The macroeconomic implications of financial volatility: the role of uncertainty and risk aversion*

Brendan Berthold†

December 2021

Abstract

Financial volatility shocks can either be due to a physical change in risk, that is an uncertainty shock, or changes in agents preferences towards risk, that is a risk aversion shock. We identify these two shocks in a proxy SVAR setting using set identification. In particular, we use volatility events and a narrative approach to develop two new proxies that can be used jointly to identify exogenous variations in uncertainty and risk aversion. We find that the Great Financial Crisis is more associated with risk aversion shocks and that the COVID recession coincides with large uncertainty shocks.

Keywords. Financial volatility, Uncertainty, Risk Aversion, Set Identification, Proxy SVAR

JEL Codes. E44, E32, D80, G1

*I am grateful to Jean-Paul Renne, Mathieu Grobéty, Philippe Bacchetta, Pascal St-Amour, and seminar participants at the University of Lausanne and the Study Center Gerzenesee for helpful comments. Any remaining errors are my own.

†University of Lausanne, CH-1015 Lausanne, brendan.berthold@unil.ch

1 Introduction

Financial volatility is a central concept in finance and the study of its macroeconomic effects has attracted a lot of interest from the literature. The VIX and other popular risk neutral measures of financial volatility depend on both the physical risk in the economy and the compensation required by investors to bear this risk, that is a risk premium term. In the macroeconomic literature, the risk component is often referred to as "uncertainty" (Bloom (2014)).[1] On the other hand, the risk premium component is often associated with (stochastic) risk aversion (e.g. Bekaert, Hoerova, and Lo Duca (2013); Drechsler and Yaron (2011); Bekaert, Engstrom, and Xu (2021)). In this context, variations in expectations of financial volatility can stem from changes in uncertainty, risk aversion, or a combination of both. The aim of this paper is to identify both types of shocks in order to shed light on the fundamental drivers of financial volatility and to better understand their macroeconomic implications. Intuitively, we want to know whether the recessionary effect of volatility is either due to the *physical* change in risk or to sudden shifts in the investors' preferences towards risk? The answer to these questions is critical to improve our understanding of the interlinkages between financial markets and the real economy but also challenging in terms of identification because risk aversion and uncertainty are likely to be closely related.

So far, the literature has mostly studied the macroeconomic effect of uncertainty shocks without explicitly distinguishing between the two concepts. For instance, Bekaert, Hoerova, and Lo Duca (2013) argue that it is not clear to which extent the results from the seminal paper of Bloom (2009) on the effect of uncertainty shocks reflect risk aversion rather than uncertainty shocks. We argue that this is an important shortcoming of existing evidences. For example, an increase in uncertainty is likely to affect risk aversion. Not controlling for

[1] Strictly speaking, risk is conceptually different than uncertainty as defined by Franck Knight. In his view, risk describes situations with a known probability distribution over a set of events; tossing a fair coin with outcome 1 and 2 is riskier than getting 1.5 with certainty. On the other hand, uncertainty is best understood as "peoples' inability to forecast the likelihood of events happening." For example, the number of coins ever produced by mankind is uncertain. This finer distinction is typically ignored in the literature such that uncertainty is best understood as a mixture of risk and "true" Knightian uncertainty (Bloom (2014)).

changes in uncertainty when identifying risk aversion shocks may thus result in mistakenly interpreting variations in uncertainty as shocks to risk aversion. On the other hand, to the extent that market based measures of uncertainty are likely to depend on risk aversion through the risk premia, it is very likely that part of the identified uncertainty shocks actually reflect risk aversion rather than uncertainty. More generally, while it is generally well established that financial volatility can have real effects on the economy, the ultimate source of these fluctuations is less clear. Our paper aims to fill this gap.

The distinction, however, appears to be necessary because uncertainty and risk aversion are different concepts with potentially different implications. In particular, uncertainty relates to the agents' *perceptions* or *beliefs* about future risk in the economy. On the other hand, (aggregate) risk aversion refers to the marginal investor's *preferences* towards said risk. For instance, the recent COVID pandemic has generated an immense amount of risk by increasing the probabilities of extreme events; this corresponds to an uncertainty shock. At the same time, the pandemic may also have impacted the marginal investor's preferences towards risk; this corresponds to a risk aversion shock. The literature investigating the channels through which uncertainty can have real effects emphasize a real-option channel; when uncertainty rises, firms have higher incentives to delay investment or hiring in order to avoid a costly mistake. This type of mechanism dates back to Bernanke (1983) and was formally analysed in Bloom (2009). For risk aversion, the emphasis is often on financial channels, for example through increases in risk premia that raise the cost of equity finance and ultimately reduce investment and output (e.g. Christiano, Motto, and Rostagno (2014)).

The distinction between uncertainty and risk aversion shocks is challenging from an econometric perspective because both types of shocks are likely to be correlated. In particular, it is likely that there exist important feedback effects between uncertainty and risk aversion such that identifying exogenous variations in one or the other may reveal particularly difficult. The identification of these shocks is challenging because they are likely to be correlated. For instance, an important literature has documented how risk aversion can vary *endogenously* following variations in risk (e.g. Heaton and Lucas (2000)).[2] Similarly, risk aversion could also affect the physical risk in the economy; in an Epstein-Zin world, the preference over the early resolution of risk can be affected by shifts in risk aversion, thereby impacting the uncertainty in the economy. The first contribution of this paper is to tackle this issue by relying on the set identification approach developed in Piffer and Podstawski (2017). Unlike traditional external instrument identification, set identification allows a given proxy to be correlated with more than one shock. In our context, set identification allows the uncertainty and risk aversion proxies to be correlated with both uncertainty and risk aversion shocks. Identification is then achieved by making the reasonable assumption that the uncertainty shock is more correlated with the uncertainty proxy than with the risk aversion proxy. Similarly, we require that the risk aversion shock is more correlated with the risk aversion proxy than with the uncertainty proxy. Given two valid proxies, the approach allows to set-identify uncertainty and risk aversion shocks with a minimal set of assumptions while remaining agnostic about the feedback effects between uncertainty and risk aversion. To the extent that these feedback effects are likely to be complex, having this flexibility appears as a key advantage of our approach.

The second contribution of this paper is to develop two new proxies that can be used to set-identify both types of shocks. Empirically, we isolate arguably exogenous variations in uncertainty and risk aversion using an event-study framework. Intuitively, we consider a variety of events that are likely to have coincided with important variations in volatility. These volatility events are selected such that they are independent from other macroeconomic shocks such that the variations that we observe can credibly be attributed to the event only. The usage of events to identify exogenous variations has long been used in the literature (e.g. Kuttner (2001); Gürkaynak, Sack, and Swanson (2004)). Uncertainty shocks are then defined

[2]If preferences exhibit CRRA, then risk aversion is a decreasing and convex function of the endowment, which is a sufficient condition for risk aversion to increase with background risk (Eeckhoudt, Gollier, and Schlesinger (1996)).

as exogenous variations in the (physical) expected volatility of stock returns around these events and are recovered empirically using a slightly adjusted GARCH(1,1) model in the spirit of Berger, Dew-Becker, and Giglio (2019) with the VIX as the dependent variable. Risk aversion shocks are recovered empirically using the daily variations of the time-varying risk aversion measure developed in Bekaert, Engstrom, and Xu (2021). The two proxies are then built by summing up those high-frequency variations at the monthly frequency. Identification is achieved in a proxy SVAR using the two proxies as external instruments. The identifying assumptions are i) that the two proxies are exogenous with regards to other macroeconomic shocks (that is all shocks except uncertainty and risk aversion) and ii) that the uncertainty (risk aversion) shocks are more correlated with the uncertainty (risk aversion) proxy than with the other proxy. The key advantage of the set identification approach is that it allows to flexibly identify two types of shocks simultaneously without having to make an inadequate strict exogeneity assumption. More generally, our approach highlights how set identification in a proxy SVAR setting can be used to identify structural shocks for which it is difficult to find purely exogenous proxies and as such could open up new interesting research questions.

To test our framework, we estimate a proxy SVAR using monthly US data spanning 1992M2 to 2021M2 and set-identify financial uncertainty and risk aversion shocks using our two newly developed proxies. We find that set-identified financial uncertainty shocks lead to a significant decline in output and asset prices but appear to be relatively short-lived, in line with theoretical predictions (e.g. Bloom et al. (2018)) and recent empirical results (e.g. Piffer and Podstawski (2017); Ludvigson, Ma, and Ng (2021)). Risk aversion shocks appear to be damaging to output and asset prices but are estimated less precisely. A forecast error variance decomposition suggests that uncertainty and risk aversion shocks are important drivers of both real and financial variables as they can explain significant shares of the variation in stock prices and output. Using historical decompositions, we find that risk aversion shocks can explain the dynamics around the Great Financial Crisis (GFC) while uncertainty does not appear to have played a major role. The result are consistent with the massive increase in risk premia that coincided with the GFC. For the COVID recession, we find that financial uncertainty shocks were important drivers of the initial decline in output and acted as a drag on the subsequent economic recovery, in line with recent evidences from Baker et al. (2020). We further find that, albeit smaller, risk aversion shocks also played a role, most notably for asset prices. Finally, we find that the effect of risk aversion shocks display significant persistence and can partly explain the boom-and-bust pattern of asset prices around the two financial crises in our sample. The pattern is less clear before the COVID recession. These findings are consistent with financial crises resulting from waves of optimism and pessimism driven by shifts in sentiment. On the other hand, the effect of uncertainty shocks appear to be more transitory, in line with theoretical predictions. Overall, our results suggest that uncertainty and risk aversion are important drivers of real and financial variables and highlight the importance of distinguishing between the two types of shocks.

By providing a clear distinction between uncertainty and risk aversion shocks, this paper addresses the shortcomings of existing evidences by shedding light on the fundamental sources of financial volatility and by quantifying their macroeconomic implications. This empirical work thus contributes to improving our understanding of the relationship between financial volatility and macroeconomic outcomes. Overall, the results can help inform policymakers about the sources of macroeconomic fluctuations, and as such provide guidance with regards to the most effective policies.

The paper is structured as follows: Section 2 reviews the related literature. Section 3 details the set-identification approach. Section 4 and 5 develop the building of the uncertainty and risk aversion proxies Section 6 discusses the validity of the proxies. Section 7 estimates the proxy SVAR and compute the IRFs to an uncertainty and risk aversion shock. Section 8 investigates the importance of the two types of shocks for business fluctuations. Section 9 concludes.

2 Related literature

Since the Great Financial Crisis, it has become clear that the financial system cannot be ignored when looking at the sources of macroeconomic instability. Since then, an important literature has aimed to identify and quantify the contribution of so-called financial shocks to business cycle fluctuations (e.g. Gilchrist and Zakrajšek (2012); Gertler and Karadi (2015)). The recent COVID pandemic has renewed the interest on the effect of (financial) uncertainty on the real economy.

Typically, the literature on the identification of uncertainty shocks has relied on timing restrictions within vector-autoregressions (VAR) models (Bachmann, Elstner, and Sims (2013); Bloom (2009); Bekaert, Hoerova, and Lo Duca (2013)). While it is a natural starting point, there is no compelling theoretical motivation to justify this recursive ordering (see Ludvigson, Ma, and Ng (2021)). To answer these limitations, the literature is increasingly relying on an external instrument approach (e.g. Baker and Bloom (2013); Piffer and Podstawski (2017); Carriero et al. (2015)). This paper is part of this expanding literature. Our paper is most closely related to Piffer and Podstawski (2017) who identify financial uncertainty shocks using variations in the price of gold around volatility events and develop the set-identification methodology used in the present paper. Recently, Caggiano et al. (2021) set-identifies financial uncertainty shocks using reasonable economic a-priori on top of traditional sign-restrictions.

Our paper connects with the literature investigating the effect of financial volatility shocks on the macroeconomy (e.g. Bloom (2009); Bloom et al. (2018); Basu and Bundick (2017); Jurado, Ludvigson, and Ng (2015); Ludvigson, Ma, and Ng (2021)). An important difference of our approach is that we distinguish between different types of volatility, namely uncertainty and risk aversion. In Chiu et al. (2018), the authors decompose financial volatility in a long-run persistent component and a short-run transitory component and investigate their respective dynamic effects. They assume that the long-term component of volatility is associated with macroeconomic fundamentals and find that the short-run component is related to the transitory determinants of volatility, such as investor sentiment. In our paper, we focus on uncertainty and risk aversion as the main determinants of financial volatility. This interpretation is inspired by work of Bekaert, Hoerova, and Lo Duca (2013) who decompose the VIX index into two components, namely an "expected volatility" term which they interpret as a measure of uncertainty and a residual term–defined as the variance risk premium–which they interpret as a proxy for risk aversion. Their focus is on the feedback effects between uncertainty, risk aversion and monetary policy. Bekaert and Hoerova (2014) finds that the variance risk premium is a reliable predictor of stock returns while the uncertainty term (conditional volatility) is a better predictor of economic activity.

Our paper is also related to the literature on time-varying risk aversion. Campbell and Cochrane (1999) develop an asset pricing model featuring habit formation: as consumption moves closer to the habit (for instance during a business cycle trough), the curvature of the utility function increases and raises risk aversion endogenously. Longstaff and Wang (2012) and Gârleanu and Panageas (2015) feature time-varying risk premia that are driven by time-varying risk aversion in a heterogeneous agents framework. Generally speaking, the focus in these papers is on the asset pricing implications of time-varying risk aversion. Changes in risk apetite are seen as important drivers of asset prices and as potential drivers of the global financial cycles (Miranda-Agrippino and Rey (2021)). Our paper closely relates to the literature on the macroeconomic effects of risk aversion shocks. Risk premium shocks in the form of preference shocks are important drivers of macroeconomic fluctuations in Smets and Wouters (2007). More recently, Di Tella and Hall (2021) investigate the macroeconomic implications of time-varying risk premiums in a flexible-price model and find that such shocks can create inefficient recessions. Our paper adds to this literature by identifying both uncertainty and risk aversion shock within a common framework and by allowing to directly compare their respective effects.

3 Identification approach

3.1 Set identification

Now that we have motivated the relevance of this paper, we turn to the empirical approach. Intuitively, set identification requires two valid proxies to achieve identification. In our context, we need a 'uncertainty proxy' and a 'risk aversion proxy,' which we denote by z_t^{Unc}, z_t^{RA}, respectively. Let k be the number of variables in the proxy SVAR. By valid proxies, we mean:

$$E(z_t^{\mathrm{Unc}}\epsilon_t^{\mathrm{Unc}}) = \alpha_1 \neq 0 \tag{1}$$

$$E(z_t^{\mathrm{RA}}\epsilon_t^{\mathrm{RA}}) = \alpha_2 \neq 0 \tag{2}$$

$$E(z_t^{\mathrm{Unc}}\tilde{\epsilon}_t) = E(z_t^{\mathrm{RA}}\tilde{\epsilon}_t) = 0_{k-2} \tag{3}$$

Where $\tilde{\epsilon}_t$ represents the $k-2$ remaining shocks (i.e. all shocks in the VAR except for uncertainty and risk aversion).

In words, we need both proxies to be significantly related with their shock of interest (equation (1) and (2)). This is a standard relevance condition. Equation (3) is the exogeneity condition and states that both the uncertainty and the risk aversion proxies are uncorrelated with the $k-2$ remaining shocks. Note that we do not impose $E(z_t^{\mathrm{Unc}}\epsilon_t^{\mathrm{RA}}) = E(z_t^{\mathrm{RA}}\epsilon_t^{\mathrm{Unc}}) = 0$. In words, we, for example, allow the uncertainty shock to be correlated with the risk aversion proxy. This assumption is the key advantage of the set identification approach. Intuitively, variations in uncertainty and risk aversion are likely to be correlated such that finding exogenous variations in either of the two variables is highly challenging and is at the heart of the lack of empirical distinction between the two concepts. With the traditional external instrument approach–that is, absent set identification–, the proxy is assumed to be exogenous with regards to *all* the other shocks contained in the model, an assumption that is difficult to make in our context. Set identification thus appears particularly well suited because it only requires mild assumptions on the relationship between the proxies and the structural shocks of interest. We now turn to these assumptions.

Given two valid proxies, the problem of identification boils down to imposing restrictions on the correlation structure between the proxies and the structural shocks. Under the normalisation $E[(\epsilon_t^{\text{RA}})^2] = E[(\epsilon_t^{\text{RA}})^2] = 1$ and $E[(z_t^{\text{Unc}})^2] = E[(z_t^{\text{RA}})^2] = 1$, the matrix (4) below can be interpreted as the correlation structure between the proxies and the shocks.

$$\Phi \equiv \begin{bmatrix} \phi_{11} & \phi_{12} \\ \phi_{21} & \phi_{22} \end{bmatrix} = \begin{bmatrix} E(\epsilon_t^{\text{Unc}} z_t^{\text{Unc}}) & E(\epsilon_t^{\text{Unc}} z_t^{\text{RA}}) \\ E(\epsilon_t^{\text{RA}} z_t^{\text{Unc}}) & E(\epsilon_t^{\text{RA}} z_t^{\text{RA}}) \end{bmatrix} \tag{4}$$

Identification is achieved by imposing $\phi_{11} - \phi_{12} > 0$, $\phi_{22} - \phi_{21} > 0$, and $\phi_{11} > 0$ and $\phi_{22} > 0$. In words, we assume that the uncertainty shocks are more correlated with the uncertainty proxy than with the risk aversion proxy. Similarly, we impose that the risk aversion shock is more correlated with the risk aversion proxy than with the uncertainty proxy. This type of restrictions departs slightly from Piffer and Podstawski (2017) who impose restrictions on the *proxies* rather than on *shocks* as we do here.[3]

In practice, identification is achieved as follows. First, we use the wild bootstrap technique from Gonçalves and Kilian (2004). In a few words, the method consists in randomly replacing the sign of the reduced form residuals of the VAR to generate alternative datasets. We follow recent papers (e.g. Davidson and Flachaire (2008)) and use a Rademacher distribution to swap the sign of the residuals. We extend the wild bootstrap to account for set identification by also swapping the sign of the proxies following Piffer and Podstawski (2017). For each bootstraped data, we then identify the model using set identification; this requires to draw a random orthonormal matrix Q from the Haar measure. We give more details about the intuition of this result in Appendix A. For each draw of pseudo-data, we check that our identified shocks satisfy the restrictions in (4). We keep the draw only if the restrictions

[3]In their paper, they assume $\phi_{11} - \phi_{21} > \psi$ and $\phi_{22} - \phi_{12} > \psi$ for ψ some positive constant. In words, they assume that the *proxy* for shock A (B) is more correlated with shock A (B) than with shock B (A). Here, we assume that the *shock* A is more correlated with proxy A than with proxy B.

are satisfied. We repeat this procedure until 1,000 generated models satisfy the restrictions. Impulse response functions and historical decompositions are recovered using the median target specification of Fry and Pagan (2011). Intuitively, this method consists in choosing *one* model whose impulse response functions to a uncertainty and risk aversion shock are the closest to the median impulse response functions of the cross section of models considered. Confidence intervals are obtained by targeting different percentiles (10th, 32th, 68th and 90th percentiles in this paper). This method is generally preferred because it better reflects sampling uncertainty than simply taking the xth percentile of the impulse responses across *all* models.

Now that we have detailed the identification procedure, we turn to the construction of valid uncertainty and risk aversion proxies.

4 Uncertainty proxy

We detail here the construction of the uncertainty proxy.

4.1 Expected volatility as a measure of uncertainty

In standard macroeconomic models of the effect of uncertainty (e.g. Bloom et al. (2018)), uncertainty shocks are defined as exogenous variations in the conditional variance of future shocks. As in Berger, Dew-Becker, and Giglio (2019), our measure is the conditional variance of future stock returns. In particular, we consider the S&P500 index such that our measure directly relates to the uncertainty faced by the aggregation of the largest firms in the economy. The measure has the further advantage to be available at high-frequency and to have been used extensively in the literature, thereby making comparisons easier. If returns are unpredictable (as is nearly true empirically), we can show that the conditional variance of stock returns is approximately equal to the expected volatility of stock returns:

$$Var_t(s_{t+n}) \approx E_t\left[\sum_{i=1}^{n} r_{t+i}^2\right] \tag{5}$$

This last equation shows that the conditional variance of stock prices n periods in the future is approximately equivalent to the expected volatility of returns over the same period. In other words, uncertainty in our context can be measured using expectations of future (physical) realized volatility. Equipped with this, we thus define our uncertainty shocks as exogenous variations in (physical) expected volatility.

4.2 A model for expected volatility

To identify exogenous variations in expected volatility, we specify a relatively standard GARCH model inspired by Berger, Dew-Becker, and Giglio (2019). In particular, we assume the following dynamics for expected volatility and realized volatility respectively:

$$\sigma_t^2 = \omega + \rho\sigma_{t-1}^2 + kr_t^2 + v_t \tag{6}$$

$$r_t^2 = \sigma_{t-1}^2\epsilon_t^2 \text{ with } \epsilon_t \sim iid, E(\epsilon_t) = 0 \;\&\; E(\epsilon_t^2) = 1 \tag{7}$$

Where σ_t^2 is a measure of expected volatility. v_t can be thought of as the uncertainty shock. In our application, t is intraday (two observations per day). r_t^2 is thus the intraday squared returns. ω, ρ, k are parameters.

The model above is a standard GARCH model as proposed in Bollerslev (1986) with two exceptions. First, we have that r_t^2 depends on σ_{t-1}^2 (and not σ_t^2 as is usual). This assumption is motivated in Berger, Dew-Becker, and Giglio (2019). Intuitively, if uncertainty shocks have a linear effect on returns, they have a quadratic effect on realized volatility, making it in fact generally uncorrelated with uncertainty shocks. Berger, Dew-Becker, and Giglio (2019) show that this simple intuition holds in fact in leading theoretical models of the effect of uncertainty.

The second assumption about the independent shock v_t is motivated by the fact that, empirically, expected volatility measures are likely to not only driven by lags of realized volatility as is usually implied by GARCH models. We show in the next section that there are enough variations in measures of expected volatility to extract shocks that are orthogonal to realized volatility shocks.

In order to identify v_t, we rewrite the model defined by (6) and (7) in a VAR format:

$$\begin{bmatrix} r_t^2 \\ \sigma_t^2 \end{bmatrix} = C + \Phi \begin{bmatrix} r_{t-1}^2 \\ \sigma_{t-1}^2 \end{bmatrix} + \Sigma \mu_t \tag{8}$$

With C a 2×2 matrix of constant with $C_{11} = 0$ and $\Phi_{:,1} = 0$ (first column of Φ) and Σ is a 2×2 lower triangular matrix (Cholesky decomposition). $\mu_t = \begin{bmatrix} \mu_{t,1} & \mu_{t,2} \end{bmatrix}'$ is a vector of "structural" shocks (to r_t^2 and σ_t^2 respectively).

Focusing on σ_t^2, we can write:

$$\sigma_t^2 = C_{21} + \Phi_{22}\sigma_{t-1}^2 + \Sigma_{21}\mu_{t,1} + \Sigma_{22}\mu_{t,2}$$

Where the lower script indicates the element of the matrix. $\mu_{t,1}$ can be thought of as a realized volatility shock and $\mu_{t,2}$ (with $\Sigma_{22}\mu_{t,2} = v_t$) as the shock to expected volatility that is orthogonal to realized volatility (orthogonality is due to the Choleski decomposition). Estimating this model with intraday data (opening and closing prices) allows us to recover two daily shocks μ_{2t}^{Open} and μ_{2t}^{Close}.

As mentioned, the identified v_t shocks are orthogonal to realized volatility shocks. Why does this matter? Intuitively, Berger, Dew-Becker, and Giglio (2019) highlight the importance of controlling for realized volatility shocks when investigating the effect of uncertainty shocks. In particular, the authors argue argue that in macroeconomic models of the effect of uncertainty such as those with wait-and-see and precautionary savings effects, uncertainty is driven by variations in the agents' subjective distribution of future shocks, as opposed to the variance of the distribution from which today's shocks were drawn. In other words, uncertainty shocks are shocks to *future* volatility whereas current realized volatility measures how large are the shocks that have just occurred. This distinction is crucial and notably implies that realized volatility shocks, while related, are a different concept than uncertainty shocks. While the distinction may appear trivial, it is noteworthy to mention that seminal papers on the effect of uncertainty such as Bloom (2009) do not make the distinction and use a mixture of realized volatility and expected volatility shocks to identify uncertainty shocks. By identifying shocks to expected volatility that are orthogonal to realized volatility shocks, our model specification is thus consistent with these insights about the nature of uncertainty.

4.3 Data specification

To be able to estimate the model defined by equations (6) and (7) above, we first need to have an adequate measure of expected volatility. A natural candidate is the VIX. Intuitively, the VIX represents the option-implied volatility of the S&P500 index with a horizon of 30 calendar days. This concept is often referred to as a 'risk neutral' measure of volatility. As opposed to physical measures which use the actual state probabilities, risk neutral measures use probabilities that are adjusted for the pricing of risk.

To check formally that the VIX indeed contains information about the future path of realized volatility, we first run the following regressions with daily data:

$$FRV_t = \alpha + \beta VIX_t + e_t \tag{9}$$

With $FRV_t \equiv \sum_{i=1}^{120} RV_{t+i}$. We choose a time-horizon of 180 days ($\approx$ 120 business days) to have a measure of the 6-month ahead expected volatility. 6-month is standard in the literature on uncertainty because it aligns well with the horizon of firms' business decisions.

The first column of Table 1 displays the results from the simple regression in (9). As we can see, the coefficient on the VIX is highly significant (at the 1 percent level). Interestingly, this simple regression with only one regressor displays a relatively large R^2 of 18%. This suggests that the VIX contains valuable information the future path of realized volatility.

Note that we are not making the assumption that option prices are statistical expectations of future volatility. In particular, the VIX is an asset price and as such depends on (time-varying) risk premia and potential measurement errors. What we want to show, however, is that the VIX is linearly and robustly related to a statistical expectation of realized volatility. An intuitive way to see this is to take the expectation on both side of the previous regression:

$$E_t\left(\sum_{i=1}^{180} RV_{t+i}\right) = \alpha + \beta VIX_t \tag{10}$$

On the right hand side, we have the conditional expectation of future realized volatility. If β is robustly different from zero, an exogenous increase in the VIX (which we aim to identify) directly maps with an exogenous increase in the expected volatility, which is precisely what we want to isolate. In other words, exogenous variations in the VIX can be used to identify exogenous variations in expected volatility, up to a scaling factor.

In a second step, we want to show that future realized volatility cannot be entirely explained by its own lags. In standard GARCH models, expected volatility is defined as an exponentially moving average of past squared returns (Andersen et al. (2013)). In other terms, standard GARCH models do not allow for the existence of other types of shocks such as v_t in equation (6). Here, we want to show that our measure of expected volatility–the VIX–can help predict future volatility above and beyond the information contained in the lags of realized volatility. To do so, we run regressions of the form:

$$FRV_t = \alpha + \beta VIX_t + \gamma \sum_{j=0}^{4} RV_{t-j} + e_t \tag{11}$$

RV_t represents daily realized volatility. Results are displayed in Table 1. Of particular interest is the last column where the VIX remains highly significant even when we add several lags of realized volatility as additional regressors. We can also note that the coefficient remains remarkably similar across the different specifications. These results suggest that the VIX can robustly predict future volatility above and beyond what is contained in its lags. Additionally, we can note that the marginal increase of including the VIX in terms of R^2 is significantly larger than lags of realized volatility. This suggests that, even though past squared returns indeed contain information about the future of path realized volatility, their forecasting power is significantly lower than that of the VIX. Once again, this suggests that option prices contain valuable information about the future path of expected volatility and that there are enough variations in expected volatility to isolate shocks that are orthogonal to realized volatility shocks.

4.4 Narrative approach

If the model in (6)(7) is correctly specified, v_t can be interpreted as an exogenous variation in expected volatility, and thus as an uncertainty shock. In our context, we caution about this interpretation for two main reasons. First, because the VIX is an asset price, it depends on time-varying risk premia which in turn implies that part of the variations we identify could reflect changes in investors' tolerance towards risk. Second, while the model defined in (6)(7) provides a good characterization of the determinants of the VIX, it is possible that other shocks affect the VIX. To minimize this risk, relying on a narrative approach seems an adequate approach and allows to interpret the identified shock with more peace of mind. The idea of the narrative approach is that, by relying on events that are likely to have generated

Table 1: Predictive power of the VIX

	Dependent variable:			
	Agg. RV			
	(1)	(2)	(3)	(4)
VIX	0.001***		0.001***	0.001***
	(0.00002)		(0.00003)	(0.00003)
RV		13.305***	5.239***	4.384***
		(0.528)	(0.567)	(0.580)
lag(RV, 1)				3.619***
				(0.567)
lag(RV, 2)				2.691***
				(0.571)
lag(RV, 3)				3.195***
				(0.563)
lag(RV, 4)				2.285***
				(0.572)
Constant	−0.004***	0.012***	−0.002***	0.001
	(0.0005)	(0.0002)	(0.001)	(0.001)
Observations	6,798	6,798	6,798	6,794
R^2	0.180	0.086	0.190	0.206
Adjusted R^2	0.180	0.085	0.190	0.205

Notes: *p<0.1; **p<0.05; ***p<0.01

Robust standard errors. RV stands for (daily) realized volatility.
Agg. RV is aggregated over 120 business days (≈ 6 months).

variations in volatility but are unrelated with other macroeconomic shocks, the identified shock should more closely reflect exogenous variations in expected volatility. As noted in the introduction, however, these variations in expected volatility could still be correlated with risk aversion shocks and this is why set identification is required to credibly interpret the identified uncertainty shocks.

For the volatility events, we rely on an updated database from Piffer and Podstawski (2017). This database contains major global events that are likely to have generated variations in financial volatility. Importantly, we have information about the timing at which the news hit the market. The events used in the baseline analysis are selected such that is is reasonable to expect that they were not anticipated and are exogenous with respect to other macroeconomic shocks. The complete list of events used in the benchmark specification can be found in the appendix (Table 4).

4.5 Building the proxy

To build our uncertainty proxy, we proceed as follows. First, we run the bivariate VAR defined in equation (8) with a Cholesky decomposition and using the VIX as our measure of expected volatility and squared returns on the SP500 as our measure of realized volatility. In particular, we have data on the opening and closing price of each variable. Running the model thus allows us to have two measures of $\mu_{t,2}$ per day. We refer to the first daily shock (denoted as $\mu_{t,2}^{\text{Open}}$) as an "opening time shock." Similarly, we refer to the second measure (denoted as $\mu_{t,2}^{\text{Close}}$) as a "closing time shock." For illustration purpose, we also define $\mu_{t+1,2}^{\text{Open}}$ as the "next day opening time shock." Intuitively, the US stock market opens at 2:30 pm and closes at 9pm UK time. If the news of the uncertainty event hits the market before it opens, the relevant shock used in the construction of our proxy is the "opening time shock." If the news hit the market between 2:30pm and 9:00pm, the relevant shock is the "closing time shock." Finally, if the news hits the market after 9:00 pm, the relevant shock used in our proxy is the "next day opening time shock." On the day of the event, the proxy takes the value of the relevant shock and zero otherwise. Consistent with existing literature (e.g. Känzig (2021)), we then aggregate the proxy at the monthly frequency by taking the sum of the uncertainty shocks within each month. The uncertainty proxy used in the baseline analysis is displayed in Figure 2. As we can see, the proxy exhibits variations in line with what could be expected. For instance, both the 9/11 and the results of the Brexit vote led to a sharp increases in uncertainty.

5 Risk aversion proxy

In this section, we discuss the construction of the risk aversion proxy.

5.1 Definition of risk aversion shocks

We define risk aversion shocks as exogenous changes in the marginal investor's risk *preferences*. In doing so, we distinguish between endogenous variations in the coefficient of risk aversion, for instance through changes in wealth and habit formation (e.g. Campbell and Cochrane (1999)) or through changes in background risk (e.g. Heaton and Lucas (2000)) In Campbell and Cochrane (1999), time variation in risk aversion is generated through habit formation: as consumption moves closer to the habit (for instance during a business cycle trough), the curvature of the utility function increases and raises risk aversion endogenously. Generally speaking, however, relying solely on consumption volatility to explain time variations is insufficient to match the empirically observed variations in the price of risk. Guiso and Paiella (2008) find that variations in background risk can only explain a small fraction of empirically observed time-variation in risk aversion and attribute most of the (unexplained) variation to 'genuine differences in taste.' In this context, other strands of the literature have modelled risk aversion as a stochastic process with shocks that are uncorrelated to economic fundamentals. Gordon and St-Amour (2000) for instance show that stochastic risk aversion in the form of preference shocks can successfully explain empirical variations in the price of risk. In the same vein, Bekaert, Engstrom, and Xing (2009) and Bekaert, Engstrom, and Grenadier (2010) model risk aversion as a stochastic process and interpret shocks to risk aversion as preference shocks.[4] More recently, evidences from Martin (2017) and Bekaert, Engstrom, and Xu (2021) imply that risk aversion is much more rapidly mean reverting than implied by standard habit models, thereby emphasizing the empirical relevance of (high-frequency) exogenous variations. In DSGE models, preference shocks are important drivers of variations (e.g. Smets and Wouters (2007)). In an experimental setting, Cohn et al. (2015) show that risk aversion *preferences* are countercyclical. Their experimental design allows to measure the psychological part of risk aversion without the confounding influence of background risk, wealth effects, changing habits, experienced gains or losses because all these variables remain unchanged across conditions. Their results thus highlight the role of (exogenous) psychological factors in driving the risk aversion parameter. In this paper, we interpret risk aversion as a stochastic process whose exogenous disturbances can be interpreted as sudden changes in agents' preferences towards risk that are uncorrelated with economic fundamentals. The ultimate sources of such variations remain, however, open to debate. Bekaert, Engstrom, and Xu (2021) argue that it could for instance reflect shifts in sentiment induced by news, or even mood swings due to the weather (e.g. Kamstra, Kramer, and Levi (2003)).

5.2 Identification of risk aversion shocks

5.2.1 Measuring risk aversion

Generally speaking, the identification of exogenous variations in risk aversion is challenging. For one, risk aversion is a latent concept and thus cannot be directly observed. The first difficulty is thus to derive a measure that credibly captures variations in risk aversion and allows for variations that are uncorrelated with economic fundamentals. One such measure is the one developed in Bekaert, Engstrom, and Xu (2021). This measure appears particularly

[4]We can also mention Xu (2016) who interpret volatility shocks as preference shocks but do not make the distinction between volatility, uncertainty and risk aversion.

adequate because it is consistent with structural dynamic asset pricing models related to the habit models of Campbell and Cochrane (1999), Menzly, Santos, and Veronesi (2004) and Wachter (2006). In a few words, the measure of risk aversion is a second factor in the pricing kernel that is not exclusively driven by economic fundamentals. The approach explicitly separates the price of risk (risk aversion) from the quantify of risk (uncertainty), thereby making the resuting measure particularly adequate in our setting. To get some more intuition, we develop here the basic idea behind the measure of risk aversion in Bekaert, Engstrom, and Xu (2021).

The starting point is a utility function that depends on both consumption ("fundamentals") and a "non-fundamentals" factor. Consider a utility function of the HARA class:

$$U(\frac{C}{Q}) = \frac{(\frac{C}{Q})^{1-\gamma}}{1-\gamma} \tag{12}$$

Where C is consumption and Q is a process that can drive variations in risk aversion. As Q rises, consumption delivers less utility and marginal utility is high. Q is defined as follows:

$$Q = \left(\frac{a}{\gamma} - \frac{b}{C}\right)^{-1} \tag{13}$$

Where a and γ are positive parameters and b is an exogenous benchmark parameter or process. Note that the curvature paramter γ is *not* equal to risk aversion in this setting. The coefficient of relative risk aversion is given by:

$$RRA = -\frac{CU''(C)}{U'C} = aQ \tag{14}$$

In words, relative risk aversion is proportional to Q. Assuming an infinitely-lived agent facing a constant discount factor β, we can derive the pricing kernel as:

$$m_{t+1} = ln(\beta) + ln\left(\frac{U'(C_{t+1})}{U'(C_t)}\right) = ln(\beta) - \gamma\Delta c_{t+1} + \gamma\Delta q_{t+1} \tag{15}$$

This formulation nests some interesting special cases: for instance Campbell and Cochrane (1999) and their external habit model fits in this specification with $Q_t = \frac{C_t}{C_t - H_t}$ and $a = \gamma$ and $b = H$. While Campbell and Cochrane (1999) also model q_t exogenously, they restrict the correlation between q_t and Δc_t to be perfect. In Bekaert, Engstrom, and Xu (2021), q_t is partly but not fully driven by fundamentals (consumption growth) as it features an indepedent shock. While the ultimate sources of this shock are open to debate, a sensible interpretation is that of a preference shock. The resulting measure of risk aversion is displayed in Figure 1.

5.2.2 Narrative approach

To identify shocks to risk aversion, our approach is similar to the one used for the uncertainty proxy. In particular, we follow a narrative approach and we rely on the same set of volatility events used for the uncertainty proxy. The idea is that, by looking at a measure that should be closely related with the aggregate market risk aversion and by relying on volatility events that are unrelated to other macroeconomic shocks, the variations that we observe around these events should be closely related to true risk aversion shocks, thereby providing the basis to work as an external instrument. Furthermore, the resulting shocks should be more correlated with the true structural risk aversion shocks than the uncertainty ones, an assumption that is required given our set identification approach.

5.3 Building the proxy

In practice, the risk aversion proxy is constructed as follows. Because it uses end-of-the-day US variables in its derivation, the risk aversion measure that we consider should be interpreted as an "end-of-the-day" measure, that is 4pm US time. For each volatility event that we consider, we check whether it hit the market before or after the closing time of the US stock market. If it is before, the variation is equal to the current value of the risk aversion measure minus the value the day before. If it is after, the variation is equal to the value of the risk aversion measure the next day minus the value of the risk aversion measure today. Similarly to the uncertainty proxy, the daily variations are then aggregated at the monthly frequency. The resulting risk aversion proxy is plotted in Figure 2.

6 Validity of the proxies

6.1 Discussion

To have room for identification, it is necessary that the uncertainty and risk aversion proxies exhibit significantly different behaviours around the volatility events. This is important to ensure that we capture different types of shocks around the volatility events that we consider. First, while we find that the uncertainty and the risk aversion proxy are positively correlated–as could be expected–we also find that this correlation is relatively limited at 0.52. Second, we find that, with a standard deviation of 0.09, the (demeaned) risk aversion proxy is significantly less volatile than the (demeaned) uncertainty proxy with a standard deviation of 0.46. This can for instance be seen on the graph where there are a several events that lead to relatively large variations in uncertainty while the risk aversion proxy remains relatively stable. This result holds when comparing a variety of "raw" series of uncertainty to the risk aversion measure that we consider; in all cases, uncertainty exhibits larger volatility than risk aversion. This finding is also consistent with the intuition, as it is reasonable to expect the physical uncertainty faced in the economy to move more abruptly than the aggregate agents' preferences towards this risk. Finally, while we find that a few events coincide with large increases both in uncertainty and risk aversion (e.g. the 9/11 attack, the COVID pandemic or the results of the Brexit vote), there are a variety of events that lead to significantly different variations: for instance, the London bombings lead to important variations in uncertainty while it left risk aversion relatively unchanged. Another example is the election of Donald Trump as the US president which is found to be a positive uncertainty shock but a negative risk aversion shock (see Figure 10 in the appendix). These heterogeneous reactions form the basis to identify independent variations in risk aversion and uncertainty.

An important potential concern could be that, to the extent that uncertainty is likely to affect risk aversion endogenously, part of the variations in risk aversion that we identify could be due to uncertainty shocks, and as such could potentially put the validity of the risk aversion proxy in question. We argue that this is not an issue in our context. For one, as mentioned above, we find significant differences between the risk aversion and uncertainty proxy, suggesting that the variations in risk aversion that we identify do not systematically stem from uncertainty shocks. Second, our identification scheme only requires that risk aversion shocks are more correlated with the risk aversion proxy than with the uncertainty proxy. Even though our risk aversion proxy is likely to reflect in part uncertainty shocks, it is reasonable to expect it to be more correlated with true risk aversion shocks as it is computed using a measure specificaly designed to capture variations in risk aversion. Similarly, requiring the uncertainty shock to be more correlated with the uncertainty proxy than with the risk aversion shock also appears reasonable as the uncertainty proxy is built in a way to capture the empirical equivalent of uncertainty shocks in theoretical models.

6.2 Relevance

As mentioned, the joint identification requires to have two valid proxies that can then be used as external instruments. The first requirement is that both proxies are relevant. This can be tested by means of first stage regressions (i.e. a regression of a proxy on its respective reduced form residuals). Table 2 displays the F-statistics of these first stage regressions for each proxy. As we can see, the F-statistics are above 10, which is usually considered as the required threshold to avoid weak instrument issues (Montiel-Olea, Stock, and Watson (2016)).

Figure 1: The risk aversion index

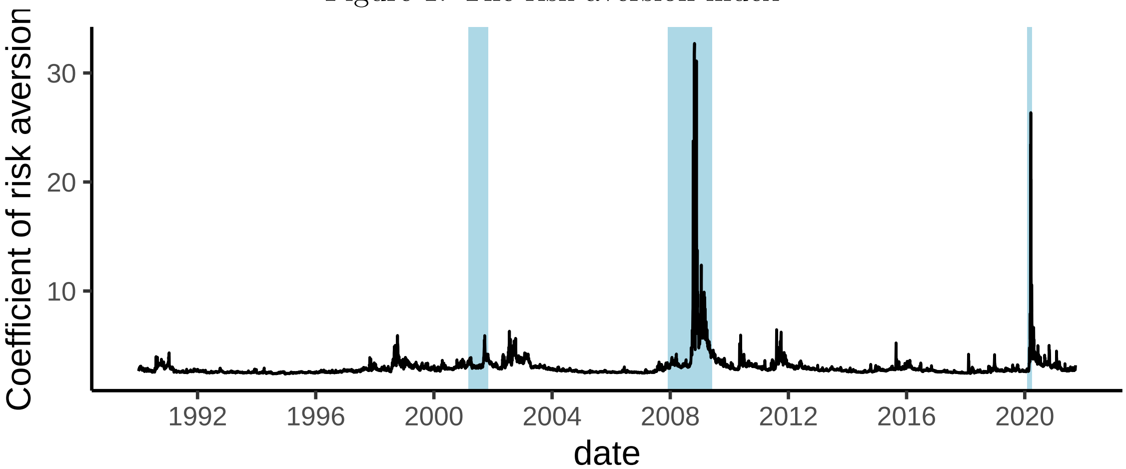

Notes:
The risk aversion index is taken from Bekaert et al. (2021). The shaded blue areas indicate the NBER recessions.

Table 2: Results of 1st stage regression

	VIX	log(SP500)
F-Stats	23.02	26.48

Notes: F-Stats are obtained by regressing the proxy on its respective reduced form residuals. For the uncertainty (risk aversion) proxy, the relevant reduced form residuals are obtained from the VIX (log(SP500)) variable in the VAR.

Figure 2: Uncertainty and risk aversion proxies

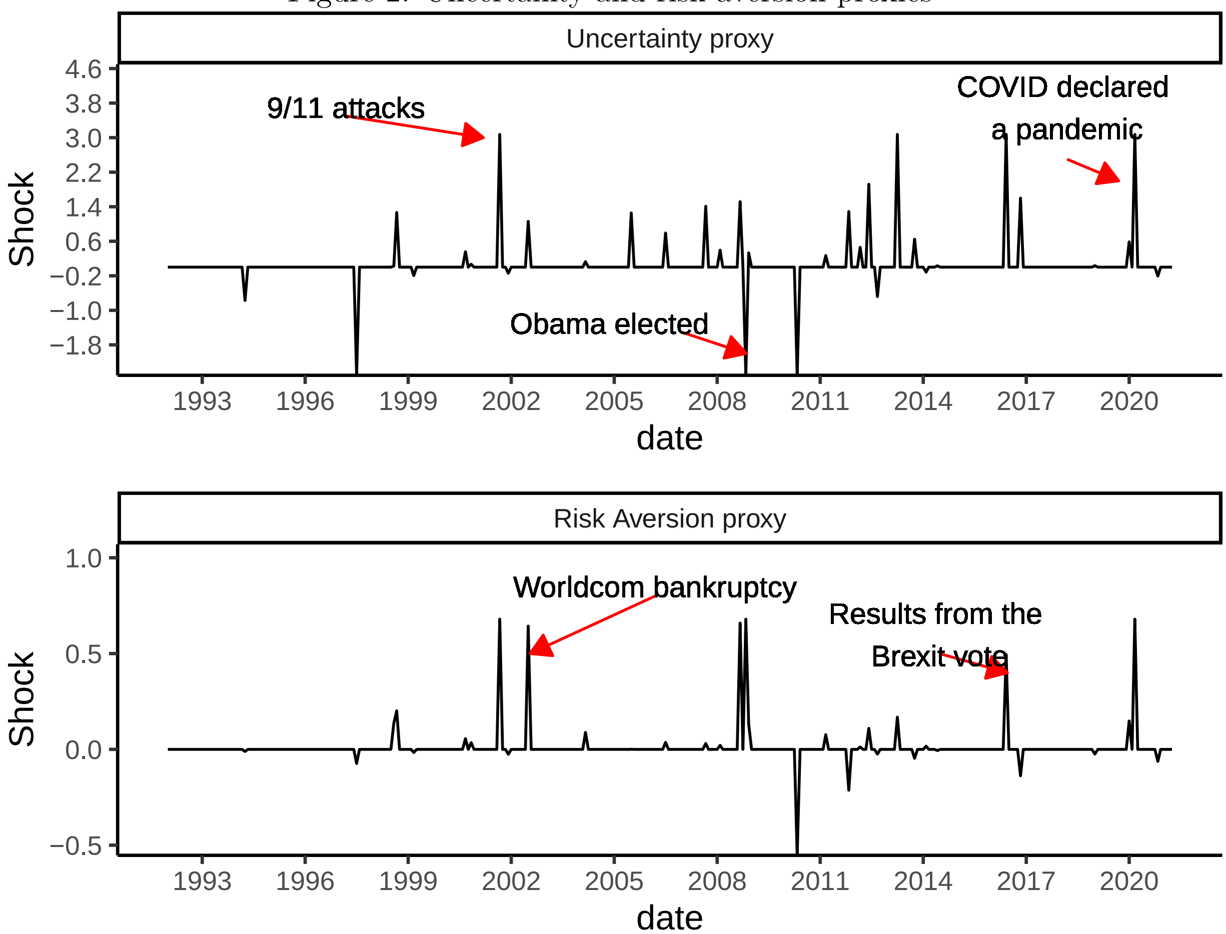

Notes:
The uncertainty shocks are recovered using a slightly adjusted GARCH(1,1) in the spirit of Berger, Dew-Becker, and Giglio (2019) and intraday data (opening and closing prices) on the VIX. The risk aversion shocks are recovered using the daily variation in the measure of risk aversion developed in Bekaert et al. (2021). Volatility events are recovered from an updated database of Piffer and Podstawski (2017). The proxies are winsorized at the 1% level.

6.3 Exogeneity

As for the exogeneity assumption, it is not possible to directly test the exogeneity of the two proxies. There are, however, indirect ways to do this. First, a good proxy should not be autocorrelated nor forecastable by past macroeconomic variables. Additionally, it should also not be correlated with other structural shocks (Ramey (2016)). Autocorrelation plots do not suggest evidences of autocorrelation. To check whether other macroeconomic variables have forecasting power over the series, we test whether any of the macroeconomic variables in the VAR Granger cause one of the proxy. We find no evidence of Granger causality. Overall, the evidences support the exogeneity assumption of our two proxies.

7 Dynamic effects of uncertainty and risk aversion shocks

7.1 Estimation and data

Now that we have detailed the estimation techniques and motivated the proxies, we can proceed to the actual estimation. To ease comparisons with existing literature, we consider a VAR similar to the seminal paper of Bloom (2009). In more details, we consider a vector of 8 endogenous variables that enter the VAR in the following order: log(SP500), VIX, FFR, log(wages), CPI, hours worked, log(employment) and log(IP). Similar to Bloom (2009), we consider most variables in percentage deviations from an HP trend ($\lambda = 129,600$). Based on the information criteria, we estimate a VAR with four lags. Our sample spans from 1990M2 to 2021M2. The sample is limited by the availability of intraday data on the VIX. Sources and graphical representation of the data used in the VAR can be found in the appendix (Figure 12). Our VAR framework differs from Bloom (2009) in three main ways: first, our sample starts later but contains data up to 2021. Second, we use the entire dynamic of the VIX (in levels) instead of an indicator series as a measure of uncertainty. Third, we take the first difference of log wages to ensure stationarity of the series. The VAR is estimated equation by equation using OLS. We report the median impulse response functions as well as the 90 and 68% confidence intervals using the wild bootstrap approach (Gonçalves and Kilian (2004)) and the target specification developed in Fry and Pagan (2011). Results are based on 1,000 bootstraped models that satisfy the restrictions detailled in Section 3.

To build the intuition, we first estimate the model using a Cholesky decomposition and ordering the stock market variable first and the VIX second. This ordering notably implies that the effect of the "uncertainty shock" on the stock market variable is constrained to zero on impact. This approach is standard in the literature (e.g. Bachmann, Elstner, and Sims (2013), Bloom (2009)). Recursive identification thus rules out contemporaneous feedback effects between uncertainty and other economic variables. Ludvigson, Ma, and Ng (2021) argue that it is an important deficiency of traditional RI schemes. Figure 11 in the appendix plots the IRFs from the 8 endogenous variables following a one standard deviation in the VIX with uncertainty shocks recursively identified. Results are in line with existing literature. In particular, we find that, following a one standard deviation shock on the VIX, the real economy reacts negatively with hours, log employment, log industrial production all declining significantly after a few quarters. This is accompanied by a looser monetary policy stance. The fact that our results align well with existing literature suggests that our sample period and specification is well suited for comparison.

In a second step, we estimate the model with our uncertainty and risk aversion proxies and the set identification approach detailled in Section 3. Without loss of generality, we order the VIX variable first and the stock market variable second. We instrument the former variable with the uncertainty proxy and the latter with the risk aversion proxy. The set identification allows to overcome the limitations of the recursive approach by allowing the stock market and the VIX to react contemporaneously to one another. Furthermore, the approach also allows to have uncertainty shocks that are cleaned of the effect of risk aversion shocks, something that is typically ignored when using the VIX as a measure of uncertainty. Finally, we are also able to identify and quantify the dynamic effects of risk aversion shocks. The structural shocks for uncertainty and risk aversion are displayed in Figure 14 in the appendix.

7.2 Uncertainty shocks

Figure 3 plots the impulse response functions to an uncertainty shock identified using the set identification approach. Generally speaking, the IRFs exhibit relatively similar dynamics than under the recursive identification scheme. An important difference, however, is that the stock market price variable drops significantly on impact whereas it is constrained to zero in the recursive scheme. This result suggests that the contemporaneous restrictions imposed by the recursive identification scheme are inappropriate.

We find that the real economy reacts negatively to a set-identified financial uncertainty shock as hours worked, employment, and industrial production all decline significantly. Compared to the recursive identification scheme, the largest drop in industrial production is 1.43 times larger under the set identification approach. This finding suggests that uncertainty shocks may be more damaging to the real economy that commonly accepted, in line with recent work such as Piffer and Podstawski (2017), Ludvigson, Ma, and Ng (2021), and Alessandri, Gazzani, and Vicondoa (2021). We find that this drop in output is accompanied by a looser monetary policy stance. In line with theoretical predictions (e.g. Bloom (2009)), the effects are relatively short-lived as they stop being statistically significant after approximately 6 months for employment, hours worked and industrial production.

7.3 Risk aversion shocks

Figure 4 plots the dynamic responses of the variables following a risk aversion shock identified under the set identification approach. Generally speaking, risk aversion shocks appear to be damaging to the economy but display larger confidence intervals such that most responses

Figure 3: IRFs following an uncertainty shock (set identification)

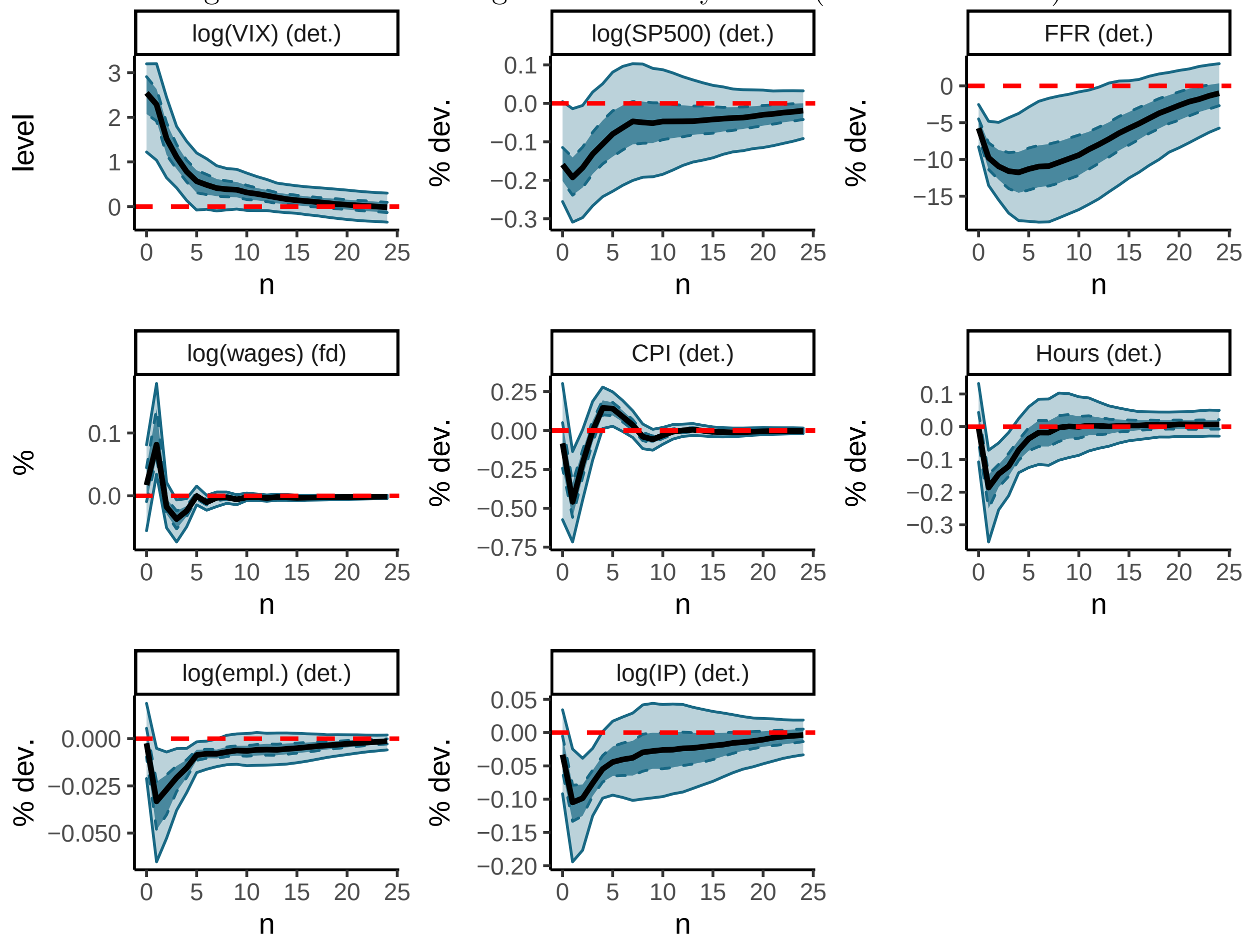

Notes:
The graphs plot the IRFs of the 8 variables contained in the VAR following an uncertainty shock that is set identified. The median IRFs is based on 1,000 bootstrap replications satisfying the restrictions discussed in Section 3 and is obtained using the median target specification of Fry and Pagan (2014). The confidence intervals are obtained by targeting the 10th, 32th, 68th and 90th percentiles, respectively. The IRFs correspond to a one standard deviation shock on the VIX. Wages and CPI are expressed in log-differences (monthly growth). The VIX is in levels. The other variables are expressed in percentage deviation from the HP trend.

fail to be statistically significant. Focusing on the median target response, we find however that risk aversion shocks coincide with persistent declines in industrial production, hours worked, and the policy rate. On impact, employment increases but this effect is short-lived. At longer horizons, it appears that risk aversion shocks do not affect employment much. Risk aversion shocks have a large and negative effect on asset prices and increase the VIX on impact. Overall, risk aversion shocks exhibit significantly different dynamics than uncertainty shocks. This highlights the importance of the distinction between the two. To get a better sense of the macroeconomic importance of both types of shocks, we will turn to forecast error variance and historical decompositions in the next section.

7.4 Robustness checks

We perform some robustness checks. Because the proxies are built using a narrative approach, a potential concern is that it is the particular choice of events that drives the results. Figure 8 and 9 in the appendix plot the IRFs from a uncertainty and risk aversion shock using different events for the construction of the proxies. Results remain similar across specifications.

8 Contribution to business cycle fluctuations

Since at least Kydland and Prescott (1982), the question of the source of macroeconomic fluctuations has been at the center of macroeconomic analysis. With the recent GFC, there has been a renewed interest in the macroeconomic importance of financial shocks as can be seen from a large body of literature aimed at identifying and quantifying the contribution of so-called 'financial shocks' to business cycle fluctuations (e.g. Gilchrist and Zakrajšek (2012), Gertler and Karadi (2015), Caldara et al. (2016)). Generally speaking, existing literature often blurs the distinction between risk aversion and uncertainty shocks by using some common measure of volatility. Because financial shocks tend to be procyclical, it is usually challenging to separate their effect from other economic cycles using recursive identification.

Figure 4: IRFs following an risk aversion shock (set identification)

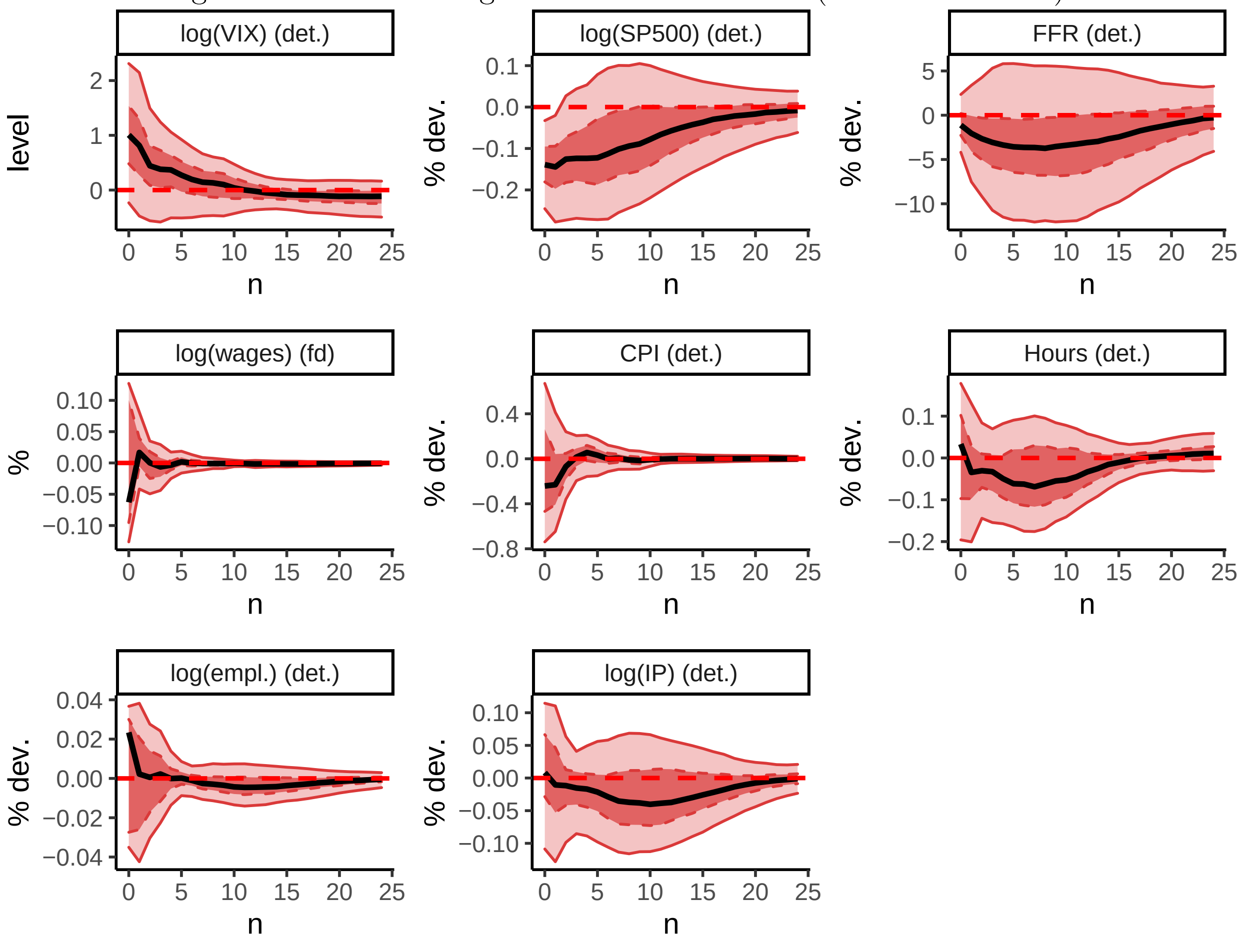

Notes:

The graphs plot the IRFs of the 8 variables contained in the VAR following a risk aversion shock that is set identified. The median IRFs is based on 1,000 bootstrap replications satisfying the restrictions discussed in Section 3 and is obtained using the median target specification of Fry and Pagan (2014). The confidence intervals are obtained by targeting the 10th, 32th, 68th and 90th percentiles, respectively. The IRFs correspond to a one standard deviation shock on the S&P500 variable. Wages and CPI are expressed in log-differences (monthly growth). The VIX is in levels. The other variables are expressed in percentage deviation from the HP trend.

Our identification approach allows to overcome these limitations and to quantify the relative importance of the two types of shocks. In this section, we investigate the contribution of two types of financial shocks, namely financial uncertainty and risk aversion, to business cycle fluctuations.

8.1 Forecast error variance decomposition

We first perform a forecast error variance decomposition to get a sense of the macroeconomic importance of uncertainty and risk aversion shocks. Results are displayed in Table 3. We find that financial uncertainty shocks can explain around 86% of variations in the VIX at a one-month horizon, with this share declining to 42% at a 24-month horizon. On the other hand, risk aversion shocks can explain around 16% at a one-month horizon, with this share declining to around 10% at longer horizons. To the best of our knowledge, this paper is the first to quantify the drivers of the VIX.[5] Second, we find that both uncertainty and risk aversion shocks can explain a non-negligible share of the forecast error of stock prices. For instance, uncertainty and risk aversion shocks can explain up to 20% and 14% of the variations, respectively. The fact that risk aversion shocks matter for the dynamics of asset prices is in line with a large macro-finance literature emphasizing the role of variation in risk aversion for asset pricing (dating back at least to Epstein (1988)). For the real economy, we find that both uncertainty and risk aversion shocks are important drivers. For industrial production, uncertainty and risk aversion shocks can explain up to 21% and 15% of the forecast error variance, respectively. Together, both types of shocks can account for around 30% of variations in output at most horizons, thereby suggesting that the two shocks are important determinants of real macroeconomic fluctuations.

[5] Bekaert, Hoerova, and Lo Duca (2013) for instance find that time-varying risk premia is an 'important driver' of the VIX, but do not quantify this statement.

Table 3: Forecast error variance decomposition

t	VIX	SP500	FFR	Wages	CPI	Hours	Empl.	IP
				Uncertainty shock				
t	**VIX**	**SP500**	**FFR**	**Wages**	**CPI**	**Hours**	**Empl.**	**IP**
1	0.84	0.2	0.18	0.05	0.02	0.02	0.05	0.06
	[0.3;1]	*[0.01;0.44]*	*[0.04;0.3]*	*[0;0.241]*	*[0;0.14]*	*[0;0.1]*	*[0;0.28]*	*[0;0.22]*
6	0.59	0.17	0.23	0.2	0.1	0.16	0.24	0.21
	[0.179;0.81]	*[0.02;0.41]*	*[0.04;0.38]*	*[0.08;0.38]*	*[0.03;0.22]*	*[0.06;0.29]*	*[0.06;0.54]*	*[0.06;0.37]*
12	0.51	0.13	0.2	0.2	0.11	0.15	0.24	0.19
	[0.15;0.72]	*[0.02;0.37]*	*[0.03;0.36]*	*[0.09;0.37]*	*[0.04;0.22]*	*[0.05;0.251]*	*[0.09;0.48]*	*[0.05;0.31]*
24	0.42	0.12	0.16	0.19	0.11	0.15	0.23	0.16
	[0.13;0.63]	*[0.03;0.33]*	*[0.03;0.34]*	*[0.09;0.37]*	*[0.04;0.22]*	*[0.05;0.25]*	*[0.09;0.43]*	*[0.04;0.3]*
				Risk aversion shock				
t	**VIX**	**SP500**	**FFR**	**Wages**	**CPI**	**Hours**	**Empl.**	**IP**
1	0.16	0.14	0.02	0.34	0.1	0.07	0.42	0.14
	[0;0.7]	*[0.01;0.43]*	*[0;0.1]*	*[0.12;0.55]*	*[0.01;0.28]*	*[0.01;0.2]*	*[0.2;0.64]*	*[0.01;0.53]*
6	0.12	0.13	0.04	0.2	0.12	0.1	0.22	0.12
	[0.02;0.47]	*[0.02;0.44]*	*[0;0.19]*	*[0.1;0.33]*	*[0.03;0.27]*	*[0.04;0.2]*	*[0.06;0.45]*	*[0.04;0.27]*
12	0.12	0.13	0.04	0.2	0.12	0.13	0.2	0.15
	[0.02;0.41]	*[0.02;0.43]*	*[0;0.19]*	*[0.1;0.33]*	*[0.03;0.27]*	*[0.05;0.25]*	*[0.07;0.39]*	*[0.05;0.28]*
24	0.11	0.11	0.05	0.2	0.12	0.13	0.19	0.14
	[0.03;0.35]	*[0.02;0.38]*	*[0.01;0.19]*	*[0.1;0.33]*	*[0.03;0.27]*	*[0.05;0.25]*	*[0.08;0.34]*	*[0.05;0.3]*

Note:
The table shows the forecast error variance decomposition of the median target at horizons 1, 6, 12 and 24 months. The 10th and 90th percentiles confidence intervals are reported in the brackets below and are based on the 1,000 bootstraped replications.

8.2 Historical decompositions

To refine our understanding of the cumulated role of uncertainty and risk aversion shocks, we perform a historical decomposition of each variable contained in our model. Figure 13 in the appendix shows this decomposition for all variables. To highlight the potentially heterogeneous dynamics of both types of shocks, we decide to investigate more closely the two arguably largest shock in our sample, namely the Great Financial Crisis and the recent COVID recession. The sources of the two crises differ greatly. On the one hand, the Great Financial Crisis is a prime example of a financial crisis where turmoils in the financial sector ultimately translated to adverse effects on the real economy. On the other hand, the COVID recession was due to the emergence of the COVID pandemic which lead to a series of lockdowns and travel restrictions that prevented the economy to function normally. To the extent that these shocks are likely to coincide with large variations in volatility but differ with regards to their origins, it appears particularly interesting to investigate their dynamics in our setting.

Great Financial Crisis

Figure 5 plots the evolution of three selected variables, namely the VIX, the stock price index and industrial production around the Great Financial Crisis. Intuitively, the historical decomposition allows to analyse the evolution of the variables in the absence of different types of shocks. The black line plots the data. The blue line plots the variables in the absence of the set-identified financial uncertainty shocks. The difference between the blue and the black line can thus directly be interpreted as the contribution of uncertainty shocks. The red line plots the variables in the absence of both set-identified uncertainty and risk aversion shocks. The difference between the red and the blue line can thus directly be interpreted as the contribution of risk aversion shocks.

As we can see, the blue line is very close to the black line for all variables of interest, thereby suggesting that uncertainty shocks were not the main drivers of the dynamics around the GFC. On the other hand, we can notice that the increase in the VIX would have been considerably smaller absent risk aversion shocks. Similarly, both the stock price index and industrial production would have declined significantly less without risk aversion shocks. This suggests that risk aversion shocks have played a major role for the dynamics around the GFC whereas financial uncertainty appears to have played a significantly smaller role. This interpretation is consistent with the fact that the GFC coincided with a large increase in risk premia across the board (see e.g. Muir (2017)) and that this can have real effects (see e.g. Smets and Wouters (2007)).

In terms of mechanism, an interpretation is that, because of its financial origin, the GFC has considerably shaken the agents' trust in the financial sector which resulted in elevated levels of risk aversion (see e.g. Guiso (2012)). This heightened risk aversion could have in turn prevented investors from undertaking high-growth but high-risk projects and thereby slow down the accumulation of capital. Additionally, because more risk averse agents require higher risk premia, it can slow down growth by raising the cost of equity investment and thereby discourage investment in innovative firms which rely disproportionately on equity finance (Guiso (2014)). In the next section, we will perform a historical decomposition around the COVID pandemic. To the extent that the cause of the shock are dramatically different (i.e. do not find their roots in the financial markets), it will be interesting to check whether risk aversion shocks play a similar role during the economic recovery.

COVID pandemic

The COVID pandemic generated a massive increase in uncertainty by its extraordinary nature. In this context, it appears interesting to investigate how this shock translated on the financial markets and the real economy. For policymakers, it is also important to understand the different channels through which it can affect the economy. Figure 6 plots the historical decomposition for the VIX, stock prices, and industrial production. Absent uncertainty

Figure 5: Historical decomposition around the GFC shock

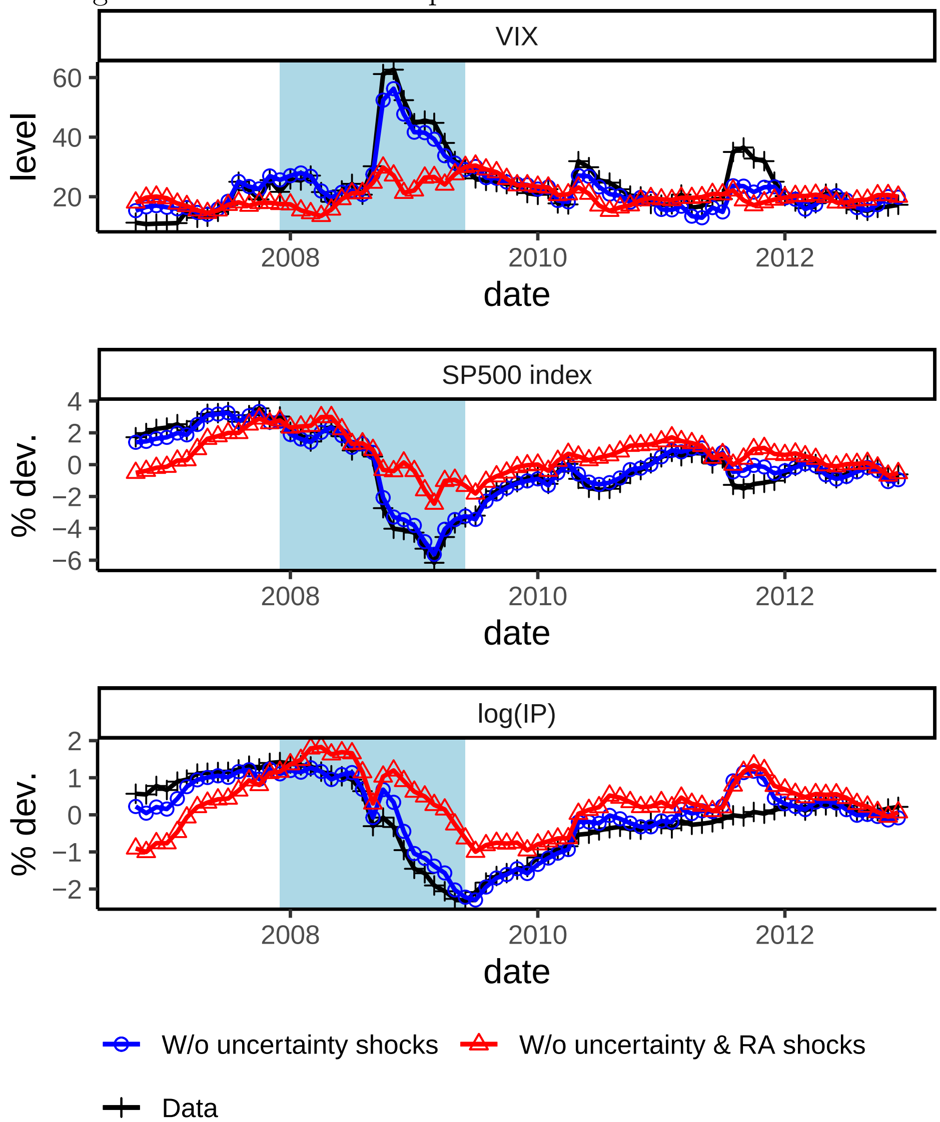

Notes:
The graph plots the historical decomposition of (detrended) log of industrial production, log of S&P500 index, and the VIX around the GCF shock (September-October 2008). The red and blue areas represent respectively the contribution of the set-identified uncertainty and risk aversion shocks in driving the variations of the three selected variables. Mathematically, one can rewrite the observed data as the cumulative sum of structural shocks:$\mathbf{Y}_2 = C + \Phi\mathbf{Y}_1 + B\epsilon_2$. Iterating forward, we can write:$\mathbf{Y}_T = \Phi^{T-1}\mathbf{Y}_1 + \sum_{j=0}^{T-2}\Phi^j(C + B\epsilon_{T-j})$. The contribution of the different types of shocks are retrieved by starting with a B matrix full of zeroes and iteratively fill it with the identified values from the median target specification.

shocks, the VIX would have increased less and remained significantly lower after the initial shock. This is consistent with uncertainty remaining high after the initial shock. We also find that stock prices and industrial production would have both decreased significantly less in the absence of uncertainty shocks. Interestingly, we find that the negative contribution of uncertainty shocks (and to a smaller extent risk aversion) remain large long after the initial shock, suggesting that the uncertainty that followed the initial COVID shock acted as an important drag on the economic recovery. For risk aversion shocks, we find that, even though the initial increase in the VIX seems to have coincided with a large risk aversion shock, risk aversion does not seem to drive the subsequent variations. Albeit smaller than financial uncertainty, risk aversion shocks also appear to reduce industrial production. We find that risk aversion and uncertainty can essentially account for the entirety of the decline in asset prices, thereby confirming that the two types of shocks are important drivers of stock prices. Overall, in line with Baker et al. (2020), we find that uncertainty was an important driver of variations in output around the COVID recession. However, we nuance this finding by emphasizing that risk aversion also played a role.

The results contrast with the application to the GFC where most of the dynamics were driven by risk aversion rather than financial uncertainty. An interpretation is that, because of its different nature, the COVID pandemic did not impact the investors' preferences to the same extent, but rather acted through an immense increase in uncertainty which resulted in a slowdown of the economy, for instant through a real-options channel (see e.g. Bloom (2014)). As such, our results do not allow to discriminate between the different potential channels through which uncertainty may affect the economy, but helps to determine the fundamental sources of the recessions.

8.3 Persistence

To conclude, we propose some evidences on the heterogeneity in terms of the persistence of both types of shocks. In the literature, it is usually expected that the effect of uncertainty

Figure 6: Historical decomposition around the COVID shock

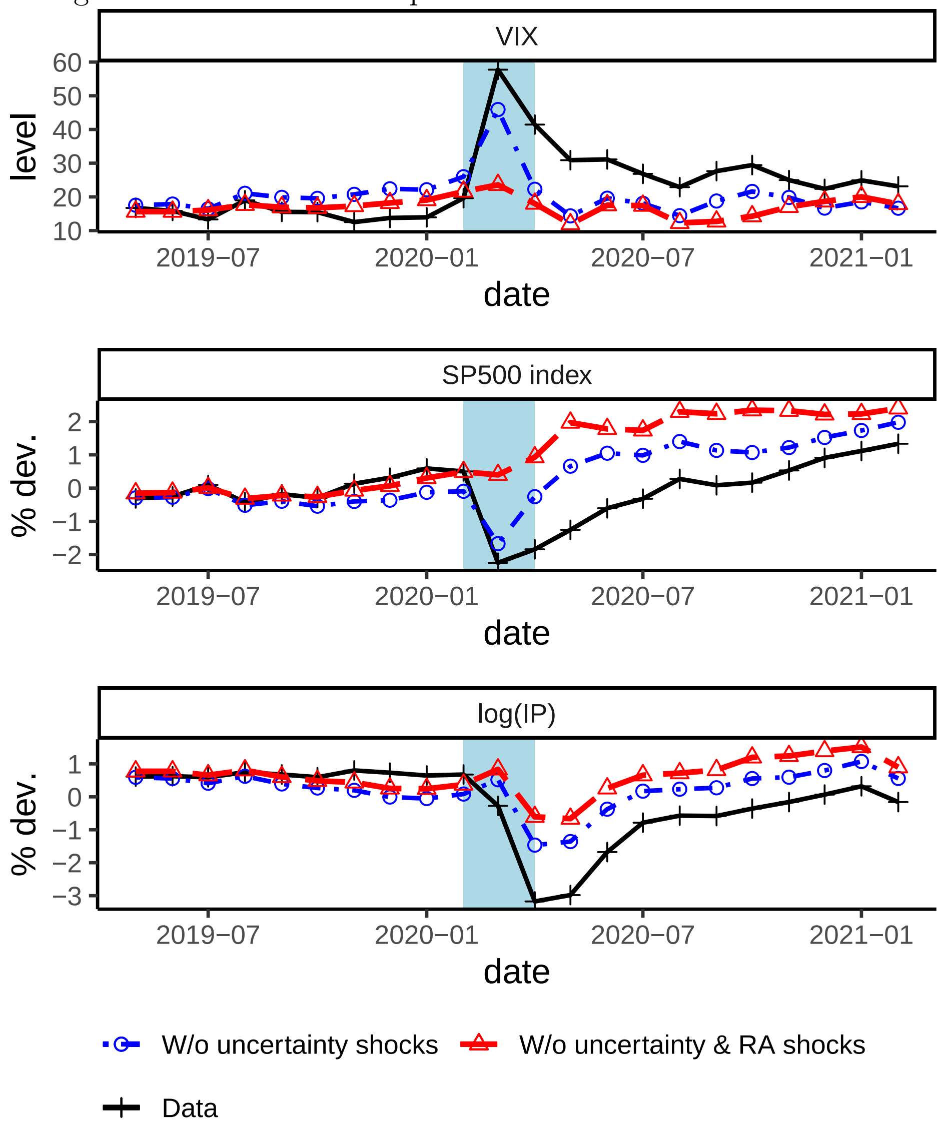

Notes:
The graph plots the historical decomposition of (detrended) log of industrial production, log of S&P500 index, and the VIX around the COVID shock (February-April 2020). The red and blue areas represent respectively the contribution of the set-identified uncertainty and risk aversion shocks in driving the variations of the three selected variables. Mathematically, one can rewrite the observed data as the cumulative sum of structural shocks:$\mathbf{Y}_2 = C + \Phi\mathbf{Y}_1 + B\epsilon_2$. Iterating forward, we can write:$\mathbf{Y}_T = \Phi^{T-1}\mathbf{Y}_1 + \sum_{j=0}^{T-2}\Phi^j(C + B\epsilon_{T-j})$. The contribution of the different types of shocks are retrieved by starting with a B matrix full of zeroes and iteratively fill it with the identified values from the median target specification.

shocks is transitory (see e.g. Bloom et al. (2018)). Due to the slow-moving nature of preferences, it is reasonable to expect the effects of risk aversion shocks to be more persistent over time. However, to the best of our knowledge, no existing studies demonstrate this. To shed light on these questions, we plot the historical contribution of the two types of shocks to the evolution of stock prices over time in Figure 7. A positive reading indicates that a given type of shock contributed positively to the dynamics of stock prices. A striking feature of the contribution of risk aversion shocks is that it exhibits a clear boom-and-bust pattern that appears to coincide with the two financial crises in the sample (the dotcom bubble and the GFC). In other terms, it appears that risk aversion shocks have a persistent effect on asset prices and can partly explain the boom-and-bust pattern of asset prices observed around financial crises. The boom-and-bust pattern is not as clear for the COVID recessions, which suggests its different nature. On the other hand, the effect of uncertainty shocks on asset prices exhibit much less persistence, suggesting that uncertainty shocks only have a transitory effect.

9 Conclusion

This paper aimed to better identify the sources of financial volatility and quantify their macroeconomic implications. The starting point was to note that financial volatility can increase either because the economy becomes physically more volatile, or simply because agents become less willing to bear the risk. The main idea of the paper was to jointly identify uncertainty and risk aversion shocks within a unified framework using set identification. Given two valid proxies, set identification only requires a minimal set of assumptions and allows to credibly identify exogenous shocks to uncertainty and risk aversion. For the proxies, we adopted a narrative approach using volatility events. In this way, we are more confident to isolate exogenous variations in uncertainty and risk aversion. Together with the set identification approach, the two proxies allow to quantify the dynamic effects of uncertainty

Figure 7: Historical contributions to stock prices

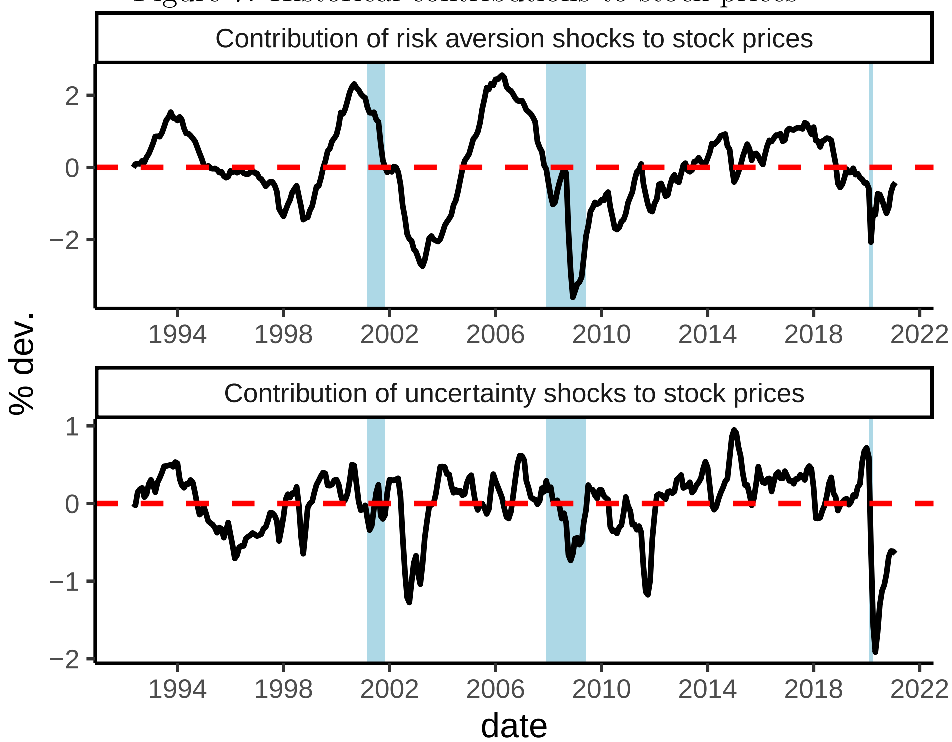

Notes:
The graph plots the historical contribution of the uncertainty and risk aversion shocks to variations in detrended log of stock prices. When the value is above (below) 0, it means that uncertainty or risk aversion shocks contribute positively (negatively) to industrial production. The contributions are retrieved by performing a historical decomposition using the median target specification.

and risk aversion shocks and their contribution to business cycle fluctuations. Consistent with a growing literature, we find that financial uncertainty shocks are significantly recessionary, and possibly more than what is commonly accepted. We find that, during the GFC, risk aversion shocks played a more important role than financial uncertainty. On the other hand, during the COVID recession, we find that financial uncertainty shocks can explain a significant part of the initial decline in production but also acted as an important drag on the economic recovery. Overall, our results show that uncertainty and risk aversion shocks are important drivers of real and financial variables and that the distinction between the two is important.

References

Alessandri, Piergiorgio, Andrea Gazzani, and Alejandro Vicondoa. 2021. "The Real Effects of Financial Uncertainty Shocks: A Daily Identification Approach."

Andersen, Torben G, Tim Bollerslev, Peter F Christoffersen, and Francis X Diebold. 2013. "Financial Risk Measurement for Financial Risk Management." In *Handbook of the Economics of Finance*, 2:1127–220. Elsevier.

Bachmann, Rüdiger, Steffen Elstner, and Eric R Sims. 2013. "Uncertainty and Economic Activity: Evidence from Business Survey Data." *American Economic Journal: Macroeconomics* 5 (2): 217–49.

Baker, Scott R, and Nicholas Bloom. 2013. "Does Uncertainty Reduce Growth? Using Disasters as Natural Experiments." Working Paper 19475. Working Paper Series. National Bureau of Economic Research. https://doi.org/10.3386/w19475.

Baker, Scott R, Nicholas Bloom, Steven J Davis, and Stephen J Terry. 2020. "Covid-Induced Economic Uncertainty." National Bureau of Economic Research.

Basu, Susanto, and Brent Bundick. 2017. "Uncertainty Shocks in a Model of Effective Demand." *Econometrica* 85 (3): 937–58. https://doi.org/https://doi.org/10.3982/ECTA13960.

Bekaert, Geert, Eric C Engstrom, and Nancy R Xu. 2021. "The Time Variation in Risk Appetite and Uncertainty." *Management Science*.

Bekaert, Geert, Eric Engstrom, and Steven R Grenadier. 2010. "Stock and Bond Returns with Moody Investors." *Journal of Empirical Finance* 17 (5): 867–94.

Bekaert, Geert, Eric Engstrom, and Yuhang Xing. 2009. "Risk, Uncertainty, and Asset Prices." *Journal of Financial Economics* 91 (1): 59–82.

Bekaert, Geert, and Marie Hoerova. 2014. "The VIX, the Variance Premium and Stock Market Volatility." *Journal of Econometrics* 183 (2): 181–92.

Bekaert, Geert, Marie Hoerova, and Marco Lo Duca. 2013. "Risk, Uncertainty and Monetary Policy." *Journal of Monetary Economics* 60 (7): 771–88. https://doi.org/https://doi.org/

10.1016/j.jmoneco.2013.06.003.

Berger, David, Ian Dew-Becker, and Stefano Giglio. 2019. "Uncertainty Shocks as Second-Moment News Shocks." *The Review of Economic Studies* 87 (1): 40–76. https://doi.org/10.1093/restud/rdz010.

Bernanke, Ben S. 1983. "Irreversibility, Uncertainty, and Cyclical Investment." *The Quarterly Journal of Economics* 98 (1): 85–106.

Bloom, Nicholas. 2009. "The Impact of Uncertainty Shocks." *Econometrica* 77 (3): 623–85. https://doi.org/https://doi.org/10.3982/ECTA6248.

———. 2014. "Fluctuations in Uncertainty." *Journal of Economic Perspectives* 28 (2): 153–76. https://doi.org/10.1257/jep.28.2.153.

Bloom, Nicholas, Max Floetotto, Nir Jaimovich, Itay Saporta-Eksten, and Stephen J. Terry. 2018. "Really Uncertain Business Cycles." *Econometrica* 86 (3): 1031–65. https://doi.org/https://doi.org/10.3982/ECTA10927.

Bollerslev, Tim. 1986. "Generalized Autoregressive Conditional Heteroskedasticity." *Journal of Econometrics* 31 (3): 307–27. https://doi.org/https://doi.org/10.1016/0304-4076(86)90063-1.

Caggiano, Giovanni, Efrem Castelnuovo, Silvia Delrio, and Richard Kima. 2021. "Financial Uncertainty and Real Activity: The Good, the Bad, and the Ugly." *European Economic Review* 136: 103750.

Caldara, Dario, Cristina Fuentes-Albero, Simon Gilchrist, and Egon Zakrajšek. 2016. "The Macroeconomic Impact of Financial and Uncertainty Shocks." *European Economic Review* 88: 185–207. https://doi.org/https://doi.org/10.1016/j.euroecorev.2016.02.020.

Campbell, John Y, and John H Cochrane. 1999. "By Force of Habit: A Consumption-Based Explanation of Aggregate Stock Market Behavior." *Journal of Political Economy* 107 (2): 205–51.

Carriero, Andrea, Haroon Mumtaz, Konstantinos Theodoridis, and Angeliki Theophilopoulou. 2015. "The Impact of Uncertainty Shocks Under Measurement Error: A Proxy SVAR

Approach." *Journal of Money, Credit and Banking* 47 (6): 1223–38. https://doi.org/https://doi.org/10.1111/jmcb.12243.

Chiu, Ching-wai Jeremy, Richard DF Harris, Evarist Stoja, and Michael Chin. 2018. "Financial Market Volatility, Macroeconomic Fundamentals and Investor Sentiment." *Journal of Banking & Finance* 92: 130–45.

Christiano, Lawrence J, Roberto Motto, and Massimo Rostagno. 2014. "Risk Shocks." *American Economic Review* 104 (1): 27–65.

Cohn, Alain, Jan Engelmann, Ernst Fehr, and Michel André Maréchal. 2015. "Evidence for Countercyclical Risk Aversion: An Experiment with Financial Professionals." *American Economic Review* 105 (2): 860–85.

Davidson, Russell, and Emmanuel Flachaire. 2008. "The Wild Bootstrap, Tamed at Last." *Journal of Econometrics* 146 (1): 162–69.

Di Tella, Sebastian, and Robert Hall. 2021. "Risk Premium Shocks Can Create Inefficient Recessions." *The Review of Economic Studies*, September. https://doi.org/10.1093/restud/rdab049.

Drechsler, Itamar, and Amir Yaron. 2011. "What's Vol Got to Do with It." *The Review of Financial Studies* 24 (1): 1–45.

Eeckhoudt, Louis, Christian Gollier, and Harris Schlesinger. 1996. "Changes in Background Risk and Risk Taking Behavior." *Econometrica: Journal of the Econometric Society*, 683–89.

Epstein, Larry G. 1988. "Risk Aversion and Asset Prices." *Journal of Monetary Economics* 22 (2): 179–92.

Fry, Renee, and Adrian Pagan. 2011. "Sign Restrictions in Structural Vector Autoregressions: A Critical Review." *Journal of Economic Literature* 49 (4): 938–60.

Gârleanu, Nicolae, and Stavros Panageas. 2015. "Young, Old, Conservative, and Bold: The Implications of Heterogeneity and Finite Lives for Asset Pricing." *Journal of Political Economy* 123 (3): 670–85.

Gertler, Mark, and Peter Karadi. 2015. "Monetary Policy Surprises, Credit Costs, and Economic Activity." *American Economic Journal: Macroeconomics* 7 (1): 44–76. https://doi.org/10.1257/mac.20130329.

Gilchrist, Simon, and Egon Zakrajšek. 2012. "Credit Spreads and Business Cycle Fluctuations." *American Economic Review* 102 (4): 1692–1720.

Gonçalves, Sılvia, and Lutz Kilian. 2004. "Bootstrapping Autoregressions with Conditional Heteroskedasticity of Unknown Form." *Journal of Econometrics* 123 (1): 89–120.

Gordon, Stephen, and Pascal St-Amour. 2000. "A Preference Regime Model of Bull and Bear Markets." *American Economic Review* 90 (4): 1019–33.

Guiso, Luigi. 2012. "Trust and Risk Aversion in the Aftermath of the Great Recession." *European Business Organization Law Review (EBOR)* 13 (2): 195–209.

———. 2014. "Risk Aversion and Financial Crisis." EIEF Working Papers Series 1412. Einaudi Institute for Economics; Finance (EIEF). https://ideas.repec.org/p/eie/wpaper/1412.html.

Guiso, Luigi, and Monica Paiella. 2008. "Risk Aversion, Wealth, and Background Risk." *Journal of the European Economic Association* 6 (6): 1109–50.

Gürkaynak, Refet S, Brian P Sack, and Eric T Swanson. 2004. "Do Actions Speak Louder Than Words? The Response of Asset Prices to Monetary Policy Actions and Statements." *The Response of Asset Prices to Monetary Policy Actions and Statements (November 2004)*.

Heaton, John, and Deborah Lucas. 2000. "Portfolio Choice in the Presence of Background Risk." *The Economic Journal* 110 (460): 1–26.

Jurado, Kyle, Sydney C. Ludvigson, and Serena Ng. 2015. "Measuring Uncertainty." *American Economic Review* 105 (3): 1177–1216. https://doi.org/10.1257/aer.20131193.

Kamstra, Mark J, Lisa A Kramer, and Maurice D Levi. 2003. "Winter Blues: A SAD Stock Market Cycle." *American Economic Review* 93 (1): 324–43.

Känzig, Diego R. 2021. "The Macroeconomic Effects of Oil Supply News: Evidence from

OPEC Announcements." *American Economic Review*.

Kuttner, Kenneth N. 2001. "Monetary Policy Surprises and Interest Rates: Evidence from the Fed Funds Futures Market." *Journal of Monetary Economics* 47 (3): 523–44.

Kydland, Finn E, and Edward C Prescott. 1982. "Time to Build and Aggregate Fluctuations." *Econometrica: Journal of the Econometric Society*, 1345–70.

Longstaff, Francis A, and Jiang Wang. 2012. "Asset Pricing and the Credit Market." *The Review of Financial Studies* 25 (11): 3169–3215.

Ludvigson, Sydney C, Sai Ma, and Serena Ng. 2021. "Uncertainty and Business Cycles: Exogenous Impulse or Endogenous Response?" *American Economic Review*.

Martin, Ian. 2017. "What Is the Expected Return on the Market?" *The Quarterly Journal of Economics* 132 (1): 367–433.

Menzly, Lior, Tano Santos, and Pietro Veronesi. 2004. "Understanding Predictability." *Journal of Political Economy* 112 (1): 1–47.

Mertens, Karel, and Morten O. Ravn. 2013. "The Dynamic Effects of Personal and Corporate Income Tax Changes in the United States." *American Economic Review* 103 (4): 1212–47. https://doi.org/10.1257/aer.103.4.1212.

Miranda-Agrippino, Silvia, and Hélène Rey. 2021. "The Global Financial Cycle." National Bureau of Economic Research.

Montiel-Olea, José L, James H Stock, and Mark W Watson. 2016. "Uniform Inference in SVARs Identified with External Instruments." *Manuscript, Columbia University*.

Muir, Tyler. 2017. "Financial Crises and Risk Premia." *The Quarterly Journal of Economics* 132 (2): 765–809.

Piffer, Michele, and Maximilian Podstawski. 2017. "Identifying Uncertainty Shocks Using the Price of Gold." *The Economic Journal* 128 (616): 3266–84. https://doi.org/10.1111/ecoj.12545.

Ramey, Valerie A. 2016. "Macroeconomic Shocks and Their Propagation." *Handbook of Macroeconomics* 2: 71–162.

Smets, Frank, and Rafael Wouters. 2007. "Shocks and Frictions in US Business Cycles: A Bayesian DSGE Approach." *American Economic Review* 97 (3): 586–606.

Stock, James H, and Mark W Watson. 2018. "Identification and Estimation of Dynamic Causal Effects in Macroeconomics Using External Instruments." *The Economic Journal* 128 (610): 917–48.

Wachter, Jessica A. 2006. "A Consumption-Based Model of the Term Structure of Interest Rates." *Journal of Financial Economics* 79 (2): 365–99.

Xu, Shaofeng. 2016. "Interpreting Volatility Shocks as Preference Shocks." *Economics Letters* 149: 135–40.

Online Appendix (not intended for publication)

A Econometric framework

A.1 Intuition

To build the intuition, we first consider the simple case with one external instrument and one shock to identify. In a second step, we consider the more general case with $l > 1$ instruments and shocks to identify. We further detail how the set identification of uncertainty and risk aversion shocks is achieved.

Simple case with one instrument: Let $\mathbf{Y}_t$ be a $k \times 1$ vector of economic and financial variables. We consider the following structural representation of the VAR:

$$\mathbf{Y}_t = C + \Phi \mathbf{Y}_{t-1} + S\epsilon_t \tag{16}$$

In particular, we model the reduced form residuals of the VAR u_t as a linear function of the structural shocks, that is $u_t = S\epsilon_t$. The structural shocks ϵ_t are white noise processes with variance Σ_ϵ. It follows that the variance-covariance matrix of the reduced form residuals is equal to $\Sigma_u \equiv E(u_t u_t') = E[S\Sigma_\epsilon S']$.

The identification of the structural effect of uncertainty shocks by means of an external instrument is achieved as follows. For expositional purpose, assume that the variable we want to instrument is ordered first. We are interested in the first column of S, which we denote s_1. This column tells us how the first structural shock $\epsilon_{1,t}$ affects all the other variables contained in the VAR. Once s_1 is known, we can easily derive several objects of interest such as IRFs and historical decompositions.

Assume there exists an external instrument z_t that is both relevant and exogenous. In other terms, it should be sufficiently correlated with the shock of interest but orthogonal to the other structural shocks of the model. Formally:

$$E(z_t\epsilon_{1t}) = \alpha \neq 0 \tag{17}$$

$$E(z_t\epsilon_{2:k,t}) = 0_{k-1} \tag{18}$$

With $\epsilon_{1,t}$ the first structural shock of ϵ_t and $\epsilon_{2:k,t}$ the remaining structural shocks. We can partition the first column as:

$$s_1 = \begin{bmatrix} s_{11} \\ \\ s_{2:n} \end{bmatrix} \tag{19}$$

Let denote $\tilde{s}_{2:n,1} \equiv s_{2:n,1}/s_{11}$ and note that:

$$\tilde{s}_{2:n,1} = E(z_t u_{2:n,t})/E(z_t u_{1,t}) \tag{20}$$

Equation (20) holds in population. In sample, it corresponds to the IV estimator of a regression of $\widehat{u}_{2:n,t}$ on $\widehat{u}_{1,t}$ using z_t as an instrument where $\widehat{u_t} = \begin{bmatrix} \widehat{u}_{1,t} & \widehat{u}'_{2:n,t} \end{bmatrix}'$ are the reduced form residuals from the VAR. To scale our coefficients, we assume that $\Sigma_\epsilon = I_n$.[6] In this way, structural shocks correspond to a one-standard deviation increase. It follows that $\Sigma_u = SS'$ and can be written as:

$$\Sigma_u \equiv \begin{bmatrix} s_{11} & \mathbf{s}_{12} \\ \\ \mathbf{s}_{21} & \mathbf{S}_{22} \end{bmatrix} \begin{bmatrix} s_{11} & \mathbf{s_{12}}' \\ \\ \mathbf{s}'_{21} & \mathbf{S}'_{22} \end{bmatrix} = \begin{bmatrix} \sigma_{11} & \boldsymbol{\sigma_{12}} \\ \\ \boldsymbol{\sigma_{21}} & \boldsymbol{\Sigma_{22}} \end{bmatrix} \tag{21}$$

We can show that:

$$s_{11} = \pm\sqrt{\sigma_{11} - s_{12}s'_{12}} \tag{22}$$

[6] Other normalizations are possible.

and that:

$$s_{12}s'_{12} = (\sigma_{21} - \tilde{s}_{21}\sigma_{11})\left[\Sigma_{22} - (\tilde{s}_{21}\sigma_{12} + \sigma_{21}\tilde{s}'_{21}) + \sigma_{11}\tilde{s}_{21}\tilde{s}'_{21}\right]^{-1}$$
$$(\sigma_{21} - \tilde{s}_{21}\sigma_{11}) \tag{23}$$

Once we have (23), we can easily compute (22). This allows us to fully characterize s_1, and thus to compute all the objects of interest such as IRFs, FEVD or historical decompositions. This external instrument approach is detailed in Stock and Watson (2018) and is applied in Gertler and Karadi (2015) in the context of monetary policy shocks. More recently, Känzig (2021) uses a similar approach to identify oil supply news shocks.

Set identification: We now turn to the description of the set identification approach. For simplicity, consider that we want to jointly identify 2 independent shocks with $l = 2$ proxies. Without loss of generality, assume that the two shocks relate to the first and second variables of the VAR. Let us rewrite the reduced form residuals as follows:

$$u_t = s^1\epsilon_t^1 + s^2\epsilon_t^2 + \tilde{S}\tilde{\epsilon}_t \tag{24}$$

With b^1 and b^2 two column vectors that collect the effect of the first and second structural shocks on all the variables contained in the VAR and $\tilde{\epsilon}_t$ a vector collecting the $k-2$ remaining shocks. Assume now that we have two valid proxies z_t^1 and z_t^2 at disposition. By valid proxies, we mean that the proxies respect the relevance and exogeneity conditions expressed in (17) and (18) with the exception that each proxy is allowed to be correlated with the shocks that we aim to identify. That is:

$$E(z_t^1\epsilon_t^1) = \alpha_1 \neq 0 \tag{25}$$

$$E(z_t^2\epsilon_t^2) = \alpha_2 \neq 0 \tag{26}$$

$$E(z_t^1\tilde{\epsilon}_t) = E(z_t^2\tilde{\epsilon}_t) = 0_{k-2} \tag{27}$$

With $E(z_t^1\epsilon_t^2)$ and $E(z_t^2\epsilon_t^1)$ not necessarily equal to zero, that is the proxy is allowed to be correlated with the other structural shock. We define:

$$\Phi \equiv \begin{bmatrix} \phi_{11} & \phi_{12} \\ \phi_{21} & \phi_{22} \end{bmatrix} = \begin{bmatrix} E(\epsilon_t^1 z_t^1) & E(\epsilon_t^1 z_t^2) \\ E(\epsilon_t^2 z_t^1) & E(\epsilon_t^2 z_t^2) \end{bmatrix} \tag{28}$$

Using the normalization $E\left((\epsilon_t^1)^2\right) = E\left((\epsilon_t^2)^2\right) = 1$ and $E\left((z_t^1)^2\right) = E\left((z_t^2)^2\right) = 1$, it follows that Φ can be interpreted as the correlation structure between the shocks of interest and the instruments. Identification of ϵ_t^1 and ϵ_t^2 is achieved by imposing restrictions on Φ. In our setting, the first two variables in our VAR are the VIX and the S&P500 index (in that order), and the two elements we aim to identify are uncertainty and risk aversion shocks, respectively. Assuming that we have one proxy for each type of shock, the identification can be achieved by assuming i) that the uncertainty shock is more correlated with the uncertainty proxy than with risk aversion proxy, that is $\phi_{11} - \phi_{12} > \psi > 0$ for ψ some positive constant, ii) that the risk aversion shock is more correlated with the risk aversion proxy than with uncertainty proxy, that is $\phi_{22} - \phi_{21} > \psi > 0$, and iii) that both proxies are positively correlated with the shocks they aim to identify, that is $\phi_{11} > 0$ and $\phi_{22} > 0$. In our benchmark, we use $\psi = 0$ but check that the results are robust to alternative values of ψ. Note that the identifiying restrictions do not require both proxies to be independent from each other. Given valid proxies, these identifying restrictions are rather mild and provide a flexibly way to identify the structural model without imposing any direct restrictions on the impulse response functions.

A.2 In practice

To estimate S, the first step is to estimate the reduced form residuals $\widehat{u}_t$ and the sample variance-covariance matrix $\widehat{\Sigma}$ using the VAR.

Second, we define $Z_t = \begin{bmatrix} z_t^{RA} \\ z_t^{U} \end{bmatrix}$ as the vector of proxies. We further define $\widehat{E} = \widehat{E}(\widehat{u}_t Z_t')$. The hat refers to the sample equivalent of the population value. Let $\widehat{E}_{11}$ be the first l elements of $\widehat{E}$. $\widehat{E}_{11}$ is thus of dimension $l \times l$. Let $\widehat{E}_{21}$ be the last $k-l$ elements, and is thus of dimension $(k-l) \times l$.

In population, we have $E(u_t Z_t') = S\epsilon_t Z_t' + 0$. Thus, $E_{11} = S_{11}\epsilon_t Z_t' = B_{11}\Phi$. Also, $E_{21} = S_{21}\epsilon_t Z_t' = S_{21}\Phi$. Recall that $(AB)^{-1} = B^{-1}A^{-1}$. Thus, $E_{21}E_{11}^{-1} = S_{21}\Phi\Phi^{-1}S_{11}^{-1} = S_{21}S_{11}^{-1}$. We can now define $\widehat{\Lambda} = \widehat{E}_{21}\widehat{E}_{11}^{-1}$ as a consistent estimator of $S_{21}S_{11}^{-1}$.

Now, we partition the matrix S as follows:

$$S = \begin{bmatrix} S_{11} & S_{12} \\ S_{21} & S_{22} \end{bmatrix} \tag{29}$$

$$\Sigma = \begin{bmatrix} \Sigma_{11} & \Sigma_{12} \\ \Sigma_{21} & \Sigma_{22} \end{bmatrix} \tag{30}$$

With S_{11} is $l \times l$, S_{22} is $(k-l) \times (k-l)$, S_{12} is $l \times (k-l)$ and S_{21} is $(k-l) \times l$. Similarly for Σ.

We can estimate $\widehat{\tilde{S}}$ (a $k \times l$ matrix) using:

$$\widehat{\tilde{S}} = \begin{bmatrix} \widehat{S}_{11} \\ \widehat{\Lambda}\widehat{S}_{11} \end{bmatrix} \tag{31}$$

We can show that:

$$\widehat{S}_{11}\widehat{S}_{11}' = \widehat{\Sigma}_{11} - \widehat{S_{12}S_{12}}' \tag{32}$$

And that:

$$\widehat{S_{12}S_{12}}' = (\widehat{\Sigma}_{21} - \widehat{\Lambda}\widehat{\Sigma}_{11})'\widehat{\Gamma}^{-1}(\widehat{\Sigma}_{21} - \widehat{\Lambda}\widehat{\Sigma}_{11}) \tag{33}$$

With:

$$\widehat{\Gamma} = \widehat{\Sigma}_{22} + \widehat{\Lambda}\widehat{\Sigma}_{11}\widehat{\Lambda}' - \widehat{\Sigma}_{21}\widehat{\Lambda}' - \widehat{\Lambda}\widehat{\Sigma}_{21}' \tag{34}$$

With only one instrument, $\widehat{S}_{11}\widehat{S}_{11}'$ is a scalar and is thus unique up to its sign. Here, as we have $l > 1$ instruments, $\widehat{S}_{11}\widehat{S}_{11}'$ is a matrix and is thus not uniquely identified. This gives rise to an identification problem that can be adressed using set identification. In particular, we have $\widehat{S}_{11} = \widehat{S}_{11}^{C}Q$ where $\widehat{S}^{C}$ is the lower Cholesky decomposition of $\widehat{\Sigma}_{11} - \widehat{S_{12}S_{12}}'$ and Q is a random orthornormal matrix according to the Haar measure. Mertens and Ravn (2013) identify B using the special case $Q = I$ and a recursive ordering. Here, identification is achieved by imposing restrictions on the correlation structure. In particular, we identify $\widehat{B}$ by imposing restrictions on the estimated moments $\Phi = \widehat{S}_{11}^{-1}\widehat{E}_{11}$ as is explained in the main text.

Without loss of generality, we order the uncertainty variable first and the S&P500 variable second. We normalize the element [1,1] of $\widehat{S}_{11}$ to be positive. Similarly, we normalize the element [2,2] of $\widehat{S}_{11}$ to be negative. This ensures that a positive uncertainty shocks leads to an *increase* of the uncertainty variable, whereas a positive risk aversion shock leads to a *decrease* in the S&P500 variable.

B Robustness checks

B.1 Different set of events

In this section, we perform a few robustness checks to check the validity of the results. Because the proxies are built using a narrative approach, a potential concern is that it is the particular choice of events that drives the results. We thus reestimate the model by using different set of events. In particular, we drop randomly 20% of the selected events and plot the resulting IRFs in Figure 8 and 9. As we can see, results remain similar with different set of events.

Figure 8: IRFs from an uncertainty shocks with different selected events

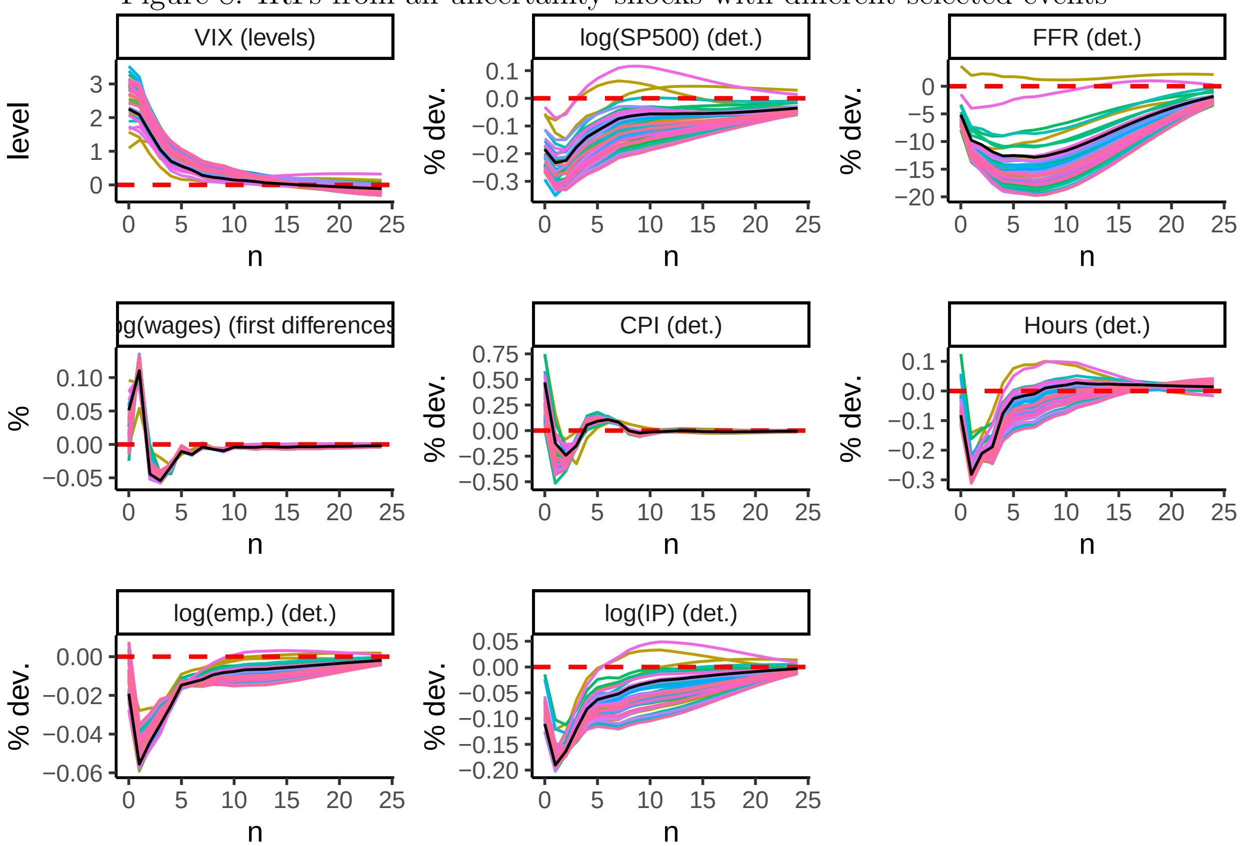

Notes:
For each of the IRFs, we randomly drop 20% of the volatility events and set identify the uncertainty shocks using the new proxies with the adjusted list of events. We repeat the procedure 60 times. The black line indicates the IRF from the baseline specification.

Figure 9: IRFs from a risk aversion shocks with different selected events

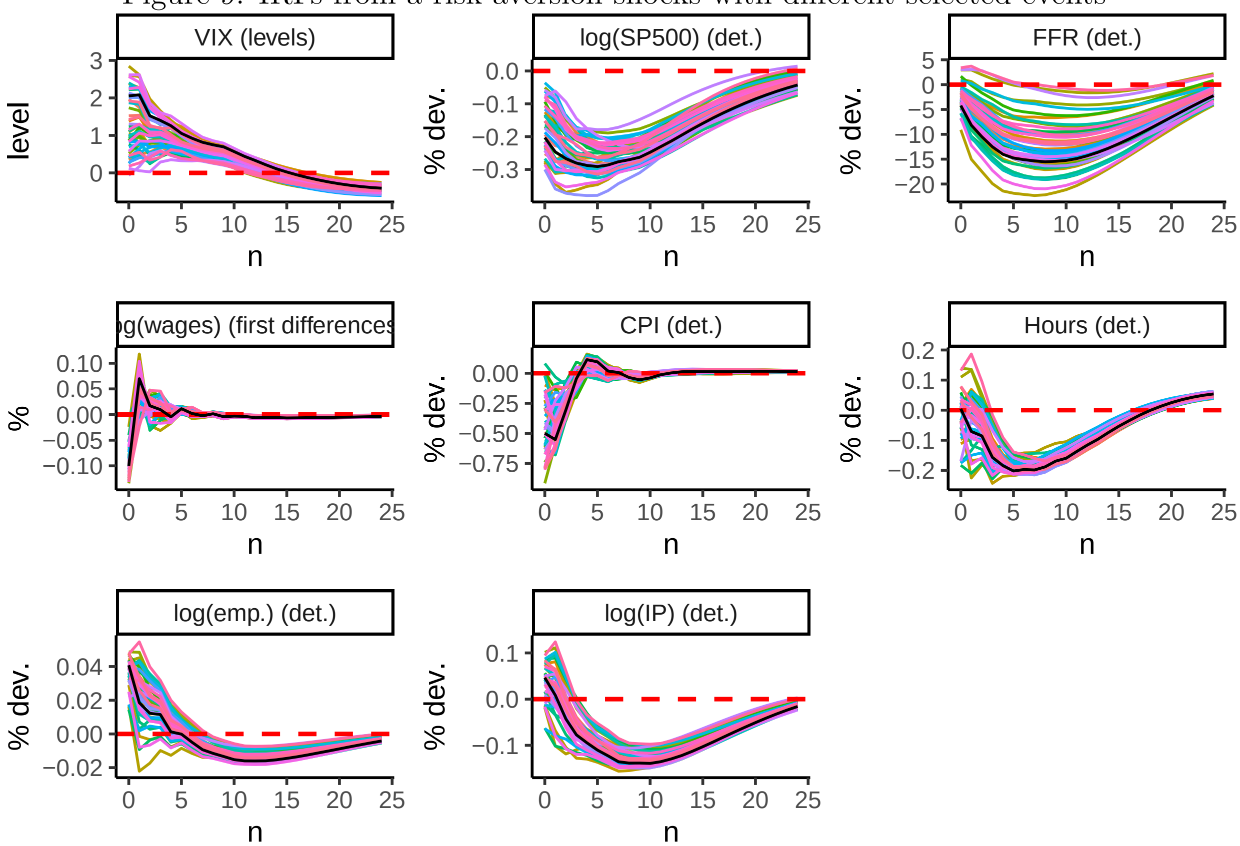

Notes:
For each of the IRFs, we randomly drop 10% of the volatility events and set identify the risk aversion shocks using the new proxies with the adjusted list of events. We repeat the procedure 60 times. The black line indicates the IRF from the baseline specification.

C Figures not included in the text

Figure 10: Variations in risk aversion and uncertainty around the election of Donald Trump.

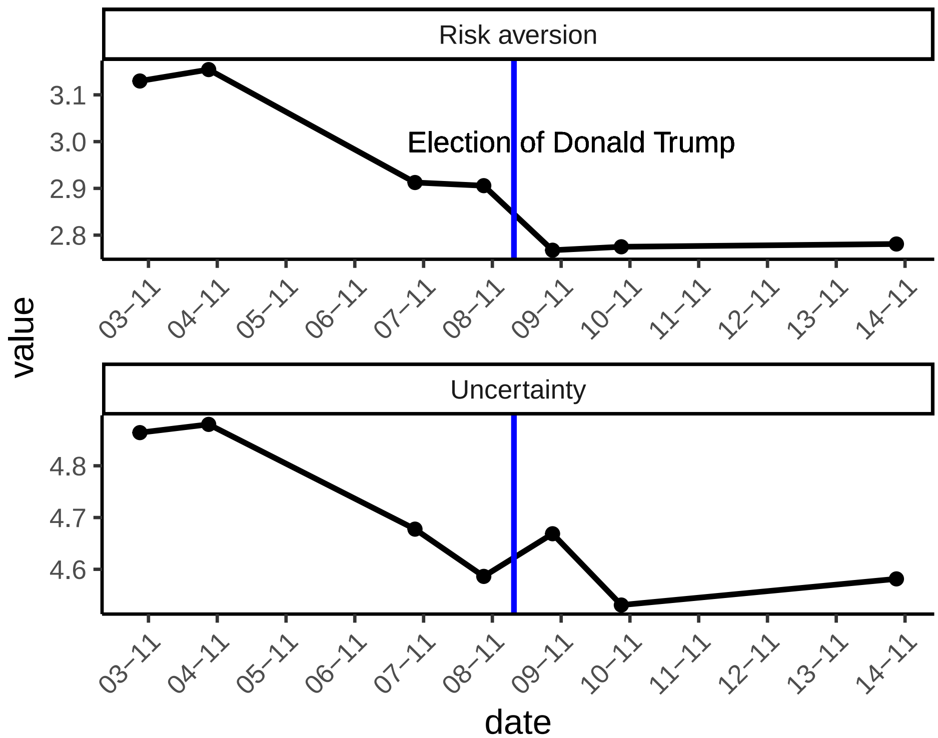

Notes: The timing of the election of Donald Trump is retrieved using Bloomberg News. The uncertainty and risk aversion indices are taken from Bekaert et al. (2021)

Table 4: List of volatility events with variations in the uncertainty and risk aversion proxies

Date	Event	RA shock	Unc. shock
1994-04-07	Genocide in Rwanda	-0.01	-0.77
1997-07-02	Thailand unpegs currency	-0.07	-2.51
1998-08-07	US embassy bombings in Kenia and Tanzania	0.14	0.02
1998-09-23	LTCM default	0.20	1.27
1999-03-24	Clinton announces US join NATO bombing in Kosovo	-0.02	-0.20
2000-09-28	Second Intifada	0.06	0.36
2000-11-08	George W. Bush declared winner in Florida	0.04	0.07
2001-09-11	9/11 attack	0.68	3.08
2001-12-23	Argentine default	-0.03	-0.14
2002-07-21	Worldcom bankruptcy	0.64	1.06
2004-03-11	Madrid train bombings	0.09	0.13
2005-07-07	London bombing	0.00	1.26
2006-07-12	2nd Lebanon War	0.04	0.79
2007-09-14	Northern Rock receives liquidity support by BoE	0.03	1.41
2008-02-17	Northern Rock in state ownership	0.02	0.40
2008-09-17	Emergency lending to AIG	0.66	1.52
2008-11-05	Obama elected	0.68	-2.51
2008-12-27	Israel-Gaza conflict	0.13	0.33
2010-05-10	EFSF adopted	-0.55	-2.51
2011-03-07	UN Security Council establishes no-fly zone in Libya	0.08	0.27
2011-11-12	Berlusconi resigning, beginning of Mario Monti	-0.21	1.29
2012-03-30	Tuareg offensive in Mali after coup starts	0.01	0.46
2012-06-24	Mursi elected president of Egypt	0.11	1.92
2012-09-12	German Court approves ESM	-0.02	-0.68
2013-04-15	Boston marathon bombing	0.17	3.08
2013-10-01	US government shutdown	-0.05	0.65
2014-02-18	Maidan riots in Ukraine	0.02	-0.12
2014-06-10	IS seizes Mosul	-0.01	0.03
2016-06-24	Result from the Brexit vote	0.49	3.08
2016-11-09	Donald Trum elected president	-0.14	1.60
2019-01-15	UK parliament rejects Theresa May's deal	-0.02	0.04
2020-01-30	WHO declares COVID-19 a Public Health Emergency of International Concern	0.15	0.59
2020-03-11	WHO declares Covid-19 a pandemic	0.68	3.08
2020-11-09	Pfizer-Biontech announce COVID-19 Vaccine	-0.06	-0.21

[a] The volatility events are recoverd from an updated database from Piffer and Podstawski (2017).

Figure 11: IRFs following an uncertainty shock (Cholesky)

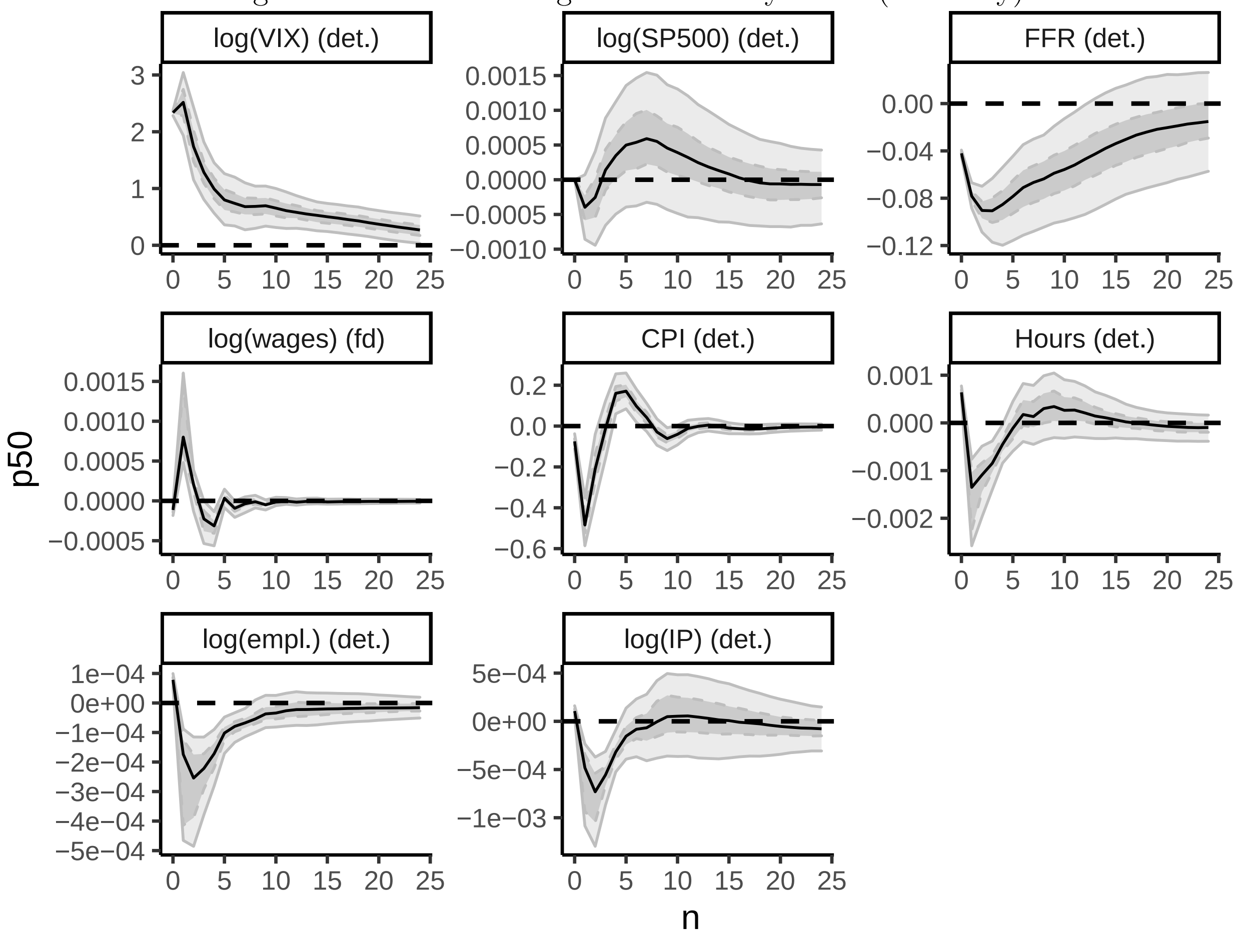

Notes:
The graphs plot the IRFs of the 8 variables contained in the VAR following an shock to the VIX that is recursively identified (VIX ordered second). The median IRFs is based on 1,000 bootstrap replications satisfying the restrictions discussed in Section 3 and is obtained using the median target specification of Fry and Pagan (2014). The confidence intervals are obtained by targeting the 10th, 32th, 68th and 90th percentiles, respectively. The IRFs correspond to a one standard deviation shock on the VIX. Wages and CPI are expressed in log-differences (monthly growth). The VIX is in levels. The other variables are expressed in percentage deviation from the HP trend.

Figure 12: Variables entering the VAR

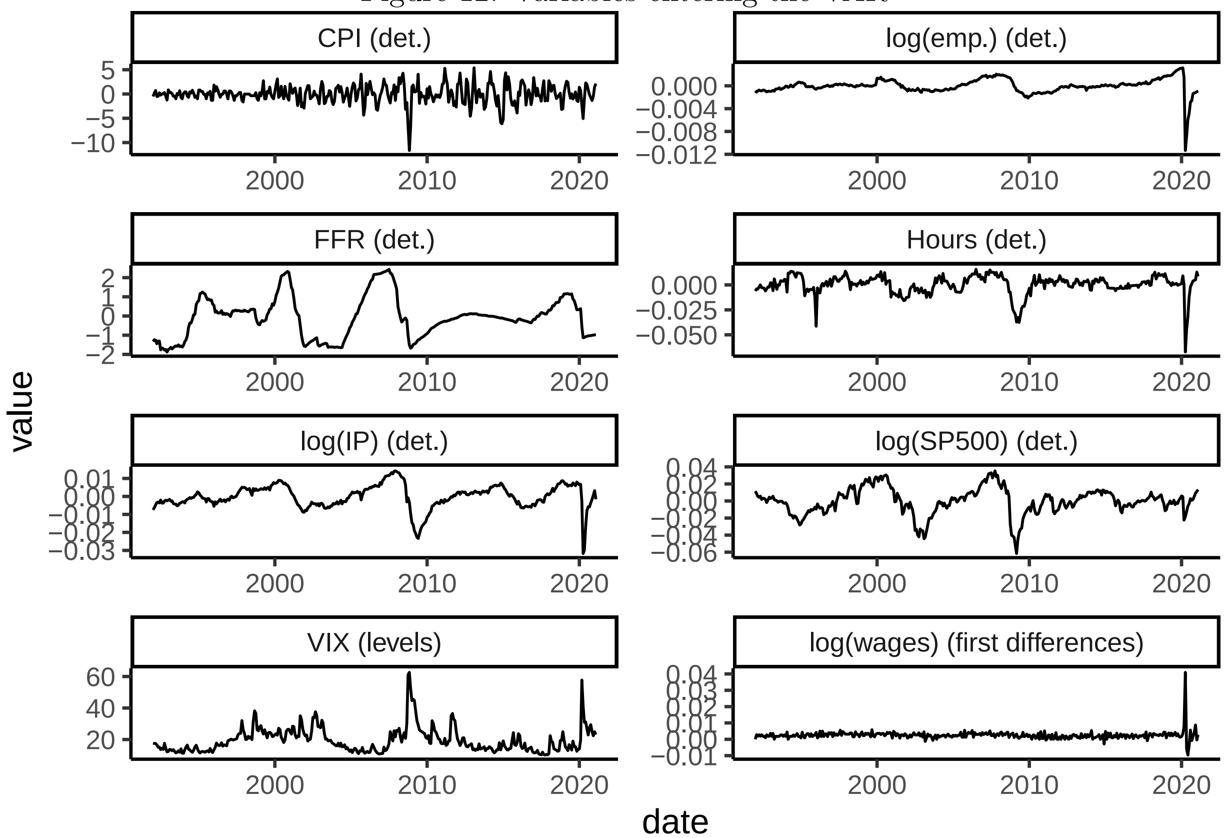

Notes:
Source: Datastream, Fred.

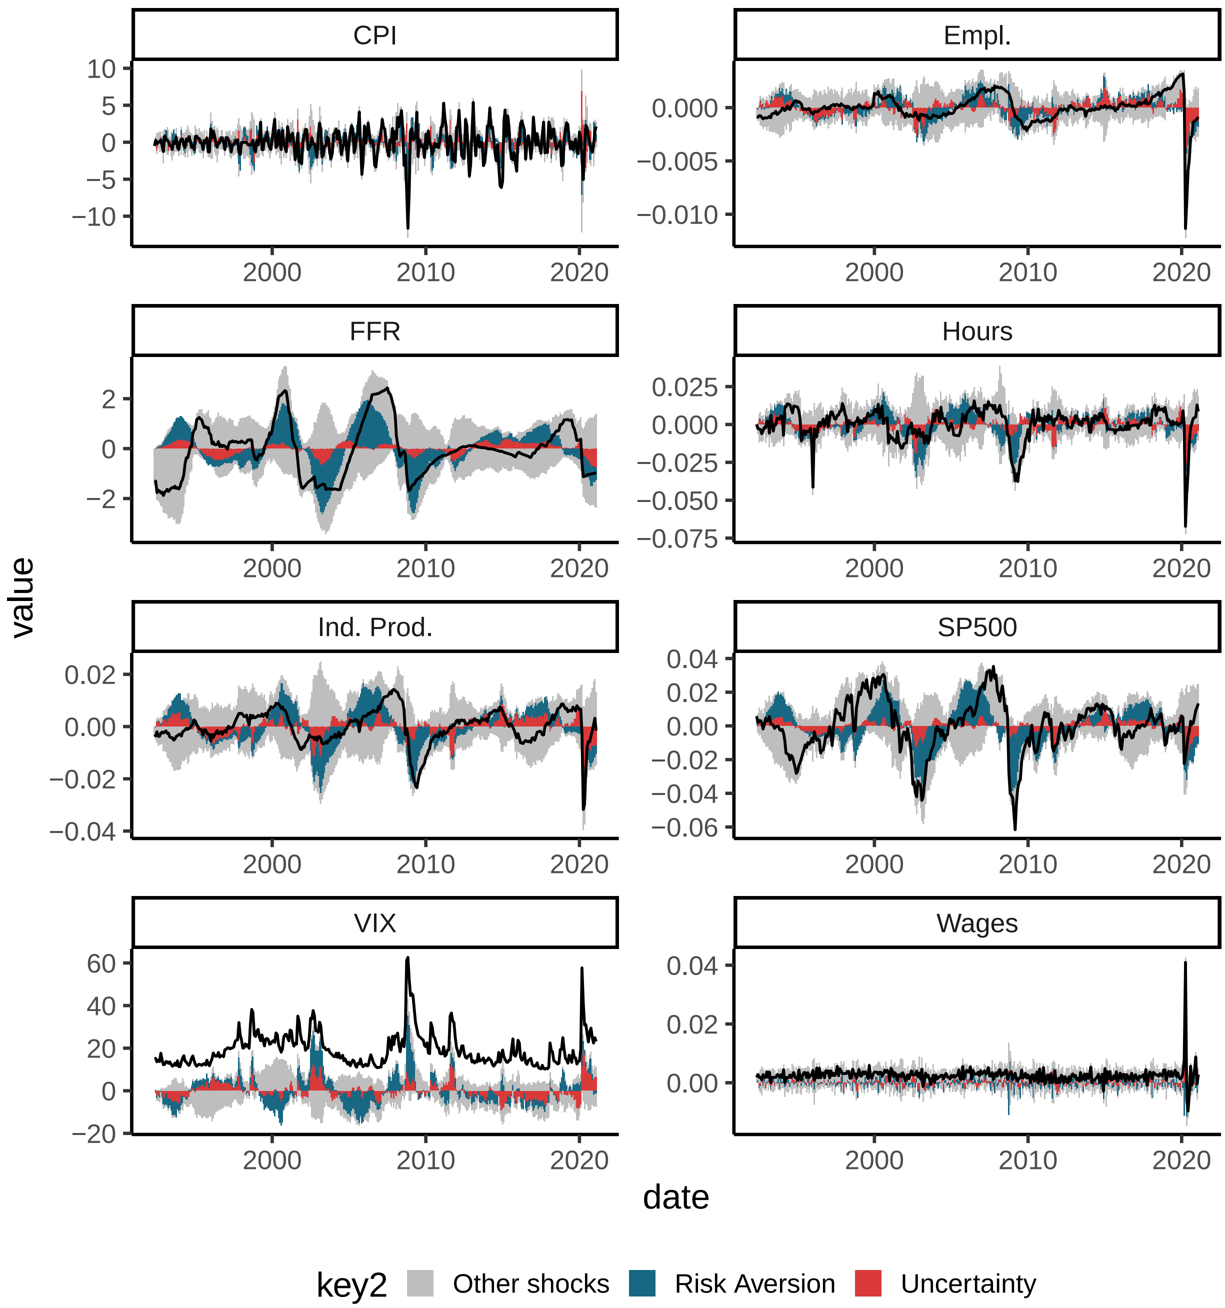

Figure 13: Historical decomposition for all endogenous variables

Figure 14: Structural shocks

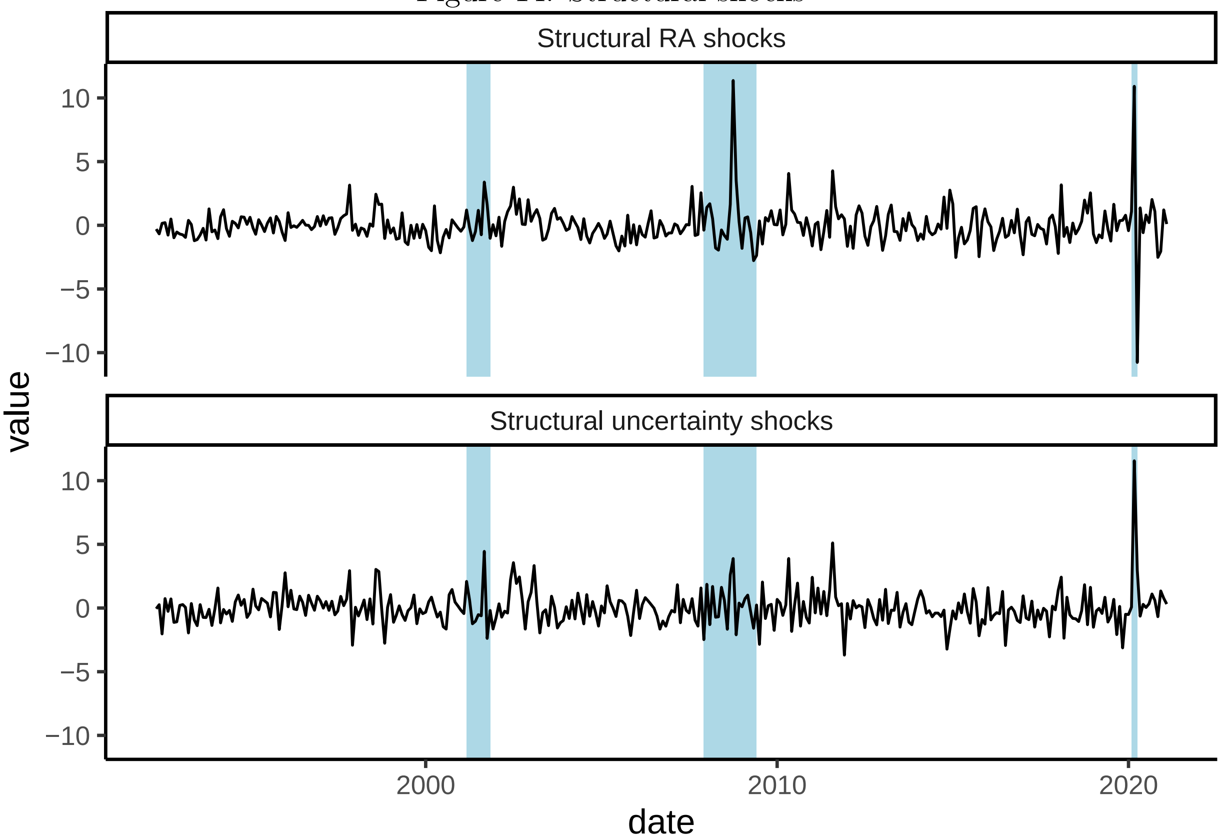

Notes:
The graph plots the set identified structural shocks for uncertainty and risk aversion using the median target specification. To facilitate the reading, a positive risk aversion and uncertainty shock is associated with a decline in asset prices and an increase in the VIX, respectively. NBER recessions are indicated in the shaded lightblue areas.