Chapter 4

Application to the Diabetic Retinopathy Study

In this chapter we apply the log hazard regression model to data collected to study the effect of laser treatment on diabetic retinopathy (see Diabetic Retinopathy Study Research Group, 1981). We give a description of this study in Section 4.1, and analyze the data in Section 4.2. Finally, in Section 4.3, we discuss analysis of this data set using other models.

4.1 Data Description

Diabetic retinopathy is a complication associated with diabetes mellitus, which consists of abnormalities in the microvasculature within the retina of the eye. It is the major cause of visual loss in many industrialized countries (Murphy and Patz, 1978).

The Diabetic Retinopathy Study (DRS) was funded by the National Eye Institute in 1971 to investigate the effectiveness of laser photocoagulation in delaying the onset of blindness for diabetic retinopathy patients. One eye of each patient was randomly chosen to receive photocoagulation and the other eye was observed without treatment.

A total of 1,742 patients was followed over several years. The endpoint used to assess the treatment effect was the occurrence of visual acuity less than 5/200 at two consecutively completed 4-month follow-ups. The only data available are from the 197 patients defined as high-risk by DRS criteria.

Of the 197 pairs of observations, approximately one-half (101/197) of the untreated eyes and one-quarter (54/197) of the treated eyes achieved the outcome after 5 years of follow-up. The histograms of the censored and uncensored data for treatment and control (Figure 4.1) show that more untreated eyes than treated eyes failed during the study and many patients left the study after about 3 years (36 months) or more from the start time.

The correlation between the uncensored observations of the treatment group and the control group is 0.28. This indicates possible dependence between the two failure times. We show the scatter plots in Figure 4.2. For this data set, the two censoring times are identical, that is, $C_{Ti} = C_{Ci}$.

The primary goal of the DRS study was to assess the effectiveness of the laser photocoagulation treatment. A secondary goal was to assess the relative risk of blindness of the untreated and the treated eyes as a function of time.

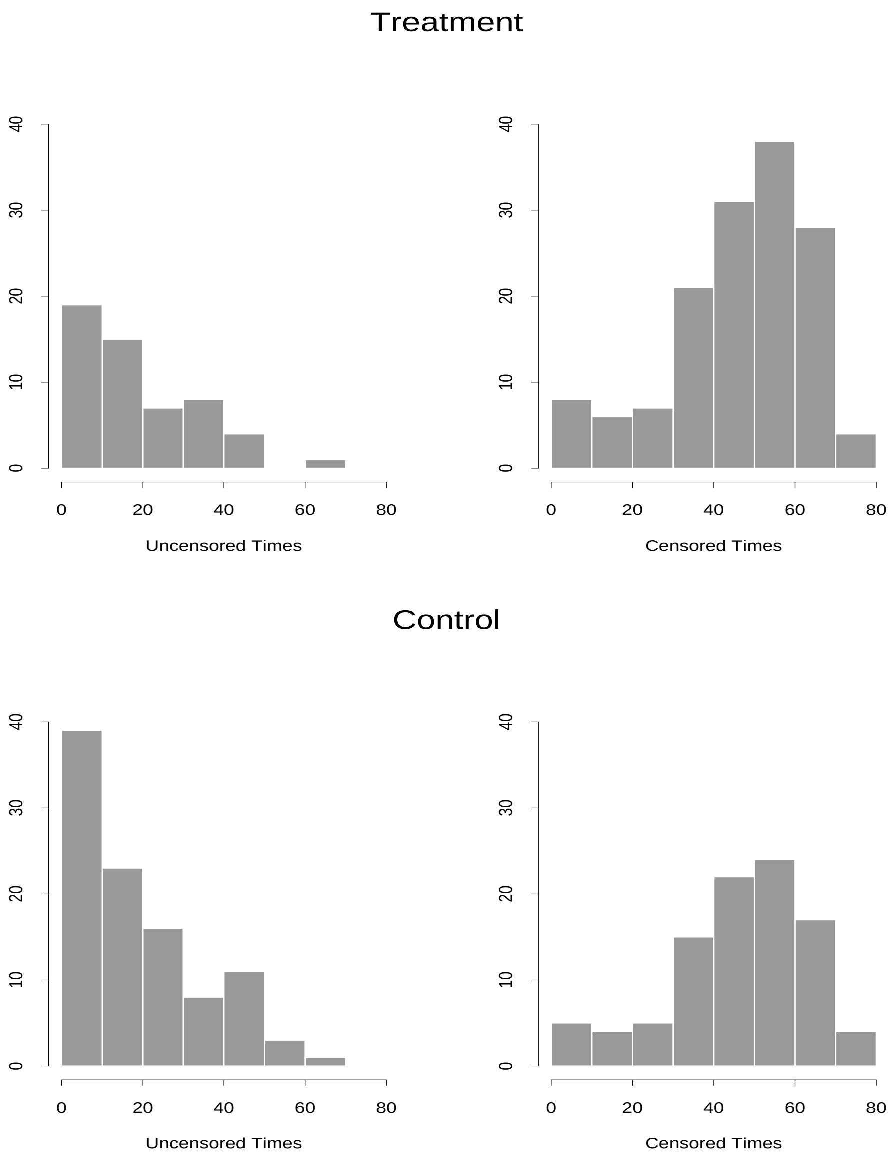

Figure 4.1: Histograms of observed times of the eye data.

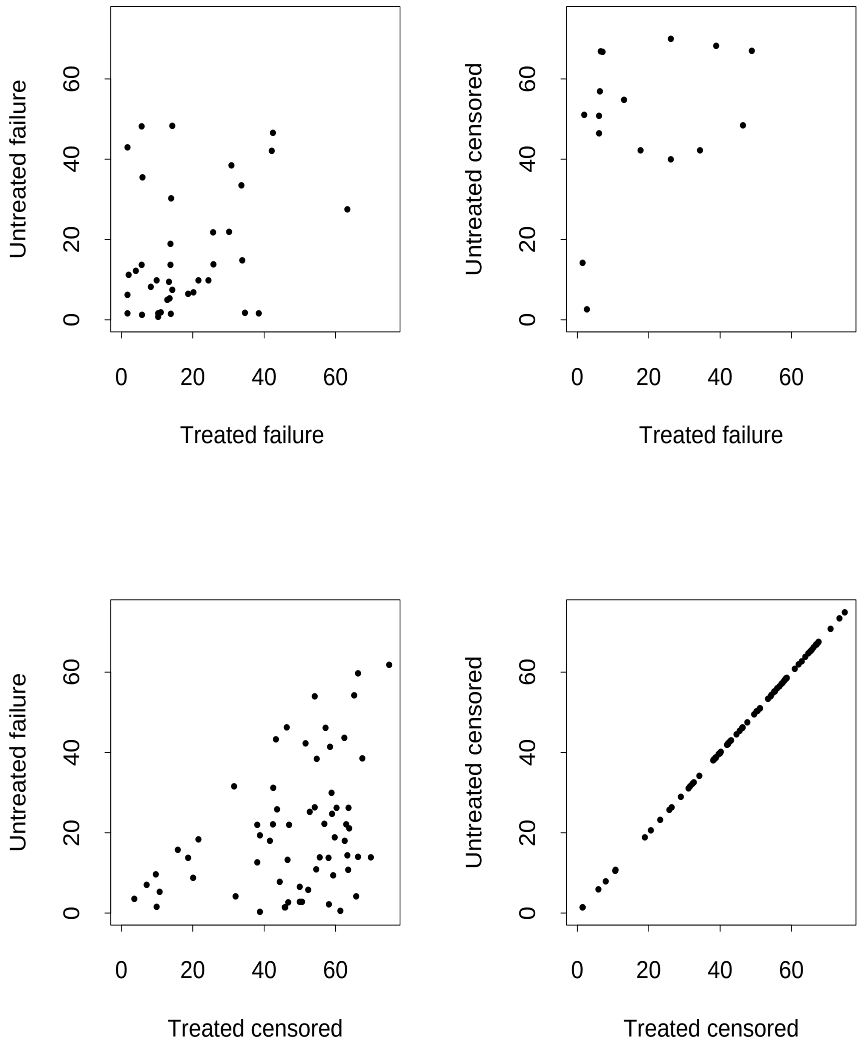

Figure 4.2: Scatter plots of the observations of the eye data.

4.2 Data Analysis

In this section we will address the following questions:

1. Do the failure times of the treated and the untreated eyes have the same distribution? That is, is there a treatment effect?

2. What is the log hazards ratio? That is, what is the relative risk of blindness of the untreated eyes and treated eyes?

We carry out two analyses of the eye data. First, as described in Section 3.4 on knot selection, we let Kooperberg's *heft* algorithm choose the knots for the log hazard regression model (Model 1). This allows different knots to be used for the log hazards estimates of the failure times of the treated and the untreated eyes. In our second analysis (Model 2), we use the same regression space for the log hazards of the two failure times. The set of knots used is the union of the knots chosen by the *heft* algorithm for each failure time. With both models we can get an estimate of the log hazards ratio with 95% pointwise conference intervals. We then test if there is a treatment effect, and if the two hazards are proportional with Model 2, as described in Section 3.5.

4.2.1 Model 1

To estimate each failure time's log hazard, we use Kooperberg's *heft* algorithm with log tail $\beta_p(t) \equiv \log(t)$ $(c = 0)$, and without specifying knots. The *heft* algorithm selects $1.5, 6.17$, and 63.33 as knots for the regression model of the treated eyes and chooses no knots for the untreated eyes, which means that *heft* chooses a Weibull model for the failure time of the untreated eye. Then the log hazard regression model, Model 1, is:

$$\alpha_T(t|\boldsymbol{\beta}_T) = \beta_{T1}b_{T1}(t) + \beta_{T2}b_{T2}(t) + \beta_{T3}b_{T3}(t), \tag{4.1}$$

$$\alpha_C(t|\boldsymbol{\beta}_C) = \beta_{C1}b_{C1}(t) + \beta_{C2}b_{C2}(t). \tag{4.2}$$

As defined in Section 3.2, b_{T1} is linear on $[0, 1.5)$, cubic on $[1.5, 6.17)$ and $[6.17, 63.33)$, constant on $[63.33, \infty)$, b_{T2} and b_{C1} are constant functions, and $b_{T3}(t) = b_{C2}(t) \equiv \log(t)$. The estimates of $\boldsymbol{\beta}_T$ and $\boldsymbol{\beta}_C$ from *heft* are

$$\hat{\boldsymbol{\beta}}_T = (0.00017, -8.13, 0.69)^T$$

and

$$\hat{\boldsymbol{\beta}}_C = (-3.68, -0.18)^T,$$

respectively. The estimates of the marginal densities, survival functions, hazard functions, and the log hazards ratio are in Figures 4.3 to 4.6.

From the plot in Figure 4.3, we find that both estimated densities have a high value in the first twenty months. But the density corresponding to the untreated eye is higher than that of the treated eye during the observed time period. The estimated density of the failure time of the untreated eyes has its maximum at time zero, while the other density achieves its maximum at about month six. Since the regression model is for the log hazard, we calculate the estimated densities based on the estimate of log hazards as in equation (2.8). A pointwise confidence interval for the estimated densities would require integration of the upper and lower bounds of confidence intervals for the hazard function. This is difficult, so these confidence intervals have not been calculated.

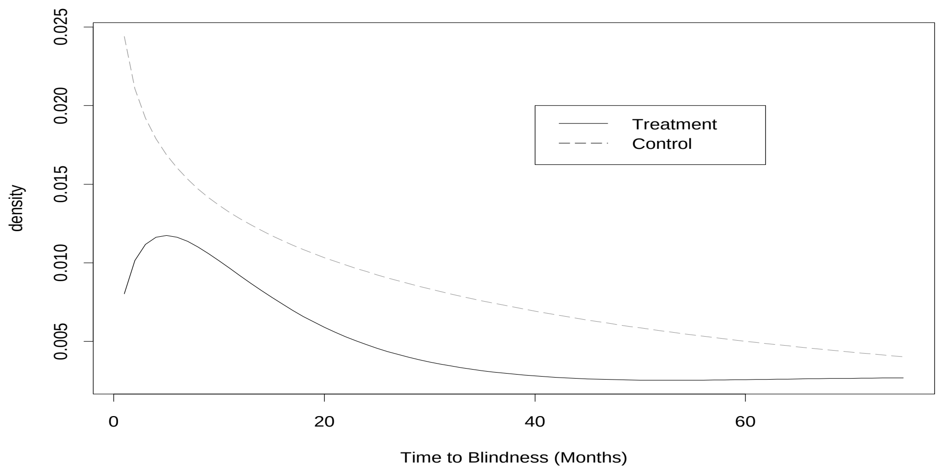

Figure 4.3: The estimated density functions of the eye data for Model 1.

The estimated survival curves, in Figure 4.4, show that the estimated survival function of the treated eyes is always greater than that of the untreated eyes. Therefore, there appears to be a large treatment effect.

From the plots of the estimated hazard functions in Figure 4.5, we see that the risk of blindness in the untreated eye is much higher than that of the treated eye. The risk of blindness in the untreated eyes decreases with time while the treated eyes may have a maximal risk at $t = 6$. Based on the estimated hazard functions, we can see that the treatment may have delayed the onset of blindness for the patients for the first couple of months after the operation.

Figure 4.6 shows the estimated log hazards ratio with pointwise 95% confidence intervals. Values greater than zero indicate a higher risk of the control group and so

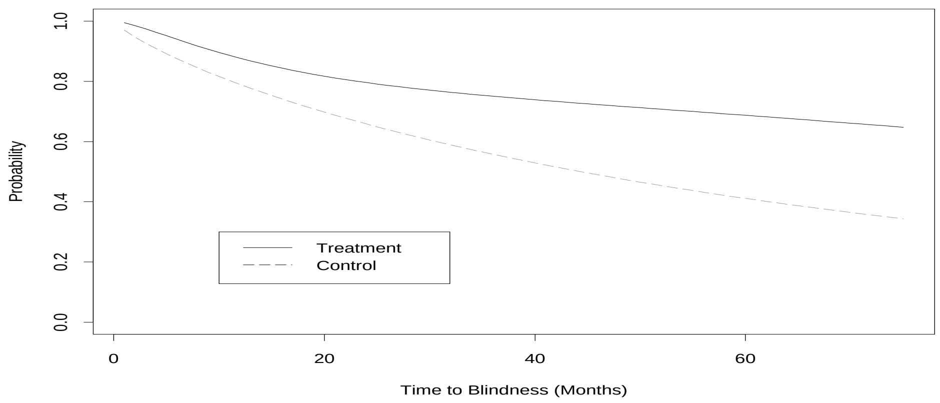

Figure 4.4: The estimated survival functions of the eye data for Model 1.

it seems that the treatment has a beneficial effect. The estimate of log relative risk of blindness of the untreated eye to the treated eye is equal to 1.15 at $t = 1$, and decreases to 0.34 at about month six, then increases smoothly. After attaining its maximum value of 1.186 at month 42, it decreases slowly.

Thus, based on Figure 4.6, the estimated log ratio of the two hazards gives an answer to questions 1 and 2 asked in the beginning of Section 4.2.

4.2.2 Model 2

To test if there is a treatment effect, we assume the log hazard functions of the treated eye and the untreated eye are from the same regression space $\mathcal{B}$. So we will use the same restricted B-spline basis plus a log tail for the two regressions. To choose knots

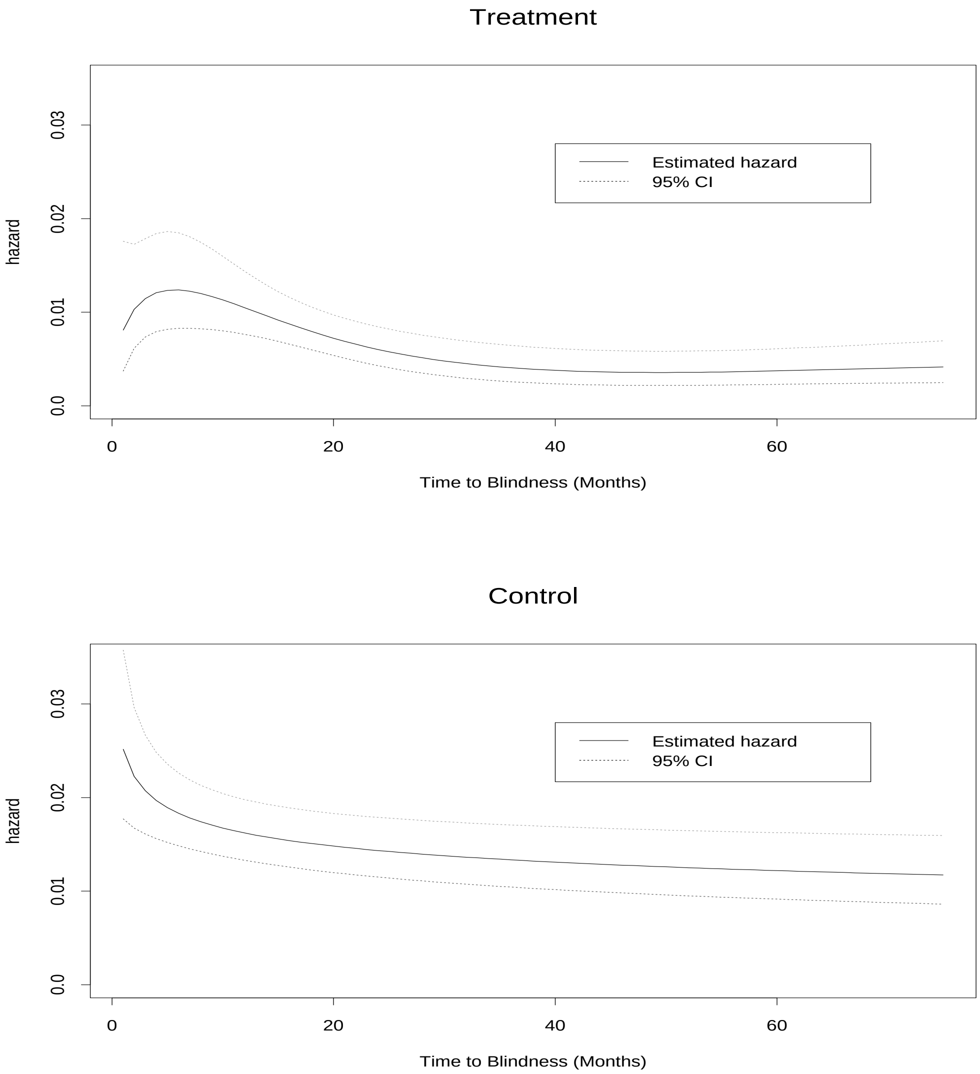

Figure 4.5: The estimated hazard functions of the eye data for Model 1 with pointwise 95% confidence intervals.

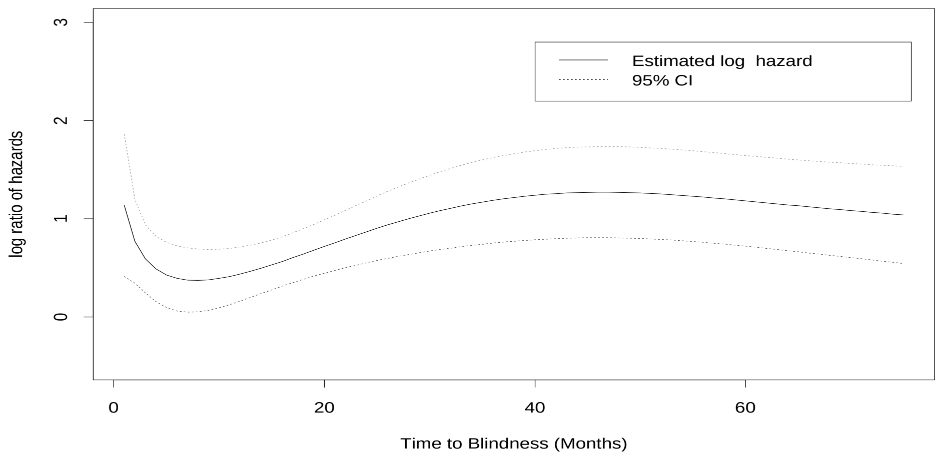

Figure 4.6: The estimated log ratio of the hazard functions of the eye data with pointwise 95% confidence intervals for Model 1.

that define $\mathcal{B}$, we refer to the knots used for Model 1. As *heft* uses $1.5, 6.17$, and 63.33 for the treatment group and no knots for the controls, we use $1.5, 6.17$, and 63.33, the union of the knots used for the two groups, as the knots to define our regression space $\mathcal{B}$. Then $\mathcal{B}$ is a three dimensional space and the log hazard regression, Model 2, is:

$$\alpha_T(t|\boldsymbol{\beta}_T) = \beta_{T1}b_1(t) + \beta_{T2}b_2(t) + \beta_{T3}b_3(t), \tag{4.3}$$

$$\alpha_C(t|\boldsymbol{\beta}_C) = \beta_{C1}b_1(t) + \beta_{C2}b_2(t) + \beta_{C3}b_3(t), \tag{4.4}$$

where $b_1 = b_{T1}$, $b_2 = b_{T2}$, and $b_3 = b_{T3}$ as defined in Section 4.2.1. Using the *heft* algorithm with fixed knots, we get the estimates

$$\hat{\boldsymbol{\beta}}_T = (0.00017, -8.13, 0.69)^T,$$

$$\hat{\boldsymbol{\beta}}_C = (0.00007, -5.15, 0.16)^T.$$

As in Section 4.2.1, we calculate the estimates of the densities, survival, hazard functions, and the log hazards ratio. We give the plots in Figures 4.7 to 4.10. The results are similar to those from Model 1.

Since Model 1 and Model 2 use the same regression space for the log hazard of the failure times of the treated eyes, the two models produce the same estimates for the failure times of the treated eyes. Thus we do not repeat the plots of these estimates. We compare the estimates from Model 1 and Model 2 for the untreated eyes in Figures 4.11 to 4.13.

From Figures 4.11 to 4.13 we see that Model 1 and Model 2 give almost the same estimates of the survival function of the untreated eyes. But the estimated densities and hazards look a little different.

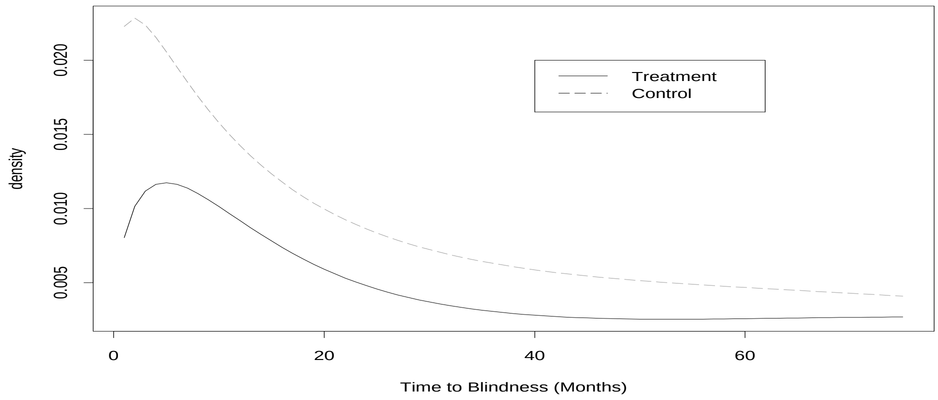

Figure 4.7: The estimated density functions of the eye data for Model 2.

Next we use Model 2 to test the hypothesis that the failure times of the treated eye and the untreated eye have the same distribution. By the definition of Model 2 from (4.3) and (4.4), this hypothesis is equivalent to

$$H_0: \ \beta_{Tj} - \beta_{Cj} = 0, \quad j = 1, 2, 3.$$

Using the test described in Section 3.5, we obtain the test statistics and p-values shown in Table 4.1.

Hence we may reject H_0 and conclude that the two distributions are different. From our plots of the log hazards ratios in Figures 4.6 and 4.10, we conclude that the laser photocoagulation treatment had a significant effect in delaying the onset of blindness in patients with diabetic retinopathy.

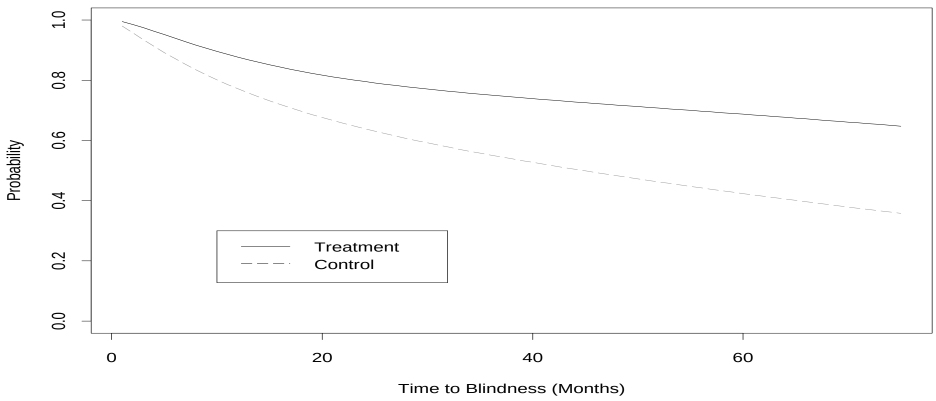

Figure 4.8: The estimated survival functions of the eye data for Model 2.

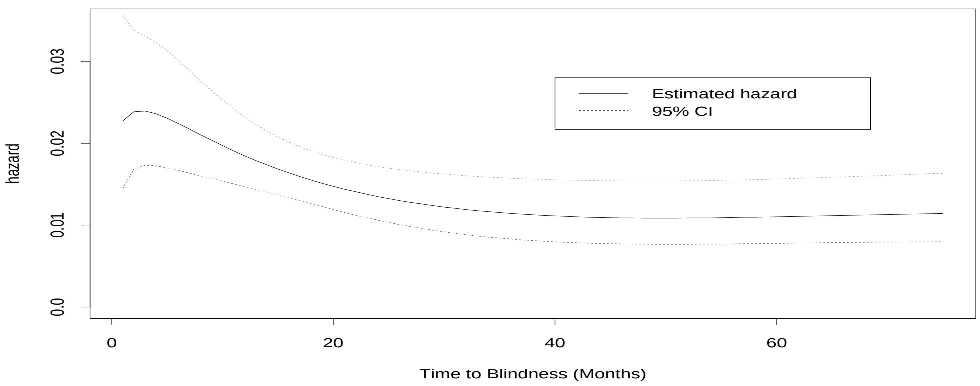

Figure 4.9: The estimated hazard function of the untreated eye for Model 2.

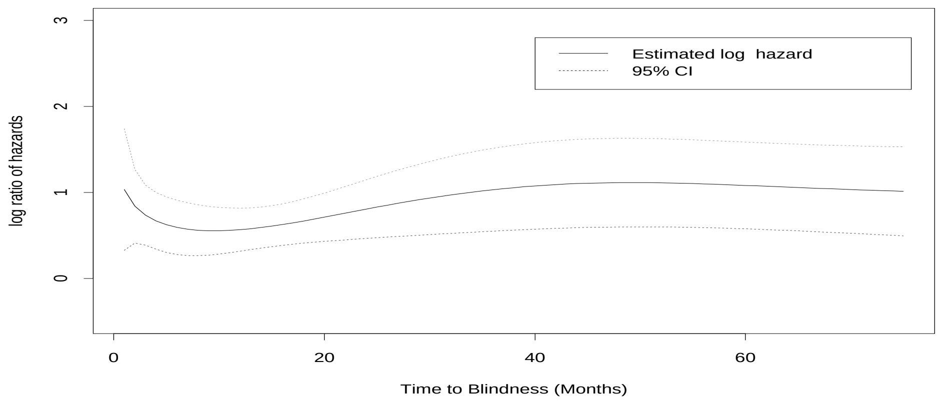

Figure 4.10: The estimated log ratio of the hazard functions of the eye data with pointwise 95% confidence intervals for Model 2.

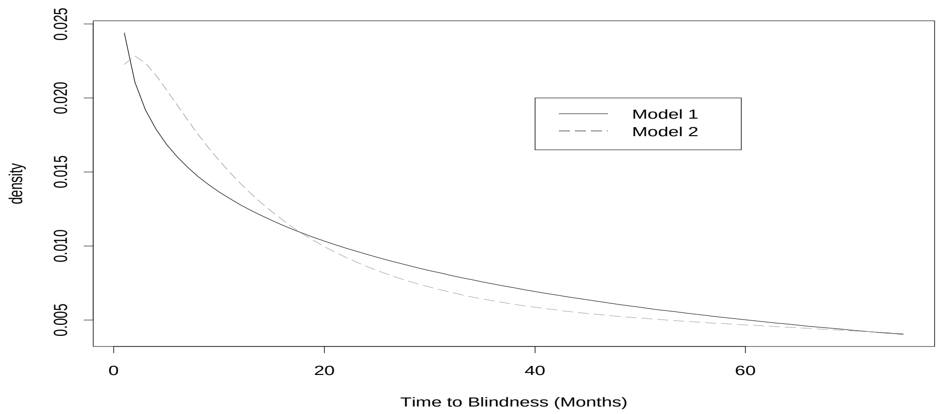

Figure 4.11: The estimated density functions of the failure times of the untreated eyes.

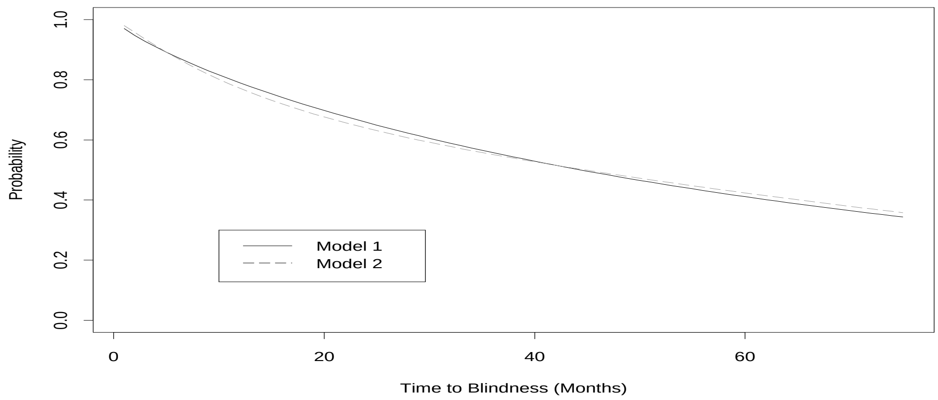

Figure 4.12: The estimated survival functions of the untreated eyes.

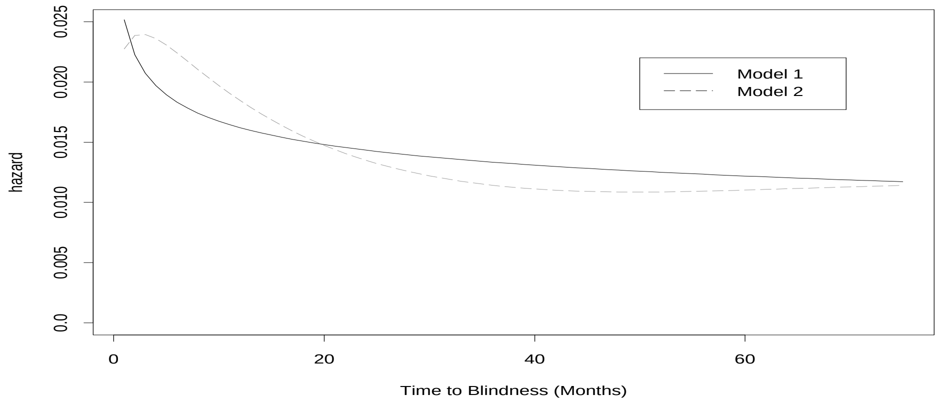

Figure 4.13: The estimated hazard functions of the failure times of the untreated eyes.

Table 4.1: Test statistic and p-value for testing that the failure times of the treated eyes and the untreated eyes have the same distribution.

Test	χ^2_3	p-value
$\beta_{Tj} - \beta_{Cj} = 0, j = 1, 2, 3$	48.56	0

4.3 Other Models Used to fit the Eye Data

Huster et. al. (1989), considered modeling the marginal distributions for the eye data as exponential and as Weibull or modeling the data via the Cox proportional hazards. We show the estimated marginal survival curves from fitting different models in Figures 4.14 and 4.15. The Kaplan-Meier estimate is also given. Note that for the treatment Model 1 and Model 2 are the same and for the control Model 1 is the Weibull. We see that all fit inside the confidence intervals of the Kaplan-Meier estimates except the estimates from the exponential model. But the estimated survival functions from our log hazard regression models are closer to the Kaplan-Meier estimates than those from the parametric models.

In this section we use Model 2, defined in Section 4.2.2, to decide which model(s) are appropriate to the eye data. Then, based on the test results, if we can choose a standard parametric model, say exponential, Weibull, or proportional hazards, we will use the selected model to test if there is a treatment effect.

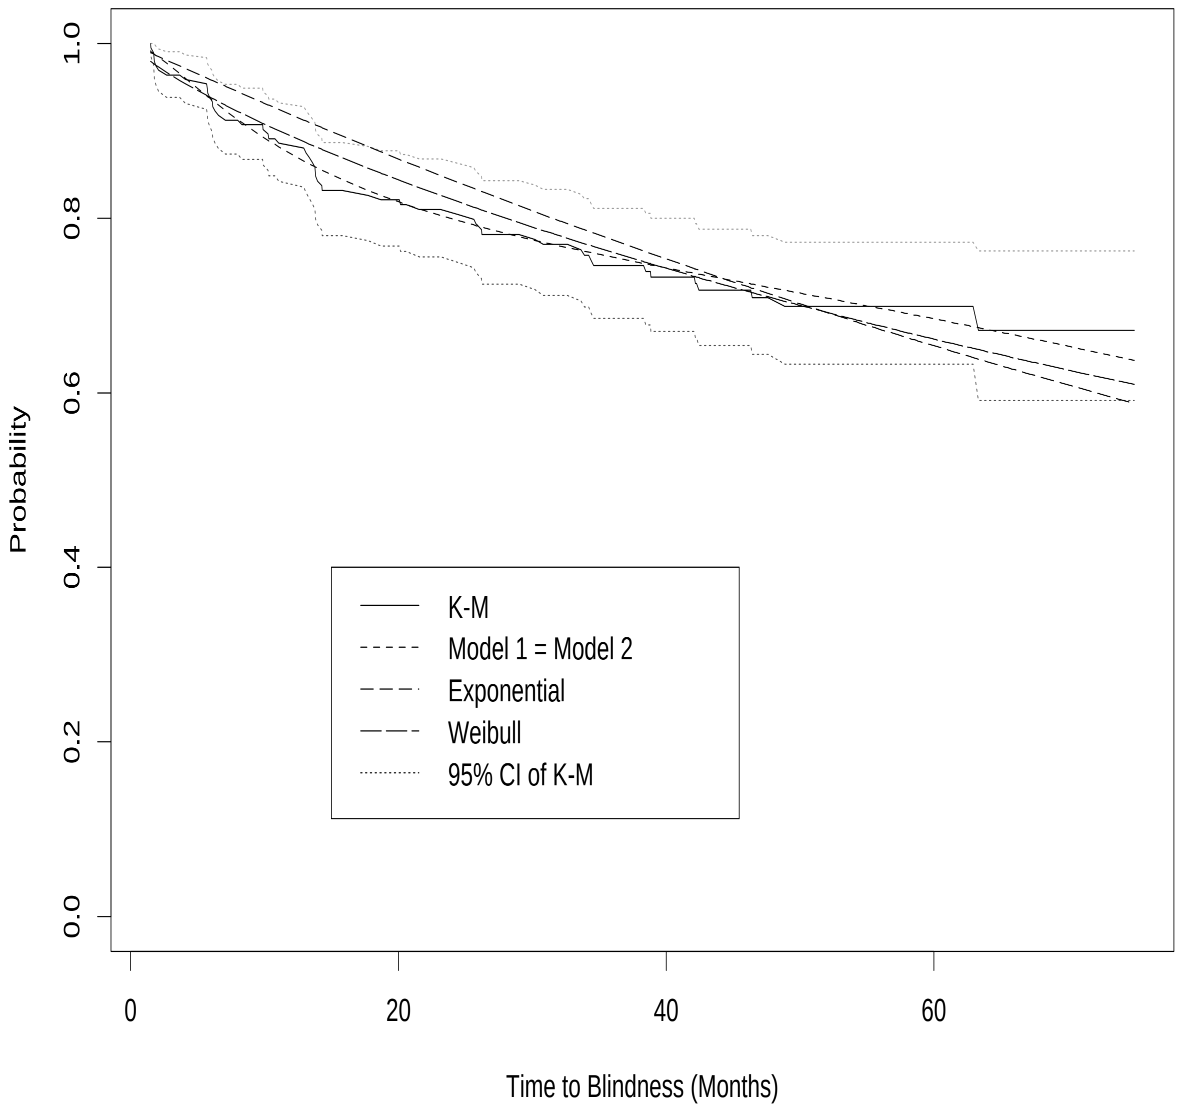

Figure 4.14: The estimated survival curves of the treated eye.

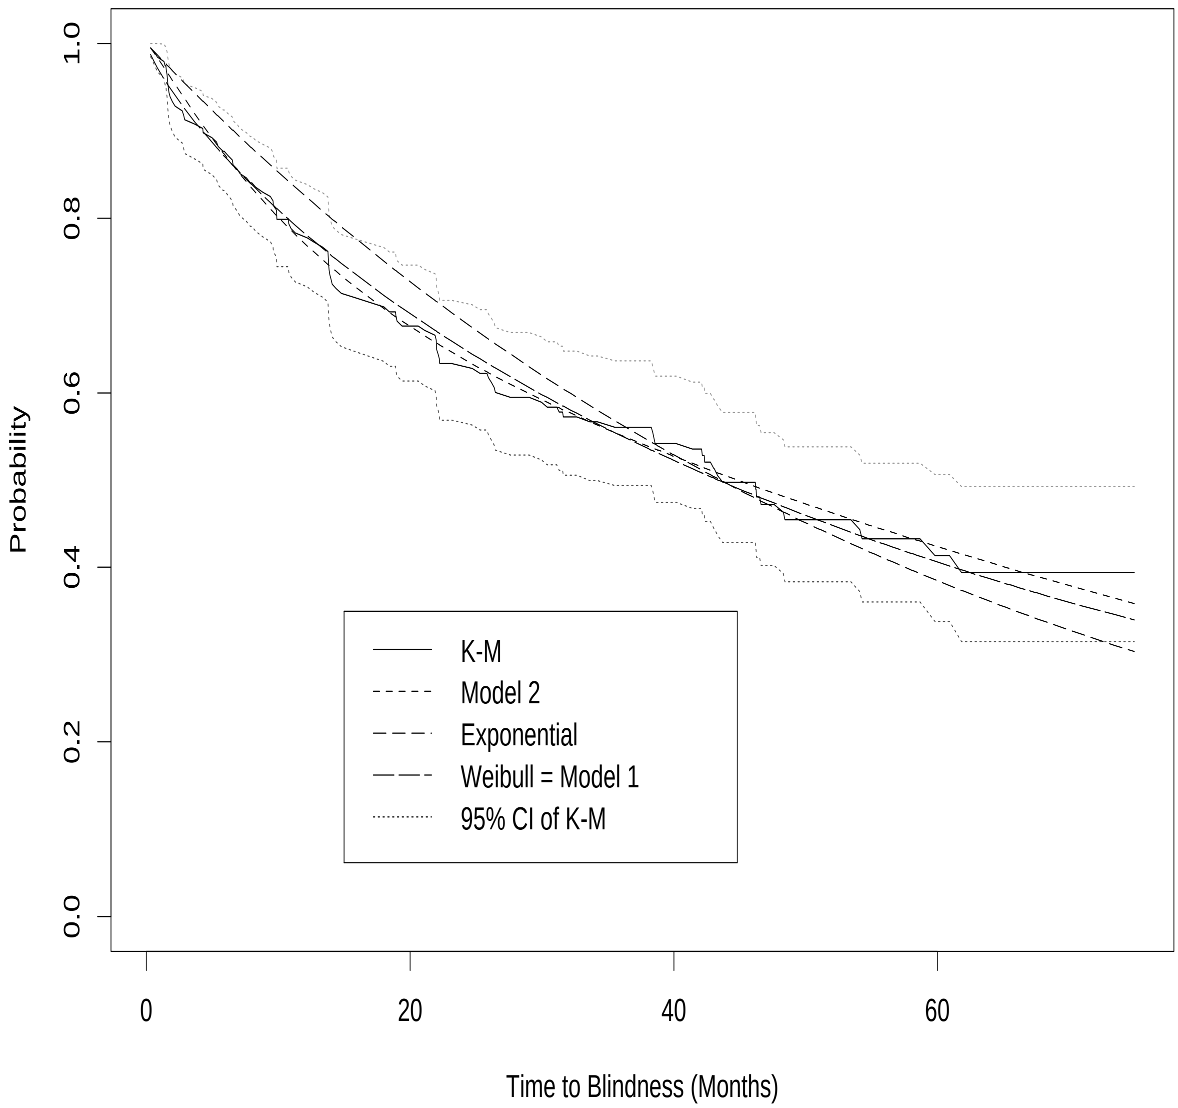

Figure 4.15: The estimated survival curves of the untreated eye.

4.3.1 Exponential Model

Since Model 2 from (4.3) and (4.4) includes the exponential distribution, we can use the model to test if the marginal distributions are exponential. As we discussed in Section 3.5, we test the two null hypotheses:

$$\beta_{Tj} = 0, j = 1, 3$$

and

$$\beta_{Cj} = 0, j = 1, 3.$$

Table 4.2 presents our test results.

Table 4.2: χ^2 test statistics and p-values for testing that failure times are exponentially distributed.

Test	χ^2_2	p-value
$\beta_{Tj} = 0, j = 1, 3$	11.26	0.004
$\beta_{Cj} = 0, j = 1, 3$	8.14	0.017

Hence we reject the hypotheses that the failure times of the treated eye and the untreated eye are exponentially distributed.

4.3.2 Weibull Model

To test the hypotheses that the marginal distributions are Weibull, we test the two hypotheses:

$$\beta_{T1} = 0$$

and

$$\beta_{C1} = 0.$$

We show the test results in Table 4.3.

Table 4.3: z statistics and p-values for testing that failure times have Weibull distributions.

Test	$\hat{\beta}$	se	z	p-value
$\beta_{T1} = 0$	1.73×10^{-4}	6.28×10^{-5}	2.75	0.006
$\beta_{C1} = 0$	6.94×10^{-5}	3.59×10^{-5}	1.93	0.053

Thus we reject the hypothesis that the failure time of the treated eye has a Weibull distribution and we can not accept the hypothesis that the failure time of the untreated eye has a Weibull distribution.

4.3.3 Cox Proportional Hazard Model

Now we use Model 2 to test the hypothesis that the two hazards are proportional. As mentioned in Section 3.5, we need to test

$$H_0: \ \beta_{Tj} - \beta_{Cj} = 0, \quad j = 1, 3,$$

since b_2 is constant. The test statistics and p-values are shown in Table 4.4.

Table 4.4: Test statistic and p-value for testing that the two hazards of the treated eyes and the untreated eyes are proportional.

Test	χ^2_2	p-value
$\beta_{Tj} - \beta_{Cj} = 0, j = 1, 3$	2.81	0.25

Thus we may assume that the two groups have proportional hazards, that is, we may assume that the relative risk of blindness of the untreated eye versus the treated eye is a constant.

Now we assume that the two failure times are independent and fit the Cox proportional hazards model (3.5) for the eye data. To test the treatment effect, we test $\gamma = 0$. The results are given in Table 4.5.

Table 4.5: Test results from the Cox proportional hazards model for the eye data.

Test	$\hat{\gamma}$	se	z-value	p-value
$\gamma = 0$	0.777	0.169	4.6	4.2×10^{-6}

The Cox proportional hazards model also indicates that there is a treatment effect. Figures 4.16 to 4.19 compare the pointwise 95% confidence intervals of the log hazards ratio from Models 1 and 2 and the Cox proportional hazards model, which assumes the two failure times are independent. We can see that the confidence interval from Model

2 is narrower than the confidence interval from the Cox proportional hazards model during the months 4 through 23, when the most failures in the both groups occurred.

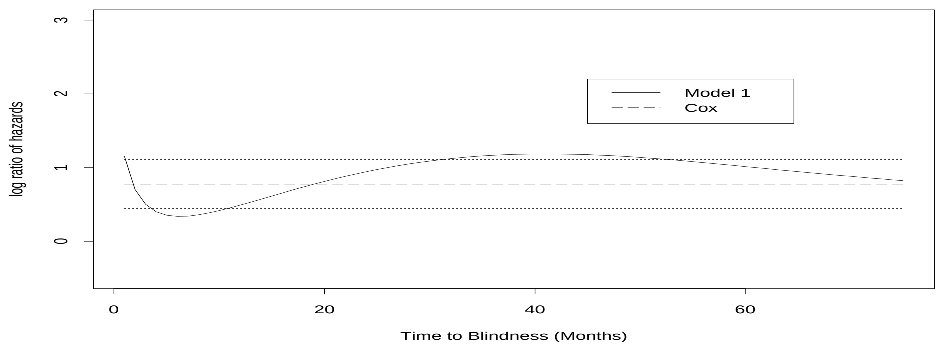

Figure 4.16: The estimated log hazards ratio from Model 1 and the Cox proportional hazards model with the pointwise 95% confidence intervals from the Cox proportional hazards model.

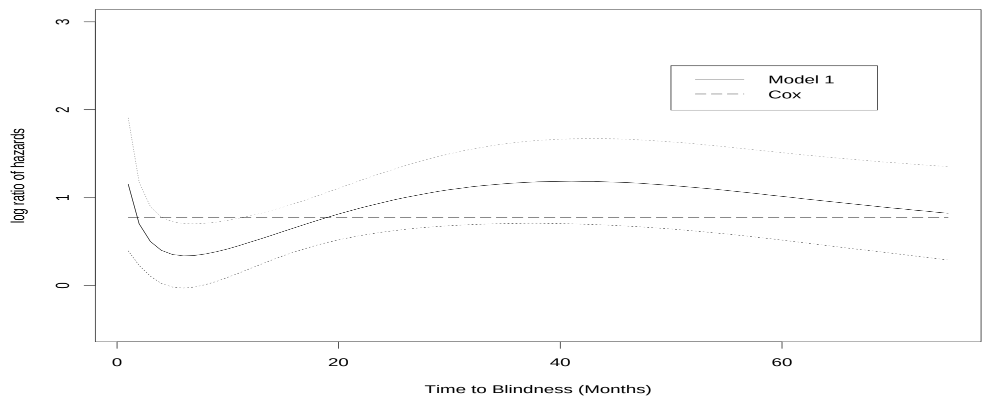

Figure 4.17: The estimated log hazards ratio from Model 1 and the Cox proportional hazards model with the pointwise 95% confidence intervals from Model 1.

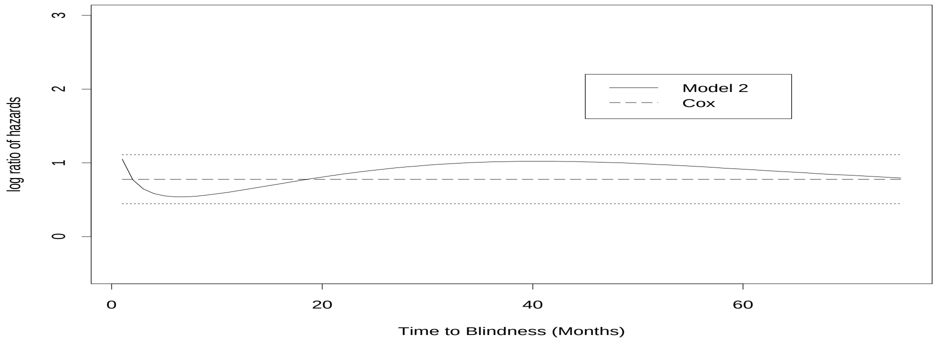

Figure 4.18: The estimated log hazards ratio from Model 2 and the Cox proportional hazards model with the pointwise 95% confidence intervals from the Cox proportional hazards model.

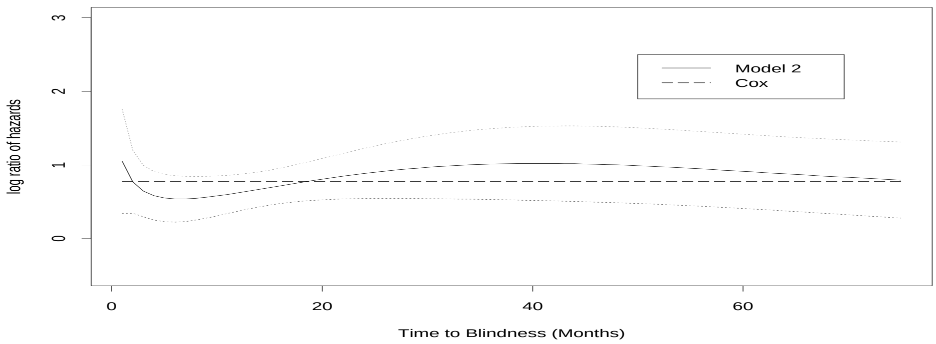

Figure 4.19: The estimated log hazards ratio from Model 2 and the Cox proportional hazards model with the pointwise 95% confidence intervals from Model 2.