ECG 590 Course Syllabus

ECG 590 – Dynamic Environmental and Resource Management

Section 002

SPRING 2017

3 Credit Hours

Course Description

This course will examine inter-temporal problems involving the management of renewable and non-renewable resources. Examples include the conservation of endangered species, the determination of optimal harvesting levels and the setting of harvest regulations, the management of pollution and the control of invasive pests. All of these issues involve important trade-offs between current and future costs and benefits. These trade-offs will be examined by casting the problems as dynamic optimization problems. Model solutions will then be obtained and interpreted. We will also examine some of the thorny issues involved in economic analysis of these topics, including model uncertainty and discounting. Students can expect to gain an appreciation of inter-temporal allocation problems and to be able to formulate and analyze dynamic optimization models. This course is designed to be of use to economics majors interested in environmental and resource economics as well as to students in biology, forestry and wildlife and natural resource management interested in modeling and decision analysis.

Learning Outcomes

By the end of this course, students will be able to:
.......***identify*** the components of Markov decision problems (MDPs)
.......***formulate*** environmental and resource management problems as MDPs
.......***be familiar with*** a range of problems that can be addressed as MDPs
.......***understand*** how alternative kinds of uncertainty can be addressed
.... ***utilize*** appropriate software for solving MDPs
.......***analyze and interpret*** model results
.......***describe*** management plans to stakeholders

Course Structure

Course activities include regular participation in the weekly (Tuesday 1:30-4:15) on-line class including participation in discussion and model development sessions. Outside of class students are expected to familiarize themselves with material provided in readings and video lectures and to complete homework assignments. Students are also expected to complete a term project involving replication of published research.

Recordings of each class will be available soon after the class is over (within the next day). URLs for the recordings will be posted to the course web site. If you must miss a class please view the recording at the earliest possible time. You can also use the recording to review material that was unclear.

Instructors

Paul L. Fackler (pfackler) - *Instructor*
Email: paul_fackler@ncsu.edu
Phone: 919-515-4535
Office Location: Nelson 4344
Office Hours: by email appointment

Course Meetings

Seminar

Days: T
Time: 1:30pm - 4:15pm
Campus: Main
Location: Internet via Blackboard Collaborate
This meeting is required.

Meeting Notes

Students must have access to a computer high speed internet and a headset.

Course Materials

Textbooks

None.

Expenses

None.

Materials

None.

Requisites and Restrictions

Prerequisites

Mathematical prerequisites for the course are minimal (no calculus is used). Students would benefit from either intermediate micro-economics and/or familiarity with resource modelling in forestry, ecology, wildlife management or a related field. In addition familiarity with basic concepts of probability will be very useful. The course will use the MATLAB based package MDPSolve to obtain model solutions; computer programming tools used will be covered in the course though some familiarity with computer programming would be beneficial.

Co-requisites

None.

Restrictions

None.

General Education Program (GEP) Information

GEP Category

This course does not fulfill a General Education Program category.

GEP Co-requisites

This course does not fulfill a General Education Program co-requisite.

Transportation

This course will not require students to provide their own transportation. Non-scheduled class time for field trips or out-of-class activities is NOT required for this class.

Safety & Risk Assumptions

None.

Grading

Grade Components

Component	Weight	Details
Mid-term exam	**NA**	Students are expected to display basic competancy on the exam.
Class participation	**30%**	
Homework	**30%**	
Term project	**40%**	

Requirements for Credit-Only (S/U) Grading

Performance in research, seminar and independent study types of courses (6xx and 8xx) is evaluated as either "S" (Satisfactory) or "U" (Unsatisfactory), and these grades are not used in computing the grade point average. For credit only courses (S/U) the requirements necessary to obtain the grade of "S" must be clearly outlined.

Requirements for Auditors (AU)

Information about and requirements for auditing a course can be found at http://policies.ncsu.edu/regulation/reg-02-20-04.

Policies on Incomplete Grades

If an extended deadline is not authorized by the Graduate School, an unfinished incomplete grade will automatically change to an F after either (a) the end of the next regular semester in which the student is enrolled (not including summer sessions), or (b) by the end of 12 months if the student is not enrolled, whichever is shorter. Incompletes that change to F will count as an attempted course on transcripts. The burden of fulfilling an incomplete grade is the responsibility of the student. The university policy on incomplete grades is located at http://policies.ncsu.edu/regulation/reg-02-50-03. Additional information relative to incomplete grades for graduate students can be found in the Graduate Administrative Handbook in Section 3.18.F at http://www.fis.ncsu.edu/grad_publicns/handbook/

Late Assignments

Late assignments will not count toward grade unless prearranged with the instructor or due to an appropriate personal emergency.

Attendance Policy

For complete attendance and excused absence policies, please see http://policies.ncsu.edu/regulation/reg-02-20-03

Attendance Policy

None.

Absences Policy

None.

Makeup Work Policy

None.

Additional Excuses Policy

None.

Academic Integrity

Academic Integrity

Students are required to comply with the university policy on academic integrity found in the Code of Student Conduct found at http://policies.ncsu.edu/policy/pol-11-35-01

All work submitted for evaluation (homework, exam and term project) should be produced by the student and should represent their own efforts. Verbal collaboration on homework and the term project (but not on the exam) is appropriate but students should submit their own work.

Academic Honesty

See http://policies.ncsu.edu/policy/pol-11-35-01 for a detailed explanation of academic honesty.

All work on the exam should be the student's own and answers should be written in their own words.

Honor Pledge

Your signature on any test or assignment indicates "I have neither given nor received unauthorized aid on this test or assignment."

Electronically-Hosted Course Components

Students may be required to disclose personally identifiable information to other students in the course, via electronic tools like email or web-postings, where relevant to the course. Examples include online discussions of class topics, and posting of student coursework. All students are expected to respect the privacy of each other by not sharing or using such information outside the course.

Electronically-hosted Components: Moodle discussion forum and recorded classroom sessions.

Accommodations for Disabilities

Reasonable accommodations will be made for students with verifiable disabilities. In order to take advantage of available accommodations, student must register with the Disability Services Office (http://www.ncsu.edu/dso), 919-515-7653. For more information on NC State's policy on working with students with disabilities, please see the Academic Accommodations for Students with Disabilities Regulation at http://policies.ncsu.edu/regulation/reg-02-20-01.

Non-Discrimination Policy

NC State University provides equality of opportunity in education and employment for all students and employees. Accordingly, NC State affirms its commitment to maintain a work environment for all employees and an academic environment for all students that is free from all forms of discrimination. Discrimination based on race, color, religion, creed, sex, national origin, age, disability, veteran status, or sexual orientation is a violation of state and federal law and/or NC State University policy and will not be tolerated. Harassment of any person (either in the form of quid pro quo or creation of a hostile environment) based on race, color, religion, creed, sex, national origin, age, disability, veteran status, or sexual orientation also is a violation of state and federal law and/or NC State University policy and will not be tolerated. Retaliation against any person who complains about discrimination is also prohibited. NC State's policies and regulations covering discrimination, harassment, and retaliation may be accessed at http://policies.ncsu.edu/policy/pol-04-25-05 or http://www.ncsu.edu/equal_op/. Any person who feels that he or she has been the subject of prohibited discrimination, harassment, or retaliation should contact the Office for Equal Opportunity (OEO) at 919-515-3148.

Detailed Outline

Topic 1: Introduction to Modeling

A mathematical approach to decision problems involves specifying a model that describes a system, a set of actions that can be taken and how those actions affect the system and one or more performance outcomes. In addition the utility obtained for alternative values of the performance variables must be specified. With these elements a best strategy can be determined that provides a rule for picking a best action given the current state of the system. An example involving a captive breeding program illustrates the ideas.

Topic 2: Markov Processes

Markov Processes are the primary way we will specify models of dynamic systems. Markov processes possess the important feature that the future behavior of the system depends only on the current state of the system and not on its past. The discussion of Markov processes will be broken into two parts. The first covers simple processes, the second covers controlled processes which are affected by actions that are under our control. Here we will learn how to specify a Markov processes via its transition matrix and how to characterize the future behavior of the system using that matrix. We will also explore the concept of a strategy which is a choice of which actions to take for each value of the system state.

Topic 3: Dynamic Programming

The combination of a system modeled as a controlled Markov process and a reward (utility) function that evaluates outcomes is a Markov Decision Problem (MDP). Dynamic programming can be used to determine an optimal strategy for an MDP. A pest management example is used to illustrate the approach.

Topic 4: Examples of DP

A number of examples of simple dynamic programming models will be developed. The goal here is to present some of the variety of problems that can be addressed and to reinforce the basic framework, which involves the specification or state and control variables, state transition probabilities, and the reward function. The examples include managing a timber stand, a mine, a reservoir and the harvest of wild biological resource.

Topic 5: Solving DP Problems and a Template for Using MDPSolve

Dynamic programming problems are solved by working backwards over time. Bellman’s equation and the basic algorithms for solving DP problems are presented. Having a standard system (a template) for specifying a decision problem can greatly aid its development and communication with others. Here such a template for using the MDPSolve software is described. It involves specifying parameter values, state and control variable values, reward function values and state transition probabilities. The solver is then called and the solution results obtained. Finally the results are analyzed via tables and figures.

Topic 6: Continuous States

In many systems it is natural to treat some variables as taking on a continuum of values. In order to use the discrete MDP framework it is necessary to discretize these variables and to obtain a transition matrix that defines a discrete process with similar characteristics to the continuous process. Transitions for continuous variables are generally defined in terms of a transition function which gives the value of the future state in terms of the current state action and one or more random noise terms. It can also be expressed in terms of a transition probability density function for the future state in terms of the current state and action. Three discretization methods are discussed, two for the former case and one for the latter. The approaches will be illustrated with an example involving the regulation of salmon harvests.

Topic 7: Multiple State and Control Variables

Most real-world problems involve multiple state variables and multiple control variables may also be used. Most of the tools developed thus far extend to multiple variables in a simple way. The main difficulty that arises in working with multiple variables is in the specification of the transition probabilities. The use of conditional probability will be illustrated with a classic decision problem involving the regulation of mallard duck hunting.

Topic 8: Analysis of Time Paths and Long-Run Probabilities

Here we take up in greater detail ways of analyzing optimal strategies. First we discuss how sensitivity analysis is affected by the choice of model parameterizations. Then the use of time paths of expected values of system variables and long run probabilities is discussed. Tools for computing long-run distributions will be demonstrated. For problems with the possibility of absorption, such as extinction models, the notion of the quasi-stationary distribution, which measures the long-run distribution given that absorption has not yet occurred, is discussed.

Topic 9: Discounting

Discounting causes utility that accrues sooner to be worth more than utility that accrues later. There are a number of reasons why this can be a useful feature in a decision model. These include a natural preference, the possibility of investing resources in alternatives that generate a positive rate of return, the conviction that future generations will be better off than the current one and as a way to address future risks. We will examine some of the alternative rationales as well as the impact of discounting on optimal decisions. We will also examine government guidelines in setting discount rates.

Topic 10: Ergodic Control

When the discount rate is 0, so the utility derived from a given state and action is the same regardless of how far into the future the utility is received, value functions may not converge. In such cases it is generally more appropriate to maximize the average reward obtained in any given future period; this is called ergodic control. Because the value function does not converge, the usual Bellman's Equation must be modified. Issues arising in this situation are examined and some models are developed. Two examples are developed, one involving species extinction and the second involving sustainable harvest.

Topic 11: Models with Stages

It is sometimes useful to be able to break a full time cycle into sub-periods and to make decisions in each of these sub-periods. Examples include situations in which seasonal transitions or rewards differ. For example, survival rates differ in summer and winter and breeding often takes place in only one season. Models with stages can also be used to describe situations with non-stationarity in which the transition probabilities are changing over time. A simple hunting example will illustrate the approach. A more involved endangered species example with both habitat management and predator control is developed to illustrate a situation in which the state variables are not the same in the different stages.

Topic 12: Explicit Spatial Models

Models with multiple non-homogeneous sites are examined. Each site can be in one of several categories. The state is a vector that specifies the current category for each site. The probabilities associated with transitioning to another category can be site specific and can depend on the entire state. In addition the rewards can be site specific. It is difficult to obtain exact solutions for such models except with a small number of sites and categories. These models can be especially useful in modeling site occupancy for threatened species having few sites remaining with suitable habitat. To illustrate explicit spatial modeling a problem involving control of a pest infestation is developed.

Topic 13: Implicit Spatial (Category Count) Models

Models with multiple homogeneous sites are examined. With homogeneous sites only the number of sites in each category need be known. Such models greatly reduce the problem size, making it possible to obtain optimal strategies for problems with more sites and/or categories. Implicit spatial models have been used to address problems including pest management and conservation of endangered species. Categories in such models include the population level or the habitat type or quality. To illustrate implicit spatial modeling a problem involving control of a pest infestation is developed.

Topic 14: Structural Uncertainty and Adaptive Management

It is generally true that system models cannot be specified with certainty. Both model parameters and functional forms are developed from data analysis and expert opinion and are subject to error. One approach to address such uncertainty specifies a set of alternative models, each of which is assigned a degree of belief. As new information becomes available the beliefs weights can be updated using Bayes Rule. It may be optimal to take actions that specifically aim at obtaining information if it contributes to better future decisions. The basic concepts of probing actions and the value of learning will be defined and illustrated. The distinction between active and passive adaptive management will be discussed. To illustrate a model of optimal harvest regulation will be developed.

Topic 15: Observational Uncertainty and POMDPs

In many situations one or more of the state variables cannot be observed without significant error. For example, populations of wild animal stocks are estimated using sampling procedures that may be subject to considerable error. Such situations are said to exhibit observational error. One way to address such problems is to replace the unobserved state variables with a probability distribution that measures the degree of belief attached to alternative values of the variables. This leads to the so-called Partially Observable MDP (POMDP). The use of POMDPs will be demonstrated and an example concerning pest management will illustrate the approach.