Macroeconomic Dynamics (2024), **29**, e43, pp. 1–22
doi:10.1017/S1365100524000555

ARTICLE

Shock persistence, uncertainty, and news-driven business cycles

Kevin Lee[1], Kalvinder Shields[2], and Guido Turnip[3]

[1]University of Nottingham, Nottingham, UK
[2]University of Melbourne, Melbourne, Australia
[3]University of Southampton, Iskandar Puteri, Malaysia
Corresponding author: Kalvinder Shields; Email: k.shields@unimelb.edu.au

Abstract

This paper distinguishes news about short-lived events from news about changes in longer term prospects using surveys of expectations. Employing a multivariate GARCH-in-Mean model for the US, the paper illustrates how the different types of news influence business cycle dynamics. The influence of transitory output shocks can be relatively large on impact but gradually diminishes over two to three years. Permanent shocks drive the business cycle, generating immediate stock price reactions and gradually building output effects, although they have more immediate output effects during recessions through the uncertainties they create. Markedly different macroeconomic dynamics are found if these explicitly identified types of news or uncertainty feedbacks are omitted from the analysis.

Keywords: News-driven business cycles; persistence; uncertainty; expectations; surveys

JEL classifications: C32; D84; E32

1. Introduction

A voluminous literature has grown over the last decade considering the extent to which business cycle fluctuations are driven by the arrival of news and the uncertainties surrounding this news.[1] The possibility of "correlated news"—where information flows involve anticipated (news shocks) and unanticipated innovations—and the distinction between "news about the near future" and 'news about the far future' is central to this literature, yet relatively little attention has been paid to the structure of information flows in the applied work exploring the news-driven business cycle.[2] This paper focuses on the business cycle consequences of news about short-lived events and news about changes in longer term prospects using survey-based measures of expectations at different forecast horizons to identify shocks with different persistence properties.

We argue that characterizing shocks according to their persistence properties provides significant insights for understanding business cycles and is at least as useful as the characterization according to source as is conventional in the literature. Examples of correlated news are widespread and have important consequences for macro dynamics. For example, discussion around the use of "forward guidance" has shown the potentially significant role of monetary policy makers' announcements on long-horizon interest rate paths in managing the economy; and on the fiscal side, where it takes time to implement plans and to balance budgets, the importance of government spending announcements for output have been well rehearsed in the literature on 'fiscal foresight' (see Leeper et al. 2013), and it is well understood that tax cuts generate a larger output response if they are expected to be longer lasting than short-lived ones (see Auerbach et al. 2010). We argue that individuals understand that shocks from different sources and with different propagation mechanisms will have more or less persistent effects. This understanding is revealed through their survey responses: for example, where a positive shock is believed to be short-lived, individuals might report that one-period-ahead expectations of a variable will be high but that two-period-ahead expectation will revert to the level observed in the absence of the shock. This insight allows us to identify and investigate the effects of the shocks with different persistence properties which are meaningful to agents irrespective of the source of the shock or propagation mechanism involved.

The aim of the paper is to highlight the role of shocks with different persistence properties in business cycle dynamics. We establish these distinctions to be important empirically. For example, we find that agents understand that a large part of contemporaneous output movements are ultimately transitory. While these transitory shocks can affect output for two or three years, their role gradually diminishes and they have little effect after that time. Agents are also able to recognize the shocks that will have permanent effects and, in "normal times," stock prices react immediately to these shocks while the effects on output build gradually over three or four years, exactly as suggested in the traditional news-driven business cycle literature. But these same permanent shocks to output, and permanent shocks to interest rates, have more immediate effects on output during recessionary times because of the uncertainties they create. These effects are quite nuanced then, and the macroeconomic dynamics and persistence properties of models are quite different if the separately identified types of news and their uncertainties are omitted from the analysis. For example, permanent shocks are found to fully take effect on output within a year, as opposed to the three years in our more sophisticated analysis.

More specifically, in this paper, we use U.S. data on outputs, interest rates, and stock prices over 1981Q1–2019Q4 to consider the way in which news—in the form of shocks with different persistence properties—and its associated uncertainties influence macroeconomic dynamics. We build on the time-series models used to investigate the news-driven business cycle by exploiting the information contained in the Survey of Professional Forecasters (SPF) to distinguish between shocks that are believed to have permanent effects on output, interest rates and stock prices, those that have transitory but long-lived effects and those that have only very short-lived effects. We further extend the model to accommodate the uncertainties surrounding these different types of shocks. This is achieved through a multivariate GARCH-in-Mean model which allows the uncertainties surrounding the various shocks to be time-varying and to influence the determination of output, interest rates and stock prices. We also describe and make use of "persistence profiles" which provide model-based measures of the effects of news on the variables at different future horizons and which can be decomposed to highlight the contributions of the shocks with different persistence properties to the macro dynamics.

Our approach acknowledges that uncertainty is potentially influential for output fluctuations and makes modeling choices to address three issues that have emerged from the literature. The first relates to the measurement of uncertainty, which can be defined by the extent to which something is unknown. In their influential papers, Jurado et al. (2015) [JLN] and Ludvigson et al., (2021) [LMN] measure uncertainty by the size of forecast errors, while we focus on the directly-observed size of the expectational errors based on variable outcomes and what was reported 'known' in the surveys.[3,4] The second issue relates to the time frame over which the uncertainty is experienced. For example, the theory and simulations of Bloom (2009) show that a short-lived and transitory shock to uncertainty can *reduce* output as firms defer investment projects, while Bar-Ilan and Strange (1996) show that a more prolonged increase in uncertainty can trigger an *increase* in investment and output if investment decisions take time to implement. JLN and LMN consider this empirically by looking at uncertainty measured over different forecast horizons, and we address the issue by looking at the uncertainty associated with shocks with different persistence properties.[5] And third, we propose to capture uncertainty through a multivariate GARCH-in-Mean model, rather than assuming that uncertainty is itself randomly distributed and represented with a stochastic volatility model. Our GARCH specification focuses on the size of the past shocks to actual and expected outputs as the drivers of the volatility in expectational errors and is perhaps best described as allowing for an 'uncertainty channel' for these shocks to affect output rather than allowing for separate uncertainty shocks. It is, of course, possible that the second moment of the expectational errors are influenced by events not captured by the first moment shocks to actual or expected outputs, but there is an appealing simplicity and internal coherence in the GARCH assumption. This issue is considered again below when the features of the GARCH-in-Mean model are described in more detail.

The layout of the remainder of the paper is as follows. Section 2 provides a brief discussion on the importance of distinguishing permanent from transitory shocks and illustrates the point with a simple algebraic example. Section 3 describes the modeling framework adopted in our empirical work. This explains how we use the direct measures of expectations, introducing the model as though output is the only variables of interest and subsequently extending the model to include the other forward-looking variables. The section also comments on the identification of the shocks with their varying persistence properties, the persistence profiles that characterize their effects and the role of uncertainty in macro dynamics. Section 4 describes the empirical work, reporting the estimated model and providing persistence profiles to illustrate the role played by the different types of news and associated uncertainties in business cycle dynamics. Section 5 concludes.

2. Persistence properties of shocks and implications for the business cycle

The persistence of a shock is related to the concept of stationarity. If a variable is driven by innovations that are partly permanent and partly temporary, the variable can be modeled as a combination of unit root and stationary series. If the unit root series dominate, or the stationary series is close to unit root, the long-range outcomes of the variable vary in line with these innovations; that is, their effect persists. But if the stationary series dominate, the effects of innovations are short lived and do not persist. In the context of a univariate series x_t, Cochrane (1988) referred to the Beveridge-Nelson decomposition of a series into its permanent, stochastic trend component and its stationary component, and he proposed the innovation variance of the stochastic trend as a measure of the importance of the random walk component. This measure focuses on the size of the effect of today's shock on the long difference $x_{t+h} - x_t$ as $h \to \infty$. More generally, we might define the persistence of shocks to the variable by the size of the response of x_{t+h} to news that becomes available at time t; that is,

$$P_t^x(h) = \mathrm{Var}(\eta_{t+h}), \tag{1}$$

$$\text{where} \quad \eta_{t+h} = \mathrm{E}[\, x_{t+h} - x_{t-1} | \varepsilon_t, I_{t-1} \,] - \mathrm{E}[\, x_{t+h} - x_{t-1} | I_{t-1} \,], \tag{2}$$

E[.] is the expectations operator, ε_t represents the news arriving at t and I_{t-1} represents the information held at $t-1$. This corresponds to Cochrane's measure at the infinite horizon, tending to a constant if x_t contains a unit root and tending to zero if x_t is stationary. But even contemporaneously and at intermediate horizons, the variance $P_t^x(h)$ describes the range of potential outcomes that might be observed for x_{t+h} depending on today's news, reflecting the accumulating influence of the persistent component of the news and the declining influence of the transitory component over time. We can think of the value $P_t^x(h)$ as describing the *persistence profile* of the series at different horizons $h = 0, 1, 2, 3, \ldots$. The profile also conveys the information content and 'importance' of time-t news for x_{t+h} and is a useful descriptive tool for characterizing a news-driven business cycle.

The persistence of shocks is also related to the uncertainty surrounding a variable. The contemporaneous measure, $P_t^x(0)$ refers to the error variance of the time-$(t-1)$ forecast of x_t and so conveys the extent to which the x_t outcome is not known; that is, its time-t uncertainty.[6] At longer horizons, $P_t^x(h)$ describes the uncertainty surrounding x_{t+h} due to the time-t shock. Hence, for example, the measure falls to zero as $h \to \infty$ if the stationary component of the time-t shock dominates and the infinite-horizon effect of the shock becomes completely certain. Of course, the *total* uncertainty surrounding x_{t+h} at time-t for $h \geq 1$ depends on the expected size of shocks between $t+1$ and $t+h$ as well as the size of η_t and, as suggested by JLN/LMN, this is typically measured by the forecast error variance.

2.1. The dangers of ignoring the persistence properties of shocks; an illustration

It is worth considering the properties of the persistence profile in the univariate case when the variable is driven by shocks with different persistence properties. A simple illustration shows that failure to take account of the persistence of different types of shocks can introduce biases in modeling output movements, mis-measurement of uncertainty and inaccurate claims on the role of uncertainty in business cycle dynamics. To illustrate these points, denote output at t by y_t, denote the expectation of y_{t+h} published in a survey at time t by ${}_t y_{t+h}$ and consider the case where output growth is driven by two shocks

$$y_t - y_{t-1} = \rho(y_{t-1} - y_{t-2}) + \epsilon_t + \omega_t - \omega_{t-1}\,. \tag{3}$$

Here $\epsilon_t \sim N(0, \sigma_\epsilon^2)$ has a permanent effect on y_t and $\omega_t \sim N(0, \sigma_\omega^2)$ is a shock whose effect on output is known to be short-lived and offset next period, providing a simple example of "correlated news". The effects of both types of shock are propagated over time through the presence of the lagged growth term. Assuming expectations are formed rationally (purely for the purpose of exposition in this illustration), the survey measure of expected growth in $t+1$, as reported in t, is

$${}_t y_{t+1} - y_t = \rho(y_t - y_{t-1}) - \omega_t$$

and, if ρ is known, the survey measure can be used to obtain a direct measure of the transitory innovation experienced in t.[7] The range of outcomes surrounding output at different horizons in the future depends on the relative size of the ϵ_t and ω_t and is described by the persistence profile of (1) which is given here by

$$P_t^y(h) = \left(\sum\nolimits_{k=0}^{h} \rho^k\right)^2 \sigma_\epsilon^2 + \left(\rho^h\right)^2 \sigma_\omega^2, \quad h = 0, 1, 2, 3, \ldots ....$$

The persistence profile is dominated by the transitory shocks at the short horizons if σ_ϵ^2 is sufficiently smaller than σ_ω^2, but tends to $\frac{\sigma_\epsilon^2}{1-\rho^2}$ at longer horizons if there are permanent effects (i.e., $\sigma_\epsilon^2 \neq 0$) and tends to zero if output is subject only to transitory shocks.

In the absence of direct measures of expectations, output growth would be modeled using only actual output data, captured by the ARMA(1,1) process

$$y_t - y_{t-1} = \rho(y_{t-1} - y_{t-2}) + \upsilon_t + \theta\upsilon_{t-1} \tag{4}$$

with $\upsilon_t \sim N(0, \sigma_\upsilon^2)$. Matching the variance/covariances of the characterizations at (3) and (4), we have $\sigma_\epsilon^2 = (1+\theta)^2\sigma_\upsilon^2$and $\sigma_\omega^2 = -\theta\sigma_\upsilon^2$, with $\theta \in [-1, 0]$ depending on the relative size of σ_ϵ^2 and σ_ω^2. The associated persistence profile is

$$\widetilde{P}_t^y(h) = \left(\sum\nolimits_{k=0}^{h} \rho^k + \theta \sum\nolimits_{k=0}^{h-1} \rho^k\right)^2 \sigma_\upsilon^2$$

which is readily used to show that

$$\widetilde{P}_t^y(h) > P_t^y(h) \quad \text{for } h = 0, 1, 2, , \ldots$$

$$\text{with} \quad \widetilde{P}_t^y(h) \to P_t^y(h) \quad \text{as} \quad h \to \infty;$$

that is the size of the response of y_{t+h} to the news conveyed by actual output alone overstates the true persistence associated with the series at all but the infinite horizon. In the absence of survey data, the shocks with different persistence properties cannot be identified separately, the consequences of known-to-be-short-lived shocks cannot be taken into account and the persistence profile measures based on the univariate ARMA representation therefore overstate the size of the effect of shock at all horizons, including the contemporaneous uncertainty measure.

In practice, the size of the shocks impacting y_t might change over time so that $\epsilon_t \sim N(0, \sigma^2_{\epsilon t})$ and $\omega_t \sim N(0, \sigma^2_{\omega t})$ with the t-subscript on the variance terms denoting time variation in the size of the underlying innovations. The persistence profiles defined at (1)—updated to include the time variation in the variances—are still calculable if direct measures of expectations are used and still reflect the influence of shocks on output at each future time horizon. But the univariate ARMA(1,1) specification at (4) is now misspecified as the θ parameter, defined by the relative size of the permanent and transitory shocks, varies over time. The extent of the overstatement of the size of the persistent effect of shocks on output will also change over time therefore, and estimated errors from the ARMA(1,1) specification will be correlated with the $\sigma^2_{\epsilon t}$ and $\sigma^2_{\omega t}$. This means that, if these variance measures had been included as additional regressors in a time-invariant ARMA(1,1) specification, a spurious relationship would be found between growth and uncertainty. The illustrative exercise shows that the use of surveys in empirical work will not only allow us to distinguish the effects of shocks with different persistence properties, capturing explicitly the influence of correlated news, but will also deliver measures of the associated time-varying uncertainties the effects of which can be badly misinterpreted in the absence of survey data.

3. News, persistence, and uncertainty in a multivariate GARCH model

The basis of our empirical work is a multivariate GARCH-in-Mean model of actual and expected variables, but we first describe the simpler GARCH model and introduce the measures of persistence in this context. The GARCH model is able to expose the news arriving about current and future outputs and to accommodate time variation in the extent of the new information arriving at any one time. For example, assume a survey published in t reports expected output growth in $t+1$ and $t+2$. Furthermore, assume that actual output growth, $y_t - y_{t-1}$, as well as expectation errors, such as $y_t - {}_{t-1}y_t$, are stationary which imply that expected growth, such as ${}_ty_{t+2} - {}_ty_{t+1}$, is also stationary. Using a VAR(1) process for exposition purposes, a trivariate GARCH model explaining the actual and expected growth in output is given by

$$\begin{bmatrix} y_t - y_{t-1} \\ {}_ty_{t+1} - y_t \\ {}_ty_{t+2} - {}_ty_{t+1} \end{bmatrix} = \mathbf{A}_0 + \mathbf{A}_1 \begin{bmatrix} y_{t-1} - y_{t-2} \\ {}_{t-1}y_t - y_{t-1} \\ {}_{t-1}y_{t+1} - {}_{t-1}y_t \end{bmatrix} + \begin{bmatrix} \varepsilon_{0t} \\ \varepsilon_{1t} \\ \varepsilon_{2t} \end{bmatrix} \tag{5}$$

with

$$\varepsilon_t \sim N(0, \Sigma_t), \text{ and } \Sigma_t = \mathbf{H}_0'\mathbf{H}_0 + \mathbf{H}_1'\Sigma_{t-1}\mathbf{H}_1 + \mathbf{H}_2'\varepsilon_{t-1}\varepsilon_{t-1}'\mathbf{H}_2. \tag{6}$$

Here, $\varepsilon_t = (\varepsilon_{0t}, \varepsilon_{1t}, \varepsilon_{2t})'$ and these shocks represent all the new information arriving at time t on the actual and expected series; that is, what is not known about the series at $t-1$. The GARCH form of (6) allows the "size" of the new information on actual and expected outputs, Σ_t, to vary over time with the time-t (co-)variances of ε_t depending on the (co-)variances in time-$t-1$ and on the sizes of the shocks observed in the previous period. The Σ_t provides the basis of a straightforward measure of the uncertainty surrounding the series.[8]

It is worth emphasising that, in its most general form, the model in (5) simply assumes that expectation errors are stationary, which is consistent with any reasonable assumption on how expectations are formed. The model can accommodate Full Information Rational Expectations, therefore, which would imply restrictions on the parameters in $\mathbf{A}_1$ so that $y_t = {}_{t-1}y_t + \varepsilon_{1t}$ and ${}_ty_{t+1} = {}_{t-1}y_{t+1} + \varepsilon_{2t}$. Or the lag order can be extended so that it can accommodate the Sticky

Information Rational Expectations form described in Coibion and Gorodnichenko (2012), by restricting the parameters so that $y_t - {}_{t-1}y_t = \frac{\lambda}{1-\lambda}({}_{t-1}y_t - {}_{t-2}y_t) + \varepsilon_{1t}$, for some parameter λ. Or the model can be left unrestricted and, in this case, no structure is imposed on the way in which shocks - whether they have permanent or transitory effects on output—impact on the differences between actual and expected output other than to ensure these differences ultimately die away. In any case, and irrespective of how the information is used in forming expectations, the ε_{it}'s represent the news becoming available at time t and the modeling framework is robust to any sensible assumption on how expectations are formed.

A further point that is worth noting is that the interactions between the actual and expected growths in (5) would be reflected by complicated dynamics in the corresponding univariate model for actual output growth. For example, the trivariate VAR of order 1 in (5) has a corresponding ARMA(3,2) model for actual growth when considered alone. This is important as the ARMA specification for the shocks reflect the idea of "correlated news" in Walker and Leeper's (2011) analysis of the news-driven business cycle, with the effects of some types of news about output (e.g. news on future productivity) potentially appearing only after a lag or accumulating only slowly.

The model (5) can be written equivalently in the difference form

$$\Delta \mathbf{Y}_t = \mathbf{B}_0 + \mathbf{B}_1 \Delta \mathbf{Y}_{t-1} + \Pi\ \mathbf{Y}_{t-1} + \varepsilon_t \tag{7}$$

where $\mathbf{Y}_t = (y_t,\ {}_ty_{t+1},\ {}_ty_{t+2})'$ and the Vector Error Correction form makes explicit that the actual and expected output series are cointegrated with $\Pi = \alpha\beta'$ and cointegrating vector $\beta = \begin{pmatrix} 1 & -1 & 0 \\ 1 & 0 & -1 \end{pmatrix}$, reflecting the assumption that expectations errors are stationary. The model can also be written in its Moving Average form

$$\Delta \mathbf{Y}_t = \mathbf{C}(L)\varepsilon_t \tag{8}$$

where $\mathbf{C}(L) = \mathbf{C}_0 + \mathbf{C}_1 L + \mathbf{C}_2 L^2 + \ldots.$ is a matrix polynomial in the lag operator L and $\mathbf{C}_0 = \mathbf{I}$. The variance of the effect of time-t news on the long-differences $(\mathbf{Y}_{t+h} - \mathbf{Y}_{t-1})$ here is given by

$$\mathbf{P}_t(h) = \widetilde{\mathbf{C}}_h\ \Sigma_t\ \widetilde{\mathbf{C}}_h' \tag{9}$$

where $\widetilde{\mathbf{C}}_h = \sum_{k=0}^{h} \mathbf{C}_k$. This variance again gives "persistence profiles" that measure the variation in $\mathbf{Y}_{t+h}$ arising from time-t news, converging now to $\mathbf{P}_t(\infty) = \mathbf{C}(1)\ \Sigma_t\ \mathbf{C}(1)'$ as the time horizon expands.[9] If output is stationary in levels, then $\mathbf{P}_t(\infty) = 0$ reflecting the fact that the effects of shocks die away at the infinite horizon. If output is nonstationary, then some part of the shocks arriving at time t persists indefinitely and the $\mathbf{P}_t(\infty)$ reflects this, converging to a constant matrix with the extent of the persistent effect dependent on Σ_t, the size of the shocks at the time. The contemporaneous value $\mathbf{P}_t(0)$ provides measures of the uncertainty surrounding actual time-t output and expected future outputs at $t+1$ and $t+2$, reflecting the range of values of current output and of the variety of paths of output to $t+2$ that might be revealed by time-t news. The persistence profile relating to actual output, taken in isolation, is defined by $P_t^y(h) = \mathbf{e}'\mathbf{P}_t(h)\mathbf{e}$ where $\mathbf{e}' = (1, 0, 0)$ is the relevant selection vector. Similarly, the size of the potential responses of the one-step-ahead expectation ${}_{t+h}y_{t+h+1}$ to the news arriving at time t is traced out using the selection vector $\mathbf{e}' = (\mathbf{0}, \mathbf{1}, \mathbf{0})$.[10]

Actual and expected output being cointegrated is an important tool for our strategy to identify the permanent output shock. Given that the three series in $\mathbf{Y}_t$ are driven by a single stochastic trend, the three rows of the $\mathbf{C}(1)$ matrix will be the same, capturing the fact that the three series will converge to the same long-run outcome and the persistent effect of shocks to the three series is the same in the long run.

3.1. Identifying the permanent and transitory shocks

The expectations series provided by the surveys are necessary to find an accurate time series representation of the actual and expected output series. But, as noted in the illustration above, they can also be used to obtain observations on the shocks distinguished by the survey respondents according to their persistence properties. Specifically, in the example at (5), we can distinguish three types of shock according to their persistence: a very short-lived transitory shock whose effects influence output one period ahead but not two periods ahead, denoted, ω_{1t} say; a transitory shock which continues to affect output for two periods and beyond but whose effects die away at the infinite horizon, denoted ω_{2t}; and a permanent shock, ω_{pt}, representing the common stochastic trend driving the three series in the long run.[11] The structural shocks are related to the reduced form shocks as follows:

$$\begin{pmatrix} k_1 & k_2 & k_3 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} \varepsilon_{0t} \\ \varepsilon_{1t} \\ \varepsilon_{2t} \end{pmatrix} = \begin{pmatrix} 1 & 0 & 0 \\ * & 1 & 0 \\ * & * & 1 \end{pmatrix} \begin{pmatrix} \omega_{pt} \\ \omega_{2t} \\ \omega_{1t} \end{pmatrix} \tag{10}$$

$$\mathbf{Q}\boldsymbol{\varepsilon}_t = \mathbf{R}\boldsymbol{\omega}_t \tag{11}$$

$$\text{and} \qquad \boldsymbol{\omega}_t = \mathbf{W}\boldsymbol{\varepsilon}_t \qquad \text{where} \qquad \mathbf{W} = \mathbf{R}^{-1}\mathbf{Q}$$

and where $\boldsymbol{\omega}_t = (\omega_{pt}, \omega_{2t}, \omega_{1t})'$, $\boldsymbol{\omega}_t \sim N(0, \Omega_t)$ and $\Omega_t = \mathbf{W}\Sigma_t\mathbf{W}'$ is a diagonal matrix. The permanent shock is identified by the (known) combination of reduced form shocks, according to the coefficients k_1, k_2, and k_3, that defines the stochastic trend as in $\mathbf{C}(1)$; the shock with the transitory but long-lived effect is identified as being that part of news on two-period-ahead expectations not explained by the permanent shock (with the short-lived shock assumed to have no effect); and the short-lived shock is that part of news on one-period-ahead expectations not explained by the other two shocks.

As mentioned, the ε_{it}'s represent all the new information arriving at time t on the actual and expected series. We do not make any assumptions on the source of the news and, indeed, consider this to be a virtue of our approach. Appendix 1 shows how the complexity of the modeling framework rises as the number of sources of information increases so that identification of shocks based on arguments relating to the source of the shocks become difficult to sustain. Our modeling framework and our identification strategy is based on how survey respondents interpret the news they receive and is unrelated to the source.

An implication of this approach is that it does not distinguish between "genuine" shocks and the effects of misperceptions. As noted, the approach accounts for the effects of, for example, sticky information, through restrictions on the dynamics of the VAR, so misperceptions of this sort (whereby agents simply do not update their information sets) are accommodated within the model.[12] However, in this case, and more generally, the effects of a genuine short-lived weather shock, say, cannot be distinguished from the noise introduced into survey responses from an incorrect understanding or interpretation of events. This noise can be assumed to be transitory and to have stable structure (in the sense that some elements might usually last only one quarter while some might be more prolonged).[13] But the identified shocks from our approach will be a hybrid of genuine shocks, and the effects of these misperceptions and our approach will not separate the effects by source.

The translation of the reduced form shocks to the structural shocks described in (10) provides a more meaningful interpretation of the model in (7) and (6), but the complexity of the dynamics remains unchanged and, in particular, the implied MA specification for the shocks in the corresponding univariate model for actual output growth will continue to display the "correlated news" property emphasized in Walker and Leeper's discussion of the news-driven business cycle. The characterization of the time variation in the variance and the persistence profiles is also unchanged. To see this, note that $\Omega_\mathbf{t} = \mathbf{W}\Sigma_t\mathbf{W}'$ so that

$$\Omega_t = \widetilde{\mathbf{H}}_0'\widetilde{\mathbf{H}}_0 + \widetilde{\mathbf{H}}_1'\Omega_{t-1}\widetilde{\mathbf{H}}_1 + \widetilde{\mathbf{H}}_2'\boldsymbol{\omega}_{t-1}\boldsymbol{\omega}_{t-1}'\widetilde{\mathbf{H}}_2$$

where $\widetilde{\mathbf{H}}_i' = \mathbf{W}\mathbf{H}_i'\mathbf{W}^{-1}$ for $i = 0, 1, 2$ and the GARCH structure is maintained (noting that the $\widetilde{\mathbf{H}}$'s are not necessarily diagonal so each of the time-varying elements of $\boldsymbol{\Omega}_t$ will respond to all of the $\boldsymbol{\omega}_{t-1}$ shocks). Similarly, the moving average representation in (8) can be written in terms of the economically meaningful shocks

$$\Delta\mathbf{Y}_t = \mathbf{C}(L)\mathbf{W}^{-1}\boldsymbol{\omega}_t = \mathbf{D}(L)\boldsymbol{\omega}_t$$

where $\mathbf{D}(L) = \mathbf{C}(L)\mathbf{W}^{-1}$ and the persistence profiles are given by

$$\begin{aligned}\mathbf{P}_t(h) &= \widetilde{\mathbf{C}}_h\ \Sigma_t\ \widetilde{\mathbf{C}}_h' \\ &= \widetilde{\mathbf{D}}_h\mathbf{W}\ \Sigma_t\ \mathbf{W}'\widetilde{\mathbf{D}}_h' = \widetilde{\mathbf{D}}_h\ \boldsymbol{\Omega}_t\ \widetilde{\mathbf{D}}_h'.\end{aligned}$$

Written in this way, the persistence profiles now highlight the contribution of the shocks of different persistence to the profiles at different future horizons. Specifically, with $\boldsymbol{\Omega}_t$ being diagonal, we can decompose the persistence profile by writing

$$\mathbf{P}_t(h) = \widetilde{\mathbf{D}}_h\boldsymbol{\Omega}_t^p\widetilde{\mathbf{D}}_h' + \widetilde{\mathbf{D}}_h\boldsymbol{\Omega}_t^2\widetilde{\mathbf{D}}_h' + \widetilde{\mathbf{D}}_h\boldsymbol{\Omega}_t^1\widetilde{\mathbf{D}}_h' \tag{12}$$

where $\boldsymbol{\Omega}_t = \boldsymbol{\Omega}_t^p + \boldsymbol{\Omega}_t^2 + \boldsymbol{\Omega}_t^1$ separates the variances into the elements relating to the permanent shock ω_{pt} and the transitory shocks ω_{2t} and ω_{1t}. Here, the persistence profile relating to the permanent shock $\widetilde{\mathbf{D}}_h\ \Omega_t^p\ \widetilde{\mathbf{D}}_h'$ will converge to the single value of $\mathbf{P}_t(\infty)$ as the time horizon grows while the profiles relating to the transitory shocks $\widetilde{\mathbf{D}}_h\ \boldsymbol{\Omega}_t^2\ \widetilde{\mathbf{D}}_h'$ and $\widetilde{\mathbf{D}}_h\ \boldsymbol{\Omega}_t^1\ \widetilde{\mathbf{D}}_h'$ will tend to zero.

3.2. News, persistence and uncertainty in a multivariate GARCH-in-mean model

The model we consider in our empirical work has the following form:

$$\Delta\mathbf{Y}_t = \mathbf{B}_0 + \mathbf{B}_1\Delta\mathbf{Y}_{t-1} + \Pi\ \mathbf{Y}_{t-1} + \mathbf{G}\ \Theta_t + \boldsymbol{\varepsilon}_t \tag{13}$$

$$\text{where} \quad \boldsymbol{\varepsilon}_t \sim N(0, \Sigma_t), \qquad \boldsymbol{\varepsilon}_t^- = \min(\boldsymbol{\varepsilon}_t, 0),$$

$$\Sigma_t = \mathbf{H}_0'\mathbf{H}_0 + \mathbf{H}_1'\Sigma_{t-1}\mathbf{H}_1 + \mathbf{H}_2'\boldsymbol{\varepsilon}_{t-1}\boldsymbol{\varepsilon}_{t-1}'\mathbf{H}_2 + \mathbf{H}_3'\boldsymbol{\varepsilon}_{t-1}^-\boldsymbol{\varepsilon}_{t-1}^{-\prime}\mathbf{H}_3, \tag{14}$$

$$\text{and} \qquad \Theta_t = diag(\Sigma_t), \tag{15}$$

with the $\mathbf{B}$'s, Π, $\mathbf{G}$, and $\mathbf{H}$'s representing model parameters. This model extends the model described in the previous subsection in a number of ways explained below.

3.2.1. Allowing uncertainty to affect business cycle dynamics

The model at (13) has a GARCH-in-Mean form, with feedbacks from the contemporaneous uncertainty measures Σ_t influencing the growth in the macroeconomic variables in $\mathbf{Y}_t$ through the inclusion of the $\Theta_t = diag(\Sigma_t)$ in the mean equation. The GARCH-in-Mean form allows us to investigate the effect of uncertainty on the macro variables of interest in a single, coherent modeling framework. This avoids the problems highlighted by JLN/LMN of using uncertainty measures based on outside metrics which could reflect uncertainties about events that have no connection with the variables in the model. The GARCH-in-Mean model also avoids the problems highlighted by Carriero et al., (2018) of taking a two-step approach where measures of uncertainty are derived in stage 1 and then used in a macro model in stage 2. Internal inconsistencies can arise in this case as the measurement of uncertainty surrounding a variable in the first stage typically does not account for the fact that uncertainty affects the variable in the second stage. Further, in modeling the variable in the second stage, it is typically assumed that the variable has homoskedastic errors even though the first step is concerned with modeling the time-varying volatility of news about the variable. The disconnect between the measurement of uncertainty and the role played by uncertainty in macro dynamics is avoided in the GARCH-in-Mean model.

As mentioned earlier, one unique feature of the GARCH model relative to a stochastic volatility model is that the GARCH does not capture the effects of systematic influences or stochastic disturbances to uncertainty beyond those from the uncertainty channel from first moment shocks to the series in the model. This potential limitation is mitigated to some extent by the inclusion of the direct survey measures and other forward-looking influences on output (as discussed below) since the additional second moment influences captured by a stochastic volatility model would have to refer to influences that affect neither expected outputs nor actual outputs themselves. And the internal consistency and simplicity of focusing on the uncertainty channel is appealing. Nevertheless, the empirical results should be read bearing in mind that these potential additional influences are not directly captured in this analysis.

A further extension in (14) allows for an asymmetry in the effects of positive and negative shocks, with $\boldsymbol{\varepsilon}_t^-$ taking a nonzero value only when the shocks are negative. The inclusion of lagged $\boldsymbol{\varepsilon}_t^-$ in (14) in addition to the lagged $\boldsymbol{\varepsilon}_t$ means that past shocks can have different effects on the time profile of the size of shocks depending on whether they are expansionary or recessionary. Allowing for this asymmetry is potentially useful because the presence of the uncertainties in (14) introduces nonlinearities to the system. For example, a particularly large positive shock might increase output on impact, and this effect could be compounded subsequently if the growth in uncertainty generated by the large shock brings forward investment decisions that also increase output. On the other hand, a large (same sized) negative shock would reduce output on impact and this effect would then be offset by the increase in uncertainty this causes. It is useful to allow positive and negative shocks to have different effects in this context to better accommodate the various direct and indirect effects of the shocks that might exist.

The nonlinearities in the system also complicate estimation and inference in the empirical work. In particular, the moving average representation of (8) will have time-varying parameters with the time-specific $\mathbf{C}_t(L)$ defined by the history of shocks experienced up to that time and by the sign and magnitude of the current shock. The problems this causes for estimation can be resolved using simulation methods based around Koop et al.'s (1996) Generalized Impulse Response functions which can characterize the time series properties of a model taking into account the history, sign, and size of shocks. Details of the approach adopted are provided in the empirical section below and the Appendix.

3.2.2. The inclusion of other forward-looking influences on output

The model of (5) and (6) can be readily extended to include other variables and in the empirical work below we consider the vector of seven variables $\mathbf{Y}_t = (y_t,\ {}_ty_{t+1},\ {}_ty_{t+4},\ r_t,\ {}_tr_{t+1},\ {}_tr_{t+4},\ s_t)'$ adding actual and expected future values of the interest rate r_t and current stock prices s_t to the output variables.[14] The inclusion of the extra variables provides a broader view of macroeconomic dynamics, but it also improves the modeling approach even if the focus of interest is just on output movements. The interpretation of the shocks as "news," and of their size as uncertainty, relies on having a relatively complete characterization of the information set dated at time $t-1$. For the news on output, it can be argued that the inclusion of the lagged expectations series as a regressor means the analysis already uses the best summary measure of the available information since, if expectations are formed rationally and with full information, then the variable ${}_{t-1}y_t - y_{t-1}$ provides a complete description of the information relating to $y_t - y_{t-1}$ available at time $t-1$. But, in the possible absence of full information rationality, it is useful to broaden the information set. The inclusion of lagged values of financial variables such as stock prices and expected future interest rates aims to accommodate the relatively rapid reaction and forward-looking nature of financial markets hoping to capture any part of the news available at $t-1$ and relevant to output determination that is not adequately captured by the lagged actual and expected output series. Of course, the inclusion of the additional variables also provides an important connection to the empirical work in the news-driven business cycle literature.

With the inclusion of the extra variables, the model explains stock price growth and actual and expected change in interest rate, with s_t and r_t assumed integrated of order one and expectational errors on interest rates assumed stationary. Similar to output, the latter assumption implies that actual and expected interest rate are cointegrated and consequently share a single common stochastic trend. The stock price, however, is assumed not to be cointegrated with either output or interest rate. All in all, among the seven variables, there are four cointegrating equations (two between the output variables and two between the interest rate variables) and three stochastic trends consisting of the output trend, the interest rate trend, and the stock price trend.

The model also includes the contemporaneous uncertainties relating to all the additional variables as explanatory variables. The structure of the identification of (10) is then elaborated to give

$$\begin{pmatrix} k_{11} & k_{12} & k_{13} & k_{14} & k_{15} & k_{16} & k_{17} \\ k_{21} & k_{22} & k_{23} & k_{24} & k_{25} & k_{26} & k_{27} \\ k_{31} & k_{32} & k_{33} & k_{34} & k_{35} & k_{36} & k_{37} \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 \end{pmatrix} \begin{pmatrix} \varepsilon^y_{0t} \\ \varepsilon^y_{1t} \\ \varepsilon^y_{4t} \\ \varepsilon^r_{0t} \\ \varepsilon^r_{1t} \\ \varepsilon^r_{4t} \\ \varepsilon^s_{0t} \end{pmatrix} = \begin{pmatrix} 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ * & 1 & 0 & 0 & 0 & 0 & 0 \\ * & * & 1 & 0 & 0 & 0 & 0 \\ * & * & * & 1 & 0 & 0 & 0 \\ * & * & * & * & 1 & 0 & 0 \\ * & * & * & * & * & 1 & 0 \\ * & * & * & * & * & * & 1 \end{pmatrix} \begin{pmatrix} \omega^y_{pt} \\ \omega^r_{pt} \\ \omega^s_{pt} \\ \omega^y_{4t} \\ \omega^r_{4t} \\ \omega^y_{1t} \\ \omega^r_{1t} \end{pmatrix}. \tag{16}$$

The arrangement in (16) relates shocks with different persistence properties to the reduced form errors obtained from the equations explaining the actual and expected series with different forecast horizons. The emphasis of the identification remains on the persistence of the shocks with three of the seven structural shocks identified by long-run restrictions and the remainder identified by short-run restrictions. In terms of the long-run restrictions, the three stochastic trends, ω^y_{pt}, ω^r_{pt}, and ω^s_{pt}, are identified by their permanent effects—as described by the $\mathbf{C}_t(1)$—assuming that ω^y_{pt} affects all variables while the ω^r_{pt} are shocks that permanently affect interest rates beyond the effect of the ω^y_{pt}, and the ω^s_{pt} are shocks that permanently affect stock prices beyond the effects of the ω^y_{pt} and ω^r_{pt}. In what follows, we call these "permanent output," "permanent interest rate" and "permanent stock price" shocks. For the short-run restrictions, we assume there are two transitory shocks that last for four periods and beyond, ω^y_{4t} and ω^r_{4t}, and two transitory shocks that last for at least one-quarter but not four quarters, ω^y_{1t} and ω^r_{1t}; we call these "long-lived" output and interest rate shocks and "short-lived" output and interest rate shocks, respectively.

An important feature of this paper's approach to identification is the idea that agents can distinguish between news that will have a permanent effect on the macroeconomy from news that has only transitory effects. Examples of shocks that might have permanent effects might include technological advances, wars, demand shocks that are sufficiently severe that, via hysteresis, they have long-term consequences, and so on. Examples of shocks that could permanently affect stock prices and interest rates might include changes in policy regimes, changes in the perception of risk, a permanent shift in the composition of demand, and so on. Such news might be relatively rare and might be dominated in most periods by the news on transitory disequilibrium phenomenon. But importantly, it is the survey respondents who judge which news will have permanent effects and which news will have transitory effects and, as a modeller, we simply identify the different types of news from the reported expectations of outcomes at different forecast horizons.[15]

The inclusion of the extra variables allows for still greater complexity in the dynamics of the individual series, which will all be subject to correlated news shocks of the form highlighted by Walker and Leeper. These authors, following Beaudry and Portier [BP] (2014) and Barsky and Sims [BS] (2011) for example, note the potential role for correlated shocks in the context of a "speculation-driven" model where business cycles are generated primarily by firms anticipating future opportunities. Here stock prices react instantly to the news, but productivity and output react only with a lag. The inclusion of the extra variables in (13) renders the model more comparable to those in BP and BS, and the extended model is equally able to capture such dynamics, although the identification of shocks comes from insights on the timing of effects expressed by survey respondents rather than the sequencing of events imposed by BP and BS (e.g., does not rely on there being instantaneous effects on productivity as in BS). On the other hand, our identification does not consider the source of the shock of the mechanisms by which the effects are propagated over time so is equally consistent with any model involving correlated news, not just "speculation-driven" variants.

4. News, uncertainty and the US business cycle, 1980q1-2019q4

The empirical work of the paper considers quarterly measures of US real GDP, the 3-month Treasury Bill rate, and the S&P500 Stock Market index between 1980q1 2019q4. The forward-looking data are the experts' forecasts on output and interest rates provided in the Survey of Professional Forecasters (SPF) at the one-quarter-ahead and four-quarter-ahead forecast horizons; the measures used are the mean of the survey respondents' expectations as reported at the Philadelphia Fed's Real Time Data Centre website. Figure 1(a)-1(c) plot the data, showing in particular the relative stability in the expected output growth series. For example, Figure 1(a) shows that the standard deviation of actual output growth is 0.60, that of the one-period-ahead expected growth is 0.23, and the standard deviation of the (quarterly-measure of) the four-quarter-ahead expected growth is smaller again, at 0.15. The largest part of the variation in output is entirely unexpected then, but around 40% of the variation relates to the expected effects of previous shocks viewed from the previous quarter, and around 25% of the variation is explained by past shocks viewed from a year earlier.[16] It is this understanding of the persistent effects of shocks, as revealed by survey respondents, that we exploit in our empirical work.

The estimation of the model of (13) and (14) is not entirely straightforward because of the nonlinearities involved in the GARCH-in-Mean model. The nonlinearity means the effects of shocks will vary over time, depending on the start position (i.e., the history of shocks to that point) and the sign and size of the shocks at that point. The estimation of the model and calculation of the impulse responses involves two stages therefore. In the first stage, a GARCH-in-Mean model is estimated for the changes in the seven variables in $\mathbf{Y}_t = (y_t,\ {}_t y_{t+1},\ {}_t y_{t+4},\ r_t,\ {}_t r_{t+1},\ {}_t r_{t+4},\ s_t)'$ allowing for feedback to the mean equations from the time-varying variance of the *reduced form* residuals Σ_t. Estimation is by maximum likelihood and subject to diagonal BEKK restrictions, which ensure the positive definiteness of the variance–covariance matrices as they evolve over time (see Engle and Kroner (1995). The second stage estimates the linear combinations of the reduced form residuals that have a permanent effect on output and stock prices at each point in time taking into account the nonlinearities introduced through the uncertainty terms. This is achieved through simulation, generating (at each point in the sample) impulse responses for simulated shock realisations based on the estimated model from the first stage and averaging across the simulations. This delivers a time-specific moving average representation corresponding to (8) but with dynamics characterized by $\mathbf{C}_t(L)$ taking into account the history of shocks to that point (see

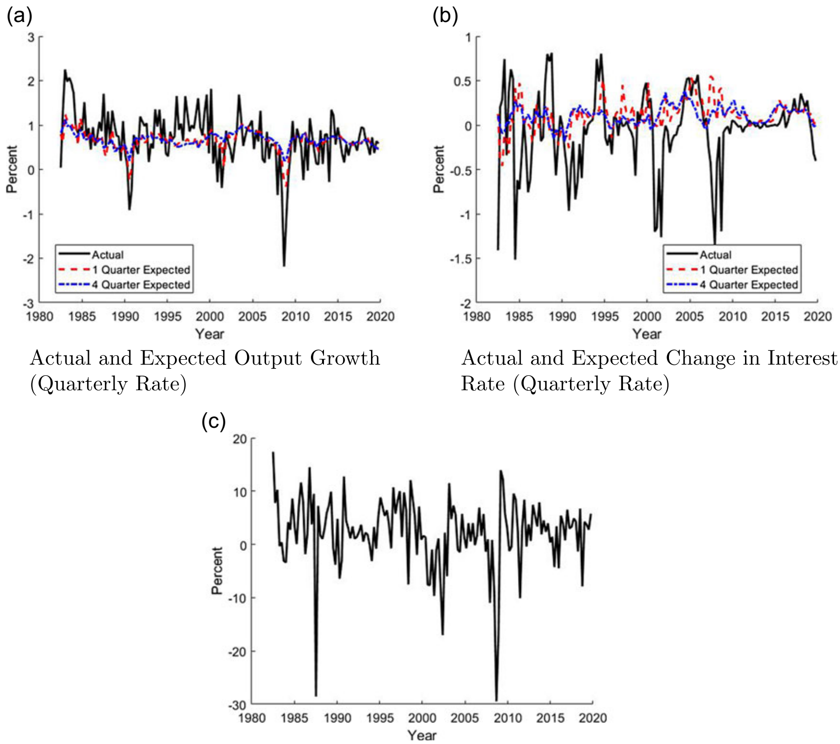

(a) Actual and Expected Output Growth (Quarterly Rate)

(b) Actual and Expected Change in Interest Rate (Quarterly Rate)

(c) Stock Return

Figure 1. Data.

the Appendix for details). Having estimated the model parameters, including an estimate of $\mathbf{C}_t(1)$ and the corresponding long-run relations, the identifying structure of (16) can be employed to investigate the effects of the structural shocks at each point in the sample.

The multivariate GARCH-in-Mean model of (13) and (14), extended to include 2 lags of the endogenous variables, was estimated following this strategy. In what follows, we refer to this baseline model as G7U (indicating the 7 variables and uncertainty feedback) and investigate the model's dynamic properties and, in particular, the dynamic effects of permanent and transitory shocks. We then explore the consequences of neglecting survey measures of expectations as well as the feedback of uncertainty in the model dynamics by estimating two alternative models—one without expectations, and one without the uncertainty feedback. For ease of exposition, the focus of the commentary below is primarily on the output effects of shocks.[17]

4.1. The dynamic effects of permanent and transitory shocks

Figure 2 illustrates the dynamics of the estimated model by showing the average response of actual and the expected output series to two system-wide shocks: one that results in a one-standard deviation *increase* in the actual output series and one that results in a respective one standard deviation *decrease*.[18] Focusing on actual output first, the figure shows that the dynamics in response to both positive and negative shocks are complex and prolonged: for the positive shock, the initial rise of 0.5 percentage points results in an overshoot at two years converging after around four years to a level in which the effect of the initial shock is doubled, with output around one percentage point higher than in the absence of the shock; for the negative shock, the convergence is achieved after six years and the long run level, at −0.5 percentage points, is roughly equal to the initial shock.

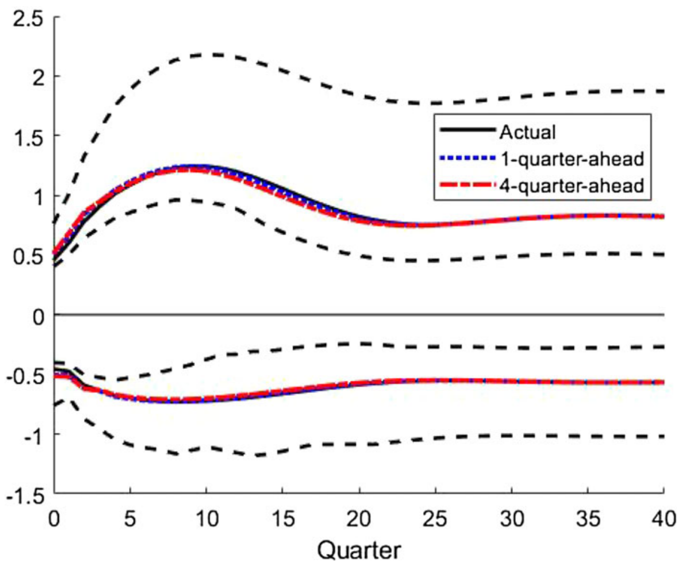

Figure 2. Response of actual and expected output growth to a positive or negative output shock.
The shock is a one standard deviation positive or negative output shock. Top Panel: Positive output shock. Bottom panel: Negative output shock. Solid lines correspond to the average responses across history. Dashed lines correspond to the 95% confidence intervals of the responses of actual output.

The figure also shows that the actual and expected output series converge to the same level eventually (by construction). However, the interactions are not straightforward and the three series are certainly not simple horizontal displacements of each other (with ${}_t y_{t+1} = y_{t+1}$ and ${}_t y_{t+4} = y_{t+4}$ for all forecast horizons t) as they would if full information rational expectations (FIRE) holds. This shows that information rigidities and/or some form of sentiment play an important role in the macro dynamics and these features are best captured through the inclusion of the survey data.[19]

It is worth exploring the nature of the prolonged dynamic response,[20] and Figures 3(y) and 4(y) provide some important insights with regard to the time-varying impact of different shocks by their relative persistence. In more detail, Figure 3(y) plots the (square root of the) persistence profile measures for actual output at one quarter ahead, at one year ahead, and at three years ahead over the sample period; that is, showing $P_t^y(h)$ for $h = 0, 3, 11$.[21] The figures illustrate the varying importance of the different types of shock at different times. Figure 3(y) shows that, at all times over the sample, the persistent effect of shocks to output grows as the future horizon increases—from the black $P_t^y(0)$ line to the blue $P_t^y(3)$ line and to the red $P_t^y(11)$ line—with considerable change occurring over one year and continuing to show for the subsequent two years. The changing balance in the relative importance of the permanent, long-lived, and short-lived shocks to the persistence profiles at different future horizons is illustrated in Figure 4(y).[22] Averaging across the whole sample, Figure 4(y) shows that, for output, the long-lived transitory shocks explain 60% of the variance on impact and continue to have a noticeable impact over the subsequent two to three years, after which the dominance of the permanent shocks is established. Short-lived shocks play a very small role at any time horizon. The prolonged impact of the identified known-to-be-transitory shocks is a key finding of the paper.

The difference between the response of output and the response of stock prices to shocks is shown by comparison with the plots in Figures 3(s) and 4(s). Here, for stock prices, more than 80% of the variance on impact is explained by the permanent shocks, and these permanent shocks explain nearly all the variance within a year. This corresponds well with the news-driven

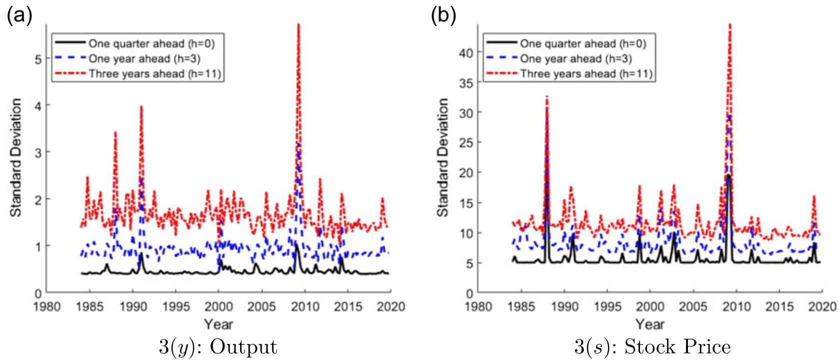

Figure 3. Output and stock price persistence profiles.
Persistence profiles from the baseline model that features both survey data and uncertainty feedback.

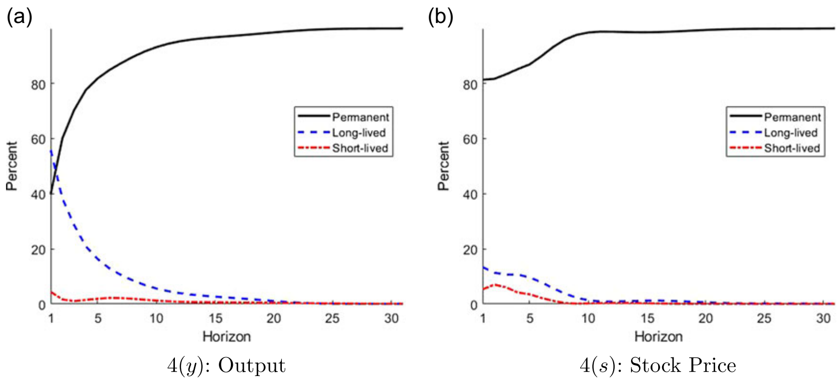

Figure 4. Average decomposition of output and stock price persistence profiles over different forecast horizons.
The contribution of each type of shock to output and stock price persistence profiles is averaged across the sample periods. The persistence profiles are constructed from the baseline model that features both survey data and uncertainty feedback.

business cycle literature where news about future, permanent output outcomes is reflected very quickly in stock prices but only gradually translates into actual output growth. The additional and novel contribution here, compared to that delivered elsewhere in the literature, is that the permanent shocks and their effects are identified explicitly and separately from the effects of transitory shocks—where the identification structure is an intuitive outcome from the survey responses of expectations.

An important part of the time variation arises through the asymmetric effect of positive and negative shocks. Figure 5 presents the equivalent figures to Figure 4(y) relating to output growth but showing the decompositions when averaged over expansions/normal times, that is, periods of positive output growth and when averaged over recessions, that is, periods of negative output growth. Output growth was positive over most of the sample, so the decomposition during expansions is similar to the averaged across the whole sample. But Figure 5 shows that the permanent shocks have a much larger impact effect on output during recessionary times than expansionary times (with 65% of the persistence profile on impact relating to permanent shocks compared to 40% in expansionary times) and the long-run effects (or at least 95% of them) are achieved much more quickly, taking around two rather than four years.[23] As we will see below, usually, the largest part of the contemporaneous uncertainty surrounding output and interest rates arises out of transitory shocks. But it is recognized that their effects will dissipate over the coming two to three

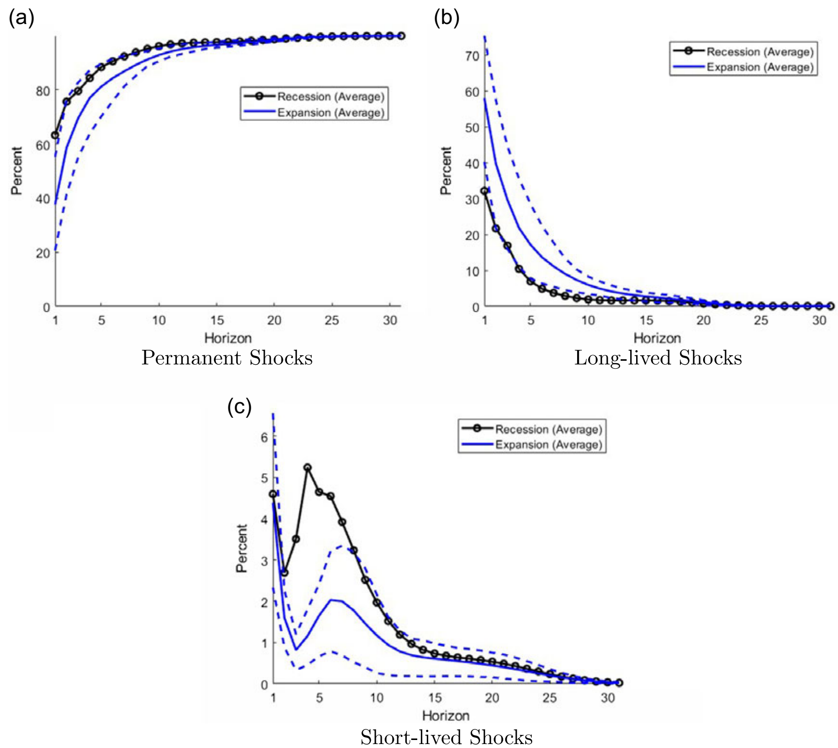

Figure 5. Contribution of different shocks to output persistence profiles in recessions and expansions over different forecast horizons.
The contribution of each type of shock to output persistence profiles in recessions vs the historical distributions of their respective contributions in expansions. The dotted lines provide the 68% historical bands over expansionary periods. The persistence profiles are constructed from the baseline model that features both survey data and uncertainty feedback.

years and the uncertainties arising from this source do not impact significantly on output growth. In contrast, the uncertainties arising surrounding permanent shocks to output and interest rates are influential, and *during recessions*, a larger part of the contemporaneous uncertainty surrounding output relates to permanent shocks; uncertainty therefore plays a particularly important role in determining output levels in recessionary times. Again, as illustrated in Figure 4, the scaling in Figure 5 reveal that short-lived shocks play a very small role.

In general, the importance of uncertainties surrounding permanent shocks to output and interest rates can also be seen from Table 1. The Table reports the contemporaneous impact of a one standard deviation increase in the uncertainties arising from the various shocks with different persistence properties on the actual and expected outputs.[24,25] The coefficients show that changes in the uncertainty arising from the permanent shocks to output and interest rates have the most influence on output; a one standard deviation increase in the uncertainty arising from permanent output shocks increases actual output by 0.17 percentage points, while a one standard deviation increase in the uncertainty arising from permanent interest rate shocks decreases actual output by 0.28 percentage points. The finding that it is uncertainty surrounding permanent shocks that impacts on the macroeconomy, and that changes in the uncertainties arising from transitory shocks have little or no impact on output, is another key result from this empirical work.[26]

Figure 6 provides a more detailed description of the output persistence profiles at $h = 0$ at different points in the sample and decomposing the profiles to see the contribution of the different types of shock. As seen earlier from Figure 4, Figure 6 shows that, for output, usually,

Table 1. Impact of uncertainty: structural shocks on actual and expected output

	θ_p^y	θ_p^r	θ_p^{sp}	θ_4^y	θ_4^r	θ_1^y	θ_1^r
Δy_t	0.165	−0.277	0.010	−0.062	−0.067	0.005	0.030
$\Delta_t y_{t+1}^e$	0.182	−0.276	0.010	−0.065	−0.077	0.006	0.050
$\Delta_t y_{t+4}^e$	0.292	−0.281	0.024	−0.069	−0.069	0.010	0.047

Impact of a one standard deviation increase in uncertainty on actual and expected output growth. Computation steps are provided in the Appendix.

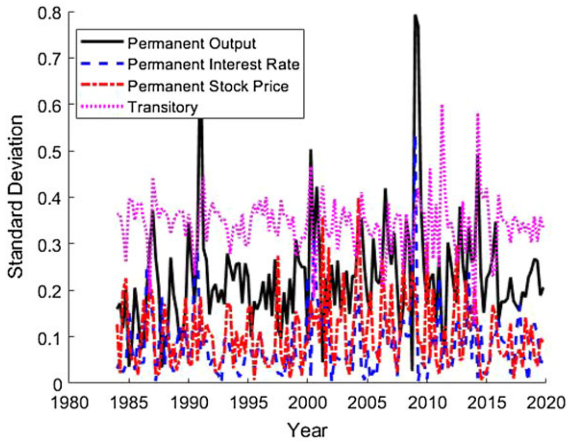

Figure 6. Decomposition of output contemporaneous persistence profile.
Decomposition of 1-quarter-ahead output persistence profile at each sample period. Transitory shocks are defined as shocks that do not have a permanent effect on all three of output, interest rate, and stock price trends.

the largest part of the contemporaneous uncertainty surrounding output arises out of transitory shocks which we know have little impact in the long run. But there are clear episodic peaks in uncertainty arising from the permanent output shocks at the start of the 1990 Gulf War recession and at the 2008 Global Financial Crisis (GFC). Relatively large values for the uncertainty arising from permanent interest rate shocks are found at these times and the two types of uncertainty also show at the time of the 1987 Stock Market crash and around the early 2000's dot-com recession. These dates are perhaps obvious candidates for special mention, but it is worth noting—like JLN/LMN before us—that the number of episodes in which these measures of uncertainty rise above 1.65 standard deviations from the mean is relatively small when compared to the very volatile VIX measures for example. As it turns out, and looking back to the profiles at $h = 11$ in Figure 3(y), it is the shocks—and associated uncertainties—of the 1987 Stock Market crash, 1990 Gulf War recession and the 2008 Global Financial Crisis (GFC) which translated into particularly extraordinary permanent output effects.

4.2. The contribution of survey expectations and the feedback of uncertainty

The central role played by the survey data and uncertainty feedback in generating the model's dynamic properties described above is demonstrated by comparing the output profiles of our model with the corresponding profiles from two alternative models: (i) "G3U" which denotes a trivariate GARCH-in-mean model employing only the actual output, actual interest rate, and stock price data but allowing for an uncertainty feedback; and (ii) "G7" which denotes the same 7-variable model as our baseline GARCH-in-Mean (G7U) model with survey data on expectations

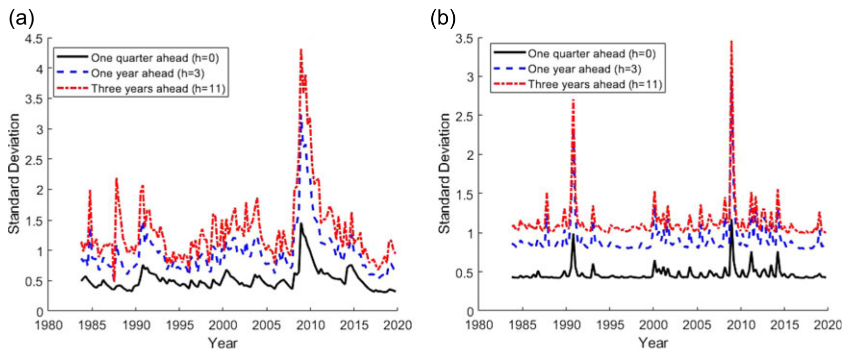

7(G3U): No Survey Data 7(G7): No Uncertainty Feedback

Figure 7. Output persistence profiles from models without survey data or uncertainty feedback.
The model without survey data is a tri-variate model of only actual variables but features the GARCH-in-mean component. The model without uncertainty feedback uses both actual and expected variables but does not feature the GARCH-in-mean component.

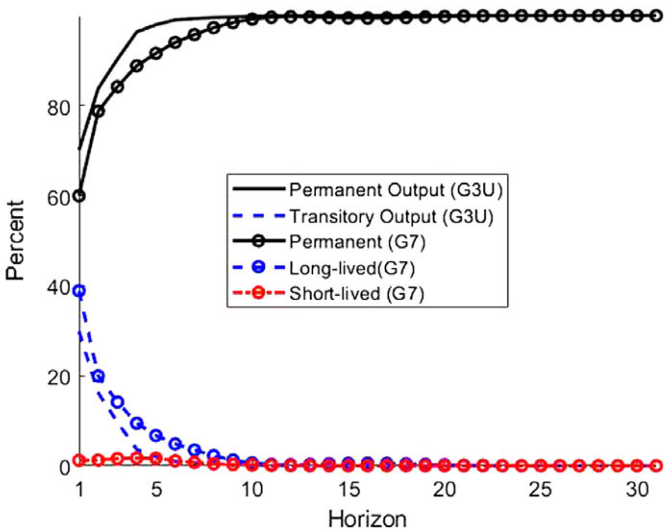

Figure 8. Average decomposition of output persistence profiles from models without survey data or uncertainty feedback.
G3U: Tri-variate model without survey data that uses only actual variables but features the GARCH-in-mean component. G7: Model without uncertainty feedback that uses both actual and expected variables but does not feature the GARCH-in-mean component.

but now without the uncertainty feedback.[27] The profiles from the alternative models are given in the sets of plots in Figures 7(G3U), 7(G7) and 8.

Focusing on the role of the survey data and the importance of explicitly identifying permanent and known-to-be-transitory shocks, in contrast to the baseline model in Figure 3(y), the output persistence profiles in Figure 7(G3U) show that there is very little gap between $P_t^y(3)$ and $P_t^y(11)$. This indicates that a large part of the response to a permanent shock takes place within the year therefore failing to capture the prolonged dynamic response over three years and beyond exhibited in the baseline model.[28] Second, in the decomposition of the persistence profiles for a given time horizon, a comparison of Figure 4(y) to Figure 8, shows that there is also a marked overstatement in the importance of the permanent output shock relative to the baseline model. It is worth noting that the extra dynamic sophistication is achieved in our baseline model not simply through the inclusion of additional variables but through the inclusion of variables capable of distinguishing and capturing the effect of known-to-be-transitory shocks.

Turning to the role of uncertainty, the plots relating to the G7 model differ in important ways from those of the baseline G7U model. First, similar to the G3U model, the G7 model also overstates the importance of permanent shocks in the short-run. In contrast to the G3U model, however, the G7 model features a similar prolonged response to shocks to our baseline model, with a similar gap between $P_t^y(0)$, $P_t^y(3)$, and $P_t^y(11)$ in Figure 7(G7) and with the long-lived transitory shocks showing for around two years in Figure 8. However, the presence of significant feedbacks from uncertainty introduces nonlinearities into the model and time variation in the effects of shocks. This time variation is particularly reflected in Figure 4(y) where there are substantial differences between the profiles at the three horizons at different points in the sample; the correlations between the output profiles $P_t^y(h)$ at $h = 0$ and the profiles at $h = 3$ and $h = 11$, calculated across the sample, are 0.71 and 0.57, respectively. Of course, some part of the time variation in the profiles is due to the changing size of the shocks captured by the GARCH effects. But the corresponding correlations in the persistence profiles for the G7 model are 0.95 and 0.94 showing that in the absence of the uncertainty feedback, the effects of the shocks simply accumulate over time and in a similar way at all points in the sample. Further, the correlation between the output profiles at $h = 0, 3$ and 11 for the G7U GARCH-in-Mean model and the G7 GARCH model—as illustrated in Figures 3(y) and 7(G7)—are 0.916 at $h = 0$ but just 0.662 and 0.578 at $h = 3$ and 11, respectively, showing the important role played by the uncertainty feedbacks in the preferred G7U model.[29]

The discussions above have elaborated how our proposed G7U model is able to capture important features of the data that are omitted in the absence of expectations data and/or its treatment of uncertainty. While the model is complex, it is useful to summarize the key findings as follows:

- The inclusion of actual and expected measures of the series allows prolonged and complex dynamics to be captured. The ability to distinguish the effects of shocks with different persistence properties allows the model to capture (i) the slow accumulation of permanent effects and (ii) the prolonged effects of some transitory shocks which are obscured in the absence of survey data. The complexity—including hump-shaped output response to shocks—is partially down to interaction of actual and expected series which is more sophisticated than would be suggested in simple FIRE models (Figure 2, Figures 3(y) and 4(y) vs. Figure 7(G3U) and Figure 8).
- The results correspond well with the news-driven business cycle literature where news about future, permanent output outcomes is reflected very quickly in stock prices but only gradually translates into actual output growth (Figures 3(y), 3(s), 4(y) and 4(s)).
- In the absence of explicitly identified permanent and known-to-be-transitory shocks, the impact of permanent shocks on output is fully realized within a year relative to more than three years and is markedly overstated on impact (Figures 3(y) and 4(y) vs. Figures 7(G3U) and 8).
- The uncertainty feedbacks introduce considerable time-variation in the output response to shocks; the uncertainty surrounding permanent shocks plays a particularly important role in determining output levels in recessionary times (Figure 5, Figure 7(G7)).
- Uncertainty plays an important role in macrodynamics episodically. The uncertainty surrounding permanent output shocks and permanent interest rates shocks was at its greatest at the time of the 1987 Stock Market crash, at the start of the 1990 Gulf War recession, around the early 2000's dot-com recession and at the 2008 Global Financial Crisis (GFC). The shocks—and associated uncertainties—of the 1987 Stock Market crash, 1990

Gulf War recession and the 2008 Global Financial Crisis (GFC) translated into particularly extraordinary permanent output effects (Figure 6).

- Uncertainty arising from the permanent shocks to output and interest rates have significant influence on output, positively in the case of permanent output uncertainty and negatively for permanent interest rate uncertainty. Changes in the uncertainties arising from transitory shocks have little or no impact on output (Table 1).

These baseline results are obtained using 2020:Q2 vintage data. However, real-time data arguably more closely reflects the information held by the forecasters when they submit their responses. While stock price and interest rate are already real-time in nature, real GDP data are often revised. To check the robustness of our main results to the use of real-time real GDP data, we estimate our baseline multivariate GARCH-in-mean model using the first-release real GDP growth series as our measure of actual GDP. The expected GDP data are constructed accordingly from the first-release actual real GDP series.

Overall, qualitatively, our main conclusion does not change—permanent shocks have a higher contribution to short-term output variations in recessions (around 40% for 1-quarter-ahead output uncertainty) than in expansions/normal times (around 20% for 1-quarter-ahead output uncertainty). Relative to the distribution of the contribution of permanent shocks to output in expansions, this difference is statistically significant for 1- to 10-quarter-ahead output uncertainty, which is a longer period than that found with revised GDP data. The mean estimates are indeed lower compared to the ones obtained with revised GDP data and the reduction can be attributed to the higher contribution of long-lived transitory shocks. The contribution of short-lived shocks on output is still small, with only a maximum of around 6% for 1-quarter-ahead output uncertainty on average. In comparison, and consistent with the results using final vintage data, stock price dynamics are primarily affected by permanent shocks in all horizons and in both expansions and recessions. Replications of Figures 4 and 5 using real time data are provided in the Appendix illustrating the robustness of the results in the paper to the use of real time data.

5. Concluding remarks

Survey data provide the means to identify shocks with different persistence properties and the influence of "correlated news" on macro dynamics. Their use in the GARCH-in-Mean model also delivers measures of uncertainty that can be used to investigate its separate role in business cycle fluctuations. The omission of survey data from time series analysis renders the work susceptible to the econometric problems relating to "foresight" but also makes it vulnerable to misspecification biases when there is time variation in the uncertainties surrounding the news arriving on current and future outcomes. The empirical work of this paper establishes the importance of using the information contained in survey data to distinguish between shocks with different persistence properties and to capture explicitly the sophisticated dynamic response of output to these shocks and their associated uncertainties. The empirical results emphasize the complexity and protracted nature of the response of output to permanent and long-lived transitory shocks, the separate contribution to macro dynamics of the process of expectation formation, the important but episodic role of uncertainty in macro dynamics, and the observation that the different types of shock play different roles in different circumstances.

We find that much of the contemporaneous fluctuation in output is recognized as transitory, and that this element of fluctuation is discounted gradually and has no persistent effect on output either directly or through the associated uncertainties. It is reassuring to know that decision-makers are aware of this short-lived element and "look through" its effects. But the earlier algebraic exercise highlighted the dangers of ignoring the element for modelers who will overstate the extent of uncertainty, and interpret it as being more influential than it actually is, if the noise is not properly accounted for.

The complementary finding is that news on permanent output change is the primary driver of macro dynamics. As emphasized in the news-driven business cycle literature, news about permanent future output change is reflected straightaway in stock price reactions, but it is understood that it will take more than three years before this is translated fully into actual long-run output outcomes. An important part of the propagation of the effect over time, at least during downturns, is the uncertainty associated with the news on permanent output and permanent interest rate change. Although these two uncertainty channels have offsetting effects, the overall effect of the uncertainty arising from the two sources on output was particularly negative in the recessions associated with the 1987 Stock Market Crash, the 1990 "Gulf war" and the 2008 GFC.

We have emphasized throughout the persistence properties of our shocks— identified from the survey responses at different forecast horizons—without reference to the source of shocks or their propagation mechanisms. We consider this to be an important feature of the approach given the frailty of some assumptions used to identify "economically meaningful" shocks when only administrative data is available. We hope this work will encourage greater use of survey data and direct measures of expectations in models used to investigate the news-driven business cycle to better accommodate the influence of information flows.

Acknowledgments. Financial support from the Australian Research Council (Project DP230100959) is gratefully acknowledged. We are also grateful to two referees and participants at the 2022 conference of the International Association of Applied Econometrics and at the 2022 Federal Reserve Bank of Cleveland conference on Survey Data Analysis for helpful comments.

Notes

1 Beaudry and Portier (2014) provide an excellent overview of the news-driven business cycle literature, and Bloom (2014) and Castelnuovo (2019) provide similarly excellent overviews of the uncertainty literature.

2 The theoretical case has been better explored. For example, Walker and Leeper's (2011) structural model shows that the often observed hump-shaped dynamic responses of macro variables to shocks is consistent with two contradictory sets of circumstances: where there are simple information flows and propagation mechanisms involving real rigidities (such as habit formation, variable capacity utilization, or costly investment adjustment); or, equally, where there is correlated news without real rigidities.

3 This contrasts with uncertainty measures based on outside metrics such as stock price volatility (as in Bloom (2009) and Barrero et al. (2017)), newspaper coverage (as in Baker et al. 2016), or the dispersion of firm-level experiences (as in Bloom et al. (2018), Bachmann et al. (2013), and Gilchrist et al. (2014)).

4 Lahiri and Liu (2006), Lahiri and Sheng (2008, 2010), Bachmann et al. (2013), Boero et al. (2015), Clements (2017), Jo and Sekkel (2019), and Binder et al. (2022), inter alia, provide further illustrations of the advantages of using surveys of expectations in deriving uncertainty measures.

5 Discussions of the role of uncertainty in business cycles focusing on the persistence of shocks are also provided in Leduc and Liu (2016) and Basu and Bundick (2017).

6 The variance is not only an intuitively reasonable measure of uncertainty but also (a scaled version of) the "differential entropy" measure of uncertainty— conveying the average surprise of a random variable—when the stochastic process underlying the variable is normally distributed.

7 Clearly the permanent ϵ_t are also retrievable - from $(1-\rho)(y_t - y_{t-1}) - ({}_t y_{t+1} - y_t) + ({}_{t-1}y_t - y_{t-1})$—and so both σ_ϵ^2 and σ_ω^2 are observable.

8 The model could be readily extended to accommodate measures of output obtained in real time, including modeling jointly the growth in first-release output measures and revisions as in Garratt et al., (2003). We use final vintage output for simplicity here as our previous work—for example, Aristidou et al. (2022) - has found that, while real-time data are important in nowcasting and forecasting, it is less influential in in-sample modeling.

9 The term "persistence profile" was introduced in the multivariate context in Lee and Pesaran (1993), drawing on the work of Lee, Pesaran and Pierse (1992, 1993), although these papers assumed homoskedasticity and were time invariant.

10 If we are concerned with the uncertainty surrounding output over the near future, we might calculate $\mathbf{e}'\mathbf{P}_t(0)\mathbf{e}$ with $\mathbf{e}' = (\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$, for example, measuring the uncertainty surrounding *average* output today and expected over the coming two quarters. The uncertainty measure still relates to the extent to which something specific is not known and it accommodates the effects of comovements between the variables in defining the uncertainty. The measure is again in line with the differential entropy concept of uncertainty surrounding outcomes involving multiple variables, which is based on the (log of the determinant of the) variance–covariance matrix of the variables when they are jointly normally distributed.

11 Clearly, the definition of the different types of shock are determined by the time-frame of the survey questions. In practice, these are usually chosen to capture economically meaningful horizons (e.g., one-quarter-ahead to consider short-term noise and one-year-ahead to consider business-cycle fluctuations).
12 Here the news that arrives is observed by those who do update their information sets to form RE and the identification scheme can be applied to the reduced form residuals as usual.
13 Binder et al., (2022) provide a useful theoretical model highlighting the effects of noise on uncertainty measures at different forecast horizons and relating this to stylized facts on density forecasts reported in surveys.
14 Note that this now uses the survey expectations of output and interest rates at one- and four-quarters-ahead forecast horizons (rather than one- and two-as in the illustration above), providing a more realistic time-frame for the shocks with different persistence properties.
15 See Garratt et al., (ch. 3, 2006) for a discussion of the difficulties in identifying economically meaningful shocks based on the short-run dynamics suggested by a structural DSGE model, or on specific adjustment costs, or on the timing and/or sequencing of decisions.
16 These are broad-brush statements: the simple decomposition of the variance of actual growth into the sum of the variances of expected and unexpected growths holds exactly only in the case of full information rational expectations where the expected and unexpected elements are orthogonal.
17 Figures depicting the stock price and interest rate effects are in the Appendix.
18 These are "system-wide" shocks because we allow for innovations in all variables alongside the output change as would be expected according to the estimated Σ_t. The "average" here is the average across the impulse response obtained at each point in the sample to see the effects of the shocks abstracting from the influence of history.
19 See Garratt et al. (2018) for further discussion on the role of sentiment or information rigidities in macro-dynamics.
20 The asymmetries in the effects of the positive and negative shocks are an important feature which we develop below.
21 Specifically, we report the relative sizes of $\mathbf{e}'\mathbf{P}_t(h)\,\mathbf{e}$ for $h = 0, 3, 11$ where $\mathbf{e} = (1, 0, 0, 0, 0, 0, 0)'$ is a selection vector picking out the first element of $\mathbf{P}_t(h)$ relating to actual output growth.
22 For example, using the notation of (12), the relative contribution of the permanent shocks to output for horizon h at time t is $\frac{\mathbf{e}'\widetilde{\mathbf{D}}_h\Omega_h^p\widetilde{\mathbf{D}}_h'\mathbf{e}}{\mathbf{e}'\mathbf{P}_t(h)\,\mathbf{e}}$, where $\mathbf{e} = (1, 0, 0, 0, 0, 0, 0)'$ is a selection vector picking out the first element of the matrices relating to actual output growth. The figures report the average values over t for each h.
23 The corresponding profiles for stock prices are similar in expansionary and recessionary times.
24 These statistics are also based on the "average" relationship between structural and reduced form errors obtained at each point in the sample; see the Appendix for details.
25 It is worth emphasizing that there are no "uncertainty shocks" here. Rather there are just shocks to the macroeconomic variables which have an effect through the uncertainty channel as well as the direct effects captured in standard linear models.
26 The results in Table 1 are based on the ordering of output being placed first and therefore allowing permanent output shocks to contemporaneously affect all variables. These results are robust to the subsequent ordering of the financial variables with respect to each other and assume that interest rates and stock prices move far quicker than output and therefore only affect output with a lag.
27 In the latter case, the trivariate VAR was estimated to include four lags to fully capture the dynamics in the three series.
28 Model G3U is driven only by (three) permanent shocks—so that 100% of the variance is explained by these even on impact—so Figure 8 is drawn to decompose the profiles between the effects of the permanent output shock (which translates into permanent output change) and the permanent stock price and interest rate shocks (which do not ultimately drive output in the long run).
29 The corresponding correlations for stock prices are 0.938, 0.639, and 0.572, while they are 0.938, 0.614, and 0.565 for interest rates showing the impact of uncertainty on the dynamics of all the variables in the model.

References

Aristidou, C., K. C. Lee and K. Shields. (2022) Fundamentals, regimes and exchange rate forecasts: Insights from a meta exchange rate model. *Journal of International Money and Finance* 123, 102601.

Auerbach, A. J., W. G. Gale and B. H. Harris. (2010) Activist fiscal policy. *Journal of Economic Perspectives* 24(4), 141–164.

Bachmann, R., S. Elstner and E. R. Sims. (2013) Uncertainty and economic activity: Evidence from business survey data. *American Economic Journal: Macroeconomics* 5(2), 217–249.

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

Bar-Ilan, A. and W. C. Strange. (1996) Investment lags. *American Economic Review* 86, 610–622.

J. M., Barrero, N. Bloom and I. Wright. (2017). Short and long run uncertainty. *NBER Working paper* 23676.

Barsky, R. B. and E. R. Sims. (2011) News shocks and business cycles. *Journal of Monetary Economics* 58(3), 273–289.

Basu, S. and B. Bundick. (2017) Uncertainty shocks in a model of effective demand. *Econometrica* 85(3), 937–958.

Beaudry, P. and F. Portier. (2014) News-driven business cycles: Insights and challenges. *Journal of Economic Literature* 52(4), 993–1074.

Binder, C., T. McElroy and X. Sheng. (2022) Term structure of uncertainty: New evidence from survey expectations. *Journal of Money, Credit and Banking* 54(1), 39–71.
Bloom, N. (2009) The impact of uncertainty shocks. *Econometrica* 77(3), 623–685.
Bloom, N. (2014) Fluctuations in uncertainty. *Journal of Economic Perspectives* 28(2), 153–176.
Bloom, N., M. Floetotto, N. Jaimovich, I. Saporta-Eksten and S. J. Terry. (2018) Really uncertain business cycles, *Econometrica* 86, 1031–1065.
Boero, G., J. Smith and K. F. Wallis. (2015) The measurement and characteristics of professional forecasters' uncertainty. *Journal of Applied Econometrics* 30(7), 1029–1046.
Castelnuovo, E. (2019) Domestic and global uncertainty: A survey and some new results. Melbourne Institute Working Paper, no. 13/19.
Clements, M. P. (2017) Assessing macro uncertainty in real-time when data are subject to revision. *Journal of Business and Economic Statistics* 35(3), 420–433.
Cochrane, J. H. (1988) How big is the random walk component in GDP? *Journal of Political Economy* 96(5), 893–920.
Engle, R. F. and K. F. Kroner. (1995) Multivariate simultaneous generalised ARCH. *Econometric Theory* 11(1), 122–150.
Garratt, A., K. C. Lee and K. Shields. (2018) The role of uncertainty, sentiment and Cross-country interactions in G7 output dynamics. *Canadian Journal of Economics* 51(2), 391–418.
Garratt, A., K. Lee, M. H. Pesaran and Y. Shin. (2006) *National and Global Macroeconometric Modelling: A Long-Run Structural Modelling Approach*, Oxford: Oxford University Press.
Gilchrist, S., J. W. Sim and E. Zakrajsek. (2014) Uncertainty, financial frictions and investment dynamics. NBER Working Paper 20038.
Jo, S. and R. Sekkel. (2019) Macroeconomic uncertainty through the lens of professional forecasters. *Journal of Business and Economic Statistics* 37(3), 436–446.
Jurado, K., S. C. Ludvigson and S. Ng. (2015) Measuring uncertainty. *American Economic Review* 105(3), 1177–1216.
Koop, G., M. H. Pesaran and S. Potter. (1996) Impulse response analysis in nonlinear multivariate models. *Journal of Econometrics* 74(1), 119–147.
Lahiri, K. and F. Liu. (2006) Modelling multi-period inflation uncertainty using a panel of density forecasts. *Journal of Applied Econometrics* 21(8), 1199–1219.
Lahiri, K. and F. Sheng. (2010) Measuring forecast uncertainty by disagreement: The missing link. *Journal of Applied Econometrics* 25(4), 514–538.
Lahiri, K. and X. Sheng. (2008) Evolution of forecast disagreement in a Bayesian learning model. *Journal of Econometrics* 144(2), 325–340.
Leduc, S. and Z. Liu . (2016) Uncertainty shocks are aggregate demand shocks. *Journal of Monetary Economics* 82, 20–35.
Lee, K. C. and M. H. Pesaran. (1993) Persistence profiles and business cycle fluctuations in a disaggregated model of U.K. output growth. *Ricerche Economiche* 47(3), 293–322.
Lee, K. C., M. H. Pesaran and R. G. Pierse. (1992) Persistence of shocks and their sources in a multisectoral model of UK output growth. *Economic Journal* 102(411), 342–356.
Lee, K. C., M. H. Pesaran and R. G. Pierse. (1993) Persistence, cointegration and aggregation: A disaggregated analysis of output fluctuations in the US economy. *Journal of Econometrics* 56(1-2), 57–88.
Leeper, E. M., T. B. Walker and S.-C. C. Yang. (2013) Fiscal foresight and information flows. *Econometrica* 81(3), 1115–1145.
Ludvigson, S. C., S. Ma and S. Ng. (2021) Uncertainty and business cycles: exogenous impulse or endogenous response? *American Economic Journal: Macroeconomics* 13(4), 369–410.
Walker, S. and E. M. Leeper. (2011) Information flows and news-driven business cycles. *Review of Economic Dynamics* 14, 55–71.

Cite this article: Lee K, Shields K and Turnip G (2024). "Shock persistence, uncertainty, and news-driven business cycles." *Macroeconomic Dynamics* **29**(e43), 1–22. https://doi.org/10.1017/S1365100524000555