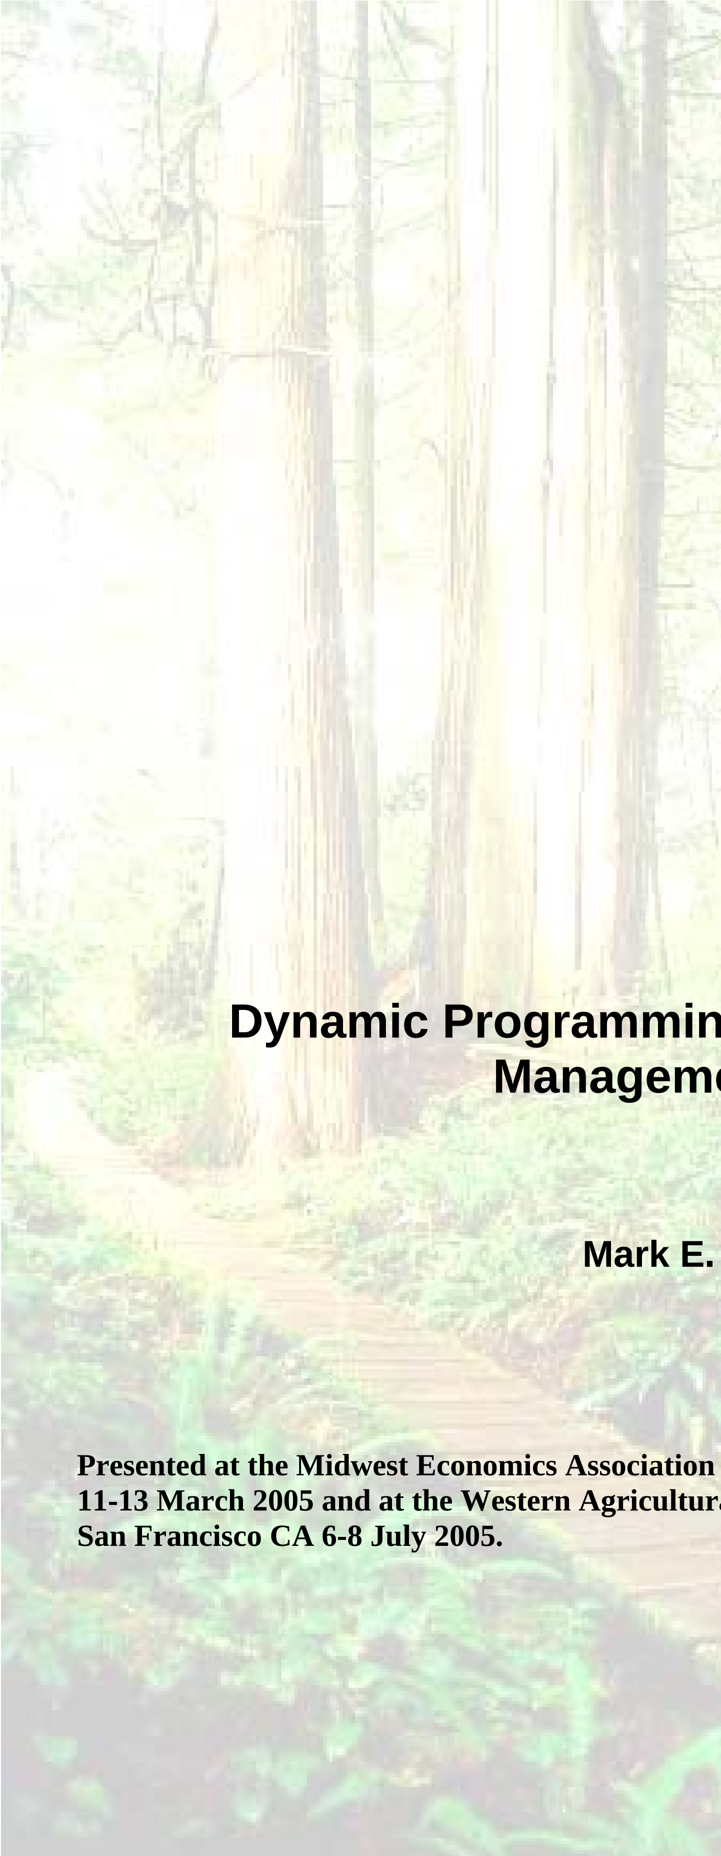

WORKING PAPER
2005-07

Resource Economics and Policy Analysis (REPA) Research Group

Department of Economics
University of Victoria

Dynamic Programming and Learning Models for Management of a Nonnative Species

Mark E. Eiswerth, G. Cornelis van Kooten, Jeff M. Lines and Alison J. Eagle

Presented at the Midwest Economics Association annual meetings in Milwaukee WI 11-13 March 2005 and at the Western Agricultural Economics Association annual meetings in San Francisco CA 6-8 July 2005.

July 2005

REPA Working Papers:

2003-01 – Compensation for Wildlife Damage: Habitat Conversion, Species Preservation and Local Welfare (Rondeau & Bulte)
2003-02 – Demand for Wildlife Hunting in British Columbia (Sun, van Kooten, & Voss)
2003-03 – Does Inclusion of Landowners' Non-Market Values Lower Costs of Creating Carbon Forest Sinks? (Shaikh, Suchánek, Sun, and van Kooten)
2003-04 – Smoke and Mirrors: The Kyoto Protocol and Beyond (van Kooten)
2003-05 – Creating Carbon Offsets in Agriculture through No-Till Cultivation: A Meta-Analysis of Costs and Carbon Benefits (Manley, van Kooten, Moeltner, and Johnson)
2003-06 – Climate Change and Forest Ecosystem Sinks: Economic Analysis (van Kooten and Eagle)
2003-07 – Resolving Range Conflict in Nevada? The Potential for Compensation via Monetary Payouts and Grazing Alternatives (Hobby and van Kooten)
2003-08 – Social Dilemmas and Public Range Management: Results from the Nevada Ranch Survey (van Kooten, Thomsen, Hobby, and Eagle)
2004-01 – How Costly are Carbon Offsets? A Meta-Analysis of Forest Carbon Sinks (van Kooten, Eagle, Manley, and Smolak)
2004-02 – Managing Forests for Multiple Tradeoffs: Compromising on Timber, Carbon and Biodiversity Objectives (Krcmar, van Kooten, and Vertinsky)
2004-03 – Tests of the EKC Hypothesis using CO2 Panel Data (Shi)
2004-04 – Are Log Markets Competitive? Empirical Evidence and Implications for Canada-U.S. Trade in Softwood Lumber (Niquidet and van Kooten)
2004-05 – Conservation Payments under Risk: A Stochastic Dominance Approach (Benítez, Kuosmanen, Olschewski and van Kooten)
2004-06 – Modeling Alternative Zoning Strategies in Forest Management (Krcmar, Vertinsky, and van Kooten)
2004-07 – Another Look at the Income Elasticity of Non-Point Source Air Pollutants: A Semiparametric Approach (Roy and van Kooten)
2004-08 – Anthropogenic and Natural Determinants of the Population of a Sensitive Species: Sage Grouse in Nevada (van Kooten, Eagle, and Eiswerth)
2004-09 – Demand for Wildlife Hunting in British Columbia (Sun, van Kooten, and Voss)
2004-10 – Viability of Carbon Offset Generating Projects in Boreal Ontario (Biggs and Laaksonen-Craig)
2004-11 – Economics of Forest and Agricultural Carbon Sinks (van Kooten)
2004-12 – Economic Dynamics of Tree Planting for Carbon Uptake on Marginal Agricultural Lands (van Kooten) (Copy of paper published in the Canadian Journal of Agricultural Economics 48(March): 51-65.)
2004-13 – Decoupling Farm Payments: Experience in the US, Canada, and Europe (Ogg & van Kooten)
2004–14 – Afforestation Generated Kyoto Compliant Carbon Offsets: A Case Study in Northeastern Ontario (Jeff Biggs)
2005–01 – Utility-scale Wind Power: Impacts of Increased Penetration (Pitt, van Kooten, Love and Djilali)
2005–02 – Integrating Wind Power in Electricity Grids: An Economic Analysis (Liu, van Kooten and Pitt)

2005–03 – Resolving Canada-U.S. Trade Disputes in Agriculture and Forestry: Lessons from Lumber (Biggs, Laaksonen-Craig, Niquidet and van Kooten)
2005–04 – Can Forest Management Strategies Sustain The Development Needs Of The Little Red River Cree First Nation? (Krcmar, Nelson, van Kooten, Vertinsky and Webb)
2005–05 – Economics of Forest and Agricultural Carbon Sinks (van Kooten)
2005–06 – Divergence Between WTA & WTP Revisited: Livestock Grazing on Public Range (Sun, van Kooten and Voss)
2005-07- Dynamic Programming and Learning Models for Management of a Nonnative Species (Eiswerth, van Kooten, Lines and Eagle)

For copies of this or other REPA working papers contact:

REPA Research Group
Department of Economics
University of Victoria PO Box 1700 STN CSC Victoria, BC V8W 2Y2 CANADA
Ph: 250.472.4415
Fax: 250.721.6214
http://repa.econ.uvic.ca

Dynamic Programming and Learning Models for Management of a Nonnative Species

Mark E. Eiswerth,[a] G. Cornelis van Kooten,[b] Jeff M. Lines[c] and Alison J. Eagle[b]

[a]Department of Economics, University of Wisconsin-Whitewater
[b]Department of Economics, University of Victoria
[c]Department of Economics, University of Toronto

Abstract

Nonnative invasive species result in sizeable economic damages and expensive control costs. Because dynamic optimization models break down if controls depend in complex ways on past controls, non-uniform or scale-dependent spatial attributes, etc., decision support systems that allow learning may be preferred. We compare three models of an invasive weed in California's grazing lands: (1) a stochastic dynamic programming model, (2) a reinforcement-based, experience-weighted attraction (EWA) learning model, and (3) an EWA model that also includes stochastic forage growth and penalties for repeated application of environmentally harmful control techniques. Results indicate that EWA learning models may be appropriate for invasive species management.

Key words: Invasive weed species; optimal control; adaptive management

Acknowledgements: The authors wish to thank Steve Schoenig of the California Department of Food and Agriculture, Joseph DiTomaso of the University of California – Davis, and Wayne Johnson of the University of Nevada for suggestions and comments in the survey design phases of this work. Earlier versions of this paper were presented at the Midwest Economics Association annual meetings in Milwaukee WI 11-13 March 2005 and at the Western Agricultural Economics Association annual meetings in San Francisco CA 6-8 July 2005.

Introduction

Non-indigenous plant species established in the United States include an estimated fifty thousand species that have escaped and now exist in natural ecosystems, threatening survival of native species and causing damage and control costs of approximately $137 billion annually (Pimentel et al. 2000). As one example, yellow starthistle (*Centaurea solstitialis L.*, hereafter YST) is a non-indigenous invasive plant species native to Eurasia that is believed to have been introduced to California as a contaminant in imported alfalfa seed in the 19th century (DiTomaso and Gerlach 2000). It has become naturalized in much of the U.S., but the heaviest infestations are in western states including California, Idaho, Oregon and Washington (USDA 2005; USGS 2005). California has become particularly heavily infested since the species' accidental introduction. Surveys of county agricultural commissioners in California reveal that the area infested by YST has increased significantly over the past five decades, from 1.2 million acres in 1958 to 1.9 million acres in 1965, 7.9 million acres in 1985, and 14.3 million acres in 2002 (Maddox and Mayfield 1985; Pitcairn et al. 2004). As an annual plant, the weed spreads through prolific seed production, with dispersion aided by birds and, more commonly, human movement and activities such as road building. Each plant is capable of producing up to 100,000 seeds (DiTomaso 2001) of which approximately 95% are viable (Lass et al. 1999); some seeds remain productive for as long as ten years, posing significant challenges to control or eradication efforts. Without efforts to control YST, potential infestation could reach upwards of 40 million acres of California grasslands (Jetter et al. 2003).

Negative impacts of YST infestation include reductions in native plant diversity, decreased grazing yields, reduced accessibility to recreational trails for hiking and horses, and decreased water for agriculture and aquatic life. Due to the scope of the infestation, eradication is impossible (Eiswerth and van Kooten 2002; Thomsen et al. 1993). Yet, the State of California has deemed worthwhile efforts to control YST as damages can be large. For example, it has been estimated that YST alone results in a water loss in the Sacramento Valley valued at $16-$56 million per year (Dudley 2000; Gerlach et al. 1998).

While the biology of YST and its control have been studied at length, surprisingly little is known about the economics of YST infestations. YST control has both a private and public good component. Biological controls (e.g., flies that lay eggs in the seed head) benefit landowners dispersed across a large area, while direct controls (e.g., burning, chemical spraying) that reduce YST infestation on one field have a spillover benefit as they can mitigate the spread to adjacent (and even more distant) lands. Therefore, public expenditures on YST control are warranted. But so are private efforts to control YST because, by reducing infestation levels on one's own fields, the landowner can increase forage production, lower veterinarian costs and reduce the costs of preventing further spread in the future. However, private efforts to control YST infestations also have negative externalities: Burning range to control YST can result in escapes that lead to wildfires, while herbicide use can lead to reduced ecosystem resiliency and human health problems.

For California, Jetter et al. (2003) calculated the costs and benefits of a YST biological control program. Their estimates of benefits were based on the views of appraisers concerning land values with and without YST infestation, which assumes that all the costs of YST infestation are reflected in land prices.[1] The researchers estimate that an effective biological control program results in benefits of some $40 million to $1.412 billion, depending on whether one considers only acres infested or those susceptible to YST.

Economists generally employ dynamic optimization models to determine the most effective means for controlling weed infestations. Depending on whether models are open or closed loop in form, their solution provides a decision maker with the optimal strategy to pursue for each level of infestation at each point in time (Eiswerth and van Kooten 2002). Dynamic optimization models break down, however, if current decisions depend not only on the current state of the system but on the state of the system and controls used in the past, on spatial attributes that are not uniform or scale-invariant, on existence of multiple stable and unstable equilibria, on potential irreversibility, and so on; uncertainty is ubiquitous and there is insufficient information to address it. These modeling problems are well known in the context of weed management (Jones and Medd 2000; Wilkerson et al. 2002) and ecology more generally (Holling and Meffe 1996). Thus, decision support systems that may be helpful in educating the public and in broader policy formation are often preferred to dynamic optimization tools. In the case of weeds such as YST, many decision-makers advocate taking a systems approach, where goals are carefully articulated and the long-term impacts of control efforts are regularly monitored (DiTomaso et al. 2000; Randall 1996). The purpose of this paper, therefore, is to compare

[1] For land that is identical in every respect except for extent of YST infestation, appraisers concluded that there would be no difference in price, only in the time taken to sell the property. The foregone interest on the property's value for one month was used by Jetter et al. (2003) as the benefit of non-infested versus infested land.

different decision-making tools that seek to implement a control strategy that is optimal from both an economic and ecological perspective.

In particular, the objective in this paper is to compare the results from three dynamic decision-making models for managing a nonnative weed species, in this case YST. Each model uses essentially the same data and the general objective of each is to maximize the future stream of discounted revenues accruing from domestic livestock grazing on YST-infested lands. We are not concerned about spillover costs or benefits. The evolving stock of YST leads to reductions in the level of the state variable (forage), but is influenced by human management interventions in stochastic fashion. Since published data on the YST stock dynamics and responses to human control are almost nonexistent, we employ data derived from an expert opinion survey administered to California land managers, weed scientists and other professionals (Eiswerth and van Kooten, unpubl. data, 2002), and data from a survey administered to agricultural producers in the state of California (Eiswerth and van Kooten, unpubl. data, 2003).

The three models that we employ are (1) a stochastic dynamic programming (SDP) model, (2) a reinforcement-based experience-weighted attraction (EWA) learning model similar to that used by Hanaki et al. (2005), and (3) a model similar to (2) except that it introduces more information and specifies a stochastic forage growth equation in which the intrinsic growth rate is a function of precipitation. EWA learning models may be an appropriate tool for this context given substantial uncertainties regarding control effectiveness and the variation in effectiveness from period to period. In the following sections, we provide more background on the biology of YST, develop the decision-making models in more detail, describe the data, and present and compare the results of the models.

Yellow Starthistle: Background to an Invasive Weed Species

YST competes vigorously with native plants, especially grasses, that are the staple for livestock grazing on rangeland, and also reduces yield and quality in non-native pasture and cultivated crops. Deep roots enable the weed to remove soil moisture at depths greater than six feet, competing with perennial grasses and causing drought conditions for native species where YST infestation is high (Gerlach et al. 1998). Although YST provides some forage value in early growth stages, the spiny nature of the weed means that livestock and wildlife avoid grazing in heavily infested areas, as the spines cause damage and discomfort to grazing animals. Pastures infested with YST contain considerably less crude protein and total digestible nutrients compared to uninfested pastures (Barry 1995). Prolonged ingestion of YST by horses causes a mostly fatal neurological disease called equine nigropallidal encephalomalacia (ENE) or "chewing disease" (Cordy 1978).

Burning, cultivation, mowing, timed grazing, application of chemical herbicides, and biological controls have been tried at various times against YST. The weed's tendency to germinate throughout the rainy season and to re-grow after mowing or grazing means that such control practices must be repeated throughout the growing season. Regular prescribed burns and controlled grazing reduce YST seedbank stocks, seedling density, and mature vegetative cover, but, because seeds can remain viable for many years, new plants can establish in subsequent years (Kyser and DiTomaso 2002; Thomsen et al. 1993). Thus, one study showed that burned grassland was not significantly different from unburned grassland after four years because of YST recovery (Kyser and DiTomaso 2002). Effective biological control agents introduced into California are seedhead attackers (flies and weevils), but seed reduction in areas where these insects are well-established has not been sufficient to reduce starthistle abundance (Jetter et al. 2003). Research into additional natural enemies is ongoing. As a result of the ineffectiveness of individual control methods and the persistence of the YST seed bank, effective management may require a combination of methods on a long-term basis.

An important consideration in adopting a particular strategy is its effectiveness and ecological impacts over time. Some controls will not work well because, if applied too often in a given period, they not only reduce the YST infestation but also damage the overall ecosystem by reducing the abundance of ‘good’ forbs relative to weeds and other forbs that are more resistant to the particular herbicides used to control YST, such as 2,4-D, triclopyr, dicamba and glyphosate (Jetter et al. 2003, p. 229). In particular, despite the economic viability of such a strategy, repeated use of chemical controls will have significant negative impacts on the surrounding ecosystem and the strategy may lose effectiveness over time as YST builds resistance to the chemical (Jetter et al. 2003, p.229).

As noted by Holling and Meffe (1996), because “ecosystems are moving targets, with multiple potential futures that are uncertain and unpredictable”, management has to be flexible and adaptive rather than based on command and control that results from most dynamic optimization approaches. Therefore, an integrated adaptive-management approach to reducing YST infestations in the long run is likely more desirable.

Decision-Making Methods

We employ three different programming approaches, each representing a decision-maker seeking to find a dynamically 'optimal' control strategy under conditions of uncertainty of varying degrees. For the current purposes, all three approaches are characterized by a common state transition probability matrix that describes the probability of a YST infestation increasing or decreasing given the application of various control methods, including no control.[2] The models differ in how they subjectively interpret this information. The objective probability of a control's true effectiveness is unknown to the decision-maker.

Stochastic dynamic programming model

The stochastic dynamic programming model offers one way to deal with uncertainty related to the effectiveness of controls (decisions about what to do) and how the state variable (weed infestation level, forage for livestock) evolves as a function of precipitation, intrinsic growth and human interventions. Since first introduced by Burt and Allison (1963), SDP has been used to determine optimal dynamic decisions for natural resource and agricultural decision making (see Kennedy 1986). The SDP model of the current study deals with optimal decision making on the part of agricultural producers faced with a harmful biological invasion.

The producer's objective is to maximize the present value of a future stream of net returns from the enterprise, where the enterprise consists of forage production and

[2] We use the same state transition probability matrices for convenience only. In practice, the transition matrices would not be markovian, depending on the level of the state variable and the controls applied in the preceding several periods. Further, we do not employ a continuous-time method for solving SDP problems because of the nature of the available data (e.g., see Miranda and Fackler 2002).

grazing on native and improved pastures. The per-acre net returns at time t are affected by the extent of infestation of YST at time t (x_t). The producer's objective function is:

$$(1) \qquad \sum_{t=0}^{T-1} \rho^{t} R_{t}(x_{t}, k_{t}) + \rho^{T} S(x_{T}),$$

where k_t is the choice of technology for YST control, $S(x_T)$ is the value of the producer's land in end period T as a function of the YST infestation level (salvage value), and ρ is the discount factor. Our analysis is partial in the sense that it focuses on the impact of the invasive species on net revenues, ignoring other determinants of R. Specifically, we assume that the net revenue is only affected by the yield (Y) reduction caused by YST and YST management costs. We thus define R as the per-acre net revenue exclusive of YST control costs (i.e., net returns include non-YST-related production costs) and $c(k_t)$ as the per-acre cost of the YST control strategy k at time t – production costs are invariant with YST infestation.

The general form of the equation of motion for the invasive species stock x_t is:

$$(2) \qquad x_{t+1} = g(x_{t}, k_{t}) + \varepsilon_{t},$$

where ε is a random variable with zero mean and variance σ^2, and the initial condition is $x_0 = \bar{x}$. Equation (2) is the Markov condition: current infestation levels are a function only of last period's infestation level and the control applied. The evolution of the YST stock over time thus depends on the level of the infestation in period t, the choice of YST control option in period t, and stochasticity given by ε.

We parameterize the model, including the stochastic state equation (2), using data collected from an expert judgment survey of weed scientists, county farm advisors, public land managers and other specialists familiar with YST (see below). The Bellman equation for the SDP problem is:

$$\text{(3)} \qquad \mathrm{V}_{\mathrm{t}}(x_{\mathrm{t}}, k_{\mathrm{t}}) = \max_{k_0, k_1, \ldots, k_{T-1}} \{\mathrm{E}[R(x_{\mathrm{t}}) - c(k_{\mathrm{t}})] + \rho \sum_{j=1}^{k} P(i, j, k_{\mathrm{t}}) \, \mathrm{V}_{\mathrm{t+1}}[x_{\mathrm{t+1}}(j)]\}$$

where $P(i,j,k_t)$ represents the probability that a biological invasion of state i (i=1,…,n) in period t will transition to state j (j=1,…,n) by period (t+1), given that control option k is chosen in period t. V_t denotes the expected discounted value of the future stream of net revenues in period t, given the level of the invasive weed stock in period t and assuming that the optimal path is taken in every future period.

To implement the SDP approach, a routine was written in MATLAB to solve the Bellman equation (3). The data used for the model are described later in this paper.

Experience-weighted attraction (EWA) model

While the SDP model allows for uncertainty in current period returns and in the evolution of the system (i.e., state equation), it neglects other types of uncertainty. In particular, it does not address uncertainty related to the control itself, and the fact that the effectiveness of a control might be diminished if it is applied for several periods in a row. For example, sequential application of herbicides will cause YST to become resistant, reducing the ability to continue to apply chemicals. Burning can destroy the roots of valuable perennials in a range ecosystem if burns follow each other too closely, while both burning and chemicals can damage ecosystems because ‘good’ forbs have been targeted by accident. These types of uncertainty are important features of the manager's problem and therefore it is desirable to test other types of tools. Reinforcement-based, experience-weighted attraction (EWA) learning models that are employed in game theory (Camerer and Ho 1999; Hanaki et al. 2004) are potential frameworks that could accommodate these types of uncertainty.

Our EWA models simulate a decision-maker who learns about the effectiveness of different strategies in a preliminary experimental phase before taking this information to a second phase. In the second phase, decisions are based on the observed results from the experimental phase as well as results from previous periods. In each period, the decision-maker has information about the net benefit from particular strategies that have been used in the past. The decision-maker tracks the average net benefit, or average payoff, that has resulted by time period τ from choosing strategy s in any given number of time periods in the past. These average payoffs are then termed the 'attractions' to strategy s by time period τ (denoted as $A_{\tau,s}$) and are calculated according to:

$$(4)\quad A_{\tau,s} = \frac{\sum_{t=1}^{\tau} R_{t,s}}{\sum_{t=1}^{\tau} d_{t,s}},$$

where $R_{t,s}$ denotes net returns in period t from selecting strategy s, and $d_{t,s}$ is a binary indicator variable equal to one if strategy s is chosen in period t, and zero otherwise.

Next, the model converts the 'attractions' to each strategy into probabilities of selecting the various strategies. The probability of selecting strategy s in time period t

depends on the attractions as follows (Camerer and Ho 1999, p.835):

$$(5)\quad p_{t,s} = \frac{e^{\lambda A_{t,s}}}{\sum_{k \in S} e^{\lambda A_{t,k}}}$$

where $A_{t,s}$ is the attraction at time t to strategy s (as before) and the parameter $\lambda \geq 0$ represents the extent to which strategies with higher attractions are favored in strategy choice. When λ=0, all strategies are equally likely to be selected. As λ increases, strategies with higher attractions increasingly have a greater probability of being selected for decreasing differences in attractions between strategies. In the experimental phase, λ is set to zero so all strategies are used and results can be obtained for each one. Then, in the second phase, λ is increased (in our case) to 0.25 so that the better strategies are chosen with greater probability. The value of λ is chosen to give the highest average performance in terms of accumulated net returns over time. Although it may be expected that a higher value for λ, where the highest paying strategy is chosen with greater probability, would better maximize profits, this is not necessarily true because the benefits of any particular strategy may be realized over several periods following the one in which the control is applied. As in SDP and other dynamic optimization approaches, strategies that do not provide high current benefits may lead to greater benefits in future periods thereby yielding a higher net present value (NPV). In that sense, EWA is like dynamic optimization, with some strategies given significant probability weight despite providing lower single-period payoffs because they lead to higher profits over the longer run.

To model the evolution of the stock of YST within the EWA decision model, we use transition probability matrices *P(i,j,s)* of the same form and embodying the same data as those used in the SDP approach (data described below). The EWA model was also solved in MATLAB using a routine that selects a YST control strategy in each time period t based on the attraction-based probabilities for each strategy s in (5). Once a strategy is selected, the stock (more precisely, percent cover) of YST (x) evolves between period t and period t+1 according to the appropriate strategy-specific transition matrix *P(i,j,s)*. Discounted net returns to agriculture for that period are then computed as $R(x)–c(s))/(1+r)^t$, in similar fashion to the SDP model. Unlike the SDP model, however, the new time period's net returns are then used to update the learner's information on the historical average payoff for that chosen strategy. Specifically, historical annual net returns, by YST control strategy, are summed in each period to compute cumulative and average net returns for each strategy. These average payoffs by strategy are then used to update the attractions and probabilities in (4) and (5) for the next round of strategy selection. In our analysis, the learning period was set to 2,000 years to establish the strategy attractions, followed by a 75-year period of strategy selection over which the final net returns are tracked. This two-phase learning/implementation process can be repeated any number of times; we conducted 30 iterations.

Enhanced EWA model

We refer to the third decision-making model as 'enhanced EWA' since it is similar to the EWA model described above, except that it provides more information to the decision-maker. Specifically, the EWA model is similar to the SDP model in that the amount of livestock forage in any period t is a simple function of the stock of YST at t, based on data from an expert judgment survey described in the next section. In contrast, the enhanced EWA model introduces more information via a forage growth equation:

$$(6)\quad F_t - F_{t-1} = \gamma \times PR_t \times \left(1 - \frac{F_{t-1}}{K_t(1-\eta(YST))}\right),$$

where PR refers to the amount of precipitation in period t relative to historical mean precipitation, K_t is the maximum forage carrying capacity or maximum animal unit months (AUMs) that can be grazed in period t in the absence of YST, γ is the intrinsic growth rate of the forage stock, and η ($0 \leq \eta < 1$) is an adjustment parameter describing the reduction in carrying capacity due to the presence of YST (see Holechek, Pieper and Herbel 1989). Precipitation is a random normal variable with mean and variance given by historical data from the Sonora, CA weather station where the climate is representative of our study area.[3] The intrinsic growth rate γ is multiplied by the ratio of current-period precipitation to mean precipitation.

To ensure adequate forage for the future and provide some resiliency of the range ecosystem, we include in the model as a management criterion that grazing can occur on a field only if available forage (biomass) exceeds half of the carrying capacity. To reflect the ecological benefits of a diversified control strategy and the problems discussed previously, we introduce penalties in the 'enhanced EWA' model when consecutive instances of burning or herbicide applications are implemented as controls. The penalties increase with the number of times the same strategy is used over a specific interval so

[3] The website is www.wrcc.dri.edu/summary/climsmnca.html.

that the agent will learn not to repeat the same control in consecutive periods. The penalty function for a given time period i (in which strategy s is chosen) is of the form:

$$(7)\quad pen_{\mathrm{i,s}} = \alpha\left(\sum_{\mathrm{t}=i-10}^{\mathrm{i}} x_{\mathrm{s}}\right)^{\beta}, \text{ where } x = \begin{cases} 1 \text{ if strategy } s \text{ is used in period } i \\ 0 \text{ otherwise} \end{cases},$$

and α and β (α, β>1) are parameters that may be adjusted to determine the degree of penalty (*pen*).

The enhanced EWA model was solved by writing an algorithm in MATLAB similar to that described for the EWA model, with an experimental learning phase of 2,000 years and second phase of 75 years. In the enhanced EWA simulations, the baseline values of all parameters were set equal to those in EWA, with additional baseline parameter values of α=1, β=2 (equation 7) and γ=1.18 (equation 6). Data for η, which in equation (6) describes the reduction in forage carrying capacity as YST increases, were collected from the expert judgment survey discussed below.

Data used in the Models

There is very little published data on the influence of various management strategies on YST. Further, the effectiveness of management strategies is sensitive to a number of site-specific characteristics. In order to collect data that is useful for decision-making models, we implemented the *Expert Judgment Questionnaire: Expansion Rate Behavior of Yellow Starthistle Infestations* in partnership with the California Department of Food and Agriculture in 2002 (see Eiswerth and van Kooten 2002). The survey sample frame included extension specialists, weed scientists, county farm advisors, public land managers and other specialists familiar with YST and YST control techniques in California. Since the relevant group of knowledgeable professionals is relatively small, we only surveyed 21 experts. Since uncertainty and high variability are salient features of the key relationships in which we are interested (e.g., state variable transition probabilities, impact of YST on usable forage), the survey was designed to allow respondents to reflect uncertainties in their estimates.

The survey collected several types of information, but for the purposes of this study two types of data are most relevant. First, we obtained estimates of the quantitative reduction in livestock forage as a function of YST infestation. For example, for rangeland grazing contexts (we assume a baseline potential of 2 AUMs per acre per year), the estimated reductions in grazing potential are 6-10% for minimal YST infestations, 22-28% for moderate, 38-50% for high, and 60-78% for very high YST infestations. The survey also collected estimates of grazing reduction for two other land profiles: irrigated tall wheatgrass forage with one growing season per year (baseline productivity of 5 AUM ac^{-1} yr^{-1}), and spring hay harvest plus summer grazing on tall fescue/irrigated orchardgrass (baseline productivity of 10 AUM ac^{-1} yr^{-1}). We used the means of the estimated ranges to estimate net grazing returns contingent upon the state of YST infestation and baseline grazing productivity.

Second, the survey elicited responses regarding different YST control strategies and the likelihood of their success. Experts provided assessments on the probabilities that, given a field has a current level of YST infestation (which was given by a descriptive category, such as ‘minimal’) and a particular control (burning, herbicide application, etc.) is employed, it will transition to any of the four possible future states (including the probability of remaining in the same state in the next period). From these responses, we generated five subjective probability transition matrices, one for each control option, which were used to determine a stochastic YST invasion process for each of the three models. The five control options available to the decision maker are as follows:

1. do nothing to control YST;
2. a one-time chemical weed control program without follow-up treatment;
3. any of a number of combinations of strategies (e.g., chemical, burning, mowing, grazing) that results in ‘successful weed management’ or ‘best practice’ (as defined individually by survey respondents) without follow-up chemical treatment;
4. as 3, but with follow-up treatment; and
5. as 3, plus a program of site revegetation with desirable species, specifically perennial grasses or forbs that provide palatable forage and good protein content for livestock.

Control costs increase as one moves from options 1 to 5, with the degree of YST control rising as well. Costs for control strategies are assumed to be \$0 for no control (NC); \$3.50 for chemical (CH); \$13.50 for best practice (BP); \$25.85 for best practice plus follow up management (BP+F); and \$43.50 for best practice plus revegetation (BP+R). The value of an AUM is set equal to \$25 (University of California Cooperative Extension, unpubl. data; Resource Concepts, Inc. 2001).

Results

The YST management strategies selected by the SDP model are indicated in Table 1. Solution of the SDP model using an infinite time horizon results in the selection of a single 'best' approach (based on maximum long-run net returns to grazing) for each combination of initial YST state condition, land productivity type, and assumed discount rate. As the baseline grazing potential of the land rises, the optimal strategy choice begins to transition from relatively low-cost techniques (CH) to higher-cost strategies (BP+F or BP+R), as expected. The optimal strategy also tends to gravitate toward higher-cost approaches as the initial state of YST infestation worsens.

A salient feature of the SDP approach is that some control should be used in each time period. This essentially is because the SDP model chooses the strategy that maximizes long-run expected net return, and not controlling YST does not maximize long-run net return under any of the conditions assumed. However, the results in Table 1 refer to the optimal strategy for each level of the state variable, regardless of time. Hence, it does not take into account the possibility that consecutive application of the same control can result in reduced effectiveness of the strategy. This begs two questions: First, how would the results of a learning model such as *EWA* differ from those of *SDP* given that the decision-maker can learn over time about strategy effectiveness? Second, how would the introduction of penalties for repeated strategies (to reflect environmental externalities), as in *EWA-enhanced* approach, alter the 'optimal' pattern of strategies over time?

The results presented in Table 2 shed some light on these questions. In the table, we provide the proportional choices of various strategies over a 75-year time horizon following the learning period of 2000 years. The number of iterations (n) is 30. The most frequently chosen strategies under *EWA* are NC (no control) and CH (chemical application), with NC chosen more frequently than CH for low land productivities and

CH chosen more frequently than NC at higher land productivities, as might be expected. The choice proportion for BP+F is relatively low but rises to 9% as the baseline productivity of the land increases to its highest level. Mean choice proportions for the other strategies (BP+F, BP+R) are relatively low but also rise as land becomes relatively more productive (e.g., mean choice proportion for BP+F is 7% assuming 10 AUM ac^{-1} yr^{-1}). Again, this is as expected since benefits of YST control increase with land productivity.

To summarize, *EWA* differs from *SDP* in a few aspects. First, NC (no control) is frequently chosen (and even more with *EWA-enhanced*). Second, it appears that the higher-cost options involving 'best practice' display attractions high enough to draw implementation after the learning period, but the relative frequency with which they would be chosen is modest. *EWA* may provide an indication of the frequency with which decision-makers may wish to apply such methods.

The results in Table 2 for the *EWA-enhanced* model display the same general patterns as the *EWA* model results. However, as expected the strategy choice proportions from *EWA-enhanced* are weighted more towards no control, reflecting the impact of the penalty functions for repeated strategy implementation. The mean choice proportions for each of the active management strategies are smaller under the *EWA-enhanced* model. In other words, accounting for potential external ecosystem costs results in a reduction in the frequency of applying control strategies.

The performance of the three models is compared in Table 3 in terms of mean net present value (NPV) of returns to agriculture, and the standard deviation of returns.[4] *SDP* yields the highest values of NPV, while the *EWA-enhanced* model yields the lowest average NPV. The information provided by the *EWA–enhanced* model results in lower agricultural returns, but also significantly reduced financial risk compared to *EWA*. However, because *EWA* and particularly *EWA-enhanced* employ fewer controls in every situation, it implicitly leads to higher long-run ecological returns, which are not explicitly taken into account in any of the models.

Conclusions

The analysis presented in this paper needs to be extended in several ways. The focus of the analysis here is primarily on exploring and comparing the properties of alternative decision-making models and the results they yield when they use the same uncertain data. That is, the three models consider only one source of uncertainty, but ignore uncertainty related to the ecosystem more generally, namely, the problems of chemical resistance, loss of 'good' forbs and perennials as a result of consecutive burning or chemical applications, and so on. Given that the *EWA* and *enhanced EWA* models are designed to address these more complicated forms of uncertainty, the comparisons that are made in this study are biased in favor of the SDP model, which also provides the greatest long-run expected net returns. However, one needs to keep in mind that the

[4] The standard deviation of returns is calculated differently for the SDP model than for the EWA and EWA-enhanced approaches. For SDP, it is calculated in the standard way from the markov solution's long-run transition probability matrix; for the experience-weighted approaches it is determined from the monte carlo simulations. As n increases, therefore, one would expect the standard deviations to decline. As a result, however, we do not compare the SDP model with the others on the basis of the dispersion of returns.

paucity of data on YST growth and spread, and its response to management, imposes a real obstacle to dynamic optimization modeling. Therefore, there is a paramount need to explore properties of decision-support systems that can (1) accommodate substantial uncertainty, which may be fuzzy, but (2) also allow the decision-maker to learn. At the same time, better data need to be generated by intermediate- and long-term scientific field studies before definitive conclusions can be reached about the optimal frequencies of various management strategies.

With the particular caveat above, we offer the preliminary results of a learning model that includes penalty functions for repeated management strategies. Thus far our focus in that model has been solely on exploring the general sensitivity of the learning model results to the incorporation of a stochastic forage growth equation and introduction of a generic penalty function. Better empirical data on the externalities of the various YST control strategies (chiefly, chemical controls and controlled burns) are necessary before one may view results of the *EWA-enhanced* model as accurately accounting for these external effects. Incorporating such data, however, would eventually allow a decision-maker to optimize over a wider set of (economic and environmental) management objectives.

Because ecosytems are incredibly complex, ‘command-and-control’ management (as provided by most dynamic optimization models) is naïve at best, because it assumes that all future contingencies can be appropriately taken into account. Because of complexity and uncertainty, adaptive management, such as that offered by the EWA models, is often preferred. Indeed, by erring on the ‘conservative’ side, the *EWA-enhanced* model presented here leads to resource management that more closely applies

Holling and Meffe's (1996) 'Golden Rule' of natural resource management: "Natural resource management should strive to retain critical types and ranges of natural variation in ecosystems. That is, management should facilitate existing processes and variabilities rather than changing or controlling them" (p.334).

References

Barry, S., 1995. Cattle Fecal Fax, October Report. UCCE Newsletter. Red Bluff, CA: University of California Cooperative Extension.

Benefield, C. B., J. M. DiTomaso, et al. 1999. "Success of mowing to control yellow starthistle depends on timing and plant's branching form." *California Agriculture* 53(2): 17-21.

Burt, O.R. and J.R. Allison, 1963. Farm Management Decisions with Dynamic Programming. *Journal of Farm Economics* 45(1): 121-136.

Camerer, C. and T-H. Ho, 1999. Experienced-Weighted Attraction Learning in Normal Form Games, *Econometrica* 67(July): 827-874.

Cordy, D.R., 1978. *Centaurea* Species and Equine Nigropallidal Encephalomalacia. In *Effects of Poisonous Plants on Livestock*, edited by R.F. Keeler, K.R. Van Kampen and L.F. James. New York: Academic Press.

DiTomaso, J.M., 2001. *Yellow Starthistle Information* University of California, Davis [accessed 27 May 2004]. Available from http://wric.ucdavis.edu/yst/.

DiTomaso, J.M., and J.D. Gerlach. 2000. *Centaurea solstitialis* L. In: C.C. Bossard, J.M. Randall, and M. Hoshovsky. Invasive Plants of California's Wildlands. University of California Press, Berkeley, CA. pp. 101-106.

DiTomaso, J. M., G. B. Kyser, et al. 1999. "New growth regulator herbicide provides excellent control of yellow starthistle." *California Agriculture* 53(2): 12-16.

DiTomaso, J.M., G.B. Kyser S.B. Orloff, and S.F. Enloe, 2000. Integrated Strategies Offer Site-Specific Control of Yellow Starthistle. *California Agriculture* 54(6): 30-36.

Dudley, D.R., 2000. Wicked Weed of the West. *California Wild* 53: 32-35.

Eiswerth, M. and G.C. van Kooten, 2002. Uncertainty, Economics, and the Spread of an Invasive Plant Species. *American Journal of Agricultural Economics* 84(5): 1317-1322.

Gerlach, J.D., A. Dyer and K.J. Rice, 1998. Grassland and Foothill Woodland Ecosystems of the Central Valley. *Fremontia* 26(4): 39-43.

Hanaki, N., R. Sethi, I. Erev and A. Peterhansel, 2005. Learning Strategies. *Journal of Economic Behavior and Organization* 56(April): 523-542.

Holechek, Jerry L., Rex D. Pieper and Carlton H. Herbel, 1989. *Range Management Principles and Practices*. Englewood Cliffs, NJ: Prentice Hall.

Holling, C.S. and Gary K. Meffe, 1996. Command and Control and the Pathology of Natural Resource Management. *Conservation Biology* 10(April): 328-337.

Jetter, K.M., J.M. DiTomaso, D.J. Drake, K.M. Klonsky, M.J. Pitcairn and D.A. Sumner, 2003. Biological Control of Yellow Starthistle. In *Exotic Pests and Diseases: Biology and Economics for Biosecurity*, edited by D.A. Sumner. Ames: Blackwell Publishers.

Jones, R.E. and R.W. Medd, 2000. Economic Thresholds and the Case for Longer Term Approaches to Population Management of Weeds. *Weed Technology* 14(2): 337-350.

Kennedy, J.O.S., 1986. *Dynamic Programming. Applications to Agriculture and Natural Resources*. London and New York: Elsevier Applied Science Publishers.

Kyser, G.B. and J.M. DiTomaso, 2002. Instability in a Grassland Community after the Control of Yellow Starthistle (*Centaurea Solstitialis*) with Prescribed Burning. *Weed Science* 50(2): 648-657.

Lass, L.W., J.P. McCaffrey, D.C. Thill and R.H. Callihan, 1999. Yellow Starthistle: Biology and Management in Pasture and Rangeland. Cooperative Extension Bulletin. Moscow, ID: University of Idaho.

Maddox, D.M. and A. Mayfield, 1985. Yellow Starthistle Infestations Are on the Increase. *California Agriculture*: 10-12.

Miranda, Mario J. and Paul L. Fackler 2002. Applied Computational Economics and Finance. Cambridge, MA: The MIT Press.

Pimentel, D., L. Lach, R. Zuniga and D. Morrison, 2000. Environmental and Economic Costs of Nonindigenous Species in the United States. *BioScience* 50(1): 53-65.

Pitcairn, M.J., S. Schoenig, J. Gendron, R. Yacoub and R. O'Connell, 2004. Yellow Starthistle Infestations Are Still Increasing. Working Paper. Sacramento, CA: California Department of Food and Agriculture.

Randall, J.M., 1996. Weed Control for the Preservation of Biological Diversity. *Weed Technology* 10(2): 370-383.

Resource Concepts, Inc., 2001. *Nevada Grazing Statistics Report and Economic Analysis for Federal Lands in Nevada*. Carson City, NV: Resource Concepts, Inc., 2001.

Thomsen, C.D., W.A. Williams, M.P. Vayssiéres, F.L. Bell and M.R. George, 1993. Controlled Grazing on Annual Grassland Decreases Yellow Starthistle. *California Agriculture* 47(6): 36-40.

USDA (U.S. Department of Agriculture). 2005. Plants profile: *Centaurea solstitialis* L. Available on: http://plants.usda.gov [accessed February 2005].

USGS (U.S. Geological Survey). 2005. *Centaurea solstitialis* L. Available on: http://el.erdc.usace.army.mil [accessed February 2005].

Wilkerson, G.G., L.J. Wiles and A.C. Bennett, 2002. Weed Management Decision Models: Pitfalls, Perceptions, and Possibilities of the Economic Threshold Approach. *Weed Science* 50(4): 411-424.

Table 1: Optimal YST strategies selected by SDP model

	Scenarios defined by productivity and discount rate					
Productivity (AUM ac^{-1} yr^{-1})	2.0	2.0	5.0	5.0	10.0	10.0
Discount rate (%)	0	5	0	5	0	5
YST states						
• Minimal	CH	CH	CH	CH	BP+F	BP+F
• Moderate	CH	CH	BP+F	BP+F	BP+F	BP+F
• High	CH	CH	BP+R	BP+R	BP+R	BP+R
• Very high	BP+F	BP+F	BP+F	BP+F	BP+F	BP+F

[a] Strategies are: CH = one-time chemical control; BP+F = best YST management practice followed by spot spray of chemicals; BP+R = best YST management practice plus revegetation.

Table 2: Control strategies resulting from learning models[a]

Model	Land productivity (AUM/ac/year)	Mean Strategy Choice Proportions (n=30)				
		NC	CH	BP	BP+F	BP+R
EWA-enhanced	2.0	0.696	0.283	0.017	0.003	0.001
		(0.062)	(0.064)	(0.016)	(0.005)	(0.004)
EWA	2.0	0.565	0.384	0.044	0.005	0.002
		(0.058)	(0.055)	(0.019)	(0.008)	(0.005)
EWA-enhanced	5.0	0.691	0.283	0.021	0.003	0.002
		(0.082)	(0.080)	(0.019)	(0.005)	(0.005)
EWA	5.0	0.438	0.484	0.059	0.014	0.004
		(0.078)	(0.075)	(0.029)	(0.013)	(0.007)
EWA-enhanced	10.0	0.724	0.248	0.021	0.003	0.003
		(0.127)	(0.127)	(0.017)	(0.006)	(0.006)
EWA	10.0	0.170	0.660	0.094	0.073	0.003
		(0.077)	(0.117)	(0.066)	(0.051)	(0.006)

[a] Standard deviations given in parentheses.

Table 3: Comparing Model Performance

Model	Max AUMs	Discount rate (%)	Years	Mean NPV ($)	Standard Deviation
Enhanced EWA	2	5	75	309.49	70.96
	2	0	75	1,044.40	155.66
	5	5	75	866.69	105.29
	5	0	75	2,763.90	659.63
	10	5	75	1,672.60	291.52
	10	0	75	5,758.50	764.69
EWA	2	5	75	572.48	106.24
	2	0	75	1,940.50	284.07
	5	5	75	1,445.50	285.04
	5	0	75	5,327.70	772.13
	10	5	75	3,408.50	367.48
	10	0	75	11,968.00	1,373.30
SDP	2	5	75	605.70	53.66
	2	0	75	2,243.22	58.78
	5	5	75	1,758.85	101.64
	5	0	75	6,583.38	108.74
	10	5	75	3,834.98	174.21
	10	0	75	14,350.50	189.98