European Economic Review 163 (2024) 104680

Contents lists available at ScienceDirect

European Economic Review

journal homepage: www.elsevier.com/locate/eer

Financial markets and legal challenges to unconventional monetary policy

Stefan Griller [a], Florian Huber [b,*,1], Michael Pfarrhofer [a]

[a] Vienna University of Economics and Business, Austria
[b] University of Salzburg, Austria

ARTICLE INFO

Dataset link: https://ars.els-cdn.com/content/image/1-s2.0-S0014292124000096-mmc1.zip

JEL classification:
E43
E44
E58
K00

Keywords:
Local projections
Euro area
Finite mixture models
Factor models
Identification through heteroskedasticity

ABSTRACT

This paper studies the empirical effects of legal challenges to monetary policy. Several policy measures of the European Central Bank have come under scrutiny before national courts and the European Court of Justice. These lawsuits have the potential to impact the scope and flexibility of central bank policies with important consequences for the economy. The number of relevant legal challenges is small and exact timestamps are often not available. Hence, we develop an econometric approach that searches for minimum variance regimes which we use to isolate and measure the effects of the rulings. Our results suggest that legal rulings addressing central bank policies have a significant effect on financial markets and can be linked to monetary policy uncertainty shocks.

1. Introduction

This paper studies the empirical effects of legal challenges to monetary policy. In the recent history of central banking, legislators have generally been improving legal architectures and procedures to free central banks from undue political pressures and influence. Measuring the economic or financial consequences of such changes is difficult — however, there exists some evidence that related shifts in central bank policies in the United States (US) and the United Kingdom have led to declining levels of inflation, decreasing inflation volatility and stabilization of inflation expectations (see, e.g., Cukierman and Meltzer, 1986; Cukierman et al., 1992; Blinder et al., 2008).

During the global financial crisis of 2007/08 and the Covid-19 pandemic, central banks like the Federal Reserve Bank (Fed) and the Bank of England (BoE) actively cooperated with governments to fight these economic downturns. Such coordinated action did not interfere with the legal structure of the Fed or the BoE. In the euro area (EA), coordinated action also took place, but subject

* Corresponding author.
E-mail address: florian.huber@plus.ac.at (F. Huber).

[1] The authors are grateful to Christiane Baumeister Marek Jarociński, two anonymous referees, an anonymous associate editor, Evi Pappa (the editor) and participants of the 27th Annual International Conference on Macroeconomic Analysis and International Finance and internal research seminars at the Vienna University of Economics and Business and the University of Salzburg for providing many helpful suggestions and parts of the dataset. In addition, we thank Andres Azqueta-Gavaldón and Luca Onorante for sending us their daily monetary policy uncertainty index. This research is supported by funds of the Oesterreichische Nationalbank (OeNB) (Austrian Central Bank, Anniversary Fund, project number: 18304). Moreover, this research was funded in part by the Austrian Science Fund (FWF) 10.55776/ZK35. For open access purposes, the author has applied a CC BY public copyright license to any author accepted manuscript version arising from this submission. Codes and replication files are available at: https://github.com/mpfarrho/idhet-mix.

https://doi.org/10.1016/j.euroecorev.2024.104680
Received 8 November 2022; Received in revised form 15 January 2024; Accepted 18 January 2024
Available online 24 January 2024

to a complicated legal framework and with member states exhibiting heterogeneous macroeconomic, financial and debt dynamics. The question whether the European Central Bank (ECB) is legally allowed to intervene in this regard naturally arose in the media and policy discussions.

Besides the aforementioned severe economic crises, ECB expertise and capacity have been challenged also by the European sovereign debt crisis. This provoked unconventional monetary policy responses (see Hartmann and Smets, 2018, for a review), and involved buying government debt securities in unprecedented amounts. Related programs have been criticized by a number of politicians, political commentators and journalists. And more importantly, they came under legal scrutiny. A key question among others was whether such programs constitute monetary or rather economic policy (a protected competence of EU member states), or could even be classified as (prohibited) monetary financing. In other words, it was unclear whether the ECB potentially exceeded its mandate. Subsequent lawsuits about the legality of these unconventional monetary policy measures could, in the worst case, severely limit the scope of the ECB's discretionary options during crises and thus damage its credibility in the context of timely and effective policy making.

We take a closer look at three rulings handed down by the Court of Justice of the European Union (CJEU) at the request of highest national courts, and at the related decisions of these courts. We scrutinize the Pringle Case which concerned the legality of establishing the European Stability Mechanism (ESM) which is a member states' measure, and was brought up by the Supreme Court of Ireland; further, we assess the Gauweiler and the Weiss Case, both raised by the Bundesverfassungsgericht (German Federal Constitutional Court, GFCC), and both concerning non-traditional ECB measures which allegedly were outside the ECB's mandate.

Empirically we develop a model that exploits variation surrounding the announcements of these rulings to pin down the effects they had on a broad panel of financial indicators on a daily frequency. Our econometric approach is thus closely related to the heteroskedasticity-based identification scheme proposed in Gürkaynak et al. (2020).[2] The key econometric challenge in our case is that the number of ruling days relative to the non-ruling days is very small and no exact timestamps of the announcements are available. We propose to contrast ruling with non-*event* (rather than non-*ruling*) days and achieve identification by assuming that the variances differ between these two sets of observations due to a factor stochastic volatility (SV) model. The validity of this assumption hinges upon selecting adequate non-event days which we obtain as a subset of the non-ruling days using a mixture model. This mixture model yields endogenously estimated clusters of observations over time that share a common variance. The non-event set is defined to comprise of days in the minimum variance cluster, which we argue indicates a lack of noteworthy events.

To estimate the effects of ruling shocks we compute local projections within a unified econometric framework. Such shocks arise from statements of a court about the previous conduct of ECB monetary policy, which in turn affects expectations about the scope and magnitude of future possible policies by clarifying the respective legal boundaries. Shifts in expectations thus likely move financial markets on announcement dates, and these are the major impacts we seek to identify. In addition, the rulings may affect monetary policy uncertainty (MPU, see, e.g., Husted et al., 2020) by broadening or constraining the ECB's leeway with implications for risks to the euro (area), and thereby likely the heterogeneity of expectations of financial market participants about, e.g., future interest rates.

Our dataset is comprised of a broad panel of daily EA-wide, country-level and international time series. Among several key variables of interest, we control for cross-country movements in long-term yield spreads by including government bond yields relative to German Bunds for several EA member states. These series are complemented by financial indicators such as measures of systemic stress, short-term interest rates, corporate bond spreads, as well as stock markets, an exchange rate, and the price of gold. To measure how rulings impact MPU, we include a news-based measure of MPU, proposed in Azqueta-Gavaldón et al. (2023), in our set of observed variables.

The findings suggest that financial conditions across EA countries improve in response to an "expansionary" ruling shock. With expansionary we mean that the respective court rejected allegations that the ECB acted outside its mandate. These effects are heterogeneous across countries, with peripheral economies exhibiting much stronger reactions than core countries. Notably, 10-year bond yields relative to German Bunds decrease for some member states, and specifically Greek government bond yield spreads decline substantially. Stock markets display a somewhat similar, mixed reaction, with stock returns in Greece increasing appreciably over the impulse response horizon. Moreover, the daily analysis shows that the euro appreciates relative to the US dollar and the price of gold declines alongside the MPU. These dynamics point towards an improved economic outlook in response to expansionary ruling shocks.

To understand whether these findings also carry over to monthly data and the aggregate EA-level, we estimate a VAR and include the identified ruling shock as an instrument to back out structural impulse responses. These results confirm the findings for daily data: financial conditions ease, stocks increase on average and spreads decline. Including price and output series also allows us to analyze whether rulings impacted the real side of the economy. Indeed, the ruling shocks decrease MPU also on a monthly frequency, and our findings are similar to the effects of negative (uncertainty-reducing) MPU shocks (see Husted et al., 2020; Bauer et al., 2022); expansionary ruling shocks lead to higher prices and production.

The paper proceeds as follows. Section 2 provides the legal background and a summary of the legal cases which form our set of events. Section 3 introduces the econometric framework used to estimate the effects of the rulings. Section 4 presents the main empirical results of the paper. The final section concludes, while an Online Appendix contains further detailed legal considerations, comprehensive information about the estimation procedure, and additional empirical results alongside robustness checks.

[2] More generally, the feature of exploiting heteroskedasticity for shock identification relates our paper to a broad literature on using time-varying volatilities in macroeconomic and financial data to pin down structural shocks statistically to enable economically meaningful interpretation (see, e.g., Sentana and Fiorentini, 2001; Rigobon, 2003; Rigobon and Sack, 2004; Lewis, 2021; Brunnermeier et al., 2021; Bertsche and Braun, 2022).

2. Fundamental legal dimensions

2.1. The European Central Bank and the law

The ECB's independence is triple guaranteed by primary law, in the Treaty on the Functioning of the European Union (TFEU) and the Statute of the European System of Central Banks (ESCB) and of the ECB.[3] Moreover, it is bolstered by flanking provisions, among those the prohibition on monetary financing, and the appointment mechanism as well as the not renewable eight-years term of office for members of the Executive Board. The underlying rationale is that delivering on the ECB's primary objective, price stability, would be made more difficult if not endangered by external, not the least political influence of any kind.

ECB operations are guaranteed by the Treaties and thus strongly protected under European Union (EU) law not only against undue influence on ECB organs but also against encroachments by member states' and other EU institutions' activities. Any violation can be brought before the CJEU. For "monetary policy for the Member States whose currency is the euro" is an exclusive competence of the EU, with the ECB as the institution entrusted to take care of; "only the Union may legislate and adopt legally binding acts, the Member States being able to do so themselves only if so empowered by the Union or for the implementation of Union acts".

This, however, does not exempt the ECB from the EU legal order. The ECB is subject to legal scrutiny by the CJEU. The latter emphasized early on that independence "does not have the consequence of separating" the ECB from every rule of EU law. First and before anything else, the ECB must act within the limits of the powers conferred on it, which is to "conduct the *monetary* policy of the Union". Acting outside that mandate, or more generally, violating EU law, would render ECB measures voidable by the CJEU.

2.2. Legal controversies about delimiting monetary and economic policy

Two aspects became controversial during and in the aftermath of the global financial crisis: the member states' respect for the ECB's powers; and, inversely, the ECB's respect for its own limits which protect the member states' economic (including fiscal) policy powers which they did not transfer to but merely coordinate within the EU. The controversy led to several high-profile judgments clarifying the limits between *economic* and *monetary* policy. The cases will be briefly sketched out below (for a more detailed analysis, see Griller, 2021). All three of them had been filed with the CJEU by highest national courts as preliminary ruling procedures, putting critical questions concerning the legality of measures taken. Only the second and third judgment (Gauweiler and Weiss) directly concerned (unconventional) ECB measures, the first (Pringle) was on a measure taken by the member states. However, all three of them dealt with the scope of monetary policy, and we therefore regard them as legal events affecting the ECB.[4] Table 1 provides a list and further details.[5]

The outcomes had been far from obvious. The CJEU's judgments included a certain margin of discretion. Indeed, all these judgments have been controversial before, during and in the aftermath of litigation, both at CJEU and national level. To put it differently, *ex ante*, these lawsuits had the potential of ousting crisis intervention measures as illegal. As a result, these measures would have had to be terminated. That would have caused a serious challenge to both the member states' and the ECB's efforts to contain the respective crises and to counteract their adverse effects. Thus, stock market and interest rate reactions to these decisions, or financial market dynamics in general, are pertinent when it comes to measuring the economic relevance of these legal decisions. The reactions reflect possible expectations or suspicions of alternative outcomes. If the court decision as such comes as a surprise to the markets, it is this unexpected event which may be considered as a shock that is provoking market reactions. In addition, the court ruling may reduce/increase uncertainty about the conduct of monetary policy, and thus provoke economic and financial effects.

Preliminary ruling procedures as established under EU law typically provide at least four, and sometimes even more, prominent stages which are worthwhile examining. First, the decision of the national court to interrupt its proceedings and ask the CJEU for an answer to a difficult question of interpretation or on the validity of secondary EU law. Second, the (non-binding) opinion of the CJEU's Advocate General whereby the latter proposes how the CJEU should decide. Third, the requested answer of the CJEU given in its judgment which is binding under EU law. Fourth, the judgment or decision of the national court based on the answer given by the CJEU.

[3] The legal exposition in this paper features many perhaps rather uncommon abbreviations and acronyms. In addition, there are several references to legal documents or treaties alongside direct quotations taken from them. To avoid disrupting the reading flow, we omit related sources in the main text. We provide an overview and more detailed information, alongside legal references, in Online Appendix A.

[4] The ECB's Outright Monetary Transactions are conditional on a memorandum of understanding between a member state and the EMS, so the Pringle case had the potential to nullify the Outright Monetary Transactions, thus directly impacting the ECB's toolkit.

[5] We have verified that the dates of the rulings do not systematically coincide with other important news. In particular, we have cross-checked with the dates of the following press releases: Eurostat releases of EA flash CPI, core CPI, produces prices, M3 money growth, trade balance, unemployment rate, business and consumer surveys, national accounts, industrial production, retail sales; Markit EA PMI indices; ECB monetary policy announcements, speeches by ECB Presidents; Federal Reserve Bank monetary policy announcements (scheduled and unscheduled FOMC press releases); US initial jobless claims and non-farm payrolls. For most releases the overlap is on less than three dates, and most often on zero dates. We discuss potential issues arising from such overlaps, and how to address them econometrically, in Section 3.

Table 1
Ruling dates and description.

Decision	Publication	Number	Court	Description
Case C-370 (Pringle), ECLI:EU:C:2012:756				
2012-07-31	2012-07-31	[2012] IESC 47	SC, Ireland	Request for preliminary ruling
2012-10-26	2012-10-26	C-317/12	CJEU, AG	Opinion of the AG
2012-11-27	2012-11-27	C-317/12	CJEU	Judgment
Case C-62/14 (Gauweiler), ECLI:EU:C:2015:400				
2014-01-14	2014-02-07	2 BvR 2728/13	GFCC	Request for preliminary ruling
2015-01-14	2015-01-14	C62/14	CJEU, AG	Opinion of the AG
2015-06-16	2015-06-16	C62/14	CJEU	Judgment
2016-06-21	2016-06-21	2 BvR 2728/13	GFCC	Judgment
		2 BvR 2729/13		
		2 BvR 2730/13		
		2 BvR 2731/13		
		2 BvE 13/13		
Case C-493/17 (Weiss), ECLI:EU:C:2018:1000				
2017-07-18	2017-08-15	2 BvR 859/15	GFCC	Request for preliminary ruling
		2 BvR 1651/15		
		2 BvR 2006/15		
		2 BvR 980/16		
2018-10-04	2018-10-04	C-493/17	CJEU, AG	Opinion of the AG
2018-12-11	2018-12-11	C-493/17	CJEU	Judgment
2020-05-05	2020-05-05	2 BvR 859/15	GFCC	Judgment
		2 BvR 1651/15		
		2 BvR 2006/15		
		2 BvR 980/16		

Notes: Abbreviations are Advocate General (AG), Court of Justice of the European Union (CJEU), European Commission (EC), German Federal Constitutional Court (GFCC), Supreme Court (SC). In our empirical work, we consider publication dates in case they differ from the decision dates for the timing of the shocks. Publication dates are those of relevance, since this is when the public is informed about the respective decision and judges are prohibited by law from disclosing their rulings prematurely.

2.3. High-profile judgments

We provide a brief summary of the three high-profile judgments that lie at the heart of this paper below. This is to avoid excessive legal jargon, but still give an informative and accurate description of what has been at stake in the respective cases. A more thorough legal discussion is provided in Online Appendix A.

The Pringle Case. The ESM treaty came under scrutiny, and was challenged before the Supreme Court of Ireland. The Irish challenge alleged that ratifying this treaty would violate provisions of key other EU treaties related to the division of competences among member states and EU institutions. The CJEU rejected related claims. The key finding of the court relevant for this paper is that it clearly stated what it considers as monetary policy. Specifically, monetary policy was defined by its objectives (i.e., mainly price stability), rather than based on the actual instruments and implementation used to achieve these objectives. In a wide sense, the CJEU thus commented on the ECB's mandate from a legal perspective, and in a narrow sense, ensured the continuation of the ESM despite concerns about its potential impact on the European Monetary Union (EMU) in general, and the ECB in particular.

The Gauweiler Case. This judgment concerned Outright Monetary Transactions (OMT) by the ECB, that is, to conduct purchases in secondary sovereign bond markets. The program announced in the context of Mario Draghi's "Whatever it takes" speech aimed at ensuring the transmission of monetary policy and the integrity of the euro. It never became operative, but was nonetheless challenged before the GFCC — which questioned the program's compatibility with the ECB's mandate. The CJEU subsequently ruled in favor of the ECB, deeming OMT legal. It emphasized that maintaining the singleness of monetary policy and an appropriate transmission would support the primary objective of price stability, confirming its assessment about what constitutes monetary policy from the Pringle case. Thus, purchasing government bonds on the secondary market might be economic policy if done by the ESM, but monetary policy if done by the ECB. The GFCC accepted the CJEU's verdict, leading to the dismissal of the constitutional complaints against the OMT program. A different outcome could have invalidated and prohibited the program.

The Weiss Case. This preliminary ruling procedure initiated by the GFCC involved the "Public Sector Purchase Programme" (PSPP), a significant part of the ECB's quantitative easing measures. The CJEU again upheld the program, stating that the ECB had not exceeded its mandate, with similar arguments delineating economic and monetary policy than before. However, different to the Gauweiler case, the GFCC refused to accept the CJEU's ruling and declared both the ECB's policies and the CJEU's judgment as *ultra vires* acts, depriving them of binding force in Germany. The GFCC held that the German government had violated constitutional rights by failing to challenge the ECB's PSPP announcement. Consequently, the Bundesbank was temporarily prohibited from participating in the PSPP. Despite subsequent efforts to seek execution of the decision, the GFCC deemed such applications inadmissible. Among other things, the GFCC stated that the ECB Governing Council's assessment of the proportionality of the PSPP had been sufficient, thereby paving the way for the full implementation of the program.

This brief discussion shows that courts have the potential to limit the possibilities and flexibility of the ECB to fight adverse economic or other events. Next, we develop an econometric model with the goal of measuring the effects associated with these rulings on a large panel of mostly financial indicators for selected EA countries.

3. Modeling and measuring ruling shocks

Our econometric framework assumes that rulings have an immediate effect on markets at high frequencies (in our case, daily), relative to days without significant economic, political or other relevant events (the non-event days). These movements are used to identify the causal effect of legal challenges on financial markets. We face two key issues. First, and most importantly, the number of rulings is relatively small and actual judgments more often than not support the ECB's monetary policy action. Second, our assumption of a difference in variances requires knowledge of the non-event days. Both issues motivate a new econometric approach that allows us to back out the effects of legal challenges to the conduct of ECB monetary policy and to control for the fact that agents may adjust their expectations by learning from previous rulings.

The approach builds on and extends the model developed in Gürkaynak et al. (2020). Our framework can be used in situations when no exact timestamps of the events are available (different to, e.g., high-frequency identification of monetary policy announcements). Let $\boldsymbol{y}_t = (y_{1t}, \dots, y_{Mt})'$ denote a panel of M variables observed over a sampling period $t = 1, \dots, T$. For simplicity, we assume that $\boldsymbol{y}_t$ has an unconditional mean equal to zero. We partition the sampling period into three sets: the (1) *ruling* days, i.e., those days/events on which rulings are made public, in the set $\mathcal{A}_{\mathrm{R}} = \{t_1, t_2, \dots, t_{\mathfrak{r}}\}$ which is formed based on the $\mathfrak{r} = 11$ dates listed in Table 1; the (2) *non-ruling* days in $\mathcal{A}_{\mathrm{NR}}$, i.e., the complementary set of $\mathcal{A}_{\mathrm{R}}$ which collects $t \notin \mathcal{A}_{\mathrm{R}}$ such that $\mathcal{A}_{\mathrm{R}} \cup \mathcal{A}_{\mathrm{NR}} = \{1, \dots, T\}$; and the so-called (3) *non-event* set $\mathcal{A}_{\mathrm{NE}}$, with $\mathcal{A}_{\mathrm{NE}} \subset \mathcal{A}_{\mathrm{NR}}$. The non-event set is comprised of days characterized by a lack of important events that move the endogenous variables $\boldsymbol{y}_t$. We discuss specifics about how this set is formed in Section 3.2. Here, it suffices to note that it will feature a cluster of observations with minimal variance.

Let d_t denote a binary indicator which is 1 if $t \in \mathcal{A}_{\mathrm{R}}$ and 0 otherwise. We discuss the more general case of $t \in \mathcal{A}_{\mathrm{NR}}$ below in the context of identifying the set $t \in \mathcal{A}_{\mathrm{NE}}$. We assume that $\boldsymbol{y}_t$ follows a partially heteroskedastic factor model:

$$\boldsymbol{y}_t = d_t \cdot \boldsymbol{\beta} f_t + \boldsymbol{\Lambda} \boldsymbol{q}_t + \boldsymbol{B} \boldsymbol{x}_t + \boldsymbol{\epsilon}_t, \quad \text{for } t \in \mathcal{A}_{\mathrm{R}} \cup \mathcal{A}_{\mathrm{NE}}. \tag{1}$$

Eq. (1) consists of several key components.[6] The arguably most important parameters are those collected in $\boldsymbol{\beta} = (\beta_1, \dots, \beta_M)'$, which denotes an $M \times 1$-vector of loadings (with time variation governed by d_t so that $\boldsymbol{\beta}_t = d_t \cdot \boldsymbol{\beta}$) associated with a static, conditionally independent, heteroskedastic ruling factor. The time-varying factor loadings imply that the ruling factor only enters the model on ruling dates:

$$f_{t_i} | h_{t_i} \sim \mathcal{N}\left(0, \exp(h_{t_i})\right) \quad \Leftrightarrow \quad f_{t_i} = \exp(h_{t_i}/2) e_{t_i}, \quad e_{t_i} \sim \mathcal{N}(0, 1).$$

The loadings capture the impact of the ruling shocks on the endogenous variables. The logarithm of the SVs evolve according to an AR(1) process:

$$h_{t_i} = \phi_h h_{t_{i-1}} + \sigma_h u_{t_i}, \quad u_{t_i} \sim \mathcal{N}(0, 1), \quad \text{for } i = 1, \dots, \mathfrak{r}, \tag{2}$$

where ϕ_h denotes a persistence parameter and σ_h^2 is the variance of the log-volatility. The unconditional mean of the h_{t_i}'s is normalized to zero to pin down the scale of the factor.

Notice that this potentially persistent process introduces sequential dependence among the otherwise independently modeled rulings, and captures uncertainty specific to each ruling. And since f_{t_i} only exists on ruling dates, the volatility process h_{t_i} can also be viewed as a learning rate about corresponding effects. In particular, $\phi_h = 0$ would imply that there is nothing to be learnt for financial market participants on a ruling-by-ruling basis, whereas $\phi_h \approx 1$ suggests that previously published legal decisions affect expectations and uncertainties surrounding decisions. The volatility of the log-volatility process, σ_h^2, regulates the amount of heterogeneity between distinct rulings.

To also provide a contemporaneous link between the observed variables and the latent volatility of the ruling factor (and vice versa), we assume a joint distribution for $(e_{t_i}, u_{t_i})'$. This corresponds to a variant of factor SV with leverage (see, e.g., Li and Scharth, 2022, for a recent discussion):

$$\begin{pmatrix} e_{t_i} \\ u_{t_i} \end{pmatrix} \sim \mathcal{N}(\boldsymbol{0}, \boldsymbol{\Sigma}_\rho), \qquad \boldsymbol{\Sigma}_\rho = \begin{pmatrix} 1 & \rho \\ \rho & 1 \end{pmatrix},$$

such that ρ measures the correlation between shocks to the ruling factor and its volatility, i.e., it connects first and potential second moment aspects of ruling shocks.

To capture cross-sectional correlation across the elements in $\boldsymbol{y}_t$, we include factors that are always-active (i.e., that are included for all t), or "ever-present" in the wording of Gürkaynak et al. (2020). These heteroskedastic factors arise from $\boldsymbol{q}_t \sim \mathcal{N}(\boldsymbol{0}, \boldsymbol{S}_t)$, feature an $M \times Q$ matrix of factor loadings $\boldsymbol{\Lambda}$ and a diagonal covariance matrix $\boldsymbol{S}_t = \mathrm{diag}(\exp(s_{1t}), \dots, \exp(s_{Qt}))$, whose elements follow independent AR(1) processes with zero unconditional mean. The presence of $\boldsymbol{q}_t$ and the assumptions about $\boldsymbol{S}_t$ imply that

[6] Eq. (1) resembles the main specification of Gürkaynak et al. (2020, see their Eq. (7)), albeit without an observed measure of surprises on ruling/announcement days, and several econometric extensions, chief among them the stochastic volatility components. It is also related to the principal components approach described in Rigobon and Sack (2008).

the elements in $\boldsymbol{y}_t$ exhibit time-varying correlation patterns irrespective of the rulings. It is worth stressing that we do not attach any structural meaning to the always-active factors. To do so, additional restrictions (such as sign restrictions, see Korobilis, 2022) would be necessary.

Furthermore, $\boldsymbol{x}_t$ is a $K \times 1$ vector of controls with an $M \times K$ matrix of regression coefficients $\boldsymbol{B}$. These control variables could include an intercept, observed factors, or other high-frequency shocks (e.g., monetary policy surprises), and thus allow for situations where additional events on a ruling day might have had an impact on $\boldsymbol{y}_t$. Finally, we let $\boldsymbol{\epsilon}_t \sim \mathcal{N}(\boldsymbol{0}, \boldsymbol{\Sigma})$ denote a set of Gaussian measurement errors with diagonal covariance matrix $\boldsymbol{\Sigma} = \text{diag}(\sigma_1^2, \dots, \sigma_M^2)$.

3.1. Implications of the model features

The assumptions laid out in the previous sub-section impose specific mostly covariance-related structures on the parameter space. The key implication of Eq. (1) is that on ruling dates (i.e., for $t \in \mathcal{A}_{\text{R}}$), the covariance matrix of $\boldsymbol{y}_t$ is:

$$\text{Var}(\boldsymbol{y}_t) = \underbrace{\exp(h_t)\boldsymbol{\beta}\boldsymbol{\beta}'}_{\text{RF}} + \underbrace{\boldsymbol{\Lambda}\boldsymbol{S}_t\boldsymbol{\Lambda}'}_{\text{AAF}} + \underbrace{\boldsymbol{\Sigma}}_{\text{ID}}, \tag{3}$$

whereas on non-event dates (i.e., for $t \in \mathcal{A}_{\text{NE}}$), it is simply $\text{Var}(\boldsymbol{y}_t) = \boldsymbol{\Lambda}\boldsymbol{S}_t\boldsymbol{\Lambda}' + \boldsymbol{\Sigma}$. The abbreviations refer to the origin of the variance component: the ruling factor (RF), the always-active factors (AAF) or idiosyncratic noise (ID).

Hence, focusing on a single variable, if $\beta_i \neq 0$, our model states that the variance of the ith element of $\boldsymbol{y}_t$ on a ruling day is $\text{Var}(y_{it}) = \exp(h_t)\beta_i^2 + \boldsymbol{\Lambda}_{i\bullet}\boldsymbol{S}_t\boldsymbol{\Lambda}_{i\bullet}' + \sigma_i^2$ with $\boldsymbol{\Lambda}_{i\bullet}$ denoting the ith row of $\boldsymbol{\Lambda}$. This implies that the variance on ruling days typically exceeds the variance on non-event days due to the rulings affecting the endogenous variables. Another implication is that the covariance structure of $\boldsymbol{y}_t$ may adjust on ruling days due to the addition of the term $\exp(h_t)\beta_i\beta_j$. Our model implies that rulings translate into shifts in volatility on these days vis-á-vis non-event dates but also in directional movements in $\boldsymbol{y}_t$. The former claim can be easily analyzed by focusing on the commonalities due to the ruling factor. If movements in $\boldsymbol{y}_t$ are dominated by legal rulings on a specific ruling date, the share of variable-specific variance explained by the ruling factor becomes large (and approaches unity):

$$\frac{\exp(h_t)\beta_i^2}{\exp(h_t)\beta_i^2 + \boldsymbol{\Lambda}_{i\bullet}\boldsymbol{S}_t\boldsymbol{\Lambda}_{i\bullet}' + \sigma_i^2} \to 1, \tag{4}$$

with the variance components to be understood according to Eq. (3).

This variance decomposition can be used to demonstrate how our approach operates when there is another significant economic event on a particular ruling day. As long as we have an adequate number of ruling days without any other significant overlaps (as is the case in this paper), the economic interpretation of the ruling factor remains, and the identification strategy is intact. However, if many or all ruling days systematically coincide with distinct events or announcements, the ruling factor may turn out to be compromised. Our approach could still accurately estimate the ruling factor, since we include additional factors that are active when a ruling day occurs — and such a situation would likely result in a decrease of the commonalities explained by the ruling factor due to an increase of the term $\boldsymbol{\Lambda}_{i\bullet}\boldsymbol{S}_t\boldsymbol{\Lambda}_{i\bullet}'$. In essence, the always-active factors provide flexibility to account for (non-ruling related) multi-dimensional shocks. The ruling shocks themselves, however, are pinned down by the binary indicator d_t, and will typically reflect the relatively highest commonality on those specific days. This process is akin to narrative identification in structural VARs, which utilizes past historical events to pinpoint a specific shock of interest (see, e.g., Mertens and Ravn, 2014; Ludvigson et al., 2021).

Estimation is carried out using Bayesian techniques. We employ a Gibbs sampler to simulate from the joint posterior of the latent quantities and the parameters of the model. This algorithm, however, relies on knowing the non-event set $\mathcal{A}_{\text{NE}}$. The next sub-section describes how this set is formed within a unified econometric framework.

3.2. Identifying non-event windows

One important assumption underlying our modeling framework is that we need to assume the variance of $\boldsymbol{y}_t$ to differ between ruling and other days. This is a rather strong assumption. Several important regular events (such as press conferences after governing council meetings, or data releases by statistical agencies) or non-regular events (such as Mario Draghi's "Whatever it takes" speech) have pronounced effects on financial variables. Having such observations in the non-event/non-ruling set would certainly lead to a situation where the assumption of a higher variance on ruling dates is difficult to maintain. This is the reason why we meticulously avoided using non-ruling sets but instead differentiated between the ruling and non-event set. Moreover, if the number of event days is small to moderate (which is the case for court rulings), simply discriminating between ruling and non-ruling dates masks differences in variances. This makes identifying the factor (and/or the loadings) difficult and renders inference imprecise.

In the related literature it is common to avoid corrupting event windows of interest with other events by specifying tight windows surrounding announcements.[7] In what follows we propose a different route which does not rely on using intra-day data and/or

[7] See, for instance, Kuttner (2001), Bernanke and Kuttner (2005), Gürkaynak et al. (2005), Gertler and Karadi (2015) or Jarociński and Karadi (2020), for prominent examples.

precise timings of the rulings. Recall that $\mathcal{A}_{\text{NR}}$ denotes the set which includes all days except the ones which feature the publication of a ruling decision (i.e., for all $t \notin \mathcal{A}_{\text{R}}$), and define $\boldsymbol{\mu}_t = \boldsymbol{\Lambda}\boldsymbol{q}_t + \boldsymbol{B}\boldsymbol{x}_t$. We assume that $\boldsymbol{y}_t$ for $t \in \mathcal{A}_{\text{NR}}$ evolves according to:

$$\boldsymbol{y}_t \sim \sum_{j=1}^{J} w_j \mathcal{N}(\boldsymbol{\mu}_t, \vartheta_j^2\ \boldsymbol{\Omega}) \quad \text{for } t \in \mathcal{A}_{\text{NR}}, \tag{5}$$

with w_j denoting a weight with $w_j \geq 0$ and $\sum_{j=1}^{J} w_j = 1$. The scalar ϑ_j^2 denotes a component-specific variance with $\vartheta_1 < \cdots < \vartheta_J$ and $\boldsymbol{\Omega} = \text{diag}(\omega_1^2, \ldots, \omega_M^2)$ is a known diagonal matrix which serves to control for differences in the scaling of $\boldsymbol{y}_t$. This implies that $\boldsymbol{\Sigma} = \vartheta_1^2\boldsymbol{\Omega}$ in terms of Eq. (1), featuring solely the observations in the minimum variance regime, i.e., the set $\mathcal{A}_{\text{NE}}$.[8] In what follows, the diagonal elements of $\boldsymbol{\Omega}$ are set equal to the empirical variances of the different columns in the dependent variables $\boldsymbol{Y} = (\boldsymbol{y}_1, \ldots, \boldsymbol{y}_T)'$.

It is noteworthy that ϑ_j serves as a common scalar factor which drives the variances on non-ruling dates. To see this more clearly, Eq. (5) can be equivalently stated in terms of auxiliary variables δ_t which select the mixture component:

$$\boldsymbol{y}_t | \delta_t = j \sim \mathcal{N}(\boldsymbol{\mu}_t, \vartheta_j^2\ \boldsymbol{\Omega}), \tag{6}$$

which implies that if $\delta_t = j$, the idiosyncratic variances of $\boldsymbol{y}_t$ are proportional to ϑ_j^2. Under the ordering restriction that the variances are increasing, the set $\mathcal{A}_{\text{NE}} = \{t : \delta_t = 1\}$ collects all days where the idiosyncratic variances of Eq. (6) are minimal. Our key assumption is that on days in the non-event set, no substantial new information is driving movements in the variables collected in $\boldsymbol{y}_t$. The full time series of this idiosyncratic scale factor is obtained by assigning the respective component-variance t-by-t, that is, $\theta_t^2 = \sum_{j=1}^{J} \mathbb{I}(\delta_t = j)\vartheta_j^2$.

This procedure relies on an appropriate specification of J. Too small values of J could lead to regimes contaminated with other/non-ruling events (and thus feature a too high variance) whereas too large values of J could translate into regimes with too few observations in the non-event set. In our case, we use a Bayesian approach in the spirit of Malsiner-Walli et al. (2016) and estimate J by setting it to a large value and then use shrinkage to empty out irrelevant clusters.

For the remaining parameters of the model we choose adequate weakly informative priors. Our priors are generally chosen to introduce little information on the implied reactions of $\boldsymbol{y}_t$ to ruling shocks. In particular, we account for different scalings of the variables, and specifically use symmetric priors across all regions of the parameter space. Online Appendix B provides comprehensive details on specifics, the corresponding posterior distributions and the Markov chain Monte Carlo sampling algorithm. Empirical results showcasing the effects of choices about key tuning and prior (hyper-)parameters, alongside several additional robustness checks, are provided in Online Appendix C.

3.3. Local projections to compute dynamic responses

So far our model was presented in its static form. But our aim is also to comment on dynamics in response to the rulings. To this end, we specify the general model in Eq. (1) as a local projection for horizons $\ell = 0, 1, 2, \ldots, \mathscr{L}$. This enables us to interpret the ruling factor loadings as impulse responses at the desired horizon by relying on the idea that ruling shocks in time t trigger a larger time $t + \ell$ variance of $\boldsymbol{y}_{t+\ell}$ (with a variance shift governed by the loadings in $\boldsymbol{\beta}$).

Note that our dataset is comprised mostly of financial variables. We thus assume that their dynamic evolution is described well with a random walk (see, e.g., Swanson, 2021), and consequently consider the dependent variables $\Delta^\ell \boldsymbol{y}_t = \boldsymbol{y}_{t+\ell} - \boldsymbol{y}_{t-1}$ in differences[9]:

$$\Delta^\ell \boldsymbol{y}_t = d_t \cdot \boldsymbol{\beta} f_t + \boldsymbol{\mu}_t + \boldsymbol{\varepsilon}_t, \quad \boldsymbol{\varepsilon}_t \sim \sum_{j=1}^{J} w_j \mathcal{N}(\boldsymbol{0}, \vartheta_j^2\ \boldsymbol{\Omega}). \tag{7}$$

Notice that the timing of the ruling shocks is fixed via the binary indicator d_t. We avoid explicit labels here for readability, but note that all parameters are thus specific to the horizon, e.g., we have a sequence of $\boldsymbol{\beta}^{(\ell)}$'s or $f_t^{(\ell)}$'s. We consider a horizon of up to $\mathscr{L} = 60$ (corresponding to three months of business day weeks). This structure shares similarities with the local projection framework for difference-in-difference event studies proposed in Dube et al. (2023).

To provide further intuition about Eq. (7), consider the following trivial local projection: $\Delta^\ell \boldsymbol{y}_t = d_t\boldsymbol{\beta} + \tilde{\boldsymbol{\mu}}_t + \boldsymbol{\varepsilon}_t$, which arises from normalizing $f_t = 1$ and denoting some variant of fixed effects by $\tilde{\boldsymbol{\mu}}_t$. This represents a rather crude way of estimating of the average effect of a ruling on the dependent variables. It neglects the fact that rulings may be expansionary or contractionary, and that some of them may be more influential than others. In fact, one interpretation of f_t in our proposed framework is that it captures both the respective size and sign (from ruling date to ruling date) of the shock. Specifically, the impulse responses are derived as:

$$\frac{\partial \boldsymbol{y}_{t+\ell}}{\partial f_t^{(\ell)}} = \mathbb{E}\left(\boldsymbol{y}_{t+\ell} |\ f_t^{(\ell)} = d_t = 1, \bullet\right) - \mathbb{E}\left(\boldsymbol{y}_{t+\ell} |\ f_t^{(\ell)} = d_t = 0, \bullet\right) = \boldsymbol{\beta}^{(\ell)}.$$

In other words, our heteroskedasticity-based approach identifies $\boldsymbol{\beta}^{(\ell)}$ up to its sign, which in turn implies that we pin down the impulse response functions to a ruling shock. Specifying the model directly as a local projection has the advantage of taking all

[8] The ordering restriction on the variances solely refers to the labeling of the regimes, and is a technical restriction that does not affect posterior distributions (see Frühwirth-Schnatter, 2006, for details).

[9] The scaling matrix $\boldsymbol{\Omega}$ thus reflects the unconditional variance of $\Delta^\ell \boldsymbol{y}_t$, which is in fact the conditional variance of $\boldsymbol{y}_{t+\ell}$ with respect to its first lag. Our results are robust to alternative choices for $\boldsymbol{\Omega}$.

sources of uncertainty into account within a unified econometric framework. By contrast, a two-stage approach, estimating the factor first and using, e.g., its posterior median to compute dynamic responses with local projections, may yield incorrect inference and too narrow credible sets for parameter estimates. Alternatively, one could stack $\Delta^{\ell} \boldsymbol{y}_t$ for all ℓ and estimate a single factor. However, this choice would introduce strong restrictions (and comes with substantial computational complexity), since it effectively implies that a single factor explains movements in the conditional variance for all impulse response horizons. Moreover, it would require restrictive assumptions about the mixture model to pin down the non-event set.

In our baseline specification, we set $J = 30$ and consider $Q = 2$ always-active factors. Empirical results for restricted versions of the model and further robustness with respect to key tuning parameters are provided in Online Appendix C.

4. The dynamic effects of ruling shocks

In this section we present our empirical findings. The first sub-section discusses the dataset. Sections 4.2 and 4.3 present results originating from key model features. Section 4.4 displays and discusses the dynamic responses to ruling shocks and Section 4.5 uses the factor as an instrument in a monthly vector autoregression (VAR) to comment on longer-run economic dynamics.

4.1. Data overview

To understand the impact of ruling shocks, we study a set of EA, country-level and international financial variables observed at a daily frequency. Our sample covers the period from March 2012 to March 2021. Among other variables we consider in our discussion of dynamic responses to ruling shocks is the Composite Indicator of Systemic Stress (CISS) in the financial system (see Hollo et al., 2012). It is a summary measure aggregating indicators of financial stress in the banking sector, money, equity, bond and foreign exchange markets. To measure sentiments in sovereign debt markets we include the yields on 10-year government bonds of EA member states as spreads relative to Germany. The risk-free yield curve, based on the overnight index swap (OIS) rates at maturities from 1 month to 2 years reflects the market expectations regarding ECB policies. European stock market developments are represented by the EuroStoxx 50 index accompanied by headline country-indices. Corporate bond market developments are reflected by the Option-Adjusted Spreads (OAS) of high-yield, BBB and AA corporate bonds over the sovereign bonds of the same maturity, constructed by ICE BofAML. The information set further includes gold prices and the USD/EUR exchange rate. Since we are interested in how ruling shocks affect monetary policy uncertainty (MPU), we also include a corresponding measure. This daily news-based MPU index in the spirit of Baker et al. (2016) is taken from Azqueta-Gavaldón et al. (2023), and reflects MPU for the four largest euro area economies.

In terms of additional right-hand side variables, we always include an intercept term in $\boldsymbol{x}_t$. Moreover, we add several high-frequency controls that have been used to identify exogenous shocks in other studies. In particular, we add oil shocks and three distinct monetary policy shocks to control for any movements induced by these events. A list of variables and further information are provided in Online Appendix D. Series and country-coverage is governed by data availability on a daily frequency.

4.2. Ruling dates from an econometric perspective

Our model yields a posterior distribution over the factors f_t for different values of ℓ. These can be interpreted as the magnitudes and signs of an unexpected shock induced by legal events which are dated with the binary indicator d_t. Since the factor is not identified econometrically up to the sign, economic interpretation needs to be carried out in light of the estimated factor loadings in $\boldsymbol{\beta}$. This is what we do in Section 4.4. In this section, we focus on the quantitative properties of the factor.

Fig. 1 shows the posterior distribution of the contemporaneous ruling factor $f_t^{(0)}$ for all ruling dates (upper panel), and the posterior of the associated volatility $\exp(h_t)$ in the middle panel. Moreover, the figure also provides evidence on how much of the variation in all time series and in the different CISS series is explained by the different latent components of the model. These shares represent averages across the indicated sets of time series based on the variance decomposition described in Eq. (3).

The top panel of the figure reveals that the estimated factor substantially varies on the first and last ruling event. In between, we observe some variation but the magnitudes (in terms of posterior medians) are substantially smaller. In the middle panel, we find that volatility on ruling days generally fluctuates around the unconditional mean, and only the first and last events lead to somewhat larger increases in volatility. Notably, the consecutive rulings between the first and the last exhibit slightly lower than average variances, rather consistently. This suggests that financial markets may have learned that rulings predominantly favor the ECB's monetary policy actions. The lack of statistically significant differences from the unconditional mean on a ruling-by-ruling basis, however, points towards the notion that a homoskedastic ruling factor would likely be sufficient — we show in robustness checks collected in Online Appendix C that this is indeed the case.

The bottom panel shows the variance decomposition across ruling days. If we focus on all series in the panel we note that the amount of variation explained is considerable, hovering at about 25 percent for most dates. In some cases (such as during the period 2016 to 2018) we observe a modest decline in the explanatory power of the ruling factor. Notice, however, that these are averages across all series in our panel. When we focus on the CISS, as a particular variable where we would expect rulings to trigger substantial variation, we find that the ruling shock explains about 75 percent of the overall variation on average. We will revisit this aggregate assessment of the commonalities on a more granular and variable-by-variable basis alongside further details in the next sub-section.

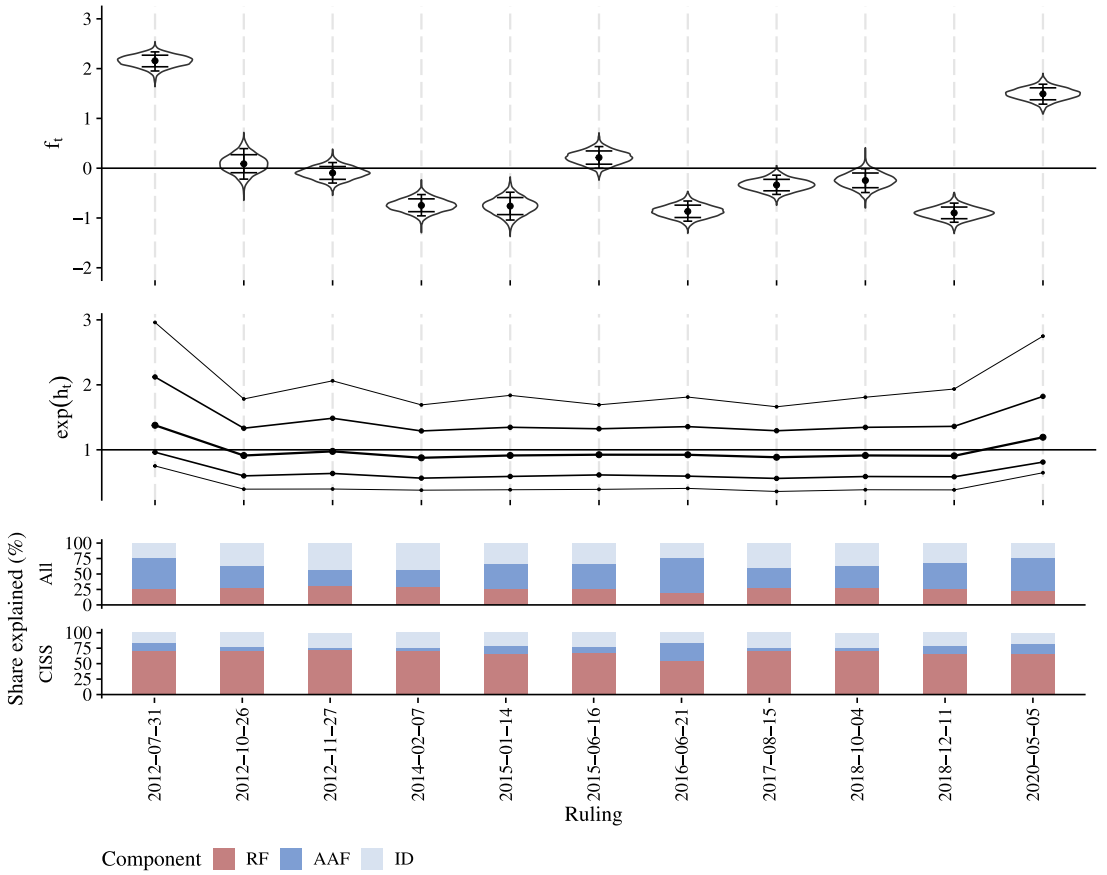

Fig. 1. Estimates of the ruling factor f_t and factor volatility $\exp(h_t)$ at ruling dates, $\ell = 0$. The indicated moments of the posteriors are the posterior median, the 68 and 90 percent posterior credible set. Ruling factor (RF), always-active factors (AAF), idiosyncratic noise (ID).

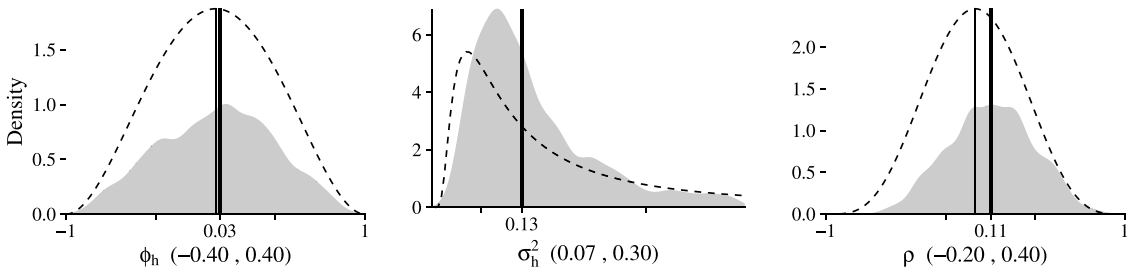

Fig. 2. State-equation parameters for the log-volatility process of the ruling factor, $\ell = 0$. The shaded area marks the posterior distribution, the dashed line indicates the respective prior. The thick black vertical line is the posterior median, the numbers in parentheses are the 16/84 posterior percentiles.

Assessing the parameter estimates of the log-volatility process of the ruling factors allows us to comment on the learning process of financial market participants. The prior (dashed lines) and posterior (gray shaded areas) of ϕ_h and σ_h^2 are depicted in Fig. 2. Before discussing the results, we note that the weakly informative priors on ϕ_h and σ_h^2 are centered on a homoskedastic specification without leverage and unit variance. When we consider the figure we find that the likelihood is only modestly informative about the specific values of ϕ_h, σ_h^2 and ρ, which is to be expected considering that these parameters are inferred from a latent variance process using $\mathrm{t} = 11$ ruling observations.

The posterior of the parameter ϕ_h indicates that the volatility estimates display almost no persistence. This provides empirical evidence in favor of the rulings being more or less independent over time. The posterior of the variance of the factor volatility slightly differs from the corresponding prior. This points towards at least some variation in volatilities across the different rulings and can be rationalized by noting that our set of event dates comprises the distinct stages of preliminary ruling procedures. Turning to the leverage parameter ρ, we find that the median is slightly shifted towards positive values, but posterior uncertainty is pervasive, rendering the parameter statistically insignificant.

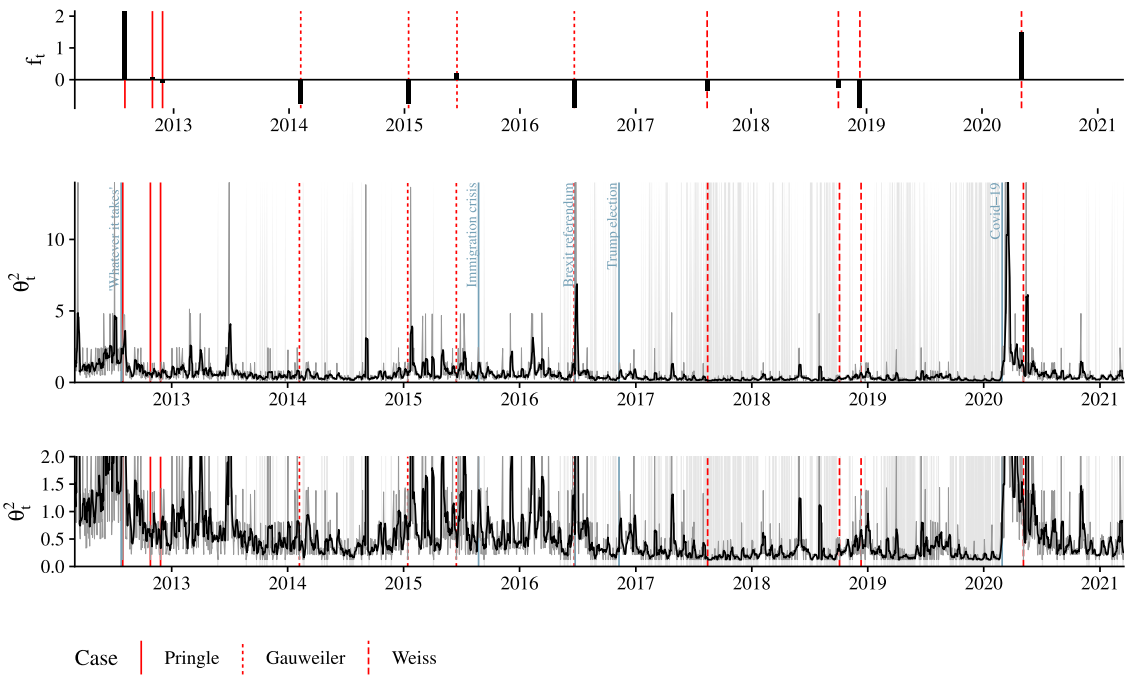

Fig. 3. Ruling factor, variance clusters and the non-event set alongside ruling dates for $\ell = 0$. The dark gray line is the period-by-period component variance, the black line is a five-day moving average of these variances. The gray shaded area indicates the posterior median of a specific period being featured in the minimal-variance regime (and thus, the non-event set). The red vertical lines are the ruling dates from Table 1, the light blue vertical lines are other known events that moved financial markets. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

4.3. Characterizing heteroskedastic features across event sets

Next we consider the estimated variance clusters as well as the non-event set. These are contrasted with the ruling dates and shown in Fig. 3 for $\ell = 0$. The gray shaded areas in the figure refer to the probability that a specific period is included in the non-event set. The red lines refer to the different ruling dates while the light blue lines mark important events that may have moved financial markets. For reference, the top panel shows the posterior median of the ruling factor over time. The middle panel shows the full range of the estimated error volatilities θ_t^2 over time (dark gray lines) alongside a five-day moving average (black line), whereas the lower panel restricts the range of the graph. For further results using a different number of mixture components, see Online Appendix C.

From this figure, it is clearly visible that contrasting ruling dates with non-ruling dates is insufficient. In this case, other important events and volatility spikes mask potential effects and may lead to imprecise inference on the effect of the ruling dates. Volatility spikes in several instances coincide with important events, and, if these were to dominate the estimates of the idiosyncratic volatilities, the explanatory power of the ruling factor would be greatly diminished (implying that there is little evidence for a shift in the conditional variance when comparing ruling with non-ruling days). Considering the periods where the mixture labels observations as non-event days reveals that this set is mostly comprised of observations from 2017 to 2019 and the first month of 2020. Several shorter periods with sequences of gray shaded bars appear before, but these relate to groups of days in various months without such a clear pattern than in the post-2017 period. When contrasted with recessions dated by the Euro Area Business Cycle Network (EABCN, eabcn.org) and general measures of volatility/uncertainty (see, e.g., policyuncertainty.com) for the EA, the regimes seem both economically and historically reasonable.

The combined variance estimates on the ruling days are higher than in the minimum variance regime, implying that the key identifying assumption of an increase in the conditional variance is fulfilled. We can drill deeper into this aspect by considering the commonalities due to the ruling factor of the different elements in $\boldsymbol{y}_t$. Recall that if our approach is successful, the commonalities stated in Eq. (4) should be sizable (i.e., close to one). By contrast, if we do not systematically control for other events on non-ruling days, the estimates of σ_i^2 plus the variance from the always-active factors will dominate β_i. We investigate this in Fig. 4. The figure depicts the share of variance explained through the ruling factor on ruling dates variable-by-variable, computed for each forecast horizon. White cells indicate values close to zero whereas red cells mark values approaching unity.

The figure reveals a great deal of white and blue colored cells (indicating commonalities induced by the rulings at or below 50 percent). At a very general level, there are differences in the amount of variation explained through the ruling factor for horizons up to 30 days (such as for the CISS indices, gold and the MPU). In this case, commonalities are large and always above 50 percent (and sometimes very close to 100 percent). For the CISS, these shares decline substantially for longer-horizon responses, suggesting that

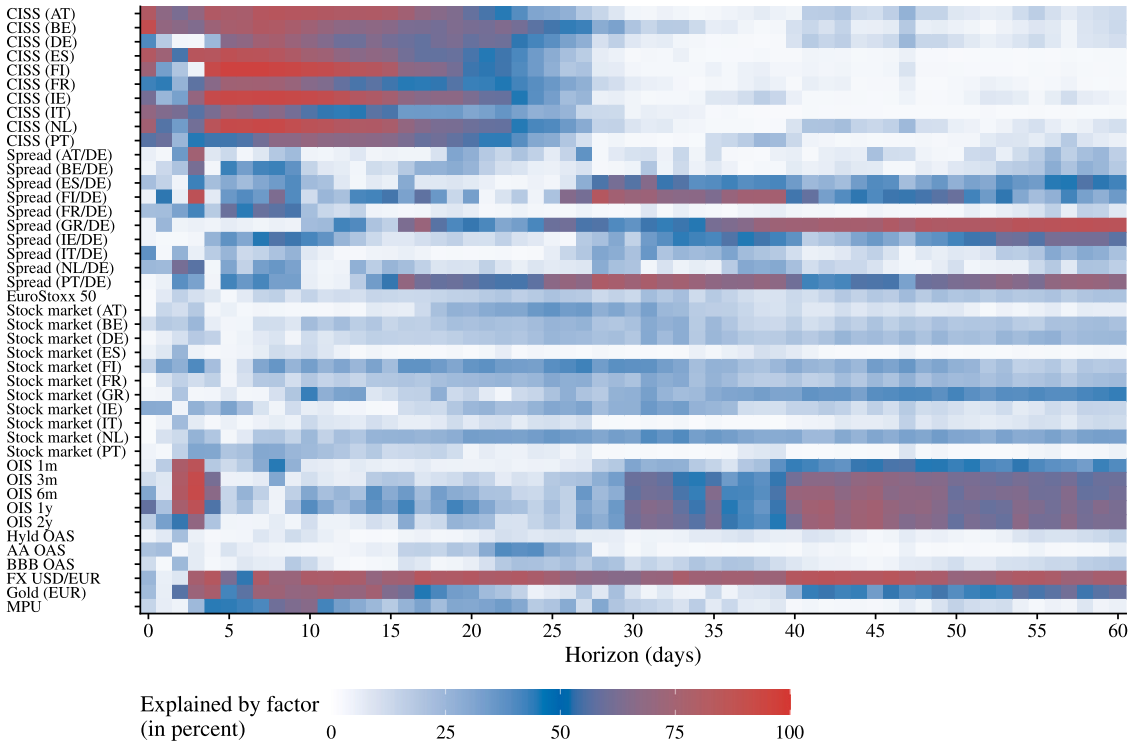

Fig. 4. Share of the variance which is explained by the ruling factor (average across event dates) per variable, as a percentage of the unconditional variance.

rulings seem to explain a substantial portion of variation in financial conditions in the short run and this effect dies out afterwards. The same is true for MPU where the ruling factor exhibits rather short-lived explanatory power, which is consistent with interpreting the ruling shocks partly as one of the multi-dimensional aspects that may cause variation in uncertainty around ECB monetary policy actions.

For some other variables (such as selected spreads or the OIS curve) we find the opposite pattern; commonalities start out at rather low levels and then approach 100 percent after around two weeks. This pattern is particularly pronounced for the spread between Greek and German government bond yields. The only variable which, across horizons, is almost entirely driven by the ruling factor on ruling days is the USD/EUR exchange rate. Finally, it is worth emphasizing that our mixture model is important for finding appropriate counterfactuals with a low variance and this leads to higher commonalities. In Online Appendix C we show that the explained share of variance due to rulings is substantially lower when disregarding distinct variance clusters for adequate comparisons between the ruling and non-event days.

In sum, for most variables considered, rulings play a role at least during certain impulse response horizons. The fact that in most cases explained variance shares are sizable, and appreciably above 30 percent, suggests that our identifying assumption of a transitory shift in variances when moving from non-event to ruling days is a reasonable one.

4.4. Impulse response functions from local projections

In this section we discuss the impulse responses to an expansionary ruling shock. With expansionary we mean a situation in which the courts rejected the challenges to the ECB's conduct.[10] We normalize the shock such that an expansionary ruling surprise decreases the average cross-country CISS by one standard deviation after 10 days.[11] Additional empirical results for other model parameters are collected in Online Appendix C.

[10] This definition of an expansionary ruling reflects the notion that the probability of further easing measures increased in light of the macroeconomic and financial circumstances during our sampling period (with the ECB clearly communicating – prior, during and after legal challenges – its intention to pursue an extremely accommodative monetary policy stance). Further, it is due to the ECB having a variety of discretionary policy options available thereby shifting expectations of market participants. By contrast, a contractionary ruling could at worst threaten the euro and imply the fragmentation of the euro area, dramatically increasing monetary policy uncertainty.

[11] The majority of rulings in our sample (as opposed to intermediate legal procedures) is in favor of policies enacted by the ECB, and thus, "expansionary" in nature. This is why we choose, for ease of interpretation, to normalize the shock producing our impulse response functions to be expansionary. It is worth noting that responses in our model framework are symmetric with respect to the sign and size of the shock.

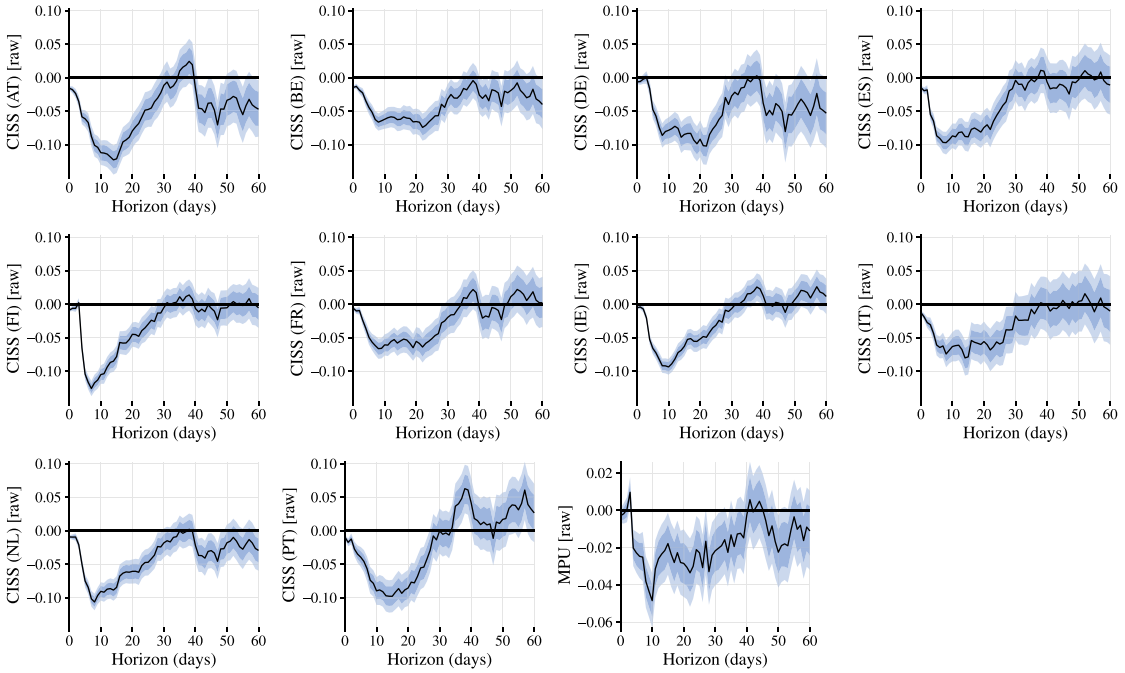

Fig. 5. Impulse responses for CISS measures across EA economies, and monetary policy uncertainty (MPU). The solid black line shows the posterior median alongside the 68 and 90 percent posterior credible sets in blue. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

Fig. 5 presents the responses to an expansionary ruling shock of financial conditions across EA member states, and the MPU index. Our economic intuition suggests that an expansionary ruling shock should decrease financial stress. This is caused by a lower likelihood of the ECB being forced to terminate its unconventional monetary policy programs that, among other goals, aim to improve stability in EA financial markets. Hence, we expect improvements in financial conditions. In addition, and relatedly, we expect MPU to decline since an expansionary ruling will likely contribute to the ECB being able to credibly follow its communicated policy path in the context of the announced unconventional policy measures, which were at stake at the courts.

Consistent with our normalization assumption and this prior reasoning, Fig. 5 reports that financial stress as measured by the CISS decreases over the first 30 days across all countries. When we focus on the shapes and magnitudes of the responses we observe many similarities across country-level CISS measures: immediate effects appear to be modest and after a few days, the CISS starts to decline, peaking at around −0.1 points 10 to 20 days after a ruling shock before increasing gradually. After around a month, the effects turn statistically insignificant for most countries. This brief discussion suggests the rulings to exhibit a broad-based effect on domestic financial markets and associated financial stability. We discuss the effects of ruling shocks on the real economy in the next sub-section in more detail, but note that an easing of financial stress relates to reduced macroeconomic downside-risk (see, e.g., Figueres and Jarociński, 2020). The shape of the MPU response is reminiscent of the CISS reactions, although it features somewhat more high-frequency movements. This, however, is consistent with the fact that uncertainty typically can be characterized by a low degree of persistence; and we find that uncertainty around ECB monetary policy actions declines after an expansionary ruling shock.

Next, we turn to the spread reactions of 10-year-government bonds relative to German treasuries in Fig. 6. We conjecture that expansionary rulings should translate into declining spreads between peripheral bonds and German Bunds. In fact, spreads between core/periphery countries in the EA are a crucial aspect considered by ECB policy makers. This is due to their impact on the smoothness of transmission of enacted policies to the member states. In fact, the avoidance of fragmentation in bond markets has been among the concerns underlying Mario Draghi's "Whatever it takes" speech, and also prominently featured in subsequent ECB communications as a potential threat to the singleness of monetary policy.

Our conjecture is confirmed empirically. Spreads for countries such as Spain, Ireland, Italy, Portugal and Greece (relative to Germany) decline substantially. Particularly strong reactions are found for Portugal and Greece, with peak declines of about 100 and 500 basis points (BPs), respectively. For the other peripheral countries, we also detect declines but these are much smaller in magnitude and peak at around 50 basis points. When we focus on core economies (such as the Netherlands, Belgium, France, Finland or Austria) we find only small reactions which are often negative and, in one instance, predominantly positive (for Finland). In sum, countries that have been heavily impacted by the EA sovereign debt crisis are also the ones that display the largest reductions in spreads whereas for the other countries they are more muted. From an economic perspective, this finding can be explained by noting that the expansionary ruling shock implies that the ECB may further intervene in secondary bond markets. Compared to a hypothetical no-policy intervention scenario, this yields shifts in relative demand for the safer core bonds relative to the riskier periphery ones — with the ECB indeed purchasing an enormous and economically significant amount of such securities.

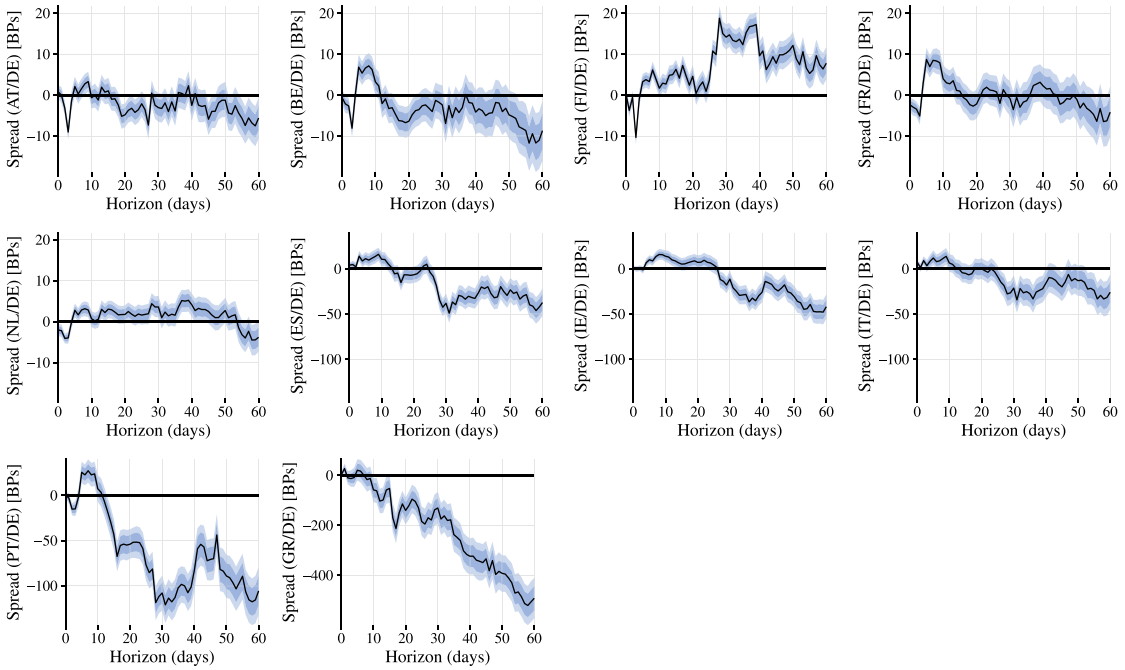

Fig. 6. Impulse responses for spreads of bond yields relative to Germany. The solid black line shows the posterior median alongside the 68 and 90 percent posterior credible sets in blue. *Caution*: *y*-axis scaling differs for core and periphery countries, and Greece. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

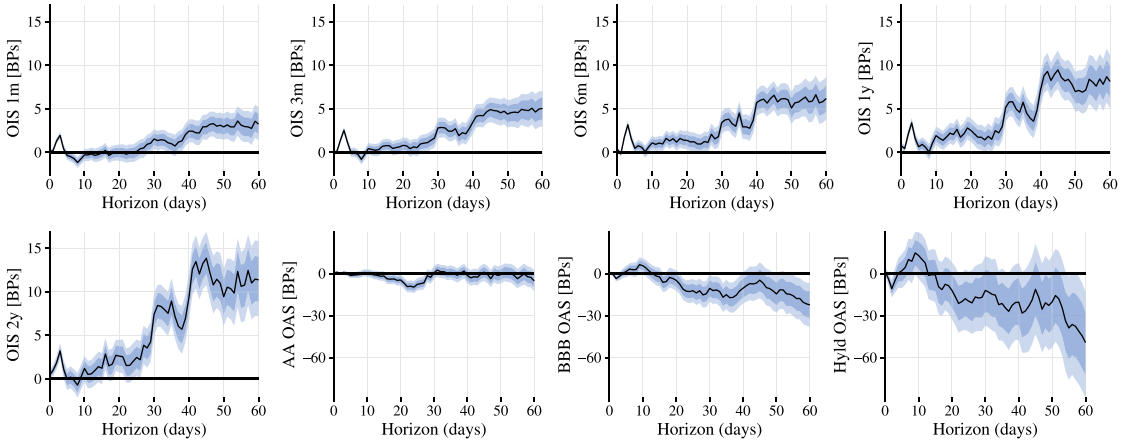

Fig. 7. Impulse responses for short-term interest rates and variables measuring financial conditions. The solid black line shows the posterior median alongside the 68 and 90 percent posterior credible sets in blue. *Caution*: *y*-axis scaling differs for interest rates and spreads. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

The previous discussion has shown that expansionary ruling shocks exert a meaningful effect on the longer end of the yield curve. In Fig. 7 we focus on the reactions at the short end of the OIS rates with maturities up to two years as well as corporate bond spread reactions. Turning to the OIS responses reveals only modest reactions for maturities up to six months. They become gradually more pronounced for maturities up to two years but with peak reactions of about 10 to 15 BPs their economic significance is modest. These subdued reactions can be explained by the fact that during our estimation period, interest rates at the short end of the yield curve have been close to zero and thus display relatively little variation. The moderately positive effects are again in line with our interpretation of the ruling shocks as affecting expectations and monetary policy uncertainty. In studies assessing monetary policy uncertainty on an aggregate level (see Husted et al., 2020), a stronger outlook due to decreased uncertainty may increase the probability of potential future rate hikes.

Focusing on the reactions of corporate bond markets across different rating categories (investment-grade and sub-investment-grade) reveals that expansionary ruling shocks trigger declines in spreads which increase with decreasing levels of credit quality.

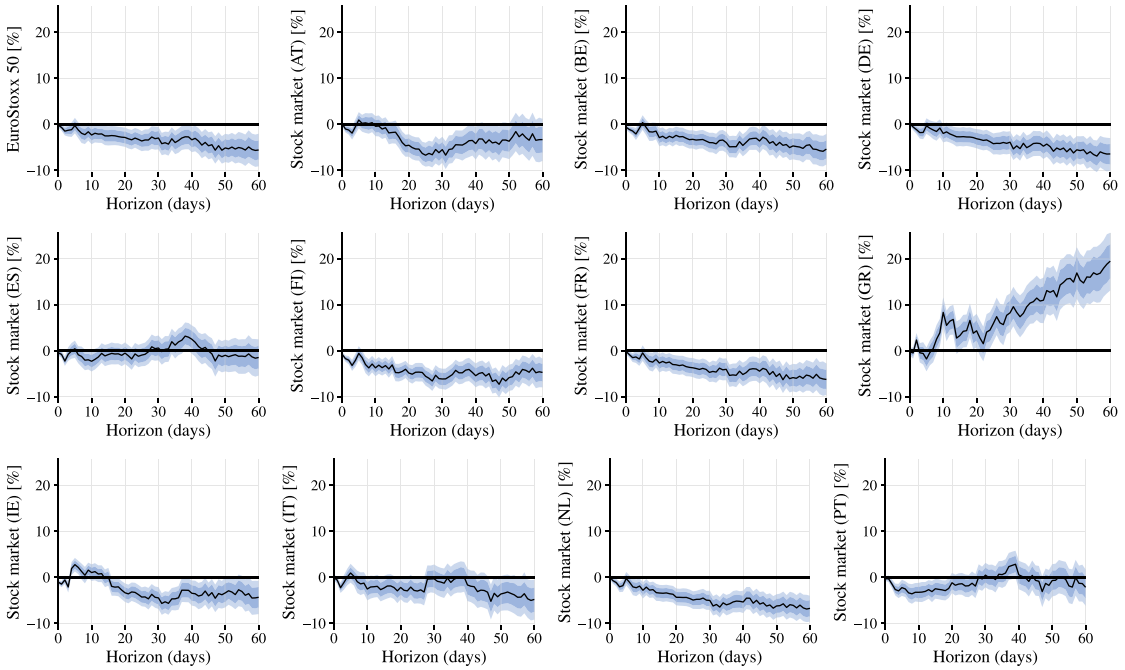

Fig. 8. Impulse responses for headline stock market indices across EA economies. The solid black line shows the posterior median alongside the 68 and 90 percent posterior credible sets in blue. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

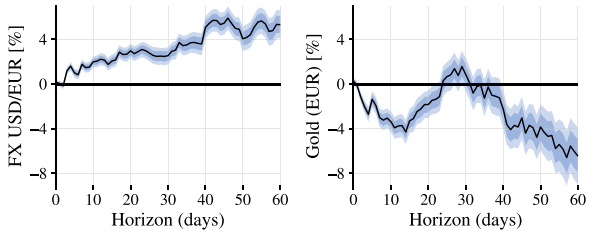

Fig. 9. Impulse responses for the exchange rate and gold prices. The solid black line shows the posterior median alongside the 68 and 90 percent posterior credible sets in blue. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

For AA-rated corporate bond spreads the model produces a decline of around 15 to 20 BPs after 20 to 30 days. When we consider BBB-rated bonds, these effects almost double in magnitude, peaking at around 30 BPs. Considering responses of sub-investment grade (high-yield) bond spreads suggests that these decrease in an even more pronounced manner, with peak declines of over 60 BPs.

These findings are in line with the general notion of looser financial conditions in response to an expansionary ruling shock, as measured by the CISS. However, they provide a more nuanced picture of the reactions of credit conditions. For instance, the larger declines of spreads for riskier categories suggest that comparative certainty about liquidity providing policies of the central bank (i.e., a reduced possibility of related programs being revoked) translates to a disproportionately reduced probability of default for firms that are likely to be financially constrained in some way. This in turn would reduce the required premia due to otherwise excessive default risk, particularly so for high-yield corporate bonds (see e.g., Christiano et al., 2014; Gilchrist et al., 2014).

Considering the reactions of stock markets across the EA (see Fig. 8) again reveals a heterogeneous pattern in terms of cross-country stock market reactions. For most countries, we find that stocks modestly decline. But these declines are in almost all cases very small and economically insignificant, particularly when viewed in light of historical fluctuations of these time series. The main exception to this pattern are stock market returns in Greece (and to a lesser degree several horizons in Ireland, Spain and Portugal). In the Greek case, we find strong increases over the impulse response horizon with peak effects of about 20 percent after three months. This result is not surprising and squares with our findings for spreads that countries which have been taking center stage in the EA sovereign debt crisis also exhibit more pronounced reactions to ruling shocks.

In the next step, we ask whether rulings trigger reactions in the USD/EUR exchange rate and the price of gold. These two variables are relevant for two simple reasons. First, in times of high (policy) uncertainty, the US dollar usually appreciates and the price of gold increases. This is due to the role of gold as a safe asset, and the US dollar being the currency in which assets deemed safe other

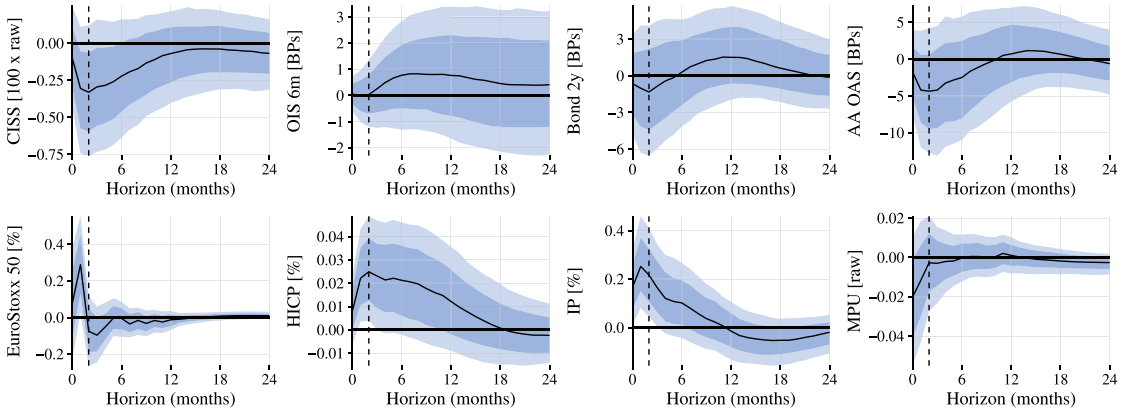

Fig. 10. Impulse responses from a VAR using monthly data. The solid black line shows the posterior median alongside the 68 and 90 percent posterior credible sets in blue. The vertical dashed line marks the 30-day, daily horizon. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

than gold are typically denominated. Second, the USD/EUR exchange rate reflects the credibility of the ECB, in the sense of whether it is in fact in a position to "do whatever it takes to preserve the euro". Hence, we have a prior intuition that an expansionary ruling shock should lower the price of gold and the dollar should depreciate. The responses are depicted in Fig. 9. Starting with the reaction of the USD/EUR exchange reveals that the euro persistently appreciates relative to the US dollar in response to an expansionary ruling shock. By contrast, gold prices decline. Both reactions suggest a narrative of decreasing macroeconomic uncertainty, but also uncertainty about the conduct of monetary policy, alongside an increase in the confidence of investors.

4.5. Investigating aggregate macroeconomic effects

In the previous sub-section, we showed that ruling shocks trigger pronounced high-frequency fluctuations in financial series. In this sub-section, our aim is to investigate whether rulings have effects on aggregate macroeconomic time series that are sampled at monthly frequencies. We answer this question through the lens of a medium-scale Bayesian VAR with 12 lags and data that spans from 2000:01 to 2021:03.[12]

The VAR includes some of the daily series from the previous sub-section to provide a direct nexus, but they are aggregated to the monthly frequency. This enables us to focus not only on reactions within the first three months after a shock, but also consider longer response horizons and gain intuition on how rulings would impact lower (i.e., monthly) frequency components in these variables. To also understand whether rulings have an effect on prices and output, we include the logarithm of the harmonized index of consumer prices (HICP), capturing headline inflation, and the logarithm of industrial production (IP). We rely on an internal instrument approach (see, e.g., Plagborg-Møller and Wolf, 2021), that is, the identified ruling shocks $f_t^{(0)}$ are included as the first endogenous variable in the VAR, followed by all others. The first column of a Cholesky decomposition of the reduced form covariance matrix of the VAR thus yields the contemporaneous impact of the ruling shocks on the endogenous variables. The dynamic VAR coefficients subsequently allow for straightforward computation of higher-order impulse response functions. These are displayed in Fig. 10.

At a general level, the reactions of the CISS and interest rate spreads are consistent with the results based on daily series. In particular, area-wide financial conditions improve. This reflects the sharp declines in daily country-specific CISS indices depicted in Fig. 5. The same consistency in results is observable for the MPU index, which exhibits its peak decline on impact and peters out quickly after a ruling shock.

Bond and interest rate responses are in most cases only borderline significant and the 16th and 84th credible intervals include zero. Interestingly, we find that the EuroStoxx 50 increases significantly for a brief period after the first month. This increase can be understood through two different channels: The asset price/wealth channel (see, e.g., Bernanke and Kuttner, 2005; Bjørnland and Leitemo, 2009) and the risk taking channel (see, e.g., Borio and Zhu, 2012). Since an expansionary ruling shock implies a decline in the probability that the ECB has to (unexpectedly) stop its unconventional monetary policy operations (thus allowing for a comparatively more accommodative stance), financial markets expect loose monetary policy in the short-run and an improved economic outlook, such that stock prices increase due to higher expected future dividends.

The remaining two variables left to discuss are the HICP and industrial production. Our working hypothesis is that expansionary ruling shocks, by not constraining the ECB's policy options, should stabilize output and exert inflationary effects — particularly in light of the unique challenges to monetary policy in the euro area, and our corresponding sampling period which mainly features

[12] We estimate the Bayesian VAR using a conjugate Minnesota prior, which avoids severe overfitting due to the size of the unrestricted parameter space. The prior is not set to be particularly tight, but appreciably reduces excessive high-frequency noise in the monthly impulse response functions.

inflation rates below the ECB's target. This is confirmed by our responses in Fig. 10. Both prices and output increase, with peak effects happening after six and two months for prices and industrial production, respectively.

In sum, the monthly VAR evidence broadly agrees with our higher frequency findings. Moreover, other recent VAR evidence shows that positive monetary policy uncertainty shocks lower output and prices, while spreads increase and financial conditions tighten (see Husted et al., 2020). Our (expansionary) ruling shocks are scaled conversely to coincide with a decrease in uncertainty about the ECB's monetary policy. This means that our responses mirror those reported in Husted et al. (2020). Thus, the observed patterns in response to such shocks are consistent with the interpretation that ruling shocks are among the several components that may induce exogenous declines in uncertainty about monetary policy.

5. Conclusions

Legal challenges to the conduct of the ECB's monetary policy have the potential to seriously impact policy making by reducing the variety of tools at the disposal of the central bank. This implies a reduction in flexibility and may also limit the scope of credible future discretionary policy options. In this paper, we analyze three lawsuits surrounding the ECB's unconventional monetary policy programs.

The number of corresponding ruling days is small and information on the precise timing of the announcements of the court decisions is not available. Standard tools in the literature have a difficult time in such a situation since the presumed non-ruling days may be contaminated with unknown events which potentially jeopardize identification of the shock of interest. To cope with this, we develop an econometric model which discriminates between ruling and non-event days. With non-event days we refer to days with minimal variance that can be regarded as days free of significant economic or political events.

We use this framework to analyze the effects of an expansionary ruling shock. The corresponding impulse responses are consistent with our prior expectations. Financial stress measures decrease and bond spreads of peripheral countries relative to German Bunds decline; the euro appreciates vis-á-vis the dollar and the price of gold declines. Using the identified shock in a standard VAR shows that European stock markets, inflation and industrial production increase while financial stress eases and bond spreads decline. The VAR-based results give rise to partly interpreting the ruling shocks as one of the drivers of monetary policy uncertainty.

Data availability

Data is available at https://ars.els-cdn.com/content/image/1-s2.0-S0014292124000096-mmc1.zip.

Appendix A. Supplementary material

Supplementary material related to this article can be found online at https://doi.org/10.1016/j.euroecorev.2024.104680.

References

Azqueta-Gavaldón, A., Hirschbühl, D., Onorante, L., Saiz, L., 2023. Sources of economic policy uncertainty in the euro area. Eur. Econ. Rev. 152, 104373.
Baker, S.R., Bloom, N., Davis, S.J., 2016. Measuring economic policy uncertainty. Q. J. Econ. 131 (4), 1593–1636.
Bauer, M.D., Lakdawala, A., Mueller, P., 2022. Market-based monetary policy uncertainty. Econ. J. 132 (644), 1290–1308.
Bernanke, B.S., Kuttner, K.N., 2005. What explains the stock market's reaction to federal reserve policy? J. Finance 60 (3), 1221–1257.
Bertsche, D., Braun, R., 2022. Identification of structural vector autoregressions by stochastic volatility. J. Bus. Econom. Statist. 40 (1), 328–341.
Bjørnland, H.C., Leitemo, K., 2009. Identifying the interdependence between US monetary policy and the stock market. J. Monetary Econ. 56 (2), 275–282.
Blinder, A.S., Ehrmann, M., Fratzscher, M., De Haan, J., Jansen, D.-J., 2008. Central bank communication and monetary policy: A survey of theory and evidence. J. Econ. Lit. 46 (4), 910–945.
Borio, C., Zhu, H., 2012. Capital regulation, risk-taking and monetary policy: A missing link in the transmission mechanism? J. Financ. Stabil. 8 (4), 236–251.
Brunnermeier, M., Palia, D., Sastry, K.A., Sims, C.A., 2021. Feedbacks: Financial markets and economic activity. Amer. Econ. Rev. 111 (6), 1845–1879.
Christiano, L.J., Motto, R., Rostagno, M., 2014. Risk shocks. Amer. Econ. Rev. 104 (1), 27–65.
Cukierman, A., Meltzer, A.H., 1986. A theory of ambiguity, credibility, and inflation under discretion and asymmetric information. Econometrica 1099–1128.
Cukierman, A., Web, S.B., Neyapti, B., 1992. Measuring the independence of central banks and its effect on policy outcomes. World Bank Econ. Rev. 6 (3), 353–398.
Dube, A., Girardi, D., Jorda, O., Taylor, A.M., 2023. A local projections approach to difference-in-differences event studies. NBER Working Paper. 31184.
Figueres, J.M., Jarociński, M., 2020. Vulnerable growth in the euro area: Measuring the financial conditions. Econom. Lett. 191, 109126.
Frühwirth-Schnatter, S., 2006. Finite Mixture and Markov Switching Models. Springer.
Gertler, M., Karadi, P., 2015. Monetary policy surprises, credit costs, and economic activity. Am. Econ. J.: Macroecon. 7 (1), 44–76.
Gilchrist, S., Sim, J.W., Zakrajšek, E., 2014. Uncertainty, financial frictions, and investment dynamics. NBER Working Paper 20038.
Griller, S., 2021. Analytical report on the legal background of member states' positions towards economic and fiscal integration. In: Griller, S., Lentsch, E. (Eds.), EMU Integration and Member States' Constitutions. Hart Publishing, Oxford.
Gürkaynak, R.S., Kisacikoğlu, B., Wright, J.H., 2020. Missing events in event studies: Identifying the effects of partially measured news surprises. Amer. Econ. Rev. 110 (12), 3871–3912.
Gürkaynak, R.S., Sack, B., Swanson, E., 2005. The sensitivity of long-term interest rates to economic news: Evidence and implications for macroeconomic models. Amer. Econ. Rev. 95 (1), 425–436.
Hartmann, P., Smets, F., 2018. The European central bank's monetary policy during its first 20 years. Brook. Pap. Econ. Act. 2018 (2), 1–146.
Hollo, D., Kremer, M., Lo Duca, M., 2012. CISS—A composite indicator of systemic stress in the financial system. ECB Working Paper Series 1426.
Husted, L., Rogers, J., Sun, B., 2020. Monetary policy uncertainty. J. Monetary Econ. 115, 20–36.
Jarociński, M., Karadi, P., 2020. Deconstructing monetary policy surprises—the role of information shocks. Am. Econ. J.: Macroecon. 12 (2), 1–43.
Korobilis, D., 2022. A new algorithm for structural restrictions in Bayesian vector autoregressions. Eur. Econ. Rev. 148, 104241.
Kuttner, K.N., 2001. Monetary policy surprises and interest rates: Evidence from the fed funds futures market. J. Monetary Econ. 47 (3), 523–544.

Lewis, D.J., 2021. Identifying shocks via time-varying volatility. Rev. Econom. Stud. 88 (6), 3086–3124.
Li, M., Scharth, M., 2022. Leverage, asymmetry, and heavy tails in the high-dimensional factor stochastic volatility model. J. Bus. Econom. Statist. 40 (1), 285–301.
Ludvigson, S.C., Ma, S., Ng, S., 2021. Uncertainty and business cycles: Exogenous impulse or endogenous response? Am. Econ. J.: Macroecon. 13 (4), 369–410.
Malsiner-Walli, G., Frühwirth-Schnatter, S., Grün, B., 2016. Model-based clustering based on sparse finite Gaussian mixtures. Stat. Comput. 26 (1–2), 303–324.
Mertens, K., Ravn, M.O., 2014. A reconciliation of SVAR and narrative estimates of tax multipliers. J. Monetary Econ. 68, S1–S19.
Plagborg-Møller, M., Wolf, C.K., 2021. Local projections and VARs estimate the same impulse responses. Econometrica 89 (2), 955–980.
Rigobon, R., 2003. Identification through heteroskedasticity. Rev. Econ. Stat. 85 (4), 777–792.
Rigobon, R., Sack, B., 2004. The impact of monetary policy on asset prices. J. Monetary Econ. 51 (8), 1553–1575.
Rigobon, R., Sack, B., 2008. Noisy macroeconomic announcements, monetary policy, and asset prices. In: Campbell, J.Y. (Ed.), Asset Prices and Monetary Policy. University of Chicago Press, pp. 335–370.
Sentana, E., Fiorentini, G., 2001. Identification, estimation and testing of conditionally heteroskedastic factor models. J. Econometrics 102 (2), 143–164.
Swanson, E.T., 2021. Measuring the effects of federal reserve forward guidance and asset purchases on financial markets. J. Monetary Econ. 118, 32–53.