



survHE: Survival Analysis for Health Economic Evaluation and Cost-Effectiveness Modeling

Gianluca Baio

University College London

Abstract

Survival analysis features heavily as an important part of health economic evaluation, an increasingly important component of medical research. In this setting, it is important to estimate the mean time to the survival endpoint using limited information (typically from randomized trials) and thus it is useful to consider parametric survival models. In this paper, we review the features of the R package **survHE**, specifically designed to wrap several tools to perform survival analysis for economic evaluation. In particular, **survHE** embeds both a standard, frequentist analysis (through the R package **flexsurv**) and a Bayesian approach, based on Hamiltonian Monte Carlo (via the R package **rstan**) or integrated nested Laplace approximation (with the R package **INLA**). Using this composite approach, we obtain maximum flexibility and are able to pre-compile a wide range of parametric models, with a view of simplifying the modelers' work and allowing them to move away from non-optimal work flows, including spreadsheets (e.g., Microsoft **Excel**).

Keywords: survival analysis, health economic evaluation, probabilistic sensitivity analysis, R.

1. Introduction

Broadly speaking, the objective of publicly funded health care systems – e.g., those in the United Kingdom (UK), Canada, Australia and many other countries around the world – is to maximize health gains across the general population, given finite monetary resources and a limited budget. Bodies such as the National Institute for Health and Care Excellence (NICE) in the UK provide guidance on decision-making on the basis of health technology assessment (HTA) and cost-effectiveness analysis (NICE 2013). These cover a suite of analytical approaches for combining costs and consequences of intervention(s) compared to a control or status quo, the purpose of which is to aid decision-making associated with resource allocation. To this aim, much of the recent research has been oriented towards building the health economic evaluation on sound and advanced statistical decision-theoretic foundations,

arguably making it a branch of applied statistics (Willan and Briggs 2006; Briggs, Sculpher, and Claxton 2006).

Interventions that impact upon survival form a high proportion of the treatments appraised by NICE (Latimer 2011). Interestingly, in order to quantify accurately the economic benefits of a new intervention, it is necessary to estimate the mean survival time (rather than usual summaries, such as the median time). Thus, it is necessary to extrapolate the observed survival curve (which often only covers a limited time frame and is subject to high degree of censoring) to a longer time horizon. Consequently, a parametric approach to the survival analysis is usually followed and it is recommended by a highly influential NICE Decision Support Unit (DSU) technical document (Latimer 2011).

In general terms, a Bayesian approach to statistical inference is particularly helpful in economic evaluation, because it allows to fully characterize uncertainty in the model parameters. This in turn is a fundamental component of the cost-effectiveness analysis, which are often made by several potentially correlated “modules”, which may be informed by different and diverse sources of evidence. For this reason, it is crucial to assess the impact of this uncertainty on the decision-making process (Briggs *et al.* 2006; Baio and Dawid 2011; Baio 2012), a process termed *probabilistic sensitivity analysis* (PSA).

In addition, the clinical evidence used to inform the survival analysis for HTA can be limited, thus emphasizing the problem of uncertainty in the extrapolation of the survival curves. The inclusion of prior information (e.g., to “anchor” cancer patients’ survival well below that of the healthy population, or from evidence on drugs with similar therapeutic mechanisms), which is instrumental to a Bayesian analysis, would again improve the model performance. However, in addition to the need of specifying suitable prior distributions that are consistent with the information available for the case at hand, Bayesian models for survival analysis fitted using standard Markov chain Monte Carlo (MCMC) algorithms can be computationally intensive and sometimes have problems with convergence. Moreover, existing Bayesian software, notably BUGS (Lunn, Spiegelhalter, Thomas, and Best 2009) or JAGS (Plummer 2003), require the full specification of the modeling assumptions, which often proves an insurmountable barrier to implementation for inexperienced researchers/modelers.

Perhaps for this reason, frequentist methods to perform survival analysis with an emphasis in economic evaluations represent the industry standard. In particular, survival analysis is often embedded in health economic evaluations using a multi-step/multi-software approach. Firstly, estimates from a survival model are obtained. These are sometimes obtained by published clinical studies presenting point and interval estimates for the model parameters; alternatively, survival models are fitted to (pseudo)individual level data (see Section 9.1). This part is by necessity based on proper statistical software (e.g., Stata, StataCorp 2019; SAS, SAS Institute Inc. 2013; or R, R Core Team 2020). Secondly, modelers produce a set of simulations for the model parameters using tools such as Cholesky decomposition and Monte Carlo (MC) procedures (Briggs *et al.* 2006). Finally, these are used to produce a large number of survival curves, which are fed to the economic model (for example to determine the benefits of a given intervention). The last two steps are typically performed using a spreadsheet, almost invariably Microsoft **Excel**.

From both the scientific and the practical point of view, this process is less than ideal. Firstly, in the MC simulations, potential correlation among the model parameters is only approximated and not fully accounted for, potentially introducing bias in the estimate for the survival

curves and thus in the outcome of the decision-making process. A full Bayesian approach, e.g., based on MCMC, would eliminate this issue because the inference would be produced directly on the full joint distribution of the model parameters. While it is in theory possible to build a frequentist model by constructing a multivariate joint likelihood of all the model parameters, this is in practice not very common. Incidentally, this exercise would be likely to increase the theoretical and computational complexity, thus *de facto* eliminating the classical objections to a Bayesian approach, in favor of a “simpler” one based on maximum likelihood (ML).

The limitations of spreadsheets calculators such as Microsoft **Excel**, in terms of statistical modeling (and particularly in survival analysis), are increasingly often recognized in the health economics literature (Baio and Heath 2016; Williams, Lewsey, Briggs, and Mackay 2017; Krijkampet, Alarid-Escudero, Enns, Jalal, Hunink, and Pechlivanoglou 2018; Incerti, Thom, Baio, and Jansen 2019). In fact, it is no accident that a number of R packages have recently been developed to use in cost-effectiveness analysis, notably including among others **BCEA** (Baio, Berardi, and Heath 2017), **SAVI** (Strong, Oakley, and Brennan 2014), **heemod** (Filipović-Pierucci, Zarca, and Durand-Zaleski 2017), **hesim** (Incerti and Jansen 2020), **EVSI** (Heath, Manolopoulou, and Baio 2017) and **missingHE** (Gabrio, Mason, and Baio 2019). With this in mind, the objective of this work is to develop a suite of functions and tools for the freely available statistical software R, specifically designed for the needs of modelers using survival analysis results to build extensive models for health economic evaluation.

2. The R package **survHE**

survHE is an R package specifically designed to aid in the process of survival analysis for health economic evaluation and cost-effectiveness analysis. In fact, **survHE** can be actually considered as a wrapper for three other R packages; the first one, **flexsurv** (Jackson 2016), is in turn a general-purpose tool for performing several types of survival analysis using *maximum likelihood estimates* (MLEs). The second one, **rstan** (Carpenter *et al.* 2015), is a relatively new R package that can be used to perform Bayesian analysis using *Hamiltonian Monte Carlo* (HMC). This is a form of MCMC algorithm, which can be used to produce samples from a joint posterior distribution of a set of model parameters and unobserved quantities. Finally, the third one, **INLA** (Martins, Simpson, Lindgren, and Rue 2013), can be used to perform fast Bayesian computations (on a limited set of survival models) using *integrated nested Laplace approximation* (Rue, Martino, and Chopin 2009). In a sense, thus, **survHE** is a very simple tool that specializes functions from other relevant packages and builds some other specific commands to simplify and standardize the process of using survival data in health economic evaluation.

Package **survHE** (Baio 2020) is available from the Comprehensive R Archive Network (CRAN) at <https://CRAN.R-project.org/package=survHE> and there is also a development version maintained under GitHub (<https://github.com/giabaio/survHE>). The design for **survHE** is modular; sets of function have specific objectives, which can be broadly categorized as “*Data preparation*”, “*Model fitting and assessment*” and “*PSA and extrapolation*”, as shown schematically in Figure 1. These are described in detail in the following sections.

Depending on the user’s instruction, **survHE** maps internally to different code, which calls either **flexsurv**, **rstan** or **INLA** in the background to produce the relevant estimates. Once

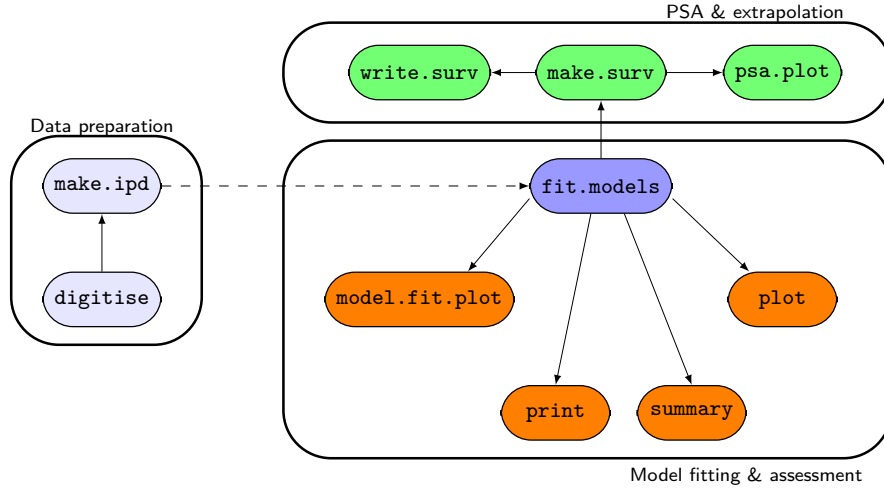


Figure 1: A schematic of the design for **survHE**. The rounded boxes represent the “modules”, a set of functions used to perform a specific task. For example, the functions `digitise` and `make.ipd` can be used to process digitized data from published papers and recreate an individual patient dataset. This can in turn be fed to `fit.models`, which estimates the parameters for a given survival model. The output of this function can be processed to assess model fit, or produce relevant summaries and plots. In addition, it can be fed to the function `make.surv`, which simulates the survival curves and effectively performs PSA. The output of this process can be plotted using the function `psa.plot`, or exported to a spreadsheet via the function `write.surv`.

these are obtained, the output is standardized and returned in a form that is comparable across the inferential methods. In other words, for the set of models that are implemented in **survHE**, practitioners can take full advantage of the modeling capabilities of the three packages and analyze survival data under both a frequentist or a computationally efficient Bayesian approach, using the same syntax. More importantly, **survHE** also includes specific built-in functions to simplify and standardize the process of using survival data in health economic evaluation.

2.1. Modeling framework

The general modeling framework considered in **survHE** can be described as follows. The observed data are at least made by the pair (t_i, d_i) , for $i = 1, \dots, n$, where $t_i > 0$ is the observed time at which the event under study occurs and d_i (for “dummy” variable) is an event indicator, taking value 1 if the event actually happens, or 0 when the i th individual is “censored”. If $d_i = 1$, then t_i is indeed observed; conversely, if $d_i = 0$, we do not know whether the event actually occurs – it may in the future, but we just do not have this information. Consequently, when $d_i = 0$, then the observed t_i does not represent the true “survival time”. Notice here that we consider for the sake of simplicity “right censoring”, which is the most common for applications in health economics.

The observed data t_i are modeled using a suitable probability distribution characterized by a density $f(t_i | \theta)$, as a function of a vector of relevant parameters θ . This can be linked to

the *survival function*

$$S(t) = 1 - F(t) = 1 - \int_0^t f(u | \boldsymbol{\theta}) du,$$

indicating the probability of an individual surviving up to time t , as well as to the *hazard function*

$$h(t) = \frac{f(t)}{S(t)},$$

which quantifies the instantaneous risk of experiencing the event. In the presence of censoring, the resulting log-likelihood function is modified to account for the possibility of partially observed data (in correspondence with censoring) and is expressed as

$$\log \mathcal{L}(\boldsymbol{\theta}) = \sum_{i=1}^n [d_i \log h(t_i) + \log S(t_i)]. \quad (1)$$

This basically models the risk of experiencing the event at any time point t , conditionally on the fact that the i th unit has in fact survived up to that time point; if they have not, then the probability of experiencing the event again is 0.

When formulating a parametric survival model, we need to specify the form of the probability distribution assumed to describe the underlying uncertainty in the observed times. As mentioned above, it is good practice to test a set of (more or less) plausible parametric models for the survival data. This is the procedure recommended by the NICE-DSU technical document (Latimer 2011).

In general terms, we can specify the vector of relevant parameters as $\boldsymbol{\theta} = (\mu(\boldsymbol{x}), \alpha(\boldsymbol{x}))$. In this notation, consistent with **flexsurv**, we consider: a vector of potential covariates \boldsymbol{x} (e.g., age, sex, trial arm, etc.); a *location* parameter $\mu(\boldsymbol{x})$, which indicates the mean or the scale of the probability distribution; and a (set of) ancillary parameters $\alpha(\boldsymbol{x})$, which describe the shape or variance of the distribution. While it is possible for both μ and α to explicitly depend on the covariates \boldsymbol{x} , usually the formulation is simplified to assume that these only affect directly the location parameter.

In addition, since $t > 0$, we typically use a generalized linear formulation

$$g(\mu_i) = \beta_0 + \sum_{j=1}^J \beta_j x_{ij} [+ \dots] \quad (2)$$

to model the location parameter. The function $g(\cdot)$ is typically the logarithm; notice that here we slightly abuse the notation and omit the dependence of μ_i on \boldsymbol{x} . Generally speaking, (2) can be extended to include additional terms – for instance, we may want to include random effects to account for repeated measurements or clustering. We indicate this possibility using the $[+ \dots]$ notation and highlight the fact that this is rather straightforward in a Bayesian context (particularly for **INLA**, but also for **rstan**).

The objective of the statistical analysis is the estimation of the parameters $\boldsymbol{\theta}$, which can then be used to obtain estimates for all the relevant quantities (e.g., the survival function), which are in turn used in the economic modeling, typically through a state-transition structure (Siebert *et al.* 2012). In a frequentist setting, the estimation procedure concerns some

relevant statistics, i.e., functions of the observed data and is performed via ML estimation. Conversely, in a full Bayesian setting, the parameters are directly modeled using a prior probability distribution, which is updated by the observed data into a posterior. It is this posterior distribution that is the object of the inferential process. Thus, when using a Bayesian framework, the model needs to be completed by specifying suitable prior distributions for the parameters θ .

Assuming that the location parameter is specified using a linear predictor form, on the scale determined by $g(\cdot)$ and as a function of J covariates, we can model $\beta = (\beta_0, \beta_1, \dots, \beta_J) \stackrel{\text{iid}}{\sim} \text{Normal}(\mu_\beta, \sigma_\beta)$, where, in general, we use the notation σ to indicate a standard deviation. Note that **survHE** expands any categorical covariates to a set of dummy variables: so if a covariate has four categories, in line with R notation, **survHE** considers three binary indicators. Thus the profile (0, 0, 0) indicates the first (reference) category, while the profiles (1, 0, 0), (0, 1, 0) and (0, 0, 1) indicate the second, third and fourth category, respectively. In **survHE**, the number of covariates J depends on this full expansion of the design matrix.

In both its Bayesian versions, **survHE** assumes by default $\mu_\beta = 0$ and $\sigma_\beta = 5$ for the models in which the linear predictor is defined on the log scale and $\sigma_\beta = 100$ for those in which the linear predictor is defined on the natural scale and thus $g(\cdot)$ is the identity function. This amounts to specifying a “minimally informative” prior on the regression coefficients that determine the location parameter – in other words, we are not including strong prior information in this aspect of our model. The observed data (and the censoring structure) will be mainly driving the update to the posterior distribution. When genuine prior knowledge is present, e.g., about the likely size of a treatment effect, it is possible to modify these priors to encode the information in the model formulation.

Notice that, particularly in case where the sample size is small or censoring is large, even these minimally informative priors generally contribute to stabilize the inference, by placing most of the probability mass for the regression coefficients on a relatively small interval. For instance, if $\beta \sim \text{Normal}(0, 5)$, we are then implying a negligible prior probability that $\beta > 10$ (recall that in this case, β is on the log scale). If the data suggested that the effect of a particular covariate is in fact large, the model would respond by updating the prior distribution accordingly. However, because large effects are not very likely in the prior, the model would respond in a “skeptical” way, thus counter-acting the potential bias implied by taking at face value the signal from the limited amount of information present in the data.

As for the ancillary parameter, the choice of prior depends on the specification of the probability distribution selected to model the observed data. Table 1 shows a summary of the models directly implemented in **survHE**. In each, by default we specify minimally informative priors on the relevant parameters; for example, in the Weibull model, we define $\alpha \sim \text{Gamma}(a, b)$, for given values of a, b .

All the models presented in Table 1 are available using MLE and HMC as inferential engines. On the other hand, **INLA** currently only handles Exponential, Weibull, log-Normal and log-Logistic models. The names presented in the rightmost column can be used when calling **survHE**. We mostly follow the notation of **flexsurv**, but also allow for some specific differences in the **INLA** notation: for example, the log-Logistic distribution can be referred to using **flexsurv** (`llogis`), or **INLA** notation (`loglogistic`). **survHE** will internally map the different strings of text and select the correct routine.

Notice also that, in line with **flexsurv** and **INLA**, **survHE** allows the two versions of the Weibull

Data model	Location parameter	Ancillary parameter*	survHE name
$t_i \sim \text{Exponential}(\mu_i)$	Rate: $\mu_i = \exp\left(\beta_0 + \sum_{j=1}^J \beta_j x_{ij}\right)$	—	exp, exponential
$t_i \sim \text{Weibull}(\mu_i, \alpha)$	Scale: $\mu_i = \exp\left(\beta_0 + \sum_{j=1}^J \beta_j x_{ij}\right)$	Shape: $\alpha \sim \text{Gamma}(0.1, 0.1)$	weibull, wei weibullPH, wph
$t_i \sim \text{logNormal}(\mu_i, \alpha)$	log-mean: $\mu_i = \beta_0 + \sum_{j=1}^J \beta_j x_{ij}$	log-sd: $\alpha \sim \text{Uniform}(0, 5)$	lnorm, lno, lognormal
$t_i \sim \text{logLogistic}(\mu_i, \alpha)$	Rate: $\mu_i = \exp\left(\beta_0 + \sum_{j=1}^J \beta_j x_{ij}\right)$	Shape: $\alpha \sim \text{Gamma}(0.1, 0.1)$	llogis, llo, loglogistic
$t_i \sim \text{Gamma}(\mu_i, \alpha)$	Rate: $\mu_i = \exp\left(\beta_0 + \sum_{j=1}^J \beta_j x_{ij}\right)$	Shape: $\alpha \sim \text{Gamma}(0.1, 0.1)$	gamma, gam
$t_i \sim \text{Gompertz}(\mu_i, \alpha)$	Rate: $\mu_i = \exp\left(\beta_0 + \sum_{j=1}^J \beta_j x_{ij}\right)$	Shape: $\alpha \sim \text{Gamma}(0.1, 0.1)$	gompertz, gom
$t_i \sim \text{Gen Gamma}(\mu_i, \boldsymbol{\alpha})$	Location: $\mu_i = \beta_0 + \sum_{j=1}^J \beta_j x_{ij}$	$\boldsymbol{\alpha} = (\sigma, q)$ Scale: $\sigma \sim \text{Gamma}(0.1, 0.1)$ Shape: $q \sim \text{Normal}(0, 100)$	gengamma, gga
$t_i \sim \text{Gen F}(\mu_i, \boldsymbol{\alpha})$	Location: $\mu_i = \beta_0 + \sum_{j=1}^J \beta_j x_{ij}$	$\boldsymbol{\alpha} = (\sigma, q, p)$ Scale: $\sigma \sim \text{Gamma}(0.1, 0.1)$ Shape(1): $\log(p) \sim \text{Normal}(0, 0.5)$ Shape(2): $q \sim \text{Normal}(0, 2.5)$	genf, gef

Table 1: The distributions supported by **survHE**. *Distributions presented for the ancillary parameters are the default used by **survHE**.

model, i.e., using an accelerated time failure (AFT, `weibull` or `wei`) or a proportional hazard (PH, `weibullPH` or `wph`) parameterization. When the interest is in estimating the effect of some covariates on the survival time, these two versions yield of course different estimates. However, in the case of health economic evaluation, the interest is really in producing an estimate of the distribution of the survival curves. **survHE** internally maps the estimated coefficients to the correct transformation so as to estimate $S(t)$ correctly, irrespective of the original parameterization. For example, when the AFT parameterization is used, then $S(t) = \exp\left(-\left(\frac{t}{\mu_i}\right)^\alpha\right)$, while for the PH parameterization $S(t) = \exp(-\mu_i t^\alpha)$.

The “original” parameterization of the Generalized F and Generalized Gamma distributions (respectively `genf.orig` and `gengamma.orig`) are currently available only when the estimation is performed using MLE in **flexsurv** (this is mostly for backward compatibility).

2.2. Example

In the following, we use a running example to present the features of **survHE**. Suppose that the user has a suitable dataset, perhaps obtained from a trial, in which data are recorded for the time at which observations are made, an “event” indicator taking value 1 if the clinical outcome of interest (e.g., progression to a cancerous state, or death) has actually occurred at that time and 0 if the individual has been censored (e.g., we have not observed any event at the end of follow up), as well as an arm indicator, specifying whether the individual to whom the data refer belongs in the control or the active treatment arm of the trial. We also have information about the individuals’ age, sex, socio-economic circumstances (as measured by an “Index of Multiple Deprivation”, typically referred to as IMD; [Noble, Wright, Smith, and Dibben 2006](#)) and ethnic background. Other variables may be observed, e.g., relevant co-morbidity.

For the moment we consider the simple case in which the data are available in the R workspace as a data-frame (say, `data`) that can be partially visualized using the following command:

```
R> rbind(head(data), tail(data))
```

to show the first 6 and the last 6 rows:

ID_patient	time	event	arm	sex	age	imd	ethnic
1	0.03	0	0	1	34	1	4
2	0.03	0	0	0	29	4	2
3	0.92	0	0	0	30	4	5
4	1.48	0	0	0	31	5	1
5	1.64	0	0	0	27	4	1
6	1.64	0	0	0	37	5	4
362	13.97	1	1	1	36	4	1
363	14.56	1	1	0	26	3	4
364	14.56	1	1	1	23	2	2
365	14.85	1	1	0	31	3	3
366	16.07	1	1	1	36	3	1
367	18.16	1	1	0	44	4	5

The dataset consists of 367 individuals in total, grouped in two arms (here `arm = 0` indicates the controls and `arm = 1` indicates the active treatment; `sex = 0` or `1` indicates Male and

Female, respectively; `imd` takes values $1, \dots, 5$ to indicate increasing levels of small area deprivation; and `ethnic` = $1, \dots, 5$ groups various ethnic backgrounds, for instance in line with census data).

2.3. Modeling survival data: To be or not to be (Bayesian) ...?

By relying on either `flexsurv`, `INLA` or `rstan`, `survHE` can fit models under the frequentist or Bayesian framework.

As already discussed, a Bayesian approach can be beneficial to propagate the underlying uncertainty in all the model parameters in a principled way. In the specific case of survival analysis, this is often helpful because it allows the formal inclusion of external information, e.g., in the form of data coming from population registries or, at the very least, expert knowledge. This is not a major benefit when the underlying time-to-event data are mature (e.g., the follow up and the underlying event rate of the study are such that most individuals are fully observed – and censoring does not affect a large proportion of the subjects). However, particularly with the recent development of immuno-therapies, the case where the observed data do not even reach the median survival time is increasingly frequent. This is clearly a major limitation, when the main objective is to extrapolate the survival curve way beyond the observed follow up. As mentioned earlier, in such cases, it is fundamental to anchor the extrapolation of the survival curves, using external information – for example by combining the trial data with life-tables or registries to ensure that the extrapolation does not artificially favor a given treatment. A Bayesian approach naturally lends itself to this aim, by allowing the modeler to include genuine prior information, e.g., on the probability that the average survival time does not exceed some set threshold (Benaglia, Jackson, and Sharples 2015; Jackson *et al.* 2017; Guyot, Ades, Beasley, Lueza, Pignon, and Welton 2017).

On the other hand, in addition to the need of specifying suitable prior distributions that are consistent with the information available for the case at hand, Bayesian models for survival analysis fitted using MCMC can be computationally intensive and sometimes have problems with convergence; perhaps for this reason, often practitioners use MLE-based routines to obtain relevant estimates from the survival component of the wider economic model. These are usually simple to obtain. In order to deal with PSA, `flexsurv` uses bootstrap (based on multivariate Normal distributions), which may be a good approximation of the underlying full joint posterior distribution of the survival parameters.

A good compromise between frequentist and fully Bayesian models is provided by `INLA`, which is effectively an alternative method of performing Bayesian inference. By using a specific (albeit rather general), clever model specification and an approximation algorithm, `INLA` typically requires a computational time that is very close to that of MLE-based routines, while also estimating an approximation to the full joint posterior distribution for the model parameters. However, in general terms, it is a bit more complex to embed `INLA` within a more complex model; in addition, currently, `INLA` can only fit a limited number of survival models.

For these reasons, we chose to design `survHE` around these three approaches: ideally, we would build the whole economic model under a Bayesian framework and take full advantage of the flexibility provided by MCMC estimation – this would be naturally obtained by using `rstan`. Note that we choose `rstan` over other software such as `OpenBUGS` or `JAGS`; the reason for this is that HMC often proves a superior mode of inference compared to Gibbs sampling (the MCMC algorithm upon which `OpenBUGS` and `JAGS` are based). In particular, for the

specific set of models that we consider here, HMC proves faster and more reliable in terms of convergence than Gibbs sampling. In addition, **rstan** models can be “pre-compiled” and thus the computational time required is totally devoted to sampling from the relevant posterior distributions, which makes a full Bayesian approach more competitive, in comparison to MLE-based methods.

However, despite these very useful features of **rstan**, we acknowledge that at times Bayesian survival models can still be challenging from the point of view of computation (especially with large datasets and for particularly complex model structures). In these cases, **INLA** is very helpful, for the range of distributions it supports and going forward, to expand the modeling to flexible structures (e.g., including smooth random effects); similarly, particularly in specific settings, some practitioners may want to use a standard approach to statistical inference and MLEs.

The combined use of **survHE** and its “dependencies” allows all of these options in a unified framework. More importantly, irrespective of the way in which the inference on the survival model is performed, **survHE** has a set of built-in functions that can be used to produce a standardized post-processing of the results, for their inclusion in the economic model and PSA.

3. MLE via flexsurv

survHE allows the user to define in R a vector of model names (in the format that **flexsurv** or **INLA** can recognize). We could for instance decide that we want to consider the Exponential, Weibull, Gamma, log-Normal, log-Logistic and Generalized Gamma models for our analysis. We can do this in R by using the following commands.

```
R> mods <- c("exp", "weibull", "gamma", "lnorm", "llogis", "gengamma")
```

This syntax defines a vector of model names to be used by **flexsurv** in the fitting process. These *must* adhere with **survHE** convention, as specified above and in Table 1.

At this point, we are almost ready to actually perform the survival analysis using the 6 models specified above; before we can do this, however, we need to specify the model “formula”, for example as the following.

```
R> formula <- Surv(time, event) ~ as.factor(arm)
```

This creates an object instructing R to analyze the data using a survival model in which the only covariate is the treatment arm, interpreted as an R “factor” (i.e., a categorical variable). Of course, it is possible to use an even less structured version of the linear predictor by simply defining `formula <- Surv(time, event) ~ 1`, which would fit a model with no covariates (i.e., intercept only). Other covariates can also be added to the formula.

The **survHE** function `fit.models` can be used to actually perform this analysis in batches, e.g., by typing the command:

```
R> m1 <- fit.models(formula = formula, data = data, distr = mods)
```

The function `fit.models` takes as mandatory inputs the analysis formula, the name of the dataset to be used and the type of distribution(s) to be fitted. Just like in this case, this may be a vector, in which case `fit.models` will store all the different analyses in the resulting object.

Executing the command above creates an object `m1` of class ‘`survHE`’, in which the results of the survival analyses are stored for each parametric model considered. The usual R command

```
R> names(m1)
```

returns the names of the several elements in the list.

```
[1] "models"          "model.fitting" "method"        "misc"
```

The object `models` is itself a list, in this case containing 6 elements (one for each of the parametric models fitted). The command

```
R> names(m1$models)
```

returns the string of names used to label each, in this case

```
[1] "Exponential"    "Weibull (AFT)" "Gamma"         "log-Normal"
[5] "log-Logistic"  "Gen. Gamma"
```

For example, the first model can be accessed using the standard R notation `m1$models[[1]]` by position (i.e., using “double square brackets”), or alternatively by using its name and `$` via `m1$models$Exponential` and can be inspected typing the commands

```
R> names(m1$models[[1]])
```

or

```
R> names(m1$models$Exponential)
```

which both output

```
[1] "call"           "dlist"          "aux"            "ncovs"
[5] "ncoveffs"      "mx"             "basepars"      "covpars"
[9] "AIC"           "data"           "datameans"     "N"
[13] "events"        "trisk"          "concat.formula" "all.formulae"
[17] "dfns"          "res"            "res.t"          "cov"
[21] "coefficients" "npars"          "fixedpars"     "optpars"
[25] "loglik"        "logliki"        "c1"             "opt"
```

The quantities included in the model objects are the standard output from `flexsurv`. Typically, the user does not need to access or manipulate them directly (that is the point of `survHE`!); in fact, other `survHE` functions will use these to produce plots or further analyses. Users familiar with R programming can however access them to post-process their results and customize even further the output provided by `survHE`.

The other elements of the object `m1` are:

- `model.fitting`: a list storing some suitable statistics that can be used to assess, well ... model fitting. These are the Akaike, Bayesian and deviance information criteria (AIC, BIC and DIC, respectively). The former two can be estimated using both the Bayesian and frequentist approach, while the latter is specific to Bayesian modeling. Thus, in this case, the R call

```
R> m1$model.fitting
```

will return the following results

```
$aic
```

```
[1] 1274.576 1203.130 1203.504 1214.984 1208.494 1204.785
```

```
$bic
```

```
[1] 1282.387 1214.846 1215.220 1226.700 1220.211 1220.406
```

```
$dic
```

```
[1] NA NA NA NA NA NA
```

Note that because we are storing the results obtained from fitting 6 models in the same object, the elements `$aic`, `$bic` and `$dic` are vectors. In general, the model associated with a lower information criterion value tends to be associated with a better fit, as we discuss in Section 6.4.

- `method`: a string describing the method used to fit the model(s). In this case the code

```
R> m1$method
```

returns the output

```
[1] "mle"
```

- `misc`: a list containing some miscellanea – these are mainly used internally by the other functions in **survHE** to do plots and tables or other calculations. Specifically, the elements of this objects are

- `time2run`: the time used to run the model(s) (in seconds);
- `formula`: the R object containing the formula used to define the model(s);
- `data`: the data-frame containing the original data used to fit the model(s);
- `model_name`: a (vector of) abbreviation(s) associated with the model(s) fitted;
- `km`: the Kaplan-Meier estimate produced automatically by **survHE** (using the function `npsurv` from the R package **rms**; Harrell Jr 2020).

4. Bayesian analysis via INLA

4.1. Integrated nested Laplace approximation

Integrated nested Laplace approximation (INLA; Rue *et al.* 2009) can be used to perform direct numerical calculation of posterior densities in a wide sub-class of Bayesian hierarchical models (called latent Gaussian models, LGMs), avoiding time-consuming Markov chain Monte Carlo simulations. The INLA implementation covers models of the form

$$\begin{aligned} y_i | \phi, \psi &\sim p(y_i | \phi, \psi), \\ \phi | \psi &\sim \text{Normal}(\mathbf{0}, \mathbf{Q}^{-1}(\psi)), \\ \psi &\sim p(\psi), \end{aligned}$$

where y_i is the observed variable, ϕ is a set of main parameters (which may and often has a very large dimension) and ψ is a set of “hyperparameters”, so that the full set of parameters is $\theta = (\phi, \psi)$. The main restrictions of the INLA formulation are the fact that the number of hyperparameters needs to be small (for computational convenience) and the form of the prior imposed on ϕ . This is a multivariate Normal distribution where the precision (i.e., inverse variance) matrix $\mathbf{Q}^{-1}(\psi)$ depends on the hyperparameters and exploits conditional independences across the parameters. This general structure is called a Gaussian Markov random field (GMRF; Rue and Held 2005).

The basic principle is to approximate the posterior density for ϕ and ψ using a series of nested Normal approximations. The algorithm uses numerical optimization to find the mode of the posterior, while the marginal posterior distributions are computed using numerical integration over the hyperparameters. INLA is a very fast method of inference and can be applied to many models that can be written in the form of LGMs – for example generalized linear models (including structured components, such as simple random effects, as well as spatial or temporal effects).

On the other hand, not all models can be easily framed within the LGM formulation. In addition, INLA’s estimates are, by definition, approximations to the full joint posterior distribution of the model parameters. Thus, there is a trade-off between the computational complexity and the accuracy of the estimation. In many general cases, INLA produces a good compromise by allowing a good level of accuracy (often comparable with simulation methods such as MCMC that, if run for long enough, are guaranteed to give the “exact” value) and running time, often in the same order as standard ML algorithms.

4.2. Using **survHE** to fit models with INLA

When fitting models using a Bayesian approach via **INLA**, **survHE** allows the user to select a vector of distributions; as mentioned above, currently, **INLA** and its R implementation allow four survival models: Exponential, Weibull, log-Normal and log-Logistic. The user can specify a vector `distr <- c("exp", "weibull", "lognormal", "loglogistic")`, or select only a subset of those models, or maybe run the `fit.models` command separately for each of them. If a distribution is specified that is not allowed in **INLA**, then **survHE** will fall back on the MLE specification and use **flexsurv** instead.

One important distinction is in the way in which **flexsurv** and **INLA** handle the names of the distributions to be fitted and the formula specifying the model. As for the latter, the correct notation is the following: “**exponential**”, “**weibull**”, “**lognormal**” and “**loglogistic**” – these do not directly match the **flexsurv** notation. As mentioned above, to avoid issues, **survHE** recodes internally the names given to the distributions. Thus, if the user specifies the additional option `method = "inla"` in the call to `fit.models`, then the string “**exp**” (which would be accepted by **flexsurv**) will be recoded to “**exponential**”, which is required by **INLA**. Similarly if a distribution that is accepted by **flexsurv** is given by the user in **INLA** terminology but the `method` is either unspecified or specifically set to “**mle**” in the call to `fit.models`, then **survHE** will recode the name of the distribution in **flexsurv** terms.

With regards to the model formula, **INLA** requires that this is specified using the notation

```
formula <- inla.surv(time, event) ~ ...
```

where `time` and `event` are the variables in which the times and the event indicator are stored

and ... is a suitable form for the combination of covariates to be used. Again, **survHE** tries to simplify the modeler's life by fixing the code provided for the formula, depending on the value specified for the option `method`. So if `method = "inla"` but the formula is specified using the **flexsurv** terminology, **survHE** will recode this to make it acceptable to **INLA**.

A suitable call to `fit.models` using **INLA** is the following.

```
m2 <- fit.models(formula = formula, data = data, distr = distr,
  method = "inla", ...)
```

where `distr` is a vector of strings containing names from the four models available in **INLA** and ... indicates optional arguments. In particular, when `method` is set to "inla", it is possible to add the following arguments.

- `dz`: defines the step length for the grid search over the hyperparameters space (default = 0.1). As mentioned above, **INLA** estimates the value of the hyperparameters in the model (e.g., the shape of a Weibull distribution), using a grid search. The finer this grid, the more accurate (but more computationally expensive!) the resulting estimates for all the parameters, e.g., for both the shape and scale of the Weibull distribution.
- `diff.logdens`: defines the difference in the log-density for the hyperparameters to stop the numerical integration used to obtain the marginal posterior distributions (default = 5). Again, this is related to how the hyperparameters are estimated in the first stage of the nested algorithm. Decreasing this difference is likely to increase the computational time, since the estimation of the hyperparameters will become more accurate.
- `control.fixed`: defines the default for the prior distributions, unless specified by the user. By default, **INLA** assumes that "fixed effects" associated with covariates are modeled using a Normal with mean 0 and variance 1000, while the overall intercept is modeled using a Normal with 0 mean and even smaller precision. **survHE** overrules this and sets the precision of the covariates in the linear predictor to $1/5^2 = 0.04$ – this is consistent with the default setting used when HMC is selected as the inferential engine.
- `control.family`: a list of options controlling the model for the observed data. If `distr.inla` is a vector, containing all or a subset of the models supported by **INLA**, this can be provided as a named list of options; for example:

```
R> m2 <- fit.models(formula = formula, data = data, distr = distr.inla,
+   method = "inla", control.family = list(weibull = list(prior =
+     "gamma", param = c(.1, .1)), lognormal = list(initial = 0.1)))
```

would instruct **INLA** to assume a Gamma(0.1,0.1) prior distribution for the shape parameter of the Weibull model and to use an initial value of 0.1 for the approximation routine of the log-Normal model. Notice that the names of the elements of the `control.family` list must adhere with the **survHE** notation shown in the right-most column of Table 1.

Using **INLA** advanced controls is very powerful and allows much versatility in fitting the models. However, some knowledge and understanding of the **INLA** syntax and philosophy is required. Guidance is provided in the **INLA** help functions as well as, for example, in [Rue](#)

et al. (2009); Blangiardo, Cameletti, Baio, and Rue (2013); Blangiardo and Cameletti (2015) and Wang, Ryan, and Faraway (2018).

Back to the running example, we may fit the models in **INLA** using the following code.

```
R> m2 <- fit.models(formula, data, c("exp", "weibull", "llogis", "lnorm"),
+   method = "inla", control.family = list(exponential = list(),
+   weibull = list(prior = "gamma", param = c(.1, .1)),
+   loglogistic = list(), lognormal = list(initial = 1)))
```

Note that it may be necessary to fiddle with the `control.family` option to successfully fit some of these models. This highlights the fact that, because the process of economic evaluation is embedded within a fundamentally statistical analysis, the modeler needs a thorough understanding of the underlying statistical tools – perhaps even more so when using a Bayesian approach, but, it is our strong belief, under any modeling circumstances.

5. Bayesian analysis via HMC

5.1. Hamiltonian Monte Carlo

Hamiltonian Monte Carlo (HMC; Radford 2011; Gelman, Carlin, Stern, Dunson, Vehtari, and Rubin 2013; Kruschke 2015; Betancourt 2012) is one of the algorithms belonging to the general class of MCMC methods. In a nutshell, HMC is based on the physical concept of Hamiltonian dynamics, which can be used to model the idealized situation of a frictionless particle sliding over a surface of varying slope. Basically, the movements of the particle depend on: (i) the *potential energy*, a function of its current location l , which is proportional to the height of the surface at the current position; and (ii) the *kinetic energy*, a function of its momentum m , depending on the mass of the particle. The way in which these movements happen can be described by a set of ordinary differential equations: this means that if we are able to compute the derivatives of these two functions and given a set of initial conditions specifying the starting location l_0 and momentum m_0 , at time t_0 , then we can predict the location and momentum of the particle at any point in time, by simulating these dynamics for a given duration.

Leaving aside all the technical difficulties, the basic intuition behind HMC is the following: the surface of interest is the *unnormalized* posterior log-density for the parameters in the model

$$\log \tilde{p}(\boldsymbol{\theta} \mid \mathbf{t}) = \log p(\boldsymbol{\theta}) + \log p(\mathbf{t} \mid \boldsymbol{\theta}). \quad (3)$$

In general, we are not in a position of knowing the target distribution $\log \tilde{p}(\boldsymbol{\theta} \mid \mathbf{t})$ exactly and in closed form¹. Moreover, even if we were, this would only be proportional to the actual posterior density for the parameters (because the expression above is computed without rescaling by the marginal log-density for the observed data \mathbf{t} ; for this reason, we use the notation \tilde{p}).

However, both log-densities on the right-hand side of (3) are known because they are part of the model specification. If we can compute the derivatives of $\log \tilde{p}(\boldsymbol{\theta} \mid \mathbf{t})$, given initial values

¹In fact, the relevant target surface in HMC is described by $-\log \tilde{p}(\boldsymbol{\theta} \mid \mathbf{t})$, but this is just a technical detail. The general argument still holds.

for the location of the parameters θ and their momentum, we can simulate Hamiltonian dynamics. As it turns out, this is extremely efficient at exploring the (negative) posterior log-density, by proposing a move to a new position that is determined by letting the “particle” slide over the density.

This means that if the current position is far away from the portion of the parametric space where most of the probability mass lies, the potential energy is large and thus the “particle” will have higher speed when sliding over a very steep surface. More specifically, unlike simpler (and far less efficient methods) such as the Metropolis algorithm, the proposed moves will not necessarily be characterized by symmetrical distributions and will tend to be pulled towards the mode of the joint posterior distribution more quickly.

Especially for models where the size of θ is very large, in comparison to other MCMC algorithms such as Gibbs sampling, HMC often proves to be very efficient in computational terms. This is mainly due to the fact that it updates the joint distribution of *all* the model parameters at once, instead of sequentially looping through each conditional distribution for one parameter given the observed data and all the other parameters. In addition, in comparison to other acceptance-rejection methods (e.g., the simpler Metropolis-Hastings), HMC is capable of adapting the proposal to non-symmetrical distributions, which translates to a faster rate of convergence to the target posterior, as well as faster decay in autocorrelation. These properties mean that Bayesian inference can be performed on models of arbitrary complexity (thus extending the limitations of INLA), at a reduced computational time and improved convergence, with respect to Gibbs samplers.

In practice, HMC can be very complex, because in addition to the specific computation of possibly complex derivatives, it requires fine tuning of several parameters. However, **rstan** provides a very clever system in which most of the adaptation is automatic. The user can still specify some of the basic inputs (and at times this is crucial to improve, or even reach convergence to the target posterior distributions), but **rstan** is a very general system to perform HMC estimation on a very wide range of models.

5.2. Using **survHE** to fit models with HMC

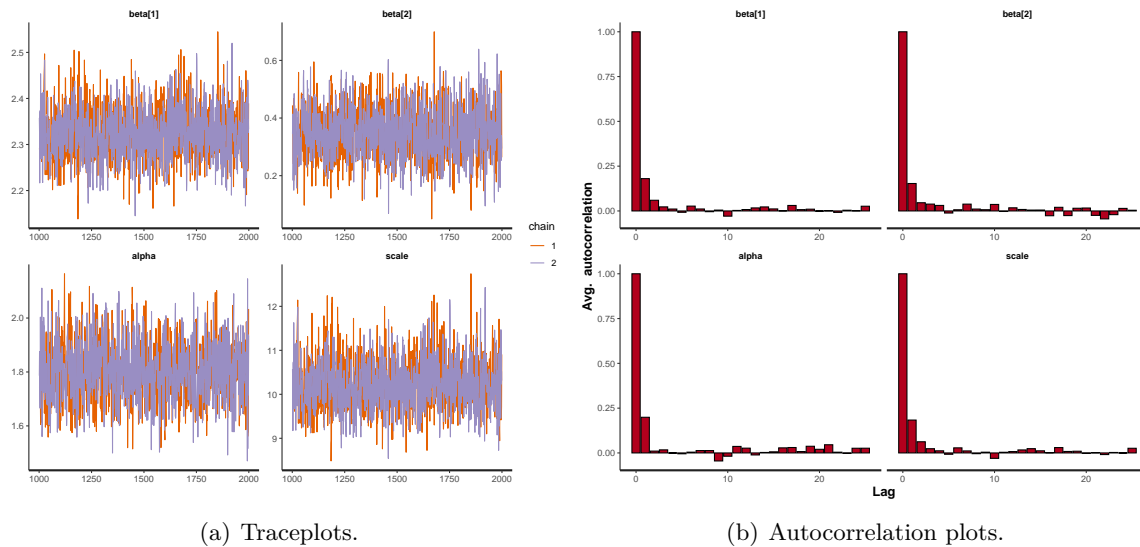
Much of the work performed by **rstan** consists in determining a set of derivatives from the model structure, that are used to apply the Hamiltonian dynamics and explore the parametric space in an efficient way. This requires a preparatory step, which **rstan** does by compiling the model in C++ (via R). This step can be quite lengthy, but interestingly it is possible to pre-compile a model – if all that changes is the data (but not the structure and the distributional assumptions), then the pre-compiled model can be used directly, thus saving substantial computation time.

This is another attractive feature of **rstan**, because it means that **survHE** can pre-compile all the standard models presented in Table 1. Thus, it is possible to estimate them by using a command such as the following.

```
R> m3 <- fit.models(formula, data, distr = mods, method = "hmc")
```

In this case, **survHE** will perform the following steps:

1. Format the original data contained in the R object **data** in a way that can be used by **rstan**;



(a) Traceplots.

(b) Autocorrelation plots.

Figure 2: Checking model convergence using the `rstan` built-in facilities, for example through inspection of the traceplots or the autocorrelation plot.

2. Select a pre-compiled model code, depending on the distributional assumptions;
3. Call `rstan` in the background to sample from the posterior distribution of the model parameters.

As with any MCMC estimation, it is important to thoroughly assess convergence – we return to this point in Section 6.1. The package `rstan` contains specialized functions to visualize the model output and assess convergence. For example the commands²

```
R> rstan::traceplot(m3$models[[2]])
R> rstan::stan_ac(m3$models[[2]])
```

would produce respectively a graph with the “traceplots” for the relevant variables, as shown in Figure 2(a) and an autocorrelation plot, shown in Figure 2(b). In this case, we can be confident that convergence is satisfactorily reached for all the variables monitored, since the traceplots show good mixing in the two chains; autocorrelation does not seem a major problem either, as the level of dependence in consecutive iterations wanes down relatively quickly. Additional model checking tools are also available in the package `shinystan` (Stan Development Team 2016), an add-on to `rstan`, which creates a web-app that the user can access locally through the default web browser. Ultimately, we reiterate that it is the utmost responsibility of the modeler to ensure that the procedure has reached convergence and the output can be used safely.

There are several additional options that can be used when the inferential method is specified to `"hmc"`, which we describe in the following.

²In order to use the functions in the package `rstan`, the user needs to either load the full package using the R command `library("rstan")`, or to prefix their name with the code `rstan::` to make them available in the current workspace.

- **chains**: the number of chains to run in the HMC (default = 2);
- **iter**: the total number of iterations (default = 2000);
- **warmup**: the number of “warm up” iterations (default = $\text{iter}/2$). The warm up is the adaptive phase in which the basic inputs of the HMC procedure are tuned. **rstan** does that automatically and once this stage is completed, the procedure is ready to immediately sample from the posterior distributions of interest;
- **thin**: the number of thinning (default = 1). For example, setting thinning to some value h , consists in selecting for inference every one in h simulations from the posteriors and can sometimes reduce the level of autocorrelation (for an equally large number of iterations used for the final estimation);
- **control**: a list specifying **rstan**-related options, e.g., `control = list(adapt_delta = 0.85)`, which can be used to tune more finely the acceptance rate in the HMC procedure (the closer this rate is to the upper limit of 1, the less likely there are to be numerically unstable simulations);
- **seed**: the random seed (to make the analysis replicable);
- **pars**: a string vector with the names of the relevant parameters. By default, **survHE** selects the location and ancillary parameters, as well as the coefficients associated with the covariates included in the model;
- **include**: a logical indicator (if set to **FALSE**, then the parameters specified in **pars** are not saved);
- **priors**: a list (of lists) specifying the values for the parameters of the prior distributions in the models (see the example below);
- **cores**: the number of CPU (cores) used to run the sampling procedure using **rstan** (default = 1);
- **save.stan**: a logical indicator (default = **FALSE**). If **TRUE**, then saves the model text file(s) to the user’s working directory. These can be used as templates to modify the basic model structure.

In practice, the user should not need to fiddle much with these optional arguments – certainly not without a clear understanding of the underlying modeling assumptions and the implications of any change to the default structure. Perhaps the default number of chains or iterations may be increased; or may be specific numeric values for the parameters of the prior distributions could be defined.

For instance, the default prior for the linear predictor coefficients is $\beta = (\beta_0, \dots, \beta_J) \sim \text{Normal}(\mu_\beta, \sigma_\beta \mathbf{I}_{J+1})$, where: μ_β and σ_β are vectors of size $(J + 1)$; and \mathbf{I}_{J+1} is the $(J + 1) \times (J + 1)$ identity matrix (see Table 2). Suppose the user wanted to select a smaller standard deviation for the Generalized Gamma model; this can be done using the following command – the list containing the custom values for the priors needs to be named, using the **survHE** terminology highlighted in Table 1.³

³It is worth mentioning that, unlike **OpenBUGS** or **JAGS** (which use the mean and precision), **rstan** parameterizes the Normal distribution in terms of the mean and the standard deviation.

```
R> m4 <- fit.models(formula, data, distr = "gengamma", method = "hmc",
+   priors = list(gengamma = list(sigma_beta = rep(5, 2))))
```

Note that we need to specify the values for the standard deviation for all the $J = 2$ covariates (the intercept and the treatment arm) and so in this case we define `sigma_beta = rep(5, 2)`, i.e., a vector of two elements, each equal to 5. Of course, there is nothing special about the value 5 and we could also select different values for the intercept and the treatment arm, e.g., `sigma_beta = rep(10, 2)`.

It is also possible to specify multiple values to modify the priors, for example

```
R> priors <- list(exp = list(sigma_beta = rep(4, 2)),
+   wei = list(mu_beta = rep(2, 2)))
R> m4 <- fit.models(formula, data, distr = mods, method = "hmc",
+   priors = priors)
```

would instruct **survHE** that the user wants to set: (a) the value for the standard deviation of the parameters β to 4, in the first model to be considered (Exponential); and (b) the value for the mean of the parameters β to 2, for the second model (Weibull).

Because the number of models in `mods` is 6, then **survHE** will complete the list `priors` by adding 4 more empty lists and **survHE** will use the default values for the remaining models. Consequently, it is important that the user specifies the required values in the correct order. For instance, if we wanted to specify $\sigma \sim \text{Gamma}(2, 4)$ for the Generalized Gamma model (the sixth in the list `mods`), we would need to make sure that this information is contained in the sixth element of the list `priors`. This could be done by using the following code

```
R> priors <- vector("list", 6)
R> names(priors) <- mods
R> priors[[6]] <- list(a_sigma = 2, b_sigma = 4)
```

which creates 5 empty lists to be associated with the first five models and the required list of values for the sixth one. Note that we still need to provide names for the elements of the list `prior`, so that **survHE** can figure out which element should be used for a given model among those specified in the call to `fit.models`. Even more succinctly, the same goal can be achieved by typing the following code

```
R> priors <- replicate(5, list())
R> priors[[6]] <- list(a_sigma = 2, b_sigma = 4)
R> names(priors) <- mods
```

– although we note that, in general terms, when it is necessary to specify complex options (such as the definition of the priors), it is perhaps a better idea to use one single distribution in the call to `fit.models`.

Then we can run **survHE** with the same command as before.

```
R> m4 <- fit.models(formula, data, distr = mods, method = "hmc",
+   priors = priors)
```

Model	Location parameters	Ancillary parameters	Natural parameters
Exponential	$\beta \sim \text{Normal}(\mu_\beta, \sigma_\beta I_{J+1})$ $\mu_\beta = \text{mu_beta} = \text{rep}(0, J + 1)$ $\sigma_\beta = \text{sigma_beta} = \text{rep}(5, J + 1)$	— —	Rate = $\exp(\beta_0)$
Weibull	$\beta \sim \text{Normal}(\mu_\beta, \sigma_\beta I_{J+1})$ $\mu_\beta = \text{mu_beta} = \text{rep}(0, J + 1)$ $\sigma_\beta = \text{sigma_beta} = \text{rep}(5, J + 1)$	$\alpha \sim \text{Gamma}(a, b)$ $a = \text{a_alpha} = 0.1$ $b = \text{b_alpha} = 0.1$	Shape = α Scale = $\exp(\beta_0)$
log-Normal	$\beta \sim \text{Normal}(\mu_\beta, \sigma_\beta I_{J+1})$ $\mu_\beta = \text{mu_beta} = \text{rep}(0, J + 1)$ $\sigma_\beta = \text{sigma_beta} = \text{rep}(100, J + 1)$	$\alpha \sim \text{Uniform}(a, b)$ $a = \text{a_alpha} = 0$ $b = \text{b_alpha} = 5$	Mean = β_0 Std. dev. = α
log-Logistic	$\beta \sim \text{Normal}(\mu_\beta, \sigma_\beta I_{J+1})$ $\mu_\beta = \text{mu_beta} = \text{rep}(0, J + 1)$ $\sigma_\beta = \text{sigma_beta} = \text{rep}(5, J + 1)$	$\alpha \sim \text{Gamma}(a, b)$ $a = \text{a_alpha} = 0.1$ $b = \text{b_alpha} = 0.1$	Shape = α Rate = $\exp(\beta_0)$
Gamma	$\beta \sim \text{Normal}(\mu_\beta, \sigma_\beta I_{J+1})$ $\mu_\beta = \text{mu_beta} = \text{rep}(0, J + 1)$ $\sigma_\beta = \text{sigma_beta} = \text{rep}(5, J + 1)$	$\alpha \sim \text{Gamma}(a, b)$ $a = \text{a_alpha} = 0.1$ $b = \text{b_alpha} = 0.1$	Shape = α Rate = $\exp(\beta_0)$
Gompertz	$\beta \sim \text{Normal}(\mu_\beta, \sigma_\beta I_{J+1})$ $\mu_\beta = \text{mu_beta} = \text{rep}(0, J + 1)$ $\sigma_\beta = \text{sigma_beta} = \text{rep}(5, J + 1)$	$\alpha \sim \text{Gamma}(a, b)$ $a = \text{a_alpha} = 0.1$ $b = \text{b_alpha} = 0.1$	Shape = α Rate = $\exp(\beta_0)$
Gen. Gamma	$\beta \sim \text{Normal}(\mu_\beta, \sigma_\beta I_{J+1})$ $\mu_\beta = \text{mu_beta} = \text{rep}(0, J + 1)$ $\sigma_\beta = \text{sigma_beta} = \text{rep}(100, J + 1)$	$\sigma \sim \text{Gamma}(a_1, b_1)$ $q \sim \text{Normal}(a_2, b_2)$ $a_1 = \text{a_sigma} = 0.1, b_1 = \text{b_sigma} = 0.1$ $a_2 = \text{mu_q} = 0, b_2 = \text{sigma_q} = 100$	Location = β_0 Scale = σ Shape = q
Gen. F	$\beta \sim \text{Normal}(\mu_\beta, \sigma_\beta I_{J+1})$ $\mu_\beta = \text{mu_beta} = \text{rep}(0, J + 1)$ $\sigma_\beta = \text{sigma_beta} = \text{rep}(100, J + 1)$	$\sigma \sim \text{Gamma}(a_1, b_1)$ $\log(p) \sim \text{Normal}(a_2, b_2)$ $q \sim \text{Normal}(a_2, b_2)$ $a_1 = \text{a_sigma} = 0.1, b_1 = \text{b_sigma} = 0.1$ $a_2 = \text{mu_P} = 0, b_2 = \text{sigma_P} = 0.5$ $a_3 = \text{mu_q} = 0, b_3 = \text{sigma_q} = 2.5$	Location = β_0 Scale = σ Shape (1) = q Shape (2) = p

Table 2: A summary of the default assumptions used for the models defined by `survHE` using `rstan`.

Table 2 shows a summary of the distributional assumptions used to define the default priors in the models implemented by **survHE** using **rstan**, together with the names assigned to the parameters of these distributions. For instance, if we wanted to specify a Normal(1,4) prior for the ancillary parameter of the Gompertz model (the fourth in the vector `mods`), we would need to specify the following command.

```
R> priors <- replicate(6, list())
R> names(priors) <- mods
R> priors[[4]] <- list(mu_alpha = 1, sigma_alpha = 4)
```

The use of informative priors is one of the examples in which external information can be included in the model, to stabilize the inference and complement the limited evidence provided by the data – although, again, it is the responsibility of the modeler to understand the specification and the implications of these assumptions into the model.

If the option `save.stan` is set to `TRUE`, then **survHE** will also save the model code as a text file (with the extension `.stan`) in the current directory. The data list formatted in a way that **rstan** can use is also automatically stored in the element `$misc$data.stan` inside the output of `fit.models`. The user can then modify the model structure starting from this template – for example it is possible to change the distributional assumptions and use, e.g., a Uniform prior for the scale σ of a Generalized F model. This will require a new compilation and, at present, the new model has to be run using **rstan** commands directly. Again, this illustrates the process of adapting relatively standard models to more complex situations, in which there is the need to bring external information to bear, although this may mean that the user needs to have a more thorough understanding of the underlying statistical problem.

6. Summarizing the results from **survHE**

Objects of class ‘**survHE**’ (such as `m1`, `m2`, `m3` and `m4` above) have associated methods such as `print`, `summary` and `plot` that can be used to summarize and visually inspect the results of the models analyzed. We describe them in the following.

6.1. Tabular form

When the models have been estimated, we usually want to summarize the estimates using a tabular format. **survHE** has a specialized function `print` that can do this, e.g., by typing

```
R> print(m1)
```

which returns the following table.

```
Model fit for the Exponential model, obtained using Flexsurvreg
(Maximum Likelihood Estimate). Running time: 0.025 seconds
```

	mean	se	L95%	U95%
rate	0.0824203	0.00828355	0.0676839	0.100365
as.factor(arm)1	-0.4656075	0.15427131	-0.7679738	-0.163241

Model fitting summaries

```
Akaike Information Criterion (AIC)....: 1274.576
Bayesian Information Criterion (BIC)...: 1282.387
```

In this case, the object `m1` contains many models; but unless the user specifies which one to print, **survHE** will assume that the first one should be used. If, for example, we wanted to visualize the estimates for the log-Logistic model (the fifth element of the string vector `distr`), then we would need to type

```
R> print(m1, mod = 5)
```

which would return the following output.

```
Model fit for the log-Logistic model, obtained using Flexsurvreg
(Maximum Likelihood Estimate). Running time: 0.024 seconds
```

	mean	se	L95%	U95%
shape	2.233748	0.1406365	1.974434	2.52712
scale	8.160865	0.5264208	7.191658	9.26069
as.factor(arm)1	0.348356	0.0943506	0.163432	0.53328

Model fitting summaries

```
Akaike Information Criterion (AIC)....: 1208.494
Bayesian Information Criterion (BIC)...: 1220.211
```

In both cases, **survHE** standardizes the format of the output, so that the results are reported for the “basic” parameters (e.g., rate, shape or scale) as well as the covariate effects. Notice that the “basic” parameters are always reported on the natural scale, while the covariate effects are in the scale defined by the linear predictor (as presented in Table 1). Thus, in the cases presented above, the value of the coefficient `as.factor(arm)1` represents the impact of the treatment arm on the log scale, because both for the Exponential and the log-Logistic the location parameter is modeled using a log link – and thus in Table 1 we write $\mu_i = \exp(\dots)$. The `print` method for ‘**survHE**’ objects has an additional option, which allows the user to visualize the summary of the model results in the original notation used by the relevant package used to perform the estimation. Thus, in this case typing

```
R> print(m1, mod = 6, original = TRUE, digits = 3)
```

would show the results for the Generalized Gamma model (the sixth in the string vector `mods`) using the original **flexsurv** formatting and using only 3 significant decimal places.

Call:

```
flexsurv::flexsurvreg(formula = formula, data = data, dist = distr)
```

Estimates:

	data	mean	est	L95%	U95%	se	exp(est)	L95%	U95%
mu	NA	2.2918	2.1357	2.4479	0.0796	NA	NA	NA	NA

```

sigma          NA 0.5871 0.4613 0.7471 0.0722      NA      NA      NA
Q              NA 0.8507 0.3620 1.3394 0.2493      NA      NA      NA
as.factor(arm)1 0.4850 0.3463 0.1743 0.5183 0.0878 1.4138 1.1904 1.6792

```

```

N = 367, Events: 172, Censored: 195
Total time at risk: 2612.07
Log-likelihood = -598.3923, df = 4
AIC = 1204.785

```

The same principles apply to ‘survHE’ objects storing models fitted using either **INLA** or **rstan**. For instance the command

```
R> print(m3, 6)
```

returns the output

```

Model fit for the Generalised Gamma model, obtained using Stan
(Bayesian inference via Hamiltonian Monte Carlo). Running time:
23.449 seconds

```

	mean	se	L95%	U95%
mu	2.296495	0.0835967	2.123717	2.455760
sigma	0.596774	0.0712512	0.462840	0.743653
Q	0.851923	0.2520732	0.364672	1.363819
as.factor(arm)1	0.347634	0.0875420	0.176135	0.525895

Model fitting summaries

```

Akaike Information Criterion (AIC)....: 7103.724
Bayesian Information Criterion (BIC)..: 7123.250
Deviance Information Criterion (DIC)..: 7612.510

```

For HMC models, even more importantly, when using the option `original = TRUE` we are able to look at helpful convergence statistics, which should be used to assess whether the MCMC procedure has been successful in exploring the relevant posterior distributions. For example, we could use the code

```
R> print(m3, mod = 2, original = TRUE, digits = 6)
```

which returns the following output.

```

Inference for Stan model: WeibullAF.
2 chains, each with iter=2000; warmup=1000; thin=1;
post-warmup draws per chain=1000, total post-warmup draws=2000.

```

	mean	se_mean	sd	2.5%	25%	50%
beta[1]	2.324760	0.001587	0.056368	2.219097	2.286221	2.323779
beta[2]	0.349667	0.002620	0.088872	0.178268	0.288323	0.349394
alpha	1.803243	0.002924	0.108722	1.601334	1.728505	1.801598

scale	10.240529	0.016460	0.580670	9.199025	9.837695	10.214201
lp__	-600.292784	0.039113	1.207529	-603.491904	-600.890441	-599.972663
	75%	97.5%	n_eff	Rhat		
beta[1]	2.360766	2.439654	1261	1.002989		
beta[2]	0.406506	0.522045	1151	1.000595		
alpha	1.873171	2.019929	1383	1.000788		
scale	10.599070	11.469076	1245	1.003169		
lp__	-599.406559	-598.904339	953	1.000773		

Samples were drawn using NUTS(diag_e) at Tue Sep 29 20:28:54 2020.
 For each parameter, `n_eff` is a crude measure of effective sample size, and `Rhat` is the potential scale reduction factor on split chains (at convergence, `Rhat=1`).

The table displays the original output from `rstan`. Here, the coefficient `beta[1]` is the intercept (log scale), while the coefficient `beta[2]` is the effect of the only covariate included in the model. The estimates are reported as the mean, Monte Carlo error (`se_mean`), standard deviation, an approximate 95% credible interval and then the *effective sample size* `n_eff` and the *potential scale reduction* (PSR) `Rhat`.

The effective sample size gives an indication of the underlying autocorrelation in the MCMC samples – values close to the total number of iterations, or at any rate not too low, indicate a low level of autocorrelation (which is what we want). The PSR is an analysis-of-variance-type of statistics, indicating for each variables whether convergence is reached. If the procedure is ran on more than one parallel chain, then `Rhat` is computed as a function of the ratio of the variance within to the variance between chains – if this is close to 1 and definitely less than 1.1, convergence can be satisfactorily declared. If not, it may be necessary for example to increase the number of iterations.

A further optional inputs to the `print` methods is `digits` (default = 6), which determines the number of digits printed in the table.

6.2. Visual interpretation

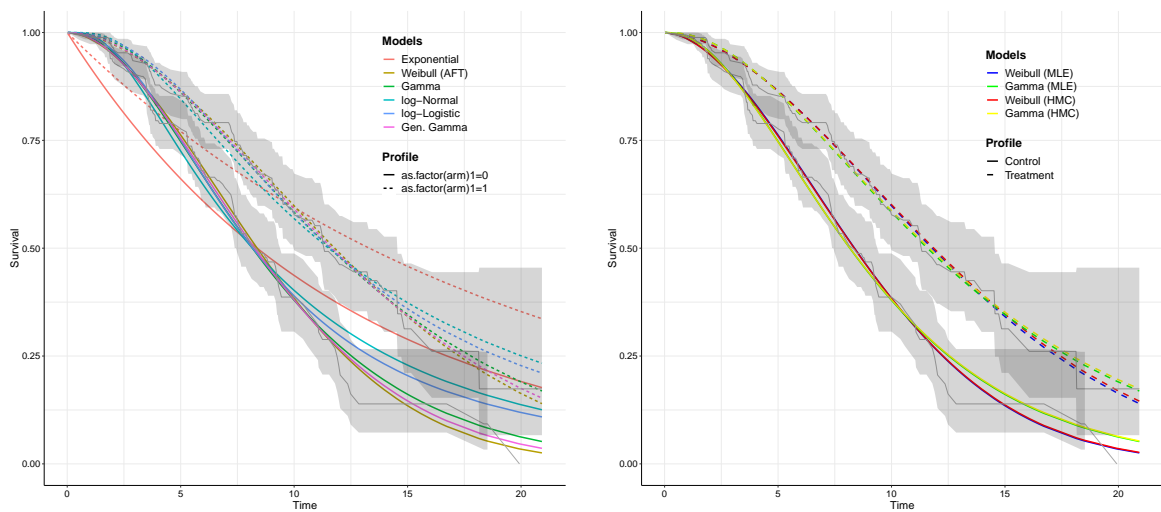
Once an object of the class ‘`survHE`’ has been created using the command `fit.models`, it is possible to visualize the resulting survival curve(s) by simply using the `plot` method. `survHE` uses `ggplot2` (Wickham 2016) as the graphical engine. This means that the graphs produced by `survHE` are highly customizable. However, `survHE` simplifies the modeler’s work as most of the code needed to produce the graphs is already implemented in the `plot` method.

For example, the command

```
R> plot(m1, add.km = TRUE)
```

displays the graph shown in Figure 3(a). The option `add.km = TRUE` instructs `survHE` to add a Kaplan-Meier estimate to the plot (in this case, one for each treatment arm is produced as the grey areas underlying the colored curves).

The `plot` method allows to display the output from different models. For example, we may be interested in comparing the survival curves from some of the models obtained using either MLE or HMC. We could do this by using the following command.



(a) Survival curves for the models fitted in the object `m1`.

(b) Survival curves for some of the models fitted in the objects `m1` and `m3`.

Figure 3: Graphs produced by the `plot` method in `survHE`.

```
R> plot(m1, m3, mods = c(2, 3, 8, 9), colour = c("blue", "green", "red",
+ "yellow"), lab.model = c("Weibull (MLE)", "Gamma (MLE)",
+ "Weibull (HMC)", "Gamma (HMC)"), lab.profile = c("Control", "Treatment"),
+ add.km = TRUE)
```

This command instructs `survHE` to stack together the objects `m1` and `m3`, which basically produces a single ‘`survHE`’ object containing 12 models (6 from `m1` and 6 from `m3`). The option `mods = c(2, 3, 8, 9)` selects the models in positions 2, 3, 8 and 9 (i.e., the second and third from `m1` and then the second and third from `m3`). The option `colour` can be used to select the colors with which to plot each curve on the graph. The options `lab.model` and `lab.profile` can be used to customize the graph legends. The former specifies the labels to be associated with the several models plotted, while the latter customizes the “profile” of covariates depicted. The resulting graph is shown in Figure 3(b).⁴

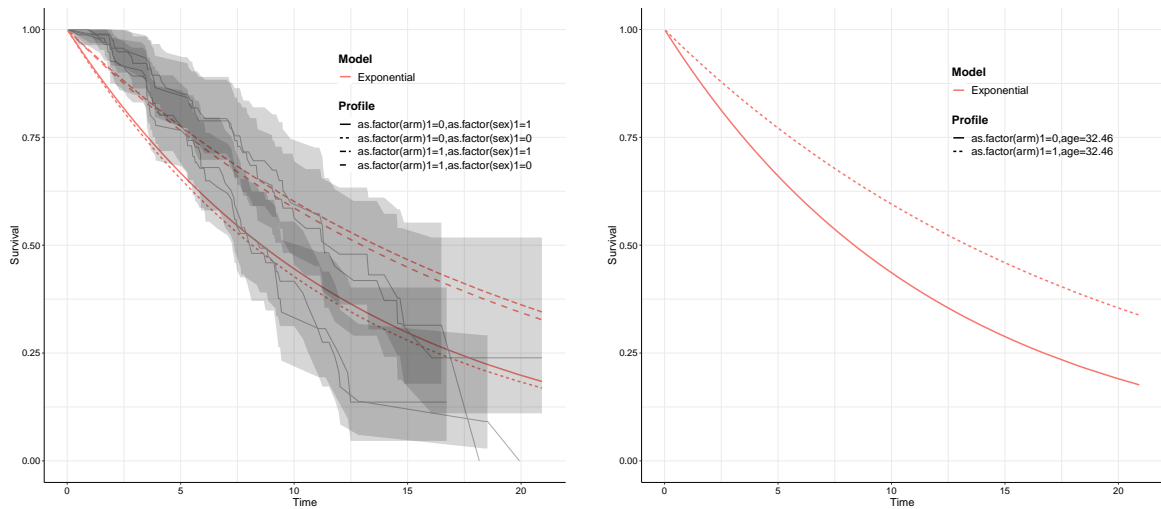
Notice that, *in this particular instance*, the differences in the frequentist and Bayesian analysis are not striking. This is due to the fact that we have considered relatively vague priors in combination with mature data (notice that the Kaplan-Meier estimates are fairly close to 0 for both arms, indicating that, in this study, most individuals have experienced the events and the number of censored subjects is not large). This means that the observed data are in fact driving the overall inferential process, thus limiting the impact of the priors and consequently the differences between the MLE and the Bayesian analysis. We note again however that, particularly when data are not mature and are subject to a large amount of censoring, the inclusion of genuine prior information does contribute more substantially to the final inference.

The full list of the options for the method `plot` in `survHE` is given below.

⁴Alternatively, the user can name the `survHE` objects in the `plot` call, for example using a command such as `plot(MLE = m1, HMC = m3, mods = c(2, 3, 8, 9))`. This would automatically create labels for the models by combining the name of the object and the name of the distribution used, e.g., `MLE:Weibull (AFT)`, `MLE:Gamma`, `HMC:Weibull (AFT)` and `HMC:Gamma`.

- **add.km**: a logical variable. If **TRUE** (the default value), then also add the Kaplan-Meier estimates of the data to the plot;
- **newdata**: a list (of lists) providing the values for the relevant covariates to stratify the survival curves. If **NULL**, then **survHE** will use the mean value for all the covariates if at least one is a continuous variable, or the full combination of the categorical covariates, in line with **flexsurv**;
- **xlab**: a string with the label for the x -axis (default = "Time");
- **ylab**: a string with the label for the y -axis (default = "Survival");
- **lab.profile**: a (vector of) string(s) indicating the labels associated with the strata defining the different survival curves to plot. This defaults to the value used by the Kaplan Meier estimate given in **fit.models**;
- **t**: a vector specifying the time horizon over which the survival curves should be computed and plotted. It defaults to the observed time range in the data, but for instance, the user can specify the extrapolation to time $t = 60$ using the code `plot(m1, t = seq(0,60))`. The range of the input **t** determines also the x axis limits in the plot;
- **colors**: a vector of characters defining the colours in which to plot the different survival curves;
- **annotate**: a logical indicator (defaults to **FALSE**), describing whether **survHE** should also add text to highlight the observed vs extrapolated data in the survival curves;
- **legend.position**: a vector of proportions to place the legend. Default to `c(.75, .9)`, which means 75% across the x -axis and 90% across the y -axis;
- **legend.title**: suitable instructions to format the title of the legend. This defaults to `element_text(size = 15, face = "bold")`. Thus, for example, the command `plot(m1, legend.title = element_text(size = 20))` would increase the font size for the title of the legends (i.e., the strings "Models" and "Profile") up to 20 points. Further customization is possible, but requires knowledge of **ggplot2**;
- **legend.text**: suitable instructions to format the text of the legend; It defaults to `element_text(colour = "black", size = 14, face = "plain")`, but again can be further customized using **ggplot2** commands;
- **nsim**: if this input is set to a value greater than 1, then **survHE** will automatically construct simulations for the entire distribution of the underlying survival curves (see Section 7 for more details). By default, **survHE** sets this to 1 and only plots the average of the distribution for the survival curves.

Most of these options are actually trivial; perhaps only **newdata** deserves a more detailed explanation. **survHE** follows the philosophy of **flexsurv** to construct the survival curves. Consequently, when the model contains categorical covariates, a single survival curve is estimated for each combination of their modalities. Consider for example the following command.



(a) Survival curves for the Exponential model fitted using MLE and controlling for sex and treatment arm. (b) Survival curves for specific values of the covariates, as selected in the list `newdata`.

Figure 4: Graphs produced by the `plot` method in **survHE** when using different combination of the covariate values.

```
R> m5 <- fit.models(Surv(time, event) ~ as.factor(arm) +
+   as.factor(sex), data = data, distr = "exp")
```

This estimates the survival times via MLE using an Exponential model and controlling for the effect of the treatment arm and the individuals' sex. The estimates are as follows.

Model fit for the Exponential model, obtained using `Flexsurvreg` (Maximum Likelihood Estimate). Running time: 0.016 seconds

	mean	se	L95%	U95%
rate	0.0843425	0.0102787	0.066422	0.107098
as.factor(arm)1	-0.4658255	0.1542727	-0.768194	-0.163457
as.factor(sex)1	-0.0505827	0.1533420	-0.351127	0.249962

Model fitting summaries

Akaike Information Criterion (AIC)....: 1276.467

Bayesian Information Criterion (BIC)..: 1288.183

As shown in Figure 4(a), this simple model already becomes complex to visualize, because there are 4 strata identified by the two covariates, each of which is binary. Things become even more complicated if we mix a continuous and a categorical covariate, e.g., treatment and age. Thus, it may be helpful then to plot the results using a different strategy. For example we can define a list in which we specify the value for the covariates that we want to use to compute the survival curves, e.g.,

```
R> m5 <- fit.models(Surv(time, event) ~ as.factor(arm) + age,
+   data = data, distr = "exp")
```

```
R> newdata <- list(list(arm = 0, age = mean(data$age)),
+   list(arm = 1, age = mean(data$age)))
```

which would create two “profiles” by varying the treatment arm and keeping the value for sex to the observed average in the data. The resulting plot can be obtained using the following code.

```
R> plot(m5, newdata = newdata)
```

which produces the graph in Figure 4(b). When the option `newdata` is used, it is probably best to plot a single model in a graph – even if the ‘`survHE`’ object contains many (just like `m1`, `m2` or `m3`), the option `mods` can be used to select only one of them, as shown above. When selecting the “profile” to be specified in the list `newdata`, the user needs to specify a value for all the covariates that appear in the model formula, although the order in which they are entered is not important.

As all `survHE` graphs are `ggplot2` objects, they can also be manipulated using `ggplot2` functions and commands. So for instance, the user can save a `survHE` plot to a named object and then add `ggplot2` customization, such as in the following example

```
R> p <- plot(m1)
R> p + theme_classic()
```

which would draw the graph in a template similar to the one used by `base R`. Again, knowledge of `ggplot2` is required in order to take full advantage of its functionalities – but we reiterate here that `survHE` also takes care of many of the basic needs of the modelers.

6.3. Estimation of the mean survival time

As mentioned earlier, in a health economic evaluation it is of interest to estimate the mean survival time – this is the quantity that is relevant to determine for example the effectiveness of a given intervention. For some of the parametric models described above, the mean survival time can be computed analytically. For example, for a Weibull model, this is

$$E[T] = \int_0^{\infty} tf(t | \boldsymbol{\theta})dt = \int_0^{\infty} S(t)dt = \mu\Gamma\left(1 + \frac{1}{\alpha}\right),$$

where μ is the scale, α is the shape and $\Gamma(\cdot)$ is the Gamma function.

More generally, it is possible to approximate this quantity using the “trapezium rule” over a large enough time horizon as

$$E[T] \approx \frac{1}{2} \sum_{t \in \mathcal{T}} \tau [S(t + \tau) - S(t)],$$

where $\mathcal{T} = \{0, (t + \tau), (t + 2\tau), \dots, (t + K\tau)\}$ is a discrete set of $(K + 1)$ time points and τ is an arbitrarily small increment. Notice that in order to approximate the mean sufficiently well, it is important to extend the range \mathcal{T} long enough so that all the survival curves actually fade out to 0. Also, the smaller the increment τ , the more trapezoids are fitted under the survival curve and thus the better the approximation. In practice, there is a trade-off between the level of approximation and the computational time required for the calculation.

`survHE` automatically performs this calculation by means of the method `summary`; so for example, the R code

```
R> summary(m1)
```

produces the following output.

```
Estimated average survival time distribution*
              mean          sd      2.5%   median   97.5%
as.factor(arm)=0  9.929582 0.6185205  8.778254  9.90262 11.11727
as.factor(arm)=1 12.753808 0.6510119 11.390461 12.79599 13.94223
```

```
*Computed over the range: [ 0.030-20.920] using 1000 simulations.
NB: Check that the survival curves tend to 0 over this range!
```

Because the user has not specified a time range over which to compute the mean survival, **survHE** assumes the observed range of times (in this case [0.03–20.92]). As is obvious from Figure 3, in this range the survival curves have not reached 0 and thus the estimated mean survival is certainly biased downwardly. To correct this, it is sufficient to use the code

```
R> summary(m1, t = seq(0, 60))
```

which instructs **survHE** to compute the means over a range of times between 0 and 60, with default unit increments. The resulting values are substantially different.

```
Estimated average survival time distribution*
              mean          sd      2.5%   median   97.5%
as.factor(arm)=0 12.15368 1.157162  9.991659 12.11087 14.47000
as.factor(arm)=1 18.39585 1.811342 15.098328 18.37225 21.95334
```

```
*Computed over the range: [ 0-60] using 1000 simulations.
NB: Check that the survival curves tend to 0 over this range!
```

Incidentally, for the Exponential model, the analytic average values are computed as $1/\text{Rate}$, which equal

$$\exp(\beta_0)^{-1} = \exp(\log(0.0824))^{-1} = 12.133$$

and

$$\exp(\beta_0 + \beta_1)^{-1} = \exp(\log(0.0824) - 0.4656)^{-1} = 19.328,$$

for the control and the treatment arm, respectively (with β_0 the intercept and β_1 the treatment effect as estimated in **m1**). The estimate produced by the **summary** command can be improved by selecting a longer time horizon and/or a lower value for the increment τ , e.g., `t = seq(0, 100, 0.1)`. It is also possible to increase the number of simulations used to characterize uncertainty in the underlying parameters (and hence survival curves – we return to this point in Section 7). The default value of 1000 can be modified by specifying the optional argument **nsim** to a different value. Similarly, it is possible to specify a specific “profile”, in terms of the covariates included in the model by using the optional argument **newdata**. So for example, the command

```
R> summary(m1, t = seq(0, 60), newdata = list(list(arm = 1)))
```

produces an estimate of the distribution of the mean survival time for the treated. Finally, the user can also include a vector `labs` in the call to `summary`, which contains a suitable number of text strings that can be used to label the values in the resulting table, for example `labs = c("Controls", "Treated")`.

6.4. Model assessment

Health economics guidelines suggest that the assessment of the models is done by complementing the visual inspection with some more formal testing. The suggested way is through the use of a specific information criterion (IC), such as Akaike IC (AIC) or the Bayesian IC (BIC). These statistics are based on the model deviance $-2 \log \mathcal{L}(\boldsymbol{\theta})$ and a penalty function, which typically depends on the number of parameters and the complexity of the model – the rationale being that models that are too complex tend to “overfit” the data; this means that they may do very well at estimating the data at hand, but usually have poor predictive ability of other data.

An additional IC specific to Bayesian models is the deviance IC (DIC), proposed by Spiegelhalter, Best, Carlin, and Van Der Linde (2002)⁵. `survHE` then computes an estimate of AIC and BIC for any inference engine and also an estimate of the DIC for models fitted in either `INLA` or `rstan`⁶.

The results of model fit can be visually inspected using the `survHE` function `model.fit.plot` as in the examples below.

```
R> model.fit.plot(m1)
R> model.fit.plot(m1, type = "bic")
R> model.fit.plot(MLE = m1, HMC = m3, stacked = TRUE,
+   mods = c(1, 2, 3, 7, 8, 9), name_legend = "Inferential method")
R> model.fit.plot(MLE = m1, HMC = m3, stacked = TRUE,
+   mods = c(1, 2, 3, 7, 8, 9), type = "dic")
```

In panel (c) of Figure 5, we specify that the bar plots should be “stacked”, which means that models using the same distribution from different objects are grouped together. The objects can be named (as for the main `plot` function), which automatically creates a label to associate with the colors with which the bars are plotted. If nothing is specified, then `survHE` uses the generic labels `Object1`, `...`, `ObjectN` (assuming there are `N` `survHE` objects to be plotted). The option `name_legend` specifies the title of the legend. If nothing is specified, then `survHE` uses the default label `survHE object`.

In panel (d) of Figure 5 for some of the models the bar is not plotted: this is because for model `m1` the DIC cannot be computed, as it is fitted using MLE.

⁵In fact, for Bayesian models obtained via HMC, `survHE` also computes the DIC using the slightly different definition suggested by Gelman *et al.* (2013); the two versions of the DIC may differ in the presence of extreme asymmetry or multi-modality in the posterior distributions. The alternative estimates for the DIC are stored in the element `$model.fitting$dic2` of a ‘`survHE`’ object, but are only present if the inferential method is set to “hmc”.

⁶Notice that `rstan` does not provide measures of model fit based on information criteria. There are several arguments to prefer other methods to assess the performance of a statistical model, for example the posterior predictive check, as discussed in Gelman *et al.* (2013). However, because NICE guidelines suggest using AIC and BIC (and, by extension, DIC), `survHE` computes the relevant model fit statistics and reports them using the `print` method and the `model.fit.plot` function.

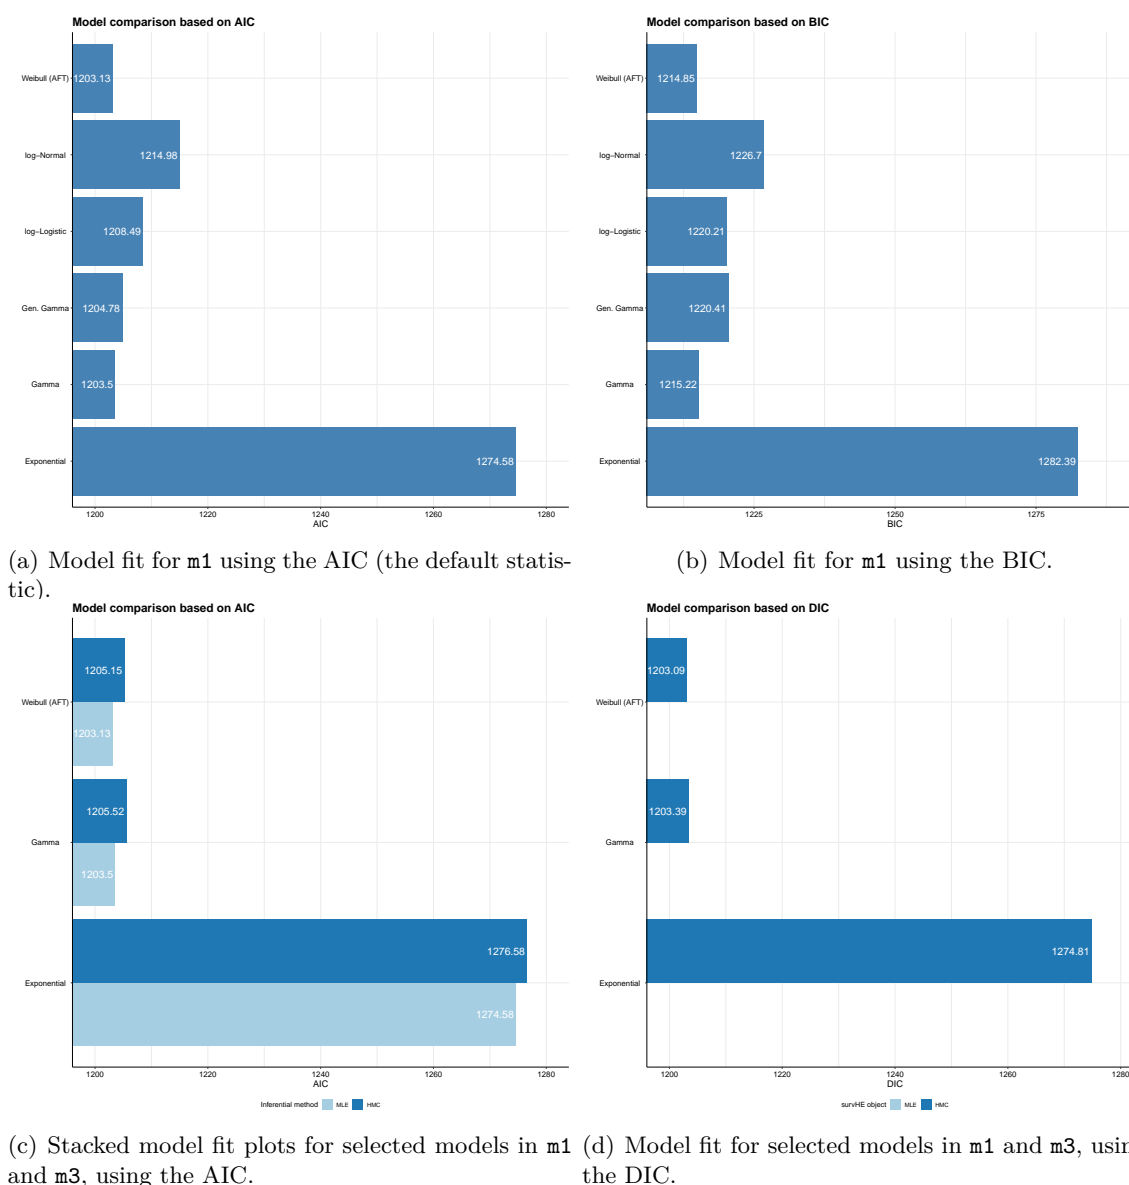


Figure 5: Example of plots to assess model fit for ‘survHE’ objects.

There are many optional arguments to the `model.fit.plot` function:

- **type**: a string specifying the statistic to be used (possible values are "aic", "bic" or "dic");
- **scale**: if `scale = "absolute"` (default), then plot the absolute value of the relevant IC. If `scale = "relative"` then plot a rescaled version taking the percentage increase in the *IC in comparison with the best-fitting model;
- **colour**: specifies a vector of strings to assign colors to the bars. In fact, the user can use the code `col` or `color` interchangeably;
- **stacked**: a logical indicator, defining whether the bars should be stacked and grouped

by ‘survHE’ objects (default = FALSE);

- **name_legend**: a string defining the title of the legend that is added to the bar plot when **stacked** is set to TRUE;
- **main**: a string specifying the title of the plot;
- **models**: a (vector of) string(s), specifying the label to associate with the models being plotted.

As a final note, we stress here that while model fitting statistics are useful to rank the plausibility of the distributional assumptions tested, they cannot inform the quality of the extrapolation of the survival curves. This is because any IC is based on the *observed* data and quantifies how well a specific model fits to them. Thus, there is no measure as to how a model performs on data that are as yet unobserved, based on such criteria. For this reason, it is recommended that plausibility of the extrapolation is judged using a combination of visual inspection, formal testing and the effectively unavoidable inclusion of external information and clinical opinions – again suggesting the usefulness of the Bayesian approach.

7. Probabilistic sensitivity analysis

Unlike standard epidemiological analysis where the objective is often to estimate the effect of some relevant covariates on the survival time, in health economic evaluation the goal is rather to produce an estimate of the entire survival curve over a long period of time (or at any rate, a longer period than the observed follow up). This estimate is then used to populate the economic model, e.g., by obtaining estimates of the transition probabilities between states in a Markov model or, perhaps even more commonly, in a partitioned survival analysis (Woods, Sideris, Palmer, Latimer, and Soares 2017). More importantly, we need to quantify the impact of uncertainty in the model parameters on the decision-making process and thus we typically repeat this exercise for a large number of times, upon varying the value of the parameters that determine the survival curves. In a fully Bayesian approach this uncertainty is induced by the full joint posterior distribution of the parameters.

survHE is designed to perform this task directly in R, through the function `make.surv`. A typical call is as follows

```
R> psa <- make.surv(fit = m3, nsim = 1000, t = seq(.1, 63))
```

which generates an object `psa` containing, among other things, `nsim = 1000` simulations for the survival curves. In the above case, `m3` contains in fact 6 different models (upon varying the distributional assumptions), but because the user has not specified a value for the input `mod` (in this case a number between 1 and 6), `make.surv` uses the first one, i.e., the Exponential, by default. Adding the option `mod = 6` to the call to `make.surv` would consider the sixth model (Generalized Gamma) instead.

The resulting output `psa` can be accessed directly by the user. The command

```
R> names(psa)
```

shows that it comprises several elements


```
[1] "S"          "sim"        "nsim"       "mat"        "des.mat"    "times"
```

each of which can be accessed using either the “dollar” or the “double bracket” notation in R e.g., `psa$$S` or `psa[[3]]`. A brief description of each of these elements is given in the following.

- **S**: a list – for each simulated value of the parameters, a list with the survival curves associated with the configuration of the covariates;
- **sim**: simulated values for the main parameters (e.g., scale, shape, rate, mean, sd) for each configuration of the covariates;
- **nsim**: the number of simulations saved;
- **mat**: a list – for each configuration of covariates a matrix with **nsim** rows and as many columns as time points with the survival curves (to be read row-wise);
- **des.mat**: a design matrix with the combination of the covariates used (each represents an element in the lists **S** and **mat**).
- **times**: the vector of times used in the computation. By default, this is the set of times in the data passed as argument to `fit.models`, but the user can specify a different set of times.

The reason why `make.surv` creates so many outputs is mainly for internal convenience. In fact, this is a central function in how **survHE** works and it is called internally by other utility functions (e.g., `plot` and `print`). By and large, the user does not need to manipulate the output directly.

The list of inputs for the `make.surv` function is as follows.

- **fit**: the result of the call to the `fit.models` function, containing the model fitting (and other relevant information), e.g., `m3` from above;
- **mod**: the index of the model. The default value is 1, but the user can choose which model to visualize, if the call to `fit.models` has a vector argument for the input `distr`;
- **t**: the vector of times to be used to create the survival curves. By default, **survHE** uses the times observed in the original data, but the user can specify a longer follow up (this is actually a useful feature for the purpose of performing a health economic evaluation and PSA);
- **newdata**: a list (of lists), specifying the values of the covariates at which the computation is performed. For example `list(list(arm = 0), list(arm = 1))` will create two survival curves, one obtained by setting the covariate `arm` to the value 0 and the other by setting it to the value 1. In line with **flexsurv** notation, the user needs to either specify the value for *all* the covariates or for none (in which case, **newdata** is set to `NULL`, which is the default). If some value is specified and at least one of the covariates is continuous, then a single survival curve will be computed in correspondence of the average values of all the covariates (including the factors, which in this case are expanded into indicators);

- `nsim`: the number of simulations from the distribution of the survival curves. The default is at `nsim = 1`, in which case the point estimate for the relevant distributional parameters is used and the resulting “average” survival curve computed.

To visualize the results of the PSA, **survHE** has a specialized function called `psa.plot`, which can be used, for example as follows.

```
R> psa.plot(psa)
```

This command produces the graph in Figure 6(a). The graph shows the average survival curve and 95% interval estimates around them. If the `method` is set to `"mle"`, then the intervals are obtained by multivariate Normal bootstrap, while for the Bayesian models they are obtained using samples from the relevant posterior distributions.

The user can specify the labels for the x - and y -axis (respectively by including the additional arguments `xlab = "..."` and `ylab = "..."`). In addition, if no color specification is given by the user, e.g., in the form `col = c("blue", "red")`, then **survHE** will randomly choose a coloring scheme. Finally, the parameter `alpha` (default at 0.1) determines the transparency of the curves; values for `alpha` close to 0 imply greater transparency, while values closer to 1 create a solid plot (with no transparency). A fully customized call to `psa.plot` is as follows and produces the graph in Figure 6(b).

```
R> psa.plot(psa, xlab = "Extrapolated time",
+   ylab = "Estimation of the survival curves",
+   alpha = 0.2, col = c("dark grey", "black"),
+   main = "PSA to survival curves", cex.txt = .95)
```

A similar graph can also be obtained directly using the `plot` method. For example, the command `plot(m3, mods = 1, nsim = 1000, t = seq(0.1, 63))` produces a graph with the average estimates for the survival curves and 95% intervals obtained using either bootstrap or the posterior distributions, depending on the method selected. This graph is shown in Figure 6(c).

7.1. Exporting the results to Microsoft Excel

Once the PSA samples are available for the survival curves, it may be easy to continue building the economic model using R. In fact, this is the strategy that we advocate (Baio and Heath 2016). However, we acknowledge that practitioners use Microsoft **Excel** to produce the economic assessment. Thus, **survHE** has a specialized function, called `write.surv`, that allows the user to export the simulations for the survival curves to a `.xls(x)` file⁷, so that they can be easily used when constructing a Markov model or a partitioned survival analysis in Microsoft **Excel**.

As mentioned above, the user can actually manipulate the output of the call to `make.surv` independently and so in a way bypass `write.surv` entirely. However, **survHE** tries to simplify the work process and can be used as follows.

⁷We recommend using the `xlsx` extension, because `xls` has a restriction to 256 columns, which may not suffice if the object being exported is the output of a call to `make.surv` with the number of simulations `nsim` set to a value larger than 256 (normally, a value of 1000 or greater is used in real cases).

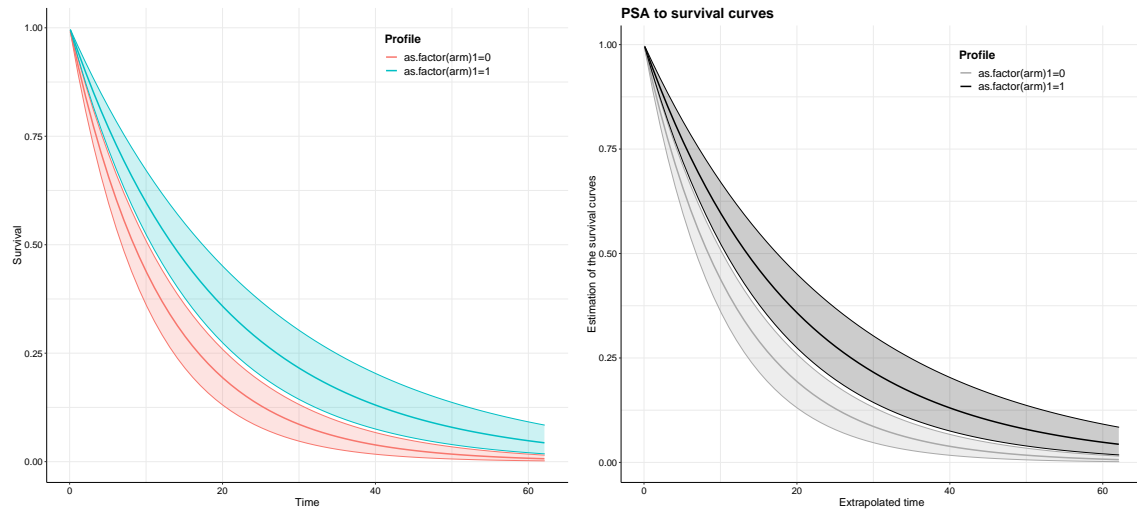

(a) Default graph produced by `psa.plot`.(b) Fully customized graph produced by `psa.plot`.(c) PSA plot obtained directly using the main plot method in `survHE`.

Figure 6: Probabilistic sensitivity analysis for the survival curves.

```
R> write.surv(psa, file = "temp.xlsx")
```

which produces the following output in R

```
1000 simulation(s) for the survival curve written to file: 'temp.xlsx'
```

```
Profile(s):
```

```
as.factor(arm)1=0
```

```
as.factor(arm)1=1
```

and creates a file containing the relevant simulations. The R output clarifies that the resulting file contains two spreadsheets: the first one is for the survival curves considering the treatment arm set to 0 (e.g., controls), while the second one is for the intervention arm (set to 1). Different model specifications would create a different number of matrices (with `nsim` rows and as many time points for columns) depending on the covariate combinations.

8. Advanced models

While the models presented in Table 1 are likely to produce at least one “good” candidate in most situations, it is possible that more complex model structures may be needed to accommodate a particular dataset or analysis. The main advantage of using these more complicated models is that they are able to define hazard functions that can better approximate the lifetime risk of the specific population under study, potentially including competing causes or external knowledge that can be used to anchor the survival probabilities below a certain threshold (e.g., the healthy population or a subset of the population affected by similar diseases or treated with drugs with similar action).

In particular, we focus here on two “flexible” models: the first one (which is already implemented in **flexsurv**) is based on cubic splines, while the second is an extension of the standard Weibull model.

8.1. Royston-Parmar splines

Splines are numeric functions typically defined as a collection of local polynomials; the main idea is to partition the x -axis into a set of intervals defined by “knots” and then within each interval use a different polynomial function to approximate the underlying “true” function on the y -axis. This construction provides great flexibility and effectively an arbitrary number of parameters, depending on how many knots (and hence on the density with which the x -axis is partitioned) are selected.

In the context of survival analysis, splines can be used to model flexibly (a suitable transformation of) one of the basic functions, e.g., the survival or the hazard (Royston and Lambert 2011). An increasingly popular model is the one developed by Royston and Parmar (2002). Basically, this defines a probability distribution to model the observed and censored times as a function of an “augmented” set of parameters $\theta = (\beta, \gamma)$ and covariates (\mathbf{X}, \mathbf{B}) . Here, β are the coefficients associated with the observed covariates \mathbf{X} , exactly as in (2). In addition, for each individual (and, hence, observed time) in the dataset, we consider a set of $(M + 2)$ “basis” function $\mathbf{B}_i = (B_{i0}, B_{i1}, \dots, B_{iM})$, where

$$B_{im} = (\log t_i - k_m)_+^3 - \lambda_m (\log t_i - k_{\min})_+^3 - (1 - \lambda_m) (\log t_i - k_{\max})_+^3, \quad \lambda_m = \frac{k_{\max} - k_m}{k_{\max} - k_{\min}}$$

and $(\log t_i - a)_+ := \max\{0, (\log t_i - a)\}$. The vector of knots is defined as $k_{\min} = 0 < k_1 < \dots < k_{\max} = \infty$; typically, the M “internal” knots are set up in terms of the quantiles of the observed distribution of the times. For example, if we set $M = 3$, **survHE** would automatically consider the three quartiles (q_1, q_2, q_3) , representing the 25%-, 50%- and 75%-percentiles of the observed times distribution.

The Royston-Parmar (RP) model defines a modified linear predictor

$$\eta_i = \sum_{m=0}^{M+2} \gamma_m B_{im} \left[+ \sum_{j=0}^J \beta_j X_{ij} \right],$$

which is used to model directly $\log(-\log S(t_i)) = \log H(t_i) = \eta_i$ (notice that we use the $[+ \dots]$ notation here to highlight the fact that the model may not include any covariate \mathbf{X}

and thus only rely on the splines structure)⁸.

Given this set up, it is possible to prove that

$$\log h(t_i) = -\log t_i + \log \eta'_i + \eta_i \quad (4)$$

and

$$\log S(t_i) = -\exp(\eta_i), \quad (5)$$

with $\eta'_i = \sum_{m=1}^{M+2} \gamma_m B'_{im}$ – the first derivatives B'_{im} are easy to compute, recalling that $\frac{\partial x^3}{\partial x} = 3x^2$. Substituting (4) and (5) into (1), we can completely characterize the resulting likelihood. The RP model is directly available in **flexsurv** and is also implemented under a pre-compiled HMC-based Bayesian framework in **survHE**. The two versions can be obtained via a specialized call to `fit.models`.

```
R> formula <- Surv(time, event) ~ as.factor(arm)
R> m6 <- fit.models(formula = formula, data = data, distr = "rps", k = 2)
R> m7 <- fit.models(formula = formula, data = data, distr = "rps", k = 2,
+   method = "hmc")
```

survHE accepts the string "rps" to indicate the “Royston-Parmar splines” distribution and also requires the input `k` to specify the number M of internal knots to be used (if this is not provided by the user, **survHE** assumes that `k` = 0, which reduces the RP model to a Weibull PH formulation). The `formula` is used to specify the “fixed” component of η_i , i.e., the set of covariates in \mathbf{X} . In this case, we use only the treatment arm.

All the usual methods are available for the resulting ‘**survHE**’ objects `m6` and `m7`. For example the commands

```
R> print(m6)
R> print(m7)
```

return the summary tables

Model fit for the Royston & Parmar splines model, obtained using
Flexsurvreg (Maximum Likelihood Estimate). Running time: 0.129 seconds

	mean	se	L95%	U95%
gamma0	-4.652836	0.474704	-5.583239	-3.722433
gamma1	2.523687	0.565148	1.416017	3.631357
gamma2	0.387949	0.331049	-0.260894	1.036792
gamma3	-0.419667	0.398441	-1.200596	0.361262
as.factor(arm)1	-0.615885	0.155792	-0.921232	-0.310538

⁸In fact, there are three different versions of the RP model; the one presented here is the “proportional hazard”, which can be seen as an extension to the basic Weibull PH model. The other two versions extend the log-Logistic model by setting $\log(1/S(t) - 1) = \eta_i$ (“proportional odds”) and the log-Normal model using $\Phi^{-1}(S(t)) = \eta_i$ (“probit model”). Currently, all are implemented in **flexsurv**, while only the PH model is implemented in **survHE** using HMC.

Model fitting summaries

Akaike Information Criterion (AIC)....: 1205.235

Bayesian Information Criterion (BIC)...: 1224.761

and

Model fit for the Royston & Parmar splines model, obtained using Stan (Bayesian inference via Hamiltonian Monte Carlo). Running time:

55.507 seconds

	mean	se	L95%	U95%
gamma0	-4.656390	0.463923	-5.605131	-3.810731
gamma1	2.502764	0.545411	1.505152	3.619227
gamma2	0.366841	0.321249	-0.235608	1.032577
gamma3	-0.394103	0.387038	-1.175527	0.330248
as.factor(arm)1	-0.621653	0.152526	-0.907300	-0.335945

Model fitting summaries

Akaike Information Criterion (AIC)....: 1205.624

Bayesian Information Criterion (BIC)...: 1225.150

Deviance Information Criterion (DIC)...: 1259.027

The Bayesian version of the RP model is much more computationally intensive, although HMC does a very good job and keeps the time to generally acceptable levels; also, it helps in this case to take advantage of the multi-processing capability of **rstan**: adding the option `cores = 2` reduces the time from 57.851 to 32.777 seconds, for 2 chains of 2000 iterations each.

In this case, there is very good agreement in the point and interval estimates for the two versions of the model, but in general the MLE may underestimate the underlying level of correlation among the γ coefficients in particular.

8.2. Poly-Weibull

In a nutshell, the Poly-Weibull model (Berger and Sun 1993; Demiris, Lunn, and Sharples 2015) extends the basic set up of a Weibull survival model by accounting for the possibility that in fact the observed times are the result of a mixed data generating process, depending on several independent Weibull components. For example, we may consider that the occurrence of the event under study depends on M independent causes and that we are willing to model each using a suitable Weibull distribution. In line with (1), the resulting density is

$$\begin{aligned} f(t_i | \boldsymbol{\theta}) &= h(t_i)^{d_i} S(t_i) \\ &= \left[\sum_{m=1}^M \alpha_m \mu_{im} t_i^{\alpha_m - 1} \right]^{d_i} \left[\exp \left(- \sum_{m=1}^M \mu_{im} t_i^{\alpha_m} \right) \right], \end{aligned}$$

where $\boldsymbol{\theta} = (\boldsymbol{\theta}_1, \dots, \boldsymbol{\theta}_M)$ and $\boldsymbol{\theta}_m = (\alpha_m, \mu_{im})$ are the shape and scale for the m th component of the mixture.

survHE implements this density as an add-on model to be estimated using HMC and **rstan** (i.e., there is only a Bayesian version for this model). While it is in principle possible to model

both the shape and the scale as functions of a set of covariates, **survHE** considers the simpler version where only the location parameter is allowed to depend on \mathbf{X} . It is fairly easy to modify this structure and implement a version of the **rstan** model in which also α_m depends on the covariates.

Practically, **survHE** has a specific function to run the Poly-Weibull model. A typical call is as in the following.

```
R> formula.pw <- list(Surv(time, event) ~ 1,
+   Surv(time, event) ~ as.factor(arm))
R> m8 <- poly.weibull(formula.pw, data, cores = 4)
```

The main difference with respect to the standard call to `fit.models` is that the formula input now needs to be made by a *list* of objects of class ‘formula’. This is because we need to specify a formula for each of the components that we want to fit to the mixture model identified by the Poly-Weibull distribution. For instance, in the above example, the object `formula.pw` is a list with two elements. This instructs R to assume a model with $M = 2$ components and the following specification for the linear predictors

$$\mu_{i1} = \exp(\beta_{01}) \quad \text{and} \quad \mu_{i2} = \exp(\beta_{02} + \beta_{12}\text{Arm}_i).$$

By default, **survHE** places minimally informative priors on the parameters β_m by using $\beta_m \stackrel{\text{iid}}{\sim} \text{Normal}(0, 10)$; the values for the mean (`mu_beta`) and the standard deviation (`sigma_beta`) can be modified using the option `prior`, as shown earlier. In addition, we need to impose an identifiability constraint on the shape parameters $\alpha = (\alpha_1, \dots, \alpha_M)$ so that they are ordered (i.e., $\alpha_1 < \dots < \alpha_M$) – see Demiris *et al.* (2015) for a discussion of this issue. The components of the vector α are then given a vague prior over their entire domain.

As is often the case for mixture models, convergence to the posterior distributions may be a crucial issue. This is essentially due to the fact that it may be very difficult (or even impossible) for the model to distinguish two or more of the components. For example, the above call to `poly.weibull` returns the following warning

```
Warning: There were 7 divergent transitions after warmup. See
http://mc-stan.org/misc/warnings.html#divergent-transitions-after-warmup
to find out why this is a problem and how to eliminate them.
```

```
Warning: Examine the pairs() plot to diagnose sampling problems
```

```
Warning: The largest R-hat is 1.12, indicating chains have not mixed.
Running the chains for more iterations may help. See
http://mc-stan.org/misc/warnings.html#r-hat
```

```
Warning: Bulk Effective Samples Size (ESS) is too low, indicating posterior
means and medians may be unreliable.
Running the chains for more iterations may help. See
http://mc-stan.org/misc/warnings.html#bulk-ess
```

```
Warning: Tail Effective Samples Size (ESS) is too low, indicating posterior
```

variances and tail quantiles may be unreliable. Running the chains for more iterations may help. See <http://mc-stan.org/misc/warnings.html#tail-ess>

indicating that the HMC algorithm has failed to fully explore the posterior density of the parameters. Possible solutions are to either include more information in the priors (e.g., by reducing the range of variation of the coefficients in the β_m), or to increase the value of the acceptance rate (see Section 5.2). For example, the command

```
R> m9 <- poly.weibull(formula.pw, data, cores = 4, control = list(
+   adapt_delta = .95, stepsize = .005, max_treedepth = 100), iter = 4000)
```

modifies the standard **rstan** settings to have a denser discrete approximation of the continuous Hamiltonian dynamics⁹. This increases the running time from 79.608 to 225.776 seconds, but successfully estimates the posterior distributions. The results can be analyzed by using the `print` method.

```
R> print(m9)
```

```
Model fit for the Poly-Weibull model, obtained using Stan
(Bayesian inference via Hamiltonian Monte Carlo). Running time:
240.589 seconds
```

	mean	se	L95%	U95%
shape_1	1.38546	0.518854	0.132859	1.939147
shape_2	2.03571	0.318739	1.652758	2.868474
(Intercept)_1	-7.59853	4.929838	-21.473366	-4.076481
(Intercept)_2	-5.24629	1.140821	-7.920003	-3.852474
as.factor(arm)1_2	-4.87825	5.614111	-19.243445	-0.385614

```
Model fitting summaries
```

```
Akaike Information Criterion (AIC)....: 1216.545
Bayesian Information Criterion (BIC)...: 1236.071
Deviance Information Criterion (DIC)...: 1195.041
```

The resulting table identifies the parameters of each component by using a suffix “_m”, so for instance `(Intercept)_1` is the intercept for the first component.

Notice that in this particular case, using a more complex model does not seem to improve things substantially: for example, the DIC for the simpler Weibull model we fitted in the

⁹In brief, a single iteration of HMC proceeds to update the value of the parameters for `max_treedepth` steps before deciding whether to accept or reject the current value. Each step is scaled by a factor `stepsize`, which determines the level of the discrete approximation to the underlying continuous Hamiltonian dynamics. Thus, increasing the target acceptance rate `adapt_delta` effectively forces smaller step sizes, meaning that the algorithm will be in principle more thorough in visiting the posterior density (at the expense of computational time). Similarly, increasing `max_treedepth` means that the algorithm will take more samples within each step, thus potentially being more accurate. In practice, it is important to strike a balance and avoid too fine a discretization of the Hamiltonian dynamics, while guaranteeing sufficient coverage of the posterior density. More technical details can be found in Hoffman and Gelman (2014) and Carpenter *et al.* (2015).

second element of the object `m3` is 1203.152 (cfr. Figure 5c) – essentially the same as for the Poly-Weibull model, indicating that a simpler version is perhaps to be preferred to *this* specification of the more complex one. This is consistent with the potential issues in convergence and identifiability, which again indicate perhaps that the Poly-Weibull model may not be appropriate for this particular set of data.

9. Other tools

9.1. Digitizing Kaplan-Meier curves from published studies

Often, the individual level data from, e.g., a clinical trial measuring the survival times are not directly available for the health economic evaluation. Perhaps for one of the treatments being considered, the sponsor of the trial is able to make the data available, but this could only cover one of the relevant interventions/drugs. To overcome this limitation, usually modelers try and use systematic reviews of the available literature to gather information on the possible comparators.

One clever way of doing this is by “digitizing” Kaplan-Meier data available from published papers. This is done by using specific software (e.g., **DigitizeIt**; Bormann 2013); the user needs to click on several points on the survival curves and the values are digitized and exported to some output files, describing the input survival times from graph reading and reported number of individuals at risk at several time points in the follow up.

Once these two files are available, **survHE** takes them as input and following the algorithm developed by Guyot, Ades, Ouwens, and Welton (2012), which can be used to map from the digitized curves back to the unobserved Kaplan-Meier data by numerical approximation. Assuming that the **DigitizeIt** outcome is saved in the two files `survival.txt` and `nrisk.txt` in the current working directory, a typical **survHE** call to perform this task is the following.

```
R> surv_inp <- "survival.txt"
R> nrisk_inp <- "nrisk.txt"
R> km_out <- "KMdata.txt"
R> ipd_out <- "IPDdata.txt"
R> digitise(surv_inp = surv_inp, nrisk_inp = nrisk_inp, km_output = km_out,
+   ipd_output = ipd_out)
```

The above code first defines the input and output files and then uses the **survHE** command `digitise` to reconstruct the original Kaplan-Meier curves and save their values, as well as a fictional dataset, which closely resembles the one that has generated the published survival curves. `digitise` also writes to the R terminal the following text.

```
Kaplan Meier data written to file: KMdata.txt
IPD data written to file: IPDdata.txt
```

These can be then used as input data to fit the survival models and then perform the PSA. The final feature worth mentioning is that often we will be in a position of creating several individual level datasets to mimic data originally used to produce Kaplan-Meier estimates for the survival curves, e.g., for many different treatments, or for the same treatment observed in

several papers. **survHE** has another specialized function that can be used to stack the different files with the individual level data into a single dataset. This function is called `make.ipd` and it takes as inputs: a list of all the names of the individual level data files created as output of `digitise`; the index of the file associated with the control arm (by default, this is the first file) and the control arm will be coded as 0; and a vector of labels for the column of the resulting data matrix. These should match the arguments to the formula specified for the function `fit.models` and should be 3 elements (each representing the time variable, the event variable and the treatment arm). Using the output of the digitization process shown above, a call to this function would look like the following.

```
R> data <- make.ipd("IPDdata.txt", ctr = 1, var.labs = c("time", "event",
+ "arm"))
```

This generates a R ‘`data.frame`’, which can then be fed to the `fit.models` function. Naturally, `make.ipd` assumes that there are no other covariates in addition to the treatment arm, because it is unlikely that digitized data are recorded for different strata, or values of additional variables.

Note also that we can use a collection of digitized data and stack them to create a single dataset made by several pseudo-data obtained in this way. For instance, if we had available three digitized datasets, say `IPD1.txt`, `IPD2.txt` and `IPD3.txt`, we could create a single dataset using a command similar to the following.

```
R> ipd_files <- list("IPD1.txt", "IPD2.txt", "IPD3.txt")
R> data <- make.ipd(ipd_files, ctr = 1, var.labs = c("time", "event", "arm"))
```

This generates a R ‘`data.frame`’, which can then be fed to the `fit.models` function. Naturally, `make.ipd` assumes that there are no other covariates in addition to the treatment arm, because it is unlikely that digitized data are recorded for different strata, or values of additional variables.

10. Limitations and current/future developments of **survHE**

In this paper we have presented a comprehensive tutorial into how **survHE** can help modelers in developing survival analyses, with the specific aim of conducting a health economic evaluation. The main aim of **survHE** is to provide a suite of standardized commands, allowing the user to perform the analysis both within a frequentist (via **flexsurv**) and a Bayesian approach (through **INLA** or **rstan**). One of the main potential advantages to using **survHE** is that much of the “entry cost” in using a Bayesian model may be reduced by how **survHE** guides the user in forming the model, using commands that are extremely similar to those applied to frequentist analyses. We note, however, that this is and should not be a substitute for a thorough understanding of the theory and assumptions underpinning either of these methods. The current version of **survHE** is robust, comprehensive and fully tested and it already allows for a wide range of model specifications. However, there are other venues that we wish to explore.

At present, the main limitations of **survHE** are related to improvements to the range of models that can be fitted within a Bayesian approach. For example, we are currently working

to expand the range of **INLA** survival models to cover all the parametric distributions described in Table 1. Even more importantly, some of the advanced tools such as RPS models are particularly suitable to an **INLA** implementation. We expect to develop all versions of RPS within **INLA**, which will allow modelers to use them under a Bayesian approach and in a computationally efficient manner. At the same time, we also plan to expand the current HMC implementation to include other versions of the RPS model, as well as different forms of censoring (including left and interval). We will also explore the functionality of alternative Bayesian approximation inferential methods, e.g., “variational Bayes” (Kucukelbir, Ranganath, Gelman, and Blei 2015), which are already implemented in **rstan**.

In terms of **rstan**, it is interesting to note the recent surge of interest for survival modeling, within the HMC community, which has also been stimulated by our work with **survHE**, as evidenced by the discussion in the **rstan** forum¹⁰. Related to this, we aim at also expanding the Poly-Weibull model by implementing a version that allows the inclusion of covariates on the ancillary parameters as well as on the location. This effort is also part of a wider plan to bring into **survHE** more models based on mixtures, e.g., in the form of cure models (Amico and Van Keilegom 2018). This may result in an efficient way of accounting for the underlying complexity of the data generating process, as well as extrapolating the survival curves, potentially including external data.

As a general principle, we see the development of **survHE** as continuously ongoing – in fact this is the reason why we have established the GitHub repository; our aim is to expand the community of users and continuously liaise with them, so as to implement more and more models to keep the package up-to-date with the latest methodological developments.

Acknowledgments

The author wishes to thank Peter Konings, William Browne, Geoff Holmes and Andrea Berardi for providing comments or help in writing part of the original code in **survHE**. I would also like to thank two anonymous reviewers for their helpful comments and suggestions. This work has been partially supported by a research grant sponsored at University College London by Mapi ICON, a consultancy company working in the area of health economic evaluation.

References

- Amico M, Van Keilegom I (2018). “Cure Models in Survival Analysis.” *Annual Review of Statistics and Its Application*, **5**(1), 311–342. doi:10.1146/annurev-statistics-031017-100101.
- Baio G (2012). *Bayesian Methods in Health Economics*. 1st edition. Chapman & Hall/CRC, Boca Raton.
- Baio G (2020). **survHE**: *Survival Analysis in Health Economic Evaluation*. R package version 1.1.1, URL <https://CRAN.R-project.org/package=survHE>.
- Baio G, Berardi A, Heath A (2017). *Bayesian Cost-Effectiveness Analysis with the R Package BCEA*. 1st edition. Springer-Verlag, New York, NY. doi:10.1007/978-3-319-55718-2.

¹⁰See <https://discourse.mc-stan.org/t/survival-models-in-rstanarm/3998>.

- Baio G, Dawid P (2011). “Probabilistic Sensitivity Analysis in Health Economics.” *Statistical Methods in Medical Research*, **24**(6), 615–634. doi:10.1177/0962280211419832.
- Baio G, Heath A (2016). “When Simple Becomes Complicated: Why **Excel** Should Lose Its Place at the Top Table.” *Global & Regional Health Technology Assessment*. doi:10.5301/grhta.5000247.
- Benaglia T, Jackson CH, Sharples LD (2015). “Survival Extrapolation in the Presence of Cause Specific Hazards.” *Statistics in Medicine*, **34**(5), 796–811. doi:10.1002/sim.6375.
- Berger J, Sun D (1993). “Bayesian Analysis for the Poly-Weibull Distribution.” *Journal of the American Statistical Association*, **88**(424), 1412–1418. doi:10.1080/01621459.1993.10476426.
- Betancourt M (2012). “Cruising the Simplex: Hamiltonian Monte Carlo and the Dirichlet Distribution.” *AIP Conference Proceedings*, **1443**(1), 157–164. doi:10.1063/1.3703631.
- Blangiardo M, Cameletti M (2015). *Spatial and Spatio-Temporal Bayesian Models with R-INLA*. 1st edition. John Wiley & Sons, Chichester. doi:10.1002/9781118950203.
- Blangiardo M, Cameletti M, Baio G, Rue H (2013). “Spatial and Spatio-Temporal Models with R-INLA.” *Spatial and Spatio-temporal Epidemiology*, **4**, 33–49. doi:10.1016/j.sste.2012.12.001.
- Bormann I (2013). *DigitizeIt 2.4*. Braunschweig, Germany. URL <https://www.digitizeit.de/>.
- Briggs A, Sculpher M, Claxton K (2006). *Decision Modelling for Health Economic Evaluation*. 1st edition. Oxford University Press, Oxford.
- Carpenter B, Gelman A, Hoffman M, Lee D, Goodrich B, Betancourt M, Brubaker MA, Guo J, Li P, A R (2015). “Stan: A Probabilistic Programming Language.” *Journal of Statistical Software*, **76**(1), 1–32. doi:10.18637/jss.v076.i01.
- Demiris N, Lunn D, Sharples L (2015). “Survival Extrapolation Using the Poly-Weibull Model.” *Statistical Methods in Medical Research*, **24**(2), 287–301. doi:10.1177/0962280211419645.
- Filipović-Pierucci A, Zarca K, Durand-Zaleski I (2017). “Markov Models for Health Economic Evaluation: The R Package **heemod**.” arXiv:1702.03252 [stat.AP], URL <https://arxiv.org/abs/1702.03252>.
- Gabrio A, Mason A, Baio G (2019). “A Full Bayesian Model to Handle Structural Ones and Missingness in Economic Evaluations from Individual-Level Data.” *Statistics in Medicine*, **38**(8), 1399–1420. doi:10.1002/sim.8045.
- Gelman A, Carlin J, Stern H, Dunson D, Vehtari A, Rubin D (2013). *Bayesian Data Analysis*. 3rd edition. Chapman & Hall/CRC, Boca Raton. doi:10.1201/b16018.
- Guyot P, Ades AE, Beasley M, Lueza B, Pignon JP, Welton NJ (2017). “Extrapolation of Survival Curves from Cancer Trials Using External Information.” *Medical Decision Making*, **37**(4), 353–366. doi:10.1177/0272989x16670604.

- Guyot P, Ades AE, Ouwens MJ, Welton NJ (2012). “Enhanced Secondary Analysis of Survival Data: Reconstructing the Data from Published Kaplan-Meier Survival Curves.” *BMC Medical Research Methodology*, **12**, 9. doi:10.1186/1471-2288-12-9.
- Harrell Jr FE (2020). *rms: Regression Modeling Strategies*. R package version 6.0-1, URL <https://CRAN.R-project.org/package=rms>.
- Heath A, Manolopoulou I, Baio G (2017). “Efficient Monte Carlo Estimation of the Expected Value of Sample Information Using Moment Matching.” *Medical Decision Making*, **38**(2), 163–173. doi:10.1177/0272989x17738515.
- Hoffman M, Gelman A (2014). “The No-U-Turn Sampler: Adaptively Setting Path Lengths in Hamiltonian Monte Carlo.” *Journal of Machine Learning Research*, **15**(1), 1593–1623.
- Incerti D, Jansen JP (2020). *hesim: Health-Economic Simulation Modeling and Decision Analysis*. R package version 0.3.1, URL <https://CRAN.R-project.org/package=hesim>.
- Incerti D, Thom H, Baio G, Jansen J (2019). “R You Still Using **Excel**? Why We Need Modern Software Tools for Health Technology Assessment.” *Value in Health*, **22**(5), 575–579. doi:10.1016/j.jval.2019.01.003.
- Jackson C (2016). “**flexsurv**: A Platform for Parametric Survival Modelling in R.” *Journal of Statistical Software*, **70**(8), 1–33. doi:10.18637/jss.v070.i08.
- Jackson C, Stevens J, Ren S, Latimer N, Bojke L, Manca A, Sharples L (2017). “Extrapolating Survival from Randomized Trials Using External Data: A Review of Methods.” *Medical Decision Making*, **37**, 377–390. doi:10.1177/0272989x16639900.
- Krijkampet E, Alarid-Escudero F, Enns E, Jalal H, Hunink M, Pechlivanoglou P (2018). “Microsimulation Modeling for Health Decision Sciences Using R: A Tutorial.” *Medical Decision Making*, **38**(3), 400–422. doi:10.1177/0272989x18754513.
- Kruschke J (2015). *Doing Bayesian Data Analysis: A Tutorial with R, JAGS, and Stan*. 2nd edition. Academic Press.
- Kucukelbir A, Ranganath R, Gelman A, Blei D (2015). “Automatic Variational Inference in Stan.” arXiv:1506.03431 [stat.ML], URL <http://arxiv.org/abs/1506.03431>.
- Latimer N (2011). “NICE DSU Technical Support Document 14.” *Technical report*, National Institute for Health and Care Excellence, Decision Support Unit. URL <http://nicedsu.org.uk/wp-content/uploads/2016/03/NICE-DSU-TSD-Survival-analysis.updated-March-2013.v2.pdf>.
- Lunn D, Spiegelhalter D, Thomas A, Best N (2009). “The BUGS Project: Evolution, Critique and Future Directions.” *Statistics in Medicine*, **28**(25), 3049–3067. doi:10.1002/sim.3680.
- Martins T, Simpson D, Lindgren F, Rue H (2013). “Bayesian Computing with **INLA**: New Features.” *Computational Statistics & Data Analysis*, **67**, 68–83. doi:10.1016/j.csda.2013.04.014.
- NICE (2013). “Guide to the Methods of Technology Appraisal 2013.” *Technical report*, National Institute for Health and Care Excellence. URL <https://www.nice.org.uk/process/pmg9/chapter/foreword>.

- Noble M, Wright G, Smith G, Dibben C (2006). “Measuring Multiple Deprivation at the Small-Area Level.” *Environment and Planning A*, **38**(1), 169–185.
- Plummer M (2003). “**JAGS**: A Program for Analysis of Bayesian Graphical Models Using Gibbs Sampling.” In K Hornik, F Leisch, A Zeileis (eds.), *Proceedings of the 3rd International Workshop on Distributed Statistical Computing (DSC 2003)*. Technische Universität Wien, Vienna, Austria. URL <https://www.R-project.org/conferences/DSC-2003/Proceedings/Plummer.pdf>.
- Stan Development Team (2016). “**shinystan**: Interactive Visual and Numerical Diagnostics and Posterior Analysis for Bayesian Models.” URL <http://mc-stan.org/>.
- Radford N (2011). “MCMC Using Hamiltonian Dynamics.” In S Brooks, A Gelman, G Jones, XL Meng (eds.), *Handbook of Markov Chain Monte Carlo*. Chapman & Hall/CRC, Boca Raton.
- R Core Team (2020). *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria. URL <https://www.R-project.org/>.
- Royston P, Lambert P (2011). *Flexible Parametric Survival Analysis Using Stata: Beyond the Cox Model*. 1st edition. Stata Press.
- Royston P, Parmar M (2002). “Flexible Parametric Proportional-Hazards and Proportional Odds Models for Censored Survival Data, with Application to Prognostic Modelling and Estimation of Treatment Effects.” *Statistics in Medicine*, **21**(1), 2175–2197. doi:10.1002/sim.1203.
- Rue H, Held L (2005). *Gaussian Markov Random Fields: Theory and Applications*. 1st edition. Chapman & Hall/CRC, Boca Raton.
- Rue H, Martino S, Chopin N (2009). “Approximate Bayesian Inference for Latent Gaussian Models Using Integrated Nested Laplace Approximations.” *Journal of the Royal Statistical Society B*, **71**(2), 319–392. doi:10.1111/j.1467-9868.2008.00700.x.
- SAS Institute Inc (2013). *The SAS System, Version 9.4*. SAS Institute Inc., Cary. URL <http://www.sas.com/>.
- Siebert U, Alagoz O, Bayoumi AM, Jahn B, Owens DK, Cohen DJ, Kuntz KM (2012). “State-Transition Modeling: A Report of the ISPOR-SMDM Modeling Good Research Practices Task Force-3.” *Medical Decision Making*, **32**(5), 690–700. doi:10.1177/0272989x12455463.
- Spiegelhalter D, Best N, Carlin B, Van Der Linde A (2002). “Bayesian Measures of Model Complexity and Fit.” *Journal of the Royal Statistical Society B*, **64**(4), 583–639. doi:10.1111/1467-9868.00353.
- StataCorp (2019). *Stata Statistical Software: Release 16*. StataCorp LLC, College Station. URL <http://www.stata.com/>.
- Strong M, Oakley J, Brennan A (2014). “Estimating Multi-Parameter Partial Expected Value of Perfect Information from a Probabilistic Sensitivity Analysis Sample: A Non-Parametric Regression Approach.” *Medical Decision Making*, **34**(3), 311–326. doi:10.1177/0272989x13505910.

- Wang X, Ryan Y, Faraway J (2018). *Bayesian Regression Modeling with INLA*. 1st edition. Chapman and Hall/CRC.
- Wickham H (2016). *ggplot2: Elegant Graphics for Data Analysis*. Springer-Verlag, New York. doi:10.1007/978-0-387-98141-3.
- Willan A, Briggs A (2006). *The Statistical Analysis of Cost-Effectiveness Data*. 1st edition. John Wiley & Sons, Chichester. doi:10.1002/0470856289.
- Williams C, Lewsey J, Briggs A, Mackay D (2017). “Estimation of Survival Probabilities for Use in Cost-Effectiveness Analysis: a Comparison of a Multi-State Modelling Survival Analysis Approach with Partitioned Survival and Markov Decision-Analytic Modelling.” *Medical Decision Making*, **37**(4), 427–439. doi:10.1177/0272989x16670617.
- Woods B, Sideris E, Palmer S, Latimer N, Soares M (2017). “NICE DSU Technical Support Document 19.” *Technical report*, National Institute for Health and Care Excellence, Decision Support Unit. URL <http://nicedsu.org.uk/wp-content/uploads/2017/06/Partitioned-Survival-Analysis-final-report.pdf>.

Affiliation:

Gianluca Baio
Department of Statistical Science
University College London
Gower Street, London, WC1E 6BT, United Kingdom
E-mail: g.baio@ucl.ac.uk
URL: <http://www.statistica.it/gianluca>