

Continuous Covariates in Mark-Recapture-Recovery Analysis: A Comparison of Methods

Simon J. Bonner^{1*}, Byron J. T. Morgan^{2**}, and Ruth King^{3***}

¹Department of Statistics, University of British Columbia

Vancouver, BC V6T 1Z2, Canada

²Institute of Mathematics, Statistics and Actuarial Science, University of Kent

Canterbury, Kent CT2 7NF, England

³School of Mathematics and Statistics, University of St Andrews

St Andrews, Fife KY16 9LZ, Scotland

**email*: s.bonner@stat.ubc.ca

***email*: B.J.T.Morgan@kent.ac.uk

****email*: ruth@mcs.st-and.ac.uk

SUMMARY: Time-varying, individual covariates are problematic in experiments with marked animals because the covariate can typically only be observed when each animal is captured. We examine three methods to incorporate time-varying, individual covariates of the survival probabilities into the analysis of data from mark-recapture-recovery experiments: deterministic imputation, a Bayesian imputation approach based on modelling the joint distribution of the covariate and the capture-history, and a conditional approach considering the events which depend only on the observed covariate data (the trinomial method). After describing the three methods, we compare results from their application to the analysis of the effect of body mass on the survival of Soay Sheep (*Ovis aries*) on the Isle of Hirta, Scotland. Simulations based on these results are then used to make further comparisons. We conclude that both the trinomial and Bayesian imputation methods perform best in different situations. If the capture and recovery probabilities are all high, then the conditional model produces precise, unbiased estimators that do not depend on any assumptions regarding the distribution of the covariate. In contrast, the Bayesian imputation method performs substantially better when capture and recovery probabilities are low, provided that the specified model of the covariate is a good approximation to the true data generating mechanism.

KEY WORDS: Bayesian Inference; Imputation; Individual Covariates; Mark-Recapture-Recovery; Missing Covariates; Time-Varying Continuous Covariates; Trinomial Model

1. Introduction

Mark-recapture-recovery (MRR) methods are widely applied in the study of wild animal populations. On a series of capture occasions, subsets of the animals are captured, marked with unique identifiers, and released back into the population. Demographic parameters of interest are then estimated by modelling the recaptures of live and/or recoveries of dead individuals on later occasions. While MRR data provide information about processes of both recruitment into and departure from the population, the most common objectives are to estimate the probabilities of survival between consecutive capture occasions and further to understand how different variables explain changes in survival among individuals and over time (e.g. Lebreton et al., 1992; Pollock, 2002).

Potential covariates for modelling the probability that a marked individual survives from one occasion to the next may be cross-classified as either constant or time-varying and either extrinsic (environmental) or intrinsic (individual). Of these four categories, time-varying, intrinsic covariates, like an individual's fitness or body mass, pose a specific problem in the analysis of data from MRR experiments. Extrinsic variables can be measured at will and constant individual covariates need only to be measured once for each individual (typically on the first occasion the individual is captured). However, time-varying intrinsic variables can only be observed for a specific individual on the occasions when that individual is captured. As a result, a potentially large proportion of the values needed to model the effects of such covariates will be unknown.

Time varying, intrinsic covariates may be further categorized as discrete or continuous. Information from those that are discrete, like breeding status, can be incorporated with the well known multi-state model (Arnason, 1973; Schwarz et al., 1993). In essence, this model constructs the distribution of the covariate over all occasions by assuming that each individual's covariate values (states) obey the Markov property. The multi-state likelihood

contribution for each individual is then formed by summing the joint probability of the capture history and covariate values over all possible configurations of the unknown covariates.

One strategy for modelling the effect of a continuous covariate is to partition its range into a finite number of values and then to apply the multi-state model. For example, Nichols et al. (1992) studied survival of meadow voles (*Microtus pennsylvanicus*) as a function of size by dividing mass into four discrete categories. However, discretizing a continuous variable in this way is subjective, results in some loss of information, and may increase the uncertainty of the estimated survival probabilities.

A second strategy is to impute the missing values: replacing the missing values by some deterministic algorithm and then proceeding with the analysis as if these values had been observed. This approach is common in the analysis of longitudinal data, and several methods of imputation exist including last observation carry forward (LOCF), linear imputation, and replacement by the observed individual or occasion means (Little and Rubin, 2002). Although these methods may produce reasonable guesses at the missing values in some situations, two problems arise: inferences may be severely biased if the chosen algorithm does not match well with the true data generating mechanism and standard errors may be artificially deflated because all uncertainty in the missing values is ignored.

Recently, Bonner and Schwarz (2006) and King et al. (2008) proposed a new method that extends the idea behind the multi-state model to the case of continuous covariates. Their Bayesian imputation approach models the (now continuous) distribution of the covariate so that the likelihood contribution for each individual is formed by the joint distribution of both its capture history and observed covariate values. The primary challenge in moving to the continuous case is that sums associated with unobserved values in the multi-state model are replaced by integrals, which makes the likelihood analytically intractable and hinders

estimation with classical methods. Instead, both Bonner and Schwarz (2006) and King et al. (2008) obtain inferences from Bayesian methods via Markov chain Monte Carlo (MCMC).

At the same time, Catchpole et al. (2008) developed an alternative approach based on conditional methods for the analysis of MRR data with missing covariates. These methods avoid missing values in the likelihood by considering only the events for which the complete covariate information is available. For models with continuous, time-varying covariates of the survival probability, the likelihood contribution for an individual includes only those events directly following the occasions on which the individual was captured and released. Each event included in the likelihood has three potential outcomes (the individual is either recaptured at the next occasion, recovered before the next occasion, or not observed at all) and, for this reason, the specific approach is referred to as the trinomial method. The corresponding likelihood constructed in this way contains no missing values, so that it is easily evaluated and maximized to estimate the effect of the covariate on survival. Simulations conducted by Catchpole et al. (2008) demonstrated that estimates from the trinomial model are unbiased and outperform estimates from a naive complete case approach. A second conditional method (the binomial method) also included in these simulations performed slightly better than the trinomial, but this method requires external estimation of the capture and recovery probabilities and is not considered further.

The primary objective of this paper is to compare the performance of the Bayesian imputation and trinomial methods for handling continuous, time-varying, individual covariates in MRR analysis. Section 2 describes these two methods as well as the deterministic imputation approach. In Section 3, we apply all three methods to analyse data obtained from a 16 year study of Soay sheep (*Ovis aries*) on the Isle of Hirta, Scotland, and Section 4 presents a simulation study based on the results. In the final section, we discuss our conclusions and provide some general recommendations for choosing between the methods.

2. Methods

2.1 MRR Experiments with Continuous Covariates

In this section, we describe the data structure and basic probability model for MRR studies of populations with only one age class (i.e., such that all individuals with the same value of the covariate share the same probabilities of capture, recovery, and survival). In analysing the effects of body mass on the survival of the Soay sheep of Hirta, Section 3, we extend this to a model with four age classes allowing for differential growth and survival between the age classes, but assuming equal capture and recovery probabilities for all of sheep.

The data collected for each individual captured in an MRR experiment are easily summarized in the capture history. With T capture occasions, the capture history for the i^{th} marked individual is a string of T values, denoted by ω_i , such that:

$$\omega_{it} = \begin{cases} 1 & \text{if the individual is captured alive on occasion } t \\ 2 & \text{if the individual is recovered dead between occasions } t-1 \text{ and } t \\ 0 & \text{otherwise} \end{cases}$$

The simplest MRR model for estimating survival probabilities without any explanatory variables defines probabilities for each capture history conditional on an individual's first release in terms of three sets of parameters:

- 1) Survival probabilities: $\phi_t = P(s_{i,t+1} = 1 | s_{it} = 1)$
- 2) Capture probabilities: $p_t = P(\omega_{it} = 1 | s_{it} = 1)$
- 3) Recovery probabilities: $\lambda_t = P(\omega_{it} = 2 | s_{it} = 0, s_{i,t-1} = 1)$

where $s_{it} = 1$ (individual i is alive on occasion t). It is assumed that these probabilities are the same for all individuals, that events for one individual are independent of the events for all other individuals, and that dead animals can only be recovered in the interval in which they die. As an example, the probability assigned to the capture history $\omega_i = 11012$ is: $\phi_1 p_2 \phi_2 (1 - p_3) \phi_3 (1 - \phi_4) \lambda_4$. The full likelihood is then formed by multiplying the conditional

probabilities assigned to the capture histories of each observed individual, and estimates of the capture, recovery, and survival probabilities may be computed by maximum likelihood (ML) or a Bayesian approach.

The methods we consider here extend the simple model by allowing for individual heterogeneity in the survival probabilities as a function of a covariate. In particular, we assume that the probability that individual i survives from occasion t to $t + 1$ is solely dependent on the value of a continuous, time-varying, individual covariate, x_{it} , which can only be observed if the individual is captured on occasion t , and model the survival probability as a linear function of x_{it} on the logit scale:

$$\text{logit}(\phi(x_{it})) = \beta_0 + \beta_1 x_{it}. \quad (1)$$

The probability assigned to the capture history $\omega_i = 11012$ becomes: $\phi(x_{i1})p_2\phi(x_{i2})(1 - p_3)\phi(x_{i3})(1 - \phi(x_{i4}))\lambda_4$. Note that this probability depends on the value of x_{i3} which cannot be observed because the individual was not captured on occasion 3, and so it is not a valid likelihood contribution. Something must be done to remove the unobserved values from the likelihood function before parameter estimates can be computed.

2.1.1 Simple Imputation. Perhaps the most basic method for analysing data with missing covariates is simple imputation. Values for the missing covariates are generated from some deterministic model and the completed data set is then analysed as if these values had been observed. Little and Rubin (2002) describe several strategies for imputing missing values in longitudinal data that could also be used with MRR data including: last observation carry forward (LOCF), linear imputation, and both conditional or unconditional mean imputation.

In our analysis of the Soay sheep data (Section 3) and the simulations (Section 4) we impute the missing values by linear interpolation between an individual's first and last capture and by LOCF after the last recapture. This strategy seems sensible for imputing the missing body masses because it allows for growth between the first and last occasions that an individual

is known to be alive, but avoids unreasonably small or large values that might arise if linear imputation is continued after an individual is last observed. Note that the values prior to the first occasion that an individual is captured do not need to be imputed because these occasions do not contribute to the likelihood. Applying this imputation method to the sample capture history above, x_{i3} is replaced by the mean of x_{i2} and x_{i4} and the likelihood contribution becomes: $\phi(x_{i1})p_2\phi(x_{i2})(1-p_3)\phi((x_{i2}+x_{i4})/2)(1-\phi(x_{i4}))\lambda_4$. The full likelihood is again formed by multiplying the probabilities assigned to each observed capture history capture for individual, once the missing covariates have been imputed, and ML or Bayesian methods can be used to make inference about the survival probabilities. In our computations, ML estimates and their asymptotic standard errors were obtained using the Quasi-Newton Broyden, Fletcher, Goldfarb, and Shanno (BFGS) optimization algorithm implemented in the function `optim` in the R software package to maximize the likelihood and compute the information matrix (R Development Core Team, 2009).

2.1.2 Bayesian Imputation Approach. A second approach to the missing value problem is to model the distribution of the covariate values. The likelihood contribution for each individual is then formed as the joint density of the capture history and the observed covariate values. If all values of the covariate had been observed, then this density would simply be the product of two components: the conditional probability of the capture history given the covariates and the joint density of the observed covariate values. When some values of the covariate are unobserved, as in our situation, the likelihood contribution is formed by integrating this product with respect to each of the missing values.

Consider again the sample capture history $\omega_i = 11012$ and let $f(\cdot)$ denote the chosen model of the covariate. The joint density of the capture history and the values of the covariate on all capture occasions is then: $\phi(x_{i1})p_2\phi(x_{i2})(1-p_3)\phi(x_{i3})(1-\phi(x_{i4}))\lambda_4 \cdot f(x_{i2}, x_{i3}, x_{i4})$. However, it is necessary to integrate across the value of x_{i3} , because this value was not observed, and

the actual likelihood contribution for this capture history becomes: $\int_{-\infty}^{\infty} \phi(x_{i1})p_2\phi(x_{i2})(1 - p_3)\phi(u)(1 - \phi(x_4))\lambda_4 \cdot f(x_{i2}, u, x_{i4}) du$. Except when very few values are missing, the number of large integrals in the likelihood formed by multiplying the contributions from all animals will make it impossible to perform ML estimation. Instead, Bonner and Schwarz (2006) and King et al. (2008) draw inference about the survival probabilities from Bayesian inference.

Following Bonner and Schwarz (2006), we specify a conditional model for \boldsymbol{x} motivated by assuming that the covariate follows a diffusion process with time-dependent mean but constant variance. Given the value of the covariate on occasion t , the distribution of the covariate on occasion $t + 1$ is defined by:

$$x_{t+1}|x_t \sim N(x_t + \mu_t, \sigma^2), \quad t = 2, \dots, T. \quad (2)$$

with μ_t varying over time and σ^2 constant.

As in most complex Bayesian problems, the dimension of the posterior distribution makes it impossible to compute summary statistics directly. Instead, inferences for the data analysis and simulation study were based on samples from the posterior distribution obtained via MCMC using the `OpenBUGS` software package (Thomas et al., 2006). To determine the necessary length of the chains and the burn-in period required to reach convergence, 8 MCMC chains were run for one simulated data-set starting at widely spaced initial values. Plots of the Brooks-Gelman-Rubin diagnostics indicated that fewer than 15,000 iterations were required to estimate quantiles within an accuracy of .005 with probability .95 (Brooks and Gelman, 1998). To ensure convergence and good coverage of the entire posterior distribution, all chains were run for 100,000 iterations and the first 25,000 iterations were discarded. Samples obtained in the simulation study were thinned 5-fold to save storage space. Standard, vague prior distributions were selected for all parameters as detailed in the Appendix. Web Appendix C contains the BUGS code for the analysis of the Soay sheep data.

2.1.3 *Trinomial Model.* The third approach, presented in Catchpole et al. (2008), avoids the difficulties of missing data by ignoring events that depend on the unobserved values of the covariate. Given that an individual's probability of surviving from occasion t to $t+1$ depends only on the value of the covariate at occasion t , and further that the capture and recovery probabilities are the same for all individuals, the likelihood contribution includes only the events on the occasions directly following the occasions on which the individual is captured and released. Each event has three possible outcomes which are assigned probabilities:

$$P(\omega_{t+1}|\omega_t = 1) = \begin{cases} 1 - \phi(x_t)p_{t+1} - (1 - \phi(x_t))\lambda_t & \omega_{t+1} = 0 \\ \phi(x_t)p_{t+1} & \omega_{t+1} = 1 \\ (1 - \phi(x_t))\lambda_t & \omega_{t+1} = 2 \end{cases}$$

regardless of when else the individual is captured. For this reason, the method has been termed the trinomial approach. An individual with the sample history $\omega_i = 11012$ was captured three times and so three events contribute to the likelihood: the recapture at occasion 2, the failure to either recapture or recover at occasion 3, and the recovery between occasions 4 and 5. The probability assigned to this history by the trinomial model is: $\phi(x_{i1})p_2 \cdot (1 - \phi(x_{i2})p_3 - (1 - \phi(x_{i2}))\lambda_2) \cdot (1 - \phi(x_{i4}))\lambda_4$. This differs from the probability assigned as if all values of the covariate were known in that: 1) there is no contribution from occasion 4 because the probability of this event depends on the unknown value x_{i3} and 2) the contribution from occasion 3 is the same as if the individual had never been observed again, even though it was later captured and must have survived from occasion 2 to 3. The likelihood for this model is once again formed by multiplying the contributions from each individual can easily be maximized to compute MLEs. As with simple imputation, maximization of the trinomial likelihood and computation of the information matrix was performed with the `optim` function in R.

3. Example: The Soay Sheep of Hirta

In this section, we compare the application of these three methods to the analysis of data obtained from the population Soay sheep (*Ovis aries*) on the isle of Hirta, in the St Kilda Archipelego off the west coast of Scotland. The sheep were introduced to the islands of the archipelego by early settlers and when the last residents left in 1932, a population of 107 sheep was moved from the smaller isle of Soay to Hirta where they established a wild population. Because of the isolated nature of the population and the ease in tracking marked individuals, the sheep of Hirta have been studied extensively as a model for understanding population ecology. Since 1986, an intensive observational study of a sub-population of the sheep living on one part of the island (the Village Bay area) has been conducted through physical captures, visual censuses, and searches for dead animals. Here we examine the effect of body mass on the sheeps' survival using the data collected from the sightings and captures of the sheep every August from 1986 to 2000.

Much is known about the complex dynamics of the sheep population, and we direct the interested reader to the papers published in Clutton-Brock and Pemberton (2004). In short, the sheep have no predators and do not compete for food with other species, and the primary cause of mortality is starvation during the late winter months. Resources on the island are not rich enough for the sheep to maintain themselves during the winter, and individuals may lose up to 30% of their body weight depending mostly on the severity of the winter weather conditions and the population density. Every 3 to 4 years the population crashes and then slowly builds again. Smaller sheep are more susceptible to starvation, males have higher mortality rates than females, and younger sheep have higher mortality rates than adults. For simplicity, we restricted to data from the female sheep and modelled the annual probability of survival as a function of body mass separately for each of four age classes: lambs (first year), yearlings (second year), adults (years 3-7), and seniors (8 years and up).

Our analysis is not intended as a comprehensive study of the factors affecting the survival of the sheep and several other variables known to affect survival have been ignored.

Capture histories were constructed from the summer resightings of marked sheep and spring recoveries of dead individuals. After restricting to records of female sheep captured and weighed at least once over the 15 year study period, the data contained capture histories for 802 individuals. The median number of observations for these sheep was 2; 302 sheep (38%) were sighted only once (primarily because of the very high mortality of lambs) and only 1 was sighted in all 15 years. The number of sheep sighted per year ranged from 83 in 1986 to 307 in 1999, with a median of 198. A total of 516 of the sheep (64%) were recovered dead during the study, and the number recovered per year ranged from only 1 in 1987 to 133 in 1999, with a median of 12.

One complication that arose in the analysis was that not all of the sheep sighted in one year were physically captured and weighed. Because of this, body masses for the 802 female sheep were missing for 37% of the observations. The median number of masses recorded per sheep was only 1 (414 (52%) were weighed only once, 142 (18%) 2 times, and 246 (31%) 3 times or more). The missing values when a sheep was sighted but not physically captured do not affect the construction of the likelihood with deterministic imputation or with the Bayesian imputation model, but do affect the trinomial likelihood. For the imputation methods, the occasions when an individual was sighted but not captured and weighed were included in the likelihood by imputing or modelling the missing body mass in exactly the same manner as the unknown body masses when an individual is not resighted. However, only the occasions when an individual was both sighted and captured contribute to the trinomial likelihood function. This essentially reduces the information available for the trinomial method.

Modifications to the Bayesian imputation model were also needed to allow for differences in the growth rates of young and adult sheep. Here we have applied a simple extension

of the diffusion based model in equation (2) that allows separate, time-dependent means and separate variances for the changes in body mass for the sheep in each of the four age categories. The new model is given by:

$$x_t|x_{t-1} \sim N(x_{t-1} + \mu_{A_{t-1},t-1}, \sigma_{A_{t-1}}^2)$$

where A_t indicates the individual's age category in year t .

Figures 1 and 2 compare the estimates of the sighting and recovery probabilities from the three different methods. Point and interval estimates from the linear and Bayesian imputation methods were remarkably similar in all years. In comparison, sighting probabilities estimated from the trinomial model tended to be slightly higher, and were negligibly different from 1.0 in several years. Recovery probability estimates from the trinomial model were also higher than those from the other two models in the majority of years. Most strikingly, the trinomial model estimated the recovery probability in 1990 to be almost 1.0, while the linear and Bayesian imputation models produced estimates near .2.

[Figure 1 about here.]

[Figure 2 about here.]

Figure 3 compares the estimated survival probabilities as a function of body mass for each of the four age categories. All three models indicated that the probability of survival increased with increasing body mass, and that this effect was strongest for the lambs and weakest for the adults. The functions produced by all 3 methods were very similar for both the lambs and yearlings, but the linear imputation model estimated less of an effect on the survival probability of the both the adult and senior sheep. This difference will be discussed further in Section 5.

[Figure 3 about here.]

Estimates of the yearly changes in weight obtained from the Bayesian imputation method

are plotted in Figure 4. In general, the posterior means were similar to direct averages of the observed changes in weight between subsequent years, but were shrunk toward the overall mean. Note that direct averages for some age classes could not be computed in some years, but the hierarchical model still produced estimates for these values. In fact, the model detected a strong decrease in the weight of senior sheep between 1988 and 1989, also evident for the adult sheep, even though no senior sheep were captured in either of these years. These findings are consistent with the previous observations from the Soay sheep study.

[Figure 4 about here.]

4. Simulation Study

Although all three methods produced similar results in Section 4, it is not possible to gauge their performance from the analysis of a single data set. To do so, we present select results from a simulation study comparing estimation of the survival probability coefficients, β_0 and β_1 . Initial simulations were based on the fit of the Bayesian imputation model to the adult female sheep (because this is the only method which provided information about the distribution of weights). Further simulation scenarios were then conducted with modifications to the parameter values or the data generating structure. The results show that all three methods produce very similar estimates when the capture and recovery probabilities are high, though estimates from the trinomial method are slightly less precise. When capture and/or recovery probabilities are low, the linear and Bayesian imputation methods perform better than the trinomial method, provided that the assumed model of the covariate is close to the truth. If the model of the covariate is not close to the truth, then both the linear and Bayesian methods produce estimates of β_0 and β_1 that are severely biased.

In each of the simulation scenarios, we generated 500 data sets comprising capture histories of $T = 10$ occasions for $n = 500$ individuals. Results for one scenario with $n = 1000$

individuals are provided to assess how the sample size affects the different methods. Web Appendix A contains details of the method for generating capture histories. In short, we assigned the occasion of first capture for each individual deterministically to ensure balance in the data and simulated the corresponding covariate from a normal distribution with constant mean and variance, μ_0 and σ_0^2 . We then generated the remaining values of the covariate from the diffusion model in equation (2) or some variation, simulated the individual's survival conditional on the covariate values, and finally simulated capture and recovery events conditional on survival.

Comparisons of the methods were based the distributions of point and interval estimates for β_0 and β_1 over the 500 data sets simulated in each scenario. Specific criteria considered were the relative bias of the point estimates $((\hat{\beta} - \beta)/\beta)$, and the coverage probability and average width of the 95% interval estimates. Maximum likelihood estimates and symmetric 95% confidence intervals based on the asymptotic distribution were computed for the imputation and the trinomial methods. Posterior means and symmetric 95% credible intervals were computed for the Bayesian imputation method. Results are provided in Table 1.

[Table 1 about here.]

4.1 Reference Scenario

Parameter values for the first scenario were based directly on the estimates obtained from the Bayesian imputation analysis of the adult female sheep. To simplify the generating model, we ignored the difference between captures and resights: capture probabilities for the simulations were set equal to the resight probabilities for the Soay sheep and data were generated as if the covariate was recorded every time an individual was captured. Table 2 lists the values for the time dependent parameters. Both the capture and recovery probabilities were relatively high ($p_t > .90$ and $\lambda_t > .50$ on 8 of 9 occasions) and the mean changes in the covariate from one year to the next were small relative to the standard deviation (i.e., $\mu_t/\sigma < 1$ in

all but two years). The remaining values needed for generating data were the parameters of the initial covariate distribution, $\mu_0^a = 22.90$ and $\sigma_0^{2a} = 8.40$, the variance of the changes in the covariate between occasions, $\sigma^{2a} = 2.99$, and the coefficients of the logistic survival probabilities, $\beta_0^a = -3.18$ and $\beta_1^a = .25$. The superscript a is used to indicate that these values were obtained from the analysis of the adult female sheep.

All three methods performed well in this scenario. The only noticeable differences were that trinomial methods produced slightly biased point estimates and interval estimates that were slightly wider than those from the linear and Bayesian imputation methods. However, these discrepancies are not large enough to be of biological concern. Summary statistics for the linear and Bayesian imputation methods differed by at most .02.

[Table 2 about here.]

4.2 Scenarios with Capture and Recovery Probabilities

The next three scenarios compared the performance of the methods with lower capture or recovery probabilities, as noted in Table 3. In Scenario 2, the capture probabilities were reduced to one third of the values estimated for the adult female sheep. This decreased the average number of captures per individual, which increased the uncertainty of the estimates for all three methods (i.e., widened the interval estimates). The effect was much larger for the trinomial model than for the other two models; interval estimates increased in width by approximately 1.8 times in comparison to approximately 1.2 times for both the linear and Bayesian imputation methods. Moreover, point estimates of β_0 and β_1 from the trinomial model were biased by approximately -20% and +5%, respectively, which lead to underestimation of the survival probabilities for all values of the covariate. The linear imputation method also produced slightly biased estimates (less than 5%) for both β_0 and β_1 . Despite these biases, coverage proportions for all three methods were greater than 92%, though lowest for the Bayesian imputation method for both β_0 and β_1 .

[Table 3 about here.]

To study the impact of sample size, we repeated Scenario 2 by simulating 100 data sets with the same parameter values but with the number of marked individuals increased from 500 to 1000 (Scenario 2b in Table 3). This greatly reduced the bias of the estimates from the trinomial method, but not from the linear imputation method. Interval estimates were also narrower for all three methods.

Decreasing only the recovery probability, Scenario 3, had very little effect on either the bias or precision of the linear and Bayesian imputation methods. Point estimates differed almost negligibly from Scenario 1 and the width of the interval estimates increased by less than 1.10. In contrast, interval estimates for the trinomial model increased in width by a factor of approximately 1.3. No significant biases were observed for any of the methods and the coverage probabilities for the interval estimates were again greater than 92% for all 3 methods, but smallest for the Bayesian imputation method.

4.3 Scenarios with Modified Generating Model

In the final two scenarios (Scenario 4 and 5), we modified the model for generating the covariate so that it no longer matched the diffusion model assumed by the Bayesian imputation method. In particular, the expected change in the covariate on each occasion was modelled as a linear function of the current value for each individual:

$$E(x_{it} - x_{i,t-1} | x_{i,t-1}) = \mu_t + b(x_{i,t-1} - E(x_{i,t-1})).$$

To select values of the slope parameter, b , we imposed the constraint that $Var(x_{it}) = Var(x_{i,t+1})$, satisfied if $b = -.20$ or $b = -1.80$ (see Web Appendix B for details).

Setting $b = .20$ produced a slight regression toward the mean such that $|x_t - E(x_t)| < |x_{t-1} - E(x_{t-1})|$ on average. This represents a minor departure from the assumptions of the Bayesian imputation model, and the changes from the results of Scenario 2 were relatively small for both the trinomial and Bayesian imputation methods. However, point estimates

from the linear imputation method were severely biased, and coverage probabilities of the interval estimates were greatly reduced.

Setting $b = -1.80$ produced oscillations so that $x_t - E(x_t)$ tended to decrease in absolute value from one occasion to the next and also to switch sign. This presented a more severe departure from the assumptions of the Bayesian imputation method and the impacts on both the linear imputation and Bayesian imputation models were extreme. Point estimates for both models were strongly biased and coverage probabilities of the interval estimates were very low. In comparison, results for the trinomial were almost the same as in Scenario 2 with bias of 20% for the point estimate of β_0 and interval estimates that were wider than those of the other two methods on average, but in this case had much higher coverage probabilities.

5. Discussion

The results of the simulation study in Section 4 clearly demonstrate the advantages and disadvantages of the different methods for handling time-varying, individual covariates, an important outstanding problem in the analysis of mark-recapture-recovery data. Deterministic imputation is simple to implement but has the potential to produce severely biased estimates if the assumed imputation model does not match the true data generating mechanism. This is most pronounced in the final simulation scenario in which the relative biases of the parameters β_0 and β_1 were 1.29 and -.66. Further simulations indicated that the bias increased as capture probabilities decreased and a larger proportion of values had to be imputed.

Modelling the joint distribution of the covariate in the Bayesian imputation method allows for some uncertainty in the missing values, within the variation of the selected model. One result is that the model is more robust to the specification of the imputation model than the deterministic imputation method, as seen in Scenario 4. In comparison with the deterministic imputation method, the Bayesian imputation model produced unbiased estimates and the

coverage of the interval estimates remained close the nominal value. However, the Bayesian imputation method still performs poorly if the selected model of the covariate is a poor approximation to the data generating model, as in the final simulation scenario. This clearly highlights the importance of selecting a good model of the covariate and the need for goodness-of-fit methods to test the model's fit.

Another advantage of the Bayesian imputation method is that it provides inference about the distribution of the covariate. For the Soay sheep, the hierarchical model shows that, averaged over all years, the growth of lambs was higher than the growth of yearlings, adults had a negligible change in mass, and the seniors lost weight (though the 95% credible interval for the hierarchical mean of μ_t does cover 0). We also see that there has been considerable variation in the mean change in body mass from year to year, with recent years being more favourable for all of the age classes.

As with any Bayesian analysis, the results of the Bayesian imputation method may be depend on the specification of the prior distributions. Several alternative prior specifications for β_0 and β_0 were tested in the simulation study, and we found that the results were fairly robust to these choices. Although the summary statistics changed slightly, the differences in the estimated survival probabilities were almost imperceptible. The prior distributions used in Sections 3 and 5 were selected because they produced unbiased posterior means and close to nominal coverage of the credible intervals when the Bayesian imputation model matched the data generating model (i.e., simulation scenarios 1,2, and 3).

The trinomial model avoids all assumptions regarding the covariate by ignoring any event that depends upon an unobserved value. The advantage of this is clear in the final simulation scenario; results from the trinomial method are almost identical to those of Scenario 2 despite the poor performance of the other two methods. Another advantage of the trinomial method is that it is easily implemented in standard software packages. The trinomial model can in

fact be implemented in the popular software **Program MARK** (White and Burnham, 1999) by using the “Both (Burnham)” data type for recapture/recovery experiments with properly formatted data and parameter settings. Instructions and code for running the model through **RMark**, the R interface to **Program MARK**, is available from the first author.

The main drawback of the trinomial model is that it potentially disregards a large amount of information in the data. This is particularly true when the capture probabilities are low such that even long lived individuals are rarely observed on two consecutive occasions. Further to this, the trinomial model depends much more on the recovery of dead individuals, as seen in Scenario 3. Our intuitive reasoning is that the trinomial requires recoveries of dead individuals to distinguish the individuals released on one occasion that die before the next occasion from those that simply evade recapture. In comparison, the linear and Bayesian imputation models make use of the information in future captures to distinguish between these individuals.

One surprising result from our simulation study was the bias of the trinomial model when the capture probabilities were low (Scenario 2). This effect was not observed in the simulation study of Catchpole et al. (2008), but their simulations were conducted with relatively high capture and recovery probabilities (.8 and .4 respectively). From the results of Scenario 1 and the comparison between Scenarios 2 and 2b we conclude that the bias is caused by small sample size – a small number of events contributing to the trinomial model likelihood either because few individuals were marked or because the numbers of recaptures per marked individual were low. Catchpole et al. (2008) have confirmed that the method is consistent in that parameter estimates do tend to the true values as the sample size increases to infinity.

If capture and recovery probabilities are both high and the animals being studied are long lived, it makes little difference which method is chosen to incorporate the effect of a time-varying, individual covariate on the probability of survival. Individuals will be captured

many times on average, their fates will be known with little uncertainty, and the proportion of missing data will be low. In this case, it matters little how the missing data are handled.

The Bayesian imputation method is to be preferred when capture and recovery probabilities are low, but only if a reasonable model of the covariate is available. Provided that the model is accurate, estimates will be much more precise than those of the trinomial method and will be less biased than estimates from deterministic imputation. However, we strongly recommend that some method be used to check the fit of the selected model to the observed data because of the potential for incorrect inference if the assumed model does not approximate truth (see for example Meng, 1994; Brooks et al., 2000, on the use of Bayesian p-values). Because of the bias introduced at small sample sizes, the trinomial model should be used when capture probabilities are low only if no adequate model of the covariate can be found.

In all cases, we recommend against deterministic imputation. If the capture probabilities are high and there are few missing values, then the trinomial model is simple to implement and provides unbiased inference without any assumptions regarding the covariate. If the proportion of missing values is high and a model of the covariate must be assumed, then the Bayesian imputation method will be less biased and produce better estimates of uncertainty.

We conclude by returning to the estimated recovery and survival probabilities for the female Soay sheep of Hirta. As noted in Section 3, the linear imputation model indicated less of an effect of body mass on the survival of the adults and seniors (see Figure 3). This occurred because the linear imputation model failed to account for the decrease in the average weight of the older sheep between 1988 and 1989. Although 60 adult sheep were captured in 1989 only 15 (25%) were weighed, and the values imputed for the remaining 45 sheep ignored the decrease. Repeating the analysis using the mean changes in weight estimated from the Bayesian imputation model to guide imputation produced estimates of the survival probabilities that matched the other two methods almost exactly (results not shown).

The second discrepancy was that the trinomial model estimated the recovery probability in 1990 to be very close to 1.0 whereas the other methods produced estimates less than .20 (see Figure 2). This occurred because the number of sheep captured and weighted in 1989 was very small. Although 112 sheep were resighted in the summer of 1989, only 27 events contributed to the trinomial likelihood. Of these, 24 were resighted the next year, 2 were found dead, and only one, a relatively large lamb, went unobserved. Given the high survival probability for large lambs, the most likely scenario is that the last individual survived but was not captured. It then appears that only 2 of the 27 sheep died and both were recovered, so the recovery probability is estimated to be near 1.0. The linear and Bayesian imputation model included the events for all 112 sheep resighted in 1989, and all other sheep marked previously marked and not yet recovered, and so produced much more reasonable estimates of the recovery probability for this year.

Supplementary Materials

The Web Appendices referenced in Sections 2 & 4 are available under the Paper Information link at the *Biometrics* website <http://www.biometrics.tibs.org>.

Acknowledgments

This work was initiated while Simon Bonner was a PhD student at Simon Fraser University and completed after he began a post-doctoral position at the University of British Columbia. During this time, Simon was funded by the National Science and Engineering Research Council of Canada. Ruth King was supported by grants from the Engineering and Physical Sciences Research Council, the Natural Environment Research Council, and the Wellcome Trust. Byron Morgan was the recipient of a Leverhulme Research Fellowship.

We thank all of the past and current members of the Soay sheep project for allowing the use of the data collected from the sheep of St. Kilda, the National Trust for Scotland and the

Scottish Natural Heritage who have supported this project, and the army units stationed on St. Kilda and Benbecula that provided assistance to the project. We are particularly grateful to Tim Coulson who answered our many questions about the Soay sheep and Ted Catchpole who provided the formatted data.

References

- Arnason, A. N. (1973). The estimation of population size, migration rates, and survival in a stratified population. *Research in Population Ecology* **15**, 1–8.
- Bonner, S. J. and Schwarz, C. J. (2006). An extension of the Cormack-Jolly-Seber model for continuous covariates with application to *Microtus pennsylvanicus*. *Biometrics* **62**, 142–149.
- Brooks, S. P., Catchpole, E. A. and Morgan, B. J. T. (2000). Bayesian animal survival estimation. *Statistical Science* **15**, 357–376.
- Brooks, S. P. and Gelman, A. (1998). General methods for monitoring convergence of iterative simulations. *Journal of Computational and Graphical Statistics* **7**, 434–455.
- Catchpole, E. A., Morgan, B. J. T. and Tavecchia, G. (2008). A new method for analysing discrete life-history data with missing values. *Journal of the Royal Statistical Society – Series B* **70**, 445–460.
- Clutton-Brock, T. H. and Pemberton, J. M., editors (2004). *Soay Sheep: Dynamics and Selection in an Island Population*. Cambridge University Press, Cambridge.
- King, R., Brooks, S. P. and Coulson, T. (2008). Analyzing complex capture-recapture data in the presence of individual and temporal covariates and model uncertainty. *Biometrics* **64**, 1187–1195.
- Lebreton, J.-D., Burnham, K. P., Clobert, J. and Anderson, D. R. (1992). Modelling survival and testing biological hypotheses using marked animals: A unified approach with case studies. *Ecological Monographs* **62**, 67–118.

- Little, R. J. A. and Rubin, D. B. (2002). *Statistical Analysis with Missing Data*. John Wiley, New York, 2nd edition.
- Meng, X.-L. (1994). Posterior predictive p-values. *The Annals of Statistics* **22**, 1142–1160.
- Nichols, J. D., Sauer, J. R., Pollock, K. H. and Hestbeck, J. B. (1992). Estimating transition probabilities for stage-based population projection matrices using capture-recapture data. *Ecology* **73**, 306–312.
- Pollock, K. H. (2002). The use of auxiliary variables in capture-recapture modelling: an overview. *Journal of Applied Statistics* **29**, 85–102.
- R Development Core Team (2009). *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria. ISBN 3-900051-07-0.
- Schwarz, C. J., Schweigert, J. F. and Arnason, A. N. (1993). Estimating migration rates using tag-recovery data. *Biometrics* **49**, 177–193.
- Thomas, A., O’Hara, B., Ligges, U. and Stutz, S. (2006). Making BUGS open. *R News* **6**, 12–17.
- White, G. C. and Burnham, K. P. (1999). Program MARK-survival estimation from populations of marked animals. *Bird Study* **46 Supplement**, 120–138.

31 October 2009

Appendix

Prior Distributions for the Bayesian Imputation Model

$$\mu_t \sim \text{Normal}(0, 10^6), \quad t = 1, \dots, T - 1$$

$$\tau \sim \text{Inverse Gamma}(.1, .1)$$

$$\beta_k \sim \text{Normal}(0, 10^6), \quad k = 0, 1$$

$$p_t \sim \text{Uniform}(0, 1), \quad t = 2, \dots, T$$

$$\lambda_t \sim \text{Uniform}(0, 1), \quad t = 2, \dots, T$$




Figure 1. Estimated yearly resighting probabilities for the female Soay sheep of Hirta. Points represent the ML estimates for the linear imputation and trinomial models, and posterior means for the Bayesian imputation method. Error bars indicate the extents of the 95% confidence/credible interval estimates.

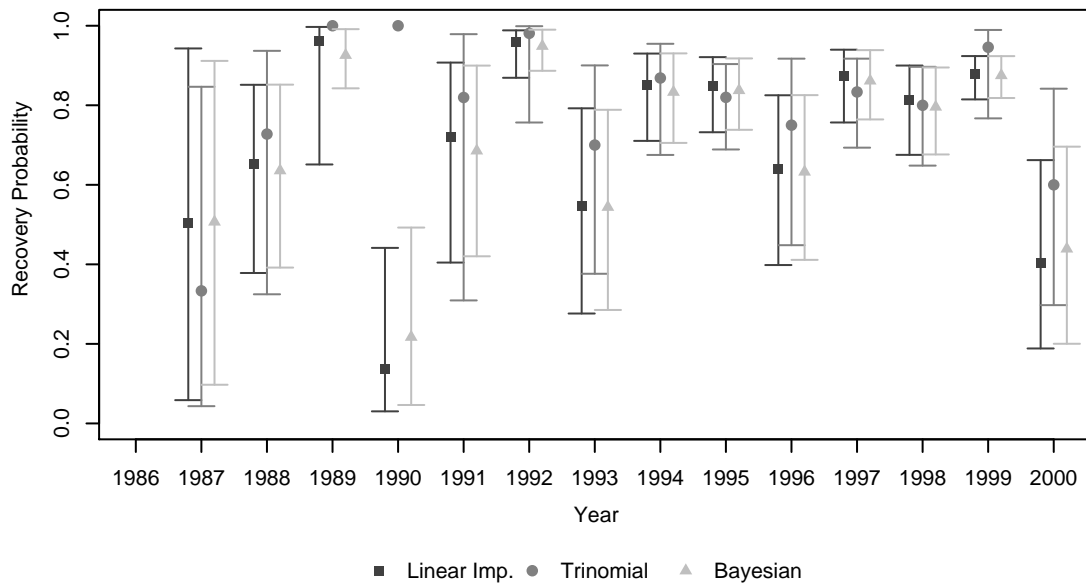


Figure 2. Estimated yearly recovery probabilities for the female Soay sheep of Hirta. Points represent the ML estimates for the linear imputation and trinomial models, and posterior means for the Bayesian imputation method. Error bars indicate the extents of the 95% confidence/credible interval estimates.

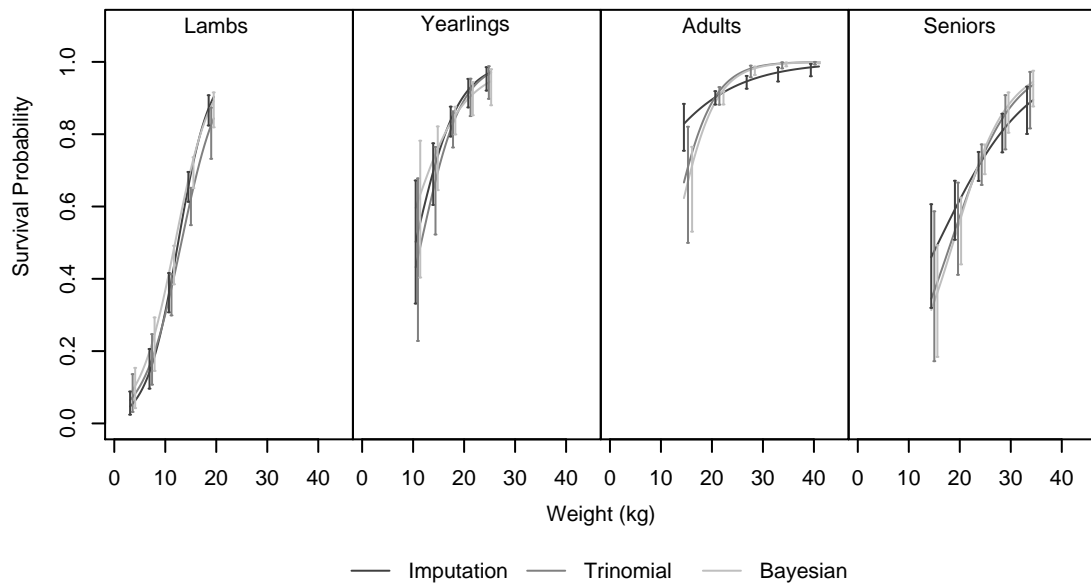


Figure 3. Estimated survival probabilities as a function of age and body mass for the female Soay sheep of Hirta. Lines represent the ML estimates for the linear imputation and trinomial models and the posterior means for the Bayesian imputation method. Error bars indicate the extents of the pointwise 95% confidence/credible interval estimates at selected masses.

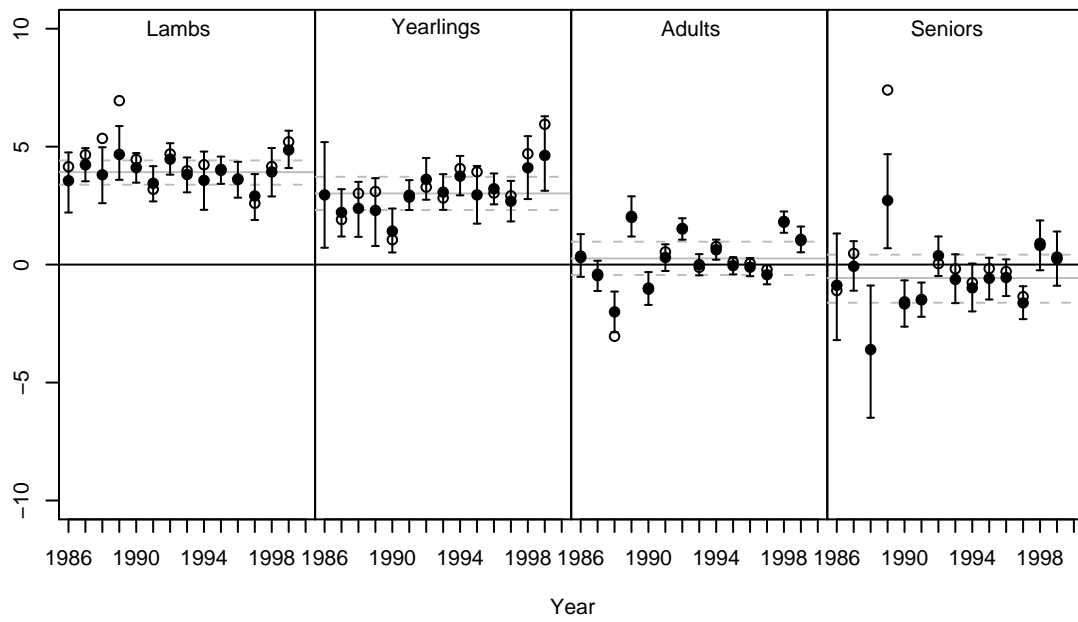


Figure 4. Estimates of the yearly mean change in mass by age for the female Soay sheep of Hirta. Open points represent the observed mean change and closed points the estimate obtained from the Bayesian imputation method with error bars extending to the limits of the 95% credible interval. The solid grey lines represent the hierarchical mean for each age class with 95% credible interval bounded by the dashed lines.

Table 1

Summary statistics for the estimation of the coefficients of the survival probability linear predictor, β_0 and β_1 . Columns in the table are: SS= Simulation Scenario, Mo=Model, ME=Mean Estimate, RB=Relative Bias (with 95% CI indicated by subscripts), CW=CI Width, and CC=CI Coverage. Models are specified as: L=linear imputation, T=trinomial method, and B=Bayesian.

		Intercept ($\beta_0 = -3.18$)						Slope ($\beta_1 = .25$)					
SS	Mo	ME	RB			CW	CC	ME	RB			CW	CC
1	L	-3.22	-.03	-.01	.00	2.15	.95	.26	.00	.01	.01	.10	.94
	T	-3.29	-.05	-.04	-.02	2.37	.94	.26	.01	.02	.03	.12	.94
	B	-3.22	-.03	-.01	.00	2.13	.95	.26	.00	.01	.02	.10	.93
2	L	-3.09	.01	.03	.05	2.59	.96	.24	-.06	-.05	-.03	.12	.93
	T	-3.77	-.22	-.19	-.15	4.48	.94	.27	.04	.05	.07	.20	.95
	B	-3.19	-.02	.00	.02	2.62	.92	.25	-.02	-.01	.00	.13	.92
2b	L	-2.99	.03	.06	.09	1.82	.93	.24	-.08	-.06	-.04	.09	.86
	T	-3.38	-.11	-.06	-.02	3.21	.95	.26	-.02	.00	.03	.14	.97
	B	-3.10	.00	.03	.06	1.82	.92	.25	-.04	-.02	.00	.09	.90
3	L	-3.16	-.01	.01	.02	2.22	.95	.25	-.02	-.01	.00	.11	.95
	T	-3.23	-.03	-.01	.01	3.06	.97	.26	.00	.01	.02	.16	.96
	B	-3.21	-.03	-.01	.01	2.20	.93	.26	.00	.01	.02	.11	.92
4	L	-2.12	.32	.33	.35	2.35	.56	.20	-.23	-.22	-.21	.11	.46
	T	-3.60	-.16	-.13	-.11	4.00	.94	.26	.02	.04	.05	.18	.94
	B	-3.46	-.11	-.09	-.07	2.62	.90	.27	.04	.05	.06	.13	.90
5	L	.65	1.27	1.29	1.31	2.19	.00	.07	-.67	-.66	-.65	.10	.00
	T	-2.69	-.23	-.20	-.18	3.65	.94	.22	.03	.05	.06	.16	.94
	D	-1.11	.47	.50	.54	2.49	.48	.16	-.26	-.24	-.22	.12	.51

Table 2

Base parameter values for the simulation study. These values were obtained from the fit of the Bayesian model to the analysis of the adult female sheep.

	Capture Occasion (t)									
	1	2	3	4	5	6	7	8	9	10
p_t^a	·	0.96	0.94	0.83	0.95	0.94	0.95	0.97	0.91	0.92
λ_t^a	·	0.50	0.64	0.93	0.21	0.69	0.95	0.54	0.83	0.84
μ_t^a	0.36	-0.47	-2.00	2.05	-0.99	0.30	1.51	0.00	0.63	·

Table 3

Parameter values for the six simulation scenarios. Superscript *a* denotes a parameter value taken from the results of the analysis for the adult female sheep, as given in Table 2. Parameter values in the first simulation were all set equal to their estimates for the adult female sheep. Bold text in the following rows highlights the parameter values that have been modified.

Description	μ_0	μ_{it}	p_t	λ_t
1) Base	μ_0^a	μ_t^a	p_t^a	λ_t^a
2) Reduced capture prob.	μ_0^a	μ_t^a	$p_t^a/3$	λ_t^a
3) Reduced recovery prob.	μ_0^a	μ_t^a	p_t^a	$\lambda_t^a/3$
4) Covariate regression 1	μ_0^a	$\mu_t^a - .20(z_{i,t-1} - E(z_{i,t-1}))$	$p_t^a/3$	λ_t^a
5) Covariate regression 2	μ_0^a	$\mu_t^a - 1.80(z_{i,t-1} - E(z_{i,t-1}))$	$p_t^a/3$	λ_t^a