Tracing the income−fertility nexus: Nonparametric Estimates for a Panel of Countries

Holger Strulik
University of Hamburg

Siddiqui Sikandar
University of Hamburg

Abstract

We apply the a nonparametric method of kernel regression on a dataset for 109 countries to estimate the income−−fertility nexus in demo−economic transition. The results suggest the existence of a critical level of per capita income above which fertility decreases exponentially with rising income. For income levels below fertility stays on a high level and its relation to income is of minor significance. The critical income threshold changes over time and is lower for more recent periods, which gives evidence of major structural shifts in the relationship under investigation.

1. **Introduction**

For a long time the relationship between fertility and income has been regarded as an issue of high importance in development economics and policy. A popular concept is the idea that the fertility-income relationship takes the form of an inverted U. This notion was first raised by Birdsall (1980) and subsequently displayed in the 1984 World Development Report. A concept that has attracted similar attention is the so-called Norm curve, remarkably also shown in the 1984 World Development Report. The Norm curve, which describes fertility as a monotonically (and possibly exponentially) declining function of per capita income, was later taken up in Birdsall's (1988) article in the Handbook of Economic Development, and the fact that it reappears in Dasgupta's (1995) survey on the population problem gives evidence of its ongoing popularity.

From a theoretical point of view, the question whether there exists a range of income levels in which fertility reacts positively on income improvements is anything but trivial, because this idea has often been used to substantiate the possibility of fertility-driven poverty traps. Apart from the well-known Malthusian and Neo-Malthusian notions of decreasing returns to scale in production (see e.g. Niehans, 1963, Strulik, 1997), models of economic growth with endogenous fertility, as developed by Becker, Murphy and Tamura (1990), Ehrlich and Lui (1991), Tamura (1996), and Galor and Weil (2000) have provided theoretical justifications for an inverted U-shaped path of fertility transition. These contributions describe the individual decisions on fertility as a tradeoff between child quality and quantity, and show the existence of a stagnation equilibrium with high fertility and low per capita income on one hand, and of a growth equilibrium with increasing per capita income and low fertility on the other.

One of the implications of these theoretical models is that it should be possible to detect an income threshold in cross-sectional estimates of the income fertility relationship. Nevertheless, rather few attempts have so far been made to produce empirical support of the resulting hypotheses. Winegarden and Wheeler (1993) and Barro and Sala-i-Martin (1995, Ch. 12) impose a quadratic income fertility relationship and find evidence for an inverted U shape of the fertility function. The purpose of this paper is to supplement the ongoing debate about the "true" nature of the income-fertility nexus with empirical evidence from a panel dataset of 109 countries observed in quinquennial time intervals during the years 1960 to 1989. Because of the obvious uncertainty about the adequate specification of the underlying functional form – which itself predetermines the set of possible results – we decide to employ nonparametric estimation techniques in order to unearth the "correct" shape of the underlying statistical relationship. The strategy of applying our estimation procedure to each of the five cross-sections in the dataset separately enables us to detect structural shifts. Our analysis, however, does not address the question of causality (which may run from income to fertility, or vice versa,

or be mutual). Moreover, the fact that the dataset is relatively short in the time series dimension impedes any attempt to estimate country-specific long-run fertility transition paths in a reasonable manner.

A possible objection to this univariate approach might be that it fails to account for a considerable variety of other socio-economic factors, which may possess significant explanatory power for human fertility. It is, of course, not our intention to deny that factors like the status of women, school enrolment rates, or the availability of contraceptives do have an impact on fertility. The purpose of our investigation, however, is to isolate the impact of the income variable as precisely as possible. To this end it has to be avoided that the power of our results to either confirm or refute the predictions of related theoretical models is reduced by the use of additional regressors, which can be reasonably expected because virtually all of these regressors are highly correlated with income and would, therefore, absorb part of the income effect. Seen from this angle the use of a univariate yet nonparametric approach is appropriate, especially since it does not contradict the idea that variation in fertility that is not due to income differentials can at least in part be attributed to other determinants.

The remainder of the paper is as follows: Section 2 contains a description of the dataset in use and how our estimation procedure was applied. The results are presented and discussed in Section 3, and Section 4 concludes the paper. An Appendix describes the used estimation procedure.

2. Data and Empirical Implementation

The database we use was initially collected by Robert Barro and Jong-Wha Lee, and subsequently utilized in their paper on "International Comparisons of Educational Attainment" (Barro and Lee, 1993) as well as in several other empirical investigations related to growth theory. In an attempt to keep as much information as possible while avoiding misjudgements about the variation in the income-fertility relationship over time, we choose to retain only those observations where the variables of interest can be observed for at least two subperiods. This led to the removal of four extreme outliers (Kuwait, Oman, Saudi Arabia, and the United Arab Emirates) from the sample. Their inclusion would lead to the erroneous impression that average fertility in high-income countries underwent a significant increase during the 1980 to –85 period. Nevertheless, it has to be kept in mind that the removal of outliers will cause the estimate of the underlying functional relationship to appear a bit more unambiguous than it is in reality.

The income variable equals real GDP per capita in 1985 prices (originally from the Summers and Heston dataset version 5.5) and fertility is measured as the total fertility rate (children per woman), which was compiled from various issues of the World Bank's *World Tables*. Descriptive statistics of the variables in use can be found in the Appendix.

2

We employ a nonparametric kernel regression according to Härdle (1990) which has the advantage that it does not necessitate a priori restrictions regarding the functional form of the relationship between income and fertility. Without being supplemented with a measure of the uncertainty prevailing with regard to the true nature of the functional relationship the informational content of the kernel estimate, however, is very limited. It has therefore become a common practice to augment the results by estimating pointwise confidence intervals via bootstrap methods. A related technique which has the favorable property of incorporating the bias inherent in kernel estimates, while at the same time allowing for heteroscedasticity and correcting for local skewness, is the so-called wild bootstrap method due to Härdle and Mammen (1993). The general estimation procedure is explained in the Appendix. A difficulty arises because, as Härdle (1990, pp. 130-133) points out, the bias of kernel regression estimates tends to be larger at the boundaries of the observation interval than in its interior. This constitutes a particular problem here since a relatively large number of observations in our dataset are concentrated at these boundaries, and because the shape of the regression line at both ends of the observation interval is of considerable interest in the context of our investigation. In order to mitigate this complication, we therefore decide to choose a common bandwidth of $h = 0.3$ for all panel waves, which is slightly smaller than the asymptotically optimal one that would have been suggested by the cross validation criterion.

3. Results

Figs. 1 to 6 present the estimated income-fertility regression lines for subperiods from 1960 to 1985. Dots display actual observations and dashed lines show the upper and lower bound of the 95% confidence interval.

The results indicate that a linear approximation of the income fertility nexus, if assumed for the entire range of the income variable would indeed veil a lot of information contained in the dataset and is bound to result in incorrect inferences. Clearly, the quadratic approximation advocated by Winegarden and Wheeler (1993) and Barro and Sala-i-Martin (1995) would perform far better on a goodness-of-fit basis; yet due to the enormous sensitivity of polynomial-based estimates to the presence of outliers situated close to the boundaries, this way of proceeding might very well produce an inverted U shape as a mere statistical artefact. Our nonparametric estimates give some (although rather limited) support to the idea that, in a situation of extreme poverty, a modest rise in income is associated with an increase in fertility: For four periods (1960, -75, -80, and -85), we find that the fertility estimate obtained at the lower end of the income scale is located outside the 95% confidence interval around the peak of the estimated fertility function. However, since in all of these cases the difference between these two points is very small, it would obviously be an exaggeration to regard this finding as a proof of the inverted U hypothesis.

Moreover, our estimates show that in every year under investigation, there exists a certain income threshold above which the correlation between income and fertility is significantly negative. Between this point and the almost horizontal "floor" which occurs at the upper end of the income scale only in the estimates for 1960 and 1965, the nature of the estimated relationship can be captured very well by a straight, downward-sloping line. Keeping in mind that we use a logarithmic scale for per capita income, this implies that in the corresponding section of the income axis, fertility falls exponentially with rising income. The slope of this line undergoes remarkably little variation over different subsections of the curve as well as over time.

Our third main finding is that while the total fertility rate coordinate of the threshold does not change much from 1960 to 1985 and is approximately 6.5, the income coordinate changes considerably. This pattern can be summarized as follows: The more recent the observed period, the lower the level of income from which on income improvement goes hand in hand with a marked and persistent reduction in fertility. The observed thresholds for fertility transition range approximately from 1340 (i.e. $e^{7.3}$) international dollars income per capita per year in 1960 to 500 (i.e. $e^{6.2}$) in 1985.

While our investigation provides only limited evidence against a monotonously declining fertility function, it strongly confirms the observation of structural shifts in the fertility pattern stated in both, Birdsall (1980) and the Development Report (1984), and theoretically supported by Tamura (1996): More recent fertility declines occur at lower income per capita levels. The most marked improvements in this direction can be observed in the 1980's. Nevertheless, even in the 1980's there exists a club of least developed countries with fertility rates as high as in the 1960's and income per capita still on a very low level, which in turn suggests the existence of a low-income-high-fertility trap.

4. Conclusion

By employing nonparametric techniques to a panel of countries we find only weak evidence for a positive income-fertility nexus at very low income levels. For developing countries having passed a certain income threshold, rising income is accompanied by exponentially falling fertility rates, and the (negative) growth rate for fertility appears to be almost constant over the whole range of developing countries and over time. Furthermore, this negative relationship could not be obtained for all income levels as suggested by the Norm-Curve. Our analysis produces evidence for the existence of a low-income-high fertility trap at very low income per capita levels, whereas the threshold for development decreases over time. Our analysis does not support simple Malthusian reasoning, but is very well in line with predictions deduced from modern models of endogenous growth with endogenous fertility.

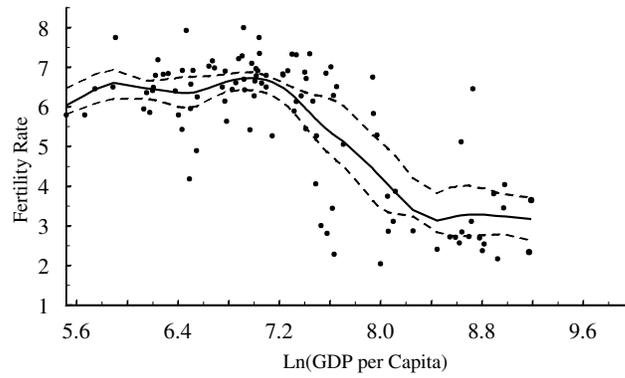Figure 1: Fertility as a Function of Per Capita Income 1960

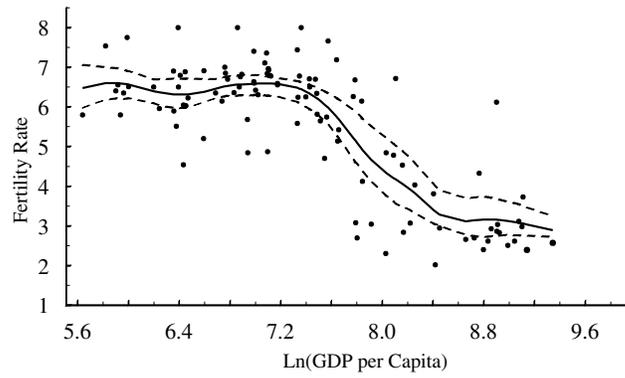Figure 2: Fertility as a Function of Per Capita Income 1965

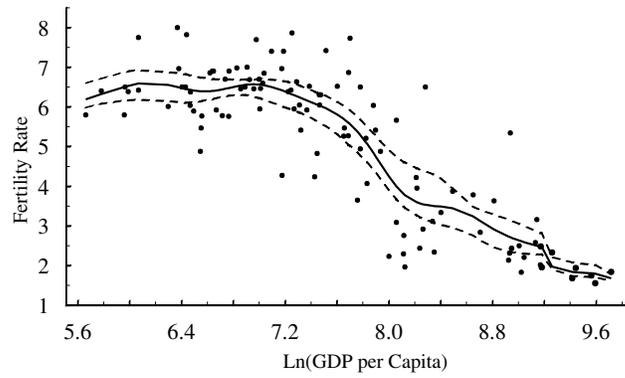Figure 3: Fertility as a Function of Per Capita Income 1970

5
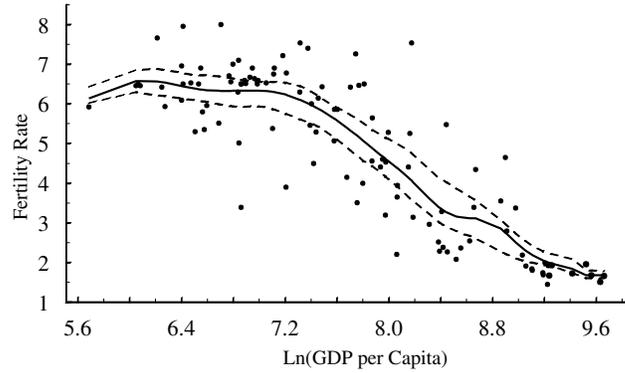
Figure 4: Fertility as a Function of Per Capita Income 1975

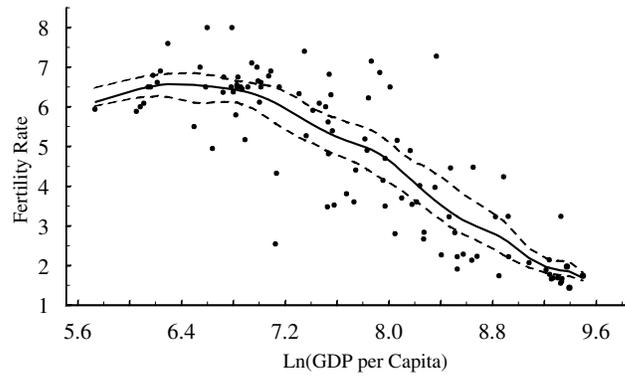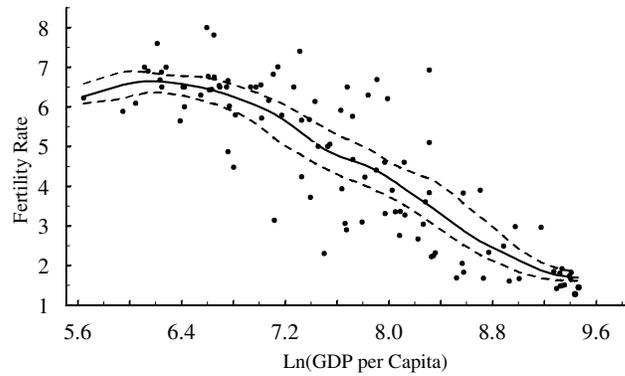Figure 5: Fertility as a Function of Per Capita Income 1980

Figure 6: Fertility as a Function of Per Capita Income 1985

References

Barro, R.J., and J.-W. Lee (1993), "International Comparisons of Educational Attainment", *Journal of Monetary Economics* **32**, 363–394 and ftp://www.nber.org/pub.barro.lee

R.J. Barro, and X. Sala-i-Martin (1995) *Economic Growth*, McGraw-Hill: New York,

Becker, G.S., K.M. Murphy, and R. Tamura (1990), "Human Capital, Fertility, and Economic Growth", *Journal of Political Economy* **98**, 12–37.

Birdsall, N. (1980), Population Growth and Poverty in the Developing World, World Bank Staff Working Paper 404, Washington D.C.

Birdsall, N. (1988), "Economic Approaches to Population Growth", in *Handbook of Development Economics*, by H. Chenery and T.N. Srinivasan, Eds, Vol. 1, Elsevier: Amsterdam.

Dasgupta, P. (1995), "The Population Problem: Theory and Evidence", *Journal of Economic Literature* **33**, 1879–1902.

Ehrlich, I., and F.T. Lui (1991), "Intergenerational Trade, Longevity, and Economic Growth", *Journal of Political Economy* **99**, 1029–1059.

Galor, O., and D. N. Weil (2000) "Population, Technology, and Growth: From Malthusian Stagnation to the Demographic Transition and Beyond", *American Economic Review* **90**, 806-828.

Härdle, W. (1990), *Applied Nonparametric Regression*, Cambridge University Press: Cambridge.

Härdle, W., and E. Mammen (1993), "Comparing Nonparametric versus Parametric Regression Fits", *Annals of Statistics* **21**, 1926–1947.

Nadaraya, E.A. (1964), "On Estimating Regression", *Theory of Probability Applied* **10**, 186–190.

Niehans, J. (1963) "Economic Growth with Two Endogenous Factors", *Quarterly Journal of Economics* **77**, 349–371.

Strulik, H. (1997), "Learning-by-doing, Population Pressure, and the Theory of Demographic Transition", *Journal of Population Economics* **10**, 285- 298.

Tamura, R. (1996) "From Decay to Growth: A Demographic Transition to Economic Growth", *Journal of Economic Dynamics and Control* **20**, 1237–1261.

Watson, G.S. (1964), "Smooth Regression Analysis", *Sankhya, Series A*, **26**, 359–372.

Winegarden, C.R., and M. Wheeler (1993), "The Role of Economic Growth in the Fertility Transition if Western Europe: Econometric Evidence", *Economica* **59**, 421–435.

The World Bank (1984), World Development Report 1984, New York (Oxford University Press).

Appendix

TABLE 1. DESCRIPTIVE STATISTICS

	FERT 1960	FERT 1965	FERT 1970	FERT 1975	FERT 1980	FERT 1985	GDP 1960	GDP 1965	GDP 1970	GDP 1975	GDP 1980	GDP 1985
Mean	5.51	5.43	5.15	4.84	4.62	4.42	2246	2664	3092	3516	4005	4117
Std. Dev	1.72	1.72	1.83	1.94	1.99	2.03	2245	2717	3141	3475	3963	4303
Min	2.02	1.81	1.55	1.45	1.44	1.28	249	281	287	294	308	283
Max	8.00	8.00	8.00	8.00	8.00	8.00	9774	11529	13299	13461	15101	16559

Estimation Procedure

Let Y denote the dependent variable and X the scalar explanatory variable. Assuming that a total of N independent random draws $\{y_i, x_i\}, i = 1, \ldots, N$, from the joint distribution of Y and X are at our disposal, and that both of these variables are continuously distributed, we can formulate the nonparametric model as

$$(1) \qquad y_i = m(x_i) + \epsilon_i, \quad E(\epsilon|X) = 0, \quad i = 1, \ldots, N \ ,$$

where ϵ represents the error term. As proposed by Nadaraya (1964) and Watson (1964), the expression $m(x)$, which characterizes the conditional mean of Y given $X = x$, can then be estimated by

$$(2) \qquad \widehat{m}_h(x) = \frac{N^{-1} \sum_{i=1}^{N} K\left[(x - x_i)/h\right] y_i}{N^{-1} \sum_{i=1}^{N} K\left[(x - x_i)/h\right]} \ ,$$

where $K(\cdot)$ denotes the *kernel*, a continuous, bounded and symmetric real function that integrates to unity. The scalar h is a smoothing parameter (referred to as the *bandwidth* that determines the "roughness" of the estimated curve. There is a considerable variety of possible choices for the function $K(\cdot)$; yet since it is well known that, compared to the selection of the bandwidth, the choice of the kernel is of minor importance for the properties of the resulting estimate, we use the Gaussian Kernel,

$$(3) \qquad K_G(u) := \frac{\exp(-0.5u^2)}{\sqrt{2\pi}} \ .$$

As is shown by Härdle (1990, p. 29), the Nadaraya-Watson method produces a consistent estimate of $m(\cdot)$ if a number of regularity conditions are fulfilled. Yet in practice, when the sample size N is fixed, its application implies a bias-efficiency tradeoff which has to be resolved by an appropriate selection of the smoothing parameter h. A popular bandwidth selection rule is to regard the particular value of h as optimal that minimizes the average

squared error criterion

$$(4) \qquad d_A := N^{-1} \sum_{i=1}^{N} \left(m(x_i) - \widehat{m}_h(x_i) \right)^2 \ ,$$

which in turn equals the sum of variance and squared bias. A consistent estimate of this unknown quantity is the cross validation criterion

$$(5) \qquad CV(h) := N^{-1} \sum_{i=1}^{N} \left(y_i - \widehat{m}_{h,i}(x_i) \right)^2 \ ,$$

where $\widehat{m}_{h,i}(x_i)$ is the *leave-one-out-estimation* obtained by deliberately skipping observation $\{x_i, y_i\}$ upon computation of the Naradaya- Watson estimate of $m(x_i)$. The mechanical application of this selection rule, however, can be misleading, especially if the sample size N is rather small. As Härdle (1990, p. 165) points out, the discrepancy between the finite-sample minimum of $CV(h)$ and the true minimum of d_A can be substantial, and the speed at which the finite-sample estimate of the optimal bandwidth tends to its true value as N increases is very slow. It follows that in practice, one might be forced to choose an appropriate bandwidth "by eye" whenever N appears to be too small for the finite-sample bias in $CV(h)$ to be negligible.

Panel 1: Summary of the Wild Bootstrap Algorithm
$b = 1$
Repeat
    Step 1: Sample ϵ_i^* from the two-point distribution with
    $P(\epsilon_i^* = \hat{\epsilon}_i \cdot \beta) = \delta; \quad P(\epsilon_i^* = \hat{\epsilon}_i \cdot \gamma) = 1 - \delta$
    with $\beta = (1 - \sqrt{5})/2$, $\gamma = (1 + \sqrt{5})/2$, $\delta = (5 + \sqrt{5})/10$
    Step2: Construct new observations $y_i^* = \hat{m}_g(x_i) + \epsilon_i^*$,
    where $\hat{m}_g(x)$ is an estimate of $m(\cdot)$ based on a bandwidth g
    that is slightly larger than h, and compute an estimate $\hat{m}^*(x)$
    from the artificially generated sample $\{y_i^*, x_i\}$, $i = 1, \ldots, N$.
    Step3: Raise b by 1
Until $b > B$ (= number of bootstrap samples).
Define the lower bound of the estimated $\alpha \cdot 100\%$ confidence interval
as the $(\alpha/2)$ quantile of the B bootstrap estimates, and the upper
bound accordingly.

The precision of the estimates obtained increases as the number of bootstrap samples, B, rises.

9