

Sources of error and bias in methods of fertility estimation contingent on the P/F ratio in a time of declining fertility and rising mortality

Tom A Moultrie and Rob Dorrington
Centre for Actuarial Research (CARE), University of Cape Town

The most commonly used indirect fertility estimation methods rely on the use of the P/F ratio, first proposed by Brass. As originally conceived, the ratio compares cumulated cohort fertility with cumulated period fertility on the basis of three, fairly strong, assumptions: First, that fertility rates have been constant over time; second, that the age distribution of fertility has been constant; and third, that the fertility of women who do not survive to report their numbers of children borne does not differ from those who do survive. The intention of this paper is to interrogate what happens to the results produced by the P/F ratio method as each of these three assumptions is violated, first independently, and then concurrently. These investigations are important given the generally poor quality of census data collected in developing countries, particularly sub-Saharan Africa, and the radically altering demographic conditions associated with a generalised HIV/AIDS epidemic in the region.

1. Introduction

Improvements in vital registration systems and census methodologies, together with the expansion of coverage of the Demographic and Health Surveys have, apparently, all but obviated the need for indirect techniques of fertility and mortality estimation in most regions of the world. Although still necessary throughout sub-Saharan Africa, the majority of these techniques were developed in the 1960s and 1970s, on the assumptions that fertility and mortality were constant (which was, roughly, the case in countries to which these techniques were applied at the time the methods were developed).

While some of the techniques were subjected to further refinement in the 1980s (for example, with the evolution of the Relational Gompertz fertility model; and the variable- r mortality estimation methods), the improvements in data (number of sources, methodologies, and surveys), combined with the shift in research priorities after the 1994 Cairo Conference on Population and Development, have meant that little attention has been paid to the continued validity of the methods under radically altering demographic conditions. This paper aims to address one aspect of this lacuna in our knowledge, at least with regard to the methods of estimating fertility from census-type data.

The most commonly used method of estimating fertility indirectly relies on the use of the P/F ratio, first proposed by Brass (1964; 1968). In essence, the ratio compares cumulated cohort fertility with cumulated period fertility on the basis of three, fairly strong, assumptions. First, that fertility rates have been constant over time. Second (and related closely to the first) that

the age distribution of fertility has been constant, and third, that the fertility of women who do not survive to report their numbers of children borne does not differ from those who do survive. The intention of this paper is to interrogate what happens to the results produced by the P/F ratio method as the first two assumptions are violated, first independently, and then concurrently. Implications of a violation of the third assumption are discussed by means of general reasoning.

The paper comprises three main analytical sections, each of which seeks to shed light on the implications of a violation of the assumptions on the results produced by the method. The first section does so by means of mathematical reasoning, using the continuous form representation of the P/F ratio at time t and age x as a starting point and developing this expression as far as possible, although this analysis, without further hugely simplifying assumptions, rapidly becomes intractable. Nonetheless, the exercise (futile as it proves) is a worthwhile one.

The second approach develops the analysis a little further by means of general reasoning, while the third approach uses a simulation model with variable inputs regarding age patterns and trends in mortality and fertility to establish how the P/F ratio method behaves when many demographic parameters are changing simultaneously. Several scenarios are presented: the first indicates the ‘ideal’ conditions of constant fertility and mortality. The results of this simulation are presented simply to show that the simulation model works. Subsequent simulations impose, independently and then together, changes in the age distribution of fertility; changes in the total fertility rate (an increase followed by a long-term decline), indicating a stylised demographic transition); changing levels and age distributions of mortality; and, finally, allowing for different fertility distributions and levels by mortality risk (as might typically be found in a country experiencing a generalised and extensive HIV/AIDS epidemic). The results from the three investigations are synthesised and interpreted in the conclusions to the paper.

At its core, the paper seeks to answer the following question; how well does the P/F ratio method estimate fertility in these simulated environments? An ancillary question asks which of several recommended variants of the P/F ratio tend to work best in different environments.

2. Brief description of the P/F ratio method

The P/F ratio method is the foundational technique for most approaches to measuring fertility indirectly. If fertility has remained constant (in terms of both level and age distribution) over an extended period of time, and is the same for women who survive and don’t survive to the

time of the survey, then the P/F ratio should, with perfect recall, be equal to one. Hence, according to the original formulation of the method, deviations in the ratio of P to F could be used diagnostically to understand the data better, and, as is the case in most circumstances, to use this information to derive a better estimate of fertility that corrects for errors in the reporting of cohort or lifetime fertility.

One of the primary uses of the P/F ratio method, then, is to compensate for errors of recall which are supposed to occur in the reporting of fertility in developing countries. The developers of the P/F ratio method further argue that the typical error is that women's recent fertility in censuses tends to be under-reported and that this is independent of their age: in other words, reported period fertility may be too low, but the shape of the fertility distribution can be assumed to be correct. By contrast, older women are presumed to be more likely to incompletely recall the number of children ever borne. If these conditions hold, then the extent of underreporting of lifetime (cohort) fertility, and hence the 'true' level of fertility can be estimated by scaling up the observed Total Fertility Rate by the ratio of P/F.

However, when fertility is falling, period fertility will be lower than lifetime fertility, and hence one would expect to see, under conditions of complete reporting of births, increasing departure of the P/F ratio from one with increasing age, which is difficult to distinguish analytically from underreporting of lifetime fertility. It is implicit in the description of the method outlined above that, if current fertility were underreported, but lifetime fertility were correct, there should be a uniform correction factor (the constant P/F ratio). But what does one do in the situation outlined above where the indications that there is both a decline in fertility underway as well as underreporting of current fertility? In this case, the P/F ratio will increase with age; and there is no obvious choice of scaling factor to be applied to the observed current fertility data to compensate for the generalised underreporting of recent fertility.

Related to this, Brass and Airey (1988:9) suggest that a "serious problem with the P/F ratio method is its vulnerability to changing fertility patterns. With changes in the level of fertility, the shape of the fertility distribution tends also to change." How much of problem this changing age distribution of fertility poses to the method is an important component of the investigations conducted here.

Manual X (United Nations 1983) seeks to get round this problem by recommending that the P/F ratio for the 20-24 year old group (which is assumed to be mostly unaffected by recall error) be used to adjust reported period fertility data when fertility is falling. An alternative also

mentioned is to use the average P/F ratio for the 20-24 and 25-29 age groups as the scaling factor*.

Yet another alternative approach to the problem of declining fertility in applications of the P/F ratio method is that suggested by Feeney (1998). Feeney's reconceptualisation of the P/F ratio method to render it compatible with declining fertility is particularly elegant in its reasoning. He makes use of an empirical observation made by Norman Ryder (1964, 1983) that the average parity (i.e. cumulated cohort fertility) of a given cohort of women presently aged at the mean age of childbearing (\bar{m}) is a close approximation to the period fertility that prevails at the time of the survey. On the basis of this observation, Feeney demonstrates how the P/F method can be reconceptualised (although mathematically unchanged) in such a fashion as to make redundant the need for the assumption of constant fertility over time, with the P/F ratio applicable at the mean age of childbearing being the preferred scaling factor for the observed fertility rates. A further benefit of this approach, Feeney points out, is that a time series of other estimates of fertility can be attained simply by applying P/F ratios applicable to each of the conventional age groups (15-19 ... 45-49) to the observed fertility rates, with those rates applying to times $\bar{m} - 17.5$; $\bar{m} - 22.5$... $\bar{m} - 47.5$ years after the survey date. How well this reconceptualisation works is one of the questions addressed in the paper.

It must be appreciated, however, that in any practical application, one is implicitly trading off the portrayal of a real fertility decline combined with a generalised underreporting of lifetime fertility against long-term constant fertility, since the former combination could generate an outcome that is indistinguishable from the latter. Since, for the most part, we will never actually know the true level (or indeed whether it is constant or not by age), this problem is more intractable than most†. This problem, which is significant, is not the subject of this paper.

3. Analytical processes

a. Mathematical reasoning

Let $f_p(a,t)$ be the (period) fertility rate of women aged a at time t . Then, the cumulative fertility of a cohort of women aged a at time t , $P(x,t)$ is given by

* For the rest of this paper, the conventional Brass P/F ratio referring to the fertility of women aged 20-24 is referred to as Brass(2). The average ratio of women 20-29 is referred to as Brass(2,3). The index refers to the age groups, where 1 refers to 15-19, 2 to 20-24 etc.

† Over time, of course, one could get a measure of the level of past fertility by applying reverse survival methods to the enumerated population at various ages. This presupposes knowledge of child survival, however, which may not be known, and accuracy of census counts, which may not be the case.

$P(x,t) = \int_0^x f_p(a,t-x+a)da$, and the cumulated period fertility, $F(x,t)$ will be

$$F(x,t) = \int_0^x f_p(a,t)da$$

$$\begin{aligned} \text{But, by definition, } f_p(a,t) &= \frac{B(a,t)}{W(a,t)} \\ &= \frac{B(a,t)}{\left(\frac{1}{1+s}\right)B(t)\exp\left[-\int_0^a r(y,t)dy\right]p(a,t)} \end{aligned}$$

Where s is the ratio of male to female births, $p(a,t)$ is the cohort probability that females born at time $t-a$ will survive to age a (at time t) and $r(y,t) = \partial \ln N(y,t) / \partial t$. Substituting this into the form of the P/F ratio, we get

$$\begin{aligned} \frac{P(x,t)}{F(x,t)} &= \frac{\int_0^x \frac{B(a,t-x+a)\exp\left[\int_0^a r(y,t-x+a)dy\right]}{B(t-x+a)p(a,t-x+a)} da}{\int_0^x \frac{B(a,t)\exp\left[\int_0^a r(y,t)dy\right]}{B(t)p(a,t)} da} \\ \text{And } B(t-(x-a)) &= B(t)\exp\left[-\int_0^{x-a} r(y,t)dy\right] \\ \frac{P(x,t)}{F(x,t)} &= \frac{\int_0^x \frac{B(a,t-x+a)\exp\left[\int_0^a r(y,t-x+a)dy\right]}{\exp\left[-\int_0^{x-a} r(y,t)dy\right]p(a,t-x+a)} da}{\int_0^x \frac{B(a,t)\exp\left[\int_0^a r(y,t)dy\right]}{p(a,t)} da} \quad (1) \end{aligned}$$

If $f_p(a,t) = f_p(a) \quad \forall t$ then the ratio $P(x,t)/F(x,t) = 1$ for all t (whether or not the population is stable).

Equation (1) offers the completely generalised formulation of the P/F ratio method, expressing the ratio under conditions of non-constant fertility, mortality and population growth. However, it is not particularly useful, as the magnitude and direction of any effects on the P/F ratio of changing those demographic assumptions cannot be resolved or identified readily. As Coale (1972: 208) notes, expressions such as this may not be conducive to further analysis or manipulation:

It is quite possible that there is a general principle at work here: The age distribution effects of certain simple specified sequences of fertility or mortality changes can be visualized, but the general case will always remain merely calculable, but not readily understandable at an intuitive level.

b. General reasoning

Mathematically, as shown above, the method requires that fertility be constant over time for period and cohort fertility to be the same. Assuming no errors in recall of fertility, a reduction in fertility in the past 35 years would imply a P/F ratio greater than one. If that fertility decline was monotonic, it would further imply a series of P/F ratios that increases monotonically with age, with the greatest departure from unity at the oldest ages. By itself, this observation is not of practical use in estimating the level of fertility. However, if one maintains the assumption of constant fertility, but allows for underreporting of both current and cohort fertility, the relationship can be used to estimate the level of fertility that would pertain if fertility were constant.

To do so, two additional assumptions inherent to the Brass method are required; first that reports of current (i.e. period) fertility are underreported by a constant factor independent of age (i.e., women of any age are as likely to report inaccurately the number of births she has had in the preceding 12 months). If fertility were constant and unchanging (and reports of lifetime fertility completely accurate) then the P/F ratios would indicate the factor by which the cumulated period fertility rates should be multiplied to provide an accurate estimate of fertility.

However, in parallel with underreported current fertility, Brass surmises that women's reports of their lifetime fertility will also tend to be underreported, but that – unlike current fertility – this will be underreported differentially by age, with older women systematically underreporting their fertility more than younger women. This is believed to occur because older women have had more children and suffer greater lapses in recall about their lifetime fertility; a greater proportion of whom may have died or no longer be living with their mother. In such cases, the P/F ratio, assuming perfectly reported current fertility will deviate systematically from unity, and will decline with age.

Underreporting of current fertility

What implications does a degree of underreporting of current fertility have on the results? For the most part, the effect is trivial. With the exception of Feeney's approach, the methods all call for using the P/F ratios pertaining to the youngest ages. At these ages, cohort fertility is underreported to the least extent, while current fertility may be underreported (possibly significantly – Moultrie and Timæus (2003) observed that after removing clearly erroneous reports of current fertility in the 1996 South African census, only approximately one in two births occurring in the twelve months before the census were actually reported to enumerators).

If cohort fertility is accurately reported (as is more likely when using the reported parities of younger women), but current fertility is underreported by x per cent, then the P/F ratio will be estimated to be $(1/(1-x))$ times higher, but off a base of fertility estimated to be x per cent too low. Thus multiplying the reported current fertility by the P/F ratio returns the true fertility rates (provided fertility is constant). Thus, other than through the substitution between underreporting and fertility decline mentioned above, underreporting of fertility in and of itself will not, after correction, have an effect on the result, provided lifetime fertility is accurately reported.

Differential mortality from HIV

The effect of violations of the third assumption (that the fertility of those who have died is not different from those who survive to report their fertility) is not investigated analytically here, the reason for not doing so is not that the matter is uninteresting (it is very interesting), but because censuses only collect fertility data of the living, any simulated assessment of the impact of differential fertility based on differential mortality by means of simulation exercises will merely reflect the assumptions used in the simulation.

If it were simply the case that it was only the assumption of constant fertility that was violated, this paper would be less relevant to contemporary demographic research. While rapidly changing mortality in and of itself does not hugely affect the results produced by applications of the P/F ratio methods, a widespread HIV/AIDS epidemic adds greatly to the complexity of population dynamics: the fertility rates of young HIV-positive women are likely to be higher than those of HIV-negative women (a consequence of the selectivity of the onset of sexual behaviour meaning that sexually active young women are at greater risk for both infection and pregnancy), while the fertility of older HIV-positive women is likely to be less than that of HIV-negative women – a consequence of secondary sterility arising from co-infection with other STDs, as well as directly due to lower fecundability and higher rates of foetal loss (Lewis, Ronsmans, Ezeh *et al.* 2004; Zaba and Gregson 1998). HIV-positive women also experience a significantly higher mortality. If, as is generally presumed, HIV-infected women have lower fertility, but are proportionally fewer at the time of censuses and surveys as a result of differentially higher mortality, the implications for the estimation of fertility directly from census data are clear: estimates of national (for example) fertility will be overstated; but probably by quite a small amount; since the missing data would relate to women of reproductive age who had died during the preceding twelve months; and given that fertility is a relatively rare event in the first instance, sub-fecundability of those who die is unlikely to be a significant determinant of the overall level of fertility (although a greater impact may well arise

from the lower aggregate fertility arising from HIV where the epidemic is generalised. However, this last aspect of the general reasoning is made more complex by the different fertility schedules for HIV+ and HIV- women, which only partly determines the ages at which HIV-infected women have lower aggregate fertility than uninfected women. While HIV positive women's fertility may indeed be higher at younger ages, and lower at older ages, the effect of HIV-related sub-fecundability is not only a function of cumulated fertility, but also the average age at death of HIV+ women. This determines whether, at death, infected women have had fewer children than uninfected women of the same age. Clearly if the death from HIV occurs at a sufficiently young age, that would not be the case. Results from the simulations below, in passing, offer some insight into this matter.

Where direct fertility estimation is not possible, the already-marginal effect of the omission of the fertility of dead women from the application of the P/F method is amplified only slightly. The effect on the denominator will be as before, while the numerator will also be overstated, but by a greater amount as the cohort effect arising from sub-fecundability will be more noticeable (even if not severely) at most ages (other than the youngest). The net result of HIV alone is that the underlying fertility rate arising from the application of the P/F ratio method will be under-estimated. However, as we shall see in the following section, since the violation of the other underlying assumptions tends to cause the P/F ratio method to overestimate fertility slightly, the error in the method when fertility is falling in the context of an HIV epidemic arising from HIV is actually attenuated.

c. Simulation

Given the analytical complexity of Equation 1, simulation models offer an alternative approach to understanding what happens to the P/F ratio under different scenarios of fertility and mortality change. This section describes one such simulation model that can be used to cast light on the errors described above.

The method is to set up a long-range single-sex cohort component population projection model. (The restriction to a single-sex population is simply for parsimony; and the assumption of a sex ratio at birth allows ready extrapolation to conventional fertility measures). In order to fully understand the implications of deviations from the assumptions on the results of the method, the projection period of necessity has to be long; the simulations described here begin with a base population, allow stationarity to be reached (at least in the population of reproductive age women), and then allow key parameters to change gradually in both directions before stability is re-attained.

The base population used is the stationary population associated with the Brass General Standard. A fertility distribution is used which is assumed to be typical of pre-fertility transition African fertility distributions (in actual fact that associated with African South African women in the 1950s (Moultrie and Timæus 2002)). Initially, the Gross Replacement Ratio is set to maintain stationarity.

Variations in fertility

In scenarios where fertility is assumed to vary, several modifications to the model are required. First, the age distribution of fertility is assumed to trend towards a fertility distribution appropriate to post transition fertility patterns (in actual fact, for convenience we assumed the pattern observed for White South Africans in 2001). The trend should not be linear (to avoid discontinuities at the beginning and end of the variation), so we assumed it to follow a Weibull mixing function of the starting and ending fertility distributions parameterised such that the cumulative distribution progresses from 0 to (infinitely close to) 1 over a period of fifty years. The assumed standardised fertility distributions are as shown in Figure 1.


The level of fertility is set arbitrarily via the GRR such that fertility is assumed to rise from the fiftieth year of the projection for 25 years, and then decline for a further 50 before a stable fertility rate is again attained. Where required, the effect of HIV is allowed for using the prevalences (and associated mortality rates) indicated by the ASSA model for South Africa (Actuarial Society of South Africa 2005). The effect of HIV on fertility is allowed for by using a (smoothed) schedule of fertility of HIV-positive women observed in Masaka by Zaba and Gregson (1998), and assuming that the fertility rate among HIV-positive women is roughly five sixths of that among HIV-negative women, a ratio that closely reproduces the Percent Attributable Change (PAC) of 0.4 found by Lewis et al (2004) (i.e. for every 10 per cent in the level of population HIV prevalence, population fertility is expected to fall by 4 percent). Caution should be taken in interpreting the results from this simulation: the PAC may not be universally applicable. A potentially more significant problem with the assumption of a constant PAC is that it assumes a level of stability in the mean duration of infection of those infected, while it is unlikely that it does remain constant and probable that the impact of HIV on fertility will be a function of how long women have been infected.

Variations in mortality

Initially, as described above, mortality is assumed to follow the Brass General standard, and then to improve. This is done by using a relational logit transformation holding beta constant ($=1$), but changing alpha linearly by -0.02 per year until infant mortality is less than 50 per 1 000 and then fixing the mortality schedule at that level. The change in mortality is assumed to start at time 50, and mortality reaches its revised stable level after 30 years.

The effect of HIV is allowed for using the (female, African South African) prevalences indicated by the ASSA model for South Africa (Actuarial Society of South Africa 2005), while additional mortality from HIV/AIDS associated with the prevalences is set to the associated single decrement life tables derived from the ASSA all-cause life tables. This mortality is then combined with the underlying mortality schedule to produce the assumed all-cause mortality.

Figure 1 Assumed standardised fertility distributions at times 50 to 100, in intervals of 10 years



The output we are principally interested in is that giving rise to the estimation of the P/F ratio, i.e., the cumulated cohort and period fertility to a given age (again, for simplicity of presentation, data by single age arising from the projections are grouped into the conventional quinquennial groups 15-19, 20-24 ... 45-49).

From these P/F ratios, we can seek to understand how well the P/F ratio method predicts fertility rates (since we have assumed both the level and age distribution of fertility a priori). In particular, here we compare and contrast the estimates of fertility produced as a result of using three different scaling factors:

- the ratio of $P(2)/F(2)$ – i.e. the P/F ratio for the age group 20-24, referred to in the results presented below as Brass(2);
- the average of $P(2)/F(2)$ and $P(3)/F(3)$, referred to as Brass(2,3);
- the ratio of $P(\bar{m})/F(\bar{m})$ – as suggested by Feeney.

4. Results

This section presents the results from some of the simulations conducted.

Scenario 1 (stationary population)

The first simulation is inherently trivial, and simply demonstrates that the projection model works, producing a stable population, with stable fertility, and a P/F ratio of one at all age groups based on the following assumptions:

Table 1 Parameterisation of Scenario 1

Base population	Stable population based on Brass General standard
Fertility level	Constant
Fertility distribution	Constant
Mortality	Constant
HIV	No

Likewise, a slight modification on this scenario, keeping fertility constant but allowing mortality to improve, produced the expected results (constant P/F ratios, and a growing population). No results are presented for this scenario; it was derived simply to verify that the model maintains a stationary and stable state when required to do so.

Scenario 2 (constant mortality; fertility changes)

This scenario is the first that considers the impact of changing fertility patterns on the results from the P/F ratio method. In Scenario 2a, mortality is assumed to remain constant, no HIV epidemic is assumed (and hence no knock-on effect on fertility), while the level of fertility is assumed to change, but not the age distribution. This, together with Scenario 2b (holding fertility levels constant, but not the age distribution) are highly artificial insofar as a change in the level of fertility unaccompanied by a change in age distribution of fertility is highly improbable. Nonetheless, these scenarios are useful in highlighting the (independent) effects of changes in the level and distribution of fertility.

Table 2 Parameterisation of Scenario 2

Base population	Stable population based on Brass General standard
Fertility level	Scenario 2a: Constant Scenario 2b: Changing Scenario 2c: Changing
Fertility distribution	Scenario 2a: Changing Scenario 2b: Constant Scenario 2c: Changing
Mortality	Constant
HIV	No

The results of these simulations are shown in Figures 2-4. The top panel (panel I) shows the P/F ratios in each age group over time. The second panel (panel II) shows the (assumed to be known) level of fertility together with the fertility levels estimated by applying the three different approaches to estimating fertility from the P/F ratio. The third panel (panel III)

shows the magnitude of the error arising from the application of the three approaches relative to the (assumed to be known) level of fertility.

In respect of Scenario 2a (only the age distribution of fertility changes), the biggest effect on the P/F ratio is concentrated in the youngest age group (panel I of Figure 2). Panel II of the same figure shows the estimated level of fertility arising from application of the three variants of the P/F method as described above. (The very slight decline in the 'true' fertility level arises because the fertility age distribution to which the distribution tends has some residual fertility after age 49, which is not captured in the P/F method). The Feeney approach to estimating fertility has the lowest error of the three variants (panel III), which is not surprising, since it uses the P/F ratio at the mean age of childbearing (in the distributions to hand, around 27.5 years), to scale the current levels of fertility. For essentially similar reasons, the estimate of fertility derived from the average P/F ratio of women aged 20-29 represents a marginal improvement in accuracy over the estimate of fertility derived from the P/F ratio of women aged 20-24.

However, while the errors in the various methods are large (upwards of 15 percent at times), these errors are far less generalised than those that arise from assuming that fertility levels are constant when they are not. Scenario 2b is designed to illustrate the errors in fertility estimation that may arise if the age distribution of fertility remained constant, but fertility levels did not. Figure 3 shows the results from the simulation conducted under Scenario 2b. Several things are readily apparent from the output. First, the effect on the P/F ratios is reversed when fertility declines; the P/F ratios are highest among older women. Second, the effect on fertility estimation of declining fertility is a lot less severe than a change in age distribution, unless one uses the Feeney method, where the errors are of a similar magnitude – under-estimating (over-estimating) by around 20 per cent when fertility is rising (falling) fastest. Errors when using the P/F ratio of women aged 20-24 are consistently lower than when using the P/F ratio averaged across women aged 20-29.

Scenario 2c (Figure 4) combines the effects of falling fertility and changing the age distribution of fertility, and – predictably – the results presented in Figure 4 can be readily interpreted as a combination of the two preceding scenarios, and reflects the magnitude of errors in fertility estimation arising from violation of the assumptions of applying different variants of the P/F ratio method to the data. Again, the Feeney method is the least accurate variant, while using the P/F ratio for ages 20-24 (as suggested by Manual X) is usually preferable to using the P/F ratio averaged across women aged 20-29. For the most part of the fertility decline, however,

the errors are relatively small (of the order of 5 per cent or less). In the context of the generally poor data which pertain when the indirect methods have to be applied, errors of this magnitude are not a major cause for concern.

Figure 2 Output from scenario 2a – No mortality change, no fertility level change; change in fertility distribution






Figure 3 Output from scenario 2b – No mortality change, fertility level change; no change in fertility distribution




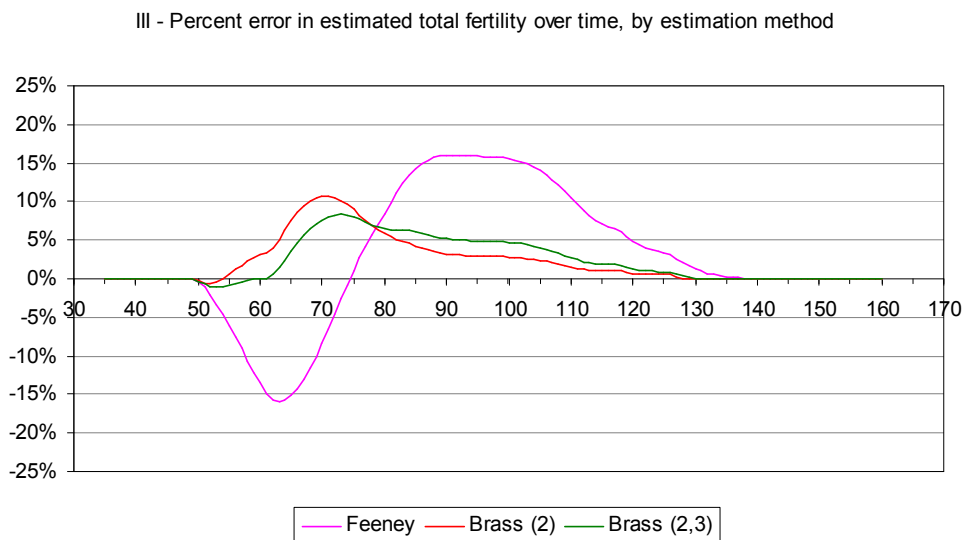

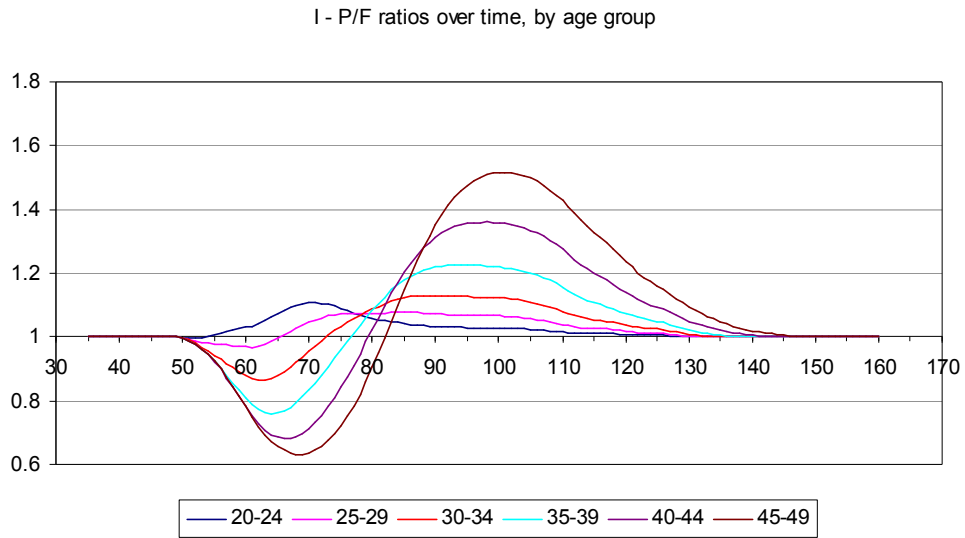


Figure 4 Output from scenario 2c – No mortality change, fertility level change; change in fertility distribution



Scenario 3 (mortality and fertility changes)

General reasoning, borne out by the results of the simulation, suggests that provided the fertility of those that die is the same as those that survive, changes in mortality will – *ceteris paribus* – have an obvious effect on the overall population size, but not on fertility estimates – which are standardised for the age distribution of the population. As can be seen from Figure 5, this is indeed the case. However, where mortality dynamics induce changes in fertility too, it stands to reason that there will be an impact on the results of the P/F ratio method. This is explored in Scenario 4.

Scenario 4 (mortality and fertility changes; HIV affects fertility)

One obvious mortality effect that might induce fertility changes of its own is, of course, the HIV epidemic. Women who are infected with HIV have been found to have lower fertility, as well as a different age distribution of fertility (Lewis, Ronsmans, Ezeh *et al.* 2004; Zaba and Gregson 1998). The effect of HIV/AIDS – on the assumptions outlined above – on the results produced by the P/F ratio method is investigated in Scenario 4; the results are shown in Figure 6. In general, the presence of an HIV epidemic slightly attenuates the overall impact of falling fertility and changing age distributions; this is largely due to the fact that the assumed age distribution of HIV fertility works to slow the increase in the mean age at childbearing. A clearer indication of this effect at work can be gained from assessing the impact of an assumed HIV epidemic on the P/F ratio in isolation from other assumed changes in fertility. This is considered in Scenario 5.

Output from this scenario allows the determination of the cross-over point (in age and time) at which the cumulated fertility of uninfected women exceeds that of infected women. This has bearing on the argument by general reasoning presented above regarding the likely direction of bias arising from differential mortality of HIV positive women and its impact on fertility estimation.


While the mean age at childbearing is much younger among HIV-infected women (as expected from the much younger age distribution of fertility of these women), the results from the simulations suggest that, beyond age 38, the cumulated fertility of uninfected women is always greater than that of infected women, while the reverse holds for all ages under 29. Between these two ages, cumulated fertility by HIV status is essentially similar and which prevails is conditioned by the overall assumed level of HIV prevalence (since this affects the aggregate PAC applied to determine the total fertility of HIV-infected women).

The implication of this finding for the general reasoning presented in the previous section is that – at a census – among women up to the age of 30, fertility rates estimated directly from the data will tend not to be overestimated since the mean age at death of infected women in a generalised epidemic is older than 30 (these women are still likely to be alive). Likewise, the reported average parities of women (used in the numerator of the P/F ratio) will not be significantly affected.


This reasoning is entirely consistent with the findings presented under this scenario, which suggests that the impact of HIV on the method is trivially small. This is shown more directly by the next scenario.

Figure 5 Output from scenario 3 –mortality, fertility level and fertility distribution change

I - P/F ratios over time, by age group



II - Actual vs estimated total fertility over time, by estimation method



III - Percent error in estimated total fertility over time, by estimation method

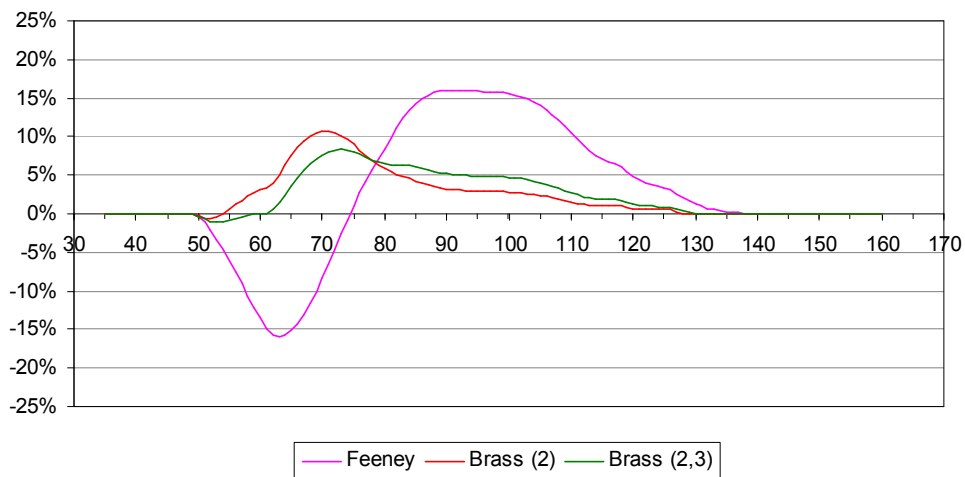

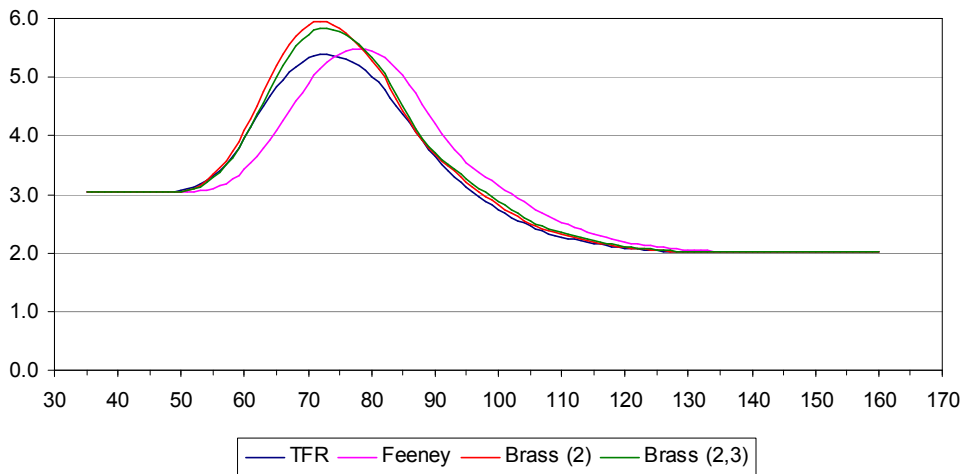


Figure 6 Output from scenario 4 –mortality, fertility level and fertility distribution change; HIV affects mortality and fertility


I - P/F ratios over time, by age group



II - Actual vs estimated total fertility over time, by estimation method



III - Percent error in estimated total fertility over time, by estimation method



Scenario 5 (mortality changes; HIV affects fertility; no other fertility changes)

The scenario (the results of which are shown in Figure 7) demonstrates that – while other violations of the assumptions have tended to lead to the P/F ratio overestimating fertility, the presence of an HIV/AIDS epidemic has a (small) counterbalancing effect on the P/F results. The drop in fertility shown in Panel II of Figure 7 arises from the assumed sub-fecundability of HIV-infected women, while the errors associated with each of the three variants of the estimation procedure clearly reflect the average age at which the scaling factor is determined. Further, while – again – the Feeney approach works least well, all errors are in the same direction, and – relative to the errors described in the preceding scenarios – are very small indeed. Nonetheless, these results would seem to indicate that – if anything – HIV has a (small) mediating effect on the errors produced by the P/F ratio method in a time of declining fertility.

One final investigation is conducted here; into the variation in the errors associated with HIV relative to the prevalence among women of childbearing age. This is best done by examining the plot (shown below) of the errors presented in Panel III of Figure 7 against the HIV prevalence at each time point in the evolution of the epidemic. From the results presented in Figure 8, it would appear that the once the epidemic has reached some kind of stability, the attenuating effect is, as might be expected, almost zero; the effect on the two Brass variants is greatest when prevalence is rising fastest.

Figure 7 Output from scenario 5 –fertility level and distribution do not change; HIV affects mortality and fertility


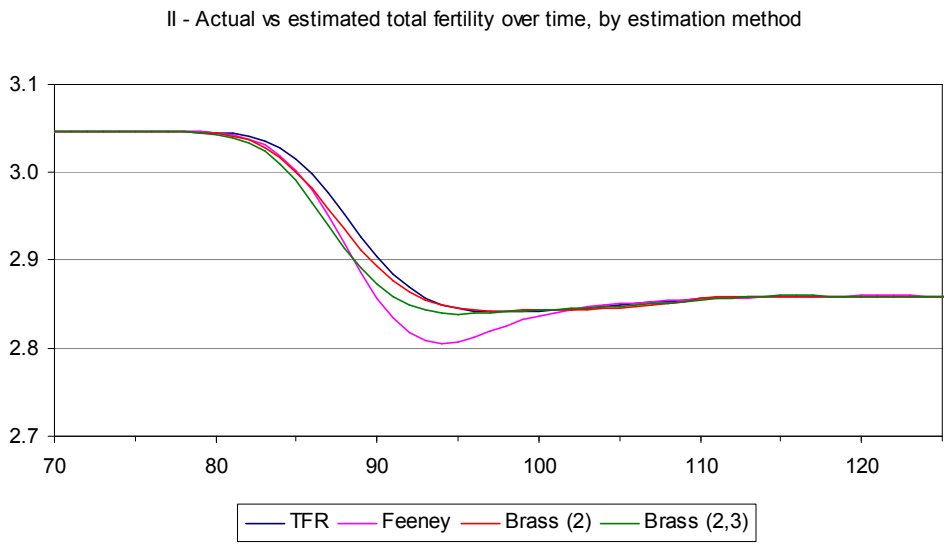

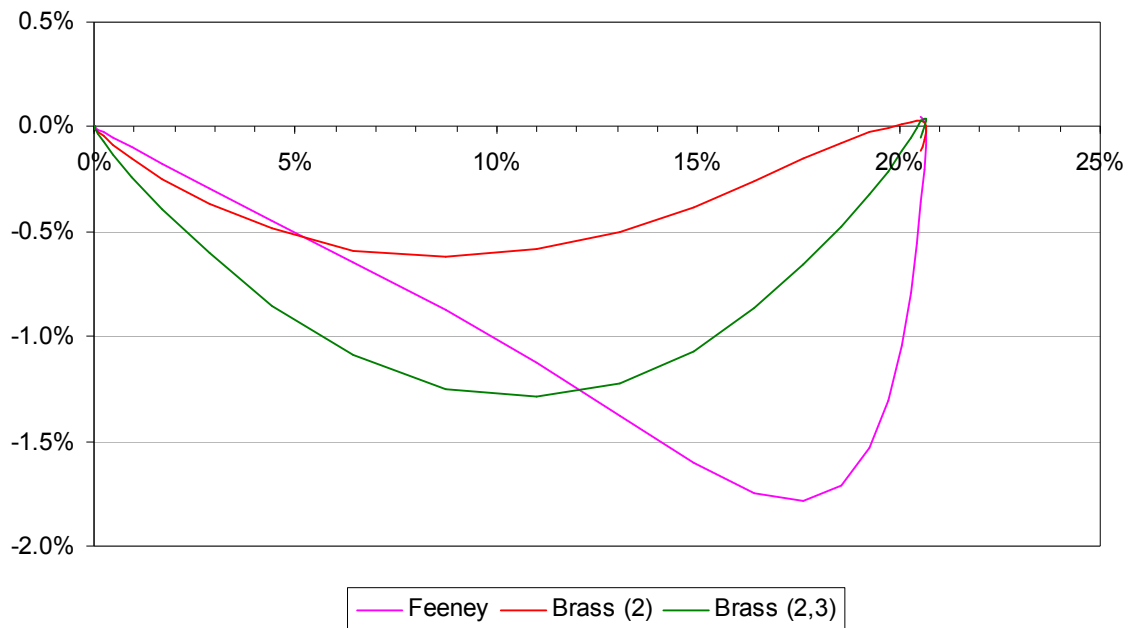


Figure 8 Percent error in different variants of the P/F method relative to HIV prevalence among women 15-49, Scenario 5



5. Conclusions

Even in sub-Saharan Africa, both demographic and health surveys, and longitudinal studies, are imperfect substitutes for the demographic data collected in censuses. While the results from such surveys and studies contribute hugely to our understanding of demographic dynamics, their small sample sizes and spatial concentration often mean that they cannot provide estimates of fertility (and even more so, mortality) that can be relied upon by policy makers. And, for as long as the quality of data collected in national censuses in the developing world (and sub-Saharan Africa in particular) is of poor quality, demographers will have to continue to rely on estimates of fertility and mortality based on the ‘indirect techniques’.

This paper has sought to understand the nature, magnitude and direction of biases arising from systematic violation of the assumptions underlying the Brass P/F ratio method. Brass speculated that changing age-distributions of fertility were a ‘serious problem’ with the method and the investigations conducted here suggest that these concerns are well-founded. However, the much greater contributors to errors in the Brass P/F ratio method are those arising from changes in the level of fertility.

The investigations suggest that the Manual X recommendation of using the P/F ratio at 20-24 to scale the reported fertility schedule is more accurate than the Feeney method, although it would overstate fertility while fertility is rising and for some time after period fertility peaks, reaching a maximum of around 10 percent at the peak of period fertility. There

is little to choose between the two conventionally recommended scaling factors (using the P/F ratio at age 20-24, or the average ratio between 20 and 29), although the former appears to be preferable – except when fertility is increasing. In general though, the errors induced by either of the conventional scaling factors are small: only after a sustained rise in fertility are the errors in estimated fertility greater than ten percent. For the most part, they are under five percent. Bearing in mind the almost certainly poor quality of the data that would necessitate the use of this method in the first place, these errors are acceptable.

By contrast, the Feeney approach appears to get the peak of the schedule about right, but the entire curve seems to be shifted to the right by about six years. Quite why this is so (or, indeed, whether or not this is even a robust finding) is something requiring further examination. Intuitively, we suspect that Ryder's equivalence (on which Feeney's method hinges) does not hold when fertility is falling.

The findings presented here suggest that differential fertility between HIV-infected and HIV-uninfected women has a trivial impact on the methods, even in an environment with a simulated highly generalised epidemic, and that effect – such as it may be – serves to attenuate the extent of over-estimation of fertility identified above.

This research is relevant for several reasons. In the first instance, an appreciation of the behaviour of different variants of the P/F ratio at different stages of the demographic transition provides for greater understanding of the level and trajectories of the fertility decline in the developing world where estimates have been derived using these methods. Second, the research helps shed light on which methods are more robust under conditions of simultaneously declining fertility and rising mortality associated with differential fertility (as is the case in most of sub-Saharan Africa, and other settings also afflicted by widespread HIV). Given that the P/F method in its various forms is still widely used in the estimation of fertility rates in the region, this will result in more certainty regarding the trajectory of fertility rates and patterns in sub-Saharan Africa.

References

- Actuarial Society of South Africa. 2005. *Aids Demographic Model 2003*. Cape Town: Actuarial Society of South Africa. www.assa.org.za.
- Brass, William. 1964. "Uses of census or survey data for the estimation of vital rates," Paper presented at African Seminar on Vital Statistics. Addis Ababa, 14-19 December 1964.
- Brass, William and Pauline Airey. 1988. *The P/F Ratio*. Unpublished lecture notes. London: Centre for Population Studies, London School of Hygiene and Tropical Medicine.

- Brass, William, Ansley J. Coale, Paul Demeny, Don F. Heisel *et al.* (eds). 1968. *The Demography of Tropical Africa*. Princeton NJ: Princeton University Press.
- Coale, Ansley. 1972. *The Growth and Structure of Human Populations*. Princeton, NJ: Princeton University Press.
- Feeney, Griffith. 1998. *A New Interpretation of Brass' P/F Ratio Method Applicable when Fertility is Declining*. <http://www.gfeeney.com/notes/pfnote/pfnote.htm>. Accessed: 11 January 2000.
- Lewis, James JC, Carine Ronsmans, Alex Ezech and Simon Gregson. 2004. "The population impact of HIV on fertility in sub-Saharan Africa", *AIDS* **18**(S2):S35-S43.
- Moultrie, Tom A. and Ian M. Timæus. 2002. *Trends in South African fertility between 1970 and 1998: An analysis of the 1996 Census and the 1998 Demographic and Health Survey*. Cape Town: Medical Research Council. <http://www.mrc.ac.za/bod/trends.pdf>.
- Moultrie, Tom A. and Ian M. Timæus. 2003. "The South African fertility decline: Evidence from two censuses and a Demographic and Health Survey", *Population Studies* **57**(3):265-283.
- Ryder, Norman B. 1964. "The process of demographic translation", *Demography* **1**(1):74-82.
- Ryder, Norman B. 1983. "Cohort and period measures of changing fertility," in Bulatao, Rodolfo A. and Ronald D. Lee (eds). *Determinants of Fertility in Developing Countries*. Vol. 2. New York: Academic Press, pp. 737-756.
- United Nations. 1983. *Manual X: Indirect Techniques for Demographic Estimation*. New York: United Nations.
- Zaba, Basia and Simon Gregson. 1998. "Measuring the impact of HIV on fertility in Africa", *AIDS* **12**(suppl 1):S41-S50.