

Comment on “Letting Different Views about Business Cycles Compete”*

Jonas D.M. Fisher
Federal Reserve Bank of Chicago
jfisher@frbchi.org

June 29, 2009

Abstract

This comment explains why the findings presented in [Beaudry and Lucke \(2009\)](#) are misleading.

Journal of Economic Literature Classification Numbers: E0, E1, E2

Keywords: Technology shocks, investment-specific shocks, neutral shocks, news shocks, business cycle

*The views expressed herein are those of the authors and do not necessarily represent those of the Federal Reserve Bank of Chicago or the Federal Reserve System. I thank Berndt Lucke for sending me his data and software.

1 Introduction

Beaudry and Lucke (BL) address one of the classic macroeconomic questions, one that the NBER Macro Annual often returns to: What shocks drive the business cycle? Their focus is on the role of technology shocks: neutral shocks which affect the production of all goods symmetrically, investment-specific shocks which affect investment production only, and news shocks which reflect information about future neutral or investment-specific technological change. BL conclude that news shocks are the primary mover of the business cycle and that the other technology shocks are essentially irrelevant. This comment explains why the findings underlying this conclusion are misleading. Once one acknowledges the limitations of BL’s empirical strategy and introduce a structural framework to address these limitations, the primacy of news shocks disappears. Contrary to BL’s assertions, the data do not point definitively to one kind of technology shock. A balanced view of the data suggests a role for both investment and news shocks. One BL finding appears to be robust: neutral shocks are not important for the business cycle.

The modern business cycle research program at its outset focused on neutral technology shocks in addition to traditional “demand” shocks such as fiscal and monetary shocks. The real business cycle (RBC) model implies a dominant role for neutral technology shocks. But VAR analysis suggests neutral shocks are not nearly as dominant, especially when nominal variables are included in the analysis.¹ Galí (1999) and Francis and Ramey (2005) argue that estimated declines in hours after an identified positive neutral technology shock are incompatible with these shocks playing an important role in the business cycle. Without other kinds of technology shocks it would seem the business cycle cannot not be driven by technology shocks. But, there are other possibilities for technological sources of fluctuations. Two such shocks, which until recently were not given much attention, are news about future changes to technology and shocks to investment-specific technology.

Beaudry and Portier (2007) and Jaimovich and Rebelo (2009) demonstrate how information about future changes in technology can drive business-cycle-like dynamics

¹Much of the RBC literature has focused on transitory technology shocks while the VAR literature emphasizes permanent shocks. Still, early RBC models with permanent technology shocks also show these shocks accounting for a significant fraction of the business cycle. See for example Christiano and Eichenbaum (1992).

in a DSGE framework. This shock passed its first VAR test in [Beaudry and Portier \(2006\)](#). They identify news shocks with innovations to stock prices that have no contemporaneous effect on TFP and a positive long run effect on TFP and find these shocks to explain a lot of the business cycle variation in consumption.

Shocks to investment-specific technical change have been studied for some time in the DSGE business cycle literature, starting with [Greenwood et al. \(1988\)](#) and later [Greenwood et al. \(2000\)](#) and [Christiano and Fisher \(2003\)](#). This work points toward a possible role for such shocks. The shock passed its first VAR test in [Fisher \(2006\)](#). That paper shows that by measuring shocks to this technology as the sole source of permanent shocks to the relative price of investment goods, investment-specific technology shocks explain a significant fraction of the business cycle variation in output and hours.

[Fisher \(2006\)](#) does not consider news shocks and [Beaudry and Portier \(2006\)](#) do not consider investment-specific shocks. So it is natural to ask what happens when you consider both at the same time. This is what BL do. They specify a vector error correction model (VECM) in total factor productivity (TFP), the real price of investment, the per capita real value of the stock market, per capita hours, and a short term interest rate. Popular tests for the number of common trends lead them to focus on specifications with three common trends. To identify shocks to neutral technology, investment technology, technology news, preference and monetary policy they propose two alternative identification schemes. With three common trends, both schemes imply news shocks are the dominant source of business cycle variation in hours and neither neutral nor investment-specific shocks are very important. The demand shocks also play a limited role.

The remainder of this comment discusses BLs findings and points out why they are misleading. The next section describes a simple model to organize thinking about the empirical problem posed in BL. This reveals some drawbacks of their identification assumptions, but does not uncover any serious flaws. The third section discusses the central role of BL's assertion that there are three common trends in the data. It reveals that BL essentially provide no support for this assertion and that their findings are extremely sensitive to the number of common trends assumed to be in the data. Section four describes and implements an alternative to BLs empirical

strategy which uses the structural model's implication that there is a single common trend. This analysis confirms that the role ascribed to news shocks by BL is vastly overstated and that investment shocks are far more important than they suggest. The penultimate section briefly discusses a key problem any researcher faces when trying to identify neutral, investment and news shocks. The final section briefly discusses directions for future research on the sources of the business cycle.²

2 Implications of a Structural Model

It is helpful to introduce a simple model to shed light on BL's analysis. This section describes a simple stochastic growth model which could be at the heart of any modern DSGE model. For simplicity, it abstracts from nominal rigidities and emphasizes technological sources of business cycle variation. The model is:

$$\begin{aligned} & \max_{\{C_t, H_t, K_{t+1}, I_t\}_{t=0}^{\infty}} \mathcal{E}_0 \sum_{t=0}^{\infty} \{\ln C_t - H_t\} \\ \text{subject to } & G(C_t, V_t I_t) \leq A_t K_t^\alpha H_t^{1-\alpha}; \\ & K_{t+1} = (1 - \delta)K_t + (1 - \theta_t s(I_t/I_{t-1}))I_t; \\ & \ln A_t = g_A + \ln Z_{t-1} + \eta_t + \sum_{i=1}^8 \xi_{t-i}^i; \\ & \ln V_t = g_V + \ln V_{t-1} + \omega_t + \sum_{i=1}^8 \zeta_{t-i}^i; \\ & \ln \theta_t = \rho \ln \theta_{t-1} + \varepsilon_t + \sum_{i=1}^8 \varphi_{t-i}^i; \\ & G(C_t, V_t I_t) = \left[\lambda_C C_t^\psi + \lambda_I (V_t I_t)^\psi \right]^{1/\psi}, \\ & K_0 \text{ given.} \end{aligned}$$

Here C_t, H_t, I_t, K_t, A_t and V_t are consumption, hours, investment, capital, neutral technology and investment-specific technology. The function $s(\cdot)$ is a standard adjustment cost function. The i.i.d. shocks driving this model are neutral technology

²This comment does not address the potential for weak instruments to make it impossible to disentangle the role of technology shocks with reduced form models. See [Watson \(2006\)](#) for an insightful discussion on this topic.

shocks, η_t , investment specific technology shocks, ω_t , a shock to the efficiency of the installation technology, ε_t , and news about future levels of the neutral, investment and installation technology, ξ_{t-i}^i , ζ_{t-i}^i and φ_{t-i}^i . This model has two stochastic trends arising from growth in the neutral and investment-specific technology, g_A and g_V .

The neutral and investment shocks are familiar. The news shocks are less familiar and so they require some discussion. The news shocks are interpreted as signals about future levels of technology. These signals may occur one to eight quarters before the level of technology is realized. A signal may be offset by the current draw of the innovation so that signals may be incorrect. Notice also that “news” in this model is in fact 24 shocks, reflecting signals of three kinds of technology in future over different horizons. So, any empirical procedure that seeks to extract one news shock, such as BL’s, is necessarily going to be picking up the effects of many different shocks that each have different dynamic effects. We should therefore be very careful interpreting impulse response functions from such a procedure, but the variance decomposition probably still has meaning.

The installation shock is even less familiar. Recently it has been studied by [Christiano et al. \(2008\)](#) and [Justiniano et al. \(2009\)](#). It is argued below that by not taking into account this shock, BL’s and similar identification procedures may actually be picking up the effects of this shock, or something like it. However, most of this comment proceeds as if this shock is turned off.

Another unusual feature of this model is the CES aggregator for consumption and investment goods, $G(\cdot)$. The CES formulation is the easiest way to introduce an endogenous relative price of investment goods into a model. This does not appear in the modern DSGE literature, but it is useful for thinking about identifying investment shocks using the relative price of investment goods.

In the competitive equilibrium the aggregate resource constraint is equivalent to

$$\frac{A_t}{G_{Ct}} K_t^\alpha H_t^{1-\alpha} = C_t + P_{It} I_t$$

where

$$P_{It} = V_t^\psi \frac{\lambda_I}{\lambda_C} \left[\frac{C_t}{I_t} \right]^{1-\psi},$$

P_{It} is the price of investment goods in consumption units, and G_{Ct} is the derivative of G with respect to C_t . The expression for the investment price reveals the obvious

point that if there is any curvature in the transformation frontier then the investment price does not identify the investment specific shock. Any shock which changes the ratio of consumption to investment changes the investment price as well. This holds in much more general settings. For example, [Basu et al. \(2009\)](#) show in an economy with multiple intermediate and final goods that the real investment price confounds technology and variation in the use of intermediate inputs and factors of production.³

This point is important for BL. Essentially, it invalidates the first of their two identifications schemes, ID1. The ID1 scheme, the one based exclusively on short run restrictions, includes the assumption that news and non-technology shocks do not effect the relative price. In light of the relative price expression this assumption is clearly not plausible. These shocks change the consumption:investment ratio and hence the relative price.

BLs second identification scheme, ID2, combines short run and long run restrictions. The short run restrictions are that only innovations to TFP can affect TFP contemporaneously and that the monetary policy shock does not affect hours contemporaneously. Here there is no restriction on what shocks may affect the relative price of investment contemporaneously. The long run restrictions are that only neutral and news shocks affect TFP in the long run and that neutral, investment-specific and news shocks are the only shocks which affect investment-specific technology in the long run. These assumptions are sufficient to just-identify five shocks to neutral technology, investment technology, technology news, preferences and monetary policy.

The only problem with the long run restrictions in ID2 is that they allow the neutral shock to have a long run impact on the investment-specific technology. This is not true in the model written above. BL consider the over-identifying restriction that the neutral shock has no long run effect on the investment-specific technology and report that it does not effect their findings. It is unclear how robust this finding is to the number of common trends assumed.

ID2's short run restrictions on TFP are more problematic. Typical measures of TFP are well-known to be plagued by difficulties in measuring the factor inputs. Since

³[Greenwood et al. \(1997\)](#) argue that one can identify the neutral technology, A_t , with a Solow residual computed using output measured in consumption units. The resource constraint written with the investment price shows that this approach is valid only if the transformation frontier is linear, that is $G_{Ct} = 1$.

these variables are likely to be influenced contemporaneously by the other shocks, the short run restrictions on TFP are implausible unless the endogenous variation in TFP can be purged. One could go some way toward this by exploiting quarterly versions of the “cleansed” TFP measure developed by [Basu et al. \(2006\)](#), but BL do not. Even so, the cleansing done by [Basu et al. \(2006\)](#) is at the annual frequency and interpolation is used to compile a quarterly series. Therefore this alternative measure is likely subject to endogeneity problems as well. Still, it is preferable to the rudimentary TFP measure used by BL.

A second concern about the short-run restrictions on TFP derives from consideration of a multi-sector growth model. An implication of results in [Basu et al. \(2006\)](#) and [Basu et al. \(2009\)](#) is that aggregate TFP is a weighted average of sectoral TFP. As an example, consider the simple two-sector growth model with investment good and consumption good sectors where production functions in each sector are subject to correlated TFP shocks. In this model measured aggregate TFP is an expenditure share weighted average of consumption and investment sector TFP. The relative price of investment is a function of both TFPs as well as factor input usage in the different sectors. As long as consumption and investment shocks are not perfectly correlated, then shocks to either sector’s TFP will influence both aggregate TFP and the relative price of investment. Consequently, the short run restriction that the relative price does not effect aggregate TFP is invalid in this model. Note however, it is valid in the model written above. In this model conventional measures of TFP correctly identify the neutral technology, up to endogenous variation in factor inputs.

While the preceding certainly raises doubts about the paper’s findings, some may view them as quibbles. And there are many things to like about the identification strategy ID2. It is simple and, setting aside the caveats just mentioned, intuitively appealing. Therefore, for the remainder of this comment I consider the implications of adopting ID2.

3 Common Trends in the Data

A central element of BL is its approach to co-integrating relationships, that is common trends. Unfortunately BL provide no guidance whatsoever on the sort of relationships

we should expect to see in the data. Essentially BL attempt to “let the data speak” alone on this issue. In the context of co-integration, this is an empirical strategy fraught with peril. This section explains why this is so and shows how sensitive BL’s results are to the assumed number of common trends.

Recall that the baseline empirical model is a VECM in five variables: logs of TFP, the relative price of investment, the per capita real value of the stock market, per capita hours and a short term interest rate. To implement the VECM BL need to settle on lag lengths and the number of co-integrating relationships. BL follow a robust strategy for choosing lag length by allowing enough lags to remove any serial correlation from the residuals. They work with five lags. For the co-integrating relationships, they rely on a much less robust approach due to Johansen. This leads them to focus on three common trends.

One advantage the authors attribute to the VECM framework is that it allows them to bypass issues regarding the stationarity of hours. In particular, while their VECM requires that hours are included in first differences, in principle one of the co-integrating vectors could put weight only on hours so that the model is equivalent to one where hours are stationary in levels. In practise, while the hours variable is stationary, the estimated error correction terms all put weight on every variable in the system. So it unclear whether the stated advantage has any practical value. However, even if it does, by including hours and a short term interest rate in the VECM another problem crops up.

In the baseline case, the Johansen procedure for determining the number of common trends involves sequentially testing the hypotheses of no more than zero, one, two or three co-integrating vectors. The first time the test is not rejected determines the number of co-integrating vectors. BL find that no more than zero, one and two co-integrating vectors are rejected at the 5% level, but that the hypothesis of no more than three is not rejected. This is their justification for focusing on three common trends. Note that, while two stochastic trends and five integrated variables are consistent with three common trends, this scenario is also consistent with one or two common trends as well.

The problem with this empirical strategy was documented by [Elliott \(1998\)](#) in a famous *Econometrica* article. He studied the performance of popular tests for the

order of co-integration when there are variables in the system that are incorrectly assumed to have an exact unit root. He found that even very small departures from a unit root lead to a severe breakdown in the size properties of the tests. Recall that “size” relates to the probability of incorrectly rejecting the null hypothesis. [Elliott \(1998\)](#) finds that the size of Johansen tests could approach unity, that is there is close to a one hundred percent chance of incorrectly rejecting the null, for even small deviations from a unit root.

Should BL be worried about this? Definitely. The variables per capita hours and short term interest rates are notorious for their unit root like behavior even though the series themselves are stationary. BL even go out of their way to point out that their per capita hours series is stationary. While working in first differences may be justified in some situations, for instance forecasting, it is clearly a drawback when testing for the number of common trends.

Given the lack of information conveyed by the statistical tests, BL should assess the robustness of their findings to the assumption of three common trends. Since they focus on three common trends, they should present a compelling argument for why this is a plausible way to view the data. They do neither of these things. Instead BL dismiss less than three common trends by appealing to the permanent response of hours to their identified monetary policy shock.

This a spurious argument. Only if one is interested in identifying monetary shocks does the permanent response of hours to an identified monetary shock matter. One can still address questions about the role of technology shocks in the business cycle without having to label other shocks in the system. This is because the identification assumptions for technology are not invalidated by the existence of some linear combination of underlying structural shocks which have a permanent impact on hours. This is all their finding about hours is indicating. So, BL do not provide a valid rationale for focusing on three common trends.

As it turns out the results are *very* sensitive to how many common trends are assumed to be in the data. To make this point, consider a version of BL’s baseline system with slightly different choices for how the variables are measured. TFP is the one developed by [Basu et al. \(2006\)](#). This measure of TFP is preferred to the one used by BL since the underlying capital and labor service series are more carefully

constructed and the treatment of endogeneity is more sophisticated.⁴ The investment price is the NIPA nonresidential fixed investment deflator divided by the chain weighted deflator for consumption of non-durables and services. Using the quality adjusted measures described by BL does not change the results substantively. The stock variable is measured with the per capita S&P 500 deflated in the same way as the investment price. Both of these measures depart from the variables used by BL. BL use the overall consumption deflator for their baseline measures of these variables. The nondurables and services deflator is preferred because it is the measure suggested by the theory. In addition, the relative price of consumer durables has a trend much like the investment price, so using the overall consumption price confounds multiple sources of technical change. The hours and interest rate variables are the same as in BL's NIPA.h system. As in BL, five lags are used. The VECM is estimated with software provided by Berndt Lucke.

Table 1 displays the percent of the forecast error decomposition of hours which can be attributed to the three kinds of technology using ID2, for one, two and three common trends. With three common trends news shocks clearly dominate the other technology shocks. From one year to eight years news shocks account for about one half of the forecast error variance. The investment shock is important after two years in this specification, but half as important as the news shock. Note how this last finding differs from BL. This shows a sensitivity to measurement choices not evident in BL. The neutral shock is irrelevant with three common trends. With two common trends news shocks continue to dominate and now the investment shocks are much less important. The neutral shocks continue to be irrelevant. With one common trend news shocks are only dominant among the technology shocks at horizons under two years. From two years to eight years, that is at horizons corresponding to business cycle frequencies, investment shocks are the most important technology shock and news shocks become much less important. Neutral shocks are still irrelevant. The findings with one common trend clearly are very different from the results emphasized in BL.

At this stage there is nothing to choose between these three scenarios. So, contrary

⁴Note that this measure of TFP is based on output prices, not consumption prices. This is the correct way to proceed in all cases where there is a non-linear trade-off between consumption and investment. See footnote 3.

to the claims in BL, within the VECM framework the data do not speak strongly in favor of one kind of technology shock. There do seem to be two robust findings however. First, technology shocks combined account for at least 50 percent of the business cycle variation in hours. Second, neutral shocks are irrelevant.

4 An Economic Rationale for One Common Trend

When the data do not speak clearly on a question, it is natural to consider whether we can use theory to provide guidance. Indeed this is precisely the kind of situation when theory can be most useful. It turns out that the model described above (and a large class of mainstream macro models) is definitive on the question of how many common trends one should expect to see in the data: exactly one. This section describes that common trend.

By considering the balanced growth path of the model described above, it is easy to verify that there is just one co-integrating relationship among the variables in BL's baseline specification. Specifically, the "error correction" variable x_t given by

$$x_t = \ln S_t - \ln A_t - \frac{\alpha}{1 - \alpha} \ln V_t$$

is stationary along the model's balanced growth path. The term S_t is equal to $P_{k't}K_{t+1}$, where $P_{k't}$ is the price of new installed capital. It is the consumption value of the capital stock and can be interpreted as the value of the representative firm.

Assume that means have been removed from the variables used in constructing x_t . Then this variable has a very simple and intuitive interpretation. When $x_t = 0$ then the value of the firm simply reflects the current state of the long run fundamentals, *i.e.* the levels of neutral and investment specific technology. When $x_t \neq 0$ then the value of the firm reflects the effects of news and installation shocks since, in this model, there are no other shocks.

Since this co-integrating relationship is predicted by a model which is at the heart of most empirical DSGE models it is worth studying its empirical counterpart to learn how it behaves. An empirical version of the error correction term is

$$\hat{x}_t = \ln \hat{S}_t - \ln \hat{A}_t + \frac{\alpha}{1 - \alpha} \ln \hat{P}_{It}$$

where \hat{S}_t is the leverage-adjusted value of the stock market, \hat{A}_t is measured neutral technology and \hat{P}_{It} is the measured consumption price of investment. Below, these variables are measured as previously described with the addition of data for leverage. The leverage adjustment is to convert the stock market value into the value of the underlying firms. For this the leverage ratio for the non-financial corporate sector taken from the Flow of Funds accounts is used. Finally, $\alpha = .33$.

Figure 1 displays \hat{x}_t after removing its mean. The series is clearly stationary and has pronounced low frequency fluctuations. These low frequency movements are suggestive of a role for news shocks. Over the sixties and early seventies the series is positive. This seem in line with the view that news was generally good about future productivity growth arising from neutral and investment technical change during this time. From the early 1970s to the mid 1990s the series is negative, indicating that the stock market undervalued the current state of technology, presumably because the news was bad about future productivity growth. The second half of the 1990s appear as a period of great optimism about future productivity growth. As of the end of the sample, 2006q4, the series indicates the stock market reflects the current state of technology well, with news not pointing for any substantial changes in rates of productivity growth going forward. Overall, the nature of this series seems consistent with the theory, in the sense that it is stationary and that its dynamics seem in line with what an empirical version of the model might predict.

If possible, it would be useful to compare \hat{x}_t with the error correction terms associated with the estimates in Table 1. This could provide guidance on which specification is most plausible. However, with the exception of the one common trend case, there is not a unique way to write the error correction terms. So, it is not possible to distinguish between the specifications in this way. In the one common trend case, the common trend is identified up to a factor of proportionality. Examining the error correction term in this case reveals that it is nothing like the one suggested by the theory.

Another approach to distinguishing among the specifications in Table 1 is to examine the corresponding impulse responses to the technology shocks. Doing so reveals something quite interesting: in all three specifications the response of hours to the individual technology shocks is permanent. The theory described above, and any

mainstream macro model consistent with balanced growth, predicts transitory effects of technology shocks. It is hard to know what to make of this finding except that it casts doubt on the VECM methodology.

5 An Alternative Empirical Strategy

The fact that mainstream theory predicts a single common trend suggests taking seriously the findings in the last three columns of Table 1. However, the impulse responses for hours are at odds with the theory so the veracity of these findings is in doubt. Fortunately there is more than one way to write down a co-integrated system and BL's identification assumptions do not depend on working within the confines of a VECM. This section considers an alternative empirical specification consistent with the theory which does not suffer from the drawbacks of the VECM uncovered above.

One particularly convenient empirical approach involves estimating a garden variety VAR. Given the implications of the theory, there is a natural VAR to consider: the system including log first differences of TFP and the investment price, the error correction variable x and hours and the interest rate in levels. This has the virtue of imposing the co-integrating relationship derived from the theory. This system is estimated using the same measures of the variables as before and by including five lags (the results are similar with four lags.)

Figure 2 displays the impulse response functions for hours associated with the three technology shocks identified using ID2. These responses are transitory and appear large for the investment and news shocks. Table 2 displays the corresponding forecast error decomposition. This shows that investment and news shocks each account for about a quarter of the forecast error variance from three to eight years out. Neutral shocks are again irrelevant. The finding that about half of the forecast error at horizons associated with the business cycle can be attributed to technology shocks holds as well.

Now consider splitting the sample. There are many reasons to do so, including the apparent trend break in the investment price emphasized by Fisher (2006). Splitting the sample makes it possible to compare the findings to those in Fisher (2006). The two subsamples are the same as in Fisher (2006): 1955q1-1979q2 and 1982q3-2006q4.

Because of the shorter samples the estimates are based on including four lags.

Figure 3 displays the impulse responses for hours for the two sub-samples also identified with ID2. In all cases the responses are transitory. The responses to the investment and news shocks are both smaller in the first sub-sample compared to the full sample. The magnitude of the response to the investment shock is now similar to the neutral case, which continues to be small. The response to the news shock is similar to the ones estimated over the first sub-sample in Fisher (2006) for neutral and investment shocks. The second sub-sample is very different. Here the response to the investment shock is large and the responses to the other two shocks are small. The response to the investment shock is similar to the one estimated in Fisher (2006).

Table 3 shows the corresponding variance decompositions. In the first sub-sample the role of technology shocks is much diminished. News shocks at best account for about a quarter of the forecast error from four to eight years. In the second subsample investment shocks account for more than fifty percent of the variance at all horizons; the other shocks are essentially irrelevant.

Overall, these findings paint a very different picture to the one BL present. Over the full sample news and investment shocks are equally important, each accounting for a quarter of hours fluctuations. For the split sample, the news shock is marginally important in the early period, but irrelevant in the later period. Investment shocks are very important in the later period. One finding which holds over from the VECM analysis is that neutral shocks are always irrelevant.

6 An Important Caveat

So far this comment has not considered installation shocks. These are included in the model to make a simple point about a potential specification error faced by any reduced form approach to identifying the different technology shocks using stock market and investment good prices. The DSGE literature has started to study installation shocks. Justiniano et al. (2009), in a model similar to the one described here, find that installation shocks account for about 3/4 of the business cycle variance of hours. News and investment shocks account for virtually nothing. These findings are with an estimation procedure that includes data on the investment price, which their prior

work has not done.

This suggests that we need to consider the implications of ignoring the installation shock. To gain intuition for why something like the installation shock could matter for the findings I have presented and those in BL, consider the following equilibrium relationship between the price of new installed capital, the price of investment, and the marginal installation product of investment:

$$P_{k't} = \frac{P_{It}(V_t, \dots)}{MPI_t(\theta_t, \dots)}$$

This close relationship between stock prices, the investment price and the marginal installation product of investment strongly suggests that if the installation shock is important, BL's identification strategy might attribute variation due to installation shocks to news or investment shocks. Note that one does not have to rely on installation shocks to make this point. Any factor which influences the investment process and so enters into this equation, including financial frictions, could lead to specification error.⁵

To verify that this is indeed an issue, consider the following Monte Carlo experiment involving a large number of datasets generated from the model in [Justiniano et al. \(2009\)](#) in which news and investment shocks account for essentially zero business cycle variation. Using this artificial data, the VAR described in the previous section is estimated many times, each time calculating the variance decomposition using ID2. Over these many samples, the mean contribution of news shocks to the forecast variance of hours at the 20 quarter horizon is 33%. Investment shocks are correctly predicted to have a small effect.

7 Conclusion

This comment has raised serious questions about the plausibility of BL's findings. The data do not support their conclusion that news shocks drive out other technology shocks and are the primary mover of the business cycle. A balanced view of the data

⁵One way to address this possibility of specification error would be to take advantage of the fact that the wedge between investment and stock prices is Hayashi's marginal q. Empirical measures of q could be incorporated into the empirical model and assumptions developed to identify installation shocks.

suggests a role for both investment and news shocks. One BL finding appears to be robust: neutral shocks are not important for the business cycle. Of course, these conclusions are subject to the caveat described in the previous section.

Where do we go from here? As we raise the number of shocks on the table, VARs (or VECMs) become untenable, and we are nearing that point. So I think progress in determining the sources of the business cycle is most likely to be within the DSGE setting and in the direct measurement of shocks. The next steps for the DSGE literature should be to incorporate stock price and data and to endogenize the investment price. In particular, the DSGE literature needs to move beyond identifying the investment price with the investment technology. Direct measurement of shocks, such as in my work with Susanto Basu, John Fernald and Miles Kimball where we use the US input-output tables and industry-level production data to identify sector-specific technology shocks, is a complementary approach and should provide information useful for assessing DSGE findings.

References

- Basu, S., J. Fernald, and J. D. F. M. Kimball (2009). Sector-specific technical change. San Francisco Fed Manuscript.
- Basu, S., J. Fernald, and M. Kimball (2006). Are technology improvements contractionary? *American Economic Review* 96(5), 1418–1448.
- Beaudry, P. and B. Lucke (2009). Letting different views about business cycles compete. In D. Acemoglu, M. Woodford, and K. Rogoff (Eds.), *NBER Macroeconomics Annual*. Cambridge, MA: NBER.
- Beaudry, P. and F. Portier (2006). Stock prices, news and economic fluctuations. *American Economic Review* 96, 1293–1307.
- Beaudry, P. and F. Portier (2007). When can changes in expectations cause business cycle fluctuations in neo-classical settings? *Journal of Economic Theory* 135, 458–477.
- Christiano, L. and M. Eichenbaum (1992). Current real business cycle theories and aggregate labor market fluctuations. *American Economic Review* 82(3), 430–450.
- Christiano, L. and J. D. Fisher (2003). Stock market and investment good prices: Implications for macroeconomics. NBER Working Paper No. 10031.
- Christiano, L., C. Ilut, R. Motto, and M. Rostagno (2008). Signals: Implications for business cycles and monetary policy. NBER Summer Institute Presentation.
- Elliott, G. (1998). The robustness of cointegration methods when regressors almost have unit roots. *Econometrica* 66, 149–158.
- Fisher, J. D. (2006). The dynamic effects of neutral and investment-specific technology shocks. *Journal of Political Economy* 114, 413–451.
- Francis, N. and V. Ramey (2005). Is the technology-driven real business cycle hypothesis dead? shocks and aggregate fluctuations revisited. *Journal of Monetary Economics* 52(8), 1379–1399.

- Galí, J. (1999). Technology, employment, and the business cycle. *American Economic Review* 89(1), 249–271.
- Greenwood, J., Z. Hercowitz, and J. Greenwood (1988). Investment, capacity utilization and the business cycle. *American Economic Review* 73(3), 402–417.
- Greenwood, J., Z. Hercowitz, and P. Krusell (1997). Long run implications of investment-specific technological change. *American Economic Review* 87(3), 342–362.
- Greenwood, J., Z. Hercowitz, and P. Krusell (2000). The role of investment-specific technological change in the business cycle. *European Economic Review* 44(1), 91–115.
- Jaimovich, N. and S. Rebelo (2009). Can news about the future drive the business cycle? *American Economic Review*. Forthcoming.
- Justiniano, A., G. Primaceri, and A. Tambalotti (2009). Investment shocks and the relative price of investment. Northwestern University Manuscript.
- Watson, M. (2006). Comment on “Assessing Structural VARs”. In D. Acemoglu, M. Woodford, and K. Rogoff (Eds.), *NBER Macroeconomics Annual*. Cambridge, MA: NBER.

Table 1: Variance Decomposition of Hours in the VECM

Horizon	Three Common Trends			Two Common Trends			One Common Trend		
	Neut.	Invest.	News	Neut.	Invest.	News	Neut.	Invest.	News
1	.00	.00	.25	.00	.00	.25	.00	.05	.46
4	.02	.06	.50	.02	.04	.50	.04	.19	.44
8	.01	.21	.56	.01	.03	.58	.04	.40	.23
12	.02	.24	.49	.02	.05	.50	.02	.55	.14
20	.02	.26	.48	.02	.10	.40	.02	.64	.10
32	.01	.29	.53	.02	.11	.39	.01	.69	.07

Table 2: Variance Decomposition of Hours in the VAR, Full Sample

Horizon	Neutral	Investment	News
1	.00	.09	.05
4	.00	.12	.03
8	.02	.16	.16
12	.03	.21	.25
20	.04	.25	.27
32	.04	.27	.26

Table 3: Variance Decomposition of Hours in the VAR, Split Sample

Horizon	1955q1-1979q2			1982q3-2006:4		
	Neutral	Investment	News	Neutral	Investment	News
1	.01	.02	.14	.05	.60	.10
4	.02	.01	.11	.03	.71	.09
8	.02	.01	.10	.02	.67	.04
12	.02	.01	.15	.02	.63	.03
20	.03	.01	.22	.02	.57	.04
32	.02	.01	.30	.02	.54	.05

Figure 1: The Error Correction Term Suggested by the Model




Figure 2: Hours Responses to Technology Shocks in the Full Sample




Figure 3: Hours Responses to Technology Shocks in the Split Sample

