

To appear in *Managing Invasive Species*,

Terry Walshe and Mark A. Burgman, eds., Cambridge University Press

Management of Invasive Species: Info-Gap Perspectives¹

Yakov Ben-Haim²

Contents

1	Introduction	1
2	Culling and Learning: An Allocation Problem	3
2.1	Problem Statement	3
2.2	Uncertainty and Robustness	3
2.3	Interpreting Robustness Curves	4
2.4	Selecting Between Management Strategies	6
3	Variable Culling Effort	7
3.1	Formulation	7
3.2	Results	8
4	Conflicting Experts	10
4.1	Two Models	11
4.2	Robustness	12
4.3	Strategy Selection	13
5	Conclusion	14
6	References	15

Abstract. Invasive species may significantly disrupt existing eco-systems, thus impairing the validity of established scientific models. This does not invalidate the use of these models in the science-based selection of management strategies, but it does require careful attention to the unknown errors of the models. We use info-gap decision theory to illustrate the treatment of model uncertainty in strategy selection for managing invasive species.

1 Introduction

Invasive species instigate difficult challenges to the scientific basis of environmental management. Science is an inductive procedure. By studying existing systems we come to understand the governing processes and thereby learn how to manage them. Invasive species may disrupt existing processes and thus reduce the relevance of established scientific models.

Scientific models are extremely useful in evaluating and selecting strategies for managing invasive species. However, because the actual ecological system may differ in unknown ways from the system for which a model was built, it is necessary to use the model cautiously. In particular, model-based predictions of optimal outcomes may be unrealistic due to mis-match between the mathematical model and the ecological system. This means that evaluating a strategy in terms of its model-based predicted outcome, and choosing a strategy whose predicted outcome is best, may be unrealistic. A management strategy which maintains that ‘This is the best model we have’ (which is probably

⁰\papers\Adaptive-Erad2010\mis06.tex 7.9.2010 © Yakov Ben-Haim 2010.

¹This paper was written in part while the author was the Alan Richards Fellow in Grey College, Durham University, Durham U.K. The author expresses his gratitude for the support of the College. The author also gratefully acknowledges the support of the U.K. Royal Academy of Engineering.

²Yitzhak Moda'i Chair in Technology and Economics, Technion—Israel Institute of Technology, Haifa, Israel, yakov@technion.ac.il

true), ‘so we should use this model to identify and strive for the best possible outcome’, is wishful thinking. The best possible outcome is unknown since the model errs in unknown ways. We simply cannot reliably identify a strategy whose outcome will be better than any other strategy.

Gigerenzer and Selton (2001) take a very different approach, by attempting to achieve adequate—rather than optimal—outcomes by frugally employing decision resources. This approach employs the concept of *satisficing* introduced by Simon (1955, 1956) and motivated by the limited information and information-processing ability of the decision maker, which Simon referred to as *bounded rationality*. To satisfice has come to mean “To decide on and pursue a course of action that will satisfy the minimum requirements necessary to achieve a particular goal.” (Oxford English Dictionary, 1989). Satisficing makes very good sense in situations of severe uncertainty such as managing invasive species. When one’s understanding is known to be very deficient, it is bad policy to exploit that understanding too extensively. Bad models should be used only as much as is necessary, and as little as possible, even when they are the best models around. That balancing act—as much as necessary and as little as possible—is what satisficing strategies attempt to do.

The goal of this paper is to provide a conceptual framework for selecting management strategies which satisfice, rather than optimize, the outcome. We demonstrate that best-model optimization is actually a special limiting case of satisficing. We show how the decision maker can choose a strategy to achieve maximal confidence of adequate outcome. We will see when best-model optimization is the strategy of choice, and when it is not. Our work is based on info-gap decision theory (Ben-Haim, 2006).

We discuss simple examples of the management of invasive species. The aim not to advocate specific models, or to develop real-life policy recommendations. Rather, the goal is to illustrate the info-gap decision methodology so the reader can understand how an info-gap robustness analysis is used in the selection of a management strategy, and what insights it provides. No attempt is made to explain the mathematical details of how the info-gap analysis is actually performed, which is described extensively elsewhere (Regan *et al.*, 2005; Burgman, 2005; Ben-Haim, 2006; <http://info-gap.com>).

We will consider three examples of progressive complexity. In section 2 the manager must allocate a limited budget between culling the invasive species and studying the disturbed ecology. The purpose of this example—which is static, simplistic and stylized—is to introduce the central ideas of info-gap uncertainty and robustness. In section 3 we consider a dynamic model and illustrate the ways in which the info-gap robustness analysis supports the choice of a time-varying culling effort. We extend this example in section 4 to include both uncertainty in the parameters of the ecological model as well as conflicting expert opinion on the structure of the model itself.

Our examples employ various versions of the Lotka-Volterra model of population dynamics. These are widely used models, and will provide ample opportunity for studying the implications of model uncertainty for the selection of management strategies. The standard Lotka-Volterra model represents interactions between species by quadratic terms—products of population sizes—with constant coefficients (Krebs, 1978). Pielou (1969, p.74) comments, however, that these models, while useful, “are certainly too simple to be realistic”. The standard Lotka-Volterra model ignores spatial and temporal non-uniformity of environmental conditions, idiosyncracies of individual species, spatial clumping and temporal variation of density, time lags of environmental impacts, immigration and emigration, and so on. As Ramsey and Veltman (2005) point out, the coefficients of the interaction terms are usually evaluated by linearization around an equilibrium, so their relevance is limited to small perturbations around that steady state. Consequently the effect of large disruptions cannot be predicted. Furthermore, as shown by Courchamp, Langlais and Sugihara (1999), very simple systems can have many equilibrium points, so a Lotka-Volterra model linearized around one will not be suitable for studying other dynamic domains. Not surprisingly, extensions of the basic Lotka-Volterra model have been explored by many authors (Royama, 1992; Kot, 2001).

2 Culling and Learning: An Allocation Problem

2.1 Problem Statement

Consider an invasive cat preying on an indigenous bird species. Our aim is to keep the bird population viable. We can exert effort to control the cat population, but we are unable to eradicate all cats; it is simply too hard to find them all. Furthermore, our understanding of the population dynamics is limited. How many cats should we cull in order to confidently maintain a healthy bird population? Should we allocate some of our management budget to study the population dynamics so that we can choose a more effective and reliable culling strategy? We will consider a simplified static example to begin to understand the info-gap approach.

We have access to a team of analysts who have built a Lotka-Volterra model for the two-species interaction between cats and birds (Krebs, 1978; Pielou, 1969; Royama, 1992). Regarding birds, the model states that the bird population would grow exponentially if they were not predated by cats. Regarding cats, the model states that the cat population would decay exponentially if they were not able to consume birds.

This model predicts two possible stable states for the bird-cat interaction. Either both species become extinct, or both species stabilize at population sizes which depend on the properties of the dynamic interactions. The equilibrium size of the bird population, B_{eq} , as it interacts with the cat population, is:

$$B_{\text{eq}} = \frac{\gamma_1}{\gamma_2} \quad (1)$$

The equilibrium size of the bird population is the ratio of the rate at which the cat population would decline in the absence of birds, γ_1 , to the rate at which the cat population grows as a result of each bird-cat encounter, γ_2 . (An analogous expression describes the equilibrium size of the cat population but it needn't concern us here.)

We can influence the size of the equilibrium bird population by management actions such as investing effort in culling cats or investing resources in studying the population dynamics so that culling efforts can be more effectively deployed. Our management goal is to keep the equilibrium bird population viable. We would like to stabilize with at least 200 individuals, which is the minimal acceptable equilibrium population size for birds, denoted B_{min} . If the equilibrium bird population is no less than B_{min} , then the management activity has succeeded. Since we are uncertain if the model is correct, we would like to know if it is feasible to equilibrate at no less than 200 birds, and what management strategy is needed to reliably achieve this.

2.2 Uncertainty and Robustness

Our analysts should be able to help answer these questions. However, their quantitative model is uncertain. Our analysts are using an adaptation of the best available model developed in similar situations, but the specifics of our situation are unique. This particular invasive species (the cat) has not been studied before in this habitat with this particular prey (the bird). And even if it has, we suspect that the habitat is changing, perhaps due to urbanization, or climate change, or other factors.

Our analysts of course are fully aware of the limitations of their model. They are able to evaluate the robustness-to-uncertainty of any proposed management strategy. By “strategy” they mean the fraction of a fixed budget which is allocated to the culling of cats, while the remaining budget is devoted to studying the eco-system and improving the population dynamics model. By the “robustness” of a strategy they mean how good the strategy will be, even if the model is wrong. The concept of robustness can be expressed more precisely in two different—but complementary—ways.

One way to understand the idea of the robustness of a strategy is as an answer to the following question: if we want to achieve an acceptable outcome (yielding a bird population no less than B_{min}), then how large an error in the estimated species-interaction model can we tolerate? Robustness is a

measure of our immunity to model-error while achieving a specified required outcome. The strategy is robust if acceptable bird population size is maintained even at large modelling error.

If the robustness of a particular strategy is large, then an acceptable outcome will be achieved even if the model is very wrong. That is, large robustness means that the bird population at equilibrium will not fall below the specified minimal value even if the real interaction dynamics are greatly different from the mathematical model. On the other hand, small robustness of a strategy means that the outcome requirements are confidently achieved only if the habitat dynamics are quite close to our analysts' model. We will tend to prefer a strategy which is more robust over a strategy which is less robust.

The other way to understand the robustness of a strategy is as an answer to the following question: how small could the bird population be if our model contains errors up to a specified magnitude? This turns things around. Rather than specifying the minimal acceptable population size, we specify the magnitude of error of the model and ask: how much could the bird population fall, given this strategy, if our model errs by this amount. The strategy is robust if bird-loss is small even with large modelling error.

A strategy with large robustness—a strategy for which, even at large modelling error, only small bird-loss results—is preferable over a strategy with small robustness. Whichever of the two interpretations of robustness which we use, the robust preference between strategies will be the same.

The robustness can also be used to evaluate the feasibility of different choices of the minimal acceptable population size, B_{\min} . For a given strategy, we can ask: what value of B_{\min} has large (or small) robustness? Values of B_{\min} that, for a given strategy, have large robustness are more feasible (with that strategy) than values of B_{\min} that have low robustness.

Let's illustrate these ideas with some graphs.

2.3 Interpreting Robustness Curves

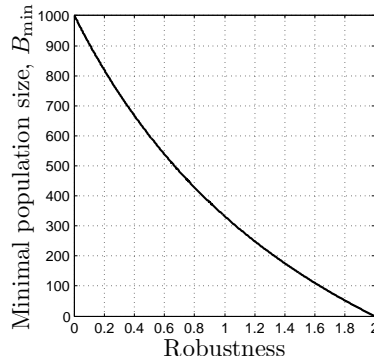


Figure 1: Robustness vs minimal bird population size, B_{\min} , for strategy S_1 .

Our analysts display the robustness of a specified strategy in terms of its **robustness curve**, as shown in fig. 1. The vertical axis is the minimal acceptable population size, while the horizontal axis is the robustness. This curve is for a specific strategy: an allocation between culling cats and studying the eco-system.

The interpretation of the scale on the robustness axis depends on how the uncertainty is modelled. Our analysts have estimated the values of the coefficients in the equation for the equilibrium size of the bird population, γ_1 and γ_2 in eq.(1), and they have estimated the errors of these estimated coefficients (e.g. standard errors, or spread of expert opinion, etc). Uncertainty is expressed as an unknown fractional error of each coefficient with respect to its estimated error. Robustness is the greatest fractional error of the estimated parameters of the model up to which all models yield acceptable outcome. Thus, for instance, a robustness of 0.5 means that any magnitude of error

in all coefficients of the model, up to 50% error, does not jeopardize the corresponding minimum bird-population size requirement, which is 600 birds for this value of robustness, as we see in fig. 1. Similarly, a robustness of 1.4 means that γ_1 and γ_2 can each err by as much as 140% (each coefficient with respect to its estimated error) if the minimal requirement is for 176 birds.

Three features of all robustness curves appear in fig. 1: trade-off, cost of robustness, and zeroing.

The negative slope of the robustness curve in fig. 1 represents the **trade off** between robustness and performance: better performance (larger minimal population size) is obtained in exchange for lower (worse) robustness to uncertainty. The negative slope in fig. 1 quantifies the ordinary intuition that higher aspirations (for large values of B_{\min}) entail greater vulnerability (lower robustness) than lower aspirations. The negative slope enables us to evaluate how much immunity against model-error must be foregone in exchange for an increase in our performance requirement. For instance, we see from fig. 1 that requiring an equilibrium population size of birds no less than 800 has a robustness of 0.22: the bird population will not fall below 800 provided no model coefficient errs by more than 22%. Allowing the equilibrium size to be as small as 400 birds entails a robustness of 0.86.

The trade off property is manifested in the slope of the robustness curve which can be understood as a **cost of robustness**. A shallow slope means that a unit decrease in the minimal population size entails a large increase in robustness. A steep slope means that a unit decrease in population size results in only a small gain in robustness. For instance, we see in fig. 1 that decreasing B_{\min} from 300 to 100 birds results in an increase in robustness of about 0.5. In contrast, a decrease from 1000 to 800 results in an increase in robustness of only 0.2. The cost of robustness is lower at the lower right than at the upper left of the curve in fig. 1. Shallow slope means low cost of robustness, while steep slope means that robustness is very costly in units of decreased performance.

The trade-off property has a further important implication for strategy selection. Recall that a robustness curve is evaluated for a specific management strategy; fig. 1 is for a strategy labeled S_1 . Using this strategy, the analysts tell us that the estimated model predicts an equilibrium population of 1000 birds. As we see in fig. 1, this is precisely the value of B_{\min} at which the robustness equals zero. This illustrates a property of all robustness curves: the robustness of a strategy is zero when attempting to achieve the outcome which is *predicted* for that strategy. Best-model predictions have zero robustness. This is called the **zeroing** property.

This means that we should not evaluate a strategy in terms of the predicted outcome of that strategy, since this prediction is an unreliable indication of how the strategy will fare. Rather, we should evaluate a strategy by its full robustness curve. Using the concept of trade off between robustness and performance, we can evaluate the robustness of poorer-than-predicted outcomes; these are the only outcomes which have positive robustness against modelling error.

A further implication has to do with best-model optimization. A common approach to selecting a management strategy is to seek the best outcome which can be achieved, based on our best model (why settle for less?). The problem with this approach is that best-model predictions have no robustness against uncertainty, so best-model optimization can be wishful thinking as we will see.

Each strategy has its own robustness curve that reflects the trade off and zeroing properties of that strategy. Why are some curves steep and some curves shallow? What does this indicate about the corresponding strategies? A strategy has large predicted equilibrium bird population (hence the robustness becomes zero at a large value of B_{\min}) if that strategy effectively exploits the properties of the system as represented by the estimated model of the population dynamics. A strategy will have low cost of robustness (and hence a shallow robustness curve) if that strategy is insensitive to the gap between the estimated and the correct model. These two properties of a strategy—predicted outcome and robustness against error in the predictive model—together determine the slope and intercept of the robustness curve.

The method for selecting a management strategy which is indicated by these consideration is called **robust-satisficing**. Identify an outcome which you can accept, and seek a management strategy which maintains at least an acceptable outcome over the widest possible range of deviation of reality from your model. That is, satisfice (rather than optimize) the outcome and robustify against modelling error. This is illustrated in our next graph.

2.4 Selecting Between Management Strategies




Figure 2: Robustness vs minimal bird population size, B_{\min} , for strategies S_1 and S_2 .

Fig. 2 shows robustness curves for two different management strategies, e.g. different allocations between culling cats and studying the population dynamics. Strategy S_1 entails more culling of cats and less modelling effort than strategy S_2 . The best-model prediction for S_1 is that the equilibrium bird population will have 1000 individuals while S_2 predicts only 500. On the other hand, since S_2 is based on a more thorough study of the bird-cat interaction, the cost of robustness is lower with S_2 than with S_1 . That is, the robustness curve for S_2 is steeper than for S_1 .

Consider strategy S_1 . The predicted equilibrium size of the bird population, based on S_1 , is 1000 individuals. However, the robustness to uncertainty in this value is zero, so we cannot rely on stabilizing at 1000 birds if we implement strategy S_1 . Only smaller populations have positive robustness against modelling error. We are aiming to maintain at least 200 individuals, and we see from the solid curve in fig. 2 that the robustness for this requirement is 1.3; the coefficients of the Lotka-Volterra model can all err up to 130% without jeopardizing the requirement that the final population not fall below 200 individuals.

Now consider strategy S_2 , whose robustness curve is the dashed line in fig. 2. The predicted outcome of S_2 is only 500 birds (since we are culling fewer cats than with S_1). Nominally—based on the best-model predictions—we would prefer S_1 over S_2 (since 1000 birds is better than 500). However, since predicted outcomes have zero robustness, this is not a good basis for comparing these strategies. Rather, we see that the robustness for achieving a population no smaller than 200 birds is 2.1 with strategy S_2 (as opposed to 1.3 with S_1). Using S_2 (with its larger share of ecosystem research), the model coefficients can all vary up to 210% without allowing the equilibrium bird population to fall below 200. Strategy S_2 is more robust to uncertainty than strategy S_1 ; the robust preference—given the requirement for at least 200 birds—is for S_2 rather than S_1 .

What has happened is a reversal of preference between these strategies, resulting from the fact that their robustness curves cross one another. Let B_{\times} denote the minimal population size, B_{\min} , at which the robustness curves in fig. 2 cross one another. This is about 333 birds, occurring at robustness of 1.0. If we require that the equilibrium population size include more than 333 birds, then strategy S_1 is more robust than, and hence preferable over, strategy S_2 . However we must recognize that, depending on the value of B_{\min} , the robustness might be rather low. On the other hand, if we are willing to accept an equilibrium population with less than 333 birds then strategy S_2 is more robust than S_1 , so we would prefer S_2 . Now the robustness might be large, as in the case of requiring at least 200 birds.

Let's summarize the robust-satisficing approach to strategy selection. Strategy S_1 is nominally preferred over S_2 : the best-model prediction is better with S_1 than with S_2 . However, best-model predictions have zero robustness. Furthermore, the robustness curves of these two strategies cross one another. If an outcome below the crossing point is acceptable (B_{\min} less than B_{\times}) then S_2 is

more robust and hence preferable over S_1 . S_1 is preferred if we need a population larger than B_x .




Figure 3: Robustness vs minimal bird population size, B_{\min} , for strategies S_1 , S_2 and 2 others.

There are many management strategies which are intermediate between S_1 and S_2 , all with the same total budget. The robustness curves of two such intermediates are shown as thin curves in fig. 3, together with the curves for S_1 and S_2 . The four strategies all show the trade off and zeroing properties, as well as extensive curve crossing. Most notably, however, we see that—to a good approximation—the robustness curves for S_1 and S_2 enclose the intermediate strategies. This means that either S_1 or S_2 is more robust than the other strategies which thus need not be considered.

3 Variable Culling Effort

We now extend the example of section 2 to include variation of the culling effort over time. Our aim, as before, is to maintain a healthy bird population threatened by an invasive cat. We have a fixed budget which determines the total culling effort which can be exerted over a specified duration. However, when the cat population is small, a large effort is needed to remove additional cats, while culling from a large population requires less effort. The question is how to allocate culling effort over time in order to reliably maintain a viable bird population in light of uncertainty in the population dynamics.

3.1 Formulation

The analysts estimate that the amount of effort, E_t , required to remove n_t cats (at any time step t) depends on the current size of the cat population, C_t , as:

$$E_t = \frac{\varepsilon}{C_t - n_t} - \frac{\varepsilon}{C_t} \quad (2)$$

where ε is a constant. Removing no cats ($n_t = 0$) requires no effort, but as n_t approaches the size of the cat population (that is, as we attempt to remove all but the last few cats) the required effort increases without bound. This equation can be inverted to express the number of cats culled as a function of the effort and the size of the cat population. We denote this as $n_t(E_t, C_t)$. (Strictly speaking, this is the greatest integer less than the value of n_t which satisfies eq.(2)). This is a ‘feedback’ strategy since the size of the cat population feeds back to influence the culling action.

The discrete-time Lotka-Volterra equations which the analysts use are:

$$B_{t+1} = (1 + \gamma_3)B_t - \gamma_4 B_t C_t \quad (3)$$

$$C_{t+1} = (1 - \gamma_1)C_t + \gamma_2 B_t C_t - n_t(E_t, C_t) \quad (4)$$

The coefficients γ_1 and γ_2 are the same as in eq.(1). γ_3 is the fractional growth rate of birds in the absence of predation, and γ_4 is the death rate of birds per bird-cat encounter. Eq.(4) contains the cat-culling term, $n_t(E_t, C_t)$, which depends on the size of the cat population and on the culling effort which is allocated for the current step.

Our management goal is to maintain the size of the bird population above a specified minimum value, B_{\min} , throughout a management duration. That is, we require that $B_t \geq B_{\min}$ for each time step t from $t = 1$ to $t = T$. Our budget determines the total culling effort which can be exerted, E , and our task is to allocate this effort over time subject to the budget constraint, $E = \sum_{t=1}^T E_t$.

We will consider three different strategies for allocating the culling effort. The effort, E_t , can be constant over time, or it can increase linearly from zero to maximal effort at time T , or it can decrease linearly to zero effort at time T . While there are other time varying strategies, these three will suffice to illustrate the decision methodology. We will also consider the impact of different total budget sizes, E .

The robustness of a strategy can be understood exactly as in section 2, where we discussed two equivalent definitions. The robustness of a specified strategy (allocation of effort over time) is the greatest uncertainty in the parameters of the model which can be tolerated while also maintaining the bird population above the minimal requirement, B_{\min} , throughout the management duration. Equivalently, the robustness of a strategy can be understood as the smallest size which the bird population could attain, given a specified level of error in the model. In this latter definition, robust strategies assure large bird populations even with large model errors.

3.2 Results

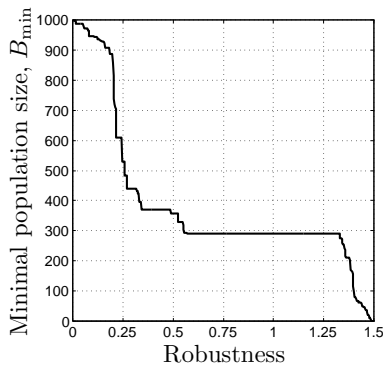


Figure 4: Robustness vs minimal bird population size, B_{\min} , for constant culling effort.

Fig. 4 shows a robustness curve for constant culling effort over a duration in which both the bird and the cat populations display repeated fluctuations. The number of cats which are culled varies up and down over time as the size of the cat population changes, in accord with eq.(2), even though the effort E_t is constant.

We see the same three generic features of the robustness curves which were observed in section 2. First, robustness trades off against performance: requiring that the bird population not fall below a large minimal size entails low robustness against modelling error. Second, the robustness can be significantly increased without substantially reducing the population minimum only if the slope is shallow. Third, the robustness equals zero at the population size which is predicted by the estimated model for this strategy.

The striking difference between the robustness curve in fig. 4 and those in section 2 is that here the robustness is jagged and piece-wise flat. The step-like shape in fig. 4 results from the discrete culling process in which small numbers of cats are culled, typically between 0 and 10 per time step. Thus for instance the long plateau in the robustness curve at a value of B_{\min} about 300 results

from both the culling, and the extremes of the fluctuating bird population, staying the same for the corresponding wide range of different models.

This long plateau is quite significant since it shifts the lower right part of the robustness curve to substantially larger robustness values. For instance, we see from fig. 4 that the robustness is 1.38 if our goal is to maintain at least 200 birds throughout the management period. This means that fractional errors in all of the model parameters of 138% (each with respect to its own estimated error) do not jeopardize this requirement. In contrast, the robustness is only 0.55 if we require a B_{\min} of 300 birds.

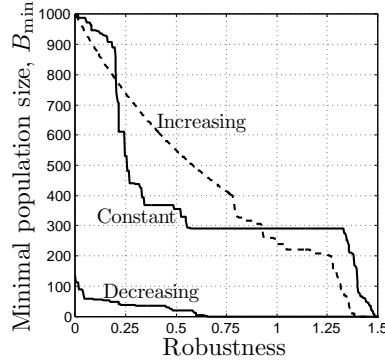


Figure 5: Robustness vs minimal bird population size, B_{\min} , for 3 different culling effort strategies.

Fig. 5 shows robustness curves for three different culling strategies, all with the same total effort. The curve with constant culling effort is reproduced from fig. 4. The dashed curve which crosses the constant-culling curve has linearly increasing effort starting with zero effort in the first time step. The robustness curve in the lower left corner has linearly decreasing culling effort reaching zero effort in the last time step. The numbers of cats which are culled in each step are different in each case, and fluctuate as the cat population fluctuates due to the feedback from population size to actual culling.

The first thing to conclude from fig. 5 is that the decreasing-effort strategy is vastly more vulnerable to error in the model than either the constant-effort or the increasing-effort strategies. From among these three options (at the total budget considered) one would never choose the decreasing effort strategy.

The choice between constant and increasing effort is more subtle, since their robustness curves cross one another, which raises the possibility of preference-reversal between them, as discussed in section 2. Nominally, these strategies are almost the same: their predicted minimal populations sizes are nearly identical but the robustnesses of these predictions are zero. Thus their nominal performance is not a reason for indifference between them. Indeed, their robustness curves are substantially different over some ranges of B_{\min} . For instance, for a minimal acceptable bird population size of 500, the robustness of the constant-culling strategy is 0.26 while the robustness of the linearly increasing strategy is 0.58, more than twice as robust. On the other hand, at B_{\min} of 200 the preferences are reversed though the robustness advantage of the constant-effort strategy is not as dramatic.

Fig. 6 shows robustness curves for constant-effort culling strategies at three different total budgets of culling effort. The over all trend is clear and unsurprising: more total effort is usually more robust. A total budget of $E = 2.0$ is strictly more robust than a total budget of $E = 1.5$, though the robustness advantage of the greater effort is not constant over the range of B_{\min} values. A budget of $E = 2.0$ is more robust than $E = 1.0$ over most, but not all, of the B_{\min} values.

Fig. 7 shows a close-up of fig. 6 in the high-performance low-robustness corner of the graph. Here we see extensive curve-crossing, with the associated potential for reversal of preferences among these management options. In particular, for B_{\min} between 800 and 900 birds, the *lower* effort has *twice*




Figure 6: Robustness vs minimal bird population size, B_{\min} , for constant culling effort at 3 different total budgets.




Figure 7: Close-up of fig. 6.

the robustness of the higher effort, when comparing $E = 1.0$ against $E = 2.0$. If a very large bird population is required—and recognizing that the associated robustness is relatively low—then one might prefer the lower total effort (from among these constant-effort strategies).

This somewhat counter-intuitive conclusion is a result of the non-linear population dynamics and the feedback between population size and actual culling. Culling too many cats can lead to a rise in the bird population which can then cause a sharp rise in the number of cats which in turn may drive the bird population below the acceptable level. We can't predict this at all precisely since the model is uncertain, but the robustness analysis reveals this mechanism and steers us away from being too trigger happy.

These examples illustrate how the choice of a dynamic state-feedback culling strategy is supported by analysis of robustness to modelling uncertainty. The properties of robustness curves—trade off, cost of robustness, and zeroing—reveal important and sometimes surprising properties of available strategies.

4 Conflicting Experts

Experts sometimes disagree among themselves. The need for two or more opinions arises in many fields. The Intergovernmental Panel on Climate Change presents suites of advanced science-based models predicting various global climatological processes (IPCC-AR4, 2007). These models, and their predictions, often differ substantially. Likewise, macro-economic processes are extraordinarily complex and can be modelled in many conceptually distinct ways. A host of macro-economic models have each received serious support by distinguished scholars. Nonetheless, the axiomatic, conceptual and predictive disparity among these models is substantial (Snowdon, Vane and Wyncarczyk, 1994).

Ecological models—and experts—are no exception. Expert opinions on matters of fact—how large a population can this habitat sustain, how low a rainfall can this species tolerate, and so on—can vary greatly. Expert's opinions depend on how the questions are formulated, how the elicitation interview is structured, and many other factors. Experts tend to be over-confident in assessing statistical confidence intervals, which can be ameliorated only by careful design of the elicitation procedure (Speirs-Bridge *et al.*, 2010). Opinions can differ on the conceptual foundations and mathematical structure of predictive models for ecological processes. Several distinct paradigms are in use for describing spatial and temporal population dynamics, including the Lotka-Volterra model, time series modelling (Turchin and Taylor, 1992), meta-population modelling (Hanski, 1999) and others. Even within the class of Lotka-Volterra models, specific realizations vary widely as mentioned in section 1.

In this section we study the choice of a time-varying culling strategy when our team of analysts is divided in two groups, each supporting a dramatically different model of the population dynamics.

Both models are uncertain, their predictions are different, and we must choose a culling strategy before the scientists are able to work out their differences. We have no reason to believe, or to disbelieve, one model more than the other. We are not able to “bet” on either model. Both models cannot be true, and in fact both are uncertain so both are likely to be false. And yet each model is founded on sound evidence and supported by credible scientific opinion. It is not due to negligence or incompetence that the scientists disagree among themselves. The specific ecological situation which must be managed is unique in important ways. We may hope and believe that the scientists, given time and resources, would work out their differences. But we must choose a management strategy for the current situation now. In this section we illustrate an info-gap robustness approach to strategy selection in this predicament.

4.1 Two Models

The two models which are used by our team of analysts are both finite-difference Lotka-Volterra models, but their structures are different and some of each model’s parameters are fractionally uncertain as described in section 2.3

The **first model** is specified in eqs.(3) and (4) with the culling quantity determined by eq.(2) in terms of the culling effort. That is, the culling strategy determines the culling effort at each time step, and the actual number of cats which are removed with this effort depends on the size of the cat population.

The **second model** is also a finite difference Lotka-Volterra model with a different structure (this model is motivated by Krebs, 1978, p.242):

$$B_{t+1} = \left(1 - \frac{B_t}{\zeta_2}\right) \zeta_3 B_t - \zeta_4 B_t C_t \quad (5)$$

$$C_{t+1} = \zeta_1 B_t C_t - n_t(E_t, C_t) \quad (6)$$

The coefficients, ζ_i , are all positive. ζ_2 is the carry capacity of birds in the habitat in the absence of predation by cats. ζ_3 is the fractional growth rate of birds in the absence of predation and far below the carrying capacity. ζ_4 is the death rate of birds per bird-cat encounter. ζ_1 is the growth rate of cats per bird-cat encounter. $n_t(E_t, C_t)$ is determined by eq.(2).

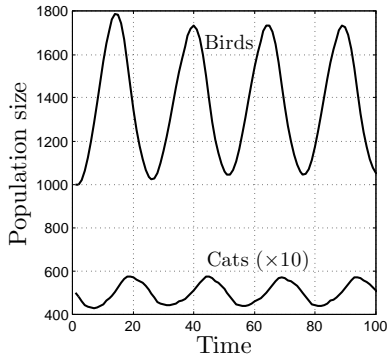


Figure 8: Bird and cat populations predicted with 1st model; constant culling effort.

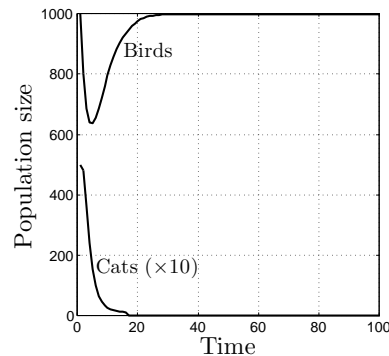


Figure 9: Bird and cat populations predicted with 2nd model; constant culling effort.

The dramatic predictive difference between the two models is illustrated in figs. 8 and 9. These figures show bird and cat population sizes vs. time for each of the two models, given a constant culling effort. The first model shows fairly stable and nearly harmonic oscillation of the populations, while the second model shows rapid and highly damped convergence to an equilibrium in which cats have been eradicated. The same total budget for culling effort is available in both cases, but the second model actually expends far less effort because the cats are eradicated fairly early. Thus the policy implications of these models are quite different, and we will see that their robustnesses to uncertainty are different as well.

Our management goal, as in section 3, is to maintain the size of the bird population above a specified minimum value, B_{\min} , for times $t = 1$ through $t = T$. We must choose the culling effort at each step, E_t , subject to the budget constraint, $E = \sum_{t=1}^T E_t$. The question is: how to choose a culling strategy? The answer is based on the idea of robustness.

4.2 Robustness

We define the robustness for each model just as we did in section 3. Indeed, the 1st model here is the same as the model used in that section so the robustness is the same as well. We denote the robustness functions for these two models as $\hat{h}_1(B_{\min})$ and $\hat{h}_2(B_{\min})$. In each case, we can understand the robustness in either of the senses introduced in section 2. If we specify a lowest acceptable bird population size, B_{\min} , then the robustness is the greatest fractional error of the model parameters at which the bird population will not fall below B_{\min} . Equivalently, if we postulate the magnitude of fractional error, then the robustness function determines the lowest size which the bird population can reach. The robustness function depends on the culling strategy. More robustness is better than less robustness, so the robustness function generates preferences on culling strategies.




Figure 10: Robustness vs minimal bird population size, B_{\min} , with constant culling effort, for two conflicting models.

Fig. 10 shows robustness functions for the two models with constant culling effort. The difference in the dynamic predictions of these models, which was illustrated in figs. 8 and 9, is transformed here into dramatic difference in robustness. The 1st model, which predicts persistence of both populations and hence persistent removal of small numbers of cats, shows the step-like robustness curve which we first encountered in section 3. The 2nd model predicts smooth and rapid extinction and its robustness curve is smooth as well.

The 2nd model is more robust than the 1st model for acceptable bird-population sizes below 600. For instance, requiring no less than $B_{\min} = 200$ birds, we are vastly more robust if we can rely on model 2 rather than model 1 ($\hat{h}_2(200)$ is off scale far to the right). But we are not in a position to favor model 2 over model 1. And furthermore, if we need a larger bird population, say $B_{\min} = 800$, then the 1st model is more robust than the 2nd.

How should we evaluate and choose among culling strategies? The approach is to extend the info-gap model of uncertainty beyond the fractional errors of the parameters and to include the structural uncertainty as well (as reflected in having two models rather than one). For any minimal acceptable population size, B_{\min} , the robustness is the greatest fractional error in the parameters of either model, so that the corresponding model satisfies the population requirement. We don't know the actual fractional errors of the parameters of either model, and we don't know which of these models we really should pay attention to. Consequently, we attend to both models and evaluate a culling strategy according to whichever model is more vulnerable at the specified minimal population size. The robustness of a culling strategy is the robustness of the more vulnerable of the two models.

We state this formally as:

$$\widehat{h}(B_{\min}) = \min \left\{ \widehat{h}_1(B_{\min}), \widehat{h}_2(B_{\min}) \right\} \quad (7)$$




Figure 11: Same as fig. 10, indicating overall robustness, $\widehat{h}(B_{\min})$.

The thick curve in fig. 11 is the overall robustness, $\widehat{h}(B_{\min})$, evaluated from the two individual model robustnesses displayed as thin curves. The constant-effort culling strategy is evaluated as the worst of the two alternatives. We don't know which model is in fact more accurate, so we evaluate the strategy in terms of whichever model is more vulnerable to error since we are unwilling to favor either model over the other.

4.3 Strategy Selection

Having established what we mean by robustness to conflicting uncertain models, we are now ready to evaluate alternative culling strategies.

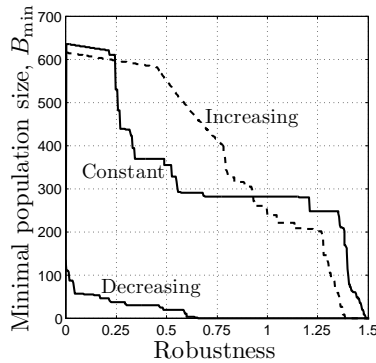


Figure 12: Robustness vs minimal bird population size, B_{\min} , for 3 different culling effort strategies.

Fig. 12 shows the overall robustnesses for the three different culling strategies described in section 3.1. The culling effort can be either constant over time, linearly increasing, or linearly decreasing. The total budget of effort is fixed.

Fig. 12 is a truncated version of fig. 5. We can understand this truncation by considering fig. 11 in which the robustness curve for the 1st model is truncated by the robustness curve for the 2nd. Much of the discussion of strategy selection in fig. 5 applies here as well. The decreasing culling effort would not be chosen since it is robust-dominated by the other two strategies. The robustness curves

for constant and linearly increasing effort cross one another so the preference between them depends on our performance requirement as discussed in section 3.2.

The policy impact of having two conflicting models is demonstrated in fig. 12. At high performance ($B_{\min} > 600$) the robustness is determined by the 2nd model for which the constant culling strategy is more robust. At lower performance ($B_{\min} < 600$) the robustness is determined by the 1st model for which either constant or increasing culling is more robust, depending on B_{\min} .

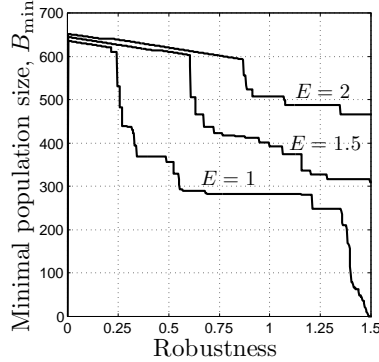


Figure 13: Robustness vs minimal bird population size, B_{\min} , for constant culling effort at 3 different total budgets.

Fig. 13 shows robustness curves for three different total budgets of culling effort, all with constant effort over time. This is a truncated version of fig. 6. The presence of two conflicting models has eliminated all curve crossing and established an unambiguous preference ordering: more culling effort is preferred over less in this specific case, unlike the situation in fig. 6.

5 Conclusion

Science-based models are useful, even when applied to systems which deviate in unknown ways from those for which the models were developed. However, it is then necessary to deal with the info-gap between model and reality. Management strategies cannot be evaluated only by their model-based predicted outcomes, since these predictions have no robustness against error in the model. We discussed three examples in which an info-gap robustness analysis is used to select a management strategy in attempting to reliably achieve an acceptable outcome.

These examples can be extended in many directions (relevant examples can be found in Ben-Haim, 2006). We could explore the robustness of the predator population, and the management implications of considering both predator and prey robustness. Our analysis can be applied to systems with three or more species. We could study how monetary allocations are translated into models of population dynamics. This would probably introduce additional uncertainties which must be modelled and managed. One might wish to consider uncertainty about the size of the smallest viable prey population. We have represented uncertainty in the population dynamics with non-probabilistic info-gap models of uncertainty. In some situations we have probabilistic information which, while valuable, may itself be uncertain. One would like to consider adaptive extensions: use part of the budget for exploratory management, improve the dynamic model from what is learned, and then invest the remaining budget. These and other extensions can be explored using info-gap techniques.

6 References

1. Ben-Haim, Yakov (2006). *Info-Gap Decision Theory: Decisions Under Severe Uncertainty*, 2nd edn. London: Academic Press.
2. Burgman, Mark (2005). *Risks and Decisions for Conservation and Environmental Management*, Cambridge: Cambridge University Press.
3. Courchamp, Franck, Michel Langlais and George Sugihara, 1999, Control of rabbits to protect island birds from cat predation, *Biological Conservation*, **89**, 219–225
4. Gigerenzer, Gerd and Reinhard Selten, eds. (2001). *Bounded Rationality: The Adaptive Toolbox*, Boston: MIT Press.
5. Hanski, Ilkka (1999). *Metapopulation Ecology*, Oxford: Oxford University Press.
6. IPCC-AR4 (2007). Global Climate Projections. In: *Climate Change 2007: The Physical Science Basis, Contribution of Working Group I to the Fourth Assessment Report of the Intergovernmental Panel on Climate Change*, Cambridge: Cambridge University Press, <http://www.ipcc.ch>.
7. Kot, Mark (2001). *Elements of Mathematical Ecology*, Cambridge: Cambridge University Press.
8. Krebs, Charles J. (1978). *Ecology: The Experimental Analysis of Distribution and Abundance*, New York: Harper & Row.
9. Oxford English Dictionary (1989). 2nd edn. Oxford: Oxford University Press.
10. Pielou, E.C. (1969). *An Introduction to Mathematical Ecology*, New York: Wiley.
11. Ramsey, Dave and Clare Veltman (2005). Predicting the effects of perturbations on ecological communities: what can qualitative models offer? *Journal of Animal Ecology*, **74**, 905–916.
12. Regan, Helen M., Yakov Ben-Haim, Bill Langford, Will G. Wilson, Per Lundberg, Sandy J. Andelman, Mark A. Burgman (2005). Robust decision making under severe uncertainty for conservation management, *Ecological Applications*, **15**(4), 1471–1477.
13. Royama, T. (1992). *Analytical Population Dynamics*, London: Chapman & Hall.
14. Simon, Herbert A. (1955). A behavioral model of rational choice, *Quarterly Journal of Economics*, **69**(1), 99–118.
15. Simon, Herbert A. (1956). Rational choice and the structure of the environment, *Psychological Review*, **63**(2), 129–138.
16. Snowdon, Brian, Howard Vane and Peter Wynarczyk (1994). *A Modern Guide to Macroeconomics: An Introduction to Competing Schools of Thought*, Cheltenham: Edward Elgar Publishing Company.
17. Speirs-Bridge, Andrew, Fiona Fidler, Marissa McBride, Louisa Flander, Geoff Cumming, and Mark Burgman (2010). Reducing overconfidence in the interval judgments of experts, *Risk Analysis*, **30**(3), 512–523.
18. Turchin, Peter, and Andrew D. Taylor (1992). Complex Dynamics in Ecological Time Series, *Ecology*, **73**(1), 289–305.