Econ 411-3 Problem Set 1

Due by noon, Monday, April 15[1]

Problem 1: a stochastic process for discount rate beta. Start with the default calibration of the partial equilibrium standard incomplete markets model that we have used in class (e.g. the calibration returned by the `example_calibration()` function in `sim_steady_state.py`). Plot the average quarterly MPC (on the vertical axis) against the aggregate ratio of assets to labor income (on the horizontal axis) as we vary β over a wide range, including high enough β to achieve an annual asset-income ratio of 10 (annual) or 40 (quarterly).[2]

Now, suppose that the exogenous state s consists of two independent components, s_1 and s_2, so that $s = (s_1, s_2)$. Assume that y is a function of s_1, and that $y(s_1)$ and the Markov process followed by s_1 are unchanged relative to above. Assume that s_2 can take two values, L and H, that $\beta(s_2)$ is given by

$$\beta(s_2) = \begin{cases} \bar{\beta} - \Delta & s_2 = L \\ \bar{\beta} + \Delta & s_2 = H \end{cases}$$

and that the transition matrix for s_2 between L and H is given by

$$\Pi_2 = \begin{pmatrix} 0.99 & 0.01 \\ 0.01 & 0.99 \end{pmatrix}$$

Here, the $s_2 = L$ state is meant to represent an "impatient" low-β household, and $s_2 = H$ is meant to represent a "patient" high-β household. The average time spent in either state is 100 quarters or 25 years—a length of time we might (loosely) interpret as a generation.

Make the same plot of average quarterly MPC vs. aggregate assets for this new model where $\Delta = 0.01$ (or 0.04 when annualized), and superimpose it on top of our earlier plot, which corresponded to $\Delta = 0$. How do the two plots differ? Explain. Does the model with stochastic β allow a better fit to MPCs and wealth? What is the intuition behind this?

Finally, solve for the $\bar{\beta}$ in each case that gives an annual wealth-to-labor income ratio of 5, and report the corresponding average MPCs for each. Also report the Gini coefficient for wealth inequality in each case, and plot the Lorenz curves together (see problem 3 below for more detail on this). How does wealth inequality differ, and why do you think this is?

Note: this problem can be done entirely using the `sim.steady_state()` function, with both `y` and `beta` now being arrays of length S, where S is the number of composite states $s = (s_1, s_2)$. You will have to figure out what the transition matrix `Pi` for s should look like. Further, for the code to work properly, `beta` should be supplied to the `sim.steady_state()` function as an $S \times 1$ array; this can be done by taking a length-S `beta` and writing `beta[:, np.newaxis]`. This last point is necessary for the product `beta * Pi` in the code to use the discount rate corresponding to the current state rather than the future state.

[1] You can work in groups of up to four; you only have to submit once per group, but remember to list all members of the group when submitting. Please email solutions to Jose Lara (joselara@u.northwestern.edu), including whatever code you used to produce them. You may want to do much or all of the problem set in a Jupyter notebook. For this problem set, you are free to reuse any code that has been posted for the lectures on Canvas as part of your solution, and indeed most of you will probably use `sim_steady_state.py` or `sim_steady_state_fast.py`.

[2] See `lecture 2 more figures.ipynb` for code that calculates average MPCs and generates a similar plot, although that code varies r rather than β.

Problem 2: wealth inequality in partial equilibrium. Return to the partial equilibrium standard incomplete markets model with our benchmark calibration (i.e. `example_calibration()`).

The *Lorenz curve*, if used to represent the distribution of wealth, is a plot of the cumulative share of wealth y% held by the bottom x% of the distribution. (See, for instance, the Wikipedia article.) It goes from $(0, 0)$ to $(1, 1)$, and lies weakly below the 45-degree line connecting $(0, 0)$ to $(1, 1)$.

The *Gini coefficient* is the area between the 45-degree line and the Lorenz curve, multiplied by 2. It is a measure of inequality: it is 0 under perfect equality (where the Lorenz curve is the 45-degree line) and 1 under perfect inequality (where the Lorenz curve stays at y% until the highest percentile).[3]

First, plot the Lorenz curve and calculate the Gini coefficient of wealth for our benchmark calibration. Note that this is a bit tricky since our wealth distribution consists of point masses on a grid rather than a continuous density. To solve this, you should think of the mass of households at a point as being uniformly distributed in the relevant percentile range. For instance, if 50% of people have wealth 1, and 50% of people have wealth 3, then you should think of the first group as being uniformly distributed between the 0 and 0.5 percentiles, and the second group as being distributed between the 0.5 and 1 percentiles, and the Lorenz curve should be pointwise linear with two segments, the first connecting $(0, 0)$ to $(0.5, 0.25)$, and the second connecting $(0.5, 0.25)$ to $(1, 1)$.[4]

Then, show how the Gini coefficient varies across steady-state calibrations as we vary r from 1% lower (4% annualized) than its benchmark level to 1% higher (4% annualized) than its benchmark level. Does it rise or fall with r? Do you have any intuition for why this might be? Plot all the Lorenz curves together.

Finally, consider the following alternate interpretation of the Gini coefficient: that if p is the average percentile (ordered by wealth) of households at which each dollar of wealth is held in the economy, which can in principle range between 0.5 (perfect equality) and 1 (perfect inequality), then $2p - 1$ is the Gini coefficient.[5] For instance, in the example above where 50% of people have wealth 1 and 50% of people have wealth 3, the average percentile at which wealth is held is

$$\frac{50\% \times 0.25 \times 1 + 50\% \times 0.75 \times 3}{50\% \times 1 + 50\% \times 3} = 0.625$$

and the Gini coefficient is therefore $2 \times 0.625 - 1 = 0.25$. (The 0.25 and 0.75 are the average percentiles of the lower and upper groups.) Verify for our benchmark calibration that this way to calculate the Gini coefficient gives the same answer as our earlier calculation.

Problem 3: effect of changing capital share in Aiyagari model. Consider the Aiyagari model as specified in lecture 4 (with zero borrowing limit and CRRA preferences, so that $a(r)$ can be defined), in the case with Cobb-Douglas technology $F(K, L) = K^\alpha L^{1-\alpha}$.

i) Derive, to first order, the effect of a change $d\alpha$ in the capital share α on the steady-state real interest rate r, the steady-state ratio K/wL of capital to wage income, and the steady-state ratio K/Y of capital to output, obtaining analytical formulas (with semelasticities) analogous to the ones we derived in class for shifts in asset demand.

ii) Starting with the steady state used in class for lecture 4 (which can be seen in the `lecture 4 figures.ipynb` notebook), which has $\alpha = 0.25$, plot the $a(r)$ vs. $k(r)/w(r)$ curves, and plot how the latter shifts, and

[3]Here x and y are actually from 0 to 1, not from 0 to 100; I'm just saying "percentile" and "percent" because they sound more natural.

[4]There is some code that does this in `lecture 2 more figures.ipynb` that you can adapt if desired.

[5]If you have some extra time, you can try to prove that this is equivalent to our definition using the Lorenz curve!

equilibrium r changes, when the capital share increases to $\alpha = 0.30$. How close was the first-order approximation from the analytical formula obtained above? How much of the "partial equilibrium" increase in $k(r)/w(r)$ from a rising capital share is offset by the change in r? Discuss.