

Human judgment vs. quantitative models for the management of ecological resources

MATTHEW H. HOLDEN^{1,2,4} AND STEPHEN P. ELLNER^{1,3}

¹*Center for Applied Mathematics, Cornell University, Ithaca, New York 14853 USA*

²*ARC Centre of Excellence for Environmental Decisions, University of Queensland, 523 Goddard, St Lucia 4072 Australia*

³*Department of Ecology and Evolutionary Biology, Cornell University, Ithaca, New York 14853 USA*

Abstract. Despite major advances in quantitative approaches to natural resource management, there has been resistance to using these tools in the actual practice of managing ecological populations. Given a managed system and a set of assumptions, translated into a model, optimization methods can be used to solve for the most cost-effective management actions. However, when the underlying assumptions are not met, such methods can potentially lead to decisions that harm the environment and economy. Managers who develop decisions based on past experience and judgment, without the aid of mathematical models, can potentially learn about the system and develop flexible management strategies. However, these strategies are often based on subjective criteria and equally invalid and often unstated assumptions. Given the drawbacks of both methods, it is unclear whether simple quantitative models improve environmental decision making over expert opinion. In this study, we explore how well students, using their experience and judgment, manage simulated fishery populations in an online computer game and compare their management outcomes to the performance of model-based decisions. We consider harvest decisions generated using four different quantitative models: (1) the model used to produce the simulated population dynamics observed in the game, with the values of all parameters known (as a control), (2) the same model, but with unknown parameter values that must be estimated during the game from observed data, (3) models that are structurally different from those used to simulate the population dynamics, and (4) a model that ignores age structure. Humans on average performed much worse than the models in cases 1–3, but in a small minority of scenarios, models produced worse outcomes than those resulting from students making decisions based on experience and judgment. When the models ignored age structure, they generated poorly performing management decisions, but still outperformed students using experience and judgment 66% of the time.

Key words: adaptive management; bioeconomics; conservation; ecological modeling; expert judgment; fisheries management; natural resource management; optimal harvest.

INTRODUCTION

In the past 50 yr, environmental management has benefited from major advances in decision science. Perhaps the most influential concept among these advances is adaptive management, the iterative process of modeling, hypothesis testing, optimization, acting, and monitoring to reduce uncertainty and maximize net benefits (Holling 1978, Walters 1986). Government agencies, scientists, and theoreticians widely agree that adaptive management is the best way to manage a biological population in cases where the benefit of different actions strongly depends on uncertain ecological processes that can be learned through observing system changes in response to management (Possingham *et al.* 2001, Stankey *et al.* 2005, Nichols and Williams 2006, Walters 2007, Williams and Brown 2012, Game *et al.* 2014).

While managers often do practice some components of adaptive management by collecting data and making decisions based on their findings, with the exception of a few large-scale management programs in fisheries, waterfowl, forestry, and conservation (e.g., Sainsbury 1988, Moore and Conroy 2006, Johnson *et al.* 2015, Nichols *et al.* 2015), managers rarely use dynamic modeling and optimization, and instead use their experience, intuition, and best judgment as a substitute for formal system analysis (Johnson and Williams 2015). This is despite the fact that many scientists have proposed management plans based on quantitative optimization methods, for a variety of ecological systems, which in theory offer managers substantial cost savings and improved environmental outcomes (e.g., McCarthy *et al.* 2001, Westphal *et al.* 2003, Gerber *et al.* 2005, Wilson *et al.* 2006, Asano *et al.* 2008, Johnson *et al.* 2011, Martin *et al.* 2011, Probert *et al.* 2011, Helmstedt *et al.* 2014, Hughes *et al.* 2014, Rout *et al.* 2014).

One potential reason for the resistance to using mathematical modeling in management is that it's unclear how much modeling and optimization actually improve

Manuscript received 13 July 2015; revised 10 November 2015; accepted 6 January 2016. Corresponding Editor: F. A. Johnson.

⁴E-mail: m.holden1@uq.edu.au

management outcomes over expert opinion. This is especially a concern when model-based decisions are calculated using passive dynamic optimization (Johnson and Williams 2015). The defining feature of passive optimization is that the method does not consider the value of information while solving for the optimal action, meaning that a manager never sacrifices expected gains, given current information, in order to learn about the system and potentially improve long-term benefits.

When the value of improved system knowledge resulting from each action is incorporated explicitly into the objective, the optimization is referred to as active adaptive management. Passive adaptive management incorporates learning based on observations, but active adaptive management also values the future benefit of knowledge resulting from decisions made in the present. Unfortunately, unless a manager is willing to simplify their description of the management problem (e.g., Hauser and Possingham 2008), active adaptive management is often computationally infeasible, and hence passive optimization is the predominant method for solving management problems (Johnson and Williams 2015).

Humans can possibly use intuition and past experience to incorporate the benefit of learning into decision making, without the aid of mathematical models. Can humans use their flexibility to learn about the system to outperform a model-based, passive adaptive management program? Unfortunately, it is difficult to answer this question because experiments in management are, in general, not repeatable. That is, once a manager makes a decision based on their expertise, it is usually impossible to compare the outcome to how well an alternative decision, aided by a mathematical model, would have done.

In this study we take a first step towards quantifying the economic benefits of using simple dynamic models and passive optimization methods to aid environmental management decisions, rather than solely relying on human judgment. To do this, students in multiple college classes played an online game where they managed a simulated fishery. The data from each game was saved on a server, and therefore we were able to compare exactly how model-based decisions would have performed, compared to the students' performance, for each unique instance of the game.

Through an analysis of the users' harvest decisions entered during each instance of the game, not only were we able to compare user performance to the performance of model-based decisions, but we could also begin to understand the strategies students deploy when managing a population without the aid of a mathematical model and compare how their strategies differ from those developed using quantitative techniques.

METHODS

Experiments

Students played two online games, accessed using a web browser, where they earned points corresponding

to the profits from managing a simulated herring and a simulated Pacific salmon fishery. We describe the experiment for the herring fishery game and then explain how the salmon game was different.

The students played the game using their laptops during the lecture period of two courses, "Environmental Conservation" at Cornell University (123 students) and "Principles of Biology" at Ithaca College (60 students), and at the "Graduate Student Science Colloquium" at Cornell University (15 students). Prior to managing each fishery, the students filled out a multiple choice survey that asked them their major, educational experience, fishing experience, and environmental management experience. See Table S1 in the online Supporting Information for a copy of the survey.

After the survey, each game showed a page of directions describing the fish stock's population dynamics. In addition, Matthew Holden, the game facilitator, read a standardized script aloud to each class, reiterating the points listed on the page. This included statements about the existence of a fishery carrying capacity, measurement error, environmental randomness out of the manager's control, and how their performance would be scored. See the online Supporting Information for a copy of the game directions. Before starting the game, each student was randomly assigned a σ between 0 and 0.25, using a uniform distribution. Students with high σ experienced large random variation in stock biomass unrelated to their management actions. Before playing the game the students played an eight-turn practice game. This served three purposes: (1) they developed experience with the fishery, (2) we used the data from the practice game to identify students who didn't understand the directions, and (3) it provided a set of "past data" for the models and students to use as a basis for making decisions in the future.

Before the user entered their first harvest decision in the practice game, they were presented with three harvest data points, and the resulting biomasses from the deterministic version of the model underlying the simulated population dynamics, to give them some context of the range of harvest values they could potentially enter. We chose to use the deterministic model for this purpose so that all users saw the exact same past data before playing the game. A description of the models used to simulate the biomass data observed during the game is presented in *Simulated population dynamics*.

The game showed the user graphs of harvest, estimated remaining biomass in the fishery, and cumulative profit at each time step. See Fig. S1 in the online Supporting Information for a picture of the game display. At the beginning of each turn of the game, the user entered an amount of biomass they wanted to harvest into a text box, clicked enter, and then the remaining biomass, post-harvest, grew according to the models that governed the simulated fishery, and the result was displayed on the screen numerically. In addition, all plots updated, adding the player's harvest choice to the harvest

plot, the resulting biomass to the biomass plot, and the new accumulated profit to the total profit plot.

The user's score was the discounted net profit accumulated over the game, with a discount rate of 0.03 and a constant profit of 10 000 dollars per ton (0.907 megagrams) of biomass caught. In addition, the user received a bonus added to their score at the end of the game, which was the discounted profit that would have been generated by harvesting all of the remaining biomass left in the fishery after the game was over. The bonus provided incentive for users to let fish escape harvest on the last turn. Without the bonus, the user's score would be highly sensitive to their last harvest decision. This bonus is explained to the user in the game directions (see Fig. S3 in the online Supporting Information for a copy of the directions).

After a student completed their last turn, the game displayed their score in addition to a leaderboard, which included the scores and initials of the top players in the class, up to that point in time. The leaderboard provided an external incentive to play well. However, the students did not receive a course grade or monetary incentives based on performance.

Throughout the game, data were stored locally on the user's computer using browser cookies. Upon exiting the game, these anonymous data were sent to a server, using PHP (a server-side programming language for web development; Welling and Thomson 2003), and stored in a database. These data included the time the user finished playing the game, an anonymous user ID number, the student's answers to the survey questions, the environmental noise variable σ , total profit (i.e., "points"), and their time series of harvest decisions, resulting biomasses, realizations of environmental noise, and measurement error, and in addition the analogous data from their practice game. By recording the environmental noise and measurement error values experienced by the user, we were able to compare how any strategy (in our case, strategies generated by optimization) would have performed playing that user's exact instance of the game.

After playing the unstructured herring game, the student was directed via a link to the salmon game. Using cookies, the anonymous user ID number from the herring game was saved and recorded along with a unique user ID number for the salmon game as well. In the salmon game, the fishery population dynamics were age-structured, so the game directions also included information on the salmon's life cycle, which consisted of juvenile (1-yr-old) and immature (2-yr-old) fish survival and growth and adult fish (3-yr-old) reproduction. On each turn of the game, the user entered the biomass of adult and immature fish they chose to harvest in two side-by-side text boxes. Plots of the student's harvest and biomass time-series data were the same as for the herring fishery, except now each plot had two curves, one for immature fish and one for adult fish. The user could not observe or harvest juvenile biomass. See Fig. S2 for a picture of the game display in the age-structured game.

The user's score in the age-structured game was similar to the unstructured game, except discounted net profit was summed over both adult and immature harvest, and the bonus was the discounted profit that would have been generated by harvesting all of the remaining adult biomass for 3 yr after the game was over (it takes 3 yr for the recruits at the end of the game to return to be harvested as adults).

Another goal of this study was to collaborate with instructors to incorporate the game into their curriculum to facilitate active learning. Therefore, while the students played each version of this game multiple times, for pedagogical reasons, students were only asked to try their hardest to score the most amount of points possible during their first game. After everyone had finished their first game, they were allowed to collaborate and experiment, to facilitate students learning the principles of conservation biology, and therefore we did not include the students' latter turns in the analysis.

While we never explicitly told the students not to perform any modeling, the game facilitator did not observe any of the students writing down models on paper or using statistical software packages on their computers during the experiment.

Simulated population dynamics

The herring fish game was governed by a simple unstructured, one-dimensional model, where the manager chooses to harvest h_t tons of biomass in year t , and the resulting biomass in year $t + 1$, B_{t+1} , is a nonlinear function of the biomass that escaped harvest in year t , $R(B_t - h_t)$, times a log-normally distributed random number, z_t , with mean one and SD σ .

$$B_{t+1} = z_t R(B_t - h_t). \quad (1)$$

We choose R to be the Beverton-Holt recruitment function, to exclude the possibility of complicated chaotic and periodic dynamics in the absence of harvest,

$$R(B) = \frac{b_1 B}{1 + b_2 B} \quad (2)$$

where b_1 is recruitment per unit biomass at low densities and b_2 controls the carrying capacity of the population (Jorgensen and Fath 2008).

The student managing the population observes a stock biomass of $m_t B_t$, in year t , where m_t is a log-normally distributed random variable with mean one and SD 0.025. The small random variation in m_t represents measurement error in assessing the current fish abundance.

The age-structured fish game is based on the life cycle of coho salmon (*Oncorhynchus kisutch*), including three independent cohorts that undergo a three-stage life cycle. Juvenile fish live in the river and survive and grow into small fish that swim downstream to the ocean where they mature, and finally swim upstream to spawn and die. The manager sets a total catch of $h_{2,t}$ for immature

fish and $h_{3,t}$ for adult fish. Adult fish harvest occurs prior to recruitment, giving population dynamics

$$\begin{aligned} B_{1,t+1} &= z_t R(B_{3,t} - h_{3,t}) \\ B_{2,t+1} &= z_t a_{21} B_{1,t} \\ B_{3,t+1} &= z_t a_{32} (B_{2,t} - h_{2,t}), \end{aligned} \quad (3)$$

where a_{ij} is the per unit biomass contribution, from age j biomass that escaped harvest, in year t , to age i biomass, in year $t+1$.

We parameterized the two models by starting with rough estimates from the literature and then adjusted the values so that the growth rate of our hypothetical herring (unstructured) and coho salmon (age-structured) populations, at low densities and in the absence of harvest, were equal. The reason for using equal growth rates is that when comparing a user's scores from the unstructured and structured population games, we wanted to make sure that any observed difference was due to demographic structure and not due to differences in the absolute growth rate.

The average 3-yr-old coho salmon weighs 8.0 pounds (1 pound = 0.45 kg) and the average 2-yr-old salmon weighs ~3.1 pounds (Marr *et al.* 1944). A typical survival probability for Pacific salmon populations is 0.8 in good years and 0.28 in bad years (Worden *et al.* 2010). Hence, we fixed $a_{32} = (8\text{lbs}/3.1\text{lbs})(0.8 + 0.28)/2 \approx 1.4$. Coho salmon are more productive than herring (*Clupea harengus*) at low densities, hence we chose to lower salmon recruitment as much as "believably" possible so that the growth rate in our salmon and herring fisheries matched. To do this, we assumed the average survival probability of juvenile salmon was equal to the estimate for bad years. Therefore, with the composite parameter of recruitment at low densities estimated in Worden *et al.* (2010) of 60 juveniles per spawner, we let the product of maximum recruitment and juvenile survival be $b_1 a_{21} = (0.28)(60\text{recruits/spawner})(\text{spawner}/8\text{lbs})(3\text{lbs}/\text{recruit}) \approx 6.6$. Because juvenile fish are not harvested or observed, the exact value of a_{21} and b_1 are unimportant individually, as they only affect the observed immature biomass through their product, and therefore we arbitrarily let them equal 4.4 and 1.5, respectively, so that their product was 6.6.

Our salmon parameters imply that at low density, the population will grow by a factor of $b_1 a_{21} a_{32} = (6.6)(1.4) = 9.24$ over 3 yr. We therefore set herring maximum population growth rate at $b_1 = 2.1$ because $2.1^3 \approx 9.24$. The growth rate reported for herring population dynamics ranges from 1.4 to 1.8 (Bjørndal and Conrad 1987, Nøstbakken and Bjørndal 2003), so while our herring growth rate is high, it is not unreasonably so. Carrying capacity is arbitrarily set to 5400 tons, which determines b_2 for both models.

The raw population data used in the game display was created by simulating the above dynamics with initial conditions of 948 tons of fish in the unstructured game

and 481, 309, and 183 tons of adult, immature, and juvenile fish, respectively, in the age-structured game.

Optimal strategies and statistical analysis

Explicit formulas for the optimal harvest strategy, as a function of the parameters, are well known for the unstructured model, and presented in Reed (1979). A similar optimal harvest rule for the age-structured model is given by Holden and Conrad (2015). In both cases the optimal harvest rule is a fixed escapement strategy, where escapement is the biomass that escapes harvest. In other words, the manager leaves a fixed amount of fish in the ocean and this fixed amount of fish is called the escapement. For the parameters in the game, the optimal escapement in the unstructured game is 2049 tons of fish, and 556 tons of adult fish (and all immature fish, i.e., no immature harvest) in the age-structured game (see case 1 in Holden and Conrad 2015).

The first goal of the experiments was to compare the performance of users to fitted models playing the exact same instance of the game. As a control, we compared both the fitted models' and users' performance to the net discounted profit generated by the optimal constant escapement rule specified previously (i.e., the optimal strategy with the true parameters known).

Because the optimal strategy is a function of the parameters, each quantitative model, managing the simulated population, needs an estimate of these parameters to calculate the optimal harvest policy. During the model's first turn of the game, a computer program estimated the parameter values by fitting the model to the data generated from the users' eight-turn practice game. It harvests the simulated population by substituting these parameter estimates into the equation for optimal escapement (i.e., harvests the population down to the optimal escapement or avoids harvest altogether if the fish stock is already below optimal escapement prior to harvest). After observing the stock biomass resulting from its previous harvest, the program re-estimates the parameters using the practice game data along with this new data point. It then harvests using the optimal escapement strategy based on the new parameter estimates, and the process is continued until the game is over. This process is the classic method of passive adaptive management described by Walters (1986).

The parameter estimation for the unstructured game is performed by minimizing sum of squared errors between the log-transformed recruitment data, $\log[m_{t+1} B_{t+1}]$, and log transformed predicted recruitment under the model, $\log[R(m_t B_t - h_t)]$, using the function "lsqcurvefit" in (MATLAB 2010). Note the raw recruitment data is generated by simulating the parameterized model described in *Simulated population dynamics*.

For the age-structured game, because juvenile biomass is unobservable, the procedure is the same as previously, except predicted recruitment is $a_{21} R(m_{3,t} B_{3,t} - h_{3,t})$ and observed recruitment is $m_{2,t+2} B_{2,t+2}$. The transition

between immature and adult biomass is estimated similarly. It should be noted that because the mean of the lognormal measurement error is not exactly zero, the above regression is slightly biased. However, correcting for this small bias did not affect the results presented in this study.

We consider fitted models with the same functional form (Beverton-Holt recruitment) as the model underlying the simulated population dynamics, and in addition models that incorrectly specify the functional form, discrete logistic (May *et al.* 1976), and Ricker recruitment (Ricker 1954). For the age-structured game we also considered escapement rules based on an unstructured Beverton-Holt recruitment model (as in Eq. 1). To estimate the parameters for this model, the computer minimizes the sum of squared error between the log-transformed aggregate biomass data, $\log[m_{2,t+1}B_{2,t+1} + m_{3,t+1}B_{3,t+1}]$, and the predicted biomass, $\log[R(m_{2,t}B_{2,t} + m_{3,t}B_{3,t} - h_{2,t} - h_{3,t})]$. It then harvests the two age classes in proportion to their respective observed biomasses.

To test whether the percent of optimal profit achieved by the user was correlated with the answers to the survey questions, the standard deviation of environmental stochasticity, and net profit generated during the practice game, not including the bonus, we fitted a linear regression model, using the function “lm” in (R Core Team 2012).

Another goal of the experiments was to analyze what strategies the users were deploying and how well different strategies performed compared to others. We compared the user's behavior to three idealized candidate strategies:

constant harvest, proportional harvest, and constant escapement. Constant harvest means the user enters the same harvest at every time step (harvest = β , where β is the user's mean harvest). Under a proportional harvest strategy the user harvests a constant proportion of the biomass (harvest = $\beta \cdot$ biomass, where β is their harvest proportion). Constant escapement, means the user lets a constant amount of biomass escape harvest (harvest = 0 if biomass $\leq \beta$, harvest = biomass - β if biomass $> \beta$, where β is the biomass they let escape harvest). We fit the three models by calculating the constant harvest, constant escapement, and harvest proportion that minimized the sum of squared errors between the user's observed harvest decision and the predicted harvest under each model, using the function “optimize” in (R Core Team 2012). Then, the users were categorized into the three strategy classes based on which model fit had the lowest sum of squared errors. For the age-structured game we repeated this analysis on adult harvest for simplicity (because exclusive adult harvest is optimal), but the results reported in the following section are similar if total harvest is used instead.

RESULTS

Unstructured population game

All human subjects achieved less discounted net profit than would be achieved using the optimal constant escapement strategy with known parameters (Fig. 1a). This was not 100% certain to occur, because the optimal

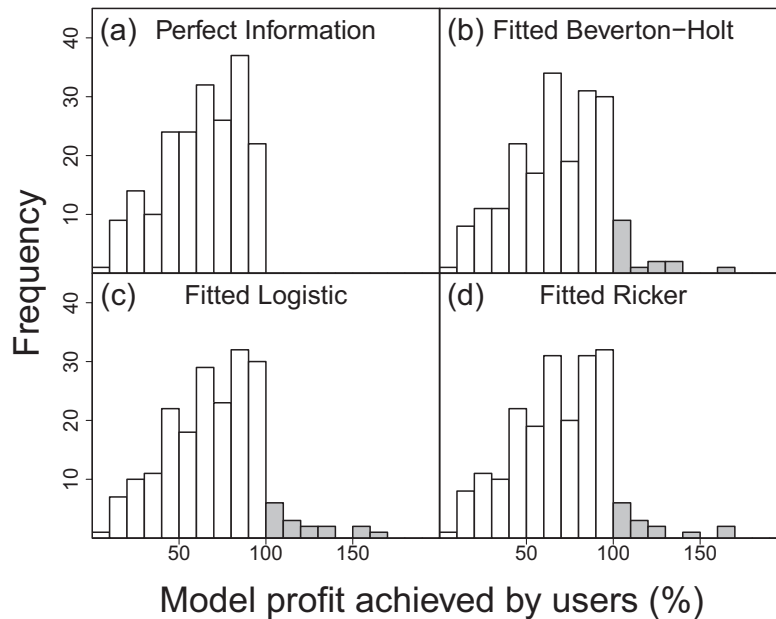


FIG. 1. The percentage of the mathematical model's total profit achieved by the user, in the unstructured game, when the mathematical model is (a) Beverton-Holt with true parameters, i.e., perfect information (b) Beverton-Holt but with parameters estimated from the data, (c) discrete logistic with parameters estimated from the data, and (d) Ricker with parameters estimated from the data. For example, a value of 50% means the user generated half the profit the mathematical model did managing the exact same instance of the game. A value of 200% means the user generated twice as much profit as the model.

strategy is only optimal in expectation, and therefore is not necessarily the most profitable strategy during a run of atypical years. On average, humans scored 65.4% of the discounted net profit generated using the optimal constant escapement strategy, and only 11.0% of students achieved over 90% of this optimal expected net profit.

Most users performed worse than the escapement rules generated from the models with parameters estimated from the historical harvest data (Fig. 1b–d) even if the model made incorrect assumptions about the underlying recruitment function (Fig. 1c, d).

The only significant predictor of the user's performance was their performance in the practice game ($P < 0.001$). Simple linear regression of the user's percent optimal profit on practice game score explained 28% of the variation in the user's percent of optimal profit ($R^2 = 0.28$, see Fig. S6).

When the practice score was removed as a predictor, the user's level of study (1st, 2nd, 3rd or 4th year undergraduate, or PhD student), academic field of study, and SD of the observed environmental stochasticity, still did not significantly correlate with the user's performance. Two predictors were significant in this model. The five

students who responded "I am considering a career in fisheries management, but have no experience" generated more profit than students that responded "I am not considering a career in fisheries management" ($P = 0.033$) and students in Cornell's "Environmental Conservation" course scored significantly higher than the students in Ithaca College's "Principles of Biology" course ($P = 0.041$). However, a linear model with just these two predictor variables only explained 4% of the variation in user performance. It should also be noted that if we group the two students who responded that they actually had fisheries management experience with those five students who indicated a career interest but no experience, the answer to the management experience question does not significantly correlate with the users' scores. This suggests that the sample size for students who were considering careers in fisheries management may be too small to draw any meaningful conclusions.

Classifying the students' harvest strategies into the three categories: constant harvest, proportional harvest, and constant escapement, people harvested a constant proportion (129 users) much more often than allowing a constant escapement (30 users) (Fig. 2). Only 39 users

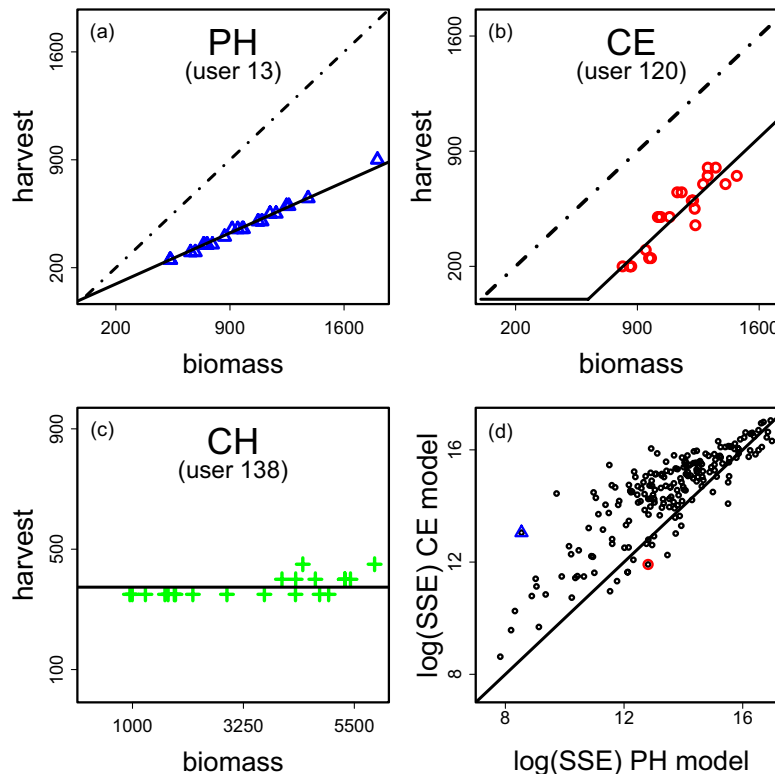


FIG. 2. (a–c) An example user's harvest decisions vs. the biomass they observed prior to making those decisions in the unstructured game. The black line is the best model fit, which for (a) user 13 is a proportional harvest strategy (PH), for (b) user 120 is a constant escapement strategy (CE), and for (c) user 138 is a constant harvest strategy (CH). (d) The sum of squared error when fitting each user's harvest data, in the unstructured game, to a constant escapement model vs. fitting a proportional harvest model on a log-log scale. Points to the right of the 1:1 line represent users whose variation in harvest is better explained by constant escapement than proportional harvest. The blue triangle and red circle in (d) correspond to the proportional harvester and constant escaper in (a) and (b) respectively.

were classified as constant harvesters. Many users repeated their harvest decision from the previous turn and the average user only entered 10 unique harvest values over the course of the 21 turn game (Fig. S7 in the Supporting Information). All five users who indicated a career interest in fisheries management were classified as proportional harvesters. Of the two users with actual management experience, one was a constant harvester and the other was a proportional harvester. The optimal policy type (constant escapement) was the only strategy not used by students with fisheries interests or experience.

Forty-five percent of students allowed less fish to escape harvest, on average, than the optimal value (Fig. 3a). In other words, 45% of users overfished the population while 55% of users underfished the population. If we were to re-classify users whose median escapement was within q percent of the optimal value as neither under nor overfishing, the result that there are roughly the same number of over and underfishers holds for all $q < 70$. So while many students harvested suboptimally, over and underfishing mistakes were equally likely.

Students who used constant escapement strategies (circles in Fig. 3a) were more likely to overfish than to underfish. Proportional harvesters (triangles in Fig. 3a)

both over- and underfished and constant harvesters (pluses in Fig. 3a) were much more likely to underfish. Note that constant harvesters underfishing makes sense, because if they were to drastically overfish during this 21-turn game, the biomass would eventually decrease to the point where their constant harvest would begin to crash the fishery, at which point they would have to abandon the constant harvest strategy.

When harvest rules generated by the fitted Beverton-Holt recruitment model performed poorly, this was most often due to over rather than underfishing (Fig. 3b). Poor model performance was due to a combination of two reasons: (1) during the practice game the user allowed similar amounts of biomass to escape harvest on every turn, generating poor data for model fitting, and (2) the SD of the environmental noise was high (Fig. 4). When both of these conditions are true, the data can misrepresent the recruitment function (Fig. 4b compared to c) and lead to a poor escapement strategy. Despite the poor escapement strategies that sometimes resulted from the fitted models, they still were less frequent and generated higher long-term discounted profit than the worst users (compare the low points in Fig. 3a to b).

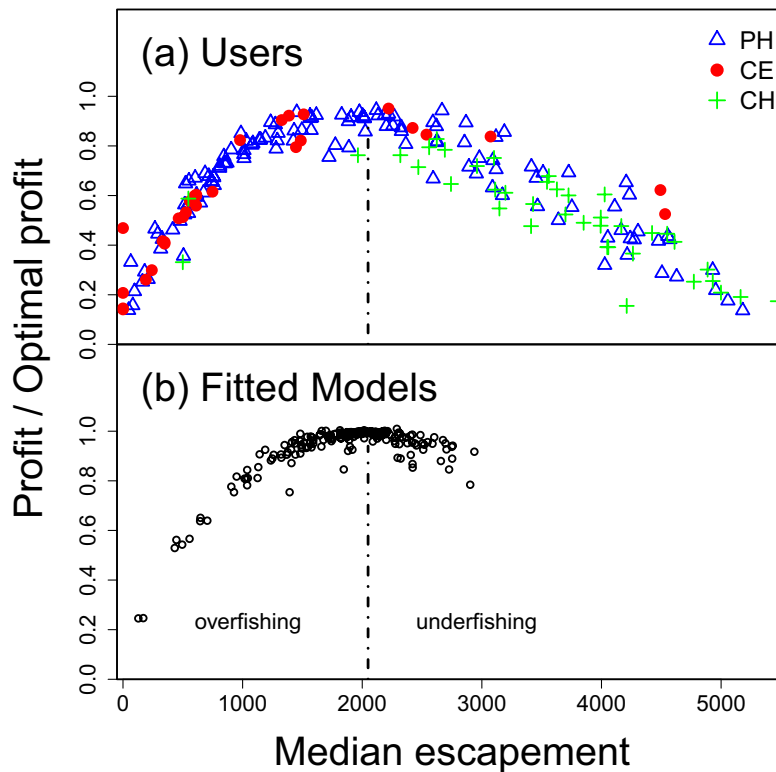


FIG. 3. (a) The profit generated by each user, in the unstructured game, relative to the net profit the optimal strategy with perfect information would generate in the corresponding instance of the game, as a function of the median amount of fish the user let escape harvest. The red circles, blue triangles, and green pluses are for users who used constant escapement (CE), proportional harvest (PH), and constant harvest (CH) strategies, respectively. (b) The analogous proportion of optimal profit generated by the fitted model vs. the median of escapements chosen by the model after it fit a recruitment function to the data during each turn of the game. The dotted line is optimal escapement under perfect information.

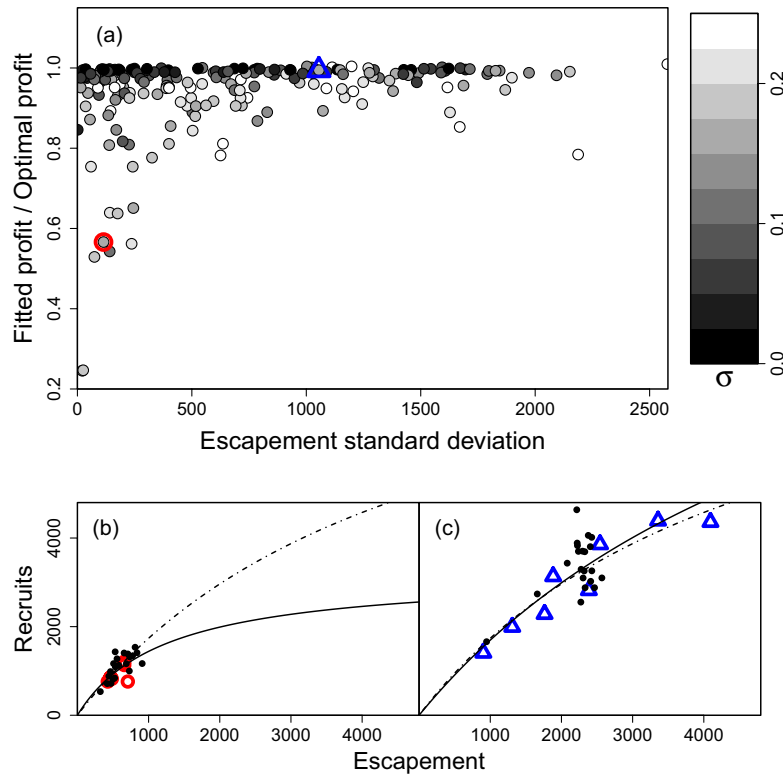


FIG. 4. (a) The profit generated from the strategies using the fitted Beverton-Holt model relative to the optimal profit under perfect information, as a function of the SD in practice game escapement, generated by the user. Dark and light circles are for instances of the games with low and high levels of environmental stochasticity respectively. (b–c) The true recruitment function (dashed line) and fitted recruitment function (solid line) for two instances of the game, (these examples are highlighted by a red circle and blue triangle in (a)), where the fitted model generates unprofitable escapement strategies (b; red circle) and profitable ones (c; blue triangle). The open symbols are recruitment data generated by the user in the practice game, whereas the smaller filled points are generated by the fitted model when playing the actual game.

In the Supporting Information we show how quickly the parameter estimates, during the Beverton-Holt model fitting, converged to the true values governing the game dynamics. In general, poor model fits were rather common. For example, when decisions were made based on the Beverton-Holt model with parameters estimated from biomass observations during the game, even after all 21 turns, the estimated value for b_2 was off by more than 25% over 40% of the time (see turn 21 in Fig. S8a). The estimation of b_1 was more accurate (Fig. S8b), and for both parameters the poor estimates improved as time moved forward (Fig. S8a, b) and as a result so did the model's harvest decisions (Fig 5a–c). Still, in many of the cases where parameter estimates were off, the fitted models led to more profitable decisions than those generated by the users, despite the poor model fit (Fig. 1b).

The users did not, on average, improve their escapement decisions through time during the 21 turn game. For example, as the users advanced through the game, the percentage of users who chose an escapement value within 50% of the optimal escapement value decreased (Fig. 5f). Changing the cut off from 50% of optimal escapement to 10% and 30% did not lead to a positive trend between the number of turns completed and the

percentage of users within the specified percentage of the optimal value (Fig. 5d–f). This suggests that while students had the opportunity to learn about the system, they did not do so in a way that improved their management decisions.

Age-structured population game

In the age-structured model, the average user achieved 63.6% of the optimal profit achieved by a model with perfect information. The most profitable user scored only 84.3% of the optimal profit, in comparison to the best performer in the unstructured population game who scored over 95% of optimal profit. On the opposite end of the spectrum, the worst users in the unstructured population game only scored 7.2% of optimal profit, while in the age-structured game the worst user scored 11.8% of optimal profit. A full distribution of the relative performance of the users compared to the optimal policy in the age-structured game is given in Fig. 6. The improved performance by the worst players, despite the age-structured game being more complex, was due to the fact that this game includes three independent cohorts. Even if one or two cohorts were driven to low

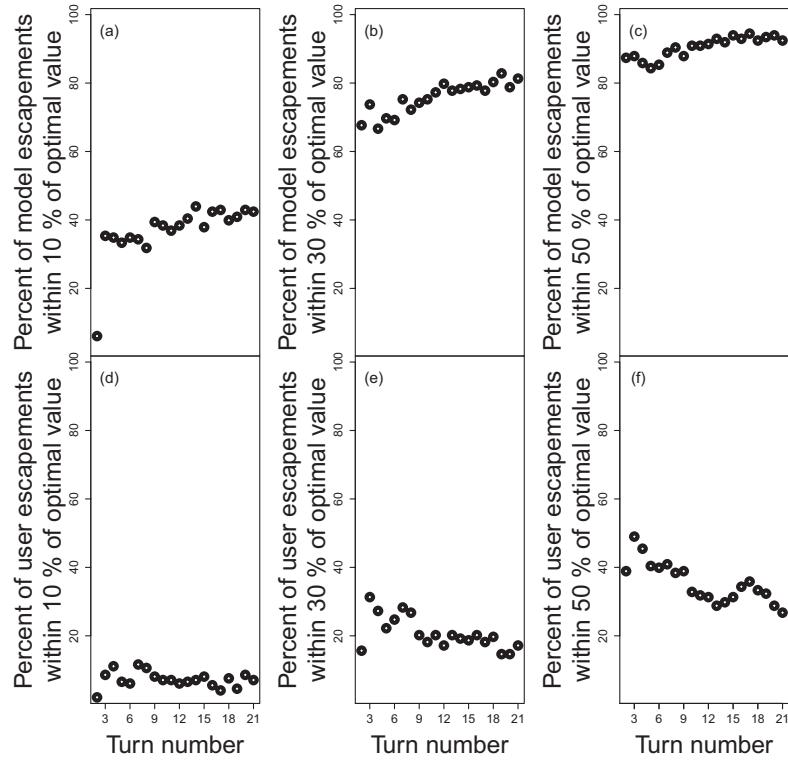


FIG. 5. Percentage of escapement decisions in the unstructured game, by the fitted Beverton-holt model, that were within (a) 10%, (b) 30%, and (c) 50% of optimal escapement as a function of the number of turns completed during the game. Percentage of users who chose escapement values within (d) 10%, (e) 30%, and (f) 50% of optimal escapement as a function of the number of turns completed during the game.

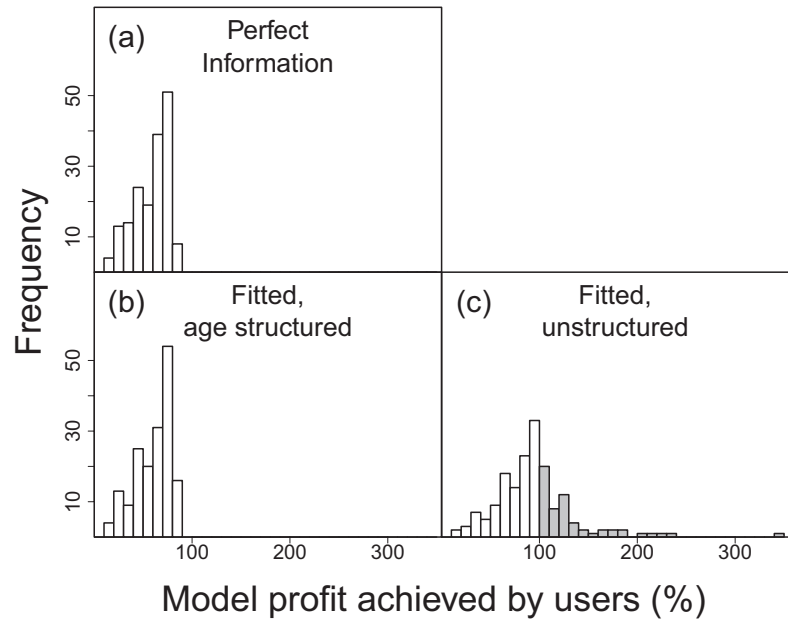


FIG. 6. The percentage of the mathematical model's total profit achieved by the user, in the age-structured game, when the mathematical model is (a) age-structured with true parameters, i.e., perfect information (b) age-structured but with parameters estimated from the data, and (c) unstructured with parameters estimated from the aggregated (immature + adult) biomass data. For example, a value of 50% means the user generated half the profit the mathematical model did, managing the exact same instance of the game. A value of 200% means the user generated twice as much profit as the model.

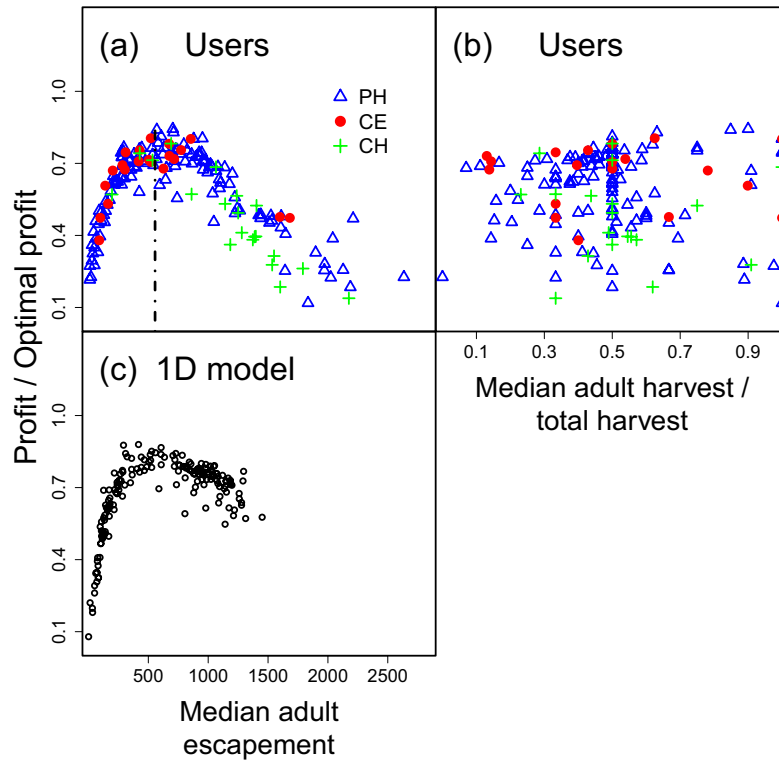


FIG. 7. (a–c) The proportion of optimal profit, in the age-structured game, generated by (a) the user and (c) fitted unstructured model, with parameters estimated from the aggregated (immature + adult) biomass data, for each game as a function of the median escapement chosen. (b) The proportion of optimal profit generated by the user as a function of the median proportion of harvest allocated to adult biomass during the game.

levels, some harvest could be achieved in the remaining turns as long as one cohort was still abundant. A player could make one very bad decision, and learn from it, without collapsing the entire fishery.

The user's performance in the age-structured game was mainly determined by their overall fishing pressure and not their decision of which age class to fish (compare Fig. 7a to b). The majority of users harvested more immature biomass than adult biomass, despite exclusive adult harvest being the optimal strategy (Fig. 7b). Similar to the simple unstructured game, users who deployed a constant escapement strategy (for adults) were more likely to over-fish (Fig. 7a).

The average escapement strategies resulting from the age-structure model with parameters estimated from historical data (fitted age structure model) achieved 98.4% of the optimal profit, even better than in the unstructured game. Even in the fitted model's lowest performing game, it achieved 78.9% of optimal discounted net profit, far better than even the median user. This is for two reasons (1) users in the practice game tended not to let the same amount of adult fish escape harvest every turn, producing good data for model fitting, and (2) the transition rate from immatures to adults was always estimated well, because it is a single parameter that can be estimated independently from recruitment, whereas the recruitment function requires two

parameters to be estimated simultaneously. The result of point (2) is that the model-based decisions always fished from the correct age class.

The average escapement strategy, generated by fitting a one-dimensional unstructured population model to the aggregate age-structured data, achieved 72.3% of the optimal profit. This represents a 13.5% gain in profit over the average student operating solely on intuition. Only 58 users, out of 172, generated more profit than would have been obtained by harvesting based on the fitted unstructured model. However, for instances of the game where the unstructured model generated low discounted net profit, the model's proposed escapement rule crashed the fishery, by letting very little biomass escape harvest. These strategies generated less discounted net profit than the least profitable users (compare the lowest points in Fig. 7a to c).

DISCUSSION

Many mathematical tools exist to improve decision making in environmental management. Yet some managers are still resistant to using these tools to develop management plans, and instead often rely on their experience and intuitive judgment (Johnson and Williams 2015). It has been shown that humans using intuition and judgment create more biased, but not necessarily

less accurate, predictions of species extinction risk when compared to humans aided by mathematical modeling (McCarthy *et al.* 2004). However, we are unaware of any study that compares human judgment and model-based strategies for developing explicit management decisions.

In this study we studied optimal escapement strategies for the management of simulated fisheries, developed using simplified models of fish stock dynamics, and tested their performance compared to students managing the simulated population using only their experience and judgment. The model-based decisions performed better than the students, on average, even when the models mis-specified recruitment or state variables. However, in the age-structured game, the worst outcomes produced by the simplified unstructured model were worse than the worst outcomes generated by the users.

Each quantitative model-based approach to managing the simulated fishery used a single equation or system of equations with unknown parameters. Alternatively, a manager could develop a set of candidate models representing alternative hypotheses about the system (e.g., an age-structured and unstructured model) and then use quantitative methods to resolve structural uncertainty (Williams 2001, Nichols *et al.* 2015). Our results on using unstructured models to manage age-structured populations are an example of a worst case scenario where a manager's candidate model set does not contain a model that approximates the system well. Even in such a case, on average, quantitative methods outperformed human judgment in our experiment.

Users and fitted models tended to make different types of mistakes. An equal number of users overfished vs. underfished the stock. However, when management based on fitted models failed, it was almost always due to overfishing.

We found no evidence to support the hypothesis that users learned about the system in such a way that helped them improve their management decisions (Fig. 5d–f). Because the subjects were students with little experience managing actual populations, it is unclear if experienced managers would perform similarly. However, the students did gain experience managing the simulated fishery through a practice game, and most of the subjects were exposed to curriculum in ecological management through the course in which the experiment was administered (concepts such as maximum sustainable yield and tragedy of the commons).

Learning did occur in model-based passive adaptive management, with many model runs quickly estimating the parameters reasonably well. However, even when the model was perfectly specified, and only needed parameter estimates from the data, it still sometimes performed worse than a student using intuition alone, especially when environmental stochasticity was high and prior management decisions were all similar (Fig. 4).

The lack of data with sufficient variability in stock abundance to estimate parameters well may be common

in fisheries management because overexploited populations are ubiquitous (Pauly 2008), and therefore the time-series data of recent fish stock abundances for many species may often contain only population sizes well below carrying capacity. In such cases, recruitment curves may often be incorrectly estimated such that simple models will naively suggest that it is optimal to keep overfishing. This suggests that passive adaptive management, choosing the best strategy based on the current knowledge of the system to optimize some objective, without any regard to the information gained by deploying that action, can potentially lead to poor performance even when model structure is correctly specified.

Our results suggest that probing the system by performing an action that is suboptimal given the manager's current belief about the system, but that will reveal information that improves management in the future, might be desirable in such scenarios. Incorporating the economic benefits of learning from experimentation explicitly into the optimal decision problem, known as active adaptive management, has been studied within the context of harvested populations. However, due to computational limitations solutions are always limited to cases with one of the three following assumptions: (1) both the probability distribution specifying environmental stochasticity and all parameters in the recruitment function are perfectly known, except for a single parameter to be estimated from the data (Walters 1981, Ludwig and Walters 1982), (2) there is a small number of candidate models, with all parameters fixed within each model (Williams 2001), or (3) only a small number of actions and system states are admissible (e.g., action = harvest or not, fishery state = robust, vulnerable, or collapsed; Hauser and Possingham 2008). Note, though, that the benefit gained from calculating an optimal solution using active adaptive management can be small depending on the problem (Walters and Green 1997).

Unfortunately, the problem of choosing an optimal escapement level in our game, using the principles of active adaptive management, is computationally infeasible given current algorithms and computing power because our game allows for an infinite set of possible actions and states, governed by multiple unknown parameters and unknown variability in environmental noise.

We used a version of passive adaptive management, as it was originally described by Walters (1986), to generate model-based decisions in this study. This method assumes the parameter values estimated from the data were the true values governing the dynamics during the optimization procedure. Alternatively, it is possible to incorporate parameter uncertainty in the optimization step of passive adaptive management (Johnson and Williams 2015). Correcting for uncertainty could potentially improve the model-based decisions in this study, without the need for more complicated active adaptive management methods.

It may be concerning that even in the most optimistic case, where the underlying dynamic model is known and parameters have to be estimated from the data, classic passive adaptive management can fail to achieve desirable

results. However, letting students manage our simulated fishery based solely on their experience and judgment typically led to much worse outcomes. Because using simple mathematical models along with the most basic passive adaptive management techniques usually improved management outcomes in our experiment, we would recommend that modeling be more widely adopted in management, even when challenges prevent the manager from using more complicated, state-of-the-art methods.

ACKNOWLEDGMENTS

The website and game development, used in this research, was funded by the National Science Foundation grant "Computational Sustainability: Computational Methods for a Sustainable Environment, Economy, and Society" (award number 0832782). We especially thank our web developer Mikhail Yahknis for building the fisheries game website and Jon Conrad, Jan Nyrop and two anonymous reviewers for helpful comments on the manuscript.

LITERATURE CITED

- Asano, E., L. J. Gross, S. Lenhart, and L. A. Real. 2008. Optimal control of vaccine distribution in a rabies metapopulation model. *Mathematical Biosciences and Engineering* 5:219–238.
- Bjørndal, T., and J. M. Conrad. 1987. The dynamics of an open access fishery. *Canadian Journal of Economics* 20(1):74–85.
- Game, E. T., E. Meijaard, D. Sheil, and E. McDonald-Madden. 2014. Conservation in a wicked complex world; challenges and solutions. *Conservation Letters* 7:271–277.
- Gerber, L. R., M. Beger, M. A. McCarthy, and H. P. Possingham. 2005. A theory for optimal monitoring of marine reserves. *Ecology Letters* 8:829–837.
- Hauser, C. E., and H. P. Possingham. 2008. Experimental or precautionary? Adaptive management over a range of time horizons. *Journal of Applied Ecology* 45:72–81.
- Helmstedt, K. J., H. P. Possingham, K. E. Brennan, J. R. Rhodes, and M. Bode. 2014. Cost-efficient fenced reserves for conservation: single large or two small? *Ecological Applications* 24:1780–1792.
- Holden, M. H., and J. M. Conrad. 2015. Optimal escapement in stage-structured fisheries with environmental stochasticity. *Mathematical Biosciences* 269:76–85.
- Holling, C. S. 1978. *Adaptive environmental assessment and management*. Wiley, London, UK.
- Hughes, M. J., E. G. Johnson, and P. R. Armsworth. 2014. Optimal spatial management of an invasive plant using a model with above-and below-ground components. *Biological Invasions* 16:1009–1020.
- Johnson, F. A., and B. K. Williams. 2015. A decision-analytic approach to adaptive resource management. Pages 61–84 in C. R. Allen, and A. S. Garmestani, editors. *Adaptive management of social-ecological systems*. Springer, Dordrecht, Netherlands.
- Johnson, F. A., D. R. Breininger, B. W. Duncan, J. D. Nichols, M. C. Runge, and B. K. Williams. 2011. A markov decision process for managing habitat for florida scrub-jays. *Journal of Fish and Wildlife Management* 2:234–246.
- Johnson, F. A., G. S. Boomer, B. K. Williams, J. D. Nichols, and D. J. Case. 2015. Multilevel learning in the adaptive management of waterfowl harvests: 20 years and counting. *Wildlife Society Bulletin* 39:9–19.
- Jorgensen, S. E., and B. Fath. 2008. *Encyclopedia of ecology, five-volume set: online version*. Elsevier, Amsterdam.
- Ludwig, D., and C. J. Walters. 1982. Optimal harvesting with imprecise parameter estimates. *Ecological Modelling* 14:273–292.
- Marr, J. C., et al. 1944. Age, length and weight studies of three species of Columbia River salmon (*Oncorhynchus keta*, *O. gorbuscha* and *O. kisutch*). Technical report. Department of Research, Fish Commission of the State of Oregon, Salem, Oregon, USA.
- Martin, J., P. L. Fackler, J. D. Nichols, M. C. Runge, C. L. McIntyre, B. L. Lubow, M. C. McCluskie, and J. A. Schmutz. 2011. An adaptive-management framework for optimal control of hiking near golden eagle nests in Denali National Park. *Conservation Biology* 25:316–323.
- MATLAB. 2010. version 7.10.0 (R2010a). The MathWorks Inc., Natick, Massachusetts, USA.
- May, R. M., et al. 1976. Simple mathematical models with very complicated dynamics. *Nature* 261:459–467.
- McCarthy, M. A., et al. 2004. Comparing predictions of extinction risk using models and subjective judgement. *Acta Oecologica* 26:67–74.
- McCarthy, M. A., H. P. Possingham, and A. Gill. 2001. Using stochastic dynamic programming to determine optimal fire management for banksia ornata. *Journal of Applied Ecology* 38:585–592.
- Moore, C. T., and M. J. Conroy. 2006. Optimal regeneration planning for old-growth forest: addressing scientific uncertainty in endangered species recovery through adaptive management. *Forest Science* 52:155–172.
- Nøstbakken, L., and T. Bjørndal. 2003. Supply functions for North Sea herring. *Marine Resource Economics* 18:345–361.
- Nichols, J. D., and B. K. Williams. 2006. Monitoring for conservation. *Trends in Ecology & Evolution* 21:668–673.
- Nichols, J. D., F. A. Johnson, B. K. Williams, and G. S. Boomer. 2015. On formally integrating science and policy: walking the walk. *Journal of Applied Ecology* 52(3):539–543.
- Pauly, D. 2008. Global fisheries; a brief review. *Journal of Biological Research-Thessaloniki* 9:3–9.
- Possingham, H., S. Andelman, B. Noon, S. Trombulak, and H. Pulliam. 2001. Making smart conservation decisions. Pages 225–244 in M. E. Soulé, and G. Orians, editors. *Conservation biology: research priorities for the next decade*. Island Press, Washington (DC).
- Probert, W. J., C. E. Hauser, E. McDonald-Madden, M. C. Runge, P. W. Baxter, and H. P. Possingham. 2011. Managing and learning with multiple models: objectives and optimization algorithms. *Biological Conservation* 144:1237–1245.
- R Core Team. 2012. *R: a language and environment for statistical computing*. R Foundation for Statistical Computing, Vienna, Austria.
- Reed, W. J. 1979. Optimal escapement levels in stochastic and deterministic harvesting models. *Journal of Environmental Economics and Management* 6:350–363.
- Ricker, W. E. 1954. Stock and recruitment. *Journal of the Fisheries Board of Canada* 11:559–623.
- Rout, T. M., J. L. Moore, and M. A. McCarthy. 2014. Prevent, search or destroy? A partially observable model for invasive species management. *Journal of Applied Ecology* 51:804–813.
- Sainsbury, K. 1988. The ecological basis of multispecies fisheries, and management of a demersal fishery in tropical Australia. *Fish Population Dynamics* 2:349–382.
- Stankey, G. H., R. N. Clark, and B. T. Bormann. 2005. *Adaptive management of natural resources: theory, concepts, and management institutions*, Vol. 654. US Department of Agriculture, Forest Service, Pacific Northwest Research Station.
- Walters, C. J. 1981. Optimum escapements in the face of alternative recruitment hypotheses. *Canadian Journal of Fisheries and Aquatic Sciences* 38:678–689.

- Walters, C. 1986. Adaptive management of renewable resources. Macmillan Pub. Co., New York, New York, USA.
- Walters, C. J. 2007. Is adaptive management helping to solve fisheries problems? *AMBIO: A Journal of the Human Environment* 36:304–307.
- Walters, C. J., and R. Green. 1997. Valuation of experimental management options for ecological systems. *The Journal of Wildlife Management* 61(4):987–1006.
- Welling, L., and L. Thomson. 2003. PHP and MySQL web development. Sams Publishing, Indianapolis, Indiana.
- Westphal, M. I., M. Pickett, W. M. Getz, and H. P. Possingham. 2003. The use of stochastic dynamic programming in optimal landscape reconstruction for metapopulations. *Ecological Applications* 13:543–555.
- Williams, B. K. 2001. Uncertainty, learning, and the optimal management of wildlife. *Environmental and Ecological Statistics* 8:269–288.
- Williams, B. K., and E. D. Brown. 2012. Adaptive management: the US Department of the Interior Applications Guide. US Department of the Interior, Adaptive Management Working Group.
- Wilson, K. A., M. F. McBride, M. Bode, and H. P. Possingham. 2006. Prioritizing global conservation efforts. *Nature* 440: 337–340.
- Worden, L., L. W. Botsford, A. Hastings, and M. D. Holland. 2010. Frequency responses of age-structured populations: Pacific salmon as an example. *Theoretical Population Biology* 78:239–249.

SUPPORTING INFORMATION

Additional supporting information may be found in the online version of this article at <http://onlinelibrary.wiley.com/doi/10.1890/15-1295/supinfo>