

Writing about Hazards Models: Practical Guidelines for Effective Presentation
Jane E. Miller, Ph.D
January 2008

Jane E. Miller, Ph.D.
Research Professor
Institute for Health, Health Care Policy and Aging Research
Rutgers University
30 College Avenue
New Brunswick NJ 08901
Voicemail: (732) 932-6730; fax (732) 932-6872
Email: jmiller@ifh.rutgers.edu

Abstract:

Economists use hazards models to analyze factors associated with the occurrence and timing of events such as exiting poverty or finding a job, but their descriptions of the analysis are often confusing, incomplete, or written in ways that don't relate to their specific research question. This article reviews hazards models concepts, discusses the types of information to report for such models, and describes effective ways to present results, illustrated with example tables, charts, and sentences from published literature. An appendix of terminology and a checklist for planning and evaluating papers are also included. By following the guidelines and examples in this article, authors can improve the completeness and clarity of their papers about applications of hazards models to topics in economics and related fields.

Keywords: communication; Cox regression; duration models; event history analysis; proportional hazards models; survival analysis.

Hazards models (also known as duration or survival models) are widely used in economics and related fields, where they are used to analyze factors associated with the occurrence and timing of events such as exiting poverty, finding a job, or dying. Most people learn about these methods in courses that emphasize understanding statistical assumptions, estimating models, interpreting statistical tests, and assessing coefficients and model fit. Textbooks communicate hazards models principally using statistical lingo and equations, but few include example sentences to illustrate how to write about coefficients or model fit in ways that relate them back to the specific research topic at hand. Moreover, journal articles are strikingly inconsistent in the information provided about the data, methods, and results related to their application of hazards analysis, often resulting in confusing or incomplete exposition.

This paper demonstrates how to write about hazards models in ways that emphasize the research question to which they are applied, and that use the statistical results as evidence in a clear story line about the question (see also author citation). The paper is intended to be used as a complement to an econometrics or statistics textbook (Cox and Oakes 1984; Greene 2002; Allison 1995; Yamaguchi 1991), ideally in courses where students also learn to prepare and analyze data using hazards models.

The paper begins with an introduction to concepts and vocabulary for hazards models, and then reviews the types of research questions for which hazards analyses are used and explains what information to include in a journal article about an application of a hazards model. To show how to convey such information, this paper includes examples of tables, charts, and sentences from published articles that have done a good job of presentation, including several from health journals, which have a longer tradition of applying these types of models. Some examples are accompanied by samples of ineffective writing based on other, less effective papers, annotated to explain weaknesses. The Appendix defines the terminology for hazards models, keyed to the example topics used here. The paper also includes a checklist of essential elements and recommended approaches to writing an effective paper about an application of hazards models.

Key concepts and terminology for hazards models

General background on hazards models

Models to analyze the time to occurrence of events are known variously as *hazards models* (including Cox proportional hazards models), *duration models*, *Cox regression*, *survival models*, *event history models*, and *failure time models* (Allison 1995; Maciejewski 2002). The dependent variable in a hazards model consists of two parts: An *event indicator* (e.g., a binary indicator of whether someone found a job) and a *measure of time* from baseline to the event or censoring. *Censoring* occurs when the event under study is not observed for a given case. Left censoring occurs when the event preceded the observation period. Right censoring arises when the event either never occurred or took place after the period of observation, and can come about due to a loss-to-follow up, refusal to participate after initial enrollment, or the end of the observation period. Reasons for censoring vary depending on the type of event under study. For instance, in an analysis of poverty spells, death would constitute censoring because it removes that person from the population at risk of leaving poverty. In a study of mortality, however, death constitutes the event, not a reason for censoring.

Hazard rates (also known as *transition rates*) measure the risk of event occurrence within a specified time interval, conditional on survival to the beginning of that time interval, where *survival* is defined as not yet having experienced the event. For example, in an analysis of finding employment, people who found a job, exited the labor market, died, or were censored in the first month are not included in calculation of the hazard rate for the second or subsequent months. Thus, the hazard rate of finding a job in the second month of observation is the number of job finders in the second month divided by the number of person-months at risk (or *exposure*) in that month.

Proportional hazards imply that the ratio of hazard rates for two groups under comparison is roughly constant at all time points since baseline. For example, if the hazard of college completion for females is roughly 1.2 times as high as for males at all durations, the assumption of proportionality is met, and that multiplier (1.2) is the hazard ratio (or relative hazard; a form of relative risk) of college completion for females compared to males. When hazard rates across time are plotted for those subgroups, proportional hazards appear as approximately parallel lines. If hazard curves appreciably diverge, converge, or cross one another, the assumption of proportionality is violated, and a specification allowing for *non-proportional hazards* is needed.

Types of hazards models

Hazards models can be used to analyze a variety of research questions, with terminology varying by discipline. In event history analysis, transitions from one mutually exclusive state to another are described in terms of event occurrence. For example, in an analysis of marriage, the *event* (becoming married) is a transition from one marital state (unmarried) to another (currently married). In failure time analysis and survival analysis, these transitions are dubbed failures. And in life table models, an event or transition may be referred to as a source of *decrement* – removal from the population at risk of the event.

The simplest type of hazards analysis involves a *single-decrement, non-repeatable event* – a single type of one-way transition such as from alive to dead, or from being single to first marriage. In a single-decrement analysis, each respondent contributes exactly one *spell* (period at risk of experiencing the event) to the analysis, with that spell ending either with the event or censoring. The event indicator is binary (dichotomous), usually with a value of 1 indicating the event and a value of 0 indicating censoring. *Repeatable events* are those that can occur more than once to the same respondent during the study period. Examples include finding a job or exiting welfare. For such topics, usually some respondents do not experience any events, others experience exactly one, and a few experience two or more events.

Competing risks analysis involves more than one type of event, as in a study contrasting risk factors for several possible reasons for leaving the ranks of the unemployed – found a job, quit looking, retired, or died. Each of these mutually exclusive causes comprises a different value of the (multichotomous) event indicator, and every case is at risk of any of these competing reasons until it experiences one or is censored. For a cause-of-death competing risks analysis, mortality from one cause (such as heart disease) comprises censoring for all of the other causes (for example, cancer and

pneumonia), and the transitions are obviously one-way, so each respondent contributes exactly one spell. In analyses of repeatable-event competing risks, such as no longer being unemployed, each time someone becomes unemployed initiates a new period at risk for that individual.

Increment-decrement or *multistate life tables* analyze transitions among different states, with movement possible in more than one direction, such as back and forth between being employed and unemployed, or between being poor and non-poor. For such types of analysis, people can be added to or subtracted from the population at risk of each type of event, rather than only subtracted as in a single-decrement analysis such as mortality.

A diagram can be a particularly effective way to identify all possible transitions among the different states in an increment-decrement or competing risks analysis. For instance, Figure 1 illustrates the possible transitions in a hypothetical study of employment status. Arrows run in both directions among the first three states (unemployed, employed, not in labor force), but only in one direction from each of those states to death. This analysis combines repeatable events and competing risks, analyzing a total of nine types of transitions, with one transition for each arrow (labeled F_{nm} on the diagram, where n and m represent different states).

Figure 1 about here

The data and methods

Several important elements of the data and methods section of a paper about a hazards analysis differ from other types of multivariate statistical analysis. As noted, the dependent variable in a hazards analysis has two parts – the event indicator and a measure of time at risk of the event, alternatively termed the *period at risk*, *spell*, *exposure*, *episode*, or *duration*. In the data section, identify the kind of event(s) under study, including the definition of each event, possible reasons for censoring, and sources of data about event occurrence and timing, the maximum possible length of the follow-up, and dates or intervals of follow-up. Report the units of measurement for the duration measure (e.g., days, months, or years), and whether it is specified as continuous or discrete. For instance, in their paper about the impact of welfare benefits on conjugal status of single mothers, Lefebvre and Merrigan (1998) wrote:

“The General Social Survey, carried out by Statistics Canada in 1990 as a supplement of the Labour Force Survey, collected the marital and parental histories of a sample of almost 13,500 individuals who were 15 years of age or older at the time of the interview. The survey incorporates retrospective information on two types of union: marriage and common-law relationships. For each union reported by respondents, the date, month and year of the beginning of the union are known, as well as the date of a dissolution, where applicable... This data set enables us to piece together all the episodes and the duration of spells of single parenthood experienced by the women reached for the survey. An episode of single parenthood is defined here as any period, irrespective of its duration, during which a woman lives without a partner and had at least one ‘dependent’ child (p. 748).”

Sample size

Another aspect of the data and methods section that differs for a hazards analysis is sample size, which is affected by length of follow-up, rates of censoring, and – for repeatable events or increment-decrement analysis – more than one spell for some respondents. These issues often cause effective sample size to differ from the number of cases at baseline. Unlike logit models where exposure is assumed to be the same for all cases (for example, each case contributes 1 unit to the denominator of the hazard rate), in hazards models, exposure varies across cases depending how long they are observed. In long-term studies, there can be wide variation in the person-time contributed by different cases. For example, in a two-decade study, some cases might be lost in the first few months due to mortality, moving away, or other reasons, while other cases might contribute as many as 20 person-years of observation. See Yamaguchi (1991) for an excellent review of questions to consider when constructing the data set and calculating periods at risk.

Hazard rates for each time interval are based on person-time at risk in that interval. In studies with long follow-up periods and high rates of attrition or event occurrence, the population at risk at later time points can be substantially smaller than the population at baseline. In multivariate models, the problem of small sample sizes at later durations is compounded because hazard rates are calculated for subgroups.

Single-event analysis

For studies of single-event topics, provide information on number of cases, exposure time, and the sources of data for measures of time and event occurrence. For example, to report sample size in their analysis of how public policy and economic conditions affect women's chances of returning to welfare, Hofferth and colleagues (2005) wrote:

“Two analysis files were created. For the descriptive analysis of spell durations, we created a spell-based file, with one observation for each person-spell off AFDC following a spell on AFDC... There were 1,085 spells based on 742 women who ever exited welfare between January 1989 and December 1996. Of the 1,085 spells of non-receipt, 546 were completed (i.e., the woman returned to welfare); the remainder were right censored... The second file, the basis of the regression analysis, consists of a separate observation for each month a person is a head or wife and is in a spell of non-participation. This file, which consists of as many observations per person as months off AFDC following an exit, contains 22,254 persons-months of former recipient data between January 1989 and December 1996 and is based on 742 women (pp. 348-9).”

Repeatable events

For studies of repeatable events, supplement information on cases and person-time with data on the distribution of spells per respondent, and mention the statistical methods use to correct for multiple spells per respondent. For instance, Hofferth et al. (2005) include a dummy variable in their multivariate analysis to identify the 25 percent of spells that were second or higher periods off of welfare.

In analyses of potentially repeatable events, sometimes researchers choose to analyze one event per respondent, which should be explained in the methods section. For example, in their study of effects of changes to Canadian welfare benefits on marriage patterns of single mothers, Lefebvre and Merrigan (1998) discussed several substantive reasons why they include only one spell for each respondent, explicitly referring to historical events and theoretical considerations that pertain to their research question:

“Our subsample includes only respondents experiencing a first spell of single parenthood after 1974. The purpose of this restriction is to eliminate the remarriages of older women that took place during the years that followed the last war and which are more likely to be the result of traditional family behavior. The restriction also eliminates single mother spells resulting from the sudden increase in the number of divorces which took place after the enactment of the 1968 Act, the official ratification for couples who had been separated for a long time. Moreover, it is only at the end of the sixties that categorical welfare programs in Canada were replaced by the much more generous and universal Canadian Assistance Plan. Finally, our sample concentrates on the period where real welfare benefits were systematically increased in all provinces (pp. 748-9).”

Competing risks analyses

For a competing risks analysis, define the events that constitute the competing risks, and report number of cases, person-time of exposure, and the number of each type of event.

Increment-decrement analyses

For multistate or increment-decrement analysis, information on exposure is complicated not only by different populations at risk of each type of event, but by the fact that cases can move both into and out of the respective “at risk” pools. For example, a person is removed from the population at risk of marriage while they are currently married, but is added back into the population at risk of marriage after divorce or

widowhood. For an increment-decrement analysis, define the events to be studied, the number of each type of event, and the number of cases and person-time at risk.

Independent variables

In a hazards analysis, the element of time introduces another dimension not only to the dependent variable, but to independent (explanatory) variables as well. Some independent variables will be *fixed* (also known as *non-time-varying* or *time-invariant covariates*), maintaining the same value throughout the observation period for a given case. The most obvious examples are variables such as gender and race, which cannot change. Other variables may be specified as fixed because they are only measured once even though they could change over time. For example, Lefebvre and Merrigan (1998) include mother's age, level of educational attainment and number of children at the start of the spell in their analysis of marriage patterns among single mothers. Definitions and statistics for fixed covariates can be presented as for other multivariate analyses.

Some independent variables may be specified as *time-varying covariates* (also known as *time-dependent covariates*, not to be confused with *time-dependent effects*; see below) if more than one measurement is available for a given case. In the data and methods section, identify time-varying variables, indicate the dates of, or intervals between, the repeated measurements, and mention the source of data for those variables. In their analysis of new-firm survival, Audretsch and Mahmood (1995) described the frequency of data collection and time-varying attributes of the covariates used in their analysis as follows:

“Some of the variables, such as the innovation rates, are specific to the industry within which a particular establishment operates and do not vary over time... Other measures, such as wages, and the price-cost margin, vary both with respect to the industry and over time, but do not vary across particular establishments within any given industry. And finally, the two macroeconomic variables, the unemployment and interest rates, are constant across both establishments and industry, and vary only with respect to time (p. 100).”

For variables that could change over time but are specified as fixed, explain why and then return in the discussion section to consider implications for the interpretation of findings.

Model specification

To describe model specification, begin by indicating whether the model is specified as a continuous-time hazards model, such as a Cox proportional hazards model or some parametric models, or as a discrete-time model such as a logistic hazards specification. For parametric models such as Weibull, Gompertz, exponential, log-logistic, or log-normal, explain how you arrived at that particular specification for your research question and data, whether precedent in the literature, graphical evaluation of functional form (Allison 1995), or empirical tests of overall model fit. For instance, in DesJardins et al (2002) study of how financial aid affects chances of a first “stopout” (temporary or permanent leave from college), both the timing of financial aid packages and the effects of financial aid on stopout were allowed to change. They compared “frontloading” scholarships or grants in the first two years of college against funding that was spread across all terms. They used log-likelihood statistics to evaluate overall model fit under several assumptions:

“Initially, a time-constant coefficient (TCC) model was estimated. Next, we estimated time-varying coefficient (TVC) models with different time effects. Based on model fit, the TVC models were preferred compared to the TCC model (DesJardins et al. 2002; p. 669).”

Describe exploratory analyses used to arrive at the final model specification, including diagnostics to evaluate proportionality of hazards. Graphs of hazard rates by subgroup can be invaluable for illustrating situations where the assumption of proportionality is not met. For example, in a study of factors related to inpatient length of stay for youths with serious mental illness, Pottick and colleagues (1999) found that hazard rates of discharge for publicly and privately insured children were far from proportional (parallel) – they actually crossed one another quite steeply (Figure 2) and the change was

highly statistically significant. If they had not tested for time-varying effects, insurance appeared to be uncorrelated with length of stay – a highly implausible result.

Figure 2 about here

The authors summarized the bivariate pattern and its consequences for model specification: “Exploratory analyses revealed that the relationship between type of insurance and rate of discharge varied across time: For the first three weeks after admission, the rate of discharge for patients with private insurance was lower than for those with other means of payment, after which privately insured patients were discharged faster. Hence all models that include insurance type also include an interaction between time less than 21 days and lack of private insurance (Pottick et al. 1999; p. 220).”

Other aspects of the data and methods section for a hazards analysis are the same as in papers involving other types of multivariate analysis; see the checklist at the end of this article or Montgomery (2003) or Miller (2005) for guidance on the contents of a data and methods section.

The results section

Effective tables and charts for presenting results of hazards models should be self-contained, so readers can interpret the meaning of every number without consulting the text. In the title, convey the topics or questions addressed in that table or chart, naming each of the major components of the relationships illustrated. For tables or charts based on multivariate hazard results, mention the type of model, the event(s) modeled, and the concepts captured by the independent variables. In the row and column labels of a table, or axis labels and legend of a chart, show the identity and units or coding for every variable; replace acronyms or other abbreviated variable names with short phrases that readers can understand without referring to the text, or define them in notes to the table or chart.

Poor title: “Hazards results with confidence intervals.”

Comment: This title fails to convey the type of event being modeled, the sample under study, the independent variables, whether reported results are coefficients (log-relative hazards) or hazard ratios, or the width of the reported confidence intervals (95% CI are most common, but other widths are sometimes used).

Better title: “Relative risks of all-cause mortality among Harvard alumni, 1962/1966 through 1988, according to body mass index” (Table 1, from Lee et al. 1993).

Comment: This title identifies the event (all-cause mortality), the measure of effect size (relative risks), the key independent variable (body mass index; BMI), and the “W’s” (who, when, where) for the sample (Harvard alumni, 1962/66-1988). The type of inferential statistical information (95% confidence interval; CI) and the list of control variables are mentioned in the column headings and footnotes, respectively, while the definition and units for BMI are given in the row heading. Confidence intervals are a commonly-used means of presenting information about statistical significance in biomedical journals.

Table 1 about here

In tables of multivariate results, create separate columns for each model, labeling each model with a name (e.g., “Age-adjusted” or “Multivariate”; Table 1). Column headings can also be used to identify different competing risks or models for different subgroups (such as “Never Married” and “Separated or Divorced” in Table 3). See (author citation) for additional guidelines on creating tables and charts to present multivariate results.

Univariate and stratified statistics

Use tables of univariate and bivariate tables to report unadjusted levels and rates of event occurrence so that readers can assess their level and compare those values with data from other sources. Descriptive statistics on the independent variables can often be combined in a table with information on

number of spells, person-years at risk, and so forth. Table 2, for example, accompanied Hoynes' (2002) analysis of the association between labor market conditions and length of welfare spells.

Table 2 about here

Univariate and stratified survival, cumulative failure, and hazard statistics are often best presented graphically because the key point often is the general shape of those patterns rather than precise numeric values. For example, Figure 3 portrays cumulative failure curves from an analysis of employment patterns among older workers, stratified by age group (Chan and Stevens, 2001). Median survival time for each group is the time at which the proportion "failing" (in this case becoming re-employed) reaches 0.5.

Figure 3 about here

If the level and shape of the hazard curve across time is of interest, consider including a graph – either of the hazard curve for all groups combined, or stratified by subgroup as in Figure 2. The associated prose should complement the graph by including a few selected numeric values that document the pattern.

Multivariate model results

In most ways, tables of multivariate results from a hazards model are similar to tables of logistic regression results. To facilitate interpretation, hazard ratios that can be explained in terms of multiples of risk are preferred to coefficients (log-relative hazards). For categorical independent variables, present the hazard ratios in a table, then explain the size of the effect compared to the reference category, which is named in the sentence. For continuous independent variables, contrast other increments in the text if a one-unit increase is not typical or of interest. For instance, an increase in annual U.S. income of \$1 is unlikely to be of interest, so instead illustrate effects of a \$1,000 increase. Indicate statistical significance of the hazard ratio for each variable using *t*-statistics, standard errors (converted to the same units as hazard ratios, by exponentiating *s.e.* _{β}), or confidence intervals, along with *p*-values, or symbols or formatting to denote level of significance.

To present hazards estimates for several independent variables net of the effects of other variables in the model, create a table such as that adapted from Lefebvre and Merrigan (1998) which reports hazard ratios for a variety of demographic and policy variables, with separate models predicting first marriage among never-married mothers and among divorced or separated mothers (Table 3).

Table 3 about here

A similar approach can be used to present a series of nested models, competing risks, or models for separate groups, places, or time periods. Create one table with results from the series of models to be compared, and then summarize the extent to which the independent variables have similar effects in each of the specifications, on outcomes, or for subgroups studied. For example, in their analysis of mortality patterns, Lee and colleagues (1993) focused on how the hazard ratio for their key independent variable – body mass index (BMI) – changes between a model that controls only for age and one that also controls for the complete set of other variables listed in the footnote to the table (Table 1). Number of deaths, subjects, and person-years of exposure for each BMI subgroup are reported in the table with the multivariate results.

To maintain a focus on your research question when writing about bivariate or multivariate patterns, incorporate information about the concepts and units of the variables. A description of the direction (sign) and magnitude of each association is more effective than simply reporting statistical significance, stating that variables were correlated, or just reporting the coefficient or hazard ratio.

Poor description: "In the age-adjusted model, the relative risk (RR) for BMI 23.5-24.5 was 0.96 (95% confidence interval (CI) 0.86, 1.06). [Paragraph continues with separate sentences for the RR and CI for each of the other three BMI groups, followed by parallel reporting of the RR and CI for the multivariate model.]"

Comment: *This description essentially replicates the table contents without specifying the dependent variable or interpreting the shape of the relationship.*

Better description: "Taking into account age alone, we observed a J-shaped relation between BMI and all-cause mortality (Table 1) (*p* for linear trend = 0.02; *p* for quadratic trend = 0.003). The lowest mortality was in men with a BMI of 23.5 to less than 24.5. Men in the heaviest fifth of

BMI (26.0 or greater) experienced a significantly higher risk of dying during follow-up than men in the lightest fifth (less than 22.5) (RR, 1.12, 95% CI, 1.03 to 1.22) (Lee et al. 1993, p. 2825).”

Comment: The authors start by naming the model to be described (“age alone”), the independent variable (BMI) and dependent variable (all-cause mortality), and using a simile (“J-shaped”) to describe the general shape of the relationship. They then identify the lowest and highest risk groups, the reference category (“lightest fifth”), and report the size and statistical significance of mortality differences between those groups.

Lee et al. go on to explain:

“When we further adjusted for cigarette smoking habit and physical activity, we again found the lowest mortality among men with a BMI of 23.5 to 24.5. Increased mortality risk among the heaviest fifth now was accentuated somewhat (RR, 1.18, 95% CI, 1.08 to 1.28; Lee et al. 1993, p. 2825).”

Comment: Rather than describing the entire shape of the BMI/mortality relationship again for their full multivariate model, the authors name that model (“further adjusted for...”) and use the word “again” to get across similarity of results with those in the previous model. In the second sentence, the phrase “accentuated somewhat” conveys that the mortality difference across BMI groups was slightly larger than in the age-adjusted one.

Time-dependent effects

Time-dependent effects (or *non-proportional hazards*) occur when the relationship between an independent variable and the dependent variable changes over time rather than remaining constant. Put differently, a time-dependent effect is an interaction between an independent variable and time. It can occur even for a fixed (non-time-varying) covariate. For example, in the study by Pottick et al. (1999), although each youth’s insurance type (the independent variable) was constant throughout the study period, the effect of insurance type on risk of discharge changed with time since admission (Figure 2).

The most complicated situation occurs when an independent variable is time-varying and has a time-dependent effect, meaning that the value of the variable changes over time and that its effect on the dependent variable also changes with time since baseline. To interpret the coefficient on a time-varying covariate, describe the shape of its association with the independent variable and how it varies across time. For instance, when explaining the estimated effects of financial aid on chances of college “stopout”, DesJardins and colleagues (2002) wrote:

“In every year, scholarships had the largest impact on retention. Work/study had the next largest impact in the first two years of college... However in later years this effect wanes and earnings from campus employment have larger effects on retention. (p. 669).”

Summary

Writing about hazards analysis shares many aspects with writing about other types of multivariate models, including the basic contents of the data and methods sections, and use of tables, charts, and prose to present statistical results as evidence for testing a hypothesis about the specific research question. However, the focus on both occurrence and timing of events also necessitates information on measurement of the temporal aspects of the dependent and independent variables and sample size, as well as the type of hazards model specification. The discussion and conclusions should include a description of the advantages of hazards models for the specific research question and data, as well as limitations of the data such as single measures of covariates that could vary over time, or retrospective recall of dates or timing of event occurrence. By following the guidelines and examples in this paper, analysts can learn to improve the completeness and clarity of papers on applications of hazards models to topics in economics and related fields.

Checklist for Effective Presentation of Hazards Models

In the abstract:

- Under methods, name the type of hazards model used in your analysis, the number of cases, and the length of follow-up.
- Under results, summarize key findings about relative hazards.

In the data section:

- Set the context for your study by reporting the W's: who, when, where for your data set.
- Mention the data sources that provided information on event occurrence and timing.
- Specify the maximum possible length of follow-up, and dates or intervals of follow-up(s).
- Define your dependent variable:
 - Indicate what event(s) you are studying and whether they are
 - One-time or repeatable
 - One-way or increment-decrement
 - Competing risks, and if so, the specific types of competing events
- Explain what constitutes censoring in your data.
- Report the following aspects of sample size
 - Number of cases (respondents)
 - Person-time at risk, including units of time.
- For repeated events, report:
 - Number of spells
 - Percentage of cases with more than one spell
 - Corrections for multiple observations per case
- Define your independent variables, indicating
 - Which are fixed (time-invariant)
 - Which are time-varying, and how often were they measured.
 -

In the methods section:

- Name the type of statistical specification (e.g., Cox proportional hazards model).
- Explain how you arrived at your model specification, including diagnostics for proportionality and other exploratory data analysis.

In the results section:

- Use tables to present detailed statistical findings, charts to show general patterns.
 - Make each table or chart self-contained, naming the concepts, types of statistics, context (W's), identity and units or coding of every variable.
- Report direction, magnitude, and statistical significance of associations between key independent variables and the event under study.
 - Explain time-dependent effects, if any.

In the discussion and conclusions:

- Describe advantages of hazards models for your specific research question and data.
- Discuss limitations of your data for a hazards analysis, such as single measures of potentially time-varying covariates or sources of data on event occurrence and timing.

Biography: Jane Miller is a professor at the Institute for Health, Health Care Policy and Aging Research, and the E.J. Bloustein School of Planning and Public Policy, Rutgers University, where she teaches in the public health program. She is the author of *The Chicago Guide to Writing about Numbers* (Chicago, 2004) and *The Chicago Guide to Writing about Multivariate Analysis* (Chicago, 2005) and several articles on quantitative communication. Her research interests include poverty, child health, and access to health care.

Acknowledgements: This paper is adapted in part from material in *The Chicago Guide to Writing about Multivariate Analysis*, Chicago: The University of Chicago Press. © 2005 by The University of Chicago. All rights reserved. I thank Louise Russell and Yana van der Muelen Rodgers for comments on earlier drafts.

References

- Allison, P.D. 1995. *Survival Analysis Using the SAS System: A Practical Guide*. Cary, NC: SAS Institute.
- Audretsch, D.B., and T. Mahmood. 1995. New Firm Survival: New Results Using a Hazard Function. *The Review of Economics and Statistics* 77(1): 97-103.
- Chan, S., and A.H. Stevens. 2001. Job Loss and Employment Patterns of Older Workers. *Journal of Labor Economics*. 19(2):484-521.
- Cox, D.R., and D. Oakes. 1984. *Analysis of Survival Data*. Monographs on Statistics and Applied Probability. New York: Chapman and Hall.
- DesJardins, S.L., D.A. Ahlburg, and B.P. McCall. 2002. Simulating the Longitudinal Effects of Changes in Financial Aid on Student Departure from College. *The Journal of Human Resources* 37(3): 653-79.
- Greene, W.H. 2002. *Econometric Analysis, 5th edition*. Upper Saddle River, NJ: Prentice Hall.
- Hofferth, S.L., S. Stanhope, and K.M. Harris. 2005. Remaining Off Welfare in the 1990s: The Influence of Public Policy and Economic Conditions. *Social Science Research* 34: 426-53.
- Hoynes, H.W. 2000. Local Labor Markets and Welfare Spells: Do Demand Conditions Matter? *The Review of Economics and Statistics* 82(3): 351-68.
- Kalbfleisch, J.D. and R.L. Prentice. 1980. *The Statistical Analysis of Failure Time Data*. New York: John Wiley and Sons, Inc.
- Lefebvre, P. and P. Merrigan. 1998. The Impact of Welfare Benefits on Conjugal Status of Single Mothers in Canada: Estimates from a Hazard Model. *Journal of Human Resources* 33(3): 742-757.
- Lee, I-M, J.E. Manson, C.H. Hennekens, and R.S. Paffenbarger. 1993. Body Weight and Mortality: A 27-year Follow-up of Middle-aged Men. *Journal of the American Medical Association* 270(23):2823-8.
- Maciejewski, M.L., P. Diehr, M.A. Smith, and P. Hebert. 2002. Common Methodological Terms in Health Services Research and their Symptoms. *Medical Care* 40: 477-84.
- Miller, J.E. 2006. How to Communicate Statistical Findings: An Expository Writing Approach. *Chance* 19(4):43-49.
- Miller, J.E. 2005. *The Chicago Guide to Writing about Multivariate Analysis*. Chicago: University of Chicago Press.
- Montgomery, S.L. 2003. *The Chicago Guide to Communicating Science*. Chicago: University of Chicago Press.
- Pottick, K.J., S. Hansell, J.E. Miller and D. Davis. 1999. Factors Associated with Inpatient Length of Stay for Children and Adolescents with Serious Mental Illness. *Social Work Research* 23(4):213-24.
- UC Data. 1994. California Work Pays Demonstration Project: Statewide Longitudinal Database – Persons: 1% Sample, 1987-1992. Berkeley CA: University of California.
- Yamaguchi, K. 1991. *Event History Analysis*. Applied Social Research Methods Series, Volume 28. Newbury Park, California: Sage Publications.

Table 1. Relative risks of all-cause mortality among Harvard alumni, 1962/1966 through 1988, according to Body Mass Index* in 1962 or 1966.

Body Mass Index (kg/m ²)	# Deaths/ #Subjects	# Person- Years	Age-Adjusted Relative Risk† (95% CI)	Multivariate‡ Relative Risk§ (95% CI)
<22.5	988/4,380	96,471	1.00	1.00
22.5- <23.5	599/2,789	62,142	0.96 (0.86-1.06)	0.99 (0.89-1.20)
23.5-<24.5	853/4,112	91,632	0.92 (0.84-1.10)	0.95 (0.87-1.05)
24.5- <26.0	890/4,008	88,517	0.95 (0.87-1.05)	1.01 (0.91-1.10)
≥26.0	1,040/4,008	86,956	1.12 (1.03-1.23)	1.18 (1.08-1.28)

* Body Mass Index (BMI) = Weight in kilograms divided by the square of height in meters

† Linear trend: p = 0.02

‡ Adjusted for age (single years), cigarette habit (never, former, or current smoker), and physical activity.

§ Linear trend: p = 0.008.

Adapted from Lee et al. (1993), Table 2.

Table 2. Distribution of Length of AFDC and Non-AFDC Spells by Demographic Group, Unconditional Estimates, California, 1987-1992

	# spells	Cumulative probability that an AFDC spell is completed in:			
		≤ 6 months	≤ 1 year	≤ 2 years	≤ 4 years
All	12,177	0.28	0.46	0.62	0.75
Family type					
Single-parent (AFDC – family group)	10,313	0.27	0.45	0.62	0.75
Two-parent (AFDC – unemployed parents)	1,864	0.31	0.48	0.63	0.72
Race/ethnicity					
White	5,835	0.31	0.51	0.67	0.79
Hispanic	2,855	0.28	0.45	0.61	0.74
Black	2,639	0.23	0.41	0.57	0.70
Asian refugee groups	458	0.10	0.19	0.31	0.43
Other	390	0.24	0.42	0.62	0.78
Household head					
Non-teen	11,081	0.28	0.47	0.63	0.76
Teen	1,096	0.20	0.37	0.51	0.67
Residence					
Urban	10,606	0.27	0.45	0.60	0.74
Non-urban	1,571	0.33	0.54	0.71	0.82

Source: Longitudinal Database of Cases (LDB) compiled by UC Data and the University of California, Berkeley in association with the California Department of Social Services (UC Data, 1994)

Adapted from Hoynes (2000), Table 2.

Table 3. Cox Proportional Hazards Estimates of the Determinants of First Marriage Among Mothers, by Type of Single Marital Status, Canada, 1990

Variable	Never Married			Separated or Divorced		
	Coeff. ^a	Hazard Ratio	t-statistic	Coeff. ^a	Hazard Ratio	t-statistic
Age of mother in yrs ^b	-0.078	0.92	3.66***	-0.062	0.94	3.55**
Age of youngest child in yrs ^b	0.229	1.26	3.10***	0.138	1.15	5.03***
Number of children ^b	-0.069	0.93	0.53	-0.213	0.81	1.97**
Mother's Education ^c						
(Grade school)						
High School	-0.580	0.56	1.39	-0.516	0.60	1.69*
Post-secondary	-0.821	0.44	1.84*	-0.621	0.54	2.01**
University	0.341	1.41	0.59	-0.689	0.50	1.91*
Annual provincial welfare benefit in \$1986 (\$100s)						
For single person w/ 2 children	-0.398	0.67	1.60	-0.250	0.78	1.02
For couple w/ 2 children	0.378	1.46	1.67*	0.210	1.23	0.93
Total # cases		240			433	
- Log likelihood		-695.0			-1102.79	
Test for proportional hazards ^c		0.27			0.31	

^a Coefficient = log(relative hazard)

^b At start of spell

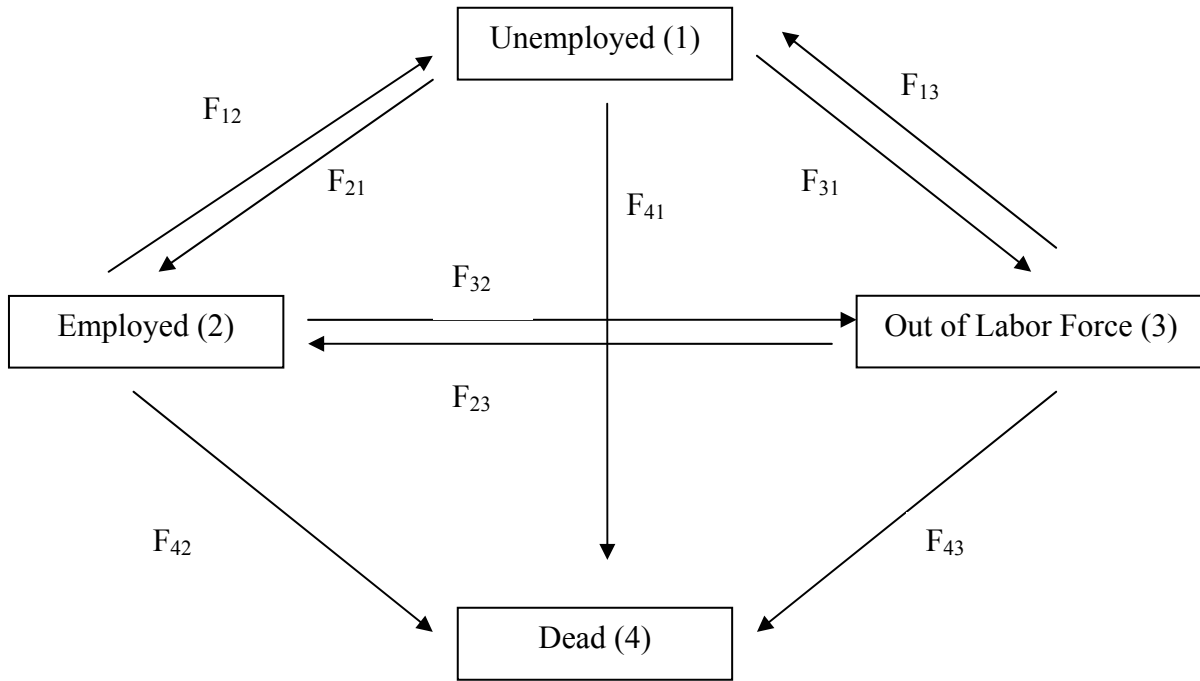
^c Reference category in parenthesis.

*** denotes p<.01; ** p<.05; * p<.10

Model also controls for province of residence.

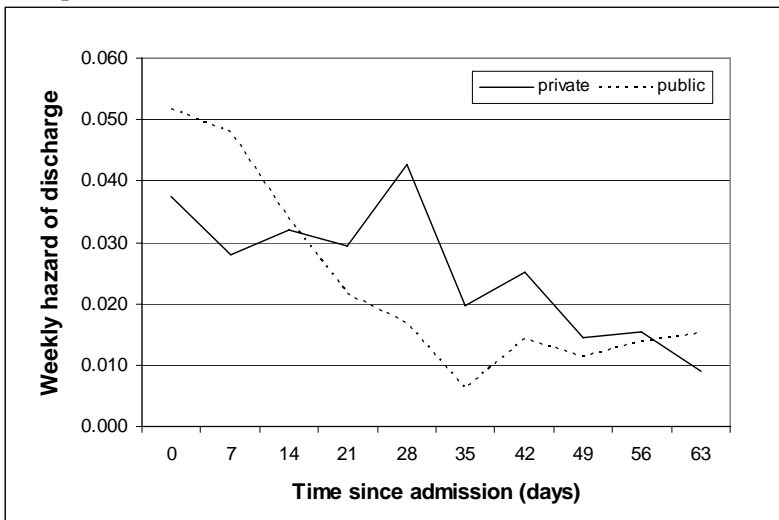
Adapted from Lefebvre and Merrigan (1998), Table 2.

Figure 1. Four-state increment-decrement life table model.




F's refer to functions for the transition between states. E.g., F_{21} refers to the transition to state 2 (employed; see number in parentheses within the "employed" box) from state 1 (unemployed). Thus F_{21} refers to finding a job - the transition from unemployed to employed.

Figure 2. Hazard of discharge from inpatient facility, by type of insurance and duration, 1986 U.S. C/PSS youth sample*



*C/PSS = Client/Patient Sample Survey; youths are <18 years of age at time of admission. Data from Pottick et al. (1999), personal communication.

Figure 3. Percentage of women who became re-employed following job displacement, by age group, Health and Retirement Survey, 1992-1996



Adapted from Chan and Stevens (2001), Figure 1b

Appendix. Terminology for hazards models

Terms and synonyms	Definition	Example topics, exhibits, and citations*	Comments
<u>General terminology for hazards models</u>			
Hazards analysis Duration models Survival analysis Event history analysis Cox regression Failure time analysis	Type of statistical analysis used to study patterns and correlates of event occurrence and timing.	Finding employment Exiting welfare	
<u>Life table modeling</u>			
Event Transition Failure Source of decrement	When a case moves from one state to another. Each type of transition is one kind of event.	“Stopout” from enrolled in college to disenrolled (DesJardins et al. 2002) Marriage (Lefebvre and Merrigan 1998)	Measured with a categorical variable indicating change from one discrete state to another. See “event indicator.”
Event indicator	Categorical variable used to specify the status of a case at the end of a spell. One component of the dependent variable in a hazards model.	<i>Binary</i> : 1 = event occurred; 0 = censored <i>Multichotomous</i> : 1 = employment; 2 = exiting labor force; 3 = death; 0 = censored	Can be binary (dichotomous) or multichotomous.

Spell Period at risk Duration Episode Exposure	Length of time since start of observation period or time at risk. Second component of the dependent variable in a hazards model.	Years from baseline to death Time to first marriage (Lefebvre and Merrigan 1998)	For repeatable events, each respondent can contribute more than one spell.
Population at risk Cohort	The group of cases that could potentially experience the event as of start of observation period.	<i>For employment:</i> Unemployed persons in the labor force. <i>For job loss:</i> Employed persons (Chan and Stevens 2001)	Population at risk depends on type of event (see examples)
Censoring	When the event is not observed for a given spell.	Loss-to-follow-up Dropping out of study End of observation period without event Mortality (for studies of events other than death)	Includes both those who have not yet and those who may never experience the event. Reasons for censoring depend on the type of event and data source(s).
Survival	When the event has not yet occurred for a given case.	<i>For mortality:</i> Remaining alive <i>For survival of a firm:</i> Remaining in business (Audretsch and Mahmood 1995).	For topics other than mortality, “survival” refers to not yet having experienced the event.
Survival curve	Chart showing the proportion of the cohort (population at risk) that has not yet experienced the event, plotted against time since baseline.	<i>For an analysis of re-employment:</i> Proportion of persons who remain unemployed (not shown)	Survival = 1 – cumulative failure. Starts at 1.0. Can only decline or remain level, cannot increase. Can be stratified to show patterns for subgroups.

Cumulative failure	The proportion of the cohort (population at risk) that has experienced the event, by time since baseline.	<p><i>For an analysis of employment:</i> Proportion employed since displacement from last job (Figure 3; from Chan and Stevens 2001)</p> <p><i>For an analysis of welfare spells:</i> Cumulative probability that an AFDC spell is completed (Table 2; from Hoynes 2000)</p>	<p>Cumulative failure = 1.0 – survival.</p> <p>Starts at 0. Can only increase or remain level, cannot decrease.</p> <p>Can be stratified to show patterns for subgroups.</p>
Hazard rate Transition rate	Rate of event occurrence within a time interval, conditional on survival to the beginning of that time interval.	Weekly rate of hospital discharge, by type of insurance (Figure 2; from Pottick et al. 1999)	<p>Number of events within a time interval divided by person-time at risk in that time interval.</p> <p>Can be stratified to show patterns for subgroups.</p>
Hazard ratio Relative hazard	<p>For continuous independent variables, the relative risk of event occurrence associated with a 1-unit increase in that independent variable.</p> <p>For categorical independent variables, the relative risk of event occurrence in a category of interest compared to the reference category.</p>	<p>Relative hazard of marriage, by age of mother in years (Table 3; from Lefebvre and Merrigan 1998)</p> <p>Relative risk of all-cause mortality for BMI ≥ 26.0 versus BMI < 22.5 (Table 1; from Lee et al. 1993)</p>	<p>Hazard ratio = $e^{\beta} = \exp(\beta)$, where β = the coefficient from a hazards model. Hence the coefficient is the log-relative-hazard.</p> <p>Confidence limits for a hazard ratio are calculated $\exp(\beta \pm [1.96 \times \text{standard error}])$.</p>
<u>Hazards models for different types of events</u>			
Single-decrement model	Type of hazards model used to analyze a single type of event.	<i>First</i> marriage out of single-mother spell (Lefebvre and Merrigan 1998)	Contrast against a competing risks model.

Repeatable event model	Type of hazards model used to analyze events that can occur more than once to each respondent.	Leaving AFDC (Hofferth et al. 2005)	Can be single- or multiple-decrement.
Competing risks model Multiple-decrement model	Type of hazards model used to analyze several alternative, mutually exclusive types of events.	Multiple causes of death, e.g., heart disease vs. cancer vs. all other causes	Contrast against a single-decrement model. Can be one-way or increment-decrement.
Increment-decrement model Multistate life table	Type of hazards model used to analyze reversible transitions or transitions among several different states or conditions.	<i>Reversible transition</i> : Poor to non-poor <i>Multistate</i> : Transitions among unemployment, employment, out of labor force, and death (Figure 1)	Can be single-decrement or competing risks model.
<u>Time dependence in hazards models†</u>			
Cox proportional hazards model Proportional hazards model	Type of hazards model where the hazard ratio for two subgroups is constant across time since baseline.	Relationship between body mass index and mortality (Lee et al. 1993).	A type of semi-parametric hazards model.
Non-proportional hazards Time-dependent effects	Type of hazards model where the hazard ratio for two subgroups changes with time since baseline. E.g., hazards curves for subgroups by converge, diverge, or cross one another.	Risk of hospital discharge for publicly versus privately insured children changes with time since admission, as evidenced by crossing (non-parallel) hazards curves (Figure 2, from Pottick et al. 1999).	Specified as an interaction between an independent variable and time since baseline. Can be specified with semi-parametric or parametric model. Can occur for fixed or time-varying covariates.

Fixed covariate Non-time-varying covariate Time-invariant covariate	An independent (explanatory) variable whose value remains the same throughout the spell.	Gender Number of children at start of unmarried spell (Table 3; from Lefebvre and Merrigan 1998)	Some variables that could theoretically change across time may be defined as fixed because of data constraints.
Time-varying covariate Time-dependent covariate	An independent variable whose value changes during the period at risk.	Wages, unemployment and interest rates in analysis of new-firm survival (Audretsch and Mahmood 1995)	Requires repeated observations or longitudinal surveillance data. Can, but does not necessarily, have a time-dependent effect on risk of event occurrence.

* Tables and figures cited refer to those in the body of the article. See list of references for full citations.

† See Allison (1995) or Yamaguchi (1991) for information on other functional forms for hazards models, including discrete-time and parametric hazards models.

See Cox and Oakes (1984), Allison (1995), or Yamaguchi (1991) for an in-depth discussion of concepts and terms for hazards analysis.