Estimation for univariate and bivariate reinforced urn processes under left-truncation and right-censoring

Luis A. Souto Arias
Pasquale Cirillo
Cornelis W. Oosterlee

21st January, 2021

Table of Contents

Motivation

- An **annuity** is financial product in which we pay a fixed amount of money and, in return, we receive periodic payments in the future.
- These products offer a fixed income stream **during the life** of the owner(s) of the contract (annuitant).
- Here we focus on **joint and last survivor** annuities, for which payments are made as long as **at least one of the annuitants is still alive**.

How much does it cost?

The pricing formula for a **joint and last-survivor** annuity is

$$C(x_0, y_0) = \sum_{k=0}^{\infty} \frac{\mathrm{P}_{LS}^{(x_0,y_0)}(k)}{(1+r)^k}, \tag{1}$$

with:

- r the interest rate.
- (x_0, y_0) the ages of the annuitants when they buy the annuity.
- $\mathrm{P}_{LS}^{(x_0,y_0)}(k)$ the probability that at least one of the two annuitants survives another k years, given their current age.

Gathering the data

- Estimating $P_{LS}^{(x_0,y_0)}$ requires large amounts of data about the joint life span of the spouses.

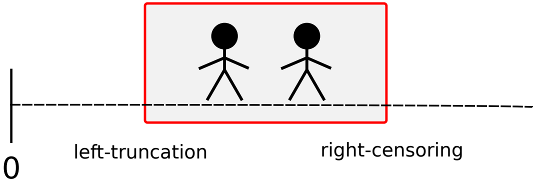

- In survival studies most of the couples are **left-truncated** and **right-censored** observations.

Topics of this talk

- Use of the **Reinforced Urn Process** (RUP) of Walker and Muliere [1997] and the bivariate extension of Bulla et al. [2007] (B-RUP).
- Estimation of the B-RUP via the **Expectation-Maximization** algorithm.
- Performance of the model with an empirical data set.
- Compare the annuity price when assuming **dependence** vs. **independence**.
- This talk is based on the following paper: `https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3593391`

Table of Contents

Urn representation

Imagine a series of Pólya urns with balls of two colors: green and red.

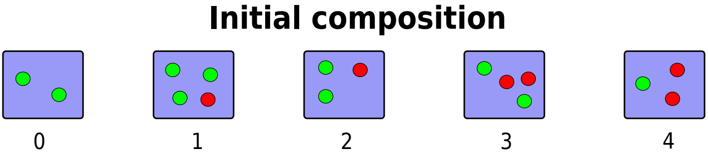

- Initial number of green balls in the j-th box: ω_j
- Initial number of red balls in the j-th box: β_j

Rules of the system

Rules:

- Sample a ball **uniformly** from the first urn.
- Add a ball of the **observed color** to the urn.
- If the ball is green go to the **next** box.
- If the ball is red **go back** to the first box.

First update I

- We sample a green ball from the first urn, put it back and add an extra green ball. We jump to the next urn.

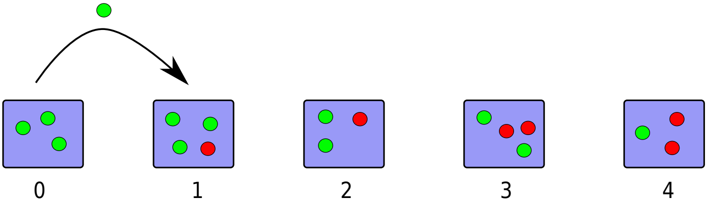

First update II

- We sample a ball from the second urn and also turns out to be green. We add an extra green ball and jump to the third urn.

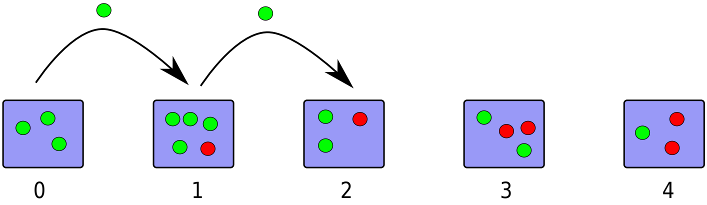

First update III

- In the third urn we sample a red ball. We add an extra red ball, set the age of death at 3, and go back to the first urn.

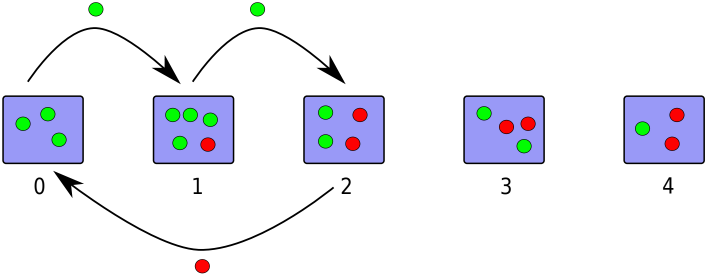

Initial composition

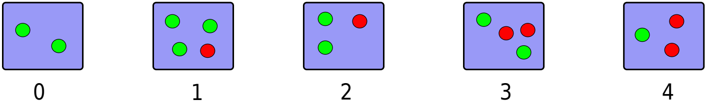

Updated composition

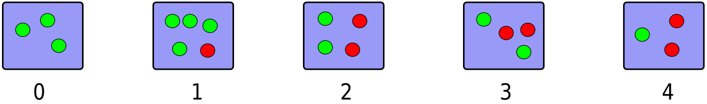

The probability of the trajectory {●, ●, ●} has been **reinforced**.

If X denotes the lifetime of an individual, and $\boldsymbol{X}_n^*$ is a series of left-truncated, right-censored observations, the posterior distribution is given by

$$\mathbb{P}(X_{n+1} > x | \boldsymbol{X}_n^*) = \prod_{j=0}^{x} \left[1 - \frac{\beta_j + m_j(\boldsymbol{X}_n^*)}{\beta_j + \omega_j + s_j(\boldsymbol{X}_n^*)} \right], \tag{2}$$

where:

- $m_j(\boldsymbol{X}_n^*)$ is the number of deaths at age $x = j$.
- $s_j(\boldsymbol{X}_n^*)$ is the number of survivors up to age j.
- $\{\beta_j, \omega_j\}$ are initial number of balls.

The role of the initial composition

The initial composition can be chosen to reflect expert's knowledge about the process:

$$\frac{\beta_j}{\beta_j + \omega_j} = \mathbb{P}_G(X = j | X \geq j), \quad j \in \mathbb{N} \tag{3}$$

where G is a prior distribution of our choice.

From balls to priors

- Possible choice is

$$\beta_j = c_j \, \mathbb{P}_G(X = j), \quad \omega_j = c_j \, \mathbb{P}_G(X > j), \tag{4}$$

with $c_j > 0$ the **strength of belief** in our prior.

- As $c_j \to 0$, the RUP tends to the **Kaplan-Meier** estimator.

RUP for coupled lifetimes

- (X, Y) denote the lifetimes of a couple of annuitants.
- Define three **independent** RUPs: A, B and C.
- Set

$$\begin{aligned} X &= A + B, \\ Y &= A + C. \end{aligned} \tag{5}$$

- The dependence between X and Y clearly relies on A, and it is **linear** and **positive**.
- We do not observe A, B and C. They need to be inferred from (X, Y).

Table of Contents

How to calibrate RUPs

Unknown composition

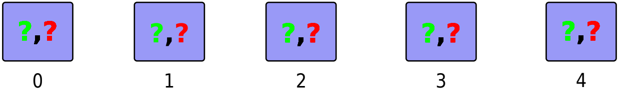

- The probability of **survival** is:

$$\mathbb{P}(X > x) = \prod_{j=0}^{x} \left(1 - \frac{R_j}{G_j + R_j}\right), \tag{6}$$

- What are the values of $\{G_j, R_j\}$ that **maximize** the likelihood of a given set of observations?

RUPs and MLE

- The optimal configuration can be obtained explicitly via MLE:

$$\frac{R_j}{G_j + R_j} = \frac{m_j(\boldsymbol{X}_n^*)}{s_j(\boldsymbol{X}_n^*)}, \tag{7}$$

  where $m_j(\boldsymbol{X}_n^*)$ and $s_j(\boldsymbol{X}_n^*)$ are defined as before.

- This is a RUP **without** a prior distribution. That is, the Kaplan-Meier estimator.

B-RUP and MLE?

- Can we use the same procedure to calibrate the B-RUP?
- Unfortunately no, since we do not have **direct observations** of A, B and C. What we observe is the pair (X, Y).
- We can calibrate the B-RUP via the *Expectation-Maximization* (EM) algorithm.

Expectation-Maximization

The EM algorithm computes the expectation of the **complete** likelihood from the **incomplete** data in an iterative procedure.

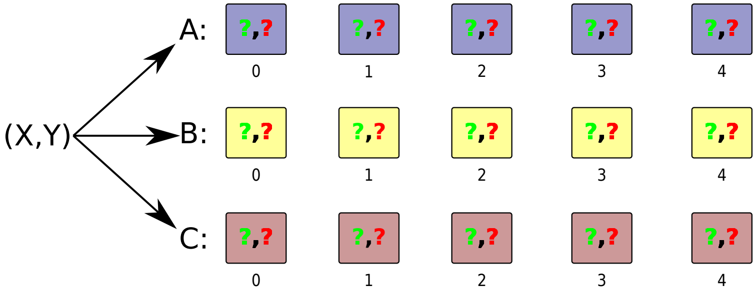

B-RUP and EM

- Since these variables are independent, we can compute the **likelihood** for each variable **separately**.
- The configuration at each iteration can be obtained **explicitly** from the optimal configuration of the **previous** iteration:

$$\frac{R_j^{A,k}}{G_j^{A,k} + R_j^{A,k}} = \frac{\sum_{i=1}^{n} \mathbb{P}[A = j | \theta^{[k-1]}, x_i, y_i]}{\sum_{i=1}^{n} \mathbb{P}[A \geq j | \theta^{[k-1]}, x_i, y_i]}, \tag{8}$$

What about the prior?

- We still would like to include **expert's knowledge**, as in the RUP.
- We can actually do that for A, B and C:

$$\mathbb{P}(A > a) = \prod_{j=0}^{a} \left[1 - \frac{\beta_j^A + R_j^A}{\beta_j^A + \omega_j^A + G_j^A + R_j^A} \right], \tag{9}$$

where the pairs $\{\beta_j^A, \omega_j^A\}$ define a prior for A.

Table of Contents

Canadian insurer data set

- 11,421 male-female couples of clients from a Canadian insurance company observed from 29th December 1988, until 31st December 1993 (see Frees et al. [1996]).
- Since most of the couples were alive by the end of the observation period, this data set contains mostly right-censored observations.
- Truncation also plays a big role because many couples joined the study at an old age.

Choosing a prior

- We start by defining the initial behaviour of our variables through the pairs $\{\beta_j, \omega_j\}$.
- In this case we choose: $G^A = \text{Poi}(35)$, $G^B = \text{Poi}(40)$ and $G^C = \text{Poi}(40)$.
- We choose a **low** strength of belief to check the "objective" behaviour of the B-RUP.

Marginal distributions

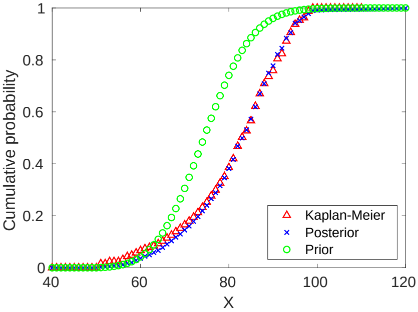

(a) Marginal of X.

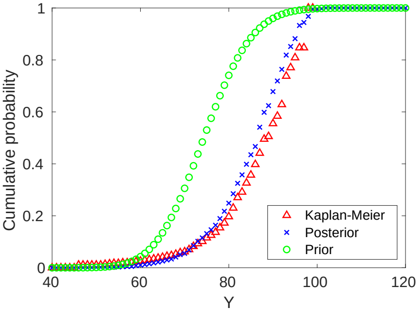

(b) Marginal of Y.

Joint distribution

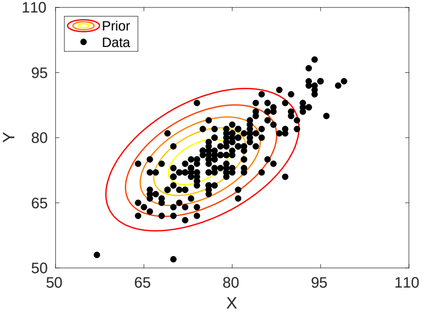

(a) Uncensored data.

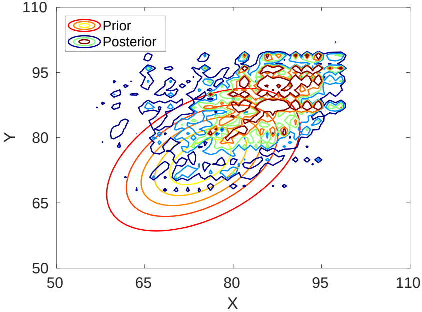

(b) Posterior distribution.

Annuity scenarios

- Now that we have the survival distribution it is time to **price** some annuities.

- We compute the **annuity ratio**:

$$AR = \frac{\text{Price assuming dependence}}{\text{Price assuming independence}} \tag{10}$$

- This will give us an idea of the **impact** of dependence.

Constant interest rate

We fix $r = 0.05$ and compute the AR for several initial ages.

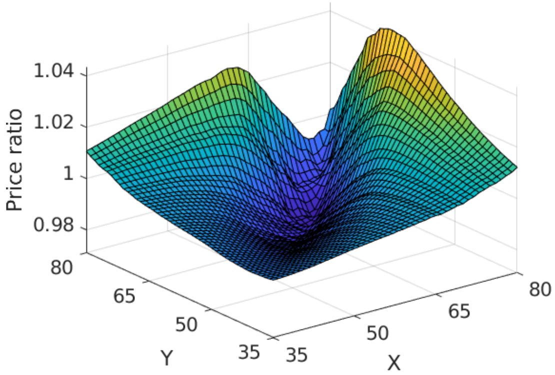

Figure: Interest rate $r = 0.05$.

Couples with same age

We set $x_0 = y_0$ and vary the interest rate.

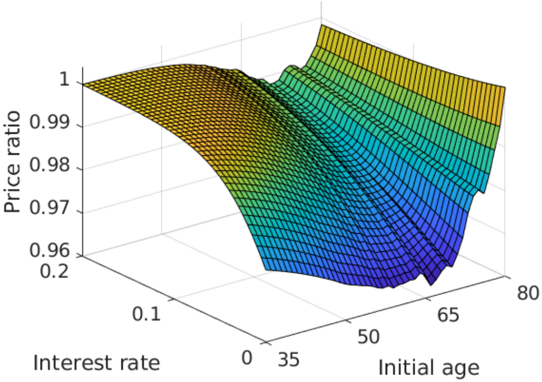

Figure: Couples with same age.

Table of Contents

Conclusions

- The B-RUP can be used in an objective way to model coupled lifetimes.
- It has the advantage of **allowing for expert's knowledge**, typical of Bayesian methods.
- Assuming independence in the lifetimes can lead to **both overpricing and underpricing**, depending on the age difference.
- Ignoring censoring and truncation **overestimates the dependence** and **underestimates life expectancy**.

References

Paolo Bulla, Pietro Muliere, and Steven G. Walker. Bayesian Nonparametric Estimation of a Bivariate Survival Function. *Statistica Sinica*, 17(3): 427–444, 2007.

Edward W. Frees, Jacques Carriere, and Emiliano Valdez. Annuity valuation with dependent mortality. *The Journal of Risk and Insurance*, 63(2): 229–261, 1996.

Stephen G. Walker and Pietro Muliere. Beta-Stacy Processes and a Generalization of the Pólya-Urn Scheme. *The Annals of Statistics*, 25(4): 1762–1780, 1997.

Thanks for your attention!