



City Research Online

City St George's, University of London

Citation: Gandrud, C. (2015). simPH: An R package for illustrating estimates from cox proportional hazard models including for interactive and nonlinear effects. *Journal of Statistical Software*, 65(3), pp. 1-20. doi: 10.18637/jss.v065.i03

This is the published version of the paper.

This version of the publication may differ from the final published version. To cite this item please consult the publisher's version.

Permanent repository link: <https://city-ref-test.eprints-hosting.org/id/eprint/12754/>

Link to published version: <https://doi.org/10.18637/jss.v065.i03>

Copyright and Reuse: Copyright and Moral Rights remain with the author(s) and/or copyright holders. Copies of full items can be used for personal research or study, educational, or not-for-profit purposes without prior permission or charge, unless otherwise indicated, provided that the authors, title and full bibliographic details are credited, a hyperlink and/or URL is given for the original metadata page and the content is not changed in any way. For full details of reuse please refer to [City Research Online policy](#).



simPH: An R Package for Illustrating Estimates from Cox Proportional Hazard Models Including for Interactive and Nonlinear Effects

Christopher Gandrud
Hertie School of Governance

Abstract

The R package **simPH** provides tools for effectively communicating results from Cox proportional hazard (PH) models, including models with interactive and nonlinear effects. The Cox (PH) model is a popular tool for examining event data. However, previously available computational tools have not made it easy to explore and communicate quantities of interest and associated uncertainty estimated from them. This is especially true when the effects are interactions or nonlinear transformations of continuous variables. These transformations are especially useful with Cox PH models because they can be employed to correctly specifying models that would otherwise violate the nonproportional hazards assumption. Package **simPH** makes it easy to simulate and then plot quantities of interest for a variety of effects estimated from Cox PH models including interactive effects, nonlinear effects, as well as standard linear effects. Package **simPH** employs visual weighting in order to effectively communicate estimation uncertainty. There are options to show either the standard central interval of the simulation's distribution or the shortest probability interval – which can be useful for asymmetrically distributed estimates. This paper uses hypothetical and empirical examples to illustrate package **simPH**'s syntax and capabilities.

Keywords: Cox proportional hazard models, hazard ratios, time interactions, time-varying, nonlinearity, splines, visual weighting, R.

1. Introduction

The [Cox \(1972\)](#) proportional hazards (PH) model is used in a wide range of disciplines, including epidemiology and political science, to study time to event data. However, many researchers poorly communicate quantities of interest and associated uncertainty estimated with these models. This is especially true when estimates are generated from interactions

and nonlinearly transformed continuous variables. These transformations can be especially useful with Cox PH models for correctly specifying models that would otherwise have non-proportional hazards, a key violation of the model’s assumptions. These effects are difficult to interpret using coefficient tables and available computational tools. In addition, estimation uncertainty around quantities of interest from Cox PH models is hard to assess because these quantities are on asymmetric and nonlinear scales.

This article aims to improve the use of Cox PH models. It first briefly discusses what a Cox PH model is and previous research on how time interactions and nonlinear effects can be helpful for correctly specifying models that would otherwise violate the proportional hazards assumption (PHA). Second, it advocates using shortest probability intervals and visual-weighting to display simulated quantities of interest from interactive, nonlinear, and even standard linear estimates from Cox PH models. Shortest probability intervals are often more appropriate for Cox PH model results. Third, it demonstrates how the new R (R Core Team 2015) package **simPH** (Gandrud 2015) makes it easy to implement these techniques. The package is freely available from the Comprehensive R Archive Network (CRAN) at <http://CRAN.R-project.org/package=simPH>. Package **simPH** can simulate and plot quantities of interest, including hazard ratios, first differences, relative hazards, marginal effects, and hazard rates from linear, linear multiplicative interactions between covariates, time interactions, and nonlinear transformations of continuous variables. Hypothetical and empirical examples are used to illustrate package **simPH**’s syntax and capabilities.

2. The Cox PH model

The Cox PH model is a semi-parametric survival model that allows us to examine how specified factors influence the rate of a particular event happening, e.g., infection, death, the adoption of a public policy, at a particular point in time given that the event has not already occurred. This rate is commonly referred to as the hazard rate ($h_i(t)$). The hazard rate for unit i at time t is estimated with the Cox PH model using:

$$h(t|\mathbf{X}_i) = h_0(t)e^{(\beta^\top \mathbf{X}_i)}, \quad (1)$$

where $h_0(t)$ is the baseline hazard, i.e., the instantaneous rate of the event of interest occurring at time t when all of the covariates are zero. β is a vector of coefficients and $\mathbf{X}_i$ is the vector of covariates for unit i .

We are often interested in how a covariate changes the rate of an event happening. In general researchers have tried to address this by looking at Cox PH coefficient estimates β . However, only examining single coefficients can lead to erroneous substantive interpretations of results. This is especially true when time interactions and nonlinearly transformed continuous variables are used to correct for violations of the PHA.

3. Violations of the proportional hazards assumption

Time interactions and nonlinear continuous variable transformations are particularly important when using Cox PH models. They can correct for violations of the PHA. The PHA is one of the most important sources of estimation bias in Cox PH models. It has been discussed at length by Licht (2011a), Box-Steffensmeier and Zorn (2001), and Box-Steffensmeier and Jones

(2004). The PHA is that the hazards of two units experiencing an event are proportional to one another and that this relationship is constant over time. Formally, for the PHA to hold the hazard for units j and l must be:

$$\frac{h_j(t)}{h_l(t)} = e^{\beta^\top (\mathbf{X}_j - \mathbf{X}_l)}. \quad (2)$$

for all points in time. This is also the equation for the hazard ratio between $\mathbf{X}_j$ and $\mathbf{X}_l$. If the PHA is violated and measures are not taken to correct for the violation, then researchers may create biased parameter estimates and statistical tests with lower power (Therneau, Grambsch, and Fleming 1990; Keele 2010). Beyond these statistical problems, not adjusting for violations of the PHA can prevent researchers from finding evidence for phenomena they are interested in studying, including how an effect changes over time and whether or not the effect of a continuous variable changes nonlinearly over the range of its values.

There are a number of widely used tests to examine if the PHA has been violated. See Grambsch and Therneau (1994), Box-Steffensmeier and Zorn (2001), and Box-Steffensmeier and Jones (2004) for discussions of various methods. Many software packages implement versions of these tests. R's `survival` package (Therneau 2015) implements Grambsch and Therneau's (1994) modified Schoenfeld residuals test with the `cox.zph` function.

3.1. Correcting nonproportional hazards with time interactions

If a covariate is determined to violate the PHA, Box-Steffensmeier and co-authors (see Box-Steffensmeier, Reiter, and Zorn 2003; Box-Steffensmeier and Jones 2004) suggest directly modeling the relationship between the variable and time. This usually entails including an interaction between the variable and some function of time, such as the natural logarithm or some exponent. The decision to use a particular functional form should be guided by theory and will likely also be influenced by findings in the data. If $f(t)$ is some function of time then a simple model estimating the hazard rate for unit i with one time interaction is given by:

$$h_i(t|x_i) = h_0(t)e^{(\beta_1 x_i + \beta_2 f(t)x_i)}. \quad (3)$$

Like any other interaction effect (see Brambor, Clark, and Golder 2006) extra care should be taken when interpreting the β_1 and β_2 parameter estimates and their associated uncertainty. We cannot simply interpret the effect by looking at β_1 or β_2 in isolation. They need to be combined. Licht (2011a) argues that post-estimation simulation techniques should be employed to substantively interpret these combined coefficients and the uncertainty surrounding them. Let's briefly look at ways to calculate combined effects. In the next section, we will look at showing our uncertainty about the combined effects using simulations.

Licht (2011a) describes two methods for calculating the combined effect of a time interaction on the hazard of an event happening in ways that are relatively easy to interpret: (a) first differences and (b) relative hazards. A first difference is the percentage change in the hazard rate at time t between two values of x :

$$\% \Delta h_i(t) = (e^{(x_j - x_l)(\beta_1 + \beta_2 f(t))} - 1) \cdot 100. \quad (4)$$

Relative hazards are given by:

$$\frac{h_j(t)}{h_l(t)} = e^{x_j(\beta_1 + \beta_2 f(t))}. \quad (5)$$

In this situation the covariate x_l is zero. Relative hazards represent the change in the hazard when x is “switched on”. As such, relative hazards are a special case of the hazard ratio (Licht 2011a, p. 231). They are the expected change in the hazard when x is fitted at a value different from zero compared to when x is zero.

3.2. Correcting nonproportional hazards with nonlinear transformations

Un-modeled time-interactive effects are not the only cause of PHA violations. Building on Grambsch and Therneau (1994) and Therneau and Grambsch (2000), Keele (2010) points out that common diagnostic tests will also indicate PHA violations if the model is misspecified for other reasons, such as omitting important covariates, using a proportional hazards model when another survival model is more appropriate, or including a continuous covariate as a linear component when its effect is actually nonlinear.

Keele suggests that *before* testing the PHA we should try to make sure that we are not omitting important variables (an issue beyond the scope of this paper) and find the covariates’ appropriate functional forms, typically using either polynomials or splines. He demonstrates this in replication studies by adding penalized splines and then using a Wald test to examine if the spline estimates have a better fit than their linear counterparts. Many studies using Cox PH models do not test for nonlinearity, but instead jump straight to testing the PHA, automatically including time interactions when it is violated. As Keele (2010) demonstrates, ascribing a time-interactive effect to a covariate when in fact the effect varies not over time, but nonlinearly over values of a continuous covariate can have major implications for substantive interpretations of results.

Because of their role in correcting for violations of the PHA, it is especially important to have accessible tools for exploring estimates and associated uncertainty from time interactions and nonlinear transformations of continuous variables.

4. Show estimation uncertainty

How can we effectively examine and communicate both the point estimates of and our uncertainty about quantities of interest from time-interactive and nonlinear effects? I advocate illustrating these measures using simulations, shortest probability intervals, and visually-weighted plots. These tools are useful for showing standard linear effects estimated from Cox PH models as well.

4.1. Post estimation simulations

Following King, Tomz, and Wittenberg (2000), Licht (2011a) proposes post-estimation simulation techniques to make it easier to estimate the uncertainty surrounding quantities of interest for time interactions like first differences and relative hazards. See King *et al.* (2000, pp. 352–353) for a discussion of alternative approaches including fully Bayesian Markov chain Monte Carlo techniques and bootstrapping. The main difference between these three approaches is how the parameters are drawn. Using the post-estimation simulation technique, we first find the parameter point estimates for $\hat{\beta}_1$ and $\hat{\beta}_2$. Second, we draw n values of β_1 and β_2 from multivariate normal distributions with means $\hat{\beta}_1$ and $\hat{\beta}_2$ and variances specified by the parameters’ estimated covariance. Third, we use these simulated values to calculate

a quantity of interest, such as the first difference or relative hazard, for a range of times as well as specified values of x_j and x_l (as appropriate). Finally, we plot the results. The simulation technique allows us to estimate full time-interactive effects, how they change over time, substantively evaluate the effects, and show uncertainty surrounding the estimates.

We can easily extend this simulation technique to quantities of interest for other estimated effect types. For example, if a nonlinear effect from a continuous variable is modeled with a second order polynomial, i.e., $\beta_1 x_i + \beta_2 x_i^2$, we can draw n simulations from the multivariate normal distribution for both β_1 and β_2 . Then we simply calculate quantities of interest for a range of values and plot the results as before. We find the first difference for a second order polynomial with:

$$\% \Delta h_i(t) = (e^{\beta_1 x_{j-l} + \beta_2 x_{j-l}^2} - 1) \cdot 100, \quad (6)$$

where $x_{j-l} = x_j - x_l$. Note we will not be showing the estimated effect over time. For this we need to estimate the hazard rate for a range of comparisons between x_j and x_l .

We can use a similar procedure for splines. Penalized splines (Eilers and Marx 1996) are a commonly used way of showing more complex nonlinear effects than polynomials (see Keele 2008). A B-spline basis for penalized spline estimation provides a means for linear combinations of spline basis functions on the complete domain of the basis, not only at the knots.

A Cox PH model with one penalized spline is given by:

$$h(t|x_i) = h_0(t)e^{g(x_i)}, \quad (7)$$

where $g(x)$ is the penalized spline function. For our post-estimation purposes $g(x)$ is a series of linearly combined coefficients such that:

$$g(x) = \beta_{k_1}(x)_{1+} + \beta_{k_2}(x)_{2+} + \beta_{k_3}(x)_{3+} + \dots + \beta_{k_n}(x)_{n+}, \quad (8)$$

where $k_1, \dots, k_n$ are the equally spaced spline knots with values inside of the range of observed x and n is the number of knots. x_{c+} indicates that:

$$(x)_{c+} = \begin{cases} x & \text{if } k_{c-1} < x \leq k_c \\ x & \text{if } x \leq k_1 \text{ and } k_c = k_1 \\ x & \text{if } x \geq k_n \text{ and } k_c = k_n \\ 0 & \text{otherwise.} \end{cases} \quad (9)$$

Note, x should be within the observed data.

We can again draw values of each $\beta_{k_1}, \dots, \beta_{k_n}$ from the multivariate normal distribution described above. We then use these simulated coefficients to estimate quantities of interest for a range of covariate values. For example, the first difference between two values x_j and x_l is:

$$\% \Delta h_i(t) = (e^{g(x_j - x_l)} - 1) \cdot 100. \quad (10)$$

Relative hazards and hazard ratios can be calculated by extension. Once we find the simulated quantities of interest, plotting the results is straightforward.

We can use this post-estimation simulation technique for virtually any quantity of interest estimated from Cox PH models including marginal effects for multiplicative interactions and plain linear effects.

4.2. Which interval to show?

If researchers go beyond usual “train timetable” coefficient tables and graphically show their parameter estimates, as package **simPH** makes easier, they usually do so by plotting lines of some measure of central tendency and confidence bands calculated from standard errors over a range of fitted values. Previous work with post-estimation simulations, e.g., [Licht \(2011a\)](#), has mirrored this approach in graphs with a line for the median or mean of the simulation distribution and lines representing the boundaries of a central interval of the distribution. For example, the central 95 percent interval could be represented by lines at the 2.5 and 97.5 percentiles of the distribution.

Many quantities of interest from Cox PH models have asymmetric probability distributions on a nonlinear scale and can therefore be poorly summarized by central intervals. Recall that most quantities of interest are on an exponential scale with a lower bound of 0 or, in the case of first differences, -100 . They can have very long and sparse upper regions relative to a tighter concentration of the distribution near the lower boundary. In these cases it can be more useful to look at highest density regions (see [Box and Tiao 1973](#); [Hyndman 1996](#)). The underlying idea is that we should be more interested in the set of quantities of interest values with the most probability ([Hyndman 1996](#), p. 120), rather than an arbitrary central interval. When the simulation has a normal distribution, the highest density region will be equivalent to the central interval with the same percentage of the simulations, e.g., 95 percent. However, when the highest density is at the boundary, for example when many of the simulated relative hazard values are close to 0, then [Liu, Gelman, and Zheng \(2013\)](#) argue that the highest density region is preferable to the central interval. In these cases “central intervals can be much longer and have the awkward property [of] cutting off a narrow high-posterior slice that happens to be near the boundary, thus ruling out a part of the distribution that is actually strongly supported by the inference” ([Liu et al. 2013](#), p. 2). If this happens [Liu et al.](#) recommend finding the shortest probability interval (SPIn). This is the shortest interval of a particular probability coverage based on the simulations. They find this to be a very efficient way of finding the shortest highest density region for unimodal distributions.

4.3. Visual weighting

Whether graphing a central or shortest probability interval, only using lines to represent the center and edges can draw the reader’s attention away from what the graph is trying to communicate. This approach overemphasizes the edges of the interval, the areas of lowest probability. Some graphs uniformly shade the interval between the upper and lower bounds. Uniform shading suggests to the reader a uniform distribution between the edges. Both of these characteristics give misleading information about the quantities of interest’s probability distributions, especially when they are on an exponential scale.

Visual weighting presents a solution to these problems. [Hsiang \(2013, p. 3\)](#) calls visual weight “the amount of a viewer’s attention that a graphical object or display region attracts, based on visual properties”. More visual weight can be created by using more “graphical ink” ([Tufte 2001](#)). Visual weight is decreased by removing graphical ink. The simplest way to increase or decrease graphical ink with our simulations is to simply plot each simulation value as a point or a series of values with a line that is semi-transparent. Areas of the distribution with many simulations will be darker. Areas with fewer simulations, often near the edges of the distribution, will be lighter. Plotting semi-transparent points or lines clearly communicates a

quantity of interest's probability distribution. When each point or line is partially transparent, areas of the chart where the points or lines are darker indicate areas of the distribution with higher probability, because more points or lines are stacked on top of one another. This approach gives more visual weight to areas of higher probability and avoids drawing the reader's attention to the edges of the distribution (the areas of lower probability) in the way that traditional confidence interval lines do.

As a practical issue if a plot shows very many simulations as individual points, and to a lesser extent lines, it may create a very large file size. This is especially true if higher quality vector graphics are desired. An alternative to summarize simulated distributions with multiple ribbons of increasing transparency the further from the central tendency a portion of the distribution is. These ribbons could stretch across given segments of the distribution such as the central 50 and 95 percentage intervals. Using ribbons rather than points conveys somewhat less information about a distribution, but can be convenient if very many simulations are plotted.

5. Package **simPH**: Tools for simulating and graphing effects

One reason that researchers have inconsistently incorporated suggestions to test for and examine time interactions and nonlinearities in their Cox PH models and show their estimation uncertainty is that there has been a lack of computational capabilities to easily do so. In R the **survival** package (Therneau 2015) has functions for testing the PHA. The **survival** and **Zelig** (Owen, Imai, King, and Lau 2013) packages can estimate models with splines and interactions. However, their capabilities for graphically showing these estimates and associated uncertainty are very limited. Package **Zelig** can plot basic estimates from Cox PH models using simulation techniques, but not time-interactive or nonlinear spline effects. Current tools for exploring spline estimates present results in difficult to interpret quantities. In general the capabilities for showing results from time interactions is very limited. Usually, showing these types of results requires considerable effort to extract estimates from model objects and then devise ways to show them graphically. See for example Licht's (2011b) Stata (StataCorp. 2009) code for replicating the time interaction plots in her paper. No previously available method allows you to easily plot shortest probability intervals and virtually none effectively uses visual weighting.

The **simPH** package for R aims to solve these problems. There are three basic steps to use the package:

1. Estimate a Cox PH model using **survival**'s `coxph` function.
2. Simulate parameter estimates and calculate quantities of interest, e.g., relative hazards, first differences, hazard ratios, marginal effects for linear interactions, or hazard rates, using the function from the **simPH** package corresponding to the effect type. See Table 1 for a summary of the simulation functions.
3. Plot the simulations using the `simGG` method.

simPH's simulation functions follow King *et al.* (2000) to simulate parameters and calculate a variety of quantities of interest. The user can specify the number of simulations to run per fitted value with the `nsim` argument. The default is 1,000. Warning: in some cases,

Simulation function	Effect type
<code>coxsimLinear</code>	linear
<code>coxsimInteract</code>	linear multiplicative interactions
<code>coxsimPoly</code>	polynomials
<code>coxsimtvc</code>	time-interactive
<code>coxsimSpline</code>	penalized splines

Table 1: **simPH** quantity of interest simulation functions.

especially with penalized splines, it is easy to ask the program to create more simulations than an average desktop computer can easily handle. The user may need to balance a desire for a clear view of the probability distribution a quantity of interest comes from with what is computationally feasible.

simPH's simulation functions allow the user to keep either the traditional central interval of the simulations' distributions or use the shortest probability interval (both are the 95 percent interval by default). In either case the range is set with the `ci` argument, e.g., to find the central 90 percent interval use `ci = 0.9`. To tell package **simPH** to find the shortest probability interval with any of **simPH**'s simulation functions simply set the argument `spin = TRUE`.¹

The **simGG** plotting method can then be used to plot these simulated values as semi-transparent points, lines, or ribbons. To choose between the plot types use the `type` argument and set it, simply enough, to `"points"`, `"lines"`, or `"ribbons"`, respectively. The default is `type = "points"`. It is important to note that while `"points"` and `"ribbons"` plot values in the selected interval for each value of the x -axis, `"lines"` plots simulations in the selected interval for all values plotted on the x -axis. This is to avoid creating jagged line plots, though it creates an interval with a slightly different interpretation.

The transparency level can be set with the argument `alpha`. If points or lines are used, `alpha` sets the transparency level for simulation values at the center of the distribution. Simulation values further from the center become more transparent the further out they are. If the user selects ribbon plots, three ribbons will be shown. The most transparent shows the furthest extent of the central or shortest probability interval. The less transparent ribbon shows the central 50 percent of this interval. And the middle line shows the interval's median.

A smoothing line of a type that can be specified by the user with the `method` argument is also plotted to summarize the distribution's central tendency. In general, any smoothing method accepted by **ggplot2** (Wickham 2009; Wickham and Chang 2015) can be used. **simGG** visually weights simulation distributions by plotting semi-transparent points for each simulated quantity of interest value or lines for the series of values from each simulation.

The **simGG** method draws on **ggplot2** to plot the simulations. In most cases **simGG** returns a **ggplot2** 'gg' object. For these plots, you can add any aesthetic attributes to the plots that **ggplot2** allows.

¹This capability was developed from Liu *et al.* (2013) and the accompanying code in the **SPIn** R package Liu (2013). It is currently unavailable for hazard rates.

6. Demonstrations

To illustrate package **simPH**'s syntax and capabilities we will first go through a simple example without interactive or nonlinear effects. Then we will replicate key figures and findings from [Licht \(2011a\)](#) and [Keele \(2010\)](#). The quantities of interest in the latter examples are from time-interactive and penalized spline effects, respectively. For examples of other quantities of interest and variable types please see package **simPH**'s documentation.

6.1. Simple non-interactive linear example

Package **simPH** can be used to show results from effects that are not explicitly modeled as interactions or with nonlinear transformations. To do this use the `coxsimLinear` function. Let's look at an example using a hypothetical data set called `hmohiv` from [Hosmer, Lemeshow, and May \(2008\)](#). The data is hosted by and the basic estimation model is from [UCLA Academic Technology Services \(2014\)](#). The data set contains 100 members of a Health Maintenance Organization (HMO). All of the members are HIV positive. The HMO wants to examine their survival times. Note: the model we will estimate below does not violate the PHA. This can be verified, for example, with the **survival** package's `cox.zph` function.

Let's estimate a simple Cox PH model that includes information on the HMO members' age and intravenous drug use history. We can load the necessary packages for this example and download the data with the following code:

```
R> library("survival")
R> library("simPH")
R> hmohiv <- read.table(
+   "http://www.ats.ucla.edu/stat/r/examples/asa/hmohiv.csv",
+   sep = ",", header = TRUE)
```

In the data set intravenous drug use is recorded as a binary variable called `drug`, where one is a history of drug use and zero otherwise. Each member's age in years is recorded in a variable called `age`. The ages range from 20 to 54 years with a median of 35.

We are first going to present results for how age is associated with survival times. Remember that quantities of interest such as relative hazards and hazard ratios are comparisons. Relative hazards are simply hazard ratios comparing a unit with a value of zero on some variable to another unit with another value of this variable. For the `age` variable this would mean a comparison between someone with an age of zero and someone with some other age. This may not be a particularly interesting comparison. In our case, it is also an out of sample comparison, as the youngest observed HMO member is 20 years old. We can easily create a much more substantively interesting comparison by making the reference value the median age (35):

```
R> hmohiv$AgeMed <- hmohiv$age - 35
```

The new variable ranges from -15 to 19 . Now all of the simulated relative hazards will be based on a comparison with a 35 year old. Let's estimate the model using the `coxph` function from the **survival** package:

```
R> M1 <- coxph(Surv(time, censor) ~ AgeMed + drug, method = "breslow",
+   data = hmohiv)
```




Figure 1: Simulated relative hazards of age on survival time for HIV positive HMO members. Default `simGG` plot (left) and plot with user-defined styling (right).

Age is clearly a time-varying quantity. Not only does it change over time in the way a person's income (probably) does, for example, but in effect it has a linear interaction with time. However, in our model we do not explicitly interact age with time as we will interact other variables with time in later examples. Because of this we can treat the variable the same way as personal income for the purposes of creating simulations with package **simPH** and so can use the `coxsimLinear` function:

```
R> Sim1 <- coxsimLinear(M1, b = "AgeMed", Xj = seq(-15, 19, by = 1))
```

We told `coxsimLinear` that we wanted to simulate relative hazards (the default quantity of interest) based on our model object `M1`. We specified which variable to find hazard ratios for using the `b` argument. The `Xj` argument sets the fitted values of x_j ; in this case a sequence of values between -15 and 19 at one unit intervals. The smaller the interval, the smoother the resulting graph will look. To graphically present the results now in the `Sim1` object we can use the `simGG` method. The following code creates the left panel of Figure 1:

```
R> simGG(Sim1)
```

We can set the x - and y -axis labels and make other aesthetic changes to create the right panel of Figure 1 using the following code:

```
R> simGG(Sim1, xlab = "\nYears of Age from the Sample Median (35)",
+   ylab = "Relative Hazard with Comparison\n to a 35 Year Old\n",
+   alpha = 0.05, type = "lines")
```

We can see in Figure 1 that the relative hazard at `AgeMed` zero (age 35) is one. A relative hazard for a unit at zero is always one, as it is a ratio of the hazards with itself. The simulated




Figure 2: Simulated relative hazards of drug use history on survival time for HIV positive HMO members.

relative hazards for ages below the median are less than one. This means that they are less likely to die at a given point in time than someone aged 35. Ages above the median have a relative hazard greater than one and so are more likely to die than 35 year olds. We can also see in this figure that we have estimated that 54 year olds (those 19 years older than the median) are about six times more likely than 35 year olds to die, all else equal.

We can use package **simPH** to find and plot quantities of interest for the binary **drug** use variable. For example, using the model object **M1** we simulate relative hazards for the variable's two levels: zero and one.

```
R> Sim2 <- coxsimLinear(M1, b = "drug", Xj = 0:1)
```

We can then use **simGG** (with some stylistic modifications) to create the plot in Figure 2. It is important to note that the **method** argument should be set to **method = "lm"** (i.e., linear model) when using binary variables. The nonlinear smoothers will not work with fewer than 10 *x*-axis values.

6.2. Time-interactive effects example

Let's now look at how to use package **simPH** to explore time-interactive effects that are explicitly created by interacting a given variable in the estimation model with some time transformation. To do this we will recreate plots from Licht (2011a). She re-examines Golub and Steunenberg's (2007) analysis of European Union legislative deliberation time. They wanted to find out what factors affected the amount of time it took the European Union to pass a new piece of legislation. Key variables they examined were the voting procedure that was used for a given piece of legislation, including the so-called qualified majority vote (QMV) procedure,² and the amount of other legislation pending, i.e., legislative backlog. Both of these variables violated the PHA and were more accurately modeled with log-time interactions.

²The procedure has been changed over time, but essentially QMV requires some voting majority greater than 50% + 1. See the European Union's website for more details: http://europa.eu/legislation_summaries/glossary/qualified_majority_en.htm (accessed January 2014).




Figure 3: Simulated first differences for the effect of qualified majority voting on the time it takes to pass legislation.

Figure 3 recreates Licht’s (2011a) figure showing the first difference of a log-time interaction for how the effect of the QMV procedure on legislative deliberation time changes as the number of days of deliberation increases (see Licht 2011a, p. 236).³ To create this figure we first load the data set `GolubEUPData`. This data set is included in package **simPH**.

```
R> data("GolubEUPData", package = "simPH")
```

Before creating the log-time interactions, let’s look at the data’s format:

```
R> head(GolubEUPData[, 2:5])
```

	caseno	begin	end	event
1	1	0	595	1
2	2	0	1246	1
3	4	0	341	1
4	5	0	335	1
5	6	0	522	1
6	7	0	1664	0

We can see that each piece of prospective legislation is identified with a case number in the `caseno` variable. The `event` variable records if the event of interest occurred (i.e., the legislation was passed) or not. Observation intervals `begin` at 0 and extend to the `end` when either the event occurs, *or* one of the covariate values change.

³Her original figure was not in terms of a percentage difference to make it more comparable to a figure in Golub and Steunenberg’s original. The results are presented in terms of percentage difference, as the first difference is commonly reported. The pre and post Single European Act time periods are not separated out for simplicity. Finally, I also examined whether or not nonlinear functional forms would be better fits than time-interactions as per our discussion above. However, no evidence of this was found.

If we create a time interaction by using the `end` variable as the time value, we will inaccurately measure the interaction for failures before this time. Imagine we build an interaction between `end` time and a dichotomous variable with the value 1 for the second unit in our data set. The interaction term would be 1,246. This is accurate at the second unit's event time t_m . However, it is inaccurate at all times t_n , where $t_n < t_m$. When another unit experiences an event at a time t_n before t_m we should use the second unit's time interaction value at t_n to find the partial likelihood, rather than the value at t_m .

To solve this problem, package **simPH** includes the `SurvExpand` function. It expands a data set out into equally spaced intervals. It keeps all intervals that end at an observed event time and when a covariate changes. These are the only time points that are needed to find the partial likelihood. Removing unneeded periods helps save memory. The resulting data frame allows us to create accurate time interactions. Here is how we use the function with the `GolubEUPData` data:

```
R> GolubEUPData <- SurvExpand(GolubEUPData, GroupVar = "caseno",
+   Time = "begin", Time2 = "end", event = "event")
R> head(GolubEUPData[, 1:4])
```

	caseno	begin	end	event
1	1	0	1	0
2	1	1	2	0
3	1	4	5	0
4	1	6	7	0
5	1	8	9	0
6	1	10	11	0

Now we can create the log-time interactions. The two time-interaction variables we will focus on are for qualified majority voting (`qmv`) and legislative backlog (`backlog`). Other time-interaction variables are also from the original model. We can use the `tv` function to create log-time interactions. The function is included with package **simPH**. The covariate we wish to interact is specified with `b` and the time variable with `tvar`. Finally, we specify the function of time with the `tfun` argument.

```
R> BaseVars <- c("qmv", "backlog", "coop", "codec", "qmvpostsea", "thatcher")
R> GolubEUPData <- tv(GolubEUPData, b = BaseVars, tvar = "end",
+   tfun = "log")
R> names(GolubEUPData)[18:23]
```

```
[1] "qmv_log"          "backlog_log"      "coop_log"
[4] "codec_log"        "qmvpostsea_log"   "thatcher_log"
```

Notice that the time interaction variables are added to the original `GolubEUPData` using the base variables' names plus a suffix denoting the log transformation.

Now estimate the model with `coxph`. Note that each log-time interaction is entered into the model along with the corresponding non-time interacted term.

```
R> M2 <- coxph(Surv(begin, end, event) ~
+   qmv + qmvpostsea + qmvpostteu + coop + codec + eu9 + eu10 + eu12 +
+   eu15 + thatcher + agenda + backlog + qmv_log + qmvpostsea_log +
+   coop_log + codec_log + thatcher_log + backlog_log,
+   data = GolubEUPData, ties = "efron")
```

In the following code chunk we take the `M2` model object created by `coxph` and use it in **simPH**'s `coxsimtvc` function to simulate the first difference of the time-interactive effect of qualified majority voting on directive deliberation time. We create the simulations for 80 through to 2000 days after deliberation begins. We specify the quantity of interest that we want to simulate with the `qi` argument. The non-time interacted term is entered with the `b` argument. The time interacted term is entered with `btvc`. The form of the time interaction is declared with the `tfun` argument, e.g., `tfun = "log"`.⁴ `Xj = 1` specifies that the difference is between values of `qmv` one and zero. This is a comparison between using QMV and using some other voting procedure. We could change the x_l value with the `Xl` argument, but this is not relevant for binary variables. The `from` and `to` arguments allow us to specify the time period over which to simulate the quantity of interest. The `by` argument allows us to specify the increment of the time sequence to simulate values at.

```
R> Sim3 <- coxsimtvc(obj = M2, b = "qmv", btvc = "qmv_log",
+   qi = "First Difference", Xj = 1, tfun = "log", from = 80, to = 2000,
+   by = 10)
```

Once we create the `Sim3` object containing the simulations, we can simply use the `simGG` method to plot the results. In this example, we adjust the median line size with the `lsize = 0.5` argument. This makes this particular plot clearer. The `legend = FALSE` argument hides the legend describing the fitted value comparisons that are being made between x_j and x_l . A legend is uninformative in this case as we are only comparing the values one and zero. The default is to show the legend. The `legend` argument follows the **ggplot2** syntax for creating plot guides.

```
R> simGG(Sim3, xlab = "\nTime in Days", type = "ribbons", lsize = 0.5,
+   legend = FALSE, alpha = 0.3)
```

We can clearly see in Figure 3⁵ that QMV increases the probability of passing a directive early in the deliberation process (almost by 250 percent at about 80 days). But as the deliberation time increases, the effect decreases and then becomes negative at around 1,000 days.

Let's look at another example, this time with multiple comparisons plotted in one figure. Figure 4⁶ shows the effect of different levels of legislative `backlog` on directive deliberation

⁴Other options include `"linear"` for linear time interactions and `"pow"` for polynomials. If `tfun = "pow"` then also set the argument `pow` to specify the power the time interaction was raised to.

⁵The figure's ribbons extend across the central probability 95 percent interval of the simulations. As in Licht's (2011a) original, the time period plotted is truncated from 80 to 2000 to make the estimates more easily interpretable.

⁶It replicates the right panel of Licht's Figure 3 (2011a, p. 237). One difference is that she estimated uncertainty from 10 draws of 1,000 simulations, whereas Figure 4 is based on one draw of 200 simulations. This reduced computation time and leads to substantively similar conclusions. The figure's ribbons extend across the middle 95 percent of the simulations.




Figure 4: Simulated relative hazards for the effect of different levels of legislative backlog on directive deliberation time.

time from 1,200 to 7,000 days after the directive was proposed. The effect shown is also modeled as a log-time interaction. The process of creating the plot is similar to what we have already seen. The only differences to note are that we entered a sequence of values for backlogged legislation for X_j in `coxsimtvc` ranging from 40 to 200 at increments of 40. This creates five separate sets of relative hazard simulations, one for each value of X_j compared to zero. Finally, we specified the legend name for the plot in the `simGG` call with the `leg.name` argument.

```
R> Sim4 <- coxsimtvc(obj = M2, b = "backlog", btvc = "backlog_log",
+   qi = "Relative Hazard", Xj = seq(40, 200, 40), tfun = "log",
+   from = 1200, to = 7000, by = 100, nsim = 200)
R> simGG(Sim4, xlab = "\nTime in Days", type = "ribbons",
+   leg.name = "Backlogged \n Items")
```

The main conclusion we can draw from Figure 4 is that if a piece of legislation is not passed in the first 1,200 or so days from when it was proposed, it is less likely that it will be passed the larger the legislative backlog is (for more details see [Licht 2011a](#), pp. 236–237).

In the previous examples, the simulation probability distributions are not strongly influenced by boundary effects or long upper tails. Because of this the central and shortest probability intervals (not shown) are largely equivalent. The distributions are also relatively tight, which is indicated by fairly even visual weight across the distributions.

6.3. Nonlinear effects: Penalized spline example

Let's now turn to look at how package **simPH** can be used to simulate and plot quantities of interest estimated from penalized splines. To do this we will build on [Keele \(2010\)](#). He demonstrated that we need to check for and explicitly model nonlinearity with a replication of [Carpenter \(2002\)](#). [Carpenter \(2002\)](#) examined the time it took the US Food and Drug

Administration (FDA) to approve a new drug. Using the steps discussed above, he found that modeling nonlinearity with penalized splines, rather than time-interactive effects was the more appropriate strategy for dealing with the covariates that violated the PHA. This allowed him to draw conclusions that, for example, the number of FDA staff reviewing a drug's application increases the likelihood that it will be accepted, but that the effect diminishes after a threshold number of staff are assigned.

It has been difficult to examine and communicate the functional form, magnitude, and uncertainty surrounding spline effects. Coefficient tables are very cumbersome, because a spline fitted effect is estimated using multiple coefficients, each for a different range of a covariate's values. Keele does not show spline coefficients in his results tables. Instead he simply denotes their overall statistical significance with standard significance stars.

In R you can plot estimated spline effects over a range of values with the `termplot` function. These plots, however, have a number of drawbacks. First, the plots show the log hazard ratio, which is not a particularly intuitive quantity to understand and is rarely reported in studies using Cox PH models. More importantly it attempts to communicate uncertainty by plotting standard errors, instead of more widely used confidence intervals. As such, casual readers could easily think the uncertainty around the spline estimates is smaller than it really is.

Package **simPH** allows us to estimate quantities that we are more interested in like relative hazards, hazard rates over time, hazard ratios, and first differences. For example, let's simulate the hazard ratios for the effect of an additional drug review staff on the time it takes for a drug to be approved. The plots in Figure 5 show the simulated hazard ratios over the full range of FDA staff per drug trial (`stafcdcr`) observed in Carpenter's data.

Let's look at the syntax used to create the left panel of Figure 5. Load Carpenter's data included with package **simPH** and estimate the model with splines using `coxph` and `pspline`. The model was originally used to create the results in Keele (2010, Table 7).

```
R> data("CarpenterFdaData", package = "simPH")
R> M3 <- coxph(Surv(acttime, censor) ~ prevgenx + lethal + deathrt1 +
+   acutediz + hosp01 + pspline(hospdisc, df = 4) +
+   pspline(hhosleng, df = 4) + mandiz01 + femdiz01 + peddiz01 +
+   orphdum + natreg + vandavg3 + wpnoavg3 + pspline(condavg3, df = 4) +
+   pspline(orderent, df = 4) + pspline(stafcdcr, df = 4),
+   data = CarpenterFdaData)
```

Now simulate the hazard ratios for a sequence of values with `coxsimSpline` and graph the results with `simGG`.

```
R> XjFit <- seq(1100, 1700, by = 10)
R> XlFit <- setXl(Xj = XjFit, diff = 1)
R> Sim5 <- coxsimSpline(M3, bspline = "pspline(stafcdcr, df = 4)",
+   bdata = CarpenterFdaData$stafcdcr, qi = "Hazard Ratio",
+   Xj = XjFit, Xl = XlFit)
R> simGG(Sim5, xlab = "\n No. of FDA Drug Review Staff",
+   title = "Central Interval\n", alpha = 0.1, type = "lines")
```

There are a few syntax points to note. First, we created our vector of x_l values using the `setXl` function. This simply takes a vector of x_j values and returns a vector of corresponding




Figure 5: Simulated hazard ratios for the effect of FDA staff on drug approval time.

x_i values that are different from x_j by an amount set with the `diff` argument. Second, we told `coxsimSpline` the term to estimate for using the `bspline` argument and the same syntax we used to enter the term as a penalized spline in the `coxph` call, i.e., `pspline(stafcdcr, df = 4)`. Third, we specified the vector containing the observed `stafcdcr` data with the `bdata` argument. This is important for `coxsimSpline` to be able to accurately find the spline knots. The simulation values in Figure 5 have been smoothed using post-simulation cubic smoothing splines. This is `simGG`'s default behavior when plotting splines. Smoothing the simulations gives a more continuous appearance. It is accomplished with `smooth.spline`, which is available in package `stats` in R. You can tell `simGG` to not smooth the simulations by setting the argument `SmoothSpline = FALSE`.

We can see in the left panel of Figure 5 that the simulated values are concentrated near the bottom of the distribution. Following Liu *et al.* we may find it useful to show not the central 95 percent interval, but the shortest 95 percent probability interval. The right panel of Figure 5 shows this interval. To create it we used all of the same syntax as before, while simply adding `spin = TRUE` to `coxsimSpline`. Using the shortest 95 percent probability interval indicates that there is actually a higher probability that the hazard ratio is 1, i.e., no effect, than the central 95 percent interval for all but the highest quartile or so of FDA drug review staff. It also deemphasizes the higher hazard ratio values, where there is a low concentration of the probability. In both of Figure 5's panels the visual weighting draws readers' attention to the lower part of the distribution where the model estimates there is a higher probability that the hazard ratio is located.

Finally, to ease comparison, the right panel plot is on the same y -axis scale as the first. To illustrate how to do this, imagine that we placed the output of `coxsimSpline` with the argument `spin = TRUE` into an object called `SimPlot`. We then use `ggplot2` to change the y -axis range so that it is comparable the left panel of Figure 5:

```
R> library("ggplot2")
R> SimPlot + scale_y_continuous(breaks = c(0, 20, 40, 60), limits = c(0, 60))
```

By simulating quantities of interest and plotting the results we can more adequately explore the estimated spline effects than we could have with previously available tools. In so doing we bring into question the substantive significance of Keele’s (2010) findings. Remember that Figure 5 plots hazard ratios between an x_j value that is only one unit, i.e., one FDA review staff member, greater than the x_l value. Given that the observed drug review staff variable `stafcdcr` ranges from 996 to 1796, a one unit change is not really substantively meaningful. However, when we create simulations using only slightly larger differences between x_j and x_l , e.g., `setXl(Xj = XjFit, diff = 5)`, we end up with very large (often infinite) simulated quantity of interest values. Package **simPH** automatically reports and drops extreme values when running `coxsim*` functions. This avoids crashes that may occur when trying to find the distributions’ middle intervals or during plotting. The user can set the argument `extremesDrop = FALSE` in the simulation function call to keep extreme values if they wish to explore them in more detail.

7. Conclusion

Exploring and communicating quantities of interest and their associated uncertainty estimated from Cox PH models presents a number of difficulties. Quantities of interest are on nonlinear and asymmetric scales. Key types of covariate transformations that help address violations of the PHA are especially difficult to interpret using coefficient tables and other previously available tools. The **Zelig** package for R comes closest to allowing us to estimate and communicate uncertainty from Cox (PH) models. However, it has limited or no ability to simulate quantities of interest for time-interactive and nonlinear estimated effects. So, this paper has demonstrated a new R package, **simPH**, that makes it easy to effectively explore and present quantities of interest for interactive and nonlinear effects. The package can also be used for standard linear effects, making it a useful all-round tool for showing results from Cox PH models. This paper has also argued for the use of visual weighting and shortest probability intervals for understanding results from these models. Package **simPH** fully supports these methods.

Acknowledgments

Thank you to Andreas Beger, Jeffrey Chwioroth, Kelly Kadera, Luke Keele, Mintao Nie, Alison Post, Meredith Wilf, participants of the International Studies Association Annual Convention (2013), two anonymous reviewers, and the editors.

References

- Box GEP, Tiao GC (1973). *Bayesian Inference in Statistical Analysis*. John Wiley & Sons, New York.
- Box-Steffensmeier JM, Jones BS (2004). *Event History Modeling: A Guide for Social Scientists*. Cambridge University Press, Cambridge.
- Box-Steffensmeier JM, Reiter D, Zorn CJ (2003). “Nonproportional Hazards and Event History Analysis in International Relations.” *Journal of Conflict Resolution*, **47**(1), 33–53.

- Box-Steffensmeier JM, Zorn CJ (2001). “Duration Models and Proportional Hazards in Political Science.” *American Journal of Political Science*, **45**(4), 972–988.
- Brambor T, Clark WR, Golder M (2006). “Understanding Interaction Models: Improving Empirical Analyses.” *Political Analysis*, **14**(1), 63–82.
- Carpenter DP (2002). “Groups, the Media, Agency Waiting Costs, and FDA Drug Approval.” *American Journal of Political Science*, **46**(3), 490–505.
- Cox DR (1972). “Regression Models and Life Tables.” *Journal of the Royal Statistical Society B*, **34**(2), 187–220.
- Eilers PHC, Marx BD (1996). “Flexible Smoothing with B-Splines and Penalties.” *Statistical Science*, **11**(2), 89–102.
- Gandrud C (2015). *simPH: Tools for Simulating and Plotting Quantities of Interest Estimated From Cox Proportional Hazards Models*. R package version 1.3.1, URL <http://CRAN.R-project.org/package=simPH>.
- Golub J, Steunenberg B (2007). “How Time Affects EU Decision-Making.” *European Union Politics*, **8**(4), 555–566.
- Grambsch PM, Therneau TM (1994). “Proportional Hazards Tests and Diagnostics Based on Weighted Residuals.” *Biometrika*, **81**(3), 515–526.
- Hosmer David W J, Lemeshow S, May S (2008). *Applied Survival Analysis: Regression Modeling of Time to Event Data*, volume 618. 2nd edition. John Wiley & Sons.
- Hsiang SM (2013). “Visually-Weighted Regression.” Available at SSRN, URL <http://ssrn.com/abstract=2265501>.
- Hyndman RJ (1996). “Computing and Graphing Highest Density Regions.” *The American Statistician*, **50**(2), 120–126.
- Keele L (2008). *Semiparametric Regression for the Social Sciences*. John Wiley & Sons, Chichester.
- Keele L (2010). “Proportionally Difficult: Testing for Nonproportional Hazards in Cox Models.” *Political Analysis*, **18**(2), 189–205.
- King G, Tomz M, Wittenberg J (2000). “Making the Most of Statistical Analyses: Improving Interpretation and Presentation.” *American Journal of Political Science*, **44**(2), 347–361.
- Licht AA (2011a). “Change Comes with Time: Substantive Interpretation of Nonproportional Hazards in Event History Analysis.” *Political Analysis*, **19**(2), 227–243.
- Licht AA (2011b). *Replication Data for: Change Comes with Time*. IQSS Dataverse Network [Distributor] V3 [Version], URL <http://hdl.handle.net/1902.1/15633>.
- Liu Y (2013). *SPIn: Simulation-Efficient Shortest Probability Intervals*. R package version 1.1, URL <http://CRAN.R-project.org/package=SPIn>.

- Liu Y, Gelman A, Zheng T (2013). “Simulation-Efficient Shortest Probability Intervals.” arXiv:1302.2142 [stat.CO], URL <http://arxiv.org/abs/1302.2142>.
- Owen M, Imai K, King G, Lau O (2013). *Zelig: Everyone’s Statistical Software*. R package version 4.2-1, URL <http://CRAN.R-project.org/package=Zelig>.
- R Core Team (2015). *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria. URL <http://www.R-project.org/>.
- StataCorp (2009). *Stata Statistical Software: Release 11*. StataCorp LP, College Station. URL <http://www.stata.com/>.
- Therneau T (2015). *survival: Survival Analysis*. R package version 2.38-1, URL <http://CRAN.R-project.org/package=survival>.
- Therneau TM, Grambsch PM (2000). *Modeling Survival Data: Extending the Cox Model*. Statistics for Biology and Health. Springer-Verlag.
- Therneau TM, Grambsch PM, Fleming TR (1990). “Martingale-Based Residuals for Survival Models.” *Biometrika*, **77**(1), 147–160.
- Tufte ER (2001). *The Visual Display of Quantitative Information*. 2nd edition. Graphics Press, Cheshire.
- UCLA Academic Technology Services (2014). *R Textbook Examples: Applied Survival Analysis, by Hosmer and Lemeshow*. URL http://www.ats.ucla.edu/stat/r/examples/asa/asa_ch4_r.htm.
- Wickham H (2009). *ggplot2: Elegant Graphics for Data Analysis*. Springer-Verlag, New York. URL <http://had.co.nz/ggplot2/book>.
- Wickham H, Chang W (2015). *ggplot2: An Implementation of the Grammar of Graphics*. R package version 1.0.1, URL <http://CRAN.R-project.org/package=ggplot2>.

Affiliation:

Christopher Gandrud
Hertie School of Governance
Friedrichstrasse 180
Berlin, 10117, Germany
E-mail: gandrud@hertie-school.org