

stci — Confidence intervals for means and percentiles of survival time

Description	Quick start	Menu	Syntax
Options	Remarks and examples	Stored results	Methods and formulas
References	Also see		

Description

`stci` computes means and percentiles of survival time, standard errors, and confidence intervals. For multiple-event data, survival time is the time until a failure.

`stci` can be used with single- or multiple-record or single- or multiple-failure `st` data.

Quick start

Median survival time with standard error and 95% confidence interval using `stset` data
`stci`

Also report medians with standard errors and confidence intervals for each level of `v1`
`stci, by(v1)`

Same as above, but report 99% confidence intervals
`stci, by(v1) level(99)`

Report the 75th percentile of survival times instead of the medians
`stci, by(v1) p(75)`

Mean survival time, computed by exponentially extending curve to zero if last follow-up time is censored
`stci, emean`

Same as above, and plot the extended survivor function
`stci, emean graph`

Menu

Statistics > Survival analysis > Summary statistics, tests, and tables > CIs for means and percentiles of survival time

Syntax

```
stci [if] [in] [, options]
```

<i>options</i>	Description
Main	
<code>by(<i>varlist</i>)</code>	perform separate calculations for each group of <i>varlist</i>
<code>median</code>	calculate median survival times; the default
<code>rmean</code>	calculate mean survival time restricted to longest follow-up time
<code>emean</code>	calculate the mean survival time by exponentially extending the survival curve to zero
<code>p(#)</code>	compute the # percentile of survival times
<code>ccorr</code>	calculate the standard error for <code>rmean</code> using a continuity correction
<code>noshow</code>	do not show st setting information
<code>dd(#)</code>	set maximum number of decimal digits to report
<code>level(#)</code>	set confidence level; default is <code>level(95)</code>
<code>graph</code>	plot exponentially extended survivor function
<code>tmax(#)</code>	set maximum analysis time of # to be plotted

Plot

`cline_options` affect rendition of the plotted lines

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`

You must `stset` your data before using `stci`; see [ST] `stset`.

`by` and `collect` are allowed; see [U] 11.1.10 **Prefix commands**.

Options

Main

`by(varlist)` specifies that separate calculations be made for each group identified by equal values of the variables in *varlist*, resulting in separate summaries and an overall total. *varlist* may contain any number of variables, each of which may be string or numeric.

`median` specifies median survival times. This is the default.

`rmean` and `emean` specify mean survival times. If the longest follow-up time is censored, `emean` (extended mean) computes the mean survival by exponentially extending the survival curve to zero, and `rmean` (restricted mean) computes the mean survival time restricted to the longest follow-up time. If the longest follow-up time is a failure, the restricted mean survival time and the extended mean survival time are equal.

`p(#)` specifies the percentile of survival time to be computed. For example, `p(25)` will compute the 25th percentile of survival times, and `p(75)` will compute the 75th percentile of survival times. Specifying `p(50)` is the same as specifying the `median` option.

`ccorr` specifies that the standard error for the restricted mean survival time be computed using a continuity correction. `ccorr` is valid only with the `rmean` option.

`noshow` prevents `stci` 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](#).

`dd(#)` specifies the maximum number of decimal digits to be reported for standard errors and confidence intervals. This option affects only how values are reported and not how they are calculated.

`level(#)` specifies the confidence level, as a percentage, for confidence intervals. The default is `level(95)` or as set by `set level`; see [U] [20.8 Specifying the width of confidence intervals](#).

`graph` specifies that the exponentially extended survivor function be plotted. This option is valid only when the `emean` option is also specified and is not valid in conjunction with the `by()` option.

`tmax(#)` is for use with the `graph` option. It specifies the maximum analysis time to be plotted.

Plot

`cline_options` affect the rendition of the plotted lines; see [G-3] [cline_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](#)).

Remarks and examples

stata.com

Remarks are presented under the following headings:

Single-failure data

Multiple-failure data

Single-failure data

Here is an example of stci with single-record survival data:

```
. use https://www.stata-press.com/data/r18/page2
. stset, noshow
. stci
```

	Number of subjects	50%	Std. err.	[95% conf. interval]	
Total	40	232	2.562933	213	239

```
. stci, by(group)
```

group	Number of subjects	50%	Std. err.	[95% conf. interval]	
1	19	216	7.661029	190	234
2	21	233	3.081611	232	280
Total	40	232	2.562933	213	239

In the example above, we obtained the median survival time, by default.

To obtain the 25th or any other percentile of survival time, specify the $p(\#)$ option.

```
. stci, p(25)
```

	Number of subjects	25%	Std. err.	[95% conf. interval]	
Total	40	198	10.76878	164	220

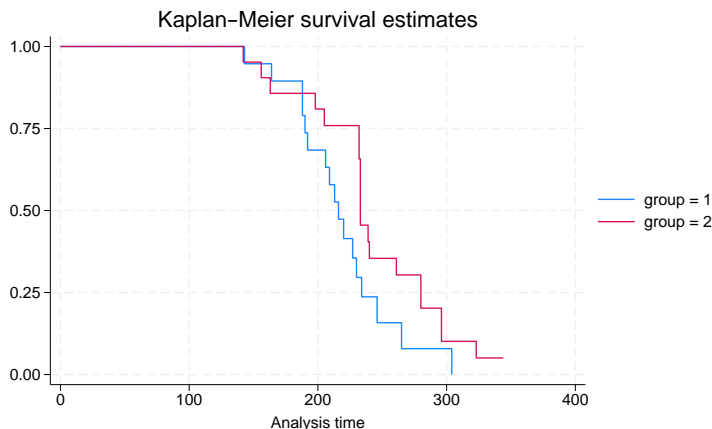
```
. stci, p(25) by(group)
```

group	Number of subjects	25%	Std. err.	[95% conf. interval]	
1	19	190	13.43601	143	213
2	21	232	19.42378	142	233
Total	40	198	10.76878	164	220

The p -percentile of survival time is the analysis time at which $p\%$ of subjects have failed and $1 - p\%$ have not. In the table above, 25% of subjects in group 1 failed by time 190, whereas 25% of subjects in group 2 failed by time 232, indicating a better survival experience for this group.

We can verify the quantities reported by `stci` by plotting and examining the Kaplan–Meier survival curves.

```
. sts graph, by(group)
```



The mean survival time reported by `rmean` is calculated as the area under the Kaplan–Meier survivor function. If the observation with the largest analysis time is censored, the survivor function does not go to zero. Consequently, the area under the curve underestimates the mean survival time.

In the graph above, the survival probability for `group = 1` goes to 0 at analysis time 344, but the survivor function for `group = 2` never goes to 0. For these data, the mean survival time for `group = 1` will be properly estimated, but it will be underestimated for `group = 2`. When we specify the `rmean` option, Stata informs us if any of the mean survival times is underestimated.

```
. stci, rmean by(group)
```

group	Number of subjects	Restricted mean	Std. err.	[95% conf. interval]
1	19	218.7566	9.122424	200.877 236.636
2	21	241.8571(*)	11.34728	219.617 264.097
Total	40	231.3522(*)	7.700819	216.259 246.446

(*) largest observed analysis time is censored, mean is underestimated

Stata flagged the mean for `group = 2` and the overall mean as being underestimated.


If the largest observed analysis time is censored, `stci`'s `emean` option extends the survivor function from the last observed time to zero by using an exponential function and computes the area under the entire curve.

```
. stci, emean
```

	Number of subjects	Extended mean
Total	40	234.2557

The resulting area must be evaluated with care because it is an ad hoc approximation that can at times be misleading. We recommend that you plot and examine the extended survivor function. This is facilitated by the use of `stci`'s `graph` option.

```
. stci, emean graph
```



stci also works with multiple-record survival data. Here is a summary of the multiple-record Stanford heart transplant data introduced in [ST] stset:

```
. use https://www.stata-press.com/data/r18/stan3
(Heart transplant data)
. stset, noshow
. stci
```

	Number of subjects	50%	Std. err.	[95% conf. interval]	
Total	103	100	38.64425	69	219

stci with the by() option may produce results with multiple-record data that you might think are in error:

```
. stci, by(posttran)
```

posttran	Number of subjects	50%	Std. err.	[95% conf. interval]	
0	103	149	43.81077	69	340
1	69	96	58.71712	45	285
Total	103	100	38.64425	69	219

For the number of subjects, $103 + 69 \neq 103$. The posttran variable is not constant for the subjects in this dataset:

```
. stvary posttran
```

Variable	Subjects for whom the variable is				
	constant	varying	never missing	always missing	sometimes missing
posttran	34	69	103	0	0

In this dataset, subjects have one or two records. All subjects were eligible for heart transplantation. They have one record if they die or are lost because of censoring before transplantation, and they have two records if the operation was performed. Then the first record records their survival up to transplantation, and the second records their subsequent survival. `posttran` is 0 in the first record and 1 in the second.

Therefore, all 103 subjects have records with `posttran = 0`, and when `stci` reported results for this group, it summarized the pretransplantation survival. The median survival time was 149 days.

The `posttran = 1` line of `stci`'s output summarizes the posttransplantation survival: 69 patients underwent transplantation, and the median survival time was 96 days. For these data, this is not 96 more days, but 96 days in total. That is, the clock was not reset on transplantation. Thus, without attributing cause, we can describe the differences between the groups as an increased hazard of death at early times followed by a decreased hazard later.

Multiple-failure data

If you simply type `stci` with multiple-failure data, the reported survival time is the survival time to the first failure, assuming that the hazard function is not indexed by number of failures.

Here we have some multiple-failure data:

```
. use https://www.stata-press.com/data/r18/mfail2
. st
-> stset t, id(id) failure(d) time0(t0) exit(time .) noshow
Survival-time data settings
      ID variable: id
      Failure event: d!=0 & d<.
Observed time interval: (t0, t]
      Exit on or before: time .
. stci
```

	Number of subjects	50%	Std. err.	[95% conf. interval]	
Total	926	420	13.42537	394	451

To understand this output, let's also obtain output for each failure separately:

```
. stgen nf = nfailures()
. stci, by(nf)
```

nf	Number of subjects	50%	Std. err.	[95% conf. interval]	
0	926	399	13.91796	381	430
1	529	503	28.53425	425	543
2	221	687	69.38412	549	817
3	58
Total	926	420	13.42537	394	451

The `stgen` command added, for each subject, a variable containing the number of previous failures. `nf` is 0 for a subject, up to and including the first failure. Then `nf` is 1 up to and including the second failure, and then it is 2, and so on; see [ST] [stgen](#).

The first line, corresponding to `nf = 0`, states that among those who had experienced no failures yet, the median time to first failure is 399.

Similarly, the second line, corresponding to `nf = 1`, is for those who have already experienced one failure. The median time of second failures is 503.

When we simply typed `stci`, we obtained the same information shown as the total line of the more detailed output. The total survival time distribution is an estimate of the distribution of the time to first failure, assuming that the hazard function, $h(t)$, is the same across failures—that the second failure is no different from the first failure. This is an odd definition of *same* because the clock, t , is not reset in $h(t)$ upon failure. The hazard of a failure—any failure—at time t is $h(t)$.

Another definition of *same* would have it that the hazard of a failure is given by $h(\tau)$, where τ is the time since the last failure—that the process resets itself. These definitions are different unless $h()$ is a constant function of t .

Let's examine this multiple-failure data, assuming that the process repeats itself. The key variables in this `st` data are `id`, `t0`, `t`, and `d`:

```
. st
-> stset t, id(id) failure(d) time0(t0) exit(time .) noshow
Survival-time data settings
      ID variable: id
      Failure event: d!=0 & d<.
Observed time interval: (t0, t]
      Exit on or before: time .
```

Our goal, for each subject, is to reset `t0` and `t` to 0 after every failure event. We must trick Stata, or at least trick `stset` because it will not let us set data where the same subject has multiple records summarizing the overlapping periods. The trick is create a new `id` variable that is different for every `id-nf` combination (remember, `nf` is the variable we previously created that records the number of prior failures). Then each of the “new” subjects can have their clock start at time 0:

```
. egen newid = group(id nf)
. sort newid t
. by newid: replace t = t - t0[1]
(808 real changes made)
. by newid: gen newt0 = t0 - t0[1]
. stset t, failure(d) id(newid) time0(newt0)
Survival-time data settings
      ID variable: newid
      Failure event: d!=0 & d<.
Observed time interval: (newt0, t]
      Exit on or before: failure
```

```
1,734 total observations
   0 exclusions
```

```
1,734 observations remaining, representing
1,734 subjects
   808 failures in single-failure-per-subject data
435,444 total analysis time at risk and under observation
                                     At risk from t =      0
                                     Earliest observed entry t =      0
                                     Last observed exit t =     797
```

`stset` no longer thinks that we have multiple-failure data. Whereas with `id`, subjects had multiple failures, `newid` gives a unique identity to each `id-nf` combination. Each “new” subject has at most one failure.


```
. stci, by(nf)
      Failure _d: d
      Analysis time _t: t
      ID variable: newid
```

nf	Number of subjects	50%	Std. err.	[95% conf. interval]	
0	926	399	13.91796	381	430
1	529	384	18.22987	359	431
2	221	444	29.80391	325	515
3	58
Total	1734	404	10.29992	386	430

Compare this table with the one we previously obtained. The number of subjects is the same, but the survival times differ because now we measure the times from one failure to the next, whereas previously we measured the time from a fixed point. The time between events in these data appears to be independent of event number.

Similarly, we can obtain the mean survival time for these data restricted to the longest follow-up time:

```
. stci, rmean by(nf)
      Failure _d: d
      Analysis time _t: t
      ID variable: newid
```

nf	Number of subjects	Restricted mean	Std. err.	[95% conf. interval]	
0	926	399.1802	8.872794	381.79	416.571
1	529	397.0077(*)	13.36058	370.821	423.194
2	221	397.8051(*)	25.78559	347.266	448.344
3	58	471(*)	0	471	471
Total	1734	404.7006	7.021657	390.938	418.463

(*) largest observed analysis time is censored, mean is underestimated

Stored results

stci stores the following in `r()`:

Scalars

<code>r(N_sub)</code>	number of subjects	<code>r(se)</code>	standard error
<code>r(p#)</code>	#th percentile	<code>r(lb)</code>	lower bound of CI
<code>r(rmean)</code>	restricted mean	<code>r(ub)</code>	upper bound of CI
<code>r(emean)</code>	extended mean		

Methods and formulas

The percentiles of survival times are obtained from $S(t)$, the Kaplan–Meier product-limit estimate of the survivor function. The 25th percentile, for instance, is obtained as the minimum value of t such that $S(t) \leq 0.75$. The restricted mean is obtained as the area under the Kaplan–Meier product-limit survivor curve. The extended mean is obtained by extending the Kaplan–Meier product-limit survivor curve to zero by using an exponentially fitted curve and then computing the area under the entire curve. If the longest follow-up time ends in failure, the Kaplan–Meier product-limit survivor curve goes to zero, and the restricted mean and extended mean are identical.

The large-sample standard error for the p th percentile of the distribution is given by Collett (2015, 38) and Klein and Moeschberger (2003, 122) as

$$\frac{\sqrt{\widehat{\text{Var}}\{\widehat{S}(t_p)\}}}{\widehat{f}(t_p)}$$

where $\widehat{\text{Var}}\{\widehat{S}(t_p)\}$ is the Greenwood pointwise variance estimate for $\widehat{S}(t_p)$ and $\widehat{f}(t_p)$ is the estimated density function at the p th percentile.

Confidence intervals, however, are not calculated based on this standard error. For a given confidence level, the upper confidence limit for the p th percentile is defined as the first time at which the upper confidence limit for $S(t)$ (based on a $\ln\{-\ln S(t)\}$ transformation) is less than or equal to $1 - p/100$, and, similarly, the lower confidence limit is defined as the first time at which the lower confidence limit of $S(t)$ is less than or equal to $1 - p/100$.

The restricted mean is obtained as the area under the Kaplan–Meier product-limit survivor curve. The extended mean is obtained by extending the Kaplan–Meier product-limit survivor curve to zero by using an exponentially fitted curve and then computing the area under the entire curve. If the longest follow-up time ends in failure, the Kaplan–Meier product-limit survivor curve goes to zero, and the restricted mean and the extended mean are identical.

The standard error for the estimated restricted mean is computed as given by Klein and Moeschberger (2003, 118) and Collett (2015, 390):

$$\widehat{\text{SE}} = \sum_{i=1}^D \widehat{A}_i \sqrt{\frac{d_i}{R_i(R_i - d_i)}}$$

where the sum is over all distinct failure times, \widehat{A}_i is the estimated area under the curve from time i to the maximum follow-up time, R_i is the number of subjects at risk at time i , and d_i is the number of failures at time i .

The $100(1 - \alpha)\%$ confidence interval for the estimated restricted mean is computed as

$$\widehat{A}_i \pm Z_{1-\alpha/2} \widehat{\text{SE}}$$

References

- Collett, D. 2015. *Modelling Survival Data in Medical Research*. 3rd ed. Boca Raton, FL: Chapman and Hall/CRC.
- Klein, J. P., and M. L. Moeschberger. 2003. *Survival Analysis: Techniques for Censored and Truncated Data*. 2nd ed. New York: Springer.

Also see

- [ST] **stdescribe** — Describe survival-time data
- [ST] **stir** — Report incidence-rate comparison
- [ST] **stptime** — Calculate person-time, incidence rates, and SMR
- [ST] **sts** — Generate, graph, list, and test the survivor and related functions

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

[ST] **stvary** — Report variables that vary over time

Stata, Stata Press, and Mata are registered trademarks of StataCorp LLC. Stata and Stata Press are registered trademarks with the World Intellectual Property Organization of the United Nations. StataNow and NetCourseNow are trademarks of StataCorp LLC. Other brand and product names are registered trademarks or trademarks of their respective companies. Copyright © 1985–2023 StataCorp LLC, College Station, TX, USA. All rights reserved.



For suggested citations, see the FAQ on [citing Stata documentation](#).