Renewable resource management with stochastic recharge and environmental threats

Arie Leizarowitz* Yacov Tsur$^{\diamond}$

August 2010

Abstract

Exploitation diminishes the capacity of renewable resources to withstand environmental stress, increasing their vulnerability to extreme conditions that may trigger abrupt changes. The onset of such events depends on the coincidence of extreme environmental conditions (environmental threat) and the resource state (determining its resilience). When the former is uncertain and the latter evolves stochastically, the uncertainty regarding the event occurrence is the result of the combined effect of these two uncertain components. The environmental threat renders the single-period discount factor policy-dependent and, as a result, the compound discount factor becomes history-dependent. We study optimal management in such a setting. Existence of an optimal Markovian-Deterministic stationary policy is established and the optimal state process is shown to converge to a steady state distribution. A numerical example illustrates these properties.

JEL classification: C62, Q20

Keywords: Stochastic stock dynamics, catastrophic event, endogenous discounting, Markov decision process, optimal stationary policy.

*Department of Mathematics, Technion, Haifa 32000, Israel. YT: Professor Leizarowitz passed away unexpectedly after the completion of this work. He was a friend, a colleague and a tutor. I miss him deeply.

$^{\diamond}$Department of Agricultural Economics and Management, The Hebrew University of Jerusalem, POB 12, Rehovot 76100, Israel. Email: tsur@agri.huji.ac.il.

1 Introduction

We study management of renewable resources with stochastic state evolution and environmental uncertainty regarding the occurrence of an abrupt catastrophic event. The effects on management policies of these two uncertain processes are highly intertwined, as the vulnerability of a resource to (uncertain) environmental stress depends critically on its (stochastic) state. Admittedly, numerous uncertain elements prevail in any given resource situation and the literature addresses many of them (see Pindyck 2007). But the combined effect of stochastic state evolution and uncertain abrupt change (or regime shift) has not been thoroughly addressed so far.

The economic literature on natural resources with stochastic state dynamics mostly ignores uncertain catastrophic events such as abrupt regime shift or ecological collapse (see Burt 1964, Reed 1974, 1979, Pindyck 1984, Knapp and Olson 1995, Pindyck 2002, Costello et al. 2001, Sethi et al. 2005, Singh et al. 2006, Mitra and Roy 2006, Wirl 2007, McGough et al. 2009, and references they cite). Some works incorporate deterministic thresholds, e.g., project investment thresholds (Pindyck 2002), extinction thresholds (Mitra and Roy 2006) and temperature thresholds (Wirl 2007), so the uncertainty emanates only from the stochastic stock dynamics. Other works allow for uncertain regime shift, such as extinction of a fishery population (Roughgarden and Smith 1996, Sethi et al. 2005, McGough et al. 2009)[1], but fall short of modeling it as a regime shift in which the extinction occurrence changes the rules of the game, since both the fishery stock and growth rate are known with certainty to equal zero from the extinction date onward. When the regime shift

[1]The uncertainty in the extinction thresholds stems from the inaccurate stock measurement, introduced by Clark and Kirkwood (1986).

is properly modeled, it turns the discount factor endogenous and this feature is consequential for the optimal policy and the ensuing steady state distribution.

The sudden occurrence of catastrophic events (regime shifts, abrupt changes) in renewable resource situations is related to nonlinear phenomena such as positive feedbacks, hysteresis and the presence of uncertain thresholds that are prevalent in environmental processes (Dasgupta and Mäler 2003, Brock and Starrett 2003). Examples include pollution-induced catastrophes (Cropper 1976, Clarke and Reed 1994, Aronsson et al. 1998, Tsur and Zemel 1998), a sudden collapse of an ecosystem or of animal and plant populations (Clark and Kirkwood 1986, Reed 1989, Tsur and Zemel 1994, Brock and Xepapadeas 2003), destruction of coastal aquifers due to seawater intrusion (Tsur and Zemel 1995, 2004), phosphorus loading into lakes inducing an irreversible transition from an oligotrophic (clear) state to a eutrophic (turbid) state (Harper 1992, Carpenter et al. 1999, Mäler 2000), and global-warming induced catastrophes (Tsur and Zemel 1996, 2009, Broecker 1997, Mastrandrea and Schneider 2001, Alley et al. 2003, Nævdal 2006, Haurie and Moresino 2006, Roe and Baker 2007, Stern 2007, Bahn et al. 2008, Weitzman 2009).[2] This literature strain assumes a deterministic evolution of the resource state.

The most pronounced effect on resource management policies of the presence of a catastrophic threat shows up in the discount factor, which becomes policy- and history-dependent. Implications of this property for climate policies under threats of global warming induced catastrophes have recently been studied by Tsur and Zemel (2008, 2009) in a deterministic resource evolu-

[2]The abrupt change may be desirable, as in Bahn et al. (2008) who consider two such events: the resolution of uncertainty regarding climate sensitivity and technological breakthrough regarding a carbon-free energy source.

tion framework.[3] Here we consider stochastic state dynamics in a general renewable resource situation. The endogeneity of the discount factor requires extending properties of Markov decision processes (MDPs), known to hold under constant discounting (see, e.g., Puterman 2005), to the present case. In particular, we establish the existence of an optimal stationary Markovian-deterministic policy and show that the optimal state process converges in the long-run to a well specified steady-state distribution. The first result implies that the search for optimal policy rules can be confined to the (relatively simple) set of stationary Markovian-deterministic policies. The steady-state distribution of the optimal stock process provides a useful reference according to which simple (even if suboptimal) management policies can be designed to avoid or reduce catastrophic threats.

The resource setup, with the stochastic state dynamics and the environmental threat, is formulated in Section 2. Section 3 formulates the resource management problem. Existence of an optimal, Markovian-deterministic, stationary policy (under the history-dependent discount factor) is established in Section 4. Asymptotic (long-run) behavior of the optimal state process is characterized in Section 5. A numerical illustration is presented in Section 6. Section 7 concludes and the appendixes contain technical details.

2 Resource setup

We consider discrete time, state and action spaces. The discrete time formulation reflects the cyclical (seasonal, annual) nature common to most

[3] There is a parallel line of macroeconomics literature, stemming from Yaari's Yaari (1965) uncertain lifetime model and its perpetual-youth-model extension (see Acemoglu 2009, p. 345 for a discrete time version), in which the discount factor is affected by an event occurrence hazard (i.e., death). The event hazard in this literature, however, is exogenous.

renewable resources. The assumption of discrete state and action spaces (finite or countable) is an abstraction, which can be viewed as an approximation of the continuous case. Practically, there are pros and cons to both approaches. On a more philosophical level, a model, by definition, is a simplification (of what it intends to model) and should be judged on the basis of its capacity to enhance our understanding of the phenomenon under study. In the present case, the consideration of discrete state and action spaces simplifies the exposition and allows for a sharper results.

2.1 States, actions and recharge

The state of the resource system at the beginning of period t is denoted $S_t = (S_t^1, S_t^2, ..., S_t^M)'$, where S_t^m is the m'th stock, $m = 1, 2, ..., M$. The resource evolves in time according to

$$S_{t+1} = S_t + R(S_t) + X_t - g_t, \; t = 1, 2, \dots, \tag{2.1}$$

where $R(S_t) = (R^1(S_t), R^2(S_t), ..., R^M(S_t))'$, $X_t = (X_t^1, X_t^2, ..., X_t^M)'$ and $g_t = (g_t^1, g_t^2, ..., g_t^M)'$ are M-dimensional vectors representing deterministic recharge, stochastic recharge and exploitation (harvest, extraction) rates, respectively.[4]

The initial time period is $t = 1$ and the initial state vector, S_1, is given. In each time period $t = 1, 2, ...$, the resource state S_t is observed. Based on S_t, the action g_t is chosen, the deterministic recharge (growth) $R(S_t)$ is determined and the stochastic recharge X_t is realized, giving rise to S_{t+1}, according to (2.1).

[4] The stochastic recharge (X_t) in the state evolution equation (2.1) enters additively. Using the multiplicative specification commonly used in many fishery models (see, e.g., Reed 1979, Roughgarden and Smith 1996, Sethi et al. 2005, McGough et al. 2009) changes the formulation of the transition probabilities below, but otherwise has no effect on the general structure.

The discrete (finite or countable) state, recharge and action spaces are denoted $\mathcal{S}$, $\mathcal{X}$ and $\mathcal{A}$, respectively. Thus, $\mathcal{S} = \{s_1, s_2, ..., s_{n_s}\}$, where $s_j \in \mathrm{I\!R}^M$, $j = 1, 2, ..., n_s$ and n_s (possibly infinite) is the number of states. With $\mathcal{X}^m(s)$ representing the support of stock m's recharge distribution at state $s \in \mathcal{S}$ and $\mathcal{X}(s) = \mathcal{X}^1(s) \times \mathcal{X}^2(s) \times \cdots \times \mathcal{X}^m(s)$, the admissible recharge support is $\mathcal{X} = \bigcup_{s\in\mathcal{S}} \mathcal{X}(s) = \{x_1, x_2, ..., x_{n_x}\}$, containing n_x (possibly infinite) feasible recharge vectors $x_j \in \mathrm{I\!R}^M_+$. The recharge probability at time t, given $S_t = s$, is denoted $p_{x|s}(\cdot)$, i.e.,

$$p_{x|s}(x) \equiv Pr\{R(S_t) + X_t = x | S_t = s\}. \tag{2.2}$$

In a similar manner we let $\mathcal{A}^m(s)$ consist of stock m's actions (exploitation rates) feasible at state $s \in \mathcal{S}$ and let $\mathcal{A}(s) = \mathcal{A}^1(s) \times \mathcal{A}^2(s) \times \cdots \times \mathcal{A}^M(s)$. The admissible action space is $\mathcal{A} = \bigcup_{s\in\mathcal{S}} \mathcal{A}(s) = \{a_1, a_2, ..., a_{n_a}\}$, where $a_j \in \mathrm{I\!R}^M$ and n_a is the number of actions (finite or countable). An action $g_t = (g_t^1, g_t^2, ..., g_t^M)'$ corresponds to exploiting (harvesting, extracting) source m at the rate g_t^m, $m = 1, 2, ..., M$, during time period t. The information available when period t's action is chosen is $H_t = \{S_1, g_1, ..., S_{t-1}, g_{t-1}, S_t\}$. The action is feasible if $g_t \in \mathcal{A}(S_t)$.

2.2 Environmental threat

The resource system is under risk of an abrupt shock (regime shift) with undesirable consequences. The conditions that trigger such events depend on the resource state and exploitation policy and are uncertain due to genuine environmental uncertainty. There is a subtle distinction between environmental threat in the form of a catastrophic event whose occurrence depends on genuine environmental uncertainty, and that associated with crossing an unknown threshold (Tsur and Zemel 2004, discuss these two types of environemntal threat in the context of groundwater management with deterministic recharge). The uncertainty in the latter case is mostly due to our own ignorance of the triggering threshold and there is plenty of room to learn from experience (as we "test the waters" and find that the world did not come to an end we gain new information about the threshold). Here we consider the former case and the stochastic nature of the environmental threat is represented by the survival function λ.[5]

We denote by κ the catastrophic state of the resource system and let $1 - \lambda(s,a)$ be the hazard probability to end up in κ at time $t+1$ when occupying state $s \neq \kappa$ and employing action a at time t. Let T denote the time period at which the event occurs. Then,

$$Pr\{T = \tau\} = [1 - \lambda(S_\tau, g_\tau)] \prod_{j=1}^{\tau-1} \lambda(S_j, g_j),\ \tau = 1, 2..., \tag{2.3}$$

where we use the convention that $\prod_{j=1}^{\tau-1} \lambda(S_j, g_j) = 1$ for $\tau = 1$. The event occurrence probability (2.3) represents the environmental uncertainty conditional on the resource state trajectory and exploitation policy. The combined effect of the event uncertainty and the stochastic evolution of the resource state shows up in the resource transition probabilities, specified next.

2.3 Transition probabilities

Let $p(j|i,a)$ represent the probability of occupying state s_j at time $t+1$ conditional on $S_t = s_i$, $g_t = a$ and $T > t$ (i.e., that the event will not interrupt):

$$p(j|i,a) = Pr\{S_{t+1} = s_j | S_t = s_i, g_t = a, T > t\}.$$

[5] An interesting future extension would be to consider a Knightian uncertainty, e.g., by assuming that the event occurrence hazard is known up to a (subjective) probability and specifying an updating learning process as new information comes along (see Epstein and Schneider 2007, Vardas and Xepapadeas 2010, for a possible approach).

In view of (2.1)-(2.2),

$$p(j|i,a) = p_{x|s_i}(s_j - s_i + a). \tag{2.4}$$

We let P_a represent the $n_s \times n_s$ matrix with $p(j|i,a)$ as the (i,j) element.

Given that the event has not occurred by time $t-1$, the probability during time t of moving from s_i to s_j and of nonoccurrence is

$$\begin{aligned} q(j|i,a) &\equiv Pr\{S_{t+1} = s_j, T > t | S_t = s_i, g_t = a\} \\ &= Pr\{S_{t+1} = s_j | S_t = s_i, g_t = a, T > t\} Pr\{T > t | T > t-1, S_t = s_i, g_t = a\} \\ &= p(j|i,a)\lambda(s_i, a). \end{aligned} \tag{2.5}$$

We denote by Q_a the $n_s \times n_s$ matrix with the (i,j) element given by $q(j|i,a)$.

3 Management policies and welfare

We begin by formulating rewards (single-period) and payoffs. The decision rules and policies are explained next and subsection 3.3 presents the welfare criterion.

3.1 Rewards and payoffs

If the event does not occur during time period t, while the resource is at state S_t and the action g_t is undertaken, period t's reward $\tilde{b}(S_t, g_t)$ is obtained, whereas if the event occurs the post-event value $v^p(S_t)$ is acquired. The latter represents the present-value, under the optimal post-event policy, of the benefit flow from the occurrence time onwards, discounted to the beginning of the occurrence period. We assume that $\tilde{b}(s,a)$ and $v^p(s)$ are bounded and that the latter is smaller than the pre-event value (defined below), as we consider undesirable events.

With $\beta \in [0, 1)$ representing the (constant) discount factor, the (uncertain) payoff is

$$\sum_{t=1}^{T-1} \tilde{b}(S_t, g_t)\beta^{t-1} + v^p(S_T)\beta^{T-1}. \tag{3.1}$$

Noting (2.3), the expected payoff (with respect to the event occurrence time T) is

$$\sum_{\tau=1}^{\infty} \left(\sum_{t=1}^{\tau-1} \tilde{b}(S_t, g_t)\beta^{t-1} + v^p(S_\tau)\beta^{\tau-1} \right) [1 - \lambda(S_\tau, g_\tau)] \prod_{j=1}^{\tau-1} \lambda(S_j, g_j) =$$
$$\sum_{\tau=1}^{\infty} \sum_{t=1}^{\tau-1} \tilde{b}(S_t, g_t)\beta^{t-1}[1 - \lambda(S_\tau, g_\tau)] \prod_{j=1}^{\tau-1} \lambda(S_j, g_j) +$$
$$\sum_{\tau=1}^{\infty} v^p(S_\tau)\beta^{\tau-1}[1 - \lambda(S_\tau, g_\tau)] \prod_{j=1}^{\tau-1} \lambda(S_j, g_j). \tag{3.2}$$

By changing the order of summation (permitted when $\tilde{b}$ is bounded), the first term on the right-hand side above is expressed as

$$\sum_{t=1}^{\infty} \tilde{b}(S_t, g_t)\beta^{t-1} \sum_{\tau=t}^{\infty} \left([1 - \lambda(S_\tau, g_\tau)] \prod_{j=1}^{\tau-1} \lambda(S_j, g_j) \right). \tag{3.3}$$

The inner sum above equals

$$\sum_{\tau=t}^{\infty} \left(\prod_{j=1}^{\tau-1} \lambda(S_j, g_j) - \prod_{j=1}^{\tau} \lambda(S_j, g_j) \right) = \prod_{j=1}^{t-1} \lambda(S_j, g_j),$$

which upon substituting back in (3.3) gives

$$\sum_{t=1}^{\infty} \left(\tilde{b}(S_t, g_t) \prod_{j=1}^{t-1} \beta\lambda(S_j, g_j) \right). \tag{3.4}$$

This expression is the present value of the benefit flow $\tilde{b}(S_t, g_t)$ discounted with the history-dependent discount factor

$$\gamma(t) = \begin{cases} 1 & t = 1 \\ \prod_{j=1}^{t-1} \beta\lambda(S_j, g_j) & t = 2, 3, ... \end{cases}, \tag{3.5}$$

corresponding to the running (single period) discount factor $\beta\lambda(S_t, g_t)$.

The second term on the right-hand side of (3.2) is expressed as

$$\sum_{t=1}^{\infty} v^p(S_t)[1 - \lambda(S_t, g_t)]\gamma(t). \tag{3.6}$$

Combining (3.4) and (3.6), the expectation of the payoff with respect to event occurrence time T is give by

$$\sum_{t=1}^{\infty} b(S_t, g_t)\gamma(t), \tag{3.7}$$

where

$$b(S_t, g_t) \equiv \tilde{b}(S_t, g_t) + v^p(S_t)[1 - \lambda(S_t, g_t)]. \tag{3.8}$$

The catastrophic environmental threat affects the payoff in two ways: First, it changes period t's benefit from $\tilde{b}(S_t, g_t)$ to $b(S_t, g_t)$. Second, it changes the running (single period) discount factor from the constant β to the state-and-action-dependent discount factor $\beta\lambda(S_t, g_t)$. The latter effect is twofold: first, it decreases the discount factor ($\beta\lambda(s, a) \leq \beta$ since $\lambda(s, a) \leq 1$), thereby inducing less conservation (since the future is discounted more heavily); second, it turns the discount factor endogenous to the exploitation policy. The policy implications of these effects were studied by Tsur and Zemel Tsur and Zemel (2008, 2009) in a deterministic state evolution model of climate-change induced catastrophes. Here they are studied in the context of a stochastic state evolution.

3.2 Decision rules and policies

A decision rule $d_t(\cdot)$ determines the action at time t given the available information $\{S_t, S_{t-1}, S_{t-2}...\}, \{g_{t-1}, g_{t-2}, ...\}$. It may be history-dependent or Markovian (depends only on the current state S_t), randomized or deterministic. Consequently, the four types of decision rules are history-dependent and randomized (HR), history-dependent and deterministic (HD), Markovian and randomized (MR), Markovian and deterministic (MD). A policy (or plan) specifies the decision rules for all time periods, $\pi = \{d_1, d_2, ...\}$, and is classified as HR, HD, MR or MD depending on the type of the decision rules d_t, $t = 1, 2, ...$ A policy is stationary if the same decision rule is repeated in all time periods, i.e., $d_t(\cdot) = \varphi(\cdot)$ for all $t = 1, 2....$ (Thus, a stationary policy is necessarily Markovian.)

The HR class of policies is the widest and contains all other classes as special cases, while the MD class is contained in all other classes. Within the MD class, stationary policies are the simplest, hence the most attractive for actual implementations.

3.3 Welfare

Under a Markovian policy $\pi = \{d_1, d_2, ...\}$, with $g_t = d_t(S_t)$, the (random) payoff, noting (3.7), is

$$\sum_{t=1}^{\infty} b(S_t, d(S_t))\gamma(t)$$

and the expected payoff given the initial state $S_1 = s$ is

$$v^{\pi}(s) = E^{\pi}\left\{\sum_{t=1}^{\infty} b(S_t, d_t(S_t))\gamma(t)\right\}. \tag{3.9}$$

The welfare (value) function is defined as

$$v^*(s) = \sup_{\pi \in \Pi^{\mathrm{HR}}} v^{\pi}(s), \;\; s \in \mathcal{S}. \tag{3.10}$$

4 Optimal policy

The optimal policy π^*, when exists, satisfies $v^{\pi^*}(s) = v^*(s)$ for all $s \in \mathcal{S}$. We denote by $v^{\varphi}(s)$ the value corresponding to the stationary policy $\pi =$

$(\varphi, \varphi, \cdots)$. As Markovian-Deterministic (MD) stationary policies are attractive for practical purposes, it is of interest to know if an optimal MD stationary policy exists, i.e., if the value v^* can be attained by an MD stationary policy. For standard Markov Decision Processes (MDPs), with a constant discount factor, the answer is in the affirmative (see, e.g., Puterman 2005, Chapter 6). Here, however, the environmental threat (catastrophic event, regime shift) turns the running (one period) discount factor $\beta\lambda(s_i, a_i)$ policy-dependent, implying that the compound discount factor $\gamma(t)$ is history-dependent (cf. equation (3.5)) and undermining the validity of this result (an example in which there is a history-dependent discount factor and where there exists no optimal MD stationary policy is presented in Appendix D).

Nonetheless, we verify that in the present case an optimal MD stationary policy does exist and specify (in Section 5) the steady state distribution to which the optimal state process converges in the long run. It turns out that the history-dependent discount factor in the present case is of a specific form that allows to specify the unconditional transition matrix Q_a (defined in (2.5)) and, in turn, replicate the analysis of the standard, constant discount factor case. For a general history-dependent discount factor we cannot perform this reduction, and the example (Appendix D) exhibits a situation where this result is false.

The existence property is stated in the theorem below. An extended version of the theorem is proven in Appendix A.

Theorem 4.1. *Suppose* (*A*1) $0 \leq \beta < 1$, (*A*2) $\mathcal{S}$ *is discrete* (*finite or countable*), (*A*3) $\tilde{b} : \mathcal{S} \times \mathcal{A} \mapsto \mathbb{R}$ *and* $v^p : \mathcal{S} \mapsto \mathbb{R}$ *are bounded and* (*A*4) $\tilde{b}(s_i, a)$ *and* $\lambda(s_i, a)p(j|i, a)$ *are continuous in a, and* $\mathcal{A}(s_i)$ *is compact for all* $s_i, s_j \in \mathcal{S}$.

Then, there exists an optimal, Markovian-Deterministic stationary policy φ^*, i.e., the policy $(\varphi^*, \varphi^*, ...)$ satisfies

$$v^{\varphi^*}(s) = v^*(s) \ \forall s \in \mathcal{S}. \tag{4.1}$$

The Theorem allows confining attention to Markovian-Deterministic stationary policies, for which a variety of algorithms exists (see Judd 1998, Puterman 2005). In the numerical example of Section 6 we calculate the optimal policy using an algorithm based on Linear Programming (see Puterman 2005, Chapter 6.9), adopted to the present case of a policy-dependent discount factor.

5 Long-run behavior

In this section we verify that the optimal state process converges in the long-run to a steady-state distribution and characterize this distribution. We also specify the event-occurrence probability for each initial state. To simplify the exposition we confine attention to the finite state case.

Recalling equations (2.4)-(2.5), $P_{\varphi^*}(i,j) = p(j|i, \varphi^*(s_i))$ gives the probability that the resource system moves from $S_t = s_i$ to $S_{t+1} = s_j$ when the optimal policy $g_t = \varphi^*(s_i)$ is employed, conditional on the event not occurring during period t. The unconditional transition probabilities are $Q_{\varphi^*}(i,j) = \lambda_i^* P_{\varphi^*}(i,j)$, where

$$\lambda_i^* \equiv \lambda(s_i, \varphi^*(s_i)), \ i = 1, 2, ..., n_s, \tag{5.1}$$

is the survival (nonoccurrence) probability under the optimal policy.

The transition matrix P_{φ^*} classifies each state in $\mathcal{S}$ as either recurrent or transient.[6] We denote by E_0 the subset containing the n_0 transient states.

[6]We assume that P_{φ^*} is aperiodic, which is the common case.

The recurrent states can be arranged in K irreducible subsets E_k, each containing n_k states, $k = 1, 2, ..., K$.[7] Recurrent, irreducible subsets are absorbing, i.e., once the state process enters E_k it stays there forever. We denote by P_k the $n_k \times n_k$ submatrix of P_{φ^*} corresponding to the states contained in E_k, $k = 0, 1, ..., K$.

It is convenient at this point to rearrange the states such that the transient states are the first n_0 states, the states in E_1 constitute the next n_1 states and so on. Thus, $\mathcal{S} = \bigcup_{k=0}^{K} E_k$ and $\mathcal{S} \bigcup \{\kappa\}$ is the state space containing also the (recurrent, absorbing) occurrence state κ.

A state $s_i \in \mathcal{S}$ is called "safe" or "unsafe" depending on whether $\lambda_i^* = 1$ or $\lambda_i^* < 1$, respectively. The subset

$$\mathcal{S}_1 = \{s_i \in \mathcal{S} | \lambda_i^* = 1\} \tag{5.2}$$

contains all "safe" states. ($\mathcal{S}_1$ may well be empty.)

If a recurrent subset E_k contains no "unsafe" states, i.e., $E_k \subseteq \mathcal{S}_1$, then entering E_k ensures that the event will never occur. This is so because the probability that the event will occur during period t given $S_t = s_i \in E_k$ is $1 - \lambda_i^* = 0$ for any $s_i \in E_k$ and E_k is absorbing. For recurrent, irreducible sets containing only "safe" states we define the limiting matrix[8]

$$\hat{P}_k = \lim_{\tau \to \infty} P_k^\tau. \tag{5.3}$$

The (i, j) element of $\hat{P}_k$ represents the probability that in the long run the system will occupy state s_j when it starts at state s_i and the optimal policy is employed for any $s_j \in E_k$. Clearly, $\hat{P}_k$ satisfies $\hat{P}_k P_k = \hat{P}_k$ (taking one extra

[7]The subset $E_k \subset \mathcal{S}$ is closed if $Pr\{S_{t+\tau} = s_j | S_t = s_i, \varphi^*(\cdot)\} = 0$ for any $s_i \in E_k$ and $s_j \notin E_k$, $\tau = 1, 2, ...$ The subset E_k is irreducible if no proper subset of it is closed.

[8]The limit exists since P_k is aperiodic.

step cannot change the limiting behavior), implying that $\hat{P}_k$ has identical rows $\hat{q}_k{}'$, where $\hat{q}_k{}' \in \mathbb{R}_+^{n_k}$ is the unique solution of the equation (see Puterman 2005, p. 592):

$$q' = q' P_k \text{ subject to } \sum_{j=1}^{n_k} q_j = 1. \tag{5.4}$$

Let $\hat{p}_k{}' = (0, ..., 0, \hat{q}_k{}', 0, ..., 0)$ be the n_s-dimensional vector with $\hat{q}_k$ at the n_k elements corresponding to $s_i \in E_k$ and 0 elsewhere. Then, when the state process departs from a recurrent set $E_k \subseteq \mathcal{S}_1$, the event occurrence probability (the probability to enter the occurrence state κ) is zero and the optimal state process converges in the long run to the steady state distribution $\hat{p}_k$.

Departing from a recurrent subset containing at least one "unsafe" state (s_u, say), implies that the event will (eventually) occur with probability one. This is so because each time the "unsafe" state s_u is visited an occurrence probability of $1 - \lambda_u^* > 0$ is inflicted and (once in $E_k \nsubseteq \mathcal{S}_1$) visits to s_u never stop prior to the event occurrence.[9] It follows that the limiting probability of all $s_i \in \mathcal{S}$ vanish and the limiting probability of κ (the occurrence state) is one. We summarize the above discussion in:

Proposition 5.1. *Suppose the state process departs from one of the recurrent sets E_k, $k = 1, 2, ..., K$.*

(i) If $E_k \subseteq \mathcal{S}_1$, then the event-occurrence probability is zero and the optimal state process converges in the long run to the steady state distribution $\hat{p}_k$.

(ii) If $E_k \nsubseteq \mathcal{S}_1$, then the long-run event-occurrence probability (the limiting probability of the occurrence state κ) is 1 and the long-run probabilities of all

[9]Suppose, without loss of generality, that s_u is the only "unsafe" state in E_k and notice that, unless interrupted by the event, the recurrent state s_u will be visited infinite number of times with probability one. Occurrence may happen on the first visit with probability $1 - \lambda_u^*$ or on the second visit with probability $\lambda_u^*(1 - \lambda_u^*)$ or on the third visit with probability $\lambda_u^{*2}(1 - \lambda_u^*)$ and so on. Summing all possibilities gives the occurrence probability $(1 - \lambda_u^*) \sum_{j=0}^{\infty} (\lambda_u^*)^j = 1$.

states in $\mathcal{S}$ vanish.

Suppose now that the state process departs from a transient state $s_j \in E_0$. The optimal state process must eventually exit E_0 to one of the recurrent sets $E_k,\ k = 1, 2, ..., K$, or to the event-occurrence set $E_{K+1} \equiv \{\kappa\}$ – a recurrent, absorbing set on its own. To specify the probability of each of these possibilities, let $\tilde{Q}_0^1$ be the $n_0 \times (K+1)$ matrix whose (j, k) elements equal the *one-period* probability of moving from $s_j \in E_0$ to a state in $E_k,\ k = 1, 2, ..., K + 1$:

$$\tilde{Q}_0^1(j, k) \equiv Pr\{S_{t+1} \in E_k | S_t = s_j \in E_0, \varphi^*(s_j)\} = \sum_{s_i \in E_k} Q_{\varphi^*}(j, i), k = 1, 2, ..., K,$$

and

$$\tilde{Q}_0^1(j, K + 1) \equiv Pr\{S_{t+1} = \kappa | S_t = s_j \in E_0, \varphi^*(s_j)\} = 1 - \lambda_j^*.$$

Define

$$\tilde{Q}_0 = (I - Q_0)^{-1} \tilde{Q}_0^1, \tag{5.5}$$

where Q_0 is the $n_0 \times n_0$ submatrix of Q_{φ^*} corresponding to the n_0 transient states. We verify in Appendix B that (i) $\tilde{Q}_0$ exists and (ii) when departing from $s_j \in E_0$, the probabilities that the optimal state process will exit the transient set E_0 into $E_k,\ k = 1, 2, ..., K$, are given by $\tilde{Q}_0(j, k),\ k = 1, 2, ..., K$, and the probability that it will exit E_0 into κ equals $\tilde{Q}_0(j, K + 1)$. Consequently, noting Proposition 5.1, when the state process departs from a transient state $s_j \in E_0$, the steady-state distribution of states in $\mathcal{S}$ is given by

$$\sum_{k=1}^{K} \tilde{Q}_0(j, k) \hat{p}_k \tag{5.6a}$$

and the steady state probability of the event occurrence state κ is

$$\sum_{E_k \not\subseteq \mathcal{S}_1} \tilde{Q}_0(j, k), \tag{5.6b}$$

where the sum in (5.6b) extends over $k = 1, 2, ..., K + 1$. We summarize the above discussion in:

Proposition 5.2. *When the state process departs from a transient state $s_j \in E_0$, the optimal state process converges in the long run to the steady state distribution specified in (5.6a)-(5.6b).*

Together, Propositions 5.1 and 5.2 establish the convergence of the optimal state process to a well-specified steady-state distribution. This long-run distribution provides a reference by which to evaluate the actual state of the resource – depending on how far off the actual state distribution has been from the optimal long-run distribution. Such information is particulary useful in the present context, as the catastrophic threat may rule some resource states prohibitive when their long-run probabilities vanish.

6 A numerical illustration

The Kinneret water basin (Lake Kinneret is also known as Lake Tiberias or the Sea of Galilee) is the largest of Israel's natural water sources, providing over 30 percent of the country's natural water supply on average. Like other moderately shallow lakes[10] (Harper 1992, Mäler 2000), it faces a threat of abrupt ecosystem collapse as the pollution loading may trigger a eutrophication process.[11] The risk of such abrupt regime-shift depends on the lake's water head (stock). This property, together with the highly volatile recharge process (Fig-

[10] Lake Kinneret's maximal and average water depths are 46 m and 25 m, respectively (Gvirtzman 2002, p. 34).

[11] A lower water-head raises the concentration of nutrients at the top layer and, in turn, increases algal activity. An aggressive algal bloom may trigger a eutrophication process (see Serruya and pollingher Serruya and Pollingher (1977) and Gvirtzman (Gvirtzman 2002, pp. 43-55)).

ure 1), render the above framework particularly suitable for demonstrating our analysis.

In the next subsection we describe the basin's recharge process and derive its distribution. Subsection 6.2 defines states and actions and subsection 6.3 derives the ensuing transition probabilities. The rewards are specified in subsection 6.4, paying special attention to the catastrophic threat associated with over-exploitation.[12] In subsection 6.5 we apply an algorithm based on Linear Programming (LP) for solving Markov decision Processes (MDPs) and derive the optimal policy and value (the algorithm is described in Appendix C). Finally, the steady state distribution under the optimal policy is calculated in subsection 6.6.

6.1 Recharge process

Figure 1 presents the Kinneret's net (accounting for evaporation) annual recharge for the period 1932 - 2008. We use the gamma distribution to approximate the recharge distribution, i.e., we assume that the recharge series consists of iid draws from a gamma distribution with parameters α and θ, satisfying

$$\alpha\theta = \text{ Mean(recharge) - Min(recharge) } = 570.38 - 157 = 413.38 \text{ MCMY}$$

and

$$\alpha\theta^2 = \text{ Var(recharge) } = 77333.8,$$

where MCMY stands for million m^3 per year (the mean, min and standard deviation of the recharge series are displayed in Figure 1). We obtain $\alpha =$

[12] All benefit and cost calculations are based on ad hoc assumptions regarding the derived demand for water and the cost of water supply, made for the illustration purposes only and should not be given any empirical connotations.

2.20967 and $\theta = 187.077$. Figure 2 depicts the empirical distribution of the recharge series (dots) and the gamma distribution with the above (α, θ) parameters.

Figure 1

Figure 2

The support of the recharge distribution is denoted $\mathcal{X} = \{x_1, x_2, ..., x_{n_x}\}$, with $x_1 = 150$ MCMY (the minimal recharge realization – see Figure 1), $x_{n_x} = 1450$ MCMY (approximately the maximal recharge realization) and $x_{\ell+1} - x_\ell = \Delta_x$, $\ell = 1, 2, ..., n_x - 1$. Thus,

$$x_\ell = 150 + (\ell - 1)\Delta_x, \ \ell = 1, 2, ..., n_x, \tag{6.1}$$

and $p_{x|s}(x_\ell)$ is calculated as

$$p_{x|s}(x_\ell) = \begin{cases} F(x_\ell + \Delta_x/2) & \text{if } \ell = 1 \\ F(x_\ell + \Delta_x/2) - F(x_\ell - \Delta_x/2) & \text{if } 2 \leq \ell \leq n_x - 1 \\ 1 - F(x_\ell - \Delta_x/2) & \text{if } \ell = n_x \end{cases} \tag{6.2}$$

where $F(\cdot)$ is the gamma distribution specified above (and depicted in Figure 2). Since n_x and Δ_x are related according to $x_{n_x} = x_1 + (n_x - 1)\Delta_x$, setting one parameter determines the other. Setting $\Delta_x = 50$ MCMY implies $n_x = 15$.

6.2 States and actions

The Kinneret water-head ranges between the altitudes 208.8 and 215 meter below sea level (-208.8 m and -215 m, respectively). Above the upper water-head (-208.8 m) the water overflows the lake's edges (flooding is avoided by opening the gates of the Degania dam at the southern outlet of the lake leading into the lower Jordan river). The lower altitude (-215 m) is the minimal water head level at which water can be pumped (due to pumping infrastructure) and is designated as the black line.[13] In between there is the so-called red line – an imaginary water-head level indicating a critical water stock below which the above-mentioned catastrophic risk increases sharply. The red line is set at -213 m.[14]

The water stock corresponding to the black line is normalized at zero and each meter of water-head above the black line is equivalent to 165 - 170 million m^3 (MCM).[15] A water state corresponds to the water stock above the black line, so $s = 0$ when the water-head level is at -215 m, $s = 300$ MCM when the water head is at the red line (-213 m) and $s = \bar{s} = 1000$ when the water-head level is at -208.8 m. The admissible state set is $\mathcal{S} = \{s_1, s_2, ..., s_{n_s}\}$, where the s_j's are evenly spread apart. Setting $s_{j+1} - s_j \equiv \Delta_s = 50$ MCM gives $n_s = 21$ states.

An action a corresponds to pumping a million m^3 per year (MCMY). The admissible action set is $\mathcal{A} = \{a_1, a_2, ..., a_{n_a}\}$ with $a_1 = 0$, $a_{n_a} = 700$ MCMY (determined by the existing pumping infrastructure) and $a_{j+1} - a_j = \Delta_a,\ j = 1, 2, ..., n_a - 1$. Setting $\Delta_a = 50$ MCMY implies $n_a = 15$.

A time period (a year) in the present case begins at the end of the rainy season (the bulk of the rain in Israel's Mediterranean weather occurs during the months of November through April) while water extraction occurs mostly during the dry season (May - October). It is therefore not feasible to extract more than the water stock available at the beginning of the period, i.e., given the water stock S_t at the beginning of period t, $g_t \leq S_t$. Thus, $\mathcal{A}(S_t) = \{a_k \in$

[13]The exact minimal water head from which pumping is feasible is -214.87 m and we round it to -215 m.

[14]The red line has been modified in the past in response to pressure to increase pumping during dry years (see Gvirtzman 2002, p. 36).

[15]The range is due to differences in the surface of the lake at different water levels.

$\mathcal{A}|a_k \leq S_t\}$. At the end of the dry season, the water stock will reach the level $S_t - g_t \geq 0$ and this level affects the catastrophic hazard, as explained in subsection 6.4.

6.3 Transition probabilities

The transition probabilities, conditional on nonoccurrence, are

$$\begin{aligned} p(j|i, a_k) &= Pr\{S_{t+1} = s_j | S_t = s_i, g_t = a_k\} \\ &= Pr\{R(S_t) + X_t = s_j - s_i + a_k\} \\ &= p_{x|s}(s_j - s_i + a_k\}, \; j, i = 1, 2, ..., n_s, \; k = 1, 2, ..., n_a, \end{aligned} \quad (6.3)$$

where $p_{x|s}(\cdot)$ is defined in (6.2).

6.4 Period-t benefit

The immediate reward at time t, specified in (3.8), is repeated here for convenience:

$$b(S_t, g_t) = \tilde{b}(S_t, g_t) + v^p(S_t)[1 - \lambda(S_t, g_t)].$$

The first term on the right-hand side is the benefit enjoyed during non-occurrence periods; the second term is the benefit under the interrupting regime-shift, namely the post-event value weighted by the occurrence probability. The former consists of the surplus water users (irrigators, households, industry) derive from the pumped water g_t net of the supply cost (extraction, conveyance, treatment, distribution); the latter stems from the forgone benefit associated with not being able to use the lake for a prolong period of time. We discuss each in turn.

6.4.1 Immediate benefits during non-occurrence periods

Let $D(\cdot)$ denote the inverse demand facing the Kinneret's water, i.e., at a water price $\$D(a)$ per million m^3 (MCM) the water demand is a million m^3 per year (MCMY). Let $C(a)$ represent the cost of supplying a MCMY. The consumer surplus, net of the supply cost, associated with the consumption of a MCMY is

$$\int_0^a D(\xi)d\xi - C(a).$$

Assuming that the derived demand for water is inversely related to the water price, i.e., $D(a) = c_1/(a+1)$, and that $C(a) = c_2 a$, the net consumer surplus becomes

$$\tilde{b}(s,a) = c_1 \ln(a+1) - c_2 a, \tag{6.4}$$

where c_1 is a positive demand parameter and c_2 is the unit cost of water supply. Assuming further that at a price of $\$0.5 \times 10^6$ per MCM ($\$0.5$ per m^3) the water demand is 600 MCMY implies $c_1 = 300 \times 10^6$. The unit cost of supply is taken at $\$0.2 \times 10^6$ per MCM ($c_2 = 0.2 \times 10^6$).

6.4.2 Post-event value and occurrence probability

We consider the case in which the event (the abrupt regime shift) renders the lake's water unusable for a very long period and take the post-event value v^p to represent the forgone consumer surplus (i.e., the benefit water users could derive had the regime shift been prevented) as well as ecological damages and loss of recreational opportunities. We estimate this forgone value by the present value of constant flow $\tilde{b}(s,a)$ evaluated at $a = 550$ MCMY (which is about the average recharge). Thus, with the discount factor $\beta = 0.9434$

(corresponding to 6% interest rate) and the above specification of $\tilde{b}$,

$$v^p = -\tilde{b}(s, 550)/(1-\beta) \approx -3 \times 10^{10}.$$

The survival probability $\lambda(S_t, g_t)$ equals one if $S_t - g_t$ (the minimal water stock during time period t) does not fall below the critical water stock $s_c = 300$ MCM corresponding to the red line. As soon as the water-head drops below the red line, the survival probability decreases and reaches $\lambda(0) = \lambda_0 \geq 0$ at $s = 0$ (the black line). We use the following specification of the survival probability:

$$\lambda(s, a) = \begin{cases} \lambda_0 + (1-\lambda_0)\exp\{\delta(s-a-s_c)/(s-a)\} & \text{if } s-a < s_c \\ 1 & \text{if } s-a \geq s_c \end{cases} \tag{6.5}$$

where δ is a (positive) shape parameter. Indeed for $a = s$, exploitation brings the water stock to the black line and $\lambda(s, s) = \lambda_0$.

The immediate benefit specializes to

$$b(s, a) = c_1 \ln(a+1) - c_2 a + v^p(1-\lambda_0)\max\{1-\exp[\delta(s-a-s_c)/(s-a)], 0\}. \tag{6.6}$$

The function specifications and parameter values are summarized in Table 1.

Table 1

6.5 Optimal policy and value

We calculate the optimal policy using an algorithm based on Linear Programming (LP). Appendix C describes the algorithm and its application in the present case. The algorithm provides the optimal policy $\varphi^*(s_i),\ i = 1, 2, ..., n_s$, depicted in Figure 3.

Figure 3

Noting (A.9) and $\tilde{v} = v^*$, the value $v^* = (v^*(s_1), ..., v^*(s_{n_s}))'$ is calculated by

$$v^* = (I - \beta Q_{\varphi^*})^{-1} b_{\varphi^*}, \tag{6.7}$$

where $b_{\varphi^*} = (b(s_1, \varphi^*(s_1)), ..., b(s_{n_s}, \varphi^*(s_{n_s}))'$ and Q_{φ^*} is the $n_s \times n_s$ matrix with $\lambda(s_i, \varphi^*(s_i))p(j|i, \varphi^*(s_i))$ as the (i, j) element. The value is depicted in Figure 4.

Figure 4

6.6 Steady state

From the optimal extraction policy in Figure 3 we conclude that there is one recurrent, irreducible subset $E_1 = \{450, 500, ..., 1000\}$, and all states below 450 MCM are transient. This is so because the optimal extraction policy is such that it is not optimal to intentionally drop the water stock below 300 MCM (the red line) at the end of the dry season, and the minimal recharge (during the rainy season) is 150 MCMY. Thus, at the end of the year the water stock will be at or above 450 MCM. Water stocks (at the end of the rainy season) below 450 can only be encountered initially and for a limited number of periods (until recharge increases the stock), hence are transient.[16]

The λ_j^* data of Figure 3 reveal that E_1 contains only "safe" states ($\lambda_i^* = 1$ for all $s_i \in E_1$). Thus, once the optimal state process enters E_1 the event will never occur (the environmental threat is removed).

The steady state probabilities, characterized in Proposition 5.1 and applied with the above E_1, are depicted in Figure 5. In the long run (steady state), under the optimal policy, the stock never drops below 450 MCM (the red line,

[16]This state classification can be reached also by applying the procedure described in (Puterman 2005, p. 590) on the transition matrix P_{φ^*}.

below which the environmental threat is activated, is at 300 MCM). This allows pumping at least 150 MCMY without drawing the water head below the red line (recall that the water head at the end of the dry season reaches $S_t - g_t$), thereby providing a buffer against bad draws (dry years).

Figure 5

The average long-run stock and extraction are, respectively,

$$\hat{s} = \sum_{j=1}^{n_s} q_j^* s_j = 834.003 \text{ MCM}$$

and

$$\hat{g} = \sum_{j=1}^{n_s} q_j^* \varphi^*(s_j) = 494.211 \text{ MCMY}.$$

If the recharge were stable at the mean $\bar{x} = 570.38$ MCMY (see Figure 1), the steady-state extraction were set at this rate and this policy could have been maintained at a much lower stock level, e.g., at 300 MCM corresponding to the threshold stock (the red line water-head level). The higher (average) stock constitutes a buffer that allows mitigating extraction fluctuations, in spite of the stochastically fluctuating recharge, by drawing down the stock during bad (low recharge) years and filling it up during good (high recharge) years. On average, extractions are slightly less than the average recharge (494 MCMY vs. 570 MCMY), while under the steady state distribution the optimal extractions' standard deviation,

$$\sqrt{\sum_{j=1}^{n_s} \hat{q}_j [\varphi^*(s_j) - \hat{g}]^2} = 117.225,$$

is substantially smaller than the recharge process' standard deviation of 278.09 (see Figure 1). The latter owes to the buffer role of the water stock (this effect is akin to the buffer value proposed in Tsur and Graham-Tomasi (1991)).

The large long-run probability of the full capacity stock (the steady-state probability of $s = 1000$ MCM is about 1/3, implying that, under the optimal policy, in the long run the lake should be filled up every third winter on average) is an outcome of the policy of maintaining a large average stock (as a buffer against a series of dry years). Thus, it pays to let more water flow into the lower Jordan river (by opening the gates of Degania dam at the lake's southern outlet during rainy years) in order to have the buffer stock available during dry years. We note that this property is linked to the particular specifications and parameter values of Table 1, set for illustration purpose only.

7 Concluding comments

Exploitation has diminished the capacity of many renewable resources to endure stress, increasing their vulnerability to extreme environmental conditions that may trigger abrupt changes. The onset of such events depends on the coincidence of extreme environmental conditions and the resource state. When both of these elements are uncertain, the uncertainty associated with the event occurrence is the result of their combined effect. We analyzed resource management in such a setting.

The environmental threat affects management policies in two ways: first, it changes the immediate benefit flow; second, it turns the running discount factor endogenous to the exploitation policy and the compound discount factor becomes history-dependent. The consequences regarding the existence of an optimal Markovian-Deterministic stationary policy can be detrimental, as demonstrated by an example. Nonetheless, we establish the existence of such a policy and show that the optimal state process converges in the long run to a well specified steady-state distribution. A numerical example illustrates these properties.

The environmental threat is manifest in our framework via the abrupt change – the regime shift or event occurrence (ecosystem collapse, biomass extinction) – and a key feature in the analysis is the distinction between the pre- and post-event regimes. Different resources have different pre-event regimes; different environmental threats entail different post-event regimes. The framework developed here provides a basis for studying a host of renewable resource situations under a wide variety of environmental threats.

Appendix

A Existence of optimal stationary policy

We prove an extended version of Theorem 4.1, which makes use of the following definitions and notation. Recall that without the catastrophic threat, i.e., when the survival probability $\lambda(s,a) = 1$ for all $s \in \mathcal{S}$ and $a \in \mathcal{A}$, the discount factor is constant and the optimality equations are

$$v(s_i) = \max_{a_i \in \mathcal{A}(s_i)} \left\{ b(s_i, a_i) + \beta \sum_{j=1}^{n_s} p(j|i, a_i) v(s_j) \right\}, \; i = 1, 2, ..., n_s,$$

or in matrix notation

$$v = \max_{a \in \mathcal{A}} \left\{ b_a + \beta P_a v \right\},$$

where $v = (v(s_1), ..., v(s_{n_s}))'$, $a = (a_1, ..., a_{n_s}) \in \mathcal{A}(s_1) \times \cdots \times \mathcal{A}(s_{n_s}) = \mathcal{A}(s)$, $b_a = (b(s_1, a_1), ..., b(s_{n_s}, a_{n_s}))'$ and P_a is the $n_s \times n_s$ matrix with the (i,j) element given by $p(j|i,a)$. In the presence of environmental threat, the discount factor $\beta\lambda(s_i, a)$ is state-and-action-dependent and the optimality equations become

$$v(s_i) = \max_{a_i \in \mathcal{A}(s_i)} \left\{ b(s_i, a_i) + \beta\lambda(s_i, a_i) \sum_{j=1}^{n_s} p(j|i, a_i) v(s_j) \right\}, \; i = 1, 2, ..., n_s, \tag{A.1}$$

or in matrix notation

$$v = \max_{a \in \mathcal{A}} \{b_a + \beta Q_a v\}, \tag{A.2}$$

where Q_a is an $n_s \times n_s$ matrix with (i,j) element given by $\lambda(s_i, a) p(j|i,a)$ (the i'th row of Q_a equals $\lambda(s_i, a)$ times the i'th row of P_a).

Let V be the space of bounded functions on $\mathcal{S}$ endowed with the supremum norm $\|v\| = \sup_{s \in \mathcal{S}} v(s)$. Define the mapping $L : V \mapsto V$:

$$L(v)_i = \max_{a_i \in \mathcal{A}(s_i)} \left\{ b(s_i, a_i) + \beta\lambda(s_i, a_i) \sum_{j=1}^{n_s} p(j|i, a_i) v(s_j) \right\}, \; i = 1, 2, ..., n_s,$$

or in matrix notation

$$L(v) = \max_{a \in \mathcal{A}} \left\{ b_a + \beta Q_a v \right\}. \tag{A.3}$$

The optimality equations can be expressed in terms of L as

$$v(s_i) = L(v)_i, \; i = 1, 2, ..., n_s,$$

or in matrix notation as

$$v = L(v). \tag{A.4}$$

We can now establish the following extended Theorem 4.1:

Theorem A.1. *Suppose that (A1) $0 \leq \beta < 1$, (A2) $\mathcal{S}$ is discrete (finite or countable), (A3) $\tilde{b} : \mathcal{S} \times \mathcal{A} \mapsto \mathrm{I\!R}$ and $v^p : \mathcal{S} \mapsto \mathrm{I\!R}$ are bounded and (A4) $\tilde{b}(s_i, a)$ and $\lambda(s_i, a)p(j|i, a)$ are continuous in a, and $\mathcal{A}(s_i)$ is compact for all $s_i, s_j \in \mathcal{S}$. Then:*
(i) the optimal value v^* is the unique fixed point of (A.4);
(ii) a stationary policy φ is optimal if and only if the actions $a_i = \varphi(s_i)$, $i = 1, 2, ..., n_s$, realize the maxima in (A.1);
(iii) there exists an optimal, Markovian-Deterministic stationary policy φ^*, i.e., the policy $(\varphi^*, \varphi^*, ...)$ satisfies

$$v^{\varphi^*}(s) = v^*(s) \; \forall s \in \mathcal{S}. \tag{A.5}$$

Proof. Assumptions (A3)-(A4) ensure that the maxima in (A.1) are attained. For a given $v \in V$, let $a_i(v)$, $i = 1, 2, ..., n_s$, denote the actions where the maxima in (A.1) are attained. Then, for any $u \in V$ we have

$$L(u)_i \geq \left\{ b(s_i, a_i(v)) + \beta\lambda(s_i, a_i(v)) \sum_{j=1}^{n_s} p(j|i, a_i(v))u_j \right\}, \; i = 1, 2, ..., n_s,$$

which together with

$$L(v)_i = b(s_i, a_i(v)) + \beta\lambda(s_i, a_i(v)) \sum_{j=1}^{n_s} p(j|i, a_i(v))v_j$$

implies

$$L(v)_i - L(u)_i \leq \beta\lambda(s_i, a_i(v)) \sum_{j=1}^{n_s} p(j|i, a_i(v))(v_j - u_j), \; i = 1, 2, ..., n_s. \quad \text{(A.6)}$$

Since $\sum_{j=1}^{n_s} p(j|i, a_i(v)) = 1$, we conclude from (A.6) that

$$L(v)_i - L(u)_i \leq \beta\lambda(s_i, a_i(v)) \max_j |v_j - u_j|, \; i = 1, 2, ..., n_s.$$

Since $\beta\lambda(s_i, a_i(v)) \leq \beta < 1$, we can further conclude that

$$\max_i \{L(v)_i - L(u)_i\} \leq \beta \max_j |v_j - u_j|.$$

Interchanging in the above inequality the roles of u and v we obtain

$$\max_i |L(v)_i - L(u)_i| \leq \beta \max_j |v_j - u_j|. \quad \text{(A.7)}$$

It follows from (A.7) and (A1) that L is a contraction, implying the existence of a unique fixed point of (A.4). Denote this fixed point by $\tilde{v}$. We next show that $\tilde{v} = v^*$.

Let a_i^*, $i = 1, 2, ..., n_s$, be the actions that realize the maxima in (A.1), and define $\varphi^*(s_i) = a_i^*$. Then,

$$\tilde{v}(s_i) = b(s_i, \varphi^*(s_i)) + \beta\lambda(s_i, \varphi^*(s_i)) \sum_{j=1}^{n_s} p(j|i, \varphi^*(s_i))\tilde{v}(s_j), \; s_i \in \mathcal{S}, \quad \text{(A.8)}$$

or in vector notation

$$\tilde{v} = b_{\varphi^*} + \beta Q_{\varphi^*}\tilde{v}, \quad \text{(A.9)}$$

where $b_{\varphi^*} = (b(s_1, \varphi^*(s_1)), ..., b(s_{n_s}, \varphi^*(s_{n_s})))'$ and Q_{φ^*} is the $n_s \times n_s$ matrix with the (i, j) element given by $\lambda(s_i, \varphi^*(s_i))p(j|i, \varphi^*(s_i))$.

Evaluating (A.8) at time t, with $s_i = S_t$ and $g_t = \varphi^*(S_t)$, gives

$$\begin{aligned}
\tilde{v}(S_t) &= b(S_t, \varphi^\star(S_t)) + \beta\lambda(S_t, \varphi^*(S_t)) \sum_{j=1}^{n_s} p(j|S_t, \varphi^\star(S_t))\tilde{v}(s_j) \\
&= b(S_t, \varphi^\star(S_t)) + \beta\lambda(S_t, \varphi^*(S_t)) E_t^{\varphi^*} \tilde{v}(S_{t+1}), \qquad \text{(A.10)}
\end{aligned}$$

where $E_t^{\varphi^*}$ denotes expectation under the $g_t = \varphi^*(S_t)$ decision rule conditional on the information available at time t (which includes S_t). Multiplying (A.10) by $\gamma^{\varphi^*}(t)$, where $\gamma(t)$ is defined in (3.5) under the $g_t = \varphi^*(S_t)$ decision rule, and rearranging gives

$$b(S_t, \varphi^\star(S_t))\gamma^{\varphi^*}(t) = \tilde{v}(S_t)\gamma^{\varphi^*}(t) - \gamma^{\varphi^*}(t+1) E_t^{\varphi^*} \tilde{v}(S_{t+1}). \qquad \text{(A.11)}$$

Since $\gamma^{\varphi^*}(t+1)$ depends only on information available at time t, the second term on the right hand side of (A.11) can be written as

$$\gamma^{\varphi^*}(t+1) E_t^{\varphi^*} \tilde{v}(S_{t+1}) = E_t^{\varphi^*} \left[\gamma^{\varphi^*}(t+1)\tilde{v}(S_{t+1})\right]$$

and (A.11) is written as

$$b(S_t, \varphi^*(S_t))\gamma^{\varphi^*}(t) = \gamma^{\varphi^*}(t)\tilde{v}(S_t) - E_t^{\varphi^*} \left[\gamma^{\varphi^*}(t+1)\tilde{v}(S_{t+1})\right].$$

Taking the unconditional expectation under the $\varphi^*(\cdot)$ decision rule yields

$$E^{\varphi^*} b(S_t, \varphi^\star(S_t))\gamma^{\varphi^*}(t) = E^{\varphi^*}\gamma^{\varphi^*}(t)\tilde{v}(S_t) - E^{\varphi^*}\gamma^{\varphi^*}(t+1)\tilde{v}(S_{t+1}).$$

Summing over $t = 1, 2, ..., \tau$ gives

$$E^{\varphi^*} \sum_{t=1}^{\tau} b(S_t, \varphi^\star(S_t))\gamma^{\varphi^*}(t) = \tilde{v}(S_1) - E^{\varphi^*}\gamma^{\varphi^*}(\tau+1)\tilde{v}(S_{\tau+1}). \qquad \text{(A.12)}$$

Since $\gamma^{\varphi^*}(\tau) \to 0$ exponentially (uniformly in the policies), letting $\tau \to \infty$ in (A.12) yields

$$E^{\varphi^*} \sum_{t=1}^{\infty} b(S_t, \varphi^*(S_t))\gamma^{\varphi^*}(t) = \tilde{v}(S_1), \qquad \text{(A.13)}$$

where we use the property that $s_i \mapsto \tilde{v}(s_i)$, $s_i \in \mathcal{S}$, is a bounded function, namely $\tilde{v}$ is a bounded solution of (A.2), which is guaranteed by (A3).

For an arbitrary policy $\varphi(\cdot)$ we can repeat the above derivation with inequalities rather than equalities, obtaining

$$\tilde{v}(S_t) \geq b(S_t, \varphi(S_t)) + \beta\lambda(S_t, \varphi(S_t)) \sum_{j=1}^{n_s} p(j|S_t, \varphi(S_t))\tilde{v}(s_j)$$

instead of (A.10) and

$$E^{\varphi} \sum_{t=1}^{\infty} b(S_t, \varphi(S_t))\gamma^{\varphi}(t) \leq \tilde{v}(S_1)$$

instead of (A.13). It follows that $\varphi^{\star}(s)$ is an optimal policy and $\tilde{v}(s) = v^*(s)$, establishing claims (i) and (ii) of the theorem. As indicated above, the only condition for the existence of $\varphi^*(\cdot)$ is that there exists a bounded solution for (A.2), which follows from condition (A3) and claim (i), establishing (iii). $\square$

B Transient states probabilities

Suppose the state process departs from a transient state $s_j \in E_0$ and consider the $n_0 \times (K+1)$ matrix $\tilde{Q}_0$ defined in equation (5.5). We verify that $\tilde{Q}_0(j,k)$, $k = 1, 2, ..., K$, are the probabilities that the optimal state process will exit the transient set E_0 into the recurrent set E_k, $k = 1, 2, ..., K$, respectively, and $\tilde{Q}_0(j, K+1)$ is the probability that it will exit E_0 into the event occurrence set $E_{K+1} \equiv \{\kappa\}$. Recall that $\tilde{Q}_0^1$ is the $n_0 \times (K+1)$ matrix whose (j,k) elements give the *one-period* probabilities of moving from $s_j \in E_0$ to E_k, $k = 1, 2, ..., K+1$. Then, $\tilde{Q}_0(j,k)$ satisfies

$$\tilde{Q}_0(j,k) = \tilde{Q}_0^1(j,k) + \sum_{\{l|s_l \in E_0\}} Q_{\varphi^*}(j,l)\tilde{Q}_0(l,k), \;\; s_j \in E_0, \;\; k = 1, 2, ..., K+1.$$

In matrix notation, the above is expressed as

$$Q_{0K} = \tilde{Q}^1_{0K} + Q_0 Q_{0K},$$

where Q_0 is the $n_0 \times n_0$ submatrix of Q_{φ^*} corresponding to the n_0 transient states. Thus, equation (5.5) follows if the inverse matrix $(I - Q_0)^{-1}$ exists. To show this, note that, since the optimal state process cannot reside in the transient set E_0 forever, it must be that $Q_0^n \to 0$ as $n \to \infty$. This implies that the eigenvalues of Q_0 are all smaller than one in absolute value (to verify this use the Jordan canonical form of Q_0), hence $Q_0^n \to 0$ exponentially and $(I - Q_0)^{-1} = \sum_{n=0}^{\infty} Q_0^n$ exists.

C The LP algorithm for calculating optimal policies of MDPs

Puterman (Puterman 2005, Chapter 6) presents a variety of algorithms for calculating optimal policies of Markov decision processes (MDPs). We use the algorithm based on Linear Programming (LP), adopted to the present case of a state-dependent discount factor. We briefly describe the algorithm and its application.

C.1 The LP approach for solving MDPs

The algorithm is based on the following property:

Proposition C.1. *If $v \in V$ satisfies $v \geq L(v)$, then $v \geq v^*$.*

Proof. The mapping L, defined in (A.3), is monotonic, i.e., for $v, u \in V$, $v \geq u$ implies $L(v) \geq L(u)$. This property follows from $\beta \geq 0$ and $Q_a(i,j) \geq 0\ \forall(i,j)$. Thus, $v \geq L(v)$ implies $L(v) \geq L(L(v)) \equiv L^2(v)$, hence $v \geq L(v)$

implies $v \geq L^2(v)$. Repeating this reasoning, we find that $v \geq L(V)$ implies $v \geq L^k(v)$ for $k = 1, 2, ...$ Letting $k \to \infty$, recalling that L is a contraction and v^* is the unique fixed point of $v = L(v)$ (Theorem 4.1), establishes the result. $\square$

It follows that the inequality $v \geq L(v)$, or in component notation

$$v_i \geq b(s_i, a_k) + \beta\lambda(s_i, a_k) \sum_j p(j|i, a_k) v_j \ \forall a_k \in \mathcal{A}(s_i),\ i = 1, 2, ..., n_s,$$

can at best hold as equality, in which case $v = v^*$. This suggests the following (primal) Linear Programming (LP) problem for finding v^*:

Set $\alpha_j > 0,\ j = 1, 2, ..., n_s$, satisfying $\sum_j \alpha_j = 1$ (any positive α_j will do but the requirement that they sum to one allows a probability interpretation) and find (unconstrained) $v_j,\ j = 1, 2, ..., n_s$, in order to minimize

$$\sum_{j=1}^{n_s} \alpha_j v_j$$

subject to

$$v_i - \beta\lambda(s_i, a_k) \sum_{j=1}^{n_s} p(j|i, a_k) v_j \geq b(s_i, a_k) \ \forall a_k \in \mathcal{A}(s_i),\ i = 1, 2, ..., n_s.$$

This LP problem has n_s unknowns (columns) and $\sum_{i=1}^{n_s} n_{a_i}$ constraints (rows), where n_{a_i} is the number of actions in $\mathcal{A}(s_i)$.

The dual to the above LP problem is formulated as follows:

Find $x(s_i, a_k) \geq 0,\ i = 1, 2, ..., n_s,\ a_k \in \mathcal{A}(s_i)$, in order to maximize

$$\sum_{i=1}^{n_s} \sum_{a_k \in \mathcal{A}(s_i)} b(s_i, a_k) x(s_i, a_k) \tag{C.1}$$

subject to

$$\sum_{a_k \in \mathcal{A}(s_j)} x(s_j, a_k) - \sum_{i=1}^{n_s} \sum_{a_k \in \mathcal{A}(s_i)} \beta\lambda(s_i, a_k) p(j|i, a_k) x(s_i, a_k) = \alpha_j,\ j = 1, 2, ..., n_s. \tag{C.2}$$

The dual LP has $\sum_{i=1}^{n_s} n_{a_i}$ unknowns (columns) and n_s constraints (rows). The number of constraints is smaller than that of the primal LP problem, which renders the dual LP more tractable. Properties of the dual LP problem, including a verification that a basic solution exists, are discussed in Puterman (2005, pp. 223-231).

Let $x^*(s_i, a_k),\ i = 1, 2, ..., n_s,\ k = 1, 2, ..., n_{a_i}$, denote the solution of the dual LP. Since the dual LP has n_s constraints, only n_s out of the $\sum_{i=1}^{n_s} n_{a_i}$ elements of x^* are positive. Moreover, for any state s_i only one $x^*(s_i, a_k) > 0$. The optimal (Markov-deterministic) stationary policy is specified as

$$\varphi^*(s_i) = \sum_{a_k \in \mathcal{A}(s_i)} \mathbf{1}(x^*(s_i, a_k) > 0)a_k,\ i = 1, 2, ..., n_s, \tag{C.3}$$

where $\mathbf{1}(\cdot)$ assumes the values 1 or 0 when its argument is true or false, respectively.

C.2 LP specification in the present case

Let $D(i, k) = 1$ or 0 as $s_i \geq a_k$ or $s_i < a_k$, respectively. Thus, $D(i, k) = 1$ if the action a_k is feasible at s_i and $D(i, k) = 0$ otherwise (see discussion in subsection 6.2). Let B be the $n_s \times n_a$ matrix with the i, k element given by $b(s_i, a_k)D(i, k)$, where $b(s, a)$ is defined in (6.6). The LP objective (C.1) can be rendered as

$$\sum_{i=1}^{n_s} \sum_{k=1}^{n_a} B(i, k)x(i, k). \tag{C.4}$$

Similarly, let $\tilde{p}(j|i, a_k) = \lambda(s_i, a_k)p(j|i, a_k)D(i, k)$, where $p(j|i, a_k)$ is defined in (6.3). Then

$$\sum_{i=1}^{n_s} \sum_{a_k \in \mathcal{A}(s_i)} p(j|i, a_k)x(i, k) = \sum_{i=1}^{n_s} \sum_{k=1}^{n_a} \tilde{p}(j|i, a_k)x(i, k)$$

and the dual LP constraints (C.2) can be expressed as

$$\sum_{i=1}^{n_s}\sum_{k=1}^{n_a} D(i,k)x(i,k) - \beta\sum_{i=1}^{n_s}\sum_{k=1}^{n_a}\tilde{p}(j|i,k)x(i,k) = 1/n_s,\ j=1,2,...,n_s, \tag{C.5}$$

where we set $\alpha_j = 1/n_s,\ j=1,2,...,n_s$.

The LP problem then is to find $x(i,k)\geq 0,\ i=1,2,...,n_s,\ k=1,2,...,n_a$, in order to maximize (C.4) subject to (C.5).

D A non-existence example

Theorem 4.1 extends a result that holds for standard state-dependent models to a history-dependent situation. The dependence on the whole history has a specific form, which enables this extension. We describe here an example in which the history-dependence of the process is such that there does not exist an optimal deterministic stationary policy.

Consider an MDP with two states, s_1 and s_2, and three actions, $a_1,\ a_2$ and a_3, such that the following holds:

$$\frac{1}{2} < p(s_1|s_2,a_1),\ p(s_2|s_1,a_1) < 1, \tag{D.1}$$

$$0 < p(s_1|s_2,a_2),\ p(s_2|s_1,a_2) < \frac{1}{2} \tag{D.2}$$

and

$$p(s_1|s_1,a_3) = p(s_2|s_2,a_3) = 1. \tag{D.3}$$

The state process is $\{S_t\}_{t=0}^{\infty}$ and the action at time t is g_t. The running (single period) rewards $c(s_1,a_1)$, $c(s_1,a_2)$, $c(s_2,a_1)$, and $c(s_2,a_2)$ are *negative and of order 1*, and

$$c(s_1,a_3) = c(s_2,a_3) = -M,\ M >> 1. \tag{D.4}$$

There are 9 possible deterministic stationary policies, and we assume that if $(\hat{p}_1, \hat{p}_2)$ is a stationary equilibrium distribution then (the above parameters are so chosen that)

$$(\hat{p}_1, \hat{p}_2) \neq (0.5, 0.5). \tag{D.5}$$

The discount factor $\gamma(H_t)$ at time t depends on the history H_t at time t:

$$H_t = (S_0, g_0, S_1, g_1, ..., S_{t-1}, g_{t-1}, S_t).$$

To define $\gamma(H_t)$ we use the empirical distribution of the state, namely

$$\nu_t(s_1) = \frac{\#\{0 \le j \le t : S_j = s_1\}}{t+1}, \ \nu_t(s_2) = 1 - \nu_t(s_1) \tag{D.6}$$

and define

$$\gamma(H_t) = \begin{cases} 0 & \text{if } \nu_t(S_t) < 0.5 \\ 1 & \text{if } \nu_t(S_t) \ge 0.5. \end{cases} \tag{D.7}$$

Thus, e.g., if at time t we have $S_t = s_1$ and $\nu_t(s_1) < 1/2$ then the reward at time t is zero, and if $\nu_t(s_1) \ge 1/2$ the reward is $\beta^t c(s_1, g_t)$ (which is a negative number).

We seek to maximize

$$C^\pi(S_0) = \sum_{t=0}^{\infty} \gamma(H_t) c(S_t, \pi_t(H_t))$$

and we claim that there exists no optimal deterministic stationary policy which maximizes C^π. Suppose to the contrary, that $\pi = \{\varphi, \varphi, ...\}$ is such a policy. It follows from (D.3) and (D.4) that the actions $\varphi(s_1)$ and $\varphi(s_2)$ belong to $\{a_1, a_2\}$. Let the equilibrium distribution under φ be $(\hat{p}_1, \hat{p}_2)$, where we recall that

$$0 < \hat{p}_1, \ \hat{p}_2 < 1 \text{ and } \hat{p}_1 \neq 0.5. \tag{D.8}$$

We will construct a policy π_0 with a higher payoff than π.

Let

$$\lambda = \min\{\hat{p}_1, \hat{p}_2\},$$

let $\epsilon > 0$ be such that $\epsilon << \lambda$ and let t_0 be a large integer such that

$$\lambda - \epsilon < \nu_{t_0}(s_1) < \lambda + \epsilon.$$

(How large t_0 should be will be determined below.) Let

$$T = [t_0(1 - 2(\lambda + \epsilon))] \tag{D.9}$$

where here $[x]$ denotes the integer part of x. It follows from (D.8) that $\lambda + \epsilon < 1/2$ and therefore $T \to \infty$ as $t_0 \to \infty$. It is then easy to see that

$$\nu_t(s_1) < 1/2 \text{ for every } t_0 \leq t \leq t_0 + T.$$

For the policy π_0 we take $g_t = a_3$ for $t_0 \leq t \leq t_0 + T$, and it follows that $S_t = s_1$ for every $t_0 \leq t \leq t_0 + T$. For $t > t_0 + T$, π_0 coincides with π. Comparing the difference $C^{\pi_0}(S_0) - C^{\pi}(S_0)$ between the payoffs of π_0 and π starting at S_0, we consider the corresponding rewards on the time interval $t_0 \leq t \leq t_0 + T$. The difference on this time interval is larger than

$$\beta^{t_0}(\mu - (\beta)^T |d_2 - d_1|)$$

where

$$\mu = \min\{|c(s_1, a_1)|, |c(s_1, a_2)|\}$$

is positive and of order 1, and where d_1 and d_2 are the payoffs $C^{\pi(}s_1)$ and $C^{\pi}(s_2)$ corresponding to s_1 and s_2 respectively. For large enough T, namely for large enough t_0 (recall (D.9)) this expression is positive, and hence the payoff under π_0 is strictly larger than that under π.

Table 1: Specifications and parameter values

Function	Form	Description
$\tilde{b}(s,a)$	$c_1 \ln(a+1) - c_2 a$	Reward under no occurrence
$v^p(s)$	Constant	Post-event value
$\lambda(s,a)$	$\min\left\{1, \lambda_0 + (1-\lambda_0)e^{\delta(s-a-s_c)/(s-a)}\right\}$	Survival probability
Parameter	**Value**	**Description**
β	0.9434	Discount factor =1/(1+0.06)
α	2.20967	Recharge dist. parameter
θ	187.077	Recharge dist. parameter
Δ_s	50 MCM	Diff between consecutive states
n_s	21	Number of admissible states
Δ_a	50 MCMY	Diff between consecutive actions
n_a	15	Number of admissible actions
Δ_x	50 MCMY	Diffe between consecutive recharge
n_x	26	Number of recharge points
c_1	300×10^6	Demand parameter
c_2	0.2×10^6	Unit supply cost
v^p	-3×10^{10}	Forgone benefit due to occurrence
s_c	300 MCM	Critical stock (at red line)
λ_0	0.5	Survival prob at $s = 0$ (black line)
δ	0.2	Hazard parameter

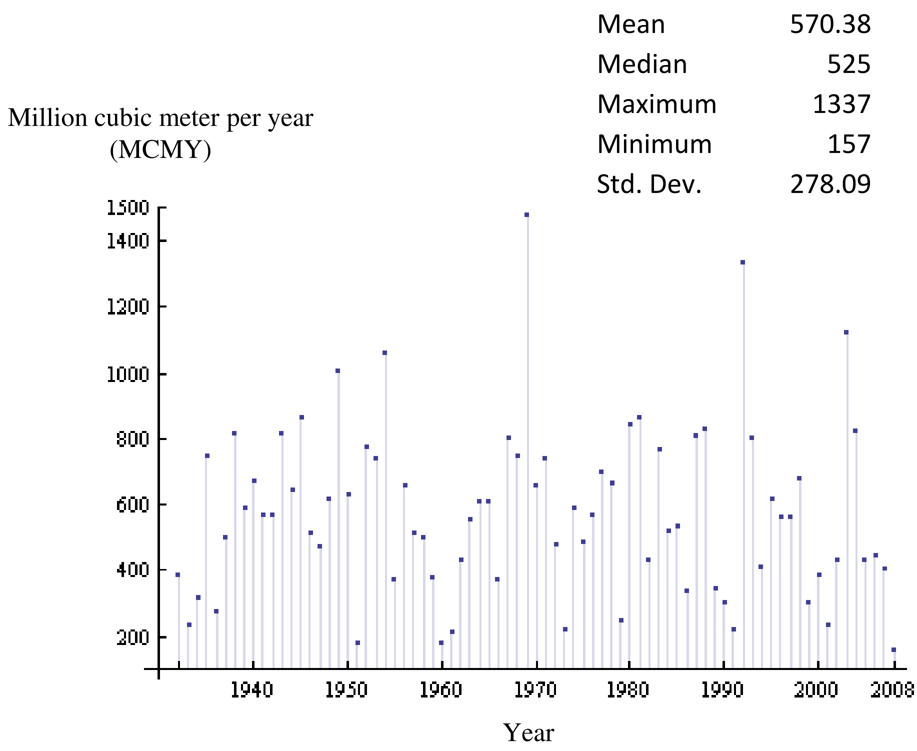

Figure 1: Lake Kinneret's recharge series during 1932 - 2008. The descriptive statistics are calculated for the 1980 - 2008 data.

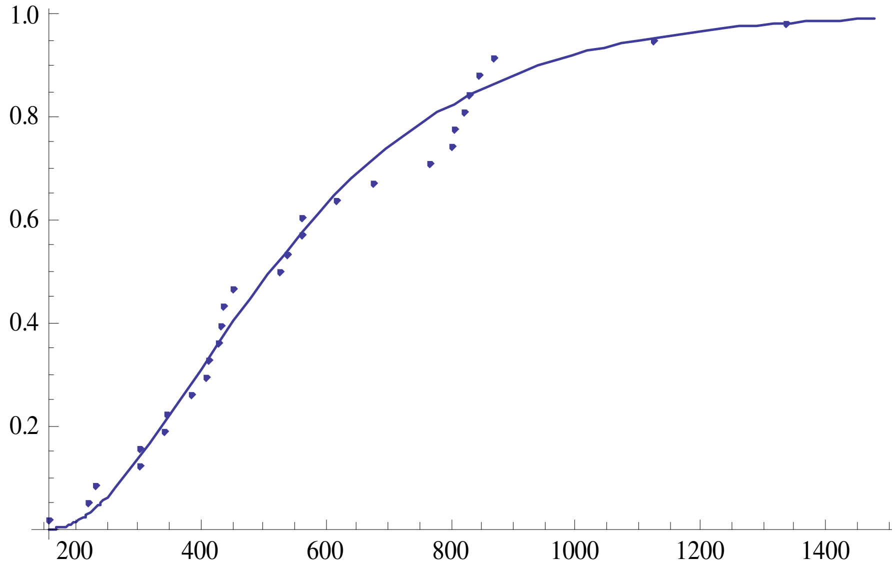

Figure 2: The gamma distribution with parameters $\alpha = 2.20967$ and $\theta = 187.077$ (solid) and the empirical distribution (dots) of the Kinneret's recharge series for the period 1980 - 2008.

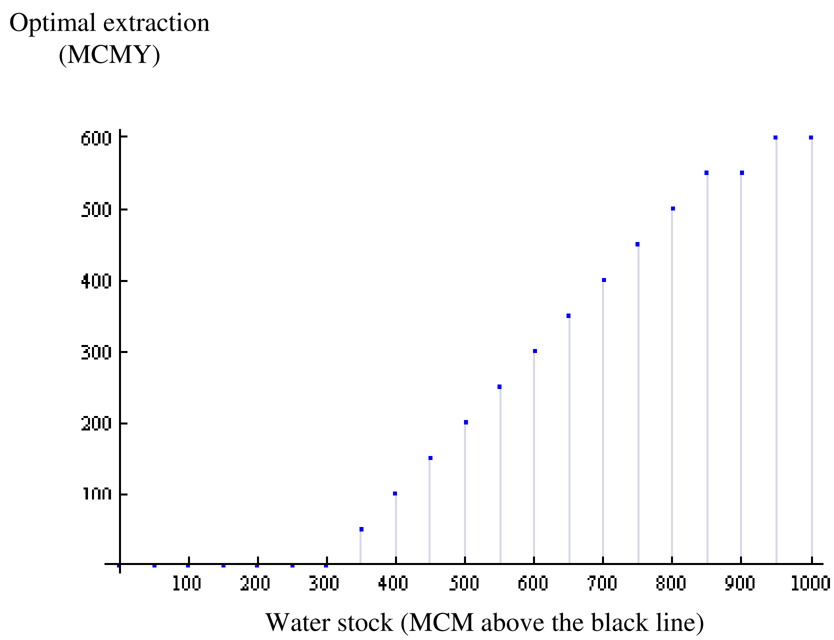

s_j	$\varphi^*(s_j)$	λ^*_j
0	0	0.5
50	0	0.68394
100	0	0.83516
150	0	0.909365
200	0	0.952419
250	0	0.980395
300	0	1
350	50	1
400	100	1
450	150	1
500	200	1
550	250	1
600	300	1
650	350	1
700	400	1
750	450	1
800	500	1
850	550	1
900	550	1
950	600	1
1000	600	1

Figure 3: The optimal stationary Markov extraction policy $\varphi^*(s)$ (MCMY) for $s = 0, 50, 100, ..., 1000$. The data are reported to the right of the figure and contain also the survival probabilities λ_j^*.

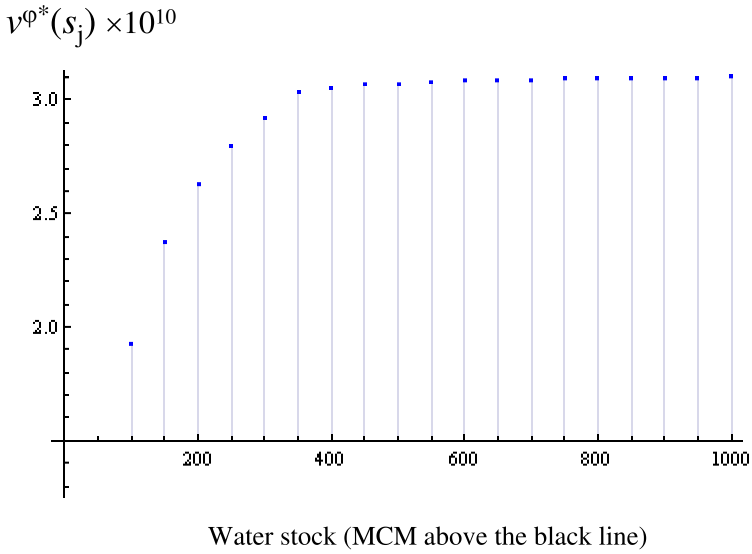

s_j	$v^{\varphi^*}(s_j)$
0	-0.0699987
50	1.02578
100	1.92885
150	2.37428
200	2.63255
250	2.8008
300	2.91903
350	3.03599
400	3.05549
450	3.06655
500	3.07413
550	3.0798
600	3.08425
650	3.08786
700	3.09085
750	3.09338
800	3.09553
850	3.09739
900	3.09919
950	3.1008
1000	3.10226

Figure 4: The value $v^{\varphi^*}(s)$ ($\times 10^{10}$ \$) for $s = 0, 50, 100, ..., 1000$ MCM.

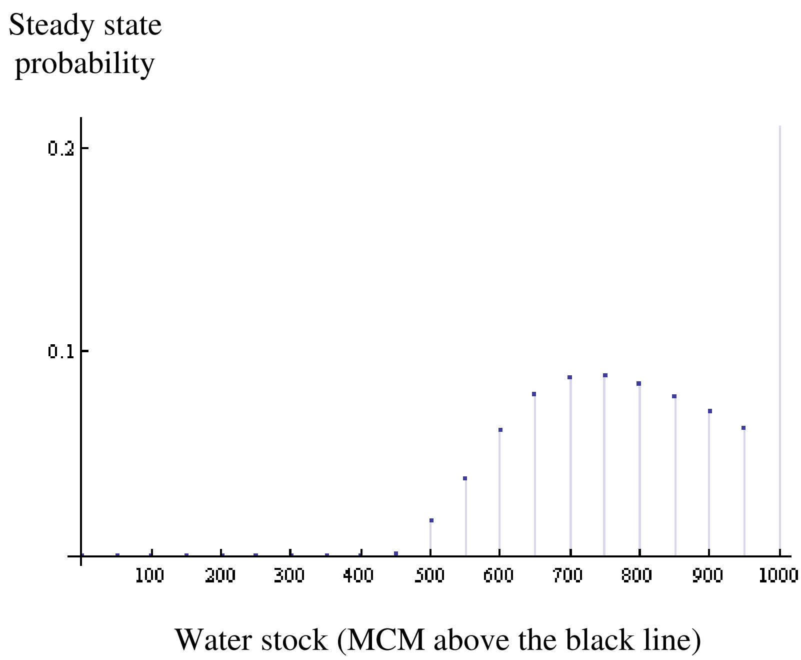

s_j	Steady state prob.
0	0
50	0
100	0
150	0
200	0
250	0
300	0
350	0
400	0
450	0.00116072
500	0.0173963
550	0.038434
600	0.0621503
650	0.0791231
700	0.0868794
750	0.0878662
800	0.0843703
850	0.0781758
900	0.0705778
950	0.0624706
1000	0.331396

Figure 5: Long run (steady state) probabilities.

References

Acemoglu, D.: 2009, *Introduction to Modern Economic Growth*, Princenton University Press, Princeton, NJ.

Alley, R. B., Marotzke, J., Nordhaus, W. D., Overpeck, J. T., Peteet, D. M., Pielke Jr., R. S., Pierrehumbert, R. T., Rhines, P. B., Stocker, T. F., Talley, L. D. and Wallace, J. M.: 2003, Abrupt climate change, *Science* **299**, 2005–2010.

Aronsson, T., Backlund, K. and Löfgren, K.-G.: 1998, Nuclear power, externalities and non-standard pigouvian taxes: a dynamic analysis under uncertainty, *Environmental & Resource Economics* **11**, 177–195.

Bahn, O., Haurie, A. and Malhamé, R.: 2008, A stochastic control model for optimal timing of climate policies, *Automatica* **44**, 1545–1558.

Brock, W. A. and Starrett, D.: 2003, Managing systems with non-convex positive feedback, *Environmental & Resource Economics* **26**, 575–602.

Brock, W. A. and Xepapadeas, A.: 2003, Valuing biodiversity from an economic perspective: a unified economic, ecological and genetic approach, *American Economic Review* **93**, 1597 – 1614.

Broecker, W. S.: 1997, Thermohaline circulation, the Achilles heel of our climate system: Will man-made CO_2 upset the current balance?, *Science* **278**, 1582–1588.

Burt, O. R.: 1964, Optimal resource use over time with an application to ground water, *Management Science* **11**(1), 80 – 93.

Carpenter, S. R., Ludwig, D. and Brock, W. A.: 1999, Management of eutrophication for lakes subject to potentially irreversible change, *Ecological Applications* **9**(3), 751 – 771.

Clark, C. W. and Kirkwood, G. P.: 1986, On uncertain renewable resource stocks: Optimal harvest policies and the value of stock surveys, *Journal of Environmental Economics and Management* **13**(3), 235–244.

Clarke, H. R. and Reed, W. J.: 1994, Consumption/pollution tradeoffs in an environment vulnerable to pollution-related catastrophic collapse, *Journal of Economic Dynamics and Control* **18**(5), 991–1010.

Costello, C., Polasky, S. and Solow, A.: 2001, Renewable resource management with environmental prediction, *Canadian Journal of Economics* **34**(1), 196 – 211.

Cropper, M. L.: 1976, Regulating activities with catastrophic environmental effects, *Journal of Environmental Economics & Management* **3**(1), 1–15.

Dasgupta, P. and Mäler, K.-G.: 2003, The economics of non-convex ecosystems: Introduction, *Environmental & Resource Economics* **26**, 499–525.

Epstein, L. G. and Schneider, M.: 2007, Learning under ambiguity, *Review of Economic Studies* **74**, 1275 – 1303.

Gvirtzman, H.: 2002, *Israel Water Resources*, Yad Ben-Zvi Press, Jerusalem (in Hebrew).

Harper, D.: 1992, *Eutrophication of freshwater*, Chapman and Hall, London.

Haurie, A. and Moresino, F.: 2006, A stochastic control model of economic growth with environmental disaster prevention, *Automatica* **42**, 1417–1428.

Judd, K. L.: 1998, *Numerical Methods in Economics*, Vol. 1, The MIT Press, Cambridge, MA.

Knapp, K. and Olson, L.: 1995, The economics of conjunctive groundwater management with stochastic surface supplies, *Journal of Environmental Economics and Management* **28**(3), 340 – 356.

Mäler, K.-G.: 2000, Development, ecological resources and their management: A study of complex dynamic systems, *European Economic Review* **44**, 645–665.

Mastrandrea, M. D. and Schneider, S. H.: 2001, Integrated assessment of abrupt climatic changes, *Climate Policy* **1**, 433–449.

McGough, B., Plantiga, A. J. and Costello, C.: 2009, Optimally managing a stochastic renewable resource under general economic conditions, *The B.E. Journal of Economic Analysis & Policy* **9**, 1–29.

Mitra, T. and Roy, S.: 2006, Optimal exploitation of renewable resources under uncertainty and the extinction of species, *Economic Theory* **28**, 1 – 23.

Nævdal, E.: 2006, Dynamic optimization in the presence of threshold effects when the location of the threshold is uncertain – with an application to a possible disintegration of the western antarctic ice sheet, *Journal of Economic Dynamics and Control* **30**, 1131–1158.

Pindyck, R. S.: 1984, Uncertainty in the theory of renewable resource markets, *Review of Economic Studies* **51**, 289 – 303.

Pindyck, R. S.: 2002, Optimal timing problems in environmental economics, *Journal of Economic Dynamics and. Control* **26**, 1677 – 1697.

Pindyck, R. S.: 2007, Uncertainty in environmental economics, *Review of Environmental Economics and Policy* **1**(1), 45 – 65.

Puterman, M. L.: 2005, *Markov Decision Processes: Discrete Stochastic Dynamic Programming*, Wiley.

Reed, W. J.: 1974, A stochastic model for the economic management of a renewable animal resource, *Mathematical Biosciences* **22**, 313 – 337.

Reed, W. J.: 1979, Optimal escapement levels in stochastic and deterministic harvesting models, *Journal of Environmental Economics and Management* **6**(4), 350–363.

Reed, W. J.: 1989, Optimal investment in the protection of a vulnerable biological resource, *Natrural Resource Modeling* **3**, 463 – 480.

Roe, G. H. and Baker, M. B.: 2007, Why is climate sensitivity so unpredictable?, *Science* **318**(5850), 629 – 632.

Roughgarden, J. and Smith, F.: 1996, Why fisheries collapse and what to do about it, *Proc. Natl. Acad. Sci.* **93**, 5078 – 5083.

Serruya, C. and Pollingher, U.: 1977, Lowering of water level and algal biomass in lake kinneret, *Hydrobiologia* **54**(1), 73–80.

Sethi, G., Costello, C., Fisher, A., Hanemann, M. and Karp, L.: 2005, Fishery management under multiple uncertainty, *Journal of Environmental Economics and Management* **50**(2), 300–318.

Singh, R., Weninger, Q. and Doyle, M.: 2006, Fisheries management with stock growth uncertainty and costly capital adjustment, *Journal of Environmental Economics and Management* **52**(2), 582 – 599.

Stern, N.: 2007, *The Economics of Climate Change*, Cambridge University Press.

Tsur, Y. and Graham-Tomasi, T.: 1991, The buffer value of groundwater with stochastic surface water supplies, *Journal of Environmental Economics and Management* **21**, 201 – 224.

Tsur, Y. and Zemel, A.: 1994, Endangered species and natural resource exploitation: Extinction vs. coexistence, *Natural Resource Modeling* **8**, 389–413.

Tsur, Y. and Zemel, A.: 1995, Uncertainty and irreversibility in groundwater resource management, *Journal of Environmental Economics & Management* **29**, 149–161.

Tsur, Y. and Zemel, A.: 1996, Accounting for global warming risks: Resource management under event uncertainty, *Journal of Economic Dynamics & Control* **20**, 1289–1305.

Tsur, Y. and Zemel, A.: 1998, Pollution control in an uncertain environment, *Journal of Economic Dynamics & Control* **22**, 967–975.

Tsur, Y. and Zemel, A.: 2004, Endangered aquifers: Groundwater management under threats of catastrophic events, *Water Resources Research* **40**, 1–10.

Tsur, Y. and Zemel, A.: 2008, Regulating environmental threats, *Environmental and Resource Economics* **39**, 297–310.

Tsur, Y. and Zemel, A.: 2009, Endogenous discounting and climate policy, *Environmental and Resource Economics* **44**, 507–520.

Vardas, G. and Xepapadeas, A.: 2010, Model uncertainty, ambiguity and the precautionary principle: Implications for biodiversity management, *Environmental and Resource Economics* **45**, 379–404.

Weitzman, M. L.: 2009, On modeling and interpreting the economics of catastrophic climate change, *Review of Economics and Statistics* **91**, 1–19.

Wirl, F.: 2007, Energy prices and carbon taxes under uncertainty about global warming, *Environmental and Resource Economics* **36**(3), 313 – 340.

Yaari, M. E.: 1965, Uncertain lifetime, life insurance, and the theory of the consumer, *Review of Economic Studies* **32**(2), 137–150.