Response to Updated Mortality Forecasts in Life Cycle Saving and Labor Supply

Niku Määttänen and Juha Alho

December 2013

Abstract

Historical evidence shows that demographic forecasts, including those of mortality, have often been grossly in error. One consequence is that forecasts are frequently updated. How should individuals or institutions react to updates, as these are, in turn, expected to be uncertain? We discuss this problem in the context of a life cycle saving and labor supply problem, in which a cohort of workers decides how much to work and how much to save for mutual pensions. Mortality is stochastic and point forecasts are updated regularly. A Markovian approximation for the predictive distribution of mortality is derived. This renders the model computationally tractable, and allows us to compare a theoretically optimal rational expectations solution to a strategy in which the cohort merely updates the life cycle plan to match each updated mortality forecast. The implications of the analyses for overlapping generations modeling of pension systems are pointed out.

1 Introduction

Demographic forecasts, including mortality forecasts, have historically been more uncertain than has been generally believed (e.g. Alho (1990)). Yet usually only point forecasts are discussed in the media. Even macroeconomic analyses that purport to address questions of sustainability of policies have often been carried out under a single population scenario.

During the past two decades attempts have been made to include stochastic demographics into generational accounting that seeks to determine if current tax and entitlement rules lead to the balancing out of public finances in the long run (e.g. Alho and Vanne (2006)), and into overlapping generations models (Samuelson (1958)), whose computational versions can provide realistic representations for national economies that include a pension system (cf. articles in Alho et al. (2008)).

Overlapping generations models involve individual or household optimization with respect to savings and labor supply. The decisions are dynamic in nature, and they depend on expectations about future demographics. A key issue is how exactly those expectations are assumed to be formed when mortality is stochastic.

In Alho et al. (2008) it was assumed, for technical simplicity, that the decision maker knows future demographics without error. In such a set-up, the decision maker formulates a life cycle saving and labor supply plan for each sample path of mortality and never revises it.

More recently, Lassila et al. (2012) applied, apparently for the first time, what they call "the updated point forecast" approach to the overlapping generations computations. In this approach, the decision maker learns the most recent point forecast every period, assumes that it is without error, and optimizes accordingly. This is a major step forward, since the perfect foresight assumption is of course problematic in the context of stochastic demographics.

In the updated point forecast problem, the feasibility of the numerical calculations hinges on the assumption that the cohort behaves as if the updated forecast were without error. One can argue that this is probably the way most societies react to changing forecasts (cf. Alho (2013)). However, it is clear that this is not the theoretically optimal approach, since one can always do at least as well by taking into account the possibility of future updates in the optimization. The latter is the so-called rational expectations approach that was first proposed by Muth (1961). However, in general, the consideration of stochastic demographics with fully rational expectations leads to very complicated, if not intractable, computations (caused by the "curse of dimensionality"). [1]

We will compare life cycle savings and labor supply decisions under different informational assumptions. In particular, we will compare decisions made under rational expectations about future mortality with those made using updated point forecasts. In the rational expectations problem, the decision maker is assumed to know exactly the conditional distribution of future mortality, given the current mortality. This information is provided to her in the form of a stochastic forecast. Since our model is Markovian, and since the state space is not overly large, we can overcome the computational problems related to rational expectations (cf., Filar and Vrieze (1997)).

The first question we aim to answer is to what extent the updated point forecast problem can be interpreted as providing a good approximation of the rational expectations solution.[2] We will consider how much the consumer loses, in welfare terms, by considering only point forecasts instead of stochastic forecasts with rational expectations. In addition, we will be able to compare the optimal savings and labor supply decisions in the two cases.

A related question of independent interest is, should individual consumers be interested in demographic uncertainty. If the answer to this question is positive, the public should be provided stochastic demographic forecasts instead of the usual point forecasts.

In Alho and Määttänen (2008) we considered a life cycle savings problem in which mortality was stochastic and followed a very simple Markovian model that was essentially a discrete approximation to the log-bilinear model of Lee and Carter (1992). In this paper, we provide a Markovian approximation to a more general mortality model that has been estimated for Finnish females as described in Alho and Spencer (2005). We then expand the earlier economic setting to include not only saving but also supply of labor. The assumption will be that individuals enjoy both leisure and consumption and the two have to be balanced to maximize utility. The analysis is carried out for a cohort that has a funded pension system, or a *tontine*, in which the cohort's savings accrue interest. Cohort survival is assumed to follow exactly the survival probabilities obtained from the Markovian approximation. Or, the model includes the so-called aggregate uncertainty in mortality, but the idiosyncratic uncertainty that can be important in actual pension schemes is left out. Therefore, the model includes the most important elements used in computational overlapping generations models, such as that of Lassila et al. (1997).

[1]See Hasanhodzic and Kotlikoff (2013) for a discussion about the curse of dimensionality in overlapping generations models.

[2]This question relates our paper to a literature on boundedly rational decisions. For a recent example, see for instance Winter et al. (2012)

We will start by discussing the cohort's decision problem in terms of a utility function, and give the full rational expectations formulation of the work/saving problem. Then, we provide the formulation of the decision problem that assumes updated forecasts in a Markovian setting. This is followed by a description of how, in general, one may obtain a Markovian approximation to a multidimensional predictive distribution. The results on expected life time utility are given in terms of welfare equivalents and compared to those that have been obtained earlier in a simpler setting. We conclude by considering the implications of our findings on the application of computational overlapping generations calculations.

2 Decisions on Work and Savings and a Pension System

2.1 Utility Considerations

Consider a cohort of homogeneous individuals who enter working age at exact time $t = 0$ and live during $t = 0, 1, ..., T$. As of $s > 0$, the individuals are retired. During $t = 0, 1, \ldots, s-1$ the individuals have a time endowment of one unit that is split between work and leisure. In the beginning of each period t the members of the cohort decide how much to work, $0 \leq w_t \leq 1$, and how much to consume, $c_t \geq 0$. The periodic utility is a function of consumption and time worked, $u_t = u(c_t, w_t)$.

We assume that both leisure and consumption are necessary for utility to be created, and the periodic utility function is of the form,

$$u(c, w) = ((c^{\alpha}(1-w)^{\beta})^{1-\lambda} - 1)/(1-\lambda), \tag{1}$$

where $\alpha > 0, \beta > 0, \lambda > 0$ and $\lambda \neq 1$. For $\lambda = 1$ we have $u(c, w) = log(c^{\alpha}(1-w)^{\beta})$. Following definitions in Arrow (1971) this utility function is often considered to represent constant relative risk aversion. I.e., given the other parameters of the utility function, λ determines the degree of aversion to relative risk.[3] Parameters α and β, in turn, determine the relative importance of consumption and leisure, respectively. Preferences are assumed to be time-separable, i.e., the utilities of different periods are additive.

Consumption is financed by wages that are proportional to the time worked. Therefore, there is no limitation in assuming that the total wage in a period is equal to the time worked w. If there were no saving for the

[3]Statisticians will notice that the functional form of the utility function is that of the so-called Box-Cox transformation (Box and Cox (1964)).

future, this would also be the consumption, or $c = w$. In this case the utility can be re-expressed (with some abuse of notation) as $u = ((w^{\alpha}(1-w)^{\beta})^{1-\lambda} - 1)/(1-\lambda)$. This is maximized at $w* = \alpha/(\alpha+\beta)$. But, during retirement ages $w_t = 0$, so $u_t = u(c_t, 0)$, where the consumption must be financed from earlier savings.

Define A_t as the total savings remaining in the beginning of the period t. They and any net savings during period t accrue interest at the rate r. For definiteness, we assume that the time worked is decided in the beginning of the period, the wage is immediately paid and the consumption occurs without delay. This leads to the cash flow identity

$$A_{t+1} = (1+r)(A_t + w_t - c_t), \tag{2}$$

where $A_t + w_t \geq c_t$ so there is no borrowing for consumption; $w_t = 0$ in retirement ages; and $A_0 = 0$.

2.2 Markovian Mortality and Pension Scheme

Only a single cohort will be considered, so we can identify age with time. The highest age is T, so that by exact age $T+1$ everybody dies. Deaths are assumed to occur at the end of each period. We assume that mortality is a stochastic process that is driven by a (non-stationary) Markov chain with a finite number of states $d = 1, \ldots, D$. The states reflect the level of mortality. Specifically, the probability of surviving from exact age t to exact age $t+1$ is $S_t(d)$ if the state of the chain is d. The mortality state of period $t = 0$ is chosen randomly with probabilities $0 < \pi_d < 1$ such that $\pi_1 + \ldots + \pi_D = 1$. Thereafter, $P_t(d, e)$ is the conditional probability that the state during age $t = 1, \ldots, T$ is e, given that the state during period $t-1$ was d. Since nobody survives to age $T+1$, we take $S_T(d) = 0$ for all d.

To match our simple cohort setting to the overlapping generations calculations, we now introduce a simple pension scheme that is financed by the savings. We assume that a cohort insures itself against mortality risk by pooling savings. Thus, the savings contributed by the deceased are redistributed to those still alive.

We will ignore the complications arising from the finiteness of any such cohort in real life, and assume that the cohort depletion exactly follows the Markovian mortality process. It follows that the cash flow identity above will be modified to the form

$$A_{t+1} = (1+r)(A_t + w_t - c_t)/S_t(e) \tag{3}$$

when mortality in age/period t happens to be $e = 1, \ldots, D$ and the pension system exists.

3 An Individual's Use of Forecasts in Decision Making

We will now consider the decision problem from the perspective of a single member of the cohort. Her work/saving plan from any year $t > 0$ onwards conditions on her being alive at t, and the state of the mortality process at t is assumed to be known. (In reality, for a finite cohort the state can only be estimated with some error.)

3.1 Rational Expectations Problem

In the rational expectations problem the decision maker is assumed to know exactly the conditional distibution of future mortality, given the current mortality. Let the discount factor be $\phi > 0$ and let $V_t = V_t(A, d)$ be the total discounted expected utility from exact age $t = 1, \ldots, T$ onwards given that the total savings at t are A and mortality during period $t-1$ was d. Then, because preferences are assumed to be time-separable, we have the recursion

$$V_t(A_t, d) = max_{c,w}\{u(c, w) + \phi \sum_{e=1}^{D} P_t(d, e)S_t(e)V_{t+1}(A_{t+1}, e)\}, \qquad (4)$$

where the future savings A_{t+1} are determined via the cash flow identity (3) for each value of e. To define the expected utility at $t = 0$, replace the transition probability $P_t(d, e)$ by π_e, on the right hand side of (4).

3.2 Updating Forecasts Problem

We compare the above rational expectations problem to two alternatives in which the decision maker has less information. (a) She only learns the initial point forecast for future mortality at $t = 0$. In each period she makes a life cycle saving and labor supply plan for the rest of her life. She is not given any further information about the evolution of mortality. However, she does update her saving and labor supply plan periodically, because the return from the tontine reflects the actual level of realized mortality and is typically different from what she expected. (b) She learns the most recent point forecast every period and updates her future consumption and labor supply plan accordingly, given that she is alive at t. In both these cases, the decision maker does not take into account that she may later change her mind.

Unlike in the case with rational expectations, the work/saving problem can now be solved in a forward looking manner, without resorting to recursive methods.

Consider a time $y \geq 0$ and let $S_{t|y,d}$ be the *point forecast* for the one-year survival probability at a future age/time $t \geq y$ given that at exact time/age y mortality was in state d. To maximize her future welfare with savings A_y at $t = y$ she needs to maximize the future discounted expected utility,

$$\sum_{t=y}^{T} \phi^{t-y} (\prod_{k=y}^{t-1} S_{k|y,d}) u(c_t, w_t) \tag{5}$$

with respect to $(c_t, w_t), y \leq t \leq T$, subject to

$$A_{t+1} = (1+r)(A_t + w_t - c_t)/S_{t|y,d}, \tag{6}$$

where $A_t + w_t \geq c_t$, and d and A_y are known at $t = y$. Here period t utility is weighted by the probability of being alive in period t conditional on being alive in period y. The maximization amounts to solving a system of non-linear equations that consists of the consumer first-order conditions and the budget constraints. [4]

Consider a consumer who uses updated point forecasts. In order to determine her life cycle consumption and labor supply path for a given sequence of mortality states d_t for ages $t = 0, \ldots, T-1$, we proceed as follows. The first state is given by probabilities π_d. Given that $d_0 = d$, the point forecast for age-specific survival probabilities $S_{t|0,d}$ at $t \geq 1$ is, by the Chapman-Kolmogorov equations (e.g., Çinlar (1975), p. 110), given by element d of the column vector

$$(\prod_{k=1}^{t} \mathbf{P}_k)\mathbf{S}_t, \tag{7}$$

where $\mathbf{P}_k = (P_k(d,e))$ is the matrix of transition probabilities for period k, and $\mathbf{S}_t = (S_t(1), \ldots, S_t(D))^T$ is the vector of one-year survival probabilities for period t under the different states of the mortality process. Given the point forecast and $A_0 = 0$, we can solve the deterministic work/saving problem described above. The result consists of consumption and labor supply for all periods. However, we only store the first period allocation. Given c_0,

[4] For reasonable parameter values, the no borrowing constraint is binding only in the last period. Before retirement, the consumer needs to save in order to have something to consume during retirement. After retirement, before the last period, if the consumer does not save anything for next period, she would not be able to afford any consumption in the following period. That cannot be optimal since the marginal utility of consumption goes to infinity as consumption goes to zero.

n_0, and d_1, we determine A_1 as well as the new point forecast. Then, we solve again the optimization problem from $t = 1$ onwards, which in turn gives us c_1 and n_1. This process is repeated until the last period.

The case in which the consumer uses only the initial point forecast is analogous. In that case, however, the need to revise the life cycle plan stems solely from surprises in the return to savings. In the absence of the tontine (that is, with budget constraint (2)), the consumer using only the initial point forecast would always find it optimal to stick to her initial life cycle plan.

The primary problem we wish to answer is how much a decision maker of type (a), and especially of type (b), is expected to loose utility as compared to a decision maker who has access to the full rational expectations solution.

3.3 Intuition from a Three-Period Model

In order to understand when and how the uncertainty about mortality matters for the consumer, it is useful to consider a three-period version of the model. In this case $T = 2$, and the first decisions relate to time $t = 0$. We assume here that the consumer makes a labor supply decision at $t = 0, 1$ and is retired during $t = 2$.

Suppose first that the tontine is not available. Let us first take the decisions at $t = 0$ as given and consider the consumer's problem at $t = 1$. Given rational expectations, and taking into account that the consumer is retired during $t = 2$, we can write the consumer's problem at $t = 1$ as

$$\max_{c_1, w_1} \left\{ u(c_1, w_1) + \phi \sum_{e=1}^{D} P_1(d, e) S_1(e) u(c_2, 0)) \right\}, \tag{8}$$

where the budget constraint is $c_2 = (1 + r)(A_1 + w_1 - c_1)$. Defining $\overline{S}_1 = \sum_{e=1}^{D} P_1(d_1, e) S_1(e)$ as the expected survival probability from age $t = 1$ to $t = 2$, the second term in brackets can be rewritten as $\phi \overline{S} u(c_2, 0)$. In other words, given current savings, an individual who makes her decision based on expected survival makes exactly the same decision at $t = 1$ as an individual who forms rational expectations (i.e. knows the values $P_1(d, e)$ and $S_1(e)$ for $e = 1 \dots D$) based on the stochastic forecast.

However, at $t = 0$ the situation is different. Since $\overline{S}_1$ depends on the state d of the mortality process during the period $t = 0$, the optimal values for c_1 and w_1 in (8) depend on it as well. For instance, one can show that if d is such that the expected survival rate is low, the consumer will choose to save relatively little. Thus, different realizations of the mortality process can lead to different levels of consumption and hours worked, and an individual who takes demographic uncertainty properly into account achieves higher expected welfare than an individual using only a point forecast.

When a tontine is part of the model, there is a more direct reason why uncertainty about aggregate mortality matters. Consider again the problem in period 1 but assuming that consumption in period 2 is determined as $c_2^e = (1+r)(A_1 + w_1 - c_1)/S_1(e)$. Since consumption in period $t = 2$ now depends on the state of mortality e during $t = 1$, the expected second period utility $\phi \sum_{e=1}^{D} P_1(d,e) S_1(e) u(c_2^e, 0))$ cannot be written using the expected survival probability only. Hence, the individual would benefit from taking the demographic uncertainty into account even in period 1.

4 Construction of the Markovian Approximation to the Predictive Distribution of Mortality

The starting point of our empirical illustration is a numerical construction of a Markovian approximation to the predictive distribution of age-specific mortality for females in Finland, with jump-off time January 1, 2009, for years 2010-2090. The lead time of the forecast is 80 years, and the ages considered are 25-104. Single year data are combined to produce survival probabilities for five-year periods corresponding to ages $25-29, \ldots, 100-104$. Thus, there are 16 ages/periods in all.

The point forecast is based on an official forecast produced by Statistics Finland, but both age and lead time have been extended by extrapolation. This is in keeping with the way the original forecast was made. The uncertainty estimates were produced as discussed in Alho and Spencer (2005), i.e., they reflect the historical error in extrapolation estimates in past age and sex-specific mortality data for ten European countries during, roughly, the 20th century, and the so-called scaled model for error was assumed.

The calculations were as follows. First, 300,000 samples from the 80×80 dimensional predictive distribution were generated and single year survival probabilities were computed. Second, five-year survival probabilities were obtained for the cohort in question by multiplying appropriate single-year survival probabilities along a diagonal of the Lexis diagram. This produced 300,000 samples from the predictive distribution of each of the 16 survival probabilities. Third, for each five-year period, the distribution was split into deciles $d = 1, \ldots, D$, with $D = 10$, so for the the first age/period we have the probabilities $\pi_d = 1/10$. Fourth, for each age/period $t = 0, \ldots, T$ with and decile d, the average survival probability $S_t(d)$ was estimated. Fifth, for $t = 1, \ldots, 15$ transition probabilities $P_t(d, e)$ were estimated from the observed transition frequencies.

5 Calibration of the Model for Labor Supply and Saving

In order to analyze the model numerically, we still need to choose meaningful values for the economic parameters $\alpha, \beta, \lambda, r$ and ϕ. In addition, we need to specify the retirement period.

We set the retirement period at $s = 8$, which corresponds to real age 65-69. We set the interest rate at $r = 0.10$. Since unit time in our model corresponds to five years, this implies an annual real interest rate of about 2%.

The parameters α, β and λ are functionally dependent in the sense that substitutions $\alpha := c\alpha$, $\beta := c\beta$, $(1-\lambda) := (1-\lambda)/c$ for any $c > 0$ only change the value of the utility function by a multiplicative factor of c. Since this does not influence optimality, we set $\alpha = 1 - \beta$, where $0 < \beta < 1$.

We are thus left with three parameters, β, λ and ϕ. The parameter β determines the utility weight of leisure relative to consumption. Given β, the parameter λ determines both the intertemporal elasticity of substitution, which measures the extent to which an increase in the real interest rate induces consumers to substitute future consumption for current consumption, and the degree of risk aversion towards consumption fluctuations (Hall (1988)). One can show that the intertemporal elasticity of substitution in consumption is given by $1/(1 + (1 - \beta)(\lambda - 1))$, or it is inversely related to λ. The degree of risk-aversion is in turn inversely related to the elasticity of substitution in consumption. Hence, risk aversion increases with λ.

The value of λ is usually determined based on empirical results about the intertemporal elasticity of substitution in consumption. However, as the empirical estimates vary widely, we will consider a wide range of values, $\lambda = 1, 5, 10, 15$. For each value of λ, we choose β and ϕ so that the time allocation between leisure and work and the amount of retirement savings are realistic in the model. As a result, in the parametrizations we consider, the intertemporal elasticity of substitution varies from 1 at $\lambda = 1$ to 0.21 at $\lambda = 15$. This range covers values typically used in related models and also most of the empirical estimates reviewed in Auerbach and Kotlikoff (1987).

Regarding time allocation, we want individuals in working age to devote 1/3 of their time endowment to work.

As for savings, we base our calibration on the Wealth Survey that was conducted by Statistics Finland in 2009. The average annual earnings of Finnish households of age 25-64 was 42,000 €. (Household age is defined as the age of the highest-earning member of the household.) The average private net wealth of households in age group 60-64, which corresponds to the last working period in the model ($t = 7$), is 230,000 €. In addition, households have pension entitlements. In the mandatory earnings related pension system the ratio of average pension to average wage income is around 60% in Finland. Given the average annual income of 42,000 €, this implies an average annual pension of approximately 25,000 €. Given the assumed interest rate $r = 0.10$ and the mortality process, the present value of an annual pension of 25,000 starting at age 65 is approximately 550,000 €. Hence, total household wealth in age group 60-64 is 780,000 €. Thus, the ratio of the average wealth in age group 60-64 and average wage income during a five year period is $780,000/(5 \times 42,000) \approx 3.7$. In terms of the model parameters, the average wage income is 1/3 (recall that the wage rate and the time endowment were normalized to 1), so the average savings at $t = 7$ should be about 1.2.

The required utility weight of leisure is (roughly) the same for all values of λ we considered. Taking $\beta = 0.73$, individuals below retirement age devote very close to 1/3 of their time endowment to work.

The discount factor ϕ, in contrast, depends markedly on λ, because λ also affects the degree of substitutability between consumption and leisure. It turns out that when λ is large, individuals choose to lower consumption abruptly at retirement because increased leisure is a good substitute for consumption. As a result, increasing λ but keeping ϕ would lead to a decrease in savings. To keep the savings at target, we need to increase ϕ with λ. The result is that for $\lambda = 1, 5, 10, 15$, the calibrated discount factors are $\phi = 0.94, 1.14, 1.45, 1.75$, respectively. [5]

Figure 1 displays the average per capita consumption and saving by age for the top quintile of life expectancies (dashed lines) and bottom quintile of life expectancies (solid lines), as computed from 1000 simulated paths. The average lifespans are 14.5 and 12.9 model periods, respectively. The calibration corresponds to $\lambda = 5$.[6]

---

[5] Admittedly, discount factors of 1.45 and 1.75 are very high. These high values reflect the fact that we want to compare different informational assumptions also under very high risk aversion. Assuming a utility function that is separable between consumption and leisure would weaken the link between retirement savings and risk aversion. However, the utility function considered here is very common in the related literature.

[6] Curves for labor supply are almost identical for the two quintiles, and are not displayed. The curves are downward sloping for model ages 0-7.

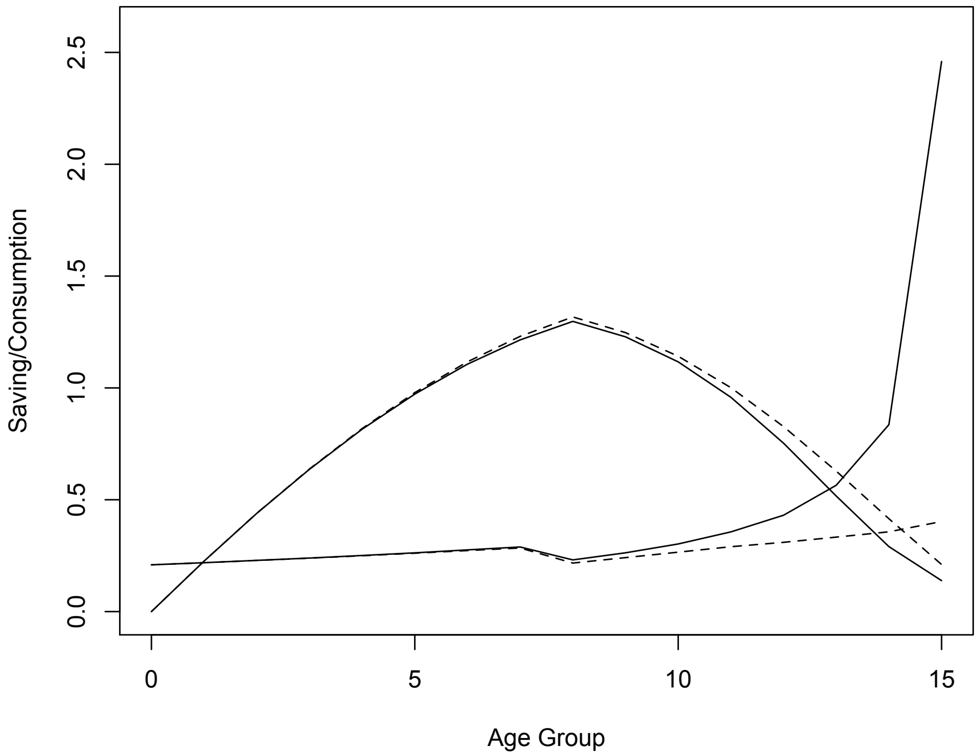

Figure 1: Average life cycle Savings and Consumption, for low mortality (dashed) and high mortality (solid). Curves for Savings start from zero.

The savings profiles reflect standard life cycle consumption smoothing. Consumers accumulate savings during working life in order to be able to finance consumption during retirement. The upward trend in the consumption profiles relates to the fact that the calibrated discount factor is relatively high. For e.g. $\lambda = 1$, the discount factor would be smaller and the consumption profiles downward sloping, except for the very last periods. The fall in consumption from $t = 7$ to $t = 8$ relates to the increase in leisure that is associated with retirement.

There is not much difference between the average profiles for low and high mortality paths until model age 9 or so. This reflects the fact that updates in mortality forecasts reveal information about the cohort's average lifespan only gradually. The differences become significant in older ages. A high average lifespan implies higher average savings and lower consumption in all ages. Partly because of differences in the return from the tontine between high and low mortality paths, differences in per capita consumption are large in very old ages. In the very last period, there is no longer uncertainty about the remaining lifetime and in the absence of a bequest motive it is optimal to consume all savings. However, the share of the cohort that survives to the last period, which corresponds to real ages 100-104, is quite small on all simulated mortality paths.

6 Effect of Informational Assumptions on Welfare and Allocation

We first analyze how important it is for the consumer to take demographic uncertainty into account, in terms of expected lifetime utility.

Our welfare measure is the so-called consumption equivalent variation. It gives the constant percentage change in periodic consumption (in all periods) that is needed to make the expected lifetime welfare in the comparison case as high as in a benchmark case. For instance, in order to compute the welfare cost of not taking aggregate mortality risk into account, we first generate a large number of aggregate mortality paths that are all consistent with $\{S_t, P_t \mid t = 1, \ldots, T\}$. We then find the optimal life cycle consumption and labor supply paths for each simulated mortality path under different informational assumptions. Finally, we ask how much we should increase the consumption of a low information consumer (in every period and in all mortality paths) in order to make her average lifetime utility across all mortality paths the same as that of a decision maker with rational expectations.

Table 1 compares expected lifetime utilities associated with the initial point forecast, updated point forecasts, and perfect foresight with the expected lifetime utility associated with rational expectations. The welfare cost of using only the initial point forecast, instead of having rational expectation based on stochastic forecasts, increases rapidly with λ and ranges from 0.3% to 7.8% in terms of consumption. Thus, depending on λ the welfare loss is potentially very large.

Interestingly, the welfare loss diminishes drastically when the individual uses updated point forecasts. Indeed, for moderate values $\lambda = 1$ and $\lambda = 5$, the welfare cost of using updated point forecasts instead of stochastic forecasts is less than 0.001%. In these cases, a consumer who makes her savings and labor supply decisions based on updated point forecasts reaches virtually the same expected lifetime utility as a consumer who is able to form

Table 1: Relative welfare cost (%) as compared to rational expectations, as a function of λ.

λ	1	5	10	15
Initial point forecast	0.3	0.8	2.8	7.8
Updated point forecasts	0.0	0.0	0.3	0.7
Perfect foresight	-0.3	-0.8	-1.2	-1.5

rational expectations based on stochastic forecasts.

Obviously, a consumer that has perfect foresight does even better than a consumer with rational expectations, as indicated by a negative welfare cost in the last row of Table 1. In the model, having perfect foresight improves expected lifetime welfare between 0.3% to 1.5% relative to rational expectations that are based on stochastic forecasts. It is interesting to note that this welfare cost increases much slower with λ than the one associated with using only the initial point forecast. Intuitively, an individual that relies on the initial point forecast only, may have a very biased view about mortality through her entire life. In contrast, the beliefs of an individual using updated forecasts (either stochastic forecasts or point forecasts), converges towards the perfect foresight information as she gets older.

Why is it the case that updated point forecasts suffice? In other words, why is it not necessary to take uncertainty into account in this set-up? Intuitively, there is not much that an individual can do to mitigate the effects of the uncertainty about future mortality. All she can can do is to increase savings so as to increase consumption in cases where average life time is exceptionally long. That, however, comes at the cost of having lower consumption earlier in life. Moreover, the uncertainty about average lifetime diminishes over time and the individual has time to adjust her savings plans as new information arrives.

In the simpler setting of Alho and Määttänen (2008), smaller welfare losses were seen also when comparing the use of initial point forecast and stochastic forecasts. One reason why some of the welfare losses are larger in the current setting is the presence of the tontine. When we eliminate the tontine, especially the welfare losses related to the use of the initial point forecast relative to rational expectations (results not shown) become much smaller than those presented above. As described in Section 3.3, an individual that uses only the initial point forecast in the presence of the tontine may have a very poor estimate not just about the survival probability, but also about the future return to her savings.

It should be carefully noted that these results are relative, given the economic setting. For example, the availability of the tontine does improve welfare, in a major way. Our calculations show that in the case of rational expectations and $\lambda = 5$, for instance, the average welfare is 9 % higher with the tontine than without it. This is in line with the literature on annuities (e.g. Brown (2001)). The tontine allows the consumers to avoid leaving 'accidental bequests'. At the same time, it helps them smooth consumption over time by effectively increasing the return on savings at old age. The increase in welfare associated with the tontine is very similar in magnitude for all informational assumptions.

It is possible that even if expected lifetime utility were almost the same with rational expectations and updated point forecasts, households decisions could be very different for some mortality paths. Figure 2 compares household savings under rational expectations to household savings under updated point forecasts, single point forecast, and perfect forecast among the simulated mortality paths. Here we are again assuming $\lambda = 5$. The comparison is between savings in each age along each path. The figure displays (a Gaussian kernel estimate of) the density of the distribution of differences in savings relative to the average savings in the rational expectations case. A positive difference means savings are higher with rational expectations.

The solid density in Figure 2 reveals that a consumer with rational expectations has *higher* savings in each age than a consumer who bases her decisions on the updated forecast only.[7] On the other hand, the latter decisions track quite closely those based on rational expectations. Most of the differences are less than 10% of average savings. These findings suggest that at least for moderate values of the risk-aversion parameter, solving a consumer's problem based on updated point forecasts provides a reasonable proxy for the the rational expectations solution even on a path by path basis.

Decisions based on the initial point forecast are characterized by the dashed density. Negative values, i.e. inadequate saving is fairly frequent, and, for some paths the values are well outside the range shown the figure. In particular, in some (rare) cases, a consumer who relies on the initial point forecast alone, saves far too much for the future, when the actual mortality is exceptionally high.

These results are complemented by the perfect forecast solution given by the dotted density. A decision maker who knows the mortality path she is on (although not her own lifetime) saves half of the time more and half the time less than the one who is "only" able to form rational expectations. Again,

[7] This is true of *all* simulated data points, which is not obvious from the smooth kernel estimate.

the largest differences are observed at very old age. The welfare implications (in terms of expected lifetime welfare) of these mistakes are not as large as one might infer from the figure, since only a small fraction of the cohort is alive in those periods.

Figure 3 compares labor supply under rational expectations and alternative informational assumptions. Since labor supply is assumed to be always zero from model age 8 onwards, the comparison is between labor supply in each model age from 0 to 7. Again, the comparison is made along each simulated path. The solid curve is the density of the relative difference in labor supply under rational expectations and updated point forecast. The largest differences are less than 2%. On the other hand, a consumer with rational expectations works in almost all cases more than an individual who bases her decision on the updated point forecast only. I.e., under rational expectations individuals both work and save more than those who only consider point forecasts, early in life. Naturally, savings and labour supply decisions are interrelated. A higher labor income allows consumers to save more.

In the case of labor supply, even the solution based on the initial point forecast is usually quite close to the rational expectations solution. This is partly because labor supply decisions are made relatively early in life when the most recent forecast is unlikely to differ much from the forecast that was made in the beginning of the cohort's working life. Again, differences between the rational expectations solution and the perfect foresight solution are on average much higher. For instance, if consumers knew already during working life that their future mortality will be very low, they would choose to work a lot.

Our model is a partial equilibrium one in the sense that the interest and the wage rate are exogenously fixed. In principle it is possible that the welfare results would be very different in a general equilibrium version of the model where the wage and the interest rate are determined endogenously. A standard way to endogenize the wage and the interest rate would be to assume that they are determined as marginal productivities of aggregate labor supply and capital stock which corresponds to aggregate savings. (Naturally, this also requires postulating a production function.) However, as shown above, both the per capita labor supply and savings that are associated with either the rational expectations solutions or the updated point forecast solution are always relatively close to each other. Hence, the dynamics of aggregate labor supply and capital stock would be very similar in both cases. This suggests that general equilibrium effects are unlikely to change the welfare results concerning the updated point forecast solution vis-a-vis rational expectations substantially.

7 Discussion

We have compared the relevance of having available a full stochastic forecast for mortality as opposed to having merely a point forecast and subsequent updates, from the point of view of an individual that needs to make life cycle saving and labor supply decisions.

Our results suggest that the individuals need to revise their work and saving plans as new information regarding longevity becomes available. Failing to use updated forecasts about future mortality leads to non-trivial welfare losses.

On the other hand, it appears that at least in our setting the individuals need not fully understand the uncertainty related to mortality, as long as frequent forecast updates are available. The welfare loss that is related to using only the updated point forecasts, rather than the stochastic forecast, is small. Moreover, the optimal savings and labor supply paths that stem from using updated forecasts appear to be close to those optimally chosen based on a full stochastic forecast.

One implication of our results is that the solution that is based on updated point forecasts may usually be considered as a good proxy for rational expectations solution in similar life cycle problems. This is of practical interest since, as discussed in the Introduction, incorporating demographic uncertainty into general equilibrium OLG models with rational expectations leads to very difficult computational challenges. Assuming that consumers make their decisions based on updated point forecasts instead of rational expectations simplify the computations drastically because one can then solve the model separately for each simulated demographic path as a sequence of perfect foresight problems.

The presence of a tontine is highly beneficial for the decision maker. Thus, from an individual's point of view it seems relevant that those designing such schemes understand the risk characteristics of the various processes involved.

Another issue is that the economic environment in our model is resilient to changes in mortality. Moderate changes in mortality do not lead to abrupt changes in individual budget constraints. This is because the tontine we considered is essentially a fully funded pension system. Non-funded pay-as-you-go pension schemes, in contrast, are likely be more vulnerable to demographic changes - although they may have other benefits not discussed here, such as protection against inflation. Even in non-funded system the mortality risks can probably be mitigated by having transparent rules that specify how the benefits and contributions are adjusted with changes in demographics. This is an other example in which stochastic demographic forecasts can be useful.

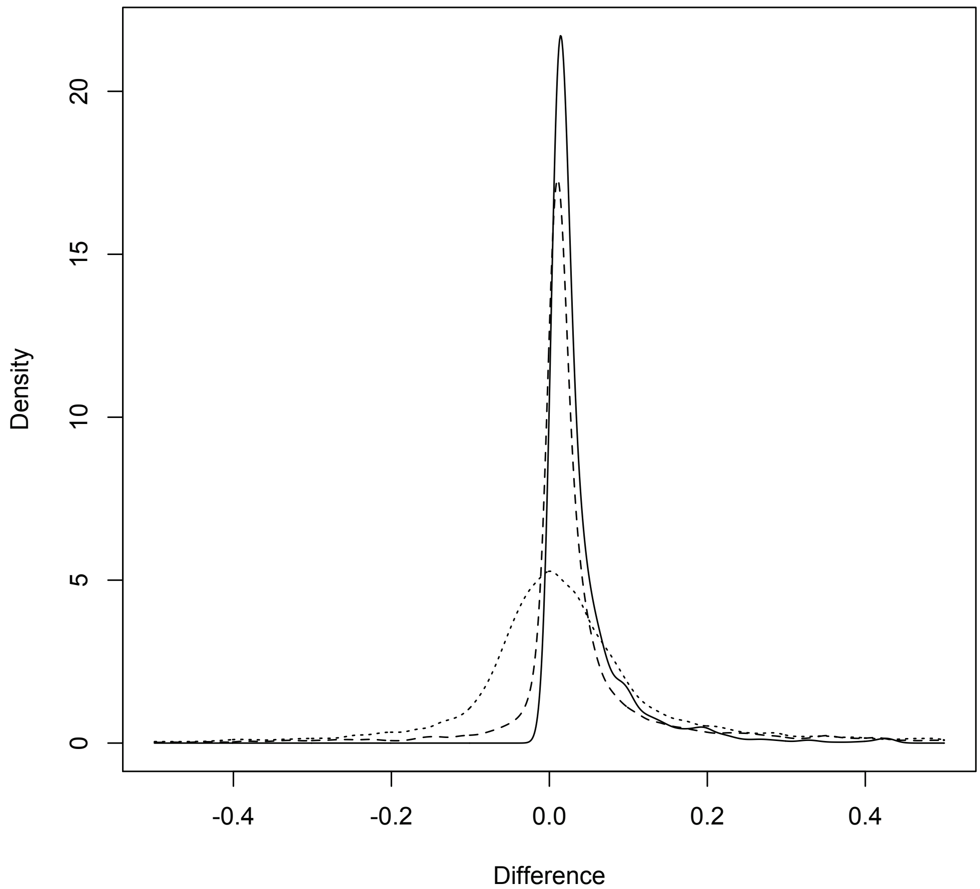

Figure 2: Relative difference in savings under rational expectations vs. updated forecasts (solid), use of a single forecast (dashed), and perfect foresight (dotted), as a proportion of savings under rational expectations.

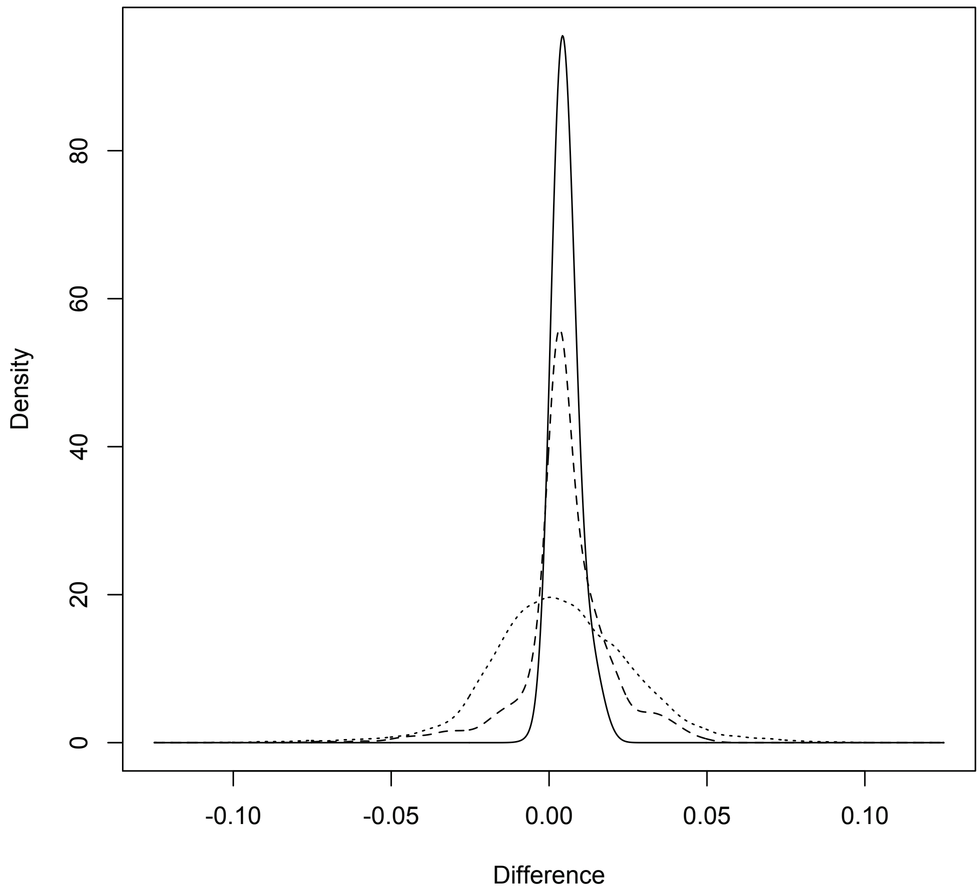

Figure 3: Relative difference in labor supply under rational expectations vs. updated forecasts (solid), use of a single forecast (dashed), and perfect foresight (dotted), as a proportion of labor supply under rational expectations.

References

Alho J.M. (1990) Stochastic methods in population forecasting. *International Journal of Forecasting*, 6, 521-530.

Alho J.M. (2013) Forecasting demographic forecasts. Submitted.

Alho J.M., Hougaard Jensen S.E. and Lassila J. (Eds.)(2008) *Uncertain Demographics and Fiscal Sustainability.* Cambridge: Cambridge University Press.

Alho J.M. and Määttänen N. (2008) Informational assumptions, aggregate mortality risk and life-cycle saving, pp. 219-238 in Alho J.M., Hougaard Jensen S.E. and Lassila J. (Eds.)(2008).

Alho J.M. and Spencer B.D. (2005) *Statistical Demography and Forecasting.* Nerw York: Springer.

Alho J.M. and Vanne R. (2006) On predictive distributions of public net liabilities. *International Journal of Forecasting*, 22, 725-733.

Arrow K.J. (1971) *Lectures on decision making under uncertainty.* Helsinki: ETLA.

Auerbach A.J. and Kotlikoff J.K. (1987) *Dynamic Fiscal Policy.* Cambridge: Cambridge University Press.

Box G.E.P. and Cox D.R. (1964) An analysis of transformations. *Journal of the Royal Statistical Society, Series B*, 211-252.

Brown J. (2001) Private Pensions, Mortality Risk, and the Decision to Annuitize. *Journal of Public Economics*, 82, 29-62.

Çinlar E. (1975) *An Introduction to Stochastic Processes.* Englewood Cliffs: Prentice-Hall.

Filar J. and Vrieze K. (1997) *Competitive Markov decision processes.* New York: Springer.

Hall, R.E. (1988) Intertemporal substitution in consumption. *Journal of Political Economy*, 96, 339-357.

Hasanhodzic J. and Kotlikoff L. (2013) Generational Risk - Is It a Big Deal?: Simulating an 80-Period OLG Model with Aggregate Shocks. *NBER Working paper* 19179.

Lassila J., Palm H. and Valkonen T. (1997) FOG: The Finnish Overlapping Generations model. *ETLA Discussion Papers* 601. Helsinki.

Lassila J., Valkonen T. and Alho J. (2012) Fiscal sustainability and policy rules under changing demographic forecasts. *ETLA Discussion Papers* 1265. Helsinki.

Lee R.D. and Carter L. (1992) Modeling and forecasting the time-series of U.S. mortality. *Journal of the American Statistical Association*, 87, 659-671.

Muth J.F. (1961) Rational expectations and the theory of price movements. *Econometrica*, 29, 315-335.

Samuelson P. (1958) An exact consumption-loan model of interest with or withour the social contrivance of money. *Journal of Political Economy*, 66, 467-482.

Winter J.,Schalfmann K. and Rodepeter R. (2012) Rules of thumb in life-cycle savings decisions. *Economic Journal*, 122, 479-501.