15 – Econometric Applications of Dynamic Programming
AGEC 642 - 2024

I. Positive analysis using a DP foundation: Econometric applications of DP

An active area of research involving numerical dynamic optimization, is the use of DP models in positive analysis. This work seeks to understand the nature of a particular problem being solved by a decision maker.[1] By specifying explicitly the optimization problem being solved, the analyst is able to estimate the parameters of a *structural* model, as opposed to the easier and more common *reduced form* approach. This approach was first developed by John Rust (1987), and I draw here directly on his original paper to explain how this is done. He provides a more complete description of his approach in Rust (1994a and 1994b).

My focus here will be quite narrow and serves only as an introduction to this literature. A more general review of approaches for the estimation of the parameters of dynamic optimization problems is provided by Keane et al. (2011) and surely more up-to-date surveys are now available.

Rust's paper is entitled "Optimal Replacement of GMC Bus Engines: An Empirical Model of Harold Zurcher." It is, in effect, an econometric estimate of the structural parameters that govern the decisions of Harold Zurcher, who was the superintendent of maintenance at the Madison (Wisconsin) Metropolitan Bus Company. The state variable x_t denotes the accumulated mileage (since last replacement) of the GMC bus engines of the bus fleet. Zurcher must make the choice as to whether to carry out routine maintenance, i=0, or replace the engine, i=1. Each period the operating costs, $c(x_t, \theta_1)$ where θ is a vector of parameters to be estimated. He assumes that mileage traveled in a period, say Δx, is distributed exponentially so that the cdf of Δx is $F(\Delta x) = 1 - \exp(\theta_2 \Delta x)$. If you replace the engine, i=1, then $x_{t+1} = \Delta x$, and if the engine is not replaced, i=0, then $x_{t+1} = x_t + \Delta x$.

An old engine has scrap value $\underline{P}$ and a new one costs $\overline{P}$. Hence, the additional cost to replace an engine is equal to $\left(\overline{P} - \underline{P}\right)$.

The problem he is solving, therefore, is what Rust calls a *regenerative optimal stopping* problem:

(3.5) $$V_\theta(x_t) = \max_{i_t \in \{0,1\}} \left[u(x_t, i_t, \theta_1) + \beta E V_\theta(x_t, i_t) \right]$$

where

[1] Adda and Cooper's textbook focuses particularly on this application of DP and the interested reader is referred to them for an introduction.

(3.6) $$EV_\theta(x_t,i_t)=\int_0^\infty V_\theta(y)\,p(y)(dy\,|\,x_t,i_t,\theta_2).$$[2]

where y is the index of the integral, i.e., what is used in place of x_{t+1} inside the integral, and $p(y)$ being a density function, i.e., $\partial F(y)/\partial y$. The econometric task is to estimate the parameter vector θ. To express this in more familiar notation, this would be equivalent to the Bellman's equation

$$V(x_t)=\max_{i_t\in\{0,1\}} u(x_t,i_t,\theta_1)+\beta E(V(x_{t+1}))$$
$$s.t. \qquad x_{t+1}=f(x_t,i_t,\theta_2,\varepsilon_t)$$

where ε_t~is a continuously distributed random variable with density equivalent to $p(y)$ in (3.6).

As Rust notes, the problem above has much in common with standard models of discrete choice in which the econometric problem is to identify the parameters θ that maximize the likelihood of observing the set of data. That is, if we let

$$\tilde{V}_\theta(x_t,i_t)=\left[u(x_t,i_t,\theta_1)+\beta EV_\theta(x_t,i_t)\right]$$

i.e., the value function with the choice i_t, then, for a given value for the parameter vector θ, we can write the probability of observing i_t=1, as the probability that $\tilde{V}_\theta(x_t,1)>\tilde{V}_\theta(x_t,0)$. This is essentially the same discrete-choice econometric problem as considered by McFadden (1973). In standard discrete choice models, however, the objective function is a static optimization problem that can be evaluated quickly. In a dynamic problem the function $\tilde{V}_\theta(x_t,1)$ is the solution to a dynamic optimization problem.

A. The approach he <u>doesn't</u> like

In earlier papers, Rust had shown that using the Bellman's equation, … there is an optimal stationary, Markovian replacement policy $\Pi=(f,f,\ldots)$ where f is given by

(3.7) $$i_t=f(x_t,\theta)=\begin{cases}1 & if \quad x_t>\gamma(\theta_1,\theta_2)\\ 0 & if \quad x_t\le\gamma(\theta_1,\theta_2)\end{cases}$$

where $\gamma(\theta_1,\theta_2)$ is the solution to

(3.8) $$\beta(\bar{P}-\underline{P})=\frac{1}{r}\int_0^{\gamma(\theta_1,\theta_2)}\left\{\left[1-\beta\exp\{-\theta_2(1-\beta)y\}\right]\frac{\partial c(y,\theta_1)}{\partial y}\right\}dy$$

where γ represents a threshold value of mileage. This equation is complicated, and we will not attempt to understand every term. However, there is clear intuition for each side of the equation. The left-hand side is the present value of the cost of replacing the engine in the next period. The right-hand side is the capitalized value (note $1/r$) of the stream of expected savings that result from having a new engine.

[2] I follow Rust's notation here. Following the notation more common in this series of lecture notes, $EV_\theta(x_t,i_t)$ would be written $EV_\theta(x_{t+1})$, $x_{t+1}=f(x_t,i_t)$.

He goes on to state:

> "The likelihood function $l\left(i_1,...,i_T,x_1,...,x_T,\theta\right)$ specifies the conditional probability density of observing the sequence of states and replacement decisions for a single bus in periods 1 to *T*. Under the assumption that monthly mileage and replacement decisions are independently distributed across buses, the likelihood function $L(\theta)$ for the full sample of data is simply the product of the individual bus likelihoods *l*. "(p. 1007)

This estimation approach, however, depends critically on functional form assumptions that are unlikely to hold. As Rust writes,

> The solution of the likelihood function depends critically on specific choice of functional form: namely, that monthly mileage (x_{t+1}-x_t) has an i.i.d. exponential distribution. Unfortunately, my sample of data flatly refutes this assumption.

The more general problem faced is provided in paragraphs on pages 1008-1009:

> A basic result in Markovian decision theory (cf. Blackwell (1968)) shows that under quite general conditions the solution to the class of infinite horizon Markovian decision problems takes the general form
>
> (3.9) $$i_t = f\left(x_t,\theta\right)$$
>
> where *f* is some deterministic function relating the agent's state variables x_t to his optimal action it. Suppose we assume that there are no unobserved state variables, i.e., that the econometrician observes all of x_t. The theory then implies that the data obey the deterministic relation (3.9) for some unknown parameter value θ^*. However, in general, real data will never exactly obey (3.9) for any value of the parameter θ: the data *contradict* the underlying optimization model. The typical solution to this problem is to "add an error term" ε_t in order to reconcile the difference between $f(x_t, \theta)$ and the observed choice i_t
>
> (3.10) $$i_t = f\left(x_t,\theta\right)+\varepsilon_t.$$
>
> By making a convenient distributional assumption for ε_t, one might use the model (3.10) to estimate θ. The difficulty with this procedure is that it is *internally inconsistent*: the structural model was formulated on the hypothesis that the agent's behavior is described by the solution of a dynamic optimization problem, yet the statistical implementation of that model implies that the agent randomly departs from this optimal solution. If error terms ε_t are to be introduced to a structural mode in an internally consistent fashion, they must be explicitly incorporated into the solution of the dynamic optimization problem. When this is done, a correct interpretation of the "error term" ε_t is that it is an *unobservable*, a state variable which is observed by the agent but not by the statistician .

B. The approach he <u>*does*</u> *like*

Rust (1987) proposes an alternative structural model. He notes that Zurcher's Bellman's equation is

(4.4) $$V_\theta(x_t,\varepsilon_t)=\max_i\left[u(x_t,i,\theta_1)+\varepsilon_t(i)+\beta EV_\theta(x_t,\varepsilon_t,i)\right],$$

which is a function of the distance that the engine has traveled, x_t, and an unobserved state variable, ε_t, which is known by Zurcher but not by the analyst. Writing out the expected value on the RHS of the Bellman's equation, we have

(4.5) $$EV_\theta(x_t,\varepsilon_t,i)=\int_y\int_\eta V_\theta(y,\eta)\,p(dy,d\eta\,|\,x_t,\varepsilon_t,i,\theta_2,\theta_3).$$

The function $p(\cdot)$ is a continuous Markov transition density function, which allows us to estimate the expected value of x_{t+1} conditional on the choice i, the current state, x_t, and the parameters θ_2 and θ_3, and y is the variable of integration for x_{t+1} and η is the variable of integration for the unobserved state variable, ε_{t+1}.
Rewriting (4.4) in way more familiar to us, we obtain

(4.4') $$V_\theta(x_t,\varepsilon_t)=\max_i\left[u(x_t,i,\theta_1)+\varepsilon_t(i)+\beta EV_\theta(x_{t+1},\varepsilon_{t+1})\right].$$

Hence, Rust's preferred approach involves having the unobserved state variable, ε_t, enter linearly into the benefit function. As we can see in (4.5), however, there is the possibility that ε_t will be correlated across time since the density function for x_{t+1} and ε_{t+1}, $p(x_{t+1},\varepsilon_{t+1}\,|\,x_t,\varepsilon_t,i,\theta_2,\theta_3)$, is conditional on x_t and ε_t. That would be a problem because the distribution at t+1, is a function of the variable in t, which is itself unknown. Hence, the distribution grows geometrically as we look from t to t+1, to t+2 and so on, making the problem intractable for both the computer and Zurcher.

To achieve an estimable model, Rust makes a critical assumption of Conditional Independence (CI) (p. 1011). This assumption "implies that any statistical dependence between ε_t and ε_{t+1} is transmitted entirely through the vector x_{t+1}. Second, the probability density of x_{t+1} depends only on x_t and not ε_t."

As an important practical matter, the CI assumption means that the transition density, can be written in a multiplicative form

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

where θ_2 and θ_3 are parameters governing the dynamics of the state variables ε and x respectively. In other words, $q(\cdot)$ and $p(\cdot)$ are independent Markov transition probabilities. Note that neither q not p are conditional on ε_t.

Making the CI assumption leads to his Theorem 1, which implies, (p. 1012) "that the conditional choice probabilities $P(i|x,\theta)$ [i.e., the conditional probability of choosing action i given the state variable x] can be computed using the same formulas used in the static case with the addition of the term $\beta EV_\theta(x,i)$ to the usual static utility term $u(x,i,\theta_1)$. Notice that McFadden's (1973), (1981) static model of discrete choice appears as a special case of Theorem 1 when $p(\cdot|x,i,\theta_3)$ is independent of i."

Following on McFadden, if we make the assumption that $q(\varepsilon|x,\theta_2)$ is multivariate extreme value, then the probability of observing choice i, $P(i|x,\theta)$, can be written as a multinomial logit:

(4.13) $$P(i\mid x,\theta)=\frac{\exp\left\{u(x_t,i,\theta_1)+\beta EV_\theta\left(x_{t+1}(x_t,i)\right)\right\}}{\sum_{j\in C(x)}\exp\left\{u(x,j,\theta_1)+\beta EV_\theta\left(x_{t+1}(x_t,j)\right)\right\}}\ .$$

The log likelihood function, therefore, can be written

$$LL=\sum_{k=1}^{n}\ln P(i_k\mid x_k,\theta),$$

where i_k and x_k are the k^{th} observation observed in the data set.

Hence, Rust suggests "the *following nested fixed-point* algorithm: an 'inner' fixed point algorithm computes the unknown function EV_θ for each value of θ and an "outer" hill climbing algorithm searches for the value of θ which maximizes the likelihood function." That is:

1) Choose a vector of parameters of $\hat{\theta}=\theta^1$
2) Solve the DP to find V_θ with $\theta=\hat{\theta}$
3) Calculate the likelihood function for your data set (and slopes of the likelihood function w.r.t the elements of θ)
4) Update $\hat{\theta}=\theta^2$ and return to step 1.

Since we will be iterating repeatedly on this, efficient solutions to the DP (step 2) is critical. Collocation methods (see Lecture 11) are frequently used for this step.

Rust is not overly optimistic about the ability to apply this approach to a wide range of problems. The computational burden is great. Each time we calculate the Bellman's equation it is necessary to solve an infinite horizon DP problem, and then you need an efficient algorithm for searching over the parameter space. He cites his own work, which showed that the contraction mapping T_θ (i.e., the successive approximation algorithm) is Fréchet differentiable.[3] This gives him two numerical advantages. He uses a method to obtain EV_θ and, as a by-product of this method, gets analytic solutions for the θ derivatives for EV_θ, which can then be used by the hill-climbing algorithm to maximize the likelihood function. (p. 1013)

C. Applications

There have been many applications of Rust's nested fixed-point approach and with the increasing speed of computers and the use of efficient solution methods, there is great scope for expanding the applications of the nested fixed-point algorithm. Some

---

[3] According Weisstein, a function is Fréchet differentiable at a if $\lim_{x\to a}\left(f(x)-f(a)\right)/(x-a)$ exists. That is Fréchet differentiability is essentially the standard concept of continuous differentiability.

applications by applied economists include papers that have studied the cow replacement problem (Miranda and Schnitkey, 1995), the value of recreationa fishing (Provencher, 1995, Baerenklau and Provencher, 2005) and the value of opportunities to hunt (Reeling, Verdier and Lupi, 2020).

An interesting extension of this approach was provided by Hendel and Nevo (2005) who consider the consumer problem of a storable good – laundry detergent.[4] In their application, the problem is further complicated by the fact that they never observe the state variable, the stock of laundry detergent in the household, and they do not observe withdrawals from the stock either, only additions when the individual makes a purchase. The bottom line is that they do find substantial differences between the static and dynamic models, with static models overestimating demand elasticities by 30%.

D. Other methods and critiques

Although these notes focus on Rust's nested fixed-point algorithm, his is not the only method that has been developed for estimating dynamic structural models. Other approaches include Keane and Wolpin (1994), Pakes et al., (2007) and Bajari et al., (2007, 2009). Rust (1994b) describes an alternative "forward" approach that "abandons the pretense of starting with an 'a priori' specification for *u* and instead conducts a specification search over the value functions *v* directly, using the estimated *v* functions to 'back out' estimates of the single-period utility functions *u* from the fixed-point condition" (p. 149). This approach is computationally much less burdensome, and he notes that in his 1988 paper, he finds that both methods generate "basically similar estimates of *u* and *v*."

The structural approach advocated by Rust, however, is not without its detractors. Heckman and Navarro (2007) describe Rust's approach as follows:

> [Rust (1994a)] shows that without additional restrictions, a class of infinite horizon dynamic discrete choice models for stationary environments is nonparametrically nonidentified. His paper has fostered the widespread belief that dynamic discrete choice models are identified only by using arbitrary functional form and exclusion restrictions. The entire dynamic discrete choice project thus appears to be without empirical content, and the evidence from it at the whim of investigator choices about functional forms of estimating equations and application of ad hoc exclusion restrictions.

Not surprisingly, they go on to offer an alternative.

---

[4] Hendel and Nevo (2005) is a testament to the lengths that researchers have gone to understand the deep mysteries behind the demand for laundry detergent, one of the most important social issues of the last century ;-).

II. References

Adda, Jérôme and Russell Cooper, 2003. *Dynamic Economics: Quantitative Methods and Application*. MIT Press: Cambridge, Mass.

Bajari, Patrick, C. Lanier Benkard and Jonathan Levin. 2007. Estimating dynamic models of imperfect competition. *Econometrica*, 75(5):1331–1370.

Baerenklau, Kenneth A. and Bill Provencher. 2005. "Static modeling of dynamic recreation behavior: implications for prediction and welfare estimation", *Journal of Environmental Economics and Management*, 50(3): 617-636.

Heckman, James J, and Salvador Navarro. 2007. Dynamic Discrete Choice and Dynamic Treatment Effects. *Journal of Econometrics* 136(2):341-96.

Hendel, Igal, and Aviv Nevo. 2006. Measuring the Implications of Sales and Consumer Inventory Behavior. *Econometrica* 74(6):1642-73.

Keane, Michael P. and Kenneth I. Wolpin. 1994. The Solution and Estimation of Discrete Choice Dynamic Programming Models by Simulation and Interpolation: Monte Carlo Evidence. *The Review of Economics and Statistics* 76(4):648–672.

Keane, Michael P., Petra E. Todd, and Kenneth I. Wolpin. 2011. The Structural Estimation of Behavioral Models: Discrete Choice Dynamic Programming Methods and Applications. *Handbooks in Economics - Handbook of Labor Economics, Edited by: Orley Ashenfelter and David Card* Vol 4(A):331-461.

McFadden, D. 1973. "Conditional Logit Analysis of Qualitative Choice Behavior," in P. Zarembka (ed.), *Frontiers in Econometrics*, New York: Academic Press.

Miranda, M.J. and G.D. Schnitkey. 1995. "Estimation of Dynamic Agricultural Decision Models: The Case of Dairy Cow Replacement." *Journal of Applied Econometrics* 10:41-56.

Pakes, Ariel, Michael Ostrovsky, and Steven Berry. 2007. Simple estimators for the parameters of discrete dynamic games (with Entry/Exit examples). *The RAND Journal of Economics* 38(2):373–399.

Provencher, Bill. 1995. Structural Estimation of the Stochastic Dynamic Decision Problems of Resource Users: An Application to the Timber Harvest Decision. *Journal of Environmental Economics and Management* 29:321-338.

Reeling, C., Verdier, V. and Lupi, F. (2020) 'Valuing goods allocated via dynamic lottery', *Journal of the Association of Environmental and Resource Economists*, 7(4), pp. 721–749. Available at: https://doi.org/10.1086/709142.

Rust, J. 1987. "Optimal Replacement of GMC Bus Engines: An Empirical Model of Harold Zurcher." *Econometrica* 55:999-1033

Rust, John. 1994a. Structural Estimation of Markov Decision Processes. *Handbook of Econometrics* IV. edited by RF Engle and DL McFadden IV (1994): 3082-3143. New York: Elsevier, 3082-3143

Rust, J. 1994b, "Estimation of dynamic structural models, problems and prospects: discrete decision processes," Chapter 4, pp. 119-170 of *Advances in Econometrics: Proceedings of the 6th World Congress of the Econometric Society*, edited by J.J. Laffont and C. Sims.

Weisstein, Eric W. "Fréchet Derivative." From MathWorld--A Wolfram Web Resource. http://mathworld.wolfram.com/FrechetDerivative.html