IGNORABLE DROPOUT IN LONGITUDINAL STUDIES

Maria Eleanor Tipa, Susan A. Murphy, Diane McLaughlin, Penn State University
Maria Eleanor Tipa, Dept. of Statistics, PSU, 326 Thomas Bldg, University Park, PA 16802

Key Words: Censoring, Missing Data

INTRODUCTION

In a longitudinal study, measurements are taken on a unit or subject repeatedly at different time points. If measurements are taken n times, then the complete vector of measurements on the i^{th} unit is $\mathbf{Y}_i = (Y_{i1}, \ldots, Y_{in})$. It is not uncommon for longitudinal data to have missing values, that is, for some i, the whole Y_i vector will not be completely observed. A special case of this is when subjects leave the study prematurely, that is, if Y_{ij} is missing then all the succeeding Y_{ik}'s, $k > j$, are also missing. This special case is commonly called a *dropout* ((Diggle, 1989), (Diggle & Kenward, 1994), (Heyting et al, 1992), (Little,1995)).

Dropouts cause the data to be unbalanced and there are statistical procedures that can handle unbalanced data. However caution is required in using these techniques in the presence of dropouts. This is because the dropout may cause bias. For instance, in a clinical trial patients who recover may tend to drop out more than those who do not. Then the group with more cured subjects will also have more dropouts. Analysis of the observed data will then be biased against the group with more recovered patients.

However there are cases where valid inference still results even in the presence of dropout. Rubin (1976) calls such dropouts *ignorable*, otherwise they are *nonignorable*. Since different types of statistical methodology require different conditions for inference to be valid, the definition of *ignorability* depends on the statistical procedure used. Ignorability has been examined by many authors ((Rubin, 1976), (Laird, 1988), (Little & Rubin, 1987), (Diggle & Kenward, 1994), (Little, 1995)). Except for the work by Rubin (1976), the complex relationship between dropout, the population of inference and the statistical methodology is not emphasized.

The main focus of this paper is to give a concise definition for ignorable dropout. This definition depends on both the population of inference and the type of statistical inference used. In Section 2 the definition of dropout is related to the population of inference. This is illustrated using the data from the Panel Study of Income Dynamics (PSID).

In section 3 different types of dropout are described and compared to those found in the literature. These are interpreted using the PSID example. Section 4 describes how the definition of ignorability depends on the type of inference. As an example, ignorability conditions for likelihood-based inference are given.

DROPOUT AND THE POPULATION OF INFERENCE

We say that a subject leaving a study prematurely is a dropout if and only if following the departure the subject remains in the population of inference. Note the explicit reference to the population. We illustrate this definition of dropout using the Panel Study of Income Dynamics (PSID) (Hill, 1992),

Suppose that we are interested in explaining transitions into and out of poverty for the nonimmigrant, noninstitutionalized elderly, those aged 55 or over. That is, our population of inference is the set of all people, aged 55 or over, residing in the United States and not living in institutions. The PSID is an ideal data set for this purpose since the subset of the elderly in the PSID sample, when properly weighted, is representative of the above population. Further, sample families have been interviewed annually since 1968 to collect information on variables such as sources of income, employment information, work hours, geographic mobility and other demographic variables (User Guide to the PSID, 1984). This makes it possible for us to trace the economic history of the respondents since entry into the study.

The PSID loses some respondents for a variety of reasons, including: 1) change of address - unknown new address; 2) refusal to be interviewed; 3) institutionalization which can include going into service, moving into a nursing home, moving into a dormitory, etc., and; 4) death. Whether or not these losses should be considered dropouts depends on the the population given above. Certainly individuals who were not interviewed for reasons like residence change or refusal to be interviewed are dropouts. We know that these respondents can be in one of two states: in poverty or not in poverty but which state is not known due to failure to obtain information on them.

A more complex question is whether people lost from the study due to death are dropouts. In our population of inference, death is a natural occurrence. Therefore, those who die are not dropouts. Similarly, a person who goes into an institution ceases to be a member of our population of inference and thus, we do not lose any information once he enters an institution. In both cases a person leaving the study due to death or institutionalization to a nursing home is not a dropout since we want to make generalizations about the poverty experiences of the noninstitutionalized elderly in the U.S. Death and institutionalization are conditions or states a person can be in aside from being in poverty or not being in poverty. In survival analysis, these are called competing risks where for a given period of time, a person can be in one of different possible states. We see then that death or institutionalization signifies the end of a complete observation and that there is no loss of information on a respondent if either of these occurs.

Next consider a conceptual population of people, aged 55 and older, who reside in the U.S. and for whom institutionalization is not permitted/possible. Then if in our sample an individual becomes institutionalized, this is a dropout. This is because we do not have information on the person's poverty status which would have been available had they not been permitted or allowed to enter an institution.

On the other hand suppose our population of inference is the set of all people, aged 55 and older who reside in the U.S. Is poverty a relevant measure for people in institutions – that is, can a person be poor and in an institution at the same time? If an institutionalized person can be poor or not poor, then after a person in our sample becomes institutionalized, we no longer have complete information about his poverty status. He is a dropout. If however a person cannot be poor or not poor in an institution, institutionalization is then a competing risk. In this case, we do not lose any information if someone leaves the study due to institutionalization. Therefore, he is not a dropout.

As we have seen a sample may be used for inference about more than one population. And a dropout for one population of inference may not be a dropout for a different population of inference. It is important that we first determine the population of inference and then define dropout for this population.

TYPES OF DROPOUT

Suppose that the units in our study are observed in discrete time, $(t_1, \ldots, t_K)$. The ideal data for a subject is represented by $\mathbf{Y} = (Y_1, \ldots, Y_K)$ and $\mathbf{X} = (X_1, \ldots, X_K)$ corresponding to the K time points. We assume that for any k the distribution of $(Y_1, \ldots, Y_k)$ given $(X_1, \ldots, X_k)$ is independent of $(X_{k+1}, \ldots, X_K)$. In other words, X may be a function of time and baseline variables, which are completely observed at time t_1, and/or X is an external covariate. Conditioning on past outcomes, the ideal data likelihood function , $L(\theta)$ for one subject can be expressed as

$$L(\theta) = f_{Y_1}(Y_1|\mathbf{X};\theta) \prod_{k=2}^{K} f_{Y_k}(Y_k|Y_1, ..., Y_{k-1}, \mathbf{X}; \theta),$$

where $f_{Y_1}(y_1|\mathbf{X};\theta)$ is the marginal density of Y_1 given $\mathbf{X}$ and the conditional density of Y_k given the past outcomes $Y_1, ..., Y_{k-1}$ and $\mathbf{X}$ is $f_{Y_k}(y_k|Y_1, ..., Y_{k-1}, \mathbf{X}; \theta)$.

We now consider the case where dropout is possible. In order to focus on a typical subject's response and dropout, we assume that observations on the subjects are independent and identically distributed. We also assume that $\mathbf{X}$ is always observed. For one subject, the complete data can be represented by $\{(Y_1, \mathbf{X}), (Y_2, D_2), \ldots, (Y_K, D_K)\}$ where $D_k = I(D > k)$ and D is the time of dropout.

In reality, when $D_k = 0$, Y_k is not observed. Let $\tilde{Y}_k = Y_k$ if $D_k = 1$ and $\tilde{Y}_k = *$ otherwise. The incomplete data can then be represented as $\{(Y_1, \mathbf{X}), (\tilde{Y}_2, D_2), \ldots, (\tilde{Y}_K, D_K,)\}$. The incomplete data likelihood for one subject is

$$\begin{aligned} L(\psi) \quad = \quad & f_{Y_1}(\tilde{Y}_1|\mathbf{X};\theta) \prod_{k=2}^{K} \{f_{Y_k, D_k}(\tilde{Y}_k, 1|\text{past}_k; \psi)\}^{D_k} \\ & \prod_{k=2}^{K} \{f_{D_k}(0|\text{past}_k; \psi)\}^{1-D_k} \qquad (1) \end{aligned}$$

where $\text{past}_k = \{\mathbf{X}, (Y_j)_{j<k}, D \geq k\}$ and $\psi = (\theta, \phi)$ and ϕ contains parameters that describe the conditional distribution of $(D|Y, \mathbf{X})$. Note that in this representation, there may be an overlap between θ and ϕ.

We now describe three conditions on dropout that appear frequently in the literature. In the following section these conditions on dropout will be used to establish ignorability for different methods of inference. The conditions will be interpreted in the context of the population of nonimmigrant, noninstitutionalized people aged 55 or over (people in late midlife through old-age) who were residing in the U.S. in 1968. As we shall see, conditional independence is a statement about subgroups of this population. For the interpretations that will be given later on, we will assume, for simplicity, that the only reason for dropout is emigration. The response is whether a person is in poverty or not.

Moreover, unless otherwise indicated, the conditions that will be given are defined for all possible samples from the population.

One condition, which we will call *independent dropout I*, is that for a person still in the study up to time t_k the probability density of his potential response Y_k is the same as the assumed distribution in the case when dropout is not a possibility. Kalbfleisch and Prentice (1980,p. 120) and Andersen et al. (1988, p. 30) referred to an analogous condition, *independent censoring*, in the context of survival analysis for continuous lifetimes. This condition can be quantified as

$$Y_k \text{ is independent of } \{D \geq k\} \text{ given } (\mathbf{X}, Y_j, j < k)$$

for $k = 1, \ldots, K$. Or in terms of the densities,

$$f_{Y_k}(y_k|\text{past}_k; \psi) = f_{Y_k}(y_k|\mathbf{X}, Y_j, j < k; \theta) \qquad (2)$$

for $k = 1, \ldots, K$.

We now illustrate this condition using the population described earlier. Consider the group S_1 of people with a similar past, say college graduate males who have not been poor from age 55 up to age t_k, and let P_k be the group of males in S_1 who are poor at age t_k. Suppose that the proportion of males in poverty at age t_k in S_1 is 0.10. That is, $\frac{n(P_k)}{n(S_1)} = 0.10$. Now, divide S_1 into S_2 and $S_1 \setminus S_2$, where S_2 contains those who have not emigrated before t_k and $S_1 \setminus S_2$ contains those who have emigrated before t_k. If the independent dropout I assumption is to be satisfied, then the proportion of poor men in S_2 should also be 0.10, that is, $\frac{n(P_k \cap S_2)}{n(S_2)} = 0.10$. We illustrate this in Figure 1. Note that this also implies that the proportion of poor men at t_k in $(S_1 \setminus S_2)$ is also 0.10. In effect, we are assuming here that those who have not emigrated before t_k are as likely to be poor at age t_k as those who have emigrated before t_k.

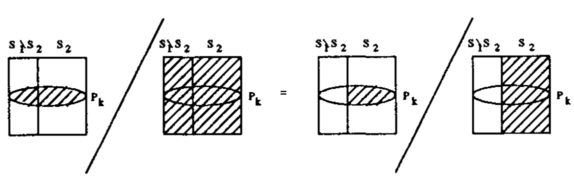

Figure 1: Independent Dropout I

Suppose that men unwilling to emigrate are more likely to become poor. In this case, relative to group S_1, group S_2 would be overly composed of men resistant to emigration and hence more likely to be poor. The independent dropout I assumption would not hold in this case.

A second condition is the *random dropout* condition discussed by Diggle and Kenward (1994, p. 53). They defined this as that the probability of a dropout at time t_k depends only on the previous observations and not on the observation at t_k. Andersen et al. (1988, p. 30) described an analogous condition that the censoring process should depend (in a functional sense) only on the past and not on future events. They referred to censoring processes that satisfy this condition as *predictable* censoring processes.

Mathematically, we can express the random dropout condition as

$$Y_k \text{ and } D_k \text{ are conditionally independent given past}_k$$

for $k = 1, \ldots, K$. In terms of densities:

1. $$f_{D_k|Y_k}(1|\text{past}_k, Y_k; \psi) = f_{D_k}(1|\text{past}_k; \psi) \qquad (3)$$

   or equivalently,

   $$f_{Y_k|D_k}(y_k|\text{past}_k, D_k = 1; \psi) = f_{Y_k}(y_k|\text{past}_k; \psi), \qquad (4)$$

   for $k = 1, \ldots, K$.

2. Given the past, the distribution of the potential response at time t_k is the same for someone who drops out at time t_k and someone who does not drop out at time t_k. This can be expressed as

   $$f_{Y_k|D_k}(y_k|\text{past}_k, D_k = 0; \psi) = f_{Y_k|D_k}(y_k|\text{past}_k, D_k = 1; \psi) \qquad (5)$$

   for $k = 1, \ldots, K$.

It can be shown that (3), (4) and (5) are equivalent to each other.

We now interpret Equation (5) using the elderly population. Define a subgroup in S_2, group E_k, composed of men who emigrate at t_k. Equation (5) requires that the proportion of poor men in S_2 among those who do not emigrate at t_k is the same as the the proportion of poor men in S_2 among those who emigrate at t_k. That is, $\frac{n(P_k \cap E_k' \cap S_2)}{n(E_k' \cap S_2)} = \frac{n(P_k \cap E_k \cap S_2)}{n(E_k \cap S_2)}$. Thus, among the men who have not emigrated prior to t_k those who emigrate at t_k are as likely to be poor at t_k as those who do not emigrate at t_k. We show this in the second and third sets of diagrams in Figure 2. By Equation (4), we have the equality with the first diagram.

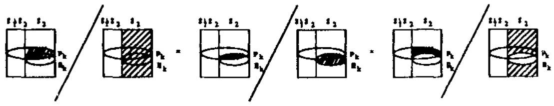

Figure 2: Random Dropout

Diggle and Kenward (1994) call the dropout *informative* if the random dropout condition is not satisfied, that is, if the probability of dropout at time t_k depends on the potential response Y_k. In estimating the rate of change over time in a continuous variable in a random effects model, Wu and Carroll (1988), Wu and Bailey (1989) and Schluchter (1992) call the censoring process *informative* if the censoring probability for each individual is related to the random effects for that individual. Little (1995) distinguishes between these two uses of the word informative by calling the first *nonignorable outcome-based dropout* in which the dropout probability depends on the missing values of the response variable and the second *nonignorable random-coefficient-based dropout* where the dropout probability depends on the random effects.

Another condition we will use to establish ignorability of the dropout is the *independent dropout II* condition,

$$Y_k \text{ is independent of } \{D > k\} \text{ given } (\mathbf{X}, Y_j, j < k)$$

for $k = 1, \ldots, K$. In terms of densities,

$$f_{Y_k|D_k}(y_k|\text{past}_k, D_k = 1; \psi) = f_{Y_k}(y_k|\mathbf{X}, Y_j, j < k; \theta) \quad (6)$$

for $k = 1, \ldots, K$. This is one possible quantification of Diggle and Kenward's (1994) assumption that "if an experimental unit is still in the study at time t_k its associated sequence of measurements $\tilde{Y}_j : j = 1, ..., k$ follows the same joint distribution as that of $Y_j : j = 1, ..., k$". Diggle and Kenward do not quantify this statement.

We are assuming here that given the past and that the subject has not dropped out at the present time, the distribution of his response at the present time is the same as in the case when there is no possibility of a dropout. In the population described earlier, we again consider the group, S_1, of college graduate males who have not been poor from age 55 up to age t_k and the following subgroups in S_1: S_2, composed of men who have not emigrated before t_k; and P_k, composed of men who are poor at t_k.

Suppose that the proportion of poor men at age t_k in S_1 is equal to 0.10. That is, $\frac{n(P_k)}{n(S_1)} = 0.10$. Now, we consider only those men who do not emigrate up to and including at age t_k ($E'_k \cap S_2$). Under Condition (6), the proportion of poor men in this set is also 0.10. That is, $\frac{n(P_k \cap E'_k \cap S_2)}{n(E'_k \cap S_2)} = 0.10$. This is shown in Figure 3.

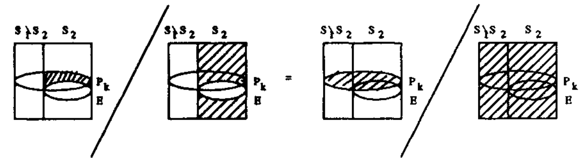

Figure 3: Independent Dropout II

From Figures 1, 2, 3 we see that any two conditions imply the third. It should be noted however that no one condition implies either of the two other conditions.

These conditions are similar to the conditions used in the literature for general missing data patterns. Rubin (1976) defined missing data as *missing at random* (MAR) if the probability of the observed pattern of missing data does not depend on the missing y values. If we apply this to longitudinal data where a dropout occurs at time t_d and we observe $\mathbf{X} = \mathbf{x}, (\tilde{Y}_1 = \tilde{y}_1, \ldots, \tilde{Y}_{d-1} = \tilde{y}_{d-1})$, we can interpret this condition as

$$P(D = d|\mathbf{X}, Y_1, \ldots, Y_K; \phi)|_{\mathbf{X}=\mathbf{x}, Y_1=\tilde{y}_1, \ldots, Y_K=y_K} =$$
$$P(D = d|\mathbf{X}, Y_1, \ldots, Y_{d-1}; \phi)|_{\mathbf{X}=\mathbf{x}, Y_1=\tilde{y}_1, \ldots, Y_{d-1}=\tilde{y}_{d-1}} \quad (7)$$

for the observed time of dropout, d, the observed past,$(\mathbf{x}, \tilde{y}_1, \ldots, \tilde{y}_{d-1})$ and for all unobserved $Y_d, \ldots, Y_K$. We will call this condition as *Bayesian MAR*. Some authors ((Laird, 1988), (Heyting et al., 1992)) use (7) and assume it to hold for all $d = 2, \ldots, K$, all $(\tilde{y}_1, \ldots, \tilde{y}_{d-1})$ and all $(Y_d, \ldots, Y_K)$. We will call this condition as *frequentist MAR*.

As is shown by Robins et al. (1995), frequentist MAR is equivalent to the following strengthening of random dropout:

$$D_k \text{ is independent of } (Y_k, \ldots, Y_K) \text{ given past}_k$$

or equivalently,

$$f_{D_k}(1|\text{past}_k; \psi) = f_{D_k|Y_k, \ldots, Y_K}(1|\text{past}_k, Y_k, \ldots, Y_K; \psi), \quad (8)$$

for $k = 2, \ldots, K$. However, Bayesian MAR is not equivalent to (8).

One may show that frequentist MAR (7 or 8) implies that independent dropout I (2) holds. However the independent dropout I condition does not imply frequentist MAR.

Frequentist MAR (7 or 8) also implies that random dropout (3) holds. However, if

$$D_k \perp Y_{k+1}, \ldots, Y_K \text{ given past}_k, Y_k. \quad (9)$$

for $k = 2, \ldots, K$, then (8) is equivalent to (3). It would seem that (9) would be satisfied in many practical cases. Thus, in situations in which (9) is plausible, frequentist MAR is not substantially stronger than random dropout. Since independent dropout I and random dropout imply independent dropout II, it follows that frequentist MAR also implies the independent dropout II condition (6).

IGNORABILITY OF THE DROPOUT

We call the dropout *ignorable* if and only if the mechanism of inference for the ideal data (but possibly unbalanced) is valid for the incomplete data subject to dropout. In large-sample frequentist inference, large sample properties, such as consistency and asymptotic normality of the estimator of the parameter are used for hypothesis testing and construction of confidence intervals. In particular, it is necessary that the estimating function for θ be unbiased in order for the estimator of θ to be consistent and asymptotically normal. We focus on this minimal property in defining ignorable dropout. We illustrate this using likelihood-based inference.

Valid inference on the parameter θ using likelihood methods depends on the correct specification of the likelihood function up to proportionality constants not depending on θ. Then, in general, the score function evaluated at the true value θ will have mean zero (will be an unbiased estimating function).

Suppose we observe, $\mathbf{X}$,D, and $\tilde{Y}_1, \ldots, \tilde{Y}_{D-1}$. Then if we pretend that D is a constant and that there are only $D-1$ observations, we will use the function

$$L^C(\theta) = f_{Y_1}(\tilde{Y}_1|\mathbf{X};\theta) \times \prod_{k=2}^{D-1} f_{Y_k|Y_1,\ldots,Y_{k-1}}(y_k|\mathbf{X}, \tilde{y}_1, \ldots, \tilde{y}_{k-1};\theta) \quad (10)$$

and

$$\frac{\partial}{\partial\theta}\log L^C(\theta) = \frac{\partial}{\partial\theta}\log f_{Y_1}(\tilde{Y}_1|\mathbf{X};\theta) + \sum_{k=1}^{K} D_k \frac{\partial}{\partial\theta}\log f_{Y_k|Y_1,\ldots,Y_{k-1}}(y_k|\mathbf{X}, \tilde{y}_1, \ldots, \tilde{y}_{k-1};\theta)$$

as one subject's contribution to the likelihood and score function respectively. In this case we are *ignoring the dropout.* Note that using this ignores the fact that the number of observed values of Y depends on the random variable D. The correct likelihood to use is the incomplete data likelihood $L(\psi)$ (Equation 1).

If we can factor $L(\psi)$ into $L^C(\theta)$ and some other factor not involving θ, then the correct score function for θ will be $\frac{\partial}{\partial\theta}L^C(\theta)$ summed over all subjects. In this case, we say that the dropout is *ignorable.* We now give a set of conditions for which the dropout is ignorable.

Theorem 1
The dropout is ignorable under likelihood-based inference if,

1. *the independent dropout II condition holds for $k = 1, \ldots, K$, and;*

2. $$f_{D_k}(1|past_k;\psi) \text{ is functionally independent of } \theta \quad (11)$$
   holds for $k = 1, \ldots, K$.

Under independent dropout II, $L^C(\theta)$ is a partial likelihood. Cox (1975) showed that under the usual regularity conditions, the score function from the partial likelihood has similar asymptotic properties as the score function from the full likelihood. This implies then that even if (11) is not satisfied, that is, even if the conditional probability of a dropout given the past depends on θ, we can still make valid large sample frequentist likelihood inference on θ based on $L^C(\theta)$ alone. However there will be a loss in efficiency if we use the partial likelihood instead of the full likelihood.

Note that for a particular observed sample the likelihood will be correctly specified, if Equation (6) and Equation (11) hold only for the particular observed data and all unobserved Y values. Thus, for direct-likelihood inference which Rubin (1976) defines as inference that "results solely from ratios of the likelihood function", assuming (6) and (11) for the observed data and all unobserved data is sufficient for ignorability. Rubin's conditions for ignorability under direct-likelihood inference are that the missing data are Bayesian MAR and that θ should be distinct from ϕ. However to consider the asymptotic properties of estimators , we assume these conditions for all possible samples from our population. Laird (1988) and Heyting et al. (1992) use the frequentist MAR condition (7/8) for ignorability and implicitly assume that θ and ϕ are distinct. We have shown in the previous section that frequentist MAR (7/8) implies both independent dropout I and independent dropout II.

Laird (1988) comments that the asymptotic variance of the estimator should be estimated by the observed information. Under the conditions of Theorem 1 or 2, the observed information matrix depends only on $L^C(\theta)$ and is thus a consistent estimator of the asymptotic variance. The expected information must be calculated relative to the density in (1). Therefore, the conditions of Theorem 1 or 2 are not sufficient for the consistency of the expected information.

DISCUSSION

We have shown that the definition of ignorable dropout depends on the population of inference and the statistical methodology used. Since one sample can be used to make inference about more than one population, it is important to carefully define the population when dealing with premature departures from a longitudinal study. Not all subjects who leave the study prematurely are dropouts – the population of inference dictates whether a premature departure is a dropout or an end of a complete observation.

Ignorability conditions depend on the type of inference. Rubin (1976) considers ignorability for Bayesian inference. Since a Bayesian analysis is conditional on the collected sample, Rubin's ignorability conditions concern only the collected sample. We have discussed likelihood-based inference. Since this method depends on large sample theory for the construction of confidence intervals and hypothesis tests, our ignorability conditions must hold for all possible samples collected from the population.

Tipa, Murphy & McLaughlin (1996) give ignorability conditions for different large-sample frequentist types of inference. They show that inference concerning the preservation of marginal moments have stronger ignorability conditions than likelihood-based inference. Ignorability conditions when interest is in the parameters of the conditional mean, on the other hand, are weaker.

Reference

Andersen, P., Borgan, O., Gill, R., & Kieding, N. (1988). Censoring, truncation and filtering in statistical models based on counting processes. *Contemporary Mathematics*, *80*, 19–60.

Cox, D. (1975). Partial likelihood. *Biometrika*, *62*(2), 269–76.

Diggle, P., & Kenward, M. G. (1994). Informative dropout in longitudinal analysis. *Applied Statistics*, *43*(1), 49–93.

Diggle, P. J. (1989). Testing for Random Dropouts in Repeated Measurement Data. *Biometrics*, *45*, 1255–1258.

Economic Behavior Program, Survey Research Center, Institute for Social Research, The University of Michigan, Ann Arbor, Michigan (1984). *User Guide to the Panel Study of Income Dynamics.*

Heyting, A., Tolboom, J., & Essers, J. (1992). Statistical Handling of Drop-outs in Longitudinal Clinical Trials. *Statistics in Medicine*, *11*, 2043–2061.

Hill, M. (1992). *The Panel Study of Income Dynamics: A User's Guide.* Sage Publications, Inc.

Kalbfleisch, J., & Prentice, R. (1980). *The Statistical Analysis of Failure Time Data.* John Wiley and Sons, Inc.

Laird, N. M. (1988). Missing Data in Longitudinal Studies. *Statistics in Medicine*, *7*, 305–315.

Little, R. (1995). Modeling the drop-out mechanism in repeated-measures studies. *Journal of the American Statistical Association*, *90*(431), 1112–1121.

Little, R., & Rubin, D. (1987). *Statistical Analysis with Missing Data.* John Wiley and Sons, Inc.

Robins, J., Rotnitzky, A., & Zhao, L. (1995). Analysis of semiparametric regression models for repeated outcomes in the presence of missing data. *Journal of the American Statistical Association*, *90*(429), 106–21.

Rubin, D. B. (1976). Inference and Missing Data. *Biometrika*, *63*(3), 581–92.

Schlucter, M. D. (1992). Methods for the Analysis of Informatively Censored Longitudinal Data. *Statistics in Medicine*, *11*, 1861–1870.

Tipa, M. E. V., Murphy, S. A., & McLaughlin, D. K. (1996). Ignorable Dropout in Longitudinal Studies. Tech. rep. 96-04, The Pennsylvania State University, University Park, PA.

Wu, M. C., & Carroll, R. J. (1988). Estimation and Comparison of Changes in the Presence of Informative Right Censoring by Modeling the Censoring Process. *Biometrics*, *44*, 175–188.

Wu, M., & Bailey, K. (1989). Estimation and Comparison of Changes in the Presence of Informative Right Censoring: Conditional Linear Model. *Biometrics*, *45*, 939–955.