

Estimation of dynamic discrete choice models

Jean-François Houde
Cornell University & NBER

February 23, 2018

Introduction: Dynamic Discrete Choices¹

- We start with an single-agent models of dynamic decisions:
 - ▶ Machine replacement and investment decisions: Rust (1987)
 - ▶ Renewal or exit decisions: Pakes (1986)
 - ▶ Inventory control: Erdem, Imai, and Keane (2003), Hendel and Nevo (2006)
 - ▶ Experience goods and bayesian learning: Erdem and Keane (1996), Akerberg (2003), Crawford and Shum (2005)
 - ▶ Demand for durable goods: Gordon (2010), Gowrisankaran and Rysman (2012), Lee (2013)
- This lecture will focus on econometrics methods, and next lecture will discuss mostly applications.
- Next, we will discuss questions related to the dynamic of industries:
 - ▶ Markov-perfect dynamic games
 - ▶ Empirical model of static and dynamic games

¹These lectures notes incorporate material from Victor Aguirregabiria's graduate IO slides at the University of Toronto.

Machine replacement and investment decisions

- Consider a firm producing a good at N plants (indexed by i) that operate independently.
- Each plant has a machine.
- Examples:
 - ▶ Rust (1987): Each plant is a Madison WI bus, and Harold Zucher is the plant operator.
 - ▶ Das (1992): Consider cement plants, where the machines are cement kiln.
 - ▶ Rust and Rothwell (1995): Study the maintenance of nuclear power plants.
- Related applications: Export decisions (Das et al. (2007)), replacement of durable goods (Adda and Cooper (2000), Gowrisankaran and Rysman (2012)).

Bus Replacement: Rust (1987)

- Profit function at time t :

$$\pi_t = \sum_{i=1}^N y_{it} - rc_{it}$$

where y_{it} is the plant's variable profit, and rc_{it} is the replacing cost of the machine.

- Replacement and depreciation:

- ▶ Replace cost:

$$rc_{it} = a_{it} \times RC(x_{it})$$

where $\partial RC(x)/\partial x \geq 0$ and $a_{it} = 1$ if the machine is replaced. In the application, $RC(x_{it}) = \theta_{R0} + \theta_{R1}x_{it}$.

- ▶ State variable: machine age x_{it} , choice-specific profit shock $\{\epsilon_{it}(0), \epsilon_{it}(1)\}$.
- ▶ Variable profits are decreasing in the age x_{it} of the aging, and increasing in profit shock $\epsilon_{it}(a_{it})$:

$$y_{ij} = Y((1 - a_{it})x_{it}, \epsilon_{it}(a_{it}))$$

where $\partial Y/\partial x < 0$.

Profits and Depreciation

- Variable profit: Step function

$$\pi_{it} = \begin{cases} Y(0, \epsilon_{it}(1)) - RC(x_{it}) & \text{If } a_{it} = 1 \\ Y(x_{it}, \epsilon_{it}(0)) & \text{Otherwise.} \end{cases}$$

- Aging/depreciation process:

$$\text{Deterministic: } x_{it+1} = (1 - a_{it})x_{it} + 1$$

$$\text{Stochastic: } x_{it+1} = (1 - a_{it})x_{it} + \xi_{t+1}$$

Note: In Rust (1987), x_{it} is bus mileage. It follows a random walk process with a log-normal distribution.

- **Assumptions:**

- 1 Additive separable (AS) profit shock:

$$Y((1 - a)x, \epsilon(a)) = \theta_{Y_0} + \theta_{Y_1}(1 - a)x + \epsilon(a)$$

- 2 Conditional independence (CI): $f(\epsilon_{t+1} | \epsilon_t, x_t) = f(\epsilon_{t+1})$
- 3 Aging follows is a discrete random-walk process: $x_{it} \in \{0, 1, \dots, M\}$ and matrix $F(x' | x, a)$ characterizes its controlled Markov transition process.

Dynamic Optimization

- Harold Zucher maximizes expected future profits:

$$V(a_{it}|x_{it}, \epsilon_{it}) = E \left(\sum_{\tau=0}^{\infty} \beta^{\tau} \pi_{it+\tau} \middle| x_{it}, \epsilon_{it}, a_{it} \right)$$

- **Recursive formulation:** Bellman equation

$$\begin{aligned} V(a|x, \epsilon) &= Y((1-a) \cdot x) - RC(a \cdot x) + \epsilon(a) \\ &\quad + \beta \sum_{x'} E_{\epsilon'} (V(x', \epsilon')) F(x'|x, a) \\ &= v(a, x) + \epsilon(a) \end{aligned}$$

where $V(x, \epsilon) \equiv \max_{a \in \{0,1\}} V(a|x, \epsilon)$.

- Optimal replacement decision:

$$a^* = \begin{cases} 1 & \text{If } v(1, x) - v(0, x) = \tilde{v}(x) > \epsilon(0) - \epsilon(1) = \tilde{\epsilon} \\ 0 & \text{Otherwise.} \end{cases}$$

- If $\{\epsilon(0), \epsilon(1)\}$ are distributed according to a T1EV distribution with unit variance:

Solution to the dynamic-programming (DP) problem

- Assumptions (1) and (2) imply that we only need numerically find a fixed-point to the “E_{max}” function $\bar{V}(x)$ (M elements):

$$\begin{aligned}\bar{V}(x) &= E_{\epsilon} \left(\max_a v(a, x) + \epsilon(a) \right) \\ &= E_{\epsilon} \left(\max_a \Pi(a, x) + \beta \sum_{x'} \bar{V}(x') F(x'|x, a) + \epsilon(a) \right) \\ &= \Gamma(x|\bar{V})\end{aligned}$$

where $\Pi(a, x) = Y((1 - a) \cdot x) - RC(a \cdot x)$, and $\Gamma(x|\bar{V})$ is a contraction mapping.

- Matrix form representation using the T1EV distribution assumption:

$$\begin{aligned}\bar{\mathbf{V}} &= \ln \left(\exp \left(\Pi(0) + \beta F(0) \bar{\mathbf{V}} \right) + \exp \left(\Pi(1) + \beta F(1) \bar{\mathbf{V}} \right) \right) + \gamma \\ &= \Gamma(\bar{\mathbf{V}})\end{aligned}$$

where γ is the Euler constant, $F(0)$ and $F(1)$ are two $M \times M$ conditional transition probability matrix.

Algorithm 1: Value Function Iteration

- Fixed objects:

- ▶ Payoffs ($M \times 1$):

$$\Pi(a) = \{\theta_0 + \theta_x(1 - a)x_i - RC(a \cdot x)\}_{i=1,\dots,M} \text{ for } a \in \{0, 1\}$$

- ▶ Conditional transition probability ($M \times M$): $F(a)$ for $a \in \{0, 1\}$

$$F_{j,k}(a) = F(x_{t+1} = x_k | x_t = x_j, a_t = a)$$

- ▶ Stopping rule: $\eta \approx 10^{-14}$.

- Value function iteration algorithm:

- 1 Guess initial value for $\bar{V}^0(x)$. Example: Static value function

$$\bar{\mathbf{V}}^0(x) = \ln(\exp(\Pi(0)) + \exp(\Pi(1))) + \gamma$$

- 2 Update value function iteration k :

$$\bar{\mathbf{V}}^k = \ln(\exp(\Pi(0) + \beta F(0)\bar{\mathbf{V}}^{k-1}) + \exp(\Pi(1) + \beta F(1)\bar{\mathbf{V}}^{k-1})) + \gamma$$

- 3 Stop if $\|\bar{\mathbf{V}}^k - \bar{\mathbf{V}}^{k-1}\| < \eta$. Otherwise, repeat steps (2)-(3).

Policy Function Representation

- Define conditional choice-probability (CCP) mapping:

$$\begin{aligned} P(x) &= \Pr \left(\begin{array}{l} \Pi(1, x) + \beta \sum_{x'} \bar{V}(x') F(x'|x, 1) + \epsilon(1) \\ \geq \Pi(0, x) + \beta \sum_{x'} \bar{V}(x') F(x'|x, 0) + \epsilon(0) \end{array} \right) \quad (1) \\ &= \exp(\tilde{v}(x) / (1 + \exp(\tilde{v}(x)))) = (1 + \exp(-\tilde{v}(x)))^{-1} \\ &\text{Where, } \tilde{v}(x) = v(1, x) - v(0, x). \end{aligned}$$

- At the “optimal” CCP, we can write the Emax function as follows:

$$\begin{aligned} \bar{V}^P(x) &= (1 - P(x)) \left[\Pi(0, x) + e(0, x) + \beta \sum_{x'} \bar{V}^P(x') F(x'|x, 0) \right] \\ &\quad + P(x) \left[\Pi(1, x) + e(1, x) + \beta \sum_{x'} \bar{V}^P(x') F(x'|x, 1) \right] \end{aligned}$$

where $e(a, x) = E(\epsilon(a) | a^* = a, x)$ is the conditional expectation $\epsilon(a)$.

Policy Function Representation (continued)

- If $\epsilon(a)$ is T1EV distributed, we can write this expectation analytically:

$$e(a, x) = \gamma - \ln P(a|x).$$

- This implicitly define the value function in terms of the CCP vector:

$$\bar{\mathbf{V}}^P = \left(I - \beta \mathbf{F}^P \right)^{-1} \begin{bmatrix} (1 - \mathbf{P}) * (\boldsymbol{\Pi}(0) + \mathbf{e}(0)) \\ + \mathbf{P} * (\boldsymbol{\Pi}(1) + \mathbf{e}(1)) \end{bmatrix} \quad (2)$$

where $\mathbf{F}^P = (1 - \mathbf{P}) * F(0) + \mathbf{P} * F(1)$ and $*$ is the element-by-element multiplication operator.

- Equations 1 and 2 define a fixed-point in \mathbf{P} :

$$\mathbf{P}^* = \Psi(\mathbf{P}^*)$$

where $\Psi(\cdot)$ is a contraction mapping.

Algorithm 2: Policy Function Iteration

- 1 Guess initial value for the CCP. Example: Static choice-probability

$$P(x) = (1 + \exp(-(\Pi(x|1) - \Pi(x|0))))^{-1}$$

- 2 Calculate expected value function:

$$\bar{\mathbf{V}}^{k-1} = \left(I - \beta \mathbf{F}^{k-1} \right)^{-1} \begin{bmatrix} (1 - \mathbf{P}^{k-1}) * (\boldsymbol{\Pi}(0) + \mathbf{e}^{k-1}(0)) \\ + \mathbf{P}^{k-1} * (\boldsymbol{\Pi}(1) + \mathbf{e}^{k-1}(1)) \end{bmatrix}$$

- 3 Update CCP:

$$P^k(x) = \Psi(P^{k-1}(x)) = \left(1 + \exp(-\tilde{v}(x)^{k-1}) \right)^{-1}$$

where $\tilde{\mathbf{v}}^{k-1} = (\boldsymbol{\Pi}(1) + \beta F(1)\bar{\mathbf{V}}^{k-1}) - (\boldsymbol{\Pi}(0) + \beta F(0)\bar{\mathbf{V}}^{k-1})$.

- 4 Stop if $\|\mathbf{P}^k - \mathbf{P}^{k-1}\| < \eta$. Otherwise, repeat steps (2)-(4)

Value-function *versus* Policy-function Algorithms

- Both algorithms are guaranteed to converge if $\beta \in (0, 1)$
- Policy-function iteration algorithms converges in fewer steps than value-function iteration.
- However, each step of the policy-function algorithm is **slower** due to the matrix inversion. M is typically very large (in the millions).
- If M is very large, it can be faster and more accurate to find $\bar{\mathbf{V}}$ using linear programming tools (e.g. linsolve in Matlab):

$$\begin{aligned}(I - \beta \mathbf{F}^{k-1}) \bar{\mathbf{V}}^{k-1} &= \begin{aligned} &(1 - \mathbf{P}^{k-1}) * (\boldsymbol{\Pi}(0) + \mathbf{e}^{k-1}(0)) \\ &+ \mathbf{P}^{k-1} * (\boldsymbol{\Pi}(1) + \mathbf{e}^{k-1}(1)) \end{aligned} \\ \Leftrightarrow \mathbf{A} \mathbf{y} &= \mathbf{b} \end{aligned}$$

- Suggested algorithm:
 - ▶ Start with value-function iteration if $\bar{\mathbf{V}}^k(x) - \bar{\mathbf{V}}^{k-1}(x) > \eta^1$
 - ▶ Switch to policy-function iteration when $\bar{\mathbf{V}}^k(x) - \bar{\mathbf{V}}^{k-1}(x) < \eta^1$
 - ▶ Where $\eta^1 < \eta$ (e.g. $\eta^1 = 10^{-2}$)

Estimation: Nested fixed-point MLE

- **Data:** Panel of choices a_{it} and observed states x_{it}
- **Parameters:** Technology parameters $\theta = \{\theta_{Y_0}, \theta_{Y_1}, \theta_{R_0}, \theta_{R_1}\}$, discount factor β , and distribution of mileage shocks $f_x(\xi_{it})$.
- **Initial step:** If the panel is long-enough, we can estimate $f_x(\xi)$ from the data. The estimated process can then be *discretized* to construct $\hat{F}(1)$ and $\hat{F}(0)$.
- Maximum likelihood problem:

$$\begin{aligned} \max_{\theta, \beta} \quad & \sum_i \sum_t a_{it} \ln P(x_{it}) + (1 - a_{it}) \ln(1 - P(x_{it})) \\ \text{s.t.} \quad & P(x_{it}) = \Psi(x_{it}) \quad \forall x_{it} \end{aligned}$$

- In practice, we need two functions:
 - ▶ *Likelihood:* Evaluate $L(\theta, \beta)$ given $P(x_{it})$.
 - ▶ *Fixed-point:* Routine that solves $P(x_{it})$ for every guess of θ, β .

Incorporating Unobserved Heterogeneity

- **Why?** Relax the conditional independence assumption.
- **Example:** Buses have heterogeneous replacement costs (K types)
 - ▶ This increases the number of parameters by $K(K-1)$: $\{\theta_{R_0}^1, \dots, \theta_{R_0}^K\} + \{\omega_1, \dots, \omega_{K-1}\}$ (probability weights).
 - ▶ E.g.: discretize a parametric distribution: $\ln \theta_{R_0}^i \sim N(\mu, \sigma^2)$
 - ▶ This changes the MLE problem:

$$\begin{aligned} \max_{\theta, \beta, \omega} \quad & \sum_i \ln \left[\sum_k g(k|x_{i1}) \prod_t P_k(x_{it})^{a_{it}} (1 - P_k(x_{it}))^{1-a_{it}} \right] \\ \text{s.t.} \quad & P_k(x_{it}) = \Psi_k(x_{it}) \quad \forall x_{it} \text{ and type } k \end{aligned}$$

Where $g(k|x_{i1})$ is the probability that bus i is type k conditional on initial mileage x_{i1} (i.e. initial condition problem).

- ▶ How to calculate $g(k|x_{i1})$?

Side note: The initial condition problem

- Unobserved heterogeneity creates a correlation between the initial state (i.e. x_{i1} mileage) and types (Heckman 1981).
- Two solutions:
 - ▶ **New buses:** Exogenous initial assignment $g(k|x_{i1}) = \omega_k$.
 - ▶ **Limiting distribution:** The bus engine replacement creates a *finite-state Markov chain* defined by

$$F_k(x'|x) = \sum_a P_k(a|x)F(x'|x, a) \text{ for each type } k$$

Under fairly general assumptions, this process generates a unique limiting distribution:

$$\pi_k(x) = \sum_{i=1}^M F_k(x_{t+1} = x | x_t = x_i) \pi_k(x_i) \leftrightarrow \pi_k = F_k^T \pi_k$$

We can use the limiting distribution to calculate the type probability conditional on initial mileage:

$$g(k|x_{i1}) = \frac{\omega_k \pi_k(x_{i1})}{\sum_{k'} \omega_{k'} \pi_{k'}(x_{i1})}$$

Identification: Residual profit

- **Assumption:** Parametric distribution function F_{ϵ} .
- Standard normalization: $\sigma_{\epsilon} = 1$.
 - ▶ This means that we cannot identify the “dollar” value of replacement costs. Only relative to variable profits.
 - ▶ True in any discrete-choice problem.
- When profits or output data are available, we can relax this normalization, and estimate σ_{ϵ} (e.g. investment and production data).

Identification: Discount Factor

- The data is summarized by the empirical hazard function:


$$h(x) = \Pr(\text{replacement}_t | \text{miles}_t = x)$$

- This corresponds to the reduced form of the model:

$$\begin{aligned} h(x) = P(x) &= F_{\tilde{\epsilon}}(\tilde{v}(x)) \\ &= F_{\tilde{\epsilon}}\left(-\beta \sum_{x'} \frac{(\Pi(1, x) - \Pi(0, x))}{V(x')(F(x'|x, 1) - F(x'|x, 0))}\right) \end{aligned}$$

- **Claim:** β is not identified, unless we parametrize payoffs: Y and RC .
 - ▶ If $\Pi(x)$ is linear in x , then non-linearity in the **observed** hazard function identifies β .
 - ▶ If $\Pi(x)$ is a non-parametric function, we cannot distinguish between a non-linear myopic model ($\beta = 0$), and a forward-looking model ($\beta > 0$).
- What would identify β ?
 - ▶ **Exclusion restriction:** The model includes a state variable z that only enters the Markov transition function (i.e. $F(x'|x, z, a)$), and not the static payoff function.

Empirical Hazard Function



Identification of β and search for the right specification

TABLE VIII
SUMMARY OF SPECIFICATION SEARCH^a

Cost Function	Bus Group		
	1, 2, 3	4	1, 2, 3, 4
Cubic $c(x, \theta_1) = \theta_{11}x + \theta_{12}x^2 + \theta_{13}x^3$	Model 1 -131.063 -131.177	Model 9 -162.885 -162.988	Model 17 -296.515 -296.411
quadratic $c(x, \theta_1) = \theta_{11}x + \theta_{12}x^2$	Model 2 -131.326 -131.534	Model 10 -163.402 -163.771	Model 18 -297.939 -299.328
linear $c(x, \theta_1) = \theta_{11}x$	Model 3 -132.389 -134.747	Model 11 -163.584 -165.458	Model 19 -300.250 -306.641
square root $c(x, \theta_1) = \theta_{11}\sqrt{x}$	Model 4 -132.104 -133.472	Model 12 -163.395 -164.143	Model 20 -299.314 -302.703
power $c(x, \theta_1) = \theta_{11}x^{\theta_{12}}$	Model 5 ^b N.C. N.C.	Model 13 ^b N.C. N.C.	Model 21 ^b N.C. N.C.
hyperbolic $c(x, \theta_1) = \theta_{11}/(91 - x)$	Model 6 -133.408 -138.894	Model 14 -165.423 -174.023	Model 22 -305.605 -325.700
mixed $c(x, \theta_1) = \theta_{11}/(91 - x) + \theta_{12}\sqrt{x}$	Model 7 -131.418 -131.612	Model 15 -163.375 -164.048	Model 23 -298.866 -301.064
nonparametric $c(x, \theta_1)$ any function	Model 8 -110.832 -110.832	Model 16 -138.556 -138.556	Model 24 -261.641 -261.641

^a First entry in each box is (partial) log likelihood value ℓ^2 in equation (5.2)) at $\beta = .9999$. Second entry is partial log likelihood value at $\beta = 0$.

^b No convergence. Optimization algorithm attempted to drive $\theta_{12} \rightarrow 0$ and $\theta_{11} \rightarrow +\infty$.

Main estimation results

TABLE IX
STRUCTURAL ESTIMATES FOR COST FUNCTION $c(x, \theta_1) = .001\theta_{11}x$
FIXED POINT DIMENSION = 90
(Standard errors in parentheses)

Parameter		Data Sample			Heterogeneity Test	
Discount Factor	Estimates/ Log-Likelihood	Groups 1, 2, 3 3864 Observations	Group 4 4292 Observations	Groups 1, 2, 3, 4 8156 Observations	LR Statistic ($df = 4$)	Marginal Significance Level
$\beta = .9999$	RC	11.7270 (2.602)	10.0750 (1.582)	9.7558 (1.227)	85.46	1.2E-17
	θ_{11}	4.8259 (1.792)	2.2930 (0.639)	2.6275 (0.618)		
	θ_{30}	.3010 (.0074)	.3919 (.0075)	.3489 (.0052)		
	θ_{31}	.6884 (.0075)	.5953 (.0075)	.6394 (.0053)		
	LL	-2708.366	-3304.155	-6055.250		
$\beta = 0$	RC	8.2985 (1.0417)	7.6358 (0.7197)	7.3055 (0.5067)	89.73	1.5E-18
	θ_{11}	109.9031 (26.163)	71.5133 (13.778)	70.2769 (10.750)		
	θ_{30}	.3010 (.0074)	.3919 (.0075)	.3488 (.0052)		
	θ_{31}	.6884 (.0075)	.5953 (.0075)	.6394 (.0053)		
	LL	-2710.746	-3306.028	-6061.641		
Myopia test:	LR Statistic ($df = 1$)	4.760	3.746	12.782		
$\beta = 0$ vs. $\beta = .9999$	Marginal Significance Level	0.0292	0.0529	0.0035		

Patents as options, Pakes (1986)

- This paper studies the value of patent protection: (i) what is the stochastic process determining the value of innovations?, (ii) how patent protection laws affect the decision to renew patents and the distribution of returns to innovation?
- The model is an example of an **optimal stopping problem**. The model is setup with a finite horizon, but it does not have to be. Other examples: retirement, firm exit decisions, technology adoption, etc.
- **Contributions:**
 - ▶ Illustrate how we can infer the implicit option value of patents (or any other dynamic investment decision) from dynamic discrete choices (i.e. principle of revealed preference).
 - ▶ This is done without actually observing profits or revenues from patents. Only the dynamic structure of renewal costs are needed.
 - ▶ More technically, the paper is one of the firsts applications of simulation methods in econometrics (very influential).

Data and Institutional Details

- Three countries: France, Germany and UK
- Renewal date for all patents: $n_{m,t}(a)$ = number of surviving patents at age a in country m from cohort t .
- Regulatory environment by country/cohort:
 - ▶ f : Number of automatic renewal years.
 - ▶ L : Expiration date on patent
 - ▶ $\mathbf{c} = \{c_1, \dots, c_T\}$: Deterministic renewal cost

TABLE I
CHARACTERISTICS OF THE DATA^a

Country Characteristic	France	U.K.	Germany
1. f	2	5	3
2. L	20	16	18
3. Application dates of cohorts	1951-79	1950-74	1952-72
4. First/last year in which renewals are observed	1970/81	1955/78	1955/74
5. Patents studied from cohort: all patents	Applied for	Applied for	Granted
6. Estimated average ratio of patents granted to patents applied for ^b	.93	.83	.35
7. $\overline{NPAT} = N/J$	36,865	37,286	21,273

Country differences in drop-out probabilities




FIGURE 2.—Average drop out proportions.^a

Country differences in renewal fee schedules




FIGURE 3.—Average of renewal fee schedules.

Model setup

- Consider the renewal problem for patent i
- Stochastic sequence of returns from patent: $r_i = \{r_{i1}, \dots, r_{iL}\}$
- Evolution of returns depend on:
 - ① initial quality level
 - ② arrival of substitutes innovations that depreciate the value of the patent
 - ③ arrival of complement innovations that increase its value.
- Model structural parameters (per country):
 - ▶ δ measures the normal obsolescence rate
 - ▶ ϕ and σ determines the arrival rate and magnitude of complementary innovations
 - ▶ λ determines to arrival rate of substitute innovations
 - ▶ μ_0 and σ_0 determines the initial quality pool of innovations
- Discount factor β is fixed.

Stochastic Process

- Markov process for returns:

$$r_{it+1} = \tau_{it+1} \max\{\delta r_{it}, \xi_{it+1}\}$$

Where, $\Pr(\tau_{it+1} = 0 | r_{it}, t) = \exp(-\lambda r_{it})$

$$p(\xi_{it+1} | r_{it}, t) = \frac{1}{\phi^t \sigma} \exp\left(-\frac{\gamma + \xi_{it+1}}{\phi^t \sigma}\right)$$

$$r_{i0} \sim \text{LN}(\mu_0, \sigma_0^2)$$

or more compactly for $t > 0$,

$$f(r_{it+1} | r_{it}, t) = \begin{cases} \exp(-\lambda r_{it}) & \text{If } r_{it+1} = 0 \\ \Pr(\xi_{it+1} < \delta r_{it} | r_{it}, t) & \text{If } r_{it+1} = \delta r_{it} \\ \frac{1}{\phi^t \sigma} \exp\left(-\frac{\gamma + \xi_{it+1}}{\phi^t \sigma}\right) & \text{If } r_{it+1} > \delta r_{it} \end{cases}$$

Optimal stopping problem

- In the last year, the renewal value depends only on c_L and r_{iL} :

$$V(L, r_{iL}) = \max\{0, r_{iL} - c_L\}$$

and therefore the patent is renewed if $r_{iL} > r_L^* = c_L$.

- At year $L - 1$, the value is defined recursively:

$$V(L, r_{iL-1}) = \max \left\{ 0, r_{iL-1} - c_{L-1} + \beta \int_{r_L^*}^{\infty} V(L, r_{iL}) f(r_{iL} | r_{iL-1}, L-1) dr_{iL} \right\}$$

- This value function is strictly increasing in r_{iL-1} (see proposition 1). Therefore, there exists a unique threshold such that the patent is renewed if

$$r_{iL-1} > r_{L-1}^* = c_{L-1} - \beta \int_{r_L^*}^{\infty} V(L, r_{iL}) f(r_{iL} | r_{L-1}^*, L-1) dr_{iL}$$

Optimal stopping problem (continued)

- Similarly, for any year $t > 0$ the value function is defined recursively as follows:

$$V(L, r_{it}) = \max\{0, r_{it} - c_t + \beta \int_{r_{t+1}^*}^{\infty} V(t+1, r_{it+1}) f(r_{it+1}|r_{it}, t) dr_{it+1}\}$$

which lead to a series of optimal stopping rules:

$$r_{it} > r_t^* = c_t - \beta \int_{r_{t+1}^*}^{\infty} V(t+1, r_{it+1}) f(r_{it+1}|r_t^*, t)$$

- Given the function form assumptions on $f(r'|r_t, t)$, the thresholds can be solved analytically by backward induction.
- **Note:** When the terminal period is stochastic the value function becomes stationary (i.e. infinite horizon). For instance, optimal stopping problems arise when studying retirement or exit decisions:

$$V(s_t) = \max\left\{0, \pi(s_t) + \beta \int (1 - \delta(s_t)) V(s_{t+1}) f(s_{t+1}|s_t) ds_{t+1}\right\}$$

Estimation Method

- Likelihood of the observed renewal sequence \mathbf{N}_m conditional on the regulation environment $\mathbf{Z}_m = \{L_m, f_m, \mathbf{c}_m\}$ in country m :

$$L(\mathbf{N}_m | \mathbf{Z}_m, \theta) = \max_{\theta} \sum_{t=1}^L n_m(t) \ln \Pr(t^* = t | \mathbf{Z}_m, \theta)$$

Where,

$$\Pr(t^* = t | \theta, \mathbf{Z}_m) = \int_{-\infty}^{\infty} \int_{r_1^*}^{\infty} \int_{r_2^*}^{\infty} \dots \int_0^{r_t^*} dF(r_{i1}, \dots, r_{it-1}, r_{it}) dF_0(r_{i0})$$

- Monte Carlo integration approximation:

0. Sample $r_{i0}^s \sim \text{LN}(\mu_0, \sigma_0^2)$

1. Period 1:

- ① Sample τ_1^s from Bernoulli with probability $\exp(-\lambda r_0^s)$
- ② If $\tau_1^s = 1$, sample ξ_1^s from exponential distribution. Otherwise, do not renew patent: $a_1^s = 0$.
- ③ Calculate r_1^1
- ④ Evaluate decision: $a_1^s = 1$ if $r_1^s > r_1^*$.

...

t. Repeat sampling for period t if patent was renewed at $t - 1$

Estimation Method (continued)

- After collecting the simulated sequences of actions, we can evaluate the simulated choice-probability at period t :

$$\tilde{P}_S(t|\theta, \mathbf{Z}_m) = \frac{1}{S} \sum_s 1(a_1^s = 1, a_2^s = 1, \dots, a_{t-1}^s = 1, a_t^s = 0)$$

- Numerical problem: $\tilde{P}_S(t|\theta, \mathbf{Z}_m)$ is not a smooth function of the parameters θ + equal to zero for some t unless $S \rightarrow \infty$.
- Smooth alternative approximation:

$$\hat{P}_S(t, \theta, \mathbf{Z}_m) = \frac{\exp(\tilde{P}_S(t|\theta, \mathbf{Z}_m)/\eta)}{1 + \sum_{t'} \exp(\tilde{P}_S(t'|\theta, \mathbf{Z}_m)/\eta)}$$

- **Note:** All the structural parameters are identified in this model (except β). The implicit normalization is that coefficient on renewal cost c_t is one: all the parameters are expressed in dollar.

TABLE II
PARAMETER ESTIMATES^a

	Country		
	France	U.K. ^b	Germany
A. Parameter			
σ	5689 (8.24)	5467 (6.09)	7460 (19.72)
γ	9162 (13.67)	6919 (10.29)	8687 (17.09)
ϕ	.5084 (5.66×10^{-4})	.4383 (2.17×10^{-3})	.4896 (1.16×10^{-3})
δ	.8475 (2.62×10^{-4})	.8102 (1.81×10^{-3})	.8861 (2.48×10^{-4})
σ_R	1.579 (2.92×10^{-3})	1.525 (3.04×10^{-3})	1.158 (2.36×10^{-3})
μ	4.705 (2.75×10^{-3})	5.425 (2.55×10^{-3})	6.718 (3.70×10^{-3})
θ	.0990 (6.36×10^{-4})	.36 ^b	.0855 (2.46×10^{-3})
B. Dimension^c			
B.1. NPAT	1,069,095	983,471	446,741
B.2. NSIM	20,000	20,000	20,000
B.3. Age: f/L	2/20	5/16	3/18
B.4. NCHRT	29	26	21
B.5. NCHRTAGE	238	272	237
C. Summary Statistic^d			
C.1. MSE[$\tilde{\pi}$]	5.42×10^{-4}	6.91×10^{-4}	1.48×10^{-4}
C.2. PDW[$\tilde{\pi}$]	1.65	2.24	1.85
C.3. V[$\tilde{\pi}$; data]	3.90×10^{-2}	1.07×10^{-2}	2.65×10^{-2}

^a Patents are assigned to cohorts by year of application. Numbers in parenthesis beside parameter estimates are their estimated standard errors.

TABLE III
THE EVOLUTION OF IMPLICIT REVENUES IN THE EARLY AGES^a

	Country	
	France	Germany
<i>Characteristic</i>		
$E_{(r_1)}[r_1 r_1 > 0]$	380.43	1608.57
Pr (Downside); Pr (Upside)	.0637; .1807	.0004; .2705
$\pi(2)$.0637	(no required renewal)
$E_{(r_2)}[r_2 r_2 > 0]$	1414.72	3400.98
Pr (Downside); Pr (Upside)	.0387; .0331	.0006; .0584
$\pi(3)$.0907	.0013
$E_{(r_3)}[r_3 r_3 > 0]$	1432.24	3224.56
Pr (Downside); Pr (Upside)	.0118; .0012	.0005; .0039
$\pi(4)$.0792	.0121
$E_{(r_4)}[r_4 r_4 > 0]$	1339.05	2899.41
Pr (Downside); Pr (Upside)	.0048; 0.00	.0003; 0.0
$\pi(5)$.0381	.0277
$E_{(r_5)}[r_5 r_5 > 0]$	1192.70	2641.40
<i>NPAT</i>	36,865	21,273

Summary of the Results

- Main differences across countries: (i) patent regulation rules, (ii) initial distribution of patent returns.
- Germany has a more selective screening system for granting new patents: higher mean and smaller variance of initial returns r_{i0} .
- Learning about complementary innovations: $\phi \approx 0.5$. Imply very fast learning/growth in returns.
- This has important policy implications: Regulator wants to keep initial renewing cost low, and increase them fast to extract rents from high value patents (low distortions after learning is over).

The distribution of realized patent value is highly skewed

TABLE V
PERCENTILES (p1) AND LORENZ CURVE COEFFICIENTS (lc) FROM THE DISTRIBUTION OF
REALIZED PATENT VALUES^a

Per cent p	Country					
	France		U.K.		Germany	
	p1 (\$)	lc per cent	p1 (\$)	lc per cent	p1 (\$)	lc per cent
.25	75.23	.544	355.55	.554	1,999.60	2.249
.50	533.96	1.833	1,516.84	3.247	6,252.93	7.341
.75	3,731.35	8.087	7,947.55	16.369	19,576.26	25.288
.85	10,292.06	19.575	15,357.09	31.721	32,428.14	41.001
.90	17,423.11	31.261	22,206.21	44.257	44,241.87	52.672
.95	31,609.59	52.461	34,740.07	62.960	65,753.61	69.223
.97	42,905.78	65.514	43,889.95	73.640	78,299.01	78.348
.98	51,215.84	73.729	51,277.22	80.072	94,842.63	83.800
.99	66,515.40	84.011	65,075.08	87.858	118,354.78	90.330
maximum	259,829.27	—	374,028.70	—	419,217.55	—
mean	5,631.03	—	7,357.05	—	16,169.48	—
<i>NPAT</i>		36,865		37,826		21,273

^a The realized value for patent i is $\sum_{\tau=1}^{\tau_i^*} \beta^{(\tau-1)}(r_{i,\tau} - c_p)$, where τ_i^* is the last age at which patent i was kept in force. See also the note to Table III.

- Implied rate of returns on R&R: France = 15.56%, UK = 111.03%, Germany = 13.83%.

Sequential estimators of DDC models

- **Key references:**

- ▶ Hotz and Miller (1993)
- ▶ Hotz, Miller, Sanders, and Smith (1994)
- ▶ Aguirregabiria and Mira (2002)
- ▶ Identification: Magnac and Thesmar (2002), Kasahara and Shimotsu (2009)

- Consider the following dynamic discrete choice model with additively separable (AS) and conditional independent (CI) errors.

- ▶ A discrete actions.
- ▶ Payoff function: $u(x|a)$
- ▶ State space: (x, ϵ) .
- ▶ Where x is a discrete state vector, and ϵ is an A -dimensions continuous vector.
- ▶ Distribution functions:
 - ★ $\Pr(x_{t+1} = x' | x_t, a) = f(x' | x, a)$
 - ★ $g(\epsilon)$ is a type-1 EV density with unit variance.

Bellman Operator

- Bellman equation:

$$\begin{aligned} V(x) &= \int \max_{a \in A} \{ u(x|a) + \epsilon(a) + \beta \sum_{x'} V(x') f(x'|x, a) \} g(\epsilon) d\epsilon \\ &= \int \max_{a \in A} \{ v(x|a) + \epsilon(a) \} g(\epsilon) d\epsilon \\ &= \ln \left(\sum_a \exp(v(x|a)) \right) + \gamma \\ &= \Gamma(V(x)) \end{aligned}$$

CCP Operator

- Express $V(x)$ as a function of $P(a|x)$.

$$V(x) = \sum_a P(a|x) * \left\{ u(x|a) + E(\epsilon(a)|x, a) + \beta \sum_{x'} V(x') f(x'|x, a) \right\}$$

Where,

$$\begin{aligned} E(\epsilon(a)|x, a) &= \frac{1}{P(a|x)} \int 1\left(v(x|a) + \epsilon(a) > v(x|a') + \epsilon(a'), a' \neq a\right) g(\epsilon) d\epsilon \\ e(a, P(a|x)) &= \gamma - \ln P(a|x) \end{aligned}$$

CCP Operator (continued)

- In Matrix form:

$$V = \sum_a P(a) * [u(a) + e(a, P) + \beta F(a)V]$$

$$[I - \beta \sum_a P(a) * F(a)] V = \sum_a P(a) * [u(a) + e(a, P)]$$

$$V(P) = [I - \beta \sum_a P(a) * F(a)]^{-1} \left[\sum_a P(a) * (u(a) + e(a, P)) \right]$$

where $F(a)$ is $|X| \times |X|$ and V is $|X| \times 1$.

CCP Operator (continued)

- In Matrix form:

$$V = \sum_a P(a) * [u(a) + e(a, P) + \beta F(a)V]$$

$$[I - \beta \sum_a P(a) * F(a)] V = \sum_a P(a) * [u(a) + e(a, P)]$$

$$V(P) = [I - \beta \sum_a P(a) * F(a)]^{-1} \left[\sum_a P(a) * (u(a) + e(a, P)) \right]$$

where $F(a)$ is $|X| \times |X|$ and V is $|X| \times 1$.

- The CCP contraction mapping is:

$$\begin{aligned} P(a|x) &= \Pr \left(v(x|a, P) + \epsilon(a) > v(x|a', P) + \epsilon(a'), a' \neq a \right) \\ &= \frac{\exp(\tilde{v}(x|a, P))}{1 + \sum_{a' \neq a} \exp(\tilde{v}(x|a', P))} \\ &= \Psi(a|x, P) \end{aligned}$$

where $\tilde{v}(x|a, P) = v(x|a, P) - v(x|1, P)$.

Two Special Cases

- ① **Linear payoff:** If $u(x|a, \theta) = x\theta$, the value function is also linear in θ .

$$V(P) = Z(P)\theta + \lambda(P)$$

Where
$$Z(P) = [I - \beta \sum_a P(a) * F(a)]^{-1} \left[\sum_a P(a) * X \right]$$

$$\lambda(P) = [I - \beta \sum_a P(a) * F(a)]^{-1} \left[\sum_a P(a) * e(a, P) \right]$$

Two Special Cases

- ① **Linear payoff:** If $u(x|a, \theta) = x\theta$, the value function is also linear in θ .

$$V(P) = Z(P)\theta + \lambda(P)$$

$$\text{Where} \quad Z(P) = [I - \beta \sum_a P(a) * F(a)]^{-1} \left[\sum_a P(a) * X \right]$$

$$\lambda(P) = [I - \beta \sum_a P(a) * F(a)]^{-1} \left[\sum_a P(a) * e(a, P) \right]$$

- ② **Absorbing state:** $v(x|0) = 0$ (e.g. Exit or retirement). This change the value function:

$$V(x, \varepsilon) = \max \left\{ u(x) + \varepsilon(1) + \beta \sum_{x'} \underbrace{E_{\varepsilon'}[V(x', \varepsilon')]}_{=\bar{V}(x')} F(x'|x), \varepsilon(0) \right\}$$

As before, the expected continuation value is:

$$\begin{aligned} \bar{V}(x) &= \log \left(\exp(0) + \exp \left(u(x) + \beta \sum_{x'} \bar{V}(x') F(x'|x) \right) \right) + \gamma \\ &= \log(1 + \exp(v(x))) + \gamma \end{aligned}$$

Two Special Cases (continued)

- The choice probability is given by:

$$\Pr(a = 1|x) = P(x) = \frac{\exp(v(x))}{1 + \exp(v(x))}$$

Note that the log of the “odds-ratio” is equal to the choice-specific value function:

$$\log \left(\frac{P(x)}{1 - P(x)} \right) = v(x)$$

- Therefore, the expected continuation value can be expressed as a function of $P(x)$:

$$\begin{aligned}\bar{V}^P(x) &= \log(1 + \exp(v(x))) + \gamma = \log \left(1 + \frac{P(x)}{1 - P(x)} \right) + \gamma \\ &= -\log(1 - P(x)) + \gamma\end{aligned}$$

- **Implication:** With an absorbing state, we don't need to invert $[I - \beta \sum_a P(a) * F(a)]$ to apply the CCP mapping.

Two-Step Estimator

- The objective is to estimate the structural parameters θ without repeatedly solving the DP problem
- **Initial step:** Reduced form of the model
 - ▶ Markov transition process: $\hat{f}(x'|x, a)$
 - ▶ Policy function: $\hat{P}(a|x)$
 - ▶ **Constraint:** Need to estimate both functions at EVERY state point x .

Two-Step Estimator

- The objective is to estimate the structural parameters θ without repeatedly solving the DP problem
- **Initial step:** Reduced form of the model
 - ▶ Markov transition process: $\hat{f}(x'|x, a)$
 - ▶ Policy function: $\hat{P}(a|x)$
 - ▶ **Constraint:** Need to estimate both functions at EVERY state point x .
- How? Ideally $\hat{P}(a|x)$ is estimated non-parametrically to avoid imposing a particular functional form on the policy function (i.e. no theory involved at this stage). This would correspond to a frequency estimator:

$$\hat{P}(a|x) = \frac{1}{n(x)} \sum_{i \in n(x)} 1(a_i = a)$$

- For finite samples, we need to impose smooth the policy function and interpolate between states are not visited (or infrequently). Kernels or local-polynomial techniques can be used.

Two-Step Estimator

- The objective is to estimate the structural parameters θ without repeatedly solving the DP problem
- **Initial step:** Reduced form of the model
 - ▶ Markov transition process: $\hat{f}(x'|x, a)$
 - ▶ Policy function: $\hat{P}(a|x)$
 - ▶ **Constraint:** Need to estimate both functions at EVERY state point x .
- How? Ideally $\hat{P}(a|x)$ is estimated non-parametrically to avoid imposing a particular functional form on the policy function (i.e. no theory involved at this stage). This would correspond to a frequency estimator:

$$\hat{P}(a|x) = \frac{1}{n(x)} \sum_{i \in n(x)} 1(a_i = a)$$

- For finite samples, we need to impose smooth the policy function and interpolate between states are not visited (or infrequently). Kernels or local-polynomial techniques can be used.
- **Second-step:** Structural parameters conditional on (\hat{P}, \hat{f})

Example: Linear payoff function, $u(x|a, \theta) = x(a)\theta$

- **1- Data Preparation:** Use (\hat{P}, \hat{F}) to calculate:

$$Z(\hat{P}, \hat{F}) = [I - \beta \sum_a \hat{P}(a) * \hat{F}(a)]^{-1} \left[\sum_a \hat{P}(a) * X(a) \right]$$

$$\lambda(\hat{P}, \hat{F}) = [I - \beta \sum_a \hat{P}(a) * F(a)]^{-1} \left[\sum_a \hat{P}(a) * e(a, \hat{P}) \right]$$

- **2- GMM:** Let W_{it} denote a vector of predetermined instruments (e.g. state-variables and their interactions). We can construct moment conditions:

$$E \left(W_{it} \left[a_{it} - \Psi(a_{it}|x_{it}, \hat{P}, \hat{F}) \right] \right) = 0$$

Where,

$$\Psi(a_{it}|x_{it}, \hat{P}, \hat{F}) = \frac{\exp \left(v(x_{it}|a_{it}, \hat{P}, \hat{F}) \right)}{\sum_{a'} \exp \left(v(x_{it}|a', \hat{P}, \hat{F}) \right)}$$

$$v(x|a, \hat{P}, \hat{F}) = x(a)\theta + \beta \sum_{x'} \underbrace{V(x|\hat{P}, \hat{F})}_{=Z(x, \hat{P}, \hat{F})\theta + \lambda(x, \hat{P}, \hat{F})} \hat{f}(x'|x, a).$$

$$v(x|a, \hat{P}, \hat{F}) = \left(x(a) + \beta \bar{Z}(x|\hat{P}, \hat{F}) \right) \theta + \beta \bar{\lambda}(x|\hat{P}, \hat{F})$$

- Therefore, the second-stage of problem is equivalent to a linear GMM (note: This also highlights the difficulty of identifying β separately from θ)

Pseudo-likelihood estimators (PML)

- **Source:** Aguirregabiria and Mira (2002)
- **Data:** Panel of n individuals of T periods:

$$(A, X) = \{a_{it}, x_{it}\}_{i=1, \dots, n; t=1, \dots, T}$$

- **2-Step estimator:**

- 1 Obtain a flexible estimator of CCPs $\hat{P}^1(a|x)$
- 2 Feasible PML estimator:

$$Q^{2S}(A, X) = \max_{\theta} \sum_t \sum_i \psi(a_{it}|x_{it}, \hat{P}^1, \hat{F}, \theta)$$

If $V(P)$ is linear, the second step is a linear probit/logit model.

Pseudo-likelihood estimators (PML)

- **NPL estimator:** The NPL repeat the PML and policy function iteration steps sequentially (i.e. swapping the fixed-point algorithm).
 - ① Obtain a flexible estimator of CCPs $\hat{P}^1(a|x)$
 - ② Feasible PML step:

$$Q^{k+1}(A, X) = \max_{\theta} \sum_t \sum_i \Psi(a_{it}|x_{it}, \hat{P}^k, \hat{F}, \theta)$$

- ③ Policy function iteration step:

$$\hat{P}^{k+1}(a|x) = \Psi(a|x, \hat{P}^k, \hat{F}, \hat{\theta}^{k+1})$$

- ④ Stop if $\|\hat{P}^{k+1} - \hat{P}^k\| < \eta$, else repeat step 2 and 3.
- In the single agent case: The NPL is guaranteed to converge to the MLE estimator (i.e. NFXP).
 - In practice, Aguirregabiria and Mira (2002) showed that 2 or 3 steps is sufficient to eliminate the small sample bias of the 2-step estimator, **and** is computationally easier to implement than the NFXP.

Simulation-based CCP estimator

Source: Hotz, Miller, Sanders, and Smith (1994)

- **Starting point:** The H&M GMM estimator suffers from a curse of dimensionality in $|X|$, since we must invert a $|X| \times |X|$ matrix to evaluate the continuation value (not true for optimal-stopping models). This is less severe for NFXP estimators, since we can use the value-function mapping to solve the policy functions.

- **Solution:**

- ▶ First insight: We only need to know the relative choice-specific value function $\tilde{v}(a|x) = v(a|x) - v(1|x)$ to predict behavior.

$$a_{it} = \begin{cases} 1 & \text{If } \tilde{v}(a|x) + \tilde{\epsilon}(a) < 0 \text{ for all } a \neq 1 \\ a & \text{If } \max\{0, \tilde{v}(a'|x) + \tilde{\epsilon}(a')\} < \tilde{v}(a|x) + \tilde{\epsilon}(a) \text{ for all } a' \neq a \end{cases}$$

- ▶ Second insight: There exists a one-to-one mapping between $\tilde{v}(a|x)$ and $P(a|x)$.

Simulation-based CCP estimator

- Logit example:

$$P(a|x) = \frac{\exp(v(a|x))}{\sum_{a'} \exp(v(a'|x))} = \frac{\exp(\tilde{v}(a|x))}{1 + \sum_{a' > 1} \exp(\tilde{v}(a'|x))}$$
$$\Leftrightarrow \tilde{v}(a|x, P) = \ln P(a|x) - \ln P(1|x)$$

- Third insight: We can approximate the model's predicted value function at any state x by simulating actions according to a policy function $P(a|x)$.

$$\hat{V}^S(x|P) = \frac{1}{S} \sum_s \sum_{\tau=0}^T \beta^\tau \{ u(x_{t+\tau}^s, a_{t+\tau}^s) + e(a_{t+\tau}^s | P(a_{t+\tau}^s | x_{t+\tau}^s)) \}$$

where (x^s, a^s) is a simulated sequence of choices and states sampled from $P(a|x)$ and $f(x'|x, a)$, and $e(a|P(a|x)) = E(\epsilon(a)|a_i = a, x, P)$ [closed-form expression]. Importantly, $\lim_{S \rightarrow \infty} \hat{V}^S(x|P) = V(s|P)$.

Estimation Procedure

- **Step 1:** Estimate $\hat{P}(a|x)$ and $\hat{f}(x'|x, a)$, and compute the “dependent variable”:

$$\tilde{v}_n(a_{it}|x_{it}, \hat{P}) = \ln \hat{P}(a_{it}|x_{it}) - \ln \hat{P}(1|x_{it})$$

- **Step 2a:** Simulation of value functions of each observed state and choice (x_{it}, a_{it}) . Each simulated sequence calculate the value of “future” choices:

- ① Calculate static value of (x_{it}, a_{it}) : $u(x_{it}, a_{it}|\theta) + e(a_{it}|\hat{P}, x_{it})$
- ② Sample new state for period $t + 1$: $x_{it+1} \sim \hat{f}(x'|x_{it}, a_{it})$
- ③ Sample new choice for period $t + 1$: $a_{it+1} \sim \hat{P}(a|x_{it})$

Repeat steps 1-3 for T periods. This gives us the net present value of one simulated sequence:

$$\begin{aligned} v^s(a_{it}|x_{it}, \hat{P}, \theta) &= u(x_{it}, a_{it}|\theta) + e(a_{it}|\hat{P}, x_{it+\tau}) \\ &+ \sum_{\tau=1}^T \beta^\tau \left[u(x_{it+\tau}^s, a_{it+\tau}^s|\theta) + e(a_{it+\tau}^s|\hat{P}, x_{it+\tau}^s) \right] \end{aligned}$$

Estimation Procedure (continued)

- Repeat this process S times.
- This gives us the simulated value of choosing a_{it} in state x_{it} :

$$v^S(a_{it}|x_{it}, \hat{P}, \theta) = \frac{1}{S} \sum_s v^s(a_{it}|x_{it}, \hat{P})$$

Let $\tilde{v}^S(a_{it}|x_{it}, \hat{P}, \theta) = v^S(a_{it}|x_{it}, \hat{P}, \theta) - v^S(1|x_{it}, \hat{P}, \theta)$.

- **Note:** If $u(x, a|\theta)$ is linear in θ , we need to do this simulation process only once.
- **Step 2b:** Moment conditions

$$E \left(W_{it} \left[\tilde{v}_n(a_{it}|x_{it}, \hat{P}) - \tilde{v}^S(a_{it}|x_{it}, \hat{P}, \theta) \right] \right) = 0$$

where W_{it} is a vector of instruments.

Estimation Procedure (continued)

- Importantly, setting up the moment conditions this way implies that the estimate will be consistent even with a finite number of simulated number of draws S .
- Why? The simulation error, $\tilde{v}(a_{it}|x_{it}, \hat{P}, \theta) - \tilde{v}^S(a_{it}|x_{it}, \hat{P}, \theta)$, is additive, and therefore vanishes as $n \rightarrow \infty$ (instead of $S \rightarrow \infty$).
- However, the small sample bias in \hat{P} enters non-linearly in the moment conditions, and can induce severe biases (same as before):

$$\ln \left(\hat{P}(a_{it}|x_{it}) + u_{it}(a) \right) - \ln \left(\hat{P}(1|x_{it}) + u_{it}(1) \right) \neq \ln \hat{P}(a_{it}|x_{it}) - \ln \hat{P}(1|x_{it}) + u_{it}$$

For instance, if $\hat{P}(a_{it}|x_{it}) = 0$, the objective function is not defined.

- HMSS presents Monte-Carlo experiment to illustrate the small-sample bias. It can be quite large.

Akerberg, D. (2003).

Advertising, learning, and consumer choice in experience good markets: A structural empirical examination.
International Economic Review 44, 1007–1040.

Adda, J. and R. Cooper (2000).

Balladurette and juppette: A discrete analysis of scrapping subsidies.
Journal of Political Economy 108(4).

Aguirregabiria, V. and P. Mira (2002).

Swapping the nested fixed point algorithm: A class of estimators for discrete markov decision models.
Econometrica 70(4), 1519–1543.

Crawford, G. and M. Shum (2005).

Uncertainty and learning in pharmaceutical demand.
Econometrica 73, 1137–1174.

Das, S. (1992).

A microeconomic model of capital utilization and retirement: The case of the us cement industry.
Review of Economic Studies 59, 277–297.

Das, S., M. Roberts, and J. Tybout (2007, May).

Market entry costs, producer heterogeneity, and export dynamics.
Econometrica.

Erdem, T., S. Imai, and M. P. Keane (2003).

Brand and quantity choice dynamics under price uncertainty.
Quantitative Marketing and Economics 1, 5–64.

Erdem, T. and M. P. Keane (1996).

Decision-making under uncertainty: Capturing dynamic brand choice processes in turbulent consumer goods markets.
Marketing Science 15(1), 1–20.

Gordon, B. (2010, September).

A dynamic model of consumer replacement cycles in the pc processor industry.
Marketing Science 28(5).

Gowrisankaran, G. and M. Rysman (2012).

Dynamics of consumer demand for new durable goods.

Journal of Political Economy 120, 1173–1219.

Heckman, J. J. (1981).

The incidental parameters problem and the problem of initial conditions in estimating a discrete time-discrete data stochastic process.

In C. F. Manski and D. McFadden (Eds.), *Structural Analysis of Discrete Data with Econometric Applications*, pp. 179–195. MIT Press.

Hendel, I. and A. Nevo (2006).

Measuring the implications of sales and consumer stockpiling behavior.

Econometrica 74(6), 1637–1673.

Hotz, V. J. and R. A. Miller (1993).

Conditional choice probabilities and the estimation of dynamic models.

The Review of Economic Studies 60(3), 497–529.

Hotz, V. J., R. A. Miller, S. Sanders, and J. Smith (1994).

A simulation estimator for dynamic models of discrete choice.

The Review of Economic Studies 61(2), 265–289.

Kasahara, H. and K. Shimotsu (2009).

Nonparametric identification of finite mixture models of dynamic discrete choices.

Econometrica 77(1).

Lee, R. (2013, December).

Vertical integration and exclusivity in platform and two-sided markets.

American Economic Review 103(7), 2960–3000.

Magnac, T. and D. Thesmar (2002).

Identifying dynamic discrete decision processes.

Econometrica 70, 801–816.

Pakes, A. (1986).

Patents as options: Some estimates of the value of holding european patent stocks.

Econometrica: Journal of the Econometric Society 54(4), 755–784.

Rust, J. (1987).

Optimal replacement of gmc bus engines: An empirical model of harold zurcher.

Econometrica: Journal of the Econometric Society 55(5), 999–1033.

Rust, J. and G. Rothwell (1995).

Optimal response to a shift in regulatory regime: The case of the us nuclear power industry.

Journal of Applied Econometrics 10(Special Issue: The Microeconometrics of Dynamic Decision Making), S75–S118.