INSTITUTO DE ECONOMÍA

DOCUMENTO de TRABAJO

559

2021

The Real Effects of Financial Uncertainty Shocks: A Daily Identification Approach

Piergiorgio Alessandri, Andrea Gazzani y Alejandro Vicondoa.

www.economia.puc.cl • ISSN (edición impresa) 0716-7334 • ISSN (edición electrónica) 0717-7593

The Real Effects of Financial Uncertainty Shocks: A Daily Identification Approach *

Piergiorgio Alessandri† Andrea Gazzani‡ Alejandro Vicondoa§

January 14, 2021

Abstract

Isolating financial uncertainty shocks is difficult because financial markets rapidly price changes in several economic fundamentals. To bypass this difficulty, we identify uncertainty shocks using daily data and use their monthly averages as an instrument in a VAR. We show that this novel approach is theoretically appealing and has dramatic implications for leading empirical studies on financial uncertainty. Daily interactions between equity returns, bond spreads and expected volatility cause previous identification schemes to fail at the monthly frequency. Once these interactions are explicitly modeled, the impact of uncertainty shocks on output and inflation is significant and similar across specifications.

JEL classification: C32; C36; E32

Keywords: uncertainty shocks; financial shocks; structural vector autoregression; high-frequency identification; external instruments

*We are grateful to Evi Pappa, Fabio Canova and Juan Dolado for fruitful discussions and suggestions. We also thank Ambrogio Cesa-Bianchi, Paul Beaudry, Danilo Cascaldi-Garcia, Maarten Dossche, Luca Gambetti, Aeimit Lakdawala, Peter Hansen, Matteo Iacoviello, Francesca Loria, Riccardo Jack Lucchetti, Kurt Lunsford, Michele Piffer, Morten Ravn, Juan Rubio-Ramirez, Andreas Tryphonides, Srečko Zimic, and many seminar participants for helpful comments and suggestions. The paper supersedes an earlier paper entitled "Bridge Proxy-SVAR: Estimating the Macroeconomic Effects of Shocks Identified at High Frequency". The views expressed in the paper are those of the authors only and do not involve the responsibility of the Bank of Italy.

†Bank of Italy. Email: piergiorgio.alessandrii@bancaditalia.it.

‡Bank of Italy. Email: andreagiovanni.gazzani@bancaditalia.it.

§Instituto de Economía, Pontificia Universidad Católica de Chile. Email: Alejandro.Vicondoa@uc.cl

1 Introduction

Isolating the role of uncertainty in the business cycle is challenging because spikes in uncertainty often coincide with a deterioration of the macroeconomic outlook. This problem is particularly acute in the case of financial uncertainty. High stock market volatility, rising credit spreads and economic slowdowns often materialize together, but researchers disagree on the interpretation of this coincidence. Ludvigson et al. (2018) argue that financial uncertainty is an important source of fluctuations for the US economy, while Berger et al. (2019) conclude that uncertainty is endogenous to (and hence inconsequential for) the evolution of the real economy.

We examine this puzzle using a novel empirical strategy specifically designed to disentangle the interactions between uncertainty and asset prices. The strategy consists of identifying the shock in a daily VAR, averaging the daily shock series to the monthly frequency, and then using this average as an instrument in a monthly VAR model. We demonstrate that this procedure delivers consistent estimates of the low-frequency impact of the shock in a broad range of models: intuitively, the linearity of a VAR guarantees that the causal effect of a sum of (within-month) shocks coincides with the sum of their individual (daily) effects. Daily data serve two crucial objectives, allowing us to (i) control accurately for the information set available to market participants at any given point in time, and (ii) leave the relation between financial and real variables unrestricted at the monthly frequency. In our framework the validity of the restrictions, especially those imposed between financial variables, can be tested on monthly data rather than assumed ex-ante. After discussing the theoretical properties of our procedure and examining its accuracy through Montecarlo exercises, we use it to revisit two influential studies on the role of financial uncertainty shocks by Caldara et al. (2016) (CFGZ) and Berger et al. (2019) (BDG). In both cases we keep the models unchanged and only shift the identification restrictions from the monthly to the daily frequency, preserving the theoretical strengths of the identification schemes while accounting for the possibility that portfolio decisions are taken daily rather than monthly. This frequency shift has dramatic implications for the results.

CFGZ identify uncertainty and financial shocks in a VAR as the innovations that maximize the responses of the VXO and the Excess Bond Premium over a six-month horizon. The identification is sequential, as one of the shocks must be identified first and the other one is subject to an orthogonality condition with respect to it. CGFZ find that the ordering is crucial: uncertainty shocks have a large impact on output and inflation if they are identified first, but become irrelevant if they are identified after (i.e. conditioning on) financial shocks. Our analysis reveals that this ambiguity is caused by the strong interactions between corporate bond spreads and the VXO volatility index at the daily frequency. Since spreads and volatility respond to both disturbances within a month, any identification achieved through monthly timing assumptions is inevitably biased and at odds with the data. Our estimates show that VXO and bond spreads increase rapidly in response to a financial tightening and to a rise in uncertainty, as one would expect on theoretical grounds. More importantly, provided that their daily interactions are adequately captured, uncertainty shocks have a significant effect on industrial production, employment and inflation irrespective of the identification order.

A similar conclusion emerges in the BDG setup. BDG draw an important distinction between the realized volatility of the stock market (a broad proxy of changes in economic fundamentals) and its option-implied volatility (which reflects changes in volatility expectations, i.e. forward-looking uncertainty). Uncertainty shocks are then identified as the linear combination of innovations that maximizes the forecast error variance of implied volatility over two years while having no contemporaneous impact on realized volatility. The rationale for this restriction is the fact the large negative news about macroeconomic fundamentals are captured by realized volatility. Our analysis shows that, if the BDG restrictions are imposed on daily data, uncertainty shocks significantly impact the realized volatility of the market within a given month. As in the CFGZ case, the monthly restrictions are invalid. Once the within-month feedback between expected and realized volatility is taken into account, an exogenous rise in uncertainty also causes a significant contraction in industrial production. Interestingly, the responses are quantitatively similar to those obtained in the CFGZ model, suggesting that – net of temporal aggregation issues – the two identification schemes isolate the same structural disturbance.

The paper makes three contributions to the literature. The first one is to provide new evidence on the broad economic implications of an exogenous shift in financial uncertainty. The second one is to demonstrate that using financial data in a VAR is a double-edged sword: asset prices provide valuable information but do not lend themselves easily to structural identification in monthly or quarterly models. The overlap between exogenous shocks and endogenous responses uncovered in our study of financial uncertainty may easily arise in other contexts. The third contribution is to provide a flexible and robust tool to identify VARs using high-frequency data, including those obtained from financial markets. The method we propose is conceptually and computationally simple, and the shock(s) obtained from high-frequency data can be employed as external or internal instruments, included in an exogenous block of the VAR, or used in a local projection setup.[1] Hence, the procedure can be used in a broad range of cases where identification restrictions imposed at "low" frequencies could bias the results.[2]

Related Literature. Uncertainty has attracted considerable attention in the business cycle literature from both a theoretical and an empirical perspective.[3] After Bloom (2009), researchers often measure aggregate uncertainty using the VIX or VXO implied volatility indices. However, different views have emerged on the identification of financial uncertainty shocks and on the implications of such shocks for the real economy (CFGZ, Carriero et al., 2018, Ludvigson et al., 2018, BDG). Our paper contributes to this debate by reassessing CFGZ and BDG through a new high-frequency identification approach, showing that temporal aggregation can completely cloud causality in this context. Few other works have employed daily or weekly data to study uncertainty shocks. Ferrara and Guerin (2018) use a mixed-frequency VAR where the VXO is included as a weekly average and find that the results are similar to those obtained from a standard monthly VAR. However, identification relies on a recursive scheme that is unlikely to

[1] The use of external information in VARs is discussed in Beaudry and Saito (1998); Faust et al. (2004), Stock and Watson (2012), Mertens and Ravn (2013), Noh (2018), Miranda Agrippino and Ricco (2018), and Paul (forth). See Plagborg-Møller and Wolf (2019b) and Herbst and Johannsen (2020) for a discussion of the relation between VARs and local projections.

[2] For a general discussion of temporal aggregation biases see Sims (1971), Christiano and Eichenbaum (1987), Marcet (1991), Hendry (1992), and Swanson and Granger (1997). Marcellino (1999) and Foroni and Marcellino (2016) show that temporal aggregation can heavily distort impulse-response functions and forecast error variance decompositions in SVAR models.

[3] See e.g. Fernandez-Villaverde et al. (2011), Jurado et al. (2015), Baker et al. (2016), Basu and Bundick (2017), Arellano et al. (2019). Extensive reviews of the literature can be found in Bloom (2009) and Fernandez-Villaverde and Guerron-Quintana (2020).

capture the feedbacks between uncertainty and macroeconomic conditions.[4] Our identification approach exploits daily rather weekly data and it can accommodate restrictions that are not subject to those limitations. Piffer and Podstawski (2018) identify uncertainty shocks as variations in gold prices around specific events and then use them as external instruments in a VAR, concluding that uncertainty shocks are a major driver of the business cycle. Our method represents a robust and flexible alternative to event-based identification strategies. Furthermore, our analysis delivers more conservative estimates of the impact of financial uncertainty shocks than those of Piffer and Podstawski (2018). The reason is that, by imposing the CFGZ or BDG restrictions in daily VAR models that include a range of financial indicators, we can better isolate uncertainty shocks from financial or real disturbances that may simultaneously hit the economy.[5]

Estimation of dynamic effects based on external information has recently spread in the empirical macroeconomic literature (Beaudry and Saito 1998; Faust et al. 2004, Stock and Watson, 2012, Mertens and Ravn, 2013). This paper makes two contributions to this research. First, it provides a novel approach to exploiting high-frequency information for identification purposes. Second, it shows how to employ proxies that are available at higher frequency than the endogenous variables of interest. In particular, we show that, as long as the data-generating process is a VAR, averaging the high-frequency proxy to a lower frequency delivers consistent estimates of the responses in a broad range of empirical setups and model specifications.[6] Our method can be promptly applied to cases where daily VAR models are used to isolate monetary policy shocks (Wright, 2012) or shocks to expected growth and risk premia (Cieslak and Pang, 2020).

[4]See e.g. Stock and Watson (2012), Baker and Bloom (2013), Baker et al. (2016), Cascaldi-Garcia and Galvao (2016) and Ludvigson et al. (2018). Another limitation of stacked mixed-frequency VARs is that they do not yield unique IRF estimates, and hence they cannot be directly compared to monthly or quarterly VAR models.

[5]In Alessandri et al. (2020) we take this argument one step further, exploiting the non-Gaussianity of the daily series to identify structural shocks that are statistically independent (and not just orthogonal) from one another. In this paper we deliberately stick to the restrictions proposed by CFGZ and BDG in order to isolate the role of daily dynamics in the context of established, well-known identification strategies.

[6]Chudik and Georgiadis (2019) propose an alternative framework, treating the shocks as observed and estimating the effect of an individual high-frequency shock on a set of lower-frequency variables. In our approach, by contrast, the high-frequency shocks are used as instruments and the aim is to recover the cumulated effects of all the shocks that occur within a low-frequency period: this allows us to compare our findings to those obtained from standard VAR models based on monthly or quarterly data.

Outline. The remainder of this paper is organized as follows. Section 2 describes the influence of temporal aggregation on the identification of VAR models, introducing our method and discussing its theoretical properties and its performance in Monte Carlo simulations. In Section 3 we apply our method to the estimation of financial uncertainty shocks in the CFGZ and BDG setups. Section 4 concludes.

2 High-Frequency Identification of Low-Frequency Causal Effects

We propose a new strategy to exploit high-frequency data (obtained for instance from financial markets) in order to estimate the effects of structural shocks on variables that are only available at lower frequencies (such as price and economic activity indicators). Our proposal is motivated by the consideration that the use of high-frequency observations can significantly improve the structural identification of a VAR model. Macroeconomists routinely use monthly or quarterly series to examine the implications of various structural shocks. Yet firms, households and investors may take their decisions at higher frequencies than a month: this implies that in a monthly or quarterly dataset exogenous shocks and endogenous responses might be inextricably mixed, and that even theoretically sound identification strategies may fail to isolate the former from the latter. The empirical analysis in Section 3 shows that this is indeed a first-order problem in the case of financial uncertainty shocks. However, the problem we highlight is entirely general and the simple method proposed in this paper can be easily exploited in other contexts. After summarizing the logic of our approach below, we develop an analytical example in Section 2.1, provide results on the general validity of the method in a VAR framework in Section 2.2, and document its performance in a Montecarlo study in Section 2.3.

The method involves three steps:

I) Identification of the shocks on high-frequency data. The first step consists of recovering the structural shock(s) of interest by applying an appropriate identification strategy to a VAR model estimated on 'high-frequency' (HF) data. We deliberately use the HF label in a loose sense: the data can be weekly, daily (as in our empirical applications) or intradaily, depending on the task at hand. All that matters is that this information is available at a frequency for which the identification restrictions are reasonable; in many cases this is higher than one month or one quarter, which is the sampling frequency of the macroeconomic aggregates on which the impact of the shock(s) must ultimately be quantified. The specification of the VAR must of course allow the identification of the shock of interest; in particular, the variables need to capture an information set that is broad enough to insure that the model is informationally sufficient. Provided this condition is met, the method can rely on any identification scheme. In the empirical analysis of Section 3 we use it in combination with the identification strategies proposed by CFGZ and BDG in two influential papers on the real effects of financial uncertainty.

II) Temporal aggregation of the shocks. The second step consists of computing low-frequency (LF, e.g. monhtly or quarterly) averages of the high-frequency shocks obtained in (I). One strength of our approach is its robustness: if the underlying HF data generating process is a VAR, then averaging is the correct temporal aggregation filter for the shocks irrespective of the gap between the two frequencies and the type of variables employed in the analysis (prices, flows, stocks, etc.; see subsections 2.1–2.3). This result on the optimal aggregation filter is also relevant for VARs identified using external information (i.e. proxies), which have recently become very popular in the applied macroeconomics literature (see e.g. Gertler and Karadi, 2015 and Ramey and Zubairy, 2018). Our analysis demonstrates that there are indeed good reasons to stick to simple within-period averaging rather than using alternative filters, such as moving averages, large shocks, or shocks that occur at the beginning or the end of each month/quarter (see Section 2.2). The averaged LF shocks should again pass the informational sufficiency test (e.g. Forni and Gambetti (2014)). In our applications we systematically test whether they are orthogonal to past information obtained from large auxiliary data sets.

III) Estimation of the causal effects of the shocks on low-frequency variables. The third and last step consists of employing the series of LF shocks obtained in (II) as a proxy in a VAR specified at low-frequency that contains the endogenous variables of interest. The LF shocks can be treated as 'external' or 'internal' instruments for the LF-VAR; the estimation of the IRFs, like the identification in step (I), can thus be carried out in a number of alternative ways. The correct way to use external information to estimate IRFs and forecast error variance decomposition has been extensively analyzed in the recent literature. The shock captured by the proxy is invertible if and only if the proxy does not Granger-cause the residuals of the LF-VAR. If the test is not passed, then the inference based on the Proxy-SVAR is not valid but the relative IRFs can be still estimated by including the proxy and its lags as an exogenous variable in the VAR (Paul, forth). Alternatively, the proxy can be used as an internal instrument in the VAR by including it as an endogenous variable in the VAR and computing the IRFs by ordering it first in a Cholesky decomposition (see, for example, Plagborg-Møller and Wolf, 2019a; Noh, 2018; and Miranda Agrippino and Ricco, 2018). The latter approach constitutes a parsimonious equivalent of the popular local projection instrumental variable approach (LP-IV) with controls. Notice that, under the (testable) assumption that the lagged endogenous variables of the LF-VAR are orthogonal to the proxy, the impact effect estimated using the proxy as an external instrument, an internal instrument or an exogenous variable in the VAR must coincide. Therefore, the partial invertibility of the shock and the strategy adopted to compute the IRFs only matter for model's dynamics over longer horizons. Finally, as discussed in Wooldridge (2010), generated instruments do not suffer the inference problem of associated regressors.

2.1 Illustrative Cases

Univariate Case. Some of the implications of temporal aggregation can be easily illustrated using a scalar AR(1) process. Let y follow an AR(1) process with white noise innovations ε at the frequency $t = 1, 2, ..., T$ (e.g. daily):

$$y_t = \alpha y_{t-1} + \varepsilon_t \tag{1}$$

If the variable is observed only at the lower frequency $\tau = 2, 4, ..., T$, i.e. half of the original frequency t, the process becomes:

$$y_\tau = \alpha^2 y_{\tau-1} + u_\tau, \tag{2}$$

where $y_\tau = y_t$, $y_{\tau-1} = y_{t-2}$ and the residual is:

$$u_\tau := y_\tau - E_{\tau-1}(y_\tau | \mathscr{F}_{\tau-1}) = \alpha\varepsilon_{t-1} + \varepsilon_t \tag{3}$$

Consider estimating the contemporaneous response of y_τ to the underlying structural shock, defined as its cumulative response to the two shocks ε_t and ε_{t-1} that are part of the time-τ residual. This response is the difference between two conditional forecasts calculated setting both shocks to 1 and to 0:

$$\Theta_0 \quad := \quad \mathbb{E}_{\tau-1}\left[y_\tau | \varepsilon_t = 1, \varepsilon_{t-1} = 1\right] - \mathbb{E}_{\tau-1}\left[y_\tau | \varepsilon_t = 0, \varepsilon_{t-1} = 0\right] = 1 + \alpha$$

Projecting y_τ (or u_τ) on the average of the shocks at frequency t over τ, $\varepsilon_\tau = \frac{1}{2}\varepsilon_t + \frac{1}{2}\varepsilon_{t-1}$, recovers Θ_0:

$$\mathbb{E}\left[\mathbb{E}\left(\varepsilon_\tau\right)' \mathbb{E}\left(\varepsilon_\tau\right)\right]^{-1} \mathbb{E}\left[\mathbb{E}\left(\varepsilon_\tau\right)' y_\tau\right] = 2\mathbb{E}\left[\left(\frac{1}{2}\varepsilon_t + \frac{1}{2}\varepsilon_{t-1}\right)\left(\alpha\varepsilon_{t-1} + \varepsilon_t\right)\right] = (1 + \alpha)$$

Therefore, even if y is observable only at the frequency τ but there the series of innovations $\{\varepsilon_t\}$ can be obtained by external sources, then it is possible to recover Θ_0 by using the shocks $\{\varepsilon_t\}$ as average over τ.[7]

Bivariate Case. In a multivariate context, temporal aggregation interacts in a non-trivial way with the identification issues that generally arise in VAR models. Assume that the true data

[7]Note that averaging and the sum of the high frequency shocks would deliver the same qualitative IRF but scaled since the average is just a linear transformation of the sum.

generating process is a bivariate VAR(1) at the high frequency t :

$$\begin{bmatrix} x_t \\ y_t \end{bmatrix} = \begin{bmatrix} a_{11} & a_{12} \\ a_{21} & a_{22} \end{bmatrix} \begin{bmatrix} x_{t-1} \\ y_{t-1} \end{bmatrix} + \begin{bmatrix} b_{11} & 0 \\ b_{21} & b_{22} \end{bmatrix} \begin{bmatrix} \varepsilon_t^x \\ \varepsilon_t^y \end{bmatrix}, \tag{4}$$

where $[x_t\, y_t]'$ are scalar endogenous variables, $A = \begin{bmatrix} a_{11} & a_{12} \\ a_{21} & a_{22} \end{bmatrix}$ is the autoregressive matrix, $B = \begin{bmatrix} b_{11} & 0 \\ b_{21} & b_{22} \end{bmatrix}$ is the impact matrix, $\varepsilon_t = [\varepsilon_t^x\, \varepsilon_t^y]'$ is the vector of structural shocks , and $\mathbb{E}\left[\varepsilon_t \varepsilon_t'\right] = \mathscr{I}_2$. For simplicity we assume B to be lower triangular, implying that ε_t^y does not affect x_t within the same period. The results can be generalized to the case where $b_{12} \neq 0$ but the derivations are more cumbersome. The reduced-form residuals $u_t = B\varepsilon_t$ are correlated, with $\mathbb{E}\left[u_t u_t'\right] = \Sigma_{u_t} = BB'$. Assume further that the goal of the empirical analysis is to identify the effect of the shock ε^y on x, but y is observed in every period $t = 1,2,..,T$ whereas x is observed every two periods, i.e. in $\tau = 2,4,...T$. This frequency mismatch generates an estimation problem as the SVAR in Eq.(4) is not observable. The problem is commonly solved by aggregating y to the lower frequency at which x is observable and identifying the shock of interest in a VAR estimated on low-frequency observations. Unfortunately this procedure can be seriously misleading. The temporally aggregated system has the following form:[8]

$$\begin{bmatrix} x_\tau \\ y_\tau \end{bmatrix} = \begin{bmatrix} a_{11}^2 + a_{12}a_{21} & a_{11}a_{12} + a_{12}a_{22} \\ a_{11}a_{21} + a_{21}a_{22} & a_{12}a_{21} + a_{22}^2 \end{bmatrix} \begin{bmatrix} x_{\tau-1} \\ y_{\tau-1} \end{bmatrix} + \begin{bmatrix} u_\tau^x \\ u_\tau^y \end{bmatrix}, \tag{5}$$

where the reduced-form residuals u_τ are linear combinations of present and past structural shocks:

$$\begin{bmatrix} u_\tau^x \\ u_\tau^y \end{bmatrix} = \begin{bmatrix} b_{11}\varepsilon_t^x + (a_{11}b_{11} + a_{12}b_{21})\,\varepsilon_{t-1}^x + a_{12}b_{22}\varepsilon_{t-1}^y \\ b_{21}\varepsilon_t^x + b_{22}\varepsilon_t^y + (a_{21}b_{11} + a_{22}b_{21})\,\varepsilon_{t-1}^x + a_{22}b_{22}\varepsilon_{t-1}^y \end{bmatrix} \tag{6}$$

[8]The example assumes that y is a 'stock', so that the aggregation from high to low frequency entails keeping the last observation within each period τ. The Appendix reports an identical example where the variables are instead averaged over time, as if they were 'flows'.

This process is still a VAR(1), but its residual variance-covariance matrix is given by $\mathbb{E}[u_\tau u_\tau] = \Sigma_{u_\tau} = (\mathscr{I}+A)BB'(\mathscr{I}+A)'$. It is simple to verify that $\Sigma_{u_\tau} \neq \Sigma_{u_t}$. In particular, Σ_{u_τ} is not lower triangular. Thus, the Cholesky structure that would correctly recover the shocks from the high-frequency residuals u_t cannot be applied to u_τ. In other words, temporal aggregation makes it generally impossible to recover the causal impact of the shocks. This conclusion holds for other identification schemes and different structures of the B matrix.[9]

We now demonstrate that, as in the univariate example, the cumulated effects of ε_t^y and ε_{t-1}^y on x_τ can be recovered using a simple average of the HF shocks occurring within period τ, i.e. $\varepsilon_\tau^y = \frac{1}{2}\varepsilon_t^y + \frac{1}{2}\varepsilon_{t-1}^y$.[10] We label Θ_0 the impact effect of the shock ε_τ^y, which is given in this particular case by:

$$\Theta_0 = \begin{bmatrix} \Theta_0^x \\ \Theta_0^y \end{bmatrix} = B_{\bullet 2} + (AB)_{\bullet 2} = \begin{bmatrix} a_{12}b_{22} \\ b_{22}(2+a_{22}) \end{bmatrix} \tag{7}$$

where $B_{\bullet 2}$ and $(AB)_{\bullet 2}$ denote the second columns of the matrices B and AB, respectively.

Our approach proposes to identify the shock from HF data and to compute the IRFs of the LF endogenous variable to the average of shocks identified at HF. One can estimate a VAR at the frequency t that includes y_t and a vector of variables Ω_t of length k, which is defined such that the VAR estimated at frequency t allows the identification of the shocks (information sufficiency):

$$\begin{bmatrix} \Omega_t \\ y_t \end{bmatrix} = \begin{bmatrix} D_{11} & D_{12} \\ D_{21} & d_{22} \end{bmatrix} \begin{bmatrix} \Omega_{t-1} \\ y_{t-1} \end{bmatrix} + \begin{bmatrix} C_{11} & 0 \\ C_{21} & c_{22} \end{bmatrix} \begin{bmatrix} \varepsilon_t^\Omega \\ \varepsilon_t^y \end{bmatrix} \tag{8}$$

where $D = \begin{bmatrix} D_{11} & D_{12} \\ D_{21} & d_{22} \end{bmatrix}$ is the autoregressive matrix and $C = \begin{bmatrix} C_{11} & 0 \\ C_{21} & c_{22} \end{bmatrix}$ is the contemporaneous one. The informational sufficiency assumption states that $\varepsilon_t^\Omega \supseteq \varepsilon_t^x$ and, under this condition and with the right identification strategy, we recover ε_t^y or at least its proxy (i.e. a noisy measure of the shock) $z_t = \varepsilon_t^y + \eta_t, \eta_t \sim wn(0, \sigma_\eta^2), \eta_t \perp \varepsilon_t^y$. In this illustrative case,

[9] See Marcellino (1999) for a more general discussion of the temporal aggregation bias.

[10] The temporal aggregation filter correctly preserves the dynamic effects of the HF shocks between horizon 1 and m and consequently also the aggregated impact effect at LF. Conversely, the LF-VAR autoregressive components are employed to compute the dynamic effects at LF.

$\{z_{t-1}, z_t\}$ are aggregated to z_τ as the average over τ:

$$z_\tau = \frac{z_{t-1} + z_t}{2} \quad \{\tau, t\} = 2, 4, ..., T \tag{9}$$

Then, we employ z_τ as a proxy for the structural shock ε^y_τ. This strategy is valid under the typical assumptions in the Proxy-SVAR literature on the exogeneity of z_t, i.e. $\mathbb{E}[z_t \varepsilon^x_t] = 0$, and its strength, i.e. $\mathbb{E}\left[z_t \varepsilon^y_t\right] \neq 0$. These two properties are translated to ε^y_τ under the correct specification of the HF-VAR ensured by our assumptions on Ω_t (and a large enough number of lags compared to the the frequency mismatch). In this way, we correctly identify the impact effect of the shock ε^y_τ up to a scale factor μ:

$$\Theta^y_0 = \mathbb{E}[z_\tau z_\tau]^{-1} \mathbb{E}\left[z_\tau u^y_\tau\right] = \mu \left(2b_{22} + a_{22} b_{22}\right) \tag{10}$$

$$\Theta^x_0 = \mathbb{E}[z_\tau z_\tau]^{-1} \mathbb{E}[z_\tau u^x_\tau] = \mu a_{12} b_{22} \tag{11}$$

In this way, the ratio $\frac{\Theta^x_0}{\Theta^y_0}$ is correctly estimated (see Eq. 7). Notice that the aggregation of the shocks in alternative ways would not yield the correct ratio $\frac{\Theta^x_0}{\Theta^y_0}$. For example, in the literature, shocks available at a daily frequency have been sometimes aggregated by using weights proportional to the days left in a month or as moving averages. Our paper shows that those approaches are inconsistent with an underlying VAR structure.

2.2 General Proposition

The results obtained in the previous section for a bivariate VAR(1) can be promptly extended to more general set-ups. Consider the general VAR process given by

$$A(L) y_t = B \varepsilon_t \tag{12}$$

where L is the lag operator such that $L^i y_t = y_{t-i}$, the lag polynomial $A(L) = \mathscr{I} - A_1 L - A_2 L^2 - ... - A_p L^p$, ... The SVAR is not observable itself, but corresponds to a multiplicity of reduced-form VAR representations of the form:

$$A(L) y_t = u_t \tag{13}$$

Temporal aggregation can be expressed as a two-step filter. First, the data are made observable only once every m periods, which represents the frequency mismatch, via the filter $D(L) = \mathscr{I} + D_1 L + D_2 L^2 + ... + D_{pm-p} L^{pm-p}$. The specification of $D(L)$ has to be such that the elements of $D(L)A(L)$ are powers of L^m, meaning that only the observable data points enter the transformed process. The conditions for the existence of such a filter, as well as the values taken by the matrices D_i are derived in Marcellino (1999). The second filter, denoted by $W(L)$, depends on the temporal aggregation scheme considered; skip-sampling (or point-in-time sampling) is usually applied to stock variables (e.g. prices) whereas averaging is typically applied to flow variables (e.g. volumes). The validity of our approach, which is based on averaging the HF shocks, is summarized by the following proposition:

Proposition I. *Let y_t follow an underlying high-frequency (HF) VAR process with structural shocks ε_t and reduced-form innovations u_t, with $t = 1, 2, ...T$; and let y_τ represent the low-frequency (LF) version of the process obtained by applying the filters $D(L)$ and $W(L)$ to y_t,* with $\tau = m, 2m, ..., T$. *A consistent estimate of the contemporaneous impact of ε_t on y_τ can be recovered by projecting the LF residuals u_τ on the averages of the HF shocks that occurred within the LF periods, i.e. $\varepsilon_\tau = \frac{\sum_{t=\tau-m+1}^{\tau} \varepsilon_t}{m}$ for every τ.*

Proof: see Appendix B. ■

Proposition I implies that the causal effects of structural shocks in a low-frequency VAR can be recovered using simple averages of the structural shocks identified using high(er) frequency data. Importantly, this procedure is appropriate irrespective of whether the underlying series are aggregate using averaging or skip sampling. This means that it can be safely applied to VARs that include stocks, flows, or any combination of the two. This result hinges on the linearity of VAR models: linearity implies that the sum of the causal effect of the HF shocks $\{\varepsilon_{t-m+1}, \varepsilon_{t-m+2}, ..., \varepsilon_t\}$ is equal to the causal effect of their average, i.e. the aggregated LF shock ε_τ.[11]

[11] Although we do not deal explicitly with shocks identified using narrative sources and event studies, our results on temporal aggregation are also relevant for this strand of literature. One implication of our work, for instance, is that the practice of aggregating HF shocks by taking a moving average or a weighted average of within-period observations is inconsistent if the true data-generating process is a HF VAR.

2.3 Monte Carlo Evidence

In this section we assess the performance of our approach through Monte Carlo tests. Using a structural high-frequency VAR as a data generating process (DGP), we simulate both high-frequency (HF) and low-frequency (LF) data sets and then compare the IRFs obtained by: (i) identifying the shocks directly in a LF-VAR; and (ii) identifying the shocks in a HF-VAR and using them as instruments in the LF-VAR as suggested by our procedure. In all experiments we use the right identification scheme, so that the comparison isolates specifically the implications of identifying the shocks on HF or LF data. For comparison we also report IRFs obtained from a "counterfactual HF VAR" that assumes that all variables are available at the highest frequency. This case is counterfactual from our perspective because our method is motivated precisely by the need to combine HF and LF information, so it has no use in a context where all series are available at HF.

Our experimental design is similar to that in Foroni and Marcellino (2016). The DGP, described in Eq.(14), is a VAR(1) process driven by Cholesky innovations:

$$\begin{pmatrix} x_t \\ y_t \end{pmatrix} = \begin{pmatrix} a_{11} & a_{12} \\ a_{22} & a_{21} \end{pmatrix} \begin{pmatrix} x_{t-1} \\ y_{t-1} \end{pmatrix} + \begin{pmatrix} b_{11} & 0 \\ b_{21} & b_{22} \end{pmatrix} \begin{pmatrix} \varepsilon_t^x \\ \varepsilon_t^y \end{pmatrix}, \tag{14}$$

where $\begin{pmatrix} e_t^x \\ e_t^y \end{pmatrix} \sim \mathcal{N}(0, \mathscr{I}_2)$. We focus on a VAR(1) process because any VAR of higher order can be expressed as a VAR(1) through its companion form. The model is parameterized as follows:

$$\begin{pmatrix} a_{11} & a_{12} \\ a_{22} & a_{21} \end{pmatrix} = \begin{pmatrix} 0.71 & -0.82 \\ 0 & 0.82 \end{pmatrix}, \quad \begin{pmatrix} b_{11} & 0 \\ b_{12} & b_{21} \end{pmatrix} = \begin{pmatrix} 0.28 & 0 \\ 0.23 & 0.95 \end{pmatrix}$$

For ease of exposition we focus on the temporal aggregation via skip-sampling, where variables are observed only once every m periods, for a monthly-quarterly case ($m = 3$) with $\mathscr{T} = 1000$. Figure 1 displays respectively the IRFs of the system for the ε_t^x shock and the ε_t^y shock.

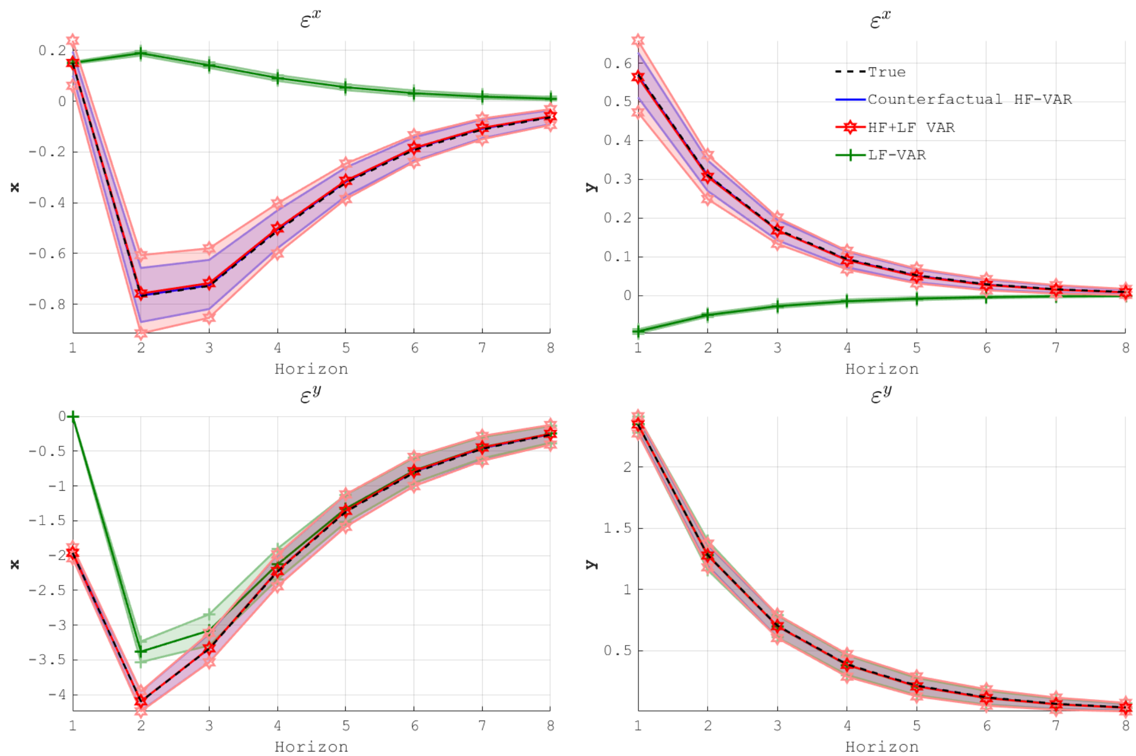

Figure 1: MONTE CARLO EXPERIMENT

IRFs estimates based on data generated using equation 14. The true IRF is represented by the dotted black line. . The remaining lines denote estimates obtained by imposing the correct recursive structure on data sampled at different frequencies: the Counterfactual HF-VAR (blue) relies entirely on high-frequeny data; the LF-VAR (green) relies on low-frequency data; in the HF+LF VAR (red) the shocks are first identified using high-frequency data and used as instruments for the residuals of a low-frequency VAR. Shock 1 and Shock 2 represent respectively shocks to x and *y*; the temporal aggregation is based on skip-sampling and it assumes HF and LF to be monthly and quarterly frequencies. Shaded areas correspond to the 90% percentile across 1000 replications. .

The true responses generated by the DGP are shown as black dashed lines.[12] The differences between the naive LF-VAR estimates and those obtained from our approach, labeled here "HF+LF VAR" are extremely stark: the HF+LF VAR (red stars) recovers the true IRFs (black dash) with great accuracy for both shocks , while the LF-VAR (green crosses) systematically misestimates sign and/or magnitude of the responses. The LF-VAR clearly misses shape and sign of the impact of ε^x on y (figure 1, top right panel); this in turn implies a wrong estimate of the dynamic effect of the shock on x itself (top left panel). Furthermore, the LF-VAR constraints the impact of e^y on x to be 0 (bottom left panel). This is a direct implication of the identification restrictions, that assume a lower-triangular structure. This restriction holds by construction in the HF data, but it is invalid in the LF data, which incorporate an endogenous response of x to ε^y in the (one

[12]Notice that these do not match the element of the B matrix because of temporal aggregation: as in the examples of the previous section, the object of interest is the cumulated low-frequency impact of shocks that materialize at a higher frequency. In particular, the time-1 responses in the figure represent the *quarterly* changes in x and y triggered by shocks whose *monthly* impact is defined by the B matrix.

or two) months that follow the occurrence of the shock within a given quarter. This initial error causes the LF-VAR response to be significantly biased up to the one-year horizon. An important conclusion of this exercise is that, although temporal aggregation only has a direct effect on the contemporaneous response of the model, this is sufficient to distort the path of the IRFs over the entire horizon, both qualitatively and quantitatively. The figure also shows the responses obtained from a Counterfactual HF-VAR : this assumes that the series are all available at high frequency, so that the responses can be estimated in HF and then aggregated to LF . These IRFs (in blue) are virtually indistinguishable from those of the DGP. This model is a useful benchmark from an inference perspective: our approach is consistent but it is also by construction less efficient than the HF-VAR because it relies on a two-stage estimation (IV versus OLS). However, the figures show that (in this representative setup) the efficiency loss is fairly small and the confidence bands are only marginally wider than those of the HF-VAR.

Table 1 displays the accuracy in the estimation of the IRFs for the three approaches described above in a more general setup. We take as benchmark the LF-VAR, whose estimation errors are consistently larger, and report the reduction in the Mean Absolute Deviations (MADs) from the true response obtained by HF-VAR and HF+LF VAR relative to the benchmark. We use 100 random parametrizations of the DGP. We select those parametrizations with real eigenvalues of the autoregressive matrix that belong to the set $(0.7, 0.95)$ to avoid non-stationarity while imposing some persistence in the IRFs.[13] The impact coefficients are drawn as $\{b_{11}, b_{22}\} \sim \mathscr{U}(0, I_2)$ and $b_{12} \sim \mathscr{U}(-1, 1)$. In order to maintain a clear mapping between shocks and variables, we impose $b_{11} > 0.1$ and $b_{22} > 0.1$ and retain only the parametrizations for which $b_{21} < b_{11}$ and $b_{21} < b_{22}$ For each parametrization we examine (i) two alternative sample sizes (T=100 versus T=1,000) and (ii) two frequency mismatches (M/Q versus D/M). Our approach is significantly superior to the LF-VAR in all cases, generating MAD gains of the order of 10% to 70% compared to the LF-VAR. Similar results hold for VAR models of higher order and/or identification schemes based on sign restrictions instead of zero (recursive) restrictions.

[13]This is important because IRFs aggregated at LF are an uninteresting case without persistence, yielding zero effect at LF independently of the impact matrix.

MAD gains over LF-VAR		
	Skip-sampling	**Averaging**
Small (LF) sample size T=100		
Frequency Mismatch: Monthly-Quarterly Case (3)		
HF-VAR	65%	83%
HF+LF VAR	59%	19%
Frequency Mismatch: Daily-Monthly Case (30)		
HF-VAR	80%	97%
HF+LF VAR	43%	42%
Large (LF) sample size T=1000		
Frequency Mismatch: Monthly-Quarterly Case (3)		
HF-VAR	87%	95%
HF+LF VAR	85%	21%
Frequency Mismatch: Daily-Monthly Case (30)		
HF-VAR	93%	90%
HF+LF VAR	79%	79%

Table 1: Monte Carlo Results

Performance comparisons across the counter-factual HF-VAR, the LF-VAR, and our approach. Performances are evaluated in terms of the Mean Absolute Distance (MAD) between the true IRFs and the estimated IRFs in 100 randomly parametrized DGPs. One summary statistic is computed as a mean across all combinations of shocks-variables in the system. The gains are expressed as percentage MAD gains over the LF-VAR. We analyze different cases for a VAR(1) DGP by varying: i) temporal aggregation scheme, either skip-sampling or averaging; ii) the frequency mismatch between HF and LF by 3 (monthly-quarterly case) or 30 (monthly-daily case); iii) sample size, either small (100 LF observations) or large (1000 LF observations).

3 Uncertainty, Volatility, and Financial Markets

The debate on the role of uncertainty in the business cycle is open and fluid. Earlier studies documented a strong impact of aggregate uncertainty on investment and output, but recent contributions have cast doubt on those conclusions showing that uncertainty is often an endogenous response to changes in fundamentals rather than an independent source of fluctuations. This ambiguity is particularly evident in the case of financial uncertainty. Recessions in the US typically coincide with spikes in the volatility of the stock market, but researchers are very much at odds on the interpretation of this coincidence. In line with the findings of Bloom (2009), Ludvigson et al. (2018) argue that shocks to financial uncertainty are a genuine and quantitatively important source of fluctuations in economic activity. Caldara et al. (2016) (CFGZ) document instead a limited role for financial uncertainty, suggesting that economic and policy uncertainty are more relevant. Carriero et al. (2018) conclude that, irrespective of their origin, uncertainty shocks give overall a modest contribution to the dynamics of the real economy. Berger et al. (2019) (BDG) argue that, once they are properly isolated from concurrent changes in fundamentals, exogenous shifts in the expected volatility of the stock market have no impact whatsoever on output and employment.[14] In this section we contribute to this debate by revisiting BDG and CFGZ through the lens of our identification approach. In both cases we take the identification schemes as given and simply shift the restrictions from the monthly frequency used in the original papers to the daily frequency. The motivation for this test is straightforward: if investors respond to uncertainty and macroeconomic news on a daily or intradaily basis, then using monthly observations can lead to biased results even if the underlying restrictions are theoretically sound. Our results forcefully corroborate this conjecture. In both cases we find that (i) the monthly identification restrictions are rejected by the data, and (ii) uncertainty shocks have a significant negative impact on economic activity if daily restrictions are used instead. Interestingly, the real impact of the shock is quantitatively similar in the two applications. This suggests that, once temporal aggregation issues are taken out of the picture, the BDG and CFGZ identification schemes isolate the same type of structural disturbance (see Section 3.3).

3.1 Berger, Dew-Becker and Giglio (2019)

Berger et al. (2019) (BDG) demonstrate that the literature on the role of financial uncertainty is plagued by a pervasive identification problem: volatility in financial markets reflects changes in fundamentals as much as uncertainty about the future, so high volatility can be a powerful predictor of economic slowdowns even if it does not cause them in any way. Following this consideration, BDG distinguish between the realized and the expected volatility of the stock market, and associate "uncertainty shocks" exclusively with the latter. In this way, uncertainty

[14] BDG and CFGZ use option-based measures of the expected volatility of the stock market as a proxy for financial uncertainty. Despite being methodologically different, the model-based proxies used by Ludvigson et al. (2018) and Carriero et al. (2018) (which are constructed using respectively forecast errors and conditional volatility estimates) bear a strong empirical resemblance to implied volatility indices. The correlation between the VXO index and the financial uncertainty proxy in Ludvigson et al. (2018), for instance, is 0.85. This suggests that the discrepancies among the authors' conclusions are caused primarily by the identification schemes rather than measurement issues (see below).

about the future is disentangled from (a broad range of) realized changes in fundamentals that can affect market dynamics at a given point in time. BDG employ a VAR model that includes realized volatility (rv), a measure of option-implied volatility constructed by the authors and conceptually similar to the VXO index (v_1), the Fed Funds rate (ffr), industrial production (ip), and employment (emp). Although the VAR is by construction linear in all variables, thanks to the presence of rv and v_1 the model accounts for the possibility that large price shocks 'today' may raise the expected volatility of the market 'tomorrow', thus capturing a (nonlinear) GARCH effect that is typically found to be important in the empirical finance literature. This feature is particularly valuable when moving from monthly to daily data. The model is estimated over the period between 1983 and 2014. Following the strategy originally used for TFP news shocks (see e.g. Barsky and Sims, 2011), BDG identify uncertainty shocks as the linear combination of the reduced-form residuals that (a) maximizes the two-year ahead forecast error variance (FEV) of realized volatility, but (b) does not affect realized volatility within the same month. The authors find that realized volatility shocks cause a significant decline in economic activity, whereas uncertainty shocks have no effects on the real economy (despite accounting for 30-60 percent of the FEV of realized volatility itself).

We revisit these results using our two-step approach to apply the BDG identification restrictions to daily rather than monthly data. To isolate the specific role of the data frequency, we use the same dataset and rely on the implied volatility measure constructed by the authors rather than the VXO index.[15] As a preliminary step we replicate on daily data the predictive regression run by BDG on their monthly sample. Following BDG, we define the daily realized volatility rv as the squared daily return on the S&P500 index, and regress the cumulative realized volatility over a 6-month horizon (504 trading days), i.e. $\sum_{i=1}^{504} rv_{t+i}$, on the current level of implied volatility, v_1. The regression yields an R^2 coefficient of 0.46, which is nearly identical to that obtained by BDG using monthly data (see Table A1 of the Annex). The correlation between v_1 and leads of rv is thus virtually unaffected by the switch to daily observations. We then follow the procedure discussed in

[15]Our sample starts in 1986 rather than 1983 because the identification strategy exploits the daily $v1$ series constructed by BDG, which, unlike its monthly counterpart, is only available from 1986 onward. This change has no impact on the comparison: BDG show indeed that their results also hold for the 1988-2014 period.

Section 2: (I) we estimate a daily VAR that includes rv and v_1, applying the identification scheme proposed by BDG (and using the same horizon for the forecast error variance maximization); (II) we calculate monthly averages of the the daily realized and implied volatility shocks; and (III) we use these averages as external instruments for the residuals of a monthly VAR model. This model follows exactly the BDG specification: it includes four lags of rv, v_1, ffr, ip and emp (all expressed in natural logarithms but for ffr) and a deterministic constant. Two modeling issues are worth commenting on. The first one is informational sufficiency. The daily 'shocks' obtained from a bivariate model based on rv_t and $v_{1,t}$ fail to pass the Forni and Gambetti (2014) test: they turn out to be correlated to (i.e. Granger-caused by) lagged factors extracted from the FRED-MD database. We consequently discard the bivariate specification in favor of a richer model that includes a range of daily indicators for bond, equity and commodity markets: the shocks identified using this specification pass the information sufficiency test (see Table A2 of the Annex).[16] The second issue relates to invertibility. The shocks obtained from high-frequency data can be used indifferently as internal or external instruments in the monthly VAR model (see Section 2). Since we find no evidence of Granger-causality running from the shocks to the reduced-form residuals of the VAR, following Paul (forth) and Noh (2018), we employ the realized volatility and expected volatility shocks as external instruments for the residuals of the rv_t and $v_{1,t}$ equations.[17]

A comparison between our estimates and those of BDG is reported in Figure 2. The left column shows the response of the US economy to an uncertainty (expected volatility) shock in our model, where the BDG restrictions are applied to daily data. The right column shows a replication of the original results from the BDG monthly estimates.[18] The daily identification reveals that rv increases significantly 'on impact', i.e. within a month, in response to a v_1 shock. This result is crucial because it shows that the zero restriction suggested by BDG cannot be imposed on a

[16]The expanded daily VAR includes (in logs) $s\&p500$ price index, Fed Funds rate, BAA corporate bond spread, euro-dollar exchange rate, Economic Policy Uncertainty index, the $s\&p$ Goldman Sachs Commodity Index (GSCI) - gold spot price. If these variables are included in a forecasting regression for rv_t, the R^2 increases from 0.46 to 0.68, see table A1.

[17]In the application of Section 3.2 we follow the alternative route, due to failure in the aforementioned test, and include the shocks directly in the VAR.

[18]In order to keep the test more general and make the models more comparable we focus on the "unrestricted" VAR specification used by BDG (where no restrictions are placed on the coefficients of the VAR model). BDG show that the results are very similar to those of their benchmark "restricted" VAR, where the coefficients of lagged interest rates, industrial production and employment in the rv_t and $v_{1,t}$ equations are set to zero.

monthly dataset: even if *rv* is assumed not to respond *within a day* to an uncertainty shock, the stock price adjustments that take place in the days after the shock cause *rv* to rise *within a month* after an exogenous increase in uncertainty. This result is extremely robust even using an "agnostic" set-identification procedure: out of 1 million random draws of the impact matrix of the daily VAR, not even one is compatible with *rv* remaining constant for 21 business days. The monthly BDG model captures the impact of *rv* on v_1, but, by failing to account for the within-month response of the stock market, it shuts down the channel that runs in the opposite direction. As the rest of Figure 2 shows, this has first-order implications for all remaining IRFs: contrary to BDG, we find that uncertainty shocks cause a large and statistically significant drop in both industrial production and employment, coupled with a (less significant) decline in interest rates.

To have a clearer view of the propagation mechanisms, in Figure 3 we compare the responses to the expected and realized volatility shocks estimated with our model. Following BDG, we normalize the v_1 and *rv* shocks so that they have the same cumulated impact on *rv* over the 2- to 24-month horizon. The responses of v_1 to the two shocks are nearly identical. In particular, v_1 rises strongly in response to an *rv* shock, which is consistent with this shock causing a large and persistent increase in the realized volatility of the stock market over the following months. As in BDG, the *rv* shock produces larger fluctuations in industrial production and employment than the v_1 shock.[19] The response of *ip* to the v_1 shock is about one-third the response to a *rv* shock. In the case of *emp* the response is roughly half as big and the confidence bands are narrower. Employment may be more sensitive to uncertainty if hiring and firing decisions are subject to non-convex adjustment costs, as demonstrated originally in Bloom (2009).

In summary, by imposing the BDG restrictions on daily rather than monthly observations we model in a more flexible way the daily interactions between volatility expectations, market returns and the realized volatility of the stock market. The exercise reveals that (i) realized volatility rises endogenously in response to positive uncertainty shocks occurring within a given month; and (ii) forcing this response to be zero biases the IRFs in the BDG framework, causing uncertainty shocks

[19]This difference is not surprising because the *rv* shock is effectively a summary statistic that captures a broad range of structural shocks (including e.g. macroeconomic news and changes monetary or fiscal policy surprises) rather than a true structural shock (see BDG for further details on the interpretation of the identified shocks).

to appear irrelevant for the the real economy.

Daily Analysis in BDG. BDG explicitly acknowledge that high-frequency dynamics could potentially invalidate their identification procedure and dedicate one of their robustness test specifically to this problem (see BDG Section 6.1.3). However, their test focuses on a different issue compared to our approach. BDG use a daily VAR model to generate a counterfactual realized volatility series that is by construction unaffected by uncertainty shocks ($rv^{iv-purged}$), and show that the IRFs do not change if this is introduced in the VAR instead of the original rv series. As the authors note, $rv^{iv-purged}$ satisfies by construction their monthly restrictions: hence, the test demonstrates that the estimated impact of rv shocks in the baseline model is not distorted by temporal aggregation.[20] Our procedure is more general (it eliminates the bias in the estimation of v_1 as well as rv shocks) and it allows us to estimate the contemporaneous monthly responses of v_1 and rv to the two shocks without restrictions (instead of assuming one of them to be zero). The results show that the critical distortion arises indeed for the v_1 rather than the rv shocks. In line with BDG, we find that rv shocks matter irrespective of the data frequency and that their quantitative relevance is not affected by the endogeneity of rv to v_1. Unlike BDG, however, we conclude that v_1 shocks are important in their own right: their economic impact appears only if the monthly response of rv is left unrestricted, which in turn requires an identification based on daily data.

[20]The logic is that $rv^{iv-purged}$ removes the within-month influence of v_1 shocks on rv. If the impact of rv shocks in the baseline model were caused by such influences, the VAR based on the counterfactual series would deliver different (i.e. smaller and less significant) estimates of the impact of these shocks. This test eliminates biases in the rv but not in the $v1$ shocks. These shocks are in fact deliberately removed from the counterfactual rv series (and hence made less powerful than they are in reality) in order to construct a robust realized volatility estimate.

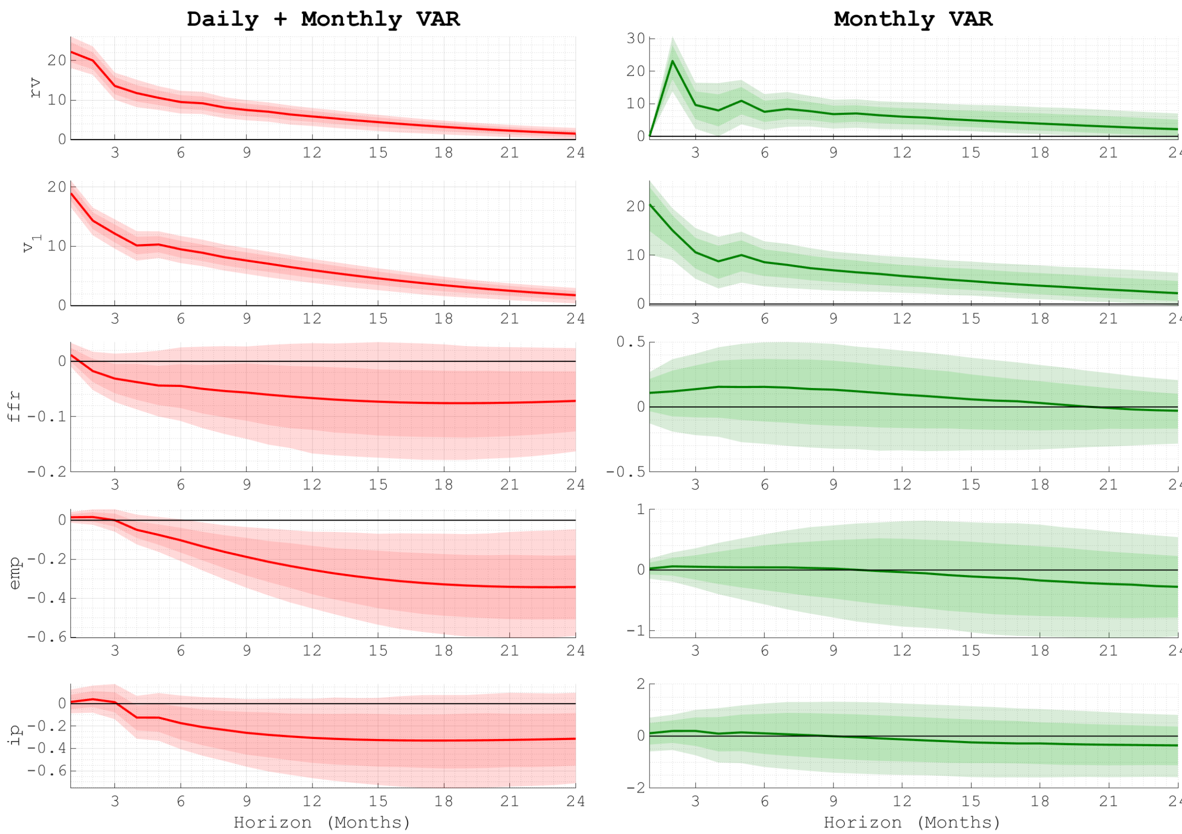

Figure 2: Comparison with BDG

Impact of uncertainty shocks under the identification strategy of Berger et al. (2019, BDG). The shocks are identified as innovations to the option-implied expected volatility of the stock market ([v_1]) that are orthogonal to the realized market volatility (rv). In the Daily+Monthly VAR (left column) the shock is identified imposing the restrictions on daily data, averaged to the monthly frequency, and then used as an external instrument in the monthly VAR model. In the Monthly VAR (right column) the restrictions are imposed directly on monthly data as in BDG. The estimation sample is January 1983-December 2014. The variables included in the VAR are: realized volatility (rv), option implied volatility (v_1), Fed Funds rate (ffr), employment (emp) and industrial production (ip). Each plot reports the median response with 68% and 90% confidence bands computed using 1000 wild bootstrap replications.

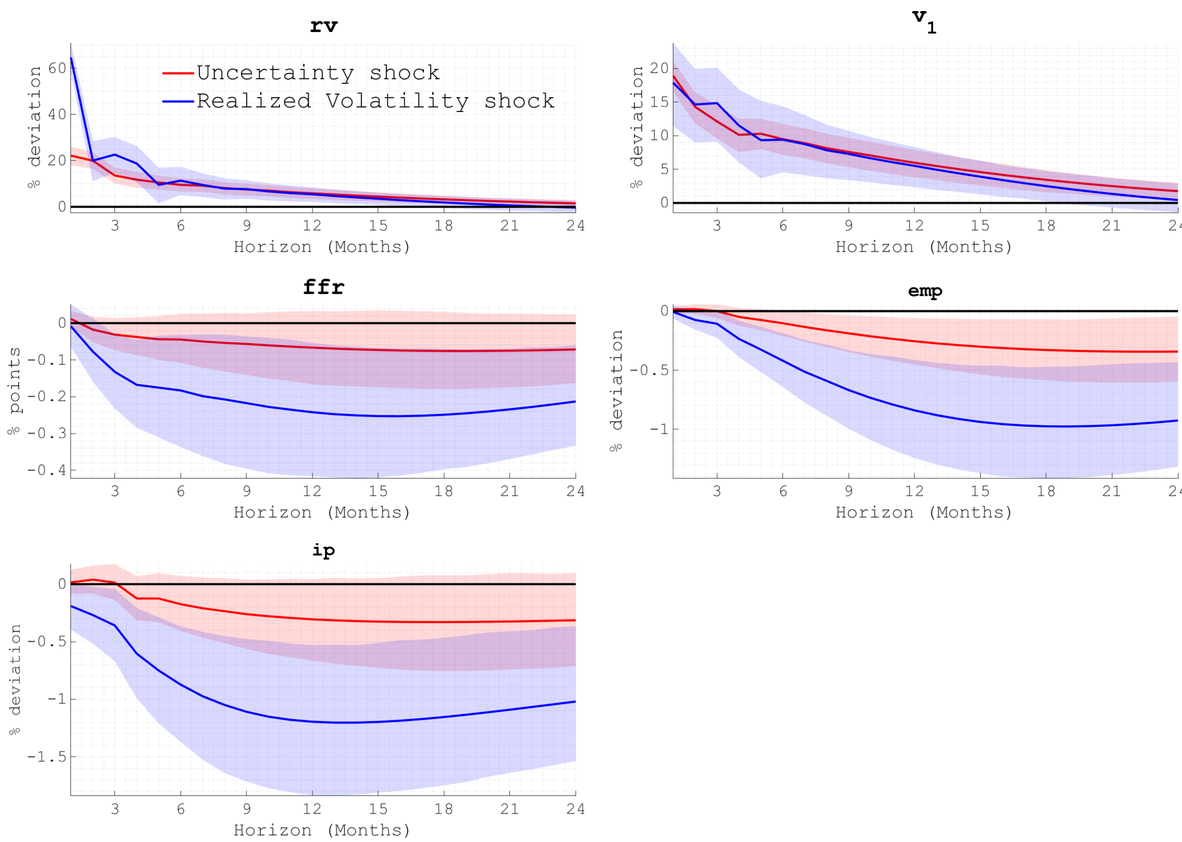

Figure 3: Comparison of IRFs: Uncertainty versus Realized Volatility Shocks

Impact of realized volatility and uncertainty (i.e. expected volatility) shocks in the Daily+Monthly VAR model. Following BDG, uncertainty shocks are identified as innovations to the option-implied expected volatility of the stock market (v_1) that are orthogonal to the realized market volatility (rv). The shocks are identified applying the BDG restrictions to a daily VAR model, averaged to the monthly frequency, and then used as external instruments in the monthly VAR model. The estimation sample is January 1986–December 2014. The variables included in the VAR are: realized volatility (rv), option implied volatility (v_1), Fed Funds rate (ffr), employment (emp) and industrial production (ip). Each plot reports the median responses with 90% bootstrapped confidence bands.

3.2 Caldara, Fuentes-Albero, Gilchrist and Zakrajsek (2016)

Caldara et al. (2016) (CFGZ) focus on a specific source of confusion between first- and second-moment shocks, namely the overlap between uncertainty and credit conditions. Spikes in volatility are typically associated to a rise in credit spreads: Stock and Watson (2012) note that this correlation is so strong that one could be tempted to see uncertainty proxies and credit spreads as different manifestations of a common underlying structural shock. Furthermore, large bond investors are an obvious example of highly reactive economic agents that respond quickly to any change in economic fundamentals. Based on these considerations, CFGZ propose an identification strategy (at the monthly frequency) that avoids contemporaneous restrictions on asset prices and allows credit and uncertainty shocks to have similar effects on the economy. In this section we show however that using plausible restrictions is not sufficient: the endogenous response of the bond market cannot be captured using low-frequency observations, and the results change significantly if the CFGZ identification is applied to daily rather than monthly data.

CFGZ estimate a range of Bayesian VAR(6) models using monthly data for the period January 1975–March 2015 and specifications that include 10 variables: industrial production, private payroll employment, real personal consumption expenditures, PCE deflator, the S&P Goldman Sachs Commodity Index, a value-weighted stock market price index, the 2-year and 10-year Treasury bond yields, the Excess Bond Premium and a proxy of economic or financial uncertainty. The first six variables are included in log differences, yields and spreads are in percentage points and the uncertainty proxy (which differs across specifications) is in logarithms. All models are estimated imposing a Minnesota prior on the reduced-form parameters. In order to account for the simultaneity between asset prices and uncertainty, CFGZ propose a novel identification strategy based on the penalty function approach (PFA) developed by Faust (1998) and Uhlig (2005). In particular, they identify uncertainty (financial) shocks as the linear combination of reduced-form residuals that maximize the response of the uncertainty (financial conditions) indicator over a predefined horizon. The identification is implemented sequentially, imposing an orthogonality condition between the two structural shocks. An important advantage of the PFA is that it leaves the impact response of uncertainty and financial conditions unrestricted while identifying shocks that cause significant and persistent changes in these variables. One of its limitations is that, as CFGZ acknowledge, the results may depend on which shock is identified first, pretty much as in a recursive identification scheme. The authors find indeed that, while macro uncertainty has a significant macroeconomic impact regardless of the ordering, financial uncertainty shocks matter if and only if they are identified first: under the alternative ordering, a rise in financial uncertainty causes a decline in the stock market but it has no implications for the real economy. This dichotomy is problematic because there are no economic arguments to favor either of the two orderings.

We run our tests focusing on the CFGZ specification where financial uncertainty is captured by the implied volatility of the stock market, i.e. the VXO index; this makes the results more comparable to those obtained in the BDG setup examined in previous section.[21] As in the BDG case, we take sample, VAR specification and identification restrictions as given, and simply shift the restrictions from the monthly to the daily frequency. We set the PFA maximization horizon to 120 business days, so to match the 6-month horizon used by CFGZ. Since the Excess Bond Premium computed by Gilchrist and Zakrajsek (2012) is not available at daily frequency, we replace it with the spread between Moody's Seasoned BAA Corporate Bond yield and the 10-Year Treasury constant maturity yield (BAA10Y).[22] We estimate the IRFs by including the shocks (aggregated at the monthly frequency) as internal instruments in the monthly VAR. The reason is twofold. First, we are consistent with the Bayesian framework employed by CFGZ, which matters for comparability reasons. Second, we cannot reject that the identified shocks Granger cause the residuals of the VAR, thus they cannot be employed as external instruments.[23]

Figure 4 displays the IRFs to an uncertainty shock identified with our approach (left column) and with the CFGZ approach (right column). The results obtained with uncertainty shocks ordered first are shown in red (VXO-BAA case), those obtained with uncertainty ordered second are in green (BAA-VXO case). To ease the comparisons across models and orderings, the shocks are normalized to generate a 3% increase in the VXO index in all models. For each variable the figure reports the mean response along with a 90% credible set. To save space we focus on the VXO

[21] Like CFGZ, we use a sample starting from January 1986 to exploit the VXO series from the beginning. We obtain similar results when replacing the VXO with the realized volatility of the stock market or the Economic Policy Uncertainty index, two alternative proxies of uncertainty used in CFGZ that are also available at the daily frequency. However, these have important limitations in our context: the EPU index measures policy rather than financial uncertainty, and the realized volatility of the stock market captures changes in fundamentals as well as uncertainty in a strict sense (BDG use it indeed as a control variable rather than an uncertainty proxy, see section [3.1]).

[22] To insure the spread has no bearing on our results, we (i) replicate the monthly analysis of CFGZ using the BAA spread instead of EBP; and (ii) estimate an additional specification where the identification is based on the daily BAA spread but the EBP is used as a proxy of financial conditions in the monthly VAR. The results are reported in Appendix C.

[23] Notice that, conversely, the shocks pass the information sufficiency test of Forni and Gambetti (2014).

index, the BAA spread and a set of key macroeconomic aggregates, namely industrial production (IPM), employment (EMPL), consumption (PCE) and the consumption deflator (PPCE): a full set of IRFs can be found in Appendix C. Our first and most important result is that using daily data is sufficient to resolve the ambiguity encountered in CFGZ. The right column of the figure confirms the CFGZ findings: uncertainty shocks affect credit spreads and economic activity if they are identified before financial shocks (green bands), but become completely irrelevant if the ordering is reversed (red bands). The left column shows that this discrepancy disappears in our model. If the identification is based on daily data, irrespective of the ordering uncertainty shocks cause a sizable and persistent increase in the spread (row 2), a decline in industrial production and employment (rows 3 and 4) and a fall in prices (bottom row). For most variables the mean responses are indeed almost indistinguishable across orderings. The credible sets tend to be wider when financial shocks are identified first, but they are still sufficiently narrow to rule out a null response of output, employment and inflation to the shock. The behavior of the credit spread is key in explaining this result. Under the BAA-VXO ordering, the identification assumes uncertainty shocks to have maximal impact on VXO only after 'netting out' the impact of financial shocks on both volatility and credit spreads. This seemingly minor restriction has a dramatic impact in the CFGZ model, where the response of the spread becomes negative on impact (a result that is hardly justifiable from a theoretical perspective) and null in the longer term.[24] It has no impact instead in our model, where the spread rises on impact by roughly the same amount under both orderings. The likely reason is that bond markets price uncertainty on a daily rather than a monthly basis, and the bulk of the adjustment takes place within a few days after a shock. The identification sequence becomes irrelevant once this component of the endogenous market response is fully taken into account. A second interesting result is the large negative response of the PCE price index. This deflationary effect is far more pronounced than in the original CFGZ model, and it corroborates the idea that uncertainty acts mainly through the demand side of the economy (Leduc and Liu, 2016 and Basu and Bundick, 2017). The responses of inflation, output and employment are indeed

[24]The negative initial response of the BAA spread in figure 4 is fully consistent with the results of CFGZ: the authors show that (i) the EBP also falls after a rise in uncertainty if financial shocks are identified first, and (ii) this result is indeed robust to replacing the VXO index with a number of alternative proxies of aggregate uncertainty.

qualitatively similar to those obtained in the BDG setup examined in Section 3.1. This suggests that the differences between BDG and CFGZ also depend on temporal aggregation issues, and that using daily data can yield results that are (also) more robust across identification schemes (see Section 3.3).

Figure 5 compares the impact of uncertainty and financial shocks in our model under the two alternative PFA orderings. In this case we do not rescale the responses, so to assess the relevance of the two mechanisms for shocks of equal size (one standard deviation). The ordering makes no difference for the financial shock (right column): a 1σ financial tightening causes a rise in the spread and in the VXO index and a decline in prices and economic activity, and the responses are virtually identical for the VXO-BAA and the BAA-VXO case. Furthermore, the response of the BAA spread is extremely similar to that in Figure 5, with a large contemporaneous increase and a peak at the three-month horizon followed by a slow decline. This similarity is presumably one of the reasons why the PFA identification is not robust to the ordering of the shocks when applied to monthly observations: changes in uncertainty and changes in risk aversion, leverage or net worth (all of which represent likely 'financial shock' candidates) have approximately the same implications for corporate bond spreads. Finally, the impact of financial and uncertainty shocks are not only qualitatively similar but also of comparable magnitudes. A 1σ financial shock causes IPM and EMP to drop respectively by -0.5% and -0.12% at the one-year horizon; for the 1σ uncertainty shock these figures are approximately -0.2% and -0.1%, with some variation depending on the ordering. PCE only responds to financial shocks, whereas PPCE is more sensitive to uncertainty shocks. These estimates confirm clearly that uncertainty shocks have significant macroeconomic effects. From a methodological perspective, they also demonstrate that the ambiguities documented by CFGZ disappear completely if their identification strategy is implemented using daily rather than monthly data.

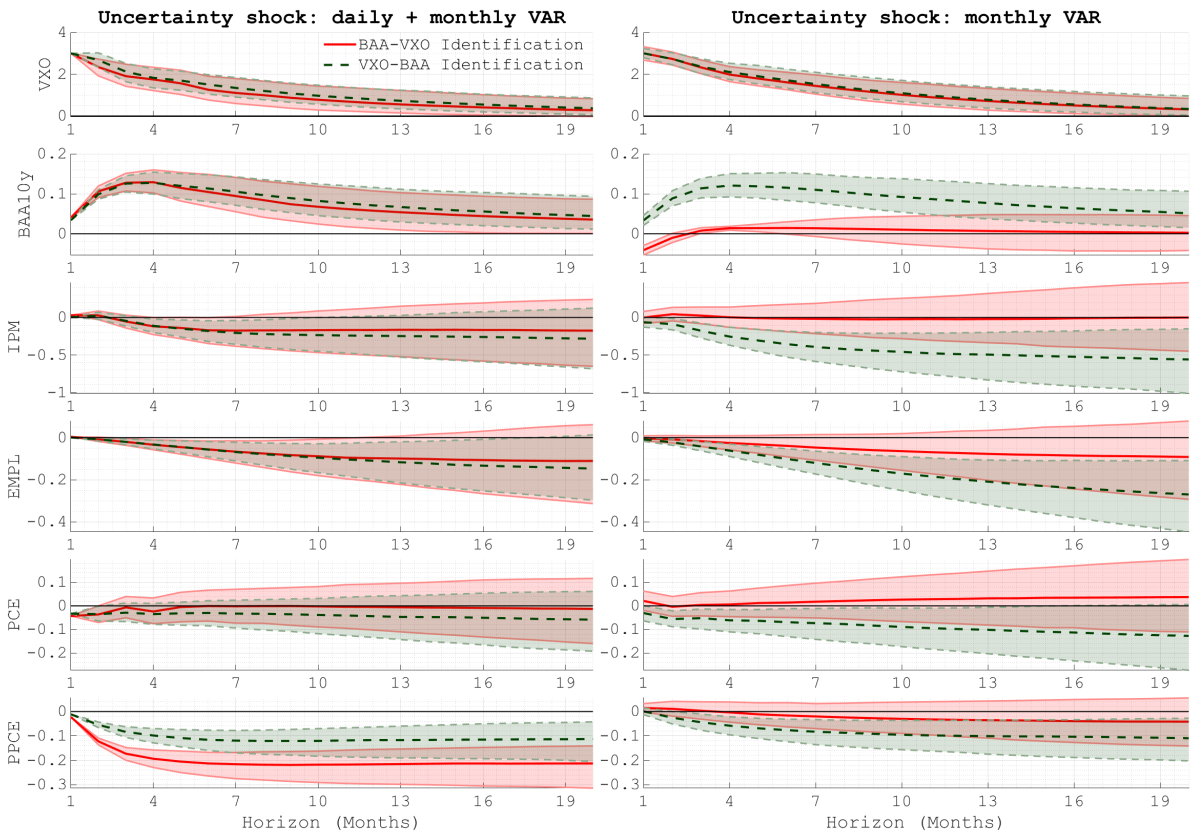

Figure 4: Impact of Uncertainty Shocks - Comparison with Caldara et al. (2016)

Impact of an uncertainty shock under the identification restrictions of Caldara et al. (2016, CFGZ). All IRFs are based on the CFGZ Penalty Function Approach (PFA). In the Daily+Monthly VAR model (left column) the shock is identified imposing the restrictions on daily data, averaged to the monthly frequency and introduced as an additional variable in the monthly VAR model. In the Monthly VAR model (right column) the restrictions are imposed directly on monthly data as in CFGZ. Green and red areas correspond to the responses estimated ordering the ucertainty shock before and after the financial shock in the penalty function maximization. The variables plotted are: VXO index, the spread between US BAA Corporate Yield and the 10Y US Treasury yield (BAA10Y), industrial production (IPM), private payroll employment (EMPL), real personal consumption expenditures (PCE) and PCE deflator (PPCE) .Each plot reports the median response with a 90% Bayesian credible set.

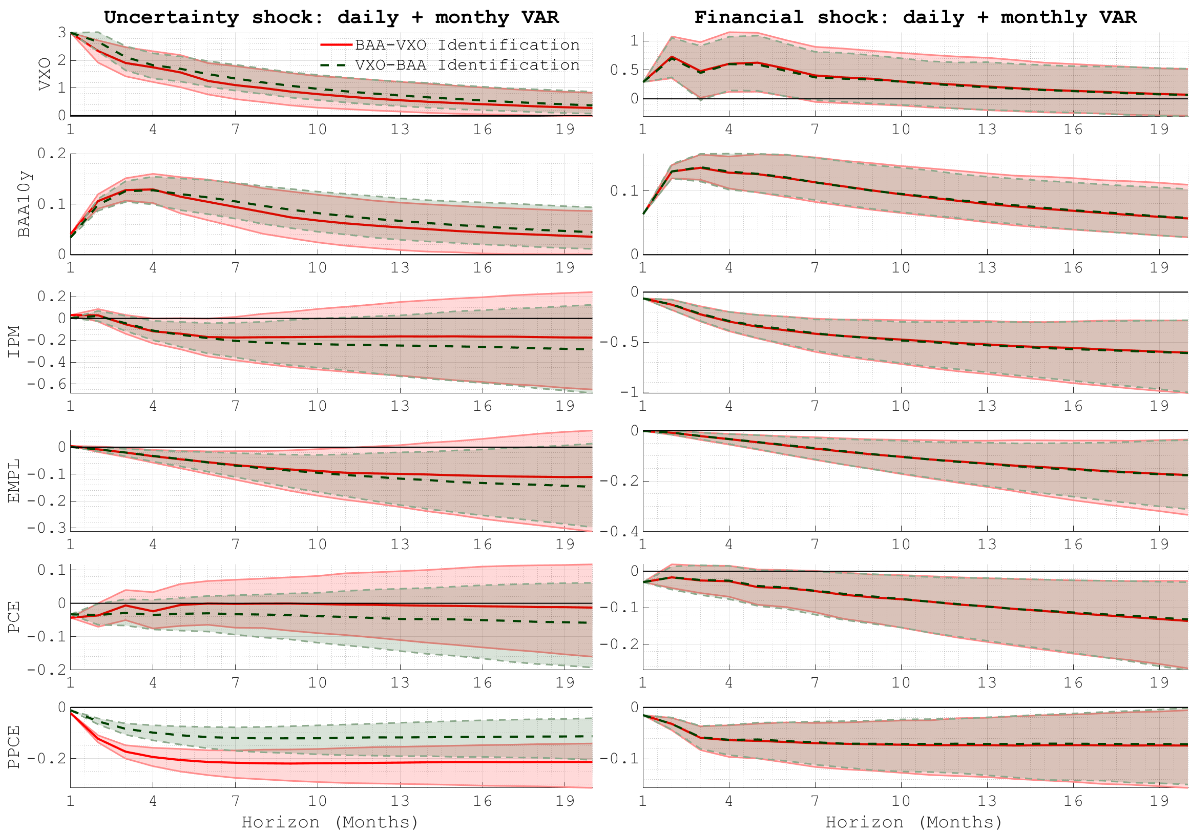

Figure 5: Impact of Uncertainty and Financial Shocks in the Daily+Monthly VAR model
The shocks are identified by applying the Penalty Function Approach (PFA) of Caldara et al. (2016) (CFGZ) to a daily VAR, averaged to the monthly frequency, and then introduced as additional variables in a monthly VAR model based on the CFGZ specification. Green and red areas correspond to the responses estimated ordering the ucertainty shock before and after the financial shock in the PFA. The variables plotted are: VXO index, the spread between US BAA Corporate Yield and the 10Y US Treasury yield (BAA10Y), industrial production (IPM), private payroll employment (EMPL), real personal consumption expenditures (PCE) and PCE deflator (PPCE) Each plot reports the median response with a 90% Bayesian credible set.

3.3 The real impact of financial uncertainty: a summary

The macroeconomic implications of financial uncertainty shocks have proved to be hard to pin down and highly variable across samples, models, and identification strategies. Our work shows that this ambiguity is largely caused by the daily interaction between stock returns, bond spreads and market volatility, which seriously complicates the identification of uncertainty shocks in VAR models based on monthly or quarterly data. In Figure 6 we bring together the results of our revisitation of BDG and CFGZ. The lines with diamonds show the cumulative response of industrial production to uncertainty shocks in the original monthly models of BDG and CFGZ; the lines without markers show the responses obtained by replicating their analysis with our method, which imposes the same identification restrictions on daily rather than monthly data. Solid and dashed lines indicate respectively statistically significant and not significant responses . To make the responses comparable, we normalize the IRFs so that the cumulative increase in financial uncertainty (i.e. implied equity volatility, which is captured respectively by VXO in CFGZ and v_1 in BDG) is the same across models. The responses thus represent the elasticities of industrial production to implied volatility. With monthly data, the responses differ widely both across models (BDG vs CFGZ) and between identification schemes within a given model (see CFGZ cases). The null hypothesis that uncertainty has no impact cannot be rejected, but the upper bound for the elasticity exceeds 30% at the one-year horizon. With daily data, the elasticities at the one-year horizon fall in an interval between 2% and 12% and the estimates are statistically significant in all cases.

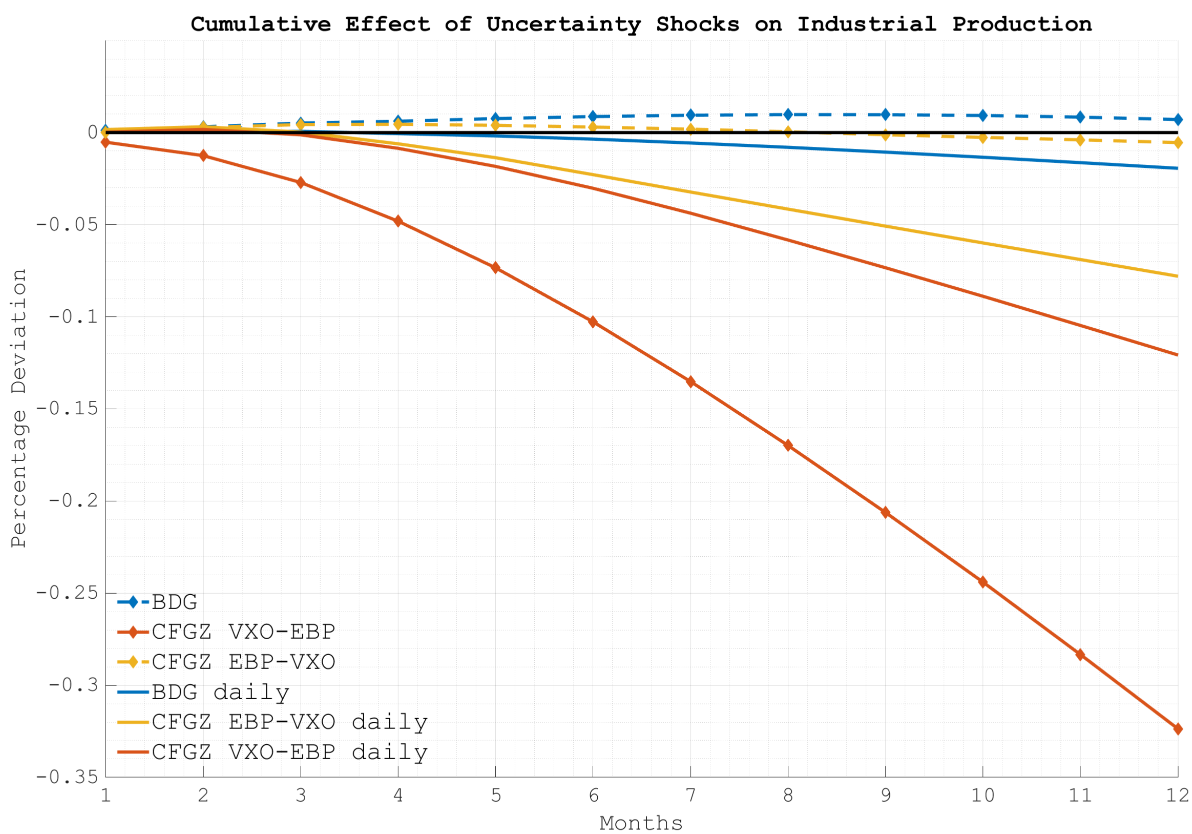

Figure 6: Cumulative Effect of Uncertainty Shocks on Industrial Production

4 Conclusions

Does financial uncertainty affect the real economy? In this paper we answer this question using a novel approach to tackle the complex interactions between uncertainty and asset prices. The approach consists of three steps: we identify uncertainty shocks applying the restrictions proposed by Caldara et al. (2016) and Berger et al. (2019) to daily observations, we aggregate the shocks to the monthly frequency, and we use the aggregated series as instruments in monthly VAR models of the US economy. The strategy is motivated by a simple consideration: if financial markets react to macroeconomic news and changes in risk on a continuous basis, then using monthly or quarterly data may by construction prevent a correct separation between exogenous shocks and endogenous responses. More generally, the use of high-frequency observations simplifies structural identification in a VAR because restrictions that hold on high-frequency data can easily be violated by the corresponding low-frequency (temporally aggregated) data.

Our methodological contribution is to show that, as long as the economy is described by a VAR at high frequency, temporal aggregation problems can be bypassed by identifying the shock of interest at that frequency and then using a low-frequency average of the shock to estimate the impulse-response functions. This simple strategy delivers consistent and unbiased estimates of the impact of the shocks in a broad range of data-generating processes. We then demonstrate that using daily data is necessary in order to correctly isolate the impact of financial uncertainty shocks. We replicate the studies by Berger et al. (2019) and Caldara et al. (2016) imposing their identification restrictions on daily rather than monthly data, and find that the impact of uncertainty shocks on economic activity is (i) always negative and significant, and (ii) fairly similar across models. This reflects daily interactions between stock prices, bond yields and volatility expectations that are captured by our method but neglected in the original analyses. In particular, our replication of Berger et al. (2019) shows that shocks to the expected volatility of the stock market also affect the realized volatility of the market within a given month, and that this endogenous response, which is shut down in the monthly model, plays a key role in the transmission of uncertainty to the real economy. A similar problem arises in the Caldara et al. (2016) set-up, where temporal aggregation distorts the relation between expected volatility and credit spreads. By using daily data for identification, and hence letting spreads and volatility interact freely at the monthly frequency, we estimate a negative and robust impact of volatility shocks on financial conditions, output and inflation. The responses are quantitatively similar in the two applications, suggesting that – once the endogenous response of equity and bond markets is accounted for – the two identification schemes isolate the same type of shock despite the differences among the underlying models.

Future work on the effects of uncertainty shocks should start from the premise that financial markets are a key link in the transmission mechanism, and that even theoretically sound identification strategies may misrepresent their role if applied to monthly or quarterly samples. More generally, researchers can resort to the approach proposed in this paper to tackle identification challenges in cases where imposing restrictions on low-frequency data is problematic and high-frequency data (e.g. from financial markets or textual sources) can be used to better isolate the shocks of interest.

References

ALESSANDRI, P., A. GAZZANI, AND A. VICONDOA (2020): "Uncertainty and the Macroeconomy: A High Frequency Identification Strategy," Temi di Discussione (Working Paper) 1284, Bank of Italy.

ARELLANO, C., Y. BAI, AND P. J. KEHOE (2019): "Financial Frictions and Fluctuations in Volatility," *Journal of Political Economy*, 127, 2049–2103.

BAKER, S. R. AND N. BLOOM (2013): "Does Uncertainty Reduce Growth? Using Disasters as Natural Experiments," Working Paper 19475, National Bureau of Economic Research.

BAKER, S. R., N. BLOOM, AND S. J. DAVIS (2016): "Measuring Economic Policy Uncertainty," *The Quarterly Journal of Economics*, 131, 1593–1636.

BARSKY, R. B. AND E. R. SIMS (2011): "News shocks and business cycles," *Journal of Monetary Economics*, 58, 273–289.

BASU, S. AND B. BUNDICK (2017): "Uncertainty Shocks in a Model of Effective Demand," *Econometrica*, 85, 937–958.

BEAUDRY, P. AND M. SAITO (1998): "Estimating the effects of monetary shocks: An evaluation of different approaches," *Journal of Monetary Economics*, 42, 241–260.

BERGER, D., I. DEW-BECKER, AND S. GIGLIO (2019): "Uncertainty Shocks as Second-Moment News Shocks," *The Review of Economic Studies*.

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

CALDARA, D., C. FUENTES-ALBERO, S. GILCHRIST, AND E. ZAKRAJŠEK (2016): "The macroeconomic impact of financial and uncertainty shocks," *European Economic Review*, 88, 185–207.

CANOVA, F. (2007): *Methods for Applied Macroeconomic Research*, Princeton University Press.

CARRIERO, A., T. E. CLARK, AND M. MARCELLINO (2018): "Measuring Uncertainty and Its Impact on the Economy," *The Review of Economics and Statistics*, 100, 799–815.

CASCALDI-GARCIA, D. AND A. B. GALVAO (2016): "News and Uncertainty Shocks," EMF Research Papers 12, Economic Modelling and Forecasting Group.

CHRISTIANO, L. J. AND M. EICHENBAUM (1987): "Temporal aggregation and structural inference in macroeconomics," *Carnegie-Rochester Conference Series on Public Policy*, 26, 63 – 130.

CHUDIK, A. AND G. GEORGIADIS (2019): "Estimation of Impulse Response Functions When Shocks are Observed at a Higher Frequency than Outcome Variables," Globalization Institute Working Papers 356, Federal Reserve Bank of Dallas.

CIESLAK, A. AND H. PANG (2020): "Common Shocks in Stocks and Bonds," Working Paper 28184, National Bureau of Economic Research.

FAUST, J. (1998): "The robustness of identified VAR conclusions about money," *Carnegie-Rochester Conference Series on Public Policy*, 49, 207 – 244.

FAUST, J., E. T. SWANSON, AND J. H. WRIGHT (2004): "Identifying VARS based on high frequency futures data," *Journal of Monetary Economics*, 51, 1107 – 1131.

FERNANDEZ-VILLAVERDE, J., P. GUERRON-QUINTANA, J. F. RUBIO-RAMIREZ, AND M. URIBE (2011): "Risk Matters: The Real Effects of Volatility Shocks," *American Economic Review*, 101, 2530–2561.

FERNANDEZ-VILLAVERDE, J. AND P. A. GUERRON-QUINTANA (2020): "Uncertainty shocks and business cycle research," *Review of Economic Dynamics*, 37, 118–146.

FERRARA, L. AND P. GUERIN (2018): "What are the macroeconomic effects of high-frequency uncertainty shocks?" *Journal of Applied Econometrics*, 33, 662–679.

FORNI, M. AND L. GAMBETTI (2014): "Sufficient Information in Structural VARs," *Journal of Monetary Economics*, 66, 124–136.

FORONI, C. AND M. MARCELLINO (2016): "Mixed Frequency Structural VARs," .

GERTLER, M. AND P. KARADI (2015): "Monetary Policy Surprises, Credit Costs, and Economic Activity," *American Economic Journal: Macroeconomics*, 7, 44–76.

GILCHRIST, S. AND E. ZAKRAJSEK (2012): "Credit Spreads and Business Cycle Fluctuations," *American Economic Review*, 102, 1692–1720.

HENDRY, D. F. (1992): "An econometric analysis of TV advertising expenditure in the United Kingdom," *Journal of Policy Modeling*, 14, 281–311.

HERBST, E. AND B. K. JOHANNSEN (2020): "Bias in Local Projections," .

JURADO, K., S. C. LUDVIGSON, AND S. NG (2015): "Measuring Uncertainty," *American Economic Review*, 105, 1177–1216.

KALMAN, D. (1982): "Generalized Fibonacci numbers by matrix methods," *The Fibonacci Quarterly*, 20 (1), 73ï¿œ76.

KILIAN, L. AND H. LUTKEPOHL (2017): *Structural Vector Autoregressive Analysis*, Themes in Modern Econometrics, Cambridge University Press.

KILIC, E. (2007): "The generalized order-k Fibonacciï¿œPell sequence by matrix methods," *Journal of Computational and Applied Mathematics*, 209, 133–145.

LEDUC, S. AND Z. LIU (2016): "Uncertainty shocks are aggregate demand shocks," *Journal of Monetary Economics*, 82, 20–35.

LUDVIGSON, S. C., S. MA, AND S. NG (2018): "Uncertainty and Business Cycles: Exogenous Impulse or Endogenous Response?" .

MARCELLINO, M. (1999): "Some Consequences of Temporal Aggregation in Empirical Analysis," *Journal of Business and Economic Statistics*, 17, 129–136.

MARCET, A. (1991): "Temporal Aggregation of Economic Time Series," in *Rational Expectations Econometrics*, ed. by L. P. Hansen, T. J. Sargent, J. Heaton, A. Marcet, and W. Roberds, Westview Press Boulder, chap. 10, 237–282.

MERTENS, K. AND M. O. RAVN (2013): "The Dynamic Effects of Personal and Corporate Income Tax Changes in the United States," *American Economic Review*, 103, 1212–47.

MIRANDA AGRIPPINO, S. AND G. RICCO (2018): "(A Note on the) Identification with External Instruments in Structural VARs under Partial Invertibility," .

NOH, E. (2018): "Impulse-response analysis with proxy variables," *Mimeo*.

PAUL, P. (forth): "The Time-Varying Effect of Monetary Policy on Asset Prices," *The Review of Economics and Statistics*.

PIFFER, M. AND M. PODSTAWSKI (2018): "Identifying Uncertainty Shocks Using the Price of Gold," *Economic Journal*, 128, 3266–3284.

PLAGBORG-MØLLER, M. AND C. K. WOLF (2019a): "Instrumental Variable Identification of Dynamic Variance Decompositions," *Unpublished paper: Department of Economics, Princeton University*.

——— (2019b): "Local projections and VARs estimate the same impulse responses," *Unpublished paper: Department of Economics, Princeton University*.

RAMEY, V. (2016): "Macroeconomic Shocks and Their Propagation," .

RAMEY, V. A. AND S. ZUBAIRY (2018): "Government Spending Multipliers in Good Times and in Bad: Evidence from US Historical Data," *Journal of Political Economy*, 126, 850–901.

SAHIN, A. (2018): "Inverse and factorization of triangular Toeplitz matrices," *Miskolc Mathematical Notes*, 19, 527.

SIMS, C. (1971): "Discrete Approximations to Continuous Time Distributed Lags in Econometrics," *Econometrica*, 39, 545–563.

STOCK, J. AND M. WATSON (2012): "Disentangling the Channels of the 2007-2009 Recession," *Brookings Papers on Economic Activity*, Spring, 81–135.

SWANSON, N. R. AND C. W. J. GRANGER (1997): "Impulse Response Functions Based on a Causal Approach to Residual Orthogonalization in Vector Autoregressions," *Journal of the American Statistical Association*, 92, 357–367.

UHLIG, H. (2005): "What are the effects of monetary policy on output? Results from an agnostic identification procedure," *Journal of Monetary Economics*, 52, 381 – 419.

WOOLDRIDGE, J. M. (2010): *Econometric Analysis of Cross Section and Panel Data*, vol. 2nd ed, The MIT Press.

WRIGHT, J. H. (2012): "What does Monetary Policy do to Long-term Interest Rates at the Zero Lower Bound?*," *The Economic Journal*, 122, F447–F466.

Appendix

A Detailed Econometric Framework

This section formalizes and generalizes the description of our approach. First, we presents the econometric framework used for the analysis, then Sections A.1 to A.3 describe the different steps of the methodology in detail. Conditions and tests for the validity of identification performed in the HF-VAR, as well as the correct procedure to conduct causal inference on the LF variables are discussed.

The framework that we consider is the following: i) the vector of innovations that drives the variables can be partitioned as $\varepsilon = \left[\varepsilon^1 \, \varepsilon^{\bar{1}}\right]$, where ε^1 denotes the structural shock of interest and $\varepsilon^{\bar{1}}$ is the vector that comprises all the other shocks of the system; ii) the objective is to identify the effect the innovation ε^1 on a vector of endogenous variables y; iii) ε^1 can be recovered at the frequency $t = 1, 2, ..., T$ higher than the frequency τ at which y is observable ($\tau = m, 2m, ..., T$ where $m > 1$ is the frequency mismatch). Our identification approach is comprised of three steps: 1) identify ε_t^1 in a daily VAR; 2) aggregate ε_t^1 transforming it into ε_τ^1; and 3) estimate a VAR at lower frequency on y_τ and use ε_τ^1 as a proxy (internal or external instrument) to estimate the causal impact effect of ε^1 on y.

Consider a vector of n time series modeled as a causal and covariance stationary SVAR of lag-length p:

$$y_t = A_1 y_{t-1} + ... + A_p y_{t-p} + B\varepsilon_t \tag{A.1}$$

where ε_t is a vector of stochastic innovations and B is a $n \times n$ matrix whose coefficients determine how ε_t contemporaneously affects the variables y_t. Such a process can be expressed via compact notation through the polynomial $A(L) = \mathscr{I} - A_1 L - A_2 L^2 - ... - A_p L^p$:

$$A(L) y_t = B\varepsilon_t \tag{A.2}$$

where L is the lag operator such that $L^i y_t = y_{t-i}$. The SVAR is not observable itself, but corresponds to a multiplicity of reduced-form VAR representations of the form:

$$A(L) y_t = u_t \tag{A.3}$$

In what follows we focus exclusively on the problem of identification of matrix B under temporal aggregation. Temporal aggregation can be expressed as a two-step filter. First, the data are made observable only once every m periods, which represents the frequency mismatch, via the filter $D(L) = \mathscr{I} + D_1 L + D_2 L^2 + ... + D_{pm-p} L^{pm-p}$. The specification of $D(L)$ has to be such that the elements of $D(L)A(L)$ are powers of L^m, meaning that only the observable data points enter the transformed process. The conditions for the existence of such a filter, as well as the values taken by the matrices D_i are derived in Marcellino (1999). The second filter, denoted by $W(L)$, depends on the temporal aggregation scheme considered; skip-sampling (or point-in-time sampling) is usually applied to stock variables (e.g. prices) whereas averaging is typically applied to flow variables (e.g. volumes). In the former case, $W(L)$ does not modify the original data, i.e. $W(L) = \mathscr{I}$. For example, consider the *VAR(1)* process analyzed in the stylized example: $y_t = A_1 y_{t-1} + B\varepsilon_t$ and $m = 2$. The filter $D(L)W(L) = \mathscr{I} + A_1 L$ transforms the original process into $(\mathscr{I} + A_1 L)\, y_t = (\mathscr{I} + A_1 L)\, A_1 y_{t-1} + (\mathscr{I} + A_1 L)\, B\varepsilon_t$, which can be rearranged as $y_\tau = A_1^2 y_{\tau-1} + B\varepsilon_t + A_1 B\varepsilon_{t-1}$. In the averaging case $W(L) = \mathscr{I} + L + L^2 + ... + L^{m-1}$ and for the specific process under consideration $W(L) = \mathscr{I} + L$ and $D(L)W(L) = (\mathscr{I} + A_1 L)\,(\mathscr{I} + L)$. The temporally aggregated process would then become $\bar{y}_\tau = A_1^2 \bar{y}_{\tau-1} + B\,(\varepsilon_t + \varepsilon_{t-1}) + A_1 B\,(\varepsilon_{t-1} + \varepsilon_{t-2})$ with $\bar{y}_\tau = y_\tau + y_{\tau-1}$.

The typical object of interest in the SVAR literature are the dynamic effects of the innovations ε_t on y_{t+k} where $k \in \mathbb{N}$ represents the horizon. The impact effects of the innovations are defined as: $\Theta_{0,t} = \mathbb{E}_{\mathscr{F}_{t-1}}\,[y_t/\varepsilon_t = 1]$ - $\mathbb{E}_{\mathscr{F}_{t-1}}\,[y_t/\varepsilon_t = 0]$ with the information set $\mathscr{F}_{t-1} = \{y_{t-1}, ..., y_{t-p}\}$. Conversely, there are two possible definitions of temporally aggregated IRFs $\Theta_{0,\tau}$ according to the information set considered at LF. A first possibility employs the relevant information set for the LF representation; the information set is $\mathscr{F}_{\tau-1} = \{y_{\tau-1}, y_{\tau-2}, ...\}$ and consequently $\Theta_{0,\tau} = \mathbb{E}_{\mathscr{F}_{\tau-1}}\,[y_\tau/\varepsilon_t = 1, ..., \varepsilon_{t-m+1} = 1]$ - $\mathbb{E}_{\mathscr{F}_{\tau-1}}\,[y_t/\varepsilon_t = 0, ..., \varepsilon_{t-m+1} = 0]$. This definition considers all the innovations occurring between $\tau - 1$ and τ. Alternatively, the LF-IRFs can be defined directly via the temporal aggregation filter, implicitly using the information set $\mathscr{F}_{t-1}$ (i.e. using the HF information set). In this case, the LF-IRFs are defined directly as $\Theta_{0,\tau} = D(L)W(L)\Theta_0^t$. Thus, this choice is key to determining how to correctly use the HF shocks ε_t. While both definitions are formally correct, we regard the first as the most interesting from a macroeconomic perspective.

A.1 First Step: Identification at High-Frequency

The first step concerns the identification of a shock or proxy ε_t^1 at the high-frequency $t = 1, 2, ..., T$. The HF-VAR has to be specified to achieve informational sufficiency, meaning that the shock ε_t^1 can be recovered as a linear combination of the reduced form residuals if the appropriate identification strategy is applied. Thus, a set of variables Ω_t are included to achieve this goal as in Eq.(A.4). $\mathring{A}(L)$ and $\mathring{B}$ denote respectively the autoregressive matrix and the impact matrix that characterize the HF system $\begin{bmatrix} \Omega & y_t^1 \end{bmatrix}'$:

$$\mathring{A}(L) \begin{bmatrix} \Omega_t \\ y_t^1 \end{bmatrix} = \mathring{B} \begin{bmatrix} \varepsilon_t^\Omega \\ \varepsilon_t^1 \end{bmatrix} \tag{A.4}$$

After estimating the reduced form HF-VAR, one can apply any of the different identification strategies previously used by the literature to identify the structural shock of interest from the reduced form residuals (i.e. recursive, sign, or narrative restrictions).[1a] The identification at HF recovers z_t^1, a potentially noisy measure of the shock ε_t^1 as described in Eq.(A.5):

$$z_t^1 = \varepsilon_t^1 + w_t \tag{A.5}$$

[1a]See Ramey (2016) or Kilian and Lutkepohl (2017) for a summary of different identification strategies

where w_t is a measurement error such that $\mathbb{E}[w_t w_t] = \sigma_w^2$ and $\mathbb{E}\left[\varepsilon_t^1 w_t\right] = 0$. Assuming that $\sigma_w^2 = 0$ implies that we are recovering the true shock. In what follows, we use z_t and ε_t interchangeably.

A.2 Second Step: Aggregation and Information Sufficiency Test

In the second step, we aggregate the shocks identified at high-frequency and test for their partial invertibility. Considering that the IRFs are defined in terms of the LF information set $\mathscr{F}_\tau$. This means that the evaluation of the IRFs takes into account all the innovations occurring between $\tau - 1$ and τ. Let $\{\Gamma_{m-1}, \Gamma_{m-2}, ..., \Gamma_0\}$ be the sequence of the LF impact effects associated respectively to the shocks $\{\varepsilon_{t-m+1}, \varepsilon_{t-m+2}, ..., \varepsilon_t\}$. For example, in the stylized process of Section 2.1 $\Gamma_0 = \Theta_{0,t} = B$ and $\Gamma_1 = \Theta_{0,t} + A_1\Theta_{0,t} = (\mathscr{I} + A_1)B$. Then the aggregated response at LF is given by $\Theta_{0,\tau} = \sum_{i=0}^{m-1} \Gamma_j$ and can be recovered according to *Proposition I* presented in Section 2.2.

The crucial property that distinguishes structural shocks from reduced form residuals, i.e. their orthogonality, is preserved when the HF shocks are aggregated to LF according to *Lemma IA*. It is worth stressing that the proposition provides sufficient conditions for the orthogonality of the LF shocks but not a necessary condition. Furthermore, *Lemma IIA* expresses that the lack of autocorrelation of shocks is preserved by the aggregation. Proofs are reported in Section B.4.

Lemma IA. *Given the shocks* $\varepsilon_t = \left[\varepsilon_t^1, \varepsilon_t^{\bar{1}}\right]$ *identified in the HF-VAR that by construction satisfy* $\mathbb{E}\left[\varepsilon_t^1 \varepsilon_t^{\bar{1}\prime}\right] = 0$, *if the HF-VAR approximates well enough the data generating process, then* $\mathbb{E}\left[\varepsilon_\tau^1 \varepsilon_\tau^{\bar{1}\prime}\right] = 0.$[2a]

An ancillary result regarding the autocorrelation of the aggregated shocks, which is based on Proposition IA, is contained in Lemma IA.

***Lemma II**A. Given the shocks ε_t identified in the HF-VAR, if the HF-VAR is a good enough approximation of the data generating process, then the shocks aggregated at the LF ε_τ do not display autocorrelation.*

A.3 Third Step: External Instruments and Identification of the LF-VAR

The last step is twofold. First, we estimate the LF-VAR of order p in Eq.(A.6) including the relevant HF variables (aggregated at LF) together with the macroeconomic variables of interest in the system:

$$y_\tau = \tilde{A_1} y_{\tau-1} + \tilde{A_2} y_{\tau-2} + ... + \tilde{A_p} y_{\tau-2} + u_\tau \tag{A.6}$$

where u_τ denotes the vector of reduced form residuals. Second, the causal impact effect of ε_τ^1 on y_τ is identified by employing z_τ^1 as external instrument. Assuming that the shock of interest ε_τ^1 is invertible, i.e. it can be expressed as a linear combination of the reduced form residuals u_τ, inference from the Proxy-SVARs is valid under three conditions:

[2a]Following Canova (2007), a HF-VAR is considered to approximate well enough the data generating process if the underlying MA representation of the process can be expressed as the linear combination of current and past values of y_τ using the HF-VAR.

i) exogeneity: $\mathbb{E}\left[\varepsilon_{\tau}^{\bar{1}} z_{\tau}^{1}\right] = 0$

ii) strength: $\mathbb{E}\left[\varepsilon_{\tau}^{1} z_{\tau}^{1}\right] \neq 0$

iii) limited lag-lead exogeneity: $\mathbb{E}\left[e_{\tau+j}^{\bar{1}} z_{\tau}^{1}\right] \neq 0$ for $j \neq 0$, where $e^{\bar{1}}$ denotes the subset of $\varepsilon_{\tau}^{\bar{1}}$ of non-invertible shocks.

See main text for a discussion of the appropriate estimation method of the IRFs, which depends on the invertibility and exogeneity of the proxy.

B Proofs

This section contains the proofs of *Proposition I* included in Section A.2. The proof relies on some intermediate results which are preliminary included in this section.

We first derive intermediate results on the aggregation of HF shocks under alternative temporal aggregation filters (B.1-B.2) and then demonstrate *Proposition I* (B.3). Finally, we report the proofs of the Lemmas previously stated in this Annex.

B.1 IRFs under Skip-Sampling

The skip-sampling case is simpler because the second step of the temporal aggregation process consists of the trivial filter $W(L) = \mathscr{I}$ that leaves the variables unaffected. Skip-sampling is usually applied by taking the last value: for example the last daily observation within the month. We focus on this skip-sampling scheme without loss of generality.[3a] Eq.(A.2) is modified by temporal aggregation as:

$$D(L)A(L)y_t = D(L)B\varepsilon_t \tag{A.7}$$

Under skip-sampling, the impact effect of ε on y is trivially given $\Theta_{0,\tau} = \Theta_{0,t} = B$. Suppose that the HF shocks are identified under the assumptions described in Equation A.4.

Proposition IA. *Given an underlying HF-VAR temporally aggregated via skip-sampling, the IRF $\Theta_{0,\tau} = B$ can be recovered by projecting the reduced form residuals estimated from the LF-VAR, u_τ, on the last HF shock within the LF period. Thus the correct filter $J(L)$ applied to ε_t is $J(L) = \mathscr{I}$ such that $\varepsilon_\tau = \varepsilon_t$ for $\{\tau, t\} = m, 2m, \ldots, T$.*

Proof. The correct impact matrix can be recovered simply projecting u_τ on $\varepsilon_{t,m-1}$. The LF reduced form residuals are given by:

$$u_\tau = D(L)u_t = D(L)B\varepsilon_t \tag{A.8}$$

[3a]Notice that the same results that we provide hold simply by using the shock corresponding to the skip-sampling scheme (e.g. take the first shock if skip-sampling is performed using the first HF value)

Recall that $D(L) = \mathscr{I} + D_1L + D_2L^2 + ... + D_{pm-p}L^{pm-p}$ always contains the identity matrix as first term. Thus

$$D(L)B\varepsilon_t = B\varepsilon_t + D_1B\varepsilon_{t-1} + D_2B\varepsilon_{t-2} + ... + D_{pm-p}B\varepsilon_{t-pm+p} \tag{A.9}$$

Independently of the values of D_i (that depend on the HF-VAR lag length p and frequency mismatch m) and considering that ε_t are uncorrelated, the following relationship holds:

$$Proj\,(u_\tau/\varepsilon_t) = B = \Theta_{0,\tau} \tag{A.10}$$

B.2 IRFs under Averaging

The averaging case is more complex as the second filter is $W(L) = \mathscr{I} + L + L^2 + ... + L^{m-1}$. We use the summing filter, which is equivalent to averaging up to a constant. Consequently, the system becomes

$$D(L)W(L)A(L)y_t = D(L)W(L)B\varepsilon_t \tag{A.11}$$

The HF impact effect is again $\Theta_0^t = B$ but the IRFs must be consistently temporally aggregated if we want to dispose of a reliable metric of comparison. Under linearity, the impact effect of ε on y at low-frequency is given by $\Theta_0^\tau = \Theta_0^t + \Theta_1^t + \ldots + \Theta_{m-1}^t$.

Proposition IIA. *Given an underlying HF-VAR temporally aggregated via averaging, the IRF $\Theta_{0,\tau}$ can be recovered by projecting the reduced form residuals estimated from the LF-VAR, u_τ, on the first HF shock within the LF period. Thus, the correct filter $J(L)$ applied to ε_t is $J(L) = L^{m-1}$ such that $\varepsilon_\tau = \varepsilon_{t-m+1}$ for $\{\tau, t\} = m, 2m, \ldots, T$.*

Proof. It is convenient to express the IRFs at horizon k employing the companion form of the VAR:

$$x_t = F\,x_{t-1} + \eta_t \tag{A.12}$$

where

$$F = \begin{bmatrix} A_1 & A_2 & A_3 & \ldots & A_p \\ \mathscr{I} & 0 & 0 & \ldots & 0 \\ 0 & \mathscr{I} & 0 & \ddots & 0 \\ \vdots & \ddots & \ddots & \ddots & \vdots \\ 0 & 0 & \ldots & \mathscr{I} & 0 \end{bmatrix} \tag{A.13}$$

$$\eta_t = \begin{bmatrix} u_t = B\varepsilon_t \\ 0 \\ 0 \\ \vdots \\ 0 \end{bmatrix} \tag{A.14}$$

$x_t = \left[y_t\, y_{t-1} \ldots y_{t-p-1}\right]'$ where p is the lag length of the VAR and $x_{t-1} = Lx_t$, such that x_t is a vector of $(n \times p)$ variables. Then, considering that $\Theta_{0,t} = B$, the dynamic effects can be written as well in companion form as:

$$\tilde{\Theta}_{k,t} = F\tilde{\Theta}_{k-1,t} \qquad k \in \mathbb{N} \tag{A.15}$$

We are interested in the matrix that contains the impulse response at horizon k, positioned in $\tilde{\Theta}_{k,t}(1,1) = \Theta_{k,t}$.

Kalman (1982) and Kilic (2007) have shown that this formulation of the IRFs corresponds to a generalized Fibonacci sequence of order p in the matrices $A_1, A_2, ..., A_p$ (also referred to as generalized order p Fibonacci polynomial), denoted by $S_p(A_1, ..., A_p)$. Given an initial condition B, $S_p(A_1, ..., A_p)$ generates a sequence whose elements are a linear combination of the previous terms in the sequence weighted by $A_1, ..., A_p$. A_1 is the weight associated to the previous element of the sequence whereas A_p multiplies the *p-th* previous element. For instance $S_0 = B$, $S_1 = A_1 S_0$, $S_2 = A_1 S_1 + A_2 S_0$, and so on and so forth.

Following Marcellino (1999), we define the following vector of matrices D^v and A^v with dimension $1 \times pm$, where m denotes the frequency mismatch, and the matrix of matrices G with dimension $(pm - p) \times pm$

$$\begin{aligned}
D^v &= \left(D_1, D_2, ..., D_{pm-p}\right) \\
A^v &= \left(A_1, A_2, ..., A_p, 0, ..., 0\right) \\
G &= \begin{bmatrix}
-\mathscr{I} & A_1 & A_2 & \dots & A_{p-1} & A_p & 0 & 0 & \dots & & & 0 \\
0 & -\mathscr{I} & A_1 & A_2 & \dots & A_{p-1} & A_p & 0 & 0 & \dots & & 0 \\
0 & 0 & -\mathscr{I} & A_1 & A_2 & \dots & A_{p-1} & A_p & 0 & 0 & \dots & 0 \\
\vdots & 0 & 0 & -\mathscr{I} & \ddots & \ddots & \ddots & \ddots & \ddots & \ddots & \ddots & \vdots \\
 & \vdots & 0 & \ddots & \ddots & & & & & & & \\
 & & \vdots & \ddots & & & & & & A_{p-1} & A_p & 0 \\
0 & 0 & & \cdots & \dots & 0 & -\mathscr{I} & A_1 & A_2 & \dots & A_{p-1} & A_p
\end{bmatrix}
\end{aligned}$$

Marcellino (1999) shows that $D(L)$ exists if $|G_{-k}| \neq 0$ where G_{-k} corresponds to the G matrix whose columns multiple of m (i.e. $m, 2m, ...$) have been deleted. Then, the coefficients of $D(L)$ are given by $D^v = -A^v_{-k}(G_{-k})^{-1}$. The intuition is that they have to be such that only the powers of L^m can have a coefficient different from 0 (since the variables are unobserved between L^m and L^{2m}).

Since our goal is the identification of the impact effect $\Theta_{0,\tau}$, we focus exclusively on the upper square block of G of dimension $(m-1) \times (m-1)$, denoted by $\tilde{G}$. Consistently, we denote $\tilde{A}^v$ and $\tilde{D}^v$ the vector containing the first $m-1$ elements of the original A^v and D^v vectors. The matrix $\tilde{G}$ is a very special matrix, being a Toeplitz upper triangular matrix with main diagonal $-\mathscr{I}$. Sahin (2018) shows that, if invertibility is satisfied, the inverse of this class of matrices, denoted in our case by $\tilde{G}^{-1}$, contains the elements of the Fibonacci sequence in the matrix $A_1, ..., A_p$. By considering that

$$\tilde{D}^v = -\tilde{A}^v \tilde{G}^{-1}$$

it follows that the elements of $\tilde{D}^v$ correspond to $S_p(A_1, ..., A_p)$, the Fibonacci sequence in

$A_1,...,A_p$. Disregarding the initial effect B, the IRFs $\Theta_0^t,\Theta_1^t,...,\Theta_{m-1}^t$ and the elements of the temporal aggregation filter $\mathscr{I},D_1,D_2,...,D_{m-1}$ are generated by the same generalized Fibonacci sequence $S_p\left(A_1,...,A_p\right)$.. Finally, the temporally aggregated residuals are given by $u_\tau = D(L)W(L)\varepsilon_t$. The filter $W(L) = \mathscr{I} + L + ... + L^{m-1}$ implies the first shock within the LF period, i.e. $\varepsilon_{t,1}$, enter u_τ through all the terms $\Theta_{0,t},\Theta_{1,t},...,\Theta_{m-1,t}$ and thus recovers the correct $\Theta_{0,\tau} = \Theta_{0,t} + \Theta_{1,t} + ... + \Theta_{m-1,t}$. Thus, the projection of u_τ on ε_{t-m+1} yield the correct IRFs $\Theta_{0,\tau}$:

$$Proj\left(u_\tau/\varepsilon_{t-m+1}\right) = \Theta_{0,\tau} \tag{A.16}$$

■

B.3 Proof of Proposition I

This proof follows directly from the proof of *Proposition II*A, which showed that the IRFs $\Theta_0^t,\Theta_1^t,...,\Theta_{m-1}^t$ and the elements of the temporal aggregation filter $\mathscr{I},D_1,D_2,...,D_{m-1}$ are generated by the same generalized Fibonacci sequence $S_p\left(A_1,...,A_p\right)$ (for the given initial condition B). This implies that the contemporaneous LF impact effect of the shock ε_{t-m+1} is given by the sum of the first m elements of the Fibonacci sequence denoted by S_p^{m-1}. The effect of ε_{t-m+2} is given by the sum of the elements in S_p^{m-2} and so on and so forth until the last HF shock ε_t that impacts through $B = S_p^0$. In the case of skip-sampling, $\Theta_{0,\tau} = S_p^0$ whereas in the case of averaging $\Theta_{0,\tau} = \sum_{i=1}^m BS_p^i$. In both cases, projecting u_τ on ε_τ recovers $\Theta_{0,\tau}$.

Consider that $\varepsilon_\tau = \left(I + L + ... + L^{m-1}\right)\varepsilon_t$ and $u_\tau = D(L)W(L)B\varepsilon_t$ where the first m elements of $D(L)$ are $I + D_1L + ... + D_{m-1}L^{m-1}$ and $W(L) = \left(I + L + ... + L^{m-1}\right)$. By the result in *Proposition IIA* on the equivalence between the recursive definition of IRFs and the matrices $\mathscr{I},D_1,D_2,...,D_{m-1}$, it is straightforward to verify that:

$$\begin{aligned} Proj\left(u_\tau/\varepsilon_\tau\right) &= \mathbb{E}\left[\left(\varepsilon_\tau\varepsilon_\tau'\right)^{-1}\right]\mathbb{E}\left[\varepsilon_\tau u_\tau'\right] \\ &= \sum_{i=0}^{m-1}\Gamma_j = \Theta_{0,\tau} \end{aligned}$$

■

B.4 Proof of Lemmas IA and IIA

Proof of Lemma IA. Under standard assumptions, the Wold Representation Theorem implies that the innovations $\{\eta_t\}$ of a time series process $\{y_t\}$ are white noise. If the underlying process is well approximated by a $VAR(p)$, then this property extends to the residuals estimated by the VAR

$\{u_t\}$ (Canova, 2007 Ch.4). Thus, it holds that:

$$\begin{aligned} \mathbb{E}\left[u_t u_{t-j}^{'}\right] &= \Sigma \quad \text{for } j = 0 \\ \mathbb{E}\left[u_t u_{t-j}^{'}\right] &= 0 \quad \text{for } \forall j \neq 0 \end{aligned} \tag{A.17}$$

Property A.17 extends to the structural shocks $\varepsilon_t = B^{-1}u_t$ because they are a linear transformation of u_t: $\mathbb{E}\left[\varepsilon_t \varepsilon_{t-j}^{'}\right] = 0$ for $j = 1, \ldots, p$. Furthermore, it follow from identification that $\mathbb{E}\left[\varepsilon_t \varepsilon_t^{'}\right] = 0$. Without loss of generality, partition $\varepsilon_t = \left[\varepsilon_t^1 \; \varepsilon_t^{\bar{1}}\right]^{'}$ and define $\varepsilon_t^1 = b_{\bullet 1} u_t$ and $\varepsilon_t^{\bar{1}} = b_{\bullet \bar{1}} u_t$, where $b_{\bullet 1}$ denotes the first column of the impact matrix B and $b_{\bullet \bar{1}}$ the remaining columns. Combining the previous properties, it holds

$$\mathbb{E}\left[\varepsilon_t^1 \varepsilon_{t-j}^{\bar{1}'}\right] = 0 \;\; \text{for } \forall j \tag{A.18}$$

Consider now the aggregated shocks $\varepsilon_\tau = \sum_{i=0}^{m-1} \varepsilon_{t-i}$. To evaluate $\mathbb{E}\left[\varepsilon_\tau^1 \varepsilon_\tau^{\bar{1}'}\right]$, we need to compute the correlations between all the elements that are summed into ε_τ^1 and $\varepsilon_\tau^{\bar{1}}$, which are all null by (A.18). Thus, $\mathbb{E}\left[\varepsilon_\tau^1 \varepsilon_\tau^{\bar{1}'}\right] = 0$.■

Proof of Lemma IIA. Notice that $\varepsilon_{1t} = b_{1\bullet} u_t$ and $\varepsilon_{2t} = b_{2\bullet} u_t$ are two structural shocks obtained as linear combination of the residuals u_t, where $b_{\bullet 1}$ denotes the first column of the impact matrix B and $b_{\bullet \bar{1}}$ the remaining columns Based on Proposition IA, $\mathbb{E}\left[u_t u_{t-j}^{'}\right] = 0$ for $\forall j \neq 0$. This extends to $\mathbb{E}\left[\varepsilon_{1t} \varepsilon_{it-j}\right]$ for $\forall j \neq 0$ and $i = 1, 2$ by the property of the linear operator $b_{1\bullet}$. Plus, $\mathbb{E}\left[\varepsilon_{1t} \varepsilon_{2t}\right] = 0$ since they are structural shocks. Thus, each element of the sum in $\varepsilon_{1\tau}$ is uncorrelated to the elements in $\varepsilon_{2\tau}$ and so it is their sum. Based on Lemma IA, the lack of autocorrelation of order 1 holds. ■

C Additional Results on the Empirical Applications

C.1 Berger, Dew Becker, and Giglio (2020)

	(1)	(2)
	$\sum_{i=1}^{504} rv_{t+i}$	$\sum_{i=1}^{504} rv_{t+i}$
$v_{1,t}$	0.68***	0.47***
	(0.01)	(0.01)
Additional predictors	$\times$	✓
Observations	7181	7181
$R^2 - adj$	0.46	0.68

Table A1: Predictive Regressions
Predictive regressions of 6-months rv. (1) includes only v_1 as regressor; (2) includes all variables included in the full VAR specification.
*Standard errors are reported in parenthesis, *** p<0.01, ** p<0.05, * p<0.1.*

	(1) Bivariate DVAR		(2) Full DVAR	
	ε_τ^u	ε_τ^{rv}	ε_τ^u	ε_τ^{rv}
$F - stat_{(8,338)}$	2.42**	1.10	1.10	0.41
$F - test\ (pval)$	0.01	0.36	0.38	0.89
Observations	347	347	347	347
$R^2 - adj$	0.06	0.07	0.01	0.01

Table A2: Invertibility Test - Uncertainty Shocks - BDG
The regressions include seven lagged factors from FRED-MD database, one lag of the dependent, and a constant. ε_τ^u and ε_t^{rv} denote the uncertainty and realized volatility shocks . F-stat denotes the value of the F test statistic and F-test is the pvalue on the joint test of all coefficients associated with the factors being 0.
*Standard errors are reported in parenthesis, *** p<0.01, ** p<0.05, * p<0.1.*

C.2 Caldara, Fuentes-Albero, Gilchrist and Zakrajsek (2016)

	$Unc-id1$	$Fin-id1$	$Unc-id2$	$Fin-id2$
$F-stat_{(8,326)}$	1.02	1.97	1.07	1.98
$F-test\,(pval)$	0.42	0.06	0.38	0.06
Observations	338	338	338	338
R^2-adj	0.01	0.02	0.01	0.02

Table A3: Invertibility Test - Uncertainty and Financial Shocks - CFGZ
*The regressions include seven lagged factors from FRED-MD database, one lag of the dependent, and a constant. Unc and Fin denote the uncertainty and financial shocks. F-stat denotes the value of the F test statistic and F-test is the pvalue on the joint test of all coefficients associated with the factors being 0. Standard errors are reported in parenthesis, *** p<0.01, ** p<0.05, * p<0.1.*

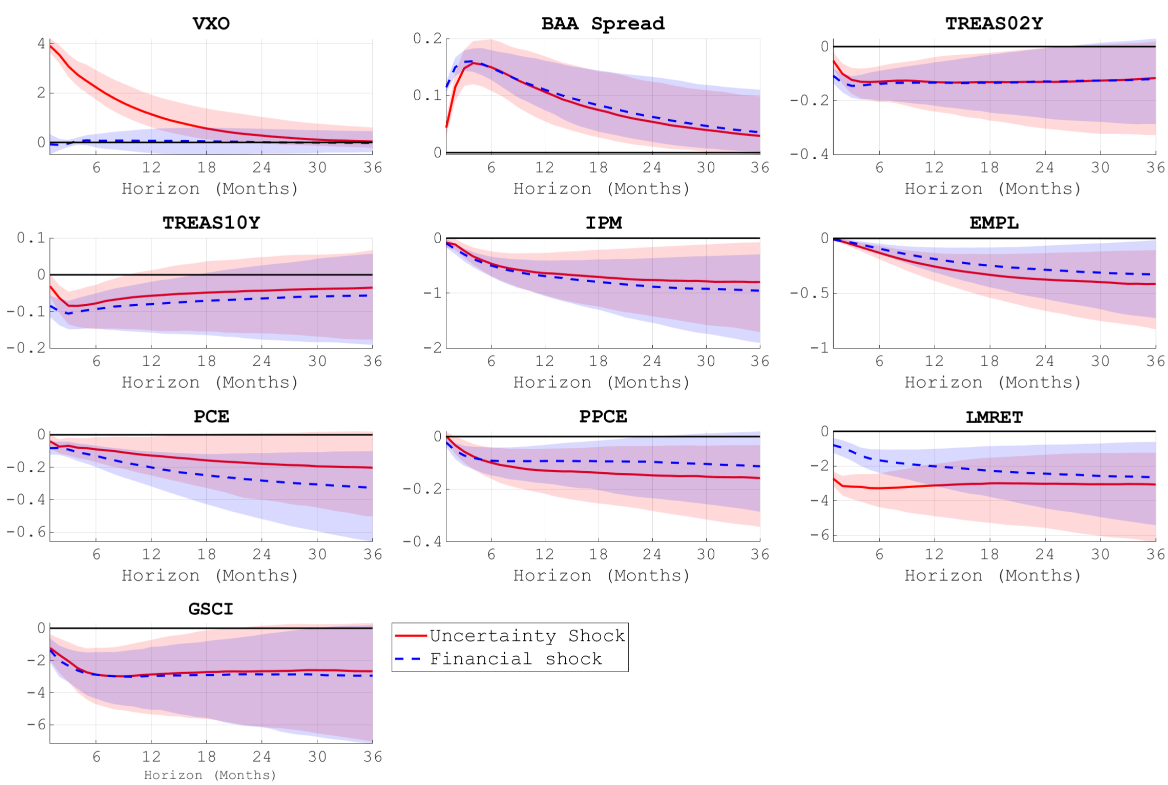

Figure A1: IRFs in monthly VAR of CFGZ with BAA spread under VXO-BAA identification

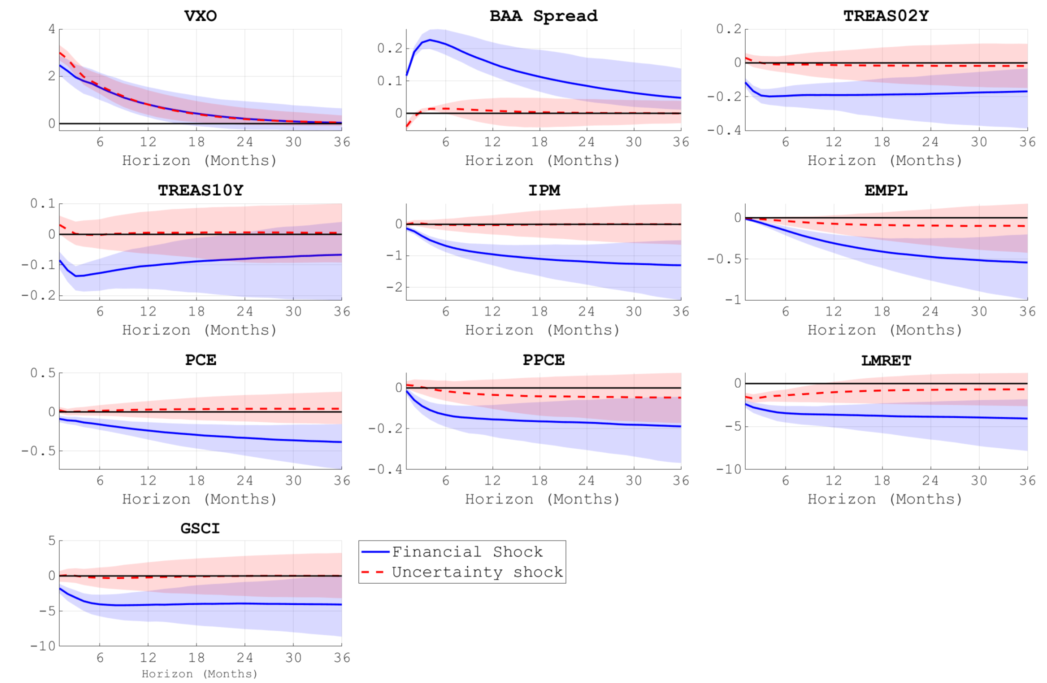

Figure A2: IRFs in monthly VAR of CFGZ with BAA spread under BAA-VXO identification

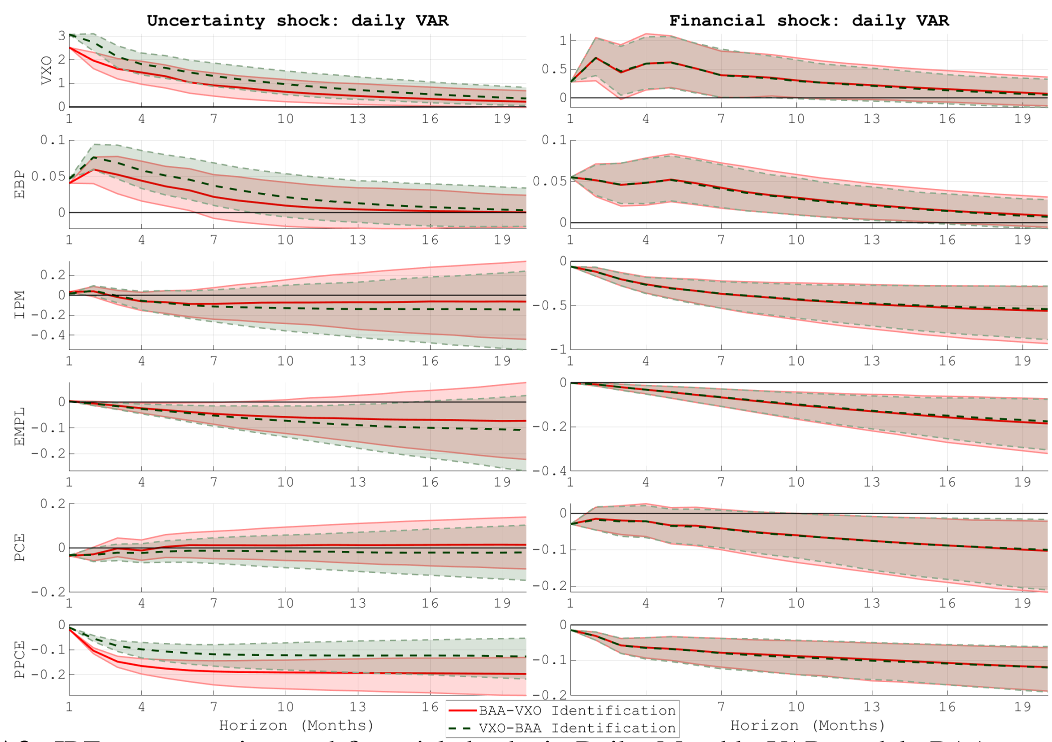

Figure A3: IRFs to uncertainty and financial shocks in Daily+Monthly VAR model. BAA spread in the daily VAR and EBP in the monthly VAR

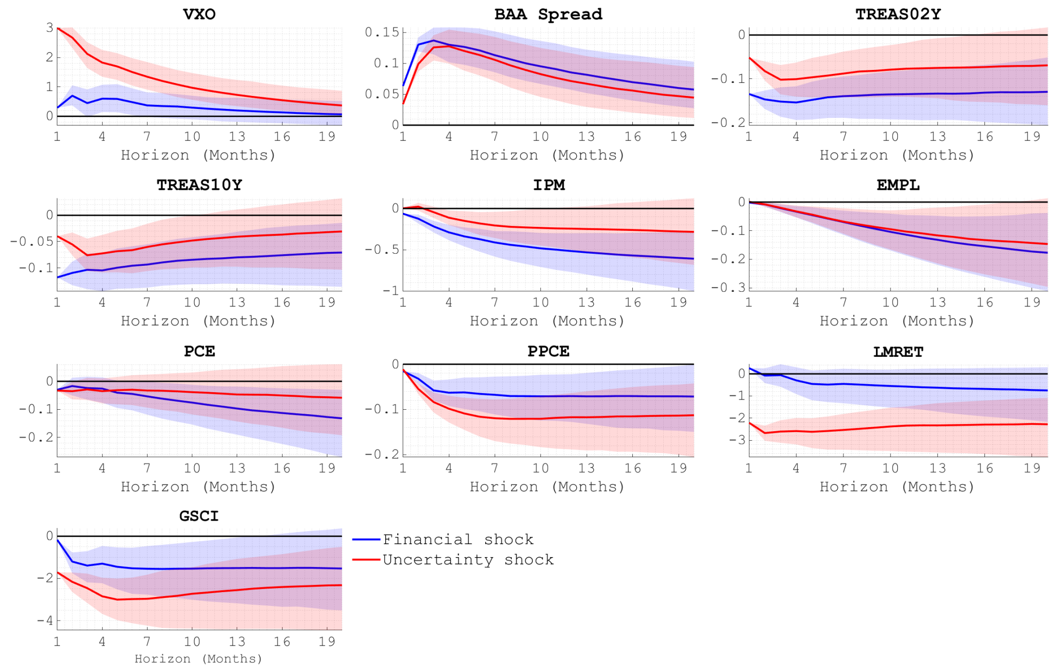

Figure A4: IRFs in baseline Daily+Monthly VAR of CFGZ VXO-BAA identification

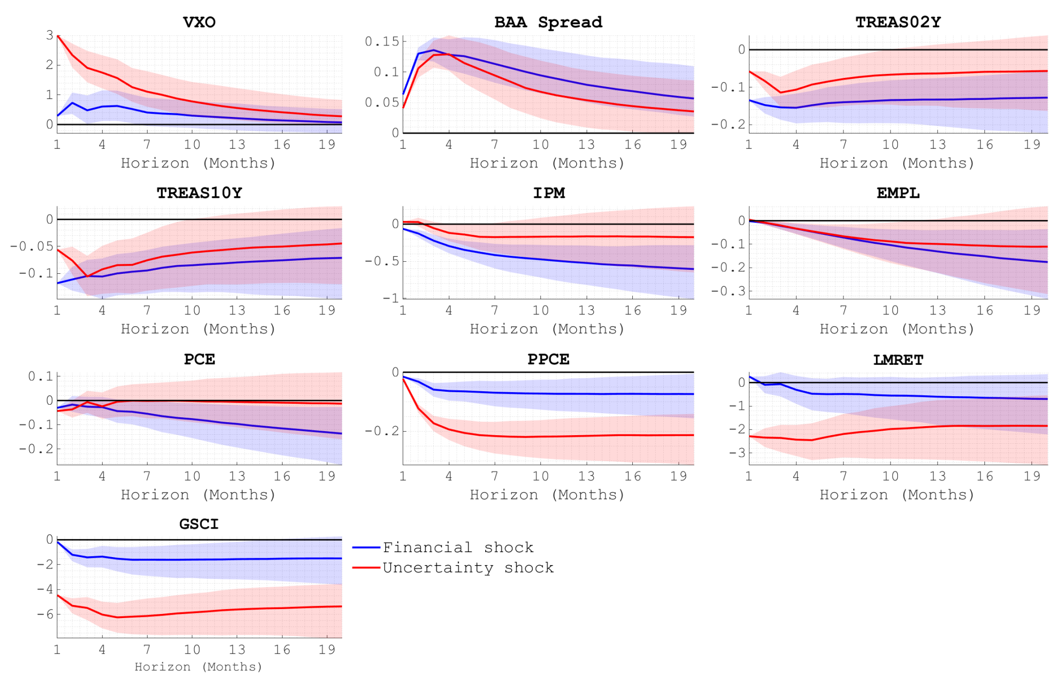

Figure A5: IRFs in baseline Daily+Monthly VAR of CFGZ under BAA-VXO identification