Discussion

Kevin Moran

Context and Summary

The task of discussing this paper was an occasion for me to learn about the debate over the effects of technology shocks on hours, in particular, and recent advances on the connection between empirical and quantitative-theoretic models, in general. I am grateful to the authors and to the conference organizers for this opportunity.

One of the research programs pursued by the authors is estimating dynamic stochastic general-equilibrium (DSGE) models by minimizing the distance between two sets of impulse responses: those implied by the model and those retrieved from data. The success of this strategy thus rests largely on the quality of the impulse responses derived from the data. There is already a voluminous literature documenting the shape of the impulse responses that follow a monetary policy shock, and these have been used to estimate DSGE models with real and nominal rigidities (see Christiano, Eichenbaum, and Evans 2001). The underlying research objective motivating the work contained in the present paper is to identify data-derived impulse responses from technology shocks and propose them as the next targets that the estimation of DGE models should try to match.

The specific contribution of the paper is to shed light on the shape of the impulse responses that follow technology shocks. This has been a topic of brisk controversy since Galí (1999) proposed an econometric procedure to identify technology shocks and, among other results, reported that hours worked appeared to decrease following persistent, positive technology shocks. Such a result is important because sticky-price models can generate it fairly easily, while standard real-business-cycle models will most likely generate the opposite result, i.e., an increase in hours. An estimation method that seeks to make the model consistent with these responses would therefore probably favour the sticky-price model.

The paper uses annual data from Canada and the United States, thus providing a comparison with parallel research that used quarterly American data (Christiano, Eichenbaum, and Vigfusson 2003). Using Galí's procedure, the authors first show that his result does not appear to be robust to the assumption regarding the stationarity of hours. When hours are treated as difference-stationary (as was the situation in Galí's benchmark case) and consequently the growth rates of hours are used in the empirical work, Galí's result obtains: hours fall following a favourable technology shock. However, if one views hours as stationary and therefore levels are used, Galí's results are reversed: hours now increase following the technology shock. Second, the authors repeat their analysis using a simulated data set, where technology shocks increase hours *by construction*. They find that assuming difference-stationarity leads to the incorrect conclusion that hours fall. This casts doubt on the ability of the difference-stationary specification to recognize any positive correlation between hours and technology shocks. Third, the authors extend their analysis to a multivariate system in which the interest rate, inflation, and money growth are added as macroeconomic variables of interest. They report that the responses using the Canadian data are consistent with the view that the Bank of Canada accommodates technology shocks. The authors conclude that technology shocks are associated with an increase—rather than a decrease—in hours, even if the contribution of these shocks to business cycles is not as strong as once thought.

1 How to (Statistically) Treat Hours

Galí (1999) identifies technology shocks by assuming they are the only source of permanent fluctuation in measures of average labour productivity, within a bivariate vector autoregression (VAR) containing hours and the productivity variable.[1] The authors confirm that when the measure of hours worked is first-differenced (Galí's benchmark case), the procedure produces a negative response of hours worked following a favourable technology shock. However, when using the level of (per-capita) hours, or a (linearly)

1. The strategy of using long-run restrictions to identify shocks in VARs originates in Blanchard and Quah (1989), where the authors identify supply shocks by assuming that within a bivariate (output, unemployment) VAR, such shocks are the only ones having a permanent effect on output. Interestingly, Blanchard and Quah report that following a favourable supply shock, unemployment initially increases before eventually decreasing, a result that can be interpreted as consistent with Galí's.

detrended version of these levels, they report that a positive technology shock is associated with an increase in hours.

My first comment concerns the possibility that the correct statistical representation of hours may be in the form of a trend-stationary process. Galí (1999) reports, within a series of checks on the robustness of his results, that the negative correlation between hours and technology shocks is still present when hours are detrended linearly (see Galí 1999, Figure 3). While confirming these results, the authors (2003) argue that the correlation is dependent on the specific manner in which the trend is introduced and, furthermore, that there is no conclusive evidence that the trend-stationary assumption is the correct way to describe the available data. It would be interesting to verify what kind of trend, if any, is appropriate for Canadian data and what the implications are of using detrended hours for the results presented here.[2]

Second, we know that the effect of the first-difference filter is to over-emphasize the high-frequency components of the series to which it is applied (i.e., the gain of that filter increases significantly as the frequency increases; see Baxter and King 1999, Figure 5). This suggests that both negative and positive correlations between hours and productivity shocks might be present in the data, with the correlation being positive at low frequencies and negative at higher frequencies. Christiano, Eichenbaum, and Vigfusson refer implicitly to this conjecture when they discuss the fact that the technology shocks they identify have a significantly greater effect on the longer horizons in the variability of hours and other economic variables. It would be interesting to explore more systematically the relevance of the conjecture using filters that can select specific frequencies of the data, such as the band-pass filters advocated by Baxter and King.

2 The Encompassing Test: Some Further Suggestions

The effect of technology shocks on hours thus seems to depend on whether hours are treated as difference- or level-stationary. After arguing that formal, statistical tests of stationarity yield no decisive conclusion when applied to hours, the authors argue that one can still develop a sense of which result is the most robust by following an encompassing strategy.

2. Note that Christiano, Eichenbaum, and Vigfusson (2003) select a quadratic trend, whereas Galí used a linear trend. Furthermore, throughout their papers, Christiano, Eichenbaum, and Vigfusson use the log of (per-capita) hours, while Galí uses the log of total hours.

To do so, they take the VAR they estimated with hours in levels (the one that yielded the positive response of hours to technology shocks) and generate simulated data from it. Next, they apply Galí's procedure to that simulated data set using the difference-stationary assumption and find that it fails to detect the positive correlation. This raises questions about the ability of the difference-stationary specification to identify a positive correlation between hours and technology shocks. Next, they use the second estimated VAR as the generator of simulated data (the one that implies a negative correlation) and re-apply the procedure using the level specification. The procedure does not detect the negative correlation when using Canadian data and reports essentially a zero response of hours when using the American data; note, however, that all confidence bands are wide.[3] The upshot of these experiments is that the level specification appears more adept at identifying the truth than the difference-stationary one is. Otherwise said, first-differencing the data may produce such strong distortions that the information in the original data is lost.

I would like to suggest that alternative artificial data sets might be used to run experiments similar to those performed by the authors. For example, a data set created by simulating a standard real-business-cycle model (one for which we *know* that hours and technology shocks would be positively correlated) could be used. Would the difference-stationary specification be able to identify that truth? Furthermore, a model with sticky prices, in which the immediate aftermath of a technology shock is characterized by a decrease in hours, could also be employed. Would the level specification be able to recognize that alternative truth?

3 Industry-Level Data: Support for the Conclusion

Applying Galí's procedure to industry-specific data, Chang and Hong (2003) report results that can be interpreted as favouring the conclusions of Christiano, Eichenbaum, and Vigfusson on the likely effect of technology shocks on hours worked. Among the many subcategories of manufacturing

---

3. It may be helpful to rewrite the encompassing test using standard notation; let two non-nested hypotheses, H_0 and H_1, be defined as follows:

H_0: Hours are I(0) and increase following technology shocks;

H_1: Hours are I(1) and decrease following technology shocks.

The testing strategy followed by the authors thus takes H_0 as given (simulates data from the first estimated VAR) and verifies whether applying H_1 (first-differencing the simulated data) leads to the conclusion that hours increase following a technology shock. The test is repeated after interchanging the two hypotheses. Written in that way, the procedure resembles some tests for non-nested hypotheses, such as the Cox test.

industries they analyze,[4] favourable technology shocks lead to increases in hours in more than three out of four cases. These results, however, are arrived at by identifying technology shocks with the permanent innovations to total factor productivity (TFP) measures, instead of the permanent innovations to labour productivity, as Galí's paper proposes.[5]

Conclusions

I very much appreciated reading this paper. I learned a great deal about the effect of technology shocks (and business cycles, more generally) on economic variables. The authors' conclusion—that, on balance, the data support the view that hours worked increases following favourable technology shocks—receives support from a variety of sources. Their paper is therefore certain to rekindle the debate that originated with Galí's result. I believe that an interesting way to enrich this debate would include a thorough examination of specific frequencies at which the authors' and Galí's results might, in turn, be supported by the data.

References

Baxter, M. and R.G. King. 1999. "Measuring Business Cycles: Approximate Band-Pass Filters for Economic Time Series." *The Review of Economics & Statistics* 81 (4): 575–93.

Blanchard, O.J. and D. Quah. 1989. "The Dynamic Effects of Aggregate Demand and Supply Disturbances." *American Economic Review* 79 (4): 655–73.

Chang, Y. and J.H. Hong. 2003. "On the Employment Effect of Technology: Evidence from US Manufacturing for 1958–1996." University of Pennsylvania. Manuscript.

Christiano, L.J., M. Eichenbaum, and C.L. Evans. 2001. "Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy." NBER Working Paper No. 8403.

Christiano, L.J., M. Eichenbaum, and R. Vigfusson. 2003. "What Happens After a Technology Shock?" Board of Governors of the Federal Reserve System International Finance Discussion Paper No. 768.

Galí, J. 1999. "Technology, Employment, and the Business Cycle: Do Technology Shocks Explain Aggregate Fluctuations?" *American Economic Review* 89 (1): 249–71.

---

4. They examine over 300 subcategories of the manufacturing industry.

5. When Chang and Hong follow Galí exactly and identify technology shocks as the permanent component in labour productivity, hours fall following technology shocks in most industries. Chang and Hong suggest, however, that using labour productivity as the basis for the identification of the technology shocks introduces effects linked to relative price movements rather than TFP changes.