TFR Predictions Based on Brownian Motion Theory[1]

NICO KEILMAN, Ph.D.
Professor
Department of Economics, University of Oslo, Norway

Abstract

In stochastic population forecasts, the predictive distribution of the TFR is of central concern. Common time series models can be used to predict the TFR and its moments on the short run (up to 10 or 20 years), but on the long run (40–50 years) they result in excessively wide prediction intervals. The aim of this study is to analyse and apply a time series model for the TFR, which restricts the predicted values to a certain pre-specified interval.

I will model the time series of log TFR-values as a Brownian motion with absorbing upper barrier. I will give and analyse
* *expressions for the predictive distribution of the log of the TFR assuming it follows a Brownian motion with absorbing ceiling;*
* *expressions for the first and second moments of the predictive distribution of the log of the TFR.*

When the log of the TFR follows a random walk with absorbing ceiling, I find that the second moment of the predictive distribution for the long-run TFR in Norway is insensitive for ceiling levels beyond a threshold of approximately 3.4 children per woman. This conclusion holds for a fairly broad range of innovation variances.

If the log of the TFR follows a random walk, sample paths that exceed approximately 3.4 children per woman may be rejected when simulating future fertility in Western countries. This will not have any major effect on the width of the long-term predictive distribution.

1. Introduction

In the last decade, a number of stochastic forecast models have been developed (Lee and Tuljapurkar 1994; Lutz and Scherbov 1998; Alho 1998; De Beer and Alders 1999; Keilman et al. 2001). The basic idea of a stochastic forecast, as opposed to the traditional deterministic forecast, is to think of the population in the future not as one number, but as a whole distribution: some numbers are more likely than others. If forecast results are presented in the form of *predictive* distributions of this kind, the user is forced to take forecast uncertainty into account. At the same time, forecasts of this nature are less vulnerable to sudden changes in fertility and mortality trends than traditional deterministic forecasts.

Stochastic population forecasting is based on the Cohort Component Method. But instead of one set of parameters for fertility, mortality, and migration, as in the traditional deterministic method (or perhaps three, when a high, a medium, and a low forecast variant are computed), one specifies the joint statistical distribution of all input parameters. Because of the large amount of input parameters (35 fertility rates, 200 death rates, and 120 parameters for net migration for each forecast year), simplifying assumptions are used. First, one assumes that fertility, mortality, and migration are independent. For industrial countries this is a reasonable assumption. Second, one focuses on the distribution of a few summary indicators (for instance total fertility rate, life expectancy at birth, level of net-immigration), thereby ignoring the statistical distributions of the detailed parameters (age-specific rates).

There are three main methods for specifying the statistical distributions of the key parameters: time-series extrapolation; expert judgement; and extrapolation of historical forecast errors (Keilman 2001). Time-series methods seldom give an accurate description of the past. If many of the historical facts remain unexplained, time-series methods result in excessively wide prediction intervals when used for long-term forecasting (Saboia 1977; Lee 1993; Keilman and Pham 2000). Judgemental methods can be used to correct or constrain such broad prediction intervals. Expert judgement is also used when expected values and corresponding prediction intervals are hard to obtain by formal methods (Lutz et al. 1998). The expert-judgement method implies that a group of experts is asked to indicate the probability that a key parameter (such as the average number of children per woman, or life expectancy) in some future year falls within a certain pre-specified range. A weakness of this approach is that experts, often being unduly confident, tend to attach an overoptimistically high probability to an interval of a certain width. Finally, empirical errors observed for past forecasts may be extrapolated to predict the expected errors for the current forecast.

This study presents a time series model, in which expert judgement is used to restrict the prediction intervals. The contribution I hope to make is to provide a sharper

distinction between two of the three methods currently in use for stochastic population forecasts. Illustrated by the case of the Total Fertility Rate (TFR), it shows mathematically how the choice for a certain upper bound to the future TFR influences its predictive distribution. For high enough values of the upper bound, the distribution is close to an unconstrained one. The lower the bound, the stronger its impact on the predictive distribution. The aim of this study is to analyse this sensitivity, and to find a critical value for the upper bound of the process in question, that is, a value above which the impact of the bound on the distribution in practice can be neglected.

Keilman and Pham (2000) present TFR-predictions for Norway for the years 1996–2050. These are based on an ARIMA (1,1,0) model for the logarithm of the annual TFR, estimated on the basis of data for the years 1945–1995. The long-term predictions for the mean and median values of the TFR appear reasonable: 2.2 and 1.9 children per woman, respectively. However, and not surprisingly, the prediction intervals are excessively wide in the long run. For example, in 2050 the 95 per cent prediction interval for the TFR ranges from 0.6 to 6.1 children per woman. Around the year 2020, the upper 95 per cent bound exceeds the level of four children per woman, clearly an unrealistic value for Norway, even if the probability is only 2.5 per cent. Thus, the model produces informative prediction intervals up to 20–30 years ahead, but in the long run the predictions are unrealistic and the prediction intervals are too wide.

The reason why the ARIMA model for the long term produces wide prediction intervals for the period TFR is that it contains no extraneous information beyond past data. Several factors are associated with the development of fertility after World War II in Norway, similar to other countries. Important elements in this respect are the introduction of modern contraception, the adoption of new norms and values with respect to childbearing and partnership, increased interest in tertiary education, and growing levels of labour force participation among women (Kravdal 1994; Lesthaeghe and Surkyn 1988). None of these factors is, or could be, explicitly modelled. It is reasonable to expect that they will be in operation in the future, too. Yet we have no idea to what extent these or other factors will constrain Norwegian fertility in the next century, be it period or cohort developments.

One has to take recourse to other methods when predictions so far ahead are required. The easiest one is to assume that in 2020, say, uncertainty is already so high that it will not increase any more (e.g. Alho and Spencer 1997). In that case prediction intervals are constant after 2020. An alternative method is to assume that there is an upper bound to fertility levels in Norway. This can be done in several ways: by means of an appropriate logit transformation of the TFR, by means of simulation in which extreme values are rejected, or by assuming that the TFR can be modelled as a time series which includes an upper bound. In Sections 2 and 3, I

will discuss earlier studies based on logit transformation and simulation. In Section 4, I will assume that the TFR can be modelled as a Brownian motion, which is absorbed at a certain ceiling. Section 5 contains an application of that model using Norwegian TFR data.

2. Logit transformation

Lee (1993) has used a logit transformation in order to restrict predicted TFR-values. He first transformed the annual TFR for the United States into

$$g_t = \ln\{(\text{TFR}_t - L)/(U - \text{TFR}_t)\}$$

where L and U are pre-specified lower and upper limits for the TFR. Next, he identified and estimated an ARMA-model for the transformed variable g_t. This logit transformation will produce a forecast for the TFR that will never exceed U or fall below L. However, as noted by Alho and Spencer (1997), such a model may have undesirable consequences. They demonstrated that when g_t follows a random walk process, then TFR_t will eventually be "absorbed" close to U or L for large enough t. This anomaly also showed up empirically in the case of Norway. In an earlier study (Keilman and Pham 2000), we selected $L=1.0$ and $U=3.1$, and identified an ARIMA (2,1,0)-process for the logit-transformed Norwegian TFR. (The maximum TFR-value in the period 1945–1995 was 3.02, in 1964. Hence, an upper bound of 3.0 would cause the transformation to break down in that year.) In 2050, the bounds of the 67 per cent prediction interval (1.12, 2.86) were very close to those of the 95 per cent interval (1.01, 3.08). In the long run, the boundaries of any interval approach the upper and lower bounds U and L arbitrarily closely. The conclusion is that a logit transformation cannot be used for constraining prediction intervals that result from a non-stationary model.

3. Rejecting extreme values

In Keilman and Hetland (1999) we simulated prediction intervals for the TFR in Norway by means of the ARIMA (1,1,0) model mentioned in Section 1, thereby rejecting all sample paths that exceeded certain extreme values. We experimented with upper bounds ranging from 2 to 5, and inspected the resulting prediction intervals. We found that setting reasonable limits to the predicted TFR has little impact on the bounds of the 67 percent prediction intervals in the short run. In the year 2020, this interval is 0.80 children per woman-wide when the TFR is limited to 3 children, and 0.92 and 0.95 children-wide for maximum TFR values of 4 and 5, respectively. In the long run however, uncertainty (as expressed by the width of the 67 percent prediction interval) grows faster. In 2050, the interval widens from

1.31 children per woman for a maximum TFR of 3, to 1.64 and 1.88 children per woman for maximum TFRs of 4 and 5, respectively. The 95 percent interval bounds are rather strongly influenced by the choice of a certain maximum TFR-value, also in the short run. In 2010, the width amounts to 1.60, 1.92, and 2.06 children per woman for maximum TFR-values of 3, 4, and 5. A maximum TFR equal to 3 would imply a median and a mean TFR that decrease over time, which was judged unreasonable. Therefore the general conclusion was that a maximum TFR equal to 4 children per woman appeared to be a reasonable choice. (Experiments with a *lower* bound to the TFR suggested that minimum values up to 0.5 children per woman hardly influenced the TFR prediction intervals.)

4. Brownian motion

Choosing upper and lower bounds for the TFR, as discussed in the previous sections, is largely an arbitrary matter. At the same time, the choice may have important consequences for the shape of the predictive distribution and the width of the prediction interval, as was argued in Sections 1 and 3. In the current section I shall discuss a method based on the theory of Brownian motion, which makes it possible to analyse how sensitive the prediction interval is for a certain choice of the upper bound. A Brownian motion (also called Wiener process) is a normal process with stationary independent increments (e.g. Lamperti 1966; Feller 1968). Assuming that the process starts at time $t=0$, the predictive distribution is that of a normal variable with a standard deviation proportional to $\sqrt{t}$. While time is a continuous variable in a Brownian motion model, its discrete-time counterpart is the Random Walk. When an absorbing upper bound is introduced to the Brownian motion, the predictive distribution of the resulting process is no longer normal.

Assume that a random variable $Z=Z_t$ follows a Brownian motion, but that it is absorbed as soon as it reaches a ceiling a. The transition density of the process at time t is (Borodin and Salminen 1996, 107)

$$f(z_t) = \frac{1}{\sigma\sqrt{2\pi t}}\left\{\exp\left(-\frac{(z_t - z_0)^2}{2\sigma^2 t}\right) - \exp\left(-\frac{(z_t - (2a - z_0))^2}{2\sigma^2 t}\right)\right\}, \quad z_t, z_0 < a. \tag{1}$$

where z_0 is the starting position at time $t=0$ of the random variable, and σ^2 is the variance of the innovation process. When a approaches infinity, the second exponential vanishes, and the density approaches that of a normal variable with expectation z_0 and variance $\sigma^2 t$, as expected. I shall derive expressions for the expected value and the variance of Z_t as a function of the ceiling a, given that the process has not been absorbed. To the best of my knowledge, these expressions have not been given before.

At time t there are two possibilities: the process moves freely on the line $-\infty < z_t < a$, or it is absorbed at the ceiling a. The probability that we encounter the first situation is $\int_{-\infty}^{a} f(z_t)dz_t$. Using expression (1) it can be shown that this non-absorption probability equals

$$\int_{-\infty}^{a} f(z_t)dz_t = 2\Phi\left(\frac{a - z_0}{\sigma\sqrt{t}}\right) - 1, \tag{2}$$

where $\Phi(u)=\Pr[U \leq u]$ denotes the cumulative distribution of a standard normally distributed variable U. The probability that the process is absorbed at a is the complement of this expression, viz. $2\{1-\Phi(.)\}$, and the transition density for the event that the process is in z_t, given that it has *not* been absorbed, is $g(z_t)=f(z_t)/[2\Phi(.)-1]$. The expected value and variance of the probability density function g can be computed using its moment generating function $M(s) = \int_{-\infty}^{a} e^{sz} g(z_t)dz_t$.

Writing u for $(a-z_0)/\sigma\sqrt{t}$, one may readily verify that the expected value of $g(z_t)$ is

$$E_g(Z_t) = M'(0) = \frac{z_0 - 2a\{1-\Phi(u)\}}{2\Phi(u)-1} = a - \frac{a - z_0}{2\Phi(u)-1}. \qquad a < \infty \tag{3}$$

For large a, the probability of non-absorption $(2\Phi(u)-1)$ approaches 1, so that the expected value approaches z_0. For finite a, $2\Phi(u)-1<1$, and the expected value is lower than z_0. $\Phi(u)$ falls when t increases, so that the expected value becomes progressively lower than z_0. All these findings are as one would expect.

After some tedious but straightforward algebra one finds for the second moment of $g(z_t)$

$$M''(0) = (z_0^2 + \sigma^2 t) - 4a(a - z_0)\frac{1-\Phi(u)}{2\Phi(u)-1} + \frac{2(a-z_0)}{2\Phi(u)-1}\sqrt{\frac{\sigma^2 t}{2\pi}}\exp\left(-\frac{u^2}{2}\right), \qquad a < \infty$$

so that its variance equals (4)

$$Var_g(Z_t) = (z_0^2 + \sigma^2 t) + \left\{-4a(a-z_0)\frac{1-\Phi(u)}{2\Phi(u)-1} - \left(a - \frac{a-z_0}{2\Phi(u)-1}\right)^2\right\} + \frac{2(a-z_0)}{2\Phi(u)-1}\sqrt{\frac{\sigma^2 t}{2\pi}}\exp\left(-\frac{u^2}{2}\right).$$

For large a, $\Phi(u)$ approaches one, while at the same time the exponential vanishes. In that case the second moment is close to $z_0^2 + \sigma^2 t$, and the variance approaches $z_0^2 + \sigma^2 t - z_0^2 = \sigma^2 t$, which is the variance of an unbounded Brownian motion.

How do the expected value in expression (3) and the variance in expression (4) change when the ceiling a changes? For fixed z_0, t, and σ^2, the expected value increases monotonically with growing a. It is lower than the starting value z_0, but approaches that value asymptotically, when a is increased. For a approaching z_0 (from above), the expected value approaches (by L'Hôpital) $z_0 - \sigma\sqrt{(\pi t/2)}$.

The variance in expression (4) is the sum of three terms, T_1, T_2, and T_3, say. $T_1 = z_0^2 + \sigma^2 t$ represents the variance of a Brownian motion – it is independent of the ceiling a. T_2 is written in curled brackets, and T_3 is the term which contains the exponential function. Figure 1 illustrates how T_2, T_3, and the variance depend on a, for the standard case $z_0 = 0$, $t = 1$, and $\sigma^2 = 1$.

Figure 1. Variance of a Brownian motion with ceiling a, and its components.
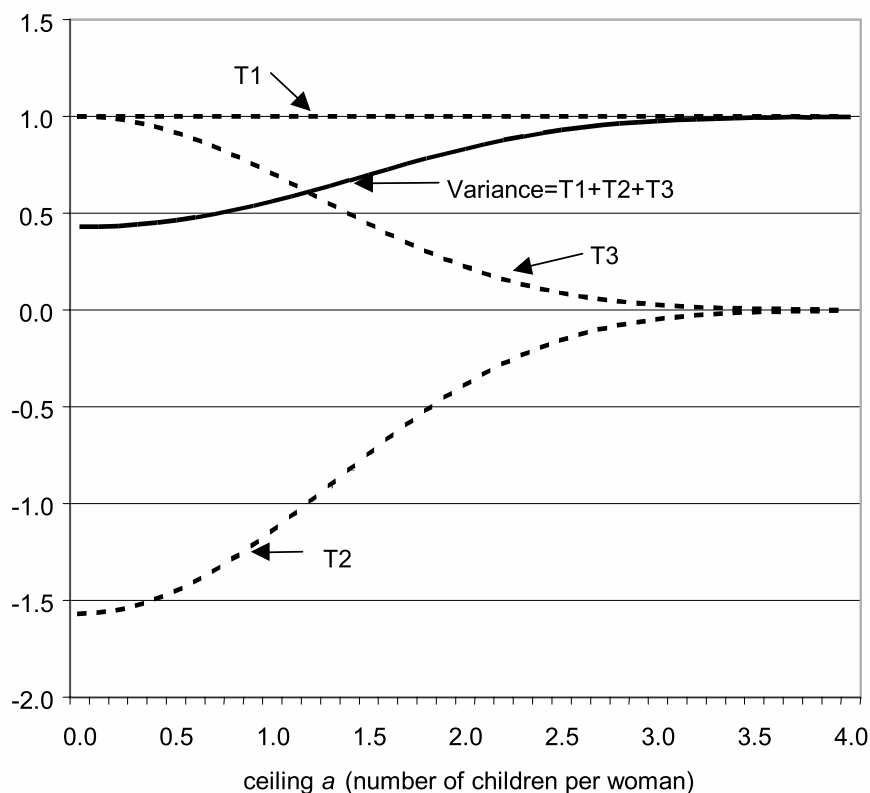

The figure shows that the variance is S-shaped. For small a, it starts a little under 0.5, and it approaches $\sigma^2 t = 1$ asymptotically when a grows. The reason is that, for large enough a, the term T_2 approaches $(-z_0^2) = 0$, while T_3 dies out.

Figure 1 holds for $t = 1$ only. When time increases, the S-curve is shifted upwards, and at the same time it is stretched to the right. For very small yet positive a, the

variance is close to $t(2-\pi/2)$, and for large a it approaches t (still assuming $z_0=0$ and $\sigma^2=1$), which confirms the earlier finding. Empirically I found that for a given ceiling a, at time $t \approx a$, the variance is close to its asymptotic level t.

Figure 2 shows how the form of the probability density function $g(z_t)$ depends on the ceiling a, again for the standard case $z_0=0$, $t=1$, and $\sigma^2=1$. Note that the curves for a equal to 3 or 4 are almost indistinguishable, while the curve for $a=2$ is very similar to that for $a=3$. This confirms what was found for the variance in Figure 1: the variance changes little when a is increased from 2 to 3, and for a larger than a "threshold-level" of around 3, it remains virtually the same. In empirical applications, the threshold level for the ceiling, i.e. the level beyond which the probability density hardly changes, will depend on the starting point z_0, and in particular on the innovation variance σ^2 (and on time t).

Figure 2. Probability density function $g(z_t)$ for a Brownian motion with absorbing ceiling a between 1 and 4.
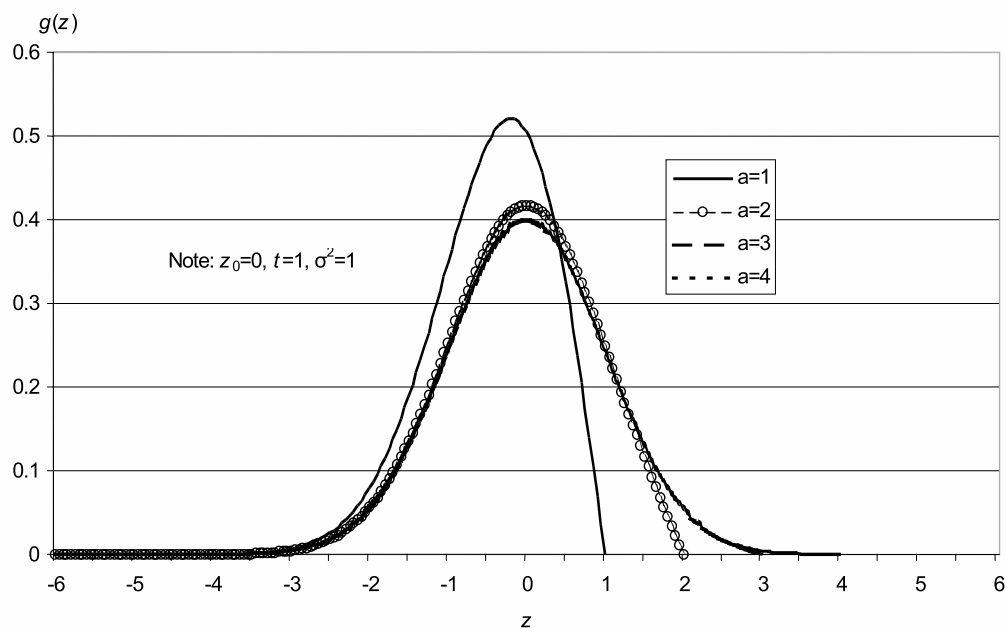

5. Application

In this section I shall work in discrete time, and hence use the term "random walk" instead of "Brownian motion". In order to illustrate the expressions of Section 4, I have assumed that the log of the TFR for Norway in the period 1900–2000 follows a random walk. This is obviously not a good model: elsewhere it was found that log-values of the TFR follow an ARIMA (1,1,0) model (Keilman and Pham 2000). But this assumption was used in order to get an impression of the order of magnitude of the innovation variance.

Estimates of the innovation variance σ^2 varied a great deal, depending upon the estimation period. I estimated a value for σ^2 equal to 0.0017 for the entire period 1900–2000, and values of 0.00084, 0.00083, 0.00102, and 0.00041 for the periods 1950–2000, 1960–2000, 1970–2000, and 1980–2000, respectively.

For large values of a, the process Z_t is an unconstrained random walk, and the TFR-process $Y_t=\exp(Z_t)$ is the exponential of a random walk. In that case, predictions for Z_t are normally distributed, and Y_t has a lognormal distribution with known expectation and variance. When the ceiling becomes effective, the distribution of Z_t is no longer normal.

I have analysed the expected value and the standard deviation of Y_t as a function of the ceiling $b=\exp(a)$ for the medium-term (2030) and the long-term (2050) situation. Since $Y_t=\exp(Z_t)$, it has expectation $E(Y_t)=\exp[E(Z_t)+\frac{1}{2}\mathrm{Var}(Z_t)]$ and variance $\mathrm{Var}(Y_t)=\{\exp[\mathrm{Var}(Z_t)]-1\}.\exp[2E(Z_t)+\mathrm{Var}(Z_t)]$. The process starts in the year 2000, with $Y_0=1.851$, which is the TFR-value observed in that year. The innovation variance σ^2 was chosen equal to 0.0008. Figure 3 shows how the expected value and the standard deviation of Y change when the ceiling b goes up.

Figure 3. Expected value and standard deviation of Y_t as a function of the ceiling b.
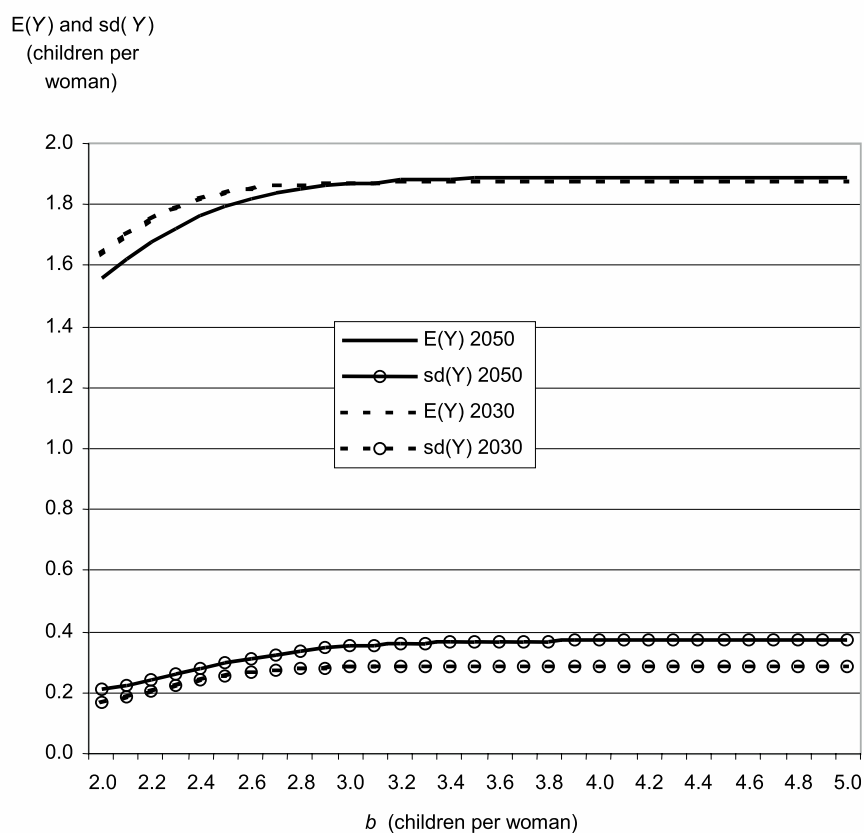
Both variables increase towards their maximum values, as the ceiling moves upwards. We see that ceiling values b higher than approximately 2.5 children per woman have little impact on the expected value. This insensitivity is explained by the fact that at $b=2.5$, the absorption probability $2(1-\Phi(u))$ is very low in 2030 (5 per cent), and still moderate (13 percent) in 2050. The higher this probability, the stronger the ceiling influences the pdf of the random walk process Z_t. For the standard deviation of Y_t, the threshold is approximately 3 children per woman, a level at which the absorption probabilities are no more than 0.2 and 1.6 percent in 2030 and 2050.

Note that when the ceiling is below 3 children per woman, the expected value in 2050 is lower than in 2030. $E(Y_t)=\exp[E(Z_t)+\tfrac{1}{2}\text{Var}(Z_t)]$ depends on the variance of Z_t, which increases with time. However, for low ceiling values, the expected value of Z_t falls as a function of time, because the process is more and more likely to be absorbed. The net result is that also $E(Y_t)$ falls with time. The higher the innovation variance, the larger the absorption probability is, and the steeper the fall in $E(Y_t)$ over time. Take as an example in Figure 3, with σ^2 equal to 0.0008, the case of $b=2.4$ children per woman. The TFR in 2050 (1.761) is 0.057 children lower than that in 2030 (1.818). But when σ^2 equals 0.0004, the difference between the 2050 TFR (1.832) and the 2030 TFR (1.852) is only 0.02 children.

In Figure 3, the expected value and the standard deviation of Y_t change very little for ceiling values beyond a threshold of 3 children per woman. The threshold depends strongly on the innovation variance σ^2. This can be analysed more formally as follows. $E(Y_t)$ varies with b, but for large enough b it is independent of the ceiling (in practice for $b>6$). The latter value of $E(Y_t)$ may be called the unconstrained value. It corresponds with the pure random walk value of the original Z-process. For a certain level of b, $E(Y_t)$ is at 95 percent of its unconstrained value. Denote this as the threshold level b^*. In Figure 3, for instance, the unconstrained value of $E(Y_t)$ in 2050 is 1.89. This leads to a threshold level b^* of 2.5. Similarly, for the standard deviation of Y_t, Figure 3 shows a threshold level of 3.0 in 2050.

Table 1. Threshold levels b^* for expected value and standard deviation of Y_t in 2050, for various levels of innovation variance.

Innovation variance σ^2	Representing the period	Expected value of Y_t		Standard deviation of Y_t	
		b^*	$E(Y_t;b^*)$	b^*	$sd(Y_t;b^*)$
0.0004	1980–2000	2.2	1.78	2.6	0.25
0.0008	1950/1960–2000	2.5	1.79	3.0	0.36
0.0012	1970–2000	2.8	1.81	3.4	0.45
0.0016	1900–2000	3.0	1.83	3.7	0.53

Table 1 shows that the threshold level for the year 2050 increases by between 0.2 and 0.4 children per woman, when the innovation variance of the underlying random walk is increased by 0.0004. The threshold for the standard deviation is higher than that for the expected value, and it increases a bit faster. Threshold levels in 2030 were found to be lower than those for 2050 (figures not shown here).

Hence one may conclude as follows. When innovation variance values in the range 0.0004–0.0012 are thought to be representative of the TFR-development in Norway after World War II (erroneously assuming that a random walk is an appropriate model for the log of the TFR), choosing a ceiling equal to 3.4 children per woman will not have any major impact on the standard deviation or the expected value of the predicted TFR up to the year 2050.

6. Conclusion

In this paper I have demonstrated that when the log of the TFR follows a random walk, choosing a ceiling equal to 3.4 children per woman will not affect the expected value and the standard deviation of the predicted TFR up to the middle of this century. This conclusion holds for a fairly broad range of innovation variances, including rather extreme variance values. A reasonable choice for that variance leads to a sharper estimate for the ceiling's threshold value, namely 3 children per woman.

This conclusion has important consequences for stochastic population forecasting. It implies that formal methods can be used when defining the bounds of prediction intervals for one of the forecast's most important input variables, e.g. the TFR (provided of course that the TFR can be modelled as a random walk, or that it is close to a random walk). The upper bound *can* be set lower than three children per woman. But that implies that the resulting prediction interval is influenced, and such a subjective decision has to be justified carefully. Thus the current paper contributes to sharpening the distinction between subjectively defined prediction intervals for the TFR predominantly based on expert opinion, and intervals that are primarily based on time series analysis (Keilman 2001).

The random walk/Brownian motion model analysed in this paper might fruitfully be applied to stochastic forecasts for other demographic processes than fertility. For instance, annual outmigration flows (in logarithmic form) from Norway in the period 1958–1997 are described well by a random walk (Keilman et al. 2001). But the estimated innovation variance was so large that predictions further ahead than approximately five years resulted in too-wide intervals. Thus it was decided to keep the intervals constant after an arbitrarily chosen period of five years.

218

A methodological extension would be to introduce, simultaneously, an upper and a lower bound. In the current application, the lower bound was set to zero, by the logarithmic transformation of the TFR. Earlier experiments with Norwegian data indicated that a minimum TFR up to 0.5 children per woman had very little impact on the resulting prediction intervals. Therefore the current analysis focused on an upper bound. But in other applications, for instance for migration flows, or life expectancy, lower bounds will be as relevant as upper bounds. In these cases one could model Z_t=log(Y_t-α), with α the intended lower bound, and investigate the sensitivity of the expected value and the variance for different choices of α. Another extension would be to analyse more complicated bounded processes, based for instance on an integrated random walk, or an autoregressive process.

References

Alho, J. M. 1998. *A stochastic forecast of the population of Finland.* Reviews 1998/4. Helsinki: Statistics Finland.

Alho, J. and B. Spencer. 1997. *Statistical demography and forecasting.* Unpublished manuscript.

Borodin, A.N. and P. Salminen. 1996. *Handbook of Brownian Motion: Facts and Formulae.* Basel: Birkhauser Verlag.

De Beer, J. and M. Alders. 1999. Probabilistic population and household forecasts for the Netherlands. Working Paper nr. 45, Joint ECE-Eurostat Work Session on Demographic Projections, Perugia, Italy, 3–7 May 1999.
Internet <www.unece.org/stats/documents/1999.05.projections.htm>

Feller, W. 1968. *An Introduction to Probability Theory and its Applications.* Volume I. (3[rd] edition) New York: John Wiley & Sons.

Keilman, N. 2001. Uncertain population forecasts. *Nature* 412, 2 August 2001, 490–491.
Internet <http://www.nature.com/cgi-taf/DynaPage.taf?file=/nature/journal/v412/n6846/full/412490a0_fs.html>

Keilman, N. and A. Hetland. 1999. Simulated confidence intervals for future period and cohort fertility. Working Paper No. 6, Joint ECE-Eurostat Work Session on Demographic Projections, Perugia, Italy, May 1999.
Internet <www.unece.org/stats/documents/1999.05.projections.htm>

Keilman, N. and D.Q. Pham. 2000. Predictive intervals for age-specific fertility. *European Journal of Population.*16(1):41–66.

Keilman, N., D.Q. Pham, and A. Hetland. 2001. Norway's Uncertain Demographic Future. *Social and Economic Studies* no. 105. Oslo: Statistics Norway.

Kravdal, Ø. 1994. The importance of economic activity, economic potential and economic resources for the timing of first births in Norway. *Population Studies* 48:249–267.

Lamperti, J. 1966. *Probability: A Survey of the Mathematical Theory*. New York: W. A. Benjamin Inc.

Lee, R. 1993. Modeling and forecasting the time series of US fertility: Age distribution, range, and ultimate level. *International Journal of Forecasting* 9:187–202.

Lee, R. and S. Tuljapurkar. 1994. "Stochastic population forecasts for the United States: Beyond High, Medium, and Low", *Journal of the American Statistical Association* 89(428): 1175–1189.

Lesthaeghe, R. and J. Surkyn. 1988. Cultural dynamics and economic theories of fertility change. *Population and Development Review* 14(2):3–45.

Lutz, W. and S. Scherbov. 1998. An expert-based framework for probabilistic national population projections: The example of Austria. *European Journal of Population* 14(1):1–17.

Saboia, J. 1977. Autoregressive Integrated Moving Average (ARIMA) models for birth forecasting. *Journal of the American Statistical Association* 72(358):264–270.