

Additive Bayesian variable selection under censoring and misspecification

David Rossell, Francisco J. Rubio

Universitat Pompeu Fabra and University College London

Abstract. We discuss the role of misspecification and censoring on Bayesian model selection in the contexts of right-censored survival and concave log-likelihood regression. Misspecification includes wrongly assuming the censoring mechanism to be non-informative. Emphasis is placed on additive accelerated failure time, Cox proportional hazards and probit models. We offer a theoretical treatment that includes local and non-local priors, and a general non-linear effect decomposition to improve power-sparsity trade-offs. We discuss a fundamental question: what solution can one hope to obtain when (inevitably) models are misspecified, and how to interpret it? Asymptotically, covariates that do not have predictive power for neither the outcome nor (for survival data) censoring times, in the sense of reducing a likelihood-associated loss, are discarded. Misspecification and censoring have an asymptotically negligible effect on false positives, but their impact on power is exponential. We show that it can be advantageous to consider simple models that are computationally practical yet attain good power to detect potentially complex effects, including the use of finite-dimensional basis to detect truly non-parametric effects. We also discuss algorithms to capitalize on sufficient statistics and fast likelihood approximations for Gaussian-based survival and binary models.

Key words and phrases: Additive regression, Generalized additive model, Misspecification, Model selection, Survival.

1. INTRODUCTION

Determining what covariates have an effect on a survival (time-to-event) outcome is an important task in many fields, including Biomedicine, Economics and Engineering. For interpretability and computational convenience it is common to use parametric and semi-parametric models such as Cox proportional hazards [7] or accelerated failure time (AFT) regression for survival outcomes, possibly with non-linear additive effects. The proportional hazards model assumes that covariates have a multiplicative effect on the baseline hazard, whereas in AFT models covariates drive the mean of the logarithmic (or other monotonic transform) times-to-event. These models can be combined with

Universitat Pompeu Fabra, Department of Business and Economics, Barcelona, Spain; Barcelona Graduate School of Economics, Data Science Center, Barcelona, Spain.(e-mail: david.rossell@upf.edu). University College London, Department of Statistical Science, London, WC1E 6BT, UK (e-mail: f.j.rubio@ucl.ac.uk).

Bayesian model selection to provide a powerful mechanism to select variables, enforce sparsity and quantify uncertainty. However, the precise consequences of model misspecification and censoring are not sufficiently understood. By misspecification we mean that the data are truly generated by a distribution outside the considered class. For instance, one may fail to record truly relevant covariates or represent their effects imperfectly, *e.g.* when using a Cox model but the true covariate effects on the hazard are non-proportional. This issue can be addressed by enriching the model, *e.g.* via non-parametric or time-dependent effects. Then, the potential concerns are that the larger number of parameters can adversely affect inference, unless the sample size is large enough, and that computation can be costlier. Censoring is also important. First, it reduces the effective sample size. Second, wrongly assuming the censoring mechanism to be non-informative, *i.e.* independent of the outcome conditionally on covariates, may affect the selected model, even asymptotically.

Our goal is to help understand the consequences of three important issues on model selection: misspecification, censoring, and trade-offs when including non-linear effects. We first consider that the data analyst assumed a non-linear additive AFT model, or an additive Cox model, but data are truly generated by a different probability distribution F_0 . We also consider probit regression, which can be formulated as a particular case of the Normal AFT model, and more general concave log-likelihood regression, which provides a unifying framework for the models we consider here.

There are many data analysis methods for survival outcomes, along with theory for well-specified models and empirical results suggesting potential issues under misspecification, but their implications for model selection have not been described in sufficient detail. We first review results on the behavior of misspecified AFT and proportional hazard models, and subsequently discuss some model selection methods for survival data. Although both models have similar asymptotic properties, and which model is more appropriate depends on the data at hand, AFT inference has been argued to be more robust and to better preserve interpretability under misspecification. More precisely, the maximum likelihood estimators under misspecified Cox and AFT models have comparable limiting distributions if censoring is absent or independent of covariates [49, 56], but not so under covariate-dependent censoring [47]. Covariate-dependent censoring also affects frequentist hypothesis tests. In misspecified Cox models it can lead to a substantial type I error inflation [8]. In misspecified AFT models the power of the tests may be affected, but simple strategies to control the type I error are available [47, 20, 15].

Another situation where both models behave differently is when omitting truly active covariates, *e.g.* because these were not recorded. A proportional hazards model with omitted variables tends to underestimate covariate effects, even for a treatment of interest that is uncorrelated with other covariates [49, 25]. Further, even if the data-generating truth has proportional hazards, the marginal model conditioning only on the observed covariates does not (except in positive stable distributions, [18]). In contrast, if the data-generating truth is an AFT model and one omits relevant covariates, the unaccounted variability is subsumed into the error term, and regression parameters remain interpretable as averaged effects across the population [20]. Note that omitting covariates is intimately connected to incorporating a covariate but misspecifying its effect: using a linear or finite-dimensional effect can be seen as omitting a subset of the columns of the basis defining a truly non-parametric effect. Thus our discussion on omitted variables applies directly to misspecifying covariate effects. To summarize, censoring and misspecification have non-trivial effects on estimation and hypothesis testing.

We now review some model selection methods for survival data, discussing the extent to which they considered misspecification. [19, 46, 34] proposed likelihood penalties for Cox and semiparametric AFT models, and [51] for broader generalized hazards models. Most of this work focused on linear covariate effects, computation and proving consistency under covariate-independent censoring. There are however empirical results on the effect of misspecification, *e.g.* [58] showed in simulations an increase in false positives of the Cox-LASSO method [50] when data truly arise from an AFT model. There are also many Bayesian variable selection methods for survival data. [10] and [44] proposed shrinkage priors for the Cox and AFT models and assessed performance via simulations where the model was well-specified. [22] studied Bayesian model selection for the Cox model, [9] for the so-called additive hazards model, and [31] for the Cox model under non-local priors [23, 24]. See [21] for a review, with a focus on the Cox model. While interesting, these Bayesian proposals do not consider misspecification. [38] did study misspecified Bayesian linear regression, showing that misspecification often reduces the power to detect active variables, but did not consider censoring.

We summarize our main messages. We show that, under mild assumptions, Bayesian model selection asymptotically discards covariates that do not predict neither the outcome nor the censoring times. By *predict*, we refer to increasing the expectation of the log-likelihood function. For any fully-specified model, said expectation is a weighted average of a reward for assigning a high probability to the observed censoring time in individuals that are censored, and a reward for predicting survival times accurately in uncensored individuals (*e.g.* mean squared error, for Normal AFT models). For the partially-specified Cox model, the reward is for assigning a high hazard to individuals who experienced the event, relative to other individuals at risk. We discuss that both censoring and wrongly specifying covariate effects have an exponential effect in power, but that asymptotically neither leads to false positive inflation. We also develop a novel non-linear effect decomposition to ameliorate the power drop, and study the consequences of using finite basis to describe covariate effects, a practical strategy to speed up computations when one considers many models. For concreteness, we outline a formulation based on a novel combination of non-local priors [23] and group-Zellner priors that induce group-level sparsity for non-linear effects. As a technical contribution, we prove the asymptotic validity of Laplace approximations to Bayes factors for concave log-likelihoods under minimal conditions, allowing for misspecification, which provides a simple basis to study Bayes factors that covers all models considered in this paper. We also provide software (R package `mombf`).

The outline is as follows. Section 2 discusses the likelihood for AFT and Cox models, priors and a non-linear effect decomposition aimed at improving power. Section 3 discusses known and novel results on asymptotic normality and Bayes factor rates, and how to interpret the Bayesian model selection solution under misspecification. Similar results are obtained for general concave log-likelihoods, see Section S10. See also Section S9 for a known but seemingly unexploited result in the literature, that probit models are a particular case of the Normal AFT model. Section 4 discusses the relative computational convenience of AFT vs. Cox models related to the use of sufficient statistics. It also discusses an approximation to the Normal log-distribution function and derivatives that significantly increases speed and accuracy, and may have some independent interest, and simple model exploration strategies. Section 5 illustrates the effect of misspecification and censoring in simulations and cancer datasets, practical power-sparsity trade-offs, and the use of finite-dimensional non-linear basis. Section 6 concludes. The supplementary

material contains derivations related to the likelihood, priors and their derivatives, and prior elicitation (Sections S1-S3). Sections S4-S6 offer detailed discussions on computational algorithms, including a novel approximation to Normal log-distribution functions that may have some independent interest. Section S7 contains all proofs for our main results, Sections S8-S9 additional propositions for the AFT model with Laplace errors and probit models, and Section S10 on the asymptotic validity of Laplace approximations to integrated likelihoods. Finally, Section S11 contains empirical results that supplement those in the main paper.

2. FORMULATION

Our discussion focuses on survival data, but see Section S9 for binary regression and Section S10 for concave log-likelihoods. Section 2.1 sets notation, reviews the AFT and proportional hazards models, their being a particular cases of the generalized hazards structure, and a non-linear effects decomposition to improve interpretability and power. Section 2.2 embeds the problem within a Bayesian model selection framework. Section 2.3 introduces prior distributions that can accommodate group and hierarchical constraints, and Section 2.4 suggests default prior parameter values.

2.1 Likelihood

Let us introduce the notation. Suppose that one is interested in studying the dependence of a survival (or time-to-event) outcome $o_i \in \mathbb{R}_+$ on a covariate vector $x_i = (x_{i1}, \dots, x_{ip})^\top \in \mathbb{R}^p$, for individuals $i = 1, \dots, n$. Suppose that there are right-censoring times $c_i \in \mathbb{R}_+$, such that one only observes the outcome for uncensored individuals, *i.e.* those for which $o_i \leq c_i$. Denote by $u_i = \mathbb{I}(o_i < c_i)$ the indicator that observation i is uncensored, $y_i = \min\{\log(o_i), \log(c_i)\}$ the observed log-times, $y = (y_1, \dots, y_n)$, $u = (u_1, \dots, u_n)$, and the number of uncensored individuals $n_o = \sum_{i=1}^n u_i$.

We review two popular models for survival data, the AFT and Cox models, and discuss a strategy to decompose non-linear effects. The AFT model postulates

$$\log(o_i) = \sum_{j=1}^p g_j(x_{ij}) + \epsilon_i,$$

where $g_j : \mathbb{R} \rightarrow \mathbb{R}$ and ϵ_i are independent across $i = 1, \dots, n$ with mean $E(\epsilon_i) = 0$ and variance $V(\epsilon_i) = \sigma^2$ (assumed finite). Typically, g_j is expressed in terms of an r -dimensional basis, *e.g.* splines or wavelets [53]. For interpretability and to gain power (see Section 3.2) it is convenient to decompose g_j into a linear and a deviation-from-linearity components. To fix ideas, the cubic splines used in our examples consider

$$(2.1) \quad \log(o_i) = x_i^\top \beta + s_i^\top \delta + \epsilon_i,$$

where $\beta = (\beta_1, \dots, \beta_p)^\top \in \mathbb{R}^p$, $\delta^\top = (\delta_1^\top, \dots, \delta_p^\top) \in \mathbb{R}^{rp}$ and $s_i^\top = (s_{i1}^\top, \dots, s_{ip}^\top)$, where $s_{ij} \in \mathbb{R}^r$ is the projection of x_{ij} onto a cubic spline basis orthogonalized to x_{ij} (and the intercept). The idea is that $x_i^\top \beta$ captures linear effects, whereas $s_i^\top \delta$ captures deviations from linearity. Even if a covariate truly has a non-linear effect, the linear term often captures a fraction of that effect using a single parameter, hence one can gain in power to detect its presence. Specifically we built s_{ij} , the i^{th} row of the $n \times r$ matrix S_j , as follows. Let X_j and \tilde{S}_j have row i equal to $(1, x_{ij})$ and the cubic spline projection of x_{ij} (equi-distant knots), then $S_j = (I - X_j(X_j^\top X_j)^{-1} X_j^\top) \tilde{S}_j$ is orthogonal to X_j . Denote by (X, S) the design matrix with (x_i^\top, s_i^\top) in its i^{th} row, and by (X_o, S_o)

and (X_c, S_c) the submatrices with the rows for uncensored and censored individuals (respectively). This formulation contains partially linear models as particular cases (*i.e.* when only some of the covariates are assumed to have a non-linear effect). We denote the parameter space by $\Gamma \subset \mathbb{R}^{p(r+1)} \times \mathbb{R}^+$.

In survival analysis it is common to pose a model only for the survival times, such as (2.1). This is because the censoring is assumed to be non-informative given the covariates, and then the censoring distribution factors out of the likelihood function (see Section S1 for a brief derivation). The likelihood and partial likelihood used by the AFT and Cox models, which we review next, embed such a non-informativeness assumption. See Section 3 for a discussion on the consequences of this assumption not holding.

Regarding the likelihood associated to (2.1), consider the particular case where the errors are Gaussian. It is convenient to reparameterize $\alpha = \beta/\sigma$, $\kappa = \delta/\sigma$ and $\tau = 1/\sigma$, as then the log-likelihood is concave, provided the number of uncensored individuals is greater than the number of model parameters ($n_o \geq p + rp$) and that (X_o, S_o) has full column rank [2, 45]. The log-likelihood is

$$\begin{aligned} \ell(\alpha, \kappa, \tau) = & -\frac{n_o}{2} \log \left(\frac{2\pi}{\tau^2} \right) - \frac{1}{2} \sum_{u_i=1} (\tau y_i - x_i^\top \alpha - s_i^\top \kappa)^2 \\ (2.2) \quad & + \sum_{u_i=0} \log \left\{ \Phi \left(x_i^\top \alpha + s_i^\top \kappa - \tau y_i \right) \right\}, \end{aligned}$$

see Supplementary (S2.1) for its gradient and hessian.

The Cox model instead assumes that the hazard function at time t takes the form

$$h_{PH}(t | x_i) = h_0(t) \exp \left\{ x_i^\top \beta + s_i^\top \delta \right\},$$

where $h_0(\cdot)$ is a baseline hazard, typically estimated non-parametrically, and (β, δ) are estimated using the log partial likelihood [7]

$$(2.3) \quad \ell_p(\beta, \delta) = \sum_{u_i=1} \log \left(\frac{\exp \{x_i^\top \beta + s_i^\top \delta\}}{\sum_{k \in \mathcal{R}(o_i)} \exp \{x_k^\top \beta + s_k^\top \delta\}} \right),$$

where $\mathcal{R}(t) = \{i : o_i \geq t\}$ denotes the set of individuals at risk at time t .

To relate both models, (2.1) can be formulated in terms of the hazard function $h_{AFT}(t) = h_0(t \exp \{x_i^\top \beta + s_i^\top \delta\}) \exp \{x_i^\top \beta + s_i^\top \delta\}$. Both models are special cases of the generalized hazards structure [5]

$$(2.4) \quad h_{GH}(t) = h_0 \left(t \exp \{x_i^\top \beta + s_i^\top \delta\} \right) \exp \{x_i^\top \theta + s_i^\top \xi\},$$

which we use in our examples to portray the behaviour of misspecified AFT and Cox models. Clearly, (2.4) contains the AFT model for $(\beta, \delta) = (\theta, \xi)$ and the proportional hazards model for $(\beta, \delta) = 0$.

2.2 Model selection

Our goal is model selection, which we formalize as choosing among three possibilities

$$\gamma_j = \begin{cases} 0, & \text{if } \beta_j = 0, \delta_j = 0, \\ 1, & \text{if } \beta_j \neq 0, \delta_j = 0, \\ 2, & \text{if } \beta_j \neq 0, \delta_j \neq 0, \end{cases}$$

corresponding to no effect, a linear and a non-linear effect of each covariate $j = 1, \dots, p$. That is, $\gamma = (\gamma_1, \dots, \gamma_p)$ determines what covariates enter the model and their effect, and there are 3^p total models to consider. We remark that by non-linear effect we refer to the specific effect coded by the chosen basis, *e.g.* B-splines in our examples. One could extend the exercise by considering other types of non-linear effects, for example by adding a fourth possibility $\gamma_j = 3$ associated to a wavelet basis. Such basis would be orthogonalized to the linear term, as described after (2.1).

This formulation has two key ingredients. First, it decomposes effects into linear and deviation from linearity components, enforcing the hierarchical desiderata that the latter are only included if the linear terms are present. This decomposition is similar to the structured additive regression of [42], the main difference is that they do not test for exact $\beta_j = 0$, $\delta_j = 0$ and that they rely on a spectral decomposition that is less general than our simpler orthogonalization of (X_j, S_j) . Our theory and results show that such decompositions improve the power to detect truly active effects. As discussed, this is because the option $\gamma_j = 1$ captures part of the effect of a variable with a single parameter. The second ingredient is considering the group inclusion of all non-linear coefficients $\delta_j \in \mathbb{R}^r$. The motivation is that including individual entries in δ_j increases the probability of false positives, *e.g.* if j truly had no effect there would be $2^r - 1$ subsets of S_j leading to including j .

Bayesian model selection proceeds as follows. Let $p_\gamma = \sum_{j=1}^p \mathbf{I}(\gamma_j \neq 0)$ be the number of active variables according to model γ , $s_\gamma = \sum_{j=1}^p \mathbf{I}(\gamma_j = 2)$ the number of non-linear effects, and $d_\gamma = p_\gamma + rs_\gamma + 1$ the total number of parameters in γ for AFT models, and $d_\gamma = p_\gamma + rs_\gamma$ for Cox and probit models. (X_γ, S_γ) and $(\beta_\gamma, \delta_\gamma)$ are the corresponding submatrices of (X, S) and subvectors of (β, δ) , and $(X_{o,\gamma}, S_{o,\gamma})$ and $(X_{c,\gamma}, S_{c,\gamma})$ the submatrices of the observed (X_o, S_o) and censored (X_c, S_c) design matrices. One then obtains posterior model probabilities

$$(2.5) \quad \pi(\gamma | y) = \frac{p(y | \gamma)\pi(\gamma)}{\sum_\gamma p(y | \gamma)\pi(\gamma)} = \left(1 + \sum_{\gamma' \neq \gamma} B_{\gamma', \gamma} \frac{\pi(\gamma')}{\pi(\gamma)} \right)^{-1},$$

where $\pi(\gamma)$ is the model prior probability, $B_{\gamma', \gamma} = p(y | \gamma')/p(y | \gamma)$ the Bayes factor between (γ', γ) and

$$p(y | \gamma) = \int p(y | \alpha_\gamma, \kappa_\gamma, \tau) \pi(\alpha_\gamma, \kappa_\gamma, \tau | \gamma) d\alpha_\gamma d\kappa_\gamma d\tau,$$

the integrated likelihood $p(y | \alpha_\gamma, \kappa_\gamma, \tau)$ with respect to a prior density $\pi(\alpha_\gamma, \kappa_\gamma, \tau | \gamma)$. One may choose the model with highest $\pi(\gamma | y)$, variables with high marginal posterior probabilities $\pi(\gamma_j \neq 0 | y)$ and, when the interest is in prediction, use Bayesian model averaging where models are weighted according to $\pi(\gamma | y)$, or alternatively choosing a sparse model giving similar predictions [13]. Either way $\pi(\gamma | y)$ are critical for inference, hence the importance to understand their behavior.

To conclude, we comment upon a practically-relevant computational issue. In additive models, it is common to either let the basis dimension r grow with n and add a regularization term (*e.g.* P-splines), or to learn r from the data (*e.g.* knot selection). Letting r grow with n is interesting theoretically and in prediction problems where one fits a single model, but less so when one considers many models. Large r increases the computational cost (*e.g.* matrix determinants require $r^3/3$ operations) and is often unneeded when the goal is just to detect if a covariate has an effect. Instead one may use a moderate

r , *e.g.* misspecify the predictive-optimal model. The question is then, what answer can one hope to obtain and what are its properties. Our theory and software allow learning r among several fixed values, but in our examples a small $r = 5$ provided better inference at lower cost (particularly for small n , *e.g.* Figure S3, bottom).

2.3 Prior distributions

Although our discussion applies to a wide class of priors, we present three concrete options.

$$\begin{aligned}\pi_L(\alpha_\gamma, \kappa_\gamma, \tau) &= \left[\prod_{\gamma_j \geq 1} N(\alpha_j; 0, g_L n / (x_j^\top x_j)) \prod_{\gamma_j = 2} N(\kappa_j; 0, g_S n (S_j^\top S_j)^{-1}) \right] \pi(\tau) \\ \pi_M(\alpha_\gamma, \kappa_\gamma, \tau) &= \left[\prod_{\gamma_j \geq 1} \frac{\alpha_j^2}{g_M} N(\alpha_j; 0, g_M) \prod_{\gamma_j = 2} N(\kappa_j; 0, g_S n (S_j^\top S_j)^{-1}) \right] \pi(\tau) \\ \pi_E(\alpha_\gamma, \kappa_\gamma, \tau) &= \left[\prod_{\gamma_j \geq 1} e^{\sqrt{2} - g_E / \alpha_j^2} N(\alpha_j; 0, g_E) \prod_{\gamma_j = 2} N(\kappa_j; 0, g_S n (S_j^\top S_j)^{-1}) \right] \pi(\tau),\end{aligned}$$

where $\pi(\tau) = 2\tau^{-3} \text{IG}(\tau^{-2}; a_\tau/2, b_\tau/2)$, and IG denotes the inverse gamma density, and $g_L, g_S, g_M, g_E, a_\gamma, b_\tau \in \mathbb{R}_+$ are given dispersion parameters, for which we propose default values in Section 2.4.

These choices include a standard Normal prior and two variations of non-local priors. The use of non-local priors can be argued from a foundational viewpoint, where one wishes to assign prior beliefs that are coherent with the parameters assumed non-zero by a given model [23]. For our purposes, however, their main role is that they lead to faster Bayes factor rates to discard spurious parameters. See [39] for further discussion. We refer to π_L as group-Zellner prior. It is a product of Zellner priors across groups of linear and non-linear terms for each covariate. This prior is local, *i.e.* it assigns non-zero density to α_γ having zeroes. The Zellner structure is chosen for simplicity, our theory can be easily extended to other local priors, provided they are continuous and positive at the asymptotically-optimal parameter values (Section 3, [23]). The priors π_M and π_E are non-local with respect to α_γ , the so-called product MOM and eMOM priors introduced in [24, 40], and a group-Zellner prior on κ_γ .

Regarding the prior on the models $\pi(\gamma)$, we consider joint group inclusion of non-linear coefficients δ_j and the hierarchical restriction that their inclusion requires that of the corresponding linear coefficient β_j . Letting $\pi(\gamma)$ depend only on the number of non-zero parameters in $(\beta_\gamma, \delta_\gamma)$, as customarily done when only linear effects are considered, would ignore such structure and hence be inadequate. Instead, we let $\pi(\gamma)$ depend on the number of variables having linear and non-linear effects, (p_γ, s_γ) . By default, we consider independent Beta-Binomial priors [43]

$$(2.6) \quad \pi(\gamma) = \frac{1}{C} \text{BetaBin}(p_\gamma; p, a_1, b_1) \binom{p}{p_\gamma}^{-1} \text{BetaBin}(s_\gamma; s, a_2, b_2) \binom{s}{s_\gamma}^{-1},$$

where $\text{BetaBin}(z; p, a, b)$ is the probability of z successes under a Beta-Binomial distribution with p trials and parameters (a, b) and C a normalizing constant that does not need to be computed explicitly. Any model such that the number of parameters is $p_\gamma + rs_\gamma > n$ is assigned $\pi(\gamma) = 0$, as it would result in data interpolation. By default we let $a_1 = b_1 = a_2 = b_2 = 1$ akin to [43], *e.g.* in the $p = 1$ case these give

$\pi(\gamma_1 = 0) = \pi(\gamma_1 = 1) = \pi(\gamma_1 = 2) = 1/3$. As alternatives to (2.6), one can also consider Binomial priors where $\text{BetaBin}(z; p, a_j, b_j)$ is replaced by $\text{Bin}(z; p, a_j)$ for a given success probability $a_j \in [0, 1]$ and Complexity priors [4] where it is replaced by $1/p^{a_j z}$ for some constant $a_j > 0$. These two alternatives are implemented in our software and covered by our theory in Section 3, but for simplicity our examples focus on (2.6).

2.4 Prior elicitation

The prior dispersion parameters (g_L, g_M, g_E, g_S) are important for variable selection. For instance, setting large dispersions helps induces sparsity, particularly when they are allowed to grow with the sample size n [30]. However such large values also reduce power, see [36] and our Propositions 3, 4 and S5, and are harder to justify from the point of view that the expected effect sizes a priori should not depend on n . We briefly discuss default values that do not depend on n , and refer the reader to Section S3 for details.

Specifying prior parameters provides an opportunity to define what effects are practically relevant. Importantly, in what follows we assume that continuous covariates were standardized to unit variance, else the parameter interpretation and default values change. Basic considerations give a fairly narrow range of values that we deem reasonable in applications. For example, in AFT and Cox models $e^{|\beta_j|}$ define the effect size, when these are say $< 15\%$ (i.e. $e^{|\beta_j|} < 1.15$) they are typically practically irrelevant. Based on these considerations, our recommended defaults for AFT and Cox models are $g_M = 0.192$, $g_E = 0.091$, $g_L = 1$, $g_S = 1/r$ and $a_\tau = b_\tau = 3$, whereas for probit regression they are $g_M = 0.139$ and $g_E = 0.048$. One should not take these defaults at their exact value, rather as defining a range of reasonable values. These ranges are discussed in Section S3. In our examples, results were robust to the prior dispersions, provided they stay within our recommended range.

We remark that if one were to change the prior dispersion arbitrarily then results would be affected, in a similar manner to how regularization parameters affect penalized likelihood results. However, in our view the prior beliefs implied by arbitrary prior dispersions would be unreasonable in most applications. We also note that there is a wide objective Bayes literature on using the data to set the prior parameters, see [6] for an excellent review. We do not argue against such strategies, but we focus on our defaults as a simple strategy that attains a fairly competitive performance in practice.

3. THEORY

This section describes the asymptotic solution returned by Bayesian model selection, when the observed data $(o_i, c_i, z_i) \sim F_0$ are independent realizations from some F_0 , where $z_i \in \mathbb{R}^{p(r+1)+q}$ for $q \geq 0$ contains the observed covariates $(x_i, s_i) \in \mathbb{R}^{p(r+1)}$, and potentially also q additional columns. These columns may contain covariates that were not recorded but are truly relevant for the outcome or the censoring, or non-linear effects and interactions missed by (x_i, s_i) . We do not assume F_0 to be parametric, rather it can be quite general, and the whole model structure assumed by the analyst (e.g. accelerated times, proportional hazards) may be wrong.

Section 3.1 shows that when one assumes the Normal AFT model (2.1) but truly $(o_i, c_i, z_i) \sim F_0$, the maximum likelihood estimator under each model γ converges to an optimal $(\alpha_\gamma^*, \kappa_\gamma^*, \tau_\gamma^*)$ and is asymptotically normally-distributed. See [16] and [17] for related asymptotic results, and Section S8 for analogous results for the Laplace AFT model. Section 3.2 shows that Bayesian model selection in the AFT model asymptotically returns the smallest γ^* such that all effects in $(\alpha_\gamma^*, \kappa_\gamma^*)$ are non-zero. Equivalently,

γ^* is defined by the zeroes in (α^*, κ^*) , the optimal value under the full model including all parameters. Section 3.3 gives analogous results for Cox models. These results are extended to probit models in Section S9, and in Section S10 to more general concave log-likelihood models. It is possible to derive similar results beyond the concave case, however this class encompasses all the models we consider here and allows simplifying the proofs and technical conditions.

Throughout we help interpret the solution and certain Bayes factors properties. Of particular relevance, Section 3.1 discusses that the asymptotic solution γ^* excludes covariates that do not help predict the outcome nor the censoring times, and offers some examples. Section 3.2 comments on potential advantages of using low-dimensional basis and non-linear decompositions to detect covariate effects.

3.1 Asymptotic solution in AFT models

As the sample size grows, Bayesian model selection recovers a model γ^* that excludes parameters that are asymptotically estimated to be zero. Under mild regularity conditions, this limiting parameter is the value maximizing the expected log-likelihood under F_0 . We start by defining the expected log-likelihood, then state the limiting result, and finally interpret its meaning and implications for model selection.

Let $\eta_\gamma = (\alpha_\gamma, \kappa_\gamma, \tau) \in \Gamma_\gamma$ be the **vector with $p_\gamma + rs_\gamma$ regression parameters under a given model γ (Section 2.2) plus the error variance, where $\Gamma_\gamma = \mathbb{R}^{p_\gamma + rs_\gamma} \times \mathbb{R}^+$ is the corresponding parameter space. Let**

$$\begin{aligned} m(\eta_\gamma) &= (1 - u_1) \left[\log \Phi \left(x_1^\top \alpha_\gamma + s_1^\top \kappa_\gamma - \tau \log(c_1) \right) \right] \\ &+ u_1 \left[\log(\tau) - \frac{1}{2} \log(2\pi) - \frac{1}{2} \left(\tau \log(o_1) - x_1^\top \alpha_\gamma - s_1^\top \kappa_\gamma \right)^2 \right], \end{aligned}$$

the contribution of one observation to the log-likelihood (2.2), and

$$\begin{aligned} M(\eta_\gamma) &= \mathbb{E}_{F_0}(m(\eta_\gamma)) = P_{F_0}(u_1 = 0) \mathbb{E}_{F_0} \left[\log \Phi \left(x_{1\gamma}^\top \alpha_\gamma + s_{1\gamma}^\top \kappa_\gamma - \tau \log(c_1) \right) \mid u_1 = 0 \right], \\ (3.1) \quad &+ P_{F_0}(u_1 = 1) \left(\log(\tau) - \frac{1}{2} \log(2\pi) - \frac{1}{2} \mathbb{E}_{F_0} \left[\left(\tau \log(o_1) - x_{1\gamma}^\top \alpha_\gamma - s_{1\gamma}^\top \kappa_\gamma \right)^2 \mid u_1 = 1 \right] \right) \end{aligned}$$

its expectation under the data-generating F_0 . Under minimal conditions, $M(\eta_\gamma)$ has a unique maximizer, denoted by $\eta_\gamma^* = (\alpha_\gamma^*, \kappa_\gamma^*, \tau_\gamma^*)$. Below we focus our interpretation on viewing (3.1) as the expectation of a likelihood-associated reward, and η_γ^* as the associated minimizer, but η_γ^* can also be viewed as minimizing the Kullback-Leibler divergence to $F_0(y, u)$ (also called generalized Kullback-Leibler divergence, see [16]).

Proposition 1 proves that the maximum likelihood estimator $\hat{\eta}_\gamma$ converges to η_γ^* , and Proposition 2 its asymptotic normality with a sandwich covariance that is standard in misspecified models, and corresponds to the smallest possible covariance for unbiased estimators under model misspecification. Such variance alteration does not affect consistency but can alter finite n false positives and asymptotic power (see Section 3.2). See also Propositions S1-S2 for analogous results on the AFT model with Laplace errors. Mild technical conditions, denoted A1-A5, that suffice for the proposition to hold are discussed in Section S7. **We remark that A3 assumes the existence and finiteness of η_γ^* and $\hat{\eta}_\gamma$ (the latter for large enough n), which implies that these optima cannot occur at the boundary of Γ_γ and must be unique (by concavity). For example, this rules out situations where η_γ^* contains infinite regression parameters or variance, or zero variance,**

which we view as pathological cases that we exclude from consideration. We thank an anonymous referee for pointing out an inconsistency in our original proof, and providing our current argument leading to A3.

PROPOSITION 1. *Assume A1-A3. Then, $\eta_\gamma^* = \operatorname{argmax}_{\Gamma_\gamma} M(\eta_\gamma)$ is unique and $\hat{\eta}_\gamma \xrightarrow{P} \eta_\gamma^*$ as $n \rightarrow \infty$.*

PROPOSITION 2. *Assume A1-A5. Then*

$$\sqrt{n}(\hat{\eta}_\gamma - \eta_\gamma^*) \xrightarrow{D} N\left(0, V_{\eta_\gamma^*}^{-1} \mathbb{E}_{F_0}[\nabla m(\eta_\gamma^*) \nabla m(\eta_\gamma^*)^\top] V_{\eta_\gamma^*}^{-1}\right),$$

where $V_{\eta_\gamma^*}$ is the Hessian matrix of $M(\eta_\gamma)$ evaluated at η_γ^* , and $m(\eta_\gamma^*) = \log p(y_1 \mid \eta_\gamma^*)$.

Proposition 1 has important implications for model selection. Let (α^*, κ^*) be the optimal parameter under the full model that includes all linear and non-linear terms. Asymptotically, one obtains the model γ^* of smallest dimension maximizing (3.1) (see Section 3.2), which is defined by zeroes in (α^*, κ^*) . Specifically, $\gamma_j^* = 0$ if both linear and non-linear coefficients (α_j^*, κ_j^*) are zero, $\gamma_j^* = 1$ if $\alpha_j^* \neq 0$ and $\kappa_j^* = 0$, and $\gamma_j^* = 2$ if $\kappa_j^* \neq 0$.

To interpret this asymptotic solution, we turn attention to (3.1). If a covariate does not contribute to improving neither of the two terms in (3.1), then its corresponding entry in (α^*, κ^*) is zero. The first term is the expected log-probability, as predicted by the model, that the individual is censored at the observed $\log(c_1)$ (conditional on being censored). Therefore, any covariate that helps the model predict more accurately the occurrence of censoring events contributes to this first term. The second term is the mean squared error in predicting the observed time $\log(o_1)$, conditional on the time being uncensored. Expression (3.1) is an average of these two components weighted by the true censoring probability $P_{F_0}(u_1 = 0)$, and averaged across covariate values under F_0 . Hence γ^* drops covariates that do not predict survival neither censoring times, but may include those that, even if truly unrelated to survival, help explain the censoring. This interpretation extends to working models other than the Normal AFT. For any other fully-specified model, the first term in (3.1) features the model log-predicted probability of censoring, and the second term the usual log-likelihood for uncensored data. For example, under a AFT model with Laplace errors the asymptotic solution is defined by the mean absolute error and the Laplace survival function (see Section S8).

We present some simple examples to illustrate our discussion.

EXAMPLE 1. *Suppose that under F_0 , truly $\log o_i \mid c_i \sim N(x_{i1} + \theta \log c_i, \sigma^2)$ and $\log c_i \sim N(x_{i2}, \sigma^2)$. The analyst adopts the model $\log o_i \sim N(\beta_1 x_{i1} + \beta_2 x_{i2}, 1/\tau^2)$, which, as discussed, assumes non-informative censoring. If $\theta = 0$, the censoring under F_0 is non-informative, and then $\alpha_2^* = \beta_2^* = 0$, hence x_{i2} is discarded asymptotically.*

However, if $\theta \neq 0$ then truly $\log o_i = x_{i1} + \theta x_{i2} + \epsilon_i$, where $\epsilon_i \sim N(0, (1 + \theta^2)\sigma^2)$. Plugging this expression into (3.1), it is easy to show that then $\alpha_2^ \neq 0$. That is, the presence of informative censoring causes x_{i2} to be asymptotically selected.*

EXAMPLE 2. *Suppose that there is a fixed administrative censoring at $\log c_i = a$ for all individuals (so it is truly non-informative under F_0), a single covariate $x_i \in \mathbb{R}$, and that the analyst adopts the model $\log o_i \sim N(\beta_1 + \beta_2 x_i, 1/\tau^2)$. Suppose that x_i truly has an effect on the outcome under F_0 , but that said effect only occurs at a time*

$b > a$. Then the effect cannot be detected from the observed data, since all individuals are censored at a . The issue is that the covariate has an effect that deviates from the assumed AFT structure. For example suppose that, under F_0 ,

$$\log o_i = z_i + \theta x_i I(z_i > b),$$

where $x_i \in \{0, 1\}$ indicates that individual i received a treatment, $z_i \sim N(0, 1)$ is the survival time for untreated individuals, and $\theta > 0$ quantifies the treatment effect.

Here the effect is only present among individuals that live longer than b and, since censoring occurs before b , for all uncensored individuals one observes $\log o_i = z_i$. Plugging this expression and $\log c_i = a$ into (3.1), and noting that the conditioning on u_1 can be removed from the expectations, one can show that $\alpha_2^* = \beta_2^* = 0$. This is an extreme example where one cannot detect an effect that strongly deviates from the assumed mean structure, even though the censoring is non-informative. One could conceive related examples where a covariate has a time-varying effect that is first positive and then negative, before administrative censoring occurs, so that the average effect is near-zero.

EXAMPLE 3. Suppose that a potentially informative censoring occurs early, so that $P_{F_0}(u_1 = 0) \approx 1$. Then (3.1) under the full model is approximately equal to

$$\mathbb{E}_{F_0} \left[\log \Phi \left(x_1^\top \alpha + s_1^\top \kappa - \tau \log(c_1) \right) \right].$$

As discussed, this term is the log-probability that the outcome occurs after the observed censoring time, as predicted by the Normal AFT model. Hence, (α^*, κ^*) are essentially chosen to predict censoring times. If the censoring is informative and depends on a set of covariates, then (α^*, κ^*) will in general assign non-zero coefficients to these covariates, which will be asymptotically selected. A similar argument can be made for late censoring where $P_{F_0}(u_1 = 1) \approx 1$, then (α^*, κ^*) is approximately the usual (population) least-squares solution. If the outcome depends on the censoring, which in turn depends on a set of covariates, then least-squares will assign a non-zero coefficient to the latter.

3.2 Bayes factor rates for misspecified AFT models

This section proves that the posterior probability of the optimal model γ^* converges to 1, under mild conditions. Recall that the posterior probability of γ^* is

$$\pi(\gamma^* | y) = \frac{p(y | \gamma^*) \pi(\gamma^*)}{\sum_{\gamma} p(y | \gamma) \pi(\gamma)} = \left(1 + \sum_{\gamma \neq \gamma^*} B_{\gamma, \gamma^*} \frac{\pi(\gamma)}{\pi(\gamma^*)} \right)^{-1}.$$

Proposition 3 gives the rate at which each B_{γ, γ^*} converges to 0 (in probability), when one assumes a potentially misspecified AFT model. Provided that each $B_{\gamma, \gamma^*} \pi(\gamma) / \pi(\gamma^*)$ converges to 0 (this follows immediately in the standard case where prior model probabilities are bounded, for example) it follows that $\pi(\gamma^* | y) \xrightarrow{P} 1$. This implies that the highest posterior probability model consistently selects γ^* , and that including covariates with marginal posterior probability $\pi(\gamma_j^* | y) > t$, for any fixed threshold t , also leads to consistent selection.

Proposition 3 clarifies the role of censoring and misspecification. The result is stated for Laplace approximations to Bayes factors, a computationally-convenient alternative

to obtaining exact marginal likelihoods, but in our setting both are asymptotically equivalent (Proposition S6). Specifically, we consider

$$(3.2) \quad B_{\gamma, \gamma^*} = \frac{\widehat{p}(y \mid \gamma)}{\widehat{p}(y \mid \gamma^*)},$$

where $\widehat{p}(y \mid \gamma)$ is obtained via a Laplace approximation:

$$\widehat{p}(y \mid \gamma) = \exp\{\ell(\tilde{\eta}_\gamma) + \log \pi(\tilde{\eta}_\gamma)\} (2\pi)^{d_\gamma/2} |H(\tilde{\eta}_\gamma) + \nabla^2 \log \pi(\tilde{\eta}_\gamma)|^{-1/2},$$

where $\tilde{\eta}_\gamma = \arg \max_{\eta_\gamma} \ell(\eta_\gamma) + \log \pi(\eta_\gamma)$ is the maximum a posteriori under prior $\pi(\eta_\gamma)$. See Section S4 for details on computing this approximation.

Proposition 3 treats separately overfitted models (containing γ^*) and non-overfitted models (not containing γ^*). Overfitted models contain all truly relevant plus a few spurious parameters, a situation where the challenge is to enforce sparsity. Non-overfitted models are missing some truly relevant parameters, there the challenge is also to have high power to detect the missing signal. By truly relevant we mean improving $M(\eta_\gamma^*)$, *i.e.* the prediction of either observed or censored times, see Section 3.1. Recall that $d_\gamma = \dim(\eta_\gamma) = p_\gamma + r s_\gamma + 1$. Intuitively the proof of Proposition 3 is based on establishing the asymptotic distribution of the likelihood-ratio test statistic $2[\ell(\tilde{\eta}_\gamma) - \ell(\tilde{\eta}_{\gamma^*})]$, which is bounded by central chi-squares in the overfitted case and non-central chi-squares in the non-overfitted case, and then finding an asymptotic approximation to the other quantities featuring in $\widehat{p}(y \mid \gamma)$.

PROPOSITION 3. *Let B_{γ, γ^*} be the Bayes factor in (3.2) under either π_L , π_M or π_E , where γ^* is the AFT model with smallest d_{γ^*} minimizing (3.1), and $\gamma \neq \gamma^*$ another AFT model. Assume that both γ^* and γ satisfy Conditions A1-A5. Suppose that (g_M, g_E, g_L, g_S) are non-decreasing in n .*

(i) *Overfitted models. If $\gamma^* \subset \gamma$, then*

$$\log B_{\gamma, \gamma^*} = \log(a_n) + \frac{r}{2}(s_{\gamma^*} - s_\gamma) \log(ng_S) + \mathcal{O}_p(1),$$

where $a_n = (ng_L)^{\frac{p_{\gamma^*} - p_\gamma}{2}}$ under π_L , $a_n = (ng_M)^{3(p_{\gamma^*} - p_\gamma)/2}$ under π_M , and $a_n = (ng_E e^{2g_E \sqrt{n}})^{(p_{\gamma^*} - p_\gamma)/2}$ under π_E .

(ii) *Non-overfitted models. If $\gamma^* \not\subset \gamma$, then*

$$\log(B_{\gamma, \gamma^*}) = -n[M(\eta_{\gamma^*}^*) - M(\eta_\gamma^*)] + \log(b_n) + \frac{r}{2}(s_{\gamma^*} - s_\gamma) \log(ng_S) + \mathcal{O}_p(1)$$

where $b_n = (ng_L)^{\frac{p_{\gamma^*} - p_\gamma}{2}}$ under π_L , $b_n = (ng_M^3)^{(p_{\gamma^*} - p_\gamma)/2}$ under π_M , and $b_n = (g_E n)^{p_{\gamma^*} - p_\gamma} e^{-g_E c}$ under π_E , for finite $c \in \mathbb{R}$.

By Proposition 3(i) the rates to discard overfitted models are unaffected by misspecification and censoring (but certain constants can affect finite n behaviour, see the proof). These sparsity rates are improved by non-local priors and by setting large prior dispersions (g_L, g_M, g_E, g_S) , extending previous results [24, 30, 39, 38] to misspecified survival models. By Proposition 3(ii) the rate to detect non-spurious effects is exponential in n with a coefficient $M(\eta_{\gamma^*}^*) - M(\eta_\gamma^*) > 0$ that measures the drop of predictive ability in γ relative to γ^* , and is hence affected by misspecification and censoring. Recall that

predictive ability can be understood as a weighted average of forecasting the outcome to occur after the censoring time (for censored individuals) and the actual outcome time (for uncensored individuals).

When one misspecifies the model family, $M(\eta_{\gamma^*}^*) - M(\eta_\gamma^*)$ is driven by the projection of F_0 onto the assumed family. Interpreting the geometry of such projections is beyond our scope, but intuitively projections usually reduce distances and hence make $M(\eta_{\gamma^*}^*) - M(\eta_\gamma^*)$ smaller than if one were to assume the correct model class. By Part (ii), this would decrease the power to detect non-zero effects in η_γ^* .

To facilitate interpretation suppose there is no censoring. Then simple algebra shows that $M(\eta_{\gamma^*}^*) - M(\eta_\gamma^*) = \mathbb{E}_{F_0} [\log(\tau_{\gamma^*}^*/\tau_\gamma^*)]$, which measures the difference in mean squared prediction errors from using model γ instead of the optimal γ^* (given by $1/(\tau_\gamma^*)^2$ and $1/(\tau_{\gamma^*}^*)^2$, respectively). For instance, omitting covariates increases $\tau_{\gamma^*}^*/\tau_\gamma^*$, causing an exponential drop in power, see our examples in Sections 5.1-5.2 for an illustration.

Proposition 3 also highlights trade-offs in modeling non-linear covariate effects. Including a truly active non-linear term is rewarded by an improved model fit $M(\eta_{\gamma^*}^*) - M(\eta_\gamma^*)$, but runs into an $r \log(ng_S)$ penalty. In contrast, including a linear effect leads to a smaller improvement in fit, but also incurs a smaller $\log(ng_S)$ penalty. Hence, decomposing effects into a linear and non-linear components can improve power.

A similar observation illustrates that for model selection purposes, the advantages of using fully non-parametric effects over a finite-dimensional basis may be small. Suppose one replaced the basis dimension r by a larger r^* maximizing $M(\eta_{\gamma^*}^*) - M(\eta_\gamma^*)$. For m -degree splines with equi-spaced knots and sufficiently smooth $M(\cdot)$ the improvement in $M(\eta_{\gamma^*}^*) - M(\eta_\gamma^*)$ associated to increasing r to r^* is at most of order $1/r^m$ [35]. For said increase to offset the complexity penalty it needs to hold that $r^{m+1}(r^* - r)/2$ is of a smaller order than $n/\log(ng_S)$. Hence by letting $r^{m+1}r^*$ grow sub-linearly with n could improve power relative to r . However for even moderate r and cubic splines ($m = 3$) the required $n > r^*r^4$ can be impractically large, *e.g.* see the examples in Section 5.1 with $r \in \{5, 10, 15\}$. Further, the computational cost of using a large r^* for each considered model γ is impractical when one wishes to consider many models.

In summary, using a small basis dimension r (*e.g.* $r = 5$, in our examples) within the non-linear effect decomposition in Section 2.1 may be practically preferable to a non-parametric basis where r grows with n , for the purpose of detecting the effect.

3.3 Bayes factor rates for misspecified additive Cox models

Our Bayes factor results under misspecified Cox models are similar to Section 3.2, but here the optimal model γ^* is defined by zeroes in the parameter $\eta^* = (\beta^*, \delta^*)$ maximizing the expected partial likelihood (2.3) under F_0 , see (S7.11) for its expression and some discussion. The interpretation of η^* is also analogous, though here (2.3) rewards predicting a higher risk for individuals who experienced the event (uncensored) than for other individuals at risk. An alternative interpretation is possible by noting that (2.3) can be approximated by a Poisson regression log-likelihood [26], where one models the mean number of uncensored events in infinitesimal intervals. Intuitively, any covariate that helps predict this mean, which depends on the distribution of the censoring and survival times, is asymptotically selected. Covariates that are unrelated both to survival and censoring are hence discarded.

We consider Bayes factors obtained by a Laplace approximation to the integrated

partial likelihood

$$(3.3) \quad p(y \mid \gamma) = \int \exp \{ \ell_p(\beta_\gamma, \delta_\gamma) \} \pi(\beta_\gamma, \delta_\gamma \mid \gamma) d\beta_\gamma d\delta_\gamma$$

these can be viewed as the integrated likelihood under a limiting non-informative non-parametric Gamma process prior on h_0 , see [21] and [31] for a discussion. We obtain Bayes factor rates analogous to Section 3.2, the proof builds upon [52] and [28] who proved that $\bar{\eta}_\gamma = (\bar{\beta}_\gamma, \bar{\delta}_\gamma)$ maximizing (2.3) are consistent and asymptotically normal under misspecification, under Conditions B1-B4 listed in Section S7.4.

PROPOSITION 4. *Let B_{γ, γ^*} be the Bayes factor based on (3.3) under π_L , π_M or π_E , γ^* the Cox model with smallest d_{γ^*} minimizing the expected log partial likelihood M_p in (S7.10), and $\gamma \neq \gamma^*$ another Cox model. Assume that (γ^*, γ) satisfy Conditions B1-B4, and that (g_M, g_E, g_L, g_S) are non-decreasing in n .*

(i) *Let a_n be as in Proposition 3. If $\gamma^* \subset \gamma$, then*

$$\log B_{\gamma, \gamma^*} = \log(a_n) + \frac{r}{2}(s_{\gamma^*} - s_\gamma) \log(n g_S) + \mathcal{O}_p(1),$$

(ii) *Let b_n be as in Proposition 3. If $\gamma^* \not\subset \gamma$, then*

$$\log(B_{\gamma, \gamma^*}) = -n[M_p(\eta_{\gamma^*}^*) - M_p(\eta_\gamma^*)] + \log(b_n) + \frac{r}{2}(s_{\gamma^*} - s_\gamma) \log(n g_S) + \mathcal{O}_p(1).$$

That is, the Bayes factors under an assumed Cox model have similar asymptotic behavior as under an assumed AFT model, hence the conclusions stated in Section 3.2 also apply to the Cox model.

4. COMPUTATION

The two main computational challenges are exploring the model space $\gamma \in \{0, 1, 2\}^p$, and approximating the integrated likelihood $p(y \mid \gamma)$ in (2.5) for each model. We first discuss relative advantages of the Normal AFT and Cox models for computing $p(y \mid \gamma)$, and how they relate to the amount of censored data in Section 4.1. We also discuss an approximation to the Normal log-distribution function derivatives that dramatically speeds up computation for the AFT and probit models. Section 4.2 discusses the model search, when one cannot enumerate all 3^p models.

4.1 Within-model calculations

When the log-likelihood is concave (or locally concave around η_γ^* , as in asymptotically Normal models), Laplace approximations to $p(y \mid \gamma)$ are one of the fastest and more accurate methods available. A practical limitation is that, when one wishes to consider many models or the sample size is large, solving the required optimization problems can still be cumbersome. This cost can be significantly ameliorated by combining convex optimization algorithms that use warm initializations, see Section S4. See also [37] for an approach based on approximate Laplace approximations that bypasses the optimization exercise altogether.

Within survival analysis, an advantage of exponential-family AFT models is admitting sufficient statistics for the uncensored part of the likelihood, e.g. $(y_o^\top y_o, X_o^\top y, X_o^\top X_o)$ for (2.2). These can be computed upfront in $n_o(1 + p + p(p + 1)/2)$ operations and re-used whenever a new model γ is considered at no extra cost, but for large p such pre-computation has significant cost and memory requirements. Since one typically visits

only a small subset of models, many elements in $X_o^\top X_o$ are never used and it would be wasteful to compute them all upfront. It is more convenient to compute the entries in $X_o^\top X_o$ when first required by any given γ and storing them for later use. Our software follows this strategy by using sparse matrices in the C++ `Armadillo` library [41].

Given these sufficient statistics the log-likelihood in (2.2) requires $\min\{nd_\gamma, (n_c + 1)d_\gamma + d_\gamma(d_\gamma + 1)/2\}$ operations, and each entry in its gradient and hessian require $n_c + 1$ further operations. In contrast the Cox model's partial likelihood has a minimum cost of $n_o d_\gamma + n_o(n_o - 1)/2$ operations when censored times precede all observed times ($\max c_i < \min o_i$), and a maximum cost $nd_\gamma + [n(n+1) - n_c(n_c - 1)]/2$ when observed times precede all censored times. That is, the AFT likelihood has a significantly lower cost than the Cox model when $n_c < n_o$ (moderate censoring) or $n > d_\gamma$ (sparse settings).

A caveat of the Normal AFT model, however, is requiring the extensive evaluation of the log-cumulative distribution $\log \Phi$ and its derivatives. Each likelihood evaluation requires n_c terms featuring Φ and, although these terms can be re-used when computing $r(z) = \phi(z)/\Phi(z)$ and $D(z) = r(-z)^2 - zr(-z)$ in the gradient and hessian, evaluating $\Phi(z)$ is costly. Briefly, the problem of approximating the inverse Mill's ratio $r(z)$ has been well-studied [11]. There are many algorithms to approximate $\Phi(z)$, but $r(z)$ is harder, *e.g.* Expression 26.2.16 in [1] (page 932) has maximum absolute error $< 7.5 \times 10^{-8}$ for $\Phi(z)$ but unbounded absolute error for $r(z)$ as $z \rightarrow -\infty$. By combining existing proposals we built a fast approximation that guarantees the small relative errors. One may combine the Taylor series and asymptotic expansions in [1] (page 932, Expressions 26.2.16 and 26.2.12) for $\Phi(z)$ with an optimized Laplace continued fraction in [27] (Expression (5.3)) for $r(z)$ as $z \rightarrow -\infty$. The resulting $\hat{r}(z)$ has maximum absolute and relative errors < 0.000185 and < 0.000102 respectively, and for $\hat{D}(z) = \hat{r}(-z)^2 - z\hat{r}(-z)$ they are < 0.000424 and < 0.000505 . See Section S5 for further details. As an empirical check, the posterior model probabilities obtained in Section 5.3 when replacing $(r(z), D(z))$ by $(\hat{r}(z), \hat{D}(z))$ remained identical to the third decimal place.

This approximation also facilitates evaluating the log-likelihood and derivatives for probit and other models involving $\log \Phi$, and may have some independent interest. Using this approximation and the warm initializations in Section S4 is practically meaningful, for the TGFB data (Section 5.3, 868 parameters) they reduced the cost of 1,000 Gibbs iterations from >4 hours to 38 seconds.

4.2 Model exploration

Recent advances in Markov Chain Monte Carlo provide model exploration strategies that perform fairly well in practice, see [57] for a tempering approach that is particularly helpful when there are multi-modalities in $p(\gamma | y)$, or [12] for adaptive methods that reduce the effort in exploring low posterior probability models. Further, as n grows and posterior probabilities concentrate on a single model, it is possible to prove quick convergence [55]. Intuitively, if $p(\gamma^* | y) \approx 1$ and the chain converges quickly, there is high probability that γ^* will be visited after a few iterations. Most iterations are spent on models with high $\pi(\gamma | y)$ which, from Proposition 3, are models with dimension close to d_{γ^*} . The main burden arises from obtaining $p(y | \gamma)$, which only needs to be computed the first time that γ is visited and can be stored for subsequent iterations. Hence, if d_{γ^*} is not too large (sparse data-generating truths) or $\pi(\gamma | y)$ is concentrated on relatively few models, the cost is manageable.

Here for simplicity we describe Algorithm 1, a Gibbs algorithm that builds upon

earlier proposals [24, 38], with the novelty that it adds a latent augmentation to enforce hierarchical restrictions (non-linear terms in S are only added if the corresponding linear term in X is in the model) in a computationally-efficient manner. The algorithm obtains B samples $\gamma^{(1)}, \dots, \gamma^{(B)}$ from $\pi(\gamma \mid y)$. It is not a naive Gibbs algorithm that sequentially samples p trinary indicators, *i.e.* sets $\gamma_j^{(b)} = k$ with probability $\pi(\gamma_j = k \mid y, \gamma_1, \dots, \gamma_{j-1}, \gamma_{j+1}, \dots, \gamma_p)$ for $k \in \{0, 1, 2\}$. Instead, it is more convenient to run an augmented-space Gibbs on $2p$ binary indicators. Specifically let $\tilde{\gamma}_j = \mathbf{I}(\gamma_j = 1)$ for $j = 1, \dots, p$ denote that covariate j only has a linear effect, and $\tilde{\gamma}_j = \mathbf{I}(\gamma_{j-p} = 2)$ for $j = p+1, \dots, 2p$ a non-linear effect. Algorithm 1 samples $\tilde{\gamma}_j$ individually but prevents $(\tilde{\gamma}_j, \tilde{\gamma}_{j+p}) = (0, 1)$, *i.e.* enforces that having a non-linear effect when $\beta_j = 0$ has zero posterior probability. The greedy initialization of $\tilde{\gamma}^{(0)}$ is analogous to that in [24] and to the heuristic optimization in [33].

We remark that Algorithm 1 may suffer from worse mixing than naive Gibbs sampling of $\gamma_j \in \{0, 1, 2\}$, but is advantageous in sparse settings. If covariate j has a small posterior probability $\pi(\gamma_j \neq 0 \mid y)$ then $\pi(\tilde{\gamma}_j = 1 \mid y)$ is small and in most iterations $\tilde{\gamma}_{j+p}$ is set to zero without the need to perform any calculation. In contrast when sampling $\gamma_j \in \{0, 1, 2\}$ one must obtain the integrated likelihood for $\gamma_j = 2$, which can be costly due to adding the r extra parameters needed to capture the non-linear effect. As an example, in Section 5.3 sampling $\gamma_j \in \{0, 1, 2\}$ took over 5 times longer to run than Algorithm 1, but provided the same effective sample size up to 2 decimal places.

Algorithm 1 Augmented-space Gibbs sampling

- 1: Set $b = 0$, $\tilde{\gamma}^{(0)} = (0, \dots, 0)$.
- 2: For $j = 1, \dots, 2p$, update $\tilde{\gamma}_j^{(0)} = \arg \max_k p(\tilde{\gamma}_j = k \mid y, \tilde{\gamma}_{-j}^{(b)})$. If an update was made across $j = 1, \dots, 2p$ go back to Step 2, else set $\gamma_j^{(0)} = \max\{\gamma_j^{(0)}, \gamma_{j+p}^{(0)}\}$ for $j = 1, \dots, p$ and go to Step 3.
- 3: Set $b = b + 1$. For $j = 1, \dots, p$ set $\tilde{\gamma}_j^{(b)} = 1$ with probability

$$P(\tilde{\gamma}_j = 1 \mid y, \tilde{\gamma}_{-j}^{(b)}) = \begin{cases} 1, & \text{if } \tilde{\gamma}_{j+p} = 1, \\ \frac{p(y|\tilde{\gamma}_j=1, \tilde{\gamma}_{-j}^{(b)})p(\tilde{\gamma}_j=1, \tilde{\gamma}_{-j}^{(b)})}{p(y|\tilde{\gamma}_j=0, \tilde{\gamma}_{-j}^{(b)})p(\tilde{\gamma}_j=0, \tilde{\gamma}_{-j}^{(b)}) + p(y|\tilde{\gamma}_j=1, \tilde{\gamma}_{-j}^{(b)})p(\tilde{\gamma}_j=1, \tilde{\gamma}_{-j}^{(b)})}, & \text{if } \tilde{\gamma}_{j+p} = 0, \end{cases}$$

and otherwise set $\tilde{\gamma}_j^{(b)} = 0$.

- 4: For $j = p+1, \dots, 2p$ set $\tilde{\gamma}_j^{(b)} = 1$ with probability

$$P(\tilde{\gamma}_j = 1 \mid y, \tilde{\gamma}_{-j}^{(b)}) = \begin{cases} 0, & \text{if } \tilde{\gamma}_{j+p} = 0, \\ \frac{p(y|\tilde{\gamma}_j=1, \tilde{\gamma}_{-j}^{(b)})p(\tilde{\gamma}_j=1, \tilde{\gamma}_{-j}^{(b)})}{p(y|\tilde{\gamma}_j=0, \tilde{\gamma}_{-j}^{(b)})p(\tilde{\gamma}_j=0, \tilde{\gamma}_{-j}^{(b)}) + p(y|\tilde{\gamma}_j=1, \tilde{\gamma}_{-j}^{(b)})p(\tilde{\gamma}_j=1, \tilde{\gamma}_{-j}^{(b)})}, & \text{if } \tilde{\gamma}_{j+p} = 1, \end{cases}$$

and otherwise set $\tilde{\gamma}_j^{(b)} = 0$. If $b = B$ stop, else go back to Step 3.

5. EMPIRICAL RESULTS

We illustrate via examples the effect of censoring, misspecification and the use of non-linear effect decompositions on model selection. Section 5.1 considers a simple simulation study with $p = 2$ variables, which Section 5.2 extends to $p = 50$. We consider different data-generating truths where the covariates have a monotone or non-monotone effect, and where the truth follows an AFT, proportional hazards, or generalized hazards structure. In Section 5.3, we analyze the effect of gene TGFB on colon cancer. Given that the data-generating truth is unknown, in Section 5.4 we study the number of false




FIG 1. Simulation truth and a simulated dataset for Scenarios 1 (left) and 2 (right).

positives via a permutation exercise. See also Supplementary Section S11.3, where we analyze the effect of the estrogen receptor on breast cancer survival.

We consider five model selection methods combining the AFT and Cox models with local and non-local priors and with LASSO. For all Bayesian methods we took the highest posterior probability model $\hat{\gamma} = \arg \max \pi(\gamma | y)$ as the selected model. We refer to the first three methods as AFT-Zellner, AFT-pMOMZ and AFT-LASSO. They all assume an AFT model and use either the block-Zellner prior π_L , the non-local pMOM-Zellner prior π_M (Section 2.3), or LASSO penalties as proposed by [34]. AFT-Zellner and AFT-pMOMZ assume the Normal AFT model in (2.1), whereas AFT-LASSO uses a semi-parametric AFT model. The remaining two methods combine the Cox model with piMOM priors (Cox-piMOM, [31]) and LASSO (Cox-LASSO, [46]). For AFT-Zellner and AFT-pMOMZ we used the function `modelSelection` in the R package `mombf` with the default prior parameters, the Beta-Binomial prior $\pi(\gamma)$ in (2.6) and $B = 10,000$ iterations in Algorithm 1. For Cox-piMOM we used the function `cov.bvs` in the R package `BMSNLP` with default parameters and prior dispersion 0.25 as recommended by [31]. For AFT-LASSO and Cox-LASSO we used the functions `AEnet.aft` and `glmnet` in the R packages `AdapEnetClass` and `glmnet`, and we set the penalization parameter via 10-fold cross-validation.

5.1 Censoring, model complexity and misspecification with $p = 2$

We consider sample sizes $n \in \{100, 500\}$, as well as censored and uncensored data. We present results for AFT-pMOMZ, as those for AFT-Zellner and Cox-piMOM were largely analogous. These methods are compared to Cox-LASSO and AFT-LASSO in Section 5.2. We consider 6 simulation scenarios. Scenarios 1-2 have a data-generating AFT model, Scenarios 3-4 a generalized hazard model and Scenarios 5-6 a proportional hazards model. The first covariate has a linear effect in all scenarios, whereas the second covariate has a non-linear effect. In Scenarios 1, 3 and 5 this effect is strongly non-linear and non-monotone, whereas in Scenarios 2, 4 and 6 it is monotone and can be roughly approximated by a linear trend, see Figure 1.

SCENARIO 1. *AFT structure with $\log o_i = x_{i1} + 0.5 \log(|x_{i2}|) + \epsilon_i$ and $c_i = 0.5$, where $x_i \sim N(0, A)$, $A_{11} = A_{22} = 1$, $A_{12} = 0.5$, $\epsilon_i \sim N(0, \sigma = 0.5)$.*

SCENARIO 2. *AFT structure with $\log o_i = x_{i1} + 0.5 \log(1 + x_{i2}) + \epsilon_i$ and $c_i = 1$,*

where $x_i = (\tilde{x}_{i1}, |\tilde{x}_{i2}|)$, $\tilde{x}_i \sim N(0, A)$ and A, ϵ_i as in Scenario 1.

SCENARIO 3. *Generalized hazards structure with*

$$h_{GH}(t) = h_0(t \exp \{-x_{i1}/3 + 0.5 \log(|x_{i2}|)\}) \exp \{-x_{i1}/3 + 0.75 \log(|x_{i2}|)\},$$

$c_i = 0.5$, h_0 being the Log-Normal(0,0.5) baseline hazard and x_i as in Scenario 1.

SCENARIO 4. *Generalized hazards structure with*

$$h_{GH}(t) = h_0(t \exp \{-x_{i1}/3 + 0.5 \log(1 + x_{i2})\}) \exp \{-x_{i1}/3 + 0.75 \log(1 + x_{i2})\},$$

$c_i = 1$, and h_0 and x_i as in Scenario 3.

SCENARIO 5. *Proportional hazards with $h(t) = h_0(t) \exp \{3x_{i1}/4 - 5 \log(|x_{i2}|)/4\}$, $c_i = 0.55$, h_0 being the Log-Normal(0,0.5) baseline hazard and x_i as in Scenario 1.*

SCENARIO 6. *Proportional hazards with $h(t) = h_0(t) \exp \{3x_{i1}/4 - 5 \log(|x_{i2}|)/4\}$, $c_i = 0.95$, and h_0 and x_i as in Scenario 5.*

In all scenarios, we first consider that there is no censoring, and then a strong administrative censoring, giving censoring probabilities $P_{F_0}(u_i = 0) \approx 0.7$.

We first discuss Scenarios 1-2 and illustrate the advantage of using our non-linear effect decomposition. We first only considered the selection of non-linear effects, *i.e.* $\gamma_j \in \{0, 2\}$. In such case, the power to detect the effects (Figure S3, top) was significantly lower than when decomposing them into linear and non-linear parts (Figure S3, middle). These findings align with Proposition 3, in the sense that the improvement in model fit needs to overcome the penalty for using a non-linear basis. By considering $\gamma_j \in \{0, 1, 2\}$, one can capture part of the effect with a single linear term. Figure S3 also shows that censoring tends to reduce the power for both covariates.

Second, we illustrate the effect of the non-linear basis dimension r . We compared the earlier results, where r was part of the model selection, to those obtained under a single fixed $r = 5, 10$ or 15 (Figure S3, bottom). Interestingly, in Scenario 1 the best performance was observed for $r = 5$, despite the data-generating truth being strongly non-linear (Figure 1). In Scenario 2 the results were highly robust to r , as one might expect from the true effect being near-linear. That is, the smaller $r = 5$ gave a good compromise between inference and computation, we thus used $r = 5$ from now on.

The results for Scenarios 3-4 are in Figure S4, and for Scenarios 5-6 in Figure S5. The effect of censoring, model complexity and misspecifying covariate effects were largely analogous to Scenarios 1-2. To explore further the effects of misspecification, we repeated the simulations in Scenarios 1-2 but now setting F_0 to have asymmetric Laplace errors $\epsilon_i \sim \text{ALaplace}(0, s, a)$, where $a = -0.5$ is the asymmetry and s the scale in the parameterization of [38]. We set s such that the error variance was equal to the Normal simulations, that is $s = \sigma^2/[2(1 + a^2)] = 0.1$. Figure S6 shows the results. These are similar to Figure S3 except for a slight drop in the power to include active covariates.

Finally, we explored the effect of omitting covariates by analyzing the data from Scenarios 1-2 but considering that only x_{i1} was actually observed, *i.e.* removing x_{i2} from the analysis. Figure S7 shows the results. Relative to Figure S3, under Scenario 1 there was a reduction in the posterior evidence for including x_{i1} . Such reduction was not observed in Scenario 2, presumably due to x_{i1} being correlated with $\log(1 + x_{i2})$ and hence picking up part of its predictive power.

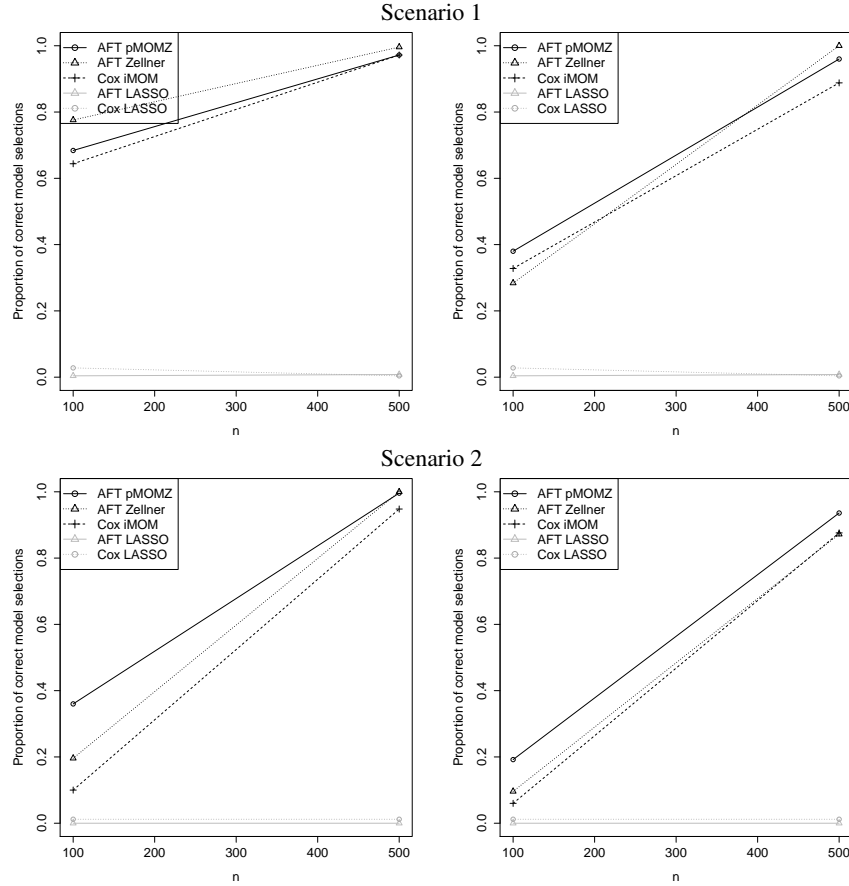


FIG 2. Scenarios 1-2, $p = 50$. Correct model selection proportion in uncensored (left) and censored (right) data.

5.2 Censoring, model complexity and misspecification with $p = 50$

We extended Scenarios 1-6 from Section 5.1 by adding 48 spurious covariates. We generated covariates $x_i \sim N(0, A)$ where A is a 50×50 matrix with unit diagonal and all off-diagonal $A_{ij} = 0.5$, and otherwise simulated data as in Section 5.1. Figure 2 shows the proportion of correct model selections by each model selection method in Scenarios 1-2, across 250 independent simulations. Figure S8 reports these results for Scenarios 3-4, and Figure S9 for Scenarios 5-6. Tables S1-S6 also display the posterior probability assigned to the optimal model γ^* and the average number of truly active and truly inactive selected covariates. All Bayesian methods exhibited a good ability to select γ^* that improved with larger n and uncensored data (as predicted by Proposition 3), and they all provided significant improvements over Cox-LASSO and AFT-LASSO, particularly in reducing the number of false positives. As expected AFT-Zellner and AFT-pMOM tended to slightly outperform Cox-piMOM under truly AFT data (Scenarios 1-2), and conversely under truly proportional hazards data (Scenarios 5-6), though the differences were relatively minor. Interestingly, under the generalized hazards model (Scenarios 3-4) again AFT-Zellner and AFT-pMOMZ achieved higher correct selection rates, presumably due to these generalized hazard settings being closer to an AFT than to an proportional hazards model.

5.3 Effect of TGFB and fibroblasts in colon cancer metastasis

[3] studied the effect of 172 genes related to fibroblasts (f-TBRS signature), a cell type producing the structural framework in animals, and a growth factor (TGFB) associated with lower colon cancer survival (time until recurrence). The authors obtained 172 genes responsive to TGFB in mice fibroblasts. They then used independent gene expression data from human patients, with tumor stages 1-3, to show that an overall high mean expression of these 172 genes was strongly associated with metastasis. We analyzed their data to provide a more detailed description of the role of TGFB and f-TBRS on survival. We used the $n = 260$ patients with available survival times, and used tumor stage (2 dummy indicators), TGFB and the 172 f-TBRS genes as covariates, for a total of $p = 175$. We first performed model selection via AFT-pMOMZ only for staging and TGFB. The top model had 0.976 posterior probability and included stage and a linear effect of TGFB, confirming that TGFB is associated with metastasis. The posterior marginal inclusion probability for a non-linear effect of TGFB was only 0.009. As an additional check, the maximum likelihood estimator under the top model gave P-values < 0.001 for stage and the linear TGFB effect. The estimated time accelerations associated to TGFB are substantial (Figure S12, left).

Next, we extended the exercise to all 175 variables, only considering linear effects. The top model contained gene FLT1 and the second top model genes ESM1 and GAS1, with respective posterior model probabilities 0.088 and 0.081. These were also the genes with highest inclusion probabilities (0.208, 0.699 and 0.567 respectively). There is plausible biology connecting FLT1, ESM1 and GAS1 to metastasis. From genecards.org [48], FLT1 is a growth and permeability factor in cell proliferation and cancer invasion. ESM1 is related to endothelium disorders, growth factor receptor binding and gastric cancer networks, and GAS1 plays a role in growth and tumor suppression. Interestingly the marginal inclusion probability for TGFB was only 0.107, that is after accounting for the top 3 genes TGFB did not show a significant effect on survival. For confirmation, we fitted via maximum likelihood the model with FLT1, ESM1, GAS1, stage and TGFB. The P-value for TGFB was 0.281 and its estimated effect was substantially reduced (Figure S12, right). Finally, we considered both linear and non-linear effects ($p(1 + r) = 1050$ columns in (X, S)). All non-linear effects had inclusion probabilities below 0.5 and the top 2 models contained FLT1, ESM1 and GAS1, as before. For comparison we run Cox-piMOM, AFT-LASSO and Cox-LASSO on the linear effects ($p = 175$). Stage and FLT1 were again selected by the top model under Cox-piMOM and by Cox-LASSO. Cox-LASSO selected 9 other genes, but only 4 had a significant P-value upon fitting a Cox model via maximum likelihood. Finally AFT-LASSO selected stage and 6 genes, two of which were also selected by Cox-LASSO. See Section S11.3 for a similar analysis of the estrogen receptor ESR1 effect on breast cancer.

Since this is a real-data application with an unknown ground truth, it is hard to assess which method performed best. As a first check, Table S7 reports the estimated predictive accuracy of each method via the leave-one-out cross-validated concordance index [14]. Cox-LASSO and AFT-pMOMZ achieved the highest concordance indexes, with the former selecting more variables than the latter on average across the cross-validation (13.6 vs. 3.9 for $p = 175$ and 11.7 vs. 4.9 for $p(1 + r) = 1050$). We remark that predictive accuracy is not our primary goal, but if a method were to miss truly active covariates then one would expect accuracy to decrease, hence it serves as a rough proxy for statistical power. To complete the exercise, we next evaluate false positive probabilities.

	Stage + TGFB ($p(r + 1) = 8$)		Stage + all genes ($p = 175$)	
	False positives	$\hat{\gamma} = 0$	False positives	$\hat{\gamma} = 0$
AFT-pMOMZ	0.0	100.0	0.0	100.0
Cox-piMOM	12.1	3.0	0.6	1.0
AFT-LASSO	35.9	31.0	2.2	45.0
Cox-LASSO	12.6	68.0	1.5	61.0

TABLE 1

Percentage of false positives and correct model selections ($\hat{\gamma} = 0$) in permuted colon cancer data (100 permutations) when the design had 8 columns (stage, linear and non-linear effect of TGFB) and 175 columns (stage and linear effect of 173 genes).

5.4 False positive assessment under colon cancer data

We did a permutation exercise to assess false positive findings in the colon cancer data. We randomly permuted the recurrence times, and left the covariates unpermuted. We obtained 100 independent permutations and recorded the model selected by each method. We first included only stage, a linear and non-linear term for TGFB as covariates, for a total of $p(r + 1) = 8$ columns in (X, S) . Next, we repeated the exercise considering linear effects for staging and the 173 genes, for a total of $p = 175$ columns.

The results are in Table 1 and Figure S10. AFT-pMOMZ achieved an excellent false positive control, it selected the null model in all permutations and assigned an average posterior probability $\pi(\gamma = 0 \mid y) = 0.846$ and 0.844 to the null model in the exercises with 8 and 175 columns (respectively). That is, AFT-pMOMZ not only selected the null model but also assigned a high confidence to that selection. All competing methods selected the null model significantly less frequently. They also showed inflated false positive percentages for the analysis with 8 columns, though interestingly these percentages were lower in the analysis with 175 columns. Figure S10 reveals an interesting pattern for Cox-piMOM, in $> 97\%$ of the permutations only 1 covariate was included. That is, although the mean false positives percentage for Cox-piMOM was similar to AFT-LASSO and Cox-LASSO, the selected model was always very close to the null model, as expected from the strong sparsity-inducing properties of non-local priors.

6. DISCUSSION

Our main contributions are describing a generic Bayesian model selection framework to incorporate non-linear effects in a data-driven fashion to balance power and sparsity and, perhaps more importantly, helping understand the interplay between censoring, misspecification and model complexity. In survival models, we showed that one asymptotically discards covariates that do not help predict the outcome neither censoring times (conditionally on other covariates), whereas in probit regression one keeps those that help reduce the probit loss function, and similarly for other concave log-likelihoods. We showed that censoring and misspecification can reduce power significantly. Understanding this phenomenon can be useful in the design of experiments, where one may increase the follow-up length to gain power. Enriching the model class, by considering semi- and non-parametric terms, to alleviate model misspecification requires some care as these additional terms can incur computational and statistical power losses. Our recommendation is to use Bayesian model selection to decide their inclusion in a data-adaptive manner, as in the proposed linear plus deviation from linearity decomposition. Although not discussed here for simplicity, one can also easily incorporate interactions between covariates into the proposed theory and computational methods.

From a technical point of view we used standard asymptotic arguments which, for concave log-likelihoods, lead to simpler proofs and technical conditions. It should be possible to extend our results, with some care, to non-concave and non-asymptotic settings (for example, using the high-dimensional framework in [32]), interval and left censored data, as well as to cure rate, recurrence or excess hazards models. We focused on fixed p to provide simpler results and intuition, under less restrictive technical conditions. While, in theory, it can be potentially interesting to allow the non-linear basis dimension r to grow with n , for actual methodology this often implies an impractical computational cost. This is critical in structural learning, where one wishes to consider many models. For this reason, in applied settings, it is common to use a finite basis.

Regarding high-dimensional settings, from recent results on misspecified penalized non-concave likelihood [29] Bayesian model selection [54, 36], we speculate that our main findings should remain valid. We remark, however, that high-dimensional formulations often incorporate stronger sparsity via the prior distribution, hence the power drop caused by censoring and misspecification could be more problematic than in our fixed p case.

We focused on model selection within additive models, but our results extend directly when one wishes to consider interactions, by adding the corresponding basis to our formulation. Our theory is valid for any given basis and also when performing selection on the basis itself, however, admittedly our examples focused on spline basis with fixed knots. We feel that a detailed study of basis selection would obscure the high-level intuition of our main results, but it represents an interesting aspect for future research.

ACKNOWLEDGMENTS

David Rossell was supported by Spanish Government grants RyC-2015-18544, Plan Estatal PGC2018-101643-B-I00, Europa Excelencia EUR2020-112096, Ayudas Fundación BBVA a Investigación en Big Data 2017, and NIH grant R01 CA158113-01.

REFERENCES

- [1] M. Abramowitz and I.A. Stegun. Handbook of mathematical functions with formulas, graphs, and mathematical table. In *US Department of Commerce*. National Bureau of Standards Applied Mathematics series 55, 1965.
- [2] J. Burrige. A note on maximum likelihood estimation for regression models using grouped data. *Journal of the Royal Statistical Society: Series B*, 43(1):41–45, 1981.
- [3] A. Calon, E. Espinet, S. Palomo-Ponce, D.V.F. Tauriello, M. Iglesias, M.V. Céspedes, M. Sevillano, C. Nadal, P. Jung, X.H.F. Zhang, D. Byrom, A. Riera, D. Rossell, R. Mangués, J. Massague, E. Sancho, and E. Batlle. Dependency of colorectal cancer on a tgf-beta-driven programme in stromal cells for metastasis initiation. *Cancer Cell*, 22(5):571–584, 2012.
- [4] I. Castillo, J. Schmidt-Hieber, and A.W. van der Vaart. Bayesian linear regression with sparse priors. *The Annals of Statistics*, 43(5):1986–2018, 2015.
- [5] Y.Q. Chen and N.P. Jewell. On a general class of semiparametric hazards regression models. *Biometrika*, 88(3):687–702, 2001.
- [6] Guido Consonni, Dimitris Fouskakis, Brunero Liseo, and Ioannis Ntzoufras. Prior distributions for objective Bayesian analysis. *Bayesian Analysis*, 13(2):627–679, 2018.
- [7] D.R. Cox. Regression models and life-tables. *Journal of the Royal Statistical Society. Series B*, 34(2):187–220, 1972.
- [8] A.G. Di Rienzo and S.W. Lagakos. Effects of model misspecification on tests of no randomized treatment effect arising from Cox’s proportional hazards model. *Journal of the Royal Statistical Society B*, 63(4):745–757, 2001.
- [9] D.B. Dunson and A.H. Herring. Bayesian model selection and averaging in additive and proportional hazards models. *Lifetime Data Analysis*, 11(2):213–232, 2005.
- [10] D. Faraggi and R. Simon. Bayesian variable selection method for censored survival data. *Biometrics*, 54(4):1475–1485, 1998.

- [11] A. Gasull and F. Utzet. Approximating Mills ratio. *Journal of Mathematical Analysis and Applications*, 420(2):1832–1853, 2014.
- [12] J.E. Griffin, Latuszynski, and M. F. J. Steel. In search of lost (mixing) time: Adaptive Markov chain monte carlo schemes for Bayesian variable selection with very large p . *Biometrika*, 108(1):53–69, 2020.
- [13] P.R. Hahn and C.M. Carvalho. Decoupling shrinkage and selection in Bayesian linear models: a posterior summary perspective. *Journal of the American Statistical Association*, 110(509):435–448, 2015.
- [14] F.E. Harrell Jr., K.L. Lee, and D.B. Mark. Multivariable prognostic models: issues in developing models, evaluating assumptions and adequacy, and measuring and reducing errors. *Statistics in Medicine*, 15(4):361–387, 1996.
- [15] S. Hattori. Testing the no-treatment effect based on a possibly misspecified accelerated failure time model. *Statistics & Probability Letters*, 82(2):371–377, 2012.
- [16] N.L. Hjort. On inference in parametric survival data models. *International Statistical Review / Revue Internationale de Statistique*, 60(3):355–387, 1992.
- [17] N.L. Hjort and D. Pollard. Asymptotics for minimisers of convex processes. *arXiv*, 1107.3806:1–24, 2011.
- [18] P. Hougaard. Frailty models for survival data. *Lifetime Data Analysis*, 1(3):255–273, 1995.
- [19] J. Huang, S. Ma, and H. Xie. Regularized estimation in the accelerated failure time model with high-dimensional covariates. *Biometrics*, 62(3):813–820, 2006.
- [20] J.L. Hutton and P.F. Monaghan. Choice of parametric accelerated life and proportional hazards models for survival data: asymptotic results. *Lifetime Data Analysis*, 8(4):375–393, 2002.
- [21] J.G. Ibrahim and M.H. Chen. *Bayesian Model Selection in Survival Analysis*. American Cancer Society, 2014.
- [22] J.G. Ibrahim, M.H. Chen, and S.N. MacEachern. Bayesian variable selection for proportional hazards models. *Canadian Journal of Statistics*, 27(4):701–717, 1999.
- [23] V.E. Johnson and D. Rossell. On the use of non-local prior densities in Bayesian hypothesis tests. *Journal of the Royal Statistical Society: Series B*, 72(2):143–170, 2010.
- [24] V.E. Johnson and D. Rossell. Bayesian model selection in high-dimensional settings. *Journal of the American Statistical Association*, 107(498):649–660, 2012.
- [25] N. Keiding, P.K. Andersen, and J.P. Klein. The role of frailty models and accelerated failure time models in describing heterogeneity due to omitted covariates. *Statistics in Medicine*, 16(2):215–224, 1997.
- [26] Nan Laird and Donald Olivier. Covariance analysis of censored survival data using log-linear analysis techniques. *Journal of the American Statistical Association*, 76(374):231–240, 1981.
- [27] C.I. C. Lee. On Laplace continued fraction for the normal integral. *Annals of the Institute of Statistical Mathematics*, 44(1):107–120, 1992.
- [28] D.Y. Lin and L.J. Wei. The robust inference for the Cox proportional hazards model. *Journal of the American statistical Association*, 84(408):1074–1078, 1989.
- [29] P.L. Loh. Statistical consistency and asymptotic normality for high-dimensional robust m -estimators. *The Annals of Statistics*, 45(2):866–896, 2017.
- [30] N.N. Narisetty and X. He. Bayesian variable selection with shrinking and diffusing priors. *The Annals of Statistics*, 42(2):789–817, 2014.
- [31] A. Nikooinnejad, W. Wang, and V.E. Johnson. Bayesian variable selection for survival data using inverse moment priors. *Annals of Applied Statistics*, 14:809–828, 2020.
- [32] M. Panov and V. Spokoiny. Finite sample Bernstein–von Mises theorem for semiparametric problems. *Bayesian Analysis*, 10(3):665–710, 2015.
- [33] N.G. Polson and L. Sun. Bayesian l0-regularized least squares. *Applied Stochastic Models in Business and Industry*, 35(3):717–731, 2018.
- [34] M.H. Rahaman-Khan and J.E.H. Shaw. Variable selection for accelerated lifetime models with synthesized estimation techniques. *Statistical Methods in Medical Research*, 28(3):937–952, 2019.
- [35] J. Ben Rosen. Minimum error bounds for multidimensional spline approximation. *Journal of Computer and System Sciences*, 5(4):430–452, 1971.
- [36] D. Rossell. A framework for posterior consistency in model selection. *Bayesian Analysis*, in press, 2021.
- [37] D. Rossell, O. Abril, and A. Bhattacharya. Approximate Laplace approximations for scalable model selection. *Journal of the Royal Statistical Society. Series B: Statistical Methodology*, 83(4):853–879, 2021.
- [38] D. Rossell and F.J. Rubio. Tractable Bayesian variable selection: beyond normality. *Journal of the*

- American Statistical Association*, 113:1742–1758, 2018.
- [39] D. Rossell and D. Telesca. Non-local priors for high-dimensional estimation. *Journal of the American Statistical Association*, 112(517):254–265, 2017.
 - [40] D. Rossell, D. Telesca, and V.E. Johnson. High-dimensional Bayesian classifiers using non-local priors. In *Statistical Models for Data Analysis XV*, pages 305–314. Springer, 2013.
 - [41] C. Sanderson and R. Curtin. Armadillo: a template-based C++ library for linear algebra. *Journal of Open Source Software*, 1(2):26, 2016.
 - [42] F. Scheipl, L. Fahrmeir, and T. Kneib. Spike-and-slab priors for function selection in structured additive regression models. *Journal of the American Statistical Association*, 107(500):1518–1532, 2012.
 - [43] J.G. Scott and J.O. Berger. Bayes and empirical-Bayes multiplicity adjustment in the variable-selection problem. *The Annals of Statistics*, 38(5):2587–2619, 2010.
 - [44] N. Sha, M.G. Tadesse, and M. Vannucci. Bayesian variable selection for the analysis of microarray data with censored outcomes. *Bioinformatics*, 22(18):2262–2268, 2006.
 - [45] M.J. Silvapulle and J. Burrridge. Existence of maximum likelihood estimates in regression models for grouped and ungrouped data. *Journal of the Royal Statistical Society: Series B*, 48(1):100–106, 1986.
 - [46] N. Simon, J. Friedman, T. Hastie, and R. Tibshirani. Regularization paths for Cox’s proportional hazards model via coordinate descent. *Journal of Statistical Software*, 39(5):1, 2011.
 - [47] P.J. Solomon. Effect of misspecification of regression models in the analysis of survival data. *Biometrika*, 71(2):291–298, 1984.
 - [48] G. Stelzer, N. Rosen, I. Plaschkes, S. Zimmerman, M. Twik, S. Fishilevich, T.I. Stein, R. Nudel, I. Lieder, Y. Mazor, S. Kaplan, D. Dahary, D. Warshawsky, Y. Guan-Golan, A. Kohn, N. Rappaport, M. Safran, and D. Lancet. The Genecards suite: from gene data mining to disease genome sequence analyses. *Current protocols in Bioinformatics*, 54(1):1–30, 2016.
 - [49] C.A. Struthers and J.D. Kalbfleisch. Misspecified proportional hazard models. *Biometrika*, 73(2):363–369, 1986.
 - [50] R. Tibshirani. The lasso method for variable selection in the Cox model. *Statistics in Medicine*, 16(4):385–395, 1997.
 - [51] X. Tong, L. Zhu, C. Leng, W. Leisenring, and L.L. Robison. A general semiparametric hazards regression model: efficient estimation and structure selection. *Statistics in Medicine*, 32(28):4980–4994, 2013.
 - [52] A.A. Tsiatis. A large sample study of Cox’s regression model. *The Annals of Statistics*, 9(1):93–108, 1981.
 - [53] S.N. Wood. *Generalized Additive Models: an introduction with R*. Chapman and Hall/CRC, New York, 2017.
 - [54] Y. Yang and D. Pati. Bayesian model selection consistency and oracle inequality with intractable marginal likelihood. *arXiv*, 1701.00311:1–38, 2017.
 - [55] Y. Yang, M.J. Wainwright, and M.I. Jordan. On the computational complexity of high-dimensional Bayesian variable selection. *The Annals of Statistics*, 44(6):2497–2532, 2016.
 - [56] Z. Ying. A large sample study of rank estimation for censored regression data. *The Annals of Statistics*, 21(1):76–99, 1993.
 - [57] G. Zanella and G. Roberts. Scalable importance tempering and Bayesian variable selection. *Journal of the Royal Statistical Society B*, 81(3):489–517, 2019.
 - [58] Z. Zhang, S. Sinha, T. Maiti, and E. Shipp. Bayesian variable selection in the accelerated failure time model with an application to the surveillance, epidemiology, and end results breast cancer data. *Statistical Methods in Medical Research*, 27(4):1–25, 2018.