An introduction to event history analyses for ecologists

JULIE LANDES [iD],[1],† SACHA C. ENGELHARDT [iD],[1,2] AND FANIE PELLETIER[1]

[1]Département de Biologie, Université de Sherbrooke, 2500, boulevard de l'Université, Sherbrooke, Quebec J1K 2R1 Canada
[2]Department of Biology, Institute of Ecology and Evolution, University of Bern, Etologische Station Hasli, Wohlenstrasse 50a, Hinterkappelen CH-3032 Switzerland

Abstract.   Efforts to understand the emergence of an event require our ability to measure and understand the dynamics between time in a state (e.g., being alive or a behavior) and the outcome of the state. Studying the main drivers that affect changes in state over time allows researchers to better understand population dynamics and evolutionary processes. Event history analyses provide a range of theoretical and empirical tools to explore the emergence of an event. Their use is still restricted in ecology; however, they are commonly used in human demography. Event history analysis is a powerful tool for measuring the probability that an event occurs at time t. Here, we provide an introductory guide for ecologists who are interested in exploring event history analyses in their research. In the first part of this article, we outline key concepts in event history analyses and present a decision tree, statistical techniques, and their applications to ecological questions. To introduce practical applications of event history analyses, we provide four detailed tutorials, stemming from observational and longitudinal records of events in mammalian and avian species, along with relevant R scripts. We then explain how to interpret and present results of such analyses. Our results show that event history analyses are useful to quantify the effect of factors on the emergence of events. We conclude by highlighting additional strengths, pitfalls, and limitations researchers should be aware of when using such methods. We foresee the use of event history analyses for ecological studies.

† **E-mail:** julie.landes@usherbrooke.ca

INTRODUCTION

Event history analyses, also known as survival analyses and failure time analyses, investigate the likelihood, also known as the risk or failure, that an event will occur and model the probability of a change in state over a time step (Klein and Moeschberger 2003, Allison 2014). The types of events that can be modeled are numerous and include death, the failure of a nest, the emergence of a disease, marriage, or machine breakage (Klein and Moeschberger 2003). From the 17th century to the 1950s, life tables were used to study mortality in humans, the reliability of equipment in industry, and the emergence of diseases (Lee and Go 1997). In the 1950s, event history analyses were further developed (Kaplan and Meier 1958, Armitage 1959), mainly for biomedical purposes, and have since been widely used in many disciplines including medical studies, demography, sociology, and engineering. Behavioral ecologists have recently used event history analyses to study mortality in gray mouse lemurs, *Microcebus murinus* (Landes et al. 2017), and rhesus macaques, *Macaca mulatta* (Brent et al. 2017), time to dispersal in the

cooperatively breeding southern pied babblers, *Turdoides bicolor* (Nelson-Flower et al. 2018), latency to donate food to a conspecific in Norway rats, *Rattus norvegicus* (Schweinfurth and Taborsky 2017), and latency to eat a reward in a novel object test in the cooperatively breeding cichlid *Neolamprologus pulcher* (Bannier et al. 2017).

Event history analyses allow researchers to consider censored and truncated data (see definitions in Fig. 1 and Table 1), thereby reducing sampling bias. These analyses also take into consideration the emergence of the event over time, avoiding averaging time across levels and excluding records with missing information. Excluding censored data not only decreases the sample size but also inflates the estimated hazard rate, which provides information about the relative risk of changing states over a time period (Griffin 1993, Klein and Moeschberger 2003). If
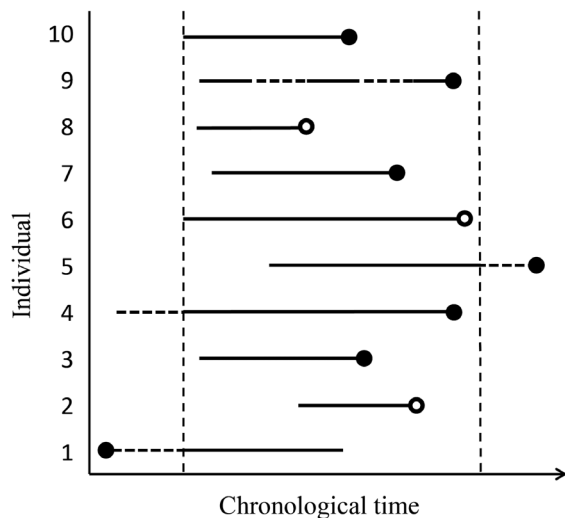

Fig. 1. Fictitious example of the monitoring of ten individuals. Each individual is followed over the studied period (bounded by the two dotted lines). Filled circles correspond to the event. Empty circles stand for the times an individual exits the study before its end without having experienced the event. Individual 1 experienced the event before the study began; the corresponding survival data is left-censored. Individuals 2, 5, 6, and 8 did not experience the event during the study; these data are right-censored. Individual 9 was only studied on repeated time intervals; the data is interval-censored.

several individuals are still alive at the end of the study period, ignoring them in the survival analyses could negatively bias the survival estimate, cause a negative temporal bias toward the end of the study, or overrepresent individuals with a greater likelihood of changing state over the observed time (Griffin 1993, Klein and Moeschberger 2003, Allison 2014). Ecological studies will often exclude cohorts with individuals who are still alive at the end of study period (e.g., Bérubé et al. 1999). For example, for long-lived species, one may have to exclude the most recent cohorts. However, applying event history analyses allows for the inclusion of censored and truncated data, which lowers potential problems with sampling bias and increases statistical power compared to likelihood analyses commonly used by ecologists (Griffin 1993, Klein and Moeschberger 2003, Allison 2014). Generalized linear models and generalized linear mixed models with a Poisson distribution can account for censoring and truncation, such as hurdle models (left truncation and right censoring; Jackman 2017) and Poisson regression models for interval-censored count data (Watson 2011); however, these models are not the focus of this article.

The objective of this article is to introduce event history analyses to ecologists who are interested in exploring the potential application of these methods to their data. The ecological questions that can be investigated using these methods are numerous. For example, researchers may want to quantify the effect of given factors on mortality, the difference in the age at first reproduction between groups of individuals, the duration between birth intervals, the effect of parental care on the age at fledging or weaning of the offspring, or the time individuals take to show a given behavior.

This article is organized in five sections. First, we begin by outlining key concepts and theory involved in models of event history analyses. Second, we describe the parameters and main functions of the event history analyses, explain why these parameters and functions are of interest to ecologists, and provide a guide for applying statistical models of event history analyses. Third, we outline the steps to follow concerning data collection, data exploration, discrete-time analyses, and continuous-time analyses. Fourth, we provide four tutorials stemming from

Table 1. Glossary of terms used in event history analyses.

Term	Definition
Event	Change of state, also known as failure, for which time-to-occurrence is analyzed. It is not necessarily death.
Truncation	Data are truncated when individuals who experience the event outside of the study period are excluded from the analysis. Left truncation occurs when all individuals enter the study before experiencing the event. Right truncation occurs when all individuals experience the event during the monitored period (and thus before it ends).
Censoring	Data are censored when the event is not observed during the study period, but happens before an individual enters (left censoring) or after he/she exits the study (right censoring). In the case an event is known to happen during a time interval, the data are interval-censored.
Absolute risk	The difference between the proportion of individuals in one group that experienced the event and the proportion of individuals in another group that experienced the event. See Appendix S1.
Hazard ratio	The ratio of the hazard rates of individuals with different values for a given variable relative to the baseline hazard function at any particular point in time. For individual i, the hazard ratio is a function of Euler's number to the power of a product of a vector of covariate values for the individual i and a vector composed of one coefficient per covariate.
Relative risk	The ratio of one minus the baseline survival function of the treatment group compared to one minus the baseline survival function of the control group. See Appendix S1.
Odds ratio	Within each group, the odds are calculated as the number of occurrences of the event divided by the number of non-occurrences of the event. The odds ratio is the quotient of the odds of the event occurring in both groups. See Appendix S1.
Remaining life expectancy	Remaining life expectancy at time t is the average time remaining before the event for individuals that have not met the event at t. See Appendix S1.
Continuous variable	Numeric variable that has an infinite number of possible values.
Discrete variable	Variable that can only have a limited number of values.
Time-varying covariate	Covariate values vary over the individual's life. The survival data of each individual are fractioned to allow the covariate to take different values on different age intervals.
Stratification	Splitting the variable into multiple levels (see decision tree for more information).
Clustering	Individuals with recurrent event and/or observations are divided into groups based on correlated common features.

Note: Some definitions were taken or adapted from Klein and Moeschberger (2003) and Moore (2016).

observational and longitudinal records of events in mammalian and avian species, R scripts, and interpretations of different event history analyses (e.g., non-parametric, semi-parametric, and parametric results). Finally, we discuss the most likely pitfalls encountered in event history analyses. As this article is intended as a primer for ecologists who are interested in applying event history analyses, we refer readers to the classical literature and propose that this guide acts as a useful stepping stone to further explore event history analyses.

THE BASICS OF EVENT HISTORY ANALYSES

Event history analyses aim to analyze the time-to-event for individuals (e.g., death or any other event of interest) over a given amount of time, which does not necessarily correspond to their lifespans. The original time is the date on which the individuals enter the study, which could be their birth date, the date of emergence of a disease, or an arbitrary date of first observation. The original time can differ for each individual. The time at last observation can correspond, for example, to the time at the event or the time the individual exits the study without having experienced the event. In the latter case, the survival data are right-censored (Fig. 1, Table 1). It is, for example, the case when individuals are still alive at the end of a study on survival, or when individuals disperse outside the studied area and are not re-captured. Under such circumstances, we are aware that the individuals did not experience the event until a known time. In the case that individuals experienced the event prior to entering in the study, the data are left-censored (Fig. 1, Table 1). It is, for example, the case when individuals already reproduced before entering a study on the age at first reproduction. This gives the information that the event occurred before a known time (the original time). Finally, in the case the event is known to occur within a time interval, the data are interval-censored (Fig. 1, Table 1). The change of status is known to occur between two times (e.g., during

the monitored or un-monitored period), but the time-to-event is not precisely known within the time period. It is also possible that individuals experiencing the event outside the observed time period are not considered. In this case, the data are truncated (Table 1). Left truncation refers to individuals that did not experience the event prior to their entrance in the study, and only those individuals will be taken into account. For example, data are left-truncated in a study of survival when all considered individuals are alive at the beginning of the study. On the contrary, right truncation refers to the case when only individuals who experienced the event during the studied period (after the start and prior to the end of the study) are taken into account. Left and right censoring and truncation are rarely accounted for in ecological studies. The structure of the survival data (truncation, censoring) must be considered when choosing which method to use.

Three functions are at the core of event history analyses. The first is the survival function (Box 1). The survival function is the probability of an individual not experiencing the event (death or any other event of interest) beyond time t (Klein and Moeschberger 2003). This monotonic function is equal to one at the beginning of the study and declines over time toward zero. This measure is often used to compare patterns of failure (i.e., emergence of the event). The remaining life expectancy at time x (i.e., the mean time before experiencing the event for an individual at time t; see Appendix S1) is estimated from the survival function. The second function, the probability density function, reflects the probability that the event will occur at time t (Box 1; Klein and Moeschberger 2003). The probability density is the negative of the derivative of the survival function. The third and last function, the hazard function, is the most frequently reported and the most useful (Table 1 and Box 1). There are several terms used to refer to the hazard function, such as the hazard rate or the force of mortality, and it is most often represented as $h(t)$ (Table 1 and Box 1; Klein and Moeschberger 2003). The hazard function is the instantaneous risk, that is, the ratio of the probability density function and the survival function, that an individual at time t will experience the event (Klein and Moeschberger 2003). The hazard function

---

Box 1.

Event history analyses main functions

The *survival function* is the probability that an individual will remain in a state until time t. The survival function has an interval spanning from 0 to 1. At time 0, the function takes the value of 1, and the value of the survival function decreases monotonically over time toward zero.

$$S(t) = \Pr(T > t), \, 0 < t < \infty$$

where T is the time at which the event occurs.

The *cumulative distribution function*, also known as the cumulative risk function in event history analysis, is the complement of the survival function.

$$F(t) = \Pr(T \le t), \, 0 < t < \infty$$

The *probability density function* is the rate of change of the cumulative distribution function. Alternatively, it can be calculated as the negative rate of change of the survival function.

$$f(t) = \frac{d}{dt}F(t) = -\frac{d}{dt}S(t)$$

The *mortality hazard function* (also known as age-specific failure rate, force of mortality, hazard rate, or risk function) is a measure of the propensity that an individual experiences the event at a given time. More precisely, it is the rate of experiencing the event between t and $t + \Delta t$ (i.e., in the time interval Δt) knowing that the individual did not experience it until age t:

$$h(t) = \lim_{\Delta t \to 0} \frac{1}{\Delta t} P(t \le T < t + \Delta t | T \ge t), \, 0 < t < \infty$$

The *cumulative mortality hazard* is the integral of the mortality hazard:

$$H(t) = \int_0^t h(u) du$$

The cumulative mortality hazard is directly linked to the survival function:

$$S(t) = e^{-H(t)}$$

---

cannot be negative and can take different shapes (Klein and Moeschberger 2003). For a categorical variable with more than two states, separate hazard functions are calculated. The hazard ratio compares the hazard rates of an individual with a covariate value to the baseline hazard function (Table 1; Appendix S1). For example, if there are two or more levels in a categorical variable, the

> ## Box 2.
>
> ### Frailty
>
> In a population, individuals of the same age can have different chances of experiencing an event depending on intrinsic properties (Vaupel et al. 1979; Fig. 2). Frailty can be used in two ways: (1) as a random effect, or (2) as an adjustment for overdispersion in univariate analysis and represents the influence of unmeasured or unobserved covariates. Frailty is most commonly an unobservable random effect shared by individuals within the same sub-group, and it has a multiplicative effect on the hazard rates of all subgroup members (Hougaard 2000, Klein and Moeschberger 2003). For example, individuals that share a high frailty value will tend to experience the event earlier (e.g., they have more risks of experiencing the event) than individuals that share a distinctly lower frailty value. Shared frailty assumes that individuals from the same subgroup share the same frailty. If the frailty variance is equal to 0, model evaluation is reduced to a fixed effect model, which implies the absence of group-level heterogeneity (Klein and Moeschberger 2003). Frailty is commonly assumed to have a gamma distribution. However, there are other distributions (e.g., inverse Gaussian), and estimates can vary based on which distribution is selected (Hougaard 1995, 2000). Failing to account for heterogeneity could cause the estimation of coefficients, standard errors (SE), and the change in the hazard rate over time to be misleading, and negative duration dependency could occur (i.e., a negative association between time and the hazard rate caused by two subpopulations with distinct failure rates; Blossfeld and Rohwer 2002).
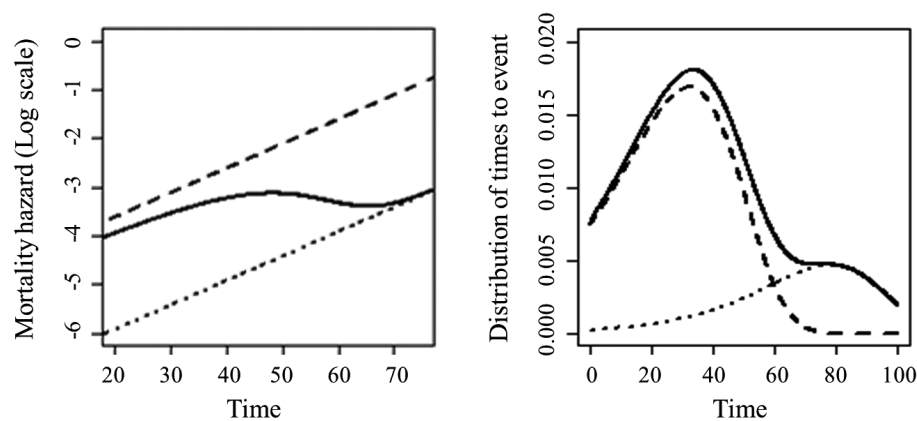>
>
>
> Fig. 2. Mortality hazard and distribution of the times to event for a population that can be divided in two sub-groups according to their frailty. The frail individuals (dashed line) die sooner than the robust individuals (dotted line). The plain line represents the entire population. Figure from Samuel Pavard.

baseline hazard function represents the hazard rate for the reference level. The hazard ratio is often reported as a relative reduction in risk (i.e., probability that an event occurs in the next instant; Table 1), and the reduction in risk is calculated as one minus the hazard ratio. For example, a hazard ratio of 0.7 is equal to a 30% reduction in risk. A hazard ratio of one means that the hazard ratio is equal between levels of a categorical covariable or across the values of a continuous covariate. For a categorical covariate

(e.g., two groups of individuals), a hazard ratio of two would mean that twice as many individuals in one group will experience the event compared to individuals in the other group during the next unit of time. For a continuous covariate, the hazard ratio is exponentiated by the units of increase (e.g., a hazard ratio of two over a 10-unit increase would be 2^{10}), which represents the multiplicative model structure. The hazard ratio only applies to individuals during a specific time frame.

If a model does not account for all sources of variation that influence the hazard rate and ratio, there could be unobserved heterogeneity (i.e., frailty, such as random effects; see Box 2). This occurs when model specificity is poor, such as when not all relevant covariates are included in the model, or when there is structure in the data such as distinct frailty groups experiencing different hazard rates. The hazard rate may not be constant over time, which could suggest unobserved heterogeneity. Alternatively, a changing hazard rate may represent the difference between the frailest individuals (i.e., they have a higher probability of experiencing the event) and the more robust individuals (i.e., they have a lower probability of experiencing the event, Box 2). The hazard may be increasing as time progresses because the frailest individuals experience the event early, while the robust individuals experience it much later (Vaupel et al. 1979). It is possible to include frailty into mixed models (i.e., models including fixed and random effects) to account for this possible inherent structure in the data, while controlling for pseudoreplication and temporal auto-correlation (Aubry et al. 2011).

Event History Analyses

In event history analyses, non-parametric, semi-parametric, and parametric models are available. Non-parametric models can neither estimate the effects of covariates nor account for time-varying covariates (Table 1), but they can be useful to visualize raw data (Klein and Moeschberger 2003). Thus, in this article, we only present non-parametric models as a tool for preliminary explorations of the survival function and to compare the distribution of the raw data to the estimated distributions of the survival function by semi-parametric and parametric models. Non-parametric models are less applicable to ecological studies (Klein and Moeschberger 2003), and, consequently, this article mainly focuses on semi-parametric and parametric models. R packages for each method are presented in Appendix S2.

Non-parametric models

The Kaplan–Meier (KM) estimation of the survival function is non-parametric and is fitted to the raw data (Kaplan and Meier 1958, Sedgwick 2014). For each time interval, the KM survival function returns the probability of remaining in a state (e.g., alive), which is calculated as the number of individuals remaining in the state divided by the number of individuals at risk, under the assumption that the instantaneous risk of a change in state remains constant across intervals. An additional assumption is that remaining in a state until time t implies not changing state at time t. Drops in the KM survival function only occur at event times, and the survival function does not drop at times when only censoring occurs. If an individual survived until the end of the study, the KM survival function will not go to zero. The interpretation and extrapolation of the survival of an individual or the cumulative survival probability (i.e., life expectancy) should not be performed beyond the end of the study. The median survival of survival functions, in the case of a variable with multiple levels, can be calculated if both survival functions cross the 50th percentile mark. Advantages of the non-parametric models are that they represent the real, observed data without assuming a distribution for the baseline, and the KM survival function helps researchers visualize preliminary comparisons of survival functions for different levels of categorical variables. Comparisons of survival functions for different levels of a categorical variable can be performed with a log rank test, which returns a χ^2 test statistic and P value.

Semi-parametric models

The effects of covariates on the response variable and the patterns of the main functions (Box 1) can be assessed with semi-parametric Cox proportional hazard models, also known as Cox's model, which estimate the hazard function without making assumptions about the distribution of the baseline hazard rate. Cox's models have no intercept and estimate covariate regression coefficients and the main functions, whereas parametric models have an intercept and estimate covariate regression coefficients, the main functions, and the change in the distribution of the baseline hazard over time, hence the term semi-parametric (Griffin 1993, Klein and Moeschberger 2003, Allison 2014). Cox's models assume proportionality of covariate effects over time (i.e., proportional hazard

assumption): for individuals with unique values for a covariate, the ratio of their hazard rates is constant (i.e., proportional) over the observed time (Klein and Moeschberger 2003, Allison 2014). Thus, the effect of a covariate would remain the same over time. For example, if sex has an effect on mortality, the amount of this effect will be the same all over the individuals' life. Violation of proportional hazard assumption implies an interaction between the covariate and time. If some variables do not satisfy the proportional hazard assumption, the variable could be stratified (Klein and Moeschberger 2003, Allison 2014) or parametric models should be performed. The mortality hazard function for a Cox's model is defined as:

$$h(t|x) = h_0(t)e^{\beta x}$$

where $h_0(t)$ is the baseline mortality hazard, for which the distribution is unknown, t is the time at the event or censoring, and the parameter β quantifies the association between covariate x and the time to the event.

Parametric models

Parametric models assume that the baseline hazard follows a distribution. This distribution can be monotonic (Gompertz and Weibull distributions), or non-monotonic (lognormal and loglogistic distributions; Klein and Moeschberger 2003). These different possible distributions describe the mortality distributions of several mammals (Ricklefs and Scheuerlein 2001, Gaillard et al. 2004, Bronikowski et al. 2011). Parametric models estimate the distribution's parameters (for further information on the distribution formulas and parameters, see Klein and Moeschberger 2003). The effects of the covariates on the hazard function can be proportional (proportional hazard models; PH; the parametric version of the semi-parametric model) and can vary over time (accelerated failure time models; AFT; the probability of changing state increases with time), or in the same model, some covariates can have proportional effects and others AFT effects (PH-AFT models; described in Landes et al. 2017). Parametric models are less flexible and powerful compared to semi-parametric models. Further, there can be convergence problems and difficulties characterizing the distribution of the hazard function with

parametric models. For example, if the chosen distribution of the mortality hazard baseline is the Gompertz distribution ($ae^{b.t}$), the model will be:

$$h(t|z) = ae^{bt}e^{\alpha'z}$$

in the case of a PH model and

$$h(t|x) = e^{\beta'x}ae^{bte^{\beta'x}}$$

in the case of an AFT model
where a and b are the shape and scale parameters of the Gompertz distribution, t is the time at the event or censoring, and $\alpha'z$ and $\beta'x$ are the parameters of the PH and AFT models, respectively, with α and β the coefficients corresponding to the covariates z and x.

EVENT HISTORY ANALYSES DECISION TREE

To introduce readers to event history analyses, we provide an overview of the different steps, procedures and recommendations to consider in the form of a decision tree (adapted from Griffin 1993; Fig. 3). We refer readers to the classical literature to further explore event history analyses, as this decision tree is not exhaustive and acts as an introduction.

Step 1: Data collection

The data collection method used can influence the type of analyses performed (Fig. 3). The collection of continuous-time data is more flexible in the types of analyses that can be performed than with discrete-time data. With continuous-time data, researchers can conduct continuous-time or discrete-time analyses, and the interval of interest can be adjusted, depending on the research question. Discrete-time data limit the type of analyses that can be performed to discrete-time analyses (Griffin 1993, Allison 2014).

Step 2: Data exploration

Data exploration should be performed prior to event history analyses (Fig. 3). Researchers should examine the distribution of durations to events over the observed time, the number of failures (i.e., change of state), the number of censored data, the sample size, and the number of ties (when two or more times to events have the same duration) per interval (Griffin 1993, Klein and Moeschberger 2003, Allison 2014). As a rule
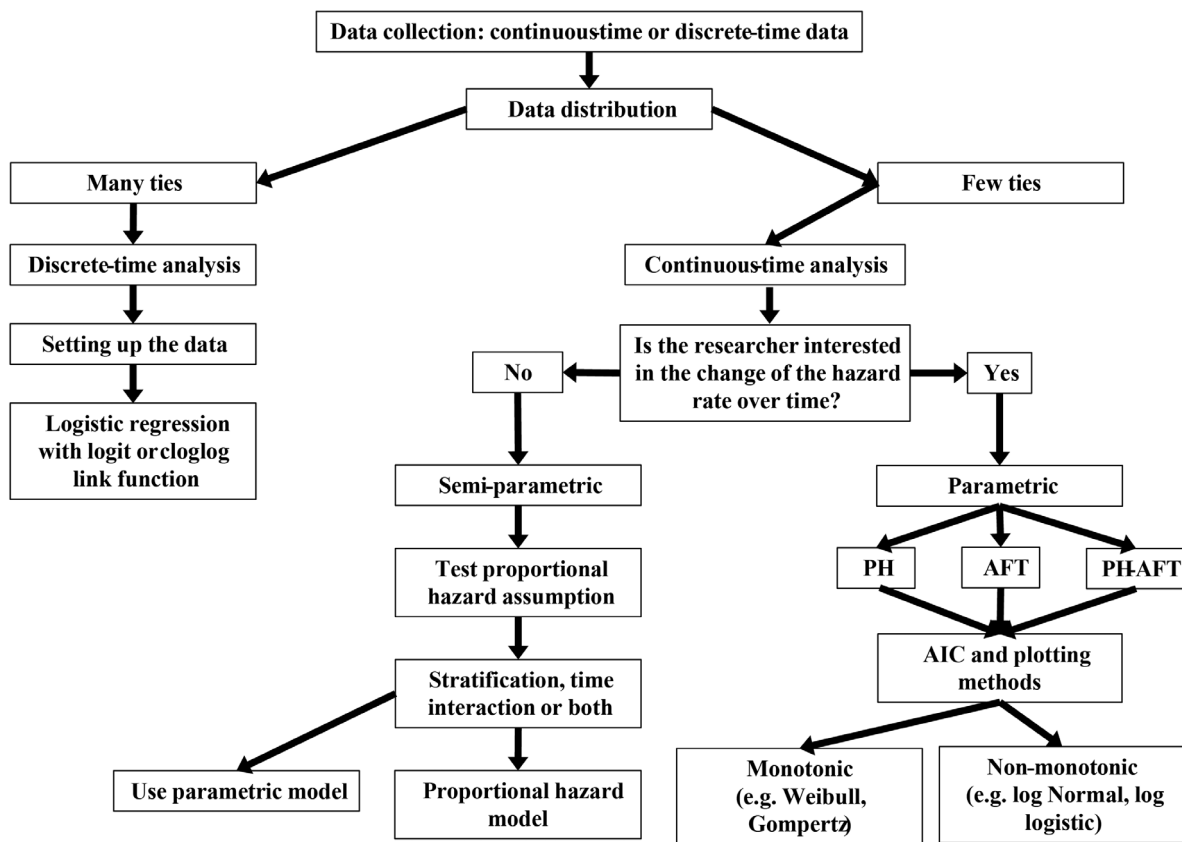
Fig. 3. Event history analyses decision tree (adapted from Griffin 1993; see methods for explanations). If the assumption of covariate proportionality is met (yes), a Cox's model can be performed. If the assumption of covariate proportionality is not met (no), either the variable is stratified and a Cox's model is performed or a parametric model should be performed.

of thumb, there should be a minimum of 10 failures per covariate. The number of ties is not the total number of ties throughout the study, rather it is the number of failures relative to the number at risk per time interval (Griffin 1993). The asymptotic properties of the estimators of the hazard rate assume that the probability for ties is zero (Griffin 1993, Klein and Moeschberger 2003). There are methods to correct for ties (e.g., Breslow and Efron approximations, Hougaard 2000, Klein and Moeschberger 2003).

Step 3: Discrete-time analyses

Logistic regression with a logit or cloglog link function is commonly used for discrete-time analyses (Allison 2014). Relative risk can be interpreted directly from the cloglog model output, whereas the log odds must be converted. For

discrete-time analyses, the data need to be reorganized. The number of records equals the number of time units spent in a state (Griffin 1993, Allison 2014). For example, if the time from hatchling to fledgling was recorded with time intervals equal to each day and offspring A became a fledgling after 30 d, the data set for offspring A would include 29 rows with the outcome equal to 0 (non-occurrence of event) and 1 row with the outcome equal to 1 (occurrence of event).

Step 4: Continuous-time analyses

If the effect of covariates on the hazard rate is of interest, independent of both the distribution of the data and the change in the hazard rate over the observation period, a semi-parametric model should be conducted (Griffin 1993, Klein

and Moeschberger 2003, Allison 2014). The semi-parametric model makes no assumption about the distribution of the baseline hazard rate, and it assumes a proportional hazard rate over time (Griffin 1993, Klein and Moeschberger 2003, Allison 2014). Each covariate should be tested for proportionality, and violation of the proportional hazard assumption implies an interaction between the covariate and time. If a covariate is non-proportional, an interaction between time and the covariate can be included in the model before re-testing for proportionality: if the interaction is significant, the variable can be stratified (Klein and Moeschberger 2003). The stratification —splitting the variable into multiple levels—of a non-proportional covariate fits a different baseline hazard for each stratum (i.e., for each level of the stratified covariate), and stratification assumes that the proportional hazard assumption is met for the other covariates within each stratum (Klein and Moeschberger 2003). Alternatively, a parametric model could be performed.

If the change in the hazard rate over the observation period (duration dependency) is of interest, a parametric model, with a distribution fitted to the baseline hazard rate, should be conducted (Griffin 1993, Klein and Moeschberger 2003, Allison 2014). To assess which distribution to use, AIC model selection and plotting methods are used.

Tutorials

We provide four tutorials to help readers learn how to perform event history analyses (see Appendix S3 for tutorials 2, 3, and 4; Appendix S4 for scripts; and Data S1 for data sets).

Tutorial 1: Monitoring data of ducks

Event history analyses can be performed on short-term (e.g., over a reproductive season) or long-term individual-level observations. The data set is published in a book (Everitt 2002; see data in Data S1). The survival of 31 hatch-year (young) and 19 after hatch-year (old) ducks (species unspecified) were monitored from the beginning of the study to the end of the follow-up. Event history analyses were conducted to assess the effect of the age group (young vs. old) on the time to death of the ducks. Some of the data are

right-censored because of individuals who exited the study before the end of the study.

First, we estimate the KM survival function. Survival times for older and younger birds do not differ ($\chi^2 = 0.10$, $P = 0.77$). In Fig. 4A, the KM survival probability of experiencing death at some time after 63 d was approximately 0.58 for the younger birds and approximately 0.62 for the older birds. The KM survival probability does not reach 0 at the end of the follow-up time (Fig. 4A), because the longest observed survival time is right-censored at the end of the follow-up time. If none of the observations were censored, the KM survival probabilities would have been absolute probabilities, and they would have represented the proportion of birds in each age group that did not experience death by the end of the study (Sedgwick 2014). However, KM survival probabilities with censored data represent conditional probabilities of experiencing death after the duration of the study (Sedgwick 2014). The median survival is not reached in both groups (Fig. 4A), such that we cannot interpret the median survival or the difference in median survival.

To estimate the effects of the variables, we conduct a semi-parametric Cox proportional hazard model. The proportional hazard assumption is not violated ($\rho = -0.43$, $\chi^2 = 2.92$, $P = 0.09$). No problem in the functional distribution of the censored and uncensored data is detected (Fig. 4B). We estimate the model's predicted survival probabilities (Fig. 4C), predicted cumulative hazard function (Fig. 4E) and predicted hazard function (Fig. 4D) for young and old birds. The smoothing spline with 6 degrees of freedom fits the cumulative hazard function better than the smoothing splines with 3, 4, or 5 degrees of freedom (Fig. 4E). Young and old birds experience the highest instantaneous risk of death (hazard function) 32 d within the follow-up time. The predicted hazard function values are below 0 from 54 d to the end of the study for both young and old birds; however, hazard function values cannot be negative, so values from 54 d on are equal to 0 (Fig. 4D). The life expectancies of young and old birds are 46.69 and 44.98 d after the beginning of the study. At any particular time, mortality rates are the same in both age groups such that the age of birds does not influence the probability of mortality ($\beta \pm SE = 0.14 \pm 0.50$, $P = 0.78$;
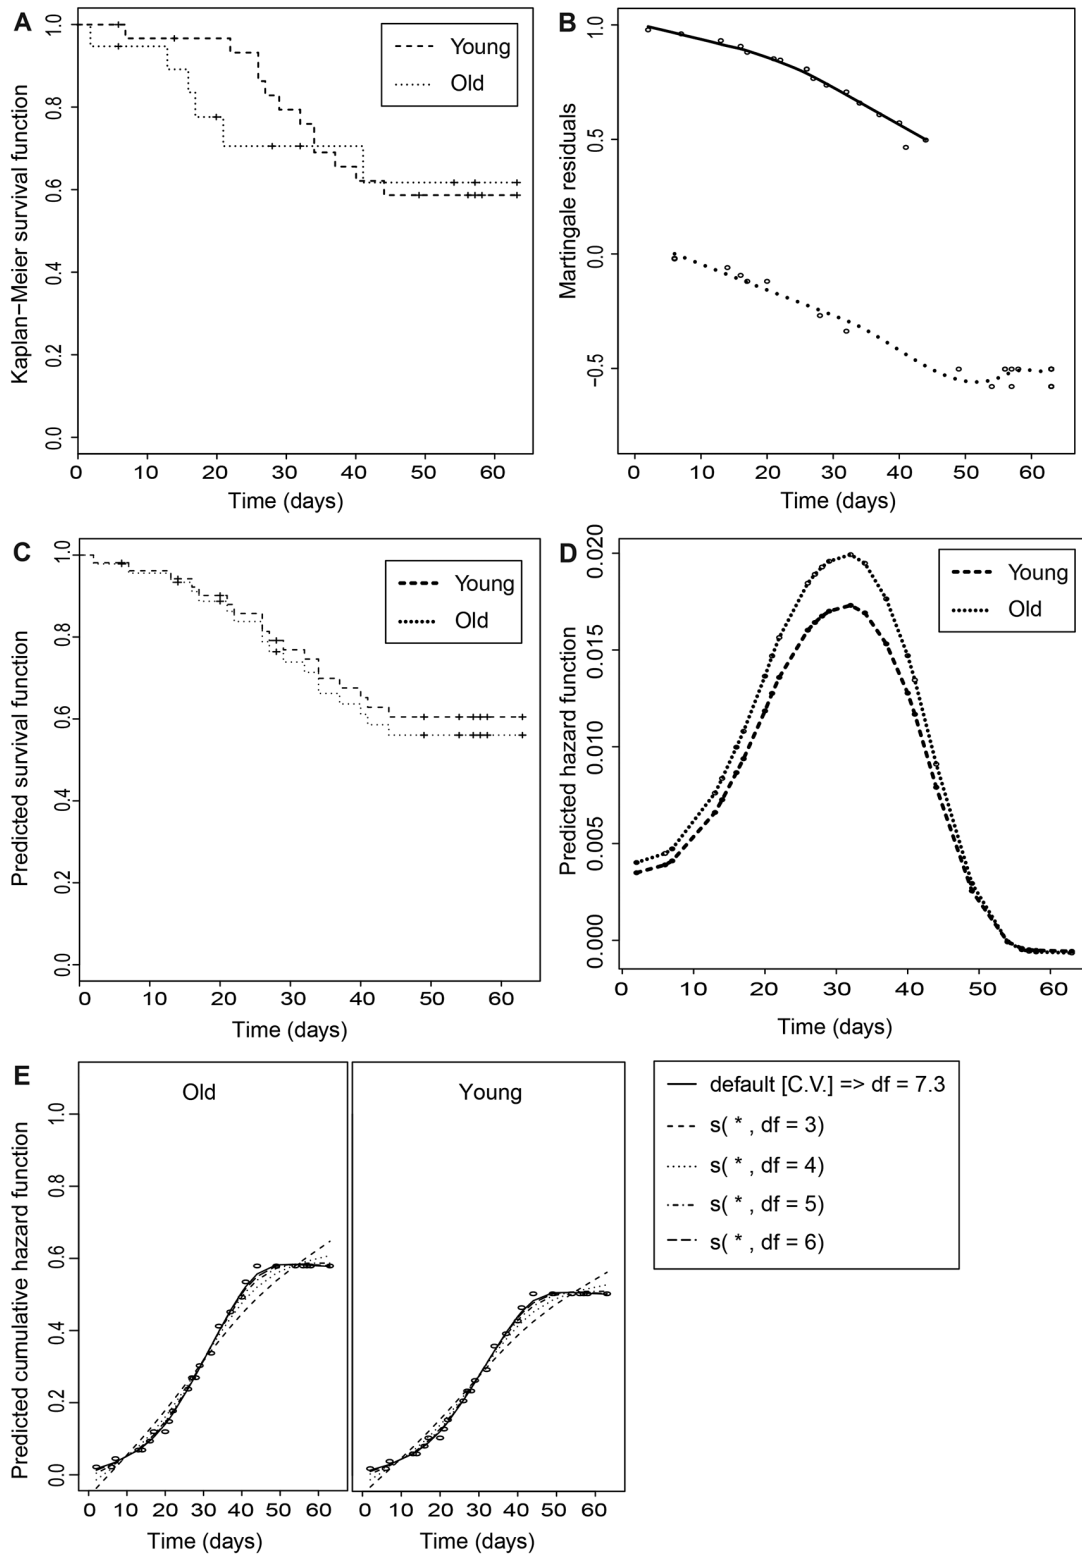
Fig. 4. (A) Kaplan–Meier survival plot representing the estimated probability of survival of ducks over time

(Fig. 4. *Continued*)

for both age groups (older, dotted line; younger, dashed line). The survival functions are represented as downward step functions, and each downward step in the survival functions represents the death of a bird in the corresponding age group. Each plus sign represents a right-censored observation in the corresponding age group. Censored observations were included when calculating survival probabilities; however, the censored observations influenced the survival probability at the next known (i.e., exact) death. (B) Diagnostic plot of model fit with Martingale residuals with smoothed lines for uncensored (solid line) and censored (dotted line) observations. (C) Predicted survival probabilities for young (dashed line) and old ducks (dotted line) from the beginning of the study to the end of the follow-up period. (D) Predicted hazard function for young (dashed line and solid circles) and old ducks (dotted line and empty circles). (E) Cumulative hazard function and smoothing splines with the degrees of freedom for each smoothing splines for old (left) and young (right) ducks.

---

hazard ratio [HR] = 1.15; 95% confidence interval [CI] of HR = 0.43–3.08).

DISCUSSION

Event history analyses have great potential to enlighten ecological studies. They provide a powerful tool to study the time to an event of interest and can be applied to most ecological domains. Their main strengths are that they allow researchers to account for time-to-event (e.g., baseline hazard) and for censored and truncated data, which highly decreases the risk of bias. Generally, semi-parametric models are to be favored to investigate research questions; however, there are research questions that can only be answered using parametric models. For example, if the aim of the study is to describe the distribution of the mortality, parametric models must be used (Vaupel et al. 1998). Additionally, event history analyses are not limited to binary states (e.g., alive/dead), and multistate modeling is available and can include intermediate states (Klein and Moeschberger 2003, Jackson 2011, Moore 2016). Competing risks occur when individuals may experience a change in state due to more than one cause and the causes preclude each other, which can be incorporated in event history analyses to estimate cause-specific hazards (Klein and Moeschberger 2003, Congdon 2007, Moore 2016). There are many different sorts of event history analyses beyond the scope of this article (Congdon 2007, Boonekamp et al. 2014).

Pitfalls

We would like to acknowledge methodological limitations and pitfalls of the previously discussed methods. One common source of bias is data selection. In the case that left truncation is due to individuals who have already experienced the event but are not included in the data, the interpretation of the results should account for this (Klein and Moeschberger 2003, Cain et al. 2011). When the event of interest is known to happen during a time interval, the precise date of the event is unknown, and the data must be considered and analyzed as interval-censored data. To simplify the analyses, it could be tempting to artificially transform these data into right-censored data by assigning arbitrary times at the event (such as the middle or the end of the time interval). This could bias the data, especially in the case of long-lasting time intervals. If there are some missing data (e.g., 30% of observations with unknown date of death of offspring who are confirmed dead), an alternative to conducting interval-censored modeling is to estimate the date of death by randomly assigning the missing data along the estimated function based on the non-missing data. When a parametric model with frailty and time-varying covariates should be performed with data that are interval-censored, the parameters of the distribution and the coefficient values for the covariates should be optimized for the highest likelihood value (or lowest negative log-likelihood value), whereas the frailty and CIs should be estimated by bootstrapping.

Biological interpretations of time-to-event analyses can be challenging, especially when first using these methods. Some examples of output interpretations are given in this article to help understand the results (see Tutorial 1, Fig. 4, and Appendix S3). The output of semi-parametric

and parametric models returns the effects of the covariates on the mortality hazard, and not survival. If the event is death, a positive effect of a covariate increases the individuals' mortality and thus decreases their survival. If the event is the fledging of offspring from the nest, a positive effect of a covariate corresponds to an earlier fledging. The hazard ratio is often reported as a reduction of risk (1 − hazard ratio × 100), and it is sometimes misinterpreted. The hazard ratio should not be confused with the relative risk, and hazard ratios should not be reported as relative risks (Sutradhar and Austin 2018). The hazard ratio can only be interpreted during a specified period of time, and the outcome should not be inferred beyond the follow-up time. For example, a hazard ratio of 0.60 comparing two levels of a categorical variable can be reported as a 40% reduction in risk. This hazard ratio of 0.60 should not be misinterpreted as individuals in a group progressing to the event at 0.60 the rate as those in the other group. It could be interpreted as individuals in one group changing state at 0.60 the rate of those in the other group, and that individuals who have not experienced the event by a specific time have 0.60 the chance of experiencing the event by the next time point compared to individuals in the other group.

Other potential interpretation problems are common to all modeling methods: coefficient estimations and significance interpretations, correct checks of the model fits, and respect of the assumptions. A particular pitfall to avoid is the violation of the proportional hazard assumption of semi-parametric models. It is incorrect to disregard this assumption and run a semi-parametric Cox non-proportional hazard model, because the estimates would be meaningless. If this assumption is not met, a categorical variable that does not meet the assumption can be stratified (Fig. 3). An interaction between the variable, which does not meet the assumption, and time can be added to test whether there is a duration-dependent effect (Fig. 3). If some, but not all, of the covariates have a duration-dependent effect and an AFT parametric model is conducted, researchers should be aware that all covariates are considered as duration dependent (Fig. 3). To avoid this problem, an AFT-PH parametric model should be performed (Landes et al. 2017; Fig. 3).

Researchers should report the predicted function values based on model predictions rather than solely reporting the KM estimates, which are estimates fitted to raw data and do not account for time-dependent covariates or frailty. To obtain the hazard function from the cumulative hazard function, we can refer researchers to two methods. First, if the time interval between two estimates of the cumulative hazard function tends to zero, we can consider that the hazard function is the difference between the cumulative hazard at the end and start of the time interval (taking into account the lengths of the time intervals; Landes et al. 2017). Second, the estimation of the predicted hazard function can be achieved by smoothing the cumulative hazard function using a spline and returning the derivative of the predicted smoothing (Roshani and Ghaderi 2016; T. Therneau, *personal communication*). Researchers are faced with the difficult question of how much smoothing to apply, and, unfortunately, there is a lack of theory about this. The steps that we followed in tutorial examples 1, 3, and 4 were as follows: (1) estimate the cumulative hazard function and (2) apply a smoothing spline of the cumulative hazard function over time, comparing results across a range of 3–6 degrees of freedom. If too many degrees of freedom are used, the smoothing spline may not be monotonic, and (3) predict the values of the derivative of the smoothing spline (T. Therneau, *personal communication*).

CONCLUSION

Event history analyses are powerful methods that can be used to assess many ecological questions. These methods have long been exclusively used by medical, demographic, and industrial studies, but ecologists have only recently started to explore the potential of these methods in their field (Aubry et al. 2011, Landes et al. 2017). These methods are flexible, powerful, and reduce the influence of various biases usually encountered in ecological studies, such as sampling biases. Semi- and fully parametric models are available, and different metrics can be estimated to describe the time of emergence of the event of interest. We expect an increasing number of studies using event history analyses in the future.

LITERATURE CITED

Allison, P. D. 2014. Event history and survival analysis. Second edition. Sage Publications, Thousand Oaks, California, USA.

Armitage, P. 1959. The comparison of survival curves. Journal of the Royal Statistical Society: Series A (General) 122:279–300.

Aubry, L. M., E. Cam, D. N. Koons, J. Y. Monnat, and S. Pavard. 2011. Drivers of age-specific survival in a long-lived seabird: contributions of observed and hidden sources of heterogeneity. Journal of Animal Ecology 80:375–383.

Bannier, F., S. Tebbich, and B. Taborsky. 2017. Early experience affects learning performance and neophobia in a cooperatively breeding cichlid. Ethology 123:712–723.

Bérubé, C. H., M. Festa-Bianchet, and J. T. Jorgenson. 1999. Individual differences, longevity, and reproductive senescence in bighorn ewes. Ecology 80:2555–2565.

Blossfeld, H., and G. Rohwer. 2002. Techniques of event history modeling: new approaches to causal analysis. Second edition. Lawrence Erlbaum Associates, Mahwah, New Jersey, USA.

Boonekamp, J. J., M. Salomons, S. Bouwhuis, C. Dijkstra, and S. Verhulst. 2014. Reproductive effort accelerates actuarial senescence in wild birds: an experimental study. Ecology Letters 17:599–605.

Brent, L. J. N., A. Ruiz-Lambides, and M. L. Platt. 2017. Family network size and survival across the lifespan of female macaques. Proceedings of the Royal Society B: Biological Sciences 284:20170515.

Bronikowski, A. M., J. Altmann, D. K. Brockman, M. Cords, L. M. Fedigan, A. Pusey, T. Stoinski, W. F. Morris, K. B. Strier, and S. C. Alberts. 2011. Aging in the natural world: comparative data reveal similar mortality patterns across primates. Science 331:1325–1328.

Cain, K. C., S. D. Harlow, R. J. Little, B. Nan, M. Yosef, J. R. Taffe, and M. R. Elliott. 2011. Bias due to left truncation and left censoring in longitudinal studies of developmental and disease processes. American Journal of Epidemiology 173:1078–1084.

Congdon, P. 2007. Survival and event history analysis. Pages 457–491 in Bayesian statistical modelling. Second edition. John Wiley & Sons, Hoboken, New Jersey, USA.

Everitt, B. S. 2002. A handbook of statistical analyses using S-PLUS. Second edition. Chapman & Hall, CRC, Boca Raton, Florida, USA.

Gaillard, J.-M., A. Viallefont, A. Loison, and M. Festa-Bianchet. 2004. Assessing senescence patterns in populations of large mammals. Animal Biodiversity and Conservation 27:47–58.

Griffin, W. A. 1993. Event history analysis of marital and family interaction: a practical introduction. Journal of Family Psychology 6:211–229.

Hougaard, P. 1995. Frailty models for survival data. Lifetime Data Analysis 1:255–323.

Hougaard, P. 2000. Analysis of multivariate survival data. Springer-Verlag, New York, New York, USA.

Jackman, S. 2017. pscl: Classes and methods for R developed in the Political Science Computational Laboratory. R package version 1.5.2. United States Studies Centre, University of Sydney, Sydney, New South Wales, Australia.

Jackson, C. H. 2011. Multi-state models for panel data: the msm package for R. Journal of Statistical Software 38:1–29.

Kaplan, E. L., and P. Meier. 1958. Nonparametric estimation from incomplete observations. Journal of the American Statistical Association 53:457–481.

Klein, J. P., and M. L. Moeschberger. 2003. Survival analysis – techniques for censored and truncated data. Second edition. Springer-Verlag, New York, New York, USA.

Landes, J., M. Perret, I. Hardy, C. G. Camarda, P.-Y. Henry, and S. Pavard. 2017. State transitions: a major mortality risk for seasonal species. Ecology Letters 20:883–891.

Lee, E. T., and O. T. Go. 1997. Survival analysis in public health research. Annual Review of Public Health 18:105–134.

Moore, D. F. 2016. Applied survival analysis using R. Use R! Springer, New York, New York, USA.

Nelson-Flower, M. J., E. M. Wiley, T. P. Flower, and A. R. Ridley. 2018. Individual dispersal delays in a cooperative breeder: ecological constraints, the benefits of philopatry and the social queue for dominance. Journal of Animal Ecology 87:1227–1238.

Ricklefs, R. E., and A. Scheuerlein. 2001. Comparison of aging-related mortality among birds and mammals. Experimental Gerontology 36:845–857.

Roshani, D., and E. Ghaderi. 2016. Comparing smoothing techniques for fitting the nonlinear effect of covariate in Cox models. Acta Informatica Medica 24:38–41.

Schweinfurth, M. K., and M. Taborsky. 2017. The transfer of alternative tasks in reciprocal cooperation. Animal Behaviour 131:35–41.

Sedgwick, P. 2014. How to read a Kaplan–Meier survival plot. British Medical Journal 349:g5608.

Sutradhar, R., and P. C. Austin. 2018. Relative rates not relative risks: addressing a widespread misinterpretation of hazard ratios. Annals of Epidemiology 28:54–57.

Vaupel, J. W., et al. 1998. Biodemographic trajectories of longevity. Science 280:855–860.

Vaupel, J., K. G. Manton, and E. Stallard. 1979. The impact of heterogeneity in individual frailty on the dynamics of mortality. Demography 16:439–454.

Watson, S. P. 2011. Poisson regression models for interval censored count data. Dissertation. Baylor University, Waco, Texas, USA.

Data Availability

The first two data sets used can be found in Data S1. The Baboons data set is available in the R package ssym. The last data set is available on Dryad (DOI: https://doi.org/10.5061/dryad.km2s0).

Supporting Information

Additional Supporting Information may be found online at: http://onlinelibrary.wiley.com/doi/10.1002/ecs2.3238/full