

Dées, Stéphane; Zimic, Srečko

Working Paper

Animal spirits, fundamental factors and business cycle fluctuations

ECB Working Paper, No. 1953

Provided in Cooperation with:

European Central Bank (ECB)

Suggested Citation: Dées, Stéphane; Zimic, Srečko (2016) : Animal spirits, fundamental factors and business cycle fluctuations, ECB Working Paper, No. 1953, ISBN 978-92-899-2201-2, European Central Bank (ECB), Frankfurt a. M., <https://doi.org/10.2866/112389>

This Version is available at:

<https://hdl.handle.net/10419/154386>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.




EUROPEAN CENTRAL BANK
EUROSYSTEM

Working Paper Series

Stephane Dées, Srečko Zimic

Animal spirits, fundamental factors
and business cycle fluctuations

No 1953 / August 2016



Note: This Working Paper should not be reported as representing the views of the European Central Bank (ECB). The views expressed are those of the authors and do not necessarily reflect those of the ECB

Abstract

This paper explores empirically the role of noisy information in cyclical developments and aims at separating fluctuations that are due to genuine changes in fundamentals from those due to temporary animal spirits or expectational errors (noise shocks). Exploiting the fact that the econometrician has a richer data-set in some dimensions than the consumers, we use a novel identification scheme in a structural vector-autoregressive (SVAR) framework. Our results show that noise shocks are more important for business cycle fluctuations than permanent (or technology) shocks. We also show that technology shocks turn negative a few years before recessions, while noise shocks are very positive at the cycle peaks. By contrast, the recovery from recessions is mostly led by technology shocks, noise shocks remaining negative for some time during this business cycle phase.

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

JEL Classification: C32, E32

Non-technical summary

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. The way households and firms form their 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 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. Although the Real Business Cycle theory does not envisage 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. In these models, although technology remains the only determinant of output in the long term, news about future fundamentals can imply a change in expectations, which affects agents' behaviours in the short term in anticipation of any fundamental change. However, economic agents receive only noisy signals about future technology, leading to expectational errors. If the information appears subsequently to be wrong (i.e. just a noise), the agents adjust their expectations and decisions accordingly. Conversely, if the information is really news, the economy adjusts gradually to the level of activity consistent with technology. These changes in expectations generate economic fluctuations, both in the short and in the long term.

This paper explores empirically the role of noisy information in cyclical developments and aims at separating fluctuations that are due to genuine changes in fundamentals (news shocks) from those due to temporary expectational errors (noise shocks). Theoretically, our approach is based on a model, where business cycle fluctuations can be driven both by news shocks (or technology shocks) and noise shocks (animal spirits shocks). The main issue we have to deal with is related to the fact that, empirically, 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 should also face the same problem. Although this point is valid in real time, the econometrician can however potentially have access to a richer dataset in hindsight.

Based on this idea, this paper claims that a structural VAR model can therefore 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 (i.e. by knowing the data for the whole sample, the econometrician has a better estimation of the technological trends than the economic agents). Second, economic agents only observe real-time data, while the econometrician has also access to revised data. By exploiting the fact that the econometricians have access to future and revised data, our methodology relies on the forecast errors consumers do when predicting the trend of GDP, which is our measure of technology. More precisely, consumers only use past data and determine trend GDP with a one-sided filter, while the econometrician also uses future data when filtering the GDP series with a two-sided filter. Moreover, consumers use real-time data, while the econometrician, having access to revised data, uses more accurate information about the current state of technology. As a result, the econometrician can in hindsight estimate the forecast error made by the consumers. This forecast error estimate can be decomposed into an error due to the use of future signals and an error due to the use of less noisy signals.

In the empirical exercise, we estimate a VAR model that includes the estimated forecast errors together with GDP, private consumption, investment, stock prices, interest rates, inflation and consumer sentiment. The estimation is conducted using US data over a period from Q1 1970 to Q2 2012. The identification of the noise and technology shocks is achieved by sign restrictions, as our theoretical model shows that the forecast error is positive for positive noise shocks (i.e. consumers are too optimistic), while it is negative for positive permanent shocks (i.e. consumers are too pessimistic).

We show empirically that the identified shocks have macroeconomic impacts that are in line with theoretical predictions. A permanent (technology) 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 more important for business cycle fluctuations. We show that noise shocks explain almost half of business cycle fluctuations

in the short term, while technology shocks explain only up to 20 percent of output variations at business cycle frequencies. We also show 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 the recovery phases. The recovery from recessions is mostly led by technology shocks, following a Schumpeterian creative-destructive dynamics. These results are robust to the size of the VAR model and to the identification scheme.

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 fundamentals (news shocks) from those due to temporary expectational errors (noise shocks). They show that identification of the shock 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

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

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 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 news 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 permanent and noise shocks have effects as predicted by theory. A permanent (technology) 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 more

important for business cycle fluctuations - depending on the model specification, noise shocks explain 30-50 percent of output variations at business cycle frequencies. On the other hand, permanent shocks drive the economy in the long run, but only account for around 20 percent of output variations at business cycle frequencies.

After a presentation of the theoretical model in Section 2, in Section 3 we explain 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 technology and noise shocks. In Section 4, we present empirical evidence on the effects of news and noise shocks, by applying our methodology to US data. We conclude in Section 5.

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:

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

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:

$$y_t = c_t \quad (2.2)$$

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:

$$\lim_{t \rightarrow \infty} E_t[c_{t+j} - a_{t+j}] = 0 \quad (2.3)$$

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:

$$c_t = \lim_{t \rightarrow \infty} E_t[a_{t+j} | \mathcal{I}_t] \quad (2.4)$$

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:

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

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 :

$$s_t = a_t + v_t \quad (2.6)$$

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

- State Equation

$$\begin{bmatrix} a_{t|t} \\ a_{t-1|t} \end{bmatrix} = A \begin{bmatrix} a_{t-1|t-1} \\ a_{t-2|t-1} \end{bmatrix} + B \begin{bmatrix} \epsilon_t \\ v_t \end{bmatrix} \quad (2.7)$$

- Observation Equation

$$s_t = C \begin{bmatrix} a_{t|t} \\ a_{t-1|t} \end{bmatrix} + D \begin{bmatrix} \epsilon_t \\ v_t \end{bmatrix} \quad (2.8)$$

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.

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:

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

²Note 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.

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:

$$y_t = c_t \quad (2.10)$$

$$u_t = a_t - y_t \quad (2.11)$$

Figure 2 shows the response of the model's variables to the technology shock, ϵ_t . The standard deviation of the shock is 0.7 percent. From equation (2.6), in the absence of noise shocks, the signal, s_t , is equal to a_t . 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 technology shock.

Figure 3 shows the responses to a noise shock. The 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 technology 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 the technology and noise shocks. Since consumers only receive a noisy signal about the shocks, it is not possible to recover technology 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:

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

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:

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

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

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

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

$$\frac{\partial s_t}{\partial \epsilon_t} = \frac{\partial s_t}{\partial v_t} \quad (3.4)$$

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:

$$\frac{\partial c_t}{\partial \epsilon_t} = \frac{\partial c_t}{\partial v_t} \quad (3.5)$$

In other words, on impact consumers respond in the same way to technology 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:

$$\frac{\partial x_t}{\partial \epsilon_t} = \frac{\partial x_t}{\partial v_t} \quad (3.6)$$

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 Technology 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 4, 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

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

on impact. Similarly, we expect that the identified transitory shock would be a mixture of the measurement error and the theoretical noise shock. Figure 4 shows that, unlike the structural shock, consumption does not respond at all to the identified transitory shock.

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

$$\frac{\partial s_{t+j}}{\partial \epsilon_t} \neq \frac{\partial s_{t+j}}{\partial v_t} \quad \forall j = 1, 2, \dots \rightarrow \frac{\partial c_{t+j}}{\partial \epsilon_t} \neq \frac{\partial c_{t+j}}{\partial v_t} \quad \forall j = 1, 2, \dots \quad (3.7)$$

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:

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

Technology responds differently to the two shocks:

$$1 = \frac{\partial a_t}{\partial \epsilon_t} \neq \frac{\partial a_t}{\partial v_t} = 0 \quad (3.9)$$

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

$$\frac{\partial \eta_t}{\partial \epsilon_t} \neq \frac{\partial \eta_t}{\partial v_t} \quad (3.10)$$

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, as is shown by the simulations presented in Figure 7.

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:

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

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: $var[v_t|\mathcal{I}_t^e] < var[v_t|\mathcal{I}_t]$
- 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:

$$\lim_{\sigma_v^2 \rightarrow 0} E[E_t[a_t|\mathcal{I}_t^e] - a_t] = 0 \quad (3.12)$$

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:

$$\sigma_v^2 < \sigma_v^{2*} \rightarrow \text{var}[\eta_t | \sigma_v^2] < \text{var}[\eta_t | \sigma_v^{2*}] \quad (3.13)$$

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.

Figure 5 shows how the estimated forecast error evolves according to the quality of the signal. The black lines show the errors implied by the theoretical models after a permanent and a noise shock. When the signal becomes less noisy, the lines move closer to the theoretical impulse responses and the correlation between the estimated and the true forecast errors increase rapidly when the variance of the noise is reduced.⁴

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 $\mathcal{I}_t^{e,j} = \{s_{t+j}, s_{t+j-1}, s_{t+j-2}, \dots\}$, we can show:⁵

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

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

⁴How 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.

⁵See Simon [2006], page 271.

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:

$$\hat{\eta}_t = E_t[a_t|\mathcal{I}_t] - E_t[a_t|\mathcal{I}_t^{e,j}] \quad (3.15)$$

Figure 6 shows how close to the actual forecast errors the estimated forecast error can become when we increase the number of leads. For the permanent shock, adding four leads is almost sufficient to reproduce the shape of the true forecast error. By contrast, it is more difficult to estimate the forecast error following a noise shock. Nevertheless, four leads imply an impulse response that is qualitatively similar to what the theoretical model gives. Overall, the correlation between the estimated and the true forecast errors is equal to 0.72 with four leads.

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:

Table 1: Forecast error estimation

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

Using simulated data, Figure 10 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 11 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 12 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.

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

	Permanent 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 8 compares the theoretical and the estimated impulse responses to a technology 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

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 B. In the following exercises we assume that the correlation between the improvements in the signal and the noise shocks is zero.

distinguish between a technology and a noise shock. The response of consumption to these shocks is also similar to what the theoretical model predicts: after a technology 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 9 shows the results of the same exercise by comparing estimated (first row) and theoretical (second row) forecast error variance decompositions (FEVD). The contributions of technological 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 permanent shocks mainly explain the variance of consumption in the long-run.

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:

$$a_t = a_{t-1} + \varepsilon_{t-1} \quad (3.16)$$

$$s_t = a_t + v_t \quad (3.17)$$

where a_t is the observable technology, s_t is a signal about technology, ε_t is the technology shock and v_t is the noise shock. The signal extraction problem comes from the fact that the technology 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 $\frac{\partial s_t}{\partial \varepsilon_t} = \frac{\partial s_t}{\partial v_t}$, but $\frac{\partial a_t}{\partial \varepsilon_t} \neq \frac{\partial a_t}{\partial v_t}$ 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 technology 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 news and noise shocks in the US. We first explain how we construct the real-time data on GDP. We then present our methodology to estimate the forecast errors, before presenting the VAR model we use to identify noise and technology shocks. We then present the results, including impulse response analysis, forecast error variance decomposition and historical decomposition, which are consistent with the predictions of our theoretical model. Finally, we show that these results are robust to a different identification strategy.

4.1 Real-time data

The first step in our procedure is to obtain the real-time data on 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:

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

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: $y_{t-1}^* = y_{t-5} (1 + \Delta^4 y_{t-1}^*)$ 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 .

4.2 Estimating forecast errors

4.2.1 Methodology

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 trend GDP, or potential output, 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 1998.⁷ Moreover, in our opinion, firms and consumers focus more on trend GDP rather than technological trend when making economic decisions.

We further assume that consumers use a Kalman filter to construct the estimate of trend GDP.⁸ The Kalman filter is set so as to represent a one-sided HP filter, and therefore has some empirical appeal. On the other hand, the econometrician can use the complete data sample, so that the econometrician's estimate of trend GDP is obtained by using a Kalman smoother. There is also a difference in the nature of the data used. Consumers can only use real-time data, while the econometrician can use revised data from later vintages. The correspondance 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 one-sided filter. The econometrician also uses future data when filtering the GDP series with a two-sided filter.

⁷See <https://www.philadelphiafed.org/research-and-data/real-time-center/real-time-data/data-files/OPH/>

⁸More detail on the methodology is given in Appendix C.1.

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

The methodology outlined previously can be applied to actual data. The possible estimates of the forecast errors can then be decomposed as in Table 1.

Figure 13 shows our estimate of $\hat{\kappa}_t$ and Figure 14 our estimate of $\hat{\lambda}_t$ on US data. These figures allows us to see how the final estimate, shown in Figure 15, is decomposed into the component due to the use of future signals and the one due to the use of less noisy signals. Of course, as the “true” forecast error is not observable in the data, our best estimate of it is the final estimate, $\hat{\eta}_t$. The correlation between $\hat{\kappa}_t$ and $\hat{\eta}_t$ is 0.68 and the correlation between $\hat{\lambda}_t$ and $\hat{\eta}_t$ is 0.92. Figure 15 also shows 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.

4.3 VAR model

We can now estimate a VAR model that includes the following variables:

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

The estimated forecast errors are those defined in the previous subsection. The estimation period is from Q1 1970 to Q2 2012. We drop 4 observations at both ends of the sample to account for the lag/lead structure. The system is estimated in differences with three lags, and standard errors are calculated using classic bootstrap methods. The results are robust to changes from differences to levels and using a different number of lags.

Compared to the theoretical model, we enrich the VAR model by including additional variables. This allows us to better identify the shocks and see their impact on these additional variables, such as stock prices or consumer sentiment. A simple two-variable VAR model, including only the estimate of the forecast error and consumption, has also been estimated. The results, available upon request, are not noticeably different from those presented below.

4.4 Identification

The identification of the noise and technology shocks is achieved 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 permanent shocks (i.e. consumers are too pessimistic).

The other variables respond similarly to both shocks according to the signs summarized in the following table:

Table 2: Imposed sign restrictions

	FE	Y	C	I	SP	IR	INF	SENT
Permanent shock	-	+	+		+			+
Noise shock	+	+	+		+			+

The restrictions are imposed for one year, though changing this does not considerably change the results. The responses of interest rates and inflation are left unrestricted. The estimation methodology is described in Appendix in Section D.

Clearly, our VAR model implicitly includes other types of shocks, such as monetary or fiscal shocks. As discussed above, our theoretical model can incorporate other shocks so long as only the shocks considered, technology and noise shocks, are not observable contemporaneously, so that other types of shocks do not affect the forecast errors. The inclusion of other shocks will be discussed as a robustness check in Section 4.7.1.

4.5 Empirical results

We now present the estimated impulse response functions, forecast error variance decompositions and historical decompositions.

First, Figure 16 shows impulse response functions (IRFs) following a 1 standard deviation positive shock to technology. 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, consumption and investment are permanently higher. As also expected from the theoretical model, the forecast error is negative with a maximum impact after 5 quarters, before returning towards the baseline. Inflation and interest rates decline on impact but are not permanently affected. Interestingly, both equity prices and consumer sentiment increase permanently following the technology shock.

Figure 17 shows IRFs following a 1 standard deviation positive noise shock. In this case, the forecast error is positive and increases up to the 10th quarter after the shock. As expected by the theoretical model, the impact on output, investment and consumption is positive in the short term but the effects are short-lived and the variables return to baseline after 8 to 10 quarters. Equity prices and sentiment are also affected positively but, after 5 quarters, the effects are no longer significant. In line with the features of a demand shock, inflation increases in the short term and interest rates are also higher.

The forecast error variance decomposition (FEVD), presented in Figure 18, shows that the noise shock explains almost half of the business cycle fluctuations in the short term. As expected from the theory, the impact of the technology shock on the business cycle builds up gradually. It explains only up to 20 percent of output variations at business cycle frequencies and half of the fluctuations in the long term, when the noise shock contribution is much smaller.

Figure 19 shows the series of extracted shocks, with technology 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 technology shock declines just before recessions and sharply increases during the recessions, which could be interpreted as a Schumpeterian creative-destruction mechanism driving the recovery phase. Second, the noise shock tends to be very positive in boom periods, like in the dot-com bubble in the early 2000s

and in the housing market boom from the mid-2000s to the financial crisis. 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.

Finally, Figures 20 and 21 show the historical decomposition of output and the contribution of the two shocks to business cycles fluctuations over time. When the permanent technology shock is removed, US GDP is much lower, especially over the period 1985-2008. This means that our VAR model attributes much of the increase in GDP to positive technology shocks during that period. At the same time, the technology shocks do not contribute so much to short-term fluctuations. By contrast, the noise shock explains most of the output fluctuations at business cycle frequencies.

4.6 Robustness check - long run restrictions

As a robustness check, we use a different identification strategy to verify to what extent the previous results depend on the sign restrictions we imposed. We use long-run restrictions instead of sign restrictions to identify the technology and noise shocks.

In particular, the noise shock is identified by imposing that it has transitory effects on consumption.⁹ 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 :

$$\Xi = \begin{bmatrix} * & 0 \\ * & * \end{bmatrix} \quad (4.2)$$

Figure 22 shows the impulse responses of consumption and the forecast error to the permanent shock, in the upper panel, and the transitory shock, in the lower panel. 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 technology shock implies negative forecast errors in the short term, while a positive transitory shock increases the forecast error being in this case positive.

⁹Note that while many shocks other than “animal spirits” are compatible with such a restriction, the purpose of this robustness check is to distinguish between the two shocks included in our theoretical model.

4.7 Robustness check - identification

In this subsection we modify the identification strategy in two ways. Firstly, we explicitly identify a third shock that we label an aggregate demand shock. Secondly, we identify shocks by using an identification scheme that uses a minimal number of restrictions needed to identify the three shocks of interest.

4.7.1 Robustness check - additional shock

Our empirical results have been so far based on the identification of the two shocks defined in our theoretical model, distinguishing the noise and technology shocks by their impact on the forecast errors. In our VAR model, which includes more variables, all other disturbances have been grouped together as “other shocks”. In this subsection we are more explicit about other shocks and separately identify an additional shock in order 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.

Theory implies two types of restrictions that we can use to distinguish a noise shock from a standard demand shock. First, other shocks except permanent and noise shocks must not have any impact on the forecast errors contemporaneously, as they are supposed to be observable. This would allow us to use zero restrictions to separate permanent and noise shocks from standard demand shocks. However, such an identification scheme is only viable in case the estimated forecast errors are measured precisely and are not 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.¹⁰

Nevertheless a third shock can be identified without imposing the zero restriction. In theoretical model agents cannot distinguish between the two shocks in the first period implying that the responses of variables to the permanent and noise shocks should be identical in the first period. This allows us to use sign restrictions to separate permanent and noise shocks from standard

¹⁰We have 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.

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 3:

Table 3: Imposed sign restrictions

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

We keep baseline restrictions and we add standard restrictions to identify demand shocks that postulate that demand shocks lead to an increase in inflation and interest rates. By contrast, we assume that noise shocks and permanent shocks lead to a decrease in inflation and interest rates, given they are perceived as shocks to technology. Demand shocks are treated as transitory shocks and therefore we assume that estimated forecast error increases after the demand shock.¹¹

Figure 25 shows the IRFs of this newly identified shock. It has positive effects on output and consumption, although these effects are only significant in the short term. Similarly, the positive impact on stock prices and consumer sentiment are significant only in the first quarters following the shock. Figures 24 and 23 show the IRFs to the permanent shock and the noise shock in this new system. The responses are relatively similar to those presented earlier. Compared with the observable demand shock, the effect of the noise shock builds up more gradually. As shown by the FEVD, in Figure 26, the observable demand shock explains a relatively larger share of fluctuations in output, inflation and interest rates, while the noise shock explains relatively more of the fluctuations in stock prices and sentiment. The share of output fluctuations explained by noise shocks decreases in comparison to the baseline identification scheme, but nevertheless noise shocks still account for 20-30 percent of output fluctuations. The main conclusions derived earlier therefore still hold in this more comprehensive shock identification exercise.

4.7.2 Robustness check - minimal restrictions

In the baseline identification scheme we have used some redundant restrictions that were not needed to separate noise, permanent and demand shocks. In this subsection we only use a

¹¹In accordance with theory, sign restrictions are imposed only in the first period.

minimal set of restrictions that are needed to identify all three shocks. Exact restrictions that we use are in Table 4:

Table 4: Imposed sign restrictions

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

Figures 28, 27 and 29 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 30, 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.

4.8 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 permanent 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.¹²

Results in Forni et al. [2013] are similar to those found in our paper. The responses of output, consumption and investment 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 4.6. 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 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 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

¹²They do not find similar patterns before the 2007 recession, as they only use data until 2008.

similar to a standard demand shock.

5 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 (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 technology 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 almost half of business cycle fluctuations in the short term, while technology shocks explain only up to 20 percent of output variations at business cycle frequencies. 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.

References


George-Marios Angeletos, Fabrice Collard, and Harris Dellas. Quantifying Confidence. NBER Working Papers 20807, National Bureau of Economic Research, Inc, December 2014.

- Philippe Bacchetta, Cédric Tille, and Eric van Wincoop. Self-fulfilling risk panics. American Economic Review, 102(7):3674–3700, December 2012.
- Robert B. Barsky and Eric R. Sims. Information, animal spirits, and the meaning of innovations in consumer confidence. American Economic Review, 102(4):1343–77, June 2012.
- P. Beaudry and F. Portier. Stock prices, news, and economic fluctuations. The American Economic Review, 96(4):1293–1307, 2006.
- Olivier J. Blanchard, Jean-Paul L’Huillier, and Guido Lorenzoni. News, noise, and fluctuations: An empirical exploration. American Economic Review, 103(7):3045–70, December 2013.
- F. Canova and G.D. Nicoló. Monetary disturbances matter for business fluctuations in the g-7. Journal of Monetary Economics, 49(6):1131–1159, 2002.
- Francis X Diebold and Glenn D Rudebusch. Forecasting output with the composite leading index: A real-time analysis. Journal of the American Statistical Association, 86(415):603–610, 1991.
- European Central Bank ECB. Confidence indicators and economic developments. Monthly Bulletin, 45(58), January 2013.
- Zeno Enders, Michael Kleemann, and Gernot Müller. Growth expectations, undue optimism, and short-run fluctuations. CESifo Working Paper Series 4548, CESifo Group Munich, 2013.
- Roger E.A. Farmer. The stock market crash of 2008 caused the great recession: Theory and evidence. Journal of Economic Dynamics and Control, 36(5):693–707, 2012.
- Mario Forni, Luca Gambetti, Mario Lippi, and Luca Sala. Noisy news in business cycle. Technical report, mimeo, 2013.
- N. Jaimovich and S. Rebelo. Can news about the future drive the business cycle? American Economic Review, 99(4):1097–1118, September 2009.
- John Maynard Keynes. General theory of employment, interest and money. Palgrave Macmillan, 1936.

- Lutz Kilian. Small-Sample Confidence Intervals For Impulse Response Functions. The Review of Economics and Statistics, 80(2):218–230, May 1998.
- Guido Lorenzoni. A theory of demand shocks. American Economic Review, 99(5):2050–84, December 2009.
- Arthur Cecil Pigou. Industrial fluctuations. Macmillan, 1929.
- Dan Simon. Optimal state estimation: Kalman, H infinity, and nonlinear approaches. Wiley.com, 2006.
- James H. Stock and Mark W. Watson. Forecasting inflation. Journal of Monetary Economics, 44(2):293–335, October 1999.
- Harald Uhlig. What are the effects of monetary policy on output? results from an agnostic identification procedure. Journal of Monetary Economics, 52(2):381–419, 2005.

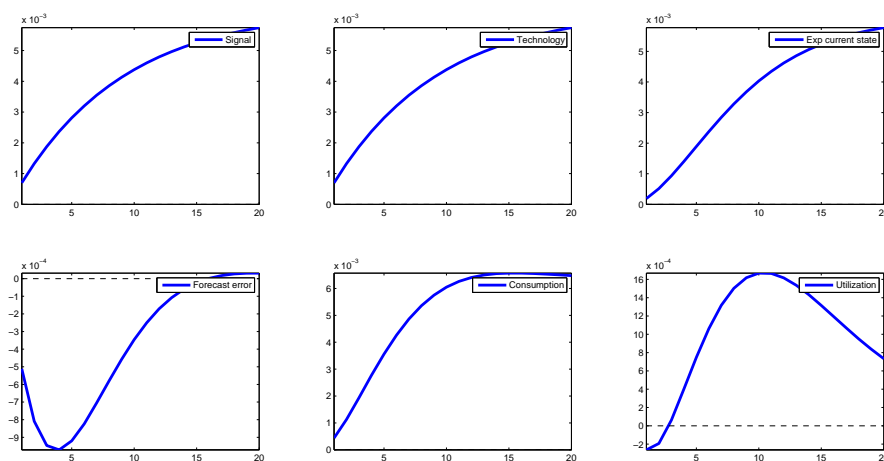
A Appendix - Model

Figure 1: Simulated Technology, Signal and Cons. expectations and Forecast error



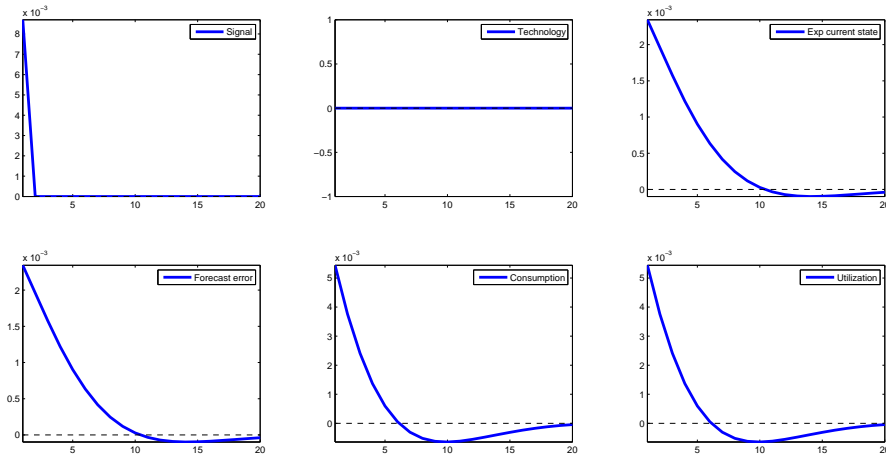
The black line represents the underlying technology process, which agents do not observe. The agents observe a signal (blue line) and via a Kalman filter, they form expectations about the technology (red line). The difference between their expectations and the true state of technology - forecast error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

Figure 2: IRFs to the Permanent shock




The graph presents the impulse responses of variables to the one-standard deviation permanent shock. The top row presents responses of signal, technology and agent's expectation. The bottom row presents responses of forecast error, consumption and utilization.

Figure 3: IRF to the Noise shock



The graph presents the impulse responses of variables to the one-standard deviation noise shock. The top row presents responses of signal, technology and agent's expectations. The bottom row presents responses of forecast error, consumption and utilization.

Figure 4: IRF of consumption



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

B 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:

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

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:

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

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:¹³

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

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

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

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

where the first row follows from equation (B.1) and the second row from equation (B.2). From equation (B.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

¹³The 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.

of the latter four variables, which implies a regression of the form:

$$\hat{x}_t^t = \beta_1 \hat{x}_{t-1}^t + \beta_2 \hat{x}_{t-2}^t + \beta_3 \hat{x}_{t-1}^{t+1} + \beta_4 \hat{x}_{t-2}^{t+1} + u_t \quad (\text{B.5})$$

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 $\hat{x}_t^{t+2}, \hat{x}_t^{t+3}, \dots$ 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 $v_t^* = av_t$, where v_t^* are noise shocks with lower variance, a is a constant and v_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 $v_t^* = av_t$ - the singularity problem does not appear. To see this, we can write equation (B.4) as:

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

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} .


Figure 5: Estimated Forecast error from less noisy signal



Correlation between Estimated and True Forecast errors

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

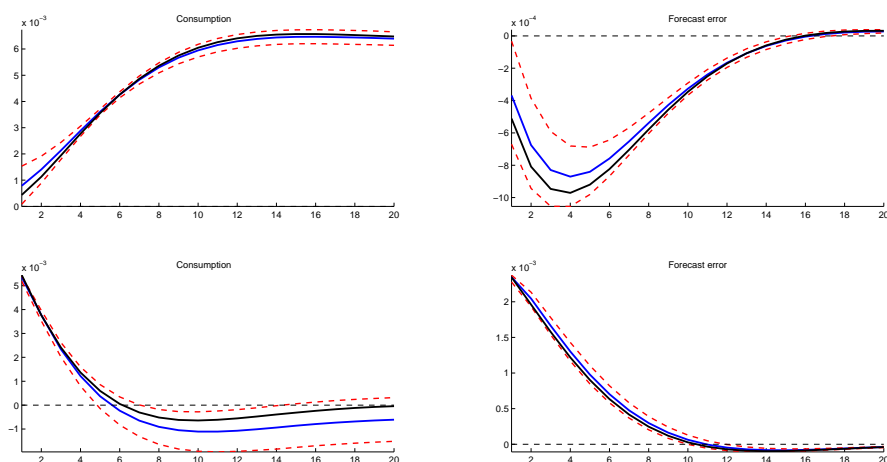
Figure 6: Estimated Forecast error from using future observations



Correlation between Estimated and True Forecast errors


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

Figure 7: Estimated IRF by using the true forecast errors
(simulated data)




The graph presents the estimated impulse responses of consumption (right) and the forecast error (left) to the permanent (above) and the noise (below) shock. The blue line is point estimate and the red lines are 90% error bands.

Figure 8: 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 permanent (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 9: 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 permanent 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.

C Appendix - Estimation of trend GDP

C.1 HP filter in a state space form

The HP-filter widely used in practice decomposes time series y_t , in our case GDP, into a trend component and a cyclical component, can be represented as:¹⁴

$$y_t = \tau_t + \epsilon_t \quad (\text{C.1})$$

$$(1 - L)^2 \tau_t = \eta_t \quad (\text{C.2})$$

where τ_t is a trend component, $\{\epsilon_t\}$ is a white noise sequence (cyclical component) and $\{\eta_t\}$ is a white noise sequence. The following process can be written as a state space:

1. State Equation

$$\begin{bmatrix} \tau_{t|t} \\ \tau_{t-1|t} \end{bmatrix} = \begin{bmatrix} 2 & -1 \\ 1 & 0 \end{bmatrix} \begin{bmatrix} \tau_{t-1|t-1} \\ \tau_{t-2|t-1} \end{bmatrix} + \begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix} \begin{bmatrix} \epsilon_t \\ 0 \end{bmatrix} \quad (\text{C.3})$$

2. Observation Equation


$$y_t = \begin{bmatrix} 1 & 0 \end{bmatrix} \begin{bmatrix} \tau_{t|t} \\ \tau_{t-1|t} \end{bmatrix} + \begin{bmatrix} 0 & 1 \end{bmatrix} \begin{bmatrix} 0 \\ \eta_t \end{bmatrix} \quad (\text{C.4})$$

where $\epsilon_t \sim \mathcal{N}(0, \sigma_\epsilon^2)$ and $\eta_t \sim \mathcal{N}(0, \sigma_\eta^2)$. The standard parameter lambda is obtained as, $\lambda = \frac{\sigma_\eta^2}{\sigma_\epsilon^2}$.

The one-sided HP filter is obtained by estimating the state space system by a Kalman filter, while non-casual estimate of the state-space system is obtained by a Kalman smoother. The alternative to obtain non-casual filter estimates of the trend and cyclical component is to use a standard two-sided HP filter.


¹⁴ See Stock and Watson [1999]

Figure 10: An estimate of forecast error due to the use of future signals - $\hat{\kappa}_t$
(simulated data)




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




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




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.

Figure 13: An estimate of forecast error due to the use of future signals - $\hat{\kappa}_t$
(US data)




The red line represents the trend GDP that is obtained via Kalman filter with the structure of HP filter. The blue line represents the trend GDP that is obtained via Kalman smoother with the structure of 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 usage of future values.

Figure 14: An estimate of forecast error due to the use of less noisy signals - $\hat{\lambda}_t$
(US data)



The red line represents the trend GDP that is obtained via Kalman filter with the structure of HP filter using real-time data. The blue line represents the trend GDP that is obtained via Kalman filter with the structure of 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 usage of more precise signals.

Figure 15: A final estimate of forecast error - $\hat{\eta}_t$
(US data)



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

D 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, \hat{p} , $VAR(\hat{p})$ is estimated by Ordinary Least Squares (OLS) to obtain 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 $B(L) = A_0 C(L)$ 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:

$$\Sigma_u = S S' \tag{D.1}$$

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:

$$\Psi(S) = \sum_{j \in J_{S,+}} \sum_{k=0}^P f\left(\frac{b_{i,j}(k)}{s_j}\right) + \sum_{j \in J_{S,-}} \sum_{k=0}^P f\left(\frac{b_{i,j}(k)}{s_j}\right) \quad (\text{D.2})$$

where P is the horizon over which restrictions should hold. 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 $c(S) = [\frac{b_{i,j}(k)}{s_j}^1, \dots, \frac{b_{i,j}(k)}{s_j}^P]$. To identify the model, we solve the following maximization problem:

$$\begin{aligned} S &= \underset{S}{\operatorname{argmax}} && \Psi(S) \\ &\text{subject to} && c(S) \geq 0 \\ &&& SS' = \Sigma_u \end{aligned} \quad (\text{D.3})$$

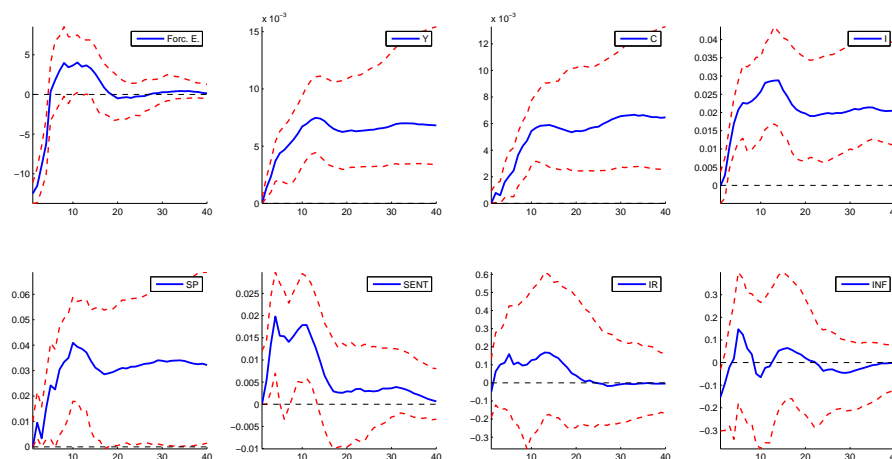
In the current paper only restrictions on forecast error enter the penalty function, while sign restrictions on other variables are applied only as non-negativity constraints. In this way we achieved that most of the forecast variance of forecast error is explained by the two identified shocks, while 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. **Estimation uncertainty:** to account for estimation uncertainty, we repeat steps 1-2 1000 times, each time with a new artificially constructed data sample, Y^* . To construct data samples, we use re-sampling of errors. To correct for small-sample bias, we use the method described in Kilian [1998]. New data sample is constructed recursively as $y_t^* = \hat{A}_1^* y_{t-1}^* + \dots + \hat{A}_N^* y_{t-N}^* + \hat{u}_t^*$, starting from initial values $[y_0, \dots, y_{N-1}]$. \hat{A}_n^* are estimated reduced form autoregressive coefficients. \hat{u}_t^* are drawn randomly with replacement from estimated reduced form errors, \hat{u}_t .

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.

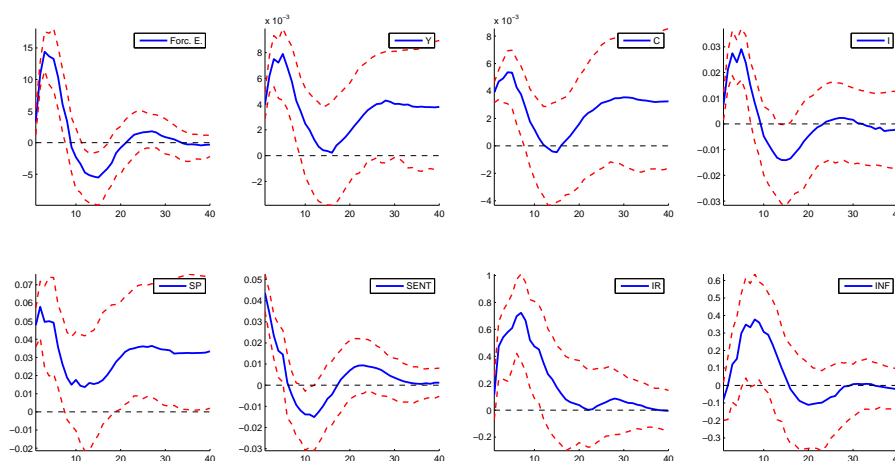
E Appendix - Baseline results

Figure 16: US data - IRFs to permanent shock




The graph presents the impulse responses to the permanent 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 17: US data - IRFs to noise shock




The graph presents the impulse responses to the 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 18: US data - FEVD




The graph presents the forecast error decomposition of the 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 permanent shocks, the green area to the median contribution of noise shocks and yellow area to the contribution of all other non-identified shocks.

Figure 19: US data - extracted shocks (smoothed)




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

Figure 20: US data - historical decomposition of output (permanent shocks off)



The graph presents the historical decomposition of output. The red line is output, while blue line is output in case when we set permanent shocks to zero in all the periods. The difference between the two series of output is presented with the green line in the graph below.

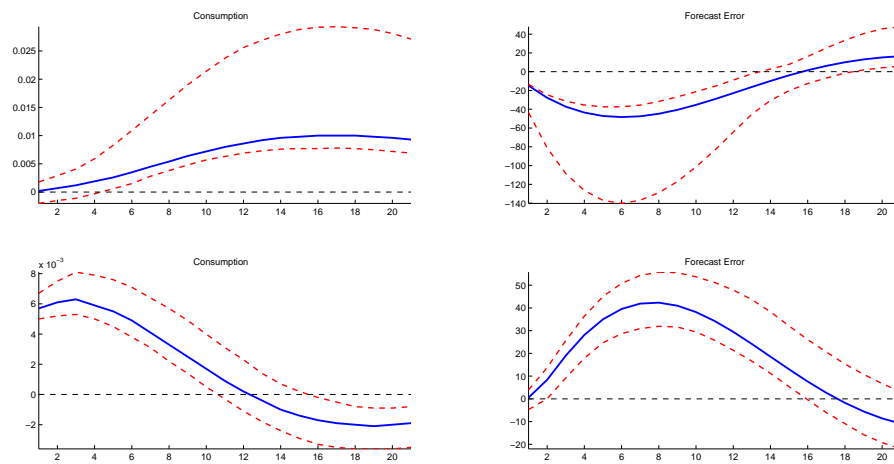
Figure 21: US data - historical decomposition of output (noise shocks off)



The graph presents the historical decomposition of output. The red line is output, while blue line is output in case when we set noise shocks to zero in all the periods. The difference between the two series of output is presented with the green line in the graph below.

F Appendix - Long-run restrictions

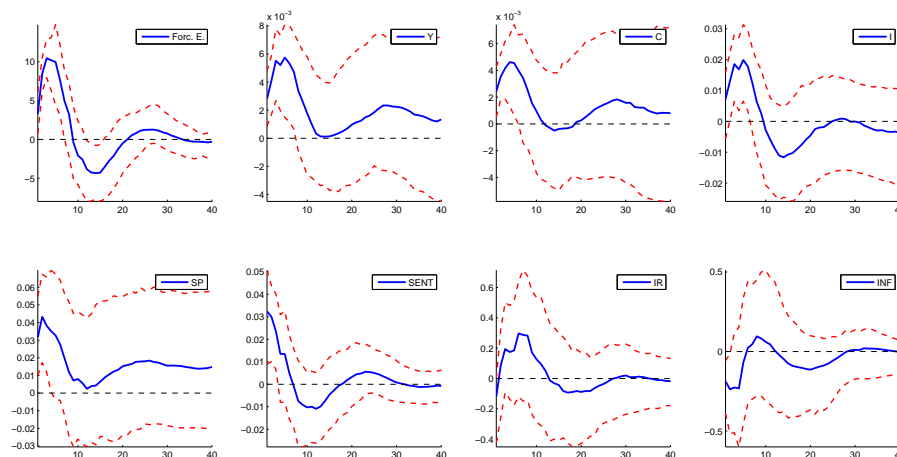
Figure 22: US data - IRFs with long-run restrictions



The graph presents the impulse responses of consumption and the estimated forecast error to the permanent shock (above) and to the transitory shock (below). The identification restrictions rely on long-run restrictions. The blue line is the point estimate and the red lines are 90% error bands.

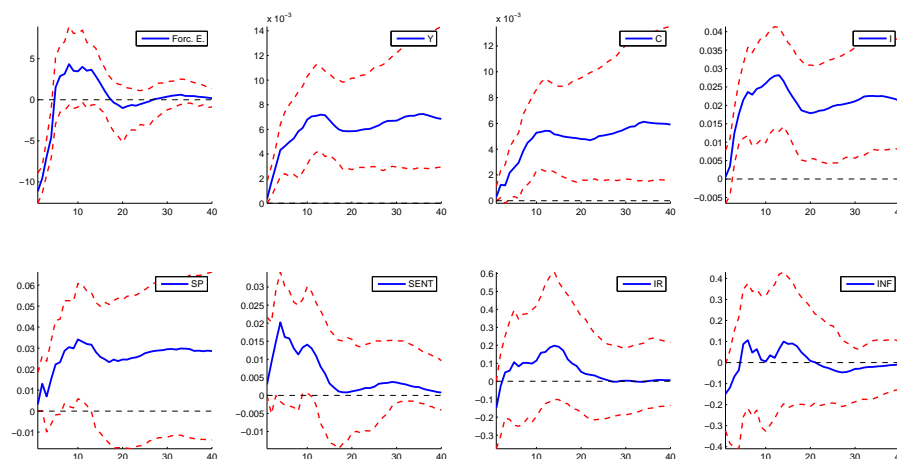
G Appendix - Extended identification

Figure 23: US data - IRFs to noise shock (extended identification)



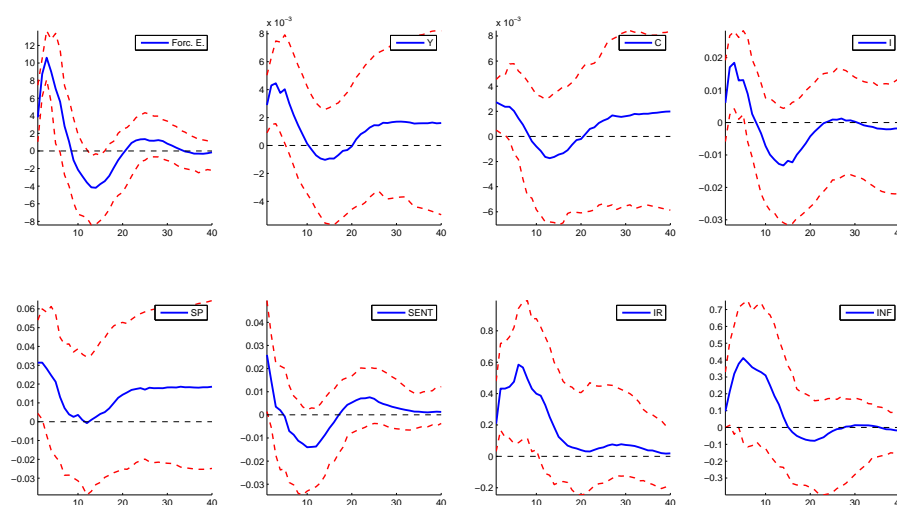
The graph presents the impulse responses to the noise shock of the estimated forecast error, output, consumption, investment (first row), stock prices, sentiment, interest rates and inflation (second row). The blue line is the point estimate and the red lines are 90% error bands.

Figure 24: US data - IRFs to permanent shock (extended identification)



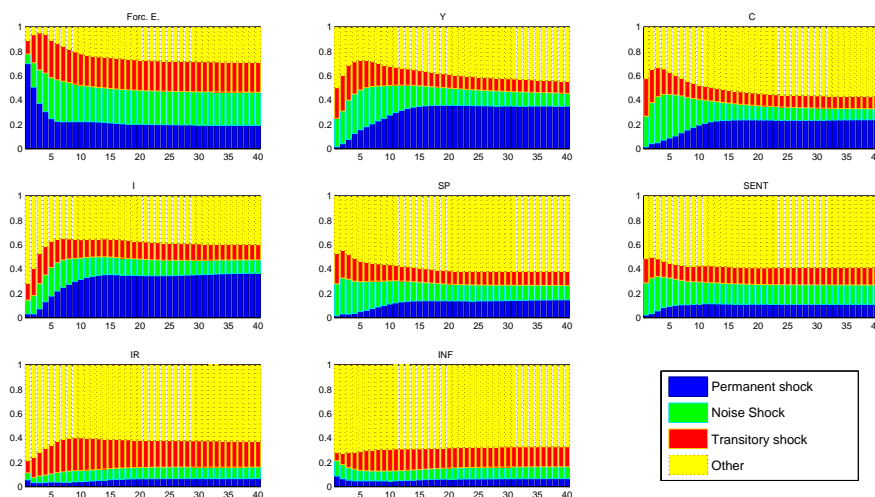
The graph presents the impulse responses to the permanent shock of the estimated forecast error, output, consumption, investment (first row), stock prices, sentiment, interest rates and inflation (second row). The blue line is the point estimate and the red lines are 90% error bands.

Figure 25: US data - IRFs to transitory shock (extended identification)



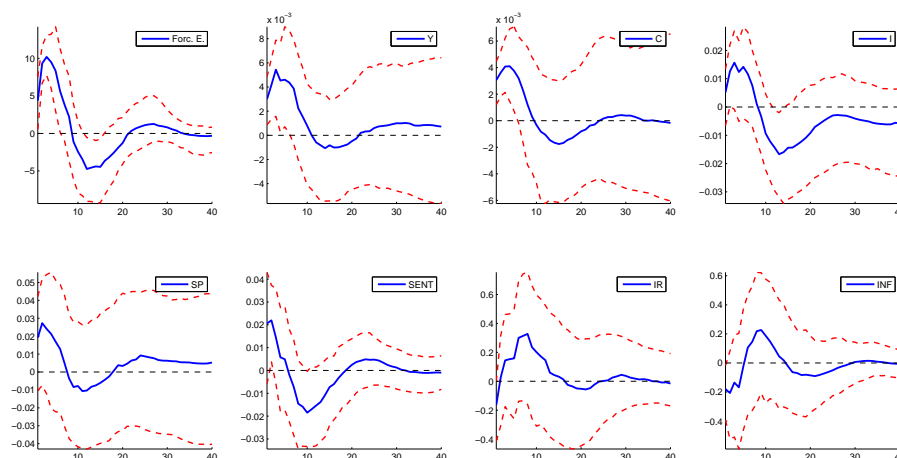
The graph presents the impulse responses to the transitory shock of the estimated forecast error, output, consumption, investment (first row), stock prices, sentiment, interest rates and inflation (second row). The blue line is the point estimate and the red lines are 90% error bands.

Figure 26: US data - FEVD (extended identification)



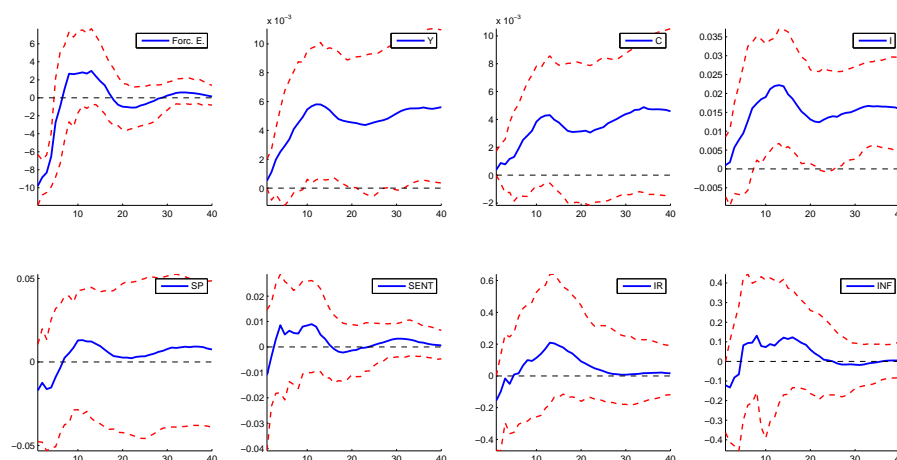
The graph presents the forecast error decomposition of estimated forecast error, output, consumption (first row), investment, stock prices, sentiment (second row), interest rates and inflation (third row). The blue area corresponds to the median contribution of permanent shocks, the green area to the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

Figure 27: US data - IRFs to noise shock (extended identification - minimal restrictions)



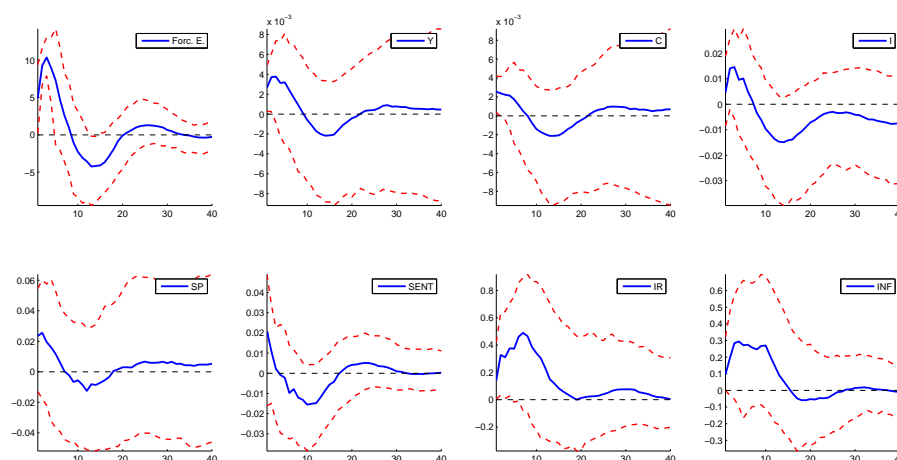
The graph presents the impulse responses to the noise shock of the estimated forecast error, output, consumption, investment (first row), stock prices, sentiment, interest rates and inflation (second row). The blue line is the point estimate and the red lines are 90% error bands.

Figure 28: US data - IRFs to permanent shock (extended identification - minimal restrictions)



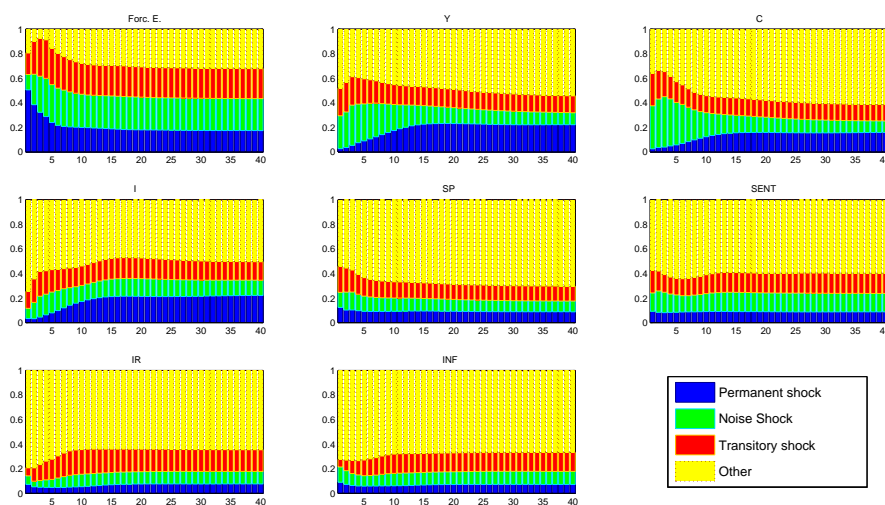
The graph presents the impulse responses to the permanent shock of the estimated forecast error, output, consumption, investment (first row), stock prices, sentiment, interest rates and inflation (second row). The blue line is the point estimate and the red lines are 90% error bands.

Figure 29: US data - IRFs to transitory shock (extended identification - minimal restrictions)



The graph presents the impulse responses to the transitory 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 30: 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 permanent shocks, the green area to the median contribution of noise shocks, the red area to the median contribution of transitory shocks and yellow area to the contribution of all other non-identified shocks.

Acknowledgements

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.

Stephane Dées

European Central Bank, Frankfurt am Main, Germany; email: stephane.dees@ecb.europa.eu

Srečko Zimic

European Central Bank, Frankfurt am Main, Germany; email: srecko.zimic@ecb.europa.eu

© European Central Bank, 2016

Postal address	60640 Frankfurt am Main, Germany
Telephone	+49 69 1344 0
Website	www.ecb.europa.eu

All rights reserved. Any reproduction, publication and reprint in the form of a different publication, whether printed or produced electronically, in whole or in part, is permitted only with the explicit written authorisation of the ECB or the authors.

This paper can be downloaded without charge from www.ecb.europa.eu, from the [Social Science Research Network](#) electronic library at or from [RePEc: Research Papers in Economics](#).

Information on all of the papers published in the ECB Working Paper Series can be found on the [ECB's website](#).

ISSN	1725-2806 (pdf)
ISBN	978-92-899-2201-2 (pdf)
DOI	10.2866/112389 (pdf)
EU catalogue No	QB-AR-16-070-EN-N (pdf)