

ADJUSTING FOR ATTRITION IN EVENT HISTORY ANALYSIS

Daniel H. Hill¹
Survey Research Institute,
University of Toledo,
Toledo, Oh 43606

¹ The research reported in this paper was conducted under contract for the WESTAT corporation (subcontract # 5-YABC-2-66025). The author would like to thank Pat Doyle and Keith Rust of the ASA/Census Bureau SIPP Working Group for their helpful comments on an earlier version of this paper. The findings, opinions and conclusions are those of the author, as is the responsibility for any errors.

ABSTRACT

The sensitivity of parameter estimates of event history models to alternative methods of correcting for panel attrition is not well understood. This paper will investigate the issue of weighting for panel attrition in event history models by comparing alternative treatments of sampling weights in marriage and divorce models for members of the 1986 Survey of Income and Program Participation (SIPP). Three distinct weighting procedures will be compared. These procedures are based, respectively, on 1) the initial selection probability weights; 2) the 1986 SIPP Panel Weights; and 3) the monthly attrition-adjusted weights. Use of these later weights will require the development of maximum likelihood algorithms for discrete time event history models which can employ time-varying weights. Finally, the weighted estimates will be compared with the estimates of a structural model in which attrition is treated as an error-correlated competing alternative to marriage or divorce. Although it is impossible to identify a "best" procedure without accurate external data, significant differences in the estimates for the various procedures are indicative of significant attrition related problems in event history models. None of the weighting adjustments are found to have any appreciable effect on the parameter estimates of the divorce hazard model examined. The reason is that all of the weighting procedures are based on the assumption of independent censoring. The competing hazards structural model relaxes this assumption and finds evidence of significant correlated unmeasured heterogeneity. Once corrections for this are made, the net divorce hazards are seen to increase by more than one-half. This suggests that in many instances divorces in the SIPP end up being recorded as attrition.

ADJUSTING FOR ATTRITION IN EVENT HISTORY ANALYSIS

I. Introduction

One of the major strengths of panel data is that change measures are derived from current reports rather than from recollected statuses. This removes recall error as a source of bias affecting parameter estimates of event-history and other dynamic models. Unfortunately, another characteristic of panel data is that respondents tend to drop out or attrite as the panel ages. This introduces a different potential source of bias in event history models. Which of these biases is most important (and implicitly, whether panel data is superior to retrospective) is a complicated question -- the answer to which probably varies from one substantive application and set of surveys to another.

The seriousness of the biases introduced to event-history models by panel attrition, for instance, will depend on the amount of attrition and on the relationship between the propensity to leave the sample and the propensity to undergo the substantive change being analyzed. If these propensities are related, then the seriousness of the bias will depend on whether the relationship is confined to the covariates included in the event-history model specification or whether there are excluded, or even unmeasured, factors which affect both propensities. If the covariates do fully account for the relationship between the propensity to change and the propensity to attrite, then the parameter estimates of the substantive model will be unaffected by attrition. If, on the other hand, there is a residual relationship, then explicit corrections for attrition will be required to obtain unbiased estimates.

This paper investigates two alternative strategies for correcting for this type of nonignorable attrition in event-history models. The first strategy explored is the use of attrition-adjusted sampling weights in the event-history model. The purpose of these weights in panel surveys is to bring the aged panel back in line with the intended population of inference with respect to the distributions of key variables. The second strategy investigated is modeling attrition and the substantive change of interest as correlated competing hazards (i.e., competing hazards with correlated unmeasured heterogeneity). The estimated net survival function for the substantive change is interpretable as the one we would obtain if attrition were eliminated.

The context of these investigations is divorce (or separation) in the 1986 Panel of the Survey of Income and Program Participation (SIPP). Although the SIPP appears to remain roughly representative of marital statuses once adjustments for attrition are made (see e.g., Singh, 1988), there is concern that it seriously under represents change in marital status, particularly divorces, over the panel period (see e.g., M. Hill, 1987 and D. Hill, 1993). Since divorce is rare, both absolutely and relative to attrition, the parameter estimates from event-history models of it are especially vulnerable to attrition bias. Divorce is also important substantively. Not only is it of interest to behavioral scientist in its own right, but it is also an important determinant of family income, program participation and economic well-being. Furthering our understanding of these conditions in the population is the fundamental reason for conducting the survey.

The paper is organized in four sections. Section 2 provides background on the SIPP and defines some key concepts and procedures used in the remainder of the paper. In the next section (3), a simple event-history model of divorce is developed and the results obtained using four alternative weight schemes in its estimation are presented. Section 4 models and presents the results of the model-based alternative to weighting. The conclusions and recommendations for future research are presented in Section 5.

2. Background and Conventions

2.1 Background

The Survey of Income and Program Participation is a large panel survey of individuals in the United States which has been in operation since 1984 (see Jabine, *et al.* (1990) for a detailed description of the SIPP). The survey is comprised of a set of panels which are fresh cross-sections introduced annually. The members of each panel are interviewed every four months for roughly two and one-half years and retrospective information on income, employment, program participation and family composition is obtained for each month of the four month reference period.

Although SIPP study procedures call for following and interviewing all panel members when they leave the original sample households, this is not always possible. In the 1986 Panel just under one-fourth of the individuals originally in interviewed households became non-response at some point in the panel period and two-thirds of these individuals were never successfully recontacted (i.e., they attrited). Early on in the survey the concern was raised that attrition was particularly problematic when individuals experienced a marital status change. This is clearly the case in Figure 2.1 which presents the final attrition patterns for husbands and wives in the 1986 SIPP panel. The sample consists of all couples who were married at some point in the panel period. In the vast majority (6,333+130) of these couples neither spouse attrited. The 24-month divorce/separation rate among these panel members was just over two percent. This rate is in sharp contrast to the 60 percent ($=100*92/(92+64)$) divorce/separation rate among those couples in which one or the other (but not both) partners attrited. The more common situation in which both partners attrite is even more problematic when it comes to estimating marital status change. The reason is that we do not know how many of the 1004 couples who were still married at the time of their last interview, divorce or separated subsequent to (or concurrently with) their attriting.

Figure 1

The original concern over the effects of differential attrition was heightened when initial analysis of the first three waves of the 1984 Panel indicated that estimated divorce rates were only some 60% as high as outside data indicated they should be (see Figure 2.2). More recent analysis confirms this initial finding and shows that the problem persists in more recent panels (see i.e., D. Hill, 1993). Among those respondents who eventually attrited from the 1986 panel, the hazard of divorce prior to attrition is some sixty percent higher than among those individuals who responded throughout the panel period. Since it is quite possible that much of the divorce among attritors occurred concurrently with or just subsequent to their leaving the study (and therefore their divorces are unrecorded in the study), this estimated divorce hazard is most likely an underestimate.

Partly in light of this concern about differential attrition among marital status changers, a series of sampling weight adjustments was developed. Most of these adjustments take the form of post-stratification adjustments which have the effect of forcing survey estimates to correspond to outside information on the distributions of key variables including Census region, age, race, ethnicity, and other demographic factors (see Jabine *et al.*, 1990, pp 85-88 for a detailed discussion of these adjustments). An additional adjustment, the family composition adjustment, was also developed to adjust for the double selection probabilities for families formed when an original sample member marries someone who was not part of the original sample (but who had a positive *ex ante* chance of having been selected). This adjustment appears to have correct the apparent under estimate of marriages originally noted by Hill, 1987.

Figure 2

With these adjustments to the weights, the SIPP sample appears to be roughly representative of the entire population with respect to the proportions of adults married, single, divorced and widowed (see McMillen, 1989). This is not the same thing, however, as the survey being representative of marital status changes since it is always possible to have good point in time estimates of the numbers of individuals in particular states (i.e. stocks) without having good representation of the flows of individuals into those states. With rare events such as divorce, such stock-flow imbalances can persist well beyond the length of the panel period.

2.2 Conventions

Since the SIPP is a survey of individuals while the appropriate unit of analysis for divorce is the married couple, some conventions need to be adopted to deal with the situation which arises when one marital partner exhibits a different pattern of marital status change or attrition than the other. This would occur, for instance, if one spouse drops out of the study while the other remains in and reports being divorced.² Since divorce is of more substantive interest than attrition, we will apply a priority coding scheme to conflicting spousal reports. The priority is: 1) divorced; 2) separated; 3) married; and 4) nonresponse.³ This means that if only one spouse is responding and providing a marital status report, then the couple will be given that status. If both spouses are responding and one spouse reports being divorced while the other reports being married but living separately (or even married), then the couple will be considered divorced. Only if neither spouse is responding will the couple be considered as having attrited.⁴

The fact that two individuals are involved also raises the question of whose weight should be used in the weighted analysis. For those couples beginning the sample period married, the differences in the initial weights are relatively minor and, as a practical matter it does not matter, whose weight is used. When, however, one respondent leaves the study or when an individual from outside the original sample marries a sample member, then the weights can be quite diverse. The convention used in this paper is:

² A similar situation occurs when one respondent reports being widowed while the other appears to be nonresponse. In this case, however, we remove both spouses from the analysis. The reason is that our concern is not with real mortality but only panel mortality. It should be noted in this regard that there were no occurrences of one spouse claiming to be widowed while the other spouse continued to respond. This is most likely a reflection of data editing rather than of actual reporting behavior.

³ Those couples who remain married but live separately are treated, here, as married. In an earlier draft of this paper (May 29, 1994) these couples were improperly treated as separated.

⁴ An alternative treatment of attrition would be to define couples as attriting if either spouse attrites. Appendix A presents the results obtained under this laxer definition of attrition.

1. if both spouses are responding and are still married to each other then the average of their weights will be used;
2. if only one spouse is responding then the larger of the two weights will be used.

This weighting convention is an example of what Kalton and Brick term a "multiplicity approach" in which the selection probability of the non-original sample individual is assumed to equal that of their sample spouse. It is similar to the procedures used in the PSID in assigning family weights to families composed of a mix of sample and nonsample spouses.

2.3 Weighting Schemes Investigated

A variety of sampling weights are provided for the SIPP. The Longitudinal Research File, upon which I rely most heavily, contains three weights--the "Panel" Weight, the 1986 Calendar Year Weight and the 1987 Calendar Year Weight. Each of these weights is intended to allow inference to the entire population from that portion of the sample which responded to each interview in the reference period. The reference period for the "Panel" weight is the entire 28-month period covered by the 1986 Panel, while for the 1986 and 1987 Calendar Year weights, it is the 12 month period from January through December of the respective years. These weights are zero for individuals who were nonrespondents at any time during the respective reference period.

Since the divorce history model I will investigate in Section 3 deals with the hazard of divorce over the entire panel period, the first weighting scheme investigated will be that based on the Panel Weight from the 1986 Full Panel Longitudinal Research File.

In addition to the weights provided on the Longitudinal Research File, the Census Bureau provides a series of monthly weights for SIPP individuals which were merged from the individual wave files to the Longitudinal Research File. The existence of these monthly weights allows the application of two additional weighting schemes to the divorce history analysis. The first scheme is to use the initial month weights for the divorce model. The advantage of this is that the initial weights reflect all the differences in selection probabilities from the original sample design with a minimum of adjustments for sample attrition. The estimates from this weighting procedure are potentially useful as a bench-mark in evaluating the efficacy of the attrition adjustments in the other weighting schemes.

The final weighting scheme investigated uses the monthly weights as the basis of a time-varying weighting adjustment in the divorce history model. The rationale here is that since each monthly weight has been adjusted for attrition to bring the sample in line with the population during that month, use of these weights brings the total number of individuals at risk of divorce in each month in line with total number at risk in the population. Similarly, to the extent that the attrition adjustments are effective, the monthly estimates of numbers of couples experiencing divorce should also be representative of the corresponding number in the population. Since the

estimated divorce risk is the ratio of the number of divorces to the number of couples at risk, application of these weights might be thought to yield unbiased estimates.

2.4 Sample and Construction of Married Couple Records

The basic sample employed in this analysis consists of all spells in which two individuals in the 1986 SIPP Panel lived together as a married couple. The first stage of constructing the married couple pair data consisted of eliminating from the entire sample of 35,792 individuals those individuals for whom none of the 28 marital status variables equaled 1 (married spouse present). The resulting 15,608 individuals were then matched on the basis of sample unit, household and person number of spouse ID variables. In all, there were 7,821 marriage spells (there were 34 cases of multiple marriages during the 24 month panel period).⁵ Of the 7,821 marriage spells, 189 ended in widowhood and were eliminated from the analysis.⁶ Of the remaining 7,632 marriage spells, 231 ended in divorce or separation and 6,397 remained intact. Both spouses attrited in the residual 1004 cases.

3. The Effects of Weighting

3.1 A Simple Divorce Model

The divorce model I use is a very simple one which assumes that for each married couple (i) there is, at each time period (t), a latent underlying divorce propensity D_{it}^* . If this propensity is greater than some threshold (τ), the marriage will end in divorce at that time. Furthermore, the model assumes that the divorce propensity can be decomposed into systematic and random components according to:

$$D_{it}^* = \alpha + \beta' X_{it} + \epsilon_{it} \quad (1)$$

where the systematic portion is composed of: X_{it} --a vector of covariates; β --a vector of parameters relating these covariates to the divorce propensity; and α --a constant. ϵ_{it} is a random disturbance representing the net effects of all excluded factors on the divorce propensity. Following Allison, 1982, I assume the ϵ are distributed according to the hyperbolic secant-square distribution. In this case, the probability of the marriage surviving up to period t, and then ending

⁵ For rotation groups 2-4 there were actually 28 months in the panel reference period. The first four of these months, however, do not contain direct measures of marital status and were removed from consideration. Rotation group 1 was only interviewed six times and provides only 20 months of measured marital status.

⁶ An alternative and possibly superior treatment of these cases would be to consider them right censored.

in divorce during time period t , is:

$$\lambda_{it}^* > \tau | D_{T < t, i}^* < \tau = \exp(\alpha + \beta' X_{it})^{w_t} \prod_{T=1}^t [1 + \exp(\alpha + \beta' \lambda)] \quad (2)$$

where w_T is the appropriate sampling weight for period T .

The covariates I will use in this example consist of four time-invariant characteristics (age at the beginning of the panel period, whether both spouses are members of traditionally Catholic ethnic groups, whether the couple received government food stamps, and whether they owned (or were buying) their home at the beginning of the panel). I also include one time-varying covariate (whether the month is a "Seam" month). Despite this last variable, this is a constant hazard model--the propensity to divorce does not increase or decrease secularly although it does (at least apparently) experience periodic but temporary jumps at the seams. The reason for excluding time itself as a covariate is that the amount of time over which divorces can be observed in the 1986 SIPP is limited to 24 months (20 months for Rotation Group 1). This is far too short a time, relative to the life of most marriages, for the underlying hazard of divorce to change appreciably. Furthermore, by assuming this form of constant hazard model, both right- and left-censored observations are useful in estimation. This specification does have important implications for the meaning of the coefficient on age of the couple since it will capture both the effects of maturity of the partners and of the marriage itself on divorce propensities.

The use of sampling weights in equation (2) is a controversial issue. Their intended function is to remove potential biases in the β 's which might result from differential effective sampling rates.⁷ In the present case, there are two sources of such differentials--initial differences in sampling rates and differences in subsequent non-response rates.

Table 3.1 presents the parameter estimates obtained for the divorce event-history model using each of the three weighting schemes introduced in Section 2 along with the estimates obtained without weighting. Perhaps the most remarkable thing to note about these estimates is their similarity across the various weighting schemes. All four treatments of the weights yield strong (and highly significant) negative estimated age and home ownership effects and strong (again highly significant) positive estimated effects of foodstamp recipiency and whether the month in question is a "seam" month. Similarly, none of the weighting treatments yield estimated effects of both spouses being members of traditionally Catholic ethnic groups which are significantly, or even noticeably, greater than zero.

⁷ Those arguing against weighting note that differential selection probabilities can only introduce bias if the model is improperly specified. In this case, it doesn't really matter if the parameter estimates are biased since they are not the desired parameters in the first place. Furthermore, if the model is properly specified then the only effect of weighting is to reduce the efficiency of the estimates by introducing weight variance.

TABLE 3.1

Event-History Parameter Estimates for Divorce
for Four Weighting Schemes

	UNWEIGHTED	PANEL WEIGHT	INITIAL WEIGHT	MONTHLY WEIGHTS
DIVORCE CONSTANT	-4.39** (.26)	-4.58** (.22)	-4.44** (.25)	-4.47** (.24)
AGE OF COUPLE	-2.58** (.36)	-2.36** (.30)	-2.56** (.36)	-2.46** (.34)
ETHNO-CATHOLIC	-.16 (.21)	-.19 (.17)	-.13 (.21)	-.17 (.18)
FOODSTAMPS	1.30** (.19)	1.38** (.15)	1.34** (.18)	1.37** (.16)
HOME OWNER	-.43** (.14)	-.43** (.11)	-.41** (.14)	-.48** (.13)
SEAM	.46** (.14)	.49** (.12)	.50** (.14)	.43** (.13)
Ln(L) (Number of Cases)	- 1635.51 (7632) - 1735.63 5.8%	- 1764.52 (7632) - 1869.45 5.6%	- 1657.75 (7632) - 1759.16 5.8%	- 1709.83 (7632) - 1816.12 5.9%

*Significant at the .95 level.

**Significant at the .99 level.

The impression that the event-history model parameter estimates in Table 3.1 are insensitive to the treatment of attrition-adjusted sampling weights is born out when we examine the cumulative hazard functions they imply. These estimated hazards were obtained by means of sample enumeration simulations--i.e., the parameter estimates were applied to the actual values of the covariates for a random sample of cases and the estimates were averaged.⁸ Figure 3.1 plots these functions for each weighting scheme. Even after 24 months of allowing the differences in hazards to accumulate, there is hardly any visually discernable difference between the variously

⁸ See Train, 1986, for an explanation of the advantages of sample enumeration in interpreting results of discreet state models. See Hill, Axinn and Thornton, 1993, for an explanation of crude and net hazards and survival functions.

weighted and unweighted estimates.

FIGURE 3 (Supplied upon request)

4. Attrition and Divorce as Correlated Competing Hazards

The robustness of the parameter estimates of the divorce event-history model to the various weighting schemes would be encouraging were it not for the fact that, judging from Vital Statistics data, all of the estimates obtained imply divorce hazards which are far too low. Apparently weighting, at least of the sort used here, is not the solution. An alternative is to model attrition and divorce as potentially correlated competing hazards. This way, the effects of attrition on the estimated divorce hazards can be removed by examining the net hazard function for divorce. The major deficiency of the weighting approaches examined in Section 3 is that they all implicitly assumed independent censoring--i.e., that attritors behaved the same way after their last interview as before. It is quite likely, however, that divorce and attrition are both symptoms of what might be called marital distress--a shared unmeasured risk factor. Thus, individuals in distressed marriages are more likely both to divorce and to attrite than are people in happier marriages. The precise timing of the SIPP interview relative to the timing of a marital disruption is certainly unimportant to these people. It is, of course, crucial to the divorce analyst using SIPP data.

4.1 The SURF Model

Most competing hazards models also assume independent censoring and, as a result, are not likely to be any more successful in removing bias than the weighting approaches. An exception is the Shared Unmeasured Risk Factor (SURF) model of Hill, Axinn and Thornton (1993). As with the simple divorce model presented in Section 3, it is most useful to formulate this model in terms of the propensities to divorce (D_{it}^*) and to leave the sample via attrition (A_{it}^*). These propensities can be represented according to:

$$\begin{aligned} D_{it}^* &= \alpha_D + \beta_D' X_{Dit} + \epsilon_{Dit} \\ A_{it}^* &= \alpha_A + \beta_A' X_{Ait} + \epsilon_{Ait} \end{aligned} \tag{3}$$

where the first equation is the same divorce propensity equation as used in Section 3 and the second represents the corresponding attrition propensity. The covariate vectors (X_{Dt} and X_{At}) may or may not have common elements and there may or may not be constraints imposed across the coefficient vectors.⁹ The dynamic mechanism assumed is that couples remain in the base state (married and responding) until such time that either D_{it}^* or A_{it}^* exceeds some threshold τ . At this time, the couple moves to whichever competing state has the highest propensity score.

Unlike most competing hazards models, the SURF model assumes that the random

⁹ If constraints are imposed across the parameter vectors, or if the stochastic components of equation (3) are correlated, then it is necessary to analyze the competing hazards explicitly. Otherwise, it is sufficient to treat one alternative as a form of right censoring in analyzing the other alternative (see Petersen, 1991).

components of the competing propensities are related via:

$$F(\epsilon_{D_t}, \epsilon_{A_t}) = \exp(-[\exp(-\epsilon_{D_t}/\rho) + \exp(-\epsilon_{A_t}/\rho)]^\rho) \quad (4)$$

where ρ , known as the index of dissimilarity, is confined to the half-open interval (0,1]. This distribution is known as Gumbel's Type B bivariate extreme-value distribution. The correlation of the ϵ 's can be shown to be:

$$r_{\epsilon_{D_t}, \epsilon_{A_t}} = 1 - \rho^2 \quad (5)$$

In the special case where $\rho = 1$, the correlation is zero and the SURF model reduces to the ordinary discrete-time competing hazards model with independent censoring discussed by Allison (1982). This special case is sufficiently important that I will devote some time to it in Section 4.2, below. First, however, I need to complete the development of the general SURF model.

With equation (4) the hazard of divorcing, conditional on the individual 1) having remained married and responding up to time $t-1$; and 2) either divorcing or leaving the sample via attrition during period t , becomes:

$$Pr(D_t | D_t \cup A_t; T=t) = \frac{\exp(\beta'_D X_{D_t}/\rho)}{\exp(\beta'_A X_{A_t}/\rho) + \exp(\beta'_D X_{D_t}/\rho)} \quad (6)$$

where the symbols D_t and A_t represent the events $D^*_t > \tau$ and $A^*_t > \tau$, respectively.

The hazard of exiting at time t (given survival through period $t-1$) via either divorce or attrition is:

$$\begin{aligned} Pr(D_t \cup A_t | T \geq t) &= \frac{[\exp(\beta'_D X_{D_t}/\rho) + \exp(\beta'_A X_{A_t}/\rho)]^\rho}{1 + [\exp(\beta'_D X_{D_t}/\rho) + \exp(\beta'_A X_{A_t}/\rho)]^\rho} \\ &= \frac{\exp(\rho \ln(I_t))}{1 + \exp(\rho \ln(I_t))} = \frac{I_t^\rho}{1 + I_t^\rho} \end{aligned} \quad (7)$$

where $I_t^\rho = [\exp(\beta'_D X_{D_t}/\rho) + \exp(\beta'_A X_{A_t}/\rho)]^\rho$ can be defined as the "inclusive" exit propensity which is, as Amemiya (1985) puts it, "a kind of weighted average" of the individual exit propensities.

The likelihood function for a sample of n couples can be expressed as:

$$L = \prod_{i=1}^n \left[\prod_{j=D}^A \left[\frac{\exp(\beta_j' X_{jti} / \rho)}{I_{ti}} I_{ti}^{\rho} \right]^{\delta_{ji}} \prod_{k=1}^{t_i} \frac{1}{1 + I_k^{\rho}} \right] \quad (8)$$

where $\delta_{ji} = 1$ if, and only if, couple i ultimately exits the base state to state j ($= D$ or A). Otherwise, $\delta_{ji} = 0$. Maximum likelihood estimates of the parameters of the SURF model can be obtained by maximizing (8) with respect to the α 's, the β 's and with respect to ρ .¹⁰

4.2 Independence of Irrelevant Alternatives

As noted above, the special case in which ρ of equation (4) is 1, is sufficiently important as to deserve some discussion now. It is also sufficiently important as to have been given a special name in the discrete choice literature. It is called the Independence of Irrelevant Alternatives (IIA) condition or--since it is most often a maintained hypothesis--assumption. The reason for this name can be seen by examining the probability of leaving the base state to one alternative relative to the probability of remaining in the base state. This relative probability for divorce is:

$$\frac{P_{Dti}}{(1 - (P_{Dti} + P_{Ati}))} = \exp(\beta_D' X_{Dti}) I^{\rho-1} \quad (9)$$

If $\rho = 1$, then the level of the propensity to **attrite** rather than divorce is completely irrelevant. This propensity affects the ratio only via the inclusive value I ,¹¹ which drops out when $\rho = 1$. This implies that if all the individuals who left the sample via attrition were somehow interviewed, they would be found to be distributed across divorce and marriage in exactly the same proportion as the rest of the sample. Only if $\rho < 1$, will the attrition propensity (and the factors affecting it uniquely) affect the relative hazards of divorce.

4.3 SURF Model Estimates of Divorce and Attrition Hazards

¹⁰ Hill, Axinn and Thornton also present a two-step procedure to obtain consistent estimates of these parameters using the ordinary logit algorithm. Because the parameters estimated directly via maximizing equation (8) are more efficient, this method was used in obtaining the parameter estimates presented in this paper.

¹¹ If there are parameter constraints imposed across β_D and β_A then the factors affecting the attrition propensity will also affect the ratio. Most competing hazards specifications assume no such constraints (see Allison, 1982, Petersen, 1991, and Yamaguchi, 1991).

Table 4.1 presents the parameter estimates obtained when equation 8 is maximized with respect to the parameter α , β and ρ . The covariates included in the attrition portion of the model consist of the union of those factors found by Rizzo, Kalton and Brick, 1993, to be important correlates of attrition (age, region, relationship to reference person, number of imputations, home and asset ownership, foodstamp reciprocity and income)¹² and the covariates included in the marriage propensity specification. The first two columns of the table present the results obtained when the IIA assumption is made. It is no coincidence that the estimates for the marriage portion of this IIA model, presented in column 2, are virtually identical to those of the simple divorce hazards model presented in Table 3.1 (column 1) above. Indeed, it can be shown that competing hazards models which assume IIA are mathematically equivalent to single hazard models which treat exits to alternatives as an independent form of right-censoring (see Petersen, 1991). The only advantage of using the competing hazards framework when the independence assumption is made is that it is, at least theoretically,¹³ possible to impose and test cross-alternative restrictions on the coefficient vectors β_D and β_A .

¹² By and large these are powerful and significant predictors of attrition and the direction of their effects are as expected. The effects of age and survey month (time in the table) are, however, positive when many other studies of panel attrition find negative effects. These positive associations also show up in simple tabular analyses of attrition as well. The time effect, for instance, is apparent in that only 195 of the 995 couples who ultimately left the sample through attrition did so in the first 12 months. This is different from the pattern observed in other analyses and is probably a result of requiring both spouses to attrite before the couple is deemed nonresponse. As in all other analyses of attrition, the most important correlate is whether the month being examined is a seam month. The second most important predictor is whether the couple was married at the beginning of the study. If so, then they were far less likely to attrite. High income employed white home owners were also less likely to attrite as were couples receiving foodstamps. Those living in the Northeast and those with imputations in the Wave 1 data were significantly more likely to become nonresponse as the study progressed.

¹³ Many canned multinomial logit models do not allow cross-destination parameter constraints.

TABLE 4.1
Structural Model of Attrition and Divorce

	IIA		SURF	
	ATTRITION	DIVORCE	ATTRITION	DIVOR
CONSTANT	-4.90** (.20)	-4.39** (.26)	-4.70** (.18)	-4.65* (.20)
AGE OF COUPLE	.39** (.12)	-2.57** (.36)	.26* (.12)	-1.82* (.28)
RELATED	-.42** (.15)	-- --	-.39** (.14)	-- --
WHITE	-.27* (.10)	-- --	-.23* (.09)	-- --
NORTH EAST	.21* (.09)	-- --	.20* (.08)	-- --
HOME OWNER	-.17* (.08)	-.44** (.14)	-.21** (.07)	-.37* (.11)
# IMPUTATIONS	.65** (.10)	-- --	.60** (.09)	-- --
EMPLOYED	-.22** (.07)	-- --	-.23** (.06)	-- --
FOODSTAMPS	-.48* (.22)	1.29** (.19)	-.26 (.19)	1.00* (.18)
INCOME	-.46** (.18)	-- --	-.40* (.16)	-- --
TIME	.17** (.06)	-- --	.14* (.06)	-- --
SEAM	2.83** (.11)	.48** (.14)	2.70** (.11)	.97** (.16)
ENTERED MARRIED	-1.59** (.12)	-- --	-1.49** (.11)	-- --
ETHNO-CATHOLIC	.46** (.09)	-.16 (.21)	.43** (.08)	-.03 (.16)

LN(L), ρ	-6667.56	1.0 (--)	-6659.74	.60** (.07)
---------------	----------	-------------	----------	----------------

*Significant at the .95 level.

**Significant at the .99 level.

Columns 3 and 4 of Table 4.1 present the results obtained when the independence assumption is relaxed. The estimated index of dissimilarity (the bottom right entry of the table) of .74 is significantly less than 1.0 and implies a correlation between the random portions of the divorce and attrition propensities of roughly .45. Relaxing the IIA assumption also has the effect of reducing the estimated effects of age, foodstamp reciprocity, and home ownership. Evidently, while they remain important and significant, part of the apparent effect of these factors in the independent specification was due to their effects on attrition propensities. The most dramatic impact of allowing for non-independent censoring in the form of attrition is to increase the estimated effect of whether the month of transition was a "seam" month by roughly fifty percent--from .48 to .79. The meaning of this is that many of the exits recorded at the seam month were attributed to attrition under the independence assumption when they were actually due to divorce. This, of course, is an almost unavoidable consequence of the survey design in which interviews are attempted only periodically--unless at least one of the ex-spouses of a new divorce remains reachable until the next SIPP interview, the case will be recorded as attrition.

It is important to note that the inclusion of the seam month as a predictor in both the attrition and divorce portions of the model is crucial to the stability of the estimates. With this variable included, all specifications of the model examined yielded dissimilarity index estimates in the .7 to .75 range. When the seam month is excluded, on the other hand, the estimated ρ ranged as high as 1.5--a value for which the Gumbel distribution is undefined. The reason this predictor is so important to the estimation of the dissimilarity index is that more than any other factor, knowing its value allows us to distinguish between attrition and divorce. Interestingly, the point estimate for ρ remains at .74 even when the seam is included as the only covariate in the model. After years of struggling with the "seam" problem it is gratifying that, in this instance at least, it proves to be useful.

Unlike the results of the alternative weighting procedures, the effects of correcting for attrition using the SURF model are readily apparent in the estimated cumulative divorce hazards functions. Figure 4.1 presents the crude and net (of attrition) cumulative hazard function estimates implied by the coefficients presented in Table 4.1. The differences between the estimated crude divorce hazards for the independent censoring (IIA) and correlated censoring (SURF) models is hardly discernable. This is not surprising given that each can be considered a multivariate summary of the same set of observed survival experiences.

The net cumulative hazards functions, however, are quite distinct both from each other and from the crude hazards functions. The twenty-four month cumulative hazard of divorce net of attrition under the independence assumption is 4.6, which is some 12 percent higher than the crude. This is exactly what we would expect from the independence of irrelevant alternatives

assumption. The corresponding net divorce hazard rate obtained when the independent censoring assumption is relaxed is 5.6%, which is 37% percent higher than the crude rate. This brings SIPP divorce rate estimates almost up to those implied by Vital Statistics data.

5. Conclusions and Recommendations for Future Research

This paper has examined the effectiveness of two methods of adjusting for attrition in event history models. Because there was some evidence that attrition in the SIPP had a significant impact on the observed divorce rates, divorce was chosen as the substantive example. The first method of adjusting for attrition consisted of using attrition-adjusted sampling weights in the likelihood function of the event-history model. This method was found to have virtually no effect on the model estimates. The second method involved modeling attrition as a competing alternative means of exiting the base state (married and responding). When the stochastic portions of the propensity to attrite was allowed to be correlated with the corresponding random component of the propensity to divorce, the estimated cumulative hazards function was found to increase significantly for a 24-month rate of roughly 4% to one of over 5.5%. This increase in implied divorce rates brings the SIPP estimates almost in line with those from outside sources.

FIGURE 4 (Supplied upon request)

The results suggest that attrition and divorce are intimately related in that there are shared, or at least correlated, unmeasured risk factors affecting each. This results in a significant stochastic dependence between them which violates the underlying assumption of independent censoring upon which the weighting adjustments are based. Only when this dependency is explicitly recognized and corrected do the estimates change appreciably.

While the results of the Shared Unmeasured Risk Factors competing hazards model are encouraging as a means of correcting event-history model estimates for attrition, more work needs to be done. First, the method needs to be applied to a wide variety of substantive events. Exits from poverty spells and from spells of participation in means tested programs in the SIPP should be investigated. Also, however, the technique should be tested on data from other panel surveys such as the PSID.

Additionally, while none of the weighting schemes investigated in this paper had any discernable effect on the parameter estimates from the divorce event-history model, there are a wide variety of weighting schemes which were not analyzed. Future research should concentrate on those weighting schemes which would allow for non-independent censoring.

References

- Allison, P. (1982). Discrete-Time Methods for the Analysis of Event Histories. In S. Leinhardt (Ed.), Sociological Methodology 1982 (pp. 61-98). San Francisco: Jossey-Bass.
- Amemiya, T. (1985). Advanced Econometrics. Cambridge, MA: Harvard University Press.
- Bureau of the Census, 1990, "Survey of Income and Program Participation (SIPP) 1986 Full Panel Microdata Research File", (Washington: Bureau of the Census).
- Hill, D., 1993, "An Investigation of "Under" Estimates in the SIPP," (Toledo, Ohio: Survey Research Institute).
- Hill, D., W. Axinn, and A. Thornton, 1993, "Competing Hazards with Shared Unmeasured Risk Factors," Sociological Methodology, 1993, (August 1993).
- Hill, M., 1987, "Marital Status and Changes Therein in SIPP," Internal Census Bureau Memorandum to Census Staff. (February 20, 1987).
- Jabine, T., K. King, and R. Petroni, 1990, "Quality Profile," Survey of Income and Program Participation, (Washington: Bureau of the Census).
- Kalton, G. and Brick, J., 1994, "Weighting Schemes for Household Panel Surveys," (Rockville, MD: Westat, Inc.).
- McMillen, D., 1989, "The Measurement of Demographic Transitions in the Survey of Income and Program Participation," Proceedings of the Social Statistics Section of the American Statistical Association.
- Petersen, T. (1991). The Statistical Analysis of Event Histories. Sociological Methods and Research, 19, 270-323.
- Rizzo, L., G. Kalton, and M. Brick, 1993, "Weighting for Panel Nonresponse in SIPP," (Rockville, MD: Westat, Inc.).
- Singh, R., 1988, "SIPP Estimates of Number of Persons Marrying Per Year," Internal Census Bureau Memorandum to Gary M. Shapiro, (February 1, 1988).
- Train, K. (1986). Qualitative Choice Analysis. Cambridge: MIT Press.
- Yamaguchi, K. (1991). Event History Analysis. (Applied Research Methods Monograph No. 28, pp. 169-171). Newbury Park, CA: Sage.