Graduate Macro Theory II: Extensions of Basic RBC Framework

Eric Sims
University of Notre Dame

Spring 2024

1 Introduction

The basic RBC model – which is just a stochastic neoclassical growth model with variable labor – is the building block of most 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. I call this set of notes "RBC Extensions" because all the extensions I assume here are "real" – I do not yet deal with nominal rigidities, which one loosely think of as "New Keynesian." Most of the extensions we consider here can be considered attempts to make the model better "fit" the data and/or make it more realistic on some dimensions.

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. There are a couple of different ways to model indivisible labor; the end result is that the Frisch labor supply elasticity becomes infinite, which helps the model generate more volatility in hours.

Suppose there is one household with a continuum of members indexed by $i \in [0, 1]$. The head of the family makes consumption and work allocation decisions for each individual member, and also makes a consumption-savings decision (via capital or bonds). The utility function of any individual household is:

$$u(C_t(i), 1 - N_t(i)) = \ln C_t(i) + \theta \ln(1 - N_t(i))$$

Individual households can work $\bar{N}$ or 0 hours – there is no in-between. In other words, labor is indivisible. The head of the household choose how *many* of the individual members of the family work. Denote this by τ_t, where $0 < \tau_t < 1$. Total labor supplied by the household is therefore $\tau_t \bar{N}$. Labor supplied by individual members is 0 or $\bar{N}$.

The head of the household wants to maximize a utilitarian objective of intratemporal utility:

$$V_t = \int_0^1 (\ln C_t(i) + \theta \ln(1 - N_t(i)))\, di$$

The households who are chosen to work are chosen randomly. This means that we can write the integral:

$$V_t = \int_0^1 \ln C_t(i) di + \tau_t \theta \ln(1 - \bar{N})$$

This is because τ_t of the individuals have labor equal to $\bar{N}$, and the others have 0, and $\ln(1-0) = 0$. But this can be written:

$$V_t = \int_0^1 \ln C_t(i) di + \theta \tau_t \bar{N} \frac{\ln(1 - \bar{N})}{\bar{N}}$$

Since $N_t = \tau_t \bar{N}$, we can define:

$$D = -\frac{\theta \ln(1 - \bar{N})}{\bar{N}}$$

Since $0 < \bar{N} < 1$, we have $D > 0$ as defined, and preferences can therefore be written:

$$V_t = \int_0^1 \ln C_t(i) di - D N_t$$

The budget constraint facing the head of the household is:

$$\int_0^1 C_t(i) di + K_{t+1} - (1 - \delta) K_t + B_{t+1} \leq \int_0^1 w_t N_t(i) di + R_t^k K_t + (1 + r_{t-1}) B_t$$

Since households are randomly chosen to work, we can distribute the integral on the RHS, noting that $\int_0^1 N_t(i) di = \tau_t \bar{N} = N_t$. So we have:

$$\int_0^1 C_t(i) di + K_{t+1} - (1 - \delta) K_t + B_{t+1} \leq w_t N_t + R_t^k + (1 + r_{t-1}) B_t$$

Hence, we have written the problem as choosing N_t (rather than τ_t, since $N_t = \tau_t \bar{N}$, and $\bar{N}$ is a technological constraint. The full problem of the head of the household is therefore:

$$\max_{C_t(i), N_t, B_{t+1}, K_{t+1}} \quad \mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \left[\int_0^1 C_t(i) di - DN_t \right]$$

s.t.

$$\int_0^1 C_t(i) di + K_{t+1} - (1-\delta)K_t + B_{t+1} \leq w_t N_t + R_t^k K_t + (1 + r_{t-1}) B_t$$

A Lagrangian is:

$$\mathbb{L} = \mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \left[\int_0^1 C_t(i) di - DN_t + \lambda_t \left(w_t N_t + R_t^k K_t + (1 + r_{t-1}) B_t - \int_0^1 C_t(i) di - K_{t+1} + (1-\delta) K_t \right) \right]$$

The FOC are:

$$\frac{\partial \mathbb{L}}{\partial C_t(i)} = 0 \iff \frac{1}{C_t(i)} = \lambda_t \ \forall i$$
$$\frac{\partial \mathbb{L}}{\partial N_t} = 0 \iff D_t = \lambda_t w_t$$
$$\frac{\partial \mathbb{L}}{\partial K_{t+1}} = 0 \iff \lambda_t = \beta \mathbb{E}_t \lambda_{t+1} (R_{t+1}^k + (1-\delta))$$
$$\frac{\partial \mathbb{L}}{\partial B_{t+1}} = 0 \iff \lambda_t = \beta \mathbb{E}_t \lambda_{t+1} (1 + r_t)$$

The first of these FOC tells us that the head of the household will choose the same consumption for each member, which, given the unit mass assumption, will equal aggregate consumption, i.e. $C_t(i) = C_t \ \forall i$. Eliminating the multipliers, the FOC are then:

$$D_t = \frac{w_t}{C_t} \tag{1}$$
$$\frac{1}{C_t} = \beta \mathbb{E}_t \frac{1}{C_{t+1}} (R_{t+1}^k + (1-\delta)) \tag{2}$$
$$\frac{1}{C_t} = \beta \mathbb{E}_t \frac{1}{C_{t+1}} (1 + r_t) \tag{3}$$

It is *as though* there is just one representative household with *linear* preferences of labor/leisure! Linear preferences mean that the Frisch labor supply elasticity is *infinite*. The full set of equilibrium conditions are:

$$\frac{1}{C_t} = \beta \mathbb{E}_t \left(\frac{1}{C_{t+1}} \left(R^k_{t+1} + (1-\delta) \right) \right) \tag{4}$$

$$\frac{1}{C_t} = \beta \mathbb{E}_t \left(\frac{1}{C_{t+1}} (1 + r_t) \right) \tag{5}$$

$$D = \frac{1}{C_t} w_t \tag{6}$$

$$w_t = (1-\alpha) A_t K_t^{\alpha} N_t^{-\alpha} \tag{7}$$

$$R^k_t = \alpha A_t K_t^{\alpha-1} N_t^{1-\alpha} \tag{8}$$

$$K_{t+1} = I_t + (1-\delta) K_t \tag{9}$$

$$Y_t = C_t + I_t \tag{10}$$

$$Y_t = A_t A_t^{\alpha} N_t^{1-\alpha} \tag{11}$$

$$\ln A_t = \rho \ln A_{t-1} + s_A \varepsilon_t \tag{12}$$

We can calibrate $\bar{N}$ and θ to be consistent with $N = \frac{1}{3}$ (or whatever we want). Again, regardless of these, the implied Frisch elasticity is infinite. The model is isomorphic to a representative agent problem with preferences over consumption and labor as follows:

$$u(C_t, 1 - N_t) = \ln C_t - D \frac{N_t^{1+\chi}}{1+\chi}$$

Where $\chi = 0$, and the Frisch elasticity is χ^{-1}. I solve the model using our "standard" calibration in both this indivisible labor setup as well as the log-log setup. Below are impulse response to a technology shock in each model:

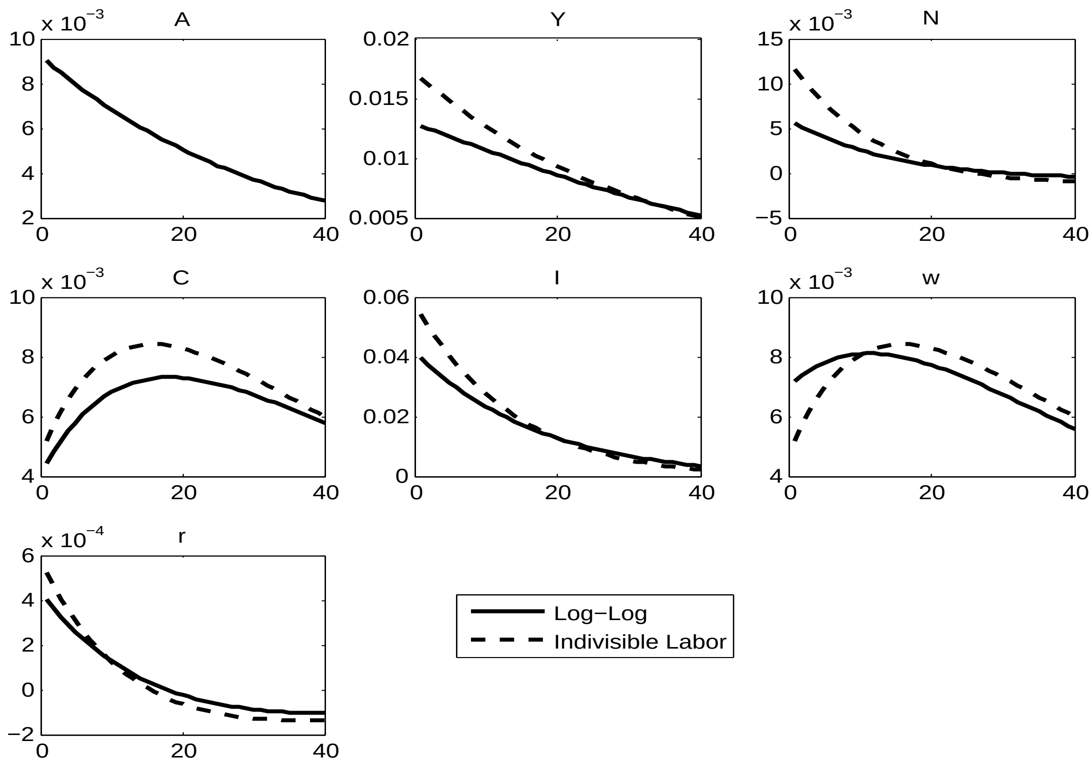

We see significantly more amplification in the indivisible labor case than in the log-log case, as shown by the differences between the dashed and solid lines. Output and hours both increase by significantly more on impact. As a result, consumption and investment both go up by more initially. Because the indivisible labor case is isomorphic to the Frisch labor supply being infinite, the labor supply curve is perfectly horizontal here. This is why we get a bigger increase in labor hours (and a smaller increase in wages) after the productivity shock.

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 productivity 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 0.69. 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 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.

2.2.1 Intratemporal Non-Separability: King, Plosser, and Rebelo (1988) and Greenwood, Hercowitz, and Hoffman (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 we typically abstract from 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:[1]

$$-u_N(C_t, 1 - N_t) = u_C(C_t, 1 - N_t)w_t$$

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 rate, 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 steady state N that is not growing.

To see a situation in which this might not work out as desired, consider a standard iso-elastic preference specification:

$$u(C_t, 1 - N_t) = \frac{C_t^{1-\sigma} - 1}{1 - \sigma} + \theta \frac{(1 - N_t)^{1-\xi} - 1}{1 - \xi}$$

The intratemporal first order condition for labor supply would then be:

$$\theta(1 - N_t)^{-\xi} = C_t^{-\sigma} w_t$$

[1] Note I am being a bit sloppy with notation; u_N refers to the partial with respect to N_t; I could use u_L to denote the partial with respect to the second argument, leisure. This would only matter insofar as there is a $-$ sign on the LHS or not.

The only way that we end up with $\frac{w_t}{C_t}$ on the right hand side is if $\sigma = 1$: in other words, for these preferences to be consistent with balanced growth, it must be the case that utility over consumption is log. This is potentially problematic, because it imposes that the coefficient of relative risk aversion is 1, which is much lower than what is needed to explain things like the equity premium.

King, Plosser, and Rebelo (1988) show that preferences must take the following form to be consistent with balanced growth (there are slightly different ways of writing this):

$$u(C_t, N_t) = \frac{(C_t v(1-N_t))^{1-\sigma} - 1}{1-\sigma} \quad \text{if} \quad \sigma \neq 1$$
$$u(C_t, N_t) = \ln C_t + \ln v(1-N_t) \quad \text{if} \quad \sigma = 1$$

The second step follows from application of L'Hopital's rule. We require that $v(1-N_t)$ be an increasing function of its argument, leisure (one minus labor). So as to make this all consistent with our original specification, suppose that $v(\cdot)$ takes the following form:[2]

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

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

[2]Note that you'll often see KPR preferences written with $v(1-N_t) = (1-N_t)^{\theta}$ or some variation thereof, so that $u(C_t, N_t) = \frac{(C_t(1-N_t)^{\theta})^{1-\sigma} - 1}{1-\sigma}$. When $\sigma \to 1$, this collapses to $\ln C_t + \theta \ln(1-N_t)$; this effectively imposes the "log-log" preference specification. By writing the $v(\cdot)$ function with the exp operator, I permit a more general specification of the utility from leisure.

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, along with N^* since we've written this in terms of leisure, 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:

$$\lambda_t = C_t^{-\sigma}\left(\exp\left(\theta\frac{(1-N_t)^{1-\xi}-1}{1-\xi}\right)\right)^{1-\sigma} \tag{13}$$

$$\lambda_t = \beta\mathbb{E}_t(\lambda_{t+1}(1+r_t)) \tag{14}$$

$$\lambda_t = \beta\mathbb{E}_t\lambda_{t+1}(R^k_{t+1}+(1-\delta)) \tag{15}$$

Log-linearize:

$$\begin{aligned}\ln\lambda_t &= \ln\beta + \ln\lambda_{t+1} + \ln(1+r_t)\\ \frac{\lambda_t-\lambda^*}{\lambda^*} &= \frac{\lambda_{t+1}-\lambda^*}{\lambda^*} + \frac{r_t-r^*}{1+r^*}\\ \widetilde{\lambda}_t &= \mathbb{E}_t\widetilde{\lambda}_{t+1} + \beta\widetilde{r}_t\end{aligned}$$

The last line follows from the fact that we define $\widetilde{r}_t$ as the actual deviation from steady state, not percentage deviation. Now log-linearize the expression for λ:

$$\begin{aligned}\ln\lambda_t &= -\sigma\ln C_t + (1-\sigma)\theta\left(\frac{(1-N_t)^{1-\xi}-1}{1-\xi}\right)\\ \frac{\lambda_t-\lambda^*}{\lambda^*} &= -\sigma\frac{C_t-c^*}{c^*} + (1-\sigma)\theta(1-N^*)^{-\xi}(N_t-N^*)\\ \widetilde{\lambda}_t &= -\sigma\widetilde{C}_t + (1-\sigma)\theta(1-N^*)^{-\xi}N^*\widetilde{N}_t\end{aligned}$$

Now combine these two expressions:

$$-\sigma\widetilde{C}_t + (1-\sigma)\theta(1-N^*)^{-\xi}N^*\widetilde{N}_t = \mathbb{E}_t\left(-\sigma\widetilde{C}_{t+1} + (1-\sigma)\theta(1-N^*)^{-\xi}N^*\widetilde{N}_{t+1}\right) + \beta\widetilde{r}_t$$

Simplify:

$$\mathbb{E}_t\left(\widetilde{C}_{t+1}-\widetilde{C}_t\right) = \frac{\beta}{\sigma}\widetilde{r}_t + \left(\frac{(1-\sigma)}{\sigma}\theta(1-N^*)^{-\xi}N^*\right)\left(\mathbb{E}_t\left(\widetilde{N}_{t+1}-\widetilde{N}_t\right)\right) \tag{16}$$

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 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: 1, 2, and 4.

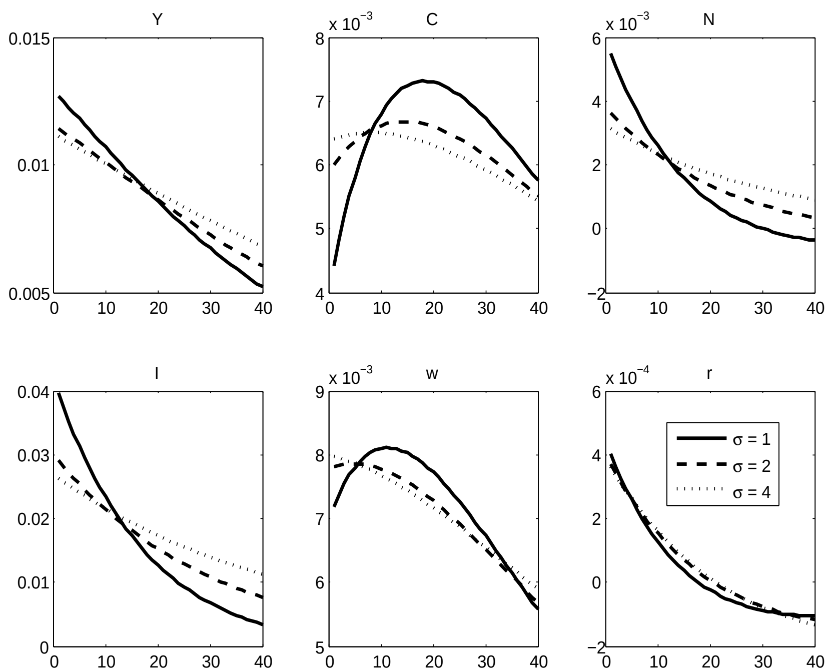

As we might expect, the initial jump in consumption is increasing in σ (note again that large σ means you want consumption to be smooth *in expectation*, not necessarily in response to a shock). This means that the jump in labor is decreasing in σ. Why is that? Think back the labor supply and demand curves. When productivity 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 – we see here that higher values of σ make the amplification problem in the model worse – we get less of a jump in labor because consumption jumps more, which means output goes up by less when productivity goes up. 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.

Greenwood, Hercowitz, and Hoffman (1988) propose another popular utility specification that features non-separability between consumption and leisure/labor. Unlike KPR preferences, GHH preferences are not consistent with balanced growth. What GHH preferences do is eliminate the wealth effect on labor supply – as we will see in a moment, this means that the FOC for labor gives labor as a function only of the wage (no consumption showing up). This means that there is no wealth effect – so, for example, when labor demand shifts out because of a technology shock, there is no inward shift of labor supply because of c_t increasing. This will result in more amplification (at the expense of not being consistent with balanced growth).

I will write utility in terms of disutility from labor instead of utility from leisure. The GHH preference specification can be written:

$$U(C_t, N_t) = \frac{1}{1-\sigma}\left(C_t - \theta\frac{N_t^{1+\chi}}{1+\chi}\right)^{1-\sigma}$$

The marginal utilities are:

$$U_C(C_t, N_t) = \left(C_t - \theta\frac{N_t^{1+\chi}}{1+\chi}\right)^{-\sigma}$$

$$u_N(C_t, N_t) = -\theta N_t^{\chi}\left(C_t - \theta\frac{N_t^{1+\chi}}{1+\chi}\right)^{-\sigma}$$

The generic first order condition for labor supply is:

$$-u_N(C_t, N_t) = u_C(C_t, N_t)w_t$$

With these marginal utilities, this works out to:

$$\theta N_t^{\chi} = w_t$$

Here we see, as promised, that the consumption term drops out altogether on the right hand side – labor is only a function of the wage. We can write the Euler equations for bonds and capital by defining auxiliary variables $\lambda_t = u_c$:

$$\lambda_t = \beta \mathbb{E}_t \lambda_{t+1}(R^k_{t+1} + (1-\delta))$$

$$\lambda_t = \beta \mathbb{E}_t \lambda_{t+1}(1 + r_t)$$

I solve the model using values $\sigma = 1$ (which corresponds to utility taking the form $\ln\left(C_t - \theta \frac{N_t^{1+\chi}}{1+\chi}\right)$) and $\chi = 1$, using standard values I've been using. Below, I compare the impulse responses under standard preferences (utility takes the form here $\ln C_t - \theta \frac{N_t^{1+\chi}}{1+\chi}$); in each case I solve for θ such that steady state labor hours are 0.33.

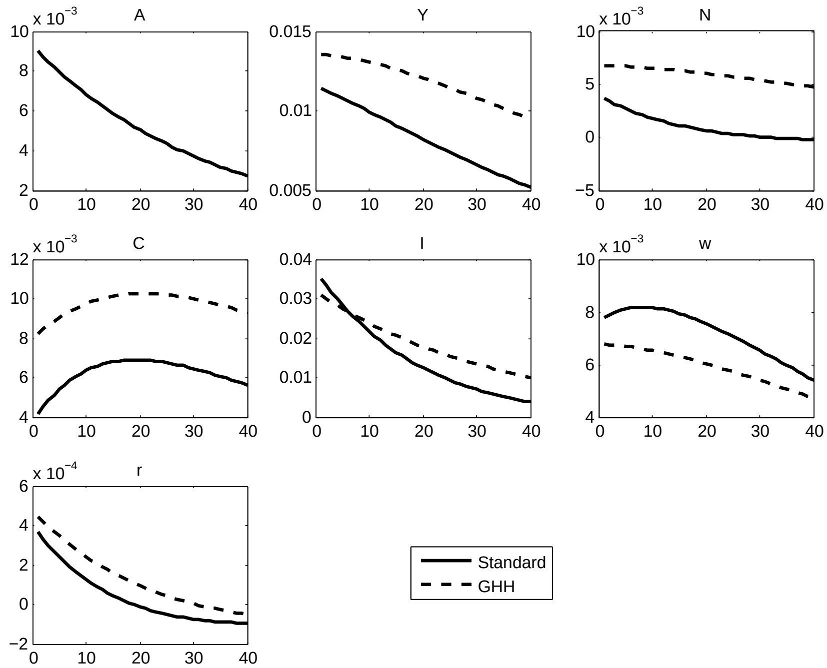

We see here that we get much larger amplification with GHH preferences – output and hours rise by substantially more (and appear to be more persistent). Naturally, the wage rises by less (this occurs because there is no inward shift of labor supply following the productivity shock). Interestingly, we see that consumption jumps up by more and actually investment increases by less. In terms of volatilities, under the standard setup I get output and labor volatilities (HP Filtered) of 0.015 and 0.005, for a relative volatility of about 1/3. For the GHH specification, I get output volatility of 0.0176 and labor volatility of 0.009, for about a relative volatility of 0.5. I could make these numbers look even better by lowering χ.

2.2.2 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:

$$\mathbb{L} = \mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \left(\ln(C_t - \phi C_{t-1}) + \theta \ln(1 - N_t) + \dots \right.$$
$$\left. \dots \lambda_t \Big(w_t N_t + R_t^k K_t + \Pi_t + (1 + r_{t-1}) B_t - C_t - K_{t+1} + (1 - \delta) K_t - B_{t+1} \Big) \right)$$

The first order conditions are:

$$\frac{\partial \mathbb{L}}{\partial C_t} = 0 \Leftrightarrow \lambda_t = \frac{1}{C_t - \phi C_{t-1}} - \beta \phi \mathbb{E}_t \frac{1}{C_{t+1} - \phi C_t} \tag{17}$$

$$\frac{\partial \mathbb{L}}{\partial N_t} = 0 \Leftrightarrow \frac{\theta}{1 - N_t} = \lambda_t w_t \tag{18}$$

$$\frac{\partial \mathbb{L}}{\partial K_{t+1}} = 0 \Leftrightarrow \lambda_t = \beta \mathbb{E}_t (\lambda_{t+1} (R_{t+1}^k + (1 - \delta)) \tag{19}$$

$$\frac{\partial \mathbb{L}}{\partial B_{t+1}} = 0 \Leftrightarrow \lambda_t = \beta \mathbb{E}_t (\lambda_{t+1} (1 + r_t)) \tag{20}$$

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 comparing a value of $\phi = 0$ (the standard case) with $\phi = 0.9$ (the habit formation case).

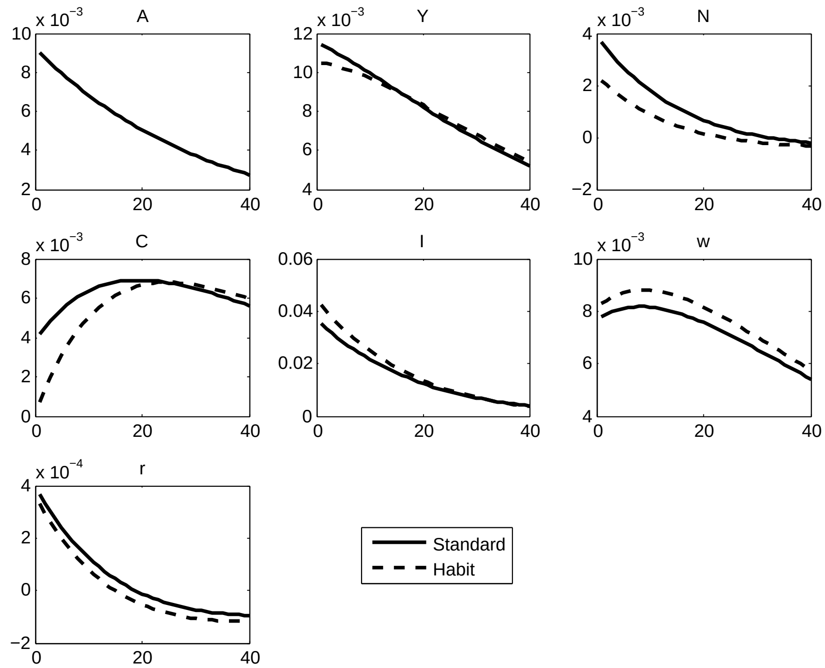

We observe that the main difference is that consumption jumps up by less on impact the bigger is ϕ. The intuition for this is high consumption today lowers utility tomorrow, other things being equal, the bigger is ϕ. Hence people will behave "cautiously" in essence by not adjusting consumption by much. The other impulse responses are reasonably similar across parameterizations, though 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.

Note that you can combine this kind of habit formation with different preferences specifications – e.g. I could embed this into the non-separable KPR or GHH preferences; I would just replace C_t with $(C_t - \phi C_{t-1})$ everywhere in both cases.

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). This ends up having the implication that the decentralized competitive equilibrium and the planner's solution do not coincide in the case of external habit – there is a non-internalized externality that the planner would take into account.

2.3 Capital and Investment Adjustment Costs

The standard RBC model has the implication that the price of capital goods relative to consumption is 1 (e.g. Hayashi's q is always 1). The standard RBC model also typically doesn't generate "hump-shaped" impulse responses to shocks, which seems to be a feature of the data. In this subsection we introduce convex costs to adjusting the capital stock. The idea is that there is some cost to adjusting the capital stock (or the level of investment) relative to what is normal. This will have the effect of breaking the $q = 1$ result in the basic model and will also impart some interesting dynamics in the model.

I'll begin with what I'll call "capital adjustment costs" because they follow the form set by Hayashi (1982). In particular, assume that the capital accumulation equation can be written:

$$K_{t+1} = I_t - \frac{\phi}{2}\left(\frac{I_t}{K_t} - \delta\right)^2 K_t + (1-\delta)K_t$$

When $\phi = 0$, this is the standard accumulation equation. If $\phi > 0$, then doing investment different than steady state (in steady state $\frac{I_t}{K_t} = \delta$) results in a cost which essentially makes your capital depreciate faster. Note that this cost is (i) denominated in units of current physical capital and (ii) is symmetric (so doing too little investment relative to steady state also costs you some capital, which may seem a bit funny).

Household preferences are standard. Instead of combining the flow budget constraint with the accumulation, let's treat them as separate (this turns out to be easier). The household problem is:

$$\max_{C_t, I_t, N_t, K_{t+1}, B_{t+1}} \quad \mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \left(\ln C_t - \theta \frac{N_t^{1+\chi}}{1+\chi}\right)$$

s.t.

$$C_t + I_t + B_{t+1} \leq w_t N_t + R_t^k K_t + \Pi_t + (1 + r_{t-1})B_t$$

$$K_{t+1} = I_t - \frac{\phi}{2}\left(\frac{I_t}{K_t} - \delta\right)^2 K_t + (1-\delta)K_t$$

Set up a Lagrangian with two constraints:

$$\mathbb{L} = \mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \Bigg\{ \ln C_t - \theta \frac{N_t^{1+\chi}}{1+\chi} + \lambda_t \left(w_t N_t + R_t^k K_t + \Pi_t + (1 + r_{t-1})B_t - C_t - I_t - B_{t+1}\right) + \dots$$
$$\dots + \mu_t \left(I_t - \frac{\phi}{2}\left(\frac{I_t}{K_t} - \delta\right)^2 K_t + (1-\delta)K_t - K_{t+1}\right) \Bigg\}$$

The first order conditions are:

$$\frac{\partial \mathbb{L}}{\partial C_t} = 0 \Leftrightarrow \frac{1}{C_t} = \lambda_t$$

$$\frac{\partial \mathbb{L}}{\partial N_t} = 0 \Leftrightarrow \theta N_t^{\chi} = \lambda_t w_t$$

$$\frac{\partial \mathbb{L}}{\partial I_t} = 0 \Leftrightarrow \lambda_t = \mu_t \left(1 - \phi \left(\frac{I_t}{K_t} - \delta\right)\right)$$

$$\frac{\partial \mathbb{L}}{\partial K_{t+1}} = 0 \Leftrightarrow \mu_t = \beta \mathbb{E}_t \left[R_{t+1}^k \lambda_{t+1} - \mu_{t+1} \frac{\phi}{2} \left(\frac{I_{t+1}}{K_{t+1}} - \delta\right)^2 + \mu_{t+1} \phi \left(\frac{I_{t+1}}{K_{t+1}} - \delta\right) \frac{I_{t+1}}{K_{t+1}} + \mu_{t+1}(1-\delta)\right]$$

$$\frac{\partial \mathbb{L}}{\partial B_{t+1}} = 0 \Leftrightarrow \lambda_t = \beta \mathbb{E}_t \lambda_{t+1}(1 + r_t)$$

Let's define $q_t \equiv \frac{\mu_t}{\lambda_t}$. μ_t is the marginal utility of having come extra installed capital (K_{t+1}), and λ_t is the marginal utility of having some extra consumption. The ratio is then how much consumption you would give up to have some extra future capital – i.e. it is the relative price of capital in terms of consumption. If $\phi = 0$, we see that $\lambda_t = \mu_t$, so $q_t = 1$ always. If $\phi > 0$, q_t can be different from 1. Using this formulation, we can write these FOC as:

$$q_t = \left(1 - \phi \left(\frac{I_t}{K_t} - \delta\right)\right)^{-1}$$

$$q_t = \beta \mathbb{E}_t \frac{C_t}{C_{t+1}} \left[R_{t+1}^k + q_{t+1} \left((1-\delta) + \phi \left(\frac{I_{t+1}}{K_{t+1}} - \delta\right) \frac{I_{t+1}}{K_{t+1}} - \frac{\phi}{2} \left(\frac{I_{t+1}}{K_{t+1}} - \delta\right)\right)\right]$$

The first expression establishes that investment to capital, $\frac{I_t}{K_t}$, is an increasing function of q_t: for this to be bigger than δ, q_t must exceed one. The second is functionally a difference equation in q_t: current q_t is a discounted value of the future marginal product of capital, future adjustment costs, and future q_{t+1}, where the discounting is by the household's stochastic discount factor, $\beta \frac{C_t}{C_{t+1}}$.

I solve the model using my standard parameter values (here $\chi = 1$) for three different values of ϕ: 0, 2, and 4. The impulse responses are below

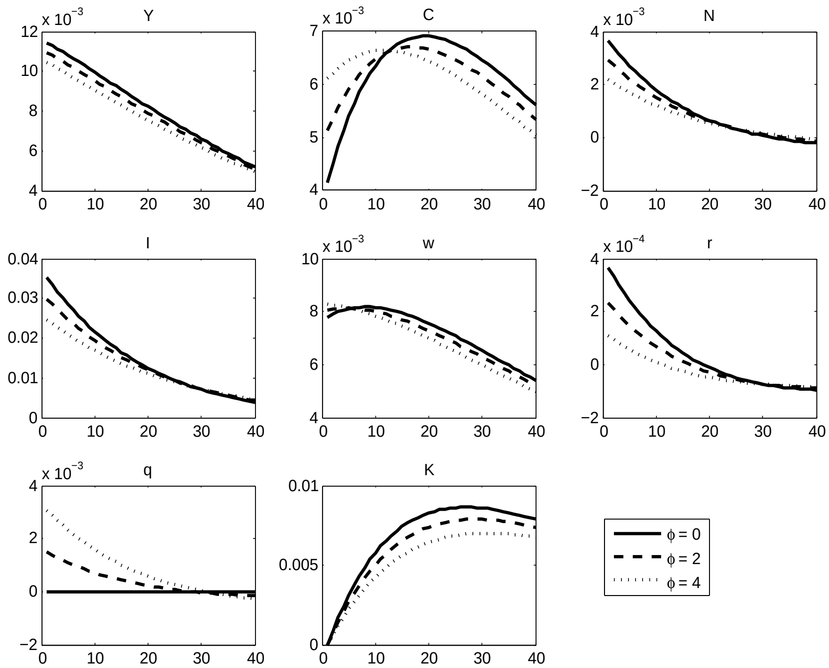

As predicted, if $\phi = 0$ then $q_t = 1$ always and the responses are the same as in the basic RBC model. As ϕ gets bigger, q_t rises more in response to the productivity shock. Quite naturally, we observe that investment goes up by less, and capital accumulates more slowly, the bigger is ϕ. Because it is costly to adjust the capital stock quickly, investment doesn't jump as much, which means consumption jumps more; this mechanically results in a smaller increase in hour than we would get because labor supply shifts in more. We also see that the real interest rate rises by less (significantly so). $\phi \neq 0$ breaks the tight connection between the real interest rate and the marginal product of capital; this is a good thing from the perspective of the model, since we'd really rather r_t not rise when hit with a productivity shock, since in the data the real interest rate is essentially acyclical.

An alternative adjustment cost specification is based on Christiano, Eichenbaum, and Evans (2005). I refer to this is an "investment adjustment cost" (as opposed to a capital adjustment cost). I call it an investment adjustment cost because (i) the adjustment cost is measured in units of investment, not units of capital as above, and (ii) the adjustment cost doesn't depend on the size of investment relative to the capital stock, but rather on the growth rate of investment. Let the capital accumulation equation be:

$$K_{t+1} = \left[1 - \frac{\phi}{2}\left(\frac{I_t}{I_{t-1}} - 1\right)^2\right] I_t + (1-\delta)K_t$$

In steady state $\frac{I_t}{I_{t-1}} = 1$, so this reverts to the standard accumulation equation. As we can see, the cost depends on the growth rate of investment and is measured in units of investment rather than units of capital.

The household problem is otherwise the same as before. As we did earlier, form a Lagrangian with two constraints:

$$\mathbb{L} = \mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \left\{ \ln C_t - \theta \frac{N_t^{1+\chi}}{1+\chi} + \lambda_t \left(w_t N_t + R_t^k K_t + \Pi_t + (1+r_{t-1})B_t - C_t - I_t - B_{t+1} \right) + \dots \right.$$

$$\left. \dots + \mu_t \left(\left[1 - \frac{\phi}{2} \left(\frac{I_t}{I_{t-1}} - 1 \right)^2 \right] I_t + (1-\delta)K_t - K_{t+1} \right) \right\}$$

The first order conditions are:

$$\frac{\partial \mathbb{L}}{\partial C_t} = 0 \Leftrightarrow \frac{1}{C_t} = \lambda_t$$

$$\frac{\partial \mathbb{L}}{\partial N_t} = 0 \Leftrightarrow \theta N_t^{\chi} = \lambda_t w_t$$

$$\frac{\partial \mathbb{L}}{\partial I_t} = 0 \Leftrightarrow \lambda_t = \mu_t \left(1 - \frac{\phi}{2} \left(\frac{I_t}{I_{t-1}} - 1 \right)^2 - \phi \left(\frac{I_t}{I_{t-1}} - 1 \right) \frac{I_t}{I_{t-1}} \right) + \beta \mathbb{E}_t \mu_{t+1} \phi \left(\frac{I_{t+1}}{I_t} - 1 \right) \left(\frac{I_{t+1}}{I_t} \right)^2$$

$$\frac{\partial \mathbb{L}}{\partial K_{t+1}} = 0 \Leftrightarrow \mu_t = \beta \mathbb{E}_t \left(\lambda_{t+1} R_{t+1}^k + (1-\delta)\mu_{t+1} \right)$$

$$\frac{\partial \mathbb{L}}{\partial B_{t+1}} = 0 \Leftrightarrow \lambda_t = \beta \mathbb{E}_t \lambda_{t+1}(1+r_t)$$

We can again define $q_t = \frac{\mu_t}{\lambda_t}$. Doing so, we can write:

$$q_t = \beta \mathbb{E}_t \frac{C_t}{C_{t+1}} \left(R_{t+1}^k + (1-\delta) q_{t+1} \right) \tag{21}$$

$$1 = q_t \left(1 - \frac{\phi}{2} \left(\frac{I_t}{I_{t-1}} - 1 \right)^2 - \phi \left(\frac{I_t}{I_{t-1}} - 1 \right) \frac{I_t}{I_{t-1}} \right) + \beta \mathbb{E}_t \frac{C_t}{C_{t+1}} q_{t+1} \phi \left(\frac{I_{t+1}}{I_t} - 1 \right) \left(\frac{I_{t+1}}{I_t} \right)^2 \tag{22}$$

Above I have made use of the fact that $\lambda_t = \frac{1}{C_t}$. Relative to the capital adjustment cost case earlier, the FOC for K_{t+1} is much simpler and defines q_t is the present discounted value of the rental rate on capital; the FOC for I_t is substantially more complicated. But you Can see that if I_t is large relative to steady state, then q_t will be greater than 1.

I solve the model quantitatively for this specification of adjustment costs for different values of the adjustment cost parameter ϕ. The impulse responses are below:

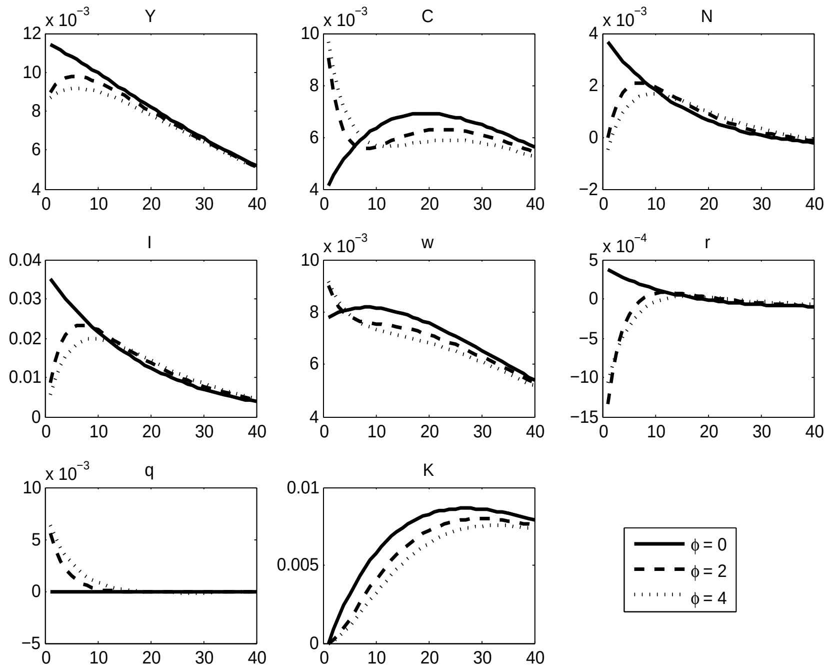

In terms of the dynamics of q_t and K_t, this specification of adjustment costs plays a fairly similar role to the capital adjustment cost specification – q_t rises more and K_t accumulates more slowly the bigger is ϕ. But the other dynamics are quite different. First, note that for positive values of ϕ both the investment and output impulse responses are "hump-shaped." This means that the *growth rates* of investment and output are autocorrelated, which is actually a feature of the data that the basic RBC model is incapable of matching. You can see why you get this in terms of investment growth rates just from the specification of the adjustment cost. Since there is a convex cost in the investment growth rate, you want to slowly adjust investment growth – – this gives you the hump-shaped investment response, which partially carries over into output. Because investment jumps up so little, consumption responds more to the productivity shock with bigger adjustment costs, which mechanically feeds into a smaller response of hours.

What is perhaps most marked in these responses is the response of the real interest rate. With these adjustment costs, the real interest rate actually declines after the productivity shock rather than increasing. As with the "capital adjustment cost" this investment adjustment cost breaks the connection between the real interest rate and the marginal product of capital – here R_t^k increases but r_t decreases. With these adjustment costs the breaking of this connection is much stronger than in the capital adjustment cost specification. This is useful because a major failure of the basic RBC model is the strong procylicality of the real interest rate predicted by the model, whereas in the data real interest rates are either acyclical or mildly countercyclical.

In DSGE models featuring many of these additions, it is now most common to include the investment adjustment cost specification in place of the Hayashi style capital adjustment cost specification, precisely because it can generate hump-shaped impulse responses (positive autocorrelation of growth rates) and works to brake the procyclicality of the real interest rate.

As an aside, there is actually a multi-sector version of this problem that gives the same results. Suppose that the household buys new capital goods, $\widehat{I}_t$, from a capital goods producing form at price P_t^k, which the household takes as given. The household's flow budget constraint (assuming that the household can only save via capital, and not bonds, since the amount of bonds are indeterminate anyway) and capital accumulation equation are:

$$C_t + P_t^k \widehat{I}_t \leq w_t N_t + R_t^k K_t + \Pi_t$$

$$K_{t+1} = \widehat{I}_t + (1-\delta)K_t$$

A Lagrangian for the household is:

$$\mathbb{L} = \mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \left[\ln C_t - \theta \frac{N_t^{1+\chi}}{1+\chi} + \lambda_t \left(w_t N_t + R_t^k K_t + \Pi_t - C_t - P_t^k \widehat{I}_t \right) + \mu_t \left(\widehat{I}_t + (1-\delta)K_t - K_{t+1} \right) \right]$$

The FOC are:

$$\frac{\partial \mathbb{L}}{\partial C_t} = 0 \iff \frac{1}{C_t} = \lambda_t$$

$$\frac{\partial \mathbb{L}}{\partial N_t} = 0 \iff \theta N_t^{\chi} = \lambda_t w_t$$

$$\frac{\partial \mathbb{L}}{\partial \widehat{I}_t} = 0 \iff \lambda_t P_t^k = \mu_t$$

$$\frac{\partial \mathbb{L}}{\partial K_{t+1}} = 0 \iff \mu_t = \beta \mathbb{E}_t \left[\lambda_{t+1} R_{t+1}^k + \mu_{t+1}(1-\delta) \right]$$

We see that $P_t^k = \frac{\mu_t}{\lambda_t}$. We can divide both sides of the last FOC by λ_t to write it:

$$\frac{\mu_t}{\lambda_t} = \beta \mathbb{E}_t \left[\frac{\lambda_{t+1}}{\lambda_t} R_{t+1}^k + \frac{\mu_{t+1}}{\lambda_t}(1-\delta) \right]$$

Or:

$$\frac{\mu_t}{\lambda_t} = \beta \mathbb{E}_t \left[\frac{\lambda_{t+1}}{\lambda_t} R_{t+1}^k + \frac{\lambda_{t+1}}{\lambda_t} \frac{\mu_{t+1}}{\lambda_{t+1}}(1-\delta) \right]$$

But then we have:

$$P_t^k = \mathbb{E}_t \Lambda_{t,t+1} \left(R_{t+1}^k + (1-\delta) P_{t+1}^k \right) \tag{23}$$

Where $\Lambda_{t,t+1} = \frac{\beta \lambda_{t+1}}{\lambda_t}$ is the stochastic discount factor. This is just a standard asset pricing condition. This is just a standard asset pricing condition.

Now consider the new capital goods producing firm. This firm uses unconsumed output, I_t, as an input to produce $\widehat{I}_t$ according to:

$$\widehat{I}_t = \left[1 - \frac{\phi}{2}\left(\frac{I_t}{I_{t-1}} - 1\right)^2\right] I_t$$

The profit maximization problem is dynamic because of I_{t-1} showing up. Future flow profits are discounted by the household's stochastic discount factor. The problem is:

$$\max_{I_t} \ \mathbb{E}_0 \sum_{t=0}^{\infty} \Lambda_{t,t+1} \left(P_t^k \left[1 - \frac{\phi}{2}\left(\frac{I_t}{I_{t-1}} - 1\right)^2\right] I_t - I_t \right)$$

The FOC is:

$$1 = P_t^k \left[1 - \frac{\phi}{2}\left(\frac{I_t}{I_{t-1}} - 1\right)^2 - \phi\left(\frac{I_t}{I_{t-1}} - 1\right)\frac{I_t}{I_{t-1}}\right] + \mathbb{E}_t \Lambda_{t,t+1} P_{t+1}^k \phi \left(\frac{I_{t+1}}{I_t} - 1\right)\left(\frac{I_{t+1}}{I_t}\right)^2 \quad (24)$$

The problem of the production firm is identical to our other setups. Total profit distributed from both types of firms to the household is:

$$\Pi_t = Y_t - w_t N_t - R_t^k K_t + P_t^k \widehat{I}_t - I_t$$

Plugging these into the household's budget constraint at equality yields the aggregate resource constraint:

$$Y_t = C_t + I_t$$

Subbing in for $\widehat{I}_t$ as a function of I_t and I_{t-1} into the capital accumulation equation yields:

$$K_{t+1} = \left[1 - \frac{\phi}{2}\left(\frac{I_t}{I_{t-1}} - 1\right)^2\right] I_t + (1-\delta)K_t$$

Now, if you compare the FOC in this setup, (23) and (24), to those form the earlier setup – (21) and (22) – you see that they are *identical* once we recognize that $q_t = \frac{\mu_t}{\lambda_t}$. The only difference is that in the former setup where there is no separate investment goods producing firm, q_t is an implicit relative price of new capital goods. In this setup with another layer to production, P_t^k is an explicit market price. But these models are observationally equivalent!

2.4 Non-Stationary Productivity

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, while the other productivity term followed a stationary AR(1) process:

$$Y_t = A_t K_t^\alpha (Z_t N_t)^{1-\alpha}$$
$$Z_t = (1+g_z)^t Z_0$$
$$\ln A_t = \rho \ln A_{t-1} + s_A \varepsilon_t$$

When writing down the model with no growth, as I have been doing, I've been implicity setting $g_z = 0$ and $z_0 = 1$. Allowing $g_z > 0$ requires re-writing the variables, but would only very mildly affect the equilibrium conditions (in essence, would have a small effect on the discount factor).

A priori, some people object to the notion of temporary productivity shocks – if one thinks of A_t as representing knowledge, does it make sense to forget things we've learned. Let's instead suppose that technology follows a stochastic trend. We can get rid of Z_t altogether and write the model as:

$$Y_t = A_t K_t^\alpha N_t^{1-\alpha}$$
$$\Delta \ln A_t = (1-\rho_A) g_z + \rho_A \Delta \ln A_{t-1} + s_A \varepsilon_t$$
$$\Delta \ln A_t = \ln A_t - \ln A_{t-1}$$

Here what I've assumed is that A_t follows an AR(1) process in the *growth rate*, with $0 \leq \rho_A < 1$. This means that shocks, ε_t, will have permanent effects on the level of A_t. g_z is the mean growth rate; if I set $\rho_A = 0$, then $\ln A_t$ would follow a random walk with drift.

We could write this process in the levels as:

$$\left(\frac{A_t}{A_{t-1}}\right) = \exp(g_z)^{1-\rho_A} \left(\frac{A_{t-1}}{A_{t-2}}\right)^{\rho_A} \exp(s_A \varepsilon_t)$$

If you take logs of this you get back the process written above. Because we only want one period of leads/lags in writing down the equilibrium conditions, it is useful to introduce a new variable, call it $g_t \equiv \frac{A_t}{A_{t-1}}$. We can then write the process above as:

$$g_t = \exp(g_z)^{1-\rho_a} g_{t-1}^{\rho_A} \exp(s_A \varepsilon_t)$$

Or in logs:

$$\ln g_t = (1-\rho_A) g_z + \rho_A \ln g_{t-1} + a_A \varepsilon_t$$

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:

$$
\begin{aligned}
Y_t &= A_t K_t^{\alpha} N_t^{1-\alpha} \\
\ln Y_t &= \ln A_t + \alpha \ln K_t + (1-\alpha)\ln N_t \\
(\ln Y_t - \ln Y_{t-1}) &= (\ln A_t - \ln A_{t-1}) + \alpha(\ln K_t - \ln K_{t-1}) + (1-\alpha)(\ln N_t - \ln N_{t-1})
\end{aligned}
$$

Along a balanced growth path hours will not grow and capital will grow at the same rate as output (capital and output must grow at the same rate because the real interest rate is constant along a balanced growth path, and the real interest rate in the long run is tied to the capital/output ratio). 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 $\frac{1}{1-\alpha}$ times 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 need to write it this way so that $\widehat{K}_t$ is a predetermined state variable in the model; scaling K_t by A_t instead would also render the model stationary, but would induce issues with the solution. 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 our change of variables, we can write this as:

$$\widehat{Y}_t = g_t^{-\frac{\alpha}{1-\alpha}} \widehat{K}_t^{\alpha} N_t^{1-\alpha}$$

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 our change of variable:

$$\widehat{K}_{t+1} = \widehat{I}_t + (1-\delta)g_t^{-\frac{1}{1-\alpha}}\widehat{K}_t$$

The accounting identity is the same in terms of the transformed variables as always: $\widehat{Y}_t = \widehat{C}_t + \widehat{I}_t$. Next, consider expressions for factor prices:

$$\begin{aligned} w_t &= (1-\alpha)A_t K_t^{\alpha} N_t^{-\alpha} \\ R_t^k &= \alpha A_t K_t^{\alpha-1} N_t^{1-\alpha} \end{aligned}$$

Transform these:

$$\frac{w_t}{A_t^{\frac{1}{1-\alpha}}} = (1-\alpha)A_t^{-\frac{\alpha}{1-\alpha}} K_t^{\alpha} N_t^{-\alpha}$$

$$\widehat{w}_t = (1-\alpha)A_t^{-\frac{\alpha}{1-\alpha}} A_{t-1}^{\frac{\alpha}{1-\alpha}} \left(\frac{K_t}{A_{t-1}^{\frac{1}{1-\alpha}}}\right)^{\alpha} N_t^{-\alpha}$$

$$\widehat{w}_t = (1-\alpha)g_t^{-\frac{\alpha}{1-\alpha}} \widehat{K}_t^{\alpha} N_t^{-\alpha}$$

For the rental rate, we have:

$$R_t^k = \alpha A_t A_{t-1}^{\frac{\alpha-1}{1-\alpha}} \left(\frac{K_t}{A_{t-1}^{\frac{1}{1-\alpha}}}\right)^{\alpha-1} N_t^{1-\alpha}$$

$$R_t^k = \alpha\left(\frac{A_t}{A_{t-1}}\right)\widehat{K}_t^{\alpha-1} N_t^{1-\alpha}$$

$$R_t^k = \alpha g_t \widehat{K}_t^{\alpha-1} N_t^{1-\alpha}$$

Now what about the Euler equations? Start with the one for bonds:

$$\frac{1}{C_t} = \beta \mathbb{E}_t \frac{1}{C_{t+1}} (1 + r_t)$$

$$\frac{A_t^{\frac{1}{1-\alpha}}}{C_t} = \beta \mathbb{E}_t \frac{A_t^{\frac{1}{1-\alpha}}}{C_{t+1}} (1 + r_t)$$

$$\frac{A_t^{\frac{1}{1-\alpha}}}{C_t} = \beta \mathbb{E}_t \frac{A_{t+1}^{\frac{1}{1-\alpha}}}{A_{t+1}^{\frac{1}{1-\alpha}}} \frac{A_t^{\frac{1}{1-\alpha}}}{C_{t+1}} (1 + r_t)$$

$$\frac{1}{\widehat{C}_t} = \beta \mathbb{E}_t g_{t+1}^{-\frac{1}{1-\alpha}} \frac{1}{\widehat{C}_{t+1}} (1 + r_t)$$

Since the rental rate on capital is stationary, the Euler equation for capital is going to look the same in terms of transformed variables:

$$\frac{1}{\widehat{C}_t} = \beta \mathbb{E}_t g_{t+1}^{-\frac{1}{1-\alpha}} \frac{1}{\widehat{C}_{t+1}} \left(R_{t+1}^k + (1 - \delta) \right)$$

The labor supply condition is straightforward to re-write in terms of transformed variables – since $\frac{w_t}{C_t}$ shows up on the right hand side, and we divide both variables by the same scaling factor, we can simply write:

$$\theta \frac{1}{1 - N_t} = \frac{1}{\widehat{C}_t} \widehat{w}_t$$

Then the full set of equilibrium conditions are:

$$\frac{1}{\widehat{C}_t} = \beta \mathbb{E}_t g_{t+1}^{-\frac{1}{1-\alpha}} \frac{1}{\widehat{C}_{t+1}} (1 + r_t)$$

$$\frac{1}{\widehat{C}_t} = \beta \mathbb{E}_t g_{t+1}^{-\frac{1}{1-\alpha}} \frac{1}{\widehat{C}_{t+1}} \left(R_{t+1}^k + (1 - \delta) \right)$$

$$\theta \frac{1}{1 - N_t} = \frac{1}{\widehat{C}_t} \widehat{w}_t$$

$$\widehat{w}_t = (1 - \alpha) g_t^{-\frac{\alpha}{1-\alpha}} \widehat{K}_t^{\alpha} N_t^{-\alpha}$$

$$R_t^k = \alpha g_t \widehat{K}_t^{\alpha - 1} N_t^{1-\alpha}$$

$$\widehat{Y}_t = \widehat{C}_t + \widehat{I}_t$$

$$\widehat{K}_{t+1} = \widehat{I}_t + (1 - \delta) g_t^{-\frac{1}{1-\alpha}} \widehat{K}_t$$

$$\widehat{Y}_t = g_t^{-\frac{\alpha}{1-\alpha}} \widehat{K}_t^{\alpha} N_t^{1-\alpha}$$

$$\ln g_t = (1 - \rho_A) g_z + \rho_A \ln g_{t-1} + s_A \varepsilon_t$$

This is nine equations in nine variables.

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 transform the responses produced by Dynare. First, I can find the response of the log level of A_t by cumulating the response of the log growth rate, $\ln g_t$. Then once I have the impulse response of A_t, I can "add back" $\frac{1}{1-\alpha}\ln A_t$ to the impulse responses of the other variables to get them back in log level form (recall that Dynare will produce the impulse response of the logs, so $\ln \widehat{x}_t$. Since $\widehat{x}_t \equiv \frac{x_t}{A_t^{\frac{1}{1-\alpha}}}$, I can add $\frac{1}{1-\alpha}\ln A_t$ to $\ln \widehat{x}_t$ to recover the log level of the variable of interest; I only need to do this for variables that are non-stationary – so not for the real interest rate, the rental rate on capital, or the level of employment).

I compute impulse responses under some different parameter configurations in order to make a few points. First, I show IRFs under a stationary shock with $\rho_A = 0.979$. Second, I show two different permanent shocks, one with $\rho_A = 0$ (random walk) and the other with $\rho_A = 0.33$ (so some positive autocorrelation in the growth rate). I set the size of the transitory shock to $s_A = 0.01$ (so productivity increases by one percent); in the case of the two permanent shocks, I set the size of the shocks such that the long run effect on productivity is also 0.01 (so one percent); this means that the size of the shock is 0.01 in the random walk case, but $(1-\rho_A)0.01$ more generally. So for the shock that is persistent in growth rates, I'm actually hitting the economy with a smaller initial shock.

Below, the solid black lines show the responses to the transitory productivity shock; the dashed black line shows the responses to the random walk productivity shock; the blue solid line shows the responses to the persistent growth rate shock.

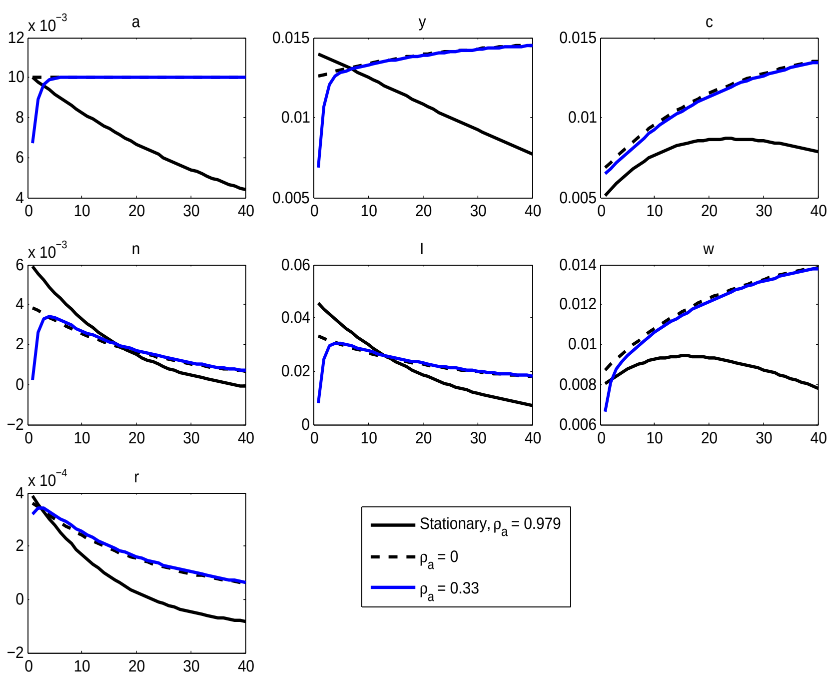

The main take-away from this picture is that the impact effects on output, hours, and investment are smaller the more persistent is the productivity shock. For the shock that is autocorrelated in growth rates, we see that labor hours essentially do not react. Also worth noting is that after about 10 periods the responses with the random walk shock and the persistent growth rate shock are very similar.

What is going on here? It's the same intuition from labor demand-supply that we had earlier. If the productivity shock is permanent, you get a much bigger initial jump in consumption than you do in the transitory shock case; this translates to a bigger inward shift of the labor supply curve, which for a given labor demand shift results in a smaller increase in N_t. A smaller increase in N_t means that output goes up by less on impact; this combined with a bigger increase in C_t means a smaller increase in investment. One thing that is interesting to note is that the impact increase in C_t is actually smaller in the case of the autocorrelated shock than in the case of the random walk shock. What is going on here? Consumption is very forward-looking; by the way I constrained the size of the shocks, the long run effects are identical in each case. Because you get more of the extra productivity sooner in the random walk case than later in the persistent growth rate case, the wealth effect is actually bigger in the random walk case, so C_t jumps by more (although not by a lot). Why then do we observe a smaller increase in N_t in the autocorrelated growth rates case than the random walk case if the wealth effect / jump in c_t is smaller? It's because the outward shift in labor demand is also smaller given that I constrained the long run effect on productivity to be the same in both cases; when the shock is autocorrelated, the immediate effect on productivity (and hence labor demand) is smaller. This combines to result in a much smaller increase in N_t, a smaller initial increase in Y_t, and a much smaller initial increase in I_t.

The bottom line here is that it is (relatively) straightforward to augment the model to allow for permanent productivity shocks, but this will only make the model fit the data worse along the dimensions of the relative volatility of hours as well as the overall lack of amplification.

2.5 Variable Capital Utilization

In this section we consider adding an amplification mechanism to the model – variable capital utilization. The idea here is that while the capital stock may be predetermined within period, but we can "work" the capital more intensively (or not) depending on conditions. One could also think about variable labor utilization, but for this to work we'd need to make labor at least partially predetermined – you'll only get a unique value of utilization if you can't completely change the factor (capital in our main case, in this case labor).

There are different ways to model capital utilization and the costs associated with it. What I'm going to do here assumes (i) the households own the capital stock, (ii) the households choose the level of utilization, and (iii) the household leases "capital services" (the product of utilization and physical capital) to firms at rental rate R_t^k. The cost of utilization is faster depreciation. There is another version of the setup where the cost of utilization is a resource cost; to a first-order approximation, this will be the same as the faster depreciation setup.

Define u_t to be utilization and $\widehat{K}_t \equiv u_t K_t$ as capital services (K_t is the physical capital stock).

The depreciation rate is now a function of capital utilization, which we want to normalize to be one in steady state. In particular:

$$\delta(u_t) = \delta_0 + \phi_1(u_t - 1) + \frac{\phi_2}{2}(u_t - 1)^2$$

The household problem is otherwise standard. We can write it as:

$$\max_{C_t, N_t, K_{t+1}, u_t} \quad E_0 \sum_{t=0}^{\infty} \beta^t \left(\ln C_t + \theta \ln(1 - N_t)\right)$$

s.t.

$$C_t + K_{t+1} - (1 - \delta(u_t))K_t + B_{t+1} \leq w_t N_t + R_t^k u_t K_t + \Pi_t + (1 + r_{t-1})B_t$$

The first order conditions for bonds/consumption and labor work out to be the same as we had earlier:

$$\frac{1}{C_t} = \beta \mathbb{E}_t \frac{1}{C_{t+1}}(1 + r_t)$$

$$\theta \frac{1}{1 - N_t} = \frac{1}{C_t} w_t$$

The first order condition for utilization works out to:

$$\delta'(u_t) = R_t^k$$

Or:

$$\phi_1 + \phi_2(u_t - 1) = R_t^k$$

The Euler equation for capital is:

$$\frac{1}{C_t} = \beta \mathbb{E}_t \frac{1}{C_{t+1}} \left(R_{t+1}^k u_{t+1} + (1 - \delta(u_{t+1})) \right)$$

The firm produces output using:

$$Y_t = A_t \widehat{K}_t^{\alpha} N_t^{1-\alpha}$$

Note that the firm acts as though it gets to choose capital services at rental rate R_t^k, even though household can separately choose utilization and capital. The first order conditions of the firm problem are otherwise standard:

$$w_t = (1 - \alpha) A_t \widehat{K}_t^{\alpha} N_t^{-\alpha}$$

$$R_t^k = \alpha A_t \widehat{K}_t^{\alpha - 1} N_t^{-\alpha}$$

The aggregate resource constraint is standard and we assume that productivity follows an AR(1) in the log. The full set of equilibrium conditions can be written (getting ride of $\widehat{K}_t$ by replacing it with $u_t K_t$):

$$\frac{1}{C_t} = \beta \mathbb{E}_t \frac{1}{C_{t+1}} (1 + r_t)$$

$$\theta \frac{1}{1 - N_t} = \frac{1}{C_t} w_t$$

$$\phi_1 + \phi_2(u_t - 1) = R_t^k$$

$$\frac{1}{C_t} = \beta \mathbb{E}_t \frac{1}{C_{t+1}} \left(R_{t+1}^k u_{t+1} + (1 - \delta(u_{t+1})) \right)$$

$$w_t = (1 - \alpha) A_t u_t^\alpha K_t^\alpha N_t^{-\alpha}$$

$$R_t^k = \alpha A_t u_t^{\alpha - 1} K_t^{\alpha - 1} N_t^{-\alpha}$$

$$Y_t = A_t u_t^\alpha K_t^\alpha N_t^{1-\alpha}$$

$$Y_t = C_t + I_t$$

$$K_{t+1} = I_t + (1 - \delta(u_t)) K_t$$

$$\delta(u_t) = \delta_0 + \phi_1(u_t - 1) + \frac{\phi_2}{2}(u_t - 1)^2$$

$$\ln A_t = \rho \ln A_{t-1} + s_A \varepsilon_t$$

This is 11 equations in 11 variables. The parameters of the utilization cost function are not entirely free. To normalize $u = 1$ in steady state, we must have: $R^k = \frac{1}{\beta} - (1 - \delta_0)$. But from the FOC for utilization this requires:

$$\frac{1}{\beta} - (1 - \delta_0) = \phi_1$$

In other words, ϕ_1 must be set to pin down steady state utilization at 1; ϕ_2 is a free parameter, and δ_0 is the steady state depreciation rate, which we can calibrate to match the steady state investment-output ratio as before.

We can think about $\phi_2 \to \infty$ as fixing $u_t = 1$, which puts us back in the standard RBC model.

This may not be obvious, because one would be tempted to look at the above expression and think this fixed R_t^k at its steady state value, $\frac{1}{\beta} - (1 - \delta_0)$. This is not so, because ∞ times 0 is something finite.

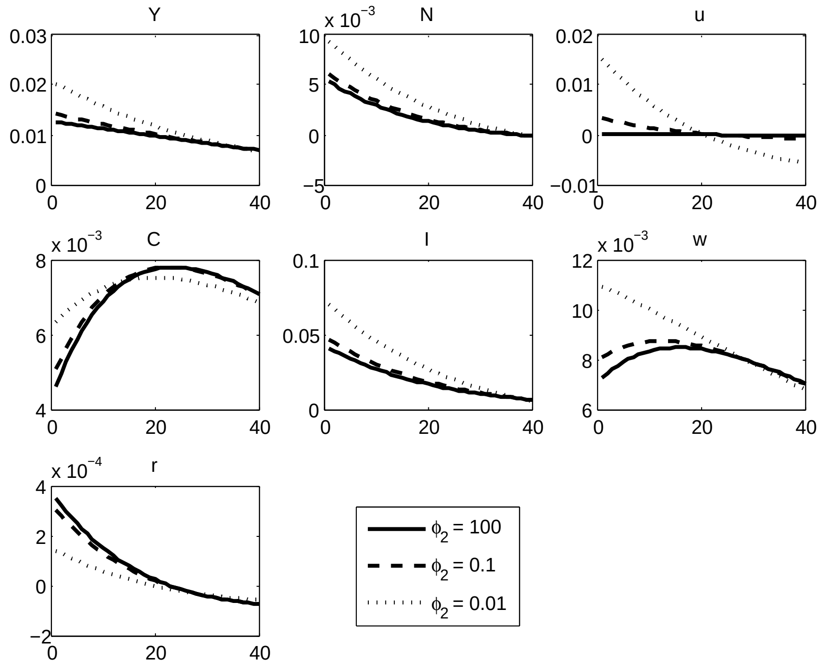

We can see that there is significantly more amplification of the productivity shock when the costs of utilization are lower (i.e. when ϕ_2 is smaller): output, employment, consumption, and investment rise by significantly more with variable utilization than without. Without variable utilization ($\phi_2 = 100$), HP filtered output volatility is about 0.017; as I decrease ϕ_2, I get output volatility of 0.019 and 0.027 ($\phi_2 = 0.1$ and $\phi_2 = 0.01$, respectively). This means that, in principle, we'd need smaller productivity shocks to generate the observed output volatility we observe in the data, which many people find attractive. We also see that the inclusion of variable utilization (slightly) increases the relative volatility of employment: its volatility relative to output goes from 0.43 in the no utilization case to 0.49 with $\phi_2 = 0.01$.

Finally, and importantly, it is worth pointing out that variable capital utilization invalidates standard "growth accounting" techniques to get measured TFP. Measured TFP is defined as $\ln \widehat{A}_t = \ln Y_t - \alpha \ln K_t - (1-\alpha) \ln N_t$. In this model, this is equal to $\ln A_t + \alpha \ln u_t$. Hence, to the extent to which utilization moves around, the volatility of measured TFP will be an overstatement of the volatility of the model counterpart, $\ln A_t$. We see this in the quantitative simulations here. The volatility of actual $\ln A_t$ is 0.0117; with no utilization the volatility of measured TFP is the same, with $\phi_2 = 0.1$ it is 0.0131, and with $\phi_2 = 0.01$ it is 0.0184. This is a big difference.

Because the model is only driven by one shock (the productivity shock), the inclusion of variable capital utilization doesn't do much to change the cyclicality of measured TFP versus true productivity – they are both very positively correlated with output. But suppose there is some shock which raises N_t (like a government spending shock, or perhaps a preference shock). Because labor and capital services are complementary in the model, anytime N_t goes up you'd like to also increase u_t. But from above, this means that measured TFP will rise (along with output) even though actual $\ln A_t$ won't change. Thus, variable capital utilization will change both the volatility and the cyclicality of measured TFP – measured TFP will be both more volatile and more positively correlated with output with variable capital utilization than the true exogenous productivity variable, $\ln A_t$, is.

2.6 Preference Shocks

The basic RBC model focuses on the productivity shock, but it is possible and fruitful to consider other sources of shocks. It is popular to write down models where there are preference shocks – shocks to how agents get utility from consumption and/or leisure/labor. I will write down the standard model with two such shocks – an intertemporal preference shock that governs how you weight current utility relative to future utility, and an intratemporal preference shock that governs how you value utility from labor/leisure.

Consider the following household problem:

$$\max_{C_t, K_{t+1}, N_t, B_{t+1}} \quad E_0 \sum_{t=0}^{\infty} \beta^t \psi_t \left(\ln C_t - \nu_t \theta \frac{N_t^{1+\chi}}{1+\chi} \right)$$

$$\text{s.t.}$$

$$C_t + K_{t+1} - (1-\delta)K_t + B_{t+1} - B_t \leq w_t N_t + R_t^k K_t + \Pi_t + r_{t-1} B_t$$

I assume that both ψ_t and ν_t follow mean zero AR(1)s in the log (so that the non-stochastic levels are unity):

$$\ln \psi_t = \rho_\psi \ln \psi_{t-1} + \varepsilon_{\psi,t}$$

$$\ln \nu_t = \rho_\nu \ln \nu_{t-1} + \varepsilon_{\nu,t}$$

The exogenous variable ψ_t is an intertemporal preference shock – it doesn't impact you value utility from consumption versus utility from leisure, but rather how you value utility today versus utility in the future. If ψ_t increases, for example you place relatively more weight on present utility than the future. ν_t is an intratemporal preference shock – it affects how you value utility from consumption relative to disutility from labor (or utility from leisure).

The first order conditions of the model can be written:

$$\frac{1}{C_t} = \beta \mathbb{E}_t \left[\frac{\psi_{t+1}}{\psi_t} \frac{1}{C_{t+1}} \left(R_{t+1}^k + (1-\delta) \right) \right]$$

$$\frac{1}{C_t} = \beta \mathbb{E}_t \left[\frac{\psi_{t+1}}{\psi_t} \frac{1}{C_{t+1}} (1 + r_t) \right]$$

$$\nu_t \theta N_t^{\chi} = \frac{1}{C_t} w_t$$

Two things to point out. First, ψ_t does not show up in the labor supply condition: higher ψ_t increases the marginal utility of both consumption and the marginal disutility of labor, but these cancel out. Second, since ψ_t is mean reverting, when ψ_t is high we will have $E_t \frac{\psi_{t+1}}{\pi_t}$ be relatively low. This means that an increase in ψ_t is isomorphic to a temporary reduction in β – it means you are relatively more impatient. Third, ν_t shows up in the labor supply condition in a way analagous to a distortionary tax rate on labor income – if you divide both sides by ν_t you see that an increase in ν_t is functionally equivalent to an increase in τ_t^n (since $(1-\tau_t^n)$ would show up on the right hand side of the condition).

I solve the model quantitatively using a first order log-linear approximation using our "standard" parameter values (a Frisch elasticity of 1, $\chi = 1$). I assume that the AR coefficients on the two preference shocks are 0.9 (i.e. $\rho_\psi = \rho_\nu = 0.9$) and that the standard deviations of each shock are one percent (i.e. 0.01). I'm not trying to parameterize these in any serious way; and for impulse response analysis in a first order approximation, the size of the shocks is irrelevant for the shape of the IRFs, thought it may have large effects on unconditional moments.

Below are the responses to the intertemporal preference shock, ψ_t. We see that output, hours, investment, and the real interest rate all decline immediately, while consumption and the real wage increase. What is going on here is the following. The increase in ψ_t is effectively like a decrease in the discount factor – households value current utility relatively more than future utility. This means they want to consume more in the present and work less – hence the increase in C_t and decline in N_t (in a mechanical sense from the FOC for labor the increase in C_t shifts the labor supply curve in). The inward shift of labor supply along a stable labor demand curve leads to an increase in the wage. Falling hours with no immediate change in A_t or K_t means that Y_t must fall. Output falling with consumption increasing means that investment must fall. The real interest rate must fall immediately. There are two ways to see this. First, since consumption is high and falling, the real interest rate must fall for the Euler equation to hold. Second, the fall in N_t lowers the marginal product of capital, and with capital fixed this means that the rental rate on capital must decline, and without adjustment costs r_t moves in the same direction as R_t^k. Overall, while this shock produces interesting dynamics, it does not produce positive co-movement between consumption and output, and hence cannot be the primary driving force of the business cycle. This is a manifestation of the "Barro-King curse."

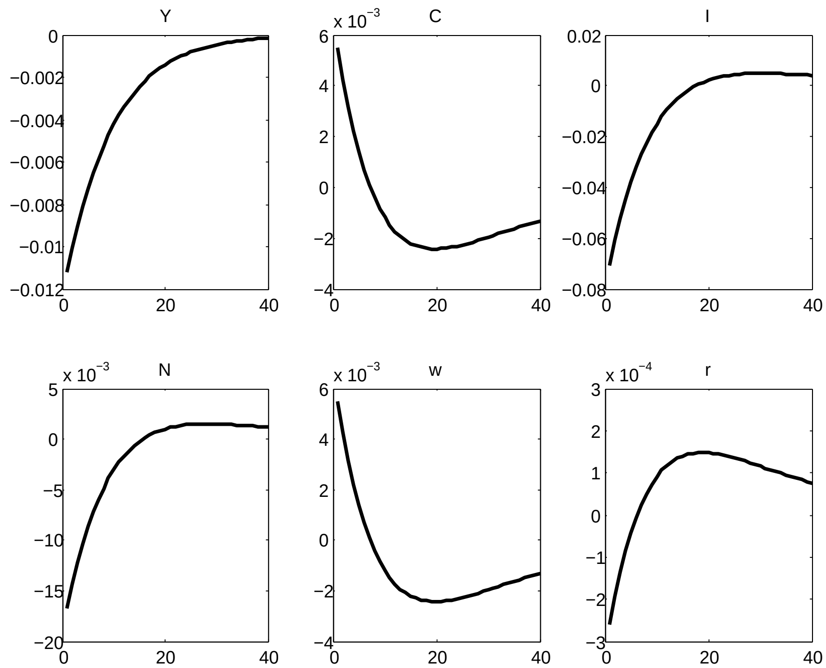

Next, consider the intratemporal preference shock. This leads to a reduction in C_t, N_t, Y_t, and I_t, with an increase in w_t. The increase in ψ_t means that people dislike labor relatively more – this means naturally that they want to work less. This shifts the labor supply curve in; along a stable labor demand curve, this means that the wage must rise. Lower employment means lower output. Consumption also falls – this occurs naturally because household income declines. Investment also falls.

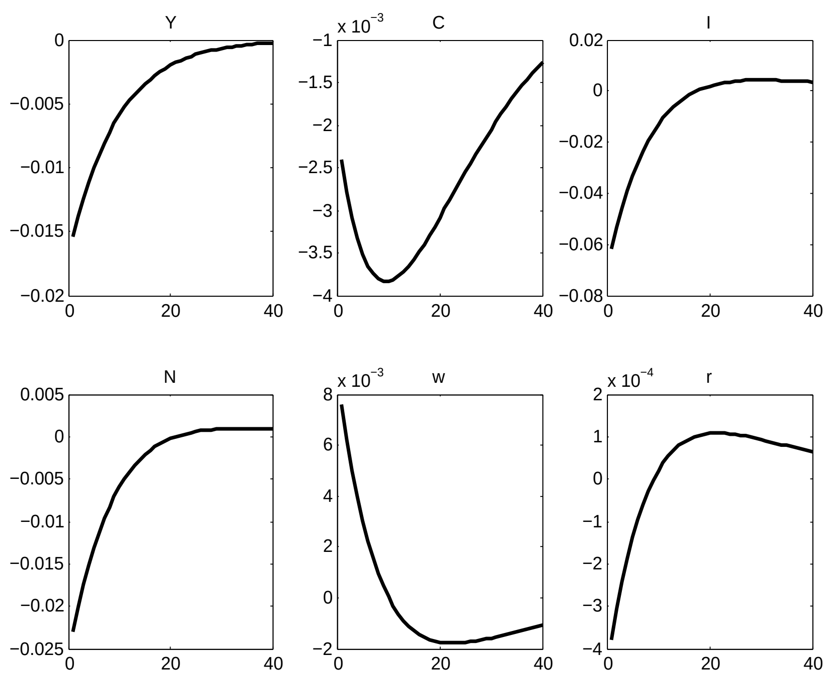

We see here that, unlike the intertemporal preference shock, the intratemporal preference shock can produce co-movement between consumption and output and employment. I've said before that the FOC for labor makes it difficult for consumption and employment to move together unless A_t changes (the Barro-King curse). In a mechanical sense here, an increase in ν_t functionally plays a similar role (although in a different way). What you need to get co-movement between C_t and N_t is for *either* the labor demand or supply curves to shift for a reason other than pure wealth effect of C_t shifting the labor supply curve. ν_t will do the trick, as would a change in the tax rate on labor income (see discussion above). This shock is kind of nice in terms of fitting the data – it generates co-movement, but a conditionally countercyclical real wage.

2.7 Investment Shocks

The standard RBC model assumes that shocks to neutral productivity are the primary (or sole) driver of business cycle fluctuations. Another kind of disturbance that has recently received a good deal of attention is a shock to the marginal efficiency of investment (MEI), or just "investment shock" for short. The investment shock makes the economy more productive at transforming investment into new physical capital (in a way somewhat analogous to how the productivity shock, A_t, makes you more productive at transforming capital and labor into output).

Let Z_t denote the investment shock. It enters the capital accumulation equation as follows:

$$K_{t+1} = Z_t I_t + (1-\delta)K_t$$

Here, an increase in Z_t means you get more K_{t+1} for a given amount of I_t – i.e. this shock increases the efficiency of investment. Some authors have argued that this shock is a reduced-form proxy for modeling the health of the financial system – the financial system essentially turns investment into capital, so the higher (or lower) Z_t is, the better (or worse) the financial system is.

Let's again assume that the household owns the capital stock and leases it to firms. The household problem can be written with two constraints:

$$\max_{C_t, N_t, K_{t+1}, B_{t+1}} \quad E_0 \sum_{t=0}^{\infty} \beta^t \left\{ \ln C_t - \theta \frac{N_t^{1+\chi}}{1+\chi} \right\}$$

s.t.

$$C_t + I_t + B_{t+1} - B_t \leq w_t N_t + R_t^k K_t + \Pi_t + r_{t-1} B_t$$

$$K_{t+1} = Z_t I_t + (1-\delta)K_t$$

These constraints can be combined into one:

$$C_t + \frac{K_{t+1}}{Z_t} - (1-\delta)\frac{K_t}{Z_t} + B_{t+1} - B_t \leq w_t N_t + R_t^k K_t + \Pi_t + r_{t-1} B_t$$

The first order conditions for bonds, labor, and capital can be written:

$$\frac{1}{C_t} = \beta \mathbb{E}_t \left(\frac{1}{C_{t+1}} (1 + r_t) \right)$$

$$\theta N_t^{\chi} = \frac{1}{C_t} w_t$$

$$\frac{1}{C_t} = \beta \mathbb{E}_t \left[\frac{1}{C_{t+1}} \frac{Z_t}{Z_{t+1}} \left(Z_{t+1} R_{t+1}^k + (1 - \delta) \right) \right]$$

The rest of the equilibrium conditions are standard:

$$Y_t = C_t + I_t$$

$$K_{t+1} = Z_t I_t + (1 - \delta) K_t$$

$$w_t = \alpha A_t K_t^{\alpha} N_t^{-\alpha}$$

$$R_t^k = (1 - \alpha) A_t K_t^{\alpha - 1} N_t^{1-\alpha}$$

$$Y_t = A_t K_t^{\alpha} N_t^{1-\alpha}$$

I assume that both A_t and Z_t follow mean zero AR(1) processes in the log:

$$\ln A_t = \rho_a \ln A_{t-1} + \varepsilon_{a,t}$$

$$\ln Z_t = \rho_z \ln Z_{t-1} + \varepsilon_{z,t}$$

I use a standard parameterization of the model. Here I assume that $\rho_z = 0.9$ and that the standard deviation of the investment shock is 0.01. The impulse responses to the investment shock are shown below.

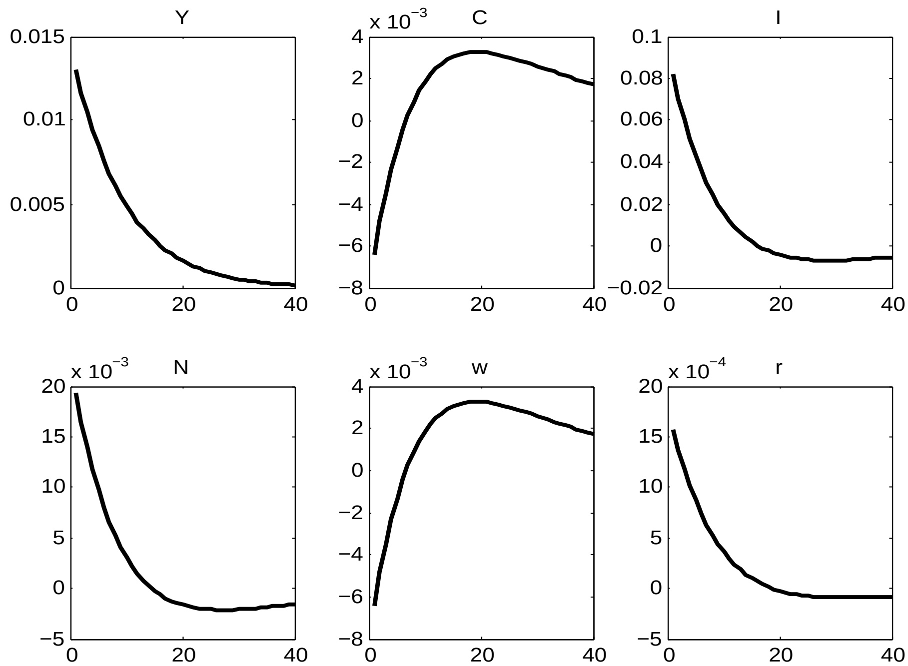

We observe that the investment shock leads to a sizeable increase in output, hours, and investment, with reductions in consumption and the real wage. The intuition for what is going on is as follows. When Z_t increases, since turning investment into capital is more efficient, it makes sense to save more through capital. Hence, consumption jumps down. This results in an inward shift of the labor supply curve (along an initially stable labor demand curve), which leads to an increase in N_t and a reduction in w_t. The increase in N_t results in an increase in Y_t, which combined with the reduction in C_t means I_t is higher. The real interest rate rises because the marginal product of capital is initially higher. As we go further out, we start to accumulate more capital, and the Z_t shock fades away, consumption beings to increase, which shifts labor supply back in, driving down N_t and the wage up. We do not generate co-movement between C_t and N_t (and Y_t and I_t) here for reasons that have been mentioned above: with our standard labor supply relationship, absent a change in A_t consumption and hours cannot move together (this holds the capital stock fixed, which is a safe approximation at short forecast horizons). Hence, while the investment shock produces interesting dynamics here, in the model as currently presented it cannot be a major source of business cycle fluctuations as it does not get the co-movement right. In difference setups where the labor supply condition is broken this is not necessarily the case.

A final caveat is in order here. The investment shock as presented is closely related to a different kind of shock that is often called "investment-specific technology" (or IST). This is a shock which affects the transformation of consumption goods into investment goods, whereas the investment shock laid out here impacts the transformation of investment goods into capital goods. Since investment is non-consumed output, in terms of the capital accumulation equation both the IST and MEI shocks show up the same way. But the two kinds of shocks have different implications for the relative price of investment and consumption. In our setup, the relative price of investment to consumption is 1; in the IST setup, the relative price of investment is the inverse of the IST shock. The relative price of investment to consumption has trend significantly down in the post-war era, suggesting that IST shocks are an important long run feature of the data. But the relative price of investment does not move a ton at cyclical frequencies, so mean-reverting shocks to IST cannot be super important business cycle shocks. Since the MEI shock here doesn't affect the relative price of investment, you can't rule out that it's important over the business cycle.

2.8 Money

We have abstracted from money thus far. A downside of doing so is that we can't say anything about price level or inflation determination.

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. Money is really liquid wealth – wealth in the sense that it serves as a store of value (i.e. you can save via holding money), but liquid in the sense that you can costless use it in exchange. 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), an efficient medium of exchange important.

In a single-good model, 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, in principle, pay interest. What differentiates money is that it does not pay interest. Hence, the "cards are stacked" against money. It *must* be an exchange motive that gets people to hold money. But in a model with one kind of good (e.g. fruit), it is difficult to formally model exchange frictions that would give rise to a natural demand for money. There is an active literature that thinks hard about this – the so-called "money search" literature. We're not going to study that here, but the basic gist is that you have to search for someone who has what you want (and wants you have) to trade. Money reduces this search cost, and can therefore be useful.

Putting aside a micro-founded model with search frictions, there are two shortcut ways to get agents to hold money. We will consider both. These shorcuts are 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 two approaches yield similar results but they are not exactly the same. We consider each in turn.

2.8.1 The Budget Constraint

Before specifying how to get agents to want to hold money, we need to write out a budget constraint that includes money. This is because money is a store of value. Let M_{t-1} denote the nominal holdings brought into period t – this is predetermined. Let M_t 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 (i.e. money serves as the unit of account). Let i_{t-1} denote the nominal interest rate on nominal bonds, B_{t+1}, observed at time t that pays off in time $t+1$. The bond pays off in dollars (i.e. units of money), which is why we call i_t a nominal (rather than real) interest rate.

In nominal terms, on the expenditure side, a household can consume (P_tC_t), invest in new physical capital (P_tI_t), accumulate new nominal bonds, $B_{t+1} - B_t$ (note that I am now thinking of these bonds as being measured in units of money, i.e. dollars, instead of units of goods as before), or accumulate new money, $M_t - M_{t-1}$. On the income side, the household earns nominal income from working, $P_tw_tN_t$ (where w_t is the real wage), nominal income from leasing capital, $P_tR_t^kK_t$, pays nominal taxes to the government, P_tT_t, receives profit from ownership in firms, P_tDIV_t, and earns nominal interest on the existing stock of bonds, $i_{t-1}B_t$. I'm going to switch to using DIV_t to denote real profit, because I'm going to want to use Π_t to measure gross inflation. Hence, the household's nominal flow budget constraint is:

$$P_tC_t + P_tI_t + B_{t+1} - B_t + M_t - M_{t-1} \leq P_tw_tN_t + P_tR_t^kK_t - P_tT_t + P_tDIV_t + i_{t-1}B_t \tag{25}$$

Note that money and bonds enter the expenditure side of the flow budget constraint in the same way – you enter with a stock, and can accumulate more (or draw down the stock). But they differ in that bonds pay interest ($i_{t-1}B_t$ appears on the income side), but money does not. To put the flow budget constraint back into real terms, as we have been doing, divide through by P_t:

$$C_t + K_{t+1} - (1-\delta)K_t + \frac{B_{t+1} - B_t}{P_t} + \frac{M_t - M_{t-1}}{P_t} \leq w_tN_t + R_t^kK_t - T_t + DIV_t + i_{t-1}\frac{B_t}{P_t} \tag{26}$$

Here, I have subbed out the accumulation equation for investment. Note also I'm not being entirely consistent with timing notations on state variables. K_t and B_t are the stocks of capital and bonds, respectively, chosen in $t-1$ that are available in period t. K_{t+1} and B_{t+1} get to be chosen in t and are available in $t+1$. For money, I'm using different timing notation – M_{t-1} is chosen in $t-1$ and brought into t¡ whereas M_t is chosen in t and brought into the future in $t+1$. This is sloppy and I'm probably being confusing. But deal with it.

Now, just looking at the budget constraint, you ought to be able to see why it is difficult to get agents to hold money. As written, money just enters the budget constraint as a store of value. Money and bonds are substitutable ways to transfer resources across time. If bonds pay interest (i.e. if $i_t > 0$, money is dominated as a store of value. Keep that in mind.

The firm problem will be the same with or without money. The representative firm just hires capital and labor to maximize profit. There is nothing dynamic about the problem, so the FOC are the same in terms of real variables. In particular, the real dividend is:

$$DIV_t = A_t F(K_t, N_t) - w_t N_t - R_t^k K_t$$

The optimal factor demand conditions are standard:

$$w_t = (1-\alpha) A_t K_t^{\alpha} N_t^{-\alpha} \tag{27}$$

$$R_t^k = \alpha A_t K_t^{\alpha-1} N_t^{1-\alpha} \tag{28}$$

The government "produces" an exogenous amount of money, M_t. There is no cost to producing money. Assume that the supply of money is set according to an AR(1) process in the growth rate:

$$g_t^M = \ln M_t - \ln M_{t-1} \tag{29}$$

$$g_t^M = (1-\rho_M) g^M + \rho_M g_{t-1}^M + s_M \varepsilon_{M,t} \tag{30}$$

The steady state growth rate is g^M. Assume that the government does no spending, so $G_t = 0$. In nominal terms, the government budget constraint is:

$$i_{t-1} D_t \leq P_t T_t + M_t - M_{t-1} + D_{t+1} - D_t \tag{31}$$

Government expenditure is just interest payments on the existing stock of debt. The *change* (or flow) in the supply of money is a revenue source for the government, $M_t - M_{t-1}$, in a way similar to how issuing debt is a revenue source. In real terms, this budget constraint is:

$$i_{t-1} \frac{D_t}{P_t} \leq T_t + \frac{M_t - M_{t-1}}{P_t} + \frac{D_{t+1} - D_t}{P_t} \tag{32}$$

Note that we are modeling one, consolidated government budget constraint here. We could separately model a central bank (who chooses M_t) and a fiscal authority, where the fiscal authority receives a transfer from the central bank (seigniorage revenue). We are abstracting from that.

Bond market clearing requires that $D_t = B_t$ each period. Go to the household's budget constraint at equality and substitute in the expression for real dividends. We get:

$$C_t + I_t + \frac{B_{t+1} - B_t}{P_t} + \frac{M_t - M_{t-1}}{P_t} \leq w_t N_t + R_t^k K_t - T_t + Y_t - w_t N_t - R_t^k K_t + i_{t-1} \frac{B_t}{P_t}$$

Which reduces to:

$$C_t + I_t + \frac{B_{t+1} - B_t}{P_t} + \frac{M_t - M_{t-1}}{P_t} \leq Y_t - T_t + i_{t-1} \frac{B_t}{P_t}$$

Now plug in for T_t using the government's budget constraint.

$$C_t + I_t + \frac{B_{t+1} - B_t}{P_t} + \frac{M_t - M_{t-1}}{P_t} \leq Y_t - \left(i_{t-1}\frac{D_t}{P_t} - \frac{M_t - M_{t-1}}{P_t} - \frac{D_{t+1} - D_t}{P_t} \right) + i_{t-1}\frac{B_t}{P_t}$$

All this other stuff cancels, leaving the normal resource constraint:

$$Y_t = C_t + I_t \tag{33}$$

Now, we have to do something to get households to willingly hold money. We'll consider two alternatives below:

2.8.2 Money in the Utility Function

In this specification households get utility from consumption, leisure, and holding real money balances – M_t/P_t. Note the timing convention here – M_t 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:

$$\max_{C_t, N_t, K_{t+1}, B_{t+1}, M_t} \quad E_0 \sum_{t=0}^{\infty} \beta^t \left\{ \ln C_t + \theta \ln(1 - N_t) + \psi \frac{\left(\frac{M_t}{P_t}\right)^{1-\zeta} - 1}{1 - \zeta} \right\}$$

s.t.

$$C_t + K_{t+1} - (1-\delta)K_t + \left(\frac{B_{t+1} - B_t}{P_t} \right) + \left(\frac{M_t - M_{t-1}}{P_t} \right) \leq w_t N_t + R_t^k K_t - T_t + \Pi_t + i_{t-1}\frac{B_t}{P_t}$$

ψ is a scaling parameter, and ζ will be an elasticity of money demand parameter.

$$\mathbb{L} = E_0 \sum_{t=0}^{\infty} \beta^t \left\{ \ln C_t + \theta \ln(1 - N_t) + \psi \frac{\left(\frac{M_t}{P_t}\right)^{1-\zeta} - 1}{1 - \zeta} + \dots \right.$$

$$\left. \dots + \lambda_t \left(w_t N_t + R_t^k K_t - T_t + \Pi_t + (1 + i_t)\frac{B_t}{P_t} - C_t - K_{t+1} + (1-\delta)K_t - \frac{M_t}{P_t} + \frac{M_{t-1}}{P_t} - \frac{B_{t+1}}{P_t} \right) \right\}$$

The FOC are:

$$\frac{\partial \mathbb{L}}{\partial C_t} = 0 \iff \frac{1}{C_t} = \lambda_t$$
$$\frac{\partial \mathbb{L}}{\partial N_t} = 0 \iff \theta \frac{1}{1 - N_t} = \lambda_t w_t$$
$$\frac{\partial \mathbb{L}}{\partial K_{t+1}} = 0 \iff \lambda_t = \mathbb{E}_t \lambda_{t+1}(R^k_{t+1} + (1-\delta))$$
$$\frac{\partial \mathbb{L}}{\partial B_{t+1}} = 0 \iff \frac{\lambda_t}{P_t} = \mathbb{E}_t \frac{\lambda_{t+1}}{P_{t+1}}(1 + i_t)$$
$$\frac{\partial \mathbb{L}}{\partial M_t} = 0 \iff \psi \left(\frac{M_t}{P_t}\right)^{-\zeta} \frac{1}{P_t} = \frac{\lambda_t}{P_t} - \beta \mathbb{E}_t \frac{\lambda_{t+1}}{P_{t+1}}$$

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

$$\frac{\theta}{1 - N_t} = \frac{1}{C_t} w_t \tag{34}$$
$$\frac{1}{C_t} = \beta \mathbb{E}_t \left(\frac{1}{C_{t+1}} (R_{t+1} + (1-\delta)) \right) \tag{35}$$
$$\frac{1}{C_t} = \beta \mathbb{E}_t \left(\frac{1}{C_{t+1}} \left((1 + i_t) \frac{P_t}{P_{t+1}} \right) \right) \tag{36}$$

We could define *gross* inflation as $\Pi_t = \frac{P_t}{P_{t-1}} = 1 + \pi_t$, where π_t is the net inflation rate. We could then write the Euler equation for bonds as:

$$\frac{1}{C_t} = \beta \mathbb{E}_t \left(\frac{1}{C_{t+1}} \left((1 + i_t) \Pi^{-1}_{t+1} \right) \right) \tag{37}$$

The *Fisher relationship* relates the real interest rate to the nominal rate and the expected rate of price inflation as follows:

$$1 + r_t = (1 + i_t) \mathbb{E}_t \Pi^{-1}_{t+1} \tag{38}$$

Why does this hold as an identity? The real rate, $1 + r_t$, tells you how many *goods* you get back in the future by foregoing one good today. The nominal interest rate tells you how many *dollars* you get in the future by foregoing one dollar today. If you save one dollar, this is equivalent to saving P_t goods today. This generates $(1 + i_t)P_t$ dollars tomorrow, which will purchase $(1 + i_t)P_t/P_{t+1}$ goods tomorrow.

We can re-arrange the FOC for money as follows:

$$\frac{1}{P_t} = \psi \left(\frac{M_t}{P_t}\right)^{-\zeta} \frac{1}{P_t} \lambda_t^{-1} + \mathbb{E}_t \beta \frac{\lambda_{t+1}}{\lambda_t} \frac{1}{P_{t+1}}$$

Now we can sub out for λ_t:

$$\frac{1}{P_t} = \psi \left(\frac{M_t}{P_t} \right)^{-\zeta} \frac{C_t}{P_t} + \mathbb{E}_t \beta \frac{C_t}{C_{t+1}} \frac{1}{P_{t+1}} \tag{39}$$

This looks a little funny, but it is just an asset pricing condition. If P_t is the price of goods in terms of money, then $1/P_t$ is the price of money in terms of goods (all prices are relative). So the left hand side is the "price" of money. What is on the right hand side? The first term is the marginal rate of substitution between money and goods – $\psi \left(\frac{M_t}{P_t} \right)^{-\zeta} \frac{C_t}{P_t}$ is the marginal utility of holding money. Dividing this by the marginal utility of consumption (which with log utility is multiplying by consumption) gives this term an MRS interpretation. The second term on the RHS is the discounted continuation value. The "price" of the asset money in $t+1$ is $1/P_{t+1}$. This is discounted by the household's stochastic discount factor, $\beta \frac{C_t}{C_{t+1}}$.

There is another way to write this asset pricing condition that is more useful in thinking about a money demand curve. In particular, we can write it:

$$\psi m_t^{-\zeta} = \lambda_t + \mathbb{E}_t \beta \lambda_{t+1} \frac{P_t}{P_{t+1}}$$

Where $m_t \equiv M_t/P_t$ is defined as real money balances. Now, from the Euler equation for bonds, we know that $\mathbb{E}_t \beta \lambda_{t+1} P_t / P_{t+1} = \lambda_t/(1+i_t)$. So we can write this as:

$$\psi m_t^{-\zeta} = \lambda_t \left(1 - \frac{1}{1+i_t} \right)$$

Or:

$$\psi m_t^{-\zeta} = \frac{1}{C_t} \left(\frac{i_t}{1+i_t} \right)$$

Or:

$$m_t = \psi^{\frac{1}{\zeta}} C_t^{\frac{1}{\zeta}} \left(\frac{1+i_t}{i_t} \right)^{\frac{1}{\zeta}} \tag{40}$$

This tells us that demand for real money balances is (i) increasing in ψ (the preference for holding money), (ii) increasing in C_t (how many transactions you are doing), and (iii) decreasing in i_t (which is the opportunity cost of holding money – if you are holding money, you aren't holding bonds and are foregoing the nominal interest you could earn on bonds).

With this money demand condition (note that one could use the asset pricing formulation above instead), the entire set of equilibrium conditions for the model is then:

$$\frac{\theta}{1-N_t} = \frac{1}{C_t} w_t \tag{41}$$

$$\frac{1}{C_t} = \beta \mathbb{E}_t \left(\frac{1}{C_{t+1}} (R^k_{t+1} + (1-\delta)) \right) \tag{42}$$

$$\frac{1}{C_t} = \beta \mathbb{E}_t \left(\frac{1}{C_{t+1}} (1+r_t) \right) \tag{43}$$

$$R^k_t = \alpha A_t K_t^{\alpha-1} N_t^{1-\alpha} \tag{44}$$

$$w_t = (1-\alpha) A_t K_t^{\alpha} N_t^{-\alpha} \tag{45}$$

$$Y_t = A_t K_t^{\alpha} N_t^{\alpha} \tag{46}$$

$$K_{t+1} = I_t + (1-\delta) K_t \tag{47}$$

$$Y_t = C_t + I_t \tag{48}$$

$$\ln A_t = \rho_A \ln A_{t-1} + s_A \varepsilon_{A,t} \tag{49}$$

$$g^M_t = (1-\rho_M) g^M + \rho_M g^M_{t-1} + s_M \varepsilon_{M,t} \tag{50}$$

$$m_t = \psi^{\frac{1}{\zeta}} C_t^{\frac{1}{\zeta}} \left(\frac{1+i_t}{i_t} \right)^{\frac{1}{\zeta}} \tag{51}$$

$$1 + r_t = (1+i_t) E_t \Pi^{-1}_{t+1} \tag{52}$$

$$g^M_t = \ln m_t - \ln m_{t-1} + \ln \Pi_t \tag{53}$$

The last condition is a stationary transformation. Formally:

$$g^M_t = \ln M_t - \ln M_{t-1}$$

M_t will be growing, so we want to write this in terms of real balances, $m_t = M_t/P_t$, or $\ln m_t = \ln M_t - \ln P_t$. So we can manipulate to get:

$$g^M_t = (\ln M_t - \ln P_t) + \ln P_t - (\ln M_{t-1} - \ln P_{t-1}) - \ln P_{t-1}$$

Which is:

$$g^M_t = \ln m_t - \ln m_{t-1} + \ln \Pi_t$$

Since $\Pi_t = P_t/P_{t-1}$, and $\ln \Pi_t = \ln P_t - \ln P_{t-1}$. Since $\Pi_t = 1 + \pi_t$ is gross inflation (π_t is net), we will often use the approximation that log of one plus a small number is the small number (so $\ln \Pi_t = \pi_t$).

Looking at the system of equations above, there 13 equations and 13 variables (C_t, N_t, K_t, Y_t, I_t, R^k_t, w_t, r_t, A_t, g^M_t, m_t, i_t, Π_t). But take note of something. The first nine of these equations, and the first nine variables, are *exactly* the same as in a RBC model without money and nominal prices. Hence, we can solve for the real side of the model *independently* of the nominal side. The *classical dichotomy* holds here: we can solve for endogenous values of real variables independently of nominal variables (the converse is not true; we need to know real variables to solve for nominal variables).

The classical dichotomy holding means that money will be *neutral*: changes in the money supply will not have any effect on real variables like output, just effects on nominal prices. Technically, it's even more than that: money is *super-neutral*. Neutrality of money means that a change in the *level* of the money supply has no effects on real variables like output. Super-neutrality means that a change in the *growth rate* of the money supply has no effect on real variables either. In the MIU model, both properties (one stronger than the other) hold.

We can solve for the steady state just as we did before. In terms of the new equations, in steady state we will have $m_t = m_{t-1} = m$, which means that $g^M = \pi$ (i.e. the steady state net inflation rate equals the growth rate of the money supply). 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 $g^M = \pi = 0.00$, so that there is no inflation in the steady state. I set $\psi = 1$ and $\zeta = 2$. The steady state growth rate of money, equal to the steady state inflation rate, does not have interesting short run dynamic effects here

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 as in the baseline RBC model 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.

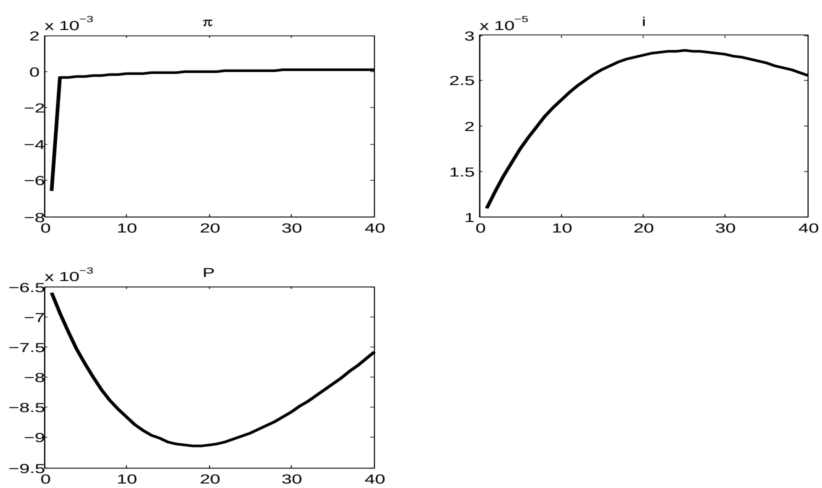

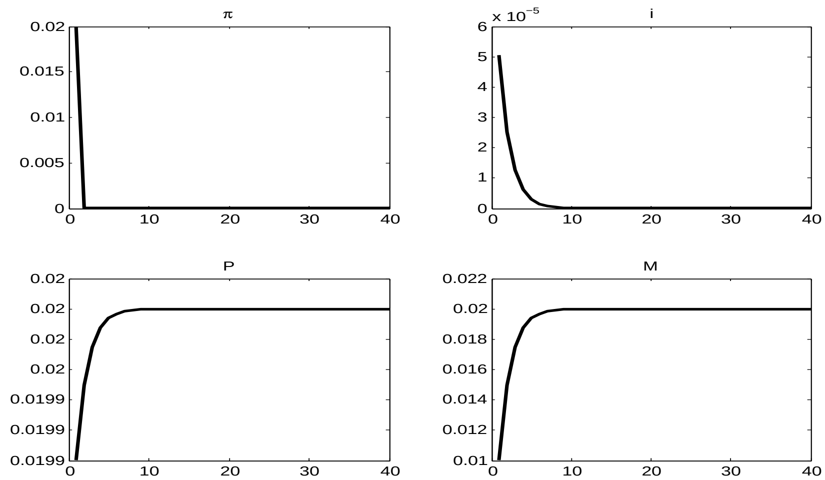

These have features we would more or less expect – inflation (and hence the price level) fall in response to a technology shock and rise in response to a monetary shock. This is consistent with the empirical fact documented earlier that the price level is countercyclical. The nominal interest rate rises when the money supply increases at an unexpectedly fast rate; the nominal interest rate also rises after the technology shock.

It's an interesting exercise to see what happens if I make $\rho_M = 0$, which means that the money supply follows a random walk (with drift, if $\pi \neq 0$). In response to a monetary shock, the only effect is for the price level to immediately jump up by the amount of the change in M – there is no change in the nominal interest rate. The logic for this is as follows. Since nothing real changes, the real interest rate, r_t, will not respond to the monetary shock. But since there is no persistence to the shock, the price level will just jump up by the increase in M_t, and since there is no further change in M_t in period $t+1$ or beyond, nothing more will happen to the price level after period t. This means that expected inflation between t and $t+1$ will not react – e.g. π_{t+1} will not change. But if π_{t+1} doesn't change, and r_t doesn't change, then from the Fisher relationship i_t doesn't change either. In response to a monetary shock, i_t will simply move with $\mathbb{E}_t\pi_{t+1}$ in such a way as to keep r_t unchanged. We can see this in the impulse responses shown below:

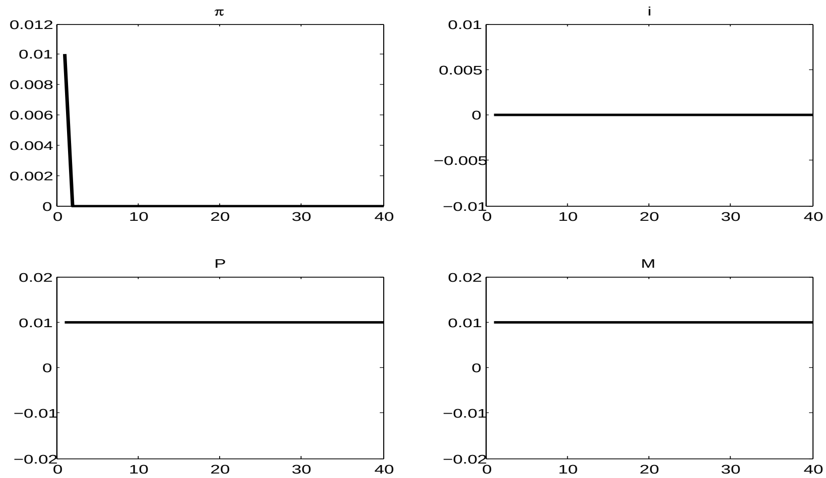

2.8.3 Cash in Advance

Another way to get money into the model is to assume a cash in advance constraint. This ends up being pretty similar, though not exactly the same, as money in the utility function. 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:

$$M_{t-1} \geq P_t C_t \tag{54}$$

Money does not appear in the utility function, but the above serves as an extra constraint. Households have to have money on hand to finance consumption. As a store of value, money enters the budget constraint in the same way. We can write out the household problem as:

$$\max_{C_t, N_t, K_{t+1}, B_{t+1}, M_t} \quad \mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \left\{ \ln C_t + \theta \ln(1 - N_t) \right\}$$

$$\text{s.t.}$$

$$C_t + K_{t+1} - (1-\delta)K_t + \left(\frac{B_{t+1} - B_t}{P_t} \right) + \left(\frac{M_t - M_{t-1}}{P_t} \right) = w_t N_t + R_t^k K_t - T_t + \Pi_t + i_{t-1} \frac{B_t}{P_t}$$
$$M_{t-1} \geq p_t C_t$$

We can form a current value Lagrangian with two constraints: the budget constraint (expenditure cannot exceed income) and the cash in advance constraint (money on hand must be used to finance consumption):

$$\mathbb{L} = \mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \left\{ \ln C_t + \theta \ln(1 - N_t) + \mu_t \left(\frac{M_{t-1}}{P_t} - C_t \right) \dots \right.$$

$$\left. \dots + \lambda_t \left(w_t N_t + R_t^k K_t - T_t + \Pi_t + (1 + i_{t-1}) \frac{B_t}{P_t} - C_t - K_{t+1} + (1 - \delta) K_t - \frac{B_{t+1}}{P_t} - \frac{M_t}{P_t} + \frac{M_{t-1}}{P_t} \right) \right.$$

$$\frac{\partial \mathbb{L}}{\partial C_t} = 0 \Leftrightarrow \frac{1}{C_t} = \lambda_t + \mu_t \tag{55}$$

$$\frac{\partial \mathbb{L}}{\partial N_t} = 0 \Leftrightarrow \frac{\theta}{1 - N_t} = \lambda_t w_t \tag{56}$$

$$\frac{\partial \mathbb{L}}{\partial K_{t+1}} = 0 \Leftrightarrow \lambda_t = \beta \mathbb{E}_t \left(\lambda_{t+1} (R_{t+1}^k + (1 - \delta) \right) \tag{57}$$

$$\frac{\partial \mathbb{L}}{\partial B_{t+1}} = 0 \Leftrightarrow \lambda_t = \beta \mathbb{E}_t \lambda_{t+1} \left((1 + i_t) \left(\frac{P_t}{P_{t+1}} \right) \right) \tag{58}$$

$$\frac{\partial \mathbb{L}}{\partial M_t} = 0 \Leftrightarrow -\frac{\lambda_t}{P_t} + \beta \mathbb{E}_t \frac{\mu_{t+1}}{P_{t+1}} + \beta \mathbb{E}_t \frac{\lambda_{t+1}}{P_{t+1}} = 0 \tag{59}$$

The first-order conditions for labor, capital, and bonds like the same as in the basic model (written in terms of λ_t, anyway). The first first-order condition looks difference. If the cash in advance constraint is binding, then $\mu_t > 0$, and C_t will be lower than it would in the absence of the constraint. The final first-order condition can be re-written in an asset pricing form as follows:

$$\frac{1}{P_t} = \mathbb{E}_t \left[\frac{\beta \lambda_t}{\lambda_{t+1}} \left(\frac{\mu_{t+1}}{P_{t+1} \lambda_{t+1}} + \frac{1}{P_{t+1}} \right) \right]$$

As we discussed above, $\frac{1}{P_t}$ is the price of money (measured in terms of goods) in period t. $\frac{\beta \lambda_t}{\lambda_{t+1}}$ is the stochastic discount factor. μ_{t+1}/P_{t+1} is the marginal utility of having extra cash in the next period. λ_{t+1} is the marginal utility of having extra goods. So $\mu_{t+1}/(P_{t+1}\lambda_{t+1}$ is the marginal rate of substitution between money and goods – this is like the flow benefit to holding money. $1/P_{t+1}$ is the price of money in $t + 1$.

If the cash in advance constraint never binds, then $\mu_t = 0$. Then $\lambda_t = 1/C_t$ and the first order conditions (with the exception of the last) would be identical to the basic RBC model (and the MIU model). If the cash in advance constraint never binds, the last first order condition (for money) would be:

$$\lambda_t = \beta \mathbb{E}_t \lambda_{t+1} \frac{P_t}{P_{t+1}}$$

The only way that both this and the bond Euler equation can hold is if $i_t = 0$. In other words, if $i_t > 0$ then $\mu_t > 0$ – i.e. the cash in advance constraint will bind if the nominal interest rate is positive. This makes sense. If the nominal interest rate were zero, there is no cost to holding cash (instead of bonds), so people will accumulate sufficient cash to make the cash in advance constraint not bind. But if $i_t > 0$, there is a cost to holding money – foregone interest on bonds. So people will economize on money holdings and the constraint will bind.

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 of real variables 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:

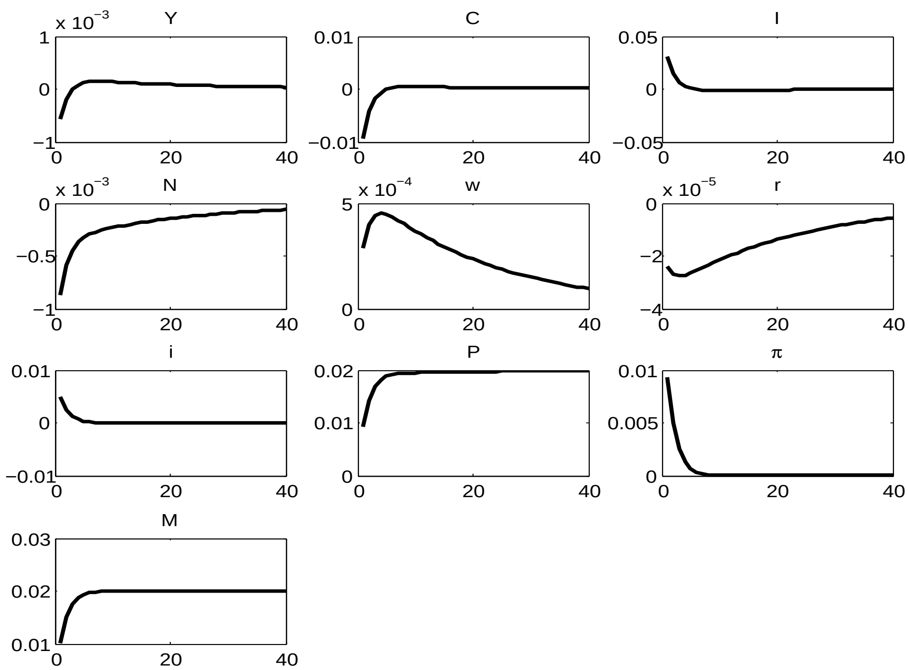

Here we see something perhaps not very intuitive. Not only does the monetary shock have real effects, it actually causes an output contraction (albeit it is very small). What is the intuition for this? Inflation is essentially a tax on holders on money. In the absence of the technological constraint 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 because it's essentially a tax on 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 pretty small (in comparison with the responses to a productivity shock) and are not very persistent. In particular, for the parameterization I used, money explains less than 1 percent of the variance of output, about 6 percent of the variance of investment, about 1 percent of the variance of horus, and about 4 percent of the variance of consumption. 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).

2.8.4 Optimal Long Run Inflation

I've thus far ignored the optimal trend growth rate of the money supply (which is in turn equal to the long run steady state growth rate of prices), simply taking it is given. What would be the optimal long run inflation rate in either of these models?

The so-called Friedman rule is to set the nominal interest rate on bonds equal to zero $i_t = 0$. 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.

At a more formal level, we can see why this is optimal in both of these specifications. It is perhaps easiest to see in the CIA model. What you would like is for the cash in advance constraint to not bind – if the constraint doesn't bing, agents have to be weakly better off than if it does bind. The constraint not binding would mean that $\mu_t = 0$. The only way for the first order conditions to all hold with $\mu_t = 0$ is if $i_t = 0$. Hence, setting $i = 0$ in the long run is optimal, as Friedman conjectured. From the Fisher relationship, since $r = \frac{1}{\beta} - 1$, $i = 0$ requires that $1 + \pi = (1/\beta - 1)^{-1}$, which means $\pi < 0$. Hence, you want deflation in the steady state.

Setting money growth such that the long-run level of the nominal interest rate is also optimal in the MIU model, even though money is completely neutral. As $i_t \to 0$, we see that real money balances go to ∞. Since the household gets utility from real money balances, bigger real money balances mean more utility. $i = 0$ again necessitates steady state deflation, so the price level is going to zero (which is how real money balances go to ∞.