CPB Netherlands Bureau for Economic Policy Analysis

CPB Background Document | October 2015

A model for joint pension risks

Sander Muns

A model for joint pension risks*

Sander Muns†

Monday 19th October, 2015

Abstract

Pension funds are faced with multiple headwinds in the coming decades: demographic risk from the continuing ageing process, wage risk from a lower productivity growth rate, and financial risk from low real interest rates. We model the joint distribution of these risks for the Netherlands by combining well-known models from the literature. The model's aim is to be a building block in evaluating a wide range of policy reforms on pension policy. Evaluations are possible at an age-specific, at a cohort-specific or at an aggregate level, and for a short run analysis as well as for a long run analysis of more than a few decades. Our results indicate that migration policy is a key determinant of the population size in the Netherlands after 2060. The short run dependence between the key risks is low thereby suggesting the presence of diversification benefits of the Dutch multi-pillar pension system.

1 Introduction

The ageing process, the low growth rates in productivity, and the low interest rate economy require a consideration of risks related to pension funds and public expenditures on social security. We therefore model demographic risk, productivity risk and financial risk with VAR(1) submodels, and find a low dependence between the annual shocks in the three different models. The joint distribution of the risk factors enables us to evaluate the long run effect of several policy measures on Dutch social security. Examples include the implementation of a new pension scheme and its transition process, or the optimal mix of pay-as-you-go (PAYG) and funded pension schemes.

Broer (2010) provides an extensive literature on the dependencies between the three key risk factors. More recently, Aro and Pennanen (2014) price mortality linked derivatives by relating mortality to several financial variables. Niu and Melenberg (2014) augment a mortality model with GDP per capita.

The impact of each risk factor on the sustainability of a pension scheme depends on the type of the scheme. Starting with demographic risk, most developed countries are confronted with declining

*I would like to thank Peter Broer, Marcel Lever, Thomas, Michielsen, Albert van der Horst and Joris de Wind for discussions and suggestions.

†email: s.muns@cpb.nl

population growth rates, and shrinking population sizes in the coming decades.[1] As a consequence, the number of retirees for each worker is expected to increase substantially during the next decades. For this fraction, Van Duin and Stoeldraijer (2014) project an increase from 29 percent in 2014 to 39 percent in 2040 if the retirement age is adjusted to life expectancy, and even 51 percent without this adjustment. Such demographic risk is of particular interest for a PAYG pension scheme because it relies on intergenerational risk sharing.

Productivity risk makes future wages uncertain. This is particularly important in a defined benefit (DB) PAYG scheme if the payments to retirees are corrected for wage inflation. Through nominal wage rigidity, an adverse productivity shock has most impact on the replacement rates of the current working population in this DB PAYG scheme. Suppose productivity growth is lower than wage growth. Thus, the current working generation faces a relatively high contribution rate to pay the wage indexed defined benefit of the old. In the long run, wages will adjust such that the current working generation is faced with lower wages, and ultimately lower pension benefits in this PAYG scheme. In a funded scheme, the current workers experience the direct effect of the productivity shock. Due to this shock, the indexation of the pension benefits will be lower at some point in time in a funded DB scheme. Therefore, the current retirees may also suffer from the productivity shock. In a funded defined contribution (DC) scheme, retirees may also suffer through their capital holdings provided capital returns are affected by the productivity shock.

Following the reasoning above, financial risk, e.g., risks in capital returns, has most impact on the replacement rates in a funded DC scheme as it affects the pension benefits of all generations. In a funded DB scheme, the response to financial shocks is smoothed by adjustments in pension contribution rates or pension benefits. Private capital holdings amplify the effects to shareholders.

Though the risks are in isolation relevant for pensions, the dependencies between the risks may amplify or mitigate certain effects. For instance, the joint effect of a low mortality, a wage growth that exceeds productivity growth, and a low capital return is very problematic in a DB scheme, particularly with wage indexing. Such risks are highly relevant for the Netherlands as the first pillar is a DB PAYG scheme, and most schemes in the second pillar are most similar to a funded DB scheme. Our results indicate a moderate instantaneous dependence between the risks.

To model the risks for the next decades, we outline an a-theoretical parsimonious VAR(1) model for the Netherlands. Demographic risk is described in Section 2, wage-productivity risk in Section 3, and financial risk in Section 4. The joint model is estimated in Section 5. Section 6 discusses potential improvements of the current model. Conclusions are in Section 7.

2 Demography

Demographic risk of pension policy consists of three components: risk in mortality, fertility, and migration. First, mortality affects the spell that an individual receives a pension (longevity risk). As such, a lower mortality is costly for the current working generation in a DB PAYG scheme. In

[1]See UN (2015) for world projections, EC (2014) for European projections, and Van Duin and Stoeldraijer (2014) for Dutch projections.

a funded DC scheme, the retirees pay the bill themselves by a reduction in their pension benefits.

Second, fertility is the key determinant for population growth, and hence the size of future generations. A small future working generation is problematic for the future funding in a PAYG scheme. In a funded DB scheme, a small working generation makes the pension system vulnerable to capital shocks. More specifically, a large negative capital shock during the working age of a small-sized generation requires a higher pension contribution to prevent cuts in pension benefits for the more sizeable generation of retirees.

Finally, migration and other changes such as administrative corrections explain the remaining changes in population size. This component mainly depends on foreign economic conditions and political unrest, which makes it highly volatile.

We model each of the three demographic submodels with age-specific time series. Each submodels contains one or two factors that are common to all age groups. The mortality submodel is discussed in Section 2.1, the fertility submodel is in Section 2.2, and the migration submodel ends this section in Section 2.3.

2.1 Mortality

The model A huge literature on mortality models has developed since the seminal paper of Lee and Carter (1992). Booth and Tickle (2008) contains an extensive literature review. Our mortality model is an extension of Lee and Carter (1992) in the spirit of Li and Lee (2005).[2] This extension assumes that the mortality of one group converges to the mortality of a larger group. In our case, Dutch mortality is assumed to converge towards the aggregate mortality of 13 European countries.[3] The data is from the Human Mortality Database if available, otherwise from Eurostat.[4] Data is available from 1960 to 2013 for the ages 0 to 99.

Let $x = 0, \ldots, 99$ and $t = 0, \ldots, T$ denote age and year, respectively. Define the mortality rate $m_{x,t} = D_{x,t}/E_{x,t}$ where $D_{x,t}$ is the number of deaths and the exposure $E_{x,t}$ is the population of age x during year t, computed following the protocol of the Human Mortality Database.[5] We model the mortality rate as

$$\ln\left(m_{x,t}^{\mathrm{EU}}\right) = a_x + b_x K_t + \varepsilon_{x,t}^{\mathrm{EU}} \tag{1}$$

$$\ln\left(m_{x,t}^{\mathrm{NL}}\right) = a_x + b_x K_t + \beta_x k_t + \varepsilon_{x,t}^{\mathrm{NL}} \tag{2}$$

where $\varepsilon_t^{\mathrm{EU}} \sim \left(0, \Sigma^{\mathrm{EU}}\right)$ and $\varepsilon_t^{\mathrm{NL}} \sim \left(0, \Sigma^{\mathrm{NL}}\right)$ are mutually independent disturbances with standard deviation σ_x^{EU} and σ_x^{NL}, respectively. The superscripts EU and NL refer to the whole group and the Netherlands, respectively.[6]

[2] Likewise, AG (2014) estimates the Li and Lee (2005) model on European data.

[3] The countries are Austria, Belgium, Germany, Denmark, Finland, France, Ireland, Iceland, Luxembourg, Netherlands, Norway, Sweden, and Switzerland.

[4] http://www.mortality.org and http://ec.europa.eu/eurostat/data/database.

[5] http://www.mortality.org/Public/Docs/MethodsProtocol.pdf

[6] In our Dutch sample, we did not find the cohort-specific mortality effects encountered in Renshaw and Haberman (2006) for mortality data of England and Wales. In addition, the dependence between cohort-and-period effects is difficult to simulate given incompletely observed cohorts.

Li and Lee (2005) include an additional country-specific intercept term in (2). However, this would represent a permanent age-specific difference between log mortality in the Netherlands and log mortality in Europe. This additional intercept term is missing here as it would contrast with our model assumption of convergence in mortality rates.[7]

Although it is straightforward to model gender-specific mortality, our unisex model aggregates mortality data for males and females. The aggregation makes the common factors K and k less sensitive to gender-specific characteristics, such as changes in smoking patterns. In addition, gender-specific models require an assessment of dependencies between the two sexes. Indeed, we do not aim to draw gender-specific policy results. Nevertheless, it is straightforward to extend our unisex model to gender-specific populations. Aggregating two linear gender-specific mortality factors gives again a linear mortality process provided the gender proportions in the total population are sufficiently stable over different ages.

Estimation Following Li and Lee (2005), we employ six steps to estimate (1) and (2):

(i) The vector a contains the age-specific means of log EU mortality. Subtract a_x from the corresponding time series $\ln\left(m_{x,t}^{\mathrm{EU}}\right)$.

(ii) The vector b and the time series K are from the first principal component of the time series obtained in (i).[8]

(iii) The vector b and the time series K are normalized such that $\sum_x b_x = 1$ (set $K := \tilde{b}K$ and $b := b/\tilde{b}$ where $\tilde{b}$ is the sum over the elements of the initial b).

(iv) Each K_t is adjusted such that it matches the life expectancy at birth in year t.[9]

(v) The covariance matrix of the residuals is Σ^{EU}.

(vi) Apply (ii)–(iv) to the residuals of Dutch mortality in (2): $\ln\left(m_{x,t}^{\mathrm{NL}}\right) - a_x - b_x K_t$. This gives the vector β, the time series k, and the covariance matrix Σ^{NL}.

Fixing a in step (i) excludes a translation of K in (1), the scaling of b in (iii) excludes a scaling of K. This identification scheme does not affect the goodness of fit nor the mortality predictions because each product $a_x + b_x K_t$ remains the same in (1). The first principal component K explains 97.2% of the variation in $\ln\left(m_{x,t}^{\mathrm{EU}}\right)$. As a consequence, the projected life expectancy in step (ii) is still close to the true life expectancy. Indeed, skipping step (iv) has no significant effect on the results. We choose to follow Li and Lee (2005) by performing this step. The monotonicity of life expectancy in K_t and k_t ensures that this step is fast.

[7]We also tested the model with a country-specific intercept in (2). Since Dutch mortality data is close to the European average, a country-specific intercept has no significant impact on the projections of life expectancy and population size.

[8]A singular value decomposition produces identical b and K if the mortality data is demeaned.

[9]We loosely refer to the period life expectancy as life expectancy. The period life expectancy in a certain period is the life expectancy of an individual facing the mortality rates in that period.

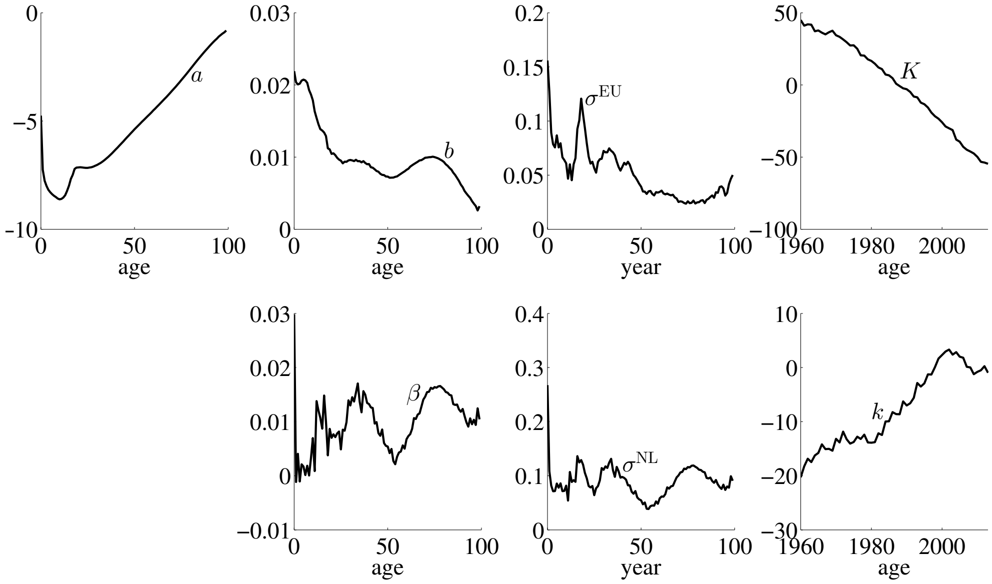

Figure 1: Estimation of (1) and (2). The intercepts a (top left), slopes b and β (second column), the standard deviation σ^{EU} and σ^{NL} of the residuals (third column), and common time series K and k (fourth column).

Figure 1 shows the estimation results. The age pattern in a shows the familiar mortality pattern. Mortality is increasing with age after a few years of living, and relatively high for individuals around twenty years old. The coefficients of b suggest that children benefited most from the decreasing pattern in K in the top right plot. Though (log) mortality is high at old ages (see a), the standard deviation of the residuals σ^{EU} is low for such ages (see third top plot). Therefore, the common factor K models old-age mortality well. This is important as most persons die at such ages.

The coefficients of β in Figure 1 suggest that the common factor k does not model the Dutch mortality deviation well for young children and persons around 50 years old. On the other hand, the standard deviation σ^{NL} is relatively low for individuals of 50 years old, thereby indicating that the common factor K models mortality quite well for this age group. The bottom right plot suggests that it is reasonable to assume that k converges to zero,

Figure 2 depicts the autocorrelation function (ACF) and the partial autocorrelation function (PACF) of K and k. Both ACFs show an exponential decay (left), while the PACFs have a cut-off after lag one (right). The ACF and PACF of the differenced series have a cut-off immediately after lag zero (not shown), which corresponds to white noise. Accordingly, the Box-Jenkins procedure suggests an AR(1) model for both K and k. The time series dynamics are

$$K_t = c + K_{t-1} + \eta_t^{\mathrm{EU}} \tag{3}$$

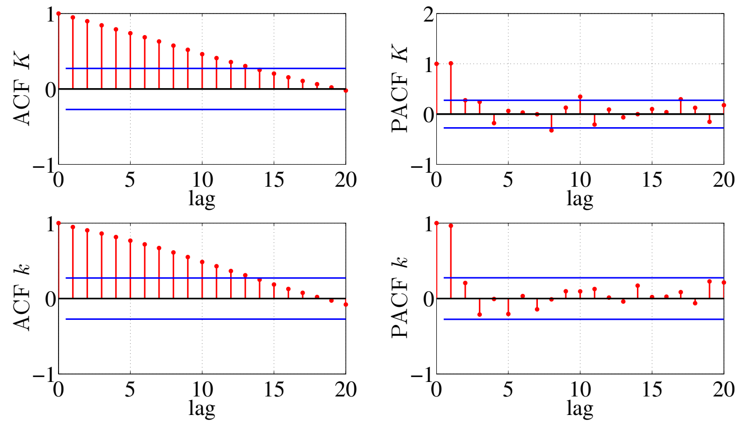

Figure 2: Autocorrelation function (left) and partial autocorrelation function (right) of K (top) and k (bottom).

$$k_t = \rho k_{t-1} + \eta_t^{\mathrm{NL}} \tag{4}$$

where $\eta_t^{\mathrm{EU}} \sim N(0, \sigma_K^2)$ and $\eta_t^{\mathrm{NL}} \sim N(0, \sigma_k^2)$ are mutually independent Gaussian disturbances. In line with the standard literature, the common factor K follows a random walk in equation (3). In contrast to Li and Lee (2005), equation (4) has no drift term. Omitting this term gives that $m_{x,t}^{\mathrm{NL}}$ converges towards $m_{x,t}^{\mathrm{EU}}$ since we excluded an age-specific intercept in (2).[10]

We estimate the parameters c, ρ, σ_K and σ_k by applying OLS to (3) and (4). Our estimation procedure ensures $\rho < 1$ (stationarity), and $\sigma_K, \sigma_k > 0$. Appendix A describes the implementation of the restrictions in more detail. We take parameter uncertainty into account by sampling from the asymptotic parameter distributions depicted in Figure 3. A negative c is consistent with declining mortality rates. The standard deviation σ_K suggests that K may increase incidentally, in line with the five years in our sample where K increases. The autocorrelation coefficient ρ of k is close to 0.9, thereby indicating a high persistence of Dutch shocks to mortality. The standard deviation σ_k tends to be lower than σ_K, consistent with the view that the common factor is the most important driver of variation in mortality rates.

High age mortality Mortality data exhibits a higher variance for high ages as mortality is more scarce at such ages. Similar to Antonio (2012), AG (2014) and Ševčíková et al. (2015), we use the regression method of Kannisto (1992) for high ages. This method linearly extrapolates the (simulated) logit of the mortality rate $m_{x,t}$:

$$\mu_{x,t} = \ln\left(\frac{m_{x,t}}{1 - m_{x,t}}\right) \tag{5}$$

[10] As a check, we added a drift term to (4), which did not change the results.

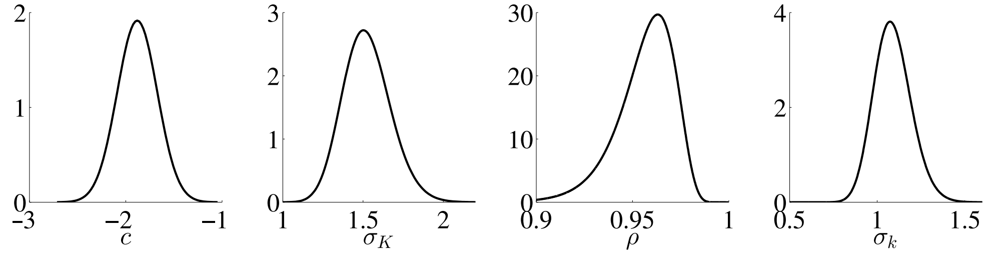

Figure 3: Asymptotic marginal parameter distribution of c, σ_K, ρ, and σ_k in (3) and (4).

In a cross-sectional regression, the logit rates $\mu_{x,t}$ are regressed on the corresponding ages $x = 80, \ldots, 95$:

$$\mu_{x,t} = \tilde{a}_t x + \tilde{b}_t + \xi \qquad\qquad \xi \sim N(0, \tilde{\sigma}_t^2) \tag{6}$$

In each simulation, we estimate for each time period t the parameters $\tilde{a}_t$, $\tilde{b}_t$, and $\tilde{\sigma}_t$ by an OLS regression. Our dataset satisfies $m_{x,t} < 1$ at these ages. The parameter estimates imply the projection

$$\hat{\mu}_{x,t} = \hat{a}_t x + \hat{b}_t + \hat{\xi}_{x,t} \qquad\qquad x > 95$$

with each $\hat{\xi}_{x,t}$ an independent draw from $N(0, \hat{\sigma}_t^2)$. The simulated mortality rate $\hat{m}_{x,t}$ follows by inverting the logit transformation (5) of $\hat{\mu}_{x,t}$:

$$\hat{m}_{x,t} = \frac{1}{1 + \exp\left(-\hat{\mu}_{x,t}\right)}$$

The following limitation applies to the method of Kannisto (1992). The mortality rate $m_{x,t}$ declines faster over time for the lower ages in $\{80, \ldots, 95\}$. The mortality projection extrapolates this divergent pattern to $\hat{\mu}_{x,t}$ for more distant simulation horizons t. When the time horizon t increases, the slope $\hat{a}_t$ increases in the cross-sectional regression (6), while the intercept $\hat{b}_t$ decreases. In most of the simulations, the increase in $\hat{a}_t$ dominates, which leads to a mortality rate that *increases* over time for very high ages. A typical example is in Figure 4 where the projected logit mortality rate $\hat{\mu}_{x,t}$ increases over time for ages $x > 110$. This limitation is solely a consequence of the extrapolation method of Kannisto (1992), regardless of the underlying mortality model at lower ages. Though the impact is relatively small for our current purposes, we stress that simulations over a horizon of several hundreds of years could be severely affected by this limitation. For instance, mortality rates may become close to one for ages above 110, while they converge to zero for lower ages.

Simulation The simulated mortality rates in a certain period generate a distribution of the period life expectancy at birth. This is the life expectancy a new born faces under the mortality rates of

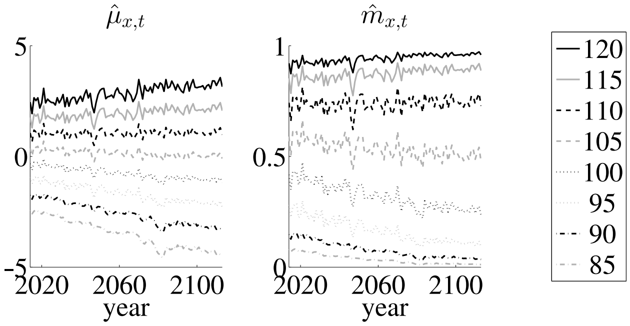

Figure 4: Example simulation for logit mortality rate $\hat{\mu}_{x,t}$ (left) and mortality rate $\hat{m}_{x,t}$ (right) for ages $x = 85, 90, 95, 100, 115, 120$. The Kannisto (1992) procedure is applied to ages above 95.

that specific period. It differs from the cohort life expectancy which aims to anticipate future developments in mortality. Though more relevant for an individual's life expectancy, the cohort life expectancy depends on highly uncertain probabilities in the more distant future. This paper focuses on the more commonly used period life expectancy to show simulations of life expectancy. That being said, policy implementations should be based on the more uncertain cohort life expectancy if necessary.

An important caveat applies to the classical Lee-Carter setup in (1) and (3). It leads to the counterintuitive result that the confidence intervals of the mortality rates decrease exponentially when the horizon increases. To see this, the distribution of the European mortality rate is for given model parameters b_x and K_t,

$$\begin{aligned} \ln\left(m_{x,t}^{\mathrm{EU}}\right) - \ln\left(m_{x,t-1}^{\mathrm{EU}}\right) &= b_x\left(K_t - K_{t-1}\right) + \varepsilon_{x,t}^{\mathrm{EU}} - \varepsilon_{x,t-1}^{\mathrm{EU}} \\ \frac{m_{x,t}^{\mathrm{EU}}}{m_{x,0}^{\mathrm{EU}}} &= \exp\left(b_x\left(K_t - K_0\right) + \varepsilon_{x,t}^{\mathrm{EU}} - \varepsilon_{x,0}^{\mathrm{EU}}\right) \end{aligned}$$

For large t,

$$t\sigma_K^2 = \mathrm{Var}\left(K_t - K_0\right) \gg \mathrm{Var}\left(\varepsilon_{x,t}^{\mathrm{EU}} - \varepsilon_{x,0}^{\mathrm{EU}}\right) = 2\left(\sigma_x^{\mathrm{EU}}\right)^2$$

such that

$$m_{x,t}^{\mathrm{EU}} \to m_{x,0}^{\mathrm{EU}} \exp\left(b_x c t + b_x \sigma_K \sqrt{t} Z_t\right) \tag{7}$$

That is, the mortality rate follows a lognormal distribution with variance

$$\begin{aligned} \mathrm{Var}\left(m_{x,t}^{\mathrm{EU}}\right) &\to \left[1 - \exp\left(-b_x^2\sigma_K^2 t\right)\right]\exp\left(2t\left[b_x c + b_x^2\sigma_K^2\right]\right) \\ &\to \exp\left(2b_x t\left[c + b_x\sigma_K^2\right]\right) \end{aligned}$$

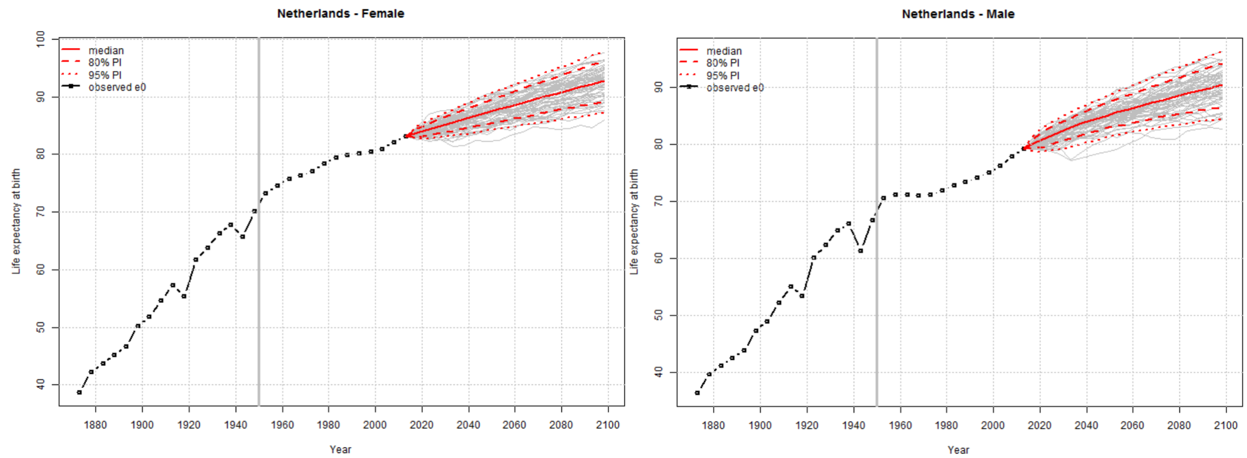

Figure 5: Projections of Dutch life expectancy at birth in UN (2015). Left: females, right: males. Source: http://esa.un.org/unpd/wpp/Graphs/

The inequality $c < -b_x \sigma_K^2$ holds for the vast majority of samples of the parameters in Figure 1 and 3. It then follows from $b_x > 0$ (Figure 1) that for high t the confidence interval of $m_{x,t}^{\mathrm{EU}}$ *decreases* in t. Accordingly, the confidence interval of the period life expectancy shrinks when the (high) simulation horizon *increases*, while intuition and the projection in UN (2015) suggest the opposite (see Figure 5).

To resolve this issue, the variance of η_t^{EU} in (3) increases if mortality rates are low, i.e., K is low. More specifically, we set

$$\sigma_K(K_t) = \sigma_K \max\left((K_0 - K_t)\,\sigma_0, 1\right) \tag{8}$$

The constant σ_0 is calibrated on the width of the 95% confidence interval of life expectancy for the year 2100 in UN (2015). Figure 5 indicates that this width is about 10 years. In Van Duin and Stoeldraijer (2014), the width of this 95% confidence interval is already 10 years in 2056.[11]

For several values of σ_0, we draw 5,000 parameter values together with a time series for the disturbances in (1) and (4). It turns out that $\sigma_0 = 0.04$ in (8) results in the desired width of 10 years in 2100, see Figure 6. On average, the adjustment $\sigma_K(K_t) > \sigma_K$ in (8) starts to work after 13 years since $\sigma_K(K_t) = \sigma_K$ if $t < (c\sigma_0)^{-1} = 13.2$ years.

Not only the width is similar to the UN model, the median Dutch unisex life expectancy in Figure 6 is also comparable to the median projection of the gender-specific distributions in Figure 5. More specifically, the median of our projection is 88 years in 2060 compared to 87.6 in the UN

[11] See Statistics Netherlands' StatLine for exact numbers, and Carolina and Van Duin (2010) for a motivation using historical gender-specific projections.

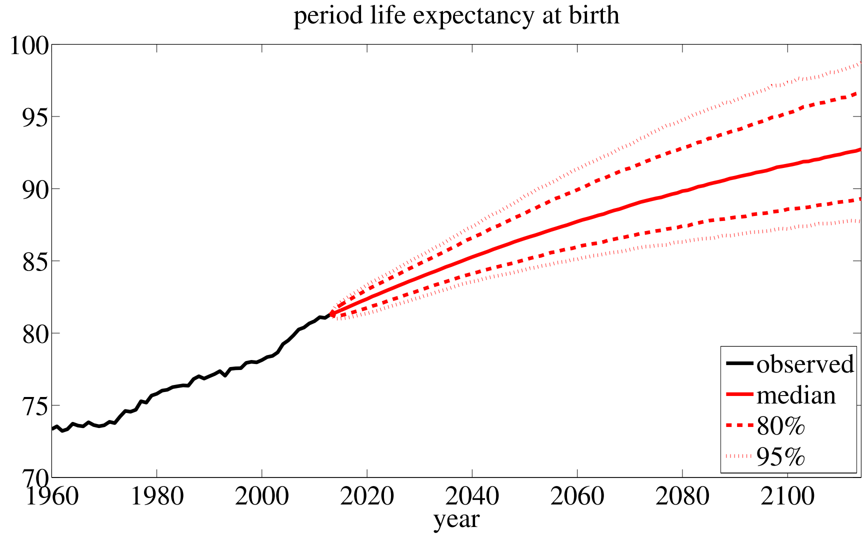

Figure 6: Projected Dutch unisex life expectancy at birth with 5,000 simulations with the model (1)–(4). Each simulation has a different sample from the parameter distribution.

model.[12] Our projection is also close to the 87 years in EC (2014, Table III.19.1), and 88.5 years in Van Duin and Stoeldraijer (2014).[13]

Instead of projecting age-specific probabilities of death, the model in UN (2015) makes a Bayesian projection of life expectancy for multiple countries (see Gerland et al. (2014) and the references therein). A potentially fruitful research avenue is to translate the Bayesian scenarios of life expectancy into scenarios of age-specific mortality. The UN is currently evaluating extensions along such lines to model mortality and fertility, see Ševčíková et al. (2015). Solely modelling life expectancy requires a substantially lower number of parameters, and is therefore more suitable for a Bayesian estimation. Nonetheless, four arguments motivate us to stick to the standard approach where the simulated age-specific mortality implies the simulated life expectancy at birth.

First, it is more intuitive that life expectancy is the result of the more fundamental age-specific probabilities of death than the other way around. Second, our focus on pension policy implies a focus on age-specific mortality rates. Third, we experimented with a simple time series specification to model the period life expectancy directly. However, this procedure cannot capture the slowdown in gains observed in Figure 5 and 6 at long horizons. For developed countries, the more sophisticated UN model implicitly imposes convergence of life expectancy towards a low rate of growth, thereby forcing a slowdown in gains of life expectancy at a high life expectancy. Fourth, we discussed that the extrapolation method of Kannisto (1992) cannot deal with improvements in mortality at high

[12] Average of the 2055-60 and 2060-65 in the Excel file of the UN Population Prospects, the 2015 Revision.

[13] See Statistics Netherlands' StatLine for exact numbers.

ages. This becomes particularly problematic under a direct simulation of life expectancy since the life expectancy increases faster. Hence, the age-specific mortality distribution would converge to a degenerate distribution at a relatively short horizon.

Next, we compare our results with the results in AG (2014). Like our paper, AG (2014) uses the method in Li and Lee (2005) to project Dutch mortality. There are some differences with our setup. First, along with the 13 countries in our dataset, that paper includes combined mortality data from England and Wales. Second, they consider gender-specific mortality models. Though the setup of both gender-specific models is identical, the processes are assumed to be mutually independent. Third, their underlying model is a Poisson process such that they model mortality intensities rather than mortality rates. Fourth, they estimate an unobserved mortality process by using a Kalman filter. Our model estimates a time series model for the time series of the first singular value (corrected for differences with the observed period life expectancy). Fifth, we take parameter uncertainty into account by sampling from their asymptotic distributions. Sixth, their data sample starts in 1970 and runs to 2013. For the latest years of some countries, data availability forces AG (2014) to use mortality projections. For each country in our sample, the data starts in 1960 and ends in 2013. Finally, AG (2014) does not anticipate the counterintuitive result on the confidence intervals that we resolve with (8).

Figure 7 shows our mortality projections of a 65 years old and the projection for females in AG (2014). The model in AG (2014) has some so-called jump-off risk from the Kalman filter. This initial mismodelling causes a too low prediction for the first simulation year. Two opposing forces explain the remarkably stable width of the confidence intervals in AG (2014). While forecasting at longer horizons should produce a wider confidence interval for the logarithm of the mortality rates, this effect is counterbalanced by inverting the log transformation of $\ln(m_{x,t})$ at lower $m_{x,t}$ in (7). In contrast, our model has additional uncertainty from (i) the parameter uncertainty, and (ii) the calibrated $\sigma_K(K_t)$ in (8) when mortality rates decline. The additional channels explain why the confidence bands widen over time in Figure 6 and the left panel in Figure 7. After 50 years of simulation, the median of the projections in Figure 7 are quite similar, particularly when taking into account the lower mortality rate of females and that our model employs unisex mortality data. In 2063, the 95% confidence intervals of mortality of a 65 years old have a width of 0.3% (left panel), and 0.2% (right panel).

2.2 Fertility

The model Figure 8 characterizes fertility rates from several perspectives. The period total fertility rate (TFR) is the expected number of births for a female experiencing at each age the birth rate in a certain period. As such, it is not affected by composition effects of the population. The two left plots and the two middle plots indicate that the period total fertility rate (TFR) per female has stabilized around 1.8. Nonetheless, the economy has some effect on the fertility rate by period (bottom middle plot). The top right panel shows that the fertility distribution is remarkably stable for the cohorts born since 1970. The bottom plots indicate some time variation in the fertility distribution, thus mainly due to cohorts born before 1970.

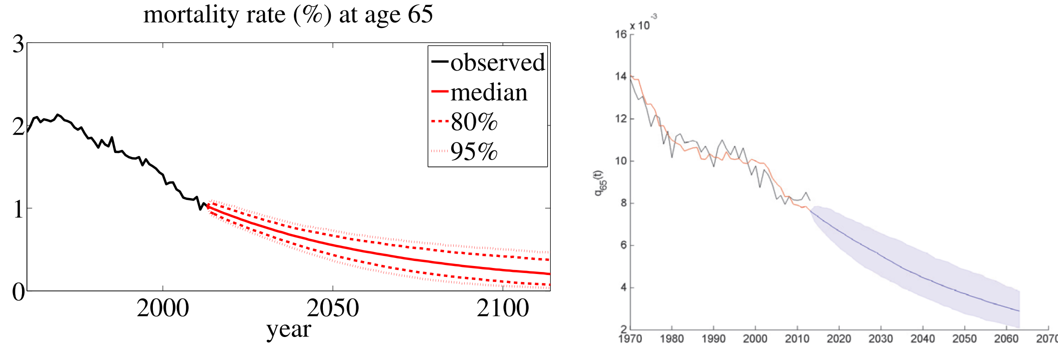

Figure 7: Mortality predictions for 65 years old with 5,000 simulations (left), and for 65 years old females from Fig.1 in AG (2014) with 95% confidence intervals (right).

Figure 8 motivates us to

- impose a long run fertility rate of $\bar{a} = 1.8$ (top middle plot). This is identical to the long run total fertility rate in EC (2014, Table III.19.1), slightly below the rate of 1.85 assumed in Raftery et al. (2014), and slightly above the rate of 1.75 assumed in Van Duin and Stoeldraijer (2014). Lee and Tuljapurkar (1994) restrict the TFR in the U.S. to converge towards the mean replacement rate of 2.1.

- set the long run fertility distribution a_x equal to the fertility distribution of the two most recent years in our sample (see right plots). We scale a_x such that it sums to the long run fertility rate: $\sum_x a_x = \bar{a}$.

The fertility model is a simplified version of the fertility model in Lee and Tuljapurkar (1994), which is based on the mortality model in Lee and Carter (1992). Let $f_{x,t}$ denote the expected number of children born in year t for a female with end-of-year age x. The (period) TFR is then $f_t = \sum_x f_{x,t}$. The dynamics of the common fertility factor F_t determine the fertility at each age x:

$$f_{x,t} = a_x + b_x F_t + \varepsilon_t \tag{9}$$

$$F_t = \rho F_{t-1} + \eta_t + \theta \eta_{t-1} \tag{10}$$

where $x = 15, \ldots, 49$, $\varepsilon_t \sim \left(0, \sigma_x^2\right)$, and $\eta_t \sim N\left(0, \sigma_F^2\right)$. The ARMA(1,1) model in (10) is also in Lee and Tuljapurkar (1994). By imposing, $\rho < 1$, the TFR has unconditional mean $\mathbb{E}[f_t] = \bar{a}$ and unconditional variance

$$\mathrm{Var}\left(f_t\right) = \sum_x b_x^2 \mathrm{Var}\left(F_t\right) + \sigma_x^2 = \sum_x \frac{1+\theta^2}{1-\rho^2} b_x^2 \sigma_F^2 + \sigma_x^2 \tag{11}$$

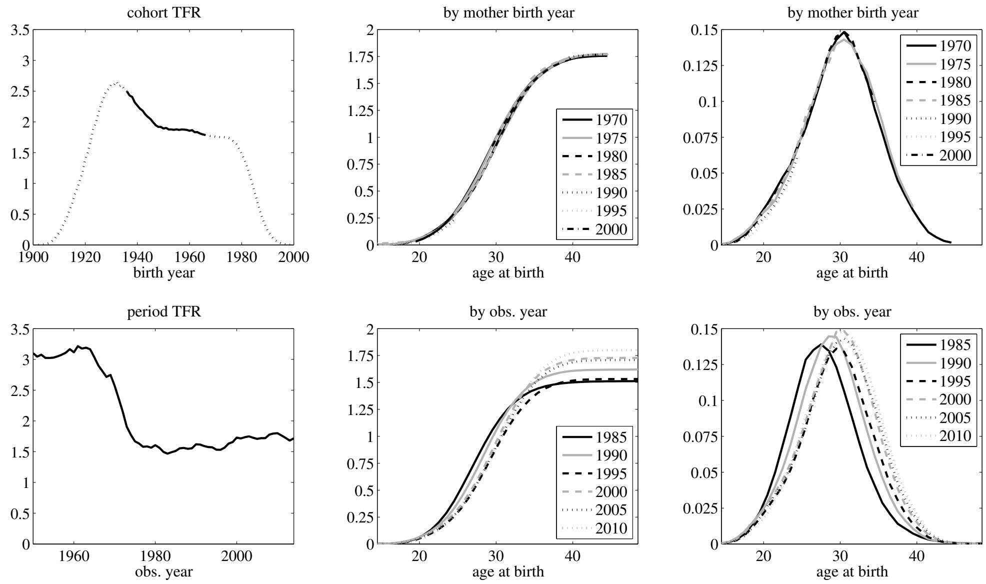

Figure 8: Fertility plots. TFR is the period total fertility rate. Age is the mother age at 31 December, obs. year stands for the observation year. In the top left plot, solid lines are without missing fertility data, and dotted lines have missing data. Since the data is by period, (i) the birth data at low ages is missing for females with earlier birth years, and (ii) the birth data at high ages is missing (still unknown) for females with more recent birth years.

To keep the model simple, we omit (i) time variation in the age pattern of a_x and b_x, (ii) an explicit modelling of cohort effects.[14] None of these effects has an effect on the long run population size nor on the age composition, the focus of our analysis.

Due to data limitations, all births with a mother's age x below 16 are in the age group labeled $x = 15$, and all births with a mother age that exceeds 48 are in the age group labeled $x = 49$. The estimation sample runs from 1960 to 2014.

Estimation Note that $F_t = f_t - \bar{a}$ represents the difference between the TFR at time t and the long run TFR. Each b_x is the slope of the regression of $f_{x,t}$ on an intercept and F_t. It can be verified that $\sum_x b_x = 1$ holds, because by construction a change in F_t leads to an identical change in the TFR f_t.

Similar to the mortality model, parameter uncertainty is taken into account by considering the asymptotic parameter distributions. Figure 9 reports the estimates of a, b, σ, and F in (9). By construction, a_x in the left plot has a similar bell shape in the age x as the right panels in Figure

[14] Broer (2010) models variation in fertility by the variation along the cohort dimension. The estimation does not take any uncertainty into account for cohorts older than 15 at the end of the sample. As a consequence, Figure 4.2 in that paper shows a remarkable small uncertainty in the period TFR for the first three decades of simulation.

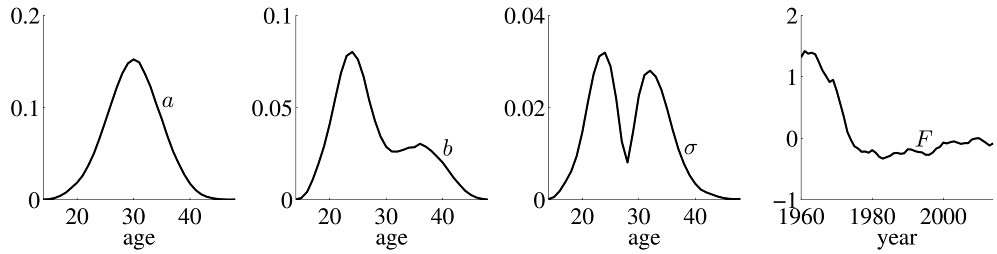

Figure 9: Estimation of (9). The intercept a, slope b, the standard deviation of the residuals σ, and the common time series F.

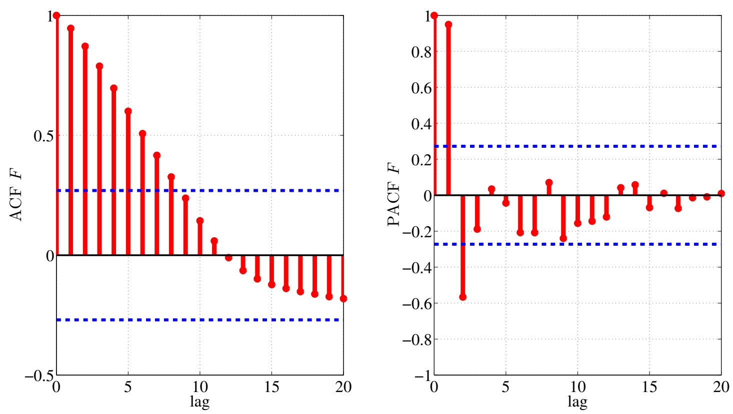

Figure 10: Autocorrelation function (ACF, left) and partial autocorrelation function (PACF, right) of F.

8 where fertility is maximal around the age of 30. In the second plot in Figure 9, the age pattern in b_x reflects that fertility has the highest exposure to F for females in their early twenties. Thus, the decrease in the time series F in the right plot may capture a lower TFR as well as the higher mother birth age due to increases in the period of education for females. The age-specific standard deviation σ in the third plot shows that the model is particularly accurate for females in their late twenties, i.e., exactly when fertility is high.

Figure 10 shows that the autocorrelation function of F is highly persistent, whereas the partial autocorrelation function is significant for the first two lags. Hence, the ARMA(1,1) model in (10) is in line with the Box-Jenkins procedure. The parameters ρ, θ, and σ_F in this equation are obtained with maximum likelihood estimation.[15] Instead of estimating the parameters directly, we estimate

[15] The ARMA(1,1) model is a Kalman filter model with two latent states. The error prediction decomposition gives

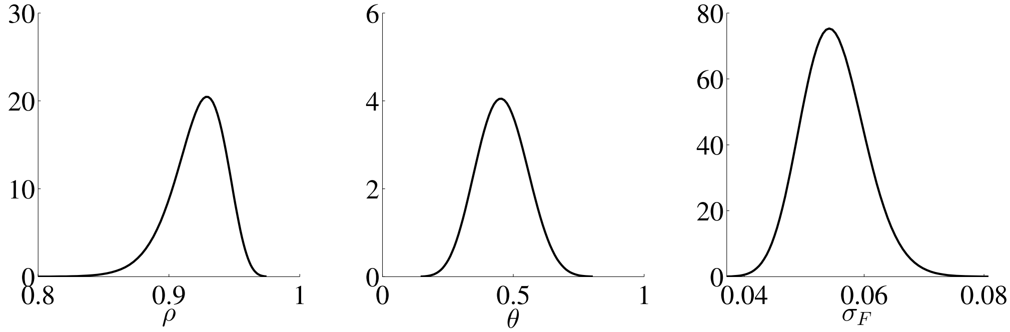

Figure 11: Estimation of (10). Asymptotic marginal parameter distribution of ρ, θ and σ_F.

$\rho = \text{logit}(\tilde{\rho})$, $\theta = \text{logit}\left(\tilde{\theta}\right)$, and $\sigma = \exp(\tilde{\sigma})$ where a tilde indicates an unrestricted parameter. This ensures stationarity of TFR ($\rho, \theta < 1$), and a positive standard deviation ($\sigma > 0$).

We address parameter uncertainty by simulating $\tilde{\rho}$, $\tilde{\theta}$, and $\tilde{\sigma}$ from their asymptotic multivariate distribution. The marginal asymptotic distributions of ρ and θ are logit-normal, while the distribution of σ is lognormal. The three parameter distributions are in Figure 11. The time series is highly persistent as the AR(1) parameter ρ is close to one, while the MA(1) parameter θ is around 0.5. Though the parameter σ_F exceeds the age-specific σ_x (see Figure 9), we find $\frac{1+\theta^2}{1-\rho^2} b_x^2 \ll 1$ in (11). Therefore, most variation in fertility rates is unrelated to the persistent F.

Simulation We draw 5,000 parameter samples for the tuple (ρ, θ, σ). For each parameter sample $(\hat{\rho}, \hat{\theta}, \hat{\sigma})$, we simulate $f_{x,t}$ in (9) by simulating a path for ε_t and η_t. The projection of the TFR is in the left panel in Figure 12. Notice the overlap with the bottom left plot in Figure 8. The projection indicates a small and slow increase of the median TFR towards the long run median of 1.8. The results are qualitatively similar to the projections in UN (2015) (right plot in Figure 8), particularly when taking into account the slightly higher unconditional TFR rate in UN (2015) (1.85 versus 1.8).

2.3 Migration

Section 2.1 and 2.2 provide mortality and fertility as explanatory variables for changes in the population. Similar to EC (2014, p.14), the remaining population changes are referred to as net migration (immigration less emigration) though administrative errors may also explain some of the changes in population.

Migration is mainly determined by exogenous factors that are hard to predict, such as migration policy, educational policy, (foreign) political unrest, climate disasters, etc. As a consequence, net

the likelihood, see for details, e.g., Durbin and Koopman (2012).

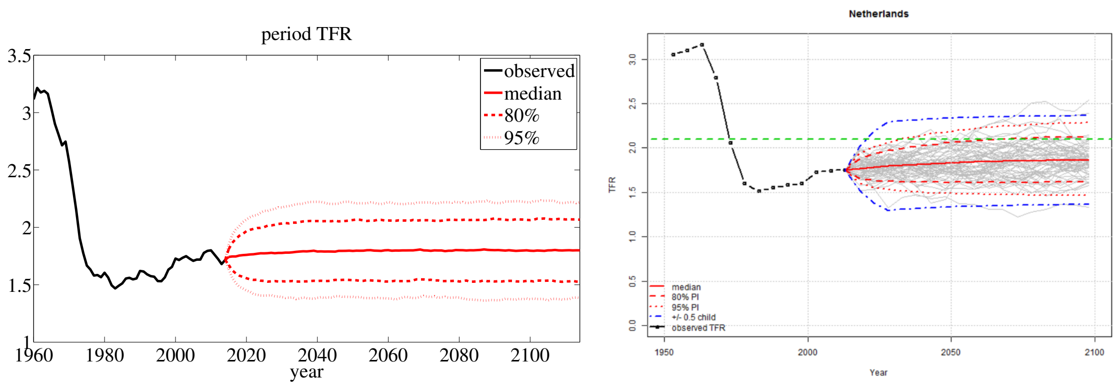

Figure 12: Left: TFR for 5,000 simulations, each from a different parameter sample
Right: TFR in the World Population Prospects of the United Nations, http://esa.un.org/unpd/wpp/Graphs/

migration is trendless. The first principal component of the age-specific net-migration series explains 64% of the variation, compared to 17% for the second component. Since the time series of migration exhibit autocorrelation, we model the common migration factor with an AR(1) factor. In line with the stationary fertility model, we estimate the model

$$n_{x,t} = a_x + b_x N_t + \varepsilon_t \tag{12}$$

$$N_t = \rho N_{t-1} + \eta_t \tag{13}$$

for end-of-year age $x = 0, \ldots, 99$, and mutually independent disturbances $\varepsilon_t \sim \left(0, \sigma_x^2\right)$, and $\eta_t \sim N(0, \sigma_N^2)$.

Estimation and simulation Each intercept term a_x in (12) is the mean age-specific annual net migration since 1970. Under $\rho < 1$, the unconditional mean annual migration $\mathbb{E}[n_{x,t}]$ is the mean migration at age x since 1970, and the unconditional total net migration is $\bar{a} := \sum_x a_x = \mathbb{E}[\sum_x n_{x,t}] = 25.8$ thousand. This substantially differs from the UN model where annual net migration is projected to decrease from 22 thousand in 2060 towards 11 thousand in 2100.[16] EC (2014, Table III.19.1) assumes a decline from 20.8 thousand in 2035 to 9.3 thousand in 2060. Given the globalization and continuing increase in the world population size, we do not consider a decline in migration as the most likely future outcome. Compared to the last four decades, an increase in migration is in our view at least as likely as a decline.

The time series $N_t = n_t - \bar{a}$ is the difference between annual net migration $n_t := \sum_x n_{x,t}$ and its long-run mean $\bar{a}$. Each slope b_x in (12) is the slope in the regression of the time series $n_{x,t}$ on an intercept and n_t. This gives $\sum_x b_x = 1$, similar to the submodels of mortality and fertility.

[16] http://esa.un.org/unpd/wpp/DVD/

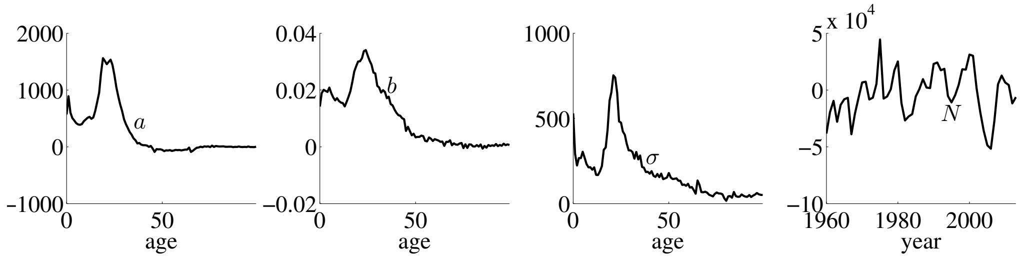

Figure 13: Estimation of (12). The intercept a, slope b, the standard deviation σ of the residuals, and the common time series N.

The age pattern in a (see left plot in Figure 13) suggests that migration is maximal at ages between 10 and 30 years. At such ages, the age pattern in b (second plot) shows that migration is most sensitive to the common factor N. However, the peak in the standard deviation σ of the residuals (third plot) reflects that a substantial part of the variation is still unexplained for this age group. The time series N shows large fluctuations which corresponds to the highly volatile $n = N + \bar{a}$ (right plot).

The small ACF and PACF of N (Figure 14) motivates us to consider a simple AR(1) model for N_t in (13). The persistence of N is relatively low (left plot Figure 15). Generally speaking, the idiosyncratic shock ε_t, and the shock from $b_x N_t$ have a similar instantaneous effect on the net-migration $n_{x,t}$ in (12) since σ_x (Figure 13) and $b_x \sigma_N$ (Figure 13 and 15) have the same order of magnitude. This further downgrades the relative impact of the shocks η_t on the common factor on net-migration. The simulations in Figure 16 confirm a quick convergence of total net-migration towards the asymptotic distribution.

2.4 Aggregate demographic model

We merge the three demographic submodels into one demographic model. Consider the correlation matrix in (14) of the three disturbance terms labeled η for mortality (8), fertility (10), and net-migration (13). The low correlations in (14) reflect a small dependence between the demographic disturbance terms. A positive shock to mortality (a higher mortality) is associated with an insignificant positive shock to fertility (a higher TFR, and a lower mother birth age), and an insignificant positive shock to migration (a higher net inward migration). Note that the mortality shocks are deviations from the trend whereas the mortality trend itself is deterministic.

$$
\begin{array}{cc}
 & \begin{array}{ccc} \text{mort} & \text{fert} & \text{migr} \end{array} \\
\begin{array}{c} \text{mort} \\ \text{fert} \\ \text{migr} \end{array} & \begin{pmatrix} 1 & 0.03 & 0.16 \\ & 1 & 0.09 \\ & & 1 \end{pmatrix}
\end{array}
\tag{14}
$$

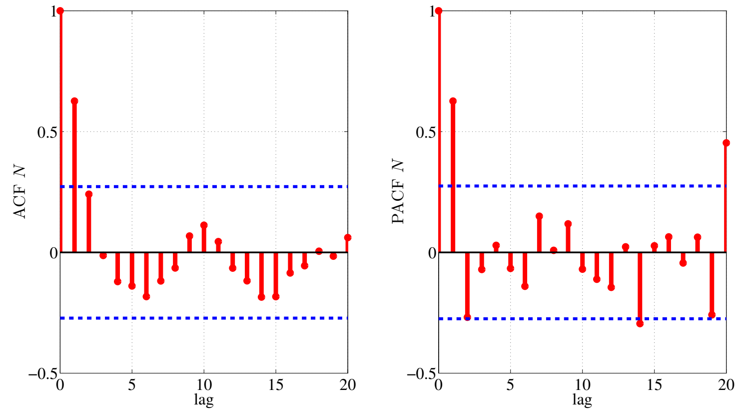

Figure 14: Autocorrelation function (ACF, left) and partial autocorrelation function (PACF, right) of N.

The corresponding projection of the population size is in the left plot in Figure 17. The plot suggests that the population size increases up to 2040, and then stabilizes at 18 million until 2065. This pattern is very similar to the projections in Van Duin and Stoeldraijer (2014, Figure 2.3). During 2040-2065, there is a balance between at one hand the below-replacement rate of fertility, and at the other hand the increase in life expectancy (Figure 6) and the positive net inward migration (Figure 16). The 95% confidence bands in Figure 17 are somewhat narrower than the bands in Van Duin and Stoeldraijer (2014). A likely explanation is the smaller confidence interval around the life expectancy at birth, see page 9.[17]

After fifty years of simulation, most individuals of the relatively large cohorts of children of the baby boomers have died out in 2065. This decrease in the number of deaths leads to a renewed increase in population size. The right panel in Figure 6 shows that the median projection of the UN model misses this renewed increase. Mortality and fertility are almost identical in both models (Figure 5 and 12). By elimination, the most important driver of the renewed increase is the higher migration in our model compared to the UN model. Using half of the projected net migration, we find in unreported results a similar median projection in 2100 as in UN (2015).

3 Productivity model

Following the standard literature on productivity (e.g., OECD (2001)), we measure productivity by output per working hour. More specifically, we measure productivity by GDP per working hour at constant prices. Productivity changes may arise from supply shocks (technological innovation, a

[17] See Statistics' Netherlands StatLine for the exact numbers of the projection of Statistics Netherlands, and Carolina and Van Duin (2010) for the employed methodology.

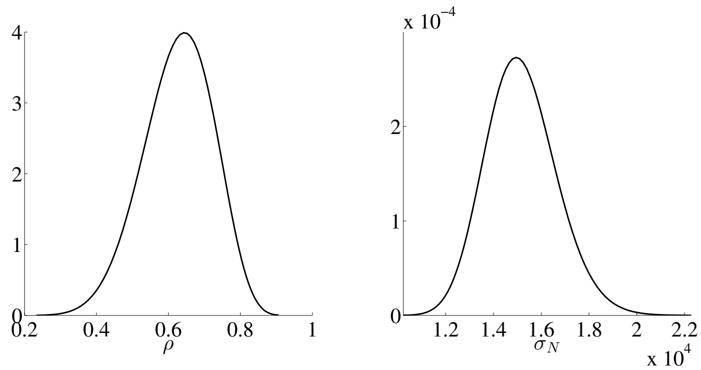

Figure 15: Estimation of (13). Asymptotic marginal parameter distribution of ρ (left) and σ_N (right).

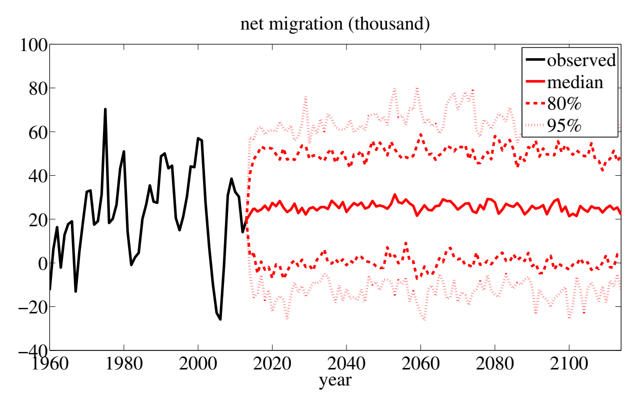

Figure 16: Net migration for 5,000 simulations, each from a different parameter sample.

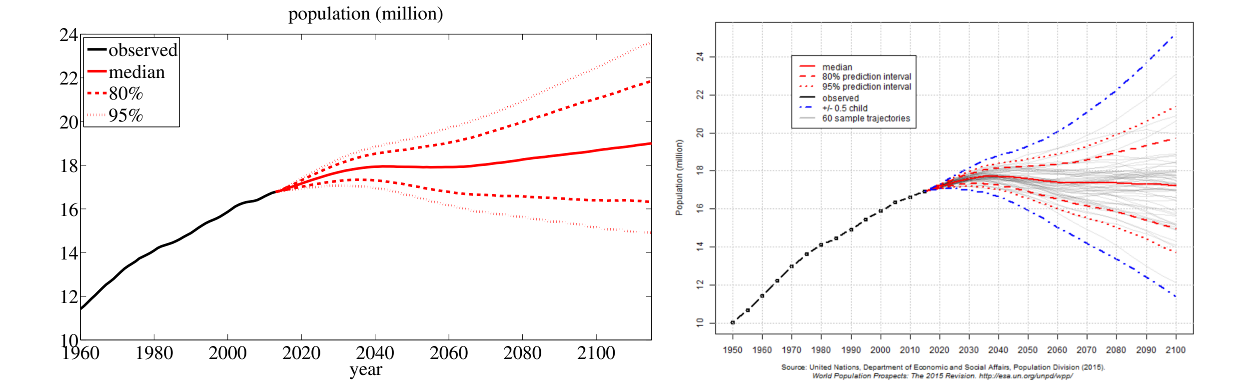

Figure 17: Population size in the Netherlands. Left: Projections for 5,000 simulations, each from a different parameter sample
Right: World Population Prospects of the United Nations, http://esa.un.org/unpd/wpp/Graphs/

change in labour market policy, etc.) or demand shocks (changes in risk perception, interest rates, reduced government spending, exchange rates, etc.). Our empirical model does not distinguish between the two causes.

Figure 18 depicts productivity growth for the Netherlands and the cross-sectional equally weighted mean productivity growth of a set X of 13 countries.[18] Starting from the end of the 70s, we see a slowly evolving trend towards a lower level. In other words, productivity growth shows some persistence and autocorrelation. Accordingly, we estimate a single aggregate productivity factor P_t from the annual (continuously compounded) productivity growth $p_{x,t}$ of each country $x \in X$. The productivity growth data is from the OECD series GDP per hour worked at constant prices. The series P_t is assumed to follow an AR(1) process

$$p_{x,t} = a_x + b_x P_t + \varepsilon_{x,t} \tag{15}$$

$$P_t = \rho P_{t-1} + \eta_t \tag{16}$$

with $x \in X$ and mutually independent disturbances $\eta_t \sim N(0, \sigma_P^2)$ and $\varepsilon_{x,t} \sim N(0, \sigma_x^2)$.[19]

Estimation and simulation

The specification of the productivity model in (15)–(16) is very similar to the fertility model (9)–(10), and the migration model (12)–(13). In (15), the intercept a_x of each country is set equal to a long-run log productivity growth $\bar{a} := \log(1 + 1.4\%) = 1.39\%$. The 1.4% growth rate is the

[18]The set consists of the following countries: Belgium, Canada, Denmark, France, Germany, Italy, Japan, the Netherlands, Spain, Sweden, Switzerland, United Kingdom, and the United States.

[19]A country-specific disturbance for the Netherlands (as in the mortality model) leads to an undesirable high volatility in productivity growth. We prefer the current setup where *real* productivity growth above 3% has a frequency of somewhat less than once in ten years.

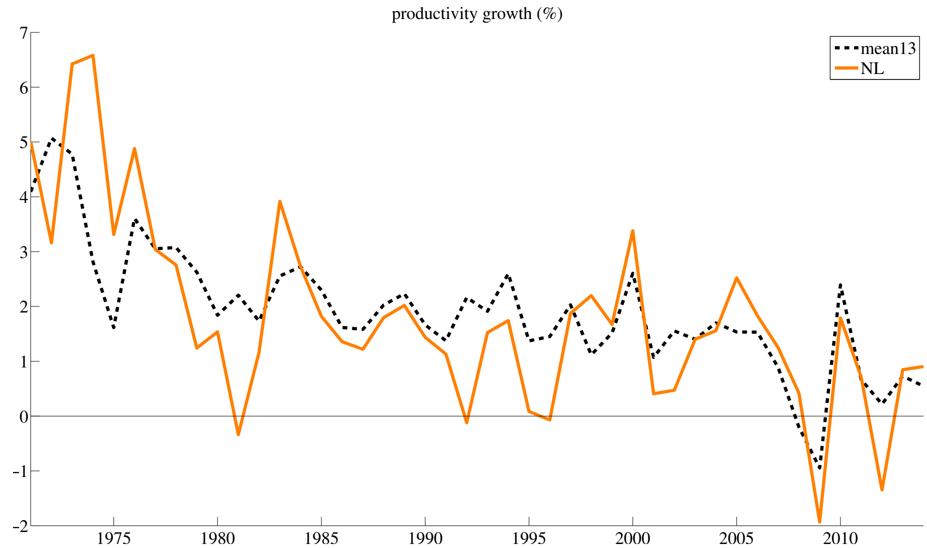

Figure 18: Productivity growth in the Netherlands and 12 other countries.

assumed growth rate of the labour productivity per hour in the EU during 2013-2060 in EC (2014, Table I.3.5). Under $\rho < 1$, the unconditional mean growth rate $\mathbb{E}[p_{x,t}]$ of country x equals $a_x = \bar{a}$. The parameter b_x in (15) is the slope in a regression of the time series $n_{x,t}$ on an intercept and the mean productivity growth $p_t := \frac{1}{|X|} \sum_x p_{x,t}$. By construction, $\frac{1}{|X|} \sum_x b_x = 1$ and $P_t = p_t - \bar{a}$ is the mean cross-sectional deviation from the asymptotic growth rate.

Figure 19 reports parameter estimates of b, σ, and P. Note from $b_{\text{US}} = 0.20$ that US productivity growth is more idiosyncratic than for the other countries. The variance of the disturbance $\sigma^2_{\text{NL}} = 1.27$ in the middle plot is below half of the variance of $p_{\text{NL},t}$ which is 3.0. Thus, the common factor P_t explains more than half of the variation in productivity growth in the Netherlands. The time series P in the right plot, which is simply the cross-sectional mean productivity growth p_t less $\bar{a}$, suggests that it is reasonable to impose that in the long-run $\mathbb{E}[P_t] = 0$. Figure 20 indicates that the AR(1) model in (16) is appropriate for P.[20] The coefficient ρ is centered around 0.67, thereby indicating some persistence in productivity growth. The coefficient σ_P is of the same order of magnitude as σ_{NL}. It can be verified that the long-run annual variance from $b_x P_t$ is $b^2_{NL} \frac{\sigma^2_P}{1-\rho^2} = 1.71$, which exceeds the unexplained variance from $\sigma^2_{\text{NL}} = 1.27$. The top plot in Figure 21 shows that simulated productivity growth in the Netherlands stabilizes around $\bar{a} = 1.39\%$.

In most countries, indexation of pension benefits and pension contributions is mainly determined by price inflation or wage inflation. Wage inflation consists of price inflation and real wage growth. The financial model in Section 4 produces the price inflation projections, whilst we base real wage growth on the productivity growth in this section.

[20]The MA(1) coefficient of an ARMA(1,1) model is close to zero when restricted to the positive half line. We did not consider negative values due to a lack of economic interpretation.

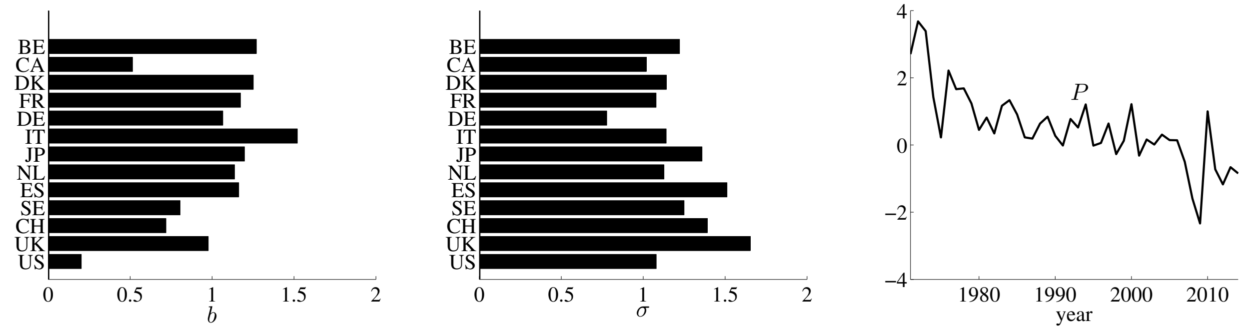

Figure 19: Estimation of (15). The slope b, the standard deviation of the residuals σ, and the common time series P. The intercept $a_x = 1.39\%$ for each country x.

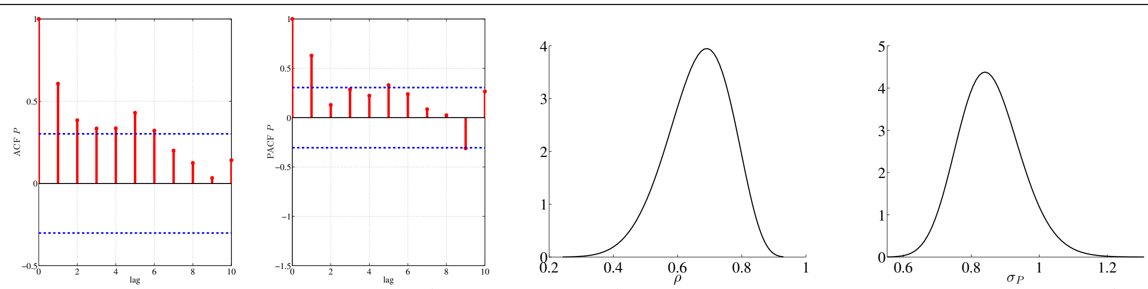

Figure 20: Autocorrelation function (ACF, leftmost) and partial autocorrelation function (PACF, second plot) of P.
Asymptotic marginal parameter distribution in in (16): ρ (third) and σ_P (rightmost).

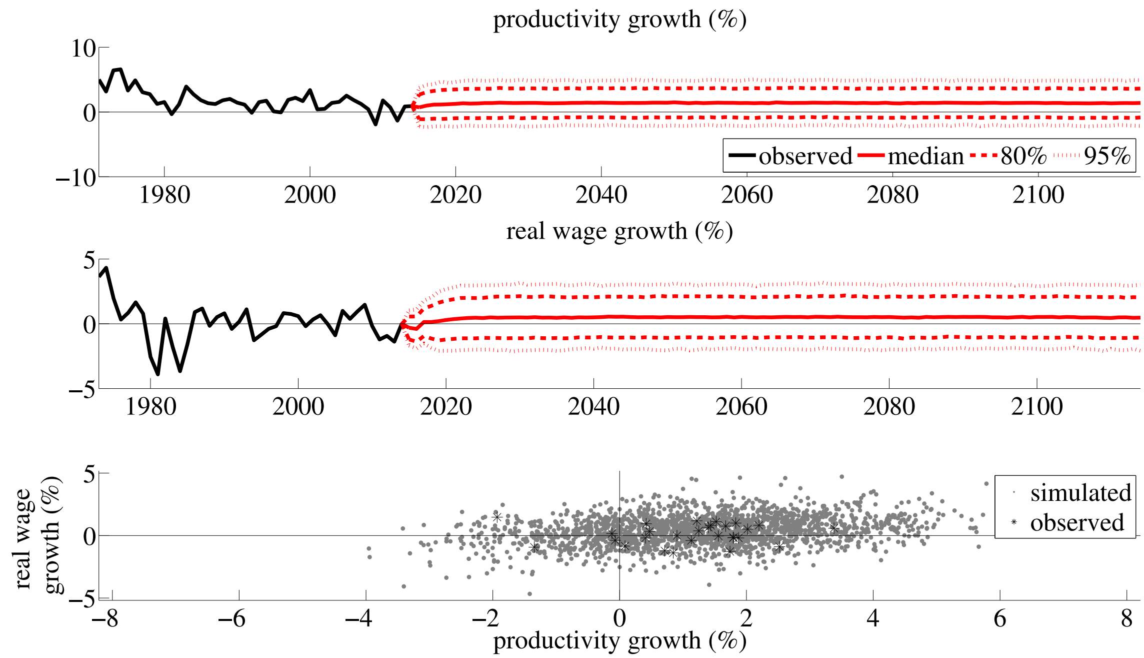

Figure 21: Simulation of productivity growth p_{NL} (top), real wage growth $\tilde{rw}_t$ (middle), and scatter plot of both (bottom).

Langejan et al. (2014) advises a long-run geometric mean growth rate of 0.5% for real contractual wage costs. We regress the continuously compounded real wage growth[21] in year t, in excess of a long run mean of 0.5%, on (i) the cumulative (continuously compounded) Dutch productivity growth rate $\tilde{p}_{NL,t}$ in excess of its long run mean ($\tilde{p}_{\mathrm{NL},t} - \bar{a}$, $\bar{a} = 1.39\%$) in year $t - u_1, \ldots, t - u_0$, and (ii) the cumulative mean common (continuously compounded) productivity growth rate $\tilde{p}_t$ in excess of its long run mean, $\tilde{P}_t = \tilde{p}_t - \bar{a}$ in year $t - v_1, \ldots, t - v_0$. The sample starts in year t_0 and ends in 2014.

To find the optimal setup for the wage regression, we minimize the standard error of the regression subject to the constraints $t_0 < 1990$, $u_0 \geq 0$, $u_1 \geq 0$, $v_0 \geq 0$, and $v_1 \geq 0$. Optimality is at $t_0 = 1986$, $u_0 = 0$, $u_1 = 4$, $v_0 = 0$, and $v_1 = 1$. This result is robust under alternative setups such as annually compounded growth rates, and country-specific long-term productivity growth rates $\bar{a}_x$. The AIC, BIC and the adjusted R^2 are by definition also optimal at these parameter values. The optimal parameter values imply that the sample starts in $t_0 - u_1 = 1982$. This coincides with the Wassenaar Agreement which restrained wage growth in return for the adoption of policies to

[21]Real wage growth is nominal wage growth corrected for CPI inflation. Both time series are from Statistics Netherlands. Nominal wage growth is the growth in employers' labour related expenditures including gratuities, holiday allowances and social charges. The long run dynamics of this growth should be comparable with wage growth in collective agreements plus a smaller incidental component. The first component is the Dutch standard measure for indexation in pension benefits.

combat unemployment and inflation.

The estimated equation is (Newey-West t-statistics in parentheses)[22]

$$\tilde{rw}_t - 0.5\% = \underset{(5.7)}{1.27}\,(\tilde{p}_{\mathrm{NL},t} - \bar{a}) - \underset{(-5.8)}{0.683}\,(\tilde{p}_t - \bar{a}) + e_t \tag{17}$$

where $\tilde{rw}_t = \ln(rw_t/rw_{t-1})$, $\tilde{p}_{\mathrm{NL},t} = \frac{1}{5}\sum_{s=0}^{4} p_{\mathrm{NL},t-s}$, $\tilde{p}_t = \frac{1}{2}\sum_{s=0}^{1} p_{t-s}$, $\bar{a} = \ln(1 + 1.4\%)$, and e_t is a residual with a standard deviation of 0.610%. By the absence of an intercept term in (17), the real wage growth converges towards 0.5%.

Notably, a lagged term $\tilde{rw}_{t-1}$ on the right hand side of (17) is insignificant. Relatedly, the ACF and the PACF of the residuals e_t have a cut-off at the first lag, thereby indicating white noise. The coefficient of the five-year average wage growth $\tilde{p}_{\mathrm{NL},t}$ exceeds one. However, the positive correlation of 0.52 between $\tilde{p}_t$ and $\tilde{p}_{\mathrm{NL},t}$ counterbalances this effect through the reverse effect of $\tilde{p}_t$ in (17). While the coefficient of the international growth factor $\tilde{p}_t$ has an insignificant positive sign in a univariate regression, it becomes significantly negative in the bivariate setup in (17). Therefore, we interpret a positive shock to $\tilde{p}_t$ (conditional on $\tilde{p}_{\mathrm{NL},t}$) as a competitive disadvantage for the Dutch economy, leading to a downward pressure on real wages.

We simulate real wage growth using (17) with a disturbance term that follows a normal distribution with a standard deviation equal to 0.610%. This value corresponds to the standard deviation of the residuals e_t in (17). The two top panels in Figure 21 show that real wage growth is a stationary series with a higher level of persistence than productivity growth. On an annual basis, productivity growth and real wage growth stabilize around 1.39% and 0.5%, respectively. The bottom plot in Figure 21 confirms that in our simulation the instantaneous correlation between productivity growth and real wage growth is similar to the observed correlation in our data sample 1986-2014.

4 Financial model

The setup of the financial submodel is based on the model in Koijen et al. (2010) that we briefly describe.[23] The term structure of continuously compounded nominal swap interest rates on (yields) $y_t(\tau)$ is modelled by an affine two-factor model

$$y_t(\tau) = -\frac{1}{\tau}\left(A(\tau) + B(\tau)' f_t^{(\mathrm{term})}\right) + \xi_{t,\tau}$$

[22] The positive correlation (0.52) of the two explanatory variables $\tilde{p}_{\mathrm{NL},t}$ and $\tilde{P}_t$ contributes to higher standard deviations, and hence lower t-statistics. Still, multicollinearity is of minor importance with a VIF of 1.37. The adjusted R^2 is 0.52 when calculated following the principles in Eisenhauer (2003, eq.(4')).

[23] The reader is referred to Draper (2014) and Muns (2015) for more details on the data set and on estimation of the model. Peter Vlaar kindly provided the dataset. The quarterly time series run from 1973 to 2014. The parameter estimates are adjusted to the annual model.

The 2-year and 5-year yields are assumed to be observed without measurement error, $\xi_{t,2} = \xi_{t,5} \equiv 0$:[24]

$$\begin{bmatrix} -2y_t\,(2) \\ -5y_t\,(5) \end{bmatrix} = \begin{bmatrix} A(2) \\ A(5) \end{bmatrix} + \begin{bmatrix} B(2)' \\ B(5)' \end{bmatrix} \begin{bmatrix} f_{1,t}^{(\text{term})} \\ f_{2,t}^{(\text{term})} \end{bmatrix} \tag{18}$$

The functions $A(\tau)$ and $B(\tau)$ follow from the model parameters, see below. The function values at $\tau = 2$ and $\tau = 5$ imply the two latent factors in $f^{(\text{term})}$. Both latent factors are stationary with a long run mean of zero. The factors jointly follow a VAR(1) process with the shocks $\varepsilon_{1,t}$ and $\varepsilon_{2,t}$ having an independent standard normal distribution:

$$\mathrm{d}f_t^{(\text{term})} = -Kf_t^{(\text{term})}\mathrm{d}t + \begin{bmatrix} \varepsilon_{1,t} & \varepsilon_{2,t} \end{bmatrix}' \tag{19}$$

The price of risk $\tilde{\Lambda}_t \in \mathbb{R}^2$ is affine in the latent factors:

$$\tilde{\Lambda}_t = \tilde{\Lambda}_0 + \tilde{\Lambda}_1 f_t^{(\text{term})}$$

Imposing no-arbitrage arguments implies that the intercept function $A : \tau \to \mathbb{R}$ and the slope function $B : \tau \to \mathbb{R}^2$ satisfy

$$\begin{aligned} A(\tau) &= \int_0^\tau \dot{A}(s)\mathrm{d}s \\ \dot{A}(\tau) &= -\delta_{0R} - \tilde{\Lambda}_0' B(\tau) + \frac{1}{2} B(\tau)' B(\tau) \\ B(\tau) &= \left(K' + \tilde{\Lambda}_1'\right)^{-1} \left[\exp\left(-\left(K + \tilde{\Lambda}_1\right)' \tau\right) - I_{2\times 2}\right] \delta_{1R} \end{aligned}$$

with $A(0) = B_1(0) = B_2(0) = 0$. The interest rate, inflation, stock return, and bond porfolio return depend on the two latent factors and possibly a variable-specific disturbance term:

$$\begin{aligned} R_t(0) &= \delta_{0R} + \delta_{1R}' f_t^{(\text{term})} && \text{(Instantaneous nom. interest rate)} \\ \pi_t &= \delta_{0\pi} + \delta_{1\pi}' f_t^{(\text{term})} && \text{(Expected inflation)} \\ \frac{\mathrm{d}\Pi_t}{\Pi_t} &= \pi_t \mathrm{d}t + \sigma_\Pi' \varepsilon_t && \text{(Realized inflation)} \\ \frac{\mathrm{d}S_t}{S_t} &= \left(R_t(0) + \eta_S\right)\mathrm{d}t + \sigma_S' \varepsilon_t && \text{(Realized stock return)} \\ \frac{\mathrm{d}P_t^B(\tau)}{P_t^B(\tau)} &= \left(R_t(0) + B(\tau)'\tilde{\Lambda}_t\right)\mathrm{d}t + B(\tau)'\varepsilon_{t(1:2)} && \text{(Realized bond portfolio return)} \end{aligned}$$

[24]This differs from the Kalman procedure described in Appendix C of Koijen et al. (2010). However, Table 5 in their paper shows a zero measurement error for the one year and five year interest rates. As such, Koijen et al. (2010) employ a similar procedure as in Draper (2014).

Latent factor nom. term structure $\mathrm{d}f_t^{(\text{term})} = -Kf_t^{(\text{term})}\mathrm{d}t + \left[\varepsilon_{1,t} \quad \varepsilon_{2,t}\right]'$		Instantaneous nominal interest rate $R_t(0) = \delta_{0R} + \delta_{1R}' f_t^{(\text{term})}$	
$K_{(1,1)}$	7.63%	δ_{0R}	2.40%
$K_{(2,1)}$	-19.00%	$\delta_{1R(1)}$	-1.48%
$K_{(2,2)}$	35.25%	$\delta_{1R(2)}$	0.53%
Expected inflation $\pi_t = \delta_{0\pi} + \delta_{1\pi}' f_t^{(\text{term})}$		Realized inflation $\frac{d\Pi_t}{\Pi_t} = \pi_t \mathrm{d}t + \sigma_\Pi' \varepsilon_t$	
$\delta_{0\pi}$	2.00%	$\sigma_{\Pi(1)}$	0.02%
$\delta_{1\pi(1)}$	-0.63%	$\sigma_{\Pi(2)}$	-5.68$\cdot10^{-3}$%
$\delta_{1\pi(2)}$	0.14%	$\sigma_{\Pi(3)}$	0.61%
Realized stock return $\frac{dS_t}{S_t} = (R_t + \eta_S)\,\mathrm{d}t + \sigma_S' \varepsilon_t$		Prices of risk $\Lambda_t = \Lambda_0 + \Lambda_1 f_t^{(\text{term})}$	
η_S	4.52%	$\Lambda_{0(1)}$	0.280
$\sigma_{S(1)}$	-0.53%	$\Lambda_{0(2)}$	0.027
$\sigma_{S(2)}$	-0.76%	$\Lambda_{1(1,1)}$	0.149
$\sigma_{S(3)}$	-2.11%	$\Lambda_{1(1,2)}$	-0.381
$\sigma_{S(4)}$	16.59%	$\Lambda_{1(2,1)}$	0.089
		$\Lambda_{1(2,2)}$	-0.083

Table 1: Parameters estimated in Draper (2014) and calibrated by DNB for the scenario set of their feasibility test of 2015Q2. The parameter symbols are identical to Koijen et al. (2010) and Muns (2015).

where $\varepsilon_t \in \mathbb{R}^4$ contains independent standard normal shocks, $\sigma_{\Pi(4)} = 0$, and η_S is the constant equity premium.

Draper (2014) estimates the model using Dutch quarterly data up to 2013, and converts the estimated model from a quarterly frequency to an annual frequency.[25] DNB calibrated the parameter set to satisfy some restrictions on the long run expectation of $y_t(\tau)$ and π_t. The resulting parameters are in Table 1.

The corresponding term structure coefficients are in Figure 22. The line $-A(\tau)/\tau$ is the unconditional expectation of the nominal interest rate $R_t(\tau) = \lim_{t\to\infty} \mathbb{E}[y_t(\tau)]$. It is concave and upward sloping. Both are a familiar property of the term structure. The functions $-B_1(\tau)/\tau$ and $-B_2(\tau)/\tau$ converge to zero for long maturities τ. This reflects that $y_t(\tau)$ is for large τ independent of the current state variables $f_t^{(\text{term})}$.

Table 2 presents long run statistics of some variables. The interest rate $y_t(\tau)$ converges to the exogenously imposted UFR of 4.18%.[26] The long run inflation rate is close to the inflation target

[25] Muns (2015) derives expressions for the ultimate forward rate (UFR), and the long run means and covariances of the included variables.

[26] DNB has recently revised the method to determine the UFR ("DNBulletin: Adjustment of UFR results in more realistic actuarial interest rate for pensions"). Instead of the previously fixed UFR at 4.2%, the UFR is now a 120-months moving average of the 1-year forward over 20-years. Using $\tau = 20$ and $\tau = 21$, this forward follows endogenously by applying (12), (37), and (39) from Muns (2015) to the simulations.

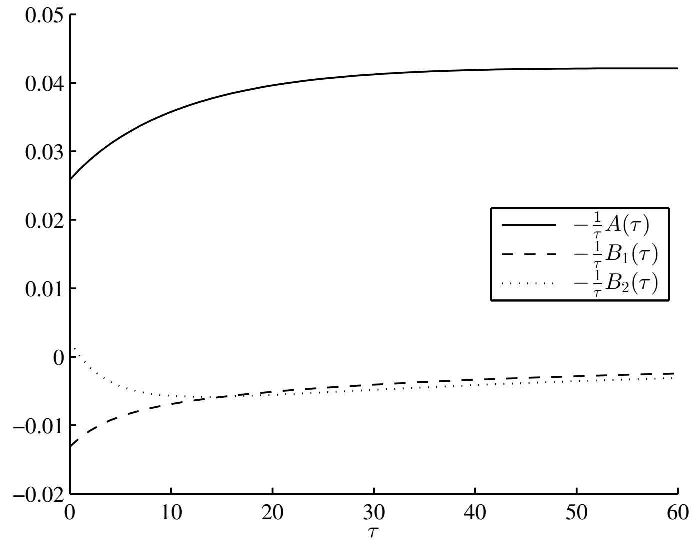

Figure 22: Coefficients of the term structure $\mathbb{E}[y_t(\tau)]$.

of 2%. The excess stock return is 3.2% on an annual basis, while a bond portfolio with a constant maturity of 5 years has an excess return of 1.3%. The real instantaneous interest rate $R_t(0) - \pi_t$ is in the long run 0.4%. In line with empirical practice, the volatility of a stock portfolio exceeds the volatility of a bond portfolio.

5 Implementation

This section merges the submodels of the previous sections into one aggregate model. In total, the annual model contains eight common risk factors. Each factor has a disturbance that follows a

	geometric mean	annual st.dev.
UFR	4.18%	NA
π	2.02%	1.59%
R^S	5.67%	18.43%
$R(0)$	2.43%	3.29%
$R^B(5)$	3.69%	5.92%

Table 2: Ultimate forward rate (UFR) and long run statistics of inflation Π, stock return R^S, instantaneous nominal interest rate $R(0)$, and bond portfolio return $R^B(5)$ with a constant duration of 5 years. Parameters are from the calibrated DNB parameter set, derivations are in Muns (2015).

multivariate normal distribution. The factors represent risks in mortality (2), fertility, migration, productivity, real wages, and financial risk (2) through the term structure.

5.1 Disturbance correlations

From a theoretical perspective, it is possible to estimate the submodels jointly. This estimation would directly produce a covariance matrix for the disturbances. However, this procedure performs very slow as the estimation of the financial submodel takes more than an hour. Moreover, we might still end up at a local optimum as the number of simultaneously estimated parameters increases substantially. Accordingly, we first estimated each submodel in isolation and estimate the disturbance correlations afterwards.

We find the correlations of the common factors from the pairwise correlation of the residuals $\hat{\eta}_t$ of the latent factor in each block. Within the financial block, we use the theoretical correlation from the discretized version of (19) regardless of the empirically observed correlation.[27] Also due to different sample periods of the submodels, this procedure results in a disturbance covariance matrix Σ that is not positive definite, a prerequisite for a covariance matrix. Next, we discuss how to adjust the matrix Σ in such a way that it is positive definite.

5.2 The latent factor disturbance covariances

We have theoretical estimates for each disturbance variance and the correlations of the disturbances within the financial model. We estimate the remaining pairwise correlations from the time series of the estimated factors. However, it turns out that the resulting matrix is not positive definite, a prerequisite for a covariance matrix. Hence, we adjust the correlation of a pair of disturbances if one disturbance is in the financial submodel and the other disturbance in one of the other submodels.

We have a set X with nonfinancial variables, and a set Y with financial variables. The covariance matrix is as follows

$$\Sigma = \begin{pmatrix} \Sigma_{XX} & \Sigma_{XY} \\ \Sigma'_{XY} & \Sigma_{YY} \end{pmatrix} \tag{20}$$

An empirical estimate is available for each pairwise correlation. The employed sample is the common sample of the two time series, which may differ across different time series.

$$\hat{\Sigma} = \begin{pmatrix} \hat{\Sigma}_{XX} & \hat{\Sigma}_{XY} \\ \hat{\Sigma}'_{XY} & \hat{\Sigma}_{YY} \end{pmatrix} \tag{21}$$

A theoretical estimate is available for Σ_{YY}, while we need to rely on an empirical estimate for Σ_{XX} and Σ_{XY}. Unfortunately, letting $\Sigma_{XX} = \hat{\Sigma}_{XX}$ and $\Sigma_{XY} = \hat{\Sigma}_{XY}$ may not lead to a positive definite covariance matrix, though $\hat{\Sigma}_{XX}$ and $\hat{\Sigma}_{YY}$ are of course positive definite. We resolve this issue by

[27] The correlation depends on the parameters and is in general unequal to the zero correlation of the continuous counterpart in (19).

setting $\Sigma_{XX} = \hat{\Sigma}_{XX}$ and adjusting the estimate for Σ_{XY}. Consider the Cholesky decompositions $\Sigma = U'U$, $\hat{\Sigma} = \hat{U}'\hat{U}$, $\Sigma_{II} = U_I'U_I$, and $\hat{\Sigma}_{II} = \hat{U}_I'\hat{U}_I$ ($I \in \{X, Y\}$). Notice that

$$U = \begin{pmatrix} U_X & U_{XY} \\ O & U_{YY} \end{pmatrix} \quad \hat{U} = \begin{pmatrix} \hat{U}_X & \hat{U}_{XY} \\ O & \hat{U}_{YY} \end{pmatrix}$$

where O is a zero matrix of appropriate dimension. It follows from

$$\Sigma_{YY} = U_{XY}'U_{XY} + U_{YY}'U_{YY} \tag{22}$$

that $U_{YY} \neq U_Y$ if $U_{XY} \neq O$. Therefore, only the block $U_X = \hat{U}_X$ of U is known while the blocks U_{XY} and U_{YY} are in general unknown. We find U_{XY} and U_{YY} by transforming their empirical counterparts $\hat{U}_{XY}$ and $\hat{U}_{YY}$ in such a way that the covariance matrix of Y equals Σ_{YY}:

$$\begin{aligned}
\Sigma_{YY} &= U_Y'U_Y \\
&= U_Y'\left(\hat{U}_Y'\right)^{-1}\hat{\Sigma}_{YY}\hat{U}_Y^{-1}U_Y \\
&= U_Y'\left(\hat{U}_Y'\right)^{-1}\left(\hat{U}_{XY}'\hat{U}_{XY} + \hat{U}_{YY}'\hat{U}_{YY}\right)\hat{U}_Y^{-1}U_Y \\
&= U_Y'\left(\hat{U}_Y'\right)^{-1}\hat{U}_{XY}'\hat{U}_{XY}\hat{U}_Y^{-1}U_Y + U_Y'\left(\hat{U}_Y'\right)^{-1}\hat{U}_{YY}'\hat{U}_{YY}\hat{U}_Y^{-1}U_Y
\end{aligned} \tag{23}$$

Combining (22) and (23), the two missing blocks of U are

$$U_{XY} = \hat{U}_{XY}\hat{U}_Y^{-1}U_Y \qquad U_{YY} = \hat{U}_{YY}\hat{U}_Y^{-1}U_Y.$$

Accordingly, we obtain the covariance matrix $\Sigma = U'U$ from

$$U = \begin{pmatrix} \hat{U}_X & \hat{U}_{XY}\hat{U}_Y^{-1}U_Y \\ O & \hat{U}_{YY}\hat{U}_Y^{-1}U_Y \end{pmatrix}.$$

5.3 Estimation results

Table 3 contains the correlations of the eight annual disturbances. It follows that most correlations are low. The small negative correlation between mortality and productivity growth translates into a small positive correlation between the period life expectancy and cumulative productivity growth. Allowing for shocks to the trend of the factors may lead to a stronger association between life expectancy and productivity. Nonetheless, it seems unlikely that productivity improvements have a substantial instantaneous relation with mortality. In addition, there is a consensus in the literature (e.g., see Acemoglu and Johnson (2007)) that income inequality rather than the level of income affects life expectancy among high-income countries.

The correlation between productivity and fertility is close to zero, while the correlation between productivity and net-migration is negative. Thus, a positive productivity shock is associated with a lower net-migration, or even a net outflow. To assess the contemporaneous relation between

	eq.	mort EU	mort NL-EU	fert	migr	prod	rw	term1	term2
mort EU	(3)	1	-0.18	0.03	0.16	-0.13	0.15	0.05	-0.23
mort NL-EU	(4)		1	0.06	0.03	0.14	0.45	0.33	0.13
fert	(10)			1	0.09	-0.06	0.17	-0.14	0.14
migr	(13)				1	-0.40	0.27	0.05	-0.15
prod	(16)					1	0.02	-0.30	-0.02
rw	(17)						1	0.18	0.04
term1	(19)							1	0.1
term2	(19)								1

Table 3: Correlations of the annual disturbances of the common factors, and the real wage rw.

productivity and demography in more detail, we run 2,000 simulations over 50 years, and find a correlation coefficient of -0.24 between productivity growth and population size. Figure 23 suggests that the impact of this negative relation is quite low. Notice the high uncertainty in cumulative productivity growth. A cumulative productivity growth of zero appears as likely as a cumulative growth of 130%.

We detect a positive correlation of 0.19 between cumulative productivity growth and the old-age dependency ratio, defined as the ratio of the population size above 65 years old and the population size between 21 and 65 years old.[28] The negative correlation between net-migration and productivity results in a negative correlation between productivity growth and the population size between 21 and 65 years. The resulting positive correlation between productivity growth and the old-age dependency ratio is beneficial in a DB PAYG scheme because a higher productivity growth partly offsets a higher dependency ratio.

By construction, the correlation of the productivity shock and the real wage residuals is low (0.02) as the productivity residuals are explanatory variables through $\tilde{p}_{\mathrm{NL},t}$ in the regression (17).

All positive maturities depend negatively on shocks to the first latent term factor through the negative coefficients $-B_1(\tau)$ (Figure 22). This means that a negative shock to the term structure is associated with a negative shock to productivity growth. This is intuitive as an economic downturn is associated with lower interest rates as well as a lower productivity growth. The small correlations between demographic shocks and the term factors indicate that demography and interest rates have a weak correlation, at least in the short-run. This is in line with the mixed evidence in the literature.[29] The low correlation between demographic shocks and financial shocks suggests diversification benefits of a multi-pillar pension system since the main driver of variation in the first

[28] The correlation decreases to 0.12 with a 100 years of simulation, i.e., until 2114.

[29] see Poterba (2001), Poterba (2004), Broer (2010), and the references in this literature.

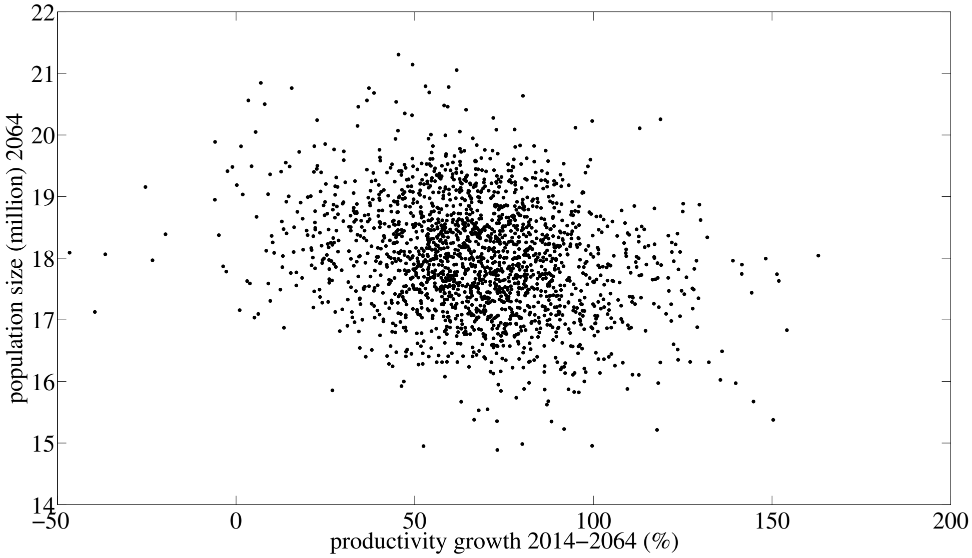

Figure 23: Scatter plot of productivity growth and population size with 2,000 simulation runs of 50 years.

and second pillar are demographic shocks and financial shocks, respectively.

6 Potential improvements

Though our model generates plausible results, improvements might be possible along the following dimensions:

(i) Our estimation technique does not exploit the implicit dependence between the coefficients of different ages in our model (see the smooth age-patterns in Figure 1). The smoothing approaches proposed in Currie et al. (2004) and Hyndman and Ullah (2007) could be helpful in reducing the number of age-specific parameters. Still, the results in Cairns et al. (2009) on English and Welsh mortality data for the ages 60-89 suggest a poor performance of the approach in Currie et al. (2004). In addition, Figure 1 indicates that this approach is even less useful at low ages since the parameters are more volatile at such ages. Booth et al. (2006) show that the accuracy of the approach in Hyndman and Ullah (2007) in modelling life expectancy is not significantly better than the standard Lee and Carter (1992) model.

(ii) The fertility model does not take the rather stable cohort effect into account (top left plot in Figure 8). This suffices for our planned purposes. Nonetheless, further research along the cohort dimension is needed if one wants to estimate the number of children per female, e.g., for variations in child allowance.

(iii) The quarterly observations of the financial submodel result in a high standard deviation for the parameter estimates (see Draper (2014) and Muns (2015)). For a similar model with monthly U.S. data, Koijen et al. (2010) obtain substantially lower standard deviations. Unfortunately, monthly Dutch data was unavailable for the financial submodel.

7 Conclusions

For the next decades, Dutch pension funds need to deal with several headwinds simultaneously. The continuing ageing process makes demography an adverse factor. A low real interest rate makes pension saving costly. In addition, a lower productivity growth may translate into lower wages, and hence reduces both pension contributions and indexation of pension benefits. The latter affects the economies of scale of pension funds.

By combining well-known models from the literature, this paper models the joint distribution of these risks for the Netherlands. The simple AR(1)-structure of each submodel intends to avoid overfitting. The model allows for different sample lengths for the submodels, and can generate an arbitrarily large number of scenarios.

The confidence bands for the demographic projections are similar to the UN model. Likewise, our estimation yields population projections with the familiar increase in total population up to 2040, and a stable size up to 2060. Our results indicate that migration policy is a key determinant of the population size after 2060.

Notably, the dependence between the annual disturbances is low. The low correlation between demographic shocks and financial shocks suggests diversification benefits of a multi-pillar pension system. The most significant correlation is a negative correlation of -0.4 between net inward migration and productivity growth. This works in a compensating way in a DB PAYG pension scheme.

The model's aim is to be a building block in evaluating a wide range of policy reforms on pension policy. Evaluations are possible at an age-specific, a cohort-specific or an aggregate level, and for short run as well as for a long run analysis of more than a few decades. Still, one should keep in mind that projections at long horizons are subject to an inherent high uncertainty (see Keyfitz (1981)).

Appendix

A Ordinary least squares with restrictions

This appendix outlines some properties of OLS regressions with the functional form $\rho(a)$ for the slope and $\sigma = \exp(\tilde{\sigma})$ for the standard error. The latter ensures $\sigma > 0$. Consider the model

$$y_t = c + \rho(a)x_t + \varepsilon_t \qquad \varepsilon_t \sim N(0, \sigma^2)$$

The time series $\{x_t\}_{t=1}^T$ and $\{y_t\}_{t=1}^T$ are observed, and the functional form of the monotonic function ρ is known. We describe the estimation procedure of the unknown parameters a, c, and $\sigma > 0$, and the corresponding asymptotic standard errors. The likelihood function is

$$L(c,a,\sigma)=\prod_{t=1}^{T}\frac{1}{\sigma\sqrt{2\pi}}\exp\left(-\frac{(y_t-c-\rho(a)x_t)^2}{2\sigma^2}\right).$$

Define $\tilde{\sigma}=\ln(\sigma)$, $\mathbf{x}=(x_1,\ldots,x_T)$, $\mathbf{y}=(y_1,\ldots,y_T)$, and $\mathbf{1}$ an all-ones vector of appropriate dimension. The loglikelihood function is

$$\begin{aligned}LL(c,a,\tilde{\sigma}) &= -\frac{T}{2}\ln(2\pi)-T\tilde{\sigma}-\frac{1}{2}\exp(-2\tilde{\sigma})\sum_{t=1}^{T}(y_t-c-\rho(a)x_t)^2.\\ &= -\frac{T}{2}\ln(2\pi)-T\tilde{\sigma}-\frac{1}{2}\exp(-2\tilde{\sigma})(\mathbf{y}-c\mathbf{1}-\rho(a)\mathbf{x})'(\mathbf{y}-c\mathbf{1}-f(a)\mathbf{x}).\end{aligned}$$

The gradient is

$$\nabla LL(c,a,\tilde{\sigma})=\begin{bmatrix}LL_a\\ LL_c\\ LL_{\tilde{\sigma}}\end{bmatrix}=\begin{bmatrix}\exp(-2\tilde{\sigma})\,\mathbf{1}'(\mathbf{y}-c\mathbf{1}-\rho(a)\mathbf{x})\\ \rho'(a)\exp(-2\tilde{\sigma})\,\mathbf{y}'(\mathbf{y}-c\mathbf{1}-\rho(a)\mathbf{x})\\ -T+\exp(-2\tilde{\sigma})(\mathbf{y}-c\mathbf{1}-\rho(a)\mathbf{x})'(\mathbf{y}-c\mathbf{1}-\rho(a)\mathbf{x})\end{bmatrix}$$

and the Hessian matrix is (use $\sigma=\exp(\tilde{\sigma})$)

$$H(c,a,\tilde{\sigma})=\begin{bmatrix}-\frac{T}{\sigma^2} & -\frac{\rho'(a)}{\sigma^2}\mathbf{1}'\mathbf{x} & -2LL_c\\ -\frac{\rho'(a)}{\sigma^2}\mathbf{1}'\mathbf{x} & \frac{\rho''(a)}{\rho'(a)}LL_a-\left(\frac{\rho'(a)}{\sigma^2}\right)^2\mathbf{x}'\mathbf{x} & -2LL_a\\ -2LL_c & -2LL_a & -2(LL_{\tilde{\sigma}}+T)\end{bmatrix}$$

The maximum likelihood estimate satisfies $\nabla LL(c^*,a^*,\tilde{\sigma}^*)=(0,0,0)'$ such that

$$\begin{bmatrix}\mathbf{1}'\mathbf{e}^*\\ \mathbf{x}'\mathbf{e}^*\\ \exp(-2\tilde{\sigma})(\mathbf{e}^*)'\mathbf{e}^*\end{bmatrix}=\begin{bmatrix}0\\0\\T\end{bmatrix}\qquad \mathbf{e}^*=\mathbf{y}-c^*\mathbf{1}-\rho(a^*)\mathbf{x}$$

Thus, c^* and $\rho(a^*)$ satisfy the orthogonality conditions $\mathbf{1}'\mathbf{e}^*=\mathbf{x}'\mathbf{e}^*=0$. Therefore, $\rho(a^*)$ and c^* are the (unbiased) OLS estimates:

$$\begin{pmatrix}c^*\\ \rho(a^*)\end{pmatrix}=(X'X)^{-1}X'\mathbf{y}\qquad X=\begin{bmatrix}\mathbf{1} & \mathbf{x}\end{bmatrix}$$

The parameter a^* follows from inverting $\rho(a^*)$. After some rewriting, we obtain the well-known biased variance estimator

$$(\sigma^*)^2=\exp(2\tilde{\sigma}^*)=\frac{(\mathbf{e}^*)'\mathbf{e}^*}{T},$$

while the unbiased variance estimator takes the number k of free parameters in X into account (e.g., Heij et al. (2004, eq.(3.22))):

$$\sigma_U^2 = \frac{(\mathbf{e}^*)' \, \mathbf{e}^*}{T-k}$$

At the unbiased parameter estimate ($\tilde{\sigma}_U^* = \ln(\sigma_U^*)$),

$$H(c^*, a^*, \tilde{\sigma}_U^*) = -\frac{1}{(\sigma_U^*)^2} \begin{bmatrix} T & \rho'(a^*)\mathbf{1}'\mathbf{x} & 0 \\ \rho'(a^*)\mathbf{1}'\mathbf{x} & (\rho'(a^*))^2 \, \mathbf{x}'\mathbf{x} & 0 \\ 0 & 0 & 2T(\sigma_U^*)^2 \end{bmatrix}$$

Thus, the estimate of σ_U is uncorrelated with the estimate for c^* and a^*. The asymptotic covariance matrix of the latter two parameters is

$$\Sigma_{a^*c^*} = \sigma_U^2 \left[(X^*)' \, X^*\right]^{-1} \qquad X^* = \begin{bmatrix}\mathbf{1} & \rho'(a^*)\mathbf{x}\end{bmatrix} \tag{24}$$

Indeed, this expression coincides with the OLS covariance matrix if $\rho(a) = a$. At the ML estimate $\tilde{\sigma}^*$, the asymptotic variance of $\tilde{\sigma}$ equals $\Sigma_{\tilde{\sigma}^*} = 1/(2T)$ regardless of c, a, and the functional form of ρ.

For an AR(1)-model, we sample $\left(\hat{c}, \hat{a}, \hat{\tilde{\sigma}}\right)$ from $N(\mu_U, \Sigma^*)$ where

$$y_t = x_{t+1} \qquad \mu_U = (c^*, a^*, \tilde{\sigma}_U) \qquad \Sigma^* = \begin{bmatrix} \Sigma_{a^*c^*} & 0_{2\times 1} \\ 0_{1\times 2} & \frac{1}{2T} \end{bmatrix}$$

the covariance matrix is as in (24), and ρ is a logit-normal distribution

$$\rho(a) = \operatorname{logit}(a) = \frac{1}{1+\exp(-a)}$$

$$\rho'(a) = \frac{\exp(-a)}{[1+\exp(-a)]^2} = \frac{1}{2+\exp(-a)+\exp(a)}$$

This functional form of $\rho(a)$ ensures that the AR(1)-coefficient $\rho(a)$ is smaller than one. It follows that

$$\rho(a) \sim (1+\exp(-a^* + \sigma_{a^*}Z))^{-1}$$

$$\sigma \sim \exp(\sigma_{\tilde{\sigma}^*}Z + \tilde{\sigma}_U) = \exp\left(\frac{1}{\sqrt{2T}}Z + \tilde{\sigma}_U\right),$$

where $Z \sim N(0,1)$. Simulation paths are

$$x_t = \hat{c} + \rho(\hat{a})x_{t-1} + \varepsilon_t \qquad \varepsilon_t \sim N(0, \hat{\sigma}^2) \qquad t = T+1, T+2, \ldots$$

It is straightforward to adjust the procedure if c or a is fixed.

References

Acemoglu, D. and Johnson, S. (2007). Disease and development: The effect of life expectancy on economic growth. *Journal of Political Economy*, 115(6):925–985.

AG (2014). *Prognosetafel AG2014*. Actuarieel Genootschap.

Antonio, K. (2012). Sluiten van de periodetafel GBM/V 2005-2010. *Bijlage bij Prognosetafel AG2012-2062*.

Aro, H. and Pennanen, T. (2014). Stochastic modelling of mortality and financial markets. *Scandinavian Actuarial Journal*, 2014(6):483–509.

Booth, H., Hyndman, R., Tickle, L., and de Jong, P. (2006). Lee-carter mortality forecasting: a multi-country comparison of variants and extensions. *Demographic Research*, 15(9):289–310.

Booth, H. and Tickle, L. (2008). Mortality modelling and forecasting: A review of methods. *Annals of Actuarial Science*, 3(1-2):3–43.

Broer, P. (2010). Macroeconomic risks and pension returns. *CPB Memorandum*, 241.

Cairns, A. J., Blake, D., Dowd, K., Coughlan, G. D., Epstein, D., Ong, A., and Balevich, I. (2009). A quantitative comparison of stochastic mortality models using data from England and Wales and the United States. *North American Actuarial Journal*, 13(1):1–35.

Carolina, N. and Van Duin, C. (2010). Onzekerheidsmarges voor de sterfteprognose van het CBS. *Bevolkingstrends*, 2(58):32–37.

Currie, I. D., Durban, M., and Eilers, P. H. (2004). Smoothing and forecasting mortality rates. *Statistical Modelling*, 4(4):279–298.

Draper, N. (2014). A financial market model for the netherlands. *CPB Background Document*.

Durbin, J. and Koopman, S. J. (2012). *Time series analysis by state space methods*. Number 38. Oxford University Press, 2nd edition.

EC (2014). *The 2015 Ageing Report: Underlying Assumptions and Projection Methodologies*. European Commission.

Eisenhauer, J. G. (2003). Regression through the origin. *Teaching Statistics*, 25(3):76–80.

Gerland, P., Raftery, A. E., Ševčíková, H., Li, N., Gu, D., Spoorenberg, T., Alkema, L., Fosdick, B. K., Chunn, J., Lalic, N., et al. (2014). World population stabilization unlikely this century. *Science*, 346(6206):234–237.

Heij, C., De Boer, P., Franses, P. H., Kloek, T., and Van Dijk, H. K. (2004). *Econometric methods with applications in business and economics*. Oxford University Press, Oxford.

Hyndman, R. J. and Ullah, M. S. (2007). Robust forecasting of mortality and fertility rates: a functional data approach. *Computational Statistics & Data Analysis*, 51(10):4942–4956.

Kannisto, V. (1992). Development of the oldest-old mortality, 1950-1980: Evidence from 28 developed countries.

Keyfitz, N. (1981). The limits of population forecasting. *Population and Development Review*, 7(4):579–593.

Koijen, R. S., Nijman, T. E., and Werker, B. J. (2010). When can life cycle investors benefit from time-varying bond risk premia? *Review of Financial Studies*, 23(2):741–780.

Langejan, T., Gelauff, G., Nijman, T., Sleijpen, O., and Steenbeek, O. (2014). Advies commissie parameters. Technical report, SZW.

Lee, R. D. and Carter, L. R. (1992). Modeling and forecasting us mortality. *Journal of the American Statistical Association*, 87(419):659–671.

Lee, R. D. and Tuljapurkar, S. (1994). Stochastic population forecasts for the United States: Beyond high, medium, and low. *Journal of the American Statistical Association*, 89(428):1175–1189.

Li, N. and Lee, R. (2005). Coherent mortality forecasts for a group of populations: An extension of the lee-carter method. *Demography*, 42(3):575–594.

Muns, S. (2015). A financial market model for the Netherlands: A methodological refinement. *CPB Background Document*.

Niu, G. and Melenberg, B. (2014). Trends in mortality decrease and economic growth. *Demography*, 51(5):1755–1773.

OECD (2001). *Measuring Productivity: Measurement of aggregate and industry-level productivity growth*. OECD.

Poterba, J. M. (2001). Demographic structure and asset returns. *Review of Economics and Statistics*, 83(4):565–584.

Poterba, J. M. (2004). Impact of population aging on financial markets in developed countries. In *Federal Reserve Bank of Kansas City Annual Symposium, Jackson Hole, WY*.

Raftery, A. E., Alkema, L., and Gerland, P. (2014). Bayesian population projections for the United Nations. *Statistical Science*, 29(1):58.

Renshaw, A. E. and Haberman, S. (2006). A cohort-based extension to the lee–carter model for mortality reduction factors. *Insurance: Mathematics and Economics*, 38(3):556–570.

Ševčíková, H., Li, N., Kantorová, V., Gerland, P., and Raftery, A. E. (2015). Age-Specific Mortality and Fertility Rates for Probabilistic Population Projections. *ArXiv e-prints*.

UN (2015). *World Population Prospects: 2015 Revision*. United Nations.

Van Duin, C. and Stoeldraijer, L. (2014). Bevolkingsprognose 2014-2060: Groei door migratie. *Centraal Bureau voor de Statistiek*.

Publisher:

CPB Netherlands Bureau for Economic Policy Analysis
P.O. Box 80510 | 2508 GM The Hague
T +31 70 3383 380

info@cpb.nl | www.cpb.nl

October 2015