Dynamic Programming: Estimation

ECON 34430: Topics in Labor Markets

T. Lamadon (U of Chicago)

Winter 2016

Agenda

1. Introduction
   - General formulation
   - Assumptions
   - Estimation in general
2. Estimation of Rust models
   - Example of Rust and Phelan
   - NFXP
   - Partial likelihood
   - Hotz and Miller

- I follow Aguirregabiria and Mira (2010)

Introduction

General formulation

- time is discrete indexed by t
- agents are indexed by i
- state of the world at time t:
  - state s_{it}
  - control variable a_{it}
- agent preferences are

$$\sum_{j=0}^{T} \beta^j U(a_{i,t+j}, s_{i,t+j})$$

- agents have beliefs about state transitions $F(s_{i,t+1} | a_{it}, s_{it})$

Decision

- The agent Bellman equation is

$$V(s_{it}) = \max_{a \in A} \left[U(a, s_{it}) + \beta \int V(s_{i,t+1}) \mathsf{d}F(s_{i,t+1} | a, s_{it}) \right]$$

- we define the **choice specific value function** or **Q-value** :

$$v(a, s_{it}) = U(a, s_{it}) + \beta \int V(s_{i,t+1}) \mathsf{d}F(s_{i,t+1} | a, s_{it})$$

- and the **policy function**:

$$\alpha(s_{it}) = \arg\max_{a \in A} v(a, s_{it})$$

Data

- a_{it} is the action
- x_{it} is a subset of $s_{it} = (x_{it}, \epsilon_{it})$
  - ϵ_{it} gives a source of variation of the individual level
  - can have structural interpretation (pref shocks)
- y_{it} is a payoff variable $y_{it} = Y(a_{it}, x_{it}, \epsilon_{it})$
  - then $U(a_{it}, s_{it}) = \widetilde{U}(y_{it}, a_{it}, s_{it})$
  - earnings is a good example
- Data $= \{a_{it}, x_{it}, y_{it} : i = 1, 2...N; t = 1, 2..T_i\}$
  - usually N is large, T_i is small

Estimation

- parameter θ affects $U(a, s_{it})$ and $F(s_{i,t+1}|a, s_{it})$
- we have an estimation criteria $g_N(\theta)$
- example is likelihood $g_N(\theta) = \sum_i l_i(\theta)$:

$$l_i(\theta) = \log Pr[\alpha(x_{it}, \epsilon_{it}, \theta){=}a_{it},\ Y(a_{it}, x_{it}, \epsilon_{it}, \theta){=}y_{it}, x_{it}, t = 1..T_i|\theta]$$

- in general, we need to solve for $\alpha(\cdot)$ for each value of θ
- the particular form of $l_i(\theta)$ depends on relation between observables and unobservables

Econometric assumptions

Assumptions

AS Additive separability

- $U(a, x_{it}, \epsilon_{it}) = u(a, x_{it}) + \epsilon_{it}(a)$
- $\epsilon_{it}(a)$ is 1-dimensional, mean 0, unbounded
- there is one per choice, ϵ_{it} is $(J+1)$-dimensional

IID IID unobservables

- ϵ_{it} are iid across agents and time
- distribution $G_\epsilon(\epsilon_{it})$

CLOGIT

- ϵ_{it} are independent across alternatives and type-1 extreme value distribution

Assumptions

CI-X Conditional independence of future x

- $x_{i,t+1} \perp \epsilon_{it} | a_{it}, x_{it}$
- θ_f describes $F(x_{i,t+1} | a_{it}, x_{it})$
- future realization of the state do not depend on the shock

CI-Y Conditional independence of y

- $y_{i,t} \perp \epsilon_{it} | a_{it}, x_{it}$
- θ_Y describes $F(y_{it})(a_{it}, x_{it})$
- rules out Heckman type selection

DIS Discrete support for x

- x_{it} is finite

Example 1

Retirement model, Rust and Phelan (1997)

Model

- consumption is $c_{it} = y_{it} - hc_{it}$ (hc_{it} is health care expenditure)
- earnings is $y_{it} = a_{it} w_{it} + (1 - a_{it}) b_{it}$
- m_{it} is marital status Markov, h_{it} is health status Markov
- pp_{it} is pension point with $F_{pp}(pp_{it+1} | w_{it}, pp_{it})$
- preferences:

$$U(a_{it}, x_{it}, \epsilon_{it}) = E[c_{it}^{\theta_{u1}} | a_{it}, x_{it}] \cdot exp\left\{\theta_{u2} + \theta_{u3} h_{it} + \theta_{u4} m_{it} + \theta_{u5} \frac{t_{it}}{1 + t_{it}}\right\} - \theta_{u6} a_{it} + \epsilon_{it}(a_{it})$$

- wages:

$$w_{it} = exp\left\{\theta_{w1} + \theta_{w2} h_{it} + \theta_{w3} m_{it} + \theta_{w4} \frac{t_{it}}{1 + t_{it}} + \theta_{w5} pp_{it} + \xi_{it}\right\}$$

Retirement model, Rust and Phelan (1997)

Assumptions

- **CI-Y** holds since ξ_{it} is
  - serially uncorrelated
  - independent of x_{it}, ϵ_{it}
  - unknown at the time of decision
- **AS**, additive separability is also assumed here
  - this implies that there is no uncertainty about future marginal utilities of consumption
- **CI-X** also holds
  - future x_{it} do not depend on the current shock, just current action (it does depend on ξ_{it} though)
- **DIS** and **IID** also hold

Retirement model, Rust and Phelan (1997)

Implications

- under **CI-X** and **IID** we get that

$$F(x_{i,t+1}, \epsilon_{i,t+1} | a_{it}, x_{it}, \epsilon_{it}) = G_\epsilon(\epsilon_{i,t+1}) F_x(x_{i,t+1} | a_{it}, x_{it})$$

- the unobserved ϵ_{it} drops from the state space, we can look at **integrated value function** or **Emax** function:

$$\bar{V}(x_{it}) = \int \max_{a \in A} \Big\{ u(a, x_{it}) + \epsilon_{it}(a) + \beta \sum_{x_{i,t+1}} \bar{V}(x_{i,t+1}) f_x(x_{i,t+1} | a, x_{it}) \Big\}$$

- the computational complexity is only driven by the size of the support of x
- we define

$$v(a, x_{it}) = u(a, x_{it}) + \beta \sum_{x_{i,t+1}} \bar{V}(x_{i,t+1}) f_x(x_{i,t+1} | a, x_{it})$$

Retirement model, Rust and Phelan (1997)

Implications

- under **CI-X** and **IID** we get that

$$F(x_{i,t+1}, \epsilon_{i,t+1} | a_{it}, x_{it}, \epsilon_{it}) = G_\epsilon(\epsilon_{i,t+1}) F_x(x_{i,t+1} | a_{it}, x_{it})$$

- the unobserved ϵ_{it} drops from the state space, we can look at **integrated value function** or **Emax** function:

$$\bar{V}(x_{it}) = \int \max_{a \in A} \Big\{ u(a, x_{it}) + \epsilon_{it}(a) + \beta \sum_{x_{i,t+1}} \bar{V}(x_{i,t+1}) f_x(x_{i,t+1} | a, x_{it}) \Big\}$$

- the computational complexity is only driven by the size of the support of x
- we define

$$v(a, x_{it}) = u(a, x_{it}) + \beta \sum_{x_{i,t+1}} \bar{V}(x_{i,t+1}) f_x(x_{i,t+1} | a, x_{it})$$

Retirement model, Rust and Phelan (1997)

Implications 2

- under **CI-X** and **IID** the log-likelihood is separable:

$$l_i(\theta) = \sum_t \log P(a_{it}|x_{it}, \theta) + \sum_t \log f_Y(y_{it}|a_{it}, x_{it}, \theta_Y)$$
$$+ \sum_t \log f_X(x_{i,t+1}|a_{it}, x_{it}, \theta_f) + \log Pr[x_{i1}|\theta]$$

- note how here each term can be tackled separately, in particular, the wage equation and the transition probabilities can be estimated directly from the data

- $P(a_{it}|x_{it}, \theta)$ is referred to as **Conditional Choice Probability** or **CCP**

Retirement model, Rust and Phelan (1997)

Implications 3

- The CPP is given by:

$$P(a_{it}=a|x_{it},\theta) = \int I[\alpha(x_{it},\epsilon_{it};\theta) = a]\mathrm{d}G_\epsilon(\epsilon_{it}) = \int I[v(a,x_{it}) + \epsilon_{it}(a) > v(a',x_{it}) \text{ for all } a']\mathrm{d}G_\epsilon(\epsilon_{it})$$

- If we add to this the **CLOGIT** assumption we get:

$$\bar{V}(x_{it}) = \log \sum_{a\in A} \exp\Big\{u(a,x_{it}) + \epsilon_{it}(a) + \beta \sum_{x_{i,t+1}} \bar{V}(x_{i,t+1}) f_x(x_{i,t+1}|a,x_{it})\Big\}$$

- we do not even need to do a maximization!

Retirement model, Rust and Phelan (1997)

Implications 4

- using:

$$v(a, x_{it}) = u(a, x_{it}) + \beta \sum_{x_{i,t+1}} \bar{V}(x_{i,t+1}) f_x(x_{i,t+1} | a, x_{it})$$

- we get our CCP:

$$P(a | x_{it}, \theta) = \frac{exp\{v(a, x_{it})\}}{\sum_j exp\{v(a_j, x_{it})\}}$$

Estimation procedures

Nested fixed point

1. pick parameter θ
2. solve for the policy $\alpha(\cdot)$

$$V(s_{it}) = \max_{a \in A} \left[U(a, s_{it}) + \beta \int V(s_{i,t+1}) \mathrm{d}F(s_{i,t+1}|a, s_{it}) \right]$$

3. compute full likelihood

$$l_i(\theta) = \log Pr[\alpha(x_{it}, \epsilon_{it}, \theta){=}a_{it}, Y(a_{it}, x_{it}, \epsilon_{it}, \theta){=}y_{it}, x_{it}, t = 1..T_i|\theta]$$

4. update θ (use gradient method or other ...)

- very intuitive
- provide MLE estimate
- very costly, sometimes you need to simulate integral (obejctive might not be smooth)
- how closely to solve inside problem

Nested fixed point

1. pick parameter θ
2. solve for the policy $\alpha(\cdot)$

$$V(s_{it}) = \max_{a \in A} \left[U(a, s_{it}) + \beta \int V(s_{i,t+1}) \mathsf{d} F(s_{i,t+1} | a, s_{it}) \right]$$

3. compute full likelihood

$$l_i(\theta) = \log Pr[\alpha(x_{it}, \epsilon_{it}, \theta){=}a_{it}, Y(a_{it}, x_{it}, \epsilon_{it}, \theta){=}y_{it}, x_{it}, t = 1..T_i | \theta]$$

4. update θ (use gradient method or other ...)

- very intuitive
- provide MLE estimate
- very costly, sometimes you need to simulate integral (obejctive might not be smooth)
- how closely to solve inside problem

Rust partial likelihood

- use the separability of the likelihood under **CI-X** and **IID**:

$$l_i(\theta) = \sum_t \log P(a_{it}|x_{it},\theta) + \sum_t \log f_Y(y_{it}|a_{it},x_{it},\theta_Y) + \sum_t \log f_X(x_{i,t+1}|a_{it},x_{it},\theta_f) + \log Pr[x_{i1}|\theta]$$

1. estimate $f_Y(y_{it}|a_{it},x_{it},\theta_Y)$ and $f_X(x_{i,t+1}|a_{it},x_{it},\theta_f)$
2. then iterate on θ_u only, using a small NFXP, and using **CLOGIT**

Rust partial likelihood

part 2

- solve Bellman, where $f_x(x_{i,t+1}|a, x_{it})$ is given :

$$\bar{V}(x_{it}) = \log \sum_{a \in A} \exp \Big\{ u(a, x_{it}) + \epsilon_{it}(a) + \beta \sum_{x_{i,t+1}} \bar{V}(x_{i,t+1}) f_x(x_{i,t+1}|a, x_{it}) \Big\}$$

- with:

$$v(a, x_{it}) = u(a, x_{it}) + \beta \sum_{x_{i,t+1}} \bar{V}(x_{i,t+1}) f_x(x_{i,t+1}|a, x_{it})$$

$$P(a|x_{it}, \theta) = \frac{exp\{v(a, x_{it})\}}{\sum_j exp\{v(a_j, x_{it})\}}$$

Rust partial likelihood

- take advantage of log-likelihood separability
- loose on estimator efficiency, huge gains in computational efficiency
- **CLOGIT** and finite T allows for exact solutions
- Rust type models are the building block to tackle dynamic games

Hotz and Miller (1993) approach

- under RUST assumptions, we don't need to solve the model
- intuition: agents in the data have already done it for us!
- the costly part of Rust partial likelihood is to solve at each θ_u:

$$\bar{V}(x_{it}) = \log \sum_{a \in A} \exp \Big\{ u(a, x_{it}) + \epsilon_{it}(a) + \beta \sum_{x_{i,t+1}} \bar{V}(x_{i,t+1}) f_x(x_{i,t+1} | a, x_{it}) \Big\}$$

- the transitions in the data are according to the correct policy, we can estimate directly the CCP.

Hotz and Miller (1993) approach

- consider a linear utility model $u(a, x, \theta_u) = z(a, x)'\theta_u$
- Hotz and Miller show that

$$v(a, x_t, \theta) = \tilde{z}(a, x_t, \theta)'\theta_u + \tilde{e}(a, x_t, \theta)$$

- where $\tilde{z}(a, x_t, \theta)$ and $\tilde{e}(a, x_t, \theta)$ depend on θ only through parameters in the transition probabilities F_x and the CCPs of the individuals.

Hotz and Miller (1993) approach

- for instance:

$$\begin{aligned}
&\tilde{z}(a, x_t, \theta) \\
&= z(a, x_t) + \sum_{\tau=0}^{T-\tau} \beta^\tau \mathbb{E}_{x_{t+\tau},\epsilon_{t+\tau}}[z(\alpha(x_{t+\tau}, \epsilon_{t+\tau}, \theta), x_{t+\tau})|a_t = a, x_t] \\
&= z(a, x_t) + \sum_{\tau=0}^{T-\tau} \beta^\tau \mathbb{E}_{x_{t+\tau}}[\sum_j P(a_{t+\tau}{=}a_j|x_{t+\tau}) \cdot z(a_{t+\tau}, x_{t+\tau})|a_t{=}a, x_t]
\end{aligned}$$

- however we can measure $P(a_{t+\tau}{=}a_j|x_{t+\tau})$ directly from the data

- we end up with logit restrictions of the form

$$I(a_{it} = a_j) - \frac{\exp\left(\tilde{z}(a_j, x_{it})\theta_u + \tilde{e}(a_j, x_{it})\right)}{\sum_j' \exp\left(\tilde{z}(a_j', x_{it})\theta_u + \tilde{e}(a_j', x_{it})\right)} = 0, \text{for each } j, t, x_{it}$$

Hotz and Miller (1993) approach

recap

1. estimate F_x and F_y using partial likelihood
2. estimate conditional choice probabilities $P(a_j|x_{it})$
3. construct $\tilde{z}$ and $\tilde{e}$
4. estimate θ_u using logistic expression

- this is computational trivial (convex problem)
- we never have to solve a Dynamic problem!
- of course it relies on quite strong assumptions

Eckstein-Keane-Wolpin type models

Occupation of young men, Keane and Wolpin (1997)

Model

- start at age 16, to age T
- chooses stay home ($a_{it} = 0$), school ($a_{it} = 4$), work among 3 occupations ($a_{it} = 1, 2, 3$)
- preferences where h_{it} is years of schooling:

$$
\begin{aligned}
U(0, s_{it}) &= \omega_i(0) + \epsilon_{it}(0) \\
U(4, s_{it}) &= \omega_i(4) - \theta_{tc1} I[h_{it} \geq 12] - \theta_{tc2} I[h_{it} \geq 16] + \epsilon_{it}(4) \\
U(a, s_{it}) &= W_{it}(a)
\end{aligned}
$$

- wages, where r_a is an occupation price, $k_{it}(a)$ is experience

$$
W_{it}(a) = r_a exp\Big\{\omega_i(a) + \theta_{a1} h_{it} + \theta_{a2} k_{it}(a) - \theta_{a3} k_{it}(a)^2 + \epsilon_{it}(a)\Big\}
$$

Occupation of young men, Keane and Wolpin (1997)

Data structure

- $\epsilon_{it}(a)$ jointly normal with unrestricted covariance structure, serially uncorrelated
- $\omega_i(a)$ choice specific permanent heterogeneity

- unobservable states: $\epsilon_{it}(a), \omega_i(a)$
- observable states: $x_{it} = \{h_{it}, t_{it}, k_{it}(a) : a = 1, 2, 3\}$
- $f_x(x_{i,t+1}|a_{it}, x_{it})$ is deterministic
- the payoff $W_{it}(a)$ is only observed at choice a_{it}

- do **IID**, **CI-Y**, **CI-X**, **AS**, **CLOGIT** hold here?

Occupation of young men, Keane and Wolpin (1997)

Assumptions

- $\epsilon_{it}(a) + \omega_i(a)$ are serially correlated, no **IID**
- $\epsilon_{it}(a)$ are normal and correlated, no **CLOGIT**
- $\epsilon_{it}(a)$ is observed before decision, no **CI-Y**
- we do have **CI-X**, but here F_x is trivial (just accumulation)
- is there anything we can do?

Occupation of young men, Keane and Wolpin (1997)

Latent structure

- Conditional on $\omega_i(a)$, $\epsilon_{it}(a)$ do satisfy **IID**
- using discrete $\Omega = \{\omega_l(a)\}_{l=1..L}$ we can factor the likelihood:

$$l_i(\theta, \Omega, \pi) = \log(\sum_l L_i(\theta, \omega^l)\pi_{j|x_1})$$

- and

$$L_i(\theta, \omega^l) = \prod_t p(a_{it}|x_{it}, \theta, \omega^l) \cdot f_Y(y_{it}|a_{it}, x_{it}, \theta, \omega^l)$$
$$\prod_t f_X(x_{i,t+1}|a_{it}, x_{it}, \theta_f, \omega^l)$$

- but we can apply an EM approach.
- however f_Y still depends on full θ.
- there is an initial condition difficulty

Occupation of young men, Keane and Wolpin (1997)

Latent structure

- Conditional on $\omega_i(a)$, $\epsilon_{it}(a)$ do satisfy **IID**
- using discrete $\Omega = \{\omega_l(a)\}_{l=1..L}$ we can factor the likelihood:

$$l_i(\theta, \Omega, \pi) = \log(\sum_l L_i(\theta, \omega^l)\pi_{j|x_1})$$

- and

$$L_i(\theta, \omega^l) = \prod_t p(a_{it}|x_{it}, \theta, \omega^l) \cdot f_Y(y_{it}|a_{it}, x_{it}, \theta, \omega^l)$$

$$\prod_t f_X(x_{i,t+1}|a_{it}, x_{it}, \theta_f, \omega^l)$$

- but we can apply an EM approach.
- however f_Y still depends on full θ.
- there is an initial condition difficulty

Occupation of young men, Keane and Wolpin (1997)

Latent structure

- conditional on ω, we have IID:

$$\bar{V}_l(x_{it}) = \int \max_{a \in A} \Big\{ u(a, x_{it}, \epsilon_{it}, \omega^l) + \beta \sum_{x_{i,t+1}} \bar{V}_l(x_{i,t+1}) f_x(x_{i,t+1} | a, x_{it}, \omega^l) \Big\}$$

- getting $\alpha(x_{it}, \epsilon_{it}, \omega^l)$ requires solving the DP L times
- integration is done by approximating $\bar{V}_l(x_{it})$ as a regression model and using monte carlo:

$$\bar{V}_l(x) = \phi(x)' \gamma_l + \nu$$

Arcidiacono and Jones

- if **CI-Y** holds conditional on ω^l, we get get separability within

$$L_i(\theta, \omega^l) = \prod_t p(a_{it}|x_{it}, \theta, \omega^l) \cdot f_Y(y_{it}|a_{it}, x_{it}, \theta_y, \omega^l)$$

$$\prod_t f_X(x_{i,t+1}|a_{it}, x_{it}, \theta_f, \omega^l)$$

- this means that within ω_l group, we can get the F_y and F_X directly when we have the posterior probabilities of ω^l

Recap

- Full nested fixed point
  - for each θ including π, ω^l, evaluate likelihood, and maximize
  - use numerical integration
- EM approach
  - start with guess of θ
  - compute posterior $\pi_i(l)$
  - update θ
  - use numerical integration
- EM approach
  - start with guess of θ including
  - compute posterior $\pi_i(l)$
  - update θ using separability!
  - use numerical integration

References

Aguirregabiria, V., and P. Mira (2010): "Dynamic discrete choice structural models: A survey," *J. Econom.*, 156(1), 38–67.