

Marginal Structural Models for Partial Exposure Regimes

Stijn Vansteelandt*

Karl Mertens†

Carl Suetens‡

Els Goetghebeur**

*Ghent University, stijn.vansteelandt@ugent.be

†Scientific Institute of Public Health

‡Scientific Institute of Public Health

**Harvard School of Public Health and Ghent University, egoetghe@hsph.harvard.edu

This working paper is hosted by The Berkeley Electronic Press (bepress) and may not be commercially reproduced without the permission of the copyright holder.

<http://biostats.bepress.com/harvardbiostat/paper78>

Copyright ©2008 by the authors.

Marginal Structural Models for Partial Exposure Regimes

Stijn Vansteelandt, Karl Mertens, Carl Suetens, and Els Goetghebeur

Abstract

Intensive care unit (ICU) patients are well known to be highly susceptible for nosocomial (i.e. hospital-acquired) infections due to their poor health and many invasive therapeutic treatments. The effects of acquiring such infections in ICU on mortality are however ill understood. Our goal is to quantify these effects using data from the National Surveillance Study of Nosocomial Infections in Intensive Care

Units (Belgium). This is a challenging problem because of the presence of time-dependent confounders (such as exposure to mechanical ventilation) which lie on the causal path from infection to mortality. Standard statistical analyses may be severely misleading in such settings and have shown contradicting results.

While inverse probability weighting for marginal structural models can be used to accommodate time-dependent confounders, inference for the effect of

ICU acquired infections on mortality under such models is further complicated (a) by the fact that marginal structural models infer the effect of acquiring infection on a given, fixed day "in ICU", which is not well defined when ICU discharge comes prior to that day; (b) by informative censoring of the survival time due to hospital discharge; and (c) by the instability of the inverse weighting estimation procedure. We accommodate these problems by developing inference under a new class of marginal structural models which describe the hazard of death for patients if, possibly contrary to fact, they stayed in the ICU for at least a given number of days s and acquired infection or not on that day. Using these models we estimate that, if patients stayed in the ICU for at least s days, the effect of acquiring infection on day s would be to multiply the subsequent hazard of death

by 2.74 (95 per cent conservative CI 1.48; 5.09).

Marginal Structural Models for Partial Exposure Regimes

STIJN VANSTEELANDT

*Department of Applied Mathematics and Computer Science, Ghent University,
Ghent, Belgium*

stijn.vansteelandt@ugent.be

KARL MERTENS, CARL SUETENS

Epidemiology Unit, Scientific Institute of Public Health, Brussels, Belgium

AND ELS GOETGHEBEUR

*Department of Applied Mathematics and Computer Science, Ghent University,
Ghent, Belgium, and Department of Biostatistics,
Harvard School of Public Health, Boston, USA*

Intensive care unit (ICU) patients are well known to be highly susceptible for nosocomial (i.e. hospital-acquired) infections due to their poor health and many invasive therapeutic treatments. The effects of acquiring such infections in ICU on mortality are however ill understood. Our goal is to quantify these effects using data from the National Surveillance Study of Nosocomial Infections in Intensive Care Units (Belgium). This is a challenging problem because of the presence of time-dependent confounders (such as exposure to mechanical ventilation) which lie on the causal path from infection to mortality. Standard statistical analyses may be severely misleading in such settings and have shown contradicting results.

While inverse probability weighting for marginal structural models can be used to accommodate time-dependent confounders, inference for the effect of ‘ICU acquired infections on mortality under such models is further complicated (a) by the fact that marginal structural models infer the effect of acquiring infection on a given, fixed day ‘in ICU’, which is not well defined when ICU discharge comes prior to that day; (b) by informative censoring of the survival time due to hospital discharge; and (c) by the instability of the inverse weighting estimation procedure. We accommodate these problems by developing inference under a new class of marginal structural models which describe the hazard of death for patients if, possibly contrary to fact, they stayed in the ICU for at least a given number of days s and acquired infection or not on that day. Using these models we estimate that, if patients stayed in the ICU for at least s days, the effect of acquiring infection on day s would be to multiply the subsequent hazard of death by 2.74 (95 per cent conservative CI 1.48; 5.09).

KEY WORDS: causal inference; direct effect; intermediate variables; marginal structural models; time-dependent confounding

1 INTRODUCTION

Intensive care unit (ICU) patients are estimated to have a 5 to 10 times higher risk of acquiring nosocomial (i.e. hospital-acquired) infections than patients in other hospital units, due to their poor health conditions and the many invasive therapeutic treatments to which they are typically subjected. Such infections are believed to account for 50% of all major complications of hospitalization. They are considered to have a substantial impact on morbidity, mortality and medical costs (Gaynes 1997) and thus to pose a major public health burden. In 1985, the SENIC study (Haley et al. 1985) demonstrated that surveillance of nosocomial infections can reduce infection rates by as much as 30%, provided that sufficient infection control staff is available and that surveillance results are used in the infection control management. Since then, surveillance of nosocomial infections has played a fundamental role in assessing and improving the quality of medical care.

In 1995, the Scientific Institute of Public Health – Louis Pasteur (Belgium) set up a national surveillance network in ICUs in collaboration with the Belgian Society for Intensive Care and Emergency Medicine (Suetens et al. 1999). The aim of this network is twofold. First, to assist individual ICUs to obtain local incidence statistics for ICU acquired Nosocomial Pneumonia (NP; i.e. one specific nosocomial infection) and Bloodstream Infection. Second, to offer national statistics in parallel to guide the interpretation of each ICU's performance. To this end, surveillance follows a standard protocol which is largely based on a consensus obtained at the European level in the HELICS (Hospitals in Europe Link for Infection Control through Surveillance) project in 1995 (Suetens et al. 2003).

In this article, we will use data collected through the network to quantify the effect of NP on subsequent mortality in ICU patients. Estimating the mortality rate attributable to NP is however a complex problem for various reasons. First, the association between infection and mortality is disturbed by time-dependent confounders (i.e. time-dependent variables which simultaneously affect mortality and infection). For

instance, daily exposure to invasive treatments such as mechanical ventilation, the presence of a central vascular catheter, parenteral feeding (i.e. feeding given by injection, bypassing the gastrointestinal tract), ... increases the risk of NP at subsequent time points and the poor health conditions leading to these treatments are also indicative of an increased mortality risk. These confounders lie on the causal path from infection to mortality because infection makes patients more prone to receive invasive therapeutic treatments. Standard adjustment approaches, such as time-dependent proportional hazards regression, will therefore usually give biased results (see for example Kalbfleisch and Prentice 1980, Robins 1986, Robins et al. 2000, Bryan et al. 2003, Vansteelandt 2007). Second, the survival time of patients is (informatively) censored upon hospital discharge. The decision to discharge a patient is closely related to his/her health status, as observed by physicians, suggesting that mortality rates may differ substantially between patients who are discharged on a given day and those who are not.

The problem of estimating the mortality rate attributable to NP has been reviewed by several researchers (Carlet 2001, Vincent 2003). Common practice is to fit logistic regression models for mortality in ICU, adjusting for NP status upon ICU discharge, for length of stay in ICU and for time-dependent variables measured prior to infection, or to fit proportional hazards models for time to death, adjusting for NP status upon ICU discharge or time-dependent NP status and for time-dependent variables measured prior to infection (Mertens et al. 2006a). These analyses ignore the aforementioned problems and empirical results have therefore remained controversial, with several studies reporting relative risk estimates for mortality ranging from neutral to severely harmful. The present study addresses the above problems by using marginal structural models (Hernán et al. 2000, van der Laan and Robins 2003, Bryan et al. 2004) for the analysis of the effects of NP on health outcomes in ICUs and applying them to the Belgian National Surveillance Study of nosocomial infections in ICUs.

We review the surveillance study in Section 2. The presence of time-dependent confounders which also lie on the causal path from infection to mortality leads us to consider marginal structural models, which we introduce in Section 3.1. However, stan-

standard inference for such models cannot be applied in this study for the following reasons. First, marginal structural models for the effect of ICU-acquired infection on death would describe the hazard of death in the possibly hypothetical situation where the patient acquires an (ICU-acquired) infection a given number of days since admission, but this is not well defined when ICU discharge comes earlier. Second, time-varying exposures (i.e. ICU acquired infection) and confounders were only recorded during the patient's stay in ICU, while survival times were recorded until hospital discharge, to alleviate the problem of informative censoring. Marginal structural proportional hazards methods cannot be directly applied to such data because they require the infection status and confounders to be observed from the start of the study until death or censoring of the survival time. Similar problems arise in observational studies with a mortality endpoint where exposures are incompletely measured due to loss to follow-up or end-of-study, but survival times are assessed much longer in time (using death registers, for instance).

To accommodate both problems, we propose a new class of marginal structural models in Section 3.2, which express the effect of acquiring infection on a given day s on the hazard of death if patients were kept in ICU for at least s days. We call the models in our class marginal structural models for partial exposure regimes as each considered 'exposure regime' (i.e. infection path) determines the 'exposures' (i.e. infections) for a given patient only up to the chosen time point s , leaving them unspecified (i.e. random and observational) afterwards. This has the added advantage of yielding more stable inferences since we merely aim to infer what the hazard of death for given patient would have been if, possibly contrary to fact, the patient's infection status on each day were as given up to a chosen time s , and not up to end-of-study time. We derive a class of estimators which are consistent and asymptotically normal under the considered model and provide a reasonably efficient estimator in that class. In Section 4, we present results obtained for the surveillance data. In Section 5, we discuss the usefulness of marginal structural models for partial exposure regimes in other settings and provide a comparison with structural nested models (Robins 1997b, Keiding et al. 1999).

2 SURVEILLANCE STUDY ON NOSOCOMIAL INFECTIONS

All ICUs in Belgian hospitals were invited to participate in this surveillance study on a voluntary basis and data were collected on all patients admitted to the ICU and staying more than 24 hours. Specifically, data were recorded on personal characteristics, reasons for ICU admission and baseline health status, as well as daily indicators of received invasive treatments and acquired infections in ICU. Nosocomial infections were defined as infections acquired by patients after the second day of ICU stay, this to exclude infections that were in incubation upon enrollment in the ICU. The third day of stay in ICU will therefore be the starting point for our analysis and, hence, patients staying less than 3 days will not be considered.

We will restrict the analysis to surveillance data collected for the year 2002 in one of the largest hospitals, because of its reputation to have accurate daily measurements of received invasive treatments and acquired infections. Our analysis will focus on the effect on mortality of nosocomial pneumonia. This is one of the main nosocomial infections, defined according to the 1995 HELICS protocol. A total of 1072 ICU patients were analysed. Of the 100 (9.3%) patients who acquired NP in ICU (and stayed more than 2 days), 41 (41%) died in hospital (of whom 27 died in ICU), as compared to 183 (18.8%) deaths among the 972 patients who remained NP-free in ICU (of whom 99 died in ICU). Among patients who stayed more than 2 days in ICU, the median length of stay in ICU was 4 days (IQR 3, 95th percentile 13) for patients without a history of NP and 16 days (IQR 13, 95th percentile 54.5) for the remaining patients. Additional background details on the surveillance study can be found in Mertens et al. (2006a).

A preliminary causal analysis (using marginal structural models, see Section 3.1) involving measurements collected in ICU only (i.e. censoring survival times upon ICU discharge) revealed highly unstable results (Mertens et al. 2006b). This was mainly due to the high censoring rates following ICU discharge. Using patient registers, the survival status of each patient was therefore assessed upon hospital discharge, which

typically comes later than ICU discharge. This will make results much more reliable, but at the same time complicates the analysis considerably because the exposure ‘ICU acquired infection’ is not well defined between ICU and hospital discharge.

Throughout this article, we will use the following notation. Let for each patient in the study, A, D, C and T be the observed times from admission in the ICU to (ICU-acquired) NP, discharge from ICU, discharge from hospital and time to death, respectively (see Figure 1). Let A_t be a counting process that indicates 1 for (ICU-acquired) infection at or prior to time t and 0 otherwise, where $A_0 = 0$ by definition and where A_t is observed for all discrete times $t < D$. Note that A can be recovered from the path $\{A_t, t = 0, \dots, D - 1\}$ (up to the resolution permitted by discrete time) when $A_{D-1} = 1$ and that A is ill defined otherwise. Likewise, let D_t (C_t) be a counting process that indicates 1 if ICU (hospital) discharge happened at or prior to time t and 0 otherwise. Further, define L_t to be a vector of time-dependent variables collected in the ICU at day t , which is observed for all $t < D$. Here, L_0 is a vector of baseline variables collected upon admission to the ICU. In our analyses, it consists of age, gender, reason for ICU admission, acute coronary care, multiple trauma, presence and type of infections upon ICU admission, prior surgery, baseline antibiotic use and the SAPS score. The latter is a severity score based on a set of 15 clinical parameters predicting the mortality risk of a patient admitted to the ICU (Le Gall et al. 1993). Further, $L_t, t > 0$ is a vector of invasive therapeutic treatment indicators collected on day t , consisting of indicators of exposure to mechanical ventilation, central vascular catheter, parenteral feeding, presence and/or feeding through naso- or oro-intestinal tube, tracheotomy intubation, nasal intubation, oral intubation, stoma feeding and surgery. Discharge from the ICU defines the end of follow-up for all measured variables, except survival time T , which is censored by discharge from the hospital. For any vector $Z = (Z_0, \dots, Z_K)$ we use \bar{Z}_t to denote the history (Z_0, \dots, Z_t) up to and including day t . Throughout, we assume that infection and discharge on day t can only be affected by time-dependent variables measured on previous days (and thus not by time-dependent variables measured on the same day).

3 MARGINAL STRUCTURAL MODELS FOR PARTIAL EXPOSURE REGIMES

3.1 Marginal Structural Models

Time-dependent multi-state models for event history analysis (Andersen and Keiding 2002) may appear naturally suited for addressing the multi-state nature (see Figure 1) of our problem. However, when the goal is to estimate the effect of acquiring infection on a given day in ICU on subsequent mortality, such approaches (as well other standard regression approaches) are typically biased whether or not one adjusts for the relevant past confounder history (Robins 1997b). For the unadjusted analysis, this is so because these analyses ignore time-varying confounders, like mechanical ventilation which increases the subsequent risk of infection and death. For the adjusted analysis, this is so when (as often) these time-varying confounders lie on the causal path from infection to mortality, because standard regression adjustment for such post-infection measurements then introduces bias. This problem of adjusting for internal time-dependent covariates has been known for a long time in the survival literature (Kalbfleisch and Prentice 1980), but solutions to it have emerged only relatively recently. One such solution, which is getting increasingly popular among statisticians and epidemiologists, is to use marginal structural Cox regression models (Hernán et al. 2001). We briefly review these models in this section.

Let $T_{\bar{a}}$ express the counterfactual survival time (Rubin 1978, Robins 1986) which an ICU patient would, possibly contrary to fact, have had under a given infection path $\bar{a} = (a_1, a_2, \dots, a_K)$ following which the patient is infected on day t since ICU admission if $a_t = 1$ and uninfected if $a_t = 0$, and where K represents end-of-study time. Then a marginal structural Cox regression model is a Cox regression model for the counterfactual survival time $T_{\bar{a}}$, possibly conditional on baseline covariates V . It thus expresses how the hazard of death would have been if, possibly contrary to fact, all subjects in the population had followed infection path \bar{a} . A simple example of a

marginal structural model is

$$\lambda_{\bar{a}}(t|V) = \lambda_0(t) \exp(\beta_1 a_t + \beta_2' V) \quad (1)$$

where $\lambda_{\bar{a}}(t|V)$ is the hazard of death at time t among subjects with baseline covariates V , had they all followed infection path \bar{a} , $\lambda_0(t)$ is an unknown baseline hazard of death at time t and β_1, β_2 are unknown parameters. In model (1), $\exp(\beta_1)$ expresses the causal rate ratio at time t of acquiring infection at time t . This represents the ratio of the mortality (hazard) rate at any time t had all patients with baseline covariates V acquired infection at time t compared to the mortality (hazard) rate at time t had these patients acquired no infection up to time t . Further, $\lambda_0(t)$ expresses the hazard of death for patients with $V = 0$ had they followed an infection path in which they acquire no infection. The model's name 'marginal' expresses that the model does not involve time-dependent confounders. Adjustment for such confounders happens by fitting the model to data from a pseudo-population in which there are no time-varying confounders (but the target effect is the same). This pseudo-population is constructed by reweighting subjects in the risk set at each time t by the inverse of the product of the conditional probabilities of the observed infection path up to time t , given the history of time-varying confounders (Hernán et al. 2001).

In our study, inference for marginal structural Cox models is not directly applicable to estimate the hazard of death under each infection path because the exposure status 'ICU acquired infection' is not well defined between ICU discharge and death or censoring of the survival time, by the same token that A is ill defined for subjects who get discharged uninfected from the ICU. Because our goal is to estimate the effect of *ICU acquired* infection on mortality, it would be natural to define patients uninfected during the entire study when they were not infected upon ICU discharge. However, this would make standard estimators for marginal structural models irregular (Robins et al. 2000) because there would be patients with certain prognostic factors (namely, patients who are discharged from ICU and were uninfected upon discharge) who are precluded from becoming infected under this definition. This follows from the assumption of experimentation in the assignment of infection (van der Laan and Robins 2003), according to

which, at each time t ,

$$0 < P(A_t = 1 | \bar{A}_{t-1}, \bar{L}_{t-1}, D_t, V) < 1 \text{ with probability 1.}$$

One solution is to view the infection and ICU discharge status of a patient at each time as a joint exposure by defining an infection path to be any regime (d, a, as) in which a patient, while alive, will be discharged from the ICU on day d and either acquire infection on a given day $s < d$ (if $a = 1$) or stay uninfected during his/her stay in ICU (if $a = 0$). Thus, patients following regime $(d, a, as) = (d, 0, 0)$ stay uninfected in ICU and are discharged at time d ; patients following regime $(d, a, as) = (d, 1, s)$ acquire infection at time $s < d$ (and not earlier) and are discharged at time d (and not earlier). The joint causal effect of discharge and infection in the ICU on the hazard of death can be expressed in function of baseline covariates V through a marginal structural Cox models for multiple interventions (Hernán et al. 2000, Robins et al. 2003). The following is a simple example of such model:

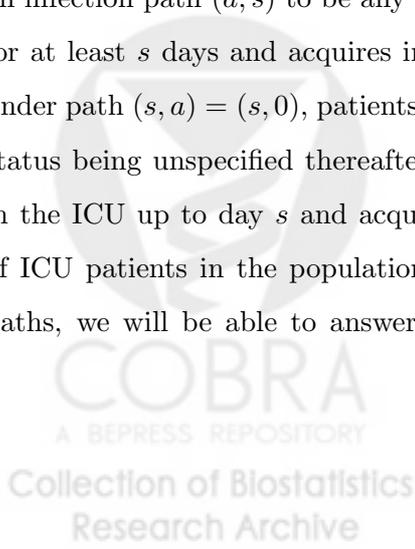
$$\lambda_{(d,a,as)}(t|V) = \lambda_0(t) \exp \{ (\beta_1 + \beta_2(t-s))aI(t \geq s) + (\beta_3 + \beta_4(t-d))I(t \geq d) + \beta_5'V \} \quad (2)$$

with $d > s$, $\lambda_{(d,a,as)}(t|V)$ the hazard of death at time t among subjects with baseline covariates V , had they all been exposed to infection path (d, a, as) , $\lambda_0(t)$ is an unknown baseline hazard of death at time t , and $\beta_1, \beta_2, \beta_3, \beta_4, \beta_5$ are unknown parameters. In particular, $\exp \{ \beta_1 + \beta_2(t-s) \}$ is the causal rate ratio at time t of acquiring infection at time $s, s \leq t, s < d$. This represents the ratio of the mortality (hazard) rate at any time t had all patients with baseline covariates V acquired infection at time s compared to the mortality (hazard) rate at time t had these patients experienced the same discharge time, but no infection up to (but not including) time t . Note that the causal effect parameter in the considered models has limited relevance from a public health perspective. First, it expresses the effect of acquiring NP at a given time on mortality in the hypothetical and unrealistic scenario where we would keep the patients in ICU until some given, later time. Second, by comparing the same group of patients under 2 possible infection histories, all other things - including time of discharge from

the ICU - being equal, $\exp\{\beta_1 + \beta_2(t - s)\}$ represents only the direct effect of acquiring infection at time s on mortality at time t under model (2). As such, it does not capture the indirect effect of infection on death that may arise when infection prolongs the time of stay in the ICU, which may itself affect mortality risk. Furthermore, estimates for the parameters in model (2) are highly unstable as a result of inverse weighting by small probabilities in the estimation procedure (Hernán et al. 2000, Robins et al. 2003). This is due to a lengthy follow-up for a limited number of patients and because many collected time-dependent variables are extremely predictive for censoring due to ICU discharge. In the next section, we develop a more appropriate solution.

3.2 Marginal Structural Models for Partial Infection Paths

We conclude from the previous discussion that inferring the effect of acquiring infection on a given day in ICU requires fixing the length of stay in ICU because this effect is not entirely well defined when ICU discharge comes prior to that day. At the same time, however, fixing the length of stay in ICU is problematic because (a) this is not feasible in practice (and hence the practical meaning of the estimated effect becomes more difficult under a hypothetical scenario which fixes the length of stay in ICU), and (b) by doing so, we miss the indirect effect of infection on mortality through modifying the length of stay in ICU. Rather than fixing the discharge time of patients up to the study end, we will therefore infer mortality rates under infection paths in which patients stay in the ICU for *at least* s days and acquire infection or not at that time (i.e. we fix the discharge status of a patient only up to the time of infection). Specifically, we redefine an infection path (a, s) to be any path in which a patient, while alive, stays in the ICU for at least s days and acquires infection (if $a = 1$) or not (if $a = 0$) on day s . Thus, under path $(s, a) = (s, 0)$, patients are uninfected in the ICU up to day s , their infection status being unspecified thereafter; under path $(s, a) = (s, 1)$, patients are uninfected in the ICU up to day s and acquire infection on day s . By analyzing mortality rates of ICU patients in the population following uniform application of different infection paths, we will be able to answer causal questions like ‘What would be the effect of



acquiring infection at time s on the mortality rate of ICU patients if they stayed in the ICU for at least s days?'. At the same time, we will be solving the problem that the infection status is unknown or ill defined after ICU discharge because we only consider regimes which fix the infection status of patients during their stay in ICU. Note that these infection paths are partially random and observational, and thus generalize the deterministic treatment regimes of Robins (1997a) which specify the treatment at each time from start till end of study.

For a given path (s, a) , let $T_{(s,a)}$ be the random variable representing the subject's time from admission in the ICU to death had he/she, possibly contrary to fact, experienced infection regime (s, a) rather than his/her own infection history, all other things being equal. We can then express the causal effect of infection in the ICU on the hazard of death through the marginal structural Cox model:

$$\lambda_{(s,a)}(t|V) = \lambda_0(t) \exp \{(\beta_1 s + \beta_2 a)I(t \geq s) + \beta_3' V\} \quad (3)$$

Here, $\lambda_{(s,a)}(t|V)$ is the hazard of death at time t among subjects with baseline covariates V , had they all been exposed to infection regime (s, a) , $\lambda_0(t)$ is an unknown baseline hazard of death at time t , and $\beta_1, \beta_2, \beta_3$ are unknown parameters. Note that $\lambda_0(t) = \lambda_{(0,0)}(t|0)$ is the hazard of death at time t among patients with $V = 0$ and is hence, in principle, directly estimable from the observed data. In addition, note that $\exp(\beta_2)$ is the causal rate ratio at time t of acquiring infection at any time $s, s \leq t$. It represents the ratio of the mortality (hazard) rate at any time t had all patients with baseline covariates V stayed in ICU up to at least time s and acquired infection at that time compared to the mortality (hazard) rate at time t had these patients also stayed in ICU up to at least time s but acquired no infection up to that time. By fixing the discharge status only up to the time of infection, $\exp(\beta_2)$ represents the overall effect of acquiring infection at time s on mortality at time t under model (3). We call (3) a marginal structural model for partial exposure regimes to express that it determines each exposure regime only for a fixed time period, contrary to the more standard marginal structural models of Section 3.1.

3.3 Inference

In this section, we develop inference for the parameters indexing marginal structural models for partial infection paths under the assumption of no-unmeasured-confounders that, for each path (s, a) and at each time $t \leq s$, survivors with prognostic factors $\bar{L}_{t-1}, \bar{A}_{t-1}, D_{t-1} = 0$ and V (compatible with the regime (s, a)) have the same hazard of infection and ICU discharge at time t regardless of their counterfactual survival time $T_{(s,a)}$; that is, $(A_t, D_t) \perp\!\!\!\perp T_{(s,a)} | \bar{L}_{t-1}, \bar{A}_{t-1}, D_{t-1} = 0, T > t$ for each time t and each infection path (s, a) . This assumption is reasonable when (a) the physician's decision to discharge a patient from the ICU at time t is based only on observed daily patient characteristics which were recorded in $\bar{L}_{t-1}, \bar{A}_{t-1}$ and V , and (b) in addition, all time-dependent confounders for the association between infection and death (i.e. all prognostic factors for death that affect a patient's susceptibility to NP) are recorded in the database (and accounted for).

Even if all patients were observed until the study end or death, analysis tools for marginal structural Cox models as described in Hernán et al. (2000) would not be directly applicable to fit model (3) under these assumptions. This is because, unlike in usual marginal structural models, each infection path is fixed for only a limited period of time and becomes observational afterwards. Below, we give a practical algorithm for obtaining consistent and asymptotically normal (CAN) estimators for the parameter $\beta = (\beta_1, \beta_2, \beta_3)'$ indexing model (3) in the absence of unmeasured time-dependent confounders. The motivation for this algorithm is given in the Appendix, where the resulting estimates are defined via weighted partial likelihood estimation.

1. First we identify, for each infection regime (s, a) , those patients in the database whose observed infection history is compatible with the regime (s, a) . Specifically, for each time t , we construct a vector of variables (S_t, A_t^*) which takes a given value (s, a) for given patient at that time if that patient's observed infection regime up to time t is compatible with the regime (s, a) (i.e. if the patient's data up to time t could have been obtained under the regime (s, a)). That is, for given s ,

$(S_t, A_t^*) = (s, 1)$ ($= (s, 0)$) for given patient at time t if that patient was in the ICU at time $s \leq t$ and acquired NP at that time (stayed NP-free up to that time). Contrary to inference for ordinary marginal structural models, the data for a given patient at a given time may be compatible with (and may hence carry information about) multiple infection regimes. This is because the considered regimes are partly observational. For instance, if a patient's data are compatible with regime $(s, 0)$ at time t (i.e. if the patient was in the ICU and uninfected at time s), then they are compatible with all regimes $(u, 0)$ for $u < s$. Henceforth, they may appear multiple times in the database corresponding to different values of S_t .

2. Next, for all infection regimes (s, a) jointly, we fit a proportional hazards model using only the data compatible with the given regime and weighting each observation by inverse the probability of following that regime to account for the selective nature of our subsample. Specifically, having added the variables S_t and A_t^* to the database, we fit the time-dependent Cox model

$$\lambda(t|S_t, A_t^*, V) = \lambda_0^*(t) \exp(\beta_1^* S_t + \beta_2^* A_t^* + \beta_3^* V) \quad (4)$$

as an analog to (3), where the contribution of a patient to the risk set at time t is weighted by the stabilized weights

$$\begin{aligned} swi(t, S_t, \bar{A}_t, \bar{D}_t, \bar{L}_{t-1}, V) &= \prod_{k=1}^{S_t} \frac{P\{A_k | A_{k-1} = D_k = 0, V\}}{P\{A_k | A_{k-1} = D_k = 0, \bar{L}_{k-1}, V\}} \\ &\times \frac{P\{D_k = 0 | A_{k-1} = D_{k-1} = 0, V\}}{P\{D_k = 0 | A_{k-1} = D_{k-1} = 0, \bar{L}_{k-1}, V\}} \end{aligned} \quad (5)$$

The impact of inversely weighting by the denominator probabilities is to remove the association between exposure (A_t, D_t) and time-varying confounders at each time t (i.e. the impact is to eliminate time-varying confounders), while leaving the causal effect of interest unchanged. The numerator probabilities in (5) are allowed to be misspecified by the fact that model (3) is conditional on V .

Because standard software for PH regression does not allow to reweight the risk sets at each time, we do this by fitting a weighted pooled logistic regression model using

generalized estimating equations, treating each patient-day as an observation and using regression splines to fit the time effect (Hernán et al. 2000). Unbiasedness of the estimating equations requires use of the independence working correlation (Vansteelandt 2007). Note that, because we wish to assess the impact of acquiring infection in ICU on a given day, information regarding the end of the infection episode and regarding the infection status outside the ICU is irrelevant.

The weights (5) differ from the usual stabilized weights for marginal structural models (Hernán et al. 2000, 2001) in that they consider the joint treatment process given by infection and discharge at each time and do this only up to the artificial time S_t . Note that they involve the discharge process to account for the fact that, at each time t , those subjects who are still in the ICU (i.e. for whom we have information on the infection history) may form a selective subset of the study population. Note also that, by using generalized estimating equations to fit model (4), we account for the potentially strong correlation arising in the augmented dataset which may contain the same observations multiple times (corresponding to different values of S_t).

To deal with censoring of the survival status due to hospital discharge, we proceed under the additional assumption of sequentially ignorable censoring (Robins 1997a). For our data, this assumption states that among subjects with a given observed past $\bar{A}_{t_D}, \bar{D}_t, \bar{L}_{t_{D-}}, V$, where $t_D = \min(t, D - 1)$ and $t_{D-} = \min(t, D) - 1$, the censored and uncensored subjects at time t have the same survival time distribution; that is, $C \perp\!\!\!\perp T | \bar{A}_{t_D}, \bar{D}_t, \bar{L}_{t_{D-}}, V, T > t, C > t$ for each time t . At a given time t , this assumption could be reasonable for short term survival chances because we have available a large and detailed collection of prognostic factors for survival that also predict time of discharge from the ICU. However, for given t , it is questionable for the longer term because we lack data monitoring the health status of patients after leaving the ICU. In our study, the median length of stay in hospital after ICU discharge was 8 days (IQR 10, 5% percentile 0, 95% percentile 50).

Under the assumption of sequential ignorability, we can correct the above analysis for sequentially ignorable censoring by further weighting the contribution of each patient

to the risk set at time t by the stabilized weights

$$swc(t, \bar{A}_{t_D}, \bar{D}_t, \bar{C}_{t-1}, \bar{L}_{t_D}) = \prod_{k=1}^t \frac{P(C_k = 0 | \bar{A}_{k_D}, \bar{D}_k, C_{k-1} = 0, V)}{P(C_k = 0 | \bar{A}_{k_D}, \bar{D}_k, C_{k-1} = 0, \bar{L}_{k_{D-}}, V)} \quad (6)$$

where the numerator and denominator probabilities equal 1 when $D_k = 0$. Here, we implicitly assume that the event of hospital discharge does not causally affect survival. We conclude that fitting model (4) and weighting each patient's contribution to the risk set at time t by the product of (5) and (6) produces a consistent estimator for the causal rate ratio provided that the measured (time-dependent) covariates are sufficient to adjust for time-dependent confounding and censoring due to hospital discharge.

4 DATA ANALYSIS

We first consider the unadjusted time-dependent PH model

$$\lambda(t|\bar{A}_t) = \lambda_0(t) \exp(\beta_1 A_t)$$

To enhance comparability with later results, we fitted this model via unweighted pooled logistic regression using generalized estimating equations with regression splines for the time effect. The estimate of the hazard ratio of death comparing patients who acquired infection prior to time t and those who did not, was 1.89 (95 per cent confidence interval (CI) 1.32; 2.71). When adding baseline covariates (SAPS score and reasons for admission to the ICU) to the model, the estimated hazard ratio was no longer significant and equalled 1.37 (95 per cent CI 0.93; 2.04). A more detailed overview of various conventional analysis results for these data and a discussion of the biases incurred by standard regression methods can be found in Mertens et al. (2006a).

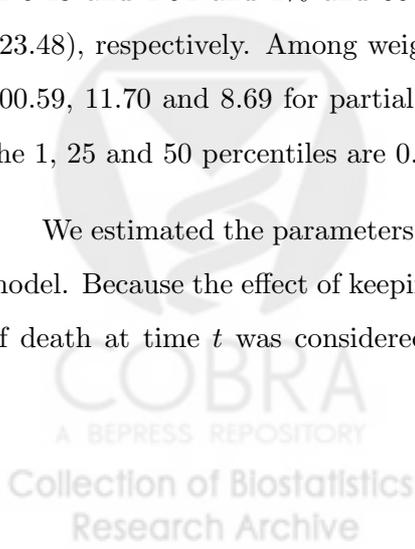
To adjust for time-dependent confounding, we extended our data set to include the variables S_t and A_t^* for each patient at each time t . Next, we calculated stabilized weights by means of 6 pooled logistic regression models for the numerator and denominator weights in (5) and (6). To build parsimonious models, we used the following conservative approach. In the first stage, all main effects were added and then sequentially removed if

non-significant at the 10% level (ignoring correlations between outcomes from the same patient). In the second stage, interaction terms considered to be plausible by clinicians were added if significant following the same criterion. Time-dependent information on exposure to invasive treatments was summarized in terms of the presence/absence of the treatment on each of the 2 previous days and by the total number of previous days on invasive treatments. Furthermore, splines were used to model the time effect in all models.

Using the estimated predicted values from these models we calculated the probability of each patient having his/her observed infection status up to time t , given baseline variables and then also given time-dependent variables \bar{L}_{t-1} . We fitted similar models for the probability of ICU and hospital discharge, the latter after also adjusting for the infection and ICU discharge history. To avoid unstable weights, we considered only infection regimes (s, a) with $3 \leq s \leq 11$. This implies that the estimated effect of infection on the hazard of death pertains only to infection regimes where infection is acquired during the first 11 days (starting from day 3). Note however that we included all observed person-days in the analysis. Further, to avoid unstable weights, we included baseline covariates in the marginal structural model as this enables the inclusion of baseline covariates in the numerator weights (and hence increases their stability).

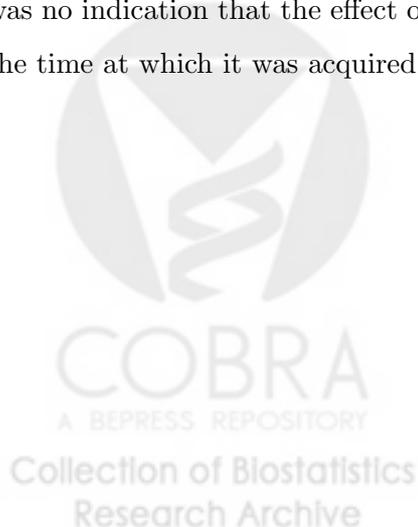
Figure 1 displays the distribution of the natural logarithm of the stabilized weights as a function of time. The boxes run from the first to the third quartile and the whiskers from the first to the 99th percentile. The overall distribution of the stabilized weights had a median and mean of 0.81 and 0.93, an interquartile range and standard deviation of 0.48 and 1.94 and 1% and 99% percentiles of 0.048 and 3.89 (min. 0.0039, max. 123.48), respectively. Among weights greater than 5, the 99, 75 and 50 percentiles are 100.59, 11.70 and 8.69 for partial exposure regimes. Among weights smaller than 0.2, the 1, 25 and 50 percentiles are 0.0066, 0.069 and 0.12 for partial exposure regimes.

We estimated the parameters of the MSM (4) by fitting a weighted GEE regression model. Because the effect of keeping the patient in the ICU up to time S_t on the hazard of death at time t was considered a nuisance, we modelled the effect of S_t in model



(4) using regression splines. Our causal estimate of the hazard ratio for infection was 2.74 (95 per cent conservative CI 1.48; 5.09). We conclude that under any regime in which patients stay in ICU for at least a given number of days s , the effect of acquiring infection on day s is to multiply the subsequent hazard of death by 2.74. Figure 2 shows estimated survival curves for the study population along with 95% confidence intervals, and predicted survival curves in the hypothetical scenario where all patients acquire infection at the third day of their stay in the ICU. It illustrates the severe estimated impact of ICU acquired infection on mortality. Confidence intervals were obtained using the robust standard error with independence working covariance matrix. By not taking into account the estimation of the weights, this yields an asymptotically conservative confidence interval for our causal parameters (Robins et al. 2000).

To examine the stability of this result to extreme weights, we additionally evaluated the effect of infection on mortality for infection regimes with $3 \leq s \leq s_{\max} = 7, 8, 9$ and 10. The weights are more stable for these analyses by the fact that the product in (5) runs over a smaller number of time points. The results are displayed in Table 1 and show that the effect size and significance stay the same despite the increasing stability of the weights. Finally, we performed an ad-hoc procedure whereby stabilized weights smaller than 0.2 or greater than 5 were truncated at 0.2 and 5, respectively. It yielded a hazard ratio of 2.50 (95 per cent conservative CI 1.45; 4.31), suggesting once more robustness to the extreme weights. Allowing for an interaction between infection status and the number of days since acquiring the infection, revealed that (on the hazard scale) the effect of acquiring infection on a given day s increases non-significantly with 2.8% (95% conservative CI -1.2%; 6.7%, P 0.17) per day since acquiring infection. Likewise, there was no indication that the effect of infection in ICU on the hazard of death depends on the time at which it was acquired (P-value 0.29).



5 DISCUSSION

The goal of this paper was to quantify the effect of acquiring NP in ICU on subsequent mortality. Because standard statistical analyses of data for ICU patients may be severely misleading due to the presence of time-varying confounders, and have shown controversial results in this context, we have proposed to use analyses of marginal structural models instead. Unlike many standard analyses, the latter take into account the time order in which infection, mortality data and time-dependent confounders were collected and succeed to correct for time-dependent confounders that lie on the causal path from NP to mortality. Inference for marginal structural models was however not directly applicable to our data because

(a) the infection status of patients was not well defined subsequent to ICU discharge, an event which lies itself on the causal path from infection to mortality.

(b) the required weights in the weighted estimating equations for standard marginal structural models would be highly unstable as there was a lengthy follow-up for several patients and because many collected time-dependent variables were extremely predictive of ICU discharge.

To accommodate these problems we have proposed to model mortality rates following partial infection paths. Patients under any such regime stay in the ICU up to a given time s , at which they may or may not acquire infection. We believe that the results obtained from such models are useful in the present context for various reasons. First, they enable us to assess the effect of acquiring infection on a given day s in the hypothetical situation where all patients stayed in the ICU for at least s days. Second, they solve the problem mentioned in (a) without fixing the discharge time after the event of infection; that is, without fixing variables that possibly lie on the causal path from infection to mortality. Third, the required weights in the weighted estimating equations are more stable under these models than under standard marginal structural models. This is because the fixed episode of the infection regime is now limited up to a given time s (rather than up to the study end) so that the weights merely involve

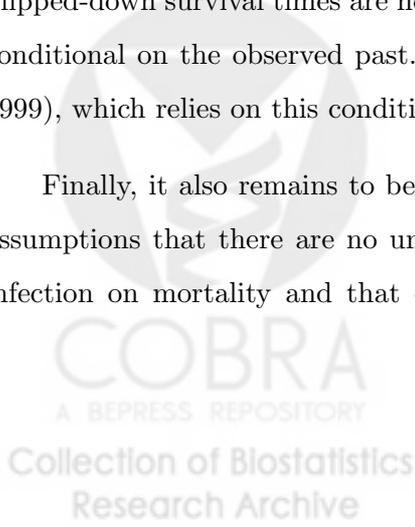
the first s time points and, hence, are less affected by lengthy follow-up with frequent infection measurements.

Our results are more generally useful for assessing the causal effect of a time-varying exposure in observational studies. First, they directly accommodate situations where exposures are not collected up to the time where outcomes are assessed. This may happen in settings where the mortality status of patients is assessed at the time of data analysis (i.e. later than end-of-follow-up) through death registers, or where each patient's treatment is closely monitored for only a limited time period. By not fixing treatment levels observed after this time period, the proposed models succeed to isolate the overall effect of treatment over the given period on outcome. Second, our results extrapolate less from the observed data by making fewer untestable assumptions. Specifically, to draw inference about infection regimes (s, a) with $s < s_{\max}$, the no-unmeasured-confounders assumption relates only to the hazard of infection and ICU discharge at time points $t < s_{\max}$. Third, they have the further advantage of yielding more stable weights in the weighted estimating equations and still yielding meaningful answers to the causal question. The former will be especially clear in studies where most patients have lengthy follow-up with frequent treatment measurements. Alternative proposals to deal with unstable weights have been made by Joffe et al. (2004) and by Yu and van der Laan (2006). Joffe et al. (2004) consider the effect of treatments received after the beginning of a moving partition. While this approach is clearly valuable in certain study settings, it does not allow to assess the impact of early infections for the long survivors and would not be successful in our setting because most infection events happened early on in the study. Yu and van der Laan (2006) use so-called doubly-robust estimators for the marginal structural model parameters. These are computationally more complex, but have the advantage of yielding CAN estimators for the causal parameters when either the model for the weights in the estimating equations holds, or some model for the counterfactual survival distribution. When reasonable models can be postulated for the latter, such doubly-robust estimators allow to truncate extreme weights in the weighted estimating equations and thus to achieve a better finite-sample performance. It remains to be seen how such doubly-robust estimators can be constructed and how they would

perform for the marginal structural models proposed in this article.

Alternatively, we could have chosen to assess the effect of avoiding infection among patients who acquired infection on a given day (with given prognostic factors). Such effect may be different from the effect identified in this article because infection might be more harmful for those who actually were infected. At the same time, it might have greater public health relevance since clinicians may be primarily interested in the effect of infection on mortality among those who actually were infected. Structural nested distribution models (Robins 1997a, Keiding et al. 1999) can be used to assess such effects. These are models for the effect of a change in infection status on survival time among patients with a given history of measured time-dependent confounders and infection. By describing causal effects within subpopulations with a given observed infection history, these models may tend to extrapolate less from the observed data and may yield more stable inferences because estimation does not involve inverse probability weighting. Because inference for these models is more complicated and does not allow the use of standard software, they have not frequently been used so far. We have therefore chosen to adopt marginal structural models in this article and plan to report on structural nested models elsewhere. Standard inference for the latter models would not be directly applicable by the same problem that the infection status is ill defined for patients who are discharged uninfected from the ICU, even if we would define those patients uninfected. Indeed, the blipped-down survival times in these models express what would be the survival time if there was no infection prior to ICU discharge, but the infection status after ICU discharge equalled the observed infection status. Because the latter may depend on the observed infection status prior to ICU discharge, these blipped-down survival times are not guaranteed to be independent of observed infection conditional on the observed past. Hence, G-estimation (Robins 1997a, Keiding et al. 1999), which relies on this conditional independence assumption, may fail.

Finally, it also remains to be seen how sensitive conclusions are to the untestable assumptions that there are no unmeasured time-varying confounders for the effect of infection on mortality and that censoring is sequentially ignorable. The former as-



sumption implies that among patients with prognostic factors $\bar{L}_{t-1}, \bar{A}_{t-1}, D_{t-1} = 0$, the causal effect of infection and ICU discharge is the same regardless of their infection and ICU discharge status at time t . This may not be entirely relevant because we anticipate the causal effect of infection to be greater among the infected and we may lack sufficient prognostic factors conditional on which this is no longer so. The assumption of sequentially ignorable censoring may also be questioned because the decision to discharge patients from hospital is intimately connected to their health status, about which no information was recorded after ICU discharge. A sensitivity analysis can and remains to be undertaken as in Brumback et al. (2004).

APPENDIX

Let us first assume there is no censoring due to hospital discharge. Under this assumption, we will construct (up to asymptotic equivalence) all unbiased estimating functions for β indexing model \mathcal{M} for the observed data $(T, D, \bar{A}_D, \bar{L}_D)$ defined by the law of the infection and discharge process under sequential randomization (SRA, i.e. the assumption of no unmeasured confounding for infection and ICU discharge):

$$f \{D_{D+1} = 1 | \bar{A}_D, \bar{D}_D = 0, \bar{L}_D, V\} \prod_{t=1}^D f \{A_t, D_t = 0 | \bar{A}_{t-1}, \bar{D}_{t-1} = 0, \bar{L}_{t-1}, V\} \quad (7)$$

where f is an unknown probability function, and by the (discrete-time) marginal multiplicative intensity model:

$$\lambda_{(s,a)}(t|V) = \lambda_0(t) \exp(\beta'W) \quad (8)$$

for $t \geq 0$, $a = 0, 1$, $s = 1, \dots, K$ for a given integer constant $K > 0$, where $W = W(a, s, t, V)$ is a known function of a, s, t and V . Note that model (3) is a special case of (8) with $\beta = (\beta_1, \beta_2, \beta_3)'$ and $W(a, s, t, V) = (sI(t \geq s), aI(t \geq s), V)$. Upon noting that the assumption of CAR-SRA equivalence (Robins, Rotnitzky and Scharfstein 1999) continues to hold when the length D of the infection period is random, the derivation is similar to the construction of CAN estimators for parameters indexing a conditional

mean model under CAR (i.e. coarsening at random). Application of Theorem 1.3 in van der Laan and Robins (2003) shows that, up to asymptotic equivalence, all CAN estimators of β under model \mathcal{M} can be obtained by solving an estimating equation based on estimating functions in the set $\{U\} + T_{SRA}$, where U is an arbitrary unbiased estimating function for β under this model and where T_{SRA} is the tangent space (Bickel et al. 1993) for the infinite-dimensional parameters indexing the infection and discharge process (7), which is assumed to satisfy the SRA assumption. A similar argument as in Theorem 1.2 (van der Laan and Robins 2003) shows that $T_{SRA} = T_{SRA,1} + T_{SRA,2}$ where

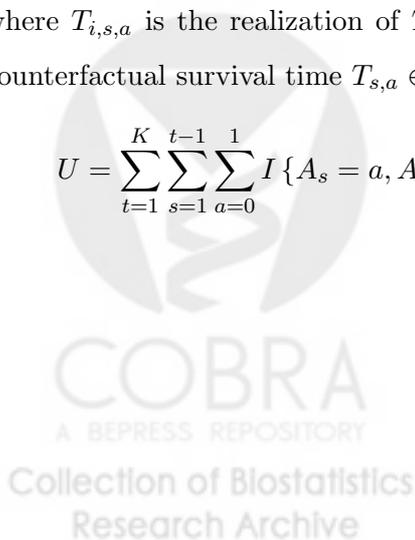
$$\begin{aligned}
 T_{SRA,1} &= \left\{ \sum_{s=1}^D Z(\bar{A}_s, \bar{D}_s, \bar{L}_{s-1}, V) \right. \\
 &\quad \left. - E[Z(\bar{A}_s, \bar{D}_s, \bar{L}_{s-1}, V) | \bar{A}_{s-1}, \bar{D}_{s-1}, \bar{L}_{s-1}, V] : Z \text{ arbitrary} \right\} \\
 T_{SRA,2} &= \left\{ Z(\bar{A}_D, \bar{D}_{D+1}, \bar{L}_D, V) \right. \\
 &\quad \left. - E[Z(\bar{A}_D, \bar{D}_{D+1}, \bar{L}_D, V) | \bar{A}_D, \bar{D}_D, \bar{L}_D, V] : Z \text{ arbitrary} \right\}
 \end{aligned}$$

To determine an unbiased estimating function U under model \mathcal{M} , let $U_{t,(s,a)}(T_{s,a}, V; \beta)$ be an unbiased estimating function for β in the full data model defined by the full data $(T_{s,a}, V)$, restriction (8) just for the given t and the given infection path (s, a) . Because (8) is a multiplicative intensity model, such estimating functions follow from standard results on such models. In particular, it may be the discrete-time partial likelihood score

$$\left\{ W(s, a, t, V) - \frac{\sum_{i=1}^n W(s, a, t, V_i) I(T_{i,s,a} \geq t) \exp(\beta' W(s, a, t, V_i))}{\sum_{i=1}^n I(T_{i,s,a} \geq t) \exp(\beta' W(s, a, t, V_i))} \right\} dN(t)$$

where $T_{i,s,a}$ is the realization of $T_{s,a}$ for the i th subject and $dN(t)$ indicates 1 if the counterfactual survival time $T_{s,a} \in]t - 1, t]$ for the considered subject. Define

$$U = \sum_{t=1}^K \sum_{s=1}^{t-1} \sum_{a=0}^1 I\{A_s = a, A_{s-1} = 0, D_s = 0\} s w_{s,a} U_{t,(s,a)}(T_{s,a}, V; \beta)$$



where we define

$$sw_{s,a}(\bar{L}_{s-1}) = \frac{P\{A_s = a, D_s = 0 | A_{s-1} = D_{s-1} = 0, V\}}{P\{A_s = a, D_s = 0 | A_{s-1} = D_{s-1} = 0, \bar{L}_{s-1}, V\}} \\ \times \prod_{k=1}^{s-1} \frac{P\{A_k = D_k = 0 | A_{k-1} = D_{k-1} = 0, V\}}{P\{A_k = D_k = 0 | A_{k-1} = D_{k-1} = 0, \bar{L}_{k-1}, V\}}$$

where we ignore the dependence of $sw_{s,a}$ on random variables for ease of notation. Then U is an unbiased estimating function in model \mathcal{M} . Indeed, first note that U is a function of the observed data because replacing $T_{s,a}$ by T yields the same full data function under the consistency assumption that we observe $T_{s,a} = T$ for subjects with $A_s = a, A_{s-1} = 0, D_s = 0$. Furthermore, for $s < t - 1$ and provided that the sequential randomization assumption holds,

$$\begin{aligned} & E [I\{A_s = a, A_{s-1} = 0, D_s = 0\} sw_{s,a}(\bar{L}_{s-1}) U_{t,(s,a)}(T_{s,a}, V; \beta)] \\ = & E [E [I\{A_s = a, D_s = 0\} sw_{s,a}(\bar{L}_{s-1}) | A_{s-1} = D_{s-1} = 0, \bar{L}_{s-1}, T_{s,a}] \\ & \times I\{A_{s-1} = D_{s-1} = 0\} U_{t,(s,a)}(T_{s,a}, V; \beta)] \\ = & E [E [I\{A_s = a, D_s = 0\} sw_{s,a}(\bar{L}_{s-1}) | A_{s-1} = D_{s-1} = 0, \bar{L}_{s-1}] \\ & \times I\{A_{s-1} = D_{s-1} = 0\} U_{t,(s,a)}(T_{s,a}, V; \beta)] \\ = & E [P\{A_s = a, D_s = 0 | A_{s-1} = D_{s-1} = 0, V\} \\ & I\{A_{s-1} = D_{s-1} = 0\} sw_{s-1,0}(\bar{L}_{s-2}) U_{t,(s,a)}(T_{s,a}, V; \beta)] \\ = & \dots = E \left[U_{t,(s,a)}(T_{s,a}, V; \beta) P\{A_s = a, D_s = 0 | A_{s-1} = D_{s-1} = 0, V\} \right. \\ & \left. \prod_{k=1}^{s-1} P\{A_k = D_k = 0 | A_{k-1} = D_{k-1} = 0, V\} \right] = 0 \end{aligned}$$

where the last equality is true because the estimating functions $U_{t,(s,a)}(T_{s,a}, V; \beta)$ are conditionally unbiased given V . We conclude that U is an unbiased estimating function. The result now follows by noting that solving an estimating equation with estimating function U is mathematically equivalent to fitting the time-dependent (discrete-time) multiplicative intensity model

$$\lambda(t|V) = \lambda_0(t) \exp(\beta' W(A_t^*, S_t, t, V))$$

where $\lambda_0(t)$ is an unknown baseline hazard, where $A_t^*, S_t, t > 0$ are defined as in Section 3.3, and where the risk set at each time is weighted by the weights (5).

It further follows from Theorem 1.2 in van der Laan and Robins (2003) that for given estimating function U , the choices $Z(\bar{A}_s, \bar{D}_s, \bar{L}_{s-1}, V) = E(U|\bar{A}_s, \bar{D}_s, \bar{L}_{s-1}, V)$ and $Z(\bar{A}_D, \bar{D}_{D+1}, \bar{L}_D, V) = E(U|\bar{A}_D, \bar{D}_{D+1}, \bar{L}_D, V)$ are optimal in the sense that they yield an efficient estimator of β under model \mathcal{M} in the class of estimators obtained by solving estimating equations in the class $\{U\} + T_{SRA}$ for given U . Finally, the above methods are easily adapted to handle sequentially ignorable censoring following the lines of Robins et al. (1994) and to account for estimation of the parameters indexing the infection and discharge process (7). It also follows from Theorem 2.4 in van der Laan and Robins (2002) that we obtain an asymptotically conservative confidence interval for our causal parameters by not taking into account estimation of the weights, provided that the unknown parameters in the models for the weights are efficiently estimated.

REFERENCES

- Andersen, P.K., and Keiding, N. (2002), Multi-state models for event history analysis. *Statistical Methods in Medical Research* **11**, 91-115.
- Bickel, P.J., Klaassen, C.A.J., Ritov, Y., and Wellner, J.A. (1993), *Efficient and Adaptive Estimation for Semiparametric Models*. Springer-Verlag: New-York.
- Brumback, B.A., Hernán, M.A., Haneuse, S.J.P.A., and Robins, J.M. (2004), “Sensitivity analyses for unmeasured confounding assuming a marginal structural model for repeated measures,” *Statistics in Medicine*, **23**, 749-767.
- Bryan, J., Yu, Z. and van der Laan, M.J. (2004), “Analysis of longitudinal marginal structural models,” *Biostatistics*, **5**, 361-380.
- Carlet, J. (2001), “Dying from or with a Nosocomial Pneumonia in the intensive care unit?,” *Critical Care Medicine*, **29**, 2392-2394.

- Gaynes, R.P. (1997), "Surveillance of Nosocomial Infections: A fundamental ingredient for quality," *Infection Control and Hospital Epidemiology*, 18, 475-478.
- Haley, R.W., Culver, D.H., White, J.W., Morgan, W.M., Emori, T.G., Munn, V.P., and Hooton, T.M. (1985), "The efficacy of infection surveillance and control programs in preventing nosocomial infections in US hospitals," *American Journal of Epidemiology*, 121, 182-205.
- Hernán, M., Brumback, B., and Robins, J.M. (2000), "Marginal structural models to estimate the causal effect of zidovudine on the survival of HIV-positive men," *Epidemiology*, 11, 561-570.
- Hernán, M., Brumback, B., and Robins, J.M. (2001), "Marginal structural models to estimate the joint causal effect of nonrandomized treatments," *Journal of the American Statistical Association*, 96, 440-448.
- Joffe, M., Knauss, J., Robinson, B. et al. (2004), "Lagged partially marginal structural models for controlling confounding," *American Journal of Epidemiology*, 159, S66-S66 Suppl. S.
- Kalbfleisch, J.D., and Prentice, R.L. (1980), *The Statistical Analysis of Failure Time Data*, New York: Wiley.
- Keiding, N., Filiberti, M., Esbjerg, S., Robins, J.M., Jacobsen, N. (1999), The graft-versus leukemia affect after bone marrow transplantation: A case study using structural nested failure time models. *Biometrics*, **55**, 23-29.
- van der Laan, M.J., and Robins, J.M. (2003), *Unified Methods for Censored Longitudinal Data and Causality*, New-York: Springer-Verlag.
- Le Gall, J.R., Lemeshow, S., and Soulnier, F. (1993), "A new simplified acute physiology score (SAPS II) based on a European / North American Multicenter Study," *Journal of the American Medical Association*, 270, 2957-2963.

- Mertens, K., Suetens, C., Vansteelandt, S., and Goetghebeur E. (2006a), "Various statistical regression models to estimate the effect of ICU-acquired Nosocomial Pneumonia on mortality," Technical Report.
- Mertens, K., Vansteelandt, S., Suetens, C., and Goetghebeur E. (2006b), "Methodological considerations when estimating the attributable effect of ICU-acquired Nosocomial Pneumonia on mortality," Technical Report.
- Robins, J.M. (1986), "A new approach to causal inference in mortality studies with sustained exposure periods - Application to control of the healthy worker survivor effect," *Mathematical Modelling*, 7, 1393-1512.
- Robins, J.M. (1997a), "Causal Inference from Complex Longitudinal Data," in *Latent Variable Modeling and Applications to Causality*, Lecture Notes in Statistics (120), Ed. M. Berkane, New-York: Springer Verlag, pp. 69-117.
- Robins, J.M. (1997b), "Structural nested failure time models in survival analysis," in *The Encyclopedia of Biostatistics*, Eds. P. Armitage and T. Colton. Chichester: John Wiley & Sons, pp. 4372-4389.
- Robins, J.M., Hernán, M.A., and Brumback, B. (2000), "Marginal Structural Models and Causal Inference in Epidemiology," *Epidemiology*, 11, 550-560.
- Robins, J., Hernán, MA, and Siebert, U.(2003), "Effects of multiple interventions," in *Population Health Metrics*, Ed by Murray C.J., Ezzati M., Lopez A.D., Rodgers A., Vander Hoorn S. Volume 2: 2191-2230.
- Suetens, C., Jans, B., and Carsauw, H. (1999), "Nosocomial infections in intensive care: results from the Belgian national surveillance, 1996-1998," *Archives of Public Health*, 57, 221-231.
- Suetens, C., Savey, A., Lepape, A., Morales, I., Carlet, J. and Fabry, J. (2003), "Surveillance of nosocomiale infections in intensive care units: towards a consensual approach in Europe," *Réanimation*, 12, 205-213.

Rubin, D.B. (1978), “Bayesian inference for causal effects: the role of randomization,” *Annals of Statistics*, 6, 34-58.

Vansteelandt, S. (2007), “On confounding, prediction and efficiency in the analysis of clustered and longitudinal data,” *The Scandinavian Journal of Statistics*. First published on 31 January 2007, doi: 10.1111/j.1467-9469.2006.00555.x.

Vincent, J.L. (2003), “Nosocomial Infections in adult intensive care units (review),” *Lancet*, 361, 2068-2077.

Yu, Z., and van der Laan, M.J. (2006), “Double Robust Estimation in Longitudinal Marginal Structural Models,” *Journal of Statistical Planning and Inference*, 136, 1061-1089.

Table 1: Distribution of the stabilized weights (1% and 99% percentiles, minimum and maximum), hazard ratios (HR) and 95% confidence intervals in marginal structural models for partial infection paths with $3 \leq s \leq s_{\max}$.

s_{\max}	1% perc.	99% perc.	min	max	HR	95% CI
7	0.065	3.50	0.0071	25.94	2.75	1.48; 5.12
8	0.064	3.49	0.0061	35.91	2.66	1.43; 4.94
9	0.060	3.51	0.0053	54.14	2.70	1.46; 5.00
10	0.054	3.72	0.0046	81.88	2.71	1.47; 5.01
11	0.048	3.89	0.0039	123.48	2.74	1.48; 5.09



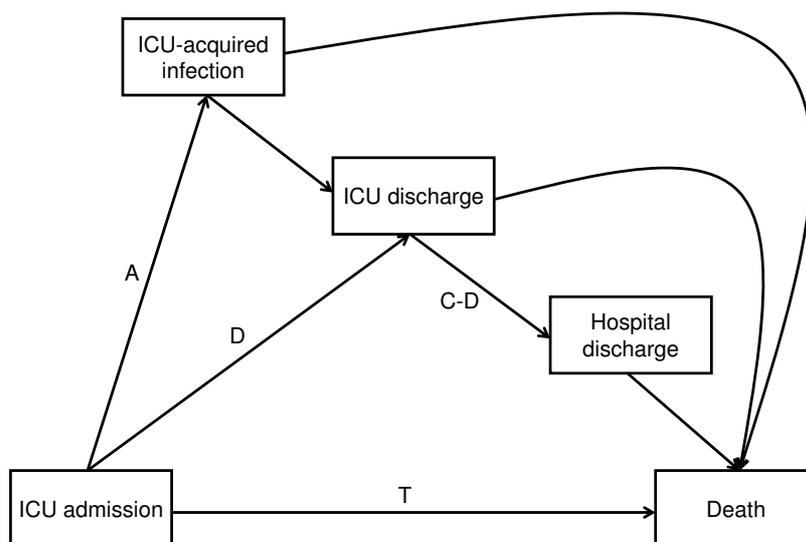


Figure 1: Multi-state model: directed arrows show the possible transitions from one state to another.

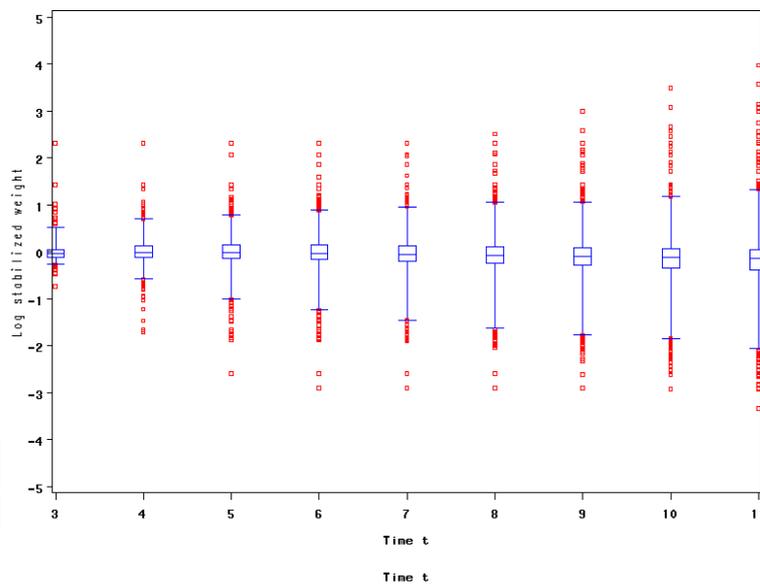


Figure 2: Boxplots of the natural logarithm of the stabilized weights in function of time t .

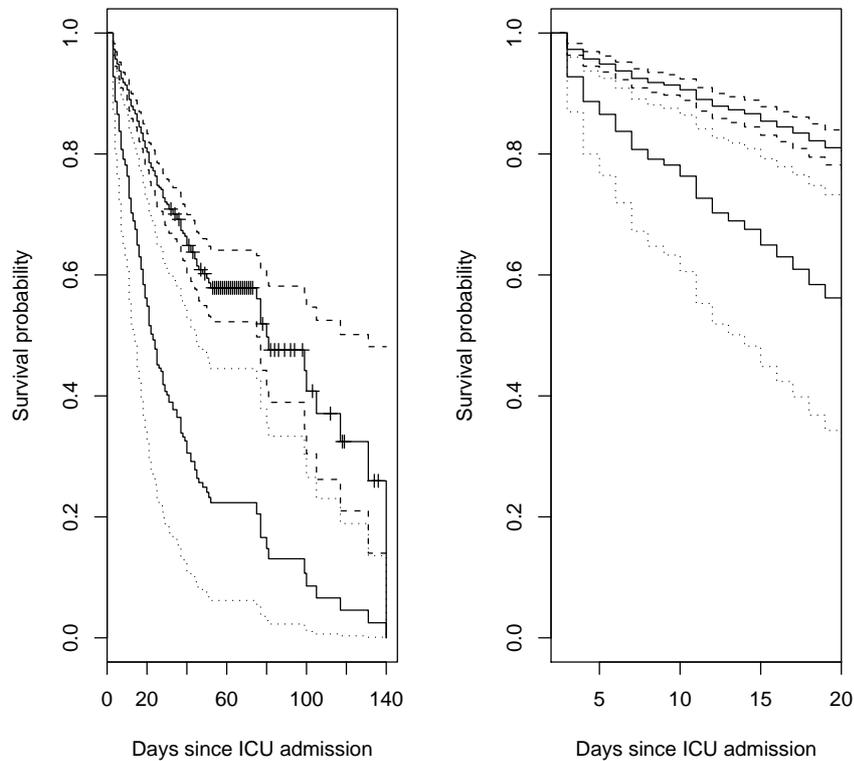


Figure 3: Marginal survival curve (directly estimated from the observed data) with 95% confidence intervals (top 3 lines) and predicted survival curves following immediate infection (based on the marginal structural model) with approximate 95% confidence intervals (bottom 3 lines) (the bottom 95% confidence interval acknowledges imprecision on the estimated causal effect, but ignores imprecision on the estimated survival curve). Left: from 3 to 140 days after ICU admission; Right: from 3 to 20 days after ICU admission.