Graduate Macro Theory II:
Extensions of Basic RBC Framework

Eric Sims
University of Notre Dame

Spring 2012

1 Introduction

The basic RBC model – which is just a stochastic neoclassical growth model with variable labor – is the building block of almost all modern DSGE models. It fits the data well on some dimensions, but less well on others. In this set of notes we consider several extensions and modifications of the basic framework.

2 Common Extensions

This section works through a number of extensions designed to make the RBC model (i) more realistic and (ii) a better fit with the data.

2.1 Indivisible Labor

One failure of the RBC model is that it fails to generate sufficient volatility in hours of work. It also models hours in a rather unrealistic way that is at odds with reality – all fluctuations in hours come from the *intensive* margin (e.g. average hours worked) as opposed to the *extensive* margin (the binary choice of whether to work or not). In the real world most people have a more or less fixed number of hours worked; it is fluctuations in bodies that drive most of the fluctuation in total hours worked.

In reality, households face two decisions: (1) work or not and (2) conditional on working, how much to work. This is difficult to model because it introduces discontinuity into the decisions household make. Hansen (1985) and Rogerson (1988) came up with a convenient technical fix. Suppose that within period preferences of any household are:

$$u(c_t, 1-n_t) = \ln c_t + \theta \frac{(1-n_t)^{1-\xi} - 1}{1-\xi} \tag{1}$$

These preferences nest the basic specification that we've used already if $\xi = 1$. Recall that the Frisch labor supply elasticity is given by $(\xi\gamma)^{-1}$, where $\gamma = \frac{n^*}{1-n^*}$. Hence, a high value of ξ would correspond to a very low Frisch elasticity.

Suppose that the structure of the world is as follows. There are a large number of identical households. Households either work or they do not. If they work, they work $\bar{n}$ hours, with $0 < \bar{n} < 1$. The amount of work, $\bar{n}$, can be interpreted as a technological constraint and is exogenous to the model. Each period, there is a probability of working, τ_t, with $0 < \tau_t < 1$. This probability is indexed by t because it is a choice variable – essentially the household can choose its probability of working, but not how much it does work if it does work. There is a lottery such that the household has a τ_t chance of being selected to work; the rule of the game is there is perfect insurance in that every household gets paid whether they work or not. Hence, in expectation households will work $n_t = \tau_t \bar{n}$ and they will have identical consumption.

We can write out the households expected flow utility function as:

$$u(c_t, 1-n_t) = \ln c_t + \tau_t \theta \frac{(1-\bar{n})^{1-\xi} - 1}{1-\xi} + (1-\tau_t)\theta \frac{(1)^{1-\xi} - 1}{1-\xi} \tag{2}$$

We observe that preferences are linear in τ_t. Collecting terms we get:

$$u(c_t, 1-n_t) = \ln c_t + \tau_t \theta \left(\frac{(1-\bar{n})^{1-\xi} - 1}{1-\xi} - \frac{(1)^{1-\xi} - 1}{1-\xi} \right) + \theta \frac{(1)^{1-\xi} - 1}{1-\xi} \tag{3}$$

Now, from above we know that $\tau_t = \frac{n_t}{\bar{n}}$. Make this substitution:

$$u(c_t, 1-n_t) = \ln c_t + \frac{n_t}{\bar{n}} \theta \left(\frac{(1-\bar{n})^{1-\xi} - 1}{1-\xi} - \frac{(1)^{1-\xi} - 1}{1-\xi} \right) + \theta \frac{(1)^{1-\xi} - 1}{1-\xi} \tag{4}$$

As long as $\xi > 0$, then $\frac{1^{1-\xi}-1}{1-\xi} > \frac{(1-\bar{n})^{1-\xi}-1}{1-\xi}$. Hence, re-write this again as:

$$u(c_t, 1-n_t) = \ln c_t - \frac{n_t}{\bar{n}} \theta \left(\frac{(1)^{1-\xi} - 1}{1-\xi} - \frac{(1-\bar{n})^{1-\xi} - 1}{1-\xi} \right) + \theta \frac{(1)^{1-\xi} - 1}{1-\xi}$$

Let's define two constants as:

$$B = \frac{\theta}{\bar{n}} \left(\frac{(1)^{1-\xi} - 1}{1-\xi} - \frac{(1-\bar{n})^{1-\xi} - 1}{1-\xi} \right)$$

$$D = \theta \frac{(1)^{1-\xi} - 1}{1-\xi}$$

We can actually just drop D altogether from the analysis – adding a constant to the utility function won't change the household's optimal choices. Then we can write the within period utility function as:

$$u(c_t, 1-n_t) = \ln c_t - B n_t \tag{5}$$

In other words, utility effectively becomes linear in labor under this indivisible labor with lotteries framework. This holds for *any* value of ξ. But indeed, it is *as if* $\xi = 0$. In other words, the aggregate labor supply elasticity is *infinite* even if the micro labor supply is very small (i.e. ξ very large). This is potentially very helpful – one can generate more hours volatility with a higher Frisch elasticity, and is (potentially) not subject to the criticisms that the labor supply elasticity is inconsistent with micro evidence.

The full model can then be written as follows. I assume that households own the capital stock and lease it to firms. Households also have access to one period bonds. I abstract from the presence of these bonds in the firm first order condition because the quantity of bonds ends up being indeterminate anyway. As such, the firm problem becomes static.

Households:

$$\max_{c_t, n_t, b_{t+1}, k_{t+1}} \quad E_0 \sum_{t=0}^{\infty} \beta^t \left(\ln c_t - B n_t \right)$$

s.t.

$$c_t + k_{t+1} - (1-\delta)k_t + b_{t+1} - b_t = w_t n_t + R_t k_t + r_t b_t$$

Firms:

$$\max_{n_t, k_t} \quad a_t k_t^{\alpha} n_t^{1-\alpha} - w_t n_t - R_t k_t$$

In a competitive equilibrium the first order conditions hold and all budget constraints hold. This gives rise to the following characterization of equilibrium:

$$\frac{1}{c_t} = \beta E_t \left(\frac{1}{c_{t+1}} \left(R_{t+1} + (1-\delta) \right) \right) \tag{6}$$

$$\frac{1}{c_t} = \beta E_t \left(\frac{1}{c_{t+1}} (1 + r_{t+1}) \right) \tag{7}$$

$$B = \frac{1}{c_t} w_t \tag{8}$$

$$w_t = (1-\alpha) a_t k_t^{\alpha} n_t^{-\alpha} \tag{9}$$

$$R_t = \alpha a_t k_t^{\alpha-1} n_t^{1-\alpha} \tag{10}$$

$$k_{t+1} = I_t + (1-\delta) k_t \tag{11}$$

$$y_t = c_t + I_t \tag{12}$$

$$y_t = a_t k_t^{\alpha} n_t^{1-\alpha} \tag{13}$$

$$\ln a_t = \rho \ln a_{t-1} + \varepsilon_t \tag{14}$$

We would like to come up with a calibration of this model that is consistent with our previous calibrations. We don't actually need to calibrate anything that goes into B, just the value of B. But how do we do that?

From (6), combined with (10), we can solve for the steady state capital to labor ratio

$$\left(\frac{k^*}{n^*}\right) = \left(\frac{\alpha}{\frac{1}{\beta} - (1-\delta)}\right)^{\frac{1}{1-\alpha}} \tag{15}$$

Now combine (8) with (9) to solve for c^* in terms of the steady state capital to labor ratio:

$$c^* = \frac{1}{B}(1-\alpha)\left(\frac{k^*}{n^*}\right)^{\alpha} \tag{16}$$

Now go to the aggregate accounting identity, which can be written as:

$$c^* = n^*\left(\left(\frac{k^*}{n^*}\right)^{\alpha} - \delta\left(\frac{k^*}{n^*}\right)\right) \tag{17}$$

Since (16) and (17) must both hold, we can set them equal to one another and solve for B, taking $n^* = \frac{1}{3}$ as a given target in the calibration. We get:

$$B = \frac{(1-\alpha)\left(\frac{k^*}{n^*}\right)^{\alpha}}{n^*\left(\left(\frac{k^*}{n^*}\right)^{\alpha} - \delta\left(\frac{k^*}{n^*}\right)\right)} \tag{18}$$

If we use our now "standard" calibrations of $\alpha = 0.33$, $\beta = 0.99$, and $\delta = 0.025$, then we see that $\frac{k^*}{n^*} = 28.35$. This implies that $B = 2.63$.

I solve both this model and the standard RBC model (with log preferences over leisure and θ calibrated to guarantee $n^* = 1/3$). All other parameters are the same across both models. Below are impulse response to a technology shock in each model:

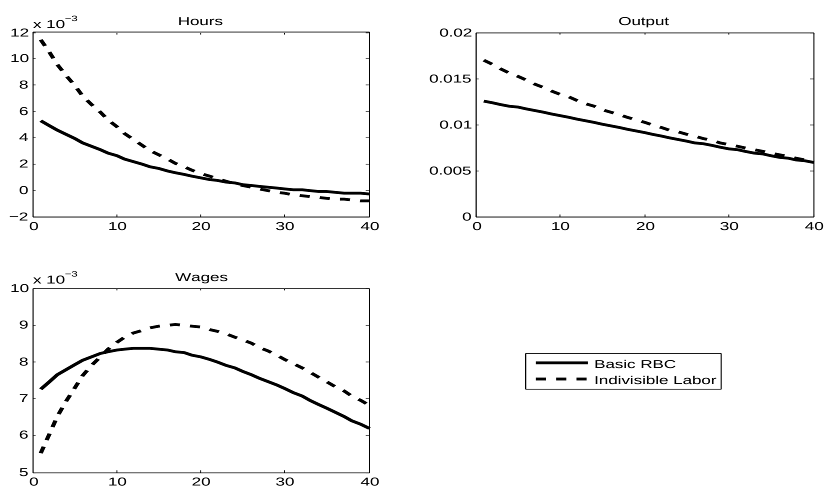

Hours and output respond by more to a technology shock on impact in the indivisible labor model relative to the basic RBC model, while wages respond by less on impact. Effectively this is because the labor supply curve is flatter (indeed, perfectly horizontal) in the indivisible labor case.

Quantitatively, indivisible labor improves the fit of the model along several dimensions. First, it provides greater amplification – I get output volatility of 2.2 percent with indivisible labor, as opposed to 1.7 percent in the standard case. This means that I can match the output volatility in US data with smaller TFP shocks. In addition, indivisible labor increases the relative volatility of hours substantially. In the benchmark RBC case, the relative volatility of hours is 0.43. In the indivisible labor case it is 0.68. This is a large improvement, though it is still quite far from the data. Furthermore, indivisible labor makes wages somewhat less volatile (volatility of 0.008 instead of 0.010) and somewhat less procyclical (correlation with output of 0.92 instead of 0.99).

2.2 Money

We have abstracted from money thus far. Isn't economics all about money? We will give money a functional definition – it is anything which is used as a medium of exchange, serves as a unit of account, and serves as a store of value. The existence of money eliminates the problems presented by the double coincidence of wants presented by a system of exchange based on barter. In the model we have presented thus far the presence of money is somewhat trivial, since there is only one good. But in a multi-good world (i.e. reality) money is obviously important.

That being said, it turns out to be fairly difficult to get agents to hold money. Agents will not willingly hold money in equilibrium for its store of value function – agents can also "save" through capital or bonds, which pay interest. What differentiates money is that it does not pay interest. Hence, the "cards are stacked" against money. It *must* be the exchange motive that gets people to hold money.

There are three ways in which to get agents to hold money, and we will consider two of them. The one we will not consider is the "money search" literature. This is a super micro-founded literature that considers money as rising endogenously in a search-theoretic framework. It is beyond the scope of this course. We will consider two "shortcuts" – cash in advance and money in the utility function. Both get at the exchange role of money. Cash in advance assumes that cash is *required* to purchase goods. This can be thought of as a "technological" constraint and is a reduced-form way of getting at the exchange role of money. The other approach is "money in the utility function". In this case we assume that agents get utility from holding money. This is also a reduced form way of getting at the idea that holding money makes conducting exchange "easier". These approaches yield similar results but they are not exactly the same. We consider each in turn.

2.2.1 The Budget Constraint

Before proceeding we need to write out a budget constraint that includes money. This is because money is a store of value. Let M_t denote the nominal holdings brought into period t – this is predetermined. Let M_{t+1} denote new money holdings (determined at time t) that will be brought into $t+1$. Let p_t denote the nominal price of goods – this is the price of goods measured in units of money. Let i_{t+1} denote the nominal interest rate on nominal bonds, B_t, observed at time t that pays off in time $t+1$. It pays off in dollars.

The household earns real income on work ($w_t n_t$), real income from leasing capital ($R_t k_t$), and nominal interest earned on bonds brought into this period ($i_t B_t$). We can convert this income from holding bonds into real terms by dividing by the price level, p_t. With this real income the household can (i) consume, c_t; (ii) purchase more capital, k_{t+1}, (iii) buy more (real) bonds, B_{t+1}/p_t, or)iv) accumulate more money, $(M_{t+1} - M_t)/p_t$:

$$c_t + k_{t+1} - (1-\delta)k_t + \left(\frac{B_{t+1} - B_t}{p_t}\right) + \left(\frac{M_{t+1} - M_t}{P_t}\right) = w_t n_t + R_t k_t + i_t \frac{B_t}{p_t} \tag{19}$$

In either of the following setups, the household can freely choose the real variables c_t and n_t. It can freely choose the bond and money holdings it carries over into the future, M_{t+1} and B_{t+1}. It takes all prices (w_t, i_t, and p_t as given). Firms are unaffected by money, since we can model their problem as completely static (they technically have the ability to operate in debt markets, but this ends up being indeterminate anyway so we can abstract from that part of the problem).

The firm problem is always standard:

$$\max_{k_t, n_t} \quad a_t f(k_t, n_t) - w_t n_t - R_t k_t$$

The first order conditions are:

$$w_t = a_t f_n(k_t, n_t)$$
$$R_t = a_t f_k(k_t, n_t)$$

There exists a central bank that sets the money supply in both set ups. Let's suppose that the exogenous process for the money supply follows an AR(1) in the growth rate (first difference of the log). This specification will generate positive trend inflation:

$$\ln M_{t+1} - \ln M_t = (1-\rho_m)\pi^* + \rho_m (\ln M_t - \ln M_{t-1}) + \varepsilon_{m,t} \tag{20}$$

Here π^* is the steady state growth rate of the money supply. This will end up being equal to steady state inflation in both models. Note that I am assuming that the Fed sets $\ln M_{t+1}$ at time t, since, under our timing convention, $\ln M_t$ is predetermined with respect to time t.

2.2.2 Money in the Utility Function

In this specification households get utility from consumption, leisure, and holding real money balances – M_{t+1}/p_t. Note the timing convention here – M_{t+1} is how much money the household chooses to hold today to carry into tomorrow. The idea here is that the more money one has (relative to the price level), the "easier" conducting transactions is. As before, we will go ahead and make functional form assumptions that permit a quantitative solution of the model. The household problem is:

$$\max_{c_t,n_t,k_{t+1},B_{t+1},M_{t+1}} \quad E_0 \sum_{t=0}^{\infty} \left\{ \ln c_t + \theta \ln(1-n_t) + \psi \frac{\left(\frac{M_{t+1}}{p_t}\right)^{1-\zeta} - 1}{1-\zeta} \right\}$$

s.t.

$$c_t + k_{t+1} - (1-\delta)k_t + \left(\frac{B_{t+1}-B_t}{p_t}\right) + \left(\frac{M_{t+1}-M_t}{p_t}\right) = w_t n_t + R_t k_t + i_t \frac{B_t}{p_t}$$

Form a current value Lagrangian:

$$\mathcal{L} = E_0 \sum_{t=0}^{\infty} \left\{ \ln c_t + \theta \ln(1-n_t) + \psi \frac{\left(\frac{M_{t+1}}{p_t}\right)^{1-\zeta} - 1}{1-\zeta} + \dots \right.$$

$$\left. \dots + \lambda_t \left(w_t n_t + R_t k_t + (1+i_t)\frac{B_t}{p_t} - c_t - k_{t+1} + (1-\delta)k_t - \frac{M_{t+1}}{p_t} + \frac{M_t}{p_t} - \frac{B_{t+1}}{p_t} \right) \right\}$$

The first order conditions are:

$$\frac{\partial \mathcal{L}}{\partial c_t} = 0 \Leftrightarrow \frac{1}{c_t} = \lambda_t \tag{21}$$

$$\frac{\partial \mathcal{L}}{\partial n_t} = 0 \Leftrightarrow \frac{\theta}{1-n_t} = \lambda_t w_t \tag{22}$$

$$\frac{\partial \mathcal{L}}{\partial k_{t+1}} = 0 \Leftrightarrow \lambda_t = \beta E_t \left(\lambda_{t+1}(R_{t+1} + (1-\delta))\right) \tag{23}$$

$$\frac{\partial \mathcal{L}}{\partial B_{t+1}} = 0 \Leftrightarrow \lambda_t = \beta E_t \left(\lambda_{t+1}(1+i_{t+1})\frac{p_t}{p_{t+1}}\right) \tag{24}$$

$$\frac{\partial \mathcal{L}}{\partial M_{t+1}} = 0 \Leftrightarrow \psi \left(\frac{M_{t+1}}{p_t}\right)^{-\zeta} \frac{1}{p_t} = \frac{\lambda_t}{p_t} - \beta E_t \frac{\lambda_{t+1}}{p_{t+1}} \tag{25}$$

The first four equations can be re-arranged to yield the *exactly* the same first order conditions which obtain in the standard RBC model:

$$\frac{\theta}{1-n_t} = \frac{1}{c_t} w_t \tag{26}$$

$$\frac{1}{c_t} = \beta E_t \left(\frac{1}{c_{t+1}}(R_{t+1} + (1-\delta))\right) \tag{27}$$

$$\frac{1}{c_t} = \beta E_t \left(\frac{1}{c_{t+1}}\left((1+i_{t+1})\frac{p_t}{p_{t+1}}\right)\right) \tag{28}$$

This is identical to the previous setup because of the Fisher relationship, which says that $1 + r_{t+1} = (1+i_{t+1})\frac{p_t}{p_{t+1}}$. We can greatly simplify the first order condition for holdings of money by using the first order condition for bonds:

$$\left(\frac{M_{t+1}}{p_t}\right)^{-\zeta} = \lambda_t - \beta E_t \lambda_{t+1} \frac{p_t}{p_{t+1}}$$

$$\beta E_t \lambda_{t+1} \frac{p_t}{p_{t+1}} = \frac{\lambda_t}{1 + i_{t+1}}$$

$$\Rightarrow$$

$$\left(\frac{M_{t+1}}{p_t}\right)^{-\zeta} = \lambda_t \left(1 - \frac{1}{1 + i_{t+1}}\right)$$

Define $m_t = \frac{M_{t+1}}{p_t}$ as "real money balances". This essentially says how much current consumption you are giving up by carrying money from today into tomorrow. Simplifying the above we have:

$$\psi m_t^{-\zeta} = \frac{1}{c_t} \left(\frac{i_{t+1}}{1 + i_{t+1}}\right)$$

This can be simplified further to yield:

$$m_t = \psi^{\zeta} c_t^{\zeta} \left(\frac{1 + i_{t+1}}{i_{t+1}}\right)^{\zeta} \tag{29}$$

This is quite intuitive. It says that the demand for money (i) increases one for one in the price level; (ii) is increasing in consumption; and (iii) is decreasing in the nominal interest rate. The nominal interest rate is the opportunity cost of holding money – if you didn't save in money, you could have saved in bonds, earning nominal interest. The fact that it is increasing in consumption essentially just says that money is a "normal" good in this setup – the wealthier you are, the more consumption you want and the more money you want to hold. The fact that the demand for money increases one for one with the price level gets at the idea that you get utility from real money balances, so an increase in the price level (which affects nothing else in the model) leads one to desire to hold more money.

To close the model out we need to deal with the non-stationarity inherent in the assumed process for the money supply. In particular, we want to write it in terms of real balances (which will be stationary). Hence, we need to play around with (20) by adding and subtracting logs of the price level at various leads and lags.

$$\ln M_{t+1} - \ln M_t = (1-\rho_m)\pi^* + \rho_m \left(\ln M_t - \ln M_{t-1}\right) + \varepsilon_{m,t}$$

$$\ln M_{t+1} - \ln p_t + \ln p_t - \ln p_{t-1} - \ln M_t + \ln p_{t-1} = (1-\rho_m)\pi^* \dots$$

$$\dots + \rho_m \left(\ln M_t - \ln p_{t-1} + \ln p_{t-1} - \ln p_{t-2} - \ln M_{t-1} + \ln p_{t-2}\right) + \varepsilon_{m,t}$$

We have $\ln m_t = \ln M_{t+1} - \ln p_t$, and define $\pi_t = \ln p_t - \ln p_{t-1}$. We can write this as:

$$\Delta \ln m_t + \pi_t = (1-\rho_m)\pi^* + \rho_m \pi_{t-1} + \rho_m \Delta \ln m_{t-1} + \varepsilon_{m,t} \tag{30}$$

The full set of conditions characterizing the model's equilibrium are then:

$$\frac{\theta}{1-n_t} = \frac{1}{c_t} w_t \tag{31}$$

$$\frac{1}{c_t} = \beta E_t \left(\frac{1}{c_{t+1}} (R_{t+1} + (1-\delta)) \right) \tag{32}$$

$$\frac{1}{c_t} = \beta E_t \left(\frac{1}{c_{t+1}} \left((1+i_{t+1}) \frac{p_t}{p_{t+1}} \right) \right) \tag{33}$$

$$R_t = \alpha a_t k_t^{\alpha-1} n_t^{1-\alpha} \tag{34}$$

$$w_t = (1-\alpha) a_t k_t^{\alpha} n_t^{-\alpha} \tag{35}$$

$$y_t = a_t k_t^{\alpha} n_t^{\alpha} \tag{36}$$

$$k_{t+1} = I_t + (1-\delta) k_t \tag{37}$$

$$y_t = c_t + I_t \tag{38}$$

$$\Delta \ln m_t = (1-\rho_m)\pi^* - \pi_t + \rho_m \pi_{t-1} + \rho_m \Delta \ln m_{t-1} + \varepsilon_{m,t} \tag{39}$$

$$m_t = \psi^{\zeta} c_t^{\zeta} \left(\frac{1+i_{t+1}}{i_{t+1}} \right)^{\zeta} \tag{40}$$

$$1 + r_{t+1} = (1+i_{t+1}) E_t \frac{p_t}{p_{t+1}} \tag{41}$$

$$\ln a_t = \rho \ln a_{t-1} + \varepsilon_t \tag{42}$$

$$\Delta \ln m_t = \ln m_t - \ln m_{t-1} \tag{43}$$

The equation determining the real variables of the model are *exactly* the same as in the basic RBC model. Intuitively, this means that the response of the real variables to a technology shock will be identical in this set up to earlier, and real variables will not respond to monetary shocks. Put differently, money is completely neutral with respect to real variables, and the classical dichotomy holds – real variables are determined first and then nominal variables are determined.

I parameterize the model exactly the same as before in the basic RBC notes. There are a few new parameters to be set, however. I set $\rho_m = 0.5$, and the standard deviation of the monetary policy shock to 0.01 (i.e. 1 percent). I set $\pi^* = 0.02$, so that there is two percent inflation in the steady state. I set $\psi = 1$ and $\zeta = 1$.

Below I show impulse response functions (just of the nominal variables, since the responses of the real variables to a technology shock are the same and their response to the monetary shock is zero) to both shocks. I construct the responses of the price level and the level of the nominal money supply using the facts that:

$$\ln p_t = \pi_t + \ln p_{t-1}$$

$$\ln M_t = \ln m_t + \ln p_t$$

Since the model is linearized, impulse responses don't depend on initial conditions, so I can normalize $\ln p_{t-1} = 0$ in constructing those responses.

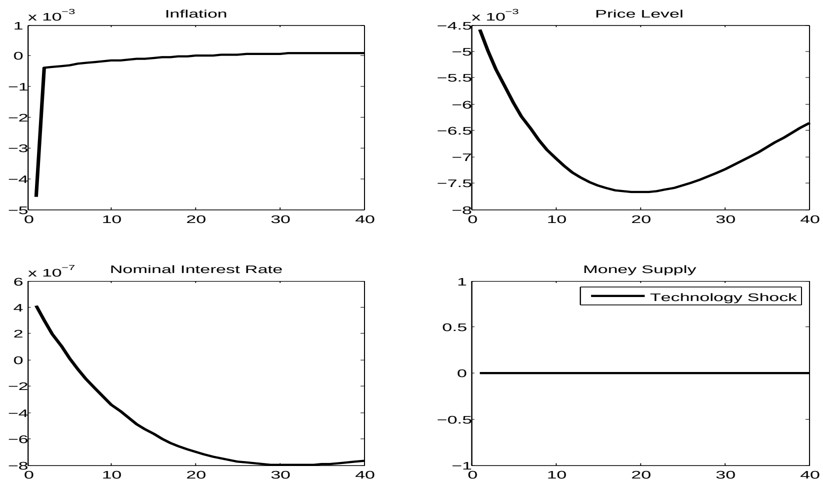

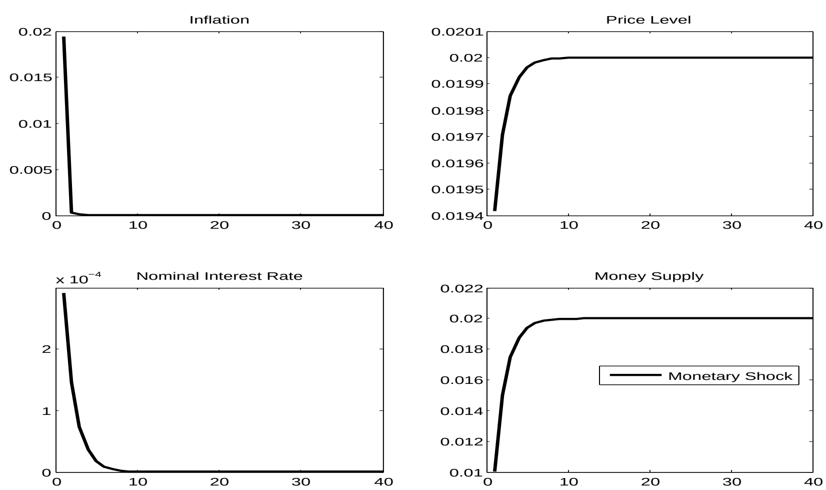

These have features we would more or less expect – inflation (and hence the price level) fall in response to a technology shock and rise in response to a monetary shock. The nominal interest rate rises when the money supply increases at an unexpectedly fast rate. Further, we can immediately tell that the price level will be countercyclical – this is because technology shocks (which raise output) lower inflation and the price level, inducing a negative correlation. That correlation is consistent with the data.

2.2.3 Cash in Advance

Now we undertake another assumption that allows us to get money into our basic RBC model. It ends up having fairly similar implications to the money in the utility model specification but it is not identical. In particular, in this framework money is not completely neutral and the classical dichotomy does not hold.

The cash in advance constraint says that one must have enough money on hand to finance all nominal purchases of consumption goods. In particular:

$$M_t \geq p_t c_t \tag{44}$$

Otherwise the problem is the standard real business cycle model, modified to have money entering the budget constraint as a store of value. We can write out the household problem as:

$$\max_{c_t, n_t, k_{t+1}, B_{t+1}, M_{t+1}} \quad E_0 \sum_{t=0}^{\infty} \{\ln c_t + \theta \ln(1 - n_t)\}$$

s.t.

$$c_t + k_{t+1} - (1-\delta)k_t + \left(\frac{B_{t+1} - B_t}{p_t}\right) + \left(\frac{M_{t+1} - M_t}{p_t}\right) = w_t n_t + R_t k_t + i_t \frac{B_t}{p_t}$$
$$M_t \geq p_t c_t$$

We can form a current value Lagrangian:

$$\mathcal{L} = E_0 \sum_{t=0}^{\infty} \beta^t \left\{ \ln c_t + \theta \ln(1-n_t) + \mu_t \left(\frac{M_t}{p_t} - c_t\right) \ldots \right.$$
$$\ldots + \lambda_t \left(w_t n_t + R_t k_t + (1+i_t)\frac{B_t}{p_t} - c_t - k_{t+1} + (1-\delta)k_t - \frac{B_{t+1}}{p_t} - \frac{M_{t+1}}{p_t} + \frac{M_t}{p_t} \right)$$

This is a similar setup to before, except now there is no money in the utility function and there is an extra constraint, with Lagrange multiplier given by μ_t. The first order conditions are:

$$\frac{\partial \mathcal{L}}{\partial c_t} = 0 \Leftrightarrow \frac{1}{c_t} = \lambda_t + \mu_t \tag{45}$$

$$\frac{\partial \mathcal{L}}{\partial n_t} = 0 \Leftrightarrow \frac{\theta}{1-n_t} = \lambda_t w_t \tag{46}$$

$$\frac{\partial \mathcal{L}}{\partial k_{t+1}} = 0 \Leftrightarrow \lambda_t = \beta E_t \left(\lambda_{t+1}(R_{t+1} + (1-\delta))\right) \tag{47}$$

$$\frac{\partial \mathcal{L}}{\partial B_{t+1}} = 0 \Leftrightarrow \lambda_t = \beta E_t \left((1+i_{t+1}) \left(\frac{p_t}{p_{t+1}}\right) \right) \tag{48}$$

$$\frac{\partial \mathcal{L}}{\partial M_{t+1}} = 0 \Leftrightarrow -\frac{\lambda_t}{p_t} + \beta E_t \frac{\mu_{t+1}}{p_{t+1}} + \beta E_t \frac{\lambda_{t+1}}{p_{t+1}} = 0 \tag{49}$$

The final first order condition can be simplified to yield:

$$\lambda_t = \beta E_t \left(\mu_{t+1} \frac{p_t}{p_{t+1}} + \lambda_{t+1} \frac{p_t}{p_{t+1}} \right) \tag{50}$$

The first four of these first order conditions are identical to the money in the utility function setup. Take a look at the last one. Suppose that one had enough money (i.e. M_t was sufficiently big relative to consumption) that the cash in advance constraint was never binding. This would mean that $E_t \mu_{t+1} = 0$ for all time. Plugging this in to the last first order condition and comparing with the previous first order condition, we see that the only way the two could simultaneously hold is if $i_{t+1} = 0$. This makes sense – as long as bonds pay non-zero interest, one would *never* want to hold money. Hence, we would be at a corner solution. In the presence of the cash in advance constraint, however, this will not be true.

I solve the model in Dynare assuming that the cash in advance constraint always binds. The firm problem, money growth rule, and stochastic process for technology are identical to above. It is helpful to *not* eliminate the Lagrange multipliers when solving this problem – that's fine, as it just introduces more variables. I also have to solve the model using inflation and real money balances, as the nominal money supply and price level are non-stationary.

The impulse responses to a technology shock are *exactly* the same in the cash in advance model as in the money in the utility function model, which are, in turn, exactly the same as in the basic RBC model. It is in this sense that abstracting from money altogether in that model is fine. It turns out here, however, that money does have real effects, although these are small. The impulse responses of the real variables to a monetary policy shock are shown below:

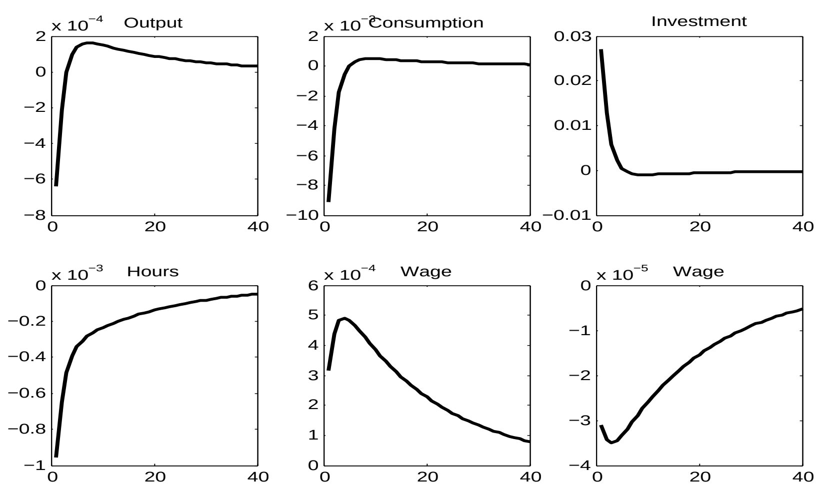

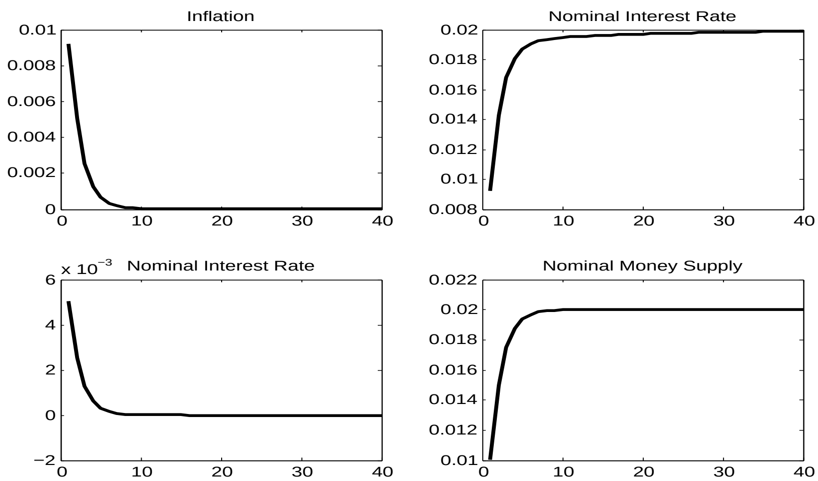

Here we see something perhaps not very intuitive. Not only does the monetary shock have real effects, it actually causes an output contraction (albeit it is very small). What is the intuition for this? Inflation is essentially a tax on holders on money. In the absence of the technological constraint requiring them to hold money (the cash in advance constraint), people thus wouldn't hold it at all. But given that they do have to hold it, an increase in the rate of growth of the money supply – which causes more inflation – makes people want to "get out of" money. Since consumption requires money, they can't substitute from money to consumption, so they substitute from money to leisure. Hence, there is a reduction in labor supply and a reduction in consumption, which leads to an output decline, real wage increase, and investment increase.

That being said, the real effects of money in this model are indeed *tiny*. In particular, for the parameterization I used, money explains less than 1 percent of the variance of output. Hence, money is *approximately* neutral in this model. To get large monetary non-neutrality, one needs to introduce other kinds of frictions (like price stickiness).

The so-called Friedman rule is to set the nominal interest rate on bonds equal to zero $i_{t+1} = 0$. It turns out that this is optimal (from the perspective of steady state welfare) in both the cash in advance and money in the utility function models. Friedman's original intuition was straightforward. Money is a "good" thing in the sense of reducing transactions frictions and therefore increases welfare. It is (essentially) costless to produce. The nominal interest rate being positive imposes a tax on the holders of money, which distorts welfare. Put differently, the private marginal cost of holding money is the nominal interest rate, while the public marginal cost of producing money is (essentially) zero. To bring about efficiency we need to bring these into equality by reducing the distortion. From the Fisher relationship, this requires setting $\pi^* = -\left(\frac{1}{\beta} - 1\right)$.

At a more formal level, we can see why this is optimal in both of these specifications. In the CIA model, having negative steady state inflation means that $m_t \to \infty$. This means that, eventually, the cash in advance constraint will not bind. If the cash in advance constraint does not bind, then we are back in the basic RBC case (just look at the FOC). The constraint not binding means $\mu_t = 0$

which requires $i_{t+1} = 0$. Intuitively, households must be (weakly) better off by having constraints not bind. In the MIU case, welfare is strictly increasing in m_t. Hence, having negative inflation, and hence driving $m_t \to \infty$, maximizes welfare.

2.3 Non-Separability in Preferences

In our basic specification we have assumed two kinds of separability in preferences – separability between leisure and consumption (intratemporal separability) and separability of both leisure and consumption across time (intertemporal separability). We consider both of these in turn.

2.3.1 Intratemporal Non-Separability: King, Plosser, and Rebelo (1988)

The generic definition of balanced growth path is a situation in which all variables growth at a constant rate over time (though this rate need not be the same across variables). A special case of a balanced growth path is a steady state, in which the growth rate of all variables is equal to zero. In our benchmark specification above there is no explicit trend growth, though we could fairly easily modify the model in such a way that we get (essentially) the same first order conditions in the redefined variables which are detrended.

In any balanced growth path, feasibility requires that hours not grow. The intuition for this is straightforward – if hours were declining, we would eventually hit zero and have no output. If hours were growing, we would eventually hit 1, which is the upper bound on hours. It is straightforward to show, under the assumptions about technology and production we have made, consumption and the real wage must grow at the same rate along the balanced growth path, irrespective of the kinds of preferences. Consider a generic, possibly non-separable within-period utility function: $u(c_t, 1 - n_t)$. The only assumptions are that it is increasing and concave in its arguments. The generic static labor supply condition is as follows;

$$-u_n(c_t, 1 - n_t) = u_c(c_t, 1 - n_t)w_t$$

This is really just an MRS = price ratio condition between consumption and leisure. To satisfy the conditions laid out above (namely that consumption and the wage grow at the same rate and hours not grow), it must be the case that this first order condition reduce to something like:

$$f(n_t) = \frac{w_t}{c_t}$$

In other words, the left hand side must be a function only of n_t, and the right hand side must feature the wage over consumption. With wages and consumption growing at the same right, the right hand side will be constant along a balanced growth path. Then with the left hand side only a function of n_t (and, of course, parameters), there will be a unique solution for n^* that is not growing.

King, Plosser, and Rebelo (1988) show that preferences must take the following form for this to be true:

$$u(c_t, n_t) = \frac{(c_t v(1-n_t))^{1-\sigma} - 1}{1-\sigma} \quad \text{if} \quad \sigma \neq 1$$

$$u(c_t, n_t) = \ln c_t + \ln v(1-n_t) \quad \text{if} \quad \sigma = 1$$

The second step follows from application of L'Hopital's rule. We require that $v(1-n_t)$ be an increasing and concave function. So as to make this all consistent with our original specification, suppose that $v(\cdot)$ takes the following form:

$$v(1-n_t) = \exp\left(\theta \frac{(1-n_t)^{1-\xi} - 1}{1-\xi}\right)$$

With this specification, if $\sigma = 1$, then we get:

$$u(c_t, 1-n_t) = \ln c_t + \theta \frac{(1-n_t)^{1-\xi} - 1}{1-\xi}$$

Then, if $\xi = 1$, by L'Hopital's rule we would get: $u(c_t, 1-n_t) = \ln c_t + \theta \ln(1-n_t)$. If $\xi = 0$, we would get $u(c_t, 1-n_t) = \ln c_t + \theta(1-n_t)$, which is essentially the indivisible labor model. Thus, we can nest all of these specifications in terms of this general functional form.

For the general case in which $\sigma \neq 1$ and $\xi \neq 1$, we can verify that these preferences will be consistent with constant labor hours in steady state. Let's find the marginal utilities:

$$u_c(c_t, 1-n_t) = \left(c_t \exp\left(\theta \frac{(1-n_t)^{1-\xi} - 1}{1-\xi}\right)\right)^{-\sigma} \exp\left(\theta \frac{(1-n_t)^{1-\xi} - 1}{1-\xi}\right)$$

$$u_n(c_t, 1-n_t) = -\left(c_t \exp\left(\theta \frac{(1-n_t)^{1-\xi} - 1}{1-\xi}\right)\right)^{-\sigma} c_t \exp\left(\theta \frac{(1-n_t)^{1-\xi} - 1}{1-\xi}\right)(1-n_t)^{-\xi}$$

Then for the generic first order condition, we get:

$$-\frac{u_n(c_t, 1-n_t)}{u_c(c_t, 1-n_t)} = \frac{\theta(1-n_t)^{-\xi}}{c_t} = w_t \Rightarrow \theta(1-n_t)^{-\xi} = \frac{w_t}{c_t}$$

In other words, the static first order condition for labor supply ends up looking *exactly* like it does in the case of log consumption with these preferences. Hours will be stationary. θ and ξ will have exactly the same interpretations as in the basic model (θ will determine n^* and ξ will determine the Frisch elasticity).

What does σ govern? It is still going to have the interpretation as the elasticity of intertemporal substitution. The first order conditions for the household side of the model for consumption and bonds can be written:

$$\lambda_t = c_t^{-\sigma}\left(\exp\left(\theta \frac{(1-n_t)^{1-\xi} - 1}{1-\xi}\right)\right)^{1-\sigma} \tag{51}$$

$$\lambda_t = \beta E_t(\lambda_{t+1}(1 + r_{t+1})) \tag{52}$$

Log-linearize (52):

$$\ln \lambda_t = \ln \beta + \ln \lambda_{t+1} + \ln(1 + r_{t+1})$$
$$\frac{\lambda_t - \lambda^*}{\lambda^*} = \frac{\lambda_{t+1} - \lambda^*}{\lambda^*} + \frac{r_{t+1} - r^*}{1 + r^*}$$
$$\widetilde{\lambda}_t = E_t \widetilde{\lambda}_{t+1} + \beta \widetilde{r}_{t+1}$$

The last line follows from the fact that we define $\widetilde{r}_{t+1}$ as the actual deviation from steady state, not percentage deviation. Now log-linearize (51):

$$\ln \lambda_t = -\sigma \ln c_t + (1 - \sigma)\theta \left(\frac{(1 - n_t)^{1-\xi} - 1}{1 - \xi} \right)$$
$$\frac{\lambda_t - \lambda^*}{\lambda^*} = -\sigma \frac{c_t - c^*}{c^*} + (1 - \sigma)\theta(1 - n^*)^{-\xi}(n_t - n^*)$$
$$\widetilde{\lambda}_t = -\sigma \widetilde{c}_t + (1 - \sigma)\theta(1 - n^*)^{-\xi} n^* \widetilde{n}_t$$

Now combine these two expressions:

$$-\sigma \widetilde{c}_t + (1 - \sigma)\theta(1 - n^*)^{-\xi} n^* \widetilde{n}_t = E_t \left(-\sigma \widetilde{c}_{t+1} + (1 - \sigma)\theta(1 - n^*)^{-\xi} n^* \widetilde{n}_{t+1} \right) + \beta \widetilde{r}_{t+1}$$

Simplify:

$$E_t \left(\widetilde{c}_{t+1} - \widetilde{c}_t \right) = \frac{\beta}{\sigma} \widetilde{r}_{t+1} + \left(\frac{(1 - \sigma)}{\sigma} \theta(1 - n^*)^{-\xi} n^* \right) \left(E_t \left(\widetilde{n}_{t+1} - \widetilde{n}_t \right) \right) \quad (53)$$

If we approximate $\beta \approx 1$, then this says that the elasticity of intertemporal substitutions is $\frac{1}{\sigma}$, just like in the case with separable utility. There is just an additional term now that depends on expected employment growth, though if $\sigma = 1$ this term drops out and we are in the normal case.

What this specification of preferences thus does is allow us to consider different parameterizations of σ different from one while still having preferences that are consistent with balanced growth. Loosely speaking, σ governs the household's desire to smooth consumption. If σ is very large, the household will want consumption (in expectation) to be very smooth, whereas if σ is quite small then the household will be quite willing to allow consumption to not be smooth (again in expectation).

Below are impulse responses to a standard technology shock for different values of σ. I fix all other parameter values at their "baseline" values. I consider the following values of σ: 0.75, 1, 2, and 4.

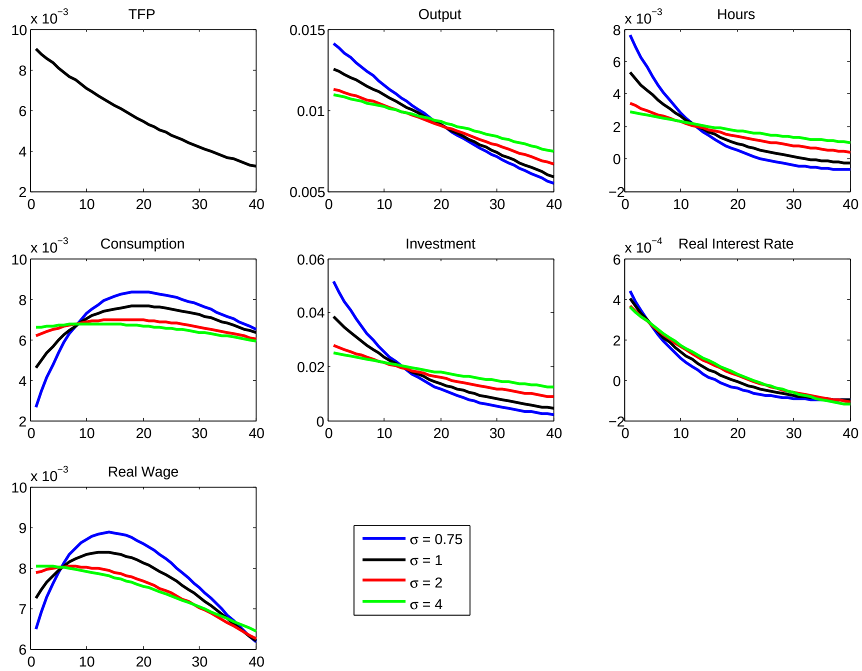

As we might expect, the initial jump in consumption is increasing in σ (note again that large σ means you want consumption to be smooth *in expectation*, not necessarily in response to a shock). This means that the jump in labor is decreasing in σ. Why is that? Think back the labor supply and demand curves:

$$\theta(1-n_t)^{-\xi} = \frac{1}{c_t} w_t$$
$$w_t = a_t(1-\alpha)k_t^{\alpha} n_t^{-\alpha}$$

When TFP increases, labor demand shifts right, the amount by which is independent of σ. When consumption increases, labor supply shifts left. The bigger is the consumption increase, the bigger is this inward shift in labor supply, and therefore the smaller is the hours response in equilibrium and the larger is the wage response. That's exactly what we see in terms of the impulse responses: when σ is bigger, the hours jump is smaller, the output jump is smaller, and the wage jump is larger.

This all suggests that one way to make the model better fit the data is to make σ smaller. In particular, we get more employment volatility and hence more amplification for $\sigma < 1$. The problem with this is that most micro evidence does not support such a claim – there estimates of σ are typically *far* greater than one. In particular, Hall (1988) says "... supporting the strong conclusion that the elasticity (of intertemporal substitution, the inverse of σ) is unlikely to be much above 0.1, and may well be zero." This would mean that $\sigma \geq 10$! A number of papers in the asset pricing literature rely upon very large values of σ in order to generate the excess returns on equity over debt that we see in the data. If we take values of σ much greater than 1, the RBC model begins to fit the data even worse than in the log case (in terms of amplification and relative volatility of hours). It is worth mentioning, however, that most of these estimates that find very large values of σ are based on time series data. Gruber (2006) finds a much smaller value of σ (more like 0.5) using micro data from looking at tax variation.

2.3.2 Intertemporal Non-Separability: Habit Formation

Another important kind of non-separability is non-separability across time. This usually goes by the name "habit formation", with the idea that people get utility not from the level of consumption, but from the level of consumption relative to past consumption. The idea is that one becomes accustomed to a certain level of consumption (i.e. a "habit") and utility becomes relative to that. Habit formation has been included in macro models for a variety of reasons. In particular, habit formation can help resolve some empirical failings of the PIH. For example, habit formation can help resolve the "excess smoothness" puzzle because, the bigger is habit formation, the smaller consumption will jump in response to news about permanent income. Another area where habit formation has gained ground is in asset pricing, in particular with regard to the equity premium puzzle. A large degree of habit formation, in essence, makes consumers behave "as if" they are extremely risk averse, and can thereby help explain a large equity premium without necessarily resorting to extremely large coefficients of relative risk aversion (see the previous subsection).

Assume intratemporal separability so that utility from consumption is logarithmic. Let the within period utility function be given by:

$$u(c_t, 1-n_t) = \ln(c_t - \phi c_{t-1}) + \theta \ln(1-n_t)$$

ϕ is the habit persistence parameter; if $\phi = 0$ we are in the "normal" case, and as $\phi \to 1$ agents get utility not from the level of consumption, but from the change in consumption. For computational purposes we need to restrict $\phi < 1$ – if it is exactly 1 then marginal utility in the steady state would be ∞.

Let's setup the household's problem using a Lagrangian. Assume that households own the capital stock:

$$\mathcal{L} = E_0 \sum_{t=0}^{\infty} \beta^t \left(\ln(c_t - \phi c_{t-1}) + \theta \ln(1-n_t) + \dots \right.$$
$$\left. \dots \lambda_t \left(w_t n_t + R_t k_t + \Pi_t + (1+r_t) b_t - c_t - k_{t+1} + (1-\delta) k_t - b_{t+1} \right) \right)$$

The first order conditions are:

$$\frac{\partial \mathcal{L}}{\partial c_t} = 0 \Leftrightarrow \lambda_t = \frac{1}{c_t - \phi c_{t-1}} - \beta \phi E_t \frac{1}{c_{t+1} - \phi c_t} \tag{54}$$

$$\frac{\partial \mathcal{L}}{\partial n_t} = 0 \Leftrightarrow \frac{\theta}{1-n_t} = \lambda_t w_t \tag{55}$$

$$\frac{\partial \mathcal{L}}{\partial k_{t+1}} = 0 \Leftrightarrow \lambda_t = \beta E_t(\lambda_{t+1}(R_{t+1} + (1-\delta))) \tag{56}$$

$$\frac{\partial \mathcal{L}}{\partial b_{t+1}} = 0 \Leftrightarrow \lambda_t = \beta E_t(\lambda_{t+1}(1 + r_{t+1})) \tag{57}$$

The only first order condition that is different is the one that defines λ_t; if $\phi = 0$ we are back in the usual case. It is easiest to solve this model by *not* substituting out for the Lagrange multiplier – just treat it as another endogenous variable. Below are impulse responses – using our otherwise standard calibration of a RBC model – for different values of ϕ:

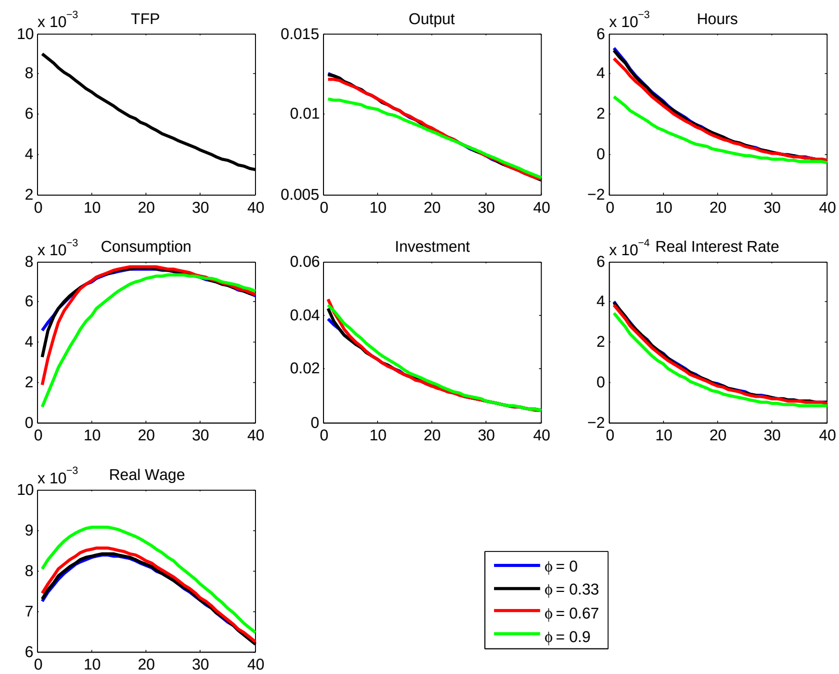

We observe that the main difference is that consumption jumps up by less on impact the bigger is ϕ. The intuition for this is high consumption today lowers utility tomorrow, other things being equal, the bigger is ϕ. Hence people will behave "cautiously" in essence by not adjusting consumption by much. The other impulse responses are reasonably similar across parameterizations, though for $\phi = 0.9$ we see that hours don't jump up by much, the real wage jumps up by a lot, and output doesn't jump up by much (i.e. this is not going to improve the fit of the model along those dimensions). The main dimension along which the inclusion of habit formation does help the model match the data is not in terms of unconditional moments, but rather in terms of conditional impulse response functions. Most estimated impulse responses to identified shocks (say, monetary policy shocks) show "hump-shaped" responses of consumption. This is difficult to generate without habit formation.

Another form of habit formation is sometimes what is called "external habit formation" or "Catching Up with the Joneses" (Abel, 1990). Here the idea is that utility from consumption depends not on consumption relative to own lagged consumption, but rather on consumption relative to lagged *aggregate* consumption – the idea being that you care about your consumption relative to that of your neighbor. Now, of course, in a representative agent framework own and aggregate end up being the same. The difference is that external habit formation simplifies the problem, because the consumer does not take into account the effect of current consumption decisions on the habit stock (essentially the second term in the expression for λ above drops out.

2.4 Non-Stationary Technology

In the specifications we have thus far looked at, we have (implicitly, most of the time) assumed that the non-stationary series of the model are *trend stationary*, because we assumed that labor augmenting technology followed a deterministic, linear time trend:

$$y_t = a_t k_t^{\alpha}(z_t n_t)^{1-\alpha}$$
$$z_t = (1+g_z)^t z_0$$

Let's instead suppose that technology follows a stochastic trend. We can get rid of z_t altogether and write the model as:

$$y_t = a_t k_t^{\alpha} n_t^{1-\alpha}$$
$$\ln a_t = g_z + \ln a_{t-1} + \varepsilon_t$$

Here we have assumed that $\ln a_t$ follows a random walk with drift, where the drift component, g_z, is the trend growth rate. We can equivalently write the process for this as:

$$\exp(\ln a_t) = \exp(g_z + \varepsilon_t)\exp(\ln a_{t-1})$$

The non-stochastic versions of this and the deterministic trend case are identical – in particular, the average growth rates will be the same. The economic content is very different, however. In the determinstic case shocks have temporary effects. In this case shocks have permanent effects. It also turns out that we need to detrend differently.

Let's figure out how to transform the variables of this model. Start with the production function, then take logs, and then first difference so as to get in growth rate form:

$$
\begin{aligned}
y_t &= a_t k_t^{\alpha} n_t^{1-\alpha} \\
\ln y_t &= \ln a_t + \alpha \ln k_t + (1-\alpha) \ln n_t \\
(\ln y_t - \ln y_{t-1}) &= (\ln a_t - \ln a_{t-1}) + \alpha(\ln k_t - \ln k_{t-1}) + (1-\alpha)(\ln n_t - \ln n_{t-1})
\end{aligned}
$$

Based on our discussion above, along a balanced growth path hours will not growth and capital will grow at the same rate as output. Using these facts, we have:

$$
\begin{aligned}
(\ln y_t - \ln y_{t-1}) &= (\ln a_t - \ln a_{t-1}) + \alpha(\ln y_t - \ln y_{t-1}) \\
\ln y_t - \ln y_{t-1} &= \frac{1}{1-\alpha}(\ln a_t - \ln a_{t-1})
\end{aligned}
$$

This says that, along a balanced growth path, output will grow at $\frac{1}{1-\alpha}$ the rate of technological progress (if we had written this as labor augmenting technological progress, as opposed to neutral, they would grow at the same rate . . . these two setups are equivalent provided we re-define the trend growth rate appropriately).

Play around with the above:

$$
\begin{aligned}
\ln\left(\frac{y_t}{y_{t-1}}\right) &= \ln\left(\frac{a_t}{a_{t-1}}\right)^{\frac{1}{1-\alpha}} \\
\frac{y_t}{y_{t-1}} &= \left(\frac{a_t}{a_{t-1}}\right)^{\frac{1}{1-\alpha}} \\
\frac{y_t}{a_t^{\frac{1}{1-\alpha}}} &= \frac{y_{t-1}}{a_{t-1}^{\frac{1}{1-\alpha}}}
\end{aligned}
$$

In other words, along the balanced growth path output divided by $a_t^{\frac{1}{1-\alpha}}$ does not grow – i.e. it is stationary. Hence, we can induce stationarity into the model by dividing through by this.

Define the following stationarity inducing transformations:

$$\widehat{y}_t \equiv \frac{y_t}{a_t^{\frac{1}{1-\alpha}}}$$
$$\widehat{k}_t \equiv \frac{k_t}{a_{t-1}^{\frac{1}{1-\alpha}}}$$
$$\widehat{I}_t \equiv \frac{I_t}{a_t^{\frac{1}{1-\alpha}}}$$
$$\widehat{w}_t \equiv \frac{w_t}{a_t^{\frac{1}{1-\alpha}}}$$
$$\widehat{c}_t \equiv \frac{c_t}{a_t^{\frac{1}{1-\alpha}}}$$

There is one very slight modification due to a timing assumption – we need to divide by k_t by a_{t-1}. Intuitively, this is because k_t is chosen at $t-1$, not t. We can use these transformations to alter the first order conditions of the basic model as needed. Begin with the production function, dividing both sides by the scaling factor $a_t^{\frac{1}{1-\alpha}}$:

$$\widehat{y}_t = a_t^{\frac{-\alpha}{1-\alpha}} k_t^{\alpha} n_t^{1-\alpha}$$

Now, multiply and divide by $a_{t-1}^{\frac{\alpha}{1-\alpha}}$ to get the capital stock in the correct terms:

$$\widehat{y}_t = a_t^{\frac{-\alpha}{1-\alpha}} \left(\frac{k_t}{a_{t-1}^{\frac{1}{1-\alpha}}} \right)^{\alpha} a_{t-1}^{\frac{\alpha}{1-\alpha}} n_t^{1-\alpha}$$
$$\widehat{y}_t = \left(\frac{a_t}{a_{t-1}} \right)^{\frac{-\alpha}{1-\alpha}} \widehat{k}_t^{\alpha} n_t^{1-\alpha}$$

Using the assume process for technology, we get:

$$\widehat{y}_t = \exp(g_z + \varepsilon_t)^{-\frac{\alpha}{1-\alpha}} \widehat{k}_t^{\alpha} n_t^{1-\alpha}$$

Next go to the capital accumulation equation, and divide both sides by the scaling factor at date t.

$$\frac{k_{t+1}}{a_t^{\frac{1}{1-\alpha}}} = \frac{I_t}{a_t^{\frac{1}{1-\alpha}}} + (1-\delta)\frac{k_t}{a_t^{\frac{1}{1-\alpha}}}$$

$$\widehat{k}_{t+1} = \widehat{I}_t + (1-\delta)\left(\frac{k_t}{a_{t-1}^{\frac{1}{1-\alpha}}}\right)\left(\frac{a_{t-1}}{a_t}\right)^{\frac{1}{1-\alpha}}$$

We can simplify this further by noting the process for technology:

$$\widehat{k}_{t+1} = \widehat{I}_t + (1-\delta)(\exp(g_z + \varepsilon_t)^{-1})\widehat{k}_t \tag{58}$$

The accounting identity is the same in terms of the transformed variables as always: $\widehat{y}_t = \widehat{c}_t + \widehat{I}_t$. If we assume log utility over consumption, then no adjustment is necessary to the discount factor. Otherwise a (small) downward adjustment needs to be made. Otherwise the remaining first order conditions are the same.

I solve the model using our standard parameter values. Dynare will produce impulse responses of the detrended variables. To construct impulse responses of the regular variables, I need to "add back" the level of technology: i.e. $\ln x_t = \ln \widehat{x}_t + \frac{1}{1-\alpha}\ln a_t$. The responses are below:

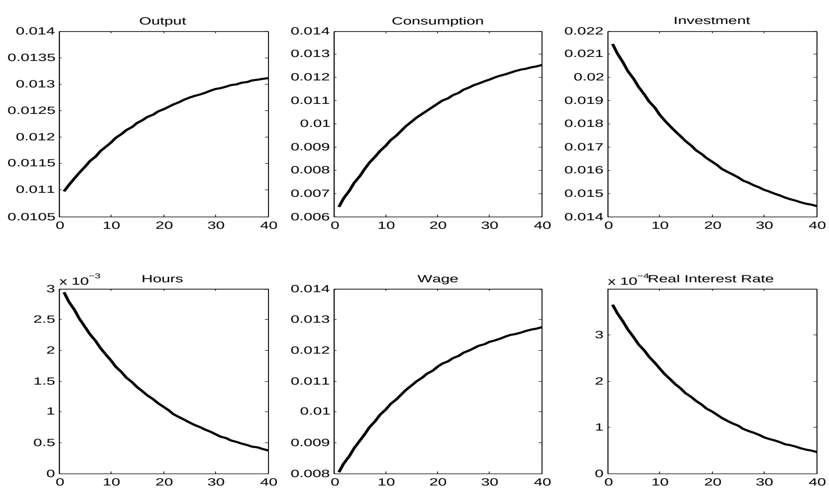

It turns out that random walk technology generate much less volatility in hours than do transitory technology shocks. The intuition for this is straightforward and is based on our permanent income intuition and our graphical analysis of the basic model. If the change in technology is permanent, consumption will jump up by more. Consumption jumping up by more means that labor supply shifts in by more. Hence, hours rise by less (and the real wage rises by more). Since hours rise by less, output rises by less. Combined with a bigger increase in consumption, investment would also rise by less on impact. Below is a graph comparing the response of hours two a technology shock in the transitory shock case ($\rho = 0.974$) and the random walk case.

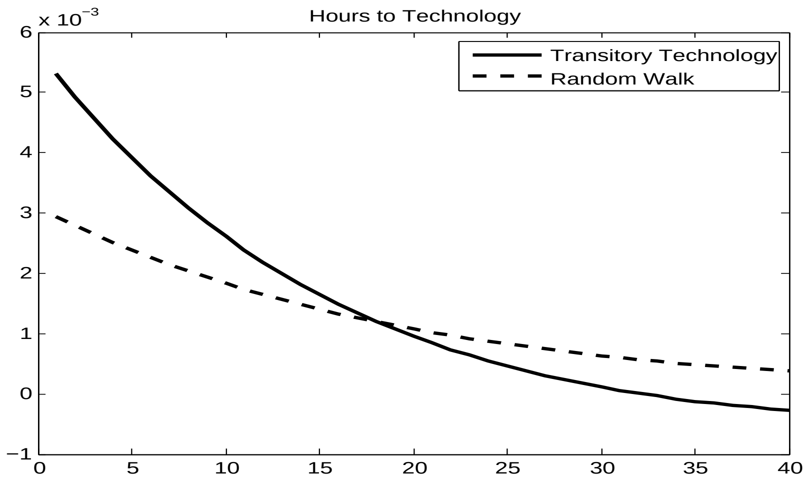

2.5 Variable Factor Utilization

As noted earlier, the basic RBC model has weak amplification mechanisms – the only thing that can make output go up by more than the driving technology shock is labor moving. The idea of variable factor utilization is to allow capital input to vary some in response to shocks. The basic idea is that while the amount of capital one enters the period with is predetermined, one can work those machines at greater or lesser intensity depending on current conditions. This is what we mean by utilization – firms can "utilize" their capital more or less intensely.

The household side of the problem is standard. Let's suppose that firms own the capital stock here. Denote the production function as follows:

$$y_t = a_t \left(u_t k_t\right)^{\alpha} n_t^{1-\alpha}$$

Here u_t denotes utilization. Capital "services" are the product of utilization and the amount of physical capital. There needs to be a cost associated with utilization for there to be a well-defined choice here. We model the cost of utilization as faster depreciation of capital. The basic idea is that if you work your capital harder, it depreciates faster. Model this cost as convex in utilization:

$$\delta_t = \delta_0 u_t^{\phi} \quad \phi > 1$$

The firm's problem can be written:

$$\max_{n_t, k_{t+1}, u_t, I_t} \qquad V_0 = E_0 \sum_{t=0}^{\infty} M_t \left(a_t \left(u_t k_t\right)^{\alpha} n_t^{1-\alpha} - w_t n_t - I_t\right)$$

s.t.

$$k_{t+1} = I_t + \left(1 - \delta_0 u_t^{\phi}\right) k_t$$

$M_t = \beta^t u'(c_t)$, the stochastic discount factor. We can impose that the constraint hold with equality and solve this as an unconstrained problem. The first order conditions for the firm's problem are:

$$\frac{\partial V_t}{\partial n_t} = 0 \Leftrightarrow (1-\alpha) a_t \left(u_t k_t\right)^{\alpha} n_t^{-\alpha} = w_t \tag{59}$$

$$\frac{\partial V_t}{\partial k_{t+1}} = 0 \Leftrightarrow u'(c_t) = \beta E_t u'(c_{t+1}) \left(\alpha a_{t+1} u_{t+1}^{\alpha} k_{t+1}^{\alpha-1} n_{t+1}^{1-\alpha} + \left(1 - \delta_0 u_{t+1}^{\phi}\right)\right) \tag{60}$$

$$\frac{\partial V_t}{\partial u_t} = 0 \Leftrightarrow \alpha a_t u_t^{\alpha-1} k_t^{\alpha} n_t^{1-\alpha} = \delta_0 \phi u_t^{\phi-1} k_t \tag{61}$$

The first order condition for utilization can equivalently be written as follows:

$$\alpha \frac{y_t}{u_t} = \delta_0 \phi u_t^{\phi-1} k_t \tag{62}$$

We can simplify this further to get:

$$u_t^{\phi} = \frac{\alpha y_t}{\phi \delta_0 k_t} \tag{63}$$

To solve for the steady state of the model we need to come up with a normalization of u^*. Let's set this to one. This normalization provides a parametric restriction on ϕ:

$$\phi = \frac{1}{\delta_0} \left(\alpha \frac{y^*}{k^*} \right) \tag{64}$$

This is convenient, because we know from the first order condition for the capital stock that, with $u^* = 1$, $\left(\alpha \frac{y^*}{k^*}\right)$ is equal to the steady state marginal product of capital. Hence, we can use that first order condition to simplify:

$$\phi = \frac{\frac{1}{\beta} - (1-\delta_0)}{\delta_0} = \frac{\frac{1}{\beta} - 1}{\delta_0} + 1 \tag{65}$$

Since, with $u^* = 1$, δ_0 is the average rate of depreciation, we can calibrate this to 0.025 to match the average share of investment in total output. The household first order conditions are the same as before, and we assume the same stochastic process for log technology as before. Calibrate the parameters of the model such that $\alpha = 0.33$, $\beta = 0.99$, $\rho = 0.974$, and $\sigma_{e,t} = 0.009$ as the standard deviation of the technology shock. These values imply that $\phi = 1.41$.

I solve the model in Dynare. For the purposes of comparison with the basic RBC model, the responses in this model are shown as solid lines, whereas the responses in the basic RBC model are dashed lines.

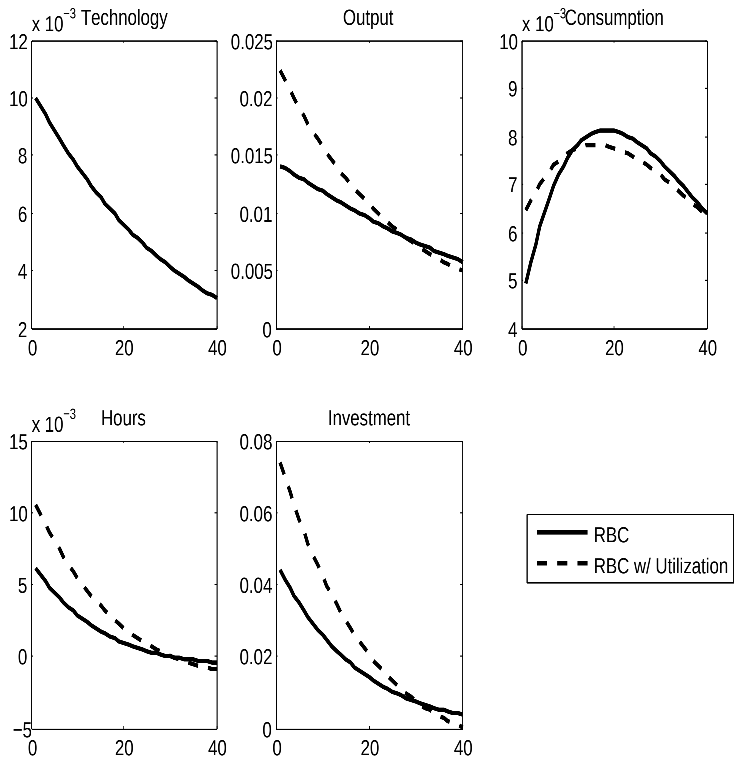

The first key thing to note is that capital utilization rises in response to the TFP shock. This makes sense – when TFP is high, firms would like more capital. They can't (in aggregate) get more capital without waiting a period, but utilizing their existing capital is almost as good, so that's what they do. This utilization provides a boost to output, which rises by more in this model than it does in the basic RBC model. In addition, we observed that investment and consumption both rise by more. Perhaps somewhat curiously, we see that hours actually rise by less here than in the model without utilization. The intuition is based on the wealth effect – since consumption jumps by more on impact, the inward shift of labor supply is bigger. Even though labor demand shifts out by more (because the marginal product of labor rises for two reasons – increased utilization and increased TFP) – the supply effect evidently dominates, and hours rise but by less than if there were no utilization. Consistent with that logic, we observe that the wage rises by more on impact. We also observe that the real interest rate rises by (significantly) less on impact here than in the case with no utilization of capital.

Utilization evidently serves as a pretty strong amplification mechanism. What this means is that we can get the same output volatility with smaller TFP shocks, which is appealing to lots of folks who wonder why we don't read about these large TFP shocks in the newspaper. Indeed, in this model I get output volatility of about 2 percent, as opposed to output volatility of 1.65 percent in the model without utilization. The presence of utilization is getting my that extra 0.35 percentage points of output volatility.

Note also that the presence of variable utilization means that measured TFP will not correspond to the "true" a_t of the model if one does not account for utilization.

2.6 Imperfect Competition

We now deviate from the assumption of perfect competition. Although it is not necessary, it is helpful to break production up into two sectors. The first is the "final goods" sector and is competitive, so we can think about there being a representative final goods firm. This firm doesn't use any factors of production, but rather "bundles" intermediate goods into a final good. The intermediate goods use capital and labor to produce. There are a continuum of intermediate goods firms who populate the unit interval. This is just a convenient normalization – the point is that there are a "lot" of intermediate good firms, but they produce differentiated goods.

The final good is a constant elasticity of substitution aggregate of intermediate goods. The "production" technology is:

$$y_t = \left(\int_0^1 y_{j,t}^{\frac{\nu-1}{\nu}} dj \right)^{\frac{\nu}{\nu-1}} \tag{66}$$

Remember that an integral is just the sum – this is the sum of each intermediate input raised to a power, with the whole sum raised to a power that is the inverse of the power on each intermediate input. ν is a parameter assumed to be positive and it governs the degree of substitutability among intermediate inputs. As it goes to infinity, this just becomes the sum of intermediate goods (i.e. goods are perfect substitutes). As it goes to zero, the production technology becomes Leontief (perfect complements). For $\nu = 1$, there is a "unit elasticity of substitution" and the production technology is Cobb-Douglas (the product of the intermediate inputs). Assume for what follows that $\nu > 1$.

The final goods firm wants to maximize (nominal) profits, given a final good price, p_t, and taking intermediate good prices, $p_{j,t}$, as given:

$$\max_{y_{j,t}} \quad \Pi_t^F = p_t \left(\int_0^1 y_{j,t}^{\frac{\nu-1}{\nu}} dj \right)^{\frac{\nu}{\nu-1}} - \int_0^1 p_{j,t} y_{j,t} dj$$

The first order conditions are found by differentiating with respect to $y_{i,t}$ and setting equal to zero:

$$\frac{\partial \Pi_t^F}{\partial y_{i,t}} = 0 \Leftrightarrow p_t \frac{\nu}{\nu-1} \left(\int_0^1 y_{j,t}^{\frac{\nu-1}{\nu}} dj \right)^{\frac{\nu}{\nu-1}-1} \frac{\nu-1}{\nu} y_{j,t}^{\frac{\nu-1}{\nu}-1} = p_{j,t}$$

$$p_t \left(\int_0^1 y_{j,t}^{\frac{\nu-1}{\nu}} dj \right)^{\frac{1}{\nu-1}} y_{j,t}^{\frac{-1}{\nu}} = p_{j,t}$$

$$y_{j,t}^{\frac{-1}{\nu}} = \left(\frac{p_{i,t}}{p_t} \right) \left(\int_0^1 y_{j,t}^{\frac{\nu-1}{\nu}} dj \right)^{\frac{-1}{\nu-1}}$$

$$y_{j,t} = \left(\frac{p_{i,t}}{p_t} \right)^{-\nu} \left(\int_0^1 y_{j,t}^{\frac{\nu-1}{\nu}} dj \right)^{\frac{\nu}{\nu-1}}$$

Using the definition of the aggregate final goods production technology, this reduces nicely to:

$$y_{j,t} = \left(\frac{p_{i,t}}{p_t}\right)^{-\nu} y_t \tag{67}$$

In words, the relative demand for differentiated intermediate good j depends on its relative price, with ν the price elasticity of demand.

We can now solve for the aggregate price index. The nominal value of the final good is just the sum of prices times quantities of intermediate goods, using the above demand specification:

$$p_t y_t = \int_0^1 p_{j,t} y_{j,t} dj = \int_0^1 p_{j,t} \left(\frac{p_{j,t}}{p_t}\right)^{-\nu} y_t dj$$
$$p_t y_t = \int_0^1 p_{j,t}^{1-\nu} p_t \nu y_t dj = p_t^{\nu} y_y \int_0^1 p_{j,t}^{1-\nu} dj$$

Simplifying, we get:

$$p_t = \left(\int_0^1 p_{j,t}^{1-\nu} dj\right)^{\frac{1}{1-\nu}} \tag{68}$$

The intermediate goods firms produce output using capital and labor, according to a standard production technology:

$$y_{j,t} = a_t k_{j,t}^{\alpha} n_{j,t}^{1-\alpha} \tag{69}$$

a_t is aggregate technology and is common across intermediate goods firms. It follows that aggregate capital and aggregate employment are just the sum of these factors across intermediate goods firms:

$$k_t = \int_0^1 k_{j,t} dj$$
$$n_t = \int_0^1 n_{j,t} dj$$

Assume that there are no debt instruments and that the intermediate goods firm rents capital from households. These firms all face the same factor prices (rental rate and wage rate). The firms do, however, have the ability to set their own price, given that they face downward sloping demand curves (as long as ν is not ∞). Hence, they want to solve the following constrained problem:

$$\max_{y_{j,t}, p_{j,t}, k_{j,t}, n_{j,t}} \quad p_{j,t} y_{j,t} - w_t n_{j,t} - R_t k_{j,t}$$

s.t.

$$y_{j,t} = a_t k_{j,t}^{\alpha} n_{j,t}^{1-\alpha}$$
$$y_{j,t} = \left(\frac{p_{i,t}}{p_t}\right)^{-\nu} y_t$$

Set the problem up using a Lagrangian, with two multipliers, $\lambda_{1,t}^j$ and $\lambda_{2,t}^j$.

$$\mathcal{L} = p_{j,t}y_{j,t} - w_t n_{j,t} - R_t k_{j,t} + \lambda_{1,t}^j \left(a_t k_{j,t}^{\alpha} n_{j,t}^{1-\alpha} - y_{j,t}\right) + \lambda_{2,t}^j \left(\left(\frac{p_{i,t}}{p_t}\right)^{-\nu} y_t - y_{j,t}\right)$$

The first order conditions are:

$$\frac{\partial \mathcal{L}}{\partial y_{j,t}} = 0 \Leftrightarrow p_{j,t} = \lambda_{1,t}^j + \lambda_{2,t}^j \tag{70}$$
$$\frac{\partial \mathcal{L}}{\partial p_{j,t}} = 0 \Leftrightarrow y_{j,t} = \nu \lambda_{2,t}^j p_{j,t}^{-\nu-1} p_t^{\nu} y_t \tag{71}$$
$$\frac{\partial \mathcal{L}}{\partial k_{j,t}} = 0 \Leftrightarrow R_t = \lambda_{1,t}^j \alpha a_t k_{j,t}^{\alpha-1} n_{j,t}^{1-\alpha} \tag{72}$$
$$\frac{\partial \mathcal{L}}{\partial n_{j,t}} = 0 \Leftrightarrow w_t = \lambda_{1,t}^j (1-\alpha) a_t k_{j,t}^{\alpha} n_{j,t}^{\alpha} \tag{73}$$

The first order condition for the price can be simplified:

$$y_{j,t} = \nu \lambda_{2,t}^j \left(\frac{p_{j,t}}{p_t}\right)^{-\nu} y_t p_{j,t}^{-1}$$
$$p_{j,t} = \nu \lambda_{2,t}^j$$
$$\lambda_{2,t}^j = \frac{p_{j,t}}{\nu}$$

Now plug this into the first order condition for output:

$$p_{j,t} = \lambda_{1,t}^j + \frac{p_{j,t}}{\nu}$$

Simplify:

$$p_{j,t} = \frac{\nu}{\nu - 1}\lambda_{1,t}^j \tag{74}$$

Now that $\frac{\nu}{\nu-1} \geq 1$. What is the interpretation of this statement? $\lambda_{1,t}^j$ can be interpreted as marginal cost. It is the shadow value on the first constraint: if I make you produce a little less, by how much do your profits go up (equivalently how much do your costs go down). Hence, this expression says that the optimal pricing rule is to set price equal to a "markup" of price over marginal cost, with the markup defined as $\varphi = \frac{\nu}{\nu-1}$. The less substitutable the intermediate goods are (i.e. the smaller is ν) the bigger the markup will be.

Plug this into the first order conditions for capital and labor; this will allow these conditions to be written in terms of the real product wage and the real product rental rate (the "product" qualifier means that we divide the nominal factor price by the price of the product, not the price level of all goods . . . this is the real factor price relevant for firm decision making):

$$\frac{w_t}{p_{j,t}} = \frac{\nu - 1}{\nu} \alpha a_t k_{j,t}^{\alpha} n_{j,t}^{-\alpha} \tag{75}$$

$$\frac{R_t}{p_{j,t}} = \frac{\nu - 1}{\nu} (1-\alpha) a_t k_{j,t}^{\alpha-1} n_{j,t}^{1-\alpha} \tag{76}$$

Because $\frac{\nu-1}{\nu} \leq 1$, factors will be paid less than their marginal products; this gives rise to economic profits for the intermediate goods firms.

Now use the first order conditions for labor and capital to eliminate $\lambda_{1,t}^{j}$:

$$\lambda_{1,t}^{j} = \frac{R_t}{\alpha a_t k_{j,t}^{\alpha-1} n_{j,t}^{1-\alpha}}$$

$$w_t = \frac{R_t}{\alpha a_t k_{j,t}^{\alpha-1} n_{j,t}^{1-\alpha}} (1-\alpha) a_t k_{j,t}^{\alpha} n_{j,t}^{-\alpha}$$

$$w_t = R_t \frac{1-\alpha}{\alpha} \frac{k_{j,t}}{n_{j,t}}$$

$$\frac{k_{j,t}}{n_{j,t}} = \frac{\alpha}{1-\alpha} \frac{w_t}{R_t}$$

This last condition is important. It says that all forms will hire capital and labor in the same ratio, since the wage, the rental rate, and α are common to all firms. Use this fact to go back to the expression for $\lambda_{1,t}^{j}$, which again has the interpretation as marginal cost:

$$\lambda_{1,t}^{j} = \frac{R_t}{\alpha a_t \left(\frac{k_{j,t}}{n_{j,t}} \right)^{\alpha-1}}$$

Since all firms will hire capital and labor in the same ratio, this means that they all have the same marginal cost. But going back to the pricing rule, if they all have the same marginal cost, then they all will charge the same price. Then using the formula for the aggregate price level, we see:

$$p_{j,t} = p_t \quad \forall \ j \tag{77}$$

In other words, all firms charge the same price, which is equal to the final goods price. From the demand specification, if all firms charge the same price, they must produce the same amount of output:

$$y_{j,t} = y_t \quad \forall \ j \tag{78}$$

This may seem a little odd, but this is the advantage of defining firms as existing over the unit interval – the output of any one firm is equal to the aggregate output which is equal to average output. The individual production function is:

$$y_{j,t} = a_t \left(\frac{k_{j,t}}{n_{j,t}} \right)^{\alpha} n_{j,t}$$

Since all firms hire capital and labor in the same ratio, and also produce the same amount of output, we can see that they must all hire the same amount of labor, and hence the same amount of capital:

$$\begin{aligned} k_{j,t} &= k_t \quad \forall \ j \\ n_{j,t} &= n_t \quad \forall \ j \end{aligned}$$

This means that we can think of there being an aggregate production function (for the final good) that is identical to the production function of any intermediate good firm:

$$y_t = a_t k_t^{\alpha} n_t^{1-\alpha} \tag{79}$$

Because all firms charge the same price, the relative price of all goods comes out to be 1 in equilibrium. The level of prices is indeterminate without specifying some process for money (i.e. we could easily do that). Hence, we can normalize all prices to be one; this means that there is no difference between real and nominal factor prices. The factor demand equations become:

$$w_t = \frac{\nu - 1}{\nu}(1-\alpha) a_t k_t^{\alpha} n_t^{-\alpha} \tag{80}$$

$$R_t = \frac{\nu - 1}{\nu} \alpha a_t k_t^{\alpha-1} n_t^{1-\alpha} \tag{81}$$

The household side of the model is the same as in our benchmark case, and has the same first order conditions. The entire set of first order conditions characterizing the equilibrium of this model are:

$$\frac{1}{c_t} = \beta E_t \left(\frac{1}{c_{t+1}} (R_{t+1} + (1-\delta) \right) \tag{82}$$

$$\frac{1}{c_t} = \beta E_t \left(\frac{1}{c_{t+1}} (1 + r_{t+1}) \right) \tag{83}$$

$$\frac{\theta}{1-n_t} = \frac{1}{c_t} w_t \tag{84}$$

$$w_t = \frac{\nu - 1}{\nu} (1-\alpha) a_t k_t^{\alpha} n_t^{-\alpha} \tag{85}$$

$$R_t = \frac{\nu - 1}{\nu} \alpha a_t k_t^{\alpha - 1} n_t^{1-\alpha} \tag{86}$$

$$y_t = c_t + I_t \tag{87}$$

$$y_t = a_t k_t^{\alpha} n_t^{1-\alpha} \tag{88}$$

$$k_{t+1} = I_t + (1-\delta) k_t \tag{89}$$

$$\ln a_t = \rho \ln a_{t-1} + \varepsilon_t \tag{90}$$

We can see that these are exactly the same first order conditions which obtain in the basic RBC model, with the exception of the inverse of the price markup in the factor demand equations. If we assume that ν is constant, then the only thing that will be different about this model is the steady state – in particular, $\nu < \infty$ will distort the steady state values. In a linearization of the model, the impulse responses will be identical. Essentially the imperfect competition is a steady state distortion; to a first order approximation it does not impact the dynamics of the model. A welfare optimizing government would want to levy taxes to restore the first best equilibrium.

We can, however, entertain fluctuations in ν. We can effectively think of these as being markup shocks. As above, define $\varphi = \frac{\nu}{\nu - 1}$. Suppose that the log of this follows a stationary AR(1):

$$\ln \varphi_t = (1 - \rho_\varphi) \varphi^* + \rho_\varphi \ln \varphi_{t-1} + \varepsilon_{\varphi,t}$$

I can re-write the factor demand equations as:

$$w_t = \frac{1}{\varphi_t} (1-\alpha) a_t k_t^{\alpha} n_t^{-\alpha} \tag{91}$$

$$R_t = \frac{1}{\varphi_t} \alpha a_t k_t^{\alpha - 1} n_t^{1-\alpha} \tag{92}$$

I'm just going to first create parameterization of the process for φ out of thin air: let's assume that $\varphi^* = 1.2$ (equivalently $\nu^* = 5$, that $\rho_\varphi = 0.9$, and that the standard deviation of the innovation is 0.01 (i.e. 1 percent). Below are impulse responses to a markup shock in the model with our benchmark parameterization (note that, to match the steady state hours of $\frac{1}{3}$, I would need to adjust θ to reflect φ^* . . . I don't do that here, as it doesn't affect the dynamics).

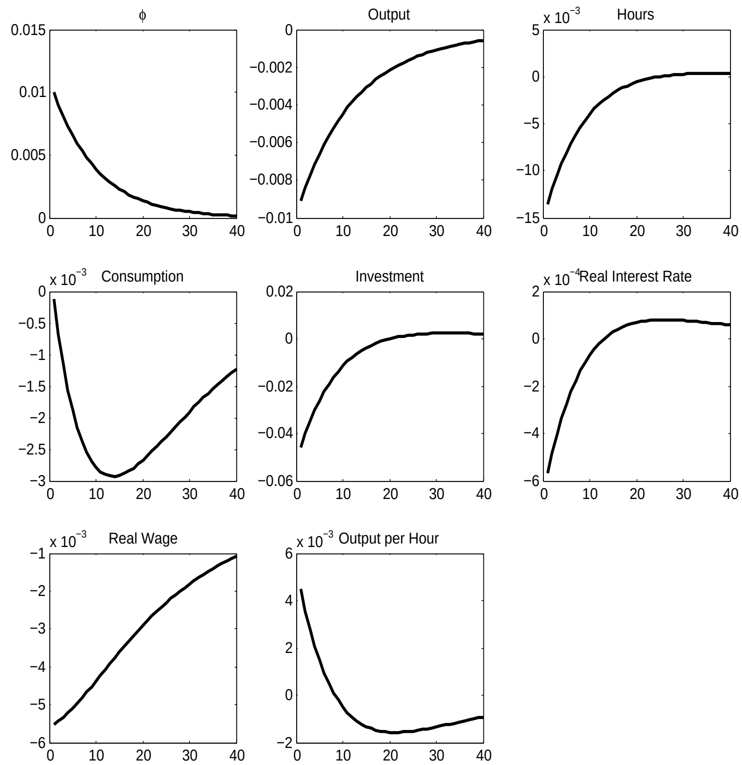

There are a couple of things evident from these responses. First, the markup shock causes consumption, hours, output, and investment to all decline together (i.e. positive co-movement). But second, average labor productivity goes up, which means that hours are falling by more than output. Adding markup shocks is thus going to do two things for us that the basic RBC model struggles with – it will reduce the cyclicality of labor productivity (which is way too high in the basic model relative to the data), and it will increase the volatility of hours relative to output (which is way too low in the model relative to the data).

Now let's think about a serious calibration of the process for the markup, φ_t. How can I measure it? Well, if I combine first order conditions from above, equilibrium in the labor market requires:

$$\frac{\theta}{1-n_t} = \frac{1}{c_t}\frac{1}{\varphi_t}(1-\alpha)\frac{y_t}{n_t} \tag{93}$$

Take logs of this and simplify, isolating φ_t:

$$\ln\varphi_t = -\ln\theta + \ln(1-n_t) - \ln c_t + \ln(1-\alpha) + \ln y_t - \ln n_t \tag{94}$$

In principle I observe everything on the right hand side (minus the parameters, which I have calibrated). I get data on per capita consumption, output, and hours. The only thing I need to be careful about is that I need to rescale the resulting per capita hours series so that its mean in the level is 0.33, in accord with our calibration. I do not need to rescale output or consumption – since these are going to get the same scale, the re-scaling would cancel out (the signs are the opposite above). Then I can back out a measure for the markup from above – essentially as a residual from the static first order condition for labor supply. Here is a picture of what I get:

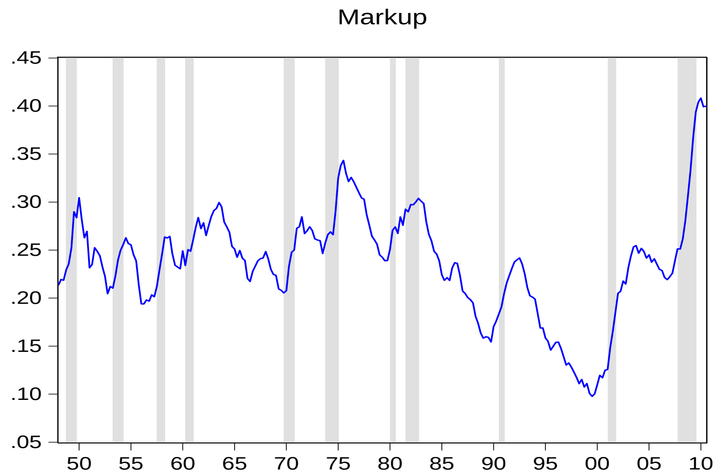

Remember that this is a log-scale graph, we see that the mean of φ is approximate $\varphi^* = 1.26$. This implies a mean value of $\nu^* = 4.75$. We can see from the graph that the markup appears to be *very* counterylical. Next I estimate an AR(1) on my empirical measure of the markup:

$$\ln \widehat{\varphi}_t = \beta_0 + \rho_\varphi \ln \widehat{\varphi}_{t-1} + e_t \tag{95}$$

I get $\rho_\varphi = 0.99$ and the standard error of the regression of 0.011. Here are impulse responses of to the markup shock under this paramterization (as opposed to the largely ad-hoc one above):

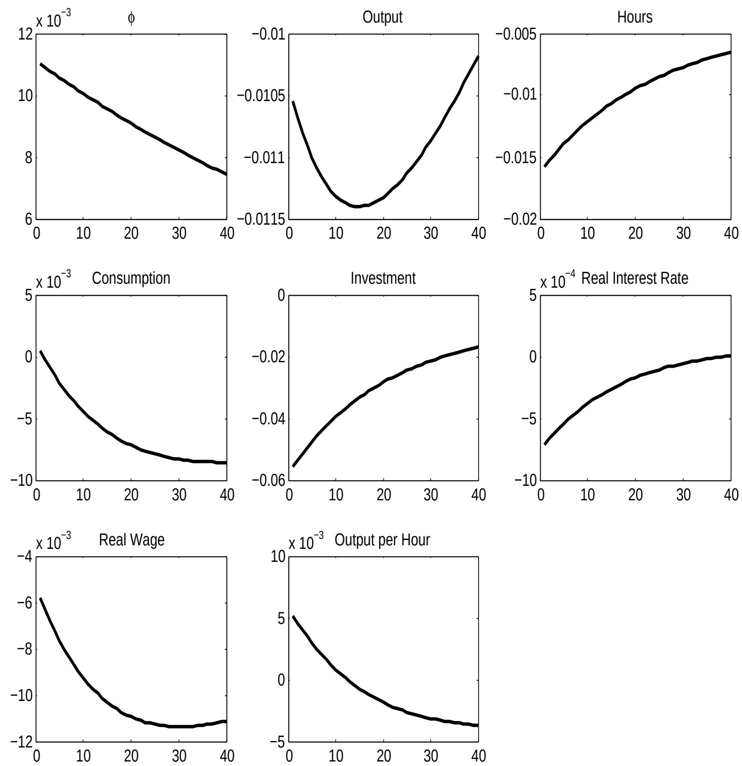

Note that consumption actually rises slightly on impact. The intuition for this is as follows. The markup essentially works like a tax on labor – the bigger the markup, the less of the marginal product workers get to keep. Thus an increase in the markup discourages work by lowering the take home wage. The substitution effect is to consume more stuff; the income effect is to consume less. These two work against one another, but here the substitution effect actually dominates.

This process for the markup substantially improves the fit of the model in terms of unconditional moments of quantities.

Series	Data		RBC		RBC w/ Markup Shocks	
	Relative S.D.	Corr. w/ y	Relative S.D.	Corr. w/ y	Relative S.D.	Corr. w/ y
y	1	1	1	1	1	1
c	0.53	0.76	0.44	0.95	0.38	0.66
I	2.76	0.79	3.13	0.99	4.24	0.96
n	1.12	0.88	0.44	0.98	1.01	0.84
y/n	0.65	0.42	0.63	0.99	0.56	0.27
TFP	0.71	0.76	0.75	0.99	0.55	0.76

These numbers are much better (relative to the basic RBC model) on two dimensions – first, the relative volatility of hours is much better. Second, the cyclicality of most of these variables (correlations with output) are much smaller – these are all essentially one in the basic model. This is particularly apparent for the cyclicality of labor productivity.

This all suggests that something like a markup shock may be an important part of business cycles. But do we really think that fluctuations in ν are really what's going on? Maybe, maybe not. Chari, Kehoe, and McGrattan (2008) propose a procedure which they call "business cycle" accounting, where they essentially identify "wedges", or residuals, from the first order conditions of a very basic RBC model. What they call the "labor wedge" is just the (negative) of what we identified as the markup above (hence the time series properties are identical). A plot of the labor wedge is shown below:

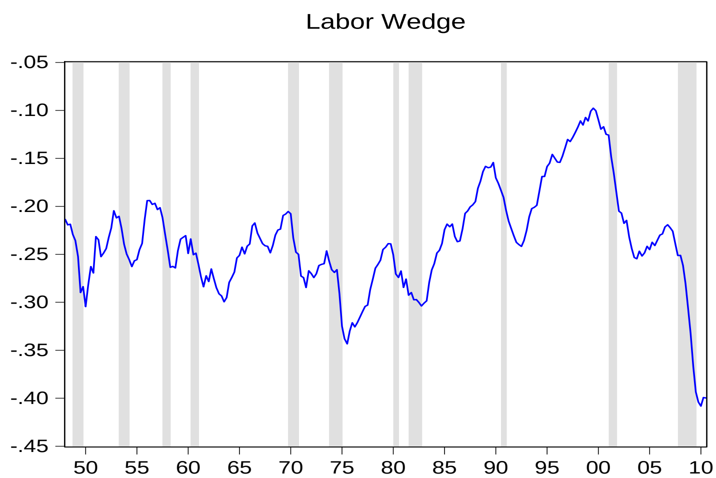

Here we see that the labor wedge is very procyclical (just the flip of the process we estimated for the markup). There are different explanations for what is going on here than just fluctuations in the markup, however. One explanation is variation in labor income tax rates, although we don't really see that at business cycle frequencies like we do here. Another would be frictions associated with reallocation. Another would be frictions associated with labor markets and matching. It could also be that the labor wedge isn't an independent shock, but is rather determined by some other shock. Any of these are possibilities. The labor wedge appears to be very important, and a lot of research in macro is currently devoted to understanding where it comes from.