Workshop BCB: Macro com agentes heterogêneos

Aula 3 e 4: HH Heterogeneity: Transition Dynamics and Aggregate Fluctuations

Tomás R. Martinez

Insper

Introduction

- At this point, we have only focus on the **Stationary Equilibrium**.
- Many questions involve solving the model beyond the Steady-State Stationary Equilibrium.
- **Aggregate Uncertainty:**
  - How the Aiyagari economy reacts to aggregate shocks.
  - Does heterogeneity matters to the business cycles?
- **Transitions Dynamics:**
  - How long it takes to the economy reach a new steady state after an economic reform.
  - How to compute the transition from one steady state to another.

References

- Boppart, Krusell and Mitman (2018, JEDC)*: Intuitive paper on how transition dynamics can be used to simulate aggregate shocks (+ history about the MIT shocks).
- Auclert, Bardóczy, Rognlie and Straub (2021, ECTA)*: State-of-the-art method to solve HA models with aggregate uncertainty.
- Krusell and Smith (1998, JPE): original paper outlining the famous algorithm.
- Krueger, Mitman and Perri (2016, Handbook of Macro): Application of the model to the great recession.
- Heer and Maussner (2009): Ch. 8 and 10; Fehr and Kindermann (2019): Ch. 11: Textbook treatment of the computational methods.
- Algan et al (2014, Handbook of Computational Economics): Entire handbook on how to solve HA economies with aggregate uncertainty. See also their special edition on the JEDC.

Aiyagari + Aggregate Uncertainty

- We are going to focus in the simplest version of the Aiyagari model with aggregate uncertainty.

- The only modification is an aggregate TFP shock in the production funcion:

$$
\begin{aligned}
Y_t &= Z_t K_t^{\alpha} L_t^{1-\alpha} \\
\log Z_t &= \rho_z \log Z_{t-1} + \sigma_z \varepsilon_t
\end{aligned}
$$

- Sometimes this is known as the **Krusell-Smith** economy.

Aiyagari + Aggregate Uncertainty

- Prices are allowed to vary over the cycles.
- To write in recursive form, we also include the aggregate variables as a state in the HH problem.
  - Individual state: (a, s), aggregate state: (Z, λ).

$$
\begin{aligned}
V(a, s; Z, \lambda) &= \max_{c,\ a' \geq 0} \{u(c) + \beta \mathbb{E}[V(a', s'; Z', \lambda') | s, Z]\} \\
\text{s.t} \qquad c + a' &= w(Z, \lambda) \exp\{s\} + (1 + r(Z, \lambda) - \delta) a \\
\lambda' &= H(Z, \lambda, Z')
\end{aligned}
$$

  - Note the dependence of prices on the distribution.
  - The function H is the law of motion/forecasting function of the distribution.

Equilibrium

- Prices are given by FOCs of firm's problem:

$$w(Z,\lambda) = Z\alpha\left(\frac{K(Z,\lambda)}{N(Z)}\right)^{1-\alpha} \quad \text{and} \quad r(Z,\lambda) = Z(1-\alpha)\left(\frac{N(Z)}{K(Z,\lambda)}\right)^{\alpha}$$

  where aggregate employment $N(Z)$ is given by the distributions of the Markov process (which may depend on Z).

- Asset market clears:

$$K(Z,\lambda) = \int a d\lambda$$

- The distribution evolves according to the function: $\lambda' = H(Z,\lambda,Z')$. In equilibrium, with rational expectation, this function is *consistent* with the individual decisions.

HA with Aggregate Uncertainty

Hard problem to solve:

- Prices are functions of the distribution, so distribution must be part of the state space.
- But the distribution is infinitesimal object with a lot of information.
- Furthermore, agents must forecast the evolution of the distribution to form expectations.
- And the forecast has to be consistent with the individuals decision (i.e., fixed point).

Computational

To solve a heterogeneous agent economy with aggregate uncertainty the main methods are:

- **State-space Methods:**
  - Krusell-Smith (1998, JPE) bounded rationality algorithm.
  - Reiter (2009, JEDC) Method.
- **Sequence-space Methods:**
  - MIT shock (Boppart, Krusell and Mitman, 2018, JEDC).
  - Auclert, Bardóczy, Rognlie and Straub (2021, ECTA) sequence space Jacobian.
- There are others/variations of algorithms and evolutions of the original ones. Check Algan et al (2014).

Krusell & Smith (1998) Bounded Rationality

- References: Krusell-Smith's original paper is easy to follow. Check also Nakajima's notes.
- Because prices are allowed to vary over the cycles and they are needed for the household problem: the aggregate state, (Z, λ), is part of the state of the HH.
  - **Problem:** the distribution, λ, is a high-dimensional object and the state space increases substantially.
- **Krussel & Smith (1998)**: instead of using the entire distribution, just use some moments of the distribution:
  - Households are "boundedly rational" on how the distribution evolves.
  - In this class of models, the **mean** (first moment) is enough to correctly forecast prices:

$$\lambda' = H(Z, \lambda, Z') \qquad \Rightarrow \qquad K' = H(Z, K, Z') \tag{1}$$

Krusell & Smith (1998)

- Substitute λ by K. Example:

$$
\begin{aligned}
V(a, s; Z, K) = \max_{c,\ a' \geq 0} & \left\{ u(c) + \beta \mathbb{E}[V(a', s'; Z', K') | s, Z] \right\} \\
\text{s.t} \qquad & c + a' = w(Z, K) \exp\{s\} + (1 + r(Z, K) - \delta) a \\
& K' = H(Z, K, Z')
\end{aligned}
$$

- **Intuition:** the mean of λ works well to forecast prices because the savings policy function is approximately linear.
  - The curvature of the policy function is close to the borrowing constraint, but these agents hold little wealth and thus do not matter to the aggregate.

- For more complex models, one may need higher moments.

Krusell-Smith Algorithm

- Suppose Z is a two-state Markov-Chain (recession and boom).
- Approximate the function forecasting function $H()$ with a log-linear form:

$$\log K' = a_l + b_l \log K \quad \text{if } Z = Z_l$$
$$\log K' = a_h + b_h \log K \quad \text{if } Z = Z_h$$

- We have to find the parameters: (a_l, a_h, b_l, b_h).
- As any continuous state, we must discretize K so we must interpolate when applying the function above.

Krusell-Smith Algorithm

Discretize the state space: (a, s, K, Z). Recover the prices $r(K, Z)$ and $w(K, Z)$ for each state space using the firm's problem.

(i) Guess the parameters of the forecast function: $(a_l^0, a_h^0, b_l^0, b_h^0)$.

(ii) Given $(a_l^0, a_h^0, b_l^0, b_h^0)$, solve the Bellman Equation of the HH for all the state space (a, s, K, Z).

(iii) Given the household policy functions, simulate T periods:

- Draw a sequence of Z_t for all T. Guess a initial distribution λ_0.
- Using the policy function and the sequence Z_t, keep updating the distribution λ_t forward.
- Compute the mean of the distribution K_t (and other moments if necessary).
- Drop the first T_0 periods. Now, we have a sequence $\{Z_t, K_t\}_{t=T_0}^{T}$.

Krusell-Smith Algorithm

(iv) Using the sequence $\{Z_t, K_t\}_{t=T_0}^{T}$, run a linear regression and recover the new coefficients: $(a_l^1, a_h^1, b_l^1, b_h^1)$.

(v) Check the distance between the guess a^0, b^0 and the new parameters a^1, b^1. If it is smaller than tol, we are done. Otherwise, update the guess and start again:

$$a^0 = d\ a^0 + (1-d)a^1$$
$$b^0 = d\ b^0 + (1-d)b^1$$

where $d \in (0,1)$ is a damping parameter.

Krusell-Smith Algorithm: Issues

- After you finish, you must check the R^2 of the forecast regression. If the R^2 is low, you must add more moments or change the function form.
  - In Krusell-Smith, $R^2 = 0.999$, so the perceived law of motion of K is very close to the actual law of motion.

- Poor initial guesses might not converge. One good guess is $a = \log K_{ss}$ and $b = 0$.

- **Good:** KS captures potential non-linearities and large shocks. For instance, asymmetries between the boom and the recession; uncertainty shocks; etc.

- **Bad:** KS can be inaccurate if there are explicitly distributional channels coming from the top of the wealth distribution. Potentially very slow.

Krussell-Smith: State of the Art

- If you need to solve a HA model using a truly global method, the state-of-the-art nowadays is to use **Deep learning/machine learning**:
- See Fernandéz-Villaverde, Hurtado & Nuño (ECTA, 2023); Azinovic et al (IER, 2022); Maliar et al (2021); and other papers by Fernandéz-Villaverde and Galo Nuño.

Reiter Method: Projection + Pertubation

- **Perturbation Methods**:
  - Generalization of the well-known linearization around the steady state.
  - Often used to solve DSGE/representative agent models.
  - They tend to be fast, but require derivatives and some stability conditions (Blanchard-Kahn).

- Standard software (i.e., dynare) uses this method.

- Reiter (2009) propose to solve for the stationary equilibrium using global methods (projection methods), and then use perturbation methods to solve for the aggregate shock.

- If you need a refresher on Pertubation methods, check Fernandez-Villaverde's notes.

Reiter's Method

- We can write the solution of DSGE models as a nonlinear system of difference equations:

$$E_t F(x_t, x_{t+1}, y_t, y_{t+1}) = 0 \tag{2}$$

  where x is the vector of predetermined variables (state), y is nonpredetermined variables (control).

- Then, we can linearize the system (either numerically or analytically) and use methods to solve the linear system of difference equations:
  - Blanchard and Kahn (1980); Uhlig (1999); Sims (2000); Rendahl (2018).

Reiter's Method

- **Example**: Stochastic Neoclassical Growth model

$$E_t F(x_t, x_{t+1}, y_t, y_{t+1}) = E_t \begin{bmatrix} c_t^{-\gamma} - \beta E_t c_{t+1}^{-\gamma}[\alpha k_{t+1}^{\alpha-1} + 1 - \delta] \\ c_t + k_{t+1} - e^{z_t} k_t^{\alpha} - (1-\delta)k_t \\ z_{t+1} - \rho z_t - \sigma \varepsilon_{t+1} \end{bmatrix} = 0$$

where $x = [k, z]'$ and $y = [c]$.

- First row is the Euler Equation, second is the feasibility constraint, and the last is the stochastic process of the shock.

Reiter's Method

- **Example**: Krusell-Smith economy.

$$E_t F(x_t, x_{t+1}, y_t, y_{t+1}) = E_t \begin{bmatrix} \lambda_{t+1} - \lambda_t \Pi_{g_{a,t}} \\ V_t - (\overline{u}_{g_{a,t}} + \beta \Pi_{g_{a,t}} V_{t+1}) \\ z_{t+1} - \rho z_t - \sigma \varepsilon_{t+1} \\ \text{ED}(g_{a,t}, \lambda_t, z_t, P_t) \end{bmatrix}$$

where $x = [\lambda, z]'$ and $y = [V P]'$.

- λ is the p.d.f of the distribution;
- P_t are the prices;
- ED() is an arbitrary excess demand function (which implicitly includes firm's foc);
- $\Pi_{g_{a,t}}$ is the transition matrix induced by the optimal policy:

$$g_{a,t} = \arg\max u(a(1 + r_t) + w_t s - a') + \beta E_t V_{t+1}(a', s', \lambda', z')$$

Reiter's Method

- Since we discretize both λ and V_t, the first two rows must hold for ALL the idiosyncratic state.
- The number of equations that we need to linearize is exponentially large.
- Linearization is often done using numerical derivatives. Nowadays people use automatic differentiation to do the job.
- Solution (up to first order) has **certainty equivalence**: no precautionary savings because of aggregate risk.
- The method cannot capture nonlinearities or sign asymmetries (again up, to first order).

Reiter's Method: State-of-the-art

- **Good:**
  - Similar to standard methods using in RA-DSGE (some people argue that is possible to do it in Dynare).
  - Easier to do second-order approximations than the sequence-space methods and faster than fully global methods.
- **Bad:**
  - Tend to be quite hard to implement because they require some type of dimensionality reduction to be fast.
  - Numerical derivatives can be unstable when mixed with discretization.
- **State-of-the-art**:
  - Bayer and Luetticke (QE, 2020): The codes are available in their website (Matlab, Python and Julia): https://www.ralphluetticke.com/.
  - Ahn, Kaplan, Moll, Winberry and Wolf (NBER Macro, 2018); Bhandari et al. (2023) - Second and higher-order approx; Winberry (QE, 2018) - HA Firms, but implemented in Dynare; Bilal (2023).

Sequence Space

- Instead of including the aggregate variables in the state-space, we can index everything through time: $\{r_t, w_t, V_t, \lambda_t, ...\}_{t=0}^{T}$.

- Then, we solve the model in the **Sequence space** from $t \in \{0, .., T\}$, where T is a large number.

- For instance, we can simulate an **impulse response function** (IRF), which is just a deterministic transition dynamics between two steady states after an unexpected aggregate shock (a MIT shock).

- **Boppart, Krusell and Mitman (2018)** show that the IRF can be used to compute equilibrium of HA with agg. uncertainty.
    - Solving for the transition dynamics is also useful if you are interested in studying the transition to a new steady state after a change in economic policy.

Transition Dynamics and MIT shocks

- **MIT shock:** an unpredictable shock to the steady-state equilibrium of an economy without shocks.
  - No shocks are expected to ever materialize but nevertheless a shock now occurs!

- We can now analyze the equilibrium transition along a **perfect-foresight path** until the economy reaches the steady state.

- Some argue that **Tom Sargent** coined the term reflecting that some researchers at MIT used the method.
  - For fresh-water economists, a MIT shock is inconsistent with rational expectations!
  - "A shock of probability zero, only at MIT they can get away with that!".

MIT shock

- Suppose a standard Ayiagari in the steady state at $t = 0$. At $t = 1$, the economy receives an (unexpected) TFP aggregate shock:

$$
\begin{aligned}
Y_t &= Z_t K_t^{\alpha} L_t^{1-\alpha} \\
\log Z_t &= \rho_z \log Z_{t-1} + \varepsilon_t
\end{aligned}
$$

  where $\varepsilon_t = 0.01$ if $t = 1$ and $\varepsilon_t = 0$ otherwise.

- If $0 < \rho_z < 1$, when $t \to \infty$, the shock vanishes and we are back to the original steady state.

- Our goal is to solve the **transition dynamics** between the two steady states.
  - Because Z_t varies in the transition, aggregate variables (prices, savings, distribution) change during the transition.

Sequential Equilibrium

- Instead of carrying the aggregate state, we index the Bellman Equation by time t.

$$V_t(a,s) = \max_{c,\ a' \geq 0} \left\{ u(c) + \beta \sum_{s' \in S} \pi(s'|s) V_{t+1}(a' s') \right\}$$

$$\text{s.t} \qquad c + a' = w_t s + (1 + r_t - \delta) a$$

- The distribution follows the L.O.M: $\lambda_{t+1} = \Pi_{g_{a,t}} \lambda_t \quad \forall t$.

- The asset market must clear for all t:

$$\int_{A \times S} a d\lambda_t(a, s; r_t) \equiv A_t(r_t) = K_t(r_t)$$

both the distribution, $\lambda_t(a,s)$, and the aggregate capital, K_t, are indexed by t.

IRF: Standard Aiyagari Economy

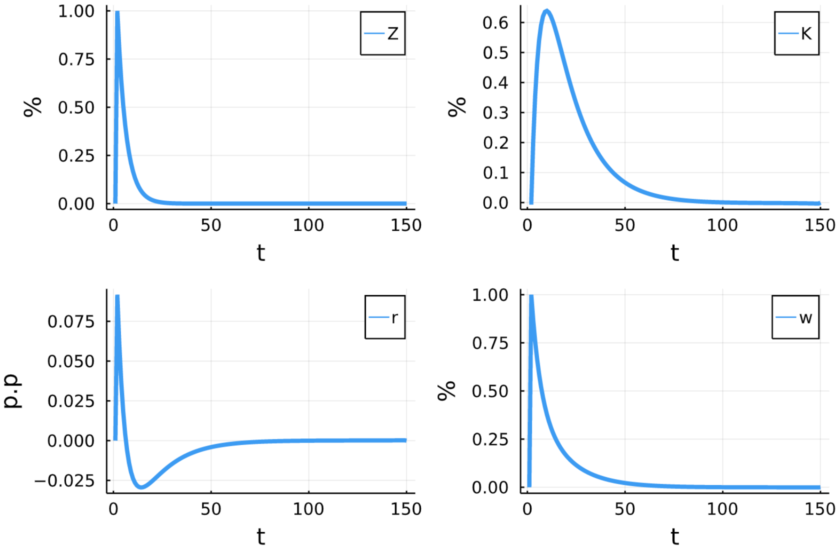

IRF: Standard Aiyagari Economy

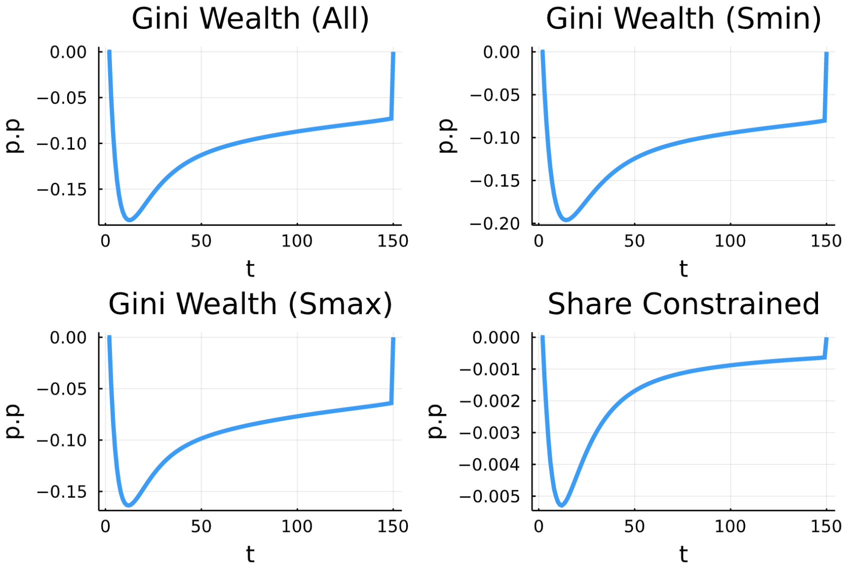

Transition Dynamics between Steady States

- The method is useful to compute transition between different steady states.

- **Example**: Suppose a labor tax, τ_l, that is used to finance a lump-sum transfer, T_t. The budget constraint:

$$c + a' = w_t s(1-\tau_l) + (1 + r_t - \delta)a + T_t.$$

  The government runs a balanced budget: $T_t = \tau_l w_t L$.

- Suppose the economy is in the SS with $\tau_l = 0$. At $t = 1$, the government decides to raise the tax rate: $\tau_l = 0.2$ (there are no aggregate shocks).

- How long does the economy take to reach the new steady state?

Transition to New SS: Labor Tax

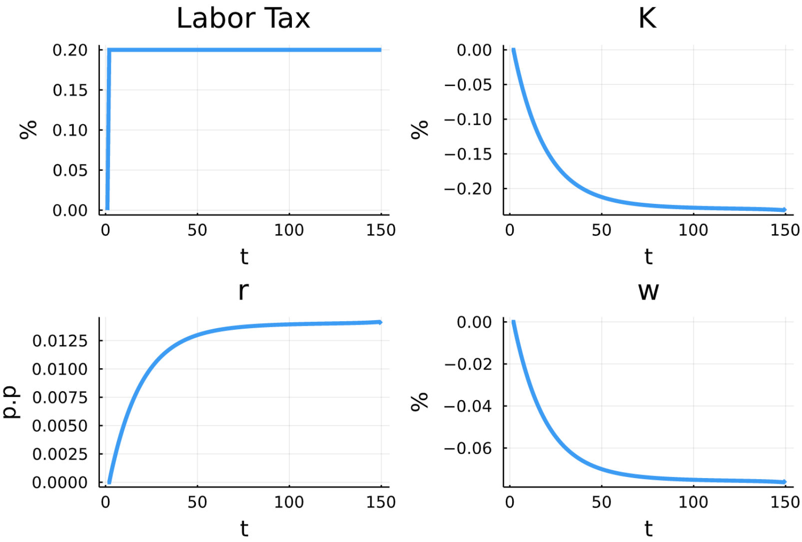

Algorithm

(i) Solve for the initial and the final steady state. Select a large number of periods T.

(ii) Guess a path of $\{K_t^g\}_{t=2}^{T-1}$. K_1 and K_T are given by the initial/final steady state. Recover the prices $\{r_t, w_t\}_{t=2}^{T-1}$ using the firm's problem and the sequence of Z_t.

(iii) Given prices, $\{r_t, w_t\}_{t=2}^{T}$, solve the value function (and policy functions) backwards from $t = T-1, ...2$ starting from the **final steady state value function**.
   - Endogenous Grid works well, but careful to use the correct prices!

(iv) Starting from the **initial steady state distribution**, simulate the distribution forward from $t = 1, ..., T-1$ using the policy functions, $g_{a,t}(a, s)$ and the Markov process of s.

Algorithm

(v) Compute aggregate savings (capital) using the distribution for all t: $\{K_t^s\}_{t=2}^{T-1}$.

(vi) Compute the maximum difference between the guess sequence, $\{K_t^g\}$, and the new sequence, $\{K_t^s\}$. If it is smaller than tol, stop. Otherwise, update the guess using the rule:

$$K_t = dK_t^s + (1-d)K_t^g \qquad \text{for } t = 2, .., T-1,$$

where $\lambda \in (0,1)$ is a dampening parameter, and return to (ii).

Algorithm

- The "shooting algorithm" does not have established convergence properties but tends to work well in practice.
- The damp parameter should not be too large, otherwise, it may not converge.
- T has to be large enough to allow the shock to fade out completely. Always check the last transition between times $T-1$ and T.
- A good initial guess is $K_{ss} = K_t$ for all t.
- If labor supply is endogenous you can guess K/L. If you need to find the eq. in other markets you have to guess an additional sequence.

Boppart-Krusell-Mitman (2018)

- Intuitively, the method uses the impulse response function as a sufficient statistic to compute the eq. of the model.
- In theory, dynamic programming says that any aggregate statistic of the model can be computed as a function of the aggregate state: $x(Z, \lambda)$.
- Instead of using aggregate state, we can also write the aggregate stats as a function of past shocks. For example, the aggregate capital at time t is:

$$K_t = K(\varepsilon_t, \varepsilon_{t-1}, \varepsilon_{t-2}, ...),$$

where ε_t is the innovation of the aggregate at time t.

Boppart-Krusell-Mitman (2018)

- If we assume that the model response to the shock is approximately linear, we can write K_t as a linear function of past shocks:

$$K_t = K(\varepsilon_t, \varepsilon_{t-1}, \varepsilon_{t-2}, ...) = \varepsilon_t K(1, 0, 0, ...) + \varepsilon_{t-1} K(0, 1, 0, ...) + \varepsilon_{t-2} K(0, 0, 1, ...) + ...$$

  where $K(0, 1, 0, ...)$ is the (non-linear) response of capital at time t to a shock (scaled to 1) that happened at $t-1$.

- Note that each K is the response of ONLY ONE shock at each point in time.

- In the notation of BKM: $K_0 = K(1, 0, 0, ...)$, $K_1 = K(0, 1, 0, ...)$, etc. Then:

$$K_t = \sum_{s=0}^{\infty} \varepsilon_{t-s} K_s$$

Boppart-Krusell-Mitman (2018)

- When we compute an impulse response function to an MIT shock, we get exactly the response of capital to a 1% shock that happened s periods before!

- That is, we have a sequence of K:

$$[K(1,0,0,...), K(0,1,0,...), K(0,0,1,...), ...]$$

- We have that for all aggregate statistics of the model: $C,\ A, ...$

- To simulate the model, we can simply draw a sequence of shocks ε and use the statistics computed by the impulse response.

Boppart-Krusell-Mitman (2018)

Figure: Simulation of Aggregate Capital using BKM

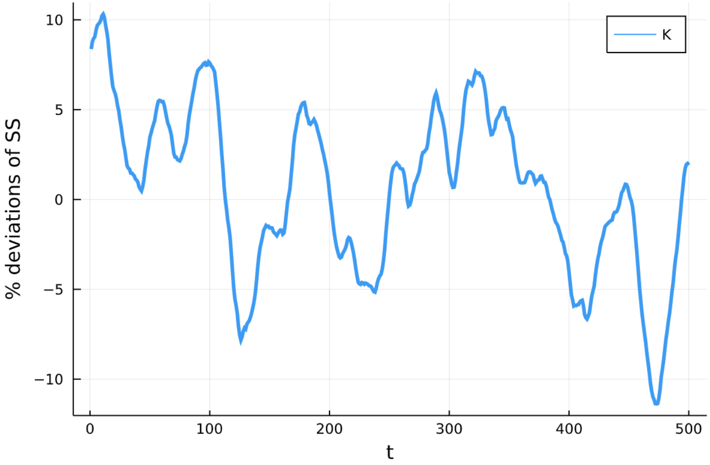

Boppart-Krusell-Mitman (2018)

- **Good:** It is easy to use. The only thing you need is an impulse response function. You can compute using standard dynamic programming methods.

- It is trivial to add more shocks. Because shocks are linear, you just need to simulate two IRF for each shock. Then, the final effect of the shocks is simply additive.

- **Bad:** If the model is highly non-linear or has sign-dependence it can be a poor approximation.

- As every other linear method, it assumes certainty equivalence. No second-order effects from aggregate risk; It may perform poorly if the shock brings you far from the steady state.

Sequence-Space Jacobian

- **Auclert, Bardóczy, Rognlie and Straub (2021)**. Using the Sequence-Space Jacobian to Solve and Estimate Heterogeneous-Agent Models.

- Instead of solving for the full transition, they show that the Jacobian (the derivatives of perfect-foresight) of the equilibrium is enough.

- They provide a very efficient method to compute these Jacobians, and show that by composing and inverting the Jacobians we can solve for the GE of the model very fast.
  - Check their NBER lecture notes at: here.
  - Also the notes of Jeppe Druedahl: here.
  - Python notebooks with plenty of examples are available: here.

Sequence-Space Jacobian

- Their idea is that we can write the model in **blocks** and draw it as directed acyclic graphs (DAGs).
- A block is a part of the model that can be solved independently of the other parts.
  Example:
  - **Household Block** $\Rightarrow$ takes as given sequences of prices/policies (interest rates, wages, tax policies) and output sequences of aggregate consumption, savings, etc.
- Every block takes a sequence of inputs and outputs.
- The model is a combination of household block, firm block, government block, equilibrium block, etc.

Sequence-Space Jacobian

- Denote sequences of variables, e.g. Z_t, as vectors $\mathbf{Z} = (Z_0, Z_1, ...)$.
- **Example:** Krusell-Smith Model $\rightarrow$ Exogenous: $\mathbf{Z}$, Endogenous: $\mathbf{K}$.
  - **Firm's Problem**: $\mathbf{Z}, \mathbf{K} \longrightarrow \mathbf{r}, \mathbf{w}$.
  - **Household's Problem**: $\mathbf{r}, \mathbf{w} \longrightarrow \mathbf{C}, \mathbf{A}$ (where $\mathbf{C}$ and $\mathbf{A}$ are vectors of aggregate consumption and savings, e.g., $C_t = \int g_{c,t}(a,s) d\Phi_t$).
  - **Market Clearing**: $\mathbf{A}, \mathbf{K} \longrightarrow \mathbf{H} \equiv \mathbf{A} - \mathbf{K}$ (assets mkt clearing, alternatively we could have used the goods mkt).
- **Equilibrium:** There is a sequence $\mathbf{K}$, that clears the market, $\mathbf{H} = 0$, in all periods t given the sequence of exogenous variable $\mathbf{Z}$.

Block Representation of Krusell-Smith Model

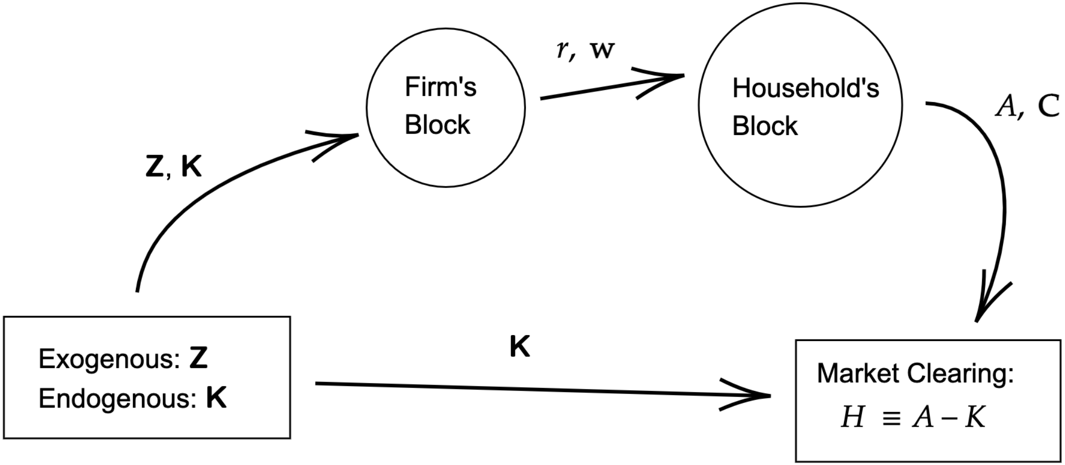

Capital Response to Shocks

- Goal is to solve for market equilibrium given a sequence of exogenous shocks. In our example: $\mathbf{H} \equiv \mathbf{A} - \mathbf{K}$.
  - The sequence of aggregate savings, $\mathbf{A} = (A_0, A_1, ...)$, is a function of the entire sequences of interest rate, $\mathbf{r}$, and, $\mathbf{w}$. Further, $\mathbf{r}$, and wage, $\mathbf{w}$ are functions of the sequences of shock, $\mathbf{Z}$, and capital, $\mathbf{K}$.
  - Also, for every t, aggregate savings is a function of the **entire sequences** $\mathbf{Z}$ and $\mathbf{K}$. Then:

$$A_t(\mathbf{r}, \mathbf{w}) = A_t(\mathbf{Z}, \mathbf{K}) \tag{3}$$

- The equilibrium condition in period t is:

$$H_t(\mathbf{Z}, \mathbf{K}) = A_t(\mathbf{Z}, \mathbf{K}) - K_t$$

- The sequences of equilibrium conditions are: $\mathbf{H}(\mathbf{Z}, \mathbf{K}) = \mathbf{A}(\mathbf{Z}, \mathbf{K}) - \mathbf{K}$.

Capital Response to Shocks

- Auclert et al (2021) $\Rightarrow$ we don't need to solve for the entire equilibrium sequence to recover the response of $\mathbf{K}$ to $\mathbf{Z}$. Just need to look **Jacobians**.

- From the implicit function theorem, the linear impulse response of $\mathbf{K}$ to a transitory technology shock $d\mathbf{Z} = (dZ_0, dZ_1, ...)'$ is:

$$d\mathbf{K} = \mathbf{H}_{\mathbf{K}}^{-1}\mathbf{H}_{\mathbf{Z}}d\mathbf{Z}$$

  where $\mathbf{H}_{\mathbf{K}}$ and $\mathbf{H}_{\mathbf{Z}}$ are the Jacobians of $\mathbf{H}$ with respect to $\mathbf{K}$ and $\mathbf{Z}$, evaluated at the steady state.

- Once we have $d\mathbf{K}$, we can easily compute the response of other variables.

The Jacobians

- To compute $\mathbf{H_K}$ and $\mathbf{H_Z}$, we may have to use the chain-rule.

- For example, the eq. response to $\mathbf{Z}$ is the response of $\mathbf{A}$ to changes in $\mathbf{r}$ and, $\mathbf{w}$, which further respond to $\mathbf{Z}$. We can write as a composite of Jacobians:

$$\mathbf{H_Z} = \mathbf{J}^{A,r} \cdot \mathbf{J}^{r,Z} + \mathbf{J}^{A,w} \cdot \mathbf{J}^{w,Z}$$

  where $\mathbf{J}^{A,r}$ is the Jacobian of $\mathbf{A}$ to $\mathbf{r}$, and so on.

- The Jacobians of $\mathbf{H}$ are just the chain-rule of each **model blocks' Jacobians** ($\mathbf{J}$).

The Jacobians

- What the Jacobians look like? Depends how complicated are model blocks.
- Some are very simple, some are complicated. The "Representative firm block" is simple.
- **Example:** $\mathbf{w}$ only depends on the contemporaneous $\mathbf{Z}$.
  - $w_t = (1-\alpha)Z_t\left(\frac{K_t}{N_t}\right)^{\alpha}$. Then, the Jacobian is:

$$\mathbf{J}^{w,Z} = \begin{bmatrix} \frac{\partial w_0}{\partial Z_0} & \frac{\partial w_0}{\partial Z_1} & \cdots & \frac{\partial w_0}{\partial Z_T} \\ \vdots & \ddots & \ddots & \vdots \\ \frac{\partial w_T}{\partial Z_0} & \frac{\partial w_T}{\partial Z_1} & \cdots & \frac{\partial w_T}{\partial Z_T} \end{bmatrix} = \begin{bmatrix} (1-\alpha)\left(\frac{K_0}{N_0}\right)^{\alpha} & 0 & \dots & 0 \\ \vdots & \ddots & \ddots & \vdots \\ 0 & 0 & \dots & (1-\alpha)\left(\frac{K_T}{N_T}\right)^{\alpha} \end{bmatrix}$$

  - Note that we can exploit the sparsity of the matrix.

Simple Jacobian

- $\mathbf{J}^{Y,Z} = K_t^{\alpha} N_t^{1-\alpha}$ in the diagonal.

```
J_{Y,Z}:
 [[2.28364649 0.         0.         0.         0.        ]
 [0.         2.28364649 0.         0.         0.        ]
 [0.         0.         2.28364649 0.         0.        ]
 [0.         0.         0.         2.28364649 0.        ]
 [0.         0.         0.         0.         2.28364649]]
```

The Jacobians

- The household Jacobian is complicated. Since the EE is forward looking, future shocks are anticipated by the household..
- **Example:** $\mathbf{A}$ depends on the entire path of $\mathbf{w}$.
  - Household changes its behavior in time t, once she understands her earnings change in time $t+s$.
  - Since A_t is aggregate savings, we just need that *some* households change their behavior to change A_t.

$$\mathbf{J}^{A,w} = \begin{bmatrix} \frac{\partial A_0}{\partial w_0} & \frac{\partial A_0}{\partial w_1} & \cdots & \frac{\partial A_0}{\partial w_T} \\ \vdots & \ddots & \ddots & \vdots \\ \frac{\partial A_T}{\partial w_0} & \frac{\partial A_T}{\partial w_1} & \cdots & \frac{\partial A_T}{\partial w_T} \end{bmatrix}$$

  - Matrix is not sparse anymore.

HA Jacobian

- $\mathbf{J}^{C,r}$ has intertemporal effects.

```
J_{C,r}:
 [[ 0.53172749 -1.22045882 -1.14979648 -1.08450169 -1.0237652 ]
 [ 0.52752393  0.600944   -1.15779184 -1.0919005  -1.03060995]
 [ 0.52334341  0.59396873  0.66245358 -1.10087911 -1.03888918]
 [ 0.51918469  0.58763628  0.65396671  0.71839458 -1.04871115]
 [ 0.51504555  0.5815717   0.64610874  0.70874635  0.76968641]]
```

- Each consumption response ($\frac{\partial C_i}{\partial r_j}$) is an element of the matrix:
  - If the increase of r happened in the past ($j \leq i$): consumption increases $\Rightarrow$ wealth effect changes the distribution.
  - If the increase of r will happen in the future ($j > i$): consumption decreases (savings increase) $\Rightarrow$ substitution effect.

HA Jacobians

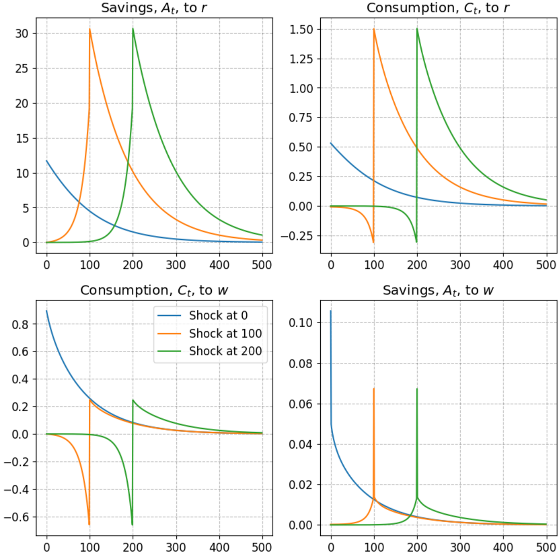

Fake News Algorithm

- **Problem:** Computing the Jacobians can be very costly $\Rightarrow$ It requires backward (policy function) and forward (distribution) for every $\mathbf{J}$.

- Auclert et al (2021) develops an algorithm based on "news shocks" (i.e., learning today that future income increases) $\Rightarrow$ Fake News Algorithm.

- **Intuition:**
  - Only the difference between two periods matter (not the actual t) for policy functions $\Rightarrow$ a single backward iteration is sufficient.
  - For the effect through the distribution, they use a "Fake News" shock: a shock in period s announced in $t = 0$ but retracted at $t = 1$.
  - Using tedious algebra and the chain-rule they can construct all the Jacobians fast.

Impulse Response Function

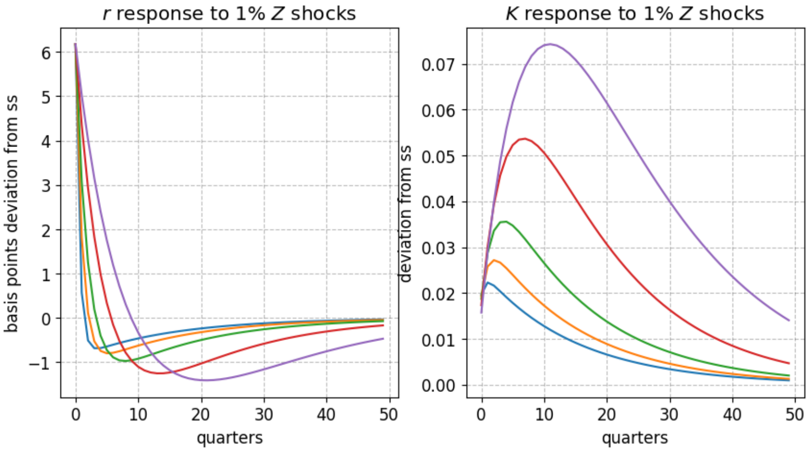

Non-linear Solutions

- The Jacobians give a linearized IRF. They can be imprecise for large shocks or in models with aggregate non-linearities.

- The package also give an algorithm to compute the **nonlinear perfect foresight dynamics** (i.e., the MIT shock).

- The idea is to use the fact that an equilibrium must solve: $H(K, Z) = 0$, iterate in a sequence of K^j, where j is the guess of K, and update using:

$$K^{j+1} = K^j - H_K(K_{ss}, Z_{ss})^{-1} H(K^j, Z)$$

- Note this is very similar to a Newton Algorithm, which in practice has very fast convergence.

Sequence-Space Jacobian

- Once we have the Jacobians of each model block, we can compute the response to any type of shocks, IRF, or transition dynamics for a new SS.
- The key is to compute the Jacobians efficiently.
- The algorithm allows us to solve even very complex HANK models.
- It can also be applied to more general models (entry-exit, discrete choices, etc), but some details must be taken care of.
- **Limitations**:
  - $\Rightarrow$ models where the Bellman equation depends directly on the distribution (e.g., wage posting search models).
  - $\Rightarrow$ solving the stationary equilibrium can be costly in some models, must apply some tricks to speed up this step.