INSTALLING AND USING
THE ADAPTIVE RESOURCE MANAGEMENT SIMULATOR

Clint Moore
U.S. Fish and Wildlife Service
January 10, 1995

1.0 INTRODUCTION

Managers of waterfowl populations face the difficult task of recommending management actions for a resource that is measured only crudely, that may associate with one or more stochastic environmental phenomena, and that responds unclearly to management actions. Furthermore, a realized management action (e.g., a particular rate of kill) may not necessarily reflect the desired action, or even be measured finely. Each manager has his own notion, or conceptual model, about the dynamics of the population and how the population does or does not respond to specific management actions. Population data are generally so coarse that often, two or more competing models may all be found consistent with a single set of data. Even if managers were in universal agreement on the structure of a population model, the trouble remains of making management choices that are optimal with respect to some criterion; optimal choices are not clear-cut for resources that are stochastic and imperfectly measured and controlled. Lacking rigorous, objective means to evaluate competing models and to make optimal choices, recommendations are prone to be based on emotions and politics rather than science, and the decision-making process can be characterized by trepidation and discord.

Adaptive Resource Management (ARM) seeks management choices that are optimal with respect to some objective despite uncertainty about population dynamics, environmental stochasticity, imprecision in population measurement, and inability to implement actions exactly as intended. ARM comprises two major mechanisms: a model-credibility, or learning, component and an optimization component. The learning component assigns weights of credibility to each competing model. These weights are functions of measures of agreement between model predictions and current-year estimated population size and of weights from the prior year. The optimization component is a procedure that prescribes a strategy for optimal management given (1) an objective function, (2) an array of management choices, and (3) a set of models describing population dynamics. Choices indicated by the strategy are time-independent, indexed only by the measured state of the population (and other quantities which affect population dynamics) and by model weights.

This simulation program demonstrates both mechanisms. The program was developed to show the user how ARM incorporates uncertainty into the decision-making process, how our ability to manage is enhanced or hindered depending on our control of variable processes, how optimal management and rapid learning are associated, and how a real implementation of ARM might look. The program simulates population dynamics of the mid-continent mallard population. The dynamics incorporate environmental and recruitment models described by Anderson (1975). Pond abundance is a stochastic element that drives the recruitment model. The assumptions of additive versus compensatory harvest mortality are encoded as competing mortality models. A limited set of harvest options are available at each time period. Partial system observability (measurement imprecision) and partial harvest controllability (inability to execute targeted harvest rate) are simulated in the program and are controlled by the user.

2.0 INSTALLATION

The ARM Simulator may be run on a 386 (with 387 math chip) or higher CPU that has at least 4 MB RAM. The program requires 3 MB of hard disk space, and it runs under DOS version 3.3 or higher. A color monitor is recommended but not required. A mouse is not required.

2.1 DOS Installation

1. Check CONFIG.SYS for presence of the ANSI.SYS driver. If the driver is not called, edit CONFIG.SYS to load it upon boot-up. Generally, ANSI.SYS is found in C:\DOS, so the command DEVICE=C:\DOS\ANSI.SYS (or DEVICEHIGH if the DOS upper memory manager is being used) should suffice. Re-boot the computer after editing CONFIG.SYS.

2. Insert installation disk in drive A and run A:\INSTALL at the DOS prompt. If the 3.5" drive is called B, then run B:\INSTALLB. In either case, the program will be installed to a directory named C:\ARM_SIM on the hard disk.

3. Change to the C:\ARM_SIM directory, and invoke the program by typing ARM_SIM at the DOS prompt.

2.2 Windows Installation

1. Follow steps 1 and 2 above to assure that ANSI.SYS is loaded and to install the software.

2. Start Windows and open a new or existing program group. Create a new program item called "ARM Simulation Program" with the following properties:

   command line: C:\ARM_SIM\ARM_SIM.BAT
   working directory: C:\ARM_SIM

3. Invoke the program by double-clicking on the new program icon.

3.0 PROGRAM OPERATION

The simulation program features two primary modes of operation. The Interactive mode mimics a management scenario in which a mallard population is monitored over time and the user must make a management (harvest) decision at each time period. Information available to the user at time T includes estimates of the population size and pond abundance, and the likelihood that the additive mortality model is driving the observed system (of course, the complement of this value is the likelihood that the system is compensatory). System partial observability (agreement between population estimate and population size) and partial controllability (agreement between target harvest rate and realized harvest rate) are set beforehand by the user. Additionally, the user may choose to select the mortality model that the program will use.

The Sensitivity mode replicates entire series of management decisions. This mode is useful for studying behavior of expected annual harvest and expected model probabilities at varying levels of (1) target harvest rate, (2) partial observability, and (3) partial controllability. Through stochasticity, no pair of identical series of management decisions is expected to produce the same outcome. The replication of series smooths out stochasticity and gives an impression of the central tendency of the outcome.

The first screen that appears when running the ARM Simulator prompts the user for the primary operating mode. Valid responses are "I" for Interactive, "S" for Sensitivity, or "Q" to Quit. Typing "Q" at this prompt ends the program. Several features are common throughout the program. Any response may be entered in upper or lower case. Typing a "Q" at any other prompt in the program returns the user immediately to the opening screen. For character responses, each prompt lists all response options available. For numeric responses, each prompt specifies a range of valid input. In either case, any invalid response causes the program to re-prompt for input. All numeric responses have a default response, indicated on the prompt. Hitting <Enter> affirms program use of the default value. For screen graphics, hitting <Esc> clears the graph. The user has an opportunity to re-display the output screens prior to starting a new session.

3.1 Interactive mode

Prompt 1: Number of years to simulate. The user enters an integer between 1 and 100 (default = 10 years).

Prompt 2: "Divinity" mode. The user may decide whether the population model is chosen by the user (the user becomes "God", response "G") or by the program (the user is a "Mortal", response "M"). If the user selects "G", the user is then prompted to select either the model of additive (response "A") or compensatory (response "C") mortality. Otherwise, the program selects one of these models covertly at random.

Prompt 3: Partial observability. The user enters a value between 0 and 100 that represents the coefficient of variability (CV, %) of population estimates.

Prompt 4: Partial controllability. The user enters a CV value (%) between 0 and 100 that represents how poorly the realized harvest rate agrees with the target harvest rate.

Following this prompt, the year-by-year simulation begins. The user is presented a full screen of information on population size, pond abundance, fall flight, likelihood of the additive model, and harvest statistics. Furthermore, information on true population size, true pond abundance, true fall flight, and realized harvest rate is not revealed if the population model was selected by the program (a mortal can't know these things).

The user is also presented with an array of target harvest rates. These rates are keyed to single-character codes: "C" (closed, target rate = 0%), "R" (restrictive, 10%), "M" (moderate, 20%), "L" (liberal, 30%), "X" (extra-liberal, 40%), or "S" (slay, 50%). The program expects a target harvest rate indicated by one of the codes. Above one or more of these targets are choices for optimal harvest assuming the additive model is correct (symbol "A"), the compensatory model is correct ("C"), or neither model is certain, but that one is correct ("D", for aDaptive). The user is free to take a target indicated by any of the 3 optimal choices, or to choose one of his own.

At the end of the time series, a screen graphic is produced. The estimated population size appears as a thin solid line, and the true population size is shown as a thin broken line. Predictions of population size under the additive and compensatory models are displayed as triangles and circles, respectively. The likelihood of the additive model is shown as a bold solid line (scaled 0.0-1.0). This line moves toward 1.0 when the estimated population size is closest to the additive prediction and toward 0.0 when closest to the compensatory prediction. When the screen is cleared, the program displays the mortality model used in the simulation. The user can then review the graphic, if desired.

3.2 Sensitivity mode

Prompt 1: Number of years to simulate. The user enters an integer between 1 and 100 (default = 20 years).

Prompt 2: Number of sensitivity iterations. The user enters an integer between 1 and 100 (default = 20 iterations). This value controls the number of iterations of each of the time series defined in prompt 1.

Prompt 3: Select model. The user selects either the additive ("A") or compensatory ("C") mortality model.

Prompt 4: Variable for sensitivity analysis. The user specifies whether target harvest rate (choice "H"), partial observability (choice "O"), or partial controllability (choice "C") is adjusted in five equally-spaced steps.

Prompt 5: If harvest rate is being analyzed, the user is asked to specify lower and upper sensitivity values for harvest rate (%). Otherwise, the user supplies a fixed value for harvest rate.

Prompts 6-7: Similar to prompt 5, but for partial observability and partial controllability, respectively.

The graphic produced is divided into two sections. The top section displays mean annual harvest, averaged over the sensitivity iterations, against year. One line is drawn for each of the five values of the sensitivity analysis variable, and line thicknesses are proportional to variable value. Seven lines appear when harvest rate is analyzed: the extra two lines represent mean annual harvest under the additive and compensatory optimal strategies. That is, the program simulates five constant-level harvest rate strategies alongside two varying-level strategies. The bottom section plots the average "learning curve", i.e., the average likelihood of the additive model, and the display is similar to that of the top. In the right-hand margin of each section is a display of mean standard deviations, averaged over the time series. This display reveals relative variability either in annual harvest or learning rate among the levels of the sensitivity variable.

When the graphic is cleared, a table of summary statistics is produced. The table summarizes the model selected, the range of the sensitivity variable, and the fixed values for the other variables. For each level of the sensitivity variable, the program provides number of iterations that the population time series crashed, mean final population size, mean annual harvest (standard error), and mean cumulative harvest (standard error). The right-hand block of the table provides the year at which the mean learning curve exceeded the 0.90, 0.95, 0.99, 0.999 cutpoints for the additive model, or descended below the 0.10, 0.05, 0.01, 0.001 cutpoints for the compensatory model. If harvest rate was analyzed, two additional rows of information for the optimal additive and compensatory strategies are provided.

The graphic and the table help quantify the price paid either in harvest or in learning for specific levels of harvest rate, partial observability, or partial controllability. At the bottom of the table, a prompt appears to allow the user to review the graphic and table.