



Stochastic Mortality Modeling

ABCs of APCs and other
advanced mortality models

Mark Spong, FSA, CERA, MAAA

May 2025

1

Poll
Story
Acknowledgements

2

Regression
Age, Period, & Cohorts

3

Hierarchy of
Mortality Models

4

Using StMoMo in R
Calibrating parameters
Mortality improvement

5

Overlapping structures
Poll
Goodness-of-Fit

6

Stochastic Forecast
Sense-check
Poll

Introductory story



Thanks for the StMoMo package and video

Spong, Mark
 To andresmauriciovillegas@gmail.com; a.villegas@unsw.edu.au 9/18/2024

Internal Use

Hi Andres,

I'm a mortality researcher learning more about APC models and I found your 20 minute video about "An R Package for Stochastic Mortality Modelling" with StMoMo to be very informative. Although I haven't used it yet, I hope to and I wanted to express my appreciation for your contribution to this area.

Hi Mark,

Many thanks for the message. I am very glad hear that you have found my work on mortality modelling useful.

Best,

-Mark

SOA Research INSTITUTE

2024 STUDENT RESEARCH Case Study Challenge SuperLife Saving Lives January 19, 2024

Finalists

- **First Place:** University of New South Wales, Australia - StatSmart Solutions (winning a \$5000 USD grant)
- **Second Place:** University of New South Wales, Australia - Luminous Lake Lifeguards (winning a \$3000 USD grant)
- **Third Place:** Salisbury University, U.S. - Mew Consulting (winning a \$2000 USD grant)
- Ball State University, U.S. - Cardinal Consulting
- University of New South Wales, Australia - Apex Consulting
- University of New South Wales, Australia - Lumarian Consulting Group

Semifinalists

We are pleased to announce that a total of 17 teams from the following universities, listed alphabetically by university, have advanced to the semi-final round of the 2024 Challenge:

- Ball State University, U.S.
- Hunan University, China
- National Economics University, Vietnam
- Salisbury University, U.S.
- Universidad Nacional Mayor de San Marcos, Peru
- University of New South Wales (12 semifinalist teams), Australia

StMoMo: Stochastic Mortality Modelling

Implementation of the family of generalised age-period-cohort stochastic mortality models. This family of models including the Lee-Carter (1992) <doi:10.2307/2290201> and the Cairns-Blake-Dowd (2006) <doi:10.1111/j.1539-6978.2006.01111.x> analysing their goodness-of-fit and performing mortality projections and simulations.

Version: 0.4.1
 Depends: R (≥ 3.2.0), gnm (≥ 1.0), forecast (≥ 6.1)
 Imports: MASS, rootSolve (≥ 1.6.5.1), fanplot (≥ 3.4), reshape2 (≥ 1.4.1), fields (≥ 8.2), RColorBrewer
 Suggests: demography, knitr, xtable, MTS
 Published: 2018-04-13
 DOI: 10.32614/CRAN.package.StMoMo
 Author: Andres Villegas, Pietro Millosovich, Vladimir Kaishev
 Maintainer: Andres Villegas <andresmauriciovillegas@gmail.com>
 BugReports: http://github.com/amvillegas/StMoMo/issues
 License: GPL-2 | GPL-3 [expanded from: GPL (≥ 2)]
 URL: http://github.com/amvillegas/StMoMo
 NeedsCompilation: no
 Citation: StMoMo citation info
 Materials: README NEWS
 In views: ActuarialScience
 CRAN checks: StMoMo results

Cass Business School
 Doctor of Philosophy (PhD), Actuarial Science
 2010 - 2015
 Thesis - Mortality: Modelling, Socio-Economic Differences and Basis Risk

PhD Candidate
 Cass Business School
 Sep 2010 - Oct 2015 · 5 yrs 2 mos
 London, United Kingdom

My work colleague

Andrés Villegas

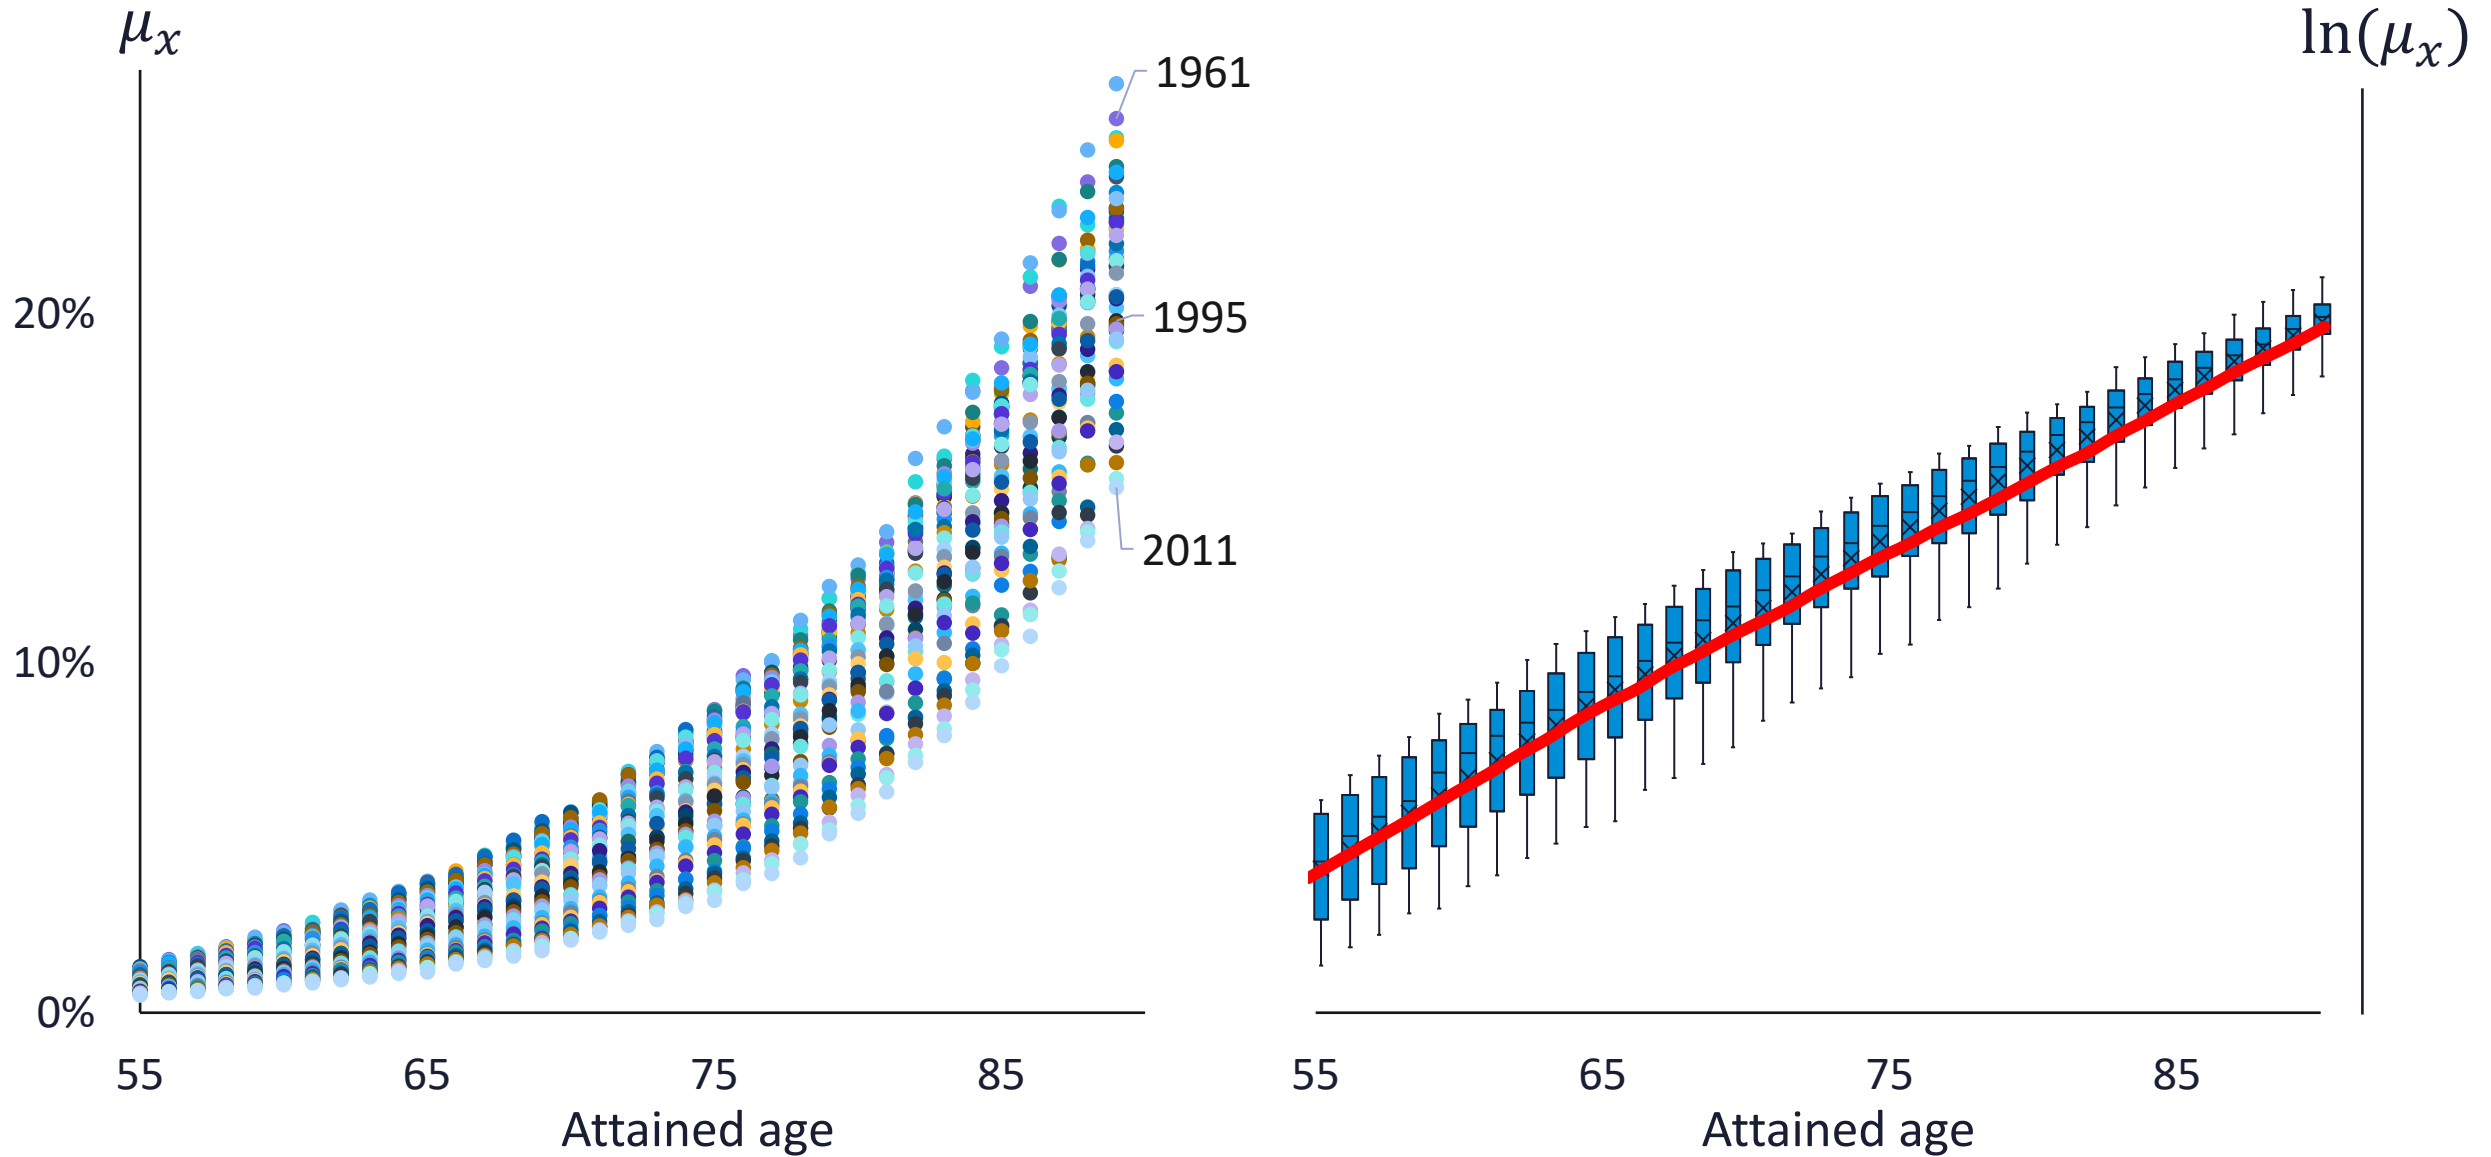
My simulated mortality dataset

Acknowledgements


Matt Smith, Pacific Life Re	Originally presented a similar presentation “The ABCs of APCs” on 6/15/2023 with the Canadian Institute of Actuaries
StMoMo	R package used to fit all models and source default dataset – Authors Andres Villegas, Pietro Millosovich , Vladimir Kaishev (2017)
Lee-Carter	“Modelling and Forecasting U.S Mortality” (1992)
Cairns-Blake-Dowd	“A Two Factor Model for Stochastic Mortality with Parameter Uncertainty: Theory and Calibration” (2006)
Renshaw-Haberman	“A Cohort Based Extension to the Lee Carter Model for Mortality Reduction Factors” (2006)
Plat	“On Stochastic Mortality Modeling” (2009)

Today’s content is based on my code which closely follows the vignette that accompanies the StMoMo R package

Mortality Modeling is hard but regression is easy



What do we mean by Age, Period, and Cohort?



Hierarchy of Mortality Models

Mortality as a function of sex and ...

Name	Link ¹	Age α_x	Period κ_t	Period:Age $\beta_x \cdot \kappa_t$	Period:Age ²	Cohort γ_{t-x}
Age only	LN/LOGIT	α_x	-	-	-	-
Age and time	LN/LOGIT	α_x	κ_t	-	-	-
Lee-Carter	LN	α_x	-	$\beta_x \cdot \kappa_t$	-	-
Age-Period-Cohort ("APC")	LN/LOGIT	α_x	κ_t	-	-	γ_{t-x}
Renshaw-Haberman ("RH")	LOGIT	α_x	-	$\beta_x^{(1)} \cdot \kappa_t$	-	γ_{t-x}
Cairns-Blake-Dowd ("CBD")	LOGIT	-	$\kappa_t^{(1)}$	$(x - \bar{x}) \cdot \kappa_t^{(2)}$	-	-
Plat	LOGIT	α_x	$\kappa_t^{(1)}$	$(x - \bar{x}) \cdot \kappa_t^{(2)}$	-	γ_{t-x}
M7	LN/LOGIT	-	$\kappa_t^{(1)}$	$(x - \bar{x}) \cdot \kappa_t^{(2)}$	$\left((x - \bar{x})^2 - \sigma_t^2 \right) \kappa_t^{(3)}$	γ_{t-x}

→ LN Deaths follow Poisson distribution $\mu_{x,t} = e^{\text{sum of predictors}}$
 → LOGIT Deaths follow Binomial distribution $\mu_{x,t} = \frac{e^{\text{sum of predictors}}}{1 + e^{\text{sum of predictors}}}$

Stochastic Mortality Modeling in R with StMoMo

```
> install.packages("StMoMo")

> library("StMoMo")

> EWMaleData
...# some data transform/truncation steps to produce EWMaleIniData

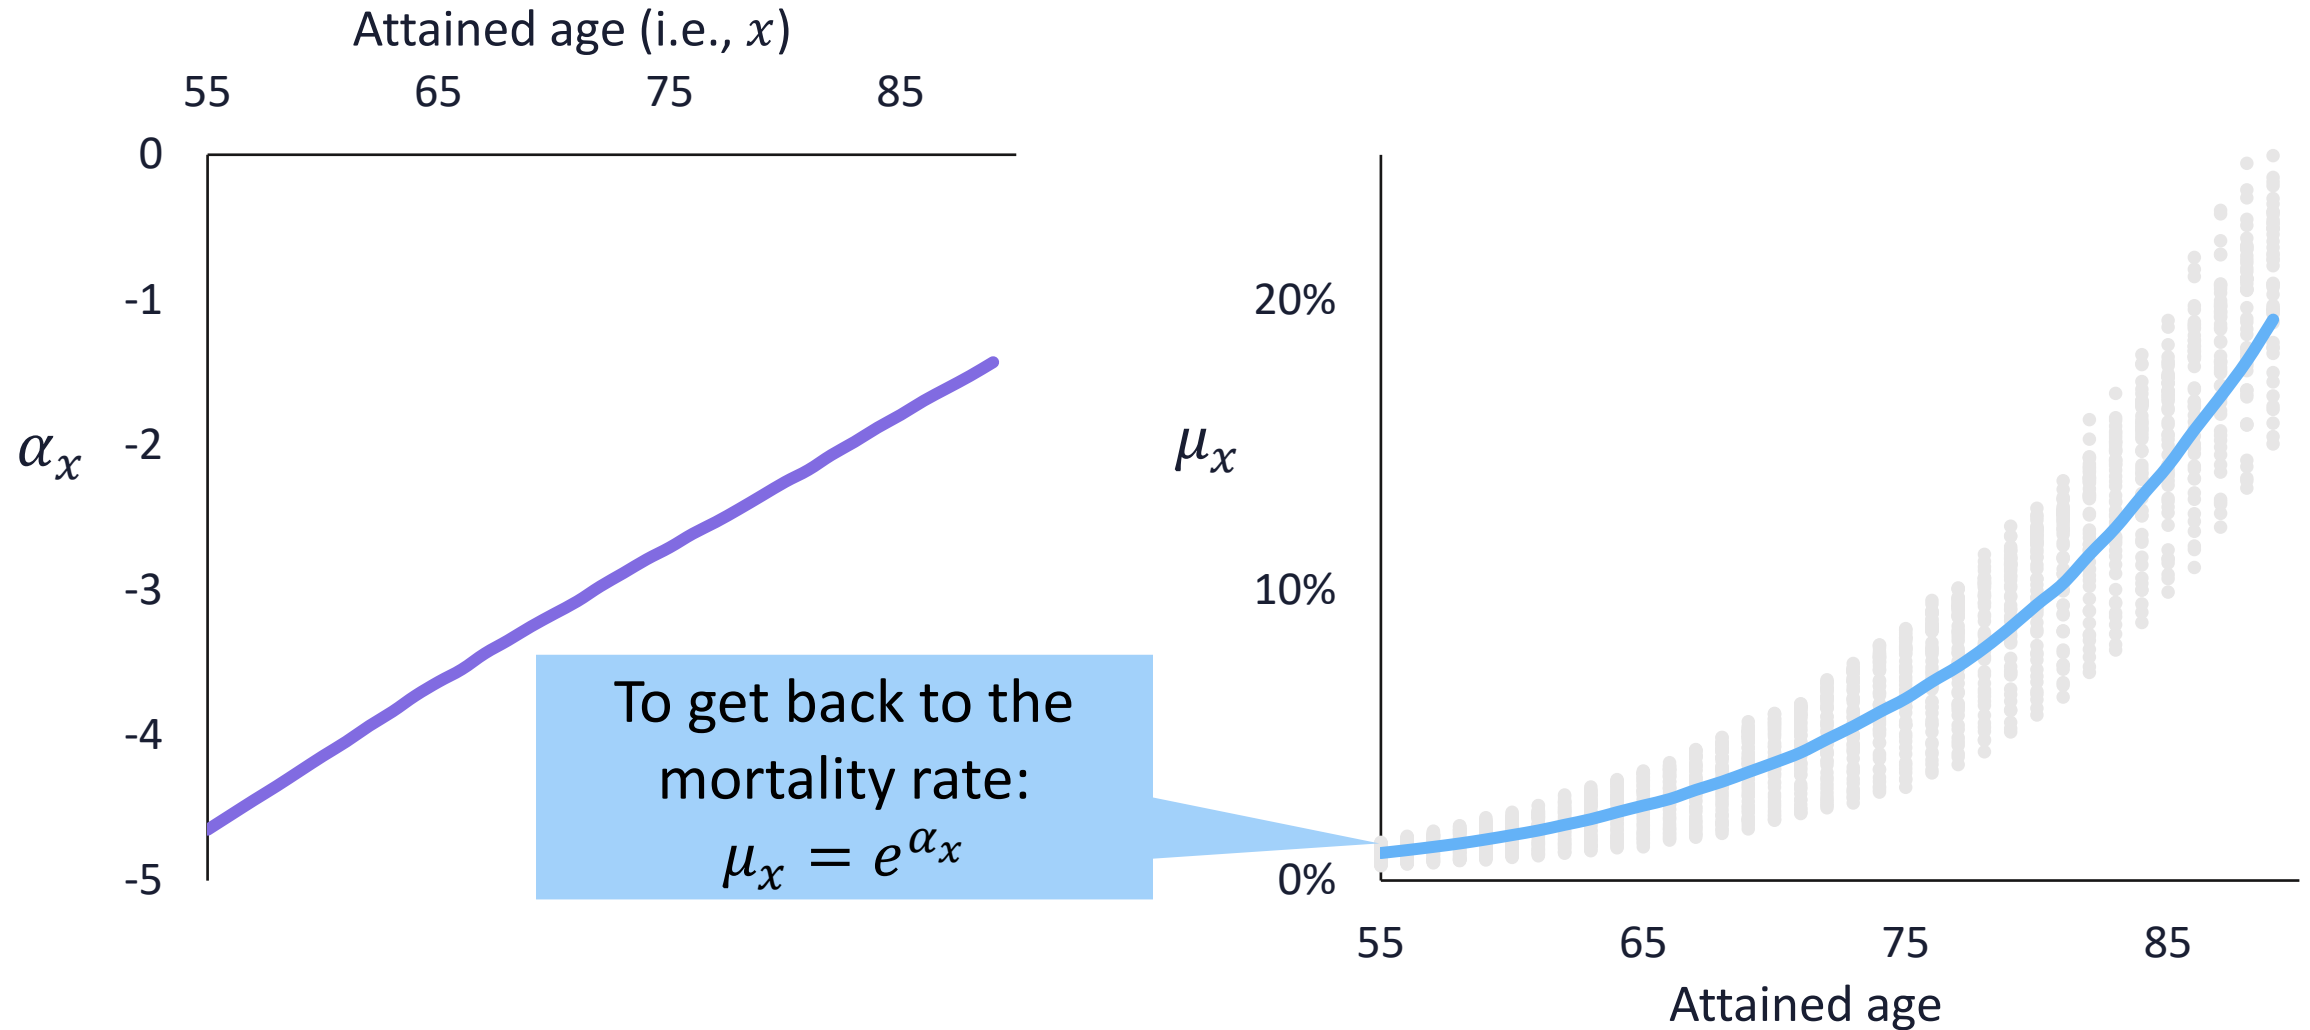
> LC_simple <- StMoMo(link = "log", staticAgeFun = TRUE,
                     periodAgeFun = NULL, cohortAgeFun = NULL)

> LC_simplefit <- fit(LC_simple, data = EWMaleIniData, ...)

> plot(LC_simplefit, nCol = 1)
```


Age only results

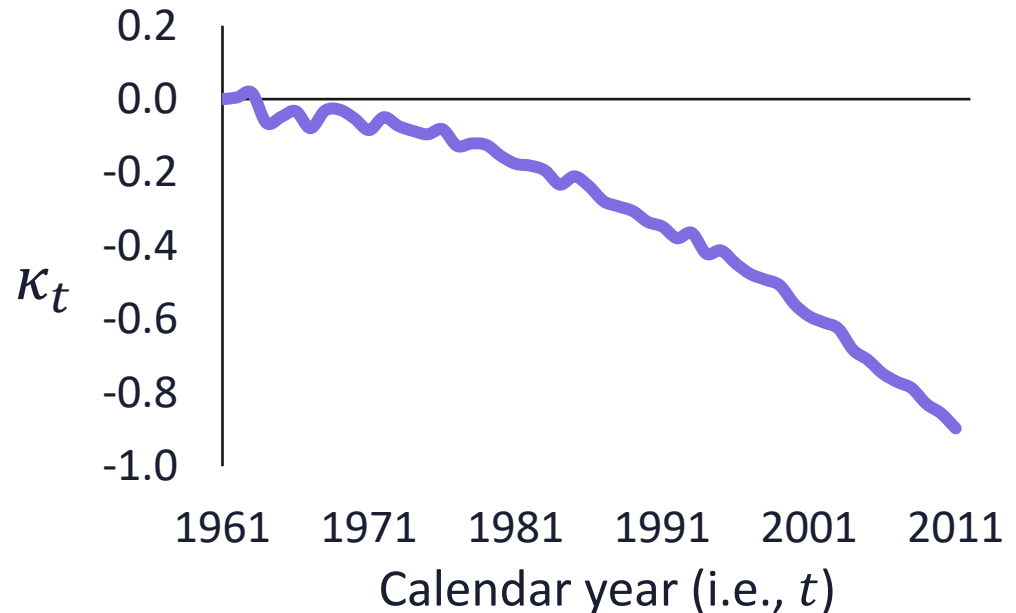

$$\ln(\mu_x) = \alpha_x$$



Age and time

$$\ln(\mu_x) = \alpha_x + \kappa_t$$

```
> LC_Agetime <- StMoMo(link = "log", staticAgeFun = TRUE,  
                        periodAgeFun = "1", cohortAgeFun = NULL)  
> LC_Agetimefit <- fit(LC_Agetime, data = EWMaleIniData, ...)  
> plot(LC_Agetimefit, nCol = 2, parametricbx = FALSE)
```



Age and time – Mortality Improvement Interlude

$$\ln(\mu_x) = \alpha_x + \kappa_t$$

Change over time ... that is mortality improvement!

$$\begin{aligned} \text{Mortality Improvement} &= 1 - \frac{\mu_{x,t}}{\mu_{x,t-1}} \\ &= 1 - \frac{e^{\alpha_x + \kappa_t}}{e^{\alpha_x + \kappa_{t-1}}} \\ &= 1 - e^{\cancel{\alpha_x + \kappa_t} - (\cancel{\alpha_x + \kappa_{t-1}})} \\ &= 1 - e^{\kappa_t - \kappa_{t-1}} \end{aligned}$$

To get back to the mortality rate:

$$\mu_{x,t} = e^{\alpha_x + \kappa_t}$$

Year	Age	α_x	κ_t	Mort.	MI	MI
1983	75	-2.4	-0.19	7.2%		
1984	75	-2.4	-0.23	6.9%	3.8%	3.8%
1985	75	-2.4	-0.21	7.1%	-2.3%	-2.3%

$$= 1 - \frac{\mu_{x,t}}{\mu_{x,t-1}}$$

$$= 1 - e^{\kappa_t - \kappa_{t-1}}$$

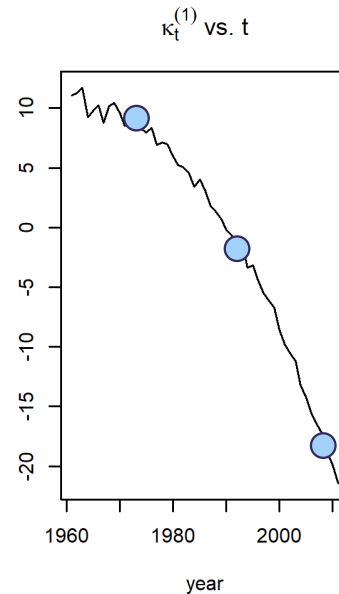


Lee Carter – Age and Period:Age

$$\ln(\mu_x) = \alpha_x + \beta_x \cdot \kappa_t$$

```
> LC <- lc(link = "log")
> LCfit <- fit(LC, data = EWMaleIniData, ...)
> plot(LCfit, nCol = 3)
```

To get back to the mortality rate:

$$\mu_{x,t} = e^{\alpha_x + \beta_x \cdot \kappa_t}$$



Year	Age	α_x	β_x	κ_t	Mort.
1970	60	-4.2	0.035	10	2.1%
1990	60	-4.2	0.035	0	1.5%
2010	60	-4.2	0.035	-20	0.8%
1970	70	-3.2	0.033	10	5.5%
1990	70	-3.2	0.033	0	4.0%
2010	70	-3.2	0.033	-20	2.1%
1970	80	-2.3	0.023	10	12.3%
1990	80	-2.3	0.023	0	9.8%
2010	80	-2.3	0.023	-20	6.2%

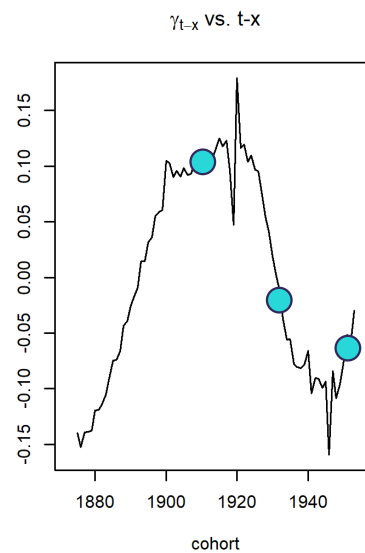


Age, Period and Cohort

$$\ln(\mu_x) = \alpha_x + \kappa_t + \gamma_{t-x}$$

```
> APC <- apc(link = "log")
> APCfit <- fit(APC, data = EWMaleIniData, ...)
> plot(APCfit, parametricbx = FALSE, nCol = 3)
```

To get back to the mortality rate:

$$\mu_{x,t} = e^{\alpha_x + \kappa_t + \gamma_{t-x}}$$



Year	Age	α_x	κ_t	γ_{t-x}	Mort.
1970	60	-4.2	0.3	0.11	2.1%
1990	60	-4.2	-0.1	0.02	1.4%
2010	60	-4.2	-0.5	-0.08	0.8%
1970	70	-3.3	0.3	0.10	5.6%
1990	70	-3.3	-0.1	0.18	4.3%
2010	70	-3.3	-0.5	-0.07	2.2%
1970	80	-2.4	0.3	-0.03	12.1%
1990	80	-2.4	-0.1	0.11	9.8%
2010	80	-2.4	-0.5	0.02	5.9%





RH

$$\text{logit}(\mu_x) = \alpha_x + \beta_x^{(1)} \cdot \kappa_t + \gamma_{t-x}$$

```
> RH <- rh(link = "logit", cohortAgeFun = "1")
> RHfit <- fit(RH, data = EWMaleIniData, ...)
> plot(RHfit, nCol = 3, parametricbx = FALSE)
```

To get back to the mortality rate:

$$\mu_{x,t} = \frac{e^{\alpha_x + \beta_x^{(1)} \cdot \kappa_t + \gamma_{t-x}}}{1 + e^{\alpha_x + \beta_x^{(1)} \cdot \kappa_t + \gamma_{t-x}}}$$



Year	Age	α_x	$\beta_x^{(1)}$	κ_t	γ_{t-x}	Mort.
1970	60	-4.2	0.02	9.3	0.2	2.1%
1990	60	-4.2	0.02	-3.1	0.0	1.4%
2010	60	-4.2	0.02	-13.3	-0.4	0.8%
1970	70	-3.3	0.03	9.3	0.2	5.6%
1990	70	-3.3	0.03	-3.1	0.3	4.3%
2010	70	-3.3	0.03	-13.3	-0.1	2.1%
1970	80	-2.3	0.04	9.3	0.0	12.0%
1990	80	-2.3	0.04	-3.1	0.2	9.8%
2010	80	-2.3	0.04	-13.3	0.0	6.0%

CBD

$$\text{logit}(\mu_x) = \kappa_t^{(1)} + (x - \bar{x}) \cdot \kappa_t^{(2)}$$


$\kappa_t^{(1)}$ = General level of mortality at time (t)

$\kappa_t^{(2)}$ = Rate of change at each time (t) in mortality at each age (x)


$$y = a + bx$$

```
> CBD <- cbd()  
> CBDfit <- fit(CBD, data = EWMaleIniData, ...)  
> plot(CBDfit, parametricbx = FALSE)
```

$\kappa_t^{(1)}$ vs. t



$\kappa_t^{(2)}$ vs. t



Year	Age	$\kappa_t^{(1)}$	$\kappa_t^{(2)}$	Mort.
1970	60	-2.7	0.09	2.2%
1990	60	-3.0	0.10	1.5%
2010	60	-3.6	0.11	0.8%
1970	70	-2.7	0.09	5.6%
1990	70	-3.0	0.10	4.1%
2010	70	-3.6	0.11	2.2%
1970	80	-2.7	0.09	13.9%
1990	80	-3.0	0.10	10.9%
2010	80	-3.6	0.11	6.5%

Hierarchy of Mortality Models

Mortality as a function of sex and ...

Name	Link ¹	Age α_x	Period κ_t	Period:Age $\beta_x \cdot \kappa_t$	Period:Age ²	Cohort γ_{t-x}
Age only	LN/LOGIT	α_x	-	-	-	-
Age and time	LN/LOGIT	α_x	κ_t	-	-	-
Lee-Carter	LN	α_x	-	$\beta_x \cdot \kappa_t$	-	-
Age-Period-Cohort ("APC")	LN/LOGIT	α_x	κ_t	-	-	γ_{t-x}
Renshaw-Haberman ("RH")	LOGIT	α_x	-	$\beta_x^{(1)} \cdot \kappa_t$	-	γ_{t-x}
Cairns-Blake-Dowd ("CBD")	LOGIT	-	$\kappa_t^{(1)}$	$(x - \bar{x}) \cdot \kappa_t^{(2)}$	-	-
Plat	LOGIT	α_x	$\kappa_t^{(1)}$	$(x - \bar{x}) \cdot \kappa_t^{(2)}$	-	γ_{t-x}
M7	LN/LOGIT	-	$\kappa_t^{(1)}$	$(x - \bar{x}) \cdot \kappa_t^{(2)}$	$\left((x - \bar{x})^2 - \sigma_t^2 \right) \kappa_t^{(3)}$	γ_{t-x}

Q: Which one is best?

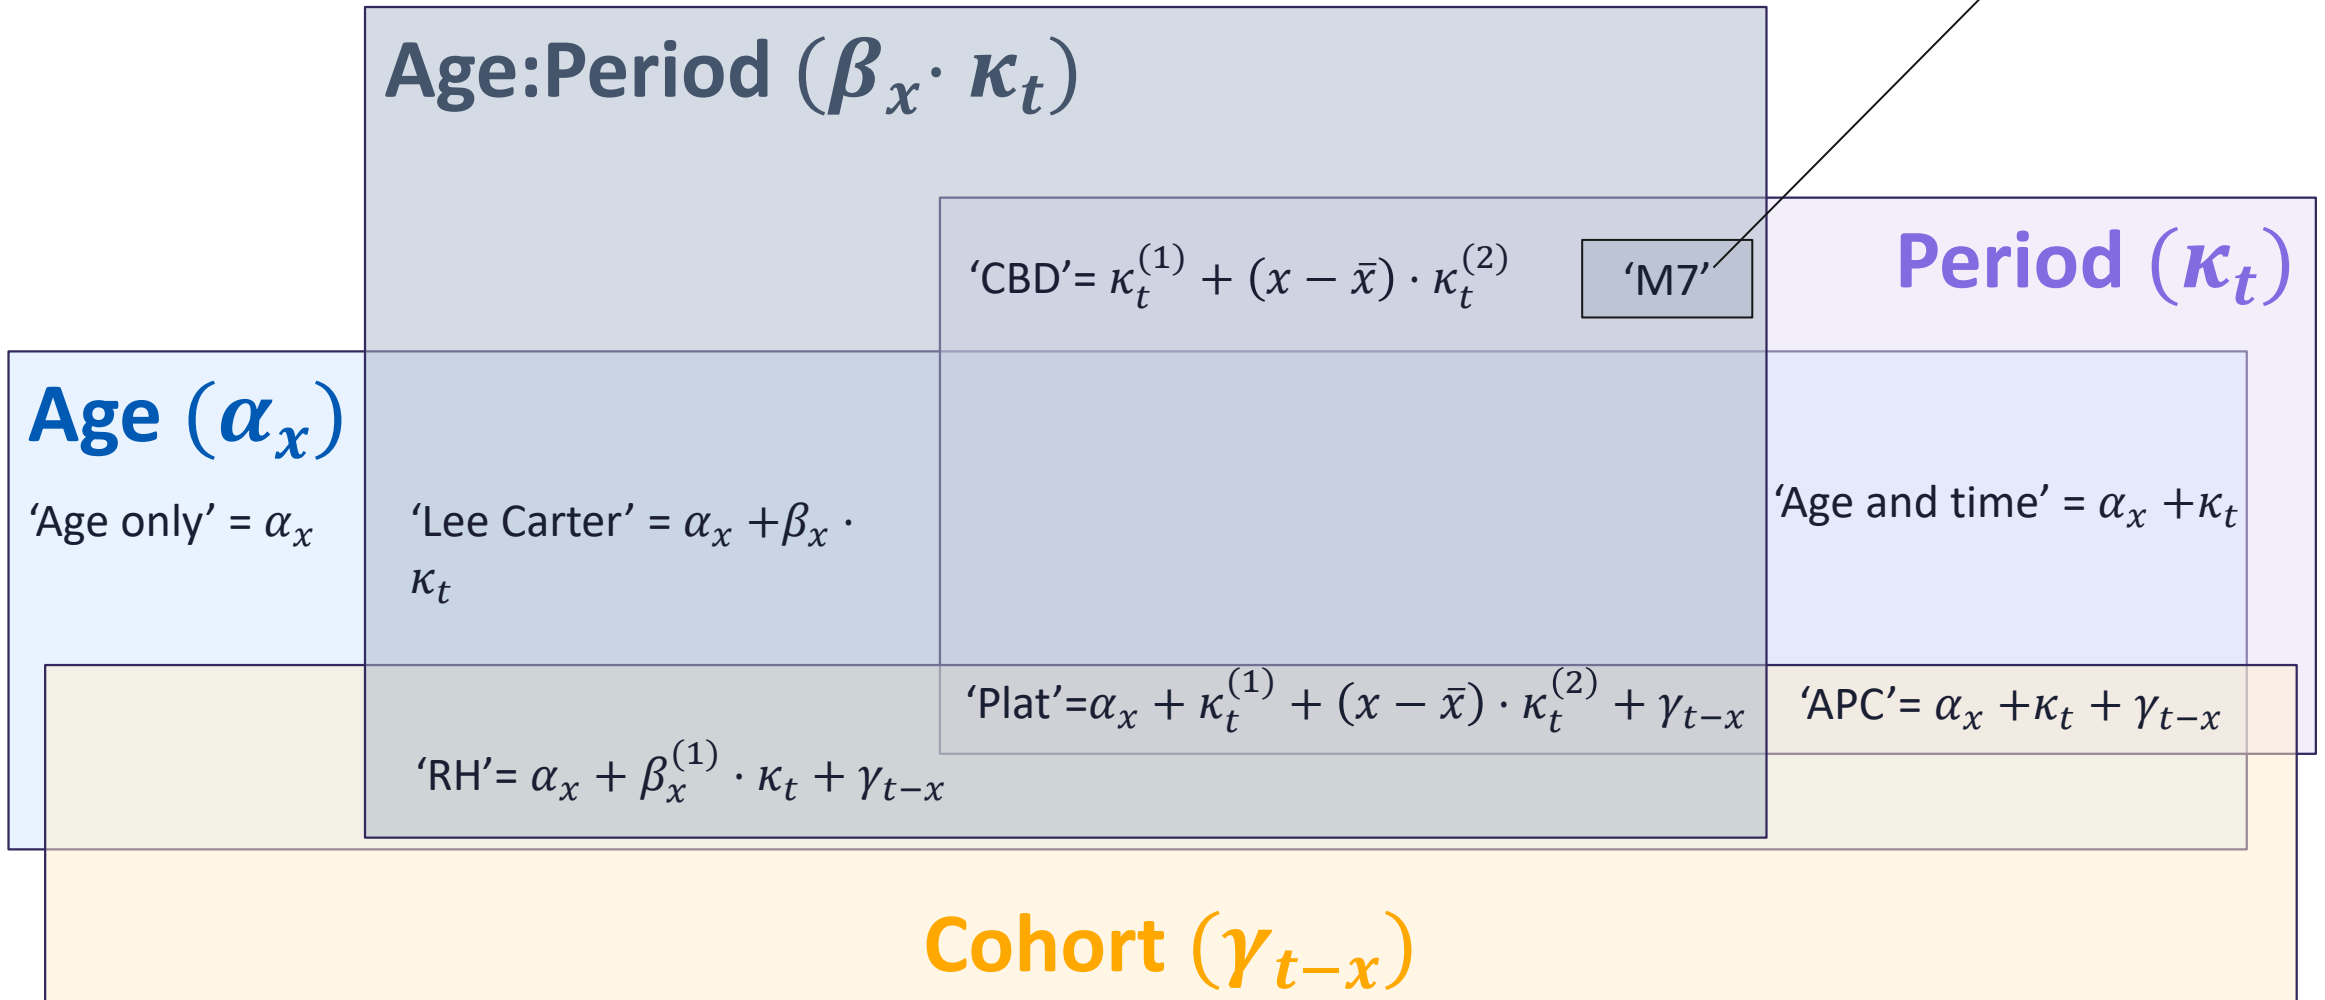
A: It depends, need to analyze goodness of fit

Q: After we choose the best one and calibrate parameters, how do we model stochastically?

A: Use an Autoregressive Integrated Moving Average models to extend κ_t and γ_{t-x}

Overlapping Model Structures

$$= \kappa_t^{(1)} + (x - \bar{x}) \cdot \kappa_t^{(2)} + \left((x - \bar{x})^2 - \sigma_t^2 \right) \kappa_t^{(3)}$$



Goodness of fit analysis – code


```
# First calculate the residuals
> LC_simpleres <- residuals(LC_simplefit)
> LC_Agetimeres <- residuals(LC_Agetimefit)
> LCres <- residuals(LCfit)
> APCres <- residuals(APCfit)
> RHres <- residuals(RHfit)
> CBDres <- residuals(CBDfit)

# Plot the residuals to see if there are any obvious patterns
> plot(LC_simpleres, type = "colourmap", reslim = c(-3.5, 3.5))
> plot(LC_Agetimeres, type = "colourmap", reslim = c(-3.5, 3.5))
...
```


```
# Plot the residuals in a scatterplot by age, period, and cohort
> plot(LC_simpleres, type = "scatter", reslim = c(-3.5, 3.5))
> plot(LC_Agetimeres, type = "scatter", reslim = c(-3.5, 3.5))
...
```


Goodness of fit analysis – heatmap of residuals


Age only




Age & time




Lee-Carter




APC



RH

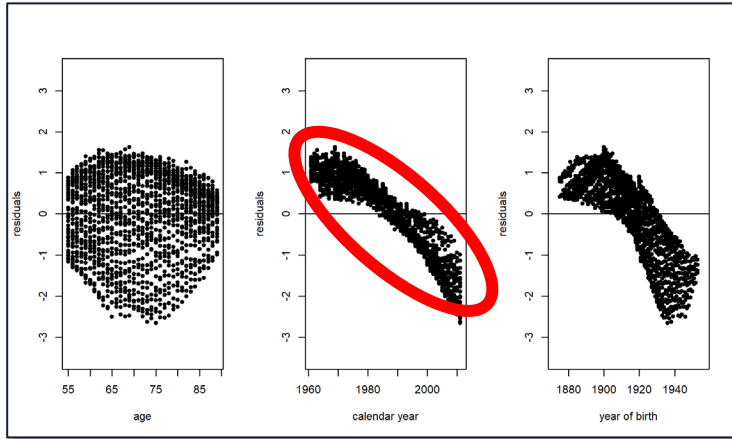


CBD




Goodness of fit analysis – residual scatterplots


Age only




Age & time



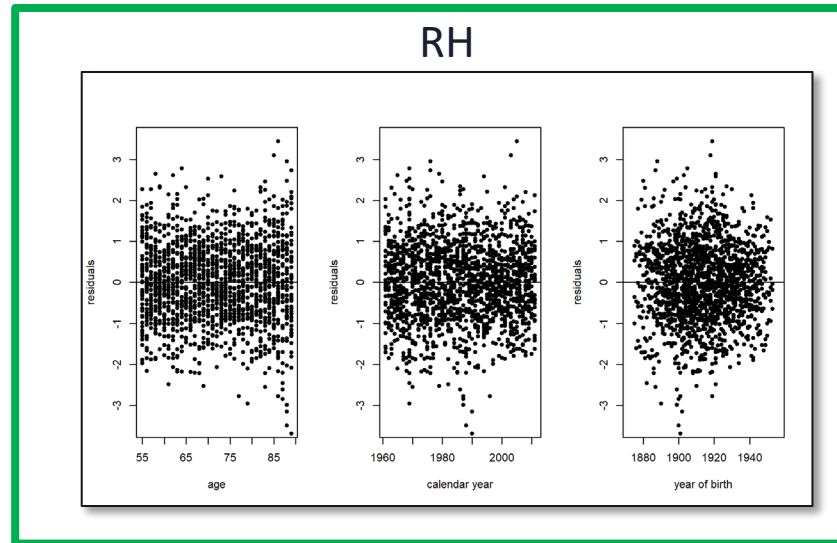
Lee-Carter




APC



RH



CBD




Stochastic forecast

```
# Make a stochastic forecast of the chosen model
```


```
> RHfor <- forecast(RHfit, h = 50, gc.order = c(1, 1, 0))
```

```
> plot(RHfor)
```


α_x vs. x




$\beta_x^{(1)}$ vs. x




$\kappa_t^{(1)}$ vs. t




γ_{t-x} vs. $t-x$



Projected Mortality based on forecasted mean



Close with a sense-check



Please fill out the survey!

Thank you

