

A Behavioral Gompertz Model for Cohort Fertility Schedules in Low and Moderate Fertility Populations

by
Joshua R. Goldstein
Max Planck Institute for Demographic Research
and
Princeton University

7 March 2008

Paper to be presented at PAA 2008

(This is a work-in-progress. Comments are welcome. Please use caution in citing as some of the results may be subject to revision.)

Abstract

In this paper, I re-introduce the Gompertz model of age-specific fertility. This model has been rejected by past authors as fitting poorly to cross-sectional, or period rates. However, I find that the model fits very well to recent medium and low fertility *cohort* schedules in France, Italy, and Japan. A distinct advantage of the Gompertz model is that it has a simple behavioral interpretation as the result of social diffusion of fertility behavior in a cohort competing with the fertility-depressing effects of older age.

The Gompertz model, and further refinements that include a better specification of the biological limits of childbearing, offer a means for forecasting future fertility, describing temporal change, and assessing the fertility-limiting effects of older entry into motherhood. In addition, the Gompertz function allows for traditional uses of model age-schedules such as smoothing and correction of data.

The model estimates the completed cohort fertility of French, Japanese, and Italian cohorts born in 1965 to be 2.0, 1.6, and 1.6, respectively. For France, this represents only a minor decline from earlier cohorts, but for Japan the decline in cohort fertility is marked. The model with infertility suggests that recent decline in Italy is largely due to shifts to older ages of childbearing.

Introduction

Low fertility is becoming a world-wide phenomenon, with nearly half the world's population with period Total Fertility Rates of less than 2.1 (Wilson, 2006). Some of this decline is due to real declines in completed family size, while other parts of it are due to the distortionary effects of the postponement of fertility. Period

measures are good indicators of the intensity of births at a particular time, but are poor measures of completed family size and thus of the underlying tendencies for long-term fertility levels in the absence of tempo effects. While much demographic research has been occupied with finding better period measures (e.g., Bongaars and Feeney, 1999; Kohler and Ortega, 2002a, 2002b), this paper takes the view that more cohort analysis is needed. In particular, attention needs to be paid to understanding the timing and level of cohort fertility in the form of model fertility schedules.

The criteria for a good model of cohort fertility are several:

(1) If possible the model should have a behavioral interpretation, and not be simply curve fitting. Otherwise, there is no reason to expect an arbitrary mathematical function that fits in one time and place fit another. Behavioral models not only offer the promise of more universality, they also allow insight when the model does not fit.

(2) The model should offer the possibility of forecasting the future fertility of incompletely observed cohorts.

(3) Goodness-of-fit should be adequate but should not be the only factor in model selection. For example, high order polynomials would produce arbitrarily close fits to any specific age schedule the data but would offer little confidence in the applicability to other age schedules. Also, there are different uses of model life tables. Fit at all ages might be desirable for some purposes, while fit at oldest ages might be desirable for others.

Past Studies

There is a large literature on model fertility schedules. Most of the literature has focused on period schedules. Thus, for example Hoem et al. (1981)'s rejection of the Gompertz model was based on fit to period schedules. Others have found mixed success using the Gompertz model (e.g., Wunsch 1966, Valkovics and Pollard 1992).

Recent literature using model fertility schedules using mixture models -- two Hadwiger distributions in the case of Chandola et al. (1999) and two normal distributions in the case of Peristera and Kostaki (2007) -- also relies on cross-sectional data. The problem with modeling cross-sectional data is that fertility at different ages has no necessary relationship to each other. This is particularly the case in recent decades, when young cohorts have been postponing births even more than older cohorts, creating cross-sectional age-distributions of fertility that have a different skewness than that of actual cohorts.¹

The social diffusion interpretation of cohort fertility presented here follows up on a large literature on the importance of social influence on individual fertility behavior. In recent decades, social diffusionist explanations have become more important in explaining demographic change (National Research Council 2001). Notable, Watkins and others from the Princeton European Fertility project have argued for the dominant role of diffusion in the first demographic transition. More recently, Kohler (2000, 2001) and others has written extensively on social interaction effects in low fertility populations. The use of social diffusion models to analyze cohort behavior was introduced by Hernes (1972) in his study of a cohort's entry into first marriage. The Hernes model was used for forecasting first marriage in the United States by Goldstein and Kenney (2001).

A Behavioral Gompertz Model

Although the Gompertz (1825) model has been used in the past, its use seems to be motivated entirely on the previous familiarity of the demographers with the mortality applications of the model rather than a specific behavioral storyline. As Hobcraft, Menken and Preston (1982) wrote, the Gompertz and other models such as

¹ There is some older work applying the Gompertz model to cohorts (Murphy and Nagnur, 1972; Denton and Spencer 1974; Wunsch, 1966).

the Hadwiger had formed the basis of “many attempts to fit various mathematical curves, *without a behavioral or theoretical interpretation.*”(p. 13) Page (1997) notes similarly, “Despite their great utility, however, models based simply on finding the function that best fits the data are not very satisfying – unless, that is, the function’s parameters can be identified with biological or social process that govern fertility, ...” (p. 85) With this in mind, I present an behavioral and biological basis for the Gompertz model for fertility analogous to Hernes's (1972) social diffusion model of first marriage.²

Let the total cumulative fertility of a cohort at age x be denoted $F(x)$. This is the average number of births to cohort members by age x . Let $f(x) = F'(x)$ be the density of births at age x (where F' is the age derivative of F).

The simplest form of social diffusion (Coleman, 1964:42) is proportionality, in which the rate of new adoption of a behavior is proportional to the number already practicing it. In the case of fertility, this translates to letting the fertility rate be proportional to the cumulative fertility. That is, the more children born to a cohort the higher the social pressure for everyone in the cohort to have children. To formalize this, one can write

$$f(x) = F'(x) = AF(x),$$

where the parameter A has the interpretation similar to the contagiousness of a disease. When it is larger, social pressure is more effective. The equation above is a

² Intriguingly, Hernes notes that he considered the Gompertz model for modeling first marriage. As he notes in his footnote 6, p. 181. “A model generating a logistic curve, for example, systematically gives a much worse fit, and hence can be eliminated. But one could make assumptions generating a Gompertz or doubly exponential curve.[Footnote 6: I first fitted a Gompertz curve directly to the cumulative percentages of first marriages; the model in (10)[the Hernes model] was developed later.” Two points are of interest. First, that despite its form of deductive theoretical presentation, the validity of the model for Hernes rested in its goodness-of-fit. Second, the Gompertz model would seem *a priori* in-appropriate for social diffusion model of a single-decrement events such as marriage, because it does not take into account the shrinking exposure to risk. This is the essential difference between the Hernes model, which has $P' = P(1-P)AB^x$, versus the Gompertz model which has $F' = FA B^x$.

differential equation that has as its solution the exponential function: $F(x) = \exp(Ax)$ or, equivalently, $f(x) = A \exp(Ax)$. In this form, we see that such a model is clearly inappropriate for the schedule of fertility at all ages, which does not increase without bound.

A more realistic form of social diffusion includes a time – or, equivalently for a cohort, an “age” -- effect. The model then has the form

$$f(x) = F'(x) = A(x)F(x)$$

For fertility it is reasonable to posit that the function $A(x)$ is declining. This happens for a number of reasons. Most well-known is an increase in secondary sterility with age. However, there are also social factors at play. For example, as people get older, they may get more set in their ways, and thus be more resistant to the example of others, making social pressure less effective (Hernes 1972). There are also other stories one could tell, involving cohort heterogeneity in desired family size and thus in the age at which targets are reached, that would be consistent with lower fertility at older ages. These three factors can be thought of as linked to “biological age”, “duration” of the process, and selection in a “heterogeneous population.”

Following Hernes, we let $A(x) = a \exp(-bx)$. Here the exponential is used as the most parsimonious way to describe a continuous decline without taking negative values. This gives us the model in the form

$$f(x) = F'(x) = a \exp(-bx)F(x).$$

Solving this differential equation gives

$$F(x) = K \exp\left[-(a/b)e^{-bx}\right]$$

for the cumulative fertility function. Letting, $A = \exp(-a/b)$ and $B = \exp(-b)$, we have

$$F(x) = KA^x(B^x),$$

the familiar Gompertz function. Taking derivatives of the cumulative functions gives us the fertility schedules,

$$f(x) = KA^x(B^x) \log(A) B^x \log(B)$$

or equivalently, substituting for A and B,

$$f(x) = Ka \exp\left[-\frac{a}{b} e^{-bx} - bx\right] \quad (*)$$

These functions can be fit to data using non-linear least squares estimates.³ In the applications that follow I minimize the squared residuals of the fertility schedules using equation (*). When trying to model fertility at late ages it is important to model the density schedule $f(x)$ rather than the cumulative fertility $F(x)$ in order to give appropriate weight to observed fertility at older ages.⁴

The parameters of the Gompertz model are directly interpretable. The completed total fertility over all ages is “ K ”. As in mortality modeling, “ a ” is a location parameter, which for a given “ b ” determines the location of the schedule along the age axis. Likewise, “ b ” is a time-scale parameter determining the rapidity of the process. Large “ b ” indicates a compressed distribution with little variance, as the cohort moves through its life relatively quickly. Small “ b ” indicates a slower time-scale of the process.⁵

³ For example, one can use the `nls()` function in R. Or in Excel, one can use the “solver” function. See Harris (1998) for a tutorial.

⁴ Fitting the cumulative schedule is problematic because it in effect weights younger ages much more than older ages, because the rates at younger ages contribute to cumulative fertility at all older ages, but not vice-versa. Fitting the density directly avoids this problem.

⁵ The interpretations given here differ slightly from those given earlier in the literature on Gompertz models of fertility. In particular, the understanding that “ b ” influences the time scale of the process is an insight directly from the Gompertz mortality literature. (See Canudas-Romo, ... for a recent exposition.)

The 3-parameter model is not infinitely flexible. Indeed, as Volkovics and Pollard show, the distribution has fixed skewness and peakedness.⁶ Thus, the fit of the Gompertz model across a wide number of cohorts tells us something informative about the character of cohort fertility schedules. This is also a reason that the model is often inappropriate for modeling period schedules.

Data and Analysis

In order to test the suitability of the Gompertz model across a range of circumstances I fit the Gompertz and related models to cohorts from three populations intended to contrast a range of social, cultural, political, and economic regimes. Published cohort schedules were available for France and Japan, representing populations with moderately high and nearly lowest-low fertility. These countries also differ in the share of childbearing within marriage, nearly universal in Japan but not so in France. Pro-natalist policies are strong in France, but not yet in Japan. As a third point of comparison, I include analysis of Italy cohorts in order to see the suitability of the Gompertz model for a Mediterranean pattern of late, low fertility.⁷ (In the future, efforts such as the Human Fertility Data Project (MPIDR) will make it possible to do more comprehensive comparisons.)

My analysis proceeds as follows. First, I look at the success of the Gompertz model applied to complete cohorts to judge its appropriateness for modeling moderate fertility populations. Second, I look at more recent, truncated cohorts, showing how the Gompertz model can be improved by incorporating an additional factor to account for declines in fecundity in the 30s and 40s. Finally, I show that the behaviorally-and-

⁶ As pointed out by Windsor and the appendix of Murphy and Nagur, the Gompertz distribution has the interesting property that the mode of the density occurs when a $1/e$ share of the cumulative distribution has been reached. In our case, this means that fertility rates begin to fall when cumulative fertility has reached a level of K/e . This highlights inflexibility of the model and the rigidity of its assumptions, for few would dare to predict K , as soon as a cohort fertility had peaked.

⁷ I have also fit Danish cohort data, but this analysis is not yet complete.

biologically inspired Gompertz approach produces fits that are comparable with the so-far theory-less Hadwiger model and that the Gompertz-infertility model appears to be more robust than the Hadwiger for forecasting truncated cohorts with a late pattern of childbearing.

The performance of the Gompertz Model

The fit of the Gompertz model for complete cohorts is shown in Figure 1 for French, Japanese, and Italian cohorts born in 1945. We see that the model fits all three schedules extremely well, despite the variety of fertility levels, location, and spreads in these three populations.

The lower panels of Figure 1 show the Gompertz model fitted to the cohort of 1965 – a cohort that was last observed at age 38 in France, 39 in Italy, and 35 in Japan according to these observations. Compared to the cohort of 1945, the cohort of 1965 has a lower level of fertility, less-concentrated fertility around the mode, and a later mode in all three countries. From the figures, we see that the Gompertz model seems to fit well over the peak ages of fertility but tends in these more recent cohorts to overestimate fertility at ages over about 35. This is because biological infertility increases more quickly than the exponential function assumed in the Gompertz model.

Improving the Gompertz Model with an Additional Infertility Term

A better fitting model for fertility above age 35 can be obtained by incorporating an additional term for secondary sterility. Leridon et al. (2004) suggest that a linear decline in fecundity from a level of 100% at age 33 to 0 at age 45 as a best estimate of population level declines. Letting the function $g(x)$ take the value of 1 before age 33 and 0 after age 45, with a linear decline in-between, we can model the fertility schedule as the Gompertz function with a multiplicative sterility effect. I define the “Gompertz model with infertility” as

$$f^*(x) = f(x)g(x)$$

or, written-out fully, substituting from equation (*),

$$f^*(x) = Ka \exp\left[-\frac{a}{b}e^{-bx} - bx\right]g(x).$$

The results of applying the Gompertz model with infertility as well as the non-theory-based Hadwiger function to the cohorts of 1965 and 1970 are shown in Figure 2. The 1970 cohort is shown in order to see the performance of the models when fewer ages are observed. The solid black lines are the Gompertz model, the dashed black line is the Gompertz model with infertility, and the solid grey line is the Hadwiger model.

We see that the incorporating the additional infertility term into the Gompertz model improves greatly the fit at older ages. In Italy and Japanese cohorts of 1965, the model now predicts nearly perfectly the oldest observed fertility rates. In France, the infertility term seems to over-compensate slightly. Speculatively, a possible cause of this unusually high late fertility in France could be pro-natalist policies toward 3rd and higher order births.

Comparison with the Hadwiger Model

The Gompertz model with infertility appears to fit moderate and low fertility cohort data very well. Moreover, it has the advantage of having a plausible behavioral and biological basis. It is difficult to distinguish between the age-related declines in the original Gompertz model and those in the extended model with infertility, since the downward pressure on fertility at older ages captured in the declining exponential term of the Gompertz model is due to several factors, including declining fecundity with age. Still, I would argue that Gompertz model with infertility a firmer behavioral

and biological basis than mathematical curves such as the Hadwiger, Gamma, Beta, which are used with appeal only to goodness-of-fit.

The Hadwiger model does not provide better fits than the Gompertz with infertility model. Indeed, in Figure 2 we see that for Italian cohorts the Hadwiger model performs poorly, much worse than the Gompertz model with infertility, and barely better than the unadjusted Gompertz. For Japanese cohorts, the Hadwiger model gives nearly identical estimates to the Gompertz model with infertility. For France, the performance is in-between. Overall, there seems little reason to prefer the Hadwiger model on the basis of goodness-of-fit. In particular, as we see in the Italian case, the Hadwiger model can seriously overestimate late fertility.

Figure 3 shows goodness of fit comparisons of the three models across a much larger number of cohorts in Japan, Italy, Denmark, and France. The measure used is the root-mean-squared-error between the estimated and observed age-specific rates for each cohort. For example, the value of about 0.004 for the Danish cohort of 1955 means that the average error (in root mean square terms) for each age-specific fertility rate was about 4/1000. Age-specific fertility rates are on the order of 0.1, so all of the models are performing very well, with no single model clearly superior. In general, the Gompertz with infertility fits better than the Gompertz without fertility (with the two estimates being identical when fertility is observed only before age 33 such is the case for the cohort of 1970 in Japan.). Goodness-of-fit by this measure does not show the Hadwiger to be superior to the Gompertz with infertility, or vice versa. The Gompertz with infertility model does slightly better for the completed cohorts (before about 1960) in France and Japan, where as in Italy neither is preferable. For truncated cohorts (after about 1960) the Hadwiger does somewhat better (except in Italy) in

fitting rates at the observed ages, but as we have seen it has the tendency of over-predicting fertility at older ages that are “out-of-sample.”

From this more comprehensive comparison, we both the Hadwiger model and Gompertz model with infertility tend to fit better to observed fertility rates than the Gompertz model without the additional infertility term. However, sum-of-squared-residuals calculations are based only on observed fertility rates, and do not factor in the tendency for the Hadwiger and Gompertz-without-infertility models to overpredict fertility at yet-to-be-observed ages.

Applications of the Gompertz Model

I now show some applications of the Gompertz model of age-specific cohort fertility. First, I show forecasts of the completed cohort fertility for Italy, Japan, and France. I then show that the infertility model can be used to estimate a counter-factual un-met need for children that would have been born if it were not for the consequences of infertility.

Figure 4 shows the forecast completed cohort fertility for the cohorts born from 1945 to 1965 in Italy, France, and Japan. The Gompertz and Gompertz-with-infertility models are used to future fertility of cohorts. The forecast complete cohort fertility is calculated by adding the fertility rates for the observed ages to the fertility rates for the forecast ages. (For example, for the cohort born in 1960, for which the last observation is in 2000, I calculate the forecast complete cohort as the sum of the observed fertility rates up to age 40 and the forecast fertility rates for older ages.)

Taking the Gompertz-with-infertility as the preferred model, we see that there has been a downward trend in cohort fertility in France and Japan since the mid-1950s, and in Italy over the entire period. Cohort fertility in France remains over 2.0 in 1965, whereas in Italy and Japan it is about 1.6. The rate of decline in cohort is

about twice as fast in Japan as in Italy, which in turn has a rate of decline about twice as fast as France.

The dashed lines in Figure 4 show the forecast cohort TFR based on the Gompertz model without the extra “infertility” term. One interesting interpretation of the results is to take the difference between the predicted completed cohort fertility between the two models as a measure of the “unmet need for children” created by the rapid decline in fecundity at older ages. This, in effect, takes the unmodified Gompertz model as an estimate of the fertility that the cohort would have had without the additional age-related infecundity. “Waiting too long” appears to account for a shortfall of about 0.1 children in the cohort of 1965 in these three countries. We see that if it were not for this effect, cohort fertility in France and Italy there would have been no decline in France and Italy from 1955 to 1965. Furthermore, it appears that the declines in cohort fertility in Italy before 1955 are not due to an increasing share of fertility be delayed to older ages but rather to a decline at all ages. Finally, biological limits of childbearing appear to play a relatively small role in Japan. Even without, cohort fertility would have declined rapidly from the cohort of 1955 onwards.

Future research

Further research has several directions. The first is to expand the range of fitted populations, incorporating cohort fertility from the United States as well as the lowest-low fertility populations in Eastern Europe.

The second line of research is to refine further the model of sterility. In particular, it would seem that a model that did not begin and end so abruptly would be more appropriate.

A third line of research could involve testing some of the assumptions of the social diffusion framework. Does a shock in fertility when a cohort is aged x appear to influence its subsequent fertility older ages? This kind of thinking also suggests that statistical analysis of cohort fertility as a time-series, albeit a short one, could be potentially revealing.⁸

Third, a closer comparison of the Gompertz models with other standard mathematical functions in use such as the Hadwiger, Gamma, and Beta functions may be productive, as well as the elaboration of behavioral assumptions that might be expected to produce these distributions. For example, the Gamma and Hadwiger functions both result as the result of waiting times in stochastic processes.

Finally, I will look at the variability in the model parameters across time both in order to provide a richer description of past fertility change and to inform forecasts of future fertility.

Conclusion

The Gompertz model appears to describe cohort age-schedules of fertility in low and moderate fertility populations. It also appears to be useful for forecasting the future fertility of cohorts that are still young, particularly if additional account is taken of the decline of fecundity with age. If the results presented here are found to be more broadly applicable, I would suggest that such cohort forecasts form part of the basis, along with tempo-adjusted period measures, and the analysis of age-specific trends across cohorts for forecasting fertility in low fertility populations.

(Reference list to be added). (Figures in separate file.)

⁸ Ron Lee suggested such an approach to me concerning the Hernes model. Likewise, the Gompertz model can be written as a recursive equation $F(t+1) = b_0 + b_1 F(t) + b_2 F(t) \cdot \log(F(t))$. Adding an error term to the right hand side, and estimating the coefficients by regressing $F(t+1)$ on $F(t)$ is potentially a promising approach. Reformulation is required, however, in order to make $F(t)$ monotonically increasing.

Figures for “A Behavioral Gompertz Model for Cohort Fertility Schedules in Low and Moderate Fertility Populations” by Joshua R. Goldstein to be presented at PAA 2008







Figure Captions







Figure 1. Age-specific cohort fertility schedules for France, Italy and Japan for women born in 1945 and 1965, with fitted estimates from the Gompertz model.(This figure shows that the Gompertz model fits well to a range of completed, moderate fertility cohorts but that it over-estimates fertility at older ages for cohorts that have not yet completed their childbearing.)



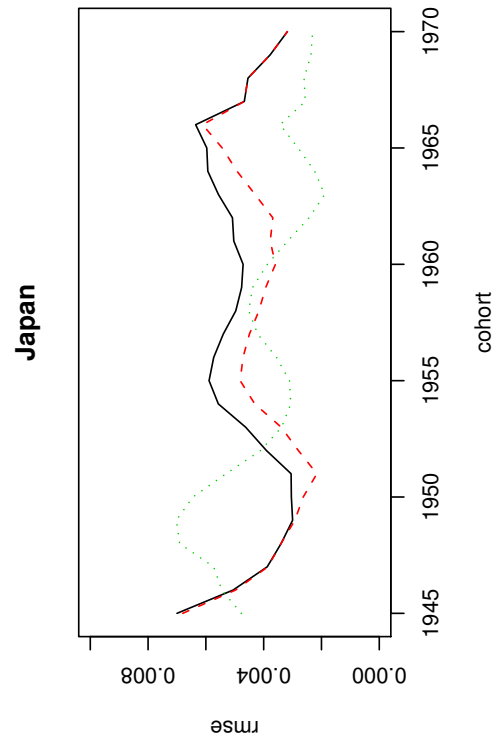

Figure 2. Age-specific cohort fertility schedules for France, Italy and Japan for women born in 1965 and 1970, with fitted estimates from the Gompertz model, the Gompertz-with-infertility model, and the Hadwiger model.(This figure shows that the Gompertz model with infertility accounts well for declines in fertility for truncated cohorts.)

Figure 3. Goodness-of-fit measures of the Gompertz, Gompertz-with-infertility, and Hadwiger models applied to single year cohorts from 1945 to 1965 for France, Italy, and Japan. (Danish estimates are preliminary.)

Figure 4. Forecast cohort completed fertility according to the Gompertz and Gompertz-with-infertility models. (These estimates include estimated data only for unobserved ages. The difference between the two models can be interpreted as the “unmet need” for children that results from fertility postponement.)







Forecasts of Completed Cohort Fertility

