Proceedings of the Fourth International
Partners in Flight Conference: Tundra to Tropics
354–364

HIERARCHICAL POPULATION MODELS FOR THE RED-COCKADED WOODPECKER

MEVIN B. HOOTEN,[1,6] CHRISTOPHER K. WIKLE,[2] LAWRENCE D. CARLILE,[3] RICHARD E. WARNER,[4] AND DONALD PITTS[5]

[1]*Department of Mathematics and Statistics, Utah State University, Logan, Utah 84322-3900, USA;*

[2]*Department of Statistics, University of Missouri, Columbia, Missouri 65211, USA;*

[3]*Fort Stewart Fish and Wildlife Branch, 1177 Frank Cochran Drive, Fort Stewart, Georgia 31314, USA;*

[4]*214 Mumford Hall, Department of Natural Resources and Environmental Sciences, University of Illinois, Urbana, Illinois 61801, USA; and*

[5]*Engineer Research and Development Center, P.O. Box 9005, Champaign, Illinois 61826-9005, USA*

Abstract. The management of wildlife populations for sustainability has evolved as a main theme in natural science over the years and sophisticated statistical models have been developed for use in the prediction and characterization of population dynamics. Though game species serve as the common focus for many studies, the evaluation of populations involving threatened and endangered animals warrants significant attention as well. Due to reduced population densities and habitat fragmentation, the formal statistical modeling of endangered species populations presents significant challenges. In many cases, much scientific expertise exists with regards to a particular endangered species but data are scarce. Herein we adopt a hierarchical statistical model that formally accounts for expert knowledge while utilizing limited data to learn about the population dynamics of the Red-Cockaded Woodpecker, an explicitly managed and endangered species in the United States. The models we present are also used to make statistical population projections while rigorously accounting for uncertainty.

Key Words: Bayesian models, dynamical systems, endangered species, population biology.

MODÈLE POPULATIONNEL HIÉRARCHIQUE DU PIC À FACE BLANCHE (*PICOIDES BOREALIS*)

Résumé. La conservation pour le maintien des populations sauvages a évolué en un thème majeur des sciences naturelles au cours des années, et des modèles statistiques sophistiqués ont été développés afin de prédire et de caractériser la dynamique de ses populations. Même si les espèces chassées servent de modèles biologiques à de nombreuses études, le suivi des populations comprenant des animaux menacés ou en voie d'extinction a également retenu l'attention. En raison d'une réduction de la densité des populations et de la fragmentation de l'habitat, la modélisation statistique classique appliquée à des espèces en voie d'extinction présente de nombreuses limitations. Dans la plupart des cas, l'expertise scientifique existe pour une espèce menacée distincte, mais les données sont limitées. Dans cette étude, nous optons pour un modèle hiérarchique statistique qui prend en compte à la fois une connaissance poussée du système et une quantité limitée de données. Ce modèle nous permet d'étudier la dynamique des populations du Pic à face blanche, une espèce sous protection et considérée comme étant en voie d'extinction aux Etats-Unis. Le modèle que nous présentons est également utilisé pour la réalisation de projections statistiques à l'échelle de la population, tout en prenant rigoureusement en compte l'incertitude associée à nos données.

INTRODUCTION

The Red-Cockaded Woodpecker (RCW, *Picoides borealis*) is a federally listed endangered species (35 Federal Register 1970) native to open, fire-maintained pine ecosystems and is endemic to the Southeastern United States. Historically, southern pine ecosystems, particularly longleaf pine (*Pinus palustris*)—wiregrass (*Aristida stricta*) ecosystems,were fire-prone due to lightning strikes from frequent thunderstorms (Frost 1993). The RCW is a cooperatively

[6] E-mail: mevin.hooten@usu.edu

breeding species that lives in family groups (Walters 1990) and requires stands of mature living pines with large heartwood/sapwood ratios in which to excavate roosting and nesting cavities (Jackson et al. 1979). Additionally, the aggregate of cavity trees used by a group (cluster) must occur in stands with little or no hardwood mid-story or over-story. The RCW also requires an abundant foraging habitat of large pines, low densities of small pines, little or no hardwood mid-story, and an abundant groundcover of native bunchgrasses and forbs (Conner et al. 2001). Major factors limiting the growth of RCW populations are a lack of suitable cavities and fire suppression, which allows for hardwood encroachment (USFWS 2003). Managers can mitigate for a shortage of suitable cavities by installing artificial cavities in younger pines with smaller heartwood/sapwood ratios (Allen 1991). A potential breeding group (PBG) is an adult male and adult female that occupy the same cluster with or without one or more adult helpers (usually males from previous nesting seasons).

Clusters of cavity trees that are occupied by a family group (or non-empty subset of a family group) of RCWs are termed ``active clusters.'' Due to the fact that some clusters are occupied by single birds, the number of active clusters is not a complete measure of a successful RCW population by itself. The ratio of PBGs to active clusters is a better gauge of RCW population vigor, therefore a population model that accommodates both active clusters and PBGs is necessary. As a side note, an inactive cluster consists of a suitable number and arrangement of cavity trees that has either been occupied previously and abandoned, or newly created by managers (i.e., artificial) and is, as of yet, unoccupied.

Statistical Population Modeling

A hierarchical modeling approach for this particular ecological process (i.e., RCW population growth) is useful because it allows for the formal introduction of prior scientific knowledge and can incorporate multiple sources of uncertainty. A hierarchical statistical model can generally be constructed in terms of three main components or models that are linked by parameters and latent state variables (Berliner 1996). These components are frequently referred to as: The data model, the process model, and the parameter model. The data model, often traditionally called the likelihood, provides a formal link between the data and the underlying process from which the data arise. As is always true in statistical analyses, the actual data are treated as observed random variables and are frequently subject to measurement error, the form of which is explicitly specified in the likelihood.

The process component of the model can contain prior scientific understanding of how the true underlying process might behave (e.g., certain populations are known to exhibit various forms of logistic growth) as well as incorporate uncertainty related to the specific model used. A hierarchical model implemented with Bayesian methods will treat all model parameters as random variables, the probability distributions of which are to be estimated and then utilized to make statistical inference. The parameter model, also called the prior distribution, is the model component where any pre-existing scientific information about the parameters that govern the process or data models can be included. In the implementation, the parameter model represents the information we have about the process under study before observing the data. After a hierarchical model is fully specified, we then seek to update our understanding of the parameter distributions given additional information provided by the data. This updated parameter distribution is referred to as a posterior distribution and can then be used to make inference.

A full discussion of the procedure used to implement hierarchical Bayesian models is beyond the scope of this paper, however, additional details can be found in Gelman et al. (2004). It should also be noted that Bayesian statistical methods are recently enjoying substantial popularity in the ecological and biological sciences (Clark 2007, chapter 4), due at least in part to their ability to accommodate preexisting scientific knowledge and for the ease in which complex models can be specified and implemented.

In the setting specifically considered here, managers have a good understanding of important aspects of population growth such as growth forms, carrying capacities, and relationships between various cluster types. The formal incorporation of such information in a dynamical population model allows for meaningful prediction in the presence of uncertainty.

METHODS

Data

RCW data were collected from Ft. Stewart, Georgia for the years 1994 through 2004. At this site, on over 113 000 ha, there are inactive clusters (clusters of cavity trees containing no RCWs) that have been either abandoned by RCWs or newly created by managers to provide for RCW population expansion. The focus of the

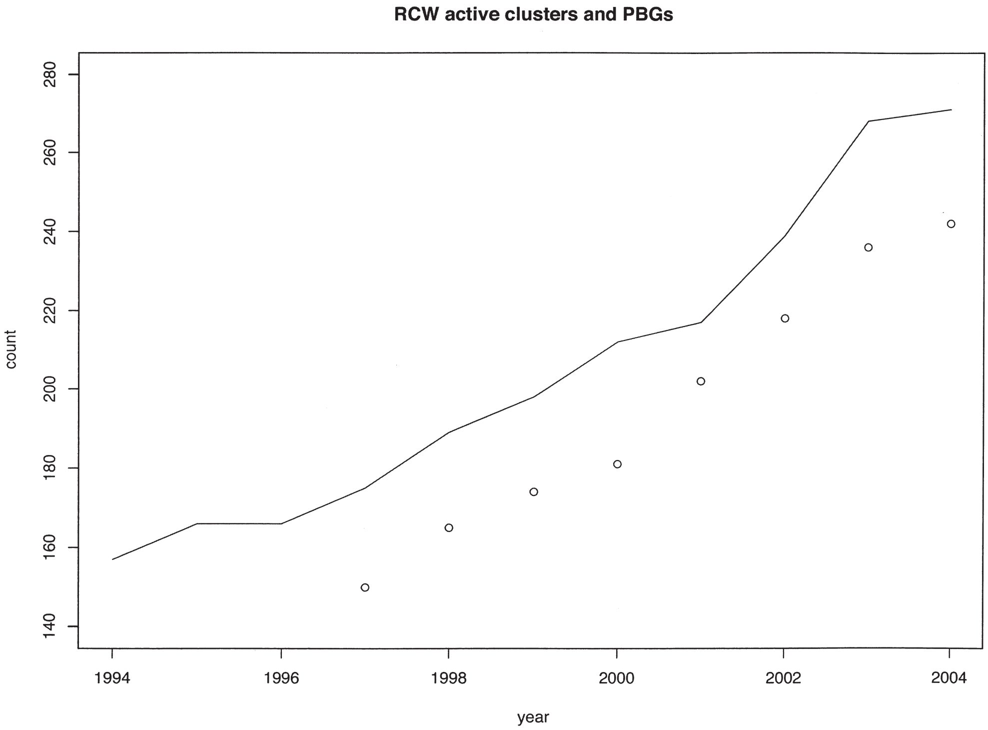

FIGURE 1. RCW active clusters (solid line) and PBGs (points) at Ft. Stewart for 1994 through 2004.

active cluster model is on two measures of RCW population status: number of active clusters and number of PBGs, while the total cluster model accommodates total clusters as well (i.e., active and inactive clusters). The number of PBGs is an estimate of the effective breeding size of the population at each time. Here, it is assumed that the number of active clusters is known at all times when data were collected at Ft. Stewart. Additionally, PBGs were not observed in every cluster in every year (Fig. 1).

Statistical Models

One way to model the growth in population over time while taking into consideration the limits imposed by density-dependence is to assume the true (underlying) growth in the population follows a conventional population biology growth equation (e.g., a Ricker growth curve). Such equations allow for ecologically realistic growth and carrying capacities while being flexible enough to allow for chaotic population behavior (though rare in nature) and possible limit-cycles. The discrete nature of the annual population data that are available for this study indicates the need for a discrete model in this situation.

A hierarchical framework allows for the incorporation of multiple sources of information about the growth of the RCW population and also accommodates the assumption of the existence of an underlying dynamical system. An additional advantage to a Bayesian implementation of the model is that it allows for various sources of uncertainty and previously existing scientific information about the RCW biology to be accounted for.

In what follows, we present two models of increasing complexity and show the difference in terms of inference about the RCW population on Ft. Stewart.

Active Cluster Model

First, consider a model for the number of active clusters and PBGs. Let N_t denote the number of active clusters at time t for $t = 1, \ldots, T$. This quantity is assumed known at each time for which data exist, although it need not be. Additionally, it is not known for times $t > T$ and is of interest at those future times. That is, we desire statistical forecasts of RCW population measures.

Let n_t denote the number of PBGs at time t for the finite set of times where PBG data were collected. Note that this may not be at every time for which the number of active clusters (N_t) is known. Now, since the number of PBGs is a subset of the active clusters, we can consider a binomial model for n_t given N_t and the probability of an active cluster being a PBG (p), at top level in hierarchical framework. Field expertise at Ft. Stewart suggests that there always exists a portion of active clusters not containing a breeding group. Thus, knowledge about the unknown probability p, in the data model described above, may be of use to managers. We have little prior information about what this probability should be at Ft. Stewart, and additionally, it could vary from one population to another, therefore a non-informative Jeffreys' prior is indicated (Carlin and Louis 2000). In a Bayesian model, the prior represents knowledge about population parameters before any data are observed, and in this case, for p, it allows us to account for a lack of specific and precise information about this parameter while still being able to estimate it.

Assuming that the dynamics of the ecological process (i.e., RCW population growth) are driven by some underlying mechanism related to cluster occupancy, we can then specify a Poisson process model as the stochastic form that gives rise to the number of active clusters given some underlying active cluster mean (i.e., $E(N_t) = L_t$). This mean, or intensity, is changing temporally and can be characterized through a mathematical formulation for population growth. Specifically, we let the Poisson intensity (L_t, the mean number of active clusters at time t) evolve over time according to the Ricker equation, a discrete density dependent population growth model (Turchin 2003, pg. 53). This specific mathematical form of growth allows the current population size to depend directly on the past population size as well as two parameters, b_1 and b_2, that control the growth rate and carrying capacity respectively. When considering the number of active clusters at Ft. Stewart, very specific prior information about the carrying capacity (b_2) exists due to the limited geographic extent of the study area and the cluster size constraints. On average, only 625 active clusters are possible within Ft. Stewart because of space limitations, so we specify a Gamma probability distribution for the prior of b_2, with a mean of 625 and standard deviation of 20, to allow for some uncertainty in the carrying capacity of active clusters due to varying cluster size and any clusters potentially overlapping the Ft. Stewart boundary.

Conversely, relatively little is known about the growth rate parameter (b_1) even though, on Ft. Stewart, much of the increase in the number of active clusters can be directly attributed to the active management of the species. That is, managers are regularly adding and removing potential clusters that may or may not become occupied by RCW in the future, though there is no way to explicitly quantify that activity in terms of the growth rate parameter (b_1). Therefore we specify a vague prior distribution over a scientifically reasonable range of values for this parameter (i.e., a Gamma with mean equal to one and standard deviation of two). Obviously, this growth rate can vary somewhat with changes in management practices and habitat. Such variation could be accounted for using a more complex time-varying growth rate model given additional covariate information. In the absence of quantifiable information about this covariate effect on Ft. Stewart, we must assume that b_1 is similar over the temporal domain and that its variability will account for some uncertainty related to changes in population growth rate. In fact, this assumption may not be unreasonable, given that recent findings suggest that RCW reproductive success is not significantly affected by military activity (Doresky et al. 2001).

Finally, to complete a specification of the active cluster model, the true population size at the initial time (L_1) in the period of interest is assumed to be an unknown random variable with a fairly informative Gamma distribution with mean equal to $N_{t=1}$ and standard deviation equal to two.

Total Cluster Model

Now, as a generalization to the active cluster model specified in the preceding section, consider a more complex model that accounts for a known number of total clusters (i.e., active and inactive clusters) as well as the observed number of active clusters, as before. Since we now have another time varying parameter to keep track of, we let the number of total clusters at time t be denoted as C_t and then retain the earlier notation from the active cluster model for the observed number of active clusters (N_t) and PBGs (n_t). Now, all of n_t, N_t, and C_t are assumed to be known at times for which data exist and have the constraint that n_t is at most N_t, while N_t itself has an upper bound of C_t. Additionally, it is believed that while the average ratio of PBGs to active clusters remains constant over time, the number of active clusters actually approaches the number of total clusters over time. This is partially due to the age of the forest becoming more suitable for RCW success, but may also be due to the utilization of all available space at Ft. Stewart, as resources becoming limiting when population size approaches carrying capacity. This constraint poses significant challenges for predicting active clusters and PBGs at future times. Therefore, if the number of total clusters is to be appropriately accommodated, a non-trivial extension to the active cluster model must be made. In doing so, an added benefit of the upgrade is a richer probabilistic characterization of all components of RCW population growth.

In this extended total cluster model, we retain the earlier specification for the number of PBGs (n_t), but now given the new information about the total number of clusters (C_t) we let the active clusters (N_t) arise from another binomial distribution depending on C_t and a time varying probability of an unoccupied cluster becoming occupied (p_t). We have only the number of total clusters left to model, and given that this number can vary as more clusters are created by the RCW as well as the Ft. Stewart managers, we allow C_t to arise from a Poisson distribution with intensity $L_{C,t}$.

Now, to allow for dynamic temporal evolution in the clusters, and hence RCW population size, we let $L_{C,t}$ be modeled by a Ricker growth equation with random growth rate ($b_{C,1}$) and carrying capacity (b_2) parameters. Recall that for the model to be realistic, N_t needs to approach C_t as C_t approaches the carrying capacity (b_2); thus, this implies that the probability (p_t) of an unoccupied cluster becoming occupied at time t must approach one. Since the rate at which p_t converges to one is unknown, it must be estimated from the data. Thus, consider a second dynamical growth process ($L_{N,t}$) that represents the average number of active clusters over time. Similarly to the model for $L_{C,t}$, we let $L_{N,t}$ evolve over time according to the same growth equation with the same carrying capacity (b_2), but different growth rate ($b_{N,1}$). Finally, setting $p_t = L_{N,t} / L_{C,t}$ allows p_t to approach one, as needed. Then to complete the hierarchical specification, all model parameters and initial states were assigned the same Gamma prior distributions as used in the active cluster model.

As in the active cluster model, posterior predictions of active clusters and PBGs, for times $t > T$, are the primary focus of our statistical inference and can be found readily by integrating over the model parameters (Gelman et al. 2004).

RESULTS

The active and total cluster models are non-conjugate in all parameters, with the exception of p, thus, when computationally implementing the models, a Gibbs sampler with Metropolis-Hastings updates was necessary to sample from the posterior distributions using the R Statistical Computing Environment (R Core Development Team 2007). Convergence was reached almost immediately for all parameters (using 100 000 Gibbs iterations with a 10 000 iteration burn-in period). Gibbs samples were systematically thinned to remove any correlation induced by the fitting algorithm.

In what follows, we present the results from each of the model fits separately and then provide a summative discussion and general conclusions in the two latter sections.

Active Cluster Model

First, for the active cluster model, consider the parameter estimates for the population growth (b_1 and b_2 in Figure 2). Their corresponding prior distributions are also provided for comparison.

The posterior distribution pertaining to the probability that an active cluster becomes occupied by a PBG (p) is given in Figure 3, while the posterior means and 95% credible intervals are provided for the PBGs at all times and for the posterior predictions of active clusters and PBGs for 50 years into the future in Figure 4. Additionally, in Figure 4, is the current PBG management objective of 350 breeding groups and its intersection with the posterior predictive distribution for n_t.

Total Cluster Model

For the model that explicitly takes into account the number of total clusters, there are two growth rate parameters ($b_{N,1}$ and $b_{C,1}$) that control the speed at which the number of active clusters and total clusters converge (Figure 5a) to the carrying capacity (b_2) (Figure 5b).

The posterior distribution for parameter relating to the probability that an active cluster will become a PBG (p) is given in Figure 6, while 95% posterior predictive credible intervals are provided for the PBGs at all times and for the total and active clusters 50 years into the future in Figure 7. Additionally, in Figure 7, is the current PBG management objective of 350 and its intersection with the posterior predictive distribution for n_t.

DISCUSSION

These models highlight the advantages of incorporating known sources of uncertainty into an ecological model as well as the effect of excluding potentially important information (i.e., number of total clusters). For the active cluster model, the fact that the prior and posterior distribution for the carrying capacity

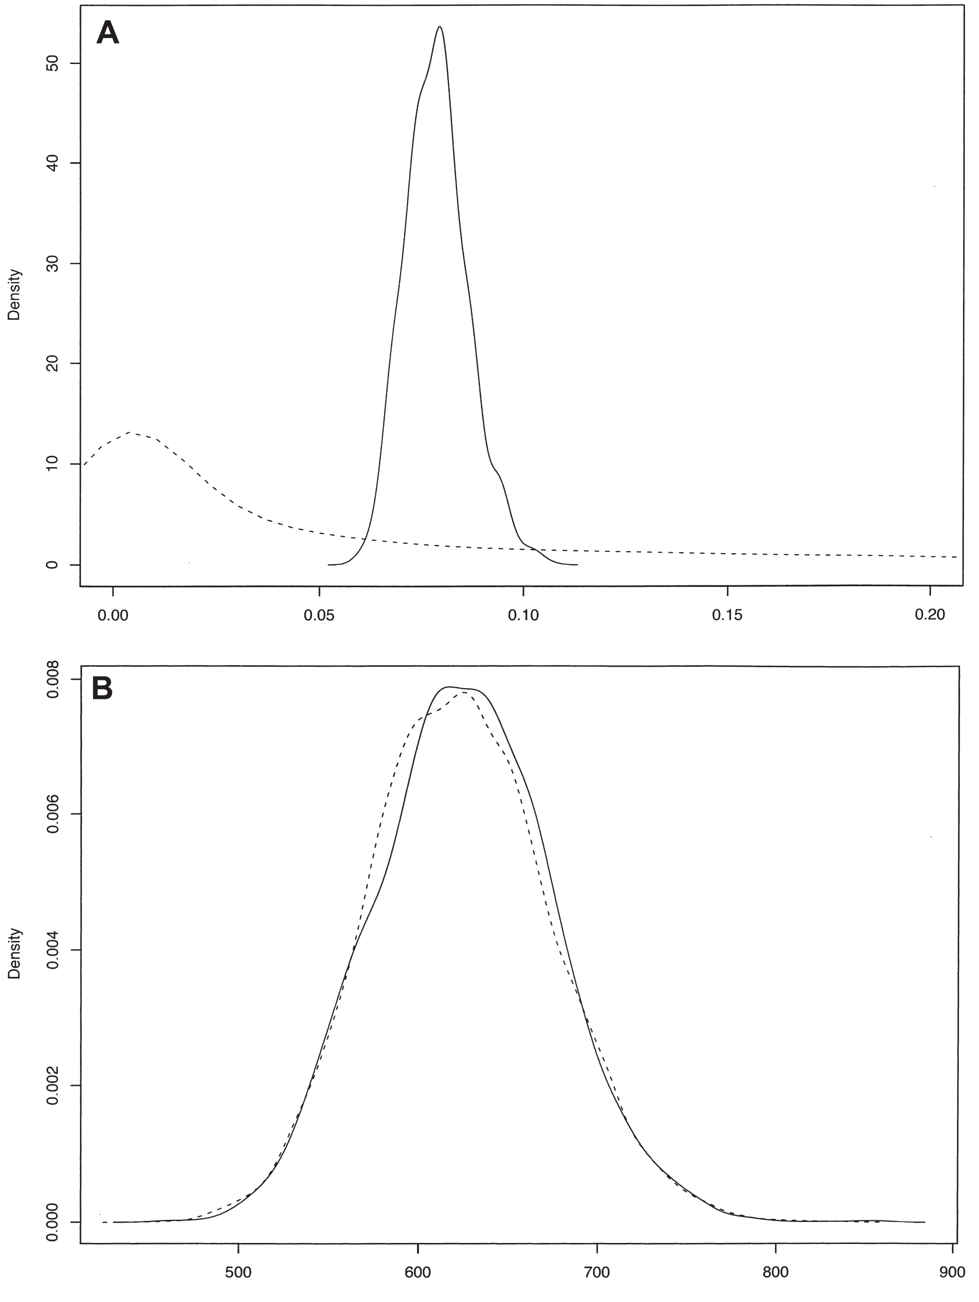

FIGURE 2. (A) Prior (dashed) and posterior (solid) distributions for the growth rate. (B) Prior (dashed) and posterior (solid) distributions for the carrying capacity.

parameter (b_2; Fig. 2b) are quite similar suggests that the data contains limited information about the RCW saturation for this study area. This is likely due to the lack of an asymptote in the data (Fig. 1). However, since specific knowledge about this parameter was already available, it allows the model to focus on the estimation of the growth rate parameter (b_1). In this case, as mentioned earlier, a very naive model is assumed for b_1. Though, in the presence of covariate information, this need not be the case. Also, notice the increase in precision for the

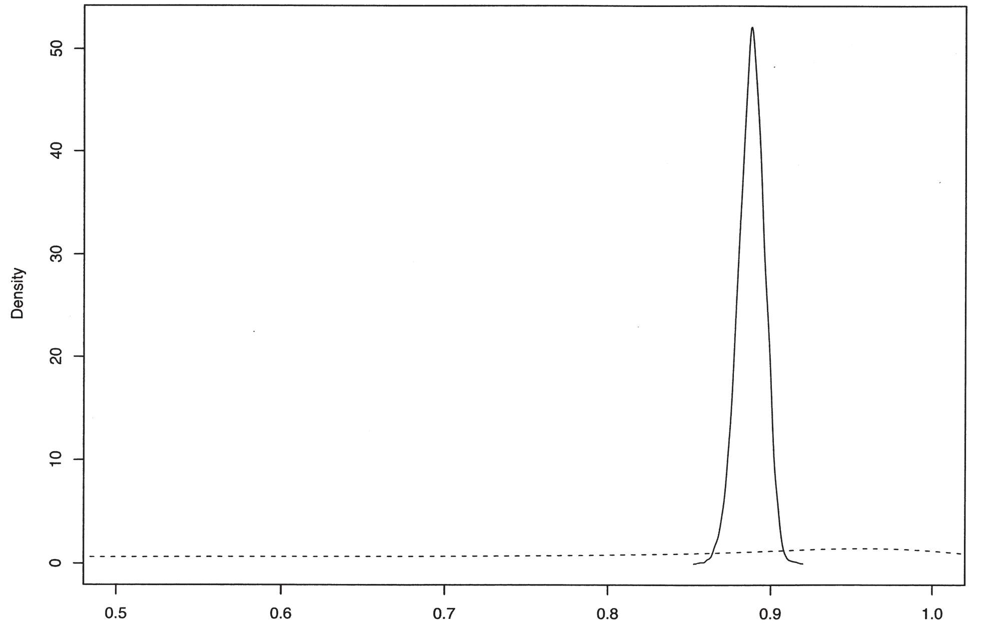

FIGURE 3. Prior (dashed) and Posterior distribution (solid) for the probability (p) of an active cluster containing a PBG.

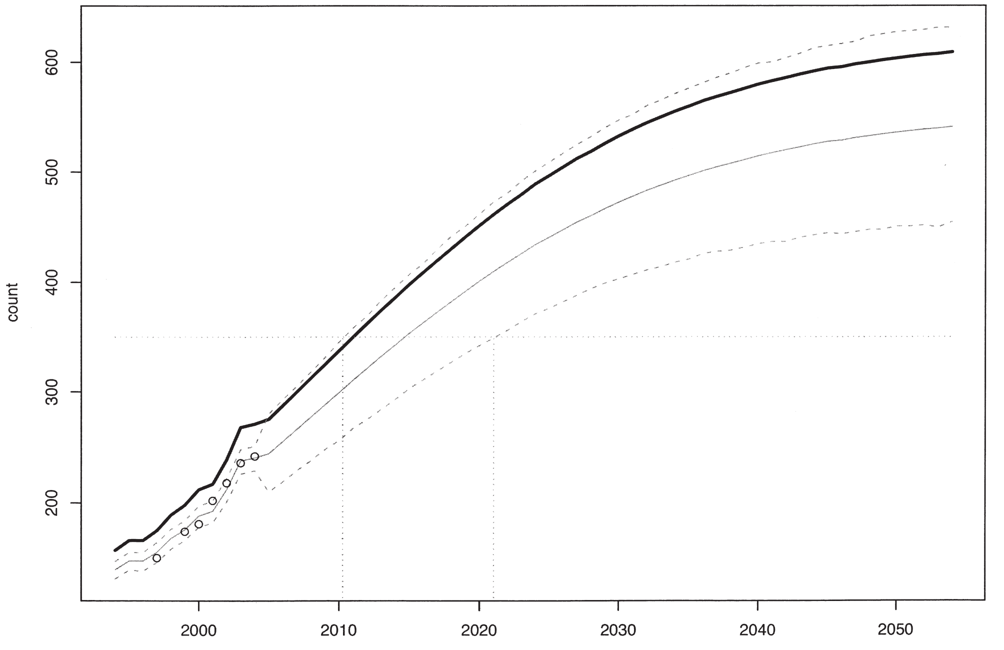

FIGURE 4. Predictions for active clusters (bold line) and posterior mean (solid, thin line) and 95% credible interval (dashed, thin line) for the number of PBGs (thin) at Ft. Stewart. Also shown are the PBG recovery goal (horizontal dotted line) and its intersection with the PBG credible interval (vertical dotted lines). PBG data are shown as points.

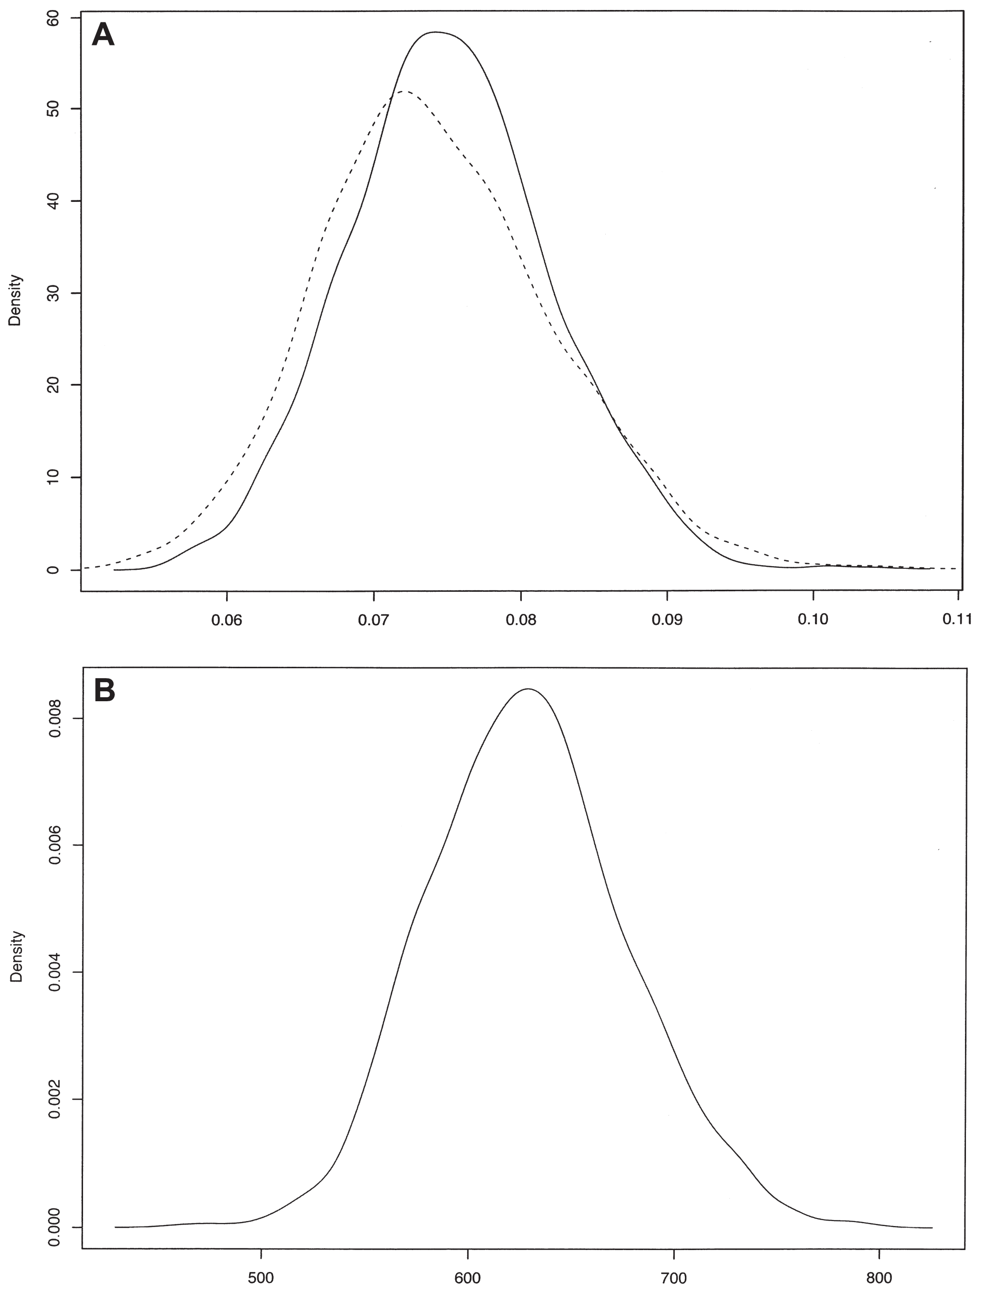

FIGURE 5. (A) Posterior distributions of the growth rates ($b_{C,1}$ = solid and $b_{N,1}$ = dashed). (B) Posterior distribution of the carrying capacity (b_2).

posterior estimate of b_1 as compared to its prior (Fig. 2a). This indicates that the data contain substantial information about this parameter.

The probability of an active cluster containing a PBG (p) is also a parameter that is estimated precisely despite its non-informative prior (Fig. 3). The posterior mean of p (i.e., $E(p \mid N, n) = 0.888$) is quite similar to point estimates from previous studies ($p = 0.89$). The uncertainty in this parameter, however, represents a little studied, but much speculated about characteristic of RCW biology and may provide useful information for setting management objectives for this species.

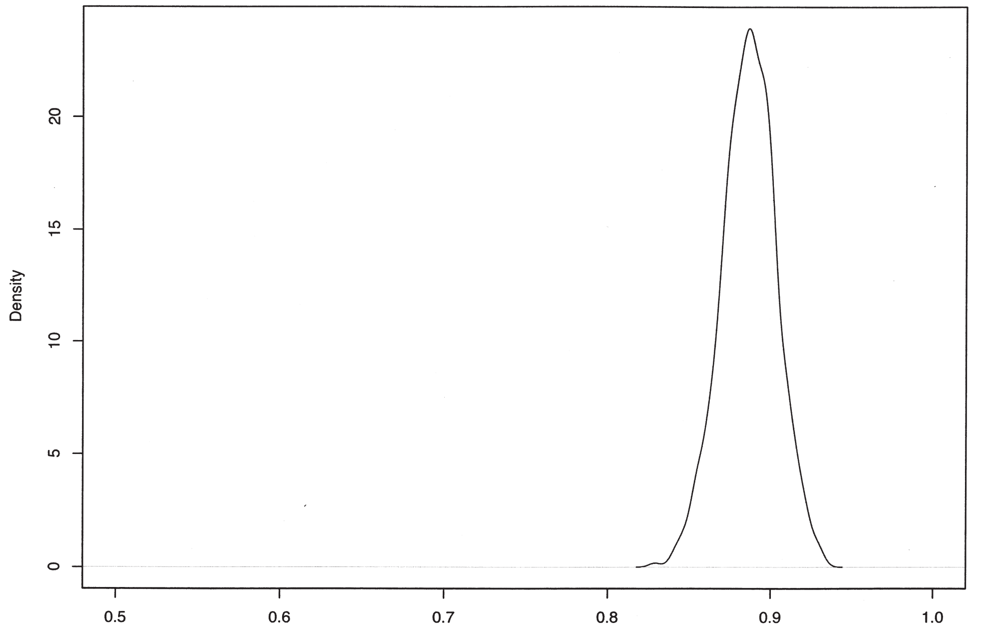

FIGURE 6. Posterior distribution for the probability (p) of an active cluster containing a PBG.

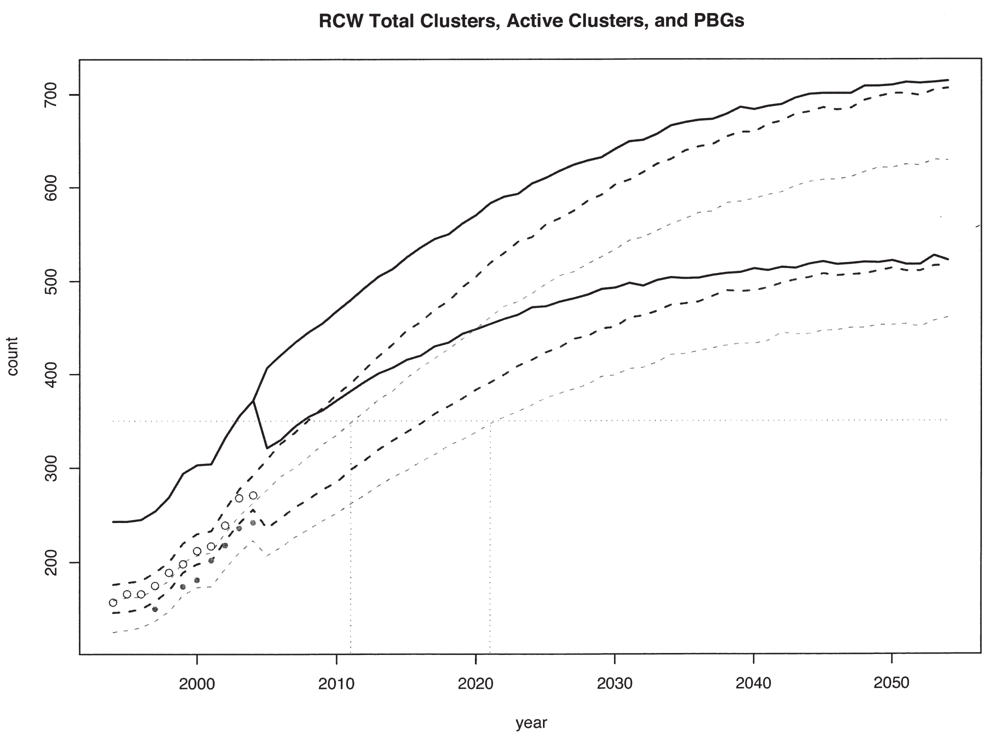

FIGURE 7. 95% predictive credible intervals for total clusters (bold lines), active clusters (dashed lines), and for the number of PBGs (thin dotted lines) at Ft. Stewart. Also shown are the PBG recovery goal (horizontal dotted line) and its intersection with the PBG credible interval (vertical dotted lines). PBG data are shown as points.

One of the primary interests here is the evolution of the true population as a dynamical system. Of specific utility, are the trajectories and forecasts for critical population measures. From Figure 4 it is apparent that for years where data exist there is little uncertainty about the number of PBGs and the observed data fall into the expected range of values. However, for times without data, especially where the number of active clusters is unknown, there exists substantial uncertainty associated with these population measures. Over the course of 50 years, the posterior predicted population levels approach saturation. The level of saturation, though, varies by about 200 clusters. Additionally, the posterior predicted date at which the current management goal of 350 PBGs will be reached is between the years 2010 and 2021 with probability 0.95 (the posterior predictive mean date at which 350 PBGs will be reached is year 2015).

The total cluster model produces similar results with subtle but possibly important differences. Based on the posterior distributions for the growth rate parameters ($b_{N,1}$ and $b_{C,1}$), it is clear that although $b_{C,1}$ is more precise, the two estimated distributions are not significantly different. This suggests that based on the data, the initial state and carrying capacity have more influence on the rate of convergence between N_t and C_t than do separate growth rates. Obviously, with more data, this influence could change.

As with the active cluster model, the posterior distribution of p is estimated using the total cluster model (Fig. 6) and is similar but slightly less precise than that of the active cluster model. This suggests that there is slightly more uncertainty associated with the probability of an active cluster containing a PBG when information exists about the total number of clusters.

The 95% posterior predictive credible envelopes for C_t, N_t, and n_t from the total cluster model (Fig. 7) appear similar to those from the active cluster model (Fig. 4). Notice, however, that the predicted date to attain the management goal of 350 PBGs is shifted back to between 2011 and 2021 with a probability of 0.95 using the total cluster model. Thus the extra information related to total clusters increases the precision of the posterior predictions but delays the earliest probable date by approximately 1 year.

Dynamic Population Growth

In the proposed model, various components of the RCW population growth (i.e., total clusters, active clusters, and thus PBGs), are modeled via a stochastic Ricker growth equation. The population biology literature contains much debate about the appropriateness of various discretized differential equation models (e.g., Kot 2001, Turchin 2003). From a purely deterministic standpoint, certain growth models are preferred over others for various reasons. For example, the Beverton-Holt growth model is sometimes preferred because it behaves similar to the logistic differential equation (i.e., the continuous time model that serves as motivation for many discrete time growth equations). The case can be made that other models (e.g., the Ricker growth model) can exhibit more complex dynamical behavior (e.g., stable limit cycles and chaotic behavior, though the latter occurs infrequently in nature) and are thus more flexible. In deterministic analyses, there may be important scientific reasons to choose one form or another, but in a stochastic setting such as the one considered here, results pertaining to deterministic properties of such models are not so easily interpreted. The hierarchical framework allows us to account for uncertainty in this two-component growth process and inference is being made on a set of models rather than just one. A stochastic Ricker model is flexible enough to exhibit behavior similar to many other forms of deterministic models and thus the choice of the form of growth is less important than the manner in which uncertainty is accounted for. Additionally, the data used in this study indicate that the population has not reached carrying capacity yet and thus nonlinear growth characteristics have not been observed. The choice of growth model may prove to be more important when the process has been observed more extensively. Another alternative is to model the growth semi-parametrically and allow the data to choose the most appropriate model; this may be an approach worth considering in future modeling efforts, when more data have been collected and nonlinearity in the process has been observed.

CONCLUSION

In summary, simple but ecologically meaningful models are proposed for estimating population growth parameters and forecasting critical population measures of RCW at Ft. Stewart via a hierarchical Bayesian framework with a latent non-linear dynamical system. Specifically, the use of a scientifically motivated population model for RCW growth enhances the knowledge gained from limited cluster occupancy data and allows for statistical population forecasts that are based both on current expert knowledge as well as recently collected data.

It is worth restating that both process growth models are very naive, in that, they do not explicitly consider management practices and (or) habitat changes over time. Thus the parameter estimates are only valid under the assumption that the model is implicitly accounting for the covariate effects. Additionally, the current models have no mechanism to account for possible future population collapses (e.g., major disturbances such as hurricanes or military activities). However, based on information available and strong but scientifically reasonable assumptions, they allow for long-range population projections and provide a means for true probabilistic inference.

ACKNOWLEDGMENTS

This work was supported by University of Missouri Subaward 2005-4159-01 from the University of Illinois in support of the University's Prime Award with the Department of Army, U.S. Army Construction Engineering Research Laboratory. The authors also thank Lise Aubry for her assistance.

LITERATURE CITED

Allen, D. 1991. Constructing artificial red-cockaded woodpecker cavities. Forest Service Technical Report SE-73, USDA.

Berliner, L. 1996. *In* Maximum Entropy and Bayesian Methods, chapter Hierarchical Bayesian time series models, pages 15–22. Kluwer Academic Publishers.

Carlin, B. P., and T. A. Louis. 2000. Bayes and Empirical Bayes methods for Data Analysis. CRC, Boca Raton, FL.

Clark, J .S. 2007. Models for Ecological Data: An Introduction. Princeton University Press, Princeton, NJ.

Conner, R., D. Rudolph, and J. Walters. 2001. The red-cockaded woodpecker: surviving in a fire-maintained ecosystem. University of Texas Press, Austin, TX.

Doresky, J., K. Morgan, L. Ragsdale, and H. Townsend. 2001. Effects of military activity on reproductive success of red-cockaded woodpeckers.

Frost, C. 1993. Four centuries of changing landscape patterns in the longleaf pine ecosystem, pp. 17–44. *In* S. Herman [ed.], The longleaf pine ecosystem: ecology, restoration, and management, number 18. Tall Timbers Research Station, Tallahassee, FL.

Gelman, A., J. Carlin, H. Stern, and D. Rubin. 2004. Bayesian Data Analysis, 2nd ed. CRC, Boca Raton, FL.

Jackson, J., M. Lennartz, and R. Hooper. 1979. Tree age and cavity initiation by red-cockaded woodpeckers. Journal of Forestry 77:59–63.

Kot, M. 2001. Elements of Mathematical Ecology. Cambridge University Press, Cambridge, UK.

R Core Development Team. 2007. R: A Language and Environment for Statistical Computing. Vienna, Austria: R Foundation for Statistical Computing. [Online.] <http://www.R-project.org>.

Turchin, P. 2003. Complex Population Dynamics. Princeton University Press, Princeton, NJ.

USFWS. 2003. Recovery plan for the red-cockaded woodpecker (*Picoides borealis*). U.S. Fish and Wildlife Service, Atlanta, Georgia.

Walters, J. 1990. Cooperative breeding in birds, chapter Red-cockaded woodpeckers: a primitive cooperative breeder. Cambridge University Press, London, UK.