An Exposition of Structural Estimation of Discrete Dynamic Decision Processes[1]

Bill Provencher
University of Wisconsin-Madison

Kenneth A. Baerenklau
University of California-Riverside

[1] *In "Applications of Simulation Methods to Environmental and Resource Economics", edited by Riccardo Scarpa and Anna Aberini, Springer 2005.*

An Exposition of Structural Estimation of Discrete Dynamic Decision Processes

I. Introduction

In the analysis of dynamic decision problems (DDP's), the vast majority of the literature has focused on the *normative* aspects of such problems: what resources managers *should* do to maximize a particular objective function. Examples are ubiquitous in the literature, covering a range of issues from the land development decision, to the optimal rate of livestock grazing on arid land, to the well-known Faustmann model (and its many derivatives) of when to harvest a timber stand. Rarely have resource economists directly employed their structural dynamic models in the pursuit of the *positive* question, What decision problem do resource managers actually solve? The typical empirical study instead presents dynamic structural models in a theoretical discussion to motivate the choice of regressors in a reduced-form statistical analysis.[2] The primary explanation for this weakness in the literature is that such estimation is computationally very difficult. It usually requires a large amount of computer processing time, and perhaps more importantly, it requires a substantial amount of difficult programming by the analyst. Another explanation is that in the estimation of a parametric structural model the analyst is required to declare specific functional forms –state variable transition functions and benefit/utility functions, for instance –for which there is very little supporting evidence. This alone may give the analyst pause. A third explanation provides the justification for this chapter: by and large, resource economists remain fundamentally unfamiliar with the methods used to estimate the structure of DDP's.

This chapter takes the perspective that structural estimation of discrete DDP's is not especially difficult, and yet is difficult enough that most analysts require some explanation of why they should bother with it at all. In the next section we present the basic logic of structural estimation of discrete dynamic decision problems by developing the general case and distilling it to the tractable form most often found in the literature. In section three we illustrate the method using as a case study a modified version of the timber harvesting problem examined by Brazee and Mendelsohn (1988) (hereafter, the modified BM model). In section four we discuss, in the context of the modified BM model, the estimation of reduced-form and static models in lieu of structural estimation of the "true" dynamic model. We show that reduced-form estimation recovers the behavior engendered by the modified BM model with surprising ease, but that such estimation lacks policy-relevant empirical content. Moreover, insofar as one can relabel reduced-form estimation as structural estimation of a static decision model, it is not possible to distinguish empirically whether microeconomic data is generated by static or dynamic behavior, without maintained structural assumptions about the form of intraperiod benefit (utility, profit) functions, and even then it may not be possible. In the last section we briefly discuss future research directions and opportunities.

---

[2] Some exceptions in resource economics are Provencher (1995a,b), Provencher and Bishop (1997), and Howitt et al. (2002).

II. Structural Estimation of a Discrete Dynamic Decision Problem

Suppose an agent can make one of *S+1* choices in each period *t*. We index these choices by *s=0,…,S*. The amount of the payoff in a period depends on the state variables $\mathbf{x}$, which evolve over time according to the probability density function $f\left(\mathbf{x}_{t+1} \mid \mathbf{x}_t, s\right)$; the decision-specific random shock ε_i ; and perhaps other variables, denoted by **y**, that are invariant over time. The random shocks are contemporaneously observed by the agent but never observed by the analyst. For simplicity we assume these shocks are additive, and identically and independently distributed over time. Letting $R^s\left(\mathbf{x}_t, \mathbf{y}\right)+\varepsilon_t^s$ denote the decision-specific payoff at time *t*, the relevant decision problem can be stated in Bellman's form,

$$v\left(\mathbf{x}_t, \boldsymbol{\varepsilon}_t, \mathbf{y}\right)=\max_s\left[R^s\left(\mathbf{x}_t, \mathbf{y}\right)+\varepsilon_t^s+\beta E_{\mathbf{x}, \boldsymbol{\varepsilon} \mid \mathbf{x}_t, s}\left\{v\left(\mathbf{x}_{t+1}, \boldsymbol{\varepsilon}_{t+1}, \mathbf{y}\right)\right\}\right], \qquad (1.1)$$

where the expectation of future value is taken over both the observed state variables **x** and unobserved state variables $\boldsymbol{\varepsilon}=\left(\varepsilon^0, \ldots, \varepsilon^S\right)$, conditional on the current values of observed state variables and the decision *s*.

The problem in (1.1) reflects Bellman's principle of optimality; in particular, it implicitly recognizes that regardless of the decision at time *t*, optimal decisions are made in the future. This problem can be solved via the recursive methods of dynamic programming. The solution is an optimal decision rule $s\left(\mathbf{x}, \boldsymbol{\varepsilon}, \mathbf{y} ; \Gamma\right)$, where Γ is the set of structural parameters associated with the decision problem. This includes the parameters implicit in $R^s\left(\mathbf{x}, \mathbf{y}\right)$, the parameters of the density function $f\left(\cdot\right)$, the parameters of the distribution of $\boldsymbol{\varepsilon}$, and the discount factor β .

Now suppose there exist observations over time *t=1,…,T*, and across agents *j=1,…,J*, of decisions s_{jt} and variables $\mathbf{x}_{jt}$ and $\mathbf{y}_j$. Then defining,

$$v_t^s=R^s\left(\mathbf{x}_t, \mathbf{y}\right)+E_{\mathbf{x}, \boldsymbol{\varepsilon} \mid \mathbf{x}_t, s}\left\{v\left(\mathbf{x}_{t+1}, \boldsymbol{\varepsilon}_{t+1}, \mathbf{y}\right)\right\},$$

the probability that agent *j* makes choice *0* at time *t* is given by,

$$\begin{aligned}\Pr\left(s_{jt}=0 \mid \mathbf{x}_{jt}, \mathbf{y}_j ; \Gamma\right) &= \Pr\left(v_t^0+\varepsilon_t^0 \geq v_t^1+\varepsilon_t^1 ; \ldots ; v_t^0+\varepsilon_t^0 \geq v_t^S+\varepsilon_t^S\right) \\ &= \int_{\varepsilon_t^0 \geq \Delta v_t^{1,0}+\varepsilon_t^1} \cdots \int_{\varepsilon_t^0 \geq \Delta v_t^{S,0}+\varepsilon_t^S} g\left(\varepsilon^1\right) \cdots g\left(\varepsilon^S\right) d \varepsilon^1 \ldots d \varepsilon^S ,\end{aligned} \qquad (1.2)$$

where $\Delta v_t^{s,0} = v_t^s - v_t^0$, and $g(\cdot)$ is the common distribution of the unobserved state variables ε_t^s. The probability that in period t agent j makes a different choice $s_{jt}=1,\ldots,S$ can be constructed in similar fashion.

Denoting by s_{jt}^O the observed decision by agent j in period t, and assuming agent decisions are independent of one another, the likelihood function is

$$L(\Gamma) = \prod_{j=1}^{J} \prod_{t=1}^{T} \Pr\left(s_{jt}^O \mid \mathbf{x}_{jt}, \mathbf{y}_j; \Gamma\right). \tag{1.3}$$

Estimation of the parameters Γ is a complicated and computationally intensive exercise, much more so than is found with typical static models. The search for the parameter vector that maximizes the likelihood function involves solving the dynamic decision problem (1.1) *each time* new parameter values are evaluated in the search; this is apparent by the presence of v_t^s in the limit of integration in (1.2). Solving the decision problem (1.1) requires dynamic programming, and so a DP simulation must be run *each time* new parameter values are evaluated in the search for the maximum likelihood value. Maximum likelihood estimation thus involves nesting an "inner" dynamic programming algorithm within an "outer" hill-climbing algorithm.

As a general matter, the inclusion of the unobserved state variables $\boldsymbol{\varepsilon}$ seriously impacts the tractability of the dynamic programming problem (1.1), as the expectation of the value function must be taken over the random components of the observed state vector **x** *and* the S-dimensional vector $\boldsymbol{\varepsilon}$. For all but the smallest problems this is not computationally feasible given that (1.1) must be solved many times during the course of the search for Γ^*, the likelihood-maxmizing value of Γ. Making matters even more difficult is the multi-dimensional integration in (1.2) associated with each observation. The reader familiar with the literature on random utility models will recognize that the properties of the Gumbel distribution can be used to resolve the difficulty of the integration in (1.2). In a seminal paper Rust (1989) shows that these same properties of the Gumbel distribution can be used to resolve the difficulty of the integration in (1.1), as follows.

We assume that $\boldsymbol{\varepsilon}$ is iid Gumbel-distributed with location parameters $\boldsymbol{\theta}_\varepsilon = \left(\theta^0, \ldots, \theta^S\right)$ and common scale parameter η_ε. Moreover, for expositional reasons we define

$$V\left(\mathbf{x}_t, s_t; \Gamma\right) = E_{\mathbf{x},\boldsymbol{\varepsilon}|\mathbf{x}_t, s_t}\left\{v\left(\mathbf{x}_{t+1}, \boldsymbol{\varepsilon}_{t+1}, \mathbf{y}\right)\right\}, \tag{1.4}$$

in which case the decision problem (1.1) can be restated,

$$v\left(\mathbf{x}_t, \boldsymbol{\varepsilon}_t, \mathbf{y}\right) = \max_s \left[R^s\left(\mathbf{x}_t, \mathbf{y}\right) + \varepsilon_t^s + \beta V\left(\mathbf{x}_t, s; \Gamma\right)\right]. \tag{1.5}$$

Then from the standard properties of the Gumbel distribution (see Ben-Akiva and Lerman (1985)), integrating both sides of (1.5) with respect to $\boldsymbol{\varepsilon}$ on day *t+1* yields,

$$E_{\boldsymbol{\varepsilon}} v\left(\mathbf{x}_{t+1}, \boldsymbol{\varepsilon}_{t+1}, \mathbf{y}\right) = \frac{1}{\eta_\varepsilon} \ln\left(\sum_{s=0}^{S} \mathrm{e}^{\eta_\varepsilon\left(R^s\left(\mathbf{x}_{t+1}, \mathbf{y}\right) + \theta^s + \beta V\left(\mathbf{x}_{t+1}, s; \Gamma\right)\right)} \right) + \frac{\gamma}{\eta_\varepsilon}, \qquad (1.6)$$

where γ is Euler's constant (~0.577). Substitution of (1.6) into (1.4) at time *t+1* yields,

$$V\left(\mathbf{x}_t, s_t; \Gamma\right) = E_{\mathbf{x}|\mathbf{x}_t, s_t} \frac{1}{\eta_\varepsilon} \ln\left(\sum_{s=0}^{S} \mathrm{e}^{\eta_\varepsilon\left(R^s\left(\mathbf{x}_{t+1}, \mathbf{y}\right) + \theta^s + \beta V\left(\mathbf{x}_{t+1}, s; \Gamma\right)\right)} \right) + \frac{\gamma}{\eta_\varepsilon}. \qquad (1.7)$$

With $v(\cdot)$ known, the determination of $V(\cdot)$ –necessary to solve the decision problem (1.5) –is now a relatively simple affair involving integration over the random elements in the observable state vector $\mathbf{x}_{t+1}$. In practice, the decision problem (1.5), in conjunction with (1.6) and (1.7), can be solved via successive approximation (backwards recursion) due to the contraction mapping properties of Bellman's form. In the first stage $V(\cdot)$ is set identically equal to zero, and the value function $v(\cdot)$ is approximated from (1.5) using standard techniques for approximating a function, such as linear or Chebychev polynomial interpolation. In the second stage the expectation of $v(\cdot)$ with respect to $\boldsymbol{\varepsilon}$ is calculated from (1.6), and the function $V(\cdot)$ is approximated from (1.7) using standard techniques of function interpolation and numerical integration. The approximation of $V(\cdot)$ is then used in (1.5) to obtain a new approximation of $v(\cdot)$. This iterative mechanism terminates under conditions for convergence, such as sufficiently small changes across iterations in the optimal decision rule, $s\left(\mathbf{x}, \boldsymbol{\varepsilon}, \mathbf{y}; \Gamma\right)$.

It deserves emphasis that the algorithm described above identifies the optimal decision rule –and, more to the point, the associated expected value function $V(\cdot | \Gamma)$ –*conditional* on a particular set of parameters Γ. With the expected value function in hand, the assumption that $\boldsymbol{\varepsilon}$ is iid Gumbel-distributed allows a restatement of (1.2) in the analytical form,

$$\begin{aligned}\Pr\left(s_{jt} = 0 \,|\, \mathbf{x}_{jt}, \mathbf{y}_j; \Gamma\right) &= \Pr\left(v_t^0 + \varepsilon_t^0 \geq v_t^1 + \varepsilon_t^1; \ldots; v_t^0 + \varepsilon_t^0 \geq v_t^S + \varepsilon_t^S\right) \\ &= \frac{e^{\eta_\varepsilon\left(R^0\left(\mathbf{x}_t, \mathbf{y}\right) + \theta^0 + \beta V\left(\mathbf{x}_t, 0; \Gamma\right)\right)}}{\sum_{s=0}^{S} e^{\eta_\varepsilon\left(R^s\left(\mathbf{x}_t, \mathbf{y}\right) + \theta^s + \beta V\left(\mathbf{x}_t, s; \Gamma\right)\right)}}. \qquad (1.8)\end{aligned}$$

The upshot is that although it remains the case that a DP algorithm must be nested within the estimation algorithm used to find Γ^*, assuming the unobserved state variables are iid Gumbel-distributed greatly simplifies the algorithm.

Examination of (1.8) also shows that a well-documented weakness of static multinomial logit models –the property of *independence of irrelevant alternatives* (IIA) –does not hold in a dynamic model employing Gumbel-distributed random variables. The IIA property derives from the fact that, from the perspective of the analyst, the odds that one alternative is chosen over another in a static model depends only on the attributes of the two alternatives. So, for instance, in the example presented by Bockstael, the IIA property implies the unlikely result that the odds of visiting a saltwater beach instead of a freshwater lake does not depend on whether a third beach is itself a saltwater or freshwater site. However, the log odds ratio for the dynamic model can be stated,

$$\log\left(\frac{\Pr\left(s_{jt}=i\right)}{\Pr\left(s_{jt}=k\right)}\right)=\eta_{\varepsilon}\left(\left(R^{i}\left(\mathbf{x}_{t},\mathbf{y}\right)+\theta^{i}+\beta V\left(\mathbf{x}_{t+1},i;\Gamma\right)\right)-\left(R^{k}\left(\mathbf{x}_{t},\mathbf{y}\right)+\theta^{k}+\beta V\left(\mathbf{x}_{t+1},k;\Gamma\right)\right)\right).$$

And so by virtue of the presence of $\mathbf{x}_t$ and $\mathbf{y}$ in $V(\cdot)$, the odds of choosing alternative *i* over alternative *k* depends on the attributes (state of nature) of all the alternatives.

III. An Illustration: The Brazee-Mendelsohn Timber Harvest Problem

In this section we attempt to clarify the discussion above by examining a modification of the Brazee-Mendelsohn (BM) model of optimal harvesting when timber prices are stochastic (Brazee and Mendelsohn 1988). This is a simple but illuminating example. In the original BM model the forest owner faces, in each period, the binary decision to either harvest a timber stand or to postpone harvest. This decision depends on two state variables: the age of the timber stand *a*, and the price of timber *p*. Timber volume at stand age *a* is given by

$$W(a)=e^{\phi_1-\frac{\phi_2}{a}} \tag{1.9}$$

Timber prices are independent and identically normally-distributed with mean price μ_p and standard deviation σ_p; we denote the probability density function by $g\left(\mu_p,\sigma_p\right)$. The forest owner solves the the problem,

$$v\left(a_t,p_t\right)=\max\left[\beta E_p v\left(a_t+1,p_{t+1}\right),\ \ p_t W\left(a_t\right)-c+\beta E_p v\left(1,p_{t+1}\right)\right], \tag{1.10}$$

where *c* is the cost to harvest and replant. The problem is easily solved by ordinary dynamic programming techniques. The optimal harvest policy $s(p,a)$ is a reservation price policy in which, for a given stand age, the forest owner harvests if and only if the observed timber price is above a reservation price. Graphically this is represented by the harvest isocline shown in Figure 1; harvest occurs for all combinations of price and stand age above the isocline.

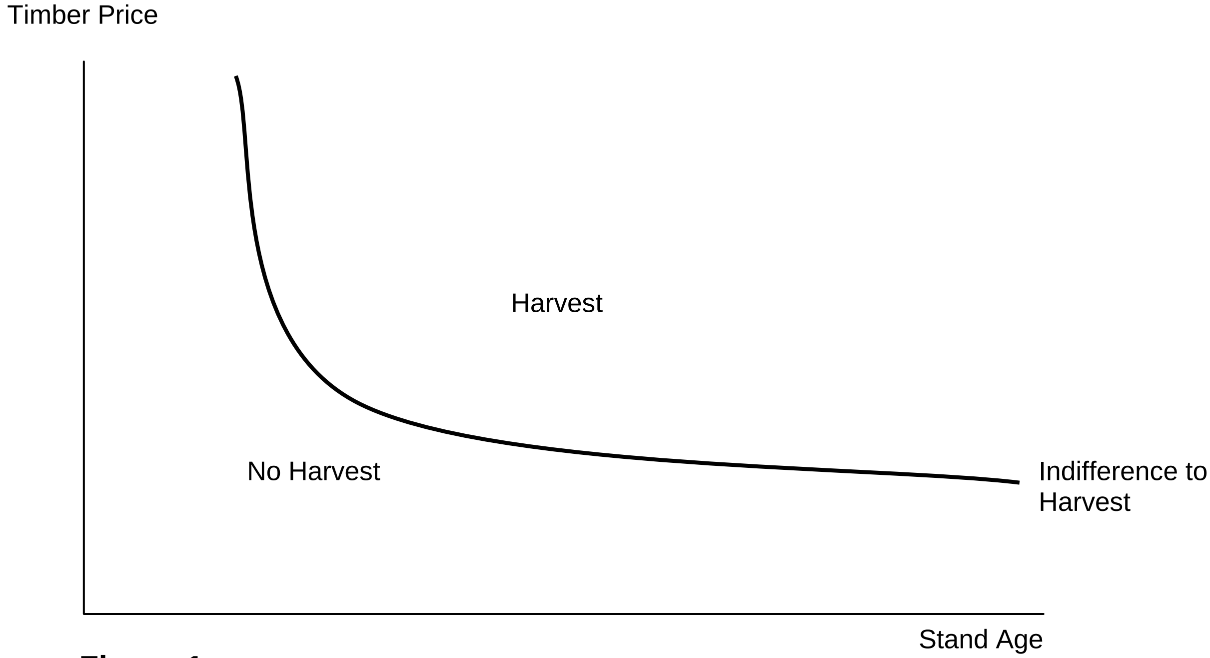

Figure 1.

When applying this model to actual data the analyst must account for the possibility that observed harvest decisions deviate from the normative decision rule. This is accomplished by introducing the decision-specific state variables $\boldsymbol{\varepsilon} = \left(\varepsilon^0, \varepsilon^1\right)$. This paradox –that the normative model does not fully describe the decision problem faced by "real world" dynamic optimizers, and so the normative model is, as a practical matter, not normative at all –is almost invariably the case for dynamic problems of resource allocation, because normative models distill a rich decision environment to a world fully described by just several state variables. The addition of the unobserved state variables can be viewed as a somewhat crude attempt to account for the richness of the real world.

The modified decision problem is,

$$\begin{aligned} v\left(a_t, p_t, \boldsymbol{\varepsilon}_t\right) &= \max\left[\varepsilon_t^0 + \beta E_p v\left(a_t + 1, p_{t+1}\right),\ \ p_t W\left(a_t\right) - c + \varepsilon_t^1 + \beta E_p v\left(1, p_{t+1}\right)\right] \\ &= \max\left[\varepsilon_t^0 + \beta V\left(a_t, 0; \Gamma\right),\ \ p_t W\left(a_t\right) - c + \varepsilon_t^1 + \beta V\left(a_t, 1; \Gamma\right)\right] \end{aligned} \quad ,(1.11)$$

where the control variable takes a value of 1 if the stand is harvested and 0 otherwise. Note that (1.11) is a special case of (1.1) in which only one of the observed state variables is stochastic (p), and the distribution of the state variable is not conditional on the current value. In this framework, one possible interpretation of ε is that it is the utility received from the standing forest (though in this simple case, the utility received is not conditional on stand age –generally an unrealistic specification).

Assuming that $\boldsymbol{\varepsilon}$ is iid Gumbel-distributed with choice-specific location parameters θ^i, *i=0,1*, and common scale parameter η_ε, the expected value function can be stated (using a modified version of (1.7)):

$$\begin{aligned} V\left(a_t,0;\Gamma\right) &= \int_{-\infty}^{\infty} \frac{1}{\eta_\varepsilon} \ln\left(\mathrm{e}^{\eta_\varepsilon\left(\theta^0+\beta V\left(a_t+1,0;\Gamma\right)\right)} + \mathrm{e}^{\eta_\varepsilon\left(p_{t+1}W\left(a_t+1\right)-c+\theta^1+\beta V\left(a_t+1,1;\Gamma\right)\right)} \right) g\left(p;\mu_p,\sigma_p\right)dp + \frac{\gamma}{\eta_\varepsilon} \\ V\left(a_t,1;\Gamma\right) &= \int_{-\infty}^{\infty} \frac{1}{\eta_\varepsilon} \ln\left(\mathrm{e}^{\eta_\varepsilon\left(\theta^0+\beta V\left(1,0;\Gamma\right)\right)} + \mathrm{e}^{\eta_\varepsilon\left(p_{t+1}W\left(1\right)-c+\theta^1+\beta V\left(1,1;\Gamma\right)\right)} \right) g\left(p;\mu_p,\sigma_p\right)dp + \frac{\gamma}{\eta_\varepsilon} \end{aligned} \qquad (1.12)$$

Given parameters $\Gamma = \left\{\beta,\phi_1,\phi_2,\mu_p,\sigma_p,\eta_\varepsilon,\theta^0,\theta^1\right\}$, $V\left(\cdot\right)$ can be approximated by a simple iterative recursion. In the first iteration, $V\left(\cdot\right)$ on the right-hand side of (1.12) is set to an arbitrary value, and an update of $V\left(\cdot\right)$ is found by solving (1.12) for each stand age *a* (this requires numerical approximation of the integral taken over timber prices). The update is then used on the right-hand side of (1.12) in the second iteration, and so on, until a convergence criterion is met.

For the analyst with observations on the state variables *p* and *x*, as well as the actual harvest decision $s_{jt} \in \left\{0,1\right\}$, the probability of the observed harvest decision by forest owner *j* at time *t* is:

$$\Pr\left(s_{jt} \mid p_t, a_{jt};\Gamma\right) = \frac{\left(1-s_{jt}\right)e^{\eta_\varepsilon\left(\theta^0+\beta V\left(a_{jt},0;\Gamma\right)\right)} + s_{jt}e^{\eta_\varepsilon\left(p_tW\left(a_{jt}\right)-c+\theta^1+\beta V\left(a_{jt},1;\Gamma\right)\right)}}{e^{\eta_\varepsilon\left(\theta^0+\beta V\left(a_{jt},0;\Gamma\right)\right)} + e^{\eta_\varepsilon\left(p_tW\left(a_{jt}\right)-c+\theta^1+\beta V\left(a_{jt},1;\Gamma\right)\right)}} \qquad (1.13)$$

And the likelihood function takes the form,

$$L\left(\Gamma\right) = \prod_{j=1}^{J}\prod_{t=1}^{T}\Pr\left(s_{jt} \mid p_t, a_{jt};\Gamma\right) \qquad (1.14)$$

It is important to understand that even though the harvest problem involves four state variables –the observed state variables *a* and *p* and the unobserved state variables ε^0 and ε^1 –equation (1.12) is the basis of the "inner" DP algorithm used in estimation, and so the only state variable affecting the dimensionality of the algorithm is stand age (*a*). For a given set of parameters Γ, $V\left(\cdot\right)$ is found by iteration over (1.12), as described above. With $V\left(\cdot\right)$ in hand, the likelihood value is calculated from (1.13) and (1.14). The outer search algorithm then chooses an alternative set of parameters, and so on, until convergence.

The most useful measure of forest value is the value conditional on stand age and observed price. From (1.11) and the properties of the Gumbel distribution, the expected value of the forest land for a given stand age and timber price is,

$$E_{\boldsymbol{\varepsilon}} v\left(a_t, p_t, \boldsymbol{\varepsilon}_t\right) = E_{\boldsymbol{\varepsilon}} \max\left[\varepsilon_t^0 + \beta V\left(a_t, 0; \Gamma\right),\ \ p_t W\left(a_t\right) - c + \varepsilon_t^1 + \beta V\left(a_t, 1; \Gamma\right)\right]$$
$$= \frac{1}{\eta_\varepsilon} \ln\left(\mathrm{e}^{\eta_\varepsilon\left(\theta^0 + \beta V\left(a_t, 0; \Gamma\right)\right)} + \mathrm{e}^{\eta_\varepsilon\left(p_t W\left(a_t\right) - c + \theta^1 + \beta V\left(a_t, 1; \Gamma\right)\right)} \right) + \frac{\gamma}{\eta_\varepsilon} \quad . \ (1.15)$$

Some Numerical Results Illustrating the Model

For a given timber price and stand age, harvest is probabilistic because it depends on the values of $\boldsymbol{\varepsilon}$ generated for the period. Figure 2a-c presents probabilities of harvest for three variations of the modified BM model (that is, the BM model with $\boldsymbol{\varepsilon}$ included). All models consider the case of Loblolly Pine on a low quality site, using parameter values found in Brazee and Mendelsohn (1988). Timber growth is implied by volume parameters $\phi_1 = 12.09$ and $\phi_2 = 52.9$ (see (1.9)). Timber prices are normally distributed with mean \$167.4 per thousand board feet, and standard deviation \$40.41. The discount rate is 3% (discount factor =.97). Harvest and replanting costs are \$147, and $\boldsymbol{\theta} = 0$.

The models underlying the panels of Figure 2 differ in the value of the scale parameter, η_ε, with this value falling across the three panels. Importantly, the scale parameter of the Gumbel distribution is inversely proportional to the variance (variance= $\pi^2 / 6\eta_\varepsilon^2$), and so the variance in the distribution of $\boldsymbol{\varepsilon}$ is *rising* across the panels. In the first panel the scale parameter is especially large in the context of the model ($\eta_\varepsilon = 20.0$). Consequently, values of $\boldsymbol{\varepsilon}$ are invariably close to zero, and so the model is, for all intents and purposes, the same as the original BM model; that is, for a given stand age, the timber is harvested if and only if the timber price is above the reservation price for the stand age. Graphically this is represented in Figure 2a by a distinct “probability ledge” tracing the harvest isocline of the original BM model: as the stand ages the reservation price falls (see Figure 1). Henceforth we refer to this model as the “virtually no variance” (VNV) model. In Figure 2b, $\eta_\varepsilon = 2.0$ (the “low variance” (LV) model), and so the reservation price policy no longer applies, as indicated by the transformation of the probability “ledge” of Figure 2a to a steep probability “hill” in Figure 2b. An unexpected result apparent in the figure is that harvest often occurs at very young stand ages. This is the case because growth at these ages is relatively low, and so if, at these young stand ages, the difference $\varepsilon^1 - \varepsilon^0$ is sufficiently high –a distinct possibility because of the relatively low value of the scale parameter –it is advantageous to harvest the stand and start over. In Figure 2c, $\eta_\varepsilon = 0.2$ (the “high variance” (HV) model), and the probability surface is noticeably smoother than in Figure 2b. Essentially the harvest decision is now heavily driven by the observed values of ε^0 and ε^1.

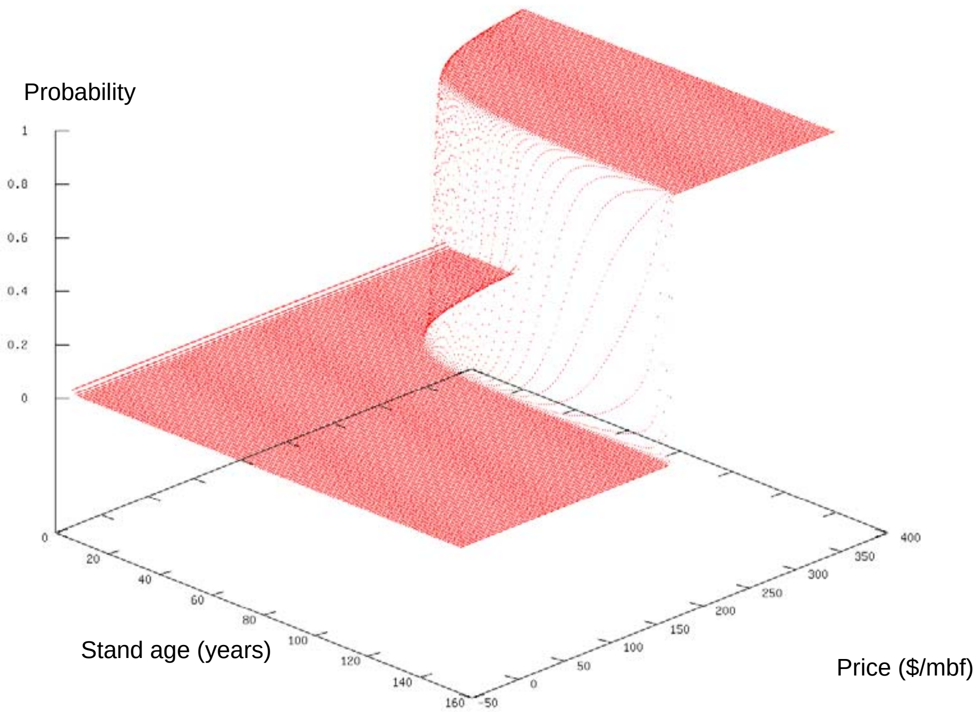

Figure 2a. Probability of harvest with $\eta_\varepsilon = 20.0$

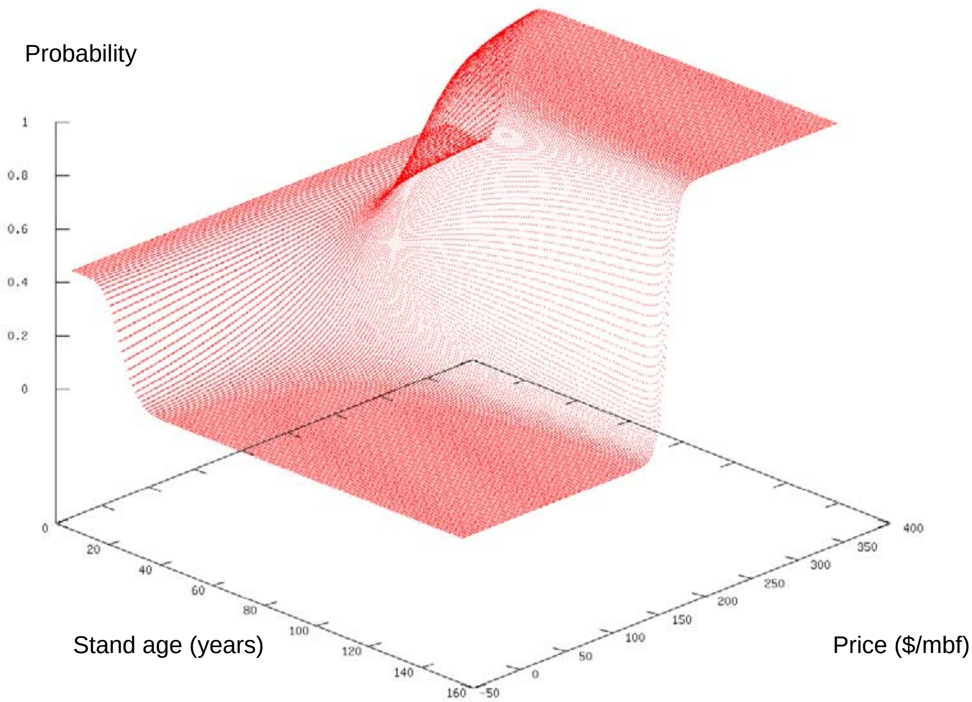

Figure 2b. Probability of harvest with $\eta_{\varepsilon} = 2.0$

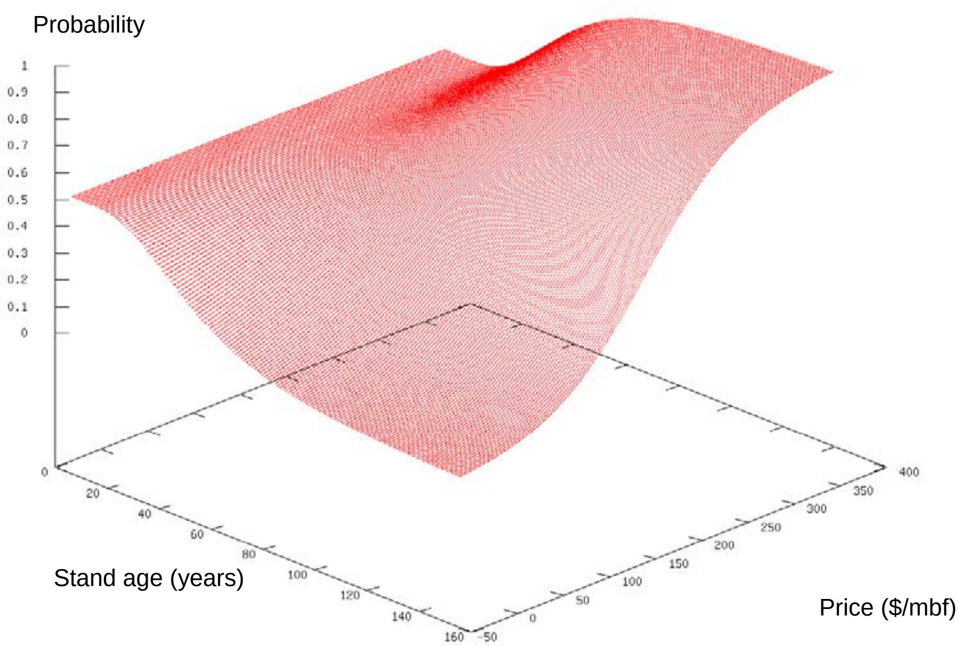

Figure 2c. Probability of harvest with $\eta_{\varepsilon} = 0.2$

Figure 3a-b presents simulated data from the VNV and LV models. The data involve 500 timber stands of varying ages observed over an 80-year sequence. Initial stand ages for the sequences are random draws from a uniform distribution in the range [1,150]. Prices for the sequences were random draws from the price distribution (a single 80-year price sequence applies to all timber stands in the figures). The figures show that the distribution of timber harvests is very different for the two models. The VNV model has relatively few harvests, and these are concentrated in several years. The HV model has many more harvests spread fairly evenly over time, reflecting the very young rotation ages engendered by the relatively high variation in $\boldsymbol{\varepsilon}$. Although this is not realistic (generally southern pine must be at least 30 years old to be harvested for sawtimber), the sharp contrast with results from the VNV is useful because it allows us to explore whether qualitative differences in data affect the structural estimation of discrete dynamic decision processes.

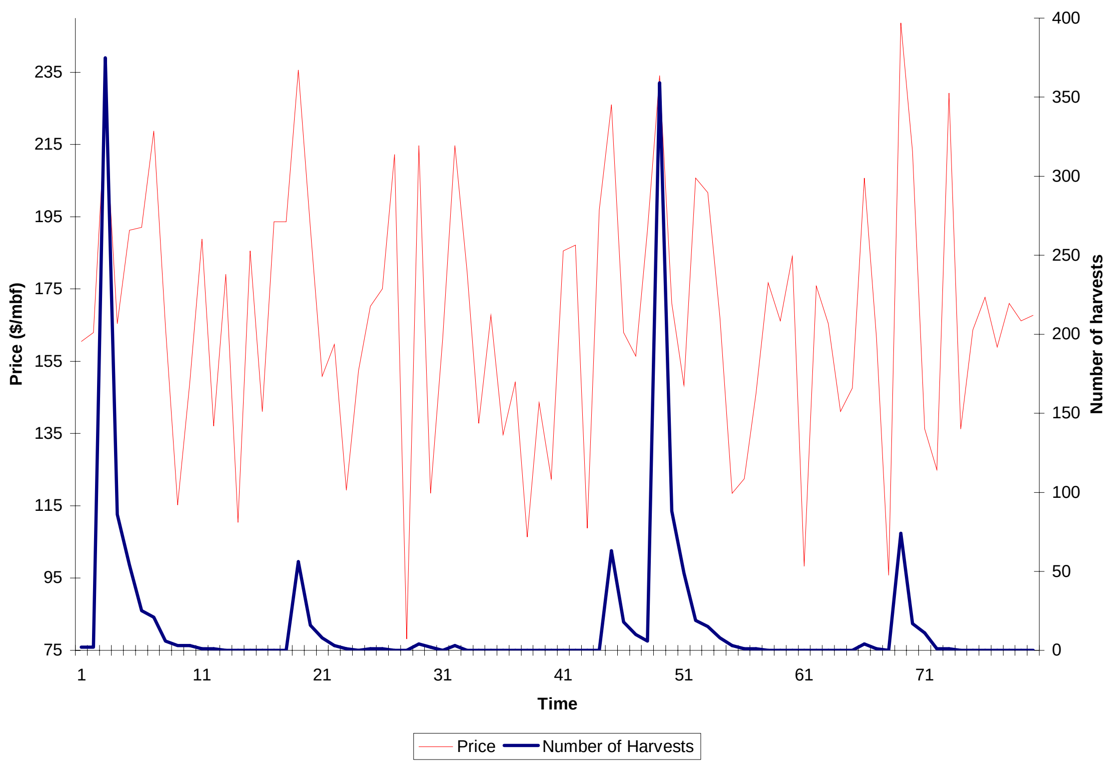

Figure 3a. Simulated Harvest Data for the VNV model (η_{ε} =20).

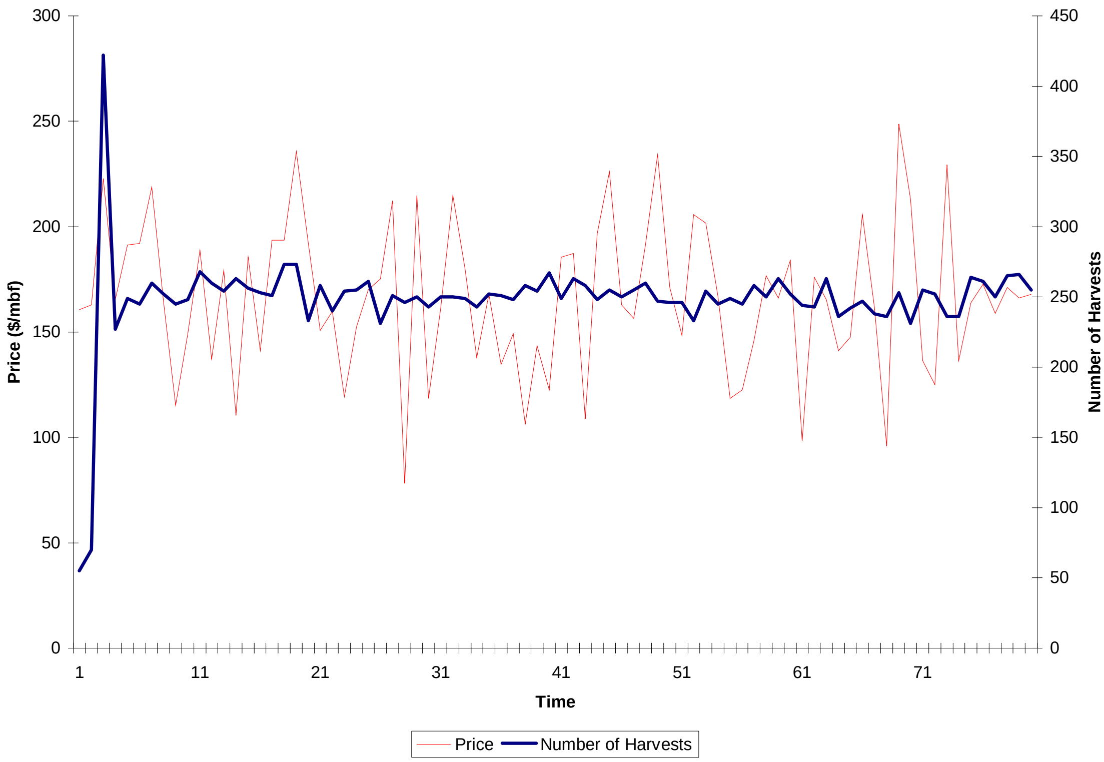

Figure 3b. Simulated Harvest Data for the LV Model (η_ε =2.0)

Table 1 presents estimation results for the simulated data. We present results for both the full simulated samples of the VNV and LV models, and for partial samples in which estimation is based on only the last 40 and 20 years of the sequences.[3] For all seven estimations, the exogenous parameters –the price parameters μ_p and σ_p and the volume parameters ϕ_1 and ϕ_2 --are fixed at their actual values.[4] This leaves four potential parameters to be estimated: $\beta, \theta^0, \theta^1, \eta_\varepsilon$. But because the values of not harvesting and harvesting are linear in ε^0 and ε^1, respectively, only the *difference* in the location parameters θ^1 and θ^0 can be identified. To see this, observe from (1.11) that the timber owner harvests if

$$p_t W(a_t) - c + \beta\left(V(a_t,1;\Gamma) - V(a_t,0;\Gamma)\right) + \varepsilon_t^1 - \varepsilon_t^0 > 0 \quad . \qquad (1.16)$$

[3] Reducing the sample in the time dimension reflects the judgment that it is easier for researchers to obtain a large cross section of timber harvest data than a large time series.

[4] In actual estimation, the corresponding approach is to estimate the growth function and timber price process exogenously (that is, outside the main estimation algorithm), and to use the estimated parameter values in the main estimation algorithm, thereby significantly reducing the size of the estimation problem. This approach of estimating exogenous processes outside of the main estimation algorithm is common in the literature. Alternatively, the analyst could decide, for instance, that timber owners may be using an incorrect price process, in which case the parameters of the price process would be included in the set of parameters to be estimated in the main estimation algorithm. In this case an iterative estimation algorithm suggested by Rust (1994b) would be useful.

Defining $\Delta\theta = \theta^1 - \theta^0$, it follows from the properties of the Gumbel distribution that the harvest decision can be cast as logistic with location parameter,

$$p_t W(a_t) - c + \beta\big(V(a_t, 1; \Gamma) - V(a_t, 0; \Gamma)\big) + \Delta\theta \ ,$$

and scale parameter η_ε; significantly, only the difference $\Delta\theta$ can be identified. This is reflected in Table 1, where we fix θ^0 at zero and estimate θ^1.

When the full sample is used, estimation results are generally excellent for both the VNV and the LV models. For all models except the LV model with half the sample (Model 5 in Table 1), the estimate of the discount rate is within 2/1000 of the actual value, and even for Model 5 the true value is within two standard deviations of the estimated value. This no doubt reflects the tremendous influence of the discount rate on the harvest decision. By comparison, in estimations using partial samples the estimates of η_ε and θ^1 tend to be *statistically* different than the true values of these parameters. Are these differences significant as a *practical* matter? Yes and no. In an investigation of timber harvest behavior, the analyst is primarily interested in two questions: the effect of stand age and timber price on the likelihood of harvest, and the expected value of bare forestland (In particular, the nontimber value of forestland). On both counts, models 2, 3, and 6 generate results very similar to the true model. The exception to these generally favorable estimation results is model 5, for which the estimated discount factor is considerably lower than the true discount factor. This lower discount factor (higher discount rate) has two significant effects. First, the estimated value of bare land is much higher for Model 5 than the actualvalue --$9282 per acre versus an actual value of $1456 per acre. Of course, the large standard error on the estimate of the discount factor in Model 5 signals the analyst that the estimate of bare land value is imprecise. Second, for timber stands that by chance mature past age 20 or so, the predicted harvest age is lower for Model 5 than for the true model. [5] Figure 4 tells the story on this. It presents the *difference* in harvest probabilities between the true model and the estimated Model 5. The deep trough in Figure 4 beginning at about stand age 25 in the price range of roughly $200-$270, indicates that for older stands Model 5 overpredicts the probability of harvest, and so underpredicts the expected harvest age. This result is entirely consistent with the usual literature on optimal timber harvesting that indicates the expected harvest age falls with an increase in the discount rate (decrease in the discount factor).

---

[5] It is worth emphasizing that in both the true model and Model 5, the odds of a timber stand reaching age 10 is extremely low. This is apparent from Figure 2b, which shows that for η_ε =2.0, the probability of harvest in each of the first ten years is roughly .4, and Figure 4, which shows that in the first ten years the probability of harvest in Model 5 is virtually the same as in the true model.

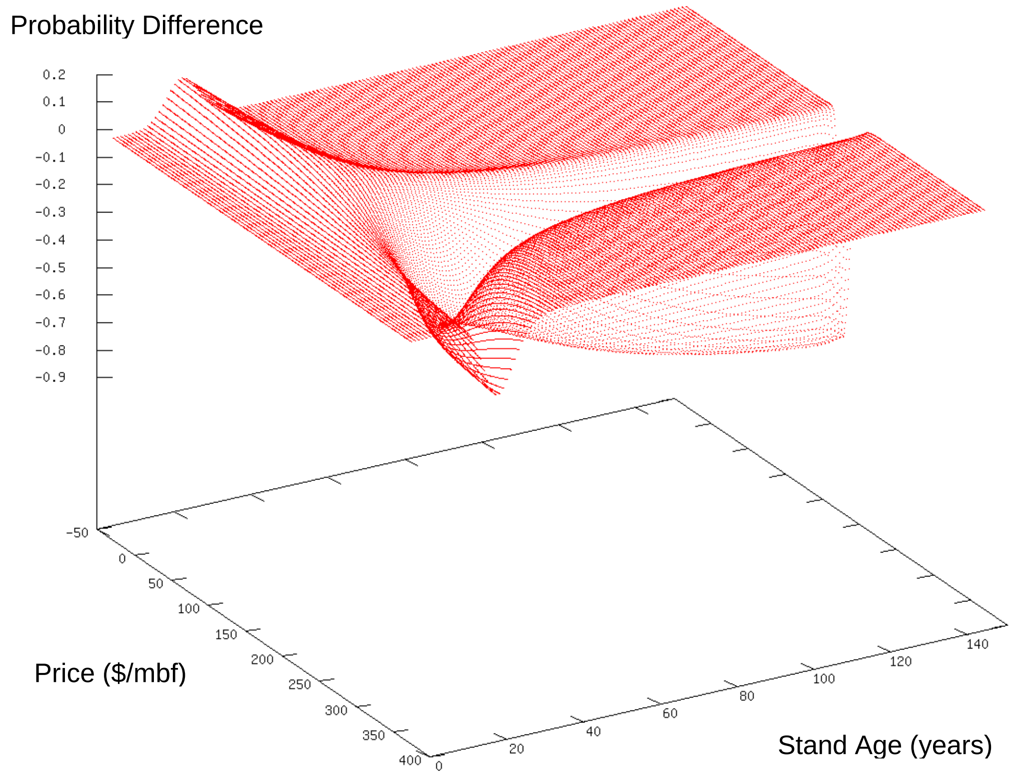

Figure 4. Harvest Probability Difference between the True Model and Model 5 ($Pr_{True} - Pr_{Model\ 5}$)

It should be emphasized that the low estimated discount rate in Model 5 is not a fluke associated with the particular starting values used in estimation. In a search across a range of starting values, including the parameters of the true model, we could not find estimates generating a higher likelihood value for the sample than those presented in Model 5. It should also be emphasized that the deep probability trough observed in Figure 4 occurs *outside* the range of the simulated data. Of the 8621 harvests observed in the data, 8561 (99.3%) occur at stand age 10 or less, and all harvests occur at a stand age less than 20. In other words, for very few observations is the stand age greater than 10, and for no observations is it greater than 20. Perusal of Figure 4 shows, then, that *in the range of the data* the harvest probability of Model 5 is virtually the same as for the true model. This teaches an old and familiar lesson, one certainly not unique to the estimation of DDP's: in the absence of a good range in the data, identification of a model can be difficult.

IV. Comments on Reduced Form (Static) Estimation of Discrete Dynamic Decision Problems

Given the difficulty of structural estimation of a discrete DDP, it is reasonable to question whether the effort is worth the gain. In particular, why not simply estimate a reduced-form version of the problem? Inspection of the timber harvesting decision (1.16) makes clear that one can specify the decision problem as one in which the forest owner harvests trees if

$$F\left(p_t, a_t\right) + \Delta\varepsilon_t > 0\,, \tag{1.17}$$

where $\Delta\varepsilon_t$ is distributed logistically. This leads to a straightforward application of logistic maximum likelihood estimation. A similar reduced form can be used to approximate the optimal decision rule of *any* discrete DDP. Of course, the presence of the value functions in (1.16) argues for a flexible form in the approximation, and even then the quality of the approximation may be poor. Figure 5 presents harvest probabilities from first-, second-, and third-order estimation of (1.17) for the case where the true model is the LV model (σ_ε=2.0), and so the figure against which the panels of Figure 5 are to be judged is Figure 2b. Estimates are based on simulated data of the same size as used in Table 1 –namely, pooled time-series, cross-sectional data of length 80 years and width 500 forest stands. The structural counterpart in estimation is Model 4 (see Table 1), which generated harvest probabilities virtually identical to the true harvest probabilities (that is, it generates a probability surface that looks exactly like the probability surface of the true model presented in Figure 2b, so we do not bother presenting the surface here). A comparison of Figure 5a and Figure 2b indicates that when $F\left(\cdot\right)$ takes the simple linear form,

$$F = \alpha_0 + \alpha_1 p + \alpha_2 a\,,$$

the approximation to the harvest probability surface is very poor, and the probability of harvest *decreases* as the stand age increases. Figure 5b indicates that when $F\left(\cdot\right)$ takes the quadratic form,

$$F = \alpha_0 + \alpha_1 p + \alpha_2 a + \alpha_3 p^2 + \alpha_4 a^2 + \alpha_5 p \cdot a$$

the approximation is considerably improved, and Figure 5c indicates that when $F\left(\cdot\right)$ takes the cubic form

$$\begin{aligned} F = {} & \alpha_0 + \alpha_1 p + \alpha_2 a + \alpha_3 p^2 + \alpha_4 a^2 + \alpha_5 p \cdot a \\ & + \alpha_6 p^3 + \alpha_7 a^3 + \alpha_8 p^2 \cdot a + \alpha_9 p \cdot a^2 \end{aligned},$$

the approximation is excellent. Seemingly and not surprisingly, even a fairly low-order polynomial will do a good job of approximating fairly complex decision rules.

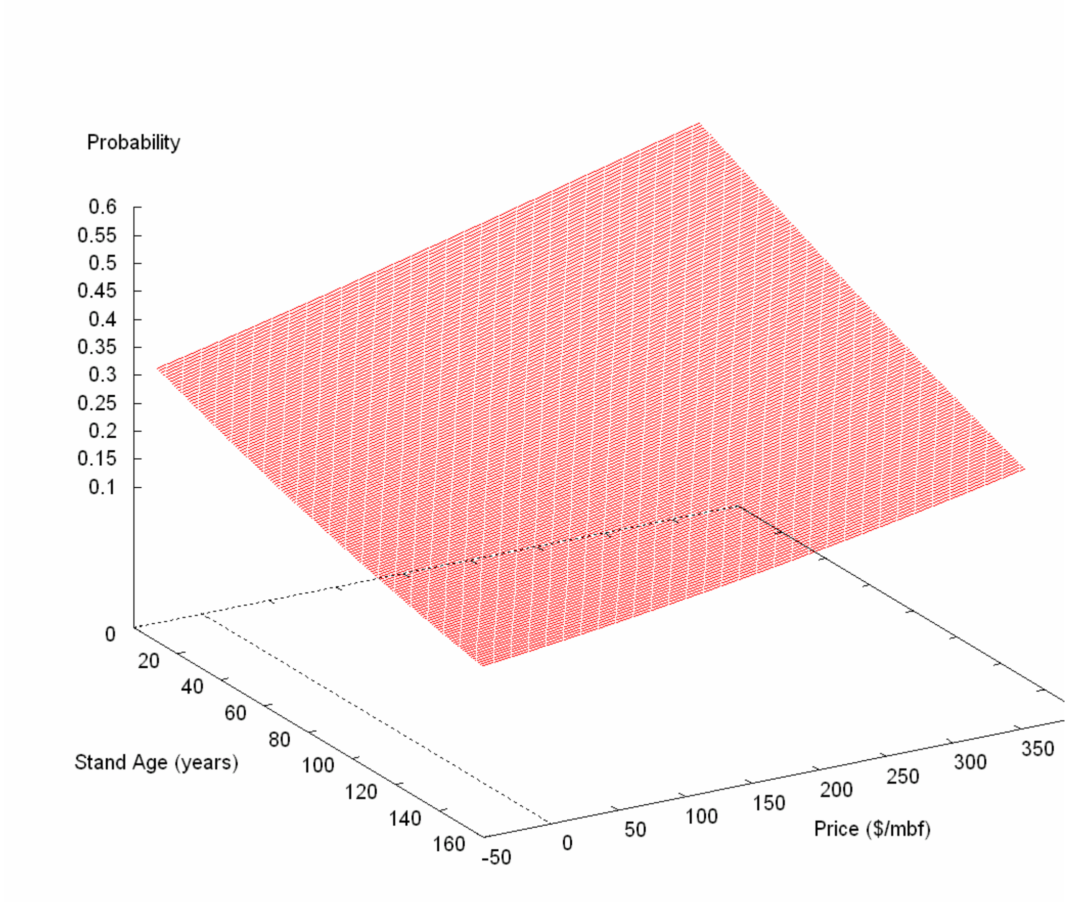

Figure 5a. Harvest Probability Surface for Linear Reduced-Form (Logit) Model

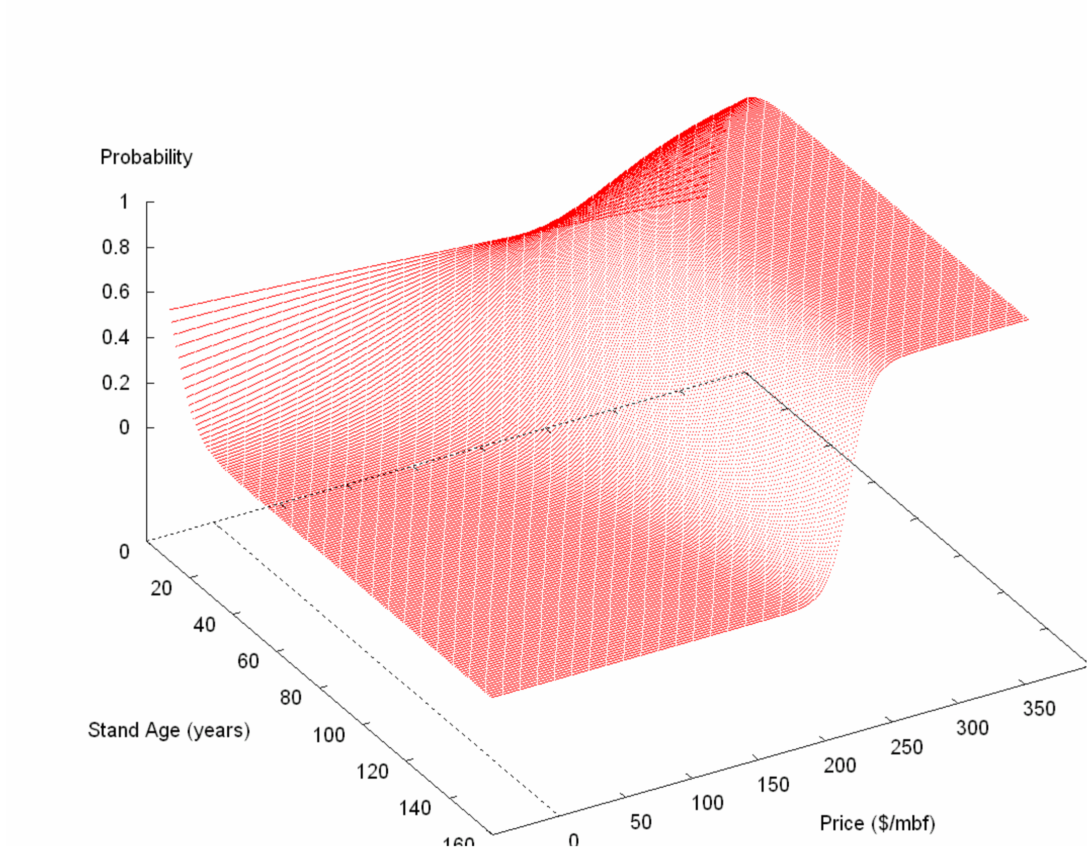

Figure 5b. Harvest Probability Surface for Quadratic Reduced-Form (Logit) Model

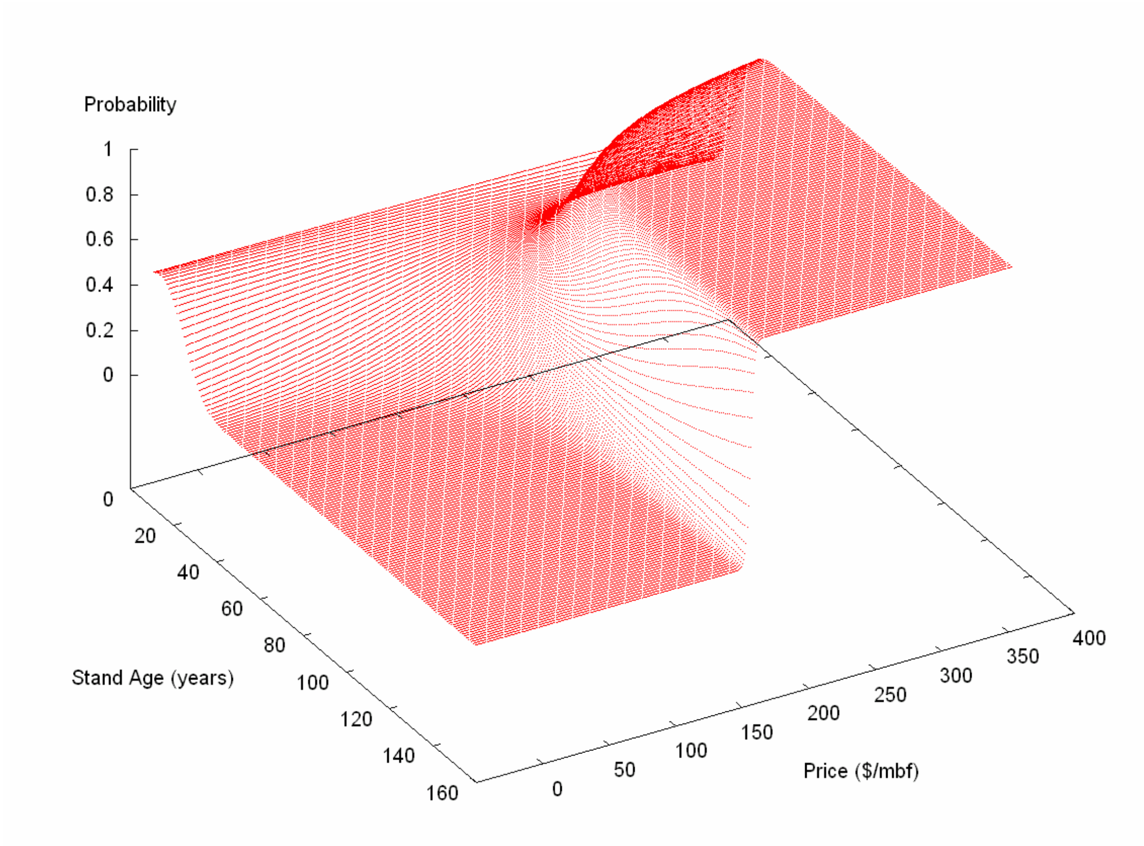

Figure 5c. Harvest Probability Surface for Cubic Reduced-Form (Logit) Model

The conceptual weakness of reduced form estimation is the same for DDPs as it is for any economic model; it does a reasonably good job of describing the effect of various state variables on decision variables, but the estimation is otherwise devoid of economic content. It identifies the variables affecting a dynamic decision, but it provides no insight about the mechanism of the relationship. In the illustrative example presented here, reduced form estimation tells the analyst that timber price and stand age do affect the harvest decision, but it is silent about the discount rate used in the harvest decision, an important issue in the long debate about whether forest owners harvest timber too soon (and therefore require incentives to take the long view); it says nothing about the nontimber benefits of forestland (as embodied in $\boldsymbol{\varepsilon}$), a major issue in the allocation of land across various uses; it says nothing about the overall value of forestland; and it says nothing about the price expectations of forest owners.[6]

A related issue is whether it is safe to assume that a decision problem is static. For our timber example, one might simply posit that in each period the forest owner harvests if the intraperiod utility from harvesting is greater than the utility of not harvesting. For

[6] In the structural models estimated in this chapter, price distribution parameters are fixed at their true values. In other estimations conducted in preparation of the chapter but not presented here, price distribution parameters were included in the set of parameters to be estimated. Generally results were excellent, though a bit sensitive to starting values. Provencher (1995a,1995b) estimates models of pulpwood harvesting in which the price process is autoregressive.

instance, letting $H\left(a_t\right)$ denote the money-metric utility derived from a standing forest at age a_t, and assuming that harvest occurs at the beginning of the period and $H\left(0\right)=0$, harvest occurs if

$$\begin{aligned} & p_t W\left(a_t\right)-c+\varepsilon_t^1>H\left(a_t\right)+\varepsilon_t^0 \\ & \Rightarrow p_t W\left(a_t\right)-c-H\left(a_t\right)>\Delta \varepsilon_t \end{aligned}, \tag{1.18}$$

which, under the same assumptions about the distribution of $\boldsymbol{\varepsilon}$ used above, is the basis of a simple logistic regression. However, this approach can be problematic. For many problems –timber harvesting surely among them –it is not reasonable to assume away dynamic behavior. For problems where dynamic behavior is an open question, things are a bit complicated. Baerenklau and Provencher (2004) examine the issue of whether recreational anglers allocate a “fishing budget” over the course of a season. We estimate both a structural dynamic model and a reduced form static model of trip taking behavior and find significantly different welfare estimates across models. More troubling, we also demonstrate theoretical inconsistencies and identification problems with the static model when behavior is truly dynamic.

Furthermore, although dynamic models encompass static ones, actually testing for static behavior is problematical for the simple reason that if one specifies a sufficiently flexible form for the static utility function, the static model will provide an excellent fit to the data. Put another way, a reduced form model can be relabeled as a structural, albeit static, model, and as already demonstrated, reduced-form models can provide an excellent fit to the data. In our timber example, a sufficiently flexible static form requires making $H\left(\cdot\right)$ a function of *p* as well as *a*, and this is difficult to justify. Such structural restrictions are ultimately necessary to test static vs. dynamic models.

For many problems there is little *a priori* knowledge about the intraperiod benefit (utility, profit) function, and so it is difficult to distinguish static from dynamic behavior. For such problems, out-of-sample forecasting may shed light on whether behavior is static or dynamic, but this is a quite expensive diagnostic. Baerenklau and Provencher (2004) conduct such a test and find that in general a dynamic model does a better job of predicting out-of-sample trip-taking behavior than does a static model.

V. The Future of Structural Estimation of Dynamic Decision Processes

For years, structural estimation of DDP’s simply was not practical because of the substantial barrier posed by the computational requirements of estimation. Rapid advances in computational speed in the mid-1990s reduced this barrier considerably, yet the literature estimating DDP’s is decidedly modest, and there is no evidence that it growing. This is partly because the econometric modeling is difficult to understand and implement; this chapter attempts to clarify the estimation method. Perhaps it also reflects two related objections.

First, estimation requires strong parametric assumptions about the decision problem generating the data. In the example considered in this chapter, we maintained the strong assumption that the forest owner knows the growth function of trees and, perhaps even more unlikely, that the forest owner knows the stochastic process generating timber prices. In the real world, such assumptions are complicated by the requirement that the analyst *correctly* specify the parametric form of various dynamic processes, where by the "correct" process we mean the process actually used by the decision-makers.

Second, the estimation maintains that agents are dynamic optimizers. In a working paper entitled, "Do People Behave According to Bellman's Principle of Optimality?", Rust (1994a) makes the important point that *any* data set of state variables and decision variables can be rationalized as the outcome of dynamically optimizing behavior. The issue becomes, then, whether the data generated can be rationalized by dynamically optimal behavior circumscribed by plausible parametric specifications of the intraperiod benefit function and dynamic processes. The issue returns, in other words, to a variation of the first objection: Is the parametric specification of the dynamic model plausible in some theoretical or practical sense?

Given that a flexible static model can fit a data set as well as a dynamic one, and that dynamic optimization *per se* does not imply testable restrictions (and thus is not refutable), how might an analyst who suspects that behavior is dynamic proceed? In our view, estimation of DDP's should be informed by agent self-reporting about expectations of future states and the relative importance of the future in current decisions. So, for instance, if agents respond in surveys that future states are unimportant to their current decisions, it would seem difficult to justify a dynamic model. If agents report that they believe the best guess of prices tomorrow is the current price, a random walk model of price expectations (or at least, a specification of prices that allows a test for a random walk) is warranted.

Yet even as economists are uneasy about the strong parametric assumptions to be made in the estimation of DDP's, there exists a longstanding uneasiness in the profession about using agent self-reports to aid in the estimation of dynamic behavior. Manski (2004) argues that with regards to agent expectations, this is not much justified. The author states,

> "Economists have long been hostile to subjective data. Caution is prudent, but hostility is not warranted. The empirical evidence cited in this article shows that, by and large, persons respond informatively to questions eliciting probabilistic expectations for personally significant events. We have learned enough for me to recommend, with some confidence, that economists should abandon their antipathy to measurement of expectations. The unattractive alternative to measurement is to make unsubstantiated assumptions" (pg. 42).

Although the role of self-reports in an analysis can be informal, giving the analyst a rough basis on which to choose a static vs. a dynamic model, to choose one parametric specification of dynamic processes over another, and so forth, a useful direction of future research is to incorporate survey data in the estimation problem itself, similar in principle to the way revealed- and stated-choice data are integrated in the estimation of random utility models. We anticipate that in the next ten years insights gained from experimental psychology, experimental economics, and survey research will provide the basis for richer and more accurate models of dynamic economic behavior.

Table 1. Estimation Results from Simulated Data (see text)[a]

	Model 1: VNV T={1…80}	*Model 2: VNV T={40…80}*	*Model 3: VNV T={60…80}*	*Model 4: LV T={1…80}*	*Model 5: LV T={40…80}*	*Model 6: LV T={60…80}*
Parameter [true value]						
β [0.97]	0.9708 (0.0030)	0.9703 (0.0023)	0.9686 (0.0018)	0.9697 (0.0036)	0.9313 (0.0366)	0.9702 (0.0502)
η_ε [20.0]	18.43 (1.12)	17.58 (0.51)	61.22 (2.46)	-	-	-
η_ε [2.0]	-	-	-	2.07 (0.14)	5.25 (0.21)	4.36 (0.40)
θ^1 [0.0]	-0.01365 (0.0186)	-0.05119 (0.0831)	0.1238 (0.0128)	0.0157 (0.0095)	0.0941 (0.0016)	0.0807 (0.0067)
Log Likelihood	-357.58	-226.84	-21.60	-26701.82	-14007.72	-7164.75

a VNV="Virtually No Variance" model, with η_ε =20.0 ; LV="Low Variance" model, with η_ε =2.0. T=[x,y] indicates that only sample observations between and including years x and y in the 80-year sequence of simulated prices are included in the estimation. Additional explanation of Table 1 is found in the text. Starting values for all estimations: β =0.95, η_ε =10.0, θ^1=0.20.

References

Baerenklau, Kenneth A., and Bill Provencher. "Static Modeling of Dynamic Recreation Behavior: Implications for Prediction and Welfare Estimation". Manuscript, University of California-Riverside, 2004.

Ben-Akiva, Mosha, and Steven R. Lerman. Discrete Choice Analysis: Theory and Application to Travel Demand. MIT Press, 1985.

Brazee, Richard, and Robert Mendelsohn. "Timber Harvesting with Fluctuating Prices", *Forest Science* 34, p. 359-372 (1988).

Howitt, Richard E., Arnaud Reynaud, Siwa Msangi, and Keith C. Knapp. "Calibrated Stochastic Dynamic Models for Resource Management". Paper presented at the World Congress of Environmental and Resource Economists, Monterey, CA. 23 p. (2002).

Manski, Charles F. "Measuring Expectations". *Econometrica* 72, 1329-1376 (2004).

Provencher, Bill. "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 (1995b).

Provencher, Bill. "An Investigation of the Harvest Decisions of Timber Firms in the Southeast United States", *Journal of Applied Econometrics* 10, S57-74 (1995a).

Provencher, Bill, and Richard C. Bishop. "An Estimable Dynamic Model of Recreation Behavior with an Application to Great Lakes Angling", *Journal of Environmental Economics and Management* 33, 107-127 (1997).

Rust, John. "Optimal Replacement of GMC Bus Engines: A Dynamic Model of Harold Zurcher", *Econometrica* 55, 999-1033 (1989).

Rust, John. "Do People Behave According to Bellman's Principle of Optimality?" Manuscript, University of Wisconsin-Madison (1994a).

Rust, John. "Structural Estimation of Markov Decision Processes". *(In) Handbook of Econometrics*, D. McFadden and R. Engle (eds.), volume 4, North-Holland (1994b).