Discussion of "*The Business Cycle and the Life Cycle*" by Gomme, Rogerson, Rupert, and Wright

Robert Shimer

Department of Economics

University of Chicago

and NBER

shimer@uchicago.edu

June 11, 2004

1 Introduction

This is an ambitious paper. The authors extend the standard real business cycle model in two directions. First, they allow for both market production and home production, as in Benhabib, Rogerson, and Wright (1991). Second, they allow for overlapping generations of finitely-lived agents, as in Rios-Rull (1996). They then compute the aggregate implications of the overlapping generations model for cyclical fluctuations, focusing on the relative volatility of 'market hours' (the number of hours that the representative agent spends working in the market sector) and 'market output' (the output that she produces in the market sector). The main question they ask is whether this generalization of the real business cycle model can explain the observed differentials in the cyclical fluctuations in hours across age groups in the United States. I will use most of my discussion to address this question, but an initial digression is useful.

2 Representative Agent Model

The representative agent model which the authors develop in Section 2 performs remarkably well. Depending on the intertemporal elasticity of labor supply $\frac{1}{\gamma-1}$ and the elasticity of substitution between home and market goods $\frac{1}{1-\xi}$, the authors can easily match the relative volatility of market hours and market output (see their Table 2).[1] A critical question is therefore which values of these parameters are reasonable?

A well-established microeconomics literature starting with MaCurdy (1981) concludes from the life cycle behavior of wages and hours worked in the market that an appropriate value for γ is close to infinity, at least for men, but recent work has questioned that finding. For example, Keane and Imai (2003) argue that MaCurdy and followers neglect an important component of the compensation of younger workers, human capital accumulated at work. After correcting for this, they find that the intertemporal elasticity is close to 4, i.e. γ is approximately 1.25. More to the point of this paper, Rupert, Rogerson, and Wright (2001) conclude that hours worked at home is an important omitted variable in MaCurdy-type regressions and show that including home work raises the estimated value of γ considerably. Both of these arguments seem quite convincing, and so values of γ not much larger than 1 are plausible.

There is much less evidence on the parameter ξ, that is, on the elasticity of substitution between home and market goods, although this parameter is also critical to the performance of the model. If home and market goods are strong complements, a decrease in market productivity induces workers to reduce the time they spend producing the complementary home goods, further reducing the cyclical fluctuations in market hours. The introduction of home goods amplifies the cyclical fluctuations in market hours only if the elasticity of substitution exceeds 1. So what is a reasonable value of ξ? To my knowledge, only two papers have tackled this question. McGrattan, Rogerson, and Wright (1997) pin down ξ using macro data, showing that

the model requires a high elasticity of substitution $\frac{1}{1-\xi}$ in order to match the behavior of important aggregate variables, including market hours and consumption and home capital. But this is analogous to saying that we know that the intertemporal elasticity of substitution is high because we observe that market hours fluctuate a lot over the business cycle. It does not provide independent evidence on the empirical relevance of the particular model. One wants to use microeconomic evidence to calibrate macroeconomic models.

Rupert, Rogerson, and Wright (1995) provide the best such evidence, but unfortunately their estimates are imprecise. They start by estimating a fairly complicated home production model, allowing for the possibility that home and market goods are imperfect substitutes, home and market hours are imperfect substitutes, and the production of home goods is a concave function of home hours. Perhaps not surprisingly, their estimates of this very general model are imprecise. Despite this, their point estimates suggest that the elasticity of substitution between home and market goods is economically indistinguishable from 1 for single men and married couples, although it is larger for single women. This would seem to be a significant blow for the usefulness of the home production model in thinking about macroeconomics. But in footnote 13, the authors impose that home and market hours are perfect substitutes and restrict the curvature of the home production function exogenously. They show that the estimates of ξ for married couples range between 0.2 and 0.3, somewhat less than the numbers that Gomme, Rogerson, Rupert, and Wright use in the current paper, but still significantly larger

than zero. Based on this evidence, or lack of evidence, I think it is fair to say that the jury is still out on the true value of ξ, a viewpoint that seems to contrast with that of the authors.

3 Life Cycle Model: Cyclical Fluctuations

The heart of this paper is the analysis of the behavior of employment volatility conditional on a worker's age. The authors use data from the March Current Population Survey (CPS) from 1962 to 2000 to document that the cyclical volatility of market hours is almost four times as high for teenagers as it is for prime age workers.[2] This volatility decreases monotonically until approximately age 50, but is twice as high for workers over age 65 compared to workers age 45 to 64. I've constructed a similar measure using times series for age-contingent employment-population ratios, constructed by the Bureau of Labor Statistics from the basic (monthly) CPS from 1948 to 2003. The results are remarkably similar to those in the paper. It is also worth noting that the relative volatility of different age groups is extremely stable over time.

Looking at this data and at the behavior of the aggregate model, the obvious hypothesis is that the aggregate model can satisfactorily explain the behavior of prime age workers but does a poor job of explaining the behavior of younger and older workers. Of course, to test this hypothesis, it is necessary to write down a model in which workers of different ages

interact. The overlapping generations model in Sections 4 and 5 is an obvious benchmark.

How well does the model perform? The authors claim that "the model's ability to account for fluctuations in [market] hours increases as we consider older age groups." At some level this is correct. In fact, the model predicts more hours volatility for workers age 55 to 64 than there is in the data. But this conclusion is misleading. The model ignores teenagers, the group with the greatest fluctuations in market hours. It imposes mandatory retirement at age 65, and so cannot hope to match the high volatility of market hours for the oldest workers. And for the age ranges that are considered in the model, the standard deviation of market hours is basically a decreasing function of age in the data and an increasing function of age in the model (Table 14). A more accurate conclusion is that the life cycle model can explain virtually none of the age pattern of fluctuations in hours. If anything, the life cycle model predicts the opposite of what is observed in the data.

It is not particularly surprising that the life cycle model predicts little of the variation in the observed fluctuations in hours. As my discussion of the aggregate model should have made clear, there are two important determinants of the volatility of market hours: the intertemporal elasticity of labor supply $\frac{1}{\gamma-1}$ and the elasticity of substitution between home and market goods $\frac{1}{1-\xi}$. Although the authors allow different age workers to have different preferences over market goods, home goods, and leisure, and they allow different age workers to be endowed with labor that is more or less

productive, they do not allow either of the elasticities to vary with age. I suppose this puts discipline on the theoretical exercise, but in light of the results, the obvious reconciliation between the model and data is to allow for the possibility that younger workers have a more elastic labor supply or are more willing to substitute between market and home goods than are prime age workers.

Of course, one would like some direct microeconomic evidence in support of this hypothesis. It seems impossible to measure age variation in the intertemporal elasticity of labor supply using MaCurdy's (1981) methodology, so this might be untestable. But the authors could have easily extended Rupert, Rogerson, and Wright (1995) to examine how the elasticity of substitution between home and market goods varies with age. Introspection suggests to me that as people age and have children, they become less willing to substitute between any goods, in particular market and home goods, a pattern that may help to reconcile the model and data.[3] Conversely, if the data do not show differential elasticities of substitution, I would again use the authors' words against them: "In looking for alternative theories to better account for aggregate labor market fluctuations, attention should be directed toward features that specifically affect individuals during the first half of their life." If the elasticity of substitution is the same over the life cycle, attention is best diverted away from home production models.

4 Life Cycle Model: Secular Trends

There are some very interesting facts lurking in the shadows of this paper, in the secular trends market hours. Although the number of hours worked per adult has not shown any trend in the U.S. during the past 55 years,[4] this is not true for particular age groups. Figure 1 shows that in 1948, the employment-population ratio for people over age 65 stood at over 26 percent. This fell steadily until around 1990, reaching as low as 10 percent before increasing slightly in the last fifteen years. The decline for older men has been even more dramatic. The same data indicate the opposite pattern for workers age 20 to 54, a secular increase in employment due to a sharp increase in women's increased labor force participation partially offset by a moderate decline in employment for prime age men. Of course, this is not news to at least one of these authors, who has written "the number of average weekly hours of market work per person in the United States since World War II ... has been roughly constant; for various groups, however, it has shifted dramatically from males to females, from older people to younger people, and from single- to married-person households." (McGrattan and Rogerson, 1998).

[Figure 1 about here.]

Why does this matter for Gomme, Rogerson, Rupert and Wright? I am not simply saying that they should have written about secular trends instead of business cycle fluctuations because the trends are more interesting than the cycle. In fact, it is conceivable that the same forces that explain

the differential cyclical fluctuations in employment are also important for understanding the differential secular trends. But the secular trends raise a major concern: how does one calibrate such a model? The authors write "as is standard, we follow the procedure of requiring that parameter values are such that the model's deterministic steady state matches the time series averages for several aggregate variables." But if the aggregate variables are trending over time, does that mean that the calibrated parameters must also trend over time? It seems they must, which casts doubt on the discipline of the calibration exercise.

Conversely, the secular trends contain a lot of information that the authors ignore. McGrattan and Rogerson (1998) claim that changes in social security benefits, in fertility rates, and in family structure are critical for understanding secular changes in the employment-population ratio. In light of these changes, should we be surprised that the cross-sectional pattern of volatility is so stable over time? For example, if the cross-sectional pattern of volatility is due to different elasticities of substitution, then why did changes in fertility and family structure not alter the age-conditional elasticity of substitution and therefore rearrange the cross-sectional pattern of volatility? In other words, why is cyclical volatility so stable over time, even in the presence of the changes that induced large secular shifts in employment? Unfortunately, the paper does not answer this question.

5 Conclusion

Let me conclude by saying what I think we learn from this exercise. First, the authors make a convincing case that any model that purports to explain employment fluctuations must be able to explain why employment fluctuates more for younger workers and workers over the age of 65 than it does for prime age workers. Second, they carefully describe and solve one particular model that, based on this criterion, cannot explain employment fluctuations: the real business cycle model extended to allow for home production and overlapping generations. The next step is to explain what type of model can explain the cross-sectional pattern of employment fluctuations. That is an interesting and important question that I suspect will continue to occupy researchers' attention for many years.

6 References

Jess Benhabib, Richard Rogerson, and Randall Wright, "Homework in Macroeconomics," *Journal of Political Economy*, 1991.

Michael Keane and Susumu Imai, "Intertemporal Labor Supply and Human Capital Accumulation," forthcoming, *International Economic Review*, 2004.

Thomas MaCurdy, "An Empirical Model of Labor Supply in a Life Cycle Setting," *Journal of Political Economy*, 1981.

Ellen McGrattan and Richard Rogerson, "Changes in Hours Worked Since 1950," *Federal Reserve Bank of Minneapolis Quarterly Review*, 1998.

Ellen McGrattan, Richard Rogerson, and Randall Wright, "An Equilibrium Model of the Business Cycle with Household Production and Fiscal Policy," *International Economic Review*, 1997.

Victor Rios-Rull, "Life-Cycle Economies and Aggregate Fluctuations," *Review of Economic Studies*, 1996.

Peter Rupert, Richard Rogerson, and Randall Wright, "Using Panel Data to Estimate Substitution Elasticities in Household Production Models," *Economic Theory*, 1995.

Notes

[1]The authors do not ask whether the model can explain the absolute volatility of the variables. Presumably the answer depends on whether cyclical fluctuations in the Solow residual represent a primitive technology shock.

[2]The authors measure the volatility of hours as the standard deviation of the detrended hours series for a particular age group projected on the detrended hours series for the overall population.

[3]Of course, this only makes the high volatility of hours for older workers more puzzling. Explaining this first requires a serious model of the retirement decision.

[4]I measure the employment-population ratio using the basic CPS and average hours worked for production workers using the Current Employment Statistics (CES). The product of these two numbers is a rough measure of average hours per person. Between 1964 (the first year when hours data are available) and 1983, the average person worked 21.1 hours per week. Over the next twenty year period, this increased to 21.5 hours per week. There is no evidence of a secular trend in this variable.

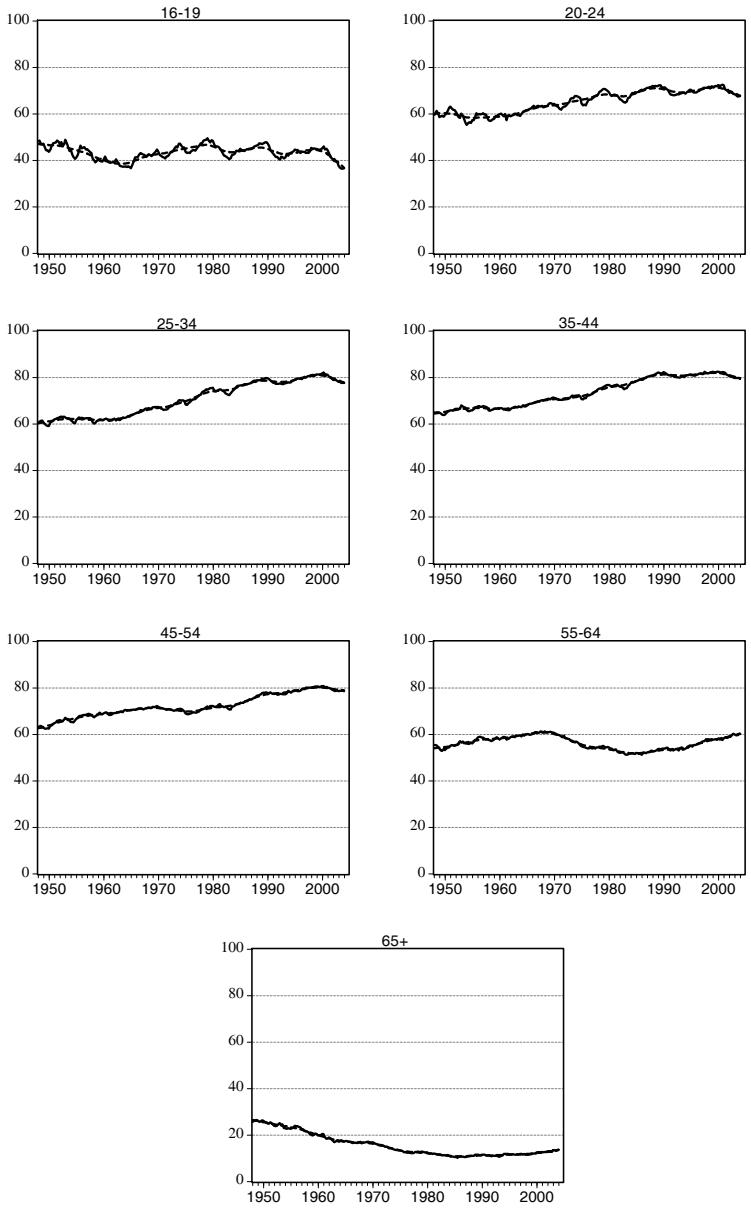

Figure 1: Employment-population ratio as a function of age, 1948–2003. The dashed line is an HP filter with smoothing parameter 1600 on quarterly data. The data were constructed by the Bureau of Labor Statistics from the CPS.