Title

stcox PH-assumption tests — Tests of proportional-hazards assumption

Description

stphplot plots $-\ln\{-\ln(\text{survival})\}$ curves for each category of a nominal or ordinal covariate versus ln(analysis time). These are often referred to as "log-log" plots. Optionally, these estimates can be adjusted for covariates. The proportional-hazards assumption is not violated when the curves are parallel.

stcoxkm plots Kaplan–Meier observed survival curves and compares them with the Cox predicted curves for the same variable. The closer the observed values are to the predicted, the less likely it is that the proportional-hazards assumption has been violated.

estat phtest tests the proportional-hazards assumption on the basis of Schoenfeld residuals after fitting a model with stcox.

Quick start

Log-log plot of survival

Check for parallel lines in plot of $-\ln\{-\ln(\text{survival})\}$ versus ln(analysis time) for each category of covariate a using stset data

```
stphplot, by(a)
```

As above, but adjust for average values of covariates x1 and x2

```
stphplot, by(a) adjustfor(x1 x2)
```

Same as above

```
stphplot, by(a) adjustfor(x1 x2, atomeans)
```

Adjust for x1 = 0 and x2 = 0

```
stphplot, by(a) adjustfor(x1 x2, atzeros)
```

Kaplan–Meier and predicted survival plot

Compare Kaplan–Meier survival curve with predicted survival from Cox model for each category of covariate a using stset data

```
stcoxkm, by(a)
```

As above, but create separate plots for each level of a

```
stcoxkm, by(a) separate
```

Test using Schoenfeld residuals

Test the proportional-hazards assumption after stcox x1 x2 x3

```
estat phtest
```

As above, and report separate test for each covariate

```
estat phtest, detail
```

Menu

stphplot

Statistics > Survival analysis > Regression models > Graphically assess proportional-hazards assumption

stcoxkm

Statistics > Survival analysis > Regression models > Kaplan-Meier versus predicted survival

estat phtest

Statistics > Survival analysis > Regression models > Test proportional-hazards assumption

Syntax

Check proportional-hazards assumption:

Log-log plot of survival

`stphplot` [*if*] , { `by(`*varname*`)` | `strata(`*varname*`)` } [*stphplot_options*]

Kaplan–Meier and predicted survival plot

`stcoxkm` [*if*] , `by(`*varname*`)` [*stcoxkm_options*]

Using Schoenfeld residuals

`estat phtest` [, *phtest_options*]

stphplot_options	Description
Main	
* `by(`*varname*`)`	fit separate Cox models; the default
* `strata(`*varname*`)`	fit stratified Cox model; requires `adjustfor()`
`adjustfor(`*varlist* [, *suboptions*]`)`	adjust the estimates to specific values of *varlist*; default is overall means
Options	
`nonegative`	plot ln{−ln(survival)}
`nolntime`	plot curves against analysis time
`noshow`	do not show st setting information
Plot	
`plot#opts(`*stphplot_plot_options*`)`	affect rendition of the #th connected line and #th plotted points
Add plots	
`addplot(`*plot*`)`	add other plots to the generated graph
Y axis, X axis, Titles, Legend, Overall	
twoway_options	any options other than `by()` documented in [G-3] ***twoway_options***

*Either `by(`*varname*`)` or `strata(`*varname*`)` is required with `stphplot`.

stphplot_plot_options	Description
cline_options	change look of lines or connecting method
marker_options	change look of markers (color, size, etc.)

stcoxkm_options	Description
Main	
* by(*varname*)	report the nominal or ordinal covariate
ties(breslow)	use Breslow method to handle tied failures
ties(efron)	use Efron method to handle tied failures
ties(exactm)	use exact marginal-likelihood method to handle tied failures
ties(exactp)	use exact partial-likelihood method to handle tied failures
separate	draw separate plot for predicted and observed curves
noshow	do not show st setting information
Observed plot	
obsopts(*stcoxkm_plot_options*)	affect rendition of the observed curve
obs#opts(*stcoxkm_plot_options*)	affect rendition of the #th observed curve; not allowed with separate
Predicted plot	
predopts(*stcoxkm_plot_options*)	affect rendition of the predicted curve
pred#opts(*stcoxkm_plot_options*)	affect rendition of the #th predicted curve; not allowed with separate
Add plots	
addplot(*plot*)	add other plots to the generated graph
Y axis, X axis, Titles, Legend, Overall	
twoway_options	any options other than by() documented in [G-3] ***twoway_options***
byopts(*byopts*)	how subgraphs are combined, labeled, etc.

* by(*varname*) is required with stcoxkm.

stcoxkm_plot_options	Description
connect_options	change look of connecting method
marker_options	change look of markers (color, size, etc.)

You must stset your data before using stphplot and stcoxkm; see [ST] **stset**.

fweights, iweights, and pweights may be specified using stset; see [ST] **stset**.

phtest_options	Description
Main	
`log`	use natural logarithm time-scaling function
`km`	use $1 -$ KM product-limit estimate as the time-scaling function
`rank`	use rank of analysis time as the time-scaling function
`time(`*varname*`)`	use *varname* containing a monotone transformation of analysis time as the time-scaling function
`plot(`*varname*`)`	plot smoothed, scaled Schoenfeld residuals versus time
`bwidth(`*#*`)`	use bandwidth of *#*; default is `bwidth(0.8)`
`detail`	test proportional-hazards assumption separately for each covariate
Scatterplot	
marker_options	change look of markers (color, size, etc.)
marker_label_options	add marker labels; change look or position
Smoothed line	
`lineopts(`*cline_options*`)`	affect rendition of the smoothed line
Y axis, X axis, Titles, Legend, Overall	
twoway_options	any options other than `by()` documented in [G-3] ***twoway_options***

`estat phtest` is not appropriate after estimation with `svy`.

Options

Options are presented under the following headings:

Options for stphplot
Options for stcoxkm
Options for estat phtest

Options for stphplot

Main

`by(`*varname*`)` specifies the nominal or ordinal covariate. Either `by()` or `strata()` is required with `stphplot`.

`strata(`*varname*`)` is an alternative to `by()`. Rather than fitting separate Cox models for each value of *varname*, `strata()` fits one stratified Cox model. You must also specify `adjustfor()` with the `strata()` option; see [ST] **sts graph**.

`adjustfor(`*varlist* [, *suboptions*]`)` adjusts the estimates of the survivor function to specific values of *varlist*. The default is to adjust to overall mean values of covariates. `adjustfor()` can be specified with `by()`; it is required with `strata()`.

suboptions are `atomeans` (the default), `atmeans`, `atzeros`, `atbase`, and `at()`; see [ST] ***adjustfor_option***.

Options

`nonegative` specifies that $\ln\{-\ln(\text{survival})\}$ be plotted instead of $-\ln\{-\ln(\text{survival})\}$.

`nolntime` specifies that curves be plotted against analysis time instead of against ln(analysis time).

`noshow` prevents `stphplot` from showing the key st variables. This option is seldom used because most people type `stset, show` or `stset, noshow` to set whether they want to see these variables mentioned at the top of the output of every st command; see [ST] **stset**.

Plot

`plot#opts(`*stphplot_plot_options*`)` affects the rendition of the #th connected line and #th plotted points; see [G-3] ***cline_options*** and [G-3] ***marker_options***.

Add plots

`addplot(`*plot*`)` provides a way to add other plots to the generated graph; see [G-3] ***addplot_option***.

Y axis, X axis, Titles, Legend, Overall

twoway_options are any of the options documented in [G-3] ***twoway_options***, excluding `by()`. These include options for titling the graph (see [G-3] ***title_options***) and for saving the graph to disk (see [G-3] ***saving_option***).

Options for stcoxkm

Main

`by(`*varname*`)` specifies the nominal or ordinal covariate. `by()` is required.

`ties(breslow|efron|exactm|exactp)` specifies one of the methods available to `stcox` for handling tied failures. If none is specified, `ties(breslow)` is assumed; see [ST] **stcox**.

`separate` produces separate plots of predicted and observed values for each value of the variable specified with `by()`.

`noshow` prevents `stcoxkm` from showing the key st variables. This option is seldom used because most people type `stset, show` or `stset, noshow` to set whether they want to see these variables mentioned at the top of the output of every st command; see [ST] **stset**.

Observed plot

`obsopts(`*stcoxkm_plot_options*`)` affects the rendition of the observed curve; see [G-3] ***connect_options*** and [G-3] ***marker_options***.

`obs#opts(`*stcoxkm_plot_options*`)` affects the rendition of the #th observed curve; see [G-3] ***connect_options*** and [G-3] ***marker_options***. This option is not allowed with `separate`.

Predicted plot

`predopts(`*stcoxkm_connect_options*`)` affects the rendition of the predicted curve; see [G-3] ***connect_options*** and [G-3] ***marker_options***.

`pred#opts(`*stcoxkm_connect_options*`)` affects the rendition of the #th predicted curve; see [G-3] ***connect_options*** and [G-3] ***marker_options***. This option is not allowed with `separate`.

Add plots

`addplot(`*plot*`)` provides a way to add other plots to the generated graph; see [G-3] ***addplot_option***.

Y axis, X axis, Titles, Legend, Overall

twoway_options are any of the options documented in [G-3] ***twoway_options***, excluding `by()`. These include options for titling the graph (see [G-3] ***title_options***) and for saving the graph to disk (see [G-3] ***saving_option***).

`byopts(`*byopts*`)` affects the appearance of the combined graph when `by()` and `separate` are specified, including the overall graph title and the organization of subgraphs. See [G-3] ***by_option***.

Options for estat phtest

Main

`log`, `km`, `rank`, and `time()` are used to specify the time scaling function.

By default, `estat phtest` performs the tests using the identity function, that is, analysis time itself.

`log` specifies that the natural log of analysis time be used.

`km` specifies that 1 minus the Kaplan–Meier product-limit estimate be used.

`rank` specifies that the rank of analysis time be used.

`time(`*varname*`)` specifies a variable containing an arbitrary monotonic transformation of analysis time. You must ensure that *varname* is a monotonic transform.

`plot(`*varname*`)` specifies that a scatterplot and smoothed plot of scaled Schoenfeld residuals versus time be produced for the covariate specified by *varname*. By default, the smoothing is performed using the running-mean method implemented in `lowess, mean noweight`; see [R] **lowess**.

`bwidth(`*#*`)` specifies the bandwidth. Centered subsets of `bwidth()` $\times$ N observations are used for calculating smoothed values for each point in the data except for endpoints, where smaller, uncentered subsets are used. The greater the `bwidth()`, the greater the smoothing. The default is `bwidth(0.8)`.

`detail` specifies that a separate test of the proportional-hazards assumption be produced for each covariate in the Cox model. By default, `estat phtest` produces only the global test.

Scatterplot

marker_options affect the rendition of markers drawn at the plotted points, including their shape, size, color, and outline; see [G-3] ***marker_options***.

marker_label_options specify if and how the markers are to be labeled; see [G-3] ***marker_label_options***.

Smoothed line

`lineopts(`*cline_options*`)` affects the rendition of the smoothed line; see [G-3] ***cline_options***.

Y axis, X axis, Titles, Legend, Overall

twoway_options are any of the options documented in [G-3] ***twoway_options***, excluding `by()`. These include options for titling the graph (see [G-3] ***title_options***) and for saving the graph to disk (see [G-3] ***saving_option***).

Remarks and examples

stata.com

Cox proportional hazards models assume that the hazard ratio is constant over time. Suppose that a group of cancer patients on an experimental treatment is monitored for 10 years. If the hazard of dying for the nontreated group is twice the rate as that of the treated group (HR = 2.0), the proportional-hazards assumption implies that this ratio is the same at 1 year, at 2 years, or at any point on the time scale. Because the Cox model, by definition, is constrained to follow this assumption, it is important to evaluate its validity. If the assumption fails, alternative modeling choices would be more appropriate (for example, a stratified Cox model, time-varying covariates). For examples of testing the proportional-hazards assumption using Stata, see Allison (2014).

`stphplot` and `stcoxkm` provide graphical methods for assessing violations of the proportional-hazards assumption. Although using graphs to assess the validity of the assumption is subjective, it can be a helpful tool.

`stphplot` plots $-\ln\{-\ln(\text{survival})\}$ curves for each category of a nominal or ordinal covariate versus $\ln(\text{analysis time})$. These are often referred to as "log–log" plots. Optionally, these estimates can be adjusted for covariates. If the plotted lines are reasonably parallel, the proportional-hazards assumption has not been violated, and it would be appropriate to base the estimate for that variable on one baseline survivor function.

Another graphical method of evaluating the proportional-hazards assumption, though less common, is to plot the Kaplan–Meier observed survival curves and compare them with the Cox predicted curves for the same variable. This plot is produced with `stcoxkm`. When the predicted and observed curves are close together, the proportional-hazards assumption has not been violated. See Garrett (1997) for more details.

Many popular tests for proportional hazards are, in fact, tests of nonzero slope in a generalized linear regression of the scaled Schoenfeld residuals on time (see Grambsch and Therneau [1994]). The `estat phtest` command tests, for individual covariates and globally, the null hypothesis of zero slope, which is equivalent to testing that the log hazard-ratio function is constant over time. Thus rejection of the null hypothesis of a zero slope indicates deviation from the proportional-hazards assumption. The `estat phtest` command allows three common time-scaling options (`log`, `km`, and `rank`) and also allows you to specify a user-defined function of time through the `time()` option. When no option is specified, the tests are performed using analysis time without further transformation.

▷ Example 1

These examples use data from a leukemia remission study (Garrett 1997). The data consist of 42 patients who are monitored over time to see how long (`weeks`) it takes them to go out of remission (`relapse`: 1 = yes, 0 = no). Half the patients receive a new experimental drug, and the other half receive a standard drug (`treatment1`: 1 = drug A, 0 = standard). White blood cell count, a strong indicator of the presence of leukemia, is divided into three categories (`wbc3cat`: 1 = normal, 2 = moderate, 3 = high).

```
. use https://www.stata-press.com/data/r16/leukemia
(Leukemia Remission Study)

. describe

Contains data from https://www.stata-press.com/data/r16/leukemia.dta
  obs:            42                          Leukemia Remission Study
 vars:             8                          23 Mar 2018 10:39
--------------------------------------------------------------------------------
              storage   display    value
variable name   type    format     label      variable label
--------------------------------------------------------------------------------
weeks           byte    %8.0g                 Weeks in Remission
relapse         byte    %8.0g      yesno      Relapse
treatment1      byte    %8.0g      trt1lbl    Treatment I
treatment2      byte    %8.0g      trt2lbl    Treatment II
wbc3cat         byte    %9.0g      wbclbl     White Blood Cell Count
wbc1            byte    %8.0g                 wbc3cat==Normal
wbc2            byte    %8.0g                 wbc3cat==Moderate
wbc3            byte    %8.0g                 wbc3cat==High
--------------------------------------------------------------------------------
Sorted by: weeks

. stset weeks, failure(relapse)

     failure event:  relapse != 0 & relapse < .
obs. time interval:  (0, weeks]
 exit on or before:  failure

------------------------------------------------------------------------------
         42  total observations
          0  exclusions
------------------------------------------------------------------------------
         42  observations remaining, representing
         30  failures in single-record/single-failure data
        541  total analysis time at risk and under observation
                                                at risk from t =         0
                                     earliest observed entry t =         0
                                          last observed exit t =        35
```

In this example, we examine whether the proportional-hazards assumption holds for drug A versus the standard drug (`treatment1`). First, we will use `stphplot`, followed by `stcoxkm`.

```
. stphplot, by(treatment1)
        failure _d:  relapse
  analysis time _t:  weeks
```

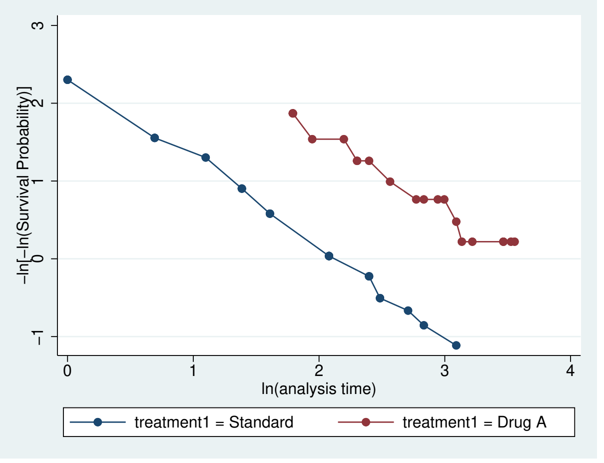

Figure 1.

```
. stcoxkm, by(treatment1) legend(cols(1))
        failure _d:  relapse
  analysis time _t:  weeks
```

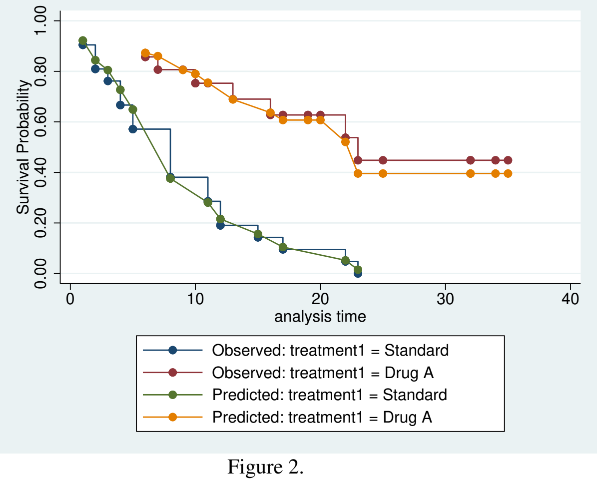

Figure 2.

Figure 1 (`stphplot`) displays lines that are parallel, implying that the proportional-hazards assumption for `treatment1` has not been violated. This is confirmed in figure 2 (`stcoxkm`), where the observed values and predicted values are close together.

The graph in figure 3 is the same as the one in figure 1, adjusted for white blood cell count. By default, this adjustment sets each level of `wbc3cat` to its overall mean. In other words, the results are adjusted based on the observed proportions of individuals having normal, moderate, and high white blood cell counts.

```
. stphplot, strata(treatment1) adjustfor(i.wbc3cat)
        failure _d:  relapse
  analysis time _t:  weeks
```

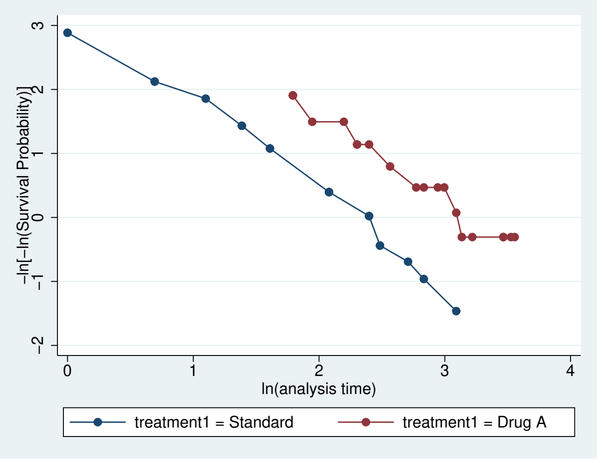

Figure 3.

The lines in figure 3 are still parallel, although they are somewhat closer together. Examining the proportional-hazards assumption on a variable without adjusting for covariates is usually adequate as a diagnostic tool before using the Cox model. However, if you know that adjustment for covariates in a final model is necessary, you may wish to reexamine whether the proportional-hazards assumption still holds.

If we wanted to adjust to the base level of the factor variable `wbc3cat` instead of the level-specific averages, we could have typed

```
. stphplot, strata(treatment1) adjustfor(i.wbc3cat, atbase)
```

Adjusting to a different value, however, would not affect our conclusion about the curves being parallel.

Another variable in this dataset measures a different drug (`treatment2`: $1 =$ drug B, $0 =$ standard). We wish to examine the proportional-hazards assumption for this variable.

```
. stphplot, by(treatment2)
        failure _d:  relapse
  analysis time _t:  weeks
```

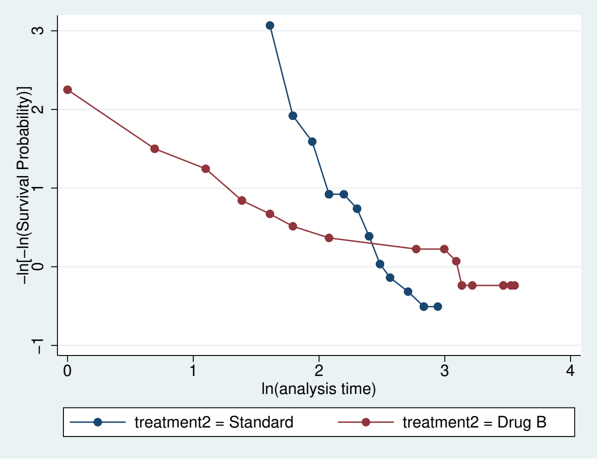

Figure 4.

```
. stcoxkm, by(treatment2) separate legend(cols(1))
        failure _d:  relapse
  analysis time _t:  weeks
```

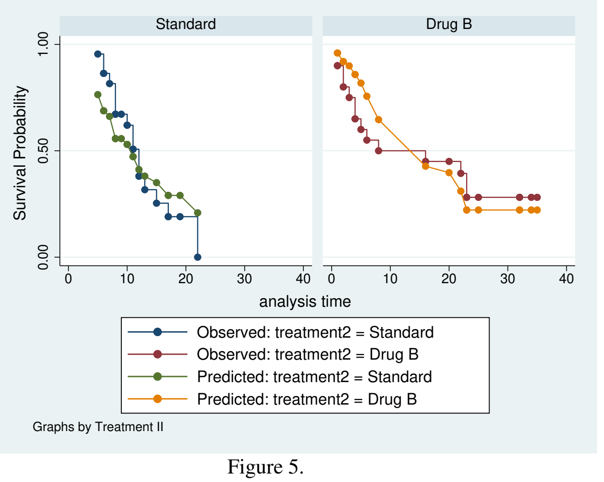

Figure 5.

This variable violates the proportional-hazards assumption. In figure 4, we see that the lines are not only nonparallel but also cross in the data region. In figure 5, we see that there are considerable differences between the observed and predicted values. We have overestimated the positive effect of drug B for the first half of the study and have underestimated it in the later weeks. One hazard ratio describing the effect of this drug would be inappropriate. We definitely would want to stratify on this variable in our Cox model.

◁

▷ Example 2: estat phtest

In this example, we use `estat phtest` to examine whether the proportional-hazards assumption holds for a model with covariates `treatment1` and `wbc3cat`. After `stset`ting the data, we first run `stcox` with these factor variables as regressors. Then we use `estat phtest`:

```
. stset weeks, failure(relapse)

     failure event:  relapse != 0 & relapse < .
obs. time interval:  (0, weeks]
 exit on or before:  failure

------------------------------------------------------------------------------
         42  total observations
          0  exclusions
------------------------------------------------------------------------------
         42  observations remaining, representing
         30  failures in single-record/single-failure data
        541  total analysis time at risk and under observation
                                                at risk from t =         0
                                     earliest observed entry t =         0
                                          last observed exit t =        35

. stcox i.treatment1 i.wbc3cat, nolog

         failure _d:  relapse
   analysis time _t:  weeks

Cox regression -- Breslow method for ties

No. of subjects =           42                  Number of obs    =         42
No. of failures =           30
Time at risk    =          541
                                                LR chi2(3)       =      33.02
Log likelihood  =   -77.476905                  Prob > chi2      =     0.0000

------------------------------------------------------------------------------
          _t | Haz. Ratio   Std. Err.      z    P>|z|     [95% Conf. Interval]
-------------+----------------------------------------------------------------
  treatment1 |
     Drug A  |   .2834551   .1229874    -2.91   0.004     .1211042    .6634517
             |
     wbc3cat |
   Moderate  |   3.637825   2.201306     2.13   0.033     1.111134    11.91015
       High  |   10.92214   7.088783     3.68   0.000      3.06093    38.97284
------------------------------------------------------------------------------

. estat phtest, detail

      Test of proportional-hazards assumption

      Time:  Time
      ----------------------------------------------------------------
                  |       rho            chi2       df       Prob>chi2
      ------------+---------------------------------------------------
      0b.treatme~1|         .               .        1             .
      1.treatment1|  -0.07019            0.15        1          0.6948
      1b.wbc3cat  |         .               .        1             .
      2.wbc3cat   |  -0.03223            0.03        1          0.8650
      3.wbc3cat   |   0.01682            0.01        1          0.9237
      ------------+---------------------------------------------------
      global test |                      0.33        3          0.9551
      ----------------------------------------------------------------
```

Because we specified the `detail` option with the `estat phtest` command, both covariate-specific and global tests were produced. We can see that there is no evidence that the proportional-hazards assumption has been violated.

Another variable in this dataset measures a different drug (`treatment2`: 1 = drug B, 0 = standard). We now wish to examine the proportional-hazards assumption for the previous model by substituting `treatment2` for `treatment1`.

We fit a new Cox model and perform the test for proportional hazards:

```
. stcox i.treatment2 i.wbc3cat, nolog

         failure _d:  relapse
   analysis time _t:  weeks

Cox regression -- Breslow method for ties

No. of subjects =           42                Number of obs    =        42
No. of failures =           30
Time at risk    =          541
                                              LR chi2(3)       =     23.93
Log likelihood  =   -82.019053                Prob > chi2      =    0.0000

------------------------------------------------------------------------------
          _t | Haz. Ratio   Std. Err.      z    P>|z|     [95% Conf. Interval]
-------------+----------------------------------------------------------------
  treatment2 |
     Drug B  |   .8483777   .3469054    -0.40   0.688     .3806529    1.890816
             |
     wbc3cat |
   Moderate  |   3.409628   2.050784     2.04   0.041     1.048905    11.08353
       High  |    14.0562   8.873693     4.19   0.000     4.078529    48.44314
------------------------------------------------------------------------------

. estat phtest, detail

      Test of proportional-hazards assumption

      Time:  Time
      ----------------------------------------------------------------
                   |       rho            chi2       df       Prob>chi2
      -------------+--------------------------------------------------
      0b.treatme~2 |         .               .        1             .
      1.treatment2 |  -0.51672           10.19        1         0.0014
      1b.wbc3cat   |         .               .        1             .
      2.wbc3cat    |  -0.09860            0.29        1         0.5903
      3.wbc3cat    |  -0.03559            0.04        1         0.8448
      -------------+--------------------------------------------------
      global test  |                     10.24        3         0.0166
      ----------------------------------------------------------------
```

`treatment2` violates the proportional-hazards assumption. A single hazard ratio describing the effect of this drug is inappropriate.

The test of the proportional-hazards assumption is based on the principle that, for a given regressor, the assumption restricts $\beta(t_j) = \beta$ for all t_j. This implies that a plot of $\beta(t_j)$ versus time will have a slope of zero. Grambsch and Therneau (1994) showed that $E(s_j^*) + \widehat{\beta} \approx \beta(t_j)$, where s_j^* is the scaled Schoenfeld residual at failure time t_j and $\widehat{\beta}$ is the estimated coefficient from the Cox model. Thus a plot of $s_j^* + \widehat{\beta}$ versus some function of time provides a graphical assessment of the assumption.

Continuing from above, if you type

```
. predict sch*, scaledsch
```

you obtain five variables—`sch1`, `sch2`, `sch3`, `sch4`, and `sch5`—corresponding to the regressors. Ignoring the base categories, `sch2` corresponds to `1.treatment2`, `sch4` corresponds to `2.wbc3cat`, and `sch5` corresponds to `3.wbc3cat`. Given the utility of $s_j^* + \widehat{\beta}$, what is stored in variable `sch2` is actually $s_{j2}^* + \widehat{\beta}_2$ and not just the scaled Schoenfeld residual for the `1.treatment2`, s_{j2}^*, itself. The estimated coefficient, $\widehat{\beta}_2$, is added automatically. The same holds true for the variable representing the next regressor, `sch4` = $s^*_{j4} + \widehat{\beta}_4$, and so on.

As such, a graphical assessment of the proportional-hazards assumption for the first regressor is as simple as

```
. scatter sch2 _t || lfit sch2 _t
```

which plots a scatter of $s^*_{j2} + \widehat{\beta}_2$ versus analysis time, `_t`, and overlays a linear fit. Is the slope zero? The answer is no for `1.treatment2`, and that agrees with our results from `estat phtest`.

◁

❑ Technical note

The tests of the proportional-hazards assumption assume homogeneity of variance across risk sets. This allows the use of the estimated overall (pooled) variance–covariance matrix in the equations. Although these tests have been shown by Grambsch and Therneau (1994) to be fairly robust to departures from this assumption, exercise care where this assumption may not hold, particularly when performing a stratified Cox analysis. In such cases, we recommend that you check the proportional-hazards assumption separately for each stratum.

❑

Video example

How to fit a Cox proportional hazards model and check proportional-hazards assumption

Stored results

`estat phtest` stores the following in `r()`:

Scalars

`r(df)`	global test degrees of freedom
`r(chi2)`	global test χ^2
`r(p)`	global test p-value

Matrices

`r(phtest)`	separate tests for each covariate

Methods and formulas

For one covariate, x, the Cox proportional hazards model reduces to

$$h(t;x) = h_0(t)\exp(x\beta)$$

where $h_0(t)$ is the baseline hazard function from the Cox model. Let $S_0(t)$ and $H_0(t)$ be the corresponding Cox baseline survivor and baseline cumulative hazard functions, respectively.

The proportional-hazards assumption implies that

$$H(t) = H_0(t)\exp(x\beta)$$

or

$$\ln H(t) = \ln H_0(t) + x\beta$$

where $H(t)$ is the cumulative hazard function. Thus, under the proportional-hazards assumption, the logs of the cumulative hazard functions at each level of the covariate have equal slope. This is the basis for the method implemented in `stphplot`.

The proportional-hazards assumption also implies that

$$S(t) = S_0(t)^{\exp(x\beta)}$$

Let $\widehat{S}(t)$ be the estimated survivor function based on the Cox model. This function is a step function like the Kaplan–Meier estimate and, in fact, reduces to the Kaplan–Meier estimate when $x = 0$. Thus for each level of the covariate of interest, we can assess violations of the proportional-hazards assumption by comparing these survival estimates with estimates calculated independently of the model. See Kalbfleisch and Prentice (2002) or Hess (1995).

`stcoxkm` plots Kaplan–Meier estimated curves for each level of the covariate together with the Cox model predicted baseline survival curve. The closer the observed values are to the predicted values, the less likely it is that the proportional-hazards assumption has been violated.

Grambsch and Therneau (1994) presented a scaled adjustment for the Schoenfeld residuals that permits the interpretation of the smoothed residuals as a nonparametric estimate of the log hazard-ratio function. These scaled Schoenfeld residuals, $\mathbf{r}^*_{S_i}$, can be obtained directly with `predict`'s `scaledsch` option; see [ST] **stcox postestimation**.

Scaled Schoenfeld residuals are centered at $\widehat{\beta}$ for each covariate and, when there is no violation of proportional hazards, should have slope zero when plotted against functions of time. The `estat phtest` command uses these residuals, tests the null hypothesis that the slope is equal to zero for each covariate in the model, and performs the global test proposed by Grambsch and Therneau (1994). The test of zero slope is equivalent to testing that the log hazard-ratio function is constant over time.

For a specified function of time, $g(t)$, the statistic for testing the pth individual covariate is, for $\overline{g}(t) = d^{-1}\sum_{i=1}^N \delta_i g(t_i)$,

$$\chi^2_c = \frac{\left[\sum_{i=1}^N \{\delta_i g(t_i) - \overline{g}(t)\} r^*_{S_{pi}}\right]^2}{d\ \mathrm{Var}(\widehat{\beta}_p) \sum_{i=1}^N \{\delta_i g(t_i) - \overline{g}(t)\}^2}$$

which is asymptotically distributed as χ^2 with 1 degree of freedom. $r^*_{S_{pi}}$ is the scaled Schoenfeld residual for observation i, and δ_i indicates failure for observation i, with $d = \sum \delta_i$.

The statistic for the global test is calculated as

$$\chi^2_g = \left[\sum_{i=1}^N \{\delta_i g(t_i) - \overline{g}(t)\}\mathbf{r}_{S_i}\right]' \left[\frac{d\ \mathrm{Var}(\widehat{\boldsymbol{\beta}})}{\sum_{i=1}^N \left\{\delta_i g(t_i) - \overline{g}(t)\right\}^2}\right] \left[\sum_{i=1}^N \{\delta_i g(t_i) - \overline{g}(t)\}\mathbf{r}_{S_i}\right]$$

for $\mathbf{r}_{S_i}$, a vector of the m (unscaled) Schoenfeld residuals for the ith observation; see [ST] **stcox postestimation**. The global test statistic is asymptotically distributed as χ^2 with m degrees of freedom.

The equations for the scaled Schoenfeld residuals and the two test statistics just described assume homogeneity of variance across risk sets. Although these tests are fairly robust to deviations from this assumption, care must be exercised, particularly when dealing with a stratified Cox model.

Acknowledgment

The original versions of `stphplot` and `stcoxkm` were written by Joanne M. Garrett at the University of North Carolina at Chapel Hill. We also thank Garrett for her contributions to the `estat phtest` command.

References

Allison, P. D. 2014. *Event History and Survival Analysis*. 2nd ed. Newbury Park, CA: SAGE.

Barthel, F. M.-S., and P. Royston. 2006. Graphical representation of interactions. *Stata Journal* 6: 348–363.

Breslow, N. E. 1974. Covariance analysis of censored survival data. *Biometrics* 30: 89–99.

Cox, D. R. 1972. Regression models and life-tables (with discussion). *Journal of the Royal Statistical Society, Series B* 34: 187–220.

——. 1975. Partial likelihood. *Biometrika* 62: 269–276.

Cox, D. R., and D. Oakes. 1984. *Analysis of Survival Data*. London: Chapman & Hall/CRC.

Cox, D. R., and E. J. Snell. 1968. A general definition of residuals (with discussion). *Journal of the Royal Statistical Society, Series B* 30: 248–275.

Garrett, J. M. 1997. sbe14: Odds ratios and confidence intervals for logistic regression models with effect modification. *Stata Technical Bulletin* 36: 15–22. Reprinted in *Stata Technical Bulletin Reprints*, vol. 6, pp. 104–114. College Station, TX: Stata Press.

——. 1998. ssa12: Predicted survival curves for the Cox proportional hazards model. *Stata Technical Bulletin* 44: 37–41. Reprinted in *Stata Technical Bulletin Reprints*, vol. 8, pp. 285–290. College Station, TX: Stata Press.

Grambsch, P. M., and T. M. Therneau. 1994. Proportional hazards tests and diagnostics based on weighted residuals. *Biometrika* 81: 515–526.

Hess, K. R. 1995. Graphical methods for assessing violations of the proportional hazards assumption in Cox regression. *Statistics in Medicine* 14: 1707–1723.

Kalbfleisch, J. D., and R. L. Prentice. 2002. *The Statistical Analysis of Failure Time Data*. 2nd ed. New York: Wiley.

Xiao, T., X. He, G. A. Whitemore, and M.-L. Ting Lee. 2012. Threshold regression for time-to-event analysis: The stthreg package. *Stata Journal* 12: 257–283.

Also see

[ST] **stcox** — Cox proportional hazards model

[ST] **sts** — Generate, graph, list, and test the survivor and related functions

[ST] **stset** — Declare data to be survival-time data

[ST] ***adjustfor_option*** — Adjust survivor and related functions for covariates at specific values

[U] **20 Estimation and postestimation commands**