Discussion

Scott Hendry

I very much enjoyed reading the paper and believe it makes useful contributions to the literature. I have no major criticisms; consequently, my comments will address details of the model and the results. I will also make a few recommendations for extensions.

The Alexopoulos paper embeds an efficiency wage model in an otherwise relatively standard limited-participation model. The goal is to show that such a model does not require the usual unrealistic assumptions of high labour supply elasticity or high price markups in order to reproduce the stylized fact from U.S. data that employment is highly variable relative to real wages. Other work in the literature has proposed indivisible labour, search frictions, or home production as frameworks for generating large employment variations and much smaller real-wage fluctuations. The proposed efficiency wage model is standard, except that workers are not fired when they are discovered shirking. Instead, firms simply do not pay workers their quarterly bonus at the end of the period. This small change to the model greatly improves its ability to replicate the relative volatilities of employment and real wages. The degree of income insurance for unemployed workers also plays an important role in the model.

The paper estimates many of the model's parameters, using a Generalized Method of Moments procedure, which is still relatively uncommon in the dynamic general-equilibrium (DGE) literature. A number of important parameters, such as the adjustment-cost parameters and some efficiency wage parameters, are still calibrated, and an attempt should be made to estimate them in future work.

Another contribution of the paper is the addition of a model of equilibrium unemployment into a standard DGE model. Although a couple of papers have completed similar exercises, particularly with search models of unemployment (see Andolfatto 1996), it is quite surprising that models with unemployment, such as this one, are still not standard for important macroeconomic analysis. Consequently, this model is an important step forward, especially for central bankers concerned with the behaviour of the Phillips curve.

The paper argues that the model with partial income insurance better replicates the desired relative volatilities of employment and real wages. One could argue, however, that the full-insurance model better replicates the data because it gets closer to the relative size of the two variances and much closer to the absolute variance for employment, while giving up only a little extra in terms of real-wage variance. The full-insurance model also seems better because it gets closer to the data for the relative volatilities of employment and output. In addition, the real wage, under partial insurance and using U.S. data,[1] has a countercyclical flavour in the response to a monetary policy shock, while the full-insurance case has a procyclical response as in the data. Given these results, one could argue that the full-insurance model better replicates the data. However, we all know that the real world is not characterized by full-insurance schemes. I would be interested in any explanation of why the full-insurance model seems to do so well. Also, given that the model is estimated, it would be good to simply let the data speak and estimate the degree of income insurance. I understand that other second moments not shown in the paper do lead to a preference for the partial-insurance model. It is still surprising, however, that the partial-insurance model does not more obviously outperform the full-insurance case.

The paper focuses on replicating employment and real-wage variability. I would also be interested in the model's ability to replicate the dynamics of unemployment. It is quite possible that, given the model has no labour supply variation so that employment changes are simply the inverse of unemployment changes, the model replicates employment well but unemployment poorly. I understand from discussions with the author that unemployment dynamics are reasonably well approximated, but I would like to see evidence in the paper.

Similarly, it would be useful to allow for variation in both the intensive and extensive labour margins. Granted, the data imply that most variation is in employment, as in the model, but there is also some variation in hours worked by each employee. Since the model is estimated, the data could be

1. As I understand it, this is not a robust feature of the model, but it does occur in some specifications.

used to decide how much variation to attribute to hours, employment, the labour supply, and possibly, effort.

Impulse responses and the relative volatilities of employment and wages are the main criteria used. I would also like to see how well the model tracked the actual movements of these series over time. This could be done by simply looking at a graph of actual and fitted data. When does the model fit well? When does it fit poorly? The estimated shock series could also be examined to determine whether they conform with our expectations of when the economy was hit by large fiscal, technology, and monetary policy shocks. A more structured comparison would be to estimate some cross-correlation coefficients along with the variances. For instance, the correlation of employment with output and real wages would be most interesting. Since the real wage is countercyclical for some specifications of the model (using U.S. data) following a monetary policy shock, the model may not replicate these correlations very well.

Another metric for evaluating the models would be to do a variance decomposition of the data to determine the sources of most of the volatility in output, inflation, and employment. There are numerous other studies with which these results could be compared in terms of output and inflation. Not too many have done this for employment, however. Finally, stability tests of the parameters would lend credence to the estimates as well as to the model's policy predictions.

I was intrigued by the model's hump-shaped real-impulse responses following a monetary policy shock. In contrast, there was no hump-shaped response for either fiscal or technology shocks. It is interesting that there is more real rigidity or delayed response for a nominal shock than for a real shock. I believe that this delayed real response comes from the portfolio-adjustment cost and not from the labour market frictions. That is why this hump shape is more evident following a monetary shock (when this cost matters most) than it is for the other shocks. This interpretation is supported by the fact that the basic limited-participation model also generates the same hump shape although with a smaller peak. The efficiency wage frictions may generate unemployment, but they do not add much in the way of persistence to either employment or real wages.

This reasoning carries over to my next point—that it would be very interesting to add a flavour of labour market search on top of the efficiency wages. Search frictions could also be used as a way of avoiding the necessity of assuming high labour supply elasticities in DGE models. With both types of frictions in the same model, one could estimate the relevant parameters and determine the relative contributions of search and efficiency wages to output, employment, and real-wage dynamics.

In the same vein, one could estimate more of the efficiency wage parameters in the model to better judge the model's applicability. As I mentioned earlier, the size of the bonus, *s*, should be estimated instead of calibrated. It would also be an improvement if this differed from the degree of insurance coverage that was also estimated.

In another efficiency wage model, Burnside, Eichenbaum, and Fisher (2000) found that variable marginal taxes (instead of lump sum taxes) made it difficult to match the dynamics of employment following a fiscal policy shock. In particular, their model generated a result that hours worked initially rose following a positive fiscal shock but quickly reversed and reached its trough at approximately the same time that the fiscal shock reached its peak. There is no evidence in the data that hours worked declines in this manner. The authors needed to fix their marginal tax rates and vary lump sum taxes to pay for the fiscal shock, much as the Alexopoulos paper does, in order to get hours worked to respond only positively. It would be worth investigating how the current model performs with marginal taxes on wages and consumption instead of lump-sum taxes.

The paper could be expanded to better explain its support for the efficiency wage explanation over the indivisible labour, search friction, or home production explanations of the volatilities of employment and real wages. It is not clear why we should believe this explanation more than the others. I therefore believe it would be worthwhile to embed a couple of these factors within the same model and then estimate it. A discussion of the overlap and differences between these explanations or models would also be useful.

One final modification of the model would be to assume that monetary policy follows a Taylor-like interest rate rule instead of an AR(1) in money growth. This would obviously be a more accurate representation of monetary policy and, in numerous studies, it has been shown to substantially affect the model's dynamics.

I would like to emphasize that this paper tackles an interesting question in an innovative manner. My enthusiasm for the work is reflected in my desire to see the results of my suggested extensions.

References

Andolfatto, D. 1996. "Business Cycles and Labor-Market Search." *American Economic Review* 86 (1): 112–32.

Burnside, C., M. Eichenbaum, and J.D.M. Fisher. 2000. "Fiscal Shocks in an Efficiency Wage Model." NBER Working Paper No. 7515.