14. Informational assumptions and aggregate mortality risk in a life cycle savings problem

Juha M. Alho and Niku Määttänen

Introduction

The primary motivation of models discussed in this volume is the need to quantify the effects of policy measures on e.g. the pension system. Stylized models are preferred for qualitative insight (Diamond 2001), but cannot provide precise estimates for policy formulation. The price one has to pay for the realism is the relative complexity of the models. The models are not analytically tractable and even a description of their computational solution is involved. We discuss here a related model that deals with decision making in the presence of uncertainty about future mortality rates. While our model does not allow for analytical solutions either, it is transparent in terms of what is being optimized and permits an analysis of more general informational assumptions than the complex models used to analyse pension problems. In discussing the model we have the following three issues in mind.

Firstly, studies of the use of population forecasts in decision making[1] suggest that while the uncertainty of forecasts is readily acknowledged by both the users and producers of forecasts (Alho, Cruijsen and Keilman, 2006), there is considerable inertia in the adoption of new methods. Cohort-component forecasts of population have been produced in many European countries since the 1920s and 1930s (Alho and Spencer, 2005). Since that time, alternative forecast variants have typically been offered, but the users have almost invariably considered the middle variant only. As discussed by Alho, Cruijsen and Keilman (2006), providing alternative variants is a deficient method for handling uncertainty. Nevertheless, a reason for considering only one variant may be the cost of added complexity in a decision process involving several decision makers. Alho and Spencer (2005) note that frequent updates may similarly be costly and lead to a lack of predictability in other parts of the economic system. Thus, it is of interest to discuss low information

[1] For a recent study in the EU context, visit http://ec.europa.eu/employment_social/social_situation/docs/lot1_projections_finalreport_en.pdf

alternatives to rational expectations. In the lowest information one, we assume that the decision maker ('consumer', in our example) only learns the initial point forecast, makes a life course plan and never reconsiders. In an intermediate information case we assume that the decision maker is periodically provided with an updated point forecast that reflects the demographic development up to that time. She then produces a new plan for the rest of her life. This method of updating is continued until death. At any update, she does not take into account that she will be offered a chance to reconsider later on. Obviously, the low information alternatives should produce lower utilities than the rational expectations solution, on average. The question is how much the decision maker loses compared to a rational expectations solution.

Secondly, the OLG models developed for Denmark, Finland, Germany, Netherlands, and the U.K. that are discussed in this volume (Lassila and Valkonen, 2006; Fehr and Haberman, 2006; Armstrong et al., 2006; Weale, 2006) have been developed in a deterministic context. Or, only one future population path is considered, and decision makers are assumed to know it exactly. The authors are fully aware that the assumption of perfect foresight is not natural in the context of stochastic population development, but numerical complications (caused by lack of Markovianity and the high dimensionality of the resulting optimization problem) currently preclude a formulation with rational decision makers that take precautions in the face of uncertainty. In contrast, our setting is simple enough so that the computational problems of determining the optimal decisions can be overcome. Thus, we are able to evaluate, how much a decision maker gains if given access to the 'crystal ball' of perfect foresight as compared to the rational expectations solution.

Thirdly, both results in this volume and other work suggest that uncertainty of population development directly influences the sustainability of pension systems (Weale 2006) and public finances as a whole (Alho and Vanne 2006). On the other hand, although it has been established that past forecasts of morality have had major errors (National Research Council, 2000) it is also shown that for the sustainability of the health care systems other factors dominate (Ahn 2006). In our example we consider *aggregate mortality risk*, i.e. the uncertainty caused by the fact that the forecast of the average mortality of a cohort is uncertain. While it is well-known that aggregate mortality risk is a major concern for public pension systems and insurance companies selling annuities (Blake and Burrows, 2001; Friedberg and Webb, 2005), less is known about its direct importance to consumers trying to make optimal savings decisions. We will quantify this effect relative to *idiosyncratic mortality risk*, by which we refer to the uncertainty of an individual's lifetime conditionally on average cohort mortality. This, of course, is a topic of classical life insurance mathematics.

The life cycle savings problem we consider involves a consumer who has to decide how much to save for retirement. She understands the nature of idiosyncratic risk, but her understanding of the uncertainty of forecasting ranges from the low information alternatives, via knowledge of the exact predictive distribution of future mortality, to perfect foresight. Formally, the predictive distribution we use is essentially equivalent to the mortality model of Lee and Carter (1992). Instead of trying to replicate the mortality of any given country, we will numerically evaluate the model so that it is representative of mortality of Europe, as assessed in the UPE project (cf. Keilman, Cruijsen and Alho, 2006). A discrete Markovian approximation will be provided that represents the uncertainty of forecasting 50 years into the future. This allows us to carry out the dynamic optimisation (cf. Filar and Vrieze, 1997) needed for the rational expectations solution. Mathematically, the other informational settings are easier.

We will start by developing the mortality model. Then, we describe the consumer's decision problem in terms of her utility function, and the welfare measure that is used to translate utilities to equivalent consumption. Results for the various informational assumptions and different degrees of risk aversion follow. We conclude by considering the implications of our findings for future work.

Demographics

In this section, we present the stochastic process for aggregate mortality. The challenge is to specify a process that is simple enough so that it we can solve the life cycle savings problem of a fully rational consumer, but that can also be calibrated to be empirically relevant.

We assume that time is discrete and index the periods as $t = 1, 2, ..., T$ and ages as $j = 1, 2, ..., T$. Death occurs at the end of the period only.

Let $m_t(j) \geq 0$ be age-specific mortality rate among consumers in age j during period $t = 1, 2, ...$ True future mortality rates are taken to be random. Following loosely the bilinear model of Lee and Carter (1992), we specify the probability law of future mortality as follows. Let $m(j)$ be the current mortality rate in age *j* and assume that it is known. Its expected rate of decline is $b(j)$, and the median mortality t years ahead is $\hat{m}_t(j) = m(j)\exp(-tb(j))$. For a simple representation of the uncertainty of the forecast, we define a Bernoulli process X_t such that $P(X_t = 1) = P(X_t = 0) = 1/2$ and then a discrete random walk $Y_t = Y_{t-1} + X_t$, where $Y_0 = 0$. This is a first order Markov process with mean (or drift) *t/2*. Let us further define a function $Z_t(y) = (y - t/2)\phi$, where $\phi > 0$ is a volatility parameter to be specified later. Then, the true age-specific mortality is assumed to be of the form

$$m_t(j) = m(j)\exp(-(t + Z_t(Y_t))b(j)) \tag{14.1}$$

Denote the probability that a consumer who is alive in age *j* at time *t* will be alive in the beginning of the next period by $s_t(j)$. For calibration purposes we assume that each period corresponds to five calendar years. Assuming that mortality is constant within a period, we have

$$s_t(j) = \exp(-5m_t(j)) \tag{14.2}$$

In the analysis below, we will consider only the problem of consumers who are of age 1 in period 1. That is, we specialize to the case $t = j$. We will take the first period to correspond to age 20-24 and the last period to age 105-109. After that everyone is assumed to die. Therefore, the latest time period we will have to consider is $T = 18$. On the other hand, it will be important to keep in mind that the probabilities (14.1) are random, since they are functions of Y_t via (14.1). To make this transparent, when $t = j$ we will write $s_t(j) = S_t(Y_t)$ for (14.2). We assume that period t mortality rates are observed at the end of period t. Due to the Markovianity of process Y_t, both the expectation and median forecast at time v for the survival probability at time $t \geq v$ depend on Y_{v-1}. We denote the median by $\hat{S}_t^v$. A technical issue that arises here is that under our model, randomness is incorporated in a nonlinear manner into survival probabilities (14.2), so setting aggregate variance to zero does not exactly yield unconditional expected values. Instead, the relevant quantities must be determined numerically.

The mortality process is parameterized by the values $m(j)$, $b(j)$ and ϕ. In calibrating these we aim to provide realistic values in a European context, rather than to replicate the values of any given country at any given time. Thus, we use combined male and female mortality data from Finland in 2002 for the $m(j)$'s,[2] and average rates of decline in mortality from eleven European countries during 1970-2000 (Alho and Spencer 2005, p. 235). These empirical data were available up to age 99. The estimates were extended to age-group 100-104 by assuming mortality to increase at the same rate as it increased from age 90-94 to age 95-99. The rate of decline over time was assumed to be the same as in age 95-99. For simplicity, we further took $m(j) = 0$, for $j < 10$. To specify the volatility of the process, we used data from nine European countries with long mortality series of good quality. The data suggest that a random walk can provide an approximate representation for the forecast error of the log of the age-specific mortality, when the standard deviation of the process increment is taken to be 0.06 (Alho and Spencer 2005, p. 256). By taking into account that an average value for the $b(j)$'s is about 0.08, we arrive at a value $\phi = 3.3541$. Table 14.1 displays $m(j)$ and $b(j)$. In addition, it displays median probabilities of surviving from one period to the next.

To see how the calculations were set up, note first that the rate of decline corresponds to five year model periods, so the annual rate of decline is about 0.08/5 = 0.016, or about 1.6%. Also, note that we are displaying cohort survival probabilities here. Thus, e.g. the median mortality of those in age 85-89 has been reduced from the value 0.190 at t = 0, during 14 model periods to 0.190*exp(-14*0.062) = 0.07976. Thus, the median survival is exp(-5*0.07976) = 0.671, the value given in the last column of Table 14.1 for this age-group.

[Table 14.1 HERE]

[2] These correspond to a life expectancy at birth of about 78 years, which is the average period life expectancy for females in Europe in 2000 (United Nations 2002). Thus, the model can serve as an approximation of the combined male-female life expectancy for *cohorts* born now. From Table 2.3 we find that the UPE assumption of the combined life expectancy in 2049 is about 86.5 years for the EEA+ region. This area has a higher life expectancy than Europe as a whole.

As an illustration of the degree of uncertainty related to future mortality, we present in Figure 14.1 the distribution of average lifetimes over 10,000 realizations of the mortality process. The standard deviation is approximately 3.3 years. Since this represents *cohort* survival 50 years into the future, it agrees well with the somewhat larger UPE values of Table 3.3 that relate to *period* survival in 2049.

[Figure 14.1 HERE]

A related aspect of the mortality model is the relationship between aggregate uncertainty and idiosyncratic uncertainty. Consider median mortality. An individual's lifetime has a standard deviation of about 6.1 years around the point forecast. (This is a somewhat too low value, because we have completely eliminated the 'outliers' that occur, in reality, before age 65.) This means that the total standard deviation that includes the aggregate uncertainty is approximately $(6.1^2 + 3.3^2)^{1/2}$ = 6.45. We see that from an individual's point of view adding the aggregate uncertainty only adds about 5% to the standard deviation, or it does not dramatically alter the prospects of longevity.

The consumer's problem

We consider a simple life cycle savings problem. The consumer receives wage income during the first periods and is retired for the last periods of her life. Let w_t denote her wage income in period *t*. During retirement, the consumer needs to finance her consumption with her own savings. The model is a partial equilibrium one in that all prices are exogenously given. Prices are also independent of demographics. The periodic utility function is denoted by $u(c)$, where $c > 0$ stands for consumption.

We assume that the only asset available to the consumer is a financial asset, k, that pays a constant interest rate, r. We also impose a borrowing constraint: $k \geq 0$. Together with lifetime uncertainty, the absence of both annuities and a bequest motive implies that the consumer will leave accidental bequests whose value is not reflected in her lifetime utility.

In the case of rational expectations about future mortality, it is the easiest to write the consumer's problem recursively. Given the simple structure of the mortality process, in the beginning of period *t*, a single integer, namely $Y_{t-1} \in \{0,1,...,t-1\}$, is sufficient to determine the best possible forecast about mortality in period t and beyond. Hence, Y_{t-1} is the only state variable needed to capture all relevant information about aggregate mortality in period *t*.[3] Denoting the period *t* value function by V_t, and the discount factor by $\beta > 0$, the problem can be written recursively as follows:

$$\begin{aligned} &V_t(k_t, Y_{t-1}) = \max\{u(c_t) + \beta E_{Y_t|Y_{t-1}}[S_t(Y_t)V_{t+1}(k_{t+1}, Y_t)]\} \\ &\text{such that} \\ &k_{t+1} = w_t + (1+r)k_t - c_t \end{aligned} \tag{14.3}$$

The solution to this problem includes a savings function $k_{t+1} = k_{t+1}(k_t, Y_{t-1})$ and a consumption function $c_t = c_t(k_t, Y_{t-1})$ that determine next period's savings and current consumption given age, current savings, and demographic development up to period *t-1*.

We will compare the above rational expectations problem to two alternate problems where the consumer has less information about the mortality process, and to the problem with perfect foresight. For the ease of comparison, we write these problems using backward recursion even though, in practice, they are more easily solved directly, in a single global optimization routine. With perfect foresight, the consumer’s problem can be written as follows:

$$\begin{aligned} &V_t(k_t, \{Y_s\}_{s=t}^T) = \max\{u(c_t) + \beta S_t(Y_t)V_{t+1}(k_{t+1}, \{Y_{s+1}\}_{s=t}^T)\} \\ &\text{such that} \\ &k_{t+1} = w_t + (1+r)k_t - c_t \end{aligned} \tag{14.4}$$

We have included the path of future Y’s as an argument of the value function to highlight the fact that the optimal decision in period t depends on all future survival

[3] Under the scaled model of error and its parametrisations discussed in Alho, Cruijsen and Keilman (2006) the situation would be more complicated. For the random walk model of age-specific fertility the state variable would include a vector with as many components as there are child-bearing ages, say, 30. For mortality the whole past history for males and females would have to be added (comprising at time *t* of, say, $t \times 2 \times 101$ components). The same number would have to be added for migration. Moreover, to the extent that any future optimizations in OLG models depend on actual population numbers, we might, in principle, need to add the full population history, as well. Thus, even though the present formulation leads to a numerical optimisation problem, it is far more transparent than the other models considered in this volume.

probabilities which are assumed to be known with uncertainty.[4] Clearly, the consumption-savings decisions depend on demographic developments in the future.

Consider then the other extreme case where the consumer learns only the initial point forecast in period $t = 0$, makes a life cycle savings plan believing that future mortality rates will be exactly as predicted by the point forecast, and never reconsiders the savings plan. Following again the recursive formulation, we can write the consumer's problem in this case as follows:

$$\begin{aligned} &V_t(k_t) = \max\{u(c_t) + \beta \hat{S}_t^1 V_{t+1}(k_{t+1})\} \\ &\text{such that} \\ &k_{t+1} = w_t + (1+r)k_t - c_t \end{aligned} \tag{14.5}$$

Recall that $\hat{S}_t^1$ is the period 1 point forecast for surviving period t. The solution to this problem is a sequence of consumption-savings decisions which are independent of the demographic path.

In the other low information case, the consumer updates her savings plan periodically based on the most recent point forecast. However, she still does not take uncertainty into account. For instance, in period 1 she makes a life cycle savings plan under the assumption that her future survival probabilities will be exactly those predicted by the period 1 point forecast. Then, in period 2 she is given a new point forecast (determined by Y_2). She makes her period 2 savings decision taking her current savings as given and under the assumption that future survival probabilities will be exactly those predicted by the new point forecast, etc. For a given aggregate mortality path, the problem the consumer faces at time t can be written down formally as a sequence of dynamic optimization problems where future mortalities are assumed to be given by point forecasts that are conditional on demographic development up to time t - 1. Thus, we have formally that

$$\begin{aligned} &V_t(k_t, Y_{t-1}, \{E[Y_s \mid Y_{t-1}]\}_{s=t}^T) = \max\{u(c_t) + \\ &\beta S_t(E[Y_t \mid Y_{t-1}]) V_t(k_t, Y_{t-1}, \{E[Y_s \mid Y_{t-1}]\}_{s=t+1}^T)\}, \end{aligned} \tag{14.6}$$

---

[4] In practice we solve this problem separately for each sequence $\{Y_t\}_{t=1}^T$ in our simulations.

if survival is evaluated at the conditional expectation (rather the conditional median). In other words, the consumer decides as if she had a 'crystal ball' like in (14.4), but she actually only has conditional expectations given Y_{t-1} to work with.

Calibration

The retirement age is set at 10, which corresponds to real age 65. Until retirement, the consumer earns a wage income of 1 every period, i.e. $w_t = 1$ for $t < 10$ and $w_t = 0$ for $t \geq 10$. Preferences are time-separable and the periodic utility function is of the constant-relative-risk-aversion form. For $\sigma \geq 0,\ \sigma \neq 1$ we have

$$u(c) = \frac{c^{1-\sigma}}{1-\sigma} \tag{14.7}$$

where σ measures the constant relative risk aversion. For $\sigma = 1$ we have $u(c) = \log(c)$.

Wages are paid and consumption occurs at the beginning of the period. For simplicity, we assume that both the interest rate and the subjective discount factor are zero. That is $\beta = 1$ and $r = 0$.

The welfare measure

Our welfare measure is the consumption equivalent variation. It gives the percentage increase in consumption in all periods that is needed in a benchmark case to make the expected lifetime welfare as high as in a comparison case. For instance, in order to compute the welfare cost of not taking aggregate mortality risk into account, we first solve the consumer problem with rational expectations. Then, we solve the consumer's problem assuming that she makes her decisions based on either mean (or median) mortality rates only. Next, we generate a large number *N* of randomly drawn aggregate mortality paths. Finally, we compute the increase in consumptions needed to make average lifetime utilities equal in the two cases.

More formally, consider simulation *i = 1,…, N* with a mortality path $\{Y_t^i\}_{t=1}^T$. Let c_t^i denote the optimal consumption level in simulation *i*, in age *t*, given rational expectations. Similarly, let $\tilde{c}_t^i$ denote the optimal consumption at age *t* when the consumer makes her decisions based on point forecasts alone.[5] The consumption equivalent variation measuring the welfare cost of using the median mortality rates instead of having rational expectations about future mortality is a scalar x such that

$$\sum_{i=1}^{N}\sum_{t=1}^{T}\prod_{s=1}^{t-1}S_s(Y_s^i)\beta^s u(c_t^i)=\sum_{i=1}^{N}\sum_{t=1}^{T}\prod_{s=1}^{t-1}S_s(Y_s^i)\beta^s u((1+x)\tilde{c}_t^i) \quad (14.8)$$

Results

The welfare cost of not taking aggregate mortality risk into account

How important is it for the consumer to take aggregate mortality risk into account? In the first two rows of Table 14.2 we display the welfare cost of making savings decisions based on point forecasts alone rather than having rational expectations about future mortality. The first row ('constant') corresponds to the case where the consumer receives only the period 1 point forecast. The second row ('adjust') corresponds to the case where the consumer updates her savings plan every period based on the most recent point forecast. The results are shown for different values of the risk aversion parameter.

[Table 14.2 HERE]

The welfare cost of using just the period 1 point forecast rather than having rational expectations varies between 0.51% and 0.62% in terms of a consumption equivalent variation depending on the degree of risk aversion. The interpretation is that in order to make the expected lifetime utility of a consumer with only period 1 point forecast equal to that of a consumer that takes the uncertainty fully into account, we need to increase the consumption in the low information case by 0.51%-0.62% in every period. As we explained above, the point forecasts are based on median mortality. We could also consider expected mortality. In experiments not reported here, we found that that would result into slightly smaller welfare costs.

The welfare cost of not taking aggregate mortality risk into account is dramatically smaller when the consumer is allowed to update her life cycle savings plan based on the most recent point forecasts. It then ranges between just 0.05% and 0.06% of consumption. This is a remarkably small welfare cost. In terms of expected

[5] In the case where the consumer learns only the period 1 point forecast, the consumption path is independent of the aggregate mortality path.

lifetime utility, it is not important for the consumer to take the uncertainty related to aggregate mortality into account. It suffices that she periodically updates her beliefs about future mortalities based on the most recent point estimates and reconsiders her savings plan accordingly.

It is perhaps worth noting that the relationship between the welfare cost of not having rational expectations and the risk aversion is non-monotonic. This is probably related to the fact that given our utility function, a high risk aversion means also a low intertemporal elasticity of substitution. A risk-averse consumer reacts to aggregate mortality risk by saving more at young ages (the so called precautionary savings motive, cf. Kimball, 1990). On the other hand, even when consumers fail to take aggregate mortality risk into account, consumers with higher risk aversion, and hence a low elasticity of intertemporal substitution, save more because they want a smoother consumption profile at old ages when the survival probabilities are low. In other words, consumers with a high risk aversion parameter act as if they had a strong precautionary savings motive, even when they do not take the risk into account at all. An interesting extension of this analysis would be to consider a more general utility function where we can separate risk aversion and the intertemporal elasticity of substitution.

So far, we have discussed the expected or *average* welfare loss from not taking aggregate mortality risk fully into account. Figure 14.2 displays the *distribution* of welfare losses resulting from making savings decisions based on the period 1 point forecast rather than rational expectations. Here we assume $\sigma = 3$. The figure shows clearly that even though for most aggregate mortality paths the welfare difference is relatively small, in some cases having rational expectations would have improved welfare substantially. The largest welfare gains are equivalent to about 5% in consumption in every period.

Figure 14.2 also shows that there is a probability of approximately 25% that the low information solution is better than the rational expectations solution (the welfare ‘cost’ is negative). This is to be expected for a range of mortality paths that are near to what we expect based on the point forecast, because adjustment to random changes in early years sometimes leads to erroneous adjustments regarding saving for later life. However, we also see that in all cases the gain is 0.6%, or less. (This result is not an absolute bound, of course, as it depends on the sample size used in simulation. However, it is clear that the probability of greater gains is less than 1%.)

[Figure 14.2 HERE]

Figure 14.3 shows how the welfare cost stemming from using only the period 1 point forecast rather than having rational expectations is related to the average lifetime of a cohort. As one would expect, the welfare cost is highest in the cases where average lifetime is very different from the expected one (which is about 14.1 model periods).

[Figure 14.3 HERE]

Figure 14.4 displays the distribution of welfare costs resulting from using updated point forecasts rather than having rational expectations. The distribution is very different from that in figure 14.2. Even the largest welfare costs are now rather small, around 0.6%, and the distribution is rather symmetric.

[Figure 14.4 HERE]

Figure 14.5 shows how this welfare cost is related to the average lifetime of a cohort. Just like in the lowest information case, the updating method is better than rational expectations in (the rather unlikely) high mortality cases. For low mortality cases the reverse is true. In cases of expected mortality both approaches can dominate, but rational expectations sometimes produce fairly substantial gains over the updating method.

[Figure 14.5 HERE]

Perfect foresight vs. rational expectations

A number of studies (e.g. Lassila and Valkonen, 2006; Fehr and Haberman, 2006) consider the implications of demographic uncertainty using numerical general equilibrium models. However, as discussed earlier most of these analyses assume perfect foresight over future demographics. Therefore, it is of interest to try to see how important, or problematic, this assumption is. To analyse this issue, we compare the consumer's savings problem with rational expectations to the one in which the consumer has access to a 'crystal ball' that tells the future aggregate mortality rates precisely, but does not reveal her own lifetime.

Table 14.3 displays the average welfare gain from having perfect foresight instead of just rational expectations. The welfare gain following from perfect foresight is perhaps surprisingly small: it ranges from 0.16% to 0.19% in terms of consumption. One can argue that the normative solution to decision problems of the type we consider is obtained via rational expectations. Hence, the welfare costs computed here are measures of the welfare cost of aggregate mortality risk. Since this cost is small, the perfect foresight assumption appears to be a reasonable proxy for rational expectations in this context.

[Table 14.3 HERE]

Figure 14.6 shows the distribution of welfare gains resulting from having perfect foresight rather than rational expectations. Comparing Figures 14.2 and 14.6 shows that while using only the initial point forecast may lead to substantial welfare losses (around 5% in consumption), having perfect foresight instead of rational expectations never increases welfare dramatically (the gain in consumption is typically less than 2.5%).

[Figure 14.6 HERE]

Figure 14.7 shows how the welfare gain of having perfect foresight over rational expectations is related to the average lifetime of a cohort. Again, the welfare gain is highest in cases where aggregate mortality is very high. However, the link between average lifetime and the welfare gain is relatively weak.

[Figure 14.7 HERE]

Bracketing Rational Expectations

A consumer who has access to a 'crystal ball' is never worse off than those with less information. This has to hold even pathwise. Similarly, the rational expectations solution must always yield a higher *average utility* than that obtainable in the low information settings. Thus, forecast updating solution and the perfect foresight solution can be used to bracket the rational expectations solution. The result can be read from Tables 14.2 and 14.3. Depending on the level of risk aversion, the rational expectations solution is in a band of width ≤ 0.25% above the updating solution, when average utility is used as the measure.

Although this goes a long way towards understanding the nature of the rational expectations solution, it is clear that other types of utility (or loss) functions may weigh subsets of the data in a manner that is different from the one we use. For example, Alho, Lassila and Valkonen (2006) consider combinations of pension contribution rates and replacement rates that may be viewed as being politically acceptable (or they belong to the *viable region*) in the sense that if contribution rates would become too high, the workers would refuse to pay them, or if replacement rates would be too low, higher pensions would be demanded by the elderly. This directs attention to the joint distribution of the contribution rates and replacement rates, not just to their expected values.

Thus, it is also of interest to see how well the updating solution and the crystal ball solution can bracket the rational expectations solution for different levels of life expectancy, for example. It would be too much to ask for a pathwise bound, but we can look for bounds for conditional expectations. In the model at hand the relevant determinant seems to be the average lifetime of the cohort, so in Figure 14.8 we display nonparametric estimates[6] of average gain of the 'crystal ball' solution, and of the updating solution, over rational expectations, as a function of the average lifetimes.

[Figure 14.8 here; to be added]

The finding is that for low average lifetimes the gap between the 'crystal ball' solution and the updating solution is wider than the average gap, near the expected average lifetimes it is narrower, and for high average lifetimes it grows again. Perhaps one implication is that for mortality paths close to the expected development the rational expectations solution can be bracketed quite closely with the

[6] Estimates were obtained by the LOWESS smoother of Cleveland (1981).

perfect foresight and updating solutions. Note, however, that for low life expectancies, even forecast updating produces utilities that are higher than those given by the rational expectations solution. The reason appears to be the precautionary savings motive. Both the consumer with rational expectations and a consumer believing in updated point forecast save, on average, too much in cases where longevity is low. However, the consumer with rational expectations saves more than the one following the forecast updating procedure because the former takes uncertainty into account with additional precautionary savings.

Conclusions

We have analyzed the importance of aggregate mortality risk using a standard life cycle model with a consumption-savings decision under different informational settings. We draw three types of conclusions.

Regarding the effect of different informational assumptions our main findings are as follows: Firstly, the expected welfare cost of aggregate mortality risk is small for a consumer who knows the probability law and takes it rationally into account. That is, her expected lifetime utility would increase only little had she access to a 'crystal ball' that reveals future mortality rates without error. Secondly, a consumer who does not take the uncertainty related to future mortality into account at all, but makes her savings decisions based on the most recent point forecasts alone, loses very little compared to a consumer who takes the predictive distribution of future mortalities rationally into account. This is probably due to the fact that idiosyncratic uncertainty dominates an individual's longevity. (In contrast, aggregate uncertainty dominates the uncertainty of a pension institution, for example.) Thirdly, a consumer who makes all her savings decisions based on the point forecast available at the beginning of her adult life, suffers substantial welfare losses for extreme mortality paths relative to a consumer having rational expectations. - Intuitively, these results are all related to the fact that the consumer's information about her life expectancy improves over time. By adjusting her savings decisions as new information arrives, she can get relatively close to the perfect foresight solution even if she does not take the uncertainty into account.

An important numerical finding is that the forecast updating formulation produced utilities that are perhaps surprisingly close to the rational expectations solution. This suggests that a promising direction of future research is to develop computational tools that allow the forecast updating approach to be implemented in the OLG settings used in pension analyses. One step in this extension has already been taken by the development of a computer program FPATH[7] that computes numerical approximations to future population forecasts that are conditional to population development up to their respective jump-off times.

An observation that arises from our having considered different informational assumptions in a stochastic setting, is that there is merit in considering the whole distribution of utility outcomes, in addition to expected utility. This can reveal patterns of utility gains that are not easily detected by intuition alone. In a similar vein, we note that rational expectations solutions can sometimes be bracketed by seemingly crude formulations, like here the perfect foresight and the forecast updating formulations for mortality outcomes that are not too far removed from the expected one.

Finally, an important task for future research is to extend the present analysis to a set-up that includes an unfunded pension system. As shown by Lassila and Valkonen (2006) and Weale (2006), demographic uncertainty makes the sustainability of public pension systems very uncertain. Although our numerical estimates suggest that consumer's welfare losses that are due to not taking into account aggregate uncertainty in mortality are not very large, the 'viability' issues we mentioned suggest that the possibility of structural change should be considered. Thus, the ultimate utility depends on the design of the pension system. If the system includes transparent rules that specify how the benefits and contributions are adjusted with changes in demographics, individuals may prepare for structural change, by adjusting their savings plans correspondingly. However, unfunded (Pay-As-You-Go) pension systems typically do not include such rules, and pressures for structural change may suddenly mount to a point when cost pressures are politically deemed unacceptable. Thus, in practice, failing to take demographic uncertainty into account (via additional precautionary saving) may result in much larger welfare losses for the consumer than the loss we have considered.

References

[7] For a description, see http://joyx.jo
ensuu.fi/~ek/pep/fpath.pdf.

Ahn (2006) 'On demographic uncertainty and health care expenditure'. P. xxx-xxx n this volume.

Alho J.M., Cruijsen H. and Keilman N. (2006). 'Empirically-based specification of forecast uncertainty'. Pp. xxx-xxx in this volume.

Alho, J.M. and Määttänen, N. (2006). 'Aggregate mortality risk
and the value of annuities', discussion paper n. 1005, The Research Institute of the Finnish Economy.

Alho, J.M., Lassila, J., and Valkonen, T. (2006) 'Demographic Uncertainty and Evaluation of Sustainability of Pension Systems', pp. 95-112, in Holzmann, R. and Palmer, E. (2006) *Pension Reform*. Washington D.C.: The World Bank.

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

Alho, J. and Vanne R. (2006). 'On Predictive Distributions of Public Net Liabilities', *International Journal of Forecasting*, 22: xxx-xxx.

Armstrong, A., N. Draper, A. Nibbelink, and E. Westerhout (2006). 'Ageing, demographic uncertainty and optimal fiscal policy'. Pp. xxx-xxx in this volume.

Diamond, P. (2001). 'Comment', in Auerbach A.J. and R.D. Lee (eds.), Demographic Change and Fiscal Policy. Cambridge: Cambridge University Press.

Blake, D. and Burrows, W. (2001). 'Survivor Bonds: Helping to Hedge Mortality Risk', *Journal of Risk and Insurance*, 68(2), 339-348.

Cleveland, W. S. (1981). 'LOWESS: A program for smoothing scatterplots by robust locally weighted regression'. *The American Statistician*, 35, 54.

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

Friedberg, L. and Webb, A. (2005). 'Life is Cheap', Working paper 2005-13, Center for Retirement Research at Boston College.

Kimball, M.S. (1990). 'Precautionary Saving in the Small and in the Large', *Econometrica*, 58(1), 53-73.

Lee, R.D. and Carter, L.R. (1992). 'Modeling and Forecasting the Time Series of U.S. Mortality'. *Journal of the American Statistical Association* 87, 659-671.

Lassila, J., and Valkonen, T. (2006). Demographic uncertainty and pension projections. Pp. xxx-xxx in this volume.

National Research Council (2000). Beyond six billion. Panel on Population Projections, National Research Council. Washington D.C.: National Academy Press.

United Nations (2002). *World Population Prospects, vol. I.* New York: United Nations.

Weale, M. (2006). 'Fiscal implications of demographic uncertainty: comparisons across the European Union'. P. xxx-xxx in this volume.

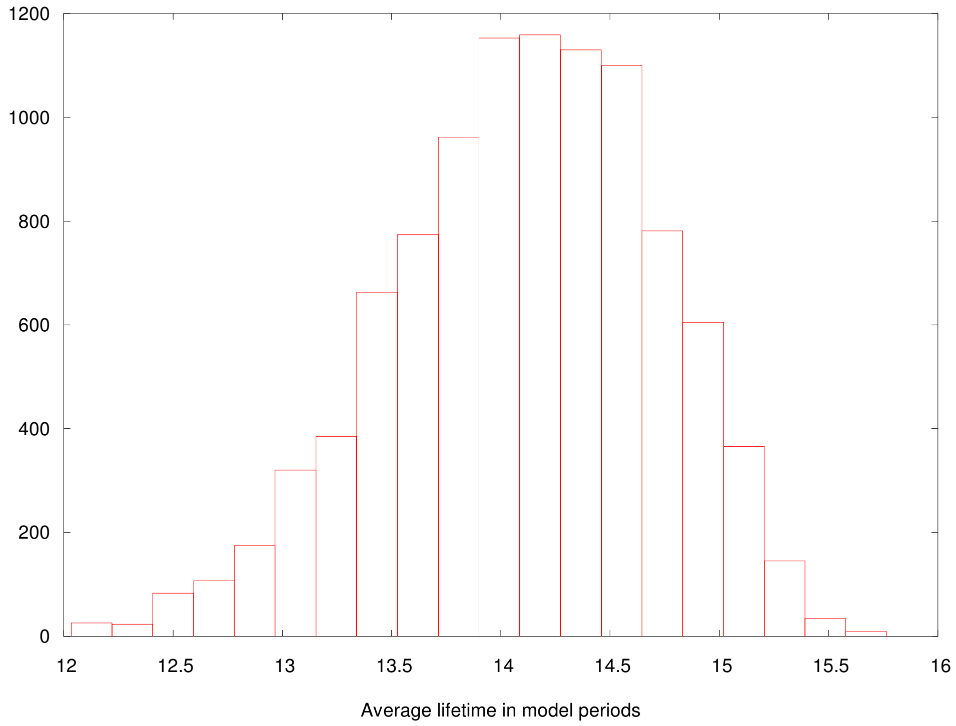

Figure 14.1 Histogram of average lifetimes in 10,000 simulations

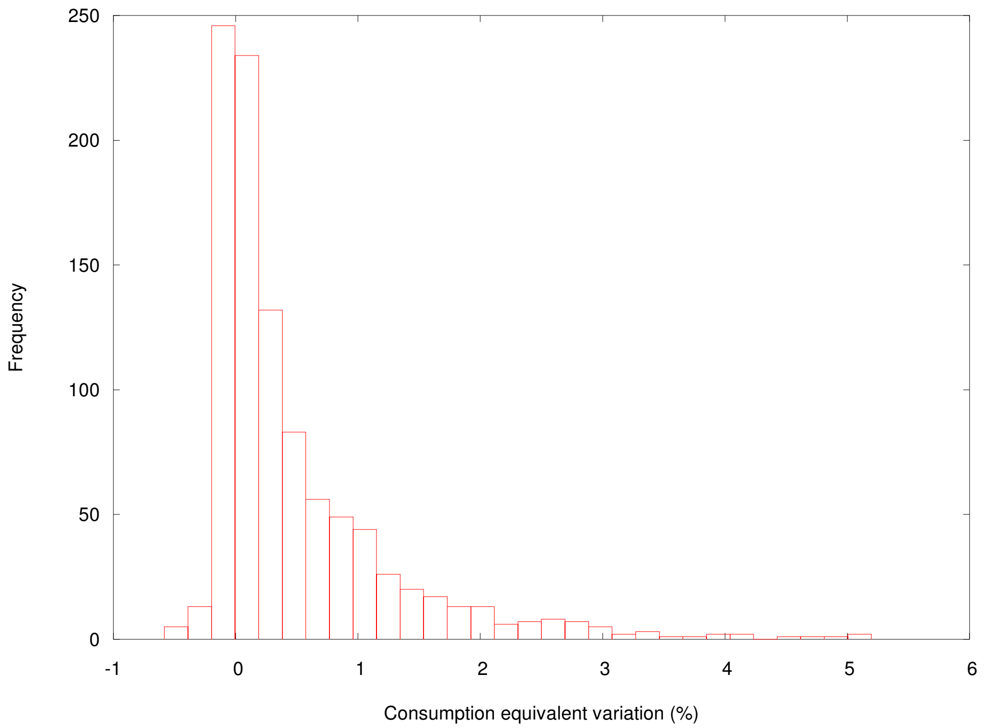

Figure 14.2 Welfare costs of using period one point forecast rather than having rational expectations

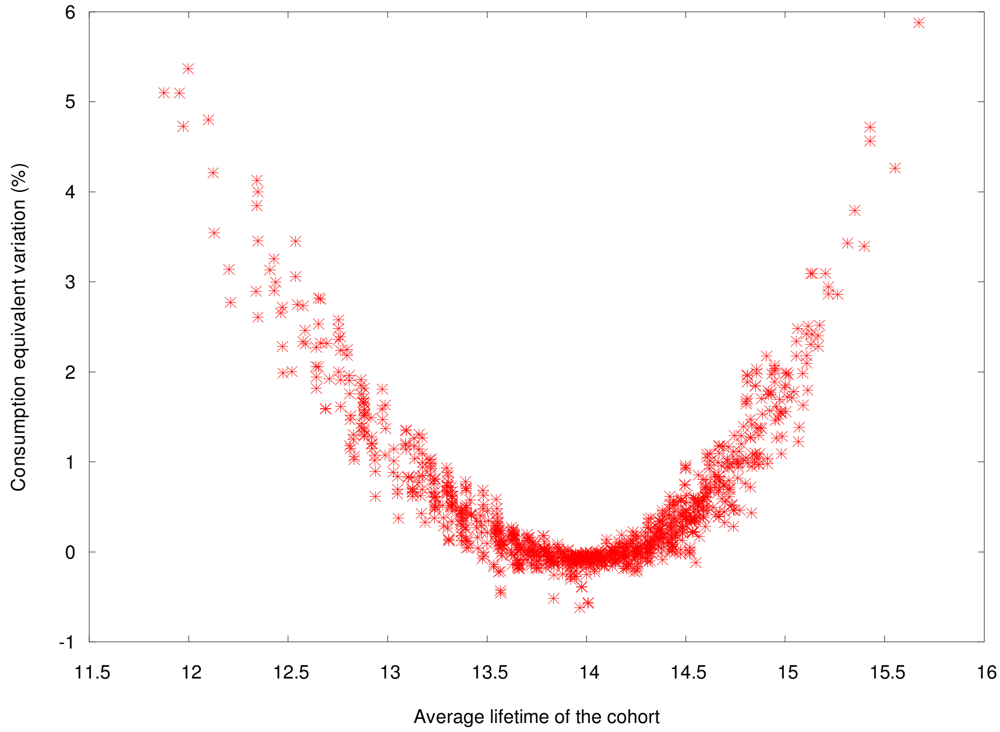

Figure 14.3 Welfare costs of constant expectations vs. rational expectations as a function of average lifetimes

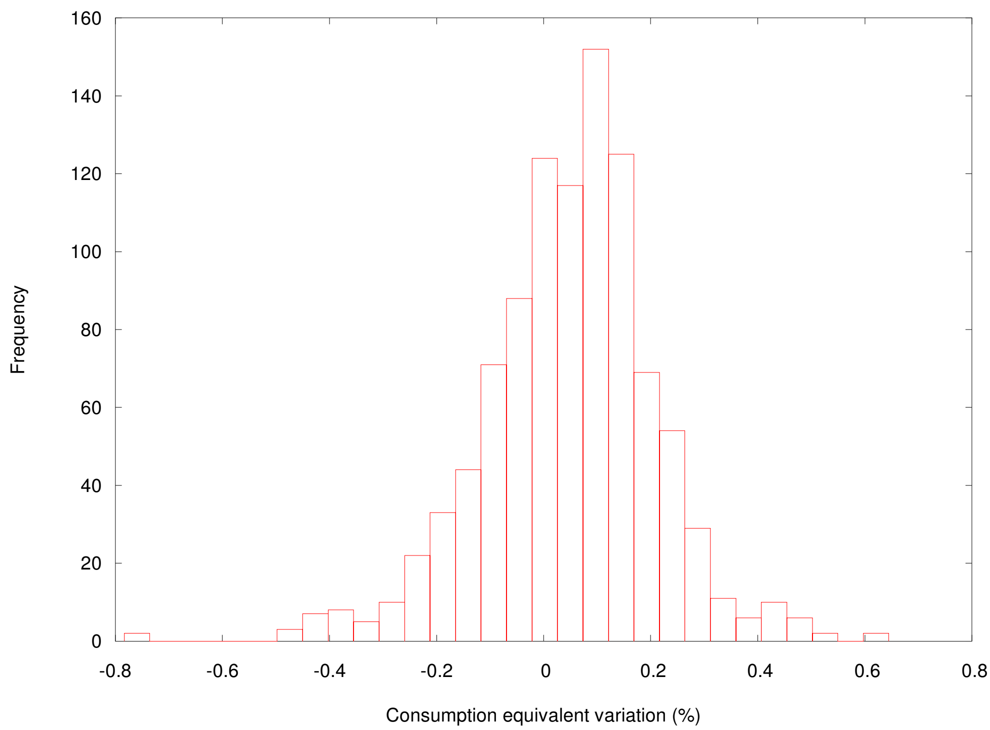

Figure 14.4 Welfare costs of updated point forecasts vs. rational expectations

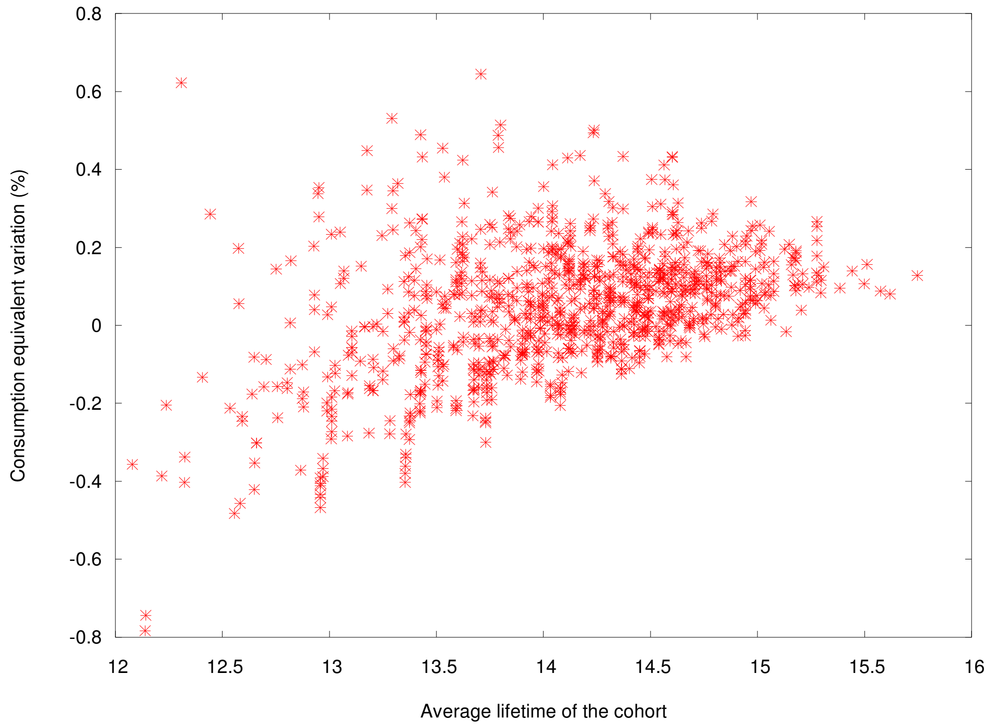

Figure 14.5 Welfare costs of updated point forecasts vs. rational expectations as a function of average lifetimes

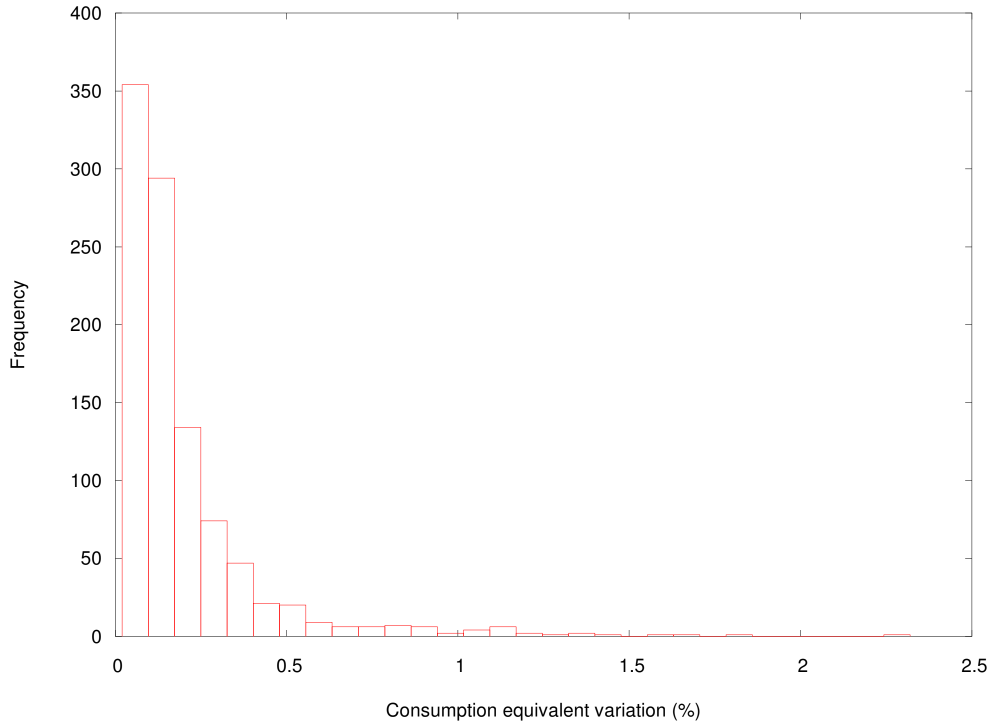

Figure 14.6 Welfare gains of having perfect foresight vs. rational expectations

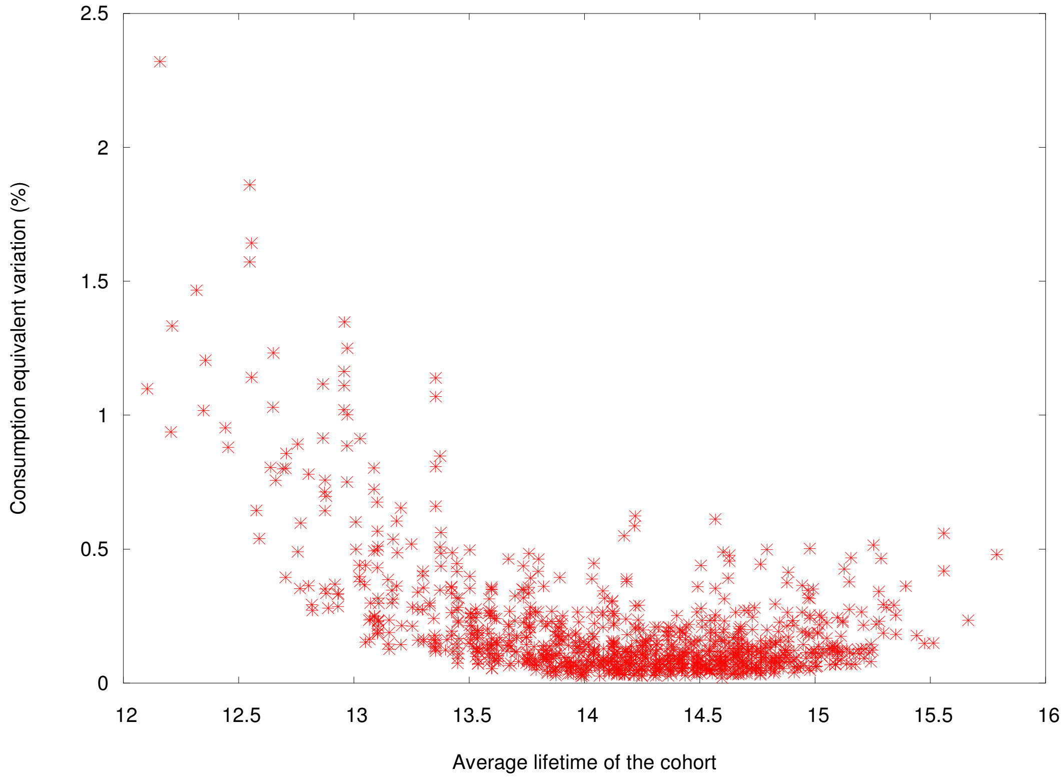

Figure 14.7 Welfare gains of perfect foresight vs. rational expectations as a function of average lifetimes

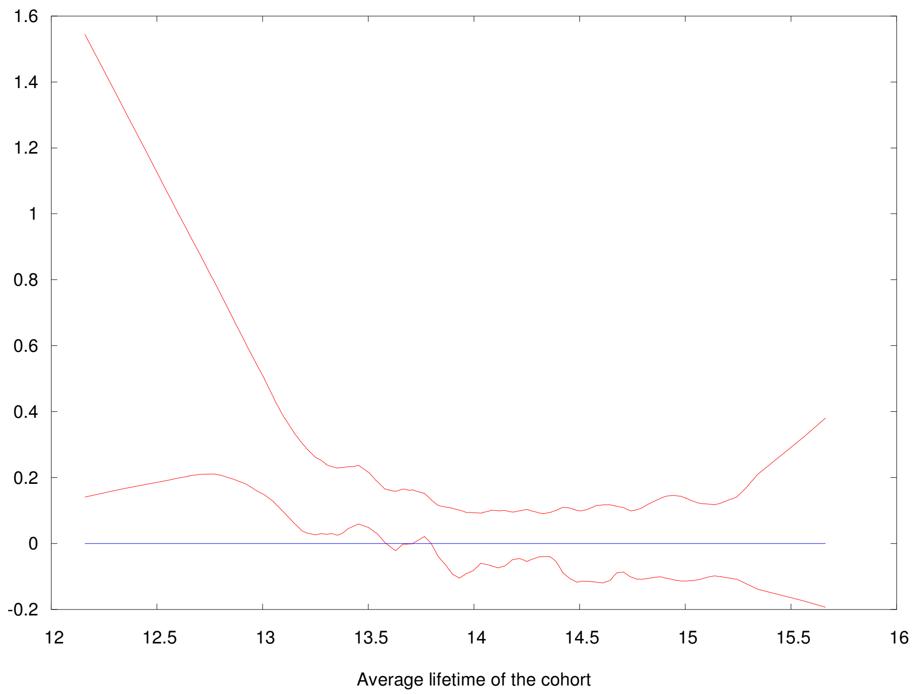

Figure 14.8 Average welfare gains of perfect foresight vs. rational expectations (upper curve), and average welfare gains of using updated point forecasts vs. rational expectations (lower curve) as a function of average lifetimes.

Age	Model age	m	b	median survival prob.
20-64	1-9	0	0	1
65-69	10	0.026	0.100	0.955
70-74	11	0.044	0.100	0.931
75-79	12	0.072	0.100	0.886
80-84	13	0.119	0.078	0.806
85-89	14	0.190	0.061	0.671
90-99	15	0.283	0.044	0.481
95-99	16	0.404	0.029	0.280
100-104	17	0.577	0.029	0.172
105-109	18	∞	0	0

Table 14.1 Parameter values for the mortality process and the median survival probabilities

	Risk aversion		
Expectations	σ=1	σ=3	σ=5
Constant	0.52%	0.62%	0.51%
Adjust	0.04%	0.06%	0.05%

Table 14.2: The welfare cost of not taking aggregate mortality risk into account under low information alternatives.

Risk aversion		
σ=1	σ=3	σ=5
0.16%	0.19%	0.16%

Table 14.3: The welfare gain of having perfect foresight vs. rational expectations