

## Heterogeneity's ruses: the neglected role of between-individual variability in longitudinal studies of COPD exacerbations

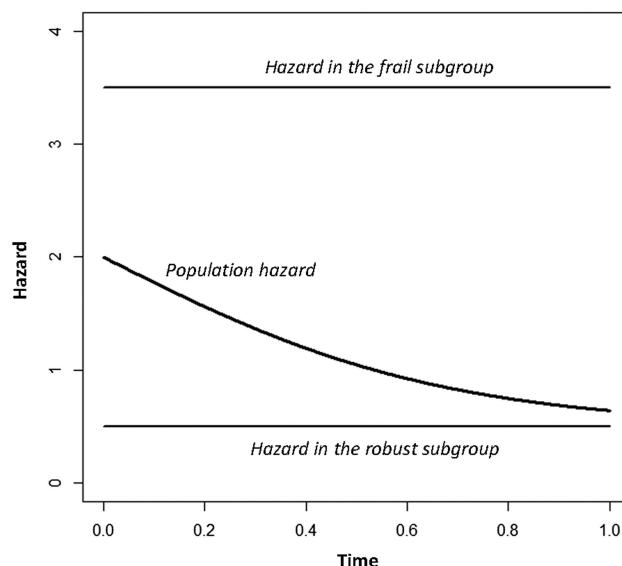
### ABSTRACT

Studying the causal and temporal association between past and future exacerbations in COPD is an active area of research. Standard survival analysis techniques often used in such studies typically produce results that pertain to the overall population, whereas the greatest interest is in the study of associations within individuals. A factor that can lead to profound discrepancies between population-level and individual-level survival patterns is the between-individual heterogeneity in the rate of exacerbations. We briefly review two studies that, while reporting valid results for the overall population, drew conclusions at the individual level that could not be supported by the observations. We caution on the distinction between population and individual-level associations in survival analysis, and recommend accounting for heterogeneity in future studies.

There is great interest in studying the impact of exacerbations on the course of COPD. Such studies often rely on survival analysis of longitudinal data. Standard survival analysis techniques produce results that pertain to the population, whereas the real interest is in within-individual associations. A crucial issue in this context is the heterogeneity in the exacerbation rate across individuals. Estimates of hazard function from survival analysis are particularly associated with surprisingly conflicting and potentially misleading patterns when heterogeneity is not taken into account.<sup>1</sup> Figure 1 provides an illustrative example.

We are concerned that the conventional wisdom about the role of exacerbations in the course of COPD might have been confounded by such phenomena. We illustrate this by replicating analyses from two previous studies using a simulated 'stylised cohort', in which there is no within-individual associations between exacerbations, but there is between-individual heterogeneity in exacerbation rates. Details of methods and additional results are provided in online supplementary material.

Figure 2A is the hazard function and median between exacerbation times after successive exacerbations from Suissa *et al.*<sup>2</sup> Based on this and similar results, the authors stated 'occurrence of every new exacerbation requiring hospitalisation



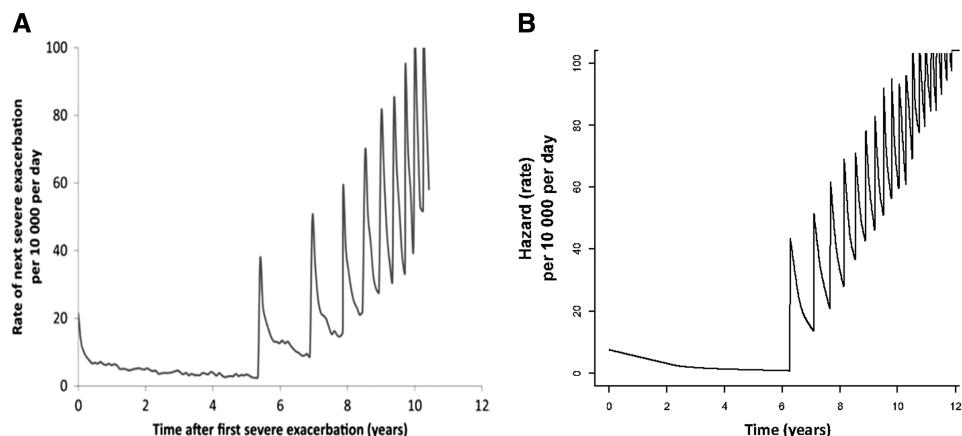
**Figure 1** Heterogeneity's ruses: an illustrative example of discrepancies between the population and individual hazards. This is a hypothetical population consisting of two subgroups (each 50% of the population) with regard to the hazard of an absorbing event (such as death or 'next' exacerbation): a *frail* subgroup that experiences the event at a high rate and a *robust* subgroup that experiences the event at a lower rate. While the hazard is constant within each subgroup, the population hazard decreases over time. This is because at any given point in time, the population hazard is the weighted average of subgroup hazards, with weights being the proportion of each subgroup in the event-free population. At first, everyone is event free, so the overall hazard is the midpoint between the subgroups' hazards. As time passes, more individuals in the *frail* than the *robust* subgroup experience the event, thus, the event-free population becomes predominantly of the *robust* subgroup, lowering the population hazard. If the population hazard is interpreted at the individual level, one might erroneously conclude that the hazard of the event decreases over time. Phenomena similar to this can be responsible for the findings of both the studies discussed in this paper. For other examples refer to reference 1.

worsens the course of the disease and increases the risk of a subsequent exacerbation', and made recommendations, such as 'delaying the second severe exacerbation as a target of COPD management'. However, as seen in figure 2B, similar patterns can be generated by repeating the analysis in the stylised cohort, where such interpretations do not hold due to a lack of causal associations or between-individual heterogeneity.

Similarly, Hurst *et al* estimated the population hazard function for time-to-next-exacerbation.<sup>3</sup> The decreasing hazard was interpreted as an increased risk of exacerbation in the period after the previous one. They concluded that 'exacerbations are not random events but cluster together in time', a finding that they believed 'has important implications for the targeting of preventative interventions'. Nevertheless, heterogeneity per se can generate decreasing population hazard through the mechanism explained in figure 1, also observed in the stylised cohort (see online supplementary material), even in the absence of temporal clustering of exacerbations within individuals.

A key aspect of both studies is that the estimated hazard functions belong to the population but were interpreted at the individual level. These interpretations are not necessarily incorrect; however, given the chosen analytical framework, it is impossible to discern to what extent the observed results were because of within-individual associations or between-individual heterogeneity. This distinction is of paramount importance, as the existence of within-individual associations can change priorities, for example, by shifting the attention from the prevention of COPD to the prevention of exacerbations.

Our intent is not to provide solutions to enable within-individual inference in this particular context. Generally, a class of survival analysis called frailty models,<sup>4</sup> or recurrent-event models that allow comparison of hazard within the same person,<sup>5</sup> might hold some promise. We invite the respiratory research community to harness the power of such techniques towards better understanding the role of exacerbations in the natural history of COPD.



**Figure 2** Hazard function and median between exacerbation times of successive exacerbations from Suissa *et al* (A) and the stylised cohort (B). The authors of the original study mentioned three features of their figure (A): high risk of subsequent exacerbations in the immediate period after each exacerbation (spikes followed by a downward curve), the decreasing between-exacerbation times (spikes getting closer to each other), and the increasing hazard of a future exacerbation after a previous one (the general upward slope). All three features can be explained by heterogeneity: the downward curves after spikes are manifestations of the mechanism explained in figure 1. As one moves along the X-axis, the observations that contribute to the curve come from individuals with higher number of previous exacerbations. In the presence of heterogeneity, these individuals are a selected subgroup of the population with high background rate of exacerbation, explaining the higher hazard and shorter time-to-next-exacerbation.

Mohsen Sadatsafavi, J Mark FitzGerald

Faculty of Medicine, Institute for Heart and Lung Health, The University of British Columbia, Vancouver, British Columbia, Canada

**Correspondence to** Dr Mohsen Sadatsafavi, Centre for Clinical Epidemiology and Evaluation, 7th Floor, 828 West 10th Avenue, Research Pavilion, Vancouver, British Columbia, Canada V5Z 1M9; msafavi@mail.ubc.ca

**Acknowledgements** The authors would like to thank Dr Mahyar Etmianan, University of British Columbia, for his helpful comments. The title of this work is inspired by a related and seminal work by Vaupel and Yashin: Heterogeneity's ruses: some surprising effects of selection on population dynamics. *Am Stat* 1985;39(3):176–85.

**Contributors** MS conceived the idea. MS and MF elaborated on the features to be discussed, and on the overall design of the simulations. MS performed the simulations and wrote the first drafts. Both authors revised the draft and agreed on the final version.

**Competing interests** None.

**Provenance and peer review** Not commissioned; externally peer reviewed.

► Additional material is published online only. To view please visit the journal online (<http://dx.doi.org/10.1136/thoraxjnl-2013-205061>).



Open Access  
Scan to access more  
free content

**To cite** Sadatsafavi M, FitzGerald JM. *Thorax* 2014;69:1043–1044.

Received 27 December 2013

Revised 14 January 2014

Accepted 17 January 2014

Published Online First 18 February 2014

*Thorax* 2014;69:1043–1044.  
doi:10.1136/thoraxjnl-2013-205061

## REFERENCES

- Vaupel JW, Yashin AI. Heterogeneity's ruses: some surprising effects of selection on population dynamics. *Am Stat* 1985;39:176–85.
- Suissa S, Dell'anielli S, Ernst P. Long-term natural history of chronic obstructive pulmonary disease: severe exacerbations and mortality. *Thorax* 2012;67:957–63.
- Hurst JR, Donaldson GC, Quint JK, et al. Temporal clustering of exacerbations in chronic obstructive pulmonary disease. *Am J Respir Crit Care Med* 2009;179:369–74.
- Hougaard P. Frailty models for survival data. *Lifetime Data Anal* 1995;1:255–73.
- Kelly PJ, Lim LL. Survival analysis for recurrent event data: an application to childhood infectious diseases. *Stat Med* 2000;19:13–33.



# Supplementary Material for "Heterogeneity's ruses: the neglected role of between-individual variability in longitudinal studies of COPD exacerbations"

---

Authors: Mohsen Sadatsafavi, J. Mark FitzGerald

## Contents

Background .....	2
Methods.....	2
Details for Suissa <i>et. al.</i> .....	3
Details for Hurst <i>et. al.</i> .....	5
What happens when heterogeneity is removed? .....	6
Details for Figure 1.....	7
References .....	8

## **Background**

The authors of the case studies have observed certain patterns in the results of their survival analysis of exacerbations in patients with COPD, and accordingly they have made conclusions on associations between exacerbations at the individual level[1,2]. The purpose of this exercise was to provide one counter example, enough to show that the observed patterns at the population level do not necessarily imply the individual-level conclusions made by the authors. We show that the heterogeneity in the rate of exacerbation can generate the same pattern of results without such conclusions holding at the individual level.

In this context, heterogeneity refers to the variability in the propensity to exacerbate across individuals. Part of such variability can be explained by variability in measurable factors such age, sex, or the degree of lung function impairment; but it is known that patients with COPD vary in their propensity to exacerbate over and above the effect of such variables[3]. We note that neither Hurst et al. nor Suissa et. al. controlled for the effect of such variables in estimating the population hazard function. Suissa et. al. did adjust for potential confounding variables in other analyses, but the administrative data used in their study does not have important variables, such as measures of lung function impairment. All in all, it is likely that in both studies there was significant (residual) heterogeneity in the population.

## **Methods**

We use a simple simulation study with very simple features (to separate and highlight the importance of heterogeneity) to show this in action. The few parameters of the study are chosen to roughly mimic the features of the case studies. We did not aim to generate the exact results. This would have required knowing the joint distribution of several variables defining the study sample (follow-up times, event times, background rate of exacerbations, demographics characteristics, to name a few). The simulation was performed in the statistical programming environment R version 3.0.2[4].

The stylized cohort consists of longitudinal data of 100,000 individuals (to minimize sampling variability). There are no demographic characteristics or competing risk, and censoring occurs at a fixed time in follow-up. With individuals in our cohort being similar in all aspects except rate of exacerbations, we did not need to model the inclusion criteria as they do not have any implications on our results (e.g., taking the first exacerbation as the cohort entry as in Suissa et al. would have only changed the distribution of the background hazard in our population, which is already arbitrary).

Within each individual, the hazard of exacerbation was fixed over time. The background rate of exacerbations was from a gamma distribution. For each individual, time to next exacerbations were generated from an exponential distribution (corresponding to a fixed hazard) with the rate equal to the background hazard of exacerbation for that individual. This process was continued till the next exacerbation time occurred past the follow-up time. Therefore, the data for each individual with  $n$  exacerbations were represented by  $(n + 1)$  time intervals adding up to follow-up time, with the first  $n$  intervals representing time to the next exacerbation, and the last interval representing time to censoring (due to end of follow-up). Note that using this scheme, the distribution of the number of exacerbations is from a Poisson distribution within each individual according to the background rate of exacerbation, and exacerbations occur completely independent of each other and with a constant hazard. The choice of the gamma distribution for modelling heterogeneity causes the distribution of between-exacerbation times across the population to follow a negative binomial distribution[5].

## **Details for Suissa et. al.**

Suissa et al. have reported on the analysis of the long-term natural history of COPD with emphasis on severe exacerbations and mortality[1]. The study used 17 years longitudinal data of 73,106 individuals with COPD.

For this study, the gamma distribution had a shape and scale parameters of 0.2. This results in an average exacerbation rate of 1 per year, with a standard deviation of 2.28. Individuals were followed for 10 years. We followed the approach by the authors in estimating the population hazard function (**Figure 3A**) of the original publication. We estimated the baseline hazard and median inter-exacerbation times as a function of the previous number of exacerbations using the life table method via the *survfit* function[6]. We used *Lowess* with default values for smoothing the hazard[7]. Results are provided in the manuscript (**Figure 2**).

In addition, we estimated the relative risk (RR) of past exacerbations on future ones (similar to Table 2 of Suissa et al.) by fitting a Cox proportional hazards model (using *coxph* function[8]) with the number of previous exacerbations as a time-dependent covariate. Results are provided in the Table below.

<b>Table E1:</b> relative risk (RR) of a subsequent exacerbation associated with each successive exacerbation, as a function of the time since the first exacerbation	
Exacerbation #	RR in the presence of heterogeneity
0 (reference)	1.000
1	5.964
2	15.791
3	20.202
4	25.583
5	29.863
6	35.780
7	39.493
8	45.485
9	54.129

These RRs are associational but not causal. They can be used for predicting the time of next exacerbations, but one cannot conclude that by preventing each exacerbation, the rate of the future one is decreased. This is because of confounding by frailty: here the background rate of exacerbation affects the risk of past and future exacerbations without itself being part of the causal pathway, thus acting as a classical confounder. Note that some other observations by the authors of the original study

can be potentially due to the same phenomenon. For example, the increased hazard of death after each exacerbation can be due to the heterogeneity and positive correlation between hazard of death and exacerbation at the population level.

## Details for Hurst *et. al.*

Hurst et al., followed 297 COPD patients for 904 patient-years to evaluate if COPD exacerbations tend to occur in clusters[2]. The investigators separated 1,923 pairs of successive exacerbations arisen in the cohort. They fitted a Weibull distribution to the time intervals between such successive exacerbations. In doing so, between-exacerbation intervals were considered as independent observations and the clustering of such intervals within individuals was ignored, an issue that has already raised concerns[9]. The temporal clustering of successive exacerbations was assessed by testing the shape parameter of the fitted Weibull distribution, with shape values of less than 1 indicating a declining hazard thus implying clustering of exacerbations. This parameter had a value of 0.966 (95%CI 0.948–0.985, P<0.001).

In simulating data for this study, the shape parameter was set to 2.71 while the scale parameter was set to 1, corresponding to the average number of exacerbation in the study. Follow-up was for 3 years (average follow-up time in the original study). We noted that in estimating the shape of the Weibull distribution, the investigators only used between-exacerbation intervals, and not the censored intervals. This *per se* can bias the results given that censored intervals were incorrectly removed from the estimation (see [9] for explanation). We therefore improved the analysis by fitting a parametric survival model (the *survreg* function[10]) with a Weibull distribution. This framework properly takes into account censoring.

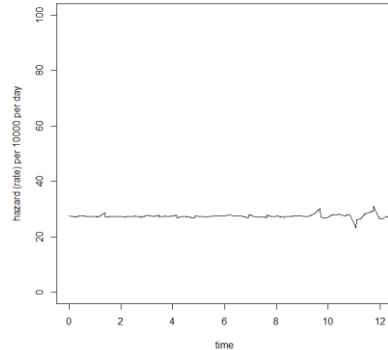
The estimated shape parameter of the Weibull distribution for the stylized cohort was 0.85, thus suggesting a strong clustering effect. This is despite the fact that within individuals, exacerbation

occurred purely at random. This is because in the presence of heterogeneity, those who exacerbate more often tend to provide more between-exacerbation periods than infrequent exacerbators. As such, there would be a preponderance of shorter exacerbation times contributed by high exacerbators, causing the population hazard to decrease over time. Further explanation of this phenomenon can be found in Keene et al.[9].

## What happens when heterogeneity is removed?

We have so far shown that the pattern of results observed in the original studies can be replicated in a stylized cohort in presence of heterogeneity without any individual-level associations among exacerbations. To fully attribute the observations to the impact of heterogeneity (and not, say, to the artefacts of data analysis), we take one additional step to show that the afore-mentioned patterns disappear in the absence of heterogeneity. Heterogeneity was removed by assigning the same rate of exacerbation to each individual which was equal to the average rate of exacerbation in the population in the respective analyses in the presence of heterogeneity. For Suissa et al., the step-wise hazard function is provided below.

**Figure E1:** Hazard function of successive exacerbations from Suissa et al. in the stylized cohort with homogenous rate of exacerbation



For Hurst et. al., the estimated shape parameter was 1.001. In attempt to mitigate the problem of mixing high and low exacerbators and taking the unit of the study as exacerbations not individuals, Hurst et al. repeated the analysis by dividing the sample into two subgroups (low and high exacerbators). As mentioned by others, this broad categorization hardly mitigates the frailty problem as individuals within subgroups still vary in their rate of exacerbation[9]. In addition, by dividing the data around an observed value (median), it is no longer guaranteed that under the null hypothesis of no temporal clustering, even in the absence of heterogeneity, the shape parameter will be 1. We nonetheless repeated this analysis as well, both with and without heterogeneity. In the presence of heterogeneity, the shape parameter was 1.05 for low- and 0.85 for high exacerbators. When we removed the impact of heterogeneity, the fitted shape parameter was 1.11 for low and 1.01 for high exacerbators, suggesting that at least for low exacerbators, dividing the sample around the median will change the value of the test statistic under the null hypothesis.

## Details for Figure 1

This figure presents unconditional (population) hazard function for a population consisting of two equal-sized subgroups (frails and robust). While the conditional (subgroup) hazards are constant, the unconditional hazard decreases over time. This is because at time  $T$  of follow-up, the population hazard is a weighted average of hazard in subgroup 1 (denoted by  $h_1$ ) and hazard in subgroup 2 (denoted by  $h_2$ ), with weights being the proportion of each subgroup still event free at time  $T$ . Because of the constant hazard, the survival curve of subgroups 1 and 2 follow an exponential curve and at time  $T$  is equal to, respectively,  $e^{-h_1 \cdot T}$  and  $e^{-h_2 \cdot T}$ . The population hazard therefore is

$$h(T) = \frac{h_1 \cdot e^{-h_1 \cdot T} + h_2 \cdot e^{-h_2 \cdot T}}{e^{-h_1 \cdot T} + e^{-h_2 \cdot T}}$$

## References

- 1 Suissa S, Dell'aniello S, Ernst P. Long-term natural history of chronic obstructive pulmonary disease: severe exacerbations and mortality. *Thorax* 2012;67:957–63. doi:10.1136/thoraxjnl-2011-201518
- 2 Hurst JR, Donaldson GC, Quint JK, et al. Temporal clustering of exacerbations in chronic obstructive pulmonary disease. *Am J Respir Crit Care Med* 2009;179:369–74. doi:10.1164/rccm.200807-1067OC
- 3 Hurst JR, Vestbo J, Anzueto A, et al. Susceptibility to exacerbation in chronic obstructive pulmonary disease. *N Engl J Med* 2010;363:1128–38. doi:10.1056/NEJMoa0909883
- 4 R Development Core Team. *R: A Language and Environment for Statistical Computing*. Vienna, Austria: : R Foundation for Statistical Computing 2011. <http://www.r-project.org/> (accessed 16 Nov2009).
- 5 Keene ON, Calverley PMA, Jones PW, et al. Statistical analysis of exacerbation rates in COPD: TRISTAN and ISOLDE revisited. *Eur Respir J* 2008;32:17–24. doi:10.1183/09031936.00161507
- 6 R: Create survival curves. <http://stat.ethz.ch/R-manual/R-devel/library/survival/html/survfit.html> (accessed 21 Dec2013).
- 7 R: Scatter Plot Smoothing. <http://stat.ethz.ch/R-manual/R-patched/library/stats/html/lowess.html> (accessed 21 Dec2013).
- 8 R: Fit Proportional Hazards Regression Model. <http://stat.ethz.ch/R-manual/R-patched/library/survival/html/coxph.html> (accessed 21 Dec2013).
- 9 Keene O, Roger J. Temporal clustering of exacerbations in chronic obstructive pulmonary disease. *Am J Respir Crit Care Med* 2010;182:984–985; author reply 985. doi:10.1164/ajrccm.182.7.984
- 10 R: Regression for a Parametric Survival Model. <http://stat.ethz.ch/R-manual/R-devel/library/survival/html/survreg.html> (accessed 21 Dec2013).

```

#R code
#Last change on 22/12/2013.
#By Mohsen Sadatsafavi

#The following packages is required. Please ensure it is installed.
library(survival)

#Figure 1
h1<-0.5
h2<-3.5
time<-1*(1:1000)/1000
h<-(exp(-h1*time)*h1+exp(-h2*time)*h2)/(exp(-h1*time)+exp(-h2*time))
plot(c(0,max(time)),c(min(0,h1,h2,h),max(h1,h2,h)),type='n',xlab="time",y
lab="hazard")
lines(time,rep(h1,length(time)),lwd=2,lty=1)
lines(time,rep(h2,length(time)),lwd=2,lty=1)
lines(time,h,lwd=3)

#genTTE generates time-to-event (tte) intervals;
genTTE<-function(n,th,shape,scale,heterogeneity=1)
{
  #bh=baseline hazard.
  if(heterogeneity==1) bh<-rgamma(n,shape,scale) else bh<-
  rep(shape/scale,n)

  #th<-rexp(n,1/(th*bh/mean(bh)))
  th<-rep(th,n)

  #A list containing the simulated time to events (ttes) in a lis
  lttes<-as.list(rep(-1,n))

  #event generator.
  maxN<<-0;
  for(i in 1:n)
  {
    #message(paste("i is ",i))
    ttes<-NULL;
    time<-0
    while(time<th[i])
    {
      if(bh[i]>0) tte<-rexp(1,bh[i]) else tte<-Inf
      if(is.nan(tte)) tte<-Inf
    }
  }
}

```

```

        #message(tte)
        time<-time+tte
        if(time<th[i])    ttes<-c(ttes,tte)
    }
    if(!is.null(ttes))
        lttes[[i]]<-ttts
    else
        lttes[[i]]<-1
    maxN<-max(maxN,length(ttes))
}

#Changes the tte list into a matrix (mttes)
mttes<-matrix(nrow=n,ncol=maxN)
for(i in 1:n)
{
    if(lttes[[i]][1]!=-1)
        mttes[i,1:length(lttes[[i]])]<-lttes[[i]]
}

return(mttts)
}

####Suisse et al.;

n<-100000
TH<-10
smooth<-1  #set to 0 if you do not want to smooth the hazard function

#Set heterogeneity to 0 to repeat the results with between-individual
#heterogeneity removed;
mttes<-genTTE(n,TH,0.2,0.2,heterogeneity=1)

#dt=time bin.
dt<-TH/1000

#x and y will hold data for the plot
x<-NULL
y<-x

#mst will hold the data for median survival times
mst<-rep(0,20)

#cumulative time
cTime<-0

```

```

#pEvent=previous number of events. The curve is generated segment by
segment per pEvent values. Up to 20 events are modelled. Change this if
changing the time horizon but otherwise does not matter
for(pEvent in 0:20)
{
  if(pEvent>0)
  {
    ttes<-mttes[!is.na(mttes[,pEvent]),pEvent+1]
    if(pEvent==1)
    {
      tl<-mttes[!is.na(mttes[,1]),1]
    }
    else
    {
      tl<-rowSums(mttes[!is.na(mttes[,pEvent]),1:pEvent])
    }
    events<-(!is.na(ttes))*1
    ttes[is.na(ttes)]<-TH-tl[is.na(ttes)]
  }
  else
  {
    ttes<-mttes[,1]
    events<-(!is.na(ttes))*1
    ttes[is.na(ttes)]<-TH
  }
}

#Fits a K-M model to estimate the median survival time and baseline
hazard
sf<-survfit(formula=Surv(ttes,events)~0)
mst[pEvent+1]<-sf$time[which.min(abs(sf$surv-0.50))]
h0<-data.frame(hazard = -log(sf$surv), time = sf$time) #see
documentation for basehaz to see why this estimated baseline hazard

#generate this segment of the curve
time<-0
tempX<-rep(0,mst[pEvent+1]/dt+1)
tempY<-tempX
counter<-1
while(counter<=length(tempX) )
{
  time<-(counter-1)*dt
  tempX[counter]<-time
  tmp<-which.min(abs(h0[,"time"]-time))
  tempY[counter]<-h0[tmp,"hazard"]
  counter<-counter+1
}

tempX<-tempX[-1]
tempY<-(tempY[-1]-tempY[-length(tempY)]) /dt

#Using Lowess for smoothing the baseline hazard
if(smooth)
{
  sm = lowess(tempX,tempY)
  xx<-sm$x
  yy<-sm$y
}
else
{
  xx<-tempX
}

```

```

        yy<-tempY
    }
x<-c(x,xx+cTime)
y<-c(y,yy)

cTime<-cTime+mst[pEvent+1]
}

plot(x,y/365*10000,type='l',xlim=c(0,12),ylim=c(0,100),
xlab="time",ylab="hazard (rate) per 10000 per day")

#HR generator (for the table in the appendix), as described by Suissa et
al. using a cox model with the number of previous exacerbation as the
time-dependent variable.
#Suggest using the first 10000 observations for this part as with
n=100000 it will take forever to fit the cox model.

n<-10000
mttes<-mttes[1:n,]

maxNEvent<-10
#Create time-dependent data;
x<-cbind(mttes,NA)
template<-t(as.matrix(rep(0,5)))
tdttes<-NULL
#creates (start, stop) data format for time-dependent Cox model;
colnames(template)<-c("id","start","stop","event","nEvent")
for(i in 1:n)
{
  time<-0
  j<-1
  nEvent<-0
  while(j<=maxN && j<=maxNEvent && !is.na(x[i,j]))
  {
    if(!is.na(x[i,j]))
    {
      template[, "id"]<-i
      template[, "start"]<-time
      time<-time+x[i,j]
      template[, "stop"]<-time
      template[, "event"]<-1
      template[, "nEvent"]<-nEvent
      nEvent<-nEvent+1
      tdttes<-rbind(tdttes,template)
    }
    j<-j+1
  }
  if(j<maxNEvent)
  {
    template[, "id"]<-i
    template[, "start"]<-time
    template[, "stop"]<-TH
    template[, "event"]<-0
    template[, "nEvent"]<-nEvent
  }
}

```

```

        tdttes<-rbind(tdttes,template)
    }
}

cm<-
coxph(Surv(start,stop,event)~factor(nEvent),data=as.data.frame(tdttes))
summary(cm)

#Hurst et al.;

#Note: Reduce the sample size if survfit fails to fit;
n<-100000
TH<-3

#Set heterogeneity to 0 to repeat the results with between-individual
heterogeneity removed;
mttes<-genTTE(n,TH,2.72,1,heterogeneity=1)

#Merge evrything in a vector as with Hurst et al. plus add the censored
intervals as well.
ft<-rowSums(mttes,na.rm=TRUE)
vttes<-as.vector(mttes)
vttes<-vttes[!is.na(vttes)]
events<-rep(1,length(vttes))
vttes<-c(vttes,TH-ft)
events<-c(events,rep(0,length(ft)))
sr<-survreg(Surv(vttes,events)~1,dist='weibull')
#Note that the shape of Weibull is 1/scale in the survreg see
http://stat.ethz.ch/R-manual/R-devel/library/survival/html/survreg.html
for details;
message(paste('shape is',1/sr$scale))

#See DISCUSSION section of Hurst et al.: subgroup analysis by dividing
the sample around median exacerbation rate;
nEvents<-rowSums(mttes*0+1,na.rm=TRUE)
cutoff<-median(nEvents)

```

```
#low exacerbators
vttes<-as.vector(mttess[nEvents & nEvents<cutoff,])
vttes<-vttes[!is.na(vttes)]
events<-rep(1,length(vttes))
vttes<-c(vttes,TH-ft)
events<-c(events,rep(0,length(ft)))
sr<-survreg(Surv(vttes,events)~1,dist='weibull')
message(paste('shape is',1/sr$scale))

#high exacerbators
vttes<-as.vector(mttess[nEvents>=cutoff,])
vttes<-vttes[!is.na(vttes)]
events<-rep(1,length(vttes))
vttes<-c(vttes,TH-ft)
events<-c(events,rep(0,length(ft)))
sr<-survreg(Surv(vttes,events)~1,dist='weibull')
message(paste('shape is',1/sr$scale))
```