

Christiano
Economics 416
Advanced Macroeconomics
Take home midterm exam.

1 Explaining Labor Market Volatility

The purpose of this question is to explore a labor market puzzle that has bedeviled business cycle researchers for years. The problem is to produce a sensible model that generates the amount of labor market volatility that we observe in the data. The first step is to show that there is a problem in the standard real business cycle model. The second step is to document a particular diagnosis of the problem, namely that it reflects excessive movement in the wage. The third step is to introduce firm/worker bargaining over the wage and show that this opens a possible route for solving the problem. The example is inspired by Hagedorn and Manovskii's 2008 AER paper in which they showed that if the unemployment compensation of the worker is high enough, then wages could be smooth enough and, hence, employment volatile enough, to match the data. Hagedorn and Manovskii's posited explanation has been criticized on the ground that real-world agents' outside option is not as great as Hagedorn and Manovskii's explanation requires. The problem is that if workers' outside option is reduced to levels that the critics argue is empirically plausible, then it is claimed (see Shimer's 2005 AER paper) that the bargaining model loses the ability to account for the volatility of labor markets. We will pursue these ideas further later in the course, by drawing attention to the observations in Hall and Milgrom's 2008 AER paper.

1.1 Real Business Cycle Model

Consider the following real business cycle model. At time t , the representative household maximizes

$$E_t \sum_{j=0}^{\infty} \beta^j [\log C_{t+j} + \psi \log(1 - N_{t+j})],$$

subject to

$$C_t + K_{t+1} - (1 - \delta) K_t \leq r_t K_t + w_t N_t,$$

for all t . Here, r_t denotes the rental rate of capital and w_t denotes the wage rate. The household is 'small' and takes the prices as given. There is a representative firm. In period t , the firm maximizes by choice of K_t and N_t , its profits:

$$Y_t - r_t K_t - w_t N_t,$$

subject to its technology:

$$Y_t = K_t^\alpha [\exp(a_t) N_t]^{1-\alpha},$$

where

$$a_t = \rho_a a_{t-1} + \varepsilon_t.$$

Here, ε_t is iid with mean zero and $E\varepsilon_t^2 = \sigma^2 = 0.01^2$, $\rho_a = 0.95$. Also, $\beta = 1.03^{-1/4}$, $\alpha = 0.36$, $\delta = 0.025$. Finally, assign a value to ψ which implies that $N_t = 1/3$ in steady state, given the setting of the other parameters. That is, the representative household works one-third of available time.

1.1.1 Questions

1. Use Dynare to solve the model and simulate 1,000 observations on log output and log employment (work with the ‘periods=1000’ command in `stoch_simul`). Detrend these two series using the HP filter. Compute the standard deviation of the result. Display the ratio of the standard deviation of (filtered) employment to the standard deviation of (filtered) output. This ratio is called the ‘relative volatility of employment to output’.
2. Go to the web-based data base of the Federal Reserve Bank of St. Louis (FRED) and retrieve data on Real Gross Domestic Product (GDP) and employment, All Employees: Total nonfarm (take these data quarterly). Do to these data what you did to the model data. Display the relative volatility of employment to output in the data.
3. You will see that the relative volatility of employment is much higher in the data than it is in the model. This failing of the model has attracted a lot of attention. One interpretation is that it reflects wages rise too much in the wake of a shock that causes output to expand. Explore this hypothesis by returning to question 1. Fix the wage rate exogenously at its steady state value and assume that firms are always on their labor demand schedule, while households always supply all the labor that is demanded (this implies that sometimes they work more than they want). What happens to the volatility of employment with this change in the model? Much of the macro labor supply literature is about trying to reproduce the properties of this sticky wage model. But, economists prefer if they can arrive at this by some endogenous mechanism.

1.2 Real Business Cycle Model with Nash Bargaining

Assume that the representative household has a unit mass of workers. Each worker goes to the labor market. A fraction, N_t , of the workers meet a firm and are employed. The complementary fraction is unemployed. Work effort is indivisible: a worker either works or not. There is perfect insurance inside the household and each worker enjoys the same level of consumption, C_t . Each employed worker brings home the wage, w_t , and each unemployed worker brings home an unemployment payment, D . The household problem is to maximize:

$$E_0 \sum_{t=0}^{\infty} \beta^t \log(C_t),$$

subject to

$$C_t + K_{t+1} - (1 - \delta) K_t \leq w_t N_t + r_t K_t + (1 - N_t) D - T_t + \pi_t.$$

Here, π_t denotes firm profits (discussed below) and T_t denotes taxes raised to finance government unemployment payments. The government budget constraint is:

$$(1 - N_t) D = T_t.$$

The flow value of labor effort corresponds to what an employed worker contributes to household utility in period t , measured in consumption good units. Because workers are assumed not to suffer any disutility from labor effort, the flow value of an employed worker is the period t wage rate, w_t . The value of being a worker, V_t , is:

$$V_t = w_t + E_t m_{t+1} [\rho V_{t+1} + (1 - \rho) (f_{t+1} \bar{V}_{t+1} + (1 - f_{t+1}) U_{t+1})]. \quad (1)$$

Here, m_{t+1} is the household's stochastic discount factor, $\beta u_{c,t+1}/u_{c,t}$, and the object in square brackets indicates the various things that can happen to the worker in period $t + 1$. Thus, with probability ρ the worker remains matched to the same firm in $t + 1$ with probability ρ , in which case the value of the worker is V_{t+1} . With probability $1 - \rho$ the worker separates from the firm, in which case there are two possibilities. With probability f_{t+1} the worker matches immediately with another firm where the worker will receive value, \bar{V}_{t+1} . Because all the firms and workers are the same, in equilibrium we have $V_{t+1} = \bar{V}_{t+1}$. With probability $1 - f_{t+1}$ the separated worker goes into unemployment in $t + 1$ and we denote the value of such a worker by U_{t+1} .

The flow value of an unemployed worker's contribution to household utility, in consumption units, is D . These are the goods that the government gives to workers that are unemployed. The value of being an unemployed worker is, then,

$$U_t = D + E_t m_{t+1} [f_{t+1} V_{t+1} + (1 - f_{t+1}) U_{t+1}]. \quad (2)$$

In period $t + 1$ the period t unemployed worker is employed with probability f_{t+1} and is unemployed with probability $1 - f_{t+1}$.

There are two types of firms in this variant of the model. There are the firms (we'll call them *RBC firms*) that look just like their cousins in the real business cycle model. They operate the Cobb-Douglas production function. To do so, they rent capital, K_t , and a second input which we denote by h_t . The RBC firms hire K_t and h_t , in competitive markets. The second type of firm, we call them the *bargaining firms*, are endowed with the knowledge of how to convert one unit of labor power into one unit of h_t . RBC firms and bargaining firms interact in competitive markets.

We now discuss bargaining firms in greater detail. The value to the firm of an employed worker is denoted J_t :

$$J_t = \vartheta_t - w_t + \rho E_t m_{t+1} J_{t+1}, \quad (3)$$

where ϑ_t denotes the (competitively determined) market value of the one unit of h_t produced by the worker and sold by the bargaining firm.

The number of employed workers in period t is denoted l_t and this evolves as follows:

$$l_t = (\rho + x_t) l_{t-1}. \quad (4)$$

Here, ρ corresponds to the exogenous rate at which employed workers are separated from their firms at the end of the period. Also, x_t denotes the hiring rate so that the number of new hires in period t is equal to $x_t l_{t-1}$. The job finding rate is defined as the ratio of the number of workers that find jobs, divided by the number of workers looking for a job:

$$f_t = \frac{x_t l_{t-1}}{1 - \rho l_{t-1}}. \quad (5)$$

The denominator term, the number of workers looking for work as the start of t , is composed of the people unemployed in period $t - 1$, $1 - l_{t-1}$, plus the number of people who were employed at $t - 1$, but were separated from their firms, $(1 - \rho) l_{t-1}$. Thus the sum of people searching for work at the end of $t - 1$ is

$$1 - l_{t-1} + (1 - \rho) l_{t-1} = 1 - \rho l_{t-1}.$$

A firm that wishes to meet with a worker can do so by paying a fixed cost, κ .¹ Free entry implies that bargaining firms cannot make profits by hiring a worker, so that

$$J_t = \kappa$$

We assume that the wage is determined by Nash bargaining:

$$J_t = \frac{1 - \eta}{\eta} (V_t - U_t).$$

Here, η is the share of the total surplus, $J_t + V_t - U_t$, given to workers. The object, J_t , is the surplus of the firm. It is what the firm gets by employing the worker (i.e., J_t), minus what it gets if it does not employ the worker (i.e., nothing). Similarly for the surplus of the workers.

The flow profits, π_t , sent by bargaining firms to the households (their owners) is:

$$\pi_t = (\vartheta_t - w_t) l_t - \kappa x_t l_{t-1}.$$

The first term is firms' period t net revenues from employing labor. The second term represents the expenses incurred by firms in period t for the purpose of

¹For a discussion of this assumption, see section 5 in Pissarides (Econometrica, Vol. 77, No. 5, September, 2009, pp. 1339–1369). For another, see Christiano, Trabandt and Walentin, Journal of Economic Dynamics & Control, vol. 35, 2011, pages 2038–2039. The assumption differs from the standard one in the DMP literature, where it is typically assumed that the costs of meeting a worker are harder when the economy is more active than when it is in recession. An important question is whether the costs of hiring have an important cyclical component or not. For an empirical study that favors the sort of assumption made here, see Yashiv, 2000, "The determinants of equilibrium unemployment," American Economic Review 90 (5), 1297–1322.

meeting workers. Typically, $\kappa > \vartheta_t - w_t$, so that a firm that meets a worker in period t sends negative dividends to its owners in that period. Households are willing to put up with these negative dividends because they expect to be compensated by positive dividends in the future, according to $J_t = \kappa$.²

The resource constraint in this economy is:

$$C_t + \kappa x_t l_{t-1} + K_{t+1} - (1 - \delta) K_t \leq K_t^\alpha [\exp(a_t) h_t]^{1-\alpha},$$

for all t . Here, $\kappa x_t l_{t-1}$ denotes the goods purchased by bargaining firms to address their hiring costs. Also, h_t is the quantity of input goods purchased by RBC firms from bargaining firms. Clearing in that market and in the labor market requires

$$l_t = h_t = N_t.$$

To save notation, you may just as well use l_t wherever h_t or N_t appears, so that the resource constraint is:

$$C_t + \kappa x_t l_{t-1} + K_{t+1} - (1 - \delta) K_t \leq K_t^\alpha [\exp(a_t) l_t]^{1-\alpha}.$$

Finally, there is the discount factor, m_{t+1} :

$$m_{t+1} = \beta \frac{C_t}{C_{t+1}}.$$

1.2.1 Questions

1. Verify Walras' law for this economy. In particular, the expression for firm profits, the government budget constraint and households' budget constraint imply the resource constraint.
2. The parameters of the RBC part of the model are, as before,

$$\alpha = 0.36, \beta = 1.03^{-1/4}, \rho_a = 0.95, E\varepsilon_t^2 = 0.01^2, \delta = 0.025.$$

Set the persistence of job matches, ρ , to 0.90. The other parameters of the bargaining part of the model,

$$D, \kappa, \eta,$$

should be set so that, given the other parameter values, the following is true in steady state:

$$\begin{array}{ll} \frac{\kappa x l}{GDP} = 0.01 & \text{hiring costs as a fraction of GDP } (\equiv C + \delta K) \\ \frac{D}{w} = 0.98 & \text{replacement ratio} \\ u = 1 - l = 0.055 & \text{unemployment rate} \end{array}$$

Display formulas for the steady state of the model, including for D, κ, η . Report the steady state values of $K, C, x, w, \vartheta, U, V, f, r$.

²In effect, meeting a worker requires an investment, much like the accumulation of physical capital. As a result, it would be interesting to explore various financial frictions that might make the employment decision more interesting.

3. Generate 1,000 observations on log GDP and log employment in the model, HP filter the result as you did in the RBC model and compute the standard deviations of the result. Does this model do a better job at accounting for the observed volatility of labor? To help answer this question, compute and display on the same graph the impulse response to technology shock implied by the RBC model and the RBC model with bargaining.
4. Discuss the role played in the analysis of the high replacement ratio. Do this by repeating 2 with replacement ratios of 0.99 and 0.97.

2 Monte Carlo Markov Chain

The idea here is to explore the accuracy of the Metropolis-Hastings algorithm for approximating a distribution. We'll specify a density function for a single random variable. Since we know the density function we're trying to approximate, we'll be able to assess the quality of the approximation provided by the MH algorithm. In applying the algorithm, we'll need to compute a second derivative. Do this numerically. Here is a formula for a function, $f(x)$:

$$f''(x) = \frac{f(x + 2\varepsilon) - 2f(x) + f(x - 2\varepsilon)}{4\varepsilon^2},$$

for ε small (for example, you could set $\varepsilon = 0.000001$.)

In this exercise we make the test of the MH algorithm pretty tough by specifying that the true density function is bimodal. At the same time, the test will be very weak because we are working with a one-dimensional random variable. We consider a mixture of normal distribution. This density function is a linear combination of two normals, where the weights in the linear combination are denoted π and $1 - \pi$. The object, π , is the probability that a variable is drawn from the first Normal distribution, and $1 - \pi$ is the probability of drawing from the second Normal distribution. Suppose the i^{th} Normal has mean and variance, μ_i and σ_i^2 , respectively, $i = 1, 2$. In addition, suppose

$$\mu_1 = -0.06, \mu_2 = 0.06, \sigma_1 = 0.02, \sigma_2 = 0.01, \pi = 1/2.$$

1. Let x denote the mixture of Normals random variable. Graph its density over the range, $x = -0.15$ to $x = 0.15$. Specify a very fine grid of values of x so that you get a very accurate graph of the mixture of Normals. If the grid is fine enough then you can pinpoint the value of x , x^* , where the density is highest by simply identifying the value of x on your grid that produces highest density (i.e., you can just use the max operator in MATLAB).
2. Compute the second derivative of the log density function around the mode, x^* . Apply the algorithm described in the handout to compute a sequence, $x^{(1)}, x^{(2)}, \dots, x^{(M)}$, where $x^{(1)} = x^*$. Choose the scalar in the jump distribution so that you get a acceptance rate of 23 percent when

$M = 1,000$. Graph the histogram of $x^{(1)}, x^{(2)}, \dots, x^{(M)}$. Scale the heights of the histogram bars so that the sum of the areas of the histogram rectangles equals unity. Only with this scaling will the histogram be comparable to the true density function, the mixture of Normals. Graph the properly scaled histogram on the same graph with the density of the true mixture of Normals you graphed in question 1. Don't make the graph a bar chart, make it a line graph so it looks like a density function. Does the MH algorithm produce a good-looking approximation?

3. In question 2, you should have found that the histogram is rather choppy. Now increase M to 100,000. You should find that you get an amazingly accurate approximation. Try different values of k , fixing M at, say 5,000. Does k affect accuracy much?
4. Center the MCMC algorithm around the second, lower, peak of the density function. Use $M = 100,000$. Does it make much difference that you did not center things on the actual mode?
5. Draw the Laplace approximation of the density function around the mode of the mixture of Normals (see the later sheets on the Bayesian inference handout for the Laplace approximation). How well does that approximate the distribution?

3 Equilibrium Indeterminacy and Sunspots in the Increasing Return Model

1. Consider the increasing returns model in homework #3. Choose a value of γ big enough so that you are in the case of indeterminacy, i.e., $q < n - l$ in the section, 'Non-Invertible a Case', in the handout, 'Solutions to Linear Expectational Difference Equations'. What are the values of q, l ? Display the matrix, \tilde{p} , and the matrix, $-D_1^{-1}D_2$. Suppose $k_0 = 0$, so that the initial capital stock is at its steady state value. One equilibrium is the steady state itself, i.e., $k_{t+1} = 0$ for $t = 0, 1, 2, \dots$. Construct an alternative deterministic equilibrium in which $k_0 = 0$ and $k_1 > 0$, and $k_t \rightarrow 0$ as $t \rightarrow \infty$ (hint: you can use the algorithm described in the handout). Graph c_t, n_t, k_{t+1} , $t = 0, 1, \dots, 10$ in the alternative equilibrium (here, c_t, n_t, k_{t+1} are all to be interpreted as deviations from steady state.) Explain the economic intuition for the existence of the two equilibria.
2. Construct a sunspot equilibrium for this model (hint: you can use the approach discussed in the handout). Draw 10,000 observations on the sunspot shock (i.e., ξ_t). Then, compute 10,000 observations on investment, consumption, employment and output, under the assumption that the initial stock of capital is in steady state. Express the simulated variables in percent deviations from steady state. Compute the standard deviations and cross correlations with output of consumption, investment

and employment (all measured in percent deviation from steady state). Does the sunspot model look like a good model of the US business cycle? Are there some dimensions on which it looks good and others on which it looks bad? (Hint: check the 1986 Minneapolis Fed Quarterly Review paper by Prescott for a quantitative characterization of the business cycle. You can interpret the percent deviations of the model data from steady state as corresponding to the log-deviation of variables from the Hodrick-Prescott trend reported in Prescott's paper, though we will criticize this procedure later in the course.)

Construct an alternative simulation (using the same computer-generated 'random' draw of ξ_t 's) by choosing a different value of the initial-period endogenous variables (recall from the handout that you have a degree of freedom on this dimension). Compare the statistics computed in the previous paragraph with the results obtained with this alternative simulation. Are the statistics in the two simulations very different? Explain. (In all cases, keep the initial stock of capital at its steady state value).

4 Sunspots in New Keynesian Model

Consider the linearized equilibrium conditions of the New Keynesian model presented in the handout, 'Solutions to Linear Expectational Difference Equations'. Consider the case where $\beta = 0.8$, so that there is a one-dimensional indeterminacy. It has been argued by Clarida-Gali-Gertler ('The Science of Monetary Policy'), that this model is a good representation of the US economy in the 1970s. They argue that the jump in inflation of that time can be interpreted as the kind of jump that can occur in response to a sunspot that is possible when $\beta = 0.8$. Generate a simulation of length 10,000 of a sunspot equilibrium. Compute the correlation between inflation, π_t , and output, y_t , in this simulation. What sign is it? Provide intuition for the indeterminacy in this model, and also for the sign of this correlation. During the 1970s in the US, output was very low while inflation was high. Does this seem like a good model of the 1970s in the US?