Estimating Single-Agent Dynamic Models

Paul T. Scott New York University

PhD Empirical IO Fall 2017

Why dynamic estimation? External validity

 Famous example: Hendel and Nevo's (2006) estimation of laundry detergent demand

Why dynamic estimation? External validity

- Famous example: Hendel and Nevo's (2006) estimation of laundry detergent demand
- The long-run demand elasticity for laundry detergent might be zero (or very close)

Why dynamic estimation? External validity

- Famous example: Hendel and Nevo's (2006) estimation of laundry detergent demand
- The long-run demand elasticity for laundry detergent might be zero (or very close)
- If detergent goes on sale periodically, we might see a nonzero short-run elasticity (perhaps even a large one) as customers might purchase during the sales and store the detergent.

Why dynamic estimation? External validity

- Famous example: Hendel and Nevo's (2006) estimation of laundry detergent demand
- The long-run demand elasticity for laundry detergent might be zero (or very close)
- If detergent goes on sale periodically, we might see a nonzero short-run elasticity (perhaps even a large one) as customers might purchase during the sales and store the detergent.
- Dynamic estimation typically involves estimating the primitives of decision makers' objective functions. We might estimate the model using short-run variation, but once we know the decision maker's objective function, we could simulate a response to long-run variation.

Why are dynamics difficult?

- The computational burden of solving dynamic problems blows up as the state space gets large. With standard dynamic estimation techniques, this is especially problematic, for estimation may involve solving the dynamic problem many times.
- Serially correlated unobservables and unobserved heterogeneity (easy to confuse with state dependence)
- Modeling expectations
- Solving for equilibria, multiplicity (dynamic games)

Outline

- Introduction to dynamic estimation: Rust (1987)
- Conditional choice probabilities: Hotz and Miller (1993)
- Euler equation estimation: Scott (2014)

Rust (1987) and NFP estimation

"Optimal Replacement of GMC Bus Engines: An Empirical Model of Harold Zurcher" John Rust (1987) The "application"

- The decision maker decides whether replace bus engines or not, minimizing expected discounted cost
- The trade-off: engine replacement is costly, but with increased use, the probability of a very costly breakdown increases
- Single agent setting: prices are exogenous, no externalities across buses

Model, part I

- state variable: x_t is the bus engine's mileage
 - For computational reasons, Rust discretizes the state space into 90 intervals.
- Action $i_t \in \{0,1\}$, where
 - $i_t = 1$ replace the engine,
 - $i_t = 0$ keep the engine and perform normal maintenance.

Model, part II

per-period profit function:

$$\pi(i_t, x_t, \theta_1) = \begin{cases} -c(x_t, \theta_1) + \varepsilon_t(0) & \text{if } i_t = 0\\ -(RC - c(0, \theta_1)) + \varepsilon_t(1) & \text{if } i_t = 1 \end{cases}$$

where

- c (x_t, θ₁) regular maintenance costs (including expected breakdown costs),
- ► *RC* the net costs of replacing an engine,
- ε payoff shocks.
- x_t is observable to both agent and econometrician, but ε is only observable to the agent.
- ε is necessary for a coherent model, for sometimes we observe the agent making different decisions for the same value of x.

Model, part III

Can define value function using Bellman equation:

$$V_{\theta}(x_{t},\varepsilon_{t}) = \max_{i} \left[\pi\left(i, x_{t}, \theta\right) + \beta EV_{\theta}\left(x_{t}, \varepsilon_{t}, i_{t}\right)\right]$$

where

$$EV_{\theta}\left(x_{t},\varepsilon_{t},i_{t}\right)=\int V_{\theta}\left(y,\eta\right)p\left(dy,d\eta|x_{t},\varepsilon_{t},i_{t},\theta_{2},\theta_{3}\right)$$

Parameters

- θ_1 parameters of cost function
- ▶ θ_2 parameters of distribution of ε (these will be assumed/normalized away)
- θ_3 parameters of *x*-state transition function
- ► *RC* replacement cost
- discount factor β will be imputed (more on this later)

Conditional Independence

Conditional Independence Assumption

The transition density of the controlled process $\{x_t, \varepsilon_t\}$ factors as:

 $p(x_{t+1},\varepsilon_{t+1}|x_t,\varepsilon_t,i_t,\theta_2,\theta_3) = q(\varepsilon_{t+1}|x_{t+1},\theta_2) p(x_{t+1}|x_t,i_t,\theta_3)$

- Cl assumption is very powerful: it means we don't have to treat \varepsilon_t as a state variable, which would be very difficult since it's unobserved.
- ▶ While it is possible to allow the distribution of *ε*_{t+1} to depend on *x*_{t+1}, authors (including Rust) typically assume that any conditionally independent error terms are also identically distributed over time.

Theorem 1 preview

- Assumption CI has two powerful implications:
 - We can write $EV_{\theta}(x_t, i_t)$ instead of $EV_{\theta}(x_t, \varepsilon_t, i_t)$,
 - We can consider a Bellman equation for V_θ (x_t), which is computationally simpler than the Bellman equation for V_θ (x_t, ε_t).

Theorem 1

Theorem 1

Given CI,

$$P(i|x,\theta) = \frac{\partial}{\partial \pi(x,i,\theta_1)} W(\pi(x,\theta_1) + \beta EV_{\theta}(x)|x,\theta_2)$$

and EV_{θ} is the unique fixed point of the contraction mapping:

$$EV_{\theta}(x,i) = \int_{y} W(\pi(y,\theta_{1}) + \beta EV_{\theta}(y) | y,\theta_{2}) p(dy|x,i,\theta_{3})$$

where

- $P(i|x,\theta)$ is the probability of action *i* conditional on state *x*
- $W(\cdot|x,\theta_2)$ is the surplus function:

$$W(v|x,\theta_2) \equiv \int_{\varepsilon} \max_{i} \left[v(i) + \varepsilon(i) \right] q(d\varepsilon|x,\theta_2)$$

Theorem 1 example: logit shocks

- ► $v_{\theta}(x, i) \equiv \pi(x, i, \theta_1) + \beta EV_{\theta}(x, i)$ the conditional value function.
- Suppose that $\varepsilon(i)$ is distributed independently across *i* with $Pr(\varepsilon(i) \le \varepsilon_0) = e^{-e^{-\varepsilon_0}} \text{logit shocks. Then,}$

$$W(v(x)) = \int \max_{i} [v(x,i) + \varepsilon(i)] \prod_{i} e^{-\varepsilon(i)} e^{-e^{-\varepsilon(i)}} d\varepsilon$$
$$= \ln \left(\sum_{i} \exp (v(x,i)) \right) + \gamma$$

where $\gamma \approx$.577216 is Euler's gamma.

It is then easy to derive expressions for conditional choice probabilities:

$$P(i|x,\theta) = \frac{\exp(v_{\theta}(x,i))}{\sum_{i'}\exp(v_{\theta}(x,i'))}$$

The conditional value function plays the same role as a static utility function when computing choice probabilities.

Some details

- He assumes ε is i.i.d with an extreme value type 1 distribution, and normalizes its mean to 0 and variance to π²/6 (i.e., the case on the previous slide).
- Transitions on observable state:

$$p(x_{t+1} - x_t = 0 |, x_t, i_t, \theta_3) = \theta_{30}$$

$$p(x_{t+1} - x_t = 1 |, x_t, i_t, \theta_3) = \theta_{31}$$

$$p(x_{t+1} - x_t = 2 |, x_t, i_t, \theta_3) = 1 - \theta_{30} - \theta_{31}$$

He tries several different specifications for the cost function and favors a linear form:

$$c(x,\theta_1)=\theta_{11}x.$$

Rust (1987) and NFP estimation

Nested Fixed Point Estimation

- Rust first considers a case with a closed-form expression for the value function, but this calls for restrictive assumptions on how mileage evolves. His nested fixed point estimation approach, however, is applicable quite generally.
- Basic idea: to evaluate objective function (likelihood) at a given θ, we should solve the value function for that θ

Nested Fixed Point Estimation

Steps:

- 1. Impute a value of the discount factor β
- 2. Estimate θ_3 the transition function for x which can be done without the behavioral model
- 3. Inner loop: search over (θ_1, RC) to maximize likelihood function. When evaluating the likelihood function for each candidate value of (θ_1, RC) :
 - 3.1 Find the fixed point of the the Bellman equation for $(\beta, \theta_1, \theta_3, RC)$. Iteration would work, but Rust uses a faster approach.
 - 3.2 Using expression for conditional choice probabilities, evaluate likelihood:

$$\prod_{t=1}^{T} P(i_t | x_t, \theta) p(x_t | x_{t-1}, i_{t-1}, \theta_3)$$

Estimates

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
β = .9999	$RC \\ \theta_{11} \\ \theta_{30} \\ \theta_{31} \\ LL$	11.7270 (2.602) 4.8259 (1.792) .3010 (.0074) .6884 (.0075) -2708.366	10.0750 (1.582) 2.2930 (0.639) .3919 (.0075) .5953 (.0075) -3304.155	9.7558 (1.227) 2.6275 (0.618) .3489 (.0052) .6394 (.0053) -6055.250	85.46	1.2E-17
eta=0	$RC \\ \theta_{11} \\ \theta_{30} \\ \theta_{31} \\ LL$	8.2985 (1.0417) 109.9031 (26.163) .3010 (.0074) .6884 (.0075) -2710.746	7.6358 (0.7197) 71.5133 (13.778) .3919 (.0075) .5953 (.0075) -3306.028	7.3055 (0.5067) 70.2769 (10.750) .3488 (.0052) .6394 (.0053) -6061.641	89.73	1.5E-18
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		

Discount factor

- While Rust finds a better fit for β = .9999 than β = 0, he finds that high levels of β basically lead to the same level of the likelihood function.
- Furthermore, the discount factor is non-parametrically non-identified. Note: He loses ability to reject $\beta = 0$ for more flexible cost function specifications.

Discount factor

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

TABLE VIII SUMMARY OF SPECIFICATION SEARCH⁸

⁴ 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 $\beta = 0$. ⁶ No convergence. Optimization algorithm attempted to drive $\theta_{1,2} \rightarrow 0$ and $\theta_{1,1} \rightarrow +\infty$.

Application



"Conditional Choice Probabilities and the Estimation of Dynamic Models" Hotz and Miller (1993)

Motivation

- A disadvantage of Rust's approach is that it can be computationally intensive
 - With a richer state space, solving value function (inner fixed point) can take a very long time, which means estimation will take a very, very long time.
- Hotz and Miller's idea is to use observable data to form an estimate of (differences in) the value function from conditional choice probabilities (CCP's)
- The central challenge of dynamic estimation is computing continuation values. In Rust, they are computed by solving the dynamic problem. With Hotz-Miller (or the CCP approach more broadly), we "measure" continuation values using a function of CCP's.

Notation

- actions a or $j \in J$, states $x \in X$
- ▶ with finite state space, state transition matrix can be represented by $|X| \times |X|$ matrices F_j (one matrix for each action)
- payoffs $\pi_j(x) + \varepsilon_j$
- distribution function G for idiosyncratic shocks ε
- ► conditional value function $v_j(x) = \pi_j(x) + \beta F_j V$, v_j denotes $|X| \times 1$ vector across states
- ex ante value function $V(x) = \int \max_j \{v_j(x) + \varepsilon_j\} dG(\varepsilon)$, V denotes $|X| \times 1$ vector across states

Rust's Theorem 1: Values to CCP's

▶ In Rust (1987), CCPs can be derived from the value function:

$$p_{j}(x) = \frac{\partial}{\partial \pi_{j}(x)} W(\pi(x) + \beta E[V(x')|x,j])$$

where $W(u) = \int \max_{j} \{u_j + \varepsilon_j\} dG(\varepsilon)$ is the surplus function.

For the logit case:

$$p_{j}(x) = \frac{\exp(v_{j}(x))}{\sum_{j' \in J} \exp(v_{j'}(x))}$$

where the conditional value function for action j in state x is

$$v_{j}(x) \equiv \pi_{j}(x) + \beta E \left[V(x') | x, j\right]$$

HM's Proposition 1: CCP's to Values

 Notice that CCP's are unchanged by subtracting some constant from every conditional value. Thus, consider

$$D_{j,0}v(x) \equiv v_j(x) - v_0(x)$$

where 0 denotes some reference action.

- ► Let $Q : \mathbb{R}^{|I|-1} \to \Delta^{|I|}$ be the mapping from the differences in conditional values to CCP's.
- Note: we're taking for granted that the distribution of ε is identical across states, otherwise Q would be different for different x.

Hotz-Miller Inversion

Q is invertible.

HM inversion with logit errors

- > Again, let's consider the case of where ε is i.i.d. extreme value type I.
- Expression for CCP's:

$$p_{j}(x) = \frac{\exp(v_{j}(x))}{\sum_{j' \in \mathbf{J}} \exp(v_{j}(x))}.$$

The HM inversion follows by taking logs and differencing across actions:

$$\ln p_{j}(x) - \ln p_{0}(x) = v_{j}(x) - v_{0}(x)$$

Thus, in the logit case

$$Q_{j}^{-1}\left(p\right) = \ln p_{j} - \ln p_{0}$$

From now on, I will use ϕ to denote Q^{-1} .

Arcidiacono and Miller's Lemma

An equivalent result to the HM inversion was introduced by Arcidiacono and Miller (2011). It's worth introducing here because it makes everything from now on much simpler and more elegant.

Arcidiacono Miller Lemma

For any action-state pair (a, x), there exists a function ψ such that

$$V(x) = v_{a}(x) + \psi_{a}(p(x))$$

Proof:

$$V(x) = \int \max_{j} \{v_{j}(x) + \varepsilon_{j}\} dG(\varepsilon_{j})$$

=
$$\int \max_{j} \{v_{j}(x) - v_{a}(x) + \varepsilon_{j}\} dG(\varepsilon_{j}) - v_{a}(x)$$
$$\int \max_{j} \{\phi_{ja}(p(x)) + \varepsilon_{j}\} dG(\varepsilon_{j}) - v_{a}(x)$$

Letting $\psi_{a}(p(x)) = \int \max_{j} \{\phi_{ja}(p(x)) + \varepsilon_{j}\} dG(\varepsilon_{j})$ completes the proof

Important relationships

The Hotz-Miller Inversion allows us to map from CCP's to differences in conditional value functions:

$$\phi_{ja}\left(p\left(x\right)\right) = v_{j}\left(x\right) - v_{a}\left(x\right)$$

The Arcidiacono and Miller Lemma allows us to relate ex ante and conditional value functions:

$$V(x) = v_j(x) + \psi_j(p(x))$$

For the logit case:

$$\phi_{ja}\left(p\left(x\right)\right) = \ln\left(p_{j}\left(x\right)\right) - \ln\left(p_{a}\left(x\right)\right)$$

$$\psi_{j}\left(p\left(x
ight)
ight) = -\ln\left(p_{j}\left(x
ight)
ight) + \gamma$$

where γ is Euler's gamma

Estimation example: finite state space I

- ► Let's suppose that X is a finite state space. Furthermore, let's "normalize" the payoffs for a reference action $\pi_0(x) = 0$ for all x.
 - We'll discuss soon whether this should really be called a "normalization"
- Using vector notation, recall the definition of the conditional value function for the reference action:

$$v_0 = \pi_0 + \beta F_0 V$$

$$v_0 = \beta F_0 V$$

Using the Arcidiacono-Miller Lemma,

Estimation example: finite state space II

Now we have an expression for the ex ante value function that only depends on objects we can estimate in a first stage:

. . . .

$$V = (I - \beta F_0)^{-1} \psi_0(p)$$

► To estimate the utility function for the other actions,

$$v_{j} = \pi_{j} + \beta F_{j} V$$

$$V - \psi_{j}(p) = \pi_{j} + \beta F_{j} V$$

$$\pi_{j} = -\psi_{j}(p) + (I - \beta F_{j}) V$$

$$\pi_{j} = -\psi_{j}(p) + (I - \beta F_{j}) (I - \beta F_{0})^{-1} \psi_{0}(p)$$

Identification of Models I

If we run through the above argument with π₀ fixed to an arbitrary vector π̃₀ rather than 0, we will arrive at the following:

$$\pi_j = A_a \widetilde{\pi}_0 + b_j$$

where A_a and b_a depend only on things we can estimate in a first stage:

$$A_j = (1 - \beta F_j) (1 - \beta F_0)^{-1}$$

$$b_{j} = A_{j}\psi_{0}\left(p\right) - \psi_{j}\left(p\right)$$

We can plug in any value for π̃₀, and each value will lead to a different utility function (different values for π_j). Each of those utility functions will be perfectly consistent with the CCP's we observe.

Identification of Models II

- ► Another way to see that the utility function is under-identified: If there are |X| states and |J| actions, the utility function has |X| |J| parameters. However, there are only |X| (|J| - 1) linearly independent choice probabilities in the data, so we have to restrict the utility function for identification.
- Magnac and Thesmar (2002) make this point as part of their broader characterization of identification of DDC models. Their main result says that we can specify a vector of utilities for the reference action *π̃*, a distribution for the idiosycratic shocks *G*, and a discount factor, and we will be able to find a model rationalizing the CCPs that features (*π̃*, *β*, *G*).

Identification of Counterfactuals

- Note that imposing a restriction like ∀x : π₀ (x) = 0 is NOT a normalization in the traditional sense. If we were talking about a static normalization, each x would represent a different utility function, and π₀ (x) = 0 would simply be a level normalization. However, in a dynamic model, the payoffs in one state affect the incentives in other states, so this is a substantive restriction.
- ► What is less clear a priori is whether these restrictions matter for counterfactuals. It turns out that some (but not all!) counterfactuals ARE identified, in spite of the under-identification of the utility function. What this means is that whatever value *π*₀ we impose for the reference action, the model will not only rationalize the observed CCP's but also predict the same counterfactual CCP's. Kalouptsidi, Scott, and Souza-Rodrigues (2016) sort out when counterfactuals of DDC models are identified and when they are not.