Using Discrete-time Event History Fertility Models to Simulate Total Fertility Rates and Other Fertility Measures

Jennifer Van Hook*

Claire E. Altman*

*The Pennsylvania State University

Word Count: 6,113

Table Count: 5

Figure Count: 6

This research was supported by an infrastructure grant provided to the Population Research Institute by the National Institutes of Health (2R24HD041025).

Using Discrete-time Event History Fertility Models to Simulate Total Fertility Rates and Other Fertility Measures

Abstract

Event history models, also known as hazard models, are commonly used in analyses of fertility. Such models offer many advantages over more simplistic OLS or Poisson models of children ever borne. One drawback of event history models is that the conditional probabilities estimated by event history models do not readily translate into summary measures, particularly for models of repeatable events, like childbirth. In this paper, we describe how to translate the results of discrete-time event history models of all births into well-known summary fertility measures: simulated age- and parity-specific fertility rates, parity progression ratios, and the total fertility rate. Our method incorporates all birth intervals, but permits the hazard functions to vary across parities. It also can simulate values for groups defined by both fixed and time-varying covariates, such as marital or employment life histories. We demonstrate the method using an example from the National Survey of Family Growth and provide an accompanying data file and Stata program.

Keywords: Event History, Fertility Measures, life tables, NSFG

Using Discrete-time Event History Fertility Models to Simulate Total Fertility Rates and Other Fertility Measures

Event history models, also known as hazard models, are commonly used in analyses of fertility. Provided that retrospective or prospective fertility histories are available, such models offer many advantages over more simplistic OLS or Poisson models of children ever borne. Event history models potentially provide richer information about the age at first birth and subsequent birth spacing because they account for both the occurrence and timing of births. They also appropriately handle right censoring, thus permitting analyses to include women at all stages of their reproductive lives rather than just older women with completed fertility. Finally, they permit researchers to incorporate critical time-varying predictors, such as employment, educational, and marital statuses, into analyses.

One drawback of event history models, however, is that they are difficult to translate into meaningful information about lifetime childbearing or birth spacing. In fact, many studies using event history analysis to model births analyze parities separately or first birth only. In this paper, we show how to convert the results of pooled discrete-time event history models of all births into well-known fertility measures, including simulated age- and parity-specific fertility rates, parity progression ratios, and total fertility rates (Newell 1990). Age- and parity-specific fertility rates provide detailed information on the number and timing of births across the reproductive life course. Parity progression ratios (PPR) are important indicators of spacing and stopping behaviors and childlessness. Finally, the total fertility rate is the most commonly-used measure of fertility. It provides an estimate of the total number of children the average woman will have in her reproductive lifetime if she experienced all current age-specific fertility rates. Although others have developed generalized methods for estimating multistate lifetables based on event history models (e.g., Cai et al. 2010; Lee and Rendall 2001), we show how to apply these ideas to fertility histories, an extension that has not yet appeared in the literature to the best of our knowledge.

The key advantage of the method we present is that it permits researchers to compare groups on these common-used fertility measures while statistically controlling for differences on other characteristics associated with childbearing. In doing so, our method provides a comprehensive picture of women's fertility experiences across the entire reproductive life course. This goes beyond most descriptive and multivariate fertility analyses, which tend to be confined to births of specific parities (e.g., first births) and only rarely consider all births simultaneously (Guzzo and Hayford 2011; Schellekens 2009; Carter 2000; Brand and Davis 2011; Goldstein, White, and Goldstein 1997; and White et al. 2008). Additionally, our method can simulate values for groups defined by both fixed and time-varying covariates. This is helpful for assessing how life course trajectories can delay (or accelerate) births and eventually reduce (or increase) completed fertility. For example, one could estimate the effects of delayed marriage on completed fertility by comparing simulated TFRs for women who married at age 35 with their peers who are similar on all characteristics except that they married at age 20.

We first describe the data and measures used in our examples. Next, we explain how single-event event history models are related to single-decrement life tables using an illustration drawn from the National Survey of Family Growth (NSFG). This section provides an orientation on the linkage between event history analysis and single-decrement life tables, but can be skipped by readers who are already familiar with these methods (see also Singer and Willett (2003) and Teachman and Hayward (1993) for more rigorous treatments of these methods). We next discuss the main contribution of this paper, the extension of single-decrement methods to multiple births. Again, we illustrate the method with an example from the NSFG. We estimate discrete-time event history models because of their clear analytic linkages to life tables. Nevertheless, it would be possible to develop a similar approach for other kinds of event history models (e.g., Cox Models). Also, although we discuss the method as relevant for fertility, it could be applied to any event history analysis in which multiple events per individual are modeled (e.g., marriages, arrests, poverty spells, etc.).

Data and Measures

To simulate fertility measures from event history model estimates, it is necessary to model the occurrence and timing of births. Therefore, a fertility history recording the precise timing of each birth is required (ideally month and year), such as is available in the National Survey of Family Growth or the Demographic and Health Surveys. In data sources like these, older women are able to report a complete fertility history, but younger women may report a partial or unfinished fertility history. Younger women may have more children in the future but they are right-censored by the survey. This is perfectly acceptable because event history models take into account the fact that older women spent more time at risk of childbearing than younger women.

To generate the examples provided in this paper, we used the continuous collection of the National Survey of Family Growth (NSFG 2006-2010). The NSFG is a nationally representative cross-sectional survey of reproductive age women in the United States conducted by the National Center for Health Statistics (NCHS). As the primary fertility survey in the United States, the NSFG collects information on the respondent's race/ethnicity, complete fertility histories, and other socio-demographic controls such as marriage history (CDC 2009 website). When available, the recoded variables, rather than the raw variables, were used in the analysis as recommended by the National Center for Health Statistics (Lepkowski et al. 2010).

We limited the sample to U.S.-born non-Hispanic black and white women with full information on the sample variables. After listwise deletion, the final sample included 7,164 women. We then organized the analytic data file into "person-year" records, with one data record for each year of age a woman is at risk of having a birth (i.e., from age 14 [the earliest reported age at first birth] to age 45 or age at censorship/survey date, whichever is younger).

The dependent variable is a dichotomous indicator of whether or not a woman had a single live birth in the age interval, x to x+1. Time-constant predictors assume the same value across all person-year records for an individual. Our analyses include race (non-Hispanic black and white) and religion (no religion, Catholic, Protestant, and other) as time-constant predictors. Values for time-varying predictors may change across ages (across person-year records) for a given individual. We included three time-varying predictors in our models: age, marital status, and educational attainment. Age is coded in five-year categories: 14-19; 20-24; 25-29; 30-34; 35-39; and 40-45. Marital status is a dichotomous indicator of whether or not a woman in married at the beginning of the interval; any status other than married is coded zero. Years of education are the number of completed years of school a woman has received at the beginning of each person-year interval. Women are assumed to begin schooling at age 6.

Section 1: Hazard Models and Life Tables for Non-Repeated Events

Event history models are conceptually linked to life tables (for examples see Allison 1989; 1995). Life tables were first developed by demographers to estimate how mortality reduces the size of a cohort as it ages. Event history models of non-repeated events (like first births) are directly analogous to single-decrement life tables (Singer and Willett 2003). In both, cohorts are conceptualized as being exposed to risk of an event. The age pattern of this risk is referred to as the hazard function in an event history analysis and notated as q_x in a life table, where q_x reports the probability of the event at age x. Once cohort members experience the event in question (such as a first birth), they exit the risk pool, leaving behind an ever-diminishing group at risk. The depletion of the cohort by age is referred to as the survival function in an event history analysis, and notated as l_x in a life table. Finally, the number of events experienced by the cohort at each age x is denoted as d_x in a life table.

The key advantage of event history models over life tables is that they estimate the associations of predictors with the timing and occurrence of events. That is, they allow us to estimate how the hazard of the event differs across subgroups (such as marital status) or varies across values of a covariate (such as education) in addition to standard errors around the estimates. Event history model coefficients specifically estimate the association of predictors with the *hazard* of the event, which is the conditional probability of the event occurring in a narrow time window (typically a year) given that it has not already occurred.

For example, in a model of first births, the hazard is the probability of a first birth occurring at age x given that the woman has not yet had a birth. Because people are often interested in lifetime patterns of fertility (e.g., median age at first birth), researchers often use event history life tables to transform event history coefficients into simulated birth histories for selected groups.

To illustrate, we estimated a simple discrete-time hazard model (Allison 1989) of first birth based on fertility histories in the National Survey of Family Growth, and then used the results to generate a life table of first birth. The hazard model of first birth is a logistic regression model predicting whether a first birth occurred in each person-year interval. The analytic sample is confined to person-year records falling within the first birth interval (i.e., from age 14 until age at first birth, age 45, or censorship by the survey, whichever comes first). The estimated model coefficients are:

logit(pr[first birth age x|no earlier birth]) = -3.88 + 0.50(age20-24) + 0.15(age25-29) + 0.09(age30-34) – 0.53(age35-39) – 2.31(age40-45) + 0.75(black) + 0.02(education) + 1.85(married) – 0.01(no religion) – 0.21(Catholic) – 0.21(other religion).

Using race as an example, the results suggest that black women are 2.11 times as likely to have a first birth as white women (exp(0.75)=2.11). To express this difference in more concrete terms, we used the model estimates to generate a life table for two hypothetical groups: black and white never-married, Protestant women with 12 years of education. This involved two steps. First, we generated predicted hazards for these two groups across all ages (shown in the q_x columns in Table 1). These were calculated by substituting values (i.e. years of education and religion category) into the event history model to obtain predicted log-odds, and then transforming the predicted log-odds into predicted hazards (i.e., q_x)[i]. Second, we used the predicted hazards to simulate the proportion of women estimated to remain childless at age x (shown in the l_x column in Table 1 and graphed in Figure 1). At very young ages (e.g., 14) all are childless, so l_{14} is set equal to 1.0. The expected proportion having births at each age x is given in the d_x column, and it is calculated as the product of l_x and q_x: $d_x = l_x * q_x$. As the cohort ages from one year to the next, the proportion remaining childless declines by the proportion who had a birth the year before: $l_{x+1} = l_x - d_x$. Finally, the proportion simulated to have had a first birth by age x is $1 – l_x$.

The results show that if women experienced the predicted hazards for never-married, Protestant women with 12 years of education, the simulated median age of first birth (i.e., the age at which half the cohort has had a first birth) would be 36 for whites and 24 for blacks. Additionally, among white women, 14.5% (1-.855) would have a birth prior to age 20, 30.9% by age 25, and 53.0% by age 45, leaving 47.0% childless by age 45. Among black women, the respective percentages would be much higher: 27.8%, 53.4%, and 79.1%, leaving only 20.9% childless by age 45.

Section 2: Hazard Models and Life Tables for Repeated Events

The methods described above work well for non-repeatable events like first births, but cannot be directly applied to repeatable events, like all births. The reason is that, in hazard models of multiple births, the underlying life table model is no longer a single-decrement life table. Rather, the underlying life table has multiple, sequenced events: women can have multiple births but do not enter the “risk set” of having the next higher order birth unless they have already had earlier parity births. Because of their underlying complexity, the hazard in models of multiple births has an even narrower interpretation than in models predicting only first births. It is conditional not only on not already having had a j^{th} birth, but also on already having had the previous (j-1^{th}) birth. The coefficients therefore provide an estimate of a group's relative risk of having a birth *among those in the same age/parity window*. Models of repeated events can be used to generate a large set of predicted hazards of having a birth by both age and parity. However, translating these predicted conditional probabilities into overall assessments of the timing or total number of births is not straightforward and can limit their utility in research. In what follows, we outline a methodology for doing this and provide an empirical example. The method involves three major steps: (1) model estimation, (2) generation of predicted hazards for selected groups, and (3) generation of fertility life tables for multiple births. The Stata programs and data file used to generate the example are available on-line at www.[link to be determined].

Step 1. Model Estimation

We first use a discrete-time event history model to model the occurrence and timing of births. As described by Allison (1989; 1995), this model uses logistic regression to estimate the log-odds of a birth occurring in each person-year interval as a function of the woman's age category, and non-time-varying and time-varying characteristics. In the case of multiple events, one may pool all birth intervals in the same analysis providing appropriate steps are taken to account for the clustering of births within individual women (Allison 1995; Cleves et al. 2008)[ii].

There are several possible ways to handle multiple birth intervals. At the simplest level, one can estimate an additive model that includes birth interval (i.e., parity) as one of the time-varying independent variables:

Additive Model: $\text{Logit}(b_x=1) = a + j_x + \underline{A}_x B_1 + \underline{C} B_2 + \underline{C}_x B_3$

where:

b_x	= birth occurred at age x
j_x	= birth interval j at age x (range: 1-J). The j^{th} birth interval starts in the person-year following the $j-1^{th}$ birth, and ends in the person-year of the j^{th} birth or censorship.
$\underline{A}_x$	= vector of dummy variables indicating membership in 5-year age category at age x (note: other age intervals or age functions are possible)
$\underline{C}$	= vector of time-fixed control variables
$\underline{C}_x$	= vector of time-varying control variables

This model allows the levels of the underlying hazard function to differ across birth intervals, but assumes that the shape of the hazard function (i.e., the age pattern given by B_1) does not vary by birth interval.

To partially relax this assumption, one can include interaction terms between birth intervals and age categories:

Partially Interactive Model: $\text{Logit}(b_x=1) = a + j_x + \underline{A}_x B_1 + \underline{C} B_2 + \underline{C}_x B_3 + (j_x{*}\ \underline{A}_x) B_6$

This model allows both the *level* and *shape* of the hazard function (i.e., age pattern of the hazard of childbearing) to vary by birth interval. For example, if first births were heavily concentrated around age 25 but subsequent births were more flatly distributed across subsequent ages, this model would detect these parity differences. A chi-square test indicates whether the partially interactive model fits the data better than the additive model (where $\chi^2 = -2LL^{additive} - (-2LL^{partially\ interactive})$ and $df = df^{additive} - df^{partially\ interactive}$).

Finally, to allow the effects of all covariates to vary across birth intervals, one can estimate a fully interactive model:

Fully Interactive Model: Logit(b_{xj}=1) = $a_j + \underline{A}_x B_{1j} + \underline{C} B_{2j} + \underline{C}_x B_{3j}$, *for each jth birth interval.*

This model is estimated separately for each birth interval; higher birth intervals (e.g., 4th or higher) can be combined if sample sizes become too small. This model is the most flexible of the three described. If racial-ethnic groups had similar first birth, but different spacing, and stopping patterns, for example, this model would capture these differences. One can use a chi-square test to test whether the fully interactive model fits better than the partially interactive model (where $\chi^2 = -2LL^{\text{partially interactive}} - (-2LL^{\text{fully interactive}})$ and df = $df^{\text{partially interactive}} - df^{\text{fully interactive}}$). The -2LL and degrees of freedom of the fully-interactive model is obtained by summing the -2LL and degrees of freedom across all of the separate birth interval models (Singer and Willett 2003: 560-561).

To illustrate, we estimated the additive, partially-interactive, and fully-interactive models for our NSFG sample while adjusting for the clustering of observations within individual women. The models are weighted with a normalized sampling weight[iii] to account for NSFG's stratified sampling design. Because of low cell sizes at higher parities, we top-coded our measure of parity at 3-or-more births. The results are shown in Table 2. The chi-square test clearly indicates that the fully-interactive model fits the data significantly better than the partially-interactive model, which in turn, fits the data better than the additive model. This is also evident by the fact that most of the age-parity interactions are significant in the partially-interactive model, and several of the coefficients vary widely across parities in the fully-interactive model.

Focusing on the effects of race and marital status in the best-fitting model (fully-interactive), black women are 2.1 times (exp(0.75)) more likely than white women to have a first birth and 1.4 times (exp(0.34)) as likely to have a third-or-higher-order birth, but are no different from white women in the likelihood of having a second birth. Additionally, married women are 6.3 times (exp(1.85)) as likely to have a first birth as unmarried women, 2.3 times (exp(0.85)) as likely to have a second birth, 1.3 times (exp(0.23)) as likely to have a third-or-higher-order birth.

Step 2. Generating Predicted Hazards

After fitting the hazard model, one can use the results to obtain predicted hazards (often called conditional probabilities) of having a birth at each year of age and parity for selected groups of interest (i.e., groups k = 1 to K). For example, if we were interested in obtaining predicted values for women with varying marital status histories, one group might be women who never married while another group might be women who married at age 25. To hold other factors constant, all other predictors must be set to the same values across all groups (often the mean). In general, the predicted probability of a birth at age x, birth interval j, for group k is obtained by inserting values for C and C_x (to identify groups 1to K) for each combination of age (x) and parity (j) into the estimated model.

The first step is to calculate the predicted log-odds for all values of age, parity, and groups (x, j, and k, respectively). For example, for each of the fully-interactive models estimated for each j^{th} birth interval, values for A_x, C, and C_x need to be replaced with values corresponding to each combination of x and k:

$$\text{Predicted logged-odds}_{xjk} = a_j + \underline{A}_x B_{1j} + \underline{C} B_{2j} + \underline{C}_x B_{3j}$$

The second step is to convert the predicted log-odds to predicted probabilities for each combination of x, j, and k;

$$q_{xjk} = \exp(\text{predicted log-odds}_{xjk}) / [1-\exp(\text{predicted log-odds}_{xjk})]$$

As an illustration, Table 3 displays the predicted hazards (q_x) for the first and second parities for white, Protestant women with 12 years of education who married at age 25. These predicted hazards serve as the key inputs for the other components of a fertility life table, which are discussed in the next section. The predicted hazard of first birth at age 20 is .042 and is obtained from the fully-interactive first parity model as follows:

Predicted log-odds = -3.88 + 0.50(age 20-24) + 0.02(12 years of education) = -3.14;

Predicted hazard = exp(-3.14)/(1+exp(-3.14)) = 0.042.

At ages 25 and older, the women in the hypothetical cohort are married, so the calculations are adjusted accordingly and the predicted hazard increases substantially:

Predicted log-odds = -3.88 + 0.15(age 25-29) + 0.02(12 years education)+1.85(married) = -1.63;

Predicted hazard = exp(-1.63)/(1+exp(-1.63)) = 0.164.

These calculations were repeated for each combination of age (x) and parity (j). The predicted hazards in Table 3 are the same for each 5-year age block (e.g., 14 to 19, 20 to 24, etc.) because we modeled age effects using a set of dummy variables for 5-year age categories. If we had modeled age as a continuous variable (or as a quadratic, for example), we would obtain different q_{xjk} values for each single year of age.

Step 3. Generating Fertility Life Tables for Multiple Births

To convert the predicted hazards to fertility measures like the total fertility rate, it is necessary to construct a sequential multi-decrement life table. The life table models report by parity how many children women in a synthetic cohort are likely to have at each age throughout their entire reproductive lifetime, *if they experienced all of the predicted probabilities produced by the event history model.* If the predicted probabilities pertain to a particular group (e.g., women who never married), then the life table will produce estimates of the simulated number and timing of births for women with these characteristics.

As noted above, Table 3 displays a portion of a fertility life table for the first two parities from age 14 to 45 for never-married, Protestant white women with 12 years of education. The complete life table is much larger with 11 parities (the maximum number of births observed in our sample).

The **x column** denotes single year of age for the synthetic cohort.

The **q_{xj} columns** show the predicated hazard of having a j^{th} birth, that is, the probability of having a j^{th} birth at age x, given that the j^{th} birth has not yet occurred but that the $j\text{-}1^{th}$ birth has occurred. For example, an 18 year old has a 0.026 probability of having a first birth if she has never had a birth before, but a 0.139 probability of having a second birth if she already had a first but not yet a second. These values are the predicted probabilities generated from the discrete-time event history models (i.e., q_{xjk}). *They are the only inputs required to produce the life table.*

The **l_{xj} columns** indicate the proportion of the entire synthetic cohort at risk of having the j^{th} birth. $l_{14(j=1)}$ is set equal to 1.0 because the entire cohort at age 14 is "at risk" of having a first birth (because none of the women in our sample had a birth at an earlier age). The proportion at risk of having a first birth declines as 1^{st} births occur. Additionally, the proportion at risk of a second birth is smaller because it includes only those who have already had a 1st birth but not yet a 2nd birth.

The **d_{xj} columns** give the unconditional probability of a jth birth at each age x. In other words, it is the proportion of all women in the synthetic cohort who have a jth birth at age x.

These three columns are related in the following ways. For the first parity:

For x = 14,

$l_{14,1} = 1.0$

$d_{14,1} = 1.0 * q_{14,1}$.

And, for x > 14,

$d_{x,1} = l_{x,1}q_{x,1}$

$l_{x+1,1} = l_{x,1} - d_{x,1}$

And for subsequent parities j:

For x = 14,

$l_{14,j} = d_{14,j-1}$

$d_{14,j} = (d_{14,j-1} /2) * q_{14,j}$

And, for x > 14,

$d_{x,j} = [(d_{x,j-1} /2) + (l_{x-1,j} - d_{x-1,j})]\, q_{x,j}$

$l_{x+1,j} = d_{x+1,j-1} + l_{x,j} - d_{x,j}$

Note: Women who had a j-1th birth in the past year ($d_{14,j-1}$ *or* $d_{x,j-1}$*) are at risk of a jth birth for only half a year because earlier births occurred on average mid-way through the age interval.*

Importantly, the d_{xj} columns can be used to estimate various summary measures of fertility. If we sum d_{xj} across all ages (summing down columns), we simulate the proportion of women in the synthetic cohort who ever had a j[th] birth in their reproductive lifetime, or the simulated parity-specific total fertility rate:

$$TFR_j = \sum_{x=10}^{49} d_{xj}$$

We can construct a series of simulated parity progression ratios (i.e., the probability of progressing to the next higher parity) from these:

$$PPR_j = \frac{TFR_j}{TFR_{j-1}}$$

where $TFR_0 = 1$.

For example, the PPR for parity 1 is the probability of ever having a first birth; PPR for parity 2 is the probability of ever having a 2[nd] birth among those who had a 1[st] birth, and PPR for parity k is the probability of having a k[th] birth among those who had k-1 births.

If we sum d_{xj} across all parities j (summing across rows), we obtain the simulated proportion of women having a birth of any parity age at x, or age-specific fertility rates:

$$f_x = \sum_{j=1}^{\infty} d_{xj}$$

Finally if we sum f_x, we obtain the simulated total fertility rate, or the expected number of children born to women if they experienced all of the predicted age- and parity-specific birth hazards:

$$TFR = \sum_{x=10}^{49} f_x$$

Standard Errors. Although alternative approaches have been used (Lynch and Brown 2010 and Lee and Rendall 2001), it is increasingly common to sample with replacement from the data to obtain bootstrapped standard error estimates of life table components (Cai et al. 2010; Rendall et al. forthcoming; Poi 2004). Here, we used Stata's bootstrap routine to draw 500 replicate samples with N observations (i.e., the same number as in the full sample) from the data. We sampled by person-level clusters (i.e., taking all person-year observations for each sampled individual) rather than sampling person-year records independently. We estimated the models for each of the 500 replicate samples and used the results to estimate the life table components. The mean across the 500 replicates provided the expected value of the life table components and the standard deviation provided an estimate of the standard error. Figure 2 plots the estimated TFR for black women by sample replicate, which clearly displays random variation across replicates. However, the average TFR and standard deviation stabilize by about the 250th replication, as shown in Figures 3 and 4, respectively.

Example: Predicted fertility measures by race and marital status. We estimated simulated fertility measures for black and white women while holding the other time-constant predictor (religion) at its sample mean and the time-varying predictors (marital status and education) at their age-specific sample means. We summarize the results in Figure 5 and Table 4, but also provide the simulated age- and parity-specific fertility rates in an on-line appendix (included here for reviewers).

Simulated age- and parity-specific fertility rates for black and white women are displayed in Figure 5. Both groups show the typical age pattern with the highest fertility rates occurring for women aged 20 to 34. However, black women exhibit much higher fertility than their white peers, especially at younger ages and at higher parities. These patterns also appear in the simulated total fertility rates and parity progression ratios shown in Table 4. For example, nearly all (95%) black women are simulated to have a birth in their lifetime compared with only 78 percent of white women. The total fertility rates and parity progression ratios further indicate that black women are more likely to continue to have subsequent births at all parities. Overall, even after controlling for marital status history, educational attainment, and religious affiliation, black women are simulated to have nearly one more birth than white women in their lifetimes (2.71 versus 1.76 births). The standard errors suggest that this difference is statistically significant; it is much larger than twice the standard error of the difference (2 * SE-diff = 2 * sqrt($.09^2$+$.04^2$) = .20).

To illustrate the method for a time-varying predictor, we generated simulated fertility measures for four groups of women with various marital status histories: (1) never married, (2) married age 20+, (3) married age 25+, and (4) married age 25-34. Again, we hold the remaining time-constant predictors (race and religion) at their sample means and the time-varying predictor (education) at its age-specific sample mean. Simulated age-, parity-specific fertility rates for the four marital status groups are displayed in Figure 6. Never-married women exhibit very low and flat age patterns, while the other groups exhibit elevated fertility rates following marriage. Women who marry at age 20 exhibit much higher lifetime fertility than women who marry at age 25 because as married women they are exposed to the highest age-specific fertility rates between ages 20 and 24. But the fertility pattern of women who were married age 25 and older is nearly identical to the pattern among those who were married between the ages of 25 and 34, largely because fertility rates are low for all groups after age 35. These patterns appear in the simulated total fertility rates and parity progression ratios shown in Table 4. For example, more than half (56%) of never-married women are simulated to ever have a birth, much less than ever-married women. Among ever-married women, age at marriage is moderately related to childlessness. Nearly all (98%) women who marry at age 20 are simulated to ever have a birth, compared with 95 percent among those marrying at age 25, and 91% among those married between the ages of 25 and 34. Interestingly, among women who already had two births, the probability of progressing to third or higher-order births is similar across all groups except those marrying at very young ages. In other words, fertility behaviors at higher parities appear to be less associated with marriage than lower-parity fertility (this is also evident in the fully-interactive models). Overall, after controlling for race, educational attainment, and religious affiliation, never-married women are simulated to have only 1.25 lifetime births compared with 2.94, 2.35, and 2.23 among women who were married age 20+, 25+, and 25-34, respectively.

Conclusions

In this paper, we presented a method for translating multiple-birth hazard model coefficients into simluated fertility measures, including age-specific fertility rates, total fertility rates, and parity progression ratios. The key advantage of the method is that it permits comparisons of the timing and number of lifetime births across groups while controlling for other characteristics. For example, one could use the method to simulate total fertility rates for various racial and ethnic groups while holding constant a large set of other fixed and time-varying characteristics, such as social class origins, educational attainment, and marriage and employment history.

It is crucial to recognize that the simulated measures and standard errors generated from the life tables are only as good as the data and model on which they are based. For example, it is important to take into account how the model predictors vary across parities. This was clear in our NSFG example as the fully-interactive model fit the data significantly better than the additive and partially-interactive models. The consequence of selecting the less optimal model can be seen in comparisons of simulated fertility measures across model specifications. For example, the simulated TFR for white women is lower (1.65) when generated from the worst-fitting additive model than the better-fitting partially-interactive model (1.71), which is in turn lower than the TFR generated from the best-fitting fully-interactive model (1.76).

Another potential concern is that many data sources, including the NSFG, include only a few time-varying covariates, so important predictors of fertility like school enrollment, labor force participation, and involvement in sexual relationships outside of marriage may be missing from the hazard model. Indeed, the differences we observed between black and white women in simulated total lifetime fertility (from the multi-state life tables) might be driven by these unmeasured factors. In other words, the summary measures of interest to many researchers are only as good as the discrete time hazard models used as the inputs for the life tables. Related to this concern, the relationships estimated by the model and the simulated fertility measures may not be causal. For example, we estimated much higher lifetime fertility for ever-married women than never-married women. However, this difference may be due to the selection of women with higher fertility intentions into marriage rather than (or in addition to) the effects of marriage on fertility.

As a related point, it is also important to avoid simulating results that extend beyond the ranges observed in the data. If one attempts to simulate results outside of the observed data, the multistate life table will not be able to reproduce the true values resulting in biased estimates (Cai et al. 2010). In general, the closer substituted values are to the sample means, the more accurate the estimates.

Finally, our NSFG example involved fertility, but the method could be applied to other repeated events. For example, criminologists are often interested in modeling the occurrence and timing of arrests. One could use multiple-event hazard models and life tables to simulate estimates of the age pattern and cumulative number of lifetime arrests for selected groups while controlling for other key variables known to be associated with arrests. Such information could be helpful for discerning between groups for whom delinquency tends to be concentrated in adolescence and groups for whom criminal activity extends into adulthood. This may provide useful information for designing interventions and deterrence programs. A similar type of analysis could be conducted on marriages, with the intent of identifying how marriage patterns differ across groups (e.g., no marriage, late marriage, multiple marriages). Researchers could use the method to produce summary measures on the total number of marriages a woman might experience in her lifetime and the likelihood of transitioning from a first marriage to a subsequent. Overall, the methodology described here offers many practical advantages for those seeking to translate the repeated or multi-event hazard model estimates into interpretable summary measures rather than relying on cumbersome conditional hazards.

---

[i] Where the predicted hazard = exp(predicted log-odds)/(1-(+exp(predicted log-odds)).

[ii] We recommend using the cluster or svy options in Stata to account for dependence among observations (Cleves et al. 2008: 191-195), or analogous options in other statistic packages. Alternatively, one could estimate such models with individual random or fixed effects.
[iii] The normalized weight is the NSFG sampling weight divided by a constant such that the sum of the weights equals the sample size.

References

Allison, Paul D. 1995. *Survival Analysis Using the SAS System*: *A Practical Guide.* Cary, NC: SAS Institute.

Allison, Paul D. 1989. *Event History Analysis: Regression for Longitudinal Event Data.*; Newbury Park, CA: Sage Publications.

Brand, Jennie E. and Dwight Davis. 2011. “The Impact of College Education on Fertility: Evidence for Heterogeneous Effects.” *Demography* 48: 863-887.

Cai, Liming, Mark D. Hayward, Yasuhiko Saito, James Lubitz, Aaron Hagedorn, and Eileen Crimmins. 2010. “Estimation of multi-state life table functions and their variability from complex survey data using the SPACE Program.” *Demographic Research* 22(6): 129-158.

Carter, Marion. 2000. “Fertility of Mexican Immigrant Women in the U.S.: A Closer Look.” *Social Science Quarterly* 81(4): 1073-1086.

Centers for Disease Control and Prevention (CDC). 2009. “About the National Survey of Family Growth.” Atlanta, GA: National Survey of Family Growth. Retrieved May 2011 (http://www.cdc.gov/nchs/nsfg/about_nsfg.htm).

Cleaves, Mario, Roberto Gutierrez, William Gould, and Yulia Marchenko. 2008. *An Introduction to Survival Analysis Using Stata.* College Station, TX: Stata Press.

Goldstein, Alice, Michael White, and Sidney Goldstein. 1997. “Migration, Fertility, and State Policy in Hubei Province, China.” *Demography* 34(4): 481-491.

Guzzo, Karen Benajmin and Sarah Hayford. 2011. “Fertility Following an Unintended First Birth.” *Demography* 48: 1493-1516.

Lee, M.A. and Michael S. Rendall. 2001. “Self-employment disadvantage in the working lives of blacks and females.” *Population Research and Policy Review* 20: 291-320.

Lepkowski, J.M., W.D. Mosher, K.E., Davis, R.M. Groves, and J. Van Hoewyk. 2010. “The 2006-2010 National Survey of Family Growth: Sample design and analysis of a continuous survey.” *National Center for Health Statistics* Vital Health Statistics 2(150).

Lynch, Scott M. and J. Scott Brown. 2010. “Obtaining Multistate Life Table Distributions From Cross-Sectional Data: A Bayesian Extension of Sullivan’s Method.” *Demography* 47(4): 1053-1077.

Newell, Colin. 1990. *Methods and Models in Demography.* New York: The Guilford Press.

Poi, Brian P. 2004. “From the help desk: Some bootstrapping techniques.” *The Stata Journal* 4(3): 312-328.

Rendall, Michael S., Margaret M. Weden, Meenakshi Fernandes, and Igor Vaynman. Forthcoming. “Hispanic and black children’s paths to higher adolescent obesity prevalence.” *Pediatric Obesity.*

Schellekens, Jona. 2009. “Family Allowances and Fertility: Socioeconomic Differences.” *Demography* 46(3): 451-468.

Singer, Judith D. and John B. Willett. 2003. *Applied Longitudinal Data Analysis: Modeling Change and Event Occurrence.* New York: Oxford Press.

Teachman, Jay and Mark Hayward. 1993. "Interpreting Hazard Rate Models." *Sociological Methods and Research* 21:340-371.

White, Michael J., Salut Muhidin, Catherine Andrzejewski, Eva Tagoe, Rodney Knight, and Holly Reed. 2008. “Urbanization and Fertility: An Event-History Analysis of Coastal Ghana.” *Demography* 45(4): 803-816.

Table 1. Life table of first birth, simulated for black and white never-married, Protestant women with 12 years of education

	White Women			Black women		
	Hazard of birth	Proportion at risk of 1st birth	Unconditional Prob. of 1st birth	Hazard of birth	Proportion at risk of 1st birth	Unconditional Prob. of 1st birth
Age	q_x	l_x	d_x	q_x	l_x	d_x
14	0.026	1.000	0.026	0.053	1.000	0.053
15	0.026	0.974	0.025	0.053	0.947	0.050
16	0.026	0.949	0.024	0.053	0.897	0.047
17	0.026	0.925	0.024	0.053	0.850	0.045
18	0.026	0.901	0.023	0.053	0.805	0.043
19	0.026	0.878	0.023	0.053	0.762	0.040
20	0.042	0.855	0.036	0.084	0.722	0.061
21	0.042	0.819	0.034	0.084	0.661	0.055
22	0.042	0.785	0.033	0.084	0.606	0.051
23	0.042	0.753	0.031	0.084	0.555	0.047
24	0.042	0.721	0.030	0.084	0.509	0.043
25	0.030	0.691	0.021	0.061	0.466	0.028
26	0.030	0.670	0.020	0.061	0.438	0.027
27	0.030	0.650	0.019	0.061	0.411	0.025
28	0.030	0.631	0.019	0.061	0.386	0.024
29	0.030	0.612	0.018	0.061	0.362	0.022
30	0.028	0.594	0.017	0.058	0.340	0.020
31	0.028	0.577	0.016	0.058	0.320	0.018
32	0.028	0.561	0.016	0.058	0.302	0.017
33	0.028	0.545	0.015	0.058	0.285	0.016
34	0.028	0.530	0.015	0.058	0.268	0.015
35	0.015	0.515	0.008	0.032	0.253	0.008
36	0.015	0.507	0.008	0.032	0.245	0.008
37	0.015	0.499	0.008	0.032	0.237	0.008
38	0.015	0.491	0.008	0.032	0.229	0.007
39	0.015	0.484	0.007	0.032	0.222	0.007
40	0.003	0.476	0.001	0.005	0.215	0.001
41	0.003	0.475	0.001	0.005	0.214	0.001
42	0.003	0.474	0.001	0.005	0.213	0.001
43	0.003	0.473	0.001	0.005	0.211	0.001
44	0.003	0.471	0.001	0.005	0.210	0.001
45		0.470			0.209	

Table 2. Additive, partially-interactive, and fully-interactive models of fertility among black and white women

	Additive Model		Partially-interactive Model		Fully-interactive Model 1st parity		2nd parity		3rd+ parity	
Age 14-19 (ref.)										
Age 20-24	0.63	***	0.66	***	0.50	***	-0.10		0.44	
Age 25-29	0.32	***	0.53	***	0.15		-0.45	**	-0.21	
Age 30-34	0.07		0.53	***	0.09		-0.66	***	-0.58	
Age 35-39	-0.73	***	-0.06		-0.53	**	-1.18	***	-1.63	***
Age 40-44	-2.57	***	-1.82	***	-2.31	***	-2.83	***	-3.67	***
First Parity (ref.)										
Second Parity	0.69	***	1.51	***	---		---		---	
x 15-19	---		-0.76	***	---		---		---	
x 20-24	---		-1.00	***	---		---		---	
x 25-29	---		-1.22	***	---		---		---	
x 30-34	---		-1.14	***	---		---		---	
x 35-39	---		-1.02		---		---		---	
Third or higher parity	-0.12		0.80	*	---		---		---	
x 15-19	---		-0.32		---		---		---	
x 20-24	---		-0.96	*	---		---		---	
x 25-29	---		-1.39	***	---		---		---	
x 30-34	---		-1.86	***	---		---		---	
x 35-39	---		-2.12	**	---		---		---	
Protestant (ref.)	---				---		---		---	
No Religion	-0.03		-0.05		-0.01		-0.11		-0.09	
Catholic	-0.10	*	-0.09		-0.21	**	0.02		0.16	
Other Religion	0.01		0.03		-0.21	*	0.06		0.71	***
Black (vs. NH-white)	0.55	***	0.52	***	0.75	***	0.12		0.34	***
Years of Education	0.02	**	0.02	*	0.02		0.03		0.03	
Married	1.24	***	1.20	***	1.85	***	0.85	***	0.23	*
Intercept	-3.77	***	-3.80	***	-3.88	***	-2.16	***	-2.74	***
N	129,166		129,166		86,765		20,097		22,304	
LL	-29,117		-28,954		-15,074		-7,762		-5,792	
-2LL	58,233		57,907		30,149		15,524		11,584	

Model Comparisons	df	Chi-square	p-value	
Partially-interactive vs. Additive	10	326	0.000	***
Fully-interactive vs. Partially-interactive	10	651	0.000	***

Table 3. Partial multi-parity fertility life table simulated for Protestant white women with 12 years of education, who married at age 25. Based on fully-interactive model (Table 2)

	Parity 1			Parity 2		
	Predicted Hazard of birth	Proportion at risk of birth	Unconditional Prob. of birth	Predicted Hazard of birth	Proportion at risk of birth	Unconditional Prob. of birth
x	$q_{x,1}$	$l_{x\ 1}$	$d_{x,1}$	$q_{x,2}$	$l_{x,2}$	$d_{x,2}$
14	0.026	1.000	0.026	0.139	0.026	0.002
15	0.026	0.974	0.025	0.139	0.049	0.005
16	0.026	0.949	0.024	0.139	0.069	0.008
17	0.026	0.925	0.024	0.139	0.085	0.010
18	0.026	0.901	0.023	0.139	0.098	0.012
19	0.026	0.878	0.023	0.139	0.108	0.014
20	0.042	0.855	0.036	0.128	0.130	0.014
21	0.042	0.819	0.034	0.128	0.150	0.017
22	0.042	0.785	0.033	0.128	0.166	0.019
23	0.042	0.753	0.031	0.128	0.178	0.021
24	0.042	0.721	0.030	0.128	0.187	0.022
25	0.164	0.691	0.113	0.194	0.278	0.043
26	0.164	0.578	0.095	0.194	0.330	0.055
27	0.164	0.483	0.079	0.194	0.354	0.061
28	0.164	0.404	0.066	0.194	0.359	0.063
29	0.164	0.338	0.055	0.194	0.351	0.063
30	0.155	0.283	0.044	0.164	0.332	0.051
31	0.155	0.239	0.037	0.164	0.318	0.049
32	0.155	0.202	0.031	0.164	0.301	0.047
33	0.155	0.171	0.026	0.164	0.280	0.044
34	0.155	0.144	0.022	0.164	0.259	0.041
35	0.090	0.122	0.011	0.104	0.230	0.023
36	0.090	0.111	0.010	0.104	0.216	0.022
37	0.090	0.101	0.009	0.104	0.203	0.021
38	0.090	0.092	0.008	0.104	0.191	0.019
39	0.090	0.083	0.008	0.104	0.179	0.018
40	0.016	0.076	0.001	0.022	0.162	0.004
41	0.016	0.075	0.001	0.022	0.160	0.003
42	0.016	0.073	0.001	0.022	0.157	0.003
43	0.016	0.072	0.001	0.022	0.155	0.003
44	0.016	0.071	0.001	0.022	0.153	0.003
45		0.070			0.150	

Table 4. Simulated total fertility rate and parity progression ratios for black and white women and by marital status history, based on fully-interactive hazard model (bootstrapped SEs in parentheses)

			Marital Status History			
	Black	White	Never Married	Married age 20+	Married age 25+	Married age 25-34
<u>Simulated Total Fertility Rate</u>						
Average number of lifetime births						
1st parity	0.95	0.78	0.56	0.98	0.95	0.91
	(0.01)	(0.02)	(0.02)	(0.00)	(0.01)	(0.01)
2nd parity	0.82	0.59	0.39	0.92	0.82	0.75
	(0.02)	(0.02)	(0.02)	(0.04)	(0.04)	(0.04)
3rd parity	0.47	0.24	0.18	0.57	0.37	0.35
	(0.02)	(0.01)	(0.01)	(0.03)	(0.02)	(0.02)
4th parity	0.25	0.10	0.08	0.28	0.14	0.14
	(0.02)	(0.01)	(0.01)	(0.02)	(0.01)	(0.01)
5th or higher parity	0.21	0.05	0.05	0.18	0.08	0.07
	(0.03)	(0.01)	(0.01)	(0.02)	(0.01)	(0.01)
All parities	2.71	1.76	1.25	2.94	2.35	2.23
	(0.09)	(0.04)	(0.05)	(0.11)	(0.07)	(0.07)
<u>Simulated Parity Progression Ratios</u>						
Probability of progression from:						
No births to 1st parity	0.95	0.78	0.56	0.98	0.95	0.91
	(0.01)	(0.02)	(0.02)	(0.00)	(0.01)	(0.01)
1st to 2nd parity	0.86	0.75	0.68	0.94	0.86	0.83
	(0.02)	(0.02)	(0.04)	(0.04)	(0.04)	(0.04)
2nd to 3rd parity	0.58	0.42	0.47	0.61	0.45	0.46
	(0.02)	(0.02)	(0.02)	(0.02)	(0.02)	(0.01)
3rd to 4th parity	0.53	0.40	0.43	0.50	0.39	0.39
	(0.02)	(0.02)	(0.02)	(0.02)	(0.01)	(0.01)
4th to 5th or higher parity	0.82	0.53	0.57	0.64	0.53	0.52
	(0.07)	(0.04)	(0.05)	(0.04)	(0.03)	(0.03)

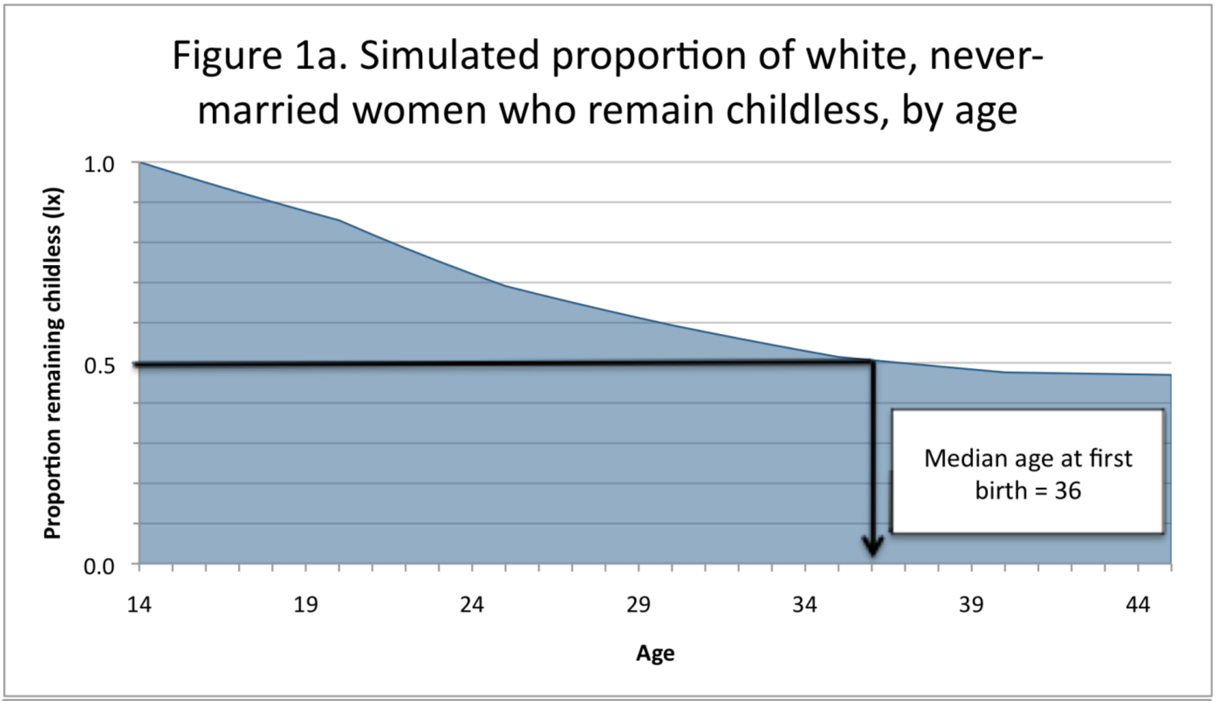
Figure 1a. Simulated proportion of white, never-married women who remain childless, by age
Proportion remaining childless (lx)
1.0
0.5
0.0
14
19
24
29
34
39
44
Age
Median age at first birth = 36

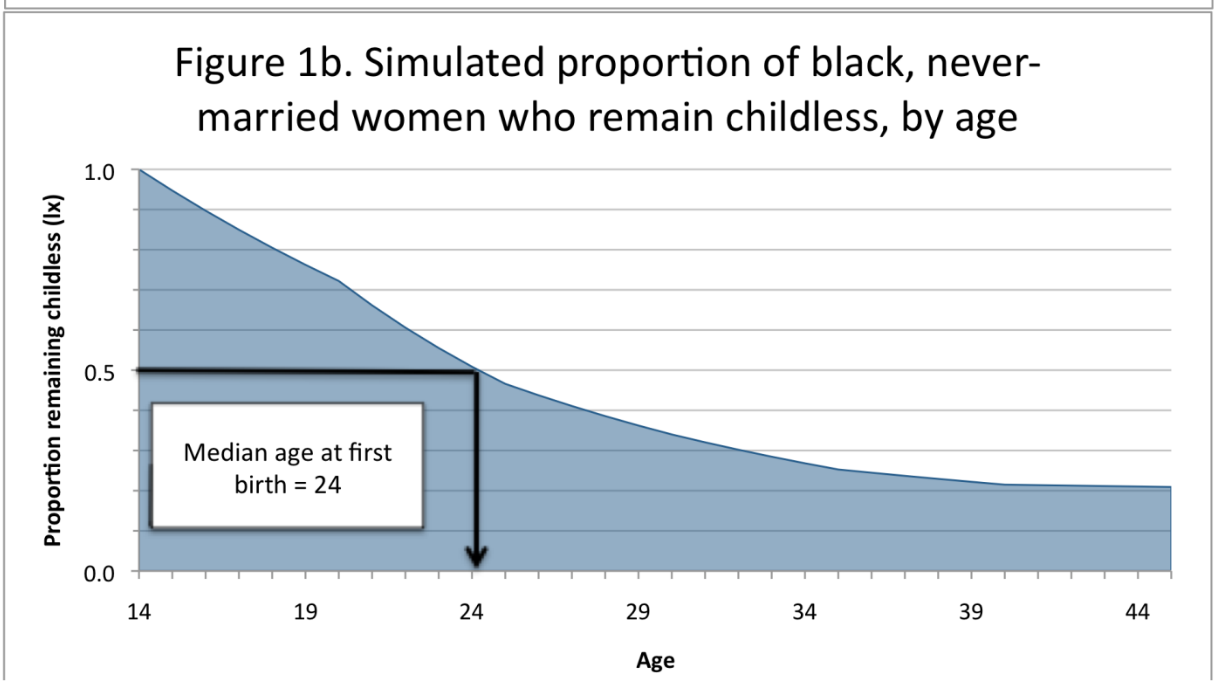
Figure 1b. Simulated proportion of black, never-married women who remain childless, by age
Proportion remaining childless (lx)
1.0
0.5
0.0
14
19
24
29
34
39
44
Age
Median age at first birth = 24

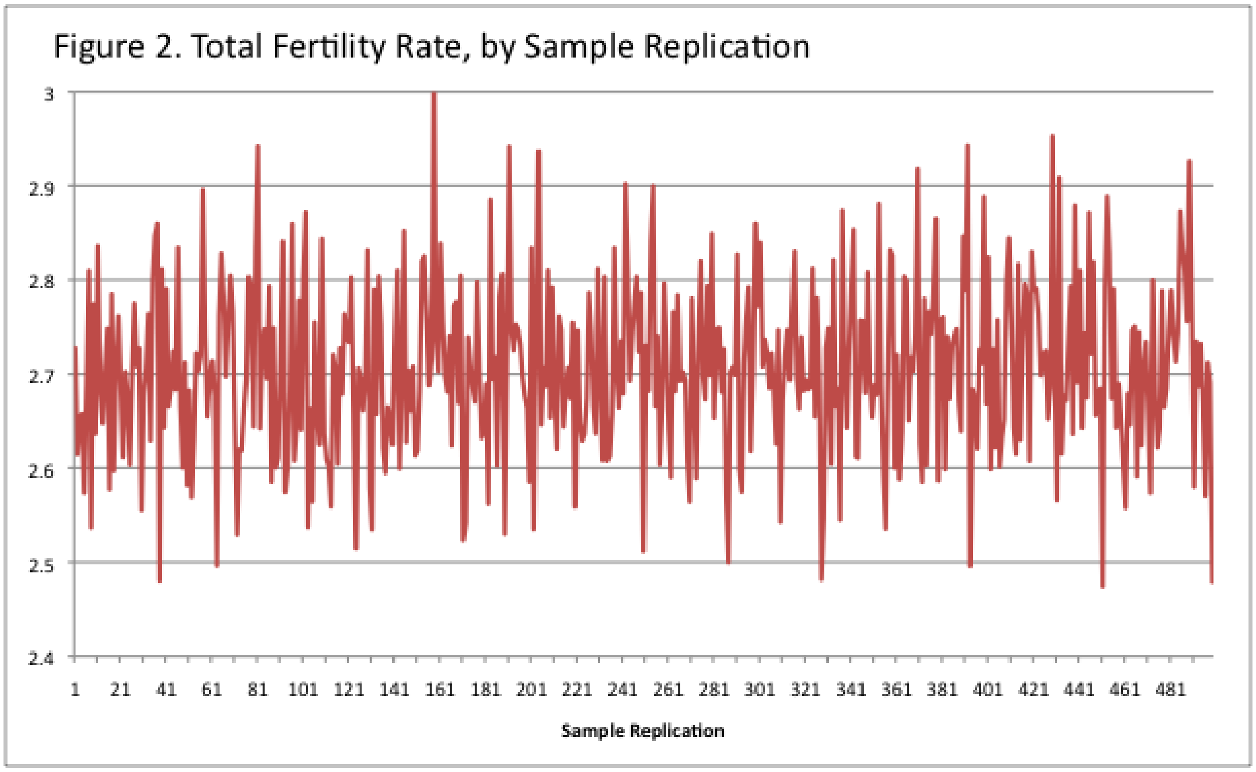
Figure 2. Total Fertility Rate, by Sample Replication
3
2.9
2.8
2.7
2.6
2.5
2.4
1 21 41 61 81 101 121 141 161 181 201 221 241 261 281 301 321 341 361 381 401 421 441 461 481
Sample Replication

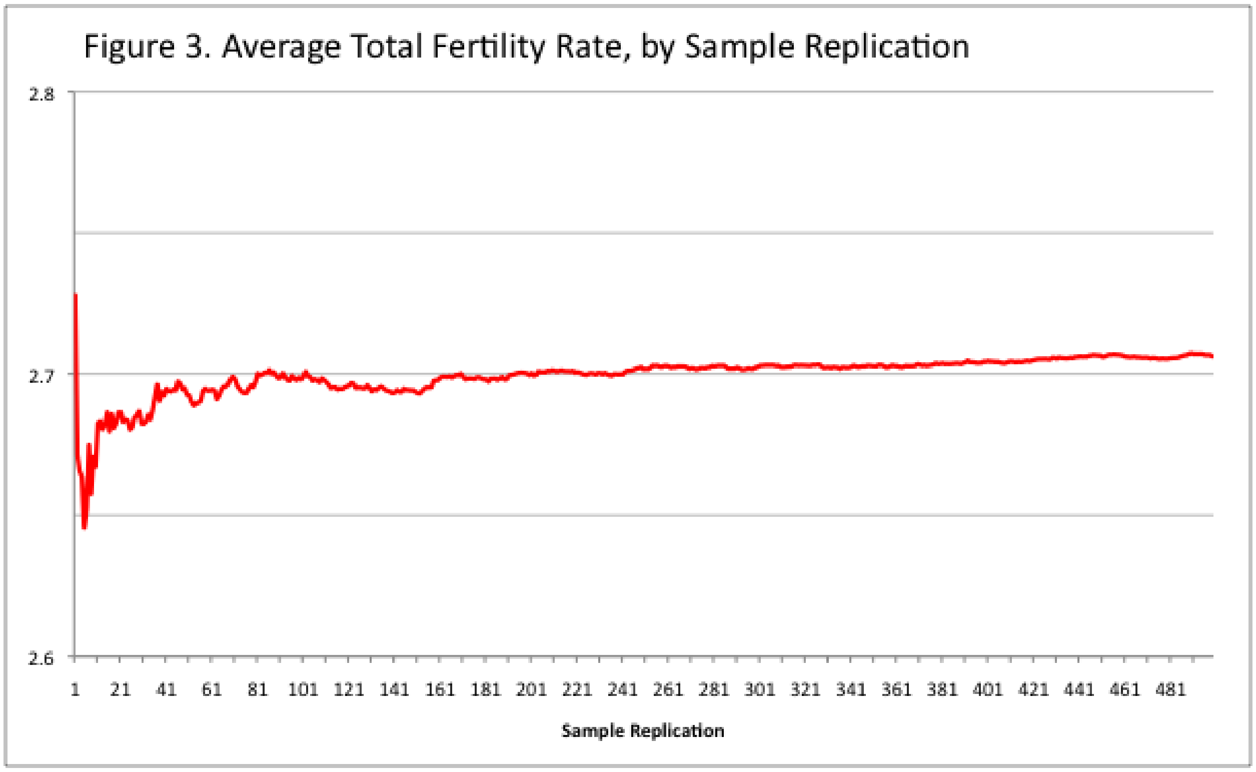
Figure 3. Average Total Fertility Rate, by Sample Replication
2.8
2.7
2.6
1 21 41 61 81 101 121 141 161 181 201 221 241 261 281 301 321 341 361 381 401 421 441 461 481
Sample Replication

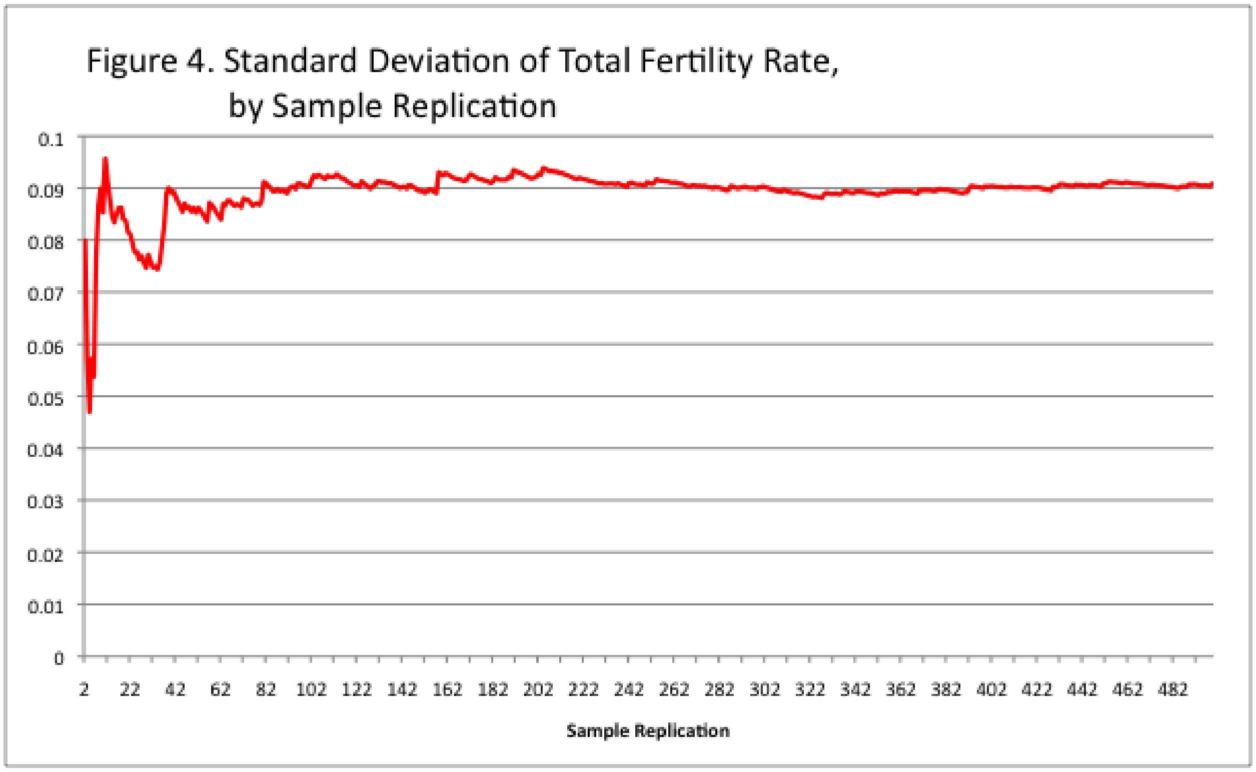
Figure 4. Standard Deviation of Total Fertility Rate,
by Sample Replication
0.1
0.09
0.08
0.07
0.06
0.05
0.04
0.03
0.02
0.01
0
2 22 42 62 82 102 122 142 162 182 202 222 242 262 282 302 322 342 362 382 402 422 442 462 482
Sample Replication

Figure 5. Simulated parity- and age-specific fertility rates, black and white women, based on fully-interactive hazard model

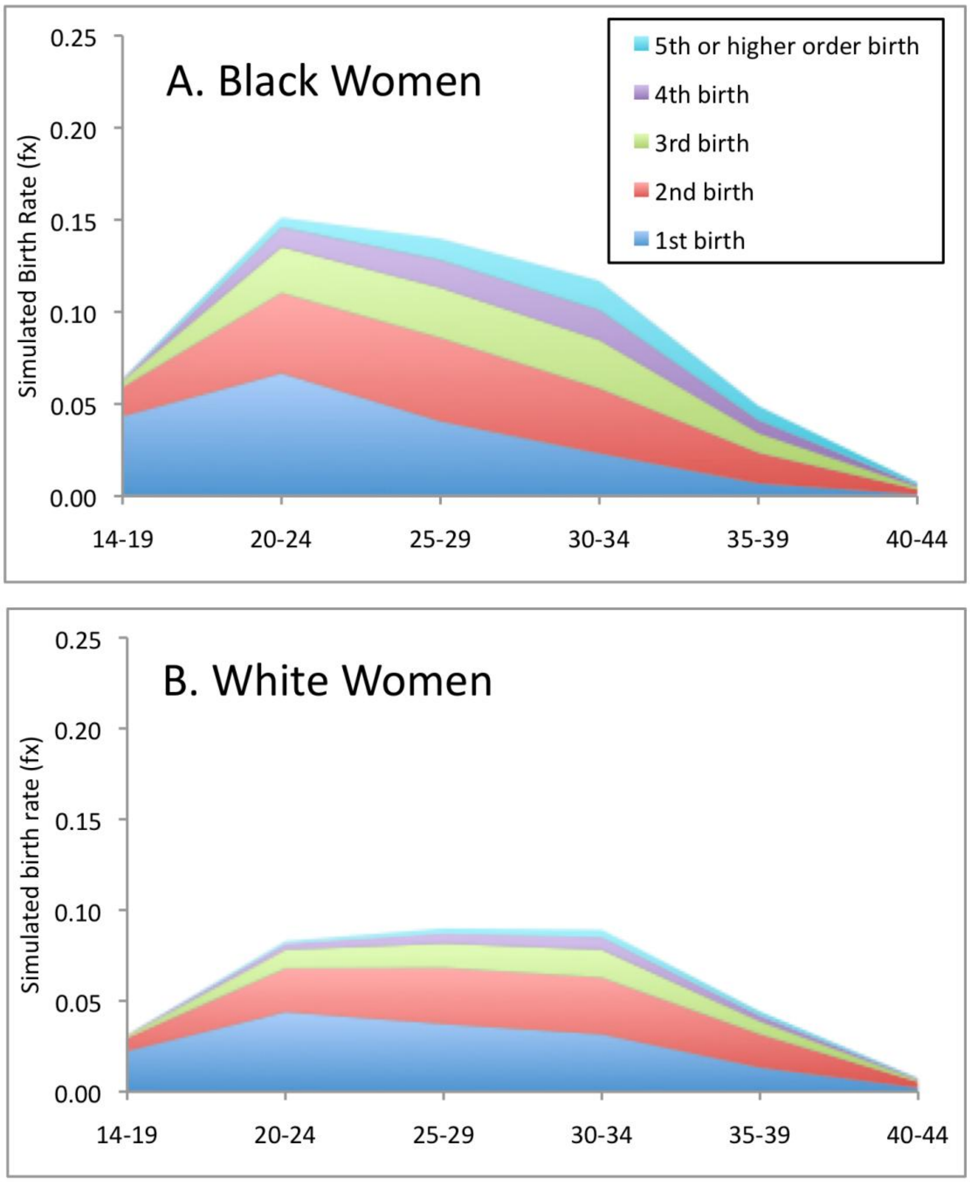

Figure 6. Simulated parity- and age-specific fertility rates, by marital status history, based on fully-interactive hazard model

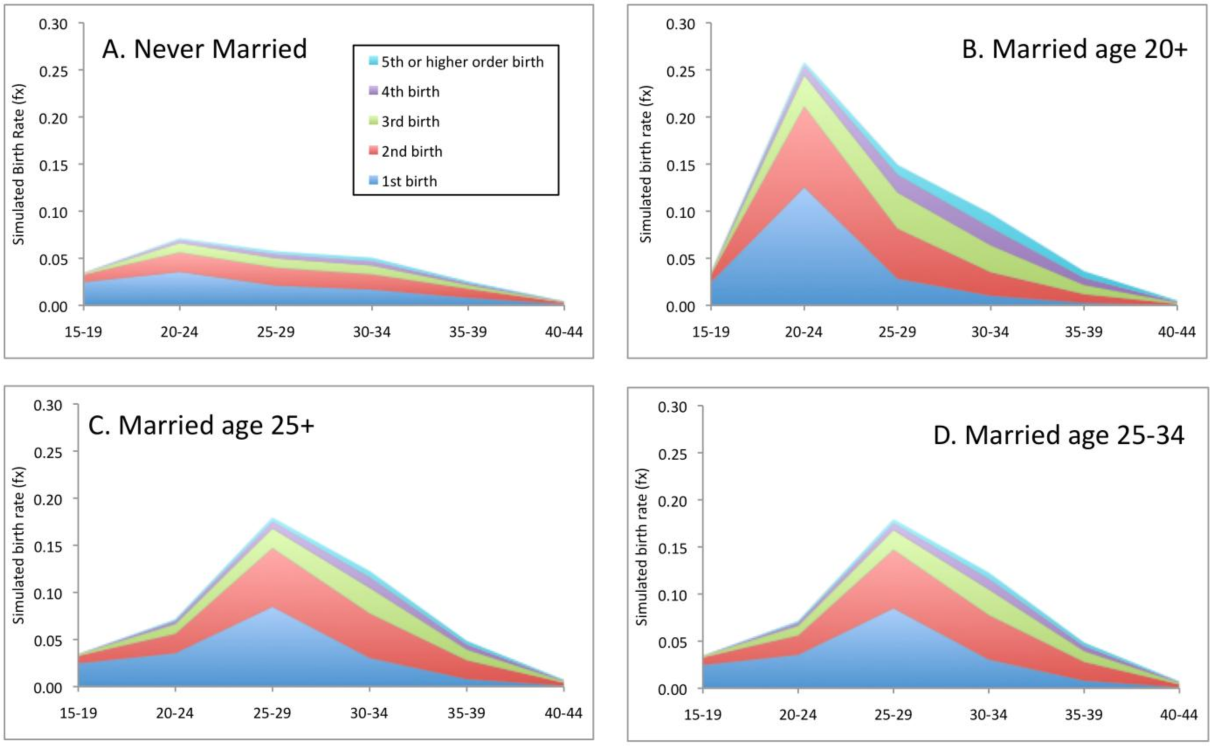

Appendix 1. Simulated parity- and age-specific fertility rates for black and white women, and by marital status history, based on fully-interactive hazard model

				Marital Status History			
		Black	White	Never Married	Married age 20+	Married age 25+	Married age 25-34
1st birth	14-19	0.043	0.022	0.025	0.025	0.025	0.025
	20-24	0.067	0.044	0.036	0.125	0.036	0.036
	25-29	0.041	0.037	0.021	0.028	0.085	0.085
	30-34	0.023	0.032	0.017	0.010	0.030	0.030
	35-39	0.007	0.013	0.008	0.003	0.008	0.002
	40-45	0.001	0.002	0.001	0.000	0.001	0.000
2nd birth	14-19	0.016	0.007	0.008	0.008	0.008	0.008
	20-24	0.044	0.024	0.021	0.086	0.021	0.021
	25-29	0.045	0.031	0.019	0.053	0.062	0.062
	30-34	0.035	0.031	0.016	0.025	0.047	0.047
	35-39	0.017	0.018	0.010	0.009	0.020	0.009
	40-45	0.003	0.003	0.002	0.001	0.003	0.002
3rd birth	14-19	0.003	0.001	0.001	0.001	0.001	0.001
	20-24	0.025	0.010	0.010	0.033	0.010	0.010
	25-29	0.027	0.013	0.010	0.038	0.021	0.021
	30-34	0.026	0.015	0.009	0.028	0.027	0.027
	35-39	0.011	0.007	0.004	0.010	0.011	0.009
	40-45	0.001	0.001	0.001	0.001	0.002	0.001
4th birth	14-19	0.001	0.000	0.000	0.000	0.000	0.000
	20-24	0.011	0.003	0.004	0.010	0.004	0.004
	25-29	0.015	0.005	0.005	0.019	0.007	0.007
	30-34	0.016	0.007	0.005	0.019	0.011	0.011
	35-39	0.007	0.003	0.002	0.007	0.006	0.004
	40-45	0.001	0.000	0.000	0.001	0.001	0.001
5th or higher order birth	14-19	0.000	0.000	0.000	0.000	0.000	0.000
	20-24	0.005	0.001	0.001	0.003	0.001	0.001
	25-29	0.012	0.003	0.003	0.010	0.004	0.004
	30-34	0.016	0.004	0.003	0.015	0.006	0.006
	35-39	0.008	0.002	0.002	0.007	0.003	0.003
	40-45	0.001	0.000	0.000	0.001	0.000	0.000
All Parities	14-19	0.063	0.031	0.035	0.035	0.035	0.035
	20-24	0.151	0.083	0.071	0.258	0.071	0.071
	25-29	0.140	0.090	0.057	0.149	0.179	0.179
	30-34	0.117	0.089	0.051	0.097	0.123	0.123
	35-39	0.049	0.044	0.025	0.036	0.048	0.027
	40-45	0.007	0.007	0.004	0.005	0.007	0.004