Estimating Intertemporal Preferences for Natural Resource Allocation

Richard E. Howitt, Siwa Msangi, Arnaud Reynaud, and Keith C. Knapp

In this article, 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 a reservoir in California, and test the data for consistency with a recursive utility model specification versus standard time-additive separability. The results show that the data are consistent with a risk-averse manager with recursive preferences. The data also reject time-additive separability, with or without risk aversion, such as the standard constant relative risk aversion utility model. The improvement in model fit when recursive preferences are used is notable.

Key words: dynamic estimation, recursive utility, stochastic dynamic programming, water management.

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. However, the tendency of researchers to use risk-neutral expected net present value (EPV) objective functions when modeling natural resource problems has caused policy makers to be somewhat skeptical of the real-world relevance of resource economics analysis.[1]

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.), few have actually tried to estimate the degree to which it enters into the decision maker's objective criterion, and none of those studies 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 toward 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 exist, it can be used to calculate empirical moments. Rust 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, Provencher, Miranda and Schnitkey). However, in the few studies that apply dynamic estimation to natural resource problems none specifically estimate intertemporal preference parameters or quantify risk aversion (Provencher, Fulton and Karp).

We seek to address this gap in the environmental and natural resource literature by

---

Richard E. Howitt is professor, Department of Agricultural and Resource Economics, University of California, Davis, CA. Siwa Msangi is research fellow, International Food Policy Research Institute, Washington, DC. Arnaud Reynaud is research fellow, LERNA-INRA, University of Social Sciences, Toulouse. Keith C. Knapp is professor, Department of Environmental Sciences, University of California, Riverside, CA.

The authors gratefully acknowledge support from a risk management grant from USDA-ERS that made this study possible, and the comments of several referees that improved the quality of the manuscript.

[1] This article was motivated by a comment made at an agency workshop in response to the presentation of results from a conventional risk-neutral stochastic dynamic programming (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.

Amer. J. Agr. Econ. 87(4) (November 2005): 969–983

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, Craddock et al.). 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 are 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 article is as follows. In the next section, we describe the general resource allocation problem and the recursive utility specification that we use. In the empirical application 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 studies use EPV as the objective, which uses the discount rate as the sole specification of time preference. In this article, recursive utility is used to specifically estimate parameters for each of these characteristics.

Model Specification

Reservoir system dynamics are given by:

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

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 and allow for precautionary flood control storage 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 network. There are two reasons to decouple a single state from the resource network.[2] The first is the reduction in the dimensionality of the 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 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 be satisfied by either 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 time[3] 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.

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 $\tilde{e}_{1t}$, 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

[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 multistate 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 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 and 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

$$(2) \qquad q_t = w_t + \tilde{e}_{2t}.$$

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

$$(3) \qquad W(q) = \int_0^q P(u)\,\mathrm{d}u.$$

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 presents, in a deterministic context, the first axiomatic presentation of recursive preferences. While Kreps and Porteus 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), to farm production behavior (Lence). In addition, Knapp and Olson (1996), Ha-Dong and Treich, and Peltola and Knapp 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 effects.[4] 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

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

where $\beta \in [0, 1]$ is the subjective discount factor, $\beta = 1/(1+\delta)$, δ the subjective rate of discount, $\alpha \in [<1, \neq 0]$ the risk-aversion parameter, $\rho \in [<1, \neq 0]$ the constant of resistance to intertemporal substitution, and $\mathrm{E}[\cdot]$ 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 EPV utility function.

In the next part of the article, 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

[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.

two recursive utility parameters. Various authors have proposed estimates of the EIS that range from 0 (Hall) 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). 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 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 article.

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:

$$(5) \qquad \max U_1$$

s.t.

$$(6a) \qquad U_t = \left\{(1-\beta)\cdot \mathrm{E}_{e_2}\left[W_t^{\rho}(q_t)\right] + \beta\left[\mathrm{E}_{e_1}\left(U_{t+1}^{\alpha}\right)\right]^{\frac{\rho}{\alpha}}\right\}^{\frac{1}{\rho}}$$

$$(6b) \qquad S_{t+1} = S_t + \tilde{e}_{1t} - w_t - \mathrm{sp}_t$$

$$(6c) \qquad q_t = w_t + \tilde{e}_{2t}$$

$$(6d) \qquad S_{t+1} \geq \underline{S}$$

$$(6e) \qquad S_{t+1} \leq \overline{S}$$

$$(6f) \qquad w_t \geq 0$$

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 assume a stationarity of preferences that we will test for later in the article. The stochastic dynamic recursive equation defining optimal natural resource management is

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

where $V(\cdot)$ is the value function representing the maximized value of recursive utility and w 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 (6a)–(6e), 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). In implementing this fixed-point procedure, we employ an orthogonal polynomial approximation to the value function as a projection method (Judd), for computational efficiency and to accommodate our continuous state variable specification. Unlike the type of collocation projection method advocated by Miranda and Fackler (1999, 2002) for the solution of continuous-state dynamic programming problems, we fit the polynomial approximation to the value function by a least squares criterion with respect to chosen node points. In the empirical part of the article, we show that this criterion allows us to achieve convergent stable results.

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 come 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:

$$(8) \quad \begin{bmatrix} \tilde{e}_1 \\ \tilde{e}_2 \end{bmatrix} \sim \mathrm{N}\left(\begin{bmatrix} \mu_1 \\ \mu_2 \end{bmatrix}, \begin{bmatrix} \sigma_1^2 & \sigma_{12} \\ \sigma_{21} & \sigma_2^2 \end{bmatrix}\right).$$

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

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

and the distribution of the rest-of-network[5] inflow conditional on the reservoir inflow, $\Phi_{2/1}(\bullet)$, is defined by

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

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, measured in millions of acre-feet (MAF) per year, are presented in table 1.

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

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

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

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

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.

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

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 CALVIN[6] 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

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

where q is the quantity of water in MAF and $P(\cdot)$ 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

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

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

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 and reduces the level of

[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" published by the Department of Water Resources of the State of California, California Department of Water Resources (1974–96).

[6] CALVIN is an economically-driven optimization model of California's statewide intertied surface and groundwater system, Jenkins et al. 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.

flood protection. Optimal reservoir management must trade-off 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 ($\tilde{e}_{1t}$) 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 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

$$\begin{aligned} (15)\quad & sp_t(\tilde{e}_{1t}, \mathrm{cap}_t) \\ & = 0.095382 \cdot \tilde{e}_{1t} + 0.005024 \cdot \tilde{e}_{1t}^2 \\ & \quad + 0.000993 \cdot \tilde{e}_{1t}^3 - 0.02305 \cdot \tilde{e}_{1t} \cdot cap_t \end{aligned}$$

with an adjusted R^2 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 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. 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, and Deaton and Laroque, we allow our decision and state variables to be continuous, in contrast to the majority of the structural estimation literature (Rust, Provencher, Miranda and Schnitkey, Wolpin 1984, 1987; Keane and Wolpin 1994, 1997; Erdem and Keane).

Several notable studies have addressed the problem of estimating the relevant parameters within a discrete choice dynamic programming problem, such as Keane and Wolpin's 1994 study, where they address the computational difficulties associated with finding the relevant functions for both the discrete-choice and the intertemporal optimization problems. Recent efforts to overcome these computational difficulties have been addressed by some authors (Geweke and Keane, Aguirregabiria and Mira, Imai et al.)—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 and Provencher, 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 comprised three stages, which can be described as follows.

In the first stage of the estimation, a combination of values for the three behavioral parameters of the model is specified in parameter space. For each set of parameter values, the coefficient values that are multiplied with the basis functions of the numerical Chebychev polynomial approximation to the carry-over value function, are found by the value-iteration (successive approximation) method. Using these polynomial coefficients and the given values of the behavioral parameters, a twenty-three-year sequence is then simulated by optimization of equation (7) 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:

$$\max_{\sigma_1,\sigma_2,\sigma_{12}} L = -n \cdot \ln(2\pi\sigma_1\sigma_2\sqrt{1-\sigma_{12}})$$

$$-\frac{1}{2(1-\sigma_{12}^2)}$$

$$\times \sum_{t=1}^{n} \left\{ \begin{array}{l} (\tilde{w}_t - w_t)^2/\sigma_1^2 + (\tilde{s}_t - s_t)^2/\sigma_2^2 \\ - \dfrac{2\sigma_{12}(\tilde{w}_t - w_t)(\tilde{s}_t - s_t)}{\sigma_1\sigma_2} \end{array} \right\}$$

$$\text{s.t.} \quad \sigma_1, \sigma_2 > 0, -1 < \sigma_{12} < 1$$

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 values of the behavioral parameters in a direction that will maximize the log-likelihood values from (16). We used the Nelder–Mead (NM) search algorithm, since it requires neither derivatives nor concavity of the log-likelihood function in the parameters. The initial iteration of the 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 (equation 7) and maximum likelihood problems (equation 16) 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 NM algorithm works are given by Dennis and Woods. Once a new set of parameter values is obtained, 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 article, 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 EIS. 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 over six points from 0.987 MAF to 2.861 MAF, and a sixth-order Chebyshev orthogonal polynomial approximation of the value function is used:[7]

$$(17) \quad V_C(S) = \sum_{i=1}^{6} a_i \cdot T_i(\hat{S}), \text{ where } \hat{S} = \mathfrak{M}(S).$$

In equation (17), the Chebyshev polynomial approximation comprised coefficients a_i, $i = 1, \ldots, 6$, which are iteratively computed by least squares regression, and then multiplied with basis functions $T_i(\cdot)$. These basis functions, which belong to a family of orthogonal polynomials, are defined over a domain, $\mathfrak{M}(S)$, that represents a mapping of the state variable, S, onto the $[-1, +1]$ interval (Judd).

For given values of the discount factor, risk-aversion, and intertemporal substitution preferences of the decision maker, we must find the values of the polynomial coefficients that allow the numerical approximation in (17) to satisfy the functional equation (7). This is done by, first, solving SDP program with initial guesses of the Chebyshev polynomial coefficient values for each chosen grid point (or node) in state space, and then computing new values of a_i by regressing the values of the nodes against the resulting objective function values of the corresponding SDP solutions. If the resulting coefficients differ from those in the previous step, the SDP problem is re-solved with new Chebyshev coefficient values over the same nodes in state space. The value-iteration algorithm ends once quasi-stabilization of a_i's is achieved, and the numerical approximation (17) is deemed to be a suitable recursive mapping for the functional equation (7). For details of this part of the solution method and its implementation using GAMS, see Howitt et al.

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.

As a final way to check that the Chebyshev approximation is a correct approximation of the Bellman's equation, we calculated the residual difference in the Bellman equation between the optimized solution value and the Chebychev value, as suggested by Judd and Miranda and Fackler (2002). The resulting residual equation oscillates between very small positive and negative values, which means that

[7] Provencher and Bishop, 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.

Table 2. Parameter Estimates for Recursive Utility SDP Model

Parameter[a]	Estimated Value	Standard Error	EIS Value $1/(1-\rho)$	Coefficient of Risk Aversion $(1-\alpha)$
ρ	−9.000	4.60	0.100	
α	−0.440	0.23		1.440
Log likelihood	−10.257			

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

the Chebyshev approximation is an acceptable solution to the Bellman's equation.[8] Moreover, the results show that the residuals did not decrease when the number of Chebychev polynomial values was increased, supporting the conclusion that the sixth-order approximation used is sufficient.[9]

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, along with their calculated standard errors, which were bootstrapped with 500 repetitions. While it is standard to use 1,000 bootstrap repetitions (Efron and Tibshirani), we found the computational time to be excessive and considered 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), resulted in unreliable estimates of the variance–covariance matrix.[10]

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 EIS is low, it is nevertheless, compatible with the results within the macroeconomic literature based on aggregate data. Hall 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 microdata—see, for example, Atkeson and Ogaki—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 the indifference contours that map between income in consecutive periods are strongly curved which implies a relative insensitivity to the discount rate. Following the intuition provided by Lence, 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 1[11] to as high as 18 in Obstfeld. Our estimate of 1.441 is well within the range of previously cited values.

---

[8] Even though the residuals do not vanish exactly at the chosen nodes of the approximation, as they would in a collocation scheme (Miranda and Fackler 2002), the least squares projection method gives sufficient accuracy for the numerical approximation.

[9] The maximum absolute residual value is 0.000725 and the mean absolute value over the twenty grid points is 0.000225. Given a mean value for the Bellman equation of 1.08, these residuals are acceptably small.

[10] The mean gradient over twenty-three 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 validity of the derivative free NM approach in this empirical example.

[11] Epstein and Zin (1991) report a value around 1, for example.

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	1.440
10%	0.91	−10.225	−9.426	0.096	1.446
15%	0.85	−10.214	−10.383	0.088	1.527

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 EIS. A low EIS means, in fact, that the decision maker is relatively insensitive to the interest rate and consequently to β.[12]

We also tested for stationarity of the recursive preferences over the 23-year data series by estimating the parameters over the first and last seven years of observed data and testing for differences between these estimates and those estimated from the full data set. We did so by constructing an F-test from the squared residuals obtained from simulating the model with parameters estimated from these two subsamples, which each represent roughly 30% of the full sample, in accordance with the procedure outlined in Theil. By comparing the test statistic with the appropriate critical F-value, we can evaluate the plausibility of a systematic change over time in the reservoir manager's preferences. Table 4 shows the results of our hypothesis test.

Table 4. Test of the Stationarity of Preferences

	1974–80	1990–96	1974–96
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 subperiods are equal, cannot be rejected by either an F-test or a likelihood ratio test. Hence, our previously stated assumption of stationary preferences over the planning horizon is validated.[13]

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 1, the recursive utility specification devolves to an EPV specification, while just setting α equal to 1 gives us a risk-neutral recursive (RNR) model. The acronyms for the three alternative objective function specifications are specified in table 5.

[12] This also explains why the discount factor does not play a strong role in the estimation procedure, as normally the discount factor has a large impact on intertemporal allocation behavior. 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 three-parameter estimation algorithm.

[13] If the F-test is performed on the residuals from two contiguous subsamples, that is splitting the data into two even subsamples, the resulting F statistic of 0.18 (df 2,19) similarly fails to reject the null hypothesis of equality of the coefficients.

Table 5. Alternative Preference Specifications

	Risk-Averse	Risk-Neutral
Recursive	RAR	RNR
Nonrecursive	CRRA	NPV

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, hence 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) using the same definition of the benefit from reservoir releases as equation (3):

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

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 (RAR), RNR, and CRRA models over that of the EPV model, is also noted.

The implicit restrictions 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 6 strongly reject both the CRRA and the 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 the comparison of the MSE values. However, the mean squared errors for both the storage and the release values from the simulations of both the recursive and the CRRA models show significant improvement over those from the EPV model and speak highly in favor of the recursive utility model specification.

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 6, the release rules can be classified into 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.

Further evidence of the superiority of the recursive utility model with risk over the EPV and CRRA models is shown in figures 2–5, which plot the observed reservoir storage and

Table 6. 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[a]		–	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	–

[a]The likelihood ratio test statistic is compared to the critical level of the χ_1^2 statistic that has a value of 6.63 at the 1% level.

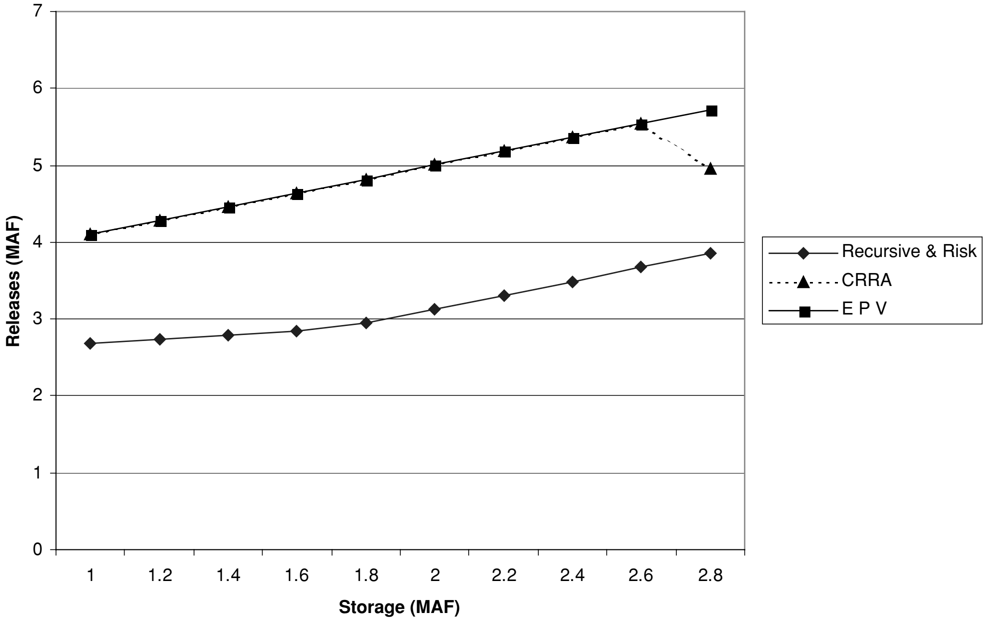

Figure 1. Comparison of release rules under alternative preference specifications

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 are very close to 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.

In comparing storage in 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

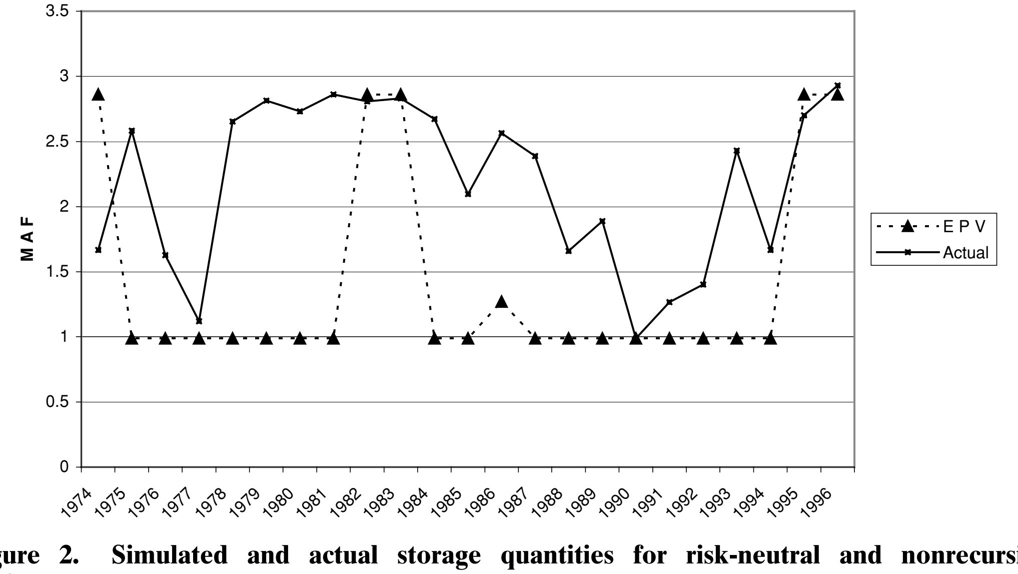

Figure 2. Simulated and actual storage quantities for risk-neutral and nonrecursive preferences (EPV)

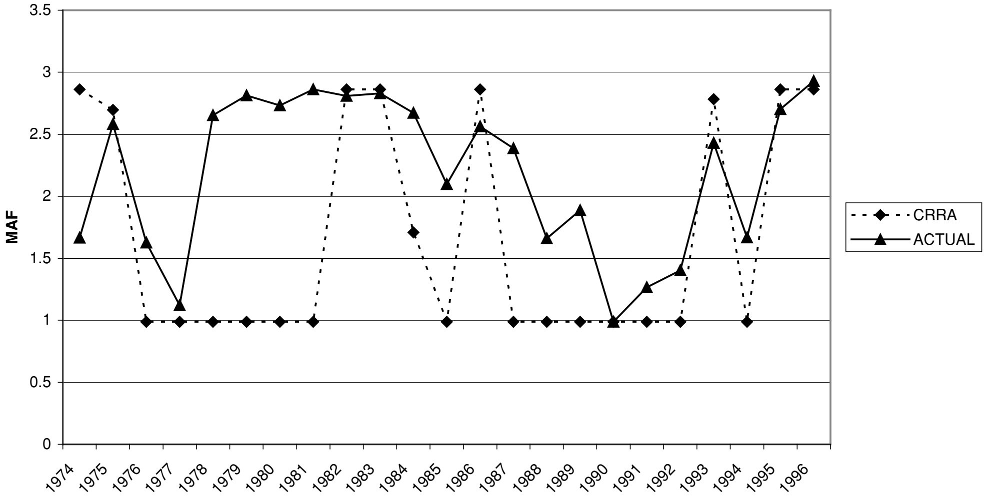

Figure 3. Simulated and actual storage quantities for nonrecursive CRRA preferences

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 the lower bounds of reservoir capacity. 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 6. 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

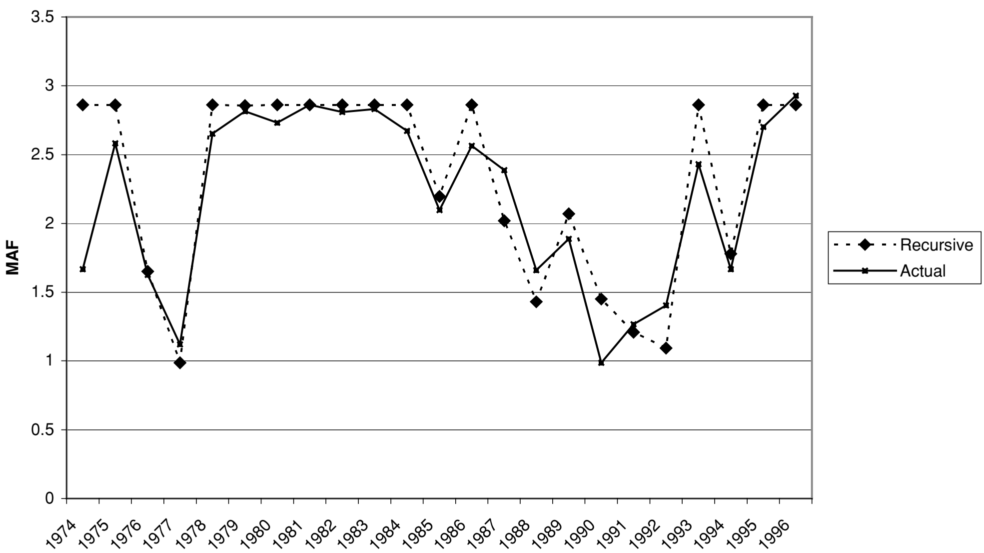

Figure 4. Simulated and actual storage for risk-averse and recursive preferences

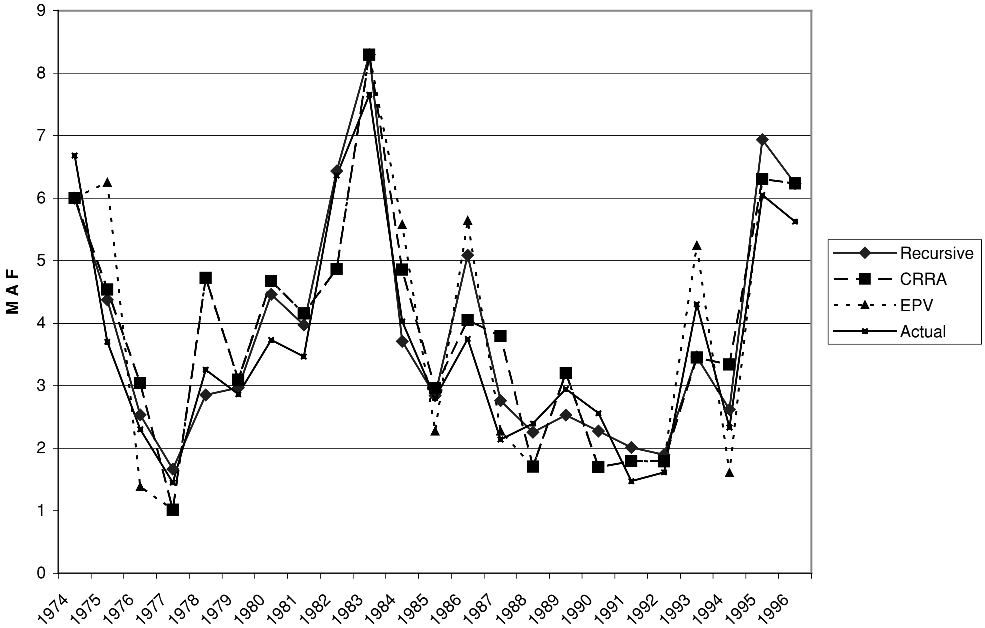

Figure 5. Simulated and actual release quantities

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 1991–92. These results reflect the major point of this article, namely that dynamic behavior seems to be best represented by a dynamic, path-dependent preference specification.

By comparison, the simulation results in figure 2 from the standard risk-neutral EPV model support the misgivings expressed by public decision makers toward 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 6. 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.

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 nonrecursive 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, suggesting 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 article, 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 the CRRA utility specifications in favor of a recursive specification with preferences for both risk and intertemporal substitutions. 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, for this empirical example, there may be little gain 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 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 intertemporal preferences in this example 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.

[Received December 2003; accepted February 2005.]

References

Aguirregabiria, V., and P. Mira. "Swapping the Nested Fixed Point Algorithm: A Class of Estimators for the Discrete Markov Decision Models." *Econometrica* 70(2002):1519–43.

Atkeson, A., and M. Ogaki. "Wealth-Varying Intertemporal Elasticities of Substitution: Evidence from Panel Data and Aggregate Data." *Journal of Monetary Economics* 38(1996):507–34.

Bertsekas, D.P. *Dynamic Programming and Stochastic Control*. New York: Academic Press, 1976.

California Department of Water Resources. *State Water Project Annual Report of Operations*. edited by the State of California, 1974–96. Available at: http://www.co.water.ca.gov/indexo.html

Craddock, M., A.D. Shaw, and B. Graydon. "Risk-Averse Reservoir Management in a Deregulated Electricity Market." *Proceedings of the 33rd Annual Conference of the Operational Research Society of New Zealand*. Auckland, New Zealand, 1998.

Deaton, A., and G. Laroque. "Competitive Storage and Commodity Price Dynamics." *Journal of Political Economy* 104(1996):896–923.

Dennis, J.E., and D.J. Woods. "Optimization on Micro-computers: The Nelder-Mead Simplex Algorithm." *Technical Report 85-9*. Houston, TX: Department of Mathematical Sciences, Rice University, 1985.

Erdem, T., and M. Keane. "Decision Making under Uncertainty: Capturing Dynamic Brand Choice Processes in Turbulent Consumer Goods Markets." *Marketing Science* 15(1996):1–20.

Efron, B., and R.J. Tibshirani. *An Introduction to the Bootstrap*. New York: Chapman and Hall, 1993.

Epstein, L.G., and S.E. Zin. "Substitution, Risk Aversion and the Temporal Behavior of Consumption and Asset Returns: A Theoretical Framework." *Econometrica* 57(1989):937–69.

——. "Substitution, Risk Aversion and the Temporal Behavior of Consumption and Asset Returns: An Empirical Analysis." *Journal of Political Economy* 99(1991):263–86.

Fafchamps, M. "Sequential Labor Decisions under Uncertainty: An Estimable Household Model of West-African Farmers." *Econometrica* 61(1993):1173–97.

Fulton, M., and L. Karp. "Estimating the Objectives of a Public Firm in a Natural Resource Industry." *Journal of Environmental Economics and Management* 16(1989):268–87.

Geweke, J., and M.P. Keane. "Bayesian Inference for Dynamic Discrete Choice Models Without the Need for Dynamic Programming." *Simulation Based Inference and Econometrics: Methods and Applications*, R. Mariano, T. Scheuermann, and M.T. Weeks, eds., Cambridge: University Press, 1995.

Ha-Duong, M., and N. Treich. "Risk Aversion, Intergenerational Equity and Climate Change." *Environmental and Resource Economics* 28(2004):195–207.

Hall, R.E. "Intertemporal Substitution in Consumption." *Journal of Political Economy* 96(1988):339–57.

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

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

Jenkins, M.W., J.R. Lund, R.E. Howitt, A.J. Draper, S.M. Msangi, S.T. Tanaka, R.S. Ritzema, and G.F. Marques. Optimization of California's Water Supply System: Results and Insights." *Journal of Water Resources Planning and Management* 130(2004):271–80.

Judd, K.L. *Numerical Methods in Economics*. Cambridge, MA: MIT Press, 1998.

Keane, M.P., and K.I. Wolpin. "The Solution and Estimation of Discrete Choice Dynamic Programming Models by Simulation and Interpolation: Monte Carlo Evidence." *The Review of Economics and Statistics* 76(1994):648–72.

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

Kerr, A., and G. Read. "Reservoir Management with Risk Aversion." *Proceedings of the 33rd Annual Conference of the Operational Research Society of New Zealand*. University of Auckland. 30 August–1 September, 1998.

Knapp, K.C., and L.J. Olson. "Dynamic Resource Management: Intertemporal Substitution and Risk Aversion." *American Journal of Agricultural Economics* 78(1996):1004–14.

Koopmans, T.C. "Stationary Ordinal Utility and Impatience." *Econometrica* 28(1960):287–309.

Krautkramer, J.A., G.C. Van Kooten, and D.L. Young. "Incorporating Risk Aversion into Dynamic Programming Models." *American Journal of Agricultural Economics* 74(1992):870–78.

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

Lence, S.H. "Using Consumption and Asset Return Data to Estimate Farmers' Time Preferences and Risk Attitudes." *American Journal of Agricultural Economics* 82(2000):934–47.

Miranda, M.J., and G.D. Schnitkey. "An Empirical Model of Asset Replacement in Dairy Production." *Journal of Applied Econometrics* 10(1995):S41–55.

Miranda, M.J., and P.L. Fackler. "Hybrid Methods for Continuous Space Dynamic Programming." *Computing in Economics and Finance 1999*, Society for Computational Economics, No. 1332, 1999.

——. *Applied Computational Economics and Finance*. Cambridge, MA: MIT Press, 2002.

Nelder, J.A., and R. Mead. "A Simplex Method for Function Minimization." *Computer Journal* 7(1965):308–13.

Normandin, M., and P. Saint-Amour. "Substitution, Risk Aversion, Taste Shocks and Equity Premia." *Journal of Applied Econometrics* 13(1998):265–91.

Obstfeld, M. "Risk-Taking, Global Diversification and Growth." *American Economic Review* 84(1994):1310–29.

Peltola, J., and K.C. Knapp. "Recursive Preferences in Forest Management." *Forest Science* 47(2001):455–65.

Provencher, B., and R.C. Bishop. "Estimable Dynamic Model of Recreation Behavior with an Application to Great Lakes Angling." *Journal of Environmental Economics and Management* 33(1997):107–27.

Provencher, B. "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(1995):321–38.

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

Theil, H. *Principles of Econometrics*. New York: John Wiley & Sons, Inc, 1971.

Weil, P. "Non-expected Utility in Macroeconomics." *Quarterly Journal of Economics* 105(1990):29–42.

Williams, B. "Review of Dynamic Optimization Methods in Renewable Natural Resource Management." *Natural Resource Modeling* 3(1989):137–216.

Wolpin, K.I. "An Estimable Dynamic Stochastic Model of Fertility and Child Mortality." *Journal of Political Economy* 92(1984):852–74.

——. "Estimating a Structural Search Model: The Transition from School to Work." *Econometrica* 55(1987):801–17.