Transitional Dynamics and Welfare Evaluation

Manuel Macera | 2026

Where we are

- Last lecture: in a stationary equilibrium the distribution is a state variable, but it does not move. Equilibrium reduces to one price that clears one market.

- Now the distribution **moves**, for one of two reasons:
  - a one-time change in policy or fundamentals moves the economy from one steady state to another **(this lecture)**;
  - aggregate shocks hit the economy every period (next lecture).

- Once λ moves, households must forecast prices, so they must forecast λ. This is the hard problem we postponed.

Plan

1. **Transitions**: an unexpected change in policy or fundamentals, equilibrium along the transition, and how to compute it.
2. **Welfare evaluation**: who gains and who loses along the transition, and why.
3. Aggregate risk: Krusell and Smith's approximate equilibrium, and approximate aggregation.
4. Modern methods: transitions as impulse responses, sequence-space Jacobians, and when linear methods fail.

This lecture covers Parts 1 and 2; the next one covers Parts 3 and 4.

Part 1: transitions

A government that taxes and transfers

- Take the Aiyagari economy and add a government. Its policy is a sequence of labor-income tax rates $\{\tau_t\}_{t=0}^{\infty}$.
- Revenue is rebated lump sum, and the **budget balances every period**: each household receives $\tau_t w_t L$, where $L = \int \theta \, \mathrm{d}\lambda_t$.
- The household's budget constraint becomes

$$c_t + q_t \, a_{t+1} = (1 - \tau_t) \, w_t \, \theta_t + \underbrace{\tau_t \, w_t \, L}_{\text{rebate}} + a_t, \qquad a_{t+1} \geq \underline{a}.$$

- The policy provides **partial insurance**: after-tax labor income is $w_t\big[(1-\tau_t)\theta_t + \tau_t L\big]$, so its dispersion shrinks with τ_t.

The household's problem along a path of prices

- There is no aggregate risk, so the paths of prices and taxes are **deterministic**. Give the household knowledge of the whole sequence $\{q_t, w_t, \tau_t\}_{t\geq 0}$:

$$V_t(a,\theta) = \max_{c,\,a'} \; u(c) + \beta\,\mathbb{E}\big[V_{t+1}(a',\theta') \mid \theta\big]$$

$$\text{subject to} \quad c + q_t\,a' = (1-\tau_t)\,w_t\,\theta + \tau_t\,w_t\,L + a, \qquad a' \geq \underline{a}.$$

- The index t is an abuse of notation. The household knows the whole sequence; t only tells it which prices and taxes apply today and which tomorrow. Strictly, $V_t(a,\theta) = V\big(a,\theta;\{q_s,w_s,\tau_s\}_{s\geq t}\big)$.

- Compare with last lecture: instead of λ and its law of motion, the household knows the price sequence. Under perfect foresight this is all it needs. The cost: V_t and the policies depend on **calendar time**, because prices do.

Equilibrium along a transition

Recursive competitive equilibrium with a given tax path

Given an initial distribution λ_0, initial capital K_0 and taxes $\{\tau_t\}$: value functions $\{V_t\}$, policies $\{c_t^\star, a_t^\star\}$, prices $\{r_t, w_t\}$ and distributions $\{\lambda_t\}$ such that, for every t:

1. given prices and taxes, V_t solves the Bellman equation and $c_t^\star, a_t^\star$ are its policies, with $q_t = 1/(1 + r_{t+1})$;
2. aggregates add up individual decisions: $K_{t+1} = q_t \int a_t^\star \, \mathrm{d}\lambda_t$ and $L = \int \theta \, \mathrm{d}\lambda_t$;
3. prices are marginal products: $r_t = F_K(K_t, L) - \delta$ and $w_t = F_L(K_t, L)$;
4. the distribution moves with the policies: $\lambda_{t+1} = Q_t^\top \lambda_t$.

- K_0 is given too: λ_0 is over face values, and $K_0 = q_{-1} \int a \, \mathrm{d}\lambda_0$ depends on last period's price.
- The government's budget holds by construction.

An unexpected reform

- Given λ_0 and K_0, the definition allows any tax path. We study one: an **unexpected, permanent** reform,

$$\tau_t = \hat{\tau} \text{ for } t < 0, \qquad \tau_t = \tilde{\tau} \text{ for } t \geq 0.$$

- Before $t = 0$, the old tax has been in place for a long time, so the economy is at the stationary equilibrium with $\hat{\tau}$: $\lambda_0 = \hat{\lambda}$ and $K_0 = \hat{K}$.
- At $t = 0$ the reform is a surprise. From then on, households know the whole path of prices.
- Eventually, the economy reaches the stationary equilibrium with $\tilde{\tau}$.
- A one-time surprise followed by perfect foresight is called an **MIT shock**. The name started as a joke about taking the easy route; it is now standard practice.

Why not just compare steady states?

- A steady-state comparison asks: in which economy would you rather be born, not knowing your place? It evaluates the reform **behind the veil of ignorance**, in the long run.
- But suppose we care about the welfare effects of a one-time change in policy: who wins, who loses, and whether the reform is politically feasible.
- The households alive at $t = 0$ are not behind the veil of ignorance. They start at specific states (a, θ), distributed as $\hat{\lambda}$. As prices move, the reform redistributes among them: changes in r_t and w_t help some and hurt others, depending on their assets and earnings.
- The comparison also ignores the cost of getting there: reaching more capital requires consuming less in the meantime; reaching less capital lets households consume part of it.

The algorithm

Algorithm: transition after an MIT shock

0. Compute the stationary equilibria with $\hat{r}$ and $\tilde{r}$. Choose a long horizon $T^\star$.
1. Guess interest rates $\{r_t\}_{t=1}^{T^\star}$ with $r_{T^\star} = \tilde{r}$ ($r_0 = \hat{r}$, since $K_0 = \hat{K}$).
2. From the firm's conditions, get the demand for capital $K(r_t)$ and the wage $w(r_t)$; set $q_t = 1/(1 + r_{t+1})$.
3. **Backward**: starting from the new steady state's solution at $T^\star$, solve the household problem one period at a time, $t = T^\star - 1, \ldots, 0$, to get $a_t^\star$.
4. **Forward**: from $\lambda_0 = \hat{\lambda}$, build Q_t and set $\lambda_{t+1} = Q_t^\top \lambda_t$ for $t = 0, \ldots, T^\star - 1$.
5. Compute excess supply, $E_t = q_{t-1} \sum a_{t-1}^\star \lambda_{t-1} - K(r_t)$, for $t = 1, \ldots, T^\star$.
6. If $\max_t |E_t| < \varepsilon$, stop. Otherwise update $r_t \leftarrow r_t - \xi_t E_t$, with steps $\xi_t > 0$ that decrease with t, and go back to step 2.

Comments on the algorithm

1. **Backward, then forward**: households look ahead, since they need future prices; the distribution looks back, since it needs past choices. So we iterate on a whole path.

2. Within an iteration there is no fixed point to solve: the solution at t follows from the one at $t+1$ in a single step, by any of last lecture's methods. An iteration costs $T^\star$ such steps and $T^\star$ histogram updates.

3. We guess **prices**, as in last lecture's search over r: households take prices as given. The alternative is to guess **quantities**, the path $\{K_t\}$, get prices from the firm, and update by averaging the guess with the implied capital. Which is easier depends on the model.

Comments on the algorithm (continued)

4. **Scale the update**. Late in the path, excess supply carries the response of saving over many periods, so it reacts more to rates than early on. Use smaller steps there: ξ_t decreasing in t. Even so, convergence is **not guaranteed**.

5. Check the horizon: the path should approach the new steady state smoothly and be flat well before $T^\star$; a longer $T^\star$ should not change it.

An experiment: a tax reform in the Aiyagari economy

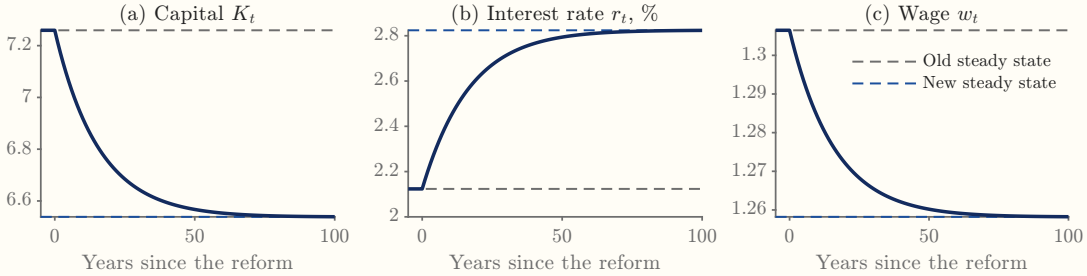
- Take the Aiyagari economy of last lecture, with the same calibration. At $t = 0$, the labor-income tax rises from $\hat{\tau} = 0$ to $\tilde{\tau} = 0.2$, unexpectedly and permanently. Revenue is rebated lump sum.

- Before looking at the results, think about a few predictions: what happens to capital, the interest rate and the wage, and to wealth and income inequality?

- Would a planner choose this reform? It insures, but lowers capital and the wage, on which the poorest depend. And would the households alive at $t = 0$ vote for it?

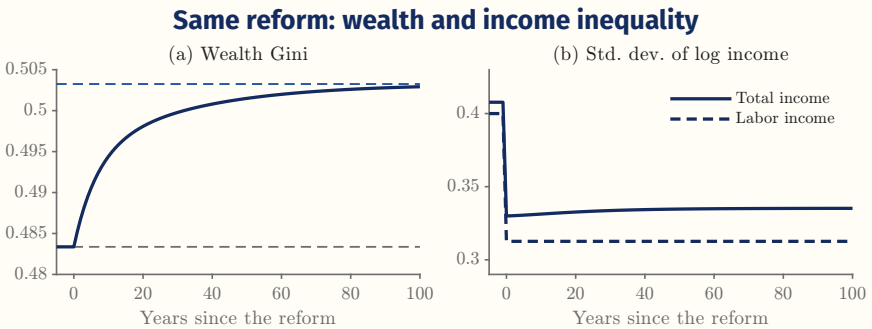
Prediction: the transition after a tax reform

Unexpected, permanent rise in τ from 0 to 0.2 at $t = 0$: aggregates



- The direction we could anticipate: it is Aiyagari's question in reverse. Less income risk means less precautionary saving, so capital falls and r rises.
- The **magnitudes** and the **timing** we could not, and that is what a quantitative model is for. In the long run capital is 10% lower, r rises from 2.1% to 2.8% and w falls by 4%. Capital is a stock, so prices cannot jump: $r_0 = \hat{r}$, and the economy takes 35 to 40 years to cover 90% of the way.

Prediction: inequality along the transition



Income: after-tax labor income plus capital income. Dashed: old and new steady states.

- Wealth inequality **rises**, slowly: the Gini goes from 0.48 to 0.50, the share with no wealth from 1.6% to 3.7%. Insurance crowds out self-insurance.
- Income inequality falls **on impact**, then rises a little. The tax compresses labor income once; then capital income, far more unequal, gains weight as r rises: its share goes from 10.6% to 12.8%.

A second application: a credit crunch

- A **liquidity trap**: the nominal rate is stuck at zero while the **natural rate** (the real rate that clears markets with flexible prices) is negative, so output falls instead of the rate.

- It needs a shock that pushes the natural rate below zero. RA models use a preference shock that makes households want to save more. Eggertsson and Krugman (2012) tighten the debt limit of impatient households who borrow from patient ones.

- Guerrieri and Lorenzoni (2017): with a whole distribution, a tighter borrowing limit works through two channels. Borrowers must repay, and everyone else saves more as a precaution.

A credit crunch: the model

- We sketch the demand side of their model: an endowment economy with bonds, as in Huggett, plus a government. Output is exogenous.
- Income: when employed, productivity θ follows a Markov chain; households also move in and out of **unemployment**, where they receive a benefit ν.
- The government supplies bonds, $B > 0$, and finances the interest and the benefit with a lump-sum tax, $T_t = (1 - q_t)B + u\,\nu$, where u is the unemployment rate:

$$c_t + q_t\, a_{t+1} = y_t - T_t + a_t, \qquad a_{t+1} \geq \underline{a}_t, \qquad \int a_t^\star \,\mathrm{d}\lambda_t = B.$$

- The limit is now indexed by t. The definition of equilibrium allows any path for it, as long as every household can meet its obligations, even when unemployed: the model has no default.
- Same algorithm, with excess supply $\int a_t^\star \,\mathrm{d}\lambda_t - B$.

An experiment: a credit crunch

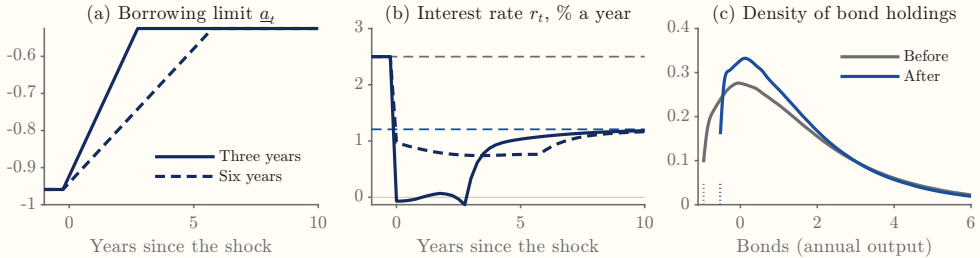
- We use their calibration. A period is a quarter and output is 1 a year. Risk aversion is 4; unemployed households receive 40% of mean labor income; B is 1.6 times annual output; β is set so that $r = 2.5\%$ a year.

- At $t = 0$, the limit tightens unexpectedly and permanently, from $\hat{\underline{a}} = -0.96$ to $\tilde{\underline{a}} = -0.53$ times annual output. It is **phased in** linearly over three years, and households know the path from $t = 0$. To see what drives the results, we also phase it in over six years.

- Why gradually? Bonds last one quarter, so a sudden tightening would require the most indebted households to repay at once, more than an unemployed household can. Over three years, an unemployed household at the old limit repays without its consumption collapsing.

Prediction: the transition after a credit crunch

Unexpected, permanent tightening of the limit, phased in over three or six years



- The interest rate **overshoots**, and the faster the tightening, the deeper it falls: to about zero while the limit tightens over three years, only to about 0.8% over six. Either way, it then rises slowly to 1.2%.
- Panel (c) compares steady states, so it holds for both paths: debt shrinks and mass piles up near zero, with the same mean, B, in a more concentrated distribution.

What we learn from the credit crunch

- Bonds are in fixed supply and output is given: aggregate saving and consumption cannot change.
- The tightening forces debtors to repay, and the rest want to save more as a precaution. At the old rate, everyone would cut consumption.
- So r must fall until those at the top are willing to save less and consume what debtors no longer buy. The distribution makes r overshoot.
- Our r is the natural rate; with the faster tightening, it turns negative. With nominal rigidities and a zero bound on the nominal rate, output falls instead: **real effects**, a liquidity trap.
- **Public liquidity can replace private liquidity**: issuing bonds to give households their lost borrowing capacity offsets the tightening exactly *(Guerrieri and Lorenzoni, Section VI)*.

Part 2: welfare evaluation along the transition

The welfare question

- At $t = 0$, ask a household with state (a, θ): what fraction of your old steady-state consumption would you give up to avoid the reform?
- The answer is the $D(a, \theta)$ that solves

$$\mathbb{E}\Big[\sum_{t=0}^{\infty} \beta^t \, u\big((1 - D(a,\theta))\, \hat{c}_t\big)\Big] = \mathbb{E}\Big[\sum_{t=0}^{\infty} \beta^t \, u(c_t^\star)\Big],$$

  where $\hat{c}_t$ is consumption under the old policy and $c_t^\star$ consumption along the transition, both from $(a_0, \theta_0) = (a, \theta)$.
- With CRRA utility, $u(c) = c^{1-\sigma}/(1-\sigma)$,

$$D(a,\theta) = 1 - \left[\frac{V_0(a,\theta)}{\hat{V}(a,\theta)}\right]^{\frac{1}{1-\sigma}}.$$

- $D > 0$: the household loses and would pay to stop the reform. $D < 0$: it gains, and would ask to be compensated if the reform were blocked.

Welfare needs the value function

- Equilibrium needs only the policies, and some of our methods, such as EGM, deliver them without ever computing V.
- But welfare evaluation **needs** V: $D(a, \theta)$ compares V_0 with $\hat{V}$. We need three value functions:
  - $\hat{V}$ and $\tilde{V}$, of the two steady states: iterate to convergence

$$V(a_n, \theta_i) = u\big(c^\star(a_n, \theta_i)\big) + \beta \sum_j \pi_{ij}\, V\big(a^\star(a_n, \theta_i), \theta_j\big);$$

  - V_0: go backward from $V_{T^\star} = \tilde{V}$ with the policies of the transition,

$$V_t(a_n, \theta_i) = u\big(c_t^\star(a_n, \theta_i)\big) + \beta \sum_j \pi_{ij}\, V_{t+1}\big(a_t^\star(a_n, \theta_i), \theta_j\big), \qquad t = T^\star - 1, \ldots, 0.$$

- In both cases the policies are known, so no maximization is needed; the value next period is interpolated between grid points.

Conditional and unconditional welfare change

- $D(a, \theta)$ is the **conditional** welfare change: it depends on the household's state. It gives a whole distribution of gains and losses.

- The **unconditional** welfare change asks the same question before the household knows its state. Since both alternatives start from $\hat{\lambda}$,

$$D = 1 - \left[\frac{\int V_0(a, \theta)\, \hat{\lambda}(\mathrm{d}a, \mathrm{d}\theta)}{\int \hat{V}(a, \theta)\, \hat{\lambda}(\mathrm{d}a, \mathrm{d}\theta)} \right]^{\frac{1}{1-\sigma}}.$$

- It evaluates the reform **behind the veil of ignorance**, or from the point of view of a planner who weighs every household equally.

Three reasons welfare changes

- The reform changes D through three channels:
  1. the **uncertainty** households face, for instance the income risk they bear;
  2. the **distribution** of resources: V_0 and $\hat{V}$ differ because the distribution moves from $\hat{\lambda}$ toward $\tilde{\lambda}$;
  3. the **level** of resources: through saving, capital and output respond to the reform.

- Flodén (2001) measures each one. With CRRA utility the three measures multiply exactly to the total.

- He compares **stationary equilibria** and sets the transition aside. We apply his measures along the transition.

Decomposing the welfare change: the ingredients

- **Certainty-equivalent consumption**: the constant consumption that gives the household its value, $c^{ce}(a,\theta) = \left[(1-\sigma)(1-\beta)\,V(a,\theta)\right]^{\frac{1}{1-\sigma}}$.
- Three levels of consumption, all averaged with the initial distribution $\hat{\lambda}$:
  - mean consumption, $\bar{C} = \int c^{\star}\,\mathrm{d}\hat{\lambda}$;
  - mean certainty-equivalent consumption, $C^{ce} = \int c^{ce}\,\mathrm{d}\hat{\lambda}$: risk is stripped from each household's stream;
  - the level that, given to everyone, yields average welfare: $\sum_t \beta^t u(C^{eq}) = \int V\,\mathrm{d}\hat{\lambda}$.
- **Cost of uncertainty**, $1-p = C^{ce}/\bar{C}$: the share of mean consumption lost to risk.
- **Cost of inequality**, $1-\omega = C^{eq}/C^{ce}$. Thought experiment: give every household the average certainty equivalent, C^{ce}. Since u is concave, average welfare rises. ω is the share of everyone's consumption we could then take away and still match the welfare of the unequal economy.

Decomposing the welfare change

- By construction, $(1-\omega)(1-p)\,\bar{C} = C^{eq} = \left[(1-\sigma)(1-\beta)\int V\,\mathrm{d}\hat{\lambda}\right]^{\frac{1}{1-\sigma}}$.

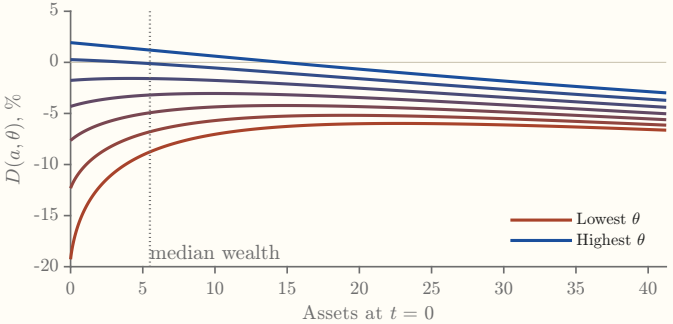
- Compute it before ($\hat{V}$) and along the transition (V_0), and divide:

$$1 - D \;=\; \underbrace{\frac{\bar{C}_0}{\hat{\bar{C}}}}_{1-D_A\text{: level}} \qquad \underbrace{\frac{1-p_0}{1-\hat{p}}}_{1-D_U\text{: uncertainty}} \qquad \underbrace{\frac{1-\omega_0}{1-\hat{\omega}}}_{1-D_I\text{: inequality}} \;.$$

- Along the transition, mean consumption $\bar{C}_t$ moves. We use for $\bar{C}_0$ the constant level with the same discounted utility, $\sum_t \beta^t u(\bar{C}_0) = \sum_t \beta^t u(\bar{C}_t)$. The product is exact for any choice of $\bar{C}_0$; the choice only splits the total between level and uncertainty.

Prediction: who gains from the reform

Welfare cost of the reform at $t = 0$, $D(a, \theta)$, % of consumption



- Almost everyone gains, the poorest most. Only 1.7% of households lose: those with high θ and little wealth, who pay the tax for years and earn a lower wage.
- At the top, gains **rise with wealth**: a higher r is a transfer to asset holders. At the old prices, all households with the highest θ would lose about 1.2%; the rise in r turns the rich among them into winners.

Prediction: steady states against the transition

Welfare cost of the reform, % of consumption

	Total D	Level D_A	Uncertainty D_U	Inequality D_I
Steady states	-2.84	1.22	-3.00	-1.08
Transition	-4.84	-0.50	-3.04	-1.23

Steady states: each with its own distribution. Transition: households alive at $t = 0$, distributed as $\hat{\lambda}$.

- Comparing steady states **understates the gain by 2 points**.
- The level effect changes sign. Long-run consumption is 1.2% lower, but along the way households eat the capital they no longer want: consumption rises 3% at $t = 0$.
- The uncertainty and inequality parts barely change: the transition mostly adds level.
- Inequality falls although the wealth Gini rises: D_I is about certainty-equivalent consumption, which the tax compresses, not wealth.

Summary

- When the distribution moves, households must forecast it. MIT shocks avoid the problem by giving households **perfect foresight** after a surprise.
- A transition is computed on a path: households **backward**, the distribution **forward**, and an update of the price path until markets clear at every date.
- Transitions make welfare analysis possible: **conditional** gains and losses, and a decomposition into level, uncertainty and inequality.
- Comparing steady states can **mislead**: in the tax reform it understates the gain by 2 points, because it ignores the capital consumed along the way.
- A tighter borrowing limit pushes the interest rate down, and the distribution makes it overshoot.
- Next lecture: aggregate risk, and how transitions become the building block of modern solution methods.

References

Eggertsson, G. B. and P. Krugman (2012). “Debt, Deleveraging, and the Liquidity Trap: A Fisher–Minsky–Koo Approach.” *Quarterly Journal of Economics* 127(3): 1469–1513.

Flodén, M. (2001). “The Effectiveness of Government Debt and Transfers as Insurance.” *Journal of Monetary Economics* 48(1): 81–108.

Guerrieri, V. and G. Lorenzoni (2017). “Credit Crises, Precautionary Savings, and the Liquidity Trap.” *Quarterly Journal of Economics* 132(3): 1427–1467.