Estimating Total Fertility Rates for Small Areas in Brazil

Suzana Cavenaghi[1], Joseph E. Potter[2], Carl P. Schmertmann[3], Renato M. Assunção[4]

[1] NEPO, State University of Campinas
[2] Population Research Center, University of Texas at Austin
[3] Center for the Study of Population, Florida State University
[4] Statistics Department, Federal University of Minas Gerais

Presented at the Annual Meeting of the Population Association of America,
Boston, Massachusetts, April 1 – 3, 2004,
Session # 63, Methodological Issues in Fertility Measurement.

Correspondence:
Joseph E. Potter
Population Research Center
University of Texas at Austin
1800 Main Building
Austin, Texas 78712 USA
(512) 471-8341
joe@prc.utexas.edu

Sources of support: NICHD R01HD041528-01A1

INTRODUCTION

In the 1960s and 1970s, indirect estimation of demographic parameters in populations where vital registration was incomplete was one of the crown jewels of population science. The techniques developed for this purpose were based on demographic theory, and often involved the ingenious combination of alternative data sources and assumptions about the distributions of the phenomena as well as the nature of any errors that might be present in the data. This work was pioneered by the authors of the Demography of Tropical Africa (Brass et al. 1968), UN Manual IV (Coale and Demeny 1967), and further developed by a second generation of demographers most of whom have now celebrated their 55th birthdays. Indirect techniques for demographic estimation was one of the first tasks addressed NAS/NRC Committee on Population and Demography when it was established in 1977, ultimately leading to the publication of UN Manual X (United Nations 1983). It was the main staple of the demography taught and further developed at the UN Regional Demographic Centers, and it was the subject of scores of journal articles, and an even larger number of more applied reports.

What might be called the golden age of indirect estimation came to an end with the widespread implementation of large-scale, nationally representative surveys in countries around the world. The World Fertility Survey and the various programs that followed it provided the impetus and financing for much of this data collection, but such internationally sponsored surveys were often complemented by similar surveys conducted by national statistical agencies, population councils, or ministries of health. The survey data proved to be reasonably reliable not only for estimating fertility, but also infant and child mortality. What indirect estimation remained was largely confined, paradoxically, to the analysis of questions included in such surveys for the purpose of estimating adult mortality.

In this paper, we estimate the level and pattern of fertility in Brazilian municipalities. The larger question raised by this analysis is whether there may now be an opportunity to take recourse to indirect estimation once again. Interest in developing reliable estimates of both fertility and mortality for small local areas in many countries seems to be growing rapidly, irrespective of the adequacy of vital registration. The reasons for this increased interest, we suspect, derives from the recent spate of efforts to either decentralize or target social policies, and the importance of demographic parameters and present and projected population sizes to such efforts. This tendency has been furthered by the interest of international organizations, perhaps especially UNDP, in developing local level indices of development such as the Human Development Index. Surveys, even the very largest, are rarely of much use for local level estimation, and if registration of birth and deaths is still incomplete, the analyst may encounter a situation very similar to the one that prevailed in the 1960s and early 1970s. The only alternative is to take advantage of what information is available from both vital registration and censuses, knowing full well that this information may be defective or incomplete.

Local level estimation in such contexts, however, poses two additional problems that were not present in the earlier applications to large populations. One derives from the sheer volume of estimates that are called for. Most methods proposed to correct data and to estimate fertility or mortality indirectly depend on the researcher's judgement and interpretation for the selection of the best estimates, and making literally thousands of such judgements would be a cumbersome task. The second difficulty is that even with census data and vital registration, the population exposed to risk may be too small to permit precise estimation.

For the purpose of estimating age-specific fertility rates and the TFR for large numbers of small areas, we have developed an empirical bayes smoothing procedure to cope with the problem of sampling variability (Assuncao et al. 2003). Here we implement that procedure for Brazilian municipalities (5,506 in number), and also extend the analysis to include an adjustment for reporting bias based on parity. We use data from the 2000 Brazilian census, and eventually compare our estimates with those now used by the UNDP as well as estimates that we derive from vital registration. While there is still much that remains to be done before we can reach a final conclusion regarding the utility of our techniques, our current reading is that there is, indeed, a promising opportunity to develop local level estimates using both spatial smoothing and indirect estimation.

DATA

Administrative records on births are available from two sources in Brazil. One is the vital registration from the official agency of the government, organized by the Bureau of the Census (IBGE), known as "Civil Records", which comes essentially from the office of notary publics. The second source, called SINASC (Information System on Live Births), available since 1994, is run by a department of the Health Ministry that collects data from the registration of births in health establishments, and completes it with data from the notary public's office for live births that occurred at home or in other settings. While there has been a notable improvement in the coverage of birth registration in the last decade, both sources still have high rates of under-registration in the less developed regions of the country.

Vital registration from the Bureau of Census covered about 70 percent of Brazilian births as of 1986. Coverage varied greatly by region: according to Rodriguez Wong (1986) the lowest coverage was around 50 percent in the North, while other regions, especially the Southeast, had almost complete registration of births. These assertions were questioned when SINASC started collecting data and improving the system by the end of the 1990's. According to SINASC data, in 2001, the North and Northeast regions had a coverage of around 80% of the live births, the Southeast and South around 90%, and the Center-West about 93% (MS/SVS, 2004). These figures are very uncertain, however, because the exact number of live births in each region is unknown.

The greatest advantage of the SINASC data on live births is that it is readily available and includes information on mother's age at the delivery, place of residence, place of occurrence of birth, sex of the baby, type of delivery, and gestational duration. Up to 2000, the SINASC system had collected more births than the Civil Records. For these reasons, to estimate fertility rates from direct measurements, we chose to utilize SINASC data for 2000.

Since vital registration is still deficient for nearly all regions of Brazil, and none of the surveys that include birth histories are representative at the state level, much less for municipalities, published estimates of Brazilian fertility rates at the state and local level available are based on indirect measures, utilizing data from Census and surveys. In accordance with international recommendations, the Brazilian Bureau of the Census included questions about the number of children ever born since 1940 among women aged

15 years and older[1] in the long form of the census. The long form was administered to a 25 percent sample of the population. In 1970 and in later years, the census long form also included questions on current fertility[2]. Current fertility questions have changed slightly from census to census.

Until 1970, the number of children ever born was phrased the number of children ever born that the woman had prior to the census date. In the next two censuses, the question was simplified, no longer referring to the census date. It was simply children ever born and a box was available to write down the number. However, different from previous years, the question was broken down by sex of the children. This procedure was intended to improve the quality of the data since it was believed that the respondent would be less likely to forget children born alive who had died or who had left home if the sex of the children had to be reported.

In 1970 children born in the preceding year was asked: of the children born alive, how many were born in the 12 months preceding the census date, and the period was specified as September 1, 1969 to August 31, 1970, the year preceding the census date. In 1980, a significant change was introduced because instead of using the one-year reference period, the questionnaire asked for the date of birth (month and year) of the last child born alive. The question was phrased as "month and year of birth of the last child (son or daughter) born alive". The information available to the researchers sometimes differs from that collected in the questionnaire. This was the case with the 1991 and 2000 data releases.

---

[1] Within the household selected in the sample, the information on fertility was collected for all women aged 15 and over that resided in the household.

[2] In addition to questions on children born alive by the date of the interview, questions on stillborn were asked since the 1940 census.

Although the date of birth was collected in the interview, in the census microdata files this information was provided in the form of the age of the last child.

Both questions on parity and recent fertility are affected by the problem of missing information. In the case where the missing data are randomly distributed among women of all parities, it will not bias fertility estimates, and all women with non-stated fertility and their children have to be excluded from the estimation of fertility rates. However, there is a well-known problem of interviewers leaving the questions for childless women blank instead of recording a zero value (El-Badry 1961). If this is the case, a more reliable proportion of women with non-stated fertility should be re-estimated, because some of the women with non-stated fertility are actually childless. Following the UN new recommendations for data collection and distribution, the Bureau of the Census decided that for the 2000 census all unknown information had to be imputed to the data before making the data set available to public use. Thus for the data presented in this paper, the potential problem of non-response for fertility questions is moot.

The long form of the 2000 demographic census was applied to a sampling fraction of either 10% or 20% (for municipalities with an estimated population larger and smaller than 15,000 inhabitants, respectively). As may be seen in Table 1, The smallest municipal sample has only 22 women 15-49 years old; the largest has 363,732. Mean sample size is 991; the 25th, 50th, and 75th percentiles are 145, 300, and 628, respectively. The distribution of sample sizes is very skewed: the 90th and 95th percentiles of sample size are 1,453 and 2,709, respectively.

The 2000 census microdata are weighted at the household level to account both for non-response and the two different sampling fractions. In order to have independent and identically distributed cases in the sample, and still have population representation in the sample, we utilized "analytical" weights that maintain the original sample size. Hence, all the figures shown above refer to the weighted (but not expanded) sample data.

The geographical boundary files for municipalities in 2000 are a third source of data for this study. There were 5,506 municipalities in Brazil in 2000 (excluding one island in the Northeast region). These units vary dramatically not only in population size (from 769 to 10,400,000 inhabitants), but also in area (from 1 to 62506 km2). (See Table 1 and map below).

Finally, we also utilize total fertility rate estimates from the United Nations Development Program. These were generated in connection with the calculation of the Human Development Index for 2000 at the municipal level (PNUD, 2003). According to the manual accompanying these estimates, the TFRs were obtained by adapting the indirect P/F technique proposed by Brass. However, just how the analysts dealt with the problem of sampling variability is not revealed in the documentation we have obtained to date.

METHODS

Shrinkage, Vector Shrinkage, and Smoothing

Shrinkage Estimation for Sets of Scalar Parameters

Even with comprehensive long-form data from the Census, municipal-level samples are frequently too small for demographic analysis. Half of the municipal samples in the 2000 Census contain fewer than 300 women, and 13% contain fewer than 100. Standard estimates of age-specific and total fertility for many municipalities will consequently be very noisy.

Empirical Bayes (EB) "shrinkage" estimation is a valuable statistical tool in such situations, when the analyst must estimate *N* different parameters from *N* samples, but each sample comes from some fundamentally similar process. The main ideas behind shrinkage are to compare each individual estimate to the overall set of estimates, and to recognize that unusual estimates are more likely to contain unusual sampling errors. Even though individual parameters are seemingly unrelated, one can improve estimation by considering the whole set. For example, above-average estimates are more likely to contain positive sampling errors, and vice versa. After 20 at bats, a baseball player hitting .650 is probably good *and* lucky. A player hitting .050 is probably bad *and* unlucky. (The baseball example was first proposed in Efron and Morris 1975.) Shrinkage procedures utilize this information by pulling each estimate back, at least slightly, toward the overall mean.

As an instructive example, suppose that a set of scalar parameters $\theta_1 \ldots \theta_N$ have a normal distribution with mean μ

$$\theta_i \sim iid\ N(\mu, \sigma^2)$$

For example θ_i might represent true values of some demographic index in each of *N* local areas. Suppose further that the parameters are estimated imprecisely from a census or survey, and that the estimates include sampling errors *u* that are also normally distributed

$$\hat{\theta}_i = \theta_i + u_i \qquad u_i \sim indep\ N(0, \omega_i^2)$$

One possible set of estimates is $\{\hat{\theta}_i\}_{i=1 \ldots N}$. However, among all cases with a particular *estimated* value $\hat{\theta}_i$, the average *true* parameter value will be closer to the grand mean μ:

$$E(\theta_i \mid \hat{\theta}_i) = \hat{\theta}_i + \left(\frac{\omega_i^2}{\sigma^2 + \omega_i^2} \right) \left(\mu - \hat{\theta}_i \right)$$

This expected value is greater than $\hat{\theta}_i$ if the estimate is below the mean (the true long run batting average .050 hitters is probably above .050), and less than $\hat{\theta}_i$ if the estimate is above the mean (the true average of .650 hitters is probably below .650).

As a consequence of this fact (if μ,σ^2, and ω_i^2s were known) a set of adjusted estimates

$$\left\{ \tilde{\theta}_i = \hat{\theta}_i + \left(\frac{\omega_i^2}{\sigma^2 + \omega_i^2} \right) \left(\mu - \hat{\theta}_i \right) \right\}_{i=1 \ldots N}$$

would have a lower mean squared error than $\{\hat{\theta}_i\}_{i=1 \ldots N}$, in the sense that $E(\tilde{\theta}_i - \theta_i)^2 \leq E(\hat{\theta}_i - \theta_i)^2$. Note that the expectation is across all possible true values of the parameter. This means that the *total* of squared errors across the *N* estimates is expected to be lower if we replace $\{\hat{\theta}_i\}$ with $\{\tilde{\theta}_i\}$; it does not guarantee a lower mean squared error for adjusted estimates of any one parameter. In practice, however, replacing $\{\hat{\theta}_i\}$ with $\{\tilde{\theta}_i\}$ would lower the expected squared error for almost all of the individual parameters in the set (see Longford 1999:233); the exceptions would be cases with the most extreme parameter values (e.g. a true .500 hitter in the baseball example).

The adjusted estimators "shrink" individual estimates toward the overall mean μ. Shrinkage is greater when the estimate is unusual ($\hat{\theta}_i$ far from μ), and when sampling errors for the parameter are large relative to the variability of true parameters (ω_i^2 large relative to σ^2). Normality is not essential for establishing that shrinkage estimation improves on separate estimation for each parameter. In fact, for any distribution of parameters and sampling errors with variances σ^2 and $\omega_1^2 \ldots, \omega_N^2$, respectively, shrinkage estimators of the form above have the minimum mean squared error among all weighted sums of $\hat{\theta}_i$ and μ (cf. Longford 1999).

In practice, of course, distributional parameters μ, σ^2, and ω^2 are unknown and must be estimated from the data. In most demographic applications the sample mean of $\hat{\theta}_i$ is a natural estimate for μ, ω_i^2 can usually be inferred from sample sizes and distributional assumptions, and σ^2 can be estimated from the relationships $N\sigma^2 = \sum_i \left[V(\hat{\theta}_i) - \omega_i^2 \right]$. These preliminary estimation steps are the *empirical* part of EB shrinkage.

Shrinkage Estimation for Sets of Vector Parameters

If the *N* parameters of interest $\theta_1 \ldots \theta_N$ are themselves vectors with *K* components (such as 7x1 schedules of age-specific fertility rates or average parities by age group), results from the previous section can be extended via matrix and vector notation. Longford (1999) demonstrated that when the *N* parameter vectors are uncorrelated with *K*x1 mean **μ**, the *K*x*K* matrix of cross-component covariances is Σ for every vector, and sampling errors are uncorrelated across both parameters and components, then the set of optimal shrinkage estimators for *K*x1 vectors $\theta_1 \ldots \theta_N$ is

$$\left\{ \tilde{\theta}_i = \hat{\theta}_i + \Omega_i \left[\Sigma + \Omega_i \right]^{-1} \left(\mu - \hat{\theta}_i \right) \right\}_{i=1 \ldots N}$$

where Ω_i is a *K*x*K* diagonal matrix of sample error variances for the components of (vector) estimate $\hat{\theta}_i$. Notice that if sample sizes are very large for estimating components of vector θ_i, then matrices Ω_i and $\Omega_i\ [\Sigma + \Omega_i]^{-1}$ will be very small, leading to very little shrinkage of the original vector estimate $\hat{\theta}_i$. Conversely, if sample sizes for estimating components of θ_i are small, $\hat{\theta}_i$ may be adjusted substantially toward a more typical vector pattern.

This is the multivariate analog of the scalar shrinkage formula. “Shrinkage” has a more complicated meaning in the vector case, because the matrix formula may move individual component estimates *away* from their global (univariate) means even as it moves a vector of estimates towards a more typical pattern. (It is useful to think of both scalar and vector shrinkage as moving individual estimates toward a fitted regression line or plane. In the scalar case the regression function is simply the global mean and shrinkage implies moving the estimate toward that mean. However, in the multivariate case, shrinking an estimate in R^K toward a regression function could mean moving some components away from their univariate means while other components move closer to theirs.) We refer readers to Longford (1999) and Assunção et al. (2003) for more details on estimating Σ and Ω_i matrices in the Empirical Bayes application of the vector shrinkage formula.

Vector Shrinkage as Curve Smoothing

The vector EB procedure shrinks individual vector estimates towards the typical vector pattern in a sample of *N* parameter vectors. The “typical vector pattern” includes information about absolute levels (e.g., the average value of component 2 is 0.15), about relative levels (e.g., component 2 is usually much larger than component 6), and about cross-component correlations (e.g., when component 2 is higher than average, component 3 is usually higher than average).

Considered as an abstract process in R^K, EB vector shrinkage is not particularly intuitive. However, if we array components along the horizontal axis of a graph and view the vectors as functions, it becomes clear that the EB procedure takes unusually shaped schedule estimates $\hat{\theta}_i$ and smooths them into more regular $\tilde{\theta}_i$ ’s.

The next figure shows examples from two municipalities, Piraju in the state of São Paulo and Barra do Jacaré in Paraná. The 2000 long-form census sample for Piraju was above median size (874 women 15-49 years old), while the sample for Barra do Jacaré was quite small (82 women). The top panels in the figure display parity data, and the bottom panels display data on fertility in the year preceding the 2000 census. All four panels show both estimates from local data for the municipality ($\hat{\theta}$ vectors), and EB vector shrinkage estimates toward typical patterns in neighborhoods around the municipality (details on neighborhood definitions are in the next subsection).

In all panels the EB procedure smooths the raw data by moving the vector of local estimates toward the typical pattern in the region around the municipality. For average parity, the EB curve rises more regularly with age than the raw census data for Piraju, particularly because it eliminates a decrease in average parity between ages 35-39 and 40-44 in the local census estimates. EB does considerably more smoothing for parity data in Barra do Jacaré: smaller sample sizes mean that local estimates are more variable, so there is more shrinkage toward the regional average pattern.

The local ASFR schedule for Piraju is somewhat noisy, with estimated ${}_5f_x$ values of zero for x=40 and 45, because no women in those age groups reported a birth in the year preceding the census. In Barra do Jacaré the very small local sample produces even more unusual ${}_5f_x$ estimates, with zero values in five of seven age groups. The EB procedure transforms the local ASFR estimates for Piraju modestly, but plausibly. Based on comparison to ASFR schedules in nearby municipalities, EB smooths Piraju's jagged-looking schedule into one with expected features (single mode in the 20s, regular decline after peak age, etc.). EB alters the ASFR schedule for Barra do Jacaré much more aggressively.

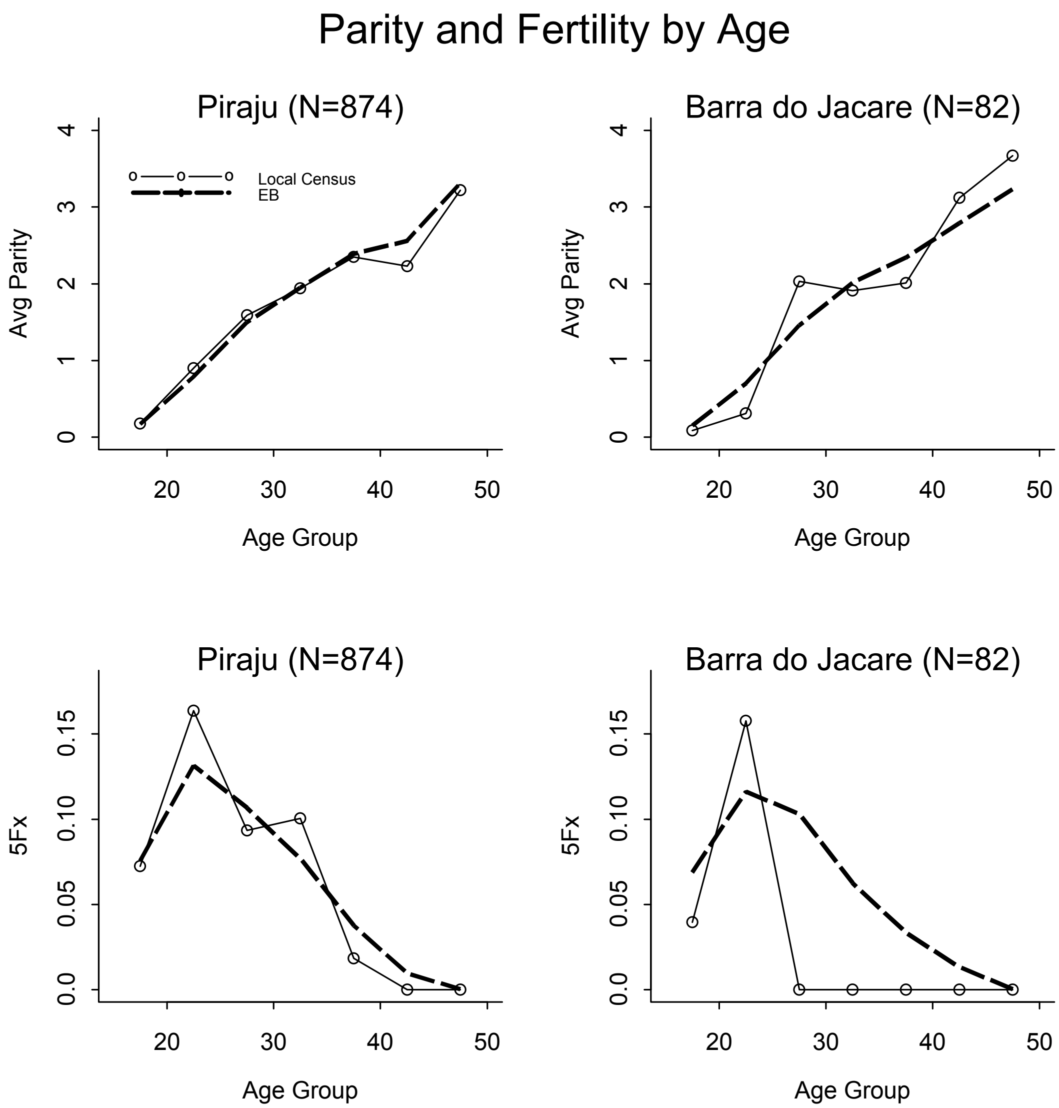

The EB schedules in the figure combine local data from Piraju and Barra do Jacaré with estimated schedules from nearby (presumably similar) municipalities. The EB procedure alters local schedules in Piraju moderately, pushing them toward patterns and levels typical of the surrounding region. It alters estimates for Barra do Jacaré much more, because sample sizes are small and local patterns are very unusual when compared to other municipalities in the region. The results are schedules that respect local information (e.g., fertility is near replacement level in both places, and completed parities suggest higher fertility in the past) while also conforming to regional patterns (e.g., parity tends to rise steadily with age, fertility is not zero after age 25, and so forth).

It is important to emphasize that the smoothing properties of the EB method come from patterns in data, *not* mathematical assumptions in a model. This is an important strength of the method.

Neighborhood Vector Shrinkage

Fertility has been falling rapidly in Brazil, but the decline has not occurred uniformly over space and time (Potter, Schmertmann, and Cavenaghi 2002). Fertility regimes remain diverse, with the extremes in rural frontier areas in the North and developed industrial areas in the South and Southeast. Consequently, vector shrinkage toward national means (as in Longford 1999) seems undesirable. As suggested by the text in the previous subsection, we have chosen to use smaller neighborhoods around each municipality, and to shrink municipal-level schedules toward neighborhood means, using neighborhood-specific covariance patterns. This moving neighborhood approach is similar to that used for univariate shrinkage by Marshall (1991).

We constructed a unique neighborhood for each of the 5,506 municipalities in the 2000 census. Our goal was to produce homogeneous neighborhoods that also had sufficiently large samples to produce good 'targets' for shrinkage. To this end, for each municipality M we

- Calculated the distance from *M*'s centroid to the centroid of every other municipality in Brazil
- Ordered municipalities 1…5506 according to increasing distance from *M*
- Selected *M* + the 7 closest municipalities as the neighborhood for *M*, *unless* the total sample size in this neighborhood of 8 was below 21,000 women age 15-49
- If the initial neighborhood had fewer than 21,000 women, we extended the neighborhood to the 8th-closest municipality, 9th-closest, etc., stopping once the extended neigborhood had at least 21,000 women.

The figure below illustrates the procedure for Piraju and Barra do Jacaré, the municipalities in the previous subsection. The figure maps a 2.5 x 2.5 degree grid of longitude and latitude in Southeastern Brazil. In more familiar terms, this grid is approximately 250 km or 150 miles on each side. Municipal centroids that fall in the grid are marked with numbers indicating their census sample sizes – for example, the southwestern-most municipality had a sample size of 264 women 15-49, while the southeastern-most had 123.

Piraju is marked with a red square, and with its sample size of 874. An initial neighborhood of Piraju+7 (dashed red lines) contained an insufficient number of women, so we expanded the neighborhood outward, municipality by municipality, until the total sample size exceeded 21,000. This produced the neighborhood inside of the solid red lines centered on Piraju. EB vector estimates for Piraju (as shown above) result from shrinking Piraju's schedules toward the mean schedule *for the Piraju neighborhood*, using local covariance information on components.

Barra do Jacaré (marked with a blue triangle) lies just outside the western edge of the Piraju neighborhood. An analogous procedure produces a different neighborhood (solid blue lines) that partially overlaps Piraju's neighborhood, but does not include Piraju itself.

EB estimates for Barra do Jacaré shrink toward the vector means in *this* neighborhood, and we repeat this procedure 5506 times.

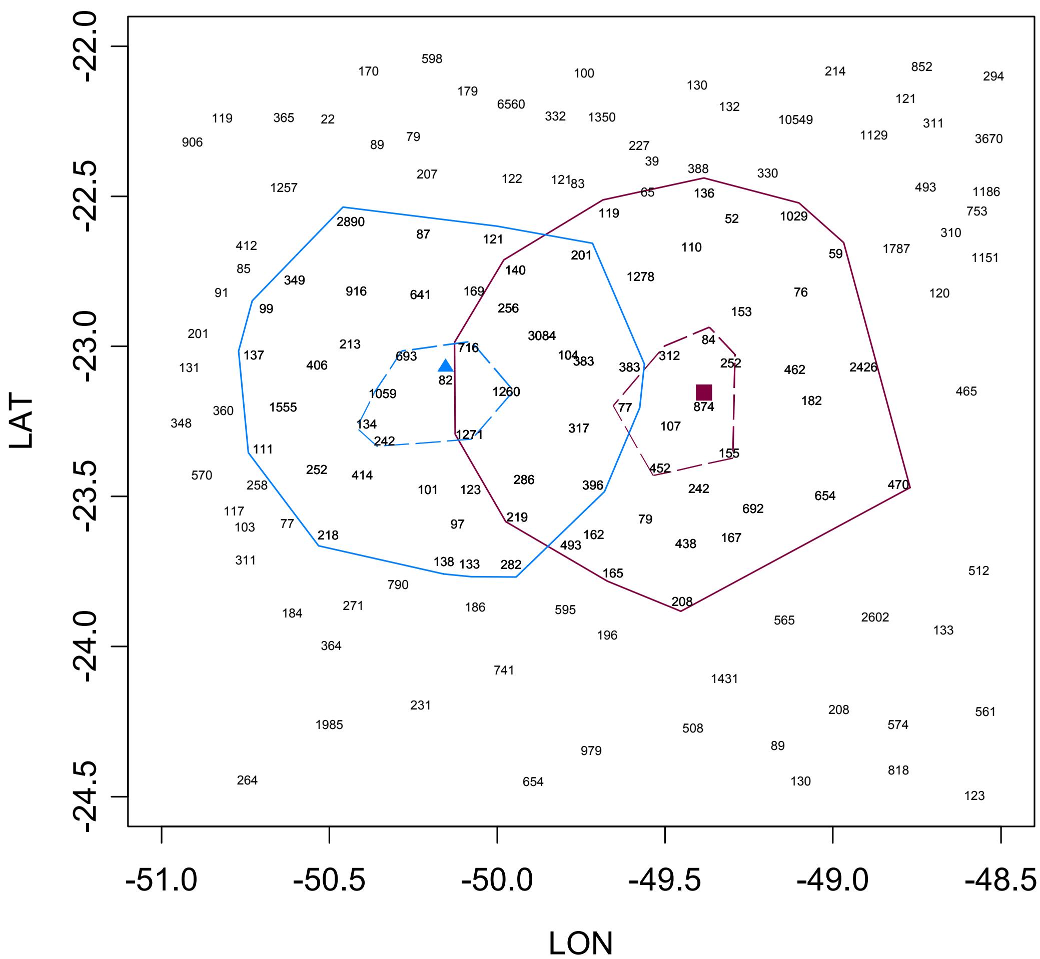

Adjusting the Census Data on Current Fertility for Reporting Errors

While the procedures outlined in the previous section provide a means of coping with the problem of sampling variability, they do nothing to adjust for any reporting errors that may be present in the responses to the question regarding date of last live birth that may be present in this information. The analyses of earlier censuses found that these were substantial even if they seemed to improve from one census to the next, and varied considerably across regions and states.[3] The classic method for adjusting for such errors, which may include both omissions and misreporting of the date of birth, is the indirect procedure first proposed by Brass that involved cumulating period rates to the point where they could be compared with the average parity of women in the respective five-year cohorts. This classic P/F methodology rested on three main assumptions: 1) reporting errors would be more or less constant across age groups; 2) census questions on the number of children ever born would be answered reliably; 3) and that there had not been enough change in either the age pattern or level of fertility so as to invalidate the comparison of period and cohort fertility. While this methodology is still widely used in Brazil and many other settings, the third assumption is, of course, completely untenable. Moreover, when applied to state or municipalities there is the additional complication posed by high rates of migration that is almost certainly selective with respect to fertility.

[3] Studies utilizing Brazilian data have usually concluded that census information on children born in the preceding year is underestimated or at least unreliable. Among them is the study by Merrick and Berquó (1983) that examined unadjusted rates with those derived after using the P/F adjustment, controlling for education and income. They concluded that the unadjusted rates were "questionable if not misleading" (1983, p. 78). Rodriguez Wong and Oliveira (1984), inspecting the progression of P/F factors, concluded that underreporting of current fertility varied by states, but that, on average, it was 12 percent for the entire country (p. 2275-78). Sawyer and Correa (1988), after comparing direct estimates of fertility rates for the state of São Paulo from vital statistics with other estimates derived from indirect methods and census data, concluded that "children born in the last year" underestimates the level of fertility in 1970 and 1980 census data3. There are no similar analyses for 1991 and 2000 census data.

In our analysis, we make use of the P/F method employing the various coefficients described in Manual X (United Nations 1983). However, we introduce a crude adjustment for the fact that Brazil has experienced rapid fertility decline over the last three or four decades that is based on birth history data collected in the last DHS survey to be conducted in Brazil. The adjustment involved calculating the P/F ratios found in the DHS for the second, third, and fourth age groups, then taking the average of these numbers (1.073), and using it to "deflate" the average adjustment factor based on these same three age groups found for each municipality after application of our EB procedures to both the current and cohort fertility data. This procedure assumes that the 1996 DHS data are free from error, and that an adjustment that might be appropriate at the national level will also work at the municipal level. We are, of course, aware that the timing and magnitude of fertility decline varies considerably across Brazilian municipalities, and we will revisit this assumption later in the analysis.

Parenthetically, the decision to use three age groups (20-24, 25-29, and 30-34) rather than just the first two arose from our realization of how unstable the P/F ratios seemed to be across cohorts in the DHS data, especially when we began to look at these ratios for regions rather than the whole country. Thus, even though using the third cohort increases the temporal distance between our period and cohort rates, the added stability seemed especially appropriate when applying the method to small areas. Second, in the course of carrying out this analysis and drafting this report, we became aware of the disjunction between the "all or none" character of the P/F adjustment, at least when applied mechanically to large numbers of municipalities, and the Bayesian weighting logic that we had used to smooth the data. We will return to this issue in the discussion section.

RESULTS

EB Estimates of ASFRs and TFRs without Brass P/F Correction

We present the results of the EB estimates of the municipal ASFRs and TFRs before any adjustment in Table 2 and Graphics 1-4. The various panels of Table 2 show the mean value and standard deviation of both the smoothed (EB) and unsmoothed estimates of the TFR for Brazil as a whole and for the five regions of the country. Not surprisingly, the variation across municipalities in the smoothed estimates, though still very large, is much less than that for the smoothed estimates. Graphics 1 and 2 show how the effect of smoothing was barely detectable for the large municipalities, but made a large difference in both shape and level of fertility in very small municipalities. The scatter plots of the TFR by the log of population size both before and after smoothing (Graphics 3 and 4) show this same effect, as well as the elimination of absurdly low and high estimates of fertility for the smaller municipalities.

Graphic 5 is a plot that compares the EB estimates of the TFR for municipalities from the census with those from the smoothed estimates derived from vital registration, ordering the cases by region and according to the size of the difference between the estimates. The graph shows that the census estimates are frequently much higher than those from vital registration, especially in the North and Northeast regions, but that there are also a sizeable fraction of cases where the vital registration estimate is higher than the census estimate. At least upon first inspection, the former are indicative of incomplete registration of births, while the latter point to omissions or misreporting of births in the census.

EB Smoothing of Parity and the P/F Correction

The EB smoothing has an effect on the census estimates of average parity by age similar to the effect on the ASFR estimates in that the range and variation of the estimates is again compressed by shrinkage toward the mean, and the age pattern of the estimates for small municipalities is greatly improved. Both effects are evident Graphics 6 and 7 which plot the estimates of average parity in the age group 20-24 (P2) against that for the age group 25-25 (P3) before and after smoothing. The plot for the smoothed estimates is much more dense, always above the diagonal, and shows a much higher correlation between the two estimated parities. Graphics 8 and 9 show the before and after "funnel plots" for the estimates of these same two parities according to the size of the municipality, again showing a greater amount of compression among the smaller municipalities.

Cumulating the EB estimates of the municipal ASFRs with the Manual X formulas yielded period estimates of average parity (F) that can be compared with the smoothed estimates based on reported number of children ever born (P) in the respective age groups. Graphic 9 shows the average P/F ratio for the third, fourth and fifth age group in each municipality ordered by region, and by the size of the ratio. These ratios vary over a wide range extending from values near or slightly below 1 to values approaching 2 in all regions. Most values may be found between 1.2 and 1.4, except in the South where there is a higher proportion of ratios under 1.2, including a recognizable fraction below 1.0.

Graphic 10 shows an estimate of what could be called the slope of the P/F estimates defined as the difference between the mean of the P/F values for the fifth and sixth age groups in a municipality minus the average for the second and third age groups. We expected this measure to be highest in municipalities that had recently experienced a large decline in fertility, and lowest in those places where the fertility decline ended or slowed a decade or more in the past. Here we may observe a pronounced difference between regions, with much the greatest slopes found in the regions with the highest fertility (the North and Northeast) and noticeably lower slopes in the regions where the bulk of the fertility decline took place in the 1970s (the South and Southeast).

As noted above, our corrected estimate of the EB TFR was obtained by multiplying the original EB estimate by the average P/F ratio (where both the numerator and denominator come from smoothed estimates), and then dividing through by the average P/F found in the 1996 DHS. The mean and standard deviation of these estimates may be seen in Table 2. While correction tends to increase the overall level of the estimate (by 0.4 children, on average), it also tends to increase the standard deviation of the estimates.

In Table 2, the mean and standard deviation of our corrected estimates may be compared both to the EB estimates based on vital registration and the estimates made by the analysts who compiled the Human Development Indices (HDI) for all municipalities. While the corrected EB estimates derived from census data are considerably higher, on average, they are quite similar in level to HDI estimates. Graphics 11 and 12 show, however, that beneath this similarity in the average value of the estimates, there are considerable differences between the two actual values for individual municipalities. These seem to be especially pronounced at higher levels of fertility, and those regions where such levels are found.

A Closer Inspection of a Few Cases

To show what is actually happening in some specific cases, Graphics 13-16 provide detailed information for four selected municipalities including the number of women in the census sample, the various estimates of the TFR, period and cohort parity, smoothed and unsmoothed ASFRs, and the complete set of P/F ratios. The first case is a small municipality, Sao Miguel de Touros, in the Northeastern state of Rio Grande do Norte. Here the smoothing has the effect of lowering the ASFRs, vital registration is miserably incomplete, and estimated parity leads to a very large correction in by way of the P/F ratio. The corrected EB estimate turns out to be half a child lower than the HDI estimate of the TFR.

In the second case (Graphic 14), a fairly large municipality in the southern state of Rio Grande do Sul, smoothing increases the ASFRs, but the P/F correction reduces them back to their original level. This level, in turn, is lower than EB estimate of the TFR from vital registration. The P/F ratios are both low and flat.

The fifth case shown in Graphic 15 is the large metropolitan city of Recife in the Northeastern state of Pernambuco. Smoothing, of course, has a negligible effect on any of the estimates, however, the P/F ratios have both size and slope, and the corrected estimate is slightly higher than the HDI estimate and the original EB estimate, and equal to that based on vital registration.

The sixth and final case is for a very small municipality in the impoverished Northeastern state of Piauí (Graphic 16). Here smoothing has a huge effect on both the ASFRs and the TFR. The P/F ratios are reasonably low to begin with but rise with age. Vital registration seems to be quite incomplete, and there is a very large difference between the corrected EB estimate from the census data and the HDI estimate of the TFR. The HDI is 1.4 children higher.

DISCUSSION

The cases just presented reveal some of the strengths and weaknesses of the methods we have applied. The first case is a tough one. The sample size is small, vital registration is very incomplete, and the neighbors seem to have lower fertility as evidenced by the effect of smoothing on both the TFR and the estimates of parity at the higher ages. Moreover, the P/F ratios are extremely high and lead to almost doubling the EB estimate from the census data on births last year. There seems to be no sure ground. We can find no reason to second-guess the corrected EB estimate except possibly for having been less than sufficiently deflated for the effect of fertility decline on the P/F, and certainly prefer this estimate to the HDI estimate.

The second case is a classic example of a situation in which no correction for P/F is called for, much less a correction that incorporates an adjustment for a fertility decline in the recent past. Here the EB estimate and the vital registration estimates are extremely close, and the only estimate that looks truly out of line is the HDI estimate, which falls below the corrected EB estimate. In contrast to the preceding case, there is little doubt about the level of fertility—it is very near replacement.

Recife, the third case, has much in common with the preceding case. The differences are that smoothing is irrelevant, but correction for omission or misreporting of births in the census appears to be warranted, and the adjustment of the correction for recent fertility decline also seems to be appropriate. The corrected EB estimate closely coincides with the vital registration estimate, and again there is little doubt about the true level of fertility.

The fourth case was, we must admit, chosen to end on a positive note. Here both smoothing and correction for census misreporting play a key role, the adjustment of the correction for fertility decline again seems to be on target. Perhaps most important, the corrected EB estimate is very different from that constructed by the HDI analysts.

One difficulty that the cases have brought out, and that was also clear from the wide range in the slopes of the P/F ratios shown in Graphic 10 is the substantial amount of "real" variation in P/F that exists across municipalities. Adjusting for fertility decline by simply dividing these ratios by the national P/F ratio found in the 1996 DHS is hardly a satisfactory or sensitive response to the different situations. It should be possible to develop algorithms for distinguishing between the more recognizable situations, and avoiding a one size fits all approach to this challenge. Moreover, there is the both good and bad news that there is additional information that we can bring to bear on the question of both the level and trend in fertility. We have yet to examine the estimates of fertility from the preceding census for the same municipality, as well as the child-woman ratio from the full sample in the 2000 census.

But these possibilities also highlight the problem of keeping the task to a manageable size, and developing procedures that are simple and clear enough that they can readily be understood, and applied in a variety of settings. As we have gained familiarity with the benefits to be obtained from shrinkage across space, it has struck us that this whole problem of resolving inconsistencies between data from different sources might also be subject to both the Bayesian vocabulary and technology of smoothing. Might this be the future of indirect estimation?

REFERENCES

RM Assunção, CP Schmertmann, JE Potter, SM Cavenaghi, 2003. “Bayesian Multivariate Spatial Estimation of Small-Area Fertility Schedules”. Paper presented at 2003 meeting of the Population Association of America, Minneapolis.

W Brass, AJ Coale, et al. (1968). The Demography of Tropical Africa. Princeton, N.J., Princeton University Press.

AJ Coale, PD Demeny (1967). Manual IV: Methods of Estimating Basic Demographic Measures from Incomplete Data. New York, United Nations.

MA El-Badry,. (1961). “Failure of enumerators to make entries of zero errors in recording childless cases in population censuses”. *Journal of the American StatisticalAssociation* **56**(296): 909-924.

B Efron and C Morris, 1975. “Data Analysis using Stein’s Estimator and its Generalizations”. *Journal of the American Statistical Association*, **70**(350):311-319.

NT Longford, 1999. “Multivariate shrinkage estimation of small area means and proportions”. *Journal of the Royal Statistical Society A* **162**: 227-245.

RJ Marshall, 1991. “Mapping disease and mortality rates using empirical Bayes estimators”. *Applied Statistics* **40:** 283-294.

T Merrick, E. Berquó (1983). The determinants of Brazil's recent rapid decline in fertility. Washington, D.C., National Academic of Sciences.

JE Potter, CP Schmertmann, and SM Cavenaghi, 2002. “Fertility and Development: Evidence from Brazil”. *Demography* **39:** 739-761.

PNUD. 2000. Atlas do Desenvolvimento Humano no Brasil, EMS Consultoria.

LR Wong, JD Oliveira. 1984. “Níveis e diferenciais de fecundidade para o Brasil segundo os dados censitarios de 1980 (notas preliminares)”. Annals of the Encontro Nacional de Estudos Populacionais, Águas de São Pedro, ABEP.

LR Wong. 1986. “A diminuição dos nascimentos e a queda da fecundidade no Brasil dos anos pós 80”. *Annals* of the Encontro Nacional de Estudos Populacionais, Águas de São Pedro, SP, ABEP.

Saude, M. d. S. S. d. A. a. 2004. “Integração de Informações do SIAB, SIM e SINASC nos Sistemas Locais de Saúde”. *Informe da Atenção Básica*. **Jan-Feb of 2004**.

DO Sawyer, IHOP Correa. 1988. “Queda de fecundidade: estimativas de níveis e tendencias”. *Annals* of the Encontro Nacional de Estudos Populacionais, Olinda, PE, ABEP.

United Nations. 1983. Manual X: Indirect Techniques for Demographic Estimation. New York, United Nations.

Table 1
Descriptive statistics on total population, sample size of women aged 15-49 and area of municipalities, according to large regions, Brazil, 2000.

Variable	Obs	Mean	Std.Dev.	Min	Max
			Total Population		
Brazil	5,506	30,839	186,767	795	10,400,000
North	449	28,732	96,212	958	1,405,835
Northeast	1,786	26,730	97,737	1,308	2,443,107
Southeast	1,666	43,465	307,989	795	10,400,000
South	1,159	21,663	73,131	1,113	1,587,315
Center-West	446	26,091	118,903	895	2,051,146
			Women 15-49 sampled		
Brazil	5,506	991	6,492	22	363,733
North	449	856	3,334	25	48,182
Northeast	1,786	826	3,520	31	91,325
Southeast	1,666	1,441	10,645	22	363,733
South	1,159	699	2,525	33	56,304
Center-West	446	868	4,345	27	75,557
			Area in Km^2		
Brazil	5,506	598	2,216	1	62,506
North	449	3,326	6,631	41	62,506
Northeast	1,786	337	540	4	6,375
Southeast	1,666	215	292	1	4,159
South	1,159	188	270	10	3,011
Center-West	446	1,395	2,320	22	25,378

Table 2

Descriptive statistics and different estimates of Total Fertility Rates (TFR) in municipalities, according to large regions, Brazil, 2000.

TFR Estimates	Obs	Mean	Std. Dev.	Min	Max
Census	5506	2.52	0.84	0.12	11.96
Census EB	5506	2.40	0.55	1.49	6.30
Census EB corrected	5506	2.81	0.73	1.35	7.16
Vital Registration EB	5506	2.08	0.53	0.19	5.75
HDI	5506	2.86	0.74	1.56	7.79
North Region					
TFR Estimates	Obs	Mean	Std. Dev.	Min	Max
Census	449	3.35	1.16	0.84	11.96
Census EB	449	3.08	0.80	1.73	6.30
Census EB corrected	449	3.77	0.97	1.87	7.16
Vital Registration EB	449	2.34	0.71	0.39	5.75
HDI	449	3.85	1.02	1.96	7.79
Northeast Region					
TRF Estimates	Obs	Mean	Std. Dev.	Min	Max
Census	1786	2.87	0.76	0.62	7.39
Census EB	1786	2.74	0.48	1.70	6.30
Census EB corrected	1786	3.25	0.63	1.77	6.20
Vital Registration EB	1786	2.19	0.64	0.19	4.90
HDI	1786	3.23	0.68	1.69	6.80
Southeast Region					
TRF Estimates	Obs	Mean	Std. Dev.	Min	Max
Census	1666	2.24	0.68	0.26	8.05
Census EB	1666	2.13	0.37	1.56	4.12
Census EB corrected	1666	2.43	0.50	1.35	5.44
Vital Registration EB	1666	1.91	0.39	0.27	4.14
HDI	1666	2.52	0.52	1.56	5.12
South Region					
TRF Estimates	Obs	Mean	Std. Dev.	Min	Max
Census	1159	2.22	0.65	0.16	4.85
Census EB	1159	2.14	0.24	1.49	3.41
Census EB corrected	1159	2.40	0.36	1.64	3.93
Vital Registration EB	1159	2.07	0.37	1.10	4.14
HDI	1159	2.49	0.38	1.68	4.30
Center-West Region					
TRF Estimates	Obs	Mean	Std. Dev.	Min	Max
Census	446	2.15	0.62	0.12	4.58
Census EB	446	2.08	0.30	1.53	3.33
Census EB corrected	446	2.51	0.37	1.81	3.71
Vital Registration EB	446	1.96	0.36	0.78	3.70
HDI	446	2.65	0.46	1.79	4.83

Municipalities' boundaries in Brazil, 2000

Graphic 1

Age-specific Fertility Rates, before and after multivariate spatial smoothing
10 largest municipalities

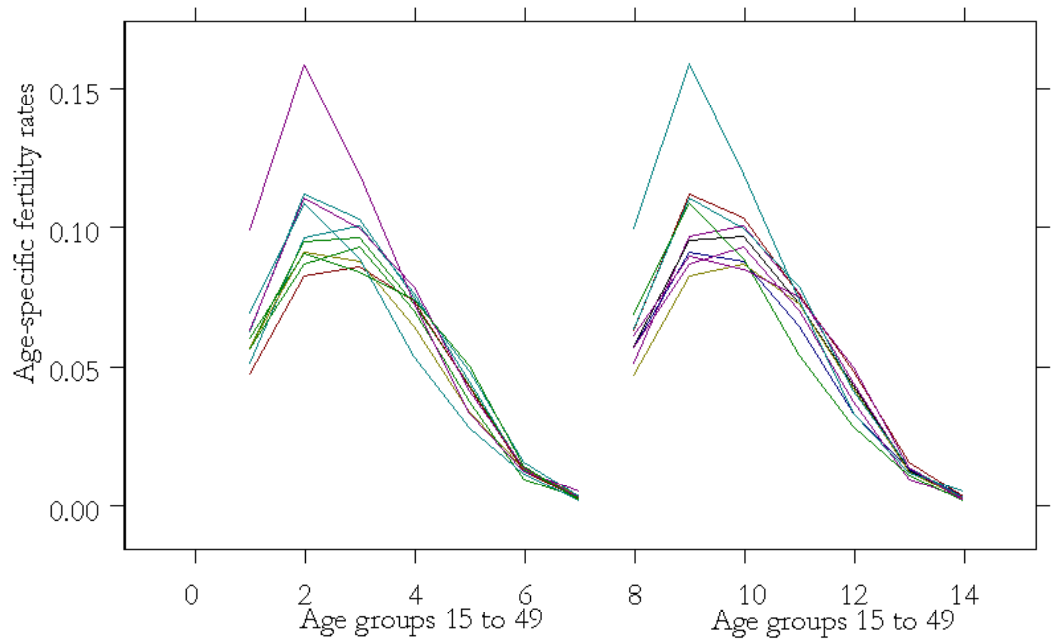

Graphic 2

Age-specific Fertility Rates, before and after multivariate spatial smoothing
10 very small municipalities

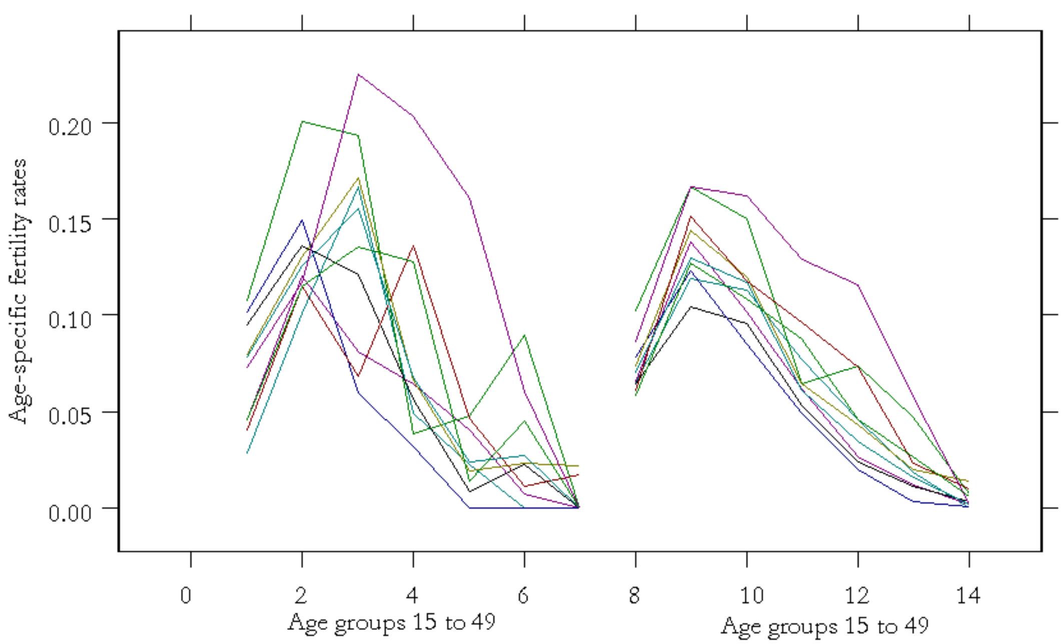

Graphic 3

Total Fertility Rate from children born in the year preceeding the census without corrections in municipalities, Brazil, 2000.

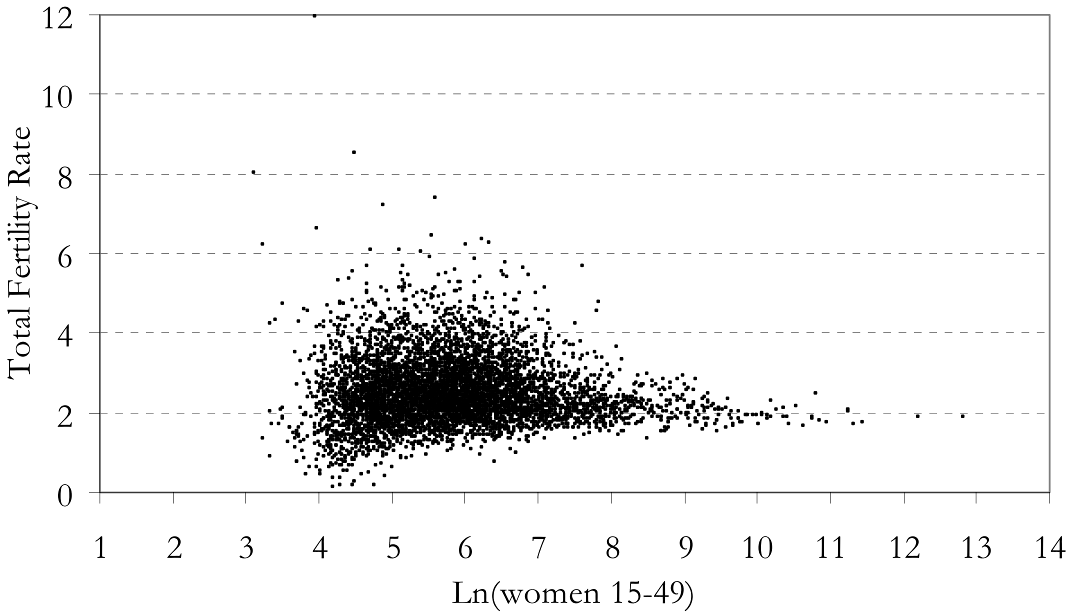

Graphic 4

Total Fertility Rate from children born in the year preceeding the census after empirical Bayes multivariate smoothing in municipalities, Brazil, 2000.

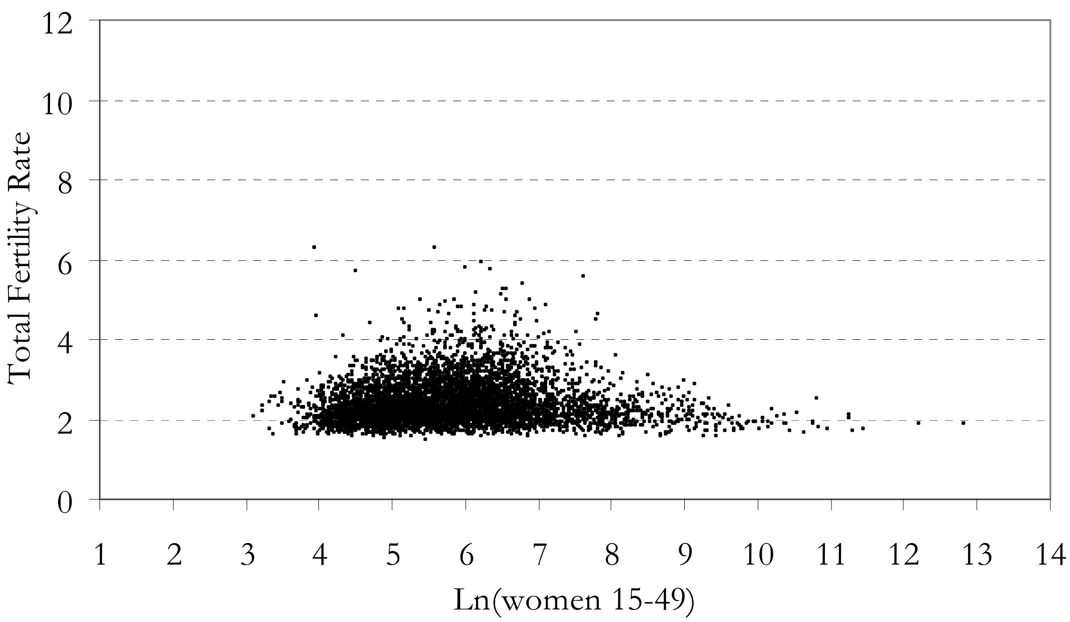

Graphic 5

Differences in TRF estimates in municipalities: Census empirical Bayes minus vital registration empirical Bayes (ordered by size in region), Brazil, 2000.

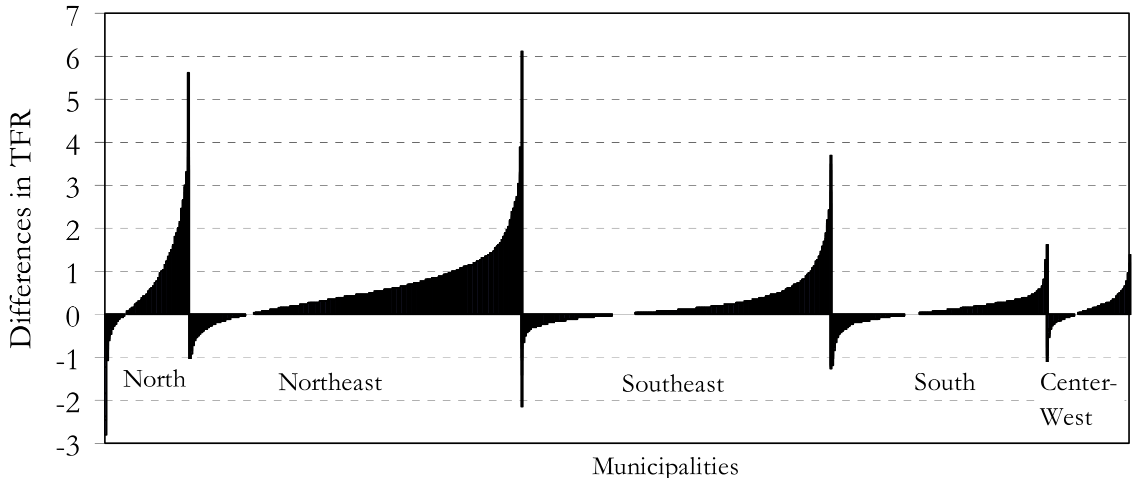

Graphic 6

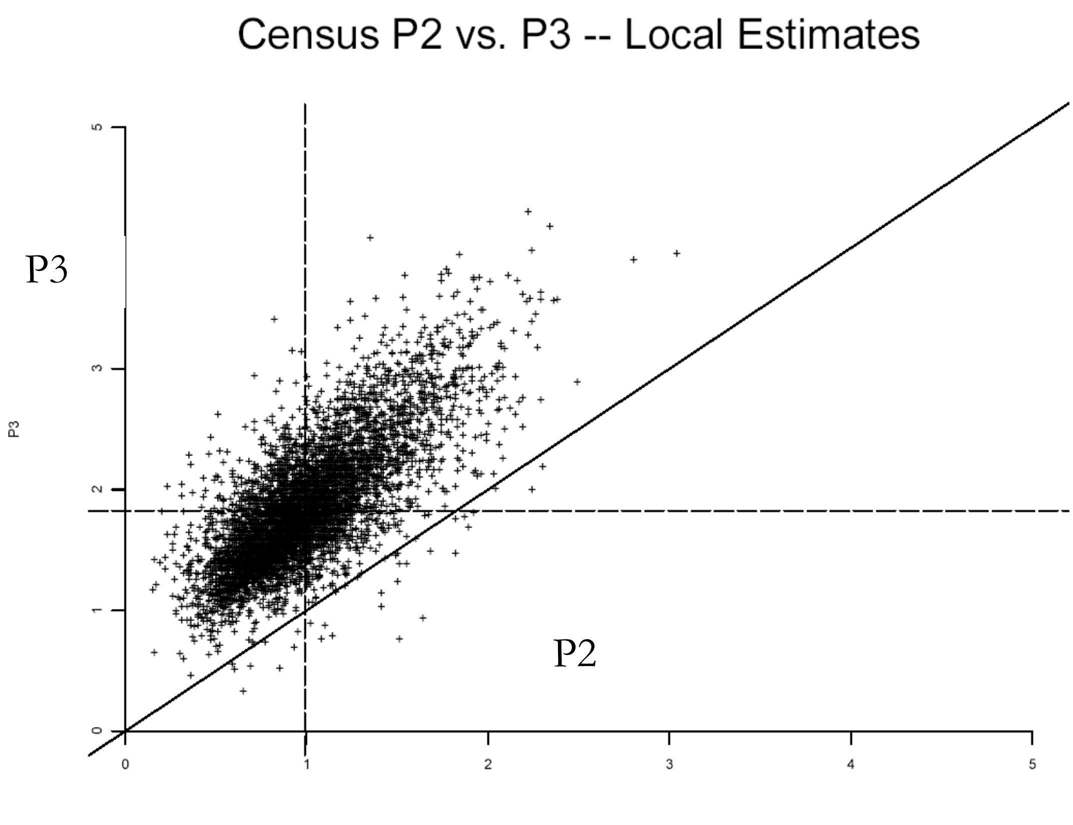

Graphic 7

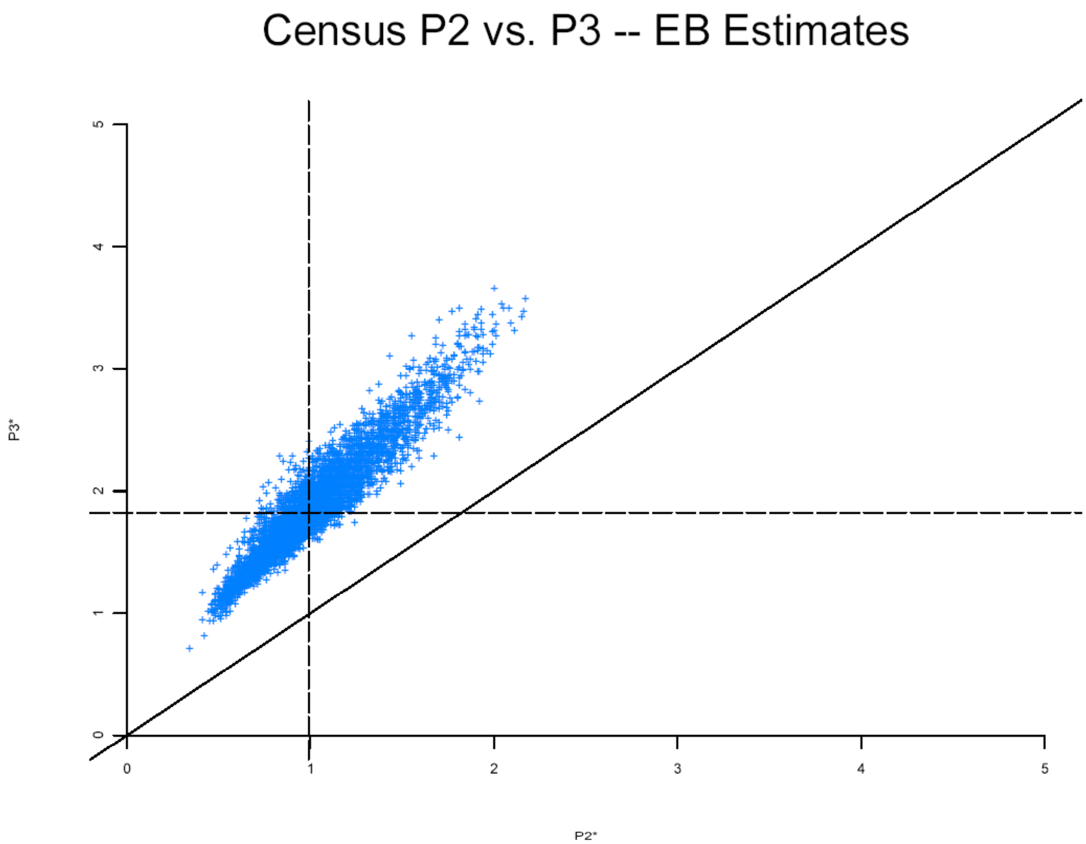

Graphic 8

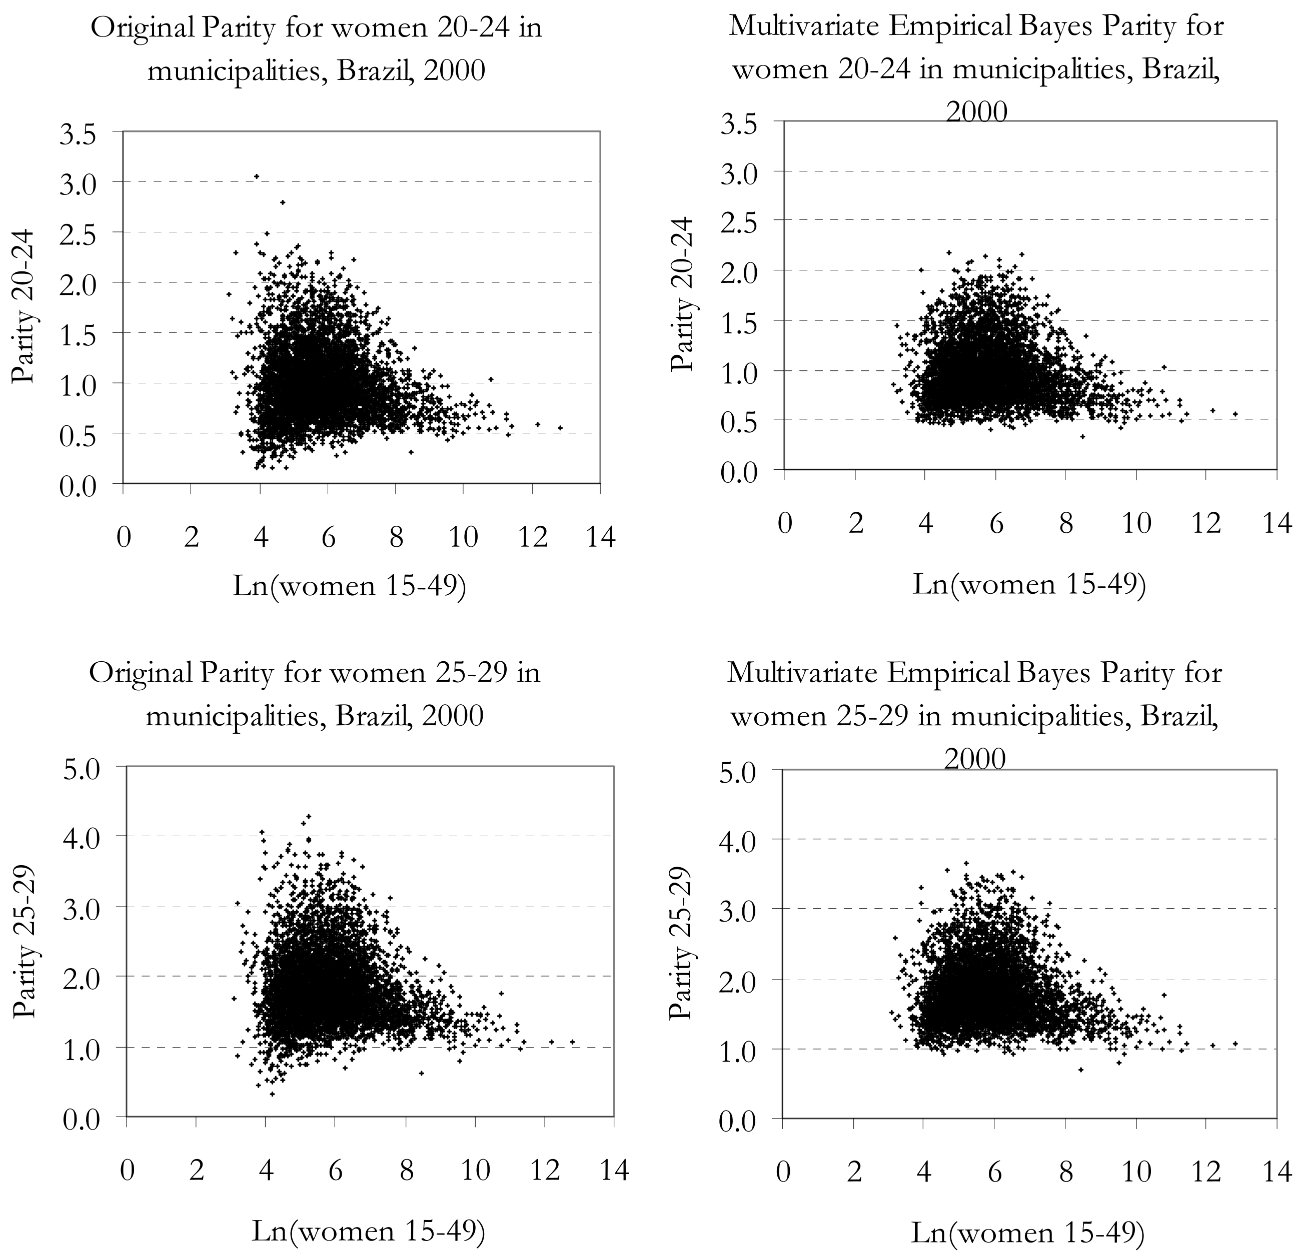
Original Parity for women 20-24 in municipalities, Brazil, 2000
Parity 20-24
Ln(women 15-49)
Multivariate Empirical Bayes Parity for women 20-24 in municipalities, Brazil, 2000
Parity 20-24
Ln(women 15-49)
Original Parity for women 25-29 in municipalities, Brazil, 2000
Parity 25-29
Ln(women 15-49)
Multivariate Empirical Bayes Parity for women 25-29 in municipalities, Brazil, 2000
Parity 25-29
Ln(women 15-49)

Graphic 9

P/F ratios, average for ages 20-24, 25-29, and 30-34 in municipalities, Brazil 2000. (ordered by P/F size and region)

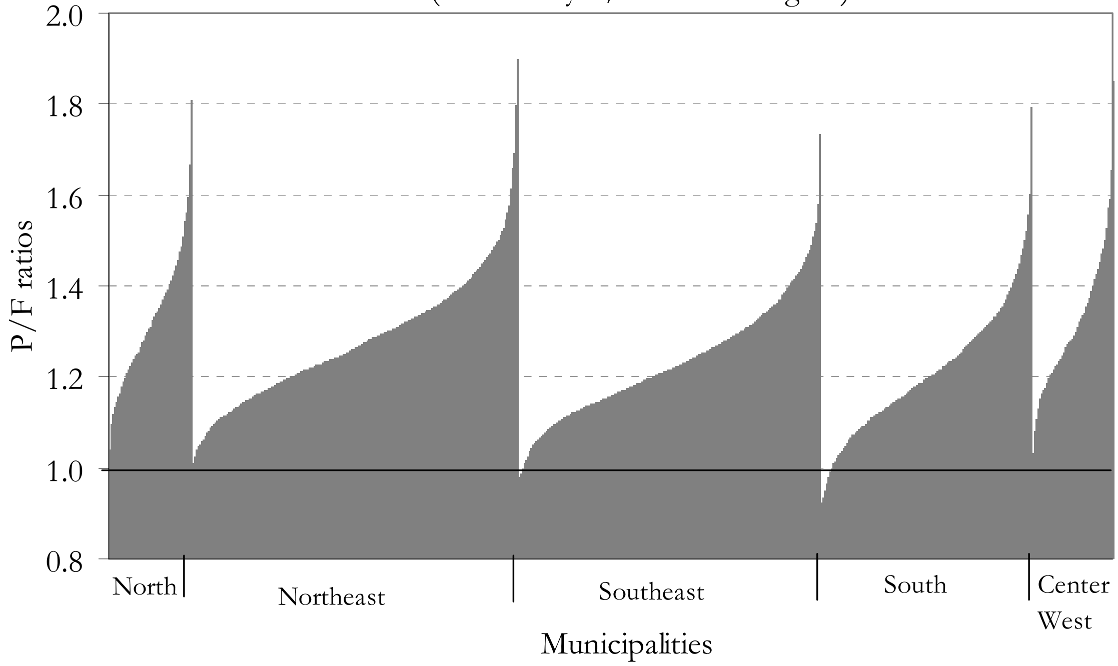

Graphic 10

Slopes of P/F (after smoothing) in municipalities (ordered by region), Brazil, 2000.

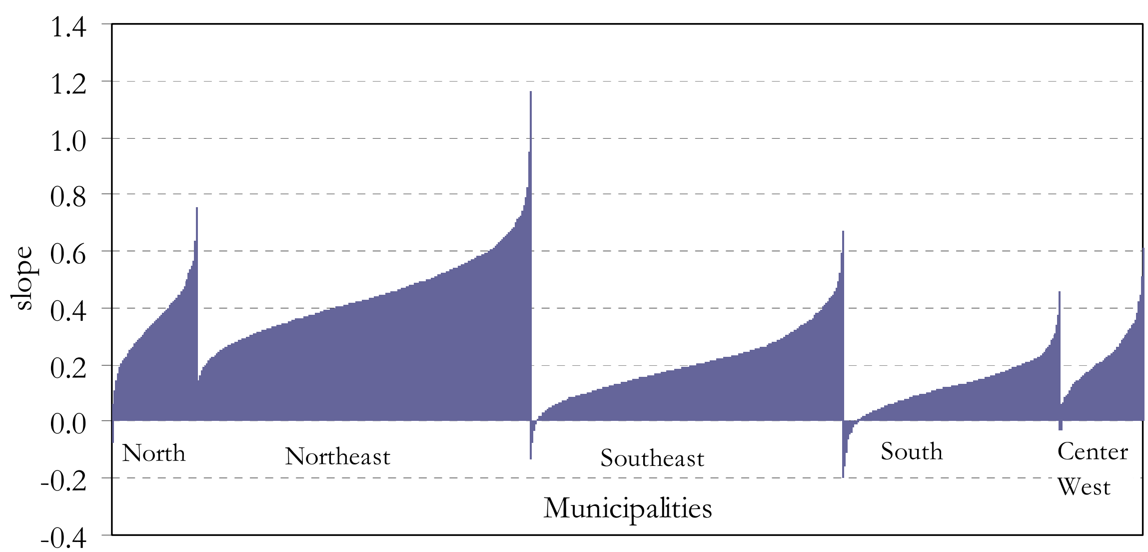

Graphic 11

Differences in TFR estimates in municipalities: Final estimates of total fertility rates minus official estimates for HDI (Human Dev. Indice), , Brazil, 2000.

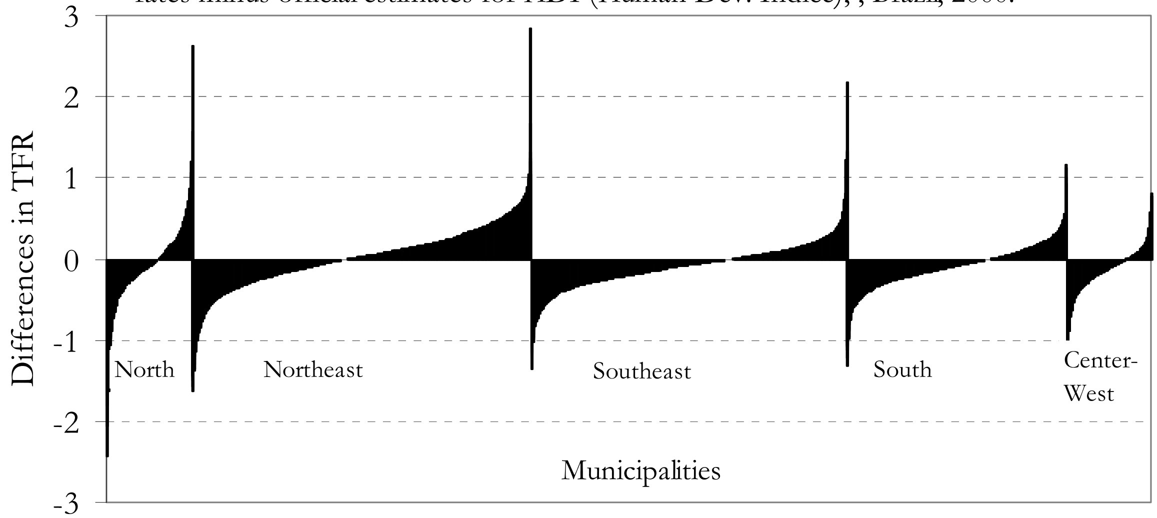

Graphic 12

Final estimates of total fertility rates and official estimates for HDI (Human Development Indice) in municipalities, Brazil, 2000.

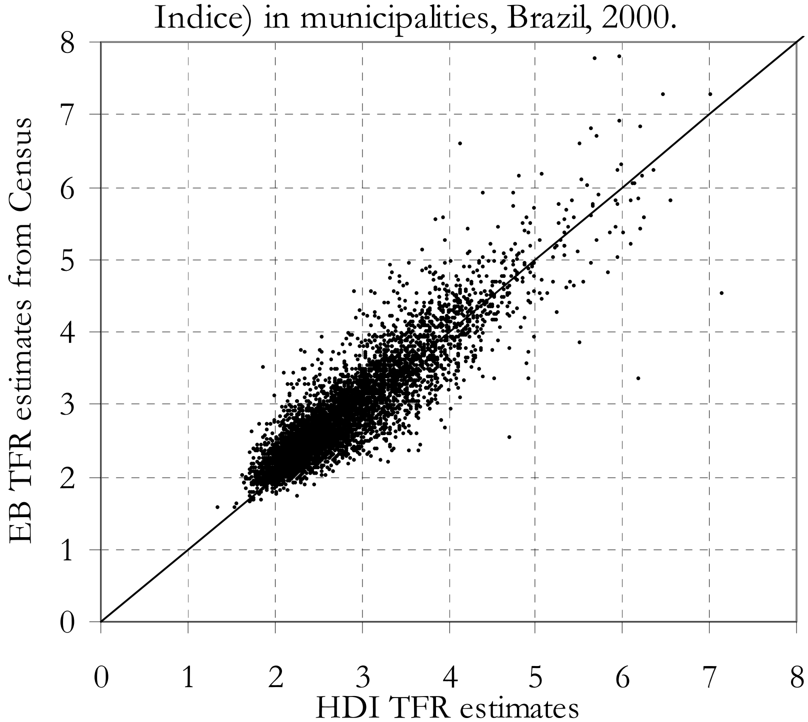

Graphic 13

Case 1

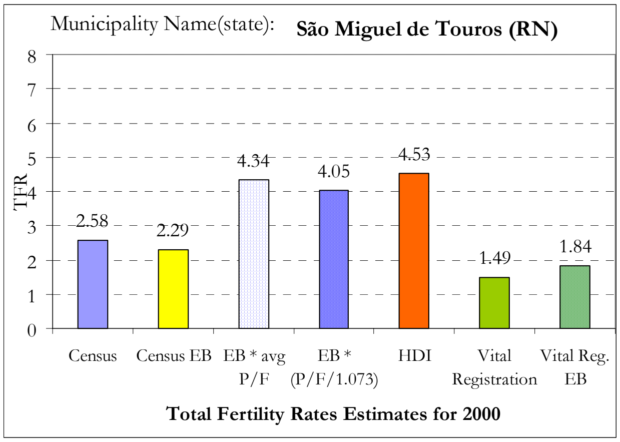

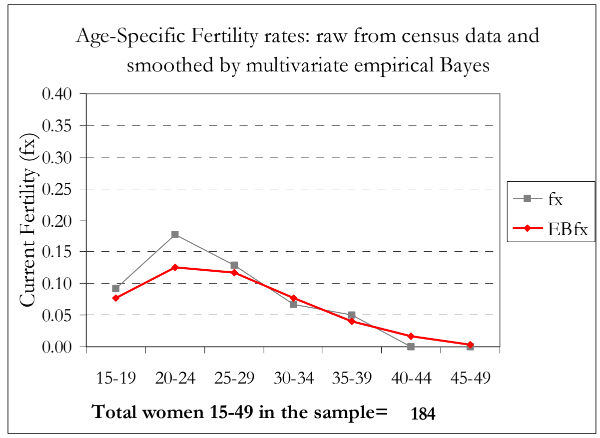

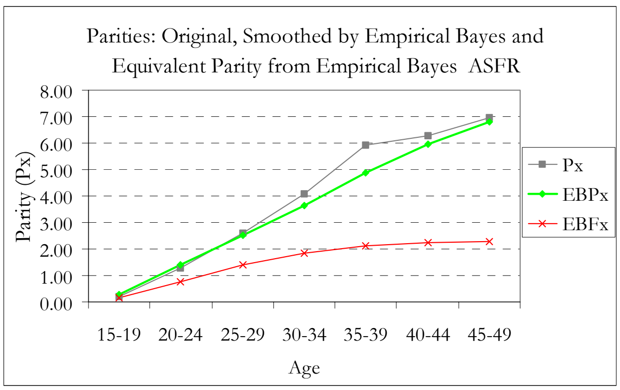

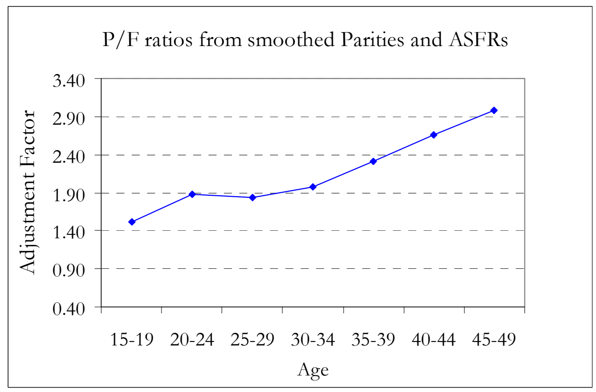

Case 2

Graphic 14

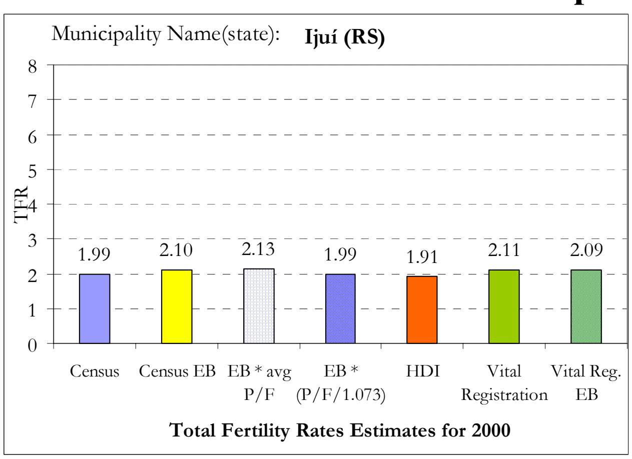

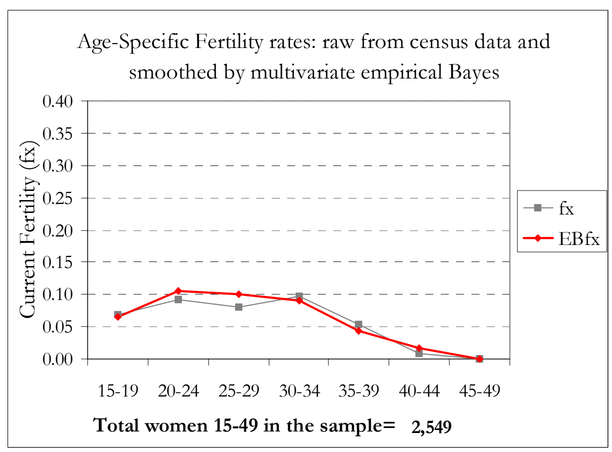

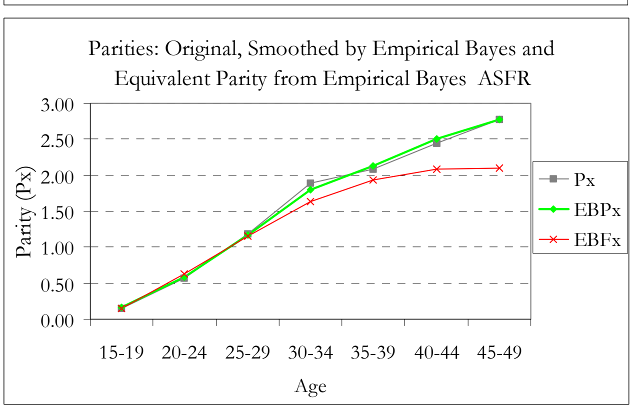

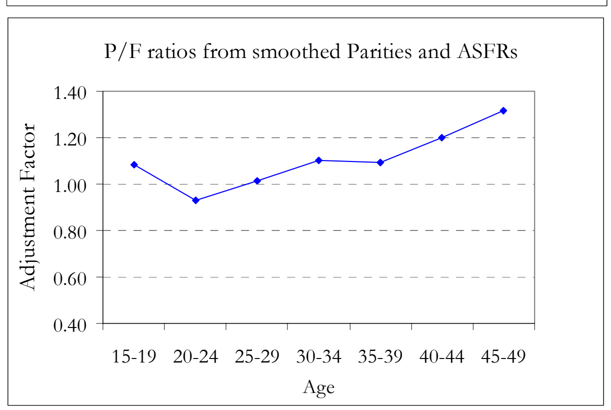

Case 3

Graphic 15

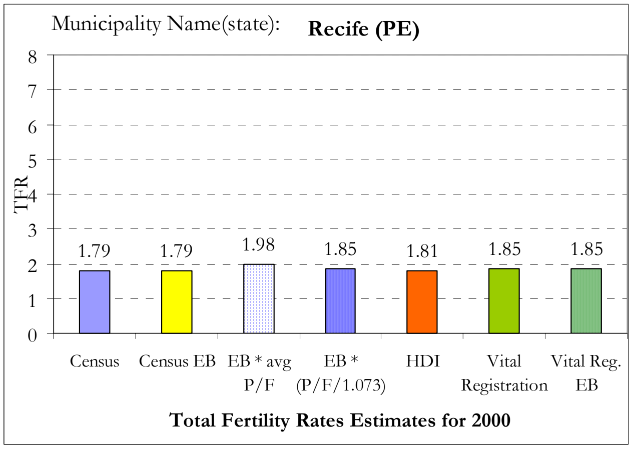

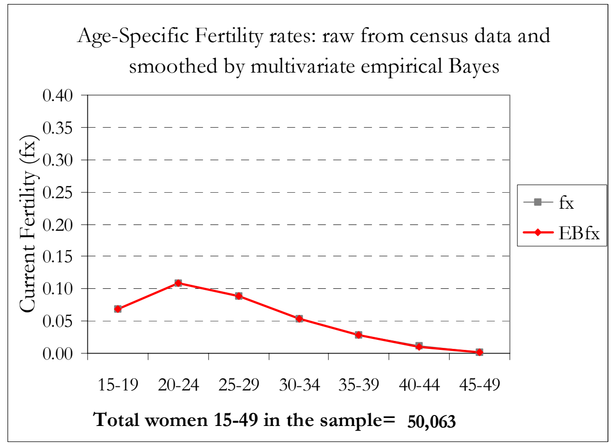

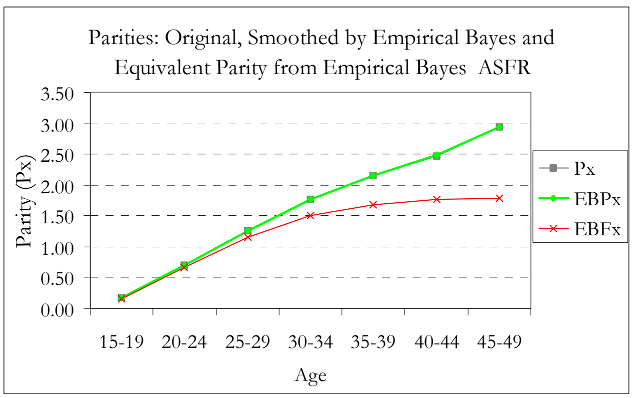

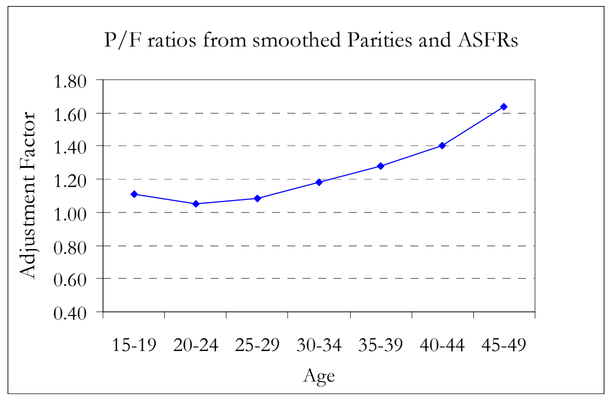

Case 4

Graphic 16