Stock Prices, News, and Economic Fluctuations: Comment*

André Kurmann
Federal Reserve Board

Elmar Mertens
Federal Reserve Board

January 2, 2013

Abstract

Beaudry and Portier (2006) propose an identification scheme to study the effects of news shocks about future productivity in Vector Error Correction Models (VECM). This comment shows that their methodology does not have a unique solution, when applied to their VECMs with more than two variables. The problem arises from the interplay of cointegration assumptions and long-run restrictions imposed by Beaudry and Portier (2006).

*The views expressed in this paper do not necessarily represent the views of the Federal Reserve System or the Federal Open Market Committee.

1 Introduction

In a highly influential paper, Beaudry and Portier (2006) estimate Vector Error Correction Models (VECMs) on U.S. data and find that shocks generating a stock market boom but no contemporaneous movement in Total Factor Productivity (TFP) — henceforth called "TFP news" — are closely related to shocks driving long-run variations in TFP. Moreover, these TFP news cause increases in consumption, investment, output and hours on impact and constitute an important source of business cycle fluctuations. These results run counter to basic dynamic stochastic general equilibrium (DSGE) models and have sparked a new literature attempting to generate news-driven positive comovement among macroeconomic aggregates.[1]

This comment shows that in the VECMs with more than two variables estimated by Beaudry and Portier (2006), their identification scheme fails to determine TFP news. Yet, these higher-dimension systems are crucial to quantify the business cycle effects of TFP news.[2] The identification problem arises from the interplay of two assumptions. First, the Beaudry-Portier identification scheme requires that one of the non-news shocks has no permanent impact on either TFP or consumption. Second, the VECMs estimated by Beaudry and Portier (2006) impose that TFP and consumption are cointegrated. This means that TFP and consumption have the same permanent component, which makes one of the two long-run restrictions redundant and leaves an infinity of candidate solutions with very different implications for the business cycle. The results reported in Beaudry and Portier (2006) represent just one arbitrary choice among these solutions.[3]

A potential way to address the identification problem is to drop the cointegration restriction between TFP and consumption. We do so by applying Beaudry and Portier's (2006) restrictions, called the "BP restrictions" from hereon, on a vector autoregressive (VAR) system in levels that does not require any a priori assumptions about cointegration. The point estimates of the BP news shock responses in the level VAR resemble closely the results reported by Beaudry and Portier

[1] See for example Beaudry and Portier (2007), DenHaan and Kaltenbrunner (2009), Jaimovich and Rebelo (2009), or Schmitt-Grohe and Uribe (2012).

[2] An equally important reason to work with systems in more than two variables is robustness. If the economy is complicated even in simple ways, then the type of bivariate systems that Beaudry and Portier (2006) use for their baseline analysis is likely to generate inaccurate answers. See Faust and Leeper (1997) for an example in another context.

[3] The replication files posted on the AER website do not include code showing how TFP news were computed by Beaudry and Portier (2006). In private communication, we learned from the authors that their computations relied on a numerical solver that arbitrarily returned one from the infinite number of viable solutions.

(2006) for their VECM systems. However, this identification is surrounded by a tremendous degree of uncertainty because the VAR estimates imply about a 50% chance that TFP and consumption are cointegrated, in which case the BP restrictions fail to identify TFP news. One can therefore not have any reasonable degree of confidence about the results obtained from the VAR in levels.

We also apply the BP restrictions to an alternative VAR system that, consistent with a large class of DSGE models, imposes absence of cointegration between TFP and consumption. In this case, the identification problem disappears but the shock implied by the BP restrictions is largely unrelated to TFP. In sum, dropping the cointegration restriction between TFP and consumption fails to solve the identification problem or generates results that are difficult to interpret as news about future productivity.

The remainder of the comment proceeds as follows. Section 2 explains the identification problem arising with the BP restrictions. Section 3 applies the BP restrictions to VAR-based systems that do not impose cointegration between TFP and consumption. Section 4 evaluates the BP restrictions in VAR systems with alternative cointegration assumptions. Section 5 concludes by briefly describing alternative identification strategies of TFP news that do not depend on cointegration restrictions between TFP and C.

2 The identification problem

Beaudry and Portier (2006) estimate bivariate, three-variable and four-variable VECMs in TFP, a real stock market price (SP), consumption (C), hours (H) and investment (I). These VECMs can be expressed in vector moving average form as

$$\begin{bmatrix} \Delta TFP_t \\ \Delta SP_t \\ X_t \end{bmatrix} = \Delta \boldsymbol{Y}_t = \boldsymbol{C}(L)\boldsymbol{\mu}_t, \tag{1}$$

where X_t is empty for the bivariate case; $X_t = [\Delta C_t]$ for the trivariate case; and $X_t = [\Delta C_t \ \Delta H_t]'$ or $X_t = [\Delta C_t \ \Delta I_t]'$ for the four-variable case. All variables are logged and detrended. The lag polynomial $\boldsymbol{C}(L) \equiv \boldsymbol{I} + \sum_{i=1}^{\infty} \boldsymbol{C}_i L^i$ is inferred from the VECM parameter estimates; the vector $\boldsymbol{\mu}_t$

contains the one-period ahead prediction errors and has variance covariance matrix $E[\boldsymbol{\mu}_t\boldsymbol{\mu}_t'] = \boldsymbol{\Omega}$.[4]

Crucially, the VECM imposes a set of cointegration restrictions $\boldsymbol{\alpha}'\boldsymbol{Y}_t \sim I(0)$, where $\boldsymbol{\alpha}$ denotes the matrix of cointegrating vectors. As discussed by King, Plosser, Stock and Watson (1991) and Hamilton (1994), cointegration imposes restrictions on $\boldsymbol{C}(L)$. In particular, since $\boldsymbol{\alpha}'\boldsymbol{Y}_t$ is stationary, $\boldsymbol{\alpha}'\boldsymbol{C}(1) = \boldsymbol{0}$ and thus, $\boldsymbol{C}(1)$ is singular. This constrains the set of linearly independent restrictions that can be imposed on the VECM to identify structural shocks. The identification problem arising with the BP restrictions stems from these constraints.

Identification maps $\boldsymbol{\mu}_t$ to structural shocks $\boldsymbol{\varepsilon}_t$ by $\boldsymbol{\mu}_t = \boldsymbol{\Gamma}_0\boldsymbol{\varepsilon}_t$, with $E[\boldsymbol{\varepsilon}_t\boldsymbol{\varepsilon}_t'] = \boldsymbol{I}$ and thus $\boldsymbol{\Gamma}_0\boldsymbol{\Gamma}_0' = \boldsymbol{\Omega}$. Impulse responses to the structural shocks are then given by $\boldsymbol{\Gamma}(L) = \boldsymbol{C}(L)\boldsymbol{\Gamma}_0$. Beaudry and Portier's (2006) original idea is that news about future TFP do not have a contemporaneous effect on measured TFP; i.e. if the TFP news innovation is the second element of $\boldsymbol{\varepsilon}_t$, that the $(1,2)$ element of $\boldsymbol{\Gamma}_0$ is zero. For the bivariate systems that Beaudry and Portier (2006) use as their baseline case, this restriction together with $\boldsymbol{\Gamma}_0\boldsymbol{\Gamma}_0' = \boldsymbol{\Omega}$ uniquely identifies TFP news.

The identification problem arises in the three- and four-variate systems where one zero restriction is no longer sufficient to identify structural shocks. Beaudry and Portier's (2006) strategy consists of adding zero restrictions until identification is achieved. In the trivariate case, these additional restrictions are that one of the non-news shocks has no permanent effect on TFP and C; so when this non-news shock is the third element of $\boldsymbol{\varepsilon}_t$, the $(1,3)$ and $(3,3)$ elements of the long-run impact matrix $\boldsymbol{\Gamma}(1) \equiv \boldsymbol{C}(1)\boldsymbol{\Gamma}_0$ are zero. In the four-variable case, the additional restrictions consist of the same two long-run restrictions plus the assumption that one of the other non-news shocks can only have a contemporaneous effect on H, respectively I; so when this other non-news shock is the fourth element of $\boldsymbol{\varepsilon}_t$, the $(1,4)$, $(2,4)$ and $(3,4)$ elements of $\boldsymbol{\Gamma}_0$ are zero.

In a typical VAR, the additional zero restrictions, together with the zero impact restriction on TFP and $\boldsymbol{\Gamma}_0\boldsymbol{\Gamma}_0' = \boldsymbol{\Omega}$, would be sufficient to uniquely identify all elements of $\boldsymbol{\Gamma}_0$ and thus TFP news. Here, this is unfortunately not the case because the three- and four-variable VECMs estimated by Beaudry and Portier (2006) are subject to two, respectively three cointegration restrictions; i.e. $\boldsymbol{\alpha}'$ is a (2×3) matrix, respectively a (3×4) matrix of linearly independent rows.[5] Since $\boldsymbol{\alpha}'\boldsymbol{C}(1) = 0$,

[4] A Web-Appendix provides details of all derivations and computations.

[5] See Footnote 8 and the notes to Figures 9 and 10 in Beaudry and Portier (2006) for the number of cointegration restrictions imposed. The notes to the Figures also state that 4-variable VECMs with 3 (or 4) cointegration restrictions correspond to VARs in levels. However, this seems to be a simple mistake. As Beaudry and Portier (2006) write themselves on page 1296, a VECM is equivalent to a VAR in levels only if the matrix of cointegrating vectors

the rows of $\boldsymbol{C}(1)$ and $\boldsymbol{\Gamma}(1)$ are linearly dependent of each other. In fact, given the number of cointegrating relationships, $\boldsymbol{C}(1)$ and $\boldsymbol{\Gamma}(1)$ are just of rank 1, and only one linearly independent restriction can be imposed on $\boldsymbol{\Gamma}(1)$. One of the two long-run zero restrictions is therefore redundant, leaving $\boldsymbol{\Gamma}_0$ and the shock that is supposed to capture TFP news under-identified.[6]

Another, perhaps more intuitive way to understand the identification problem is to realize that the imposed cointegration relationships imply for TFP and C to share a common trend. But then, when a particular shock, the third element of $\boldsymbol{\varepsilon}_t$ in this case, is restricted to have zero long-run effect on TFP, it automatically also has zero long run effect on C.

The identification problem implies that there exists an infinity of solutions consistent with the BP restrictions. The results reported in Beaudry and Portier (2006) represent one particular solution but there is no economic justification for why this solution should be preferred over any of the other solutions. As we show in the Web-Appendix, some of these solutions are not correlated with the shock driving long-run movements in TFP and generate very different impulse responses. In the context of the three- and four-variable VECMs estimated by Beaudry and Portier (2006), it is therefore impossible to draw any conclusions about TFP news based on the BP restrictions.

3 Dropping the cointegration restriction

A seemingly natural way to address the identification problem while keeping with the BP restrictions is to drop the cointegration restriction between TFP and C. Indeed, as Beaudry and Portier (2006) note themselves, the econometric evidence in favor of two versus one cointegration relationship between TFP, SP and C is not clear-cut, which leaves open the door that TFP and C do not share a common trend. Beaudry and Portier (2006) entertain this possibility in the NBER working paper version of their paper where they report results for one of their baseline bivariate systems estimated as a VAR in levels; i.e. with no cointegration restrictions imposed. However, they do not report any results for level VARs with more than two variables.

One important challenge with implementing the BP restrictions in a VAR in levels is that for

$\boldsymbol{\alpha}$ is of full rank (also see Hamilton, 1994, chapter 19).

[6] Technically, the $(1,3)$ and the $(3,3)$ equation of $\boldsymbol{\Gamma}(1) = \boldsymbol{C}(1)\boldsymbol{\Gamma}_0$ on which the long-run restrictions are imposed are the same. This leaves the system short of one equation to identify $\boldsymbol{\Gamma}_0$. Nothing would change about this identification problem if Beaudry and Portier (2006) had imposed cointegrating restrictions only on TFP and C but not on any of the other variables (i.e. if $\boldsymbol{\alpha}'$ was a row-vector with non-zero entries only in the positions related to TFP and C).

the type of non-stationary variables involved in the estimation, there is no finite-valued solution for the long-run impact matrix of the different shocks. Hence, the long-run zero restrictions on which Beaudry and Portier's (2006) identification scheme relies cannot be imposed exactly.[7] We resolve this issue by first computing the linear combination of VAR residuals that account for most of the forecast error variance (FEV) of TFP, respectively C, at a long but *finite* horizon of 400 quarters; and then using a projection-based procedure to implement the BP restrictions.[8]

We estimate the three- and four-variable level VAR equivalents of Beaudry and Portier's (2006) VECMs using their original data with the number of lags set to four based on traditional information criteria and Portmanteau tests.[9] The first row of Figure 1 reports the results for the four-variable level VAR in (TFP, SP, C, H); the second row reports the results for the level VAR in (TFP, SP, C, I). Very similar results obtain for the three-variable case and are therefore not reported. The red solid lines and the blue dashed lines display, respectively, the impulse responses — generated by the point estimates — to the shock identified by the BP restrictions and the shock driving long-run variations in TFP. The grey intervals represent a measure of uncertainty about the identification implied by the BP restrictions, which will be discussed further below.

Figure 1 about here

The impulse responses derived from the point estimates of both level VARs come surprisingly close to the results reported in Beaudry and Portier (2006) for their VECM systems.[10] In particular, the shocks identified from the BP restrictions and the long-run TFP shock lead to almost identical impulse responses and account for a large fraction of movements in TFP at longer-run frequencies and C, H and I at business cycle frequencies.

[7] Formally, let the VAR in levels be defined as $\boldsymbol{Y}_t = \sum_{i=1}^{p} \bar{\boldsymbol{F}}_{\boldsymbol{i}} \boldsymbol{Y}_{\boldsymbol{t-i}} + \bar{\boldsymbol{\mu}}_{\boldsymbol{t}} = \bar{\boldsymbol{F}}(L)\boldsymbol{Y}_t + \bar{\boldsymbol{\mu}}_{\boldsymbol{t}}$. Then, the vector-moving average representation in (1) can be recovered as $\Delta \boldsymbol{Y}_t = (1-L)(\boldsymbol{I} - \bar{\boldsymbol{F}}(L))^{-1}\bar{\boldsymbol{\mu}}_{\boldsymbol{t}} = \bar{\boldsymbol{C}}(L)\bar{\boldsymbol{\mu}}_{\boldsymbol{t}}$. Non-stationarity of the variables in $\boldsymbol{Y}_t$ implies that the roots of $(\boldsymbol{I} - \bar{\boldsymbol{F}}(L))$ lie strictly inside the unit circle. In this case, the long-run impact matrix $\bar{\boldsymbol{C}}(1)$ does not converge to a finite-valued solution.

[8] Details of the procedure, which to our knowledge is new, are provided in the Web-Appendix. Our approach of first computing shocks that account for most of the FEV at long but finite horizons is reminiscent of Francis, Owyang, Roush and DiCecio's (2012) method of imposing long-run restrictions. While approximately, the thus identified shocks account for more than 95% of movements in TFP, respectively C, at the 400 quarters horizon.

[9] The TFP measure from Beaudry and Portier (2006) that we use is the Solow Residual adjusted with BLS's capacity utilization index. See Section III.B of their paper. Results would be very similar if we instead used a quarterly interpolation of the TFP measure in Basu, Fernald and Kimball (2006), as provided by Fernald (2012).

[10] See Figure 9 in the AER paper and Figure 20 in the NBER working paper.

At first sight, one could thus be led to conclude that dropping the cointegration assumption by estimating VARs in levels addresses the identification problem and resurrects the results reported in Beaudry and Portier (2006). However, the reported impulse responses reflect just the point estimates of the level VARs. The problem is that when sampling confidence sets from the estimated level VARs, about 50% of all draws imply that TFP and C share a common trend.[11] But then, as described in the previous section, the BP restrictions do not identify TFP news and one is left instead with an infinity of candidate solutions.

To illustrate this uncertainty about the BP identification, we take each draw that implies a common trend between TFP and C and compute all candidate solutions that are consistent with the BP restrictions and generate a positive impact response of SP.[12] The grey envelopes in Figure 1 show the resulting range of impulse responses. Clearly, the range is very wide, encompassing the zero line for all variables and frequently extending far beyond the displayed scale. Hence, one cannot have any confidence in the impulse responses generated from the BP restrictions when evaluating the level VARs at their point estimates.

In principle, the lack of identification found in the VECMs could be addressed by estimating level systems, that do not impose the common trend assumption on TFP and C. For example, the point estimates of the level VARs generate a unique solution. But draws generated from the level VARs place sufficient odds in favor of the common trend assumption, such that this approach does not successfully address the identification problem.

4 Alternative cointegrating restrictions

Alternatively, the identification can be addressed by estimating systems which impose that TFP and C have separate trends. Fisher (2010), for example, notes that DSGE models with neutral and

[11] Specifically, for about 50% of the draws in each level VAR, the two shocks driving the long-term components of TFP and C — as identified by maximizing the FEV share over 400 quarters — are so highly collinear that their variance-covariance matrix is ill-conditioned. In these cases, the estimated trends in TFP and C cannot be reliably distinguished from each other, which is a key prerequisite for unique identification under the BP restrictions. Further details are described in the web-appendix.

[12] More specifically, for each draw that implies cointegration between TFP and C, we apply Givens rotations to obtain all possible impulse vectors consistent with the BP restrictions. Any rotation with a negative impact response of SP is eliminated so as not to include simple 180 degree rotations of candidate solutions. See the Web-Appendix for details. We could have instead eliminated rotations with a negative long-run effect on TFP. None of the conclusions would have changed.

investment-specific technology shocks imply that C is not cointegrated with TFP, while sharing a common trend with SP and I.[13] These balanced growth assumptions are straightforward to implement by estimating a stationary VAR in ΔTFP, ΔC, $SP - C$ and $C - I$, respectively H.[14] Since TFP is no longer cointegrated with C, the BP restrictions imply a unique identification across all draws.

Figure 2 about here

We estimate this stationary VAR specification with Beaudry and Portier's (2006) data and apply the BP restrictions. As shown in Figure 2, the resulting point estimates are very different from the ones reported in Beaudry and Portier (2006). In particular, the identified shock generates a *drop* in TFP that lasts for 10 years or more and accounts for only a very small fraction of future movements in TFP. This makes it difficult to interpret the identified shock as news about future productivity.

5 Conclusion

This comment shows that the results reported in Beaudry and Portier (2006) are subject to an important identification problem. The problem arises from the interplay of long-run restrictions and cointegration assumptions that Beaudry and Portier (2006) impose with respect to TFP and C. Dropping the cointegration restriction between TFP and C by estimating a VAR in levels fails to address the identification problem because there is about a 50% probability that TFP and C share a common trend. Alternatively, imposing that TFP and C are not cointegrated by estimating a stationary VAR generate dynamics for TFP that look very different from the ones reported in Beaudry-Portier (2006) and are difficult to interpret as news about future productivity.

The results raise the important question of how to identify TFP news in alternative ways. One example is Beaudry and Lucke (2010) who invoke short- and long-run zero restrictions for non-news shocks that do not depend on cointegration between TFP and C. As Fisher (2010) shows, however,

[13] Other possible causes for absence of cointegration between TFP and C are (permanent) changes in distortionary tax rates or labor force participation.

[14] Equivalently, the balanced growth assumptions can be implemented in Beaudry and Portier's (2006) VECMs by requiring the matrix of cointegrating vectors $\boldsymbol{\alpha}$ to contain only 1s and 0s in the appropriate positions.

the implications for TFP news coming out of this identification crucially depend on the number of cointegration relationships imposed.

Another strategy, recently proposed by Barsky and Sims (2011), is to identify TFP news as the shock orthogonal to contemporaneous TFP movements that accounts for the maximum share of unpredictable future movements in TFP. This strategy, which is consistent with Beaudry and Portier's (2006) original idea that TFP is driven by a contemporaneous component and a slowly diffusing news component, has the advantage that it does not rely on additional zero restrictions about other non-news shocks. As a result, it is robust to different assumptions about cointegration and can be applied to arbitrary vector moving-average systems. Interestingly, Barsky and Sims (2011) find that the TFP news resulting from their identification accounts for a substantial share of TFP and macroeconomic aggregates at medium- and long-run horizons. However, their TFP news shock does not generate the type of joint increase in real macroeconomic aggregates on impact that Beaudry and Portier (2006) report and that generated a lot of interest in the literature.

References

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

Basu, S., J. G. Fernald, and M. S. Kimball (2006, December). Are technology improvements contractionary? *American Economic Review 96*(5), 1418–1448.

Beaudry, P., D. Nam, and J. Wang (2011, December). Do mood swings drive business cycles and is it rational? NBER Working Papers 17651, National Bureau of Economic Research, Inc.

Beaudry, P. and F. Portier (2004, June). Stock prices, news and economic fluctuations. NBER Working Papers 10548, National Bureau of Economic Research, Inc.

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

Beaudry, P. and F. Portier (2007, July). When can changes in expectations cause business cycle fluctuations in neo-classical settings? *Journal of Economic Theory 135*(1), 458–477.

Cochrane, J. H. (1994, December). Shocks. *Carnegie-Rochester Conference Series on Public Policy 41*(1), 295–364.

Den Haan, W. J. and G. Kaltenbrunner (2009, April). Anticipated growth and business cycles in matching models. *Journal of Monetary Economics 56*(3), 309–327.

Faust, J. and E. M. Leeper (1997, July). When do long-run identifying restrictions give reliable results? *Journal of Business & Economic Statistics 15*(3), 345–53.

Fernald, J. (2012, September). A quarterly, utilization-adjusted series on total factor productivity. Working Paper Series 2012-19, Federal Reserve Bank of San Francisco.

Fisher, J. D. (2006, October). The dynamic effects of neutral and investment specific technology shocks. *The Journal of Political Economy 114*(3), 413–451.

Francis, N., M. T. Owyang, J. E. Roush, and R. DiCecio (2012, June). A flexible finite-horizon alternative to long-run restrictions with an application to technology shocks. Working Papers 2005-024, Federal Reserve Bank of St. Louis.

Hamilton, J. D. (1994). *Time-Series Analysis*. Princeton, NJ: Princeton University Press.

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

Kilian, L. (1998, May). Small-sample confidence intervals for impulse response functions. *The Review of Economics and Statistics 80*(2), 218–230.

King, R. G., C. I. Plosser, J. H. Stock, and M. W. Watson (1991). Stochastic trends and economic fluctuations. *The American Economic Review 81*(4), 819–40.

Schmitt-Grohe, S. and M. Uribe (2012). What's news in business cycles. *Econometrica 80*(6), 2733–2764.

Figure 1: Impulse Responses from Level VARs

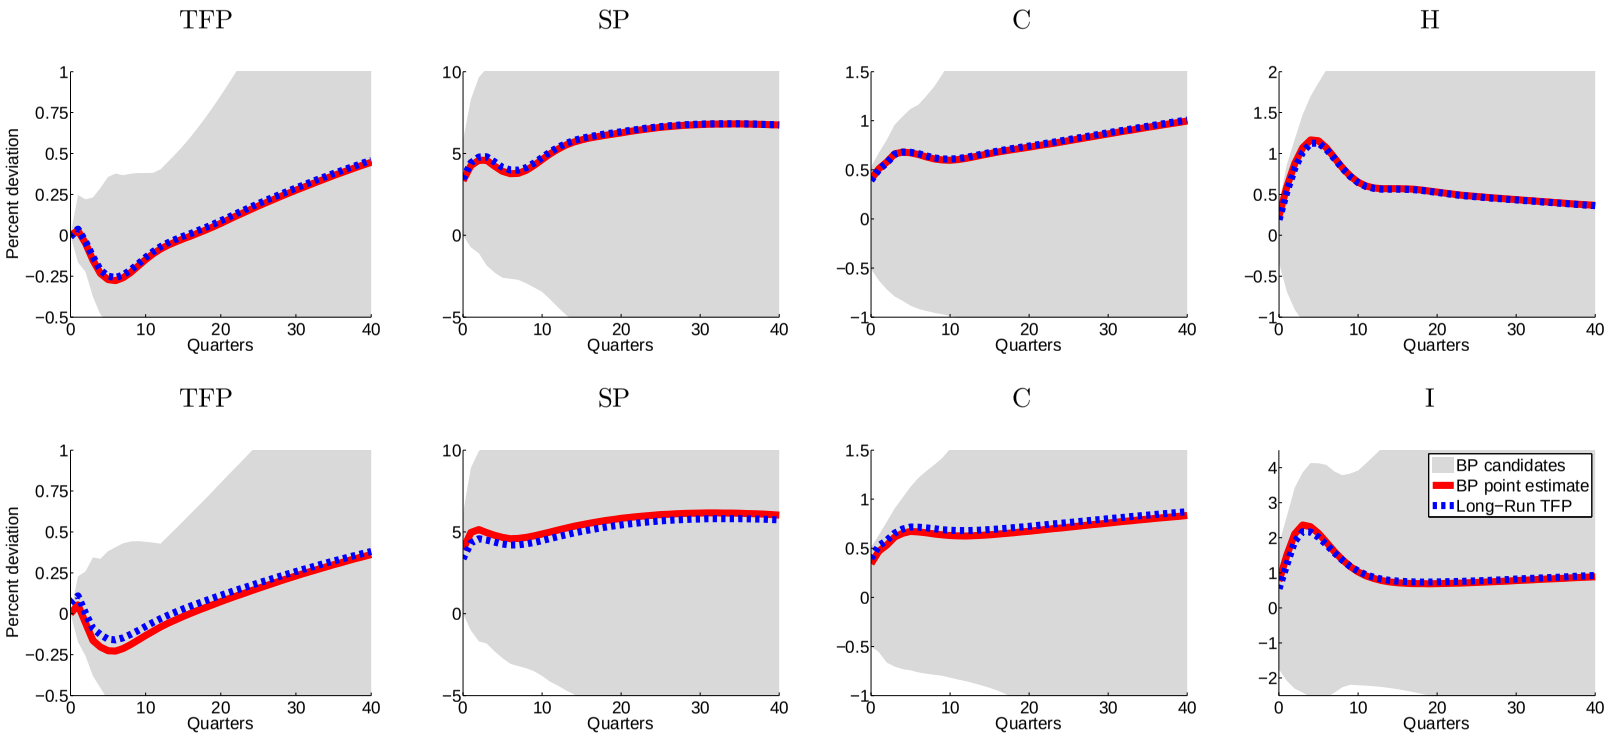

Note: The top row depicts estimates generated by the VAR in TFP, C, SP and H. The bottom row shows estimates from the VAR in in TFP, C, SP and I. Each VAR has 4 lags. In each panel, the solid red line shows point estimates for impulse responses to a news shock identified by the BP restrictions, and the dashed blue line depicts estimates for the long-run shock to TFP. While the BP shocks are uniquely identified when evaluating the level VAR at its point estimates, this is not the case in 45% (upper panels) and 58% (lower panels), respectively, of the draws generated by bootstraps of the level VAR, since the estimated trends in TFP and C are perfectly correlated (up to machine accuracy) for these draws. The grey shaded areas depict the set of impulse responses consistent with the BP restrictions across all these draws. This area also comprises also the distribution of impulse responses generated from those bootstrap draws, where the BP shocks are just identified. For better readability of the point estimates, y-axis limits in each panel have been limited, even though this does not display the full extent of the grey shaded areas. Full pictures are available in the web-appendix accompanying this paper.

Figure 2: Impulse Responses from Stationary VARs

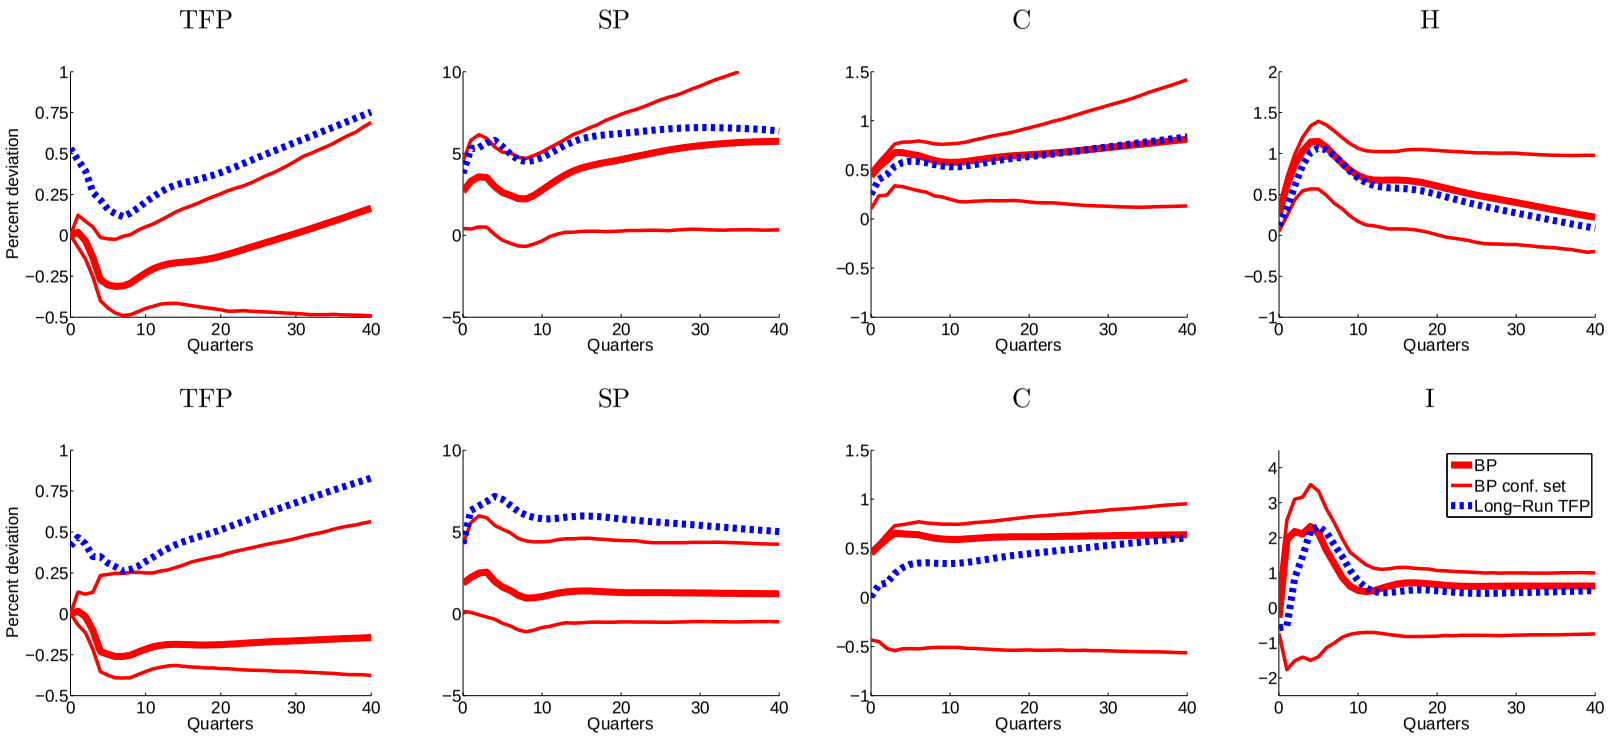

Note: The top row depicts estimates generated by the stationary VAR in ΔTFP, ΔC, $C - SP$ and H. The bottom row shows estimates from the stationary VAR in ΔTFP, ΔC, $C - SP$ and $C - I$. Each VAR has 4 lags. In each panel, the solid red lines show point estimates (thick line) and 80% confidence intervals (thin lines) for the impulse response to news identified by the BP restrictions. The confidence sets have been simulated using Kilian's (1998) adjustment for small-sample bias. The dashed blue line depicts point estimates for the impulse response to a long-run TFP shock.

APPENDIX
FOR ONLINE PUBLICATION

Stock Prices, News, and Economic Fluctuations: Comment*

André Kurmann
Federal Reserve Board

Elmar Mertens
Federal Reserve Board

January 2, 2013

Abstract

This web appendix provides some more analytical details as well as additional results to our main paper.

*The views expressed in this paper do not necessarily represent the views of the Federal Reserve System or the Federal Open Market Committee.

Contents

List of Figures

List of Tables

A The VECM's VMA representation

This appendix derives the vector-moving average (VMA) representation for the VECM systems and shows that the matrix of (non-structural) long-run coefficients, $\boldsymbol{C}(1)$, in equation (2) of the main paper, is singular when derived from the VECM systems estimated by Beaudry and Portier (2006). This relationship holds not only in population but also for any set of sample estimates of the underlying VECM coefficients. Moreover, $\boldsymbol{C}(1)$ has only rank 1, implying that only one (independent) long-run restriction can be imposed on $\boldsymbol{C}(1)$. Since $\boldsymbol{\Gamma}_0$ is assumed to be non-singular, the same properties hold for the sum of the structural VMA coefficients $\boldsymbol{\Gamma}(1) = \boldsymbol{C}(1)\boldsymbol{\Gamma}_0$.

Let $\boldsymbol{Y}_t$ be a vector of non-stationary $I(1)$ variables, which are cointegrated such that $\boldsymbol{\alpha}'\boldsymbol{Y}_t \sim I(0)$ for some matrix of cointegrating vectors $\boldsymbol{\alpha}$. There is then a VECM representation

$$\Delta\boldsymbol{Y}_t = \boldsymbol{F}\Delta\boldsymbol{Y}_{t-1} + \boldsymbol{G}(\boldsymbol{\alpha}'\boldsymbol{Y}_{t-1}) + \boldsymbol{\mu}_t \qquad \boldsymbol{\mu}_t \sim iid(\boldsymbol{0}, \boldsymbol{\Omega}) \tag{1}$$

For the sake of brevity it is assumed that there is only a first-order lag dependence in the VECM, which can be easily generalized to higher order cases. In addition, the notation abstracts from constants and other deterministic components of the data. (Our estimated VECM systems include a constant.)

The associated state-space representation is:

$$\boldsymbol{X}_t = \begin{bmatrix} \Delta \boldsymbol{Y}_t \\ \boldsymbol{\alpha}' \boldsymbol{Y}_t \end{bmatrix} = \begin{bmatrix} \boldsymbol{F} & \boldsymbol{G} \\ \boldsymbol{\alpha}' \boldsymbol{F} & \boldsymbol{I} + \boldsymbol{\alpha}' \boldsymbol{G} \end{bmatrix} \boldsymbol{X}_{t-1} + \begin{bmatrix} \boldsymbol{I} \\ \boldsymbol{\alpha}' \end{bmatrix} \boldsymbol{\mu}_t = \boldsymbol{\mathcal{A}} \boldsymbol{X}_{t-1} + \boldsymbol{\mathcal{B}} \boldsymbol{\mu}_t \quad (2)$$

And it follows the VMA representation:

$$\Delta \boldsymbol{Y}_t = \begin{bmatrix} \boldsymbol{I} & \boldsymbol{0} \end{bmatrix} (\boldsymbol{I} - \boldsymbol{\mathcal{A}} L)^{-1} \boldsymbol{\mathcal{B}} \boldsymbol{\mu}_t = \boldsymbol{C}(L) \boldsymbol{\mu}_t$$

A.1 $\boldsymbol{C}(1)$ has rank 1

As will be shown below, the matrix of long-run coefficients $\boldsymbol{C}(1)$ is singular because of the assumed cointegrating relationships. In particular we have $\boldsymbol{\alpha}' \boldsymbol{C}(1) = 0$, since $\boldsymbol{\alpha}' \boldsymbol{C}(1)$ measures the long-run effect of a shock on the cointegrating vectors, which are stationary and thus their long-run responses are zero (Hamilton, 1994). In the VECM systems used by Beaudry and Portier (2006), there are $N - 1$ cointegrating relationships, and $\boldsymbol{\alpha}$ has $N - 1$ columns, when the VECM has N variables. Thus, $\boldsymbol{C}(1)$ has only rank 1.

The same holds also in sample, for any point estimates of $\boldsymbol{F}$, $\boldsymbol{G}$ and $\boldsymbol{\alpha}$ — provided that $\boldsymbol{\mathcal{A}}$ is stable. This can be verified by computing the partitioned inverse of $\boldsymbol{I} - \boldsymbol{\mathcal{A}}$:

$$(\boldsymbol{I}-\boldsymbol{\mathcal{A}})^{-1}=\begin{bmatrix}\boldsymbol{I}-\boldsymbol{F} & -\boldsymbol{G}\\ -\boldsymbol{\alpha}'\boldsymbol{F} & -\boldsymbol{\alpha}'\boldsymbol{G}\end{bmatrix}^{-1}=\begin{bmatrix}\boldsymbol{M}^{11} & \boldsymbol{M}^{12}\\ \boldsymbol{M}^{21} & \boldsymbol{M}^{22}\end{bmatrix} \tag{3}$$

The standard formulas for the inverse of a partitioned matrix imply in this case $\boldsymbol{M}^{12}=-\boldsymbol{M}^{11}G(\boldsymbol{\alpha}'G)^{-1}$. Further, it follows that

$$\boldsymbol{C}(1)=\boldsymbol{M}^{11}\left(\boldsymbol{I}-\boldsymbol{G}(\boldsymbol{\alpha}'\boldsymbol{G})^{-1}\boldsymbol{\alpha}'\right) \tag{4}$$

And Sylvester's determinant theorem yields:

$$|\boldsymbol{C}(1)|=|\boldsymbol{M}^{11}||(\boldsymbol{\alpha}'\boldsymbol{G})^{-1}||\left(\boldsymbol{\alpha}'\boldsymbol{G}-\boldsymbol{\alpha}'\boldsymbol{G}\right)|=\boldsymbol{0} \tag{5}$$

Furthermore, it is straightforward to show that $\boldsymbol{\alpha}'\boldsymbol{C}(1)=\boldsymbol{0}$ for any point estimates of $\boldsymbol{\alpha}$, $\boldsymbol{F}$ and $\boldsymbol{G}$. To see this, notice that

$$\begin{aligned}
\boldsymbol{M}^{11} &= \left((\boldsymbol{I}-\boldsymbol{F})+\boldsymbol{G}(\boldsymbol{\alpha}'\boldsymbol{G})^{-1}\boldsymbol{\alpha}'\boldsymbol{F}\right)^{-1}\\
&= (\boldsymbol{I}-\boldsymbol{F})^{-1}-(\boldsymbol{I}-\boldsymbol{F})^{-1}\boldsymbol{G}\left(\boldsymbol{\alpha}'(\boldsymbol{I}-\boldsymbol{F})^{-1}\boldsymbol{G}\right)^{-1}\boldsymbol{\alpha}'\boldsymbol{F}(\boldsymbol{I}-\boldsymbol{F})^{-1}\\
\Rightarrow \boldsymbol{\alpha}'\boldsymbol{M}^{11} &= \boldsymbol{\alpha}'(\boldsymbol{I}-\boldsymbol{F})^{-1}-\boldsymbol{\alpha}'\boldsymbol{F}(\boldsymbol{I}-\boldsymbol{F})^{-1}\\
&= \boldsymbol{\alpha}'\\
\Rightarrow \boldsymbol{\alpha}'\boldsymbol{C}(1) &= \boldsymbol{\alpha}'\boldsymbol{M}^{11}\left(\boldsymbol{I}-\boldsymbol{G}(\boldsymbol{\alpha}'\boldsymbol{G})^{-1}\boldsymbol{\alpha}'\right)\\
&= \boldsymbol{0}
\end{aligned}$$

When $\boldsymbol{\alpha}$ has $N-1$ columns and $\boldsymbol{C}(1)$ is a $N \times N$ matrix, it follows that $\boldsymbol{C}(1)$ has rank 1.

A.2 Long-run shocks to TFP in the VECM

This section shows how to implement the identification of long-run shocks to TFP in the VECM systems. Throughout, a one-to-one mapping is assumed between forecast errors $\boldsymbol{\mu}_t$ and structural shocks $\boldsymbol{\varepsilon}_t$, $\boldsymbol{\mu}_t = \boldsymbol{\Gamma}_0 \boldsymbol{\varepsilon}_t$ which must obviously satisfy $\boldsymbol{\Gamma}_0 \boldsymbol{\Gamma}_0' = \boldsymbol{\Omega} = E[\boldsymbol{\mu}_t \boldsymbol{\mu}_t']$.

For the VECMs considered by Beaudry and Portier (2006), there is a single common trend driving the permanent component of all variables, since there are $N-1$ cointegrating relationships when the system has N variables. For the sake of convenience, the shock driving this trend will be referred to as long-run shocks to TFP, while it should be understood that the same shock also accounts for all long-run movements in C, SP and potential other variables, denoted X. This section describes how to construct these long-run shocks from the reduced form parameters of the VECM.

Consider the matrix of structural long-run responses $\boldsymbol{\Gamma}(1) = \boldsymbol{C}(1)\boldsymbol{\Gamma}_0$, and let the first column of $\boldsymbol{\Gamma}_0$ be the responses of forecast errors to the long-run shock. Since no other shock is issued to have a permanent effect on any of the VECM's variables, it follows that

$$\boldsymbol{\Gamma}(1) = \begin{bmatrix} \boldsymbol{x} & \mathbf{0} \end{bmatrix} \tag{6}$$

where $\boldsymbol{x}$ denotes the column vector of long-run responses of $\boldsymbol{Y}_t$ to the long-run shock.

A singular-value decomposition of $\boldsymbol{C}(1)$ yields

$$\boldsymbol{C}(1) = \boldsymbol{V} \begin{bmatrix} \boldsymbol{S}_1 & \boldsymbol{0} \\ \boldsymbol{0} & \boldsymbol{0} \end{bmatrix} \boldsymbol{W}' = \boldsymbol{V}_1 \boldsymbol{S}_1 \boldsymbol{W}_1' \tag{7}$$

where $\boldsymbol{V} = \begin{bmatrix} \boldsymbol{V}_1 & \boldsymbol{V}_2 \end{bmatrix}$ and $\boldsymbol{W} = \begin{bmatrix} \boldsymbol{W}_1 & \boldsymbol{W}_2 \end{bmatrix}$ are conformably partitioned, unitary matrices, $\boldsymbol{V}\boldsymbol{V}' = \boldsymbol{I}$ and $\boldsymbol{W}\boldsymbol{W}' = \boldsymbol{I}$.

Without loss of generality, $\boldsymbol{\Gamma}_0$ can be written as the product of $\boldsymbol{W}$ and another matrix $\tilde{\boldsymbol{B}}$. As will be seen next, the long-run restriction requires that $\tilde{\boldsymbol{B}}$ is (block-) triangular:

$$\tilde{\boldsymbol{B}} = \begin{bmatrix} \tilde{\boldsymbol{B}}_{11} & \tilde{\boldsymbol{B}}_{12} \\ \tilde{\boldsymbol{B}}_{21} & \tilde{\boldsymbol{B}}_{22} \end{bmatrix} = \begin{bmatrix} \tilde{\boldsymbol{B}}_{11} & 0 \\ \tilde{\boldsymbol{B}}_{21} & \tilde{\boldsymbol{B}}_{22} \end{bmatrix} \tag{8}$$

The restriction $\tilde{\boldsymbol{B}}_{12} = 0$ follows from (6) and (7), since it ensures that

$$\boldsymbol{W}_1' \boldsymbol{\Gamma}_0 = \begin{bmatrix} \boldsymbol{z} & \boldsymbol{0} \end{bmatrix} \tag{9}$$

where z denotes an arbitrary column vector.

$\tilde{\boldsymbol{B}}$ factorizes $\tilde{\boldsymbol{\Omega}} = \boldsymbol{W}'\boldsymbol{\Omega}\boldsymbol{W}$. A factorization of $\tilde{\boldsymbol{\Omega}}$ that satisfies the long-run restriction (6) is the Choleski factorization. The first column of $\boldsymbol{\Gamma}_0$ — the column associated with the long-run shock — is then given by the first

column of

$$\boldsymbol{\Gamma}_0 = \boldsymbol{W} \operatorname{chol}(\tilde{\boldsymbol{\Omega}}) \tag{10}$$

and the long-run shocks are the first element of

$$\boldsymbol{\varepsilon}_t = \boldsymbol{\Gamma}_0^{-1} \boldsymbol{\mu}_t \tag{11}$$

where the remaining column of $\boldsymbol{\Gamma}_0$, and thus also the remaining elements of $\boldsymbol{\varepsilon}_t$, reflect an arbitrary permutation of the remaining shocks, without structural interpretation. For future use, the long-run shocks will be denoted $\bar{\varepsilon}_t$.

B Multiple BP shock candidates

The BP scheme for identifying news shocks hinges on two long-restrictions, namely that one of the non-news shocks has zero effect on TFP and C in the long-run. But as shown above, the matrix of long-run responses in the VECM's VMA representation is singular, with a rank of 1, and one of these long-run restrictions is superfluous, and news shocks are not uniquely identified by the BP scheme. This section describes how to compute the set of candidate shocks in the VECM systems, that are all consistent with the BP restrictions.

As an illustration, we reestimate Beaudry and Portier's (2006) four-variable VECMs with their original data and apply the procedure described

here to obtain all possible impulse vectors that respect the BP restrictions and generate a positive impact response of the stock market. The results are reported in Section B.2 below.

B.1 The entire set of solutions the BP scheme

To recap, the BP restrictions for the four-variable case are

1. There is a measurement error shock, which affects only the fourth variable in $\boldsymbol{Y}_t$ on impact; depending on the VECM specification this variable is either H or I. The shock is denoted ε_t^4.

2. The "news shock", denoted ε_t^2 is orthogonal to TFP on impact.

3. There is a pure demand shock, denoted ε_t^3, which has no permanent effect on TFP and C. (As argued above, this shock has thus no permanent effect on any of the VECM variables.)

In addition, all structural shocks are orthogonal to each other and have unit variance. Since the VECM has four variables, the three structural shocks also imply a fourth "residual" structural shock, ϵ_t^1, without any particular interpretation.

A candidate vector of structural shocks can simply be constructed by applying a series of projections using the forecast errors $\boldsymbol{\mu}_t$ and long-run shocks $\bar{\varepsilon}_t$ (see Appendix A.2) as follows:

1. ε_t^4 is the standardized residual in a regression of the fourth VECM residual, μ_t^4 onto the other three residuals.

2. A "news shock" candidate can then be constructed as any linear combination of the VECM residuals, which is orthogonal to the forecast error for TFP, μ_t^1, and the measurement error shocks ε_t^4. As will be shown below, it is then always possible to construct ε_t^3 with the desired properties. Because of the two orthogonality restrictions, only linear combinations in μ_t^2 and μ_t^3 need to be considered when constructing the news shock candidate. Specifically, we use a Givens rotation to construct $e_t = \sin(\theta)\mu_t^2 + \sin(\theta)\mu_t^3$ and compute the news shock candidate as the standardized residual in regressing e_t onto μ_t^1 and ε_t^4. Different news shock candidates are thus indexed by the angle $\theta \in 0, \pi$, denoted $\varepsilon_t^2(\theta)$ (Only the half circle is considered, since the sign of the shock is determined by the restriction that it generates a positive stock market response on impact.)

3. For a given $\varepsilon_t^2(\theta)$ it is straightforward to compute a demand shock candidate, $\varepsilon_t^3(\theta)$, which has no permanent effect on the VECM variables. To ensure this long-run restriction, the demand shock must be orthogonal to $\bar{\varepsilon}_t$, as constructed in Appendix A.2, since $\bar{\varepsilon}_t$ is the sole driver of the permanent component in $\boldsymbol{Y}_t$. In addition, the demand shock has to be orthogonal to ε_t^4 and $\varepsilon_t^2(\theta)$. In sum, the demand shock candidate can be constructed as any linear combination of the VECM residuals

which is orthogonal to $\varepsilon_t^2(\theta)$, $\bar{\varepsilon}_t$ and ε_t^2. Since there are only four VECM residuals and there are three orthogonality constraints, any linear combination of the VECM residuals yields the same projection residual (up to scale and sign) — unless this linear combination should lie in the span of the three orthogonality restrictions, which is easy to check.

For a given candidate vector of shocks $\boldsymbol{\varepsilon}_t(\theta)$ the corresponding candidate matrix $\mathbf{\Gamma}_0(\theta)$ is equal to the covariance matrix $E[\boldsymbol{\mu}_t \boldsymbol{\varepsilon}_t(\theta)]$, which satisfies the BP restrictions by construction. All these computations hold both for population and sample moments.

For the trivariate VECMs, the procedure is identical, except for the absence of ε_t^4. The set of BP candidate shocks is then described by any linear combination of the VECM residuals that is orthogonal to TFP on impact. Again, up to scale and sign, candidate shocks can be computed by projecting of any linear combination of the residuals of SP and C, denoted μ_t^2 and μ_t^3, off μ_t^1.

B.2 Application to the BP-VECMs

The first row of Figure A.1 reports the results for Beaudry and Portier's (2006) four-variable VECM in (TFP, SP, C, H).[1] The second row of Figure A.1 reports equivalent results for the four-variable VECM in (TFP, SP, C, I). Results for the trivariate VECM in (TFP, SP, C) are very

[1]The TFP measure from Beaudry and Portier (2006) that we use is the one adjusted with BLS' capacity utilization index. See Section III.B in their paper.

similar and are available upon request. The blue solid lines replicate impulse responses for the long-run TFP shock reported in Figure 8 of Beaudry and Portier (2006). The grey intervals show the range of candidate solutions consistent with the BP restrictions. Finally, Example 1 (dash-dotted black lines) and Example 2 (dotted red lines) display the impulse responses for two particular solutions. Example 1 corresponds to the solution that fits the impulse response of TFP to the long-run shock best in a least square sense; Example 2 corresponds to the solution that generates a near-zero response of TFP at the 40 quarter forecast horizon.

Consistent with the BP restrictions, none of the candidate solutions affect TFP on impact. Likewise but not shown here, none of the corresponding non-news shocks in the third position of $\boldsymbol{\varepsilon}_t$ have a permanent effect on either TFP or C; and none of the corresponding non-news shocks in the fourth position of $\boldsymbol{\varepsilon}_t$ have a contemporaneous effect on TFP, SP and C. This confirms numerically that there is an infinity of candidate solutions satisfying the BP restrictions.

The grey intervals and the two examples show that the candidate solutions have very different implications. As Example 1 shows, there exists a solution that appears very close to the impulse responses reported in Figure 8 of Beaudry and Portier (2006). By contrast, as Example 2 shows, another solution that is equally consistent with the BP restrictions generates almost no reaction in TFP but a persistent drop in consumption and hours, respectively investment.

Given the very different results across rotations, it should not come as a surprise that the range of correlation coefficients between the shocks satisfying the BP restrictions and the long-run TFP shock is wide for both VECMs, ranging from about -0.50 to 0.99. Likewise, as Table A.1 shows, the forecast error variance (FEV) shares of the different variables attributable to shocks consistent with the BP restrictions extends from basically 0% to above 80% for certain forecast horizons.

Each of these candidate solutions also implies different responses to the "demand shock", ε_t^3. As required, all of these solutions have zero effect on TFP and C, and — by virtue of the assumed common trend in all variables — neither on SP and H. This is illustrated in Figure A.2, which depicts the set of impulse responses the demand shock in each VECM at very long horizons. These results provide a computational consistency check, that the BP restrictions indeed hold for the entire range of shock responses shown in Figure A.1.

Figure A.1: Impulse Responses: VECMs

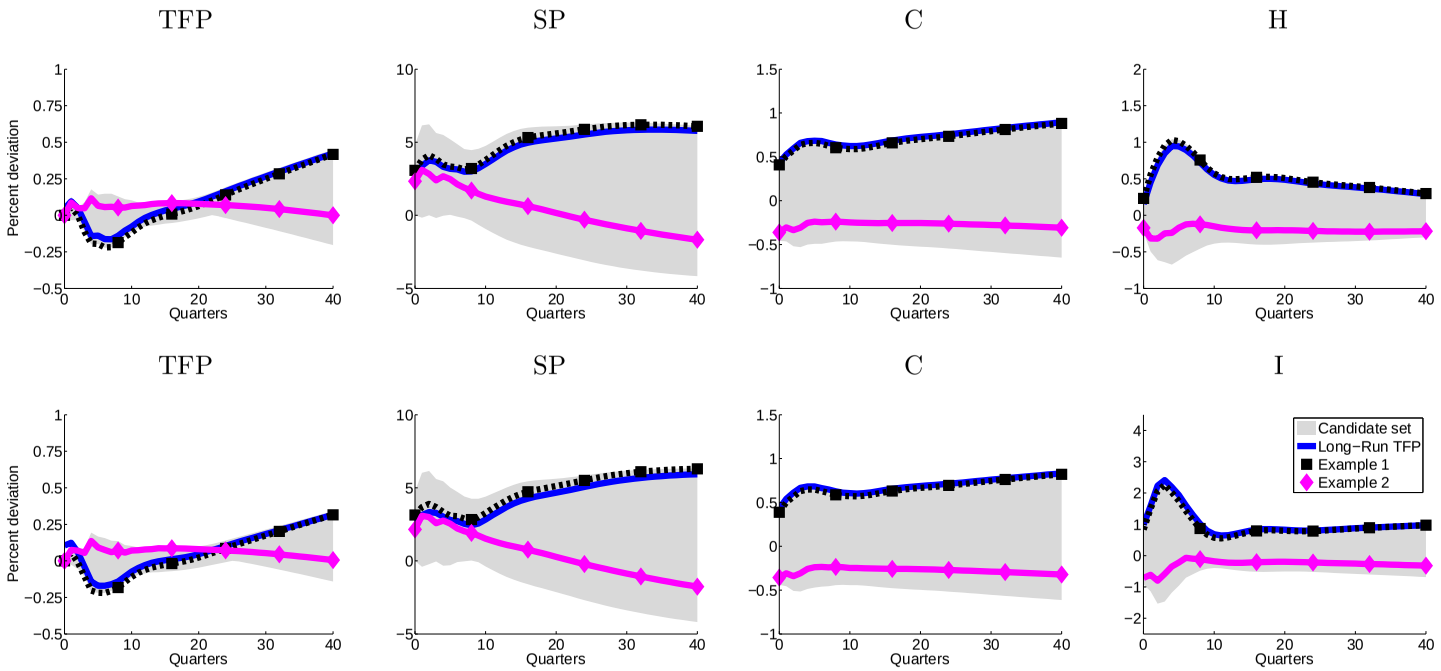

Note: The top row depicts estimates generated by a VECM in TFP, SP, C and H. Bottom row shows estimates from VECM in TFP, SP, C and I. Both VECMs are estimated with 5 lags and 3 cointegrating vectors, identical to what has been used by Beaudry and Portier (2006, "BP"). In each panel, the solid blue line shows impulse response to a long-run shock in TFP; the grey shaded area depicts the set of all impulse responses consistent with the BP restrictions.; the dashed (yellow) and dash-dotted (magenta) lines show two particular impulse responses consistent with the BP restrictions. Example 1 is as close as possible to the responses generated by the long-run TFP shocks, while satisfying the BP restrictions. Example 2 has been chosen to generate a near-zero response of TFP, while also satisfying the BP restrictions.

Table A.1: Range of FEV Shares generated by VECM Estimates

	Panel (a): VECM with $X = H$																							
	Forecast horizons																							
	1			4			8			16			40			120							
TFP	0.00	–	0.00	0.78	–	1.56	0.94	–	9.09	1.47	–	9.03	2.04	–	28.04	2.53	–	84.41					
SP	0.00	–	95.28	0.12	–	88.05	0.10	–	83.00	1.90	–	86.52	5.45	–	89.77	12.50	–	81.08					
C	0.00	–	95.43	1.50	–	88.29	1.73	–	83.25	0.94	–	82.65	0.46	–	88.02	0.34	–	95.64					
H	0.00	–	17.01	1.05	–	54.30	1.26	–	71.22	1.22	–	77.12	2.38	–	80.04	6.51	–	75.81					
	Panel (b): VECM with $X = I$																							
	Forecast horizons																							
	1			4			8			16			40			120							
TFP	0.00	–	0.00	0.82	–	1.24	0.91	–	8.84	1.29	–	8.78	2.87	–	18.38	2.41	–	81.92					
SP	0.00	–	94.66	0.06	–	85.92	0.05	–	82.25	1.51	–	86.33	6.42	–	88.54	12.67	–	80.48					
C	0.00	–	94.52	1.54	–	84.44	1.76	–	79.27	0.97	–	80.10	0.41	–	87.31	0.32	–	95.69					
I	0.00	–	15.66	0.83	–	38.15	1.05	–	36.71	0.95	–	36.39	0.78	–	47.15	1.30	–	72.98					

Note: Range of forecast error variance shares explained by shocks satisfying the BP restrictions in the VECM systems at different forecast horizons. Both VECMs are estimated with 5 lags and 3 cointegrating vectors, identical to what has been used by Beaudry and Portier. Since the BP shocks are underidentified in these systems, each column reports the lowest and highest shares found amongst all possible shocks, satisfying the BP restrictions.

Figure A.2: Sets of Impulse Responses for Demand Shock Candidates in the VECMs

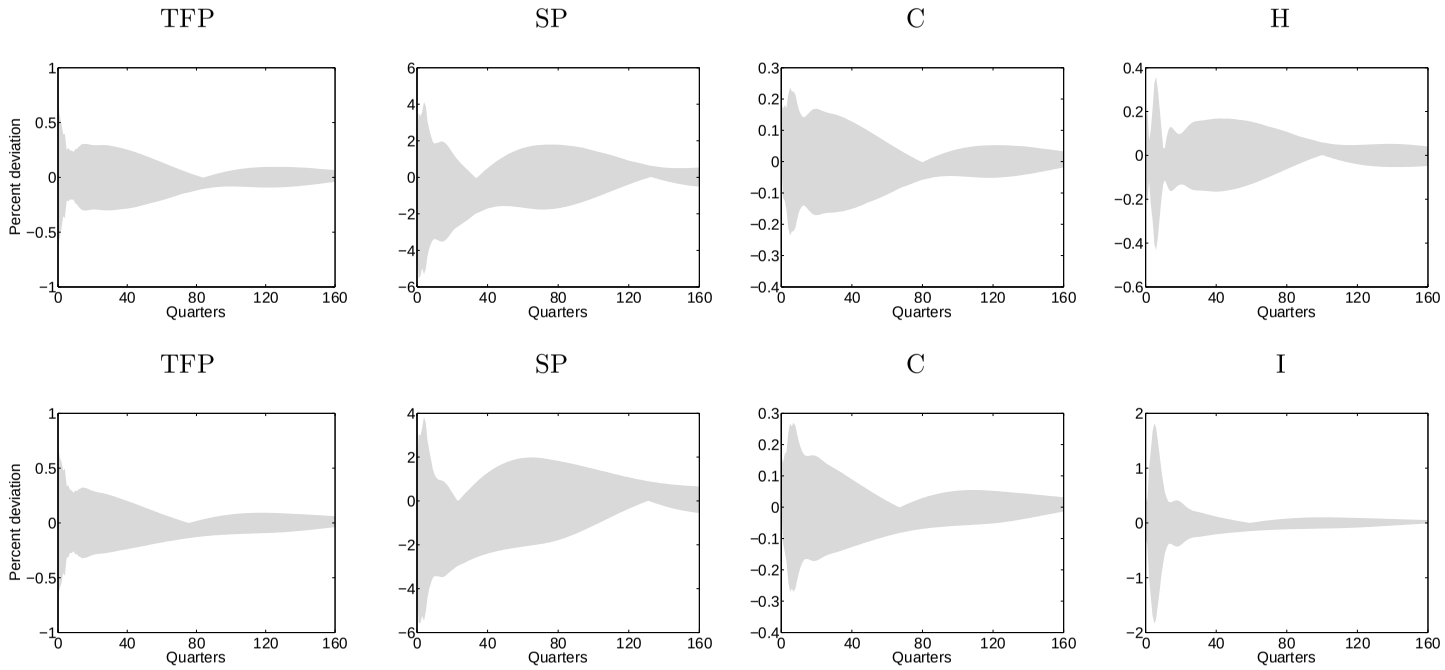

Note: The top row depicts estimates generated by a VECM in TFP, SP, C and H. Bottom row shows estimates from VECM in TFP, SP, C and I. Both VECMs are estimated with 5 lags and 3 cointegrating vectors, identical to what has been used by Beaudry and Portier (2006, "BP"). In each panel, the grey shaded area depicts the set of all impulse responses to the "demand" shock consistent with the BP restrictions. By construction, this shock has no long-run effect on either TFP or C, and by virtue of the assumed cointegrating relationships, neither on SP.

C BP restrictions in the stationary VAR

This section describes the identification of BP shocks in the stationary VAR. The implementation is fairly similar to the VECM case described in Appendix B above. The major difference is that there is now a unique solution for the BP identification, since the stationary VAR allows for distinct trends in TFP and C.

The BP news shock is constructed by projecting a linear combination of $\tilde{\boldsymbol{\mu}}_t$ off the measurement error shock ε_t^4, the demand shock ε_t^3 and the forecast error in TFP μ_t^1. As before, ε_t^4 is given by projecting μ_t^4 off $\mu_t^1, \mu_t^2, \mu_t^3$. (The construction of the demand shock will be described further below.) Let these three innovations be stacked in a vector

$$\boldsymbol{z}_t = \begin{bmatrix} \varepsilon_t^4 \\ \varepsilon_t^3 \\ \mu_t^1 \end{bmatrix}$$

and notice that $\boldsymbol{z}_t$ is entirely spanned by $\tilde{\boldsymbol{\mu}}_t$. Since $\tilde{\boldsymbol{\mu}}_t$ has four elements and $\boldsymbol{z}_t$ has three elements, the residuals of projecting any linear combination $\boldsymbol{w}'\tilde{\boldsymbol{\mu}}_t$ off $\tilde{\boldsymbol{\mu}}_t$ are perfectly correlated (provided the linear combination is not perfectly spanned by $\boldsymbol{z}_t$). For example, we can project μ_t^2 off $\boldsymbol{z}_t$ to construct the BP shock (up to sign and scale). The sign of the news shock is then determined by the condition that $E[\mu_t^3 \varepsilon_t^2] > 0$ and the scale is identified from $E[\varepsilon_t^2] = 1$.

What remains to be shown is the construction of the demand shock ε_t^3, which in turn will depend on constructing two shocks, that drive the permanent components of TFP and C; denoted $\bar{\varepsilon}_t^{TFP}$ and $\bar{\varepsilon}_t^C$. These two shocks can be constructed using the conventional procedure of Blanchard and Quah (1989) for long-run identification. Notice that these two shocks have no structural interpretation in this context, they are merely sufficient statistics for implementing the long-run restrictions on the demand shock. Specifically, the long-run restrictions amount to require that ε_t^3 is orthogonal to $\bar{\varepsilon}_t^{TFP}$ and $\bar{\varepsilon}_t^C$.

The long-run "innovations" $\bar{\varepsilon}_t^{TFP}$ and $\bar{\varepsilon}_t^C$, are constructed by factorizing the long-run variance of $\Delta\tilde{\boldsymbol{Y}}_t$, denoted $\boldsymbol{S}$ as follows:

$$\begin{aligned}
\boldsymbol{S} &= \left(\boldsymbol{I} - \tilde{\boldsymbol{F}}(1)\right)^{-1} \boldsymbol{\Omega} \left(\left(\boldsymbol{I} - \tilde{\boldsymbol{F}}(1)\right)^{-1}\right)' \\
\bar{\boldsymbol{B}} &= \left(\boldsymbol{I} - \tilde{\boldsymbol{F}}(1)\right) \mathrm{chol}\left(\boldsymbol{S}\right) \\
\begin{bmatrix} \bar{\varepsilon}_t^{TFP} \\ \bar{\varepsilon}_t^C \end{bmatrix} &= \begin{bmatrix} \boldsymbol{I} & \boldsymbol{0} \end{bmatrix} \bar{\boldsymbol{B}}^{-1} \tilde{\boldsymbol{\mu}}_t
\end{aligned}$$

In this implementation, $\bar{\varepsilon}_t^{TFP}$ accounts entirely for fluctuations in the permanent component of TFP, as well as for some of the permanent component in C, while $\bar{\varepsilon}_t^C$ explains fluctuations in the stochastic trend in C, which are orthogonal to trend movements in TFP.

Given ε_t^4, $\bar{\varepsilon}_t^{TFP}$ and $\bar{\varepsilon}_t^C$, the demand shock can be constructed as the

standardized residual from projecting any linear combination of $\tilde{\boldsymbol{\mu}}_t$ onto $\begin{bmatrix} \varepsilon_t^4 & \bar{\varepsilon}_t^{TFP} & \bar{\varepsilon}_t^C \end{bmatrix}'$. Using similar reasoning as before, any linear combination yields the same standardized residuals (except for the degenerate cases where the linear combination is completely spanned and the residuals are all zero).

As before, the matrix of impact coefficients $\boldsymbol{\Gamma}_0$ is identical to the matrix of covariances between VAR residuals and structural shocks, and these relationships hold in population as well as for sample moments.

D BP restrictions in the level VAR

Our implementation of the BP restrictions in the level VAR is very similar to the procedure for the stationary VAR outlined in Appendix C. For given shocks ε_t^4, $\bar{\varepsilon}_t^{TFP}$ and $\bar{\varepsilon}_t^C$, the news shock can be estimated as the projection residual between any linear combination of the VAR's forecast errors, $\bar{\boldsymbol{\mu}}_t$, and the above-mentioned three shocks. As before, the measurement error shock ε_t^4, can be obtained by projecting the fourth VAR residual off the other three VAR residuals.

The only special feature of our implementation for the level VAR, is the identification of the long-run shocks. Since point estimates of the level VAR typically imply explosive behavior, the sum of the estimated VMA coefficients does not converge to a finite number, and long-run shocks cannot be constructed as in Blanchard and Quah (1989) by factorizing the long-run variance (see also Appendix D).

We follow Francis et al. (2012) and identify the long-run shocks based on their explanatory power for variations in TFP and C at long but *finite* horizons. Specifically, we construct $\bar{\varepsilon}_t^{TFP}$, to explain as much as possible of the forecast-error variance of TFP at $h = 400$ lags, and similarly for $\bar{\varepsilon}_t^{C}$ and C.

For this method it is convenient to express the identification in terms of an orthonormal matrix $\boldsymbol{Q}$ ($\boldsymbol{Q}\boldsymbol{Q}' = \boldsymbol{I}$). and not in terms of the matrix of impact coefficients $\boldsymbol{\Gamma}_0$, where both are assumed to be related via the Cholesky decomposition of the VAR's forecast error variance, $\boldsymbol{\Gamma}_0 = \text{chol}\,(\boldsymbol{\Omega})\boldsymbol{Q}$.

We seek the column of $\boldsymbol{Q}$, associated with a long-run shock to TFP. Denoting this column $\boldsymbol{q}$, it solves the following variance maximization problem

$$\max_{\boldsymbol{q}} \quad \boldsymbol{h}_1' \left(\sum_{k=0}^{400} \boldsymbol{C}_k \,\text{chol}\,(\boldsymbol{\Omega})\boldsymbol{q}\,\boldsymbol{q}'\,\text{chol}\,(\boldsymbol{\Omega})'\boldsymbol{C}_k' \right) \boldsymbol{h}_1 = \boldsymbol{q}' \underbrace{\left(\sum_{k=0}^{400} \boldsymbol{C}_k' \,\text{chol}\,(\boldsymbol{\Omega})'\boldsymbol{h}_1'\,\boldsymbol{h}_1 \,\text{chol}\,(\boldsymbol{\Omega})\boldsymbol{C}_k \right)}_{\equiv S} \boldsymbol{q}$$

$$\text{subject to} \quad \boldsymbol{q}'\boldsymbol{q} = \boldsymbol{1}$$

where $\boldsymbol{C}_k$ are the coefficients of the VAR's vector moving average representation, $\boldsymbol{C}(L) = \left(\boldsymbol{I} - \bar{\boldsymbol{F}}(L)\right)^{-1}$, $\boldsymbol{h}_1$ selects TFP from the vector of variables in the VAR. Shocks $\bar{\varepsilon}_t^{TFP}$ are constructed using

$$\bar{\varepsilon}_t^{TFP} = \boldsymbol{h}_1 \begin{bmatrix} \boldsymbol{q} & \boldsymbol{N} \end{bmatrix}^{-1} \bar{\boldsymbol{\mu}}_t$$

where $\boldsymbol{N}$ spans the null space of $\boldsymbol{q}$ such that $\begin{bmatrix} \boldsymbol{q} & \boldsymbol{N} \end{bmatrix}$ is orthonormal.

The procedure is analogous for $\bar{\varepsilon}_t^C$, using instead of $\boldsymbol{h}_1$ a vector $\boldsymbol{h}_2$, which selects C from the vector of VAR variables.

A similar procedures is also used to identify news shocks as defined by Barsky and Sims (2011) and Beaudry et al. (2011). There are just two differences: First, both procedures uses different forecast horizons. Beaudry et al. consider forecast horizons of of 40, 80 or 120 leads; and our paper reports results for 120 leads. Barsky and Sims average over the forecast error variances at leads one to 40. Second, both approaches impose the additional requirement that the maximizing shock vector $\boldsymbol{q}$ is orthogonal to a vector which selects TFP from the set of VAR variables; in the present context, this requirement amounts to the first element of $\boldsymbol{q}$ being zero.

D.1 Lack of identification when long-term shocks are collinear

As a necessary condition, $\bar{\varepsilon}_t^C$ and $\bar{\varepsilon}_t^{TFP}$ must not be perfectly correlated, to obtain a unique solution to the projection-based procedure described in Appendix C. When both long-run shocks are perfectly correlated, the orthogonal complement to the space spanned by $\begin{bmatrix} \varepsilon_t^4 & \bar{\varepsilon}_t^{TFP} & \bar{\varepsilon}_t^C \end{bmatrix}'$ is not anymore one-dimensional. (A similar issue would arise, if one of the two long-run shocks were perfectly correlated with ε_t^4, the measurement error shock to the fourth variable.)

When simulating confidence sets for the level VARs, we found that for about 50% of the draws, $\bar{\varepsilon}_t^C$ and $\bar{\varepsilon}_t^{TFP}$ are so highly collinear, that their variance covariance matrix is ill-conditioned. As a consequence, the variance-covariance matrix of $\begin{bmatrix} \varepsilon_t^4 & \bar{\varepsilon}_t^{TFP} & \bar{\varepsilon}_t^C \end{bmatrix}'$ is ill-conditioned. In these cases, we treat $\bar{\varepsilon}_t^C$ and $\bar{\varepsilon}_t^{TFP}$ as perfectly correlated, such that TFP and C share the same common trend. The news shocks are then underidentified, and an infinite number of solutions can be traced out, using a procedure analagously to what is described in Appendix B.

E Additional Results

This appendix provides the following supplemental results: Figure A.3 reports impulse-responses to the BP shocks in the level VARs. The results are identical to those shown in Figure 1 of the paper, except that Figure A.3 displays the results at full scale. Table A.2 reports the shares of forecast error variances explained by the BP shocks at different horizons in the level VARs, and Table A.3 reports the analogous results for the stationary VARs.

Figure A.3: Impulse Responses: Level VAR (larger scale)

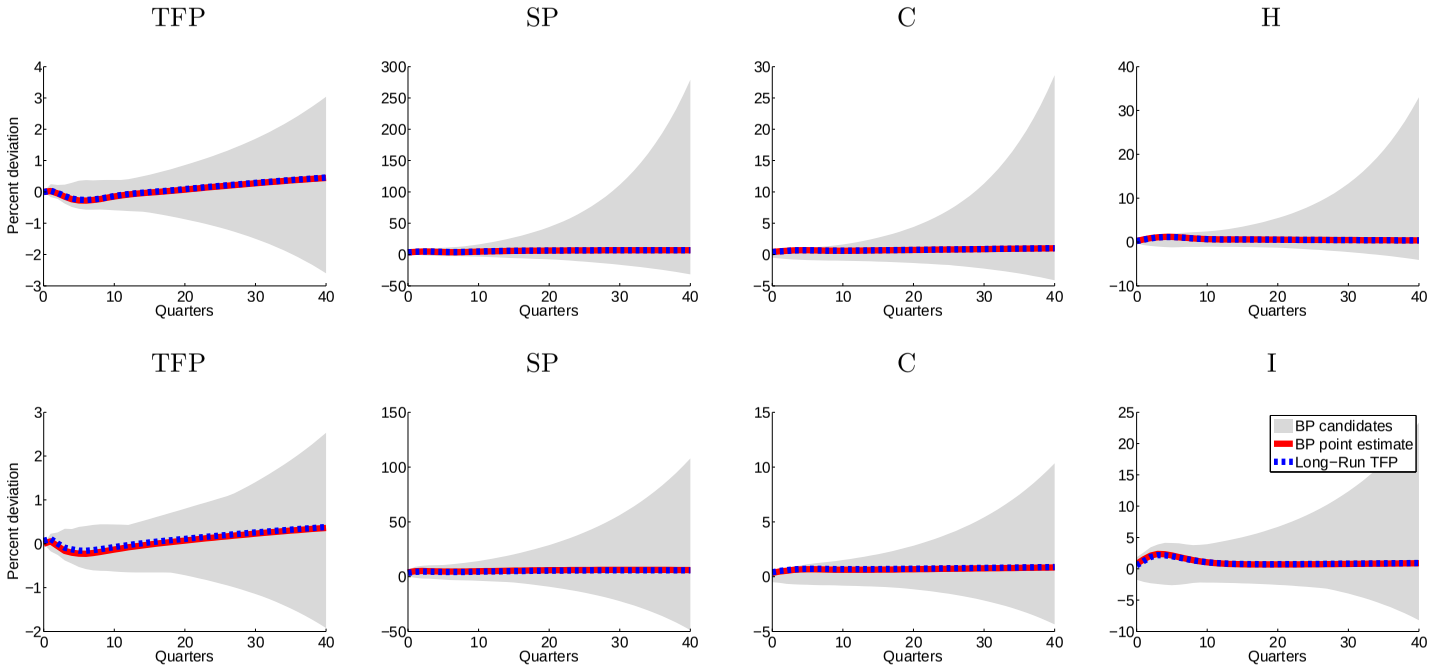

Note: The results shown in this figure are identical to Figure 1 of the main paper, except for the scale of each plot. The top row depicts estimates generated by the VAR in TFP, C, SP and H. The bottom row shows estimates from the VAR in in TFP, C, SP and I. Each VAR has 4 lags. In each panel, the solid red line shows point estimates for impulse responses to a news shock identified by the BP restrictions, and the dashed blue line depicts estimates for the long-run shock to TFP. While the BP shocks are uniquely identified when evaluating the level VAR at its point estimates, this is not the case in 45% (upper panels) and 58% (lower panels), respectivelym, of the draws generated by bootstraps of the level VAR, since the estimated trends in TFP and C are perfectly correlated (up to machine accuracy) for these draws. The grey shaded areas depict the set of impulse responses consistent with the BP restrictions across all these draws. This area also comprises also any confidence set of impulse responses generated from the bootstrap draws, where the BP shocks are just identified.

Table A.2: FEV Shares of BP shocks generated by Level VARs

	Panel (a): VAR with $X = H$					
	Forecast horizons					
	1	4	8	16	40	120
TFP	0.00	0.29	4.85	8.63	2.89	76.25
	(0.00 – 99.08)	(0.02 – 97.58)	(0.13 – 98.47)	(0.06 – 98.44)	(0.06 – 99.16)	(0.22 – 99.99)
C	70.24	82.35	85.37	84.95	88.53	96.46
	(0.00 – 99.28)	(0.06 – 98.69)	(0.02 – 98.53)	(0.05 – 98.70)	(0.16 – 99.18)	(0.51 – 100.00)
SP	43.39	43.23	44.26	51.92	73.33	91.08
	(0.00 – 99.76)	(0.01 – 99.75)	(0.06 – 99.29)	(0.03 – 99.62)	(0.07 – 99.93)	(0.42 – 100.00)
H	14.64	46.76	68.85	82.18	90.07	94.28
	(0.00 – 86.73)	(0.16 – 84.87)	(0.18 – 94.73)	(0.12 – 98.96)	(0.14 – 99.91)	(0.10 – 100.00)
	Panel (b): VAR with $X = I$					
	Forecast horizons					
	1	4	8	16	40	120
TFP	0.00	0.24	3.39	6.29	1.99	63.44
	(0.00 – 99.55)	(0.01 – 99.21)	(0.01 – 96.75)	(0.01 – 96.54)	(0.11 – 98.86)	(0.43 – 99.96)
C	53.26	67.86	72.42	74.24	79.47	89.43
	(0.00 – 95.56)	(0.04 – 96.82)	(0.09 – 97.51)	(0.03 – 98.78)	(0.11 – 99.52)	(0.31 – 100.00)
SP	59.11	53.94	54.37	60.25	78.62	94.74
	(0.00 – 99.90)	(0.01 – 99.38)	(0.02 – 99.80)	(0.05 – 99.97)	(0.09 – 100.00)	(0.24 – 100.00)
I	10.02	31.71	42.06	48.08	61.17	86.73
	(0.00 – 74.91)	(0.01 – 80.80)	(0.00 – 90.66)	(0.01 – 95.07)	(0.05 – 98.75)	(1.18 – 100.00)

Note: Shares of forecast error variances explained by shocks identified from the Beaudry-Portier restrictions in the level VARs at different forecast horizons. Results are generated from VARs with 4 lags. Since the restrictions leave the solution underidentified in about 50% of the bootstraps, the entire range of potential solutions is reported in parentheses.

Table A.3: FEV Shares of BP shocks generated by Stationary VARs

	Panel (a): VAR with $X = H$					
	Forecast horizons					
	1	4	8	16	40	120
TFP	0.00	1.15	11.34	12.21	5.15	24.64
	(0.00 – 0.00)	(0.52 – 6.94)	(3.98 – 27.20)	(4.52 – 36.54)	(3.79 – 40.14)	(5.93 – 66.27)
C	81.81	87.63	83.20	79.69	75.86	76.25
	(50.99 – 97.66)	(64.76 – 95.13)	(60.10 – 92.69)	(56.00 – 91.38)	(49.39 – 92.08)	(45.03 – 92.74)
SP	28.34	26.25	21.60	25.94	45.08	57.35
	(3.89 – 90.18)	(4.40 – 81.52)	(3.71 – 71.62)	(4.53 – 69.78)	(12.08 – 79.89)	(24.61 – 84.33)
H	15.76	51.14	63.22	64.94	64.51	60.38
	(7.67 – 27.50)	(33.15 – 66.25)	(39.68 – 79.52)	(41.03 – 83.46)	(40.11 – 84.32)	(33.53 – 83.12)
	Panel (b): VAR with $X = C - I$					
	Forecast horizons					
	1	4	8	16	40	120
TFP	0.00	0.67	8.31	10.41	6.63	1.25
	(0.00 – 0.00)	(0.41 – 5.98)	(2.31 – 22.82)	(2.85 – 30.85)	(2.51 – 35.91)	(1.69 – 47.34)
C	89.55	85.04	78.64	74.27	63.75	43.72
	(50.50 – 96.63)	(48.97 – 91.64)	(43.31 – 89.72)	(42.58 – 89.78)	(37.11 – 90.05)	(25.49 – 91.50)
SP	13.49	11.48	8.20	6.29	5.60	6.17
	(1.72 – 88.66)	(1.97 – 82.49)	(2.29 – 70.79)	(2.41 – 67.68)	(2.09 – 67.15)	(3.16 – 66.81)
I	0.47	15.07	18.05	18.03	21.62	27.97
	(0.04 – 4.57)	(4.98 – 32.61)	(4.70 – 44.22)	(5.61 – 45.39)	(9.08 – 49.15)	(15.63 – 61.45)

Note: Shares of forecast error variances at different horizons explained by the shocks identified by BP restrictions in the stationary VAR. Results are generated from VARs with 4 lags. 80 percent confidence sets reported in parentheses below the point estimates.

References

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

Beaudry, P., D. Nam, and J. Wang (2011, December). Do mood swings drive business cycles and is it rational? NBER Working Papers 17651, National Bureau of Economic Research, Inc.

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

Blanchard, O. J. and D. Quah (1989, September). The dynamic effects of aggregate demand and supply disturbances. *The American Economic Review 79*(4), 655–673.

Francis, N., M. T. Owyang, J. E. Roush, and R. DiCecio (2012, June). A flexible finite-horizon alternative to long-run restrictions with an application to technology shocks. Working Papers 2005-024, Federal Reserve Bank of St. Louis.

Hamilton, J. D. (1994). *Time-Series Analysis*. Princeton, NJ: Princeton University Press.