Long-Range Population Projections Made Simple

JOSHUA R. GOLDSTEIN
GUY STECKLOV

LONG-RANGE POPULATION PROJECTIONS are among the most arresting scientific results that demographers produce. Currently, at least three agencies make projections that extend more than a century into the future: the World Bank (Bos et al. 1994), the United Nations (1998a), and a more recent entrant in the field, the International Institute for Applied Systems Analysis (IIASA) (Lutz 1996).[1]

In this analysis, we offer a simpler and potentially more useful means of producing long-range population projections. Our method relies on new developments in mathematical demography by Li and Tuljapurkar (1999). We offer it as a supplement to and a potential replacement for the traditional and comparatively complex approach to long-range population projections. At present analysts use the same method for producing long-term projections as for short-term projections: the cohort-component method that projects each age group according to a specified time path of age-specific fertility and mortality rates. The specification of the pattern and path of age-specific rates is inevitably complicated, resulting in projections that are hard to replicate and difficult to discuss critically. Even when extensive documentation is provided, it is difficult to replicate the calculations without access to proprietary computer software used by the team that prepared the projection.

Li and Tuljapurkar provide a simple solution for computing the eventual size of populations undergoing gradual transitions to replacement fertility.[2] We extend their results to allow for increased longevity and for gradual transitions to nonreplacement fertility levels. We find that these simple analytic formulas yield population numbers and age distributions that are quite similar to those produced by existing projections by the World Bank and the United Nations. Application of the formulas requires only a few parameters: the current population size, the current birth rate, the current intrinsic population growth rate, the ultimate life expectancy, the ultimate intrinsic population growth rate, and length of the transition period from the current to the ultimate growth rate. No age-specific rates or current age distributions are needed.

Simpler projection methods offer a number of advantages. First, the calculations take little time and are easy to replicate. Second, only a few parameters are used, reducing the number of needed assumptions. Third, the sensitivity of projected population size to parameter values is easily assessed, and alternative specifications are easily estimated. Finally, the back-of-the-envelope character of analytic projections has an important subjective advantage: rather than providing an illusion of computer-calculated exactness, the analytic estimates are transparently conjectural, an appropriate characteristic given the uncertainty about vital rates in the distant future.

Such simpler methods should enhance the usefulness of long-term projections. In high-fertility contexts, they can be used to evaluate the consequences of different rates of fertility decline in terms of eventual population size. In low-fertility contexts, they can be used to project eventual population sizes and age structures. While long-range projections typically exceed the planning horizon of most agencies interested in the consequences of demographic change—the US Social Security Administration, for example, considers actuarial balance over a 75-year time horizon—they highlight the longer-term implications of population aging. Long-range population projections can also be useful for global and regional energy, resource, and climate models that use population as an input.

Simpler and more transparent projections will provide better tools for policymakers interested in analyzing or explaining the implications of different scenarios. Bongaarts and Bulatao (1999) have shown in their analysis of the World Bank population projections that most of the variation in projections can be accounted for using simple models with only a few parameters: population momentum, the time path of fertility change, and ultimate mortality levels. Our present analysis builds on their work. Specifically, we provide new simple methods for projecting populations with long-run fertility above and below replacement levels (e.g., the United Nations high- and low-fertility scenarios). We argue that there is little advantage to producing complex projection scenarios when the underlying uncertainty of future demographic rates is so large. Our results suggest that the analytical approach should complement, and perhaps even replace, the traditional cohort-component method in generating long-range population projections.

Finally, the analytic models introduced here can be used as a complement to, or even a substitute for, simulation-based studies, such as the sensitivity analysis recently presented by O'Neill, Scherbov, and Lutz (1999). They conclude that "the difference between the initial and the eventual fertility rate over the course of the [fertility] transition is the key factor in determining the impact of the path on population size" (p. 751). As will be seen, the formulas we use to model gradual fertility transitions yield this finding directly.[3]

The state of the art: Current approaches to long-range population projections

In addition to the three major producers of global long-range population projections, several national governments (e.g., the U.S. Census Bureau) also produce such calculations. All existing projections are based on the cohort-component approach, which spells out the consequences of future patterns of age-specific mortality and fertility rates for population size and age distribution.[4] The cohort-component approach begins with a classification of the starting population into age groups. In each subsequent projection step, each age group survives and produces offspring according to posited schedules of future age-specific rates.

The age-specific nature of the projections is both an advantage and a drawback. In the short run, when the effect of the starting age structure on population growth is relatively large and the variability in vital rates relatively small, the procedure performs very well (Keyfitz 1981; Stoto 1983; Bongaarts and Bulatao 2000). For longer-term projections, however, the uncertainty in the assumed path of fertility and mortality rates eventually overwhelms the additional precision gained from using the exact starting age structure of the population. In the very long run, the age structure and growth rate of the population depend entirely on the assumed path of vital rates, and not at all on the starting population. The size and structure of the starting population will express themselves only in the form of population momentum, that is, the changes that would ensue if fertility instantaneously dropped to replacement level (Keyfitz 1971).

Thus what really matters in a long-range projection is the future path of vital rates and the momentum implied by the starting age structure. Each of the published long-range population projections relies on a different method to generate vital rates estimates. The precise details of the methodology are complex and rarely described in full. For example, a set of long-range population projections of the United Nations (1998a) refers readers to a separate volume where parts of the methodology are more fully explained (United Nations 1998b). The World Bank provides a seven-page description of its methods (Bos et al. 1992: 11–17), while IIASA describes its methods in a chapter of a book (Lutz 1996). Despite the length of these detailed methodological descriptions, it is difficult to reproduce the projections without the actual computer program used to generate them.

Despite differences in details, these projections share the same basic logic. First, the trajectories of summary indicators like the total fertility rate (TFR) and life expectancy at birth, $e(0)$, are specified for each scenario. Second, these summary indicators are translated into age-specific rates for each period in the projection. Third, cohort-component methods are used to project a starting population, distributed in five-year age groups and by sex, for the duration of the projection.

Although simple in theory, each of the preceding steps can become complex in practice. For example, scenarios are specified in a variety of ways. At the heart of many long-range projections lies the assumption that fertility levels will converge to some constant level. The World Bank (e.g., Bos et al. 1994), for example, assumes that every country will eventually have replacement-level fertility, but its approach still requires a method for deciding when and how replacement fertility will be reached. It assumes that fertility for each country follows one of six patterns over time, depending on the initial value of fertility and, in some cases, on mortality levels. In the United Nations long-range projections (1992, 1998a) the "medium" scenario assumed that fertility would, by 2050, converge to replacement levels. But the paths to replacement were sometimes complicated, following curved rather than straight-line paths and including reversals of current trends. More recently, the United Nations (2000) presented a "medium" scenario in which fertility levels of countries currently below replacement were no longer assumed to rise to replacement by 2050 but do so by the middle of the twenty-second century. In all of the United Nations projections, alternative scenarios are based on proportional departures from the "medium" scenario.[5] In specifying its scenarios, IIASA (e.g., Lutz 1996) relies on a survey of expert opinion to help assess high and low scenarios, with the "most likely" being the average of the two extremes.

The translation of summary indicators into age-specific rate schedules also turns out to be a nontrivial task. Both the World Bank and the United Nations, for example, use three model age-specific schedules for fertility to translate the TFR into age-specific rates. The age-specific fertility schedule for each country may be chosen from the country's own statistics or, when those data are unavailable or of questionable quality, by comparing the most recently available age-specific fertility pattern to model schedules (United Nations 1998b). Furthermore, age-specific profiles are allowed to vary over time. IIASA, on the other hand, assumes no change in the shape of the fertility age schedule and applies only proportional adjustments (Lutz 1996: 374). Mortality also presents the challenge of translating aggregate indexes into age-specific rates. While model life tables are useful for specifying age-specific mortality schedules for the near future, little empirical basis exists for life tables for the very low levels of mortality that are projected for the more distant future.

Finally, the projection method itself involves a number of issues. Since Whelpton (1936) produced the first cohort-component projections, various formulas have been developed for converting age-specific rates into transition coefficients in the projection matrix. Each of the three producers uses different software for producing its projections, none of which is publicly available.

In some ways, the computerization of population projections is itself responsible for the complexity of the assumptions. The software is designed to allow specification of any time path of rates. The value added by producers of projections is the specification of time trajectories considered of particular interest or plausibility. It does not require expertise to assume that every country in the world will reach replacement fertility by some arbitrary target year, but it does take expertise to justify a variety of country-specific patterns over time. Not surprisingly, producers of projections want to take advantage of the flexibility that their computer programs provide.

Despite the careful attention paid to the details of population projection, the uncertainty of long-term projections is extremely large, primarily owing to the multiplicative effect of small differences in the governing assumptions. The UN's projections highlight this phenomenon. The high and low fertility scenarios differ by just one child per couple, half a child above and below replacement levels, but produce projected populations in 2150 ranging from 3.6 billion to 27.0 billion (United Nations 1998a: vii). The question, we believe, is not whether the difference between such alternative scenarios accurately reflects the underlying uncertainty, but rather what is gained from the application of the exacting method of cohort-component projections given the intrinsic uncertainty in any results.

One very simple alternative to cohort-component methods for long-range projections is the exponential growth model. In this case, the effect of a given change in total fertility on future population size is measured by the effect of this change on the stable exponential rate of growth. Such an approach would completely ignore the population's age structure.[6] The analytical projection method developed in the next section offers an intermediate solution. Unlike the simple stable population model, it incorporates the role of age structure by approximating the effect of population momentum. However, in contrast to the cohort-component approach, our proposed method does not use all of the information about the current age structure nor does it require age-specific estimates of future demographic rates.

A new approach: Analytic projections

Our analytical approach is based on results from a recent article by Li and Tuljapurkar (1999), which provides a general formula for calculating the population momentum implied by gradual transitions to replacement fertility. Their results build on Keyfitz's (1971) formula for the population momentum implied by instantaneous transitions to replacement fertility, that is, to a level of fertility which, in combination with prevailing mortality levels, yields a net reproduction rate of 1. Keyfitz found that (M_k)—the ratio of a population's ultimate size after the transition to its size at the time of the transition—was

$$M_k = \frac{be(0)}{r\mu}\frac{NRR-1}{NRR}, \tag{1}$$

where b is the birth rate before the transition, $e(0)$ is life expectancy at birth (assumed to remain constant in Keyfitz's formula), r is the intrinsic growth rate of the population before the transition, μ is the mean age of childbearing, and the NRR is the net reproduction rate before the transition. Keyfitz's hypothesized instant drop to replacement fertility has become a standard measure of population momentum.[7] Both the World Bank and the United Nations include comparisons of their projections with Keyfitz's scenario, although they implement it using projection software rather than analytically (Bos et al. 1992; United Nations 1998a).

It was long thought impossible to construct a simple formula for gradual demographic transitions. Using results from the statistical theory of renewal processes, however, Li and Tuljapurkar developed a general formula applicable to a population that follows any time path to replacement fertility.[8] They derived this formula using their solution of time-dependent population models (Li and Tuljapurkar 2000). Among their results is a simple formula for the eventual (long-term) size of the population in the case of a linear decline (or increase) in fertility to replacement level:

$$M_{LT} = M_k\left(\frac{e^{r\gamma}-1}{r\gamma}\right), \tag{2}$$

where M_k is Keyfitz's momentum as in equation (1), r is again the intrinsic growth rate in the population before the onset of the transition, and γ is the number of years it takes for the transition to occur. This simple formula holds exactly only for relatively short transitions. Using simulations, however, they show that the formula provides "accurate approximation" for transitions that last as long as 40 years. Their simulations also show that differences between the ultimate population sizes that result from linear and logistic fertility declines are small (Li and Tuljapurkar 1999).

We employ an approximation of Li and Tuljapurkar's result that is simpler in form and slightly more accurate for longer transitions.[9] Because $\left(e^x-1\right)/x \approx e^{x/2}$ for values of x relevant for population growth rates, we can write

$$M_{LT} \approx M_k \exp\left(\frac{r\gamma}{2}\right). \tag{3}$$

This formulation has an intuitive appeal. It shows that the momentum of a gradual transition is about equal to that of a Keyfitz-style instantaneous transition that occurs in the middle of the gradual transition period, by which time the initial population has grown by a factor of $e^{r\gamma/2}$. For example, a gradual fertility transition of a population from a net reproduction rate of 2.0 in 2000 to a net reproduction rate of 1.0 in 2030 produces about the same long-run population size as if it underwent an instantaneous transition from 2.0 to 1.0 in 2015.

In the next section, we compare estimates based on equation (3) to the population projections that have been made by the World Bank and the United Nations using the cohort-component method. Before doing so, we need to extend Li and Tuljapurkar's results slightly. First, we need to allow for projected changes in mortality conditions, which are assumed constant in the Keyfitz and the Li and Tuljapurkar formulas. We do this by replacing the initial life expectancy value ($e(0)$ in equation (1)) with its ultimate value for the projection period. This gives

$$M_{GS} = M_{LT} \frac{e(0)^{ultimate}}{e(0)^{initial}} . \tag{4}$$

Thus, to produce an estimate of India's population size in 2150 comparable to that projected by the World Bank, we assume that life expectancy will have increased to about 85 years—as the World Bank posits—rather than using the country's current value of roughly 65. This involves inflating Li and Tuljapurkar's result by a factor of 85/65, or 1.31. This simple adjustment is possible because the direct effect of mortality change at all ages can be captured by using the life expectancy ratio given in equation (4). The indirect effects of mortality change via net fertility need not be captured in the ratio because they are already accounted for in the time path of net reproduction (NRR).

Our second extension is to include transitions to stable demographic rates that are not at replacement level (i.e., populations that continue to grow or shrink indefinitely at a fixed exponential rate after the fertility transition came to an end).[10] Such scenarios are used by the United Nations (1998a) to generate all their scenarios other than "medium."[11] We denote the ultimate intrinsic growth rate of the population as r^* and the difference between the initial and the ultimate intrinsic growth rate as $r - r^* = \rho$. For example, if the starting stable growth rate were 3 percent and the eventual stable growth rate were 1 percent, then ρ would be 0.02. We find that the equilibrium path of births grows exponentially at the new rate, with the level of births at any time t given by

$$B(t) = B(0)\frac{e^{r^*t}}{\mu\rho}\left(\frac{e^{\rho\gamma}-1}{\rho\gamma}\right)\left(1-e^{-\mu\rho}\right) . \tag{5}$$

This formula can be simplified by employing the approximation $\left(e^x - 1\right)/x \approx e^{x/2}$ used earlier. In this case, we find that

$$B(t) \approx B(0)e^{r^*t}e^{\rho(\gamma-\mu)/2} . \tag{6}$$

The long-range birth sequence thus grows exponentially at rate r^*, with a one-time adjustment that incorporates the magnitude of fertility change ρ, the length of time over which fertility change occurs γ, and generation length μ.

Using the formula for $B(t)$, the total population size, K, as well as the numbers and proportions in each age group, can be estimated using stable population theory (see, e.g., Keyfitz 1985):

$$K(t) = B(t)\int e^{-r^*t}l^*(a)da , \tag{7}$$

where $l^*(a)$ is the life table survivorship to age a in the new stable regime. In all of our projections we have used a common life table based on the United Nations female life table with life expectancy at birth equal to 87.5 years, since in the UN projections all of the regions are assumed to have an eventual life expectancy in the range of 85 to 90 years by the year 2150.

These analytical results take into account the effects of population momentum, but allow demographic rates to change gradually instead of assuming an instantaneous change. Although it is possible to estimate analytically population age structures and sizes in both the short and long run (see Li and Tuljapurkar 2000), the solutions given here apply only after the transient population waves, caused by the difference of the initial age distribution from the age distribution of the ultimate stable population, have for all practical purposes died out, which typically takes several generations. We focus on long-term projections of at least a century, by which time any remaining transient effects are negligible.

Results: Empirical comparison of analytic and cohort-component projections

Using only a few parameters, our analytical projection method can replicate any set of projections and evaluate the effects of variants of the initial assumptions. In this section, we repeat the basic long-range population projections of the World Bank and the United Nations using the analytical projection method. We try to minimize the effects of factors other than the projection methodology itself. Thus, we compare the results of the two projection procedures when they use starting values that are as much alike as possible.

What yardstick should be used to assess the performance of analytical projections? How close is close enough? The recent analysis by Bongaarts and Bulatao (2000) of the accuracy level of past population projections by the United Nations and the World Bank found that simple misestimation of base populations at the start of projections accounts for about 3 percent error in the projected population size. As projection duration lengthens, the magnitude of the error grows fairly steadily at about 2.5 percent for each additional five years forecast, as the rates specified in the projection diverge from what actually occurred. For individual countries within each world region, the mean absolute proportional error[12] is between 5 and 23 percent for projections of 30 years' duration. The average error for all countries after 30 years is 17 percent. Extrapolating these historical findings into the future—an admittedly speculative exercise—we would expect the cohort-component country projections to err, on average, by about 75 percent after 150 years.

In the next section, we begin with projections of transitions to replacement-level fertility, which include all of the World Bank projections and the medium-fertility UN scenario. We then compare the formula to the nonreplacement-level fertility scenarios produced by the United Nations. In the last part of this section, we discuss the various factors that may explain discrepancies between the analytic and cohort-component projections.

The transition to a stationary population

Analytical projections based on World Bank long-range projections are particularly easy to prepare since most of the needed parameters are given in Bos et al. (1992).[13] The two projections in terms of projected total population size in 2150 by world regions are compared in Figure 1. The analytic method produces very nearly the same predictions of population size as the World Bank's cohort-component method. The absolute value of the proportional difference between the results of the projection methods is never larger than 16 percent for any single region, and the average of the absolute values of the proportional difference for the ten separate regions is 11 percent.

We also compared projections based on the analytic method with the cohort-component projections for each individual country. The analytical approach replicates quite closely the results obtained through cohort-component projections, although there is greater variability. Figure 2 shows a scatterplot of the projected country population sizes in 2150 according to the World Bank and according to the analytic projection approach. If the two methods produced identical projections, all points would lie on the diagonal line. The left-hand panel plots the predicted population sizes in 2150 according to the two methods. The overall r-squared estimate for the correlation between population sizes predicted for 2150 according to the two methods is 0.998. This value is nearly identical with the estimate produced in Bongaarts and Bulatao (1999), suggesting that both approaches are able to capture practically all of the variation in the projected population sizes. The right-hand panel plots the growth factor—the ratio of the projected population in 2150 to the population in 1990—calculated for each country

FIGURE 1 Population size in 2150 by world regions, as projected by the World Bank using the cohort-component method and as projected by the analytic method discussed in the text. Both projections assume long-term convergence to a stationary population.

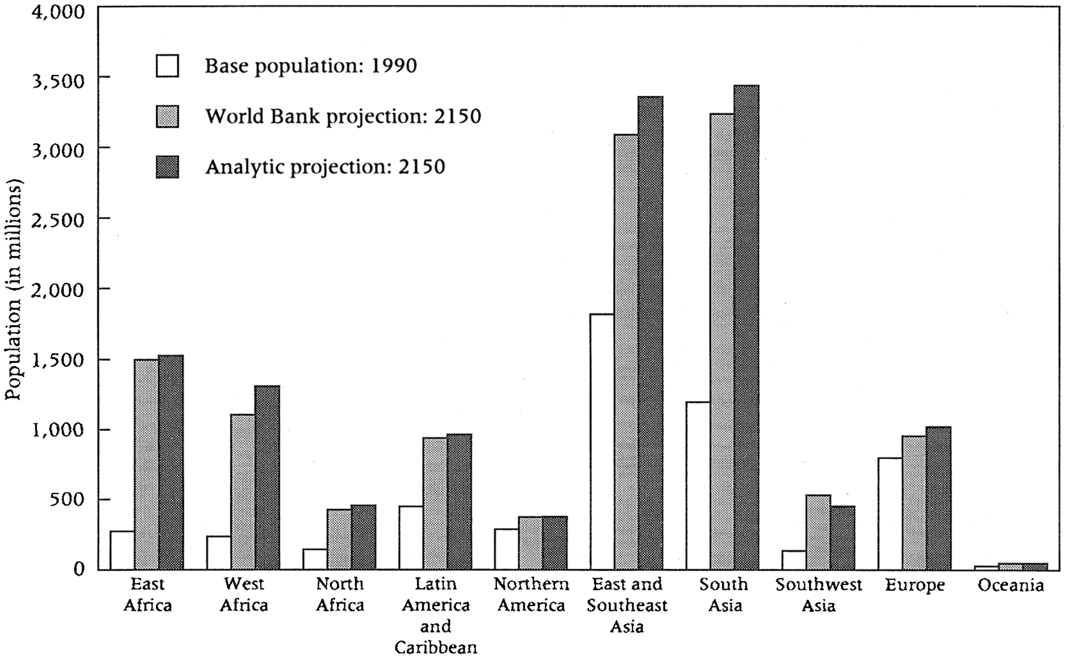

NOTE: "Europe" includes the former Soviet Union. For definitions of other regions see Bos et al. 1992.
SOURCE: World Bank projections from Bos et al. 1992.

according to the two methods. A log-log scale is used in the figure in order to show detail among the many countries with growth factors in the 1.0 to 2.0 range. Projected population sizes obtained through the analytic method are quite close to those yielded by the cohort-component method. The same conclusion can be drawn from the growth-factor plot. However, this latter plot also suggests that the analytic method tends to slightly overestimate population sizes for the fastest-growing countries—a finding that is not easily drawn from the population size plot on the left. This apparent bias is consistent with Li and Tuljapurkar's simulations and, as they have shown, is attributable to the first-order nature of their approximation (Li and Tuljapurkar 1999).

We also compared our analytic method with the UN medium-fertility projections, which assume that every country will reach replacement fertility by 2050–55. Our analytical projection estimates from the UN data are based on United Nations (1998a, 1998b).[14] The results in Figure 3 again highlight the general consistency between the analytical projection estimates

FIGURE 2 Population size in 2150 and ratio of population in 2150 and population in 1990 in 194 countries as projected by the World Bank using the cohort-component method and as projected by the analytic method discussed in the text. Both projections assume long-term convergence to a stationary population.

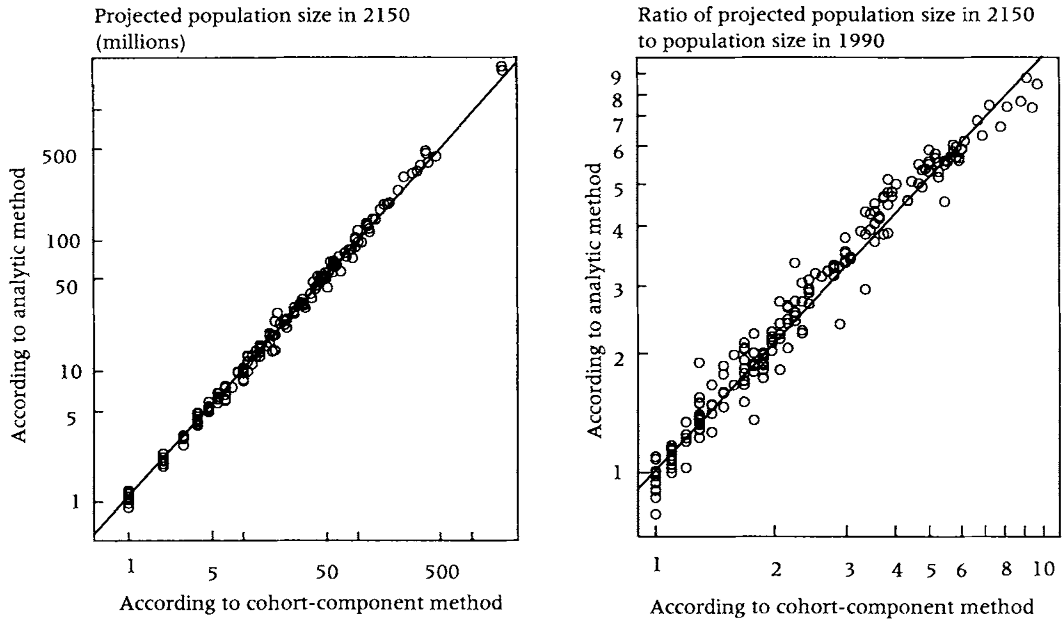

SOURCE: World Bank projections from Bos et al. 1992.

and the cohort-component-based estimates. The deviations between the results obtained from the two approaches appear slightly larger in the case of the UN projections as compared to the World Bank projections. Overall, the average of the absolute values of the proportional errors between the analytical and UN projections is 13 percent. In this case, however, the measure of error with the exception of China and India is made for regions only, since the United Nations did not prepare long-range projections country-by-country. We find India's proportional error at about 20 percent and Africa's at 17 percent. The analytical projections for the other relatively young and growing regions, namely Other Asia and Latin America, and also for China, are all somewhat higher than the UN projections. The analytical projections for Europe are essentially identical with the UN projections, and our projections for Northern America and for Oceania are about 6 and 7 percent lower.

Are these discrepancies large or small? We have already argued that such errors are small relative to the overall uncertainty in projections. An-

FIGURE 3 Population size in 2150 by world regions and in China and India, as projected by the United Nations using the cohort-component method and as projected by the analytic method discussed in the text. Both projections assume long-term convergence to a stationary population.

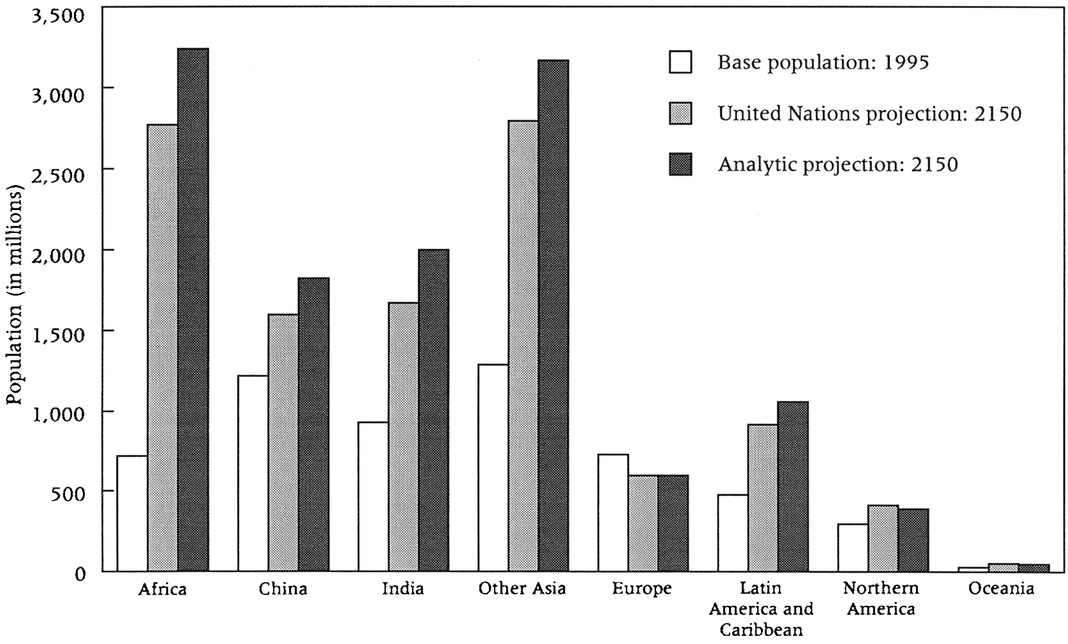

NOTE: "Europe" includes Russia and the republics of the former Soviet Union in Europe. Other Asia includes Asia outside China, India, and Russia.
SOURCE: United Nations projections from United Nations 1998a.

other way to put the magnitude of such errors into context is to translate them into how much a given parameter in the projections would have to change in order to make the two projections agree exactly. For example, were we to assume that the time required for Africa to reach replacement fertility will be not 50 years but 38 years, the difference between the analytical and UN projections would entirely disappear.

The transition to nonstationary demographic regimes

Here we compare the results of our analytical method with the UN's high-fertility and low-fertility projections (United Nations 1998a). Both of these UN scenarios, as well as the medium-fertility scenario referred to earlier, assume the same improvement in life expectancy over time. Figures 4 and 5 present the projected population in 2150 for the low-fertility and high-fertility scenarios according to the cohort-component and analytical projection approaches.

FIGURE 4 Population size in 2150 by world regions and in China and India, as projected by the United Nations using the cohort-component method and as projected by the analytic method discussed in the text. Both projections assume convergence to long-term fertility levels as assumed in the United Nations "low" scenario.

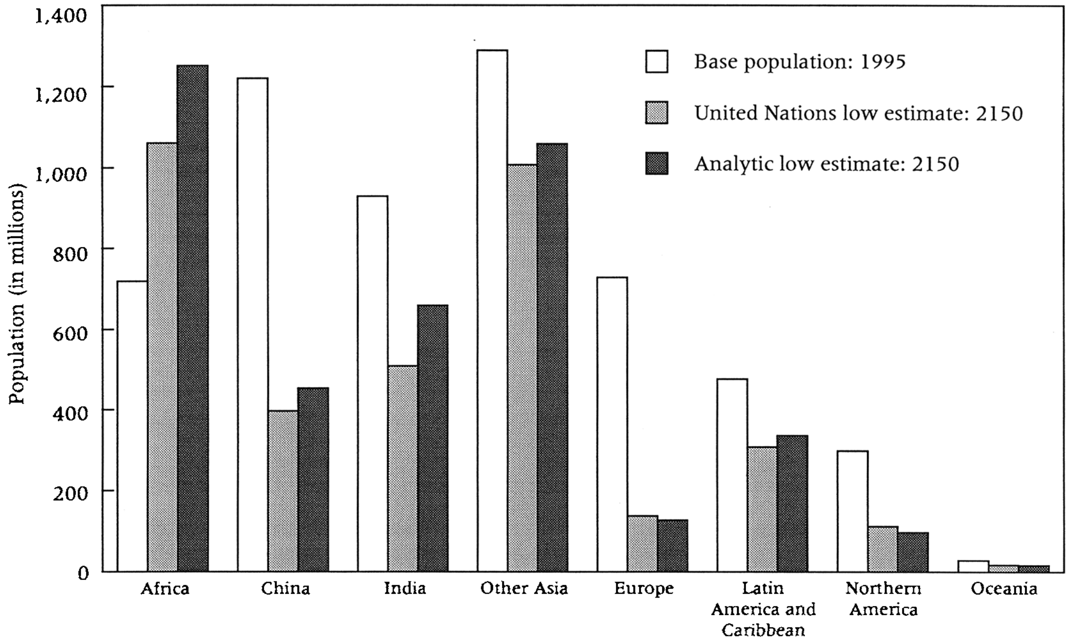

NOTE: See Figure 3 on geographic definitions.
SOURCE: United Nations projections from United Nations 1998a.

In both high and low scenarios the analytical method closely replicates the population size projected by the cohort-component approach. Overall, the average of the absolute values of the proportional differences is 13 percent for the low-fertility scenario and 8 percent for the high-fertility scenario. Comparing all three UN scenarios, the least good fit is obtained for the low-fertility scenario. Yet, even in this case, the largest proportional differences are 18 percent for Africa and 29 percent in India. As we noted earlier, the analytical projections tend to overestimate the population in 2150 relative to cohort-component projections, and this effect is larger in the low-fertility scenario.

Relative to the uncertainty inherent in specifying any demographic rate 155 years into the future, the differences between the projections are small. In fact, even the largest proportional difference between the two approaches—that of 29 percent obtained for the low-fertility scenario for India—would be eliminated if a long-term net reproduction rate of 0.72 were used in the analytic projection instead of a value of 0.77 posited by the UN.

FIGURE 5 Population size in 2150 by world regions and in China and India, as projected by the United Nations using the cohort-component method and as projected by the analytic method discussed in the text. Both projections assume convergence to long-term fertility levels as assumed in the United Nations "high" scenario.

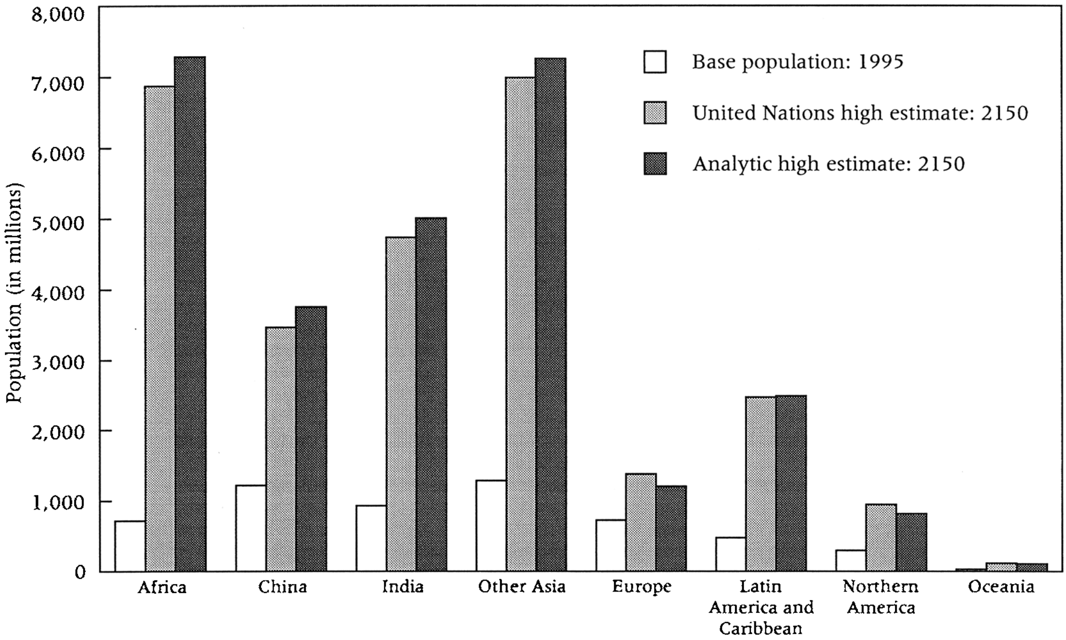

NOTE: See Figure 3 on geographic definitions.
SOURCE: United Nations projections from United Nations 1998a.

Possible sources of disagreement

In this section we discuss the possible sources of the discrepancies between our analytic results and the cohort-component projections published by the World Bank and the United Nations. We also try to distinguish between those sources of disagreement that may arise whenever one tries to replicate another set of population projections and those sources of disagreement that are due to the differences in projection methods.

Some of the differences between the population sizes in 2150 projected by the two approaches are the consequence of comparing any two sets of projections using different methods. Despite our efforts to ensure comparability, the parameters used as initial values in the two projections may not be in perfect agreement, or we may have unknowingly introduced errors of rounding, copying, or calculation. It is also possible, given the complexity of the cohort projection method, which relies on arrays of data and extensive computing, that errors may have been introduced into the cohort-component projections.

A second factor that is unrelated to the difference in methods but may affect the comparisons presented here is the influence of different levels of aggregation. The UN long-term projections are based on individual country projections through 2050 and beyond that date on projection of the population of two large countries and six regions. For simplicity, we projected these units for the entire period, beginning in 1995. Thus, some of the differences are probably attributable to the different levels of aggregation.

The differences intrinsic to the two approaches are (1) the approximate nature of the analytic formula for very slow fertility transitions; (2) different specifications of the time path of net fertility change during the transition; (3) different specifications of the age pattern of fertility; (4) the presence of international migration, for which the UN projections make allowance between 1995 and 2050; and (5) different specifications of the starting age structure.

The result we use from Li and Tuljapurkar (1999) is itself an approximation for fertility transitions lasting longer than the minimum age of childbearing, about 15 years. As we noted above, using simulation, Li and Tuljapurkar show that the approximation is accurate for transitions lasting as long as 40 years. They find that the contribution of generations born several generations after the transition begins is relatively small and that this explains why the longer transitions can be approximated with the formula for shorter transitions. In general, the analytic formula (equation (2) above) produces a slight overestimate of the eventual population size. However, our approximation (equation (3)) will tend to counterbalance this tendency. One avenue to attain additional precision in the projections is to employ Li and Tuljapurkar's (1999: 258) second-order approximation (formula not shown here), which is exact for transitions as long as 30 years but also relatively accurate for transitions as long as 100 years. However, given the added complexity of their second-order approximation and since the magnitude of the approximation error is quite small—on the order of a few percent for transitions lasting 60 years—the simpler formula can be used at a relatively small price in accuracy.

The analytic formulas we use assume linear declines in net fertility over the course of the transition. The path of fertility decline in the UN scenarios is essentially a logistic function. The World Bank trajectories are also quite complicated, depending as they do on the recent history of fertility change, but in the default case are equivalent to geometric decline. Simulations by Li and Tuljapurkar (1999) suggest that the specific path of fertility decline does not greatly influence the results. Alternative analytic formulas can be derived. For example, Li and Tuljapurkar present a result for a logistic pattern of net fertility change. Here again, however, as with the approximation for longer transitions, the effect of specifying the exact time path is quite small.

The age pattern of fertility may also account for some of the variation in the projections. For example, the assumed age pattern of fertility for different countries and regions may take one of three possible patterns in the UN projections (United Nations 1998b), whereas our model assumes that the mean age of childbearing is fixed at 29 years for all regions. While this factor may be responsible for some variation, it is likely that, as regional fertility levels converge, the assumption of a single mean age of childbearing (although not necessarily its actual assumed value) becomes increasingly reasonable.

International migration is a factor in the UN projections but not in those of the World Bank. Our projections have, again for simplicity, ignored migration. The UN projections assume that net migration between major world areas drops to zero by 2025. Differences in international migration assumptions may also partly explain why our projections appear closer to those of the World Bank projections than to those of the United Nations.

Perhaps the most important difference between the analytic method and the cohort-component method is the starting age structure. Our analytic method assumes that the age structure of the populations at the beginning of the forecast is that of a stable population with the observed crude growth and birth rates, while the cohort-component projections are based on the observed age distributions of each country/region, which in most cases are far from the stable state. While in theory this could produce large differences in ultimate population size, it appears that the actual differences due to this fact are quite small. One reason for this is that developed countries, which differ from stability by having more people in prime childbearing ages, tend to have fewer individuals in the pre-childbearing ages, and therefore any extra population growth early on tends to be canceled out when the smaller cohorts enter childbearing age.

The relative influence of these various factors on the differences between the populations projected by the two methods remains a subject for further analysis. Overall, the net magnitude of the differences—on the order of 10 to 20 percent—seems small enough to warrant the application of the simpler and more transparent analytic method.

Conclusion

We have argued that in preparing long-range population projections it may be better to rely on simple back-of-the-envelope calculations based on an analytic model than to use cohort-component projections. The two techniques produce similar projection estimates, with the difference well below the subjective margin of error that long-term estimates about fertility and mortality 100 years or more into the future warrant. But the simpler calculations are easier to carry out, their assumptions are easier to understand, and they convey a clearer and, we would argue, more accurate sense of the uncertainty inherent in long-range projections.

The relative simplicity of the analytic method we presented and illustrated is achieved by circumventing age-and-sex-composition–based traditional population projections by using recently developed results from mathematical demography (Li and Tuljapurkar 1999). The formulas apply to age-structured populations, because they include the dynamics of population momentum; however, simplifying assumptions permit the derivation of the relatively simple expressions we have presented here. These assumptions are that (1) the age structure of the starting population is stable, (2) net fertility changes proportionally at all ages, and (3) the time path of net fertility change is linear. All of these assumptions can be relaxed, but at the price of making the resulting analytical formulas more complicated.[15]

Although the fact that actual populations may be at variance with these assumptions could conceivably have large effects, in practice this is not the case. It appears that for practical purposes most contemporary populations are close enough to their stable equivalents to produce estimates that differ little whether one uses the observed age structure as the starting point of the projection or the stable age structure implied by the initial vital rates. Li and Tuljapurkar found, through simulation, that introducing some variation into the age-specific pattern and the time path of fertility change did not have large effects on projected long-run population size. A comparison of our analytic projections with those produced by the World Bank and the United Nations offers additional confirmation of this contention.

Analysts have little empirical or theoretical basis for predicting long-term demographic futures. They are not able to foretell, for example, when replacement fertility will be reached in a given population, let alone specify the trajectories of all age-specific rates. But long-range projections are not forecasts; they merely aim to illustrate the long-term consequences of specified "interesting" future trends in fertility and mortality. In painting an illustrative "big picture" of the long-term evolution of the size of populations, it is convenient to use broad strokes—relying only on general indicators like life expectancy and the net reproduction rate. Fortunately, our results indicate that one does not lose much by leaving the finer brushes out of the tool kit.

Analytic methods for projecting populations may offer some promise for developing stochastic projections, which not only present central projections but also seek to quantify the uncertainty surrounding them (Lee 1999). The production of such projections is a key recommendation of the recently published Bongaarts and Bulatao (2000) report on population projection. The potential advantage of the analytic approach for stochastic projections is that it requires specification of only a few parameters and allows developers of probabilistic models to focus on the components that will make the most difference to long-term population size: the long-term level of net fertility, the speed of any fertility transition that might occur, and the long-term level of life expectancy.

What then is the role for cohort-component projections? First, the projection method is well suited to short-term forecasts because the initial population age structure plays a large role in shaping the near-term demographic future and changes in vital rates play a relatively small role. Second, in the longer term, cohort-component projection offers an easy way (requiring no mathematics) for exploring the long-term consequences of irregular scenarios such as the effect of epidemics like AIDS, of massive in- or out-migration, or of nonmonotonic trends in fertility. The two methods can also be used to complement one another. For example, the implementation of medium-term cohort-component forecasts can be checked by extending the projection period to the long term and by comparison with analytic estimates. Disagreement is not a sure sign that something is wrong with the cohort-component projection, but it is a warning sign. Finally, the analytic estimates can be used as a baseline against which the consequences of other scenarios, such as increased AIDS mortality, can be compared.

Appendix

Long-run population projections can be easily generated using equations (1) and (3). As an illustration, we present a projection that replicates the World Bank projection for East Africa. All the values are obtained from the World Bank (Bos et al. 1992) except for the initial intrinsic growth rate, r, which we approximate using $\ln(\mathrm{NRR})/\mu$. The time until replacement fertility is reached, γ, is taken as the population-weighted average of the γ assumed by the World Bank for the individual countries in East Africa. All the values for East Africa are provided below in Table A1. The last two rows present a comparison between the ratio of the projected 2150 populations to the 1990 population according to the World Bank and according to the analytical method. The two projections for East Africa in 2150 are reasonably similar: 1,488 million according to the World Bank compared to 1,410 million obtained through the analytic method.

TABLE A1 Population size of East Africa in 2150, projected from the 1990 population according to the analytic method and according to the cohort-component method used by the World Bank

Population in 1990 (P_{1990})	272 million
Crude birth rate in 1990 (b)	0.045
Net reproduction rate in 1990	2.4
Life expectancy ($e(0)$) in 2150	84 years
Mean age of childbearing (μ)	29 years
Crude growth rate (r) in 1990	0.029
Length of transition (γ)	47 years
Population in 2150 (P_{2150}) according to World Bank	1,488 million
Population in 2150 (P_{2150}) according to analytic method	1,410 million
Growth factor (M) according to World Bank	5.47
Growth factor (M) according to analytic method	5.18

The analytical growth factor in this example is calculated as

$$M_{GS} = \frac{be(0)^{ultimate}}{r\mu}\left(\frac{NRR-1}{NRR}\right)\exp\left(\frac{r\gamma}{2}\right) = \frac{(0.045)(84)}{(0.029)(29)}\left(\frac{2.4-1}{2.4}\right)\exp(0.029\times 47/2) = 5.18,$$

where the formula for M_{GS} is obtained by combining equations (1), (2), and (4) in the text.

Projection through the analytic method requires only a hand calculator, although a spreadsheet program is useful if one wishes to examine the implications of a variety of scenarios for long-run population size.

Notes

The authors thank Henry N. Goldstein, Ronald D. Lee, Nan Li, and Charles F. Westoff for useful comments on an earlier version and Estella Rivero Fuentes for research assistance.

1 The IIASA results, predicting that humans have completed their final population doubling, became one of the rare pieces of demographic research to be published in the international scientific journal *Nature* (Lutz et al. 1997. See also Lutz et al. (1998), Lutz et al. (2001), and United Nations (2000) for long-range projections that were unavailable at the time of writing.

2 An alternative derivation for some of these results is given by Goldstein (2002).

3 The difference between current and eventual growth rates is a parameter (ρ) in the analytic expression for eventual population size.

4 This "demographic" approach contrasts with systems approaches that include the consequences of environmental and resource constraints. The systems approach, however, has produced little consensus about the carrying capacity of the planet and proves extremely sensitive to small changes in assumptions (Cohen 1995; Sanderson 1995).

5 In all cases, the United Nations explains that its scenarios are not intended as forecasts but as population outcomes if vital rates progress along a particular trajectory.

6 Another simple alternative is to assume a linear decline in population growth rates. Ronald Lee has shown that this procedure, extrapolating the recent trend, was able to reproduce cohort-component projections of population size by the United Nations and the U.S. Census Bureau with discrepancies below 1.5 percent through 2050 (Lee 1990: 48). Such an approach does not explicitly include the effects of population momentum, but does so implicitly in assuming a gradual change in population growth rates.

7 An alternative formula for estimating momentum is $M_k \cong be(0)/\sqrt{NRR}$ (Keyfitz 1985). Following Preston (1986), Bongaarts and Bulatao (1999) suggest using the ratio of the proportion of the female population under age 30 at the start and end of the projection period as a robust measure of momentum. See the debate between Preston (1988) and Wachter (1988) on the generalizability of the Preston approach.

8 Schoen and Kim (1998) offer an alternative approach conditional on the leveling off of the birth stream. They show that if the growth rate in births linearly declines over a time period Y, then the ultimate level of births will be $B(0)\ \exp[r_B(0)Y/2]$, where $B(0)$ is the initial number of births and $r_B(0)$ is the stable growth rate of births prior to the transition. We use the Li and Tuljapurkar approach because it follows from a given trajectory in the net reproduction rate rather than a given trajectory of births and thus resembles most closely the approach to population projections taken by international agencies.

9 We thank an anonymous reviewer for suggesting this simplifying approximation.

10 We derive this result by following Li and Tuljapurkar's original derivation, replac-

ing the stationary renewal equation results with those for net maternity functions implying net reproduction rates other than 1.0. Further details are available upon request from the authors.

11 IIASA also includes a number of alternative scenarios, some of which could be approximated using this method.

12 The mean of the absolute value of proportional error is a useful measure of differences that we employ in our analysis. If P_{i1} is the population estimate for area i based on method 1 and P_{i2} is the population estimate for area i based on method 2, used as the benchmark method, the mean absolute proportional error for a population of N units is equal to

$$\frac{1}{N}\sum_{i=1}^{N}\left|\frac{P_{i1}-P_{i2}}{P_{i2}}\right|.$$

13 The birth rate and net reproduction rate in 1990 are estimated from the mean of the adjacent five-year periods. We estimated the intrinsic rate by dividing the log of the NRR by the mean age of childbearing, which was assumed to be 29 years. Life expectancy, $e(0)$, represents the projected level at the end of the projection period. The length of the transition for each region was estimated as the average of the individual country transition lengths, weighted by 1990 population size. Because the regions are actually aggregates of individual countries, which are separately projected, the regional value is an approximation. The approximation is based on an average of the date at which the replacement fertility level is reached in each country within a region, weighted by the 1990 population size of each country.

14 Unlike the case of the World Bank projections, the UN comparison is based on projections from the year 1995. All rates in the UN publications refer to five-year periods; we estimated birth rates and NRR values for 1995 by averaging the 1990–95 and 1995–2000 rates. We estimated the intrinsic rate of natural increase by dividing the log of the NRR by the mean age of childbearing, which we again assumed to be 29 years. The length of the transition, γ, we estimated by examining the specified fertility trajectories for each of the world regions from the United Nations medium-fertility scenario (United Nations 1998a).

15 Indeed, Li and Tuljapurkar (1999) present general formulas that hold even when the second and third assumptions are relaxed.

References

Bongaarts, John and Rodolfo A. Bulatao. 1999. "Completing the demographic transition," *Population and Development Review* 25(3): 515–529.

——— (eds.). 2000. *Beyond Six Billion: Forecasting the World's Population.* Washington, DC: National Academy Press.

Bos, Eduard, My T. Vu, Ann Levin, and Rodolfo A. Bulatao. 1992. *World Population Projections 1992–93 Edition: Estimates and Projections with Related Demographic Statistics.* Baltimore: Published for the World Bank by Johns Hopkins University Press.

Bos, Eduard, My T. Vu, Ernest Massiah, and Rodolfo A. Bulatao. 1994. *World Population Projections 1994–95 Edition: Estimates and Projections with Related Demographic Statistics.* Baltimore: Published for the World Bank by Johns Hopkins University Press.

Bulatao, Rodolfo A. and Ann Elwan. 1985. "Fertility and mortality transition: Patterns, projections, and interdependence," World Bank Staff Working Paper No. 681. Washington, DC.

Cohen, Joel E. 1995. *How Many People Can the Earth Support?* New York: W.W. Norton.

Goldstein, Joshua R. 2002. "Population momentum for gradual demographic transitions: An alternative approach," *Demography* 39(1): 65–73.

Keyfitz, Nathan. 1971. "On the momentum of population growth," *Demography* 8(1): 71–80.

———. 1981. "The limits of population forecasting," *Population and Development Review* 7(4): 579–593.

———. 1985. *Applied Mathematical Demography*. 2nd ed. New York: Springer.

Lee, Ronald D. 1990. "Long-run global population forecasts: A critical appraisal," *Population and Development Review* 16(Supp): 44–71.

———. 1999. "Probabilistic approaches to population forecasting," *Population and Development Review* 24(Supp): 156–190.

Li, Nan and Shripad Tuljapurkar. 1999. "Population momentum for gradual demographic transitions," *Population Studies* 53(2): 255–262.

———. 2000. "The solution of time-dependent population models," *Mathematical Population Studies* 74: 311–329.

Lutz, Wolfgang (ed.). 1996. *The Future Population of the World: What Can We Assume Today?* Revised Edition. London: Earthscan.

Lutz, Wolfgang, Warren C. Sanderson, and Sergei Scherbov. 1997. "Doubling of world population unlikely," *Nature* 387(6635): 803–805.

———. 2001. "The end of world population growth," *Nature* 412: 543–545.

Lutz, Wolfgang, James W. Vaupel, and Dennis A. Ahlburg (eds.). 1998. *Frontiers of Population Forecasting*. Supplement to *Population and Development Review*, Vol. 24. New York: Population Council.

O'Neill, Brian C., Sergei Scherbov, and Wolfgang Lutz. 1999. "The long-term effect of the timing of fertility decline on population size," *Population and Development Review* 25(4): 749–756.

Preston, Samuel H. 1986. "The relation between actual and intrinsic growth rates," *Population Studies* 40(3): 343–351.

———. 1988. "Reply to Wachter," *Population Studies* 42(3): 495–501.

Sanderson, Warren C. 1995. "Predictability, complexity, and catastrophe in a collapsible model of population, development, and environmental interactions," *Mathematical Population Studies* 5(3): 259–279.

Schoen, Robert and Young J. Kim. 1998. "Momentum under a gradual approach to zero growth," *Population Studies* 52(3): 295–299.

Stoto, Michael A. 1983. "The accuracy of population projections," *Journal of the American Statistical Association* 78: 13–20.

United Nations, Department of International Economic and Social Affairs. 1992. *Long-Range World Population Projections: Two Centuries of Population Growth 1950–2150*. New York.

———. 1998a. *World Population Projections to 2150*. New York.

———. 1998b. *World Population Prospects: The 1996 Revision*. New York.

———. 2000. *Long-Range World Population Projections: Based on the 1998 Revision*. New York.

Wachter, Kenneth W. 1988. "Age group growth rates and population momentum," *Population Studies* 42(3): 487–494.

Whelpton, P. K. 1936. "An empirical method of calculating future population," *Journal of the American Statistical Association* 31: 457–473.