Cause-Specific Analysis of Competing Risks Using the PHREG Procedure

Changbin Guo and Ying So, SAS Institute Inc.

ABSTRACT

Competing-risks analysis extends the capabilities of conventional survival analysis to deal with time-to-event data that have multiple causes of failure. Two regression modeling approaches can be used: one focuses on the cumulative incidence function (CIF) from a particular cause, and the other focuses on the cause-specific hazard function. These two quantities, unlike the hazard function and the survival function in conventional survival settings, are not connected through a simple one-to-one relationship. The Fine and Gray model extends the Cox model to analyze the cumulative incidence function but is often mistakenly assumed to be the only modeling technique available. The cause-specific approach that simultaneously models all the cause-specific hazard functions offers a more natural interpretation. SAS/STAT$^{®}$ 14.3 includes updates to the PHREG procedure to perform the cause-specific analysis of competing risks. This paper describes how cause-specific hazard regression works and compares it to the Fine and Gray method. Examples illustrate how to interpret the models appropriately and how to obtain predicted cumulative incidence functions.

INTRODUCTION

Competing risks arise in studies in which individuals are subject to a number of potential failure events and the occurrence of one event might impede the occurrence of other events. For example, after a bone marrow transplant, a patient might experience a relapse, or the patient might die while in remission. For relapse, death is a competing risk because the event of relapse can no longer occur after the patient dies. The survival function and the hazard function, two familiar concepts in conventional survival analysis, become elusive under competing risks. To meet the analysis needs, concepts that pertain to a specific cause of the event have been developed. The cause-specific hazard function generalizes the classical concept of the hazard function to the competing-risks setting, and it describes the rate of failure from one event type in the presence of others. The cumulative incidence function quantifies the risk of failure from a particular event type when there are competing risks.

Competing-risks analysis extends conventional survival analysis. For the latter, statistical methods such as the Kaplan-Meier estimate, the log-rank test, and the Cox regression are widely used in many applications. Although the log-rank test and the Cox regression can be adapted with minimal effort to make inferences about the cause-specific hazard function, they do not automatically lead to a correct analysis of the cumulative incidence function because of the presence of competing risks. To overcome this hurdle, specialized methods that target the cumulative incidence function have been developed (Gray 1988; Fine and Gray 1999).

The method of Fine and Gray (1999) extends the Cox regression to model the cumulative incidence function and marks a milestone in the development of modeling techniques for competing-risks data. Despite achieving great popularity, the Fine and Gray model has been criticized for its weakness of interpretation. The model covariates, although they can be interpreted as having an effect on the cumulative incidence function, do not directly link to an underlying event rate in the real world (Andersen and Keiding 2012). On the other hand, proportional hazards models for cause-specific hazards are easy to fit and offer a more natural interpretation in terms of rate ratios.

The LIFETEST and PHREG procedures in SAS/STAT software provide a comprehensive set of techniques for the analysis of competing-risks data. The LIFETEST procedure focuses on nonparametric analysis, and recent updates include specialized features for analyzing the cumulative incidence function. For example, you can compute the nonparametric estimate of the cumulative incidence function and use Gray's (1988) test to investigate group differences. On the other hand, the PHREG procedure provides two regression approaches for analyzing competing-risks data. You can apply Fine and Gray's method to directly model the cumulative incidence function; alternatively, you can fit Cox proportional hazards models to cause-specific hazard functions. So, Lin, and Johnston (2015) provide a tutorial on how to apply these techniques to study single causes of failure by using PROC PHREG.

Starting in SAS/STAT 14.3, you can use the EVENTCODE(COX)= option in the PHREG procedure to perform the cause-specific analysis of competing risks by fitting the cause-specific Cox models to different causes of failure

simultaneously. Moreover, you can make predictions about the cumulative incidence function based on the fitted models.

This paper first reviews the basic concepts of competing-risks analysis. It then discusses regression modeling strategies and uses a real-world data example of bone marrow transplantation to illustrate how to perform competing-risks regression by using Fine and Gray's method as well as the cause-specific Cox models. Connections and differences between these two methods are discussed. The last section demonstrates how to use the two approaches to predict the cumulative incidence function.

BASIC QUANTITIES IN COMPETING RISKS

Let T and C denote the failure time and censoring time, respectively. For data that have K competing risks, the pair (X, δ) is observed, where $X = \min(T, C)$ and $\delta = 1, \ldots, K$ is an indicator that has values of 0 for censoring and other values that designate specific failure causes. For analyzing competing-risks data, two useful quantities are the cause-specific hazard function and the cumulative incidence function:

- The cause-specific hazard function $h_k(t)$ at time t is the instantaneous rate of failure due to cause k conditional on survival until time t or later. It is defined as

$$h_k(t) = \lim_{\Delta t -> 0} \frac{P(t < T < T + \Delta t, \delta = k | T > t)}{\Delta t}, \quad k = 1, \ldots, K$$

- The cumulative incidence function, denoted by $F_k(t)$, is the probability of failure due to cause k prior to time t. It is defined as

$$F_k(t) = P(T \leq t, \delta = k), \quad k = 1, \ldots, K$$

  The cumulative incidence function is also referred to as the subdistribution function, because it is not a true probability distribution.

Mathematically, the kth cumulative incidence function can be expressed in terms of *all* the cause-specific hazard functions via the integral as follows,

$$F_k(t) = \int_0^t S(u)h_k(u)du = \int_0^t S(u)dH_k(u), \quad k = 1, \ldots, K$$

where $H_k(t) = \int_0^t h_k(u)du$ is the cause-specific cumulative hazard function and $S(t) = \exp\left(-\sum_{k=1}^K H_k(t)\right)$ is the overall survival function, which is the probability of surviving beyond time t. It is clear that $F_k(t)$ involves not only $h_k(t)$ but also all the competing cause-specific hazard functions when $K > 1$. When $K = 1$, the subdistribution function degenerates to $F_1(t) = 1 - \exp(-H_1(t))$ and becomes a function of only $h_1(t)$.

The one-to-one relationship does not exist between the cause-specific hazard function and the corresponding cumulative incidence function when competing risks are present. This feature, as pointed out by Andersen et al. (2012), is the key to understanding the finesse required in competing-risks analysis and to choosing appropriate techniques.

AN EXAMPLE OF COMPETING-RISKS DATA

Bone marrow transplant is a standard treatment for acute leukemia. Klein and Moeschberger (2003) present a set of bone marrow transplant data for 137 patients, divided into three disease groups based on their diagnosis at the time of transplantation: acute lymphoblastic leukemia (ALL), acute myelocytic leukemia (AML) low-risk, and AML high-risk. Among the 137 patients in the study, 38 patients were diagnosed as having ALL, 54 patients were diagnosed as having AML low-risk, and 45 patients were diagnosed as having AML high-risk. There are a number of concomitant variables in the data set; for simplicity, only the waiting time for transplant is included here.

During the follow-up period, some patients might experience a relapse of the leukemia, and some patients might die while in remission. The comparison of the disease groups focuses on the occurrence of relapse.

The following statements provide the data. The variable **Group** designates the disease group of a patient, which is either ALL, AML low-risk, or AML high-risk. The variable **T** is the disease-free survival time in days, which is either the time to censoring, the time to relapse, or the time to death while in remission, whichever occurs first. The indicator

variable **Status** has three values: 0 for censored observations, 1 for patients who relapse, and 2 for patients who die before experiencing a relapse. The concomitant variable **WaitTime** is the waiting time for transplant, in days. Because this variable has a very large variation, a log transform is applied to stabilize the variance.

```
proc format;
    value DiseaseGroup 1='ALL'
                       2='AML-Low Risk'
                       3='AML-High Risk';

data bmt;
    input Group T Status WaitTime @@;
    logWaittime=log(WaitTime);
    format Group DiseaseGroup.;
    datalines;
1  2081  0   98  1  1602  0  1720  1  1496  0   127  1  1462  0   168
1  1433  0   93  1  1377  0  2187  1  1330  0  1006  1   996  0  1319
1   226  0  208  1  1199  0   174  1  1111  0   236  1   530  0   151
1  1182  0  203  1  1167  0   191  1   418  2   110  1   383  1   824
1   276  2  146  1   104  1    85  1   609  1   187  1   172  2   129
1   487  2  128  1   662  1    84  1   194  2   329  1   230  1   147
1   526  2  943  1   122  2  2616  1   129  1   937  1    74  1   303
1   122  1  170  1    86  2   239  1   466  2   508  1   192  1    74
1   109  1  393  1    55  1   331  1     1  2   196  1   107  2   178
1   110  1  361  1   332  2   834  2  2569  0   270  2  2506  0    60
2  2409  0  120  2  2218  0    60  2  1857  0    90  2  1829  0   210
2  1562  0   90  2  1470  0   240  2  1363  0    90  2  1030  0   210
2   860  0  180  2  1258  0   180  2  2246  0   105  2  1870  0   225
2  1799  0  120  2  1709  0    90  2  1674  0    60  2  1568  0    90
2  1527  0  450  2  1324  0    75  2   957  0    90  2   932  0    60
2   847  0   75  2   848  0   180  2  1850  0   180  2  1843  0   270
2  1535  0  180  2  1447  0   150  2  1384  0   120  2   414  2   120
2  2204  2   60  2  1063  2   270  2   481  2    90  2   105  2   120
2   641  2   90  2   390  2   120  2   288  2    90  2   421  1    90
2    79  2   90  2   748  1    60  2   486  1   120  2    48  2   150
2   272  1  120  2  1074  2   150  2   381  1   120  2    10  2   240
2    53  2  180  2    80  2   150  2    35  2   150  2   248  1    30
2   704  2  105  2   211  1    90  2   219  1   120  2   606  1   210
3  2640  0  750  3  2430  0    24  3  2252  0   120  3  2140  0   210
3  2133  0  240  3  1238  0   240  3  1631  0   690  3  2024  0   105
3  1345  0  120  3  1136  0   900  3   845  0   210  3   422  1   210
3   162  2  300  3    84  1   105  3   100  1   210  3     2  2    75
3    47  1   90  3   242  1   180  3   456  1   630  3   268  1   180
3   318  2  300  3    32  1    90  3   467  1   120  3    47  1   135
3   390  1  210  3   183  2   120  3   105  2   150  3   115  1   270
3   164  2  285  3    93  1   240  3   120  1   510  3    80  2   780
3   677  2  150  3    64  1   180  3   168  2   150  3    74  2   750
3    16  2  180  3   157  1   180  3   625  1   150  3    48  1   210
3   273  1  240  3    63  2   360  3    76  1   330  3   113  1   240
3   363  2  180
;
```

To study the occurrence of relapse, you can estimate the cumulative incidence function of each disease group. Furthermore, you can use Gray's (1988) test to compare the disease groups. However, if you also want to adjust for some concomitant variables, such as the effect of the waiting time for transplant, you need to perform a regression analysis. The next two sections show how to use the PHREG procedure to fit the standard Cox model, the cause-specific Cox models, and the Fine and Gray model to the competing-risks data.

STANDARD COX REGRESSION

Suppose a random sample of n subjects is collected. For the ith subject, $i = 1, \ldots, n$, let X_i, δ_i, and $\mathbf{Z}_i$ be the observed time, cause of failure, and covariate vector at time t, respectively. Assume that K causes of failure are observable ($\delta_i \in \{1, \ldots, K\}$); $\delta = 0$ indicates a censored observation.

For competing-risks data, it is always possible to fit a standard Cox regression model to the overall hazard function $h(t) = \sum_{k=1}^{K} h_k(t)$ by disregarding the different event types. This is sometimes referred to as a composite endpoint analysis (Wolbers et al. 2014), and it is still widely used in many applications, especially for studying event-free survival in clinical research.

The Cox proportional hazards model assumes

$$h(t|\mathbf{Z}) = h_0(t)\exp(\boldsymbol{\beta}'\mathbf{Z})$$

where $h_0(t)$ is the baseline hazard function and the vector $\boldsymbol{\beta}$ represents the covariate effects.

The regression coefficients $\boldsymbol{\beta}$ are obtained by maximizing the partial likelihood $L(\boldsymbol{\beta})$,

$$L(\boldsymbol{\beta}) = \prod_i \left(\frac{\exp(\boldsymbol{\beta}'\mathbf{Z}_i)}{\sum_{j \in \mathcal{R}_i} \exp(\boldsymbol{\beta}'\mathbf{Z}_j)} \right)^{I(\delta_i > 0)}$$

where $\mathcal{R}_i$ denotes the risk set at X_i and contains individuals who do not fail from any causes or are not censored before X_i.

The following statements fit the proportional hazards model for event-free survival by treating relapse (**Status=1**) or death (**Status=2**) as a failure event:

```
proc phreg data=Bmt;
   title 'Composite Endpoint Analysis';
   class Group (order=internal ref=first);
   model T*Status(0)=Group logWaitTime;
run;
```

Figure 1 displays the estimated effects of the risk factor **Group** on event-free survival.

Figure 1 Parameter Estimates for Event-Free Survival

Composite Endpoint Analysis

The PHREG Procedure

Parameter		DF	Parameter Estimate	Standard Error	Chi-Square	Pr > ChiSq	Hazard Ratio	Label
Group	AML-Low Risk	1	-0.62604	0.30802	4.1309	0.0421	0.535	Group AML-Low Risk
Group	AML-High Risk	1	0.36482	0.27002	1.8255	0.1767	1.440	Group AML-High Risk
logWaittime		1	-0.06243	0.13780	0.2052	0.6505	0.939	

Assuming that the Cox model is correct, you can interpret the hazard ratio as the relative risk of an event (relapse or death). The results show that the risk of the event occurring among the AML high-risk patients relative to the risk of the event occurring among the ALL patients is 0.5, whereas the risk of the event occurring among the AML low-risk patients relative to the risk of the event occurring among the ALL patients is 1.4.

Performing the composite endpoint analysis is straightforward for competing-risks data, but it has obvious limitations. When you combine different event types, the information from the individual causes is lost. For the bone marrow transplant data, the composite endpoint contains all the events due to relapse or death. The composite endpoint analysis cannot show the effects of disease group on relapse in the real world, where death can also occur.

COMPETING-RISKS REGRESSION

To study the effect of covariates on the cause-specific outcome, you can model the cause-specific hazard function or you can model the subdistribution function. The cause-specific approach has a long history of application, and proportional hazards models are typically assumed (Prentice et al. 1978). In contrast, the model of Fine and Gray (1999) targets the subdistribution function directly and has become a very popular choice in practice. The Fine and Gray model is often seen as the default choice to model competing-risks data, but in fact the two approaches focus on different aspects of analysis and complement each other. Latouche et al. (2013) argue that a proper competing-risks analysis should report results on all the cause-specific hazards and cumulative incidence functions.

Modeling the Cause-Specific Hazard

The Cox proportional hazards model can be conveniently adapted to analyze the cause-specific hazard functions from different event types. For cause k, a separate proportional hazards model can be assumed,

$$h_k(t|\mathbf{Z}) = h_{k0}(t) \exp(\boldsymbol{\beta}_k' \mathbf{Z})$$

where $h_{k0}(t)$ is the baseline of the cause-specific hazard function and the vector $\boldsymbol{\beta}_k$ represents the covariate effects on the event of interest.

Because the K models are mutually exclusive, they can be analyzed independently. For cause k, the regression coefficients $\boldsymbol{\beta}_k$ can be estimated by maximizing the modified partial likelihood,

$$L(\boldsymbol{\beta}_k) = \prod_i \left(\frac{\exp(\boldsymbol{\beta}_k' \mathbf{Z}_i)}{\sum_{j \in \mathcal{R}_i} \exp(\boldsymbol{\beta}_k' \mathbf{Z}_j)} \right)^{\delta_i = k}$$

where $\mathcal{R}_i$ denotes the risk set at X_i and contains individuals who do not fail from any causes or are not censored before X_i.

The specified cause-specific hazard model has a simple structure and is easy to understand and interpret. But it has the following limitations:

1. The regression coefficients are assumed to be different between the causes. This does not have to be the case in general, and in certain applications it might be more natural to set up the model so that certain covariates have the same effect across different causes. See Andersen et al. (1995) for an in-depth discussion of these more parsimonious models, and see Gardiner (2016) for an exposition of how to fit them using PROC PHREG by carefully preparing the necessary data set.

2. The model covariates are assumed to be time-invariant. This does not need to be true in more general settings. However, the difficulty lies in how to interpret such models in a simple fashion as well as how to obtain model-based predictions. See Cortese and Andersen (2010) for a comprehensive discussion of how to incorporate time-dependent covariates into cause-specific regression models.

For a particular cause of interest, you can fit cause-specific Cox models as regular Cox models by using PROC PHREG. A well-known strategy is to treat all the competing events as censored observations in the analysis. Doing so ensures that the risk set at each event time contains only those subjects who have not failed from competing events or are truly censored. This leads to the correct calculation of the partial likelihood for the cause-specific Cox model.

As an illustration, the following statements fit the cause-specific Cox model to relapse (**Status=1**) by treating death (**Status=2**) as censored:

```
proc phreg data=Bmt;
   title 'Cause-Specific Hazard Regression for Relapse';
   class Group (order=internal ref=first);
   model T*Status(0,2)=Group logWaitTime;
run;
```

Note that the designated codes for censoring include the value 0, which indicates true censoring, and the value 2, which represents death.

Figure 2 displays the estimated effects of the risk factor **Group** on relapse.

Figure 2 Parameter Estimates for Relapse (**Status**=1)

Cause-Specific Hazard Regression for Relapse

The PHREG Procedure

							Hazard	
			Parameter	Standard				
Parameter		DF	Estimate	Error	Chi-Square	Pr > ChiSq	Ratio	Label
Group	AML-Low Risk	1	-1.07299	0.46245	5.3836	0.0203	0.342	Group AML-Low Risk
Group	AML-High Risk	1	0.55111	0.36464	2.2842	0.1307	1.735	Group AML-High Risk
logWaittime		1	-0.23061	0.19440	1.4071	0.2355	0.794	

In large part, the cause-specific hazard ratio resembles the hazard ratio in a standard Cox regression, and it quantifies the strength of association between the cause-specific hazard rate and the covariates. A value greater than 1 implies an elevated rate, and a value less than 1 implies otherwise. However, because competing risks are present, it is no longer appropriate to interpret the cause-specific hazard ratios in terms of a relative risk type of measure as regular hazard ratios. As Andersen et al. (2012) point out, this is a direct consequence of the loss of the one-to-one relationship between the cause-specific hazard function and the cumulative incidence function.

Starting in SAS/STAT 14.3, you don't have to call PROC PHREG multiple times to fit all the cause-specific Cox models. Instead, you can simply specify the EVENTCODE(COX)= option in the MODEL statement. This option enables you to fit cause-specific Cox models to all the causes simultaneously. As a syntax requirement, you must specify a particular value that indicates which event type is the event of interest. But this piece of information does not interfere with fitting the cause-specific Cox models and is used only to select the cumulative incidence function for inferences and predictions.

The following statements fit the cause-specific Cox model to all the causes by using the EVENTCODE(COX)= option:

```
proc phreg data=Bmt;
   title 'Cause-Specific Hazard Regression for Relapse and Death';
   class Group (order=internal ref=first);
   model T*Status(0)=Group logWaitTime / eventcode(cox)=1;
run;
```

By default, PROC PHREG finds all the unique code values that have not been designated as censoring from the data set and fits a separate cause-specific Cox models to each of them.

Figure 3 displays the estimated effects of the risk factor **Group** on relapse and death, respectively.

Figure 3 Parameter Estimates for Relapse (**Status**=1) and Death (**Status**=2)

Cause-Specific Hazard Regression for Relapse and Death

The PHREG Procedure

Analysis of Maximum Likelihood Estimates								
Parameter		**DF**	**Parameter Estimate**	**Standard Error**	**Chi-Square**	**Pr > ChiSq**	**Hazard Ratio**	**Label**
Group	AML-Low Risk	1	-1.07294	0.46245	5.3830	0.0203	0.342	Group AML-Low Risk
Group	AML-High Risk	1	0.55118	0.36465	2.2847	0.1307	1.735	Group AML-High Risk
logWaittime		1	-0.23061	0.19440	1.4071	0.2355	0.794	

Analysis of Maximum Likelihood Estimates								
Parameter		**DF**	**Parameter Estimate**	**Standard Error**	**Chi-Square**	**Pr > ChiSq**	**Hazard Ratio**	**Label**
Group	AML-Low Risk	1	-0.23553	0.42702	0.3042	0.5812	0.790	Group AML-Low Risk
Group	AML-High Risk	1	0.13265	0.40784	0.1058	0.7450	1.142	Group AML-High Risk
logWaittime		1	0.11202	0.19536	0.3288	0.5664	1.119	

As you can see, the risk factor **Group** has different effects on relapse and death.

It is crucial to check the proportional hazards assumption for the cause-specific Cox model. You can use existing techniques for checking standard Cox regression. The following statements fit the cause-specific Cox model for the event of relapse and use the ASSESS statement to check the proportionality assumption:

```
proc phreg data=Bmt;
   class Group (order=internal ref=first);
   model T*Status(0,2)=Group logWaitTime;
   assess ph / resample seed=1234;
run;
```

Figure 4 displays the observed and simulated score processes for checking the variable **Group**.
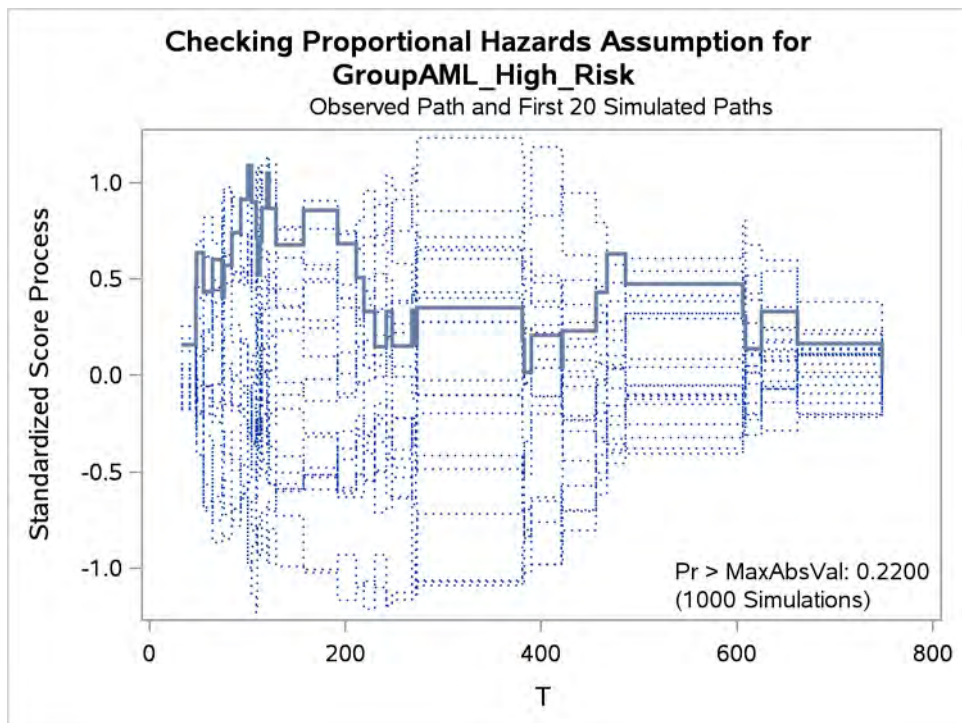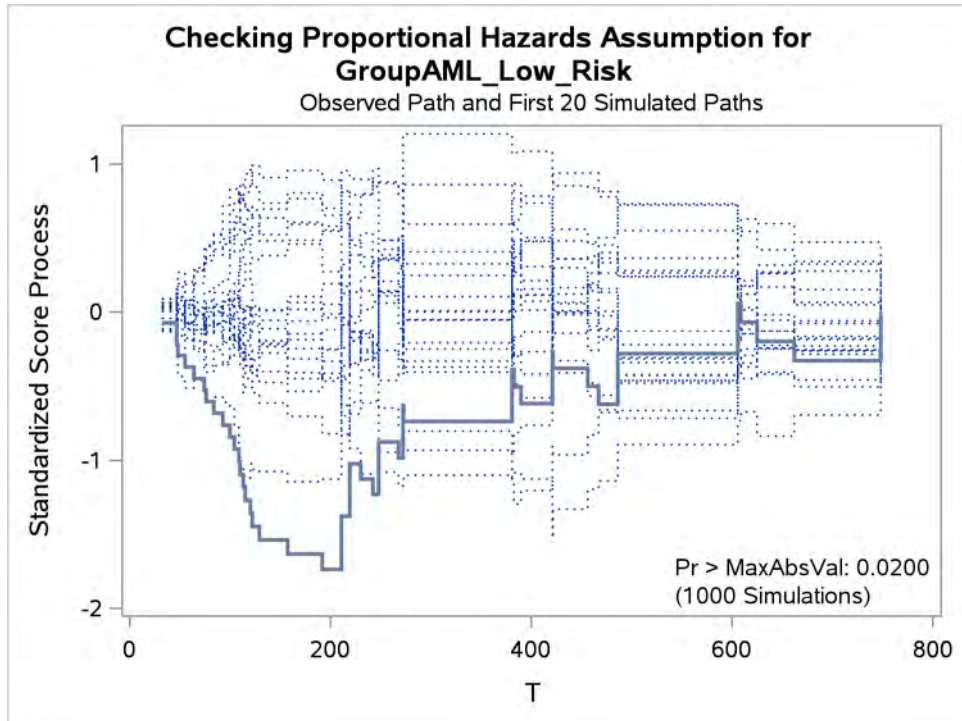
Figure 4 Score Processes for the Variable **Group**

Checking Proportional Hazards Assumption for
GroupAML_Low_Risk
Observed Path and First 20 Simulated Paths

Pr > MaxAbsVal: 0.0200
(1000 Simulations)

Checking Proportional Hazards Assumption for
GroupAML_High_Risk
Observed Path and First 20 Simulated Paths

Pr > MaxAbsVal: 0.2200
(1000 Simulations)

Figure 5 displays the results from assessing all model covariates on the basis of the supremum test of Lin, Wei, and Ying (1993).

Figure 5 Proportionality Tests

	Supremum Test for Proportionals Hazards Assumption			
Variable	**Maximum Absolute Value**	**Replications**	**Seed**	**Pr > MaxAbsVal**
GroupAML_Low_Risk	1.7335	1000	1234	0.0200
GroupAML_High_Risk	1.0871	1000	1234	0.2200
logWaittime	0.7022	1000	1234	0.3110

It appears that the proportionality assumption for the covariate **Group** is questionable because one of its two p-values is less than 0.05.

Modeling the Cumulative Incidence

Analogous to the relationship between the hazard function and the survival function in the absence of competing risks, Fine and Gray (1999) define the subdistribution hazard, which is the hazard of the cumulative incidence function $F_k(t)$,

$$\tilde{h}_k(t) = \frac{d}{dt} \log(1 - F_k(t))$$

Note that in the presence of competing risks, the subdistribution hazard $\tilde{h}_k(t)$ is not the same as the cause-specific hazard $h_k(t)$.

In the spirit of the Cox model, Fine and Gray (1999) introduce a way to model the cumulative incidence function by imposing the proportional hazards assumption on the subdistribution hazards,

$$\tilde{h}_k(t|\mathbf{Z}) = \tilde{h}_{k0}(t) \exp(\boldsymbol{\beta}_k' \mathbf{Z})$$

where $\tilde{h}_{k0}(t)$ is the baseline of the subdistribution hazard of cause k. The Fine and Gray model is also known as the proportional subdistribution hazards model.

The regression coefficients $\boldsymbol{\beta}_k$ can be estimated by maximizing the following pseudo-likelihood function,

$$\tilde{L}(\boldsymbol{\beta}_k) = \prod_i \left(\frac{\exp(\boldsymbol{\beta}_k' \mathbf{Z}_i)}{\sum_{j \in \tilde{\mathcal{R}}_i} w_{ij} \exp(\boldsymbol{\beta}_k' \mathbf{Z}_j)} \right)^{\delta_i = k}$$

where $\tilde{\mathcal{R}}_i$ is the risk set, which includes subjects who are still under observation and those who have experienced a competing event before X_i. The covariance matrix of the parameter estimator is computed as a sandwich estimate.

The subject-specific weights w_{ij} are calculated as follows. Subjects who experience no event of interest before X_i are given a weight $w_{ij} = 1$, and subjects who experience competing events before X_i are given a weight w_{ij} that decreases with time,

$$w_{ij} = \frac{\hat{G}(X_i)}{\hat{G}(\min(X_j, X_i))}$$

where $\hat{G}(t)$ is the Kaplan-Meier estimate of the survival function of the censoring distribution, which is the cumulative probability that a subject is still being followed at time t.

In more general settings, it is possible to allow the weights to depend on a set of covariates, such as via a proportional hazards model. Although the PHREG procedure does not currently support this feature, you can follow the discussion of Geskus (2016) to create an intermediate data set that contains the estimated weights and then use PROC PHREG to fit the model.

The new risk set $\tilde{\mathcal{R}}_i$ is modified from the risk set definition that comes from the regular partial likelihood in order to accommodate the subdistribution hazard. The difference is that the new risk set includes subjects who have experienced a competing event before X_i and are therefore *immortal*. Because of this unnatural formulation, you must take extra care in interpreting the Fine and Gray model, since the subdistribution hazard function might not have the proper interpretation of an instantaneous risk of failure (Andersen and Keiding 2012). Austin and Fine (2017) provide practical recommendations for interpreting the regression coefficients and reporting the results properly.

The Fine and Gray method has been available in the PHREG procedure since SAS/STAT 13.1. The following statements use PROC PHREG to fit the Fine and Gray model:

```
proc phreg data=Bmt;
   title 'Fine and Gray Analysis of Relapse';
   class Group / order=internal ref=first param=glm;
   model T*Status(0) = Group logWaitTime / eventcode=1;
   hazardratio 'Subdistribution Hazards' Group / diff=pairwise;
run;
```

To designate relapse (**Status** = 1) as the event of interest, you specify EVENTCODE=1 in the MODEL statement. The HAZARDRATIO statement requests the subdistribution hazard ratio for each pair of disease groups. The DIFF=PAIRWISE option performs an analysis of all the pointwise subdistribution hazard ratios.

Figure 6 shows a significant effect on relapse between the disease groups (p = 0.0009).

Figure 6 Subdistribution Hazard Regression

Fine and Gray Analysis for Relapse

The PHREG Procedure

Effect	DF	Wald Chi-Square	Pr > ChiSq
		Type 3 Tests	
Group	2	13.6866	0.0011
logWaittime	1	2.1283	0.1446

The subdistribution hazard ratio estimates of one disease group relative to another are displayed in Figure 7.

Figure 7 Pairwise Subdistribution Hazard Ratios

Subdistribution Hazards: Hazard Ratios for Group

Description	Point Estimate	95% Wald Confidence Limits	
Group AML-Low Risk vs AML-High Risk	0.231	0.106	0.503
Group AML-High Risk vs AML-Low Risk	4.323	1.990	9.394
Group AML-Low Risk vs ALL	0.362	0.155	0.843
Group ALL vs AML-Low Risk	2.765	1.186	6.445
Group AML-High Risk vs ALL	1.564	0.763	3.203
Group ALL vs AML-High Risk	0.640	0.312	1.310

The subdistribution hazard of relapse for the AML high-risk patients is 4.3 times that of the AML low-risk patients (95% CI 1.99 to 9.39) and 1.6 times that of the ALL patients (95% CI 0.76 to 3.20). The subdistribution hazard of relapse of the ALL patients is 2.8 times that of the AML low-risk patients (95% CI 1.19 to 6.45).

Although the Fine and Gray model was developed specifically to analyze a single cause of failure, it is not uncommon in practice to apply it to each cause of failure. You should perform such exercises with caution because in general the subdistribution models for different causes cannot hold together over the entire time horizon. The subdistribution models might be consistent in restricted time ranges, but such an assumption needs to be formally justified (Latouche et al. 2013).

There are various ways to check the proportionality assumption of the Fine and Gray model. Although PROC PHREG has no specialized model-checking feature, you can use the log minus log of the subdistribution hazard or the Schoenfeld residuals for this purpose (Fine and Gray 1999).

CUMULATIVE INCIDENCE PREDICTION

After fitting the competing-risks models, you can use the models to make predictions about future cumulative incidence functions. For the Fine and Gray model, predicting CIFs for the event you are modeling is a straightforward task, because the subdistribution hazard is being modeled directly and the CIF is only one transformation away. In contrast, predicting CIFs on the basis of the cause-specific Cox models involves more tedious calculations.

The Fine and Gray Model

Suppose that a Fine and Gray model has been fitted to the kth cause by using the covariates **Z**. With the time-invariant covariates $\mathbf{Z} = \mathbf{z_0}$, the cumulative incidence function can be estimated by

$$\hat{F}_k(t|\mathbf{z_0}) = 1 - \exp[-\hat{\Lambda}_{k0}(t) \exp(\hat{\boldsymbol{\beta}}'_k \mathbf{z_0})]$$

where $\hat{\Lambda}_{k0}(t)$ is the Breslow estimator of the baseline cumulative subdistribution hazard function derived from the pseudo-likelihood function. Its calculation incorporates the modified risk sets and the gradual reduction of weights for the *immortal* subjects.

You can predict the cumulative incidence by using the BASELINE statement in PROC PHREG. You use the COVARIATES= option to specify a data set that contains settings for predicting the cumulative incidence functions. You can also display the cumulative incidence curves by specifying the PLOTS=CIF option in the PROC PHREG statement.

The following statements use the PHREG procedure to plot the predicted cumulative incidence function for each disease group at **logWaitTime** = 5.2, the median value of the log of the waiting times:

```
Data Risk;
   logWaitTime=5.2;
   Group=1; output;
   Group=2; output;
   Group=3; output;
   format Group DiseaseGroup.;
   run;

ods graphics on;
proc phreg data=bmt plots(overlay=stratum)=cif;
   title 'Fine and Gray Analysis of Relapse';
   class Group / param=glm order=internal ref=first;
   model T*Status(0) = Group logWaitTime / eventcode=1;
   baseline covariates=risk out=_null_ / rowid=Group;
run;
```

The OUT= option in the BASELINE statement produces an output data set that contains the predicted values, but specifying the keyword _NULL_ disables it. When the ROWID= option is specified, the formatted names of **Group** are used to identify the predicted CIF curves in the plot.

Figure 8 displays the predicted CIF curves for all three disease groups. You can see that at any given time after the transplant, an AML high-risk patient is more likely to relapse than an ALL patient, and an ALL patient is more likely to relapse than an AML low-risk patient.

**Figure 8** CIFs of the Three Disease Groups at **logTimeWait** = 5.2
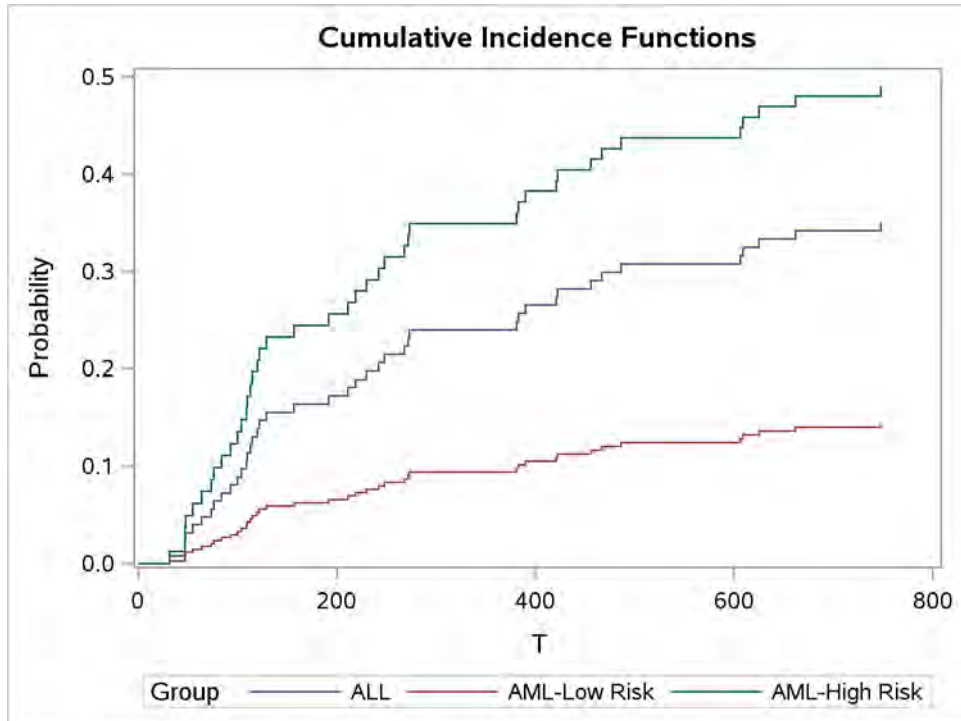
The Fine and Gray model focuses on only one cause of failure. To obtain predictions for the competing causes, you need to fit separate models. As an example, the following statements fit the Fine and Gray model to death without relapse by treating relapse as a competing risk and plot the predicted CIF curves at given covariate patterns:

```
proc phreg data=bmt plots(overlay=stratum)=cif;
   title 'Fine and Gray Analysis of Death';
   class Group / param=glm order=internal ref=first;
   model T*Status(0) = Group logWaitTime / eventcode=2;
   baseline covariates=risk out=_null_ / rowid=Group;
run;
```

Figure 9 displays the predicted CIF curves for all three disease groups.

Figure 9 CIFs of the Three Disease Groups at **logTimeWait** = 5.2
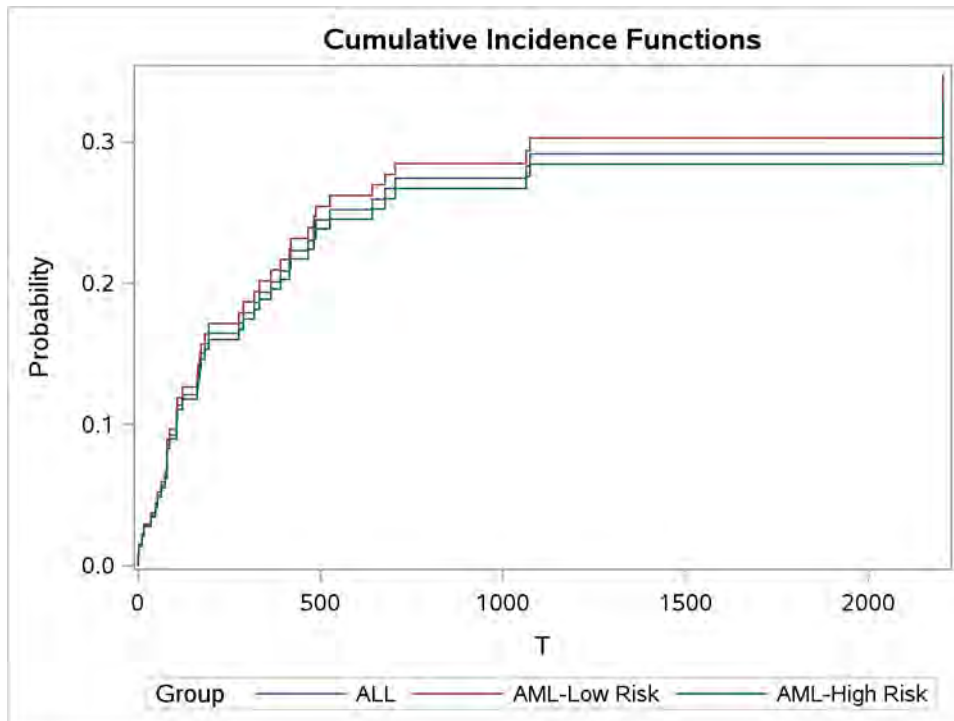
As discussed in the previous section, you must exercise caution when presenting the results from the Fine and Gray models for different event types, because no more than one of the fitted models can be correct.

The Cause-Specific Cox Models

As the previous section shows, predicting the CIF on the basis of the Fine and Gray model is straightforward. Technically, things work quite differently under the cause-specific survival analysis framework. To predict the CIF for a particular cause of failure, you need to use the fitted cause-specific Cox model for each cause of failure.

Assume that separate models have been fitted to each of the K causes by using the covariates **Z**. For the kth cause-specific Cox model, the predicted cumulative cause-specific hazard function estimate for an individual with the covariates $\mathbf{Z} = \mathbf{z_0}$ is

$$\hat{\Lambda}_k(t|\mathbf{z_0}) = \exp(\hat{\boldsymbol{\beta}}_k' \mathbf{z_0})\hat{\Lambda}_{k0}(t)$$

where $\hat{\boldsymbol{\beta}}_k$ is the maximum partial likelihood estimate and $\hat{\Lambda}_{k0}(t)$ is the Breslow estimator of the baseline cumulative cause-specific hazard function.

The predicted survival function is

$$\hat{S}(t|\mathbf{z_0}) = \prod_{s:X_i=s\leq t, \Delta_i >0} \left[1 - \hat{\Lambda}_0(s|\mathbf{z_0})\right]$$

where $\hat{\Lambda}_0(t|\mathbf{z_0}) = \sum_{l=1}^{K} \hat{\Lambda}_k(t|\mathbf{z_0})$

The predicted cumulative incidence is

$$\hat{F}_k(t|\mathbf{z_0}) = \int_0^t \hat{S}(s^-|\mathbf{z_0})d\,\hat{\Lambda}_k(s|\mathbf{z_0})$$

The cause-specific Cox models enable you to predict cumulative incidence for different causes of failure simultaneously. Furthermore, the estimator $\{\hat{F}_k(t|\mathbf{z_0}): k = 1, \ldots, K\}$ has the following property:

$$\sum_1^K \hat{F}_k(t|\mathbf{z_0}) = 1 - \hat{S}(t|\mathbf{z_0})$$

Based on the counting process and martingale theory (Andersen et al. 1995), the variance estimator of $\hat{F}_k(t|\mathbf{z_0})$ is a sum of two terms,

$$\hat{\sigma}_k^2(t|\mathbf{z_0}) = \hat{\sigma}_{k1}^2(t|\mathbf{z_0}) + \hat{\sigma}_{k2}^2(t|\mathbf{z_0})$$

where the two summands quantify different sources of uncertainty associated with the estimation process and have very technical expressions. See Rosthøj, Andersen, and Abildstrom (2004) for a complete presentation of the details.

In PROC PHREG, the syntax to obtain predictions of the CIFs is almost identical between the cause-specific survival analysis and the Fine and Gray analysis. As a requirement, you must designate a cause of interest in the EVENTCODE option. The following statements use PROC PHREG to perform the cause-specific survival analysis and estimate the CIF for the event of relapse (**Status=1**):

```
proc phreg data=Bmt plots(overlay)=cif;
   title 'Cause-Specific Survival Analysis';
   class Group (order=internal ref=first);
   model T*Status(0)=Group logWaitTime / eventcode(cox)=1;
   baseline covariates=Risk out=Out1 cif=_all_ timelist=300 600;
run;
```

The COVARIATES= option in the BASELINE statement specifies the data set that contains the covariate settings for predicting cumulative incidence functions for relapse, and the OUT= option saves the prediction results in a SAS® data set. The TIMELIST= option specifies a list of time points at which CIF predictions are performed for the output data set. The PLOTS= option in the PROC PHREG statement displays the cumulative incidence curves for relapse.

Figure 10 displays the CIF of the event Relapse (**Status**=1) for the three **Group** groups. The curves resemble those that are derived from the Fine and Gray model (Figure 8).

Figure 10 CIFs of the Three Disease Groups at **logTimeWait** = 5.2 (Cause-Specific Survival Analysis)
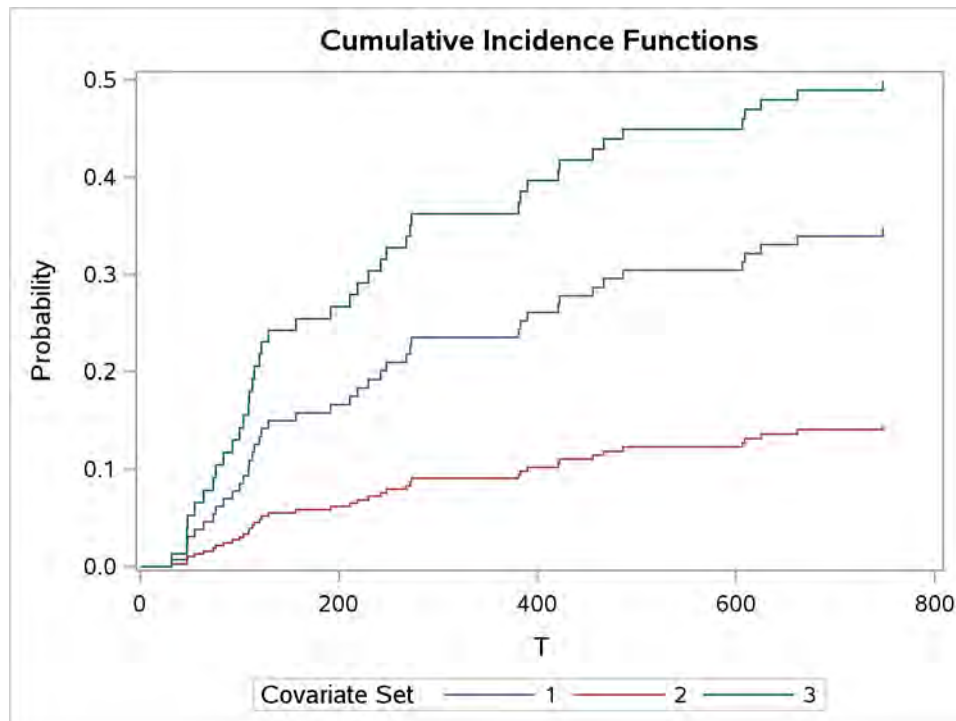

Because the cause-specific Cox models are compatible with one another, you can make CIF predictions for different causes simultaneously. The following statements repeat the fitting of the cause-specific Cox models but predict the CIFs for death without relapse (**Status=2**):

```
proc phreg data=Bmt plots(overlay)=cif;
   title 'Cause-Specific Survival Analysis';
   class Group (order=internal ref=first);
   model T*Status(0)=Group logWaitTime / eventcode(cox)=2;
```

```
        baseline covariates=Risk out=Out2 cif=_all_ timelist=300 600;
    run;
```
Figure 11 displays the CIFs of death without relapse (**Status**=2) for the three groups.

Figure 11 CIFs of the Three Disease Groups at **logTimeWait** = 5.2 (Cause-Specific Survival Analysis)
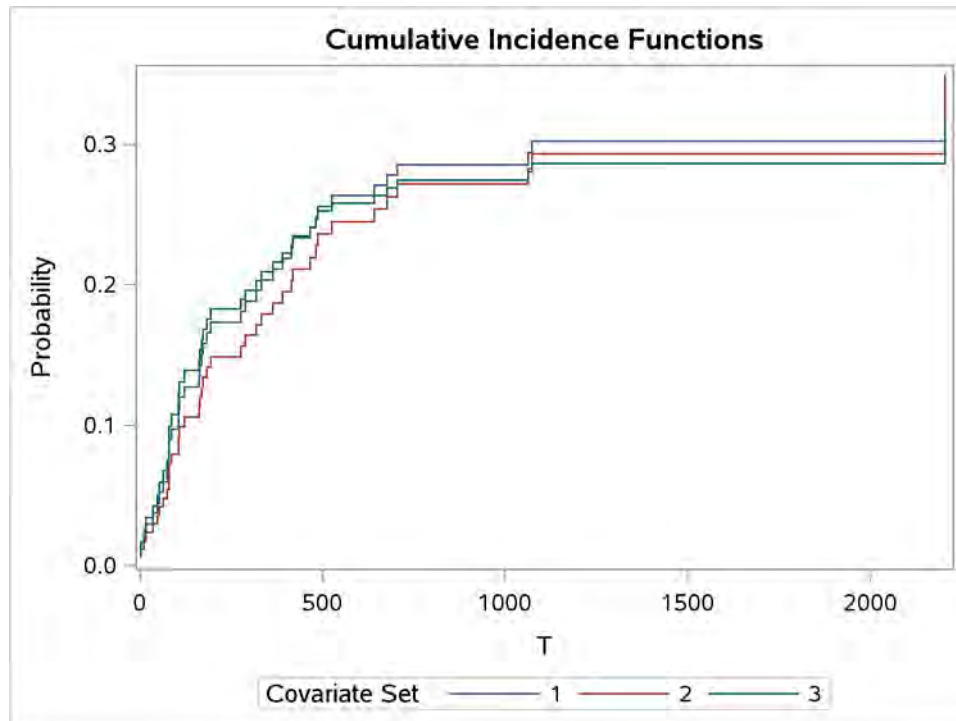
You use the following statements to display the CIF estimates for the ALL risk group (**Group**=1):

```
proc print data=Out1(where=(Group=1));
    title 'CIF Estimates for Relapse and 95% Confidence limits for the ALL Group';
run;

proc print data=Out2(where=(Group=1));
    title 'CIF Estimates for Death and 95% Confidence limits for the ALL Group';
run;
```

Figure 12 shows the point estimate and the confidence limits for the cumulative incidence at selected times for the ALL patients.

Figure 12 Cumulative Incidence Prediction (Cause-Specific Survival Analysis) for the ALL Group

CIF Estimates for Relapse and 95% Confidence limits for the ALL Group

Obs	logWaitTime	Group	T	CIF	StdErrCIF	LowerCIF	UpperCIF
1	5.2	ALL	300	0.23501	0.063654	0.13821	0.39961
2	5.2	ALL	600	0.30423	0.075133	0.18749	0.49364

CIF Estimates for Death and 95% Confidence limits for the ALL Group

Obs	logWaitTime	Group	T	CIF	StdErrCIF	LowerCIF	UpperCIF
1	5.2	ALL	300	0.18845	0.056040	0.10522	0.33754
2	5.2	ALL	600	0.26346	0.071022	0.15533	0.44687

A useful way to display the predicted CIFs in a cause-specific survival analysis is to make a stacked CIF plot. PROC PHREG does not provide this plot in SAS/STAT 14.3, but you can produce it by using the predicted CIFs. Suppose

you want to generate the stacked CIF plot for the first covariate pattern used previously. First, you need to perform the cause-specific survival analysis by using PROC PHREG and save the predicted CIF for each cause to separate data sets. Multiple calls of PROC PHREG are needed because each call can produce CIF predictions for only one single cause at a time. The following statements generate two data sets, **O1** and **O2**, for relapse and death, respectively:

```
Data Risk;
   set Risk;
   if _N_ = 1;
run;

proc phreg data=Bmt plots(overlay)=cif;
   class Group (order=internal ref=first);
   model T*Status(0)=Group logWaitTime / eventcode(cox)=1;
   baseline covariates=Risk out=O1 cif=cif;
run;

proc phreg data=Bmt plots(overlay)=cif;
   class Group (order=internal ref=first);
   model T*Status(0)=Group logWaitTime / eventcode(cox)=2;
   baseline covariates=Risk out=O2 cif=cif2;
run;
```

By default, the predicted CIF for the cause of interest is a step function with jumps at the observed event times. To generate a stacked plot, you need to align the predicted CIFs for different causes with the same set of time points. The following statements obtain the predicted CIFs for the combined set of event times by using simple DATA step operations:

```
date temp;
   t = 1200;
run;

data O12;
   merge O1 O2(keep=t) temp;
   by t;
   retain a 0;
   if cif ne . then a=cif;
   else cif=a;
   drop a;
run;

data O21;
   merge O2 O1(keep=t) temp;
   by t;
   retain a 0;
   if cif2 ne . then a=cif2;
   else cif2=a;
   drop a;
run;
```

Suppose you want to stack the predicted CIF for death onto the predicted CIF for relapse. You need to calculate the sum of the two CIFs at each time point. The following statements perform these calculations for a time horizon of 1,200 days:

```
data Stacked;
   merge O12 O21;
   by t;
   if t<=1200;
   one = 1;
   zero = 0;
   sum = cif+cif2;
run;
```

The following statements create the stacked plot of the CIFs by using PROC SGPLOT:

```
proc sgplot data=Stacked;
   title 'Stacked Plot of the CIFs';
   band x=t upper=one lower=sum  / fillattrs=graphdata1 name='a' transparency=0.5 type=step
        legendlabel='Event-free';
   band x=t upper=sum lower=cif  / fillattrs=graphdata2 name='b' transparency=0.5 type=step
        legendlabel='Death';
   band x=t upper=cif lower=zero / fillattrs=graphdata3 name='c' transparency=0.5 type=step
        legendlabel='Relapse';
   keylegend 'a' 'b' 'c';
   xaxis grid label='T' offsetmin=0 offsetmax=0 values=(0 to 1200 by 200);
   yaxis grid label='Probability' values=(0 to 1 by 0.2) offsetmin=0 offsetmax=0;
run;
```

Figure 13 displays the stacked plot of the CIFs for the ALL group.

Figure 13 Stacked CIF Plot for the ALL Group at **logTimeWait** = 5.2 (Cause-Specific Survival Analysis)
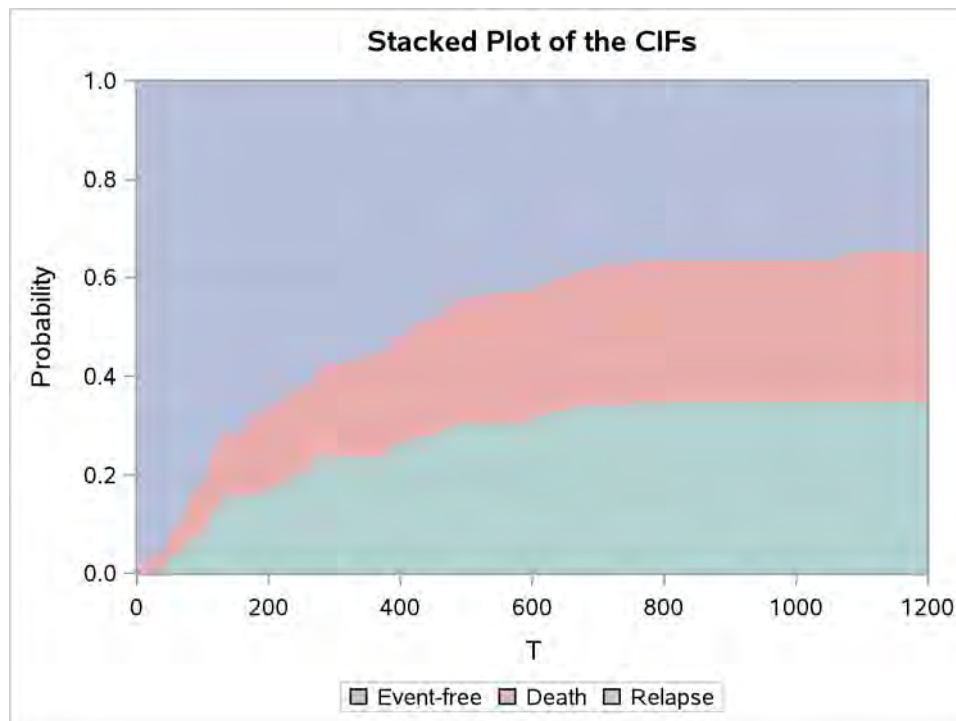
As you can see, the stacked plot is effective in displaying the progression of failures from different event types over time. Note that the area labeled "Event-free" corresponds exactly to the predicted overall survival function $\hat{S}(t\,|\,\mathbf{z}_0)$, which is the complement of the sum of all predicted CIFs.

SUMMARY

Competing-risks analysis extends conventional survival analysis to handle data that have multiple event types. The cause-specific hazard function and the cumulative incidence function are two important quantities of interest. You can apply the Fine and Gray method to model the cumulative incidence function, or you can perform the cause-specific survival analysis by modeling the cause-specific hazards. The former approach models the CIF for one particular cause of failure, and the latter approach models the cause-specific hazard of every competing cause.

You can use PROC PHREG in SAS/STAT 14.3 to analyze competing-risks data by specifying the EVENTCODE option in the MODEL statement. To carry out the Fine and Gray analysis, you specify EVENTCODE(FG)= or simply EVENTCODE=. You use the BASELINE statement to predict the CIFs. To perform the cause-specific survival analysis, you specify EVENTCODE(COX)=. PROC PHREG fits a cause-specific Cox model for each competing cause and predicts the CIF for the cause of interest if the BASELINE statement is specified.

Both the Fine and Gray model and the cause-specific Cox models are extensions of the classical Cox model and inherit its proportionality assumption. Model checking is a necessary step to ensure model validity. It is convenient to check the proportionality assumption in the cause-specific Cox models because you can use all the existing tools for assessing the Cox model.

REFERENCES

Andersen, P. K., Borgan, Ø., Gill, R. D., and Keiding, N. (1995). *Statistical Models Based on Counting Processes.* Corrected 3rd printing. New York: Springer.

Andersen, P. K., Geskus, R. B., de Witte, T., and Putter, H. (2012). "Competing Risks in Epidemiology: Possibilities and Pitfalls." *International Journal of Epidemiology* 41(3):861–870. doi:$10.1093/\mathrm{ije}/\mathrm{dyr213}$.

Andersen, P. K., and Keiding, N. (2012). "Interpretability and Importance of Functionals in Competing Risks and Multistate Models." *Statistics in Medicine* 31(11-12):1074–1088.

Austin, P. C., and Fine, J. P. (2017). "Practical Recommendations for Reporting Fine-Gray Model Analyses for Competing Risk Data." *Statistics in Medicine* 36(27):4391–4400.

Cortese, G., and Andersen, P. K. (2010). "Competing Risks and Time-Dependent Covariates." *Biometrical Journal* 52(1):138–158.

Fine, J. P., and Gray, R. J. (1999). "A Proportional Hazards Model for the Subdistribution of a Competing Risk." *Journal of the American Statistical Association* 94:496–509.

Gardiner, J. C. (2016). "Competing Risks Analyses: Overview of Regression Models." In *Proceedings of the SAS Global Forum 2016 Conference*. Cary, NC: SAS Institute Inc. http://support.sas.com/resources/papers/proceedings16/8620-2016.pdf.

Geskus, R. B. (2016). *Data Analysis with Competing Risks and Intermediate States*. Boca Raton, FL: Chapman & Hall/CRC.

Gray, R. J. (1988). "A Class of K-Sample Tests for Comparing the Cumulative Incidence of a Competing Risk." *Annals of Statistics* 16:1141–1154.

Klein, J. P., and Moeschberger, M. L. (2003). *Survival Analysis: Techniques for Censored and Truncated Data*. 2nd ed. New York: Springer-Verlag.

Latouche, A., Allignol, A., Beyersmann, J., Labopin, M., and Fine, J. P. (2013). "A Competing Risks Analysis Should Report Results on All Cause-specific Hazards and Cumulative Incidence Functions." *Journal of Clinical Epidemiology* 66(6):648–653.

Lin, D. Y., Wei, L. J., and Ying, Z. (1993). "Checking the Cox Model with Cumulative Sums of Martingale-Based Residuals." *Biometrika* 80:557–572.

Prentice, R. L., Kalbfleisch, J. D., Peterson, J., A. V., Flournoy, N., Farewell, V. T., and Breslow, N. E. (1978). "The Analysis of Failure Times in the Presence of Competing Risks." *Biometrics* 34:541–554.

Rosthøj, S., Andersen, P. K., and Abildstrom, S. Z. (2004). "SAS Macros for Estimation of the Cumulative Incidence Functions Based on a Cox Regression Model for Competing Risks Survival Data." *Computer Methods and Programs in Biomedicine* 74(1):69–75.

So, Y., Lin, G., and Johnston, G. (2015). "Using the PHREG Procedure to Analyze Competing-Risks Data." In *Proceedings of the SAS Global Forum 2015 Conference*. Cary, NC: SAS Institute Inc. http://support.sas.com/resources/papers/proceedings15/SAS1855-2015.pdf.

Wolbers, M., Koller, M. T., Stel, V. S., Schaer, B., Jager, K. J., Leffondré, K., and Heinze, G. (2014). "Competing Risks Analyses: Objectives and Approaches." *European Heart Journal* 35(42):2936–2941.

ACKNOWLEDGMENTS

The authors thank Ed Huddleston for his valuable editorial assistance in preparing this paper.

CONTACT INFORMATION

Your comments and questions are valued and encouraged. Contact the authors:

Changbin Guo	Ying So
SAS Institute Inc.	SAS Institute Inc.
SAS Campus Drive	SAS Campus Drive
Cary, NC 27513	Cary, NC 27513
changbin.guo@sas.com	ying.so@sas.com

SAS and all other SAS Institute Inc. product or service names are registered trademarks or trademarks of SAS Institute Inc. in the USA and other countries. [®] indicates USA registration.

Other brand and product names are trademarks of their respective companies.