The Intrinsic Total Fertility Rate: A New Approach to the Measurement of Fertility[1]

Peter McDonald and Rebecca Kippen
Australian Demographic and Social Research Institute
The Australian National University

Demographers measure and explain fertility at both the individual and the aggregate levels. Individual level analysis makes use of standard statistical techniques with some measure of fertility as the dependent variable. In individual analysis, the main 'demographic' variables related to fertility, age and parity, will either be a defining component of the dependent variable or will be employed as control variables in the analysis. Given the nature of available data, individual level analysis of fertility is usually cross-sectional, although when multiple cross-sections or life history data are available, comparisons will be made about the variation in observed associations across time. Again because of the nature of the available data, it is rare indeed that full fertility histories of birth cohorts are examined in individual level analysis. In aggregate level analysis of fertility, the available data are usually restricted to a very small number of characteristics such as age, place of residence and parity. The simple Period Total Fertility Rate (PTFR) for a geographic area remains the most commonly available aggregate measure. Often it is the only available aggregate measure of the fertility trend.

The main purpose of both individual level analysis and aggregate level analysis is to explain behavior; one at the individual level, the other at the aggregate level. It is not unusual for associations found at the individual level to be projected onto the aggregate level but, in the context of below replacement fertility, this type of projection has become risky indeed. At the country level, associations between labor force participation and fertility or between education and fertility tend now to be positive whereas individual level associations remain negative in most contexts. The conventional explanation of this paradox (an ecological fallacy) is that institutional factors (society-wide), generally absent from individual level analysis, play a significant role in explaining fertility when it is below replacement. Demography has always had a severe aggregation problem; the problem is even more severe, it seems, when fertility is below replacement.

In the work that has led to this paper, our aim is to try to determine better ways to project a country's fertility in the short term, defined as the next ten years or so. In general, statistical agencies still use very crude methods to do this, for example, extrapolation of trends in the TFR. Stochastic extrapolation has been propounded as a major advance, but, effectively, the approach is the same. Recently also, we have seen the emergence of a plethora of measures termed 'tempo-adjusted' Total Fertility Rates (Bongaarts and Feeney 1998; Kohler and Ortega 2002; Zeng and Land 2002; Sobotka, Lutz and Philipov 2005; Rodríguez 2006). The main purpose of these measures is to make the point that, when age at first birth is rising, the current PTFR may understate the 'underlying' or future PTFR. However, none of the proponents of adjusted measures

[1] The authors wish to acknowledge the contribution of the Australian Bureau of Statistics in making detailed census tables available to us. We would also like to acknowledge helpful comments provide by Tomas Frejka and Heather Booth.

is very explicit about precisely what is being measured through the adjustment. Is it a predicted future PTFR and, if so, to what future year will it apply? Beyond providing a sneaking sense that, indeed, the current PTFR may be an underestimate of the future PTFR, these measures have not proven to be useful in projection of fertility. This is essentially because these measures are contingent upon an assumption that observed change is due only to delay being exercised by successive cohorts; that is, that cohort quantum is invariant. These measures have been criticized because they make no allowance for the highly likely eventuality that completed cohort fertility (CCF) is a function of the timing of births. Schoen (2004) has provided a good review of critiques of tempo-adjusted measures and concludes that, despite criticism, these measures carry a weight of support that they do not deserve. He goes on to define a new measure, the Average Cohort Fertility rate (ACF) that decomposes an observed period fertility rate into tempo and quantum components. For projection purposes, however, this measure has the fatal disadvantage that we need to know the future in order to explain the present. Like van Imhoff (2001) and Ryder (1980), Schoen concludes that cohort behavior cannot be accurately measured until that behavior has been completed.

This last conclusion has not stopped the analysis of fertility trends by birth cohort (Frejka and Calot 2001a, 2001b; Frejka and Sardon 2004). These studies focus upon very detailed analyses of trends in cumulated cohort fertility, including incomplete cumulated cohort fertility. Because the measures used are cumulated measures, these analyses abstract from period effects effectively prioritizing the cohort in behavioral explanation. In this respect, they are similar to the tempo-adjusted measures that also prioritize the cohort over the cross-section. Cumulative cohort analysis has given rise to the very important concept of recuperation, the extent to which a cohort that has experienced delay of the first birth is able to recuperate its quantum of fertility through higher rates of childbearing at older ages (Lesthaeghe and Willems 1999; Lesthaeghe and Moors 2000; Lesthaeghe 2001). Using this form of analysis, with each year of new data, the extent of cohort recuperation can be assessed. The aim is to make assessments or inferences about likely future CCF taken as a summary measure of a society's behavior.

The primacy of cohort or period measurement

We accept that the tempo-adjusted measures provide a cogent argument that future fertility may well be somewhat higher than current fertility. We also accept that the cumulated cohort measures indicate the variation across societies in the extent to which CCF will be influenced by timing effects. However, we suggest that neither approach is very useful in projecting births in the short term, our central aim. Individual level analysis provides very useful indications of associations between individual characteristics and fertility that might be used in projections but there is a great deal of evidence that these associations shift over time in ways that are difficult to predict. Furthermore, it is apparent that institutional or society-wide factors also affect fertility substantially in low fertility settings.

Each birth has its own explanation and decision-making framework. While members of the same birth cohort, to some degree, may have similar ideals about the number of children that they would like to have, following others as reviewed by Ni Bhrolchain (1992), our view is that completed cohort fertility (CCF) is primarily an outcome of many cross-sectionally determined decisions (or non-decisions) rather than being determined primarily by the cohort's ideals. Trends in CCF are smoother than trends in the Total Fertility Rate only because cross-sectional fluctuations average out across many years of experience. Nevertheless, there is no doubt that the cumulative cohort approach to measurement is very revealing of trends and underlying behavior.

Thus, we take the intermediate position that, in a situation of below replacement fertility, both cohort and period effects are potentially important. We assert that, when fertility is below replacement, the essential determinants of both cohort and period fertility rates are the probability distribution of first births by age and the probability distribution of second births by age and duration since the first birth. In the case study that we describe below, we observe that probability distributions for higher order births (3+) tend to be relatively invariant. Furthermore, in very low fertility societies, probabilities of higher order births are so low that some error in their estimation will have little effect on aggregate fertility.

The timing of the first birth has a cohort influence that is related not so much to the particular circumstances of the cohort members but to the cohort's prior achievement and intentions. If a relatively high proportion of the cohort have not yet had a first birth, we can be certain that most will do so in the future because data on intentions tend to show that, in most social settings, most women want to have a first birth. This is consistent with the logic of tempo-adjusted methods. However, behavioural change in respect of the first birth is likely to be cross-sectional. If there is a continuation of a trend to later childbearing, it is likely to be spread across cohorts. If there should be a reversal of the trend (through changes in policy or fashion), again it is likely to affect all cohorts simultaneously, but, because of cohort achievement, may have a larger impact for more recent birth cohorts. This is because those who have already had a first birth cannot reverse the experience and have their first birth at a younger age.

Below in our case study, we show that the probability of a second birth is strongly influenced by the time elapsed since the first birth. Age has an additional but smaller effect. On this basis, we would argue that the occurrence of the second birth also has a strong cohort influence: a given probability of duration following the first birth. Again, the main driving forces are achievement and intention of the cohort rather than the circumstances of the cohort. Once more, however, if there should be a change of behavior, a shift to a shorter or longer interval between the first and the second birth, this will have a cross-sectional effect. A classic example was the introduction of the 'speed bonus' in Sweden where women of different birth cohorts reacted simultaneously to the policy initiative. Again, however, because of achievement, the capacity to react is higher for later birth cohorts because a higher percentage of women in the older cohorts have already had their second birth.

To summarise, where behavior is not changing, there will be cohort effects deriving from the momentum effects of past changes in behavior (eg. the making up of delayed first births). On the other hand, changes in behavior tend to be cross-sectional. However, because of achievement, changes in behavior are likely to have differential effects by birth cohort with the effects being greater for younger birth cohorts. This logic leads us to the following framework for fertility decision-making.

A framework for fertility decision-making

A birth can be considered to be the outcome of three, broadly-defined factors related to the mother: achievement, intention and circumstances.

Achievement encapsulates the number of children that the woman already has, her current age and the time elapsed since her last birth.

Intention means the intentions that the woman has about the number and timing of her births. Intentions are determined by values that are continually modified by achievement and by changes in circumstances. In relation to achievement, as a woman ages and/or the interval since her last birth increases, her intentions are likely to be modified downwards.

Circumstances refers to everything else that influences the decision to have a birth. This includes the woman's characteristics such as education, work experience, religion, cultural origins, place of residence and relationship status and quality. It also includes economic circumstances such as actual or potential wage rate, income and wealth. It includes her psychological state and her capacity to avoid unintended pregnancy. It may include the state of the economy and other macro-influences. It includes the nature of the institutional setting in which she lives, an important component of which is the supportive or conflictual nature of social institutions and policies. Importantly, her circumstances include not only whether she is partnered, but also the achievement, intention and circumstances of her partner.

Measurement issues

At the aggregate level, only achievement is readily observable, and even then, full observation is highly data intensive. Achievement is incremental. It changes uni-directionally only through the elapse of time or through an additional baby. On the other hand, intentions and most circumstances are volatile. Achievement as we define it above is consistent with Ni Bhrolchain's preferred approach to the study of period fertility:

> period measures should be standardized for parity so as to take account of fertility in previous years. Moreover, the same arguments leading to a parity control also imply the need to control for duration of interval since previous birth. In a comprehensive approach, a control for age may also be necessary (Ni Bhrolchain 1992: 613).

Several authors have examined fertility rates by age and parity (for example, Kohler and Ortega 2002, Kippen 2003). Ni Bhrolchain (1987) has examined period fertility in China standardising for parity and time since the last birth. Rallu and Toulemon (1994) provide one of very few studies that simultaneously standardizes period fertility for age, parity and interval since the last birth. Their measure is the average number of births per woman a cohort would have over their reproductive lives if they experienced the age-parity-duration-specific fertility of a particular year. This is the approach we follow in this paper using data for Australia from 1981 onwards.

Our strategy is:

1. to examine the effects of unchanged behavior: how many first births are yet to occur?; how many first and higher order births will not occur?

2. through close observation, to examine the data for changes of behavior.

In implementing this strategy, we focus upon two probability distributions: the probability of first birth by age and the probability of second birth by age and interval since first birth.

Our approach has been to seek measures that provide early indication of behavioral change at the aggregate level. We argue that a measure proposed by Rallu and Toulemon (1994) and its components enable us to gain insights into continuing behavior and changing behavior. We have named this measure the 'Intrinsic Total Fertility Rate'.

The intrinsic growth rate of a population is the growth rate resulting if age-specific fertility and mortality is constant over a long period of time. In like manner, the Intrinsic Total Fertility Rate is the Total Fertility Rate a population would experience if its age-parity-duration-specific fertility were constant over a long period of time.

Sources of data

Studies have not standardized simultaneously for all three aspects of achievement because the data have been unavailable or the effort involved in obtaining the data has been prohibitive. Because it is important for our purposes to work with data by single years of age, single parity and single year of duration since the previous birth, conventional birth histories available from sample surveys are not large enough to provide reliable data. While parity data are collected at some Australian censuses, the interval since the most recent birth is not collected. Furthermore, this information is not collected in complete form in the Australian vital statistics collections (Corr and Kippen 2006).

Data for this research are derived from complete counts from the 1981, 1986, 1991, 1996 and 2001 Australian Censuses of Population and Housing. For each census, Australian-resident women in each household are matched to their children living in the same household at the time of the census. This is done using the 'relationship in the household' census variable, which describes the relationship of each person in the household to the household reference person (generally Person 1 or Person 2 on the census household form).

Available characteristics from these data are age of woman, number of own children in the household, and age of each of her children resident in the household at the time of each census. These data are used to construct two sets of tables for each census. The first set consists of counts of women by age, by age of oldest child, second child, third child, and so on. The second set is counts of women by age with *j* or more own children, classified by age of child *j* by age of child *j*–1.

These data are used to construct distributions of women by single year of age by number of own children in the household for each year 1980–2001. For example, the distribution for 2001 is calculated by tabulating women by number of own children from the 2001 census count. The distribution for 2000 is calculated by taking the 2001 distribution, subtracting one year from the age of each woman, and adjusting down the number of her own children if she had a child aged zero years at the 2001 census. The distribution for 1999 is calculated by ‘younging’ women in the 2000 distribution, and subtracting one from the number of own children for every child aged one year at the 2001 census, and so on back to 1996. The same process is then repeated for the other censuses.

This ‘reverse survival’ results in four sets of annual distributions: 1981–86, 1986–91, 1991–96 and 1996–2001. These are used to calculate, for each year 1981–2000, the probability that a woman aged x in year y with j own children will progress to $j+1$ children in ageing to $x+1$ years in year $y+1$.

The number of own children in the household is not equal to parity, since children may have died or left the maternal home. The transition probabilities are adjusted so that they refer to parity rather than to own children. This is done by comparing the own-children transition probabilities for each year 1991–2000 to the parity transition probabilities for 1991–2000 calculated previously by Kippen (2003, 2004). We find that the ratio of transition probabilities for each age of woman and own children/parity is relatively constant across the ten years. We therefore use ratios averaged across the decade to inflate the own-children transition probabilities to parity transition probabilities for the transition to first, second, third, and fourth birth.

These parity transition probabilities—the probability that a woman aged x in year y with parity j shifts to parity $j+1$ in ageing to $x+1$ years in $y+1$—are used to calculate parity distributions for each age of woman 15–49 and each year 1981–2001, beginning with the parity distribution calculated from the 1981 census (which asked ‘For each woman, how many babies has she ever had? Do not include still-births’). The calculated parity distributions for 1986 and 1996 are compared to those derived from the 1986 and 1996 censuses (which also asked women their parity) as a check for accuracy. They are very close. The implied annual age-specific fertility rates are also compared to those derived from birth registers. Again they are very close.

The annual parity distributions and transition probabilities are used to calculate annual births by birth order and age of mother assuming cohorts of 100,000 women at each age, 1981–2000. Births of order 2, 3, 4 for each year are then classified by duration since previous birth. This is done using the tables from each census of counts of women by age with j or more own children, classified by age of child j by age of child j–1. For example, there are 3,658 second births to women in ageing from 30 years in 1982 to 31 years in 1983. The distribution of these second births by interval from first birth can be calculated by considering women aged 34 years in the 1986 census with two or more children, second child aged 3 years, by age of first child. The age interval between the second and first child is the relevant birth interval.[2]

Synthetic transition probabilities are then calculated:

---

2 We find that there is no need to adjust birth intervals to take account of children missing from the household. A comparison of interval distributions for all women, and women for whom number of children is equal to parity, for the 1981, 1986 and 1996 censuses shows that they are virtually identical. This suggests that if mothers have children missing from their household, it is likely that all their children are missing, or that the oldest or youngest are missing.

$$q_{j,x,y,i} = \frac{P_{j,y-i,x-i;j+1,x,y}}{P_{j,x-i,y-i} - \sum_{n=1}^{i} P_{j,x-i,y-i;j+1,x-n,y-n}} \quad ,$$

$1 \le j \le 3,$

$15 \le x \le 48,$

$1981 \le y \le 2000,$

$0 \le i \le 14,$

where $q_{j,x,y,i}$ is the probability that a woman aged *x* in year *y* of parity *j* with her *j*th birth *i* years ago progresses to birth *j*+1 in ageing from *x* years in year *y* to *x*+1 years in year *y*+1,
$P_{j,x,y}$ is the number of women who had a birth of order *j* between age *x* and *x*+1 and year *y* and *y*+1, and
$P_{j,x-i,y-i;j+1,x,y}$ is the number of women who had a birth of order *j* between age *x–i* and *x–i*+1 and year *y–i* and *y–i*+1 and a birth of order *j*+1 between age *x* and *x*+1 and year *y* and *y*+1.

The data collection provides us with full information on achievement: the distribution of the female population of Australia by age, parity and interval since last birth and annual births with the same description. It then becomes a straightforward exercise to obtain birth probabilities simultaneously by single years of age, single parities and single years of duration since the last birth. The population data summarize achievement up to the year to which they apply. They are an indication of the 'momentum' inherent in the population for future births if behavior remains unchanged (see next section).

Probabilities for the most recent year indicate the combined influence of intentions and circumstances given achievement. We are able to examine whether women in a given achievement cell make decisions in the same way as women in the same achievement cell in the year before. If they do, this is an indication of stability and predictability. If they don't, then we are able to examine very precisely what changes have occurred. In particular, we are able to examine if the changes are systematic across cells of the distribution and therefore indicative of behavioural change. If so, these changes may also be able to be projected into the short term.

The nature of the measurement problem

The incidence of any phenomenon relating to a population is the result of a set of behavioral probabilities being applied to a given population composition. The outcome is affected by both the probabilities and the composition. When the PTFR is adjusted to take account of delay in the timing of births, the implied argument is that there is an inherent feature of the population composition that will in future lead to higher total fertility rates. This can be explained more clearly using an example. Figure 1 provides an example where the probability distribution of having a first birth shifts to older ages, a behavioral change. The example is Australia between 1990/91 and 2000/01 (1 July to

30 June time periods). Figure 2 shows the proportion of women at each age that were childless in both years. It indicates that, during the decade, the population 'at risk' of a first birth also shifted.

Now, suppose that the 2000/01 age-specific probabilities were to remain constant after 2000 for a long period. Although no behavioral change is assumed, in subsequent years, the first-birth rate (the first birth component of PTFR) would increase because constant probabilities of first birth would be applied to a higher proportion of women with zero parity at each age. The change in population composition is indicated by the dotted line in Figure 2 that shows the ultimate distribution in the proportion of women of zero parity at each age if the 2000/01 first birth probabilities were to remain constant over a period of approximately 35 years. Under this scenario, the first birth rate (the contribution of first births to the PTFR) would rise from 0.75 births per woman in 2000 to an eventual stable level of 0.81.

Figure 1. Probability of a first birth (given zero parity) by age, Australia, 1990/91 and 2000/01

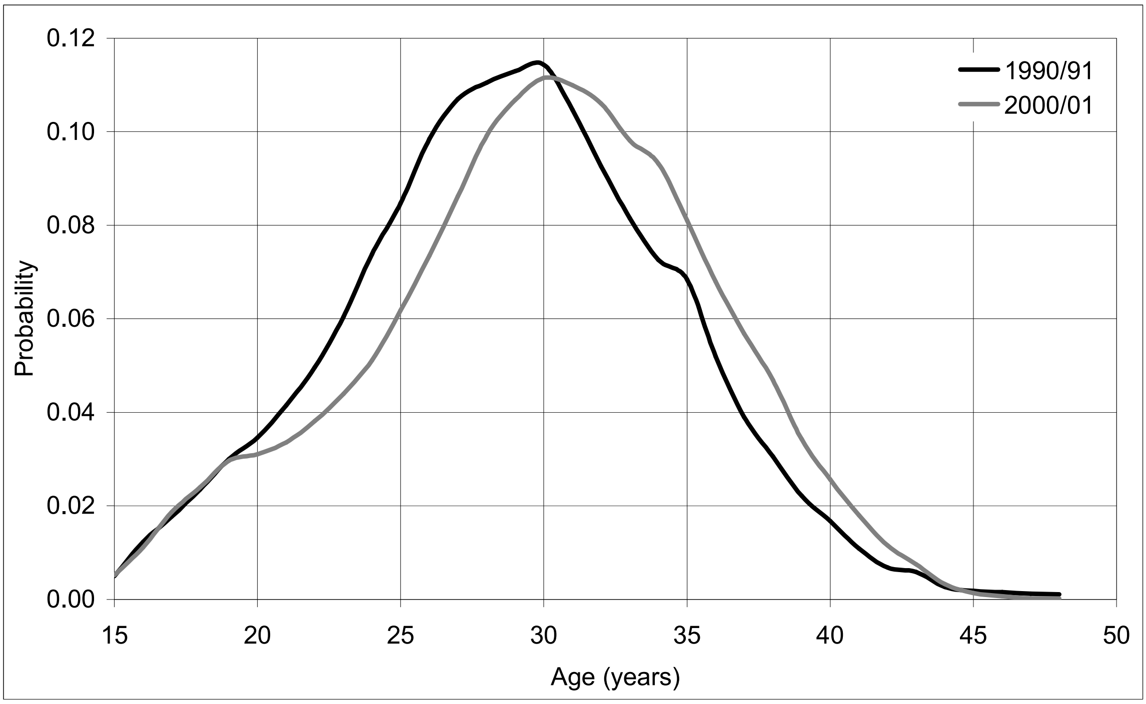

Figure 2. Proportion of women with zero parity by age, Australia, 1990 and 2000

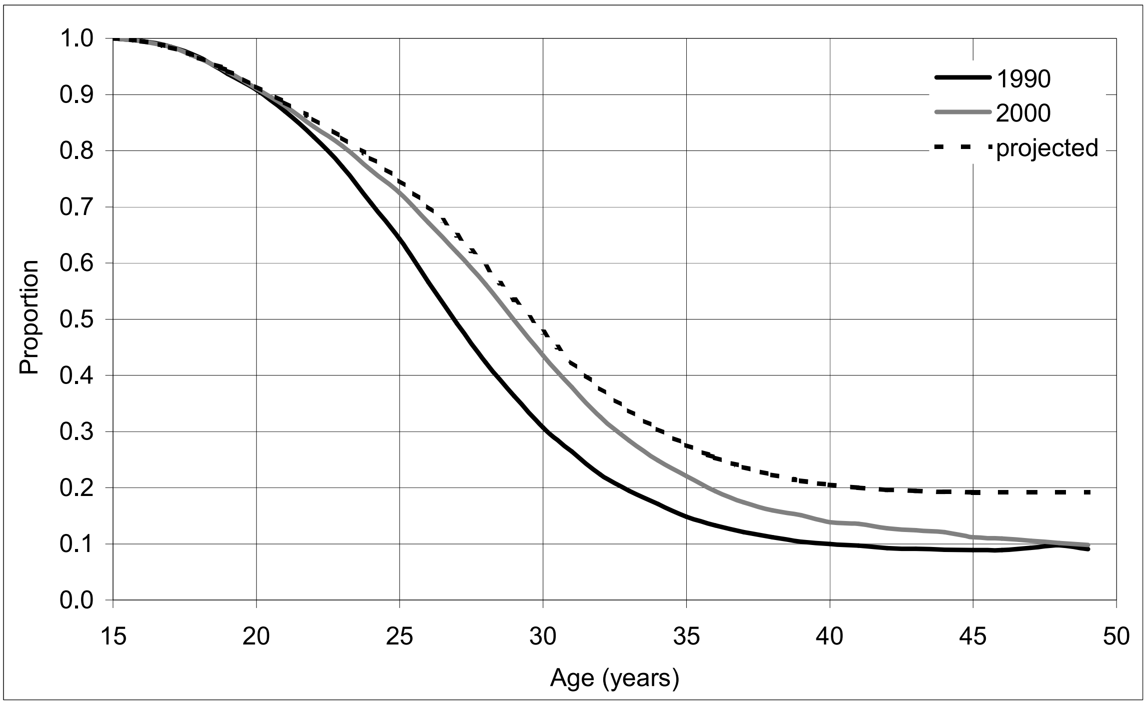

In the example above, the compositional factor that we have taken into account is the change in the proportion of women at each age with parity zero. For higher order births, not only will parity be an important dimension of composition but also the interval since the previous birth. Figure 3, using Australian data for the year 2000/01, shows the probabilities of having a second birth by interval since the first birth for three different ages of women, ages 25, 30 and 35. The figure shows the expected result that the interval since the first birth has a very large impact on the probability of having a second birth. Thus if there were to be a shift in the composition of the population that moved the population more towards the peak of this distribution (first birth two to three years ago), we would expect the second birth rate to rise. Simultaneously, the figure shows that the age of the woman also makes a small difference to her likelihood of having a second birth for a given interval since first birth. Specifically, the second birth is more likely to occur if she is aged 30 than if she is aged 25 or 35. Thus, a shift in the age composition of women of parity one would also have an impact on future second births.

Figure 3. Probability of a second birth (given parity one) by years since first birth, ages 25, 30 and 35, Australia, 2000/01

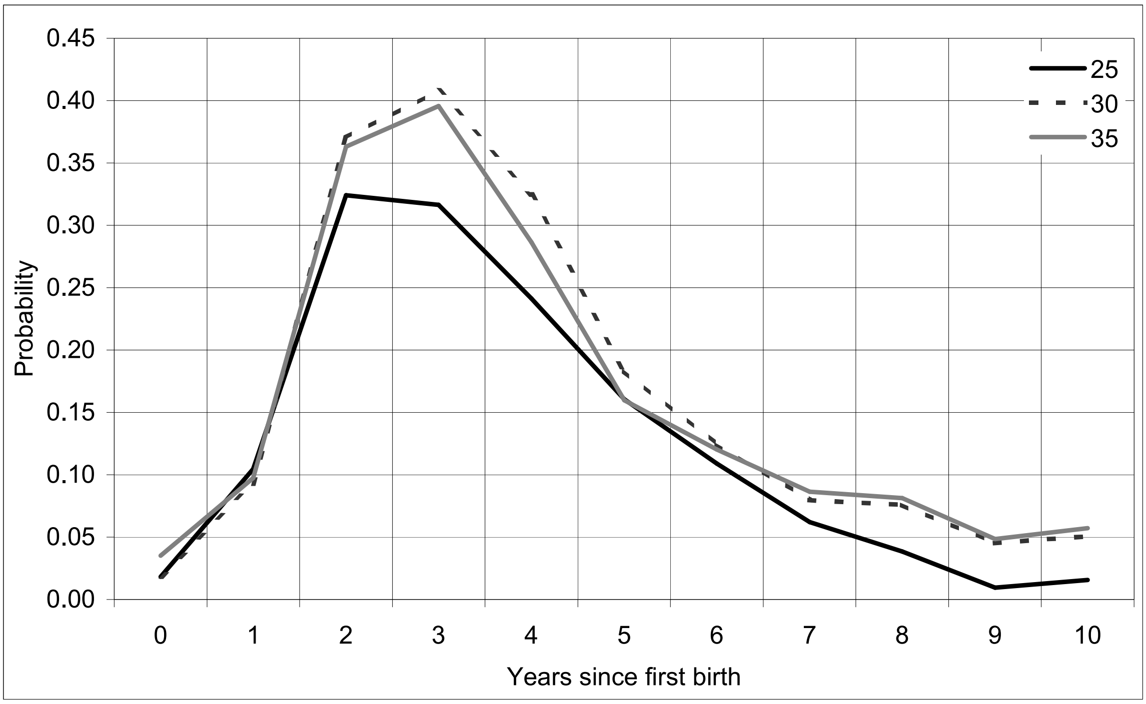

Figures 4 and 5 show how the composition by age and interval since last birth changed for women of parity one in Australia between 1990 and 2000. By 2000, by age, the peak of the distribution had shifted to the ages where the probabilities of having a second birth are highest (around age 30). Also, at ages 30–34, the proportions that had had their first birth either two or three years earlier (the intervals with the highest probabilities of progression to the second birth) were higher in 2000 than in 1990. Now suppose that the probabilities of second birth by age and interval since first birth in 2000/01 were to remain constant after 2000 for a long period. The second-birth rate (the second birth contribution to PTFR) would rise from 0.60 births per woman to 0.66 because the compositional shift was favorable to second births occurring.

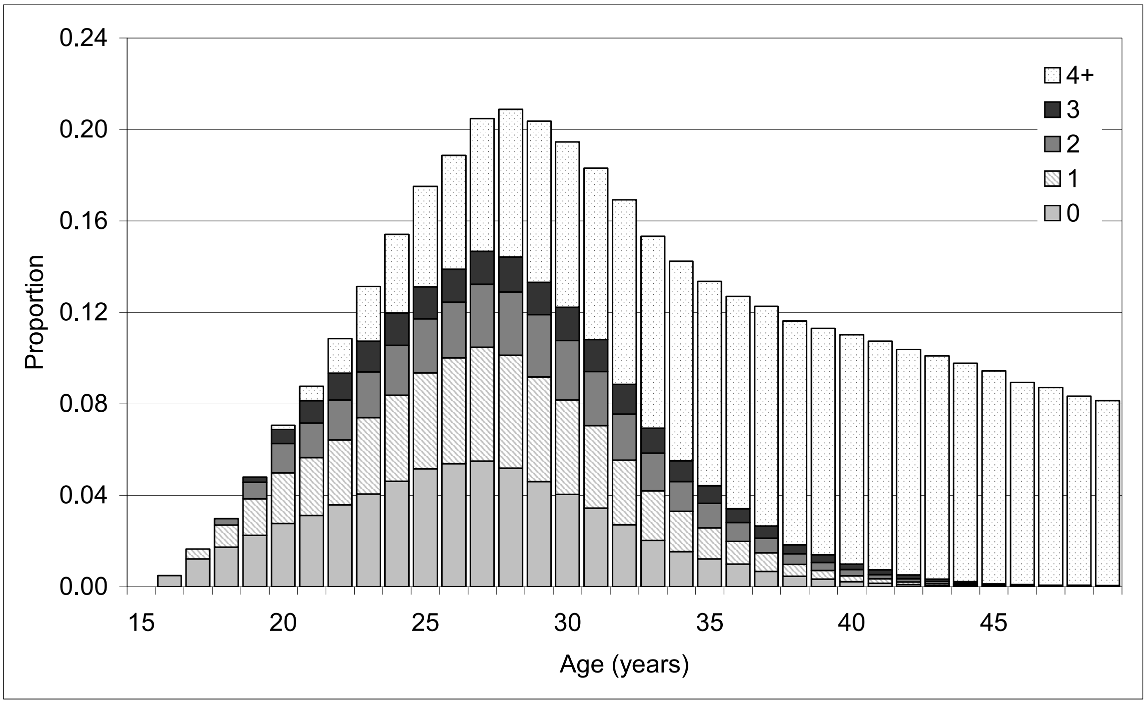

Figure 4. Proportion of women with parity one by age and years since first birth, Australia, 1990

Figure 5. Proportion of women with parity one by age and years since first birth, Australia, 2000

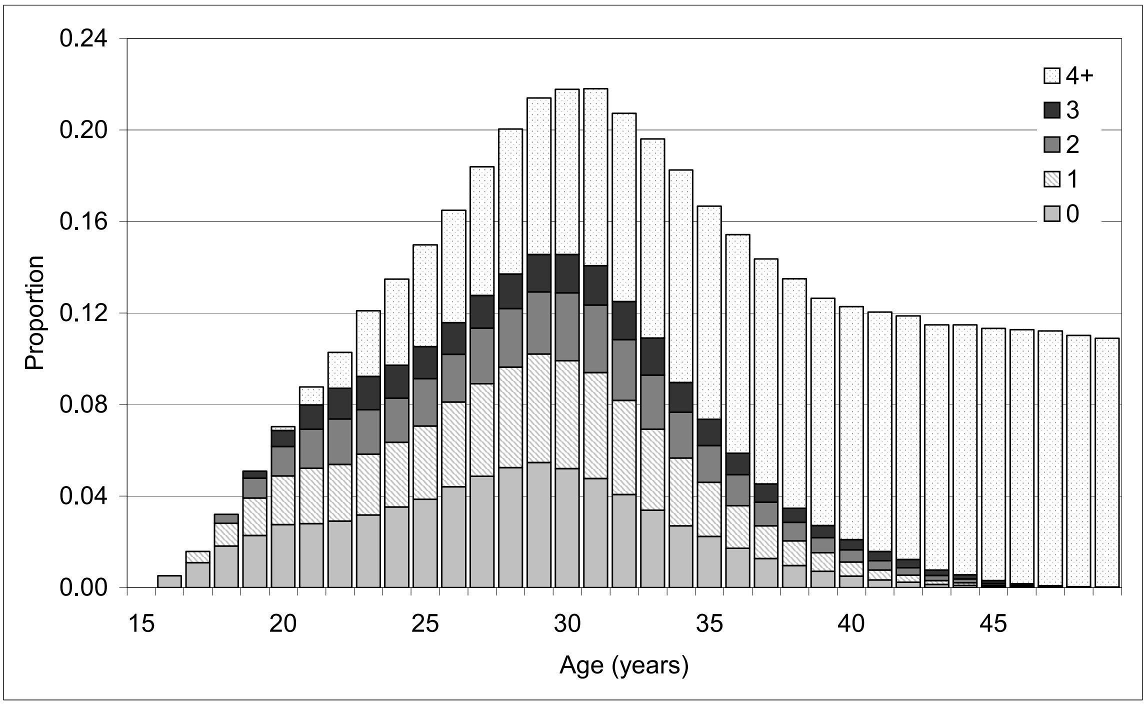

These examples demonstrate that, even when probabilities of birth by age, parity and interval since the previous birth are held constant, fertility rates may change in the future because of changes in the composition of the population according to these characteristics. If the population of parity zero has shifted towards ages that have high progressions to the first birth, the first birth rate will rise. If the population of parity one has shifted to ages and/or intervals since first birth that have higher progressions to the second birth, then the second birth rate will rise. The above empirical examples can be extended to higher order parities and we have done this for Australia. We observed that age-specific birth probabilities by parity and interval since the previous birth had hardly changed at all over an extended period of time in Australia for all parities from the second onwards, confirming the result previously observed in relation to age and parity (Kippen 2004).

Thus, almost all of the Australian action in relation to changes in the timing of births has been captured in age-specific progressions from parity zero to parity one and age-specific progressions from parity one to parity two by interval since the first birth. However, this empirical result for Australia may not apply to other countries or to other periods of time. In general, we conclude that past changes in birth timing are fully captured by a time series of matrices of age-specific birth probabilities by parity and interval since the previous birth. Furthermore, the compositional impact of changes in the timing of births in the past is captured in the current age-specific composition of the female population by parity and interval since the previous birth. Just as the current age distribution of a closed population can be derived from its history of age-specific fertility and mortality rates over approximately the past 80 years, the current age-specific composition of the female population by parity and interval since previous birth is derived wholly from matrices of age-specific birth rates by parity and interval since previous birth over approximately the past 35 years—ignoring the negligible impacts of differential mortality and migration.

The Intrinsic Total Fertility Rate

It follows from the above that if a constant matrix of age-specific birth probabilities by parity and interval since last birth is applied to the current composition of the female population by the same characteristics, each year over about 35 years, the composition of the population will change to the composition that is intrinsic to the set of rates in the matrix and this composition will become stable. The impacts of the original composition will be lost gradually and the PTFR will become constant. As we have argued above, the future potential effects of past changes in the timing of births are captured in the current age-specific composition of the female population by parity and interval since last birth. Thus, we can measure the effects of past timing changes by applying the current matrix of age-specific probabilities by parity and interval since last birth to the current composition of the population and holding the matrix constant for about 35 years. The PTFR that results at the end of the 35-year period, we call the Intrinsic Total Fertility Rate (ITFR). This is a direct measure of the effects of past timing chnages that has a very explicit meaning and interpretation. It shows what the PTFR would be if current behavior (by age, parity and interval since last birth) were to be held constant in time over a long period. Another interpretation of ITFR, consistent with the arguments of Rallu and Toulemon (1994) and Ni Bhrolchain (1992), is that it is simply a better period measure of fertility than the age-based PTFR because it controls not only age composition but also parity and interval since the previous birth.

An empirical example

To demonstrate that the concept of the ITFR is of more than purely theoretical value, we have calculated the ITFR for Australia from 1990 to 2005. Figure 6 shows the Period Total Fertility Rate and the Intrinsic Total Fertility Rate for each year from 1990/91 to 2004/05. At the beginning of the period, 1990/91, the two rates differ by 0.15 births per woman indicating that the composition of the population in this year was such that PTFR would rise by this amount if the probability matrix (behavior) were held constant. In fact, the PTFR remained constant for two years and then fell for the next eight years. This can be interpreted to mean that the behavioral change in fertility during these years was more strongly negative than is indicated by the observed fall in the PTFR. This interpretation is confirmed by the observation that, during these years, the ITFR fell more sharply than did the PTFR. The trend in ITFR provides a better indication of behavioral change because the temporal impact of population composition has been removed. In other words, we can conclude that timing changes did not stop in 1990 but continued through most of the 1990s.

For a short period, 1997-2003, the ITFR tends to flatten out, a result that indicates that behavior was relatively constant during these years and that the difference between the two rates was due only to the compositional differences. In the most recent year, both PTFR and ITFR rose with the rise in PTFR seemingly sharper than the rise in ITFR. This is suggestive of a turning point in behavior, probably a slight shift to earlier first births for younger cohorts.

If the 2000/01 probability matrix was to remain constant over a 35-year period, the PTFR would rise by only 0.07 births per woman indicating that the changes in behavior that occurred during the 1990s had shifted the population composition in such a way that it was closer to the intrinsic composition than had been the case ten years earlier.

This interpretation is confirmed by a comparison of Figures 7 and 8. These figures show the difference between the PTFR and the ITFR decomposed by age and parity. For 1990/91, most of the difference between the PTFR and the ITFR is made up of first and second births that were yet to occur (as would be expected when the first birth is delayed). With constant behavior, first births would have increased at all ages but second births would have continued to fall slightly up to age 25 but to rise above age 25. Figure 7 also shows the falls in third and higher order births that would occur at younger ages but the rises in the contribution of these birth orders at higher ages, again assuming constant behavior. In contrast, the differences between the age-parity distributions of the 2000/01 PTFR and the 2000/01 ITFR are much smaller overall and this applies to all birth orders. Compared to Figure 7, the distribution of the differences between PTFR and ITFR shifts strongly to older ages with the positive impact being concentrated at ages 35-39 rather than at 30-34. Delay of the second birth at younger ages seems to have largely disappeared by 2000/01. We can conclude from these two figures that the delay of first births at younger ages had already ceased by 1990/91 and, by 2000/01, the delay of second births at younger ages had also ceased. However, in 2000/01, there was still an impetus for future increase in fertility related to first and second births previously delayed, but at a much lower level than was the case in 1990/91. There was little positive impetus beyond 2000/01 related to delayed third or higher order births and, in fact, with constant behavior, the future contribution of births of these orders is a net negative as these birth orders would be falling for third births under age 33, for fourth births under age 38 and for fifth and higher order births at all ages.

Importantly, this methodology is explicit about the impact of delays of lower order births on future births at higher orders. This is the central area in which the parameterized methods for estimating tempo-adjusted TFR have been criticized. The example clearly shows the lower likelihood of higher order births in Australia in the future because of the delay of first and second births. Also, the methodology indicates that it is important to consider the full distribution of births by age, not simply the mean age. For example, we have concluded that the *behavioral* phenomenon of delay of the first birth ended early in the process of change in PTFR, by 1990. Subsequent increases in the mean age at first birth were due not to any further movement to delay but to previously delayed births occurring at older ages. This is a subtle distinction, but an important one because it suggests that the use of mean ages as parameters in tempo-adjusted measures may be problematic.

Overall, we conclude that the long-term delay of births in Australia has run its course and that we are now in a period that reflects only its consequences—more first and second births at older ages and fewer higher parity births. For the time being, we can expect the impact of the former (more first and second births) on PTFR will be higher than the impact of the latter (fewer births of parity three and above). In the absence of further behavioral change, we expect the Australian PTFR to settle around 1.85 births per woman, not very different from its 2005 level. We also expect that the PTFR and the ITFR will be very close to each other, a situation indicative of longer-term stability. Experience shows, however, that birth timing does not remain still for very long and there are initial signs that there may have been a small fall in the age at first birth in Australia in 2005, a new direction for behavior.

Comparison of ITFR, the tempo-adjusted TFR and cohort completed fertility

Figure 9 compares the trends in the ITFR and the PTFR with the Bongaarts and Feeney tempo-adjusted measure for Australia from 1990 to 2000. From the previous comparison of ITFR and PTFR, we concluded that there was a distinct change in behavior between 1992 and 1997: a downward shift in fertility that was evident from both measures but was more clearly evident from the trend in ITFR. In contrast, the Bongaarts and Feeney measure is relatively flat during this period, a trend that is difficult to relate to observed fertility behavior in these years. Completed cohort fertility (Figure 10) has been falling almost linearly for cohorts born from 1948 to 1970 at the rate of about 0.02 births per woman per cohort. Based on the above analysis, we would expect cohort completed fertility to continue to fall as there are fewer third and higher order births in the future consequent upon previous delays of first and second births. This continuous decline also seems to be inconsistent with the relative constancy of the Bongaarts and Feeney measure. It should be noted that our assessment of the future of cohort fertility is based upon our analysis of period fertility trends including an analysis of the changing composition of the population.

Concluding remark

The Intrinsic Total Fertility Rate is a powerful tool to interpret the impacts on period fertility of past changes in the timing of births. It is especially powerful when it and its components are compared with trends in the PTFR and its components. This is because movements in the ITFR and its breakdown by age and parity are readily interpretable. This is less the case with previously defined tempo-adjusted measures of the Total Fertility Rate although their purpose is somewhat different. In reality, the ITFR is simply a better measure of period fertility as it takes account of the simultaneous composition of the population by age, parity and interval since the previous birth. This is an advance on the PTFR that controls only for age composition. The probability matrix that underlies the ITFR is also a better basis for projection of births. While the methodology is very data intensive – the probability matrix has over 2000 cells – computation is straightforward for countries that have the necessary data.

Figure 6. Total Fertility Rate and Intrinsic Total Fertility Rate, Australia, 1990/91–2004/05

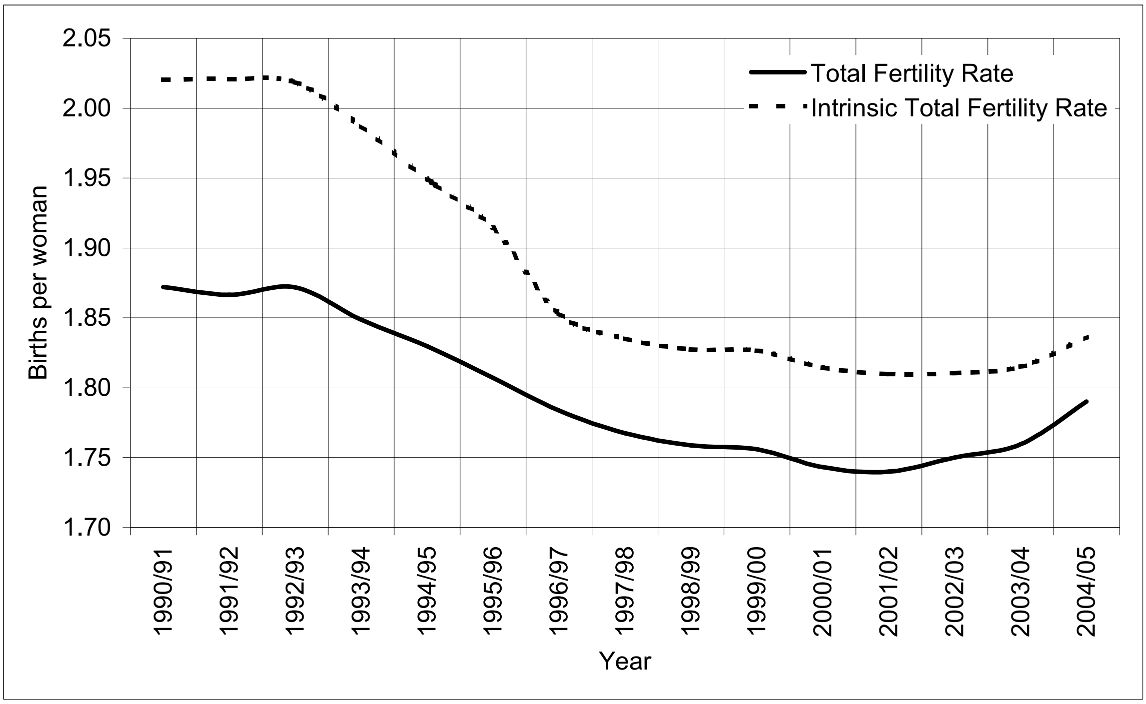

Figure 7. Difference between Total Fertility Rate and Intrinsic Total Fertility Rate decomposed by age and birth order, Australia, 1990/91

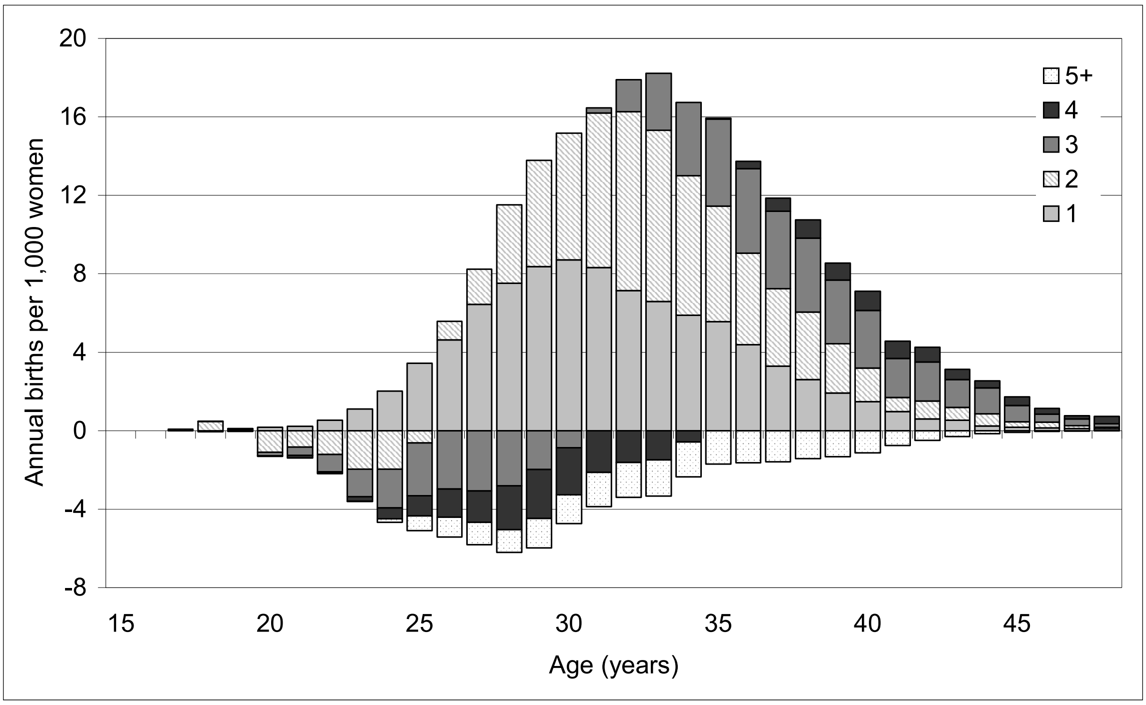

Figure 8. Difference between Total Fertility Rate and Intrinsic Total Fertility Rate decomposed by age and birth order, Australia, 2000/01

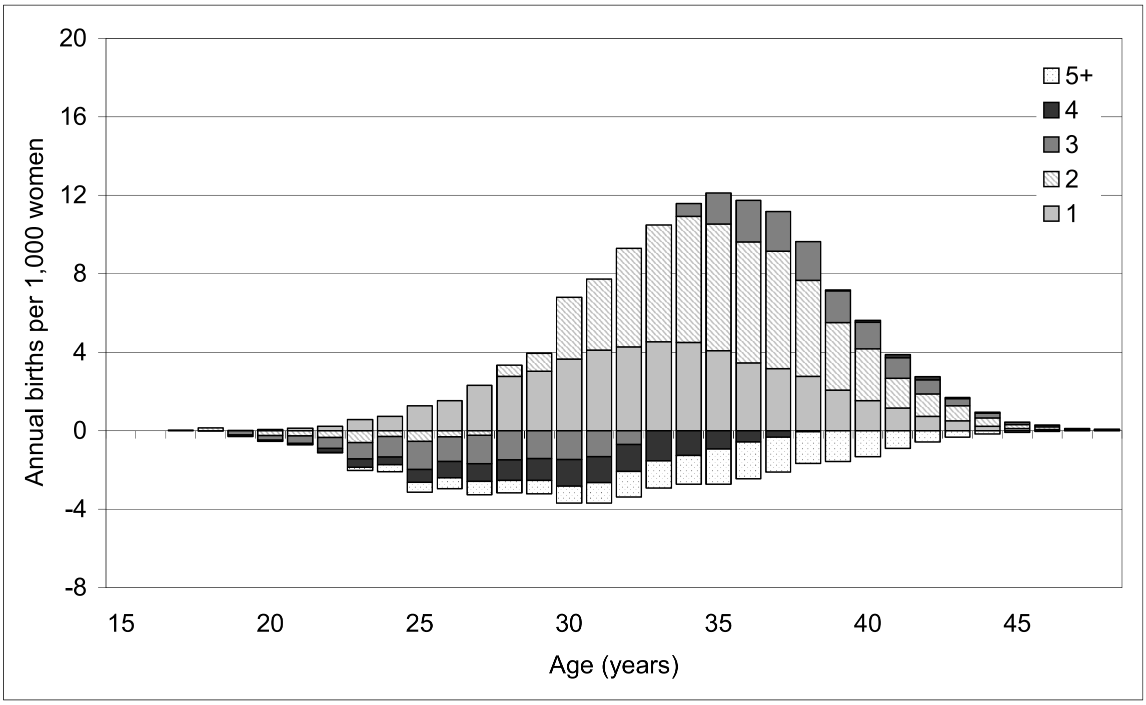

Figure 9. Total Fertility Rate, Intrinsic Total Fertility and Tempo-adjusted Total Fertility Rate (Bongaarts and Feeney), Australia, 1990/91–2000/01

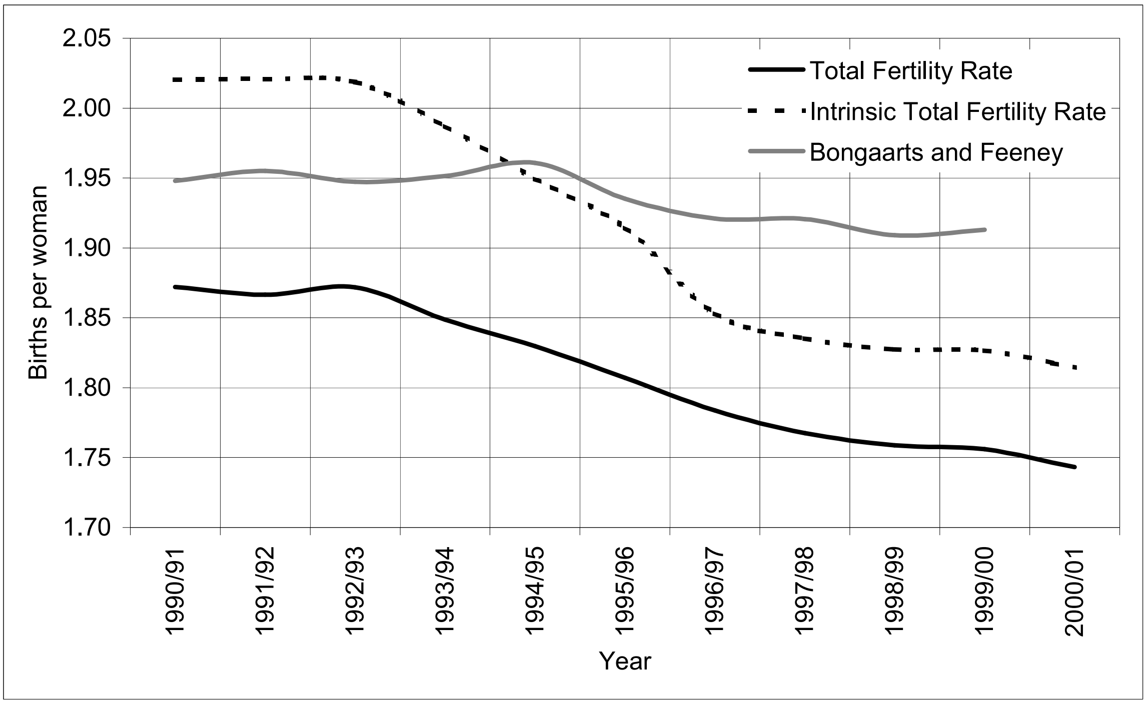

Figure 10. Cohort Fertility Rate, Australian women born 1930–1970

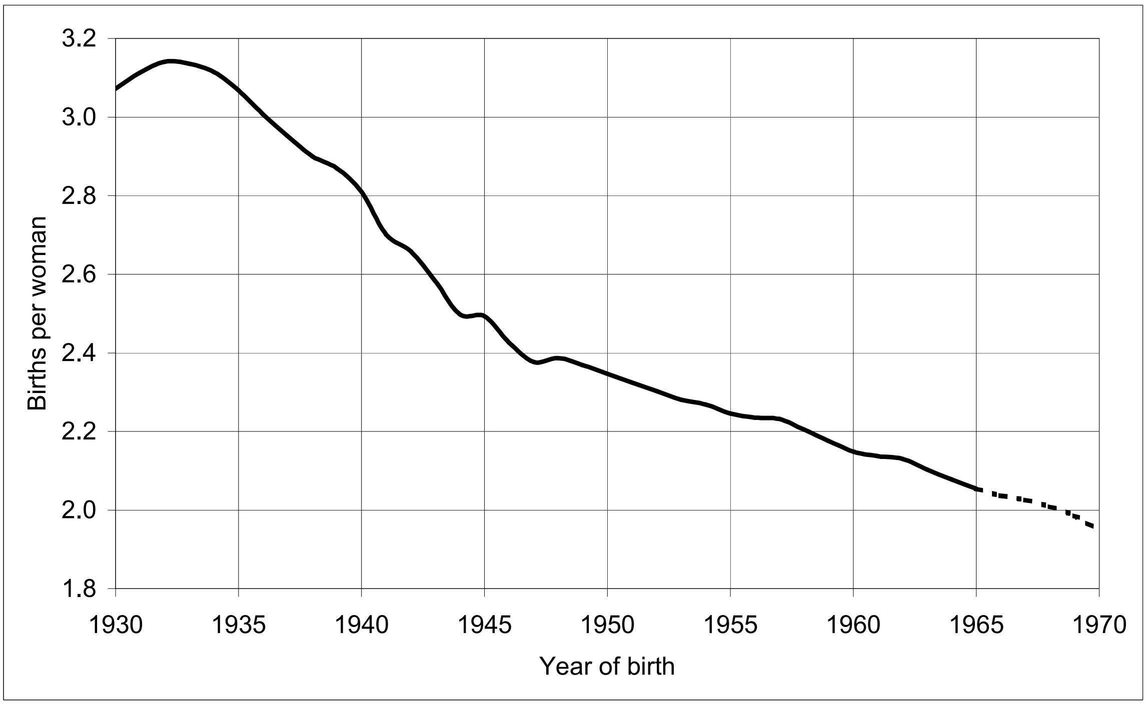

References

Bongaarts, J. and G. Feeney. 1998. "On the quantum and tempo of fertility," *Population and Development Review* 24: 271–291.

Corr, P. and R. Kippen. 2006. "The case for parity and birth-order statistics," *Australian and New Zealand Journal of Statistic* 48(2): 171–200.

Frejka, T. and G. Calot. 2001a. "Cohort reproductive patterns in low-fertility countries," *Population and Development Review* 27: 103–132.

Frejka, T. and G. Calot. 2001b. "Cohort reproductive patterns in the Nordic countries," *Demographic Research* 5: 125–186.

Frejka, T. and J-P. Sardon. 2004. *Childbearing Trends and Prospects in Low-fertility Countries: A Cohort Analysis*. Dordrecht: Kluwer Academic Publishers.

Kippen, R. 2003. "Trends in age- and parity-specific fertility in Australia," *Working Papers in Demography*, No. 91, Demography and Sociology Program, The Australian National University, Canberra.

Kippen, R. 2004. "Declines in first and second birth rates and their effect on levels of fertility," *People and Place* 12(1): 27-36.

Kohler, H-P. and J. A. Ortega. 2002. "Tempo-adjusted period parity progression measures, fertility postponement and completed cohort fertility," *Demographic Research* 6: 91–144.

Lesthaeghe, R. 2001. "Postponement and recuperation: recent fertility trends and forecasts in six Western European countries," paper presented at the International Perspectives on Low Fertility Seminar, International Union for the Scientific Study of Population, Tokyo, 21–23 March.

Lesthaeghe, R. and G. Moors. 2000. "Recent trends in fertility and household formation in the industrialized world," *Review of Population and Social Policy* 9: 121–170.

Lesthaeghe R. and P. Willems. 1999. "Is low fertility a temporary phenomenon in the European Union?" *Population and Development Review* 25: 211–228.

Ni Bhrolchain, M. 1987. "Period parity progression ratios and birth intervals in England and Wales, 1941-1971," *Population Studies* 41:103-125.

Ni Bhrolchain, M. 1992. "Period, paramount? A critique of the cohort approach to fertility," *Population and Development Review* 18: 599-629.

Rallu, J-L. and L. Toulemon. 1994. "Period fertility measures: the construction of different indices and their application to France, 1946–89," *Population: An English Selection* 6: 59–93.

Rodríguez, G. 2006. “Demographic translation and tempo effects: an accelerated failure time perspective,” *Demographic Research* 14: 85–110.

Ryder, N. 1980. “Components of temporal variations in American fertility,” in R. Hiorns (ed.), *Demographic Patterns in Developed Societies*. London: Taylor and Francis, pp. 15-54.

Schoen, R. 2004. “Timing effects and the interpretation of period fertility,” *Demography* 41(4): 801-819.

Sobotka, T., W. Lutz and D. Philipov. 2005. “‘Missing births’: decomposing the declining number of births in Europe into tempo, quantum, and age structure effects,” *European Demographic Research Paper*, Vienna Institute of Demography.

Van Imhoff, E. 2001. “On the impossibility of inferring cohort fertility measures from period fertility measures,” *Demographic Research* 5: 23-64.

Zeng, Y. and K. C. Land. 2002. “Adjusting period tempo changes with an extension of Ryder’s basic translation equation,” *Demography* 39: 269–285.