Received: 15 November 2017 | Revised: 10 August 2018 | Accepted: 14 August 2018

DOI: 10.1002/bimj.201700259

RESEARCH PAPER

Biometrical Journal

Dynamic clinical prediction models for discrete time-to-event data with competing risks—A case study on the OUTCOMEREA database

Rachel Heyard[1] | Jean-François Timsit[2] | Wafa Ibn Essaied[2] | Leonhard Held[1] | on behalf of the COMBACTE-MAGNET consortium

[1]Department of Biostatistics at the Epidemiology, Biostatistics and Prevention Institute, University of Zurich, Hirschengraben 84, Zurich, Switzerland

[2]UMR 1137, IAME, University Paris-Diderot, Inserm, Paris, France

Correspondence
Rachel Heyard, Department of Biostatistics at the Epidemiology, Biostatistics and Prevention Institute, University of Zurich, Hirschengraben 84, 8001 Zurich, Switzerland.
Email: rachel.heyard@uzh.ch

Funding information
Innovative Medicines Initiative, Grant/Award Number: 115737-2

Abstract

The development of clinical prediction models requires the selection of suitable predictor variables. Techniques to perform objective Bayesian variable selection in the linear model are well developed and have been extended to the generalized linear model setting as well as to the Cox proportional hazards model. Here, we consider discrete time-to-event data with competing risks and propose methodology to develop a clinical prediction model for the daily risk of acquiring a ventilator-associated pneumonia (VAP) attributed to *P. aeruginosa* (PA) in intensive care units. The competing events for a PA VAP are extubation, death, and VAP due to other bacteria. Baseline variables are potentially important to predict the outcome at the start of ventilation, but may lose some of their predictive power after a certain time. Therefore, we use a landmark approach for dynamic Bayesian variable selection where the set of relevant predictors depends on the time already spent at risk. We finally determine the direct impact of a variable on each competing event through cause-specific variable selection.

KEYWORDS

cause-specific variable selection, Bayesian variable selection, competing events, discrete time-to-event model, dynamic prediction models, landmarking

1 | INTRODUCTION

Ventilator-associated pneumonia (VAP) is the leading nosocomial infection in critically ill ventilated patients treated in an intensive care unit (ICU) and *P. aeruginosa* (PA) is the most common organism associated with VAPs (Hunter, 2012). We are interested in the time until a first occurrence of a VAP attributed to PA. However, once ventilator-assisted patients are extubated or dead they are not at risk for a VAP PA anymore. Furthermore, the occurrence of a VAP PA and a VAP attributed to a different bacterium (VAP noPA) are mutually exclusive events. We are thus in the presence of competing risks. Data records on VAP acquisition in intensive care units (ICUs) are usually given on a daily basis. In such situations, discrete-time methods for the analysis of event history data (Allison, 1982; Efron, 1988; Singer & Willett, 1993) are a natural choice and equivalent to continuous time-methods (D'Agostino et al., 1990). Further these methods can be easily extended to competing events (Barnett & Graves, 2008; Barnett et al., 2009) for which the discrete cause-specific hazard function can be estimated using multinomial regression models with a time-dependent intercept (Tutz, 1995). In addition, lagged time-dependent variables are straightforward to incorporate; no additional reformatting of the data is necessary as opposed to handling such variables with continuous methods. In clinical prediction models, a treatment or a certain medical measurement may not have an immediate

 © 2018 WILEY-VCH Verlag GmbH & Co. KGaA, Weinheim

effect on the outcome on the same day but rather a delayed influence. It is then of interest to select the optimal lag for a time-dependent variable.

Developing good prediction models by selecting the most relevant predictors is of central importance in clinical research (Steyerberg, 2009). Objective Bayesian variable selection has been well developed for the normal linear model (Liang, Paulo, Molina, Clyde, & Berger, 2008), the generalised linear model (Held, Sabanés Bové, & Gravestock, 2015) and the Cox model (Held, Gravestock, & Sabanés Bové, 2016) for a single outcome. This methodology can easily be used to proceed to variable selection in discrete-time models with a single outcome by allowing the reference or null model to account for the time-dependency of the baseline hazard. However, these methods have not been developed for multiple outcomes. This article extends the methodology to discrete time-to-event models with competing risks selecting predictors that are relevant for all outcomes combined. However, Gustafson and Lefebvre (2008) argue that in many applications the subset of variables characterising a specific cause, or event, will change from cause to cause. They conclude that class-specific predictor selection is more "scientifically plausible" than ordinary variable selection and propose a rather complex method based on a hierarchically structured prior with a hyperparameter controlling the difference with ordinary predictor selection. We introduce a different, easily implemented, method for cause-specific variable selection that detects the variables that are of importance for a specific cause and sets the irrelevant coefficients to zero. Our approach is based on Bayes factors and closely related to standard Bayesian variable selection. Further, some variables may be more relevant for prediction during the early observation period whereas others may only be important after some time has been spent at risk. Hence we construct dynamic prediction models by applying the proposed variable selection methodology to landmark-specific data subsets. The landmarking approach has been introduced by Anderson, Cain, and Gelber (1983) and used in the context of Bayesian variable selection by Held et al. (2016).

This paper is structured as follows: Section 2 first introduces the OUTCOMEREA database as well as our research question that led to the development of the methods that will be discussed in this paper. Section 3 describes discrete time-to-event methods for one event of interest, generalises the approach to the setting of multiple outcomes and explains how the regression coefficients of these models can be estimated within the framework of multivariate generalized linear models. The procedure to develop a clinical prediction model for discrete time competing risks data is introduced in Section 4. As a first step, we concentrate on the selection of the baseline hazard in Section 4.1; then extend standard objective Bayesian variable selection methodology to our type of models in Section 4.2. We present a method for cause-specific variable selection in Section 4.3. Section 4.4 explains how a landmarking approach can be used to acquire dynamic prediction models. Section 4.5 illustrates how we can predict specific events of interest using the selected prediction model for the hazard. We apply this methodology in Section 5 and close with some discussion on the advantages and limitations of our method as well as possible ways to validate the prediction models in Section 6.

2 | THE OUTCOMEREA DATABASE

In this article, we work on the prospective multicenter observational database OUTCOMEREA; see Bekaert et al. (2011), Bouadma et al. (2015), and Truche et al. (2016) for more information on the database as well as for a detailed definition of VAP. Data collection started in January 1997 and is still ongoing. In our subset of the database we have information on patients admitted to 32 French ICUs until August 2015. We excluded patients who were never ventilated as well as patients under 18 or with missing birth date. In order to be at risk for a VAP, patients need to be ventilated for at least 48 hours; we excluded the ones who were not from the analysis. Total of 7,319 patients were retained. Daily ICU records are available, so that the discrete time units are days.

Extubation and death without VAP diagnosis and VAP due to a different pathogen than PA (VAP noPA) are considered to be competing events for VAP PA. See Figure 1 of the Supplementary Material for a schematic representation of the setting. Discharge of the ICU or the hospital does not figure on the list of competing events since a patient needs to be extubated prior to leaving the ICU. To keep things as simple as possible we concentrate on the time until a first event occurs. So once the patient is extubated or acquires a first VAP they are not modeled anymore. If a patient dies on the same day as a VAP is diagnosed we take VAP as the final event since the main focus of our application is VAP acquisition. The time scale will be time since start of ventilation, the observation period then starts on the first day of ventilation and ends with the occurrence of a first event.

In order to allow an intervention to take place we need to predict two days ahead such that patients at high risk of acquiring a VAP PA in two days can be treated with, for example, specific antibiotics, today. Therefore, we predict the outcome at day $t + 2$ using baseline as well as daily information. Furthermore, we decide to use time-dependent information lagged by up to two days ($t - 2$, $t - 1$, t). The choice to go back in time for at most two days allows us to use some data from the patient's ventilation history without losing too much information from the early stages. However, this prevents us to predict very early

TABLE 1 Total number of distinct events. The patients are only analyzed until the occurrence of a first event

	Dead	Extubated	VAP noPA	VAP PA
Number of events	896	3,251	635	341

TABLE 2 Potential predictors and their definition

	Baseline variables	
x_1	Patient type	Admission type (1 = surgical, 0 = medical)
x_2	Gender	Gender (1 = male, 0 = female)
x_3	Baseline SAPS2	Simplified Acute Physiology Score II at first day of admission (0–123)
x_4	Pneumonia at admission	Admitted with a pneumonia (yes or no)
x_5	Sepsis at admission	Sepsis at the admission to the ICU (yes or no)
x_6	Symptom	ICU admission motif (main symptom): factor with 5 levels
		1: Multiorgan failure—different shocks
		2: Acute respiratory distress syndrome—COPD exacerbation
		3: Acute renal failure
		4: Coma
		5: Continuous monitoring - Scheduled surgery—Trauma (= Reference category)
x_7	Diabetes	Diabetes (yes or no)
x_8	Comorbidity	At least one comorbidity (yes or no)
	Time-dependent variables	
x_{9t}	Hemodialysis	Usage of hemodialysis
x_{10t}	Any catheter	Presence of a catheter
x_{11t}	SOFA	Daily sequential organ failure assessment score (0–24)
x_{12t}	AB against PA	Use of at least one antibiotic (AB) against PA
		Aminoglycosides / Penems / Fosfomycin / Ceftazidime
		Fluoroquinolones / Ureido-carboxypenicillins / Cefpirome/cefepime
x_{13t}	DNR	Do not resuscitate
x_{14t}	Multiresistant PA	Colonization or infection with *P. aeruginosa* resistant to at least two molecules out of the 3: ticarcillin, ceftazidime, or imipenem

onset infections (before the fifth day of ventilation) and we simply exclude patients ventilated less than five days from the analysis. In the OUTCOMEREA data a total of 253 VAPs were diagnosed between day one and day four of ventilation among which 47 where caused by PA.

For the proposed method, the dataset needs to be presented in a long data format where each row contains a day of observation, here ventilation, of a specific patient. A patient is represented in the dataset by as many rows as they are ventilated. After applying all the inclusion criteria we retain a total number of 64,164 ventilation-days for 5,123 distinct ICU stays. On average it takes 12.5 ventilation days for an event to occur. Table 1 gives the number of distinct events. None of the observations are censored; all patients have had an event before end of August 2015. Table 2 lists the potential predictors for which data has been collected, if applicable on a daily basis, and which are judged clinically relevant. The SAPS2 and SOFA scores have been validated to predict hospital death in ICU patients (Le Gall et al., 1993; Timsit et al., 2002).

3 | DISCRETE TIME-TO-EVENT METHODS

3.1 | Modeling the discrete hazard function

Let the discrete random variable T represent the event time for one event of interest, then, for a given vector of explanatory variables $\mathbf{x}$, an individual's discrete hazard at time t is defined by

$$\lambda(t \mid \mathbf{x}) = \Pr(T = t \mid T \geq t, \mathbf{x}),\ t = 1, 2, \ldots \tag{1}$$

Time-varying and lagged covariates can simply be included in (1)

$$\lambda(t \mid \mathbf{x}_t) = \Pr(T = t \mid T \geq t, \mathbf{x}_t),\ t = 1, 2, \ldots \tag{2}$$

where $\mathbf{x}_t$ may comprise time-constant or (possibly lagged) time-dependent explanatory variables observed at time t. The discrete hazard function can directly be interpreted as the conditional probability that the event of interest is observed at t given the event has not been observed before t and given $\mathbf{x}_t$. The discrete survival function is given by

$$S(t \mid \mathbf{x}_t) = \Pr(T > t \mid \mathbf{x}_t) = \prod_{i=1}^{t}(1 - \lambda(i \mid \mathbf{x}_i)).$$

For a fixed t, the discrete hazard function models a binary response indicating whether an event took place exactly at time t or not:

$$\lambda(t \mid \mathbf{x}_t) = h(\beta_{0t} + \mathbf{x}_t^\top \boldsymbol{\beta}). \tag{3}$$

$h(\cdot)$ is a strictly monotonically increasing response function and $g = h^{-1}$ the link function. A complementary log-log (clog-log) link leads to the "grouped proportional hazards model" that can be seen as a discretized version of Cox's proportional hazards model (Tutz & Schmid, 2016, Chapter 3). β_{0t} is the time-dependent intercept and $\boldsymbol{\beta}$ is the vector of regression coefficients. The main extension compared to standard binary regression is that the intercept β_{0t} depends on time. Further information on discrete-time survival methods can be found in Singer and Willett (1993) and Tutz and Schmid (2016).

3.2 | Competing events models

In many applications, including the one presented in this paper, other events preclude the outcome of interest. Models dealing with this kind of problem are called competing risks models (Wolkewitz et al., 2014). Suppose $R \in \{1, \ldots, m\}$ denotes the distinct (terminating) causes and let T now be a discrete random variable representing the time until the occurrence of a first event, then the discrete cause-specific hazard function for cause r is

$$\begin{aligned}\lambda_r(t \mid \mathbf{x}_t) &= \Pr(T = t, R = r \mid T \geq t, \mathbf{x}_t) = \\ &= g^{-1}(\beta_{0tr} + \mathbf{x}_t^\top \boldsymbol{\beta}_r),\ t = 1, 2, \ldots\end{aligned} \tag{4}$$

where $g(\cdot)$ is a specific link function, β_{0tr} represents the cause-specific time-dependent intercept and $\boldsymbol{\beta}_r$ is the cause-specific vector of coefficients. The overall hazard function is defined as follows:

$$\lambda(t \mid \mathbf{x}_t) = \Pr(T = t \mid T \geq t, \mathbf{x}_t) = \sum_{r=1}^{m} \lambda_r(t \mid \mathbf{x}_t). \tag{5}$$

To deal with competing events, the most commonly used model is the multinomial logit model (Tutz, 1995). Using this model, the discrete cause-specific hazard function for a specific event r is defined by

$$\lambda_r(t \mid \mathbf{x}_t) = \frac{\exp(\beta_{0tr} + \mathbf{x}_t^\top \boldsymbol{\beta}_r)}{1 + \sum_{i=1}^{m} \exp(\beta_{0ti} + \mathbf{x}_t^\top \boldsymbol{\beta}_i)},\ \text{for } r = 1, \ldots, m,\ t = 1, 2, \ldots \tag{6}$$

where β_{0tr} can be interpreted as the cause-specific baseline hazard. Note that a reference category ($R = 0$) needs to be defined, meaning that there are actually $m + 1$ different events. In most applications, this refers to "staying at risk," with discrete hazard function

$$\lambda_0(t \mid \mathbf{x}_t) = \Pr(T > t \mid T \geq t, \mathbf{x}_t) = 1 - \sum_{r=1}^{m} \lambda_r(t \mid \mathbf{x}_t) = \frac{1}{1 + \sum_{i=1}^{m} \exp(\beta_{0ti} + \mathbf{x}_t^\top \boldsymbol{\beta}_i)}\ .$$

If p variables are considered, $\boldsymbol{\beta}_r^\top = (\beta_{r1}, \ldots, \beta_{rp})$, we have

$$\log\left(\frac{\lambda_r(t \mid \mathbf{x}_t)}{\lambda_0(t \mid \mathbf{x}_t)}\right) = \beta_{0tr} + \mathbf{x}_t^\top \boldsymbol{\beta}_r = \eta_{rt} \tag{7}$$

$$\frac{\lambda_r(t \mid \mathbf{x}_t)}{\lambda_0(t \mid \mathbf{x}_t)} = \exp(\beta_{0tr}) \cdot \exp(\beta_{r1})^{x_1} \ \ldots \ \exp(\beta_{rp})^{x_p}. \tag{8}$$

So, the linear predictor η_{rt} in (7) can be interpreted as the cause-specific log-odds. A one unit increase of the variable x_k has a multiplicative effect of size $\exp(\beta_{rk})$ on the cause-specific odds (8) of the outcome r.

3.3 | Maximum likelihood estimation under noninformative censoring

Time-to-event data are often subject to censoring; that is the event time T_i is unknown for some patients $i \in \{1, \ldots, n\}$. It is then only known that the event time T_i exceeds the censoring time C_i. In practice the observed event times $t_i = \min\{T_i, C_i\}$, $i = 1, \ldots, n$ are reported together with censoring indicators

$$\delta_i = \begin{cases} 1 & T_i \leq C_i \\ 0 & T_i > C_i \end{cases} \tag{9}$$

and the cause $r_i \in \{0, 1, \ldots, m\}$ of the event, here $r_i = 0$ if and only if $\delta_i = 0$.

The likelihood contribution of patient i is then (Möst, Pößnecker, & Tutz, 2016):

$$\mathcal{L}_i = \Pr(T_i = t_i, R_i = r_i)^{\delta_i} \Pr(T_i > t_i)^{1-\delta_i} \Pr(C_i \geq t_i)^{\delta_i} \Pr(C_i = t_i)^{1-\delta_i}. \tag{10}$$

Under the assumption of noninformative censoring, the censoring mechanism does not depend on the time-dependent covariate vector $\mathbf{x}_{is}$, $s = 1, \ldots, t_i$, and the factor $\Pr(C_i \geq t_i)^{\delta_i} \Pr(C_i = t_i)^{1-\delta_i}$ in (10) can be omitted:

$$\begin{aligned} \mathcal{L}_i &= \Pr(T_i = t_i, R_i = r_i \mid \mathbf{x}_{it_i})^{\delta_i} \Pr(T_i > t_i \mid \mathbf{x}_{it_i})^{1-\delta_i} = \\ &= \lambda_{r_i}(t_i \mid \mathbf{x}_{it_i})^{\delta_i} \left\{1 - \lambda(t_i \mid \mathbf{x}_{it_i})\right\}^{1-\delta_i} \prod_{s=1}^{t_i-1} \left\{1 - \lambda(s \mid \mathbf{x}_{is})\right\}. \end{aligned} \tag{11}$$

For each patient i we now define $\mathbf{y}_{is} = (y_{is0}, y_{is1}, \ldots, y_{ism}) = (1, 0, \ldots, 0)$ for all $s = 1, \ldots, t_i - 1$. If the event r_i is observed for observation i at time t_i (so $\delta_i = 1$), we define

$$\mathbf{y}_{it_i} = (y_{it_i0}, y_{it_i1}, \ldots, y_{it_im}) = (0, \ldots, 1, \ldots, 0),$$

with $y_{it_ir_i} = 1$ and all other elements of $\mathbf{y}_{it_i}$ set to zero. If the event r_i is censored at time t_i (so $r_i = 0$ and $\delta_i = 0$), we define

$$\mathbf{y}_{it_i} = (y_{it_i0}, y_{it_i1}, \ldots, y_{it_im}) = (1, 0, \ldots, 0).$$

Using these indicator variables y_{isr}, $s = 1, \ldots, t_i$, $r = 0, \ldots, m$, Equation (11) can be rewritten with (5) as a product of multinomial likelihood contributions,

$$\begin{aligned} \mathcal{L}_i &= \prod_{s=1}^{t_i} \left\{ \prod_{r=1}^{m} \lambda_r(s \mid \mathbf{x}_{is})^{y_{isr}} \right\} \left\{1 - \lambda(s \mid \mathbf{x}_{is})\right\}^{y_{is0}} = \\ &= \prod_{s=1}^{t_i} \left\{ \prod_{r=1}^{m} \lambda_r(s \mid \mathbf{x}_{is})^{y_{isr}} \right\} \left\{ 1 - \sum_{r=1}^{m} \lambda_r(s \mid \mathbf{x}_{is}) \right\}^{y_{is0}}, \end{aligned} \tag{12}$$

and the total log-likelihood is $\sum_{i=1}^{n} \log \mathcal{L}_i$. The number of multinomial contributions depends on the observed censoring pattern and the event times.

4 | DEVELOPMENT OF CLINICAL PREDICTION MODELS FOR DISCRETE COMPETING RISKS DATA

We are considering discrete time-to-event models and need to account for the time-dependency of the cause-specific baseline hazard β_{0tr} in Equation (4). A method to determine which time-specification best fits the data is discussed in Section 4.1. Then we

aim to select relevant predictors for our model using objective Bayesian variable selection as presented in Section 4.2. Choosing among baseline and time-dependent variables with different lags can result in a vast model space that needs to be considered and drastically slows down the computation of crucial quantities. Therefore we decide to use a sequential approach by first electing the best lags for the different time-dependent variables and then adding potential baseline predictors in the selection process. To further improve the model we proceed to cause-specific variable selection (CSVS), explained in Section 4.3, by dropping the outcome-specific coefficients that are not meaningful and correcting the relevant ones. In a last step, we repeat the variable selection and the CSVS steps on landmark-specific datasets that are conditioned on the fact that patients are still at risk at the landmark. Finally, we obtain dynamic prediction models for the timing of a VAP PA, see Section 4.4. The proposed methodology is implemented in `R` in the package `TBFmultinomial` available from `R`-CRAN. A vignette is provided with the package explaining its use. The `multinom()` function from the `nnet` package as well as the `vglm()` function from the `VGAM` package can be used to fit multinomial regression models with time-dependent intercepts resulting in discrete-time competing risks models.

4.1 | Baseline hazard selection

A major difference between standard multinomial regression models and discrete time-to-event competing risks models is the time-dependency of the intercept β_{0tr} in Equation (4). Let $t \in \{1, \dots, q\}$, then, for a fixed cause r, the baseline hazard can directly be represented by the parameters $\beta_{01r}, \beta_{02r}, \dots, \beta_{0qr}$ attributing a different intercept for each t. So, if the number q of distinct discrete time points is large, the number of parameters will be large too. This is why Barnett et al. (2009) and Beyersmann, Allignol, and Schumacher (2012, Chapter 7.3) suggest to include a variable for time as a factor for the first z time points, with $z \leq q$ based on clinical knowledge and on the size of the dataset. Even though this approach is very easy to apply and to interpret we still end up with a rather large number of unstable parameters, especially if a low number of events is observed at certain points in time. Another method to account for the time-dependency of the baseline hazard would be to assume that the baseline hazard is a smooth function of time (Tutz & Schmid, 2016, Chapter 5). To do so (Tutz & Schmid, 2016) use basis functions like polynomial splines or B-splines. Alternatively, the hazard can also be specified as a linear function of time. In order to find out which baseline hazard model fits our data best, we graphically compare the discrete cause-specific cumulative baseline hazard $A_{0tr} = \sum_{s=1}^{t} \beta_{0sr}, t = 1, \dots, q$ fitted using different time-specifications to the cause-specific Nelson–Aalen (N-A) estimator:

$$\hat{A}_{0tr} = \sum_{s=1}^{t} \frac{\text{number of observed type } r \text{ events at } s}{\text{number of individuals at risk just prior to } s}, t = 1, \dots, q\,,$$

here t are the distinct cause-specific event times. Then we select the time-specification whose cumulative baseline hazard is closest to the Nelson–Aalen estimator. The N-A estimator has originally been defined for continuous time-to-event models. Its definition is however directly translatable to the discrete setting.

4.2 | Objective Bayesian variable selection for discrete competing risks models

Let us consider a specific discrete time-to-event competing risks model $\mathcal{M}_j$ with linear predictor $\eta_{jtr} = \beta_{0jtr} + \mathbf{x}_t^T \boldsymbol{\beta}_{jr}$, with $r = 1, \dots, m$, as described in Section 3.2. Suppose $\boldsymbol{\theta}_{jr}$ regroups the cause-specific time-dependent intercept β_{0jtr} and the cause-specific regression coefficients vector $\boldsymbol{\beta}_{jr}$ of $\mathcal{M}_j$. An additional class, $r = 0$, is added to represent the reference category "staying at risk" or in our application "staying ventilated."

Bayesian model or variable selection is based on the posterior model probability (PMP) of a particular model $\mathcal{M}_j$ with parameters $\boldsymbol{\theta}_{jr}$:

$$\Pr(\mathcal{M}_j \mid \text{data}) = \frac{\text{p}(\text{data} \mid \mathcal{M}_j)\Pr(\mathcal{M}_j)}{\sum_{\mathcal{M}_k \in \Delta} \text{p}(\text{data} \mid \mathcal{M}_k)\Pr(\mathcal{M}_k)}, \tag{13}$$

where $\text{p}(\text{data} \mid \mathcal{M}_j) = \int \text{p}(\text{data} \mid \boldsymbol{\theta}_{jr}, \mathcal{M}_j)\, \text{p}(\boldsymbol{\theta}_{jr} \mid \mathcal{M}_j)\, \text{d}\boldsymbol{\theta}_{jr}$ is the marginal likelihood of $\mathcal{M}_j$, $\Pr(\mathcal{M}_j)$ its prior probability, Δ is the model space containing all considered candidate models and "data" regroups the multinomial outcome variable as well as the potential censoring indicator. As discussed in Section 3.3, the likelihood of a discrete competing risks model (with censoring) is equivalent to the likelihood of a multinomial response model (see also Chapter 3 of Tutz and Schmid (2016) for single outcome models). If an individual is censored they simply stay in the reference category ($r = 0$). Using this approach, the censoring mechanism is accounted for in the multinomial observations $\mathbf{y}_{it}$ in Equation (12) and so "data" in Equation (13) corresponds to

these multinomial observations with $m+1$ levels. If the selection is made among p variables Δ contains 2^p models. $\mathcal{M}_0$ will be the reference model; the simplest model without any of the p potential predictors, but containing the time-dependent intercept selected by the procedure described in Section 4.1. Most often $\mathcal{M}_0$ is the null model; in some situations however, we might want to be able to specify variables that should be included by default. $\Pr(\mathcal{M}_j)$ can be a uniform prior, meaning that each candidate model has the same probability a priori. In all that follows we however use a multiplicity-corrected model prior, where the probability depends on the complexity of the model (Scott & Berger, 2010). If p_j is the number of additional variables included in $\mathcal{M}_j$ (not yet present in $\mathcal{M}_0$), the multiplicity-corrected prior is defined as $\Pr(\mathcal{M}_j) = \frac{1}{p+1}\binom{p}{p_j}^{-1}$.

Once, the PMPs for all models in Δ are calculated, we compute the posterior inclusion probability (PIP) of the p variables $x_k, k \in \{1, \dots p\}$:

$$\Pr(x_k \in \mathcal{M} \mid \text{data}) = \sum_{\mathcal{M}_j \in \Delta} \Pr(\mathcal{M}_j \mid \text{data}) \mathbb{1}_{[x_k \in \mathcal{M}_j]}. \tag{14}$$

The PIPs are used to decide upon the inclusion of a variable as a predictor. Including only the variables with PIP ≥ 0.5 leads to the median probability model (MPM) proven to be optimal for prediction in the linear model (Barbieri & Berger, 2004). A motivation for the 0.5 threshold is that the multiplicity-corrected model prior assumes a mean of 0.5 on the prior inclusion probabilities and so, a posterior inclusion probability higher (or equal) to 0.5 strengthens the evidence that the variable belongs in the final model. Heyard and Held (2018) discuss different inclusion thresholds to be used on the PIPs, but we restrict ourselves to the traditional MPM approach in this article; see the discussion for more details on this topic.

To calculate the quantities (13) and (14) we need to set a prior on the regression coefficients $\boldsymbol{\beta}_j = (\boldsymbol{\beta}_{j1}^\top, \dots, \boldsymbol{\beta}_{jm}^\top)^\top \in \mathbb{R}^{mp}$ of all candidate models, with $j \in \{1, \dots, 2^p\}$ and $\boldsymbol{\beta}_j$ referring to a vector of size $m \cdot p$ combining all cause-specific coefficients of $\mathcal{M}_j$. This is a tedious task without subjective prior information. Therefore we follow Held et al. (2015) and use a generalized g-prior for the regression coefficients $\boldsymbol{\beta}_j$ of a particular model $\mathcal{M}_j$,

$$\boldsymbol{\beta}_j \mid g, \mathcal{M}_j \sim \mathcal{N}\left(\mathbf{0}, g \cdot \mathcal{I}^{-1}_{\boldsymbol{\beta}_j, \boldsymbol{\beta}_j}\right), \tag{15}$$

where $\mathcal{I}_{\boldsymbol{\beta}_j, \boldsymbol{\beta}_j}$ is the Fisher information matrix of the vector $\boldsymbol{\beta}_j$. Due to the properties of the g-prior, the regression coefficients are shrunken toward their prior mean using the shrinkage factor $t = g/(g+1)$. To compute the PMP in (13), Held et al. (2015) use test-based Bayes factors (TBF) (Johnson, 2008; Hu & Johnson, 2009) based on the deviance statistic. Let z_j be the deviance statistic of model $\mathcal{M}_j$ with d_j degrees of freedom, then the TBF of model $\mathcal{M}_j$ against the reference model $\mathcal{M}_0$ can be written in closed form (Johnson, 2008):

$$\text{TBF}_{j,0} = \frac{\text{p}(z_j \mid \mathcal{M}_j)}{\text{p}(z_j \mid \mathcal{M}_0)} = (g+1)^{-d_j/2} \exp\left(\frac{g}{g+1}\frac{z_j}{2}\right). \tag{16}$$

Consequently, we can rewrite the posterior probability of $\mathcal{M}_j$ as

$$\Pr(\mathcal{M}_j \mid \text{data}) = \frac{\text{TBF}_{j,0} \Pr(\mathcal{M}_j)}{\sum_{\mathcal{M}_k \in \Delta} \text{TBF}_{k,0} \Pr(\mathcal{M}_j)}.$$

To estimate the factor g, we use an empirical Bayes (EB) approach maximizing Equation (16) with respect to g:

$$\hat{g}_{\text{EB}} = \max\{z_j/d_j - 1, 0\}.$$

Note that the estimated shrinkage factor for model $\mathcal{M}_j$ under the EB approach,

$$\hat{t}_{\text{EB}} = \frac{\hat{g}_{\text{EB}}}{\hat{g}_{\text{EB}} + 1} = \max\{1 - d_j/z_j, 0\},$$

is the same as the one proposed by Copas (1997) for optimal prediction. There are other methods to estimate g (Liang et al., 2008; Held et al., 2015), we restrict ourselves to the EB approach in this paper though.

To eventually estimate the coefficients $\boldsymbol{\beta}_j$, the maximum likelihood estimates (MLEs) $\hat{\boldsymbol{\beta}}_j^{\text{ML}}$ of model $\mathcal{M}_j$ need to be computed by maximizing the total log likelihood (see end of Section 3.3). Then, the MLEs as well as the standard errors are shrunken to obtain the shrunken vector of estimated coefficients (Held et al., 2015, Section 3.1):

$$\boldsymbol{\beta}_j \mid \text{data}, g, \mathcal{M}_j \sim \mathcal{N}\left(\hat{t}_{\text{EB}} \cdot \hat{\boldsymbol{\beta}}_j^{\text{ML}}, \hat{t}_{\text{EB}} \cdot \mathcal{I}^{-1}_{\boldsymbol{\beta}_j,\boldsymbol{\beta}_j}\right). \tag{17}$$

If we consider one particular model $\mathcal{M}_j$ with cause-specific regression coefficients $\boldsymbol{\beta}_j$, we can define the standardized shrunken estimate of a specific component β_i, say, of $\boldsymbol{\beta}_j$:

$$\frac{\hat{t}_{\text{EB}} \cdot \hat{\beta}_i^{\text{ML}}}{\sqrt{\hat{t}_{\text{EB}} \cdot \sigma_i^2}} = \sqrt{\hat{t}_{\text{EB}}} \cdot \frac{\hat{\beta}_i^{\text{ML}}}{\sigma_i},$$

where σ_i^2 is the diagonal element of $\mathcal{I}^{-1}_{\boldsymbol{\beta}_j,\boldsymbol{\beta}_j}$ corresponding to the coefficient β_i. The analysis by maximum likelihood additionally provides the deviance statistic with its degrees of freedom that is needed to compute the TBF and to estimate g by EB.

Suppose that for our prediction models, we select among $\bar{p}$ baseline as well as $\bar{\bar{p}}$ time-dependent variables. Further we need to choose among l possible lagged versions of the $\bar{\bar{p}}$ time-varying variables. This can cause a large model space ($2^{\bar{p}+l\cdot\bar{\bar{p}}}$) and make the computation of the inclusion probabilities very slow. Hence, we suggest a sequential approach where we first select the lag(s) for the time-dependent variables and then choose the baseline variables. The decision on the maximum lag length l should be based on clinical knowledge but also on the availability of the data. One option would be to include the whole history of the variable up to l time units in the past but this is not feasible in clinical prediction where we seek parsimonious models. Hence, specific lags that are relevant for prediction at time t should be selected. First we use the TBF-methodology presented above to compute the $l \cdot \bar{\bar{p}}$ inclusion probabilities for each of the lagged variables. The model space has dimension $2^{l\cdot\bar{\bar{p}}}$. The reference model is the prediction model including the time-dependent intercept selected by the procedure described in Section 4.1 as well as all possible baseline variables. We include the baseline variables by default to avoid that certain lagged variables are judged relevant in this first step and are then dropped from the final MPM because some kind of interaction exists with the baseline variables. The $p_{\text{lags}} \leq l \cdot \bar{\bar{p}}$ selected lags have PIP ≥ 0.5 since we aim for the MPM.

Afterwards, we include the $\bar{p}$ potential baseline variables in the selection process and compute the $\bar{p} + p_{\text{lags}}$ inclusion probabilities. The reference model is now the null or baseline model including only the time-dependent intercept. Again, we include the variables having PIP ≥ 0.5 to find the MPM. We recompute the PIPs of the p_{lags} lagged variables to make their relevance comparable to the one of the baseline variables. They are most likely included again but their inclusion probability changes since the considered model space is different.

4.3 | Cause-specific variable selection

The posterior inclusion probability in (14) represents the importance of a specific variable for the prediction of all $m + 1$ causes together. However, as already discussed in Gustafson and Lefebvre (2008), a certain variable might only be relevant for some events whereas there is no real effect on others. In this case, it might be scientifically more plausible to include this variable only for certain events, leading to the problem of cause-specific variable selection (CSVS).

First, the estimated cause-specific coefficients of our selected discrete time-to-event competing risks model $\mathcal{M}$ are retrieved: $\hat{\beta}_{kr}$, $k = 1, \ldots, p$ and $r = 1, \ldots, m$ with $\hat{\beta}_{kr} \sim \mathcal{N}(\beta_{kr}, \sigma^2)$, where $\hat{\beta}_{kr}$ is the ML estimate and σ^2 is the corresponding diagonal element of the inverse Fisher Information $\mathcal{I}^{-1}_{\boldsymbol{\beta},\boldsymbol{\beta}}$ in (15). Then, the Bayes factor (BF) for inclusion against no inclusion of each cause-specific coefficient is computed. Our null (H_0) and alternative (H_1) hypotheses are defined as follows:

$$H_0\colon \beta_{kr} = 0 \text{ and } H_1\colon \beta_{kr} \sim \mathcal{N}\left(0, g \cdot \sigma^2\right),$$

since we use the g-prior on the regression coefficient as described in Section 4.2. H_0 refers to the situation where the coefficient is set to zero whereas, under H_1, it is included. Under the null hypothesis, the marginal likelihood is normal with mean 0 and variance σ^2; under the alternative, it is normal with mean 0 and variance $\sigma^2(1 + g)$. So we can derive the BF of the alternative against the null for each coefficient:

$$\text{BF}_{10} = \frac{\text{p}(\hat{\beta}_{kr} \mid H_1)}{\text{p}(\hat{\beta}_{kr} \mid H_0)} = \frac{\frac{1}{\sigma\sqrt{(1+g)}}\,\varphi\left(\frac{\hat{\beta}_{kr}}{\sigma\sqrt{(1+g)}}\right)}{\frac{1}{\sigma}\,\varphi\left(\frac{\hat{\beta}_{kr}}{\sigma}\right)} = \frac{\alpha\,\varphi(\alpha t)}{\varphi(t)},$$

with $t = \hat{\beta}_{kr}/\sigma$ being the standardized coefficient, $\varphi(.)$ denotes the density of the standard normal distribution and $\alpha = 1/\sqrt{1+g}$. Replacing g with its EB estimate $\hat{g}_{\text{EB}}$ we obtain

$$\text{BF}_{10} = \alpha \cdot \exp\left(\frac{1}{2}(1-\alpha^2)\, t^2\right) = \frac{1}{\sqrt{1+\hat{g}_{\text{EB}}}} \exp\left(\frac{1}{2}\, \frac{\hat{g}_{\text{EB}}}{1+\hat{g}_{\text{EB}}} t^2\right).$$

The coefficients with $\text{BF}_{10} > 1$ are included ($\boldsymbol{\beta}_{\text{kept}}$), whereas the ones with $\text{BF}_{10} \leq 1$ are set to zero ($\boldsymbol{\beta}_{\text{dropped}}$). This approach is inspired by the median probability model, where we include a variable only if its posterior inclusion probability exceeds the prior inclusion probability of 0.5. For the CSVS we only include the coefficients if the marginal likelihood under the alternative H_1 is larger than the marginal likelihood under the null H_0.

However, the coefficients $\boldsymbol{\beta}_{\text{kept}}$ need to be corrected using the conditional posterior distribution given the other coefficients $\boldsymbol{\beta}_{\text{dropped}}$ are zero. Assuming normality of the posterior of the regression coefficients

$$\boldsymbol{\beta} \mid \text{data} = \begin{pmatrix} \boldsymbol{\beta}_{\text{kept}} \\ \boldsymbol{\beta}_{\text{dropped}} \end{pmatrix} \mid \text{data} \sim \mathcal{N}\left(\begin{pmatrix} \hat{\boldsymbol{\beta}}_{\text{kept}} \\ \hat{\boldsymbol{\beta}}_{\text{dropped}} \end{pmatrix}, \begin{pmatrix} \hat{\Sigma}_{11} \hat{\Sigma}_{12} \\ \hat{\Sigma}_{21} \hat{\Sigma}_{22} \end{pmatrix}\right),$$

we have

$$\begin{aligned} \boldsymbol{\beta}_{\text{kept}} \mid \text{data} &\sim \mathcal{N}\left(\hat{\boldsymbol{\beta}}_{\text{kept}}, \hat{\Sigma}_{11}\right) \text{ and} \\ \boldsymbol{\beta}_{\text{dropped}} \mid \text{data} &\sim \mathcal{N}\left(\hat{\boldsymbol{\beta}}_{\text{dropped}}, \hat{\Sigma}_{22}\right). \end{aligned}$$

We can correct the retained coefficients:

$$\boldsymbol{\beta}_{\text{kept}} \mid (\boldsymbol{\beta}_{\text{dropped}} = \mathbf{0}, \text{data}) \sim \mathcal{N}(\mathbf{m}, \mathbf{v})$$

with

$$\begin{aligned} \mathbf{m} &= \hat{\boldsymbol{\beta}}_{\text{kept}} - \hat{\Sigma}_{12}\, \hat{\Sigma}_{22}^{-1}\, \hat{\boldsymbol{\beta}}_{\text{dropped}} \\ \mathbf{v} &= \hat{\Sigma}_{11} - \hat{\Sigma}_{12}\, \hat{\Sigma}_{22}^{-1}\, \hat{\Sigma}_{21}. \end{aligned}$$

The CSVS method makes our prediction model sparser. However, if we want to predict the cause-specific hazard for a particular individual for example, we still need the information on the variables relevant for the other causes, see Equation (6). For the cause-specific odds in Equation (8), on the other hand, we only need the variables that are relevant for the specific cause. Using these odds, we can compute the risk of moving from state "staying at risk" to event r, using only the variables included for outcome r.

4.4 | Landmarking for dynamic variable selection

Baseline variables are potentially important to predict the outcome at the start of observation, but may lose some of their predictive power after a certain time. Time-dependent variables, on the other hand, are more likely to be relevant over the whole observation period. In order to investigate the importance of variables over time we apply a landmark approach (Anderson et al., 1983; van Houwelingen, 2007; van Houwelingen & Putter, 2012) and proceed to dynamic Bayesian variable selection resulting in dynamic prediction models. Landmarking enables us to predict an event conditional on being at risk until a fixed point in time, the landmark. The approach used in this paper is similar to the one presented in Held et al. (2016). We first define discrete landmarks, for example each day of observation. Then, for each landmark we extract the data from the landmark until the end of observation. This creates as many subsets as there are landmarks defined. In a next step, the posterior inclusion probabilities are calculated for all potential variables using each landmark-specific data subset. We calculate PIPs for each landmark and for each variable, which are conditional on the fact that an individual is still at risk at a specific landmark. Now, to predict the outcome of an individual at a certain time-point, we include only the variables with PIP ≥ 0.5, at that particular landmark. In this manner, dynamic median probability models are defined; depending on the time point considered for the prediction we may use different variables. In addition, we apply the cause-specific variable selection procedure introduced in the previous section on each dynamic MPM. Finally, dynamic cause-specific MPMs are constructed.

4.5 | Prediction

The probability that we are interested in the application is $\Pr(T = t+2, R = r \mid T \geq t, \mathbf{x}_t)$, the risk of an event r happening at $t+2$ given that the patient is still at risk at t, where $\mathbf{x}_t$ contains all the information used for the prediction that can be baseline variables as well as time-dependent and even lagged information. This is not exactly a hazard. In fact, with

$$\begin{aligned}\lambda_r(t+2 \mid \mathbf{x}_t) &= \Pr(T = t+2, R = r \mid T \geq t+2, \mathbf{x}_t) = \\ &= \Pr(T = t+2, R = r \mid T \geq t+2, T \geq t, \mathbf{x}_t) = \\ &= \frac{\Pr(T = t+2, R = r, T \geq t+2 \mid T \geq t, \mathbf{x}_t)}{\Pr(T \geq t+2 \mid T \geq t, \mathbf{x}_t)} = \\ &= \frac{\Pr(T = t+2, R = r \mid T \geq t, \mathbf{x}_t)}{1 - \Pr(T < t+2 \mid T \geq t, \mathbf{x}_t)},\end{aligned}$$

the conditional probability of interest is

$$\begin{aligned}\Pr(T = t+2, R = r \mid T \geq t, \mathbf{x}_t) &= \lambda_r(t+2 \mid \mathbf{x}_t) \cdot [1 - \Pr(T < t+2 \mid T \geq t, \mathbf{x}_t)] = \\ &= \lambda_r(t+2 \mid \mathbf{x}_t) \cdot [1 - \Pr(T = t \mid T \geq t, \mathbf{x}_t] + \\ &\quad - \Pr(T = t+1 \mid T \geq t, \mathbf{x}_t)].\end{aligned}$$

Since, we have also $\Pr(T = t+1 \mid T \geq t, \mathbf{x}_t) = \lambda(t+1 \mid \mathbf{x}_t) \cdot [1 - \Pr(T = t \mid T \geq t, \mathrm{x}_t)]$, where $\lambda(t \mid \mathbf{x}_t) = \sum_{r=1}^{m} \lambda_r(t \mid \mathbf{x}_t)$ is the overall hazard, we obtain

$$\begin{aligned}\Pr(T = t+2, R = r \mid T \geq t, \mathbf{x}_t) &= \lambda_r(t+2 \mid \mathbf{x}_t) \cdot [1 - \Pr(T = t \mid T \geq t, \mathbf{x}_t) + \\ &\quad - \lambda(t+1 \mid \mathbf{x}_t) \cdot (1 - \Pr(T = t \mid T \geq t, \mathbf{x}_t))] = \\ &= \lambda_r(t+2 \mid \mathbf{x}_t) \cdot [1 - \lambda(t \mid \mathbf{x}_t)] \cdot [1 - \lambda(t+1 \mid \mathbf{x}_t)]. \qquad (18)\end{aligned}$$

To compute the quantity in (18) one option would be to simply ignore $\mathbf{x}_t$ in $\lambda(t \mid \mathbf{x}_t)$ and $\lambda(t+1 \mid \mathbf{x}_t)$, and use the life-table estimator of the overall hazard $\lambda(s)$ with $s \in \{t, t+1\}$:

$$\hat{\lambda}(s) = \frac{\#\text{ of events of any cause at } s}{\#\text{ of events of patients still at risk just prior to } s}.$$

A second option is to also develop a prediction model for the cause-specific hazards $\lambda_r(t \mid \mathbf{x}_t)$ and $\lambda_r(t \mid \mathbf{x}_t)$ and then compute their overall hazards at t and $t+1$, respectively, or directly find a prediction model for the overall hazards. The simplest way would be to consider the variables that were selected for $\lambda_r(t+2 \mid \mathbf{x}_t)$.

5 | APPLICATION ON THE OUTCOMEREA DATA

We apply the proposed methodology to the OUTCOMEREA data for the prediction of the timing of a VAP PA in an intensive care unit for critically ill patients. Our time units are days, $t \in \{1, 2, \dots\}$. Here, the number of outcomes is $m = 4$ with levels {1: death, 2: extubation, 3: VAP noPA , 4: VAP PA} and $r = 0$ is the reference category, staying ventilated.

The probability which we are interested in predicting in this application is $\Pr(T = t+2, R = r \mid T \geq t, \mathbf{x}_{t,t-1,t-2,0})$, which is the risk of an event r happening at $t+2$ given that the patient is still at risk at t, with $\mathbf{x}_{t,t-1,t-2,0}$ being the information potentially used for the prediction: baseline as well as time-dependent lagged variables. As discussed in the previous Section 4.5 we will need to find an optimal prediction model for the hazards $\lambda_r(t+2 \mid \mathbf{x}_t)$, $\lambda(t \mid \mathbf{x}_t)$ and $\lambda(t+1 \mid \mathbf{x}_t)$. In the following sections, we will illustrate the model development algorithm for the hazard $\lambda_r(t+2 \mid \mathbf{x}_t)$ step by step. To find a prediction model for the overall hazards $\lambda(t \mid \mathbf{x}_t)$ and $\lambda(t+1 \mid \mathbf{x}_t)$ the methods described in Held et al. (2015) and implemented in the `glmBfp`-package can be used.

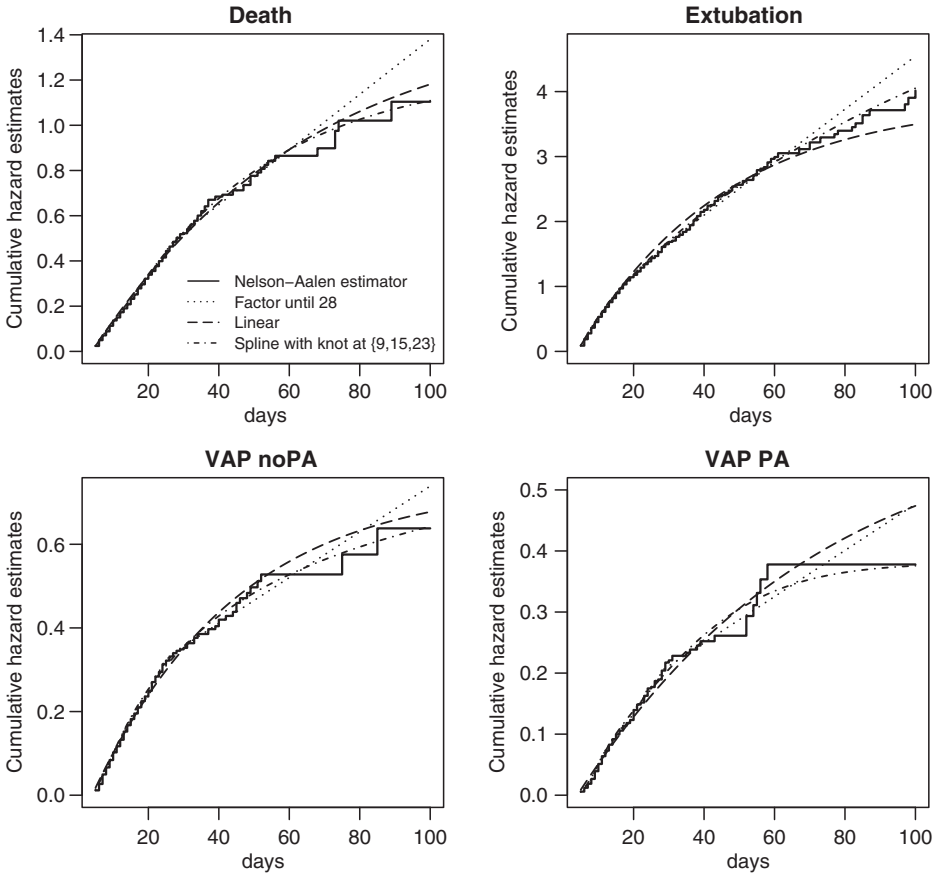

FIGURE 1 Cumulative baseline hazard estimated for the first 100 days under ventilation and for the four outcomes; depending on the definition of the time-dependency of the baseline hazard. They are compared to the Nelson–Aalen estimator of the cumulative hazard

5.1 | Baseline hazard selection

There are numerous ways to account for the time-dependency of the intercept, the baseline hazard β_{0tr}, in Equation (6). We will however restrict ourselves to the following options:

A: treat the day of ventilation as a linear variable: $\beta_{0tr} = \beta_{0r} + \beta_{1r}t$,

B: treat the day of ventilation as a factor variable until a certain time has passed, here, for example, 28 days: $\beta_{0tr} = \beta_{0r} + \beta_{1r}\mathbb{1}_{[t=1]} + \cdots + \beta_{28r}\mathbb{1}_{[t=28]} + \beta_{t>28\ r}\ \mathbb{1}_{[t>28]}$,

C: use natural cubic splines to model the time-dependency of the hazard: $\beta_{0tr} = \sum_{v=1}^{K} \beta_{0vr} N_v(t)$, where N_v are natural cubic spline basis functions and K is the total number of knots. We chose the knots at the 50%, 75%, and 90% quantile of the length of the observation period $\{9, 15, 23\}$.

Figure 1 shows the cumulative hazard estimated by the Nelson–Aalen estimator and using the above baseline hazard for each outcome individually for the first 100 days at risk. Approach B, does not mimic the Nelson-Aalen estimator well once the 28 first days passed. The cause-specific cumulative hazard is systematically overestimated in the later ventilation period. Approach B works sufficiently good for extubation, death, and VAP noPA, but not for our event of interest VAP PA. The spline approach C performs well for all outcomes, so that we choose to model the baseline hazard using a spline with three knots. Of course, there are numerous other ways to define the baseline hazard, for example, a different intercept for each possible time point. This is however not feasible in this application, because the model would be overfitted and the computation would be slowed down.

5.2 | Variable selection

Since we want to select among all baseline variables as well as time-dependent lagged variables in Table 2, the model space will be large; $2^{8+3\cdot6}$ models need to be fitted. As discussed before we consider to lag the time-dependent variables by one and two days but also consider their daily value (at the day of prediction) and opted for a sequential variable selection approach.

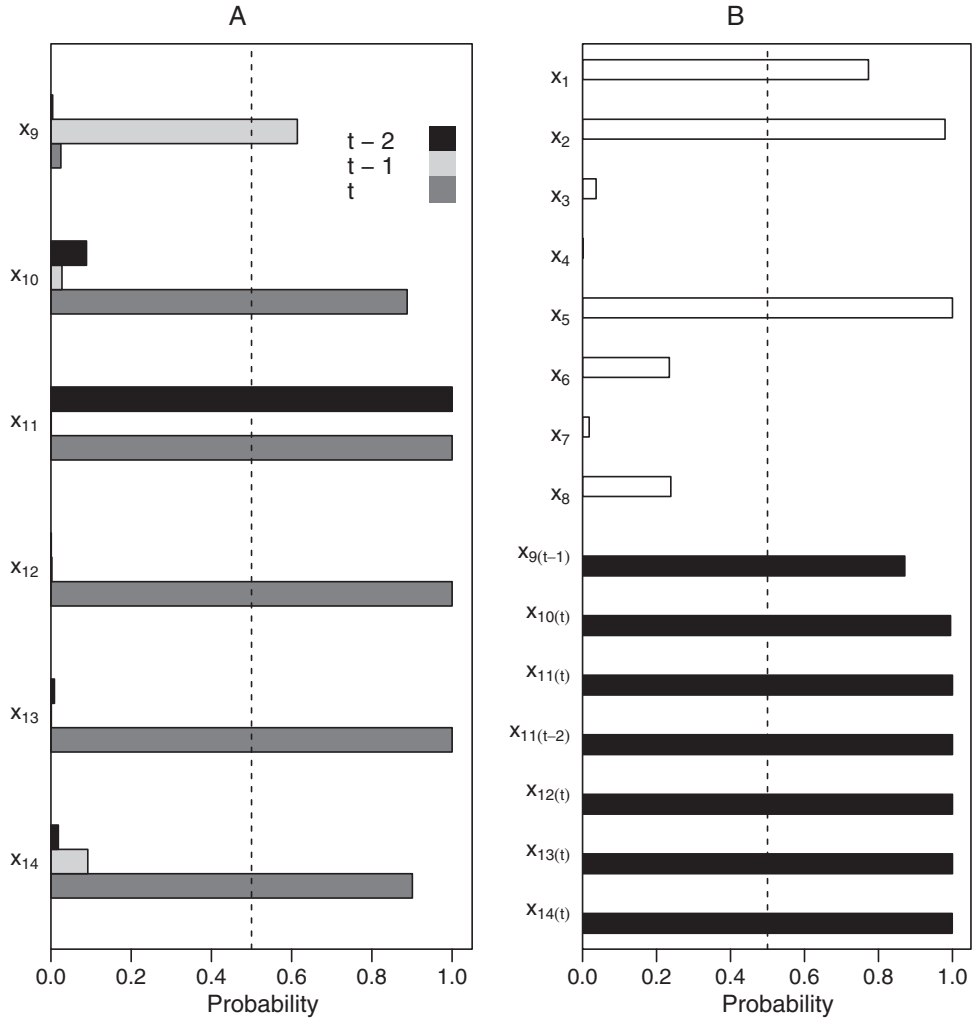

FIGURE 2 (A) The posterior inclusion probabilities for the different lags $(t-2, t-1, t)$ of the six time-dependent variables. (B) The posterior inclusion probabilities for the eight baseline variables (in white) and the seven previously selected lagged variables (in black). The inclusion threshold of 0.5 is used for the median probability model

First, we select the lag for all time-dependent variables simultaneously. The reference model includes all baseline variables in Table 2 and the time-dependent intercept selected before. Then the selection is among the six time-dependent variables with three different lags $(t-2, t-1, t)$ resulting in $3 \cdot 6 = 18$ inclusion probabilities (see Figure 2A, which shows the PIPs for the different lagged daily variables). Using the MPM inclusion threshold of 0.5, we select lags of zero and two days for x_{11}, a lag of one day for x_9 and versions without time lag for $x_{10}, x_{12}, x_{13,}$ and x_{14}. The fact that we select two specifications for the SOFA score (x_{11}) will be discussed later.

Next, we select the baseline variables. The reference model contains only the chosen baseline hazard. We recompute the PIPs of the time-dependent variables to be able to compare their relevance to the baseline variables' importance. Figure 2B shows the PIPs of the eight baseline variables (in white) as well as the PIPs of the seven time-dependent lagged variables (in black) that have been retained in the previous step. Our median probability model contains ten predictors. The seven lagged variables remain in the MPM; their PIPs even increased compared to the ones in Figure 2A. Of course, the sequential steps can be reversed. However, first choosing the lags seemed more natural here. The fact that the baseline variables were added in the reference model for the lag selection step ensures that the order of the sequential steps does not change the final results.

Figure 3 shows how the standardized shrunken coefficients evolve if we include one variable after the other in increasing order of their inclusion probabilities until we reach the MPM with ten predictors. Each plot represents one outcome. Interesting is the evolution of the coefficients of $x_{11(t)}$ and $x_{11(t-2)}$. For each cause, they evolve in opposite directions, meaning that if one of them has a positive effect the other one has a negative impact on the outcome. A possible reason could be that it is actually the trend in the evolution of the SOFA score over time that is important and the difference of $x_{11(t)}$ and $x_{11(t-2)}$ should be used

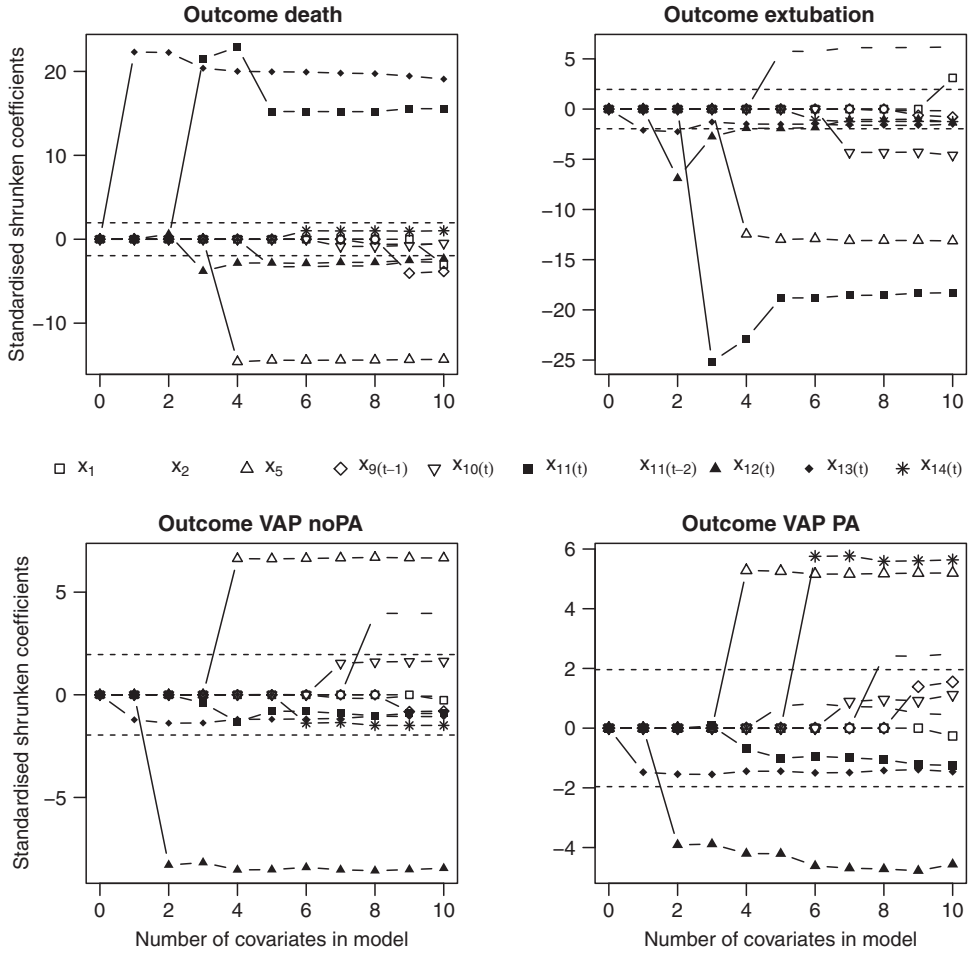

FIGURE 3 The paths of the shrunken standardized cause-specific coefficients depending on the number of variables in the model for each of the outcomes individually. The variables are added one by one in an order determined by their posterior inclusion probabilities. The most complex model here is the MPM including the ten most relevant variables. The dashed lines are at −1.96 and 1.96 to visually inspect significance of the coefficients at the conventional 5% level

to further simplify the model. However, this implies that we ignore the information of the initial value of the SOFA score. We decided that keeping both scores, at day t and with a two days lag, yields more information than their difference.

5.2.1 | Cause-specific variable selection

Since some variables may be, for example, relevant for the outcome death while they do not directly effect outcome VAP PA, we proceed to CSVS as described in Section 4.3. Figure 4 shows the absolute values of the standardized shrunken coefficients of the median probability model determined in the previous sections before and after CSVS. From the lower panel of the figure we can read off, for example, that the predictor x_1, the admission type, which was considered relevant for prediction with an inclusion probability of 0.77, is actually only important for the outcomes "dead" and "extubated" while our methodology suggests to set its cause-specific coefficients to zero for causes "VAP PA" and "VAP noPA." On the other hand, the gender of the patient, x_2, does not seem to have an important impact on the causes "dead" and "extubated" but changes the risk of acquiring a VAP caused by any pathogen. The variable x_5, sepsis at admission, is relevant for all outcomes. To be highlighted is the fact that the SOFA score x_{11} of day t is only relevant for the upper two outcomes while $x_{11(t-2)}$ influences all outcomes apart from "extubated." $x_{14(t)}$, which tells us whether a patient is infected by or carries resistant PA at the day of prediction, is only important for outcome "VAP PA." We refer to Figure 4 of the Supplementary Material for the same figure with printed values of the standardized coefficients. All of the remaining and corrected coefficients have standardized values larger than 1.96 that indicate their significance at the conventional 5% level. A variable that has been identified as a strong predictor for one outcome and not the others still influence the cause-specific hazards of the remaining outcomes, but not their log-odds (see Equation (8)). Moreover, knowing which variables have a direct effect on the different outcomes can permit more precise interventions.

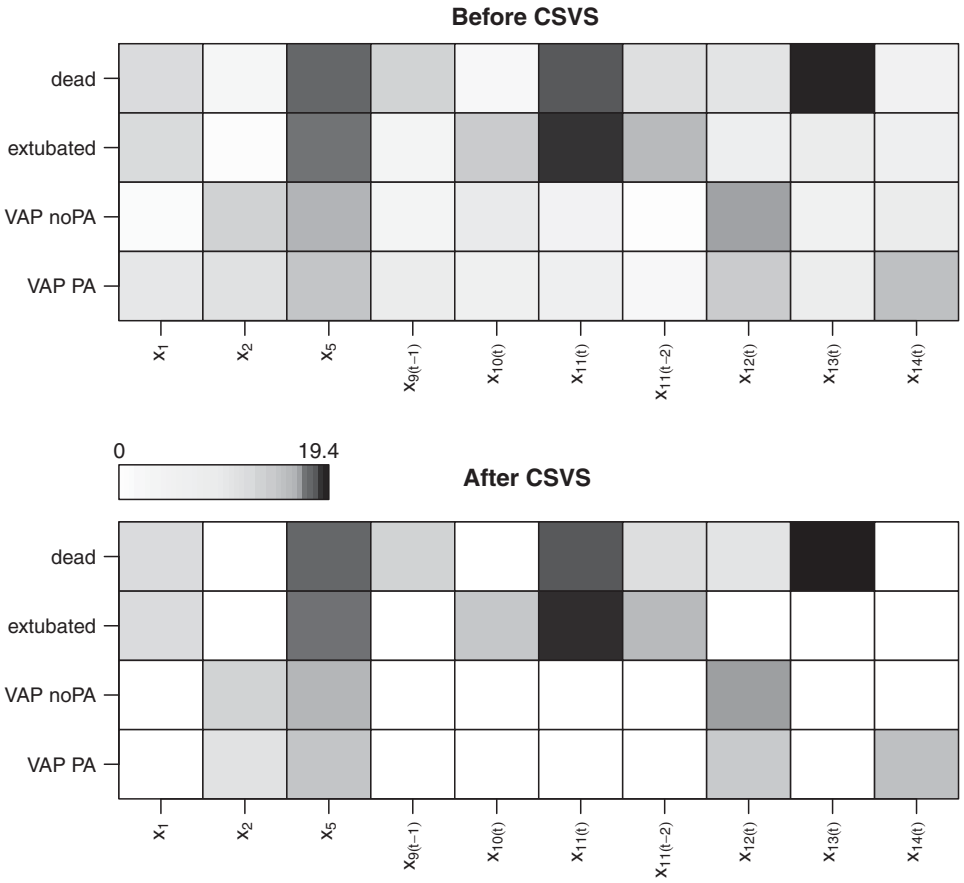

FIGURE 4 Absolute values of the shrunken standardized coefficients of the median probability model before and after cause-specific variable selection (CSVS)

5.2.2 | Dynamic cause-specific variable selection

In this last part, we reapply the proposed variable selection steps on landmark-specific data subsets as described in Section 4.4. The considered landmarks are the days of ventilation. Starting with day 5 as a first landmark, we drop past information and consider only data from landmark onward. Thus, the considered data sets are getting smaller as the landmark increases. We will only present the results up to landmark 14 to ensure their stability. Most of the events occur in the four first weeks since start of ventilation, see Figure 2 in the Supplementary Material. Considering only two weeks for the dynamic prediction models assures that we do not run into (quasi-) complete separation problems (Heinze & Schemper, 2002). We need however to be careful with the baseline hazard which we model as a natural cubic spline with knots at day nine, 15 and 23. This time-specification is practicable until landmark nine, afterwards it does not work anymore. Therefore we decide to still use a natural cubic spline with three knots as baseline hazard but leave the placement of the knots undefined.

Figure 5 shows the PIPs for the six time-dependent variables with their potential lags for each landmark. For the binary daily variable x_9, hemodialysis, we select a lag of one day for the first two landmarks as well as for landmark 11 and 12. At other points in time, the information whether hemodialysis was conducted on a patient is not relevant. No lag is used on x_{10}, presence of a catheter, at the beginning of the ventilation period, while it is better to use a lag of two days once ten days were spent ventilated. For x_{11}, the SOFA score, we use no lag and a two day lag for the two first landmarks while the two day lag is judged most important afterwards. The information at day of prediction t is assessed to be best for x_{12}, use of antibiotics, and x_{13}, DNR. The colonization or infection with a resistant PA, x_{14}, is an irrelevant variable after seven days of ventilation. At each landmark, we select the lags with PIP ≥ 0.5 in Figure 5 and add the baseline variables from Table 2 to the selection procedure. Figure 6 shows the PIPs of the baseline variables as well as the recalculated PIPs of the lagged variables as selected before, for each landmark. One can easily see that the baseline variables (blank circles) are either irrelevant right from the start of ventilation, like x_3, the baseline SAPS II score, or they lose their predictive power quite quickly (apart from x_5, sepsis at admission). The PIPs of the time-dependent lagged variables (filled circles) did not change much compared to the ones in Figure 5. For some of them the PIP is missing at particular landmarks because they were not computed due to an inclusion probability smaller than

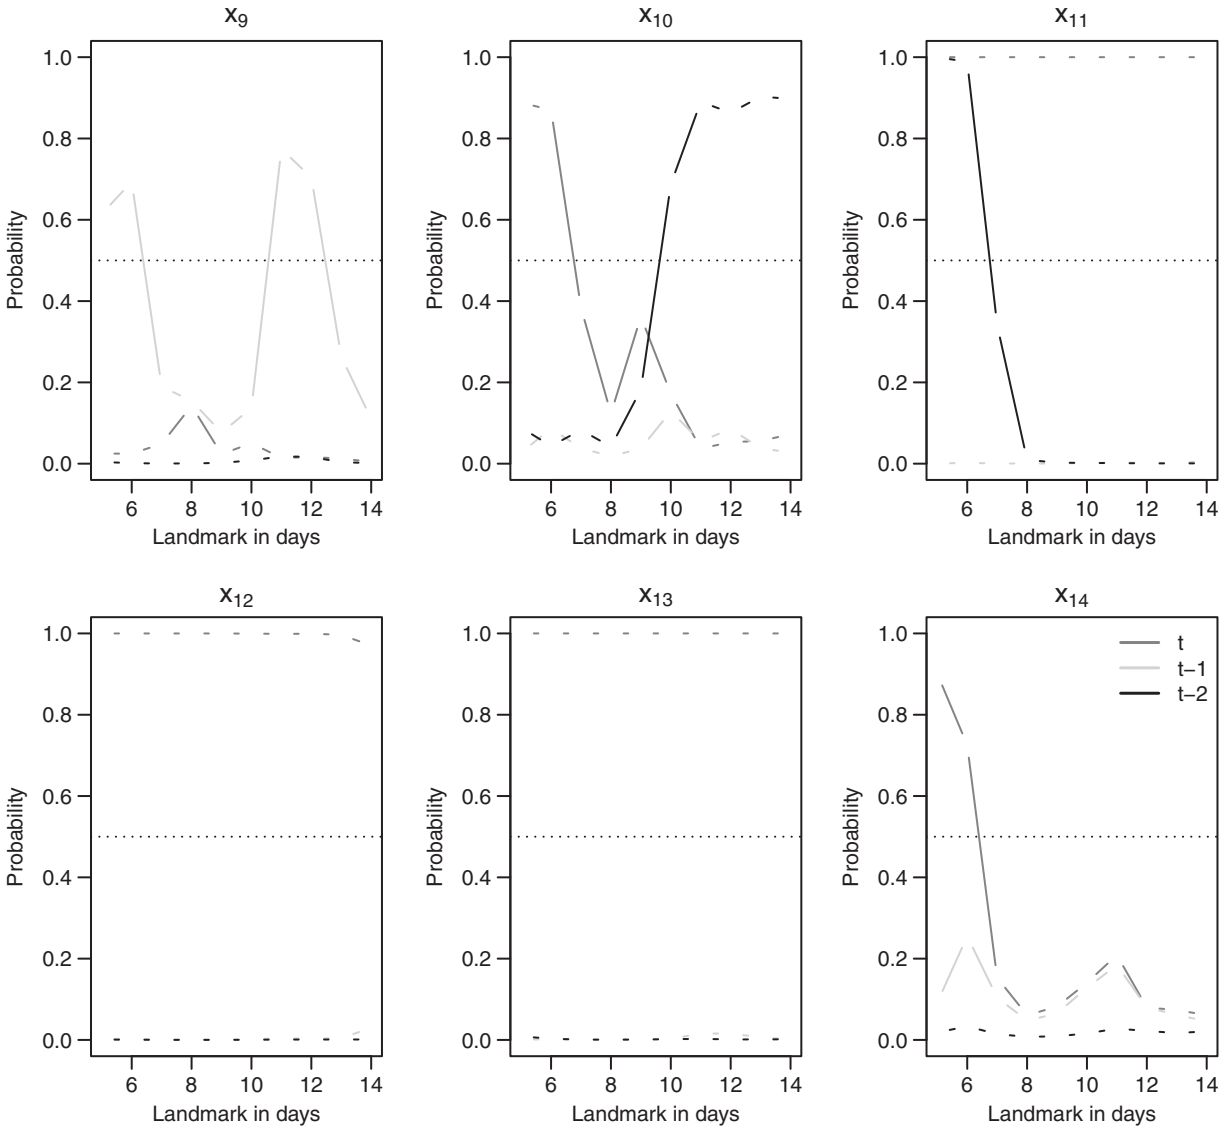

FIGURE 5 Posterior inclusion probabilities for the lagged time-dependent variables for each daily variable at each landmark. The inclusion threshold is set to 0.5

0.5 in the first step. To obtain dynamic median probability models we select the variables with PIP ≥ 0.5 at each landmark. Figure 7 represents these dynamic MPMs after we applied the cause-specific variable selection methodology for each landmark: the shrunken coefficients are dependent on the days since start of ventilation, meaning the landmarks. At a fixed landmark they are set to zero if their variable's inclusion probability in Figure 6 is smaller than 0.5. The CSVS methodology is applied on each landmark-specific MPM. We refer to Figure 3 of the Supplementary Material for a representation of the same dynamic MPMs before cause-specific variable selection. Table 3 summarizes which coefficient of which variables are finally selected to compute the cause-specific odds defined in Equation (8) at different landmarks.

6 | DISCUSSION

In this article, we presented an efficient approach to define dynamic prediction models for discrete-time competing risks data. We first extended methodology for objective Bayesian variable selection based on TBFs to discrete time-to-event models with competing risks. Then we applied the methodology to landmark-specific datasets in order to obtain dynamic median probability models. In a last step, we proceeded to cause-specific variable selection by updating the coefficients of the dynamic MPMs depending on the importance of the variables on each outcome individually. Using this approach we were able to define simple,

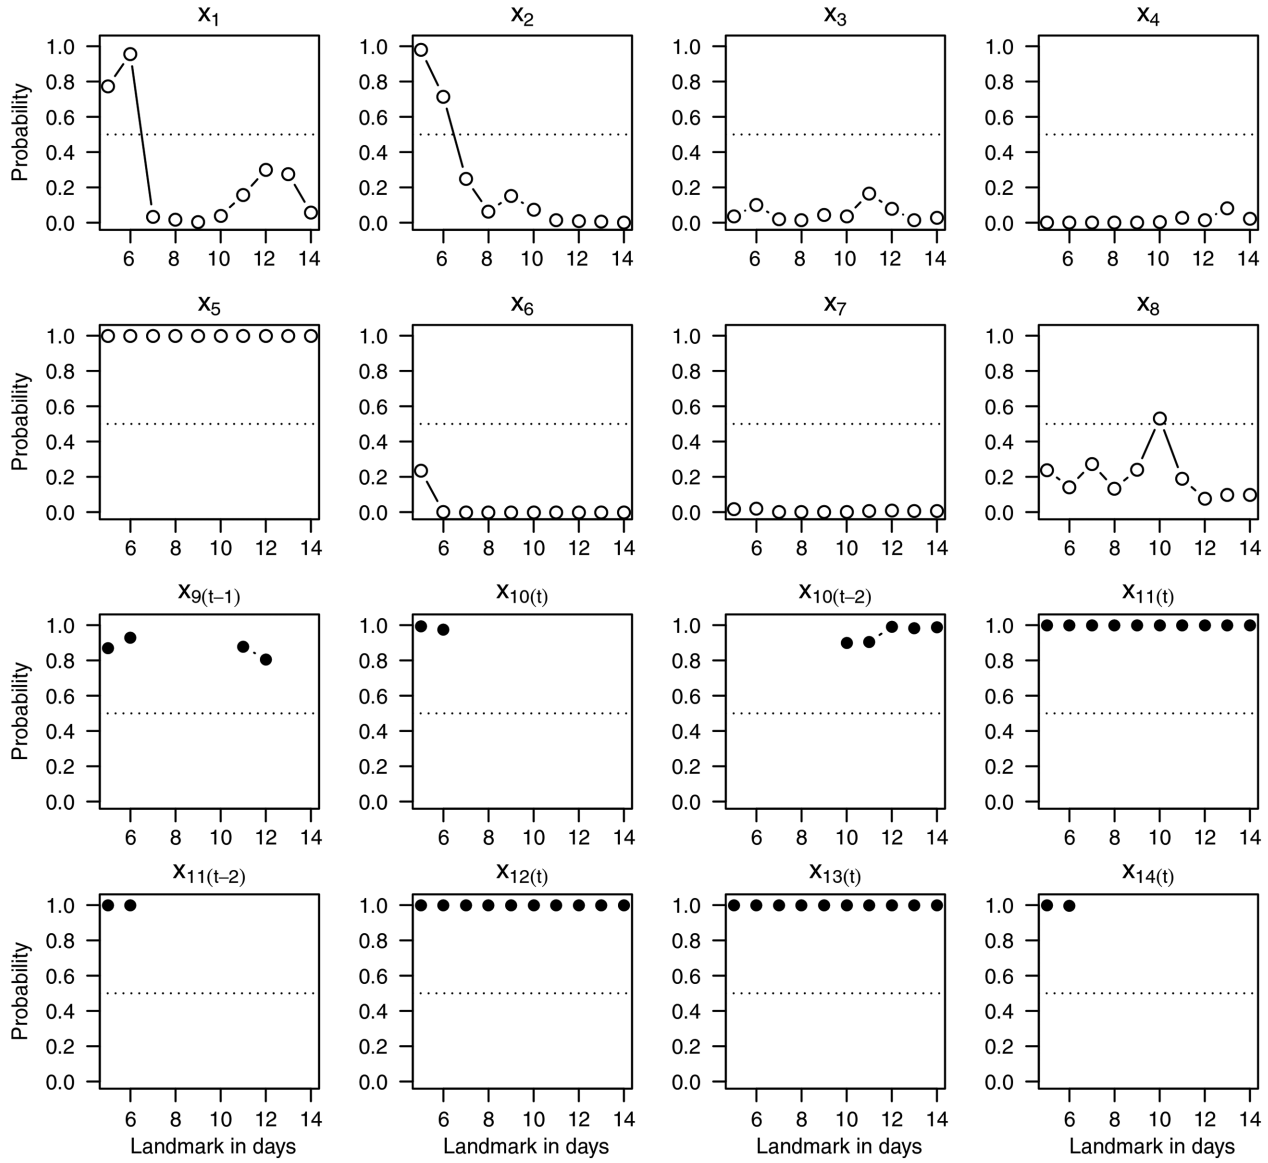

FIGURE 6 Posterior inclusion probabilities for the different baseline variables (blank circles) at each landmark and for the time-dependent lagged variables (filled circles) at each landmark for which the lagged variable was selected in the previous step

parsimonious prediction models for the risk of a VAP PA diagnosis on each day of ventilation taking the competing events, "VAP noPA," "dead," and "extubation," into account. To predict the odds of a specific outcome we only use the predictors judged relevant for that particular cause. The proposed method enables us to choose the best lag that should be used for a time-dependent variable. This issue has not been sufficiently discussed before; lags were mostly chosen based on clinical knowledge. In our application time-dependent lagged variables were identified as particularly important for prediction. From eight baseline variables only three remained in the final model whereas all the time-varying variables were included with different lags at least for early landmarks. We consider our approach easy, comprehensible and reproducible, leading to relatively parsimonious dynamic prediction models in a complicated clinical setting. Furthermore, the methods discussed here induce optimal shrinkage that improves the predictions.

Once a prediction model has been developed, its prediction performance has to be assessed (Steyerberg, 2009). This step is outside the scope of this paper but the methods that will be used in future work to validate the models developed in the previous sections are discussed here. In order to assess the ability of the prediction model to discriminate between patients experiencing specific events, cause-specific time-dependent area under the receiver operating characteristic curve (AUC) can be computed. The AUC defined by Blanche, Dartigues, and Jacqmin-Gadda (2013) for the continuous setting can easily be extended to the discrete framework. Moreover, to know how well the observed outcomes and the predictions agree, cause-specific calibration

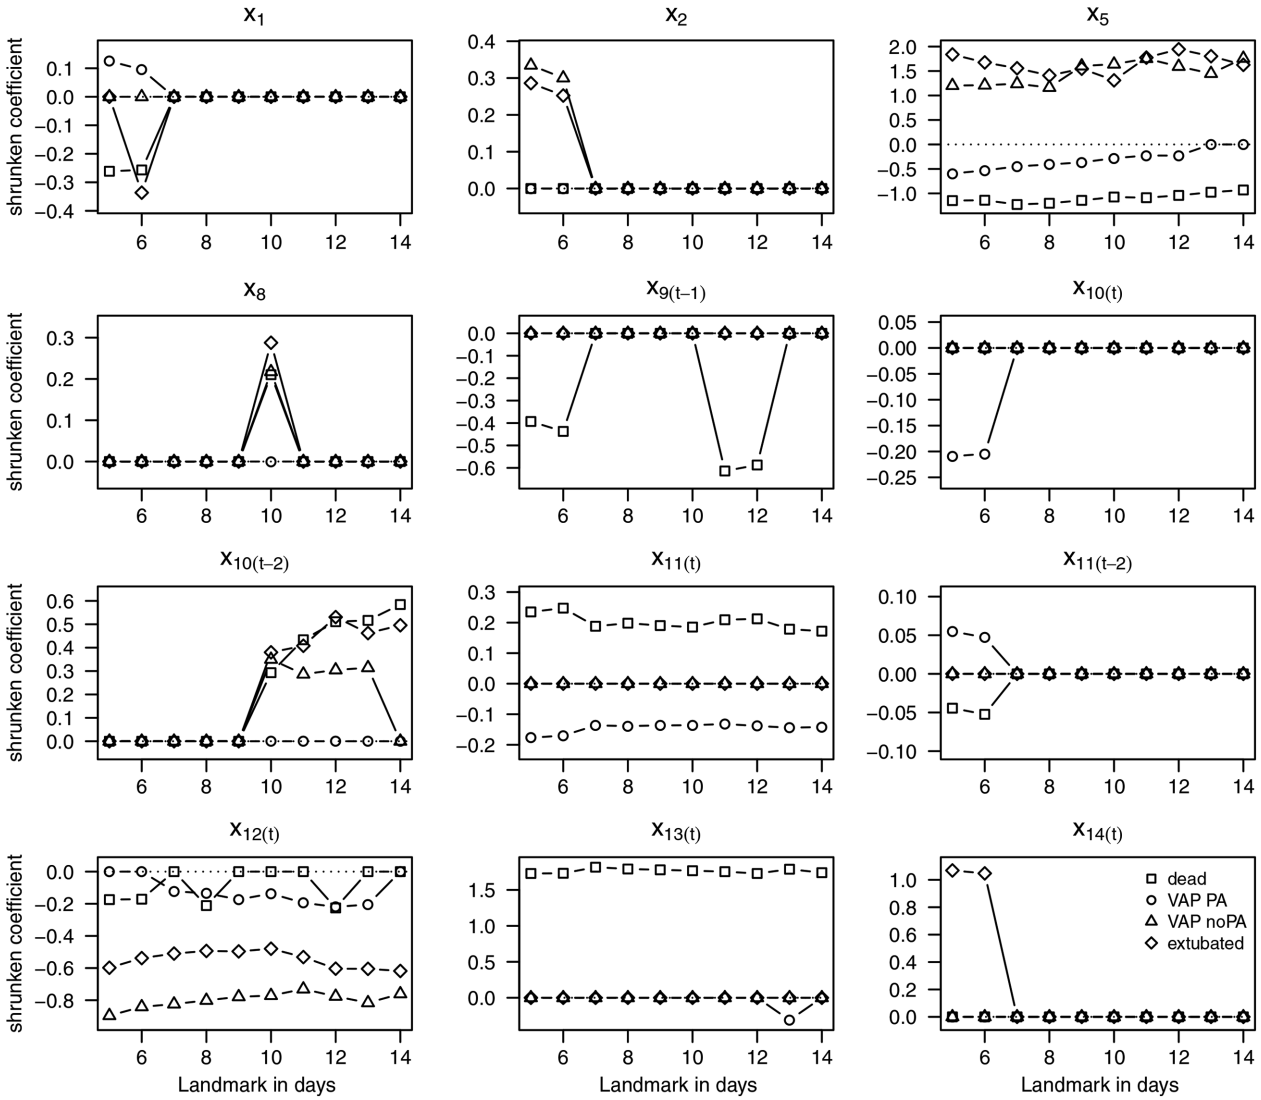

FIGURE 7 Shrunken cause-specific regression coefficients of the dynamic median probability models after dynamic variable selection via landmarking and cause-specific variable selection for all the variables with coefficients not always equal to zero. If the inclusion probability is smaller than the 0.5 threshold, the variable is excluded and the coefficient is set to zero

slopes can be obtained (Van Hoorde et al., 2014). This approach can be used for our class of models by adjusting for the time-dependency of the intercept. Finally, to summarize calibration and discrimination, proper scoring rules can be very useful to assess and compare the performance of different prediction models. Here, the Brier score, also referred to as prediction error, is the most commonly used score in survival analysis with one or more outcomes (Gerds & Schumacher, 2006; Schoop, Beyersmann, Schumacher, & Binder, 2011). Again these methods can be extended to the discrete-time setting.

For our dynamic prediction models we always considered the MPM in which variables are included if their PIP was higher or equal than the threshold 0.5. Heyard and Held (2018) propose a generalization of the MPM, the quantile probability model in which the inclusion threshold (0.5 for the MPM) is not predefined but selected such that the final model minimizes a certain model selection criterion; for example the deviance information criterion. However, this approach slows down the computation considerably in the dynamic setting of this paper such that we do not use it here.

Our methodology is directly extendable to discrete multistate model. Steele, Goldstein, and Browne (2004) define the transition hazards that are of interest in this model class very similar to the cause-specific hazards modeled in a competing risks scenario. The discrete transition hazard can be reduced to the probability of an individual moving to a specific state of interest at a fixed time point given that the individual is in an other state just prior to the fixed time point.

The methodology is not directly applicable for a large number of potential covariates p, since computation time of the inclusion probabilities would become prohibitive. Another limitation of the method is that we may run into (quasi-) complete separation

TABLE 3 The variables included for the calculation of the cause-specific odds of each outcome at each landmark

Outcome	Landmark	Variables
VAP PA	5	$x_2 + x_5 + x_{12(t)} + x_{14(t)}$
	6	$x_1 + x_2 + x_5 + x_{12(t)} + x_{14(t)}$
	7	$x_5 + x_{12(t)}$
	8	$x_5 + x_{12(t)}$
	9	$x_5 + x_{12(t)}$
	10	$x_5 + x_8 + x_{10(t-2)} + x_{12(t)}$
	11	$x_5 + x_{10(t-2)} + x_{12(t)}$
	12	$x_5 + x_{10(t-2)} + x_{12(t)}$
	13	$x_5 + x_{10(t-2)} + x_{12(t)}$
	14	$x_5 + x_{10(t-2)} + x_{12(t)}$
VAP noPA	5	$x_2 + x_5 + x_{12(t)}$
	6	$x_2 + x_5 + x_{12(t)}$
	7	$x_5 + x_{12(t)}$
	8	$x_5 + x_{12(t)}$
	9	$x_5 + x_{12(t)}$
	10	$x_5 + x_8 + x_{10(t-2)} + x_{12(t)}$
	11	$x_5 + x_{10(t-2)} + x_{12(t)}$
	12	$x_5 + x_{10(t-2)} + x_{12(t)}$
	13	$x_5 + x_{10(t-2)} + x_{12(t)}$
	14	$x_5 + x_{12(t)}$
dead	5	$x_1 + x_5 + x_{9(t-1)} + x_{11(t)} + x_{11(t-2)} + x_{12(t)} + x_{13(t)}$
	6	$x_1 + x_5 + x_{9(t-1)} + x_{11(t)} + x_{11(t-2)} + x_{12(t)} + x_{13(t)}$
	7	$x_5 + x_{11(t)} + x_{13(t)}$
	8	$x_5 + x_{11(t)} + x_{12(t)} + x_{13(t)}$
	9	$x_5 + x_{11(t)} + x_{13(t)}$
	10	$x_5 + x_8 + x_{10(t-2)} + x_{11(t)} + x_{13(t)}$
	11	$x_5 + x_{9(t-1)} + x_{10(t-2)} + x_{11(t)} + x_{13(t)}$
	12	$x_5 + x_{9(t-1)} + x_{10(t-2)} + x_{11(t)} + x_{12(t)} + x_{13(t)}$
	13	$x_5 + x_{10(t-2)} + x_{11(t)} + x_{13(t)}$
	14	$x_5 + x_{10(t-2)} + x_{11(t)} + x_{13(t)}$
extubated	5	$x_1 + x_5 + x_{10(t)} + x_{11(t)} + x_{11(t-2)}$
	6	$x_1 + x_5 + x_{10(t)} + x_{11(t)} + x_{11(t-2)}$
	7	$x_5 + x_{11(t)} + x_{12(t)}$
	8	$x_5 + x_{11(t)} + x_{12(t)}$
	9	$x_5 + x_{11(t)} + x_{12(t)}$
	10	$x_5 + x_{11(t)} + x_{12(t)}$
	11	$x_5 + x_{11(t)} + x_{12(t)}$
	12	$x_5 + x_{11(t)} + x_{12(t)}$
	13	$x_{11(t)} + x_{12(t)} + x_{13(t)}$
	14	$x_{11(t)}$

problems (Heinze & Schemper, 2002). In our application most of the events occur during the two to three first weeks. Setting a spline (with tree knots) on the day of ventilation partly resolves this potential problem since we take the timing of the events into account as we do not estimate baseline hazards at distinct days, but for a sequence of days. However, with the landmarking approach the problem of separation reoccurs because there are only few data left at higher landmarks.

ACKNOWLEDGMENT

We thank Eva Furrer for proofreading an earlier version of this manuscript. We acknowledge helpful comments by two referees on earlier versions of this article. The research leading to these results has received support from the Innovative Medicines Initiative Joint Undertaking under grant agreement no. 115737-2 (Combatting bacterial resistance in Europe—molecules against gram negative infections [COMBACTE-MAGNET]). The OUTCOMEREA database was used in the application; members of the OUTCOMEREA study group are listed in the appendix.

ORCID

Rachel Heyard http://orcid.org/0000-0002-7531-4333
Leonhard Held http://orcid.org/0000-0002-8686-5325

REFERENCES

Allison, P. D. (1982). Discrete-time methods for the analysis of event histories. *Sociological Methodology*, *13*, 61–98.

Anderson, J. R., Cain, K. C., & Gelber, R. D. (1983). Analysis of survival by tumor response. *Journal of Clinical Oncology*, *1*, 710–719.

Barbieri, M. M., & Berger, J. O. (2004). Optimal predictive model selection. *Annals of Statistics*, *32*, 870–897.

Barnett, A. G., & Graves, N. (2008). Competing risks models and time-dependent covariates. *Critical Care*, *12*, 134.

Barnett, A. G., Batra, R., Graves, N., Edgeworth, J., Robotham, J., & Cooper, B. (2009). Using a longitudinal model to estimate the effect of methicillin-resistant *Staphylococcus aureus* infection on length of stay in an intensive care unit. *American Journal of Epidemiology*, *170*, 1186–1194.

Bekaert, M., Timsit, J.-F., Vansteelandt, S., Depuydt, P., Vésin, A., Garrouste-Orgeas, M., ..., Benoit, D. (2011). Attributable mortality of ventilator-associated pneumonia: A reappraisal using causal analysis. *American Journal of Respiratory and Critical Care Medicine*, *184*, 1133–1139.

Beyersmann, J., Allignol, A., & Schumacher, M. *Competing Risks and Multistate Models with R*. New York, NY: Springer, 2012.

Blanche, P., Dartigues, J.-F., & Jacqmin-Gadda, H. (2013). Estimating and comparing time-dependent areas under receiver operating characteristic curves for censored event times with competing risks. *Statistics in Medicine*, *32*, 5381–5397.

Bouadma, L., Sonneville, R., Garrouste-Orgeas, M., Darmon, M., Souweine, B., Voiriot, G., ..., Timsit, J.-F. (2015). Ventilator-associated events: Prevalence, outcome, & relationship with ventilator-associated pneumonia. *Critical Care Medicine*, *43*, 1798–1806.

Copas, J. B. (1997). Using regression models for prediction: Shrinkage and regression to the mean. *Statistical Methods in Medical Research*, *6*, 167–183.

D'Agostino, R. B., Lee, M.-L., Belanger, A. J., Cupples, L. A., Anderson, K., & Kannel, W. B. (1990). Relation of pooled logistic regression to time-dependent Cox regression analysis: The Framingham heart study. *Statistics in Medicine*, *9*, 1501–1515.

Efron, B. (1988). Logistic regression, survival analysis, & the Kaplan-Meier curve. *Journal of the American Statistical Association*, *83*, 414–425.

Gerds, T. A., & Schumacher, M. (2006). Consistent estimation of the expected Brier score in general survival models with right-censored event times. *Biometrical Journal*, *48*, 1029–1040.

Gustafson, P., & Lefebvre, G. (2008). Bayesian multinomial regression with class-specific predictor selection. *The Annals of Applied Statistics*, *2*, 1478–1502.

Heinze, G., & Schemper, M. (2002). A solution to the problem of separation in logistic regression. *Statistics in Medicine*, *21*, 2409–2419.

Held, L., Sabanés Bové, D., & Gravestock, I. (2015). Approximate Bayesian model selection with the deviance statistic. *Statistical Science*, *30*, 242–257.

Held, L., Gravestock, I., & Sabanés Bové, D. (2016). Objective Bayesian model selection for Cox regression. *Statistics in Medicine*, *35*, 5376–5390.

Heyard, R., & Held, L. (2018). The quantile probability model. *Revised for Computational Statistics and Data Analysis*, in press, https://doi.org/10.1016/j.csda.2018.08.022.

Hu, J., & Johnson, V. E. (2009). Bayesian model selection using test statistics. *Royal Statisical Society. Series B*, *71*, 143–158.

Hunter, J. D. (2012). Ventilator associated pneumonia. *BMJ*, *344*, e3325.

Johnson, V. E. (2008). Properties of Bayes factors based on test statistics. *Scandinavian Journal of Statistics*, *35*, 354–368.

Le Gall, J., Lemeshow, S., & Saulnier, F. (1993). A new simplified acute physiology score (SAPS II) based on a European/North American multicenter study. *JAMA*, *270*, 2957–2963.

Liang, F., Paulo, R., Molina, G., Clyde, M. A., & Berger, J. O. (2008). Mixtures of *g* priors for Bayesian variable selection. *Journal of the American Statistical Association*, *103*, 410–423.

Möst, S., Pößnecker, W., & Tutz, G. (2016). Variable selection for discrete competing risks models. *Quality & Quantity*, *50*, 1589–1610.

Schoop, R., Beyersmann, J., Schumacher, M., & Binder, H. (2011). Quantifying the predictive accuracy of time-to-event models in the presence of competing risks. *Biometrical Journal*, *53*, 88–112.

Scott, J. G., & Berger, J. O. (2010). Bayes and empirical-Bayes multiplicity adjustment in the variable-selection problem. *Annals of Statistics*, *38*, 2587–2619.

Singer, J. D., & Willett, J. B. (1993). It's about time: Using discrete-time survival analysis to study duration and the timing of event. *Journal of Educational Statistics*, *18*, 155–195.

Steele, F., Goldstein, H., & Browne, W. (2004). A general multilevel multistate competing risks model for event history data, with an application to a study of contraceptive use dynamics. *Statistical Modelling*, *4*, 145–159.

Steyerberg, E. W. *Clinical Prediction Models*. New York, NY: Springer, 2009.

Timsit, J.-F., Fosse, J.-P., Troché, G., De Lassence, A., Alberti, C., Garrouste-Orgeas, M., ..., OUTCOMEREA Study Group. (2002). Calibration and discrimination by daily logistic organ dysfunction scoring comparatively with daily sequential organ failure assessment scoring for predicting hospital mortality in critically ill patients. *Critical Care Medicine*, *30*, 2003–2013.

Truche, A.-S., Darmon, M., Bailly, S., Clec'h, C., Dupuis, C., Misset, B., ..., Timsit, J.-F. (2016). Continuous renal replacement therapy versus intermittent hemodialysis in intensive care patients: Impact on mortality and renal recovery. *Intensive Care Medicine*, *42*, 1408–1417.

Tutz, G. (1995). Competing risks models in discrete time with nominal or ordinal categories of response. *Quality and Quantity*, *29*, 405–420.

Tutz, G., & Schmid, M. *Modeling discrete time-to-event data*. New York, NY: Springer, 2016.

Van Hoorde, K., Vergouwe, Y., Timmerman, D., Van Huffel, S., Steyerberg, E. W., & Van Calster, B. (2014). Assessing calibration of multinomial risk prediction models. *Statistics in Medicine*, *33*, 2585–2596.

van Houwelingen, H. C. (2007). Dynamic prediction by landmarking in event history analysis. *Scandinavian Journal of Statistics*, *34*, 70–85.

van Houwelingen, H. C. & Putter, H. (2012). *Dynamic prediction in clinical survival analysis*. Boca Raton, FL: CRC Press.

Wolkewitz, M., Cooper, B. S., Bonten, M. J. M., Barnett, A. G., & Schumacher, M. (2014). Interpreting and comparing risks in the presence of competing events. *BMJ*, *349*, 5060.

SUPPORTING INFORMATION

Additional Supporting Information including source code to reproduce the results may be found online in the supporting information tab for this article.

How to cite this article: Heyard R, Timsit J-F, Essaied WI, Held L. Dynamic clinical prediction models for discrete time-to-event data with competing risks—A case study on the OUTCOMEREA database. *Biometrical Journal*. 2018;1–21. https://doi.org/10.1002/bimj.201700259

APPENDIX A: MEMBERS OF THE OUTCOMEREA STUDY GROUP

Scientific Committee: Jean-François Timsit (Medical and Infectious Diseases ICU, Bichat-Claude Bernard Hospital, Paris, France; UMR 1137 Inserm - Paris Diderot university IAME, F75018, Paris); Elie Azoulay (Medical ICU, Saint Louis Hospital, Paris, France); Maïté Garrouste-Orgeas (ICU, Saint-Joseph Hospital, Paris, France); Jean-Ralph Zahar (Infection Control Unit, Angers Hospital, Angers, France); Christophe Adrie (ICU, Delafontaine Hospital, Saint Denis, and Physiology, Cochin Hospital, Paris, France); Michael Darmon (Medical ICU, Saint Etienne University Hospital, St Etienne, France); and Christophe Clec'h (ICU, Avicenne Hospital, Bobigny, and UMR 1137 Inserm—Paris Diderot university IAME, F75018, Paris, France).

Biostatistical and Information System Expertise: Jean-Francois Timsit (Medical and Infectious Diseases ICU, Bichat-Claude Bernard Hospital, Paris, France; UMR 1137 Inserm - Paris Diderot university IAME, F75018, Paris); Corinne Alberti (Medical Computer Sciences and Biostatistics Department, Robert Debré Hospital, Paris, France); Adrien Français (Integrated Research Center U823, Grenoble, France); Aurélien Vesin (OUTCOMEREA organization and Integrated Research Center U823, Grenoble, France); Stephane Ruckly (OUTCOMEREA organization and Inserm UMR 1137 IAME, F75018, Paris); Sébastien Bailly (Grenoble university hospital Inserm UMR 1137 IAME, F75018, Paris); Frederik Lecorre (Supelec, France); Didier Nakache (Conservatoire National des Arts et Métiers, Paris, France); and Aurélien Vannieuwenhuyze (Tourcoing, France).

Investigators of the OUTCOMEREA Database: Christophe Adrie (ICU, Delafontaine Hospital, Saint Denis, and Physiology, Cochin Hospital, Paris, France); Bernard Allaouchiche (ICU, Pierre benite Hospital, Lyon, France); Laurent Argaud (Medical ICU, Hospices Civils de Lyon, Lyon, France); Claire Ara-Somohano (Medical ICU, University Hospital, Grenoble, France); Elie Azoulay (Medical ICU, Saint Louis Hospital, Paris, France); Francois Barbier (medical-surgical ICU, Orleans, France), Jean-Pierre Bedos (ICU, Versailles Hospital, Versailles, France); Julien Bohé (ICU, Hôpital Pierre Benite, Lyon France), Lila

Bouadma (ICU, Bichat Hospital, Paris, France); Christine Cheval (ICU, Hyeres Hospital, Hyeres, France); Christophe Clec'h (ICU, Avicenne Hospital, Bobigny, France); Michael Darmon (ICU, Saint Etienne Hospital, Saint Etienne, France); Anne-Sylvie Dumenil (Antoine Béclère Hospital, Clamart, France); Claire Dupuis (Bichat hospital and UMR 1137 Inserm - Paris Diderot university IAME, F75018, Paris, France), Jean-Marc Forel (AP HM, Medical ICU, Hôpital Nord Marseille), Marc Gainier (la Timone hospital, Marseille, France), Akim Haouache (Surgical ICU, H Mondor Hospital, Creteil, France); Samir Jamali (ICU, Dourdan, Dourdan Hospital, Dourdan, France); Hatem Kallel (ICU, Cayenne General Hospital, Cayenne, France); Alexandre Lautrette (ICU, G Montpied Hospital, Clermont-Ferrand, France); Guillaume Marcotte (Surgical ICU, Hospices Civils de Lyon, Lyon, France); Benoît Misset (ICU, Saint-Joseph Hospital, Paris, France); Delphine Moreau (ICU, Saint-Louis Hospital, Paris, France); Bruno Mourvillier (ICU, Bichat Hospital, Paris, France); Laurent Papazian (Hopital Nord, Marseille, France), Benjamin Planquette (pulmonology ICU, George Pompidou hospital Hospital, Versailles, France); Bertrand Souweine (ICU, G Montpied Hospital, Clermont-Ferrand, France); Carole Schwebel (ICU, A Michallon Hospital, Grenoble, France); Gilles Troché (ICU, Antoine Béclère Hospital, Clamart, France); Marie Thuong (ICU, Delafontaine Hospital, Saint Denis, France); Guillaume Thierry (ICU, Saint-Louis Hospital, Paris, France); Dany Toledano (ICU, Gonesse Hospital, Gonesse, France); and Eric Vantalon (SICU, Saint-Joseph Hospital, Paris, France).

Study Monitors: Julien Fournier, Caroline Tournegros, Stéphanie Bagur, Mireille Adda, Vanessa Vindrieux, Loic Ferrand, Nadira Kaddour, Boris Berthe, Samir Bekkhouche, Kaouttar Mellouk, Sylvie Conrozier, Igor Theodose, Veronique Deiler, and Sophie Letrou.