Animal Spirits, Fundamental Factors and Business Cycle Fluctuations*

Stephane Dées† Srečko Zimic‡

December 29, 2017

Abstract

We explore empirically the role of noisy information in cyclical developments to separate fluctuations due to genuine changes in fundamentals from those driven by temporary animal spirits (or noise shocks). Exploiting the fact that the econometrician has a richer dataset in some dimensions than the consumers, we use a novel identification scheme in a structural VAR framework and show that noise shocks are important drivers of business cycle fluctuations. In particular, noise shocks play a large role in consumption expenditures showing how false perceptions about future fundamentals influence consumer behaviours. By contrast, interest rates are much less impacted by noise shocks.

Keywords: Technology shocks, Noise shocks, Animal spirits, Business Cycles, Identification, Structural Vector Autoregression, Kalman Filter, Signal-extraction problem

JEL Classification: C32, E32

*Any views expressed represent those of the authors and not necessarily those of the European Central Bank or the Eurosystem. We are especially grateful to Fabio Canova and Evi Pappa. We would also like to thank Luca Benati, Luca Gambetti, Gianni Amisano, Marek Jarocinski, Jochen Guentner and participants at the IAAE 2014 Annual Conference as well as seminar participants at the ECB for helpful comments. Any remaining errors are the sole responsibility of the authors.

†Banque de France

‡European Central Bank

1 Introduction

The 2008/09 financial crisis and the Great Recession that followed has led many observers and academics to interpret the recession as a sharp decline in aggregate demand resulting from a collapse in confidence. This gave rise to a self-fulfilling shock to expectations.[1] This suggests that the way that households and firms form expectations of the future may therefore be an independent driving force of the business cycle.

The idea that agents' beliefs may be a source of economic fluctuations has a long history. Pigou (1929) was among the first to stress that expectations were key in explaining business cycles, as psychological factors (i.e. undue optimism and pessimism) lead entrepreneurs to make errors when forming their expectations about future profits. These errors generate cycles through rises and falls in investment. These psychological factors are also very often called "animal spirits", following Keynes (1936). Although the Real Business Cycle theory does not incorporate such psychological factors in its explanation of economic fluctuations, Pigou's ideas have recently been reintroduced into the theory of cycles in the context of equilibrium business cycle models, notably by Beaudry and Portier (2006) or Jaimovich and Rebelo (2009). In these models, although technology remains the only determinant of output in the long run, news about future fundamentals can imply a change in expectations, which affects agents' behaviors in the short run in anticipation of the fundamental change. However, economic agents receive only noisy signals about future technology, leading to expectational errors (Lorenzoni (2009)). If the information subsequently appears to have been wrong, i.e. it was just noise, the agents readjust their expectations and decisions accordingly. Conversely, if the information proves correct, the economy adjusts gradually to the level of activity consistent with technology. These changes in expectations generate economic fluctuations, both in the short and the long term.

Blanchard et al. (2013) explore the role of noisy information in cyclical developments empirically, separating fluctuations that are due to genuine changes in supply (news shocks) from those due to temporary expectational errors (noise shocks). They show that identification of those shocks is only possible via the estimation of a full structural model, such as the one in Barsky and Sims (2012). Since economic agents face a signal extraction problem when separating news from noise shocks, the econometrician, using the same data, cannot use structural VARs to recover such shocks. However, although this point is valid in real time, the econometrician can potentially have access to a richer

[1] See e.g. ECB (2013),Farmer (2012) and Bacchetta et al. (2012)

dataset in hindsight.

Based on this idea, this paper shows that a structural VAR model can be used to identify news and noise shocks. First, while economic agents can observe only current and past data, the econometrician can also observe "future" data. In other words, by using the data from the whole sample, the econometrician can have a better estimate of the technological trends than the economic agents. Second, economic agents only observe real-time data, while the econometrician also has access to revised data.

Recent papers have also proposed alternative ways to solve the issue put forward by Blanchard et al. (2013). Forni et al. (2013) use a modification of the structural VAR method to disentangle real from noise shocks, using future data and future residuals. However, their methodology is only applicable in economies in which the true state of the economy can eventually be exactly retrieved. In this paper we show that our methodology can be used to approximately identify supply and noise shocks in more general models, such as that of Blanchard et al. (2013), in which the true state of economy can never be retrieved. Enders et al. (2013) identify noise shocks in a standard VAR model by including 'nowcast errors', defined as the difference between actual output growth and growth estimated contemporaneously by professional forecasters. Here we also use nowcast errors of output growth to improve the estimates of potential output, but, unlike those authors, we also use the fact that econometricians have access to the later realizations of the time series.

We start our analysis by designing a slightly modified version of the model by Blanchard et al. (2013) and show how a structural VAR can be used to disentangle news and noise shocks. The methodology relies on the forecast errors consumers make when predicting the trend of GDP. These forecast errors are estimated by exploiting the fact that econometricians have access to 'future' and revised data. Applying this method to US data, we find that identified supply and noise shocks have effects as predicted by theory. A permanent supply shock has an expansionary effect on the economy, which builds through time until variables settle at a new, higher value. A noise shock also has an expansionary effect on the economy, but the impact fades away over time until all variables settle at their initial value. Nevertheless, noise shocks are as important as permanent supply shocks for business cycle fluctuations, explaining 15-20 percent of output variations at business cycle frequencies. On the other hand, permanent shocks drive the economy in the long run, but also account for around 20 percent of output variations at business cycle frequencies. Interestingly, consumption is even more affected by noise shocks (25-30 percent of variations in consumers' expenditures at business cycle frequencies), showing the extent to which false perceptions about the future fundamentals play a role in consumers' behaviours. By contrast, interest rates do not seem to be strongly affected by noise shocks, as central banks may be more immune to animal spirits.

After a presentation of the theoretical model in Section 2, we explain in Section 3 the problems related to identifying noise shocks in the data with SVAR models. We then show that, by using future observations and revisions of data, we can circumvent those problems and still use SVAR models to extract supply and noise shocks. In Section 4, we present empirical evidence on the effects of supply and noise shocks, by applying our methodology to US data. Several robustness checks are performed in Section 5 and they all confirm the main results. We conclude in Section 6.

2 Model

This section presents a simple model, similar to the model proposed by Blanchard et al. (2013), in which consumers decide their level of consumption based on their expectations about the economy's long-run fundamentals. Long-run economic fundamentals are driven by productivity developments (i.e. technology), which depend on a structural shock with permanent effects and which builds up gradually. Consumers do not observe the structural shock but only a noisy signal. This additional source of fluctuations is called a noise shock (or an "animal spirits" shock). As in Blanchard et al. (2013), consumers solve a signal extraction problem and decide their level of consumption on the basis of their expectations about future technology.

2.1 The structure of the model

We assume first that consumption is determined by the following Euler equation:

$$c_t = E_t[c_{t+1}|\mathcal{I}_t] \tag{2.1}$$

where c_t is consumption and $E_t[c_{t+1}|\mathcal{I}_t]$ is expected consumption in period $t+1$ based on the information set at time t, denoted $\mathcal{I}_t$. The supply side of the economy is completely determined by the demand side, which implies that output, y_t, equals consumption:

$$y_t = c_t \tag{2.2}$$

Output depends on utilization, u_t and the level of technology, a_t, in a linear fashion, $y_t = a_t + u_t$. Given the level of technology and consumption, utilization adjusts to produce the demanded level of output. However, in the long-run, output is equal to its natural level, and utilization is equal to zero, implying:

$$\lim_{t\to\infty} E_t[c_{t+j} - a_{t+j}] = 0 \tag{2.3}$$

As is shown in Blanchard et al. (2013), equation (2.3) can be derived from a standard New-Keynesian model with Calvo pricing, when the frequency of price adjustment goes to zero. Combining (2.1) and (2.3) gives:

$$c_t = \lim_{t\to\infty} E_t[a_{t+j}|\mathcal{I}_t] \tag{2.4}$$

which implies that consumption depends on expectations about long-run productivity.

The relevant state of the economy is productivity, a_t, which follows the process:

$$a_t = (1+\rho)a_{t-1} - \rho a_{t-2} + \epsilon_t \tag{2.5}$$

where $\epsilon_t \sim \mathcal{N}(0, \sigma_\epsilon^2)$ is a technology shock. Given that productivity is modeled as a process with a stochastic trend, the technology shock ϵ_t will have a permanent effect on productivity.

2.2 Information structure

The crucial difference with a standard DSGE model lies in the information structure. Consumers do not observe productivity directly, but observe only a noisy signal of productivity, s_t:

$$s_t \;\; = \;\; a_t + \upsilon_t \tag{2.6}$$

where $\upsilon_t \sim \mathcal{N}(0, \sigma_\upsilon^2)$ is a noise shock. The signal extraction problem can be rewritten as a state-space model:

- State Equation

$$\begin{bmatrix} a_{t|t} \\ a_{t-1|t} \end{bmatrix} = A \begin{bmatrix} a_{t-1|t-1} \\ a_{t-2|t-1} \end{bmatrix} + B \begin{bmatrix} \epsilon_t \\ \upsilon_t \end{bmatrix} \tag{2.7}$$

- Observation Equation

$$s_t = C \begin{bmatrix} a_{t|t} \\ a_{t-1|t} \end{bmatrix} + D \begin{bmatrix} \epsilon_t \\ \upsilon_t \end{bmatrix} \tag{2.8}$$

where $a_{t|t} = E_t[a_t|\mathcal{I}_t]$ and the matrices A, B, C and D depend on parameters. We assume that consumers know the underlying parameters of the economy and the distributions of shocks, in other words they know the matrices A, B, C and D, and thus they can use a Kalman filter to form their expectations about the current state of technology.[2]

Figure 1 shows, for simulated data, the developments in productivity, i.e. potential output growth, displayed by the smoothed line in the upper panel. Consumers receive a noisy signal, which is by definition very volatile and form their expectations by solving the signal extraction problem. Consumers' expectations about technology are not accurate as they are affected by the noise shock. The lower panel shows the forecast errors made by consumers when predicting the state of the economy.

Figure 1: Simulated technology, signal and consumer's expectations (upper panel) and forecast error (lower panel)

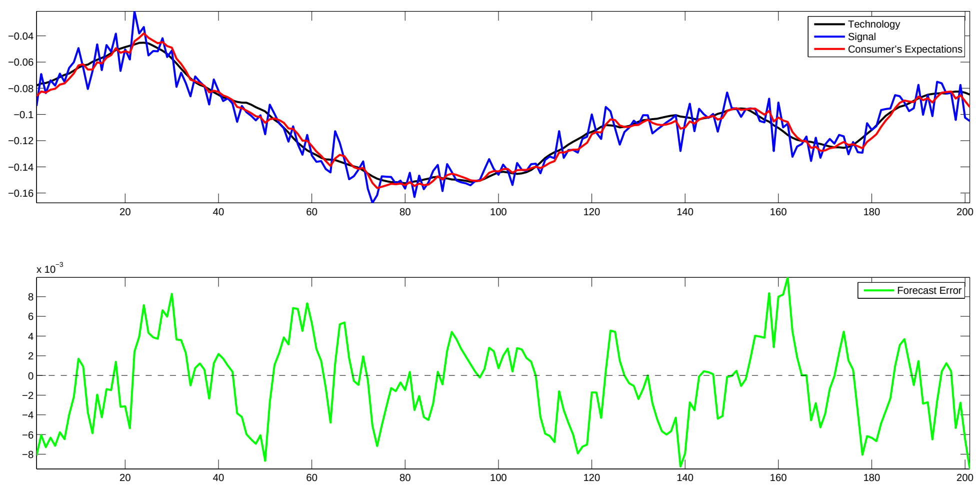

The black line represents the underlying technology process, which agents do not observe. The agents observe a signal (blue line) and via a Kalman filter, they form their expectations about technology (red line). The difference between their expectations and the true state of technology - forecast error - is represented in the graph below with the green line.

[2]Note that we are assuming that consumers do not observe utilization, otherwise agents could differentiate between technology shock and noise shock by observing utilization. Utilization is anyway difficult to observe at the macroeconomic level, even though proxies for utilization of labour can be found (e.g. hours worked). Moreover, we could also remove this assumption by introducing a transitory technological shock, as in Blanchard et al. (2013). By ignoring this extension, we aimed at keeping the presentation of the signal-extraction problem as simple as possible.

2.3 Model solution

Once we obtain the agent's expectations about the state of technology, by standard Kalman filtering, the solution of the model is straightforward. From Equation (2.4) and (2.5) we can derive the solution of consumption in terms of the agent's expectations about the state of technology:

$$c_t = \frac{1}{1-\rho}\left(a_{t|t} - \rho a_{t-1|t}\right) \tag{2.9}$$

where we have used the same notation as in the setup of the Kalman filter, $a_{t|t} = E_t[a_t|\mathcal{I}_t]$ and $a_{t-1|t} = E_t[a_{t-1}|\mathcal{I}_t]$. The other variables are a linear function of technology:

$$y_t = c_t \tag{2.10}$$

$$u_t = a_t - y_t \tag{2.11}$$

Following a supply (or technology) shock, ϵ_t, in the absence of noise shocks, the signal, s_t, is equal to a_t (see Equation (2.6)). Consumers underpredict the actual technological improvement in the short term, leading to a negative forecast error. Over time, as the signal confirms the increase in productivity, the expectations converge gradually to the new state of the economy and the forecast error fades away. Consumption improves faster in anticipation of the long-term effects of the supply shock.

Conversely, a noise shock is characterized by a one-off change in the signal and no change in productivity. In the short term, consumers believe that the positive signal could potentially be related to a change in technology and their expectations about the state of the economy are positive. Over time, they realize that the signal was just noise and the expectations gradually adjust to the initial state. As a result, the consumers' forecast error is positive in the short term and returns gradually to zero over time. This short-term optimism also leads consumption to be positively affected by the noise shock, as the shock has been unduly interpreted as a supply shock. Over time, when the information becomes more accurate, they realize their forecast error and readjust their consumption expenditures.

It is worth noting that, as in Lorenzoni (2009), the only source of exogenous uncertainty is the productivity process a_t and noise shocks have the features of aggregate demand shocks. As with productivity shocks, noise shocks are unobservable because they are related to a noisy signal of productivity. As a result, only these two shocks can lead to forecast errors in the same period about the current state of technology. In a more complex model, we could therefore assume that all other sources of disturbances are observable and do not lead to such forecast errors. Moreover, the general idea developed here can easily be carried over in more complex models with capital accumulation, nominal rigidities or habit persistence. For instance, Blanchard et al. (2013) show that such results are robust when they embed the same productivity process and information structure in a small-scale DSGE model, including investment and capital accumulation, nominal rigidities and a monetary policy rule.

3 From the model to a structural VAR

In this section, we focus on the way to identify and estimate supply and noise shocks. Since consumers only receive a noisy signal about the shocks, it is not possible to recover supply and noise shocks from actual data on c_t and s_t. First, we show that a VAR representation of the model faces an issue of non-invertibility and non-fundamentalness. Second, as the issue is mainly related to the fact that the information set used by consumers is not accurate enough to recover the shocks, we show that having a superior information set can help solve the problem. Third, we show how a structural VAR approach can be used in this context.

3.1 Singularity of VAR models

Let us consider running a VAR with consumption, c_t, and the signal, s_t, to obtain the structural shocks. From Equation (2.9) we know that consumption is a linear function of expectations about the current and past states of the economy:

$$c_t = f\left(E_t[a_t|\mathcal{I}_t]\right) \tag{3.1}$$

As the expectations about the current and past states are formed by a Kalman filtering approach, which is a linear filter, we also know that these expectations are a linear function of current and past signals:

$$E_t[a_t|\mathcal{I}_t] = g(s_t, s_{t-1}, ...) \tag{3.2}$$

Equations (3.1) and (3.2) imply that consumption is a linear function of current and past signals:

$$c_t = f(g(s_t, s_{t-1}, ...)) \tag{3.3}$$

Combining Equations (2.5) and (2.6), we can rewrite the signal process as $s_t = (1+\rho)a_{t-1} - \rho a_{t-2} + \epsilon_t + \upsilon_t$. We can notice that both shocks - supply and noise shocks - affect the signal. Moreover, both shocks affect the signal in the same way, or more formally:

$$\frac{\partial s_t}{\partial \epsilon_t} = \frac{\partial s_t}{\partial \upsilon_t} \tag{3.4}$$

For the same size shock, the signal will increase by the same amount to both shocks. As consumption is a linear function of the signal, consumption also responds in the same way to both shocks:

$$\frac{\partial c_t}{\partial \epsilon_t} = \frac{\partial c_t}{\partial \upsilon_t} \tag{3.5}$$

In other words, on impact consumers respond in the same way to supply and noise shocks, as they only observe the increase in the signal and are not able to differentiate between the two shocks.

This implies that running a VAR model with consumption and the signal results in a singular system. Moreover, if we extended our model with additional observables, such as stock prices, consumer sentiment or growth forecasts, this would not help to identify the shocks. In particular, so long as those observables are a linear function of the signal, we will have that:

$$\frac{\partial x_t}{\partial \epsilon_t} = \frac{\partial x_t}{\partial \upsilon_t} \tag{3.6}$$

To show the problems related to the use of a VAR model in such a case, we perform the following exercise: we use simulated data to estimate a non-singular VAR model with consumption and the signal by adding a measurement error to the signal.[3] Supply and noise shocks are identified by theoretically consistent long-run restrictions: the technology shock is identified as a shock with permanent effects on consumption and the noise shock is a shock with only transitory effects.

The impulse responses of consumption to those shocks are plotted in Figure 2, which shows that such an approach is not able to correctly identify the two shocks. While the first shock is permanent, it does not build up slowly, as with our theory-based technology shock, but jumps on impact. Similarly, we expect that the identified transitory shock would be a mixture of the measurement error and the theoretical noise shock. Figure 2 shows that, unlike the permanent supply shock, consumption does not respond at all to the identified transitory shock.

[3]The measurement error is assumed to be a white noise process with $\mathcal{N}(0, 0.0001\sigma_{\upsilon}^2)$.

Figure 2: IRF of consumption to a supply and noise shocks

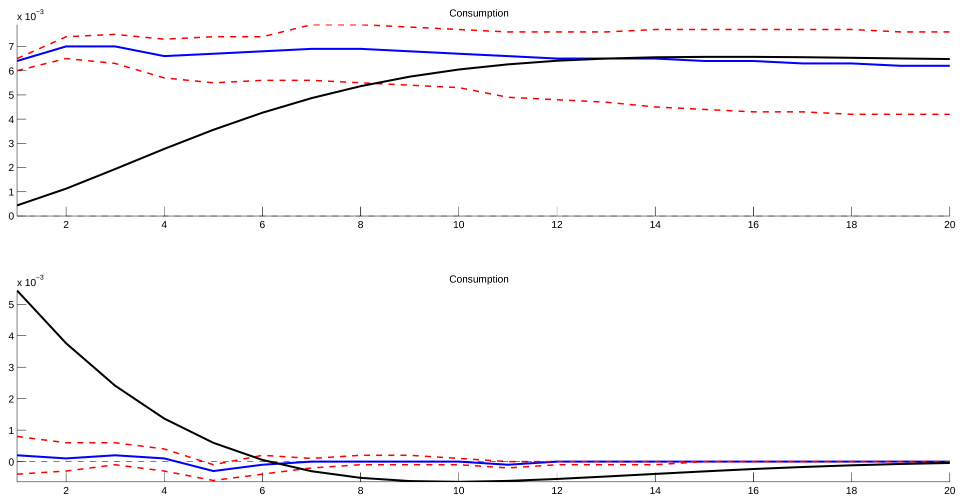

The graph presents the estimated impulse responses of consumption to a supply (above) and a noise (below) shock. The black lines are theoretical responses. The blue line is point estimate and the red lines are 90% error bands.

The above principles do not hold for the reaction of consumption in later periods. In other words, we have that:

$$\frac{\partial s_{t+j}}{\partial \epsilon_t} \neq \frac{\partial s_{t+j}}{\partial \upsilon_t} \quad \forall\ j = 1, 2, \ldots \quad \rightarrow \quad \frac{\partial c_{t+j}}{\partial \epsilon_t} \neq \frac{\partial c_{t+j}}{\partial \upsilon_t} \quad \forall\ j = 1, 2, \ldots \tag{3.7}$$

As consumers begin to differentiate between the two shocks, via the Kalman filter, the response of consumption differs in the periods following the shocks.

3.2 Consumers vs. Econometricians

The preceding discussion shows that to identify shocks we have to use a superior information set than that available to consumers. To start with an extreme situation, let us assume that the econometrician has access to time series of technology and can infer the consumers' forecast error:

$$\eta_t = E_t[a_t|\mathcal{I}_t] - a_t \tag{3.8}$$

Technology responds differently to the two shocks:

$$1 = \frac{\partial a_t}{\partial \epsilon_t} \neq \frac{\partial a_t}{\partial \upsilon_t} = 0 \tag{3.9}$$

and since the forecast error is a linear function of technology and the signal, it responds differently to the two shocks:

$$\frac{\partial \eta_t}{\partial \epsilon_t} \neq \frac{\partial \eta_t}{\partial \upsilon_t} \tag{3.10}$$

This fact is implicitly used in Enders et al. (2013) and Forni et al. (2013) to justify the identification of noise shocks. In the former paper, the authors use the assumption that forecast errors are observable at the end of the period and therefore the econometrician has access to forecast errors that can be directly included in a VAR. Similarly, Forni et al. (2013) assume that potential output is revealed after one period and therefore the econometrician can retrieve forecast errors. Had we assumed that the econometrician has access to forecast errors, for example by observing utilization, we would also be able to exactly identify the shocks.

However, in reality, the econometrician cannot observe the true technology and the related forecast errors. In a less extreme situation, we can assume that the econometrician does not have access to perfect information about the state of the economy, but can have access to superior information than the agents.

How well can the econometrician approximate the forecast error by using a superior information set, $\mathcal{I}_t^e$? We can define the forecast errors and the econometrician's estimate of the forecast errors as:

$$\eta_t = E_t[a_t|\mathcal{I}_t] - a_t \qquad \rightarrow \qquad \hat{\eta}_t = E_t[a_t|\mathcal{I}_t] - E_t[a_t|\mathcal{I}_t^e] \tag{3.11}$$

In the extreme case described above, when the econometrician has perfect information, $E_t[a_t|\mathcal{I}_t^e] = a_t$, and the econometrician would be able to exactly recover the forecast error, $\hat{\eta}_t = \eta_t$. Note that, in the opposite extreme, if the econometrician has the same information set as the agents, their estimate of the forecast error is 0.

With a more plausible situation, there are two ways in which the econometrician can use superior information to achieve a more precise estimate of the state of economy, i.e. achieve $var(E_t[a_t|\mathcal{I}_t^e] - a_t) < var(E_t[a_t|\mathcal{I}_t] - a_t)$:

- Using more accurate signals: $var[\upsilon_t|\mathcal{I}_t^e] < var[\upsilon_t|\mathcal{I}_t]$

- Using future observations: $\mathcal{I}_t \subset \mathcal{I}_t^e$

3.2.1 Using more accurate signals

In the real world, consumers and firms have to base their decisions on real-time data. As shown for instance by Diebold and Rudebusch (1991), forecast errors tend to be larger when using real-time data compared with those based on revised data.

It can be shown that:

$$\lim_{\sigma_v^2 \to 0} E[E_t[a_t|\mathcal{I}_t^e] - a_t] = 0 \tag{3.12}$$

which implies that in the limit, when the signal becomes perfectly informative, we can exactly recover the state and the forecast error, η_t. We can also show that:

$$\sigma_v^2 < \sigma_v^{2*} \;\; \to var[\eta_t|\sigma_v^2] < var[\eta_t|\sigma_v^{2*}] \tag{3.13}$$

which implies that more informative signals help to decrease the forecast error and the more precise the signal, the greater the decrease in the forecast error.

In order to see how more accurate signals help to estimate the forecast error, we run a simulation where consumers have a signal with the baseline variance, σ_v^2, while the econometrician can use a signal with smaller variance to predict the state. The estimate of the forecast error is then constructed as $\hat{\eta}_t = E_t[a_t|\mathcal{I}_t] - E_t[a_t|\mathcal{I}_t^e]$, where the difference in consumers' and the econometrician's information sets, $\mathcal{I}_t$ and $\mathcal{I}_t^e$, relates to the difference in the variance of the signal.

Table 1 shows how the estimated forecast error evolves according to the quality of the signal. When the signal becomes less noisy, the correlation between the estimated and the true forecast errors increase rapidly when the variance of the noise is reduced.[4]

Table 1: Correlation between estimated and true forecast errors

	σ_v^2	$0.75\sigma_v^2$	$0.5\sigma_v^2$	$0.25\sigma_v^2$	$0\sigma_v^2$
corr(η_t, $\hat{\eta}_t$)	0	0.88	0.93	0.94	1

[4] How much the reduction in the variance of the noise shocks contributes to the better estimate of the forecast error depends on the correlation between the improvements in the signal and the noise shock. When newly defined noise shocks with lower variance can be written as $v_t^* = av_t$, implying that the correlation between improvement in the signal and the noise shocks is perfect, the improvement in the signal contributes the most to the better estimate of the forecast error. In this subsection we assumed correlation between the improvements in the signal and noise shocks is perfect. See also footnote 6.

3.2.2 Using future observations

The econometrician can also potentially observe a larger dataset, including 'future' data. By contrast, consumers cannot observe 'future' realizations. If we define $\mathcal{I}_t^{e.j} = \{s_{t+j}, s_{t+j-1}, s_{t+j-2}, ...\}$, we can show:[5]

$$var[E_t[a_t|\mathcal{I}_t^{e.(j+h)}] - a_t] < var[E_t[a_t|\mathcal{I}_t^{e.j}] - a_t] \; \forall \; h = 1, 2, ... \tag{3.14}$$

which implies that having access to future signals helps to decrease the forecast error and the more leads are available, the greater the decrease.

In general, future data does not perfectly reveal the true state of technology and therefore the true forecast error, η_t. The gain one obtains by using future observations depends on the variance of the signal - the lower is the variance of the signal, the greater is the improvement in estimation accuracy that we can obtain by including future observations.

We perform another simulation allowing the econometrician to observe different number of leads of the signal. The forecast error is then estimated as:

$$\hat{\eta}_t = E_t[a_t|\mathcal{I}_t] - E_t[a_t|\mathcal{I}_t^{e.j}] \tag{3.15}$$

Table 2 shows how close to the actual forecast errors the estimated forecast error can become when we increase the number of leads. With four leads, the correlation between the estimated and the true forecast errors increases to 0.72.

Table 2: Correlation between estimated and true forecast errors

	0 leads	1 lead	2 leads	3 leads	4 leads
corr(η_t, $\hat{\eta}_t$)	0	0.44	0.58	0.67	0.72

It is important to note that it is this improvement in the estimate of a_t that allows the econometrician to estimate the forecast error. In fact, from the definition of $\hat{\eta}_t$, we see that the econometrician's estimate of the forecast error is equal to their correction, relative to the agents, of the estimate the true state. Thus, we can decompose $\hat{\eta}_t$ into a component due to the use of revised data, which we denote by $\hat{\lambda}_t$, and a component due to the use of future data, denoted by $\hat{\kappa}_t$. All the possible differences can be summarized in the following table:

[5] See Simon (2006), page 271.

Table 3: Forecast error estimation

	Kalman Filter	**Kalman Smoother**
Real time data	Consumers' expectation	$\hat{\kappa}_t$
Revised data	$\hat{\lambda}_t$	Econometrician's estimate ($\hat{\eta}_t$)

Using simulated data, Figure 3 shows the estimate of the forecast error using future signals, and its difference with the true forecast error. As shown in the lower panel, the estimated forecast error matches the patterns of the actual forecast error and the correlation between the two is relatively high (0.8). Similarly, Figure 4 shows the same graphs for an estimate of the forecast error using a less noisy signal. The correlation between the true forecast error and the estimated one is again high (0.8). Finally, Figure 5 shows the estimate of the forecast error that combines the previous two estimates. The correlation between this final estimate of the forecast error and the true one is very high (0.93), which supports our strategy.

Figure 3: An estimate of forecast errors due to the use of future signals - $\hat{\kappa}_t$ (simulated data)

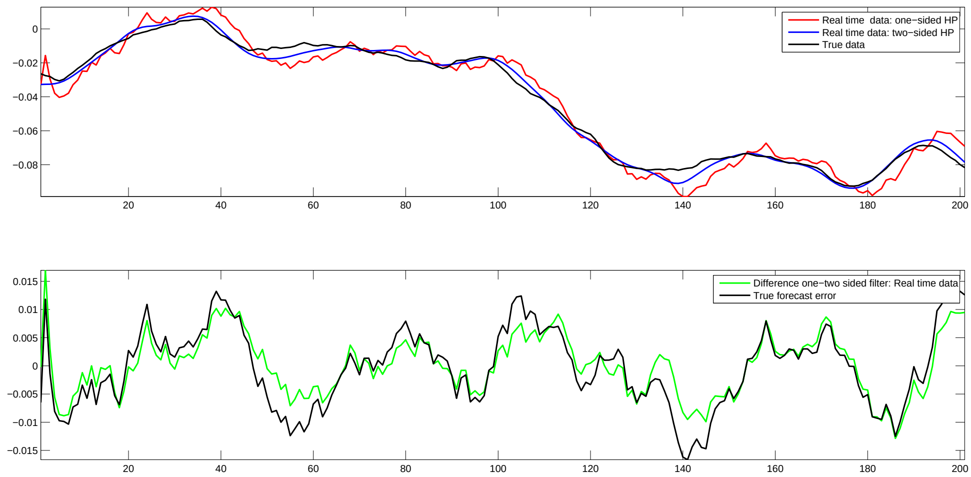

$$corr(\hat{\kappa}_t, \eta_t) = 0.8$$

The black line represents the true technology. The red line represents the expectations that are obtained via Kalman filter with the structure of HP filter. The blue line represents the expectations that are obtained via Kalman smoother with the structure of HP filter. The green line represents the difference between the two series of expectations, which is a proxy for the forecast error estimated due to the usage of future values. The black line are the true forecast errors.

Figure 4: An estimate of forecast errors due to the use of less noisy signals - $\hat{\lambda}_t$ (simulated data)

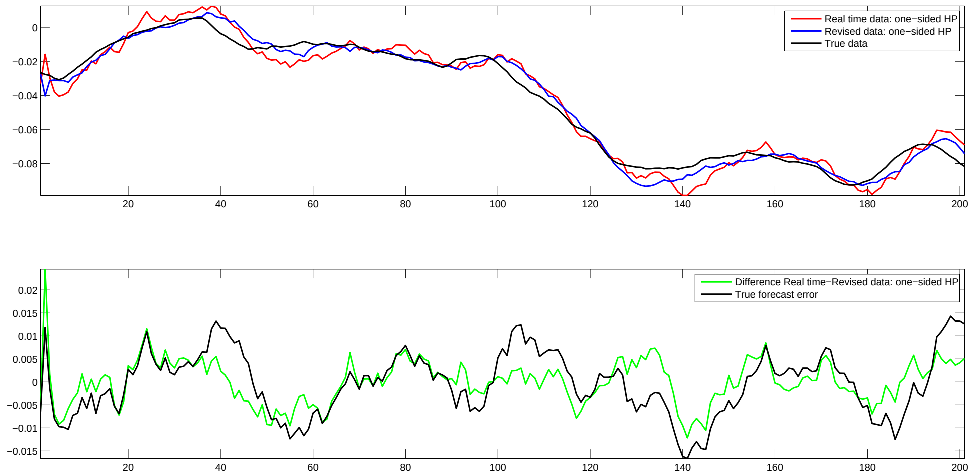

$$corr(\hat{\lambda}_t, \eta_t) = 0.8$$

The black line represents the true technology. The red line represents the expectations that are obtained via Kalman filter with the structure of HP filter using the real-time data. The blue line represents the expectations that are obtained via Kalman filter with the structure of HP filter using the revised data (less noisy signal). The green line represents the difference between the two series of expectations, which is a proxy for the forecast error estimated due to the usage of revised data. The black line are the true forecast errors.

Figure 5: A final estimate of forecast error - $\hat{\eta}_t$
(simulated data)

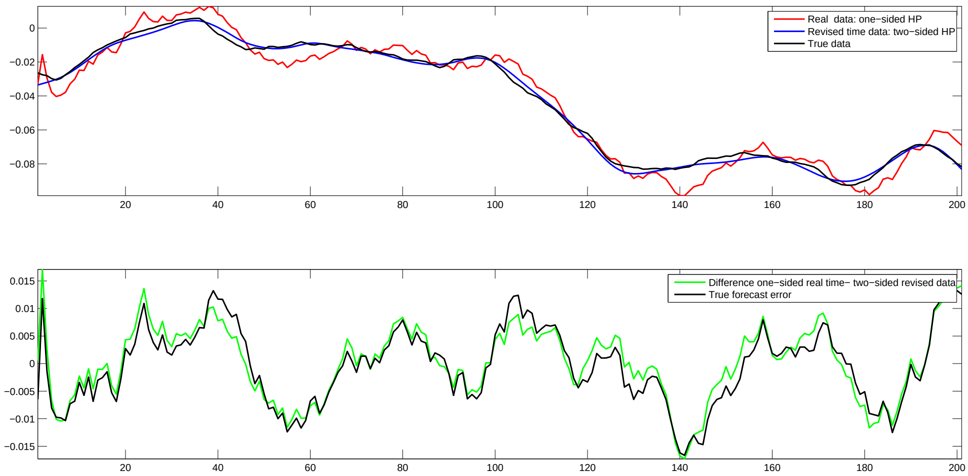

$$corr(\hat{\eta}_t, \eta_t) = 0.93$$

The black line represents the true technology. The red line represents the expectations that are obtained via Kalman filter with the structure of HP filter using the real-time data. The blue line represents the expectations that are obtained via Kalman smoother with the structure of HP filter using the revised data (less noisy signal). The green line represents the difference between the two series of expectations, which is a proxy for the forecast error estimated due to the usage of revised data. The black line are the true forecast errors.

3.3 Approximation with a VAR model

We have seen above that we can approximate the forecast error, $\hat{\eta}_t$, by using more precise data and future values of the signal. We have also shown that if we had access to the series of forecast errors, η_t, we could use a SVAR model to obtain the series of noise and supply shocks. Can we still use a SVAR model if we replace the original forecast errors by approximated forecast errors?

To answer this question, we estimate a SVAR model that includes consumption and instead of the true forecast error, as in Section 2, we use estimates of the forecast errors, $\hat{\eta}_t$. The estimates of forecast errors are obtained by using both future observations and more precise signals.[6]

The identification of the shocks is obtained by sign restrictions. From theory, we know that the forecast error responds negatively to the supply shock and positively to the noise

[6] The use of future observations may cause the VAR with forecast errors to be singular. However, when the improvements in the signal are not perfectly correlated with the noise shocks itself, i.e. the newly defined noise shocks with lower variance cannot be written as $v_t^* = av_t$, the problem of singularity does not appear. The detailed explanation can be found in Appendix A. In the following exercises we assume that the correlation between the improvements in the signal and the noise shocks is zero.

shock, while consumption responds positively to both. This implies that we can use the following sign restrictions to identify supply and noise shocks:

	Supply shock	Noise shock
Consumption	+	+
Forecast error ($\hat{\eta}_t$)	-	+

In the following example, we use a 50 percent less noisy signal and four leads of the signal to produce the econometrician's estimate of the forecast error, $\hat{\eta}_t$.

Figure 6 compares the theoretical and the estimated impulse responses to a supply shock (first row) and a noise shock (second row). The shape and size of the two forecast errors are very similar to the theoretical ones, implying that our sign restriction strategy is now able to distinguish between a supply and a noise shock. The response of consumption to these shocks is also similar to what the theoretical model predicts: after a supply shock, consumption increases gradually and is permanently affected by the shock; after a noise shock, consumption increases only in the short term and returns gradually to the baseline state after a few periods. Figure 7 shows the results of the same exercise by comparing estimated (first row) and theoretical (second row) forecast error variance decompositions (FEVD). The contributions of supply and noise shocks to the forecast error variance are very similar when comparing the estimated and the theoretical decompositions. Noise shocks explain most of the variance of consumption and forecast error in the short-run, while supply shocks mainly explain the variance of consumption in the long-run.

Figure 6: Estimated IRF by using estimated forecast errors (simulated data)

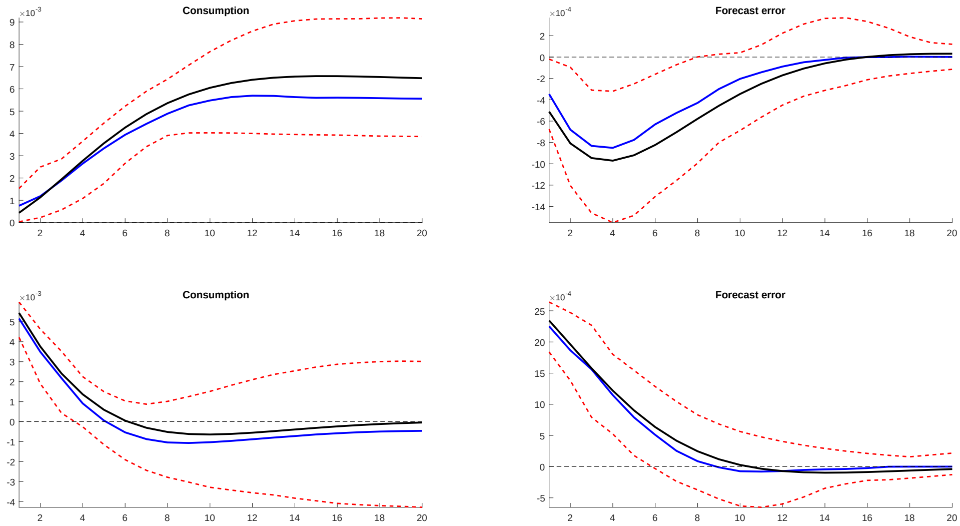

The graph presents the estimated impulse responses of consumption (right) and the forecast error (left) to the supply (above) and the noise (below) shock. The estimated forecast errors are obtained by using 4 leads of the signal and 50 percent less noisy signal. The black lines are theoretical responses. The blue line is point estimate and the red lines are 90% error bands.

Figure 7: Simulated data: Estimated FEVD (above) and theoretical FEVD (below)

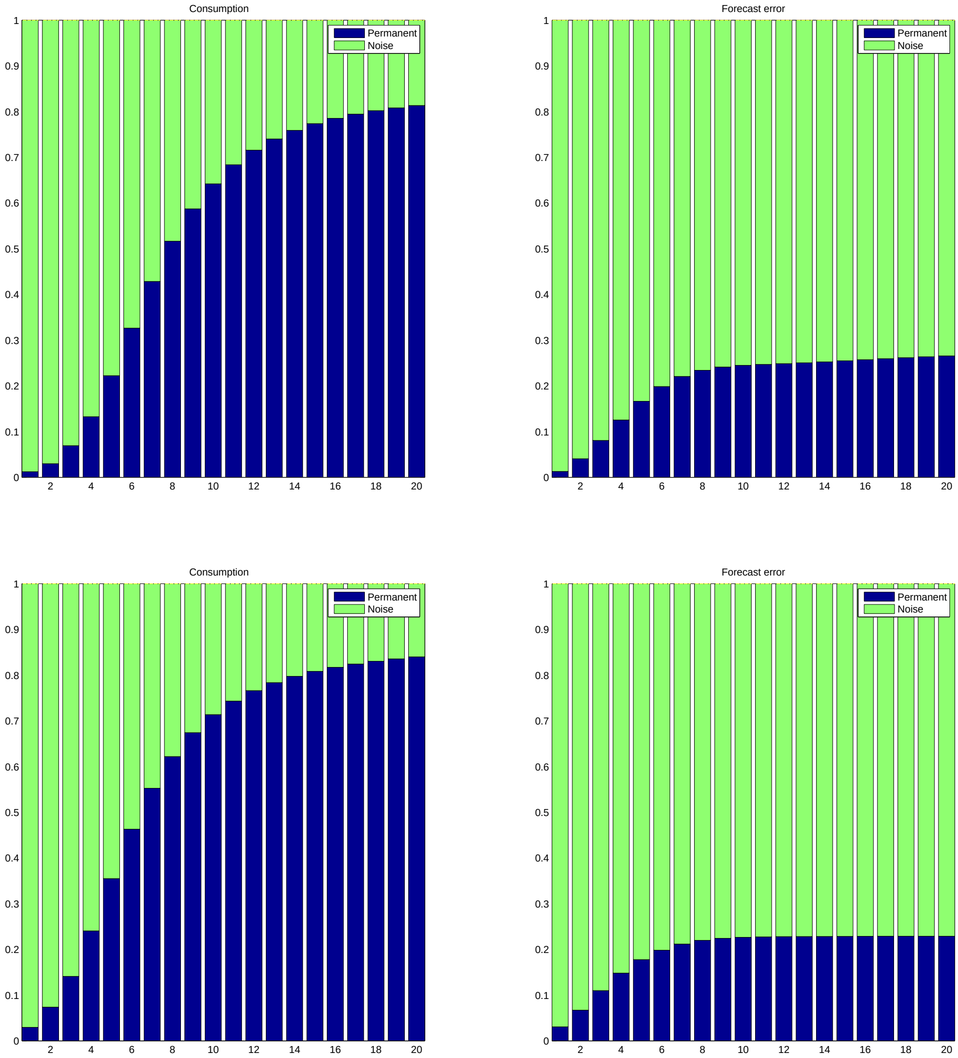

The first row presents the estimated forecast error variance decomposition (FEVD) of consumption (right) and the forecast error (left). The second row presents the theoretical forecast error variance decomposition (FEVD) of consumption (right) and the forecast error (left). The blue area corresponds to the median contribution of the supply and the green area corresponds to the median contribution of the noise shock. The estimated forecast errors are obtained by using 4 leads of the signal and 50 percent less noisy signal.

3.4 Comparison with Forni et al. (2013)

In this subsection we show that the methodology proposed in Forni et al. (2013) cannot be applied to a general signal extraction model as used in this paper and in Blanchard et al. (2013). Moreover, we show that our approach can easily accommodate their methodology.

The two crucial equations of the model in Forni et al. (2013) are:

$$a_t = a_{t-1} + \varepsilon_{t-1} \tag{3.16}$$

$$s_t = a_t + \upsilon_t \tag{3.17}$$

where a_t is the observable technology, s_t is a signal about technology, ε_t is the supply shock and υ_t is the noise shock. The signal extraction problem comes from the fact that the supply shock has delayed effects on technology $\Delta a_t = \varepsilon_{t-1}$. Importantly, contrary to the general model used in this paper, this framework implies that a VAR including the current observable technology and the signal is not singular. This is easy to see as $\frac{\partial s_t}{\partial \varepsilon_t} = \frac{\partial s_t}{\partial \upsilon_t}$, but $\frac{\partial a_t}{\partial \varepsilon_t} \neq \frac{\partial a_t}{\partial \upsilon_t}$ and therefore the problem of singularity of the VAR that was discussed in Section 3.1 does not appear in this case.

It is also straightforward to understand why our methodology can be applied to the model described by (3.16) and (3.17). Indeed, using one lead of technology enables to exactly recover the supply shock, $\Delta a_{t+1} = \varepsilon_t$. This implies that the estimated forecast error is equal to the true forecast error, $\hat{\eta}_t = \eta_t$, and the shocks can be exactly recovered.

4 Empirical evidence on US data

Given the positive results of the simulations above, we use our methodology to estimate the effects of supply and noise shocks in the US. We first explain how we construct the real-world counterpart of the forecast error. Second, we present the VAR model we use to identify supply and noise shocks. The empirical analysis is performed with two different VAR systems. First we use a small-scale VAR that includes only the forecast error and consumption and corresponds to the one used in the theoretical section. We show that results are consistent with the predictions of our theoretical model and we show that these results are robust to a different identification strategy.

In the next subsection we extend the VAR with additional variables. The main reasons to extend the VAR is that many shocks other than "animal spirits" can be compatible with a small set of restrictions used in the small-scale VAR. The bigger system allows us to isolate noise shocks from demand shocks and also allows us to study the effects of noise shocks on other variables, like stock prices, interest rates and inflation.

4.1 Constructing forecast errors

The forecast errors are constructed by taking the difference between a real-time estimate of potential GDP and the actual/official estimate provided by the CBO. To compute our estimate of real-time potential GDP, we use the real-time estimate and projections of the output gap available in the Greenbook (see https://www.philadelphiafed.org/research-and-data/real-time-center/greenbook-data) and add it to the real-time level of GDP. 'Standard' real-time data - the first estimates of GDP available in period t from statistical offices - is normally available only after at least one quarter. In order to be consistent with the model, we have to construct the measure of real GDP in period t that was available exactly in period t. To this end, we use GDP forecasts for the current period from the Survey of Professional Forecasters.

The real time data for real GDP, y_t^{rt}, is constructed as:

$$y_t^{rt} = y_{t-1}^* \left(1 + \Delta \hat{y}_t\right) \tag{4.1}$$

where y_{t-1}^* is the first estimate of real GDP in period $t-1$ as provided by the statistical office and $\Delta \hat{y}_t$ is the forecast for quarterly real GDP growth rate in period t as produced by the Survey of Professional Forecasters.

The calculation of real-time GDP is further complicated by the fact that we are interested in the GDP level, but a consistent series for the GDP level is not available from the SPF. Namely, the level forecasts from SPF are affected by changes in national accounts - for example changes in base years - and it is therefore hard to reconstruct a consistent measure of the real-time level of GDP. To circumvent this problem, we use a consistent revised series and use the growth rates based on the first vintages of the data to move to real-time GDP levels in period $t-1$. Specifically, the first estimate of real GDP in period $t-1$ as provided by the statistical office, y_{t-1}^*, is constructed in the following way: $y_{t-1}^* = y_{t-5}\left(1 + \Delta^4 y_{t-1}^*\right)$ where y_{t-5} is the revised GDP data in period $t-5$ (based on the last available vintage in Q2 2012) and $\Delta^4 y_{t-1}^*$ is the yearly growth rate of real GDP between $t-5$ and $t-1$ calculated from the first available vintage for period $t-1$, which is the one available in period t.[7]

[7] With this approximation we do not take into account the revisions that happen after the 5-quarters, as our approximation assumes $y_{t-5}^* = y_{t-5}$. Our measure can therefore underestimate the error due to the use of real time data, however, the revisions after five quarters are not substantial and therefore we opt for this approximation. We could also 'went further' into the history by assuming for example, $y_{t-9}^* = y_{t-9}$. However, in that case we would loose additional one year of data, while the gain is not large as revisions after 5-quarters are not large.

4.2 Small-scale VAR

We start our empirical analyis with a small-scale VAR that includes the estimated forecast error about the growth rate of potential output and real consumption. The estimated forecast errors are those defined in the previous subsection. The estimation period is from Q3 1977 to Q4 2011. The system is estimated in the level of forecast error and the log of consumption with 8 lags.

4.2.1 Identification

The identification of supply and noise shocks can be achieved in two ways. Firstly, we can impose sign restrictions that are based on the different response of the forecast error to the supply and noise shocks. The second way is to rather use the behavior of consumption - we know that noise shocks should lead to transitory response of consumption, while supply shocks should affect consumption permanently. Therefore, we can identify the supply and noise shocks by imposing long-run restrictions la Blanchard and Quah (1989).

We start by identifying the two shocks by sign restrictions. From the theoretical model we know that:

- The forecast error is positive for positive noise shocks (i.e. consumers are too optimistic).
- The forecast error is negative for positive supply shocks (i.e. consumers are too pessimistic).

We also know that consumption will increase after following both shocks. Hence, the sign restrictions that we use to separate the noise from supply shock are the following:

Table 4: Imposed sign restrictions

	Forecast error	Consumption
Supply shock	-	+
Noise shock	+	+

To cross-check whether results are robust to alternative identification strategy, we also use long-run restrictions proposed in Blanchard and Quah (1989). In particular, the noise shock is identified by imposing that it has transitory effects on consumption, while the supply shock are shock that has permanent effect on consumption. In a small-scale VAR with only consumption and the forecast error, we identify the permanent and transitory shocks to consumption using the following restrictions on the matrix of long-run effects, Ξ:

$$\Xi = \begin{bmatrix} * & 0 \\ * & * \end{bmatrix} \tag{4.2}$$

implying the second shock does not have permanent effect on consumption and therefore can be considered as a noise shock in our setup.

4.2.2 Empirical results - sign and long-run restrictions

Figure 8 shows the impulse responses of consumption and the forecast error to a noise shock, using the sign restriction methodology described above. The estimation algorithm is described in Appendix B. The impacts are positive and significant in the short run (up to 10-15 quarters). Once the consumer realises that the shock is not related to an improvement in productivity, the forecast error returns to 0 and the impact on consumption becomes non-significant. Figure 9 shows the impulse responses of the same variables following a supply shock. As expected, the impact on the forecast error is negative in the short run, while the impact on consumption is limited. Over time, as the consumer realises that the shock is a permanent supply shock, the forecast error returns to 0 and consumption improves permanantly.

To cross-check these results with the identification strategy based on long-run restrictions, we perform again the same IRFs (see Figures 10 and 11). The responses are qualitatively similar to those obtained with the sign restriction approach. In particular, without imposing the sign restrictions, we again find that a positive permanent supply shock implies negative forecast errors in the short term, while a positive transitory noise shock increases the forecast error, being in this case positive.

Figure 8: Small-scale VAR - IRFs to a noise shock (sign restrictions)

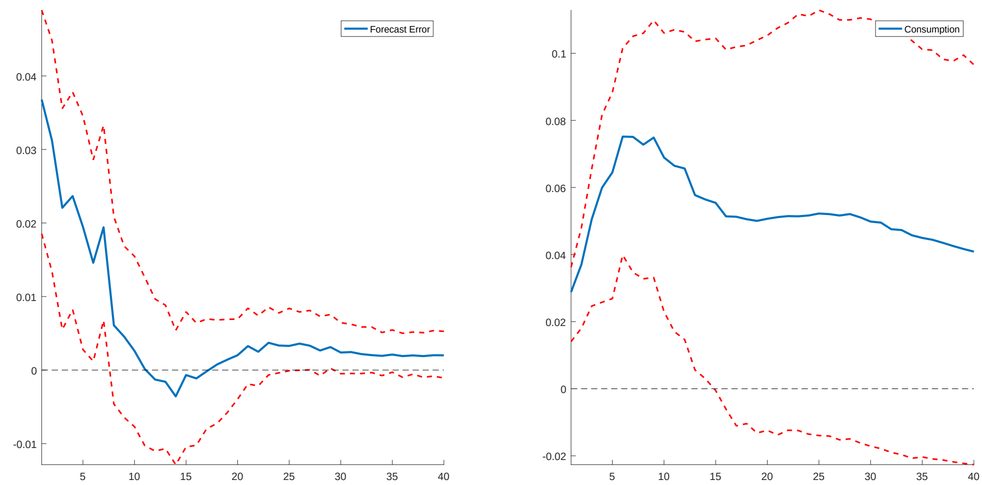

The graph presents the impulse responses to a noise shock of the estimated forecast error and consumption.The blue line is the point estimate and the red lines are 90% error bands.

Figure 9: Smal-scale VAR - IRFs to a supply shock (sign restrictions)

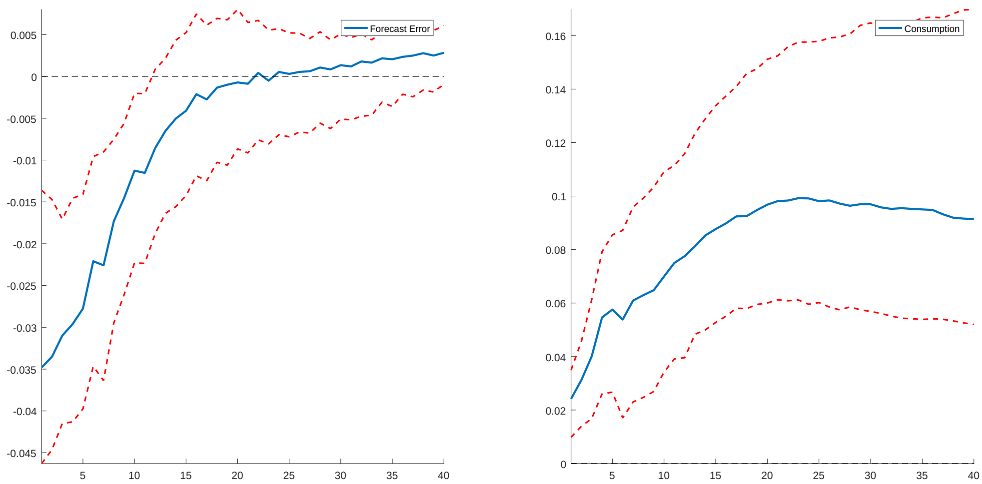

The graph presents the impulse responses to a supply shock of the estimated forecast error and consumption. The blue line is the point estimate and the red lines are 90% error bands.

Figure 10: Smal-scale VAR - IRFs to a noise shock (long-run restrictions)

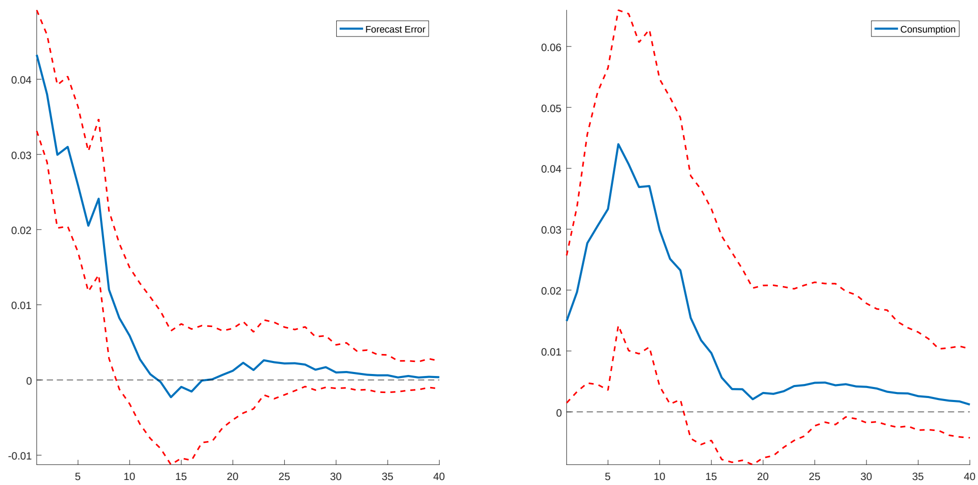

The graph presents the impulse responses to a noise shock of the estimated forecast error and consumption.The blue line is the point estimate and the red lines are 90% error bands.

Figure 11: Small scale-VAR - IRFs to a supply shock (long-run restrictions)

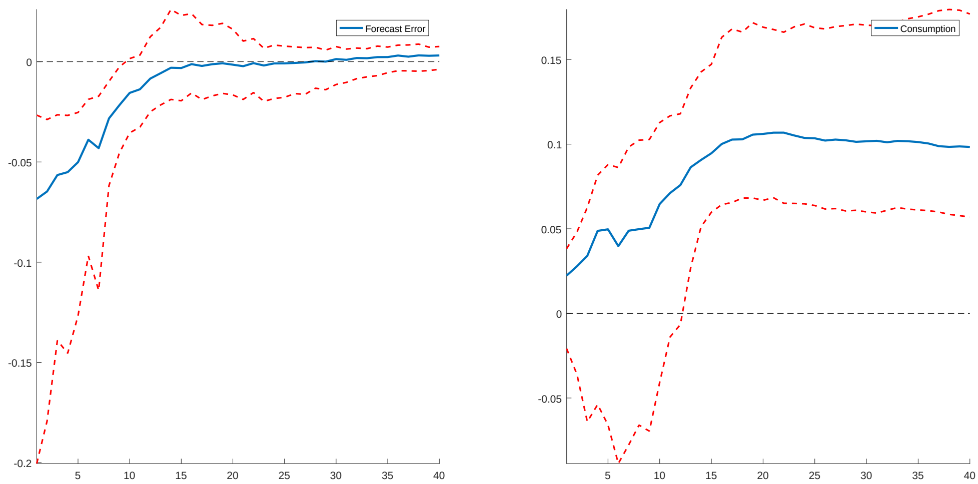

The graph presents the impulse responses to a supply shock of the estimated forecast error and consumption. The blue line is the point estimate and the red lines are 90% error bands.

4.3 Large-scale VAR

The smal-scale VAR used in the previous section has served as a test of the theory on the data, but might not have empirical appeal. The drawback of small-scale VARs is that they use little information in the data and, with a two-variable VAR, we assume that only two structural shocks drive the data. Therefore, in this section we enrich the VAR model by including additional variables and add additional identifying restrictions that mainly rely on the above theory. The extended setup allows us to identify possible additional sources of shocks and see the impact of noise shocks on additional variables, such as stock prices or consumer sentiment. The extended VAR model includes the following variables:

- Estimated Forecast Errors (FE)
- GDP (Y)
- Private Consumption (C)
- Stock prices (S&P500 - SP)
- Fed funds rate (IR)
- CPI inflation (INF)
- Consumer Sentiment (SENT)

The estimation period is still from Q3 1977 to Q4 2011. The system is estimated in log-levels with four lags, except for inflation rate and interest rate that enter in level. The results are robust to changes to differences and using a different number of lags.

4.3.1 Identification

The main identification assumption to separate the noise and the permanent supply shocks is similar to the section above - we use sign restrictions on the forecast error and consumption to separate the two shocks. Moreover, a larger VAR allows us to distinguish, among the transitory shocks, those that are related with “animal spirits” from those that are related to more standard shocks, in particular aggregate demand shocks.

In the theoretical model, the agents cannot distinguish between noise and supply shocks. This implies that in the first period the responses of variables to a supply or a noise shock should be identical. This allows us, with some additional restrictions commonly used in the literature, to use sign restrictions to separate supply and noise shocks from standard demand shocks. We can use this insight to separately identify a standard demand shock with additional restrictions on inflation and interest rates. Exact restrictions can be seen in Table 5:

Table 5: Imposed sign restrictions

	FE	Y	C	SP	IR	INF	SENT
Supply shock	-	+	+	+	-	-	+
Noise shock	+	+	+	+	-	-	+
Demand shock	+	+	+	+	+	+	+

The identification strategy in the empirical model assumes the permanent supply shock and the noise shock are both perceived as shocks from the supply side. Therefore it is assumed that all three shocks lead to an increase in output, consumption, stock prices and sentiment. To identify demand shocks, we follow the literature and postulate that demand shocks lead to an increase in inflation and interest rates (see for instance Peersman and Straub, 2009). By contrast, we assume that noise shocks and supply shocks lead to a decrease in inflation and interest rates, given they are perceived as shocks to technology.[8] Demand shocks are treated as transitory shocks and therefore we assume that estimated forecast error increases after the demand shock.[9] In accordance with theory, sign restrictions are imposed only in the first period. Robustness checks with respect to identification assumptions are performed in Section 5.

4.4 Empirical results

We now present the estimated impulse response functions, forecast error variance decompositions and extracted shocks of our large-scale VAR model.

Figure 12 shows the IRFs following a 1-standard deviation positive noise shock. In this case, the forecast error is positive and increases up to the 5th quarter after the shock. As expected by the theoretical model, the impact on output and consumption is positive and long lasting, remaining significantly above the baseline after 20 to 25 quarters. Equity prices and sentiment are also affected positively but, after 5 quarters, the effects are no

[8] We postulate that supply shocks lead to a decline in interest rates in line with sticky-price models la Ireland (2002), although in the literature this restrictions is not generally imposed. In the robusteness section below, we show that our results are robust to lifting this restriction.

[9] In case the forecast errors correspond to the theory counterpart, they should be un-correlated with other observable shocks. This is an unlikely case in our implementation of the filter-smoother procedure as output is the only observable and therefore we can expect that estimated forecast errors are correlated with other observable shocks. We have also performed an exercise where the third shock is identified with the zero restriction on the estimated forecast errors, but this shock was found to have almost no effect on business cycle fluctuations. This is additional empirical evidence that estimated forecast errors are correlated with other observable shocks. Results are available on request from authors.

longer significant. Inflation and interest rates slightly decline on impact.

Figure 13 shows impulse response functions (IRFs) following a 1 standard deviation positive supply shock. Although we only impose sign restrictions in the short term, the technology shock implies permanent effects on macroeconomic variables, as expected from theory. Indeed, output and consumption are permanently higher. As also expected from the theoretical model, the forecast error is significantly negative for 10 quarters, before returning towards the baseline. Inflation and interest rates decline on impact but are not permanently affected. Interestingly, the increase in equity prices and consumer sentiment following the supply shock is rather long lasting.

Figure 14 shows the IRFs of a demand shock. As imposed by the identification procedure, the forecast errors are positive in the short term. The impacts on the various variables are short lived, all of them coming back to baseline levels after 5-10 quarters following the shock.

The forecast error variance decomposition (FEVD), presented in Figure 15, shows that the noise shock explains 15-20 percent of the business cycle fluctuations in the short term. As expected from the theory, the impact of the supply shock on the business cycle builds up gradually. It also explains up to 20 percent of output variations at business cycle frequencies and almost 40 percent of the fluctuations in the long term, when the noise shock contribution is much smaller. It is worth noting that the noise shock explains a large share of fluctuations in consumption in the short term showing the impact of forecast errors made by consumers on their expenditure behavious. Conversely, interest rates are less sensitive to noise shocks compared with demand shocks, as central banks take more time to react to supply shocks in general and may be more immune to errors made through animal spirits.

Figure 16 shows the series of extracted shocks, with supply shocks in the upper panel and noise shocks in the lower panel. The patterns between these two shocks are very different. A few interesting observations are worth pointing out. First, the supply shock declines just before recessions and increases during the recessions, which could be interpreted as a Schumpeterian creative-destruction mechanism driving the recovery phase. Second, the noise shock was very positive during the dot-com bubble in the early 2000s and went down when optimism about future technologies was affected by the 2001 terrorist attacks, the Enron scandal and the stock market decline in 2001-02. Third, noise shocks tend to remain negative for a prolonged period of time after a recession, underlining the role of excess pessimism as a dampening factor in the recovery phases. Similarly, noise shocks were very negative during the global financial crisis in 2009.

Figure 12: Large-scale VAR - IRFs to a noise shock (extended identification)

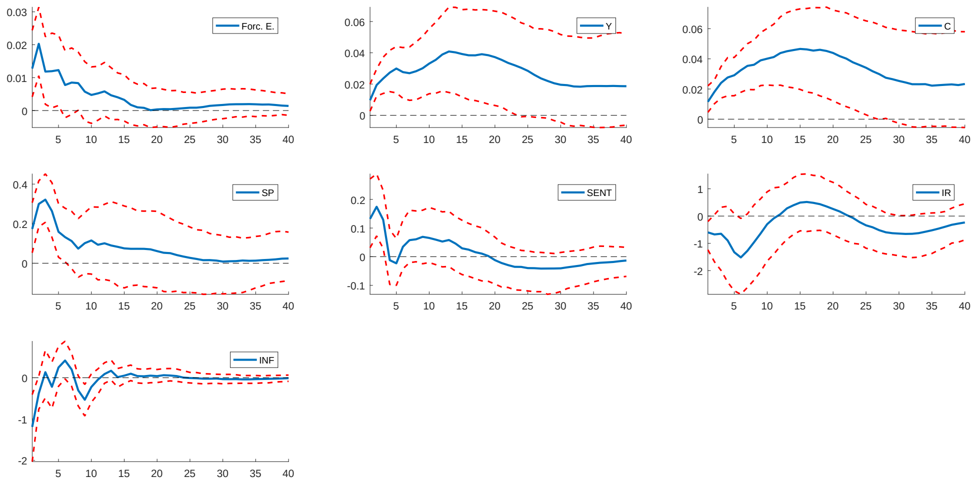

The graph presents the impulse responses to a noise shock of the estimated forecast error, output, consumption, investment (first row), stock prices, sentiment, interest rates and inflation (second row).The blue line is the point estimate and the red lines are 90% error bands.

Figure 13: Large-scale VAR - IRFs to a supply shock (extended identification)

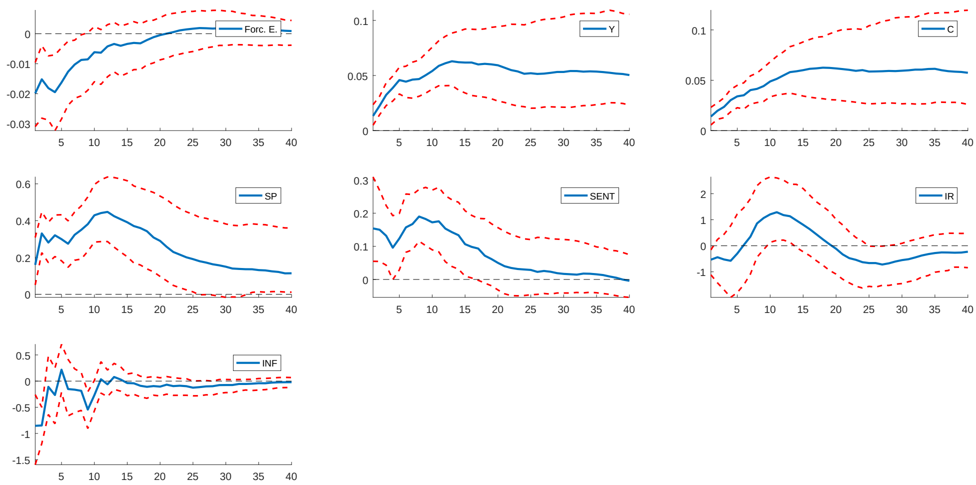

The graph presents the impulse responses to a supply shock of the estimated forecast error, output, consumption, investment (first row), stock prices, sentiment, interest rates and inflation (second row). The blue line is the point estimate and the red lines are 90% error bands.

Figure 14: Large-scale VAR - IRFs to a transitory demand shock (extended identification)

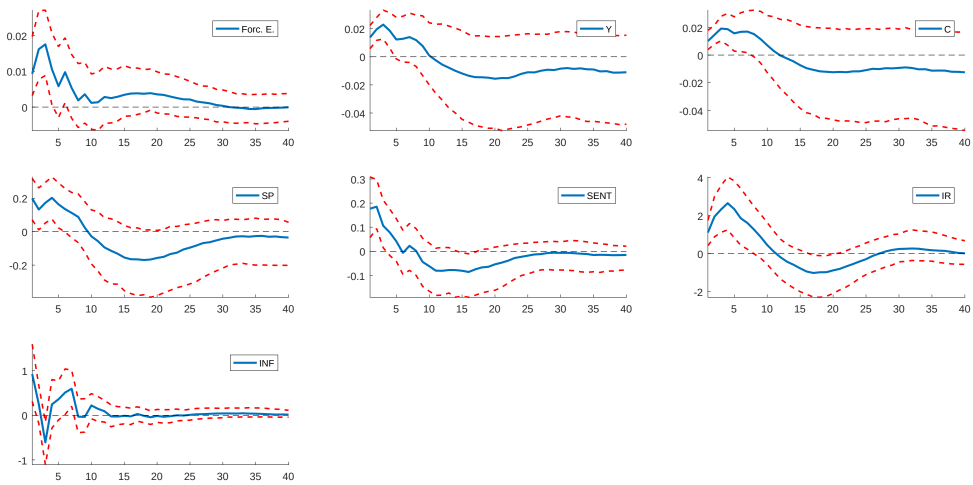

The graph presents the impulse responses to a transitory demand shock of the estimated forecast error, output, consumption, investment (first row), stock prices, sentiment, interest rates and inflation (second row). The blue line is the point estimate and the red lines are 90% error bands.

Figure 15: Large-scale VAR - FEVD (extended identification)

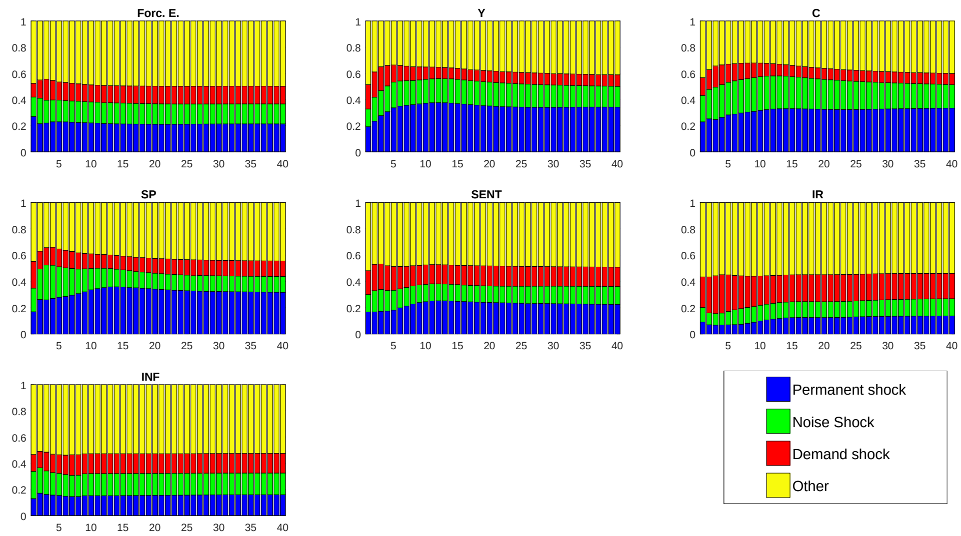

The graph presents the forecast error decomposition of estimated forecast error, output, consumption (first row), investment, stock prices, sentiment (second row), interest rates and inflation (third row). The blue area corresponds to the median contribution of supply shocks, the green area to the median contribution of noise shocks, the red area to the median contribution of transitory demand shocks and yellow area to the contribution of all other non-identified shocks.

Figure 16: Large-scale VAR - extracted shocks (smoothed)

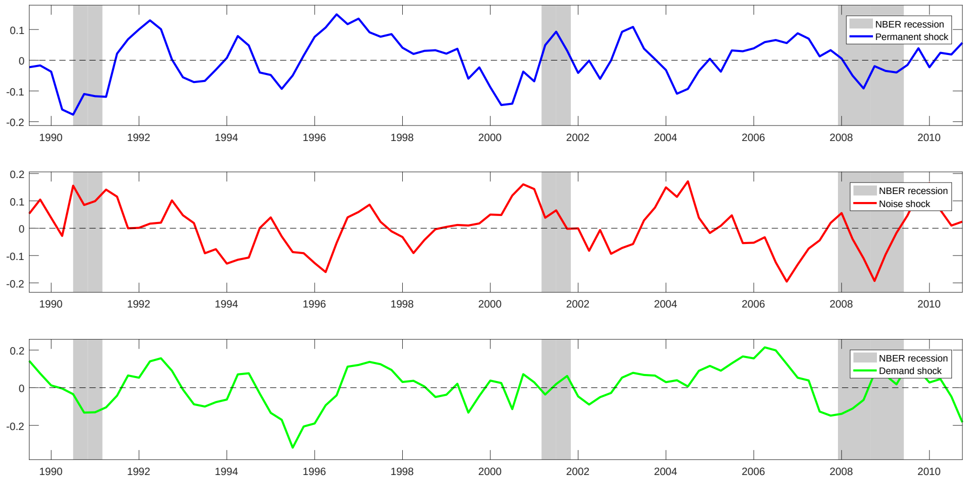

The graph presents the extracted supply (blue line above), noise shocks (red line in the middle) and demand shocks (green line below). The series of shocks are smoothed with moving average filter with 5-period window. The gray areas correspond to NBER-defined recessions.

5 Robustness checks

To confirm the robustness of our results we perform two additional exercises. First, instead of constructing the forecast error from real time data from Greenbook and revised data from CBO, we estimate forecast errors by using a similar filter to the one used in theoretical section. Secondly, we check whether our results are driven by a specific identification strategy.

5.1 Alternative Forecast Error

Because we estimate forecast errors, we do not rely on data availability on potential GDP and can therefore extend the sample to 1970Q1-2012Q2. We drop 4 observations at both ends of the sample to account for the bias in the estimation of the forecast error at the beginning and at the end of the sample.

5.1.1 Estimating the forecast error

Instead of using the real-time data from the Greenbook and the revised data from the CBO, we use an alternative way to estimate the forecast errors by relying on the filter detailed in the theoretical part.

The relevant state variable in the model is technology. However, although it would be possible to exactly map the model to the empirical applications by assuming technology as the relevant state, we opt for using the growth rate of the potential GDP as the relevant state.

The main reason behind this choice is that real-time data for productivity does not exist, except in the case of output per hour worked in the business sector, which is anyway only available since 1998.[10] Moreover, in our opinion, firms and consumers focus more on trend GDP rather than technological trend when making economic decisions. As GDP is the best available indicator of economic activity, trend GDP is also the usual measure of potential output monetary policy authorities typically use when setting monetary policy (Gavin 2012).

To estimate the potential GDP we use unobserved component model similar to the one in theoretical section above. The model is:

1. State Equation

$$\begin{bmatrix} \tau_{t|t} \\ \tau_{t-1|t} \end{bmatrix} = \begin{bmatrix} 1+\rho & -\rho \\ 1 & 0 \end{bmatrix} \begin{bmatrix} \tau_{t-1|t-1} \\ \tau_{t-2|t-1} \end{bmatrix} + \begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix} \begin{bmatrix} \epsilon_t \\ 0 \end{bmatrix} \tag{5.1}$$

2. Observation Equation

$$y_t = \begin{bmatrix} 1 & 0 \end{bmatrix} \begin{bmatrix} \tau_{t|t} \\ \tau_{t-1|t} \end{bmatrix} + \begin{bmatrix} 0 & 1 \end{bmatrix} \begin{bmatrix} 0 \\ \eta_t \end{bmatrix} \tag{5.2}$$

where $\epsilon_t \sim \mathcal{N}(0, \sigma_\epsilon^2)$ and $\eta_t \sim \mathcal{N}(0, \sigma_\eta^2)$. In this model y_t is real GDP, τ_t is its trend component, ϵ_t is its cyclical component and η_t is a white noise sequence. Normalizing the variance of the noise component to one, $\sigma_\eta^2 = 1$, we can write the signal-to-noise ratio as $\zeta = \frac{\sigma_\eta^2}{\sigma_\epsilon^2}$. To calibrate the two parameters of the filter, ρ and ζ, we minimize the distance between the state estimate of the potential output using (5.1) with the estimate

[10] See https://www.philadelphiafed.org/research-and-data/real-time-center/real-time-data/data-files/OPH/

of the potential output by Congressional Budget Office (CBO). The minimization results in setting $\rho = 1.04$ and $\zeta = 2808$.

Once we have parameterized the state space system we can estimate the forecast error in the same fashion as explained in Section 3.2. The correspondence between the model and the empirical application can be summarized as follows:

- **Use of future data** - Consumers only use past data and determine trend GDP with a Kalman filter, which is a one-sided filter. The econometrician also can use 'future' data when filtering the GDP series with a Kalman smoother, which is a two-sided filter.[11]

- **Better signal** - Consumers can only use real-time data. The econometrician, having access to revised data, uses more accurate information about the state in a given period.

The estimates of the forecast errors can then be decomposed as in Table 3. Figures 17 and 18 show the estimated forecast errors coming respectively from the use of future data and the improvement in the signal. Figure 19 shows the final estimated forecast error. We can see the forecast error is pro-cyclical: it is positive and increases in the expansion phase, peaks at the start of a recession (as defined by the NBER) and then falls, reaching a minimum around two years after the end of the recession phase. A positive, rising error can then be interpreted as an increase in consumers' optimism up to the end of the expansion phase. The error becomes negative during the recession and stays negative as long as consumers remain pessimistic about the economy, thus underpredicting the actual state of the economy.[12]

[11] We have also experimented with a Kalman filter that includes 4 forward leads of data instead of using a Kalman smoother. The results were not much affected by this.

[12] The forecast error can be further decomposed into the component due to the use of future signals and the one due to the use of less noisy signals. We have not included the graphs here, but are available on request form the authors.

Figure 17: An estimate of forecast error due to the use of future signals - $\hat{\kappa}_t$ (US data)

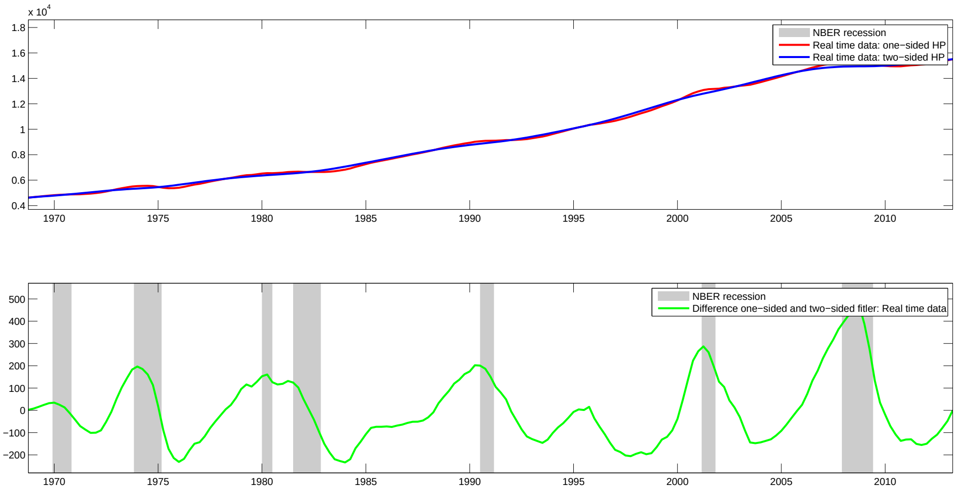

$$corr(\hat{\kappa}_t, \eta_t) = 0.68$$

The red line represents the trend GDP that is obtained via a Kalman filter with the structure of a HP filter. The blue line represents the trend GDP that is obtained via a Kalman smoother with the structure of a HP filter. The green line represents the difference between the two estimated trends, which is a proxy for the forecast error estimated due to the use of future values.

Figure 18: An estimate of forecast error due to the use of less noisy signals - $\hat{\lambda}_t$ (US data)

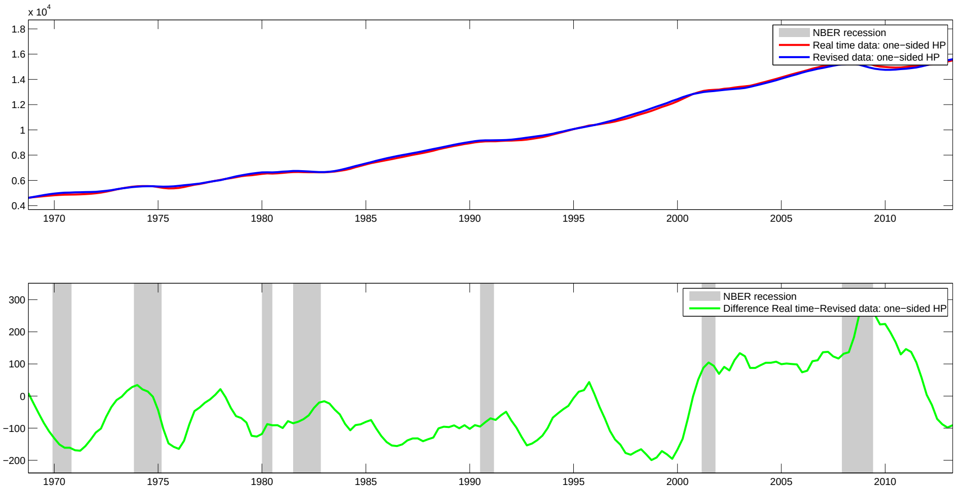

$$corr(\hat{\lambda}_t, \eta_t) = 0.92$$

The red line represents the trend GDP that is obtained via a Kalman filter with the structure of a HP filter using real-time data. The blue line represents the trend GDP that is obtained via a Kalman filter with the structure of a HP filter using revised data. The green line represents the difference between the two estimated trends, which is a proxy for the forecast error estimated due to the use of more precise signals.

Figure 19: A final estimate of forecast error - $\hat{\eta}_t$
(US data)

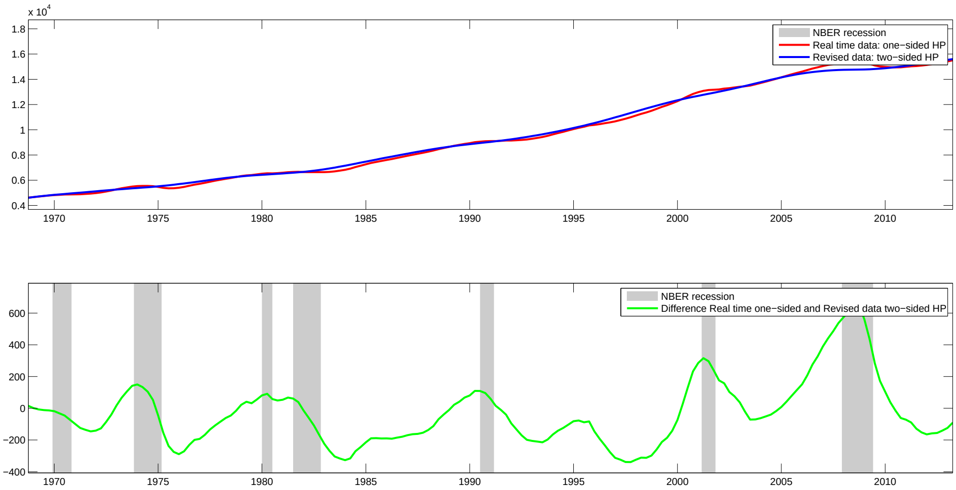

$$corr(\hat{\eta}_t, \eta_t) = 1$$

The red line represents the trend GDP that is obtained via a Kalman filter with the structure of a HP filter using real-time data. The blue line represents the trend GDP that is obtained via a Kalman smoother with the structure of a HP filter using revised data. The green line represents the difference between the two estimated trends, which is a proxy for the forecast error.

5.1.2 Results - estimated forecast error

Figures 20, 21 and 22 show the impulse responses of noise, supply and demand shocks. Overall, the results found before are confirmed by this alternative approach. In the case of the noise shock, the impact on the forecast error is also positive as expected, but more persistent than with the approach based on available data. By contrast, the impacts on output and consumption are significant only for 5-10 quarters. The other variables react broadly similarly to the base case. Also, the impuse responses corresponding to supply and demand shocks are very similar to those shown previously, confirming the robustness of the results found earlier. The FEVD is also similar (see Figures 23), the only notable difference is the larger contribution of the noise shock to changes in the forecast errors.

Figure 20: Large-scale VAR - IRFs to a noise shock (estimated forecast errors)

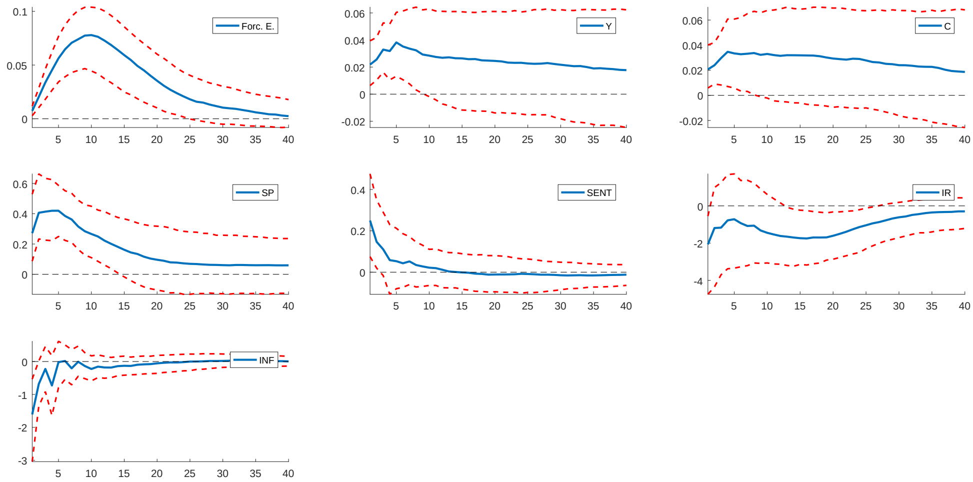

The graph presents the impulse responses to a noise shock of the estimated forecast error, output, consumption, investment (first row), stock prices, sentiment, interest rates and inflation (second row).The blue line is the point estimate and the red lines are 90% error bands.

Figure 21: Large-scale VAR - IRFs to a supply shock (estimated forecast errors)

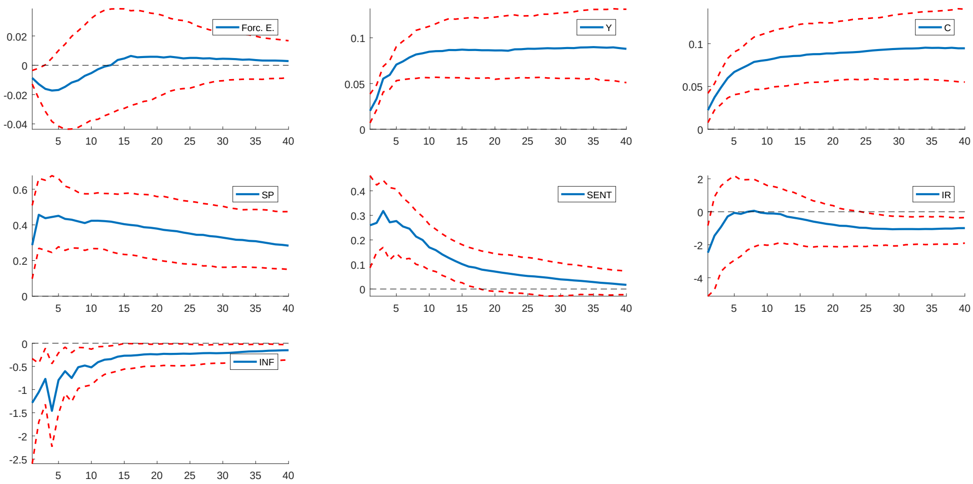

The graph presents the impulse responses to a supply shock of the estimated forecast error, output, consumption, investment (first row), stock prices, sentiment, interest rates and inflation (second row). The blue line is the point estimate and the red lines are 90% error bands.

Figure 22: Large-scale VAR - IRFs to a transitory demand shock (estimated forecast errors)

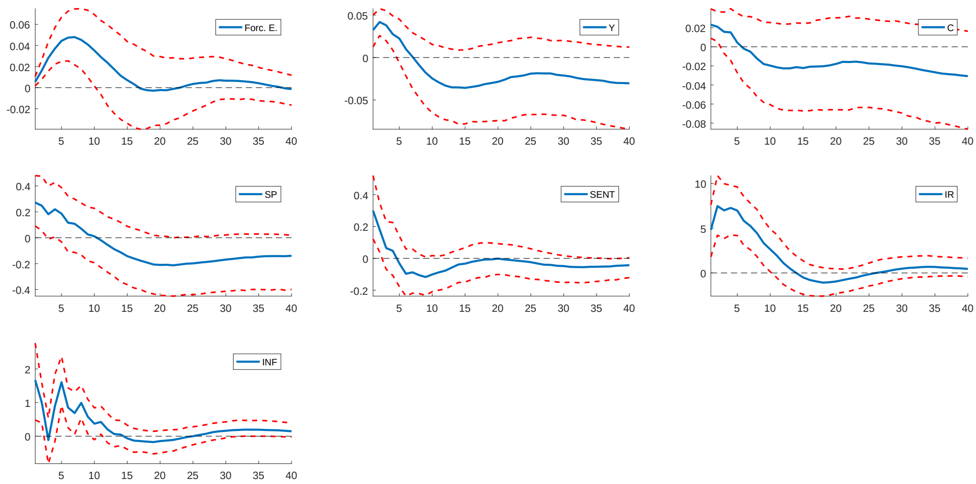

The graph presents the impulse responses to a transitory demand shock of the estimated forecast error, output, consumption, investment (first row), stock prices, sentiment, interest rates and inflation (second row). The blue line is the point estimate and the red lines are 90% error bands.

Figure 23: Large-scale VAR - FEVD (estimated forecast errors)

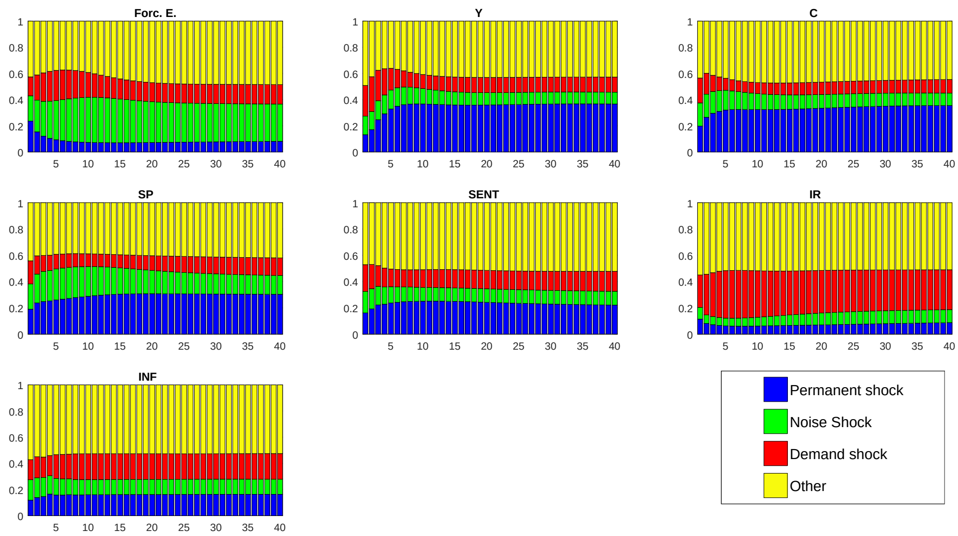

The graph presents the forecast error decomposition of estimated forecast error, output, consumption (first row), investment, stock prices, sentiment (second row), interest rates and inflation (third row). The blue area corresponds to the median contribution of supply shocks, the green area to the median contribution of noise shocks, the red area to the median contribution of transitory demand shocks and the yellow area to the contribution of all other non-identified shocks.

5.2 Alternative identification strategies - minimal restrictions

In the baseline identification scheme, we have used some redundant restrictions that were not needed to separate noise, supply and demand shocks. In this subsection we only use a minimal set of restrictions that are needed to identify all three shocks. Exact restrictions that we use are in Table 6:

Table 6: Imposed sign restrictions

	FE	Y	C	SP	IR	INF	SENT
Supply shock	-		+			-	
Noise shock	+		+			-	
Demand shock	+		+			+	

Figures 24, 25 and 26 show the IRFs to the permanent, noise and demand shocks in this new system. The main conclusion is that the responses are relatively similar to those obtained with the baseline restrictions. As it can be seen from the FEVD shown in Figure 27, the main difference is that the share of fluctuations explained by the three shocks falls, but noise shocks nevertheless still explain a relatively high share of output fluctuations. The main conclusions derived earlier therefore are not too dependent on additional restrictions used in the baseline.

Figure 24: Large-scale VAR - IRFs to a noise shock (extended identification - minimal restrictions)

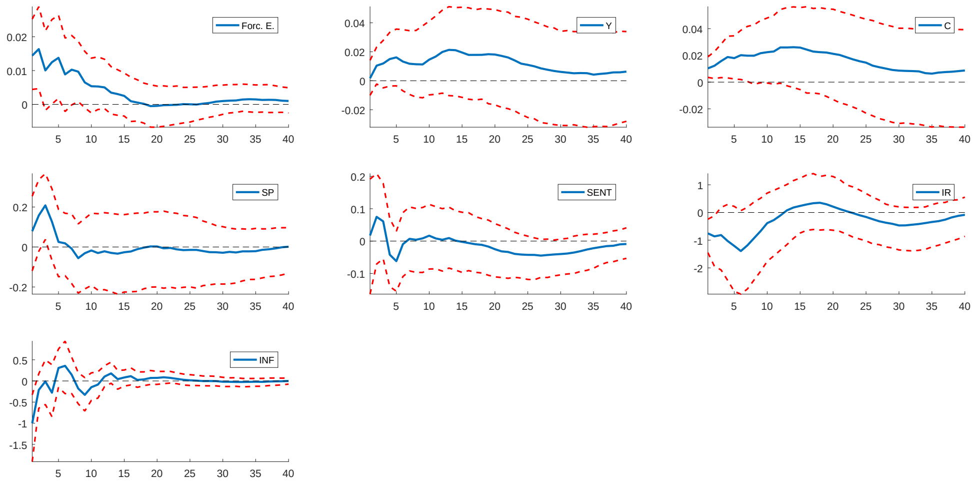

The graph presents the impulse responses to a noise shock of the estimated forecast error, output, consumption, investment (first row), stock prices, sentiment, interest rates and inflation (second row). The blue line is the point estimate and the red lines are 90% error bands.

Figure 25: US data - IRFs to a supply shock (extended identification - minimal restrictions)

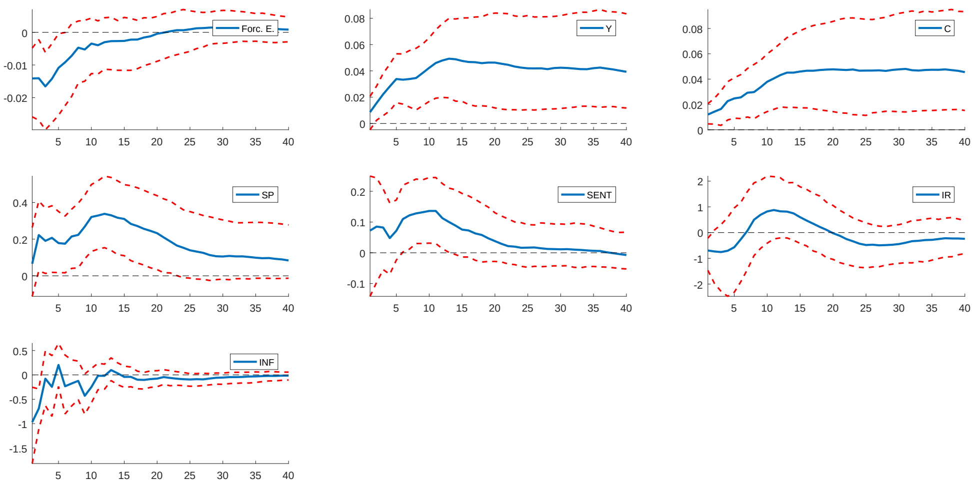

The graph presents the impulse responses to a supply shock of the estimated forecast error, output, consumption, investment (first row), stock prices, sentiment, interest rates and inflation (second row). The blue line is the point estimate and the red lines are 90% error bands.

Figure 26: US data - IRFs to a transitory demand shock (extended identification - minimal restrictions)

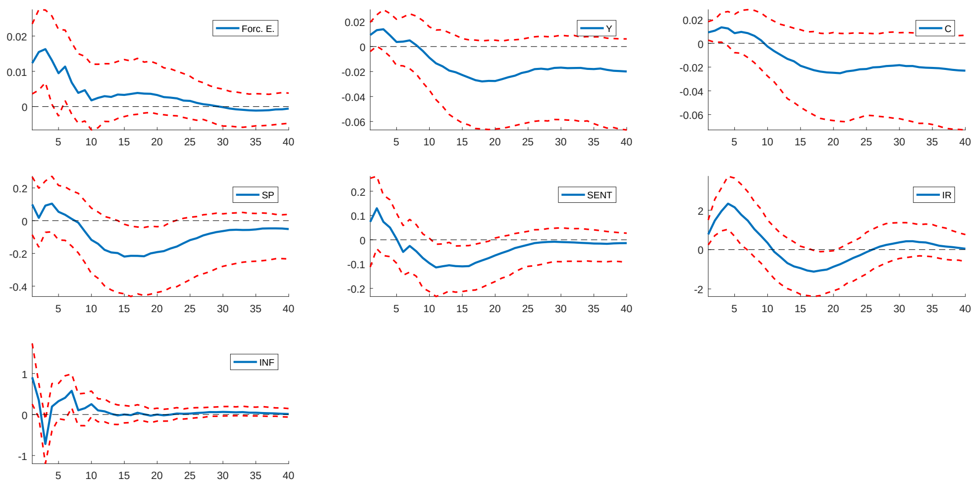

The graph presents the impulse responses to a transitory demand shock of the estimated forecast error, output, consumption, investment (first row), stock prices, sentiment, interest rates and inflation (second row). The blue line is the point estimate and the red lines are 90% error bands.

Figure 27: US data - FEVD (extended identification - minimal restrictions)

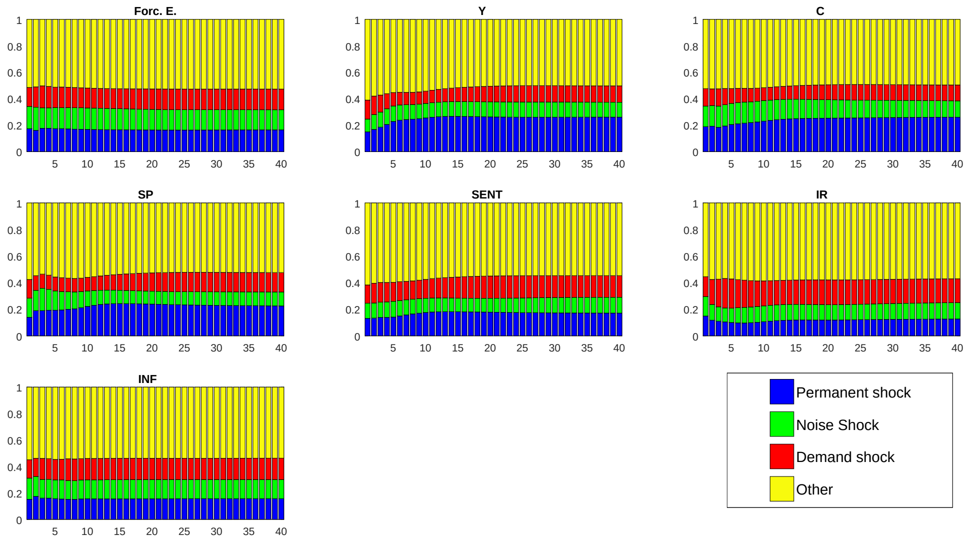

The graph presents the forecast error decomposition of estimated forecast error, output, consumption (first row), investment, stock prices, sentiment (second row), interest rates and inflation (third row). The blue area corresponds to the median contribution of supply shocks, the green area to the median contribution of noise shocks, the red area to the median contribution of transitory demand shocks and the yellow area to the contribution of all other non-identified shocks.

5.3 Comparison with the current literature

Our results correspond to a large extent with those found in the related literature. In particular, most papers that study the effects of noise versus supply shocks find that noise shocks could be interpreted as demand shocks and contribute to a large extent to economic fluctuations at the business cycle frequency. Noise shocks are therefore more important in explaining business cycles compared to permanent or technology shocks.

Blanchard et al. (2013) estimate a structural model and find that noise shocks account for more than 50% of short-run volatility of consumption, while permanent technology shocks play a smaller role, having almost no effect on quarterly volatility and explaining less than 30% at a 4-quarter horizon. Similar to this paper, they find spells of positive permanent shocks in the first half of 1980s and second half of 1990s. They estimate a structural model, thus putting structure on the data that can have important effects on the final results. By contrast, we opted for a more parsimonious SVAR model, which suggests why the smoothed series of noise shocks are more informative in our case - we obtain a clearer pattern of positive noise shocks before recessions and negative noise shocks at the start and during recessions. Nevertheless, they also find a succession of negative noise shocks around the recession in the early 1990s, and a spell of positive noise shocks before the 2001 recession.[13]

Results in Forni et al. (2013) are similar to those found in our paper. The responses of output and consumption have similar shapes. In the case of the noise shock, the responses are hump-shaped with a relatively small, although significant, impact effect; they reach a maximum after about two years, then decline towards zero after about five years. As predicted by the model, noise shocks spur a wave of private consumption and investment which vanishes once economic agents realize that the signal was just noise. They find that the responses to real shocks are permanent, although real shocks sometimes do not have significantly permanent effects as in our paper.

Enders et al. (2013) use a similar empirical strategy to our paper, but they construct forecast errors only by exploring the fact that real-time data is less precise than revised data. Nevertheless, they also find that nowcast errors react as predicted by theory, as we found in the exercise in Section 4.2.2. Given that they only use the difference between real-time and revised data to improve the estimates of the state, the distinction between productivity shocks as permanent and optimism shocks as transitory is not as evident as in this and other studies.

Barsky and Sims (2012) found the most contrasting results compared to this and other papers in the literature. They find that "animal spirits" effects are very weak and thus account for essentially none of the relationship between confidence and future consumption or income. The reason for such contrasting results is mostly related to the different methodologies - they estimate a structural model by matching theoretical impulse responses with empirical impulse responses. As discussed above, estimating a structural model imposes a structure on the data that may have important effects on the empirical results. While they find noise shocks to be unimportant, they find that permanent shocks are important drivers of business cycles, especially at longer horizons.

Finally, Angeletos et al. (2014) provides a model with only two sources of volatility: the usual technology shock and a confidence shock that drives the agents' beliefs about the state of the economy. The confidence shock therefore has a similar interpretation to our noise shock. Interestingly, although they use a different model, they also find that confidence shocks can account for about half of GDP volatility at business-cycle frequencies and that their effects look similar to a standard demand shock.

[13]They do not find similar patterns before the 2007 recession, as they only use data until 2008.

6 Conclusion

This paper has presented a model in which consumers receive noisy signals about future economic fundamentals. In this model, business cycle fluctuations can be driven both by news shocks (supply or technology shocks) and noise shocks ("animal spirits" shocks). We have shown that standard structural VAR models cannot be applied in principle to this model to identify the two types of shocks, as the VAR model faces invertibility issues. In other words, if consumers cannot distinguish between the two shocks, the econometrician also faces the same problem. However, by considering that the econometrician can potentially have a richer and more accurate information set, we have shown that a standard SVAR model can recover both supply and noise shocks. Richer information sets relate to the fact that the econometrician has access to revised data (while consumers take decisions with real-time data) and can include "future" data when estimating the state of the economy (while consumers take decisions with past and current data only).

In the empirical exercise, we have shown that the identified shocks have macroeconomic impacts that are in line with theoretical predictions. We have also shown that noise shocks explain a significant share of business cycle fluctuations in the short term, while supply shocks explain most of output variations in the long term. We have also shown that technology shocks turn negative a few years before recessions, while noise shocks are very positive at the cycle peaks and remain negative for some time during recovery phases. The recovery from recessions is mostly led by technology shocks, following Schumpeterian creative-destructive dynamics. These results are robust to the size of the VAR model and to the identification scheme. Finally, we have shown that consumption is even more affected by noise shocks than output at business cycle frequencies, showing the extent to which false perceptions about the future fundamentals play a role in consumers' behaviours. By contrast, interest rates do not seem to be strongly affected by noise shocks, as central banks may be more immune to animal spirits.

References

Angeletos, G.-M., Collard, F., and Dellas, H. (2014). Quantifying Confidence. NBER Working Papers 20807, National Bureau of Economic Research, Inc.

Bacchetta, P., Tille, C., and van Wincoop, E. (2012). Self-fulfilling risk panics. American Economic Review, 102(7):3674–3700.

Barsky, R. B. and Sims, E. R. (2012). Information, animal spirits, and the meaning of innovations in consumer confidence. American Economic Review, 102(4):1343–77.

Beaudry, P. and Portier, F. (2006). Stock prices, news, and economic fluctuations. The American Economic Review, 96(4):1293–1307.

Blanchard, O. J., L'Huillier, J.-P., and Lorenzoni, G. (2013). News, noise, and fluctuations: An empirical exploration. American Economic Review, 103(7):3045–70.

Blanchard, O. J. and Quah, D. (1989). The Dynamic Effects of Aggregate Demand and Supply Disturbances. American Economic Review, 79(4):655–673.

Canova, F. and Nicoló, G. (2002). Monetary disturbances matter for business fluctuations in the g-7. Journal of Monetary Economics, 49(6):1131–1159.

Diebold, F. X. and Rudebusch, G. D. (1991). Forecasting output with the composite leading index: A real-time analysis. Journal of the American Statistical Association, 86(415):603–610.

ECB, E. C. B. (2013). Confidence indicators and economic developments. Monthly Bulletin, 45(58).

Enders, Z., Kleemann, M., and Müller, G. (2013). Growth expectations, undue optimism, and short-run fluctuations. CESifo Working Paper Series 4548, CESifo Group Munich.

Farmer, R. E. (2012). The stock market crash of 2008 caused the great recession: Theory and evidence. Journal of Economic Dynamics and Control, 36(5):693–707.

Forni, M., Gambetti, L., Lippi, M., and Sala, L. (2013). Noisy news in business cycle. Technical report, mimeo.

Jaimovich, N. and Rebelo, S. (2009). Can news about the future drive the business cycle? American Economic Review, 99(4):1097–1118.

Keynes, J. M. (1936). General theory of employment, interest and money. Palgrave Macmillian.

Lorenzoni, G. (2009). A theory of demand shocks. American Economic Review, 99(5):2050–84.

Pigou, A. C. (1929). Industrial fluctuations. Macmillan.

Simon, D. (2006). Optimal state estimation: Kalman, H infinity, and nonlinear approaches. Wiley. com.

Uhlig, H. (2005). What are the effects of monetary policy on output? results from an agnostic identification procedure. Journal of Monetary Economics, 52(2):381–419.

A Appendix - Singularity of VAR with leads

To see the problem related to the singularity of VAR models when the future values of signals are used, we define an estimate of the state by a Kalman filter as:

$$\begin{aligned} \hat{x}_t^t &= \hat{x}_{t-1}^t + \mathcal{K}\left(s_t - A\hat{x}_{t-1}^t\right) \\ &= (1-\mathcal{K}A)\hat{x}_{t-1}^t + \mathcal{K}s_t \end{aligned} \tag{A.1}$$

where $\hat{x}_t^t$ is a state estimate in period t using signals up to period t - an estimate derived with a standard Kalman filter. $\mathcal{K}$ is the Kalman gain and s_t is the signal in period t. Similarly, the state estimate in period t using signals up to period $t+1$ can be written as:

$$\begin{aligned} \hat{x}_t^{t+1} &= \hat{x}_{t-1}^{t+1} + \mathcal{K}^0\left(s_t - A^0\hat{x}_{t-1}^{t+1}\right) + \mathcal{K}^1\left(s_{t+1} - A^1\hat{x}_{t-1}^{t+1}\right) \\ &= (1-\mathcal{K}^0A^0 - \mathcal{K}^1A^1)\hat{x}_{t-1}^{t+1} + \mathcal{K}^0 s_t + \mathcal{K}^1 s_{t+1} \end{aligned} \tag{A.2}$$

where $\hat{x}_t^{t+1}$ is a state estimate in period t using signals up to period $t+1$. $\mathcal{K}^0$ is the Kalman gain related to the signal in period t and $\mathcal{K}^1$ is the Kalman gain related to the signal in period $t+1$.

Consider now that the state estimate, $\hat{x}_t^t$ follows an autoregressive process:[14]

$$\hat{x}_t^t = \beta_1\hat{x}_{t-1}^t + u_t \tag{A.3}$$

where u_t is the error term. Comparing equation (A.1) with equation (A.3), we can see that s_t and u_t span the same linear space.

Shifting by one period and rearranging equations (A.1) and (A.2) we have:

$$\begin{aligned} \hat{x}_{t-1}^t - (1-\mathcal{K}A)\hat{x}_{t-2}^t &= \mathcal{K}s_{t-1} \\ \hat{x}_{t-1}^{t+1} - (1-\mathcal{K}^0A^0-\mathcal{K}^1A^1)\hat{x}_{t-2}^{t+1} &= \mathcal{K}^0 s_{t-1} + \mathcal{K}^1 s_t \end{aligned} \tag{A.4}$$

where the first row follows from equation (A.1) and the second row from equation (A.2). From equation (A.4), we can see that s_t can be expressed as a linear combination of $\hat{x}_{t-1}^t$, $\hat{x}_{t-2}^t$, $\hat{x}_{t-1}^{t+1}$ and $\hat{x}_{t-2}^{t+1}$. As s_t and u_t span the same linear space, u_t can also be expressed as a linear combination of the latter four variables, which implies a regression of the form:

$$\hat{x}_t^t = \beta_1\hat{x}_{t-1}^t + \beta_2\hat{x}_{t-2}^t + \beta_3\hat{x}_{t-1}^{t+1} + \beta_4\hat{x}_{t-2}^{t+1} + u_t \tag{A.5}$$

[14]The results in this section are presented only for auto-regressions with the state estimate. However, the forecast error, η_t, and consumption, c_t, are a linear function of the state estimate $\hat{x}_t^i$ and therefore results also hold for auto-regressions with the forecast error and consumption.

is characterized by a perfect linear relation between the independent and dependent variables, and the corresponding VAR is therefore singular. A similar reasoning applies also to regressions involving state estimates $\hat{x}_t^{t+2}$, $\hat{x}_t^{t+3}$,... that are constructed by using more leads of the signal. The only difference is that the more leads of the signal are used to construct state estimates, the more lags are needed to achieve a perfect linear correlation.

The improvements in the signal - lower variance of noise shocks - can reduce the problem of singularity. Whenever improvements in the signals are perfectly correlated with the noise shocks itself - newly defined noise shocks with lower variance can be written as $\upsilon_t^* = a\upsilon_t$, where υ_t^* are noise shocks with lower variance, a is a constant and υ_t are old shocks - the reduced variance of the signals does not alter the singularity problem. Namely, s_t can still be expressed as a (different) linear combination of $\hat{x}_{t-1}^t$, $\hat{x}_{t-2}^t$, $\hat{x}_{t-1}^{t+1}$ and $\hat{x}_{t-2}^{t+1}$.

On the other hand, when correlation is not perfect - noise shocks with lower variance cannot be written as $\upsilon_t^* = a\upsilon_t$ - the singularity problem does not appear. To see this, we can write equation (A.4) as:

$$\begin{aligned} \hat{x}_{t-1}^t - (1 - \mathcal{K}A)\hat{x}_{t-2}^t &= \mathcal{K}s_{t-1} \\ \hat{x}_{t-1}^{t+1} - (1 - \mathcal{K}^0A^0 - \mathcal{K}^1A^1)\hat{x}_{t-2}^{t+1} &= \mathcal{K}^0 s_{t-1}^* + \mathcal{K}^1 s_t^* \end{aligned} \tag{A.6}$$

where we use the fact that we have access to a different signal in the future, $s_t^* \neq s_t$ (second row). The new signal cannot be written as $s_t^* = cs_t$, where c is some constant. Therefore, it is not possible to form a perfect linear relation between s_t and $\hat{x}_{t-1}^t$, $\hat{x}_{t-2}^t$, $\hat{x}_{t-1}^{t+1}$ and $\hat{x}_{t-2}^{t+1}$.

B Appendix - Estimation algorithm

The estimation procedure is based on a modified version of the sign restriction approach presented in Uhlig (2005). The modification is that we restrict responses of most variables only by non-negativity constraints - in the spirit of sign restrictions approach proposed in Canova and Nicoló (2002).

The estimation procedure consists of three steps. In the first step, we estimate the reduced form VAR model. In the second step, we identify the structural shocks and take into account identification uncertainty. The third step serves to take into account estimation uncertainty. The steps are:

1. **Estimate reduced-form VAR:** Given the number of chosen lags, $\widehat{p}$, $VAR(\widehat{p})$ is estimated assuming a non-informative Normal-Wishart prior, see Uhlig (2005) for details. A draw from posterior gives as an estimate of autoregressive coefficients and the variance-covariance of reduced form errors, $\widehat{\Sigma_u}$.

2. **Identification restrictions:** The non-structural impulse responses function, $C(L)$, is related to the structural impulse responses function as $B(L) = A_0C(L)$ and reduced form errors, u_t, are related to structural errors as $u_t = A_0^{-1}B\varepsilon_t$. Impact matrix, $S = A_0^{-1}B$, must satisfy:
$$\Sigma_u = SS' \tag{B.1}$$

   The first estimate of impact matrix, $\hat{S}$, is obtained by a Cholesky decomposition of the variance-covariance matrix of reduced form errors, $\hat{S} = chol(\hat{\Sigma_u})$. The full set of permissible impact matrices can be construct as, $S* = \hat{S}Q$, where Q is an orthonormal matrix such that, $QQ' = I$.

   Define $b_{i,j}(k)$ to be a response of variable i to shock j in period k that follows from the structural lag-polynomial $B(L)B$. Define the function f on the real line per $f(x) = x$ if $x \geq 0$ and $f(x) = 100x$ if $x \leq 0$. Let s_j be the standard error of variable j. Let $J_{S,+}$ be the index set of variables for which identification restricts response to be positive and let $J_{S,-}$ be the index set of variables for which identification restricts response to be negative. Define the penalty function as:
$$\Psi(S) = \sum_{i}^{I} \sum_{j \in J_{S,+}} \sum_{k=0}^{P} f\left(\frac{-b_{i,j}(k)}{s_j}\right) + \sum_{i}^{I} \sum_{j \in J_{S,-}} \sum_{k=0}^{P} f\left(\frac{-b_{i,j}(k)}{s_j}\right) \tag{B.2}$$

   where P is the horizon over which restrictions should hold and I is a set of variables over which we impose identifying restrictions. Let $C_{S,+}$ be the index set of variables for which identification restricts response to be non-negative and let $C_{S,-}$ be the index set of variables for which identification restricts response to be non-positive. Define the non-negativity constraints vector as $c(S) = [\frac{b_{i,j}(k)}{s_j}^1, ..., \frac{b_{i,j}(k)}{s_j}^P]$. To identify the model, we solve the following minimization problem:
$$\begin{aligned} S \quad = \quad & \underset{S}{argmin} \quad & \Psi(S) \\ & subject\ to \quad & c(S) \geq 0 \\ & & SS' = \Sigma_u \end{aligned} \tag{B.3}$$

   In the current paper only restrictions on forecast error enter the penalty function, while sign restrictions on other variables are applied only as non-negativity constraints. In this way we achieved that most of the forecast variance of forecast error is explained by the two identified shocks, while at the same time achieving that identification is not too restrictive - our methodology restricts only the sign of responses and does not maximize the magnitude - in relation to responses of other variables.

3. **Distribution** We repeat steps 1-2 1000 times. The IRF's point estimates and the related confidence bands are constructed by retaining the relevant percentiles of a distribution of retained IRFs. The same procedure is used to construct FEVD's point estimates and the related confidence bands.