Ecology, 67(1), 1986, pp. 246–250

ECOLOGICAL USE OF FAILURE TIME ANALYSIS[1]

GAYLE MUENCHOW[2]
EPO Biology Department, University of Colorado, Boulder, Colorado 80309 USA

Abstract. Failure time analysis (FTA), or survival analysis, addresses data of the form "time until an event occurs." The survival times of medical patients or industrial products have been the usual subjects of FTA, but data from a wide variety of ecological studies may be cast in these terms, including survival times of organisms or parts of organisms and times until certain behaviors are exhibited. FTA can accommodate censored data: cases in which the actual time of occurrence is not known but a minimum time during which the event did not occur is known. As an example, I used FTA probability distributions and the Cox Model, a nonparametric multiple regression analog, to analyze data in which the event of interest was the arrival of any flying insect onto a male versus a female flower of *Clematis ligusticifolia*. It was determined that males were visited significantly faster than females.

Key words: Clematis ligusticifolia; *Cox Model; failure time analysis; proportional hazards; survival analysis; survivorship function.*

INTRODUCTION

This paper discusses the analysis of ecological data of the form "time until an event occurs." In industry, this has been termed "failure time analysis," since the event of interest is commonly the failure of an industrial product. In the biomedical context, it has been called "survival analysis," since the event is commonly the death of a patient, so the time until the event is the survival time. I present an analysis of ecological data in which the event is the arrival of an insect onto a flower.

Failure time analysis accommodates "censored" data. Censored data points are those in which the event was not observed, perhaps because the study ended before the event (failure, death, insect visit) happened to some of the individuals under observation or because some of the individuals were lost track of before the event occurred during the study. For these censored data points, the actual time of occurrence is not known. Instead we know a minimum length of time during which the event did not occur. Failure time analysis allows use of such censored data for their partial information. This feature is apt to be useful in field biology, where identification markers may be lost, external conditions may cause the premature end of observations, or the observation period may be too brief for all possible events to occur.

In the first section of this paper, I describe failure time distributions and the probability functions that are used in statistical tests. In the second section, I discuss and illustrate tests for comparing the failure time distributions of two groups, and, in the third section, I describe the Cox Model, a proportional hazards statistical model that is analogous to multiple regression analysis. The tests I discuss are nonparametric. If one knows a priori that the failure times fit a particular distribution, then the proportional hazards model can be modified to use that information, i.e., it can be made parametric. Also, one can choose the procedure for comparing two groups that is the most powerful given a known underlying distribution. These techniques are discussed in several recent texts (Kalbfleisch and Prentice 1980, Lee 1980, Nelson 1982).

A number of ecological questions can be phrased in terms of "time until an event occurs." Events of interest might include the arrival of a migrant or parasite, the display of a particular behavior, the dispersal of a fruit or offspring, the germination of a seed, the abscission of a flower, or the death of an organism or a particular part of an organism. One pertinent way to use failure time analysis (FTA) is in determining whether male or female flowers of a species are visited more frequently by insects. The classical approach to answering this question does not use FTA: count the number of visits to a fixed number of flowers during a fixed interval and compare the average numbers for each gender. The FTA approach is slightly different: measure the length of time from an arbitrary starting point until the first observed visit and compare the distributions of the lengths of time for each gender.

The FTA approach has the advantage of using the entire probability distributions of visits for comparisons between the two genders. The classical approach uses the average number of visits per specified unit time and makes comparisons at this scale only. Also, with the classical approach, if more than one visit occurs during the fixed interval the question of whether the probability of a second visit within a certain time is independent of the first visit becomes a problem. Since the FTA approach scores only the first visit, it does not have this potential problem. Finally, of course, the FTA approach permits the analysis of censored data.

A FIELD EXAMPLE

Male plants of dioecious species are often more floriferous than female plants (Lloyd and Webb 1977). This is true of *Clematis ligusticifolia* Nutt., the species

[1] Manuscript received 21 September 1984; revised and accepted 27 February 1985.

[2] Present address: Department of Botany, Ohio University, Athens, Ohio 45701-2979 USA.

TABLE 1. Distribution of waiting times for insect visits to *Clematis ligusticifolia* flowers.

Waiting time (min)*					
Male flowers			Female flowers		
1	9	27	1	19	57
1	9	27	2	23	59
2	9	30	4	23	67
2	11	31	4	26	71
4	11	35	5	28	75
4	14	36	6	29	75†
5	14	40	7	29	78†
5	14	43	7	29	81
6	16	54	8	30	90†
6	16	61	8	32	94†
6	17	68	8	35	96
7	17	69	9	35	96†
7	18	70	14	37	100†
8	19	83	15	39	102†
8	19	95	18	43	105†
8	19	102†	18	56	
		104†			

* The single datum for each of the 96 pairs of flowers is the time from the start of observations of the selected pair to the first visit of a flying insect to one flower of the pair.
† Censored datum.

in which I tested whether males and females are equally attractive to insects against the alternative hypothesis that males are more attractive. The data were collected in Matthews-Winters Park, Jefferson County, Colorado. In 1983 the plants bloomed from about 15 July to 15 August. Between 0700 and 1500 on 15 days throughout the blooming period, I recorded waiting times until insects arrived, as follows. I chose two target flowers on the same plant, at least 30 cm apart but close enough to allow careful simultaneous observation. The choice of two target flowers rather than one was necessitated by the overall slow arrival rate of insects. When watching a single flower, I was in danger of observing so few visits that any analysis would be compromised by small sample size. The event was defined as the arrival of any flying insect at one of the flowers. The flowers are shallow and capable of being pollinated by a wide variety of insects, as is typical of dioecious flowers (Bawa 1980, G. Muenchow, *personal observation*), so I did not restrict my attention to any one insect group. The waiting times were recorded to the nearest minute. I used ≈ 20 plants of each gender and usually collected 6–8 data points in a day. Altogether, 96 cases (pairs of target flowers) were observed during a total of 3180 min (53 h). In 10 of these cases I observed no insects; these data were censored after 75 or more minutes of observation (see Table 1). Other variables recorded were the time of day, the temperature, and an estimate of the flower density within ≈ 1 m of the targets (categorized as <50, 50–100, 100–200, or >200 flowers in a circle of ≈ 1 m radius).

FAILURE TIME DISTRIBUTIONS

Let T be a random variable that represents the failure time of an individual. The probability distribution of T can be represented in many ways, but two are particularly useful in failure time analysis. They are the survivorship function and the hazard function. Given one of these, the other can be derived.

The survivorship function, $S(t)$, is the probability that the event occurs at some time T at least as great as time t, i.e.,

$$S(t) = P(T \geq t).$$

The hazard function, $h(t)$, is the *conditional* probability that the event occurs exactly at time t, given that it has not occurred before then, i.e.,

$$h(t) = P(T = t \mid T \geq t) \quad \text{for the discrete case,}$$

$$h(t) = \lim_{\Delta t \to 0} \frac{P(t \leq T \leq t + \Delta t \mid T \geq t)}{\Delta t} \quad \text{for the continuous case.}$$

The continuous case can be treated just like the discrete case, letting the time intervals be infinitesimal.

The survivorship function, $S(t)$, is the probability that the event does not occur before some time t. One can estimate it from a data set by calculating the frequency of cases in which the event had not occurred by time t, but censoring causes problems with such a

TABLE 2. The "product-limit" estimate of the survivorship function, $S(t)$,* at each of the first five occurrence times for each gender.

Male flowers†			Female flowers‡		
t (min)	r§	$S(t)$§	t (min)	r	$S(t)$
1	1	0.9796	1	1	**0.9787**
1	2	**0.9592**	2	2	**0.9574**
2	3	0.9388	4	3	0.9361
2	4	**0.9184**	4	4	**0.9148**
4	5	0.8980	5	5	**0.8935**
4	6	**0.8776**	6	6	**0.8722**
5	7	0.8572			
5	8	**0.8368**			
6	9	0.8164			
6	10	0.7960			
6	11	**0.7756**			

* $S(t)$ is estimated at each individual occurrence time t by:

$$S(t) = \prod_{t_{(r)} \leq t} \frac{N - r}{N - r + 1},$$

where N is the sample size, r is the rank, and $t_{(r)}$ is an uncensored datum (i.e., the observation interval was terminated because the event of interest occurred, rather than being terminated for some other reason). $S(t)$ corresponds to "the probability that the event did not occur in the first interval" times "the probability that it did not occur in the second interval given that it did not occur in the first," etc.

† $N = 49$ flowers.
‡ $N = 47$ flowers.
§ Within a group of tied occurrence times, the order of ranks is arbitrary; $S(t)$ was estimated for each r, but in practice the $S(t)$ value used for each of the tied observations in the group is the lowest of the group's $S(t)$ values (depicted here in boldface type), because this value is the most conservative.

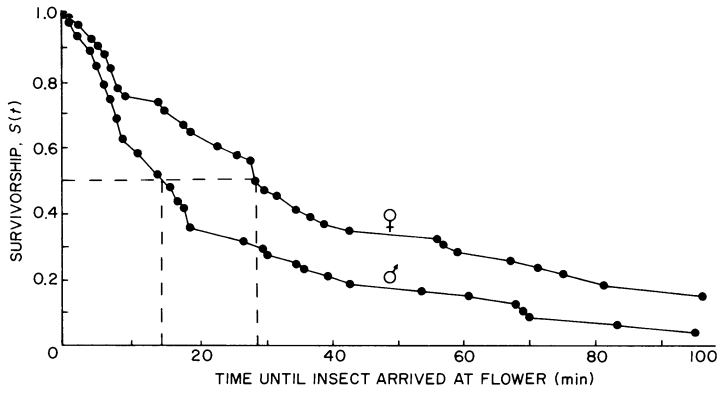

FIG. 1. "Product-limit" estimates of the survivorship functions, $S(t)$, for the male and female flower pairs of *Clematis ligusticifolia*. $S(t)$ is the probability that a flower has not yet been visited by an insect; "time until insect arrived" refers to the time elapsed since the start of an observation period.

calculation. Cases censored before time t cannot be counted as failures, because we do not know whether the actual failures occurred before, at, or after time t. To accommodate this, the total time span is broken into intervals. Each censored case is counted among those at risk during each interval that took place before the censoring. When the censored cases are lost from the study, the succeeding intervals simply have a smaller number of individuals at risk. The probability, $S(t)$, that an event does not occur before the interval beginning at time t is considered to be the product of conditional survival times $S(0)\cdot S(1\,|\,0)\cdot S(2\,|\,1)\cdot\ldots\cdot S(t-1\,|\,t-2)$, that is, "the probability that failure does not occur in the first interval" times "the probability that failure does not occur in the second interval given that it did not occur in the first" . . . and so on up to interval $t-1$. For each interval, the conditional probability is estimated by the proportion of individuals entering the interval who do not fail during that interval. The initial value is taken to be 1.0. Censoring may cause the survivorship function to remain greater than zero at the end of the observation period.

The product-limit method of estimating the survivorship function (Kaplan and Meier 1958) calculates the estimate at each individual occurrence time, whereas the life-table method groups occurrence times into intervals. The latter may be more convenient for large sets of data. SPSS (Hull and Nie 1981:205–219) has a procedure called SURVIVAL that does the latter. Lee (1980) gives computer programs for both, as well as hand calculating methods for both.

Table 2 shows the first several calculations of the product-limit estimates of the survivorship functions for the two genders in the example, and Fig. 1 illustrates these estimates for the male and female pairs of *Clematis ligusticifolia* flowers. In this case, "survivorship" is the probability that an insect has not yet visited either of the target flowers a given number of minutes after the start of observation. It is apparent from Fig. 1 that females are consistently more likely not to have been visited by a given time. One may also read percentiles directly off a graph of survivorship. The time that corresponds to $S(t) = 0.5$ is the time at which 50% of the population is estimated to have been visited. In Fig. 1, 50% of the male pairs have been visited by 14 min, and 50% of the female pairs by 29 min.

The survivorship function is commonly used to compare groups. The distribution of times until first visit to males and females can be seen in Table 1. Permutation theory, which is nonparametric, allows one to ask: If the males and females are randomly drawn from the same occurrence time distribution, what is the probability of this or a more extreme arrangement? There are several nonparametric test statistics, all rank statistics, that summarize the differences between whole survivorship curves (rather than the difference at any prespecified time). Lee (1980) shows how these are calculated by hand, and she also provides a FORTRAN listing for a program that calculates five of the two-sample tests. The tests differ slightly in power (the ability to reject a false hypothesis), depending upon whether there is censoring, how large the sample size is, and what the true underlying distribution is. They also weight early and late points a little differently. Lee (1980) summarizes the differences. In my example, the samples seem drawn from exponential distributions, because a plot of $\ln[S(t)]$ vs. t gives a straight line with a negative slope. In such a case the Cox-Mantel test can be expected to be more powerful than some. The Cox-Mantel test rejected the null hypothesis of equal visitation probabilities at $P = .0085$. I conclude that male flowers were visited at a significantly faster rate than were female flowers.

These statistics depend upon one's being able to rank the events in their true order. If some of the data points are tied in rank, i.e., the lengths of time until the event occurred are the same, those points cannot be ordered with respect to each other. They are, therefore, not fully informative, and this reduces the power of the tests. Ties can be accommodated in the tests, but much is

gained by minimizing them. My pollination data have, at worst, six failure events at the eighth minute; because such a high number of ties begins to compromise the power with which I can make inferences, I should have recorded times in seconds instead of minutes.

The Cox Model, A Proportional Hazards Model

Additive statistical models for failure time analysis, such as

$$\text{Time of occurrence} = \beta_1(\text{covariate 1}) + \beta_2\,(\text{covariate 2}) + \ldots,$$

need modification to deal with censored observations. Satisfactory FTA models usually involve the assumption of "proportional hazards," explained below. The Cox Model (Cox 1972) is a general nonparametric model that can accommodate censoring. It can be written as

$$h_i(t, \boldsymbol{z}) = h_0(t) \exp\left(\sum_j \beta_j z_{ji}\right).$$

In the example, the covariates (the z's) are gender, flower density category, temperature, and time at which the observation period started. The Cox Model for this is $h_i(t) = h_0(t) \exp[\beta_1(\text{gender}) + \beta_2\,(\text{flower density}) + \beta_3\,(\text{temperature}) + \beta_4\,(\text{start time})]$. Recall that $h(t)$ is the hazard function and is the probability that the event happens at time t given that it has not happened before time t. What this model says, then, is that the hazard function at time t for individual i, who has a particular set of values for the covariates, is equal to $h_0(t)$, which is some unknown base hazard function, multiplied by $\exp[\beta_1\,(\text{gender}) + \beta_2\,(\text{flower density}) + \beta_3\,(\text{temperature}) + \beta_4\,(\text{start time})]$. The last term multiplies the base hazard rate by some number that depends upon the j coefficients (the β's) and upon the values that the covariates take for individual i. If the coefficients are all zero, i.e., the covariates have no effect, the term exp(0) equals 1.0 and the hazard function for individual i just equals the base hazard function $h_0(t)$. If the coefficients are different from zero, the exp term is positive and the hazard function for individual i is some multiple of the base hazard function. That multiple is determined for individual i by its set of covariate values. In other words, individual i's hazard function is proportional to the base hazard function, and, by extension, proportional to the hazard functions of the other individuals in the sample. That is why it is called a proportional hazards model. In the example, the model says that the probability of a visit to a flower pair, given that it has not yet been visited, is a function of some underlying insect visitation rate; the probability is increased (exp term >1.0) or decreased (exp term between 0.0 and 1.0) depending upon the covariate values of pair i. One of the covariates is gender. The test whether the coefficient of the gender term is different from zero tests, in the presence of other covariates, whether gender influences the rate at which flowers are visited.

The base hazard function does not necessarily make biological sense by itself. One would usually not be interested in estimating it. We want to estimate the β's and test whether they are significantly different from zero. A Cox Model computer program will simply print the estimates and their P values, so it is quite simple to use this statistical procedure.

Recall that the data are a string of times at which we know events occurred or, in the censored cases, we know that events had not yet occurred. The program finds the values of the β's that maximize the probability of observing that string of events. It handles both ties and censored data by computing all the strings of event times associated with all the possible rankings given those tied and censored points. Since this rapidly gets to be a huge number of possible strings of events, the program works best when there are relatively few ties or censored data points.

The probability statement of the string of events is

$$\prod_{t=1}^{N} \frac{h_0(t_i)\exp\left(\sum_j \beta_j z_{ji}\right)}{\sum_{l \in R(t)} h_0(t_i)\exp\left(\sum_j \beta_j z_{jl}\right)},$$

where N is the total number of observed waiting times and $R(t)$ is the set of cases l in which the event has not yet been observed. This is written in terms of the hazard function. The fractional part is the hazard for individual i at time t divided by the sum of the hazards for every individual still at risk at time t. That is the probability that it was individual i who experienced the event at time t. Then, to get the probability of the string of events over all the times, the probabilities at each of the times are multiplied together. The derivative of this function is set equal to zero, and the corresponding values for the β's found. It is not a linear function, so its solution requires iterative estimation techniques, such as the Newton-Raphson technique employed by the computer program I used. The development of this area of statistics has depended upon the advent of computer technology.

The Cox Model computer program I used in this analysis (Cox Model, a Proportional Hazards Model Analysis Package for SPSS Users) was written by Lawrence J. Emrich and Peter A. Reese at the Computer Center of Roswell Park Memorial Center, Buffalo, New York, and by John D. Kalbfleisch of the Department of Statistics, University of Waterloo, Waterloo, Ontario, Canada. It is written in FORTRAN IV. I used it on a TANDEM computer at the University of Colorado Health Sciences Center.

The covariates in the example are starting time, temperature, flower density, and gender. Starting time and temperature did not influence the effect of gender on the event times. Gender and flower density were correlated. Male plants bore more flowers in denser inflo-

rescences. I stratified to separate the influences of gender and density: within a gender, denser flower groups attracted insects at a greater rate ($P = .039$); within density categories, males were visited faster than females ($P = .049$). Thus, males were more attractive to insects both because they bore flowers more densely and because they had some other (unknown) attractive character(s), perhaps the reward of pollen.

ACKNOWLEDGMENTS

I am grateful for the helpful comments of James Murphy, Richard Hopkins, Jane Bock, Jeffrey Mitton, Stephen Bernstein, Michael Grant, Yan Linhart, Andrew Stephenson, Duane Meeter, Donald Strong, and an anonymous reviewer.

LITERATURE CITED

Bawa, K. S. 1980. Evolution of dioecy in flowering plants. Annual Review of Ecology and Systematics **11**:15–40.

Cox, D. R. 1972. Regression models and life tables. Journal of the Royal Statistical Society (London) Series B **26**:103–110.

Hull, C. H., and N. H. Nie. 1981. SPSS Update 7–9. McGraw-Hill, New York, New York, USA.

Kalbfleisch, J. D., and R. L. Prentice. 1980. The statistical analysis of failure time data. John Wiley and Sons, New York, New York, USA.

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

Lee, E. T. 1980. Statistical methods for survival data analysis. Lifetime Learning Publications, Belmont, California, USA.

Lloyd, D. C and C. J. Webb. 1977. Secondary sex characters in plants. Botanical Review **43**:177–216.

Nelson, W. B. 1982. Applied life data analysis. Wiley, New York, New York, USA.