Do News Shocks Drive Business Cycles? Evidence from German Data

Thomas Haertel and Bernd Lucke

University of Hamburg

Abstract:

We study the Beaudry and Portier (2006)-hypothesis of delayed-technology diffusion and news-driven business cycles. For German data on TFP and stock prices we find qualitatively similar empirical evidence. Quantitatively, however, an impulse response analysis suggests that a substantial part of the total TFP response is immediate rather than delayed. We relate this to disembodied technological change and noisy data on TFP. Nevertheless, we confirm the technology interpretation of structural shocks by showing that they are Granger-causal for data on patents granted by the German patent agency. We also show that these shocks generate comovement of macro variables at business cycle horizons and account for a sizable share of the forecast error variance of these variables in the medium and long run.

Correspondence:
Thomas Haertel, University of Hamburg, von-Melle-Park 5, D-20146 Hamburg, email: Thomas.Haertel@ wiso.uni-hamburg.de
Bernd Lucke, University of Hamburg, von-Melle-Park 5, D-20146 Hamburg, email: lucke@econ.uni-hamburg.de

1   Introduction

In a recent paper, Beaudry and Portier (2006) have emphasized that stock prices may have relevant informational content for understanding macroeconomic fluctuations. New information, their argument goes, may alter expectations about future fundamentals. Forward looking variables such as stock prices will react to changes in expectations much earlier than the foreseen changes in fundamentals affect other macroeconomic time series. In particular, news about technological innovations may affect stock prices instantaneously, but due to an implementation lag, it may take some time until they actually alter total factor productivity (TFP). Thus, our understanding of expectations-driven macroeconomic fluctuations could be enhanced if news about expected changes in fundamentals could be properly identified from, among others, stock market data.

For this purpose, Beaudry and Portier (BP) suggest to use structural vector autoregressions (SVAR). Imposing different identifying restrictions on the estimated lag polynomial of a moving average (MA) representation yields a set of structural shocks for ready comparison. If similar shocks are found under different identifying assumptions, then the type of identifying assumption reveals important information about the way a specific shock hits the economy. This, in turn, allows inference on the validity of competing models and their assumptions.

For instance, in a bivariate vector autoregression for US TFP and stock prices, BP find two almost co-linear shocks under two polar identifying assumptions. The first identifying assumption imposes that there exists a shock which does not alter TFP in the short run, the alternative identifying assumption imposes that there exists a shock which does not alter TFP in the very long-run. The two co-linear shocks are consistent with the former but inconsistent with the latter. Hence, BP conjecture that they represent a technological innovation which affects TFP with considerable delay. However, this technological innovation affects stock prices immediately and may therefore cause expectations-driven fluctuations in consumption and investment.

The idea of a prominent role for technology-related news in macroeconomic fluctuations has recently given rise to quite a few other papers, e.g. Lorenzoni (2006) and Jaimovich and Rebelo (2006). It is therefore very interesting to investigate if the empirical finding of BP in favor of news-driven business cycles is a robust business cycle fact which can be documented for other countries and samples as well. Moreover, one would like to know if there is any kind of direct evidence which supports the interpretation of the identified shocks as being technology shocks.

In the first line of research, Beaudry and Portier (2005) took the lead by repeating their analysis with Japanese data. Here they came up with essentially the same finding as for the US: Two almost co-linear "technology" shocks under alternative identifying assumptions. In this paper, we look at Germany as a third country and present qualitatively similar evidence, with, however, some important quantitative differences: While there is evidence of a gradually increasing response of TFP to "news" shocks, there is also a clearly positive effect on impact. Both types of shocks seem to account for important fractions of the total variance of macroeconomic variables: Shocks on impact play an important role for output, while delayed-diffusion shocks seem to dominate hours.

Going further, we test whether the identified shocks are rightly considered as technology shocks by confronting them with data on patents granted by or applied for at the

German patent agency. For various measures of TFP and patents, we have very robust evidence that the identified shocks Granger-cause patents. Conversely, the identified non-persistent shock in the SVAR-approach is not Granger-causal in any of the specifications we test. This seems to be fairly strong evidence for the hypothesis that the identified shock, which affects TFP on impact and with a delay, is indeed a technology shock.

The sequel of the paper is organized as follows: In Section 2, we illustrate that the BP approach can be used to infer key model properties. We present a modified Long and Plosser (1983) model which allows for multi-period usage of capital goods. We compute stock prices as the discounted sum of expected returns to capital. We then show that a standard specification of TFP shocks has very different implications for structural identification than a delayed-implementation specification. In Section 3, we briefly illustrate the econometric approach and apply it to German data. We use three different measures of TFP in order to check the robustness of our results. In Section 4, we check if these shocks generate macroeconomic comovement and and assess their contribution to forecast error variance of macro variables, Section 5 analyzes Granger-causality between the identified shocks and different measures of patents. Section 6 concludes.

2   The Model

We will illustrate the potential of the BP approach by considering two versions of the same model: One with a standard, instantaneous reaction of TFP to a technological innovation, the other with a delayed response. This is similar to BP (2005). However, their model assumes 100% depreciation on physical capital, which makes it difficult to model stock prices, since, essentially, firms shut down each period. Hence, BP (2005) do not consider stock prices but rather focus on bonds whose price is inversely related to the return on (one-period) capital goods. By contrast, we use a model where the productive use of capital goods extends over many periods and stock prices are computed as the discounted sum of expected returns to capital.

The model is taken from Long and Plosser (1983). We aggregate their model to just a single sector, but extend its production technology to a multi-period setting. Specifically, investment goods I_t can be used for $p+1$ periods until they are completely worn out. The production elasticities of investment of period t-τ is given by a_τ and we can allow for any kind of depreciation schedule by securing $a_\tau > a_{\tau+1} \quad \forall \tau < p$. Labor input is L_t with production elasticity $b>0$, so that constant returns imply $\sum_{\tau=0}^{p} a_\tau + b = 1$. Λ_{t+1} is TFP of period $t+1$ and production is given by

$$Y_{t+1} = \Lambda_{t+1} L_t^b \prod_{\tau=0}^{p} I_{t-\tau}^{a_\tau} \tag{1}$$

or (using small letters to denote logs)

$$y_{t+1} = bl_t + \sum_{\tau=0}^{p} a_\tau i_{t-\tau} + \lambda_{t+1}. \tag{2}$$

The representative agent has a standard intertemporal utility function with subjective discount factor $0 < \beta < 1$ and risk aversion captured by $\sigma > 1$. $\theta_t = e^{\eta_{2t}}$ is a stationary preference shock and the innovation η_{2t} is, for simplicity, white noise with unit variance.

$$U = E_0 \left[\sum_{t=0}^{\infty} \beta^t \left(\ln C_t - \theta_t \frac{L_t^{\sigma}}{\sigma} \right) \right]$$

Maximizing utility under the budget constraint

$$C_t + I_t = Y_t \tag{3}$$

yields first order conditions

$$\frac{I_t}{C_t} = \sum_{\tau=0}^{p} \beta^{\tau+1} a_\tau E_t \left[1 + \frac{I_{t+1+\tau}}{C_{t+1+\tau}} \right] =: \gamma + \sum_{\tau=0}^{p} \beta^{\tau+1} a_\tau E_t \left[\frac{I_{t+1+\tau}}{C_{t+1+\tau}} \right] \tag{4}$$

and

$$\theta_t L_t^{\sigma} = \beta b E_t \left[1 + \frac{I_{t+1}}{C_{t+1}} \right]. \tag{5}$$

Imposing stability, solving (4) forward and using (3) yields

$$C_t = (1 - \gamma) Y_t, \quad I_t = \gamma Y_t. \tag{6}$$

This is the policy function, since Y_t is a state variable. For labor, we compute

$$l_t = \ln L_t = \frac{1}{\sigma} \left(\ln \frac{\beta b}{1 - \gamma} - \eta_{2t} \right). \tag{7}$$

Inserting (6) into (1) we get

$$Y_{t+1} = \Lambda_{t+1} \gamma^a L_t^b \prod_{\tau=0}^{p} Y_{t-\tau}^{a_\tau}, \tag{8}$$

where $a := \sum_{\tau=0}^{p} a_\tau$. Setting $\kappa := a \ln \gamma + \frac{b}{\sigma} \ln \frac{\beta b}{(1 - \gamma)}$ and taking logs we get:

$$y_{t+1} = \kappa + \sum_{\tau=0}^{p} a_\tau y_{t-\tau} + \lambda_{t+1} - \frac{b}{\sigma} \eta_{2t} =: \kappa + a(L) y_{t+1} + \lambda_{t+1} - \frac{b}{\sigma} \eta_{2t} \tag{9}$$

$$y_t = \frac{\kappa}{1-a} + \frac{\left(\lambda_t - \frac{b}{\sigma} \eta_{2t-1} \right)}{1 - a(L)} := \frac{\kappa}{1-a} + A(L) \left(\lambda_t - \frac{b}{\sigma} \eta_{2t-1} \right), \tag{10}$$

where $a(L), A(L)$ are polynomials in the lag operator L and $a(0) = 0, A(0) = 1$.

Net profits are output minus labor and investment costs:

$$\Pi_{t+1} := Y_{t+1} - \frac{\partial Y_{t+1}}{\partial L_t} L_t - I_{t+1} = (1 - b - \gamma) Y_{t+1} \tag{11}$$

Let $SP_t := E_t \left[\sum_{\tau=1}^{\infty} \beta^\tau \Pi_{t+\tau} \right]$ be the stock price[1], i.e. the discounted sum of profits expected in period t. By definition, we have

$$SP_t := E_t \left[\sum_{\tau=1}^{\infty} \beta^\tau \Pi_{t+\tau} \right] = \beta E_t [\Pi_{t+1}] + E_t \left[\beta \sum_{\tau=1}^{\infty} \beta^\tau \Pi_{t+1+\tau} \right] = \beta E_t [\Pi_{t+1}] + \beta E_t [SP_{t+1}].$$

Log-linearizing this equation we get:

$$
\begin{aligned}
sp_t &\approx cnst + (1 - \beta) E_t \pi_{t+1} + \beta E_t sp_{t+1} \\
&= \frac{cnst}{(1 - \beta)} + (1 - \beta) E_t \left(\sum_{\tau=0}^{\infty} \beta^\tau \pi_{t+1+\tau} \right) \\
&=: cnst_1 + (1 - \beta) \sum_{\tau=0}^{\infty} \beta^\tau E_t \left(A(L) \left(\lambda_{t+1+\tau} - \frac{b}{\sigma} \eta_{2t+\tau} \right) \right) \\
&=: cnst_1 - \delta(L) \eta_{2t} + (1 - \beta) \sum_{\tau=0}^{\infty} \beta^\tau \left(E_t A(L) \lambda_{t+1+\tau} \right),
\end{aligned}
\tag{12}
$$

where $\delta(0) \neq 0$.

Let us now consider two different specifications for the stochastic processes driving the model. The standard specification would be a random walk for log TFP:

$$\lambda_t = \lambda_{t-1} + \eta_{1t} \tag{13}$$

Here, η_{1t} is assumed to be white noise with unit variance.

The alternative specification would specify a delayed response of log TFP to permanent technological innovations. Assume that log TFP is the sum of a random walk component ζ and a stationary process v. There are orthogonal, unit-variance-white noise innovations η_1 and η_3 to ζ and v, respectively. The η_3 innovation affects log TFP in the same period in which it becomes known, while we assume that the η_1 innovation affects TFP with a delay of one period. Thus, the stochastics are described by

$$
\begin{aligned}
\lambda_t &= \zeta_{t-1} + v_t \\
\zeta_t &= \zeta_{t-1} + \eta_{1t} \\
v_t &= \rho v_{t-1} + \eta_{3t}, \quad \rho < 1
\end{aligned}
\tag{14}
$$

To preserve the same number of shocks as under the standard specification, we assume that there is no preference shock in the delayed-response specification, $\eta_{2t} = 0 \quad \forall t$.

_____

[1] Stock prices would not depend on future expected profits if we discounted with marginal utilities of consumption.

Under the standard specification we derive from (12)

$$sp_t \approx cnst_1 - \delta(L)\eta_{2t} + \tilde{A}(L)\lambda_t$$
$$\Delta sp_t \approx \tilde{A}(L)(1-L)\lambda_t - \delta(L)(1-L)\eta_{2t}$$
$$= \tilde{A}(L)\eta_{1t} - \delta(L)(1-L)\eta_{2t},$$

where $\tilde{A}(L)$ is a lag polynomial with $\tilde{A}(0) \neq 0 \neq \tilde{A}(1)$.

Hence, the moving average representation is given by

$$\begin{pmatrix} \Delta\lambda_t \\ \Delta sp_t \end{pmatrix} = \begin{pmatrix} 1 & 0 \\ \tilde{A}(L) & -\delta(L)(1-L) \end{pmatrix} \begin{pmatrix} \eta_{1t} \\ \eta_{2t} \end{pmatrix} =: C_1(L) \begin{pmatrix} \eta_{1t} \\ \eta_{2t} \end{pmatrix}. \tag{15}$$

Under the delayed-response specification we know

$$sp_t \approx cnst_2 + \tilde{A}(L)\zeta_t + \tilde{B}(L)v_t$$
$$= cnst_2 + \tilde{A}(L)\zeta_t + \frac{\tilde{B}(L)}{1-\rho L}\eta_{3t}$$

with $\tilde{B}(0) \neq 0 \neq \tilde{B}(1)$. The moving average representation for the first differences is then given by

$$\begin{pmatrix} \Delta\lambda_t \\ \Delta sp_t \end{pmatrix} = \begin{pmatrix} L & \dfrac{(1-L)}{1-\rho L} \\ \tilde{A}(L) & \dfrac{\tilde{B}(L)}{1-\rho L}(1-L) \end{pmatrix} \begin{pmatrix} \eta_{1t} \\ \eta_{3t} \end{pmatrix} =: C_2(L) \begin{pmatrix} \eta_{1t} \\ \eta_{3t} \end{pmatrix}. \tag{16}$$

From (15) and (16) we infer that the first row of both $C_1(1)$ and $C_2(1)$ is $(1,0)$, while the first row of $C_1(0)$ and $C_2(0)$ are very different: $(1,0)$ and $(0,1)$, respectively. This implies that the identification of the structural shocks gives rise to different patterns, too, and these patterns can be used to infer what the true underlying model is. We will discuss this in detail in the next section.

3   The Econometric Approach

Consider empirical time series for log TFP and log stock prices, denoted λ_t and sp_t as before. We assume these are integrated of order one and cointegrated with each other, i.e. $(\Delta\lambda_t, \Delta sp_t)'$ is I(0). Using Wold's decomposition theorem, $(\Delta\lambda_t, \Delta sp_t)'$ can be written in reduced form

$$\begin{pmatrix} \Delta\lambda_t \\ \Delta sp_t \end{pmatrix} = C(L) \begin{pmatrix} u_{1t} \\ u_{2t} \end{pmatrix} \quad \text{with} \quad C(L) := I + \sum_{i=1}^{\infty} C_i L^i \tag{17}$$

and in structural form

$$\begin{pmatrix} \Delta\lambda_t \\ \Delta sp_t \end{pmatrix} = D(L)\begin{pmatrix} \varepsilon_{1t} \\ \varepsilon_{2t} \end{pmatrix} \quad \text{with} \quad D(L) := \sum_{i=0}^{\infty} D_i L^i \,. \tag{18}$$

Identifying the structural shocks ε requires knowledge of D_0. This matrix can be recovered from the estimated $C(L)$ matrices of the reduced form (17), if one restriction is imposed on the parameters of $D(L)$. We follow Beaudry and Portier (2005, 2006) by using two alternative assumptions, which we call the short-run and the long-run restriction. The former postulates that the (1,2) element of $D(0)$ is zero, i.e. the stock market shock ε_{2t} has no effect on TFP on impact. The latter postulates that the (1,2) element of $D(1)$ is zero, i.e. the stock market shock ε_{2t} has no long-run effect on TFP. Let us think of (18) as the representation obtained under the short-run restriction and let (19) be the representation obtained under the long-run restriction:

$$\begin{pmatrix} \Delta\lambda_t \\ \Delta sp_t \end{pmatrix} = \tilde{D}(L) \cdot \begin{pmatrix} \tilde{\varepsilon}_{1,t} \\ \tilde{\varepsilon}_{2,t} \end{pmatrix} \quad \text{with} \quad \tilde{D}(L) := \sum_{i=0}^{\infty} \tilde{D}_i \cdot L^i \tag{19}$$

If the empirical data were generated by the model of Section 2 with standard specification, i.e. by equation (15), the impact matrix

$$C_1(0) = \begin{pmatrix} 1 & 0 \\ \tilde{A}(0) & -\delta(0) \end{pmatrix}$$

would already fulfill the short-run identifying assumption, hence, as structural shocks we would identify $\varepsilon_1 = \eta_1, \quad \varepsilon_2 = \eta_2$. On the other hand, the long-run matrix is

$$C_1(1) = \begin{pmatrix} 1 & 0 \\ \tilde{A}(1) & 0 \end{pmatrix},$$

thus, we would immediately have $\tilde{\varepsilon}_1 = \eta_1, \quad \tilde{\varepsilon}_2 = \eta_2$. The important point is that under both identifying assumptions we would find the same result for ε_1 and $\tilde{\varepsilon}_1$.

If, conversely, the empirical data were generated by the delayed technology specification, the impact matrix would be

$$C_2(0) = \begin{pmatrix} 0 & 1 \\ \tilde{A}(0) & \tilde{B}(0) \end{pmatrix}.$$

In this case, the identifying assumptions imply that $\varepsilon_{1t} = \eta_{3t}, \quad \varepsilon_{2t} = \eta_{1t}$, whereas under the long-run restriction we have

$$C_2(1) = \begin{pmatrix} 1 & 0 \\ \tilde{A}(1) & 0 \end{pmatrix}$$

and hence $\tilde{\varepsilon}_{1t} = \eta_{1t}, \quad \tilde{\varepsilon}_{2t} = \eta_{3t}$. Under this model, we would thus not find the same result for ε_1 and $\tilde{\varepsilon}_1$, but rather we would find that $\varepsilon_2 = \tilde{\varepsilon}_1$. Therefore, if the empirical analysis suggests that $\varepsilon_2 \approx \tilde{\varepsilon}_1$ we may infer that a model with delayed technology response is more appropriate than a standard specification. We now turn to an investigation of this issue for German data.

3.1 Data Description

As in Beaudry and Portier (2006), three different TFP variables are calculated: the standard Solow residual, the Solow residual adjusted for variable capital utilization and a TFP measure following the methodology of Groth et al. (2004) and Oulton (2001).

We have quarterly data from 1970(1) to 2005(2). The simple TFP measure (without capital utilization) is computed from data on GDP, hours worked and annual capital stock data interpolated with constant within-year quarterly growth rates. Under the assumption of constant returns to scale, the observed quarterly labor share is used as the production elasticity of labor. The log of this measure is denoted TFP_D1. Modifying the capital stock data by multiplying with the capacity utilization rate in manufacturing gives a second measure of TFP whose log is denoted TFP_D2.

Computation of the third TFP measure (TFP_D3) takes several criticisms of the standard Solow residual into account. Quality aspects are considered when measuring labor input. Under the assumption of perfect competition, the quality of work is reflected by wages. Therefore, quality adjusted labor input L_t can be constructed as

$$\hat{L}_t = \sum_{i=1}^{n} \left(\frac{s_{it} + s_{it-1}}{2} \right) \cdot \hat{h}_{it} , \tag{20}$$

where hatted variables are growth rates, n constitutes the number of employee categories, s_{it} describes output contribution and h_{it} working hours of group i in period t. In the case of Germany, relevant data in terms of gross earnings exist for salaried employees of the service and manufacturing sectors, and in terms of gross wages for wage earners of manufacturing and agriculture. From microdata, four categories of labor input can be distinguished.

The concept of capital input does not refer to the capital stock but uses a measure of capital services. Different types of assets are weighed by their rental prices to represent the value of services which can be realized under perfect competition. Rental prices MP of asset type j in period t are, in principle, computed as

$$MP_{j,t} = T_{j,t} \cdot \left[\left(r_{j,t} + \delta_j - \hat{p}_{j,t} \right) \cdot p_{j,t} \right], \tag{21}$$

where $r_{j,t}$ are opportunity costs, δ_j is depreciation and $T_{j,t}$ expresses taxation and investment allowances[2].

Stock prices are approximated using the German stock market index DAX deflated by the GDP deflator. We follow Beaudry and Portier (2006) and transform into per-capita values by dividing through the number of employees. The log of the resulting series is denoted DAX1. See the appendix for time series graphs of the constructed series TFP_D1, TFP_D2, TFP_D3 and DAX1.

---

[2] As in Oulton (2001) dwellings are excluded from buildings because dwellings may not conform with strict profit maximizing behavior. Depreciation rates for machinery and buildings are assumed to be 13% and 2.5%. Market prices of both assets result from the ratio of nominal and real values of the respective gross fixed capital formation. Rates of return are computed by the ratio of gross operating surplus and capital stock value. The tax factor is disregarded due to lack of adequate data.

3.2   Testing, Estimation, Identification

The analysis of German data requires particular care due to the change in the territorial definition which took effect in the first quarter of 1991. As unit root and cointegration tests are severely affected by structural breaks, we exploit the fact that the break point is known and that time series for West Germany and reunified Germany overlap in the first years of the nineties. In particular, the level of West German time series is still available in the first quarter of 1991, while growth rates for unified Germany are available since the second quarter of the same year. We construct time series without a unification break by applying the growth rates for unified Germany to West German levels in the first quarter of 1991 and denote the break-adjusted series TFP_D1b, TFP_D2b, TFP_D3b and DAX1b.

Given the absence of structural breaks, a standard Augmented Dickey Fuller (ADF) test is used to check for unit roots. We cannot reject the null hypothesis in any case, thus, we will treat all variables as I(1). To test for a long-run equilibrium relationship of technology and stock market prices we apply the Johansen trace test with an orthogonal trend. We can reject the null of no cointegration with DAX1b for all versions of the TFP variable. This gives rise to the specification of bivariate vector error correction models (VECM) whose coefficient estimates are used to compute the reduced form moving average representation.[3]

Going through the identification under both the short- and the long-run restriction, we find for all three measures of TFP a pronounced positive correlation between ε_2 and $\tilde{\varepsilon}_1$, cf. Figure 1. The correlation is about 0.78 for TFP_D1b and TFPD2b; it is 0.68 for TFP_D3b. While this falls short of the almost perfect co-linearity found by Beaudry and Portier (2005, 2006) for Japan and the US, it certainly tends to favor a delayed-technology-diffusion hypothesis over the standard model.

To explore this issue further, let us look at impulse response functions. We focus here on the basic TFP-version TFP_D1b (without capital utilization) – the results for TFP_D2b (with capital utilization are very similar. The quality-adjusted TFP-measure TFP_D3b will be discussed separately.

The upper panel in Figure 2 shows the impulse responses of the levels of TFP and stock prices to ε_2, i.e. the structural shock identified by the requirement that it does not have an immediate effect on TFP. This shock has a strong immediate effect on stock prices, but about 60% of this effect melts away over the next ten years. If this shock was a gradually diffusing technology shock, then the interpretation of the upper panel would state that stock markets receive new information on a technological innovation, which slowly increases TFP. Stock markets anticipate future profits and prices rise. But as the innovation diffuses through the economy, competition reduces profits again and stock prices adjust to the lower level of remaining future profits. (Note that in BP (2006) the impulse response of stock prices rises quickly in the first two or three quarters and stays virtually constant thereafter. It is not clear how a higher level of profits can be maintained over the long run if the economy is competitive).

_____

[3] All estimation results as well as the results of unit root and cointegration tests are available on request from the authors.

Figure 1: Identified structural residuals: SVAR analysis for a measure of TFP and DAX1b
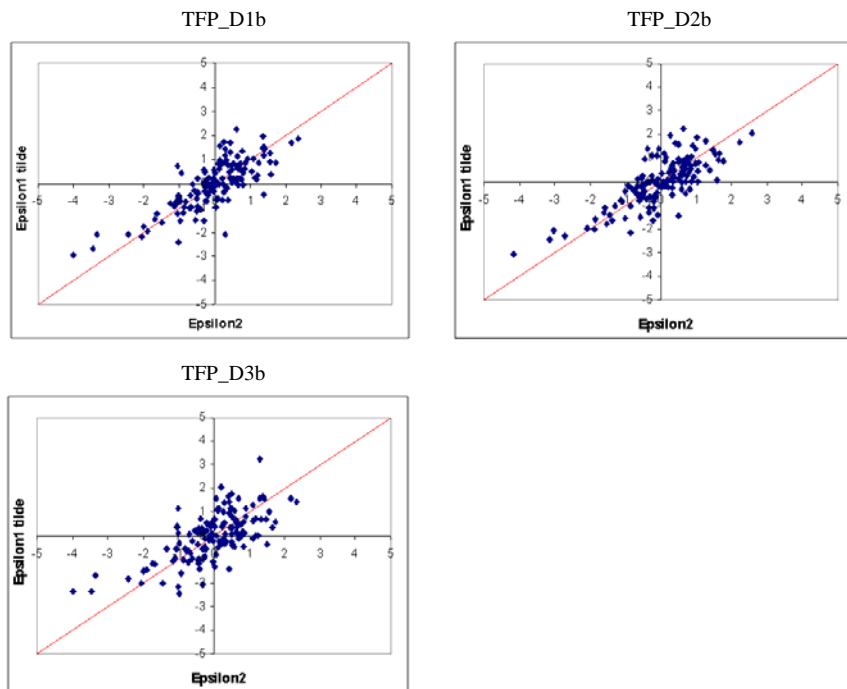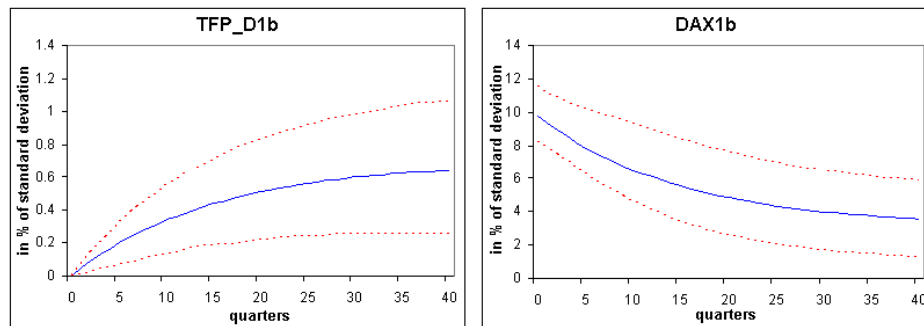

Figure 2: Impulse-response functions for TFP_D1b and DAX1b[a]

a) Short-run restriction (response to ε_2)
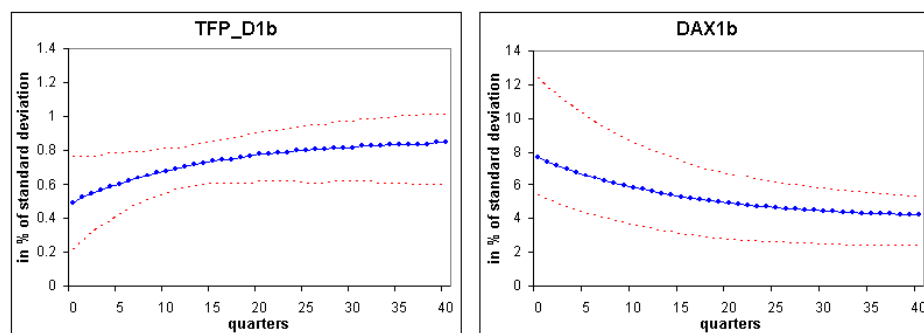
b) Long-run restriction (response to $\tilde{\varepsilon}_1$)

[a]Confidence intervals are obtained by 2,500 replications with the bootstrapping procedure of Hall (1992). They represent the 95% quantiles.

In the lower panel of Figure 2 we see the response to the shock $\tilde{\varepsilon}_1$ identified under the restriction that the other shock $\tilde{\varepsilon}_2$ must not affect TFP in the long-run. If $\tilde{\varepsilon}_1$ and ε_2 were perfectly co-linear, then the impulse responses in the lower panel of Figure 2 should be equal to those in the upper panel. This, however, is not the case. Unlike ε_2, the shock $\tilde{\varepsilon}_1$ is allowed to have an instantaneous effect on TFP. If we consider merely the point estimates, the instantaneous effect is substantial and equals about 60% of the long-run effect. Nevertheless, about 40% of the innovation seems to diffuse gradually. The impulse response of stock prices is slightly weaker on impact – a result which is well in line with the interpretation that part of the technological innovation diffuses rather quickly so that competition brings profits back to normal faster than with a more delayed technology diffusion.

Results for TFP_D2b are very similar, see the appendix. For the quality-adjusted TFP-measure TFP_D3b, however, we find that there is virtually no gradual diffusion of technology to TFP, the effect of $\tilde{\varepsilon}_1$ on TFP and stock prices (lower panel of Figure 3) is almost instantaneous and virtually constant over time. Consequently, there is a large difference between the impulse responses in the upper and lower panel of Figure 3. Given our discussion on Figure 2, this is not completely surprising, since we knew already that the correlation between ε_2 and $\tilde{\varepsilon}_1$ is lower for the quality adjusted measures than for the two standard measures. So does this latter finding, which is based on methodologically more appealing construction of TFP, constitute evidence against the delayed-diffusion model?

Figure 3: Impulse-response functions for TFP_D3b and DAX1b[a]

a) Short-run restriction (response to ε_2)

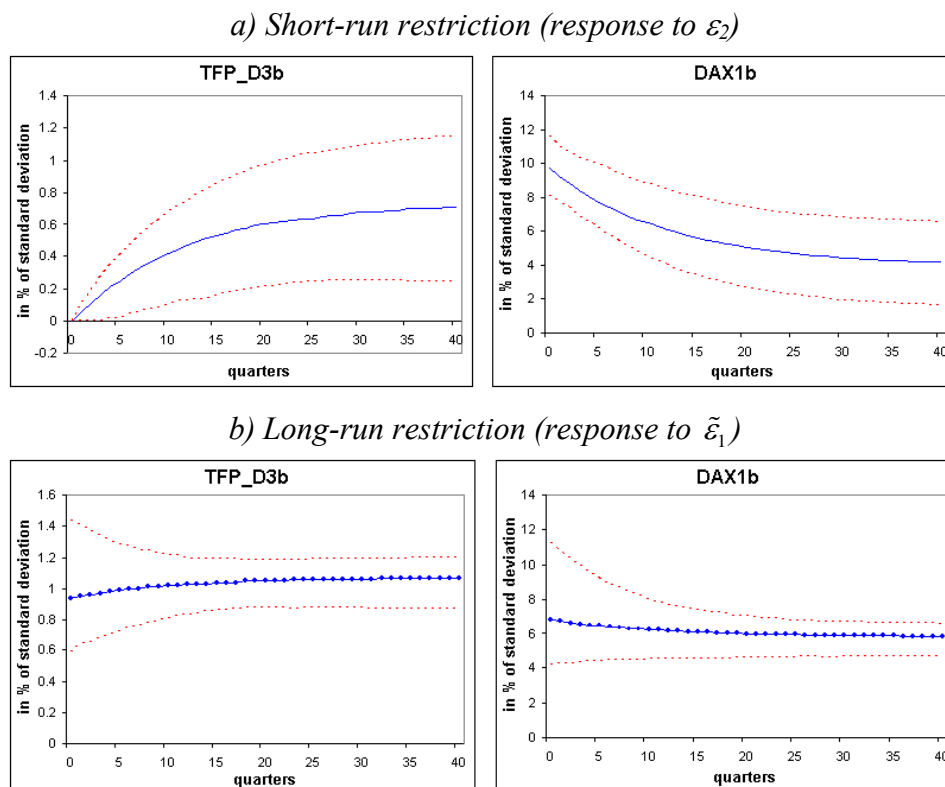*b) Long-run restriction (response to $\tilde{\varepsilon}_1$)*

[a]Confidence intervals are obtained by 2,500 replications with the bootstrapping procedure of Hall (1992). They represent the 95% quantiles.

Not necessarily. Note that technological innovations are typically embodied in either physical or human capital. In fact, this may precisely be the reason why technology diffuses somewhat slowly: It is not the mere existence of an idea, which increases standard measures of TFP, but its implementation in new machinery or its communication in terms of schooling and training and both requires time. Now, as long as a standard measure of TFP is used, advances in the quality of labor and capital are incorrectly attributed to TFP and our SVAR exercises show the gradual diffusion of a new idea into the production factors. With a quality-adjusted measure of TFP, however, our measure comprises just those technological innovations which need not be embodied in either production factor. As far as we can see, there is no compelling reason why such innovations should not be able to diffuse much faster (and possibly almost instantaneously) than embodied technical progress.

One might object to the above argument on the grounds that the quality-adjusted TFP-measure is different from standard measures, but captures nevertheless many movements in TFP in a similar way. If most technological innovations are indeed embodied in either labor or capital (or both), then one would expect that a quality-adjusted TFP-measure has little resemblance with standard measures. However, two counterarguments may be advanced: On the one hand, sticky factor prices may induce incomplete quality adjustment. On the other hand, TFP, being computed as a residual, is a measure of our ignorance (Abramovitz (1962)) and probably includes many non-technology developments which also affect production. Both standard and quality-adjusted measures of TFP share these unknown components and may therefore look more similar than they should.

4   News Shocks and Macroeconomic Fluctuations

The main finding of the preceding section states that stock prices may contain information about technological innovations which diffuse slowly in the economy and affect standard measures of TFP only with considerable lags because they must be embodied in new units of capital. In this section we check if these type of shocks are relevant for macroeconomic fluctuations. We focus on the basic TFP_D1b measure to save space[4].

There are two standard approaches to assess the relevance of a shock for macroeconomic fluctuations. On the one hand, impulse responses may be used to check if the shock generates comovement between the main aggregates such as output, consumption, investment and hours[5]. On the other hand, forecast error variance decompositions (FEVD) may be used to check how important the shock is (in terms of the variance of the dependent variable) relative to other shocks. We will use both approaches here.

In order to do so, the systems we have analysed in Section 3 must be extended to a trivariate setting in which the third variable is one of the main macro aggregates. We

---

[4] TFP_D2b has very similar properties, while TFP_D3b, if correctly computed, does not capture innovations which require embodiment.

[5] The data source is the Statistisches Bundesamt. We eliminated the reunification break by combining level information for unified Germany with growth rates for West Germany. The data are available upon request.

first checked the cointegration properties of such systems. Running trivariate Johansen tests with lags indicated by the Schwarz Criterion we found evidence of two cointegrating vectors if we added output, consumption or hours as a third variable. For investment, we did not find any significant statistic in the trivariate system, but we found evidence of bivariate cointegration between investment and TFP_D1b at the 5% level. In light of our previous finding of cointegration between TFP_D1b and DAX1b we conclude that the trivariate Johansen test suffers from insufficient power. Hence, we treat all trivariate systems as having two cointegrating vectors.[6]
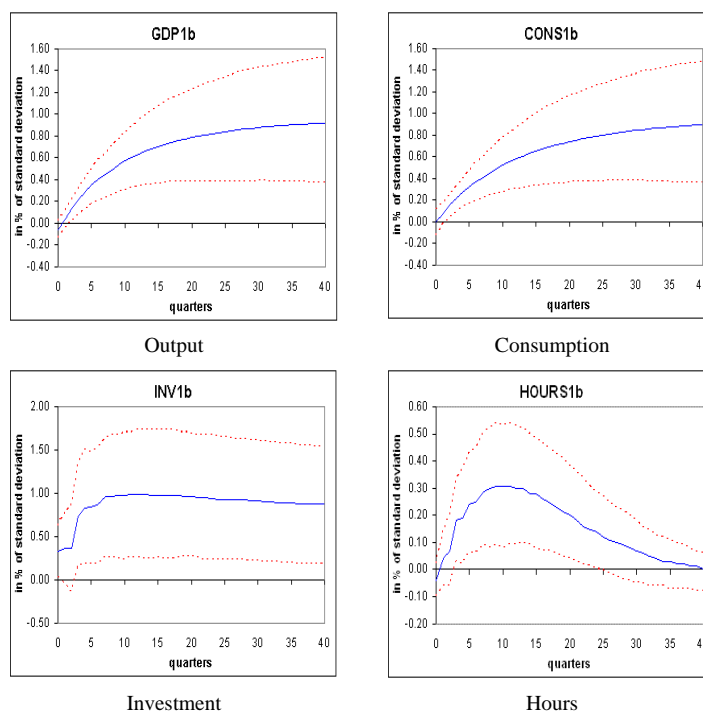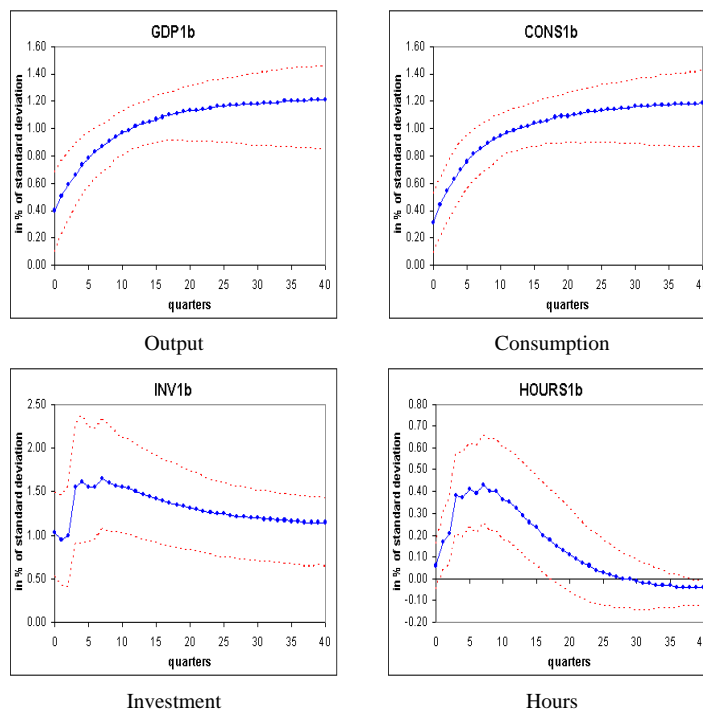
4.1   Comovement

Next we want to find out if the identified shocks ε_2 and/or $\tilde{\varepsilon}_1$ generate macroeconomic comovement and how much they contribute to the total variance of the third variable. Since estimating a trivariate VECM will in general give rise to identified structural residuals which differ from the residuals retrieved in a bivariate systems we start by imposing restrictions on the trivariate system which make the system recursive in the sense that the (TFP_D1b, DAX1b) subsystem is independent of the third variable. Essentially, this amounts to requiring that the equations for technology and stock prices are independent of the cointegrating relation between technology and the third variable and of lagged growth rates of the third variable. The equation of the third variable, by contrast, is completely unrestricted and will, in general, depend on both error correction terms and on lagged growth rates of all three variables. With this approach, we retrieve the shocks ε_2 and $\tilde{\varepsilon}_1$ in the short and long run identification scheme, respectively[7].

The impulse responses of the third variable to both shocks are displayed in Figure 4. Output, consumption, investment and hours all respond positively to ε_2 and $\tilde{\varepsilon}_1$. The responses are much stronger for $\tilde{\varepsilon}_1$ than for ε_2 – a property we will discuss below. For the short-run identification the response is insignificant on impact. (Note that this finding is not due to imposing an identifying assumption). However, for all variables, the response is significantly different from zero in the medium term, i.e. for business cycle horizons. For hours and investment the long run effects are again insignificant. Under the long-run identification, the responses are significant throughout for output, consumption and investment. Hours have a significant positive response in the medium run. Hence, the identified shocks clearly cause comovement of the main macro variables.

---

[6] Results of the cointegration test as well as all estimation results are available on request from the authors.

[7] Identification in the trivariate system (ordered as technology, stock prices, third variable) is obatined as follows: In the short run identification the (1,2), (1,3) and (2,3) element of $D(0)$ are set to zero. In the long run identification the analogous assumption is made for the corresponding long run matrix. But since there are two cointegrating vectors (implying two columns of zeros in this matrix), one more restriction is required. Since in all trivariate SVECMs with short run identification the estimate of the (unrestricted) immediate impact of innovations in the stock price equation on the third variable was virtually zero we also set the (3,2) element of $D(0)$ equal to zero.

Figure 4: Impulse responses of third variable in trivariate systems[a]

a) Short-run identification: Responses to ε_2

Output	Consumption
Investment	Hours

b) Long-run identification: Responses to $\tilde{\varepsilon}_1$

Output	Consumption
Investment	Hours

[a]Confidence intervals are obtained by 2,500 replications with the bootstrapping procedure of Hall (1992). They represent the 95% quantiles.

4.2 Variance Decompositions

Turning to the corresponding FEVDs, we compute the share of variance explained by ε_2 and $\tilde{\varepsilon}_1$ at different horizons. The rest of the variance is, of course, due to the variance of the third shock. The results, cf. Table 1, show very clearly that for all dependent variables the explanatory potential of ε_2 and $\tilde{\varepsilon}_1$ increases gradually over time, as would be expected under the delayed diffusion hypothesis. Often, the share of variance explained on impact is close to zero, while it is at least 25% in the long term (e.g. for hours) and mostly much larger. Thus, these type of shocks seem to be quite important for medium to long-term fluctuations.

 It is interesting to note that the variance decompositions are quite different under short run and long run identification. For all variables (least pronounced for hours) a much larger share of the variance is explained by $\tilde{\varepsilon}_1$ than by ε_2. This matches well our previous finding that the impulse responses to $\tilde{\varepsilon}_1$ are much stronger than to ε_2. The finding is easy to explain: ε_2 is restricted to have no technological effect on impact, i.e. ε_2 can only capture gradually diffusing technical change. $\tilde{\varepsilon}_1$, on the other hand, captures both instantaneous and gradually diffusing technological change. As the amount of technical change with an impact effect seems to be considerable, cf. Figure 2, it is not surprising that $\tilde{\varepsilon}_1$ explains a larger fraction of the variance and that its effects set in earlier and stronger than the effects of ε_2.

 In fact, one may think of the difference between the variance shares explained by $\tilde{\varepsilon}_1$ and ε_2 as the variance attributable to disembodied technical change while ε_2 captures

Table 1: FEVDs of SVECMs based on restricted VECMs

a) Short-run identification

Forecast horizon	Output		Consumption		Investment		Hours	
	ε_1	ε_2	ε_1	ε_2	ε_1	ε_2	ε_1	ε_2
1	0.63	0.00	0.26	0.00	0.25	0.02	0.09	0.01
4	0.74	0.02	0.50	0.02	0.35	0.05	0.26	0.04
8	0.76	0.09	0.65	0.08	0.39	0.11	0.28	0.11
16	0.70	0.24	0.68	0.20	0.41	0.19	0.23	0.24
24	0.63	0.33	0.64	0.29	0.42	0.25	0.21	0.30
40	0.54	0.44	0.56	0.40	0.42	0.32	0.23	0.30
80	0.46	0.53	0.47	0.51	0.42	0.41	0.24	0.30

b) Long-run identification

Forecast horizon	Output		Consumption		Investment		Hours	
	$\tilde{\varepsilon}_1$	$\tilde{\varepsilon}_2$	$\tilde{\varepsilon}_1$	$\tilde{\varepsilon}_2$	$\tilde{\varepsilon}_1$	$\tilde{\varepsilon}_2$	$\tilde{\varepsilon}_1$	$\tilde{\varepsilon}_2$
1	0.19	0.00	0.10	0.00	0.19	0.00	0.02	0.00
4	0.38	0.01	0.30	0.03	0.31	0.00	0.22	0.00
8	0.57	0.01	0.51	0.07	0.43	0.00	0.34	0.00
16	0.77	0.01	0.72	0.07	0.56	0.00	0.42	0.05
24	0.85	0.01	0.82	0.06	0.63	0.00	0.42	0.10
40	0.92	0.01	0.90	0.03	0.71	0.00	0.40	0.13
80	0.96	0.00	0.95	0.01	0.81	0.00	0.42	0.13

embodied technical change. Under this interpretation, both types of technical progress explain sizable shares of variance for all variables. But it seems that disembodied technical change dominates output, consumption and investment, while embodied technical change is relatively more important for hours. Possibly firms react to disembodied technical change primarily by increasing the utilization rate of capital and delay major increases in hours until they have decided to acquire new and better machinery which embodies the technical innovation.

The preceding analysis may be criticized because of the restrictions imposed on the trivariate VECMs. While these are necessary to reproduce the structural residuals from the bivariate SVECMs, one might suspect that unrestricted trivariate SVECMs do not support the same conclusions. Hence, we also estimate unrestricted trivariate VECMs, imposing two cointegrating vectors in each of them. We then use the same short run and long run identification as in the case of the restricted VECMs to identify structural residuals. While these are obviously not the same as in the restricted case, we continue to use the notation $\varepsilon_1, \tilde{\varepsilon}_1, \varepsilon_2, \tilde{\varepsilon}_2$ for simplicity.

The results are given in Table 2. Interestingly, the unrestricted estimates reinforce the conclusions drawn from the restricted estimates. Under the long run identification, $\tilde{\varepsilon}_2$ shocks are of almost no importance for all horizons. By contrast, $\tilde{\varepsilon}_1$ shocks increase in importance as time goes by and in the long run explain nearly all of the variance of output and consumption, two thirds of the variance of investment and slightly less than one third of the variance of hours.

Table 2: FEVDs of SVECMs based on unrestricted VECMs

a) Short-run identification

Forecast horizon	Output		Consumption		Investment		Hours	
	ε_1	ε_2	ε_1	ε_2	ε_1	ε_2	ε_1	ε_2
1	0.63	0.00	0.25	0.00	0.23	0.02	0.08	0.01
4	0.69	0.01	0.39	0.02	0.28	0.07	0.21	0.04
8	0.70	0.08	0.49	0.08	0.31	0.15	0.26	0.14
16	0.65	0.23	0.52	0.22	0.35	0.25	0.25	0.27
24	0.59	0.33	0.49	0.33	0.35	0.31	0.24	0.30
40	0.52	0.44	0.44	0.46	0.34	0.37	0.25	0.30
80	0.46	0.53	0.38	0.57	0.33	0.44	0.27	0.30

b) Long-run identification

Forecast horizon	Output		Consumption		Investment		Hours	
	$\tilde{\varepsilon}_1$	$\tilde{\varepsilon}_2$	$\tilde{\varepsilon}_1$	$\tilde{\varepsilon}_2$	$\tilde{\varepsilon}_1$	$\tilde{\varepsilon}_2$	$\tilde{\varepsilon}_1$	$\tilde{\varepsilon}_2$
1	0.21	0.00	0.13	0.00	0.01	0.00	0.01	0.00
4	0.34	0.00	0.27	0.01	0.07	0.00	0.08	0.00
8	0.51	0.00	0.44	0.02	0.16	0.01	0.20	0.02
16	0.73	0.00	0.66	0.02	0.30	0.01	0.29	0.07
24	0.83	0.00	0.78	0.02	0.40	0.01	0.28	0.10
40	0.91	0.00	0.88	0.01	0.52	0.00	0.28	0.12
80	0.96	0.00	0.95	0.00	0.68	0.00	0.29	0.12

Under the short run identification, however, ε_2 shocks have more explanatory power in the long run than ε_1 for all four macro variables. Also, the explanatory power of ε_2 is zero in the short run (note that this feature is not imposed!) and increases steadily over time. This finding is the same for all four macro variables and seems to be very much in line with the news view and the hypothesis of delayed technology diffusion. In fact, the analysis of the unrestricted VECMs makes this point even more forcefully than the analysis of the restricted VECMs.

Again, there is a sizable difference between the variance shares explained by $\tilde{\varepsilon}_1$ and ε_2 – for output and consumption. This is not the case for hours and investment. For hours, the shares are roughly the same and for investment, notable differences emerge only at the low frequencies. For both variables, embodied technological change seems to be more important than disembodied technical change – which seems to make a lot of sense for investment. As hours and investment are the variables related to production inputs, the fact that output and consumption have a large share of their variance explained by $\tilde{\varepsilon}_1$ shocks again raises the possibility that this shock not only captures disembodied technical change but also insufficient quality adjustments in the production inputs and other measures we ignore.

4.3   Comparison to Beaudry and Portier (2005, 2006)

How do our results compare to Beaudry and Portier (2005, 2006)? The analysis of German data suggests that here, as in the United States and Japan, delayed technology diffusion plays an important role in macroeconomic fluctuations. But unlike BP's results for both the US and Japan we do not find that more or less all of TFP has the delayed diffusion property. Rather, a major and perhaps dominant fraction of TFP seems to have effects on impact in the German sample.

Nevertheless, the delayed diffusion part of TFP accounts for between 25% and 33% of the variance of the main German macro variables at, e.g., 24 quarters, i.e. at a typical business cycle frequency. Its importance increases gradually over time so that it seems to be an important factor for medium and long run macro dynamics. However, its importance for short run fluctuations (two years or less) is apparently minor.

Since $\tilde{\varepsilon}_1$ does not have the impact restriction, it is clear that all three studies find a larger share of macro variance explained by $\tilde{\varepsilon}_1$ than by ε_2. But the difference is much smaller in the United States than in Japan and Germany. In fact, as far as the variance decompositions are concerned, BP's results for Japan seem to be much more in line with the German findings than with the United States'. Thus, while BP's results for the US seem to suggest that embodied technological change accounts for most of the TFP dynamics, the impression obtained from German and Japanese data is more mixed: Both types of technical progress seem to be important and, at least in the German case, the relative importance differs across macro variables.

5   Granger Causality

We are thus left with the question if the shocks ε_2 and $\tilde{\varepsilon}_1$ from the bivariate SVECMs are predominantly technological in nature. So far this is merely an interpretation sug-

gested from a simple economic model which assumed the existence of technology shocks (and assumed away all kinds of other shocks). We will now confront our candidates for technology shocks with data on technological innovations, namely the number of patents granted by the German patent agency. We will also look at data on patent applications.

These data have recently been published by Jungmittag and Grupp (2006) for 1850-1913 and 1950-1999. The data are yearly and given our earlier analysis we are just interested in the subsample 1970-1999. We aggregate the SVAR shocks by summing the quarterly values to annual frequency. We will first look at the number of granted patents. There are two measures available: the number of patents granted to German applicants (PGG) and the total number of patents granted (PGT).

It is, of course, an open issue, when exactly technological innovations become publicly known. Any invention which seeks patent protection has to be described in the patent application. By German law, this description must be sufficiently detailed to be understood by a knowledgeable person. The patent agency will publish the patent application not later than 18 months after it has been submitted. This can be long before a decision has been reached and a patent granted or refused. Hence, technological innovations are, in general, publicly known before a patent is granted.

The natural hypothesis to test for is thus a possibly causal effect of innovations on granted patents. If the technology-interpretation of the identified shocks ε_2 and $\tilde{\varepsilon}_1$ is correct, then these shocks might Granger-cause patents. Also, patents should not Granger-cause ε_2 and $\tilde{\varepsilon}_1$ since the patent applications (with technical descriptions) are already known before the patent agency decides on granting a patent.

We proceed by estimating bivariate VARs for one of the two measures of granted patents and one of the six ε_2 and $\tilde{\varepsilon}_1$ shocks identified under the three TFP measures. The lag length of the VAR is determined by the Schwarz criterion. Throughout, the chosen lag length is one.

Table 3 contains the main results. In each cell, there are two p-values. In the upper left corner we give the p-value for the hypothesis that the shock does not Granger-cause granted patents, while in the lower right corner we test the null that granted patents do not Granger-cause the shock. For all specifications, we find that the shock is indeed Granger-causal for granted patents (ususally at the 1% level of significance), while the converse seems not to be true.

Note that the tests in Table 3 are not independent. In fact, the statistics are probably highly correlated, and hence the evidence is not quite as suggestive as it may appear. But nevertheless we can safely state that Granger-causality seems to be a robust feature of the data which is found across all specifications of TFP, for both measures of patents and for both identifying assumptions.

The shocks ε_1 and $\tilde{\varepsilon}_2$ were assumed to have only a transitory effect on TFP. Thus they should not be interpreted as inventions but rather as temporary effects such as changes in motivation, strikes, variation of capital utilization etc. Hence we would expect that ε_1 and $\tilde{\varepsilon}_2$ do not Granger-cause patents, nor are Granger-caused by patents. Table 4 shows that indeed all tests give insignificant test statistics.

Apart from granted patents, data on patent applications are available: Applications from German applicants (PAG) and total applications (PAT) These data are certainly less informative, because only between one-third and one-fifth of the applications are

Table 3: Granger Causality Tests: Shocks suspected to be technological

		PGG		PGT	
TFP_D1b	ε_2	0.006		0.006	
			0.945		0.120
	$\tilde{\varepsilon}_1$	0.003		0.009	
			0.774		0.162
TFP_D2b	ε_2	0.013		0.009	
			0.952		0.179
	$\tilde{\varepsilon}_1$	0.008		0.010	
			0.897		0.144
TFP_D3b	ε_2	0.007		0.005	
			0.878		0.152
	$\tilde{\varepsilon}_1$	0.002		0.005	
			0.752		0.052

Upper left corner: P-value for null: Row variable does not Granger-cause column variable.

Lower right corner: P-value for null: Column variable does not Granger-cause row variable.

Table 4: Granger Causality Tests: Shocks suspected to have no permanent effect on TFP

		PGG		PGT	
TFP_D1b	ε_1	0.378		0.815	
			0.801		0.963
	$\tilde{\varepsilon}_2$	0.172		0.083	
			0.889		0.358
TFP_D2b	ε_1	0.544		0.715	
			0.697		0.639
	$\tilde{\varepsilon}_2$	0.341		0.249	
			0.846		0.367
TFP_D3b	ε_1	0.086		0.267	
			0.748		0.354
	$\tilde{\varepsilon}_2$	0.199		0.086	
			0.838		0.619

Upper left corner: P-value for null: Row variable does not Granger-cause column variable.

Lower right corner: P-value for null: Column variable does not Granger-cause row variable.

successful, hence many of the applications do apparently not represent inventions. Also, PAG is very strongly skewed due to a few observations in the late 1990s, when patent applications were increasing tremendously. Therefore, any kind of inference is very difficult for PAG and we have focused on PAT. We do not report the results in detail, but they are available upon request. Suffice it to say that no Granger-causality can be found between PAT and any of the identified shocks. This is what one would expect if inventors try to keep their inventions secret as long a possible, but the inventions are dis-

closed after submitting the patent application. (Recall that the time lag between submission and disclosure is at most 18 months, if it is shorter, it may well fall within the same year.)

6 Conclusions

The purpose of this study was to provide further evidence on the BP-hypothesis of delayed-technology diffusion and news-driven business cycles. For German data on TFP and stock prices we find qualitatively the same result as BP do: A high correlation between a shock with permanent effects on TFP in the long run and – under a different identification scheme – a shock which has an immediate effect on stock prices but does not affect TFP on impact.

The correlation is less pronounced, though, as in BP's analysis for US and Japanese data. Also, the impulse response analysis suggests that for Germany a substantial part of the total TFP response is immediate rather than delayed. Using a quality-adjusted measure of TFP, there is almost no delayed diffusion any more. This suggests that the delayed diffusion is confined to embodied technological change. Disembodied technological innovations seem to have immediate effects on TFP. But since the share of disembodied technological progress in total technological progress may be small, the relatively large size of the quality-adjusted TFP measure suggests that each measure of German TFP may actually contain a fairly large part of unexplained non-technology influences, i.e. our ignorance about the true nature of what we measure as TFP may be fairly large. It may therefore be the case that noise in TFP data is responsible for the immediate reaction of standard TFP measures in the impulse response analysis.

Focusing on the standard measure of TFP, we find that the shocks suspected to be technological generate comovement for output, consumption, investment and hours. Also, a sizable share of forecast errors for these variables is explained by these shocks in the medium and long run. Using unrestricted VECMs, this conclusion is even stronger: Under the short run identification shocks which do not have an instantaneous effect on technology are more important at long horizons than than those which do. At business cycle horizons, however, the former are unimportant in the short run and account for roughly a third or a fourth of the macro variance at frequencies associated with six years or more. This share is smaller than in the US data, but approximately the same size as found in Japan.

We finally checked how well the identification of technology shocks in the bivariate SVAR approach worked. The answer seems to be: Surprisingly well. Shocks suspected to be technology shocks are Granger-causal for the number of patents granted by the German patent agency, while shocks without permanent effect on technology are not. This result is very robust across different specifications, measures and identification schemes. It may therefore be the case that the SVAR approach is able to separate the true, permanent technology shocks from transitory noise which also affects measured TFP. Under this interpretation, our results are quite supportive for BP's news-driven business cycle hypothesis.
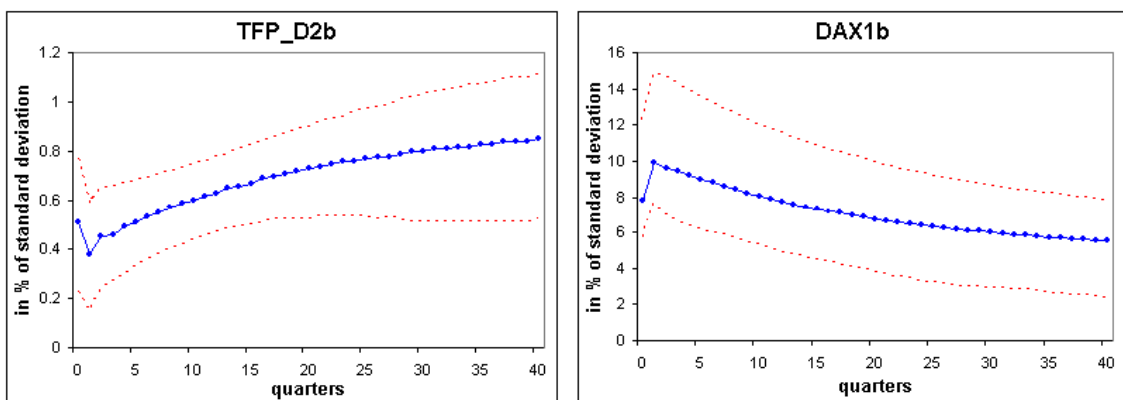
Appendix
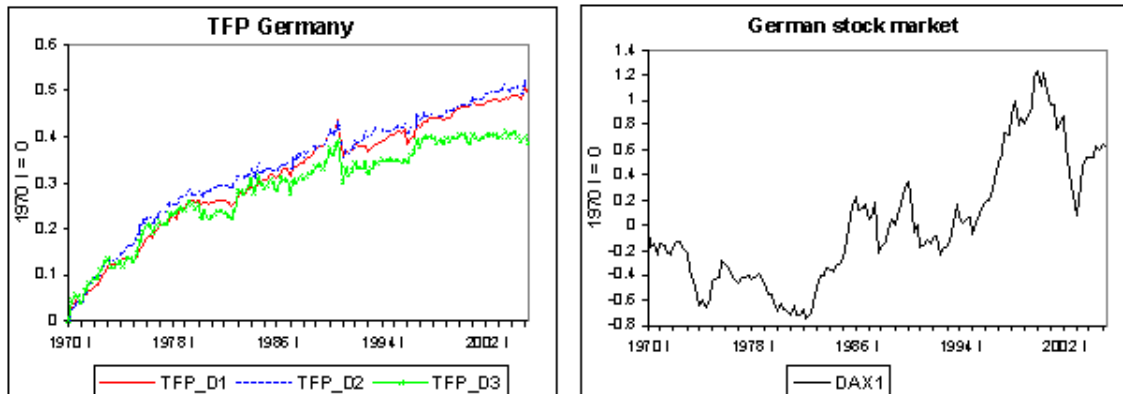
Impulse-response functions in VECM {TFP_D2b, DAX1b}[a]

a) Short-run identification (response to ε_2)
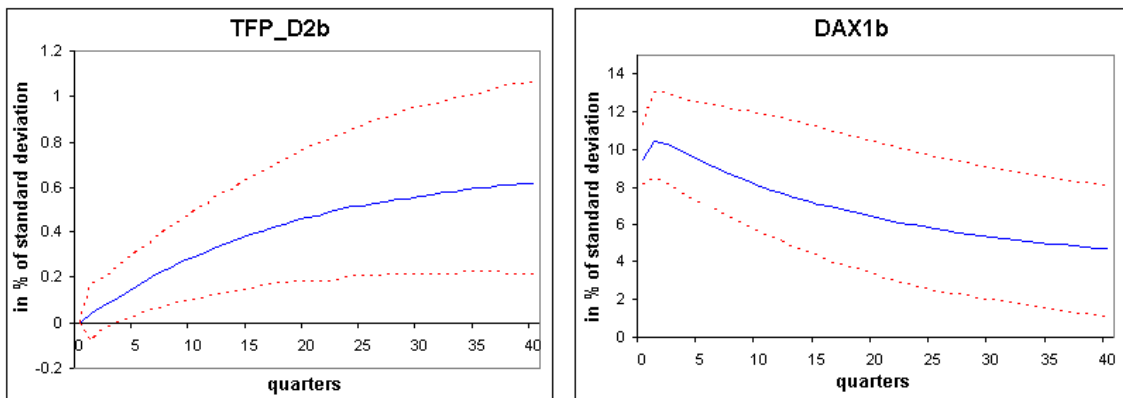
b) Long-run identification (response to $\tilde{\varepsilon}_1$)

[a]Confidence intervals are obtained by 2,500 replications with the bootstrapping procedure of Hall (1992). They represent the 95% quantiles.

References

Abramovitz, Moses (1962). Economic Growth in the United States. *American Economic Review* 52 (4) 762-782.
  www.books.google.de

Beaudry, Paul, and Portier, Franck (2005). The 'News' View of Economic Fluctuations: Evidence from Aggregate Japanese Data and Sectoral US Data. *Journal of Japanese International Economies* 19: 635-652.
  http://ideas.repec.org/p/nbr/nberwo/11496.html

Beaudry, Paul, and Portier, Franck (2006). Stock Prices, News and Economic Fluctuations. *American Economic Review* 96 (4): 1293-1307
  http://ideas.repec.org/p/nbr/nberwo/10548.html

Groth, Charlotta; Gutierrez-Domenech, Maria, and Srinivasan, Sylaja (2004). Measuring Total Factor Productivity for the United Kingdom. *Bank of England Quarterly Bulletin* Spring: 63-73,
  http://www.bankofengland.co.uk/publications/quarterlybulletin/qb040106.pdf

Hall, Peter (1992). *The Bootstrap and Edgeworth Expansion*. New York: Springer.
  www.books.google.de

Jaimovich, Nir, and Rebelo, Sergio (2006). Can News about the Future Drive the Business Cycle? Mimeo.
  http://www.kellogg.northwestern.edu/faculty/rebelo/htm/news-drive-business-cycle6.pdf

Jungmittag, Andre, and Grupp, Hariolf (2006). Dynamic Relationships between Innovation Activities and per Capita Income in Germany 1951-1999 in Comparison to 1850-1913. *Jahrbücher für Nationalökonomie und Statistik* 226 (2): 180-207.
  http://publica.fraunhofer.de/starweb/servlet.starweb?path=pub.web&search=N-42875

Long, John B., and Plosser, Charles I. (1983). Real Business Cycles. J*ournal of Political Economy* 91 (1): 39-69.

Lorenzoni, Guido (2006). A Theory of Demand Shocks. Working Paper 12477. National Bureau of Economic Research, Washington, D.C.
  http://ideas.repec.org/p/nbr/nberwo/12477.html

Oulton, Nicholas (2001). Measuring Capital Services in the United Kingdom. *Bank of England Quarterly Bulletin* Autumn:. 295-309.
  http://www.bankofengland.co.uk/publications/quarterlybulletin/qb010302.pdf

Please note:

You are most sincerely encouraged to participate in the open assessment of this article. You can do so by either rating the article on a scale from 5 (excellent) to 1 (bad) or by posting your comments.

Please go to:

www.economics-ejournal.org/economics/journalarticles/2008-10

The Editor