

The Morality of Euthanasia

Whatever the specific objections to relaxation of prohibitions on physician-assisted suicide and active voluntary euthanasia may be, opponents have fostered fear of a "slippery slope" as one of the lurking dangers of such a course. This means a fear that voluntary euthanasia shall soon pass to involuntary euthanasia; that competent patients requesting death shall soon pass to incompetent persons who have no idea what they are requesting (or for whom euthanasia is requested by someone else); that letting physicians prescribe pills for competent patients who voluntarily request assistance in dying shall soon pass to physicians killing off patients in all kinds of situation, for all sorts of nefarious reasons; that assisted death and euthanasia as socially approved options shall soon pass to circumstances in which patients feel pressured for various reasons to choose one of these options, even when they would otherwise want to prolong their lives.

Opponents of change almost always urge slippery-slope objections not only against social changes but against even thinking about changes, and this has long been the case with active voluntary euthanasia. In fact, good evidence is needed to believe such claims. However, for many opponents of active voluntary euthanasia, slippery-slope concerns are at bottom not the main argument. In circumstances where it was clear that slippery-slope effects were exceedingly unlikely to arise, few opponents to active voluntary euthanasia would remove their objections.

Most religions continue to insist that life is a gift from God that can only be justly taken by God. To many adherents it follows that quality of life concerns are not centrally germane to the whole issue of the prolongation of life. Advocates believe the fact that physician-assisted suicide and active voluntary euthanasia give voice to the value of autonomy, to the demands of patients to have a say in their own treatment, and to the wish of many to be able to decide not only how they will live but also how they will die, is unlikely to carry the day as a moral argument.

See also: *Religions, Population Doctrines of; Reproductive Technologies: Ethical Issues.*

BIBLIOGRAPHY

- Beauchamp, Tom. L., ed. 1996. *Intending Death: The Ethics of Assisted Suicide and Euthanasia*. Englewood Cliffs, NJ: Prentice-Hall.
- Brock, Dan. 1993. *Life and Death: Philosophical Essays in Biomedical Ethics*. Cambridge, Eng.: Cambridge University Press.
- Callahan, Daniel. 1993. *The Troubled Dream of Life: Living with Mortality*. New York: Simon and Schuster.
- Dworkin, Gerald., R. G. Frey, and Sissela Bok. 1998. *Euthanasia and Physician-Assisted Suicide*. Cambridge, Eng.: Cambridge University Press.
- Rachels, James. 1986. *The End of Life: Euthanasia and Morality*. Oxford, Eng.: Oxford University Press.
- VanDerMass, P. J., G. van der Wal, I. Haverkate, et al. 1996. "Euthanasia, Physician-Assisted Suicide, and Other Medical Practices Involving the End of Life in The Netherlands." *New England Journal of Medicine* 335: 1699-1705.
- Verhovek, S. H. 1999. "Oregon Reporting 15 Deaths in 1998 Under Suicide Law." *New York Times*, February 18, 1998, A1.

R. G. FREY

EVENT-HISTORY ANALYSIS

Event-history analysis is a set of statistical methods designed to analyze categorical or discrete data on processes or events that are time-dependent (i.e., for which the timing of occurrence is as meaningful as whether they occurred or not). One example of such time-dependent processes is mortality: variation across individuals is not captured by the lifetime probability of dying (which is one for every individual), but by differences in the age at which death occurs. Another example is marriage: here, variation across individuals is captured by both the lifetime probability of getting married and differences in age at marriage.

Event-history analysis, sometimes called survival analysis, has applications in many fields, including

sociology, economics, biology, medicine, and engineering. Applications in demography are particularly numerous, given demography's focus on age and cohorts. In addition to mortality, demographic events that can be investigated with event-history analysis include marriage, divorce, birth, migration, and household formation.

Comparison to Life Table Analysis

Event-history analysis has its roots in classical life table analysis. In fact, life table analysis is one of the methods covered by event-history analysis, and many of the concepts of event-history analysis, such as survival curves and hazard rates, have equivalents in a conventional life table. One difference from life table analysis is that event-history analysis is based on data at the individual level and aims at describing processes operating at that level. Also, whereas conventional life table analysis is deterministic, event-history analysis is probabilistic. Hence, many event-history analysis outcomes will have confidence intervals attached to them. Another feature of event-history analysis relative to conventional life table analysis is the use of covariates. Event-history analysis makes it possible to identify factors associated with timing of events. These factors can be fixed through time (such as ethnicity or parents' education), or vary with time (such as income and marital status).

Whereas conventional life table analysis can be applied to both longitudinal and cross-sectional data, event-history analysis requires longitudinal data. Longitudinal data can be collected either in a prospective fashion by following individuals through time, or retrospectively by asking individuals about past events.

Censored Data and Time-Varying Covariates

Because of its longitudinal nature, event history data have some features which make traditional statistical techniques inadequate. One such feature is *censoring*, which means that information on events and exposure to the risk of experiencing them is incomplete. *Right censoring*, the most common type of censoring in event-history analysis, occurs when recording of events is discontinued before the process is completed. For example, in longitudinal data collection, individuals previously included in a sample may stop contributing information, either because the study is discontinued before they experience the

event of interest, or because they discontinue their participation in the study before they experience the event. Another, less common, type of censoring is *left censoring*, which occurs when recording is initiated after the process has started. In the remainder of this article, censoring will refer to right censoring.

It is important to include censored individuals in event-history analysis, because the fact that they did not experience the event of interest in spite of their exposure is in itself meaningful. Censoring can be handled adequately as long as it is independent—that is, as long as the risk of being censored is not related to the risk of experiencing the event, or, equivalently, provided that individuals censored at any given time are representative of all other individuals. If the two risks are related, however, the estimates obtained can be seriously biased.

Another particular feature of survival data is the potential presence of time-varying covariates. For example, an individual's income may vary over time, and these variations may have an effect on the risk of experiencing events. If this is the case, it is important to include information on these variations in the analysis.

Unlike traditional statistical techniques such as ordinary least squares (OLS), event-history analysis can handle both censoring and time-varying covariates, using the method of maximum likelihood estimation. With the maximum likelihood approach, the estimated regression coefficients are the ones that maximize the likelihood of the observations being what they are. That is, the set of estimated coefficients are more likely than any other coefficient values to have given rise to the observed set of events and censored cases.

Hazard Rates

An important concept in event-history analysis is the hazard rate, $h(t)$. The hazard rate is the risk or hazard that an event will occur during a small time interval, $(t, t+dt)$. It corresponds to the rate of occurrence of an event (number of occurrences/amount of exposure to the risk of occurrence) during an infinitesimal time or age interval. If the event under study is death, then the hazard rate is called the force of mortality, $\mu(x)$, where x is age. Event-history analysis can be used to explore how hazard rates vary with time, or how certain covariates affect the level of the hazard rate.

Types of Analysis

Methods of event-history analysis fall into three categories:

1. Nonparametric, in which no assumption is made about the shape of the hazard function;
2. Parametric, requiring an assumption about how the hazard rate varies with time; and
3. Semiparametric, requiring an assumption about how the hazard rate varies across individuals but no assumption about its overall shape.

Nonparametric Models

The life table approach to analyzing event history data is a nonparametric method. It is very similar to traditional life table construction in demography, although it is based on cohort rather than period data. The logic behind the life table approach is to calculate $Q(t_i)$, the probability of "failing" (for instance, dying) in the interval $[t_i, t_i+n]$, from data on $N(t_i)$, the number of individuals at risk of failing at time t_i , and $D(t_i)$, the number of failures between t_i and t_i+n . The number of individuals at risk needs to be adjusted for the fact that some individuals, $C(t_i)$, will be censored—that is, removed from the risk of experiencing the event during the interval. Hence $Q(t_i)$ can be expressed as:

$$Q(t_i) = \frac{D(t_i)}{N(t_i) - 0.5 \cdot C(t_i)}$$

The proportion of persons surviving at time t_i , $S(t_i)$, is then obtained as the product of the probabilities of surviving over all earlier time intervals as shown below.

$$\hat{S}(t_i) = \prod_{j=0}^{i-1} [1 - Q(t_j)]$$

Another output of the life table method is the hazard rate, $h(t_i)$, which is simply calculated by dividing the number of events experienced during the interval t_i by the number of person-years lived during the interval. The number of person-years is estimated by assuming that both failures and censored cases occur uniformly through the interval. Hence $h(t_i)$ is given by:

$$h(t_i) = \frac{D(t_i)}{n \cdot (N(t_i) - 0.5 \cdot D(t_i) - 0.5 \cdot C(t_i))}$$

The above equations can produce biased results when time intervals are large relative to the rate at which events occur. If failures and censored cases are recorded with exact time, it is possible to correct for these biases by use of what is known as the Kaplan-Meier method. Suppose that d_j is the number of deaths at exact time t_j , and that N_j is the number of persons at risk at time t_j . The Kaplan-Meier estimator of the survival curve $S(t)$ is defined as:

$$\hat{S}(t) = \prod_{j: t_j \leq t} [1 - (d_j / N_j)]$$

where N_j is obtained by subtracting all failures and censored cases that occurred before t_j from the initial size of the cohort. Compared to the life table method, the Kaplan-Meier method produces a more detailed contour of the survival curve. It is more appropriate than the life table approach when the recording of events is precise. The Kaplan-Meier method permits calculation of confidence intervals around the survival curve and the hazard rate. It also makes it possible to calculate survival curves for two or more groups with different characteristics, and to test the null hypothesis that survival functions are identical for these groups.

Parametric and Semiparametric Models

Although nonparametric life table approaches can perform some tests across groups, they do not permit direct estimation of the effect of specific variables on the timing of events or on the hazard rate. In order to estimate such effects, one needs to use regression models that fall into the category of fully parametric or semiparametric methods.

Accelerated failure-time models. The most common fully parametric models are called accelerated failure-time models. They postulate that covariates have multiplicative effects both on the hazard rate and on timing of events. They commonly take T_i , the time at which the event occurs, as a dependent variable. A general representation of accelerated failure-time models is:

$$\log T_i = \beta_0 + \beta_1 x_{i1} + \dots + \beta_k x_{ik} + \sigma \varepsilon_i$$

where T_i is the time at which the event of interest occurs for individual i , and x_{i1}, \dots, x_{ik} is a set of k explanatory variables with coefficients β , ϵ_i is an error term, and σ is a scale parameter. (Taking the logarithm of T_i ensures that the timing of events will be positive whatever the values of the covariates for a specific individual.)

This model can be adapted to various situations by choosing a specific distribution for the error term ϵ_i . Common distributions chosen include normal (when the distribution of T_i is log-normal), extreme value (when the distribution of T_i is Weibull), logistic (when the distribution of T_i is log-logistic), and log-gamma (when the distribution of T_i is gamma). Accelerated failure-time models are fully parametric precisely because they require the choice of a model distribution of failure times. Although the above equation resembles that of an OLS regression, the estimation must be performed using the maximum likelihood procedure in order to accommodate the presence of censored cases. Regression coefficients in accelerated failure time models can be interpreted by calculating $100(e^{\beta}-1)$, which is an estimate of the percentage change in the time at which the event occurs for a one-unit increase in a particular independent variable.

Proportional hazard models. Another type of regression model in event-history analysis is the proportional hazard model. Such models postulate that the set of covariates acts in a multiplicative way on the hazard rate. A general formulation of proportional hazard models is:

$$\log h_i(t) = \log h_0(t) + \beta_1 x_{i1} + \dots + \beta_k x_{ik}$$

where $h_0(t)$ is the baseline hazard that is increased or decreased by the effects of the covariates.

This model is called proportional hazard because for any two individuals the ratio of the risk of the hazard is constant over time. If the form for $h_0(t)$ is specified, the result is a fully parametric model. The most common specifications for $h_0(t)$ are the exponential, Weibull, and Gompertz models. Like accelerated failure time models, fully-parametric proportional hazard models are estimated using the maximum likelihood procedure.

Proportional hazard models can also be estimated without specifying the shape of $h_0(t)$. In an influ-

ential paper, D.R. Cox (1972) showed that if one assumes that the ratio of the hazards for any two individuals is constant over time, one can estimate the effect of covariates on hazard rates with no assumption regarding the shape of $h_0(t)$, using a "partial likelihood" approach. These models, commonly called Cox regression models, are semiparametric because of the absence of any assumption regarding the time structure of the baseline hazard rate. In order to interpret the coefficients (β_i) of such regressions, one can calculate the percent change in the hazard rate for a one-unit increase in the variable, using again the transformation $100(e^{\beta}-1)$. Cox regression models, which also can be easily adapted to accommodate time-varying covariates, are probably the most popular of available event history models.

Generalizations

In some cases it is important to distinguish among different kinds of events. For example, in demography it is sometimes necessary to focus on deaths from particular causes rather than on deaths from all causes. In such situations, individuals are being exposed to "competing risks," which means that at any time they face the risk of experiencing two or more alternative events. All the methods described above can be adapted to handle multiple events by estimating separate models for each alternative event, treating other events as censored cases. As in the case of censoring, the assumption is that risks of experiencing alternative events are independent of one another; violation of this assumption leads to biased estimates.

There are cases where the event of interest occurs in discrete time intervals. This can happen because of the nature of the event, or because the timing of events is not exactly recorded. Event-history analysis includes methods that are specifically designed for dealing with discrete time. The basic principle behind these models is to use discrete time units rather than individuals as the unit of observation. By breaking down each individual's survival history into discrete time units and pooling these observations, it is possible to estimate a model predicting the probability that the event occurs during a time interval, given that it has not occurred before. Such models are easy to implement and are computationally efficient. Also, since the unit of observation is a time interval, it is easy to include covariates taking different values for different time intervals.

All the models presented here assume that two individuals with identical values of covariates have identical risks of experiencing the event of interest. If there are no covariates in the model, the assumption is that risks are identical for all individuals. Such assumptions can be problematic in survival analysis. In fact, if some important characteristics are not accounted for, the aggregate risk may appear to decrease with time because the proportion of individuals with lower risks increases as time passes. Thus, in the presence of unobserved heterogeneity, it may be erroneous to use survival analysis to make inferences about individuals' risks. Although there are solutions to handle this potential bias, options for dealing with unobserved heterogeneity are limited and are highly sensitive to the underlying assumptions of the models.

Another implicit assumption in all the models discussed above is that events can be experienced only once, which implies that individuals are removed from the population "at risk" after they experience the event. There are many situations, however, in which events are repeatable. For example, a person who had a child or changed jobs can experience those events again. Under these circumstances, it is still possible to use single-event methods by analyzing each successive event separately, or by using a discrete-time analysis where the unit of observation is a time interval and where all time intervals, assumed to be independent for a single individual, are pooled together. However, these strategies are unsatisfactory for many reasons, and specific methods exist to deal with repeatable events. As in the case of unobserved heterogeneity, options for dealing with repeatable events are still limited.

See also: *Cohort Analysis; Estimation Methods, Demographic; Life Tables; Multistate Demography; Stochastic Population Theory.*

BIBLIOGRAPHY

- Allison, Paul D. 1995. *Survival Analysis Using the SAS System: A Practical Guide*. Cary, NC: SAS Institute.
- Cleves, Mario, William W. Gould, and Roberto Gutierrez. 2002. *An Introduction to Survival Analysis Using Stata*. College Station, TX: Stata Corporation.
- Collett, David. 1994. *Modelling Survival Data in Medical Research*. London: Chapman and Hill.
- Courgeau, Daniel, and Eva Lelièvre. 1992. *Event History Analysis in Demography*. Oxford, Eng.: Clarendon Press.
- Cox, David R. 1972. "Regression Models and Life Tables." *Journal of the Royal Statistical Society B*(34): 187-220.
- Manton, Kenneth, Eric Stallard, and James W. Vaupel. 1986. "Alternative Models for the Heterogeneity of Mortality Risks among the Aged." *Journal of the American Statistical Association* 81: 635-44.
- Palloni, Alberto, and Aage B. Sorensen. 1990. "Methods for the Analysis of Event History Data: A Didactic Overview." In *Life Span Development and Behavior*, ed. Paul B. Baltes, David L. Featherman, and Richard M. Lerner. Hillsdale, NJ: Erlbaum.
- Trussell, James, Richard K. B. Hankinson, and Judith Tilton. 1992. *Demographic Applications of Event History Analysis*. Oxford, Eng.: Clarendon Press.
- Wu, Lawrence L. 2003. "Event History Models for Life Course Analysis." In *Handbook of the Life Course*, ed. Jeylan Mortimer and Michael Shanahan. New York: Plenum.

MICHEL GUILLOT

EVOLUTIONARY DEMOGRAPHY

The fields of human demography and population biology share intellectual roots and a common set of methodological tools for describing and analyzing population processes. The two disciplines, however, have developed independently with very little cross-fertilization. They developed independently probably because human populations experienced very rapid changes in patterns of mortality and fertility from the mid-nineteenth century to the early twenty-first century, suggesting to demographers that explanations of human population processes must be inherently social rather than biological.

Evolutionary demography analyzes population processes as reflecting the optimizing force of natural selection, the process by which alternative geno-