Technical Appendix to "Fiscal Policy with Heterogeneous Agents and Incomplete Markets"

Jonathan Heathcote
Georgetown University

November 26th 2003

1. *The recursive version of the household's problem*

Recall that in the revised household problem, the aggregate state of the economy is summarized by the vector $Z = (K, B, \tau)$. The notation is as follows: K denotes the aggregate capital in place at the start of the period, B is the quantity of government debt that matures in the current period, and τ is the current period tax rate. A household's individual state is given by $s = (e, a)$, where e is the household-specific productivity shock that arrives at the start of the period, and a is the stock of assets carried over from the previous period.

With the problem described recursively, households (in the benchmark incomplete-markets model) solve

$$V(s,Z) = \max_{c,a',n} \left\{u(c,n) + \beta E\left[V\left(s', Z'\right)|s, Z\right]\right\}$$

subject to

$$c + a' = \left[1 + r(Z)(1-\tau(Z))\right]a(s) + (1-\tau(Z))w(Z)e(s)n$$

$$c \geq 0,\ \ a' \in A,\ \ n \in [0,1]$$

taking as given:

1. the Markov processes Π and π_τ
2. the standard functions for factor prices,

$$r(Z) = \alpha K(Z)^{\alpha-1} N(Z)^{1-\alpha} - \delta \tag{1.1}$$

$$w(Z) = (1-\alpha)K(Z)^{\alpha}N(Z)^{-\alpha} \tag{1.2}$$

3. an aggregate decision rule for labor (see the derivation for this expression in the next sub-section)

$$N(Z) = \left(\sum_{i=1}^{l} p_i^* e_i^{1+\varepsilon}\left[\frac{(1-\alpha)K(Z)^{\alpha}(1-\tau(Z))}{\psi}\right]^{\varepsilon}\right)^{\frac{1}{1+\alpha\varepsilon}} \tag{1.3}$$

4. the law of motion for government debt,

$$B(Z') = (1+r(Z))B(Z) + G - \tau(Z)\left[r(Z)\left(K(Z)+B(Z)\right) + w(Z)N(Z)\right] \tag{1.4}$$

5. a law of motion for aggregate capital of the form

$$\ln\left(K(Z')\right) = \alpha_0 + \alpha_1 \ln\left(K(Z)\right) + \alpha_2 \ln\left(B(Z)\right) + \alpha_3 \ln\left(\tau(Z)\right). \tag{1.5}$$

α_i, $i = 0, 1, 2, 3 \in \mathrm{R}$.

1.1. *Derivation of aggregate labor supply expression*

$$\begin{aligned} N(Z) &= \int_{A\times E}\sum_{e^t\in E^t}\mu^t(e^t)e_t(e^t)n(s,Z)d\lambda \\ &= \int_{A\times E}\sum_{e^t\in E^t}\mu^t(e^t)e_t(e^t)\left[\frac{w(Z)e_t(e^t)(1-\tau(Z))}{\psi}\right]^{\varepsilon} d\lambda \\ &= \sum_{i=1}^{l} p_i^* e_i\left[\frac{(1-\alpha)K(Z)^{\alpha}N(Z)^{-\alpha}e_i(1-\tau(Z))}{\psi}\right]^{\varepsilon} \\ &= \left(\sum_{i=1}^{l} p_i^* e_i^{1+\varepsilon}\left[\frac{(1-\alpha)K(Z)^{\alpha}(1-\tau(Z))}{\psi}\right]^{\varepsilon}\right)^{\frac{1}{1+\alpha\varepsilon}} \end{aligned}$$

2. *Solving the household's problem*

The state space for this problem is large. Moreover, given that markets are incomplete and agents face a borrowing constraint that may or may not be binding, it is clear that a global solution method is appropriate. I solve the consumer's problem using the finite element method as outlined by McGrattan (1999). The method involves searching for piecewise linear decision rules that imply small Euler equation residuals across the state space. The grid I use has 30 points in the individual assets dimension, and 3 points in both the aggregate capital and aggregate debt dimensions.[1] Given 3 possible values for household productivity, and 2 values for the tax rate, the total number of grid points is 1620.

The first order conditions to the household's problem are:

$$\psi n(s,Z)^{\frac{1}{\varepsilon}} = w(Z)e(s)(1-\tau(Z)) \tag{2.1}$$

and

$$\begin{aligned} u_c(c(s,Z),n(s,Z)) &\geq \beta E\left[u_c(c(s',Z'),n(s',Z'))\left(1+r(Z')(1-\tau(Z'))\right)\right] \\ &= \text{if } c(s,Z) < a(s) + (1-\tau(Z))\left[r(Z)a(s) + w(Z)e(s)n(s,Z)\right]. \end{aligned}$$

Given a guess for the consumption decision rule, $\widehat{c}$, the residual for the household's intertemporal Euler equation may be evaluated as follows. Equations 1.1, 1.2 and 1.3 may be used to compute values for $r(Z)$, $w(Z)$ and $n(s,Z)$. Next period household wealth, $a(s') = a'(s,Z)$ follows from the household budget constraint, given $\widehat{c}(s,Z)$. The probability distribution over next period productivity, $e(s')$, is given by the matrix Π. Next period government debt, $B(Z')$, is given by 1.4. Next period capital, $K(Z')$ is given by 1.5. The probability distribtion over next period tax rates is given by the function $\pi_\tau\left((\tau(Z),B(Z')),\tau(Z')\right)$. For each possible value for $s' = (e(s'),a(s'))$ and $Z' = (K(Z'),B(Z'),\tau(Z'))$, equations 1.1, 1.2 and 1.3 may be used once more to compute values for $r(Z')$, $W(Z')$ and $n(s',Z')$. For each possible (s',Z'), next period consumption is given by $\widehat{c}(s',Z')$.

There are various possible ways to deal with the no-borrowing constraint. In an earlier version of the paper, I iterated on a function describing the expected marginal utility of consumption (rather than the consumption function itself) in an approach reminiscent of the parameterized expectations algorithm described by Christiano and Fischer (2000). In this version of the paper I follow the penalty function approach advocated by McGrattan, which involves imposing a utility penalty for violating the non-negative asset holding constraint instead of thinking of the constraint as limiting the feasible choice set. Once the problem has been solved given a large

[1] I experimented with more grid points in the aggregate capital and debt dimensions, but found that this had virtually no effect on the results. The reason is that household decision rules are very close to linear in aggregate capital and aggregate debt. Of course, the presence of the no-borrowing constraint means that decision rules are highly non-linear in household wealth. Thus it is important to have more grid points in the household wealth dimension, and to space grid points unevenly so that the grid is especially fine at low levels of wealth.

value for the penalty, I record the grid points at which the penalty is being imposed, and then resolve the household's problem imposing additional boundary conditions at these points.

3. *The economy without aggregate risk*

Prior to solving the model with aggregate tax shocks, I first solve for the stationary distribution of wealth in an economy with no aggregate uncertainty. In this economy agents assume that aggregate variables and factor prices will be constant through time. There are three reasons for first considering this economy.

1. Since the state space is much smaller (there are no aggregate states) it is simple to solve for optimal household decision rules. These decision rules are then used to construct the first guess for decision rules in the full-blown economy with aggregate risk.

2. In the economy without aggregate risk, there is a unique stationary asset holding distribution. Since the state space is small (relative to the benchmark economy) this distribution may be rapidly computed given a certain parameterization. Thus I use the no-aggregate-risk economy to begin the search for productivity process parameters that generate an equilibrium wealth distribution resembling that in the United States (see the parameterization section).

3. In the economy without aggregate risk, aggregate capital, debt and tax revenue are constant through time. Thus these values provide a first guess for $\overline{K}$, $\overline{B}$ and $\overline{\tau}$, parameters which appear in the law of motion for capital (see the next section).

To ensure that in the economy *without* aggregate risk agents face a similar overall amount of risk as agents in the benchmark economy *with* aggregate tax shocks, I compute the steady state allocations for an economy in which households face no aggregate tax shocks, but where they do face additional idiosyncratic income shocks with the same magnitude and persistence as the tax shocks in the benchmark economy. Thus in each period, half of households pay taxes at the low rate, while the remainder are taxed at the high rate. I find that the constant values for aggregate capital, government debt and tax revenue in this economy are very similar to their average values over a long simulation of the benchmark economy. Furthermore, statistics characterizing the stationary distribution of wealth are very similar to their average values over a long simulation of the benchmark economy.

4. *Keeping track of the joint distribution over wealth and productivity*

To simulate either the benchmark model economy or the economy with no aggregate shocks, I create a fine grid on asset holdings, and represent the joint distribution over asset holdings and productivity (and taxes in the no-aggregate-risk economy) as a collection of piecewise linear cumulative distribution functions, one for each value for productivity (see Rios-Rull 1999). Given a grid $\{a_1, ..., a_m, ..., a_M\}$, and a collection of distrubution functions in the current period given by $\{F_i(\cdot)\}_{i=1}^{l}$, the value of the distribution in the following period at $e(s') = e_j$, $a(s') = a_m$ is given by

$$F_j'(a_m) = \sum_{i=1}^{l} F_i\left(a^{-1}((e_i, a_m), Z)\right)\Pi_{i,j} \tag{4.1}$$

where $a^{-1}((e_i, a_m), Z)$ denotes the value for assets a^* such that $a'((e_i, a^*), Z) = a_m$. Updating update the entire joint distribution one period is simply a matter of looping over j and m. I use $2,000$ unevenly-spaced points in the grid on asset holdings. Thus this joint distribution is represented by $6,000$ numbers in the benchmark economy, and $12,000$ numbers in the no-aggregate-risk economy. Cumulative distribution functions describing the average simulation distribution of asset holdings across the entire population and conditional distributions given particular values for household productivity are shown in figure 1 at the end of this appendix.

5. *The law of motion for aggregate capital*

Given a method for computing decision rules $c(\cdot,\cdot)$, I implement the Krusell and Smith (1998) method for deriving a forecasting rule for aggregate capital as follows:

1. Use the economy without aggregate shocks to get estimates for $\overline{K}$, $\overline{B}$ and $\overline{\tau}$, the simulation averages for capital, debt and the tax rate.

2. Specify an initial vector of α parameters. I set $\alpha_1 = 1$, $\alpha_2 = \alpha_3 = 0$, and $\alpha_0 = (1-\alpha_1)\ln(\overline{K}) - \alpha_2\ln(\overline{B}) - \alpha_3\ln(\overline{\tau})$. Set an initial joint distribution over wealth and

productivity $\left\{F_i^0\left(\cdot\right)\right\}_{i=1}^l$ (e.g. the unique stationary distribution for the economy with no aggregate shocks). This joint distribution should have the property that

$$\sum_{i=1}^{l}\left[\sum_{m=2}^{M}\left(F_i^0\left(a_m\right)-F_i^0(a_{m-1})\right)\frac{a_m-a_{m-1}}{2}+F_i^0(a_1)a_1\right]=\overline{K}+\overline{B}. \tag{5.1}$$

3. Set $Z_0=(K_0,B_0,\tau_0)=\left(\overline{K},\overline{B},\tau_l\right)$

4. Compute household decision rules across the state space (see above).

5. Update the joint distribution over wealth and productivity using eq. 4.1 to generate $\left\{F_i^1\left(\cdot\right)\right\}_{i=1}^l$.

6. Apply eq. 5.1 to $\left\{F_i^1\left(\cdot\right)\right\}_{i=1}^l$ to compute A_1, aggregate wealth carried into period 1. Use the government budget constraint to compute B_1, and set $K_1=A_1-B_1$. Use the transition function for taxes to draw a value for τ_1 given τ_0 and B_1. Create $Z_1=(K_1,B_1,\tau_1)$.

7. Repeat this procedure for a large number of periods to generate time series for Z (I simulate for $11,000$ periods, and discard the first $1,000$ observations).

8. Use this time series to update the guesses for $(\alpha_1,\alpha_2,\alpha_3)$ vector by running an ordinary least squares regression of the form given in eq. 1.5. Set $\alpha_0=(1-\alpha_1)\ln(\overline{K})-\alpha_2\ln(\overline{B})-\alpha_3\ln(\overline{\tau})$, where $\overline{K}$, $\overline{B}$ and $\overline{\tau}$ are the initial values (see step 1).

9. Repeat the entire process of solving the households' problems, simulating the economy and updating parameter values until the parameters $(\alpha_1,\alpha_2,\alpha_3)$ defining the forecasting rule for aggregate capital converge.

10. At this point, compute $\overline{K}$, $\overline{B}$ and $\overline{\tau}$ over a long simulation. If these differ from the initial guessed values, update a_0, and repeat the previous steps to recompute the remaining elements of α.

I find that the two-step procedure described above in which I first look for convergence in $(\alpha_1,\alpha_2,\alpha_3)$ prior to updating the terms $\overline{K}$, $\overline{B}$ and $\overline{\tau}$ which define α_0 is much more stable in this environment than the alternative of simply letting the regression define the new intercept coefficient α_0 (as in Krusell and Smith). Note that at the end of entire procedure, the converged

α vector is such that the forecasting rule households use when solving their problems is such that their behavior generates a law of motion for capital for which the best (log) linear predictor is the same forecasting rule.[2]

For the calibrated benchmark model described in the paper, the converged value for α_1 is close to but less than 1. The converged value for α_2 is negative, indicating that *ceteris paribus* more debt implies less investment and consequently a smaller next period capital stock. The converged value for α_3 is also negative, indicating that investment tends to be lower when the tax level is high. The coefficients for the benchmark model, the incomplete-markets economy with lump-sum taxes, and the economy with stochastic aging are reported in table 1 at the end of this appendix.

6. *The size of forecasting errors*

Figure 2 at the end of these notes contains histograms showing the 1 year and 10 year forecast errors for aggregate capital and for the real interest rate across a $10,000$ period simulation of the benchmark economy.[3] The cumulative forecasting error for capital 10 years hence is rarely more than a tenth of a percentage point of the capital stock. Forecasting errors for the real interest rate (the marginal product of capital minus the depreciation rate) are slightly larger. A second measure of the accuracy of the forecasting rule is given by the R^2 implied when the forecasting rule is used to fit a long sample of simulated data. For the benchmark model the R^2 is 0.99982, suggesting a high degree of accuracy.

Since decision rules should be continuous with respect to predicted values for future capital, I argue that improving forecasting accuracy would not lead to large changes in individual decision

[2] I also experimented with using two forecasting equations for aggregate capital, one for each value of the tax rate (a la Krusell and Smith 1998). This means estimating more parameter values with fewer observations for each regression (holding constant the total length of the simulation). I found that the co-efficients on aggregate capital and government debt were extremely similar for both equations, and therefore decided to estimate a single forecasting rule.

[3] The errors at a 1 year horizon are simply the percentage difference between the predicted and realized capital stocks. The errors at a 10 year horizon are computed similarly, except that the predicted capital stock is now what a household would predict capital to be 10 years hence given perfect foresight regarding the time path for the tax rate over that horizon.

rules or the aggregate behavior of the economy. Thus the Krusell and Smith computational technique is appropriate for this particular model economy.

7. *Calibrating the tax process*

The conditions that ensure that debt remain bounded in the benchmark model are

$$\tau_h \geq \frac{r\left(K_l, N(K_l,\tau_h)\right) D_h + G}{r\left(K_l, N(K_l,\tau_h)\right)\left(D_h + K_l\right) + w\left(K_l, N(K_l,\tau_h)\right) N\left(K_l,\tau_h\right)} \tag{7.1}$$

$$\overline{D} \leq \frac{D_h - G + \tau_l\left[w\left(K_l, N(K_l,\tau_l)\right) N_l(K_l,\tau_l) + r\left(K_l, N(K_l,\tau_l)\right) K_l\right]}{1 + r\left(K_l, N(K_l,\tau_l)\right)\left(1-\tau_l\right)}. \tag{7.2}$$

$$\tau_l \leq \frac{r\left(K_h, N(K_h,\tau_l)\right) D_l + G}{r\left(K_h, N(K_h,\tau_l)\right)\left(D_l + K_h\right) + w\left(K_h, N(K_h,\tau_l)\right) N\left(K_h,\tau_l\right)} \tag{7.3}$$

and

$$\underline{D} \geq \frac{D_l - G + \tau_h\left[w\left(K_h, N(K_h,\tau_h)\right) N(K_h,\tau_h) + r\left(K_h, N(K_h,\tau_h)\right) K_h\right]}{1 + r\left(K_h, N(K_h,\tau_h)\right)\left(1-\tau_h\right)}. \tag{7.4}$$

The algorithm used to pick values for tax parameters is as follows:

1. Guess a value for $D_h/\overline{Y}$ and a value for λ, the persistence parameter.

2. Given the historical mean debt to GDP ratio $\overline{B/Y}$ set $D_l/\overline{Y} = 2 \times \overline{B/Y} - D_h/\overline{Y}$.

3. Guess a value for government consumption G, a value for the mean of the limiting distribution for aggregate capital, $\overline{K}$ and a value for η to define the bounds on aggregate capital where $K_l = \overline{K} - \eta$ and $K_h = \overline{K} + \eta$.

4. Solve for values for τ_l, τ_h and $\overline{N}$ such that:

    1. $\overline{N}$ is the value for aggregate labor supply generated by eq. 1.3 given $\overline{K}$ and $\overline{\tau} = (\tau_l + \tau_h)/2$.

    2. τ_h is given by 7.1 at equality

    3. Given $\overline{B} = (D_h + D_l)/2$ where $D_h = \left(D_h/\overline{Y}\right) \times F(\overline{K}, \overline{N})$ and $D_l = \left(D_l/\overline{Y}\right) \times F(\overline{K}, \overline{N})$, set τ_l such that at $B = \overline{B}$ and $K = \overline{K}$ the absolute change in the stock of debt implied by eq. 1.4 is equalized across tax rates, *i.e.* $\Delta B|_{\tau=\tau_l} = -\Delta B|_{\tau=\tau_h}$. This condition effectively guarantees that low and high tax regimes are equally persistent, and thus

that the mean of the limiting distribution over the tax rate is in fact equal to $\overline{\tau}$. It is important to verify that the implied value for τ_l does not violate 7.3.

5. Given τ_l, τ_h, D_h, D_l, K_l, K_h, G compute $\overline{D}$, $\underline{D}$ using 7.2 and 7.4.

6. Simulate the economy for a very large number of periods and compute maximum and minimum values for aggregate capital. Revise η appropriately and resolve for τ_l, τ_h, D_h, D_l

7. Given bounds K_l and K_h that are not violated within a long simulation, compute the mean of the limiting distribution for tax revenue as a fraction of GDP. Adjust G appropriately, resolve for τ_l, τ_h, D_h, D_l, K_l, K_h and resimulate.

8. Given a value for G such that the model reproduces on average the historical ratio of tax revenue to GDP, check the variance of the tax revenue to GDP ratio in the limiting distribution and compare to the historical variance. Adjust $D_h/\overline{Y}$ appropriately, resolve for all other parameter values, and resimulate.

9. Given the correct variance, check the persistence of tax revenue to GDP across a simulation, and adjust λ appropriately.

8. *The process for idiosyncratic wages*

Many papers in the labor literature have argued that earnings are well represented by an ARIMA(0,1,1) or an ARIMA(1,1,1) process (for some recent examples, see Dickens 2000, Gottschalk and Moffitt 2002, or Meghir and Pistaferri 2002). There are two problems with trying to implement such a process for earnings (or wages) in this type of model. The first is that shocks with a permanent component do not fit very naturally into a model with infinitely-lived agents (I return to this point below). The second is that richer shock processes imply larger state spaces, presenting computational problems. If to construct probability distributions over future wages agents need to know (i) the current value of the random walk piece, (ii) the current value of the transitory shock (for the AR bit), and (iii) the transitory innovation (for the MA bit), then the state space is unmanageably large. Thus, in contrast to the labor literature, the applied (quantitative) macro literature has typically adopted much more parsimonious specifications.

In deciding on the right earnings model, one thing people look at is the auto-covariance of log earnings at different lags, and the auto-covariance of changes in log earnings at different lags. I therefore simulated the simple AR(1) process used to construct the wage process in the model, and compared this to an ARIMA(0,1,1) process from Meghir and Pistaferri (table 3, bottom panel, left column in their paper). Simulating lots of artificial households, each for 30 years, the ARIMA(0,1,1) process essentially reproduces the numbers Meghir and Pistaferri report for their pooled sample in table 1. Comparing the simple AR(1) process with theirs (without their measurement error piece), the auto-covariances (in (log) levels and log first differences), the two processes look roughly similar (see table 2).

The main difference I see here is that the ARIMA(0,1,1) is slightly more volatile that the one I use. If I scale up the variance of the innovations to the AR(1) process to give a variance of log earnings equal to 0.182, the simple AR(1) model generates the same auto-covariances at lags 1 and 2 as the ARIMA(0,1,1) process.

If you simulate people for more years, the unit root component of the ARIMA(0,1,1) process starts to dominate. The variance of earnings rises unboundedly, and auto-covariances do not decline as the lag length increases. Thus the ARIMA(0,1,1) process and the infinite horizon model do not fit together very happily. Still it is reassuring that the AR(1) process and the ARIMA(0,1,1) process look broadly similar when you set the length of working life to a reasonable number.

Note that the wage process as implemented in my model is unavoidably pretty crude simply by virtue of being a three state Markov process characterized by only four independent parameters. It is not even designed to be the best possible three state approximation to a continuous AR(1) process: the evidence from empirical AR(1) models is just used to generate target values for the persistence and variance of wage shocks.

9. *The process for wages in the stochastic-aging economy*

I assume that the probability of transiting from state e_i via the mechanism identified as aging is equal to $1/(\widehat{p}_i L)$, where $\widehat{p}_i$ is the fraction of the population with productivity e_i in the ergodic distribution over E, and L is a constant equal to expected lifetime. The events of aging and receiving a productivity shock are assumed mutually exclusive. The overall probability of moving from state i to state j, denoted $\widehat{\Pi}_{ij}$, is therefore equal to the probability of transiting from i to j via aging, plus the probability of transiting from i to j via a productivity shock, conditional on not aging.

$$\widehat{\Pi} = \begin{pmatrix} 0 & \frac{1}{\widehat{p}_1 L} & 0 \\ 0 & 0 & \frac{1}{\widehat{p}_2 L} \\ \frac{1}{\widehat{p}_3 L} & 0 & 0 \end{pmatrix} + \begin{pmatrix} (1-1/\widehat{p}_1 L) & 0 & 0 \\ 0 & (1-1/\widehat{p}_2 L) & 0 \\ 0 & 0 & (1-1/\widehat{p}_3 L) \end{pmatrix} \Pi. \tag{9.1}$$

The fractions $\widehat{p}_i$ are the solutions to the system of equations $\widehat{p} = \widehat{p}\widehat{\Pi}$. To calibrate the stochastic-aging economy I generate $\widehat{\Pi}$ using exactly the same Π matrix as in the benchmark model. I do this because I was unable to find an alternative data-consistent specification for Π that generates realistic wealth inequality when households take as given the process associated with $\widehat{\Pi}$. The average Gini co-efficient for wealth is only 0.63 compared to a value of 0.78 for the other economies (and for the U.S.). The main reason for the low Gini is that the stochastic-aging economy fails to capture concentration at the top end of the wealth distribution; the wealth-poorest 40 percent of households still account for less than two percent of total wealth. The median household by wealth in the stochastic-aging economy has 49 percent of mean wealth, compared to only 6 percent in the benchmark model. The figure for the U.S. is 28 percent.

Note that the estimates for the variance and persistence of wages used to calibrate the benchmark model are based on household level data that has been purged of variation attributable to age and education. Thus these estimates should be compared to the properties of the process for productivity implied by the Π matrix (rather than the $\widehat{\Pi}$ matrix).

Note also that while the aging / productivity shock distinction is a convenient conceptual device, it is irrelevant for agents in the model who only care about the implied transition probability matrix $\widehat{\Pi}$.

10. *More details on results*

10.1. *The role of labor supply with lump-sum taxes when markets are incomplete*

In section 5.1 of the paper I consider a heterogeneous-agent incomplete-markets model with lump-sum taxes and exogenous labor supply. In section 5.2 I consider an identical economy except that labor supply is now endogenous. Tax changes have slightly larger effects on aggregate consumption when labor supply is endogenous; the PCT is 14.9 versus 13.5 cents per dollar. The reason is that with endogenous labor supply, hours are now positively correlated with wages, so that individuals face more volatile earnings than when labor supply is fixed. Consequently they find it slightly harder to adjust savings in order to consumption-smooth through lump-sum tax shocks. This also translates into more precautionary saving, a high capital to output ratio, and a lower interest rate (see table 3 in the paper).

10.2. *Alternative preferences with proportional taxes when markets are complete*

Suppose that $u(c,n) = \left[c^{\mu}(1-n)^{1-\mu}\right]^{1-\sigma}/(1-\sigma)$. Given a value for μ such that on average $n = 0.3$, and recalibrating the tax process parameters, I compute the average impact effect of tax cuts for various values for σ. When $\sigma = 1$ (implying preferences are separable in consumption and leisure), the average propensity to consume out a tax cut and the average percentage changes in labor supply, consumption and investment are respectively 21.0 cents per dollar, 4.9 percent, 0.6 percent and 12.2 percent. Thus compared to the baseline Greenwood-Hercowitz-Huffman utility function used in the paper, the log-log alternative implies a similar response of consumption to tax cuts, but a larger response of hours and output. Note that this example illustrates that non-separability between consumption and leisure is not a necessary condition for tax cuts to increase consumption. Increasing σ above one increases the consumption response and reduces slightly the hours response to tax changes. However, for any Cobb-Douglas specification, the Frisch elasticity of labor supply is much higher than most empirical estimates from micro data.

10.3. *Wealth inequality and the response of aggregate consumption to tax changes*

Figure 3 in the paper indicates that tax cuts generally have a greater stimulate effect on aggregate consumption the smaller the fraction of total wealth held by the wealth-poorest 40 percent of households (see the top-left panel). During a prolonged interval of high tax rates the ratio of debt to GDP gradually falls. At the same time, households gradually run down their asset holdings and more households pile up against the no-borrowing constraint. The longer the tax rate remains high before eventually falling, the stronger the response of aggregate consumption, both because more poor households have exhausted their assets, and also because the lower is debt, the longer the tax cut is expected to last.

By contrast, tax increases appear to have more powerful effects in periods when wealth-poor households are relatively wealthy, though greater scatter in this case indicates a weaker relationship between the consumption effect and this measure of wealth inequality. To understand this finding, consider a sustained period of low taxes. Expecting a tax increase in the future, households gradually accumulate more precautionary asset holdings. The longer taxes remain low, the larger is the ratio of debt to GDP when the tax rate eventually rises, and the larger the asset holdings of the wealth-poorest households. There are, however, two off-setting effects on the size of the aggregate consumption response when taxes eventially rise. On the one hand, if taxes have been low for a long time, fewer households are close to the borrowing constraint and unable to smooth through a tax increase. On the other hand, the ratio of debt to GDP will tend to be high, implying that the tax increase is likely to be quite persistent and should therefore have a large effect on aggregate consumption.

10.4. *The persistence of tax changes and the response of aggregate consumption to tax changes*

One might suspect that one way to generate larger real effects is to increase the persistence of tax shocks. I therefore experiment with lowering the persistence parameter λ so that the first order autocorrelation of tax revenue to GDP is 0.77 instead of 0.63 (which implies an average duration of tax changes of 7.7 years rather than 5 years). In this case the average PCT in the incomplete markets version rises very slightly (from 28.8 to 29.4 cents per dollar) while the average PCT in the complete markets version declines slightly (from 23.2 to 22.1 cents). I conclude that the impact effect of tax changes is not particularly sensitive to their expected duration.

References

[1] Christiano, Lawrence J. and Jonas D. M. Fisher. 2000. Algorithms for solving dynamic models with occasionally binding constraints. *Journal of Economic Dynamics and Control*, 24, no. 8, 1179-1232.

[2] Dickens, Richard. 2000. The evolution of individual male earnings in Great Britain: 1975-95. *Economic Journal*, 110, no. 460, 27-49.

[3] Gottschalk, Peter and Robert A. Moffitt. 2002. Trends in the transitory variance of earnings in the United States. *Economic Journal*, 112, no. 478, C68-C73.

[4] Krusell, Per and Anthony A. Smith Jr. 1998. Income and wealth heterogeneity in the macroeconomy. *Journal of Political Economy*, 106, no. 5, 867-896.

[5] McGrattan, Ellen. 1999. Applying weighted residual methods to dynamic economic models. Chapter 6 in Marimon and Scott eds. *Computational Methods for the Study of Dynamic Economics*, Oxford University Press.

[6] Meghir, Costas and Luigi Pistaferri. 2002. Income variance dynamics and heterogeneity. Institute for Fiscal Studies Working Paper W01/07.

[7] Rios-Rull, Jose-Victor. 1999. Computation of equilibria in heterogeneous-agent models. Chapter 11 in Marimon and Scott eds. *Computational Methods for the Study of Dynamic Economics*, Oxford University Press.

Table 1: Law of Motion for capital

		ECONOMY		
		Benchmark Incomplete-markets	Lump-sum taxes, endogenous labor, prices	Stochastic-aging economy
Co-efficients	α_1	0.918	0.995	0.922
	α_2	-0.0083	-0.0009	-0.0087
	α_3	-0.0060	0.0123	-0.0045
R-squared	R^2	0.9998	0.9984	0.9997

Table 2: Comparison of alternative shock processes

Lag	ARIMA(0,1,1), levels	Simple AR(1), levels	ARIMA(0,1,1), first differences	Simple AR(1), first differences
0	0.182	0.131	0.077	0.052
1	0.142	0.106	-0.018	-0.003
2	0.116	0.081	-0.009	-0.003
3	0.097	0.060	-0.001	-0.003
4	0.081	0.041	-0.001	-0.003

Figure 1: Distribution over asset holdings
(average over 10,000 period simulation)

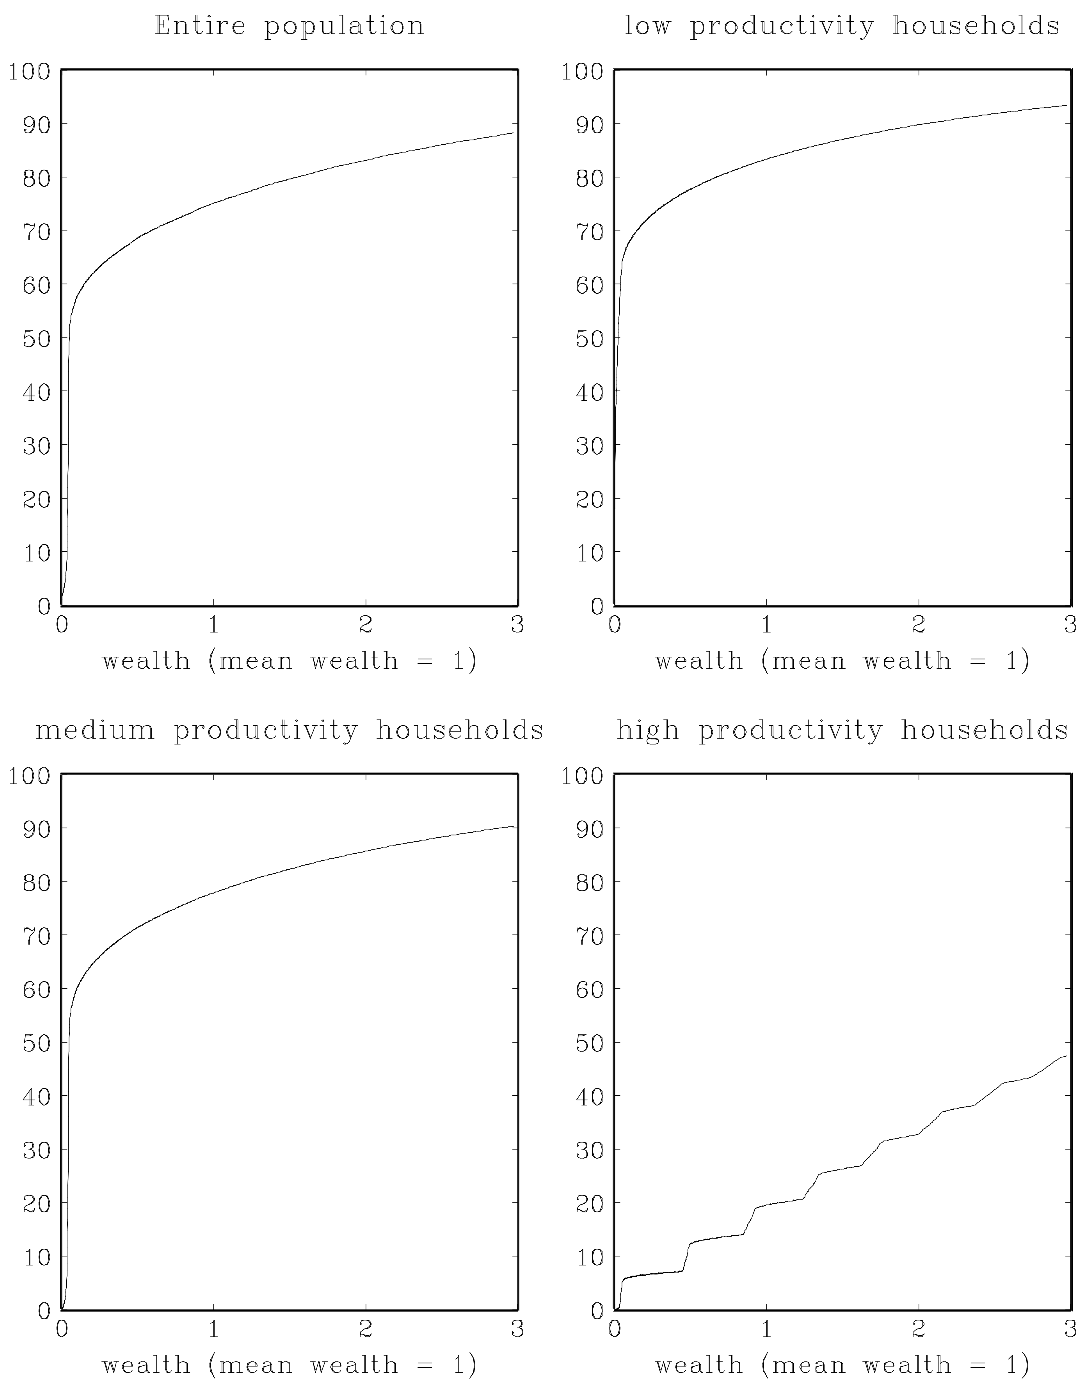

Figure 2: Forecast errors for capital and the pre-tax net real interest rate