1 Non-Stationary Equilibria

Recall that last time we have studied a problem of the form:

$$\begin{aligned} v(a,\varepsilon) &= \max_{c,a'} \left\{ u(c) + \beta \sum_{\varepsilon' \in E} v(a',\varepsilon')\pi(\varepsilon',\varepsilon) \right\} \\ & \quad s.t. \\ c + a' &= (1+r)\,a + w\varepsilon \\ a' &\geq -b \end{aligned} \tag{1}$$

We said that a stationary equilibrium is a value function $v : S \to \mathbb{R}$; policy functions for the household $a' : S \to \mathbb{R}$, and $c : S \to \mathbb{R}_+$; policies for the firm H and K; prices r and w; and, a stationary measure λ such that all agents optimize, markets clear, and the invariant distribution λ satisfies

$$\lambda\left(\mathcal{A} \times \mathcal{E}\right) = \int_{A \times E} Q\left(\left(a,\varepsilon\right), \mathcal{A} \times \mathcal{E}\right) d\lambda\left(a,\varepsilon\right)$$

where Q is the transition function:

$$Q\left(\left(a,\varepsilon\right), \mathcal{A} \times \mathcal{E}\right) = I_{\{a'(a,\varepsilon)\in\mathcal{A}\}} \sum_{\varepsilon' \in \mathcal{E}} \pi(\varepsilon',\varepsilon)$$

What makes this problem tractable is 1) the assumption that ε is iid across households (no aggregate uncertainty) and thus 2) our focus on a stationary equilibrium in which λ does not change over time (even though individual agents are subject to shocks that change their idiosyncratic a and ε from one period to another). Because the distribution is stationary, r and w are constant and the algorithm we use to solve for a stationary equilibrium is simply a search over the value of r that clears the asset market.

We would next like to drop the stationarity assumption. In particular, we want to allow for aggregate shocks to, say, ε so as to study business cycle fluctuations and the role aggregate uncertainty plays in this environment. In addition, we would like to study the effect of changes in the tax code and other policies that will affect the measure λ. This turns out to be a very difficult problem. To see why, consider an economy in which the household's endowment of efficiency units is equal to $\varepsilon + \varepsilon^a$ where ε is as earlier iid across agents, but ε^a in an aggregate shock common to all agents (e.g. aggregate productivity shock). Assume ε^a is also a Markov process with $\pi^a\left(\varepsilon^{a\prime}|\varepsilon^a\right)$.

Clearly, as ε^a changes over time, so does the mean (and other moments) of the distribution of asset holdings of the agents. For example, in booms (high ε^a) agents will accumulate capital to smooth fluctuations in consumption and the average a will increase. The distribution λ will thus change over time in response to changes in ε^a. Let $\lambda' = \Gamma(\lambda, \varepsilon^a)$ be the law of motion for this distribution and let us assume that λ and ε^a are sufficient[1] to describe the aggregate state of this economy. What this means in practice is that we can compute aggregate prices depend solely on these two objects: $w(\lambda, \varepsilon^a)$ and $r(\lambda, \varepsilon^a)$. We can thus write the consumer's problem recursively as:

$$
\begin{aligned}
v(a, \varepsilon; \lambda, \varepsilon^a) &= \max_{c,a'} \left\{ u(c) + \beta \sum_{\varepsilon'} \sum_{\varepsilon^{a\prime}} v(a', \varepsilon'; \lambda', \varepsilon^{a\prime}) \pi(\varepsilon', \varepsilon) \pi^a(\varepsilon^{a\prime}, \varepsilon^a) \right\} \qquad (2) \\
& \quad s.t. \\
c + a' &= (1 + r(\lambda, \varepsilon^a)) a + w(\lambda, \varepsilon^a) \varepsilon \\
a' &\geq -b \\
\lambda' &= \Gamma(\lambda, \varepsilon^a) \qquad (3)
\end{aligned}
$$

A recursive equilibrium in this economy is, in principle, straightforward to define (we will do so below) and Γ (the law of motion for the distribution λ) would be part of its definition. Characterizing this equilibrium, even with the aid of a computer, is in most cases impossible, because λ is an infinite-dimensional object.

What follows are 2 different techniques to solve this problem.

In the first one we will consider one-time (finite-number) of changes in ε^a (or other aggregate variable) that are anticipated. To fix ideas, we will think of one-time changes in the tax rate the government levies on labor earnings. By focusing one one-time (or finite number) of changes, we can exploit the fact that the economy converges to a stationary distribution eventually, conjecture that it does so in a finite number of periods after the shock, and search for the prices that clear the markets and the value functions and household decision rules that are consistent with these prices. Because we converge in a finite number of periods, this is a finite number of equilibrium objects to search for.

The second technique reduces to replacing λ with a finite number of moments of this distribution. We describe each method separately.

[1] This is true in many applications, but not always. See Kubler and Schmedders (2002).

2 Unanticipated one-time changes

(This can be easily modified to allow for a finite number of changes). We assume the goverment taxes a fraction τ of labor income, so the budget constraint reads:

$$c + a' = (1 + r)\, a + w\,(1 - \tau)\,\varepsilon$$

We want to study the welfare effect of a tax reform, e.g. an *unexpected* permanent rise in the labor income tax rate from τ^* to τ^{**}.[2] Suppose that taxes are used to finance government transfers ϕ which are lump-sum.

Steady-state comparison– The simplest way to do this is computing a stationary RCE for these two levels of the tax rate and compare aggregate variables and welfare between the two steady-states. However, this approach is not fully satisfactory. It can only be used to assess whether a household would prefer to live in the stationary equilibrium of an economy with tax rate equal to τ^* or in the stationary equilibrium of an economy with tax rate equal to τ^{**}.

Transition– A more interesting and relevant policy question is: consider a household living in the stationary equilibrium of an economy with initial tax rate τ^*. What is the welfare change (gain or loss), for this household, associated to a rise in the tax rate from τ^* to τ^{**}? To answer this question properly, one needs to compute the whole transition: the new policy will change the individual consumption/saving and labor supply decision, hence aggregate prices and will induce dynamics away from the current steady-state towards the new one (assuming the system has stable dynamics). See Heathcote (2004) for a recent example that applies these techniques.

How do we attack this problem? Since the transition is characterized by a sequence of aggregate prices and quantities, we need to modify appropriately the definition of recursive competitive equilibrium.

2.1 Definition of Equilibrium with Transition

First, let's define the household problem at time t still in recursive form

[2] It is important that the policy change is unexpected. If it was anticipated, agents would take actions in advance, for example to guarantee that their consumption path throughout the transition is as smooth as possible.

$$\begin{aligned} v_t(a,\varepsilon) &= \max_{c_t,a_{t+1}} \left\{ u(c_t(a,\varepsilon)) + \beta \sum_{\varepsilon_{t+1}\in E} v_{t+1}(a_{t+1}(a,\varepsilon),\varepsilon_{t+1})\pi(\varepsilon_{t+1},\varepsilon) \right\} \quad (4) \\ & s.t. \\ c_t(a,\varepsilon) + a_{t+1}(a,\varepsilon) &= (1+r_t)\, a + w_t\,(1-\tau_t)\,\varepsilon + \phi_t \\ a_{t+1}(a,\varepsilon) &\geq -b \end{aligned}$$

Note now that value functions and policies are also a function of time since aggregate prices (r_t, w_t) are time-varying. Let's denote the initial stationary distribution with λ^*.

Given an initial distribution λ^*, and a sequence of tax rates $\{\tau_t\}_{t=0}^{\infty}$, a *recursive competitive equilibrium* is a sequence of value functions $\{v_t\}_{t=0}^{\infty}$ and decision rules for households $\{c_t, a_{t+1}\}_{t=0}^{\infty}$, firm choices $\{H_t, K_t\}_{t=0}^{\infty}$, prices $\{w_t, r_t\}_{t=0}^{\infty}$, government transfers $\{\phi_t\}_{t=0}^{\infty}$ and distributions $\{\lambda_t\}_{t=0}^{\infty}$ such that, for all t:

- given prices $\{r_t, w_t\}$ and policies $\{\tau_t, \phi_t\}$, the decision rules $a_{t+1}(a,\varepsilon)$ and $c_t(a,\varepsilon)$ solve the household's problem (4) and $v_t(a,\varepsilon)$ is the associated value function,
- given prices $\{r_t, w_t\}$, the firm chooses optimally its capital K_t and its labor H_t, i.e. $r_t + \delta = F_K(K_t, H_t)$ and $w_t = F_H(K_t, H_t)$,
- the labor market clears: $H_t = \int_{A\times E} \varepsilon d\lambda_t = H$,
- the asset market clears: $K_{t+1} = \int_{A\times E} a_{t+1}(a,\varepsilon) d\lambda_t$,
- the goods market clears: $\int_{A\times E} c_t(a,\varepsilon) d\lambda_t + K_{t+1} - (1-\delta)\, K_t = F(K_t, H_t)$,
- the government budget constraint is balanced: $\phi_t = \tau_t w_t H$,
- for all $(\mathcal{A}\times\mathcal{E}) \in \mathcal{S}$, the probability measure λ_{t+1} satisfies

  $$\lambda_{t+1}(\mathcal{A}\times\mathcal{E}) = \int_{A\times E} Q_t((a,\varepsilon), \mathcal{A}\times\mathcal{E})\, d\lambda_t,$$

  where Q_t is the transition function defined as

  $$Q_t((a,\varepsilon), \mathcal{A}\times\mathcal{E}) = I_{\{a_{t+1}(a,\varepsilon)\in\mathcal{A}\}} \sum_{\varepsilon_{t+1}\in\mathcal{E}} \pi(\varepsilon_{t+1},\varepsilon). \quad (5)$$

2.2 Numerical Computation of the Transition Path

The economy at time $t = 0$ is in steady-state with stationary distribution $\lambda_0 = \lambda^*$ over assets and individual productivities and tax rate τ^*. At the end of period $t = 0$, the government announces that from $t = 1$ onward the tax policy will change to $\tau^{**} > \tau^*$ and that the additional revenues will augment the lump-sum transfer ϕ_t. Hence, the relevant tax sequence needed to compute the equilibrium is

$$\tau_t = \begin{cases} \tau^*, & \text{for } t = 0 \\ \tau^{**}, & \text{for } t \geq 1. \end{cases}$$

Next, we will assume that after T periods, with T arbitrarily large but finite, the economy will settle to the final steady-state. This assumption is helpful because it allows us to guess a finite sequence of aggregate capital stocks and use backward induction for the solution of the household problem.

To compute the equilibrium follows these steps:

1. Fix T (say $T = 200$)

2. Compute the initial steady state objects $\{v^*, c^*, a^*, K^*\}$ corresponding to $\tau = \tau^*$ and the final steady state objects $\{v^{**}, c^{**}, a^{**}, K^{**}\}$ corresponding to $\tau = \tau^{**}$.

3. Guess a sequence of aggregate capital stocks $\left\{\hat{K}_t\right\}_{t=1}^{T}$ of length T such that $\hat{K}_1 = K^*$ (capital at time 1 is predetermined at time $t = 0$ which is a steady-state) and $\hat{K}_T = K^{**}$. Note that $H_t = H$ (i.e. constant) for every t in absence of endogenous labor supply. Hence, it is easy to determine, for each t,

$$\begin{aligned} \hat{w}_t &= F_H(\hat{K}_t, H), \\ \hat{r}_t &= F_K(\hat{K}_t, H), \\ \hat{\phi}_t &= \tau_t \hat{w}_t H, \end{aligned}$$

   which are all the elements we need in the budget constraint of the household to solve the household problem at time t.[3]

4. Since we know that $v_T(a, \varepsilon) = v^{**}(a, \varepsilon)$, we can solve the household problem *by backward induction* and derive $\{\hat{v}_t(a, \varepsilon)\}_{t=1}^{T-1}$ from (4) and the associated policy functions $\{\hat{a}_{t+1}(a, \varepsilon)\}_{t=1}^{T-1}$.

[3]In this step, one can equally guess a path for the interest rate or for wages. Then, from the FOC's of the firm, one would recover aggregate capital at each t.

5. Given the policy functions, we can reconstruct the sequence of transition functions $\left\{\hat{Q}_t\right\}_{t=1}^{T}$ and, since we know that $\lambda_0 = \lambda^*$, we can recover the whole sequence of measures $\left\{\hat{\lambda}_t\left(a,\varepsilon\right)\right\}_{t=1}^{T}$ and calculate

$$\hat{A}_{t+1} = \int_{A\times E} \hat{a}_{t+1}(a,\varepsilon)d\hat{\lambda}_t.$$

   To compute this integral, we can use simulation techniques. We simulate histories of length T for N workers (say $N = 10,000$) starting from the steady-state distribution at $t = 1$ (distribution that we also obtain by simulation). Note that when computing the optimal consumption and saving choices of each of the N individuals in our sample at time t, we must use the time t decision rules computed in the previous step.

6. Check market clearing in the asset market in every period t, i.e. check if the guess of equilibrium capital stocks $\left\{\hat{K}_t\right\}_{t=1}^{T}$ is consistent with aggregate wealth $\left\{\hat{A}_t\right\}_{t=1}^{T}$ that households would accumulate when facing the sequence of prices induced by the proposed sequence of aggregate capital. In other words, choose a convergence criterion ε and check whether

$$\max_{1\leq t\leq T} |\hat{A}_t - \hat{K}_t| < \varepsilon. \tag{6}$$

   Note that if $|\hat{A}_T - K^{**}| < \varepsilon$ is satisfied, we have implicitly also checked that T is large enough.

7. If inequality (6) is not satisfied at every t, we need a new guess of the capital stock, for example

$$\hat{K}_t^{new} = 0.5\left(\hat{K}_t^{old} + \hat{A}_t\right),$$

   and go back to step 3. with this new guess.

2.3 Computing the Welfare Change from the Tax Reform

The crucial question to ask, from a policy perspective, is: how much agents gain/lose from the tax reform? This question takes us deep into welfare analysis.

In the first steady-state, an agent with initial individual state (a, ε) has expected lifetime utility associated with the stationary decision rule $c^*(a, \varepsilon)$ given by

$$v^*(a, \varepsilon) = E_0 \left[\sum_{t=0}^{\infty} \beta^t u(c_t^*) \,|(a_0 = a, \varepsilon_0 = \varepsilon) \right],$$

where the conditional expectation E_0 is taken over histories of the shocks conditional on a time-zero realization of the shock equal to ε (as made clear by the second equality) and conditional to an agent's wealth level equal to a.

Define the expected discounted utility of an agent with initial state (a, ε) at date $t = 0$ going through the transition as

$$\tilde{v}(a, \varepsilon) = E_0 \left[\sum_{t=0}^{\infty} \beta^t u(\tilde{c}_t) \,|(a_0 = a, \varepsilon_0 = \varepsilon) \right],$$

where $\tilde{v}$ differs from v^* because it is computed using the sequence of decision rules $\{\tilde{c}_t(a, \varepsilon)\}_{t=0}^{\infty}$ along the transition path.

Note that both discounted utilities can be easily computed by simulating a sample of agents in the initial steady state (for v^*) and along the transition (for $\tilde{v}$), based on the fact that for N large enough we can approximate the expectation operator (using the law of large numbers).

Conditional welfare change– The first question we can ask is: how much would an agent with initial state (a, ε) gain, in percentage terms of lifetime consumption, if he went through the transition induced by the policy reform, compared to the no-reform scenario where he lives in the initial steady-state forever? So, welfare changes are expressed in terms of consumption-equivalent variation.[4]

The answer to this question is a function $\omega(a, \varepsilon)$ that solves the equation

$$E_0 \sum_{t=0}^{\infty} \beta^t u([1 + \omega(a, \varepsilon)]\, c_t^*) = E_0 \sum_{t=0}^{\infty} \beta^t u(\tilde{c}_t).$$

For the case of power-utility this calculation is really easy to make. When $u(c) = c^{1-\sigma}$, we can exploit the homogeneity of the utility function and the equation above becomes

$$\begin{aligned} [1 + \omega(a, \varepsilon)]^{1-\sigma} v^*(a, \varepsilon) &= \tilde{v}(a, \varepsilon), \\ \omega(a, \varepsilon) &= \left[\frac{\tilde{v}(a, \varepsilon)}{v^*(a, \varepsilon)} \right]^{\frac{1}{1-\sigma}} - 1. \end{aligned} \tag{7}$$

[4] Precisely, we ask: "how much do we need to increase consumption of the agent in every state in the stationary equilibrium so that he'd be indifferent between the reformed world and the no-reform world?".

This welfare change is called *conditional welfare change*, because it is computed for an individual that is in a particular state (a, ε). Thus, we can compute the welfare change for the rich household, the poor household, the productive household, the unproductive household, etc...Moreover, we can compute the entire distribution of welfare changes and study whether the reform would be politically feasible, e.g. whether the majority of agents have positive welfare gains, hence they would support the reform.

Utilitarian social welfare change– The second type of welfare calculation is based on a Benthamian social welfare function that puts equal weight to every household in the initial steady-state (which is also the initial period of the reform), i.e. it uses the weighting criterion $\lambda^*(a, \varepsilon)$. The solution to this welfare calculation, for the power utility case is *one number only*, ω^U that solves

$$\omega^U = \left[\frac{\int_{A\times E} \tilde{v}(a, \varepsilon)\, d\lambda^*}{\int_{A\times E} v^*(a, \varepsilon)\, d\lambda^*}\right]^{\frac{1}{1-\sigma}} - 1.$$

So, ω^U computes the welfare change for "society", where every agent in society is given equal weight: some will lose, some will gain and we average across everyone with equal weights.

An alternative interpretation of this welfare criterion is that of an *ex-ante* welfare gain, or welfare gain *under the veil of ignorance*. In other words, $\int_{A\times E} \tilde{v}(a, \varepsilon)\, d\lambda^*(a, \varepsilon)$ represents the expected discounted utility of a newborn agent who is dropped at random in the first steady-state without knowing at which point in the distribution she will be, i.e. under the veil of ignorance.

2.4 A Welfare Change Decomposition

Sources of changes in social welfare– The typical utilitarian equal-weight social welfare function is:

$$\int_{A\times E} v(a, \varepsilon)\, d\lambda = \int_{A\times E} E_0 \sum_{t=0}^{\infty} \beta^t u(c_t)\, d\lambda \tag{8}$$

can increase for three reasons:

1. If average consumption $E(c_t)$ increases, since utility is monotone. It is easy to see that in presence of additional resources, at least one individual can be made better off, while the others receive the same utility. We can call this effect, the *level effect*.

For example, by reallocating resources more efficiently a tax reform can increase average consumption.

2. If the uncertainty/volatility of each individual consumption path $\{c_t\}_{t=0}^{\infty}$ is reduced, since agents are risk averse. We can call this the *uncertainty effect*. For example, by redistributing from the lucky to the unlucky, the tax reform can provide additional insurance.

3. If inequality across individuals at any point in time is reduced, since the value function is concave. This is easily seen by just applying Jensen's inequality to the left hand side of (8). We can call this effect, the *egalitarian effect*. Note the key difference between 2) and 3): even when there is no uncertainty in consumption sequences, a policy that redistributes initial wealth more equally across agents would achieve a welfare gain, under this social welfare function. This makes the social welfare function not a fully desirable criterion when studying the welfare implication of a policy reform because it *mixes concern for risk/uncertainty with concern for interpersonal equality*.

Conditional welfare is preferable– For this reason, conditional welfare is a somewhat more satisfactory welfare criterion because only the level and uncertainty effect play a role. We now show that, starting from the conditional welfare criterion, the welfare change $\omega(a,\varepsilon)$ can be decomposed additively into 1 and 2. This is a useful result because, when evaluating the welfare implication of a policy reform, one can compute separately the welfare change due to the fact that the policy is 1) generating more/less aggregate consumption, and 2) increasing/decreasing consumption insurance.

Decomposition– For simplicity, we focus on a welfare comparison between steady-states, but the methodology can be extended to transitions. Let the ex-ante welfare change between economy A (e.g. the initial steady state) and economy B (e.g. the final steady-state) be

$$E_0 \sum_{t=0}^{\infty} \beta^t u\left((1+\omega)\, c_t^A\right) = E_0 \sum_{t=0}^{\infty} \beta^t u\left(c_t^B\right), \tag{9}$$

where E_0 is the expectation taken at date $t = 0$ conditional on an initial value for the pair (a,ε). To simplify the notation, we have omitted the dependence of ω from the pair (a,ε).

Let C^j denote the average consumption in economy $j = A, B$, i.e.

$$C^j = \int c^j (a, \varepsilon) \, d\lambda^j, \text{ with } j = A, B.$$

Then the *welfare gain of increased consumption levels* between A and B ω^{lev} is defined by

$$\left(1 + \omega^{lev}\right) C^A \equiv C^B. \tag{10}$$

Next, let the certainty equivalent consumption bundle be defined by $\bar{C}^j$ that solves

$$E_0 \sum_{t=0}^{\infty} \beta^t u \left(c_t^j\right) = \sum_{t=0}^{\infty} \beta^t u \left(\bar{C}^j\right). \tag{11}$$

Then, we can define the cost of uncertainty p^j as

$$\sum_{t=0}^{\infty} \beta^t u \left(\left(1 - p^j\right) C^j\right) \equiv \sum_{t=0}^{\infty} \beta^t u \left(\bar{C}^j\right), \tag{12}$$

which is the fraction of average consumption that an individual in economy j would be willing to give up to avoid all the risk associated to productivity fluctuations.

Then, the *welfare gain of reduced uncertainty* between economy A and economy B is

$$\omega^{unc} \equiv \frac{1 - p^B}{1 - p^A} - 1. \tag{13}$$

We are now ready to state:

Proposition (Floden, 2001): Assume that u is "homogenous" in the sense that $u(xc) = g(x) u(c)$, then

$$1 + \omega = \left(1 + \omega^{unc}\right) \left(1 + \omega^{lev}\right) \Rightarrow \omega \simeq \omega^{unc} + \omega^{lev}.$$

Proof: The total welfare change is given by that value for ω that solves (9). Consider the expected utility in economy B, i.e. the R.H.S. of equation (9):

$$E_0 \sum_{t=0}^{\infty} \beta^t u \left(c_t^B\right) = \sum_{t=0}^{\infty} \beta^t u \left(\bar{C}^B\right) = \sum_{t=0}^{\infty} \beta^t u \left(\left(1 - p^B\right) C^B\right) = g \left(1 - p^B\right) \sum_{t=0}^{\infty} \beta^t u \left(C^B\right),$$

where the first equality follows from (11), the second from (12) and the third from the homogeneity assumption. Then,

$$g \left(1 - p^B\right) \sum_{t=0}^{\infty} \beta^t u \left(C^B\right) = g \left(1 - p^B\right) \sum_{t=0}^{\infty} \beta^t u \left(\left(1 + \omega^{lev}\right) C^A\right),$$

where the equality follows from definition (10). Next,

$$
\begin{aligned}
g\left(1-p^{B}\right)\sum_{t=0}^{\infty}\beta^{t}u\left(\left(1+\omega^{lev}\right)C^{A}\right) &= g\left(1-p^{B}\right)g\left(1+\omega^{lev}\right)\sum_{t=0}^{\infty}\beta^{t}u\left(C^{A}\right)\\
&= g\left(\frac{1-p^{B}}{1-p^{A}}\right)g\left(1+\omega^{lev}\right)g\left(1-p^{A}\right)\sum_{t=0}^{\infty}\beta^{t}u\left(C^{A}\right)\\
&= g\left(\frac{1-p^{B}}{1-p^{A}}\right)g\left(1+\omega^{lev}\right)\sum_{t=0}^{\infty}\beta^{t}u\left(\left(1-p^{A}\right)C^{A}\right)
\end{aligned}
$$

The line above follows from the homogeneity assumption. Using definition (12),

$$
g\left(\frac{1-p^{B}}{1-p^{A}}\right)g\left(1+\omega^{lev}\right)\sum_{t=0}^{\infty}\beta^{t}u\left(\left(1-p^{A}\right)C^{A}\right)=g\left(\frac{1-p^{B}}{1-p^{A}}\right)g\left(1+\omega^{lev}\right)\sum_{t=0}^{\infty}\beta^{t}u\left(\bar{C}^{A}\right).
$$

From the definition of certainty equivalent consumption

$$
g\left(\frac{1-p^{B}}{1-p^{A}}\right)g\left(1+\omega^{lev}\right)\sum_{t=0}^{\infty}\beta^{t}u\left(\bar{C}^{A}\right)=g\left(\frac{1-p^{B}}{1-p^{A}}\right)g\left(1+\omega^{lev}\right)E_{0}\sum_{t=0}^{\infty}\beta^{t}u\left(c_{t}^{A}\right).
$$

And, finally, from homogeneity

$$
g\left(\frac{1-p^{B}}{1-p^{A}}\right)g\left(1+\omega^{lev}\right)E_{0}\sum_{t=0}^{\infty}\beta^{t}u\left(c_{t}^{A}\right)=E_{0}\sum_{t=0}^{\infty}\beta^{t}u\left(\left(1+\omega^{lev}\right)\left(1+\omega^{unc}\right)c_{t}^{A}\right). \textbf{ QED}
$$

Floden (2001) contains a more general proof of additive decomposition for preferences which also depend on leisure, and for the notion of ex-ante welfare which also includes a third component, what we called the egalitarian effect.

3 Bewley Economies with Aggregate Uncertainty

So far we have assumed away aggregate fluctuations (i.e. business cycles) in our description of the incomplete-markets economies à la Bewley-Aiyagari. In the first part of the course, we studied RBC models within a representative agent framework. In this section, the objective is to combine aggregate and idiosyncratic risk into an equilibrium model.

The good news is that we can use the recursive language to do this. The bad news is that solving exactly for the equilibrium allocations of this economy is *impossible*. The reason is that the measure of agents across states becomes an aggregate state of the economy, since households need to know it to forecast prices. But a measure is an infinitely dimensional object, i.e. a function: how do we keep track of such a monster? The answer is that we will approximate the exact equilibrium and argue that the approximation is very good for standard parameterizations.

Aggregate and Idiosyncratic Risk– We introduce aggregate fluctuations through an aggregate productivity shock z_t that shifts the production function, i.e.

$$Y_t = z_t F\left(K_t, H_t\right),$$

and assume that the aggregate shock can take only two values, $z_t \in Z = \{z_b, z_g\}$ with $z_b < z_g$. To keep things simple, we also assume only two values for the individual productivity shock, $\varepsilon_t \in E = \{\varepsilon_b, \varepsilon_g\}$ with $\varepsilon_b < \varepsilon_g$. For example, if $\varepsilon_b = 0$, then it's as if the worker was unemployed for a period.

Let

$$\pi\left(z', \varepsilon' | z, \varepsilon\right) = \Pr\left(z_{t+1} = z', \varepsilon_{t+1} = \varepsilon' | z_t = z, \varepsilon_t = \varepsilon\right)$$

be the Markov chain that describes the joint evolution of the exogenous states. This notation allows the transition probabilities for ε to depend on z (the dependence of z on ε is also allowed in principle but does not make sense!). For example, one should expect that

$$\pi\left(z_b, \varepsilon_g | z_b, \varepsilon_b\right) < \pi\left(z_g, \varepsilon_g | z_b, \varepsilon_b\right), \text{ and } \pi\left(z_b, \varepsilon_g | z_g, \varepsilon_g\right) < \pi\left(z_g, \varepsilon_g | z_g, \varepsilon_g\right)$$

i.e. finding a job is easier if the economy is exiting from a recession and remaining employed is harder if the economy is entering a recession.

State variables– The two individual states are $(a, \varepsilon) \in S$ and the two aggregate states are $(z, \lambda) \in Z \times \Lambda$ where $\lambda(a, \varepsilon)$ is the measure of households across states. The individual states are directly budget relevant, whereas the aggregate states are needed to compute and forecast prices. Note that, although it seems reasonable that λ is enough to complete the description of the state (i.e. a recursive equilibrium exists), this is not obvious at all and there are counterexamples of economies for which one needs to keep track of a longer history of distributions.

Household Problem– The household problem can be written in recursive form as:

$$
\begin{aligned}
v(a,\varepsilon;z,\lambda) &= \max_{c,a'}\left\{u(c)+\beta\sum_{\varepsilon'\in E,z'\in Z}v(a',\varepsilon';z',\lambda')\pi(z',\varepsilon'|z,\varepsilon)\right\} \qquad (14)\\
& \quad s.t.\\
c+a' &= w\left(z,K\right)\varepsilon+\left(1+r\left(z,K\right)\right)a\\
a' &\geq 0\\
\lambda' &= G\left(z,\lambda,z'\right)
\end{aligned}
$$

where $G\left(z,\lambda,z'\right)$ is the law of motion of the endogenous aggregate state and depends on z' since the fraction of agents with $\varepsilon'=\varepsilon_b$ and $\varepsilon'=\varepsilon_g$ next period, given that the current aggregate productivity level is z, depend also on z'.

The key complication is that the value function v depends on λ which is a distribution, i.e. a huge object. Why? To solve the problem, the household needs to compute current prices and, most importantly, forecast future prices. Prices depend on aggregate capital K, and aggregate capital depends on how assets are distributed in the population λ because in equilibrium

$$K=\int_{A\times E}ad\lambda.$$

Since λ is a state variable, we need to know its equilibrium law of motion G which is a monstrous object!!!

A **Recursive Competitive Equilibrium** for this economy is a value function v; decision rules for the household a', and c; policies for the firm H and K; pricing functions r and w; and, a law of motion G such that:

- given the pricing functions $r\left(z,K\right)$ and $w\left(z,K\right)$, the decision rules a' and c solve the household's problem (14) and v is the associated value function,
- given prices, the firm chooses optimally its capital K and its labor H, i.e.

$$
\begin{aligned}
r\left(z,K\right)+\delta &= zF_K(K,H), \qquad (15)\\
w\left(z,K\right) &= zF_H(K,H),
\end{aligned}
$$

- the labor market clears[5]: $H=\int_{A\times E}\varepsilon d\lambda$,

[5] Note that H is not constant any longer. However, the dynamics of H can be perfectly forecasted through π because labor supply is exogenous.

- the asset market clears: $K = \int_{A\times E} a d\lambda$,
- the goods market clears:

$$\int_{A\times E} c(a,\varepsilon;z,\lambda)d\lambda + \int_{A\times E} a'(a,\varepsilon;z,\lambda)d\lambda = zF(K,H) + (1-\delta)K,$$

- The aggregate law of motion G is generated by the exogenous Markov chain π and the policy function a' as follows:

$$\lambda'(\mathcal{A}\times\mathcal{E}) = G_{(\mathcal{A}\times\mathcal{E})}(z,\lambda,z') = \int_{A\times E} Q_{z,z'}((a,\varepsilon),\mathcal{A}\times\mathcal{E})\,d\lambda, \tag{16}$$

where $Q_{z,z'}$ is the transition function between two periods where the aggregate shock goes from z to z' and is defined by

$$Q_{z,z'}((a,\varepsilon),\mathcal{A}\times\mathcal{E}) = I_{\{a'(a,\varepsilon;z,\lambda)\in\mathcal{A}\}}\sum_{\varepsilon'\in\mathcal{E}}\pi(z',\varepsilon'|z,\varepsilon), \tag{17}$$

where I is the indicator function, and $a'(a,\varepsilon;z,\lambda)$ is the optimal saving policy.

3.1 Computation of an Approximate Equilibrium

The state space of the problem of the household is, technically, infinite-dimensional because it contains a distribution. The problem is to find an efficient way to compute the law of motion

$$\lambda' = G(z,\lambda,z').$$

Krusell and Smith (1998) contains the insight that, since we cannot work with an infinitely dimensional distribution, we need to approximate the distribution with a finite-dimensional object. Any distribution can be represented by its entire (in general, infinite) set of moments. Let $\bar{m}$ be a M dimensional vector of the first M moments (mean, variance, skewness, kurtosis,...) of the *wealth distribution*, i.e. the marginal of λ with respect to a. Our new state is exactly the vector $\bar{m} = \{m_1, m_2, ..., m_M\}$ with law of motion

$$\bar{m}' = \Phi(z,\bar{m},z') = \begin{cases} m_1' = \phi_1(z,\bar{m},z') \\ \quad ... \\ m_M' = \phi_M(z,\bar{m},z') \end{cases}. \tag{18}$$

This method is based on the idea that households have *partial information* about λ. They don't know every detail about that measure, but only a set of moments, e.g. its mean, its variance, the Gini coefficient, the share held by the top 5% and so on. Hence, they use these M statistics to approximate the true distribution and form forecasts.

To make this approach operational, one needs to: 1) fix M and 2) specify a functional form for Φ. Krusell and Smith (and this is their main finding) show that one obtains an excellent forecasting rule by simply setting $M = 1$ and by specifying a law of motion of the form:

$$\ln K' = b_z^0 + b_z^1 \ln K,$$

where only the first moment $m^1 = K$ would matter to predict the first moment next period. The new partial-information problem of the agent becomes

$$\begin{aligned} v(a,\varepsilon;z,K) &= \max_{c,a'} \left\{ u(c) + \beta \sum_{\varepsilon',z'} v(a',\varepsilon';z',K')\pi(z',\varepsilon'|z,\varepsilon) \right\} \qquad (19) \\ & \quad s.t. \\ c + a' &= w(z,K)\varepsilon + (1 + r(z,K))\, a \\ a' &\geq 0 \\ \ln K' &= b_z^0 + b_z^1 \ln K. \end{aligned}$$

Note that this state space is definitely manageable: we collapsed an infinitely dimensional distribution into one number. Now, we also know how to solve this problem. It will be a fixed-point algorithm over the law of motion for K: recall that in equilibrium the law of motion used by the agents has to be consistent with the aggregation of the optimal individual decisions ("aggregate consistency" of rational expectation equilibrium).[6]

3.1.1 Algorithm

The algorithm to solve this problem (and the associated equilibrium) is the following:

1. Guess the coefficients of the law of motion $\{b_z^0, b_z^1\}$

2. Solve the household problem and obtain the decision rules $a'(a,\varepsilon;z,K)$, $c(a,\varepsilon;z,K)$.

   Note that with the law of motion for K in hand, we have all we need to solve for

[6] Recall that to solve for the stationary equilibrium, since the distribution is time-invariant, we guess only one value for the capital stock (or the interest rate). To compute the equilibrium transitional dynamics, we guess a deterministic sequence of capital stocks. With aggregate uncertainty, we need to guess a law of motion for the aggregate capital stock.

decision rules. For example, if we iterate on the Euler equation, then we need to compute the rule $a' = g(a, \varepsilon; z, K)$ that solves

$$u_c(R(z,K)a + w(z,K)\varepsilon - g(a,\varepsilon;z,K)) \geq$$
$$\beta E\{R(z',K')u_c(R(z',K')g(a,\varepsilon;z,K) + w(z',K')\varepsilon' - g(a',\varepsilon';z',K'))\}.$$

Thus we can use standard methods to obtain $g(\cdot)$.

3. Simulate the economy for N individuals and T periods. For example, $N = 10,000$ and $T = 2000$. Draw first a random sequence for the aggregate shocks. Next one for the individual productivity shocks for each $i = 1, ..N$, conditional on the time-path for the aggregate shocks. Use the decision rules to generate sequences of asset holdings $\{a_t^i\}_{t=1,i=1}^{T,N}$ and in each period compute the average capital stock

$$K_t = \frac{1}{N}\sum_{i=1}^{N} a_t^i.$$

4. Discard the first T^0 periods (e.g. $T^0 = 500$) to avoid dependence from the initial conditions. Using the remaining sequence, run the regression

$$\ln K_{t+1} = \beta_z^0 + \beta_z^1 \ln K_t \tag{20}$$

and estimate the coefficients $\left(\beta_z^0, \beta_z^1\right)$.

5. If $\left(\beta_z^0, \beta_z^1\right) \neq (b_z^0, b_z^1)$, then try a new guess and go back to step 1. If the two pairs are equal for each $z \in \{z_g, z_b\}$, then it means that the approximate law of motion used by the agents is consistent with the one generated in equilibrium by aggregating individual choices.

6. Recall that this equilibrium computation is approximate: we still need to verify how good this approximation is to the fully rational-expectation equilibrium. For this purpose, compute a measure the fit of the regression in step 4), for example by using R^2. Next, try augmenting the state space with another moment, for example using $m^2 = E(a_i^2)$. Repeat steps 1)-5) until convergence. If the R^2 of the new equation (20) has improved significantly, keep adding moments until R^2 is large and does not respond to addition of new explanatory moments. Otherwise, stop: it means that additional moments do not add new useful information in forecasting prices.

3.2 A Near-Aggregation Result in the Krusell-Smith Economy

Krusell and Smith's main finding is that a law of motion based only on the mean, i.e.,

$$\ln K' = \begin{cases} 0.095 + 0.962 \ln K, & \text{for } z = z_g \\ 0.085 + 0.965 \ln K, & \text{for } z = z_b \end{cases}$$

delivers an $R^2 = 0.999998$ which means that the agents with this simple forecasting rule make very small errors, for example the maximal error in forecasting the interest rate 25 years into the future is around 0.1%. This result is called *near-aggregation* in the sense that in equilibrium, the evolution of aggregate quantities and prices approximately depend only on the aggregate shock and aggregate capital. Hence, it is almost like in a complete-markets economy, where aggregation of heterogenous individuals holds perfectly.

What is the intuition for the fact that keeping track of the mean of the distribution of assets is enough? Recall that if policy functions are of the form

$$a'(a, \varepsilon, z, \lambda) = b_z^0 + b_z^1 a + b_z^2 \varepsilon,$$

then

$$K' = \int_{A \times E} a'(a, \varepsilon, z, \lambda) d\lambda = b_z^0 + b_z^1 K + b_z^2 H_z = \tilde{b}_z^0 + b_z^1 K$$

which would explain why the mean is a sufficient statistic. But saving functions are in general, not linear with uninsurable idiosyncratic shocks. They're exactly linear only with complete markets (recall the exact aggregation result of the Chatterjee economy with homothetic preferences?), where we showed that the distribution does not affect the dynamics of aggregate variables.

So, why do we get *near-aggregation* in practice? For three reasons. First, the saving functions $a'(a, \varepsilon, z, \lambda)$ for this class of problems usually display lots of curvature for low levels of ε and low levels of assets a, but beyond this region they're *almost linear*.Second, the agents with this high curvature are few and have low wealth, so they matter very little in determining aggregate wealth. What matters for the determination of the aggregate capital stock are the ones who hold a lot of capital, i.e., the rich, not the poor! Third, aggregate productivity shocks move the asset distribution only very slightly, and the mass of the distribution is always where the saving functions are linear.

But why do agents have linear saving functions, i.e. a constant marginal propensity to save out of wealth, for a very wide range of the asset space? After all, if they save for precautionary reasons (as they do in these economies) they should do so more when they hold few assets and less when they hold large assets, so the saving function should be nonlinear. The answer is that most of the consumers in this economy can smooth consumption very effectively through self-insurance, by cumulating a relatively small amount of wealth. Thus their saving behavior is guided mostly by their intertemporal motive rather than their insurance motive.

To understand why the risk-free asset is such a good vehicle of self-insurance in the Krusell-Smith model, we discuss here two theoretical results in the literature that can help explain it.

First, Yaari (1976) analyzed the optimal consumption path of a perfectly patient household ($\beta = 1$) with general concave preferences (hence with prudence) who lives for T periods, faces *iid* endowment shocks and saves and borrows at rate $r = 0$. Yaari shows that as $T \to \infty$, the optimal consumption plan converges to that of a consumer who eats a constant fraction of his wealth every period. In this sense, these households behave like certainty-equivalent consumers who are not concerned about future risk. Recall that an agent with quadratic preferences has consumption determined by

$$c_t = \frac{r}{1+r}\left[a_t + E_0 \sum_{j=0}^{\infty}\left(\frac{1}{1+r}\right)^j y_{t+j}\right] = \frac{r}{1+r}a_t + H(\mathbf{y},r),$$

where $\mathbf{y}$ is its whole future income sequence. His assets next period are determined by

$$\begin{aligned} a_{t+1} &= (1+r)(a_t - c_t + y_t) \\ &= (1+r)\left[a_t - \frac{r}{1+r}a_t + H(\mathbf{y},r)\right] + (1+r)y_t \\ &= a_t + \tilde{H}(\mathbf{y},r) \end{aligned}$$

and note that the coefficient on past wealth is exactly one, which is very close to the one computed by Krusell-Smith.

Second, Levine and Zame (2001) analyze an economy populated by infinitely-lived consumers with standard preferences satisfying $u''' > 0$ who face stationary individual endowment shocks (i.e., not random-walk) and trade a risk-free asset in zero net supply. They prove that, as $\beta \to 1$ and the individuals become perfectly patient, "market incompleteness will not matter" in the sense that the welfare of the optimal consumption plan in this economy tends to the welfare of a complete markets economy where every agent consumes her average endowment every period. In other words, a great deal of risk-sharing may take place even in absence of a complicated structure of financial markets.[7]

Finally, at this point it is not surprising that Krusell and Smith find that the cyclical properties of aggregates in their model economy (i.e. volatility of output, consumption, investment, cross-correlations, etc...) are very similar to those of the standard representtaive agent model.

References

[1] Krusell P., and A. Smith (1998);. "Income and Wealth Heterogeneity in the Macroeconomy," *Journal of Political Economy*, vol. 106(5), 867-896.

[2] Yaari, M. (1976); "A Law of Large Numbers in the Theory of Consumer's Choice under Uncertainty," *Journal of Economic Theory*, 12, 202-217.

[3] Levine, D., and W. Zame (2001); "Does Market Incompleteness Matter?," *Econometrica*, vol. 70(5), 1085-1839.

[4] Floden, Martin (2001), The Effectiveness of Government Debt and Transfers as Insurance, *Journal of Monetary Economics*

[5] Heathcote, Jonathan (2004), Fiscal Policy with Heterogeneous Agents and Incomplete Markets, *Review of Economic Studies*

[6] Kubler, Felix and Karl Schmedders, (2002), Recursive Equilibria in Economies with Incomplete Markets, Macroeconomic Dynamics, 6: 284-306.

[7] Note however, that Levine and Zame show that their result holds in the presence of aggregate uncertainty only if markets are complete with respect to aggregate risk. In the Krusell-Smith economy there is no insurance against aggregate risk.