
France	1.7233	1.2541	0.9965	2.1549
Germany		1.7696	1.8607	2.6571
Italy	1.0234	0.9305	4.3796*	2.1549
Netherlands	2.0572	0.7916	2.2181	2.6581
Spain	8.2791*	2.4232	0.7440	2.6581
Greece	3.0175*	0.9094	3.9638*	2.6581
Ireland	0.1110	0.9401	1.5052	2.1549
Portugal	1.9515	3.4332*		2.1539
Luxemburg	7.1009*	2.4442		2.6571
Slovakia	3.8868*	5.2872*		2.6571
Slovenia	8.8777*	3.8107*		2.6571
Baltics	0.6399	3.0100*		2.6571

Note: The F-statistics is computed, by country and by equation, restricting the coefficients of the three ECB proxy variables to be exactly zero. The degrees of freedom associated with the F-statistics vary, depending on the lag structure of the foreign variables in each country-specific VARX model. The * denotes cases where the F-statistics is higher than the critical value, and therefore the null can be rejected at a 5% significance level.

[34] https://www.ecb.europa.eu/stats/policy_and_exchange_rates/key_ecb_interest_rates/html/index.en.html.

Again, our main insights are derived based on the analysis of the IRFs to structural uncertainty shocks identified using magnitude restrictions. This time, however, the extended GVAR specification allows us to analyse more results with a richer dynamics, including the different ECB monetary policy responses to different uncertainty shocks. Detailed IFRs for some of the most representative countries can be found in Appendix 4, in Figures 4.1 – 4.6, and in Appendix 5 in Figure 5.1.

The most striking difference between the baseline and the extended GVAR is the reduction in persistency, and in some cases statistical significance, for uncertainty responses to uncertainty shocks. Most EPU reactions to either CISS or EPU shocks become less statistically significant compared to the baseline, and this happens particularly in EA periphery; meanwhile, CISS reactions remain broadly unchanged. In other words, with the inclusion of ECB in the extended GVAR, EPU dynamics become more muted – an idea that matches the positive views held on ECB policies (particularly UMP) that dominate policy discussions in Italy for example, and corroborates with findings from Table 5 above. As more significant differences arise for EPU responses rather than for CISS responses, we can conjecture that ECB actions must have had a more important impact on policy uncertainty dynamics rather than financial uncertainty.

The extended specification allow us to infer the responses of ECB monetary policy to structural uncertainty shocks as well, providing a complementary perspective to our results discussed above (see Appendix 5). Firstly, ECB assets temporarily increase (for about 6 months) in reaction to CISS shocks, while reactions to EPU shocks last longer and are not easily reversed. In other words, ECB might be just accommodating higher liquidity needs rather than adopting new (un)conventional policy measures in reaction to CISS shocks. This conclusion is reinforced by the fact that increases in the spread between EONIA and the ECB main policy rate, which are indicative for liquidity shortages, are more intense in reaction to CISS shocks than to EPU shocks. Secondly, the response of ECB's policy rate to CISS shocks is almost immediate, always negative and statistically significant for about 12-18 months. In contrast, responses to EPU shocks are more sluggish and weaker. Overall, these new insights only serve to strengthen our conjecture above: that ECB had adopted a more pro-active stance towards policy uncertainty shocks, and a more accommodative (i.e. passive) stance towards financial uncertainty shocks. In line with the theoretical literature mentioned in section 2.2 (e.g. Freixas and Holthausen, 2004), we might posit that ECB has tried to reduce policy uncertainty in order to prevent a "segmentation equilibrium", or at least ensure that an "integration equilibrium" remains feasible.

Lastly, with respect to sovereign spreads' reactions to structural uncertainty shocks, our results point to some similarities compared to the baseline. The IFRs in panels C and F from Figure 4.1 – 4.6 in Appendix 4 illustrate the same idea consistent with portfolio rebalancing away from EA periphery and mainly into US and UK, i.e. outside EA. However, this time around the IRFs in the extended GVAR specification show a larger undershooting occurring between 9-24 months in response to CISS shocks, highlighting the positive role ECB must have played on EA sovereign yields, and therefore on CISS dynamics, over a medium-to-long term.

4.4. Event study back-testing

According to De Santis and Zimic (2018), their magnitude restrictions are inspired by event study techniques. This type of techniques require a good understanding of the historical patterns and causality chains that facilitate the identification of the consequences of some unique shocks. As already mentioned before, magnitude restrictions are just a mathematical formulation that conveniently associates a causality order among some highly correlated variables. This section provides evidence on the suitability of our identification approach by associating an event timeline with the time-series of the identified structural uncertainty shocks. We draw on various data sources to identify a set of unique country-specific events, many of whom have shaped the recent history of Europe. Figure 2 plots the identified structural CISS and EPU shocks for Italy, Spain, Greece, Ireland, France and UK, along with a timetable that provides details on the most likely event that can be associated with some outliers.

Figure 2: Time-series of structural uncertainty shocks and timetable of major events and headlines

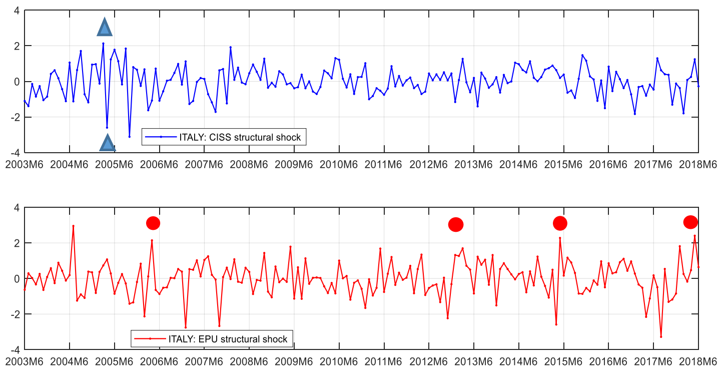

<u>Italy timeline</u>:

▲ Mar. - Apr. 2005: ABN Amro bids for Antonveneta in a first cross-border acquisition of an Italian bank. BBVA bids for BNL, another Italian bank. Italian authorities oppose, but European Commission insists.

● Apr. 2006: Berlusconi squares off against Prodi in general elections. The Italian Supreme Court confirms Prodi's election win, but Berlusconi doesn't concede

● Jan. 2013: Monti's gov't resigns

● Apr-May 2015: EU Migration crisis begins with a tragedy in the Mediterranean, with Italy taking centre stage

● May 2018: Anti-system parties M5 and Lega Nord win general elections and will form the gov't

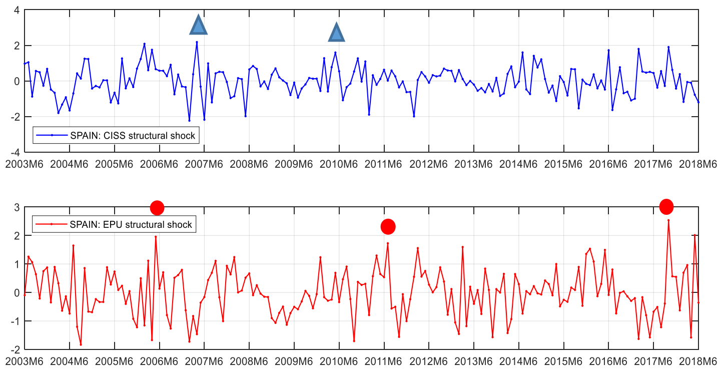

<u>Spain timeline</u>:

● May-Jun. 2006: Campaign and referendum on Catalonia autonomy.

▲ Apr. 2007: Spanish bank Santander bids with a consortium to acquire ABN Amro, in the world's biggest bank takeover to date. Markets fear financial troubles due to oversized debt.

▲ May 2010: Bank of Spain seizes CajaSur, a savings bank run by the Roman Catholic Church, making investors worried about the ongoing restructuring of the Spanish banking sector

● Mar. 2012: Strikes against austerity and labour market reform spread throughout Spain.

● Oct 2017: Madrid imposes direct rule in Catalonia after voters in a referendum back separation from Spain.

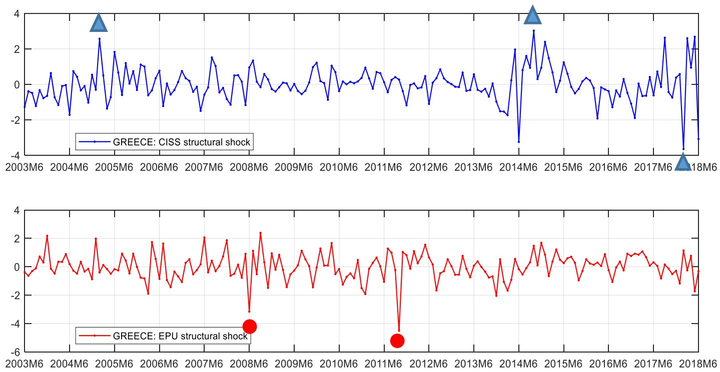

<u>Greece timeline:</u>

▲ Feb. 2005: EU finance ministers warned Greece to get its finances in order and bring its annual budget deficit in line with EU spending rules or face hefty fines

● Jun. 2008. Greece ratifies the Treaty of Lisbon.

● Oct. 2011: Eurozone leaders agree on an exceptional package of measures to address crisis, including a 50% debt write-off for Greece in return for further austerity measures.

▲ Oct. 2014: concerns over Greek gov't collapse triggers massive sell-off in the stock and bond markets; main stock index down 9.8 %. EBA stress test results published, showing 3 Greek banks failed the test.

▲ Feb. 2018: gov't eases capital controls from 2015. Greece returns to international markets with a 7-year bonds auction

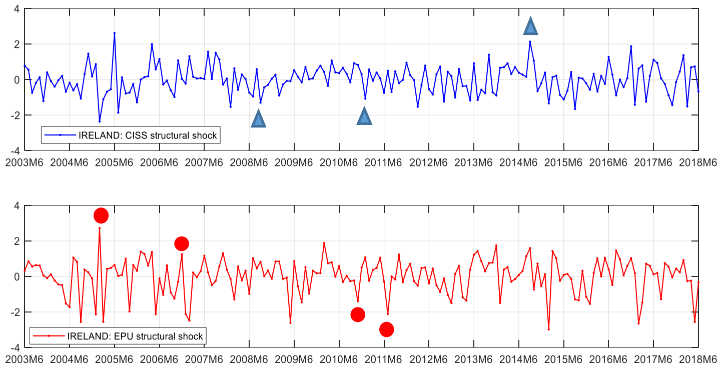

<u>Ireland timeline:</u>

● Feb. 2005: Northern Bank robbery investigation reveals money laundering and leads to political scandal; Gov't accuses top members of Sinn Fein are part of IRA leadership; Bank of Scotland chairman, associate with prime-minister, steps down.

● Dec. 2006: Irelands' biggest political scandal related to former Prime-Minister Haughey

▲ Sep 2008: Gov't to guarantee all deposits in Ireland

● Nov 2010: European ministers agree a bailout for Ireland

▲ Dec 2010 – Jan 2011: Gov't nationalizes Allied Irish Banks, the 4th bank taken over in the crisis.

●Jul 2011: Jul 21, EU leaders cut interest rates on Irish bailout

▲ Sep 2014: EU warns Ireland that the country had granted Apple tax advantages as illegal state aid.

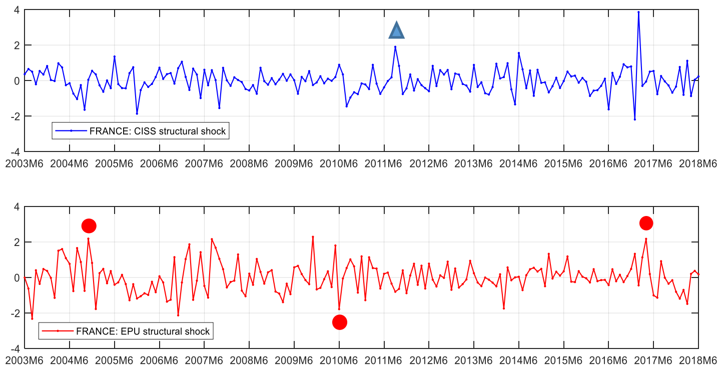

<u>France timeline</u>:

●Nov 2004: an armed conflict starts between France and Ivory Cost. Nearly 5000 foreign nationals evacuated.

●Jun 2010: Gov't announces huger public spending cuts of 45 bn. EUR to reduce public debt.

▲ Sep 2011: Moody's downgrades the two biggest French banks, Credit Agricole and Societe Generale, because of their exposure to Greek debt.

●Apr. 2017: First round presidential elections between E. Macron and M. Le Pen.

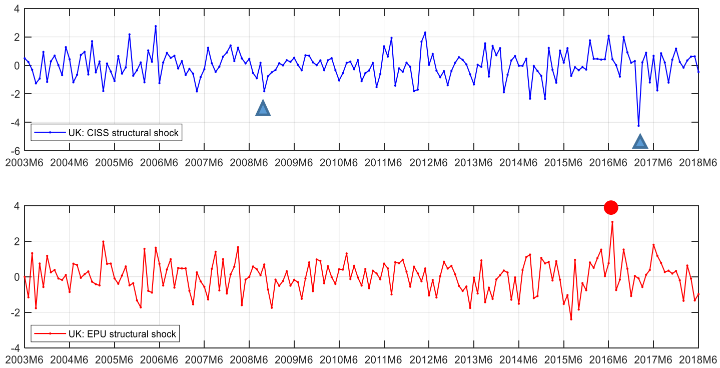

UK timeline:

▲ Oct 2008: Gov't increases to GBP 50,000 its guarantees of British bank deposits. British gov't bails out several banks, including the Royal Bank of Scotland, Lloyds TSB, and HBOS.

● Jul. 2016: dire policy implications emerge regarding the surprising results on the UK referendum for EU membership, held at the end of June. PM David Cameron resigns.

▲ Feb 2017: Parliament votes to invoke Article 50. Gov't publishes its first blueprint on Brexit, removing some of the existing uncertainty.

Note: News and headline sources include the following webpages: www.FT.com, www.bbc.com/news, www.cnn.com, www.wikipedia.org, www.imf.org, www.timelines.ws, https://ftalphaville.ft.com, https://en.wikipedia.org/wiki/Portal:Current_events.

4.5. Robustness checks

Three different robustness checks are performed in order to verify the consistency of our main findings. We next provide the main technical details behind their implementation and discuss the main results in comparison with the results of the baseline and extended specifications.

As a first robustness check, we add a measure of global liquidity risk, i.e. the TED spread, which is the spread between the 3-Month LIBOR based on US dollars and the 3-Month Treasury Bills (see Brunnermeier, 2009). Although our extended GVAR already includes a liquidity proxy relevant for EA markets (i.e. the spread between EONIA and the main ECB policy rate), the US dollar-denominated funding costs of European banks play a key role within the literature on global financial cycles[35] (see Rey, 2015; Bruno and Shin, 2014). When uncertainty raises, banks charge themselves higher interest rates for uncollateralised loans (i.e. LIBOR rate is the reference rate for interbank lending) compared to the yield of risk-free US Treasuries, and therefore the TED spread is actually a global liquidity proxy. The cost of US dollar funding has been a central element of the policy reactions during the peaks of the financial crisis from 2007/2008. All major central banks, including ECB, set up direct currency swap lines with the US Federal Reserve System, precisely to alleviate pressures from the US dollar funding.[36] By adding the TED spread as an endogenous variable to the US model in our GVAR, we account for changes in global liquidity and US dollar funding, providing a consistency check to our main findings from the previous sections. We find that the main results from sections 4.2 and 4.3 remain qualitatively unchanged.

[35] This is because the US dollar is the world's main reserve currency. Similar to the global financial cycle literature, the international bank lending channel, exposed in Schmidt, Caccavaio, Carpinelli and Marinelli (2018), highlights the importance of US dollar funding costs on lending in Europe, particularly in France and Italy.
[36] See https://www.federalreserve.gov/monetarypolicy/bst_liquidityswaps.htm.

As a second robustness check, we compute the sovereign spreads against the 10 year U.S. sovereign yield, which is the global benchmark, rather than against Germany, which is the European benchmark. The only technical change required in the model specification given by equation (2) is that the German VARX now includes the sovereign spreads against US, while the US VARX includes only EPU and VIX as endogenous variables. Main findings are again qualitatively unchanged.

In a third robustness check, we re-estimate the extended GVAR specification with a different weighting matrix based on BIS Locational Banking Statistics (LBS) data. Appendix 2 provides more details on the constructions of weights in this case. The GVAR estimated in Eickmeier and Ng (2015) fits the data better when using weights based on BIS LBS banking exposures for financial variables (along with trade weights for their model's real variables). Such weights based on BIS LBS data are also employed in Bicu and Candelon, (2013), Feldkircher and Huber (2016) among others. Yet, capital flows driving bank cross-border exposures are generally more volatile than flows driving portfolio exposures according to balance of payments statistics, and our statistical evidence points to the same result. Despite some important differences in weighting between IMF CPIS data and BIS LBS data[37], estimating the extended GVAR specification with weights based on the latter dataset delivers qualitatively similar results.

5. CONCLUSIONS

From a theoretical perspective, the sovereign-bank nexus entails strong feedback loops between financial and policy realms within a single-country setting. Instead, in the case of Euro Area, with its rather incomplete institutional architecture, more complex interactions between the two realms can be expected. This is especially so during uncertain times, since EA lacks institutional leadership to deal with several uncertainty sources, either financial or policy-related. Moreover, as the European financial markets swing between integration and fragmentation with each passing crisis, it becomes highly relevant to investigate the sources of various information frictions, the spill-over potential of country-specific uncertainty shocks, and the specific role that a key EA institution such as the ECB might play in mitigating the effects of such shocks.

We approach these relevant questions from an empirical perspective, employing a GVAR specification that efficiently summarises the time-series dynamics of our multi-country dataset focused on the Euro Area. We use the CISS composite indicator, proposed by Hollo et al., (2012), as a proxy for financial uncertainty and the EPU index, proposed in Baker et al. (2016), as a proxy for economic policy uncertainty. Despite substantial correlations, the methodological and data set differences

[37] For example, the IMF CPIS data show that most countries in our sample have out-weighted exposures towards US, UK and LU, in this order. Instead, according to BIS LBS data, most countries have out-weighted exposures towards UK and LU.

between our uncertainty proxies allow us to use a recently proposed structural identification approach based on magnitude restrictions as in De Santis and Zimic (2018). One of the main contributions we bring to the uncertainty-related literature is therefore to identify both financial uncertainty and policy uncertainty shocks in a multi-country setting, allowing for a variety of contemporaneous spill-overs. In a similar vein, Caldara et al. (2016) employ a penalty function approach, but are not able to deal with reverse causality issues between the two uncertainty shocks they uncover – a weakness that we avoid with our approach. We further discuss the main advantages of our identification over other methods available in the relevant literature. In a convincing proof, our identified structural shocks match the dates of some remarkable events that marked the recent history of some European Union countries.

The empirical findings confirm there are statistically significant and persistent effects for both financial- and policy-driven uncertainty shocks. Besides domestic effects and cross-influences shaped according to the prevalence of uncertainty shocks in each country, cross-border uncertainty spill-overs are also statistically significant. In terms of cross-influences between the two uncertainty proxies, we find that policy uncertainty reacts stronger to financial uncertainty shocks than vice-versa. When we include proxies for ECB monetary policy into the model, we find that these reactions persist although they are reduced to some extent, thus reinforcing the previous finding that causality influences are running stronger in one direction, rather than in both directions. In other words, it is more likely that financial frictions and stress amplify policy uncertainty than vice-versa – a result that is in line with much of the existing empirical (e.g. Stock and Watson 2012; Caldara et al., 2016; Alessandri and Mumtaz 2019) and theoretical literature (e.g. Arellano, et al. 2010; Christiano, et al., 2014; Bloom, 2014). In addition, ECB policy reactions to uncertainty are stronger, but less persistent, for CISS shocks than for EPU shocks. Our findings further suggest that ECB adopted a more pro-active stance towards policy uncertainty shocks (and the variety and range of ECB unconventional policy measures stands as an additional proof), but a more passive or accommodative stance towards financial uncertainty shocks. All these empirical findings withstand multiple robustness checks.

Our analysis has policy implications as well. The insights we derive on ECB policy preferences are only indirect, but robust and revealing. Firstly, we find that in reaction to financial uncertainty shocks ECB prefers to accommodate the higher liquidity demand of banks by deploying its conventional monetary policy tools (i.e. short term interest rates) with direct effects on yield curve and sovereign spreads. Although this type of reaction might appear limited in an environment where the zero lower bound is binding, empirical evidence suggests that ECB has in fact been quite effective in lowering the short end of the EA yield curve further into negative territory (see Wu and Xia, 2017). Secondly, we find that ECB prefers to deploy its unconventional toolkit and, in the same time, be less likely to reverse course, when reacting to policy uncertainty shocks, whose prevalence has increased in some EA countries (e.g. France and Italy, and mostly due to political turmoil). Yet, prevalence and political turmoil cannot be used as a justification, so arguments must be looked for elsewhere. Policy uncertainty shocks also seem not to pose too great of a challenge to area-wide financial stability, which does in fact represent one of the ECB core policy objectives. According to our discussion in section 2, policy uncertainty shocks are more likely to lead to a 'segmented equilibrium' within the EA financial system, therefore justifying ECB actions. Given the current incomplete institutional architecture of the Euro Area, we can expect ECB to remain the 'only game in town' and therefore assume leadership in reacting to both policy and financial uncertainty shocks.

REFERENCES

Acharya, V. V., & Steffen, S. (2015). The "greatest" carry trade ever? Understanding eurozone bank risks. Journal of Financial Economics, 115(2), 215-236.

Acharya, V., Drechsler, I., & Schnabl, P. (2014). A pyrrhic victory? Bank bailouts and sovereign credit risk. The Journal of Finance, 69(6), 2689-2739.

Acharya, V. V., & Yorulmazer, T. (2008). Information contagion and bank herding. Journal of money, credit and Banking, 40(1), 215-231.

Alessandri, P., & Mumtaz, H. (2019). Financial regimes and uncertainty shocks. Journal of Monetary Economics, 101, 31-46.

Allen, F., Beck, T., Carletti, E., Lane, P., Schoenmaker, D. & Wagner, W. (2011). Cross-Border Banking in Europe: Implications for Financial Stability and Macroeconomic Policies. CEPR report, http://www.cepr.org/pubs/books/CEPR/cross-border_banking.pdf.

Allen, F., Carletti, E., Goldstein, I., & Leonello, A. (2018). Government guarantees and financial stability. Journal of Economic Theory, 177, 518-557.

Anaya, P., Hachula, M., & Offermanns, C. J. (2017). Spillovers of US unconventional monetary policy to emerging markets: The role of capital flows. Journal of International Money and Finance, 73, 275-295.

Angelini, G., Bacchiocchi, E., Caggiano, G., & Fanelli, L. (2019). Uncertainty across volatility regimes. Journal of Applied Econometrics, 34(3), 437-455.

Arellano, C., Bai, Y. & Kehoe, P. (2010). Financial Markets and Fluctuations in Uncertainty. Federal Reserve Bank of Minneapolis, Research Department Staff Report.

Attinasi, M. G., Checherita, C., & Nickel, C. (2010). What explains the surge in euro area sovereign spreads during the financial crisis of 2007-09? Public Finance & Management, 10(4).

Bacchiocchi, E. (2017). On the Identification of Interdependence and Contagion of Financial Crises. Oxford Bulletin of Economics and Statistics, 79(6), 1148-1175.

Baker, S. R., Bloom, N., & Davis, S. J. (2016). Measuring economic policy uncertainty. The Quarterly Journal of Economics, 131(4), 1593-1636.

Battistini, N., Pagano, M., & Simonelli, S. (2014). Systemic risk, sovereign yields and bank exposures in the euro crisis. Economic Policy, 29(78), 203-251.

Bekaert, G., Hoerova, M., & Duca, M. L. (2013). Risk, uncertainty and monetary policy. Journal of Monetary Economics, 60(7), 771-788.

Bettendorf, T. (2019). Spillover effects of credit default risk in the euro area and the effects on the Euro: A GVAR approach. International Journal of Finance & Economics, 24(1), 296-312.

Bicu, A., & Candelon, B. (2013). On the importance of indirect banking vulnerabilities in the Eurozone. Journal of Banking & Finance, 37(12), 5007-5024.

Bloom, N. (2014). Fluctuations in uncertainty. Journal of Economic Perspectives, 28(2), 153-76.

Bloom, N., Floetotto, M., Jaimovich, N., Saporta-Eksten, I., & Terry, S. J. (2018). Really uncertain business cycles. Econometrica, 86(3), 1031-1065.

Boeckx, J., Dossche, M., & Peersman, G. (2017). Effectiveness and Transmission of the ECB's Balance Sheet Policies. International Journal of Central Banking, 13(1), 297-333.

Bruno, V., & Shin, H. S. (2014). Cross-border banking and global liquidity. The Review of Economic Studies, 82(2), 535-564.

Brunnermeier, M. K. (2009). Deciphering the liquidity and credit crunch 2007-2008. Journal of Economic perspectives, 23(1), 77-100.

Burriel, P., & Galesi, A. (2018). Uncovering the heterogeneous effects of ECB unconventional monetary policies across euro area countries. European Economic Review, 101, 210-229.

Caldara, D., Fuentes-Albero, C., Gilchrist, S., & Zakrajšek, E. (2016). The macroeconomic impact of financial and uncertainty shocks. European Economic Review, 88, 185-207.

Cesa-Bianchi, A. (2013). Housing cycles and macroeconomic fluctuations: A global perspective. Journal of International Money and Finance, 37, 215-238.

Chudik, A., & Fratzscher, M. (2011). Identifying the global transmission of the 2007–2009 financial crisis in a GVAR model. European Economic Review, 55(3), 325-339.

Cooper, R., & Nikolov, K. (2018). Government debt and banking fragility: The spreading of strategic uncertainty. International Economic Review, 59(4), 1905-1925.

Cerutti, E., Claessens, S., & Ratnovski, L. (2017). Global liquidity and cross-border bank flows. Economic Policy, 32(89), 81-125.

Christiano, L. J., Motto, R., & Rostagno, M. (2014). Risk shocks. American Economic Review, 104(1), 27-65.

De Grauwe, P., & Ji, Y. (2013). Self-fulfilling crises in the Eurozone: An empirical test. Journal of International Money and Finance, 34, 15-36.

De Santis, R. A., & Zimic, S. (2018). Spillovers among sovereign debt markets: Identification through absolute magnitude restrictions. Journal of Applied Econometrics, 33(5), 727-747.

De Santis, R. A. (2019). Impact of the asset purchase programme on euro area government bond yields using market news. Economic Modelling. Forthcoming. Doi.org/10.1016/j.econmod.2019.06.011.

Dees, S., Mauro, F. D., Pesaran, M. H., & Smith, L. V. (2007). Exploring the international linkages of the euro area: a global VAR analysis. Journal of applied econometrics, 22(1), 1-38.

Dragomirescu-Gaina, C., & Philippas, D. (2015). Strategic interactions of fiscal policies in Europe: A global VAR perspective. Journal of international Money and Finance, 59, 49-76.

Dungey, M., & Osborn, D. R. (2014). Modelling large open economies with international linkages: the USA and Euro area. Journal of Applied Econometrics, 29(3), 377-393.

Eickmeier, S., & Ng, T. (2015). How do US credit supply shocks propagate internationally? A GVAR approach. European Economic Review, 74, 128-145.

Elliott, G. (1998). On the robustness of cointegration methods when regressors almost have unit roots. Econometrica, 66(1), 149-158.

Faia, E. (2017). Sovereign risk, bank funding and investors' pessimism. Journal of Economic Dynamics and Control, 79, 79-96.

Farhi, E., & Tirole, J. (2017). Deadly embrace: Sovereign and financial balance sheets doom loops. The Review of Economic Studies, 85(3), 1781-1823.

Feldkircher, M., & Huber, F. (2016). The international transmission of US shocks—evidence from Bayesian global vector autoregressions. European Economic Review, 81, 167-188.

Forbes, K. J., & Warnock, F. E. (2012). Capital flow waves: Surges, stops, flight, and retrenchment. Journal of International Economics, 88(2), 235-251.

Foresti, P. (2018). Monetary and fiscal policies interaction in monetary unions. Journal of Economic Surveys, 32(1), 226-248.

Freixas, X., & Holthausen, C. (2004). Interbank market integration under asymmetric information. The Review of Financial Studies, 18(2), 459-490.

Fratzscher, M., Duca, M. L., & Straub, R. (2016). ECB unconventional monetary policy: Market impact and international spillovers. IMF Economic Review, 64(1), 36-74.

Feldkircher, M., & Huber, F. (2016). The international transmission of US shocks - evidence from Bayesian global vector autoregressions. European Economic Review, 81, 167-188.

Gambacorta, L., Hofmann, B., & Peersman, G. (2014). The effectiveness of unconventional monetary policy at the zero lower bound: A cross-country analysis. Journal of Money, Credit and Banking, 46(4), 615-642.

Gârleanu, N., Panageas, S., & Yu, J. (2015). Financial entanglement: A theory of incomplete integration, leverage, crashes, and contagion. American Economic Review, 105(7), 1979-2010.

Georgiadis, G. (2015). Examining asymmetries in the transmission of monetary policy in the euro area: Evidence from a mixed cross-section global VAR model. European Economic Review, 75, 195-215.

Gilchrist, S., Sim, J. W., & Zakrajšek, E. (2014). Uncertainty, financial frictions, and investment dynamics. National Bureau of Economic Research, No. 20038.

Greenwood-Nimmo, M., Huang, J., & Nguyen, V. H. (2019). Financial sector bailouts, sovereign bailouts, and the transfer of credit risk. Journal of Financial Markets, 42, 121-142.

Hebous, S., & Zimmermann, T. (2013). Estimating the effects of coordinated fiscal actions in the euro area. European Economic Review, 58, 110-121.

Hollo, D., Kremer, M., & Lo Duca, M. (2012). CISS - a composite indicator of systemic stress in the financial system. ECB Working Paper Series, No. 1426.

Kapetanios, G., Pagan, A., & Scott, A. (2007). Making a match: Combining theory and evidence in policy-oriented macroeconomic modeling. Journal of Econometrics, 136(2), 565-594.

Koijen, R. S., Koulischer, F., Nguyen, B., & Yogo, M. (2017). Euro-area quantitative easing and portfolio rebalancing. American Economic Review, 107(5), 621-27.

Kucharčuková, O. B., Claeys, P., & Vašíček, B. (2016). Spillover of the ECB's monetary policy outside the euro area: How different is conventional from unconventional policy?. Journal of Policy Modeling, 38(2), 199-225.

Leonello, A. (2018). Government guarantees and the two-way feedback between banking and sovereign debt crises. Journal of Financial Economics, 130(3), 592-619.

Ludvigson, S. C., Ma, S., & Ng, S. (2019). Uncertainty and business cycles: exogenous impulse or endogenous response? National Bureau of Economic Research. Working Paper No. 21803.

Jurado, K., Ludvigson, S. C., & Ng, S. (2015). Measuring uncertainty. American Economic Review, 105(3), 1177-1216.

Miranda-Agrippino, S., & Rey, H. (2015). US monetary policy and the global financial cycle. NBER working paper, 21722.

Moder, I. (2017). Spillovers from the ECB's non-standard monetary policy measures on south-eastern Europe. ECB Working Paper, No. 2095.

Pesaran, M. H., Schuermann, T., & Weiner, S. M. (2004). Modeling regional interdependencies using a global error-correcting macroeconometric model. Journal of Business & Economic Statistics, 22(2), 129-162.

Rigobon, R. (2003). Identification through heteroskedasticity. Review of Economics and Statistics, 85(4), 777-792.

Schmidt, J., Caccavaio, M., Carpinelli, L., & Marinelli, G. (2018). International spillovers of monetary policy: Evidence from France and Italy. Journal of International Money and Finance, 89, 50-66.

Sims, C. A., Stock, J. H., & Watson, M. W. (1990). Inference in linear time series models with some unit roots. Econometrica, 58(1), 113-144.

Shin, M., & Zhong, M. (2018). A new approach to identifying the real effects of uncertainty shocks. Journal of Business & Economic Statistics, 1-13.

Shin, H. S. (2012). Global banking glut and loan risk premium. IMF Economic Review, 60(2), 155-192.

Stângă, I. M. (2014). Bank bailouts and bank-sovereign risk contagion channels. Journal of International Money and Finance, 48, 17-40.

Stock, J. H., & Watson, M. W. (2012). Disentangling the Channels of the 2007-2009 Recession. National Bureau of Economic Research. NBER WP No. 18094.

Rey, H. (2015). Dilemma not trilemma: the global financial cycle and monetary policy independence. National Bureau of Economic Research. NBER WP No. 21162.

Wu, J. C., & Xia, F. D. (2017). Time-varying lower bound of interest rates in Europe. Chicago Booth Research Paper No. 17-06.

Appendix 1: Data description, sources and definitions

CISS – Composite Indicator for Systemic Risk. Frequency: monthly averages. Transformation: natural logarithm. Adjustment: seasonally adjusted using X-12 procedure. Source: ECB warehouse (https://sdw.ecb.europa.eu/browse.do?node=9689686). See also Hollo et al. (2012) for the methodology. For Baltics, we compute the average of CISS indexes for all three countries.

EPU – economic policy uncertainty index, computed based on Baker et al. (2016). Transformation: natural logarithm. Adjustment: seasonally adjusted using X-12 procedure. Source: data and methodology available from www.policyuncertainty.com.

VIX – the Chicago Board Options Exchange (CBOE) Volatility Index. Frequency: monthly averages. Transformation: natural logarithm. Source: Federal Reserve Bank of St. Louis database.

Spread – the difference between 10-year sovereign bond yields for each country and Germany (or US). Transformation: before computing the spreads, we apply the following transformation of yields: $yield^{adjusted} = \frac{1}{12} * \ln(1 + \frac{yield}{100})$ to smooth spikes in the time-series. Source: Eurostat.

Main Refinancing Operations rate – is the short term interest rate at which ECB provides the bulk of liquidity to the banking system of the Euro Area.[38] Source: ECB warehouse.

EONIA – is the Euro Overnight Index average or the Euro Interbank Offered Rate defined as the weighted rate for the overnight maturity, calculated by collecting data on unsecured overnight lending in the EA provided by banks belonging to the EONIA panel.[39] Frequency: monthly averages. Transformation: $yield^{adjusted} = \frac{1}{12} * \ln(1 + \frac{yield}{100})$. The liquidity proxy used in the extended GVAR is the spread (difference) between EONIA and the Main Refinancing Operations rate.

ECB assets – defined as central bank assets for Euro Area (11-19 Countries). Frequency: monthly, end of month. Transformation: natural logarithm. Adjustment: seasonally adjusted using X-12 procedure. Source: Federal Reserve Bank of St. Louis database.

TED spreads – defined as the spread between the 3-Month LIBOR based on US dollars and the 3-Month US Treasury Bills. Frequency: monthly averages. Transformation: none. Source: Federal Reserve Bank of St. Louis database.

[38] See https://www.ecb.europa.eu/stats/policy_and_exchange_rates/key_ecb_interest_rates/html/index.en.html
[39] See also the conclusions of the public consultation on euro risk-free rates at https://www.ecb.europa.eu/paym/pdf/cons/euro_risk-free_rates/ecb.consultation_details_201905.en.pdf.

Appendix 2: GVAR weighting matrixes

Panel A below displays the weighting matrix, W, based on IMF CPIS data that is used in the baseline and extended GVAR specifications. Weights reflect portfolio allocations from countries mentioned on rows towards countries on mentioned on columns (country labels are according to Table 1). The colour of each cell indicates the share of country's portfolio allocation towards other countries, based on the scale displayed on the right of the figure. Each row sums to 1, as countries on the column represent the entire investable universe for the country specified at the start of each row.

Panel A:

IMF CPIS weights

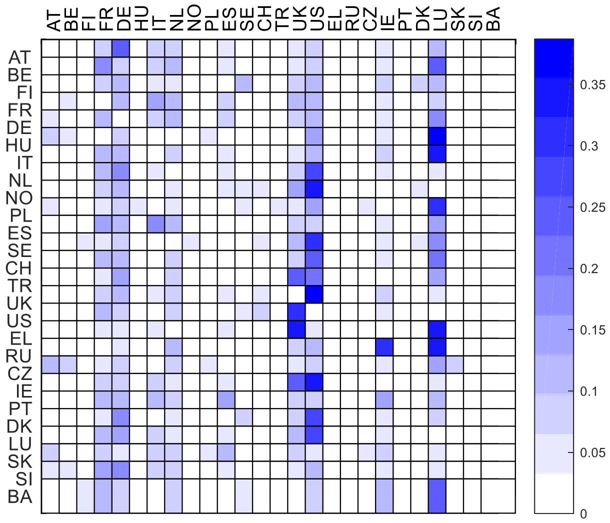

Panel B displays the weighting matrix used as robustness check in section 4.5, based on data from BIS Locational Banking Statistics, tables A6.2.[40] These tables contain data on cross-border positions in mil. USD, by counterparty's country of residency, and by location of the reporting bank. Since not all 28 countries (i.e. 25 individual countries and the 3 Baltics) are reporting to BIS, cross-border positions for banks located in other countries are only indirectly available as the reverse balance sheet positions of banks located in BIS reporting countries; for example, outward claims of banks located in Poland can be inferred as inward liabilities of banks located in BIS reporting countries with Polish resident banks as their counterparties. Moreover, for banks located in BIS reporting countries, there might be some differences between what banks from country X reports as outward claims in country Y, and what banks from country Y reports as inward liabilities from country X. To mitigate the impact of such inconsistencies, we average between (outward) claims and (inward) liabilities for all country pairs, and use bank-to-all sectors rather than just bank-to-bank positions. Further to reduce the impact of time-variation, we average the end of year (4th quarter) exposures over a 7-year period from 2010 to 2016. The colour of each cell indicates the share of country's outward exposures (i.e. claims) towards other

[40] See https://stats.bis.org/statx/toc/LBS.html.

countries, based on the scale displayed on the right of the figure. Each row sums to 1, as countries on the column represent the entire investable universe for the country specified at the start of each row.

Panel B:

BIS LBS weights

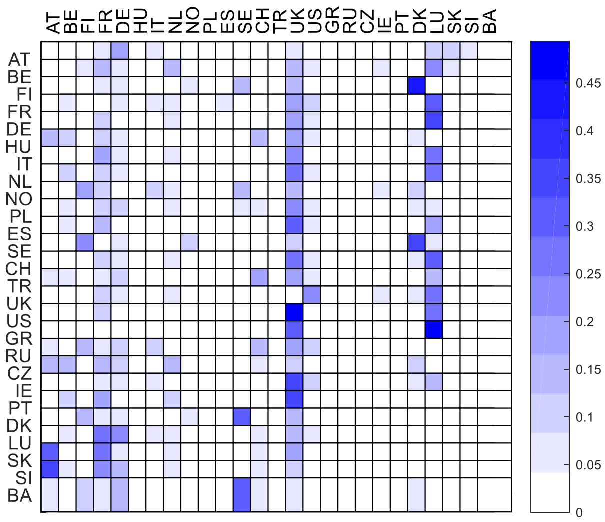

Appendix 3: The algorithm used for structural identification

The estimation algorithm consists in the following steps:

1. Bootstrap the reduced-form GVAR model given in equation (4) to obtain the variance-covariance matrix of reduced-form errors, $\Omega_u{}^{(b)}$. An initial estimate of $S_{2x2}{}^{(b)}$ is obtained as a Choleski decomposition of the upper block of $\Omega_u{}^{(b)}$, meaning $S_{2x2}{}^{(b)} = chol(\,\Omega_{u,2x2}{}^{(b)})$.

2. Obtain a candidate matrix $U^{(i)}$ that satisfies $U^{(i)}U^{(i)\prime} = \Omega_{u,2x2}{}^{(b)}$ and the identifying restrictions.

(2a) We draw a 2x2 matrix T from a standard normal distribution and obtain its QR decomposition, such that $T = \mathrm{QR}$, where Q is orthonormal, i.e. $\mathrm{QQ}' = \mathrm{I}$.

(2b) We check whether the matrix $U^{(i)} = S_{2x2}{}^{(b)}\mathrm{Q}$ satisfies the magnitude restrictions. If it does, we keep this draw (i). If not, we go back to step (2a). We repeat this process for a maximum of 200 times, such that we obtain a sufficient number of successful draws.

4. Repeat step 1 and 2 for 200 (bootstraps) times; compute the 68% confidence bands for IFRs considering all successful candidate matrices $U = S_{2x2}{}^{(b)}\mathrm{Q}$ from step 2b).

Appendix 4: Figures

Figure 4.1: IRFs to Italian uncertainty shocks in the baseline and extended specifications

Panel A

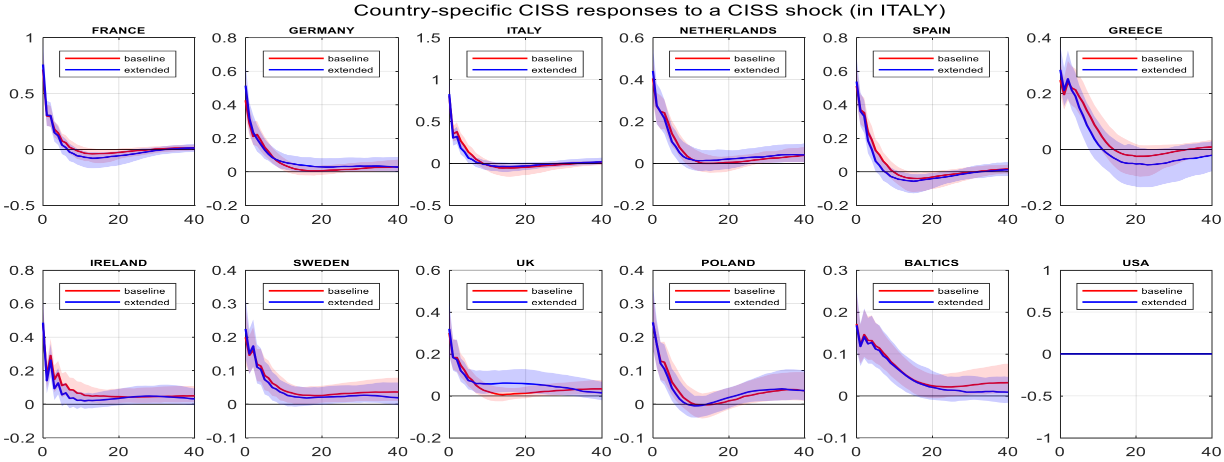

Panel B

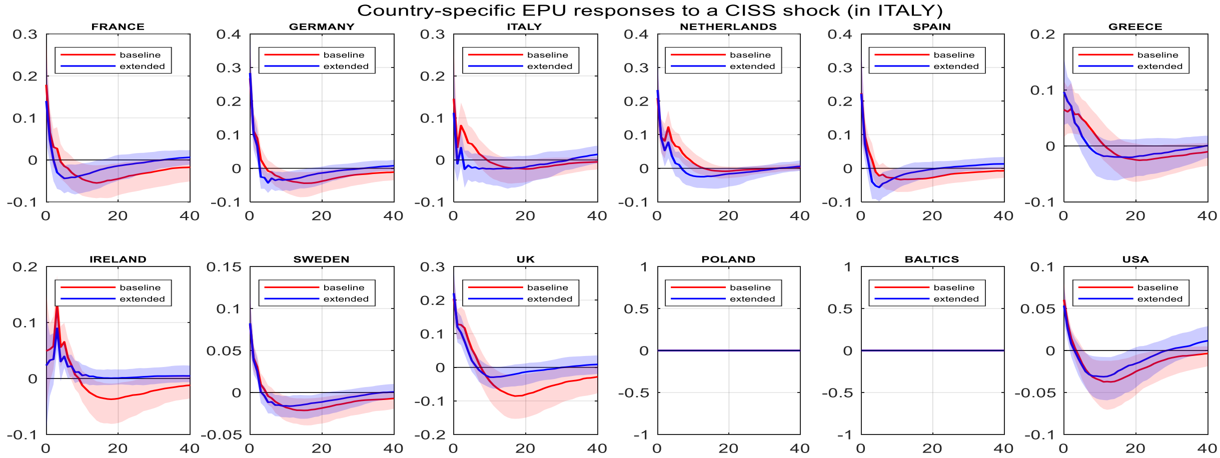

Panel C

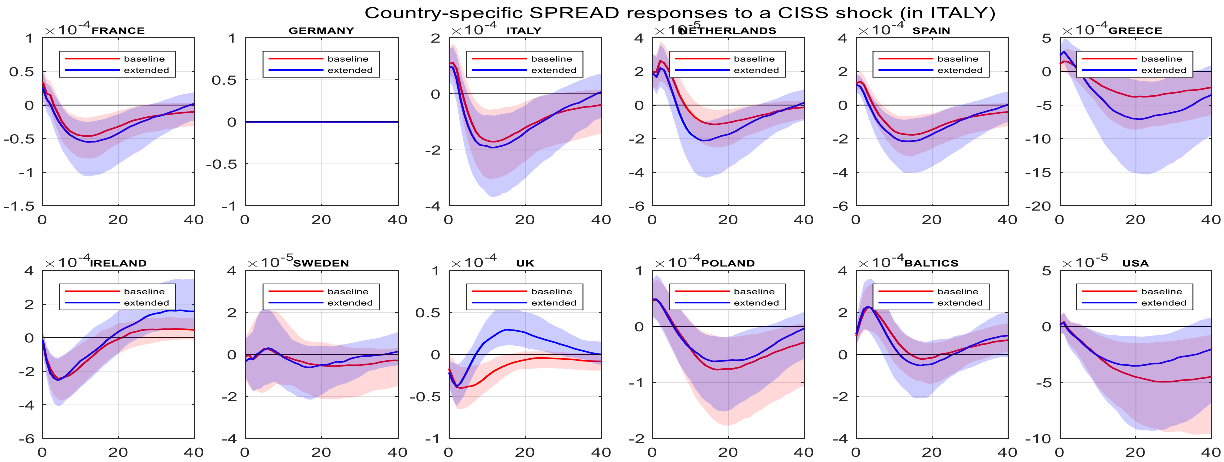

Panel D

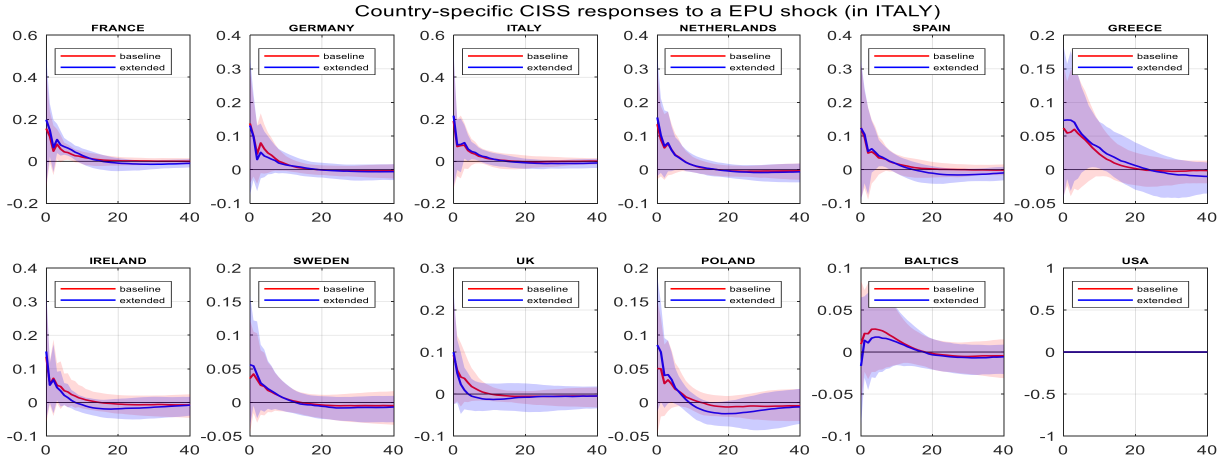

Panel E

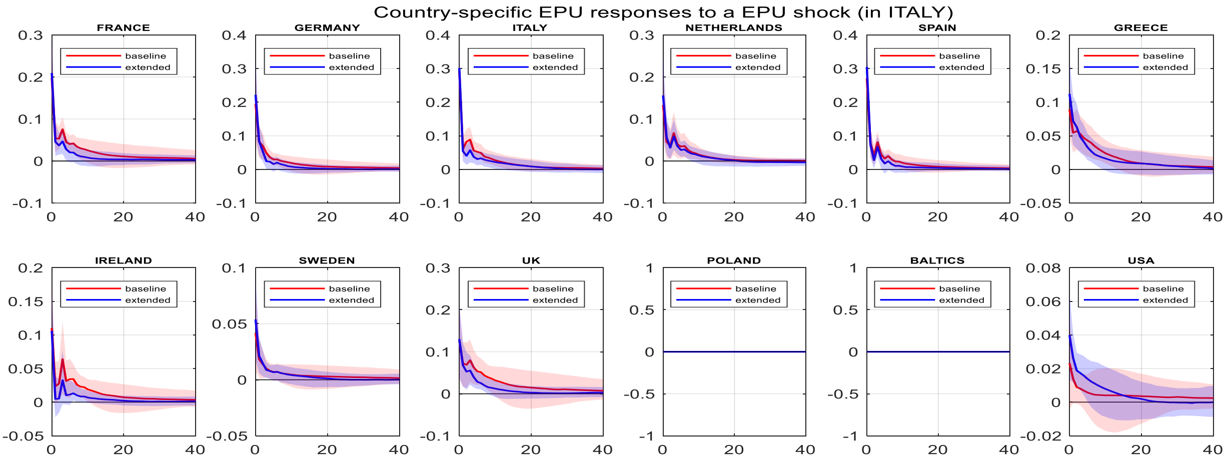

Panel F

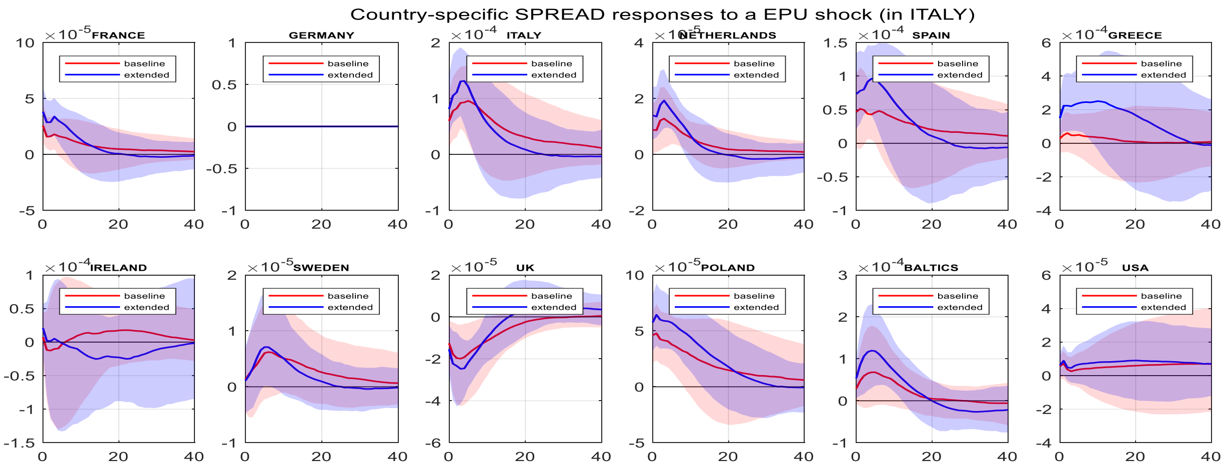

Note: The legend displays the corresponding GVAR specification; ‘baseline’ stands for the baseline GVAR; ‘extended’ stands for the extended GVAR with both conventional and unconventional monetary policy proxies included. The 68% confidence bands are constructed from 200 bootstrapped replications of the GVAR, each with 200 draws for the orthonormal matrix that insures identification.

Figure 4.2: IRFs to Spanish uncertainty shocks in the baseline and extended specifications

Panel A

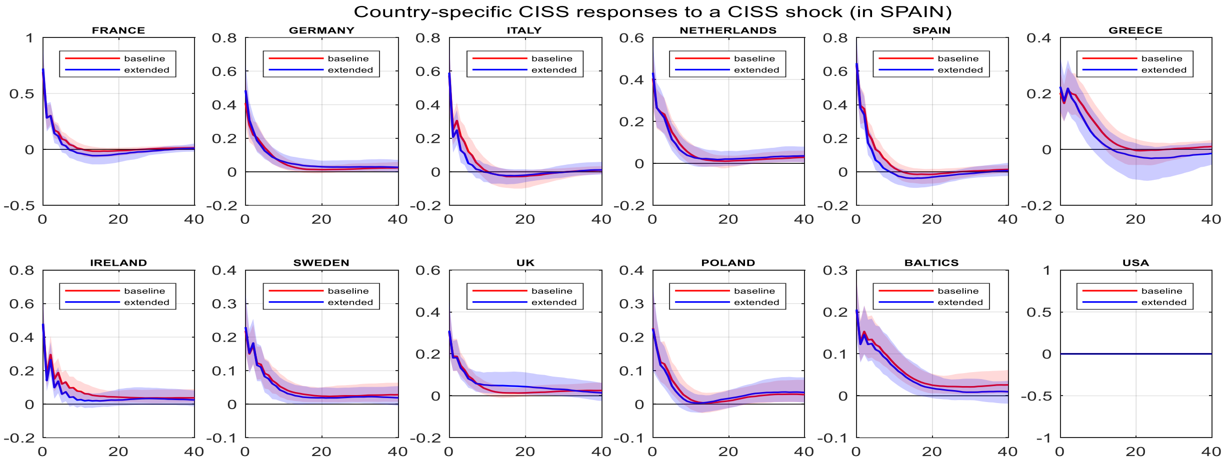

Panel B

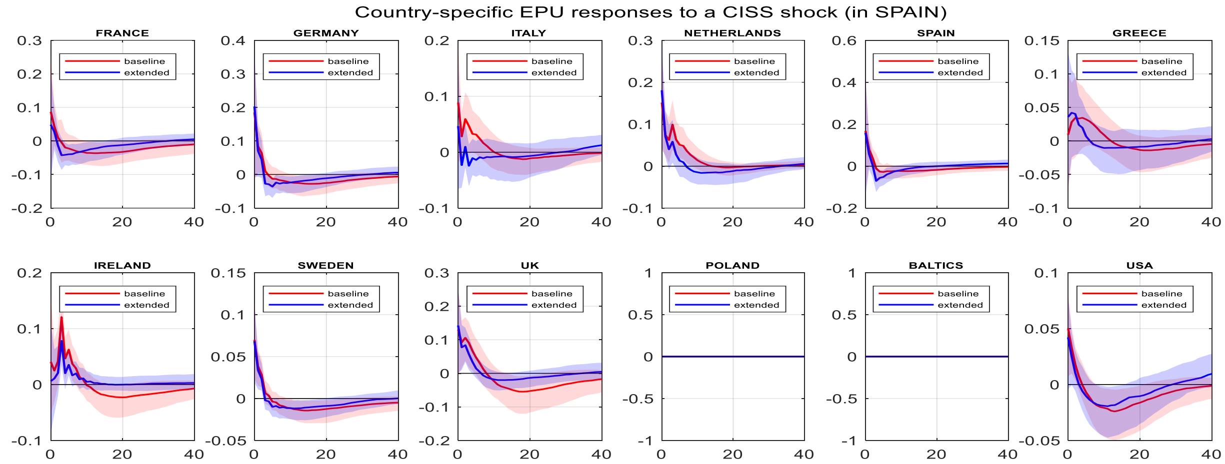

Panel C

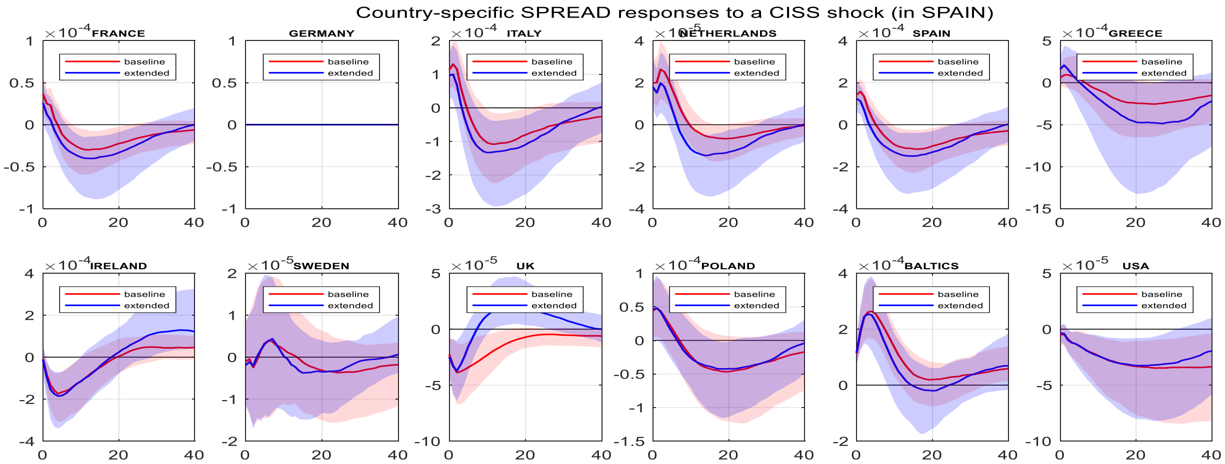

Panel D

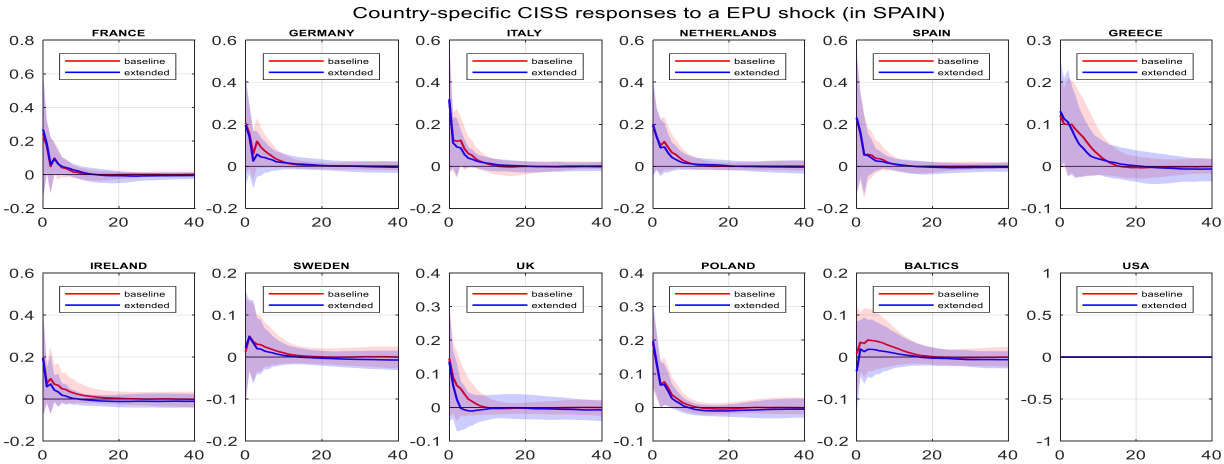

Panel E

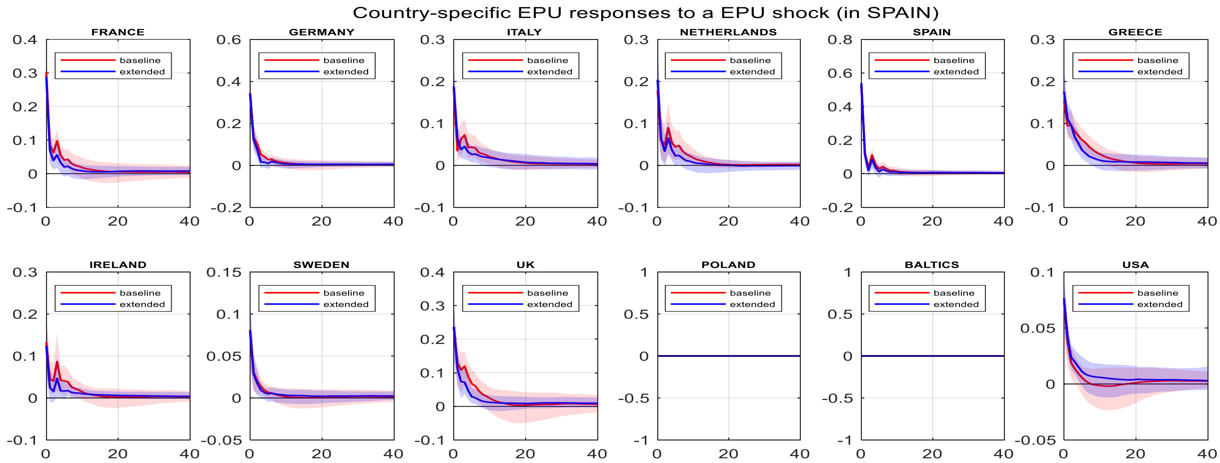

Panel F

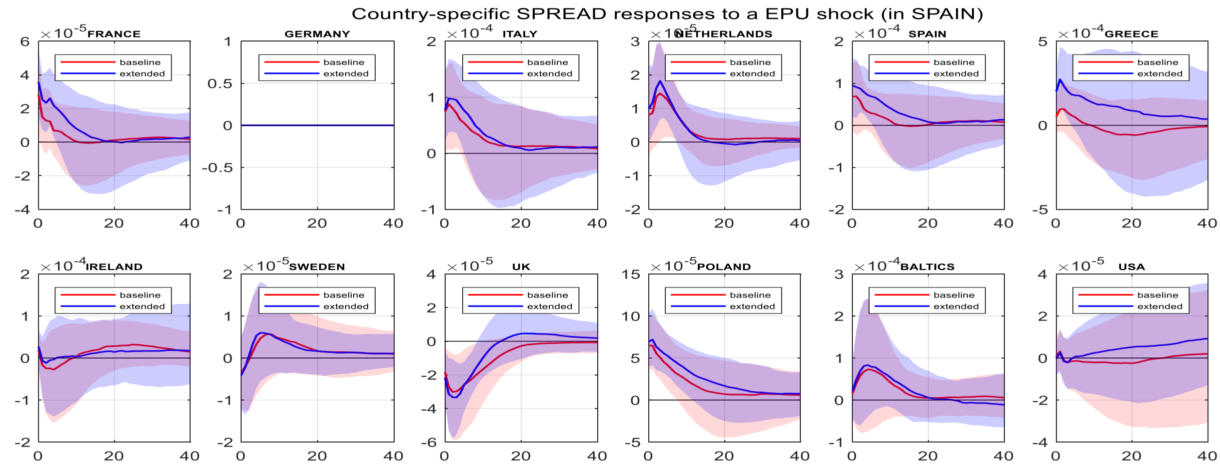

Note: The legend displays the corresponding GVAR specification; ‘baseline’ stands for the baseline GVAR; ‘extended’ stands for the extended GVAR with both conventional and unconventional monetary policy proxies included. The 68% confidence bands are constructed from 200 bootstrapped replications of the GVAR, each with 200 draws for the orthonormal matrix that insures identification.

Figure 4.3: IRFs to Greek uncertainty shocks in the baseline and extended specifications

Panel A

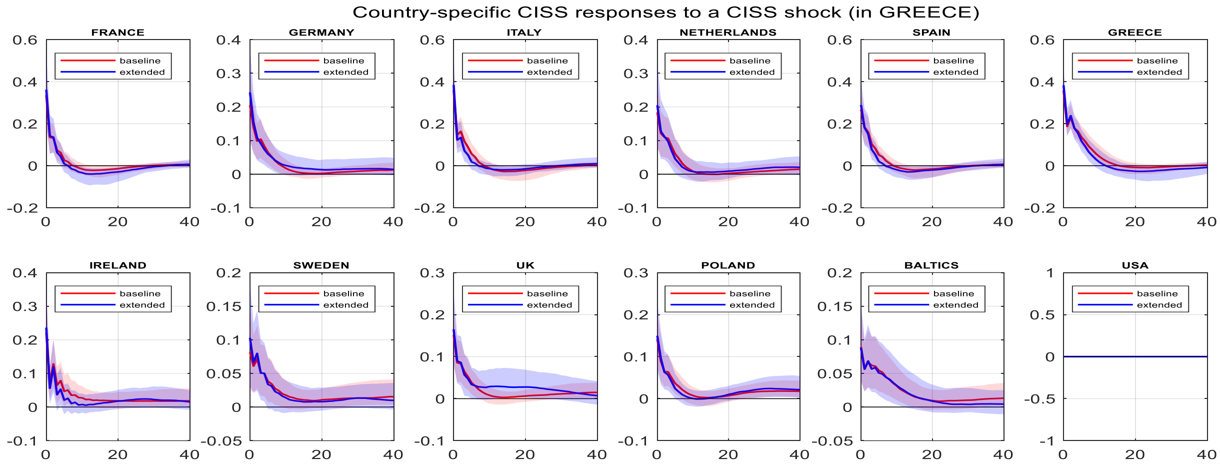

Panel B

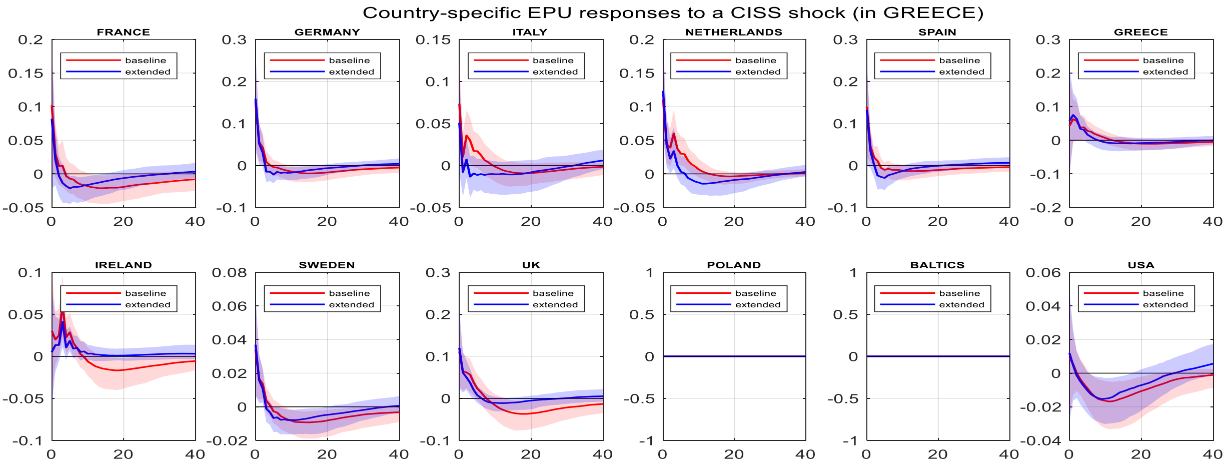

Panel C

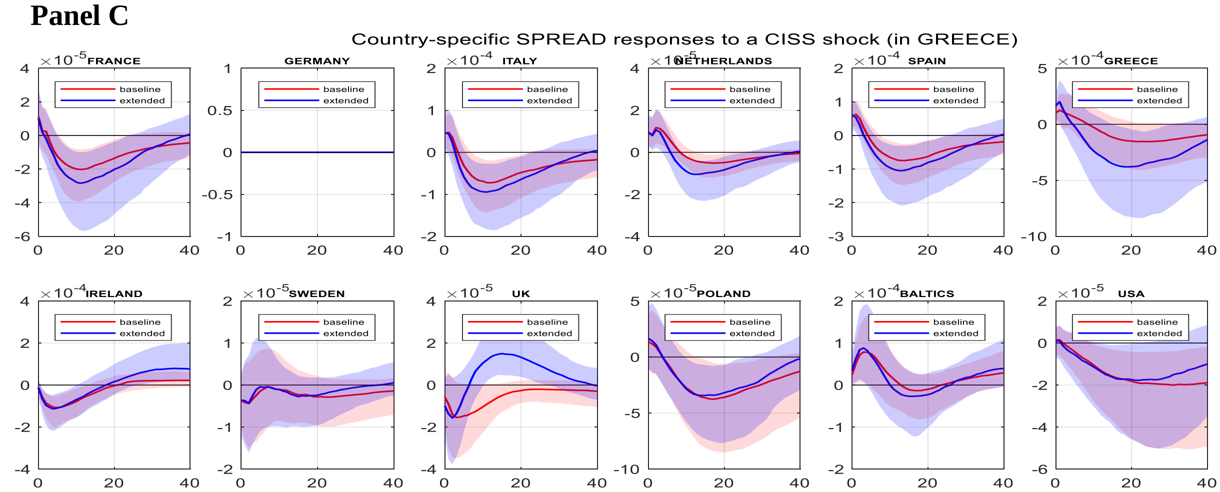

Panel D

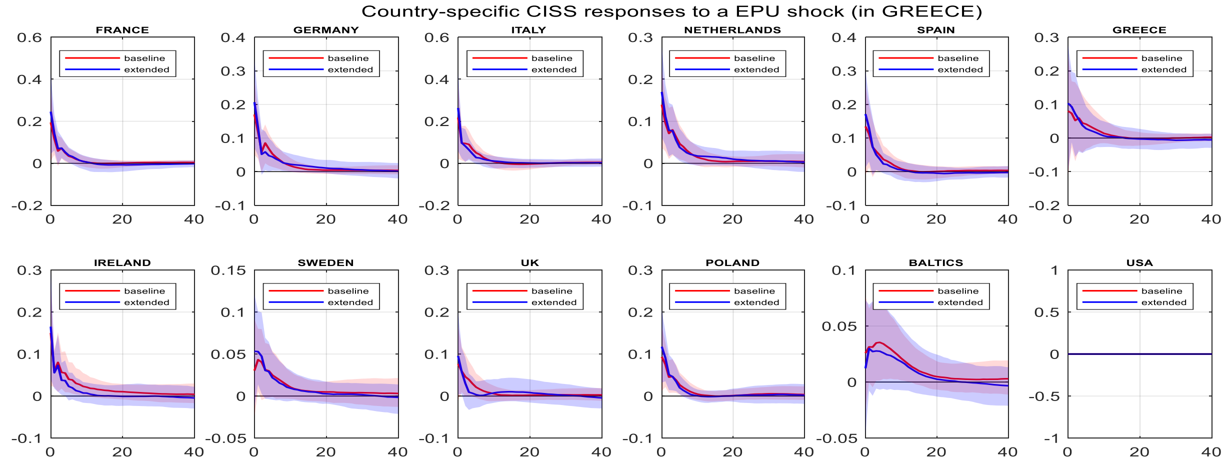

Panel E

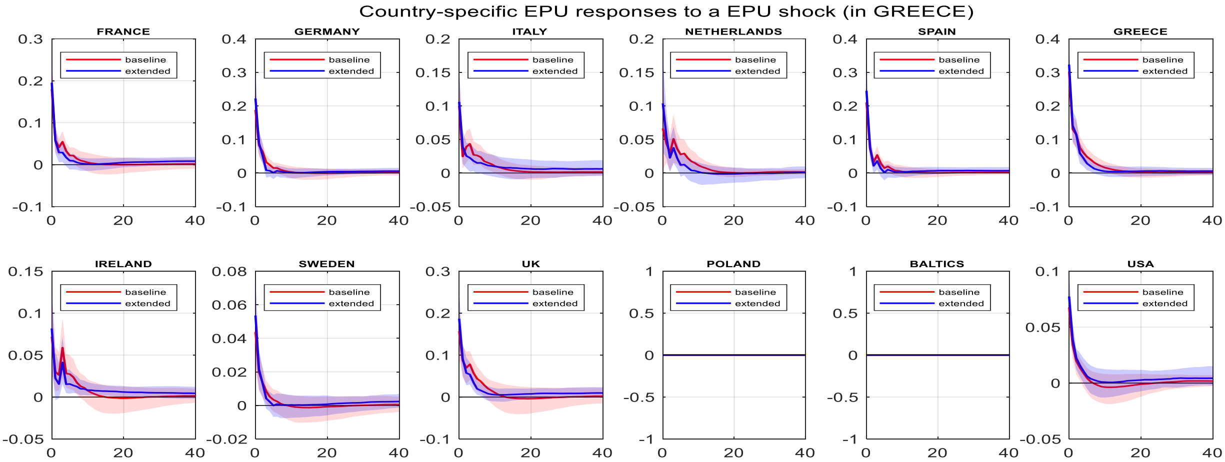

Panel F

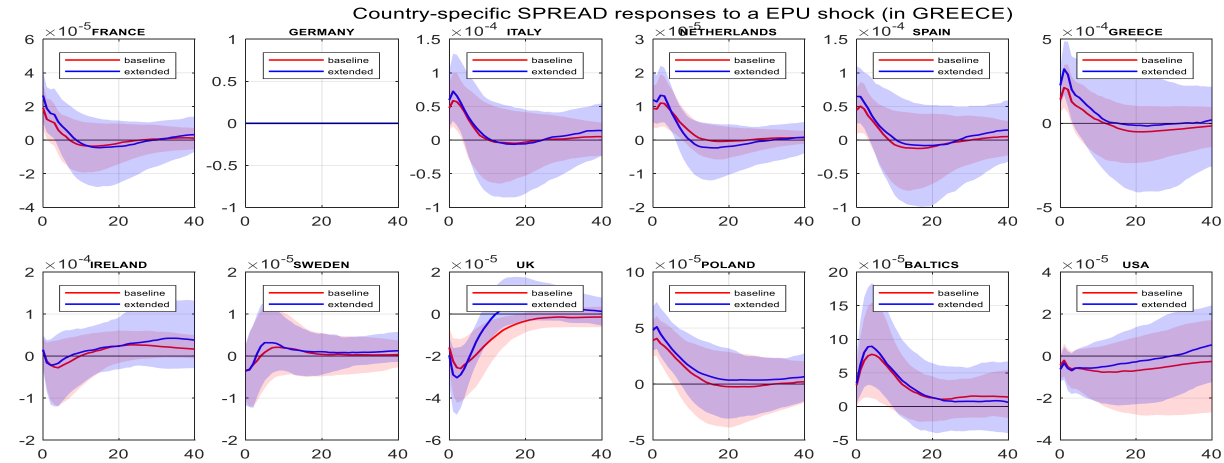

Note: The legend displays the corresponding GVAR specification; 'baseline' stands for the baseline GVAR; 'extended' stands for the extended GVAR with both conventional and unconventional monetary policy proxies included. The 68% confidence bands are constructed from 200 bootstrapped replications of the GVAR, each with 200 draws for the orthonormal matrix that insures identification.

Figure 4.4: IRFs to Irish uncertainty shocks in the baseline and extended specifications

Panel A

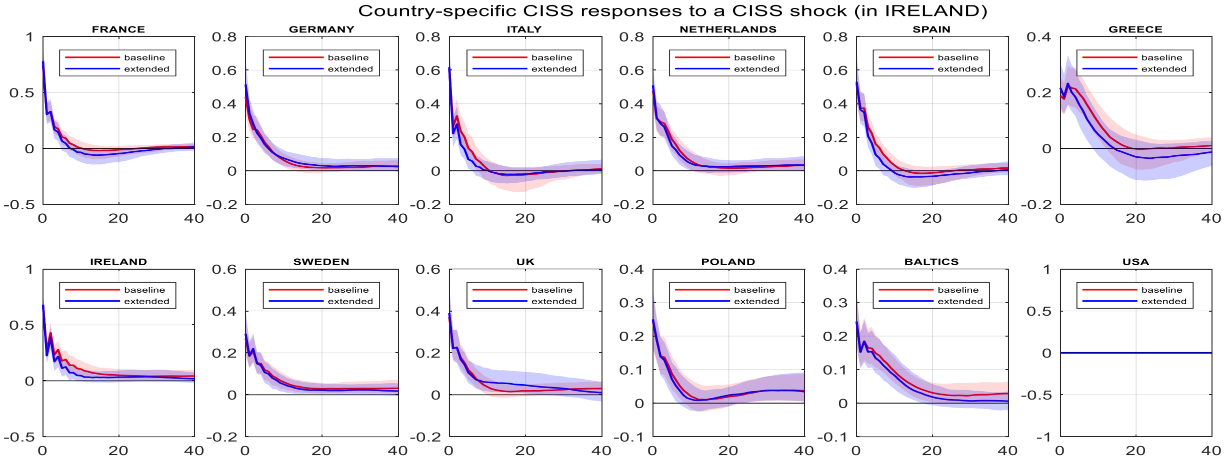

Panel B

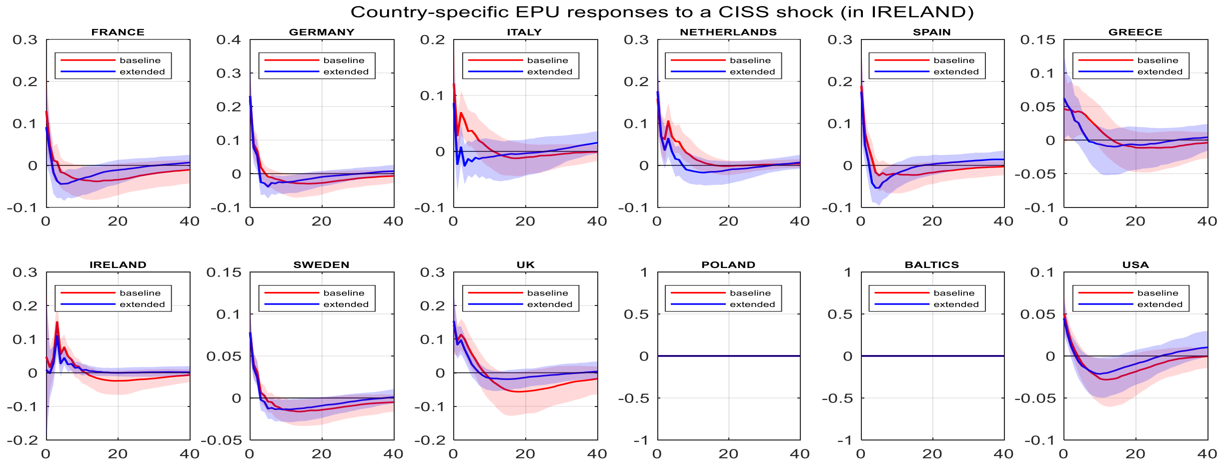

Panel C

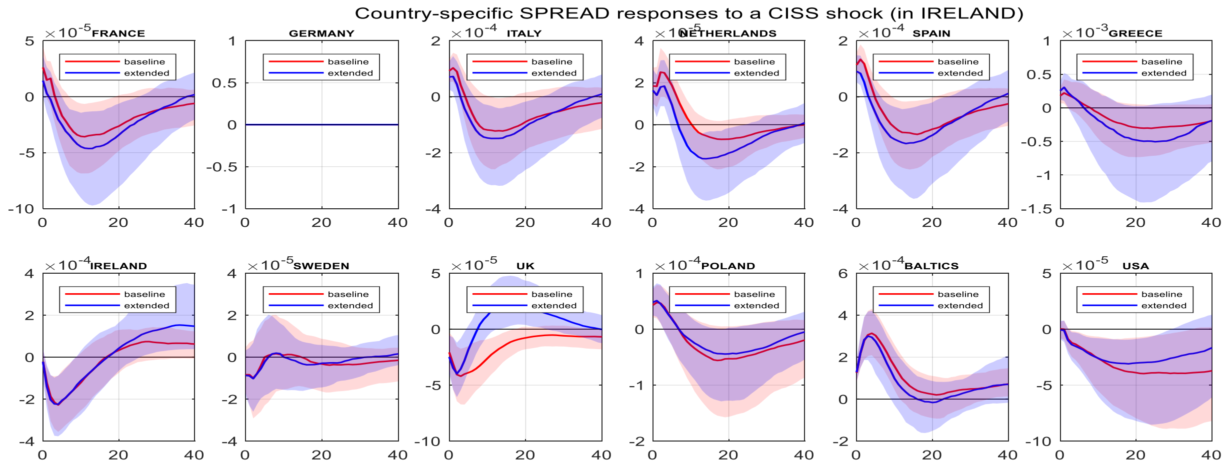

Panel D

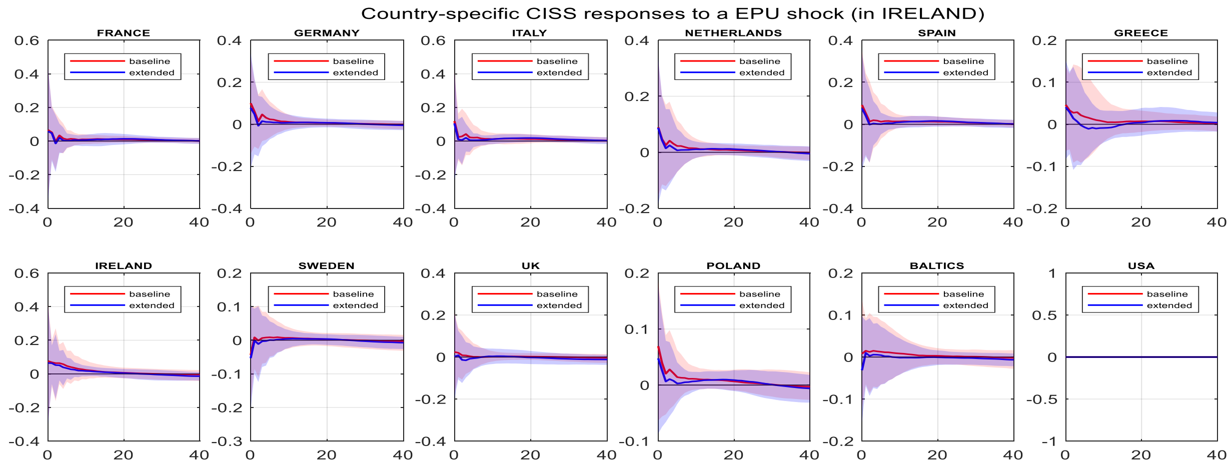

Panel E

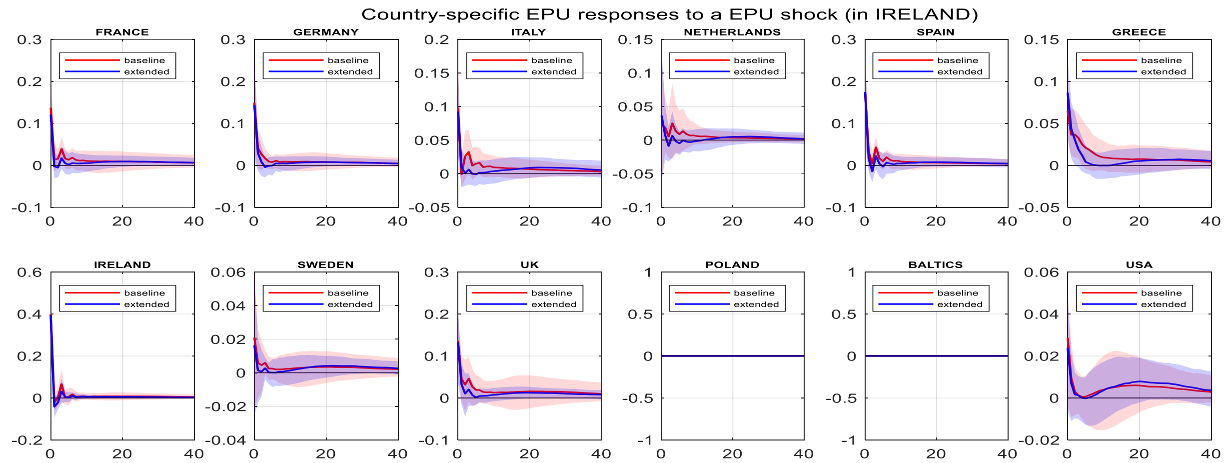

Panel F

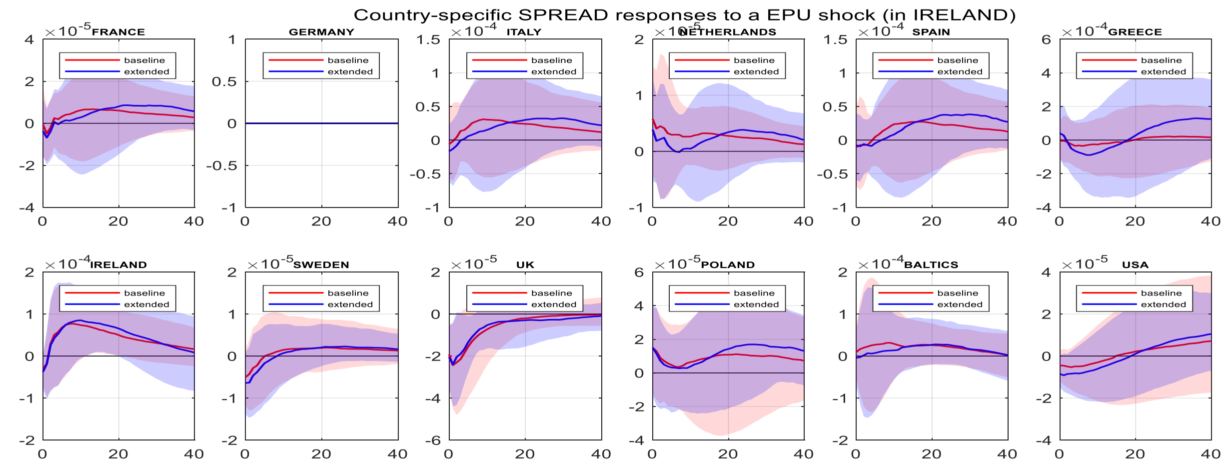

Note: The legend displays the corresponding GVAR specification; ‘baseline’ stands for the baseline GVAR; ‘extended’ stands for the extended GVAR with both conventional and unconventional monetary policy proxies included. The 68% confidence bands are constructed from 200 bootstrapped replications of the GVAR, each with 200 draws for the orthonormal matrix that insures identification.

Figure 4.5: IRFs to French uncertainty shocks in the baseline and extended specifications

Panel A

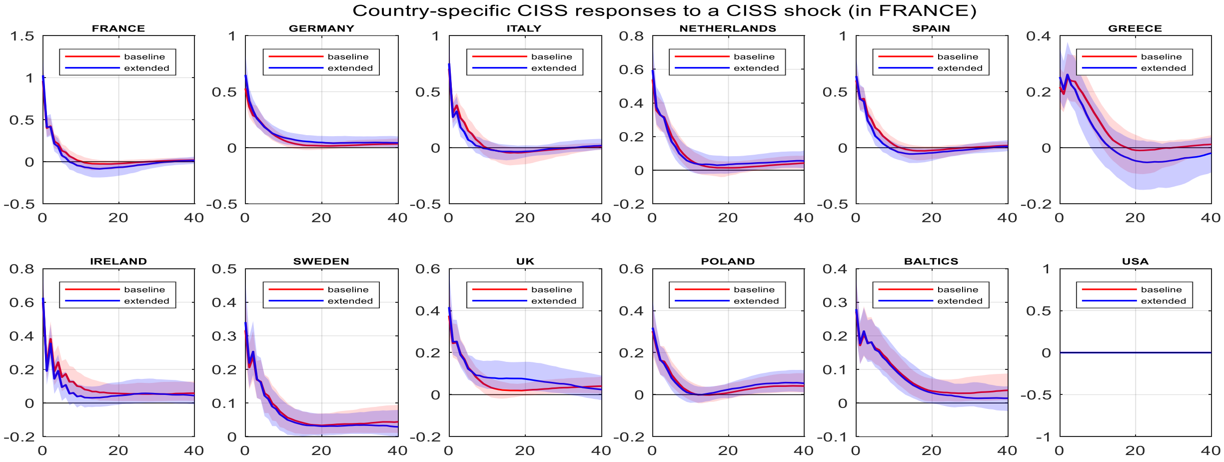

Panel B

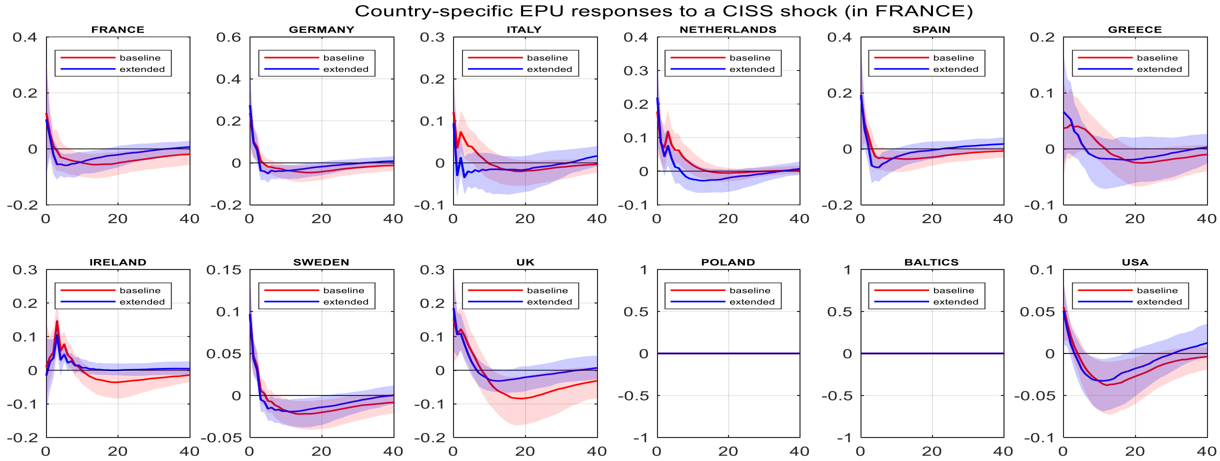

Panel C

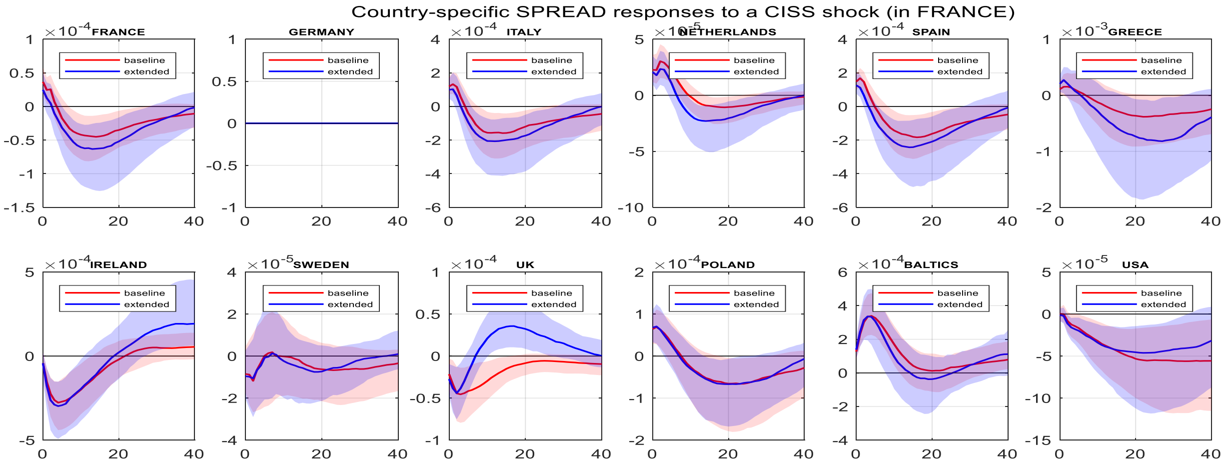

Panel D

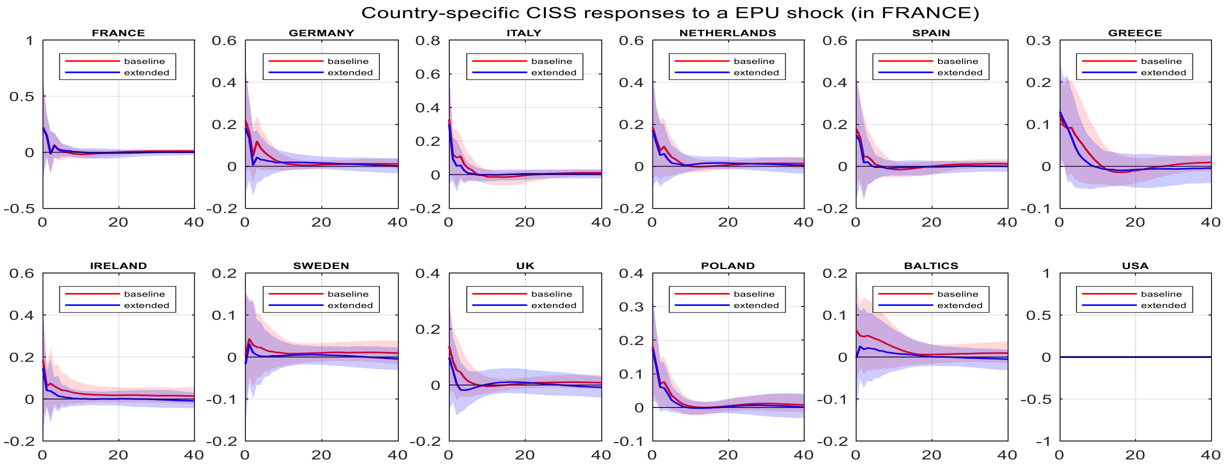

Panel E

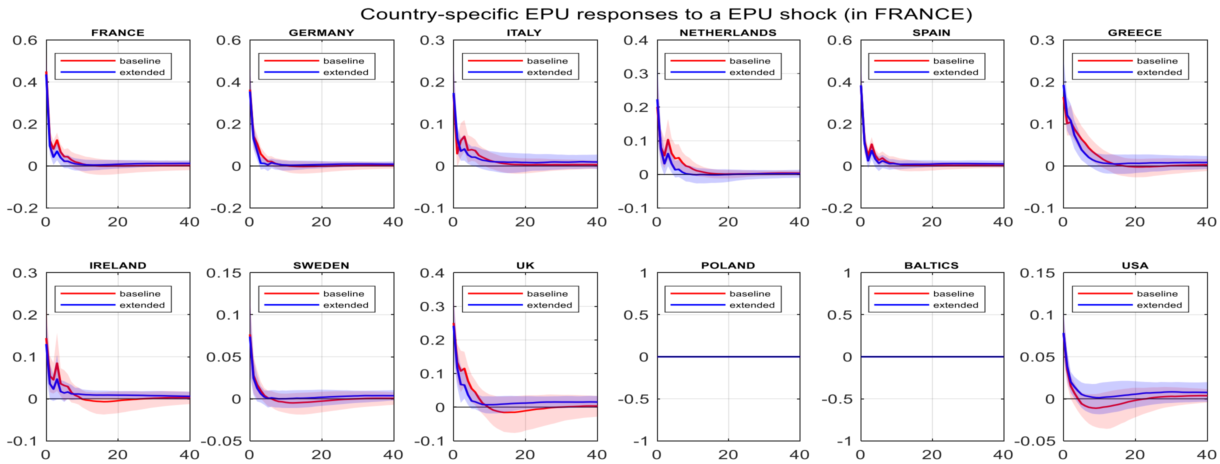

Panel F

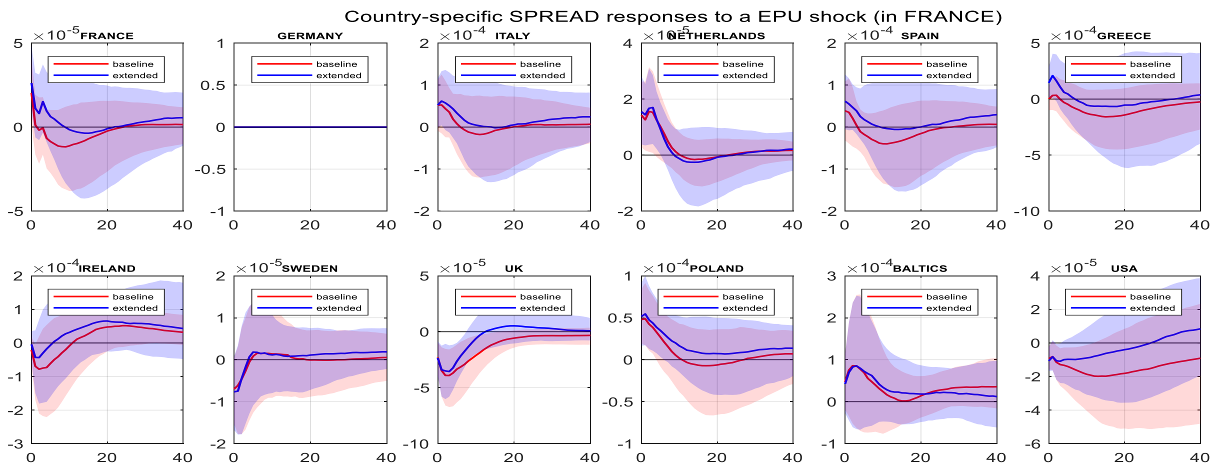

Note: The legend displays the corresponding GVAR specification; ‘baseline’ stands for the baseline GVAR; ‘extended’ stands for the extended GVAR with both conventional and unconventional monetary policy proxies included. The 68% confidence bands are constructed from 200 bootstrapped replications of the GVAR, each with 200 draws for the orthonormal matrix that insures identification.

Figure 4.6: IRFs to British uncertainty shocks in the baseline and extended specifications

Panel A

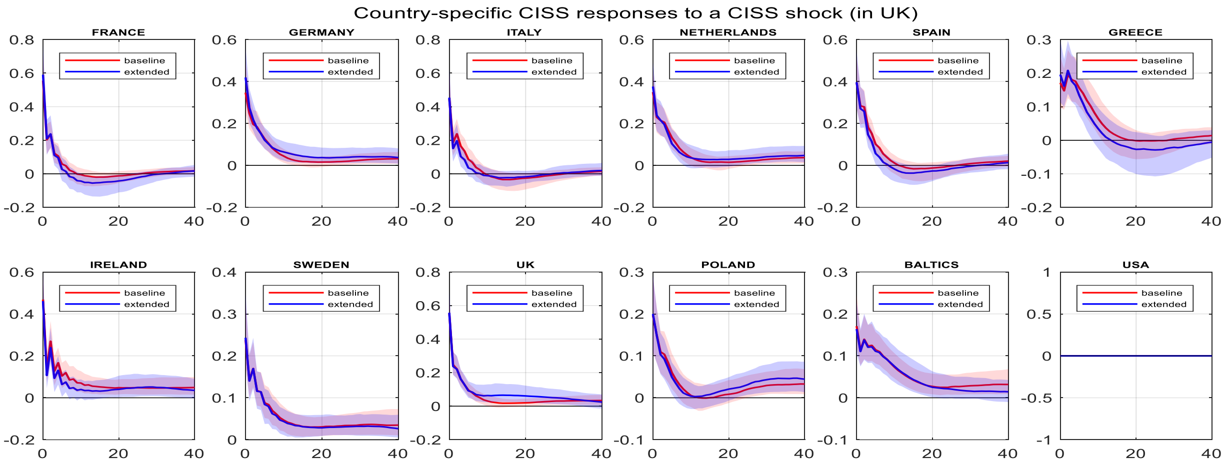

Panel B

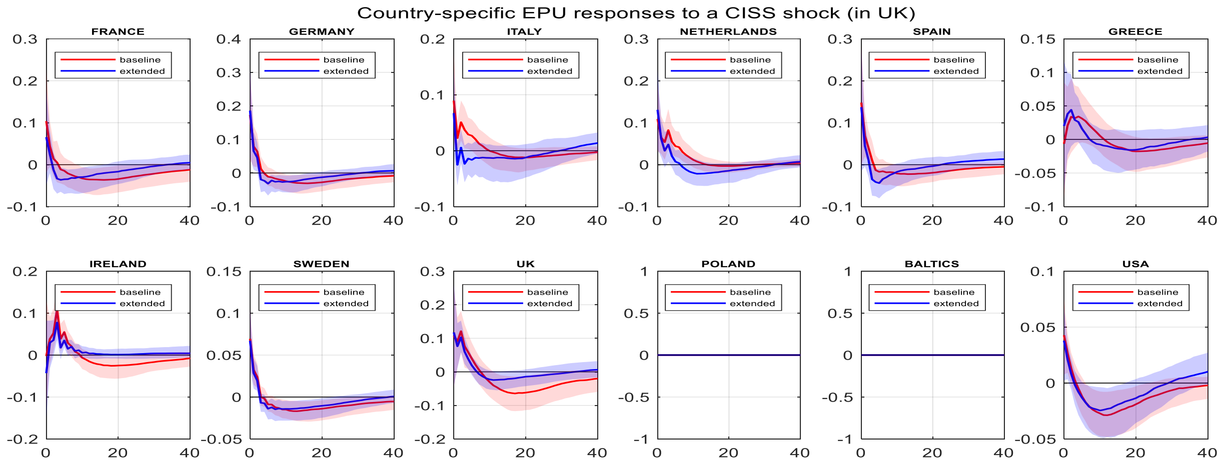

Panel C

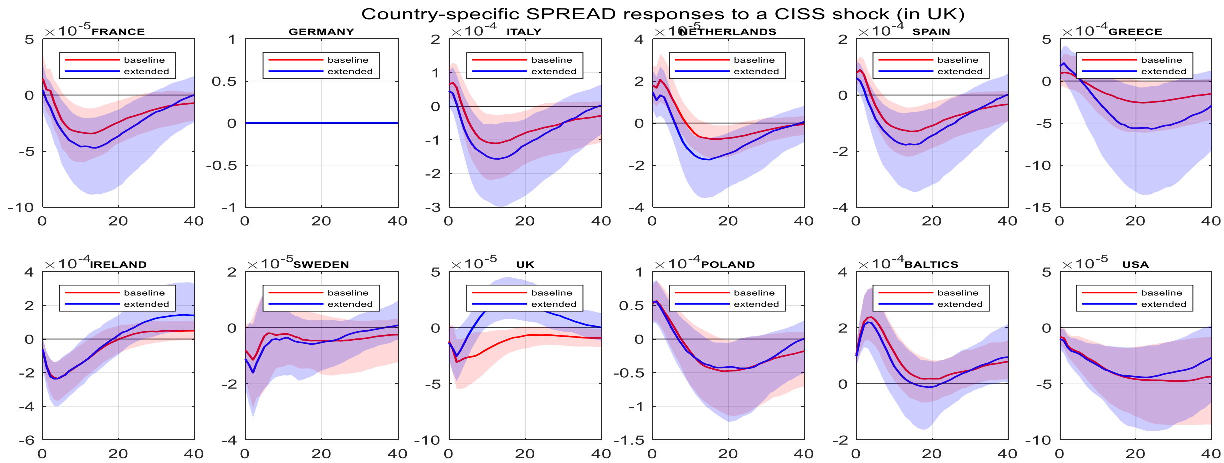

Panel D

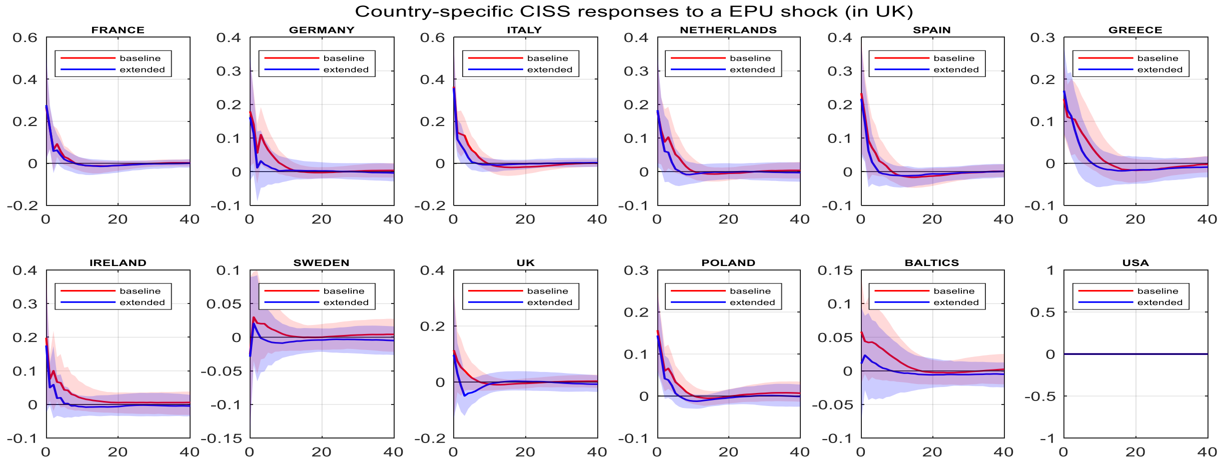

Panel E

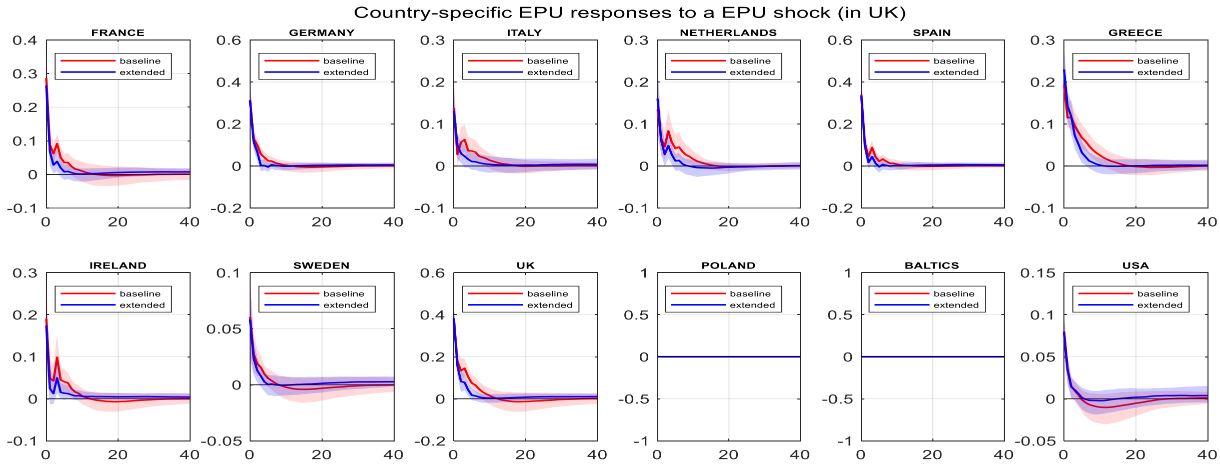

Panel F

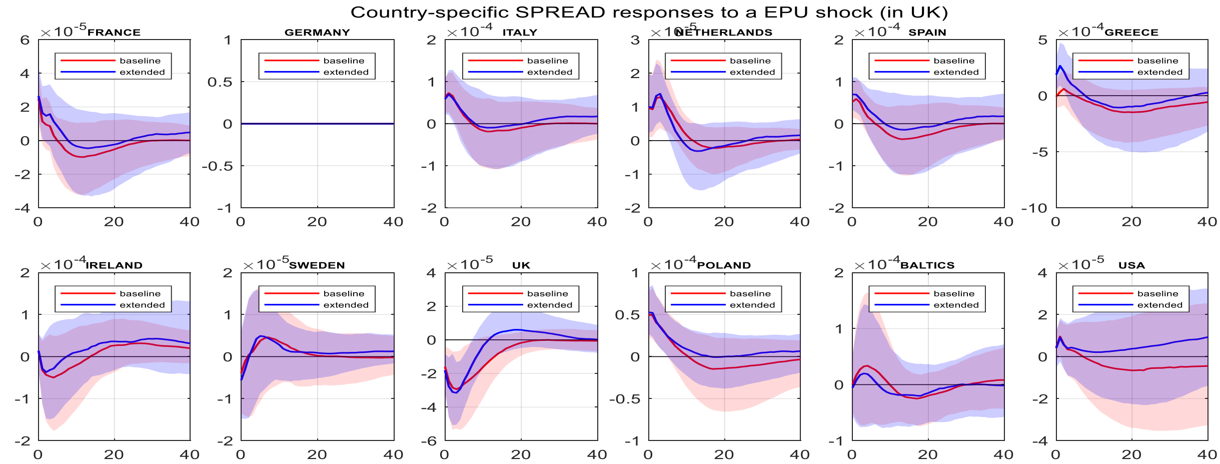

Note: The legend displays the corresponding GVAR specification; 'baseline' stands for the baseline GVAR; 'extended' stands for the extended GVAR with both conventional and unconventional monetary policy proxies included. The 68% confidence bands are constructed from 200 bootstrapped replications of the GVAR, each with 200 draws for the orthonormal matrix that insures identification.

Appendix 5

Figure 5.1: IFRs for ECB monetary policy proxies to uncertainty shocks

Panel A

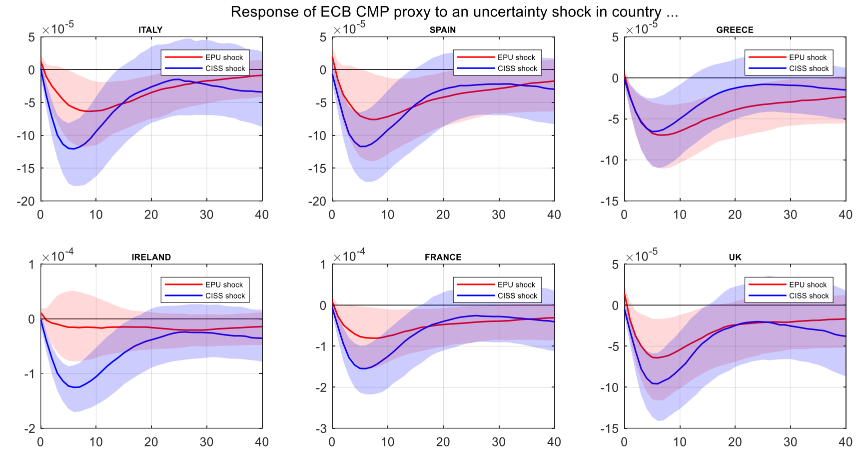

Panel B

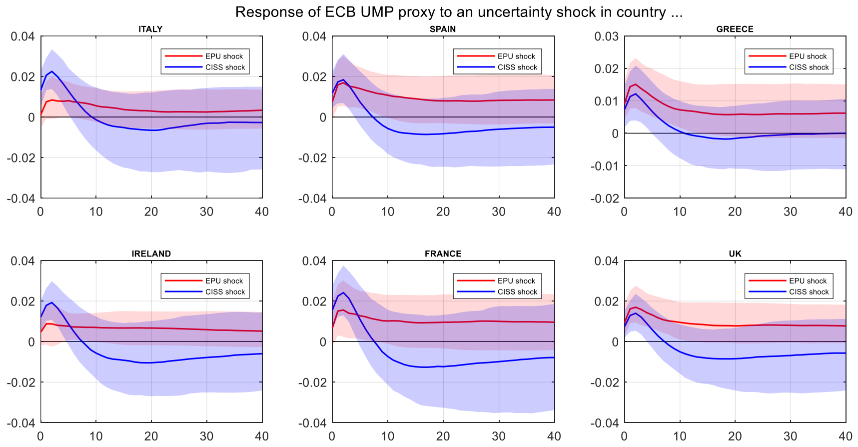

Panel C

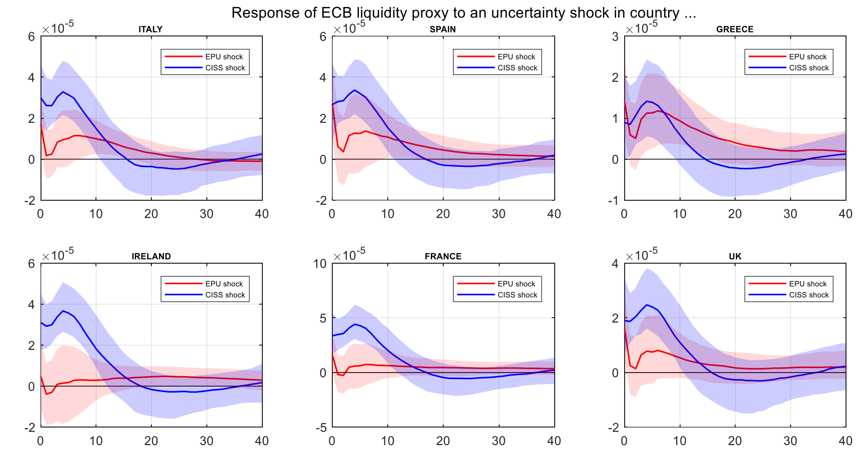

Note: The legend displays the corresponding uncertainty shock that is being simulated. The title of each plot displays the origin country of the shock. The 68% confidence bands are constructed from 200 bootstrapped replications of the GVAR, each with 200 draws for the orthonormal matrix (see algorithm in Appendix 3).

Chapter 2

Trading Off Accuracy for Speed: Hedge Funds' Decision-Making under Uncertainty†

Abstract

Hedge funds that operate quick portfolio adjustments, backed only by some rough estimates and loose predictions, can improve their market timing performances and benefit in turbulent markets, but oversimplification can lead to an inability to profit from opportunities in calm markets. The paper presents this trade-off between prediction accuracy and reaction speed, capturing some key aspects of decision-making under uncertainty. We select the accuracy levels upfront, through different data-filtering techniques, and investigate their empirical consequences on decision-making. Across different hedge funds' investment styles, our analysis shows that less accurate predictions can speed up reactions to unexpected changes in a large set of uncertainty and risk measures. We justify these empirical findings in a simulation exercise, highlighting the importance of market timing abilities for active players like hedge funds.

Keywords: hedge funds; dynamic portfolio exposure; time-varying beta

JEL classification: C14; C24; G11

† This chapter is co-authored with Prof. Dionisis Philippas from ESSCA School of Management, 55 Quai Alphonse Le Gallo, 92513, Boulogne, Paris, France; and Prof. Mike G. Tsionas from Lancaster University Management School, LA1 4YX, United Kingdom.

Acknowledgements: We are grateful to Emanuele Bacchiocchi, Daniel Andrei, Stefano Grillini, Leandro Elia, Eduardo Rossi, and Giulio Palomba for their many comments and suggestions that have greatly improved this chapter.

1. INTRODUCTION

Hedge funds (HFs) can exploit market opportunities that are not available to other market participants, being able to adjust their leverage, portfolio exposure and trading strategy, sometimes on a moment's whim. HFs implement their dynamic exposures to various pricing factors following a decision-making process based on a series of estimates and predictions that are continuously updated with new information. Obviously, in a market environment dominated by multiple sources of risk and uncertainty, such predictions cannot avoid rough approximations and simplifications of the available information set. Quick portfolio adjustments, backed only by some rough estimates and loose predictions, could increase HFs' reaction speed, which is key for their market timing performances (Cao et al., 2013; Bali et al., 2014). However, information losses due to oversimplification can lead to an inability to uncover and profit from opportunities, especially in calm markets. We hope to shed light on these key aspects of the HFs' decision-making process by drawing on some recent theoretical work where information acquisition behaviour determines both investors' attention and risky investment choices (see Huang and Liu, 2007; Andrei and Hasler, 2015, 2019; Kacperczyk et al., 2016). This growing literature strand on rational inattention builds on mostly psychological and experimental work, but also empirical evidence, on which we aim to contribute with new insights.

A representative HF is normally required to provide an investment alternative that offers diversification benefits to its clients. In fact, most HFs claim to generate excess returns that are uncorrelated, that is, have a low or even zero *beta* with some widely used benchmarks (Blocher and Molyboga, 2017; Agarwal et al., 2018). There is a large body of research examining HFs' exposure to various pricing factors, which essentially act as proxies for different (primitive) investment strategies. These studies employ various methods, ranging from the most straightforward (e.g. Fung and Hsieh, 1997; Agarwal and Naik, 2004) to complex techniques that are able to precisely infer the dynamics of the factor loadings or, in other words, the time-varying correlations between the HF portfolio returns and the pricing factors (see, among others, Aragon, 2007; Kessler and Scherer, 2011; Billio et al., 2012; Patton and Ramadorai, 2013; Savona, 2014a and 2014b; Racicot and Theoret, 2016). Moreover, there is a long and growing list of pricing factors (i.e. Fama and French, 1993, 2015; Carhart et al., 2014) able to give a better fit of realised returns, many of them specifically addressing the non-normal distribution of HF returns (Agarwal and Naik, 2004; Goyenko et al., 2009).

Such increase in the number and variety of pricing factors, and therefore of potential investment strategies, further complicates the decision-making process. In fluid markets, active investors such as HFs might decide about a change in strategy based on heuristics or intuition rather than on the basis of some accurate, real-time estimates of the portfolio impact of sudden market moves. For many HFs that claim to maintain a low *beta*, and in the same time generate profits, market timing is essential, and therefore reaction speed can be an advantage. In this context, it seems reasonable to assume that they might prefer a faster, but less accurate approach to adjusting their desired portfolio exposures to the targeted pricing factors, thus trading off accuracy for speed in this process. We address this trade-off upfront by employing data-filtering techniques that proxy for different accuracy levels and allow us to gain insights into the HFs' decision-making process under uncertainty.

With the advent of algorithmic trading, about two decades ago, transaction speed has collapsed to within milliseconds. However, what we are more interested in here is not the time required for trade execution, which is the final step in the portfolio adjustment process, but the speed through which changes in tactical allocations occur. On the one hand, the decision process leading to such changes is time consuming, even in today's HF industry, as it involves selecting among alternative strategies, evaluating and back-testing, and finally implementing the trade(s).[41] On the other hand, many HFs hold illiquid exposures that take time to reverse (Getmansky et al., 2004); for example, Aragon et al., (2013) find that HFs seek confidential treatment (and delay disclosure) for their illiquid positions to avoid front-running by other investors, implying that such positions are very important for HFs.

Consider the market factor as the most relevant pricing factor for a HF manager and think about her having to choose between a "moving-with-the-market" and a "moving-against-the-market" strategy, or else between a "high-*beta*" and a "low-*beta*" strategy in reaction to a sudden change in her information set. *Beta* is key for hedging effectiveness in portfolio management, and therefore changes in *beta* can be a good proxy for changes in strategy. A binary decision, formulated as a discrete prediction of *beta*, would require focusing her attention on a smaller subset of data or indicators, and therefore be much easier to take in an uncertain market environment (section 2 provides a detailed description of our conceptual framework). This simplification, however, is likely to be inefficient in normal times, when higher accuracy and attention levels are required to uncover smaller profit opportunities.

Allowing for time-varying levels of accuracy in HFs' predictions with respect to their portfolio exposures (i.e. *betas*) should be key for understanding HFs risky choices and behaviours under uncertainty. Whether this assumption is acceptable or not is the main question we address in the following sections. A battery of empirical models shows that strategies based on less accurate predictions can speed up HFs' reactions to relevant shocks in their information set, exposing a trade-off between accuracy and speed during market high-stress periods. To justify HFs' swift adjustments, we run a simulation exercise where we show that a portfolio switching from high to low-level accuracy *betas* during uncertain or risky periods is likely to outperform other strategies. Market timing requires higher reaction speed for HFs and, according to our empirical analysis and simulation exercise, this implies accepting lower accuracy levels in setting their portfolio exposures (*betas*).

[41] With algorithmic trading, execution can indeed take milliseconds, but it might take days as well, depending on whether the order is split in smaller sizes and allocated over time and across different market venues.

Although we consider a representative HF set-up when introducing and formulating the accuracy–speed trade-off, in our empirical application we differentiate HFs by investment style to capture the idea that HFs specialise in analysing and profiting from different information flows (e.g. equity hedge, macro, and relative arbitrage). While this specialization remains true in general, it cannot explain the prominent role that HFs have come to play in the cross-sectional transmission of systemic stress (Fung and Hsieh, 2006). In Cipriani and Guarino (2008), information spill-overs can lead to contagion when trading activity is correlated across markets, although fundamentals are not necessarily related. In the early model of King and Wadhwani (1990), contagion occurs when investors, despite investing in different markets, mistakenly interpret an idiosyncratic signal as carrying common, i.e. systemic, information. We follow this theoretical strand to explain some of our results from the empirical section, in which we find that various HF styles might respond simultaneously to unexpected shocks in some market-wide risk and uncertainty measures.[42]

We make two important contributions in this chapter. Firstly, we formulate and empirically examine the accuracy–speed trade-off, which characterises HFs investment strategies best during market stress episodes. We select the precision or accuracy of our *beta* estimates upfront, and then identify speed gains from within our empirical analyses. As discussed above, we let HFs formulate both accurate and less accurate predictions (i.e. estimates) about their portfolio *betas*, which are relabelled hereafter for convenience as *realised beta*, and *expected beta* respectively.[43] Intuitively, *accuracy* here refers to the full range of possible values taken by the prediction errors for the two *betas*, i.e. either a continuous interval or some range of discrete values. Technically, the derivation of the *realised beta* is based on a time-varying coefficients (TVC) Kalman filter, which is more precise in estimating *beta*, but less flexible, particularly in case of sudden changes in the data. By contrast, the *expected beta* is derived from a Markov Switching (MS) model, which can be considered a discrete version of the Kalman filter. Note that the recursive nature of the two filters represents an essential ingredient in an empirical analysis of decision-making. The main difference though between the two filters is that the unobservable nature of the hidden Markov Chain process behind the MS model requires approximations obtained by collapsing some of the terms (therefore, implying information losses) in the derivation of its likelihood function (see Kim, 1994). From an empirical perspective, this rather technical particularity of the MS filter should reflect the lower-level *accuracy* associated with filtering out *expected betas* from the data.

Secondly, we find that the less accurate predictions, measured by *expected betas*, could lead to simultaneous portfolio adjustments across different HF styles in reaction to large shocks that also have

---

[42] The difference between uncertainty and risk, though important, is less relevant to our discussion here. However, in section 3, we provide a full description and discuss in detail the list of risk and uncertainty measures employed in the present analysis.

[43] We label the more accurate prediction of beta as *realised beta* to deliberately suggest its higher informational content, and the less accurate prediction as *expected beta* to suggest a lower informational content. To some extent, the two *betas* might be also seen as the *ex-post* and *ex-ante* predictions derived from information processing.

high propensity to capture systemic stress and contagion risk. Our findings thus are in line with the theoretical models developed in King and Wadhwani (1990), and Cipriani and Guarino (2008), predicting that information spill-overs across various market segments can lead to contagion (see also Hasler and Ornthanalai, 2018). Two early warning indicators that might signal rapid shifts in HFs' risk appetite are identified from our analysis: the CBOE Volatility index (VIX) and the *Composite Indicator of Systemic Stress* (CISS) – a highly relevant policy indicator for the European financial sector.

The remainder of the chapter is organised as follows. Section 2 reviews the relevant literature, and provides the conceptual and mathematical formulation of the accuracy–speed trade-off. Section 3 presents a detailed description of the data. Section 4 lays out the empirical approach and discusses its main findings. Section 5 presents a simulation exercise that compares two hypothetical portfolios built on two strategies with different accuracy levels and market timing. Finally, section 6 concludes. More detailed results of our analyses are provided in three Appendixes at the end of the chapter.

2. CONCEPTUAL FRAMEWORK

Skilled, active investors like HFs can reallocate their attention[44] over time in order to extract the most relevant information from many diverse and noisy sources, including market moves, asset prices, news and rumours etc. In the general equilibrium model of Kacperczyk et al., (2016), attention allocation is optimally determined along with asset prices and portfolio allocations; skilled investors prefer to learn more about idiosyncratic risks in expansions, and more about aggregate risks during recessions, because such risks affect a higher share of their portfolios during market turmoil (given the increase in correlations across different asset classes). Similar theoretical mechanisms can be found in the recent literature[45] investigating the relationship between uncertainty and market volatility on the one side, and investors' (in)attention and information acquisition behaviour on the other side (see, Huang and Liu, 2007; Andrei and Hasler, 2015, 2019). In line with these theoretical predictions, we want to understand whether allowing for time-varying levels of accuracy in HFs' predictions can explain HFs' shifts in portfolio exposures. Assuming prediction accuracy is an increasing function of attention, we can suppose that a HF manager focuses only on aggregate risks during market turmoil and prefers to formulate her decision in terms of a simple binary choice, for example in terms of choosing between a "high-beta" versus a "low-beta" strategy.[46] Intuitively, during volatile periods her prediction *accuracy*

[44] Time-varying attention features in the studies of Da et al., (2014), Yuan (2015), Lu et al., (2016) among many others.

[45] There is a larger and growing literature strand on rational inattention in general, encompassing both psychology and economics fields, as recently summarised in Gabaix (2017), and Spiliopoulos and Ortmann, (2018). Woodford (2014), and Steiner et al., (2017) present significant theoretical contributions in this area.

[46] We focus on *beta*, which is the coefficient of the market factor, because of its prime role in portfolio management. Moreover, Blocher and Molyboga (2017) and Agarwal et al., (2018) find recent evidence that HF clients use an overall market index as

will be low; to simplify, it can be restricted to some discrete values to which her portfolio exposures will be set accordingly. In less volatile periods, her *accuracy* instead might increase in order to exploit a wider range of market opportunities.

Similar ideas can be found in the "category-learning behaviour" described by Peng and Xiong (2006), or in the "simple forecasts and paradigm shifts" described by Hong, Stein, and Yu (2007). Peng and Xiong (2006) rely on psychological evidence that attention is a scarce cognitive resource to motivate learning in their asset-pricing model. They find that investors tend to focus more on market-level rather than firm-specific information when the information-processing efficiency is low (e.g. in turbulent times). Hong et al. (2007) also use arguments rooted in psychology that provide evidence on how people tend to simplify complex problems due to limited attention and the cognitive costs associated with information processing.

From an empirical perspective, we use two recursive filters to reflect two different levels of prediction accuracy with respect to HFs portfolios' *beta*, and then investigate the consequences of time-varying accuracy on decision-making. Bollen and Whaley (2009) also compare the performances of two estimated *betas*, one discrete and one time-continuous version, concluding that the discrete *beta* (i.e. a changepoint regression in their case) has superior statistical power in revealing HFs' time-varying exposures. Unsurprisingly, the use of Kalman filters has been common in the recent empirical literature investigating HFs' dynamic exposures to pricing factors (e.g. Billio et al., 2012; Racicot and Theoret, 2016). To the best of our knowledge, this is the first attempt that relies on (both discrete and continuous) Kalman filters to expose the accuracy-speed trade-off, which is novel in the empirical finance literature, though there is plenty of experimental evidence in both psychology and economics (some recent reviews are Gabaix, 2017; Spiliopoulos and Ortmann, 2018).

A handful of empirical approaches bear some similarities to our paper. Brogaard and Detzel (2015) study the asset-pricing implications of uncertainty (i.e. economic policy uncertainty, which is a type of uncertainty based on news and keywords, as proposed by Baker et al., (2016), which is also covered here) in a static CAPM framework. The studies of Ferson and Schadt (1996), Patton and Ramadorai (2013), Bali et al., (2014), Savona (2014a and 2014b), and Amisano and Savona (2017) model dynamic portfolio exposures using macroeconomic predictors. In particular, Patton and Ramadorai (2013) investigate how implied volatility influences the dynamics of HF exposures and find that HFs reduce their market exposure only during highly volatile periods. Billio et al., (2012) also study the time-varying non-linear HF exposure to pricing factors during different market volatility regimes and propose a measure of contagion based on the joint probability that all HFs are in the high-volatility regime; other papers dealing with contagion across HFs are those by Boyson et al., (2010) and Dudley

a performance evaluation benchmark; therefore, we can presume that HF managers also attend most closely to market risk to attract clients and justify their high management fees.

and Nimalendran (2011). Compared with these papers, our approach investigates in greater depth how the decision process is affected when we allow for different accuracy levels in the real-time estimation of portfolio exposures.

Another related strand in the empirical literature focuses on forecast combinations as a means of improving forecasting accuracy in the presence of structural changes, implying a well-known trade-off between *bias* and *variance* (Pesaran and Timmermann, 2007; Clark and McCracken, 2009). Our emphasis instead is not on improving HFs' (possible) forecasting models, but on evaluating the consequences of varying prediction accuracy levels on decisions.

Our work relates to the experimental literature as well, as recently summarised in Spiliopoulos and Ortmann (2018). A closely related strand is the literature dealing with "learning-to-forecast" experiments, as described by Hommes et al. (2005) and Heemeijer et al. (2009). Pastor and Stambaugh's (2009) show how a Bayesian investor can exploit model design and misspecification to improve prediction accuracy, even when "imperfect predictors" are available. In contrast to these papers, the ex-post forecasting ability of HFs is not the main focus of our analysis.

Finally, we touch on the literature strand concerned with contagion and information spill-overs. King and Wadhawani's (1990) model provides a theoretical channel through which a signal extraction "mistake" in one market can be transmitted to all other markets (i.e. separated by the non-overlapping trading hours in the original model) because agents cannot clearly differentiate between systemic and idiosyncratic information signals, or shocks (see also Cipriani and Guarino, 2008). In a similar vein, Hasler and Ornthanalai (2018) highlight the role of information spill-overs in amplifying contagion because increased investors' attention can lead to correlated trades executed across separate markets.

2.1 Mathematical formulation

Modelling the reactions of every HF in a precise way is, naturally, an impossible task. Consider the problem of a representative HF manager updating her prior belief distribution to a posterior, when the parameters of interest are given via a loss function, say $L(R,V;\theta)$, where R is a vector of the characteristics related to returns, V is a vector of the characteristics related to volatility and θ is a parameter vector unknown to the HF. Bissiri et al. (2016) show that rational agents can update θ under such circumstances without full information on the data-generating process. Following Bissiri et al. (2016), we assume that $p(\theta)$ represents prior beliefs about the parameters θ and, hence, a Bayesian posterior/update of the beliefs about the parameters can be made using:

$$p(\theta|R,V) \propto p(\theta)\ \exp(-L(R,V;\theta)) \tag{1}$$

In other words, we have a well-defined parameter of interest θ and an initial belief distribution about the location of the parameter $p(\theta)$; in this case the loss function defines a likelihood on which the HF manager can rely to update information. Clearly, this provides a basis for Bayesian learning by using belief probability distributions based on a solid foundation (loss) for which managers, generally, have a clear notion. Bayesian learning features in the theoretical models of Andrei and Hasler (2015) and Kacperczyk et al., (2016) as well. Thus, we can say that for some function ψ, we must have:

$$p(\theta|R,V) = \psi\,[L(R,V;\theta),p(\theta)] \tag{2}$$

Consider now the updating of manager's subjective beliefs, $p(\theta|R,V)$, as an action made under uncertainty using decision theory to guide the optimal action. To assess her ability to outperform the market, assume the HF manager focuses on just two parameters, say ζ and ξ, that proxy for her selectivity and market timing abilities (Ferson and Schadt, 1996). Consequently, there is a subjective loss function $L(\zeta_t,\xi_t;\theta)$ with selectivity (ζ_t) and market timing (ξ_t) both specified as time-varying. Given a prior $p(\theta)$, an optimal posterior distribution, say $v^*(\theta)$, is obtained by minimizing expected loss, in which case we have a general form: $v^*(\theta)$= $argmin_v[L(\zeta_t,\xi_t;\theta)]$.

Based on the same information set and by imposing the same basic model structure, let us consider now the loss function L_i , $i = \{f,\ p\}$, as follows: (i) L_f is an approximation of the manager's loss function under full information, and can be taken to be the current-period expectation of an inter-temporal loss function; and (ii) L_p is the approximation of the manager's loss function under partial (incomplete) information. Denote as I^{t-1} the information set up to period $t-1$; therefore, updated beliefs are given as:

$$p(\zeta_t,\xi_t,\theta|I^{t-1}) \propto p(\zeta_{t-1},\xi_{t-1},\theta)\ \exp(-L_i(\zeta_t,\xi_t;\theta)) \tag{3}$$

Here, we update information on selectivity, market timing and any structural parameters, θ, based on a prior $p(\theta)$ and a subjective loss function, $L_i(\zeta_t,\xi_t;\theta))$ under a Bayesian process. Although $p(\zeta_{t-1},\xi_{t-1},\theta)$ can be any prior, it can be defined more reasonably as the posterior from the previous period. This specification implies that we can formulate a suitable model on selectivity and market timing and proceed with the usual methods of Bayesian inference. If prior $p(\theta)$ is diffused or "loose" (relative to the likelihood), standard frequent methods can be used to estimate θ as well as ζ_t,ξ_t, for example through the use of a Kalman filter. Depending on the prior specification, we make a distinction between the *realised* and *expected* betas, such that market timing will be dealt with separately in the estimation. With respect to selectivity, our formulation implies that HF managers draw on the market factor when analysing portfolios' strategy.

3. DATA

Our data set comes from different sources and includes weekly observations for a period spanning from the beginning of July 2004 until the end of May 2017. The timespan is rich enough to include many of the most significant economic and financial stress events that have affected the global financial industry and reverberated across the investable universe of a representative HF, including events originating in the US, Europe, and elsewhere.[47] Moreover, besides financial triggers, the sample includes recent social and political events that were relevant to the HF industry and have generated significant market reactions, such as the Brexit referendum (June 2016), the latest US presidential election (November 2016), terrorist attacks (e.g. in Paris, London, etc.), and so on.

To compute HFs' weekly average returns, we use data provided by Hedge Fund Research (henceforth HFR), which has constructed a robust classification system that includes a strategy, sub-strategy, and regional investment focus. Billio et al. (2009) provide a full description of the various statistical aspects of the data provided by different HF data sets, including HFR; when analysing differences in the distributional properties of HF returns at both daily and monthly frequencies, they find larger deviations from normality for monthly rather than daily returns. Patton and Ramadorai (2013) analyse similar issues and find that the intra-month variation in HFs' portfolio exposures is significant, because, as long as HFs report on a voluntary basis (normally at the end of the month), they have a strong incentive to engage in window dressing. Based on the above, and the fact that some of the most important indicators on which we condition the decision-making process in our empirical analysis are available only on a weekly basis, we choose to use weekly observations in the empirical analysis.

Self-selection, backfilling, and survivorship biases are quite common in the HF industry, but some of them, especially the last two, could be mitigated by using investable (rather than non-investable) benchmarks. Our empirical analysis focuses mainly on four investable HF main styles, which are captured by the following indexes: (i) the HFRX Event Driven Index, denoted by *ED*; (ii) the HFRX Equity Hedge Index, denoted by *EH*; (iii) the HFRX Macro/CTA Index, denoted by *M*; and (iv) the HFRX Relative Value Arbitrage Index, denoted by *RVA*. Since we are using aggregated (index) data, the returns can be considered the returns of a representative HF that follows a given investment style,

[47] Among the most important economic stress events covered in our data set, we can recount here the US subprime mortgage crisis that erupted in August 2007, the Lehman Brothers' moment of September 2008, the start of the European sovereign debt crisis in May 2010, the three quantitative easing (QE) programmes implemented by the FED (e.g. 2008, 2009, and 2012), the liquidity support programmes of the ECB, and so on.

for example *ED*, *EH*, *M*, or *RVA*. In a robustness check, we employ nine HF sub-styles (of the four main styles mentioned above), for which weekly data are available from the same data source.[48]

We compute the excess HF weekly returns using a risk-free rate proxy. To better reflect the HFs' investable universe and have a basis for comparison among them, we estimate specifications with *global* (and not just US-based) risk-free rates and *global* pricing factors, which we download from Kenneth R. French's website.[49] For simplicity, we use the CAPM model as a workhorse specification throughout the paper, but our approach is robust and can accommodate other factors as well.[50]

There is a continuous and growing interest in the literature in how to measure risk and uncertainty. We take no stand on which measure is best and use a range of indicators that are already available in the literature. We split the various available risk and uncertainty measures into three broad groups. A majority of these measures refer to the US market, as it is the most liquid and sophisticated financial market worldwide. Nevertheless, we include risk and uncertainty measures that have global or European coverage, the latter being particularly useful in exposing some important stress events originating in Europe during our sample period.

Group A: Uncertainty measures based on media sources

There has been a significant increase in the number of available indicators measuring uncertainty based on (text) information accessible via the Internet from various newspapers and other media sources. The success of these indicators seems to come from their ability to reflect, with high frequencies, agents' behaviour in relation to news, events, and other media-related factors (see, among others, Pastor and Veronesi, 2012; Da et al., 2014; Brogaard and Detzel, 2015; Baker et al., 2016). For our empirical analysis, we focus on three measures, which come from Baker et al. (2016) and are based on the frequency of some relevant keywords appearing in newspapers and major media sources in the US and across the globe. The first two measures refer to the US and the third is a global one; the three

[48] HFR utilises a methodology based on certain well-defined, predetermined rules and objective criteria to select and rebalance index components and maximise the representation of the HF investable universe. The construction of each index employs state-of-the-art quantitative techniques and qualitative analysis (i.e. multi-level screening, cluster analysis, Monte Carlo simulations, optimisation techniques, etc.), which ensure that each index is a pure representation of its corresponding HF investment style. A detailed description of HFR styles and sub-styles can be found in Table B1, in Appendix B. More details of the description of the HFR investment styles can be found at: https://www.hedgefundresearch.com/hfrx-index-characteristics.

[49] Using US-based pricing factors does not significantly alter the results or the conclusions of our analysis.
Source : http://mba.tuck.dartmouth.edu/pages/faculty/ken.french/data_library.html.

[50] As an alternative specification, the three-factor model of Fama and French (1993) produces very similar results (which we do not present here to save space), because betas are not significantly altered by the inclusion of additional (pricing) factors.

indicators are: economic policy uncertainty, denoted by *EPU*; US equity market uncertainty, denoted by *EQU*; and global political risk, denoted by *GPR*.[51]

Group B: Market-based risk measures

This second group refers to market-based indicators of risk, which are commonly used in many empirical exercises. We firstly use the *VIX* (i.e. Volatility Index) and *TYVIX* (i.e. Treasury Yields Volatility Index), which are provided by the Chicago Board Options Exchange (CBOE) and are available from the Bloomberg database. Even though the *VIX* is the most followed measure of implied volatility, commonly referred to as "the fear index", the *TYVIX* is equally important for active players like HFs, because it covers the most liquid segment of the financial market, that is, the fixed-income market. While equity volatility (*VIX*) can be specified exogenously, government bond volatility needs to fulfil "no-arbitrage" restrictions and to be consistent with the dynamics of the whole yield curve.

We also include in this group the *volatility risk premium* (or variance premium) denoted VP, which is defined as the difference between the ex-ante risk-neutral expectation of the future return variation and the ex-post realised return variation over a specific period (we chose $n_m = 10$ days), in line with Bollerslev et al. (2009). The variance premium is given by $VP_{i,t} = IV_{i,t} - RV_{i,t}$. The first term ($IV_{i,t}$) is proxied by the square of the respective model-free implied volatility index, in our case the *VIX*, while the second term is proxied by the ex-post realised return variation of the underlying index, in our case the S&P 500 Index.[52] The term $IV_{i,t}$ denotes the recorded closing value of the implied volatility index squared for the last trading day of period t, which also represents the market participants' expectation of the future realised variance of the underlying benchmark index in time period t. The term $RV_{i,t}$ denotes the ex-post realised variance, calculated as $RV_{i,t} = \frac{252}{n_m}\sum_{t=1}^{n_m}(r_i)^2$, where r is the daily return of the underlying S&P 500 Equity Index and $n_m = 10$ is the number of trading days in a year (i.e. we use the 252-day counting convention). Higher levels of implied volatility refer to upcoming volatile periods. However, historical observations show that implied volatility tends to overestimate future realised volatility, as most portfolio managers generally dislike variance; the volatility risk premium is shown to be a good proxy for market sentiment (Bollerslev et al., 2009).

Group C: Constructed measures of (systemic) risk

For the US market, a commonly used measure to track systemic stress is the *Financial Stress Index* (*FSI*), constructed by the St. Louis Fed and available on a weekly basis. It is an equal-variance weighted

[51] All the data and methodology notes are available at www.policyuncertainty.com. Only *EQU* is available on a daily basis; therefore, we interpolate the other two, that is, *EPU* and *GPR*, using the Denton method from a monthly into a weekly frequency, relying on the *EPU* intra-month variation.

[52] We compute the *VP* in relation only to the *VIX* and not to the *TYVIX*, as the market for derivatives based on the *TYVIX* is less liquid than the market for derivatives based on the *VIX*.

average of eighteen explanatory variables, capturing various aspects of risk and uncertainty in different segments of the market.[53] The *FSI* is constructed using principal component analysis, in which it is assumed that financial stress is the primary factor influencing the co-movement of all these variables. A similar financial stress indicator (with weekly availability) is the *Composite Indicator of Systemic Stress* (*CISS*), which captures instability in the financial system of the euro area (Hollo et al., 2012). The aggregation method takes into account the time-varying cross-correlations between the sub-indices. Therefore, the *CISS* puts relatively more weight on situations in which stress prevails in several market segments simultaneously.[54]

Finally, Diebold and Yilmaz (2009, 2014) propose a set of financial stress measures that, within the relevant literature, are normally referred to with the label *connectedness*. Grounded in modern network theories but drawing on variance decomposition methods, these measures quantify the spillovers arising between financial intermediaries, and between financial markets for various instruments. The three measures of *connectedness* that we retain for our analysis[55] are computed for: (i) global equity markets (denoted *ConnEQ*); (ii) global foreign exchange markets (denoted *ConnFX*); and (iii) global sovereign bond markets (denoted *ConnSB*).

The list of uncertainty and risk measures described above is inherently limited. However, the relevance of our empirical approach is not restricted by the list of selected indicators, which can obviously be expanded with new additions. Nevertheless, the selected measures cover a broad range of indicators from different sources and based on different methodologies. Moreover, all these risk and uncertainty measures have high (absolute) values for skewness and kurtosis, confirming their high sensitivity to stress events.[56]

4. EMPIRICAL ANALYSIS

Using the theory-based CAPM, one can specify the excess (over the risk-free rate) HF returns as:

$$\left[RHF - r_f\right]_t = \alpha + \beta * \left[Mkt - r_f\right]_t + \epsilon_t \tag{4}$$

[53] Three main categories of indicators are included: (a) interest rates (e.g. federal funds rates, short- and long-term Treasury rates, corporate bond yields, etc.); (b) yield spreads; and (c) other indicators (e.g. market volatility indices). Source: https://fred.stlouisfed.org/series/STLFSI.

[54] The *CISS* includes fifteen market-based financial stress measures for the financial intermediary sector, money markets, equity markets, bond markets, and foreign exchange markets.

[55] Daily data are available since 2004 and are downloaded from http://financialconnectedness.org/data.html.

[56] A detailed description of the summary statistics can be found in Table B2 from Appendix B.

where RHF is the weekly return on a HF-style index portfolio, r_f is the risk-free rate return, Mkt is the return on the market portfolio, t is the time index and ϵ_t is the error term. The factor loading, denoted by β, would capture the portfolio's exposure to the market portfolio, while α is a measure of the portfolio's abnormal returns.

Coefficient β is a well-known measure used in portfolio management. Blocher and Molyboga (2017) and Agarwal et al. (2018) argue that HF clients prefer to use simple models to evaluate HFs performance rather than models with many (and more complex) pricing factors. Intuitively, it can be inferred simply by comparing the observed returns of a specific investment strategy with the returns of a broad market index. There is a wide consensus that most actively managed investment funds, particularly HFs, face very few investment constraints, and hence their exposures to pricing factors are essentially dynamic; that is, the factor loadings are not constant over time (see Billio et al., 2012; Savona, 2014a and 2014b; Racicot and Theoret, 2016; Amisano and Savona, 2017). In this context, our empirical framework can be split into two steps. In the first step, we apply data filters to extract the time-varying counterparts of β's from equation (4), the so-called *betas*, from the observed (excess) HF returns. In the second step, we employ a series of multivariate models that allows us to analyse the complex interactions between *betas* and changes in the information set, which we proxy using various risk and uncertainty measures. Appendix A relaxes the assumptions required for inference in a CAPM structure by considering a non-parametric filter for *beta;* according to our findings, in the case of HF returns any estimate of *beta* is likely to have nonlinear interactions with some of the risk/uncertainty measures described in section 3.

4.1 Discrete and time-continuous filters of betas

In this section we rely on the mathematical formulation of our conceptual idea (as detailed in section 2) casting the representative HF decision-making problem in a Bayesian framework. The implementation is a juxtaposition of: (i) a discrete filter, implemented as a Markov Switching model, used to proxy for the less accurate *beta* predictions; and (ii) a continuous-time filter, implemented as a Kalman filter on a time-varying coefficients version of CAPM, used to proxy for more accurate predictions. The recursive nature of the two filters is an essential property that allows us to use them in HF decision-making analysis.

As a first filter, we adopt a MS specification of the CAPM that allows the estimation of the hidden Markov chain process driving the parameters of the model between some discrete, unobserved states or regimes. We then use the regime-dependant filtered probabilities and estimated coefficients to compute a time-varying *beta* as a proxy for the *expected beta*. For the second filter, we use a TVC Kalman filter on the same data set and model structure to derive a time-varying *beta* that will reflect the *realised beta*, or the high-level accuracy prediction of *beta*. The observation equation of the filter specifies the excess HF returns as a function of market excess returns, while the state equations define the time-varying coefficients specified as pure unit root processes. We intentionally keep the state equations as simple as possible to avoid making any additional assumptions at this point about HF managers' ability to time their strategies to market conditions (e.g. as in Ferson and Schadt, 1996; Racicot and Theoret, 2016).

The first step of our empirical approach, therefore, consists in applying two data filters on the time-varying version of equation (4):

Filter 1: *Expected beta* to be inferred from a Markov Switching (MS) specification

$$\left[RHF - r_f\right]_t = (alpha)_{s(t)} + (beta)_{s(t)} * \left[Mkt - r_f\right]_t + \varepsilon_{s(t)} \tag{5}$$

where s_t is a first-order unobserved Markov chain with two regimes and a transition matrix P (with elements on each row summing to 1) such that $s_t = P * s_{t-1}$, which also acts as a (discrete) state equation in the case that the model is considered in its state–space form. The $\varepsilon_{s(t)}$ are error terms with $\varepsilon_{s(t)=1} \sim N(0, \boldsymbol{\sigma}_{\boldsymbol{s=1}}{}^2)$ and $\varepsilon_{s(t)=2} \sim N(0, \boldsymbol{\sigma}_{\boldsymbol{s=2}}{}^2)$. Due to its discrete nature and the approximations inherent in its maximum likelihood derivation (see Kim, 1994), the MS implies information losses in the data inference process and should be a poorer fit to the observed HF returns, indirectly providing us with a measure HFs' *less accurate* predictions.

Filter 2: TVC Kalman filter for the *realised beta*

$$observation\ eq: \left[RHF - r_f\right]_t = (alpha)_t + (beta)_t * \left[Mkt - r_f\right]_t + \varepsilon_t \tag{6}$$

$$state\ eq: (beta)_t = (beta)_{t-1} + \varepsilon_{b,t}, \tag{7}$$

$$state\ eq: (alpha)_t = (alpha)_{t-1} + \varepsilon_{a,t} \tag{8}$$

where $\varepsilon_t \sim N(0, \boldsymbol{\sigma}^2)$. We use the filtered state $(beta)_t$ from equation (7) as a measure of the *realised beta*. Table 1 reports the estimation results of the two filters.

Table 1. CAPM model specification estimated on excess HF weekly returns

Style	ED		EH		M		RVA	
	Expected *beta* filter	Realised *beta* filter	Expected *beta* filter	Realised *beta* filter	Expected *beta* filter	Realised *beta* filter	Expected *beta* filter	Realised *beta* filter
Panel A: Markov Switching model estimates used to derive *expected betas*								
LogL	-470.06		-610.65		-689.53		-301.95	
AIC	3.69		3.17		2.93		4.58	
Regime 1: (more persistent and lower risk regime)								
$(alpha)_{s=1}$	0.146***		0.103***		0.053*		0.052***	
$(beta)_{s=1}$	0.103***		0.220***		-0.013		0.039***	
$\sigma_{s=1}$	0.363***		0.384***		0.560***		0.244***	
p_{11}	0.966***		0.929***		0.970***		0.981***	
Regime 2 (alternative regime)								
$(alpha)_{s=2}$	-0.422***		-0.326***		-0.481*		-0.221**	
$(beta)_{s=2}$	0.154***		0.194***		0.281*		0.146***	
$\sigma_{s=2}$	0.895*		0.964		1.642***		1.147**	
p_{22}	0.883***		0.871***		0.729*		0.934***	
Panel B: Kalman filter, used to derive *realised betas*								
LogL		-591.7		-721.9		-800.5		-512.5
AIC		-6.77		-7.16		-7.37		-6.48
σ		0.269***		0.436***		0.550***		0.152***

Note: Estimation results from equation (5) are displayed in panel A, while estimation results for the system of equations (6)-(8) are displayed in panel B. Both panels include the log-likelihood value at the optimum (LogL) and the Akaike Information Criterion (AIC). The $\sigma_{s=1}$ and $\sigma_{s=2}$ in panel A are regime-specific standard deviations of the model estimated in equation (5), while p_{11} and p_{22} denote the diagonal elements of transition matrix P. Panel B displays only the standard deviation of the observation equation, σ, which is comparable to the standard deviation of the model in panel A; the other coefficients are not displayed to save space. The four HF investment styles are specified in the first row of the table: Event Driven, denoted by ED; Equity Hedge, denoted by EH; Macro/CTA, denoted by M; and Relative Value Arbitrage, denoted by RVA. The estimation sample runs from the first week of July 2004 to the last week of May 2017. The (*), (**), and (***) denote coefficients' statistical significance at the 10%, 5%, and 1% levels, respectively.

It is worth noting that the MS specification allows both the coefficients and the equation's variance to vary over time across the two regimes, meaning that we are explicitly aiming at getting a good fit for the HF returns; the TVC Kalman filter instead allows only the coefficients (not the variance) to vary over time. For the MS estimates, with only one exception (style *M*, regime 1), the coefficient of the market factor is always statistically significant in both regimes. The constant term is always statistically significant, positive in regime 1 but negative in regime 2; the persistence of regime 1 is always greater that the persistence of regime 2, according to the estimated probabilities, p_{11} and p_{22}, respectively (which are the diagonal elements of the matrix P). Note that, based on the AIC criterion, the MS model always provides a worse fit than the TVC Kalman filter, a finding in line with the idea that MS provides a smoother (i.e. less accurate) perspective on the data. Figure 1 plots the *realised* and *expected betas* for each HF style, highlighting the partial overlap between the two measures. Moreover, there is interesting overlapping in terms of timing, with periods from 2007 to 2008 being identified as belonging to the high-volatility regime by all HF styles.

Figure 1. *Expected* and *realised betas*, by HF investment style

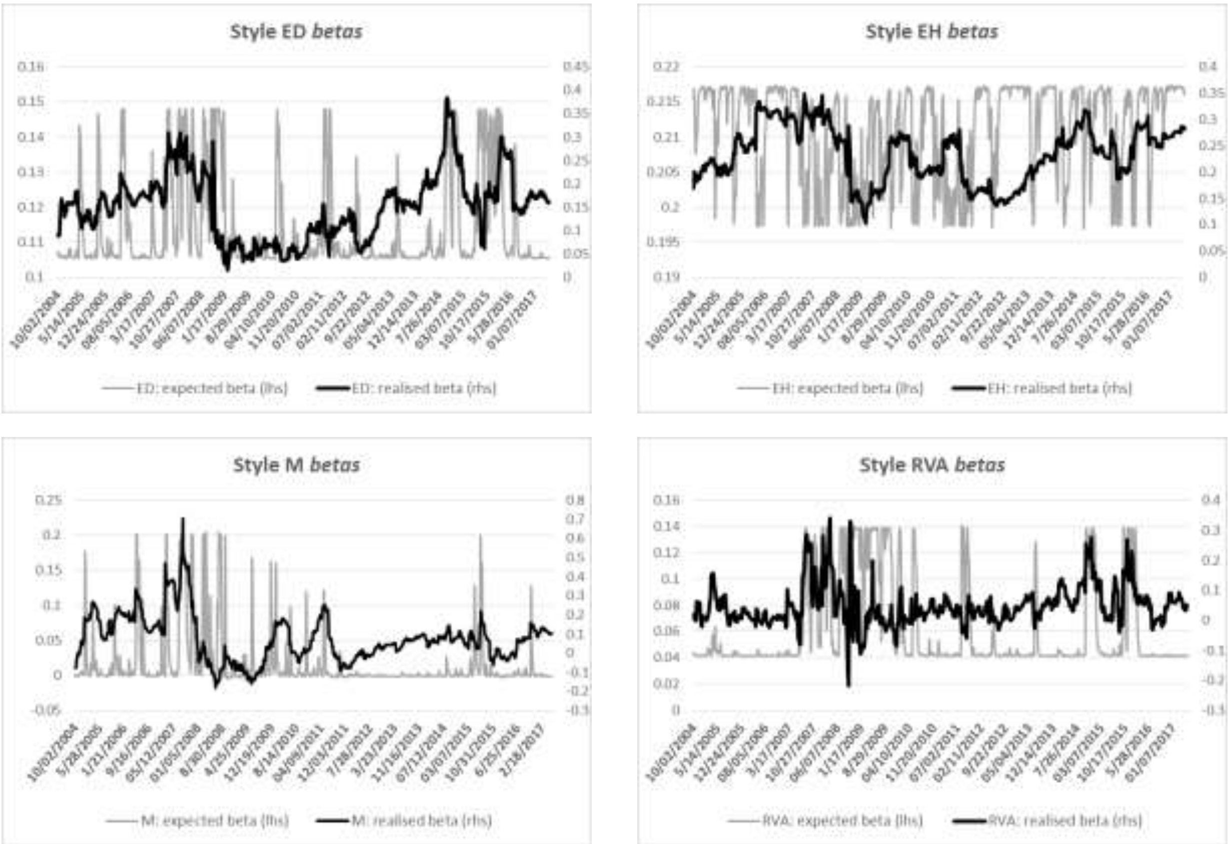

Note: The figure displays the *expected* and *realised betas* filtered using (i) a CAPM specification with coefficients switching between two unobserved regimes that follow a Markov Chain and (ii) a CAPM specification with time-varying coefficients set in a state-space and filtered using a standard Kalman filter, respectively. The four HF investment styles are specified in the in the titles: Event Driven, denoted by ED; Equity Hedge, denoted by EH; Macro/CTA, denoted by M; and Relative Value Arbitrage, denoted by RVA. We discard the first three months of data as a burn-in period, given the well-known erratic dynamics of the filtered states in a Kalman filter during the initial periods. Accordingly, the effective estimation sample that will be used in the empirical section starts with the first week of October 2004 and ends with the last week of May 2017, just as displayed in the figure above.

4.2 Multivariate analysis: Causal influences

The relationships between *betas* and various risk and uncertainty measures are neither simple, nor unidirectional, as highlighted by the application of the non-parametric filter ANOVA presented in Appendix A. Market moves give rise to risk-taking, hedging, and safe-haven motivations that can differ from one investor to another. Moreover, HFs can amplify uncertainty and increase the market risk levels through their trading strategies, leveraged bets and portfolio exposures vis-à-vis other investors (Fung and Hsieh, 2006). In addition, using various data filters to measure the *betas* (which is a standard procedure in the literature) does not necessarily help in disentangling the actual causality influences, due to the possible impact of unobserved factors on realised portfolio returns.

To mitigate these concerns in an analysis of causal inferences, we employ the partial Granger causality (GC) approach set in the time domain, pioneered by Guo et al. (2008), and applied for example in Philippas and Dragomirescu-Gaina (2016). This approach, which is more robust to model misspecification and the omission of other relevant factors, allows us to isolate all traces of common exogenous (measured) factors and latent (unmeasured) factors, assuming that they all have simultaneous effects on all the observed components of the system. Hence, we account for any exogenous and latent (endogenous) factors that can produce misleading results or inaccurate causal inferences in a multivariate setting.

We briefly present the partial GC general framework. Without loss of generality, consider a multivariate process $\boldsymbol{W_t}$ with the following autoregressive formulation:

$$\boldsymbol{B}(L)\boldsymbol{W}_t = u_t \tag{9}$$

where L is the lag operator, and $\boldsymbol{B}$ is a polynomial matrix of L; in particular $\boldsymbol{B}(0) = \boldsymbol{I}$, the identity matrix; $E(u_t) = \boldsymbol{0}$ and $var(u_t) = \boldsymbol{\Sigma}$. The process $\boldsymbol{W}_t$ is an aggregation of three components (measured variables or groups of variables), denoted by $\boldsymbol{x}_t$, $\boldsymbol{y}_t$, and $\boldsymbol{z}_t$. Accordingly, $\boldsymbol{W}_t = [\boldsymbol{x}_t \quad \boldsymbol{y}_t \quad \boldsymbol{z}_t]'$ can be used to model the causality influences arising between the first two components, for example from $\boldsymbol{y}_t$ to $\boldsymbol{x}_t$, conditional on the third one, for example $\boldsymbol{z}_t$.

The error term u_t can also be decomposed into a noise term, e_t, together with an exogenous term, denoted as E_t, and a latent variable term, Λ_t, which depends on the process $\boldsymbol{W_t}$. Thus, the unrestricted multivariate VAR model of the process $\boldsymbol{W}_t$ involving the factors $\boldsymbol{x}_t$, $\boldsymbol{y}_t$, and $\boldsymbol{z}_t$ can be written (in matrix form) as:

$$\boldsymbol{B}(L)\boldsymbol{W}_t = e_t + E_t + \boldsymbol{B}^*(L)\Lambda_t \tag{10}$$

where $\boldsymbol{B}^* = [B_x^* \quad B_y^* \quad B_z^*]'$ is a matrix of polynomials in the lag operator L, $e_t = [e_{x,t} \quad e_{y,t} \quad e_{z,t}]'$ is the noise term, $E_t = [E_{x,t} \quad E_{y,t} \quad E_{z,t}]'$ is the exogenous input, and $\Lambda_t = [\varepsilon_x^{\Lambda} \quad \varepsilon_y^{\Lambda} \quad \varepsilon_z^{\Lambda}]'$ represents the latent variables that cannot be measured in the system, but which are normally distributed random vectors, with the vectors E_t and Λ_t independent of e_t, and the variance–covariance matrix of the vector autoregressive unrestricted model denoted by $\boldsymbol{\Sigma}_{(x,y,z),t}$. Using the Wold representation, the latent variables can be represented as the summation of normally distributed random inputs; therefore, using an R superscript to denote the restricted VAR model for $\boldsymbol{W}_t^R = [\boldsymbol{x}_t\ \boldsymbol{z}_t]$, we can write:

$$\boldsymbol{B}^R(L)\boldsymbol{W}_t^R = e_t^R + E_t^R + \boldsymbol{B}^R(L)\Lambda_t^R \tag{11}$$

where the noise variance–covariance matrix of the restricted model is $\boldsymbol{S}_{(x,z),t}$.

In the time domain setting, the partial GC shows the causal influence of one component, x, on another, y, conditioned on the other component z, using a partial correlation specification that eliminates the influence of the common exogenous inputs and any latent variables. Thus, the test statistic takes the following expression:

$$F_{y\to x/z} = ln\left[\frac{\boldsymbol{S}_{xx} - \boldsymbol{S}_{xz}\boldsymbol{S}_{zz}^{-1}\boldsymbol{S}_{zx}}{\boldsymbol{\Sigma}_{xx} - \boldsymbol{\Sigma}_{xz}\Sigma_{zz}^{-1}\boldsymbol{\Sigma}_{zx}}\right] \tag{12}$$

where $\boldsymbol{S}_{xx}$, $\boldsymbol{S}_{xz}$, $\boldsymbol{S}_{zz}$, and $\boldsymbol{S}_{zx}$ are corresponding elements (or partitions in the multivariate case) of the $\boldsymbol{S}_{(x,z)}$ matrix, while $\boldsymbol{\Sigma}_{xx}$, $\boldsymbol{\Sigma}_{xz}$, $\boldsymbol{\Sigma}_{zz}$ and $\boldsymbol{\Sigma}_{zx}$ are elements (or partitions) of the $\boldsymbol{\Sigma}_{(x,y,z)}$ matrix.

Table 2 presents the results derived from applying the partial Granger causality approach to a multivariate specification that includes all four (*expected* or *realised*) *betas* and one uncertainty/risk measure, which is displayed in the first column under the general label *factor*. Including all four *betas* should accommodate all the interactions that might arise from the fact that different HFs might simultaneously implement changes to their investment strategies in response to unexpected market moves and shocks. Similar results are obtained even if we consider a bivariate setting that pairs each *beta* with each risk or uncertainty measure.

Table 2. Partial Granger causality influences between *expected/realised betas* and factors

Factor	Expected beta → *Factor*	Expected beta → *Factor*	Expected beta → *Factor*	Expected beta → *Factor*	*Factor* → Expected beta	*Factor* → Expected beta	*Factor* → Expected beta	*Factor* → Expected beta
Style	ED	EH	M	RVA	ED	EH	M	RVA
EQU		*Yes*		*Yes*		*Yes*	*Yes*	
EPU			*Yes*			*Yes*		
GPR		*Yes*					*Yes*	
VIX			*Yes*	*Yes*	*Yes*	*Yes*	*Yes*	*Yes*
TYVIX				*Yes*		*Yes*		*Yes*
VP	*Yes*	*Yes*	*Yes*	*Yes*			*Yes*	*Yes*
CISS	*Yes*	*Yes*			*Yes*		*Yes*	
FSI	*Yes*	*Yes*	*Yes*	*Yes*	*Yes*		*Yes*	*Yes*
ConnEQ		*Yes*	*Yes*	*Yes*			*Yes*	
ConnFX		*Yes*	*Yes*	*Yes*	*Yes*		*Yes*	*Yes*
ConnSB	*Yes*	*Yes*	*Yes*				*Yes*	

Factor	Realised beta → *Factor*	Realised beta → *Factor*	Realised beta → *Factor*	Realised beta → *Factor*	*Factor* → Realised beta	*Factor* → Realised beta	*Factor* → Realised beta	*Factor* → Realised beta
Style	ED	EH	M	RVA	ED	EH	M	RVA
EQU			*Yes*		*Yes*	*Yes*	*Yes*	
EPU			*Yes*	*Yes*	*Yes*		*Yes*	*Yes*
GPR	*Yes*	*Yes*	*Yes*		*Yes*		*Yes*	*Yes*
VIX			*Yes*		*Yes*	*Yes*	*Yes*	*Yes*
TYVIX			*Yes*					
VP								
CISS	*Yes*	*Yes*	*Yes*		*Yes*	*Yes*		*Yes*
FSI			*Yes*		*Yes*		*Yes*	*Yes*
ConnEQ					*Yes*	*Yes*	*Yes*	*Yes*
ConnFX		*Yes*	*Yes*	*Yes*	*Yes*	*Yes*	*Yes*	*Yes*
ConnSB		*Yes*	*Yes*	*Yes*	*Yes*	*Yes*	*Yes*	*Yes*

Note: Cells display the label "Yes" when we cannot reject the existence of partial Granger causality using 5% as the confidence level; a grey cell is displayed otherwise. The arrow in the first row indicates the direction of causality influences. The four HF investment styles are specified in the second row of the table: Event Driven, denoted by ED; Equity Hedge, denoted by EH; Macro/CTA, denoted by M; and Relative Value Arbitrage, denoted by RVA.

The results show interesting patterns of influence that differ, sometimes significantly, between *expected* and *realised betas*. We cannot infer a unique causal ordering that holds for all possible *beta*–uncertainty/risk pairs, but some qualitative results stand out. Firstly, only *VIX* affects all *expected* and *realised betas*, providing us with evidence in favour of its prominent role (see also the results obtained in Appendix A). Secondly, starting with the left half of the table, we see many more significant outgoing causality influences from *expected betas* toward uncertainty/risk factors, rather than outgoing causality influences from *realised betas* toward uncertainty/risk factors, except those from group A (based on media sources). This finding gives us a first hint that less accurate predictions can lead to (actions with) negative feedbacks on (market-based and systemic) risk/uncertainty measures. Thirdly, as we move on to the right half of the table, we find many more significant incoming causality influences toward *realised betas*, rather than toward *expected betas*, stemming from uncertainty/risk factors, especially those from groups A and C. This reveals the higher information content of *realised betas*, pointing to their ex-post, rather than ex-ante, nature. These last two findings are reassuring, because the data too seems to reflect the particular differences in the construction of the two filters. Note that none of the uncertainty/risk factors was included in the dataset used to filter the *betas* (i.e. a standard CAPM-based specification for excess HFs returns).

Using the same approach, we bundle together all of our risk and uncertainty measures into a model specification that can help us to understand better the existing patterns of causality influences among them. Table 3 below summarises the results of this exercise. Some of our indicators are very sensitive to incoming causality influences, but at the same time indifferent to influences from most others; for example, the *FSI* is influenced by almost every factor, except *ConnEQ* and *ConnFX*, but influences none. On the contrary, indicators like *TYVIX*, *VIX*, *CISS*, and *ConnSB* seem to influence many other indicators but are influenced by only a few others, thus being some of the most "exogenous" indicators in our list. This is an important finding in light of the results obtained and discussed in the rest of the empirical section.

Table 3. Partial Granger causality influences among risk and uncertainty measures

Factor	→ EQU	→ EPU	→ GPR	→ VIX	→ TYVIX	→ VP	→ CISS	→ FSI	→ Conn EQ	→ Conn FX	→ Conn SB
EQU→	-							*Yes*	*Yes*	*Yes*	
EPU→	*Yes*	-						*Yes*			
GPR→	*Yes*		-					*Yes*			
VIX→	*Yes*			-		*Yes*	*Yes*	*Yes*	*Yes*	*Yes*	
TYVIX→	*Yes*	*Yes*	*Yes*		-	*Yes*	*Yes*	*Yes*	*Yes*	*Yes*	*Yes*
VP→						-	*Yes*	*Yes*	*Yes*	*Yes*	*Yes*
CISS→	*Yes*	*Yes*		*Yes*		*Yes*	-	*Yes*	*Yes*	*Yes*	
FSI→								-			
ConnEQ→									-	*Yes*	*Yes*
ConnFX→									*Yes*	-	*Yes*
ConnSB→	*Yes*		*Yes*		*Yes*	*Yes*		*Yes*	*Yes*	*Yes*	-

Note: The first column displays the "senders", which represent the origin of the estimated causality influence, while the first row displays the "receivers". Cells display the label "Yes" when we cannot reject the existence of partial Granger causality using 5% as a confidence level; a grey cell is displayed otherwise. The four HF investment styles are specified in the second row of the table: Event Driven, denoted by ED; Equity Hedge, denoted by EH; Macro/CTA, denoted by M; and Relative Value Arbitrage, denoted by RVA.

4.3 Multivariate analysis: Vector Autoregressive (VAR) models

We claim that the decision-making process of a representative HF can be framed as a trade-off between prediction accuracy and reaction speed, a trade-off that becomes binding particularly during turbulent market periods. We apply this idea to data by specifying a vector autoregressive model with three variables: the *expected beta* denoted as $(beta)_{s(t)}$, the *realised beta* or $(beta)_t$, and one risk/uncertainty measure or $factor_t$. Our intention is to reveal the main differences in the adjustment speed of the two *betas* to sudden changes in the information set, which we proxy using unexpected shocks in the $factor_t$. The model includes both *betas* to reflect the fact that both strategies are (hypothetically) available to any HF at any given moment; therefore, the endogenous vector y_t is specified as: $y_t = [factor_t \quad (beta)_t \quad (beta)_{s(t)}]'$.

We estimate several VAR specifications and derive our findings based on the analysis of generalised impulse response functions (GIRFs) to unexpected shocks in each uncertainty/risk measure (Koop et al., 1996; Pesaran and Shin, 1998). Order-invariant GIRFs are better suited to tracking the dynamics of shocks through a system of simultaneous equations, especially when there is no prior understanding of the exogeneity rankings between the endogenous variables of the model. If we were to accept that estimating the empirical counterpart of an unobserved mental process implies that changes in the information set are exogenous[57], i.e. *factor* ordered first, then there would still be no difference between Choleski-based impulse responses and GIRFs to a shock in the first equation, which is our main focus.

Lag lengths correspond to approximately one month and were chosen based on standard selection criteria (e.g. Akaike Information Criterion); if necessary, the number of lags was increased to insure lack of serial correlation in residuals. The effective estimation sample runs from the first week of October 2004 to the last week of May 2017. All estimated models are stable, with roots inside the unit circle. It should be noted that since our multivariate models include generated regressors (i.e. the two *betas*), the confidence intervals might be inaccurate. To obtain robust confidence intervals we bootstrap the estimated VARs for 5000 times following the approach proposed in Kilian (1998); in addition, we use a rather conservative confidence level of 95% (or +/- 2 standard deviations) for GIRFs to gauge the statistical significance of the results.

Table 4 displays the sign and the horizon intervals for which the GIRFs are statistically significant. These estimates can help us gauge the timing and direction of changes in *expected* and *realised betas* in reaction to uncertainty/risk shocks.

[57] Most readers would disagree though, and rightly so, given the extensive evidence showing that HFs play an essential role in the transmission of systemic shocks (Fung and Hsieh, 2006; Racicot and Theoret, 2016). In addition, our evidence in section 4.2 does not support a unique causal ordering that can be applied to all beta-to-uncertainty/risk pairs.

Table 4. Estimated VARs: GIRFs to an unexpected positive uncertainty/risk shock

Significant GIRFs (horizon)	**Style ED**		**Style EH**		**Style M**		**Style RVA**	
	Expected *beta*	Realised *beta*	Expected *beta*	Realised *beta*	Expected *beta*	Realised *beta*	Expected *beta*	Realised *beta*
Group A: Uncertainty measures								
EPU	1-2 (+)	5-50 (-)	1-8 (-)	4-50 (-)	n.s.	2-50 (-)	n.s.	n.s.
EQU	1-11 (+)	12-50 (-)	1-15 (-)	4-50 (-)	n.s.	6-50 (-)	n.s.	n.s.
GPR	n.s.	6-50 (+)	n.s.	6-50 (+)	n.s.	n.s.	n.s.	8-9 (+)
Group B: Market-based risk indicators								
VIX	2-21 (+)	12-50 (-)	3-35 (-)	8-50 (-)	2-8 (+)	12-50 (-)	2-50 (+)	n.s.
TYVIX	1-14 (+)	13-50 (-)	1-45 (-)	12-50 (-)	2-3 (+)	11-50 (-)	1-50 (+)	3-4 (+)
VP	1-5 (+)	n.s.	1-2 (-)	n.s.	n.s.	n.s.	n.s.	n.s.
Group C: Computed measures of (systemic) risk								
ConnEQ	1-11 (+)	n.s.	2-13 (-)	n.s.	2-4 (+)	20-50 (-)	2-5 (+)	4-5 (+)
ConnFX	2-7 (+)	7-47 (-)	2-7 (-)	8-50 (-)	11-48 (-)	10-50 (-)	2-4 (+)	n.s.
ConnSB	n.s.	n.s.	n.s.	n.s.	9-50 (+)	n.s.	1-3 (+)	n.s.
FSI	1-13 (+)	35-50 (-)	1-8 (-)	6-37 (-)	n.s.	n.s.	1-50 (+)	2-4 (+)
CISS	1-13(+)	19-50 (-)	1-28 (-)	17-50 (-)	1-3 (+)	8-50 (-)	1-50 (+)	3-4 (+)

Note: The numbers displayed in the table denote the horizon intervals for which the GIRFs are statistically significant, with bootstrapped confidence bands set at +/- 2 standard deviations; 5000 bootstrap replications of the estimated model are used for the confidence interval, following Kilian (1998) approach. The (-) or (+) denotes the sign or direction of the GIRFs in the specified interval. The **n.s.** label in the table means that, given the confidence bands, the GIRFs are not significant for (at least) two consecutive observations. The maximum horizon is truncated at 50 weeks (approximately 1 year). The four HF investment styles are specified in the first row of the table: Event Driven, denoted by ED; Equity Hedge, denoted by EH; Macro/CTA, denoted by M; and Relative Value Arbitrage, denoted by RVA. The corresponding GIRF figures are reported in Figure B2, Appendix B.

The results summarised in Table 4 show, with very few exceptions, a stronger and faster reaction for the *expected betas*, but a weaker, slower (or delayed) reaction for the *realised betas*. Therefore, relying on less accurate predictions implies faster portfolio adjustments in reaction to sudden changes in the information set. This is our first important result that survives across different estimation and robustness checks. In fact speed is key for market timing, and the HFs literature provides rich empirical evidence that (at least) some HFs have such abilities that allow them to earn extra profits, i.e. positive alpha (see Cao et al., 2013; Bali et al., 2014). With the remarkable exception of the EH style, the generally positive responses seen for *expected betas* suggest that most HFs tend to increase portfolio exposures in volatile markets; this finding in line with Billio et al., (2012), but in contrast to Patton and

Ramadorai (2013), suggesting that differences in HF investment styles are important to consider when analysing the direction of HFs exposure changes during turmoil. In the case of those HFs pursuing Equity Hedge (EH) strategies, for example, *expected betas* show reactions that are generally negative, most likely due to an over-reliance on hedge positions that explore idiosyncratic rather than general market trends. Notwithstanding differences in the direction of their bets, our results show that all HFs would be able to gain speed and thus improve their market timing abilities by adopting strategies with lower levels of accuracy, or to put it differently, by allowing their exposures to change more swiftly during turmoil (e.g. due to shifts in leverage, or exposures to option-like payoffs).

Looking at Table 4 we see further that all HFs styles react to innovations in some particular indicators like *VIX*, *TYVIX*, *ConnEQ*, and *CISS*; more specifically, we find that all *expected betas* react to these same shocks, though we cannot find a similar result in the case of *realised betas*. Although HFs' market exposures can increase or decrease, depending on the adopted investment style, having a common sensitivity to any single specific factor might lead to *simultaneous* reactions in case of large shocks. This is a good enough reason for including these risk indicators on supervisors' watch lists. In fact, *VIX* and *CISS* were also two of the most 'exogenous' indicators already identified in section 4.2 and Appendix A, supporting the idea that they contain valuable information to signal early shifts in HFs' risk appetite and market timing efforts. This potential *simultaneity* in case of large shocks is our second important result from the empirical analysis, with implications for contagion and market stress. King and Wadhwani (1990), and more recently Cipriani and Guarino (2008) advocate the importance of information spill-overs that can lead to contagion when trading activity is correlated across markets, although fundamentals are not necessarily related. Our empirical results are in line with this idea: due to their possible simultaneous actions and reliance on the same signal indicators for market timing, HFs can play a key role in the cross-sectional transmission of market stress and, therefore, contagion.

4.4 Extensions of the model to HF investment sub-styles

To better reflect the heterogeneity of HF investment styles, we replicate the analysis using HFRX indexes at different levels of aggregation, both *below* and *above* the one used so far. More precisely, besides the four main HF investment styles (i.e. *ED*, *EH*, *M*, and *RVA*), we use nine HF *sub*-styles and one *global* index.[58] Our previous first main result is re-confirmed: *expected betas* react more quickly and/or strongly while *realised betas* are slower and weaker in response to an uncertainty/risk shock.

[58] The correspondence between the nine sub-styles and the four main styles can be found in Appendix B, Table B1. Figure B1, in the same Appendix B, plots the *realised* and *expected betas* for each HF sub-style and for the *Global* style. A summary of the empirical results pertaining to this section are available in Appendix B, Table B4.

With respect to the second main result, we showed above that all HF investment styles can react simultaneously to certain shocks, despite important differences in terms of their particular investment strategy, geographical focus, financial markets, assets, and instruments employed. However, when differentiating HFs further into sub-styles, this result becomes less clear, though by a very small margin: there is only one exception out of nine sub-styles, for both *VIX* and *CISS*, while the other early-warning indicators drop out of our narrow list. Suppose all HF styles process information flows from non-overlapping sources; however, in periods with high volatility, signal extraction "mistakes" can be transmitted to all other markets, generating contagion (see King and Wadhwani, 1990; Cipriani and Guarino, 2008; Hasler and Ornthanalai, 2018). Therefore, while having only four main HF styles could guarantee that the non-overlapping assumption holds exactly, a more granular approach could pose methodological challenges, but not as strong as to weaken the main implication about increasing contagion risks.

4.5 Robustness checks

To overcome some concerns regarding the specification used to explain HF excess returns and infer the *betas*, we employ a series of robustness checks. Much of the HF literature explains HF excess returns based on various pricing factors (e.g. Agarwal and Naik, 2004; Carhart et al., 2014). We replicate the empirical analysis above by using a model specification with the three Fama-French factors (the only readily available at a weekly frequency), and even by adding a forth pricing factor, which we proxy by one of the uncertainty/risk measures in our list. Conclusions are similar. In fact, the more pricing factors we include, the better the explanatory power of the model, a dimension on which the TVP Kalman already proved superior compared to MS. Therefore, the success of *expected betas* over *realised betas* in terms of reaction speed does not lie in the explanatory power of the specification used to model HF returns. Instead, it seems to be dependent on the inability of the TVP Kalman filter to reflect the time-varying nature of volatility, which the MS filter deals with directly during its estimation.

To add more information into the filtered *betas*, Savona (2014a and 2014b) proposes a system estimation where time-varying *betas* are a function of some primitive risk signals (i.e. in essence, volatility proxies). We take a more direct approach here by changing the signal-to-noise ratio at the TVP Kalman filter stage. Note that we have specifically adopted a random walk (rather than AR(1) or mean reverting as in Savona, 2014a and 2014b) specification for the state equations to allow for a higher contemporaneous pass-through of volatility signals into the filtered *beta* state (i.e. equation 7). As a robustness check, we allow the signal-to-noise ratio in the TVP filter go to infinity (i.e. recent observations receive more weight during signal extraction), but we end up with noisier *betas* that do not react faster than *expected betas*. These findings highlight the advantages of a discrete filter (see also the results in Bollen and Whaley, 2009) in striking the right balance between inferring a time-varying *beta* (needed to control portfolio exposures) and picking up relevant volatility jumps, which provide the best signals that can help managers in timing the market (Bali et al., 2014; Kacperczyk et al., 2016).

5. SIMULATION EXERCISE

This section presents a simulation exercise where we compare the performances of two hypothetical portfolios, both built based on two benchmark strategies that correspond to extreme accuracy levels. More technical details are relegated to Appendix C. Here we just provide the main intuition behind the construction of these portfolios and strategies, and describe the main results.

The first benchmark strategy, denoted as S1, covers the "perfect accuracy" case and consists in a portfolio constructed such as its (excess) returns track the (excess) market returns at all times, assuming a small but positive and constant *beta* of 0.1 at all times.[59] The second benchmark strategy, S2, corresponds to the "no accuracy" case where we assume managers just place random bets on the market based on a simulated Markov Chain (MC) variable that governs the direction of their bets ('on' or 'against' the market), therefore, allowing for some degree of persistency in the strategy being followed; the value of *beta* for strategy S2 is necessarily time-varying, switching between two discrete values, but its average equals the same value of 0.1 just as for strategy S1.[60] With some inherent simplifications, these two benchmark strategies S1 and S2 are consistent with the two investment strategies discussed so far; to see this, note that the prediction errors for the two *betas* are either zero (for *realised beta*) or a discrete range of values (for *expected beta*), in line with the intuition provided in the introduction.

Next, we construct two hypothetical portfolios that alternate their investing strategy between the two benchmark strategies above, S1 and S2, depending on some information signal that reflects market timing. The empirical results we obtained in the previous section have identified the accuracy–speed trade-off during extreme market stress periods, which were proxied by the uncertainty/risk shocks in the multivariate (VAR) models. To keep things simple in this simulation exercise, we assume that the signal can be extracted from our uncertainty/risk measures, such that one strategy is replaced by the other one whenever the (standardised) value of some uncertainty/risk indicator crosses a certain level (determined according to its distributional properties). We assume that portfolio P^A adopts strategy S1 for most of the time, except for when signals identify volatile periods, prompting a switch to strategy S2; similarly, portfolio P^B adopts strategy S2 most of the time, except for volatile periods when it

[59] A small *beta* value of 0.1 has been selected to approximate the sample average of estimated *realised betas* of the four HFs styles. In the same time, such a low *beta* reflects the idea that HFs aim at having a low correlation with the market in order to attract client flows and present themselves as effective diversification instruments.

[60] We set the upper and lower bounds of *beta* as (0.05, 0.15) to approximate the estimated variation interval between the two states in Table 1, panel A. The transition matrix for the MC is calibrated with an equal persistence of 0.95 for both regimes.

switches to strategy S1. Although these two mixed-strategy portfolios are complementary with each other by construction, their simulated performances are completely different (see Appendix C for more technical details).

There is strong empirical evidence that some HFs are able to time market liquidity (Cao et al., 2013) and uncertainty or, more generally, volatility spikes in financial and macroeconomic variables (Bali et al., 2014; Kacperczyk et al., 2016). Depending on the uncertainty/risk factor we select as a signal source, the timing of switching between the two benchmark strategies will change, as well as their performances. Using the real (excess) market returns over the October 2004 – May 2017 period and a random normally distributed noise added to each period returns, we repeatedly simulate the two portfolios P^A and P^B over a sample that matches the estimation sample length (i.e. 661 weeks/periods). Simulation results based on 5000 replications are displayed in Table 5 below. We report the median value (computed across all simulations) for the various summary statistics instead of the mean value, because the former suffers less impact from any possible outlier (i.e. simulated portfolio with extreme outcomes). We also report the Kolmogorov-Smirnov (henceforth KS) test statistics to facilitate the comparison of any two simulated distributions.

Table 5. Simulation results

Factor	**End of period portfolio value: median (KS test)**	**Standard deviation: median (KS test)**	**Skewness: median (KS test)**	**Kurtosis: median (KS test)**	**Sharpe Ratio: median (KS test)**
Benchmark strategies	S1: 9.4943 S2: 9.5410 ($KS^{S1=S2}$: 0.0160)	S1: 0.5550 S2: 0.5675 ($KS^{S1=S2}$: 0.290*)	S1: -0.1097 S2: -0.1397 ($KS^{S1=S2}$: 0.198*)	S1: 3.3285 S2: 3.5809 ($KS^{S1=S2}$: 0.308*)	S1: 0.0259 S2: 0.0256 ($KS^{S1=S2}$: 0.0192)
Group A: Uncertainty measures					
EPU	P^A: 9.5937 ($KS^{A=S1}$: 0.0112)	P^A: 0.5592 ($KS^{A=S1}$: 0.117*)	P^A: -0.1466 ($KS^{A=S1}$: 0.230*)	P^A: 3.5009 ($KS^{A=S1}$: 0.301*)	P^A: 0.0259 ($KS^{A=S1}$: 0.0122)
	P^B: 9.5449 ($KS^{A=B}$: 0.0128)	P^B: 0.5639 ($KS^{A=B}$: 0.114*)	P^B: -0.1015 ($KS^{A=B}$: 0.236*)	P^B: 3.4102 ($KS^{A=B}$: 0.285*)	P^B: 0.0255 ($KS^{A=B}$: 0.0138)
EQU	P^A: 9.6241 ($KS^{A=S1}$: 0.0082)	P^A: 0.5587 ($KS^{A=S1}$: 0.121*)	P^A: -0.1309 ($KS^{A=S1}$: 0.195*)	P^A: 3.4758 ($KS^{A=S1}$: 0.305*)	P^A: 0.0258 ($KS^{A=S1}$: 0.0086)
	P^B: 9.3694 ($KS^{A=B}$: 0.0126)	P^B: 0.5641 ($KS^{A=B}$: 0.127*)	P^B: -0.1175 ($KS^{A=B}$: 0.178*)	P^B: 3.4093 ($KS^{A=B}$: 0.288*)	P^B: 0.0251 ($KS^{A=B}$: 0.0134)
GPR	P^A: 9.4783 ($KS^{A=S1}$: 0.0074)	P^A: 0.5561 ($KS^{A=S1}$: 0.032*)	P^A: -0.1165 ($KS^{A=S1}$: 0.028*)	P^A: 3.3438 ($KS^{A=S1}$: 0.028*)	P^A: 0.0258 ($KS^{A=S1}$: 0.0088)
	P^B: 9.5139 ($KS^{A=B}$: 0.0108)	P^B: 0.5665 ($KS^{A=B}$: 0.253*)	P^B: -0.1366 ($KS^{A=B}$: 0.179*)	P^B: 3.5638 ($KS^{A=B}$: 0.299*)	P^B: 0.0254 ($KS^{A=B}$: 0.013)
Group B: Market-based indicators					
VIX	P^A: 9.6021 ($KS^{A=S1}$: 0.0108)	P^A: 0.5605 ($KS^{A=S1}$: 0.161*)	P^A: -0.1421 ($KS^{A=S1}$: 0.213*)	P^A: 3.5496 ($KS^{A=S1}$: 0.313*)	P^A: 0.0259 ($KS^{A=S1}$: 0.009)
	P^B: 9.3944 ($KS^{A=B}$: 0.0124)	P^B: 0.5626 ($KS^{A=B}$: 0.103*)	P^B: -0.1054 ($KS^{A=B}$: 0.216*)	P^B: 3.3628 ($KS^{A=B}$: 0.311*)	P^B: 0.0252 ($KS^{A=B}$: 0.012)

TYVIX	P^A: 9.5547 ($KS^{A=S1}$: 0.0106)	P^A: 0.5599 ($KS^{A=S1}$: 0.156*)	P^A: -0.1315 ($KS^{A=S1}$: 0.198*)	P^A: 3.519 ($KS^{A=S1}$: 0.309*)	P^A: 0.0257 ($KS^{A=S1}$: 0.0122)
	P^B: 9.5097 ($KS^{A=B}$: 0.0116)	P^B: 0.5631 ($KS^{A=B}$: 0.107*)	P^B: -0.1165 ($KS^{A=B}$: 0.186*)	P^B: 3.385 ($KS^{A=B}$: 0.304*)	P^B: 0.0256 ($KS^{A=B}$: 0.0128)
VP	P^A: 9.6443 ($KS^{A=S1}$: 0.0094)	P^A: 0.5585 ($KS^{A=S1}$: 0.108*)	P^A: -0.1358 ($KS^{A=S1}$: 0.211*)	P^A: 3.4666 ($KS^{A=S1}$: 0.295*)	P^A: 0.0259 ($KS^{A=S1}$: 0.0088)
	P^B: 9.4396 ($KS^{A=B}$: 0.0106)	P^B: 0.5646 ($KS^{A=B}$: 0.141*)	P^B: -0.1115 ($KS^{A=B}$: 0.204*)	P^B: 3.4274 ($KS^{A=B}$: 0.273*)	P^B: 0.0251 ($KS^{A=B}$: 0.0128)
Group C: Computed measures of (systemic) risk					
connEQ	P^A: 9.5435 ($KS^{A=S1}$: 0.004)	P^A: 0.5554 ($KS^{A=S1}$: 0.0106)	P^A: -0.1101 ($KS^{A=S1}$: 0.0054)	P^A: 3.3278 ($KS^{A=S1}$: 0.0052)	P^A: 0.0259 ($KS^{A=S1}$: 0.004)
	P^B: 9.5683 ($KS^{A=B}$: 0.0174)	P^B: 0.5673 ($KS^{A=B}$: 0.283*)	P^B: -0.1393 ($KS^{A=B}$: 0.198*)	P^B: 3.580 ($KS^{A=B}$: 0.311*)	P^B: 0.0255 ($KS^{A=B}$: 0.0206)
connFX	P^A: 9.5886 ($KS^{A=S1}$: 0.0062)	P^A: 0.5560 ($KS^{A=S1}$: 0.029*)	P^A: -0.1070 ($KS^{A=S1}$: 0.0124)	P^A: 3.3382 ($KS^{A=S1}$: 0.015)	P^A: 0.0259 ($KS^{A=S1}$: 0.0054)
	P^B: 9.4291 ($KS^{A=B}$: 0.016)	P^B: 0.5666 ($KS^{A=B}$: 0.258*)	P^B: -0.1418 ($KS^{A=B}$: 0.204*)	P^B: 3.5772 ($KS^{A=B}$: 0.305*)	P^B: 0.0252 ($KS^{A=B}$: 0.018)
connSB	P^A: 9.5682 ($KS^{A=S1}$: 0.006)	P^A: 0.5554 ($KS^{A=S1}$: 0.0182)	P^A: -0.1104 ($KS^{A=S1}$: 0.0072)	P^A: 3.33 ($KS^{A=S1}$: 0.0044)	P^A: 0.0261 ($KS^{A=S1}$: 0.0062)
	P^B: 9.4584 ($KS^{A=B}$: 0.0154)	P^B: 0.5672 ($KS^{A=B}$: 0.276*)	P^B: -0.1390 ($KS^{A=B}$: 0.196*)	P^B: 3.5808 ($KS^{A=B}$: 0.310*)	P^B: 0.0253 ($KS^{A=B}$: 0.0184)
FSI	P^A: 9.6293 ($KS^{A=S1}$: 0.0094)	P^A: 0.5599 ($KS^{A=S1}$: 0.152*)	P^A: -0.1314 ($KS^{A=S1}$: 0.196*)	P^A: 3.5209 ($KS^{A=S1}$: 0.310*)	P^A: 0.0259 ($KS^{A=S1}$: 0.0094)
	P^B: 9.4888 ($KS^{A=B}$: 0.012)	P^B: 0.5632 ($KS^{A=B}$: 0.110*)	P^B: -0.1166 ($KS^{A=B}$: 0.186*)	P^B: 3.384 ($KS^{A=B}$: 0.304*)	P^B: 0.0255 ($KS^{A=B}$: 0.012)
CISS	P^A: 9.6558 ($KS^{A=S1}$: 0.0078)	P^A: 0.5596 ($KS^{A=S1}$: 0.135*)	P^A: -0.1317 ($KS^{A=S1}$: 0.185*)	P^A: 3.5209 ($KS^{A=S1}$: 0.306*)	P^A: 0.0258 ($KS^{A=S1}$: 0.0116)
	P^B: 9.5223 ($KS^{A=B}$: 0.0132)	P^B: 0.5637 ($KS^{A=B}$: 0.114*)	P^B: -0.1189 ($KS^{A=B}$: 0.168*)	P^B: 3.4012 ($KS^{A=B}$: 0.293*)	P^B: 0.0256 ($KS^{A=B}$: 0.0146)

Note: Table presents simulation results based on 5000 replications of the two hypothetical portfolios P^A and P^B; the two benchmark strategies S1 and S2 are displayed on the second row for comparability. The first column indicates the uncertainty/risk indicator used to determine the level that triggers the switch (i.e. market timing) between the two benchmark strategies. We report the median values, computed across all simulations, for the distribution of the following summary statistics: end of period portfolio value (i.e. cumulated returns), returns' standard deviation, returns' skewness, returns' kurtosis and Sharpe ratio. In parentheses, we report the Kolmogorov-Smirnov test statistics, denoted $KS^{X=Y}$, on the equality of two simulated distributions X and Y; the * denotes rejection of the null that the two distributions are equal at the 5% statistical significance level. More technical details can be found in Appendix C.

Firstly, notice that compared to strategy S1, S2 is more likely to display characteristics associated with HF returns (see Appendix B, Table B1), meaning higher risk (or standard deviation), lower skew and higher kurtosis; these differences are statistically significant according to the KS tests. Higher kurtosis and more negative skew arise from a higher probability of extreme returns, and particularly positive returns, something HFs are aiming to achieve; Dijk et al., (2014) explain how competition for social status can explain a preference for negative skew assets by over-performers, who want to preserve their status. Therefore, it is interesting to see how a simple strategy that randomly bets on the market direction can deliver a distribution of payoffs with higher kurtosis and more negative skew, reflecting therefore some important characteristics of the actual HF returns' distribution.

Secondly, HFs (along with other active players in general) face strong incentives deriving from the possibly large, but asymmetric payoffs they can obtain in times of market stress. Since the previous section revealed speed gains from switching to lower accuracy during market stress, we concentrate on making a comparison (in probability terms) between P^A on the one side, and P^B (as well as S1) on the other side; this is to say that we are only interested in comparing strategies that differ during turmoil periods, while ignoring strategy differences, if any, during calm periods. In all 11 cases, portfolio P^A displays lower risk than P^B (as well as S1), a difference that is statistically significant according to KS tests. More importantly, P^A displays statistically significant higher kurtosis and more negative skewness in 7 cases if we compare it to P^B (and in 8 cases if we compare it to S1). As a confirmation of our previous findings, both VIX and CISS lie among the indicators delivering the best outcomes and market timing; the *connectedness* indicators instead do not seem to provide the best timing, given their much higher persistency and measurement focus. Clearly, HFs would prefer strategies with a lower risk, higher kurtosis and more negative skew, since the probability of high positive returns is much larger (see Dijk et al., 2014). In a majority of cases, therefore, our simulation demonstrates the stochastic dominance of P^A over P^B (as well as over S1) in terms of lower standard deviation and skew, and higher kurtosis; this finding is in line with the trade-off uncovered in the previous section, in which we find that periods of extreme uncertainty/risk are likely to be associated with low accuracy strategies where *beta* dynamics is discrete as in the case of P^A. Overall, the simulation exercise shows that switching exposures from a high to a low accuracy *beta* only during volatile times could deliver payoffs with distributions that are preferable by most HF managers, i.e. negative skew and excess kurtosis.

The simulation results are robust to a series of sensitivity checks we report in Appendix C at the end of this chapter. Obviously, our simulation exercise is an inherent simplification of reality, ignoring many important aspects for HF profitability such as liquidity and trading costs, margin requirements, clients' outflows (redemptions) or inflows etc. In addition, it includes a very simple signal extraction mechanism for market timing, randomly selected direction of the bets (i.e. changes in *beta*), symmetric transitions between regimes, etc. However, it helps in revealing some very clear incentives that HFs have to switch to a low accuracy strategy by timing the market and profiting from its extreme moves; with P^A being a dominating mixed-strategy portfolio (in probability terms), HFs will tend to implement it as fast as possible in order to profit from initial market moves, which are usually the biggest. The fear of missing an opportunity is probably stronger than the fear of losing a bet, also because HFs portray themselves as 'low *beta*' investment vehicles that might face scarce opportunities in trending markets; on the contrary, the increase in correlation across different asset classes could prevent more refined and carefully designed strategies from being implemented in extreme market conditions, leaving therefore many HFs with few choices.

6. DISCUSSION AND CONCLUDING REMARKS

We provide new insights into the investment decision-making process of hedge funds, some of the most active and astute investors. Drawing on the HFR database, we differentiate HFs according to several investment styles, and sub-styles, which define particularities with respect to investment horizon and strategy, preferred asset classes and market segments, etc. Despite large differences in style, their decision-making process boils down to a series of estimates and predictions that are continuously updated with new information. We assume this process can be framed as a trade-off between prediction accuracy and reaction speed, a trade-off that is best revealed during turbulent markets. We cast this process in a Bayesian framework and present a series of empirical analyses and a simulation exercise that concur in providing evidence in favour of this trade-off.

Most HFs claim to generate excess returns that have a low, or even zero *beta* with a broad market index. Moreover, hedging effectiveness in portfolio management relies heavily on the same *beta*, and therefore changes in *beta* can be a good proxy for changes in investment strategy. Using the same data set and a common CAPM model structure, we apply one discrete filter and one time-continuous filter to extract two separate measures of *beta* that entail different levels of prediction accuracy. More specifically, a low-level accuracy prediction we label *expected beta* is filtered from a two-state Markov switching specification, which is more flexible, but provides a worse fit to HF returns than a time-varying coefficient Kalman filter used to infer *realised beta* – the high-level accuracy prediction.

The empirical analysis presented in section 4 shows that less accurate portfolio strategies (implemented as *expected betas*) would adjust more quickly to a series of uncertainty/risk shocks, which we use as proxies for changes in the relevant information set. Meantime, more accurate portfolio strategies (implemented as *realised betas*) would be slow in adjusting to similar shocks. Therefore, we highlight the accuracy–speed trade-off with respect to extreme market conditions, when unknown shocks are most likely to disturb the information set on which investors rely for valuation purposes. In section 5 we justify this result by emphasizing the dominance of a portfolio that switches its *beta* from a high-level to a low-level accuracy in times of extreme market moves, which we identify based on our uncertainty/risk factors. Our simulation exercise proves that return distributions with lower risk, more negative skew and higher kurtosis are associated with a mixed strategy that switches to a low accuracy *beta* during extreme market moves. It is easy to see how these alternating patterns in market exposure (i.e. *beta*) can also deliver the option-like payoff structure outlined in the HFs literature (e.g. in Agarwal and Naik, 2004; Billio et al., 2012). Therefore, market timing remains an essential ingredient for success in the HF industry. Moreover, since opportunities are scarce and might disappear quickly, HFs have strong incentives to precipitate implementation and gain more speed, which implies lower accuracy according to our results. These findings align with theoretical predictions from Kacperczyk et al., (2016) where skilled managers allocate their (time-varying) attention to analysing aggregate risks during turmoil, and idiosyncratic risk during calm periods.

Our empirical results also show that changes in some specific risk measures, mainly *VIX* and *CISS*, contain relevant information that helps HFs to better time the market. Although, in reality, the diversity of methods used by HFs to improve their market timing abilities might be hard to quantify, the two indicators we identify can summarize relevant information to provide regulators with early warnings regarding the upcoming (and possibly *simultaneous*) shifts in risk-appetite across a heterogeneous HF sector. Better counter measures that rely on the same indicators and thresholds could then be designed by market operators and regulators. Our analysis underlines the importance of proper regulations and market designs to prevent the negative consequences stemming from sudden shifts in risk-appetite. Many HFs (and active players as well) face strong incentives deriving from the possibly large, but asymmetric payoffs they can obtain in times of market stress. Market regulators and supervisory authorities have long considered ways to restrict this type of behaviour, and the literature on this topic is extremely rich. More effective early warning indicators and market circuit breakers, counter-cyclical margins and collateral requirements that restrict HFs' ability to place highly leveraged bets during market stress are just some possible examples of intervention tools. Some negative consequences stemming from the cross-sectional transmission of market stress can be reduced as long as such risky strategies are discouraged or simply delayed by means of intervention tools that rely on the same indicators used by active players. In this context, proper identification of such indicators remains key for determining intervention effectiveness.

REFERENCES

Agarwal, V., Green, T.C., Ren, H., 2018. Alpha or Beta in the eye of the beholder: What drives hedge fund flows?. *Journal of Financial Economics* 127, 417–434.

Agarwal, V., Naik, N., 2004. Risks and portfolio decisions involving hedge funds. *Review of Financial Studies* 17, 63–98.

Amisano, G., Savona, R., 2017. Mutual funds dynamics and economic predictors. *Journal of Financial Econometrics* 15, 302–330.

Andrei, D., Hasler, M., 2015. Investor attention and stock market volatility. *Review of Financial Studies* 28, 33–72.

Andrei, D., Hasler, M., 2019. Dynamic attention behavior under return predictability. *Management Science*. *Forthcoming*. Doi.org/10.1287/mnsc.2019.3328.

Aragon, G., 2007. Share restrictions and asset pricing: Evidence from the hedge fund industry. *Journal of Financial Economics* 83, 33–58.

Aragon, G., Hertzel, M., Shi, Z., 2013. Why do hedge funds avoid disclosure? Evidence from confidential 13F filings. *Journal of Financial and Quantitative Analysis* 48, 1499–1518.

Baker, S.R., Bloom, N., Davis, S.J., 2016. Measuring Economic Policy Uncertainty. *Quarterly Journal of Economics* 131, 1593–1636.

Bali, T.G., Brown, S.J., Caglayan, M.O., 2014. Macroeconomic risk and hedge fund returns. *Journal of Financial Economics* 114, 1–19.

Billio, M., Getmansky, M., Pelizzon, L., 2009. Non-parametric analysis of hedge fund returns: New insights from high frequency data. *Journal of Alternative Investments* 12, 21–38.

Billio, M., Getmansky, M., Pelizzon, L., 2012. Dynamic risk exposures in hedge funds. *Computational Statistics & Data Analysis* 56, 3517–3532.

Bissiri, P.G., Holmes, C.C., Walker, S.G., 2016. A general framework for updating belief distributions. *Journal of Royal Statistical Society* – Series B 78, 1103–1130.

Blocher, J., Molyboga, M., 2017. The revealed preference of sophisticated investors. *European Financial Management* 23, 839–872.

Bollen, N.P., Whaley, R.E., 2009. Hedge fund risk dynamics: Implications for performance appraisal. *Journal of Finance* 64, 985–1035.

Bollerslev, T., Tauchen, G., Zhou, H., 2009. Expected stock returns and variance risk premia. *Review of Financial Studies* 22, 4463–4492.

Boyson, N.M., Stahel, C.W., Stulz, R.M., 2010. Hedge fund contagion and liquidity shocks. *Journal of Finance* 65, 1789–1816.

Brogaard, J., Detzel, A., 2015. The asset-pricing implications of Government Economic Policy Uncertainty. *Management Science* 61, 3–18.

Cao, C., Chen, Y., Liang, B., Lo, A.W., 2013. Can hedge funds time market liquidity?. *Journal of Financial Economics* 109, 493–516.

Carhart, M., Cheah, U.W., De Santis, G., Farrell, H., Litterman, R., 2014. Exotic Beta revisited. *Financial Analysts Journal* 70, 24–52.

Cipriani, M., Guarino, A., 2008. Herd behavior and contagion in financial markets. *The BE Journal of Theoretical Economics* 8, 1–56.

Clark, T.E., McCracken, M.W., 2009. Improving forecast accuracy by combining recursive and rolling forecasts. *International Economic Review* 50, 363–395.

Da, Z., Engelberg, J., Gao, P., 2014. The sum of all FEARS investor sentiment and asset prices. *Review of Financial Studies* 28, 1–32.

Diebold, F.X., Yilmaz, K., 2009. Measuring financial asset return and volatility spillovers, with application to global equity markets. *Economic Journal* 119, 158–171.

Diebold, F.X., Yilmaz, K., 2014. On the network topology of variance decompositions: Measuring the connectedness of financial firms. *Journal of Econometrics* 182, 119–134.

Dijk, O., Holmen, M., Kirchler, M., 2014. Rank matters–The impact of social competition on portfolio choice. *European Economic Review* 66, 97–110.

Dudley, E., Nimalendran, M., 2011. Margins and hedge fund contagion. *Journal of Financial Quantitative Analysis* 46, 1227–1257.

Fama, E.F., K.R. French, K.R., 1993. Common risk factors in the returns on stocks and bonds. *Journal of Financial Economics* 33, 3–56.

Fama, E.F., K.R. French, K.R., 2015. A five-factor asset pricing model. *Journal of Financial Economics* 116, 1–22.

Ferson, W.E., Schadt, R.W., 1996. Measuring fund strategy and performance in changing economic conditions. *Journal of Finance* 51, 425–461.

Fung, W.K., Hsieh, D.A., 1997. Empirical characteristics of dynamic trading strategies: The case of hedge funds. *Review of Financial Studies* 10, 275–302.

Fung, W.K., Hsieh, D.A., 2006. Hedge funds: An industry in its adolescence. *Economic Review* 91, 1–22.

Gabaix, X., 2017. Behavioral inattention. NBER Working Paper No. 24096.

Getmansky, M., Lo, A.W., Makarov, I., 2004. An econometric model of serial correlation and illiquidity in hedge fund returns. *Journal of Financial Economics*, 74, 529–609.

Goyenko, R.Y., Holden, C.W., Trzcinka, C.A., 2009. Do liquidity measures measure liquidity?. *Journal of Financial Economics* 92, 153–181.

Guo, S., Seth, A.K., Kendrick, K.M., Zhou, C., Feng, J.F., 2008. Partial Granger causality–eliminating exogenous inputs and latent variables. *Journal of Neuroscience Methods* 172, 79–93.

Hasler, M., Ornthanalai, C., 2018. Fluctuating attention and financial contagion. *Journal of Monetary Economics* 99, 106–123.

Heemeijer, P., Hommes, C.H., Sonnemans, J., Tuinstra, J., 2009. Price stability and volatility in markets with positive and negative expectations feedback: An experimental investigation. *Journal of Economics Dynamics and Control* 33, 1052–1072.

Hollo, D., Kremer, M., Lo Duca, M., 2012. CISS – A Composite Indicator of Systemic Stress in the Financial System. European Central Bank Working Paper Series No. 1426.

Hommes, C.H., Sonnemans, J.H., Tuinstra, J., van de Velden, H., 2005. Coordination of expectations in asset pricing experiments. *Review of Financial Studies* 18, 955–980.

Hong, H., Stein, J.C., Yu, J., 2007. Simple forecasts and paradigm shifts. *Journal of Finance* 62, 1207–1242.

Huang, L., Liu, H., 2007. Rational inattention and portfolio selection. *Journal of Finance* 62, 1999–2040.

Kacperczyk, M., Van Nieuwerburgh, S., Veldkamp, L., 2016. A rational theory of mutual funds' attention allocation. *Econometrica* 84, 571–626.

Kessler, S., Scherer, B., 2011. Hedge fund return sensitivity to global liquidity. *Journal of Financial Markets* 14, 301–322.

Kilian, L., 1998. Small-sample confidence intervals for impulse response functions. *Review of economics and statistics* 80, 218–230.

Kim, C.J., 1994. Dynamic linear models with Markov-Switching. *Journal of Econometrics* 60, 1–22.

King, M.A., Wadhwani, S., 1990. Transmission of volatility between stock markets. *Review of Financial Studies* 3, 5–33.

Koop, G., Pesaran, M., Potter, S., 1996. Impulse Response analysis in nonlinear Multivariate Models. *Journal of Econometrics* 74, 119–148.

Lu, Y., Ray, S., Teo, M., 2016. Limited attention, marital events and hedge funds. *Journal of Financial Economics* 122, 607–624.

Pastor, L., Stambaugh, R.F., 2009. Predictive systems: Living with imperfect predictors. *Journal of Finance* 64, 1583–1628.

Pastor, L., Veronesi, P., 2012. Uncertainty about government policy and stock prices. *Journal of Finance* 67, 1219–1264.

Patton, A., Ramadorai, T., 2013. On the high-frequency dynamics of hedge fund risk exposures. *Journal of Finance* 63, 597–635.

Peng, L., Xiong, W., 2006. Investor attention, overconfidence and category learning. *Journal of Financial Economics* 80, 563–602.

Pesaran, M.H., Shin, Y., 1998. Generalized Impulse Response analysis in linear multivariate models. *Economics Letters* 58, 17–29.

Pesaran, M.H., Timmermann, A., 2007. Selection of estimation window in the presence of breaks. *Journal of Econometrics* 137, 134–161.

Philippas, D., Dragomirescu-Gaina, C., 2016. Exposing volatility spillovers: a comparative analysis based on vector autoregressive models. *Finance Research Letters* 18, 302–305.

Racicot, F.E., Theoret, R., 2016. Macroeconomic shocks, forward-looking dynamics, and the behavior of hedge funds. *Journal of Banking and Finance* 62, 41–61.

Ratto, M., Pagano, A., 2010. Using recursive algorithms for the efficient identification of smoothing spline ANOVA models. *AStA Advances in Statistical Analysis* 94, 367–388.

Ratto, M., Pagano, A., Young, P., 2007. State dependent parameter metamodelling and sensitivity analysis. *Computer Physics Communication* 177, 863–876.

Savona, R., 2014a. Risk and beta anatomy in the hedge fund industry. *European Journal of Finance* 20, 1–32.

Savona, R., 2014b. Hedge fund systemic risk signals. *European Journal of Operational Research* 236, 282–291.

Spiliopoulos, L., Ortmann, A., 2018. The BCD of response time analysis in experimental economics. *Experimental economics* 21, 383–433.

Steiner, J., Stewart, C. Matjka, F., 2017. Rational inattention dynamics: Inertia and delay in decision-making. *Econometrica* 85, 521–553.

Woodford, M., 2014. Stochastic choice: An optimizing neuroeconomic model. *American Economic Revie*w 104, 495–500.

Yuan, Y., 2015. Market-wide attention, trading, and stock returns. *Journal of Financial Economics* 116, 548–564.

Appendix A

A Non-parametric filter: A smoothing spline ANOVA

A parametric model imposes a formal structure on the underlying data-generating process, which can be described fully by a finite number of parameters, allowing easy computation of summary statistics. However, given the well-known non-standard characteristics of the distribution of HF returns, some of the hypotheses required for standard statistical inference (e.g. linear dependence, normality, etc.) would be hard to satisfy, but can be relaxed using non-parametric techniques, also called smoothing methods (e.g. Billio et al. (2009) use kernel smoothing methods for HFs). In this appendix, we use the non-parametric filter introduced by Ratto et al. (2007) and Ratto and Pagano (2010), which combines the Kalman filter with fixed interval smoothing. The main advantages of using this non-parametric filter come from the easiness of interpreting its results, improvements in the fitness, and flexibility of the estimation approach. We only provide an intuitive description below and refer the interested reader to the studies referenced above.[61]

Any model output, y, can be seen as a mapping on a set of inputs, $X = [x_1,\ x_2,\ \dots]$. Allowing for both first-order and second-order interactions between these inputs, and ignoring the possible time subscripts, the output can be specified as follows:

$$y = f(\mathrm{x}_1, \mathrm{x}_2, \dots) = \mathrm{g}_0 + \mathrm{g}_1(\mathrm{x}_1) + \mathrm{g}_2(\mathrm{x}_2) + \mathrm{g}_3(\mathrm{x}_1 * \mathrm{x}_2) + \cdots \qquad \text{(A*)}$$

where g_i are functions that need to be identified.

Here we consider the one-factor model of Ferson and Schadt (1996) for excess HF returns, where a time-varying *beta* is allowed to depend on some lagged predetermined variables denoted by Z, according to:

$$\left[RHF - r_f\right]_t = a + \beta(Z_{t-1}) * \left[Mkt - r_f\right]_t + e_t$$

$$\beta(Z_{t-1}) = b + B * (Z_{t-1})$$

with a, b and B coefficients, $E(e_t \mid Z_{t-1}) = 0$ and $E(e_t * \left[Mkt - r_f\right]_t \mid Z_{t-1}) = 0$.

In compact form, we have:

$$\left[RHF - r_f\right]_t = a + b * \left[Mkt - r_f\right]_t + B * (Z_{t-1}) * \left[Mkt - r_f\right]_t + e_t \qquad \text{(A**)}$$

where $\left(Mkt - r_f\right)_t$ is the market factor, and e_t is an error term.

[61] The codes for running the filter are part of the global sensitivity analysis (GSA interface) toolbox, which is integrated into the Dynare platform, developed by the Joint Research Centre of the European Commission (see http://www.dynare.org).

We apply the non-parametric filter to our case assuming that HF excess returns can be explained by a set of inputs summarized by a vector $X_t = \big[[Mkt - r_f]_t,\, Z_{t-1}\big]$. The filter can show the potential explanatory power of our inputs, allowing us to gauge the ones that make the biggest contribution to mapping the realised excess returns, $[RHF - r_f]$, for each HF style. The main (or first-order) effects are computed as the percentage of variance explained by the first-order terms, while the second-order effects reflect the percentage of variance explained by the second-order (i.e. interactions) terms. The total effects include both the first- and the second-order interaction effects (the latter being double counted by construction).

Table A1 displays the results obtained from filtering the excess HF weekly returns using the smoothing spline ANOVA model. Three main findings emerge from the table. The first is that the market factor, $[Mkt - r_f]$, has always an important first-order effect, although most of the time it is dominated by other factors' contributions, especially by the *VIX* (the *EH* style could be seen as an exception, though marginally).[62] The second finding is that some of our measures, especially those from groups B and C (e.g. *VIX*, *CISS*) are very important in explaining HF excess returns. The third finding refers to the important increase in contributions once we account for second-order terms (or interactions), illustrating the non-linearities that one needs to account for. Particularly, the second order effects related to the market factor hint at some direct influences stemming from various uncertainty/risk factors, even when using lags, onto any *beta* one might wish to estimate in a CAPM settings when using HFs returns.[63]

[62] Bali et al., (2014) find that variation in uncertainty *betas* can explain a significant share of the cross-sectional variation in HF returns; although their results are not directly comparable with ours in terms of empirical design, we make a similar argument here.

[63] Second order effects related to the market factor are much higher if we use contemporaneous uncertainty/risk factors. Since the Ferson and Schadt (1996) model specification includes the lagged Z_{t-1} term, we maintain consistency with this paper in order to show that, even in this case, any estimate of *beta* would be non-linearly depending on some uncertainty/risk factors.

Table A1. A smoothing spline ANOVA model for excess HF returns

ANOVA decomposition	Style ED		Style EH		Style M		Style RVA	
	Main effects	Total effects	Main effects	Total effects	Main effects	Total effects	Main effects	Total effects
$Mkt - r_f$	0.227	0.236	0.319	0.319	0.026	0.079	0.049	0.057
Group A								
EPU	0	0	0	0	0	0	0	0
EQU	0.005	0.005	0	0	0	0	0	0
GPR	0.000	0.007	0	0	0	0	0	0
Group B								
VIX	0.496	0.525	0.304	0.315	0.147	0.242	0.122	0.122
TYVIX	0.016	0.016	0	0	0.031	0.047	0	0.109
VP	0.005	0.021	0	0.012	0	0	0	0.082
Group C								
ConnEQ	0.015	0.015	0	0	0	0	0	0.081
ConnFX	0	0.031	0.019	0.029	0.057	0.189	0	0.028
ConnSB	0	0	0	0	0	0.007	0	0
FSI	0	0.033	0	0.012	0	0.007	0	0.074
CISS	0.072	0.089	0.122	0.122	0	0	0.123	0.123
Total	**0.836**	**0.979**	**0.765**	**0.810**	**0.261**	**0.571**	**0.294**	**0.677**

Note: The four HF investment styles are specified in the first row of the table: Event Driven, denoted by ED; Equity Hedge, denoted by EH; Macro/CTA, denoted by M; and Relative Value Arbitrage, denoted by RVA. The main (or the first-order) effects are calculated as the percentage of variance explained by the first-order terms in equation (A*); second-order effects follow the same logic for the second-order terms in equation (A*). The total effects include both first- and second-order effects, the latter being double counted by the construction of the interaction terms; the first-order effects can lie between 0 and 1, but the second-order effects can exceed 1. Note that according to specification (A**) all the uncertainty/risk factors are lagged by one week.

Appendix B

Table B1. HFRX Hedge Fund style and sub-style classifications. Summary Statistics.

HF Investment Style	Ticker	Freq.	Availability	Mean	St. dev.	Skew	Kurt
Event Driven (ED)	**HFRXED**	**D**	Apr. 2003	0.0509	0.6849	-1.9174	11.79
Activist	-	M					
Credit Arbitrage	-	M					
Distressed/Restructuring	HFRXDS	D	Apr. 2003	-0.0128	0.6115	-1.3669	8.26
Merger Arbitrage	HFRXMA	D	Apr. 2003	0.0871	0.4729	-1.4751	34.27
Special Situations	HFRXSS	D	Jan. 2009	0.0889	0.6648	-1.0777	6.47
Multi-Strategy	-	M					
Equity Hedge (EH)	**HFRXEH**	**D**	Apr. 2003	0.0090	0.8619	-1.3714	7.54
Equity Market Neutral	HFRXEMN	D	Apr. 2003	0.0040	0.4718	-1.5632	14.48
Fundamental Growth	HFRXEHG	D	Jan. 2009	0.0357	1.0575	-0.5972	5.16
Fundamental Value	HFRXEHV	D	Jan. 2009	0.0372	0.8111	-1.1604	7.42
Quantitative Directional	-	M					
Sector: Energy/Basic Materials	-	M					
Sector: Healthcare	-	M					
Sector: Technology	-	M					
Short Bias	-	M					
Multi-Strategy	-	M					
Macro (M)	**HFRXM**	**D**	Apr. 2003	0.0111	0.7980	-1.3533	10.08
Active Trading	-	M					
Commodity: Agriculture	-	M					
Commodity: Energy	-	M					
Commodity: Metals	-	M					
Commodity: Multi	-	M					
Currency: Discretionary	-	M					
Currency: Systematic	-	M					
Discretionary Thematic	-	M					
Systematic Diversified	HFRXSDV	D	Jan. 2009	-0.0329	0.8792	-0.3246	3.90
Multi-Strategy	-	M					
Relative Value (RVA)	**HFRXRVA**	**D**	Apr. 2003	0.0129	0.6688	-4.6999	52.12
Fixed Income-Asset Backed	-	M					
Fixed Income-Convertible Arbitrage	HFRXCA	D	Apr. 2003	-0.0432	1.1503	-7.066	69.93
Fixed Income-Corporate	-	M					
Fixed Income-Sovereign	-	M					
Volatility	-	M					
Yield Alternatives: Energy Infrastructure	-	M					
Yield Alternatives: Real Estate	-	M					
Multi-Strategy	HFRXRVMS	D	Jan. 2009	0.1044	0.4659	0.4156	6.42

Note: Time series frequency is reported as M for monthly and D for daily. Availability refers to the month and year of the first observation in the HFR database. The mean, standard deviation (st. dev.), skewness (skew) and kurtosis (kurt) are computed with respect to average weekly returns and over the sample used in the multivariate analysis, i.e. first week of October 2004 to last week of May 2017; where sample availability is shorter, statistics are computed starting with the first weekly observation from April 2009 until the last week of May 2017.

Table B2. Summary statistics of risk and uncertainty measures

Descriptive statistics

	EPU	*EQU*	*GPR*	*VIX*	*TYVIX*	*VP*	*CISS*	*FSI*	*Conn EQ*	*Conn FX*	*Conn SB*
Mean	4.476	3.511	4.310	2.856	1.801	-0.038	-1.836	-0.432	63.86	61.35	60.41
St.dev.	0.527	0.749	0.456	0.384	0.273	0.050	0.847	1.106	9.11	5.58	8.45
Skew	0.058	0.467	0.387	1.098	0.846	2.551	0.108	2.287	-0.171	-0.317	-0.651
Kurt	2.78	3.16	3.18	4.12	3.24	18.9	1.92	9.17	2.24	2.81	3.02
N	661	661	661	661	661	661	661	661	661	661	661

Contemporaneous correlations

	EPU	*EQU*	*GPR*	*VIX*	*TYVIX*	*VP*	*CISS*	*FSI*	*Conn EQ*	*Conn FX*	*Conn SB*
EPU	1										
EQU	0.58	1									
GPR	-0.05	-0.11	1								
VIX	0.53	0.50	-0.28	1							
TYVIX	0.48	0.40	-0.27	0.78	1						
VP	0.08	0.28	0.05	0.14	0.10	1					
CISS	0.55	0.42	-0.32	0.78	0.70	0.05	1				
FSI	0.30	0.42	-0.27	0.74	0.76	0.23	0.63	1			
ConnEQ	0.44	0.35	-0.10	0.63	0.52	0.01	0.68	0.46	1		
ConnFX	0.44	0.20	0.02	0.35	0.31	-0.05	0.38	0.04	0.63	1	
ConnSB	-0.27	0.06	-0.01	0.16	0.27	0.14	0.03	0.52	0.10	-0.30	1

Note: The table presents the mean, standard deviation (st.dev.), skewness (skew), excess kurtosis (kurt), number of weekly observations (N) and correlation coefficients for all the uncertainty/risk measures included in our sample, which runs from the first week of October 2004 until last week of May 2017. All variables are in log terms, except for VP and FSI.

Table B3. Partial Granger Causality influences

	Expected beta → Factor	*Expected beta → Factor*	*Expected beta → Factor*	*Expected beta → Factor*	*Expected beta → Factor*	*Expected beta → Factor*	*Expected beta → Factor*	*Expected beta → Factor*	*Expected beta → Factor*
Sub-style	CA	DS	EHG	EHV	EMN	MA	RVMS	SDV	SS
EQU	*Yes*	*Yes*	*Yes*	*Yes*		*Yes*	*Yes*	*Yes*	*Yes*
EPU	*Yes*	*Yes*		*Yes*		*Yes*	*Yes*	*Yes*	*Yes*
GPR	*Yes*			*Yes*				*Yes*	
VIX	*Yes*		*Yes*	*Yes*					
TYVIX	*Yes*			*Yes*		*Yes*		*Yes*	
VP	*Yes*			*Yes*	*Yes*			*Yes*	
CISS	*Yes*	*Yes*		*Yes*	*Yes*	*Yes*	*Yes*		*Yes*
FSI	*Yes*			*Yes*	*Yes*	*Yes*	*Yes*	*Yes*	*Yes*
ConnEQ	*Yes*	*Yes*		*Yes*	*Yes*		*Yes*	*Yes*	*Yes*
ConnFX	*Yes*	*Yes*		*Yes*	*Yes*		*Yes*		*Yes*
ConnSB	*Yes*	*Yes*		*Yes*	*Yes*		*Yes*	*Yes*	*Yes*

Factor	*Factor → Expected beta*	*Factor → Expected beta*	*Factor → Expected beta*	*Factor → Expected beta*	*Factor → Expected beta*	*Factor → Expected beta*	*Factor → Expected beta*	*Factor → Expected beta*	*Factor → Expected beta*
Sub-style	CA	DS	EHG	EHV	EMN	MA	RVMS	SDV	SS
EQU	*Yes*	*Yes*	*Yes*		*Yes*	*Yes*	*Yes*	*Yes*	*Yes*
EPU	*Yes*	*Yes*			*Yes*		*Yes*	*Yes*	*Yes*
GPR	*Yes*	*Yes*		*Yes*	*Yes*	*Yes*	*Yes*	*Yes*	*Yes*
VIX	*Yes*	*Yes*	*Yes*	*Yes*	*Yes*	*Yes*	*Yes*	*Yes*	*Yes*
TYVIX	*Yes*	*Yes*		*Yes*	*Yes*	*Yes*	*Yes*	*Yes*	*Yes*
VP				*Yes*					
CISS	*Yes*	*Yes*	*Yes*	*Yes*	*Yes*	*Yes*	*Yes*	*Yes*	*Yes*
FSI	*Yes*	*Yes*			*Yes*	*Yes*	*Yes*	*Yes*	*Yes*
ConnEQ	*Yes*	*Yes*		*Yes*			*Yes*	*Yes*	*Yes*
ConnFX	*Yes*	*Yes*	*Yes*		*Yes*		*Yes*		
ConnSB	*Yes*	*Yes*	*Yes*		*Yes*		*Yes*	*Yes*	*Yes*

Note: Cells display the label “Yes” when we cannot reject the existence of partial Granger causality using 5% as a confidence level; a grey cell is displayed otherwise. The arrow in the first row indicates the direction of causality influences.

Table B4. Bootstrapped GIRFs to an unexpected uncertainty/risk shock

Significant GIRFs (horizon)	**Main Style ED**					
	Sub-style DS		**Sub-style MA**		**Sub-style SS**	
	Expected *beta*	Realised *beta*	Expected *beta*	Realised *beta*	Expected *beta*	Realised *beta*
Group A: Uncertainty measures						
EPU	1-8 (-)	n.s.	2-7 (+)	5-6 (-)	n.s.	n.s.
EQU	2-8 (-)	7-40 (-)	1-10 (+)	n.s.	2-6 (+)	n.s.
GPR	n.s.	n.s.	n.s.	n.s.	n.s.	7-50 (+)
Group B: Market-based risk indicators						
VIX	3-50 (-)	n.s.	2-13 (+)	n.s.	2-13 (+)	n.s.
TYVIX	1-50 (-)	n.s.	1-11 (+)	n.s.	n.s.	n.s.
VP	1-4 (-)	n.s.	2-9 (+)	n.s.	1-4 (+)	n.s.
Group C: Computed measures of (systemic) risk						
ConnEQ	2-6 (-)	n.s.	1-7 (+)	10-11(+)	2-7 (+)	2-18 (-)
ConnFX	3-10 (-)	n.s.	n.s.	11-45 (-)	n.s.	1-19 (-)
ConnSB	n.s.	n.s.	n.s.	15-30 (+)	n.s.	n.s.
FSI	1-11 (-)	1-5 (+)	2-18 (+)	1-4 (+)	1-16 (+)	n.s.
CISS	1-50 (-)	n.s.	1-13 (+)	n.s.	1-4 (+)	n.s.

Significant GIRFs (horizon)	**Main Style EH**					
	Sub-style EMN		**Sub-style EHG**		**Sub-style EHV**	
	Expected *beta*	Realised *beta*	Expected *beta*	Realised *beta*	Expected *beta*	Realised *beta*
Group A: Uncertainty measures						
EPU	5-13 (-)	n.s.	n.s.	n.s.	1-5 (+)	n.s.
EQU	2-21 (-)	n.s.	n.s.	n.s.	1-6 (+)	n.s.
GPR	n.s.	n.s.	3-12 (-)	n.s.	n.s.	5-50 (+)
Group B: Market-based risk indicators						
VIX	2-31 (-)	5-50 (-)	2-21 (+)	n.s.	3-11 (+)	n.s.
TYVIX	2-33 (-)	n.s.	1-16 (+)	n.s.	1-5 (+)	n.s.
VP	6-19 (-)	8-34 (-)	n.s.	1-2 (-)	1-4 (+)	n.s.
Group C: Computed measures of (systemic) risk						
ConnEQ	2-50 (-)	n.s.	1-6 (+)	7-33 (-)	2-8 (+)	3-4 (-)
ConnFX	n.s.	13-50(-)	n.s.	4-22 (-)	7-30 (-)	3-11 (-)
ConnSB	23-46 (-)	n.s.	n.s.	2-4 (-)	n.s.	n.s.
FSI	8-23 (-)	9-50 (-)	1-10 (+)	1-9 (-)	1-11 (+)	n.s.
CISS	1-34 (-)	18-50 (-)	1-13 (+)	n.s.	1-6 (+)	21-50 (-)

Significant GIRFs (horizon)	Main Style M	Main Style RVA						
	Sub-style SDV		Sub-style CA		Sub-style RVMS		Global HF index	
	Expected *beta*	Realised *beta*	Expected *beta*	Realised *beta*	Expected *beta*	Realised *beta*	Expected *beta*	Realised *beta*
Group A: Uncertainty measures								
EPU	n.s.	n.s.	n.s.	n.s.	n.s.	3-22 (-)	1-2 (+)	6-48 (-)
EQU	n.s.	6-15 (-)	n.s.	n.s.	n.s.	1-8 (-)	1-8 (+)	5-37 (-)
GPR	n.s.	2-7 (+)	n.s.	n.s.	n.s.	5-48 (+)	n.s.	5-28 (+)
Group B: Market-based risk indicators								
VIX	5-33 (-)	6-23 (-)	3-50(+)	n.s.	3-27 (+)	2-11 (-)	2-20 (+)	8-50 (-)
TYVIX	4-22 (-)	n.s.	1-50 (+)	n.s.	2-8 (+)	n.s.	1-20 (+)	8-50 (-)
VP	n.s.	3-22 (-)	1-30 (+)	7-13 (-)	n.s.	n.s.	1-4 (+)	n.s.
Group C: Computed measures of (systemic) risk								
ConnEQ	5-35 (-)	n.s.	n.s.	3-4 (+)	2-7 (+)	n.s.	2-11 (+)	20-45 (-)
ConnFX	n.s.	1-4 (+)	n.s.	n.s.	n.s.	1-7 (-)	2-5(+)	4-50 (-)
ConnSB	7-22 (-)	n.s.	n.s.	n.s.	1-3 (+)	n.s.	17-26 (+)	n.s.
FSI	10-37 (-)	n.s.	2-50 (+)	9-50 (-)	1-44 (+)	n.s.	1-14 (+)	1-4 (+); 20-50 (-)
CISS	4-45 (-)	n.s.	1-30 (+)	n.s.	2-8 (+)	2-7 (-)	1-6 (+)	12-50 (-)

Note: Numbers displayed in the table denote the horizon intervals for which the bootstrapped GIRFs are statistically significant at +/- 2 standard deviations. The (-) or (+) denotes the sign or direction of the GIRFs in the specified interval. The label n.s. in the table means that, given the bootstrapped confidence bands, GIRFs are not significant for (at least) two consecutive observations. The maximum horizon is truncated at 50 weeks (approximately one year).

Figure B1. *Expected* and *realised betas* for HF sub-styles and for global HF style index GL

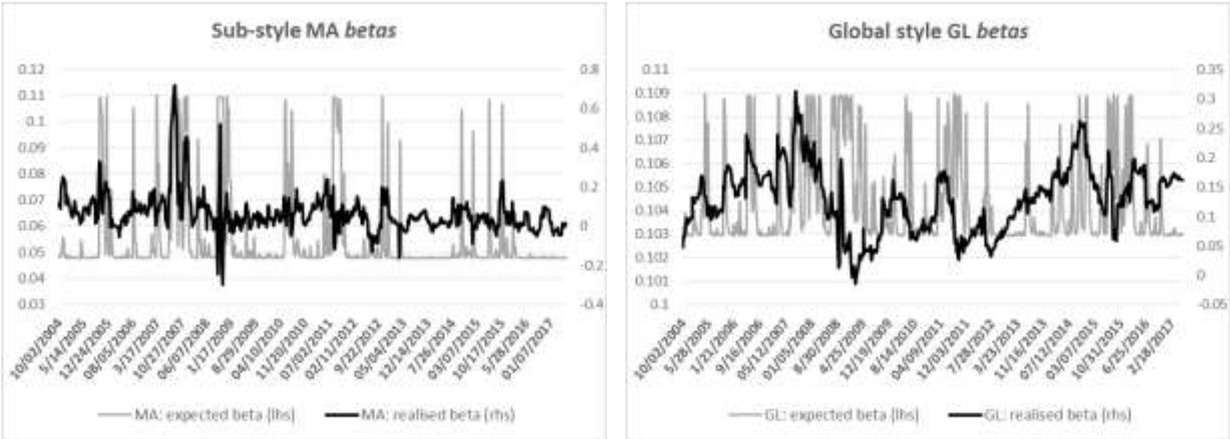

Note: We discard the first three months of data as a burn-in period, given the well-known erratic dynamics of the filtered states in a Kalman filter during the initial period. Accordingly, the effective estimation sample for the VARs starts with the first week of October 2004, and ends with last week of May 2017; for EHG, EHV, SDV, RVMS, and SS the effective estimation sample starts from the first week of April 2009 and ends with the last week of May 2017.

Figure B2. Bootstrapped GIRFs to an unexpected positive uncertainty/risk shock

Note: The figure presents the bootstrapped GIRFs to a one standard deviation positive shock in different uncertainty/risk measures, for each of the four main HF styles, in panels 1 to 4. The central line represents the median estimate of the bootstrapped impulses, while the dotted lines around the median denote the confidence bands (set at +/- 2 standard deviations); 5000 bootstrapped replications of the estimated VAR are used. The maximum horizon is truncated at 50 weeks (approximately one year).

<u>Panel 1: HF style ED</u>

Group A: Uncertainty measures

Expected *beta* response to:	Realised *beta* response to:
1 std. shock in **EPU**	1 std. shock in **EPU**
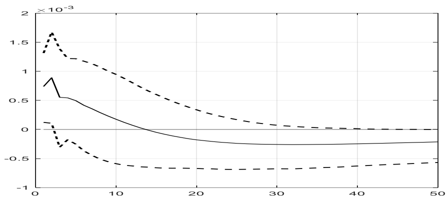	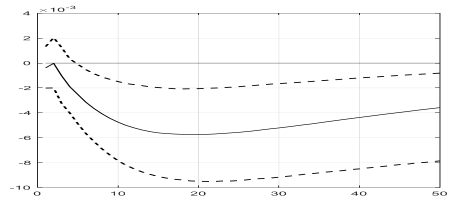
1 std. shock in **EQU**	1 std. shock in **EQU**
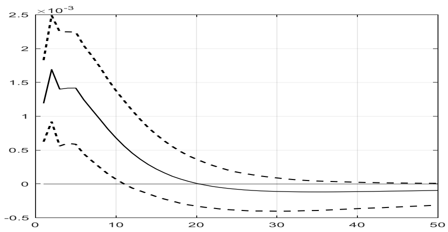	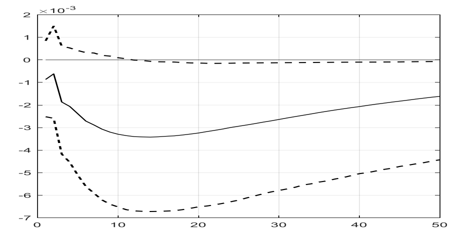
1 std. shock in **GPR**	1 std. shock in **GPR**
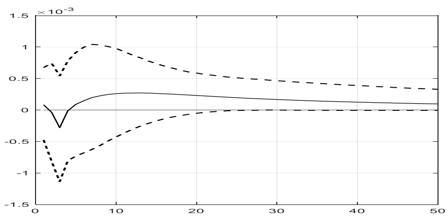	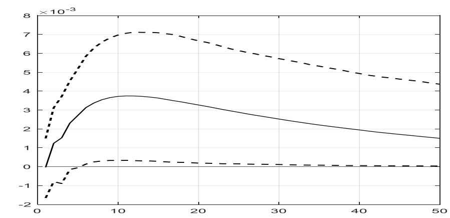

Group B: Market-based risk indicators

Expected *beta* response to:	Realised *beta* response to:
1 std. shock in **VIX**	1 std. shock in **VIX**
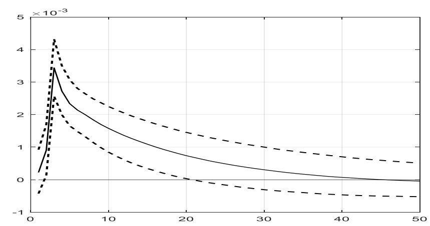	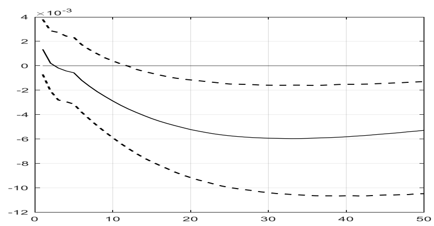

1 std. shock in **TYVIX**	1 std. shock in **TYVIX**
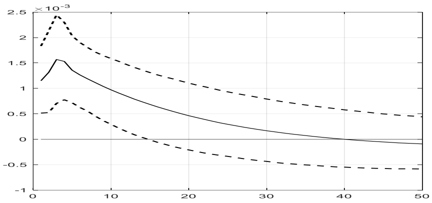	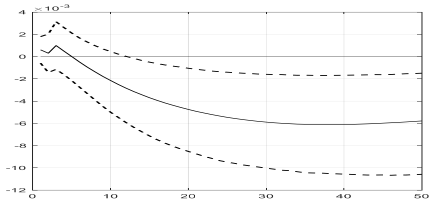
1 std. shock in **VP**	1 std. shock in **VP**
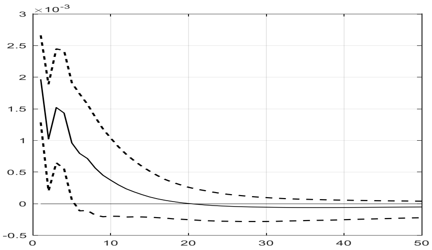	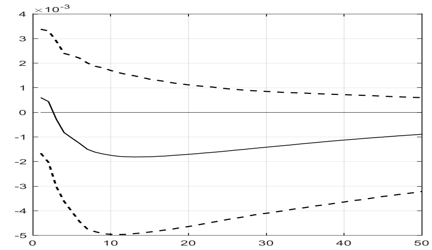

Group C: Computed measures of (systemic) risk

Expected *beta* response to:	Realised *beta* response to:
1 std. shock in **ConnEQ**	1 std. shock in **ConnEQ**
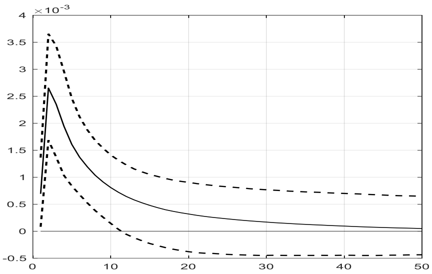	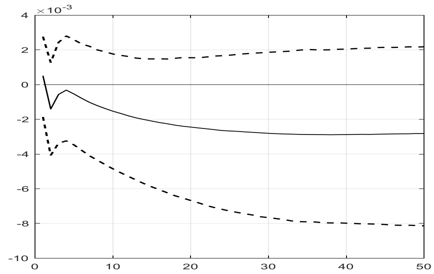
1 std. shock in **ConnFX**	1 std. shock in **ConnFX**
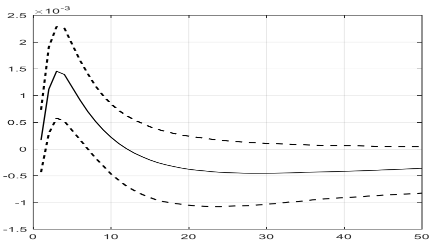	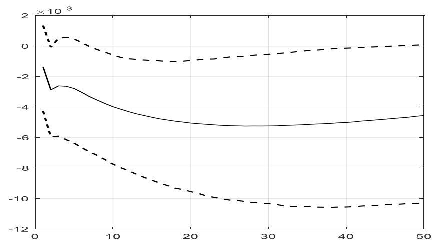
1 std. shock in **ConnSB**	1 std. shock in **ConnSB**
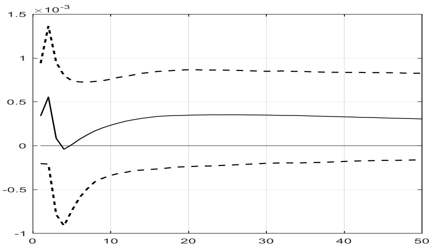	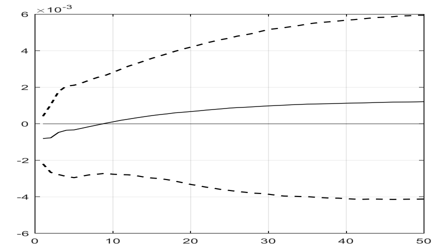

1 std. shock in **FSI**	1 std. shock in **FSI**
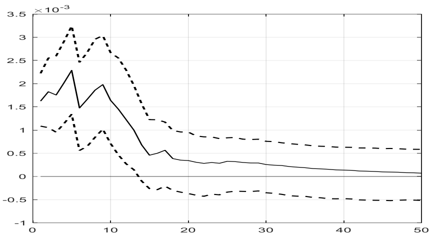	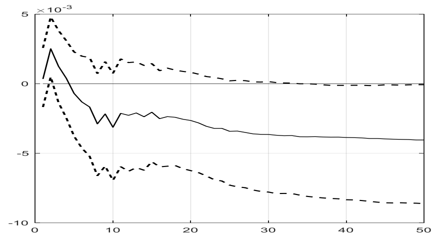
1 std. shock in **CISS**	1 std. shock in **CISS**
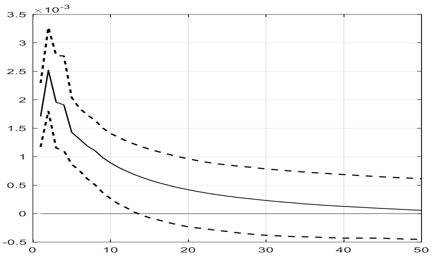	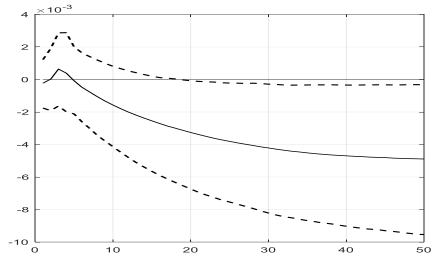

Panel 2: HF style EH

Group A: Uncertainty measures

Expected *beta* response to:	Realised *beta* response to:
1 std. shock in **EPU**	1 std. shock in **EPU**
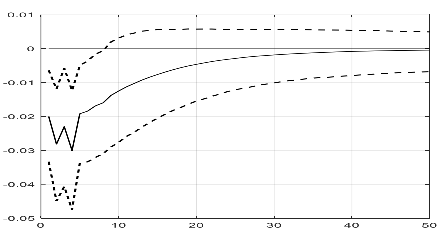	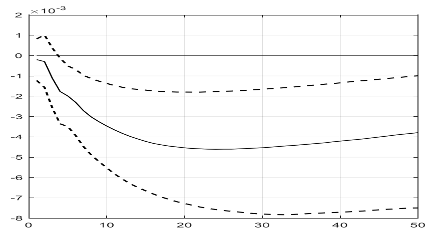
1 std. shock in **EQU**	1 std. shock in **EQU**
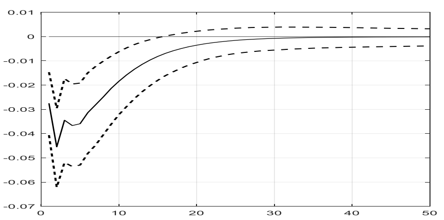	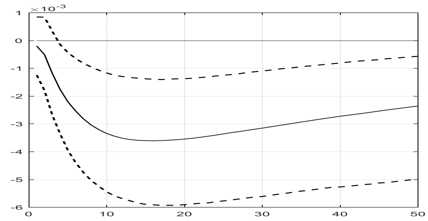
1 std. shock in **GPR**	1 std. shock in **GPR**
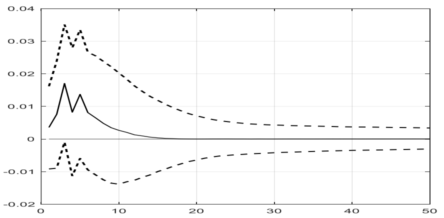	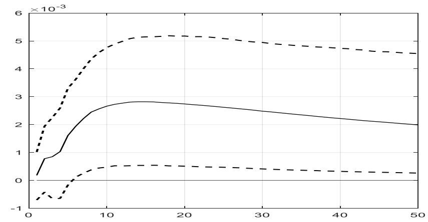

Group B: Market-based risk indicators

Expected *beta* response to:	Realised *beta* response to:
1 std. shock in **VIX**	1 std. shock in **VIX**
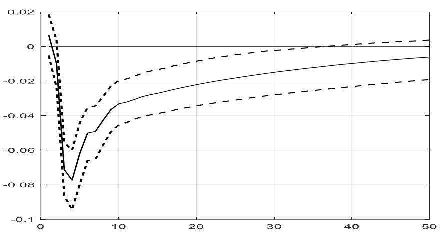	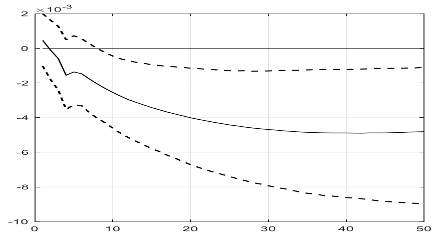
1 std. shock in **TYVIX**	1 std. shock in **TYVIX**
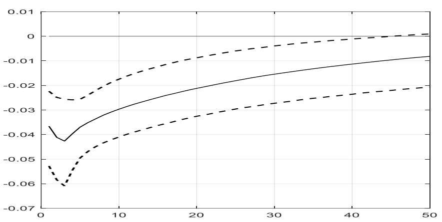	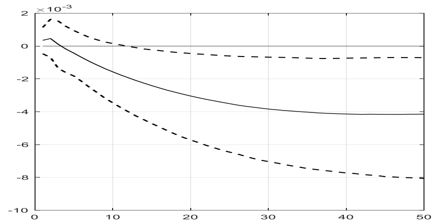
1 std. shock in **VP**	1 std. shock in **VP**
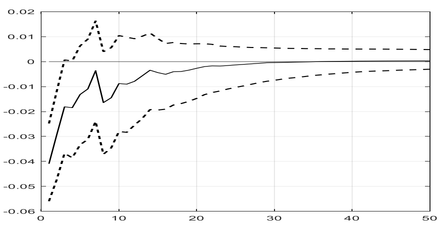	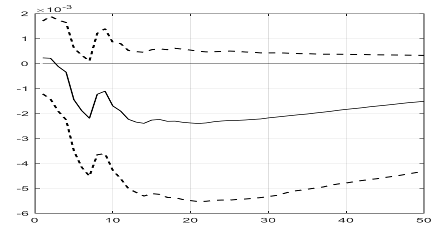

Group C: Computed measures of (systemic) risk

Expected *beta* response to:	Realised *beta* response to:
1 std. shock in **ConnEQ**	1 std. shock in **ConnEQ**
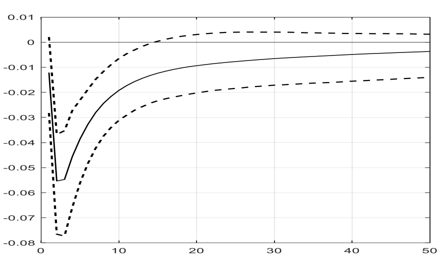	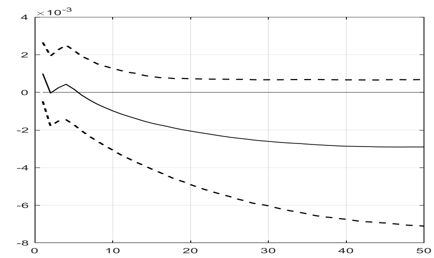

1 std. shock in **ConnFX**	1 std. shock in **ConnFX**
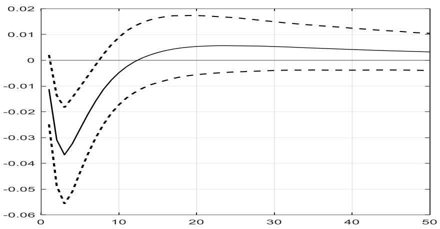	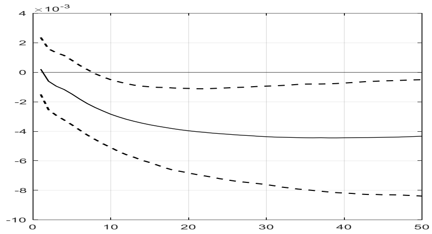
1 std. shock in **ConnSB**	1 std. shock in **ConnSB**
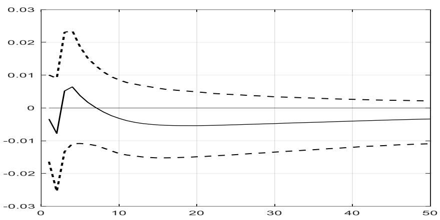	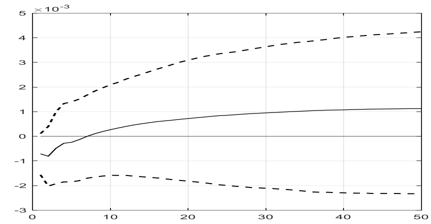
1 std. shock in **FSI**	1 std. shock in **FSI**
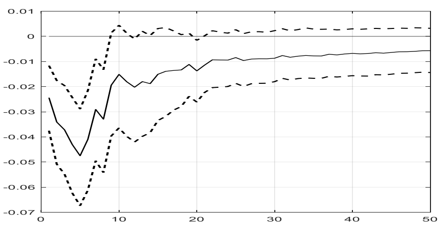	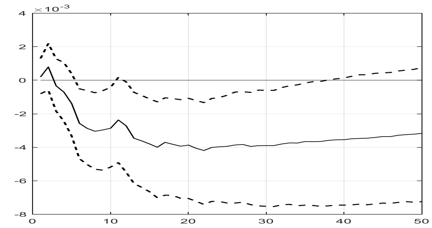
1 std. shock in **CISS**	1 std. shock in **CISS**
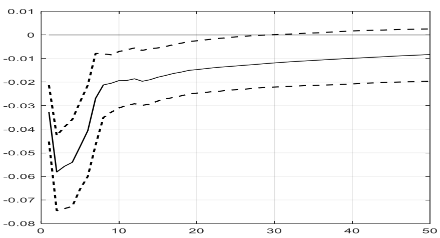	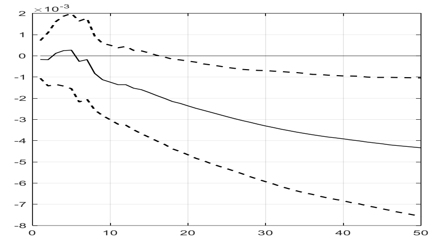

<u>Panel 3: HF style M</u>

Group A: Uncertainty measures

Expected *beta* response to:	Realised *beta* response to:
1 std. shock in **EPU**	1 std. shock in **EPU**
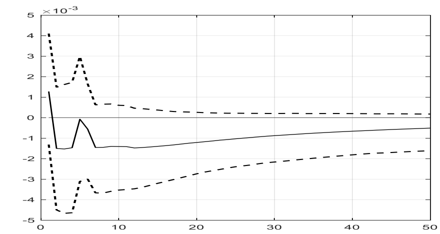	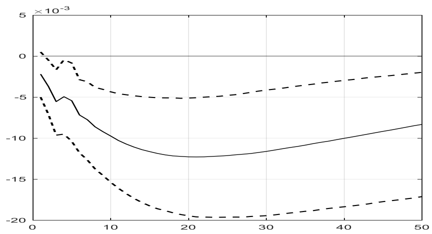

1 std. shock in **EQU**	1 std. shock in **EQU**

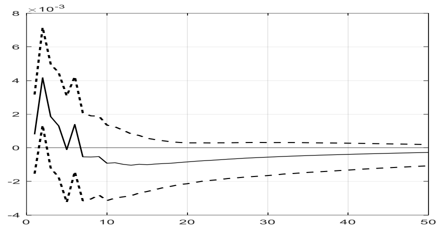

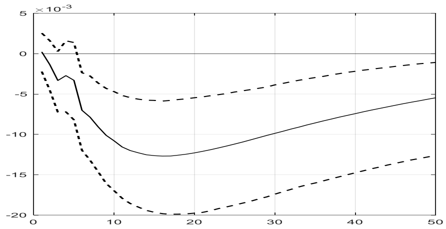

1 std. shock in **GPR**	1 std. shock in **GPR**

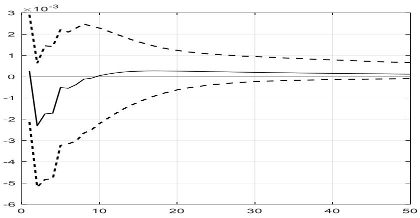

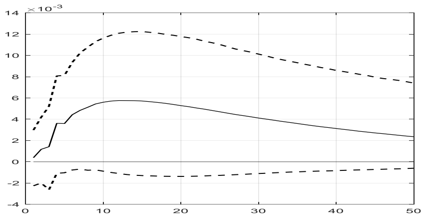

Group B: Market-based risk indicators

Expected *beta* response to:	Realised *beta* response to:
1 std. shock in **VIX**	1 std. shock in **VIX**

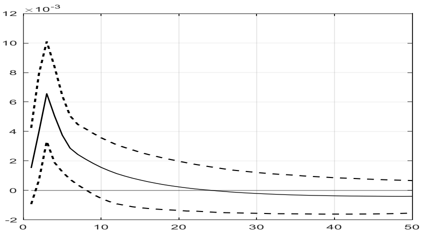

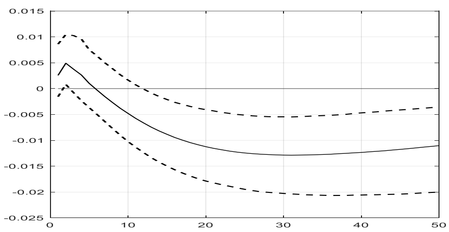

1 std. shock in **TYVIX**	1 std. shock in **TYVIX**

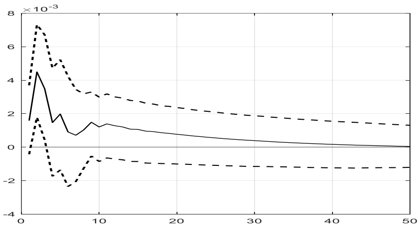

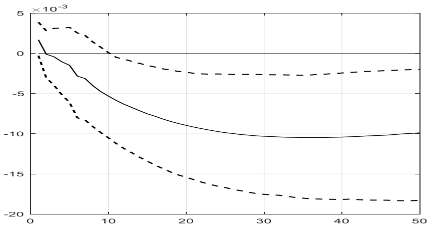

1 std. shock in **VP**	1 std. shock in **VP**

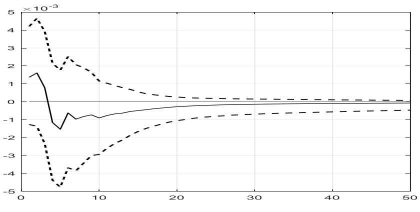

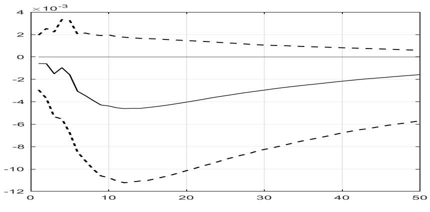

Group C: Computed measures of (systemic) risk

Expected *beta* response to:	Realised *beta* response to:
1 std. shock in **ConnEQ**	1 std. shock in **ConnEQ**
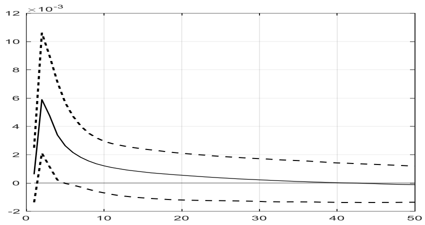	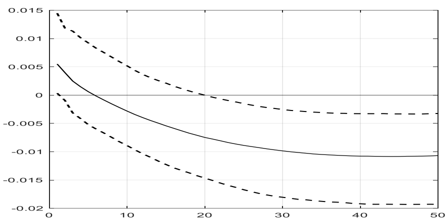
1 std. shock in **ConnFX**	1 std. shock in **ConnFX**
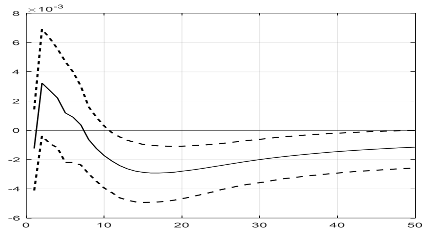	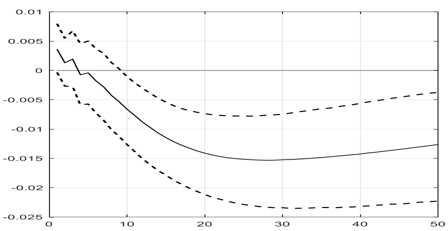
1 std. shock in **ConnSB**	1 std. shock in **ConnSB**
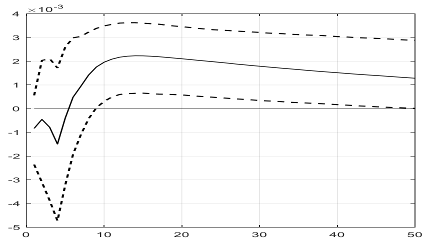	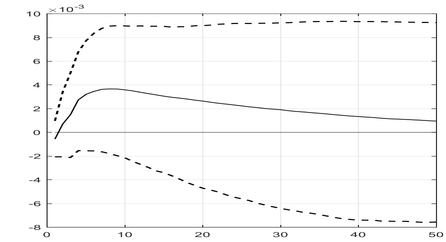
1 std. shock in **FSI**	1 std. shock in **FSI**
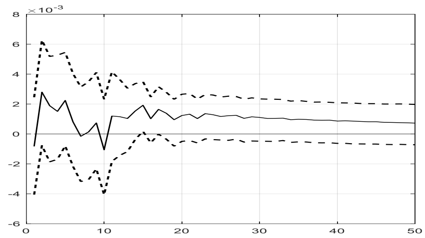	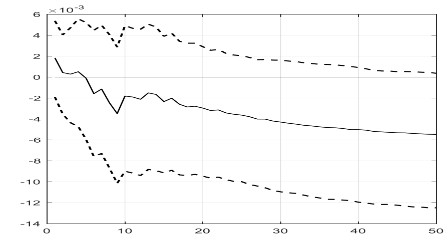
1 std. shock in **CISS**	1 std. shock in **CISS**
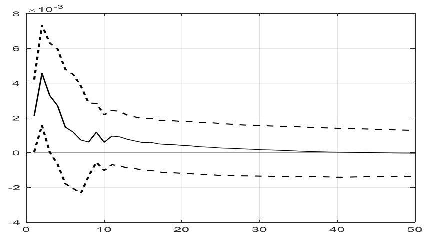	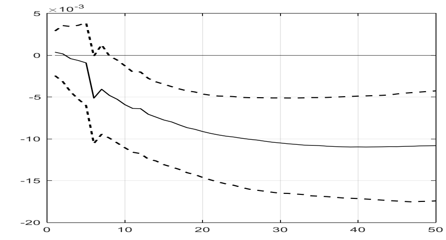

Panel 4: HF style RVA

Group A: Uncertainty measures

Expected *beta* response to:	Realised *beta* response to:
1 std. shock in **EPU**	1 std. shock in **EPU**
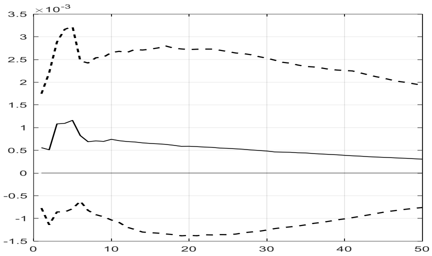	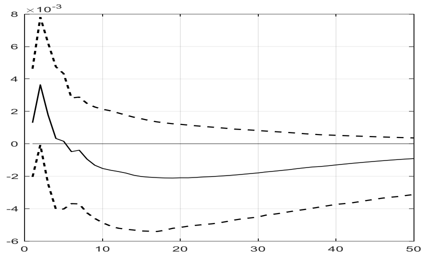
1 std. shock in **EQU**	1 std. shock in **EQU**
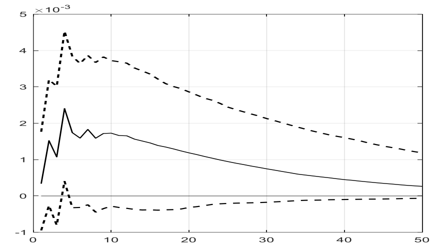	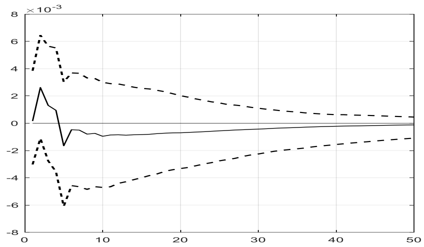
1 std. shock in **GPR**	1 std. shock in **GPR**
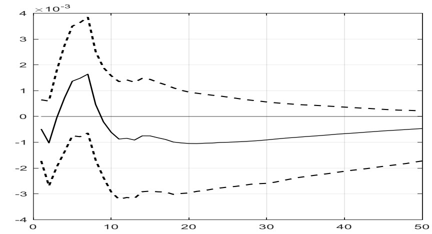	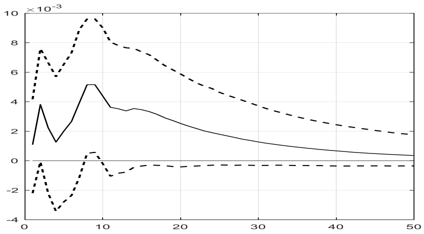

Group B: Market-based risk indicators

Expected *beta* response to:	Realised *beta* response to:
1 std. shock in **VIX**	1 std. shock in **VIX**
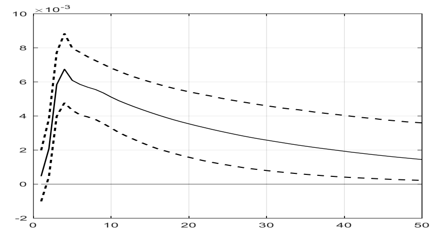	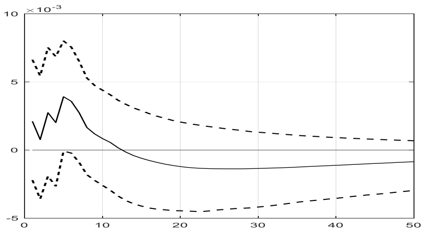

1 std. shock in **TYVIX**	1 std. shock in **TYVIX**
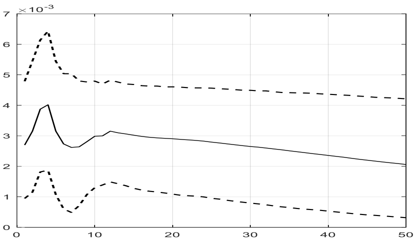	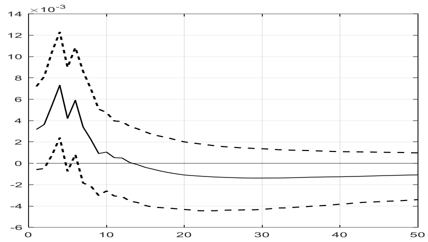
1 std. shock in **VP**	1 std. shock in **VP**
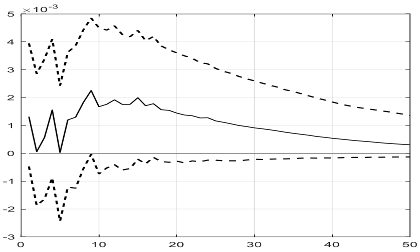	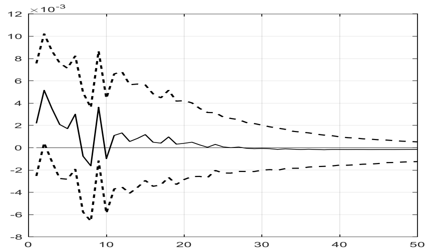

Group C: Computed measures of (systemic) risk

Expected *beta* response to:	Realised *beta* response to:
1 std. shock in **ConnEQ**	1 std. shock in **ConnEQ**
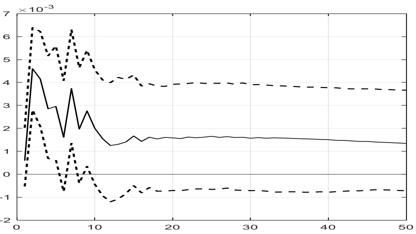	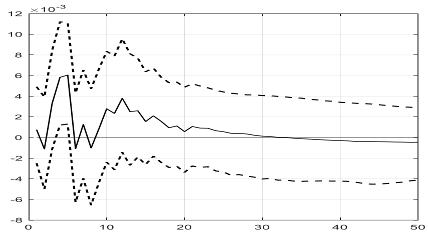
1 std. shock in **ConnFX**	1 std. shock in **ConnFX**
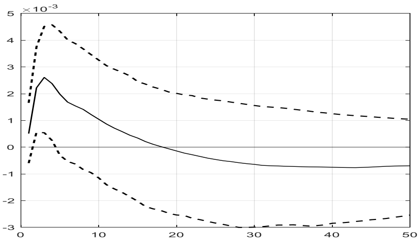	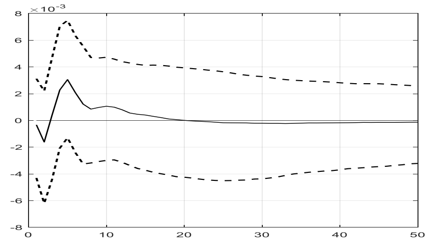
1 std. shock in **ConnSB**	1 std. shock in **ConnSB**
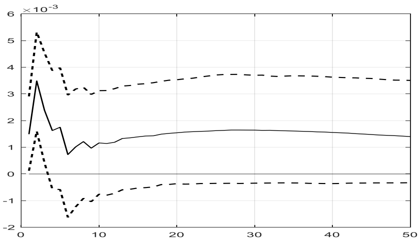	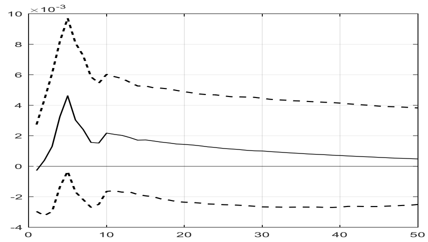

1 std. shock in **FSI**	1 std. shock in **FSI**
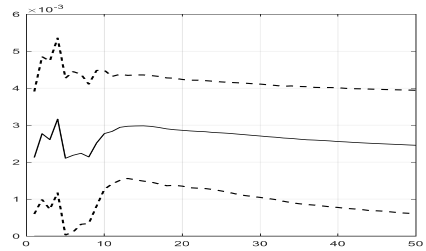	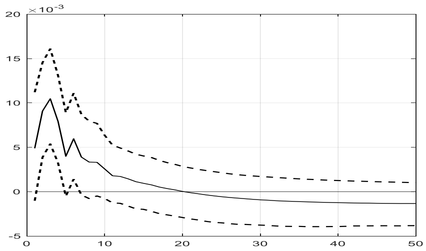
1 std. shock in **CISS**	1 std. shock in **CISS**
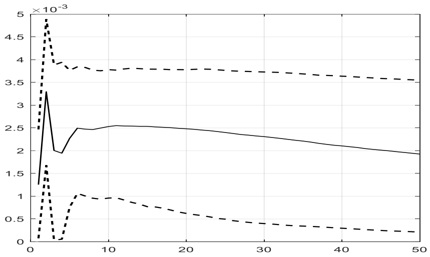	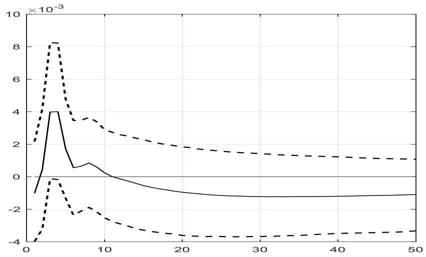

Appendix C

This Appendix provides more technical details about the simulation exercise performed in section 5. We first present the construction of two benchmark strategies: S1 – the "perfect accuracy" strategy, and S2 – the "no accuracy" strategy. *Accuracy* here refers to the range of possible values for the prediction error associated with a one-step-ahead forecast of *beta*.

Strategy S1 is constructed by cumulating all profits and losses according to the following equation:

$$S1_t = S1_{t-1} + \beta * \left[Mkt - r_f\right]_t + \epsilon_t, \quad \text{with } t = 1,2, \dots 661 \text{ and } S1_0 = 0$$

We calibrate $\beta = 0.1$ and $\epsilon_t \sim N(0,\ 0.5)$ to match the average estimates presented in Table 1 in the text. Since β is constant, prediction errors are zero, meaning that the manager has perfect accuracy regarding her portfolio strategy in any market context. Note that all simulated returns are implicitly *excess* returns. Using the same value of ϵ_t we also construct strategy S2 by cumulating all profits and losses according to the following equation:

$$S2_t = S2_{t-1} + \beta_{s(t)} * \left[Mkt - r_f\right]_t + \epsilon_t, \quad \text{with } t = 1,2, \dots 661 \text{ and } S2_0 = 0$$

This time, coefficient $\beta_{s(t)}$ is not constant but switches randomly between two extreme values, set at 0.05 and 0.15, following to a Markov Chain process: $s(t) = s(t-1) * P$, with symmetric transition matrix given by $P = \begin{pmatrix} 0.95 & 0.05 \\ 0.05 & 0.95 \end{pmatrix}$, which matches the average of the estimates provided in Table 1.

Given a signal extracted from a variable X_t, the two mixed-strategy portfolios P^A and P^B are generated according to the following rules:

$${P^A}_t = \begin{cases} S1_t, & if\ X_t < x \\ S2_t, & if\ X_t \geq x \end{cases} \quad \text{and} \quad {P^B}_t = \begin{cases} S2_t, & if\ X_t < x \\ S1_t, & if\ X_t \geq x \end{cases}$$

where X_t is the standardised value of some risk/uncertainty factor, and $x = 1.65$ such that we separate between calm (90%) and volatile periods (10%) according to the distribution properties of X_t.

We generate 5000 simulations of 661-long time series for both the noise term ϵ_t and the Markov Chain $s(t)$ processes. No exit is assumed even for negative portfolio values (since the initial value can be set arbitrarily high); there are no client inflows or outflows (redemptions). We report the median value of the following summary statistics: end of period portfolio value or cumulated (excess) returns, returns' standard deviation, returns' skewness, returns' kurtosis and Sharpe ratios. To facilitate comparison across simulated statistics, we report the Kolmogorov-Smirnov test statistics for the equality of two distributions in Table 5 (see the text).

Our calibration of the simulation inputs generally follows the estimates reported in Table 1 (see the text). To check the robustness of the simulation results, we perform a sensitivity analysis with respect to the main inputs. Our main findings, i.e. the stochastic dominance of P^A over P^B (as well as S1) in terms of lower risk, higher kurtosis and more negative skew, remain qualitatively similar (and statistically significant according to KS tests) in the following cases. Firstly, we vary the standard deviation of the random noise process $\epsilon_t \sim N(0,\ \sigma^{\epsilon})$ from $\sigma^{\epsilon} = 0.4$ to $\sigma^{\epsilon} = 0.7$ Secondly, we calibrate the persistence of the MC process both lower to $P = \begin{pmatrix} 0.9 & 0.1 \\ 0.1 & 0.9 \end{pmatrix}$ and higher to $P = \begin{pmatrix} 0.97 & 0.03 \\ 0.03 & 0.97 \end{pmatrix}$. Thirdly, we keep the mean at 0.1 but change the range of values allowed for $\beta_{s(t)}$ by narrowing the interval to [0.07; 0.13] or widening it to [0.025; 0.175]; wider intervals can lead to bimodal distributions for skew and/or kurtosis (driven by the MC process) making it hard to interpret and qualify the simulation results.

Our simulation faces several constraints that refer to the inherent simplifications we impose, such as a symmetric transition matrix and symmetric variation interval for $\beta_{s(t)}$, equal variance across the MC generated regimes, etc. Obviously, some of these constraints can be relaxed in other simulations, but the interpretation and robustness of the results might suffer as well due to increased complexity.

Chapter 3
On herding behaviour, 'green' energy and uncertainty‡

Abstract

The transition to a low-carbon economy poses significant challenges, entailing higher uncertainty, not just higher risks, for investors in energy markets. Given the current hype around 'green' investing, and climate change in general, investors should worry for price distortions driven by behavioural biases, which arise particularly in markets characterised by uncertainty and information frictions. We provide evidence on herding behaviour, and therefore social learning, in a context where investors can opt between investing in an old technology (i.e. oil) and a new (i.e. 'green') one. Based on herding dynamics and its responses to various shocks, our findings suggest that: (i) investment strategies into newer opportunities require a better information set than into older, more established ones; and (ii) policy uncertainty is a better indicator than financial risk proxies, like VIX, in reflecting the multidimensional nature of risks associated with 'green' investing today.

Keywords*:* herding; green energy; crude oil.

JEL codes: C24; G14; G15; Q40

‡ This chapter is co-authored with Prof. Dionisis Philippas from ESSCA School of Management, address: 55 Quai Alphonse Le Gallo, 92513, Boulogne, Paris, France; and Prof. Emilios Galariotis from Audencia Nantes School of Management, address: 8 Route de la Jonelière, 31222, Nantes, France. *Acknowledgements*: We would like to thank Emanuele Bacchiocchi, Rosario Crinò, Andrea Bastianin, Panagiotis Tziogkidis, Eduardo Rossi, and Giulio Palomba for insightful comments and suggestions that have greatly improved this chapter.

1. INTRODUCTION

An increasing number of institutional investors are divesting[64] from fossil fuel stocks and shifting billions of dollars into alternative assets belonging to a new sector, which claims promoting a 'green' or environmental-friendly and socially responsible investment agenda (Kaminker and Steward, 2012). Allured by wide media coverage, retail investors are jumping onto this bandwagon as well, for fear of missing out on an investment opportunity that seems to align better with the ongoing shifts in societal preferences. According to Morningstar, net inflows into the 'sustainable' sector during 2019 stand at USD 20.6 bn., nearly four times the USD 5.5 bn. record registered for the previous year.[65]

When prices do not efficiently aggregate information, a trending market can enable investors to gain more from trading rather than from acting on their private information signals, raising the probability of information cascades and, thus, herding behaviour (Cipriani and Guarino, 2008). Since herding is usually associated with information frictions and volatility spikes, its impact on financial markets and prices can be substantial and persistent (Park and Sabourian, 2011; Schmitt and Westerhoff, 2017). We aim at addressing this problem upfront in this chapter, by analysing social learning and investors' group behaviour with respect to the ongoing 'greening' in investment preferences.

Given the current hype growing around 'green' investing, and climate change topics in general, financial investors should fear for price distortions which arise particularly in markets characterised by information frictions. In the case of 'green' investing, information is often costly or limited as reflected in the high uncertainty surrounding the long-term economic viability prospects of many 'green' technologies (Kaminker and Steward, 2012; Andersson et al., 2016). As long as this uncertainty prevails, it might be that investing in an old technology (i.e. oil or fossil fuels) is less risky[66] than investing a new (i.e. 'green') technology (see Hall and Khan, 2003). High uncertainty delays the necessary learning process of identifying the most profitable market opportunities evaluated in risk-adjusted terms, but it might also incentivise (particularly skilled) investors to pay more attention and learn more about the newer opportunities (Kacperczyk et al., 2016).

It is also the case that the economic success of new technologies depends on regulations and government policies that greatly impact on the innovation process itself (Wustenhagen and Menichetti,

[64] As of December 2019, various institutions ranging from NGOs, Philanthropic Foundations and Educational Institutions to Corporations have publicly announced and committed to divest from at least one type of fossil fuel almost $11.94 trillion. Source: http://gofossilfree.org/commitments, accessed on February 20, 2020.

[65] Source: "Sustainable Fund Flows in 2019 Smash Previous Records", Morningstar, January 10, 2020, see https://www.morningstar.com/articles/961765/sustainable-fund-flows-in-2019-smash-previous-records.

[66] During the transition to a low-carbon economy, 'green' investing entails dealing with higher uncertainty, not just higher risks. See the discussion in Thoma and Chenet (2017) on the distinction between uncertainty and risk in relation to the financial implications of climate change.

2012; Andersson et al., 2016; Mazzucato and Semieniuk 2018). Mazzucato and Semieniuk (2018) find that public financial institutions (e.g. state banks) invest in higher risk technologies, and therefore can create a direction for change; in fact, Mazzucato and Semieniuk (2017) show that successful past policies in the innovation sector have been more about incentivising and shaping new markets, rather than addressing market failures.

With no certainty regarding long-term prospects and without a coherent global policy response (to which the recent U.S. withdrawal from the Paris climate accord is the latest proof), financial markets remain key to financing our society's responses to climate-related challenges[67] (Kaminker and Steward, 2012; Andersson et al., 2016; Baker et al., 2019). As multidimensional uncertainty interacts with investors' own behavioural biases, the social learning process becomes more complicated, possibly driving prices away from their fundamental values and increasing volatility (Avery and Zemsky, 1998).

Within this market environment characterized by costly or limited information and constant regulatory challenges, the paper aims at understanding how financial investors deal with the inherent uncertainty that surrounds their current portfolio allocations in energy assets. We focus on U.S. energy stocks, which are the most likely to be affected by climate-related challenges. We are most interested though in investors' collective or group dynamics, which is likely to display herding and other similar behavioural biases usually associated with information frictions and volatility spikes in financial markets. A 2015 survey conducted by the CFA Institute amongst professional portfolio managers places herding on top of a list including several behavioural biases.[68] Herding arises when investors choose to suppress their own private information and instead mimic the actions of others, leading to information cascades (Bikhchandani et al., 1992; Banerjee, 1992). Many theories predict that herding leads to higher volatility and deviations from stocks' fundamental values, i.e. asset booms and busts (Froot et al., 1992; Avery and Zemsky, 1998; Avramov et al., 2006; Park and Sabourian, 2011; Schmitt and Westerhoff, 2017). The main questions we address in this context are: What drives herding behaviour in U.S. energy sector and how does social learning occur? What should investors learn before abandoning an established investment strategy, to chase for newer (i.e. 'greener') investment opportunities?

To address these questions, we look for the main determinants of herding, considering uncertainty and the risk-return trade-offs relevant for investors in energy assets. Our main contribution is to provide

[67] It is encouraging to see banks increasingly becoming aware of their contribution to fighting climate risks, particularly through their financing decisions. The European Investment Bank (EIB) is reported to consider changing its mandate and adapting its lending policies in order to fight climate change (Source: "EIB begins metamorphosis into climate bank", Euractiv, September 9, 2019). The Dutch financial group ING Groep NV says it is allocating resources to estimate its overall carbon footprint, based on asset types and lending transactions (source: "Banks Are Finally Starting to Account for Climate Change Risk", Bloomberg Businessweek, September 12, 2019).

[68] Survey results are available at: https://blogs.cfainstitute.org/investor/2015/08/06/the-herding-mentality-behavioral-finance-and-investor-biases.

evidence on herding behaviour, and therefore social learning, in a market context where investors can opt between investing in an old technology (i.e. oil) versus a new (i.e. 'green') one. We find that investors in U.S. energy stocks herd more in response to shocks in oil returns, but not to shocks in oil volatility; investing in an old technology thus requires little besides information on returns. In contrast, the same investors herd less in response to 'green' volatility shocks, but seem immune to shocks in 'green' returns; opting for a newer investment opportunity, therefore, requires a better information set. Another contribution of our analysis regards portfolio allocations in the current market context. Environmental, social and governance (ESG) criteria are gaining in popularity among investors and companies alike, but despite this euphoria, 8 of the 10 biggest ESG funds in U.S. own substantial equity shares in big oil companies (e.g. ExxonMobil).[69] Our findings help explain this allocation strategy as well, by providing empirical evidence on the lack of sensitivity for crude oil to policy uncertainty that highly affects 'green' portfolio allocations.

From a methodological perspective, we first employ a time-varying coefficient version of the original empirical specification proposed in Chang et al., (2000) in order to expose herding towards the market consensus in U.S. energy stocks. Next, we set up vector autoregressive (VAR) models including a dynamic herding metric along with returns, volatility and uncertainty that proxy for the relevant information set available to investors. Last, we derive our main insights based on empirical tests and impulse responses from several estimated VAR models.

This chapter of the thesis is organized as follows. Section 2 provides an overview of the relevant literature, while section 3 describes our dataset used in the empirical analysis. Section 4 presents the empirical methodology and a discussion of the main results. Finally, section 5 concludes.

2. LITERATURE REVIEW

The current hype in 'greening' portfolio allocations looks exciting, although great risks and uncertainty dominate 'green' assets' valuations and their long-term prospects. Despite these inherent challenges, and despite some previous disappointing returns, 'green' assets are considered by portfolio managers for diversification motives (Miralles-Quiros and Miralles-Quiros, 2019), but also as a way to attract client inflows in the current social context. Not the same can be said about oil portfolio allocations. Andersson et al., (2016), Batten et al., (2018) and Engle et al., (2019) propose risk management techniques to hedge climate-related risks requiring portfolio allocations to both crude oil and (global) stocks. Currently, crude oil serves as a hedge against various uncertainty sources stemming mainly from

[69] Source: "ESG Funds Enjoy Record Inflows, Still Back Big Oil and Gas", Wall Street Journal, November 11, 2019, see https://www.wsj.com/articles/top-esg-funds-are-all-still-invested-in-oil-and-gas-companies-11573468200.

the policy and/or political realms (Chkili et al., 2014; Omar et al., 2017); energy equities can serve a similar scope by providing their investors with exposure to oil fluctuations and substantial dividends as well. These differences in motivations carry a significant importance for those investors deciding their portfolio allocations, particularly during periods of market stress and uncertainty.

Given the complex nature of the associated risks and uncertainties facing investors today, it is useful to review the early study of Avery and Zemsky (1998) discussing financial investors' behavioural biases, and in particular herding. Avery and Zemsky (1998) structure their discussion around (i) value uncertainty about price signals, under which herding does not occur, (ii) event uncertainty, which makes herding possible but also supportive in the price discovery process, and (iii) composition uncertainty, under which the proportion of informed versus uninformed investors is not known, thus complicating social learning and incentivizing herding behaviour that obscures rather than support price discovery.[70]

The concept of herding in financial markets was initially presented in the studies of Bikhchandani et al. (1992) and Banerjee (1992), who define it as a tendency for imitation that leads to correlated investing (e.g. buy, sell) patterns. The early literature (surveys can be found in Devenow and Welch, 1996; Bikhchandani and Sharma, 2000; Hirshleifer and Hong Teoh, 2003) has concentrated on explaining rational herding behaviour based on: (i) payoff externalities, when an individual agent's payoff depends on the number of other agents adopting the same action; (ii) principal-agent or reputational mechanisms, when failing together is less costly than failing alone; and (iii) information cascade mechanisms, when investors ignore their own beliefs and private information in order to follow the market consensus. The more recent literature adds elements drawing from the physiology and neuroeconomics fields, emphasising individuals' emotional, psychological and/or social traits (e.g. Rubinstein, 2001; Shiller, 2002; Baddeley et al., 2012).

Relying mostly on statistical measures and constructs, a vast empirical literature has analysed the presence of herding amongst several financial actors (e.g., institutional investors, fund managers, financial analysts), as well as in various financial markets (e.g., stock and bond markets, mutual funds, foreign exchange markets). The first literature strand uses micro-data to detect herding amongst (mostly institutional) investors (e.g. Sias, 2004; Blasco et al., 2012; Cipriani and Guarino, 2014). The second literature strand, to which our approach belongs as well, investigates herding towards the market consensus using aggregate market data (e.g. Christie and Huang 1995; Chang et al, 2000).

The earliest empirical specification designed for herding detection based on aggregate market data is provided by Christie and Huang (1995); alternative model specifications and extensions are included in Chang et al. (2000); Chiang and Zheng, (2010); Economou et al., (2011); Yao et al. (2014); Galariotis et al., (2015; 2016); Demirer et al., (2015); Litimi et al., (2016) etc. Even though herding is expected to

[70] This is because the trading patterns in a market with many uninformed traders that exhibit herding behaviour can be similar to the trading patterns observed in a market with many informed investors that trade based on news related to fundamentals.

be more pronounced during down or declining markets (Chang et al., 2000; Chiang and Zheng, 2010), there is evidence of significant asymmetric herding behaviour (Philippas et al, 2013) and during up markets as well (Tan et al., 2008; BenMabrouk and Litimi, 2018). We contribute to this literature by proposing a time-varying continuous herding proxy, whose inference is more data-efficient than existing approaches (e.g. Chiang et al., 2013; Babalos et al., 2015) and which preserves the theoretical consistency of the original model (Chang et al., 2000).

Cipriani and Guarino (2014) prove experimentally that rational herding behaviour increases with uncertainty. In a companion paper, they present a theoretical model where investors' private information becomes less important as trading (and volatility) increases (see Cipriani and Guarino, 2008). Sias (2004) finds that it is more likely for investors to herd in case of small cap stocks, where there is less information available and the degree of information asymmetries is higher.

Sudden changes or regime shifts, such as those triggered by (de-)regulation reforms, can speed up learning and facilitate action convergence among market participants, without necessarily being considered as herding. In this context, Wustenhagen and Menichetti (2012) discuss how changes in government policies and regulations can affect the current trade-offs associated with investing in the energy sector. We believe our analysis contributes to this debate by disentangling among different information types that affect investors' learning and portfolio allocations across the energy sector.

Finally, our paper relates to the recent booming financial literature dealing with attention allocation. Kacperczyk et al., (2016) predict that attention and information acquisition behaviour, i.e. learning, determine investors' risky portfolio choices, and present empirical evidence for their claim. Similarly, Andrei and Hasler (2019) present a model where optimal attention increase with future returns' uncertainty, treating attention as a non-financial allocation in investors' portfolios. In this context, we interpret herding as a temporary failure to allocate attention by investors.

3. DATA

Our data set comes from different sources, which we discuss in detail in this section. Data on the 31 constituent shares of the S&P 500 Energy Index come from Thomson Reuters Eikon.[71] The full sample spans from January 2011 to December 2018.[72] All our indicators are constructed as weekly averages of daily observations. Although herding behaviour is more likely to be identified with high-frequency data (see Christoffersen and Tang, 2009), daily figures would capture too much of the trading noise.

[71] Source: https://us.spindices.com/indices/equity/sp-500-energy-sector.

[72] The chosen period includes several events originating in the USA, Europe, Middle East etc., with potential impact for the U.S. and global energy industry. Some examples refer to Brexit referendum, a series of declarations and political actions by the U.S. president Donald Trump on the Iranian nuclear deal, OPEC agreements with Russia to limit oil supply, the ongoing crisis in Venezuela, political upheaval in the Middle East related to the Syrian war, etc.

Moreover, decision-making in finance is a lengthy process, and significant changes in the strategic allocations of investors are not visible on a high frequency (e.g. daily) basis, e.g. the reallocation from conventional to 'green' energy assets. In fact, most changes in portfolio strategy and asset allocations are normally associated with phases of the economic cycle, or changes in risk appetite or in the existing correlations amongst various asset classes (see Batten et al., 2018). Our results obtained at weekly frequency should be viewed therefore as being rather conservative in terms of herding detection.

The constituents of the energy sector in the S&P 500 Index are some of the biggest oil and gas companies, not only in U.S., but globally (e.g. Chevron, Exxon Mobil, Halliburton, ConocoPhillips). Some of these energy companies are already positioning as leaders in the transition towards 'greener' energy sources (Pickl, 2019). The sector includes companies active in various segments of the oil and gas industry (e.g. upstream, midstream, downstream), with different levels of vertical integration, being thus very heterogeneous. At the end of 2018, the S&P energy sector had a market representation of 5.3% in the total S&P 500 index, although this value has varied significantly over time. At its lowest point in terms of market valuation, i.e. March 2009, the energy sector had a share as high as 14.3%. In general, the market valuation share is a function of the capitalization of the constituent stocks, which may depend on the economic cycle, geopolitics, risk appetite, etc. Empirical research shows that the energy sector's performance depends mainly on the oil price (Baffes et al., 2015; Ahmadi et al., 2016; BenMabrouk and Litimi, 2018), which is the most relevant reference as well as a leading indicator of the global economic cycle. We therefore use the WTI crude oil prices as the main proxy for conventional energy investing in the following sections.

To capture the financial performance of 'green' assets instead, we rely on various available indexes and datasets. Firstly, we use 16 'green' Exchange Traded Funds (ETFs), which are some of the most representative ETFs in this sector, and are traded on the New York Stock Exchange and Nasdaq.[73] The ETFs are an attractive option for many investors, offering diversification benefits and indirect access to (sometimes illiquid or inaccessible) international equities or exotic asset classes. ETFs generally have low fees, high transparency and liquidity, and trade similarly to stocks, i.e. throughout the day, meaning that investors can employ leverage and/or short selling in order to take advantage of market moves in real time. Some ETFs might invest exclusively in particular industries (e.g. solar), or focus on specific geographical regions (e.g. South America), but many try to hold a diversified portfolio of 'green' investments in order to counter the impact of low profit margins that characterise this sector, especially during its infancy about a decade ago.[74] Several empirical papers focusing on the energy sector, use 'green' ETFs in order to study the transition to a low-carbon economy (Andersson et al., 2016; Miralles-

[73] The data source is Bloomberg. The detailed list of 'green' ETFs used in this paper is provided in Appendix A.

[74] Despite growing volumes and capacity, and sometimes generous subsidies from governments, most companies operating in the 'green' energy sector risk declining prices and margins as the technology improves and investment costs add up, new competitors enter the market, or new regulation constraints become binding.

Quiros and Miralles-Quiros, 2019). To summarize the 16 ETFs performances with a single aggregate indicator, we employ two measures: (i) a time-series of returns derived from an equally weighted (in US dollar terms) portfolio based on all 16 ETFs, and (ii) a time-series of synthetic returns derived from the first principal component of the 16 ETFs' weekly returns.[75]

Secondly, we use four indexes that are more commonly employed in empirical studies on financial aspects of climate change (Rahdari et al., 2015; Baker et al., 2018). These indexes are available from Thomson Reuters Eikon and serve to monitor the financial performance of companies, whose investing and operating principles are sensitive to environmental and climate-related risks. The first two indexes are based on the leading U.S. stock market index, i.e. S&P500, but with different allocations in order overweight (underweight) companies that fulfil (do not fulfil) certain criteria. More specifically, we use (i) the *S&P500 ESG Index*, which over-(under-)weights companies that have high (low) ESG scores, and (ii) the *S&P500 Carbon Efficient Index*, which over-(under-)weights companies with a low (high) carbon footprint. The last two indexes are (iii) the *S&P Global Clean Energy Index*, which provides investors with exposures to 30 global companies with businesses in clean energy production and equipment; and (iv) the *S&P Global Water Index*, which tracks a portfolio of about 50 global companies that do water-related businesses in utilities, infrastructure, equipment and materials.

Lastly, we take a broader view on risk, expanding beyond standard financial risk measures (e.g. volatility), and include (Knightian) uncertainty as well. To reflect risks and uncertainty affecting the investment decisions of financial investors in energy stocks, we use the CBOE Volatility index (VIX) to proxy for financial markets risk, and the *Economic Policy Uncertainty* (EPU) index[76] developed in Baker et al. (2016) for the U.S. The former is the most widely used indicator of financial risk in the empirical literature. The latter is a composite index based on the frequency of some relevant keywords in leading U.S. newspapers and is available as a daily time-series for U.S. In terms of explaining economic dynamics, Baker et al., (2016) show that the EPU is orthogonal to market volatility or risk indexes (such as VIX), despite some overlaps and correlation between the two. To refine the content of the index, the authors have defined several categorical sub-indexes[77] that pertain to different policy domains (e.g. monetary, fiscal, trade etc.). Given our research questions, in addition to the more general EPU index, and mainly as a robustness check, we use the EPU *Regulation* index which includes several

[75] The main goal of principal component analysis (PCA) in our case is to summarize the correlations among the 16 'green' ETFs returns with a smaller set of linear combinations. Considering PCA as an aggregation tool, there is an implicit requirement that the ETFs are correlated. Our analysis shows that the correlation between the returns of any two ETFs in our list is above 0.4 and statistically significant.

[76] Data and details regarding the methodology are available from www.policyuncertainty.com.

[77] For the complete list of keywords used in the construction of each categorical sub-index, see http://www.policyuncertainty.com/categorical_terms.html.

climate-related terms within its long list of keywords, e.g. carbon tax, drilling restrictions, offshore drilling, pollution controls, environmental restrictions, clean air act, clean water act.

4. EMPIRICAL APPROACH

This section presents our empirical framework, which is divided in two main parts. Firstly, we use regressions and data filtering techniques to detect the presence of herding behaviour among the constituents of the S&P 500 Energy Index. Secondly, we estimate several VAR models specified in herding, uncertainty and risk-return proxies. Finally, we draw our main insights based on impulse response functions for herding, oil and 'green' assets, and discuss their robustness.

4.1. Herding detection

We start with herding detection under what has become the common approach in the empirical literature, following the seminal paper of Chang et al. (2000). It is standard to use the cross-sectional absolute deviation (CSAD) as a proxy for assets' return dispersion:

$$CSAD_t = \frac{\sum_{i=1}^{n} \left| r_{i,t} - r_{m,t} \right|}{n-1} \tag{1}$$

where $r_{i,t}$ is the return of asset i at time t, $r_{m,t}$ is return on the market portfolio at time t, and n is the number of all assets traded on that market.

Chang et al., (2000) demonstrates that in the presence of herding, the linear relationship that capital asset pricing models (CAPM) predict between the dispersion of individual asset returns and the absolute market return, $|r_m|$, would be violated. Herding behaviour introduces nonlinearities, as some investors may trade closer to the market consensus (i.e. low CSAD values) when faced with extreme market moves (i.e. extreme r_m values). Therefore, herding can be detected in the following model:

$$CSAD_t = \beta_0 + \beta_1 \left| r_{m,t} \right| + \beta_2 r_{m,t}^2 + u_t \tag{2}$$

where $\left| r_{m,t} \right|$ denotes the absolute value of market returns, $r_{m,t}^2$ denotes the squared market returns, β_0, β_1 and β_2 are coefficients to be estimated, and u_t is an error term. For herding detection it is sufficient to let the market portfolio (used to compute the returns series, $r_{m,t}$) simply be an equally weighted portfolio constructed from all stocks in our sample, i.e. the constituent stocks of the S&P 500 Energy Index.

The mathematical derivation of eq. (2) implies that the first derivative of cross sectional dispersion with respect to the market portfolio is simply a constant term, and that the second derivative is null. Therefore, in the absence of herding, the CAPM-based model predicts that variations in market returns (in either direction) should be linearly associated with CSAD, demonstrated by a positive and statistically significant β_1 coefficient.[78] If herding exists then investors will ignore private information and switch from their own strategies to following the market consensus, thus pulling individual asset returns towards the market returns. This obscures the linear relationship between CSAD and market returns, and it is reflected in a statistically significant and negative β_2 coefficient (see Chang et al., 2000).

We estimate eq. (2) with weekly data for the constituent stocks of the S&P 500 Energy Index (see section 3). We use OLS, which provides us with consistent estimators, despite larger estimated standard errors, to gain some preliminary insights. Table 1 shows that the β_2 coefficient is positive and statistically insignificant, and therefore offers no evidence at this point of herding behaviour within the constituent stocks of the S&P500 Energy Index.

Table 1. Standard herding model and estimates by OLS

	β_0	β_1	β_2
Coefficient	0.364^{***}	0.182^{***}	0.013
Standard error	0.013	0.036	0.017

Notes: The table presents the results from the estimation of Eq. (2) for the period Jan. 2011 to Dec. 2018, at weekly frequency. Newey-West Heteroscedasticity and Autocorrelation consistent (HAC) standard errors are reported below the estimated coefficients. The Adjusted R^2 of the equation is 0.31. Three stars (***), two stars (**) and one star (*) denote significance at the 1%, 5% and 10% level, respectively.

Herding has been repeatedly shown to be market-dependent, non-linear, asymmetric and therefore essentially time-varying. Several papers have tried to identify what specific market conditions and factors, beyond pure volatility, are more likely to be associated with herding towards the market consensus. Chang et al. (2000) introduce dummies for up and down markets; Chiang and Zheng (2010), Economou et al., (2011) add spill-overs from international markets; Galariotis et al., (2015) and Hwang and Salmon (2014) separate between the role of fundamental and non-fundamental factors in driving herding; Demirer et al. (2015) include volatility persistency; Yao et al (2014) include traded volume as an additional indicator. At this point we prefer not to impose any given model structure, e.g. by adding dummies and other explanatory variables, as explained below.

Given the conditional, non-linear and time-varying nature of herding behaviour, we relax the constant coefficient assumption implicit in Eq. (2) and employ a Kalman filter to estimate it. We let only β_2 be time-varying and keep the other two coefficients as time-invariant to reflect the linear relationship between CSAD and market returns embedded in the CAPM theory. We re-specify eq. (2) in state-space form as:

[78] See Hwang and Salmon (2014) for a model of herding detection when the CAPM does not hold.

$$observation\ equation: CSAD_t = \beta'_0 + \beta'_1 |r_{m,t}| + \beta_{2,t} r_{m,t}^2 + e_t$$

$$state\ equation: \beta_{2,t} = \beta_{2,t-1} + \varepsilon_t \qquad (3)$$

where $e_t \sim N(0, \sigma^2)$, β'_0 and β'_1 are coefficients, while $\varepsilon_t \sim N(0, \sigma_{\beta 2}{}^2)$. The filtered time-varying $\beta_{2,t}$ state will be denoted as our dynamic 'herding proxy' in the remaining of the paper, because it reflects the time-varying nature of the non-linear relation between cross-sectional dispersion and market extreme returns, $r_{m,t}^2$. Although herding is an exceptional market state, a continuous and time-varying proxy can help us gain more insights into its potential determinants, without additional assumptions that impose a given structure on the estimated equation.

Estimation of the model given in eq. (3) over the Jan. 2011 – Dec. 2018 sample, using weekly observations, leads to statistically significant β'_0 and β'_1, but very close to the values of β_0 and β_1 from Table 1, for which reasons we do not report them separately. Figure 1 displays the time-varying $\beta_{2,t}$ along with its confidence interval set at +/- 2 standard deviations.

It is important to see our 'herding proxy' as reflecting real market phenomena rather than being a statistical construct. Our herding proxy generally displays mostly negative (though statistically insignificant) values, except some short periods of time at the end of 2015 – beginning of 2016 when it becomes positive. Our estimated $\beta_{2,t}$ is significantly negative during most of 2013 and 2014, coinciding with a time period when the energy market was peaking and the oil prices were hovering above 100 USD per barrel; other short periods of negative and significant values for $\beta_{2,t}$ occurred around mid-2015 and mid-2017. Overall, this observation shows that herding in energy stocks is probably more prevalent during rising (or peak) energy stocks and oil prices. This is in line with the recent empirical evidence provided in BenMabrouk and Litimi (2018), who find that herding in energy sector is more likely during rising rather than declining oil markets. Moreover, given that the market capitalisation share of the energy sector in the S&P 500 index is negatively correlated with the business cycle, many investors use energy equities and oil for portfolio diversification and protection against recession (i.e. macroeconomic) risks. There is also ample evidence on the importance of crude oil as a hedge against various other risks, including (geo)political risks (Chkili et al. 2014; Omar et al., 2017; Antonakakis et al. 2017); in this context, energy stocks allow portfolio investors not only to gain an exposure to oil as a hedge, but also access to (sometimes) substantial dividends and stock buybacks.

Figure 1: ‘Herding proxy’ in energy stocks along with S&P 500 Energy and WTI oil prices

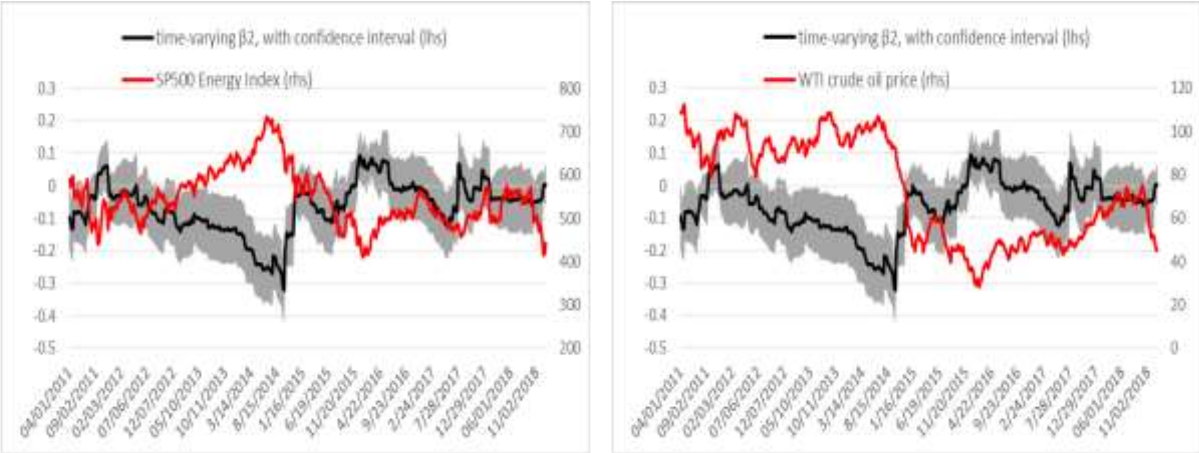

Notes: The graph displays the time-varying estimate of $\beta_{2,t}$, i.e. our ‘herding proxy’, as a solid black line along with a confidence interval set at +/-2 standard deviations depicted in grey. Sample runs from Apr. 2011 to Dec. 2018, at weekly frequency; the first 3 months of 2011 were dropped due to the known erratic behaviour of the filtered states from the Kalman filter. The left panel includes the S&P 500 Energy Index in red, while the right panel includes the crude oil prices in USD/barrel (on the right hand scale) in red.

We believe our approach to estimating a dynamic ‘herding proxy’ is more data-efficient than the existing alternatives, as it allows for more degrees of freedom, and fits some relevant real market phenomena. For example, Chiang et al., (2013), Babalos et al., (2015), and Balcilar et al., (2017) estimate herding models with time-varying coefficients using complex filtering techniques, but do not discuss whether and how their estimates match certain market dynamics or events. In additional, Chang et al. (2000) derive their testing model based on the long-run linear relationship between CSAD and market return predicted by the CAPM (i.e. constant β_0 and β_1 coefficients); therefore, estimating only $\beta_{2,t}$ as time-varying is data-efficient and balances the need for a dynamic estimate with the theoretical consistency of the original model.

4.2. Herding in a multivariate model

To capture various sources of friction that might be relevant to understanding information cascades and herding in the U.S. energy sector, we oppose oil and ‘green’ assets by looking at their returns and conditional volatility proxies. To keep things simple, we derive conditional volatility from estimating a simple GARCH (1, 1) model specified in weekly returns. For investors, the returns would convey information that is relevant from a short-term investing perspective (particularly in trending markets), while volatility proxies would carry information about the medium- to long-term portfolio implications.

While much of the empirical literature concentrates on the impact of volatility (or its proxies) on herding behaviour, there is only scarce or less clear evidence available on the inverse relationship. For example, using an intraday measure of herding intensity proposed by Patterson and Sharma (2006), Blasco et al. (2012) find that herding impacts positively on the volatility of Spanish stocks. Instead, Holmes et al. (2013) show that market volatility impacts negatively on herding, while Litimi et al. (2016) and BenSaida (2016) find evidence for U.S. that herding, particularly when trading volume is high, reduces aggregate volatility due to the presence of inactive stocks.

To address the main questions formulated in the introduction, we seek to understand herding and thus social learning when investors are facing investment choices that imply trading off an established investment strategy (into an old production technology that is oil-dependent) for a new one. To do so, we set up a simple VAR model in the following variables: one pair of information-relevant variables for each investment option (both oil and 'green' assets), the log of VIX, log of EPU, and the estimated dynamic 'herding proxy' from the previous section (albeit with a negative sign to facilitate interpretation). We have included information-rich variables like VIX and EPU indexes as a way to filter out market-related and policy-related noise, and thus reflect information from a broader context.

Compared to other estimation methods (e.g. GARCH-family models for volatility, including multivariate specifications), a simple VAR is able to expose the most important linkages and, in the same time, to remain flexible in allowing for a larger number of endogenous variables without compromising on its inference efficiency. There are 4 or 5 lags (i.e. approximately one month) included in all VAR specifications, based on selection criteria and residual autocorrelation tests. All estimated VARs are stable, with roots inside the unit circle.

We choose the equally-weighted portfolio built out of 16 ETFs, henceforth ETF_EQW, as our 'green' proxy in the remaining of this section. All the other 5 'green' proxies described in section 3 are used as robustness checks and discussed in the next sub-section.[79]

Granger causality tests are a standard check in applied econometrics and can give an overview of the existing linkages between variables included in the VAR. We report only the most relevant Granger causality results in Table 2 below. To summarize, we find that: (i) conditional volatility in the 'green' proxy Granger-cause herding in energy stocks; (ii) EPU Granger-causes 'green', but not oil conditional volatility; and (iii) VIX Granger-cases both conditional volatilities in oil and 'green' proxy. We do not go into more details here, since these results will be discussed later in this section.

[79] Appendix B to this chapter provides the summary statistics for oil and 'green' returns.

Table 2. Selected Granger causality tests

Panel A:

Dependent variable	DLOG_CRUDEOIL	
Excluded:	*Chi-square*	*Prob.*
DLOG_ETF_EQW	6.84	0.1442
LOG(VIX)	0.7257	0.9481
LOG(EPU)	1.9266	0.7492
-BETA2	8.8479	0.0650

Dependent variable	DLOG_ETF_EQW	
Excluded:	*Chi-square*	*Prob.*
DLOG_CRUDEOIL	1.3670	0.8499
LOG(VIX)	3.9260	0.4161
LOG(EPU)	5.4794	0.2415
-BETA2	4.6439	0.3258

Dependent variable	-BETA2	
Excluded:	*Chi-square*	*Prob.*
DLOG_CRUDEOIL	0.6466	0.9577
DLOG_ETF_EQW	3.9377	0.4145
LOG(VIX)	4.4310	0.3508
LOG(EPU)	1.5772	0.8129

Panel B:

Dependent variable	GARCH_CRUDEOIL	
Excluded:	*Chi-square*	*Prob.*
GARCH_ETF_EQW	11.8774	0.0365
LOG(VIX)	31.9261	0.0000
LOG(EPU)	4.1309	0.5307
-BETA2	18.5617	0.0023

Dependent variable	GARCH_ETF_EQW	
Excluded:	*Chi-square*	*Prob.*
GARCH_CRUDEOIL	10.4965	0.0623
LOG(VIX)	132.2029	0.0000
LOG(EPU)	9.3513	0.0958
-BETA2	23.9310	0.0002

Dependent variable	-BETA2	
Excluded:	*Chi-square*	*Prob.*
GARCH_CRUDEOIL	5.6640	0.3403
GARCH_ETF_EQW	14.9603	0.0105
LOG(VIX)	8.7699	0.1186
LOG(EPU)	2.6793	0.7493

Notes: For panel A, a 4-lag 5-variable VAR is used, specified in the following variables: weekly WTI crude oil returns denoted as DLOG_CRUDEOIL; weekly returns of an equally-weighted portfolio of 16 'green' ETFs, denoted as DLOG_ETF_EQW; log of VIX index; log of EPU index; and the negative of the 'herding proxy' derived in previous section denoted as -BETA2. We report the *Chi-square* statistics and the associated probability (*Prob.*) for 4 degrees of freedom. For panel B, a 5-lag 5-variable VAR is used, specified in the following variables: conditional volatility of WTI crude oil returns denoted as GARCH_CRUDEOIL; conditional volatility of an equally-weighted portfolio of 16 'green' ETFs, denoted as GARCH_ETF_EQW; log of VIX index; log of EPU index; and the negative of the 'herding proxy' derived in the previous section, denoted as -BETA2. We report the Chi-square statistics and the associated probability (*Prob.*) for 5 degrees of freedom. Estimation sample for both panels includes weekly observations from Apr. 2011 - Dec. 2018.

In terms of identification[80], we adopt an agnostic approach and employ generalised impulse response functions (or GIRFs) to derive our main insights (see Koop et al., 1996; Pesaran and Shin 1998). We believe GIRFs are more appropriate in our case, given that it is hard to impose *a-priori* restrictions on contemporaneous responses or establish a specific ordering in a model featuring fast-moving variables like volatilities, along with herding. The GIRFs, however, do not hold any structural interpretation, and only provide the average impact expected after a shock.

Figure 2: Selected bootstrapped GIRFs

Panel A

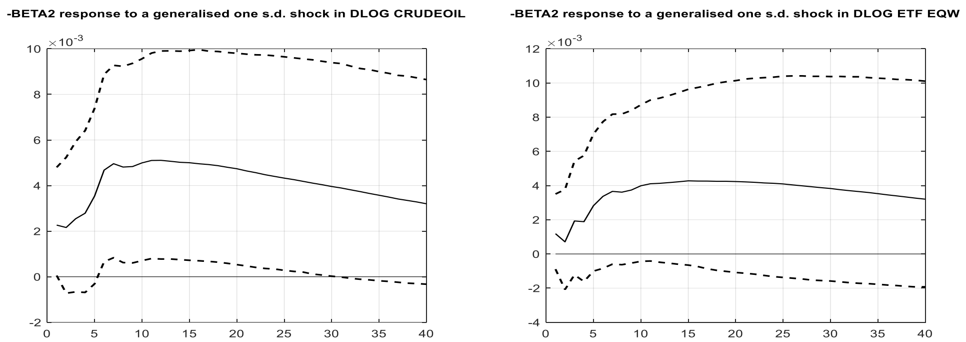

Panel B

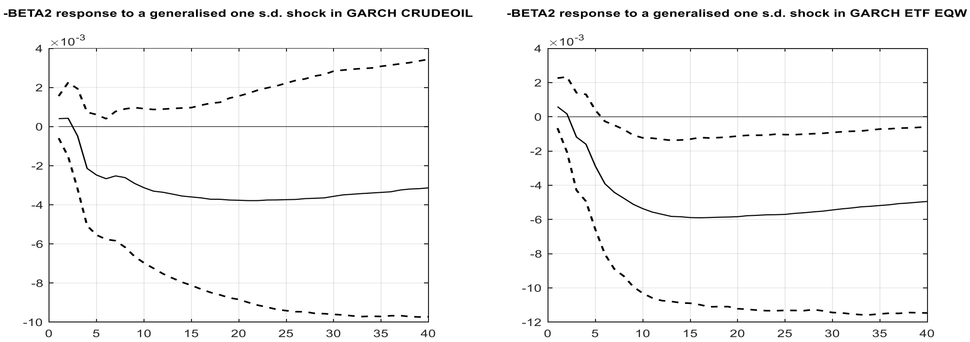

Note: The figure displays the GIRFs with bootstrapped confidence bands (set at +/- 2 standard deviations) and median responses to shocks in returns (i.e. DLOG terms) and conditional volatility (i.e. GARCH terms). For panel A, a 4-lag 5-variable VAR is used, specified in the following variables: weekly WTI crude oil returns denoted as DLOG_CRUDEOIL; weekly returns of an equally-weighted portfolio of 16 'green' ETFs, denoted as DLOG_ETF_EQW; log of VIX index; log of EPU index; and the negative of the 'herding proxy' derived in the previous section denoted as -BETA2. For panel B, a 5-lag 5-variable VAR is used, specified in the following variables: conditional volatility of WTI crude oil returns denoted as GARCH_CRUDEOIL; conditional volatility of an equally-weighted portfolio of 16 'green' ETFs, denoted as GARCH_ETF_EQW; log of VIX index; log of EPU index; and the negative of the 'herding proxy' derived in the previous section, denoted as -BETA2. Estimation sample for both panels includes weekly observations Apr. 2011 - Dec. 2018. The 'herding proxy' is taken with a negative sign in the VARs, such that an increase in herding is associated with positive values.

[80] Structural decomposition methods applied to VAR models designed to study oil and market volatility can be found among others in Kilian (2009), Bastianin and Manera (2018). Compared to these studies, our focus is not on identifying structural oil shocks, but on analysing the main possible determinants of herding behaviour.

Figure 2 provides a selection of the most relevant GIRFs from two VARs, one specified for returns and the other one specified for conditional volatility. Given that the VARs include some generated regressors (i.e. the conditional volatility, or the GARCH terms, and the herding proxy), robust confidence intervals are obtained following Kilian (1998), by bootstrapping the estimated VARs. In addition, we use a rather conservative interval of 95% for these confidence bands, or 2 standard deviations on either sides of the median impulse responses, in order to draw our main insights.

Two main results emerge from the analysis of GIRFs, confirming some of our previous insights from the Granger causality tests. Firstly, we find that our 'herding proxy' does not respond to shocks in 'green' returns, but increases in response to positive shocks in oil returns; given that oil is an established investment strategy, into an industry reliant on old production technologies, there is no need for more than information on returns. Secondly, our 'herding proxy' responds (albeit with a lag that reflects the learning process) only to shocks in 'green' (but not oil) volatility, which can be best related to the medium- and long-term prospects and the quality of the information set available for 'green' assets and technologies. Therefore, unexpected spikes in 'green' assets' volatility reduce herding intensity amongst equity investors, suggesting that some of them might be considering additional information is needed before joining the trend.

What these GIRFs from Figure 2 imply is that unexpected *negative* oil returns and/or *positive* spikes in 'green' volatility might reduce herding intensity among investors in energy equities. The results offer a perspective on what type of information is required to resolve uncertainty and ease information frictions that are conducive to herding in the first place. Kacperczyk et al., (2016) suggest that skilled investors allocate more attention and prefer to learn more about the most uncertain outcomes (see Proposition 1);[81] in other words, high volatility in 'green' assets would incentivize (at least) some investors to learn more (i.e. implying less herding) about 'green' opportunities. Andrei and Hasler (2019) arrive at a similar conclusion in a model where optimal attention represents a non-financial allocation in investors' portfolios. From this perspective, we can interpret herding as a temporary failure in allocating attention. While much of the literature surveyed in section 2 suggests that higher volatility drives more herding, we provide evidence that an opposite channel might be at work in the 'green' sector, where higher volatility/uncertainty pushes (at least some) investors towards more learning, and therefore less herding. The arrival of better informed investors can break the information cascade and reduce herding incidence (Bikhchandani and Sharma, 2000). To summarise, social learning with respect to old technologies requires nothing more than information on returns, but newer investment opportunities require a better information set in the first place.

[81] Grossman and Stiglitz (1980) made a similar prediction much earlier, i.e. when many investors are informed about some risk, market prices are being informative, therefore decreasing one's incentives to learn about the very same risk.

Appendix C at the end of this chapter includes the full set of GIRF derived from the two VARs discussed in this section; these plots provide additional evidence in line with our previous Granger causality insights where we saw that ‘green’ assets (not oil) are sensitive to policy uncertainty. We find that unexpected shocks in EPU lead to higher volatility and lower returns in the case of ‘green’ assets, but not in the case of crude oil. In the same time, unexpected VIX shocks drive higher volatility and lower returns in both oil and ‘green’ assets. Policy uncertainty is therefore a better indicator than financial risk proxies, like VIX, in reflecting the multidimensional nature of risks associated with ‘green’ investing today. The importance of this result should not be underrated: if policy uncertainty drives ‘green’ assets volatility, but not oil volatility, then investors might be able to use oil as a hedge against policy uncertainty shocks, and therefore as a hedge for their ‘green’ portfolio allocations against policy-related and regulatory changes. This surprising result might thus justify the substantial equity shares in big oil companies (e.g. ExxonMobil) reported by some of the biggest ESG funds in U.S. Moreover, it is in line with the recent results provided in Andersson et al., (2016), Batten et al., (2018) and Engle et al., (2019) regarding the role of oil allocations for financial portfolios that need to hedge climate-related risks.

4.3. Robustness

As robustness checks we implement a series of modifications to the VAR models estimated in the previous section. Firstly, we replace the equal-weighted ETF portfolio, ETF_EQW, with each of the other 5 ‘green’ proxies mentioned in section 3. Results are qualitatively similar, except for the S&P Global Clean Energy Index, where herding responses are not statistically significant based on the bootstrapped confidence intervals (set at +/-2 standard deviations around the median response).

Secondly, we use futures oil prices instead of spot prices. Avery and Zemsky (1998) claim that the presence of derivatives makes herding and price bubbles less pronounced as such instruments are better at reflecting multidimensional uncertainty. Futures prices provide a link between current prices and the expected spot prices after including all relevant uncertainty sources into the price formation process. However, in our case herding occurs in a different market segment (i.e. energy equities) than the one where we measure volatility (i.e. oil futures market), although the two would be tightly linked. The main results remain mostly unchanged, but in addition we find that herding response to an unexpected shock in futures oil price volatility is negative (just as in response to a ‘green’ volatility shock); this new result validates the interpretation provided in Avery and Zemsky (1998) that derivatives are better at reflecting multidimensional uncertainty. Figure D1 in Appendix D provides a selected set of relevant GIRFs in this case.

Thirdly, we replace EPU with one of its domain-specific sub-indexes, i.e. EPU-*Regulation*, which is expected to be more relevant for the 'green' sector given the strong feedbacks between regulation, innovation and investing (Mazzucato and Semieniuk, 2018). In addition, we might like to clean EPU from those elements that could be unrelated to climate policy and regulatory risks, but instead reflect international political crises or conflicts (e.g. U.S. fight against terrorism in Middle East), which are known to increase investors' demand for oil as a hedge (see Omar et al., 2017). Results are qualitatively unchanged (see figure D2, Appendix D).

Fourthly, we re-estimate our herding proxy from section 4.1, eq. (3), this time assuming its dynamics follows an autoregressive of order 1, i.e. AR(1), process instead of a unit root, i.e. I(1), process. While the persistency of the new estimated state variable, $\beta_{2,t}$, is high and close to 0.9 (statistically significant), and the two herding proxies are correlated (i.e. at 0.72), the confidence interval associated with the new estimates implies a lack of herding over the sample period. In addition, we believe that the new herding proxy is less intuitive and less relevant in terms of matching real market phenomena. However, the GIRFs that depict herding responses to 'green' and oil shocks are statistically significant, though slightly less persistent than in the benchmark case presented in section 4.2 (see Figure D3, in Appendix D). Acknowledging that the 'herding proxy' remains a model generated variable that might depend on specific inference techniques, none of our main findings with respect to herding and social learning are affected by the modelling choice adopted for herding persistency.

Fifthly, we include a 6-th variable in our VAR to reflect mass media coverage of climate-related topics, an indicator that complements the market information available to investors from returns or volatility proxies. Boykoff et al., (2020) collect data on media coverage from 55 countries, in several languages and from different sources, e.g. newspapers, wire services etc.[82] The authors show that coverage tends to increase around global policy events, such as the United Nations Framework Conventions on Climate Change (COPs), or during natural disasters that are most likely to raise awareness among the public and investors. Data is provided as a simple count of media news, articles and mentions; we remove the upward trend from the data (in order not to bias the results) using a HP filter and denote this variable as NEWS. The impulse responses from the extended 6-variable VAR (see Appendix D, Figure D4) confirm the main findings from the previous section, but also show a negative response of herding to a positive shock in NEWS. Despite any possible drawbacks that come with such (simple count) data, this new result proves that increased media coverage can improve social learning, and therefore have a favourable impact on investors' behaviour with respect to energy stocks – an idea in line with our previous discussion.

[82] See https://sciencepolicy.colorado.edu/icecaps/research/media_coverage/world/index.html. Data is provided as a simple count with no additional transformation, and is available on a monthly frequency since 2004, by source, and by region. We use the worldwide aggregated count as our NEWS proxy in the analysis.

5. CONCLUSIONS

Given the current hype associated with 'green' investing, and climate change in general, the transition to a low-carbon economy entails significant challenges for financial investors trading in energy stocks. Whether and how these investors, as a group, learn to deal with various information frictions and to navigate through swings in risk appetite and uncertainty is important for their investment choices and strategies. Our broader view on risk expands beyond standard financial risk measures, like volatility or the common VIX index, to include Knightian uncertainty as well, and particularly *policy uncertainty* (see Baker et al., 2016). In a market context dominated by high risks, uncertainty and information frictions, we focus on investors' herding behaviour, which is associated with situations where social learning can lead to significant price distortions. We thus seek to understand investors' behaviour when facing choices that essentially boil down to trading off a conventional investment strategy that is centred on oil, for a new one that is 'greener', but less predictable.

We first estimate a proxy for herding towards the market consensus derived from the time-varying version of the standard herding detection model of Chang et al., (2000) and next use it to understand how herding behaviour in U.S. energy stocks might interact with asset returns, volatility and uncertainty proxies. We set up a VAR model and use standard tools such as Granger causality tests and generalised impulse response functions to draw conclusions. We find that herding among investors in energy stocks responds to 'green' volatility shocks, but not to 'green' return shocks. In contrast, herding responds to shocks in oil returns, but not in oil volatility. Therefore, opting for an investment strategy into an old and established technology requires little besides information on returns, but newer investment opportunities would require a better information set. This conclusion is further supported by the result that, while both oil and 'green' assets are sensitive to VIX shocks, only 'green' assets are sensitive to EPU shocks, which encompass more uncertainty dimensions, spanning across policy-related and regulatory realms. The importance of reducing uncertainty and improving information is highlighted again in the robustness checks section, where we find that better media coverage of climate-related topics can reduce herding in energy sector.

In the case of 'green' assets, uncertainty and information frictions can prevent market prices from being fully informative about the associated risk-return trade-offs. This creates a potential for future gains because higher uncertainty incentivises (at least some) investors to allocate more attention in learning about new risks and opportunities (Kacperczyk et al., 2016; Andrei and Hasler, 2019); such an information acquisition behaviour, instead, reduces herding incidence, which can be viewed as a (temporary) failure in allocating attention. Our empirical results confirm this interpretation, contributing new evidence to the recent booming literature dealing with investors' attention allocation.

Other implications of our findings concern hedging options, which are important mostly for portfolio investors. While negative returns on conventional energy assets, such as oil, might look bad for the overall portfolio performance, we find that equity investors herd less and are more likely to refocus and search for value within the energy sector itself. In addition, since oil is also used as a hedge against various shocks and risks, negative oil returns are can be acceptable for some portfolio investors. When it comes to 'green' investing, unfortunately, and given that 'green' volatility is significantly affected by policy uncertainty, investors might have no easy hedge to rely on, except again oil, which appears less sensitive to this type of uncertainty. It is not surprising therefore to see that many 'green' funds today are invested in oil, or holding equity shares in big oil companies.

REFERENCES

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

Andersson, M., Bolton, P., Samama, F., 2016. Hedging climate risk. *Financial Analysts Journal*, 72, 13–32.

Andrei, D., Hasler, M., 2019. Dynamic attention behavior under return predictability. *Management Science*. Forthcoming. Doi.org/10.1287/mnsc.2019.3328.

Antonakakis, N., Gupta, R., Kollias, C., Papadamou, S., 2017. Geopolitical risks and the oil-stock nexus over 1899–2016. *Finance Research Letters* 23, 165–173.

Avery, C., Zemsky, P., 1998. Multidimensional uncertainty and herd behavior in financial markets. *American Economic Review*, 88, 724–748.

Avramov, D., Chordia, T., Goyal, A., 2006. Liquidity and autocorrelations in individual stock returns. *The Journal of Finance* 61, 2365–2394.

Babalos, V., Stavroyiannis, S., Gupta, R., 2015. Do commodity investors herd? Evidence from a time-varying stochastic volatility model. *Resources Policy* 46, 281–287.

Baddeley, M., Burke, C., Schultz, W., Tobler, P., 2012. Herding in Financial Behaviour: A Behavioural and Neuroeconomic Analysis of Individual Differences. *Cambridge Working Papers in Economics*. https://doi.org/10.17863/CAM.1041.

Baffes, J., Kose, M. A., Ohnsorge, F., Stocker, M., 2015. The great plunge in oil prices: Causes, consequences, and policy responses. *World Bank Group, Policy Research Note* 94725.

Balcilar, M., Demirer, R., Ulussever, T., 2017. Does speculation in the oil market drive investor herding in emerging stock markets? *Energy Economics* 65, 50–63.

Baker, S. R., Bloom, N., Davis, S. J., 2016. Measuring economic policy uncertainty. *The Quarterly Journal of Economics* 131, 1593-1636.

Bastianin, A., Manera, M., 2018. How does stock market volatility react to oil price shocks? *Macroeconomic Dynamics* 22, 666–682.

Batten, J. A., Kinateder, H., Szilagyi, P. G., Wagner, N. F., 2018. Addressing COP21 using a stock and oil market integration index. *Energy Policy* 116, 127–136.

BenSaida, A., 2017. Herding effect on idiosyncratic volatility in US industries. *Finance Research Letters* 23, 121–132.

BenMabrouk, H., Litimi, H., 2018. Cross herding between American industries and the oil market. *The North American Journal of Economics and Finance* 45, 196–205.

Bikhchandani, S., Hirshleifer, D., Welch, I., 1992. A theory of fads, fashion, custom, and cultural change as informational cascades. *Journal of Political Economy* 100, 992–1026.

Bikhchandani, S., Sharma, S., 2000. Herd behaviour in financial markets: A review. IMF Working Paper 48, *International Monetary Fund*, Wahshington, D.C.

Blasco, N., Corredor, P., Ferreruela, S., 2012. Does herding affect volatility? Implications for the Spanish stock market. *Quantitative Finance* 12, 311–327.

Boykoff, M., Aoyagi, M., Benham, A., Chandler, P., Daly, M., Doi, K., Fernández-Reyes, R., McAllister, L., McNatt, M., Nacu-Schmidt, A., Oonk, D., Osborne-Gowey, J., Pearman, O., Simonsen, A.H., and Ytterstad, A., 2020. World Newspaper Coverage of Climate Change or Global Warming, 2004-2020. Media and Climate Change Observatory Data Sets. Center for Science and Technology Policy Research, Cooperative Institute for Research in Environmental Sciences, University of Colorado. doi.org/10.25810/4c3b-b819.

Chang, E., Cheng, J., Khorana, A., 2000. An examination of herd behavior in equity markets: An international perspective. *Journal of Banking and Finance* 24, 1651–1679.

Chiang, T., Zheng, D., 2010. An empirical analysis of herd behavior in global stock markets. *Journal of Banking and Finance* 34, 1911–1921.

Chkili, W., Aloui, C., Nguyen, D. K., 2014. Instabilities in the relationships and hedging strategies between crude oil and US stock markets: do long memory and asymmetry matter? *Journal of International Financial Markets, Institutions and Money* 33, 354–366.

Christie, W., Huang, R., 1995. Following the pied piper: Do individual returns herd around the market? *Financial Analysts Journal* 4, 31–37.

Christoffersen, S. K., Tang, Y., 2009, Institutional Herding and Information Cascades: Evidence from Daily Trades. Working Paper, McGill University.

Cipriani, M., Guarino, A., 2008. Herd Behavior and Contagion in Financial Markets. *The B.E. Journal of Theoretical Economics (Contributions)* 8(1), Article 24.

Cipriani, M., Guarino, A., 2014. Estimating a Structural Model of Herd Behavior in Financial Markets. *American Economic Review* 104, 224–251.

Demirer, R., Lee, H., Lien, D., 2015. Does the stock market drive herd behavior in commodity futures markets? *International Review of Financial Analysis* 39, 32–44.

Devenow, A., Welch, I., 1996. Rational herding in financial economics. *European Economic Review* 40, 603–615.

Economou, F., Kostakis, A., Philippas, N., 2011. Cross-country effects in herding behaviour: Evidence from four south European markets. *Journal of International Financial Markets, Institutions and Money* 21, 443–460.

Engle, R. F., Giglio, S., Kelly, B. T., Lee, H., Stroebel, J., 2019. Hedging climate change news. *National Bureau of Economic Research*. WP No. w25734.

Froot, K. A., Scharfstein, D. S., Stein, J. C., 1992. Herd on the street: Informational inefficiencies in a market with short-term speculation. *The Journal of Finance* 47, 1461–1484.

Galariotis, E. C., Rong, W., Spyrou, S. I., 2015. Herding on fundamental information: A comparative study. *Journal of Banking & Finance* 50, 589–598.

Galariotis, E., Krokida, S., Spyrou, S., 2016. Herd behavior and equity market liquidity: Evidence from major markets. *International Review of Financial Analysis* 48, 140–149.

Grossman, S., Stiglitz, J., 1980. On the impossibility of informationally efficient markets. *American Economic Review* 70, 393–408.

Hall, B. H., Khan, B., 2003. Adoption of new technology. *National Bureau of Economic Research*, Working Paper No. w9730.

Hirshleifer, D., Hong Teoh, S., 2003. Herd behaviour and cascading in capital markets: A review and synthesis. *European Financial Management* 9, 25–66.

Holmes, P., Kallinterakis, V., Ferreira, M. L., 2013. Herding in a concentrated market: a question of intent. *European Financial Management* 19, 497–520.

Hwang, S., Salmon, M. 2004. Market stress and herding. *Journal of Empirical Finance* 11, 585–616.

Kacperczyk, M., Van Nieuwerburgh, S., Veldkamp, L., 2016. A rational theory of mutual funds' attention allocation. *Econometrica* 84, 571–626.

Kaminker, C., Stewart, F., 2012. The Role of Institutional Investors in Financing Clean Energy, OECD Working Papers on Finance, Insurance and Private Pensions, No.23, OECD Publishing.

Kilian, L., 1998. Small-sample confidence intervals for impulse response functions. *Review of economics and statistics* 80, 218–230.

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

Koop, G., Pesaran, M., Potter, S., 1996. Impulse Response analysis in nonlinear Multivariate Models. *Journal of Econometrics* 74, 119–148.

Litimi, H., BenSaida, A., Bouraoui, O., 2016. Herding and excessive risk in the American stock market: A sectoral analysis. *Research in International Business and Finance* 38, 6–21.

Mazzucato, M., Semieniuk, G., 2017. Public financing of innovation: new questions. *Oxford Review of Economic Policy* 33, 24–48.

Mazzucato, M., Semieniuk, G., 2018. Financing renewable energy: Who is financing what and why it matters. *Technological Forecasting and Social Change* 127, 8–22.

Miralles-Quiros, J. L., Miralles-Quiros, M. M., 2019. Are alternative energies a real alternative for investors? *Energy Economics* 78, 535–545.

Omar, A. M., Wisniewski, T. P., Nolte, S., 2017. Diversifying away the risk of war and cross-border political crisis. *Energy Economics* 64, 494–510.

Park, A., Saburian, H., 2011. Herding and contrarian behavior in financial markets. *Econometrica* 79, 973–1026.

Patterson, D.M., Sharma, V., 2006. Do traders follow each other at the NYSE? University of Michigan–Dearborn Working Paper.

Pesaran, M.H., Shin, Y., 1998. Generalized Impulse Response analysis in linear multivariate models. *Economics Letters* 58, 17–29.

Philippas, N., Economou, F., Babalos, V., Kostakis, A., 2013. Herding behavior in REITs: Novel tests and the role of financial crisis. *International Review of Financial Analysis* 29, 166–174.

Pickl, M. J., 2019. The renewable energy strategies of oil majors–From oil to energy? *Energy Strategy Reviews* 26, 100370.

Rahdari, A. H., Rostamy, A. A. A., 2015. Designing a general set of sustainability indicators at the corporate level. *Journal of Cleaner Production* 108, 757–771.

Rubinstein, M., 2001. Rational markets: yes or no? The affirmative case. *Financial Analysts Journal* 57, 15–29.

Schmitt, N., Westerhoff, F., 2017. Herding behaviour and volatility clustering in financial markets. *Quantitative Finance* 17, 1187–1203.

Sias, R.W., 2004. Institutional herding. *Review of Financial Studies* 17, 165–206.

Shiller, R., 2002. Bubbles, Human Judgement, and Expert Opinion. *Financial Analysts Journal* 58, 18–26.

Tan, L., Chiang, T.C., Mason, J.R., Nelling, E., 2008. Herding behavior in Chinese stock markets: An examination of A and B shares. *Pacific-Basin Finance Journal* 16, 61–77.

Thoma, J., Chenet, H., 2017. Transition risks and market failure: a theoretical discourse on why financial models and economic agents may misprice risk related to the transition to a low-carbon economy. *Journal of sustainable finance & investment* 7, 82–98.

Welch, I., 2000. Herding among security analysts. *Journal of Financial Economics* 58, 369–396.

Wustenhagen, R., Menichetti, E., 2012. Strategic choices for renewable energy investment: Conceptual framework and opportunities for further research. *Energy Policy* 40, 1–10.

Yao, J., Ma, C., He, W. P., 2014. Investor herding behaviour of Chinese stock market. *International Review of Economics & Finance* 29, 12–29.

Appendix A

The list below provides details on all the ‘green’ ETFs included in this paper. All these ‘green’ ETFs are traded on the New York Stock Exchange, with the exception of Global Energy Efficient Transport Index (PTRP) and Clean Edge Global Wind Energy Index (PWND), which are traded on the Nasdaq exchange. Data source is Thomson Reuters Eikon.

Table A1. The list of ‘green’ ETFs used in the empirical section

Symbol (start date)	Investment objective
CGW (May 2007)	It replicates the performance of the S&P Global Water Index, which is comprised of 50 securities selected based on the relative importance of the global water industry within the company’s business model.
TAN (April 2008)	It replicates the performance of the MAC Global Solar Energy Index, which is comprised of 25 securities selected based on the relative importance of the solar power within the company’s business model.
FAN (June 2008)	It replicates the performance of the ISE Global Wind Energy Index, which is comprises companies selected based on their actively engagement with the wind energy industry, such as the production of distribution of electricity generated by wind power, and so on.
DSI (November 2006)	It replicates the performance of the FTSE KLD 400 Social Index, which provides the exposure to companies' common stocks that exhibits positive ESG characteristics.
KLD (January 2005)	It replicates the performance of the FTSE KLD Select Social Index, which maximizes the exposure on large capitalization companies that exhibit positive ESG characteristics.
GEX (May 2007)	It replicates the Ardour Global Index, which provides the exposure of publicly traded companies that derive over 50% of total revenues from the alternative energy industry, in the globe.
KWT (April 2008)	It replicates the Ardour Solar Energy Index, which shows the exposure to publicly traded companies that derive at least 66% of their revenues from solar energy, in the global.
PZD (October 2006)	It replicates the Cleantech Index, which is comprised of companies that produces any knowledge-based product, service that improves operation, performance, productivity, efficiency, and, meanwhile they reduce costs, inputs, energy consumption and pollution.
PBD (June 2007)	It replicates the WilderHill New Energy Global Innovation Index, which is comprised companies that focus on green renewable sources of energy and technologies facilitating cleaner energy.

PTRP (September 2008)	It replicates the Wilder NASDAQ OMX Global Energy Efficient Transport Index, which is comprised companies engaged in societal transition toward using cleaner, less costly, and more efficient means of transportation, worldwide.
PIO (June 2007)	It replicates the Palisades Global Water Index, which is comprised international companies engaged with the provision of potable water, the treatment of water and the technology/services directly related to global water consumption.
PWND (July 2008)	It replicates the NASDAQ OMX Clean Edge Global Wind Energy Index, which is comprised of companies engaged with the wind energy industry, such as manufacturers, developers, distributors, installers and so on.
PHO (December 2005)	It replicates the Palisades Global Water Index, which is comprised international companies engaged with the provision of potable water, the treatment of water and the technology/services directly related to global water consumption.
PBW (March 2005)	It replicates the WilderHill Clean Energy Index, which is comprised of companies focusing on green renewable sources of energy and technologies facilitating cleaner energy.
PUW (October 2006)	It replicates the WilderHill Progressive Energy Index, which is comprised of companies that exhibit transitional energy technologies.

Appendix B

This Appendix presents the descriptive statistics for crude oil and various ‘green’ assets with respect to the S&P 500 Energy Index, which is the benchmark index for energy stocks.

Table B1. Descriptive statistics

Relative returns for:	Crude oil	S&P500 Carbon	S&P500 ESG	Global Clean Energy	Global Water	ETF_EQW	ETF_PC1
Mean	-0.135758	0.201691	0.199435	-0.120802	0.13068	-0.007396	0.047149
Median	-0.125961	0.19578	0.227077	-0.031001	0.203419	0.054875	0.271420
Maximum	19.03458	13.06619	13.01171	13.05091	12.72652	11.95351	13.08835
Minimum	-20.5408	-11.32934	-11.45179	-12.4112	-10.3813	-12.11665	-16.91626
Std. dev.	3.9212	2.734804	2.748969	3.434494	2.773523	2.924592	3.64903
Skewness	-0.226564	0.081546	0.08116	-0.294353	0.166531	-0.34964	-0.75756
Kurtosis	5.921579	5.708811	5.63433	4.409422	5.070002	4.982961	5.60703
Jarque-Bera	152.2382	128.2608	121.3252	40.6339	76.56083	77.0013	158.3562
Observations	418	418	418	418	418	418	418

Note: Sample refers to weekly observations between Jan. 2011 and Dec. 2018. All data in the tabel pertains to returns, in log terms, and taken with respect to S&P 500 Energy Index, i.e. relative returns. All indicators listed on the top row are described in Data section 3. ETF_EQW denotes the equally weighted (in USD terms) portfolio comprising all 16 ETFs mentioned in Data section 3; ETF_PC1 denotes the first principal component of the returns series computed from all 16 ETFs.

Appendix C

This appendix displays the bootstrapped GIRFs for the two VARs discussed in section 4.2. Figure C1 presents the GIRFs from a VAR with oil and 'green' returns; figure C2 presents the GIRFs from a VAR with oil and 'green' conditional volatility. The GIRFs depicted in figure 2 are highlighted with red.

Figure C1. Full set of bootstrapped GIRFs from a VAR with oil and 'green' returns

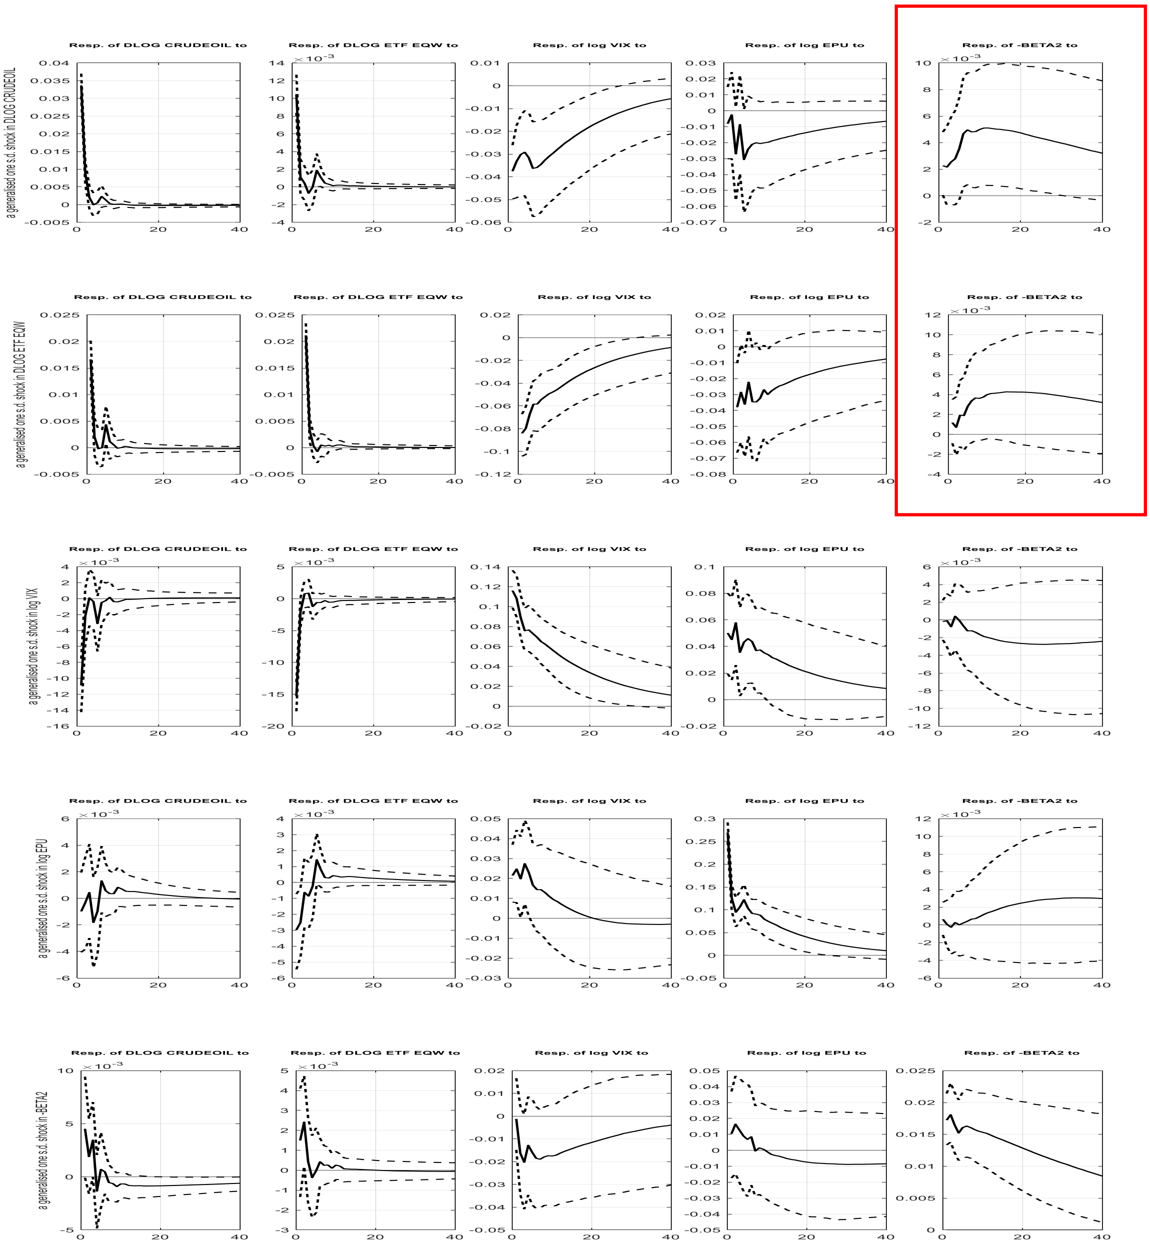

Note: The figure displays the bootstrapped GIRFs from a 5-variable VAR specified in the following variables: weekly returns for WTI crude oil, denoted as DLOG_CRUDEOIL; weekly returns of an equally-weighted portfolio of 16 'green' ETFs, denoted as DLOG_ETF_EQW; log of VIX index; log of EPU index; and the negative of the herding proxy, denoted as -BETA2. Estimation sample includes weekly observations from Apr. 2011 to Dec. 2018. A number of 5000 bootstrap replications are used to derive median responses along with (+/- 2 standard deviations) confidence bands.

Figure C2. Full set of bootstrapped GIRFs from a VAR with oil and 'green' conditional volatility

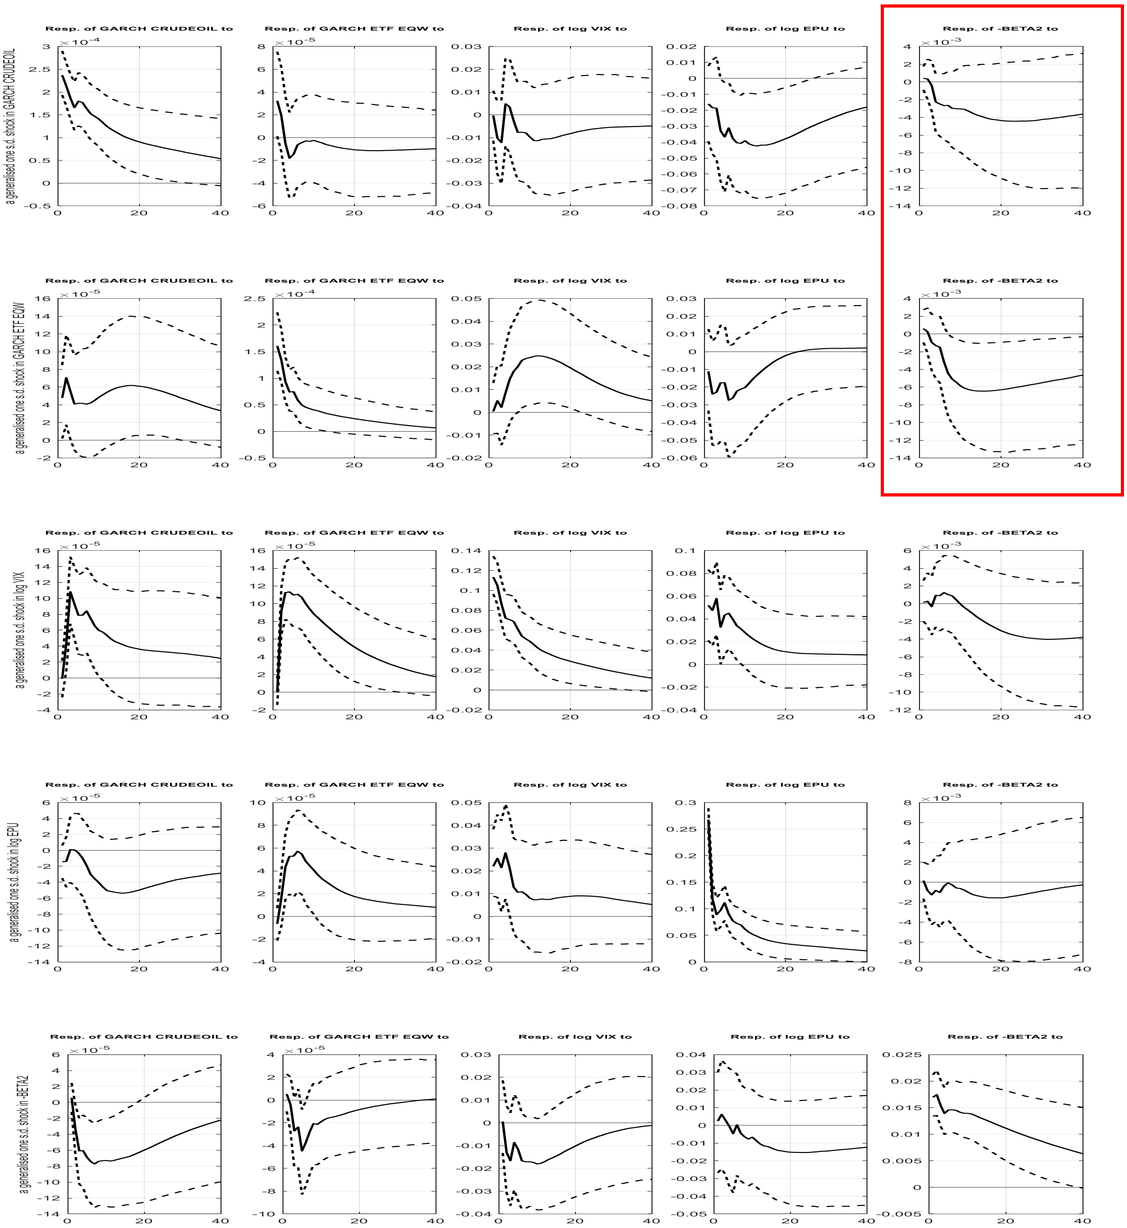

Note: The figure displays the bootstrapped GIRFs from a 5-variable VAR specified in the following variables: conditional volatility for WTI crude oil denoted as GARCH_CRUDEOIL; conditional volatility of an equally-weighted portfolio of 16 'green' ETFs, denoted as GARCH_ETF_EQW; log of VIX index; log of EPU index; and the negative of the herding proxy, denoted as -BETA2. Estimation sample includes weekly observations from Apr. 2011 to Dec. 2018. A number of 5000 bootstrap replications are used to derive median responses along with (+/- 2 standard deviations) confidence bands.

Appendix D

This appendix presents the bootstrapped GIRFs for some of the VARs discussed in section 4.3.

Figure D1. Selected bootstrapped GIRFs from VARs with spot crude and futures oil

Panel A

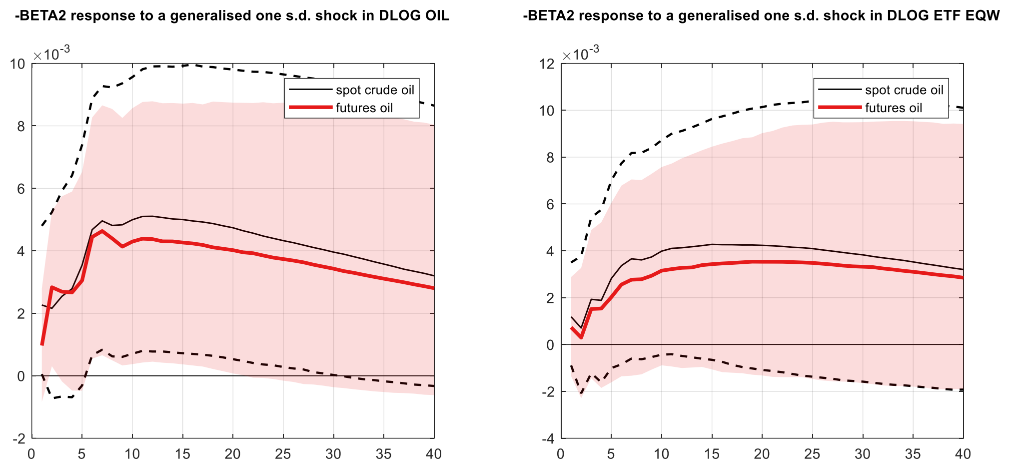

Panel B

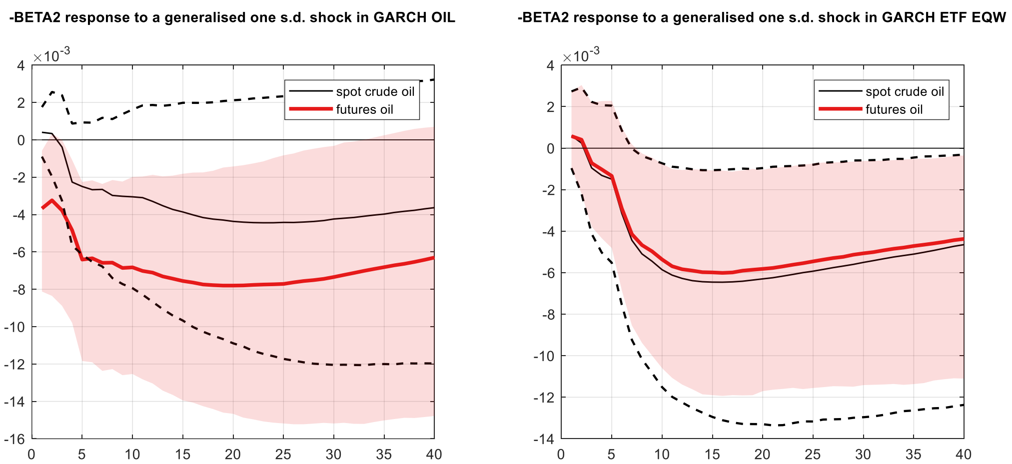

Note: The figure displays the GIRFs with bootstrapped confidence bands and median responses to shocks in returns (i.e. DLOG terms) and conditional volatilities (i.e. GARCH terms). For *panel A*, a 5-variable VAR is estimated, specified in the following variables: weekly returns for WTI crude or futures oil; weekly returns for an equally-weighted portfolio of 16 'green' ETFs; log of VIX index; log of EPU index; and the negative of the herding proxy, denoted as -BETA2. For *panel B*, a 5-variable VAR is estimated, specified in the following variables: conditional volatility of WTI crude or futures oil; conditional volatility of an equally-weighted portfolio of 16 'green' ETFs; log of EPU index; and the negative of the herding proxy, denoted as -BETA2. In black, the VAR includes spot crude oil; in red, the VAR includes futures oil. The herding proxy is always taken with a negative sign in the VARs, such that an increase in herding is associated with positive values. Estimation sample includes weekly observations from Apr. 2011 to Dec. 2018. A number of 5000 bootstrap replications are used to derive median responses along with (+/- 2 standard deviations) confidence bands.

Figure D2. Selected bootstrapped GIRFs from VARs with EPU index and EPU *Regulation* sub-index

Panel A

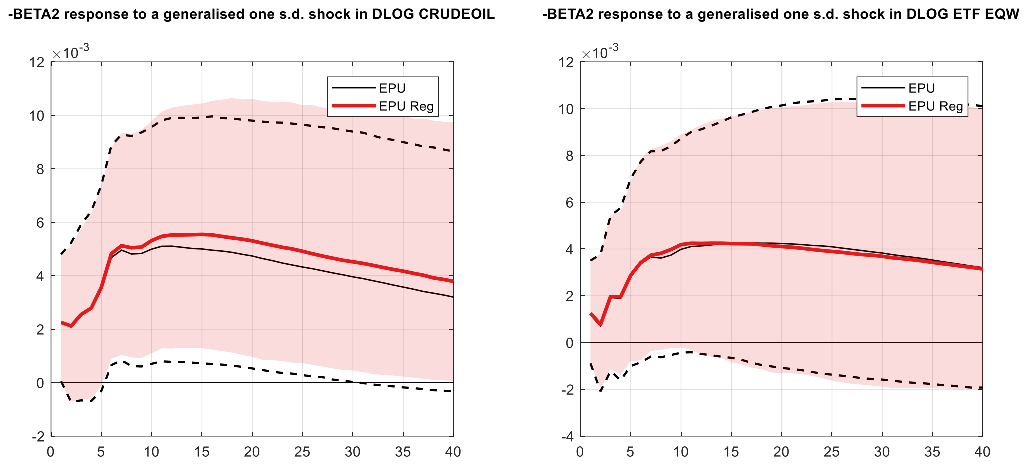

Panel B

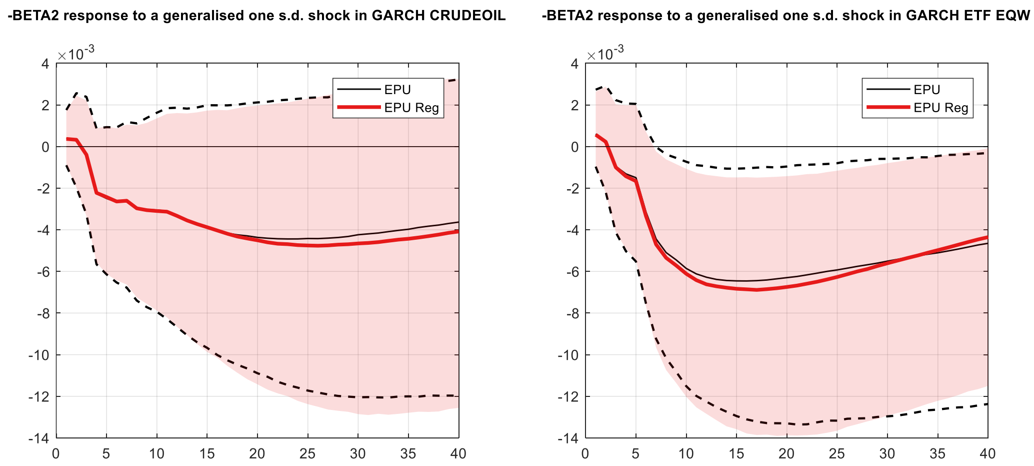

Note: The figure displays the GIRFs with bootstrapped confidence bands and median responses to shocks in returns (i.e. DLOG terms) and conditional volatilities (i.e. GARCH terms). For *panel A*, a 5-variable VAR is estimated, specified in the following variables: weekly returns for WTI crude oil; weekly returns for an equally-weighted portfolio of 16 'green' ETFs; log of VIX index; log of EPU index or EPU Regulation sub-index; and the negative of the herding proxy, denoted as -BETA2. For *panel B*, a 5-variable VAR is estimated, specified in the following variables: conditional volatility for WTI crude oil; conditional volatility of an equally-weighted portfolio of 16 'green' ETFs; log of EPU index or EPU Regulation sub-index; and the negative of the herding proxy, denoted as -BETA2. In black, the VAR includes the log of EPU index; in red, the VAR includes the EPU *Regulation* sub-index. The herding proxy is always taken with a negative sign in the VARs, such that an increase in herding is associated with positive values. Estimation sample includes weekly observations from Apr. 2011 to Dec. 2018. A number of 5000 bootstrap replications are used to derive median responses along with (+/- 2 standard deviations) confidence bands.

Figure D3. Selected bootstrapped GIRFs from VARs with the herding proxy filtered out both as AR(1) and I(1) processes

Panel A

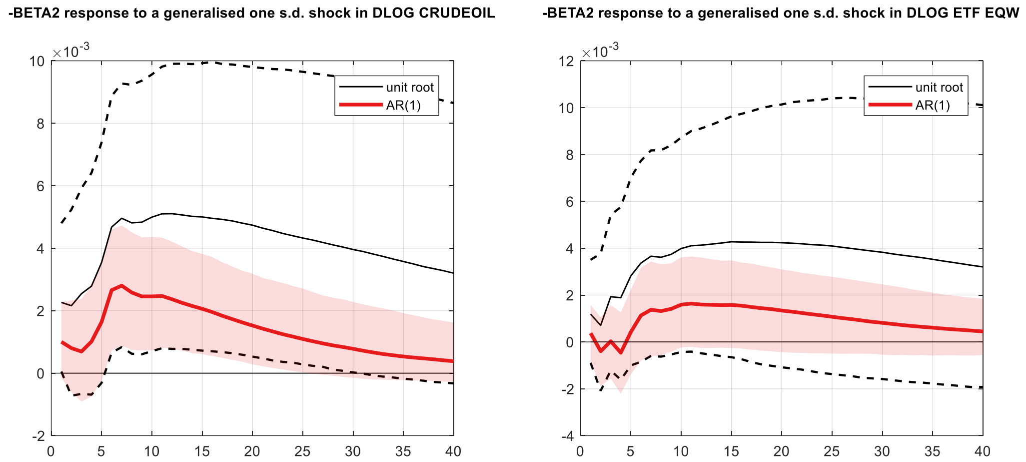

Panel B

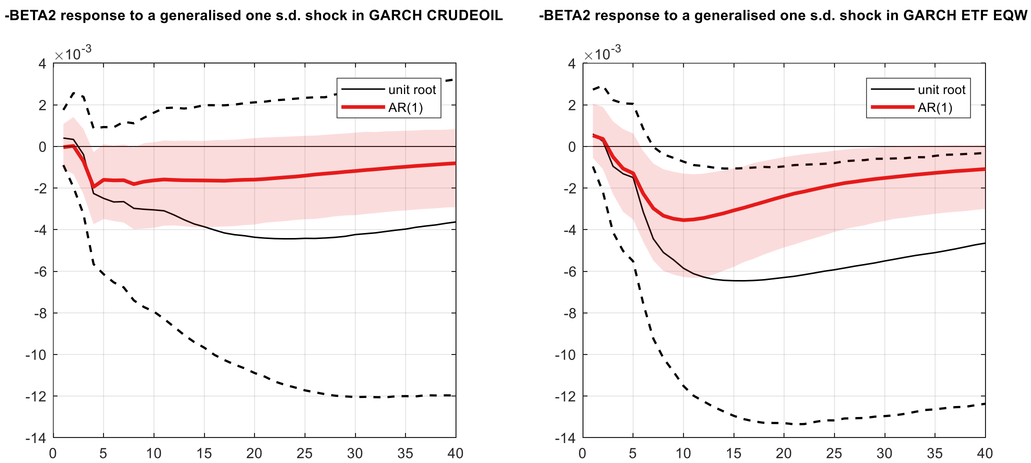

Note: The figure displays the GIRFs with bootstrapped confidence bands and median responses to shocks in returns (i.e. DLOG terms) and conditional volatilities (i.e. GARCH terms). For *panel A*, a 5-variable VAR is estimated, specified in the following variables: weekly returns for WTI crude oil; weekly returns for an equally-weighted portfolio of 16 'green' ETFs; log of VIX index; log of EPU index; and the negative of the herding proxy. For *panel B*, a 5-variable VAR is estimated, specified in the following variables: conditional volatility for WTI crude oil; conditional volatility of an equally-weighted portfolio of 16 'green' ETFs; log of EPU index; and the negative of the herding proxy. In black, the VAR includes the herding proxy filtered out as a unit root process, or I(1), as in section 4.1 (thus corresponding to Figure 2 in the main text); in red, the VAR includes the herding proxy filtered out as an AR(1) process. The herding proxy is always taken with a negative sign in the VARs, such that an increase in herding is associated with positive values. Estimation sample includes weekly observations from Apr. 2011 to Dec. 2018. A number of 5000 bootstrap replications are used to derive median responses along with (+/- 2 standard deviations) confidence bands.

Figure D4. Bootstrapped GIRFs from a 5-variable and 6-variable VAR including the NEWS variable

Panel A

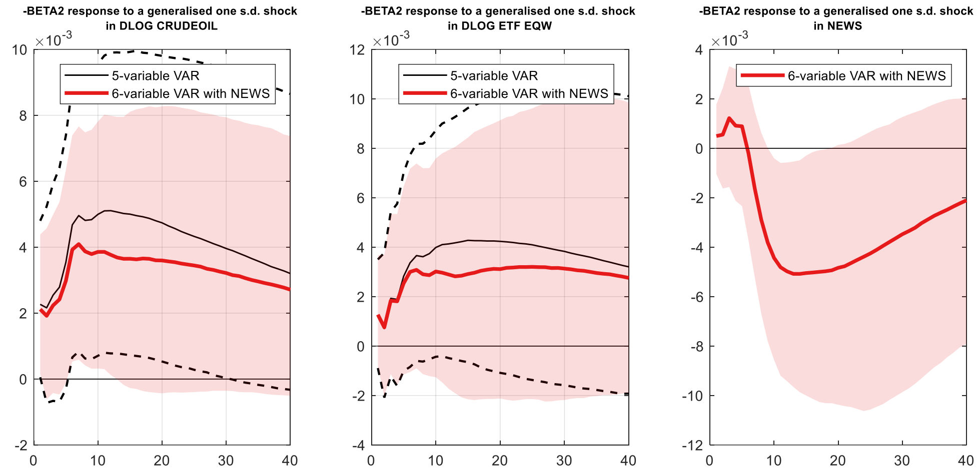

Panel B

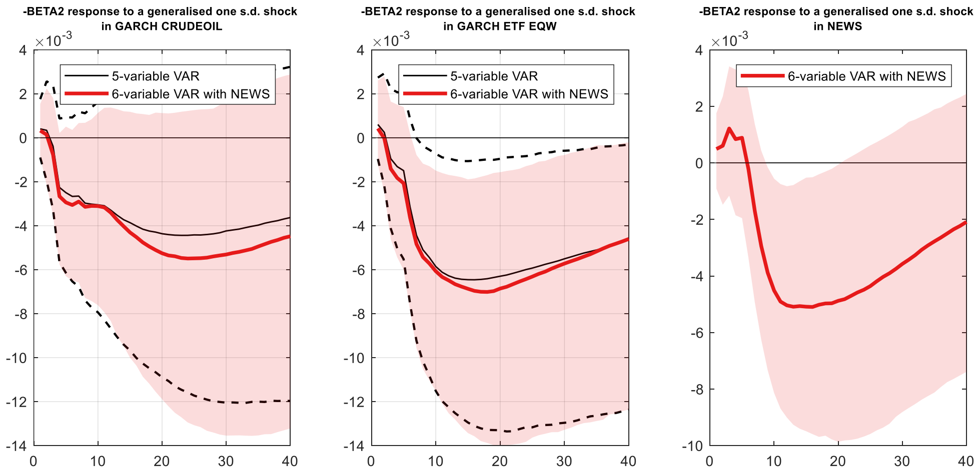

Note: The figure displays the GIRFs with bootstrapped confidence bands and median responses to shocks in returns (i.e. DLOG terms), conditional volatilities (i.e. GARCH terms) and media coverage of climate-related topics – a variable we denote as NEWS. For *panel A*, a 5-variable VAR and a 6-variable VAR are estimated, where the latter additionally includes the NEWS variable; both VARs share the following variables: weekly WTI crude oil returns; weekly returns for an equally-weighted portfolio of 16 ‘green’ ETFs; log of VIX index; log of EPU index; and the negative of the herding proxy denoted as -BETA2. For *panel B*, a 5-variable VAR and a 6-variable VAR are estimated, where the latter additionally includes the NEWS variable; both VARs share the following variables: conditional volatility for WTI crude oil; conditional volatility of an equally weighted portfolio of 16 ‘green’ ETFs; log of VIX index; log of EPU index; and the negative of the herding proxy denoted as -BETA2. In black, we display the GIRFs from the 5-variables VAR (thus corresponding to Figure 2 in the main text); in red, we display the GIRFs from the 6-variable VAR, which includes the NEWS variable. The herding proxy is always taken with a negative sign in the VARs, such that an increase in herding is associated with positive values. Estimation sample includes weekly observations from Apr. 2011 to Dec. 2018. A number of 5000 bootstrap replications are used to derive median responses along with (+/- 2 standard deviations) confidence bands.