

Tempo Effects

A hot topic in Demography concerns the need to adjust period measures of fertility, nuptiality and mortality for so-called “tempo distortions”. We will review quickly some of the main ideas, including Ryder’s demographic translation formula and the Bongaarts-Feeney tempo-adjusted measures.

Fertility

Imagine a surface $f(a, t)$ of fertility rates by age and period. A period summary is obtained by summing over ages for a fixed time. In particular, the period TFR for year t is

$$TFR(t) = \int f(a, t) da$$

It is also useful to define the mean age $\mu_p(t)$ of the fertility schedule as

$$\mu_p(t) = \int af(a, t) da / TFR_p(t)$$

Cohort summaries are obtained by summing across a diagonal, where age and time vary together. In particular, the cohort TFR for the cohort born in year t is

$$TFR_c(t) = \int f(a, t + a) da$$

and the cohort mean age of childbearing is

$$\mu_c(t) = \int af(a, t + a) da / TFR_c(t)$$

Ryder

Ryder’s chief concern was that period summaries provide a distorted view of the behavior of cohorts when fertility is changing. In particular, if women delay childbearing the period TFR will drop even if the cohorts have the same number of children as before, so the cohort TFR stays constant. Thus, a cohort change in tempo would look from the period perspective as a change in the quantum of fertility!

The following artificial example may help fix ideas. Consider cohorts having children in three age groups as follows:

	3.0	3.0	3.0	3.0	3.0	3.0	← Cohort TFR
35-44	0.6	0.6	0.6	0.7	0.8	0.8	
25-34	1.6	1.6	1.7	1.8	1.8	1.8	
15-24	0.8	0.6	0.4	0.4	0.4	0.4	
Period TFR →	3.0	2.8	2.7	2.9	3.0	3.0	

Initially cohorts have an average of 0.8, 1.6 and 0.6 births in each age group, for a total of three children. But then a cohort delays childbearing and has 0.6, 1.7 and 0.7 birth per age group, for a total of three.

Subsequent cohorts further delay childbearing and have 0.4, 1.8 and 0.8 in each age group, for a total of three. The cohort TFR is always 3.0. But as shown in the table the period TFR drops from 3.0 to 2.8 and 2.7 before it recovers to 2.9 and bounces back to 3.0. A more accurate cohort-to-period conversion would use Lexis triangles, but the simple calculation shown here suffices to show that a cohort change in tempo appears as a period change in quantum!

Ryder used a first order Taylor series expansion to relate period and cohort TFRs. He showed that for the cohort that reaches its mean childbearing age μ at time t (the cohort born at $t - \mu$)

$$\text{TFR}_c(t - \mu) \approx \frac{\text{TFR}_p(t)}{1 - r_c(t - \mu)}$$

where $r_c(t - \mu)$ is rate of change or time derivative of cohort mean age of childbearing for the cohort reaching its mean childbearing age at time t . This remarkable formula shows that, to a first order of approximation, if cohorts postpone childbearing the period TFR will fall *below* the cohort TFR by an amount that depends on how fast the mean age of childbearing was increasing. This actually happened during the baby boom, as you can see in the computing logs.

Bongaarts-Feeney

In 1998 Bongaarts and Feeney proposed a tempo-adjusted total fertility rate, usually denoted TFR^* , based on an expression that looks remarkably like Ryder's translation formula

$$\text{TFR}^*(t) = \frac{\text{TFR}(t)}{1 - r_p(t)}$$

There are, however, a few important differences. First, $r_p(t)$ is the rate of change or time derivative of the *period* mean age of childbearing at time t . This is much easier to calculate from available data. It is usually estimated by averaging the "in and out" changes between $t - 1$ and t and between t and $t + 1$. This turns out to be exactly the same as half the change between $t - 1$ and $t + 1$.

Second, TFR^* is not a cohort rate, but rather a pure-period measure representing tempo-corrected fertility. This raises issues of interpretation that we discuss below.

A third difference is that B-F recommend applying the procedure separately by birth order, using rates that divide births of a given order by all women. The reasoning behind this approach is that as women have fewer high-order births the overall mean age of childbearing will decline without any changes in the timing of earlier births, so order-specific means provide a better measure of tempo changes. In my own opinion, order-specific fertility is best analyzed using true hazard rates, a point made by van Imhoff and Keilman in comments to the original B-F paper.

There has been a lot of discussion of TFR^* and some confusion about its meaning. B-F argue that they are not trying to estimate the TFR for any particular cohort and that TFR^* is just a "period measure purged of tempo distortions". The best way to think about this is as a counterfactual estimate of what the period TFR would be if women were not delaying childbearing. Zeng and Lang show that it can also

be interpreted at the TFR that would be observed if women followed a period schedule that is shifting constantly to older ages. A simple derivation of these results may be found in my tempo paper, which also provides a result for the mean age of childbearing under the implied period-shift model, which is shown as equivalent to an accelerated failure time model of cohort behavior.

The online computing logs show an application of these ideas to U.S. fertility using the Heuser cohort fertility tables. For a much more detailed analysis see the paper by Schoen in *Demography* in 2004.

Nuptiality

The same phenomenon we have noted with fertility can happen with nuptiality. If women postpone first marriage, then period estimates of the proportion who eventually marries will decline even if the same fraction of each cohort ends up marrying. Thus, a cohort change in tempo can masquerade as a change in period quantum.

Bongaarts and Feeney apply their procedure working with period marriage frequencies, obtained by dividing first marriages by the total number of women in an age group (not just those single). They accumulate these frequencies to obtain a Total First Marriage rate (TFMR), and also use them to calculate a period Mean Age at Marriage. They then define a tempo-adjusted TFMR as

$$\text{TFMR}^* = \frac{\text{TFMR}(t)}{1 - r_p(t)}$$

where $r_p(t)$ is the rate of change or time derivative of period mean age at first marriage. The procedure is formally identical to the adjustment used for fertility.

Note that this approach relies on frequencies rather than true event-exposure rates, and this causes some technical difficulties. If 60% of a cohort marries before age A, and 60% of the next cohort marries after age A, and we combine these frequencies, we would get a synthetic cohort where 120% marry! This couldn't happen with hazard rates, but the model is predicated on a shift of the period schedule of frequencies (or equivalently, cumulative proportions married).

The approach also assumes that women postpone first marriage by the same amount of time at all ages. If we observe fewer women marrying in a given year it could be because some will marry later and/or because some will forego marriage, and it is hard to determine the relative weight of these two explanations. Bongaarts and Feeney can separate the two effects by assuming a uniform delay at all ages. The quality of the adjustment depends on the validity of this assumption.

Mortality

In a more recent series of papers Bongaarts and Feeney extended their proposed tempo adjustment to mortality. They claim that conventional period life expectancy is a biased measure of longevity when mortality is declining, with a bias of up to 2 years in developed countries. Needless to say, they created quite a stir in the demographic community. With fertility (and nuptiality) we could all understand the risk of confusing changes in quantum and tempo, but with mortality the quantum is fixed, only tempo

can change, and no one would mistake one for the other. In other words if mortality rates decline, we know it is because people are delaying death.

B-F claim, however, that period measures of mortality suffer from the same “tempo distortions” as period measures of fertility, and propose an adjustment based on an estimate of the rate at which mortality is declining. To motivate the need for adjustment they use an example along the following lines. Suppose in a given year we all took a pill that made us immune to death (and aging) for three months. Clearly such a pill would add exactly three months to our life. Yet the death rates for that year would decline 25%, and in a country such as the U.S. in 2002 period life expectancy would rise by about 3.6 years, overestimating the gain in longevity.

We should remember, however, that conventional life expectancy is a counterfactual estimate of how long we would live if the rates observed in a given year remained in effect through our lives. Effectively that assumes that we would get a magic pill every year, in which case we would indeed live quite a bit longer. However, the example serves to illustrate a key feature of the B-F approach, the assumption that adult mortality declines because we all receive “increments to life”, not because “rates decline”. The underlying model is formally identical to the model for fertility and nuptiality, assuming a uniform shift in the survival curve to older ages. This is not realistic for the youngest ages, but Bongaarts and Feeney restrict their discussion of tempo effects in mortality to adult ages, say above age 25 or 30.

The gist of the method relies on death frequencies, computed by dividing deaths in an age group by the original size of the cohort (not just those alive). Accumulating these leads to the Total Mortality Rate (TMR). The rate of change or time derivative of the TMR is used to compute an adjustment factor, that is then used to inflate the age-specific mortality rates before calculating life expectancy. The result is the B-F tempo-adjusted life expectancy. Other interesting measures that come up are the cohort average length of life (CAL), and the standardized mean age at death; in addition, of course, to conventional life expectancy.

We will not discuss these further, as the dust has not settled. The book *How long do we live?* edited by Barbi, Bongaarts and Vaupel, has a collection of papers giving different views on this issue, including my own. It turns out that when adult mortality follows a Gompertz model it is impossible to distinguish a period shift to older ages from a proportionate decline in rates at every age. However, the two models—“reduction in rates” as opposed to “increments to life”—have different implications for the future, with the latter implying that if gains in longevity were to stop age-specific mortality rates would rise.

Conclusion

Everyone agrees that tempo effects exist. Mortality is a pure tempo phenomenon, as death is bound to occur and the only question is when. Things are different with fertility and nuptiality because the event in question may or may not occur. Under these circumstances a period change in quantum may reflect a change in cohort quantum, a change in cohort tempo, or an unknown mixture of the two. Distinguishing the two while the cohorts are still “making up their minds” is a tall order. The B-F adjustment removes the tempo component of the change under a model that assumes a uniform delay at all ages. Whether this adjustment is meaningful in the case of mortality remains a hotly debated issue.