

Uncovering life-course patterns with causal discovery and survival analysis

Bojan Kostic^{a,*}, Romain Crastes dit Sourd^b, Stephane Hess^c, Joachim Scheiner^d, Christian Holz-Rau^d, Francisco C. Pereira^a

^aMachine Learning for Smart Mobility (MLSM) Group, Transport Division, Department of Technology, Management and Economics, Technical University of Denmark, Lyngby, Denmark

^bCentre for Decision Research, Leeds University Business School, University of Leeds, Leeds, UK

^cChoice Modelling Centre, Institute for Transport Studies, University of Leeds, Leeds, UK

^dDepartment of Transport Planning, Faculty of Spatial Planning, TU Dortmund, Dortmund, Germany

Abstract. We provide a novel approach and an exploratory study for modelling life event choices and occurrence from a probabilistic perspective through causal discovery and survival analysis. Our approach is formulated as a bi-level problem. In the upper level, we build the life events graph, using causal discovery tools. In the lower level, for the pairs of life events, time-to-event modelling through survival analysis is applied to model time-dependent transition probabilities. Several life events were analysed, such as getting married, buying a new car, child birth, home relocation and divorce, together with the socio-demographic attributes for survival modelling, some of which are age, nationality, number of children, number of cars and home ownership. The data originates from a survey conducted in Dortmund, Germany, with the questionnaire containing a series of retrospective questions about residential and employment biography, travel behaviour and holiday trips, as well as socio-economic characteristic. Although survival analysis has been used in the past to analyse life-course data, this is the first time that a bi-level model has been formulated. The inclusion of a causal discovery algorithm in the upper-level allows us to first identify causal relationships between life-course events and then understand the factors that might influence transition rates between events. This is very different from more classic choice models where causal relationships are subject to expert interpretations based on model results.

Keywords: Life course calendar, Causal discovery, Survival analysis

1. Introduction and Background

Inspired by a few seminal studies [1, 2] highlighting how different life-course decisions are interrelated and should be analysed jointly, several studies [3] have analysed travel behaviour as a long-term life-course dimension, similar to employment spells or residential location choices. Retrospective surveys such as the one proposed by Beige and Axhausen [4] seek to uncover how life events, such as the birth of a child or marriage, are connected to one another as well as to long term mobility decisions such as commuting mode or relocation. Retrospective surveys are

*Corresponding author. E-mail address: boko@dtu.dk (B. Kostic).

generally based on a life-course calendar, where the respondents are asked a series of questions for each year in retrospect (which is a grid analogous to Figure A1). Life-course calendars are commonly used to collect such information where panel data is unavailable or too costly. Life-course calendars have become popular because they have two main advantages over traditional panel surveys: the visual and mental relation of different kinds of events improves the quality of the data by providing reference points and preventing time inconsistencies. Moreover, the visual aid of the calendar tool eases the task of listing a potentially high number of different short events, a task that would be more difficult in traditional surveys [5]. However, while scholars agree on the fact that life-course calendars provide a helpful tool to facilitate the recall of life events, many acknowledge the issue of recall of all relevant events and the effort required to do so. In particular, Manzoni et al. [6] and Manzoni [7] find that the recall depends not only on respondents' socio-demographic characteristics but also on the type of events that are recalled. For example, people tend to omit or reduce the duration of spells of unemployment. In addition, it has been shown by Schoenduwe et al. [8] that participants find it easiest to remember "emotionally meaningful, momentous, unique or unexpected events". It is precisely these kinds of events that we consider in our analysis.

The data for our research originates from a retrospective survey that was carried out between 2007 and 2012 at the Department of Transport Planning of the Faculty of Spatial Planning at TU Dortmund. The questionnaire contains a series of retrospective questions about residential and employment biography, travel behaviour and holiday trips, as well as socio-economic characteristics.

In our work, we try to investigate if causal relations between life events can be identified in our dataset, such that these can then be used for choice modelling and provide valuable insights into understanding people's behaviour. We do this through causal discovery, a tool that deals with the process of automatically finding, from available data, the causal relations between the available (random) variables. This typically follows a graphical representation, where nodes are the variables, and arcs are the causal relationships between them. A causal discovery algorithm thus has as main objective to find directed graphical causal models (DGCMs) that reflect the data. There is also research on acyclic graphs and combinations of both, which we will not focus on here. For a broad review, we recommend Glymour et al. [9].

The ideal way to discover causal relations is through interventions or randomized experiments. However, in many cases, such as ours, this is in fact impossible. Having only pure observational data, causal discovery tools search for systematic statistical properties, such as conditional independence, to establish (possible) causal relations. This implies also that we are limited to what the data *can tell* us about causality. If the true causality structure is not observable from the data, such methods will at most find correlations disguised as causation, as happens in many other non-causal methods, such as regression. In such cases, a possible solution relies on domain knowledge, such as postulating the existence of common causes (e.g. a latent variable), or the addition of constraints.

The concept of causal discovery thus consists of algorithms to discover causal relations by analyzing statistical properties of purely observational data. In this paper, we intend to apply such methods to automatically uncover causality through temporal sequence data. Time or sequence are hardly a true cause in decision making (e.g. one will not have children or buy a car just because X time has passed, but due to some deeper cause). Therefore, it is certainly a bit too bold to claim that there is inherent causality in all the relationships that we will find. Instead, we see this process as uncovering the most likely graph of sequential patterns throughout the life of individuals.

Given a resulting causal graph, together with a dataset of temporal sequence records, we can now estimate transition probabilities for each event pair from the graph. We model the transition

probabilities through survival analysis, which is typically used for time-to-event (TTE) modelling (or similarly for time-to-failure (TTF) or durations) [10]. Classical survival analysis is widely applied in medical studies to describe the survival of patients or cancer recurrence [11, 12, 13, 14, 15]. Beside medical applications, it is now applied far beyond this field, such as predictive maintenance, customer churn, credit risk, etc.

Through survival analysis, given that an initial event occurred, we can obtain probabilistic temporal estimates of another event occurrence, with the survival values defined as $S(t) = P(T > t)$. It represents the probability of the event happening after time t from the initial event (i.e. the probability of surviving past time t), given that the event did not happen before time t .

In our case, survival probabilities will give us estimates of how likely it is that a certain life event will occur in each coming year after another specific life event. This leads to a better understanding of people’s life-course calendars and how these depend on their socio-demographics. Examples include how the number of cars affects a decision of moving after getting married, the effect of the number of children on buying a car after moving, the difference between homeowners and non-homeowners when it comes to having a child after wedding, to name a few.

2. Methods

Our approach is formulated as a bi-level problem: in the upper level, we build the life events graph, with causal discovery tools, where each node represents a life event; in the lower level, for selected pairs of life events, time-to-event modelling through survival analysis was applied to model the time-dependent probabilities of event occurrence.

2.1. Causal Discovery

Two causal discovery algorithms used in our research are Parents and Children (PC) and Greedy Equivalence Search (GES). In a resulting life events graph, we should have a set of key events as nodes, while the arcs correspond to temporal “causality” between pairs of events, where $A \rightarrow B$ implies that A is precedent to B . As discussed earlier, this is not assumed as true causality; instead, we see it as representing a high likelihood of observing a certain temporal order in such events.

The concept of conditional independence is key here:

$$X \perp\!\!\!\perp Y|S \iff P(X, Y|S) = P(X|S)P(Y|S) \text{ for all dataset}$$

The idea is quite intuitive: if two variables X and Y are independent when controlling for the (set of) variable(s) S , then X and Y are conditionally independent given S . This means that S is likely to be a common cause for X and Y . Of course, just as testing for independence in general, testing for conditional independence is a hard problem [16], but approximate techniques exist to determine it from observational data [17].

Algorithms such as Parents and Children (PC) or its derivatives such as the Fast Causal Inference (FCI) [18] are based on this principle to generate a causal graph. They essentially start from a fully connected undirected graph, and gradually “cut” links between conditionally independent nodes.

Another approach starts from the opposite direction, where we start from an empty graph, and gradually add nodes and arcs. The most well-known algorithm is the Greedy Equivalence Search (GES) [19]. At each step in GES, a decision is made as to whether adding a directed edge

to the graph will increase fit measured by some fitting score (e.g. Bayesian information criterion (BIC), z-score). When the score can no longer be improved, the process reverses, by checking whether removing edges will improve the score. It is difficult to compare this with PC in terms of the end quality of the discovered graph, but it is known that, as the sample size increases, the two algorithms converge to the same Markov Equivalence Class [9]. The vast majority of causal discovery algorithms are variations or combinations of the two approaches just mentioned.

In our dataset, we have four events, namely *child birth*, *wedding*, *new car* and *moving*, and two state variables, namely *married* (True or False) and *children* (numeric). While the final life events graph shall have only events (the state variables will be useful later, but not as nodes in the graph), PC and GES allow for inclusion of these variables, therefore we applied both PC and GES algorithms [20].

2.2. Survival Analysis

In survival analysis, survival function $S(t)$ can be defined from the cumulative distribution function $F(t)$, as $S(t) = 1 - F(t)$, which in terms of probabilities can be formulated as $S(t) = P(T > t)$. The survival function can be interpreted in terms of the gradient (slope) of a line, with steeper lines meaning an event is more likely to happen in the corresponding interval. In contrast, a flat line means there is a zero probability of an event happening in that interval.

Most commonly used statistical methods for survival analysis are the Kaplan-Meier method [21] and the Nelson-Aalen method [22, 23], which are univariate counting-based methods for general insights. To estimate the effect of covariates, a survival regression model can be formalized through the Cox proportional hazards model (CPH) [24]. It assumes that all individuals share the same baseline hazard function, and the effect of covariates acts as its scaler. It provides interpretability but lacks the accuracy with larger datasets.

In survival analysis, time-to-event (TTE) is expressed through two events: the *birth* event and the *death* event (hence the term “survival”). Time between the two events is survival time, denoted by T , which is treated as a non-negative random variable. Given the *birth* event, it represents the probability distribution of the *death* event happening in time.

Observations of true survival times may not always be available, thus creating distinction between uncensored (true survival times available) and censored observations (true survival times unknown, but only the censorship time). This information is expressed through an event indicator E , with binary outcome. The censorship time is assumed to be non-informative, and T and E are independent.

Two main outputs of survival analysis are the survival function and the (cumulative) hazard function. The survival function $S(t)$ can be formulated starting from the cumulative distribution function (CDF), denoted by F , which can be defined as the probability (P) of an event occurring up to time t :

$$F(t) = P(T \leq t) = \int_0^t f(x)dx \quad (1)$$

with properties of being a non-decreasing function of time t , having bounds on the probability as $0 \leq F(t) \leq 1$, and its relation with the survival function as $F(t) = 1 - S(t)$. Given this, we can define the survival function as the probability of surviving past time t , i.e. the probability of an event happening after time t given that the event has not happened before or at time t as:

$$S(t) = P(T > t) = \int_t^\infty f(x)dx \quad (2)$$

with properties of being a non-increasing function of t , and having bounds on the survival probability as $0 \leq S(t) \leq 1$ and $S(\infty) = 0$, given its relation with the CDF as $S(t) = 1 - F(t)$. The probability density function (PDF) of T , $f(t)$, can be formalized as:

$$f(t) = \lim_{\Delta t \rightarrow 0} \frac{P(t < T \leq t + \Delta t)}{\Delta t} \quad (3)$$

The hazard function $h(t)$ defines the instantaneous hazard rates of the event happening at time t , given that the event has not happened before time t . It can be defined as a limit when a small interval Δt approaches 0, with a probability of the event happening between time t and $t + \Delta t$, as in (4). In addition, there is a direct correlation between the hazard function and the survival function.

$$h(t) = \lim_{\Delta t \rightarrow 0} \frac{P(t < T \leq t + \Delta t \mid T > t)}{\Delta t} = \frac{f(t)}{S(t)} = -\frac{d}{dt} \log(S(t)) \quad (4)$$

The cumulative hazard function $H(t)$ is also used, and is especially important for model estimation (5). The survival function can therefore be derived by solving the differential equation above, which gives the relation with the hazard and cumulative hazard functions and is expressed in (6).

$$H(t) = \int_0^t h(x) dx = -\log(S(t)) \quad (5)$$

$$S(t) = \exp\left(-\int_0^t h(z) dz\right) = \exp(-H(t)) \quad (6)$$

Survival analysis is typically conducted through univariate survival models or through survival regression. The former takes into account only survival times, while the latter also includes covariates. The most widely used regression model is the Cox proportional hazards (CPH) model [24]. In addition, there have been recent developments in applying machine learning models in survival analysis to improve prediction accuracy, but compromising on model interpretability. Machine learning algorithms, such as support vector machines [25] and random forests [26] are common examples of such techniques. Beside better accuracy, they also demonstrated a superior efficiency over existing training algorithms, especially for large datasets.

To exploit non-linear relationship among the covariates, neural networks (NN) [27] are used. The first time a simple neural network was proposed was by Faraggi and Simon [28]. The approach is similar to the CPH, but optimizes a slightly different partial likelihood function. More recent developments are the application of deep neural networks (DNNs), with models like DeepSurv [29], DeepHit [30] and Nnet-survival [31]. DeepSurv is a proportional hazards deep neural network. It is an extension of the method proposed by [28], used to connect more than one layer in the network structure. While CPH and DeepSurv rely on the strong proportional hazards assumption, other models relax this condition. DeepHit uses a DNN to learn the survival distribution directly. It imposes no assumptions regarding the underlying processes and allows for the effect of the covariates to change over time. Nnet-survival, on the other hand, is a discrete time survival model. The hazard rate, which for the previously discussed methods is calculated as the instantaneous risk, here is defined for a specific time interval.

Because of the relatively small dataset for deep learning applications and to focus on the interpretability, in this paper we used the univariate models and the CPH regression model.

2.2.1. Univariate survival models

The estimation of the survival and cumulative hazard functions can be performed using non-parametric univariate models. Survival and cumulative hazard functions are functions of survival times with event indicator, i.e. $S(t) = f(\mathbf{T}, \mathbf{E})$ and $H(t) = g(\mathbf{T}, \mathbf{E})$. Given the ordered survival times, for each instant of interest on the timeline, the number of individuals at risk and the number of death events occurring is determined. The survival function is usually estimated using the Kaplan-Meier method [21], a product-limit estimator. It takes as input survival times and estimates the survival function as:

$$\hat{S}(t) = \prod_{i:t_i \leq t} \left(1 - \frac{d_i}{n_i}\right) \quad (7)$$

where d_i is the number of death events at time t and n_i is the number of individuals at risk of death just prior and at time t . Similarly, the cumulative hazard function can be estimated using the Nelson-Aalen method [22, 23] as:

$$\hat{H}(t) = \sum_{i:t_i \leq t} \frac{d_i}{n_i} \quad (8)$$

Hazard and cumulative hazard values cannot be directly interpreted, but are used for relative comparison among themselves in discovering which timeline instants are more likely to encounter a death event than others.

Univariate non-parametric survival models are simple and can provide flexible survival and hazard curves. They provide population-level insight into survival behaviour, but not at the individual level.

2.2.2. Survival regression

To account for the effect of covariates on the output functions, the Cox proportional hazards model is typically used for survival regression. It is a semi-parametric regression model that estimates the risk, i.e. the log of the hazard rate, of an event occurring as a linear combination of the covariates as $\boldsymbol{\beta}^\top \mathbf{x}$, with $\boldsymbol{\beta}$ being a vector of coefficients and \mathbf{x} denoting a vector of time-constant covariates. It maximizes the Cox's partial likelihood function (of the coefficients) [24]:

$$L(\boldsymbol{\beta}) = \prod_{i:E_i=1} \frac{\exp(\boldsymbol{\beta}^\top \mathbf{x}_i)}{\sum_{j:T_j \geq T_i} \exp(\boldsymbol{\beta}^\top \mathbf{x}_j)} \quad (9)$$

where only uncensored data ($i : E_i = 1$) are used (hence the term "partial"). The corresponding negative log-likelihood function for minimization is given by (10). As implemented in [32], the optimization problem is solved using maximum partial likelihood estimation (MPLE) and Newton-Raphson iterative search method [33]. Ties in survival times are handled using the Efron method [34]. The L2 regularization is used with a scale value of 0.1.

$$l(\boldsymbol{\beta}) = -\log L(\boldsymbol{\beta}) = - \sum_{i:E_i=1} \left(\boldsymbol{\beta}^\top \mathbf{x}_i - \log \sum_{j:T_j \geq T_i} \exp(\boldsymbol{\beta}^\top \mathbf{x}_j) \right) \quad (10)$$

The hazard function is computed for each individual as a product of the population-level time-dependent baseline hazard $h_0(t)$ and the individual-level time-constant partial hazard expressed through the exponent of a linear combination of the covariates. It is formulated as:

$$h(t|\mathbf{x}) = h_0(t) \exp(\boldsymbol{\beta}^\top \mathbf{x}) \quad (11)$$

All individuals share the same baseline hazard, and the effect of covariates acts as a scaler through the partial hazard. The baseline hazard function is estimated using the Breslow method [35]. The exponential form ensures that the hazard rates are always positive. The ratio of the hazard functions of two individuals is always constant and it is given by the ratio of their partial hazards (hence the term “proportional”). Given the hazard function, cumulative hazard function and survival function can be computed using (5) and (6), respectively.

The goodness-of-fit measure is the concordance index (C-index) [36], a commonly used metric to validate the predictive capability of survival models (analogous to the R^2 score in regression). It is a rank correlation score, comparing the order of observed survival times from the data and order of estimated survival times (or negative hazard rates) from the model. It is formulated as:

$$C = \frac{1}{\varepsilon} \sum_{i:E_i=1} \sum_{j:T_j \geq T_i} \mathbf{1}_{f(\mathbf{x}_i) < f(\mathbf{x}_j)} \quad (12)$$

where ε is the number of samples used in calculation and $\mathbf{1}$ is the indicator function with binary outcome based on the given condition. The C-index is equivalent to the area under a ROC (Receiver Operating Characteristic) curve, and ranges from 0–1, with values less than 0.5 representing a poor model, 0.5 is equivalent to a random guess, 0.65–0.75 meaning a good model, and greater than 0.75 indicating a strong model, with 1 being the perfect concordance and 0 being the perfect anti-concordance.

3. Data elaboration and model fitting

The data used in model estimation contains both uncensored and censored observations. Uncensored observations for an individual are survival times available directly as the number of years between two specific events, with the event observed indicator set to 1. Censored observations are created when only the birth event happened, with survival times set to 20 years and the event observed indicator set to 0.

To fit the models and validate their performance, the dataset, which comprises 1,487 individuals, was split into the training set, containing 60 % of the samples, and the test set, as the hold-out set containing the remaining 40 % of the samples. Stratified splitting was used to ensure a balanced presence of covariates’ values in both datasets.

Data elaboration and modelling was performed in Python [37], using packages for numerical and scientific computing, data manipulation and machine learning [38, 39, 40, 41]. The visualizations were created using Python plotting libraries [42]. Survival modelling was done in the lifelines library [32].

An example of elaborated time-dependent data for one person from the dataset used for survival analysis is shown in Table A2.

3.1. Covariates description

The covariates are divided into five groups: general attributes, independent state attributes, dependent state attributes, transition attributes and indicator transition attributes. Their description, type and feasible values are shown in Table A3.

3.2. Model fitting

In survival regression, survival and cumulative hazard functions are functions of survival times with event indicator and covariates with corresponding coefficients, $S(t) = f(\mathbf{T}, \mathbf{E}, \mathbf{x}, \boldsymbol{\beta})$ and $H(t) = g(\mathbf{T}, \mathbf{E}, \mathbf{x}, \boldsymbol{\beta})$. Transition attributes are used as states for the analysis, whereas the remaining ones are used for conditioning. Table A4 shows results of fitting a CPH model for each pair of life events from the causality graph. The table also contains information on the number of existing samples for each pair, out of which, as mentioned above, 60 % were used for fitting and 40 % as hold-out data for validation.

4. Results and Discussion

4.1. Life Events Graph

To build our graphs, we ran Parents and Children (PC) and Greedy Equivalence Search (GES) algorithms with prepared data. Not surprisingly, we verified that their final graphs were not mutually consistent in several components. Fig. 1 shows the results. We can also see that many event pairs have bidirectional arrows. This is neither surprising nor necessarily problematic, it only means that temporal precedence can flow in either direction, and there may be cyclic stages in life (child birth \rightarrow moving \rightarrow new car...).

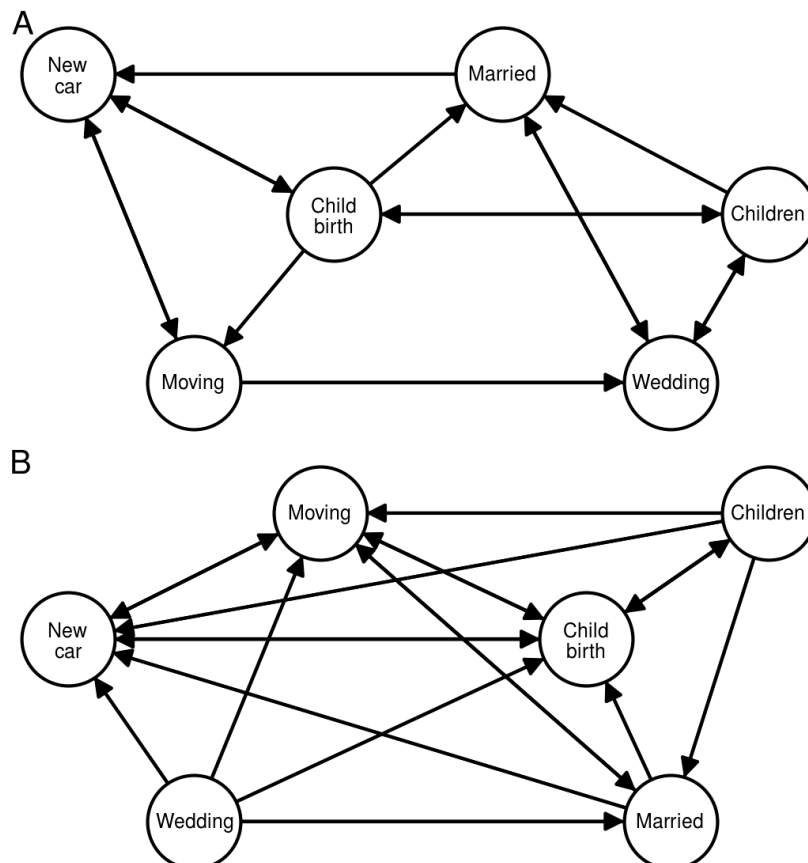


Fig. 1. Life events graphs created using PC and GES algorithms. (A) PC. (B) GES.

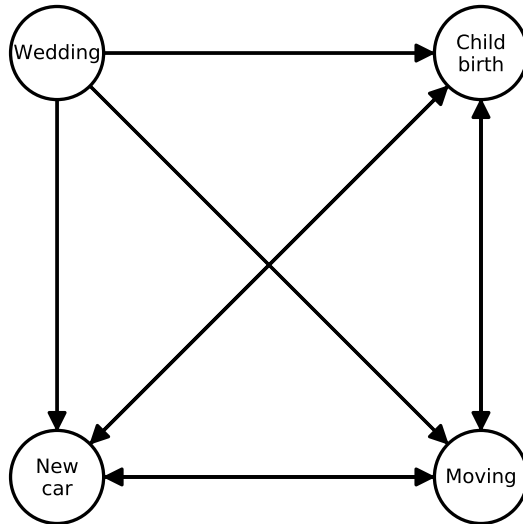


Fig. 2. Final events graph containing four life events.

The PC graph shows inconsistencies with respect to *wedding* and *child birth*, that we could not find in the dataset or in our personal intuition, particularly that *child birth* precedes being *married* (but not directly *wedding*, which becomes an isolated event, only preceded by *moving*). The child birth event almost certainly happens after other events (as can be seen in Figure A2), which was not uncovered by PC. It is apparent that the *married* state variable confuses the algorithm. On the other hand, in the GES case, the relationships do not show contradiction with our intuition and the graph proposes more relationships. This is valuable from the perspective of our work, because it allows for more diversity in life-course trajectories. We therefore decided to keep working with the GES algorithm for the remainder of this paper. We recognize that this decision is certainly quite arbitrary, and it also reflects the well-known challenge of comparability of graphs in causal discovery [43]. These algorithms themselves use different non-comparable score schemes [9].

The last step in this process was to reduce the graph to only events variables; henceforth we will only use the graph in Fig. 2.

4.2. Empirical Life Event Occurrence Distribution

When a certain life event is most likely to happen in a person’s life can be analysed from observed life-course data. Empirical histograms are a good way to understand occurrence distribution of events throughout the lifetime. Since in this case we work with completely observed data (only uncensored observations), the corresponding cumulative distribution functions (CDF), and thus the survival functions, can be directly derived. Fig. 3 shows histograms and survival functions for five life events: new car, moving, child birth, wedding and divorce.

The results show that buying a new car is likely to happen by the age of 40, with decreasing probability in time, with the age of 18 having the highest probability, which can be related to a legal minimum age for obtaining a driving license for unrestricted car use in Germany (Fig. 3A). Home relocation is most likely to happen in the later 20s, with the probability increasing up to the age of 25, and then decreasing up to the age of 50, after which people seem unlikely to move again (Fig. 3B). Having a child follows a normal distribution peaking around the age of 30, with a small chance of happening after the age of 45 (Fig. 3C). Getting married roughly follows a skewed distribution with a heavier tail on the increasing age side, and is most likely to happen in the 20s,

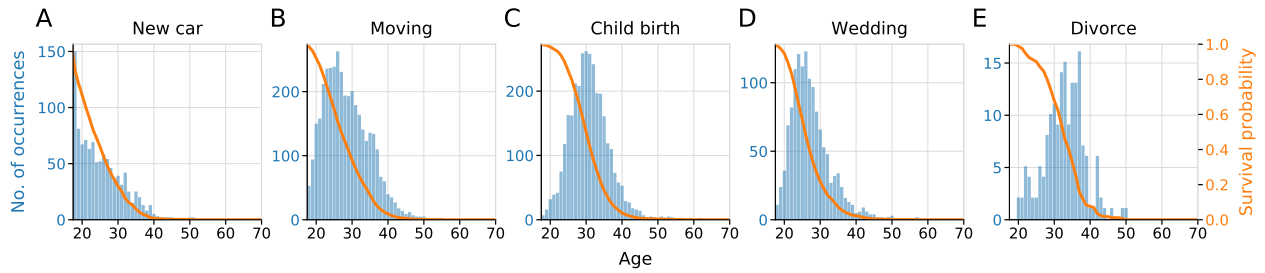


Fig. 3. Empirical histograms and corresponding survival functions. The x-axis shows age, the left-hand side y-axis shows observed counts of specific event in the dataset (depicted in blue) and the right-hand side y-axis shows derived survival functions from the data (depicted in orange). Since all of our samples are fully observed (i.e. uncensored), all survival functions converge to a value of zero. (A) New car event. (B) Moving event. (C) Child birth event. (D) Wedding event. (E) Divorce event.

specifically around mid 20s (Fig. 3D). Relating the events of getting married and having a child, having a child is shifted a couple of years after getting married, implying that in general having a child will likely happen soon after the wedding. Divorce is most likely to happen in the 30s, and generally before the age of 50, however the small sample size indicates an overall low probability of this happening (Fig. 3E).

4.3. Describing General Life Event Transitions

Survival and cumulative hazard functions are estimated using Kaplan-Meier and Nelson-Aalen methods ((7) and (8), respectively), and are shown in Fig. 4. Kaplan-Meier estimates of the survival function are shown in Fig. 4A, where colors represent survival probabilities ranging from 0–1. The more the color at time t is on the green side, the higher the survival probability $S(t)$ is, which indicates that there is a smaller probability that the event will happen before or at time t . This can also be interpreted as showing that the event has a higher probability of surviving up to time t . Normalized hazard functions are derived from the cumulative hazard function estimates and are shown in Fig. 4B, with a range from 0–1.

Looking at Fig. 4A, we can see that in the first 2-3 years after getting married, it is more likely that moving will happen next, rather than child birth or having a new car, which will then be superseded by having a child, as the survival probabilities decrease much faster after 3 years. Having a new car after getting married does not seem to be as important as child birth or moving. After having a child, moving is more likely than buying a new car. After moving, buying a car is unlikely compared to having a child, which is likely to happen within several years after moving. After buying a new car, there is a similar chance of moving as having a child.

Fig. 4B gives interesting insights regarding specific intervals when the event has a higher risk of happening on a timeline. For example, after getting married, the new car or moving event is very likely to happen within one year, whereas the child birth event has a higher risk of happening 10 or 11 years after getting married than in other periods. These estimates should be taken with caution: comparison of hazard rates among intervals gives individual difference between the rates. However, the hazard rates add up along the timeline, hence the cumulative risk of an event happening increases in time.

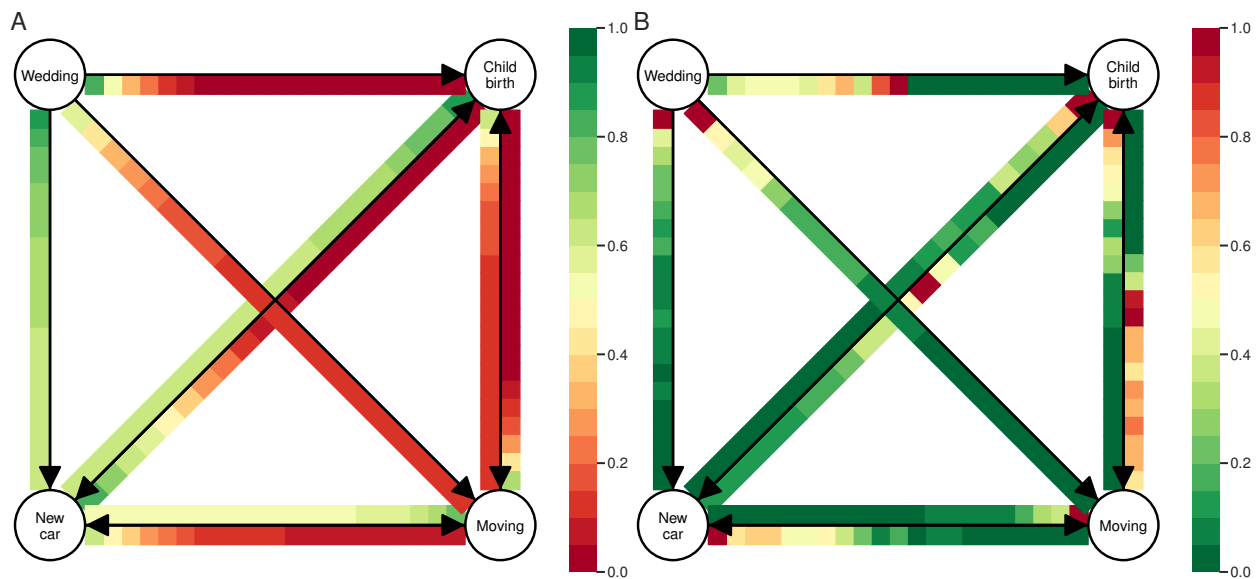


Fig. 4. General life event transitions, represented in a life event graph. Nodes represent life events and directed edges represent life event pairs, with results shown on the right-hand side of the edges. Each edge is divided into 21 equal-length parts, representing a period of 20 years and the base year, with its color depicting the function value shown in the scale from the color bar on the side of each graph. (A) Survival functions for each pair of life events. The function for each pair can be interpreted either individually or jointly compared among each other, as they all come from the same scale from 0–1. (B) Normalized hazard functions derived from the cumulative hazard functions computed using the Nelson-Aalen method. These functions should be interpreted only individually, to compare hazard rates between intervals, and not regarded jointly, as they are normalized to a range from 0–1 coming from different scales.

4.4. The Effect of Covariates on Transition Probabilities

The Cox proportional hazards models were trained and validated for each pair of life events from the life events graph. To understand the relationships in the data and prevent drawing false interpretations from the findings, detailed data preparation and description, as well as models training and validation, can be found in the previous section and in Appendix. Important to note is that the covariates obtained from time-dependent variables (e.g. age, number of children, number of cars, etc.) are measured at the time of the initial baseline event, which may help understand some of the less intuitive results shown later in the text (e.g. those without (or fewer) children at the time of wedding are less likely to get divorced, which can be explained by the fact that those are the ones who do not bring children into a partnership, and this results in the lower divorce probability, as shown in Fig. 6W). Another explanation comes from the small sample size of some covariates values (see Figure A4), e.g. people with more cars tend to buy cars faster than the ones with less or no cars, as shown later in Fig. 6N.

4.4.1. Average life event transitions behaviour

Average transition behaviour can give us an insightful overview of the survival behaviour within each pair and among the pairs, as shown in Fig. 5. Having a child is very likely to happen within 5 years from getting married, and almost certainly within ten years, with a probability close

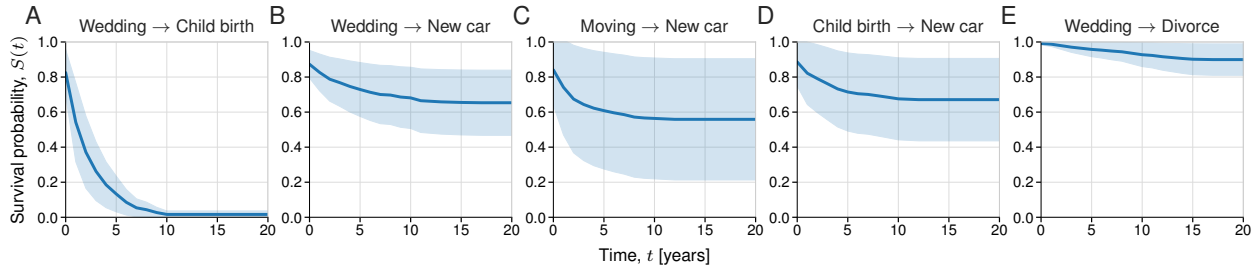


Fig. 5. Average life event transitions behaviour for five pairs of life events. Predicted survival functions are computed for all dataset samples for each pair separately, which are then averaged, thus marginalizing over all the covariates to account for their empirical distributions. The blue line represents the survival function, and the light blue area around it is the confidence belt (± 1 standard deviation). The average survival functions have different shapes, with different confidence belts at different times. The narrower the confidence belt is, the less impact covariates have on the survival probability. Since not all of our samples are fully observed (some are censored as in some cases the event does not happen within the study period), survival functions do not end up at 0. (A) Wedding to child birth transition. (B) Wedding to new car transition. (C) Moving to new car transition. (D) Child birth to new car transition. (E) Wedding to divorce transition.

to 0 after 10 years. There is a small probability of having a child in the same year the wedding happens (Fig. 5A). Buying a new car after getting married, as depicted in Fig. 5B, has almost identical probability as buying a new car after having a child (Fig. 5D), which will happen within the first 10 years from the initial event, and is unlikely after that. There is big variability in buying a new car after moving, demonstrated by the wide confidence interval in Fig. 5C, with a similar switch point at around 10 years. The divorce event is unlikely to happen, with survival probability slightly decreasing in time over the period of within 15 years (Fig. 5E).

It is important to note that general survival functions shown in Fig. 4A and average survival functions in Fig. 5 differ in the way that the former are obtained directly from the data using univariate models, while the latter are obtained by averaging over all predicted survival functions using fitted multivariate CPH regression models, thus including the effect of covariates and the model error.

4.4.2. Life event transitions behaviour

To examine the effect of socio-demographics on transition probabilities, we condition the prediction on one covariate for each transition pair (Fig. 6).

Wedding to Child birth

The decision to have a child after getting married is mainly driven by two factors: age and the number of children (Fig. 6A-E). In terms of age, after wedding younger people tend to have children later, where the age groups of [36, 45] and [46, 55] tend to have children before the others, with a high chance of having a child in the same year as the wedding or likely within 3 years from it. Regardless of age, if a couple does not have a child within 10 years of getting married, they will probably not have a child at all. In terms of nationality, Germans tend to have children slightly later than other nationals. Interestingly, couples with no previous children tend to have children later than couples with children, who are likely to have a child within 1 to 2 years after getting married.

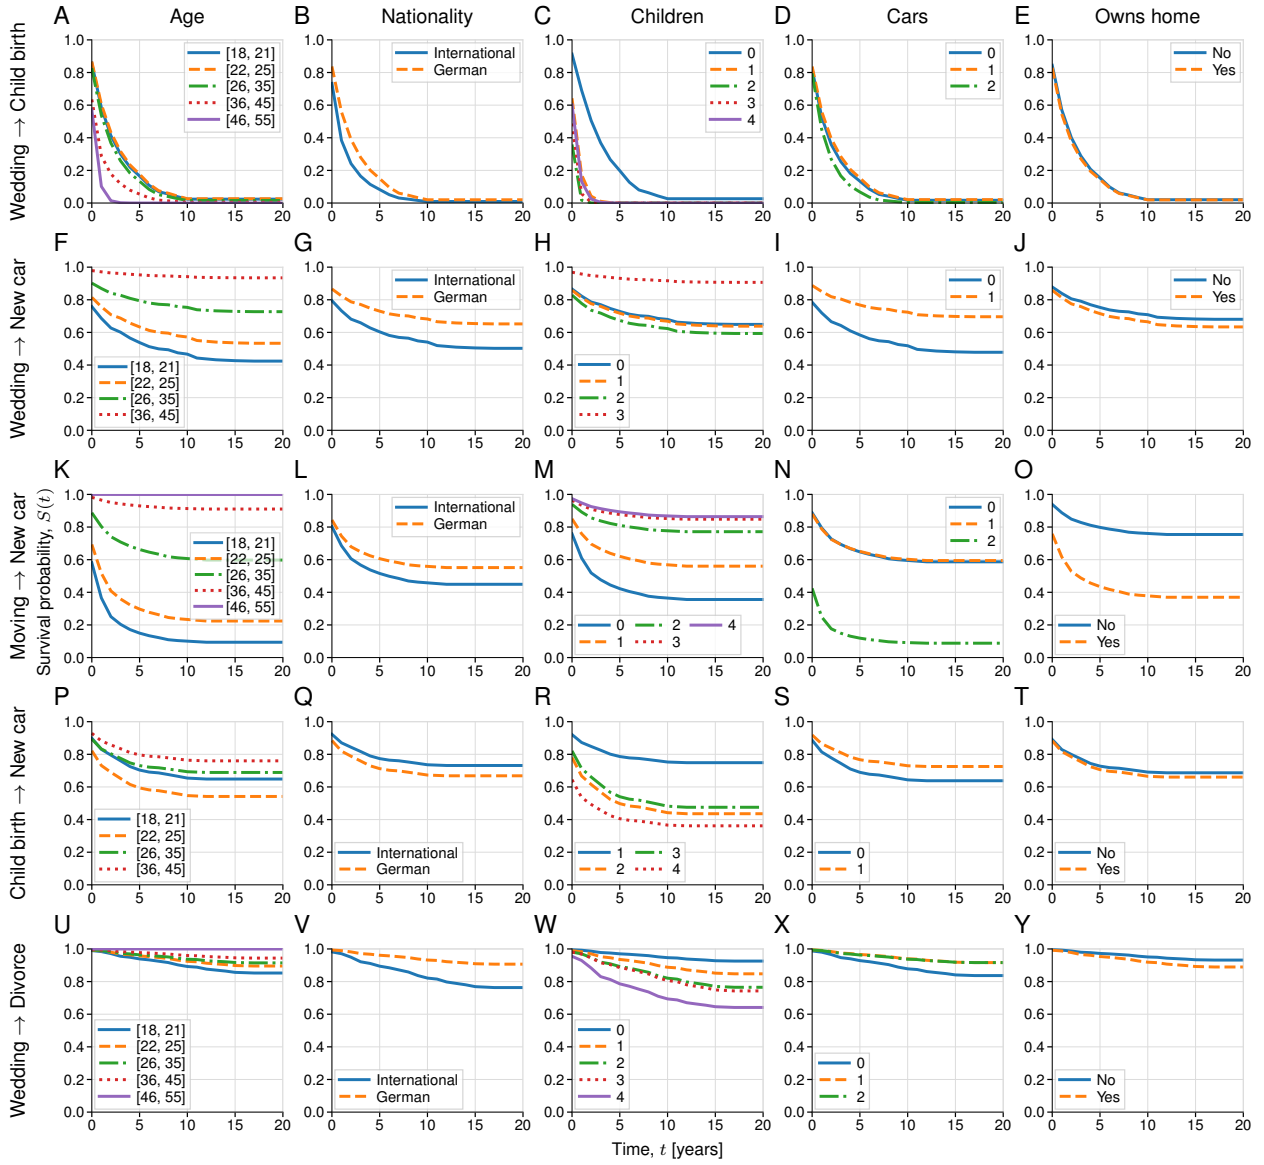


Fig. 6. Life event transitions behaviour for five selected pairs of life events expressed through survival functions. Each row with five graphs represent behaviour for a specific life events pair, with five columns indicating which socio-demographic covariates the transition behaviour is conditioned on. For each graph, the x-axis represents time in years passed from the initial event, modelling the probability of another event happening. The covariates are age, nationality, number of children, number of cars and owning a home, respectively in order of plotting. (A-E) Wedding to child birth transitions. (F-J) Wedding to new car transitions. (K-O) Moving to new car transitions. (P-T) Child birth to new car transitions. (U-Y) Wedding to divorce transitions.

Couples who already have children before the wedding could be couples who had children together before getting married or who had children with previous partners. The number of cars does not play a significant role in the decision to have a child. Similarly, whether a couple owns a home, which can be related to their economic situation, does not seem to be relevant for having a child.

Wedding to New car

The decision to have a new car after getting married is driven by several factors, with the most influential being age and the current number of cars (Fig. 6F-J). The youngest age group [18,21] tend to have a new car sooner than the others, with the probability gradually decreasing with increases in age, where people in the age group [36,45] are very unlikely to have a new car after getting married. Germans tend to have a new car later than other nationals. Those who have no children or one child exhibit similar behaviour, while having two children speeds up the decision to buy a car, whereas couples with three children postpone it. Intuitively, couples with no car tend to buy a car much sooner than couples who already possess a car. Couples who own a home tend to buy a new car slightly sooner than couples who do not own a home, which may be related to a better economic situation.

Moving to New car

The decision to have a new car after moving is driven by all factors except nationality (Fig. 6K-O). The difference in behaviour between age groups is very distinguishable - the younger the person the more likely he/she is to have a new car after moving, with the youngest group [18,21] tending to have a car within 2 years from moving. Germans tend to buy a new car later than other nationals. The number of children plays an important role, where people with fewer children tend to buy a car sooner, whereas people with 3 or more children are less likely to buy a new car after moving. For people who had no children at the time of moving, they are more likely (compared to others) to have a new (first) child after moving, which in turn may explain the faster car acquisition. Surprisingly, people living in households with two cars or more are very likely to buy or replace a car soon after moving, most of them within two years, including the year of moving. The magnitude of this is very strong, which may be an artifact of a small sample group, or the different possibilities in temporal coding. People who own a home are significantly more likely to get a car before people who do not own a home.

Child birth to New car

The decision of having a new car after having a child is driven by several factors, of which the most influential are age and the number of cars (Fig. 6P-T). With the exception of the youngest age group [18,21], the increase in age postpones buying a car after having a child. Germans are likely to have a new car slightly sooner than other nationals. People with one child will probably buy a car much later than people with more than one child. In line with intuition, after having a child people without a car tend to buy a car sooner than people with a car. Similarly, people who own a home are more likely to buy a car than people not owning a home are.

Wedding to Divorce

The overall probability of getting divorced is low, but some variability is still present (Fig. 6U-Y). In general, based on observed divorces, if a divorce does not happen within 15 years from the wedding, it will probably not happen. The youngest age group tend to get divorced sooner than others. Germans tend to get divorced later than other nationals. People with fewer children have lower divorce probability than people with more children, with a sensitive group being people with four children. An intuitive interpretation for this could be that those without (or fewer) children at

the time of getting married are those that do not bring children into a partnership, and this results in the lower divorce probability. The number of cars can also affect divorce probability, where people with fewer cars are more likely to get divorced. Homeowners are more likely to get a divorce than non-homeowners.

The effect of socio-demographics on transitions

Regarding the general effect of socio-demographics on transitions, younger age groups generally prefer to have a new car and divorce sooner, while transitions to having a child is faster as age increases, with the oldest group tending to postpone the decision of buying a car (Fig. 6A,F,K,P,U). Regarding nationality, other nationals make faster transitions than Germans, except when it comes to buying a car after having a child (Fig. 6B,G,L,Q,V). The number of children seems to play a role in all transitions (Fig. 6C,H,M,R,W). The number of cars, as expected, tends to have an effect in the decision to buy a new car (Fig. 6D,I,N,S,X). Homeowners generally have faster transitions than people not owning a home (Fig. 6E,J,O,T,Y).

5. Conclusion

Numerous studies in transport research and sociology seek to model the structure of the life course and the inter-connection between life-course events (birth of a child, home relocation, car change, etc.). The use of traditional statistical models does not generally allow an analyst to infer causality without overrelying on expert knowledge. Indeed, statistical models of choice, which are predominantly used in the field, provide correlation measures rather than causation measures. By introducing a bi-level model where the upper level is a causal discovery algorithm and the lower level is composed of a series of survival analyses modelling the transition rate between life course events, we introduce a more rigorous data driven assessment of causality in life-course analysis. The results of this research can be used for high-level activity modelling for mobility applications, to better understand how life events and socio-demographic characteristics affects and constraints people's mobility and travel choices.

Acknowledgements

The research leading to these results has received funding from the European Union's Horizon 2020 research and innovation program under the Marie Skłodowska-Curie grant agreement no. 754462. The Leeds authors acknowledge the financial support by the European Research Council through the consolidator grant 615596-DECISIONS.

References

- [1] T. Hägerstrand, “What about people in regional science?,” *Papers in regional science*, vol. 24, no. 1, pp. 6–21, 1970.
- [2] F. S. Chapin, *Human Activity Patterns in the City: Things People Do In Time and in Space*. New York: John Wiley & Sons, 1974.
- [3] H. Müggenburg, A. Busch-Geertsema, and M. Lanzendorf, “Mobility biographies: A review of achievements and challenges of the mobility biographies approach and a framework for further research,” *Journal of Transport Geography*, vol. 46, pp. 151–163, 2015.
- [4] S. Beige and K. W. Axhausen, “Interdependencies between turning points in life and long-term mobility decisions,” *Transportation*, vol. 39, no. 4, pp. 857–872, 2012.
- [5] D. Freedman, A. Thornton, D. Camburn, D. Alwin, and L. Young-DeMarco, “The life history calendar: A technique for collecting retrospective data,” *Sociological Methodology*, vol. 18, pp. 37–68, 1988.
- [6] A. Manzoni, J. K. Vermunt, R. Luijkx, and R. Muffels, “2. memory bias in retrospectively collected employment careers: A model-based approach to correct for measurement error,” *Sociological Methodology*, vol. 40, no. 1, pp. 39–73, 2010.
- [7] A. Manzoni, “In and out of employment: Effects in panel and life-history data,” *Advances in Life Course Research*, vol. 17, no. 1, pp. 11–24, 2012.
- [8] R. Schoenduwe, M. G. Mueller, A. Peters, and M. Lanzendorf, “Analysing mobility biographies with the life course calendar: a retrospective survey methodology for longitudinal data collection,” *Journal of Transport Geography*, vol. 42, pp. 98–109, 2015.
- [9] C. Glymour, K. Zhang, and P. Spirtes, “Review of causal discovery methods based on graphical models,” *Frontiers in Genetics*, vol. 10, p. 524, 2019.
- [10] M. J. Bradburn, T. G. Clark, S. B. Love, and D. G. Altman, “Survival analysis part II: Multivariate data analysis – an introduction to concepts and methods,” *British Journal of Cancer*, vol. 89, no. 3, pp. 431–436, 2003.
- [11] E. Bair and R. Tibshirani, “Semi-supervised methods to predict patient survival from gene expression data,” *PLOS Biology*, vol. 2, no. 4, p. e108, 2004.
- [12] O. O. Aalen, P. K. Andersen, Ø. Borgan, R. D. Gill, and N. Keiding, “History of applications of martingales in survival analysis,” *Electronic Journal for History of Probability and Statistics*, vol. 5, no. 1, pp. 1–28, 2009.
- [13] W.-Y. Cheng, T.-H. O. Yang, and D. Anastassiou, “Development of a prognostic model for breast cancer survival in an open challenge environment,” *Science Translational Medicine*, vol. 5, no. 181, p. 181ra50, 2013.
- [14] P. Royston and D. G. Altman, “External validation of a Cox prognostic model: principles and methods,” *BMC Medical Research Methodology*, vol. 13, no. 33, pp. 1471–2288, 2013.

- [15] R. W. Yeh, E. A. Secemsky, D. J. Kereiakes, S.-L. T. Normand, A. H. Gershlick, D. J. Cohen, J. A. Spertus, P. G. Steg, D. E. Cutlip, M. J. Rinaldi, E. Camenzind, W. Wijns, P. K. Apruzzese, Y. Song, J. M. Massaro, L. Mauri, and for the DAPT Study Investigators, “Development and validation of a prediction rule for benefit and harm of dual antiplatelet therapy beyond 1 year after percutaneous coronary intervention,” *JAMA*, vol. 315, no. 16, pp. 1735–1749, 2016.
- [16] R. D. Shah and J. Peters, “The hardness of conditional independence testing and the generalised covariance measure,” *arXiv e-prints*, no. arXiv:1804.07203 [math.ST], 2018.
- [17] A. P. Dawid, “Conditional independence in statistical theory,” *Journal of the Royal Statistical Society. Series B (Methodological)*, vol. 41, no. 1, pp. 1–31, 1979.
- [18] P. Spirtes, C. N. Glymour, R. Scheines, D. Heckerman, C. Meek, G. Cooper, and T. Richardson, *Causation, Prediction, and Search*. MIT Press, 2000.
- [19] D. M. Chickering, “Optimal structure identification with greedy search,” *Journal of Machine Learning Research*, vol. 3, no. 11, pp. 507–554, 2002.
- [20] D. Kalainathan and O. Goudet, “Causal Discovery Toolbox: Uncover causal relationships in Python,” *arXiv e-prints*, no. arXiv:1903.02278 [stat.CO], 2019.
- [21] E. L. Kaplan and P. Meier, “Nonparametric estimation from incomplete observations,” *Journal of the American Statistical Association*, vol. 53, no. 282, pp. 457–481, 1958.
- [22] W. Nelson, “Theory and applications of hazard plotting for censored failure data,” *Technometrics*, vol. 14, no. 4, pp. 945–966, 1972.
- [23] O. O. Aalen, “Nonparametric inference for a family of counting processes,” *The Annals of Statistics*, vol. 6, no. 4, pp. 701–726, 1978.
- [24] D. R. Cox, “Regression models and life-tables,” *Journal of the Royal Statistical Society. Series B (Methodological)*, vol. 34, no. 2, pp. 187–220, 1972.
- [25] S. Pölsterl, N. Navab, and A. Katouzian, “Fast training of support vector machines for survival analysis,” in *Machine Learning and Knowledge Discovery in Databases* (A. Appice, P. P. Rodrigues, V. Santos Costa, J. Gama, A. Jorge, and C. Soares, eds.), (Cham, Switzerland), pp. 243–259, Springer International Publishing, 2015.
- [26] S. Pölsterl, P. Gupta, L. Wang, S. Conjeti, A. Katouzian, and N. Navab, “Heterogeneous ensembles for predicting survival of metastatic, castrate-resistant prostate cancer patients [version 3; peer review: 2 approved, 1 approved with reservations],” *F1000Research*, vol. 5, no. 2676, 2017.
- [27] F. C. Pereira and S. S. Borysov, “Machine learning fundamentals,” in *Mobility Patterns, Big Data and Transport Analytics*, ch. 2, pp. 9–29, Amsterdam, The Netherlands: Elsevier, 2018.
- [28] D. Faraggi and R. Simon, “A neural network model for survival data,” *Statistics in Medicine*, vol. 14, no. 1, pp. 73–82, 1995.

- [29] J. L. Katzman, U. Shaham, A. Cloninger, J. Bates, T. Jiang, and Y. Kluger, “DeepSurv: personalized treatment recommender system using a Cox proportional hazards deep neural network,” *BMC Medical Research Methodology*, vol. 18, no. 24, 2018.
- [30] C. Lee, W. R. Zame, J. Yoon, and M. van der Schaar, “DeepHit: A deep learning approach to survival analysis with competing risks,” in *Thirty-Second AAAI Conference on Artificial Intelligence*, pp. 2314–2321, AAAI, 2018.
- [31] M. F. Gensheimer and B. Narasimhan, “A scalable discrete-time survival model for neural networks,” *PeerJ*, vol. 7, p. e6257, 2019.
- [32] C. Davidson-Pilon, J. Kalderstam, P. Zivich, B. Kuhn, A. Fiore-Gartland, L. Moneda, Gabriel, D. Wilson, A. Parij, K. Stark, S. Anton, L. Besson, Jona, H. Gadgil, D. Golland, S. Hussey, R. Kumar, J. Noorbakhsh, A. Klintberg, E. Ochoa, D. Albrecht, dhuynh, D. Medvinsky, D. Zgonjanin, D. S. Katz, D. Chen, C. Ahern, C. Fournier, Arturo, and A. F. Rendeiro, “CamDavidsonPilon/lifelines: v0.22.2,” 2019.
- [33] E. Süli and D. F. Mayers, *An Introduction to Numerical Analysis*. Cambridge, UK: Cambridge University Press, 2003.
- [34] B. Efron, “The efficiency of Cox’s likelihood function for censored data,” *Journal of the American Statistical Association*, vol. 72, no. 359, pp. 557–565, 1977.
- [35] N. Breslow, “Covariance analysis of censored survival data,” *Biometrics*, vol. 30, no. 1, pp. 89–99, 1974.
- [36] F. Harrell, *Regression Modeling Strategies*. Springer Series in Statistics, Switzerland: Springer International Publishing, 2 ed., 2015.
- [37] G. van Rossum and F. L. Drake, *The Python Language Reference Manual*. Boston, MA: Network Theory Ltd., 2011.
- [38] W. McKinney, “Data structures for statistical computing in Python,” in *Proceedings of the 9th Python in Science Conference* (S. van der Walt and J. Millman, eds.), pp. 51–56, 2010.
- [39] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and É. Duchesnay, “Scikit-learn: Machine learning in Python,” *Journal of Machine Learning Research*, vol. 12, pp. 2825–2830, 2011.
- [40] S. van der Walt, S. C. Colbert, and G. Varoquaux, “The NumPy array: A structure for efficient numerical computation,” *Computing in Science & Engineering*, vol. 13, no. 2, pp. 22–30, 2011.
- [41] P. Virtanen, R. Gommers, T. E. Oliphant, M. Haberland, T. Reddy, D. Cournapeau, E. Burovski, P. Peterson, W. Weckesser, J. Bright, S. J. van der Walt, M. Brett, J. Wilson, K. Jarrod Millman, N. Mayorov, A. R. J. Nelson, E. Jones, R. Kern, E. Larson, C. Carey, Í. Polat, Y. Feng, E. W. Moore, J. Vand erPlas, D. Laxalde, J. Perktold, R. Cimrman, I. Henriksen, E. A. Quintero, C. R. Harris, A. M. Archibald, A. H. Ribeiro, F. Pedregosa, P. van Mulbregt, and S. . . Contributors, “SciPy 1.0–Fundamental Algorithms for Scientific Computing in Python,” *arXiv e-prints*, no. arXiv:1907.10121 [cs.MS], 2019.

- [42] J. D. Hunter, “Matplotlib: A 2D graphics environment,” *Computing in Science & Engineering*, vol. 9, no. 3, pp. 90–95, 2007.
- [43] K. Singh, G. Gupta, V. Tewari, and G. Shroff, “Comparative benchmarking of causal discovery techniques,” *arXiv e-prints*, no. arXiv:1708.06246 [cs.AI], 2017.

Appendix

Data collection

A screenshot of the life-course calendar used in this study is shown in Fig. A1.

	Dwelling	Year	Reason for relocation	Place	City size	Distance to previous apartment in km	Rented or owned
1							
2							
3							
4							
...							

Fig. A1. A screenshot of the life-course calendar used in this study (translated from German into English).

Causal discovery

In order to prepare the data for the algorithms, we created an observation for each event pair occurrence, as illustrated in Table A1. Observation 2 can be read this way: “at some moment in time, a married individual with 2 children had a child birth event, which was followed by a moving event later in life”. Due to the nature of the dataset, an “event occurrence” always matches a year, so one may have multiple events preceding multiple events, as happens in observation 4.

Regarding the choice of the causal graph, Fig. A2 shows the survival probability marginalized over all other events (new car, moving, wedding) leading to the child birth event for the GES graph. There is a clear indication that the child birth happens almost certainly after each one of these events. These relations are not present in the PC graph, where only the new car event precedes the child birth event.

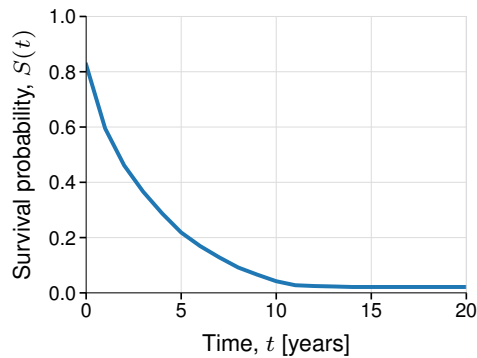


Fig. A2. Average survival curve from all other events (new car, moving, wedding) to the child birth event.

Table A1. An example of four *event pair* observations.

ID	New car	Moving	Child birth	Wedding	Married	Children	New car next event	Moving next event	Child birth next event	Wedding next event
1	0	1	0	0	1	2	1	0	0	0
2	0	0	1	0	1	2	0	1	0	0
3	1	0	1	0	1	1	0	0	1	0
4	1	1	0	1	1	0	1	0	1	0

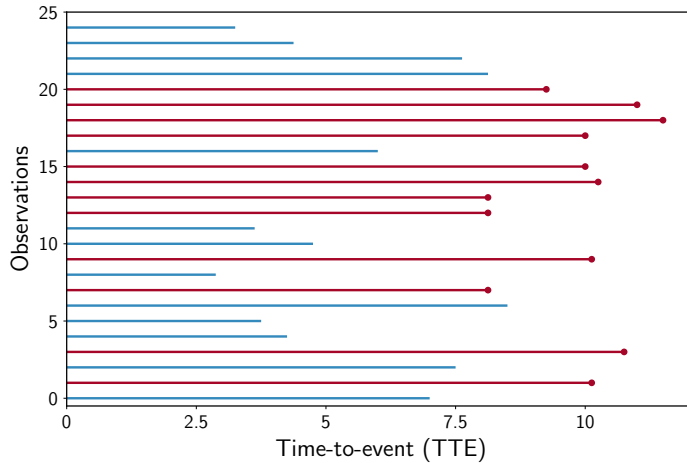


Fig. A3. General example of survival analysis data. Uncensored observations (red) and censored observations (blue).

Survival analysis

Fig. A3 shows a general overview of observations in survival analysis.

Data elaboration

An example of elaborated time-dependent data for one person from the dataset used for survival analysis is shown in Table A2.

Covariates description

The covariates are divided into five groups: general attributes, independent state attributes, dependent state attributes, transition attributes and indicator transition attributes. Their description, type and feasible values are shown in Table A3.

General attributes are *Gender*, *Nationality*, *No. of relocations* and *City size*. They represent general characteristics of individuals. Their value neither depends on the year (they are time-constant), nor on the life events pair. *Gender* and *Nationality* are categorical attributes, *No. of relocations* and *City size* are discrete.

Independent state attributes are *Age*, *Owns home*, *Distance to work* and *Rides car*. These are time-dependent attributes. Given a life events pair, an independent state attribute takes a value it has in the year the cause event occurs. *Age*, *Owns home* and *Rides car* are categorical, whereas *Distance to work* is discrete (discretised to 1 km).

Dependent state attributes are *No. of cars*, *No. of children* and *Married*. They are affected by the transition attributes. Their value is computed in the same way as for the independent state attributes. They are all categorical.

Transition attributes are *New car*, *Moving*, *Child birth*, *Wedding* and *Divorce*. These are binary variables, having values of mostly 0, and 1 when the event happens. They increment a dependent state attribute (in case of *No. of cars* and *No. of children*), or alternate the value of *Married*. These attributes are treated in a different way. To be used in estimations, they are expanded in a way how many years they happened before the cause event. Therefore, they take

Table A2. An example of time-dependent data for one person from the dataset.

Attribute	Year																			
	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20
Age	18	19	20	21	22	23	24	25	26	27	28	29	30	31	32	33	34	35	36	37
Owens home	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1
Cars	0	0	0	0	0	0	0	0	1	1	1	1	1	1	1	1	1	1	1	1
Distance to work	0	0	8	7	7	5	5	5	5	5	5	7	7	7	7	7	7	7	7	7
Rides car	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
Children	0	0	0	0	0	0	0	1	1	1	1	1	1	1	1	1	1	1	1	1
Married	0	0	0	0	0	0	1	1	1	1	1	1	1	1	1	1	1	1	1	1
New car	0	0	0	0	0	0	0	0	1	0	0	0	0	0	0	0	0	0	0	0
Moving	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
Child birth	0	0	0	0	0	0	0	1	0	0	0	0	0	0	0	0	0	0	0	0
Wedding	0	0	0	0	0	0	1	0	0	0	0	0	0	0	0	0	0	0	0	0
Divorce	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

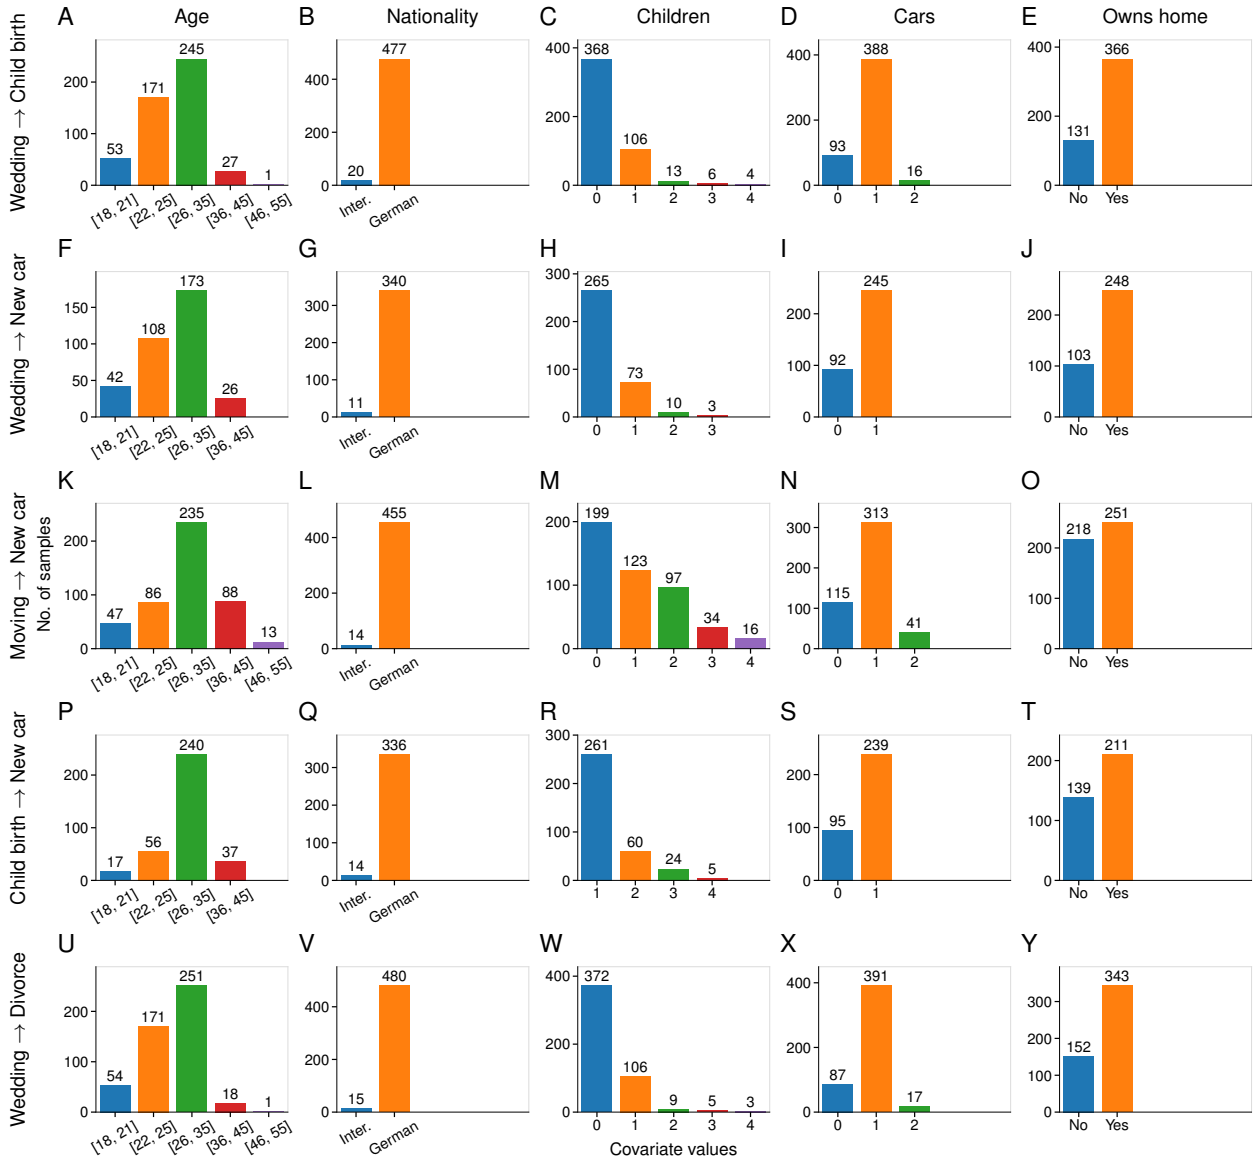


Fig. A4. Distribution of covariates values.

values in the range $[0, -19]$, since the dataset for each person covers a period of 20 years. They are all discrete.

Indicator transition attributes are dummy variable used to account for whether the transition attribute happened or not before the cause effect. They are all categorical.

Distribution of covariates values is shown in Fig. A4.

Model fitting

Model fitting results are shown in Table A4.

Table A3. Covariates description

Group	Name	Description	Type	Domain
General attributes	Gender	gender, is the person male or female	categorical	{0, 1}
	Nationality	nationality, is the person German or international	categorical	{0, 1}
	No. of relocations	no. of times the person relocated with parents	discrete	[0, 9]
	City size	the size of the city the person lives in	discrete	[0, 7]
Independent state attributes	Age	age, a range of 20 years for every person	categorical	5 age groups
	Owens home	does the person own a home or not	categorical	{0, 1}
	Distance to work	travelling distance to work [km]	discrete	[0, 100]
	Rides car	if the person uses the car to go to work	categorical	{0, 1}
Dependent state attributes	No. of cars	the number of cars the person has	categorical	[0, 2]
	No. of children	the number of children the person has	categorical	[0, 3]
	Married	marital status, is the person married or not	categorical	{0, 1}
Transition attributes	New car	the event of buying a new car	discrete	[0, -19]
	Moving	the event of moving	discrete	[0, -19]
	Child birth	the event of having a child	discrete	[0, -19]
	Wedding	the event of having a wedding (getting married)	discrete	[0, -19]
	Divorce	the event of having a divorce	discrete	[0, -19]
Indicator transition attributes	Indicator - New car	indicator for the event of buying a new car	categorical	{0, 1}
	Indicator - Moving	indicator for the event of moving	categorical	{0, 1}
	Indicator - Child birth	indicator for the event of having a child	categorical	{0, 1}
	Indicator - Wedding	indicator for the event of having a wedding (getting married)	categorical	{0, 1}
	Indicator - Divorce	indicator for the event of having a divorce	categorical	{0, 1}

Table A4. Fitted CPH models

Cause	Effect	Sample size	C-index - train	C-index - test
New car	Moving	819	0.687	0.605
New car	Child birth	736	0.801	0.757
Moving	New car	1172	0.863	0.864
Moving	Child birth	1704	0.831	0.824
Child birth	Moving	1481	0.670	0.679
Child birth	New car	873	0.784	0.809
Wedding	Moving	1175	0.591	0.604
Wedding	Child birth	1241	0.733	0.748
Wedding	Divorce	1237	0.719	0.682
Wedding	New car	877	0.608	0.652