Paul Beaudry and Bernd Lucke: "Letting Different Views about Business Cycles Compete"

Discussion prepared by Stephanie Schmitt-Grohé
Columbia University*

June 17, 2009

This paper is a contribution to the empirical literature on the estimation of the sources of business cycles. It uses various five-variable vector-error-correction models (VECM) to estimate the relative importance of anticipated total factor productivity (TFP) shocks, unanticipated TFP shocks, investment specific technology shocks, preference shocks, and monetary policy shocks. An innovation relative to the large related literature on structural vector autoregression based estimation of the sources of fluctuations is the focus on anticipated TFP shocks and on imposing cointegration relationships. Further, Beaudry and Lucke use time series on TFP, the relative price of investment, stock prices, federal funds rates, and a measure of aggregate activity in their estimation. This set of observables is slightly different than that used in the related literature. The main finding of the paper is that anticipated TFP shocks explain the majority of fluctuations in aggregate activity and stock prices at business cycle frequencies in the United States.

Many authors have studied the question what the sources of short-run fluctuations are. Yet this fundamental question in macroeconomics remains largely unresolved. John Cochrane (1994) in a piece written for the Carnegie Rochester Conference Series on Public Policy starts his article on this topic as follows: *"What shocks are responsible for economic fluctuations? Despite at least two hundred years in which economists have observed fluctuations in economic activity, we still are not sure."* Fifteen years later in business cycle research this statement

is still a valid description of the state of the literature.

Cochrane interpreted the findings of his (1994) study as suggesting that contemporaneous shocks to technology, money, credit, and oil cannot account for the majority of observed aggregate fluctuations.[1] More recent SVAR based papers using long-run restrictions such as Galí and Rabanal (2004) find, like Cochrane, a small role for permanent technology shocks in accounting for business cycle variations in hours and output. In table 2 of their paper Galí and Rabanal report that the share of variance due to technology shocks lies between 7 and 37 percent for output and between 5 and 36 percent for hours. Most importantly under their favored interpretation, the technology shock accounts for less than 10 percent of the variance of output and hours. They therefore conclude that *"Nevertheless, it is safe to state that the bulk of the evidence reviewed in the present paper provides little support for the initial claims of the RBC literature on the central role of technological change as a source of business cycles."*

On the other hand, there are papers presenting evidence that suggests that technology shocks are the major source of fluctuations, and the Beaudry and Lucke paper fits into this group. For example, the empirical paper of Fisher (2006), using SVAR methods, comes to the conclusion that neutral and investment specific *"technology shocks account for 73 percent of hours and 44 percent of outputs business cycle variation before 1982, and 38 percent and 80 percent afterward. The shocks also account for more than 40 percent of hours and 58 percent of outputs forecast errors over a three- to eight-year horizon in both samples. The majority of these effects are driven by the investment shocks."* Using Bayesian methods to estimate a dynamic stochastic general equilibrium model, Smets and Wouters (2007) find that at least 30 percent of the forecasting error variance of output is attributable to a combination of neutral and investment specific technology shocks, with the majority of this share explained by neutral technology shocks. Justiniano et al. (2008), like Smets

---

[1]At the same time, Cochrane showed that VARs estimated using artificial data from an RBC model driven by contemporaneous and news shocks to technology produce responses to consumption shocks that resemble the corresponding responses implied by VARs estimated on actual U.S. data. And thus his paper is often cited as one of the first to revive the idea of Pigou or news-driven business cycles.

and Wouters using Bayesian estimation of a dynamic stochastic general equilibrium model, find an even larger share of fluctuations driven by technology shocks. Contrary to Smets and Wouters, however, Justiniano et al. find that most of the output variance is accounted for by the investment specific technology shock rather than the neutral technology shock. Justiniano et al. attribute their finding of a larger role for the investment specific shock to data differences, such as differences in the treatment of inventories and consumer durables. These differences in the definition of the data can increase the estimated share of the variance of output due to investment specific technology shocks at business cycle frequencies from 18 to 53 percent for output and from 21 to 61 percent for hours.[2]

The paper of Beaudry and Lucke is most closely related to Beaudry and Portier (2006). In that paper, Beaudry and Portier introduce a novel identification scheme to estimate (in the context of a VECM framework) anticipated TFP shocks. Most of the analysis is carried out for bivariate systems of TFP and stock prices. Under one identification scheme the news TFP shock is that shock that does not affect TFP contemporaneously under the other scheme the news TFP shocks is the one that has a long-run effect on TFP. Beaudry and Portier show that the correlation between the news TFP shock series identified by these two alternative schemes is very high and that impulse responses to them of measures of economic activity are quite similar. Therefore, Beaudry and Portier conclude that the common component of these two shocks represents an anticipated TFP shock. Most importantly for the relation to the paper of Beaudry and Lucke is the fact that Beaudry and Portier show that the so identified news TFP shock explains more than 50 percent of the forecast error variance of consumption, hours, investment, and output (measured as the sum of investment and consumption).

---

[2]One caveat to the results of Justiniano et al. is that their estimates imply a volatility for the relative price of investment, which they exclude from the set of observables, that is significantly larger than the observed standard deviation of this variable.

1 Interpretation of Structural Disturbances

The current paper by Beaudry and Lucke extends the work of Beaudry and Portier by moving away from bivariate SVAR systems to larger ones. Within a larger SVAR/VECM system the identification assumption of Beaudry and Portier must be modified. Specifically, Beaudry and Lucke estimate a VECM model of the form:

$$\Delta y_t = \alpha \beta' y_{t-1} + \Gamma(L) \Delta y_{t-1} + B \epsilon_t,$$

where the vector y_t contains period t observations for TFP, the relative price of investment, stock prices, hours, and the Federal Funds rate, β denotes the co-integration vector, $\Gamma(L)$ denotes a lag-polynomial, and $\epsilon_t = [\epsilon_t^1; \epsilon_t^2; \epsilon_t^3; \epsilon_t^4; \epsilon_t^5]$ denotes the vector of structural shocks which are the focus of interest. To identify the VECM, in particular the matrix B and the vector ϵ_t, Beaudry and Lucke impose the following identification restrictions. Identification restriction A1 says that only ϵ_t^1 may have a contemporaneous effect on TFP. Therefore, ϵ_t^1 is labeled the TFP shock. Implicitly it is therefore assumed that TFP is measured without error and that TFP is exogenous. Identification restriction A3 says that ϵ_t^5 does not affect economic activity contemporaneously and is therefore interpreted as a monetary policy shock. Identification restriction A2 imposes that ϵ_t^4 and ϵ_t^5 have no long run effect on TFP. Under identification scheme 1, denoted ID1, ϵ_t^3, ϵ_t^4, and ϵ_t^5 are assumed to have not contemporaneous effect on the price of investment. Identification assumptions A1 and B1 then imply that ϵ_t^2 must be the contemporaneous innovation to the relative price of investment. While in principle this identification scheme allows for ϵ_t^4 to represent an anticipated temporary TFP shock or an anticipated temporary or permanent shock to the relative price of investment, the estimation results show that ϵ_t^4 has very little effect on either TFP or the relative price of investment and thus it is unlikely that it represents a technology shock. Beaudry and Lucke therefore interpret it as a preference shock.

Under identification scheme 2, ID2, assumption B1 is replaced by imposing that ϵ_t^2 has

Figure 1: Share of Forecast Error Variance Due to $\epsilon_{3,t}$, VECM with ID 2
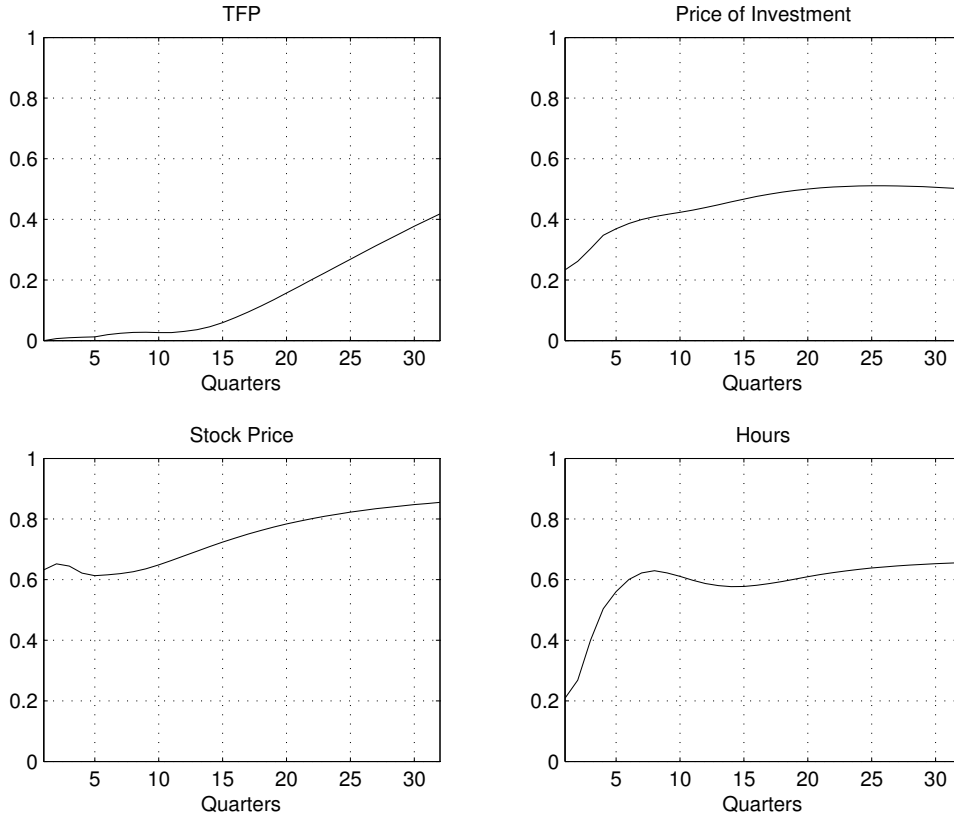
no long run effect on TFP. This would still allow for the possibility that ϵ_t^2 is an anticipated temporary TFP shock. But the estimation assigns almost no role to ϵ_t^2 in accounting for the variance of TFP. Identification scheme 2 leaves open the possibility that ϵ_t^2, ϵ_t^3, or ϵ_t^4 affect the price of investment contemporaneously and thus could be called investment specific shocks. Beaudry and Lucke, however, interpret ϵ_t^3 as an anticipated TFP shock. The reason is that the identification assumption imposes that ϵ_t^3 does not affect TFP on impact—thus it could not be an unanticipated TFP shock,— and that the estimation yields that (at horizons not shown in Figures 3 or 5, namely 60 quarters) ϵ_t^3 explains about 3/4 of the forecasting error variance of TFP under ID2. However, for horizons of 32 quarters or less (the time horizon shown in figure 5) ϵ_t^3 explains less than 20 percent of TFP and thus the interpretation as a TFP shock is less immediate. I want to entertain whether one could with equal plausibility interpret ϵ_t^3 as an investment specific shock. As shown in figure 1, under identification scheme

2, ϵ_t^3, explains between 40 and 60 percent of the forecasting error variance of the price of investment for forecasting horizons between 8 and 32 quarters. And ϵ_t^3 explains less of the FEV of TFP than that of the price of investment at any of these forecasting horizons. This might one lead to interpret ϵ_t^3 as an investment specific technology shock rather than as maintained by Beaudry and Lucke a TFP shock. The figure further shows that ϵ_t^3 explains 60 percent of the forecasting error variance of hours and stock prices for forecasting horizons greater than 4 quarters. And thus one might be lead to conclude that an investment specific technology shock is the most important source of fluctuations in stock prices and real activity. This interpretation of ϵ_t^3 would therefore be less at odds with the findings of Fisher (2006) on the importance of investment specific technology shocks.

# 2 Are anticipated shocks identified in the vector error correction model?

To be convinced by the interpretation given to the papers findings regarding the importance of anticipated shocks one needs to be convinced that the empirical strategy employed indeed is able to identify such shocks. For it is not immediately obvious that this is the case, in what follows I will discuss some concerns one may have regarding the ability of SVAR/VECM methods to identify anticipated shocks.

Beaudry and Lucke address the question of identification by presenting a theoretical model of the business cycle and checking whether the empirical identification strategy they employ, i.e., a VECM analysis, would uncover the true structural shocks from data generated by this theoretical model. In particular, Figure 4 (of the appendix of Beaudry and Lucke) shows the population FEVD of hours in the theoretical model with respect to the four structural shocks of the theoretical model: the unanticipated innovation to the growth rate of TFP, $\epsilon_{A,t}$, the eight-quarter anticipated innovation to TFP, $\epsilon_{NA,t}$, the unanticipated innovation to the growth rate of the price of investment, $\epsilon_{Z,t}$, and the unanticipated innovation to

the preference shock, $\epsilon_{\psi,t}$. (This is a real model and hence the fifth structural shock, which had the interpretation of a monetary policy shock, is dropped.) Then Figures 4 through 6 of the Beaudry and Lucke appendix show the FEVD one would obtain were one to feed data generated by the calibrated real-business-cycle model through the VECM machinery and impose the various identification schemes labeled ID1 through ID3. In figure 2, I repeat this exercise for the case of identification scheme ID1. One difference between my figure 2 and Beaudry and Lucke's figures is that I show the population FEVD implied by the calibrated theoretical model and the FEVD stemming from applying identification scheme ID1 to artificial model generated data in the same graph and for all four variables, that is, total factor productivity, the price of investment, stock prices and hours, whereas Beaudry and Lucke show this only for hours and in two different graphs.[3] The purpose of this exercise is to check whether the VECM identified innovation ϵ_t^3 does indeed explain the same share of variance in all four observables as the anticipated TFP shock, $\epsilon_{NA,t}$, that it is meant to identify. A convincing case that the ID1 scheme, or any other of the identification schemes considered, is able to recover the true structural shocks is incomplete unless it does so for all four variables considered. After all, the fact that ϵ_t^3 is interpreted as an anticipated TFP shock by Beaudry and Lucke is based on the finding that at very long forecasting horizons (60 quarters), longer than those shown in the graphs, it explains a large fraction (60 percent) of the forecasting error variance of TFP. (In the artificial economy, given the calibration of Beaudry and Lucke, the anticipated TFP shock explains 99 percent of the FEV of TFP for horizons greater than 8 quarters.) It follows that one needs to show that the ID1 scheme also picks up a similar share of the variance of these other variables as is true in population. The horizontal axis of each panel of figure 2 shows the forecasting horizon, which takes values between 1 and 32

---

[3]The FEVD from the VECM shown in my figure 2 is the mean of FEVDs performed on 500 simulated data sets with 210 observations each. The simulated time series are length 1210 and I drop the first 1000 observations. I follow the calibration of Beaudry and Lucke by setting $\alpha = 0.64$, $\beta = 0.985$, $\delta = 0.025$, $\delta_0 = 0.035$, $\mu^A = \mu^Z = 1.002$, $\rho_\psi = 0.5$, $\phi_2 = 1.3$, $\phi_0 = \phi_1 = 0$ and $\sigma^{\epsilon_Z} = \sigma^{\epsilon_A} = 0.1$, $\sigma^{\epsilon_{NA}} = 1$, and $\sigma^{\epsilon_\psi} = 0.02$. I measure the stock price as the value of the firm. Letting the value of the firm be denoted by V_t, output by Y_t, wages by w_t, and the marginal utility of income by Λ_t, stock prices can recursively be expressed as: $V_t = Y_t - w_t h_t - I_t/Z_t + \beta E_t \frac{\Lambda_{t+1}}{\Lambda_t} V_{t+1}$.

Figure 2: Forecast Error Variance Decompositions in the Baseline Model: Theoretical versus VECM Estimates
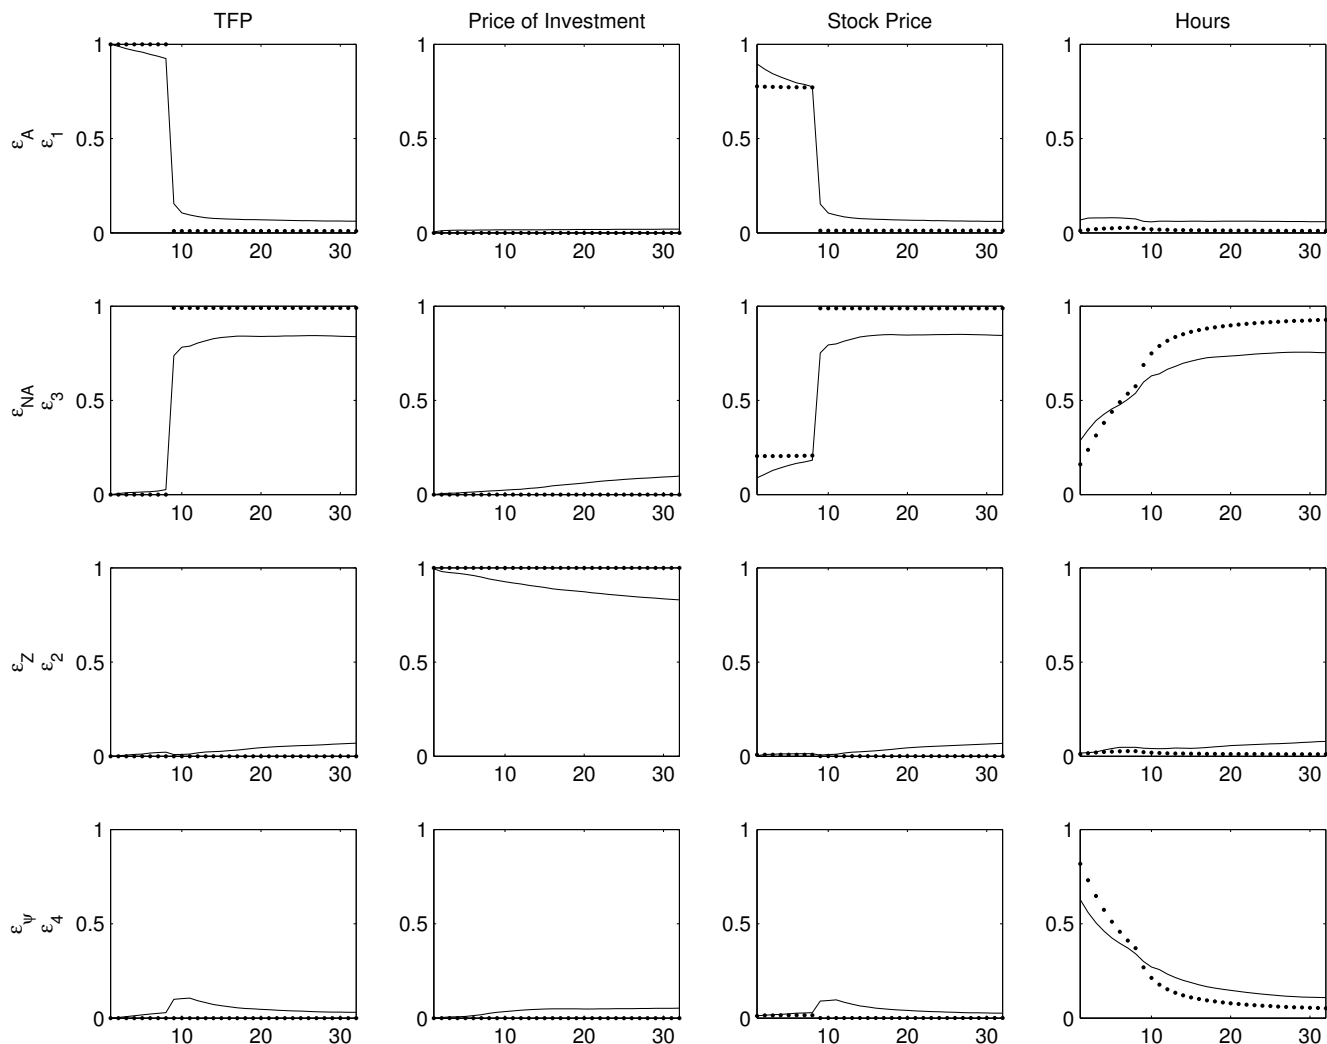
Note. Solid lines show the share of the forecasting error variance for horizons 1 through 32 quarters due to ϵ_t^i, for $i = 1, 2, 3, 4$, which are the error terms identified with scheme ID1 by estimating a VECM on artificial time series simulated from the calibrated theoretical model. Dotted lines show the population forecasting error variance shares due to the true structural shocks $\epsilon_{A,t}$, $\epsilon_{NA,t}$, $\epsilon_{Z,t}$ and $\epsilon_{\psi,t}$ respectively, and were computed from the log-linear approximation to the baseline model.

quarters, the vertical axis measures the share of variance explained by the particular shock considered. The solid line corresponds to the FEVD implied by the SVECM and the dotted line corresponds to the population FEVD implied by the log-linearized approximation to the calibrated model. If the identification strategy was perfect, the solid line and the dotted line should be identical to each other. The figure shows that the SVECM delivers FEVDs that are very close to the population ones and hence suggests that the SVECM, with identification scheme ID1, is able to identify the contribution of all four structural shocks quite closely—as argued by Beaudry and Lucke. In particular, in the theoretical model most of the variance of hours of work, the measure of economic activity used by Beaudry and Lucke, at short horizons is due to the preference shock, $\epsilon_{\psi,t}$, and ϵ_t^4 of the VECM reproduces this fact. Further, at longer forecasting horizons the most important source of fluctuations in hours are 8-quarter anticipated TFP shocks and the SVECM identified innovation, ϵ_t^3, is consistent with this feature of the theoretical model.

I next consider a small variation in the model to see how well the SVECM methodology identifies the structural shocks in a slightly more complicated but empirically equally realistic environment. The only change I introduce is that the relative price of investment now is also subject to anticipated disturbances. And to keep it similar to the structure assumed by Beaudry and Lucke for anticipated TFP shocks, I will assume that the innovations to the investment price growth rate are also anticipated 8 quarters. Formally, this yields a process for the relative price of investment of the form.

$$\ln Z_t - \ln Z_{t-1} = \ln \mu^Z + \epsilon_{Z,t} + \epsilon_{NZ,t-8},$$

where $\epsilon_{NZ,t-8}$ denotes the eight-quarter anticipated innovation to the growth rate of investment. The innovation $\epsilon_{NZ,t-8}$ enters the information set of private agents in period $t-8$ and thus will lead to changes in the endogenous variables included as observables, namely, the logarithm of hours and the growth rate of the stock price, already in period $t-8$, but will

only materialize in an observed change in the price of investment eight periods after agents learn about it. I calibrate the structural parameters of model as before changing only the standard deviations of the exogenous shocks as follows: $\sigma^{\epsilon Z} = \sigma^{\epsilon NZ} = \sigma^{\epsilon A} = \sigma^{\epsilon NA} = 1$ and $\sigma^{\epsilon_\psi} = 0.1$. Under this calibration of the relative volatilities, TFP is in equal parts due to surprise and anticipated shocks and the same is true for the relative price of investment. As we have seen in the previous exercise, stock prices respond mainly to TFP shocks (under the assumed calibration) and hence stock prices will be almost in equal parts be explained by surprise and anticipated TFP movements. I chose this calibration so that hours are in the long run almost in equal parts driven by all 5 shocks. It turns out that under this calibration in the short run preference shocks are the most important source of fluctuations. As before, I create 500 artificial time series of length 1,210, drop the first 1000 observations, and subject each of the 500 data sets to the SVECM procedure with the ID1 identification scheme and perform the FEVD. Now that there are five structural shocks and the VECM methodology only can identify four it is less clear what the identification restrictions will uncover. Identification assumption A1 of Beaudry and Lucke imposes that ϵ_t^1 is the only shock affecting TFP contemporaneously suggesting that it identifies $\epsilon_{A,t}$. By identification assumption 2, ϵ_t^4 cannot have a long run effect on TFP, thus leaving only the possibility that it is $\epsilon_{Z,t}$, $\epsilon_{NZ,t}$ or $\epsilon_{\psi,t}$, or a combination thereof. Further identification assumption B1 ensures that neither ϵ_t^3 nor ϵ_t^4 have a contemporaneous effect on the price of investment. It follows that ϵ_t^2 is likely to identify $\epsilon_{Z,t}$ and because ϵ_t^4 cannot have a long-run effect on TFP, only ϵ_t^3 has a chance of identifying $\epsilon_{NA,t}$. Finally, this leaves ϵ_t^4 to identify either $\epsilon_{NZ,t}$ or $\epsilon_{\psi,t}$ or some combination thereof. Figure 3 presents the FEVD results. As in figure 2, each panel presents with a solid line the mean of the FEVD obtained from applying the VECM methodology to the artificial data sets and with a dotted line the population FEVD implied by the theoretical model. The figure shows that in this economy, it is no longer the case that the structural disturbances identified by means of the VECM methodology identify the structural shocks of the real-business-cycle model well. The VECM methodology delivers large discrepancies

Figure 3: Forecast Error Variance Decompositions in Model with Anticipated Investment Specific Shocks: Theoretical versus VECM Estimates
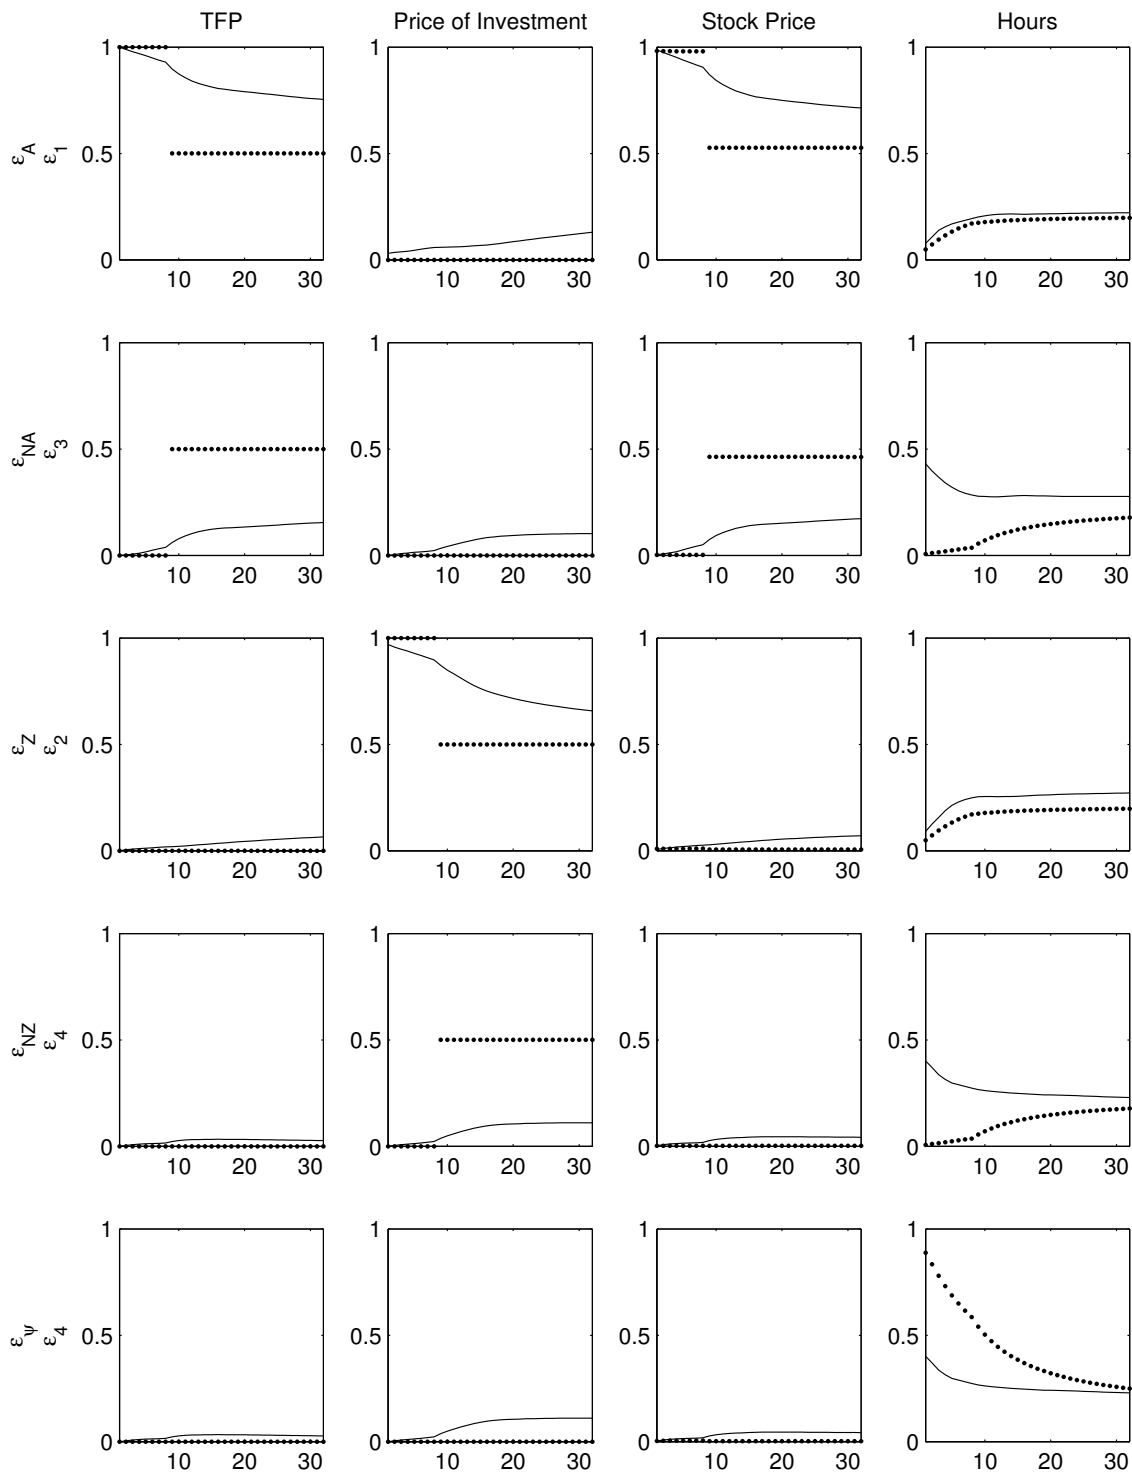
Note. Solid lines show the mean share of the forecasting error variance for horizons 1 through 32 quarters due to ϵ_t^i, for $i = 1, 2, 3, 4$, which are the error terms identified with scheme ID1 by estimating a VECM on artificial time series simulated from the calibrated theoretical model. Dotted lines show the population forecasting error variance shares due to the true structural shocks $\epsilon_{A,t}$, $\epsilon_{NA,t}$, $\epsilon_{Z,t}$, $\epsilon_{NZ,t}$ and $\epsilon_{\psi,t}$ respectively, and were computed from the log-linear approximation to the theoretical model with anticipated investment specific shocks.

in the FEVD of TFP, the price of investment, in particular to news shocks. Interestingly, in this example, it happens that the share of variations in TFP explained by anticipated TFP shocks are estimated by the VECM methodology to be much smaller than the population one. But most importantly the figure shows that the size of the contribution of news TFP shocks and news investment price shocks to economic activity, when identified using the VECM methodology, is very different than the true or population one. The VECM methodology fails to capture that the share of anticipated investment specific shocks in the FEV of the relative price of investment is fifty percent. The VECM assigns equal importance to the anticipated TFP shock and the anticipated investment specific shock in explaining the FEV of the relative price of investment. This case provides an example of a situation in which the VECM methodology fails to correctly identify the importance of competing sources of business cycles.

Identification and invertibility

Rather than comparing VECM and true FEVD compositions one could check for identification by asking whether the theoretical model with anticipated shocks gives rise to a VAR representation in the observable variables. Consider the baseline model without anticipated investment specific shocks shown in figure 2. Note that even figure 2 contains some, albeit small, differences between the true population variance decompositions and those implied by the SVECM methodology. This discrepancy could be due to sampling uncertainty or due to the fact that the particular theoretical model considered fails to have a representation of the type implicitly assumed in the VECM analysis and given in equations (1) and (2) of the body of the Beaudry and Lucke paper. In particular, letting y_t denote the vector of observables, that is, the logarithm of TFP, the logarithm of the relative price of investment, the logarithm of hours, and the logarithm of the stock price, ($y_t = [\ln A_t; \ln Z_t; \ln SP_t; \ln N_t]$), underlying the VECM analysis is the assumption that the vector y_t has a VAR representation. A log-linear approximation to the solution of the theoretical model takes the form

$\hat{\tilde{y}}_t = g_x \hat{x}_t$, where g_x is a 4 by nx matrix relating a stationary transformation of the observable variables, denoted $\tilde{y}_t$, to the vector of stationary state variables, denoted x_t, which in turn consists of observable and unobservable variables and has length n_x. The state vector evolves over time according to $\hat{x}_{t+1} = h_x \hat{x}_t + \eta \epsilon_{t+1}$, where h_x is an n_x by n_x matrix and η an n_x by 4 matrix. The 4×1 vector ϵ_t contains the four structural shocks. In the economy considered here $\epsilon_t = [\epsilon_{A,t}; \epsilon_{NA,t}; \epsilon_{Z,t}; \epsilon_{\psi,t}]$. A hat over a variable denotes log-deviations from the steady state.

One possible starting point would be to let the vector of observables $\tilde{y}_t$ consists of the growth rate of TFP, the growth rate of the price of investment, the growth rate of the stock price, and the logarithm of the level of hours, that is, $\tilde{y}_t = [\Delta \ln A_t; \Delta \ln Z_t; \Delta \ln SP_t; \ln N_t]$, . At first sight a natural strategy appears to inquire whether there exists a VAR representation for $\tilde{y}_t$ in which the VAR errors are indeed ϵ_t. One can answer this question applying the methods described in for example, Fernández-Villaverde et al. (2007). But the answer to this question should be no, for in the theoretical model there is a co-integrating relationship between the levels of the stock price, the price of investment and TFP. Therefore, the differences of these three variables, that is, $[\Delta \ln A_t; \Delta \ln Z_t; \Delta \ln SP_t;]$, should not have a VAR representation. This is the reason afterall why Beaudry and Lucke adopt a VECM rather than a VAR model in differences.

Alternatively, consider the following vector of stationary transformations of our four observables, $\tilde{y}_t = [\Delta \ln A_t; \Delta \ln Z_t; \ln SP_t/X_t^Y; \ln N_t]$, where X_t^Y denotes the trend in output, which is given by $X_t^Y \equiv Z_t^{\frac{1-\alpha}{\alpha}} A_t^{\frac{1}{\alpha}}$. Let $sp_t \equiv SP_t/X_t^Y$. Then sp_t is stationary and its natural logarithm is equal to $\ln(sp_t) = \ln sp^* + \ln SP_t - \frac{1-\alpha}{\alpha} \ln Z_t + \frac{1}{\alpha} A_t$, where sp^* is the non-stochastic steady state value of sp_t. In this case, if we were able to show that $\tilde{y}_t$ has a VAR representation, then we would conclude that the levels of the observables also have a VAR representation, and thus we would have shown that indeed estimates from a VECM model should be able to recover the true structural shocks, ϵ_t. Following Fernández-Villaverde et al. (2007), a model with this structure is invertible, i.e., has a VAR representation of the

form $\tilde{y}_{t+1} = A(L)\tilde{y}_t + B\epsilon_t$ only if all the eigenvalues of the matrix $h_x - \eta(g_x\eta)^{-1}g_xh_x$ are less than one in modulus. I perform a numerical check of this condition for the model economy under consideration and find that the invertibility condition is violated. In particular, I find that more than six eigenvalues of this matrix are greater than one, thus implying that the model fails to have a VAR representation. But in the absence of invertibility, it is impossible to interpret the residuals of the VECM as the true shocks hitting the model economy. Note that invertibility fails here despite the fact that we have four observables and four structural shocks and further that the matrix $(g_x\eta)$ is invertible.

One reason for the failure of invertibility could be the large number of unobservable state variables that emerge when an 8-quarter anticipated innovation is considered. If this were the case, this would support the view that VECM methods are not well suited to identify news shocks. I explore this hypothesis by eliminating the anticipated innovation to TFP by setting $\sigma_{\epsilon_{NA}} = 0$. Then the model is driven by three shocks only and we have $\epsilon_t = [\epsilon_{A,t}; \epsilon_{Z,t}; \epsilon_{\psi,t}]$. To have any hope of the model having a VAR representation in levels, we need to thus consider only 3 observables. I eliminate hours from the vector of observables and set: $\tilde{y}_t = [\Delta \ln A_t; \Delta \ln Z_t; \ln SP_t/X_t^Y]$. For this economy I again compute the matrices: h_x, g_x, and η. I first check whether $g_x\eta$ is full rank and find that it is. I then construct, as before following Fernández Villaverde et al. (2007), $h_x - \eta(g_x\eta)^{-1}g_xh_x$ and calculate its eigenvalues. I find that all eigenvalues are less than one in modulus. It follows that the model without news shock has a VAR representation in levels and therefore the VECM methodology should be able to identify the true structural shocks. Similarly, when I eliminate stock prices from the vector of observables and set: $\tilde{y}_t = [\Delta \ln A_t; \Delta \ln Z_t; \ln N_t]$, I can show that the theoretical model is invertible, that is, it has a VAR representation in $\tilde{y}_t$.[4] These results demonstrate that at least in the example economy considered here it is the presence of news shocks that lead to the violation of the invertibility condition. I regard these findings as further evidence

---

[4]One could also eliminate the preference shock, $\epsilon_{\psi,t}$, and let $\epsilon_t = [\epsilon_{A,t}; \epsilon_{Z,t}; \epsilon_{NA,t}]$. Again one can show that the theoretical model fails to have a VAR representation for the case that the observables are $\tilde{y}_t = [\Delta \ln A_t; \Delta \ln Z_t; \ln SP_t/X_t^Y]$ as well as for the case that $\tilde{y}_t = [\Delta \ln A_t; \Delta \ln Z_t; \ln N_t]$.

that VECM methods may not be well suited to the identification of news shocks even in environments where they provide a valid identification of unanticipated shocks.

3 Alternative empirical strategies for the estimation of anticipated shocks

Given the econometric challenges to the identification of news shocks by means of SVAR/VECM methods that I have just documented, some recent authors have pursued alternative empirical strategies to estimate the importance of anticipated shocks as a source of business cycles. Schmitt-Grohé and Uribe (2008), for example, argue that likelihood based methods provide a promising approach to the estimation of the importance of anticipated shocks. Likelihood-based methods avoid the problems that the VECM/SVAR based empirical literature on the importance of news shocks has run into for it does not require the underlying dynamic stochastic general equilibrium model to have a VAR representation in the observable variables. That is, it can be applied even when invertibility fails. Furthermore, likelihood-based methods allow to estimate what type of anticipated shock is important (as we saw above the VECM approach could not tell apart well anticipated TFP and anticipated investment specific shocks) and they allow to estimate how many quarters in advance the main drivers of business cycles are anticipated. In the VECM approach all we have is the distinction between an innovation that affects the exogenous fundamental contemporaneously (an unanticipated shock) and innovations that are learned today and that will affect the exogenous fundamental in the future (an anticipated shock). But the VECM methodology is by construction mute about the anticipation horizon. In Schmitt-Grohé and Uribe (2008) we perform a structural Bayesian estimation of the contribution of anticipated shocks to business cycles in the postwar United States in the context of a real-business-cycle model, which is slightly more complex than that considered by Beaudry and Lucke. We assume four real rigidities: investment adjustment costs, variable capacity utilization, habit formation in consumption,

Table 1: Share of Variance Explained by Anticipated Shocks

	Output Growth	Consumption Growth	Investment Growth	Hours
Mean Share	0.70	0.85	0.58	0.68
90-percent interval				
5 Percent	0.63	0.76	0.50	0.58
95 Percent	0.77	0.90	0.66	0.76

Source: Schmitt-Grohé and Uribe (2008), Table 4.

and habit formation in leisure and allow business cycles to be driven by permanent and stationary neutral productivity shocks, permanent investment-specific shocks, and government spending shocks. Each of these driving forces is buffeted by four types of structural innovations: unanticipated innovations and innovations anticipated one, two, and three quarters in advance. We find that anticipated shocks account for more than two thirds of predicted aggregate fluctuations. Table 1 summarizes our findings. Our estimation uses U.S. data on output, hours, investment, consumption, government purchases, and the relative price of investment for the period 1955:Q1 to 2006:Q4, which is very similar to the sample period considered in Beaudry and Lucke. Table 1 shows that according to our estimates 68 percent of the population variance of hours is due to anticipated shocks. We further show that the forecasting error variance of hours explained by anticipated shocks increases with the forecasting horizon from 20 percent at a forecasting horizon of 2 quarters to 60 percent at a forecasting horizon of 32 quarters, which is similar to the numbers reported in Beaudry and Lucke. Table 2 compares the decomposition of forecasting error variances at horizon 32 quarters found by Beaudry and Lucke and those found by Schmitt-Grohé and Uribe (2008). There are many differences between the two studies, the most important one being that Beaudry and Lucke apply an atheoretical VECM estimation whereas Schmitt-Grohé and Uribe estimate, using Bayesian methods, a dynamic stochastic general equilibrium model. Both studies use U.S. data on the price of investment, output, consumption, investment and hours. Beaudry and Lucke use in addition data on total factor productivity, stock prices,

Table 2: Percent Share of Variance of 32-qtr Forecasting Error Due to Anticipated Shocks

Variable	Beaudry & Lucke	Schmitt-Grohé & Uribe
Output	63	70
Consumption	59	85
Investment	61	57
Hours	71	59

Note: Variance decompositions for the column labeled Beaudry & Lucke are based on author's VECM estimation and should match the information contained in Figure 11 of Beaudry and Lucke. Variance decompositions for the column labeled Schmitt-Grohé & Uribe are taken from table 7 of Schmitt-Grohé and Uribe (2008). THESE AUTHORS report FEVD for growth rates, with the exception of hours, which is in log-levels, and perform FEVD at the mean of the posterior distribution of the estimated structural parameters.

and interest rates. Further, Beaudry and Lucke allow one measure of aggregate activity to enter the estimated system at the time. In Schmitt-Grohé and Uribe information on output, investment, consumption, and hours is used at the same time. In addition, Schmitt-Grohé and Uribe use data on government purchases. Table 2 shows that despite these many differences the estimated shares of forecast error variances explained by anticipated shocks are rather similar across the two studies. Both studies suggest that anticipated technology shocks explain the majority of short-run fluctuations in U.S. postwar quarterly time series.

References

Beaudry, Paul, and Franck Portier, "Stock Prices, News, and Economic Fluctuations," *American Economic Review 96,* September 2006, 1293-1307.

Cochrane, John H, "Shocks," *Carnegie-Rochester Conference Series on Public Policy 41,* December 1994, 295-364.

Fernández-Villaverde, Jesús, Juan F. Rubio-Ramírez, Thomas J. Sargent, and Mark W. Watson, "ABCs (and Ds) of Understanding VARs," *American Economic Review 97,* June 2007, 1021-1026.

Fisher, Jonas D. M, "The Dynamic Effects of Neutral and Investment- Specific Technology

Shocks," *Journal of Political Economy 114,* 2006, 413-451.

Galí, Jordi and Pau Rabanal, "Technology Shocks and Aggregate Fluctuations: How Well Does the Real Business Cycle Model Fit Postwar U.S. Data?," *NBER Macroeconomics Annual,* 2004, 225-288.

Justiniano, Alejandro, Giorgio Primiceri, and Andrea Tambalotti, "Investment Shocks and Business Cycles," manuscript, 2008.

Smets, Frank and Rafael Wouters, "Shocks and Frictions in U.S. Business Cycles: A Bayesian DSGE Approach," *American Economic Review 97,* June 2007, 586-606.

Schmitt-Grohé, Stephanie and Martin Uribe, "What's News in Business Cycles?," mimeo, 2008.