

power exponential — Power analysis for the exponential test

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

Description

`power exponential` computes sample size or power for survival analysis comparing two exponential survivor functions by using parametric tests for the difference between hazards or, optionally, for the difference between log hazards. It accommodates unequal allocation between the two groups, flexible accrual of subjects into the study, and group-specific losses to follow-up. The accrual distribution may be chosen to be uniform or truncated exponential over a fixed accrual period. Losses to follow-up are assumed to be exponentially distributed. Also the computations may be performed using the conditional or the unconditional approach.

Quick start

Sample size for a test of exponential hazard rates $H_0: \lambda_2 = \lambda_1$ versus $H_a: \lambda_2 \neq \lambda_1$ given control-group hazard rate $h_1 = 0.4$, experimental-group hazard rate $h_2 = 0.2$, equal group sizes, and no censoring using default power of 0.8 and significance level $\alpha = 0.05$

```
power exponential .4 .2
```

Same as above, specified using a hazard ratio of 0.5 instead of the experimental-group hazard rate

```
power exponential .4, hratio(.5)
```

As above, but specify hazard ratios of 0.4, 0.45, 0.5, and 0.55

```
power exponential .4, hratio(.4(.05).55)
```

As above, but display results in a graph

```
power exponential .4, hratio(.4(.05).55) graph
```

Total and per group sample sizes given twice as many observations in the experimental group as in the control group

```
power exponential .4 .2, nratio(2)
```

Sample size with survival probabilities $s_1 = 0.65$ and $s_2 = 0.8$ and reference survival time 2

```
power exponential .65 .8, time(2)
```

As above, specified using survival probability s_1 and a hazard ratio

```
power exponential .65, time(2) hratio(.52)
```

As above, for a one-sided test with power of 0.95

```
power exponential .65, time(2) hratio(.52) onesided power(.95)
```

Sample size for a test of a log hazard-ratio given $h_1 = 0.4$ and hazard ratio of 0.5

```
power exponential .4, hratio(.5) loghazard
```

As above, specifying corresponding survival probabilities s_1 and s_2 at reference time 2

```
power exponential .45 .67, time(2) loghazard
```

Sample size for a design with a 10-year follow-up period and a 1-year accrual period

```
power exponential .4 .2, fperiod(10) aperiod(1)
```

Power for a test of $H_0: \lambda_2 = \lambda_1$, with $h_1 = 0.4$, $h_2 = 0.2$, a sample size of 80, and default $\alpha = 0.05$

```
power exponential .4 .2, n(80)
```

Power for a test of the log hazard-ratio with $\alpha = 0.01$

```
power exponential .4, hratio(.5) loghazard n(200) alpha(.01)
```

Menu

Statistics > Power and sample size

Syntax

Compute sample size

Specify hazard rates

```
power exponential [ $h_1$  [ $h_2$ ]] [, power(numlist) options]
```

Specify survival probabilities

```
power exponential  $s_1$  [ $s_2$ ], time(#) [ power(numlist) options]
```

Compute power

Specify hazard rates

```
power exponential [ $h_1$  [ $h_2$ ]], n(numlist) [options]
```

Specify survival probabilities

```
power exponential  $s_1$  [ $s_2$ ], time(#) n(numlist) [options]
```

where

h_1 is the hazard rate in the control group;

h_2 is the hazard rate in the experimental group;

s_1 is the survival probability in the control group at reference (base) time t ; and

s_2 is the survival probability in the experimental group at reference (base) time t .

h_1 , h_2 and s_1 , s_2 may each be specified either as one number or as a list of values in parentheses (see [U] 11.1.8 **numlist**).

<i>options</i>	Description
Main	
* <u>t</u> ime(<i>numlist</i>)	reference time t for survival probabilities s_1 and s_2
* <u>a</u> lpha(<i>numlist</i>)	significance level; default is <code>alpha(0.05)</code>
* <u>p</u> ower(<i>numlist</i>)	power; default is <code>power(0.8)</code>
* <u>b</u> eta(<i>numlist</i>)	probability of type II error; default is <code>beta(0.2)</code>
* <u>n</u> (<i>numlist</i>)	total sample size; required to compute power
* <u>n</u> 1(<i>numlist</i>)	sample size of the control group
* <u>n</u> 2(<i>numlist</i>)	sample size of the experimental group
* <u>n</u> ratio(<i>numlist</i>)	ratio of sample sizes, $N2/N1$; default is <code>nratio(1)</code> , meaning equal group sizes
<u>n</u> fractional	allow fractional sample sizes
* <u>h</u> ratio(<i>numlist</i>)	hazard ratio (effect size) of the experimental to the control group; default is <code>hratio(0.5)</code> ; may not be combined with <code>lnhratio()</code> or <code>hdifference()</code>
* <u>l</u> nhratio(<i>numlist</i>)	log hazard-ratio (effect size) of the experimental to the control group; may not be combined with <code>hratio()</code> or <code>hdifference()</code>
* <u>h</u> difference(<i>numlist</i>)	difference between the experimental-group and control-group hazard rates (effect size); may not be combined with <code>hratio()</code> or <code>lnhratio()</code>
* <u>d</u> ifference(<i>numlist</i>)	synonym for <code>hdifference()</code>
<u>l</u> oghazard	power or sample-size computation for the test of the difference between log hazards; default is the test of the difference between hazards
<u>u</u> nconditional	power or sample-size computation using the unconditional approach
<u>e</u> ffect(<i>effect</i>)	specify the type of effect to display; default is method specific
<u>o</u> nesided	one-sided test; default is two sided
<u>p</u> arallel	treat number lists in starred options or in command arguments as parallel when multiple values per option or argument are specified (do not enumerate all possible combinations of values)

Accrual/Follow-up

- * studytime(*numlist*) duration of the study; if not specified, the study is assumed to continue until all subjects experience an event (fail)
- * fperiod(*numlist*) length of the follow-up period; if not specified, the study is assumed to continue until all subjects experience an event (fail)
- * aperiod(*numlist*) length of the accrual period; default is `aperiod(0)`, meaning no accrual
- * aprob(*numlist*) proportion of subjects accrued by time t_a under truncated exponential accrual; default is `aprob(0.5)`
- * aptime(*numlist*) proportion of the accrual period, $t_a/\text{aperiod}()$, by which proportion of subjects in `aprob()` is accrued; default is `aptime(0.5)`
- * atime(*numlist*) reference accrual time t_a by which the proportion of subjects in `aprob()` is accrued; default value is $0.5 \times \text{aperiod}()$
- * ashape(*numlist*) shape of the truncated exponential accrual distribution; default is `ashape(0)`, meaning uniform accrual
- * lossprob(*numlist*) proportion of subjects lost to follow-up by time `losstime()` in the control and the experimental groups; default is `lossprob(0)`, meaning no losses to follow-up
- * lossprob1(*numlist*) proportion of subjects lost to follow-up by time `losstime()` in the control group; default is `lossprob1(0)`, meaning no losses to follow-up in the control group
- * lossprob2(*numlist*) proportion of subjects lost to follow-up by time `losstime()` in the experimental group; default is `lossprob2(0)`, meaning no losses to follow-up in the experimental group
- * losstime(*numlist*) reference time t_L by which the proportion of subjects specified in `lossprob()`, `lossprob1()`, or `lossprob2()` is lost to follow-up; default is `losstime(1)`
- * losshaz(*numlist*) loss hazard rates in the control and the experimental groups; default is `losshaz(0)`, meaning no losses to follow-up
- * losshaz1(*numlist*) loss hazard rates in the control group; default is `losshaz1(0)`, meaning no losses to follow-up in the control group
- * losshaz2(*numlist*) loss hazard rates in the experimental group; default is `losshaz2(0)`, meaning no losses to follow-up in the experimental group

Table	
<code>[no]table [(tablespec)]</code>	suppress table or display results as a table; see [PSS] power, table
<code>saving(filename [, replace])</code>	save the table data to <i>filename</i> ; use <code>replace</code> to overwrite existing <i>filename</i>
Graph	
<code>graph [(graphopts)]</code>	graph results; see [PSS] power, graph
Reporting	
<code>show</code>	display group-specific numbers of events and, in the presence of loss to follow-up, numbers of losses
<code>show(showspec)</code>	display group-specific numbers of events, numbers of losses, and event probabilities
<code>notitle</code>	suppress the title

*Specifying a list of values in at least two starred options, or at least two command arguments, or at least one starred option and one argument results in computations for all possible combinations of the values; see [U] 11.1.8 **numlist**. Also see the `parallel` option.

`notitle` does not appear in the dialog box.

<i>effect</i>	Description
<code>hratio</code>	hazard ratio
<code>lnhratio</code>	log hazard-ratio
<code>hdifference</code>	difference between hazard rates
<code>lnhdifference</code>	difference between log hazard-rates (equivalent to log hazard-ratio)
<code>difference</code>	synonym for <code>hdifference</code>

<i>showspec</i>	Description
<code>events</code>	numbers of events
<code>losses</code>	numbers of losses
<code>eventprobs</code>	event probabilities
<code>all</code>	all the above

where *tablespec* is

```
column[:label] [column[:label] [...]] [ , tableopts ]
```

column is one of the columns defined below, and *label* is a column label (may contain quotes and compound quotes).

<i>column</i>	Description	Symbol
alpha	significance level	α
power	power	$1 - \beta$
beta	type II error probability	β
N	total number of subjects	N
N1	number of subjects in the control group	N_1
N2	number of subjects in the experimental group	N_2
nratio	ratio of sample sizes, experimental to control	N_2/N_1
delta	effect size	δ
s1	survival probability in the control group	$S_1(t)$
s2	survival probability in the experimental group	$S_2(t)$
time	reference survival time	t
h1	hazard rate in the control group	λ_1
h2	hazard rate in the experimental group	λ_2
hdiff	difference between hazard rates	$\lambda_2 - \lambda_1$
hratio	hazard ratio	Δ
lnhratio	log hazard-ratio	$\ln(\Delta)$
studytime	duration of a study	T
fperiod	follow-up period	f
aperiod	accrual period	r
aprob	proportion of subjects accrued by time <i>atime</i> (or by <i>aptime</i> ×100% of accrual period)	p_a
aptime	proportion of an accrual period by which <i>aprob</i> ×100% of subjects are accrued	t_a/r
atime	reference accrual time	t_a
ashape	shape of the accrual distribution	γ
E0	total number of events under H_0	E_0
E01	number of events in the control group under H_0	E_{01}
E02	number of events in the experimental group under H_0	E_{02}
Ea	total number of events under H_a	E_a
Ea1	number of events in the control group under H_a	E_{a1}
Ea2	number of events in the experimental group under H_a	E_{a2}
Pr_E01	control-group probability of an event under H_0	$\text{Pr}_{E_{01}}$
Pr_E02	experimental-group probability of an event under H_0	$\text{Pr}_{E_{02}}$
Pr_Ea1	control-group probability of an event under H_a	$\text{Pr}_{E_{a1}}$
Pr_Ea2	experimental-group probability of an event under H_a	$\text{Pr}_{E_{a2}}$
lossprob	proportion of subjects lost to follow-up in the control and experimental groups	$L(t_L)$
lossprob1	proportion of subjects lost to follow-up in the control group	$L_1(t_L)$
lossprob2	proportion of subjects lost to follow-up in the experimental group	$L_2(t_L)$

<code>losstime</code>	reference loss-to-follow-up time	t_L
<code>losshaz</code>	loss hazard rate in the control and experimental groups	η
<code>losshaz1</code>	loss hazard rate in the control group	η_1
<code>losshaz2</code>	loss hazard rate in the experimental group	η_2
<code>L0</code>	total number of losses under H_0	L_0
<code>L01</code>	number of losses in the control group under H_0	L_{01}
<code>L02</code>	number of losses in the experimental group under H_0	L_{02}
<code>La</code>	total number of losses under H_a	L_a
<code>La1</code>	number of losses in the control group under H_a	L_{a1}
<code>La2</code>	number of losses in the experimental group under H_a	L_{a2}
<code>target</code>	target parameter; synonym for <code>h2</code> or <code>hratio</code>	
<code>_all</code>	display all supported columns	

Column `beta` is shown in the default table in place of column `power` if option `beta()` is specified.

Column `hratio` is shown in the default table if option `hratio()` is specified or implied by the command.

Columns `nratio` and `lnhratio` are shown in the default table if the corresponding options are specified.

Columns `h1`, `h2`, `s1`, and `s2` are available and are shown in the default table when the corresponding command arguments are specified.

Columns `time`, `studytime`, `fperiod`, `aperiod`, `apro`, `aptime`, `atime`, `ashape`, `losshaz`, `losshaz1`, `losshaz2`, `lossprob`, `lossprob1`, `lossprob2`, and `losstime` are available and are shown in the default table when the corresponding options are specified.

Columns containing numbers of events, numbers of losses, and probabilities of an event are displayed if specified or if respective options `show(events)`, `show(losses)`, or `show(eventprobs)` are specified. If `show` is specified, numbers of events and losses are displayed. If `show(all)` is specified, numbers of events, numbers of losses, and probabilities are displayed.

Options

Main

`time(#)` specifies a fixed time t (reference survival time) such that the proportions of subjects in the control and experimental groups still alive past this time point are as specified in s_1 and s_2 . If this option is specified, the input parameters s_1 and s_2 are the survival probabilities $S_1(t)$ and $S_2(t)$. Otherwise, the input parameters are assumed to be hazard rates λ_1 and λ_2 given as h_1 and h_2 , respectively.

`alpha()`, `power()`, `beta()`, `n()`, `n1()`, `n2()`, `nratio()`, `nfractional`; see [PSS] [power](#).

`hratio(numlist)` specifies the hazard ratio (effect size) of the experimental group to the control group. The default is `hratio(0.5)`. This value typically defines the clinically significant improvement of the experimental procedure over the control procedure desired to be detected by a test with a certain power. If h_1 and h_2 (or s_1 and s_2) are given, `hratio()` is not allowed and the hazard ratio is computed as h_2/h_1 [or $\ln(s_2)/\ln(s_1)$]. Also see [Alternative ways of specifying effect](#) for various specifications of an effect size.

This option is not allowed with the effect-size determination and may not be combined with `lnhratio()` or `hdifference()`.

`lnhratio(numlist)` specifies the log hazard-ratio (effect size) of the experimental group to the control group. This value typically defines the clinically significant improvement of the experimental procedure over the control procedure desired to be detected by a test with a certain power. If h_1 and h_2 (or s_1 and s_2) are given, `lnhratio()` is not allowed and the log hazard-ratio is computed as $\ln(h_2/h_1)$ [or $\ln\{\ln(s_2)/\ln(s_1)\}$]. Also see [Alternative ways of specifying effect](#) for various specifications of an effect size.

This option is not allowed with the effect-size determination and may not be combined with `hratio()` or `hdifference()`.

`hdifference(numlist)` specifies the difference between the experimental-group hazard rate and the control-group hazard rate. It requires that the control-group hazard rate, the command argument h_1 , is specified. `hdifference()` provides a way of specifying an effect size; see *Alternative ways of specifying effect* for details.

This option is not allowed with the effect-size determination and may not be combined with `hratio()` or `lnhratio()`.

`loghazard` requests sample-size or power computation for the test of the difference between log hazards (or the log hazard-ratio test). This option implies uniform accrual. By default, the test of the difference between hazards is assumed.

`unconditional` requests that the unconditional approach be used for sample-size or power computation; see *The conditional versus unconditional approaches* and *Methods and formulas* for details.

`effect(effect)` specifies the type of the effect size to be reported in the output as `delta`. *effect* is one of `hratio`, `lnhratio`, `hdifference`, or `lnhdifference`. By default, the effect size `delta` is a hazard ratio, `effect(hratio)`, for a hazard-ratio test and a log hazard-ratio, `effect(lnhratio)`, for a log hazard-ratio test (when `schoenfeld` is specified).

`onesided`, `parallel`; see [PSS] **power**.

Accrual/Follow-up

`studytime(numlist)` specifies the duration of the study, T . By default, it is assumed that subjects are followed up until the last subject experiences an event (fails). The (minimal) follow-up period is defined as the length of the period after the recruitment of the last subject to the study until the end of the study. If r is the length of an accrual period and f is the length of the follow-up period, then $T = r + f$. You can specify only two of the three options `studytime()`, `fperiod()`, and `aperiod()`.

`fperiod(numlist)` specifies the follow-up period of the study, f . By default, it is assumed that subjects are followed up until the last subject experiences an event (fails). The (minimal) follow-up period is defined as the length of the period after the recruitment of the last subject to the study until the end of the study. If T is the duration of a study and r is the length of an accrual period, then the follow-up period is $f = T - r$. You can specify only two of the three options `studytime()`, `fperiod()`, and `aperiod()`.

`aperiod(numlist)` specifies the accrual period, r , during which subjects are to be recruited into the study. The default is `aperiod(0)`, meaning no accrual. You can specify only two of the three options `studytime()`, `fperiod()`, and `aperiod()`.

`aprob(numlist)` specifies the proportion of subjects expected to be accrued by time t^* according to the truncated exponential distribution. The default is `aprob(0.5)`. This option is useful when the shape parameter is unknown but the proportion of accrued subjects at a certain time is known. `aprob()` is often used in conjunction with `aptime()` or `atime()`. This option may not be specified with `ashape()` or `loghazard` and requires specifying a nonzero accrual period in `aperiod()`.

`aptime(numlist)` specifies the proportion of the accrual period, t^*/r , by which the proportion of subjects specified in `aprob()` is expected to be accrued according to the truncated exponential distribution. The default is `aptime(0.5)`. This option may not be combined with `atime()`, `ashape()`, or `loghazard` and requires specifying a nonzero accrual period in `aperiod()`.

- `atime(numlist)` specifies the time point t^* , reference accrual time, by which the proportion of subjects specified in `aprob()` is expected to be accrued according to the truncated exponential distribution. The default value is $0.5 \times r$. This option may not be combined with `aptime()`, `ashape()`, or `loghazard` and requires specifying a nonzero accrual period in `aperiod()`. The value in `atime()` may not exceed the value in `aperiod()`.
- `ashape(numlist)` specifies the shape, γ , of the truncated exponential accrual distribution. The default is `ashape(0)`, meaning uniform accrual. This option is not allowed in conjunction with `loghazard` and requires specifying a nonzero accrual period in `aperiod()`.
- `lossprob(numlist)` specifies the proportion of subjects lost to follow-up by time `losstime()` in the control and the experimental groups. The default is `lossprob(0)`, meaning no losses to follow-up. This option requires specifying `aperiod()` or `fperiod()` and may not be combined with `lossprob1()`, `lossprob2()`, `losshaz()`, `losshaz1()`, or `losshaz2()`.
- `lossprob1(numlist)` specifies the proportion of subjects lost to follow-up by time `losstime()` in the control group. The default is `lossprob1(0)`, meaning no losses to follow-up in the control group. This option requires specifying `aperiod()` or `fperiod()` and may not be combined with `lossprob()`, `losshaz()`, `losshaz1()`, or `losshaz2()`.
- `lossprob2(numlist)` specifies the proportion of subjects lost to follow-up by time `losstime()` in the experimental group. The default is `lossprob2(0)`, meaning no losses to follow-up in the experimental group. This option requires specifying `aperiod()` or `fperiod()` and may not be combined with `lossprob()`, `losshaz()`, `losshaz1()`, or `losshaz2()`.
- `losstime(numlist)` specifies the time at which the proportion of subjects specified in `lossprob()` or `lossprob1()` and `lossprob2()` is lost to follow-up, also referred to as the reference loss to follow-up time. The default is `losstime(1)`. This option requires specifying `lossprob()`, `lossprob1()`, or `lossprob2()`.
- `losshaz(numlist)` specifies an exponential hazard rate of losses to follow-up common to both the control and the experimental groups. The default is `losshaz(0)`, meaning no losses to follow-up. This option requires specifying `aperiod()` or `fperiod()` and may not be combined with `lossprob()`, `lossprob1()`, `lossprob2()`, `losshaz1()`, or `losshaz2()`.
- `losshaz1(numlist)` specifies an exponential hazard rate of losses to follow-up, η_1 , in the control group. The default is `losshaz1(0)`, meaning no losses to follow-up in the control group. This option requires specifying `aperiod()` or `fperiod()` and may not be combined with `lossprob()`, `lossprob1()`, `lossprob2()`, or `losshaz()`.
- `losshaz2(numlist)` specifies an exponential hazard rate of losses to follow-up, η_2 , in the experimental group. The default is `losshaz2(0)`, meaning no losses to follow-up in the experimental group. This option requires specifying `aperiod()` or `fperiod()` and may not be combined with `lossprob()`, `lossprob1()`, `lossprob2()`, or `losshaz()`.

Table

`table`, `table()`, `notable`; see [PSS] [power, table](#).

`saving()`; see [PSS] [power](#).

Graph

`graph`, `graph()`; see [PSS] [power, graph](#). Also see the [column](#) table for a list of symbols used by the graphs.

`show` and `show(showspec)` specify to display additional output containing the numbers of events, losses to follow-up, and event probabilities. If `show` is specified, group-specific numbers of events and, in the presence of losses to follow-up, group-specific numbers of losses to follow-up are displayed for the null and alternative hypotheses. With the table output, the numbers are displayed as additional columns.

`showspec` may contain any combination of `events`, `losses`, `eventprobs`, and `all.events`. `events` displays the group-specific numbers of events under the null and alternative hypotheses. `losses`, if present, displays group-specific numbers of losses under the null and alternative hypotheses. `eventprobs` displays group-specific event probabilities under the null and alternative hypotheses. `all` displays all the above.

The following option is available with `power exponential` but is not shown in the dialog box: `notitle`; see [PSS] `power`.

Remarks and examples

[stata.com](http://www.stata.com)

Remarks are presented under the following headings:

Introduction

Using power exponential

Alternative ways of specifying effect

Computing sample size

Computing sample size in the absence of censoring

Computing sample size in the presence of censoring

Nonuniform accrual

Exponential losses to follow-up

The conditional versus unconditional approaches

Link to the sample-size and power computation for the log-rank test

Computing power

Testing hypotheses about two exponential survivor functions

This entry describes the `power exponential` command and the methodology for power and sample-size analysis for a two-sample comparison of exponential survivor functions. See [PSS] `intro` for a general introduction to power and sample-size analysis and [PSS] `power` for a general introduction to the `power` command using hypothesis tests. See *Survival data* in [PSS] `intro` for an introduction to power and sample-size analysis for survival data.

Introduction

Let $S_1(t)$ and $S_2(t)$ be the exponential survivor functions with hazard rates λ_1 and λ_2 in the control and experimental groups, respectively. Define δ to be the treatment effect that can be expressed as a difference, $\psi = \lambda_2 - \lambda_1$, between hazard rates or as the log of the hazard ratio (a difference between log hazard-rates), $\ln(\Delta) = \ln(\lambda_2/\lambda_1) = \ln(\lambda_2) - \ln(\lambda_1)$. Negative values of the treatment effect δ imply the superiority of the experimental treatment over the standard (control) treatment. Denote r and T to be the length of the accrual period and the total duration of the study, respectively. Then, the follow-up period f is $f = T - r$.

Consider a study designed to compare the exponential survivor functions, $S_1(t) = e^{-\lambda_1 t}$ and $S_2(t) = e^{-\lambda_2 t}$, of the two treatment groups. The disparity in survivor functions may be tested using the hazards λ_1 and λ_2 for the exponential model. Depending on the definition of the treatment effect δ , two test statistics based on the difference and on the log ratio of the hazards may be used to conduct tests of the difference between survivor functions using respective null hypotheses, $H_0: \psi = 0$ and $H_0: \ln(\Delta) = 0$.

The basic formula for the sample-size and power calculations for the test of $H_0: \psi = 0$ is proposed by [Lachin \(1981\)](#). He also derives the equation relating the sample size and power allowing for uniform accrual of subjects into the study over the period from 0 to r . [Lachin and Foulkes \(1986\)](#) extend this formula to truncated exponential accrual over the interval 0 to r and exponential losses to follow-up over the interval 0 to T .

The simplest method for the sample-size and power calculations for the test of $H_0: \ln(\Delta) = 0$ is presented by [George and Desu \(1974\)](#). [Rubinstein, Gail, and Santner \(1981\)](#) extend their method to account for uniform accrual and exponential losses to follow-up and apply it to planning the duration of a survival study. The formula that relates the sample size and power for this test and takes into account the uniform accrual and exponential losses to follow-up is formulated by [Lakatos and Lan \(1992\)](#), based on the derivations of [Rubinstein, Gail, and Santner \(1981\)](#).

You can use `power exponential` to

- compute required number of events and sample size when you know power and effect size; or
- compute power when you know sample size (number of events) and effect size.

You can also supply effect size as hazard rates, survival probabilities, hazard ratio, or log hazard-ratio; adjust results for censoring; adjust results for uniform or exponential accrual; adjust results for group-specific exponentially distributed losses to follow-up; and compute results using the conditional or unconditional approach.

Using `power exponential`

`power exponential` computes sample size or power for a test comparing two exponential survivor functions. All computations are performed for a two-sided hypothesis test where, by default, the significance level is set to 0.05. You may change the significance level by specifying the `alpha()` option. You can specify the `onesided` option to request a one-sided test. By default, all computations assume a balanced- or equal-allocation design; see [\[PSS\] unbalanced designs](#) for a description of how to specify an unbalanced design.

To compute a total sample size, you specify an effect size and, optionally, the power of the test in the `power()` option. The default power is set to 0.8. By default, the computed sample size is rounded up. You can specify the `nfractional` option to see the corresponding fractional sample size; see [Fractional sample sizes in \[PSS\] unbalanced designs](#) for an example. The `nfractional` option is allowed only for sample-size determination.

To compute power, you must specify the total sample size in the `n()` option and an effect size.

An effect size may be specified as a hazard ratio in option `hratio()`, as a log hazard-ratio in option `lnhratio()`, or as a difference between hazard rates in option `hdifference()`. By default, a hazard ratio of 0.5 is assumed. For a fixed-duration study, the control-group hazard rate h_1 or the control-group survival probability s_1 must also be specified. See [Alternative ways of specifying effect](#) below for details.

Instead of the total sample size `n()`, you can specify individual group sizes in `n1()` and `n2()` or specify one of the group sizes and `nratio()` when computing power or effect size. See [Two samples in \[PSS\] unbalanced designs](#) for more details.

If the `time()` option is specified, the command's input parameters are the values of survival probabilities in the control (or the less favorable) group, $S_1(t)$, and in the experimental group, $S_2(t)$, at a fixed time, t (reference survival time), specified in `time()`, given as s_1 and s_2 , respectively. Otherwise, the input parameters are assumed to be the values of the hazard rates in the control group, λ_1 , and in the experimental group, λ_2 , given as h_1 and h_2 , respectively. If survival probabilities are

specified, they are converted to hazard rates by using the formula for the exponential survivor function and the value of time t in `t()`.

By default, the estimates of sample sizes or power for the test of the difference between hazards are reported. This may be changed to the test versus the difference between log hazards by using the `loghazard` option. The default conditional approach may be replaced with the unconditional approach by using `unconditional`; see *The conditional versus unconditional approaches*.

If the duration of a study (T) in option `studytime()`, the length of a follow-up period (f) in option `fperiod()`, or the length of an accrual period (r) in option `aperiod()` is not specified, then the study is assumed to continue until all subjects experience an event (failure), regardless of how much time is required. If only `studytime()` is specified or only `fperiod()` is specified, the length of the accrual period is assumed to be zero and the follow-up period equals the duration of the study. If only `aperiod()` is specified, the length of the follow-up is assumed to be zero and the duration of the study equals the length of the accrual period (continuous accrual until the end of the study). If either `aperiod()` or `fperiod()` is specified with `studytime()`, the other one is computed using the relationship $T = r + f$. If both `aperiod()` or `fperiod()` are specified, a fixed-duration study of length $T = r + f$ is assumed.

If an accrual period of length r is specified in the `aperiod()` option, uniform accrual over the period $[0, r]$ is assumed. The accrual distribution may be changed to truncated exponential when the shape parameter is specified in `ashape()`. The combination of the `aprob()` and `aptime()` (or `atime()`) options may be used in place of the `ashape()` option to request the desired shape of the truncated exponential accrual. For examples, see *Nonuniform accrual*.

To take into account exponential losses to follow-up, the `losshaz()` or `lossprob()` and `losstime()` options may be used. Instead of specifying losses common to both groups, you can use options `losshaz1()` and `losshaz2()` or `lossprob1()` and `lossprob2()` to specify group-specific losses to follow-up. For examples, see *Exponential losses to follow-up*.

Alternative ways of specifying effect

`power exponential` provides several ways to specify the disparity between the control-group and experimental-group survivor functions for sample-size and power determinations. You can specify group hazard rates or group survival probabilities at a fixed time t directly. If survival probabilities are specified, they are converted to hazard rates by using the formula for the exponential survivor function and the value of time t . Alternatively, you can specify the control-group hazard rate or the control-group survival probability and an effect size expressed as a hazard ratio, a log hazard-ratio, or a difference between the two hazard rates. The corresponding experimental-group hazard rate will then be computed using the specified values of the control-group hazard rate and effect size.

By default, `power exponential` performs computation assuming a hazard ratio of 0.5. You can use the `hratio()` option to specify a different value for the hazard ratio or you can use the `lnhratio()` option to specify an effect size as a log hazard-ratio. If a control-group hazard rate or survival probability is specified, you can also specify an effect size as a difference between the experimental-group and control-group hazard rates in option `hdifference()`.

For a fixed-duration study when not all subjects experience an event by the end of the study, a control-group hazard rate or a control-group survival probability at time t must be specified in addition to an effect size.

You specify the control-group hazard rate h_1 following the command name. You can use any of the three options mentioned above to specify an effect size. The experimental-group hazard rate is then computed using the specified values of the control-group hazard rate and effect size.

```
power exponential  $h_1$  [ , hratio() | lnhratio() | hdifference() ... ]
```

Alternatively, you can specify the experimental-group hazard rate h_2 directly.

```
power exponential  $h_1$   $h_2$  [ , ... ]
```

Instead of the control-group hazard rate, you can specify the control-group survival probability s_1 at time t ; the reference time t must be specified in option `time()`.

```
power exponential  $s_1$  , time(#) [hratio() | lnhratio() | hdifference() ... ]
```

Similarly to hazard rates, you can specify the experimental-group survival probability at time t instead of an effect size.

```
power exponential  $s_1$   $s_2$  , time(#) [ ... ]
```

The displayed effect size `delta` corresponds to the difference between hazard rates (or the hazard ratio if the control-group hazard is not specified) for the hazard-difference test and to the log hazard-ratio for the log hazard-ratio (or log hazard-difference) test when the `loghazard` option is specified. You can change this by specifying the `effect()` option: `effect(hratio)` reports the hazard ratio, `effect(lnhratio)` reports the log hazard-ratio, and `effect(hdifference)` reports the difference between the experimental-group and control-group hazard rates.

In the following sections, we describe the use of `power exponential` accompanied by examples for computing sample size and power.

Computing sample size

To compute sample size and number of events, you must specify an effect size and, optionally, the power of the test in the `power()` option. A default power of 0.8 is assumed if `power()` is not specified. A hazard ratio of 0.5 is assumed if an effect size is not specified. See [Alternative ways of specifying effect](#) for various ways of specifying an effect size.

Consider the following two types of survival studies: the first type, a type I study, is when investigators have enough resources to monitor the subjects until all of them experience an event (failure) and the second type, a type II study, is when the study terminates after a fixed period of time, regardless of whether all subjects experienced an event by that time.

Computing sample size in the absence of censoring

In this subsection we explore sample-size estimates using different approximations for a type I study. Examples of sample-size determination for a type II study are presented in the next subsection.

In survival studies, the requirement for the sample size is based on the requirement to observe a certain number of events (failures) to ensure a prespecified power of a test to detect a difference in survivor functions. For a type I study, the number of subjects required for the study is the same as the number of events required to be observed in the study because all subjects experience an event by the end of the study.

▷ Example 1: Sample size using the Lachin method

Consider an example from [Lachin \(1981, 107\)](#). A clinical trial is to be conducted to compare the survivor functions in the control and the experimental groups with a one-sided exponential test, based on the difference between hazards, of the superiority of a new treatment ($H_a: \psi < 0$) for a disease with moderate levels of mortality. Subjects in the control group receive a standard treatment and subjects in the experimental group receive a new treatment. From previous studies the yearly hazard rate for the standard treatment was found to be $\lambda_1 = 0.3$, corresponding to 50% survival after 2.3 years. The investigators would like to know how many subjects are required to detect a reduction in hazard to $\lambda_2 = 0.2$ ($H_a: \psi = -0.1$), which corresponds to an increase in survival to 63% at 2.3 years, with 90% power, equal-sized groups, and a significance level, α , of 0.05.

To obtain the estimate of the sample size for the above study, we supply hazard rates 0.3 and 0.2 as arguments and specify the `power(0.9)` option for 90% power and the `onesided` option for a one-sided test.

```
. power exponential 0.3 0.2, power(0.9) onesided
note: input parameters are hazard rates

Estimated sample sizes for two-sample comparison of survivor functions
Exponential test, hazard difference, conditional
Ho: h2 = h1 versus Ha: h2 < h1

Study parameters:
      alpha =    0.0500
      power =    0.9000
      delta =   -0.1000 (hazard difference)

Survival information:
      h1 =    0.3000
      h2 =    0.2000

Estimated sample sizes:
      N =          218
      N per group =    109
```

From the output, a total of 218 events (subjects) must be observed (recruited) in a study to ensure a power of 90% of a one-sided exponential test to detect a 13% increase in survival probability of subjects in the experimental group with $\alpha = 0.05$. Our estimate of 218 of the total number of subjects (109 per group) required for the study is the same as the one reported in [Lachin \(1981, 107\)](#). ◀

▷ Example 2: Sample size using the George–Desu method

[Example 1](#) reports the sample size obtained using the approximation of [Lachin \(1981\)](#) for the test based on the hazard difference. To obtain the sample size using the approximation of [George and Desu \(1974\)](#), for the equivalent alternative $H_a: \ln(\Delta) = -0.4055$ (a test based on the log of the hazard ratio), we need to specify the `loghazard` option.

```

. power exponential 0.3 0.2, power(0.9) onesided loghazard
note: input parameters are hazard rates

Estimated sample sizes for two-sample comparison of survivor functions
Exponential test, log hazard-ratio, conditional
Ho: ln(HR) = 0 versus Ha: ln(HR) < 0

Study parameters:
      alpha =    0.0500
      power =    0.9000
      delta =   -0.4055 (log hazard-ratio)

Survival information:
      h1 =    0.3000
      h2 =    0.2000

Estimated sample sizes:
      N =      210
      N per group = 105

```

The George–Desu method yields a slightly smaller estimate (210) of the total number of events (subjects). [George and Desu \(1974\)](#) studied the accuracy of the two approximations based on ψ and $\ln(\Delta)$ and concluded that the former is slightly conservative; that is, it gives slightly larger sample-size estimates. The latter was found to be accurate to one or two units of the exact solution for equal-sized groups.

◀

□ Technical note

The approach from [example 2](#) may also be used to obtain an approximation to the sample size or power for the exact F test of equality of two exponential mean analysis (life) times (using the relation between a mean and a hazard rate of the exponential distribution, $\mu = 1/\lambda$).

For example, the sample size of 210 obtained above may be used as an approximation to the number of subjects required in a study of which the goal is to detect an increase in a mean analysis (life) time of the experimental group from $3.33 = 1/0.3$ to $5 = 1/0.2$ by using the one-sided 5%-level F test with 90% power.

The test statistic of the F test is a ratio of two sample means from two exponential distributions that has an exact F distribution. The George–Desu method is based on the normal approximation of the distribution of the log of this test statistic. [George and Desu \(1974\)](#) studied this approximation for equal-sized groups and some common values of significance levels, powers, and hazard ratios and found it to be accurate to one or two units of the exact solution.

□

▷ Example 3: Alternative ways of specifying effect

In [Alternative ways of specifying effect](#), we described various ways in which the survival information of the groups can be supplied to `power exponential`. Here we demonstrate several examples.

In [example 1](#), we specified the survival information by supplying the control-group and experimental-group hazard rates.

Instead of the experimental-group hazard rate, we can specify the difference between hazards in the `hdifference()` option and obtain identical results.

```

. power exponential 0.3, power(0.9) onesided hdifference(-0.1)
note: input parameters are hazard rates

Estimated sample sizes for two-sample comparison of survivor functions
Exponential test, hazard difference, conditional
Ho: h2 = h1 versus Ha: h2 < h1

Study parameters:
      alpha =    0.0500
      power =    0.9000
      delta =   -0.1000 (hazard difference)

Survival information:
      h1 =    0.3000
      h2 =    0.2000
      h2 - h1 =   -0.1000

Estimated sample sizes:
      N =        218
      N per group =    109

```

We can redisplay the effect size delta as a hazard ratio instead of the hazard difference:

```

. power exponential 0.3, power(0.9) onesided hdifference(-0.1) effect(hratio)
note: input parameters are hazard rates

Estimated sample sizes for two-sample comparison of survivor functions
Exponential test, hazard difference, conditional
Ho: h2 = h1 versus Ha: h2 < h1

Study parameters:
      alpha =    0.0500
      power =    0.9000
      delta =    0.6667 (hazard ratio)

Survival information:
      h1 =    0.3000
      h2 =    0.2000
      h2 - h1 =   -0.1000

Estimated sample sizes:
      N =        218
      N per group =    109

```

We can specify the hazard ratio of $0.2/0.3 = 0.66667$ in the `hratio()` option instead of `hdifference(-0.1)`.

```

. power exponential 0.3, power(0.9) onesided hratio(0.6667)
note: input parameters are hazard rates

Estimated sample sizes for two-sample comparison of survivor functions
Exponential test, hazard difference, conditional
Ho: h2 = h1 versus Ha: h2 < h1

Study parameters:
      alpha =    0.0500
      power =    0.9000
      delta =   -0.1000 (hazard difference)

Survival information:
      h1 =    0.3000
      h2 =    0.2000
      hratio =    0.6667

Estimated sample sizes:
      N =        218
      N per group =    109

```

We can obtain the same results from `power exponential` if we specify the control-group survival probability of 0.5 at time $t = 2.3$.

```
. power exponential 0.5, time(2.3) power(0.9) onesided hratio(0.6667)
note: input parameters are survival probabilities
Estimated sample sizes for two-sample comparison of survivor functions
Exponential test, hazard difference, conditional
Ho: h2 = h1 versus Ha: h2 < h1
Study parameters:
      alpha =    0.0500
      power =    0.9000
      delta =   -0.1004 (hazard difference)
Survival information:
      h1 =    0.3014          s1 =    0.5000
      h2 =    0.2009          s2 =    0.6299
      hratio = 0.6667          t =    2.3000
Estimated sample sizes:
      N =          218
      N per group =    109
```

We can also specify the experimental-group survival probability of 0.63 at time $t = 2.3$ directly instead of specifying the hazard ratio.

```
. power exponential 0.5 0.63, time(2.3) power(0.9) onesided
note: input parameters are survival probabilities
Estimated sample sizes for two-sample comparison of survivor functions
Exponential test, hazard difference, conditional
Ho: h2 = h1 versus Ha: h2 < h1
Study parameters:
      alpha =    0.0500
      power =    0.9000
      delta =   -0.1005 (hazard difference)
Survival information:
      h1 =    0.3014          s1 =    0.5000
      h2 =    0.2009          s2 =    0.6300
      t =    2.3000
Estimated sample sizes:
      N =          218
      N per group =    109
```

◀

► Example 4: Unbalanced design

By default, `power exponential` computes sample size for a balanced- or equal-allocation design. If we know the allocation ratio of subjects between the groups, we can compute the required sample size for an unbalanced design by specifying the `nratio()` option.

In [example 1](#), we assumed the same numbers of subjects in the two groups. Suppose that we anticipate to recruit twice as many subjects in the experimental group, that is, $n_2/n_1 = 2$. We specify the `nratio(2)` option to compute the required sample size for the specified unbalanced design.

```

. power exponential 0.3 0.2, power(0.9) onesided nratio(2)
note: input parameters are hazard rates

Estimated sample sizes for two-sample comparison of survivor functions
Exponential test, hazard difference, conditional
Ho: h2 = h1 versus Ha: h2 < h1

Study parameters:
      alpha =    0.0500
      power =    0.9000
      delta =   -0.1000 (hazard difference)
      N2/N1 =    2.0000

Survival information:
      h1 =    0.3000
      h2 =    0.2000

Estimated sample sizes:
      N =      242
      N1 =     81
      N2 =     161
      N2/N1 =   1.9877

```

We need a total of 242 subjects—81 in the control group and 161 in the experimental group.

When different from the specified allocation rate, `power exponential` also displays the actual allocation rate corresponding to the reported rounded group sample sizes. If you wish, you can specify the `nfractional` option to see sample sizes without rounding; see [Fractional sample sizes](#) in [PSS] [unbalanced designs](#) for more information.

Also see [Two samples](#) in [PSS] [unbalanced designs](#) for more examples of unbalanced designs for two-sample tests.

◀

Computing sample size in the presence of censoring

Often in practice, investigators may not have enough resources to continue a study until all subjects experience an event and, therefore, plan to terminate the study after a fixed period, T . Some subjects may not experience an event by the end of the study, in which case the (administrative) censoring of subjects occurs. In the presence of censoring, the number of subjects required in a study will be larger than the number of events required to be observed in the study.

We investigate how terminating the study after some fixed period, T , before all subjects experience an event affects the requirements for the sample size. The duration of a study is divided into two phases: an accrual phase of a length r , during which subjects are recruited to the study, and a follow-up phase of a length f , during which subjects are followed up until the end of the study and no new subjects enter the study. The duration of a study, T , is the sum of the lengths of the two phases.

Consider the following study designs. In the first study design, *A*, each subject is followed up for a length of time T . Here the minimum follow-up time f is equal to T , and, consequently, $r = 0$. In practice, however, subjects will often enter the study at random times and will be followed up until the end of a study at time T , in which case the subjects observed later will have a shorter follow-up than subjects who entered the study at the beginning. Therefore, the minimum follow-up time f will be less than T , and r will be equal to $T - f$. In this case the length of the accrual period, r , must be taken into account in the computations. In the presence of an accrual period, subjects may be recruited continuously during a period of length T ($r = T$, $f = 0$) for the second study design, *B*. Or subjects may be recruited for a fixed period, r , and then followed up for a period of time, f ,

during which no new subjects enter the trial, so that the total duration of study is $T = r + f$ (the third design, C).

▷ Example 5: Sample size in the presence of accrual and follow-up periods

Continuing with [example 1](#), assume that the investigators have resources to continue the study for only 5 years, $T = 5$. We specify the duration of the study in the `studytime()` option, and we tabulate sample-size values for different lengths of an accrual period specified as a list (see [\[U\] 11.1.8 numlist](#)) in `aperiod()`. For simplicity, we use the `table()` option to obtain a table containing only columns `power`, `N`, `aperiod`, `fperiod`, `h1`, `h2`, and `alpha`.

```
. power exponential 0.3 0.2, power(0.9) onesided aperiod(0(1)5) studytime(5)
> table(power N aperiod fperiod h1 h2 alpha)
note: input parameters are hazard rates

Estimated sample sizes for two-sample comparison of survivor functions
Exponential test, hazard difference, conditional
Ho: h2 = h1 versus Ha: h2 < h1
```

power	N	aperiod	fperiod	h1	h2	alpha
.9	304	0	5	.3	.2	.05
.9	322	1	4	.3	.2	.05
.9	344	2	3	.3	.2	.05
.9	378	3	2	.3	.2	.05
.9	426	4	1	.3	.2	.05
.9	502	5	0	.3	.2	.05

Note: Uniform accrual; 50% accrued by 50% of accrual period.

For multiple values of parameters, the results are automatically displayed in a table, as we see above. For more examples of tables, see [\[PSS\] power, table](#). If you wish to produce a power plot, see [\[PSS\] power, graph](#).

The first and the last entries of the above table correspond to the extreme cases of no accrual (design *A*) and no follow-up (design *B*), respectively. When `aperiod()` is specified, a uniform accrual is assumed which implies, for example, that 50% of the subjects will be recruited once 50% of the accrual period has elapsed.

For design *A*, the estimate of the sample size, 304, is larger than the earlier estimate of 218 from [example 1](#). That is, if the study in [example 1](#) terminates after 5 years, the requirement for the sample size increases by 39% to ensure that the same number of 218 events is observed.

By trying different values of the follow-up period, we may find that a 30-year follow-up is required if the investigators can recruit no more than 218 subjects: 30 years is required to observe an event for all subjects in this study.

```

. power exponential 0.3 0.2, power(0.9) onesided fperiod(30)
note: input parameters are hazard rates

Estimated sample sizes for two-sample comparison of survivor functions
Exponential test, hazard difference, conditional
Ho: h2 = h1 versus Ha: h2 < h1

Study parameters:
      alpha =    0.0500
      power =    0.9000
      delta =   -0.1000 (hazard difference)

Accrual and follow-up information:
      duration =   30.0000
      follow-up =   30.0000

Survival information:
      h1 =    0.3000
      h2 =    0.2000

Estimated sample sizes:
      N =      218
      N per group =    109

```

Returning to our [table](#), for design *B*, instead of being monitored for 5 years, subjects are continuously recruited throughout those 5 years; the total sample size increases from 304 to 502. The reason for such an increase is that the average analysis time (the time when a subject is at risk of a failure) decreases from 5 to 2.5 and, therefore, reduces the probability of a subject failing by the end of the study.

In general, the estimates of the total sample size steadily increase as the length of the follow-up decreases. That is, the presence of a follow-up period reduces the requirement for the number of subjects in the study. For example, a clinical trial with a 3-year uniform accrual and a 2-year follow-up needs a total of 378 subjects (189 per group) compared with the total of 502 subjects required for a study with no follow-up and a 5-year accrual.

◀

▷ Example 6: Uniform accrual

In [example 5](#), we investigated the effect of the length of accrual on sample size for a type II study when not all subjects experience an event by the end of the study. We specified the length of the accrual period in option `aperiod()` and the duration of the study in option `studytime()`. When `aperiod()` is specified, the accrual distribution is assumed to be uniform, that is, 10% of the subjects are expected to be recruited once 10% of the accrual period has elapsed, 25% of subjects are expected to be recruited once 25% of the accrual period has elapsed, 50% of subjects are expected to be recruited once 50% of the accrual period has elapsed, and so on. Let's compute the sample size for a study with a 3-year uniform accrual and a 2-year follow-up. We use options `aperiod()` and `fperiod()` to specify the accrual and follow-up periods, respectively.

```

. power exponential 0.3 0.2, power(0.9) onesided aperiod(3) fperiod(2)
note: input parameters are hazard rates

Estimated sample sizes for two-sample comparison of survivor functions
Exponential test, hazard difference, conditional
Ho: h2 = h1 versus Ha: h2 < h1

Study parameters:
    alpha =    0.0500
    power =    0.9000
    delta =   -0.1000 (hazard difference)

Accrual and follow-up information:
    duration =    5.0000
    follow-up =    2.0000
    accrual =    3.0000 (uniform)

Survival information:
    h1 =    0.3000
    h2 =    0.2000

Estimated sample sizes:
    N =          378
    N per group =    189

```

The required total sample size is 378 with 189 subjects per group. This is the same sample size we obtained in the table from [example 5](#) with the corresponding values of the accrual and follow-up periods.

◀

Nonuniform accrual

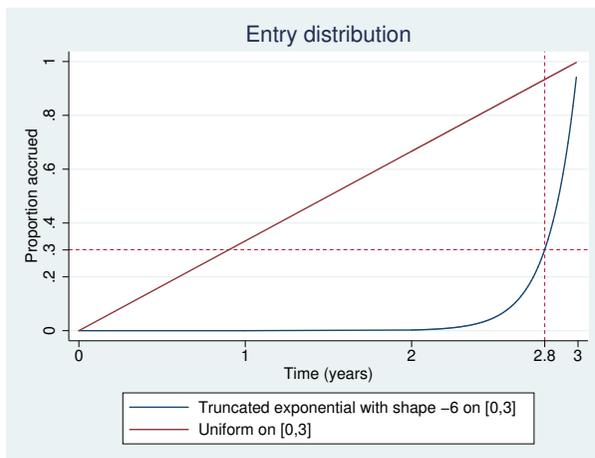
In the presence of an accrual period, `power exponential` performs computations assuming uniform accrual over the period of time r , specified in `aperiod()`. The assumption of uniform accrual may be relaxed by requesting a truncated exponential accrual over the interval 0 to r with shape γ as specified in `ashape(#)`. If an estimate of γ is unavailable, the proportion of subjects expected to be recruited, $G(t^*)$, may be specified in `aprob()` along with either the fixed time by which the subjects were recruited, t^* , in option `atime()` or the elapsed proportion of the accrual period, t^*/r , in option `aptime()`. This information is used to find the corresponding γ by using

$$G(t^*) = \{1 - \exp(-\gamma t^*)\} / \{1 - \exp(-\gamma r)\}$$

Also see [Cleves, Gould, and Marchenko \(2016, sec. 16.2\)](#) for more information, and see [Methods and formulas](#) for technical details.

► Example 7: Truncated exponential entry distribution

Continuing with [example 6](#), we investigate the influence of nonuniform accrual on the estimate of the sample size for a study with a 3-year accrual and a 2-year follow-up. Suppose that the recruitment of subjects to the study is slow for most of the accrual period and increases rapidly toward the end of the recruitment. Consider an extreme case of such an accrual corresponding to shape parameter -6 . The graph of a uniform entry distribution and an exponential entry distribution with shape -6 truncated over $[0, 3]$ is given below.



From the above graph, the accrual of subjects is extremely slow during most of the recruitment period, with 70% of subjects being recruited within the last few months of a 3-year accrual period. Stated another way, according to the graph, only 30% of subjects are expected to be recruited during the first 2.8 years.

To obtain the estimate of the sample size for this study, we type

```
. power exponential 0.3 0.2, power(0.9) onesided aperiod(3) fperiod(2) ashape(-6)
note: input parameters are hazard rates
```

```
Estimated sample sizes for two-sample comparison of survivor functions
Exponential test, hazard difference, conditional
Ho: h2 = h1 versus Ha: h2 < h1
```

Study parameters:

```
alpha = 0.0500
power = 0.9000
delta = -0.1000 (hazard difference)
```

Accrual and follow-up information:

```
duration = 5.0000
follow-up = 2.0000
accrual = 3.0000 (exponential)
accrual(%) = 50.00 (by time t*)
t* = 2.8845 (96.15% of accrual)
```

Survival information:

```
h1 = 0.3000
h2 = 0.2000
```

Estimated sample sizes:

```
N = 516
N per group = 258
```

and conclude that 516 subjects have to be recruited to this study. This sample size ensures 90% power of a one-sided, 5%-level test to detect a reduction in hazard from 0.3 to 0.2 when the accrual of subjects follows the considered truncated exponential distribution. For this extreme case of a negative truncated exponential entry distribution (the concave entry distribution), the estimate of the sample, 516, increases substantially compared with an estimate of 378 from [example 6](#), which assumes a uniform entry distribution. On the other hand, a truncated exponential distribution with positive values of the shape parameter (convex entry distribution) will reduce the requirement for the sample size when compared with uniform accrual.

Suppose that we do not know (or do not wish to guess) the value of the shape parameter. The only information available to us from the above graph is that 30% of the subjects are expected to be recruited in the first 2.8 years. We submit this information in the `aprob()` and `atime()` options, as shown below, and obtain the same estimate of 516 for sample size.

```
. power exponential 0.3 0.2, power(0.9) onesided aperiod(3) fperiod(2)
> aprob(0.3) atime(2.8)
note: input parameters are hazard rates
Estimated sample sizes for two-sample comparison of survivor functions
Exponential test, hazard difference, conditional
Ho: h2 = h1 versus Ha: h2 < h1
Study parameters:
      alpha = 0.0500
      power = 0.9000
      delta = -0.1000 (hazard difference)
Accrual and follow-up information:
      duration = 5.0000
      follow-up = 2.0000
      accrual = 3.0000 (exponential)
      accrual(%) = 30.00 (by time t*)
      t* = 2.8000 (93.33% of accrual)
Survival information:
      h1 = 0.3000
      h2 = 0.2000
Estimated sample sizes:
      N = 516
      N per group = 258
```

Another way we can supply the information about accrual is by specifying a percentage of subjects expected to be recruited by a certain percentage of the accrual period. For example, and equivalent to the above specification, 30% of subjects are expected to be recruited after 93.33% of the accrual period has elapsed. We submit this information in the `aprob()` and `aptime()` options, and we again obtain the same estimate of 516 for sample size.

```
. power exponential 0.3 0.2, power(0.9) onesided aperiod(3) fperiod(2)
> aprob(0.3) aptime(0.9333)
note: input parameters are hazard rates
Estimated sample sizes for two-sample comparison of survivor functions
Exponential test, hazard difference, conditional
Ho: h2 = h1 versus Ha: h2 < h1
Study parameters:
      alpha = 0.0500
      power = 0.9000
      delta = -0.1000 (hazard difference)
Accrual and follow-up information:
      duration = 5.0000
      follow-up = 2.0000
      accrual = 3.0000 (exponential)
      accrual(%) = 30.00 (by time t*)
      t* = 2.7999 (93.33% of accrual)
Survival information:
      h1 = 0.3000
      h2 = 0.2000
Estimated sample sizes:
      N = 516
      N per group = 258
```

Exponential losses to follow-up

Apart from administrative censoring, subjects may not experience an event by the end of the study because of being lost to follow-up for various reasons. See *Survival data* in [PSS] [intro](#) and [PSS] [Glossary](#) for a more detailed description. [Rubinstein, Gail, and Santner \(1981\)](#) and [Lachin and Foulkes \(1986\)](#) extend sample-size and power computations to take into account exponentially distributed losses to follow-up. In addition to being exponentially distributed, losses to follow-up are assumed to be independent of the survival times.

► Example 8: Exponential losses to follow-up

Suppose that in [example 6](#), in the study with a 3-year uniform accrual and a 2-year follow-up, yearly loss hazards in the control and the experimental groups are 0.2. A loss hazard rate common to both groups can be specified in option `losshaz()`.

```
. power exponential 0.3 0.2, power(0.9) onesided aperiod(3) fperiod(2)
> lossshaz(0.2) show
note: input parameters are hazard rates

Estimated sample sizes for two-sample comparison of survivor functions
Exponential test, hazard difference, conditional
Ho: h2 = h1 versus Ha: h2 < h1

Study parameters:
      alpha =    0.0500
      power =    0.9000
      delta =   -0.1000 (hazard difference)

Accrual and follow-up information:
      duration =    5.0000
      follow-up =    2.0000
      accrual =    3.0000 (uniform)

Survival information:
      h1 =    0.3000
      h2 =    0.2000

Loss-to-follow-up information:
      lh1 =    0.2000
      lh2 =    0.2000

Estimated expected number of events:
      E|Ha =    213      E|Ho =    216
      E1|Ha =    121      E1|Ho =    108
      E2|Ha =    92       E2|Ho =    108

Estimated expected number of losses to follow-up:
      L|Ha =    173      L|Ho =    172
      L1|Ha =    81      L1|Ho =    86
      L2|Ha =    92      L2|Ho =    86

Estimated sample sizes:
      N =    500
      N per group =    250
```

The sample size required for a one-sided, 5%-level test to detect a reduction in hazard from 0.3 to 0.2 with 90% power increases from 378 (see [example 6](#)) to 500. We observe that for the extreme case of losses to follow-up, sample size increases significantly. A conservative adjustment commonly applied in practice is $n(1 + p_L)$, where p_L is the expected proportion of losses to follow-up in both groups combined. For this example, p_L may be computed as $0.5(0.369 + 0.324) \approx 0.35$ from table 2 of [Lachin and Foulkes \(1986\)](#). Then the conservative estimate of the sample size is $378(1 + 0.35) = 510$, which is slightly greater than 500, the actual required sample size.

We also requested that additional information about the expected number of events and losses to follow-up under the null and under the alternative hypothesis be displayed by using the `show` option. From the above output, a total of 173 subjects (81 from the control group and 92 from the experimental group) are expected to be lost in the study with exponentially distributed losses with yearly rates of 0.2 in each group under the alternative hypothesis.

If the proportion of subjects lost to follow-up by a fixed period in each group is available, it can be supplied by using the `lossprob()` and `losstime()` options rather than loss to follow-up rates. For example, in the above study approximately 33%, $1 - \exp(-0.2 \times 2) \approx 0.33$, of subjects in each group are lost at time 2 (years). We can obtain the same estimates of sample sizes by typing

```
. power exponential 0.3 0.2, power(0.9) onesided aperiod(3) fperiod(2)
> lossprob(0.33) losstime(2)
(output omitted)
```

4

The conditional versus unconditional approaches

Denote δ to be the effect size, and denote $\hat{\lambda}_1$ and $\hat{\lambda}_2$ to be the maximum likelihood estimates of the respective hazard-rate parameters. Consider the two effect-size estimators based on the difference between the hazard rates, $\hat{\lambda}_2 - \hat{\lambda}_1$, and based on the log of the hazard ratio, $\ln(\hat{\lambda}_2/\hat{\lambda}_1)$. Both estimators are asymptotically normal under the null and under the alternative hypothesis.

We adopt [Chow, Shao, and Wang \(2008, 173\)](#) terminology when referring to the conditional and unconditional tests. The conditional test is the test that uses the constraint $\lambda_2 = \lambda_1$ (conditional on H_0) when computing the variance of the effect-size estimator under the null. The unconditional test is the test that does not use the above constraint when computing the variance of the effect-size estimator under the null. The score and the Wald tests are each one of the examples of conditional and unconditional tests, respectively. [Chow, Shao, and Wang \(2008\)](#) note that neither of the two tests (conditional or unconditional) is always more powerful than the other under the alternative hypothesis. Therefore, there is no definite recommendation of which one is preferable in practice.

The conditional approach relies on the following relationship between sample size and power (given in [Lachin 1981](#)) to compute estimates of required sample size or power,

$$|\delta| = z_{1-\alpha} \{\text{Var}(\delta, H_0)\}^{1/2} + z_{1-\beta} \{\text{Var}(\delta, H_a)\}^{1/2}$$

where $z_{1-\alpha}$ and $z_{1-\beta}$ are the $(1-\alpha)$ th and the $(1-\beta)$ th quantiles of the standard normal distribution, and $\text{Var}(\delta, H_0)$ and $\text{Var}(\delta, H_a)$ are the asymptotic variances under the null and under the alternative, respectively, of the effect-size estimator, $\hat{\delta}$. This approach uses the variance of the estimator conditional on the hypothesis type.

The unconditional approach replaces $\text{Var}(\delta, H_0)$ with $\text{Var}(\delta, H_a)$ in the above and uses the variance under the alternative to compute the estimates of sample size and power:

$$|\delta| = (z_{1-\alpha} + z_{1-\beta}) \{\text{Var}(\delta, H_a)\}^{1/2}$$

Therefore, the resulting formulas based on the two approaches are different.

[Lakatos and Lan \(1992\)](#) formulate the sample-size formula for the log hazard-ratio test based on the method of [Rubinstein, Gail, and Santner \(1981\)](#). This formula is based on the unconditional approach. [Lachin and Foulkes \(1986\)](#) provide the sample-size formula for the test of the log of the hazard ratio that uses the conditional approach. They also present both conditional and unconditional

versions of formulas for the test based on the difference between hazards. As noted by [Lachin and Foulkes \(1986\)](#), sample sizes estimated based on the unconditional approach will be larger than the estimates based on the conditional approach for equal-sized groups.

Both approaches are available with `power exponential`; the conditional is the default and the unconditional may be requested by specifying the `unconditional` option. Refer to [Methods and formulas](#) for the formulas underlying these approaches.

► Example 9: Sample size using the Rubinstein–Gail–Santner method

Consider the following scenario in [Lakatos and Lan \(1992, table I\)](#). A 10-year survival study with a 1-year accrual period and a 9-year follow-up is conducted to compare the survivor functions of the two groups by using a two-sided, 0.05 exponential test based on the log of the hazard ratio. The probability of surviving to the end of a study for subjects in the control group is 0.8 [$S_1(t) = 0.8, t = 10$]. Subjects are recruited uniformly over the interval $[0, 1]$. [Lakatos and Lan \(1992\)](#) report an estimate of 664 for the sample size required to detect a change in the hazard of the experimental group corresponding to the hazard ratio $\Delta = 0.5$ with 90% power by using the Rubinstein–Gail–Santner (1981) method. To obtain the estimates according to this method, we need to specify both `loghazard` and `unconditional`.

```
. power exponential 0.8, t(10) power(0.9) aperiod(1) fperiod(9) loghazard
> unconditional
note: input parameters are survival probabilities
Estimated sample sizes for two-sample comparison of survivor functions
Exponential test, log hazard-ratio, unconditional
Ho: ln(HR) = 0 versus Ha: ln(HR) != 0
Study parameters:
      alpha =    0.0500
      power =    0.9000
      delta =   -0.6931 (log hazard-ratio)
Accrual and follow-up information:
      duration =   10.0000
      follow-up =    9.0000
      accrual =    1.0000 (uniform)
Survival information:
      h1 =    0.0223      s1 =    0.8000
      h2 =    0.0112      s2 =    0.8944
      hratio =   0.5000      t =   10.0000
Estimated sample sizes:
      N =      664
      N per group =   332
```

Because the default value of the hazard ratio is 0.5, we omit the `hratio(0.5)` option in the above. From the output, we obtain the same estimate of 664 of the sample size as reported in [Lakatos and Lan \(1992\)](#).

◀

In the absence of censoring, the estimates of the sample size or power based on the test of log of the hazard ratio are the same for the conditional and the unconditional approaches. For example, both

```
. power exponential 0.8, t(10) power(0.9) loghazard
(output omitted)
```

and

```
. power exponential 0.8, t(10) power(0.9) loghazard unconditional
(output omitted)
```

produce the same estimate of the sample size (88). The asymptotic variance of maximum likelihood estimates of the log of the hazard ratio does not depend on hazard rates when there is no censoring and, therefore, does not depend on the type of hypothesis, $\text{Var}(\widehat{\delta}, H_0) = \text{Var}(\widehat{\delta}, H_a) = 2/N$.

Link to the sample-size and power computation for the log-rank test

► Example 10: Sample size using the Freedman and the Schoenfeld methods

Continuing with examples 1 and 2, Lachin (1981, 106) gives another approximation to obtain the estimate of the sample size under the equal-group allocation. This approximation coincides with the formula derived by Freedman (1982) for the number of events in the context of the log-rank test. We can obtain such an estimate by using power logrank and by specifying the hazard ratio of 0.66667 computed earlier.

```
. power logrank, hratio(0.66667) power(0.9) onesided
Estimated sample sizes for two-sample comparison of survivor functions
Log-rank test, Freedman method
Ho: HR = 1 versus Ha: HR < 1
Study parameters:
      alpha =    0.0500
      power =    0.9000
      delta =    0.6667 (hazard ratio)
      hratio =    0.6667
Censoring:
      Pr_E =    1.0000
Estimated number of events and sample sizes:
      E =      216
      N =      216
      N per group =    108
```

The estimate, 216, of the sample size is the same as given in Lachin (1981, 107) and is slightly smaller than the estimate, 218, obtained in example 1 and larger than the estimate, 210, obtained using the George–Desu method in example 2.

The approximation due to George and Desu (1974) is the same as the approximation to the number of events derived by Schoenfeld (1981) in application to the log-rank test. We can confirm that by typing

```
. power logrank, hratio(0.66667) power(0.9) onesided schoenfeld
Estimated sample sizes for two-sample comparison of survivor functions
Log-rank test, Schoenfeld method
Ho: ln(HR) = 0 versus Ha: ln(HR) < 0
Study parameters:
      alpha =    0.0500
      power =    0.9000
      delta =   -0.4055 (log hazard-ratio)
      hratio =    0.6667
Censoring:
      Pr_E =    1.0000
Estimated number of events and sample sizes:
      E =      210
      N =      210
      N per group =    105
```

We obtain the same estimate of 210 as when using `power exponential` with the `loghazard` option in [example 2](#).



Computing power

Sometimes the number of subjects available for enrollment into the study is limited. In such cases, the researchers may want to investigate with what power they can detect a desired treatment effect for a given sample size.

To compute power, you must specify the sample size in the `n()` option and an effect size. A hazard ratio of 0.5 is assumed if an effect size is not specified. Also see [Alternative ways of specifying effect](#) for various ways of specifying an effect size.

► Example 11: Power determination

We verify the power computation for the study from [example 9](#). We expect the power estimate to be close to 0.9.

The only thing we change in the `power exponential` command from [example 9](#) is replacing the `power(0.9)` option with the `n(664)` option.

```
. power exponential 0.8, t(10) n(664) aperiod(1) fperiod(9)
> loghazard unconditional
note: input parameters are survival probabilities
Estimated power for two-sample comparison of survivor functions
Exponential test, log hazard-ratio, unconditional
Ho: ln(HR) = 0 versus Ha: ln(HR) != 0
Study parameters:
      alpha =    0.0500
        N =     664
  N per group =    332
      delta =   -0.6931 (log hazard-ratio)
Accrual and follow-up information:
      duration =   10.0000
      follow-up =    9.0000
      accrual =    1.0000 (uniform)
Survival information:
      h1 =    0.0223          s1 =    0.8000
      h2 =    0.0112          s2 =    0.8944
      hratio =   0.5000          t =   10.0000
Estimated power:
      power =    0.9000
```

We obtain the estimate of power 0.9.



Testing hypotheses about two exponential survivor functions

► Example 12: Using streg to perform the log hazard-ratio test

In this example, we demonstrate the importance of sample-size computations to ensure a high power of a test to detect a difference between exponential survivor functions. We consider an asymptotic Wald (or normal z) test to test whether the log of the hazard ratio is zero.

Continuing with [example 11](#), suppose that the investigators have only 100 subjects available for the study. As we see below, the power to detect a 50% risk reduction in a hazard of the experimental group (the hazard ratio of 0.5) decreases from 90% to 24%:

```
. power exponential 0.8, t(10) n(100) aperiod(1) fperiod(9)
> loghazard unconditional
note: input parameters are survival probabilities

Estimated power for two-sample comparison of survivor functions
Exponential test, log hazard-ratio, unconditional
Ho: ln(HR) = 0 versus Ha: ln(HR) != 0

Study parameters:
      alpha =    0.0500
        N   =    100
N per group =     50
      delta =   -0.6931 (log hazard-ratio)

Accrual and follow-up information:
      duration = 10.0000
      follow-up =  9.0000
      accrual   =  1.0000 (uniform)

Survival information:
      h1 =    0.0223      s1 =    0.8000
      h2 =    0.0112      s2 =    0.8944
      hratio = 0.5000      t   = 10.0000

Estimated power:
      power =    0.2414
```

To demonstrate the implication of this reduction, consider the following example. We generate the data according to the study from [example 11](#) with the following code:

```
program simdata
  args n h1 h2 r
  set obs `n'
  generate double entry = `r'*runiform()
  generate double u = runiform()
  /* random allocation to two groups of equal sizes */
  generate double u1 = runiform()
  generate double u2 = runiform()
  sort u1 u2, stable
  generate byte drug = (_n<=`n'/2)
  /* exponential failure times with rates h1 and h2 */
  generate double failtime = entry - ln(1-u)/`h1' if drug==0
  replace failtime = entry - ln(1-u)/`h2' if drug==1
end

. clear
. set seed 234
. quietly simdata 100 0.0223 0.0112 1
```

The entry times of subjects are generated from a uniform $[0, 1)$ distribution and stored in variable `entry`. The subjects are randomized to two groups of equal size of 50 subjects each. The survival times are generated from exponential distribution with the hazard rate of $-\ln(0.8)/10 = 0.0223$ in the control group, `drug = 0`, and the hazard rate of $0.5 \times 0.0223 = 0.0112$ in the experimental group, `drug = 1`, conditional on subjects' entry times in `entry`.

Before analyzing these survival data, we need to set up the data properly using `stset`. The failure-time variable is `failtime`. The study terminates at $t = 10$, so we use `exit(time 10)` with `stset` to specify that all failure times past 10 are to be treated as censored. Because subjects enter the study at random times (`entry`) and become at risk of a failure upon entering the study, we also specify the `origin(entry)` option to ensure that the analysis time is adjusted for the entry times. For more details, see [ST] [stset](#).

```
. stset failtime, exit(time 10) origin(entry)
      failure event: (assumed to fail at time=failtime)
obs. time interval: (origin, failtime]
exit on or before: time 10
      t for analysis: (time-origin)
              origin: time entry
```

```
100 total observations
  0 exclusions
```

```
100 observations remaining, representing
  7 failures in single-record/single-failure data
921.825 total analysis time at risk and under observation
              at risk from t =          0
earliest observed entry t =          0
              last observed exit t = 9.990494
```

To perform the log hazard-ratio test, we fit an exponential regression model on `drug` by using `streg` (see [ST] [streg](#)). We can express the log of the hazard ratio in terms of regression coefficients as follows: $\ln(\Delta) = \ln(\lambda_2/\lambda_1) = \ln\{\exp(\beta_0 + \beta_1)/\exp(\beta_0)\} = \beta_1$, where β_0 and β_1 are the estimated coefficients for the constant and `drug` in the regression model. Then the test of $H_0: \ln(\Delta) = 0$ may be rewritten in terms of a coefficient on `drug` as $H_0: \beta_1 = 0$. This test is part of the standard output after `streg`.

```
. streg drug, distribution(exponential) nohr
      failure _d: 1 (meaning all fail)
      analysis time _t: (failtime-origin)
              origin: time entry
      exit on or before: time 10
Iteration 0: log likelihood = -29.718762
Iteration 1: log likelihood = -29.049311
Iteration 2: log likelihood = -29.014323
Iteration 3: log likelihood = -29.014222
Iteration 4: log likelihood = -29.014222
Exponential PH regression
No. of subjects =          100          Number of obs =          100
No. of failures =           7
Time at risk = 921.824925
Log likelihood = -29.014222          LR chi2(1) =          1.41
                                      Prob > chi2 =          0.2352
```

_t	Coef.	Std. Err.	z	P> z	[95% Conf. Interval]	
drug	.9428118	.83666	1.13	0.260	-.6970117	2.582635
_cons	-5.453234	.7071068	-7.71	0.000	-6.839137	-4.06733

From the output table above, the p -value for a two-sided test of the coefficient for drug, 0.260, is greater than 0.05. On that basis, we do not have evidence to reject the null hypothesis of no difference between the two exponential survivor functions. Therefore, we make an incorrect decision because we simulated the data with different group hazard rates. If we were to repeat this, say, 100 times, using different datasets simulated according to the alternative $H_a: \ln(\Delta) = \ln(0.5) = -0.6931$ (see [R] [simulate](#)), for roughly 76 of them we would have failed to reject the null hypothesis of no difference (a type II error). Therefore, more subjects are required to be able to detect the log of the hazard ratio of -0.4055 in this study.

◀

▶ Example 13: Using results from `streg` to perform the Wald test of hazard difference

We obtain the power of the test based on the difference between hazards for the study in [example 12](#) (omit the `loghazard` option from the syntax of `power exponential`).

```
. power exponential 0.8, t(10) n(100) aperiod(1) fperiod(9) unconditional
note: input parameters are survival probabilities
Estimated power for two-sample comparison of survivor functions
Exponential test, hazard difference, unconditional
Ho: h2 = h1 versus Ha: h2 != h1
Study parameters:
      alpha =    0.0500
          N =     100
  N per group =    50
      delta =   -0.0112 (hazard difference)
Accrual and follow-up information:
      duration =   10.0000
      follow-up =    9.0000
      accrual =    1.0000 (uniform)
Survival information:
      h1 =    0.0223          s1 =    0.8000
      h2 =    0.0112          s2 =    0.8944
      hratio = 0.5000          t =   10.0000
Estimated power:
      power =    0.2458
```

We obtain a power estimate of 0.2458, which is close to 0.2414 from [example 12](#).

To test the difference between hazard rates by using the Wald test, we express this difference in terms of coefficients, $\lambda_2 - \lambda_1 = \exp(\beta_0)\{\exp(\beta_1) - 1\}$, and we use `testnl` ([R] [testnl](#)) after `streg` to perform the nonlinear hypothesis test of $H_0: \exp(\beta_0)\{\exp(\beta_1) - 1\} = 0$.

```
. testnl exp(_b[_cons])*(exp(_b[drug])-1) = 0
      (1) exp(_b[_cons])*(exp(_b[drug])-1) = 0
           chi2(1) =          1.35
           Prob > chi2 =        0.2451
```

We obtain the same conclusions from the Wald test based on the difference between hazards as in [example 12](#). That is, based on the p -value of 0.2451, we fail to reject the null hypothesis of no difference between hazards of two groups (or miss the alternative $H_a: \psi = -0.0112$ corresponding to reduction in hazard from roughly 0.02 to 0.01) for the data from [example 12](#).

◀

Often in practice, to test the disparity in two exponential survivor functions, the log-rank test is used instead of the hazard-difference test. Also the Wald (or the score) test from the Cox model is used instead of the exponential log hazard-ratio test. Refer to [ST] [sts test](#) and [ST] [stcox](#) for examples on how to perform these tests (also see [PSS] [power logrank](#) and [PSS] [power cox](#)).

Sometimes the estimates of sample size and power obtained under the assumption of the exponential model are used as an approximation to the results used in a more general context of the log-rank test or the Cox proportional hazards model. Refer to [Lachin \(2011, 483–484\)](#) for the rationale behind this. Also see [Lakatos and Lan \(1992\)](#) for a discussion of the circumstances under which sample-size estimates obtained assuming the exponential model may be inaccurate when used with more general proportional hazards models.

Stored results

`power exponential` stores the following in `r()`:

Scalars

<code>r(alpha)</code>	significance level
<code>r(power)</code>	power
<code>r(beta)</code>	probability of a type II error
<code>r(delta)</code>	effect size
<code>r(N)</code>	total sample size
<code>r(N_a)</code>	actual sample size
<code>r(N1)</code>	sample size of the control group
<code>r(N2)</code>	sample size of the experimental group
<code>r(nratio)</code>	ratio of sample sizes, $N2/N1$
<code>r(nratio_a)</code>	actual ratio of sample sizes
<code>r(nfractional)</code>	1 if <code>nfractional</code> is specified, 0 otherwise
<code>r(onesided)</code>	1 for a one-sided test, 0 otherwise
<code>r(hratio)</code>	hazard ratio
<code>r(lnhratio)</code>	log hazard-ratio
<code>r(hdiff)</code>	difference between hazard rates
<code>r(h1)</code>	hazard in the control group (if specified)
<code>r(h2)</code>	hazard in the experimental group
<code>r(s1)</code>	survival probability in the control group (if specified)
<code>r(s2)</code>	survival probability in the experimental group (if specified)
<code>r(time)</code>	reference survival time (if <code>time()</code> is specified)
<code>r(aperiod)</code>	length of the accrual period (if specified)
<code>r(fperiod)</code>	length of the follow-up period (if specified)
<code>r(studytime)</code>	duration of the study (if specified)
<code>r(ashape)</code>	shape parameter (if <code>aperiod()</code> is specified)
<code>r(aprob)</code>	shape parameter (if <code>aproband()</code> is specified)
<code>r(aptime)</code>	proportion of accrual period (if <code>aptime()</code> is specified)
<code>r(ptime)</code>	reference accrual time (if <code>ptime()</code> is specified)
<code>r(losshaz)</code>	loss hazard rate in both groups (if specified)
<code>r(losshaz1)</code>	loss hazard in the control group (if specified)
<code>r(losshaz2)</code>	loss hazard in the experimental group (if specified)
<code>r(lossprob)</code>	proportions of subjects lost to follow-up in both groups (if <code>lossprob()</code> is specified)
<code>r(losstime)</code>	reference loss to follow-up time (if <code>losstime()</code> is specified)
<code>r(unconditional)</code>	1 if <code>unconditional</code> is specified, 0 otherwise
<code>r(separator)</code>	number of lines between separator lines in the table
<code>r(divider)</code>	1 if <code>divider</code> is requested in the table; 0 otherwise

Macros

<code>r(type)</code>	test
<code>r(method)</code>	exponential
<code>r(test)</code>	hazard difference or log-hazard difference
<code>r(accrual)</code>	uniform or exponential
<code>r(effect)</code>	ratio, <code>lnhratio</code> , <code>hdifference</code> , or <code>lnhdifference</code>
<code>r(columns)</code>	displayed table columns
<code>r(labels)</code>	table column labels
<code>r(widths)</code>	table column widths
<code>r(formats)</code>	table column formats

Matrices

<code>r(pss_table)</code>	table of results
<code>r(Pr_vec)</code>	1×4 matrix of probabilities of an event (when computed)
<code>r(Ea_vec)</code>	1×3 matrix of expected number of events under the alternative (when computed)
<code>r(EO_vec)</code>	1×3 matrix of expected number of events under the null (when computed)
<code>r(La_vec)</code>	1×3 matrix of expected number of losses under the alternative (when computed)
<code>r(LO_vec)</code>	1×3 matrix of expected number of losses under the null (when computed)

Methods and formulas

By default, `power exponential` computes the sample size required to achieve a specified power to detect a difference between hazard rates, $\psi_a = \lambda_{2a} - \lambda_{1a}$, using the method of [Lachin \(1981\)](#). If `loghazard` is specified, the sample size required to detect a log of the hazard ratio $\ln(\Delta_a) = \ln(\lambda_{2a}/\lambda_{1a})$ with specified power is reported using the formula derived by [George and Desu \(1974\)](#). In the presence of an accrual period, the methods of [Lachin and Foulkes \(1986\)](#) or (for uniform accrual only) [Rubinstein, Gail, and Santner \(1981\)](#) (if `loghazard` and `unconditional` are specified) are used.

In addition to the notation given in [Introduction](#), denote n , n_1 , and n_2 to be the total number of subjects required for the study, the number of subjects in the control group, and the number of subjects in the experimental group, respectively. Let $R = n_2/n_1$ denote the ratio of sample sizes of the experimental group to the control group. Let $p_1 = n_1/n = 1/(1 + R)$ and $p_2 = n_2/n = 1 - p_1 = R/(1 + R)$ be the proportions of subjects allocated to the control and the experimental groups; γ be the shape parameter of the truncated exponential distribution with p.d.f. $g(z) = \gamma \exp(-\gamma z) / \{1 - \exp(-\gamma r)\}$, $0 \leq z \leq r$, $\gamma \neq 0$; η_1 and η_2 be the loss hazards in the control and the experimental groups; and $z_{(1-\alpha/k)}$ and $z_{(1-\beta)}$ be the $(1 - \alpha/k)$ th and the $(1 - \beta)$ th quantiles of the standard normal distribution, with $k = 1$ for the one-sided test and $k = 2$ for the two-sided test. Denote $\bar{\lambda} = p_1\lambda_1 + p_2\lambda_2$. Recall that the difference between hazards is denoted by $\psi = \lambda_2 - \lambda_1$ and the hazard ratio is denoted by $\Delta = \lambda_2/\lambda_1$.

If survival probabilities $S_1(t)$ and $S_2(t)$ at a fixed time t are specified rather than hazard rates, the hazard rates are computed as $\lambda_i = -\ln\{S_i(t)\}/t$, $i = 1, 2$. If loss to follow-up probabilities $L_1(t_L)$ and $L_2(t_L)$ at a fixed time t_L are given instead of loss to follow-up hazard rates, the loss hazard rates are computed as $\eta_i = -\ln\{1 - L_i(t_L)\}/t_L$, $i = 1, 2$.

All formulas below are derived under the assumption of exponential survival distributions with hazard rates in the control and the experimental groups λ_1 and λ_2 , respectively, and rely on large-sample properties of the maximum likelihood estimates of λ_1 and λ_2 .

Denote $\xi_o = \zeta(\bar{\lambda}, \gamma, \eta_1)p_1^{-1} + \zeta(\bar{\lambda}, \gamma, \eta_2)p_2^{-1}$ and $\xi_a = \zeta(\lambda_1, \gamma, \eta_1)p_1^{-1} + \zeta(\lambda_2, \gamma, \eta_2)p_2^{-1}$.

The formula for the sample-size calculation using the conditional approach is

$$n = \frac{\left(z_{1-\alpha/k} \xi_o^{1/2} + z_{1-\beta} \xi_a^{1/2} \right)^2}{\delta^2}$$

and using the unconditional approach is

$$n = \frac{(z_{1-\alpha/k} + z_{1-\beta})^2 \xi_a}{\delta^2}$$

where $\zeta(\lambda, \gamma, \eta) = \lambda^2/p_E$ if $\delta = \psi$, $\zeta(\lambda, \gamma, \eta) = 1/p_E$ if $\delta = \ln(\Delta)$, and p_E is to be defined later. λ and η denote a failure hazard rate and a loss to follow-up hazard rate.

In the absence of censoring, the overall probability of an event (failure), p_E , is set to 1. Here the resulting formula for the sample size for the log hazard-ratio test depends only on the ratio of hazards and not on the individual group hazard rates. The resulting sample size formula for the test of the difference may also be rewritten as a function of the ratio of hazards only. Therefore, under no censoring, for a fixed value of the hazard ratio $\Delta = \lambda_2/\lambda_1$, the estimates of the sample size (or power) will be constant with respect to varying hazard rates λ_1 and λ_2 .

In the presence of censoring, when each subject is followed up for a fixed period $f = T$,

$$p_E = p_E(\lambda, \eta) = \frac{\lambda}{\lambda + \eta} [1 - \exp\{-(\lambda + \eta)T\}]$$

In the presence of an accrual period, the probability of an event is defined as

$$p_E = p_E(\lambda, \eta) = \frac{\lambda}{(\lambda + \eta)} \left[1 - \frac{\exp\{-(\lambda + \eta)(T - r)\} - \exp\{-(\lambda + \eta)T\}}{(\lambda + \eta)r} \right]$$

or

$$p_E = p_E(\lambda, \gamma, \eta) = \frac{\lambda}{(\lambda + \eta)} \left(1 + \frac{\gamma \exp\{-(\lambda + \eta)T\} [1 - \exp\{(\lambda + \eta - \gamma)r\}]}{(\lambda + \eta - \gamma) \{1 - \exp(-\gamma r)\}} \right)$$

under uniform or truncated exponential accrual with shape γ over $[0, r]$, respectively. Uniform accrual is assumed for $|\gamma| < 10^{-6}$.

The formulas are obtained from Lachin (1981), Lachin and Foulkes (1986), and Lakatos and Lan (1992). To avoid division by 0 in the case $\lambda + \eta = \gamma$, the probability of an event is taken to be the limit of the above expression, $p_E = \lim_{\lambda + \eta \rightarrow \gamma} p_E(\lambda, \gamma, \eta)$.

The number of subjects required to be recruited in each group is obtained as $n_1 = n/(1 + R)$ and $n_2 = nR/(1 + R)$. If `nfractional` is not specified, sample sizes are rounded to integer values; see *Fractional sample sizes* in [PSS] **unbalanced designs** for details.

The expected number of events and losses to follow-up are computed as suggested by Lachin and Foulkes (1986). Under the null hypothesis,

$$\begin{aligned} E_{H_0} &= n_1 p_E(\bar{\lambda}, \gamma, \eta_1) + n_2 p_E(\bar{\lambda}, \gamma, \eta_2) \\ L_{H_0} &= n_1 (\eta_1 / \bar{\lambda}) p_E(\bar{\lambda}, \gamma, \eta_1) + n_2 (\eta_2 / \bar{\lambda}) p_E(\bar{\lambda}, \gamma, \eta_2) \end{aligned}$$

and under the alternative hypothesis,

$$\begin{aligned} E_{H_a} &= n_1 p_E(\lambda_1, \gamma, \eta_1) + n_2 p_E(\lambda_2, \gamma, \eta_2) \\ L_{H_a} &= n_1 (\eta_1 / \lambda_1) p_E(\lambda_1, \gamma, \eta_1) + n_2 (\eta_2 / \lambda_2) p_E(\lambda_2, \gamma, \eta_2) \end{aligned}$$

For unconditional tests, the expected number of events and losses to follow-up under the null is computed by setting $\bar{\lambda} = \lambda_1$. The estimates of the expected number of events and losses to follow-up in each group are rounded to the nearest integer.

To obtain the estimate of the power, $1 - \beta$, the formulas for the sample size are solved for $z_{(1-\beta)}$ and the normal cumulative distribution function is used to obtain the corresponding probability $1 - \beta$.

To obtain the unknown shape parameter, γ , of a truncated exponential entry distribution, an iterative procedure is used to solve the equation

$$p_a = G(t_a) = \frac{1 - \exp(-\gamma t_a)}{1 - \exp(-\gamma r)}$$

for a given proportion of subjects p_a recruited at a given time, t_a , for $t_a \in [0, r]$.

References

- Chow, S.-C., J. Shao, and H. Wang. 2008. *Sample Size Calculations in Clinical Research*. 2nd ed. New York: Dekker.
- Cleves, M. A., W. W. Gould, and Y. V. Marchenko. 2016. *An Introduction to Survival Analysis Using Stata*. Rev. 3rd ed. College Station, TX: Stata Press.
- Freedman, L. S. 1982. Tables of the number of patients required in clinical trials using the logrank test. *Statistics in Medicine* 1: 121–129.
- George, S. L., and M. M. Desu. 1974. Planning the size and duration of a clinical trial studying the time to some critical event. *Journal of Chronic Diseases* 27: 15–24.
- Lachin, J. M. 1981. Introduction to sample size determination and power analysis for clinical trials. *Controlled Clinical Trials* 2: 93–113.
- . 2011. *Biostatistical Methods: The Assessment of Relative Risks*. 2nd ed. Hoboken, NJ: Wiley.
- Lachin, J. M., and M. A. Foulkes. 1986. Evaluation of sample size and power for analyses of survival with allowance for nonuniform patient entry, losses to follow-up, noncompliance, and stratification. *Biometrics* 42: 507–519.
- Lakatos, E., and K. K. G. Lan. 1992. A comparison of sample size methods for the logrank statistic. *Statistics in Medicine* 11: 179–191.
- Rubinstein, L. V., M. H. Gail, and T. J. Santner. 1981. Planning the duration of a comparative clinical trial with loss to follow-up and a period of continued observation. *Journal of Chronic Diseases* 34: 469–479.
- Schoenfeld, D. A. 1981. The asymptotic properties of nonparametric tests for comparing survival distributions. *Biometrika* 68: 316–319.

Also see [PSS] [intro](#) for more references.

Also see

- [PSS] [power](#) — Power and sample-size analysis for hypothesis tests
- [PSS] [power cox](#) — Power analysis for the Cox proportional hazards model
- [PSS] [power logrank](#) — Power analysis for the log-rank test
- [PSS] [power, graph](#) — Graph results from the power command
- [PSS] [power, table](#) — Produce table of results from the power command
- [PSS] [Glossary](#)
- [R] [test](#) — Test linear hypotheses after estimation
- [ST] [streg](#) — Parametric survival models