Recent Rises in Cohort Fertility in the Industrialized World: Using Bayesian Methods to Extrapolate Trends while Preserving Cohort Features

Joshua Goldstein
Max Planck Institute for Demographic Research

Mikko Myrskyla
Max Planck Institute for Demographic Research

Carl Schmertmann
Center for Demography and Population Health, Florida State University

Emilio Zagheni
Max Planck Institute for Demographic Research

September 18, 2010

Abstract

There are already several documented examples of recent increases in cohort fertility in Scandinavia, but for most countries, cohorts are too young to see if cohort fertility has increased. We produce new estimates of completed cohort fertility for cohorts born in the 1970s. We combine the best of previous efforts, using cohort forecasting methods to preserve what demographers know about the age-pattern of fertility, and using trends in the age-period-cohort Lexis surface to tell us as much as possible about the way in which fertility appears to be changing over time. Our preliminary findings suggest that cohort fertility has stopped its long-term secular decline in the majority of low fertility countries around the world. In some cases, there is a clear suggestion of increase. As we further develop our models we expect to be able to make more precise statements about further trends and the certainty of our knowledge.

EXTENDED ABSTRACT

Introduction

Several recent studies have documented increases in period fertility for many developed countries (Goldstein, Sobotka and Jasilioniene 2009; Luci and Thevenon 2010; Myrskyla, Kohler and Billari 2009; Sobotka 2008; Trovato 2010). Many of these studies face an interpretational challenge as the increases may be due to increases in the quantum of fertility, or decreases in the pace of postponement (see, for example, Goldstein et al. 2009). Consequently, authors commonly attempt to control for postponement and recover a cohort perspective from period data by using tempo-adjustments. Such an approach may not always help, as tempo-adjustments require hard-to-get data and may introduce bias or add noise to the time series (Cheng, Grace and Goldstein 2010).

An alternative to tempo adjustments would be to directly use cohort measures of fertility, and study trends in cohort fertility. Indeed, there are already several documented examples of recent increases in cohort fertility in Scandinavia (Andersson et al. 2009). For most countries, however, cohorts are still too young to see if cohort fertility has increased. We aim to produce new estimates of completed cohort fertility for cohorts born in the 1970s by using new methods and unprecedentedly large database of fertility rates by age and year, covering 34 countries or regions. Our methods combine the best of previous efforts, using cohort forecasting methods to preserve what demographers know about the age-pattern of fertility, and using trends in the age-period-cohort Lexis surface to tell us as much as possible about the way in which fertility appears to be changing over time.

Figure 1 illustrates the recent cohort fertility trends and motivates further research on the topic. Here we have used a database of single year age and single year period fertility rates covering 34 countries, and extrapolated recent age-specific fertility rates to the future to allow completion of cohorts[1]. More specifically, for each single year age group from age 15 to 44 we take the last five years of observation, estimate the trend, and use the estimated trend to extrapolate fertility to the future. We use these age-specific fertility projections to complete cohort fertility. For most countries, the data is available up to year 2008, so that our forecasts for the youngest cohorts (1978) are based on observations up to age 30 and forecasts for ages 31-44.

The results of Figure 1 are striking: cohort fertility seems to have stopped its long-term secular decline in the vast majority of low fertility countries around the world. Only in 5 of the 34 countries or regions do we continue to see decreasing cohort fertility[2]. Even in countries like Russia and Korea, which are well-known for their fertility crisis, fertility has stopped to decline and shows weak signs of increase. Stronger increases are observed in Italy, Spain, Greece, Russia, and former Western and Eastern parts of Germany.

It would, however, be premature to conclude based on Figure 1 that the world-wide decline in cohort fertility has come to its end and possibly reversed. This is because of two crucial factors that are missing from the naive extrapolation exercise presented in Figure 1: First, the results do not contain any quantitative information on the credibility of the finding that the decline may have stopped. It is crucial to produce standard errors for the forecasts, in addition to the point estimates in the figures.

Second, naive extrapolation ignores existing knowledge that demographers have about age patterns of fertility. The extrapolated trends may imply future period or cohort schedules that demog-

[1]The data sources are described in later sections.

[2]Slovakia, Hungary, Poland, Portugal, and Korea.

raphers would immediately recognize as unlikely or even impossible. In contrast, forecasts requiring that the underlying Lexis surface of age- and time-specific rates conform to plausible demographic specifications might prove more accurate. In this paper we aim to improve forecasts by imposing such requirements.

Data

Our analysis relies mainly on data from the newly constructed Human Fertility Database[3], from Eurostat[4] and national statistical agencies[5]. We use age-specific fertility rates for age groups of one year and for time intervals of single calendar years.

Methods and Preliminary Results

Bayesian methods for demographic analysis and forecasting are gaining increased attention. For example, Lynch and colleagues (2001, 2003, 2005, 2007) have used Bayesian approaches to model intertemporal mortality patterns, health-education relationships, and life table indices. Raftery and colleagues (Alkema et al. 2009) have investigated the use of Bayesian models for smoothing of noisy TFR time series from multiple sources. Most importantly for our work, Girosi and King (2008) present a strong case for the utility of Bayesian models in mortality forecasting.

Our work is largely inspired by Girosi and King's emphasis on incorporating demographic knowledge into forecasts. Many of their estimation techniques are well suited to fertility studies. In short, we aim to do for fertility what they did for mortality. The key to their Bayesian approach is to incorporate demographers' prior knowledge about features of age schedules that do not directly concern the outcome of interest (cohort fertility levels, in our case), but which may nevertheless make some future Lexis surfaces far more likely than others. For example, demographers may have considerable information about the smoothness of such surfaces over age and time, about the unimodal shapes of cohort fertility schedules, and so forth. Girosi and King's Bayesian approach allows analysts to incorporate these features probabilistically into forecasts – i.e., by stating *a priori* that forecast surfaces with these properties are more probable than others. This approach contrasts with other methods that directly incorporate expert opinion about the forecast variable itself.

There are several existing approaches to cohort fertility forecasts. By far the most common approach is to assume that fertility rates for incomplete cohorts will resemble the age-specific experience of the most recent period (e.g., Frejka and Calot (2001)). This approach, although it has been long applied by the Council of Europe, will produce underestimates of the completed fertility of current childbearing cohorts, due to well-known tempo effects on period measures, when the fertility of older women is increasingly delayed. More sophisticated extrapolation approaches include Li Nan (2003), whose work is based on Lee-Carter approaches that use singular value decomposition of the observed Lexis surface of fertility rates. Single-cohort extrapolation methods based on parametric cohort models have also been tried (Goldstein 2008). Chen (2010) finds that both the extrapolatory trend method and the parametric methods produce unbiased forecasts, but that there are relatively large errors.

[3]Norway, Sweden, Switzerland, United States, Austria, Czech Republic, Slovakia, Russian Federation. Data available at: www.humanfertility.org

[4]Netherlands, France, Finland, Denmark, Belgium, Italy, Greece, Hungary, Portugal, Bulgaria, Iceland, Luxembourg, Lithuania, Spain, United Kingdom, Romania, Slovenia, Ireland, Estonia, Poland. Data available at: http://epp.eurostat.ec.europa.eu/portal/page/portal/eurostat/home

[5]Australia (www.abs.gov.au), New Zealand (www.stats.govt.nz).

Our goal is to use Bayesian methods with explicit priors that capture the knowledge on which existing forecasts methods are implicitly based. A forecast should respect trends in observed data, but without violating existing knowledge about age, period, and cohort fertility patterns. The main task of our research will be to devise appropriate specifications of prior knowledge, and appropriate computational methods, to carry out this idea.

As an example, meant to illustrate our approach, assume that:

1. Time trends of fertility rates at each age tend to be approximately linear. We operationalize this prior by penalizing curvature in age-specific time trends. (Details are in the Appendix.)

2. The shapes of future period schedules resemble those of past schedules. We operationalize this prior by assuming that period schedules should be well approximated by linear combinations of the first two principal components of a singular value decomposition of past rates. (Details are in the Appendix.)

In this case, by using quadratic penalties we can compute analytically the maximum a posteriori estimator for estimates and forecasts of fertility rates (see the appendix for details on the procedure). Figures 2 and 3 show the results for two countries: in the case of the Netherlands, the forecast looks plausible and consistent with the results of Figure 1, obtained using the naive linear extrapolation method. In the case of Switzerland, the method does not work well. For example, the series of fertility rates at age 20 is not plausible, as the rates become negative after 2020.

The method and the two examples are shown for illustrative purposes. In practice, Bayesian priors can be applied more fruitfully to cohort, rather than period, age schedules, in order to incorporate the information that demographers have about typical shapes of cohort fertility. The possibility of period shocks might also be included. Although we have used the quadratic penalties approach for this illustrative forecasts, we plan to use more flexible approaches that rely on MCMC methods in order to incorporate a broader family of priors in a direct manner, albeit at some computational expense.

A final methodological note is that, for example, an alternative to using local derivatives to penalize violations of smoothness might be to use a more global measure of linearity (e.g., R^2). Another example is that we can introduce parametric fertility age-schedules rather than the SVD approach. Each of these alternatives may prove easier to implement using a more general computational framework than the quadratic penalties approach.

References

[1] Alkema L., A. Raftery, P. Gerland, and F. Pelletier. 2009. Estimating the Total Fertility Rate from Multiple Imperfect Data Sources and Assessing Its Uncertainty. In *Population Association of America Annual Meeting 2009*, Detroit MI. http://tinyurl.com/lhvco9

[2] Andersson, G., M. Ronsen, L.B. Knudsen, T. Lappegard, G. Neyer, K. Skrede, K. Teschner, and A. Vikat. 2009. Cohort fertility patterns in the nordic countries. *Demographic Research* 20(14):313-352.

[3] Cheng, Y.-H.A., K. Grace, and J.R. Goldstein. 2010. The Susceptibility of Fertility Measures to Random Fluctuations. In *European Population Conference 2010.* Vienna, Austria.

[4] Chen, Y-H. A. 2010. The Future of Completed Cohort Fertility in Low Fertility Countries. In *Population Association of America Annual Meeting 2010* Dallas, TX.

[5] Frejka, T. and G. Calot. 2001. Cohort Reproductive Patterns in Low-fertility Countries. *Population and Development Review* 27:103-132.

[6] Girosi, F. and G. King. 2008. *Demographic Forecasting.* Princeton University Press.

[7] Goldstein, J.R. 2008. A Behavioral Gompertz Model for Cohort Fertility Schedules in Low and Moderate Fertility Populations. In *Population Association of America Annual Meeting 2008* New Orleans, LA.

[8] Goldstein, J.R., T. Sobotka, and A. Jasilioniene. 2009. The End of Lowest-low Fertility?. *Population and Development Review* 35(4):663-699.

[9] Luci, A.and O. Thevenon. 2010. Does Economic Development Drive Fertility Rebound in OECD Countries? in *Population Association of America Annual Meeting 2010.* Dallas, TX.

[10] Li, N. and Zheng W.. 2003. Forecasting Cohort Incomplete Fertility: A Method and an Application. *Population Studies* 57(3): 303-320.

[11] Lynch S.M. and J.S. Brown. 2001. Reconsidering Mortality Compression and Deceleration: An Alternative Model of Mortality Rates. *Demography* 38(1):79-95.

[12] Lynch S.M. 2003. Cohort and Life-Course Patterns in the Relationship between Education and Health: A Hierarchical Approach. *Demography* 40(2):309-331

[13] Lynch S.M. and J.S. Brown. 2005. A New Approach to Estimating Life Tables with Covariates and Constructing Interval Estimates of Life Table Quantities. *Sociological Methodology* 35:189-237.

[14] Lynch S.M. 2007. *Introduction to Applied Bayesian Statistics and Estimation for Social Scientists.* Springer: New York.

[15] Myrskyla, M., H.-P. Kohler, and F.C. Billari. 2009. Advances in Development Reverse Fertility Declines. *Nature* 460(7256):741-743.

[16] Sobotka, T. 2008. The Rising Importance of Migrants for Childbearing in Europe. *Demographic Research* 19(9):Overview Chapter 7.

[17] Trovato, F. 2010. Fertility in Alberta (Canada) in a Context of Rapid Economic Growth, 1997-2007. in *Population Association of America Annual Meeting.* Dallas, TX.

Figures

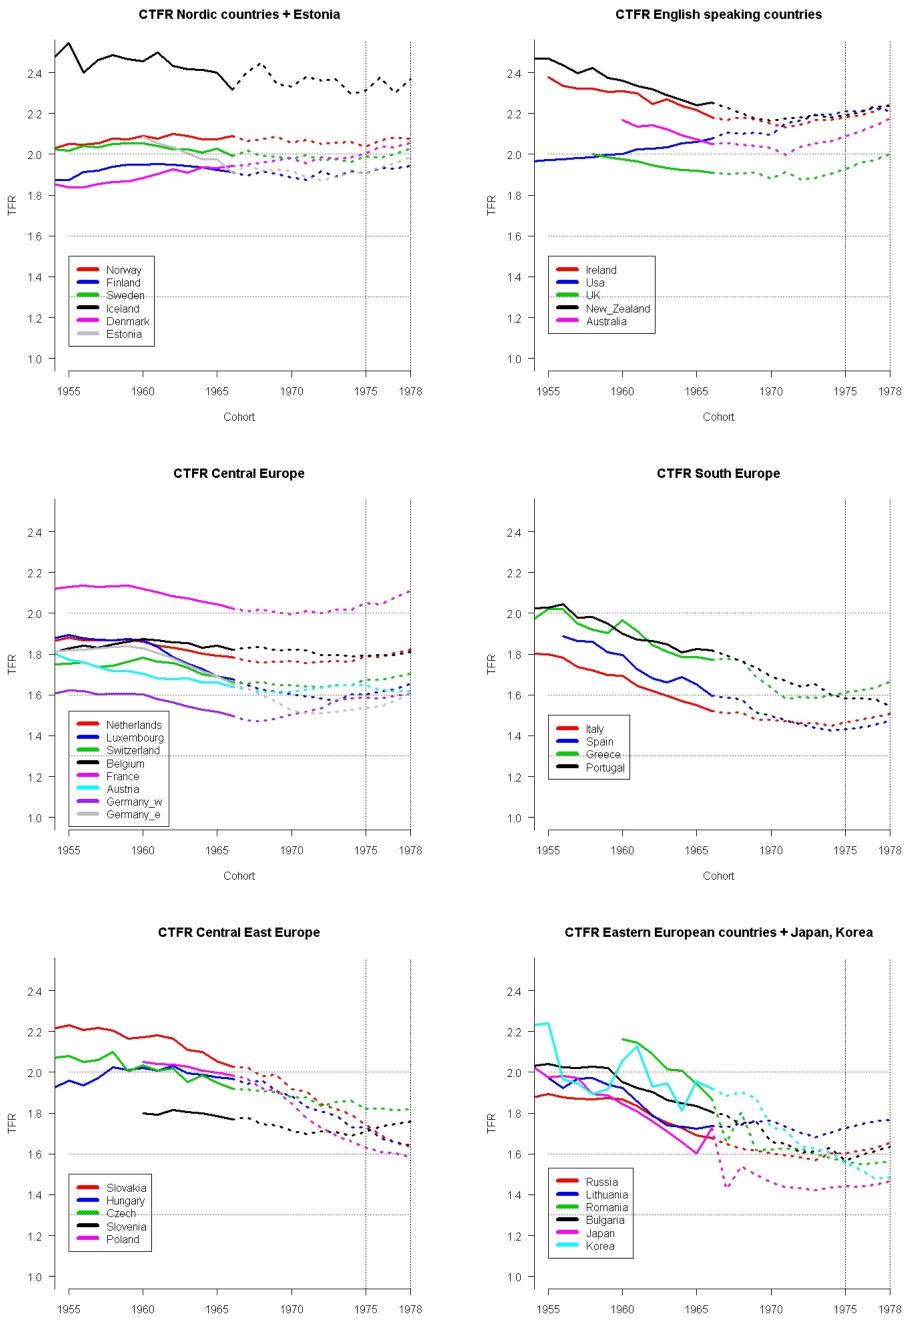

Figure 1: Recent cohort fertility trends and extrapolations based on the last five years' trends in age-specific fertility rates.

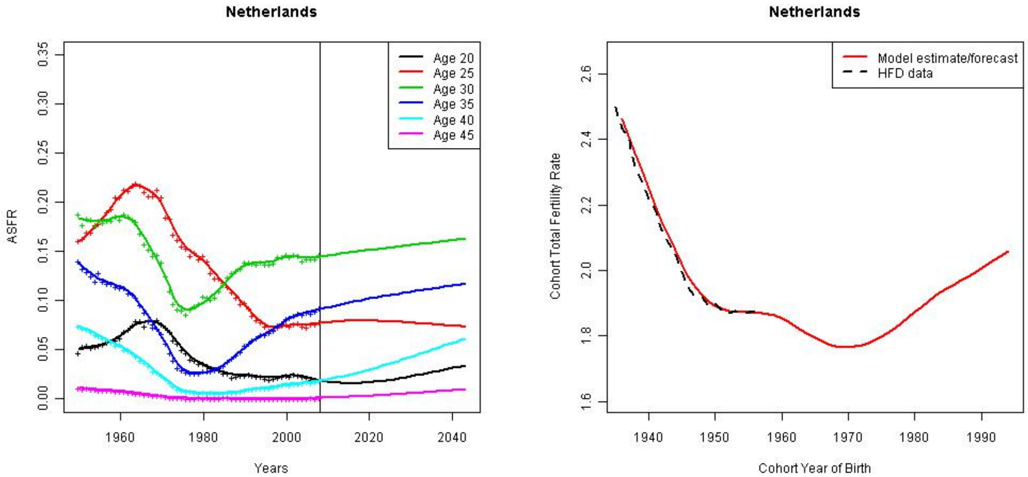

Figure 2: Estimates and forecasts of period age-specific fertility rates and cohort total fertility rates for the Netherlands, a country for which the illustrative method works well.

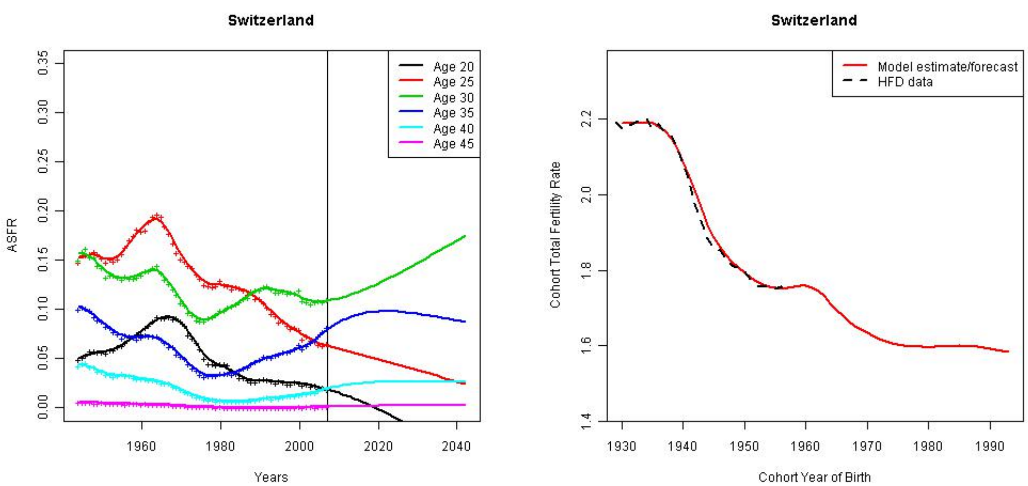

Figure 3: Estimates and forecasts of period age-specific fertility rates and cohort total fertility rates for Switzerland, a country for which the illustrative method does not work well.

Appendix

As a preliminary example of our approach, we consider a (discrete) Lexis surface of fertility rates θ_{at} over ages $A = 12, 13, ...54$ and times $T = -\eta, (-\eta + 1), \ldots, -1, 0, 1, 2, \ldots \tau$. Suppose that there are estimates f_{at} available for the $\eta + 1$ periods up through time t=0, but that the remaining τ period schedules are unavailable. The figure below illustrates for $\eta = 70, \tau = 44$.

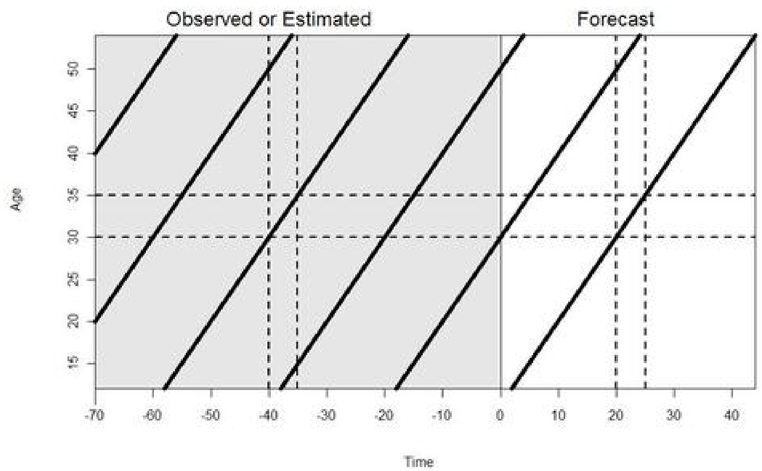

Figure 4: Illustrative Lexis diagram for the case of $\eta = 70$ and $\tau = 44$.

We convert Lexis surfaces into rate vectors by stacking periods, so that rates are sorted by increasing age within increasing period. Thus

$$\theta = [(\theta_{-70,12} \ldots \theta_{-70,54}) \ldots (\theta_{44,12} \ldots \theta_{44,54})]' \tag{1}$$

$$f = [(f_{-70,12} \ldots f_{-70,54}) \ldots (f_{0,12} \ldots f_{0,54})]' \tag{2}$$

We also define a matrix $V = [I; 0]$ such that $V\theta$ is the subset of θ for the observation period in which we have f estimates.

For illustrative purposes we model observed ASFR estimates f_{at} for the shaded area as independent normal random variables centered on the true rates θ_{at}:

$$(f_{at}|\theta_{at}) \sim N(\theta_{at}, \sigma^2) \text{ for } t \leq 0 \tag{3}$$

This implies that for the vector of rates f the log likelihood is

$$lnL(f|\theta) = -\frac{1}{2\sigma^2}(V\theta - f)'(V\theta - f) + constant \tag{4}$$

We want to find a set of rates θ that not only fits the f data well, but also produces a Lexis surface – over both the observed and forecast portions – that conforms to our prior knowledge about demographic regularities in rates. For the exploratory example in this abstract, we chose two relatively simple priors, based on age and period regularities, respectively. (We emphasize that these exploratory priors are for proof-of-concept only; much of our effort will be to make the priors more realistic and more useful.)

We construct priors from several matrix building blocks:

- G_a such that $G_a\theta$ is the 115x1 time series of rates at age a
- H_t such that $H_t\theta$ is the 43x1 period schedule of rates at time t
- D_2 such that $D_2 G_a\theta$ is a 115x1 vector of empirically estimated second derivatives of $G_a\theta$ (Girosi and King 2008)
- M_2 such that $M_2 H_t\theta$ is the 43x1 vector of residuals when projecting $H_t\theta$ onto the first 2 principal vectors of a singular value decomposition of the estimated rate matrix f_{at}

Our experimental priors are

- Time series f(a,.) should be approximately linear along horizontal slices of the Lexis surface that pass through both the observation and forecast periods
- Period schedules f(.,t) along vertical slices of the Lexis surface, in both the observation and forecast periods, should have shapes that are well approximated by the first two principal components of the matrix of estimates f_{at}.

The first priors implies that parameters sets θ are more likely when vectors of second derivatives $D_2 G_a\theta$ are 'small', or equivalently when the scalar penalty

$$P_1 = w_1\theta'(\sum_a G_a' D_2' D_2 G_a)\theta = w_1\theta' Q_1\theta \tag{5}$$

is near zero.

The second prior implies that θ values are more likely when period schedules have plausible shapes, as defined by low residuals $M_2 H_t\theta$ when projected onto components that best describe the shapes of the estimated schedules f(.,t). Thus θ is more likely when the overall scalar penalty

$$P_2 = w_2\theta'(\sum_t H_t' M_2 H_t)\theta = w_2\theta' Q_2\theta \tag{6}$$

is near zero.

Combining the quadratic penalties from likelihood and the priors produces an overall "badness of fit" measure that is identical to the log likelihood from a normal posterior distribution, namely

$$lnPost(\theta|f) = -\frac{1}{2\sigma^2}(V\theta - f)'(V\theta - f) - w_1\theta' Q_1\theta - w_2\theta' Q_2\theta + constant \tag{7}$$

which implies that $(\theta|f) \sim N(\mu, \Omega)$ with

$$\Omega = [\frac{1}{\sigma^2}V'V + w_1 Q_1 + w_2 Q_2]^{-1} \tag{8}$$

and

$$\mu = [\frac{1}{\sigma^2}V'V + w_1 Q_1 + w_2 Q_2]^{-1} V'Vf \tag{9}$$

Conditional on σ, w_1, and w_2 (which we have set at fixed values for this preliminary experiment), the maximum a posteriori (MAP) estimator of the Lexis surface parameters θ is

$$\theta^* = \mu = [\frac{1}{\sigma^2}V'V + w_1 Q_1 + w_2 Q_2]^{-1} V'V f \tag{10}$$

and the cohort TFR forecasts reported in the abstract are constructed by summing θ^* over the diagonal cohort lines that extend beyond t=0.

Notice that cohort forecasts are constructed from empirical data from the observation period, together with (explicitly stated) prior assumptions about general qualitative properties of Lexis surfaces.