

The Gift of the Dying:
The Tragedy of AIDS and
the Welfare of Future African Generations*

Alwyn Young
Graduate School of Business
University of Chicago
This Draft: May 2004

(*) For Lillie, who gave more than she received.


I. Introduction

In coming years, the AIDS epidemic will visit upon the peoples of southern Africa a plague of immeasurable horror. The growth of HIV infection amongst African adults during the 1990s far exceeded most forecasts, with infection rates amongst pregnant women in, for example, South Africa reaching 25 percent by the year 2000 (Figure 1.1). While adults with HIV typically remain asymptomatic for 8 or 9 years, they develop AIDS soon thereafter and quickly succumb to the torturous infections and cancers which overwhelm their degraded immune systems. Thus, existing infection rates, in and of themselves, presage the imminent death of at least five million young South African adults aged 15-49 in the coming decade,¹ without consideration of the millions among future cohorts who will be infected at birth or as young adults.

In considering the economic consequences of the AIDS epidemic, one is drawn to historical examples of similar demographic catastrophes, perhaps the most well researched of which is the Black Death in Britain in the late 14th century. Reaching Britain in 1348, and recurring in a series of waves, the Black Death is estimated to have reduced Britain's population to about half of its pre-plague level by 1377 (Figure 1.2). As in the case of AIDS today, contemporary accounts of the human suffering brought about by the plague are depressing in the extreme. Economic data, however, paint a surprisingly different picture. With a declining labour force, real wages rose rapidly during the plague years, and then *remained high* throughout the 15th century, as Britain's population stagnated (Figure 1.3). While early scholars (e.g. Saltmarsh 1941, Hatcher 1977) argued that the 15th century was a period of continued high mortality, more recent work (e.g. Poos 1991, Goldberg 1992) has emphasized the positive role played by the high post-plague wages in increasing female labour market participation and lowering fertility, generating a self-reinforcing cycle of high incomes and labour scarcity that lasted for

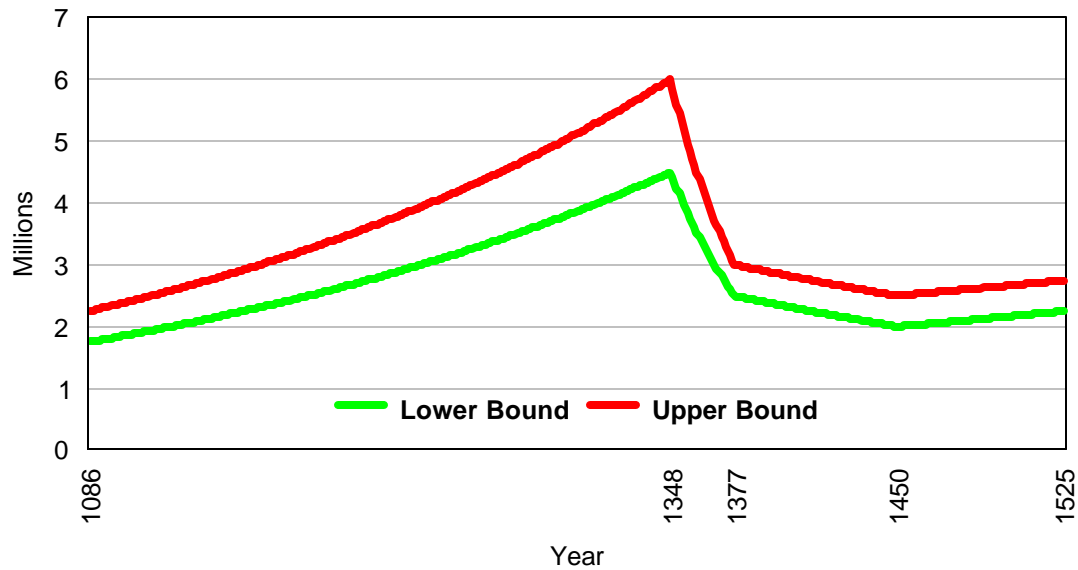
¹ I weight the age specific seroprevalence rates by the population distribution, yielding an average infection rate of 21% amongst 15-49 year olds. As explained later, the seroprevalence rates may substantially understate the true population infection rates, as medical studies indicate that HIV lowers the fertility of infected women by as much as a half.

Figure 1.1: HIV Prevalence - South Africa
(seroprevalence tests of pregnant women)




S: Department of Health (2001).

Figure 1.2: Population of England, 1086-1525



S: Hatcher (1977), pp. 68-69.

Figure 1.3: Real Wage of English Craftsmen



S: Brown & Hopkins (1956).

more than a century. Regardless of the precise mechanism, it is clear that the Black Death, in a purely economic sense, was a boon to the generations which survived and succeeded it, who, for a sustained period of time, experienced living standards not seen again until the late 19th century.

In this paper I try to calculate the impact of the AIDS epidemic on future living standards in South Africa. Drawing inspiration from the recent historical work on the Black Death, I do so in the context of a Beckerian model which endogenizes participation, fertility, and education decisions with behavioral equations estimated off of household survey data. While exploring the role the reduced supply of current and future labour plays in raising living standards, I also consider the detrimental impact of the epidemic on the accumulation of the human capital of orphaned children. I find that, even with extremely pessimistic assumptions concerning reductions in educational attainment, the labour supply effect dominates. From the perspective of per capita living standards, the AIDS epidemic endows society with additional resources. These resources could be used to care for the afflicted and provide higher living standards to future generations.

The paper proceeds as follows: Section II below details the model, a simple Beckerian household model embedded in a Solovian constant-savings macroeconomic framework. Section III explains the sources of data, while Section IV estimates the household behavioral coefficients. Section V calibrates the macroeconomy and lays out my assumptions concerning the length of the epidemic and its impact on realized fertility and children's education. Section VI simulates the economy under a variety of scenarios, showing the importance of the fertility and education assumptions, and the role the Beckerian behavior plays in amplifying and extending the model's dynamics. Section VII concludes.

II. Model: Becker and Solow

I model household behavior in the Beckerian tradition, assuming that couples, endowed with human capital E_m (male) and E_f (female) and financial resources Y , derive utility from the quantity of children n , the "quality" of children q (as measured by their human capital), individual leisure l_m and l_f , and material consumption C_m . In a one-period framework, each couple maximizes

$$(2.1) \quad U(n, q, l_m, l_f, C_m)$$

subject to the non-linear budget constraint

$$(2.2) \quad l_m W_m + l_f W_f + n t_f W_f + n q P_q + P_c C_m \leq Y + T^*(W_m + W_f)$$

where the W_i are the returns to labour market participation (functions of E_i), P_c and P_q the financial costs of consumption and quality, T the individual endowment of time, and t_f the time cost of fertility, which I assume is borne principally by women. I will use household surveys to estimate a simple utility-consistent system of demand and then use it to predict behavior.

The strictures of data severely limit what can be estimated. It is difficult to measure variation across households in the cost of education, so I shall assume, perforce, that the cost of education is proportional to the cost of material consumption ($P_q = a_q P_c$). In third world household surveys, data on assets is generally unavailable and, in any case, is endogenously determined by individual productivities. Consequently, I ignore income effects. The only remaining plausibly exogenous variation across households then lies in each couple's educational endowments, E_m and E_f . In principal, this allows the estimation of a maximum of two price effects in each demand equation. In practice, however, spousal educational attainment is unavailable for many individuals (who are either divorced, widowed or as yet unmarried) and, even when available, moves quite closely with own education. For children, mothers are identified much more fre-

quently than fathers. These considerations lead me to specify each component of household demand as a function of only one relative price:

$$(2.3) \quad \text{Male Labour Supply} = LM(W_m/P_c)$$

$$\text{Female Labour Supply} = LF(W_f/P_c)$$

$$\text{Fertility} = F(W_f/P_c)$$

$$\text{Children's Education} = Q(W_f/P_c)$$

As I will show further below, this simplistic system performs extraordinarily well in predicting out of sample household behavior.

The Hicksian composite commodity theorem states that the expenditure function and underlying preferences can be reexpressed as functions of the aggregates of products with common costs (Deaton and Muellbauer 1980). Despite the non-linearity of the budget constraint, this theorem can be applied to the problem above. Thus, the only requirement for consistency with utility maximization is that the matrix of compensated substitution terms for the negative of male labour supply ($-h_m = l_m - T$), the negative of female labour supply ($-h_f = nt_f + l_f - T$), and overall financial expenditures ($C = nqP_q + P_c C_m$),

$$(2.4) \quad \frac{\partial \mathbf{X}}{\partial \mathbf{P}'} + \frac{\partial \mathbf{X}}{\partial Y} \mathbf{X}', \text{ where } \mathbf{X}' = (-h_m, -h_f, C) \text{ and } \mathbf{P}' = (W_m, W_f, P_c),$$

be a symmetric negative semidefinite matrix.² In the context of the demand system (2.3), this reduces to requiring that LM' and LF' be nonnegative. Given this, one can integrate to derive the money metric indirect utility function

$$(2.5) \quad V(Y, w_m, w_f) = Y + \int_0^{w_m} LM(x) dx + \int_0^{w_f} LF(x) dx,$$

$$\text{where } w_i = W_i/P_c,$$

²This is a straightforward application of theorem 8 in Epstein (1981).

and associated household utility function

$$(2.6) \quad U(C, h_m, h_f) = C - \int_0^{h_m} w_m(x) dx - \int_0^{h_f} w_f(x) dx,$$

$$\text{where } w_m() = LM()^{-1} \text{ and } w_f() = LF()^{-1}.$$

Household utility is simply given by total consumption expenditures minus the disutility of labour, i.e. the area under the individual labour supply curves.

The preceeding one period model can be extended to a multiperiod lifetime setting by, in the standard macro-theorist's fashion, assuming time separable utility, so that (2.3) and (2.6) represent the period by period household demands and flow of utility. Aggregating (2.6) across all individuals and dividing by the total population, one sees that average utility per capita is given by consumption per capita minus the disutility of labour per capita. If labour supply is of the isoelastic form $h_i = \gamma w_i^{\varepsilon_i}$, per capita utility is then given by:

$$(2.7) \quad u = y \left[(1-s) - \left(\frac{\varepsilon_m}{1+\varepsilon_m} \right) \theta_{L_m} - \left(\frac{\varepsilon_f}{1+\varepsilon_f} \right) \theta_{L_f} \right]$$

where y is output per capita, s the savings rate, and θ_i the share of each factor in total income.

For constant savings rates and factor shares, this justifies the popular fixation on output per capita as a measure of welfare.

To close the model, one must specify the evolution of the macroeconomy and, in particular, the capital stock. I follow Solow and assume that the savings rate, by some happenstance of events and forces, remains fixed. This creates a simple benchmark for the evolution of the economy in the absence of AIDS. Keeping the same savings rate in the presence of the epidemic allows one to focus on the mechanisms, e.g. population growth, participation rates and educational levels, through which it influences the development of output per capita. One can then also consider what compensation, in the form of initial capital, the economy would need to

deliver the non-AIDS benchmark flow of welfare to future generations. The sign and magnitude of this compensation indicates the resources or burdens which, through the savings rate, will be distributed across future generations.³

³Thus, one can see the presentation as part of a Ramsey problem in which the planner maximizes the discounted flow of per capita utility by changing national savings and investment through the usual tax and expenditure mechanisms, but with no direct control over household decisions. For appropriate choices of the discount parameter and elasticity of intertemporal substitution, the planner will choose the same initial and steady state savings rate, with some (hopefully minimal) dynamics in between. Taking this as the benchmark, one then evaluates the change in the planner's effective resources brought about by the epidemic, which are distributed by the planner across future generations through changes in the savings rate. Since my emphasis is on the impact of the epidemic, and not the appropriate choice of distribution across future generations, I do not pursue this further.

III. Data and Estimation Strategy

I estimate the model using the microdata files of the South African 1995 October Household Survey (OHS) and the 1998 Demographic and Health Survey (DHS). The OHS provides a variety of personal and behavioral information on each household member (e.g. age, sex, education, labour force participation, income, fertility, etc), as well as recording recent deaths in the household. Most of the data reported in this survey are consistent with other sources,⁴ but the reported retrospective fertility histories are very low (i.e. appear to involve incomplete reporting) and cannot match the historical population distribution. Thus, I use the DHS, whose primary focus is fertility, to estimate that aspect of behavior.⁵

Estimation proceeds in a simple two-step process: I first use the data of the October Household Survey to estimate incomes as a function of age, sex and education, and then use the predicted relative incomes by educational attainment as the independent variable in the household behavioral equations. Exogenous variation in individual education levels then allows me to identify the price elasticity in each demand equation. I focus only on the behavior of individuals 25 and over, whose education is taken as completed. In the data, however, both fertility and labour supply begin as early as age 12 (albeit, at extremely low levels). The retrospective fertility histories in the Demographic and Health Survey allow me to identify the behavior of fertility at earlier ages, under the assumption that each woman, when young, knew what her final educational attainment would be. For participation, I have only a single cross section, so I sim-

⁴For example, predicted employee earnings (estimated off of interval income data) are just 12 percent below national accounts compensation of employees, while the reported mortality rates provide a good match to long term (1970 to 1996) intercensal cohort survival rates.

⁵Official South African data on births, based upon voluntary registration, are woefully inaccurate. Births can be registered years after the fact. Thus, 238 thousand births were registered in 1991 as occurring in that year, but by the year 2000 the total number of births registered for 1991 had risen to 680 thousand (Statistics South Africa 2001). Yet, the 1996 census found about 960 thousand five year olds (microdata files). The primary incentives for birth registration are medical insurance requirements and access to social welfare benefits, and these do not appear to be sufficient to ensure timely or comprehensive coverage. Official mortality data face similar problems.

ply extend the quadratic age profile estimated off of adults 25 and over.

I estimate both fertility and hours of work using a Poisson count model, where the probability of observing a count value of y_i is given by:

$$(3.1) \quad P(Y_i = y_i) = \frac{e^{-\lambda_i} \lambda_i^{y_i}}{y_i!} \quad \text{where} \quad \ln \lambda_i = \mathbf{B}' \mathbf{x}_i$$

As the expected count is given by λ_i , with a quadratic in age and the ln relative wage as regressors, this produces isoelastic demand and supply functions. The Poisson, while extremely easy to estimate, has the undesirable feature that the variance of the process is equal to its mean.

There are also additional potential error correlations posed by sample design (clustering) or the use of panel data (in the fertility equation). However, as shown by Gourieroux et al (1984), provided that the mean of Y conditional on \mathbf{x}_i is equal to λ_i , maximum likelihood estimates based upon the Poisson distribution are consistent, whatever the process generating the variance of Y may be. Consequently, I use pseudo-maximum likelihood methods, estimating the coefficients using the maximum likelihood Poisson model, but calculating robust standard errors allowing for an arbitrary variance process between observations within survey primary sampling units and across multiple observations for individuals, as well as adjusting for the two-step procedures involved in using parameter estimates from the income equation.⁶ This allows me to simplify the presentation of the results. As a sensitivity test, I have also estimated the various equations using

⁶ Let \mathbf{B}_1 and \mathbf{B}_2 be the vectors of parameters estimated in the first and second step, respectively, across sample sizes n_1 and n_2 . Further, let N , the union of the two samples, be divided into M independent clusters G_1, G_2, \dots, G_M and let l_{ik} denote the log likelihood of observation k in equation i . Then, accounting for the two-step procedure and arbitrary correlation within clusters, asymptotically

$$n_1^{1/2}(\hat{\mathbf{B}}_1 - \mathbf{B}_1) \sim N(0, n_1 \mathbf{H}_{11}^{-1} \mathbf{g}_{11} \mathbf{H}_{11}^{-1}) \quad \text{and} \quad n_2^{1/2}(\hat{\mathbf{B}}_2 - \mathbf{B}_2) \sim N(0, n_2 \mathbf{A} \mathbf{B} \mathbf{A}')$$

$$\text{where} \quad \mathbf{H}_{ij} = \sum_{t=1}^{n_i} \left(\frac{\partial l_{it}}{\partial \mathbf{B}_i} \right) \left(\frac{\partial l_{it}}{\partial \mathbf{B}_j'} \right), \quad \mathbf{g}_{ij} = \sum_{m=1}^M \left(\sum_{k \in G_m} \frac{\partial l_{ik}}{\partial \mathbf{B}_i} \right) \left(\sum_{k \in G_m} \frac{\partial l_{jk}}{\partial \mathbf{B}_j'} \right), \quad \mathbf{A} = [\mathbf{H}_{22}^{-1} \quad -\mathbf{H}_{22}^{-1} \mathbf{H}_{21} \mathbf{H}_{11}^{-1}] \quad \text{and} \quad \mathbf{B} = \begin{bmatrix} \mathbf{g}_{22} & \mathbf{g}_{21} \\ \mathbf{g}_{12} & \mathbf{g}_{11} \end{bmatrix}.$$

The proof is a simple extension of the standard proof of the asymptotic distribution of the one or two step MLE using a first order Taylor expansion (e.g. Davidson and MacKinnon 1993, Murphy and Topel 1985) with the outer-product-of-the-gradient estimator of the information

constant, linear and generalized variance-mean ratio negative binomial, as well as generalized event count (Winkelmann & Zimmermann 1991), cross-section models and, where applicable, random effects poisson and negative binomial panel data models (Hausman et al 1984).⁷ The results, in almost every case, are very similar, and are reported in footnotes.

I depart from the Poisson model in the determination of children's educational attainment, where I used an ordered probit. Less educated women tend to have more children, who are in turn given relatively less human capital. In this sense, the variance of educational outcomes is quite important, as a mean preserving spread in initial educational attainment will lower average educational levels and welfare in the next generation. Furthermore, bad outcomes tend to be perpetuated from one generation to another. This becomes relevant in assessing the impact of the AIDS epidemic, where I will assume that the education of orphaned children is permanently interrupted at the time of their parents' death. I use an ordered probit model, dividing the population into a finite state space of 0 to 16 individual years of education, to keep track of the variance of educational attainment and allow poor outcomes to have disproportionate dynamic effects.

matrix (in a later draft I hope to substitute the less convenient empirical Hessian which, in finite samples, tends to produce larger standard errors).

⁷With the conditional mean of Y_i given by $E[Y_i|\mathbf{x}_i] = \exp(\mathbf{B}'\mathbf{x}_i)$, in the Poisson model the variance-mean ratio is 1. In the constant and linear variance-mean ratio negative binomial models, it is given by $1 + a$ and $1 + a E[Y_i|\mathbf{x}_i]$, respectively. In the generalized event count model it is given by $1 + a E[Y_i|\mathbf{x}_i]^k$, while in the generalized variance-mean negative binomial model it is given by $1 + \exp(\mathbf{B}_2'\mathbf{z}_i)E[Y_i|\mathbf{x}_i]$ and I use the same set of regressors to determine the variance-mean relationship ($\mathbf{z}_i = \mathbf{x}_i$). The Poisson and constant and linear variance-mean binomial models are special cases of the generalized event count model, and the generalized negative binomial model subsumes the Poisson and linear variance-mean negative binomial as a special case or limiting distribution. The random effects panel data Poisson and negative binomial models have linear and constant variance-mean ratios (respectively), and take into account each panel member's average realized value of the dependent variable. As is usually the case, the various non-Poisson models do better in predicting the distribution of the dependent variable (e.g. the number of zeros). However, as my emphasis will be on the mean, which is fairly insensitive to the choice of distribution, I opt for the simple Poisson formulation, with robust standard errors.

IV. Estimates

Table I below presents the estimation of the Beckerian elements of the model. I begin, in columns (1) and (2), by predicting the ln of before and after-tax hourly wages as a function of quadratics in age and education and a dummy variable for sex.⁸ Figures 4.1 and 4.2 graph the implied age and educational profile of after tax relative wages, and compare these with similar estimates for the United States based upon data in the 2000 Consumer Expenditure Survey.⁹ As shown, the age profile of hourly wages in South Africa is quite close to that present in the United

⁸These estimates do not adjust for selectivity bias (see Killingsworth and Heckman 1986). I have run more complicated (joint maximum likelihood) systems, augmenting the wage equations with probit selection equations, where the probability of working is determined by the variables in the wage equation plus marital status, children per adult in the household, and the total number of adults in the household (all interacted with sex, i.e. a separate participation equation by sex). First, the point estimate of the correlation between the error terms in the participation and wage equations is negative, which goes against the standard presumption in this literature (this holds even when I allow sex-based covariance between the errors in the participation equations and the joint wage equation, or when I estimate the entire wage/participation model separately by sex). Second, and more importantly, while the age and sex profiles are somewhat flattened, the educational profile of incomes, which is used to identify the price effects in all the subsequent equations, is virtually unchanged. Consequently, to keep things as simple as possible, I ignore selection issues and estimate relative labour incomes using data on workers alone, as reported above. I should note that the data are interval coded (hence the pseudo-R²) and that I restrict the sample to employees, i.e. exclude the self-employed.

⁹The before tax profiles (used later in the paper to calculate the effective labour supply and determine wages in general equilibrium) are similar, albeit somewhat steeper. As after tax incomes affect household behavior, I focus on these.

In estimating separate equations for before and after tax hourly incomes, my approach produces an unrealistic linear budget constraint by worker characteristic and, also, implies negative effective tax rates for the least educated young and old age groups. I have run a more realistic model, in which I take into account the non-linearities in the tax code, work opportunities and supply decisions (due to fixed costs, fatigue effects, etc) by modeling participation as a probit zero-one decision subject to a lump sum before and after tax income. I calculate after tax income in two ways, using the tax code and then with a separate regression (mirroring the hourly estimates of Table I). First, I find that the predicted labour supply response to a change in the after tax wage is virtually the same whether I use the tax code or the after tax regression, i.e. the after tax regressions are a close approximation to the actual tax code. Second, the coefficient estimates of the probit model predict that, at current levels, a 1 percent rise in wages will increase the supply of effective male and female labour in the South African sample by .18 and .44 percent, respectively. This is quite close to the Poisson estimates presented in Table I. Since the Poisson constant elasticity functional form allows for easy interpretation of the coefficients and calculation of the disutility of labour, and the linear before and after tax budget constraint avoids the complexity of calculating the tax burden at each level of income during the simulations, I opt for the simplified formulation presented above.

Table I: Estimation of the Beckerian Elements of the Model						
	(1)	(2)	(3)	(4)	(5)	(6)
	Before-tax ln Wages	After-tax ln Wages	Male Hours of Work	Female Hours of Work	Fertility	Children's Completed Education
Model	interval regression	interval regression	Poisson	Poisson	Poisson	ordered probit
age	.0820 (.0025)	.0666 (.0024)	.1737 (.0002)	.2125 (.0002)	.5134 (.0091)	.0354 (.0083)
age ²	-.0008 (.0000)	-.0006 (.0000)	-.0021 (.0000)	-.0026 (.0000)	-.0096 (.0002)	-.0006 (.0001)
E	.0379 (.0038)	.0341 (.0036)				
E ²	.0082 (.0002)	.0070 (.0002)				
sex	-.2190 (.0099)	-.1768 (.0095)				
Wage Index			.1710 (.0016)	.4399 (.0038)	-.3538 (.0152)	1.1382 (.0327)
year of birth					-.0031 (.0015)	.0292 (.0015)
HIV					-1.633 (.2091)	
constant	2.932 (.0520)	3.225 (.0506)	-.0791 (.0038)	-1.600 (.0054)	-1.847 (2.950)	
N	26289	25967	27253	33238	171206	10135
Pseudo R ²	.1975	.1810	.1620	.1495	.0551	.0549
Notes: Wage Index = $B_E E_i + B_{E^2} E_i^2$, with B's from column (2). E_i equals own education in columns (3)-(5) and mother's education in column (6). Age in column (6) equals mother's age at time of child's birth. The sample in column (5) is made up of a panel of 7276 individuals. Standard errors are adjusted for clustering and the two-step procedure.						

Figure 4.1: Age Profile of Ln Wages

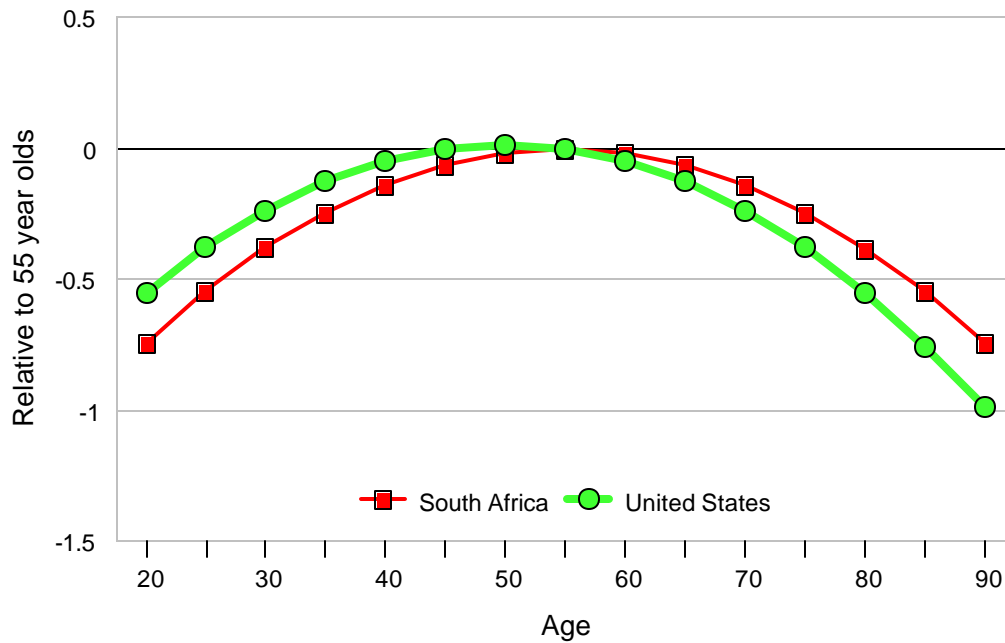
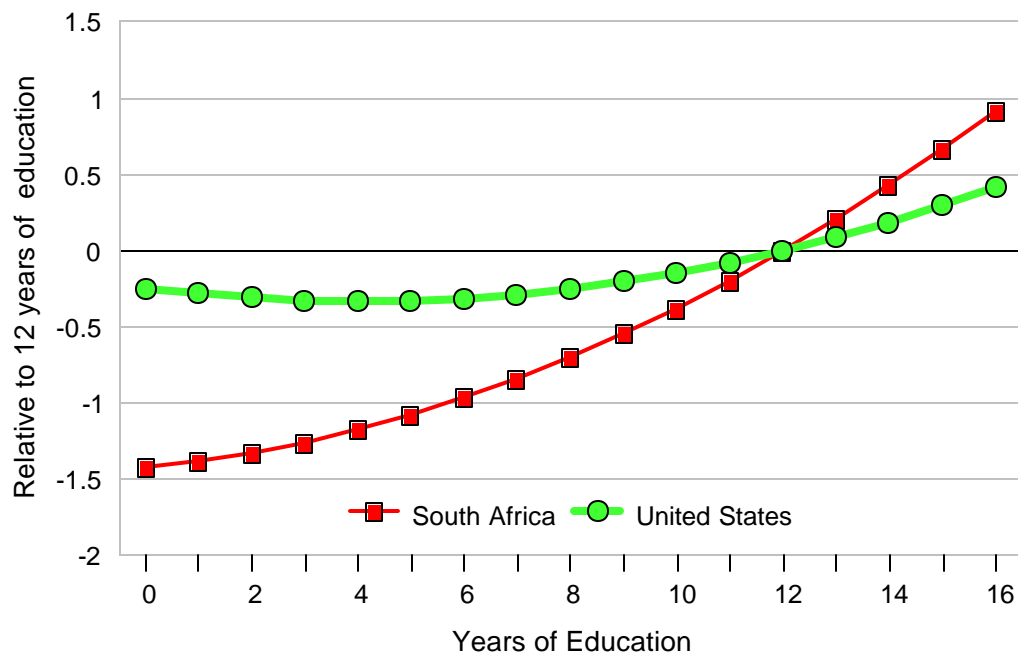


Figure 4.2: Education Profile of Ln Wages



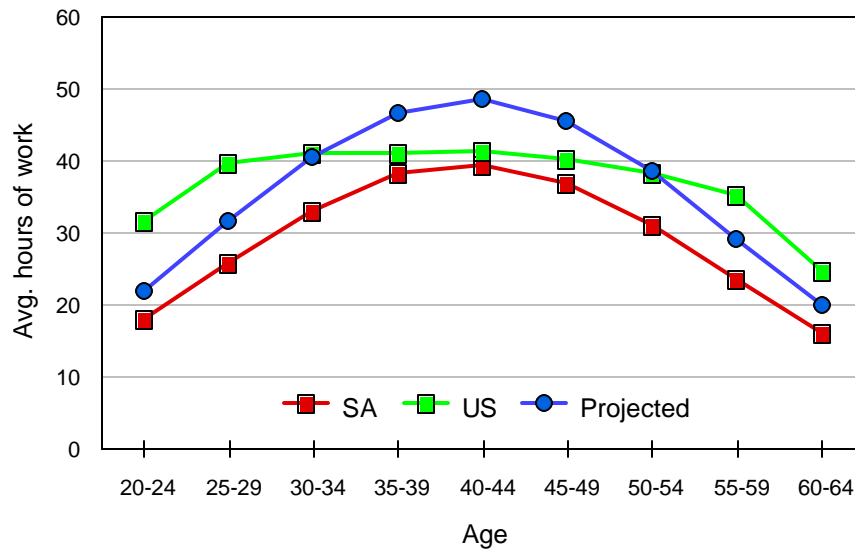
States, but the return to education is considerably greater, with a 92% premium to tertiary education and a full 234% increase in after-tax lifetime incomes associated with a movement from 0 to 16 years of education. Further, while 85% of the US sample over age 25 has at least 12 years of education, 22% of the comparable South African sample has 2 years of education or less and 24% has 12 years or more. The extraordinary dispersion in educational outcomes and average earnings by educational group present in South Africa identifies the price effects in the demand equations of the model. I use the coefficients on education in column (2) to construct an index of relative lifetime wages, Wage Index, which is used as an independent variable in the remaining columns of the table.

Columns (3) and (4) of Table I estimate the labour supply response of males and females to a change in the value of their time. As is frequently the case in the labour literature, I find that female labour supply is more elastic than men's. With every 1% increase in wages, male labour supply rises .17% and female labour supply rises .44%.¹⁰ A better sense of the implications of the coefficient estimates is provided by Figures 4.3 and 4.4, where I graph the average hours of work by age group in the prime working years in South Africa and the United States, and then predict the behavior of the United States population using the coefficient estimates of the model.¹¹ Average male hours of work between the ages of 20 and 65 are 29.3 in South Africa and 37.3 in the United States. Using US wage values, the model predicts that the United States

¹⁰ Alternative distributional assumptions yield mostly similar results. The linear and generalized variance-mean ratio negative binomial models produce male/female wage coefficients of .17/.42 and .19/.45, respectively, while the generalized event count model provides estimates of .22/.49. Only the estimates of the constant variance-mean ratio negative binomial model (.44/.77) differ substantially. However, this model is a restricted version of the generalized event count model and is strongly rejected in the data.


¹¹ I use a purchasing power exchange rate of 1.74, based upon Penn World Table 6.1 (Heston et al 2002), to convert the year 2000 US dollar wages to 1995 South African Rand. Table I uses the education slope of the wage equation to estimate the supply price response, ignoring the age profile of wages, which is subsumed in the quadratic age term. To predict US behavior, allowing for the slightly different US age profile of incomes, I reinterpret the age specific constant terms to exclude the age profile of wages, which is then captured in the wage supply response term.

Figure 4.3: Male Hours of Work



Lifetime: SA 29.3
US 37.3
Projected 36.0

Figure 4.4: Female Hours of Work



Lifetime: SA 14.7
US 28.5
Projected 24.4

male population should work an average of 36.0 hours. The model produces a profile of labour supply that is somewhat more hump-shaped than that present in the United States, reflecting patterns seen in South Africa. Part time employment opportunities may be better rewarded in a more service oriented economy, and this structural, demand consideration, might explain the higher relative participation of younger and older age groups, who have lower intrinsic labour supply, in the U.S. economy.¹² The female labour supply equation produces a similar pattern, predicting a substantial rise in participation, but not quite attaining the levels or shape seen in the US economy. While not perfect, the model does not produce obviously embarrassing results (e.g. predicting 100 hours of work at US income levels).

Column (5) of Table I uses the retrospective fertility histories of the Demographic and Health Survey to calculate the response of fertility to a change in the price of a woman's time.¹³ I include a time trend to account for trends in wages and the price of education. As discussed in the next section, a growing body of field evidence indicates that HIV has strong negative effects on fertility. I find support for this, as the historical HIV infection rate for each woman's quinquennial age group, as recorded in the maternity clinic seroprevalence surveys cited in the intro-

¹² Amongst workers aged 29 or less or 55 or more in the South African economy, 90% worked 35 hours or more. The comparable statistic for the United States sample is 67%.

¹³ The reader might worry that there is a bit of a logical inconsistency in using a count model to predict births since, as the time interval gets shorter, more than one birth event becomes biologically impossible. However, for the time frame of Table I, i.e. one year, multiple birth events may occur. The frequency of multiple pregnancies in one birth event (twins, etc.) in the South African DHS by age group is almost identical to that reported in US data (except for women over 40 where, perhaps due to fertility drugs, the US numbers are much higher than those reported in South Africa). However, 2/3 of reported multiple birth events within a 12 month period in the South African data do not involve multiple pregnancies, but rather represent children born 9 to 12 months apart. In a high fertility environment, back to back pregnancies are not that rare. I should note that the predicted Poisson probabilities of more than 2 births in a given year are negligible.

duction, has a strong negative effect on predicted fertility.¹⁴ Figure 4.5 graphs year 1995 and 2000 fertility rates in South Africa and the United States, respectively, and then projects US fertility using the coefficients of the model.¹⁵ The model comes close to matching both the life-cycle pattern and the overall level of US fertility.

I estimate the parental demand for quality in the upbringing of their children, in column (6) of Table I, by running an ordered probit of the educational attainment of adults ages 25 or more, whose education is completed, on their year of birth (to capture trends in wages and prices) and the wage index and the age at the time of birth of their mothers. Maternal income has a strong effect on realized education. As mother's educational attainment varies from 0 to 16 years, the average educational attainment of a child born in 1970 (aged 25 in 1995) to a 30 year old mother varies from 7.9 to 14.6 years.¹⁶ While persons aged 25 and over living with their mothers constitute a somewhat rarified sample in the developed world, raising issues of sample selection, this is not a problem in South Africa, where I am able to identify the mother of about 1/6 of all


¹⁴ Absent the time trend and seroprevalence rates, the coefficient on Wage Index is -.38. The three negative binomial cross section models all give exactly the same coefficient estimates as column (5) of the table, as each likelihood is maximized at the limiting Poisson distribution. The generalized event count cross section model yields coefficients of -.35, -.003, and -1.59 on the Wage Index, year of birth and the HIV seroprevalence rate, which are virtually identical to the Poisson estimates reported in the table. Similarly, the random effects poisson and negative binomial panel models both yield coefficients of -.35, -.003, and -1.65.

Since the impact of HIV on fertility plays an important role in the subsequent analysis, it is worth reporting additional sensitivity checks. Since HIV infection rates may be proxying for a non-linear time trend, I tried including higher order polynomials in time. With a quadratic, the coefficient on HIV falls to -.71, but with cubic or quartic time trends it rises to -2.93 and -2.75, respectively (in each case, remaining highly significant). Thus, a coefficient of -1.63 is not, for various formulations of the time trend, an extreme possibility.

¹⁵ I fix the time trend at 1995, which is what I use in the later simulations of the South African economy. The US data are from CDC (2002).

¹⁶ The same equation run on the wage index of fathers yields a coefficient of 1.06, with average children's attainment going from 8.1 to 14.3 years as their father's attainment varies from 0 to 16 years. I opt for mother's attainment as it is available twice as often as father's attainment. I include a quadratic in mother's age at the time of birth for consistency with the other household demand equations, but its practical impact is not large. For a woman with 12 years of education, children's education attainment ranges from a minimum of 12.1 to a maximum of 12.6 years as the mother's age at the time of birth varies.

Figure 4.5: Fertility by Age Group



Lifetime: SA 3.57
US 2.13
Projected 2.37

individuals over 25 and 1/3 of individuals aged 25 to 40. In Figures 4.6 and 4.7 I use the fertility and education estimates of Table I and the OHS female population distribution aged 62 to 99 to predict the size and educational attainment of the 1995 South African cohorts aged 25 to 50, and then compare these with the data of the OHS. As the reader can see, the fertility estimates match the population cohort size and the education estimates match historical levels of education attainment. Projection to US income levels is somewhat less meaningful in this case, given international variation in both public support of education and the human capital interpretation of a year of education. For what it's worth, individuals aged 25 had an average educational attainment of 9.5 years in South Africa in 1995 and 13.3 years in the United States in 2000. The model's prediction for the 2000 25 year old US cohort is 13.5 years.¹⁷

To summarize, the enormous variation in educational attainment and mean incomes by educational category in South Africa allows me to meaningfully estimate the impact of predictable differences in lifetime income on household decisions. The model matches historical fertility and educational distributions in South Africa and, when projected substantially out of sample, produces patterns of behavior that are not altogether inconsistent with those of the U.S.

¹⁷ This prediction is based upon the 1/5 of the U.S. sample for which I can identify mother's education. I set the cohort trend at 1970, matching the value for the 25 year old 1995 South African cohort which I use in the simulations which follow, and adjust 2000 US wages to the 1975 birth year using the 1975-2000 .9% per annum growth in real, composition adjusted, wages implied by the BLS estimates of the growth of nominal labor compensation minus the growth of effective labour input and the GDP personal consumption deflator (data from <http://www.bls.gov/web/prod3.supp.toc.htm> and Economic Report of the President 2003). Without adjustment for the trend in U.S. wages, the predicted value is 13.9 years of education. Official measures of real wage growth in South Africa (discussed below), show close to zero growth in the 1970-1995 period.

Figure 4.6: Backward Projection of Cohort Size
(based on fertility eqn)





Figure 4.7: Backward Projection of Educational Attainment
(based on fertility & education eqns)



economy.¹⁸ In the following section I calibrate the macro-economy to match historical South African data and lay out my assumptions concerning the evolution and impact of the AIDS epidemic.

¹⁸ In the interest of science, I should note that this should not be interpreted as evidence that the simple demand system I posit, with no cross-price effects, is an accurate description of reality. Since spousal education varies quite closely with own education, the responses reported in Table I really represent the response to a proportional increase in all of the labour prices faced by the household. I confirm this when I estimate cross-price effects, e.g. in the participation and fertility equations, using spousal education, in the fraction of cases where it is available, and find that the sum of the coefficients on both prices is virtually identical to the own price effects reported in Table I. The problem is that the estimates are not naturally consistent, e.g. partner's education has no effect on male labour supply but has a strong effect on female labour supply. To be consistent with utility maximization (i.e. allow calculation of indirect and direct utility functions) the demand system including cross-price effects would have to be severely restricted, reducing to something close to what I report in Table I.

In the table above I have estimated the response to changes in the overall lifetime labour prices faced by households. Since spousal education moves closely with own education, and since I shall be exploring simulations in which the overall level of wages moves up and down for decades at a time, the coefficients reported in Table I arguably provide a good approximation of likely behavior and a proper metric for calculating indirect utility functions and compensating variations, even if their interpretation as pure, contemporaneous, own price effects is something of a fiction.

V. Calibration

I calibrate the macroeconomy using standard aggregate data. The investment to GDP ratio in the South African economy has fluctuated, but shows no appreciable trend (figure 5.1). I fix the Solovian savings rate at 17.5 percent, the average investment to GDP ratio in the postwar period. I initialize the capital stock using the Federal Reserve Bank of South Africa's estimate of the 1995 capital stock and set the depreciation rate at .06, the reported ratio of capital consumption to capital for that year. I set the labour share at .62, the average ratio of compensation of employees¹⁹ to value added at factor cost in the period 1960-2000, and assume a Cobb-Douglas production function in capital and effective labour:

$$(5.1) \quad Y = AK^\alpha EL^{1-\alpha}, \quad \text{where} \quad EL = \sum_i W_i L_i.$$

Effective labour is given by the sum of labour supplies differentiated by sex, age and education, with fixed relative weights given by my estimates, in column (1) of Table I, of before-tax wages in 1995. To focus on the role played by changes in factor supplies, I assume zero future total factor productivity growth.²⁰

I will consider two scenarios for the intertemporal constraints faced by the macroeconomy. In my baseline formulation, I assume a closed-economy, so that the evolution of incomes is constrained by the unit elasticities of the Cobb-Douglas production function. As an alternative, I also consider the case of perfect capital mobility, so that both the capital rental and the wage per

¹⁹ Adjusted, using census data, for the implicit labour income of the self-employed.

²⁰ Historical estimates of total factor productivity growth are prohibited by the poor quality of South African census data (discussed below), whose coverage is inconsistent and where changes in the questionnaire design have induced massive changes in measured participation. For what its worth, the official index of real remuneration per worker grows .4% per annum between 1970-1996 (<http://www.reservebank.co.za/>). I have found it impossible to construct a believable index of worker quality for this period, but this measure suggests close to zero growth in effective wages per worker and, by implication, total factor productivity.

Figure 5.1: Investment over GDP (1995 prices)





Figure 5.2: HIV Prevalence by Age Group



unit of effective labour are fixed at their initial values. In this case, despite the curvature of the domestic production function, the economy's income is of the linear form

$$(5.2) \quad Y = \alpha_K K + \alpha_{EL} EL,$$

where α_K and α_{EL} are the fixed marginal products. This framework excludes most Beckerian effects, as the aggregate wage is fixed.²¹ It allows insight, however, into how the ability to substitute intertemporally at fixed rates, i.e. to move outside the constraints imposed by the curvature of domestic production, influences the costs of the AIDS epidemic.

Turning to the HIV epidemic, the South African maternity clinic seroprevalence data presented at the beginning of this paper suggest that overall infection rates had begun to plateau by the year 2000. Across age groups, infection rates peak at ages 25-29 (figure 5.2), and decline thereafter. I assume that new cohorts will experience lifecycle infection profiles similar to those of the year 2000 cross-section of cohorts under age 30, but if not infected by age 30, will face negligible risk thereafter. As such, I take the current infection rates of persons over 30 not as random risk, but rather as reflecting the infection of women who, were they 25-29 in 2000, would form part of the current cohort of infected women. It seems reasonable to assume that there is a segment of the population whose interaction with other individuals inevitably, but perhaps unknowingly, leads to infection and that, given the extraordinarily high infection rates already present in the population, the size of this group is given by the maximal infection rate at age 30.²² I assume that, with medical breakthroughs or behavioral change, post-2000 cohort

²¹ There are still changes in behavior brought about by changes in the distribution of educational attainment and associated incomes.

²² In this, I follow the WHO-UNAIDS modelling recommendations (UNAIDS Reference Group 2002). The UNAIDS Group suggests using aggregate infection rates to estimate the overall "at risk" group and communicability parameters, using an ad hoc procedure to initialize the epidemic. As I need historical age-specific infection rates, I opt for simply using the existing seroprevalence data, with the "at risk" proportion set by the maximal infection rate in 2000 and the lifecycle pattern of infection given by the under-30 cross section. I smooth the historical age specific seroprevalence data by running the ln of the infection to non-infection rate (a logit means model) on polynomials in age and year, weighting the observations by their sample size.

infection rates decline, sinusoidally, to zero over 50 years. I also assume that male and female infection rates are evenly matched.²³ Where available, seroprevalance data indicate fairly equal infection rates across educational groups (figure 5.3), which I shall assume as well.

Regarding mortality, adult individuals infected with HIV typically experience a short period of flu-like symptoms, after which they remain, superficially, asymptomatic, until their immune system collapses and quickly succumbs to opportunistic infections and cancers.²⁴ I use the WHO-UNAIDS (UNAIDS Reference Group 2002) recommendations on adult survival times absent retroviral therapy (figure 5.4), which suggest a median survival after infection of 9 years.²⁵ While mother-to-child transmission can be reduced through antiretroviral therapy, cesarean delivery and avoidance of breastfeeding, these are all costly and currently about 1/3 of children born of HIV positive mothers in African countries, and South Africa in particular, are infected,²⁶ a proportion that I assume will remain constant for the foreseeable future. Pediatric AIDS progresses horribly rapidly, with 1/3 of HIV infected South African children dying from the virus within one year of birth (Bobat et al 1999). I use the recommended children's mortality profile of the UNAIDS Group (figure 5.4), which predicts that all infected children die by age 12.

²³ African studies of couples consistently find that about 30% to 40% of male partners of HIV infected females and 30% to 40% of female partners of HIV infected males are themselves HIV negative (Carpenter et al 1999, Gray et al 1998, Hira et al 1990, and Serwadda et al 1995). This suggests relatively equal infection rates once the virus spreads to the general community. As the best data on infection rates come from antenatal clinic attendees, knowledge of infection rates among men is fairly limited.

²⁴ <http://www.avert.org/hivstages.htm> (11/3/2002).

²⁵ The UNAIDS Group estimates that retroviral therapy would add about three years to median survival times. In November 2003 the South African government announced it would be distributing free retroviral drugs, but as of March 2004 only 1500 people had received such support, the government had only begun to ask for bids from pharmaceutical companies, and the health minister was extolling the benefits of consuming olive oil, lemon juice and garlic (Goering 2004).

²⁶ <http://www.ama-assn.org/special/hiv/newsline/briefing/mother.htm> (11/3/2002), <http://www.avert.org/motherchild.htm> (11/3/2002), Unicef (2002) and Bobat et al (1999).

Figure 5.3: HIV Prevalence by Educational Level


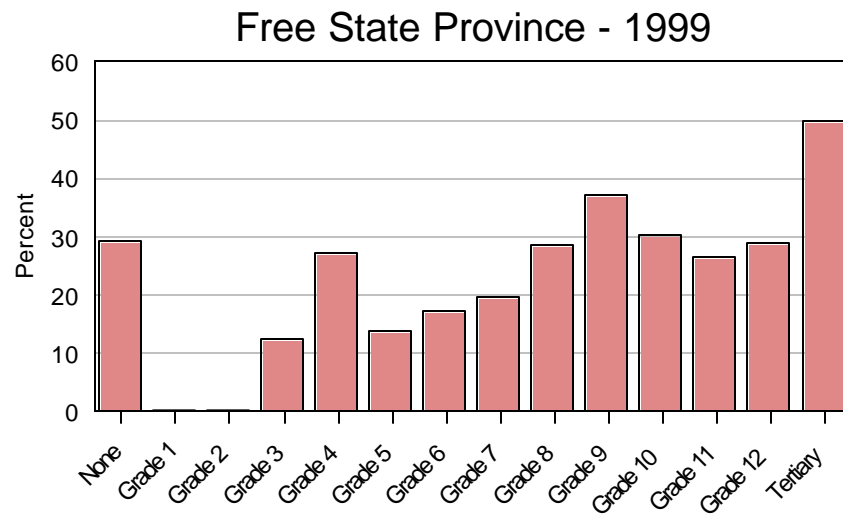



Figure 5.4: Cumulative Survival Rates



Turning to Beckerian behavior, as HIV infected individuals are generally unaware of their status and remain superficially healthy for long periods of time, I assume that adult economic participation and children's education proceeds normally until the final year of each adult's life. To incorporate the impact of orphanhood I assume that the education of the surviving children of adults who die of AIDS is interrupted at the moment of their parents' (assumed joint) death. Thus, children borne of parents with HIV will get minimal education, as their parents generally die before they reach 9 years of age, while children whose parents contract HIV after their birth will terminate their education if and when their parents die before they reach 25 years of age.²⁷ This almost certainly overstates the impact of the HIV epidemic. There is considerable discordance in the HIV status of couples (see footnote above), so that many orphaned children will have a surviving parent, while others may receive some support from relatives.²⁸ Nevertheless, as this is clearly a mechanism through which the AIDS epidemic lowers the welfare of future generations, I take this as my baseline assumption. As an alternative, I assume that orphanhood has no impact, whatsoever, on the realized education of children. Reality, presumably, lies somewhere between these two extremes.

A growing body of evidence indicates that HIV has strong negative effects on the fertility of women who are otherwise mostly asymptomatic and unaware of their infection status. Gray et

²⁷ To incorporate this effect, I use the OHS to estimate an additional ordered probit educational equation, representing attainment at ages 6-24 as a function of a quadratic term in age and the mother's wage index (interacted with the quadratic in age). Individuals who reach age 25 without their parents dying from AIDS are endowed with their full human capital, as determined by column (6) in Table I earlier.

²⁸ Running a probit of current school attendance on a quadratic in age, district dummies and the survival status of each parent for individuals 6 to 24, I find that the death of both parents reduces the average annual attendance rate from .77 to .70. South African data on attendance are not synonymous with achievement, as there is frequent repetition of grades (see Strauss 1999). As an alternative measure, I run an ordered probit of educational attainment on age, district dummies and the survival status of each parent for individuals 15-19 and 20-24 in the OHS, finding that the death of both parents reduces educational attainment an average of .6 years (from 8.6 with dual survival) for the 15-19 age group and .9 years (from 9.9 with dual survival) for the 20-24 age group. Both procedures suggest that about 10% of cumulated years of education are lost with the death of both parents.

al (1998), in a cross-sectional analysis of a Ugandan community, find that after adjustment for factors such as age, marital status, contraception, and frequency of intercourse, HIV infected women have a pregnancy odds-ratio equal to .45 that of HIV negative women. Carpenter et al (1997) and Hunter et al (2003) find similar effects on the fertility of women over the age of 20 in cohort studies in Uganda and Tanzania, respectively.²⁹ Birth interval studies on pregnant women at antenatal clinics (Desgrées du Lou et al 1999, Glynn et al 2000) find that HIV infected women have .7 to .8 times the pregnancy risk of uninfected women, which, given that these studies focus on successfully pregnant women, almost certainly understates the impact of HIV infection on fertility.³⁰ Comparisons of antenatal clinic attendees with women of the surrounding community find community infection rates that are frequently 50% higher than those present among the clinic attendees, at least for women over the age of 20.^{31,32}

²⁹ Both studies report average age adjusted fertility odds ratios of about .7, but if one examines the tables and reported data one sees that the odds ratios for women over 20 are on the order of .5. Endogeneity tends to reduce the estimated impact of HIV on fertility, as unprotected sex with multiple partners leads to both pregnancy and HIV infection. Risky sexual activity is probably a greater proportion of total sexual activity in females aged 15 to 19, who make up a large fraction of the sample in both studies and show HIV adjusted fertility odds ratios of greater than one. Gray et al (1998) is the only study to include a wide variety of controls (e.g. frequency of intercourse, number of partners, marital status, and contraception) for the frequency and type of intercourse.

³⁰ It is worth reemphasizing that the women in these studies generally did not know their HIV status (and, in fact, only a small fraction of the subjects request their test results when these are made available to the population studied). Physiologically, the current literature focuses on early term miscarriages and amenorrhea (as the disease progresses) as partial explanations (Gray et al 1998).

³¹ Kigadye et al (1993) find about 50% higher infection rates in all age groups in Mwanza Municipality, Tanzania. Chagalucha et al (2002) find infection rates 1/3 to 1/2 higher for women over 20 in Mwanza, but claim that adjustment for parity (number of births) eliminates the difference (which makes little sense and, in any case, is not borne out by the data in their Table 1). Glynn et al (2001) examine three cities in Cameroon, Kenya and Zambia and find small differences in two locales and large differences in one. All studies are hampered by high non-response rates (15 to 35 percent) in the general population and sampling issues in the choice of clinics.

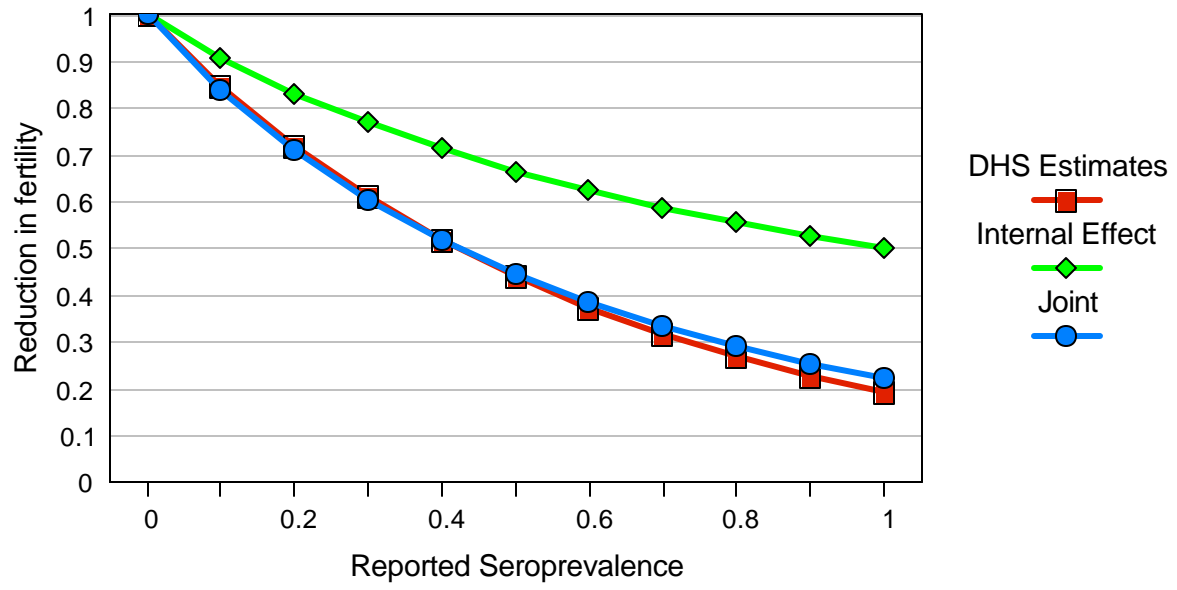
³² Another intriguing piece of evidence on this issue is the fact that, as noted by UNAIDS (2002), "projected maternal and dual orphan numbers...are consistently 40-110% higher than is found by the DHS [demographic and health survey]."

Figure 5.5 below graphs my estimates, from Table I earlier, of the impact of historical HIV seroprevalence rates on average fertility in the South African 1998 Demographic and Health Survey sample. In the same figure I consider what would be observed if HIV lowered the fertility of infected women by 50%, but did not alter the fertility of uninfected women.³³ As the reader can see, this curve cannot match the extraordinary effects present in the DHS data. I interpret the difference between these two curves as representing an external effect, i.e. a reduction in the demand for fertility (or unprotected sexual activity) on the part of all households. In my simulations, I parse Table I's estimated coefficient of -1.6 on seroprevalence rates into an individual 50% reduction in fertility and a community wide -.8 effect (in the Poisson equation) on the demand for fertility. As shown in figure 5.5, this combination matches the estimated profile.³⁴ The decline in fertility associated with the HIV epidemic plays an important role, raising the capital-labour ratio faced by future cohorts and offsetting the reductions in human capital brought about by premature parental death.

³³ If h is the true seroprevalence rate, then the seroprevalence rate observed in antenatal clinics (on the horizontal axis in the figure) will be given by $s = .5h / (.5h + 1-h)$, while the average realized fertility rate, as a fraction of fertility absent HIV, will be $.5h + (1-h) = .5 / (.5(1-s) + s)$.

³⁴ I have also run an alternative simulation, in which I assume that there is no individual reduction and only the external effect estimated in Table I earlier. The results, concerning equilibrium wages, cohort educational levels, output per capita and compensating variations are quite similar. Assuming that HIV infected women have lower individual fertility reduces the predicted number of AIDS orphans per infected woman, but raises the estimated underlying infection rate (adjustment of the seroprevalence data), increasing the number of orphans and adult deaths. These effects are largely offsetting.

Figure 5.5: HIV & Fertility



VI. Simulation

In the pages that follow I describe the evolution of the South African economy under a variety of circumstances. With regards to the HIV epidemic, I consider five scenarios: (1) "No HIV" - the path taken absent the epidemic; (2) "HIV" - the economy with the HIV epidemic; (3) "No Becker" - the economy with the epidemic, but with no endogenous response to changes in wages, i.e. all education, fertility, and participation decisions by educational class kept at their values along the No HIV path;³⁵ (4) "No Fertility" - the No-Becker scenario, but with the added dimension that HIV does not have any of the individual or community wide negative effects on fertility estimated and described in sections IV and V above; and (5) "Full Education" - the HIV path, but without the assumption that children's education is interrupted at the time of their parents' death. At the macroeconomic level, I will consider both the closed and open economy. Before delving into economic calculations, it is worth recalling the human cost of the epidemic which, as shown in Figure 6.1, could cumulatively claim 25 millions lives (in the baseline scenario). What follows is a statement about economic conditions and the welfare of survivors, not about the desirability of the wholesale slaughter of innocents.

I begin by describing the development of the closed economy. Figure 6.2 below depicts the evolution of the wage per unit of effective labour. Absent the HIV epidemic, the wage initially declines, as better educated young cohorts put pressure on the ratio of capital to effective workers, but eventually begins to climb, as higher levels of education lead to lower fertility and population growth. This pattern is reversed by the epidemic, as high mortality in the first decades of the 21st century drives up the wage, which then comes back down as the epidemic wanes and the labour force is built back up. In the baseline HIV scenario, the wage rises to 15 percent above the No HIV path by 2040. Although most of the mortality associated with the

³⁵ This simulation still contains Beckerian effects in that changes in the distribution of educational attainment lead to changes in average behavior and the behavior by educational class responds to the value of the real wage along the No HIV path.

Figure 6.1: Cumulative AIDS Deaths




Figure 6.2: Wage per Unit of Effective Labour (1995 = 1)





Figure 6.3: Educational Attainment by Birth Year



epidemic is resolved by 2040 (figure 6.1), it takes some 60 years for the wage to return to the No HIV transition path, as the reduced fertility brought about by the changed population structure and higher wages sustains a higher capital labour ratio. In the No Becker scenario, the failure of fertility to respond to temporarily higher wages results in a more rapid decline in the wage, which ultimately falls below the levels achieved along the No HIV path. When, in addition, the HIV epidemic has no impact on fertility, the wage falls well below the No HIV transition path. Finally, in the Full Education scenario, the high wages brought about by the epidemic are completely self sustaining. High wages lead to lower fertility and better educated children, which in turn sustain higher wages.

The impact of the epidemic on the accumulation of human capital is summarized in Figure 6.3, which depicts average educational attainment by cohort birth year. In the HIV, No Becker and No Fertility scenarios, I assume that the orphaning of children leads to an immediate, and permanent, interruption of their education. As shown in the figure, this lowers the educational attainment of the 1995 cohort by 2 years and substantially reduces the human capital of cohorts for the next 100 years.³⁶ Uneducated individuals have higher fertility and less educated children, so the losses from orphaning early in the epidemic persist for generations. The lower individual fertility of HIV infected women, and the positive effect of higher wages on the demand for quality in children, work against this effect. Consequently, the reductions in educational attainment are smaller in the HIV scenario than they are in the No Becker scenario, where there is no endogenous response of fertility or education demand to the higher temporary wages, and substantially better than in the No Fertility scenario, where HIV does not reduce individual fertility.

³⁶ About 1/3 of all HIV free (i.e. surviving) children born in 1995 are orphaned before completing their education, with the ratio falling to 11% and 4% for the cohorts born in 2020 and 2045, respectively. The education acquired by orphaned children depends upon how long their parents survive. Children born in 1995 who are eventually orphaned get about 1/2 of the years of education of non-orphaned children, as the epidemic in 1995 was still spreading and much of the adult population was as yet uninfected. Orphaned children born after the year 2000, when the epidemic is established and on-going, get only about 1/4 of the education of non-orphaned children.

The loss of human capital explains why the wage along the No Becker and No Fertility paths, shown earlier in Figure 6.2, falls below the original, No HIV, transition. Lower educational attainment, for a given real wage per effective worker, leads to higher average community fertility, which drives the real wage down, perpetuating low levels of education and further high fertility. In contrast to the above, in the Full Education scenario I assume that the death of parents does not interrupt their children's education. In this case, the rise in the wage (figure 6.2), leads to higher demand for children's education, which in turn supports continued lower fertility, sustaining both wages and the level of educational attainment.

Figure 6.4 graphs the GDP per capita sustained along the various epidemic paths relative to the levels sustained along the No HIV transition to the steady state.³⁷ In the baseline HIV simulation, the initial rise in the wage per effective worker more than offsets the losses in cohort educational attainment, so that output per capita remains, for some 40 years, above the No HIV path. Eventually, however, the reduced educational attainment becomes dominant and output per capita falls slightly below the No HIV path. In the No Becker scenario, neither participation, fertility nor children's education respond to the high temporary wages, so output per capita rises by less during the wage boom and then falls considerably below the No HIV path, as the economy is burdened by large uneducated cohorts. Conditions are worse in the No Fertility case, where, despite the wage boom, rising dependency ratios, brought about by adult mortality combined with continued fertility, drag the economy down from the very beginning. In this case, high fertility in an uneducated population ultimately leads to an output per capita only 2/3 of

³⁷ Absent the epidemic, GDP per capita rises .8% per annum for the first century, in the gradual transition to a steady state 3.5 times greater than its value in 1995. As the real wage per effective worker rises only 10% in the transition to the steady state, almost all of this growth is driven by the accumulation of labour, and not capital. The transition is driven by increases in educational attainment, as better educated cohorts raise better and better educated offspring. Higher levels of education increase effective labour per capita directly, through productivity, and indirectly, by raising age specific participation rates and lowering fertility, which in turn shifts the population structure in favour of age groups with higher participation. The steady state involves an average adult educational attainment of 12.9 years, well above the cross sectional average of 8.5 in 1995.

Figure 6.4: GDP per Capita (relative to No HIV)





Figure 6.5: GNP per Capita with Capital Mobility (relative to No HIV)



what would have been experienced absent the epidemic. At the other extreme, in the Full Education scenario, with no adverse effects on children's education, high wages and low fertility lead to higher output per capita everywhere along the path. This scenario is decidedly unrealistic, but so is the baseline HIV scenario, which both overstates the frequency of dual parent orphaning and its likely impact on children's education.³⁸

The various epidemic scenarios for the economy with perfect capital mobility are easily summarized. With the wage per effective unit of labour fixed at its initial value, there is no wage boom, no endogenous response in behavior, and the HIV and No Becker scenarios are identical.³⁹ However, losses of educational capital (similar to those shown earlier in figure 6.3) are offset by higher capital per capita, i.e. incomes. Figure 6.5 summarizes the resulting evolution of GNP per capita relative to the No HIV scenario. These paths are similar to those described earlier for the closed economy, with the exception of the, not surprising, result that the baseline HIV and Full Education scenarios are somewhat worse than in the closed economy case, as there are no positive behavioral responses associated with changes in wages. The true significance of the open economy lies in its implications for the cost of responding to the epidemic, as will be seen further below.

As noted earlier in section II, for a given savings rate and factor shares, GDP or GNP per capita is proportional to the flow of utility per capita in the model economy,⁴⁰ so that figures 6.4

³⁸ As noted earlier in section V, there is considerable discordance in the HIV status of married couples, so that many children would have a surviving parent. More significantly, the death of even both parents only reduces educational capital by about 10% in the OHS sample while, in the simulations above, orphaned children generally average about 25% of the years of education of non-orphaned children.

³⁹ Although Beckerian elements remain in all the HIV and non-HIV simulations as income differences across educational groups give rise to differences in behavior.

⁴⁰ Obviously factor shares are not constant in the open economy and small changes in the factor shares of male and female labour occur along the closed economy equilibrium path, but the figures provide a good approximation. In the calculations which follow, I don't use these measures, but instead explicitly calculate the flow of utility from consumption and the disutility of male and female labour, as in equation (2.7) earlier.

and 6.5 provide some insight into changes in the flow value of utility per capita. One can aggregate these flow values into a summary statistic by asking the following question: What addition to the capital stock would the economy need, in 1995, to be able to deliver, along the entire path, the flow of utility per capita it would have enjoyed absent the HIV epidemic? In this calculation, one adjusts the capital stock in 1995 and then modifies savings year by year to deliver the No HIV flow of utility,⁴¹ subject to the transversality condition that the savings rate eventually return to its Solovian value. This measure is, in essence, a compensating variation and summarizes whether, from the point of view of maintaining the flow of per capita welfare, the epidemic adds or removes resources from the economy.

Table II below provides the compensating variations, as a percent of the initial capital stock, necessary to deliver the No HIV utility path under the various epidemic and macroeconomic scenarios. Starting with the Full Education scenario for the closed economy, we see that the economy could give up 15.5 percent of its capital stock in 1995 and still be able to deliver the same per capita flow of utility, in each period, as would have been available along the No HIV path. This is not surprising, given that this path yields a higher flow of per capita utility at all points (figure 6.4 earlier). As noted earlier, the baseline HIV and No Becker scenarios yield higher output per capita initially, but lower levels in later years. The compensating variations for both scenarios are negative, i.e. the economy could give up capital initially, as the mechanisms leading to the early boom in output per capita described above allow the economy to offset any initial compensating loss (removal of capital) and still accumulate enough capital to sustain welfare in later years.

In the case of the No Fertility scenario, an extraordinary compensation, equal to 25 times the initial capital stock, is necessary to carry the economy through the epidemic. This stems from two elements. First, output per capita begins to decline immediately (figure 6.4 earlier), so

⁴¹ Taking into account the behavioral response to the wage and the disutility of labour along the compensated epidemic paths.

Table II: Compensating Variation (as percent of original capital)		
	Closed Economy	Perfect Capital Mobility
Full Education	-15.5	-22.3
HIV	-12.4	-14.4
No Becker	-10.3	↑
No Fert	2474.4	9.0
Retroviral Drugs	-5.8	-4.4

there is no opportunity to accumulate resources to offset future difficulties. Second, the trough of output per capita is extremely deep, and is reached decades into the process, so that, given the concavity of the production function, a great deal of capital, to be invested and husbanded at low rates of return, is needed at the beginning of the epidemic. Were the South African economy to go through this type of epidemic on its own, enormous up front economic sacrifices would be needed to reduce the cost to future generations. The second column of the table shows the gains to the economy from being able to step outside of the concavity of its production function. If South Africa could borrow and lend freely at the initial marginal product of capital,⁴² it would need only 9% additional capital in 1995 to carry it through the depths of the No Fertility scenario. The compensating variations for the other scenarios are more negative, as the economy can give up more initial capital without increasingly adverse (i.e. rising marginal product) consequences.

With the exception of the No Fertility scenario, the calculations in Table II show that the HIV epidemic actually endows the South African economy with additional resources, allowing it to maintain the future flow of per capita utilities and still have resources available for other uses. One obvious application of such resources would be the support of the afflicted. In the bottom row of Table II I consider the compensating variation necessary for the South African economy to maintain, under the baseline HIV scenario, the No HIV flow of per capita welfare *and* provide retroviral drugs to all AIDS patients. In this calculation, I follow UNAIDS (2002) and assume that retroviral drugs will extend the life of terminal AIDS patients by three years, i.e. each patient receives three years of treatment. I assume that these persons do not contribute in any way to economic activity, but count in the total population numbers so that, in addition to retroviral

⁴² The initial 1995 net (of depreciation) domestic marginal product of capital, i.e. the real return, is .117.

drugs, they receive the average per capita flow of consumption.⁴³ The number of persons receiving retroviral therapy peaks at around 6% of the population 2008 and remains above 1% until 2050. As shown in the table, despite the considerable expenditure on drug therapy and the support of the invalided, the compensating variation is still negative. The AIDS epidemic endows the South African economy with enough resources to care for the victims of the epidemic and still have means left over to raise the standard of living of future generations.⁴⁴

⁴³ I use a figure of \$268 (2004 dollars) per patient per year, the average generic cost of combination retroviral therapy reported in McNeil (2004), to price the drugs.

⁴⁴ I should note that recent South African data strongly support the arguments presented in the simulations above. The official index of real remuneration per non-agricultural worker (<http://www.reservebank.co.za/>) stagnates from 1970 to 1995, moving from 92.6 to 100, but then surges upward, reaching 115 by 2000. The estimated ratio of 0-4 year olds to 15-19 year olds in the 1991 census was 1.25. By the time of the 1996 census this ratio had fallen to 1.06, with a further decline to .89 in the 2001 census. Thus, aggregate data indicate a sharp rise in real wages and incredibly rapid decline in fertility, as posited in the simulations above.


However, these data are easily dismissed. The movement in real wages is consistent with the 8.6% growth in real aggregate compensation of employees (nominal compensation deflated by the household consumption deflator) in 1995-2000, if one believes the reported 6% and 12% declines in public and private sector employment, respectively. However, these movements are too soon and too large to be associated with the HIV epidemic, and are more likely business cycle phenomena. As for the census data, they are corrected for "the undercount", i.e. the failure to survey all individuals. The 1991 census was corrected using a population extrapolation of the 1970 census and estimated fertility and mortality rates (based on extremely poor data), as well as estimates of the TBVC territories. The 1996 census was heralded as the most accurate census ever, as it used resampling to estimate the undercount, an approach which was also used in the 2001 census. In principle, this should inspire confidence in the numbers. In practice, one can see that cohorts aged 0 to 14, when suitably aged, "grow" about 7% to 9% from the 1996 to 2001 census. Thus, despite the resampling, the correction for the undercount amongst younger age groups in the 1996 census was inaccurate, as may also be the case in the 2001 census. Given this, one cannot use the census estimates of the size of young cohorts to draw any sensible conclusions regarding recent fertility.

VII. Conclusion

Figure 7.1 below graphs my projections of the South African population with and without the HIV epidemic. In the absence of the epidemic, the population would have grown rapidly, reaching 150 million persons by 2070. With the epidemic, the combination of adult mortality, HIV infant mortality, the powerful effects of HIV infection on individual and communal fertility, and the endogenous response of fertility to higher wages contrive to keep the population below 50 million for the next 70 years. As shown in this paper, the positive effects of lower population growth are strong enough to counteract the most pessimistic forecasts of the human capital losses of AIDS orphaned children, implicitly endowing the economy with extra resources which can be used to extend the lifespan of the afflicted and still leave reserves to raise the per capita welfare of future generations.

An emphasis on per capita outcomes is a particularly calloused notion of "welfare", one that does not care about the number of individuals, dead or unborn. It is by no means the way I would evaluate the agonized deaths of millions of young adults and children invaded by infections and cancers. There is, however, a long standing emphasis on GDP per capita in the economics profession as a measure of economic well-being. In the classroom, this has led to the Solovian based lessons on the merits of lower population growth taught to generations of graduate students, while, in the real world, it has supported the encouragement of fertility control through subsidies and, in some cases, draconian coercive measures. One cannot endlessly lament the scourge of high population growth in the developing world and then conclude that a reversal of such processes is an equal economic disaster. The AIDS epidemic is a humanitarian disaster of millennial proportions, one that cries for assistance. It is not, however, an economic disaster.

Figure 7.1: Population



VIII. Bibliography

- Bobat, Raziya, Hoosen Coovadia, Dhayendre Moodley, et al. "Mortality in a cohort of children from to HIV-1 infected women from Durban, South Africa." South African Medical Journal 89, No. 6 (1999): 646-648.
- Carpenter, Lucy M., Anatoli Kamali, Anthony Roberantwari, et al. "Rates of HIV-1 transmission within marriage in rural Uganda in relation to the HIV sero-status of the partners." AIDS 13 (1999): 1083-1089.
- Carpenter, Lucy M., Jessica S. Nakiyingi, Anthony Ruberantwari, et al. "Estimates of the impact of HIV infection on fertility in a rural Ugandan population cohort." Health Transition Review 7, Supplement 2 (1997): 113-126.
- Centers for Disease Control and Prevention. "Births: Final Data for 2000." National Vital Statistics Report 50, No. 5 (February 12, 2002).
- Changalucha, John, Heiner Grosskurth, Wambura Mwita, et al. "Comparison of HIV prevalences in community-based and antenatal clinic surveys in rural Mwanza, Tanzania." AIDS 16 (2002): 661-665.
- Davidson, Russell and James G. MacKinnon. Estimation and Inference in Econometrics. New York: Oxford University Press, 1993.
- Deaton, Angus and John Muellbauer. Economics and Consumer Behavior. Cambridge: Cambridge University Press, 1980.
- Desgrées du Lou, Annabel, Philippe Msellati, Angèle Yao, et al. "Impaired fertility in HIV-1-infected pregnant women: a clinic-based survey in Abidjan, Cote d'Ivoire, 1997." AIDS 13 (1999): 517-521.
- Epstein, Larry G. "Generalized Duality and Integrability." Econometrica 49, No. 3 (May 1981): 655-678.
- Glynn, Judith R., Anne Buvé, Michel Caraël, et al. "Decreased Fertility Among HIV-1-Infected Women Attending Antenatal Clinics in Three African Cities." Journal of Acquired Immune Deficiency Syndrome 25 (2000): 345-352.
- Glynn, Judith R., Anne Buvé, Michel Caraël, et al. "Factors influencing the difference in HIV prevalence between antenatal clinic and general population in sub-Saharan Africa." AIDS 15 (2001): 1717-1725.
- Goering, Laurie. "Official AIDS remedy: Garlic, etc." Chicago Tribune. 15 March 2004. Section 1, p. 4.
- Goldberg, P.J.P. Women, Work, and Life Cycle in a Medieval Economy: Women in York and Yorkshire. New York: Oxford University Press, 1992.
- Gourieroux, C., A. Monfort and A. Trognon. "Pseudo Maximum Likelihood Methods: Theory." Econometrica 52, No. 3 (May 1984): 681-700.

- Gray, Ronald H., Maria J. Wawer, David Serwadda, et al. "Population-based study of fertility of women with HIV-1 infection in Uganda." Lancet 351 (1998): 98-103.
- Hatcher, John. Plague, Population and the English Economy 1348-1530. London: Macmillan Press, 1977.
- Hausman, Jerry, Bronwyn H. Hall and Zvi Griliches. "Econometric Models for Count Data with an Application to the Patents R&D Relationship." Econometrica 52, No. 4 (July 1984): 909-938.
- Heston, Alan, Robert Summers and Bettina Aten. Penn World Table Version 6.1. Center for International Comparisons at the University of Pennsylvania (CICUP), October 2002. <http://pwt.econ.upenn.edu>.
- Hira, S.K., B.M. Nkowane, J. Kamanga, et al. "Epidemiology of Human Immunodeficiency Virus in Families in Lusaka, Zambia." Journal of Acquired Immune Deficiency Syndromes 3 (1990): 83-86.
- Hunter, Susan-Catherine, Raphael Isingo, J. Ties Boerma, et al. "The Association Between HIV and Fertility in a Cohort Study in Rural Tanzania." Journal of Biosocial Science 35 (2003): 189-199.
- Kigadye, Rose-Mary, Arnoud Klokke, Angus Nicoll, et al. "Sentinal surveillance for HIV-1 among pregnant women in a developing country: 3 years' experience and comparison with a population serosurvey." AIDS 7 (1993): 849-855.
- Killingsworth, Mark R. and Heckman, James J. "Female Labor Supply: A Survey." In Handbook of Labor Economics, Vol. I, ed. Orley Ashenfelter and Richard Layard, 103-204. Amsterdam: North-Holland, 1986.
- McNeil, Donald G. Jr. "Plan to Battle AIDS Worldwide Is Falling Short." New York Times. 28 March 2004.
- Murphy, Kevin M. and Robert H. Topel. "Estimation and Inference in Two-Step Econometric Models." Journal of Business and Economic Statistics 3, No. 4 (1985): 370-379.
- Phelps-Brown, E.H. and Sheila V. Hopkins. "Seven Centuries of the Prices of Consumables, Compared with Builders' Wage-Rates." Economica 23, No. 92 (1956): 296-314.
- Poos, L.R. A rural society after the Black Death: Essex 1350-1525. Cambridge: Cambridge University Press, 1991.
- Saltmarsh, John. "Plague and Economic Decline in England in the Later Middle Ages." Cambridge Historical Journal 7 (1941): 23-41.
- Serwadda, David, Ronald H. Gray, Maria J. Wawer, et al. "The social dynamics of HIV transmission as reflected through discordant couples in rural Uganda." AIDS 9 (1995): 745-750.
- Statistics South Africa. Recorded live births 1998-2000. Pretoria: Statistics South Africa, 2001.

Strauss, J.P. Results of the Monitoring Learning Achievement Project. Bloemfontein, South Africa: Research Institute for Education and Planning, 1999.

UNAIDS Reference Group on Estimates, Modelling and Projections. "Improved methods and assumptions for estimation of the HIV/AIDS epidemic and its impact." AIDS 16 (2002): W1-W14.

Unicef. "Mother-to-Child Transmission of HIV." New York: United Nations Children's Fund, 2002.

Winkelmann, Rainer and Klaus F. Zimmerman. "A new approach for modeling economic count data." Economics Letters 37 (1991): 139-143.