ECON-GA 1025 Macroeconomic Theory I

Lecture 8

John Stachurski

Fall Semester 2018

Today's Lecture

- Comments on assessment
- Finish: Numerical methods for tracking distributions
- Start: Job Search

Shifting office hours to Wed 4:00–5:00

Style of exam:

- broad understanding of all course material
- applications of ideas

Exam prep:

- Weekly assignments
- Other **Ex.** in slides
- Review logic and applications in slides

Wealth Distributions: Estimation by Monte Carlo (Continued)

In the last lecture we studied estimation of the time t distribution Ψ_t using Monte Carlo

Method:

1. Compute sample $\{w_t^m\}$, time t wealth of m independent households

2. Calculate the empirical distribution

$$F_t^m(x) := \frac{1}{m} \sum_{i=1}^{m} \mathbb{1}\{w_t^i \leqslant x\}$$

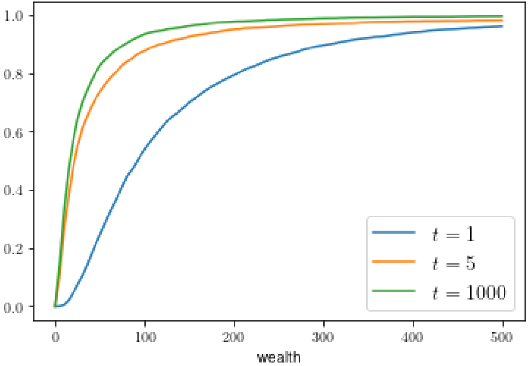

Figure: The empirical distribution F_m^t for different values of t

But we know that Ψ_t can be represented by a density ψ_t

This is structure that we would like to exploit

- helps when we get to high dimensional problems
- helps extract information from the tails

Unfortunately there is no natural estimator of densities that

- works in every setting (like the empirical distribution does)
- is always unbiased and consistent

Why?

- Empirical distributions just reflect the sample
- Density estimates must make statements about probability mass in the **neighborhood** of each observation

Let's look at our options

Option 1. **Nonparametric kernel density estimation**, where

$$\hat{f}_t^m(x) = \frac{1}{mh} \sum_{i=1}^{m} K\left(\frac{x - w_t^i}{h}\right)$$

Here

- K is a density, called the **kernel**
- h is the **bandwidth** of the estimator

Idea:

- Put a smooth bump on each data point and then sum
- Larger h means smoother estimate

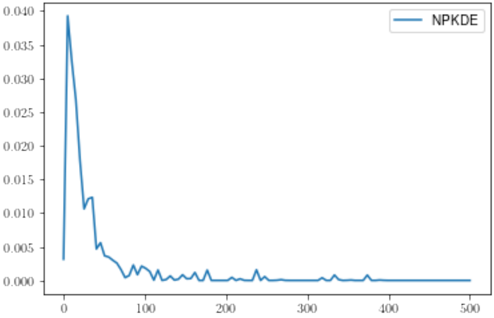

Figure: NPKDE of ψ_t using Scikit Learn ($t = 100$, $m = 500$)

Option 2. The **look ahead estimator**

$$\ell_t^m(w') := \frac{1}{m} \sum_{i=1}^{m} \pi(w_{t-1}^i, w')$$

Notes:

- sample $\{w_{t-1}^i\}$ is from time $t - 1$
- $\pi(w, w') = \int \varphi(w' - zs(w))\nu(\mathrm{d}z)$

Observe that we are combining data and model

- more information than just the sample

The estimator

$$\ell_t^m(w') := \frac{1}{m} \sum_{i=1}^{m} \pi(w_{t-1}^i, w')$$

is **unbiased**:

$$\mathbb{E}[\ell_t^m(w')] = \frac{1}{m} \sum_{i=1}^{m} \mathbb{E}[\pi(w_{t-1}^i, w')]$$

$$= \int \pi(w, w') \psi_{t-1}(w) \, \mathrm{d}w = \psi_t(w')$$

From the SLLN, we also have

$$\ell_t^m(w') \to \mathbb{E}[\pi(w_{t-1}^i, w')] = \psi_t(w') \quad \text{as} \quad m \to \infty$$

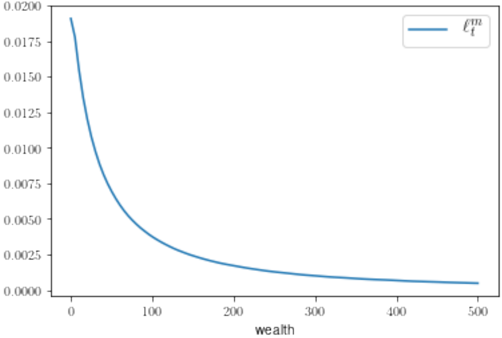

Figure: The look ahead estimate of ψ_t ($t = 100$, $m = 500$)

Stability of the Wealth Process

Lemma. The dynamical system $(\mathcal{D}, \Pi)$ corresponding to the wealth process

$$w_{t+1} = R_{t+1} s(w_t) + y_{t+1}$$

is globally stable whenever

(a) y_t has finite first moment, $\varphi \gg 0$ and

(b) $\mathbb{E}[R_t] s(w) \leqslant \lambda w + L$ for some $\lambda < 1$ and $L < \infty$

If ψ^* is the stationary density and $\int |h(w)| \psi^*(w) \, \mathrm{d}w < \infty$, then, with prob one,

$$\lim_{n \to \infty} \frac{1}{n} \sum_{t=1}^{n} h(w_t) = \int h(w) \psi^*(w) \, \mathrm{d}w$$

Proof: Follows from our stability result for

$$X_{t+1} = \zeta_{t+1} g(X_t) + \eta_{t+1}$$

Ex. Apply the last result to the case

$$s(w) = \mathbb{1}\{w > \bar{w}\} s_0 w \qquad (w \geqslant 0)$$

- Here s_0 and $\bar{w}$ are positive parameters
- What conditions do you need to impose on the parameters in the model in order to get global stability?
- Can you give some interpretation?

The **stationary density look ahead estimator**:

$$\ell_n^*(w') := \frac{1}{n} \sum_{t=1}^{n} \pi(w_t, w')$$

- sample is a **single** time series $\{w_t\}$ generated by simulation

Consistent for $\psi^*(w')$, since, with probability one as $n \to \infty$,

$$\ell_n^*(w') = \frac{1}{n} \sum_{t=1}^{n} \pi(w_t, w') \to \int \pi(w, w') \psi^*(w) \, \mathrm{d}w = \psi^*(w')$$

- Is it unbiased?

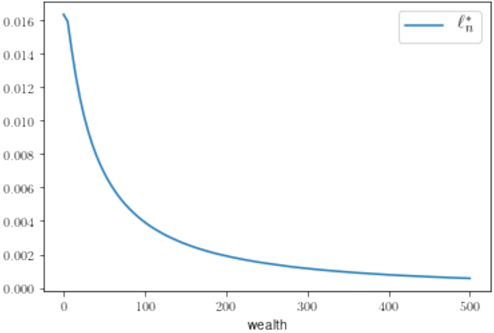

Figure: The stationary density look ahead estimator of the wealth distribution

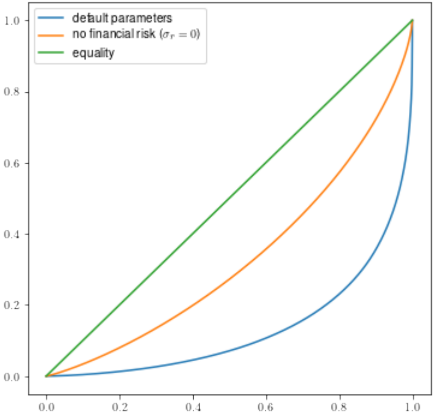

Figure: Lorenz curve, wealth distribution at default parameters

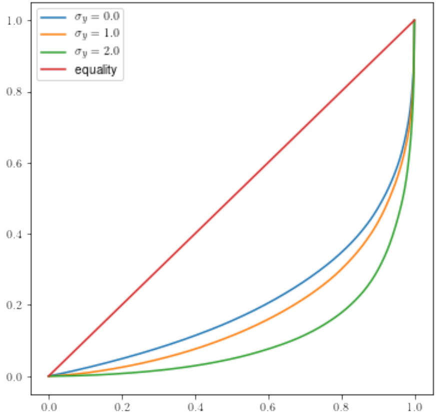

Figure: Lorenz curves with increasing variance in labor income

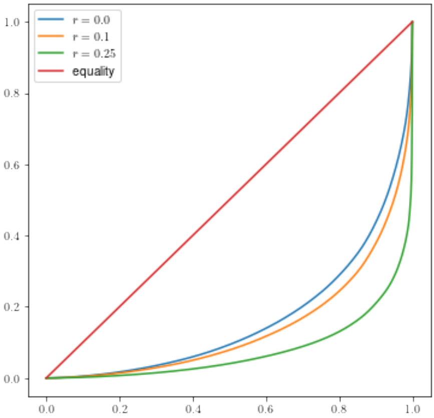

Figure: Lorenz curves with increasing rate of return on wealth

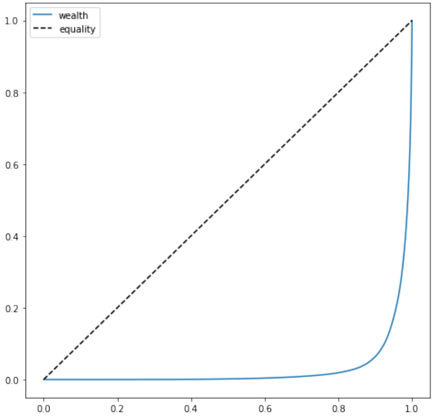

Figure: For comparison: wealth distribution in the US (SCF 2016)

See notebooks

- wealth_sk_plots.ipynb
- wealth_ineq_plots.ipynb

New Topic: Job Search

Our first deep dive into dynamic programming

- An integral part of labor and macroeconomics
- Relatively simple (binary choice)

Related to

- Optimal stopping
- Firm entry and exit decisions
- Pricing American options
- etc.

As discussed earlier

- Unemployed agent seeks to maximize

$$\mathbb{E} \sum_{t=0}^{\infty} \beta^t y_t$$

- Observes an employment opportunity with wage offer w_t
- Wage offers are IID and drawn from distribution φ
- Acceptance means lifetime value $w_t / (1 - \beta)$
- Rejection yields unemployment compensation $c \geqslant 0$ and a new offer next period

Overview

The **value function** $v^*(w) :=$ the maximal value that can be extracted from any given state w

We **will prove that** it satisfies the **Bellman equation**

$$v^*(w) = \max \left\{ \frac{w}{1-\beta}, \, c + \beta \int v^*(w') \varphi(\mathrm{d}w') \right\} \qquad (w \in \mathbb{R}_+)$$

Optimal policy is then obtained via

$$\sigma^*(w) = \mathbb{1} \left\{ \frac{w}{1-\beta} \geqslant c + \beta \int v^*(w') \varphi(\mathrm{d}w') \right\}$$

To calculate the optimal policy we need to evaluate $\int v^*(w')\varphi(\mathrm{d}w')$

To compute v^*, we introduce the **Bellman operator**

$$Tv(w) := \max\left\{\frac{w}{1-\beta}, c + \beta \int v(w')\varphi(\mathrm{d}w')\right\}$$

Fixed points of T exactly coincide with solutions to the Bellman equation

$$v(w) = \max\left\{\frac{w}{1-\beta}, c + \beta \int v(w')\varphi(\mathrm{d}w')\right\}$$

Simplifying assumption:

- There exists an $M \in \mathbb{R}_+$ such that $\int_0^M \varphi(\mathrm{d}w) = 1$

Later we will show this assumption can be weakened

But for now it's convenient...

Case 1: Continuous Wage Draws

Assumption. The offer distribution φ is a **density** supported on $[0, M]$

Any w in $[0, M]$ is possible so v^* needs to be defined on $[0, M]$

Leads us to seek a fixed point of T in $\mathscr{C} :=$ all continuous functions on $[0, M]$ paired with

$$d_\infty(f, g) := \|f - g\|_\infty, \qquad \|g\|_\infty := \sup_{w \in [0,M]} |g(w)|$$

- $(\mathscr{C}, d_\infty)$ is a complete metric space

Question: Why restrict ourselves to continuous functions?

Proposition. In this setting, T is a contraction of modulus β on $\mathscr{C}$

In particular,

1. T has a unique fixed point in $\mathscr{C}$
2. that fixed point is equal to the value function v^* and
3. if $v \in \mathscr{C}$, then $\|T^n v - v^*\|_\infty \leqslant O(\beta^n)$

For now let's take (2) as given — we'll prove it soon

- Remainder will be verified if we show T is a contraction of modulus β on $(\mathscr{C}, d_\infty)$

We use the elementary bound

$$|\alpha \vee x - \alpha \vee y| \leqslant |x - y| \qquad (\alpha, x, y \in \mathbb{R})$$

Fixing f, g in $\mathscr{C}$ and $w \in [0, M]$,

$$\begin{aligned} |Tf(w) - Tg(w)| &\leqslant \left| \beta \int f(w')\varphi(w')\, \mathrm{d}w' - \beta \int g(w')\varphi(w')\, \mathrm{d}w' \right| \\ &= \beta \left| \int [f(w') - g(w')]\varphi(w')\, \mathrm{d}w' \right| \\ &\leqslant \beta \int |f(w') - g(w')|\varphi(w')\, \mathrm{d}w' \leqslant \|f - g\|_\infty \end{aligned}$$

Taking the supremum over all $w \in [0, M]$ leads to

$$\|Tf - Tg\|_\infty \leqslant \beta \|f - g\|_\infty$$

Ex. Show that T maps the set of increasing continuous convex functions on the interval $[0, M]$ to itself

Ex. Show that v^* is increasing and convex on $[0, M]$

Case 2: Discrete Wage Draws

Let's swap the density assumption for a discrete distribution

Assumption. The offer distribution φ is supported on finite set W with probabilities $\varphi(w)$, $w \in W$

- Now v^* need only be defined on these points

Hence we define T on $\mathbb{R}^W$ by

$$Tv(w) = \max\left\{\frac{w}{1-\beta},\, c + \beta \sum_{w' \in W} v(w')\varphi(w')\right\} \qquad (w \in W)$$

- $(\mathbb{R}^W, d_\infty)$ is a complete metric space

Proposition. T is a contraction of modulus β on $\mathbb{R}^W$

In particular,

1. T has a unique fixed point in $\mathbb{R}^W$,
2. that fixed point is equal to the value function v^* and
3. if $v \in \mathbb{R}^W$, then $\|T^n v - v^*\|_\infty \leqslant O(\beta^n)$

Ex. Prove that T is a contraction of modulus β on $(\mathbb{R}^W, d_\infty)$

To compute the optimal policy we can use **value function iteration**

1. Start with arbitrary $v \in \mathbb{R}^W$
2. iterate with T until $v_k := T^k v$ is a good approximation to v^*

Then compute

$$\sigma_k(w) := \mathbb{1}\left\{\frac{w}{1-\beta} \geqslant c + \beta \sum_{w'} v_k(w')\varphi(w')\right\}$$

Approximately optimal when v_k is close to v^*

Error bounds available...

Rearranging the Bellman Equation

Actually, for this particular problem, there's an easier solution method

- involves a "rearrangement" of the Bellman equation
- shifts us to a lower dimensional problem

Recall: a function v satisfies the Bellman equation if

$$v(w) := \max\left\{\frac{w}{1-\beta}, c + \beta \int v(w')\varphi(\mathrm{d}w')\right\}$$

Taking v as given, consider

$$h := c + \beta \int v(w')\varphi(\mathrm{d}w')$$

Using h to eliminate v from the Bellman equation yields

$$h = c + \beta \int \max\left\{\frac{w'}{1-\beta}, h\right\} \varphi(\mathrm{d}w')$$

Ex. Verify this

We now seek an $h \in \mathbb{R}_+$ satisfying

$$h = c + \beta \int \max \left\{ \frac{w'}{1-\beta}, h \right\} \varphi(\mathrm{d}w')$$

Solution h^* is the continuation value

Optimal policy can be written as

$$\sigma^*(w) = \mathbb{1} \left\{ \frac{w}{1-\beta} \geqslant h^* \right\} \qquad (w \in \mathbb{R}_+)$$

Alternatively,

$$\sigma^*(w) = \mathbb{1}\{w \geqslant w^*\} \quad \text{where} \quad w^* := (1-\beta)h^*$$

The term w^* is called the **reservation wage**

To solve for h^* we introduce the mapping

$$g(h) = c + \beta \int \max \left\{ \frac{w'}{1-\beta}, h \right\} \varphi(\mathrm{d}w') \qquad (h \in \mathbb{R}_+)$$

Any solution to $h = c + \beta \int \max \left\{ \frac{w'}{1-\beta}, h \right\} \varphi(\mathrm{d}w')$ is a fixed point of g and vice versa

Assumption. The distribution φ has finite first moment

Ex. Confirm that

- g is a well defined map from $\mathbb{R}_+$ to itself
- g is a contraction map on $\mathbb{R}_+$ under the usual Euclidean distance

Conclude that g has a unique fixed point in $\mathbb{R}_+$

For computation it is somewhat easier to work with the case where wages are bounded

Ex. Suppose that $\mathbb{P}\{w_t \leqslant M\} = 1$ for some positive constant M

- Confirm that g maps $[0, K]$ to itself, where

$$K := \frac{\max\{M, c\}}{1 - \beta}$$

- Conclude that g has a fixed point in $[0, K]$, which is the unique fixed point of g in $\mathbb{R}_+$

See notebook `iid_job_search.ipynb`

Parametric Monotonicity

Recall this result:

Fact. If (M, g_1) and (M, g_2) are dynamical systems such that

1. g_2 is isotone and dominates g_1 on M
2. (M, g_2) is globally stable with unique fixed point u_2,

then $u_1 \preceq u_2$ for every fixed point u_1 of g_1

Now consider

$$g(h) = c + \beta \int \max \left\{ \frac{w'}{1 - \beta}, h \right\} \varphi(\mathrm{d}w')$$

This map is

1. globally stable
2. an isotone self-map on $\mathbb{R}_+$

Hence any parameter that shifts up the function g pointwise on $\mathbb{R}_+$ also shifts up h^*

Ex. Show that

1. the continuation value h^* is increasing in unemployment compensation c
2. the reservation wage w^* is increasing in c

Interpret

Shifting the Offer Distribution

How do shifts in this distribution affect the reservation wage?

Intuition: “more favorable” wage distribution would tend to increase the reservation wage

- the agent can expect better offers

What does “more favorable” mean for offer distributions?

One possible answer: **(first order) stochastic dominance**

First Order Stochastic Dominance

Definition. Distribution φ is **stochastically dominated** by distribution ψ (write $\varphi \preceq_{\text{SD}} \psi$) if

$$\int u(x)\varphi(\mathrm{d}x) \leqslant \int u(x)\psi(\mathrm{d}x) \text{ for all } u \in ibc\mathbb{R}_+$$

With $ibm\mathbb{R}_+$ as the increasing bounded Borel measurable functions, this is equivalent:

$$\int u(x)\varphi(\mathrm{d}x) \leqslant \int u(x)\psi(\mathrm{d}x) \text{ for all } u \in ibm\mathbb{R}_+$$

Interpretation: Anyone with increasing utility likes ψ better

Let φ and ψ be two wage distributions on $\mathbb{R}_+$ with finite first moment

Let

- w^*_φ and w^*_ψ be the associated reservation wages
- h^*_φ and h^*_ψ be the associated continuation values

Assume both are supported on $[0, M]$

Lemma. If $\varphi \preceq_{\mathrm{SD}} \psi$, then $w^*_\varphi \leqslant w^*_\psi$

Proof: Let ψ and φ have the stated properties

It suffices to show that $h^*_\varphi \leqslant h^*_\psi$

We aim to show that

$$g(h) = c + \beta \int \max \left\{ \frac{w'}{1-\beta}, h \right\} \varphi(\mathrm{d}w')$$

increases at any h if we shift up the offer distribution in $\preceq_{\mathrm{SD}}$

Sufficient: given $\varphi \preceq_{\mathrm{SD}} \psi$ and $h \geqslant 0$,

$$\int \max \left\{ \frac{w'}{1-\beta}, h \right\} \varphi(\mathrm{d}w') \leqslant \int \max \left\{ \frac{w'}{1-\beta}, h \right\} \psi(\mathrm{d}w')$$

This follows directly from the definition of stochastic dominance (why?)