

Modeling heterogeneity via Multi-state models

Jorit Delen
Mathematics and Applications
bachelor thesis

Supervisor:
Cristian Spitoni
15 january 2018



Utrecht University

Abstract

In this thesis a dataset of 271 intensive-care patients with automatic respiratory distress syndrome that are CMV seropositive is analyzed. The data, obtained from two different hospitals, is analyzed using multi-state models. The first thing we calculated was if there was heterogeneity on the hospital level when it comes to dying and the effect of covariates. To do this we used a frailty-cox model. We found a frailty variance of $1.98 * 10^{-16}$, which implies there is no noticeable heterogeneity on hospital level. After this we used a different multi-state model to calculate the population attributable fraction of CMV reactivation on dying, and found a value of about 0.23, which implies a significant effect.

Contents

1	Introduction	3
2	Theoretical background	3
2.1	Survival analysis	3
2.1.1	Markov assumption	3
2.1.2	Right-censoring	4
2.1.3	The simplest multi-state model	4
2.2	Competing risks	5
2.3	Multi-state models	6
2.4	Cox model	8
2.5	Frailty model	9
2.6	PAF	10
3	Description of data	10
4	Results	11
4.1	Frailty and covariates	11
4.2	PAF of CMV reactivation	14
5	Discussion	16
A	transition probabilities	18

1 Introduction

Multi-state models are a valuable tool of mathematicians and biologists to analyze and give a mathematical meaning to data obtained from clinical trials and observational studies. An example of a way in which multi-state models can be used is making predictions for patients in similar situations based on collected data. Another use of this system is estimating the effect covariates such as age, weight etc. can have on the probability of experiencing an event, for instance dying.

2 Theoretical background

In this section we will explain the theoretical background that is required to understand the rest of this thesis. We shall start with giving an introduction, explaining the Markov assumption and censoring. After this we will begin with one of the simplest multi-state models possible, after which we will move on to competing endpoints models, and finally to general multi-state models. In that section we shall also talk about transition probabilities and the Aalen-Johansen estimator. After that we will explain the Cox-model, a way of calculating the influence of covariates. After this we will give an introduction to frailties and give a short explanation of the *Population Attributable Fraction (PAF)*. Readers already aware of the theory behind multi-state models and the Cox-model are advised to skip to section 2.5.

2.1 Survival analysis

A multi-state model is a stochastic process $(X(t))$ with a finite state space $S = \{1, \dots, N\}$. We can use this to model a process in which a subject can be in one of a finite number of states, and can make transitions between those states. In a process with N different states, we say that if $X(t) = k$ where $k \in 1, 2, \dots, N$, then the subject is in state k at time t . A *transition into state n* is made if $\{X_t = n | X_{t-} \neq n\}$, where $t-$ implies the time just before time t . A transition between states is often referred to as a *event*, and the time at which this happens is called the event-time and is often denoted by T . States from which a transition into another state is possible are called *transient*, while states that can not be exited once entered are called *absorbing* [1].

2.1.1 Markov assumption

For now and in the rest of this thesis, we shall make the assumption that this process can be described by a Markov process: we define the transition probability from state h at time s to state j at time t as

$$p_{hj}(s, t) = P(X(t) = j | X(s) = h, H_{s-}) \quad (1)$$

In this equation H_{s-} represent the history of the process until just before time s . Then the Markov assumption implies:

$$P(X(t) = j | X(s) = h, H_{s-}) = P(X(t) = j | X(s) = h) \quad (2)$$

So the history has no influence on the transition probability [2].

2.1.2 Right-censoring

It is possible that the data collected is be incomplete. In a medical survey, this can for instance happen if the study in which the data was collected has a set end date, and at that point not all the subjects have had a event yet. This is called *right-censoring*. In this case there will be a second value called the *censoring time* C . If $C < T$ the event will not be observed and the event will be called censored. It is often the case with censoring that if the event gets censored, the real event time T is unknown, and all that is known about it is that it is larger than the censoring time C .

2.1.3 The simplest multi-state model

The simplest type of multi-state model is given in figure 1.




Figure 1: The simplest multi-state model with only 2 states. The arrow represents a possible transition from state 1 to state 2

In this figure state 1 is a transient state, while state 2 is an absorbing state. A example of a multi-state model like this is an alive-dead model, where state 1 represents "alive", while state 2 represents "death", and the transition represents dying.

The *hazard function* or *hazard rate* $\lambda(t)$ is given by the chance to have the event of interest at time t given that the event has not yet occurred before time t . This hazard function is defined as

$$\lambda(t) = \lim_{\Delta t \searrow 0} \frac{P(T \in [t, t + \Delta t) | T \geq t)}{\Delta t}. \quad (3)$$

The *cumulative hazard* is given by:

$$\Lambda(t) = \int_0^t \lambda(u) du \quad (4)$$

Both the hazard function and the cumulative hazard can be estimated from observations of the event time T . It is important to take into consideration that the estimator still works even in the presence of censored events. One of these estimators is called the *Nelson-Aalen estimator*, which is defined as follows:

Let, for a population of D individuals: $0 < t_1 < t_2 < \dots < t_D$ be the ordered event times, and let d_i be the number of events that happened at t_i . Finally, let n_i be the amount of people that are just before time t_i at risk of having a event. Than an estimator of the hazard function is given by:

$$\hat{\lambda}(t_j) = \frac{d_j}{n_j} \quad (5)$$

And the Nelson-Aalen estimator of the cumulative hazard is given by:

$$\hat{\Lambda}(t) = \sum_{t_j \leq t} \frac{d_j}{n_j} \quad (6)$$

The variance of this estimator can also be estimated in the following way:

$$\hat{\sigma}^2(t) = \sum_{t_j \leq t} \frac{d_j}{n_j^2} \quad (7)$$

Combining these estimators can give us a 95% confidence interval for the cumulative hazard, given by:

$$\hat{\Lambda}_k(t) \pm \hat{\sigma}_k^2(t) * 1.96 \quad (8)$$

2.2 Competing risks

A slightly more difficult version of the model in figure 1 is called the *competing risks* model. In this type of model there are several possible events that can happen from the initial state. An example of this type of model is given in figure 2. In this example there are three possible events that can happen from state 0, those being the transitions into state 1, 2, or 3. For every subject there are two possible situations, either only one of the possible events will happen, or the observation will be censored.

With the competing risks model, it is quite often useful to define the *cause specific* hazard, which is the chance of having a transition to one specific end




Figure 2: A model with multiple competing endpoints, all endpoints are absorbing

state given that no events have happened yet. In a model with initial state 0 and competing endpoints $\{1, 2, \dots, N\}$, with event-time T and event-cause D we define the cause specific hazard as:

$$\lambda_k(t) = \lim_{\Delta t \searrow 0} \frac{P(T \in [t, t + \Delta t), D = k | T \geq t)}{\Delta t}, k \in \{1, 2, \dots, N\} \quad (9)$$

The Nelson-Aalen estimator can be calculated almost identical to how it was done in equation 5, but here we define d_{ik} as the number of events of type k at time t_i . Then the Nelson-Aalen estimators are given by:

$$\hat{\Lambda}_k(t) = \sum_{t_j \leq t} \frac{d_{jk}}{n_j}, \quad k \in \{1, 2, \dots, N\} \quad (10)$$

And the variances by:

$$\widehat{\sigma}_k^2(t) = \sum_{t_j \leq t} \frac{d_{jk}}{n_j^2} \quad (11)$$

2.3 Multi-state models

Multi-state models can be different from competing risks models in that, depending on the model, it is possible to make a transition from the initial state into a different transient state, so that at a later point another event can happen and a individual can experience multiple events. A example of a multi-state model that is not a competing risks model is given in figure 3




Figure 3: A multi-state model with two transient and two absorbing states.

In this example states 0 and 2 are transient, while state 1 and 3 are absorbing. It is possible for someone in state 0 to first make a transition to state 2, and then at a later time make a transition from state 2 to state 3, something that was not possible in the competing risks example.

Once again, we will make a slight adaptation to the hazard function and to the Nelson-Aalen estimator to make it fit with the now slightly more complicated situation. The cause specific hazard function for a transition from state k to state l is given by:

$$\lambda_{kl}(t) = \lim_{\Delta t \searrow 0} \frac{P(T \in [t, t + \Delta t), X_{t+\Delta t} = l | T \geq t, X_{t-} = k)}{\Delta t}, \quad k, l \in \{0, 1, \dots, N\} \quad (12)$$

For the Nelson-Aalen estimator for the transition from state k to state l , we now define d_{ikl} as the number of transitions from state k to state l at time t_i , and we define n_{ik} as the amount of people in state k just before time t_i . Then the Nelson-Aalen estimators are given by:

$$\hat{\Lambda}_{kl}(t) = \sum_{t_j \leq t} \frac{d_{jkl}}{n_{jk}}, \quad k, l \in \{1, 2, \dots, N\} \quad (13)$$

And the variances by:

$$\widehat{\sigma}_{kl}^2(t) = \sum_{t_j \leq t} \frac{d_{jkl}}{n_{jk}^2} \quad (14)$$

Another important estimator is the Aalen-Johansen estimator, which can be defined as follows:

We define the matrix

$$\mathbf{P}(s, t) := (P_{lj}(s, t))_{l, j}, l, j \in \{0, 1, 2, \dots, J\}$$

where each element represents the chance for any random subject to be in state j at time t , given that that subject was in state l at time s .

We also define the matrix

$$\mathbf{\Lambda}(t) := (\Lambda_{lj}(t))_{l, j}, l, j \in \{0, 1, 2, \dots, J\}$$

where each element represents the cumulative transition hazard from state l to state j at time t .

The final thing we will define for now is $d\mathbf{\Lambda}(u)$, which is given element-wise by:

$$d(\Lambda_{lj}(u))_{l, j} = P(T \in [t, t + \Delta t), X_{t+\Delta t} = l | T \geq t, X_{t-} = k) \quad \text{if} \quad l \neq j$$

$$d(\Lambda_{ll}(u)) = - \sum_{j=0, j \neq l}^J d(\Lambda_{lj}(t)) \quad \text{if} \quad l = j$$

An estimator of the matrix \mathbf{P} is given by the Aalen-Johansen estimator. This estimator is given by the product over all event times of the sum of the identity matrix plus $d\mathbf{\Lambda}(u)$:

$$\hat{\mathbf{P}}(s, t) = \prod_{t_j \in (s, t]} (\mathbf{I} + d\hat{\mathbf{\Lambda}}(t_j)) \quad (15)$$

2.4 Cox model

Sometimes, it could prove useful to study the effects of other covariates on the hazard function for certain transitions, for example studying the effect of age or gender on the probability of dying in a clinical test. An example of a model to do this with is the *Cox proportional hazard model*, where the hazard function for transition from state l to state k at time t for subject i with covariate vector \mathbf{X}_i is given as:

$$\lambda^{kl}(t|\mathbf{X}_i) = \lambda_0^{kl}(t) \exp(\beta^T \mathbf{X}_i) \quad (16)$$

In this equation, λ_0^{kl} is the baseline hazard function for transition from state k to state l . This is the same for all subjects in the study and represents an individual for who all covariate values are zero.

The cumulative hazard for transition from state l to state k at time t for subject i with covariate vector \mathbf{X}_i is given similarly, with:

$$\Lambda^{kl}(t|\mathbf{X}_i) = \Lambda_0^{kl}(t)\exp(\beta^T \mathbf{X}_i) \quad (17)$$

2.5 Frailty model

Almost always, there is variation in treatment outcome between patients. A way of reducing this variation is to incorporate patient characteristics in the model, by for example using a Cox model as described in the previous section to incorporate patient covariates. However, often not all of the variation can be explained by this, and there are effects that can not be captured by observable covariates. A way of reducing this variation further is by incorporating *frailty*, in this context otherwise known as heterogeneity, in the model, to introduce random effects into the model. The variance of this random term can then be used to determine if there is heterogeneity in the model and the order of magnitude.

In medical studies from which the data comes from more than one center, we can consider heterogeneity on a center level. We will use a Cox multi-state model with a random effect, otherwise known as a frailty model, to study this.

The hazard function of a frailty model is defined as follows: for a transition from state k to state l , for subject i with covariate vector \mathbf{X}_i and frailty W_i :

$$\lambda^{kl}(t|\mathbf{X}_i, W_i) = W_i \lambda_0^{kl}(t) \exp(\beta^T \mathbf{X}_i) \quad (18)$$

A choice that is most often used as the frailty term is a random variable picked from a gamma distribution with shape parameter $\frac{1}{\theta}$, and scale parameter θ . This has as a result that:

$$\mathbf{E}(W_i) = 1$$

and:

$$\text{Var}(W_i) = \theta$$

As a result of this, we can make a lot of inferences about the amount of heterogeneity based on the calculated value of θ . If θ is significantly different from 0, it implies there is heterogeneity in some way. The larger this value is, the more significant [3].

2.6 PAF

The *population-attributable fraction*, or PAF for short, is a useful tool in researching the effects that intermediate events can have on the final outcome of a model. An example of an effect that can be researched this way is the effect that diseases acquired on the ICU can have on the chance of dying in the hospital (this is the example we will study later in this paper). We will define two events, event D will be the event of interest, and event E will be the event of which we want to know what the effect on event D is. If we define $P(D|t)$ as the chance of the event of interest having already taken place on time t , and $P(\bar{D}|t)$ as the chance of the event of interest not having taken place before time t . With the same type of notation, we define $P(E|t)$ as the chance of the other event having already taken place on time t , and $P(\bar{E}|t)$ as the chance of the other event not having taken place before time t . Then the PAF is given by:

$$PAF(t) = 1 - \frac{P(D|\bar{E}, t)}{P(D|t)} \quad (19)$$

The interpretation of this is as follows: the PAF gives the fraction of the events of interest D that would not have taken place had the interfering event E not existed. Back to our example of hospital data, the PAF gives the fractions of deaths that would not have occurred was it not for diseases picked up in the ICU.

3 Description of data

The dataset that was studied was an adaptation of a dataset collected from 2011 to 2013 in two tertiary intensive care units in the Netherlands, collected by the MARS consortium. The original data described 399 patients with acute respiratory distress syndrome (ARDS). The reason the data was adapted was to study the effect of Cytomegalovirus(CMV) reactivation in patients with ARDS, so in the adapted dataset only the 271 patients that were CMV seropositive are included [4]. We shall use this data to study the heterogeneity between the two centers, and also the effect CMV reactivation has on the probability of death using the PAF.

Of the 271 patients, 52 died, 209 were discharged, and 10 were censored. The data came from two hospitals, labeled hospital 0 and hospital 1. The next table shows these 271 patient, but divided per hospital. The data from this table will be used in section 4.1.

Hospital	Total patients	Deceased	Discharged	censored
0	179	32	143	4
1	92	20	66	6

In section 4.2 we shall look at the effects of CMV reactivation. the next table once again shows the same patients, but divided over whether or not CMV

reactivation took place. Nine subjects entered the study with CMV already reactivated.

CMV Reactivation	Total patients	Deceased	Discharged	censored
No	204	29	170	5
Yes	67	23	39	5

4 Results

In this section we shall present the results of data analysis that has been done on the data that has been described in the previous section. The analysis has been done using the software environment R. The presentation of the results of the analysis shall be split in two parts, both of which use their own multi-state model. In section 4.1 we will research the frailty variance of the chance of dying from the two hospitals that the data came from, and whether there is heterogeneity at the level of those hospitals for this event. To do this we shall use a competing risks frailty Cox model as described in section 2.2, section 2.4, and section 2.5. We shall also look at some of the covariates that were available with the original data, and the effect they have on the probability of dying. In section 4.2 we shall look at the PAF of CMV reactivation in the data set, as described in section 2.6. To do this we shall use a multi-state model that is a slightly more complicated version of the competing risks model, with CMV reactivation as an intermediate effect.

4.1 Frailty and covariates

In this section we shall calculate the heterogeneity on hospital level of the data set for the probability of dying. To do this we shall use the competing risks frailty Cox model shown in figure 4.




Figure 4: The competing risks model with which we shall calculate the frailty factor of the hospitals

As mentioned in section 3, the data comes from two different hospitals, referred to as "hospital 0", and "hospital 1". We shall look at the transition from state 0 to state 1, which is the transition that represents dying.

The covariates we shall be looking at are "APACHE IV score", which is a system to measure health problems. A higher score implies more problems. Other covariates are "gender", "age", "BMI", "SOFA", which is a method of calculating organ failure, where a higher score implies more failure. The final covariate is "diabetes". Except for SOFA which was taken on the day the patient entered the study, all the covariate values are taken on the day an event happened. However before we look at the calculated variance of the frailty of this model, we shall first compare the Nelson-Aalen estimators of the transition for two subsets of the data, one containing all the subject from hospital 0 and the other containing all subjects from hospital 1. If these are significantly different, it could imply that the variance of the frailty will be significantly larger than zero, and there will be a large heterogeneity. These estimators and their 95% confidence intervals are displayed in figure 5.




Figure 5: The Nelson-Aalen estimators for $\Lambda_{01}(t)$ for both subsets of the data, given with 95% confidence interval.

We see that there does not appear to be a large difference between the subsets, which implies that the variance in the heterogeneity on center level is probably not very different from zero.

As explained in section 2.5, we shall use a frailty taken from a gamma distribution with shape parameter $\frac{1}{\theta}$, and rate parameter θ , so that the expected value is 1 with variance θ . After doing the calculations, we came to the following value for θ , and the following values and 95% confidence intervals for the chosen covariates:

$$\theta = 1.98 * 10^{-16} \quad (20)$$

Covariate	Value	Confidence interval
APACHE IV score	1.009	(1.00 , 1.02)
Gender	1.250	(0.69 , 2.27)
Age	1.051	(1.03 , 1.08)
BMI	1.018	(0.97, 1.07)
SOFA	1.069	(0.96, 1.19)
Diabetes	0.545	(0.22, 1.32)

The value of the variance of the frailty term is in no way significantly different from zero, so as we expected from the Nelson-Aalen estimators, the heterogeneity between centers is pretty much nonexistent.

A covariate is considered significant if the confidence interval does not contain 1. The effect of the covariate is called significantly increasing if the entire confidence interval is larger than 1, and significantly decreasing if the confidence interval is smaller than 1. In the table of covariates we can see the following:

- While the value given in the table for APACHE IV score is barely significantly increasing with the confidence interval only slightly above 1, keep in mind that this represents an increase per point of APACHE IV score. Considering that APACHE IV is a scale that ranges from 0 to 240, a high score in this aspect does imply a higher chance of dying.
- Gender, represented here as a binary value (0,1), appears to have a severe impact on mortality of about 25%, however the confidence interval of this covariate is rather large, and contains both values far above and far below 1, so the covariate is not significant and no definitive answers about the effect can be given at this point.
- Age has a significant increasing effect per year on mortality, which because of the high range of ages of the subject, means that an older subject has a severely higher chance of dying when compared to a younger subject. This would make sense as the majority of people in the study are older than 25 years, above which an increase of age is often associated with more health related problems.
- BMI has a very small increasing effect per point, however it is not significant. This makes sense considering a larger BMI is often associated with health problems. While in this study BMI has a smaller range than for instance age and APACHE IV score, it is still a broad enough range such that a patient with a for the study high BMI has a much higher chance of dying than a patient with a lower BMI.
- SOFA has the second largest increasing effect on mortality of the covariates studied, however it is not significant. This would be logical as SOFA is

a rating system defined to monitor organ failure, and entering the study with a high score in this would logically increase the chance of dying. This is slightly offset by the fact that SOFA has a smaller range than covariates such as age or APACHE IV score, however due to the fact that the increase per point is quite large, a relatively high score in this still implies a severely increased chance of dying. The reason it is not significant might be explained by the sample size of the study, and possibly with a larger sample size it would be found that it causes a significant increase.

- Diabetes, which is like gender represented by a binary variable (0,1) shows some of the weaknesses of this system, as it appears to have a very large decreasing effect on mortality, while intuition would say that it should have a increasing effect. This is possibly caused by the fact that there were only very few people in the study with diabetes, so the confidence interval is rather wide and like gender, contains values of both far above and far below one, so it is not significant.

4.2 PAF of CMV reactivation

As mentioned in the introduction of this section, we shall use a different model than what has been used in previous sections. This model is given in figure 6.

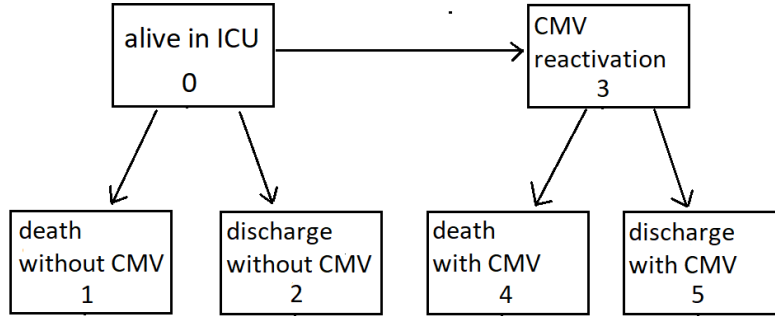


Figure 6: The model used to calculate the PAF of CMV reactivation

This time, it is not only important if death or discharge have happened, but also whether or not CMV reactivation has taken place before either of these events. This has been displayed with the addition of states 3, 4 and 5. Another difference with previous situations is that not every subject starts in state 0. As mentioned in section 3, there are nine subjects that start in state 3. Before we shall calculate the PAF of dying with CMV reactivated versus dying without it, we shall first compare the Nelson-Aalen estimators of the two transitions that imply dying, those being $\Lambda_{01}(t)$ and $\Lambda_{34}(t)$. These two estimators are displayed in figure 7.




Figure 7: The Nelson-Aalen estimators for $\Lambda_{01}(t)$ and $\Lambda_{34}(t)$, given with 95% confidence interval.

As opposed to the situation of the previous section, we do see a significant difference in the two estimators. While this does not give a definitive statement about what the PAF will be, it does make it more likely for the PAF to be significantly different from 0.

Now on to the calculations of the PAF. In this situation, the chances are given as follows:

$$P(D|\bar{E}, t) = \frac{\mathbf{P}_{01}(0, t)}{\mathbf{P}_{00}(0, t) + \mathbf{P}_{01}(0, t) + \mathbf{P}_{02}(0, t)} \quad (21)$$

This is the chance of dying while in state 0 in time $(0, t)$, divided by the chance of not having gone to state 3 in time $(0, t)$. Furthermore:

$$P(D|t) = \frac{262}{271}\mathbf{P}_{01}(0, t) + \frac{9}{271}\mathbf{P}_{34}(0, t) + \frac{262}{271}\mathbf{P}_{04}(0, t) \quad (22)$$

This can be interpreted as the chance of starting in either state 0 multiplied by the chance of dying from state 1 (either through state 3 or not), plus the chance of starting in state 3 multiplied by the chance of dying when starting from state 3. With these two equations we can finally give:

$$PAF(t) = 1 - \frac{(\frac{\mathbf{P}_{01}(0,t)}{\mathbf{P}_{00}(0,t)+\mathbf{P}_{01}(0,t)+\mathbf{P}_{02}(0,t)})}{\frac{262}{271}\mathbf{P}_{01}(0,t) + \frac{9}{271}\mathbf{P}_{34}(0,t) + \frac{262}{271}\mathbf{P}_{04}(0,t)} \quad (23)$$

Unfortunately we do not know the exact transition probabilities, but we can use the Aalen-Johansen estimators as described in section 2.3 to make an estimation of the PAF at every time t . Graphs of the Aalen-Johansen estimators of the transitions are shown in appendix A. The results of calculation the PAF at every relevant time are displayed in figure 8.




Figure 8: The estimation of the PAF for every time t . We cut of the first ten seconds as there had not been any transitions into state 4 yet so the PAF would not have made sense for those times.

As we expected, the PAF is at every time significantly larger than zero, appearing to go towards a value of about 0.23. This implies that a significant part of the deaths in this data-set are in fact caused by, or at least severely influenced by the reactivation of CMV.

5 Discussion

In this thesis we researched a dataset of 271 patients, coming from two hospitals. While we were unable to find sufficient evidence of heterogeneity with the method that we used, this does not have to imply that there is no heterogeneity

at all.

Due to there only being two hospitals in the study, the absence of heterogeneity in the center could have been coincidental. If data from more hospitals was added, maybe the variance term θ could have been significantly higher.

Another possible reason of why no heterogeneity was found was because only one method was used. While taking a gamma parameter as frailty term is the most popular option in this type of situation, it is by no means the only one. In future studies of this data other options could be researched, which might give different results.

Also important to note is that while in this study there was no heterogeneity found by this method, does not mean that using a frailty model in a similar situation is unnecessary, as unaccounted heterogeneity in a model can lead to bias.

References

- [1] J. Beyersmann, M. Schumacher and A. Allignol, *Competing Risks and Multistate Models with R*, Springer, New York, 2012
- [2] L. Meira-Machado, J. de Uña-Álvarez, C. Cadarso-Suárez and P. K. Andersen, *Multi-state models for the analysis of time-to-event data*, Statistical methods in medical research, **18**(2), 195-222, 2008
- [3] V. Rondeau, Y. Mazroui and J. R. Gonzalez, *frailtypack: An R Package for the Analysis of Correlated Survival Data with Frailty Models Using Penalized Likelihood Estimation or Parametrical Estimation*, Journal of Statistical Software, **47** (4), 2012
- [4] D. S. Y. Ong, C. Spitoni, *et al.*, *Cytomegalovirus reactivation and mortality in patients with acute respiratory distress syndrome*, Intensive care med, **42** 333-341, 2015
- [5] A. J. Rüten-Budde, *Investigating hospital heterogeneity with a competing risks frailty model*, Statistical Science for the Life and Behavioral Sciences thesis, Leiden university, Leiden, 2015
- [6] B. Liqueur, J. Timsit and V. Rondeau, *Investigating hospital heterogeneity with a multi state frailty model: application to nosocomial pneumonia disease in intensive care units*, BMC Medical Research Methodology, **12**(79), 2012

A transition probabilities

In this section we will display graphs of the transition probabilities of the model used in section 4.2.

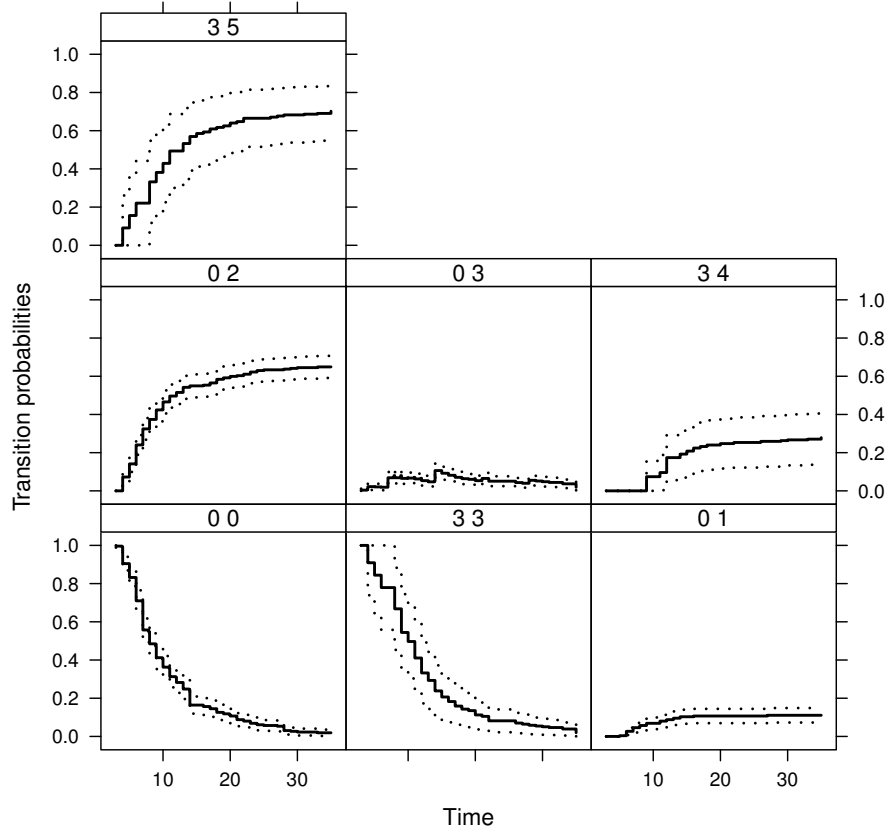


Figure 9: transition probabilities of the model used in section 4.2.