

Estimating Intertemporal Preferences for Natural Resource Allocation

Richard E. Howitt^a, Siwa Msangi^a, Arnaud Reynaud^b, and Keith C. Knapp^c

Abstract

In this paper we show how the degree of risk aversion, discounting, and preference for intertemporal substitution for a natural resource manager can be structurally estimated within a recursive utility framework. We focus on the management of Oroville Reservoir, in Northern California, and test the data to see if they are more consistent with a recursive utility model specification than one with standard time-additive separability, and estimate the implied degree of risk aversion. The results show that the data on dam storage and releases are consistent with a risk-averse manager with recursive preferences, and that preferences are stationary over the observed period. The data also rejects time-additive separability, whether specified with or without risk-aversion, such as the standard CRRA utility model. The improvement in model fit when recursive preferences are used is notable.

Contact author

Richard Howitt
Department of Agricultural and Resource Economics,
University of California
Davis, Ca 95616
Phone 530 752 1521 fax 530 752 5614
Email howitt@primal.ucdavis.edu

^a Department of Agricultural and Resource Economics, University of California at Davis.

^b LERNA-INRA, University of Social Sciences Toulouse.

^c Department of Environmental Sciences, University of California at Riverside.

The authors gratefully acknowledge support from a risk management grant from USDA-ERS that made this study possible.

Introduction

Natural resource management problems are typically stochastic and dynamic in nature, by virtue of the characteristics of the underlying physical or biological processes that govern the evolution of the resource. This has been the reason for a number of empirical applications of Dynamic Programming within the natural resource literature, as chronicled by Williams (1989). However, the tendency of researchers to use risk-neutral expected net present value objective functions when modeling natural resource problems has caused policy-makers to be somewhat skeptical of the real-world relevance of resource economics analysis¹.

Given the uncertainty facing the decision-maker in each period of the planning horizon, due to the realization of stochastic shocks, we would expect risk-aversion and intertemporal substitution to feature prominently in the characterization of intertemporal preferences. While a number of authors have incorporated risk-aversion into analytical and numerical models in the economic literature (Knapp and Olson, 1996; Krautkramer et al., 1992), few have actually tried to estimate the degree to which it enters into the decision-makers objective criterion, and none of those papers consider natural resource management problems. Most of the resource literature has imposed severe restrictions on the preferences for intertemporal substitution by adopting a time-additive separable formulation of the objective function. We avoid this problem by using a recursive utility specification that is more general and allows preferences towards risk and intertemporal substitution to be decoupled (Epstein and Zin, 1989).

The estimation of dynamic preferences has been implemented in a number of settings to elicit the underlying parameters of the decision-maker's problem. Where adequate time-series data exists, it can be used to calculate empirical moments (Hansen and Singleton, 1982, 1983). Rust (1987) has applied alternative dynamic estimation techniques to analyze the actions of a single decision-maker. This approach has laid the foundation for several analyses of this type (Provencher and Bishop, 1997; Provencher, 1995; Miranda and Schnitkey, 1995). However in the few papers that apply dynamic

¹ This paper was motivated by a comment made at an agency workshop in response to the presentation of results from a conventional risk-neutral SDP solution. The commentator was Dr. Francis Cheung of the California Department of Water Resources. He pointed out that optimization models tend to be discounted by decision-makers because they ignore the presence of risk in the objective function. We gratefully acknowledge support from USDA ERS grant "Measuring, Costing and Mitigating Institutional Risk in Californian Irrigation Water Supplies"

estimation to natural resource problems none specifically estimate intertemporal preference parameters, or quantify risk aversion (Provencher, 1995, Fulton and Karp, 1989).

We seek to address this gap in the environmental and natural resource literature by applying dynamic estimation techniques to elicit the recursive intertemporal preferences with continuous state and control variables. We use the example of reservoir management, but take a different approach from authors in the water resources engineering literature who have only discussed the comparative dynamics of increasing risk-aversion in reservoir management (Kerr and Read, 1998; Craddock et al., 1998). We identify the degree of risk-aversion that is exhibited by the decision-maker's actions, and employ likelihood ratios to test whether the data is consistent with a risk-averse decision-maker, and whether the preferences of the decision-maker can be better characterized by a recursive utility function.

The outline of the paper is as follows. In the next section, we describe the general resource allocation problem and the recursive utility specification that we use. In the following section we discuss the specific empirical application of our problem to Oroville Reservoir and describe the dynamic estimation methodology. The succeeding section presents the estimation results, followed by a section of concluding remarks.

Resource Allocation and Recursive Utility

This section develops the specification for an intertemporal natural resource management problem for a stochastic flow resource with storage. The model is defined with continuous state and control variables, and uses reservoir management as an example. Three preference characteristics that are important for resource management are the discount rate that defines the relative weight of current period benefits to those accrued in the future; resistance to intertemporal substitution which embodies the decision-maker's preferences for consumption smoothing; and risk aversion which represents the agent's degree of dislike for variability in future returns. The vast majority of dynamic resource economic papers use expected net present value (EPV) as the objective, which only specifically incorporates the discount rate. In this paper, recursive utility is used to specifically estimate parameters for each of these characteristics.

Model Specification

Reservoir system dynamics are given by:

$$S_{t+1} = S_t + \tilde{e}_{1t} - w_t - sp_t \quad (1)$$

Where the parameters are: water release w_t , stochastic inflow \tilde{e}_{1t} , storage S_t , and spills from the reservoir, sp_t at each date t . The index t in (1) denotes the time period – in our example, a year. This implies that the change in the reservoir storage equals stochastic inflow net of water release w_t , and the spills. The spills balance the system in times of high flows, but have no economic value in the model.

A general characteristic of natural resource management is that decision-makers do not operate in a closed system. They have to take into account the uncertainty in the rest of the system. We assume that we can decouple management of the unit of natural resource being modeled from the rest of the network². There are two reasons to decouple a single state from the resource network. The first is the reduction in the dimensionality of the stochastic dynamic programming (SDP) problem and an increase in its empirical tractability. The second reason is that a central aim of our approach is to estimate the recursive utility parameters from a time series of observed decisions. In most resource systems, decisions are split between agencies or levels of agencies. An estimation of preference parameters must be focused on a single decision maker (or unit) who is cognizant of, but decoupled from, the rest of the system. An example of this is in Rust (1987) where he models the actions of a single individual – Mr. Harold Zurcher.

Inflows from the rest of the system are characterized by a stochastic inflow \tilde{e}_{2t} . Final demand for water may either be satisfied by water release w_t or by \tilde{e}_{2t} . We assume that exogenous stochastic variables, in the reservoir management example, water inflows $(\tilde{e}_1, \tilde{e}_2)$, are i.i.d over time³ on a compact space and subject to a common joint distribution $\Phi(\bullet)$. $\Phi_1(\bullet)$ and $\Phi_2(\bullet)$ respectively represent the marginal distribution of the reservoir inflow and of the rest-of-network flows.

² Notice that this decoupling is not a necessary requirement of the model but just a convenient assumption to simplify it. Our model could be easily extended to a multi-state variable problem.

³ This assumption is clearly difficult to justify on a daily or monthly basis. It is more likely to hold at the yearly basis used in this model, and in the absence of any long-term trend.

The timing of management information and controls is important. First, the decision-maker observes the stock of stored water S_t and the realization of the exogenous stochastic variable $\tilde{\epsilon}_{1t}$, in this example, the local stochastic inflow. Second, the decision-maker chooses the control w_t , the level of water released for beneficial consumption. This decision is a function of the future local stochastic inflow and the current stock of water in the reservoir. The natural resource available for consumption is, at each date, composed of the resource release and the realized rest-of-network inflow. Thus the value of the natural resource stock is a function of the stochastic flow in the rest of the network. We assume that the decision-maker cannot directly observe the rest-of-network inflow, but knows its distribution. Usually, resource networks are complex, and it may be the case that the decision-maker, for a given part of the system, is not aware of the state of the system in the rest of the network. This is especially true if different authorities (state versus federal level, private versus public) manage different parts of the water network, or if the network is managed on a large spatial scale. A direct consequence of this information structure is that a decoupled decision-maker, when computing the optimal release, should take into account the realized local inflow and the distribution of rest-of-network inflow that is conditional on this realized local inflow. We denote the distribution of the inflow to the rest of the network, conditional on the local inflow, by $\Phi_{2/1}(\bullet)$.

The objective function

Natural resource demand may either be satisfied by flows from the single decoupled system or by flows from the rest of the system. At each date, the consumption of resource flows is q_t defined as:

$$q_t = w_t + \tilde{\epsilon}_{2t} \quad (2)$$

Resource demand is defined by the inverse demand function $P(q)$. The net surplus, $W(q)$, derived from resource consumption is denoted by:

$$W(q) = \int_0^q P(u) du \quad (3)$$

Note that the net surplus of resource consumption is a concave and increasing function of q .

We use a recursive utility specification to represent decision-maker preferences. Koopmans (1960) presents, in a deterministic context, the first axiomatic presentation of recursive preferences. While Kreps

and Porteus (1978) generalized this structure to stochastic models, Epstein and Zin (1989) later developed an isoelastic formulation of Kreps and Porteus preferences. This formulation has been used in applications ranging from macroeconomic modeling (Weil,1990), to farm production behavior (Lence, 2000). In addition, Knapp and Olson (1996), Ha-Dong and Treich (2000) and Peltola and Knapp (2001) have used recursive specifications in resource management problems. Three main arguments are advanced in favor of utilizing this class of preferences in theoretical work. First, it encompasses a wide range of preferences (expected utility, Kreps and Porteus specification among others). Second, it enables a distinction to be drawn between risk and intertemporal substitution effects⁴. Third, this specification satisfies the properties of intertemporal consistency and stationarity of preferences. Following Epstein and Zin (1991), we use an isoelastic formulation of Kreps and Porteus preferences. Given a current net profit W_t resulting from natural resource use in period t , recursive utility is given by:

$$U_t = \left\{ (1-\beta) \cdot W_t^\rho + \beta \left(E[U_{t+1}^\alpha] \right)^{\frac{\rho}{\alpha}} \right\}^{\frac{1}{\rho}} \quad (4)$$

where $\beta \in [0, 1]$ is the subjective discount factor, $\beta = 1/(1+\delta)$, δ is the subjective rate of discount, $\alpha \in [< 1, \neq 0]$ is the risk-aversion parameter, and $\rho \in [< 1, \neq 0]$ the constant of resistance to intertemporal substitution, and $E[.]$ is the expectation operator. Given this specification, the elasticity of intertemporal substitution (EIS), σ , is equal to $1/(1-\rho)$, $\sigma \in [0, +\infty)$. It follows that a decrease of the intertemporal substitution resistance parameter, ρ , below 1 results in a lower intertemporal elasticity of substitution. Finally, note that recursive preferences nest expected utility as a special case, by setting $\alpha = \rho$, substituting in future values of U_{t+j} and cancelling the ρ and $1/\rho$ exponents, equation (4) simplifies to the intertemporal expected utility function (Epstein and Zin 1991). In addition, if $\alpha = \rho = 1$, we get the familiar expected present value (EPV) utility function.

In the next part of the paper, we estimate the decision-maker's risk aversion, discount factor, and resistance to intertemporal substitution. Three main reasons support the estimation of these parameters. First, there is no consensus in the economic literature on the level of the two recursive utility parameters.

⁴ Attitudes toward variations in consumption across states of the world can be characterized by risk aversion. Attitudes toward variations in consumption across time are represented by the degree of intertemporal substitutability. With the usual expected utility preferences (intertemporally additive and homogeneous von Neuman-Morgenstern utility index) these two notions cannot be identified separately. Recursive preferences allow risk attitudes to be disentangled from the degree of intertemporal substitutability.

Various authors have proposed estimates of the EIS that range from zero (Hall, 1988) all the way to 0.87 (Epstein and Zin, 1991), while estimates of the risk aversion coefficient $(1-\alpha)$ range from 0.82 (Epstein and Zin, 1991) to 1.5 (Normandin and Saint-Amour, 1998). Second, the impact of risk-related parameters on optimal policies is known to be important. Knapp and Olson (1996) show that increasing risk-aversion results in more conservative decision rules. In contrast, Ha-Duong and Treich (2000) show that larger risk aversion strengthens optimal pollution control. They also find that a larger resistance to intertemporal substitution rotates the optimal control path toward less pollution control in the current period and more control in the future. However, none of these studies actually estimate these parameters, which is what we do in this paper.

Our problem maximizes the manager's recursive utility subject to the equation of motion for the natural resource stock and the feasibility constraints and is as follows:

$$\begin{aligned} & \text{Max } U_1 & (5) \\ \text{s.t. } & \left\{ \begin{aligned} U_t &= \left\{ (1-\beta) \cdot E_{e_2} [W_t^\rho(q_t)] + \beta \left(E_{e_1} [U_{t+1}^\alpha] \right)^\frac{\rho}{\alpha} \right\}^\frac{1}{\rho} & (6a) \\ S_{t+1} &= S_t + \tilde{e}_{1t} - w_t - sp_t & (6b) \\ q_t &= w_t + \tilde{e}_{2t} & (6c) \\ S_{t+1} &\geq \underline{S} & (6d) \\ S_{t+1} &\leq \overline{S} & (6e) \\ w_t &\geq 0 & (6f) \end{aligned} \right. \end{aligned}$$

The stochastic control problem consists of choosing a sequence of decision rules for resource flows that maximize the objective function (5) subject to (6a)-(6f). The problem is solved over an infinite horizon, and at each date, the current net surplus depends on the resource allocation and the stochastic level of flows in the rest of the network. The current net return $(W(q_t))$ is recursively linked through equation (6a) to all future controls and the resulting states. Thus the objective function in (5) is defined as U_1 , and reflects all future utility functions through equation (6a). All model parameters and functions are the same for all decision stages, which assumes a stationarity of preferences that we will test for later in the paper. The stochastic dynamic recursive equation defining optimal natural resource management is:

$$V(S, \tilde{e}_1) = \text{Max}_w \left\{ (1 - \beta) \cdot \int W_t^\rho (w + e_2) d\Phi_{2/1} + \beta \left[\int V^\alpha (S, \tilde{e}) d\Phi_1 \right]^\frac{\rho}{\alpha} \right\}^\frac{1}{\rho} \quad (7)$$

where $V(\cdot)$ is the value function representing the maximized value of recursive utility and w is the feasible allocation of water. We now have a standard SDP problem that we can solve by recursive solution methods, standard to the dynamic programming literature. The value iteration method that is used to solve (7), subject to 6(a)-6(e), consists of assigning an initial value for the value function, and then recursively solving the maximization problem until the implied carry-over value function converges to an invariant approximation (Bertsekas, 1976). In implementing this fixed-point procedure, we employ an orthogonal polynomial approximation to the value function, for computational efficiency and to accommodate our continuous state variable specification (Judd, 1998). This type of functional approximation has also been proposed by Miranda and Fackler (1999, 2002) for the solution of continuous-state dynamic programming problems.

An Empirical Application to Oroville reservoir

Oroville Reservoir is located on the Feather River in Northern California. The California Department of Water Resources operates this reservoir within the State Water Project. Water releases from Oroville reservoir are used for electrical power generation, irrigated agriculture and to satisfy domestic and industrial user demands. Oroville also provides flood control and enhancement of sport fisheries and wildlife habitat in the Delta area. Most of the hydrologic data used comes from the 'State Water Project Annual Report of Operations' published each year by the California Department of Water Resources from 1974 to 1996.

Specification of the problem

We consider the optimal annual use of Oroville reservoir and limit our analysis to the inter-year management problem.

We assume that yearly inflows $(\tilde{e}_1, \tilde{e}_2)$ are i.i.d over time with a Gaussian joint distribution:

$$\begin{bmatrix} \tilde{e}_1 \\ \tilde{e}_2 \end{bmatrix} \sim N\left(\begin{bmatrix} \mu_1 \\ \mu_2 \end{bmatrix}, \begin{bmatrix} \sigma_1^2 & \sigma_{12} \\ \sigma_{12} & \sigma_2^2 \end{bmatrix}\right). \quad (8)$$

It follows that the marginal distributions $\Phi_i(\bullet)$, $i = 1, 2$, are defined by:

$$\tilde{e}_i \sim N(\mu_i, \sigma_i^2) \quad (9)$$

and the distribution of the rest-of-network⁵ inflow conditional on the reservoir inflow, $\Phi_{2/1}(\bullet)$, by:

$$\tilde{e}_2 | e_1 \sim N\left(\mu_2 + \frac{\sigma_{12}}{\sigma_1^2}(e_1 - \mu_1), \sigma_2^2 - \frac{\sigma_{12}^2}{\sigma_1^2}\right). \quad (10)$$

The joint distribution of inflows is estimated by maximum likelihood using GAUSS. The estimate is based on nineteen years of observed flows into Oroville and the rest of the network. Inflow parameter estimates are presented in Table 1, below.

Table 1: Estimate of inflow distribution

Parameter	Estimates	Standard Error	Student t
μ_1	3.7957	0.6009	6.317
μ_2	15.7583	2.2742	6.929
σ_1^2	6.8594	2.2257	3.082
σ_2^2	98.2635	31.8850	3.082
σ_{12}	24.1569	8.1346	2.970

From Table 1, the marginal distribution of Lake Oroville inflow is given by:

$$\tilde{e}_1 \sim N(3.7957, 6.8594) \quad (11)$$

and $\Phi_{2/1}(\bullet)$, the distribution of the rest-of-network inflow is conditioned on the reservoir inflow by:

$$\tilde{e}_2 | e_1 \sim N(2.3910 + 3.5217 \cdot e_1, 13.1896). \quad (12)$$

⁵ “Rest-of-network” inflow are the total Sacramento-San Joaquin Delta inflow from the Sacramento Basin, less the water effectively released by Oroville dam. These “rest-of-network” inflows (or Delta inflows) consist of those flows from other sources outside Oroville going from the North into the Delta region. This information comes directly from the State Water Project Annual Reports of Operations 1974-96 published by the Department of Water Resources of the State of California, California Department of Water Resources (1974-96).

The reservoir inflow and rest-of-network conditional inflow distributions are represented by eight discrete intervals evenly spaced around the mean value, with an associated probability for the inflow falling within each interval and a mean inflow value for each interval.

The demand function

As previously mentioned, the demand for water is represented by an aggregate inverse demand function. The inverse demand function was adopted from the one used in the CALVIN⁶ model. CALVIN is run for a seventy two year hydrologic sequence and reflects the current level of development of the water system. The inverse demand is computed using annual inflows to the Delta from the Sacramento basin, and the implied scarcity values associated with them. A quadratic form is fitted to the data generated by CALVIN. The resulting inverse demand function is:

$$P(q) = 150 - 2.9 \cdot q + 0.02 \cdot q^2 \quad (13)$$

where q is the quantity of water in millions of acre-feet (MAF) and $P(.)$ is the associated marginal value in dollars per acre-feet. When water quantity varies from 10 MAF to 40 MAF, the resulting demand price per acre-feet varies from \$123 to \$66, an acceptable price range for California.

The resulting net benefit function from water consumption may be written as:

$$W(q) = 150 \cdot q - 1.45 \cdot q^2 + 0.0067 \cdot q^3 \quad (14)$$

which is increasing and concave in water consumption for q within the relevant ranges of value.

Reservoir Spills

Optimal management of a reservoir aims to minimize the occurrences of both shortages and spills. By keeping a high storage level of water from year-to-year, the decision-maker can smooth water consumption over dry years. However, keeping a high level of water storage increases the probability of important spills in the case of a wet year. Optimal reservoir management must tradeoff between these two effects. We assume that the spill during year t , defined as sp_t in equation (1), is a function of the realized inflow during this period (\tilde{e}_t) and the available storage capacity at the beginning of the period (cap_t). The available storage capacity in time “ t ” is defined as the difference between the maximum storage

⁶ CALVIN is an economically-driven optimization model of California’s statewide inter-tied surface and groundwater system, Jenkins *et al* (2001). CALVIN optimizes the operations of system resources over a given hydrologic sequence to maximize statewide net willingness-to-pay of urban consumers and agricultural producers for additional water.

capacity of the reservoir (\bar{S}) and the storage at the beginning of the year (S_t). Different functional forms were tested in the estimation of this relationship. The one giving the best fit for the realized spills is:

$$sp_t(\tilde{e}_{1t}, cap_t) = 0.095382 \cdot \tilde{e}_{1t} + 0.005024 \cdot \tilde{e}_{1t}^2 + 0.000993 \cdot \tilde{e}_{1t}^3 - 0.02305 \cdot \tilde{e}_{1t} \cdot cap_t \quad (15)$$

with an adjusted R-square of 0.657. Spill is an increasing function of inflow and decreasing in the available storage capacity. However, the greater the inflows, the more important storage capacity becomes in reducing spills. Finally, we assume that decision-maker knows the relationship in equation (15) that links spills, inflows and storage capacity.

Estimating the Intertemporal Preference Parameters

Given flood control constraints, the maximum storage capacity in Lake Oroville is determined on January first of each year and is 2.861 million acre-feet (MAF). We assume a minimum storage constraint equal to 0.987 MAF. This value corresponds to the minimum storage observed from 1974 to 1996. The model assumes that decision-makers maximize their utility subject to the equation of motion for the reservoir stock and the feasibility constraints. The stochastic dynamic optimization program is defined by equations (5), (6a) – (6f), and (14), where the spill function is given by equation (15).

We use a dynamic estimation approach to estimate the primitive parameters of the decision-maker's objective function, in a similar vein to that of Fulton and Karp (1989). However, unlike these authors, we do not restrict ourselves to the linear-quadratic case that is usually adopted for computational ease, at the expense of imposing severe behavioral restrictions on the model. Like Fafchamps (1993) and Deaton and Laroque (1996), we allow our decision and state variables to be continuous, in contrast to the majority of the structural estimation literature (Rust, 1989; Provencher, 1995; Miranda and Schnitkey, 1995; Wolpin, 1984, 1985; Keane and Wolpin, 1994, 1997; Erdem and Keane, 1996).

Several notable papers have addressed the problem of estimating the relevant parameters within a discrete choice dynamic programming problem, such as Keane and Wolpin's 1994 paper, where they address the computational difficulties associated with finding the relevant functions for both the discrete-choice and intertemporal optimization problems. Recent efforts to overcome these computational difficulties have been addressed by some authors (Geweke and Keane, 1995; Aguirreagabria and Mira, 2001; Imai et al., 2002) – however they only deal with the discrete-choice case. Since our decision

problem is a continuous one, we concentrate our efforts on developing a reliable method that can be handled within a standard software package.

The estimation procedure used to find the ‘best-fit’ model parameters, corresponds closely with the procedure described by Rust (1987) and Provencher (1995), who use an SDP optimization procedure nested within an outer ‘hill-climbing’ algorithm that perturbs the parameter values in a direction that maximizes the likelihood. This iterative procedure is comprised of three stages, which can be described as follows.

In the first stage of the parameter estimation a set of values is specified in parameter space, which is a combination of values for the three model parameters. For each set of parameter values the Chebychev polynomial values for the carry-over value function are found by the value-iteration method, and a twenty three year sequence is then simulated and compared with the actual sequence of storage and releases. In the second stage, the log likelihood value for the joint sequence of storage and releases is calculated by solving the following problem for each simulation

$$\text{Max}_{\sigma_1, \sigma_2, \sigma_{12}} L = -n \cdot \text{Ln}(2\pi\sigma_1\sigma_2\sqrt{1-\sigma_{12}}) - \frac{1}{2(1-\sigma_{12}^2)} \sum_{t=1}^n \left\{ \begin{array}{l} (\tilde{w}_t - w_t)^2 / \sigma_1^2 + (\tilde{s}_t - s_t)^2 / \sigma_2^2 \\ - \frac{2\sigma_{12}(\tilde{w}_t - w_t)(\tilde{s}_t - s_t)}{\sigma_1\sigma_2} \end{array} \right\} \quad (16)$$

subject to $\sigma_1, \sigma_2 > 0, -1 < \sigma_{12} < 1$

Where \tilde{w}_t and \tilde{s}_t are the observed water releases and storage, w_t and s_t are the calculated water releases and storage, and σ_1^2 , σ_2^2 , and σ_{12} are the unknown variances and covariances for releases and storage.

The third and final stage employs a search procedure that perturbs the parameter values in a direction that will maximize the log likelihood values from (16). We used the Nelder-Mead (Nelder and Mead, 1965) search algorithm, since it requires neither derivatives nor concavity of the log-likelihood function in the parameters. The initial iteration of the Nelder-Mead (NM) search requires the likelihood function for three sets of parameters. It follows that each of these parameter sets requires the solution of the stochastic dynamic program (SDP) and maximum likelihood problems for each iteration of the search algorithm. This computational burden points to the need for a solution method that is both rapid and stable. Further details of how the Nelder-Mead algorithm works are given by Dennis and Woods (1985).

Once a new set of parameter values is obtained, then the procedure returns to stage 1 and repeats iteratively until convergence is determined by a ‘stop’ criterion.

We found that solving simultaneously for the three unknown parameters in the recursive utility function led to instability in the algorithm and a failure to converge in likelihood values. Later in the paper it will be seen that the likelihood function is relatively unresponsive (flat) with respect to the discount rate and that this is due to a low elasticity of intertemporal substitution. Accordingly, we initially solve for the risk aversion and intertemporal substitution parameters conditional on a 5% discount rate, and then later we report sensitivity of parameter estimation to changes in the discount rate.

The SDP solution and Value Iteration process

The state variable (reservoir storage) is discretized in eight points from 0.987 MAF to 2.861 MAF. We use a 6th-order Chebyshev orthogonal polynomial approximation of the value function⁷:

$$V_c(S) = \sum_{i=0}^5 a_i \cdot T_i(\hat{S}), \text{ where } \hat{S} = \mathfrak{M}(S). \quad (17)$$

The Chebyshev polynomial coefficients a_i $i = 1, \dots, 5$ are iteratively computed using the Chebyshev regression algorithm, and $\mathfrak{M}(S)$ is a mapping of S onto the $[-1, -1]$ interval, Judd (1998). For each possible value of the discount factor and intertemporal substitution preferences of decision-maker, the SDP program is first solved with some initial values for Chebyshev polynomial coefficients. The resulting SDP solution allows us to compute new a_i values. If the resulting coefficients differ from those in the previous step, the SDP is re-solved with new Chebyshev coefficients. The SDP program ends once quasi-stabilization of a_i 's is achieved. For details of the solution method and its implementation using GAMS, see Howitt et al. (2002).

The NM search procedure continued until the likelihood values for the selected simplex of parameter sets converged to within 0.4% difference. The starting point of the algorithm was also perturbed sufficiently to ensure robustness of the estimation results. In addition, we calculated the residual difference in the Bellman equation between the optimized solution value and the Chebychev value, suggested by Judd(1998) and Miranda & Fackler(2002) as a test of the Chebychev fit. The results show

⁷ Provencher and Bishop (1997), in a different context, also use such a polynomial approximation to the value function. They nest the dynamic programming approach within a maximum likelihood procedure.

that the residuals did not decrease when the number of Chebychev polynomial values was increased, supporting the conclusion that the 6th order approximation used is sufficient⁸.

Results

The nested SDP and likelihood problems were run until convergence in the parameter estimates was achieved. To improve the numerical stability of the search procedure, the value of the discount factor, β , was set to 0.95, as mentioned previously. The resulting parameter estimates are shown in Table 2 below, along with their calculated standard errors, which were bootstrapped with 500 repetitions. While it is standard to use 1000 bootstrap repetitions (Efron and Tibshirani, 1993), we found the computational time to be excessive, and consider these to be upper-bound estimates of the standard errors. Nonetheless, our estimates are still significant at the 95% level, even with these standard error estimates, so we are confident as to the robustness of our estimates. Efforts to recover a consistent estimate of the Fisher Information Matrix through approximation of the Hessian of the likelihood function (as described in Fafchamps, 1993), resulted in unreliable estimates of the variance-covariance matrix⁹.

Table 2: Parameter Estimates for Recursive Utility SDP Model.

Parameter	Estimated Value	Standard Error	EIS value $1/(1-\rho)$	Coeff. of Risk Aversion $(1-\alpha)$
ρ	-9.000	4.60	0.100	
α	-0.440	0.23		1.440
Log Likelihood	-10.257			

These parameters were calculated with a fixed discount rate of $\beta = 0.95$. Standard errors are based on 500 bootstrap repetitions

⁸ The maximum absolute residual value is 0.000725 and the mean absolute value over the 20 grid points is 0.000225. Given a mean value for the Bellman equation of 1.08 these residuals are acceptably small.

⁹ The mean gradient over 23 years for changes in the rho (substitution parameter) value is -0.0425, and for the alpha (risk aversion) parameter is - 0.0152. The annual gradients vary in sign, but given the optimum likelihood value of -10.5, these mean gradients seem acceptable. The Hessian calculated by differences was not negative definite. To explore why the Nelder routine stopped with this Hessian value, we calculated the likelihood values over a grid of 225 points. The resulting plot showed strong curvature on the intertemporal substitution parameter, but a relatively flat likelihood surface for changes in the risk aversion parameter. This explains why the solution is optimal despite the Hessian values, and also confirms the value of the derivative free Nelder-Mead approach in this empirical example.

Because of the few empirical examples of recursive utility in the agricultural and resource literature, in the discussion of the estimates for the risk and time-preference parameters, we refer to two main fields in the economics literature that make use of recursive preferences, namely, macroeconomics and financial economics. While our estimate of the elasticity of intertemporal substitution (EIS) is low, it is nevertheless, compatible with the results within the macroeconomic literature based on aggregate data. Hall (1998) concludes, for example, that most of the studies “*support the strong conclusion that the elasticity is unlikely to be much above 0.1, and may well be zero*”. Although this conclusion has been recently challenged by empirical studies based on micro-data – see, for example, Atkeson and Ogaki (1995) – it seems that there is still consensus among macroeconomists that the intertemporal elasticity of consumption is very low. A low estimate of the decision-maker’s EIS means that she is relatively insensitive to the discount rate, and that the indifference contours that map between income in consecutive periods have strong curvature. Following the intuition provided by Lence (2000), we can interpret the low EIS value as meaning that smoothing the income benefits from water releases over time appears to be an important objective of the decision-maker that we observe. One explanation might lie in the fact that we have annual data and are focusing on the year-to-year management of the reservoir, whereas an estimation based on monthly data on water releases and storage could result in a different value of EIS.

Estimates of the risk aversion parameter have been discussed less in the macroeconomics and finance literature, and range widely. Estimates of the coefficient of relative risk aversion range from around one¹⁰ to as high as 18 in Obstfeld (1994). Our estimate of 1.441 is well within the range of admissible values.

¹⁰ Epstein and Zin (1991) report a value around 1, for example.

Sensitivity of the recursive utility parameter estimates

Since the recursive utility parameter estimates are conditional on the specified β values, we re-estimated the parameters under a range of discount rates to test for the sensitivity of the results to changes in the discount rate. Table 3 shows the changes in likelihood values and parameter estimates for discount rates that vary from 3% to 15% either side of the 5% rate used in the above results. The results show that changing β from 0.97 to 0.85 – which is equivalent to changing the discount rate by 80% (from 3% to 15%) – resulted in only a 0.85% increase in the log likelihood value, and a corresponding 9.3% decrease in the estimated value of ρ and a 10.9% decrease in the value of α . The results demonstrate that the likelihood surface that maps onto the parameter values of ρ , α and β is relatively flat with respect to β and, therefore, that the optimal parameter estimates for ρ and α are relatively insensitive to the prior value specified for β . This is consistent with a low estimate for the elasticity of intertemporal substitution (EIS). A low EIS means, in fact, that the decision maker is relatively insensitive to the interest rate, and consequently to β .

The following intuitive example illustrates the relationship between the EIS value, the discount rate and price differences. Suppose we think of the following two period problem:

$$\begin{aligned} \text{Max} \quad & u(c_1) + \beta u(c_2) \\ \text{subject to} \quad & c_1 + pc_2 = 1 \end{aligned}$$

where the price of c_1 is normalized to 1. The first order conditions result in.

$$\frac{\partial u(c_1)}{\partial c_1} = \beta \left(\frac{\partial u(c_2)}{\partial c_2} \right) \left(\frac{1}{p} \right)$$

Now suppose that the EIS is low, implying that the utility functions are strongly curved at the optimal consumption level. Then a change in the price p will only induce a small change in consumption since

this is all that is needed to adjust marginal utilities to satisfy the first order conditions. This illustrates how substitution is driven by curvature and not the discount factor β , at least over plausible discount rates.

This also explains why the discount factor doesn't play a strong role in the estimation procedure, as normally the discount factor has a large impact on intertemporal allocation behavior. The likely answer is, again, provided by the example above. For exactly the same reasons as before, if the utility functions are strongly curved (low EIS), then a change in the discount factor will have a minor effect on the consumption levels, hence the estimation likelihood surface will be relatively flat in that area. Thus a low EIS renders the discount factor less important, resulting in the flat response surface with respect to the discount factor and the observed instability in the 3-parameter estimation algorithm.

Table 3: Sensitivity of Recursive Estimates to Discount Rate.

Discount rate	β	Log Likelihood	ρ	EIS, $1/(1-\rho)$	$1-\alpha$
3%	0.97	-10.301	-9.5	0.095	1.475
5%	0.95	-10.257	-9	0.100	1.44
10%	0.91	-10.225	-9.426	0.096	1.446
15%	0.85	-10.214	-10.383	0.0879	1.527

Test for Stationarity of Preferences

We also tested for stationarity of the recursive preferences over the 23-year data series by estimating the parameters over first and last 7 years of observed data and testing for differences between these estimates and those estimated from the full data set. By testing over the two subsamples comprised of the first and

last third of the full sample we will detect any systematic change over time in the controller's preferences.

Table 4, below, shows the results.

Table 4: Test of the Stationarity of Preferences

	1974 – 1980	1990 - 1996	1974 - 1996
ρ estimate	-8.375	-3.6348	- 9.009
α estimate	-0.4470	-0.5122	-0.441
Likelihood value	-10.2690	-10.8801	-10.257
Regression Sum Squares	3.32	3.57	9.77
F Test Statistic			2.09 (df (2,10))

From these results, we see that the null hypothesis, stating that the parameter values estimated from the two sub-periods are equal, cannot be rejected by either an F-test or a likelihood ratio test. So our previously stated assumption of stationary preferences over the planning horizon is validated.

Model Specification Tests

To further evaluate the fit of our model to the observed data, we tested three other objective function specifications against the recursive utility specification. It can be seen from equation (4) that by setting the value of both ρ and α equal to one, the recursive utility specification devolves to a expected present-value (EPV) specification, while just setting α equal to one gives us a risk-neutral recursive model (RNR). The RNR model is fit to the data by fixing the β value to the same 0.95 value as in the Recursive Risk model, setting $\alpha = 1$, and estimating the ρ parameter – which was found to be $- 8.875$ (giving an

EIS value of 0.1013). We found the estimation of the EPV model to give a very implausible value of β , so we also fixed it to the same 0.95 value, and compared the simulated results with the other models.

Another alternative objective function is the familiar Constant Relative Risk Aversion (CRRA) utility function, shown in equation (18) below using the same definition of the benefit from reservoir releases as equation (3):

$$U_t = \frac{W_t^{(1-\alpha)}}{1-\alpha} + \beta E[U_{t+1}] \quad (18)$$

The unknown parameters within the CRRA specification are α , the level of risk aversion and β , the discount factor. α and β were estimated using the same procedure that is used to estimate the recursive utility parameters. The resulting optimal estimates are, $\beta = 0.975$ and $\alpha = 0.896$. We note that the estimated value for β is somewhat lower than the fixed value chosen for the estimation of the recursive utility model, but is not an unreasonable rate of discount (2.6%). The four models (with their respectively estimated parameters) are then compared in terms of their maximized likelihood values, as well as with respect to the mean squared error calculated from the resulting fit of the simulated storage and releases with observed data. The improvement of the MSE values for the Recursive Risk, Recursive Risk-Neutral and CRRA models over that of the EPV model, is also noted.

The restrictions implicit in the RNR, CRRA and EPV specifications are tested against the recursive specification with risk by using the maximized likelihood values for each model to calculate a likelihood ratio test statistic. Since, as noted earlier, the CRRA and EPV are nested within the recursive utility model, we use the simple likelihood ratio test statistic.

The likelihood ratio tests shown in Table 5, below, strongly reject both the CRRA and EPV specifications compared to the recursive specification with risk. We are unable to reject the RNR in favor of the recursive risk model, based on both the Likelihood ratio test and comparison of the MSE values.

However, the mean squared errors for both the storage and release values from the simulations of both the two Recursive and CRRA models show significant improvement over those from the EPV model, and speak highly in favor of the recursive utility model specification.


Table 5: Comparison of Alternative Objective Function Specifications

	Recursive Utility With Risk	Recursive Utility Risk-Neutral	Constant Relative Risk Aversion	Expected Net Present Value
Log Likelihood	-10.257	-10.204	-30.514	- 57.633
Likelihood ratio*		--	40.511	94.750
Storage MSE	0.1080	0.1079	0.6529	1.0851
% Improvement	91.63	91.64	49.42	--
Release MSE	0.3170	0.3151	0.6561	1.0851
% Improvement	70.79	70.96	39.54	--

(* The Likelihood ratio test statistic is compared to the critical level of the χ_1^2 statistic, which has a value of 6.63 at the 1% level)


The different model specifications result in different release (control) rules. The optimal release rules as a function of the current state variable (storage) are shown in Figure 1. As would be expected from the results in table 5, the release rules can be classified in two groups, those with a recursive preference structure and those with a separable structure. Figure 1 shows that the release rule under the recursive structure is lower for all storage levels, and at lower stock levels it responds less to changes in the storage. Both these characteristics will lead to a smoother release pattern. In addition, when risk aversion is added to either preference structure it only makes a slight difference to the optimal release rule.

Figure 1: Release Rules Under Alternative Preference Specification.



Further evidence of the superiority of the recursive utility model with risk over the EPV and CRRA models is shown in figures 2 to 5, which plot the observed reservoir storage and releases alongside the simulated storage and release values from the three models. The results from the RNR model were omitted from these figures for clarity, as they would overlap very closely with those from the recursive model with risk. Note that, consistent with the low EIS, the decision rule with a recursive specification is more conservative in its release behavior.


Figure 2: Simulated and Actual Storage Quantities- EPV Preferences



In comparing figures 2 – 4, and the releases in figure 5, it should be emphasized that the only difference in the simulations is in the specification of the intertemporal preferences and how the carry-over value function is nested within its structure. All other parameters, namely the inflows, spill function, probabilities, Chebychev collocation nodes, and the net benefit function are identical. In addition, each simulation is the outcome of the sequential annual optimization of the expected Bellman equation conditional on the current storage and capacity constraints. In this solution, the model only has the information available to the real decision maker, namely the preferences for current releases and future storage, the inflow probabilities, and the current state of the system being managed. Since the average releases are 165% of the average storage in this example, relatively small changes in the release policy will make substantial changes in the resulting storage.

Figure 2 shows that the model with EPV preferences cannot reproduce the actual storage decisions of the operators. With few exceptions, this preference structure results in storage levels that jump between the upper and lower bounds of reservoir capacity.

Figure 3: Simulated and Actual Storage Quantities- CRRA Preferences




The CRRA preference results in figure 3 show some improvement over the EPV results that is also reflected in the lower storage MSE value in table 5. However the model misses many of the annual turning points and is still subject to swings from the upper to lower bounds on storage.

In contrast to the poor EPV and CRRA storage results, the results for the recursive preference specification shown in Figure 4 are quite remarkable in the way that the annual dynamic optimization model catches nearly all the turning points of the operator. The only significant turning point that is missed is that which occurs during the major drought of 1991-1992. These results reflect the major point

of this paper, namely that dynamic behavior seems to be best represented by a dynamic, path dependent, preference specification.

Figure 4: Simulated and Actual Storage Quantities- Recursive Preferences




In contrast, the simulation results from the standard risk-neutral EPV model support the misgivings expressed by public decision makers towards models that ignore their intertemporal preferences – namely, that they do not reasonably represent the behavior of the true decision maker operating the reservoir.

Figure 5 shows the optimal releases under the four preference specifications. The release results show more similarity than the storage results because of their greater volume. However, years such as 1986 show clearly that there are significant differences between releases. The differences are reflected in the release MSE values in table 5. There is a clear reduction in MSE between the EPV solution and the

CRRA, but the improvement is even more striking when comparing the recursive specifications with the CRRA results.

Figure 5: Simulated and Actual Release Quantities



The close similarity of the Recursive models suggests that the recursive specification improves the fit much more than the addition of risk to the model. In a non-recursive setting, the addition of risk causes a rather large improvement in model fit, as seen from the Likelihood ratio test statistics and the improvement in MSE values. However this improvement is diminished once one has already incorporated

the intertemporal substitution into the behavioral model, which may suggest that the consequences of omitting risk from policy models may not be as severe, in terms of explanatory and predictive ability, once one accounts for intertemporal substitution.

Conclusion

In this paper, we demonstrate how the underlying behavioral parameters of a risk-averse natural resource manager can be dynamically estimated within the more general theoretical framework of recursive-utility preferences. We used the example of the operators of Oroville reservoir in California and were able to reject both the EPV and CRRA utility specifications in favor of a recursive specification with preferences for both risk and intertemporal substitution. The fit of the estimated recursive-utility model was much closer to the observed data on dam storage and releases, than that of either the EPV or CRRA models. We also demonstrate that these preferences remain stationary over the period we observe. However, when we compare the fit of our recursive model with one that imposes risk-neutrality (while still allowing for resistance to intertemporal substitution), we see that there is negligible improvement in fit, which suggests that there may be little gained by adding risk to a recursive model specification.

We have argued that policy-makers are reluctant to accept the results of policy models that ignore the importance of risk and intertemporal preferences to the decision-making process, especially when dealing with the management of important public utilities. However, there may be less reason to be concerned with risk-neutral models once recursivity of intertemporal preferences is taken into account, and may even suggest that reservoir managers are more concerned with their inability to trade-off public benefits over time, than with the risk they face. We note in passing that the preference specification that

performs the worst in this example, namely expected present value (EPV), is the standard specification for the majority of economic resource allocation models. With the improved empirical methods available, it now seems reasonable to hold dynamic resource and environmental models to the same standards as static demand and supply models, and use formal measures to test alternative preference specifications. The severity of imposing the assumption of time-additive separability in inter-temporal preferences is clearly demonstrated by our results.

Using our framework, the behavioral parameters of any dynamic decision process that can be modeled with continuous state and decision variables, can be elicited, without resorting to the restrictive behavioral assumptions of the linear-quadratic or time-additive separable utility frameworks. We find that a more generalized approach to the representation of intertemporal preferences offers a richer theoretical framework and more precise prediction of management behavior under uncertainty.

References

- Aguirreagabiria, V. and P. Mira (2001):** “Swapping the Nested Fixed Point Algorithm: A Class of Estimators for the Discrete Markov Decision Models.” *mimeo*.
- Atkeson, A., and M. Ogaki (1995):** “Wealth-varying Intertemporal Elasticities of Substitution: Evidence from panel Data and Aggregate data.” *Journal of Monetary Economics*, 38, 507-534.
- Bertsekas, D.P. (1976):** “Dynamic Programming and Stochastic Control.” Academic Press, New York.
- California Department of Water Resources (1974-96):** “State Water Project Annual Report of Operations.” edited by the State of California.
- Craddock, M., A.D. Shaw and B. Graydon (1998):** “Risk-Averse Reservoir Management in a Deregulated Electricity Market.” In *Proceedings of the 33rd Annual Conference of the Operational Research Society of New Zealand*, Auckland, New Zealand, 157-166.
- Deaton, A., and G. Laroque (1996):** “Competitive Storage and Commodity Price Dynamics.” *Journal of Political Economy*, 104(5): 896-923.

- Dennis, J.E., and D.J. Woods (1985):** “Optimization on Micro-computers: The Nelder-Mead Simplex Algorithm.” *Technical Report 85-9*, Department of Mathematical Sciences, Rice University, Houston, Texas.
- Erdem, T., and M. Keane (1996):** “Decision Making Under Uncertainty: Capturing Dynamic Brand Choice Processes in Turbulent Consumer Goods Markets.” *Marketing Science*, 51(1):1-20.
- Efron, B., and R.J. Tibshirani (1993):** “An Introduction to the Bootstrap.” Chapman and Hall, New York.
- Epstein, L.G., and S.E. Zin (1989):** “Substitution, risk aversion and the temporal behavior of consumption and asset returns: A theoretical framework.” *Econometrica*, 57(July), 937-969.
- Epstein, L.G., and S.E. Zin (1991):** “Substitution, risk aversion and the temporal behavior of consumption and asset returns: An empirical analysis.” *Journal of Political Economy* 99(April), 263-286.
- Epstein, L.G. (1992):** “Behavior Under Risk: Recent Developments in Theory and Applications.” In *Advances in Economic Theory, Vol II*, edited by Jean-Jacques Laffont. Cambridge: Cambridge University Press, 1-63.
- Fafchamps, M. (1993):** “Sequential Labor Decisions Under Uncertainty: An Estimable Household Model of West-African Farmers.” *Econometrica*, 61(Sep): 1173-97.
- Fulton, M., and L. Karp (1989):** “Estimating the Objectives of a Public Firm in a Natural Resource Industry.” *Journal of Environmental Economics and Management*, 16, 268-287.
- Geweke, J., and M.P. Keane (1995):** “Bayesian Inference for Dynamic Discrete Choice Models without the Need for Dynamic Programming.” In *Simulation Based Inference and Econometrics: Methods and Applications* (Mariano, Scheuermann and Weeks, eds.), Cambridge University Press.
- Ha-Dong, M. and N. Treich (2000):** “Recursive Intergenerational Utility in Global Climate Risk Modeling.” Working paper LEERNA-INRA. University of Toulouse 1.
- Hall, R.E. (1988):** “Intertemporal Substitution in Consumption.” *Journal of Political Economy*, 96(April), 339-357.

Howitt, R.E., S. Msangi, A. Reynaud, and K.C. Knapp (2002): “Using Polynomial Approximations to Solve Stochastic Dynamic Programming Problems.” Working Paper, Department of Agricultural & Resource Economics, University of California, Davis.

Imai, S., N. Jain, and A. Ching (2002): “Bayesian Estimation of Dynamic Discrete Choice Models.” In *Proceedings of the 2002 North American Summer Meetings of Econometric Society: Applied Econometrics* (D.Levine, W.Zame and S.Berry, eds.).

Jenkins, M.W, A.J. Draper, J.R. Lund, R.E. Howitt, S. Tanaka, R. Ritzema, G. Marques, S.M. Msangi, B.D. Newlin, B.J. van Lienden, M.D. Davis, K.B. Ward (2001): “Improving California Water Management: Optimizing Value and Flexibility.” Center of Environmental and Water Engineering, Report No 01-1, University of California, Davis

Judd, K.L. (1998): “Numerical Methods in Economics.” M.I.T Press. Cambridge.

Keane, M.P., and K.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(Nov.) 648-672.

Keane, M.P., and K.I. Wolpin (1997): “The Career Decisions of Young Men.” *Journal of Political Economy*, 105: 473-521.

Kerr, A., and G. Read (1998): “Reservoir Management with Risk Aversion.” In *Proceedings of the 33rd Annual Conference of the Operational Research Society of New Zealand*, 30 Aug – 1 Sept, 1998. University of Auckland.

Knapp, K.C. and L.J. Olson (1995): “The Economics of Conjunctive Groundwater management with Stochastic Surface Supplies.” *Journal of Environmental Economics and Management*, 28(May), 340-356.

Knapp, K.C. and L.J. Olson (1996): “Dynamic Resource Management: Intertemporal Substitution and Risk Aversion.” *American Journal of Agricultural Economics*, 78(November), 1004-1014.

Koopmans, T.C. (1960): “Stationary Ordinal Utility and Impatience.” *Econometrica*, 28(April), 287-309.

Krautkramer, J.A., G.C Van Kooten, and D.L. Young (1992): “Incorporating Risk Aversion into Dynamic Programming Models.” *American Journal of Agricultural Economics*, 74(November), 870-878.

Kreps, D.M. and E.L. Porteus (1978): “Temporal Resolution of Uncertainty and Dynamic Choice Theory.” *Econometrica*, 28(January) 185-200.

Lence, S.H. (2000): “Using Consumption and Asset Return Data to Estimate Farmers’ Time Preferences and Risk Attitudes.” *American Journal of Agricultural Economics*, 82(4): 934-947.

- Miranda, M.J., and G.D. Schnitkey (1995):** “An Empirical Model of Asset Replacement in Dairy Production.” *Journal of Applied Econometrics*, 10(Dec, Special Issue):S41-S55.
- Miranda, M.J., and P.L. Fackler (1999):** “Hybrid Methods for Continuous Space Dynamic Programming.” In *Computing in Economics and Finance 1999*, Society for Computational Economics, no. 1332.
- Miranda, M.J., and P.L. Fackler (2002):** “Applied Computational Economics and Finance.” MIT Press, Cambridge, Massachusetts.
- Nelder, J.A. and R. Mead (1965)** “A Simplex Method for Function Minimization.” *Computer Journal*, 7(3): 308-313.
- Normandin, M., and P. Saint-Amour (1998)** “Substitution, Risk Aversion, Taste Shocks and Equity Premia.” *Journal of Applied Econometrics*, 13(3), 265-291.
- Obstfeld, M., (1994):** “Risk-taking, Global Diversification and Growth.” *American Economic Review*, 84, 1310-1329.
- Peltola, J., and K.C. Knapp (2001):** “Recursive Preferences in Forest Management.” *Forest Science*, 47(Nov.):455-465.
- Provencher, B., and R.C. Bishop (1997):** “Estimable Dynamic Model of Recreation Behavior with an Application to Great Lakes Angling.” *Journal of Environmental Economics and Management*, 33, 107-127.
- Provencher, B. (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.
- Rust, J. (1987):** “Optimal Replacement of GMC Bus Engines: An Empirical Model of Harold Zurcher.” *Econometrica*, 55(Sep.) 999-1033.
- Weil, P. (1990):** “Non-Expected Utility in Macroeconomics.” *Quarterly Journal of Economics*, 105(February), 29-42.
- Williams, B. (1989):** “Review of Dynamic Optimization Methods in Renewable Natural Resource Management.” *Natural Resource Modeling*, 3(Spring): 137-216.
- Wolpin, K.I. (1984):** “An Estimable Dynamic Stochastic Model of Fertility and Child Mortality.” *Journal of Political Economy*, 92(5): 852-874.
- Wolpin, K.I. (1987):** “Estimating a Structural Search Model: The Transition from School to Work.” *Econometrica*, 55(4): 801-817.