

Lecture 10: Heterogeneous Agents

Jacob Adenbaum and Albert Rodriguez-Sala

University of Edinburgh

Spring 2024

Where we've been

- ▶ So far, we've learned how to write economic models **recursively**.
- ▶ Our prototypical example was the Neoclassical Growth Model:

$$\begin{aligned} v(k) &= \max_{c, k'} u(c) + \beta v(k') \\ \text{s.t.} \quad & c + k' \leq F(k) + (1 - \delta)k \end{aligned}$$

- ▶ Once they're written recursively, we've learned how to **solve** them (find a function that satisfies the recursive relationship)
- ▶ Once they're solved, we learned how to **simulate** them, and use the simulated data to **estimate** parameters
- ▶ With representative agents, an equilibrium in these models is not very complicated
 - ▶ If firms rent capital from household, we get $r = F'(k)$, etc...

Where we've been

- ▶ So far, we've learned how to write economic models **recursively**.

- ▶ Our prototypical example was the Neoclassical Growth Model:

$$\begin{aligned} v(k) &= \max_{c, k'} u(c) + \beta v(k') \\ \text{s.t.} \quad & c + k' \leq F(k) + (1 - \delta)k \end{aligned}$$

- ▶ Once they're written recursively, we've learned how to **solve** them (find a function that satisfies the recursive relationship)
- ▶ Once they're solved, we learned how to **simulate** them, and use the simulated data to **estimate** parameters
- ▶ With representative agents, an equilibrium in these models is not very complicated
 - ▶ If firms rent capital from household, we get $r = F'(k)$, etc...

Where we've been

- ▶ So far, we've learned how to write economic models **recursively**.

- ▶ Our prototypical example was the Neoclassical Growth Model:

$$\begin{aligned} v(k) &= \max_{c, k'} u(c) + \beta v(k') \\ \text{s.t.} \quad & c + k' \leq F(k) + (1 - \delta)k \end{aligned}$$

- ▶ Once they're written recursively, we've learned how to **solve** them (find a function that satisfies the recursive relationship)
- ▶ Once they're solved, we learned how to **simulate** them, and use the simulated data to **estimate** parameters
- ▶ With representative agents, an equilibrium in these models is not very complicated
 - ▶ If firms rent capital from household, we get $r = F'(k)$, etc...

Where we're going

- ▶ Computing with General Equilibrium
 - ▶ Many interesting models do not feature a representative household
 - ▶ When there are many heterogeneous agents in our models, there are several special concerns
 - ▶ Mostly about how we compute the **market clearing prices**
 - ▶ How do we approach these models computationally?
- ▶ Policy Analysis
 - ▶ The models we've worked on so far have all been efficient (Limited role for policy)
 - ▶ It's hard to even think about redistribution in a model with just a representative household
 - ▶ In many models, the government can step in to correct market failures, but we need to know: what is the **optimal policy**?

Where we're going

- ▶ Computing with General Equilibrium

- ▶ Many interesting models do not feature a representative household
- ▶ When there are many heterogeneous agents in our models, there are several special concerns
- ▶ Mostly about how we compute the **market clearing prices**
- ▶ How do we approach these models computationally?

- ▶ Policy Analysis

- ▶ The models we've worked on so far have all been efficient (Limited role for policy)
- ▶ It's hard to even think about redistribution in a model with just a representative household
- ▶ In many models, the government can step in to correct market failures, but we need to know: what is the **optimal policy**?

Where we're going

- ▶ Computing with General Equilibrium
 - ▶ Many interesting models do not feature a representative household
 - ▶ When there are many heterogeneous agents in our models, there are several special concerns
 - ▶ Mostly about how we compute the **market clearing prices**
 - ▶ How do we approach these models computationally?
- ▶ Policy Analysis
 - ▶ The models we've worked on so far have all been efficient (Limited role for policy)
 - ▶ It's hard to even think about redistribution in a model with just a representative household
 - ▶ In many models, the government can step in to correct market failures, but we need to know: what is the **optimal policy**?

Where we're going

- ▶ Computing with General Equilibrium
 - ▶ Many interesting models do not feature a representative household
 - ▶ When there are many heterogeneous agents in our models, there are several special concerns
 - ▶ Mostly about how we compute the **market clearing prices**
 - ▶ How do we approach these models computationally?
- ▶ Policy Analysis
 - ▶ The models we've worked on so far have all been efficient (Limited role for policy)
 - ▶ It's hard to even think about redistribution in a model with just a representative household
 - ▶ In many models, the government can step in to correct market failures, but we need to know: what is the **optimal policy**?

Section 1

Heterogeneous Agent Models: Aiyagari (1994, QJE)

Aiyagari (1994, QJE): Prototypical Heterogeneous Agent Model

$$\begin{aligned} v(a, y) = \max_{c, a' \geq 0} \quad & u(c) + \beta \mathbb{E}[v(a', y')|y] \\ \text{s.t.} \quad & c + a' \leq (1 + r)a + y \\ & \log(y') = \rho \log(y) + \epsilon \\ & \epsilon \sim N(0, \sigma) \end{aligned} \tag{1}$$

- ▶ Consider the problem of a large group of households who must save for the future
- ▶ They are heterogeneous in their current income y , and in their level of assets a .
 - ▶ Log income follows an AR(1) process
 - ▶ Labor supplied inelastically (no choice of how much to work)
- ▶ Derive flow utility $u(c)$ from consumption, and discount the future at rate β
- ▶ Can save for the future at a rate $1 + r$, but cannot borrow.
 - ▶ **Markets are incomplete** (There are certain risks that they cannot insure against)
- ▶ So far, this should look very familiar from your problem set...

Aiyagari (1994, QJE): Prototypical Heterogeneous Agent Model

$$\begin{aligned} v(a, y) = \max_{c, a' \geq 0} \quad & u(c) + \beta \mathbb{E}[v(a', y')|y] \\ \text{s.t.} \quad & c + a' \leq (1 + r)a + y \\ & \log(y') = \rho \log(y) + \epsilon \\ & \epsilon \sim N(0, \sigma) \end{aligned} \tag{1}$$

- ▶ Consider the problem of a large group of households who must save for the future
- ▶ They are heterogeneous in their current income y , and in their level of assets a .
 - ▶ Log income follows an AR(1) process
 - ▶ Labor supplied inelastically (no choice of how much to work)
- ▶ Derive flow utility $u(c)$ from consumption, and discount the future at rate β
- ▶ Can save for the future at a rate $1 + r$, but cannot borrow.
 - ▶ **Markets are incomplete** (There are certain risks that they cannot insure against)
- ▶ So far, this should look very familiar from your problem set...

Aiyagari (1994, QJE): Prototypical Heterogeneous Agent Model

$$\begin{aligned} v(a, y) = \max_{c, a' \geq 0} \quad & u(c) + \beta \mathbb{E}[v(a', y')|y] \\ \text{s.t.} \quad & c + a' \leq (1 + r)a + y \\ & \log(y') = \rho \log(y) + \epsilon \\ & \epsilon \sim N(0, \sigma) \end{aligned} \tag{1}$$

- ▶ Consider the problem of a large group of households who must save for the future
- ▶ They are heterogeneous in their current income y , and in their level of assets a .
 - ▶ Log income follows an AR(1) process
 - ▶ Labor supplied inelastically (no choice of how much to work)
- ▶ Derive flow utility $u(c)$ from consumption, and discount the future at rate β
- ▶ Can save for the future at a rate $1 + r$, but cannot borrow.
 - ▶ **Markets are incomplete** (There are certain risks that they cannot insure against)
- ▶ So far, this should look very familiar from your problem set...

Extending Bewley to Aiyagari

- ▶ If we take r as given, and just consider the households' consumption savings problem, then we know how to solve
 - ▶ We saw that it's not much more complicated than the neoclassical growth model with stochastic productivity
 - ▶ But r is a price: we want it to be set, in equilibrium, to clear the market for assets
- ▶ Supply Side: Suppose we have a representative firm, with production function $F(k)$, who rents capital from the households at a price r .

$$\max_k F(k) - rk \implies F'(k) = r \implies k = K(r) \quad (2)$$

For some function $K(r)$. If $F(k) = k^\alpha$, then $K(r) = \left(\frac{\alpha}{r}\right)^{\frac{1}{1-\alpha}}$

- ▶ Distribution of agents: let $\Lambda(a, y)$ be the cumulative distribution function of assets and income in the economy (with pdf λ)

Extending Bewley to Aiyagari

- ▶ If we take r as given, and just consider the households' consumption savings problem, then we know how to solve
 - ▶ We saw that it's not much more complicated than the neoclassical growth model with stochastic productivity
 - ▶ But r is a price: we want it to be set, in equilibrium, to clear the market for assets
- ▶ Supply Side: Suppose we have a representative firm, with production function $F(k)$, who rents capital from the households at a price r .

$$\max_k F(k) - rk \implies F'(k) = r \implies k = K(r) \quad (2)$$

For some function $K(r)$. If $F(k) = k^\alpha$, then $K(r) = \left(\frac{\alpha}{r}\right)^{\frac{1}{1-\alpha}}$

- ▶ Distribution of agents: let $\Lambda(a, y)$ be the cumulative distribution function of assets and income in the economy (with pdf λ)

Equilibrium

► A **recursive stationary equilibrium** in this model is a set of

1. Consumption and savings policy functions $g_c(a, y)$ and $g_a(a, y)$,
2. An interest rate r , and
3. A distribution $\Lambda(a, y)$ over assets and income levels

► Such that:

1. **Optimality:** g_c and g_a solve the household's consumption/savings problem, given r
2. **Market Clearing:** The interest rate r clears the market for capital

$$K(r) = \int a d\Lambda(a, y) = \int \int a \lambda(a, y) da dy \quad (3)$$

3. **Stationarity:** Given the policy functions g_c and g_a , and the interest rate r , the distribution Λ is unchanging over time

- We know what optimality means – need to solve the household's dynamic program as we have been doing
- Need to spend a little bit of time thinking through market clearing and stationarity

Equilibrium

- ▶ A **recursive stationary equilibrium** in this model is a set of
 1. Consumption and savings policy functions $g_c(a, y)$ and $g_a(a, y)$,
 2. An interest rate r , and
 3. A distribution $\Lambda(a, y)$ over assets and income levels
- ▶ Such that:
 1. **Optimality**: g_c and g_a solve the household's consumption/savings problem, given r
 2. **Market Clearing**: The interest rate r clears the market for capital

$$K(r) = \int a d\Lambda(a, y) = \int \int a \lambda(a, y) da dy \quad (3)$$

3. **Stationarity**: Given the policy functions g_c and g_a , and the interest rate r , the distribution Λ is unchanging over time
- ▶ We know what optimality means – need to solve the household's dynamic program as we have been doing
 - ▶ Need to spend a little bit of time thinking through market clearing and stationarity

Equilibrium

- ▶ A **recursive stationary equilibrium** in this model is a set of
 1. Consumption and savings policy functions $g_c(a, y)$ and $g_a(a, y)$,
 2. An interest rate r , and
 3. A distribution $\Lambda(a, y)$ over assets and income levels
- ▶ Such that:
 1. **Optimality:** g_c and g_a solve the household's consumption/savings problem, given r
 2. **Market Clearing:** The interest rate r clears the market for capital

$$K(r) = \int a d\Lambda(a, y) = \int \int a \lambda(a, y) da dy \quad (3)$$

3. **Stationarity:** Given the policy functions g_c and g_a , and the interest rate r , the distribution Λ is unchanging over time
- ▶ We know what optimality means – need to solve the household's dynamic program as we have been doing
 - ▶ Need to spend a little bit of time thinking through market clearing and stationarity

Equilibrium

- ▶ A **recursive stationary equilibrium** in this model is a set of
 1. Consumption and savings policy functions $g_c(a, y)$ and $g_a(a, y)$,
 2. An interest rate r , and
 3. A distribution $\Lambda(a, y)$ over assets and income levels
- ▶ Such that:
 1. **Optimality:** g_c and g_a solve the household's consumption/savings problem, given r
 2. **Market Clearing:** The interest rate r clears the market for capital

$$K(r) = \int a d\Lambda(a, y) = \int \int a \lambda(a, y) da dy \quad (3)$$

3. **Stationarity:** Given the policy functions g_c and g_a , and the interest rate r , the distribution Λ is unchanging over time
- ▶ We know what optimality means – need to solve the household's dynamic program as we have been doing
 - ▶ Need to spend a little bit of time thinking through market clearing and stationarity

Equilibrium

- ▶ A **recursive stationary equilibrium** in this model is a set of
 1. Consumption and savings policy functions $g_c(a, y)$ and $g_a(a, y)$,
 2. An interest rate r , and
 3. A distribution $\Lambda(a, y)$ over assets and income levels
- ▶ Such that:
 1. **Optimality:** g_c and g_a solve the household's consumption/savings problem, given r
 2. **Market Clearing:** The interest rate r clears the market for capital

$$K(r) = \int a d\Lambda(a, y) = \int \int a \lambda(a, y) da dy \quad (3)$$

3. **Stationarity:** Given the policy functions g_c and g_a , and the interest rate r , the distribution Λ is unchanging over time
- ▶ We know what optimality means – need to solve the household's dynamic program as we have been doing
 - ▶ Need to spend a little bit of time thinking through market clearing and stationarity

Stationarity

- ▶ A stationary distribution is one that is not changing over time
- ▶ If we step the distribution forward one time period, using our policy rules, we should get the same distribution back out again
- ▶ Let $\pi(y'|y)$ denote the conditional pdf of income tomorrow given that income today is y .
- ▶ Then we can write the **law of motion** for Λ as

$$\Lambda(a, y) = \int_{-\infty}^{\infty} \int_{-\infty}^y \int_0^{\infty} \mathbb{1} \{g_a(a_0, y_0) \leq a'\} \pi(y'|y_0) \lambda(a_0, y_0) da_0 dy' dy_0$$

- ▶ This is just fancy math for: if I step my simulated distribution of agents forward one period, the overall distribution should not change
- ▶ Each agent is moving around through the distribution, but on average it stays the same
- ▶ In this class, we will never compute those integrals directly – we will always be approximating the distribution using a simulated set with a discrete number of agents

Stationarity

- ▶ A stationary distribution is one that is not changing over time
- ▶ If we step the distribution forward one time period, using our policy rules, we should get the same distribution back out again
- ▶ Let $\pi(y'|y)$ denote the conditional pdf of income tomorrow given that income today is y .
- ▶ Then we can write the **law of motion** for Λ as

$$\Lambda(a, y) = \int_{-\infty}^{\infty} \int_{-\infty}^y \int_0^{\infty} \mathbb{1} \{g_a(a_0, y_0) \leq a'\} \pi(y'|y_0) \lambda(a_0, y_0) da_0 dy' dy_0$$

- ▶ This is just fancy math for: if I step my simulated distribution of agents forward one period, the overall distribution should not change
- ▶ Each agent is moving around through the distribution, but on average it stays the same
- ▶ In this class, we will never compute those integrals directly – we will always be approximating the distribution using a simulated set with a discrete number of agents

Stationarity

- ▶ A stationary distribution is one that is not changing over time
- ▶ If we step the distribution forward one time period, using our policy rules, we should get the same distribution back out again
- ▶ Let $\pi(y'|y)$ denote the conditional pdf of income tomorrow given that income today is y .
- ▶ Then we can write the **law of motion** for Λ as

$$\Lambda(a, y) = \int_{-\infty}^{\infty} \int_{-\infty}^y \int_0^{\infty} \mathbb{1} \{g_a(a_0, y_0) \leq a'\} \pi(y'|y_0) \lambda(a_0, y_0) da_0 dy' dy_0$$

- ▶ This is just fancy math for: if I step my simulated distribution of agents forward one period, the overall distribution should not change
- ▶ Each agent is moving around through the distribution, but on average it stays the same
- ▶ In this class, we will never compute those integrals directly – we will always be approximating the distribution using a simulated set with a discrete number of agents

Stationarity

- ▶ A stationary distribution is one that is not changing over time
- ▶ If we step the distribution forward one time period, using our policy rules, we should get the same distribution back out again
- ▶ Let $\pi(y'|y)$ denote the conditional pdf of income tomorrow given that income today is y .
- ▶ Then we can write the **law of motion** for Λ as

$$\Lambda(a, y) = \int_{-\infty}^{\infty} \int_{-\infty}^y \int_0^{\infty} \mathbb{1} \{g_a(a_0, y_0) \leq a'\} \pi(y'|y_0) \lambda(a_0, y_0) da_0 dy' dy_0$$

- ▶ This is just fancy math for: if I step my simulated distribution of agents forward one period, the overall distribution should not change
- ▶ Each agent is moving around through the distribution, but on average it stays the same
- ▶ In this class, we will never compute those integrals directly – we will always be approximating the distribution using a simulated set with a discrete number of agents

Stationarity

- ▶ A stationary distribution is one that is not changing over time
- ▶ If we step the distribution forward one time period, using our policy rules, we should get the same distribution back out again
- ▶ Let $\pi(y'|y)$ denote the conditional pdf of income tomorrow given that income today is y .
- ▶ Then we can write the **law of motion** for Λ as

$$\Lambda(a, y) = \int_{-\infty}^{\infty} \int_{-\infty}^y \int_0^{\infty} \mathbb{1} \{g_a(a_0, y_0) \leq a'\} \pi(y'|y_0) \lambda(a_0, y_0) da_0 dy' dy_0$$

- ▶ This is just fancy math for: if I step my simulated distribution of agents forward one period, the overall distribution should not change
- ▶ Each agent is moving around through the distribution, but on average it stays the same
- ▶ In this class, we will never compute those integrals directly – we will always be approximating the distribution using a simulated set with a discrete number of agents

Stationarity


- ▶ In general, your distributions will usually converge to a single, stationary distribution
As long as it's possible to move from every point in the state space to every other point in the state space (full mixing)
- ▶ We say that the distribution has converged if the histogram of assets and income has stopped changing
- ▶ The thing we actually want is to calculate the total assets in the economy:

$$A(\Lambda) = \int \int a \lambda(a, y) da dy$$

- ▶ Take the average of the assets of our agents in our simulated distribution:

$$A(\Lambda) \approx \frac{1}{N} \sum_{i=1}^N a_i$$

Stationary Distribution




Market Clearing

- ▶ Remember that our stationary distribution is calculated using policy functions g_c and g_a that take r as given.
- ▶ That means we can really write $\Lambda(r)$: the stationary distribution of assets depends on the interest rate
- ▶ Our simulation results also depend on r : average assets are $A(\Lambda(r))$
- ▶ Supply and Demand:
 - ▶ $A(\Lambda(r))$ is our upward sloping supply curve of assets
 - ▶ $K(r)$ is our downward sloping demand curve for capital
 - ▶ The market clearing price is the r that sets

$$K(r) = A(\Lambda(r))$$

Asset Market Clearing




Market Clearing: Root finding approach

- ▶ Recast problem as root finding on excess demand:

$$ED(r) = K(r) - A(\Lambda(r)) \quad (4)$$

- ▶ With a sensible root finding procedure, you will typically converge within 10 iterations for a 1D problem
- ▶ If you have multiple markets to clear, then it's a multivariate root finding problem – harder to do
- ▶ Be careful of tolerances in your root finding procedure Simulations are noisy, and so you may not be able to solve your root finding problem accurately beyond a tolerance of 10^{-3} without a prohibitively large computational cost

Asset Market Clearing



There are more clever approaches to simulating the distribution of assets, but they tend to be less intuitive

Aiyagari: Wrapping Up

- ▶ In general, our computational tools allow us to analyze these types of heterogeneous agent problems
- ▶ When we do, we will have to think more carefully about how to deal with market clearing and other equilibrium conditions
- ▶ Very few limits (other than computational cost) on which dynamic models we can solve
- ▶ Especially when you move into the world of models with many agents, and nontrivial dispersion in wealth/human capital/income/etc..., these models are *not* amenable to being solved on pen and paper
- ▶ For many problems, VFI is the slowest, but most robust solution
 - ▶ There are other approaches, but they all tend to be more situational (although they often obtain large speed gains)
 - ▶ There are approaches (like policy function iteration, and others) that can speed up VFI
- ▶ Oftentimes, without a smarter approach, the majority of your time will be spent in the simulation code, rather than solving the model

Aiyagari: Wrapping Up

- ▶ In general, our computational tools allow us to analyze these types of heterogeneous agent problems
- ▶ When we do, we will have to think more carefully about how to deal with market clearing and other equilibrium conditions
- ▶ Very few limits (other than computational cost) on which dynamic models we can solve
- ▶ Especially when you move into the world of models with many agents, and nontrivial dispersion in wealth/human capital/income/etc..., these models are *not* amenable to being solved on pen and paper
- ▶ For many problems, VFI is the slowest, but most robust solution
 - ▶ There are other approaches, but they all tend to be more situational (although they often obtain large speed gains)
 - ▶ There are approaches (like policy function iteration, and others) that can speed up VFI
- ▶ Oftentimes, without a smarter approach, the majority of your time will be spent in the simulation code, rather than solving the model

Aiyagari: Wrapping Up

- ▶ In general, our computational tools allow us to analyze these types of heterogeneous agent problems
- ▶ When we do, we will have to think more carefully about how to deal with market clearing and other equilibrium conditions
- ▶ Very few limits (other than computational cost) on which dynamic models we can solve
- ▶ Especially when you move into the world of models with many agents, and nontrivial dispersion in wealth/human capital/income/etc..., these models are *not* amenable to being solved on pen and paper
- ▶ For many problems, VFI is the slowest, but most robust solution
 - ▶ There are other approaches, but they all tend to be more situational (although they often obtain large speed gains)
 - ▶ There are approaches (like policy function iteration, and others) that can speed up VFI
- ▶ Oftentimes, without a smarter approach, the majority of your time will be spent in the simulation code, rather than solving the model

Section 2

Policy Experiments

Why do we need a model for policy evaluation?

Predictive accuracy

There are many cases where the reduced-form elasticities you get from running a regression (even a well-identified regression) are *not* good predictors of how people will behave if you make changes to policy

- ▶ People who are forward-looking are much more responsive to permanent changes than temporary changes

You have to be careful about which elasticities you're actually measuring

- ▶ People can respond to changes in policy in unexpected ways

E.g. Changes in inflation expectations in the 1970s

- ▶ Predictions that are not grounded in a model of people's underlying choices are vulnerable to the Lucas Critique: behavior rules estimated in the data are not invariant to policy
- ▶ Making sense of the data available: Indirect Inference
- ▶ Often, economic models have good predictions, even out of sample.

Why do we need a model for policy evaluation?

Predictive accuracy

There are many cases where the reduced-form elasticities you get from running a regression (even a well-identified regression) are *not* good predictors of how people will behave if you make changes to policy

- ▶ People who are forward-looking are much more responsive to permanent changes than temporary changes

You have to be careful about which elasticities you're actually measuring

- ▶ People can respond to changes in policy in unexpected ways

E.g. Changes in inflation expectations in the 1970s

- ▶ Predictions that are not grounded in a model of people's underlying choices are vulnerable to the Lucas Critique: behavior rules estimated in the data are not invariant to policy
- ▶ Making sense of the data available: Indirect Inference
- ▶ Often, economic models have good predictions, even out of sample.

Why do we need a model for policy evaluation?

Predictive accuracy

There are many cases where the reduced-form elasticities you get from running a regression (even a well-identified regression) are *not* good predictors of how people will behave if you make changes to policy

- ▶ People who are forward-looking are much more responsive to permanent changes than temporary changes

You have to be careful about which elasticities you're actually measuring

- ▶ People can respond to changes in policy in unexpected ways

E.g. Changes in inflation expectations in the 1970s

- ▶ Predictions that are not grounded in a model of people's underlying choices are vulnerable to the Lucas Critique: behavior rules estimated in the data are not invariant to policy

- ▶ Making sense of the data available: Indirect Inference

- ▶ Often, economic models have good predictions, even out of sample.

Why do we need a model for policy evaluation?

Predictive accuracy

There are many cases where the reduced-form elasticities you get from running a regression (even a well-identified regression) are *not* good predictors of how people will behave if you make changes to policy

- ▶ People who are forward-looking are much more responsive to permanent changes than temporary changes

You have to be careful about which elasticities you're actually measuring

- ▶ People can respond to changes in policy in unexpected ways

E.g. Changes in inflation expectations in the 1970s

- ▶ Predictions that are not grounded in a model of people's underlying choices are vulnerable to the Lucas Critique: behavior rules estimated in the data are not invariant to policy
- ▶ Making sense of the data available: Indirect Inference

- ▶ Often, economic models have good predictions, even out of sample.

Why do we need a model for policy evaluation?

Predictive accuracy

There are many cases where the reduced-form elasticities you get from running a regression (even a well-identified regression) are *not* good predictors of how people will behave if you make changes to policy

- ▶ People who are forward-looking are much more responsive to permanent changes than temporary changes

You have to be careful about which elasticities you're actually measuring

- ▶ People can respond to changes in policy in unexpected ways

E.g. Changes in inflation expectations in the 1970s

- ▶ Predictions that are not grounded in a model of people's underlying choices are vulnerable to the Lucas Critique: behavior rules estimated in the data are not invariant to policy
- ▶ Making sense of the data available: Indirect Inference
- ▶ Often, economic models have good predictions, even out of sample.

Why do we need a model for policy evaluation?

Counterfactuals and Welfare Analysis

- ▶ Counterfactuals are at the heart of the questions we want to answer:
 - ▶ How will people's behavior in response to a policy that has never been implemented?
 - ▶ How would they have behaved if we hadn't implemented some policy?
- ▶ In nontrivial models, we need a model in order to evaluate the welfare impacts of a change in policy
 - ▶ Will people be better off on average after a tax reform?
 - ▶ By how much?
 - ▶ Will this reform increase or decrease inequality?
 - ▶ How are the gains distributed?
- ▶ Without a model, you cannot hope to answer these kinds of questions

Tax Reform in Aiyagari

- ▶ In Aiyagari models, generally people tend to over-save relative to what the social planner would choose
 - ▶ Fear of a sequence of very many negative shocks
 - ▶ If you hit your borrowing constraint, you may wind up with very low consumption
 - ▶ Strong precautionary motive for savings, at the individual level, to self-insure against income risk
- ▶ Suppose that the government imposes a tax on capital income τ , and redistributes the money with a lump sum tax T (let's say, $\tau = 30\%$)
- ▶ How can we model the effects of this tax reform?

Tax Reform in Aiyagari

- ▶ In Aiyagari models, generally people tend to over-save relative to what the social planner would choose
 - ▶ Fear of a sequence of very many negative shocks
 - ▶ If you hit your borrowing constraint, you may wind up with very low consumption
 - ▶ Strong precautionary motive for savings, at the individual level, to self-insure against income risk
- ▶ Suppose that the government imposes a tax on capital income τ , and redistributes the money with a lump sum tax T (let's say, $\tau = 30\%$)
- ▶ How can we model the effects of this tax reform?

Tax Reform in Aiyagari: Updated Model

$$\begin{aligned} v(a, y) = \max_{c, a' \geq 0} \quad & u(c) + \beta \mathbb{E}[v(a', y')|y] \\ \text{s.t.} \quad & c + a' \leq [1 + r(1 - \tau)]a + y + T \\ & \log(y') = \rho \log(y) + \epsilon \\ & \epsilon \sim N(0, \sigma) \end{aligned} \tag{5}$$

- ▶ Only change to the Bellman equations are in the budget constraint: consumers take τ and T as given
- ▶ New considerations:
 - ▶ Taxes distort savings behavior $\implies$ different r in equilibrium
 - ▶ Government needs to balance its budget $\implies$ find T such that

$$\int \int \tau r a \lambda(a, y) da dy = T$$

- ▶ Both of these will change consumer behavior – we have to solve for all of them jointly

Tax Reform in Aiyagari: Solution Algorithm

Fix $\tau = 30\%$. Treat these market clearing conditions as nested problems:

- ▶ Define $D(r, T)$ to be the government's budget deficit
- ▶ Define $ED(r, T)$ to be the excess demand for capital
- ▶ Algorithm:

1. For any given r , solve for the T that balances the government's budget (solving and simulating the model). That is, solve the root finding problem

$$D(r, T) = 0$$

as a function of T , holding r fixed. Call the results $T^*(r)$

2. Solve the root finding problem for

$$ED(r, T^*(r)) = 0$$

Call the result r^*

- ▶ Our final (r, T) are $(r^*, T^*(r^*))$.

We'll go over code for how to do this in tutorial

Tax Reform in Aiyagari: Solution Algorithm

Fix $\tau = 30\%$. Treat these market clearing conditions as nested problems:

- ▶ Define $D(r, T)$ to be the government's budget deficit
- ▶ Define $ED(r, T)$ to be the excess demand for capital
- ▶ Algorithm:
 1. For any given r , solve for the T that balances the government's budget (solving and simulating the model). That is, solve the root finding problem

$$D(r, T) = 0$$

as a function of T , holding r fixed. Call the results $T^*(r)$

2. Solve the root finding problem for

$$ED(r, T^*(r)) = 0$$

Call the result r^*

- ▶ Our final (r, T) are $(r^*, T^*(r^*))$.

We'll go over code for how to do this in tutorial

Tax Reform in Aiyagari: Solution Algorithm

Fix $\tau = 30\%$. Treat these market clearing conditions as nested problems:

- ▶ Define $D(r, T)$ to be the government's budget deficit
- ▶ Define $ED(r, T)$ to be the excess demand for capital
- ▶ Algorithm:
 1. For any given r , solve for the T that balances the government's budget (solving and simulating the model). That is, solve the root finding problem

$$D(r, T) = 0$$

as a function of T , holding r fixed. Call the results $T^*(r)$

2. Solve the root finding problem for

$$ED(r, T^*(r)) = 0$$

Call the result r^*

- ▶ Our final (r, T) are $(r^*, T^*(r^*))$.

We'll go over code for how to do this in tutorial

Tax Reform in Aiyagari: Solution Algorithm

Fix $\tau = 30\%$. Treat these market clearing conditions as nested problems:

- ▶ Define $D(r, T)$ to be the government's budget deficit
- ▶ Define $ED(r, T)$ to be the excess demand for capital
- ▶ Algorithm:
 1. For any given r , solve for the T that balances the government's budget (solving and simulating the model). That is, solve the root finding problem

$$D(r, T) = 0$$

as a function of T , holding r fixed. Call the results $T^*(r)$

2. Solve the root finding problem for

$$ED(r, T^*(r)) = 0$$

Call the result r^*

- ▶ Our final (r, T) are $(r^*, T^*(r^*))$.

We'll go over code for how to do this in tutorial

Tax Reform in Aiyagari: Solution Algorithm

Fix $\tau = 30\%$. Treat these market clearing conditions as nested problems:

- ▶ Define $D(r, T)$ to be the government's budget deficit
- ▶ Define $ED(r, T)$ to be the excess demand for capital
- ▶ Algorithm:
 1. For any given r , solve for the T that balances the government's budget (solving and simulating the model). That is, solve the root finding problem

$$D(r, T) = 0$$

as a function of T , holding r fixed. Call the results $T^*(r)$

2. Solve the root finding problem for

$$ED(r, T^*(r)) = 0$$

Call the result r^*

- ▶ Our final (r, T) are $(r^*, T^*(r^*))$.

We'll go over code for how to do this in tutorial