



**Bridging the micro-macro gap in
population forecasting**
Contract no. SP23-CT-2005-006637

Deliverable D8 (first part)

**Description of the micro-simulation model
(Continuous-time micro-simulation)**

Work Package 2
Microsimulation

Author:

Frans Willekens



Netherlands Interdisciplinary Demographic Institute
P.O. Box 11650
2502 AR the Hague
The Netherlands

Continuous-time microsimulation *

Frans Willekens

24 September 2006

-
- * The research is part of the MicMac project, an international (European) project funded by the European Commission in the context of the Sixth Framework Programme. The project started in 2005 and aims at the development of a new demographic projection model that bridges the gap between aggregate projections of cohorts (Mac) and projections of the life courses of individual cohort members (Mic). The model is expected to be used to prepare detailed demographic projections for sustainable pension and health care systems. The project is carried out by a consortium of researchers of the Netherlands Interdisciplinary Demographic Institute (NIDI), Vienna Institute of Demography (VID), Institut National d'Etudes Demographiques (INED) (Paris), Bocconi University (Milan), Erasmus Medical Centre (Rotterdam), Max Planck Institute for Demographic Research (Rostock, Germany), the International Institute for Applied System Analysis (IIASA) (Laxenburg, Austria) and the University of Rostock. For details, see Van der Gaag et al. (2005) and the MicMac website www.micmac-projections.org

Abstract

Most dynamic microsimulation models treat time as a discrete variable and sample a uniform distribution to determine whether and when events or transitions occur. The shortcomings of discrete-time microsimulation models are removed if time is a continuous variable. In that case, the timing of transitions is determined by drawing a random waiting time from a waiting time distribution. Waiting time distributions are parametric models of time dependence of transition rates. They include the exponential model, the Gompertz model and the Weibull model. These models are studied extensively in survival analysis and event history modelling. Waiting time distributions in microsimulation is the subject of this paper. Drawing random waiting times from a distribution relies on the inverse distribution function or *quantile function*. Once the quantile function associated with a waiting time distribution is identified, continuous-time microsimulation is as straightforward as discrete-time microsimulation. The paper illustrates the quantile function approach to continuous-time microsimulation using the exponential model and the Gompertz model and shows that the approach can be extended to other models of time dependence of transition rates, such as the Cox proportional hazard model. The quantile function approach is an effective way to integrate dynamic microsimulation in survival analysis and event history modelling and to resolve several issues that trouble discrete-time microsimulations.

Content

1.	Introduction	1
2.	Continuous-time microsimulation: an overview	3
3.	Exponential waiting times	7
3.1.	Single origin and single destination	7
3.2.	Multiple origins and multiple destinations	13
3.3.	Covariates and interventions	17
4.	Other distributions of waiting times	20
5.	Conclusion	22
	References	23
	Annex A: Dates of events	27

1. Introduction

In his book *Foundations of social theory*, Coleman (1990, p. 2) advances two modes of explanation of behaviour of a system. One depends on observation of the behaviour of the system as a whole over a period of time. The other examines processes internal to the system involving components of the system and individual units. The rationale for studying the behaviour of a system by examining individual behaviour is that individual-level actions and system-level behaviour are linked. Colemans's two modes of explanation or analysis are present in simulations of population systems.

Macrosimulation studies the system as a whole. In microsimulation the individual is the unit of analysis. Microsimulation models differ in many ways, one is the treatment of time. Static models disregard time and assume that adjustments following a change or an intervention are immediate. These models have been used principally to calculate the impact of institutional changes in the tax and benefit system (Zaidi and Rake, 2001). Dynamic models incorporate time. The adjustments are not immediate but take time and the effect may vary over time. Dynamic models are used to explore the short-term and long-term impact of social and economic policies. They are also able to operate prospectively and, as a result, play an important role in informing social scientific thinking about the future. Recent overviews of dynamic microsimulation models include O'Donoghue (1999), Zaidi and Rake (2001) and Dupont et al. (2003).

According to Wolf (2000a) the essential ingredients of microsimulation are an analysis conducted at the level of the individual and the use of computer-based sampling. The sampling perspective on microsimulation is particularly interesting. Wolf asserts that microsimulation is essentially an exercise in sampling. It consists of drawing a sample from a virtual population. The observation may be a cross-section of the population at one point in time or a longitudinal study covering an extended period of time. The virtual population and its dynamics are fully described by one or several probability models. If the models are realistic, the virtual population closely resembles the real population and a sample of the virtual population cannot be distinguished from a sample of the real population. A particular class of probability models is used to describe the occurrence, timing and sequence of events that members of a population experience. Models in that class are known as transition models, duration models, survival models, hazard models, event history models and waiting time models, depending on the field of study. I refer to the models as transition models. They are regression models that relate the likelihood of an event at a given time to the duration of the process under study since a reference event (process time) and several other factors/predictors that may affect the event occurrence in addition to the process time. The likelihood of an event is expressed as a rate or a probability. Rates and probabilities differ in exposed population. In case of probabilities, the exposed population is the population at risk at the beginning of an interval. In case of rates, individuals remain exposed throughout the interval unless they leave the population at risk during the interval or enter the population late. The distinction between rates and

probabilities is particularly useful in differentiating microsimulation in continuous time from that in discrete time¹.

Dynamic microsimulation models are distinguished on the basis of the treatment of time. Most models treat time as a discrete variable. They determine which individuals experience particular events in given time intervals. Other models treat time as a continuous variable and determine the exact time to event. Transition models that are formulated in discrete time typically predict transition probabilities. Transition models that are formulated in continuous time predict transition rates. Discrete-time microsimulation models incorporate probability models, such as the logit model, the probit model and the logistic regression model. Continuous-time microsimulation models incorporate transition rate models, such as the basic exponential model, the Gompertz model, the Weibull model and the Cox regression model.

The aim of this paper is twofold. The first aim is to describe microsimulation from a sampling perspective. The second is to make explicit the transition rate models that determine the exact times to events (continuous time) experienced by members of the virtual population. In pursuing these aims, a more distant goal is pursued: *to embed microsimulation in the statistical theory of transition data and duration data*. That theory is well established (see e.g. Blossfeld and Rohwer, 2002; Klein and Moeschberger, 2003). The view of continuous-time microsimulation models as duration models is not entirely new. Galler (1997) advocated the link. This paper documents the links that exist and, as a result, aims at contributing to the development of an integrative framework that combines statistical techniques of data *analysis* and simulation techniques of data *synthesis*. The main thesis of this paper is that transition models that are commonly used for the analysis of (sample) observations may also be used for the synthesis of a virtual population that is an accurate picture of a real population.

In continuous-time microsimulation, the variable of interest is the time to event. The time to event is also referred to as the survival time, the failure time and the waiting time to the event. The time to event is a random variable with characteristic properties. The random variable may take on a range of values. The distribution of the values is described by functions, such as the distribution function, the survival function, the probability density function and the hazard function. One characteristic function is an essential ingredient of continuous-time microsimulation. It is the *inverse distribution function* or *quantile function*. The quantile function translates a probability into a real number, whereas the more commonly used distribution function and survival function translate a real number into a probability. The real number is the time to event. Microsimulation in continuous time resolves three important problems of discrete-time microsimulation. The first is how to determine the sequence of events when during an interval multiple events

¹ In the literature probabilities are often not adequately distinguished from rates. For instance Holmer et al. (2006, pp. 101-102) uses death rates from the Social Security Administration's Office of the Chief Actuary, and treat these rates as probabilities of dying. The mortality rate is used "to determine stochastically whether or not death actually occurs." (Holmer et al., 2006, p. 102). In PENSIM, a potential death is scheduled just before each birthday and discrete-time simulation is used to determine the individuals who die.

may occur. In discrete-time microsimulation, the sequence is determined exogenously. In continuous-time microsimulation, the simulation model determines the sequence. The second problem is how to determine the lengths of episodes between events. In discrete-time microsimulation the durations can be determined only approximately, whereas they can be determined precisely in continuous-time microsimulation. The third problem is the estimation of the number of events during an interval.

The paper is organized as follows. Section 2 contrasts continuous-time dynamic microsimulation to the more common discrete-time microsimulation. It lists several continuous-time microsimulation models that are operational today. It also contrasts microsimulation and macrosimulation. Most microsimulation models apply alignment techniques to bridge the gap between macro-level and micro-level analysis. The MicMac model adopts a different approach. It specifies the same transition rate model for the macrosimulation (Mac) and the microsimulation (Mic), although the models may differ in the number of explanatory variables included. The virtual population that is described by the transition rate models has many applications. It serves as a simulator, a device that incorporates important features of the real population and can be used for training, research, impact assessment and forecasting. The occurrence, timing and sequence of events in the virtual population are determined by transition rate models. Section 3 presents a simple time to event model. It is the basic exponential transition rate model. In it, the transition rate is time-invariant and the time to event follows an exponential distribution. The quantile function of the exponential model is used to determine the exact time to event from a random probability drawn from a uniform distribution. Section 4 discusses a transition rate model that incorporates time dependence of transition rates. It is the Gompertz model. Section 5 concludes the paper.

2. Continuous-time microsimulation: an overview

Transitions in the life course occur in continuous time but are often studied in discrete time. Most microsimulation models treat time as a discrete variable. Transitions in discrete time are represented by a comparison of the state occupied at the beginning and the end of a discrete time interval. The exact timing of a transition is not recorded and multiple transitions during an interval are either omitted or assumptions about the ordering and the timing of the events are imposed exogenously. Zaidi and Rake summarize the major problem of being not able to properly handle multiple events during an interval: “More importantly, the complexity and interdependence of events cannot always be modelled using such an approach. The annual cycle of events through which individuals pass has to be ordered, with the ordering implying a certain line of causality between events.” (Zaidi and Rake, 2001, p. 19). Similar problems of ordering events in discrete-time modeling arise in macrosimulation models such as the ProFamy model developed by Zeng Yi and others². For discussions of the ordering of events in discrete-

² In ProFamy, it is assumed that births occur throughout the first half and the second half of the year and that the other events occur in the middle of the year. The birth probabilities used refer to the corresponding half-year; they depend on the status at the beginning and middle of the year, respectively. The authors also assume that a person would not give a birth in the second half of the year if she or he had a birth in the first

time simulations, see e.g. Fredriksen (1998, p. 47) on Norway's MOSART model and Bonnet et al. (1999) on France's DESTINIE model. The problem of ordering multiple events in the same interval also occurs in transition data analysis when exact times to events are missing but intervals are observed. This situation is known as interval censoring (see e.g. Klein and Moeschberger, 2003). Compare that discussion with the approach adopted in continuous-time models (e.g. Holmer et al., 2006, p. 104 on the interdependence between leaving college and first job).

When transitions are measured in continuous time, events can occur at any arbitrary moment. In addition, multiple events may occur during a discrete time interval. Galler (1997), who compared in a systematic way the continuous-time approach to dynamic microsimulation modelling and approaches based on a discrete-time framework, views continuous-time microsimulation models as duration models. As a consequence, statistical and demographic methods of duration analysis may be used in the development of microsimulation models. Continuous-time microsimulation involves the generation of times to events (survival times, waiting times). The benefit of using continuous rather than discrete time is that the time between events can be measured precisely and that it allows modellers to more accurately study causation (Zaidi and Rake, 2001, p. 15).

Few continuous-time microsimulation models exist. They include the SOCSIM model developed by Hammel et al. at Berkeley (Hammel et al., 1976; Hammel 1990)^{3,4}, the demographic PopSim part of the DYNAMOD model developed at NATSEM (Antcliff 1993), MICROHUS of Uppsala University (Klevmarken and Olovsson, 1996), LifePaths of Statistics Canada (Gribble, 1997; Statistics Canada, 2001) and PENSIM of the US Department of Labor (Holmer et al., 2006). PENSIM uses the same algorithm as LifePaths (Holmer et al., 2006, p.3). The algorithm consists of drawing a sample of waiting times to event and comparing waiting times, generated by hazard models, to determine the timing and sequence of events. For a description of several of these models, including LifePaths, see Zaidi and Rake (2001). Researchers at Statistics Canada also developed a general-purpose microsimulation environment, called Model Generator (ModGen). This environment provides a common code-base for modellers which they can use to generate microsimulation models that are variants of LifePaths. Statistics Canada uses this environment to generate several special-purpose models such as the Population Health Model (POHEM) that uses the demographic module of LifePaths but replaces the mortality equations with a highly detailed model of morbidity and mortality. Dynamic microsimulation models generate life histories of individuals and generally use longitudinal data. In LifePaths, the life histories are for members of different birth

half of the year. Deaths, migrations, changes in status of co-residence with parents, marital status transitions, and changes in the number of surviving and co-residing children due to children's death or leaving or returning home are assumed to occur at the middle of the year. These probabilities of occurrence of the events refer to the whole year and depend on the status at the beginning of the year (Zeng et al., 2006). For details, see Zeng (1991).

³ For an extensive bibliography, see www.demog.berkeley.edu/~wachter/socrefs.html

⁴ In SOCSIM, time is measured in integral months (Wachter et al., 1998, p. 10). The same approach is used in DYNAMOD (Kelly, 2003, p. 4).

cohorts. For a nice illustration of the use of transition (hazard) models in microsimulation, see Rowe and Nguyen (2004).

Zaidi and Rake assert that “The LifePaths’s choice of the continuous time is definitely desirable from a theoretical point of view, although the use of continuous time puts heavy demand on the underlying data and computer resources.” (Zaidi and Rake, 2001, p. 16). In continuous-time simulation models, waiting times to events must be generated from transition rate models that are estimated from empirical data on timing of events. The description of LifePaths by Rowe and Nguyen (2004) and PENSIM by Holmer et al. (2006) are illustrative for the treatment of the interface between empirical transition rate models (based on observations) and transition rate models used in simulation. The generation of waiting times involves several practical issues. Galler (1997), who reviewed several models, concludes that a discrete-time approach based on comparatively short periods appears to be better suited to dynamic microsimulation modelling than a continuous-time framework.

Microsimulation is often contrasted with macrosimulation. In macrosimulation, the fundamental unit of analysis is a group of individuals with similar attributes (e.g. age, sex, race, marital status, health status) and the aim of dynamic models is to update the sizes of the groups. The life table and demographic projection models (cohort-component model) are macrosimulation models and the multistate life table and the multistate projection model are transition models. The life table updates the sizes of groups along a cohort path. The outcome is a (*synthetic*) *cohort biography* that indicates how the composition of a birth cohort changes with age. In microsimulation, the unit of analysis is the individual, and in dynamic models the characteristics of individuals and links between them are updated. The outcome of updating the characteristics of an individual is a (*synthetic*) *individual biography* that indicates how the attributes of an individual change with age. The updating is usually implemented using transition models that explicitly consider the transfers between groups (macro-level) that are associated with changes in attributes (micro-level). Examples of macrosimulation models that distinguish population groups and explicitly account for transitions between groups are LIPRO (van Imhoff and Keilman, 1991), MUDEA (Willekens and Drewe, 1984; Willekens, 1995) and ProFamy (Zeng et al., 1997, 1998, 2006). Bartlema (1989) and Lubitz et al. (2003) incorporate a microsimulation component into a macrosimulation model. The MicMac model fully integrates microsimulation into macrosimulation. Mac is a transition rate model for macrosimulation, similar to LIPRO and MUDEA. The parameters of the model are rates of transition that measure the propensity of members of a group to transfer between states. Mic is a transition rate model for microsimulation. It starts from the same transition rates as Mac but introduces heterogeneity (additional attributes). Members of a group differ in several ways and the differences are accounted for in the microsimulation.

In the sampling perspective on microsimulation, microsimulation involves drawing a sample from an artificial or virtual population that is very similar to a real population. Demographic and statistical models describe the virtual population. The parameters of these models are estimated from observations on a real population. The observations may include censuses, surveys and administrative records of the population. The more

information on the real population is incorporated in the virtual population, the more realistic is the virtual population. Several sources of data may be combined to estimate the parameters of the models that describe the real population. The virtual population is referred to as a *synthetic* population since the models that describe the population may be based on observational data from different sources. The more accurate the model(s), the more realistic is the synthetic population. If the virtual population incorporates detailed and accurate information, it cannot be distinguished from the real population, i.e. the real population and the virtual population have the same parameters. In that case, a sample drawn from that virtual population is undistinguishable from a sample drawn from the real population. The virtual population serves as a simulator, a device that incorporates important features of the real population and can be used for training, research, impact assessment and forecasting. For instance, the simulator (population ‘in vitro’ or ‘in silico’) may reveal characteristics of the population that are difficult to observe ‘in vivo’, such as pathways (sequences of events) and times between events. The simulator may also be used to assess the consequences on group characteristics and attributes of members of the group of changes in the transition rates that are the parameters of the model.

The resemblance between the virtual population and the real population is limited to the variables included in the model and the number of categories of each variable. Attributes not incorporated in the model are also omitted from the virtual population generated by the model. As a consequence, these attributes cannot be studied in the virtual population. The main variables of transition models are known as *state variables* and the possible values as the state space. Note that the state space determines the attributes of the virtual population at a given point in time whereas the transition rates determine the population dynamics. Several indicators may be used to describe the population groups (macro) and members of the group (micro). At the individual level, the indicators include the state occupied at a given age (state occupancy), the length of the interval or episode between two events, the sequence of events or, equivalently, the sequence of episodes. At the population level, the indicators include the size of a group at a given point in time or at a given age (state occupancy), the proportion of individuals in a given state, the average waiting time to an event, the average length of the interval between events, and the distribution of the lengths of episodes, which is a waiting time distribution. If the same state space is used at the macro level and the micro level, the characteristics at the macro and the micro levels are consistent: the group size and the number of events in the sample (virtual population) are equal to the expected group size and the expected number of events produced by the macrosimulation model (provided the sample is sufficiently large). If the macrosimulation model and the microsimulation model have different state spaces, or if the assumptions underlying the models are different, the updates are not consistent. In that case, consistency may be enforced by alignment. For extensive discussion of alignment and related methods in microsimulation the reader is referred to Fredriksen (1998, pp. 115ff) and Baekgaard (2002).

A critical aspect of dynamic microsimulation is the generation of events at the individual level, i.e. to determine when a member of the virtual population with a given attribute experiences a change in attribute and what the attribute is after the change. That problem

is solved by drawing a random number from a probability distribution. The method is commonly referred to as the Monte Carlo technique. In dynamic simulation models that treat time as a discrete variable, time intervals are distinguished and the problem is to determine whether a given member of the virtual population at the beginning of a given interval experiences a change of attribute (event) during that interval or not. The probability of an event is known for the group to which the individual belongs. To solve that problem the computer generates a random number between 0 and 1 with a uniform distribution. If the number is smaller than the probability, the event occurs. Most simulation models treat time as a discrete variable and follow this procedure. When time is treated as a continuous variable, the uniform distribution is generally not suitable. The time dependence of the transition rates determines the probability distribution to be used⁵. Once the probability distribution is selected, the inverse distribution function or quantile function is used to translate a probability of transition into a waiting time to transition that is used to determine whether an event occurs during a given interval. The purpose of this paper is to review and illustrate methods to simulate events and sequences of events in continuous time. The basic exponential transition rate model is used for that purpose. The model is discussed in the next section.

3. Exponential waiting times

Individuals are described by attributes and attributes change when events occur. An event involves a transition from an origin state to a destination state. Any event may be characterized by the time at occurrence, the origin state and the destination state. First we consider a single origin and a single destination. Next we consider multiple destinations. Subsequently, multiple origins and multiple destinations are considered. We assume that the events occur at a constant rate. The interest is in events that occur during a given interval from 0 to h .

3.1 Single origin and single destination

The single event occurs at a constant rate μ . Let the random variable T denote the waiting time to the event or transition. Observations on times to event are manifestations of T . A simple model of T is the basic exponential transition rate model (see e.g. Blossfeld and Rohwer, 2002, Chapter 4). It has a single parameter, which does not vary in time. The parameter is the transition rate; it is estimated from observed times to event.

If the rate of transition is constant then the time to event follows an exponential distribution. The exponential distribution is one of the most widely used distributions in statistical practice. The distribution is thoroughly documented by Balakrishnan and Basu (1996) and Evans et al. (2000). The survival function $S(t)$ is the probability that the event does not occur during the interval from 0 to t . It is

⁵ The transition probability density may be used instead of the transition rate. Methods of transition data analysis that are generally referred to as non-parametric such as the life table, are in fact methods that implicitly use models with piecewise constant transition probability densities. The survival function that results is a linear function of time from the beginning to the end of an interval.

$$S(t) = \exp[-\mu t]$$

The distribution function $F(t)$ is the probability that the event occurs (at least once) during an interval of length t (from 0 to t). The distribution function is $1 - S(t)$.

$$F(t) = 1 - \exp[-\mu t]$$

A distribution function of a random variable maps a real number (a particular value of that random variable or observation on that random variable) into a probability. The real number that is mapped into a probability is referred to as the *quantile* of the random variable (see Evans et al., 2000, p. 5). In the above distribution function, the real number is time. Hence, t is the quantile of T . With t is associated a probability, α say, and the distribution function gives the probability that the time to event is less than t . The inverse distribution function or *quantile function* maps a probability (α) into a real number (t). The inverse distribution function, denoted by $G(\alpha)$, is the value of t (quantile) such that the probability that T takes on a value less than or equal to it, is α :

$\Pr\{T \leq G(\alpha)\} = \alpha = F(G(\alpha))$ where $G(\alpha)$ is a value of T . $G(\alpha)$ is the 100α percentile. The inverse survival function $Z(\alpha)$ is the quantile that is exceeded with probability α :

$\Pr\{T > Z(\alpha)\} = \alpha = S(Z(\alpha))$. Inverse distribution functions are widely used in statistics, for instance to determine confidence intervals. Note that $Z(\alpha) = G(1-\alpha)$.

The inverse distribution function is used to generate random numbers from a distribution. If T is a random variable denoting time and $G(T)$ is the inverse distribution function of T , then $U = G(T)$ follows a uniform distribution on the interval from 0 to 1. That distribution is abbreviated as $U \sim U[0,1]$. In a Monte Carlo microsimulation a random draw from a distribution function involves two steps. First, a random number α is drawn from a uniform distribution over $[0,1]$. That random number is the probability (α) that the event occurs in an interval of length h . Second, using the inverse distribution function $G(\alpha)$, the probability is mapped into a real number t . If t is less than the time till the end of the interval, the event occurs at t . Otherwise the event does not occur.

Consider the distribution of waiting times to events when the transition intensity is constant. If the value of μ is fixed, T is exponentially distributed. The inverse distribution function of T is

$$F^{-1}(t) = G(\alpha) = -\frac{\ln[1-\alpha]}{\mu}$$

For a given value of α , $G(\alpha)$ gives the value t for which $F(t) = \alpha$. Once the exponential distribution is specified, we can draw many waiting times t from that theoretical distribution. Each draw represents an event or direct transition at t . If t is beyond the interval specified, the event does not occur. Consider an interval of length h . If the randomly drawn value of t is less than h , the event occurs during the interval, and the time to event is t . If t exceeds h , the event is not counted. If the event is a repeatable event, it may occur more than once during an interval of length h . Let t_1 denote the time at first occurrence. The probability that the event occurs a second time during the interval is the probability that it occurs during an interval of length $h - t_1$. That probability is $q(0, h - t_1) = 1 - \exp[-\mu(h - t_1)]$

A random number α is drawn from the uniform distribution $U \sim U[0,1]$ and the value of t_2 is determined that is consistent with that random number α . It is

$$t_2 = -\frac{\ln[1-\alpha]}{\mu}$$

If $t_2 < h - t_1$, the event occurs a second time during the interval h , otherwise it does not. The sequence of events during the interval can be simulated in a similar way.

The assumption of fixed transition rates is for presentation only. Hazard rates are generally assumed to be piecewise constant, i.e. constant during intervals of a given length, usually one or five years. In that case, the exponential distribution is a step function with parameters that differ between intervals. For a particular interval, the random draws from the particular exponential distribution are kept only if the event time is in the interval. This specification has been adopted in DYNAMOD, SOCSIM, LifePaths and PENSIM.

To illustrate continuous-time microsimulation, consider a random variable T the values of which follow an exponential distribution with event rate 0.2. Let time be measured in years and fractions of a year. First we present several indicators that characterize the distribution. The indicators are probability measures and duration measures. The values of these indicators are *expected* values that are characteristic for the distribution. They are determined analytically. Next we present the same indicators but obtained by sampling the probability distribution, i.e. by obtaining random draws from the distribution. The expected values serve as benchmarks to assess the results of the microsimulation. In the absence of Monte Carlo variation, the sample values should coincide with the expected values.

The probability of an event within a year (time unit), given the event rate of 0.2, is $1 - \exp(-0.2) = 0.1813$. The expected waiting time to an event is $E[T] = 1/0.2 = 5$ years. It is

$$E[T] = \int_0^{\infty} \exp(-\mu\tau) d\tau = \int_0^{\infty} \exp(-0.2\tau) d\tau = \left[-\frac{1}{0.2} \exp(-0.2\tau) \right]_0^{\infty} = \frac{1}{0.2} = 5$$

The median waiting time to the event, i.e. the time at which there is a 50 percent chance that the event occurred, is given by the inverse distribution function:

$$G(0.5) = -\frac{\ln[1-0.5]}{0.2} = \frac{\ln(2)}{0.2} = 3.466 \text{ years. The probability of an event reaches 25}$$

percent at the upper quartile [$G(0.25)$] which is 1.44 years; the probability of an event reaches 75 percent at the lower quartile [$G(0.75)$] which is 6.93 years. If the event rate is 0.2, there is a 90 percent probability that the event occurs before 11.5 years. Note that the probability that the event occurs before the expected waiting time is 63.2 percent. The probability that no event occurs between the start of the interval and τ ($0 \leq \tau \leq 1$) is the survival function $S(\tau) = \exp[-0.2\tau]$.

The expected waiting time to event during a period of one year is the total time expected to be spent in the origin state during one year. It is

$$\begin{aligned}
E[T] &= \int_0^1 \exp(-\mu\tau) d\tau = \int_0^1 \exp(-0.2\tau) d\tau \\
&= -\frac{1}{0.2} \left[\exp(-0.2\tau) \right]_0^1 = -\frac{1}{0.2} [\exp(-0.2) - 1] = 0.9063 \text{ years}
\end{aligned}$$

The time spent in the origin state during one year depends on whether or not the event occurs. To determine the sojourn time in the origin state during one year, provided the event occurs during that year, $E[T]$ is written as the weighted sum of two sojourn times. The first is the sojourn time in the absence of the event and the second is the sojourn time provided the event occurs:

$E_e[T] = 1 * S(1) + E_e[T][1 - S(1)]$ where $E_e[T]$ is the expected waiting time to the event provided the event occurs during the year. It is

$$E_e[T] = \frac{E[T] - S(1)}{1 - S(1)} = \frac{0.9063 - 0.8187}{1 - 0.8187} = 0.4832 \text{ years}^6.$$

Under the exponential model, the events are concentrated in the first half of the year. If an event occurs during a year, the probability that it occurs in the first half of the year is

$$F(0.5)/F(1) = [1 - \exp(-\frac{1}{2} * 0.1)] / [1 - \exp(0.1)] = 0.525, \text{ which is 52.5 percent.}$$

If the event is repeatable, it may occur more than once during a year. Suppose the rate is fixed at 0.2 and does not vary with number of occurrences, i.e. occurrence-dependence is absent. The probability of *at least* two occurrences within a year is 1.75 percent and the probability of at least three occurrences is 0.11 percent. The probability of no occurrence during a period of one year is the survival function $\exp(-\mu) = \exp(-0.2) = 0.8187$ or 81.87 percent. The probability of precisely one occurrence is 16.37 percent and the probability of exactly two occurrences is 1.64 percent. The probability of a precise number of occurrences during the period from 0 to t is given by the Poisson distribution:

$$\Pr\{N(t) = n(t)\} = \frac{(\mu t)^{n(t)} \exp[-\mu t]}{n(t)!}$$

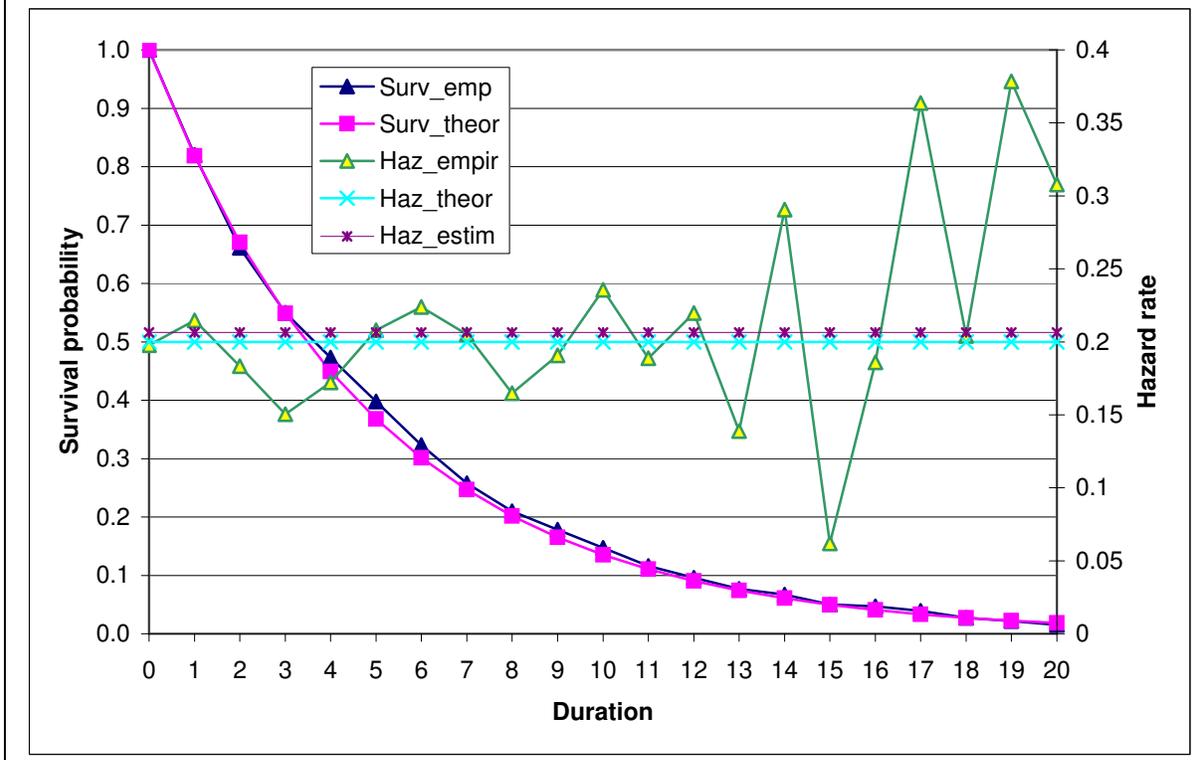
where $n(t)!$ is factorial t which is the product $1 * 2 * 3 \dots * n(t)$.

The time intervals between occurrences are independent and exponentially distributed. Let D_n denote the duration between the $(n-1)^{\text{st}}$ and the n -th occurrence. The expected length of the interval between any two occurrences is $1/\mu$. The time to the occurrence of rank n , T_n , is the sum of independent exponentially distributed variables D_n and is a gamma-distributed random variable. Hence the probability of n occurrences during an interval from 0 to t is the gamma distribution with parameters μ and n .

Consider a random sample of 1000 subjects from the virtual population that experiences a single, non repeatable event at a constant rate of 0.2 per year. The (waiting) time to event is drawn from an exponential distribution with parameters $\mu = 0.2$, using the method described above. Figure 1 shows the theoretical and *empirical* survival function and hazard function.

⁶ From the waiting time, the exact date of the event may be determined (Annex A). An event that occurs 0.4832 years since 1st January, occurs on 25th June.

Figure 1
The exponential model ($\mu = 0.2$)



The empirical survival function (Surv_emp) is close to the theoretical function as expected. The empirical transition rates vary erratically around the theoretical value of 0.2, especially at higher durations (times to events), which are drawn less frequently than lower durations. The theoretical transition rate (Haz_theor) is constant at 0.2. The transition rate is estimated from simulated waiting time data using the basic exponential transition rate model described by Blossfeld and Rohwer (2002, Chapter 4) and implemented in TDA (available at <http://www.stat.ruhr-uni-bochum.de/tda.html>). The estimated transition rate is 0.2041⁷ (Haz_estim), which is slightly different from the theoretical value of 0.2.

Suppose the event is repeatable and we want to determine for each subject the number of occurrences during a period of one year. Each subject in the virtual sample is followed (observed) for a period of one year and event occurrences are recorded. For each event that occurs during the year, the time to event and the rank of the event are recorded. The values of the indicators described above are determined from the sample and the sample

⁷ The basic exponential transition rate model is $\ln(\mu) = a$ and the estimate of a is -1.5895 (standard error is 0.0100). The coefficient a is estimated from a sample of 1000 random draws from an exponential distribution.

values are compared to expected values that are based on the theoretical distribution. Table 1 shows the results for three samples and the expected values derived from the exponential distribution and the Poisson distribution, whatever distribution applies. The expected values of the number of events are obtained by multiplying the sample size by the probabilities predicted by the Poisson distribution. In the first sample, 829 subjects did not experience an event during the year and 71 experienced at least one event. Most of the subjects who experienced at least one event experienced a single event, 18 experienced 2 events and 1 experienced 3 events. The total number of events experienced by the 1000 subjects during that one year is 191 ($152 + 18*2 + 1*3$). In the second random sample, 203 subjects experienced at least one event. Of them, 189 experienced a single event, 12 experienced two events and 2 had 3 occurrences during the year. The total number of events in that year is 219. The *expected* distribution of the subjects by number of events is given in the last column. The number of events that the 1000 subjects may expect to experience during the interval of one year, given the event rate of 0.2, is 200. ($=1000 * 0.2$).

The *expected* time to the first occurrence during the year, provided the event occurs, is 0.483 years (see above). The expected times to subsequent occurrences is more difficult to obtain analytically. The sample mean of the waiting times to the event follows directly from the microsimulation (sampling). In Table 1 the values are shown for the three samples. Note that in Sample 3, the time to the third occurrence is less than the time to the second occurrence. The reason is that the two subjects that experience three occurrences during the year experience the first and the second occurrence at earlier than the other subjects. Note also that the difference between the times at two consecutive occurrences does not yield the interval between events. To determine the intervals, the times to events must be conditioned on event occurrences. For instance, subjects that do not experience a second event during the interval should be excluded from the calculation of the time between the first and second event. Consider sample 3. The 17 subjects that experience at least two events, experience the first and the second event earlier than subjects that experience a single occurrence. To illustrate the relation between the times to events and the number of events, consider Sample 1. The 18 subjects that experience two events during the year experience the first event at 0.448 years, which is earlier than the overall average of 0.504 years. The 152 subjects with a single event during the year experience the event at 0.511 years, on average. The subject with 3 occurrences during the year experience the occurrences at 0.51 years, 0.60 years and 0.96 years. The observation that subjects with more events during a given period experience the first event earlier than other subjects is a general one. For instance, it is well-known that women of a given age with several children started childbearing earlier, on average, than women with one or two children. Similarly, women who start childbearing early are more likely to have more children than women who start late. The observation may at least in part be attributed to probabilistic factors and cannot be attributed entirely to behavioural factors. Using Monte Carlo microsimulation, it is possible to determine how much of the relation between family size and age at childbearing can be attributed to random factors and how much to behavioural factors.

Table 1				
Number of occurrences and times to event				
Random samples of 1000 events and expected values				
Number of subjects by number of occurrences within a year	Random sample 1	Random sample 2	Random sample 3	Expected values
0	829	797	828	819
1	152	189	153	164
2	18	12	17	16
3	1	2	2	1
4	0	0	0	0
5	0	0	0	0
Total	1000	1000	1000	1000
Total number of occurrences within a year	191	219	193	200
Time to event				
1	0.504	0.478	0.483	0.483
2	0.672	0.705	0.700	
3	0.960	0.740	0.596	
4	-	-	-	
5	-	-	-	

3.2 Multiple origins and multiple destinations

If an origin state may have several exits or if an event may result in multiple destinations, each exit or destination may be viewed as competing to be the exit or destination. In other words, the exits or destinations represent competing risks. Three examples illustrate the competing risks. First, the event of death may be due to various causes and each cause of death represents a competing risk. Second, marriage may be dissolved because of divorce or death of the spouse. Divorce and death of the spouse are competing risks of marriage dissolution. Third, when an individual migrates, different destinations are possible and they are competing to be the final destination. Destination-specific migrations represent competing risks. In the presence of competing risks, the type of event must be determined in addition to the time to event. The time to event is an exponential random variable, as in the previous section. The type of event is a multinomial random variable (or binomial in case of 2 competing risks). The competing risk model is generally formulated in terms of latent times to event (see e.g. Klein and Moeschberger, 2003, pp. 50ff). The approach is also used in the presentations of LifePaths. Let T_j , $j = 1, 2, \dots, J$ be the unobservable time to occurrence of the j -th event-type or competing risk, where J is the total number of event-types (exits, destinations). T_j is a latent variable. In the theory of competing risks, observations on events consist of (1) the shortest time to event, i.e. $T = \min(T_1, T_2, \dots, T_J)$ and (2) the risk that causes the event to occur. That risk is represented by an indicator

δ . If the j -th risk causes the event to occur, then $T = T_j$ and $\delta = j$. The basic competing risk parameter is the hazard rate for risk j

$$\mu_j(t) = \lim_{\Delta t} \frac{P[t \leq T < t + \Delta t, \delta = j | T \geq t]}{\Delta t}$$

which states that the cause-specific or destination-specific hazard rate is the joint probability of the event occurring in the small interval Δt and the probability that the cause, exit or destination is j , provided that the event has not occurred before t . The total hazard rate is

$$\mu(t) = \sum_{j=1}^J \mu_j(t)$$

The cause- or destination-specific hazard rate at time t may be written as the product of the total hazard rate $\mu(t)$ and a probability $p_j(t)$ that the cause of the event or the destination state after the event is j , provided the event occurs at time t :

$$\mu_j(t) = \mu(t) p_j(t)$$

In continuous-time microsimulation, two approaches may be distinguished to determine the time to event and the type of event. The first, used in e.g. LifePaths, uses the cause- or destination-specific hazard rates and generates times to events for every cause or destination. The shortest waiting time in the presence of competing events is selected. If the shortest waiting time is t_j , then j is a candidate event. It occurs if t_j is in the interval. The second approach uses two random variables, one to denote the time to event and the other to denote which of the competing risks causes the event to occur or is the destination. Two random numbers are generated. One is a random draw from an exponential distribution to determine the timing of the event. If a random time generated is in the interval, the event occurs (see previous section). The type of event is determined by a random draw from a uniform distribution $U \sim U[0,1]$. Let the draw be denoted by u . If u is less than $p_1(t)$, the first event-type (cause of event or destination) occurs, if $p_1(t) \leq u < p_1(t) + p_2(t)$, the second event-type occurs, if $p_1(t) + p_2(t) \leq u < p_1(t) + p_2(t) + p_3(t)$, the third event-type occurs, etc.

The competing risk model may easily be extended to a multistate model with hazard rates depending on state of origin and state of destination. The transition rate $\mu_{ij}(t)$ is the rate at which individuals, who occupy state i at time t , make a transition to state j . If the *sample* consists of occupants of state i , the transition rates are conditioned on the state of origin, and the multistate model resembles the competing risk model. If the sample consists of individuals irrespective of the state occupied at t , the transition rate is not conditioned on the state of origin. As a consequence, the state of origin must be determined, in addition to the state of destination and the time to event. The transition rate by origin and destination may be written as the product of the rate of leaving a state and destination probabilities, conditional on leaving. For instance, the rate of a direct transition from state 1 to state 2 may be written as: $\mu_{12} = \mu_{1+} p_{12}$, where μ_{1+} is the exit rate from 1 and p_{12} is the probability that a subject leaving 1 transits to destination 2.

By way of example suppose the state space consists of 3 states: healthy (1), disabled (2) and dead (3). At the start of the process being simulated, all subjects are in state 1. The process is simulated for a period of 10 years (0 – 10). Suppose the transition rates are constant and equal to: $\mu_{12} = 0.12$, $\mu_{13} = 0.03$, $\mu_{21} = 0.06$ and $\mu_{23} = 0.06$. The rate of leaving state 1 (μ_{1+}) is 0.15 and the rate of leaving state 2 (μ_{2+}) is 0.12. The transition rates imply that 80 percent of the subjects leaving state 1 (healthy) move to state 2 (disability) and 20 percent move to state 3 (dead). Consider a sample of 1000 subjects in state 1 at the start of the process. The *expected* state occupancies at different years are shown in Table 2. They are calculated by the following equation:

$$\mathbf{K}(t + 1) = \mathbf{P}\mathbf{K}(t) = \exp[-\mathbf{M}]\mathbf{K}(t)$$

where $\mathbf{K}(t)$ is a vector of state occupancies indicating the number of subjects in each of the 3 states at time t . \mathbf{P} is the matrix of transition probabilities. In this illustration, they are estimated from the transition rates using the linear approximation of the exponential model (for the derivation, see Willekens, 2006): $\exp[-\mathbf{M}] \approx [\mathbf{I} + \frac{1}{2}\mathbf{M}]^{-1} [\mathbf{I} - \frac{1}{2}\mathbf{M}]$

The transition rates are assembled in the transition matrix \mathbf{M} :

$$\mathbf{M} = \begin{bmatrix} 0.15 & -0.06 & 0 \\ -0.12 & 0.12 & 0 \\ -0.03 & -0.06 & 0 \end{bmatrix}$$

The matrix of transition probabilities for one-year intervals is

$$\mathbf{P} = \begin{bmatrix} 0.863 & 0.053 & 0.000 \\ 0.106 & 0.899 & 0.000 \\ 0.031 & 0.058 & 1.000 \end{bmatrix}$$

The expected state occupancies at the beginning of each year from 0 to 10 are given in Table 2.

Year	Healthy	Disabled	Dead
0	1000	0	0
1	863	106	31
2	751	187	64
3	658	247	98
4	582	292	133
5	518	324	168
6	464	346	203
7	419	360	237
8	381	368	271
9	348	371	304
10	332	358	310

In the hypothetical example, after a period of 10 years, the subjects are relatively evenly distributed between the states.

Consider again the sample of 1000 subjects in state 1 at the beginning of the process. Assume that the subjects do not differ with respect to the transitions between healthy, disabled and dead. All healthy people experience the same incidence rate of disability and the same death rate. Disabled persons experience the same recovery rate and death rate. We construct the lifepath of the 1000 subjects during a 10-year period. Continuous-time microsimulation is used to determine, for each subject, the time to transition to disability, recovery, or death. Two random variables are generated. The first is the time to transition drawn from an exponential distribution that is characteristic for the state of origin. Healthy subjects leave the state of being healthy at a time that is determined by $\mu_{1+} = 0.15$. For each subject, the exit time is drawn from an exponential distribution with parameter $\mu_{1+} = 0.15$. The destination state is determined by drawing a random variable from a uniform distribution. If the random variable is between 0 and 0.8, the healthy subject who discontinues to be healthy becomes disabled. If the random variable is between 0.8 and 1.0, the subject dies. The recovery and death of subjects with disability are determined in a similar way. The time to transition or exit time is drawn from an exponential distribution with parameter $\mu_{2+} = 0.12$ and the direction is determined by drawing a random variable from a uniform distribution. Recovery occurs when the random variable is between 0 and 0.6. If the value of the variable exceeds 0.6, the subject dies. Table 3 shows the state occupancies every year of the period from 0 to 10. The sample observations are close to the expected values. Note that all virtual subjects are followed for a period of 10 years. Censoring is at year 10.

Year	Healthy	Disabled	Dead
0	1000	0	0
1	860	105	35
2	759	174	67
3	663	241	96
4	592	283	125
5	530	315	155
6	472	338	190
7	428	356	216
8	398	361	241
9	364	364	272
10	344	350	306

Of the 1,000 healthy subjects at the start of the process, 218 remain healthy for the entire period of 10 years since they do not experience an event. 782 becomes disabled for at least some period and/or die. At the end of the observation at year 10, 344 are healthy, 350 are disabled and 306 are dead. Of those who are healthy, 126 (=344-218) experience at least one episode of disability and recovered.

The sampling from the virtual population (microsimulation) gives information on the population that cannot be obtained otherwise. The 1000 subjects experience 1228 transitions (Table 4). Most are from healthy to disability. 710 transitions initiated an episode of disability. Of the 710 episodes of disability, 211 end in recovery, 149 in dead and 350 are truncated when year 10 is reached. The table also shows the health status at death: 158 are healthy and 149 disabled.

Table 4 Transitions, sample				
ORIGIN	DESTINATION			Total
	Healthy	Disable	Dead	
Healthy	0	710	158	868
Disable	211	0	149	360
Dead	0	0	0	0
Total	211	710	307	1228

The healthy subjects that become disabled during the observation window of 10 years, become disabled after 4.1 years, on average. Disabled subjects that recover, recover at 5.9 years, on average. Although disabled subjects have a mortality that is twice that of healthy subject, they die later than healthy subjects: 6.1 years versus 3.9 years. The difference is due to the relatively late onset of disability and the low recovery rate. As a consequence, many healthy subjects die before the mean age at onset of disability.

A major advantage of continuous-time microsimulation is the possibility of multiple transitions within a year. The number of multiple transitions is relatively rare: 6.5 percent of the subjects experience two transitions and 0.7 percent three transitions.

The total number of years spent alive during the observation period is 8.45 years. An average subject spends 5.7 years healthy and 2.7 years disabled.

3.3 Covariates and interventions

Hazard rates generally depend on personal attributes or covariates. Time-invariant covariates include sex and place of birth. Time-varying covariates include level of education, marital status, employment status, place of residence and health status. The general specification of the exponential transition model is

$$\mu_{ij}(t) = \exp[\alpha_{ij} + \beta_{ij1}X_{ij1} + \beta_{ij2}X_{ij2} + \beta_{ij3}X_{ij3} \dots + \gamma_{ij1}Z_{ij1}(t) + \gamma_{ij2}Z_{ij2}(t) + \gamma_{ij3}Z_{ij3}(t) \dots]$$

where X_{ijk} denotes the value of the k-th time-invariant covariate and $Z_{ijk}(t)$ the value of the k-th time-varying covariate at time t. When the time-varying covariates are discrete variables, as is generally the case in social sciences, the coefficients γ_{ijk} vary with t. In the presence of time-varying covariates, the values of the covariates must be given at all time point, i.e. the entire *covariate path* must be supplied (concomitant observations).

Exponential transition models with time-varying covariates are discussed by Blossfeld and Rohwer (2002, Chapter 6).

The time-path of time-varying covariates is a continuous-time process. It may be approximated by a discrete-time process when piecewise constant hazard rates are used and the covariates are allowed to change at the beginning of time intervals only. In that case, the values of the covariates are updated at the beginning of each interval and the hazard rate is obtained depending on the new values of the covariates. When covariates are allowed to change at any time during the interval, i.e. in continuous time, the interval is split in two or more subintervals and the hazard rates are derived for each subinterval. In other words, the hazard rates are updated whenever covariate values change. That procedure of interval splitting is similar to episode-splitting in event-history modeling (Blossfeld and Rohwer, 2002, pp. 140ff). The technique involves the splitting of episodes at every point in time where one of the time-varying covariates changes its value. Each of the original episodes is replaced by a contiguous set of subepisodes (splits) with appropriate values of the covariates. Interval splitting is implemented in LifePaths.

By way of example, consider the disability model and consider an intervention programme that reduces the incidence of disability and the rate of recovery once disability has struck. Subjects enroll in the programme at different ages and they remain enrolled till the end of the study period, which is 10 years. If a healthy subject enrolls in the programme, the incidence rate of disability drops by 50 percent. Hence the rate after enrollment is 0.5 times the rate before enrollment. Subjects that are disabled and enrolled in the programme have a higher rate of recovery. The rate is assumed to be 3 times the rate for subjects not enrolled. Since subjects may enroll at any age, being enrolled is a time-varying covariate. The transition rates after enrollment may be written as

$${}_k m_{12}^e(x) = m_{12}(x) \exp[(\ln 0.5) * {}_k X(x)]$$

$${}_k m_{21}^e(x) = m_{21}(x) \exp[(\ln 3.0) * {}_k X(x)]$$

where ${}_k X(x)$ is the time-varying covariate that is equal to one if subject k is enrolled at age x and is 0 otherwise.

To determine who enrolled in the programme and at what age, i.e. to determine the values of ${}_k X(x)$, a random sample is drawn from the virtual population. It is assumed that, at the population level, 10 percent of the subjects not yet enrolled at the beginning of a year enroll in the programme. The enrollment rate is independent of the health status, but of course depends on the enrollment status. If the treatment programme is conditional on participation in the prevention programme, only healthy subject may enroll and the enrollment rate is dependent on the health status. Since the enrollment rate is independent of the health status, the expected proportion of subjects enrolled after a period of 10 years is 61.4 percent ($=100 * [1 - 1/(1+0.10)^{10}]$). In the sample, 712 subjects enrolled, i.e. 71.4 percent. Table 5 shows the number of new yearly enrollments in the sample of the virtual population.

Table 5 Number of new enrollments by year	
Year	Enrollment
0	124
1	84
2	80
3	83
4	62
5	58
6	51
7	64
8	34
9	42
10	30
TOTAL	712

In the presence of the intervention programme, less subjects in the *sample* of the virtual population enter disability (613 versus 710) and more recover from disability (231 versus 211). More subjects die before the end of the observation period (324 versus 307) and more die healthy (180 versus 158) because more subjects are healthy and are healthy longer. In the sample, healthy subjects die later (4.6 years versus 3.9 years) and disabled subjects die earlier (5.7 years versus 6.1 years). In the sample population, more subjects may expect to be healthy after 10 years (437 versus 344) and less are disabled (239 versus 350). Because of the intervention, subjects spend more years healthy (6.3 years versus 5.7 years) and less in disability (2.2 years versus 2.7 years). The expected number of years spent alive during the observation period decreases a little (8.41 years versus 8.45).

If the entire population enrolls in the intervention programme at the start of the observation period, the effect is more significant. In that case the number of healthy years is 7.2 and the number of years in disability is 1.3. The total number of years lived during the observation period is 8.45.

A slight variation of the previous model is a model with a baseline hazard and an exponential function of covariates. Let $\mu(t) = \mu^0 \exp[\boldsymbol{\beta}' \mathbf{X}]$, where μ^0 is the constant baseline hazard rate, \mathbf{X} is the vector of covariates and $\boldsymbol{\beta}$ is the vector of regression coefficients. The quantile function is (Bender et al., 2005)

$$t = -\frac{\ln(1 - \alpha)}{\mu^0 \exp[\boldsymbol{\beta}' \mathbf{X}]}$$

with α a probability drawn from a uniform distribution. The quantile

function determines the waiting time to an event that is consistent with a probability of the event when the probability depends on a constant transition rate and a set of time-independent covariates.

4. Other distributions of waiting times

The exponential distribution of waiting times implies (piecewise) constant hazard rates. When the hazard rate depends on time and the time-dependence is defined by a parametric model, the generation of random waiting times depends on the parametric model. Blossfeld and Rohwer (2002, Chapter 7) describe several parametric models of duration-dependence. The Gompertz, the Weibull and the Cox proportional hazard model are common. The Gompertz model is considered in this paper. For a review of the Gompertz model and a comparison with related duration models, see Willekens (2001). Mueller et al. (1995, p. 558) also generate a sample from a virtual population with Gompertz transition rates. They use discrete-time microsimulation, however. The generation of survival times for simulation studies is not possible when the baseline transition rates are not specified, as in the Cox model. Most simulation studies of the Cox model assume a parameterized baseline hazard, such as a constant hazard (exponential model) and an exponentially declining hazard (Gompertz model). Recently, Bender et al. (2005) developed a general formula for generating waiting times to simulate Cox proportional hazard models when the baseline hazard can be parameterized. Two cases are covered in this paper. The case of a constant baseline hazard was presented in the previous section. The case of an exponentially declining baseline hazard is presented in this section.

The Gompertz distribution of waiting times has 2 parameters, a scale parameter (μ) and a shape parameter (ν). The hazard rate changes exponentially; $r(t) = \mu \exp(\nu t)$ (with $\mu \geq 0$). If $\nu = 0$, the Gompertz distribution reduces to the exponential distribution. The survival function is

$$S(t) = \exp\left[\frac{\mu}{\nu}(1 - \exp(\nu t))\right]$$

The distribution function is $1 - S(t)$ and the quantile function is

$$T = \frac{1}{\nu} \left[\ln \left(1 - \frac{\nu}{\mu} \ln(1 - U) \right) \right]$$

where $U \sim U[0,1]$ is a random variable the values of which are uniformly distributed in the range from 0 to 1. A random draw from a Gompertz distribution is obtained in two steps. First, a random number α is drawn from a uniform distribution over $[0,1]$. Second, the value of t is derived from the quantile function. The function translates the probability α into a waiting time to event t . The value α is a realization of U and t is a realization of T .

Consider a sample of 1000 draws from a Gompertz distribution with a scale parameter (μ) of 0.20 and a shape parameter (ν) of -0.01 . The theoretical and empirical survival function and hazard function are shown in Figure 2. The figure also shows for every duration the theoretical transition rate and the transition rate estimated using the Gompertz mode. The estimates are shown in Table 6. They are obtained using TDA. The

estimated scale parameter is $\exp(-1.6536) = 0.1914$. It is the transition rate at duration 0, i.e. at the onset of the process.

The quantile function of a Gompertz distributed random variable in the presence of time-invariant covariates, is (Bender et al., 2005)

$$T = \frac{1}{\nu} \left[\ln \left(1 - \frac{\nu}{\mu^0 \exp[\boldsymbol{\beta}' \mathbf{X}]} \ln(1 - U) \right) \right]$$

It is the Cox model with a Gompertz-distributed baseline hazard.

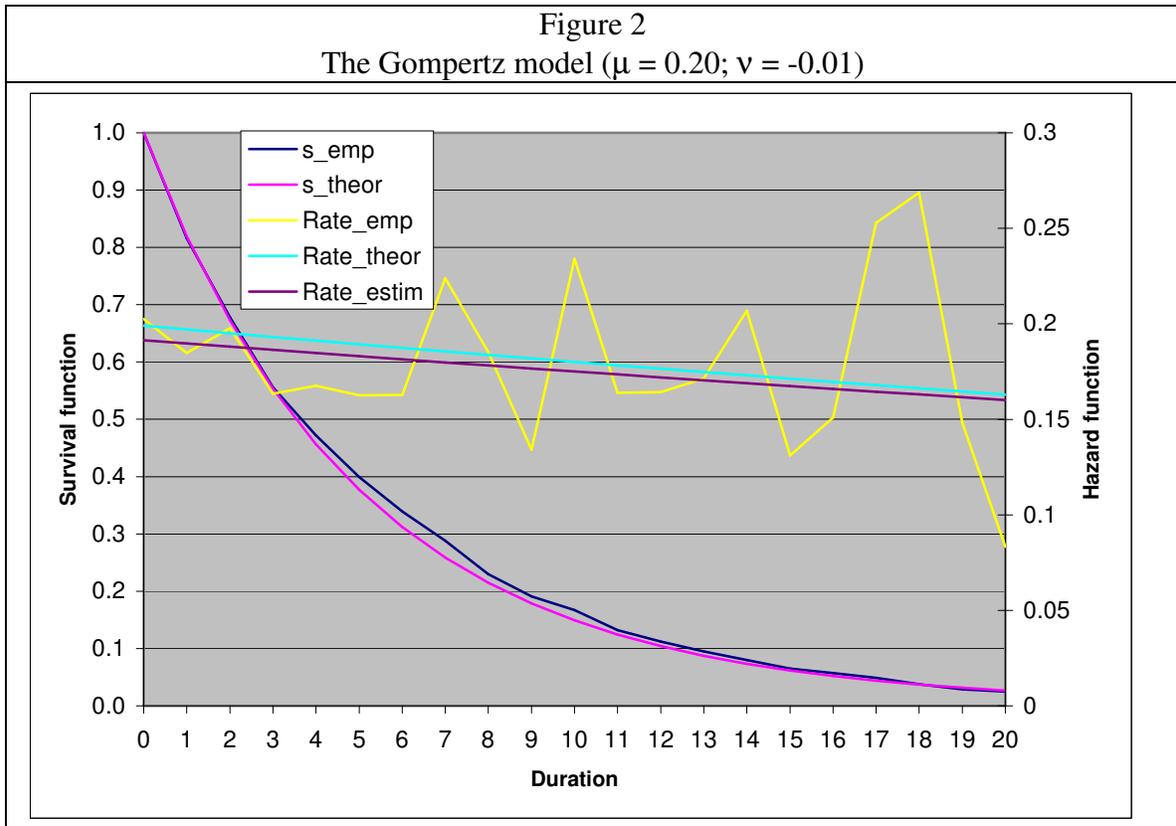


Table 6
Parameters of the Gompertz model, estimated from the sample of the virtual population

Idx	Param	Coeff	Error
1	$\ln(\mu)$	-1.6536	0.0420
2	ν	-0.0089	0.0050

5. Conclusion

Of the many dynamic microsimulation models that are operational today, few models use continuous-time microsimulation. LifePaths of Statistics Canada is a good example. It has inspired others, including the developers of MicMac. Although the algorithms that are used are adequately documented, the statistical theory that underlies continuous-time microsimulation is not. That theory is the focus of this paper. The generation of random times to events is approached as a particular application of duration analysis. The time-to-event or waiting time is a random variable that is characterized by a number of distributions, such as the distribution function, the survival function, the probability density function and the hazard function. The inverse distribution function or quantile function is the main tool for continuous-time microsimulation. The function translates the probability of an event during an interval to a waiting time to the event. In other words, it determines the waiting time to event under the assumption that the parametric model of time dependence holds. The paper illustrates the method using the basic exponential model and the Gompertz model. Both models are extensively documented in the literature. Other duration models for which quantile functions can be defined, may be applied in continuous-time microsimulation. The method based on the quantile function is a general method that applies to all waiting time models and other models as well.

Continuous-time microsimulation has several advantages over discrete-time microsimulation. First, since the exact times of events are determined, the length of intervals between events can be determined accurately and not approximately as in discrete-time microsimulation. Second, the problem of multiple events during a same interval is resolved. Since the exact dates of events are determined, the sequence of events in a same interval is also determined. Third, the problem of competing events or competing destinations is resolved by the theory of competing risks rather than by any exogenously imposed rule. The theory states that in the presence of competing events, the event with the smallest waiting time occurs. Finally, continuous-time microsimulation paves the way to an integrative framework that combines the *analysis* of life history data and the *synthesis* of life histories. That framework is rooted in probability theory and statistical theory. For many years, microsimulation has been used to study complex dynamic systems and to detect characteristics that cannot be identified by the conventional analytical methods. The sampling perspective on microsimulation points to virtual populations that resemble real populations but that have fully documented characteristics and provide a basis for controlled experiments *in silico*.

References

- Antcliff, S. (1993) An Introduction to DYNAMOD: A Dynamic Microsimulation Model, DYNAMOD Technical Paper no. 1, National Centre for Social and Economic Modelling, University of Canberra.
- Bartlema, J. (1989) Modelling stepfamilies: First results. NIDI Report no. 3, NIDI, The Hague.
- Baekgaard H. (2002). Micro-Macro Linkage and the Alignment of Transition Processes, *Technical Paper n°5*, National Centre for Social and Economic Modelling (NATSEM), Faculty of Management, University of Canberra, Canberra.
- Balakrishnan, K. and A.P. Basu (1996) The exponential distribution. Theory, methods and applications. Gordon and Breach Publishers, New York.
- Bender, R., T. Augustin and M. Blettner (2005) Generating survival times to simulate Cox proportional hazard models. *Statistics in Medicine*, 24:1713-1723.
- Blossfeld, H.-P. and G. Rohwer (2002) Techniques of event history modeling. New approaches to causal analysis. Second edition. Lawrence Erlbaum Associates, Mahwah, New Jersey.
- Bongaarts, J. (1987) The projection of family composition over the life course with family status life tables. In Bongaarts, J., T. Burch and K.W. Wachter eds. *Family Demography: Methods and Applications*. Clarendon Press, Oxford.
- Bonnet, C., C. Burricand, C. Colin et al., (1999) Le modèle de microsimulation dynamique DESTINIE. Série des documents de travail de la Direction des Etudes en Synthèses Economiques G 9913, Institut National de la Statistique en des Etudes Economiques (INSEE), Paris.
- Coleman, J.S. (1990) *Foundations of social theory*. The Belknap Press of Harvard University Press, Cambridge, Mass.
- Dupont, G., C. Hagneré and Touzé (2003) Les modèles de microsimulation dynamique dans l'analyse des réformes des systèmes de retraites : une tentative de bilan. N°2003-10. Paper presented at the Commission des finances de l'Assemblée nationale.
- Evans, M., N. Hastings and B. Peacock (2000) *Statistical distributions*. Third edition. Wiley, New York.
- Fredriksen, D. (1998) Projections of population, education, labour supply and public pension benefits. Analysis with the dynamic micro-simulation model MOSART. *Social and Economic Studies* no. 101, Statistics Norway, Oslo.

Galler, H.P. (1997) Discrete-time and continuous-time approaches to dynamic microsimulation reconsidered. Technical Paper no. 13, National Centre for Social and Economic Modelling (NATSEM), Faculty of Management, University of Canberra.

Gribble, S. (1997) LifePaths: A longitudinal microsimulation model using a synthetic approach. Paper presented in the conference "Microsimulation in Government Policy and Forecasting: International Conference on Combinatorics, Information Theory and Statistics" Maine Portland, Maine, USA, July 18-20, 1997

Hammel, E.A. (1990) *SOCSIM II*, Working Paper no. 29, Graduate Group in Demography, University of California, Berkeley.

Hammel, E.A., D. Hutchinson, K. Wachter, R. Lundy, and R. Deuel (1976) The SOCSIM Demographic-Sociological Microsimulation Program Operating Manual. Institute of International Studies Research Monograph No. 27, University of California, Berkeley, California.

Holmer M., A. Janney and B. Cohen (2006). PENSIM overview. Policy Simulation Group, U.S. Department of Labor, Washington, D.C.

Van Imhoff, E. and N. Keilman (1991) LIPRO 2.0: An application of a dynamic demographic projection model to household structure in the Netherlands. Sweets&Zeitlinger, Amsterdam.

Kelly, S. (2003) Forecasting wealth in an ageing Australia. An approach using dynamic microsimulation. Paper presented at the 7th Nordic Seminar on Microsimulation Models, Helsinki, June 2003.

Klein, J.P. and M.L. Moeschberger (2003) Survival analysis. Techniques for censored and truncated data. Second edition. Springer Verlag, New York.

Klevmarken N. A. et P. Olovsson (1996). Direct and behavioral effects of income tax changes - simulations with the Swedish model MICROHUS, in: A. Harding ed. Microsimulation and Public Policy, North-Holland-Elsevier, Amsterdam.

Lubitz, J., L. Cai, E. Kramarow and H. Lentzner (2003) Health, life expectancy and health care spending among the elderly. *New England Journal of Medicine*, 349:1048-1055

Mamun AA. 2001. Multistate models in public health – review and application to the Framingham Heart study. Population Research Centre. Masters Thesis Series 01-3, December 2001. University of Groningen, The Netherlands

Mueller, L.D., T.J. Nusbaum and M.R. Rose (1995) The Gompertz equation as a predictive tool in demography. *Experimental Gerontology*, 30(6):553-569.

- O'Donoghue C. (1999). Dynamic Microsimulation: A Methodological Survey, *Brazilian Electronic Economic Journal*. Available at http://www.ecomod.net/conferences/ecomod2001/papers_web/O'Donoghue2.pdf.
- Rowe, G. and H. Nguyen (2004) Longitudinal analysis of labour force survey data. *Survey Methodology (Statistics Canada)*, 30(1):105-114.
- Statistics Canada (2001). The LifePaths Microsimulation Model: An Overview. http://statcan.ca/english/spsd/LifePathsOverview_E.pdf
- Van Imhoff, E. and N. Keilman (1991) LIPRO 2.0: An application of a dynamic demographic projection model to household structure in the Netherlands. Amsterdam: Sweets&Zeitlinger.
- Van Imhoff, E. and W. Post (1998), Microsimulation methods for population projection. *Population: An English Selection*, special issue New Methodological Approaches in the Social Sciences, pp. 97-138.
- Wachter, K.W., D. Blackwell and E.A. Hammel (1998) Testing the validity of kinship microsimulation: an update. Manuscript. Available at www.demog.berkeley.edu/~wachter/
- Willekens, F.J. (1995) MUDEA. Version 2.0. Manual and tutorial. Report. Faculty of Spatial Sciences, University of Groningen.
- Willekens, F.J. (2001) Gompertz in context: the Gompertz and related distributions. In: E. Tabeau, A. van den Berg Jeths and C. Heathcote eds. *Forecasting mortality in developed countries. Insights from a statistical, demographic and epidemiological perspective*. Kluwer Academic Publishers, Dordrecht, pp. 105 –126.
- Willekens, F.J. (2006) The multistate model for biographic projections. Manuscript.
- Willekens, F.J. and P. Drewe (1984) A Multiregional model for regional demographic projection. In: H. ter Heide and F. Willekens eds. *Demographic research and spatial policy*. London: Academic Press, 1984, pp. 309-334. Reprinted in: D.J. Bogue, E.E. Arriaga and D.L. Anderton eds. *Readings in population research methodology*. Social Development Center, Chicago, and UNFPA, New York, 1993.) pp. 17.49 - 17.57.
- Zaidi, A. and K. Rake (2001) Dynamic microsimulation models: a review and some lessons for SAGE. SAGE Discussion Paper no. 2 [SAGEDP/02]. ESRC SAGE Research Group, The London School of Economics, London (available at www.lse.ac.uk/depts/sage)
- Zeng, Y. (1991) *Family dynamics in China: a life table analysis*. The University of Wisconsin Press, Madison, Wisconsin.

Zeng Yi, J.W. Vaupel and W. Zhenglian (1997) A multidimensional model for projecting family households. With an illustrative numerical application. *Mathematical Population Studies*, 6(3):187-216

Zeng, Yi., Vaupel, J.W. & Wang, Z. (1998), Household projection using conventional demographic data, *Population and Development Review*, supplementary issue: Frontiers of Population Forecasting 24: 59–87.

Zeng, Yi, K.C. Land, Zhenglian Wang and Danan Gu (2006) U.S. family household momentum and dynamics: an extension and application of the ProFamy method. *Population Research and Policy Review*, 25:1-41

Annex A

Dates of events

If the time to event is expressed in years, the fractions of a year may be converted into a date. First we convert a date (day, month, year) into a real number that denotes a date in years since a reference year and a fraction of a year. The conversion of a real number into a date is considered next.

Let the reference date be January 1, 1900. Suppose an event occurs on May 4, 1988. The exact number of years since 1st January 1900 is obtained by the following transformation (Mamun, 2001, p. 98):

$$EY = YEAR + (MONTH-1)/12 + (DAY-1)/(30.437*12)$$

May 4, 1988 is 88.3415 years since the beginning of the 20th century:

$$88 + (5-1)/12 + (4-1)/(30.437*12) = 88.34154702$$

The date in exact years may be converted back in year, month and day of occurrence, using the following formula:

$$\begin{aligned} \text{YEAR} &= \text{TRUNC}(EY) \\ \text{MONTH} &= \text{TRUNC}[(EY - \text{YEAR}) * 12] + 1 \\ \text{DAY} &= \text{TRUNC}[(EY - \text{YEAR} - (\text{MONTH}-1)/12) * 30.437 * 12] + 1 \end{aligned}$$

For instance, 88.3415470 is

$$\begin{aligned} \text{YEAR} &= \text{TRUNC}[88.3415470] = 88 \\ \text{MONTH} &= \text{TRUNC}[(88.3415470-88)*12] + 1 = 5 \text{ (MAY)} \\ \text{DAY} &= \text{ROUND}[(88.34154702-88 - (5-1)/12)*30.437*12] + 1 = 4 \end{aligned}$$

The conversion is not always perfect because it does not account for the different numbers of days in a month and the changing number of days in the month of February. Table A.1 illustrates the method.

In the text, the random draw from an exponential distribution resulted in a time to event, provided the event occurs during the year, of 0.4832. If the reference date is 1st January, the event occurs at

$$\begin{aligned} \text{MONTH} &= \text{TRUNC}[0.4832 * 12] + 1 = 6 \text{ (June)} \\ \text{DAY} &= \text{ROUND}[(0.4832 - (6-1)/12)*30.437*12] + 1 = 25 \end{aligned}$$

Hence the event occurs on the 25th June. Note that, using the equations, 0.5 years results in 1st July as it should be.

The dates in exact years may be converted in dates in Century Month Code (CMC):

$$\text{DATECMC} = (EY - 1900) * 12$$

Note that numeric DATECMC is not an integer value but a real value. Consider May 4, 1988. The date in CMC is $88.34154702 * 12 = 1060.098564$. The month is CMC 1060 and the day is $\text{ROUND}[0.098564 * 30.437] + 1 = 4$.

Sometimes, the reference date is different. For instance, in the Framingham Heart Study, which is a longitudinal study that started in 1948-50 and is widely used in epidemiology, the dates of the exams are measured in number of days since 1st January 1960. An event that occurs on 30th January 1960 occurs at 29 days (1+29). The date can be negative. For instance, an event that occurred on 3rd December 1959 occurred at day -29 (December 31 is day -1 and December 3 is day -29). These figures can be converted into exact number of years since the beginning of the 20th century:

$$\text{EY} = 1960 + \text{DATE}/365.25$$

where DATE is the date of the event in days since 1st January 1960 and EY is the date of the event in exact years (Mamun, 2001, p. 97). For instance, if the DATE of an event is -460, the event occurs in 1958 and more specifically at 1958.741.

Table A.1
Conversion of dates in exact years and vice versa

Year	Exact		Years	Re-estimated			Difference
	Month	Day		Year	Month	Day	
88	5	4	88.34155	88	5	4	0
80	1	1	80.00000	80	1	1	0
80	12	31	80.99880	80	12	31	0
16	3	2	16.16940	16	3	1	1
13	9	19	13.71595	13	9	19	0
27	12	23	27.97690	27	12	23	0
47	7	28	47.57392	47	7	27	1
13	11	8	13.85250	13	11	8	0
17	8	29	17.65999	17	8	29	0
23	3	27	23.23785	23	3	27	0
21	6	11	21.44405	21	6	11	0
26	2	9	26.10524	26	2	8	1
20	8	22	20.64083	20	8	21	1
15	7	9	15.52190	15	7	9	0
14	2	22	14.14083	14	2	21	1
14	9	16	14.70774	14	9	16	0
14	12	4	14.92488	14	12	3	1
19	5	16	19.37440	19	5	15	1
48	6	29	48.49333	48	6	28	1
11	12	27	11.98785	11	12	26	1
9	10	4	9.75821	9	10	4	0
13	10	19	13.79928	13	10	19	0
25	9	29	25.74333	25	9	29	0
26	6	5	26.42762	26	6	5	0
10	12	2	10.9194046	10	12	1	1