Persistence models for mark-recapture

Fred L. Ramsey · Paul M. Severns

Abstract   The stable of models available for analyzing mark-recapture data (Otis et al. Wild Momogr 66:135, 1978) includes those having behavioral characteristics, time variation, heterogeneity, along with combinations of those characteristics. This paper proposes use of a series of models based on the persistence model of Ramsey and Usner (Biometrics 59:331–339, 2003). We show that persistence can be modeled in combination with behavior and with time variation. We apply the persistence model to situations in which capture occasions are not equally-spaced in time. Two case studies illustrate the use of these extended persistence models.

1 Introduction

The class of models available for mark-recapture studies of closed populations includes the equal catchability model (M0), the behavior model (Mb), the time model (Mt), the heterogeneity model (Mh), and various combinations of those features (Otis et al. 1978; White et al. 1982). This paper suggests including a persistence model (Mp) - a Markov-chain model in which the conditional probability of capturing an animal on one occasion, given the animal's past capture history, depends only on whether the animal was or was not captured on the previous occasion. Mp is a two-state

F. L. Ramsey (✉)
Department of Statistics, Oregon State University, Corvallis, OR, USA
e-mail: flramsey5@comcast.net

P. M. Severns
Department of Botany and Plant Pathology, Oregon State University, Corvallis, OR, USA

chain which is a special case of the more general persistence model suggested by Ramsey and Usner (2003) for use in radio tracking studies. Section 2 presents the basic persistence model for mark-recapture. Subsequent sections describe extensions designed to account for departures from three features that the basic persistence model possesses: (i) the steady-state (unconditional) probability of capture of an animal is the same on all occasions; (ii) animals have diminishing memories of past captures (implicit in the Markovian model); and (iii) capture occasions are equally-spaced in time. Combinations Mpt of persistence and time and Mpb of persistence and behavior are proposed (Sect. 3), and they are applied to a mark-recapture study of a recently discovered population of Great Copper butterflies (*Lycaena xanthoides*) in Oregon's Willamette Valley (Sect. 5). Then we show how the basic persistence model can be applied to studies with unequally-spaced capture occasions (Sect. 6), using mark-recapture data from a three-year study of a population of wetland snails (Sect. 7). Section 4 provides a short description of our model-fitting strategy.

2 The basic persistence model

A trapping study of a closed population estimates the population size, N, of animals, which are here indexed by the subscript i. There are T trapping occasions, indexed by the subscript t. The variable $x_{it}(= 1/0)$ indicates whether animal i is captured on occasion t. The notation $\boldsymbol{x}_i(t) = (x_{i1}, \ldots, x_{it})$ is used to denote the capture history of animal i up to and including occasion t; so $\boldsymbol{x}_i(T)$ is the animal's full capture history. Let p_{i1} be the probability (unconditional) of animal i being captured on the first occasion:

$$p_{i1} = \Pr\{x_{i1} = 1\}.$$

And for $t > 1$, let p_{it} be the conditional probability of animal i being captured on occasion t, given the animal's prior capture history; i.e.

$$p_{it} = \Pr\{x_{it} = 1 | \boldsymbol{x_i}(t - 1)\}.$$

With this definition, $P_{i1} P_{i2} \ldots P_{iT}$ is the probability of animal i's full capture history. There are $M = 2^T$ different capture histories, and these will be indexed by the subscript j, with $j = 1$ denoting the history of animals that are not captured on any occasion. We assume that animals' capture histories are determined independently and according to the same probabilities. The result is that the probability associated with the full data set, $X = \{\boldsymbol{x}_i(T) : i = 1, \ldots, N\}$, is multinomial.

Let m_j be the number of animals that have the jth capture history. Note that $m_2 + \cdots + m_M = n$ is the number of individual animals captured at least once, while $m_1 + m_2 + \ldots m_M = N$. That is, $m_1 = N - n$. And,

$$\Pr\{X\} = \prod_{j=1}^{M} P_j^{mj} / mj! = [n!/(m_2! \ldots m_M!)] \binom{N}{n} \prod_{j=1}^{M} P_j^{mj}$$

For the purposes of parameter estimation and model selection, the first factor in the right-hand expression is ignored.

In Mp, the parametric structure has the following form.

$$p_1 = \pi$$

and for $t > 1$,

$$
\begin{aligned}
p_t &= \theta\pi, && \text{if } x_{t-1} = 0 \\
&= 1 - \theta(1 - \pi), && \text{if } x_{t-1} = 1.
\end{aligned}
$$

This defines a two-state Markov chain model with a single parameter, θ, measuring the tendency of animals to repeat their previous state. If $\theta < 1$, animals captured on the previous occasion tend to be captured again on the current occasion and animals not captured on the previous occasion tend not to be captured on the current occasion.

Section 5 will discuss an application of the persistence model to a mark-recapture study of butterflies which illustrates a biological interpretation of persistence. While flowers are producing nectar, butterflies tend to revisit them and be relatively easy to recapture. When the flowers no longer produce nectar, the butterflies leave them for other sources so that non-capture tends to be followed by non-captures. The persistence is not a behavioral response to capture. And unlike the behavioral model, the Markovian nature of the persistence model entails a diminishing influence with time of past conditions (see Sect. 4). The persistence model can be reparameterized in logit form:

$$\text{logit}(p_t) = \beta_0 + \beta_1 x_{t-1}.$$

However, it differs from the Markov model MM1(b) of Yang and Chao (2005). In MM1(b), the unconditional probability of capture either increases steadily to an asymptotic value as t increases or fluctuates about a decrease to an asymptotic value, depending on the difference between the probabilities of recapture following capture or non-capture. This model is oriented more specifically to a behavioral response. It also requires that the unconditional probability of capture on the first occasion is the same as the conditional probability of capture following a previous non-capture. Model Mp has constant unconditional probability of capture and may or may not be oriented specifically to a behavioral response.

The log-likelihood function for the parameters of the persistence model is

$$
\begin{aligned}
\log(L) = {}& \log\binom{N}{n} + (C1 - K01)\log(\pi) + (N - C1 + K10)\log(1 - \pi) \\
& + (K01 + K10)\log(\theta) + [(N - n)(T - 1) + K00]\log(1 - \theta\pi) \\
& + K11\log[1 - \theta(1 - \pi)]
\end{aligned}
$$

where $C1$ is the number of captures on the first occasion and where Kab is the total number of transitions from state a to state b by all animals that are captured at least once.

3 Persistence with time variation and with behavior

3.1 Persistence and behavior

Combining Mp with Mb yields model Mpb, where the parametric structure is

$$p_1 = \pi$$

and for $t > 1$,

$$
\begin{aligned}
p_t &= \theta\pi, & &\text{if } M_t = \max(x_s : s < t) = 0 \\
&= \theta\gamma, & &\text{if } x_{t-1} = 0 \text{ and } M_{t-1} = 1 \\
&= 1 - \theta(1 - \gamma), & &\text{if } x_{t-1} = 1
\end{aligned}
$$

The log-likelihood function for Mpb is

$$
\begin{aligned}
\log(L) = {}& \log\binom{N}{n} + n\log(\pi) + (N - C1)\log(1 - \pi) + (K01 + K10)\log(\theta) \\
& + L01\,\log(\gamma) + K10\log(1 - \gamma) + [(N - n)(T - 1) + J00)]\log(1 - \theta\pi) \\
& + L00\log(1 - \theta\gamma) + K11\,\log[1 - \theta(1 - \gamma)]
\end{aligned}
$$

where

$C1 =$ number of animals captured on the first trapping occasion
$Jab =$ number of a-to-b transitions prior to initial capture, by all animals
$Kab =$ number of a-to-b transitions by all animals and
$Lab =$ number of a-to-b transitions after initial capture, by all animals.

3.2 Persistence and time

Combining Mp with Mt gives Mpt, where the parametric structure is

$$p_1 = \pi_1$$

and for $t > 1$,

$$
\begin{aligned}
p_t &= \theta\pi_t, & &\text{if } x_{t-1} = 0 \\
&= 1 - \theta(1 - \pi_t), & &\text{if } x_{t-1} = 1
\end{aligned}
$$

The log-likelihood function for Mpt is

$$
\begin{aligned}
\log(L) = {}& \log\binom{N}{n} + C1\log(\pi_1) + (N - C1)\log(1 - \pi_1) + [K01 + K10]\log(\theta) \\
& + \sum\{K01t\log(\pi_t) + K10t\log(1 - \pi_t) + K11t\log[1 - \theta(1 - \pi_t)] \\
& + [N - n + K00t]\log(1 - \theta\pi_t)\}
\end{aligned}
$$

where the sums run from $t = 2$ to $t = T$ and where $Kabt$ is the total number of transitions of all animals from a on occasion $t - 1$ to b on occasion t. The other statistics are as defined above.

4 Model fitting

In the studies that follow, the above models admit straightforward likelihood solutions. For each $N \geq n$, Newton-Raphson iterations maximize the likelihood with respect to probability parameters. Convergence is rapid. Substitution yields a profile likelihood for N. We report the maximizing values of N along with intervals containing 95% of the normalized profile likelihood (NPL). Estimates and standard errors for the probability parameters incorporate uncertainty about N. If, for example, $\theta(N)$ and $V(\theta|N)$ are maximum likelihood estimate of the parameter θ and its estimated variance for fixed N, then we estimate the parameter by

$$\theta^* = \sum_{N \geq n} \mathrm{NPL}(N)\theta(N)$$

and its variance by

$$\mathrm{Var}(\theta^*) = \sum_{N \geq n} \mathrm{NPL}(N)[\theta(N) - \theta^*]^2 + \sum_{N \geq n} \mathrm{NPL}(N)V(\theta|N).$$

This methodology incorporates uncertainty about N in a manner similar to the model-averaging procedures of Burnham and Anderson (2002).

5 Great Copper study

In the summer of 2004, Paul Severns discovered in the southern Willamette Valley of Oregon a population of Great Copper butterflies (*Lycaena xanthoides*)- a species absent from the valley for some 35 years (Severns and Villegas 2005). In 2005, he conducted a mark-recapture study on $T = 8$ occasions, every third day through mid-July. On each occasion, he walked through the wetland habitat, netting butterflies and marking their ventral hindwings with individual codes. Capture sites were selected where adult butterflies were most abundant. Capture effort was standardized across dates and sites by time-constrained searches. Great Coppers are swift fliers with a keen sense of predators. Individuals elude capture if possible, but marked individuals tended to be netted again because their locations had been noted.

Darwin (1876) hypothesized that nectar-feeding insects would exhibit flower constancy—the tendency to return to plants of the same species. Lewis (1986) demonstrated flower constancy for Cabbage White butterflies (*Pieris rapae*). Coppers have a tight relationship with their preferred nectar sources. They congregate near favorable habitat, which is readily recognized by the observer (Severns et al. 2006). If both butterfly and researcher return to the same plants, one would expect persistence in the

capture histories. This persistence would leave open the possibility that a behavioral response to capture also exists.

Forty-five (n) butterflies were netted at least once. Appendix A displays their capture histories. Results of fitting M0, Mb, Mt, Mp, Mpb, and Mpt appear in Table 1.

Based upon the maximum log-likelihood values, Mp is preferred to Mb, and Mt, which have at least as many parameters. M0 is a special case of Mp; the likelihood ratio test chi-square statistic (13.88, with 1 degree of freedom) indicates Mp is also preferred to M0 (p-value $= .0002$). On the other hand, Mp is a special case of both Mpb and Mpt, so likelihood ratio tests apply. Mpb provides no substantial improvement (p-value $= 0.39$ from chi-square $= 0.75$ with 1 d.f.); nor does Mpt (p-value $= 0.84$ from chi-square $= 3.47$ with 7 d.f.). If, then, one is to choose the 'best fitting' of the models we consider here, it would be the basic persistence model. The estimate of the persistence effect, is 0.79 (± 0.12). This is significant persistence, with the result that the probability of capture following no capture is estimated to be 0.632 and the probability of capture following capture is estimated to be 0.842.

6 Persistence with unequally-spaced trapping occasions

The basic persistence model has the following transition probability matrix (TPM).

$$\prod = \begin{bmatrix} [\Pr\{x_t = 0 | x_{t-1} = 0\} & \Pr\{x_t = 1 | x_{t-1} = 0\}] = [1 - \theta\pi & \theta\pi] \\ [\Pr\{x_t = 0 | x_{t-1} = 1\} & \Pr\{x_t = 1 | x_{t-1} = 1\}] & [\theta(1 - \pi) & 1 - \theta(1 - \pi)] \end{bmatrix}$$

Assume that this TPM describes an underlying sequence of equally-spaced potential capture occasions. The unequally-spaced occasions study can be viewed as an irregular sampling of the underlying sequence. A convenient choice of time unit is the greatest common divisor of all actual spacings.

If $\prod_0$ is the TPM for the equal catchability model ($\theta = 1$), we can write the persistence TPM as follows:

$$\prod = \prod_0 + (1 - \theta)\left(I - \prod_0\right)$$

This is the spectral decomposition of the TPM. Therefore, the d-step TPM is

$$\prod_d = \prod_0 + (1 - \theta)^d \left(I - \prod_0\right)$$
$$= \begin{bmatrix} [1 - \pi + \pi\xi^d & \pi(1 - \xi^d)] \\ [(1 - \pi)(1 - \xi^d) & \pi + (1 - \pi)\xi^d] \end{bmatrix}$$

where $\xi = 1 - \theta$.

Let dt be the number of time units between occasion $t - 1$ and occasion t. Thus $t = 1$ is the first occasion; $t = 2$ is the occasion d2 time units later; $t = 3$ is the occasion

Table 1 Analysis summaries for the Great Copper study

Model	Est N	Interval	log(Lmax)	Capture	Probabilities							Others
M0	64	(58, 96)	−169.40	.133 (.023)	.196							
Mb	62	(53, 762)	−169.39	.057 (.054)								γ: 0.137 (0.025)
Mt	64	(53, 88)	−167.42	.120	.196	.136	.136	.151	.136	.105	.090	
Mp	97	(70, 215)	−162.46	.080 (.025)								θ: 0.790 (0.059)
Mpb	69	(60,1006)	−162.08	.044 (.050)								γ: 0.081 (0.026) θ: 0.787 (0.059)
Mpt	64	(54, 103)	−160.73	.117 (.043)	.269 (.072)	.155 (.067)	.181 (.062)	.198 (.066)	.181 (.062)	.143 (.055)	.113 (.053)	θ: 0.754 (0.072)

d3 time units after that; and so on. The observed capture history process maintains a Markovian character; i.e.

$$\Pr\{x_t | x_{t-1}, x_{t-2}, \ldots, x_1\} = \Pr\{x_t | x_{t-1}\}.$$

The process is no longer time homogeneous, as $\Pr\{x_t | x_{t-1}\}$ comes from the dt-step TPM.

A set of sufficient statistics for $\{N, \pi, \xi\}$ is:

$n =$ number of animals captured on one or more occasions

$C1 =$ number of animals captured on the first occasion

and for each $t = 2, \ldots, T$,

$K01t =$ number of $x_{t-1} = 0$ to $x_t = 1$ transitions by all animals,

$K10t =$ number of $x_{t-1} = 1$ to $x_t = 0$ transitions by all animals,

$K11t =$ number of $x_{t-1} = 1$ to $x_t = 1$ transitions by all animals, and

$K00t =$ number of $x_{t-1} = 0$ to $x_t = 0$ transitions by all captured animals.

The dot notation indicates summation over $t = 2, \ldots, T$. The logarithm of the likelihood function is:

$$\log(L) = \log \binom{N}{n} + (C1 + K01.)\log(\pi) + (N - C1 + K10.)\log(1 - \pi)$$

$$+ \sum_{t=2}^{T} \left\{ (N - n + K00t)\log(1 - \pi + \pi\xi^{dt}) + K11t \log[\pi + (1 - \pi)\xi^{dt}] \right.$$

$$\left. + (K01t + K10t)\log(1 - \xi^{dt}) \right\}$$

Observe that models M0, Mb, Mt, Mh, and their combinations (Otis et al. 1978) do not require elaboration to account for unequally-spaced capture occasion.

7 Snail study

In a 3-year mark-recapture study, terrestrial snails (*Vespericola* cf *depressa*) were captured at coverboards (traps consisting of equally-sized pieces of plywood, in place for over two years) throughout a 10 hectare wetland site in the southern Willamette Valley (Severns 2005). The timing of trapping occasions was opportunistic, depending on weather conditions that were suitable for mollusk activity. Some were on consecutive days, while others were separated by up to 20 days. Some initial information appears in Table 2. Capture histories typically contained strings of consecutive captures and strings of consecutive no-captures, again suggesting the persistence model.

Behavioral model fits (Mb) were attempted in all three years, but the profile likelihoods increased apparently without limit in the first and third years and did not decline rapidly from its maximum in the second year. Maximum likelihood estimation for the equal catchability (M0), the time (Mt) and the persistence (Mp) models were successful in all three years. Table 2 also contains the resulting estimates of population totals

Table 2 Snail study summary information

	Year												
	1999–2000				2000–2001				2001–2002				
Number of occasions	45				40				29				
Trapping spacings (days)	1–58				2–14				4–20				
Snails captured	171				80				86				
Number of captures	563				319				270				
	M0	Mb	Mt	Mp	M0	Mb	Mt	Mp	M0	Mb	Mt	Mp	
Population total:	177	–	177	200	81	187	84	86	89	–	89	111	
Lower 95% limit:	173	–	173	188	80	133	80	82	86	–	86	102	
Upper 95% limit:	183	–	184	218	83	–	85	94	94	–	95	133	
-log likelihood max:	2010	1845	1925	1775	1038	985	975	914	853	789	831	695	

and the achieved maximum of the likelihoods. The maxima for Mb in the first and third years are the values at $N = 1,000$.

To illustrate, the following paragraphs summarize the analysis of the 2000–2001 study period, where all four models successfully produced population total estimates.

Sufficient statistics for fitting the persistence model are presented in Appendix B. In addition to the summary of Table 2, the capture probability estimate is 0.0907 with a standard error of 0.0074. The estimate of the persistence parameter, $\theta = 1 - \xi$, is 0.1503 with a standard error of 0.0073. A comparison of the persistence parameter to its standard error provides overwhelming evidence of persistence.

In the fit of M0, it is estimated that the probability of capture was 0.0985, with a standard error of 0.0052. In the fit of model Mt, capture probability estimates vary from 0 to 0.1304 with standard errors around 0.0065. (The numbers of captures by occasion appear in Appendix B). For the fit of Mb, the number of no-captures prior to initial capture is 1,433, the number of recaptures is 239, and the number of no-recaptures is 1,447. The estimate of initial capture probability is 0.0138, with a standard error of 0.0015; the estimate of recapture probability is 0.1418 with a standard error of 0.0085.

A likelihood ratio test can compare the persistence model with the equal catchability model. The single degree of freedom chi-square statistic, 247.44, leaves little doubt that the persistence model provides a better fit. Model selection statistics - AIC, BIC, etc. - also indicate strong preference for the persistence model over the time and behavior models.

8 Discussion

All models entail questionable assumptions. Questioning is particularly appropriate where different models produce very different estimates. The butterfly study illustrates this strongly. The persistence model estimates the population total to be roughly 50% higher than the equal catchability model, the behavior model, and the time model. Comparisons based solely on the log-likelihood are important, but not necessarily convincing. The ability to include behavior and time variation along with persistence allows us to directly test whether those features are indeed important. In the butterfly study, they were not.

Persistence is an understandable phenomenon. The coverboards used to trap snails offer shelter for terrestrial mollusks in a habitat where coarse woody debris is absent, thereby providing a habitat wherein marked individuals are likely to persist. Persistence in the Great Copper study may be a function both of the butterflies and the observer. Butterflies tend to congregate near favorable habitat, which is readily recognized by the observer. In both of these cases that the persistence is related to the characteristics of the subject animals, the environment, and/or the observational pattern. Persistence is not necessarily a reaction of the animal to its being captured, as is the case in the behavioral model.

Other extensions to the basic persistence model are possible. One example is Mpbt, combining all three features in a model with parametric structure

$$p_1 = \pi_1$$

and for $t > 1$,

$$
\begin{aligned}
p_t &= \theta \pi_t, && \text{if } M_t = \max(x_s : s < t) = 0 \\
&= \theta \gamma_t, && \text{if } x_{t-1} = 0 \text{ and } M_{t-1} = 1 \\
&= 1 - \theta(1 - \gamma_t), && \text{if } x_{t-1} = 1
\end{aligned}
$$

where having $\text{logit}(\gamma_t) = \text{logit}(\pi_t) + \lambda$ for all t retains the single-parameter representation of behavior. We have not, however, been able to establish a method for fitting this model.

A referee and an associate editor have provided valuable assistance. The latter points out that a conditional likelihood approach to model fitting would be preferable to the general likelihood method if the probabilities in our models are related to other factors as covariates.

Appendix A: Great Copper data

This is a listing of capture histories for the $n = 45$ Great Copper butterflies on $T = 8$ occasions.

History	Butterflies
11111100	1
01111011	1
01010110	1
11100000	1
01101000	1
11000000	1
01100000	1
01000010	1
00011000	1
00010100	2
00001100	1
00001001	1
00000011	1
10000000	5
01000000	5
00100000	4
00010000	3
00001000	4
00000100	4
00000010	3
00000001	3

Appendix B: sufficient statistics from the 2000–2001 snail study

Occasion	Interval (days)	$K01$	$K10$	$K11$	$K00$	Captures
1	–	–	–	–	–	1
2	3	3	1	0	76	3
3	7	1	3	0	76	1
4	4	14	1	0	65	14
5	3	0	11	3	66	3
6	4	10	2	1	67	11
7	2	7	10	1	62	8
8	6	3	7	1	69	4
9	6	2	3	1	74	3
10	6	3	2	1	74	4
11	7	5	4	0	71	5
12	5	1	2	3	74	4
13	9	4	3	1	72	5
14	9	7	2	3	68	10
15	9	5	9	1	65	6
16	6	5	4	2	69	7
17	6	1	2	5	72	6
18	9	9	3	3	65	12
19	4	2	9	3	66	5
20	7	7	3	2	68	9
21	7	6	3	6	65	12
22	6	4	5	7	64	11
23	14	4	7	4	65	8
24	4	8	3	5	64	13
25	6	1	6	7	66	8
26	6	10	2	6	62	16
27	5	3	7	9	61	12
28	8	8	1	11	60	19
29	9	4	9	10	57	14
30	6	4	6	8	62	12
31	5	4	5	7	64	11
32	7	4	5	6	65	10
33	9	10	4	6	60	16
34	9	5	11	5	59	10
35	5	7	8	2	63	9
36	9	0	8	1	71	1
37	8	7	0	1	72	8
38	4	2	6	2	70	4

Appendix B continued

Occasion	Interval (days)	K01	K10	K11	K00	Captures
39	8	1	3	1	75	2
40	13	1	1	1	77	2
Totals:		182	181	136	2,621	320

$T = 40$ occasions, $n = 80$ snails, $C1 = 1$ capture

References

Burnham KP, Anderson DR (2002) Model selection and multimodel inference: a practical information-theoretic approach. 2nd ed. Springer, New York

Darwin C (1876) The effects of cross- and self-fertilization in the animal kingdom. Murray, London

Lewis AC (1986) Memory constraints and flower choice in *Pieris rapae*. Science (New Series) 232(4752):863–865

Otis DL, Burnham KP, White GC, Anderson DR (1978) Statistical inference from capture data on closed animal populations. Wild Monogr 62:135

Ramsey FL, Usner D (2003) Persistence and heterogeneity in habitat association studies using radio tracking. Biometrics 59:331–339

Severns PM (2005) Response of a terrestrial mollusk community to an autumn prescribed burn in a rare wetland prairie of western Oregon, USA. J Molluscan Stud 71:181–187

Severns PM, Villegas S (2005) Butterflies hanging on to existence in the Willamette Valley: a relict population of the great copper (*Lycaena xanthoides* Boisduval). Northwest Sci 79:77–80

Severns PM, Boldt L, Villegas S (2006) Conserving a wetland butterfly: quantifying early lifestage survival through seasonal flooding, adult nectar, and habitat preference. J Insect Conserv 10:361–370

White GC, Anderson DR, Burnham KP, Otis DL (1982) Capture-recapture and removal methods for sampling closed populations. Los Alamos National Laboratory, LA- 8787-NERP, 235 pp

Yang HC, Chao A (2005) Modeling animals' behavioral response by Markov chain models for capture-recapture experiments. Biometrics 61:1010–1017

Author Biographies

Fred L. Ramsey is emeritus professor of Statistics at Oregon State University. He has also taught Statistics at the Oregon Health Sciences University, the University of Copenhagen Denmark, Wollongong University in New South Wales, Australia, and Murdoch University in Perth, Western Australia. Fred's interests are in applications of Statistics to wildlife studies, with special emphasis on birds.

Paul M. Severns is a PhD candidate interested in the relationship between rapid evolution and ecological systems, and the impact of evolutionary processes on effective biological conservation. Paul likes butterflies and a few plants.