AgEcon SEARCH
RESEARCH IN AGRICULTURAL & APPLIED ECONOMICS

The World's Largest Open Access Agricultural & Applied Economics Digital Library

This document is discoverable and free to researchers across the globe due to the work of AgEcon Search.

Help ensure our sustainability.

Give to AgEcon Search

AgEcon Search
http://ageconsearch.umn.edu
aesearch@umn.edu

*Papers downloaded from **AgEcon Search** may be used for non-commercial purposes and personal study only. No other use, including posting to another Internet site, is permitted without permission from the copyright owner (not AgEcon Search), or as allowed under the provisions of Fair Use, U.S. Copyright Act, Title 17 U.S.C.*

Calibrated Stochastic Dynamic Models For Resource Management*

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

Abstract

In this paper we develop a positive calibrated approach to stochastic dynamic programming. Risk aversion, discount rate, and intertemporal substitution preferences of the decision-maker are calibrated by a procedure that minimizes the mean squared error from data on past decisions. We apply this framework to managing stochastic water supplies from Oroville Reservoir, located in Northern California. The calibrated positive SDP closely reproduces the historical storage and releases from the dam and shows sensitivity of optimal decisions to a decision-maker's risk aversion and intertemporal preferences. The calibrated model has average prediction errors that are substantially lower than those from the model with a risk-neutral expected net present value objective.

Keywords: Stochastic Dynamic Programming, Positive Calibration, Natural resource management

JEL Codes C61, D78, Q25

---

* Paper to be presented at the World Congress of Environmental and Resource Economists – Monterey California. June 24-27 2002 and American Agricultural Economics Association Annual Meeting - Long Beach California, July 27-31 2002.

[a] Department of Agricultural and Resource Economics, University of California at Davis.
[b] LEERNA-INRA, University of Social Sciences Toulouse.
[c] Department of Environmental Sciences, University of California at Riverside.

1. Introduction

Most natural resource management problems are inherently dynamic and stochastic in their evolution over time. Stochastic dynamic programming (SDP) provides a unifying framework for the economic analysis of natural resource management, since it is able to formally integrate stochastic biophysical relationships with the imputed value functions of economic users of the resources. In most cases, constraints, non-linearity in the equations of motion, and non-gaussian stochastic processes make certainty equivalent approximations inappropriate (Bertsekas, 1976). In addition, many resource management problems and policies are concerned with the effect of changes in the higher moments of the dynamic distributions, and not just their expected value. For example, preliminary investigations of the hydrologic impact of global warming indicate that the change in the distribution of precipitation will be more pronounced than changes in the mean. While a few natural resource papers show risk included within an SDP formulation (Knapp & Olson, 1996; Krautkramer et al., 1992), none are calibrated to observed decisions to show 'revealed' level of risk-aversion. The approach in this paper elicits those parameter values that are most likely to reflect the underlying inter-temporal preferences of the decision-maker. Models that do not account for risk-driven behavior will tend to be trusted less by public decision-makers who seek reliable analytical tools for application to the problems they face[1].

Despite its methodological appeal, SDP has not been as widely used for the empirical analysis of natural resource problems as was initially expected (Burt & Allison, 1963). Forty years after its introduction, there are still relatively few published economic resource management studies that use SDP as an empirical tool. One reason that may influence the application of SDP to policy problems is that it is specified, in most applications, with a normative objective function. In other words SDP is often viewed as a method for normative policy optimization under uncertainty. The ability of SDP to be used as a positive analytical method that can reproduce the historical behavior of decision-makers has been somewhat neglected in policy analysis applications, and confined mostly to discrete-choice econometric applications (Keane & Wolpin, 1994; Provencher & Bishop, 1997; Provencher, 1995; Miranda & Schnitkey, 1995). Yet, demonstrated prediction precision is important, especially if SDP is to become a decision tool for managing natural resources. We believe that the specification of SDP models on a positive rather than a normative basis will reassure decision makers that the base model

[1] This paper was motivated by a comment made at an agency workshop in response to the presentation of results from a conventional SDP solution. The commentator was Dr. Francis Cheung of the California Department of Water Resources.

incorporates their decision preferences, which may help to make SDP a more useful tool for managing natural resources under uncertainty.

Fulton and Karp (1989) used an inverse control problem approach to infer the objective function parameters of a dynamically-optimizing mining firm. By using the optimal feedback rule of their stochastic programming problem jointly with the underlying equation of motion, they obtain estimating equations that recover the parameters of their linear quadratic utility function. They then simulate alternative economic scenarios, and infer the relative importance of profit maximization with respect to other possible objectives of the firm, such as maintaining levels of employment, output or mineral reserves. In our case, we are weighing the relative importance of risk aversion and intertemporal substitution in explaining the observed behavior of our agent, within the context of a more general set of preferences.

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 problem and the inter-temporal optimization problem. The fact that our decision problem is a continuous one, allows us to concentrate our efforts on developing a robust and efficient algorithm to obtain the value function, so as to solve the dynamic problem over the possible combinations of the decision-maker's risk-aversion and resistance to inter-temporal substitution. The focus of their paper, however, is on developing a tractable numerical procedure to approximate the value functions that span the possible discrete choice set of the decision-maker at each stage of the dynamic problem. In an earlier, seminal paper (Rust, 1987) chose to simplify the discrete-choice mechanism for numerical tractability, in order to solve the stochastic dynamic problem and use it to explain the observed behavior of an agent, as we aim to do in this paper.

In this paper we specify SDP as a positive analytical method, rather than a normative prescriptive tool. Like all positive methods, the calibration of SDP solutions requires a set of observed actions by decision makers, and one or more parameters that can be calibrated to improve the fit of the model to the observed past actions. Analysis of recursive utility has shown that risk aversion, the subjective discount rate, and intertemporal substitution all influence dynamic economic allocation (Knapp & Olson, 1996). Accordingly we solve the SDP problem over a grid of parameter values and select the parameter set that has a minimum mean squared error from the observed decisions. Of course, resource decisions that are optimal for the decision-makers are not guaranteed to be optimal

---

He pointed out that optimization models tend to be discounted by decision-makers because they ignore the presence of risk in the objective function.

for their clients. However given the accountability of democratic institutions, it is reasonable to assume that decision-makers will not consistently depart from their constituent's preferences over a long period of time; in our case twenty three years.

2. An SDP Specification of Natural Resource Management

This section develops the calibrated SDP framework for a hypothetical natural resource management problem with reservoir management as an example. Two important points should be emphasized. First, the majority of natural resource management problems require the specification of an interdependent multi-state model, and any simplification to a single state must take into account the rest of the resource network. Second, managing risk and inter-temporal substitution is an integral part of resource management. This requires a correct specification of the risk and substitution preferences of the decision-maker.

2.1 Decoupling network states

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 natural resource being modeled from the rest of the network. There are two reasons to decouple a single state from the resource network. 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 the approach is to calibrate the SDP model to a historic series of observed decisions. In most resource systems, decisions are split between agencies or levels of agencies. A convincing calibration that reflects past decision 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 (1989) where he models the actions of a single individual.

Since resource states are usually interconnected, the management of a given resource or location depends on the state of the rest of the system.

Figure 1: Natural resource management model

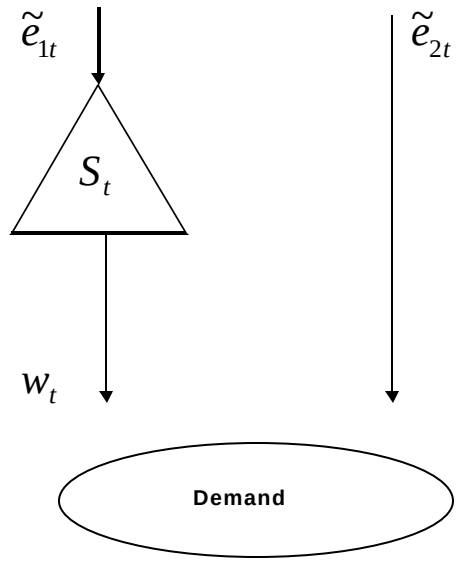

In Figure 1, we consider the simplest representation of a resource network based on a single state representation in a complex network system. The single state is decoupled from the network by approximating the network by two elements: a natural resource, with stochastic inflow $\tilde{e}_{1t}$ and storage S_t at each date *t*, and the rest of the system characterized by a stochastic inflow $\tilde{e}_{2t}$. The system dynamics are given by:

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

The change in natural resource stock must balance the local inflow and the release. For reservoir management, equation (1) states that the variation of the reservoir storage plus the stochastic inflow must be equal to the water release w_t. The index *t* in (1) denotes time period. Final demand for water in Figure 1 may either be satisfied by water release w_t or by flows from the rest of the system $\tilde{e}_{2t}$.

2.2 Exogenous stochastic variables

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[2] 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.

[2] 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.

We define the following timing of information and controls. First, the decision-maker observes S_t and the realization of the exogenous stochastic variable $\tilde{e}_{1t}$, in the reservoir management example, the local stochastic inflow. Second, the decision-maker chooses the control w_t, the level of water release. This choice is a function of the future local stochastic inflow. The natural resource available for consumption is, at each date, made of 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 $\Phi_{2/1}(\bullet)$.

2.3 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:

$$q_t = w_t + \tilde{e}_{2t} \tag{2}$$

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

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

The net surplus of resource consumption is a concave 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), and more recently to resource management by Knapp and Olson (1996), Ha-Dong and Treich (2000) and Peltola and Knapp (2001). Three main arguments are advanced in favor of this class of preferences. 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[3]. 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:

$$U_t = \left\{ (1-b)\cdot W_t^r + b\left[\mathrm{E}\left(U_{t+1}^a\right)\right]^{\frac{r}{a}} \right\}^{\frac{1}{r}} \tag{4}$$

where $b \in [0,1]$ is the subjective discount factor, $b = 1/(1+d)$, d is the subjective rate of discount, $a \in [<1, \neq 0]$ is the risk-aversion parameter, and $r \in [<1, \neq 0]$ the constant resistance to intertemporal substitution. Given this specification, the elasticity of intertemporal substitution (EIS), s, is equal to $1/(1-r)$, $s \in [0,+\infty)$. It follows that a decrease of the constant inter-temporal substitution resistance, r, below 1 results in a lower inter-temporal elasticity of substitution. Finally, note that recursive preferences nest expected utility as a special case: by setting $a = r$ we get the familiar constant relative risk aversion expected utility function. In what follows, we endogenously calibrate the decision-maker's risk aversion, discount factor, and resistance to inter-temporal substitution. Three main reasons support the endogenous calibration of these parameters.

Reason 1: there is no consensus in the economic literature on the level of the two recursive utility parameters. Various authors have proposed estimates of the EIS that range from 0 (Hall, 1988) all the way to 0.87 (Epstein and Zin, 1991), while estimates of the risk aversion coefficient (1-a) range from 0.82 (Epstein and Zin, 1991) to 1.5 (Normandin and Saint-Amour, 1998).

Reason 2: 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, while Ha-Duong and Treich (2000) show that larger risk aversion strengthens optimal pollution

[3] Attitude toward variations in consumption across states of the world can be characterized by risk aversion. Attitude toward variations in consumption across time is 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 are unattractively linked. Recursive preferences allows risk attitudes to be disentangled from the degree of intertemporal substitutability.

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.

Reason 3: It is very important to have a model that closely represents behavior of decision-makers. By endogenously calibrating some parameters to fit past decisions, we move from a normative perspective of SDP to a more positive analytical approach.

The approach that we adopt to obtain these parameters differs from that of other authors dealing with recursive utility (Lence, 2000; Epstein & Zin, 1991), since we do not observe the data that would allow us to statistically fit our model parameters to empirical moments using GMM or a similar procedure. Instead we adopt a least-squares criterion to search over the space of parameter values, such that we obtain the parameter set that is most consistent with the observed behavior . We prefer to call this an "endogenous calibration" procedure, rather than an estimation procedure, since we don't claim any statistical properties for our "best-fit" parameters. This also differs from the type of 'calibration' done in the traditional macroeconomics literature (Kydland & Prescott, 1991; King, Plosser & Rebelo, 1988), where the parameters in the model are chosen to match moments observed in empirical studies.

2.4 The SDP for resource management

We assume that the decision-maker wishes to maximize utility subject to the equation of motion for the natural resource stock and the feasibility constraints:

$$\underset{w}{\text{Max}} \;\; U_t = \left\{ (1-b)\cdot \mathrm{E}_{e_2} \;\; W_t^{\,r}(q_t) + b\left[\mathrm{E}_{e_1} \;\; U_{t+1}^{a}\right]^{\frac{r}{a}} \right\}^{\frac{1}{r}} \tag{5}$$

$$\text{s.t.} \begin{cases} S_{t+1} = S_t + \tilde{e}_{1t} - w_t & \text{(6a)} \\ q_t = w_t + \tilde{e}_{2t} & \text{(6b)} \\ S_{t+1} \geq \underline{S} & \text{(6c)} \\ S_{t+1} \leq \overline{S} & \text{(6d)} \\ w_t \geq 0 & \text{(6e)} \end{cases}$$

The stochastic control problem consists of choosing a sequence of decision rules for resource flows that maximize the objective function (5) subject to (6a)-(6e). 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 objective function is therefore the expected current net surplus. All model parameters and functions

are the same for all decision stages, which means that the problem is stationary. If the planning horizon is infinite, the optimal decision vector in state space for any decision stage is the same for all stages. The optimized value of the system at any stage is the same for all stages, and is finite even though the planning horizon is infinite, because stage returns are discounted. The stochastic dynamic recursive equation defining optimal natural resource management is:

$$V\left(S,\tilde{e}\right)=\underset{w}{\text{Max}}\ \left\{(1-b)\cdot\int W_t^{\,r}(w+e_2)\mathrm{d}F_{2/1}+b\left[\int V^a(S,\tilde{e})\mathrm{d}F_1\right]^{\frac{r}{a}}\right\}^{\frac{1}{r}} \tag{7}$$

where *V(.)* is the value function and *w* must be feasible. We have to solve a standard SDP problem. The value iteration method used for solving (7) 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.

2.5 Solving the Calibrated SDP problem

Given that we don't want to specify *a priori* values for the Arrow-Pratt constant relative risk-aversion, the discount factor, and the constant resistance to intertemporal substitution, a better formulation of (7) is:

$$V\left(S,\tilde{e};a,b,r\right)=\underset{w}{\text{Max}}\ \left\{(1-b)\cdot\int W_t^{\,r}(w+e_2)\mathrm{d}F_{2/1}+b\left[\int V^a(S,\tilde{e};a,b,r)\mathrm{d}F_1\right]^{\frac{r}{a}}\right\}^{\frac{1}{r}} \tag{8}$$

We use a grid-search in order to determine the optimal discount factor and the resistance to intertemporal substitution conditional on a grid of 11 risk aversion values from 1 to – 5. Since initial empirical solutions suggested that α, the risk aversion parameter, had the least effect on the predicted MSE of the model, we search over the discount factor and intertemporal substitution parameters conditional on 11 rates of risk aversion. Defining the limits of the parameter space as $b\in[\underline{b},\overline{b}]$ and $r\in[\underline{r},\overline{r}]$, we discretize these intervals into n_b and n_r values and solve the $n_b\times n_r$ corresponding stochastic dynamic programs. This results in $n_b\times n_r$ value functions. Given the historical data of the local and the rest-of-network inflows, we obtain the resulting $n_b\times n_r$ optimal policies, $w_t^*(b,r)$ for $t=1,\ldots,T$. We can now compute a measure of the difference between the predicted model allocations $w_t^*(b,r)$ and the historical observed allocations w_t^o.

Finally, the optimal discount factor and the constant resistance to intertemporal substitution parameters are those values that minimize the sum of squared prediction errors. In other words, the selected parameters are those that best reproduce historical actions of the decision-makers.

3. A Calibrated SDP for Oroville reservoir

Oroville Reservoir is located on the Feather River in Northern California. The State of California 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.

3.1 Specification of the problem

We consider the optimal annual use of Oroville reservoir and limit our analysis to the inter-year management problem. In the following paragraphs we determine the annual optimal water releases and carryovers[4].

The water network considered is similar to that in Figure 1 with the addition of a spill flow from the stock S_t. The change in the reservoir storage plus the stochastic inflow must be equal to the water release w_t, and the spills from the reservoir, sp_t. The spills balance the system in times of high flows, but have no economic value in the model.

Distribution of inflows

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

[4] Focusing only on inter-year reservoir management does not mean that intra-year management is without interest. Intra-year management would, however, require some model important changes. First, the temporal independence assumption of inflows would not hold, and inflows should be modeled as an auto regressive process. Second, adaptive stochastic dynamic programming should be used to take into account any updated information, which would make the resulting model much less tractable.

$$\begin{bmatrix} \tilde{e}_1 \\ \tilde{e}_2 \end{bmatrix} \sim \mathrm{N}\left(\begin{bmatrix} m_1 \\ m_2 \end{bmatrix}, \begin{bmatrix} s_1^2 & s_{12} \\ s_{12} & s_2^2 \end{bmatrix} \right). \tag{9}$$

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

$$\tilde{e}_i \sim \mathrm{N}\left(m_i, s_i^2\right) \tag{10}$$

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

$$\tilde{e}_2 | e_1 \sim \mathrm{N}\left(m_2 + \frac{s_{12}}{s_1^2}\left(e_1 - m_1\right), s_2^2 - \frac{s_{12}^2}{s_1^2} \right). \tag{11}$$

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
m_1	3.7957	0.6009	6.317
m_2	15.7583	2.2742	6.929
s_1^2	6.8594	2.2257	3.082
s_2^2	98.2635	31.8850	3.082
s_{12}	24.1569	8.1346	2.970

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

$$\tilde{e}_1 \sim \mathrm{N}\left(3.7957,\ 6.8594\right) \tag{12}$$

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

$$\tilde{e}_2 | e_1 \sim \mathrm{N}\left(2.3910 + 3.5217 \cdot e_1,\ 13.1896\right). \tag{13}$$

The reservoir inflow and rest-of-network conditional inflow distributions are discretized over 8 points.

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 function used in the CALVIN[5] 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 total inflow to the Delta per year and associated shadow values. A quadratic form is fitted to the data generated by CALVIN. The resulting inverse demand function is:

$$P(q) = 150 - 2.9 \cdot q + 0.02 \cdot q^2 \qquad (14)$$

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:

$$W(q) = 150 \cdot q - 1.45 \cdot q^2 + 0.0067 \cdot q^3 . \qquad (15)$$

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

Spillways

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, (sp_t), 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 t is defined as the difference between the maximum storage capacity of the reservoir ($\overline{S}$) 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:

$$sp_t(\tilde{e}_{1t}, cap_t) = 0.095382 \cdot \tilde{e}_{1t} + 0.005024 \cdot \tilde{e}_{1t}^2 + 0.000993 \cdot \tilde{e}_{1t}^3 - 0.02305 \cdot \tilde{e}_{1t} \cdot cap_t \qquad (16)$$

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 the relationship in equation (16) that links spills, inflows and storage capacity is known by the decision-maker.

---

[5] CALVIN is an economically-driven optimization model of California's statewide inter-tied surface and groundwater system, Jenkins *et al* (1999). CALVIN optimizes the operations of system resources over a given hydrologic sequence to

Discount rate and inter-temporal substitution preferences

We consider 11 possible values for the discount factor b. The values are uniformly distributed on the interval $[0.645, 0.943]$ which represents discount rates from 6% to 55%. Similarly, we consider 11 possible values for ρ, the (constant) resistance to intertemporal substitution which are uniformly distributed on the $[1, -19]$ interval. As a result, 121 SDPs must be solved for each level of risk aversion, one for each possible pair (b, r), to generate a grid from which the optimal calibration parameters are selected. We selected the risk aversion parameter to be exogenous to the grid search based on the results of other empirical studies of macro economic and natural resource problems that indicate that the optimal controls are least sensitive to the risk aversion parameter value (e.g. Peltola and Knapp (2001)).

The SDP formulation

Given flood control constraints, the maximum storage capacity in Lake Oroville is 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 optimization program is:

$$\underset{w}{\text{Max}} \quad U_t = \left\{ (1-b)\cdot \mathrm{E}_{e_2} \quad W_t^r(q) + b\left[\mathrm{E}_{e1} \; U_{t+1}^a\right]^{\frac{r}{a}} \right\}^{\frac{1}{r}} \tag{17}$$

$$\text{s.t.} \begin{cases} S_{t+1} = S_t + \tilde{e}_{1t} - sp_t(\tilde{e}_{1t}, cap_t) - w_t & \text{(18a)} \\ q_t = w_t + \tilde{e}_{2t} & \text{(18b)} \\ S_{t+1} \geq 0.98709 & \text{(18c)} \\ S_{t+1} \leq 2.86141 & \text{(18d)} \\ w_t \geq 0 & \text{(18e)} \end{cases}$$

where the spill function is given by equation (16).

---

maximize statewide net willingness-to-pay of urban consumers and agricultural producers for additional water.

3.2 Solving the model

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 function[6]:

$$V_C(S) = \sum_{i=0}^{5} a_i \cdot T_i\left(\hat{S}\right), \text{ where } \hat{S} = \mathrm{M}(S). \tag{19}$$

The Chebyshev polynomial coefficients a_i $i = 1,\ldots,5$ are iteratively computed using the Chebyshev regression algorithm, and $\mathrm{M}(S)$ is a mapping of S onto the [–1, –1] interval, Judd (1998). For each possible value of the discount factor and inter-temporal 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's. If the resulting coefficients differ from those in the previous step, the SDP is re-solved with new Chebyshev coefficients. The 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).

Figure 2: Parameter Calibration Surface

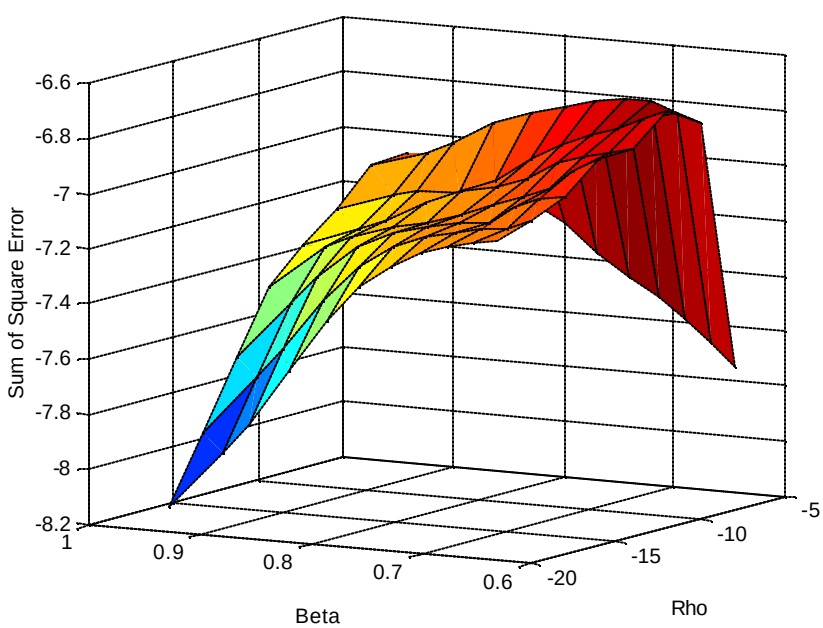

[6] Provencher and Bishop (1997), in a different context, also use such a polynomial approximation of the value function. They nest the dynamic programming approach with a maximum of likelihood procedure.

Positive calibration of parameters

The next step of the analysis consists of selecting the set of discount factor and resistance to inter-temporal substitution parameters, conditional on risk aversion, that best fit the historic realizations of the decision-maker. Figure 2 gives the sum of squared prediction errors[7] for a risk aversion parameter = 1 (risk neutral) , as a function of the two other parameters. For simplicity, the ordinate in Figure 2 is the negative of the sum of squared prediction errors. The higher the ordinate, the smaller the sum of squared prediction errors. The unique minimum sum of squared error is achieved for at r = -9 and b = 0.645. This corresponds to an EIS equal to 0.1 and a subjective discount rate of 55%. These values are admissible but extreme. Moreover, as shown in Figure 2, the sum of square errors is influenced much more by the intertemporal substitution preferences than by the discount factor. For a given value of either r or b, the sum of squares is a concave function of the other parameter. However the MSE measure hardly changes when the risk aversion parameter is changed over the range of 1 to –1. This is consistent with the results reported by Peltola and Knapp (2001)

To evaluate the quality of fit of the SDP, we simulate the optimal predicted releases and storage for Oroville Lake Reservoir under two scenarios. The first simulation is the calibrated SDP with the intertemporal substitution parameter (r = –9.0) set at the value that minimizes the mean squared error, while the discount factor (b) and the risk aversion parameter (a) are set to 0.893 and –1, respectively. While these are not the values of a and b that minimize the mean squared error, they are very close, and conform more closely to parameter values that are more commonly observed in the literature. These parameter values imply mild risk aversion and a subjective discount rate of utility of 12.5%.

The second simulation has the same discount rate (b = 0.893), while the risk aversion and substitution parameters are set such that a = r = 1. These values result in a risk neutral objective function that simply maximizes the expected net present value (ENPV) from operating the reservoir. The results for these two simulations are compared for water releases (the control variable), and storage (the state variable) with the observed releases and storage. Both the releases and storage show significant differences in the average prediction error between the calibrated preferences and the risk neutral ENPV specifications. Figure 3 presents the results of the two policy simulations plotted with

[7]The sum of prediction square errors is equal to $\sum_{t=1974}^{1996}\left[w_t^*(b,r)-w_t^0\right]^2$ where w_t^0 is the observed water release from Oroville Lake Reservoir at time t and $w_t^*(b,r)$ is the optimal release predicted by the model

the observed levels for reservoir releases. It shows the optimal releases that result from solving the annual optimization problem under the two sets of alternative parameters. The average absolute percent error differs substantially between the runs. The ENPV objective has an average error of 23.13%, while the calibrated solutions with a substitution parameter ($r = -9$) and risk aversion ($a = -1$) reduces the error by thirty seven percent to an average error of 14.51 %.

The ENPV objective function run over-shoots the turning points both on high and low releases, as one would expect due to the absence of a precautionary cost incorporated into the calibrated simulation. For example, in 1984 when actual releases were 4.30 MAF the calibrated model predicts 3.71 MAF while the ENPV model predicts 5.58 MAF. The equivalent storage results (Figure 4) show that this high release by the ENPV model drew the storage level down to its lower limit of 0.987 MAF at the end of 1984, while the calibrated model carried over two and a half times as much at 2.86 MAF. The actual storage level was 2.67 MAF. In the following year (1985) the storage under ENPV was again at its lower bound. Due to the lower inflow that year, the ENPV releases are forced to be lower that year than the calibrated run (2.27 MAF vs 3.12 MAF), and the actual releases which were 2.80 MAF. This shows that ignoring intertemporal preferences causes the ENPV model to both over and under shoot the observed releases.

Figure 3: Simulation of Water Releases

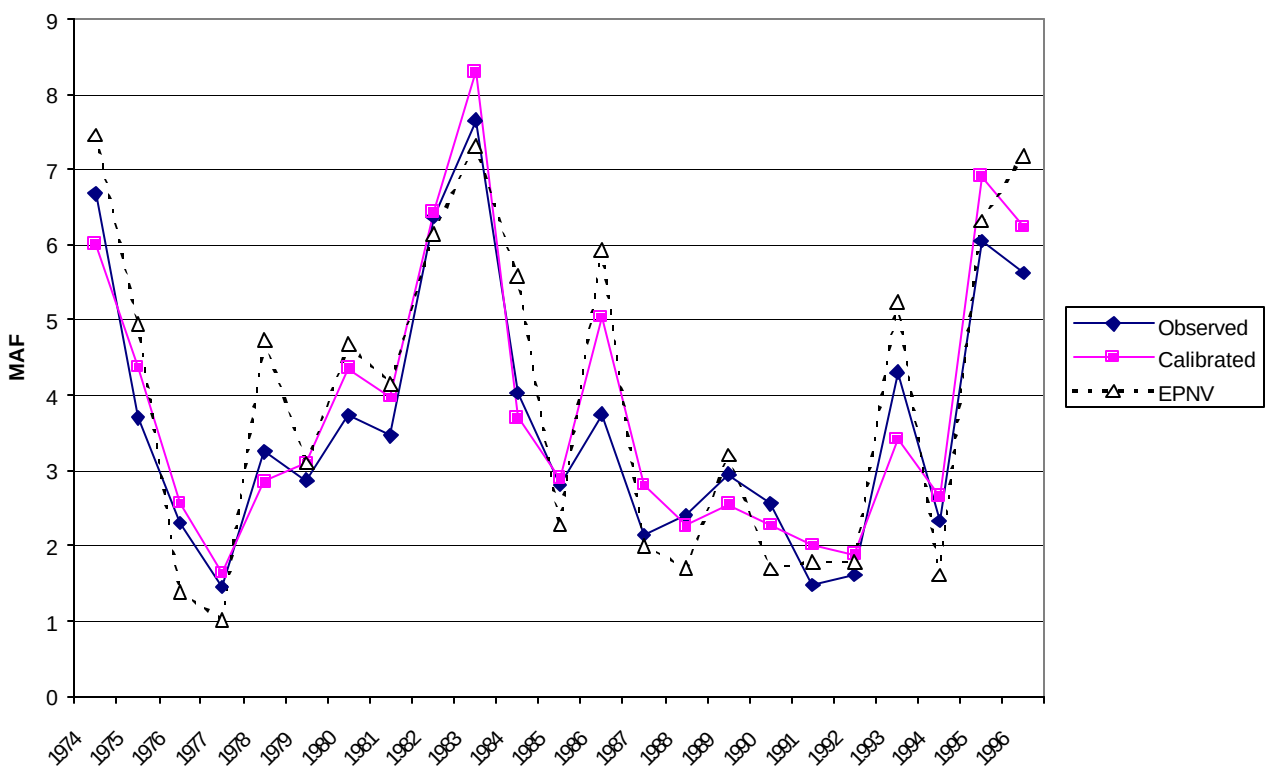

In other years, the prediction precision of both models is influenced by the spill equation (16). In 1986, both models overshot the actual release of 3.75 MAF with the calibrated model at 5.04 MAF and the ENPV model at 5.93 MAF. The difference is due to the high level of spills of 2.10 MAF, the second highest level from 1974 to 1996. This level was not accurately captured by the spill model which predicted 0.90 MAF using equation (16).

Figure 4: Simulation of Storage

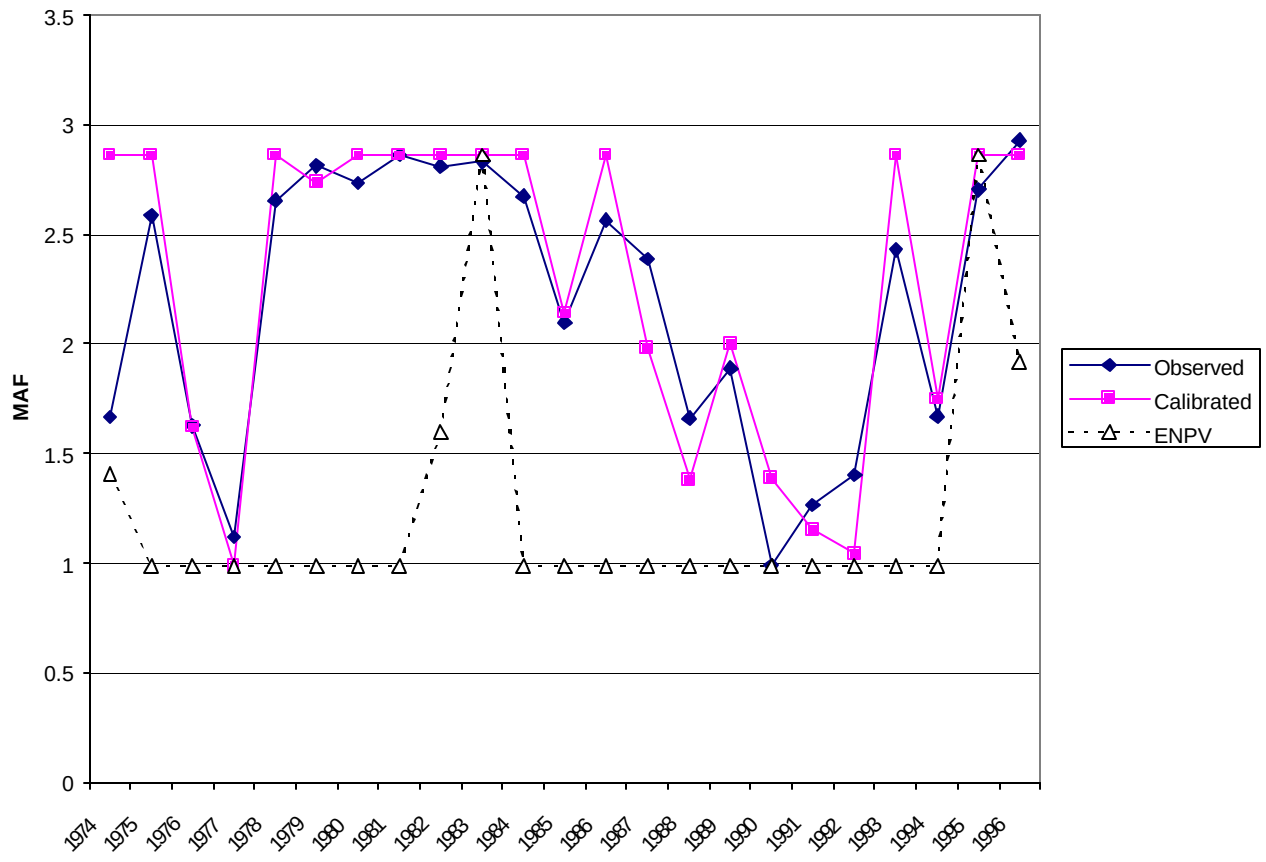

The quantitative differences between the simulations are even more marked in water storage behavior, where the ENPV solution hits the lower bound of reservoir capacity in 18 out of 23 years, in contrast to the calibrated simulation which tracks the observed storage levels fairly closely, hitting the lower bound once, and the upper bound thirteen times. The average storage error for the calibrated model is 12.14%, while the ENPV model has an average error of 41.17%. Significantly, the calibrated model correctly simulates nearly all of the major turning points (1977, 1984, 1985, 1993 and 1994).

Without the lower bound constraint on storage, both the releases and storage would have fluctuated more wildly under the ENPV preference model.

3.3 Using the SDP as a policy tool

This last section shows how the calibrated SDP framework can be used as a policy-oriented tool. There are several ways in which the model can lead to a better understanding of the optimal policy rule.

Figure 5: Optimal water release policy

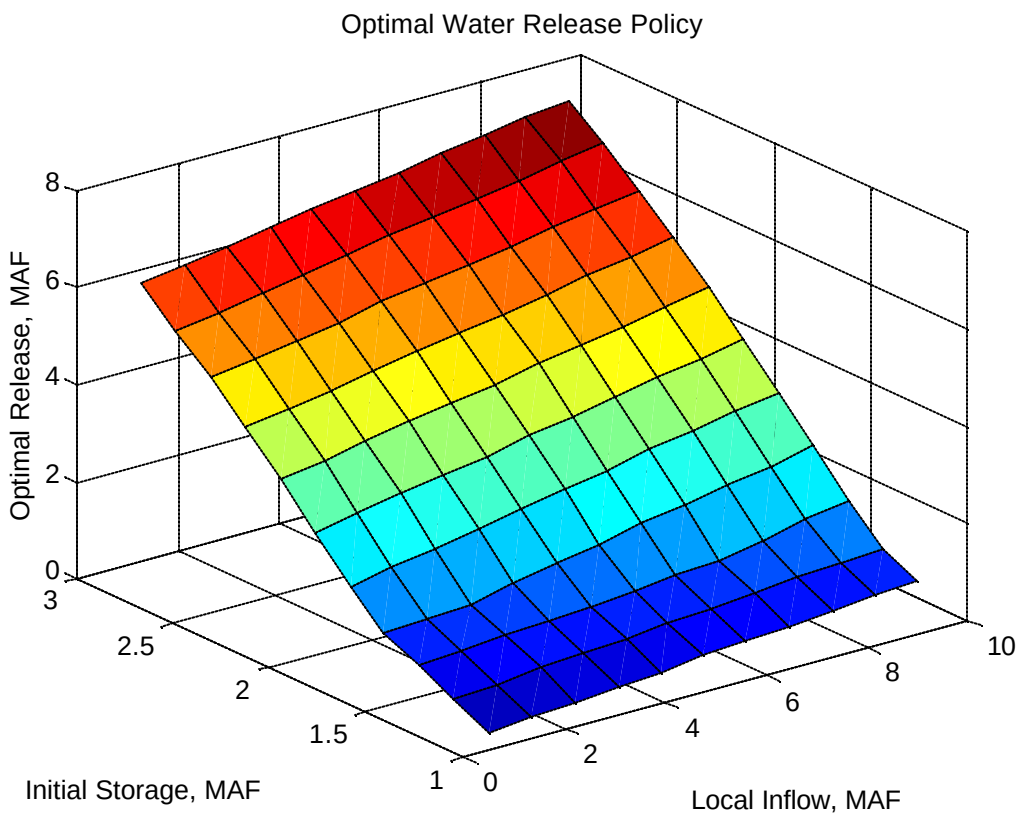

First, the SDP solution defines the optimal policy rule for the decision-makers. The response surface gives the optimal policy for the *current* operating environment known to the decision-maker, namely, storage level constraints, the uncertainty of inflows and the level of demands. For example, given the net benefit function (15) and the conditional distribution for the inflow from the rest of the system (13), Figure 5 gives the optimal water release as a function of the initial water storage and the local inflow realization. The level of storage in the reservoir seems to have a greater effect on the optimal release than the level of inflow. The shape of the surface is steeper as a function of initial

storage than as a function of local inflow. This is not at all surprising, given that a good portion of high inflow levels results in spillage, as they can neither be consumed nor stored.

Second, given that the calibrated parameters (subjective discount factor, constant relative risk aversion and constant resistance to intertemporal substitution) reflect the characteristics of the decision-maker, these parameters can be used to define the optimal policy in a marginally *modified* operating environment, such as a small exogenous shock affecting the water demand, or a change in the reservoir flood control limit. For example, let us assume that the maximum storage increases from the observed level 2.861 MAF to 3.5 MAF. Given the calibrated risk and intertemporal substitution preferences of the decision-maker ($r = -9.0$, $a = -1$, and $b = 0.893$), we can estimate the value function corresponding to these new parameter values and simulate the new optimal reservoir management policy over time. The simulation results are presented in Table 2 where we also compute the annual net benefits from water release and from the rest of the network inflow.

Table 2: Simulation of optimal policy, calibrated case

Year	Maximum Storage 2.86 MAF			Maximum Storage 3.5 MAF			Profit
	Release MAF	Storage MAF	Profit*	Release MAF	Storage MAF	Profit*	Change %
1974	6.00	2.86	3.62	5.46	3.50	3.58	-1.14
1975	4.38	2.86	2.79	4.38	3.50	2.79	0.00
1976	2.57	1.62	1.40	2.89	1.93	1.44	2.82
1977	1.64	0.99	0.89	1.96	0.99	0.93	4.75
1978	2.85	2.86	2.71	2.74	3.05	2.70	-0.40
1979	3.09	2.74	2.12	3.16	2.88	2.13	0.38
1980	4.36	2.86	2.96	3.92	3.50	2.92	-1.34
1981	3.98	2.86	2.18	3.98	3.50	2.18	0.00
1982	6.43	2.86	4.14	6.43	3.50	4.14	0.00
1983	8.30	2.86	4.69	8.30	3.50	4.69	0.00
1984	3.71	2.86	1.78	3.71	3.50	1.78	0.00
1985	2.90	2.14	2.70	3.17	2.50	2.72	0.98
1986	5.04	2.86	3.39	4.80	3.50	3.37	-0.57
1987	2.80	1.98	1.69	3.10	2.31	1.73	2.07
1988	2.26	1.38	1.54	2.44	1.55	1.56	1.41
1989	2.55	2.00	1.90	2.65	2.10	1.91	0.59
1990	2.27	1.39	1.39	2.34	1.44	1.40	0.59
1991	2.01	1.15	1.27	2.06	1.18	1.28	0.46
1992	1.89	1.04	1.30	1.93	1.06	1.31	0.36
1993	3.42	2.86	2.75	2.87	3.50	2.70	-1.87
1994	2.65	1.75	1.48	2.97	2.06	1.52	2.60
1995	6.91	2.86	3.82	6.66	3.50	3.80	-0.48
1996	6.23	2.86	3.40	6.23	3.50	3.40	0.00
Mean	**3.84**	**2.28**	**2.43**	**3.83**	**2.68**	**2.43**	**0.49**
Std. Dev.	**1.83**	**0.72**	**1.05**	**1.72**	**0.94**	**1.03**	**1.47**

*** Profit: profits (in billion $) current value.**

Table 2 shows that increasing the maximum storage limit from 2.86 MAF to 3.5 MAF (+23.7%) results in an average increase in the annual benefits that is less than 0.5%. The direct gains to be expected from an increase of the reservoir capacity are likely to be small. However, a risk-averse decision-maker will consider two other positive effects resulting from the reservoir capacity expansion. The first effect is the reduction of water release variability. The increase in the maximum storage limit allows additional smoothing of water consumption from year to year. The standard deviation of water releases is, on average, reduced from 1.83 to 1.72. This reduction of release variability results from the increase in average water storage levels from 2.28 to 2.68 MAF. This reduction in variability is valued by risk-averse decision-makers. Second, for a correct financial assessment of a discrete project such as this, it is important to know the distribution of the gains from reservoir expansion as well as the average benefit. In this example the average benefit increase corresponds to 0.5% but the stochastic realizations show that the benefits vary with a maximum change as high as 4.82%. Of course high levels of profits occur after drought years when water storage levels are low and the additional storage capacity is most useful.

We have run the same policy simulations in the case where the decision-maker maximizes the expected net present value. Complete results are available from the authors upon request. As expected, maximizing the expected net present value results in a less conservative management of the resource. The average storage of water goes down from 2.28 with calibrated parameters to 1.23 MAF (for a dam capacity equal to 2.86 MAF). Both average current profits and water releases increase. For a dam capacity equal to 3.5 MAF, they are respectively to $2.445 billion and a mean release of 4.01 MAF . Notice that to get these higher yields, the decision-maker has to accept a higher level of uncertainty. The standard deviation of releases and current profits are higher than in the calibrated case (2.17 versus 1.72 for releases and 1.08 versus 1.03 for profits).

The measure of the net benefit gains from the simulation results could be used in traditional cost / benefit analysis based on a single measure of expected net benefits from the reservoir expansion. However, using the stochastic properties and higher moments of the benefit stream enables comparison using more sophisticated financial instruments. The same type of computation could be done for any other parameter changes, such as a change in minimum storage, or a change in the distribution of inflows due to environmental restrictions on river operations.

4. Conclusion

The calibration of risk and substitution parameters add a positive component to the SDP approach that should help to reassure decision makers that their preferences are being incorporated into the basic model. In addition, the calibration process improves the fit of the model to historic data. The resulting model is more likely to be accepted as a policy tool than results from normative optimization. In addition, the empirical solution of the model uses the standard GAMS optimization routine, which means that the method is easily applied to different resource policy problems.

We have applied the calibrated SDP framework to model the management of Oroville Reservoir in California. The results show that inter-temporal substitution preferences have a significant impact on optimal policies and are more critical than the level of risk aversion or the discount factor. These empirical results underscore the importance of using a more general specification of inter-temporal preferences, such as recursive utility, in the objective function rather than relying solely on risk aversion to explain observed dynamic behavior under uncertainty.

References

BELLMAN, R. (1961): "Adaptive Control Processes: A Guided Tour", Princeton University Press.

BERTSEKAS, D.P. (1976): "Dynamic Programming and Stochastic Control", Academic Press, New York.

BURT, O., and J.R. ALLISON (1963): "Farm Management Decisions with Dynamic Programming," *Journal of Farm Economics*, 45(1): 1-22.

CALIFORNIA DEPARTMENT OF WATER RESOURCES (1974-96): "State Water Project Annual Report of Operations, " edited by the State of California.

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 Advance in Economic Theory, Vol II, edited by Jean-Jacques Laffont. Cambridge: Cambridge University Press, 1-63.

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.

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.

JENKINS, M.W., et al. (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.

KING, R., C. PLOSSER & S. REBELO (1988): "Production Growth and Business Cycles I: The Basic Neoclassical Model," *Journal of Monetary Economics*, 21: 195-232.

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.

KYDLAND, F.E. & E.C. PRESCOTT (1991): " The Econometrics of the General Equilibrium Approach to Business Cycles, " *Scandinavian Journal of Economics*, 93 (2):161-178.

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.

NORMANDIN, M., and P. SAINT-AMOUR (1998) "Substitution, Risk Aversion, Taste Shocks and Equity Premia," Journal of Applied Econometrics, 13(3), 265-291.

PELTOLA, J, and K.C. KNAPP (2001): " Recursive Preferences in Forest Management," *Forest Science*, Vol 47, 5: November 2001, 455-465.

PROVENCHER , B, and R.C. BISHOP (1997): "Estimable Dynamic Model of Recreation Behavior with an Appication 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. (1994): "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.