



Norwegian
Business School

This file was downloaded from BI Open, the institutional repository at BI Norwegian Business School <https://biopen.bi.no>.

It contains the accepted and peer reviewed manuscript to the article cited below. It may contain minor differences from the journal's pdf version.

Tierens, H., Dries, N., Smet, M., & Sels, L. (2021). Multiple-Membership Survival Analysis and Its Applications in Organizational Behavior and Management Research. *Organizational Research Methods*, 24(2), 412–442. <https://doi.org/10.1177/1094428119877452>

Copyright policy of *SAGE*, the publisher of this journal:

Authors “may post the accepted version of the article on their own personal website, their department’s website or the repository of their institution without any restrictions.”

<https://us.sagepub.com/en-us/nam/journal-author-archiving-policies-and-re-use>

**Multiple-Membership Survival Analysis and its Applications in Organizational
Behavior and Management Research**

ABSTRACT

Multilevel paradigms have permeated organizational research in recent years, greatly advancing our understanding of organizational behavior and management decisions. Despite the advancements made in multilevel modeling, taking into account complex hierarchical structures in data remains challenging. This is particularly the case for models used for predicting the occurrence and timing of events and decisions—often referred to as survival models. In this study, the authors construct a multilevel survival model that takes into account subjects being nested in multiple environments—known as a multiple-membership structure. Through this article, the authors provide a step-by-step guide to building a multiple-membership survival model, illustrating each step with an application on a real-life large scale archival dataset. Easy to use R code is provided for each model building step. The article concludes with an illustration of potential applications of the model to answer alternative research questions in the organizational behavior and management fields.

Keywords: multilevel, multiple membership, survival analysis, events, early-career turnover

Multiple-Membership Survival Analysis and its Applications in Organizational Behavior and Management Research

The analysis and prediction of events constitutes a major segment of organizational behavior and management research (Morita, Lee, & Mowday, 1989, 1993). In addition, most research that focuses on events as outcomes are not only interested in whether an event occurs, but also in when it occurs or how long it takes to occur (Morgeson, Mitchell, & Liu, 2015; Morita et al., 1993). Examining research questions that take into account the temporal dimension of an event (i.e., its timing and/or duration; Morgeson et al., 2015) typically requires the use of survival analysis techniques, also known as event-history analysis in sociology, duration analysis in economics, and reliability analysis in engineering (Morita et al., 1989, 1993). Survival analysis techniques can be applied to any phenomenon that can be conceptualized as the occurrence of an event on a timescale, whether it is calendar time or elapsed time since a previous event, or the start date of an observation period or study.

Examples of events (and their timing) that have been studied extensively in the field of organizational behavior are turnover in relation to different career stages (Allen, Hancock, Vardaman, & McKee, 2014; Cheramie, Sturman, & Walsh, 2007; Griffeth, Hom, & Gaertner, 2000; Kammeyer-Mueller, Wanberg, Rubenstein, & Song, 2013; Nyberg, 2010); time to re-employment (Wanberg, Hough, & Song, 2002); employee promotion order and speed (Barnett & Miner, 1992; Claussen, Grohsjean, Luger, & Probst, 2014); and occurrence and duration of psychological contract breaches (Eckerd, Boyer, Qi, Eckerd, & Hill, 2016; Griep, Vantilborgh, Baillien, & Pepermans, 2016; Solinger, Hofmans, Bal, & Jansen, 2016; Vantilborgh, Bidee, Pepermans, Griep, & Hofmans, 2016). More ‘macro’ examples from the management and strategy field are the speed at which organizations adapt to an industry versus disappear from it (Uzunca, 2018); the speed of success of new business ventures (Bakker & Shepherd, 2017;

Eisenhardt & Schoonhoven, 1990); and the speed of product innovation and new product adoption rates (Kessler & Chakrabarti, 1996; Schoonhoven, Eisenhardt, & Lyman, 1990).

In their work on Event System Theory, Morgeson and his colleagues (2015) state that, even though events can and do occur at different levels of an organization—from the micro employee level, to the team or department level, to the organizational level and even to the broader macro levels such as sectors and markets—, scholars have largely failed to understand the multilevel nature of both predictors and outcomes of the events. While Morgeson et al. (2015) focus primarily on events as antecedent to organizational behavior and change, they also acknowledge that events often trigger subsequent events. The latter might have a similar or distinct nature from the former; an unexpected but interesting job offer can cause an employee to quit their current job, for instance, or a co-worker quitting their job can cause another employee to quit their job as well—a phenomenon referred to as turnover contagion (Felps et al., 2009). Event System Theory, thus, allows for both the study of events as antecedents to behavioral reactions and as relevant outcomes driven by processes occurring at multiple levels of an organization. In the present paper, we propose that an integration of survival analysis and multilevel modeling is necessary to come to a better understanding of (the temporal nature of) events in organizational behavior and management research, as requested by Morgeson et al. (2015). Moreover, we also address the common issue of multiple membership—where lower-level units (e.g., individual respondents) are nested in different higher-level units (e.g., teams or organizations) simultaneously or sequentially—something that is more often than not ignored in both multilevel modeling and in survival analysis (Mathieu & Chen, 2011).

To the best of our knowledge, to date, no published organizational behavior and management research has applied multiple-membership survival analysis. In fact, the most promising advances so far have been made in other scientific fields, most notably veterinary medicine (e.g., Elghafghuf, Dufour, Reyher, Dohoo, & Stryhn, 2014a; Elghafghuf & Stryhn,

2017; Elghafghuf, Stryhn, & Waldner, 2014b). Such studies have analyzed, for instance, the (literal) survival of heifer cows ‘nested’ in several different barns and veterinary clinics (Duchateau & Janssen, 2008). Although the field of veterinary medicine is unallied to the disciplines of organizational behavior and management, applying the underlying models to our field is actually not that big of a leap—cows, then, become employees, and barns and clinics become departments and organizations. The goal of this article is to outline a step-by-step approach to the application of multiple-membership survival models in organizational behavior and management research. Since an excellent introduction to, and discussion of survival regression for management researchers published by Morita et al. (1989, 1993) is available for reference, throughout the rest of the paper we will assume some degree of familiarity with survival analysis on the part of the reader.

The intended contribution of this article therefore lies in demonstrating the applicability of multiple-membership survival analysis techniques to our field. To encourage broad readership, we build our model starting from a count regression model¹ (Blevins, Tsang, & Spain, 2015), which we assume is a relatively familiar technique to researchers in organizational behavior and management research. Computational details and mathematical derivations are placed in the supplementary materials of this article for the convenience of readers interested in the methodological specifications.

The article has a tutorial-like format, which outlines the step-by-step process required to build a multiple-membership survival model. The detailed tutorial-like format of the article illustrates procedural nuances (e.g., how to prepare and (re)structure datasets to allow for a count regression approach) by using concrete examples and sample data. We also offer guidelines for interpretation and visualization of model results.

Following the example of Bliese and Ployhart (2002), we encourage the readers to adopt the model proposed in this article by using the underlying R code (provided herewith), using

the *survival* package (Therneau & Grambsch, 2000), the *mgcv* package (Wood, 2017; Wood, Goude, & Shaw, 2015) and the *brms* package (Bürkner, 2017) that uses STAN at the back-end to run estimation procedures. Both R (R Core Team, 2017) and STAN software (Carpenter et al., 2017) are available online free of charge. Finally, we use a large archival dataset on occupational turnover of 1,642 early-career teachers (spanning five years after the start of their teaching career) working in 1,038 different schools over time, to demonstrate different aspects of our model.

This dataset is particularly useful to illustrate our modeling process due to its specific characteristics: (i) it contains archival, objective measures of events (i.e., leaving the teaching profession) over an extended period of time; (ii) it has a multiple-membership data structure (i.e., individuals teaching at multiple schools simultaneously, which is often required in early career teaching to achieve a full-time appointment; Sun & Pan, 2014), objectively observed from government records; and (iii) the teachers' membership intensity for a given school (both in terms of teaching there at all and in terms of teaching hours relative to other schools) in many cases varies over time, as early-career teachers often start off as substitutes who are not guaranteed an (equally large) assignment in the next term (Clandinin et al., 2015). Although readers might consider early-career attrition a niche topic from a content point of view, from a methodological viewpoint and for the purposes of this article, we consider teachers as being both theoretically and empirically equivalent to employees more generally, and schools to organizations or teams. For macro researchers, the teacher-level data could be equated to organizations, while the school level could be reframed as sectors or markets. We discuss both potential micro and macro research applications of our model in the Discussion. However, it is noteworthy that these examples are in no way intended to constitute an exhaustive list of potential applications. Instead, they represent active research areas in the fields of

organizational behavior and strategic management that we argue would benefit from using the proposed methodology.

Multiple Memberships: Problematization

The multilevel paradigm has permeated organizational behavior and management theory and methods in recent years, greatly advancing our understanding of organizational behavior (Mathieu & Chen, 2011) while at the same time revealing additional challenges that continue to plague multilevel research (Aguinis, Boyd, Pierce, & Short, 2011; Hitt, Beamish, Jackson, & Mathieu, 2007; Morgeson et al., 2015). Specifically, multilevel models—also known as hierarchical or mixed-effects models—allow researchers to adhere to the prescription of statistical theory that inference should take into account how a study was designed, and how the data was collected (Cafri, Hedeker, & Aarons, 2015; Gelman & Hill, 2006; Goldstein, 2011; Hox, 2010; Snijders & Bosker, 2012; Sun & Pan, 2014; Zappa & Lomi, 2015). To this end, multilevel modeling encourages a more realistic modeling of behavior and strategic management decisions by taking into account specific contexts that influence the concerned behavior or the decision (Scott, Simonoff, & Marx, 2013). Here, the underlying assumption is that individuals or organizations within the same context (e.g., employees within the same team or organizations operating in the same industry or sector) will make more similar decisions than individuals or organizations embedded in different contexts, given that actors within the same context influence each other, while simultaneously facing similar context-specific circumstances. The consequences of ignoring such non-independence of observations include; attribution errors (Ross, 1977), ecological fallacies (Robinson, 1950), and biased standard errors (Bliese & Hanges, 2004), which distort researchers' ability to make valid inferences and predict behavior in a precise manner (Bliese & Hanges, 2004; Gelman & Hill, 2006; Goldstein, 2011; Hox, 2010; Snijders & Bosker, 2012).

The current multilevel paradigm has as a central assumption that each lower-level unit is a member of only one higher-level unit—i.e., the classical ‘pure’ hierarchical nesting assumption— which is likely to misrepresent the reality of most organizational contexts (Hitt et al., 2007; Mathieu & Chen, 2011). For instance, O’Leary, Mortensen and Woolley (2011) observed that individual employees often work in multiple teams simultaneously. The implication thereof is that each team context is likely to influence the behavior and decision-making of each of its members (see, for instance, Chen et al.’s 2019 study on how leadership style experienced in one team can spill over to individual performance in another team). The influence of each team context can be proportional to the time an employee spends working in that team. At the more macro level, organizations can operate in multiple industries simultaneously (Bakker & Shepherd, 2017; Uzunca, 2018), or belong to so-called ‘hybrid’ strategic groups, each with their own strategic recipe (e.g., DeSarbo & Grewal, 2008). Such multiple-membership data structures thus violate the pure nesting assumption, for which the current methodological toolbox of organizational behavior and management has not yet provided a satisfying solution.

A commonly applied solution is the selection of a single ‘principal’ higher-level unit, which will most likely be the environment where the membership intensity (in terms of time or salience) is the highest. This approach, however, creates more conceptual problems than it is able to solve. For instance, the selection not only becomes arbitrary when there are ‘tied’ memberships, it also omits a great deal of potential variance. Imagine an employee who works in four teams and spends 30% of the total working time in Team A, 30% in Team B, 20% in Team C and the remaining 20% in Team D. Following the most commonly used approach, the employee will be assumed to be part of either Team A or Team B, which potentially represent fundamentally different contexts. Consequently, 70% of the total working time of this employee is not accounted for but assumed to be spent in the principal team as well. Additionally, all

teams with minority memberships are ignored, thereby omitting the contextual information that could be used to understand the employee's behavior. Serious potential for bias is thus introduced in the team-level effects since the total effect of omitted teams gets ascribed to the remaining higher-order units in an analytically intractable way (Cafri et al., 2015; Mathieu & Chen, 2011). The team-level variance is further reduced in case entire teams become omitted from the analysis, which occurs if none of the employees spend the majority of their working time in those specific teams.

To complicate matters further, a subject's (multiple-)group membership can change over time (Aguinis et al., 2011). Such changes can arise when either subjects switch groups (Mathieu & Chen, 2011) or when the population of groups changes (i.e., new teams emerging, teams changing size and/or composition, or teams disappearing over time). To avoid succumbing to attribution bias while interpreting the employee's behavior, the multilevel structure should thus be allowed to vary over time.

Mathieu and Chen (2011: 625) correctly state that “these different membership dynamics would wreak havoc with traditional multilevel theories and analytic techniques”. A series of simulation studies (Chung & Beretvas, 2012; Grady & Beretvas, 2010; Luo & Kwok, 2012; Meyers & Beretvas, 2006; Smith & Beretvas, 2014, 2017) have shown that failure to recognize and appropriately specify the hierarchical structure in a given dataset leads to invalid inference, such as inaccurate standard errors of regression coefficients and biased higher-order unit effects (called random effects in multilevel models).

In sum, analyzing (survival) data where the subjects of interest—whether they are employees or organizations—are subjected to a multitude of different environments requires more advanced techniques. As multiple-membership data structures arise when lower-level units are either simultaneously or sequentially members of multiple higher-level units, such techniques should ideally be able to take all environments into account simultaneously and

allow for dynamic membership (Grady & Beretvas, 2010; Lamote et al., 2013; Smith & Beretvas, 2017; Sun & Pan, 2014).

An Introduction to Multiple-Membership Survival Analysis

In this section, we conceptualize our multiple-membership survival model and show how it can be useful to answer research questions that imply or require multiple-membership data structures. In what follows, we will first explain how this model can be estimated using a count regression alternative. To this end, we start by explaining how to restructure survival data using our sample dataset of teachers, to demonstrate the count regression approach and the modeling of time-varying information. We then build the multiple-membership survival model by expanding on the survival regression approach.

Data Manipulation

Table 1 shows what most survival data looks like in practice. In the upper half of Table 1, information that was observed only once during a specified period in time is recorded. For each subject (ID), the last time the corresponding subject was observed (Time) is listed. This time may correspond to either an ‘event’ time (Event = 1; the subject had experienced the event during the study/observation period) or a ‘censoring’ time (Event = 0; the subject had not experienced the event at the end of the study/observation period). The lower half of Table 1 shows time-varying information—i.e., information that had been recorded at multiple times during the study/observation period. Notice that this time-varying information includes both covariate information and membership information. In case the membership structure for a subject is assumed to be time-invariant—that is, if there is no way in which subjects can change their membership or their intensity of participation in each of their environments—, we can as easily record this information in the time-invariant data format. Notice that each subject appears multiple times in this dataset, namely as many times as there is a change in the value of a covariate or a membership. In our example, each subject participates in at most six² different

environments (see S#-columns for environment identifiers; e.g., identification codes) during the specified period. The membership weights (see W#-columns) indicate how much working time each subject spends in each of these environments, proportional to the total amount of working time.

--- INSERT TABLE 1 ABOUT HERE ---

In the data example, the membership information is allowed to vary over time in several ways: (i) the higher-level units may change (e.g., subject 1 – time 13: environment E is traded for environment F); (ii) the number of higher-level units may change (e.g., subject 1 – time 7: environments C and D are added to the membership set, taking the total number of environments from two to four); and (iii) the membership intensity may change (e.g., subject 1 – times 1 to 6: the proportion of total working time spent in each environment varies). Note that, in this data example, the total number of environments in the membership set is not always equal to six—the maximum number of environments observed in the dataset. Therefore, nonexistent memberships are imputed as ‘missing’ membership observations for the sole purpose of working with a complete matrix. In the analysis procedure, the missing memberships are ignored and only the right number of environments in each subject-period observation is taken into account.

In order to take time-varying information into account in survival models, we should merge and convert our dataset into a stacked-form dataset, also called the ‘long-format’. To this end, we need to split our time dimension into intervals. Usually, the time horizon is split into as many intervals as there are unique event times, but smaller time increments are possible too, for example when covariates change between two distinct event times. In R, we can use the function *survSplit* from the *survival*-package (Moore, 2016; Therneau & Grambsch, 2000) to organize our dataset in the long format:

```
> survSplit(Surv(time, event)~., data=eventdata, cut=1:60, start='tstart',
end='tstop')
```

Upon running this command, the survival information contained in our sample dataset (data=eventdata) will be split up into sixty intervals (we take the number of intervals equal to the number of observed months in our sample dataset). Each interval will now start (t_{start}) where the previous interval ended (t_{stop}). Thus, each row will now correspond to a certain time-interval in which the subject could have been observed in the dataset. Congruent with traditional survival analysis, as soon as a subject experiences the event of interest (i.e., event equals 1), the subject will not occur in the dataset any longer. Time-varying information can now be merged with the long-format dataset containing the same survival information, using the subject identifiers and the time-stamped information in both datasets (code provided in Appendix 1). Table 2 shows a simplified version of the resulting dataset in the long-format, which can be used for regression modeling with time-varying information. Sample descriptives of the sample dataset can be found in Table 3.

--- INSERT TABLE 2 ABOUT HERE ---

--- INSERT TABLE 3 ABOUT HERE ---

Now that we have transformed our data into a long-format, we can start applying survival analysis techniques to predict the occurrence and timing of the event of interest using both time-invariant and time-varying information on the subject.

Survival Model: Proportional Hazards Regression

In survival analysis, one is mostly interested in the hazard rate or the ‘instantaneous event rate’, denoted as $\lambda(t)$, which represents the probability of experiencing the event exactly at time t , assuming that the event has not yet been experienced in the past. As such, this analysis aims to explain and predict the probability of experiencing the event-of-interest at each moment in time. For instance, it allows us to very specifically analyze the probability of an employee quitting his or her job in the organization after 12 months, provided that the employee has not already left the organization.

The most frequently used model to explain and predict the hazard function is the proportional hazards model (Duchateau & Janssen, 2008), which is specified as

$$\lambda_i(t) = \lambda_0(t) \cdot \exp(\eta_i) \quad (1)$$

$$\eta_i = \beta_1 x_1 + \beta_2 x_2 + \dots + \beta_k x_k \quad (2)$$

In short, $\lambda_0(t)$ represents the baseline hazard which functions as an intercept. However, instead of being a constant term, as is common in linear regression, the baseline hazard is now a function of time, reflecting how the instantaneous event rate changes over time for all subjects. The influence of the covariates on this baseline hazard is captured by the hazard ratio, $\exp(\eta_i)$, which contains the linear predictor η_i . This linear predictor is specified as a linear regression-like function of all covariates.

Consider, now, a proportional hazards model with time-varying covariates. In the long-format of our data, we can count the number of events—which is maximally one in standard survival modeling—for each subject-period observation (i.e., each row in the long-format data). The event rate can then be computed by dividing the number of observed events for the subject by the length of the time interval at hand. Hence, the counting process of events over time enables us to use a count regression technique to estimate the proportional hazards regression (Moore, 2016; see Appendix 2 for mathematical details). To show this equivalence empirically, we estimate the model using the standard *coxph* function, available in the *survival* package:

```
> coxph(Surv(tstart, tstop, event) ~ StartAge + factor(Gender) + wrkHours +
factor(Contract) + ..., data, ties = 'breslow')
```

The argument `ties = 'breslow'` is included by default to correct the estimation for tied event times (i.e., events happening at the same time).

Alternatively, we can use the count regression approach, which can be estimated using the *bam* function available from the *mgcv* package (*bam* is a large-dataset alternative function for *gam*, so results are identical; Wood et al., 2015):

```
> bam(event~offset(log(tstop-tstart))+s(tstop,k=60)+StartAge+
factor(Gender)+wrkHours+factor(Contract)+... ,data,family='poisson',
discrete=TRUE, nthreads=2)
```

The arguments `discrete=TRUE` and `nthreads=2` are included with the sole purpose of speeding up the estimation procedure. Notice that the event indicator (`event`) is regressed on: (i) the logarithm of the time duration of each interval as an offset, `offset(log(tstop-tstart))`—the regression coefficient of this offset term is set to one (by default) which turns the event indicator into an event rate over time; (ii) a smooth and flexible baseline hazard function (which will be discussed later), and (iii) all covariates, which include both time-invariant and time-varying covariates. Selected results from this analysis are presented in Table 4.

--- INSERT TABLE 4 ABOUT HERE ---

The results in Table 4 show that the coefficients of the traditional proportional hazards (*coxph*) model and the count regression alternative (*bam*) model are not very different from each other. Hence, the count regression alternative is a valuable (and perhaps more accessible) alternative technique for survival analysis. The coefficients, which can be interpreted as hazard ratios after exponentially transforming them, are in line with expectations.

One important aspect of the interpretation of a hazard regression model is that covariate effects indicate increases or decreases of the event rate. However, they do not give any information about the shape of the event rate over time. This information is captured by the baseline hazard. Hence, to fully understand the hazard regression model, the specification of the baseline hazard function should be taken into consideration.

Survival Model: Specification of the Baseline Hazard

The baseline hazard in a hazard regression model is represented by a function of time, which means that this functional form/shape needs to be appropriately specified. Misspecification of the baseline hazard function has been shown to result in (severely) inconsistent estimates of the regression coefficients (Duchateau & Janssen, 2008), which lead

researchers to wrong conclusions. However, there exist multiple alternatives for specifying the baseline hazard.

The first alternative is that researchers make an assumption about a parametric distribution for the underlying baseline hazard. Each distribution implies a different shape of the hazard function. Popular distributions are the Weibull, Exponential, Log-Normal, Log-Logistic, Gamma, Pareto, Gompertz distributions, and many more (Duchateau & Janssen, 2008; Moore, 2016). Unfortunately, in practice, a single parametric distribution often fails to fit the baseline hazard nicely to (finite) sample data. Therefore, as second alternative, the hazard distribution is often left unspecified and replaced by a non-parametric (data-driven) estimate. The proportional hazard model then becomes the ‘semi-parametric’ model by Cox (1972). A last—and preferred—alternative is the estimation of a smooth and flexible baseline hazard function. Instead of estimating either a specific baseline hazard for each time interval—which results in a piecewise constant hazard (see Moore, 2016)—, or using a fourth-order polynomial of time (e.g., Elghafghuf et al., 2014a; Elghafghuf et al., 2014b), we use a smoothing spline of time, $s(t_{\text{stop}}, k=60)$. We allow a high degree of flexibility or smoothness—this can be achieved by setting the hyperparameter k close to the number of distinct observation times (Hastie, Tibshirani, & Friedman, 2009)—, which yields a function which is similar to the non-parametric baseline hazard of a traditional Cox model. However, a smooth spline function yields a smooth trajectory over time which is more interpretable (and usable for predictions) on times where no true events are observed than a discontinuous staircase function yielded by a Cox model (see Figure 1).

--- INSERT FIGURE 1 ABOUT HERE ---

In some cases it is reasonable to assume that the baseline hazard will be (disproportionally) different for different clusters in the data. For instance, male and female subjects have a fundamentally different baseline job attrition pattern. A survival model can take

this into account by stratifying the baseline hazard, which means that different baseline hazards are estimated for each cluster. In R, this can be achieved using the following code, which estimates different baseline hazards for male and female subjects:

```
> coxph(Surv(tstart, tstop, event)~strata(Gender)+StartAge+WrkHours+
factor(Contract)+..., data, ties='breslow')

> bam(event~offset(log(tstop-tstart))+s(tstop, k=60, by=factor(Gender))+
StartAge+WrkHours+factor(Contract)+..., data, family='poisson', discrete=TRUE,
nthread=2)
```

In the regression output (not shown) there is no parameter of gender anymore, since the effect is entirely taken into account by estimation of different gender-specific baseline hazard functions.

We have now shown how the traditional survival model can be correctly and flexibly estimated using a count regression technique. Unfortunately, the proportional hazards model still fails to take the multilevel data structure of our sample data into account. At this point in our model, all subjects are assumed to be independent, even though in reality participation in any common environment creates a degree of non-independence among subjects. This means that there exists correlation between the event times of subjects nested in the same environment. To account for this non-independence, multilevel analytical techniques should be adopted—in survival analysis, these are called frailty models (Duchateau & Janssen, 2008).

Frailty Models

Traditionally, multilevel models quantify the (higher-level) environment's influence on the outcome of each of the subjects that participate in the environment. For instance, if turnover is the event of interest, the likelihood of employees quitting their jobs is likely to vary between organizations, due to differences in work climate, work-conditions or promotion opportunities, among many other factors. Nonetheless, all employees within the same organizational environment share the same frailty or risk factor, since they are exposed to the same

organizational conditions. In this way, the event rates of employees are non-independent. Taking the unobserved environmental influence into account solves the non-independence problem. The (multilevel) frailty model (Duchateau & Janssen, 2008) is a straightforward extension of the well-established multilevel models to survival regression, which can be specified as

$$\lambda_{ij}(t) = \lambda_0(t) \cdot \exp(\eta_{ij} + w_j) \quad (3)$$

where $\lambda_{ij}(t)$ expresses the hazard function for subject i in environment j , which consists of a baseline hazard, $\lambda_0(t)$, and a hazard ratio, $\exp(\eta_{ij} + w_j)$, which now contains both a linear predictor, η_{ij} , and the unobserved environmental influence, w_j . Usually, it is assumed that this environmental influence is random across environments—hence, the term random effects—and that these influences follow a Gaussian (normal) distribution (Gelman & Hill, 2006; Goldstein, 2011; Hox, 2010; Scott et al., 2013; Snijders & Bosker, 2012). The interpretation of this random effect is intuitive: individuals in an environment with an environmental effect larger than zero (i.e., $w_j > 0$) will have a higher hazard and are, thus, more ‘frail’. In such an environment, subjects have a higher risk of experiencing the event. Conversely, individuals within an environment with an environmental effect smaller than zero (i.e., $w_j < 0$) are, in general, better able to ‘resist’ the event of interest.

The estimation of frailty models in R is a straightforward extension of the previous models, using the following commands:

```
> coxph(Surv(tstart, tstop, event) ~ StartAge + factor(Gender) + wrkHours +
factor(Contract) + ... + frailty(GROUP, dist = 'Gaussian', method = 'REML'),
data, ties = 'breslow')

> bam(event ~ offset(log(tstop - tstart)) + s(tstop, k = 60) + StartAge +
factor(Gender) + wrkHours + factor(Contract) + ... + s(GROUP, bs = 're'),
data, family = 'poisson', discrete = TRUE, nthreads = 2)
```

In short, we add a frailty term, specifying the nesting variable `GROUP`, which takes the identification number of the environment to which the subject belongs. In the *bam* model, however, the random effect/frailty distribution is by default Gaussian. It is important to note that a factor variable has to be used as grouping variable, which can be easily accomplished by running the following code before fitting the model:

```
> data$GROUP<-factor(data$GROUP)
```

In the *coxph* model, we specify that our frailty term must follow a Gaussian (normal) distribution. Otherwise, the frailty term follows the default distribution, which is the gamma distribution (Duchateau & Janssen, 2008). In this article, we chose the normal distribution to ease the explanations and to maximize the congruence with more established multilevel models.

Frailty Models: Comparing Different Alternatives

In multilevel analysis, it is an ongoing challenge to identify the most salient higher-order unit to study (Mathieu & Chen, 2011). This identification may not always be very straightforward as boundaries between higher-order units are not always clearly defined. For instance, work teams can be formed based on (i) a project in which multiple employees are collaborating; (ii) having a common supervisor or leader; (iii) a shared area of expertise; or (iv) strong social network ties among colleagues. Since it is likely that these different ways of identifying subjects' work teams may not result in the same team structures, it could raise ambiguity about the most appropriate higher-level unit to model (Ashforth, Harrison, & Corley, 2008; Rousseau, 1998; Wright, 2009). Anecdotally, many organizational behavior researchers have experienced firsthand that self- versus organization-reported team structures often differ, creating challenges for multilevel survey design. This issue may even require a solution where multiple higher-order units have to be taken into account (Chen et al., 2019).

One can, for instance, employ a model pitting approach (Joo, Aguinis, & Bradley, 2017) to fit multiple competing models and ex-post select the model which provides the best fit for

the data at hand. The same strategy can be adopted as a work-around for the multiple-membership structure of our sample dataset: we can fit multiple frailty models which select only a single ‘principal’ higher-level unit and then compare which model fits the data best.

To complicate matters further, there are different ways to select the ‘principal’ higher-order unit; among other solutions, one could (i) opt for the higher-order unit with the highest membership weight for each time-period observation; (ii) opt for the higher-level unit in which the subject has spent most time throughout the career/observation time; (iii) identify the first higher-level unit the subject participated in, assuming that first impressions matter (most); or (iv) identify the last higher-level unit the subject participated in, since this is the same environment in which the subject had experienced the event. In Table 5, we reported the results of four frailty models, fitted using both the *coxph* code as well as the *bam* code.

We can see that the traditional frailty model does not yield substantially different estimates than the alternative count regression model. Comparing these results with the parameter estimates from the traditional survival models, we can conclude that the parameter estimates remain somewhat unchanged and the standard errors of the parameter estimates do remain stable as well. Formally testing whether adding a random effect term has had an impact on the statistical performance is often difficult, because the null hypothesis restricts the variance parameter to zero, which lies on the boundary of all plausible values. The *mgcv* package, however, does provide an approximate test based on Wood (2013). Based on these tests (not reported) the statistical added value of the random effects is negligible.

Apart from purely statistical motives, the main advantage of the frailty model is that, by using it substantive scholars can now determine what and how large the impact of a higher-level unit is on the event rates of its members is. Without any covariates—which is what we call the unconditional baseline (or null) model—the higher-order effects in our sample dataset are normally distributed with a standard deviation ranging between .37 and .45 (depending on

which nesting structure we used). In approximate terms, this means that the frailties/random effects themselves, will range between -0.74 (i.e., $-2*sd$) and 0.74 (i.e., $2*sd$) using a 95% interval. Although this might not seem to reflect a large effect size, its impact on the event rate is of substantial magnitude. The hazard ratio of the random effect can be computed by taking the exponential transformation of this effect. In doing so, the multiplicative impact on the event rate ranges between .48 and 2.10. In practice, this means that the probability of experiencing the event at each moment in time can be doubled or cut in half depending on the higher-order unit in which the subject is assumed to participate.

--- INSERT TABLE 5 ABOUT HERE ---

However, it is important to keep in mind that, at this point in the model, our higher-order unit identification is still seriously flawed as our selection is both arbitrary and incomplete. Most frailty models in Table 5 only make predictions for about 720 out of 1,038 higher-order units, due to the omission of minority-membership units. It would be safer to take all higher-order units into account, both for generalizability and for avoiding serious bias in higher-order effects (Cafri et al., 2015; Mathieu & Chen, 2011). This bias arises largely because the effect of omitted teams is absorbed by the principal environment, which may bias the true environmental effect. A compelling example of this caveat has been recently published by Chen and colleagues (2019) who found that for employees who are members of multiple teams simultaneously, negative team influences (i.e., a lower degree of empowering leadership) can be buffered by positive influenced in another team (i.e., a higher degree of empowering leadership). This finding would have gone undetected if they had neglected the fact that these employees were nested in multiple teams simultaneously. As a result of such unobserved buffering between adverse team effects, the consequences of neglecting multiple-memberships might even cause net environmental effects to remain entirely unobserved.

Multiple-Membership Frailty Model

Articles reporting the use of a multiple-membership survival model are generally scarce and are certainly not yet found in the field of management or organizational behavior. To the best of our knowledge, at present, its applications are only found in veterinary medicine (e.g., studying calf mortality where calves are nested in herds, which are in turn nested in regions and multiple clinics; Elghafghuf et al., 2014b), epidemiology (e.g., studying marriage and childbirth of individuals in neighbouring regions; Manda, Feltbower, & Gilthorpe, 2012) and education (e.g., studying high school dropout where students have migrated between schools; Lamote et al., 2013). In line with these applications, the multiple-membership (shared) frailty model can be written as

$$\lambda_{i\{j\}}(t) = \lambda_0(t) \cdot \exp\left(\eta_i + \sum_{j \in C} \pi_{ij} w_j\right) \quad (4)$$

where $\lambda_{i\{j\}}(t)$ is the hazard function for any subject i nested in a subset C of higher-order environments. Thus, each subject can now be nested in multiple environments at the same time. The total higher-level effect of the environments is expressed as a weighted sum of all the higher-level effects in the membership set. The weights, π_{ij} , are equivalent to the intensity of the membership (Elghafghuf et al., 2014b). Simply put, the multiple-membership frailty model is equivalent to the traditional frailty model, but a single frailty term has now been replaced by a weighted sum of multiple frailty terms.

Conventionally, the weights are operationalized as proportional memberships, which means that they sum to one for each individual, $\sum_{j \in C(i)} \pi_{ij} = 1$. There are many scenarios to assign weights in a multiple-membership model (Chung & Beretvas, 2012). Either equal weights can be used for all environments in the membership set, or weights can be made proportional to the time spent in—and, thus, exposure to—each of the environments. More complicated weighting schemes can be used as well (e.g., weights which take the order of exposure to each environment into account). Ultimately, the optimal choice of weights remains

unknown (Chung & Beretvas, 2012; Smith & Beretvas, 2014). The study of Smith and Beretvas (2014) indicates that as long as the multiple-membership structure of the data is taken into account, the choice of the weight pattern, given that it is based on reasonable arguments, does not have a high impact on the parameter estimates for both fixed and random effects.

There are no straightforward techniques to estimate parameters for a multiple-membership survival model using the traditional survival likelihood techniques known in our field. However, using the count regression alternative, we are able to specify a relatively simple regression model to which we add a multiple-membership random effect (i.e., frailty), as is currently implemented in the *brms* package (Bürkner, 2017):

```
> brm(event~offset(log(tstop-tstart))+s(tstop,k=60)+
  StartAge+factor(Gender)+wrkHours+factor(Contract)+...+
  (1|mm(S1,S2,S3,S4,S5,S6,weights=cbind(w1,w2,w3,w4,w5,w6))),
  data=data,family=poisson,chains=4,cores=4)
```

In short, we use the same formulaic equation as in the previous models, except for the frailty term which is now specified using a (1|GROUP) format (note to readers who are familiar with the package *lme4*: *brms* uses the same mixed-effects notation as this package). The grouping factor is converted to a (weighted) sum of random effects by means of the wrapper function `mm()`. The arguments of this wrapper function are: a sequence of membership indicators (cfr. S#-columns in Table 1 and Table 2), followed by an optional weighting matrix. This matrix consists of all membership intensities (cfr. W#-columns in Table 1 and Table 2).

The *brms* packages uses a Bayesian estimation procedure, sourcing the STAN software (Carpenter et al., 2017) at the back-end. To the best of our knowledge, there are no readily implemented or broadly accessible statistical tools to fit a multiple-membership model without the use of Bayesian Markov Chain estimations. This means that fitting a multiple-membership model requires a bigger (computing) time investment. Luckily, *brms* has implemented a quick and efficient estimation procedure of Hoffman and Gelman (2014). Using the `chains=4` and

`cores=4` arguments of the estimation command above, four sampling chains are estimated simultaneously at four cores—note that these numbers depend on the number of physical computing cores available in the computing device (see Bürkner, 2017)—which enhances the quality of our estimates, while simultaneously keeping the computing time in check.

In the model building process, it is important to compare nested and non-nested models with each other so as to check whether each modeling step adds value. The *brms* package prescribes that models be compared using the Widely-Applicable Information Criterion (WAIC) values (Vehtari, Gelman, & Gabry, 2017; Watanabe, 2013). Information criteria express a penalized deviance—i.e., minus two times the log-likelihood of the model—where the penalty increases with model complexity, which equals the number of estimated parameters in the model. Since the numeric value of this deviance depends on the data at hand, information criteria cannot be interpreted by themselves, nor can they be judged using a threshold value. Their virtue, however, lies in their comparability across models (Claeskens & Hjort, 2008). In general, ‘lower is better’ as the model with the smallest information criteria reflects the best fitting model within a set of competing models. (Please note that it is not possible to compare different information criteria with each other; e.g., AIC cannot be compared with WAIC, since they result from fundamentally different estimation techniques.) WAIC can be computed for any fitted model using the following R-code:

```
> WAIC(fittedmodel, pointwise=T)
```

Multiple Memberships: Problematization Revisited

Conceptually, the multiple-membership survival model takes all different higher-level units into account, which avoids the selection of a principal higher-order unit. As discussed earlier, the latter practice creates more problems than it solves—i.e., (i) misrepresentation of the higher-order nesting structure (i.e., all minority memberships are ascribed to the single, principal higher-order unit; (ii) multiple higher-order units are omitted from the analysis

altogether if none of the subjects has its principal membership in these respective higher-order units; and (iii) higher-order effects carry the potential of being seriously biased because the true higher-order effects are consolidated into the principal units' effects. The question remains, however; to what extent these conceptual problems weigh on the interpretation of empirical findings. To provide an answer to this question, we compare the multiple-membership frailty model with the single frailty model identifying the first school as the principal higher-order unit (cfr. the third column in Table 5). (Selection of other principle higher-order units yielded comparable findings when a single-frailty model was applied.) The empirical results for the multiple-membership model can be found in the third column of Table 6.

In our sample dataset, results indicate that the coefficient estimates do not differ much in magnitude. The standard errors of the estimates increase only marginally, which might be due to the big sample size of the sample dataset (i.e., more than 66,000 observations over time). In the higher-level part of the model (i.e., random effects), we can see that the standard deviation of the frailty terms decreased only marginally. This means that the single frailty model potentially overestimates the differences in unobserved school-level effects.

In fact, the higher-level frailty estimates are estimated quite differently from the multiple-membership frailties (see Figure 2). The correlation between the frailty estimates is only .58 (based on 712 paired school frailty estimates, since 326 other schools were not taken into account in the single frailty model). In Figure 2 we can see that the single frailty model yields quite divergent frailty estimates, both in sign and magnitude. About 28% of all higher-order units were found to have an opposite sign and the magnitude of the frailty were, on average, more than four times as large as the estimate from the multiple-membership survival model.

--- INSERT FIGURE 2 ABOUT HERE ---

We can also look at the total impact on the higher-level effect of the frailties on the event rate over time (i.e., for each subject-period observation in the sample dataset). For the single-frailty model, the frailties exert their full influence at each subject-period observation in the sample dataset. However, for the multiple-membership model we need to compute the weighted sum of multiple frailties manually. This can be done by multiplying the weights with the respective frailties at each subject-period observation and summing them for each subject-period observation. In Figure 3, we can see how these total higher-level effects on the event rate differ between both models. Again, many of total higher-level effects diverge quite strongly from each other. The correlation between the total effects of both models, taking all subject-period observations into account, is only .42.

Specifically, in the multiple-membership survival model, the total higher-level effect ranges between $-.17$ and $.23$, indicating that the 16% decrease or the 26% increase in the event rate is caused by the environment in which the subject is nested. For the single-frailty model, this environmental effect ranges between $-.21$ and $.40$, which indicates a 20% decrease or a 48% increase in event rates. This may mean that in many cases, the effects of one environment on the event rate can be buffered by other environments. This is conceptually close to the observations made by Chen et al. (2019) with respect to the spill-over of leadership empowerment in one team on individual-level performance in another. In the most extreme cases where the multiple-membership model predicts a near zero total environmental effect (due to buffering environments), the single-frailty model can predict a nearly maximal total environmental effect on the event rate (i.e., a 48% increase).

--- INSERT FIGURE 3 ABOUT HERE ---

In sum, a single principal higher-order unit for a sample of subjects characterized by multiple memberships—thus ignoring all other memberships—creates both conceptual and empirical problems, and seriously biases the higher-level effects (i.e., frailties) as well as the

total environmental effects on the event rate. Hence, many researchers could draw incorrect conclusions with respect to the effect of the environment on the subjects event rates over time.

Step-By-Step Guide a Multiple-Membership Survival Model

In this section, we conduct a step-by-step estimation of a multiple-membership survival model. We will start with a simple hazard regression model without random effects, working our way up to a complete hazard regression model by adding more model complexity. At each step, we will compare WAIC values between models to decide whether the more complicated model adds value to the more parsimonious model. We will build the model using the following four steps: (1) adding the multiple-membership frailty terms to an empty model to check whether the multiple-membership data structure adequately fits the data structure; (2) adding subject-level covariates to the multiple-membership frailty model; (3) adding higher-level covariates to the multiple-membership frailty model; and, finally, (4) transforming the covariates' effects into flexible fits which do not necessarily assume linear effects.

STEP 1: From Baseline to Frailty Model

As a first step for progressing from a traditional survival model to a multiple-membership survival model, we need to compare some baseline models—more specifically an intercept-only model without frailty terms and a multiple-membership model without predictors. We fit both models using the *brms* package:

```
> No_Frailty<-brm(event~offset(log(tstop-tstart))+s(tstop,k=60),
data=data,family=poisson,chains=4,cores=4)

> MM_Frailty<-brm(event~offset(log(tstop-tstart))+s(tstop,k=60)+
(1|mm(S1,S2,S3,S4,S5,S6,weights=cbind(w1,w2,w3,w4,w5,w6))),
data=data,family=poisson,chains=4,cores=4)
```

The results of both models can be found in the two first columns of Table 6. The WAIC value of the multiple-membership model is lower than the model without frailty terms, which shows that adding the frailty terms adds more value to the model relative to the added model

complexity (i.e., adding 1038 higher-level effect predictions and a variance parameter estimate). The standard deviation of the frailty term is estimated at .43. Note that this standard deviation is estimated without taking any covariate information into account.

STEP 2: Add Subject-Level Covariates

As a second step in the modeling process, we can include subject-level covariate information to explain and predict the subjects' event rates. The inclusion of covariates can change the estimated frailty dispersion (i.e., standard deviation). Using our sample dataset, we estimate the following model:

```
> MMFrailty_11cov<-brm(event~offset(log(tstop-tstart))+s(tstop,k=60)+
StartAge+factor(Gender)+wrkHours+factor(Contract)+
(1|mm(S1,S2,S3,S4,S5,S6,weights=cbind(w1,w2,w3,w4,w5,w6))),
data=data,family=poisson,chains=4,cores=4)
```

The model's results are added in the third column of Table 6. We can see that the WAIC value (WAIC = 4894) is much lower than the baseline multiple-membership survival model (WAIC = 5240), which indicates that the covariates add valuable information to the model.

Moreover, the coefficient estimates are quite consistent across models. We can infer from the coefficients that teachers who started teaching later in their career are more likely to leave the profession quickly since each additional year of age increases the turnover probability by 7%. Conversely, each additional teaching hour decreases the attrition probability by about 11%. Tenured teachers experience a significantly lower attrition probability compared to non-tenured teachers, who in turn have a lower attrition probability than interim teachers. No significant gender differences were found. With respect to the dispersion of frailty terms across higher-order units, we observe that adding the subject-level covariates leads to a decrease in standard deviation of the school-effects.

--- INSERT TABLE 6 ABOUT HERE ---

STEP 3: Add Environment-Level Covariates

In order to explain more of the higher-level variance, we can add covariates at the higher level. In our sample data, this means that we should add a covariate that is school-specific. Our archival dataset did not contain school-level predictors. However, for illustrative purposes, we can construct a higher-level covariate: the early-career attrition climate at the school level (Van Droogenbroeck, Spruyt, & Vanroelen, 2014; Wang, Hall, & Rahimi, 2015). This variable is constructed by calculating the proportion of early-career teachers who quit the teaching profession while working at the respective school (irrespective of the membership intensities of the teachers working in that school).

Adding a higher-level covariate as a predictor in an ordinary multilevel model is as easy as appending the school-level data to the dataset and adding the covariate to the model. However, in a multiple-membership model, we have multiple schools to which we need to append the school-level covariate. Additionally, we also need to take into account that all memberships are weighted according to their membership intensities. Since these membership intensities indicate the degree of exposure to the specific school context, it is only logical that each school's covariate is taken into account accordingly. In this article, we will propose one—and, in our opinion, also the easiest—approach. At first, we append the school-level covariate data using each membership column (S#). Then, we multiply each higher-level covariate vector with the respective membership intensity vector. Finally, we sum up all the weighted school-level covariate information to create a weighted combination of school-level covariates. We, then, add the higher-level covariate to the model:

```
> MMFrailty_l2cov<-brm(event~offset(log(tstop-tstart))+s(tstop,k=60)+
StartAge+factor(Gender)+wrkHours+factor(Contract)+ECAC_sch1+
(1|mm(S1,S2,S3,S4,S5,S6,weights=cbind(w1,w2,w3,w4,w5,w6))),
data=data,family=poisson,chains=4,cores=4)
```

The modeling results show that the WAIC value reduces further (WAIC = 4480) indicating that the school-level covariate again added valuable information to the model (see

fourth column in Table 6). All the coefficients decreased slightly in magnitude, but the significance and direction of the significant effects did not change. The school-level covariate indicates that if more colleagues quit the teaching profession at (one or more) of the schools, the remaining teachers in those schools will also be more likely to quit the teaching profession. This indicates that turnover contagion, coined by Felps et al. (2009), is also found in our sample data. We can also see that the standard deviation of the multiple-membership frailty terms decreased only slightly. This could mean that the school-level covariate does not explain a lot of the school-level effects in our data. Another possibility is that we may have disregarded the functional specification of our covariates. We tackle that possibility in the next modeling step.

STEP 4: Add More Flexibility (Optional)

Traditional modeling approaches model the impact of covariates on an outcome in the most simple, linear way by assuming that the impact of a covariate monotonically increases or decreases. Such an assumption thereby disregards the presence of potential inflection points where the effect levels out or even changes direction (Pierce & Aguinis, 2013). Hastie, Tibshirani and Friedman (2009) stated that non-linear effects are actually quite prevalent in real life, which was further confirmed by recent studies in our field (e.g., congruence in organizational research; Edwards & Parry, 2018; congruence in management research; Nikolaeva, Bhatnagar, & Ghose, 2015; 'too-much-of-a-good-thing' effects; Vergauwe, Wille, Hofmans, Kaiser, & De Fruyt, 2017). If potential non-linearity of effects (e.g., quadratic, cubic or even higher-order polynomial functions) is ignored and the relationship is specified as a simple linear effect, the model “may be parsimonious, but it will be wrong” (Bliese & Ployhart, 2002: 383).

In our model, adding non-linear terms is not too much of a leap forward. Recall that we modeled the baseline hazard function as a flexible function of time, which produced a smooth and flexible (non-linear) regression function over time. Hence, we can allow the effects of the

continuous covariates in our model—more specifically starting age, working hours and early-career attrition climate—to be modeled using a smooth spline function, as follows:

```
> MMFrailty_smooth<-brm(event~offset(log(tstop-tstart))+s(tstop,k=60)+
s(StartAge)+factor(Gender)+s(WrkHours)+factor(Contract)+s(ECAC_schl)+
(1|mm(S1,S2,S3,S4,S5,S6,weights=cbind(w1,w2,w3,w4,w5,w6))),
data=data,family=poisson,chains=4,cores=4)
```

Doing so decreases the WAIC value (WAIC = 4192), which indicates that changing the functional form of our covariates has improved the model fit. In the last column of Table 6, the model results are reported. The coefficients of the continuous covariates, which are now estimated by smooth functions, are hard to interpret. Essentially, the population-level effect gives an indication of the general direction of the effect (i.e., whether the effect on average is positive or negative), while the standard deviation of the smooth effect gives only an indication of how ‘wiggly’ the smooth effect behaves around the population-level effect. In order to be able to interpret the smooth functions over time, it is easier and more insightful to plot the smooth functions and visually explore the effects. Hence, smooth trajectories of the marginal effects should be plotted. These plots show the effects of each covariate along its range, often showing nonlinear smoothed lines. Statistical significance can be assessed using error regions on the plots. This can be done using the *marginal_smooths* function available from the *brms* package. The command for it in R is as follows:

```
> marginal_smooths(MMFrailty_smooth, rug=F, ask=F)
```

The resultant smooth functions are plotted in Figure 4. From these results we can see that our continuous covariates cannot be fitted well using a straight regression line. Instead, teachers who start their teaching career between the age of 20 and the age of 22 are less likely to leave the teaching profession quickly. After that age and up to the age of 30, starting teachers are more likely to quit soon. Thereafter, the impact levels out around a zero impact. Note that the effect of starting age is not statistically significant after the age of about 28 years. The

sudden increase in the grey-shaded area, which represents the 95% confidence interval, is due to few observations in that range of the variable.

With respect to the number of working hours, we can observe a mild form of a ‘too-much-of-a-good-thing’-effect for this covariate. If a teacher is not employed as a full-time employee (i.e., about 20 teaching hours), the probability of quitting the teaching profession is very high. If a teacher is employed for more hours, it has a protective effect on teachers. All hours that exceed full-time employment can serve as a slack or a buffer against unexpected changes in their employment contract (e.g., the teacher can easily give up some hours and still retain his or her full-time employment). However, if a teacher has to perform more than 26 hours a week, the protective effect wanes. If the workload exceeds 35 hours a week, the smooth effect indicates that the workload starts to drive teachers out of the profession again. It should also be noted that the confidence bands indicate that beyond 27 working hours the effect is not significantly different from zero any more.

Finally, if a teacher is employed in a set of schools where the total proportion of early-career teachers who quit the profession remains under 7%, teachers are less likely to quit the teaching profession. Conversely, this effect is positive if the school set’s early-career attrition climate is characterized by more attrition. In fact, the effect becomes very large when more than 50% of the early-career teachers leave from a given school set. The event rate of early-career teachers in such environments can be well above 50 times (i.e., $\exp(4)$) as high as an early-career teacher’s working in a very stable environment characterized by the absence of any early-career attrition among colleagues.

--- INSERT TABLE 6 ABOUT HERE ---

Following these four steps allows researchers to run a multiple-membership survival model that is capable of (i) estimating the effect of subject-level covariates on the subjects’ event rate over time; (ii) estimating the effect of higher-level covariates on the subjects’ event

rate over time; (iii) while parsimoniously allowing these effects to be non-linear, and (iv) predicting a better estimate of the higher-order contextual influence on the event rate—that is, estimating effects that take into account all simultaneous memberships of a subject as well as the membership intensities.

Discussion

The main objective of this article has been to advance general understanding of the multiple-membership analysis methods used to analyze the occurrence and timing of events in complex organizational environments. In doing so, it adds the resulting multiple-membership survival model to the methodological toolbox of the field of organizational behavior and management. Throughout the article we have outlined a step-by-step-approach for implementing our methodology such that it can be used by researchers working on events embedded in multiple-membership data structures. Essentially, the analysis we propose requires two distinct steps; first, gathering traditional survival data and transforming these into the long-format, after which the data is merged with both time-invariant and time-varying covariate information (as well as membership information); second, using a stepwise procedure to fit the most parsimonious and correct multiple-membership survival model and interpreting its results. In each of these steps, we emphasized the accessibility of the model by building the model from a simple count regression model and provide the estimation code for the model using a freely accessible, open-source software.

Throughout this article, a sample dataset on early-career attrition among teachers was used to illustrate the multiple-membership survival model. However, application-wise, the scope of the model is much broader. As previously mentioned, both theoretically and empirically, teachers can be understood as being equivalent to employees (or in a macro study: organizations), and schools being equivalent to organizations or teams (or in a macro study: sectors or markets). In the next section, we will discuss several active research areas in the fields of organizational behavior

and strategic management that would benefit from the proposed methodology. These examples are in no way intended to constitute an exhaustive list of potential applications.

Application of Multiple-Membership Survival Models to Organizational Research

As mentioned earlier, survival analysis can be applied to any phenomenon that can be conceptualized as the occurrence of an event situated on a timescale, albeit calendar time, or time elapsed since a previous event or since the start date of a study. For instance, a study of the speed of finding employment in a given labor market can be operationalized as starting at the moment of graduation, the start of the job search, or the end of a previous employment period (e.g., lay-off or job loss). Additionally, many organizational contexts in our field are characterized by multiple-membership structures. Such nesting structures emerge when lower-level units are simultaneously member of multiple higher-level units, which makes it relevant for many areas within the domain of organizational behavior and management. The application of some of these are discussed below.

Micro research examples. Many contemporary organizations have a matrix structure where employees can belong to multiple teams (Chen et al., 2019; O'Leary et al., 2011; Zaccaro, Marks, & DeChurch, 2012). Additionally, different teams are usually supervised by different leaders, which means that data collected at the employee level can be thought of in terms of employees nested in multiple teams/under multiple supervisors at the same time. Any event prevalent at the employee level that may be driven by accumulated effects across the set of teams to which an employee belongs, can be analyzed using a multiple-membership survival model. In doing so, each team's contribution to the probability and speed of the event occurring at the individual level can be estimated. We will now discuss some examples of research questions that could adopt this approach.

First, the occurrence and timing of psychological contract breach (see Vantilborgh et al., 2016), defined as “a violation of an individual's beliefs regarding the terms and conditions

of a reciprocal exchange agreement between that focal person and another party” (Rousseau, 1989: 123). This phenomenon can be studied using our multiple-membership survival model to reach more accurate and valid inferences about how each team/supervisor adds to the risk and likelihood of a subject perceiving a psychological contract breach. Similarly, the duration of reactions to, and recovery from, a psychological contract breach (see Solinger et al., 2016)—taking the moment of the breach as the starting point of analysis—could be driven by team/supervisor-provided resources (Griep et al., 2016).

Second, the multiple-membership survival model can also be used in samples of employees working in a multi-site organization (McCabe & Trevino, 1997), by which we mean that the different divisions of the organization are geographically dispersed and employees’ tasks are spread across these divisions (Alexander, Hubers, Schwanen, Dijst, & Ettema, 2011). Another similar example relating to task dispersion is multi-tasking (Leroy, 2009) and its impact on outcome variables that can be operationalized as events—such as exhaustion, burnout, and organizational turnover, among others.

Third, in addition to multiple-membership organizational structures, we can also investigate individual’s decisions as nested in spatial environments (Manda et al., 2012). For instance, we can investigate the speed of finding a (new/first) job, which will most likely depend on the availability of job opportunities in the subject’s region. This (geographical) environment can be defined as the set of regions (e.g., municipalities, provinces or states) that are included in the perimeter in which an individual has searched for a job. As an extension, geographical regions can be characterized by different commuting times (i.e., distance to the individual) as well as macro-economic structural characteristics (e.g., regional unemployment rates), all of which can be taken into account when studying (re)employment success (Wanberg, Kanfer, Hamann, & Zhang, 2016).

Macro research examples. The multiple-membership survival model can be applied to decisions made by organizations (e.g., entry or withdrawal decisions in an industry or market; internal corporate restructuring decisions), or to events occurring at the organizational level (e.g., organizational defaulting, business venture success, innovation development and/or success) within a cooperative and/or competitive business environment. A business environment can be thought of as a set of competing organizations. Since every organization has its own market power, it can be expected that each organization influences the decisions of all other organizations in the same strategic group, industry or market.

Firstly, in such a setting, researchers can apply the multiple-membership survival model to the study of market/industry entry decisions (see Uzunca, 2018). Upon entry into an industry, each organization is confronted with a set of incumbent organizations which are already rooted within the industry. By adopting the perspective of idiosyncratic competitive influences of incumbent firms in the competitive environment of the industry, each competitor can exert its own influence on the entrant's potential to gain foothold in the targeted industry. Hence, the strategic decisions of an entrant firm are, in fact, influenced by this set of competitor-specific influences. The multiple-membership model can be used to take all 'competing firms' into account simultaneously. In a similar fashion, organizations in co-operative business environments may see their chances of establishing a successful business venture change depending on the firms which are open to constituting a business venture.

Secondly, in the field of innovation, research questions are often aimed at discovering the antecedents of innovation speed or time needed to develop a successful innovation (Gittelman, 2007; Kessler & Chakrabarti, 1996), the occurrence of knowledge spillovers between organizations or industries (Thompson & Fox-Kean, 2005), or the adoption rates of new technologies (Toole, Cha, & González, 2012). In this research field, the multiple-

membership survival model can be used to detect industry-specific influences on the likelihood of developing innovations, and even geographical clustering of innovation potential.

Limitations and Suggestions for Future Research

Firstly, a continuing challenge of multiple-membership analysis lies in the computation of membership weights. In our example, the membership weights were framed as membership intensities and operationalized as the proportion of total working time spent in each organization. Even though our model implicitly accounted for mobility across higher-level units—by allowing time-varying membership sets and membership intensities—it leaves scope for an investigation that explores whether the mobility process between organizations is indeed driven by organizational influences. It is not unthinkable that employees who are driven towards attrition due to untenable organizational conditions, are more likely to become mobile, and will most likely want to move to organizations with a more rewarding climate. The use of membership transitions could add interesting information to the role of multiple memberships in survival models. For instance, it could allow us to study how contextual mobility might serve as a protective mechanism against job attrition or turnover. However, adding higher-order unit transition matrices still poses as a methodologically daunting task.

Secondly, in our multiple-membership survival model the frailty terms are predicted using the complete dataset, which causes them to be stable over time. However, this might not always be a realistic assumption given that organizations are dynamic entities which can, for example, change their policy to prevent the occurrence of less desirable events (e.g., early-career attrition). In our sample data, the time dimension in our data (five years) was too short to efficiently take time-dependent frailty terms into account. However, given survival data with a larger time horizon, frailty terms can be made time-specific which in turn could yield more insights about how the higher-order frailty factors evolve across time.

Thirdly, this article solely focuses on a traditional survival regression setting with a single, terminal event, which means that once the event has been experienced, it cannot be experienced again. The survival analysis discipline more generally has developed several extensions of this model, however, allowing for settings where (i) the same event can be experienced multiple times, such as promotions and successful innovations; (ii) there are multiple event types which can serve as terminal event—for instance turnover decisions leading either to internal role-transitions (i.e., other profession, same employer), inter-organizational mobility (i.e., same profession, changing employer), attrition (i.e., changing profession and employer), or retirement/lay-off (i.e., job quitting resulting in unemployment); and (iii) the existence of transition states between the start of the study and the terminal endpoint, for instance employees' experiencing (one or several episodes of) burnout-related experiences, absenteeism, or counter-productive work behavior before the attrition or job change event. The former extension is known as recurrent event modeling (for a detailed description and implementation in R, see Cook & Lawless, 2007), whereas the latter two are, respectively, called competing risk modeling and multistate modeling (Beyersmann, Allignol, & Schumacher, 2012). The multiple-membership survival regression model presented in this article can be used for several of these extended survival regressions, due to the count regression set-up. However, there is an ongoing need to incorporate some of the specific assumptions related to these extended models into our modeling framework.

It is our hope that other researchers will extend our model further along these lines, adding more relevant applications to the methodological toolbox of our field. For the time being, we hope that the broader use of the multiple-membership survival analysis methodology provided in this article will substantially improve our ability to form contextualized insights relating to event occurrences in the fields of organizational behavior and management.

Appendix 1: The Equivalence between Survival Regression and Poisson Likelihoods

An interval-search function to link each interval with the most recent time-observation. It is assumed that the covariate change did occur at the start of the interval in which this change has been observed. The following R-code can be used to achieve the merge of both datasets:

```
> tintmerge<-function(data1,data2){
  search<-function(i,data1,data2,id='id',time='time'){
    line<-(data2[,id]==data1[i,id] &
      data2[,time] <= data1[i,'tstop'])
    cbind(data1[i,],tail(data2[line,c(setdiff(names(data2),
      c('id','time')))]),1))
  }
  out<-lapply(1:dim(dat)[1],search,data1,data2,time='time')
  newdat<-do.call(rbind,out)
  return(newdat)
}
```

Appendix 2: The Equivalence between Survival Regression and Poisson Likelihoods

The likelihood contributions for survival regression can be written as

$$\mathcal{L}_i = f(t_i) = \lambda(t_i)S(t_i) \quad (5)$$

for those individuals experiencing the events and

$$\mathcal{L}_i = S(t_i) \quad (6)$$

for those who did not experience the event. The combination of both parts leads to the definition of the full likelihood contribution

$$\mathcal{L}_i = [\lambda(t_i)S(t_i)]^{\delta_i} [S(t_i)]^{1-\delta_i} \quad (7)$$

where δ_i is the event indicator which takes the value one if subject i has experienced the event and zero otherwise. Using the fact that the survival density can be expressed as

$$S(t_i) = \exp \left\{ - \int_0^{t_i} \lambda(s) ds \right\} \quad (8)$$

we can rewrite the likelihood contribution as:

$$\mathcal{L}_i = \lambda(t)^{\delta_i} \exp \left\{ - \int_0^{t_i} \lambda(s) ds \right\}. \quad (9)$$

For a piecewise exponential survival distribution, the time axis is partitioned in I mutually exclusive and exhaustive intervals. We choose I to be equal to Q , being the total number of distinct observed failure times. This approach increases the flexibility of the baseline hazard estimation as it approximates the non-parametric baseline hazard using piecewise constant hazards. Let $l_q := t_q - t_{q-1}$ be the interval length. Modeling the linear predictor term as $x_i^t \beta$, where x_i^t is a vector of covariates and β is a parameter vector. Hence, we can rewrite the likelihood contribution of subject i , whom experienced the event or was censored in the interval k_i , as:

$$\mathcal{L}_i(\beta) = [\lambda_{0q} \exp(x_i^t \beta)]^{\delta_i} \exp \left\{ - \int_0^{k_i} \lambda_{0q} l_q \exp(x_i^t \beta) dq \right\} \quad (10)$$

which can be rewritten as:

$$\mathcal{L}_i(\beta) = \prod_{q=0}^{k_i} [\lambda_{0q} \exp(x_i^t \beta)]^{\delta_i} \exp \{ - \lambda_{0q} l_q \exp(x_i^t \beta) \} \quad (11)$$

The full likelihood then becomes:

$$\mathcal{L}(\beta) = \prod_{i=1}^N \prod_{q=0}^{k_i} [\lambda_{0q} \exp(x_i^t \beta)]^{\delta_i} \exp \{ - \lambda_{0q} l_q \exp(x_i^t \beta) \} \quad (12)$$

$$\mathcal{L}(\beta) \propto \prod_{i=1}^N \prod_{q=0}^{k_i} [\lambda_{0q} l_q \exp(x_i^t \beta)]^{\delta_i} \exp \{ - \lambda_{0q} l_q \exp(x_i^t \beta) \} \quad (13)$$

Since the Poisson likelihood can be written as

$$\mathcal{L}(y_i) = \prod_{i=1}^N \frac{\lambda^{y_i} \exp(-\lambda)}{y_i!} \quad (14)$$

the proportionality of the survival likelihood and the Poisson likelihood with the event indicator δ_i as the dependent variable can be easily seen. Hence, a proportional hazards survival model can be estimated using a Poisson likelihood with δ_i as a binary outcome variable and l_q as an offset. Rabe-Hesketh & Skrondal (2012) suggested to model the baseline hazard, $\lambda_0(t)$, as a smooth function of time. In practice, a fourth-order polynomial of time has been used to model the baseline hazard function (Elghafghuf et al., 2014a; Elghafghuf et al., 2014b).

Authors' Note

We gratefully acknowledge the Flemish Department of Education for the permission to use the administrative databases of higher education and teaching assignments in secondary education.

Declaration of Conflicting Interests

The author(s) declared no potential conflicts of interest with respect to the research, authorship, and/or publication of this article.

Funding

This research was partially funded by an FWO project grant (G074418N) and a C1 grant from the KULeuven (DON-D2880-C14/17/014).

Notes

1. Herein, we rely on the established mathematical equivalence of the likelihood expressions of the survival model and the count regression model with a Poisson distribution—first established by Laird and Olivier (1981) and extended later to include multilevel designs (Feng, Nie, & Wolfe, 2009; Feng, Wolfe, & Port, 2005; Ma, Krewski, & Burnett, 2003)—to cut back on the computational details of the model.

References

- Aguinis, H., Boyd, B. K., Pierce, C. A., & Short, J. C. 2011. Walking New Avenues in Management Research Methods and Theories: Bridging Micro and Macro Domains. *Journal of Management*, 37(2): 395-403.
- Alexander, B., Hubers, C., Schwanen, T., Dijst, M., & Ettema, D. 2011. Anything, anywhere, anytime? Developing indicators to assess the spatial and temporal fragmentation of activities. *Environment and Planning B: Planning and Design*, 38(4): 678-705.
- Allen, D. G., Hancock, J. I., Vardaman, J. M., & McKee, D. L. N. 2014. Analytical mindsets in turnover research. *Journal of Organizational Behavior*, 35(1): 61-86.
- Ashforth, B. E., Harrison, S. H., & Corley, K. G. 2008. Identification in Organizations: An Examination of Four Fundamental Questions. *Journal of Management*, 34(3): 325-374.
- Bakker, R. M., & Shepherd, D. A. 2017. Pull the Plug or Take the Plunge: Multiple Opportunities and the Speed of Venturing Decisions in the Australian Mining Industry. *Academy of Management Journal*, 60(1): 130-155.
- Barnett, W. P., & Miner, A. S. 1992. Standing on the Shoulders of Others: Career Interdependence in Job Mobility. *Administrative Science Quarterly*, 37(2): 262-281.
- Beyersmann, J., Allignol, A., & Schumacher, M. 2012. *Competing Risks and Multistate Models with R*. New York: Springer.
- Blevins, D. P., Tsang, E. W. K., & Spain, S. M. 2015. Count-Based Research in Management: Suggestions for Improvement. *Organizational Research Methods*, 18(1): 47-69.
- Bliese, P. D., & Hanges, P. J. 2004. Being both too liberal and too conservative: The perils of treating grouped data as though they were independent. *Organizational Research Methods*, 7(4): 400-417.
- Bliese, P. D., & Ployhart, R. E. 2002. Growth Modeling Using Random Coefficient Models: Model Building, Testing, and Illustrations. *Organizational Research Methods*, 5(4): 362-387.

- Bürkner, P.-C. 2017. brms: An R Package for Bayesian Multilevel Models using Stan. *Journal of Statistical Software*, 80(1): 1-28.
- Cafri, G., Hedeker, D., & Aarons, G. A. 2015. An Introduction and Integration of Cross-Classified, Multiple Membership, and Dynamic Group Random-Effects Models. *Psychological Methods*, 20(4): 407-421.
- Carpenter, B., Gelman, A., Hoffman, M. D., Lee, D., Goodrich, B., Betancourt, M., Brubaker, M., Guo, J., Li, P., & Riddell, A. 2017. Stan: A Probabilistic Programming Language. *Journal of Statistical Software*, 76(1): 1-32.
- Chen, G., Smith, T. A., Kirkman, B. L., Zhang, P., Lemoine, G. J., & Farh, J.-L. 2019. Multiple Team Membership and Empowerment Spillover Effects: Can Empowerment Processes Cross Team Boundaries? *Journal of Applied Psychology*, 104(3): 321-340.
- Cheramie, R. A., Sturman, M. C., & Walsh, K. 2007. Executive career management: Switching organizations and the boundaryless career. *Journal of Vocational Behavior*, 71(3): 359-374.
- Chung, H., & Beretvas, S. N. 2012. The impact of ignoring multiple membership data structures in multilevel models. *British Journal of Mathematical and Statistical Psychology*, 65(2): 186-200.
- Claeskens, G., & Hjort, N. L. 2008. *Model selection and model averaging*: Cambridge University Press Cambridge.
- Claussen, J., Grohsjean, T., Luger, J., & Probst, G. 2014. Talent management and career development: What it takes to get promoted. *Journal of World Business*, 49(2): 236-244.
- Cook, R. J., & Lawless, J. F. 2007. *The Statistical Analysis of Recurrent Events*. New York: Springer.
- Cox, D. R. 1972. Regression Models and Life-Tables (with discussion). *Journal of the Royal Statistical Society. Series B (Methodological)*, 34(2): 187-220.
- DeSarbo, W. S., & Grewal, R. 2008. Hybrid strategic groups. *Strategic Management Journal*, 29(3): 293-317.
- Duchateau, L., & Janssen, P. 2008. *The Frailty Model*: Springer.
- Eckerd, S., Boyer, K. K., Qi, Y., Eckerd, A., & Hill, J. A. 2016. Supply Chain Psychological Contract Breach: An Experimental Study across National Cultures. *Journal of Supply Chain Management*, 52(3): 68-82.
- Edwards, J. R., & Parry, M. E. 2018. On the Use of Spline Regression in the Study of Congruence in Organizational Research. *Organizational Research Methods*, 21(1): 68-110.
- Eisenhardt, K. M., & Schoonhoven, C. B. 1990. Organizational Growth: Linking Founding Team, Strategy, Environment, and Growth Among U.S. Semiconductor Ventures, 1978-1988. *Administrative Science Quarterly*, 35(3): 504-529.
- Elghafghuf, A., Dufour, S., Reyher, K., Dohoo, I., & Stryhn, H. 2014a. Survival analysis of clinical mastitis data using a nested frailty Cox model fit as a mixed-effects Poisson model. *Preventive Veterinary Medicine*, 117(3): 456-468.

- Elghafghuf, A., & Stryhn, H. 2017. Robust Poisson likelihood estimation for frailty Cox models: A simulation study. *Communications in Statistics - Simulation and Computation*, 46(4): 2907-2923.
- Elghafghuf, A., Stryhn, H., & Waldner, C. 2014b. A cross-classified and multiple membership Cox model applied to calf mortality data. *Preventive Veterinary Medicine*, 115(1): 29-38.
- Felps, W., Mitchell, T. R., Hekman, D. R., Lee, T. W., Holtom, B. C., & Harman, W. S. 2009. Turnover contagion: How coworkers' job embeddedness and job search behaviors influence quitting. *Academy of Management Journal*, 52(3): 545-561.
- Feng, S., Nie, L., & Wolfe, R. A. 2009. Laplace's approximation for relative risk frailty models. *Lifetime data analysis*, 15(3): 343-356.
- Feng, S., Wolfe, R. A., & Port, F. K. 2005. Frailty survival model analysis of the national deceased donor kidney transplant dataset using Poisson variance structures. *Journal of the American Statistical Association*, 100(471): 728-735.
- Gelman, A., & Hill, J. 2006. *Data analysis using regression and multilevel/hierarchical models*. New York: Cambridge University Press.
- Gittelman, M. 2007. Does Geography Matter for Science-Based Firms? Epistemic Communities and the Geography of Research and Patenting in Biotechnology. *Organization Science*, 18(4): 724-741.
- Goldstein, H. 2011. *Multilevel Statistical Models* (4 ed.): John Wiley & Sons.
- Grady, M. W., & Beretvas, S. N. 2010. Incorporating student mobility in achievement growth modeling: A cross-classified multiple membership growth curve model. *Multivariate Behavioral Research*, 45(3): 393-419.
- Griep, Y., Vantilborgh, T., Baillien, E., & Pepermans, R. 2016. The mitigating role of leader-member exchange when perceiving psychological contract violation: a diary survey study among volunteers. *European Journal of Work and Organizational Psychology*, 25(2): 254-271.
- Griffeth, R. W., Hom, P. W., & Gaertner, S. 2000. A meta-analysis of antecedents and correlates of employee turnover: Update, moderator tests, and research implications for the next millennium. *Journal of Management*, 26(3): 463-488.
- Hastie, T., Tibshirani, R., & Friedman, J. 2009. *The Elements of Statistical Learning: Data Mining, Inference, and Prediction* (2 ed.): Springer.
- Hitt, M. A., Beamish, P. W., Jackson, S. E., & Mathieu, J. E. 2007. Building theoretical and empirical bridges across levels: Multilevel research in management. *Academy of Management Journal*, 50(6): 1385-1399.
- Hoffman, M. D., & Gelman, A. 2014. The No-U-Turn Sampler: Adaptively Setting Path Lengths in Hamiltonian Monte Carlo. *Journal of Machine Learning Research*, 15(1): 1593-1623.
- Hox, J. J. 2010. *Multilevel Analysis: Techniques and Applications* (2 ed.). UK: Routledge.
- Joo, H., Aguinis, H., & Bradley, K. J. 2017. Not All Nonnormal Distributions Are Created Equal: Improved Theoretical and Measurement Precision. *Journal of Applied Psychology*, 102(7): 1022-1053.

- Kammeyer-Mueller, J., Wanberg, C. R., Rubenstein, A., & Song, Z. 2013. Support, undermining, and newcomer socialization: Fitting in during the first 90 days. *Academy of Management Journal*, 56(4): 1104-1124.
- Kessler, E. H., & Chakrabarti, A. K. 1996. Innovation Speed: A Conceptual Model of Context, Antecedents, and Outcomes. *The Academy of Management Review*, 21(4): 1143-1191.
- Laird, N., & Olivier, D. 1981. Covariance analysis of censored survival data using log-linear analysis techniques. *Journal of the American Statistical Association*, 76(374): 231-240.
- Lamote, C., Van Damme, J., Van Den Noortgate, W., Speybroeck, S., Boonen, T., & de Bilde, J. 2013. Dropout in secondary education: an application of a multilevel discrete-time hazard model accounting for school changes. *Quality & Quantity*, 47(5): 2425-2446.
- Leroy, S. 2009. Why is it so hard to do my work? The challenge of attention residue when switching between work tasks. *Organizational Behavior and Human Decision Processes*, 109(2): 168-181.
- Luo, W., & Kwok, O.-m. 2012. The consequences of ignoring individuals' mobility in multilevel growth models: A Monte Carlo Study. *Journal of Educational and Behavioral Statistics*, 37(1): 31-56.
- Ma, R., Krewski, D., & Burnett, R. T. 2003. Random effects Cox models: a Poisson modelling approach. *Biometrika*, 90(1): 157-169.
- Manda, S. O. M., Feltbower, R. G., & Gilthorpe, M. S. 2012. A Multivariate Random Frailty Effects Model for Multiple Spatially Dependent Survival Data. In Y.-K. Tu, & D. C. Greenwood (Eds.), *Modern Methods for Epidemiology*: 157-172. Dordrecht: Springer Netherlands.
- Mathieu, J. E., & Chen, G. 2011. The Etiology of the Multilevel Paradigm in Management Research. *Journal of Management*, 37(2): 610-641.
- McCabe, D. L., & Trevino, L. K. 1997. Individual and contextual influences on academic dishonesty: A multicampus investigation. *Research in higher education*, 38(3): 379-396.
- Meyers, J. L., & Beretvas, S. N. 2006. The impact of inappropriate modeling of cross-classified data structures. *Multivariate Behavioral Research*, 41(4): 473-497.
- Moore, D. 2016. *Applied survival analysis using R*. Switzerland: Springer.
- Morgeson, F. P., Mitchell, T. R., & Liu, D. 2015. Event System Theory: An Event-Oriented Approach to the Organizational Sciences. *Academy of Management Review*, 40(4): 515-537.
- Morita, J. G., Lee, T. W., & Mowday, R. T. 1989. Introducing Survival Analysis To Organizational Researchers: A Selected Application To Turnover Research. *Journal of Applied Psychology*, 74(2): 280-292.
- Morita, J. G., Lee, T. W., & Mowday, R. T. 1993. The Regression-Analog to Survival Analysis: A Selected Application to Turnover Research. *Academy of Management Journal*, 36(6): 1430-1464.

- Nikolaeva, R., Bhatnagar, A., & Ghose, S. 2015. Exploring Curvilinearity Through Fractional Polynomials in Management Research. *Organizational Research Methods*, 18(4): 738-760.
- Nyberg, A. 2010. Retaining Your High Performers: Moderators of the Performance-Job Satisfaction-Voluntary Turnover Relationship. *Journal of Applied Psychology*, 95(3): 440-453.
- O'Leary, M. B., Mortensen, M., & Woolley, A. W. 2011. Multiple Team Membership: A Theoretical Model of its Effects on Productivity and Learning for Individuals and Teams. *Academy of Management Review*, 36(3): 461-478.
- Pierce, J. R., & Aguinis, H. 2013. The Too-Much-of-a-Good-Thing Effect in Management. *Journal of Management*, 39(2): 313-338.
- R Core Team. 2017. R: A Language and Environment for Statistical Computing. R Foundation for Statistical Computing.
- Robinson, W. S. 1950. Ecological Correlations and the Behavior of Individuals. *American Sociological Review*, 15(3): 351-357.
- Ross, L. 1977. The Intuitive Psychologist And His Shortcomings: Distortions in the Attribution Process. In B. Leonard (Ed.), *Advances in Experimental Social Psychology*, Vol. Volume 10: 173-220: Academic Press.
- Rousseau, D. M. 1989. Psychological and implied contracts in organizations. *Employee Responsibilities and Rights Journal*, 2(2): 121-139.
- Rousseau, D. M. 1998. Why Workers Still Identify with Organizations. *Journal of Organizational Behavior*, 19(3): 217-233.
- Schoonhoven, C. B., Eisenhardt, K. M., & Lyman, K. 1990. Speeding Products to Market: Waiting Time to First Product Introduction in New Firms. *Administrative Science Quarterly*, 35(1): 177-207.
- Scott, M. A., Simonoff, J. S., & Marx, B. D. (Eds.). 2013. *The SAGE handbook of multilevel modeling*. London: SAGE.
- Smith, L. J. W., & Beretvas, S. N. 2014. The Impact of Using Incorrect Weights With the Multiple Membership Random Effects Model. *Methodology*, 10(1): 31-42.
- Smith, L. J. W., & Beretvas, S. N. 2017. A comparison of techniques for handling and assessing the influence of mobility on student achievement. *The Journal of Experimental Education*, 85(1): 3-23.
- Snijders, T. A. B., & Bosker, R. J. 2012. *Multilevel Analysis: An Introduction to Basic and Advanced Multilevel Modeling* (2 ed.). London: Sage Publications.
- Solinger, O. N., Hofmans, J., Bal, P. M., & Jansen, P. G. W. 2016. Bouncing back from psychological contract breach: How commitment recovers over time. *Journal of Organizational Behavior*, 37(4): 494-514.
- Sun, S., & Pan, W. 2014. A methodological review of statistical methods for handling multilevel non-nested longitudinal data in educational research. *International Journal of Research & Method in Education*, 37(3): 285-308.
- Therneau, T. M., & Grambsch, P. M. 2000. *Modeling Survival Data: Extending the Cox Model*. New York: Springer.

- Thompson, P., & Fox-Kean, M. 2005. Patent Citations and the Geography of Knowledge Spillovers: A Reassessment. *American Economic Review*, 95(1): 450-460.
- Toole, J. L., Cha, M., & González, M. C. 2012. Modeling the Adoption of Innovations in the Presence of Geographic and Media Influences. *PLoS ONE*, 7(1): 1-9.
- Uzunca, B. 2018. A Competence-Based View of Industry Evolution: The Impact of Submarket Convergence on Incumbent-Entrant Dynamics. *Academy of Management Journal*, 61(2): 738-768.
- Van Droogenbroeck, F., Spruyt, B., & Vanroelen, C. 2014. Burnout among senior teachers: Investigating the role of workload and interpersonal relationships at work. *Teaching and Teacher Education*, 43(1): 99-109.
- Vantilborgh, T., Bidee, J., Pepermans, R., Griep, Y., & Hofmans, J. 2016. Antecedents of Psychological Contract Breach: The Role of Job Demands, Job Resources, and Affect. *PLoS ONE*, 11(5): 1-22.
- Vehtari, A., Gelman, A., & Gabry, J. 2017. Practical Bayesian model evaluation using leave-one-out cross-validation and WAIC. *Statistics and Computing*, 27(5): 1413-1432.
- Vergauwe, J., Wille, B., Hofmans, J., Kaiser, R. B., & De Fruyt, F. 2017. The Too Little/Too Much Scale: A New Rating Format for Detecting Curvilinear Effects. *Organizational Research Methods*, 20(3): 518-544.
- Wanberg, C. R., Hough, L. M., & Song, Z. 2002. Predictive Validity of a Multidisciplinary Model of Reemployment Success. *Journal of Applied Psychology*, 87(6): 1100-1120.
- Wanberg, C. R., Kanfer, R., Hamann, D. J., & Zhang, Z. 2016. Age and Reemployment Success After Job Loss: An Integrative Model and Meta-Analysis. *Psychological Bulletin*, 142(4): 400-426.
- Wang, H., Hall, N. C., & Rahimi, S. 2015. Self-efficacy and causal attributions in teachers: Effects on burnout, job satisfaction, illness, and quitting intentions. *Teaching and Teacher Education*, 47(1): 120-130.
- Watanabe, S. 2013. A Widely Applicable Bayesian Information Criterion. *Journal of Machine Learning Research*, 14(1): 867-897.
- Wood, S. N. 2013. A simple test for random effects in regression models. *Biometrika*, 100(4): 1005-1010.
- Wood, S. N. 2017. *Generalized additive models: an introduction with R* (2 ed.). New York: Chapman Hall/CRC.
- Wood, S. N., Goude, Y., & Shaw, S. 2015. Generalized additive models for large data sets. *Journal of the Royal Statistical Society: Series C (Applied Statistics)*, 64(1): 139-155.
- Wright, C. 2009. Inside Out? Organizational Membership, Ambiguity and the Ambivalent Identity of the Internal Consultant. *British Journal of Management*, 20(3): 309-322.
- Zaccaro, S. J., Marks, M. A., & DeChurch, L. 2012. *Multiteam systems: An organization form for dynamic and complex environments*. London: Routledge Academic.
- Zappa, P., & Lomi, A. 2015. The analysis of multilevel networks in organizations: models and empirical tests. *Organizational Research Methods*, 18(3): 542-569.




Figure 1. Baseline survival function for teachers during the first five years in the profession.

Notes: The grey staircase function indicates the (data-driven) Kaplan-Meier survival curve (equivalent to an intercept-only Cox model). The shaded area corresponds to the 95% confidence interval of the data-driven baseline survival.

The full line links the estimated survival rates from the count regression model, with an underlying smoothing-spline-based baseline hazard function.




Figure 2. Higher-Level (Random) Effects: Comparison between Single Frailty estimates and Multiple-Membership estimates.

Notes: The black line corresponds to the ascendingly ranked multiple-membership random effects (1038 schools). Grey points are the random effects for the single-frailty model taking the first school into account (712 schools).




Figure 3. Total Higher-Level Effects: Comparison between Single Frailties and Multiple-Membership Frailties.




Figure 4. Smooth Covariate Effects on the Log-Hazard.

Note: The shaded region represents the 95% pointwise confidence interval.

Table 1. Illustration of Event Data: time-invariant event data and time-varying information.

<i>Time-Invariant Information</i>																
ID	Time	Event	StartAge	Gender												
1	21	1	23	0												
2	4	0	24	1												
...	...															
<i>Time-Varying Information</i>																
ID	Time	WrkHour	Contract	S1	S2	S3	S4	S5	S6	W1	W2	W3	W4	W5	W6	
1	0	10	Interim	A	B	NA	NA	NA	NA	.60	.40	NA	NA	NA	NA	
1	1	22	Interim	A	B	NA	NA	NA	NA	.50	.50	NA	NA	NA	NA	
1	4	19	Interim	A	B	NA	NA	NA	NA	.47	.53	NA	NA	NA	NA	
1	6	12	Interim	A	B	NA	NA	NA	NA	.42	.58	NA	NA	NA	NA	
1	7	28	Interim	A	B	C	D	NA	NA	.18	.25	.32	.25	NA	NA	
1	8	16	Interim	C	D	NA	NA	NA	NA	.56	.44	NA	NA	NA	NA	
1	12	20	Interim	E	NA	NA	NA	NA	NA	1	NA	NA	NA	NA	NA	
1	13	20	Interim	F	NA	NA	NA	NA	NA	1	NA	NA	NA	NA	NA	
1	18	20	Vacant	F	NA	NA	NA	NA	NA	1	NA	NA	NA	NA	NA	
1	19	20	Interim	A	B	NA	NA	NA	NA	.30	.70	NA	NA	NA	NA	
2	0	20	Interim	G	H	NA	NA	NA	NA	.40	.60	NA	NA	NA	NA	
2	2	12	Interim	H	NA	NA	NA	NA	NA	1	NA	NA	NA	NA	NA	
...																

Notes: The values in the S#-columns (i.e., starting with an S) list the identification numbers of each environment and have no numerical interpretation.

The W#-columns (i.e., starting with a W) are the proportions of total working time spent in each of the respective environments.

Table 3. Descriptive Statistics and Correlations for All Turnover Predictors.

		N	M	SD	1	2	3	4a
1	Gender (ref=female)	1,642	.29	-				
2	Starting Age	1,642	23.00	2.13	.10			
3	Number of hours taught	80,357	20.60	4.74	.02	-.16		
4a	Contract type: interim (ref: vacant)	80,357	.55	-	-.03	.02	.00	
4b	Contract type: tenure (ref: vacant)	80,357	.12	-	.00	-.05	.07	-.40

Note: Correlations with an absolute magnitude of .01 or greater are significant at a 99% confidence level.

Table 4. Hazard Regression Model: Comparison between Cox model and Count Regression.

Variable	Cox Model		Count Regression	
	Estimate	Std.Error	Estimate	Std.Error
StartAge	.069	(.016)***	.066	(.016)***
Gender (ref=female)	-.017	(.106)	-.000	(.106)
WorkHours	-.113	(.008)***	-.103	(.008)***
Contract type: interim (ref: vacant)	.603	(.111)***	.540	(.111)***
Contract type: tenure (ref: vacant)	-2.142	(.730)***	-2.107	(.729)***

Notes: p-values: ‘ ’ 0.10 ‘*’ 0.05 ‘***’ 0.01 ‘****’

Table 5. Frailty Models: Results and Comparison between Different Nesting Structures.

	Most Prominent School (at that time)		Most Prominent School (during whole career)		First School Ever		Last School	
Number of schools	958 (92.3%)		720 (69.4%)		712 (68.6%)		710 (68.4%)	
<i>COXPH</i>	Estimate	Std.Error	Estimate	Std.Error	Estimate	Std.Error	Estimate	Std.Error
StartAge	.070	(.016)***	.072	(.017)***	.071	(.017)***	.071	(.017)***
Gender (ref=female)	-.013	(.106)	-.009	(.107)	-.019	(.108)	-.011	(.107)
WorkHours	-.114	(.008)***	-.114	(.008)***	-.114	(.008)***	-.114	(.008)***
Contract type: interim (ref: vacant)	.600	(.112)***	.602	(.112)***	.596	(.112)***	.604	(.112)***
Contract type: tenure (ref: vacant)	-2.144	(.730)***	-2.147	(.730)***	-2.155	(.730)***	-2.144	(.730)***
Standard Deviation (SD) Frailty	.255		.286		.346		.285	
Null Model: SD Frailty	.427		.377		.457		.395	
AIC (REML)	6,061		6,060		6,057		6,060	
<i>COUNT REGRESSION</i>	Estimate	Std.Error	Estimate	Std.Error	Estimate	Std.Error	Estimate	Std.Error
StartAge	.071	(.016)***	.073	(.017)***	.073	(.017)***	.072	(.017)***
Gender (ref=female)	-.010	(.106)	-.007	(.107)	-.015	(.108)	-.008	(.107)
Workhours	-.112	(.008)***	-.112	(.008)***	-.112	(.008)***	-.112	(.008)***
Contract type: interim (ref: vacant)	0.612	(.111)***	.614	(.111)***	.611	(.112)***	.617	(.111)***
Contract type: tenure (ref: vacant)	-2.218	(.721)***	-2.220	(.721)***	-2.224	(.721)***	-2.218	(.721)***
Standard Deviation (SD) Frailty	.221		.266		.339		.283	
Null Model: SD Frailty	.421		.369		.445		.414	
AIC (REML)	4,899		4,898		4,895		4,898	

Notes: p-values: ‘ ’ 0.10 ‘*’ 0.05 ‘***’ 0.01 ‘****’

Table 6. Step-By-Step Model Building of a Multiple-Membership Survival Model.

	No Frailty	Multiple-Membership Frailty	MM Frailty with covariates	MM Frailty with higher-level covariates	MM Frailty with Smooth Effects
Subject-level covariates	B (SE)	B (SE)	B (SE)	B (SE)	B (SE)
StartAge	–	–	.07 (.02)*	.04 (.02)*	.19 (0.40) SD 3.43*
Gender (ref=female)	–	–	–.01 (.10)	.03 (.10)	–.06 (.10)
WorkHours	–	–	–.11 (.01)*	–.08 (.01)*	–.10 (.35) SD 2.44*
Contract type: interim (ref: vacant)	–	–	.62 (.12)*	.52 (.11)*	.50 (.11)*
Contract type: tenure (ref: vacant)	–	–	–2.56 (.83)*	–2.49 (.83)*	–2.49 (.82)*
Environment-level covariates					
Early-career Attrition Climate	–	–	–	.04 (.00)*	3.85 (.90)* SD 10.79*
Random Effect					
Standard Deviation (SD) Frailty	–	.43 (.16)*	.26 (.15)*	.23 (.11)*	.06 (.05)
Model Fit Statistics					
WAIC	5,244	5,240	4894	4480	4192

Notes: ‘*’ indicates that the 95% credible interval does not contain zero.