AgEcon SEARCH
RESEARCH IN AGRICULTURAL & APPLIED ECONOMICS

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

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

Help ensure our sustainability.

Give to AgEcon Search

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

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

A Bioeconomic Analysis of the Duration of Conservation Contracts.

Bronwyn Crowe & Ben White.

Paper presented at the Australian Agricultural and Resource Economists Society 51st Annual Conference 2007

Title: A Bioeconomic Analysis of the Duration of Conservation Contracts.

Authors: Bronwyn Crowe & Ben White.

Abstract

Conservation and restoration of native vegetation is often a gradual process which may require many years to transform an ecosystem from one vegetative state to a target ecosystem. This process is stochastic, with some changes potentially irreversible. In contrast, contracts with landholders to undertake conservation measures on their property are typically for less than ten years and often make no contingencies for re-contracting at the end of the contract period. The risk to land holders and conservation agencies of contracts not being renewed and the consequent potential loss of previous investment means including covenants in conservation contracts may be attractive to both parties. A model is developed to empirically examine the optimal dynamic conservation contract and the possible role of covenants in the costs and benefits of contracts.

Key words: POMDP, biodiversity, contracts, monitoring

Contact details for authors
School of Agricultural & Resource Economics
Faculty of Natural and Agricultural Sciences
University of Western Australia, Stirling Hwy, Crawley, WA 6009
Email: croweb@are.uwa.edu.au
Phone: 08 6488 4653
Fax:08 6488 1098

1. Introduction

A large number of schemes exist around the world to conserve or establish target vegetation communities. Examples include the USA Conservation Reserve Program and the Wetlands Reserve Program (Hanrahan and Zinn 2005), as well as the UK Countryside Stewardship Scheme and Environmental Stewardship (NE 2006). Typically schemes contract with farmers to take actions which increase the probability of establishing or conserving the target vegetation community. A regulator who is paying farmers compensation for retiring land and/or undertaking revegetation actions is expected to measure the output of the scheme by monitoring the vegetation succession. To do this they must decide on monitoring frequency and monitoring accuracy. For instance monitoring may be by satellite images such as Landsat, aerial photographs, or a ground survey. Reviews of agri-environmental policy monitoring in the UK and elsewhere (Hooper 1992, National Audit Office 1997, World Bank 1998) conclude that monitoring to assess ecosystem change incurs significant costs and is prone to inaccuracy in the form of mis-classifications of vegetation types.

The aim of this paper is to assess the problem faced by the regulator in terms of the optimal frequency of monitoring , the best monitoring technique and, on the basis of the results of monitoring, should a contract be initiated, stopped or continued. The unit of analysis is an area of land which either had or has the potential to establish the target vegetation community. This analysis draws upon the ecology literature on how vegetation successions are modelled, the economic analysis of monitoring and irreversible environmental change and the operations research analysis of dynamic monitoring and control problems. Each of these strands is discussed.

In ecology, Markov chains have been proposed as a representation of vegetation successions (Usher 1979). Anderson and Goodman (1957) developed methods for estimating transition probabilities from observations of the states of a system through time. Their methods have been widely developed and applied (Cox 1972, Kalbfleisch and Lawless 1985). This has led to applications in vegetation succession, see for instance, Rushton et al. (1996), Balzar (2000) and Logofet and Lesnaya, (2000). Monitoring and decision making in ecology and natural resource management is the focus of the adaptive control literature reviewed by Walters and Holling (1990) and White (2000). However, this literature focuses on the control of state variables, such as biomass which can be represented by continuous variables, for instance Williams' (1996) model of wildfowl harvesting. Williams acknowledges the problems of solving general stochastic control problems with parameter updating and monitoring. The adaptive control literature, in ecology, does not explicitly deal with the case where the states are categorical rather than continuous variables.

Operations research has approached monitoring as part of a general stochastic control literature, see for instance, Bertsekas and Shreve (1978). The partially observed Markov decision process (POMDP) model (Monahan 1982, Smallwood and Sondik 1973) is a tractable approach to stochastic control when the states follow a Markov chain, but the decision maker is unable to observe the current state of the system. To date POMDP has had relatively few applications in environmental and natural resource economics, although the paper on salmon fishing by Lane (1989) is a notable exception.

The economics of environmental monitoring has its origins with Becker's (1968) model of crime and punishment which predicts that the decision to offend depends upon a

comparison of the expected benefits with the expected costs (Heyes 2000). Becker's model predicts that regulators will fix fines as high as possible, monitor infrequently to reduces costs (if fines are punitive) and always prosecute transgressors. Harrington (1988) uses a dynamic model to explain why the stylized facts of observed environmental monitoring practice are at odds with Becker's model. These facts are first that firms are rarely fined and fines are small, second monitoring frequency is low and third most firms comply most of the time. Harrington's model, which is based on a model of tax regulation (Greenberg 1984), analyses a dynamic game between the regulator and the firm where the regulator places firms in different groups in response to observed behaviour. These groups determine the frequency of monitoring with a non-compliant firm, if detected, being moved to a group where they are subject to more frequent monitoring. Heyes and Rickman (1999) attribute Harrington's stylised facts to the need for regulators to control a range of pollutants, thus non-compliance for one pollutant is offset by compliance in another pollutant or at another plant. Malik (1990) considers the problem of monitoring a stochastic pollutant where firms can reduce their probability of detection by costly activities. The result from this model is that it is not always optimal to have high fines for non-compliance as this increases the amount spent on activities which avoid detection.

The monitoring problem described here differs from most previous contributions to the literature in a number of fundamental respects. First the variable monitored is a categorical variable classifying the state of the vegetation community into a finite number of classes. Most previous economic studies describe monitoring an emission variable where standards are in terms of quantities or concentrations. The monitoring problem here is dynamic and extends from 2 periods up to an infinite time horizon. Given this added complexity the strategic interaction between the firm and the regulator is not modelled explicitly, instead it is characterised as 'nature' which determines if a conservation scheme succeeds or fails.

The decision to conserve is similar to the decision to delay development in the Arrow-Fisher-Henry quasi-option value model (Arrow and Fisher 1974, Fisher 2000, Henry 1974) and the optimal stopping model of Batabayal (1998). The Arrow-Fisher-Henry model calculates the value (quasi-option value) of delaying a development decision on the basis that more precise information will become available on the value of conservation. POMDP may be viewed as a generalisation of this model in that it allows for imperfect information and costly monitoring instead of information arriving 'passively' through time. The aim of this paper is to explore, through a simplified case study, how POMDP might be applied to ecological monitoring. The remainder of the paper is organised as follows. The next section introduces the POMDP model. Section 3 describes the case study for the conservation and restoration of Salmon Gum woodland in the Western Australian Wheatbelt. Section 4 Concludes.

2. Methods

A regulator wishes to maximise the private and public value of a piece of land where vegetation communities are described by *N* discrete states $s_i = 1, ..., N$. The vegetation community changes through time according to a Markov process and the (*NxN*) matrix of transition probabilities are a function of the level of conservation effort, for instance for three vegetation states we have:

$$P(e_t) = \begin{bmatrix} p_{11}(e_t) & p_{12}(e_t) & p_{13}(e_t) \\ p_{21}(e_t) & p_{22}(e_t) & p_{23}(e_t) \\ p_{31}(e_t) & p_{32}(e_t) & p_{33}(e_t) \end{bmatrix} \tag{1}$$

The elements $p_{ij}(e_t)$ give the probability of the land in state *i* being in state *j* after a single period *t*. Conservation effort, e_t, is a measure of resources allocated to conservation, in the example it is based on the work of Yates and Hobbs (1997) and Gibbons and Freudenberger (2006). The regulator offers a contract that stipulates conservation effort e_t. Conservation effort increases or decreases the probability of a transition to the target vegetation community.

The regulator has a prior probability of the current vegetation community given by the (*1xN*) vector π known in the POMDP literature as the *belief state*. For many ecosystems this is a realistic assumption: vegetation classifications are uncertain or the vegetation may be a mosaic of different vegetation classes. Often the high cost of a definitive vegetation survey means that conservation schemes are initialised with incomplete knowledge of the initial vegetation community across the whole area. The observation matrix, which is a function of monitoring effort u_t determines the accuracy of monitoring. For three states the *(NxN)* observation matrix is given by:

$$\Theta(u_t) = \begin{bmatrix} r_{11}(u_t) & r_{12}(u_t) & r_{13}(u_t) \\ r_{21}(u_t) & r_{22}(u_t) & r_{23}(u_t) \\ r_{31}(u_t) & r_{32}(u_t) & r_{33}(u_t) \end{bmatrix} \tag{2}$$

where the element $r_{j\theta}(u_t)$ is the probability that if state θ is observed the vegetation at the end of period *t* is *j*. If $\Theta(u_t)$ is an identity matrix then monitoring is perfectly accurate, if it is uniform it is uninformative. Increased monitoring effort raises the probability of a correct observation.

Monitoring reduces the uncertainty about which state the land is in and updates the prior probability to a posterior probability by Bayes rule:

$$\pi_{jt} = \frac{\sum_i \pi_{it-1} p_{ij}(e_t) r_{j\theta}(u_t)}{\sum_{i,j} \pi_{it-1} p_{ij}(e_t) r_{j\theta}(u_t)} \tag{3}$$

The new belief state is a $1 \times N$ vector of probabilities. In vector form, (3) can be rewritten as:

$$\pi_t = T(\pi_{t-1} \mid e_t, u_t, \theta) \tag{4}$$

where *T(.)* is the belief transformation function. The belief state captures the history of all past observations and actions.

Monitoring costs

Heyes (2002) draws a useful distinction between inspecting an environmental variable which generates a noisy signal and an environmental audit which is definitive. Methods for monitoring vegetation community change range from low cost remote sensing methods such as aerial photographs and satellite images, to relatively high cost field surveys (World Bank 1998). We assume that from past ‘ground truthing’, these methods have established observation matrices. For instance remote sensing methods are known

for relatively high probabilities of misclassification (Hooper 1992), while intensive field surveys are accurate.

We assume that the cost of monitoring depends on the observation matrix thus the quasi-convex monitoring cost function $c^{v}(u_t)$ is at a maximum when $\Theta(u)$ is an identity matrix, that is the state is observed with perfect accuracy, and $c^{v}(u_t)=0$ when $u_t=0$ and $\Theta(u)$ is a uniform matrix with all elements equal to 1/N.

The Regulator's problem

The regulator maximises the expected present-value of the welfare function in relation to conserving an area of land by choosing conservation effort and monitoring effort. The regulator's problem can be represented by the following POMDP problem represented in a mathematical programming problem

$$V[\pi_t] = \underset{wrt\ e_t,u_t}{Maximise} \sum_t \sum_i \pi_{it}[g_i(e_t) - c_i(e_t) - c^{v}(u_t)]\delta^t \qquad (5a)$$

Subject to

$$\pi_t = T(\pi_{t-1} \mid e_t, u_t, \theta) \qquad (5b)$$

$$\pi_0 = \tilde{\pi} \qquad (5c)$$

The first term $g_i(e_t)$ in (5a) gives the non-market net benefits of vegetation community i, it is given as a function of e_t as conservation effort may enhance the benefits of a particular state. The term $c_i(e_t)$ gives the resource cost to the farmer of conservation effort in state e_t. The farmer may be compensated by the regulator for these costs, but as these are transfer payments they do not appear in the objective function. Monitoring costs depend upon the monitoring effort and are given by $c^{v}(u_t)$. The term δ^t is the discount factor which converts net benefits generated at time *t* to their present-value at *t=0*, *g* is the discount rate. Equation (5b) gives the updating equation for the belief state (4) and (5c) gives the belief state (prior probabilities of states) at the start of the planning horizon when t=0 as $\tilde{\pi}$. To simplify the notation in later sections we define net-benefit as

$$w_i(e_t, u_t) = g_i(e_t) - c_i(e_t) - c^{v}(u_t) \qquad (6)$$

Dynamic optimisation

Unlike a Markov Decision Problem (MDP) which has a standard dynamic programming solution (Puterman, 1994), the solution to a POMDP problem is more difficult because the probability of the system being in a particular state depends upon past monitoring and the resulting observations. The original solution by Smallwood and Sondik (1973) introduces the notion of a *belief state* where the conventional states of MDP, namely s_i, are replaced by a *belief state* π_t which is the vector of probabilities of being in the states. The solution entails finding a set of actions which are optimal across the belief state (Cassandra 1995). In a simplified form the optimisation problem is to solve the following version of Bellman's equation:

$$V_t[\pi_t] = \underset{e_t,u_t}{\text{maximise}} \sum_{\mathrm{i}} \pi_{it}\{w_i(e_t,u_t) + \sum_j \sum_\theta p_{ij}(e_t) r_{j\theta}(u_t) V_{t+1}[T(\pi_t | e_t, u_t, \theta)]\}. \qquad (7)$$

where $V_t(\pi_t)$ is the optimal value from optimizing across the time horizon from t to T starting in belief state π_t. The optimal value comprises two components, the first term is the expected immediate reward and the second term is the expected reward for the remaining periods, the term $p_{ij}(e_t)r_{j\theta}(u_t)$ gives the joint probability of observing state θ when the previous state is *i* and the current state *j*. Equation (7) is similar in construction to a standard stochastic dynamic programming model the only difference is in the presence of the belief state. For instance if the initial state was known with certainty and there was no monitoring, optimization would proceed by maximizing the current net-benefit whilst accounting for the effect the action has on the expected value across the remaining periods. This principle of optimality still holds in POMDP except it has to solve the problem for all possible belief states. This involves defining the optimal solution as a set of action vectors which are optimal in some belief state. This is illustrated and discussed in greater detail in the contest of the case study.
Solving the dynamic optimization presented in (7) is not trivial due to the problems of determining $V_t[\pi_t]$. However, if we restrict e_t and u_t to a discrete set of values we can make use of the result that $V_t[\pi_t]$ is always piecewise linear and convex (Smallwood and Sondik 1973), thus a modified dynamic programming algorithm can determine $V_t[\pi_t]$ as a set of vectors generated from different actions. This allows us to rewrite (7) as:

$$V_t[\pi_t] = \underset{e_t,u_t}{\text{maximise}} \sum_{\mathrm{i}} \pi_{it}\{w_i(e_t,u_t) + \sum_j \sum_\theta p_{ij}(e_t)r_{j\theta}(u_t)\alpha_j^{\iota(\pi_t,e_t,u_t,\theta)}(t+1)\} \qquad (8)$$

where $\alpha_j^k(t)$ is a (1xN) policy vector which gives the expected payoff from an action across all the states. The superscript on the policy vector gives the optimal vector for a particular belief state and is formally defined as follows:

$$\iota(\pi_t,e_t,u_t,\theta) = \arg\max_k \left[\sum_i \sum_j \pi_{it} p_{ij}(e_t)r_{j\theta}(u_t)\alpha_j^k(t+1)\right] \qquad (9)$$

that is it selects the vector, by the superscript *k*, which gives the highest expected value for the belief state resulting from the prior probability, action and observation.

Case study

Background

The western Australian wheatbelt, and particularly the Northeastern Wheatbelt Regional Organisation of Councils (NEWROC), has received attention recently due to its agricultural and environmental importance. The area is of agricultural significance as well as biodiversity and under threat from salinity and large scale clearing. The NEWROC comprises the shires of Koorda, Mount Marshall, Mukinbudin, Nungarin, Trayning, Westonia and Wyalkatchem. The area was 75% zoned for clearing, with 12% of the cleared area remnant vegetation in 2002. Within each shire the area of remnant vegetation in cleared areas ranged from 5% in the south west shirt of Wyalkatchem to 21% in the eastern most shire of Westonia.
Yates and Hobbs (1997) detail the state of *Eucalyptus* woodlands in southeast and southwest Australia. Woodlands have been extensively cleared and are often badly degraded due to livestock grazing. Currently it is estimated only 10% of *Eucalyptus*

loxophleba (York gum) and 20% of *Eucalyptus salmonophloia/Eucalyptus salubris* (salmon gum/gimlet) woodlands remain. A similar situation exists on the east coast of Australia, where 0.01% of *eucalyptus albens* (white box) woodland remains relatively unmodified.

The removal of degrading factors such as grazing and weeds may be insufficient to restore the woodland, with revegetation action required. Yates and Hobbs (1997) go on to identify the stable woodland states that exist in *Eucalyptus salmonophloia* woodlands currently and the transitions required to shift the woodland areas from one state to anotherFigure 1. Remnant vegetation in the NEWROC area is highly fragmented due to agricultural clearing, and degraded due to weeds, livestock grazing and firewood collection. Together with the impact of dryland salinity this means high levels of habitat loss, with the remaining vegetation severely degraded. The works required and probability of their success is largely determined by the current state of the woodland and its ability to shift to another state. The fencing of remnant vegetation to remove livestock and feral grazing may be insufficient to return a degraded woodland to an undegraded state. Extensive revegetation and weed control would likely be required to achieve this shift.

This case study uses POMDP methodology to investigate the optimal conservation action for growers, regulators and wider society in the NEWROC area in light of the stable woodland states and transitions approach of Yates and Hobbs (1997) and to the Auctions for Landscape Recovery project findings XX.

Data

The states and transitions for salmon gum woodland as defined by Yates and Hobbs (1997) are simplified to the diagram given in Figure 1 for this case study. Agriculture (Agric) refers to a stable state of annual rotations of crop or livestock production on the land. Degraded Woodland (Degwood) is a woodland with few perennial understory species, a ground layer of annual weeds and compacted soil. Undegraded Woodland (Undegwood) has an intact understory of shrubs, a layer of plant litter across the ground and good soil. The transition from Undegwood to Degwood and Degwood to Agric through the action Sheep consists of periodic grazing by livestock or wild animals such as rabbits. Revegetation is the fencing of the area to exclude livestock and other animals, and the planting of woodland species. No Action refers to abandonment or not using the land for any specific purpose. Changing the action has two effects, first it changes the net benefit $w_i(e_t, u_t)$ and second it changes the transition probabilities. The probable impact of these actions on a specific state is given in Table 1.

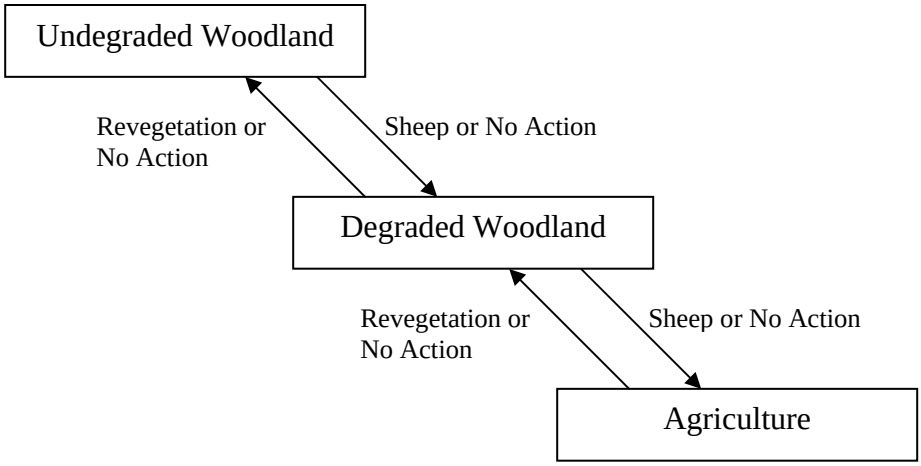

Figure 1 Transition and stable states of salmon gum woodland (based on Yates and Hobbs, 1997)

The action, state of the land and monitoring determine the net benefit to the grower, regulator and wider society of the land for each period of the analysis. The benefit of sheep grazing on any given land type (Sheep) is assumed to be $5 per hectare per year (DAFWA 2005). The cost of revegetation (Reveg) is $935 per hectare (Gole *et al.* 2005). While not undertaking any activity (No Action) does not incur a cost or provide a benefit. Land being in the state of Undegwood provides a benefit to the grower, regulator and wider society or non-market value. The community willingness to pay for teatree woodland in Queensland is used as an estimate of the benefit to the grower, regulator and wider society of salmon gum woodland, $18 per hectare (Mallawaarachchi *et al.* 2001). Monitoring the land to determine its current state requires engaging a local expert and is initially estimated to cost (c^m) $8 per hectare (Gole *et al.* 2005).
The grower or regulator is able to engage an expert to monitor the land and estimate its current state to inform their future decisions. The probability this monitoring correctly estimates the current state of the land for each action is given in Table 2. Increasing monitoring effort makes the observation matrix 'more informative' no monitoring gives a uniform observation matrix. Combinations of conservation effort and monitoring effort give six different actions in all.

Table 1 Probability of transition between states given selected action

	Sheep		End state	
		Undegwood	Degwood	Agric
	Undegwood	0.8	0.2	0
Start state	Degwood	0	0.8	0.2
	Agric	0	0	1

	Reveg		End state	
		Undegwood	Degwood	Agric
	Undegwood	1	0	0
Start state	Degwood	0.8	0.2	0
	Agric	0.5	0.5	0

	No Action	End state		
		Undegwood	Degwood	Agric
	Undegwood	0.9	1	0
Start state	Degwood	0	0.9	0.1
	Agric	0	0	1

Table 2 Observation probabilities when monitoring occurs

	Sheep	Start State		
		Undegwood	Degwood	Agric
	Undegwood	0.85	0.15	0
Observation	Degwood	0	0.85	0.15
	Agric	0	0	1

	Reveg	Start State		
		Undegwood	Degwood	Agric
	Undegwood	1	0	0
Observation	Degwood	0.1	0.8	0.1
	Agric	0	0.2	0.8

	No Action	Start State		
		Undegwood	Degwood	Agric
	Undegwood	0.9	1	0
Observation	Degwood	0.05	0.9	0.05
	Agric	0	0.1	0.9

Results

The POMDP analysis was run over a specified number of periods to compare the optimal decision for grower, regulator and wider society in the short term and the long term. A discount factor of $\delta = 0.95$ is assumed for all analysis. The estimated cost of monitoring agricultural land in the NEWROC area was $8 per ha (Gole *et al.* 2005). However, as this monitoring cost proved too high for growers or regulators to undertake monitoring, a reduced cost $2 per ha was also analysed

2 Period Example

If the scheme based on the above runs for just 2 periods the only expected action for both periods is Sheep. The policy graph in Figure 2 shows how after the first period of taking the action Sheep & No Monitoring is complete, Sheep & No Monitoring is repeated as the next action in the sequence. Sheep & No Monitoring for 2 periods gives a net benefit of $38.26 were the land Undegwood, and $9.75 if it were Degwood or Agric. The solution to the 2 period problem follows from two dynamic programming iterations using the Smallwood and Sondik algorithm (1973) as developed by Cassandra (1995). The expected payoff is calculated from the term in brackets on the right hand side of equation (7). For instance the net benefit of sheep when the initial state of the land is Undegwood is calculated as (0.8*23+0.2*14+0*14) + [0.8*(0.8*23+0.2*14+0*14) + 0.2*(0*14+0.8*5+0.2*5) + 0*(0*14+0*5+1*5)]*0.95.

Altering the cost of monitoring does not alter the action decision of the grower/regulator, undertaking Sheep & No Monitoring in both period (Figure 2). The net benefit of this action is the same as with a monitoring cost of $8 per ha, $39.16 were the land Undegwood, and $10.00 if it were Degwood or Agric.

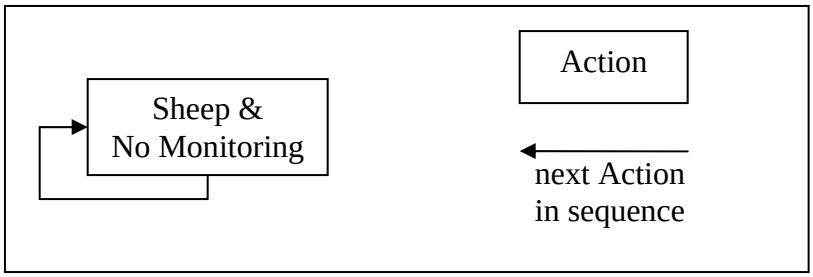

Figure 2 Policy graph when monitoring costs $8 per ha or $2 per ha, and the time horizon is 2 periods.

5 Period Example

Extending the number of time periods to 5 periods alters the actions of the grower/regulator under certain prior probabilities of the land state. When the probability of land being Undegwood is higher the action No Action & No Monitoring becomes optimal. Figure 3 shows how with a monitoring cost of $8 per ha and 5 period horizon the grower or regulator would undertake either (1) act as in the 2 period example, beginning with Sheep & No Monitoring or (2) begin with a period of No Action & No Monitoring prior to continuous Sheep & No Monitoring. The net benefit of 5 periods of Sheep & No Monitoring is $73.01 for Undegwood and $22.62 for Degwood or Agric. An initial period of No Action and No Monitoring increases the net benefit when the land is Undegwood to $73.18 but decreases it if the land is Degwood or Agric to $17.62.
A lower monitoring cost of $2 per ha does not alter the grower/regulator's actions in a 5 period to those with an $8 per ha monitoring cost. The net benefit of the actions is the same as with an $8 per ha monitoring cost also.

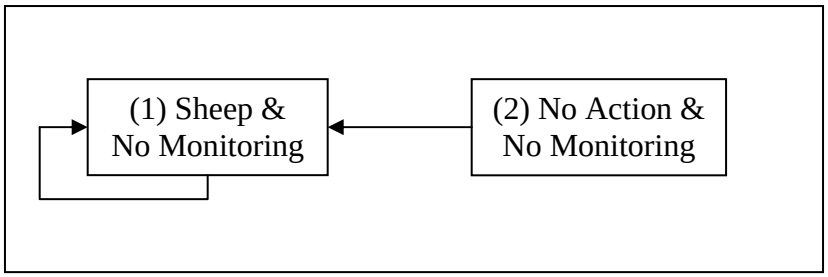

Figure 3 Policy graph when monitoring costs $8 per ha and the time horizon is 5 periods, and when monitoring is $2 per ha and the time horizon is 5 periods.

10 Period Example

The action decision does not alter when the monitoring cost is $8 per ha and the analysis is for 10 periods, Figure 1. The net benefit of continuous Sheep and No Monitoring Action 1 in Figure 4, for 10 periods is $103.29 for Undegwood and $40.13 for Degwood

or Agric. An initial period of No Action & No Monitoring prior to continuous Sheep & No Monitoring (Action 2), has a net benefit of $105.06 for Undegwood and $35.13 for Degwood and Agric. Two periods of No Action & No Monitoring prior to continuous Sheep & No Monitoring (Action 3), has a net benefit of $105.95 for Undegwood and $30.38 for Degwood and Agric. While three periods of No Action & No Monitoring prior to continuous Sheep & No Monitoring (Action 4), has a net benefit of $106.10 for Undegwood and $25.86 for Degwood and Agric. The number of initial periods of No Action & No Monitoring is higher the greater the probability of the land being Undegwood.

With a monitoring cost of $2 per ha the grower/regulator may undertake Sheep or No Action, with Monitoring or No Monitoring, depending on the probability of the initial state of the land being Undegwood, Degwood or Agric, Figure 5. The five actions optimal at the various combinations of the probability of land being Undegwood, Degwood and Agric are; (1)continuous Sheep & No Monitoring, (2) Sheep & Monitoring followed by continuous Sheep & No Monitoring if Degwood or Agric is observed or No Action & No Monitoring if Undegwood is observed, (3) No Action & Monitoring followed by continuous Sheep & No Monitoring if Degwood or Agric is observed or No Action & No Monitoring if Undegwood is observed, (4) and (5) continuous No Action & No Monitoring. The net benefit of Actions 1 to 5 is given in Table 1. Monitoring is undertaken when uncertainty is high, when the current state is believed to be Agric Sheep is the action, while when it is believed to be Undegwood No Action is taken.

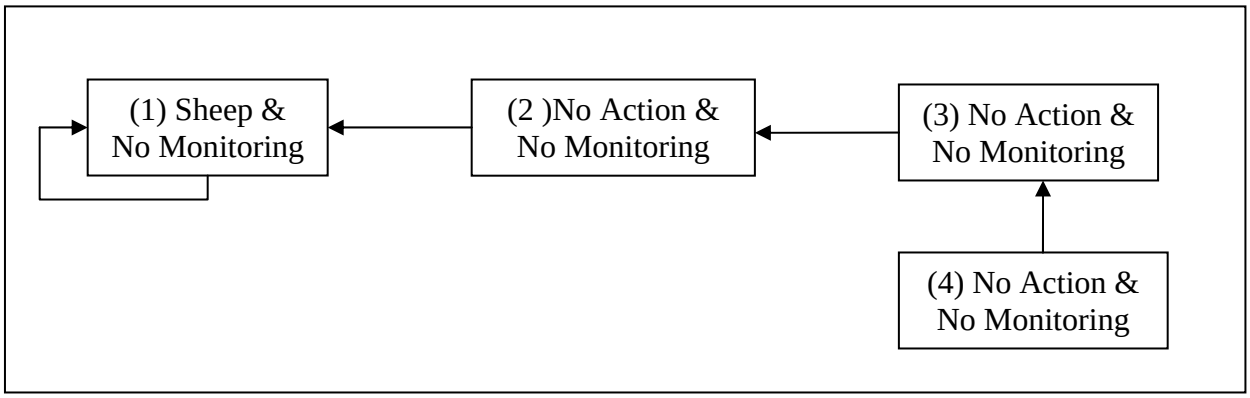

Figure 4 Policy graph when monitoring costs $8 per ha and the time horizon is 5 periods.

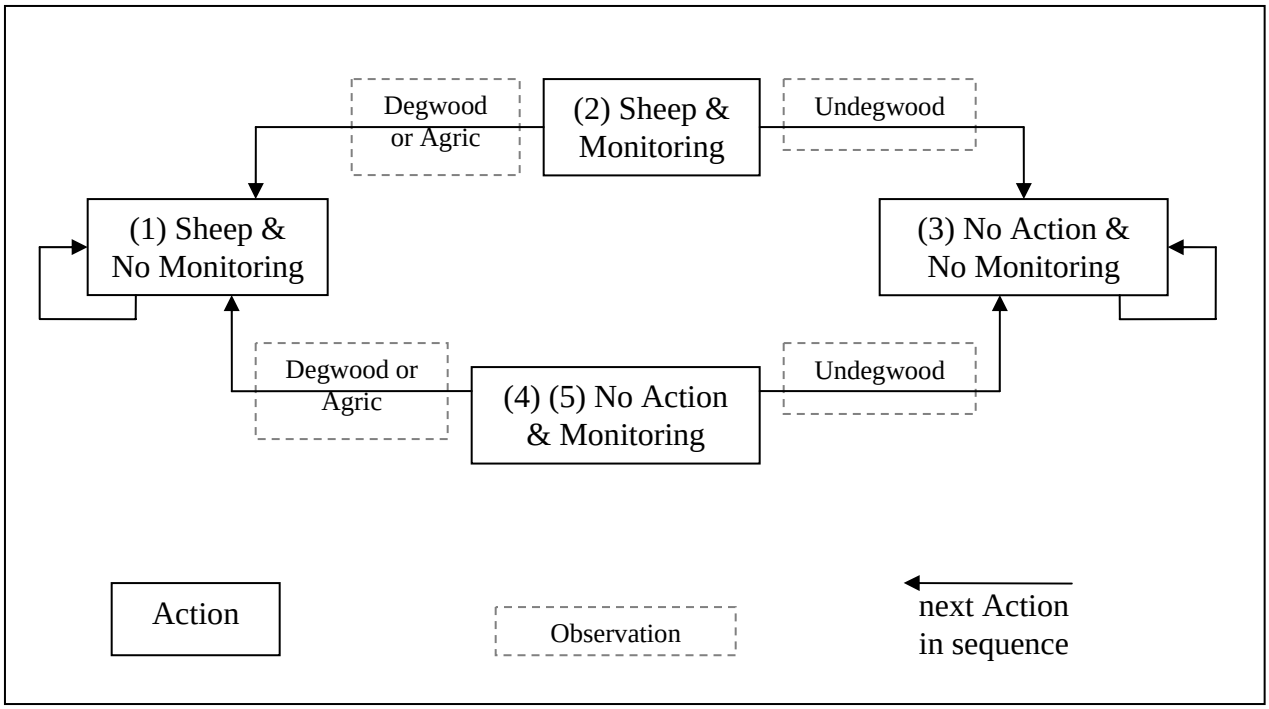

Figure 5 Policy graph when monitoring costs $2 per ha and the time horizon is 10 periods.

Table 3 Net benefit of Actions 1 to 5 depending on the initial state of the land, when monitoring is $2/ha and the horizon is 10 periods.

	Initial State		
	Undegwood	**Degwood**	**Agric**
Action 1	$122.30	$50.00	$50.00
Action 2	$123.55	$48.00	$48.00
Action 3	$128.01	$42.33	$43.00
Action 4	$128.01	$35.00	$35.00
Action 5	$128.02	$30.00	$30.00

20 Period Example

With a monitoring cost of $8 per ha the only change to the action choice of the grower/regulator when analysed over 20 periods, from that of the 5 or 10 period example, is an increase in the number of periods of No Action & No Monitoring prior to continuous Sheep & No Monitroing at high probabilities of Undegwood, Figure 6. The net benefit of continuous Sheep & No Monitoring (Action 1) is $131.37 if Undegwood and $64.15 if Degwood or Agric, refer to Table 4. One initial period of No Action & No Monitoring (Action 2) increases the net benefit when Undegwood to $133.65 and decreases to $59.15 if Degwood or Agric. The net benefit of further increasing the number of periods of No Action and No Monitoring prior to continuous Sheep & No Monitoring for Undegwood/Degwood and Agric (Actions 3, 4 and 5) are shown in Table 4.

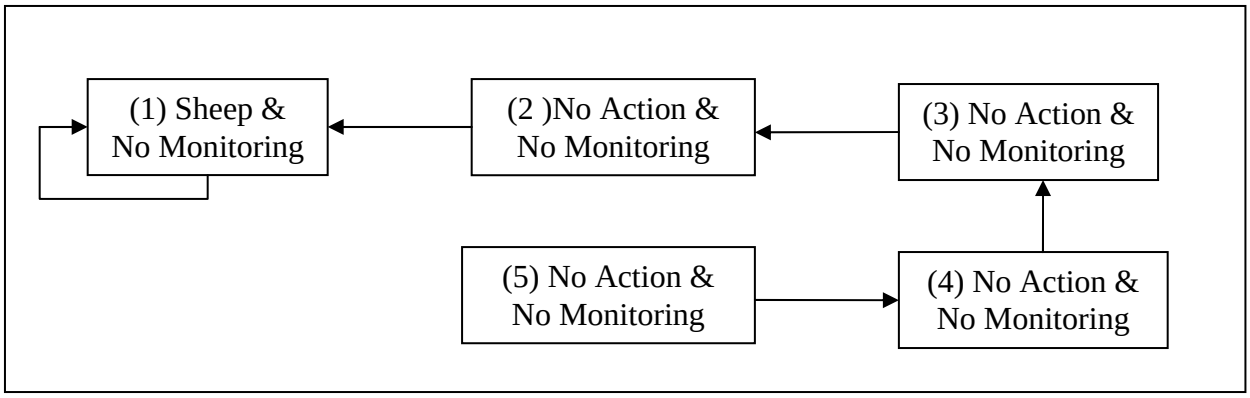

Figure 6 Policy graph when monitoring costs $8 per ha and the time horizon is 20 periods.

Table 4 Net benefit of Actions 1 to 5 depending on the initial state of the land when the monitoring cost is $8 per ha and the time horizon is 20 periods.

	Initial State		
	Undegwood	**Degwood**	**Agric**
Action 1	$131.37	$64.15	$64.15
Action 2	$133.65	$59.15	$59.15
Action 3	$135.11	$54.40	$54.40
Action 4	$135.90	$49.89	$49.89
Action 5	$136.14	$45.60	$45.60

Figure 7 shows the action sequence for the 20 period example when the monitoring cost is $2 per ha. There are nine initial actions which match to the various probabilities of the land being initially Undegwood, Degwood or Agric; (1) continuous Sheep & No Monitoring, (2) Sheep & Monitoring followed by continuous Sheep & No Monitoring if Degwood or Agric is observed or Action 7 if Undegwood is observed, (3) Sheep & Monitoring followed by continuous Sheep & No Monitoring if Agric is observed, Action 2 if Degwood or Action 8 if Undegwood, (4) No Action & Monitoring followed by continuous Sheep & No Monitoring if Degwood or Agric is observed or Action 7 if Undegwood is observed, (5) No Action & Monitoring followed by continuous Sheep & No Monitoring if Agric is observed, Action 2 if Degwood or Action 8 if Undegwood, (6) No Action & No Monitoring followed by Action 4, (7) No Action & No Monitoring followed by Action 6, (8) No Action & No Monitoring followed by Action 5, (9) No Action & No Monitoring followed by Action 8. The net benefit of each action with each initial land state is given in Table 5.

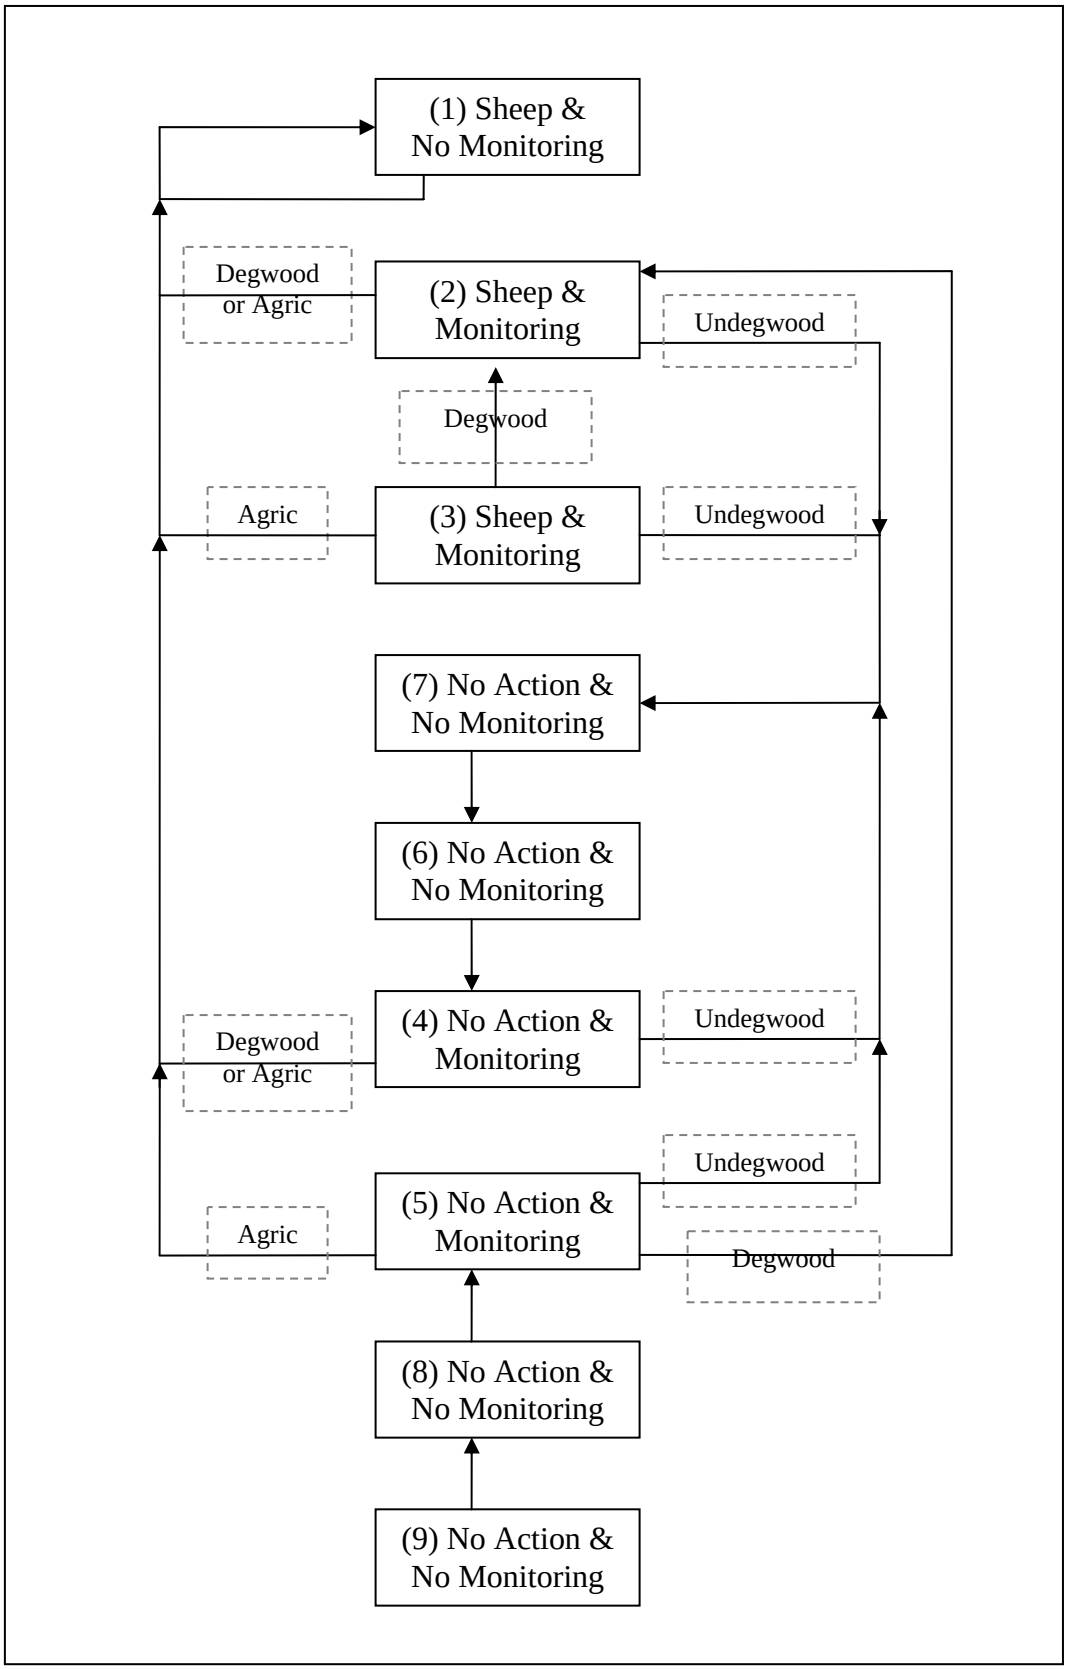

Figure 7 Policy graph when monitoring costs $2 per ha and the time horizon is 20 periods.

Table 5 Net benefit of Actions 1 to 9 depending on the initial state of the land when the monitoring cost is $2 per ha and the time horizon is 20 periods.

	Initial State		
	Undegwood	**Degwood**	**Agric**
Action 1	$180.07	$100.00	$100.00
Action 2	$188.03	$98.00	$98.00
Action 3	$188.59	$96.64	$98.00
Action 4	$193.73	$92.21	$93.00
Action 5	$194.22	$90.57	$92.80
Action 6	$194.90	$87.29	$88.00
Action 7	$195.13	$85.79	$87.80
Action 8	$195.38	$82.36	$83.00
Action 9	$195.40	$80.99	$82.80

Infinite Horizon Problem

The infinite horizon solution is generated by running the POMDP algorithm until the solution converges to a steady state solution, where $V_t[\pi_t]$ is constant and the sequence of actions is the same in each period. Convergence occurs when the discount factor is less than one, in this case $\delta = 0.95$. An infinite time horizon mimics entering into a covenant or permanent agreement.

Assuming monitoring costs $8 per ha the optimal action is Sheep & No Monitoring when the probability of Agric is high. when the probability of Undegwood is higher an initial period(s) of No Action & No Monitoring occur, to a maximum of 4 (Figure 8). The net benefit of each action for Undegwood, Degwood and Agric is given in Table 6.

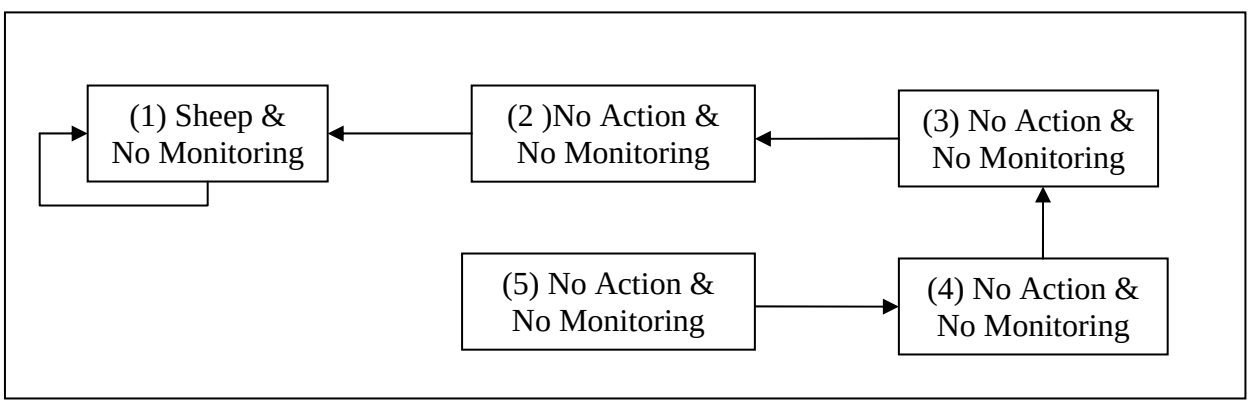

Figure 8 Policy graph when monitoring costs $8 per ha and the time horizon is infinite.

Table 6 Net benefit of Actions 1 to 5 depending on the initial state of the land when monitoring costs $8 per ha and the time horizon is infinite.

	Initial State		
	Undegwood	**Degwood**	**Agric**

Action 1	$167.50	$100.00	$100.00
Action 2	$169.81	$95.00	$95.00
Action 3	$171.31	$90.25	$90.25
Action 4	$172.15	$85.74	$85.74
Action 5	$172.43	$81.45	$81.45

When the time horizon is infinite, reducing the monitoring cost from $8 per ha to $2 per ha increases the number of actions possibly undertaken, but to the same quantity as when the analysis is for 20 period. Figure 9 shows the seven possible initial actions undertaken depending on the probability of Undegwood, Degwood and Agric; (1) continuous Sheep & No Monitoring, (2) Sheep & Monitoring followed by continuous Sheep & No Monitoring if Degwood or Agric is observed or Action 7 if Undegwood is observed, (3) Sheep & Monitoring followed by continuous Sheep & No Monitoring if Agric is observed, Action 2 if Degwood or Action 8 if Undegwood, (4) No Action & Monitoring followed by continuous Sheep & No Monitoring if Degwood or Agric is observed or Action 7 if Undegwood is observed, (5) No Action & Monitoring followed by continuous Sheep & No Monitoring if Agric is observed, Action 2 if Degwood or Action 8 if Undegwood, (6) No Action & No Monitoring followed by Action 5, (7) No Action & No Monitoring followed by Action 6. the net benefit of these action sequence is given in Table 7. The optimal action for each probability of Undegwood, Degwood and Agric is shown in Figure 10. Growers/regulators are choose do Sheep & No monitoring when the probability of Undegwood is 0 or very low. Sheep & Monitoring occurs when the possibility of the land being Undegwood increases, with grower/regulators distinguishing a separate following action for all three states when the probability of Degwood is higher. No Action & Monitoring occurs when the grower/regulator is very uncertain about the initial state of the land. No Action & No Monitoring occurs when the probability of Undegwood is very high, meaning the probability of Degwood or Agric is very low. When the probability of Undegwood approaches 1 the number of period of No Action & No Monitoring increases.

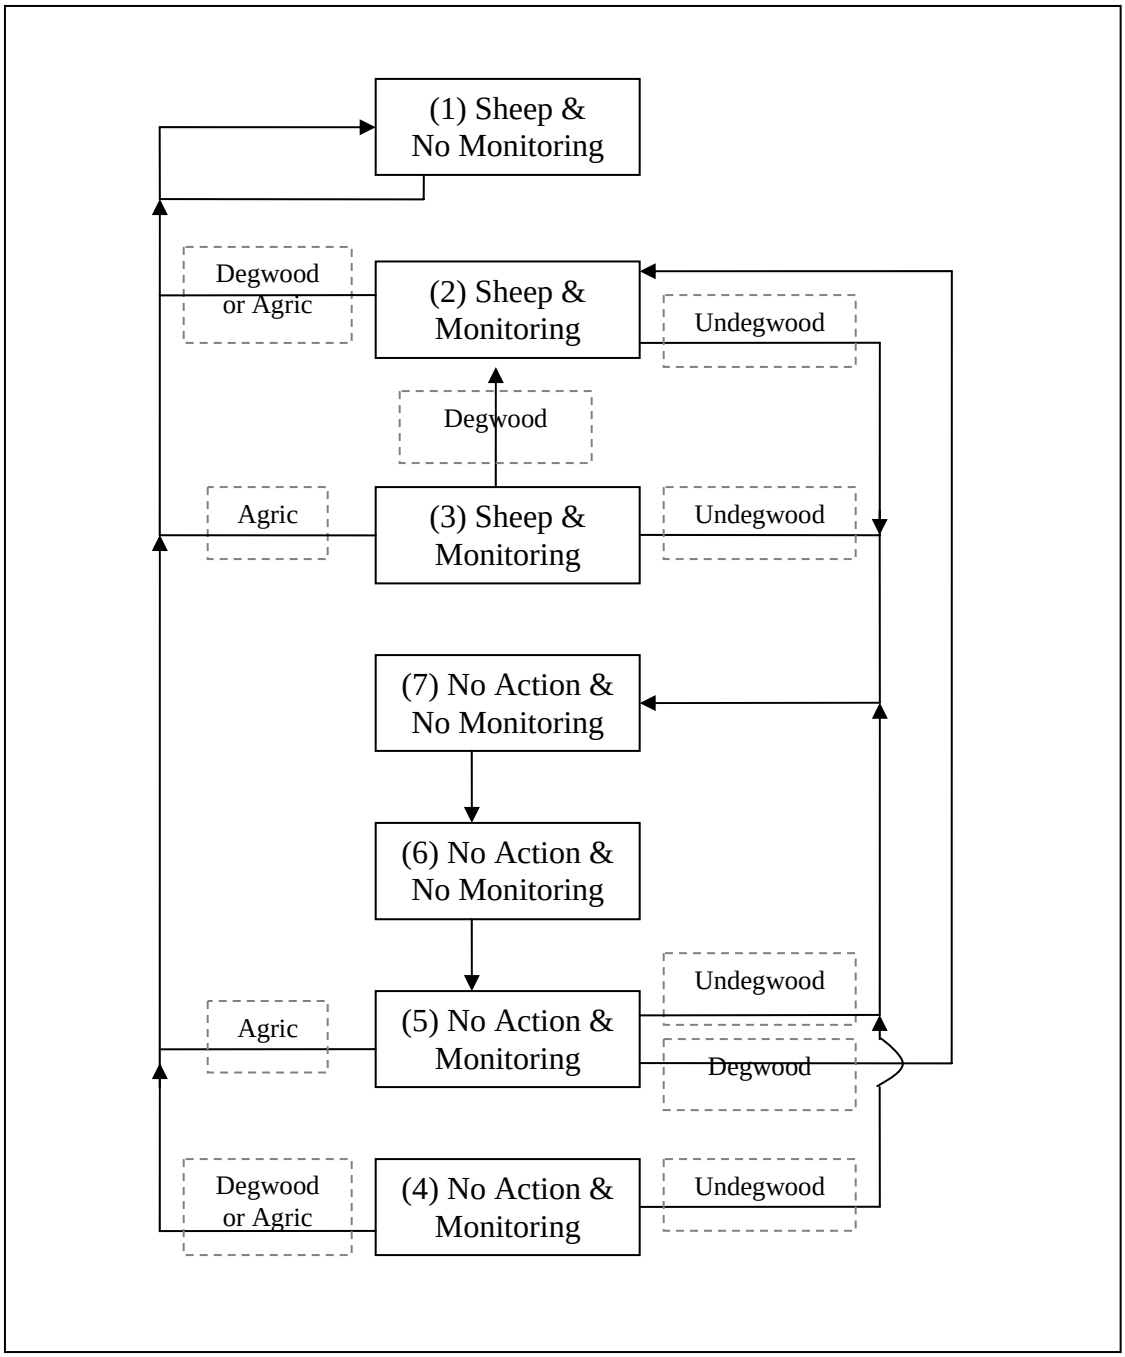

Figure 9 Policy graph when monitoring costs $2 per ha and the time horizon is infinite.

Table 7 Net benefit of initial Actions 1 to 7 depending on the initial state of the land when the monitoring cost is $2 per ha and the time horizon is infinite.

	Initial State		
	Undegwood	**Degwood**	**Agric**
Action 1	$5,081.00	$5,000.00	$5,000.00
Action 2	$5,094.18	$4,998.00	$4,998.00

Action 3	$5,095.42	$4,996.64	$4,998.00
Action 4	$5,100.98	$4,992.14	$4,993.00
Action 5	$5,101.99	$4,990.50	$4,992.80
Action 6	$5,102.94	$4,985.73	$4,987.80
Action 7	$5,103.32	$4,980.94	$4,982.80
Action 8	$5,081.00	$5,000.00	$5,000.00
Action 9	$5,094.18	$4,998.00	$4,998.00

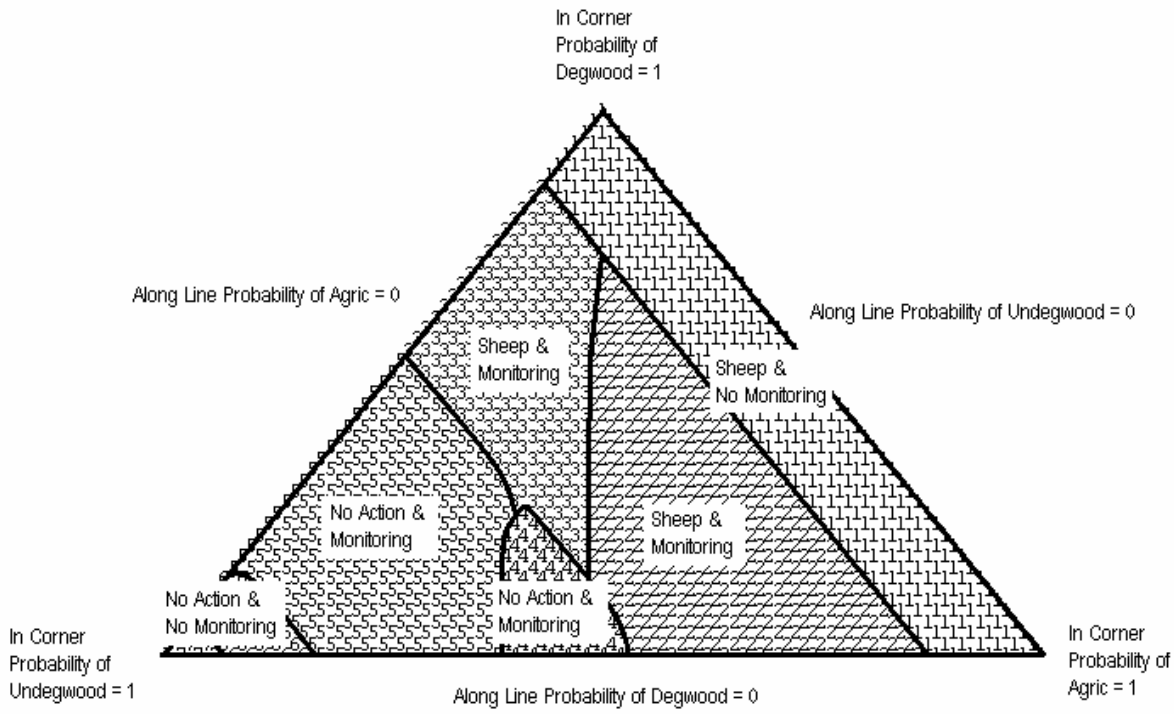

Figure 10 Decision of grower/regulator to undertake Actions 1 to 7 at given probabilities of Undegwood, Degwood and Agric.

Concluding comments

Monitoring is an essential component of schemes designed to conserve an ecosystem and such schemes are now a central part of environmental policy. A characteristic of these policies is that they offer generally short term contracts to farms to conserve or reinstate a particular vegetation community. To date the literature has had relatively few contributions on the subject of optimal dynamic monitoring. The original Arrow-Fisher-Henry quasi-option value model is concerned with information gathering, but as a passive function of time passing. The POMDP framework presents a flexible approach to determine optimal actions where the stochastic process is represented by a Markov chain. In this paper we propose a POMDP as a tractable framework for analysing optimal levels of monitoring accuracy and frequency at varying contract lengths. The results show that monitoring decision depend upon the prior probabilities of the states of the system and cost of monitoring. A general result is that monitoring tends to be optimal where there is

more uncertainty about the target state and repeated monitoring might be used before making an irreversible decision. Reducing monitoring cost increases the quality and frequency of monitoring.

The fact that the Smallwood and Sondik (1973) algorithm is reasonably robust and that Markov chains are familiar to ecologists as a method for modelling environmental change means that this approach has the potential to contribute to the analysis of monitoring systems and may lead to significant savings in monitoring costs. Currently monitoring is often undertaken as a matter of routine rather than relating monitoring to the actual predicted rates of vegetation change.

This paper has only presented a small set of results on the impact on optimal monitoring of different assumptions. One avenue for further research is to include a POMDP analysis for ecological studies of conservation schemes so that the monitoring decisions can be based on costs and the rates of vegetation change. Further research will address the main focus of the economic literature on monitoring, namely the incentives for compliance and cheating to reflect the reality that conservation schemes fail due to both non-compliance and natural variation.

Acknowledgements

This work is part of a larger Land & Water Australia Project entitled 'A Bioeconomic Analysis of the Duration of Conservation Contracts'. We are grateful for the assistance of Richard Hobbs of Murdoch University and the World Wide Fund for Nature.

This work builds on previous analysis presented in White (2000).

References

Anderson, T. W. and Goodman, L. A. (1957). Statistical Inference about Markov Chains, *The Annals of Mathematical Statistics* 28, 89-110.

Arrow, K. J. and Fisher, A. C. (1974). Environmental preservation, uncertainty, and irreversibility, *The Quarterly Journal of Economics* 88, 312-319.

Balzter, H. (2000). Markov chain models for vegetation dynamics, *Ecological Modelling* 126, 139-154.

Batabyal, A. A. (1998). An optimal stopping approach to the conservation of biodiversity, *Ecological Modelling* 105, 293-298.

Becker, G. S. (1968). Crime and punishment an economic approach, *Journal of Political Economy* 76, 169-217.

Bertsekas, D. (1978). *Stochastic Optimal Control: The Discrete Time Case*. Academic, New York.

Cassandra, A. R. (1995). *Optimal policies for partially observable markov decision processes*. Department of Computer Science, Brown University,

Cox, D. R. (1972). Regression Models and Life-Tables, *Journal of the Royal Statistical Society. Series B (Methodological)* 34, 187-220.

DAFWA. (2005). *Gross Margins Guide 2005 Western Australia*. Department of Agriculture and Food Western Australia, Government of Western Australia, Perth.

Fisher, A. C. (2000). Investment under uncertainty and option value in environmental economics, *Resource and Energy Economics* 22, 197-204.

Gibbons, P. and Freudenberger, D. (2006). An overview of methods used to assess vegetation condition at the scale of the site, *Ecological Management & Restoration* 7, S10-S17.
Gole, C., Burton, M., Williams, K. J., Clayton, H., Faith, D. P., White, B., Huggett, A. and Margules, C. (2005). *Auctions for Landscape Recovery Final Report*. World Wide Fund for Nature - Australia, Sydney.
Greenberg, J. (1984). Avoiding tax avoidance: A (repeated) game-theoretic approach, *Journal of Economic Theory* 32, 1-13.
Hanrahan, C. E. and Zinn, J. (2005). *Green payments in U.S. and European Union agricultural policy*. Congressional Research Service report for Congress, Washington.
Harrington, W. (1988). Enforcement Leverage When Penalties Are Restricted, *Journal of Public Economics* 37, 240.
Henry, C. (1974). Investment decisions under uncertainty, *The American Economic Review* 37, 240.
Heyes, A. G. (2000). Implementing environmental regulation: compliance and enforcement, *Journal of Regulatory Economics* 17, 107-129.
Heyes, A. G. (2002). A Theory of Filtered Enforcement, *Journal of Environmental Economics and Management* 43, 34-46.
Heyes, A. G. and Rickman, N. (1999). Regulatory Dealing - Revisiting the Harrington Paradox, *Journal of Public Economics* 72, 361-378.
Hooper, J. (1992). Measurement and perception of change: field monitoring of environmental change in the Environmentally Sensitive Areas, in Whitby, M. C. (eds.), *Land Use Change and it Consequences*. HMSO, London.
Kalbfleisch, J. D. and Lawless, J. F. (1985). The Analysis of Panel Data Under a Markov Assumption, *Journal of the American Statistical Association* 80, 863-871.
Lane, D. (1989). A Partially Observable Model Of Decision Making By Fishermen, *Operations Research* 37, 240.
Logofet, D. O. and Lesnaya, E. V. (2000). The mathematics of Markov models: What Markov chains can really predict in forest successions, *Ecological Modelling* 126, 285-298.
Malik, A. S. (1990). Avoidance, Screening and Optimum Enforcement, *The Rand Journal of Economics* 37, 240.
Mallawaarachchi, T., Blamey, R. K., Morrison, M. D., Johnson, A. K. L. and Bennett, J. W. (2001). Community values for environmental protection in a cane farming catchment in Northern Australia: A choice modelling study, *Journal of Environmental Management* 62, 301-316.
Monahan, G. E. (1982). A Survey of Partially Observable Markov Decision Processes: Theory, Models, and Algorithms, *Management Science* 37, 240.
National Audit Office. (1997). *Protecting Environmentally Sensitive Areas*. HMSO, London.
NE. (2006). *Strategic Direction 2006-2009*. Natural England, London.
Rushton, S. P., Sanderson, R. A., Wildig, J. and Byrne, J. P. (1996). The effects of grazing management on moorland vegetation: a comparison of farm unit, grazing paddock and plot experiments using a community modelling approach, *Aspects of Applied Biology* 44, 211-219.

Smallwood, R. D. and Sondik, E. J. (1973). The optimal control of partially observable markov processes over a finite horizon, *Operations Research* 21, 1071-1088.
Usher, M. B. (1979). Markovian Approaches to Ecological Succession, *Journal of Animal Ecology* 48, 413-426.
Walters, C. J. and Holling, C. S. (1990). Large-Scale Management Experiments and Learning by Doing, *Operations Research* 37, 240.
White, B. (2000). An economic analysis of ecological monitoring, *Ecological Modelling* 189, 241-250.
Williams, B. K. (1996). Adaptive optimization and the harvest of biological populations, *Mathematical Biosciences* 136, 1-20.
World Bank. (1998). *Guidelines for Monitoring and Evaluation of Biodiversity Projects*. World Bank Global Environmental Division, Washington.
Yates, C. J. and Hobbs, R. J. (1997). Woodland restoration in the Western Australian wheatbelt: A conceptual framework using a state and transition model, *Restoration Ecology* 5, 28-35.