CICRED'S SEMINAR, Paper for presentation at the Global Seminar on "Age-Structural Transitions: Demographic Bonuses, but Emerging Challenges for Population and Sustainable Development" organized by Committee for International Cupertino in National Research in Demography (CICRED), Paris, Republic France, 23-26 February 2004.

Demographically based global income forecasts up to the year 2050

Bo Malmberg and Thomas Lindh
Institute for Futures Studies
Stockholm

Preliminary version to be presented at the Symposium on Demographic Change and Global Income Growth in the 21st Century, Krusenberg 7-9 January 2004

Abstract

Demographic projections of age structure provide the best information available on long-term human resources and demand. In current data fairly robust correlations with GDP and GDP growth have been discovered. In this paper we use these two facts and study the forecasting properties of demographically based models. Extending the forecasts to 2050 suggests that due to fertility decreases poor countries of today will start to catch up with developed economies in which the growth process will stagnate due to the growth of the retired population.

1 Forecasting global growth by age structure projections

Will the current income gaps between the developed and the developing world persist or will we have a long-run convergence between different parts of the world? This question has been thorougly researched in cross-country convergence studies over the last two decades where a rather pessimistic view has emerged. Given current structures the indication is that convergence only takes place within certain clubs of countries like the OECD, while the poorest part of the world is actually diverging relative to the developed countries. Of course, the conclusion is not that the poverty of some countries is ordained by destiny, but rather that drastic internal structural change combined with international aid and pressures is needed in order to bring an end to the extreme poverty in the world. Based on rather fragile empirical evidence factors such as technological change, human capital, free market institutions and strict budgetary discipline, just to mention a few, have been more or less emphasised in the debate. Long-term projections of income have been hard to achieve from these hypotheses, since the implied driving forces are hard to measure and even harder to project into the future.

This paper draws on another fundamental force driving growth: demographic change. This is a variable that is relatively easy to measure, and can be projected into the future with at least reasonably low uncertainty for long periods of time. The focus in this paper is on evaluating the forecasting performance of demographically based models. We estimate demographically based models of income as a function of age structure, urbanization and life expectancy, and test their statistical properties both in and out of sample. Finally we use UN projections of age structure to derive forecasts of global income levels and the evolution up to 2050.

Our conclusion is that such forecasts behave reasonably well in predicting income and in the long term shows promise to be a much safer method than any available alternative method. The long-run forecasts also indicate that, as a consequence of the demographic age transition, primarily decreasing shares of children and increasing longevity, developing countries (with the possible exception of AIDS inflicted countries in sub-Saharan Africa) will tend to grow faster and thus start converging to the income levels of the developed world.

Although the contribution of this paper is intended to be in forecasting it is of some importance to make clear that the use of age structure variables for income forecasting does have an extensive theoretical and empirical foundation, so the first section of the paper is devoted to an account of the facts and theories behind the demographic age transition and its impact on the economy. A brief discussion of econometric problems in Section 3 prepares our choice of a two pronged strategy with one homogeneous model in Section 4 and one heterogeneous approach in Section 5. The concluding section argues that our forecasts inspire confidence in a rather bright global future.

2 Starting point: The demographic transition

What we now know as the demographic transition was discovered in the second quarter of the 20th century by demographers observing the variation of death rates and birth rates across countries and over time. The finding was that in the Western cultural area there has been a general process of change from the relative population stability at high levels of mortality and fertility to the slowing growth or actual decline of population numbers at low levels of mortality and fertility which characterized the interwar period. The initial effect was a decrease in mortality, leaving largely intact the large family pattern. For reasons that are not entirely clear a small family pattern became not only an ideal but a goal realized by increasing proportions of the population. Birth rates fell rapidly, at first in cities, later in the surrounding rural areas.

The prerequisite for this observation was increasing efforts to assemble demographic data from the 19th century and onwards that had made available time-series stretching over multiple decades or longer. Researchers in the inter-war era, thus, could base their analysis of population trends on a comparatively rich empirical material. However, the concept of demographic transition was not only used as a name for an empirical pattern. Instead the researchers who introduced the term saw these trends in mortality and fertility as the expression of an underlying social transformation process. Moreover, their conclusion from observing similar demographic trends in different countries was that the process of demographic transition would not be restricted to Europe but would spread also to other parts of the world. In fact, Irene B. Taeuber's original formulation of the transition idea appears in a paper that analyses the population development of Southern and Eastern Asia.

In 1945 the idea that the world would witness demographic transitions throughout Asia, Africa and America was basically a scientific hypothesis yet to be verified or falsified. The key assumption behind this hypothesis was essentially a belief that the technologies giving rise to lower mortality would continue to spread to those parts of the world that still suffered from high mortality levels. Developments after 1945 have demonstrated that this assumption was well-founded, although even proponents of the transition hypothesis soon became surprised by the "amazing decline of mortality in underdeveloped areas" (Kingsley Davis, 1956).

Today, almost 60 years later we can conclude that the demographic transition hypothesis has proved to be correct. First, the mortality decline observed in the Western world and in parts of Asia before 1945 has indeed spread to practically all parts of the world. Before 1945 the life expectancy at birth in many non-Western populations were appreciably below 40 years. In India 1941-50 and also in Mauritius 1942-46, for example, the expectation of life at birth has been estimated to be only 32 years. Similarly, the Bantu population of South Africa had an estimated life expectancy at birth of 33.6 years in 1936-46. Life expectancy in China was lower still: data on the Chinese farmer population in 1929-31 indicate values of 28.1 years in North China and 24.2 years in South China. In other parts of the world the picture was somewhat brighter with male life expectancy at birth at 40.9 years in Chile (1940), 41.1 years in Taiwan (1936-41), 42.1 years in Egypt (1937) , 44.5 years in Trinidad and Tobago (1930-32), 45.1 years in Puerto Rico (1939-41), 45.8 years in Venezuela (1942-42), and 46.9 years in Japan (1935-36).

After the Second World War life expectancy in these countries has risen dramatically. The most impressive gains have been made in China. In 1999 the World Bank estimated a life expectancy at birth of 70 years. That is a gain of 44 years compared with the life expectancy for the Chinese farmer population around 1930. Japan, Taiwan and Chile have gains in life expectancy around 35 year. The gains in India, Puerto Rico, and Trinidad Tobago are around 30 years whereas Egypt and Venezuela have gained about 25 years. Also in South Africa there was a 23 years gain in life-expectancy up to 1980 but nine years of this gain has been lost between 1980 and 1999 primarily because of AIDS related deaths. Of the large world regions (East Asia & Pacific, Europe & Central Asia, Latin America & Caribbean, Middle East & North Africa, South Asia, and Sub-Saharan Africa) only Sub-Saharan Africa now, despite a substantial increase, has a life-expectancy at birth below 50 years. In the rest of the world life expectancy is above 60 years.

Also the second prediction of the demographic transition hypothesis—that birth rates for an extended period will stay high when the death rates have begun to decline—has proved correct. Using UN data (World population prospects—the 2000 revision) this process can be most easily traced in African countries like Kenya, Ethiopia and Nigeria where falling death rates and relatively stable birth rates lead to accelerated population growth during the 1950s, 1960s, and 1970s. However, the starting point of the UN data is the 1950-54 period which means that this data set doesn't capture the dramatic decline in death rates that followed in the first few years after 1945. According to a study by Kingsley Davis, based on "18 underdeveloped areas, not chosen because they had unusual declines in mortality but because they were representative of different areas and had fairly constant boundaries and a relatively continuous series of registered death statistics" the total decline in the crude death rate between 1935 and the 1950-54 period can be estimated to 58%, with the largest part of the decline coming after 1945. This large decline up to the 1950-54 period implies that the UN data gives us a picture of what is happening to the population directly after a sharp mortality reduction. If Kingsley Davis' estimate is correct—and assuming stable birth rates— the population growth rate in the less developed countries (UN, 2001 definition) would have accelerated from 0.3% annually in 1935 to 2.1% in 1950-54. This growth rate was the result of a crude birth rate of 44.6 and a crude death rate coming down from 41.2 to 24.1. Declining death rates, thus, led to accelerated population growth exactly as the Princeton demographers had predicted.

Towards the end of the 1960s the death rate of the less developed countries had declined a further 9 points but the birth rate had only declined by half us much. Population growth in the less developed countries, thus, had accelerated further and now stood at 2.5% annually. The end of the 1960s, then, was a time when the transition hypothesis did not seem to be correct. Had not the hypothesis assumed that declining death rates would be followed by lower birth rates? And, what would happen if birth rates remained high? This, thus, was also a good time for dooms-day predictions concerning the future of the earth and its population. Not everyone lost their nerves, though. The UN forecasters continued to base their projections on the assumption of a slow but continuing fertility decline that would bring down the population growth rate. And—as it turned out—the UN forecasters were right. Between 1965-69 and 1970-74 there came an important shift in the population trend as growth rates in Eastern Asia and South-Eastern Asia started to decline. This had happened in Korea and in Latin America already some years earlier but it was with the Asian turn-around that the shifting trend became globally significant. Later in the 1970s population growth rates started to decline in Southern Africa and in the mid 1980s, population growth began to slow also in North Africa, Western and South Central Asia. In the early 1990s, it happened in Western Africa and in Eastern Africa too, although in the latter case this slowdown was partly due to an increase in the death rates. And at the end of the 1990s, Middle Africa also seems to have passed its peak in the population growth rate.

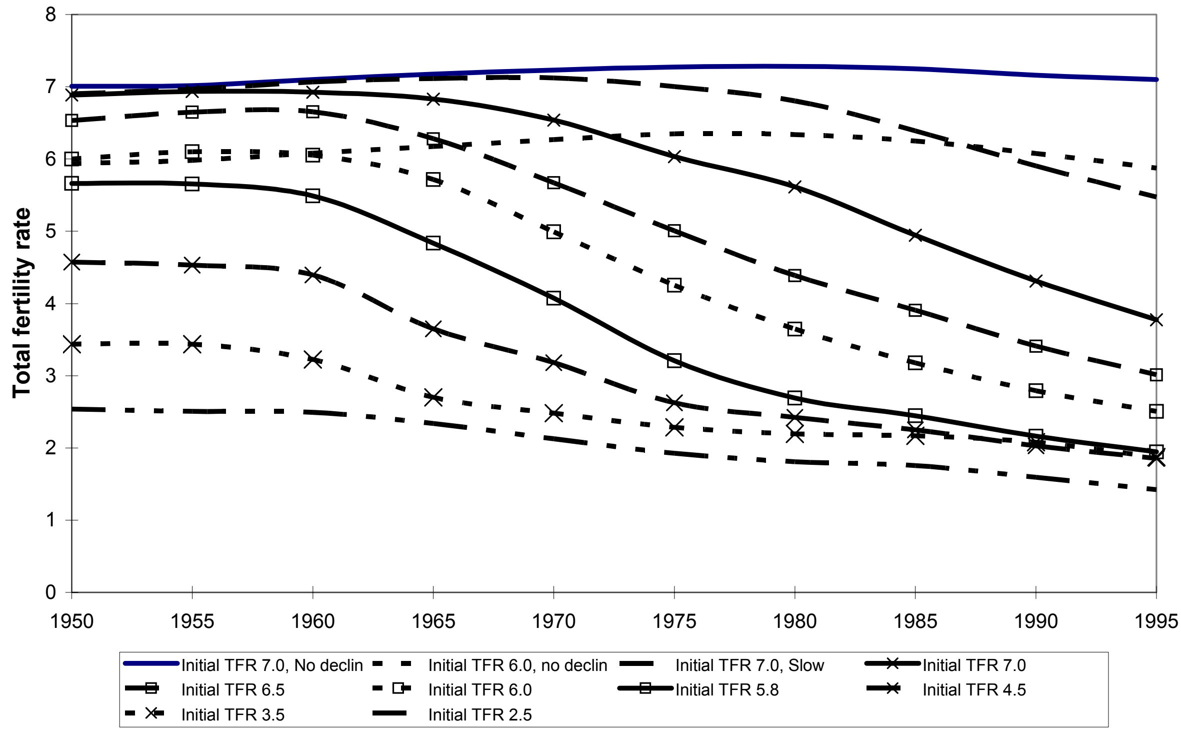

Figure 1: The evolution of total fertility rates 1950 to 1995 from different initial positions. Source United Nations Population Statistics.

A more detailed view of factors behind the downward trend in population growth comes from a look at how the total fertility rates have developed after 1950. In Figure 1, the most important patterns of fertility change in the world between 1950 and 2000 are outlined. Countries with a similar pattern of fertility change have been grouped together and for each year the mean fertility rate in each group has been calculated. Figure 1 shows clearly that fertility rates differ much across the world and that there are different patterns of fertility change. However, there are also strong common trends. In all groups, except two, fertility has gone down since 1950 although at different rates.

Only two groups have a record of no significant decline in mortality. In these groups we find a number of very poor countries many of which have experienced strong social disruption: Afghanistan, Angola, Burkina Faso, Burundi, Ethiopia, Liberia, Malawi, Mali, Niger, Somalia, Uganda, and Yemen all belong to the group where the total fertility rate still is close to 7 children per woman. Also for Bhutan, Cameroon, the Central African Republic, Chad, Congo, the Democratic Republic of the Congo, Equatorial Guinea , Gambia, Guinea-Bissau, Lao People's Democratic Republic, Mauritania, Mozambique, Namibia, and Pakistan there has been very little change in the total fertility rate, almost constant around 6 children per woman.

Another group for which the fertility rate in 1995-1999 still is very high, around 5.4, is made up by Benin, Comoros, Côte d'Ivoire, Djibouti, Eritrea, Ghana, Guatemala, Guinea, Iraq, Jordan, Kenya, Libya, Madagascar, Maldives, Nigeria, the Occupied Palestinian Territory, Oman, Rwanda, Saudi Arabia, Senegal, Solomon Islands, Sudan, Swaziland, the Syrian Arab Republic, Togo, the United Republic of Tanzania, Zambia, and Zimbabwe. For this group starting with a TFR of 7 in 1950, there has, however, been an appreciable reduction in the fertility rate, especially since 1980. A trend which, if we are to judge from the development of other groups, can be expected to continue.

The group consisting of Algeria, Bangladesh, Belize, Bolivia, Botswana, Cape Verde, Egypt, Honduras, Iran (Islamic Republic of), Mongolia, Morocco, Nicaragua, Philippines, Qatar, Samoa, Tajikistan, the United Arab Emirates, Vanuatu, and Western Sahara in 1950 also had an average TFR of 7 but by the end of the 1990s the average rate had been reduced to 3.8, a very considerable reduction

The three groups Bahrain, Brunei Darussalam, Colombia, Costa Rica, the Dominican Republic, Ecuador, El Salvador, Fiji, India, Kuwait, Malaysia, Mexico, Myanmar, Peru, Saint Lucia, South Africa, Tunisia, Turkey, Turkmenistan, Uzbekistan, Venezuela, and Viet Nam starting at a TFR of 6.5, followed by Albania, Brazil, French Polynesia, Guyana, Indonesia, Lebanon, Panama, Suriname, and Thailand starting at TFR at 6.0, and finally Azerbaijan, Chile, China, Guadeloupe, Martinique, Mauritius, the Republic of Korea, Réunion, Sri Lanka, Trinidad and Tobago also starting at a TFR of 6.0 have essentially tracked one another with respect to the fertility decline and arrived at mean TFR rates of 3.0, 2.5, and 1.9 respectively in the late 1990s.

In the group Armenia, Bahamas, Barbados, Bosnia and Herzegovina, Hong Kong, Macao, Cuba, Ireland, Israel, Kazakhstan, the Netherlands Antilles, Puerto Rico, and Singapore with a mean TFR in 1950 of about 4.6 has fallen to 1.9. The last two groups, with 1950 TFRs of 3.5 and 2.5 consists largely of developed countries that already in the beginning of the period had completed the classical fertility transition.

What we can see in Figure 1, then, is a quite impressive corroboration of the demographic transition hypothesis. To be sure, some countries have yet to enter the phase of fertility decline and others have a long way to go before they reach replacement fertility, that is, a level of TFR slightly above 2 that in the long run keeps a population constant. But the patterns of fertility change during the last fifty year strongly suggests that a continuing fertility decline in today's high fertility countries is to be expected.

The Princeton hypothesis of a demographic transition affecting countries all over the world, thus, has turned out to be one of the most successful predictions made by social scientists in the 20th century.

2.1 The age transition

In the 1990s different scholars have pointed out that population growth will not be evenly distributed across the different age groups during the demographic transition. On the contrary, as the demographic transition unfolds the growth rate of different age groups will follow an uneven pattern. In some periods during the transition population growth will be concentrated to the youngest part of the population. In other periods the young adult, middle aged, or old age population segments are increasing most rapidly. This implies that population growth will give coherent effects on per capita income only if it is correct that the economic effects of population growth are the same irrespective of whether it is children, young adults, middle aged adults, or old age adults that are increasing in numbers. Even as a scientific simplification that seems to be a rather far fetched hypothesis.

2.1.1 How does age structure change during a demographic transition?

This question can be answered both empirically and through the use of demographic models. As it turns out, both methods give very similar results, at least over long periods with moderate levels of immigration and emigration. The typical pattern of a mortality decline followed—after a lag—by declining fertility will generate a very distinct pattern of age structure change. This pattern can be observed in most countries affected by a demographic transition. One way to summarize is to distinguish four different phases of population growth during the demographic transition:

The first phase following the onset of mortality decline is characterized by an increase in the number of children. The primary reason for this is that in high mortality regimes it is among newborns, infants and young children that the death toll is especially high. So when mortality comes down it is to a large extent the lives of the very youngest that are spared. In time, this increase in the number of children will also increase the young adult fertile population and—as long as fertility rates are unchanged—this will, by inducing more births, further accelerate the increase in the number of children. This first child-rich phase of what can be called the age transition will continue as long as the fertility rate remains high. The fertility decline eventually slow down the increase of the child population.

The second phase is characterized by an expansion of the young adult population. The mechanism behind this expansion is simply that with declining mortality, and later an increase in the number of births, the number of surviving individuals in each cohort will increase. As these ever larger cohorts reach adult ages the young adult population will start to expand. Because it takes time for newborns to reach adult age the young adult phase will start 15-20 years later than the child-rich phase and it will continue 15-20 years after the expansion of the child population has stopped.

The third phase is characterized by an expansion of the middle aged population. This phase starts when the cohorts enlarged by mortality decline and increases in the number of births reach the middle ages. Depending on how one defines "middle age" this expansion is initiated 20-30 years after the young adult phase starts. It thus takes four to five decades or more before the mortality decline of a demographic transition produces an appreciable increase in the number of middle aged.

The fourth phase, finally, is reached when the enlarged cohorts reach retirement and is characterized by an expansion of the old age population. Often, this expansion starts after the fertility rates have dropped to a low level and this means that the old age phase, and also at least part of the middle age phase, lies outside the time span that we normally consider when we talk about the demographic transition. In Sweden, for example, the classic transition was complete in the 1930s when the total fertility rate had come down to, or even below, the replacement level. The middle age phase of the age transition, however did not end until about 1970 and the old age phase continued for another two decades.

Two things should be noted about the age transition. First, as pointed out above, the age transition as a phenomenon extends for a considerably longer period than the classical transition. This implies that the demographic transition, because of population momentum, will have social and economic effects also when there are no longer any current changes in the vital rates. Second, it is important to remember that, to some extent, the phases of the age transition will overlap. Depending on for how long later cohorts keep getting larger than earlier cohorts this time of overlap may be short or long.

The fact that countries with large differences in terms of historical traditions, climate, and geography have gone through demographic transitions that in terms of changing mortality and birth rates have been similar is certainly a challenge to social science. In this study—where the purpose is not to explain why demographic transitions have occurred but to utilize the effect the transitions have on income growth—the existence of numerous instances of demographic transitions in very different contexts is, however, a great advantage creating variation and correlations that can be exploited to forecast income.

2.1.2 The global geography of age transitions

How then, has the age transition affected the population structure in different parts of the world. One way to illustrate this is to use the UN data that gives both an account of estimated age structures for the 1959-2000 period and population structure forecasts for the 2000-2050 period. Given that the UN provides data for more than 200 different regions and 17 to 20 five year age groups for a 100 years period implies that it can be hard to get a grip on how the age structure is changing. However, by the use of cluster analysis it is possible to extract a smaller number of characteristic age structures that will capture the typical age distributions that different population will demonstrate as they go through the age transition.

These typical age structures can be associated with different phases of the age transition. Figure 2 shows, using these typical age patterns, how age structure in the world will change in the next 50 years according to the UN forecast. The remarkable thing is that there will be very drastic changes in the age structure during the next 50 years, especially in many regions with a large share of developing countries.

The question now is if these projected changes in the age structure can be used to forecast future levels of per capita income? One indication of this possibility is given in Figure 4 that shows the per capita income level in countries assigned to the different age structure types of the cluster analysis. Old countries are rich countries whereas young countries are poor countries.

The UN predicts that many of the young countries will become old countries.Does this imply that they also will become rich? If we believe that the income patterns in Figure 4 will remain true also in the next 50 years then Figure 3, showing how regions of the world passes from one type of age pattern to another, also gives an indication of how per capita incomes will change.

From recent empirical growth literature we also know that there is a robust association between age structure and economic growth, where dependent groups tend to depress growth and the working age groups tend to enhance it. At present it is not clear which mechanisms these associations rely on. Obviously there should be a direct labor supply effect making the economy grow faster as the working age groups increase the fastest, but just

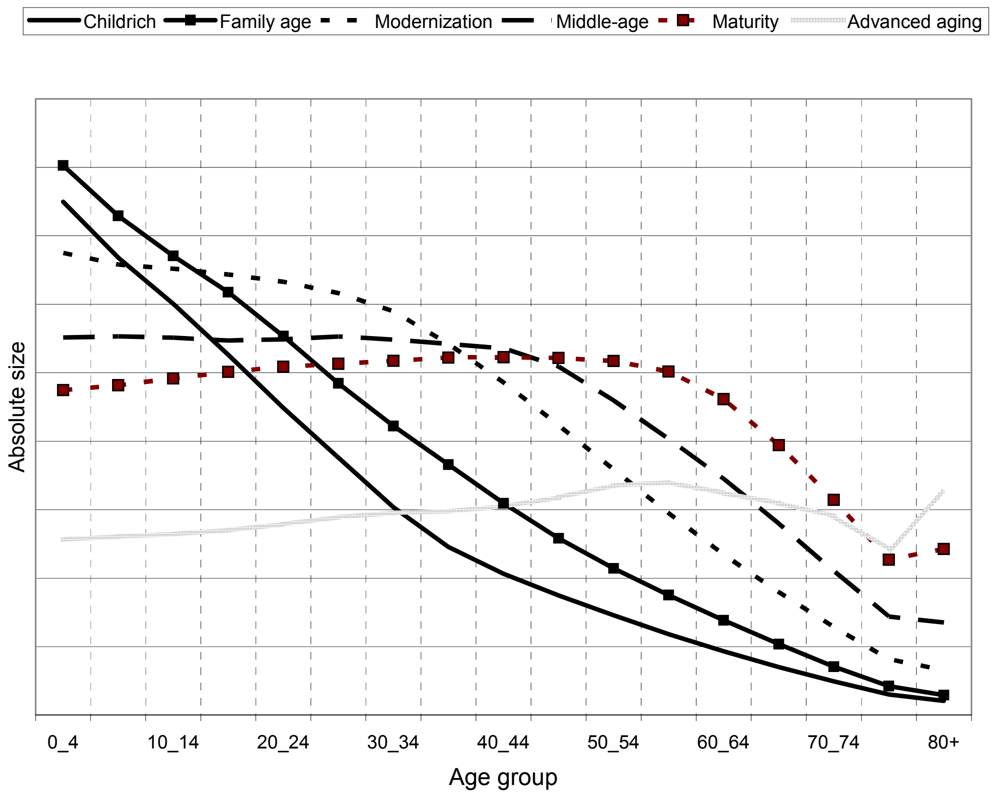

Figure 2: Five typical age structures corresponding to different phases in the demographic transition. The sixth grey line is the projected age structure for an advanced aging society.

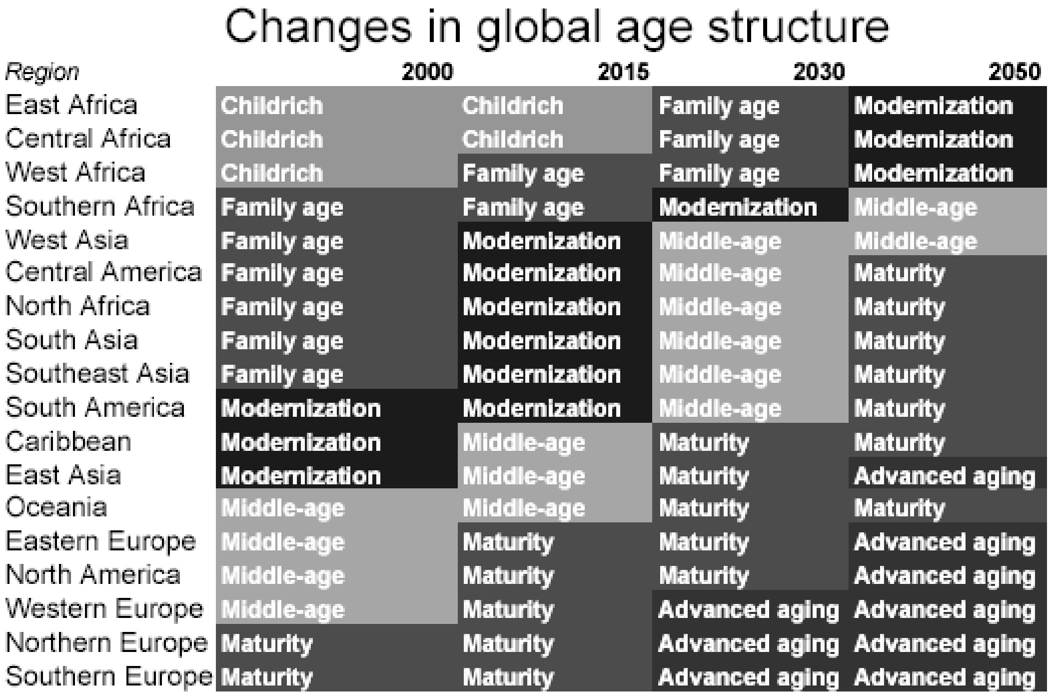

Changes in global age structure

Region	2000	2015	2030	2050
East Africa	Childrich	Childrich	Family age	Modernization
Central Africa	Childrich	Childrich	Family age	Modernization
West Africa	Childrich	Family age	Family age	Modernization
Southern Africa	Family age	Family age	Modernization	Middle-age
West Asia	Family age	Modernization	Middle-age	Middle-age
Central America	Family age	Modernization	Middle-age	Maturity
North Africa	Family age	Modernization	Middle-age	Maturity
South Asia	Family age	Modernization	Middle-age	Maturity
Southeast Asia	Family age	Modernization	Middle-age	Maturity
South America	Modernization	Modernization	Middle-age	Maturity
Caribbean	Modernization	Middle-age	Maturity	Maturity
East Asia	Modernization	Middle-age	Maturity	Advanced aging
Oceania	Middle-age	Middle-age	Maturity	Maturity
Eastern Europe	Middle-age	Maturity	Maturity	Advanced aging
North America	Middle-age	Maturity	Maturity	Advanced aging
Western Europe	Middle-age	Maturity	Advanced aging	Advanced aging
Northern Europe	Maturity	Maturity	Advanced aging	Advanced aging
Southern Europe	Maturity	Maturity	Advanced aging	Advanced aging

Figure 3: Different regions are categorized according to the typical age structure they will have achieved in the future according to UN projections.

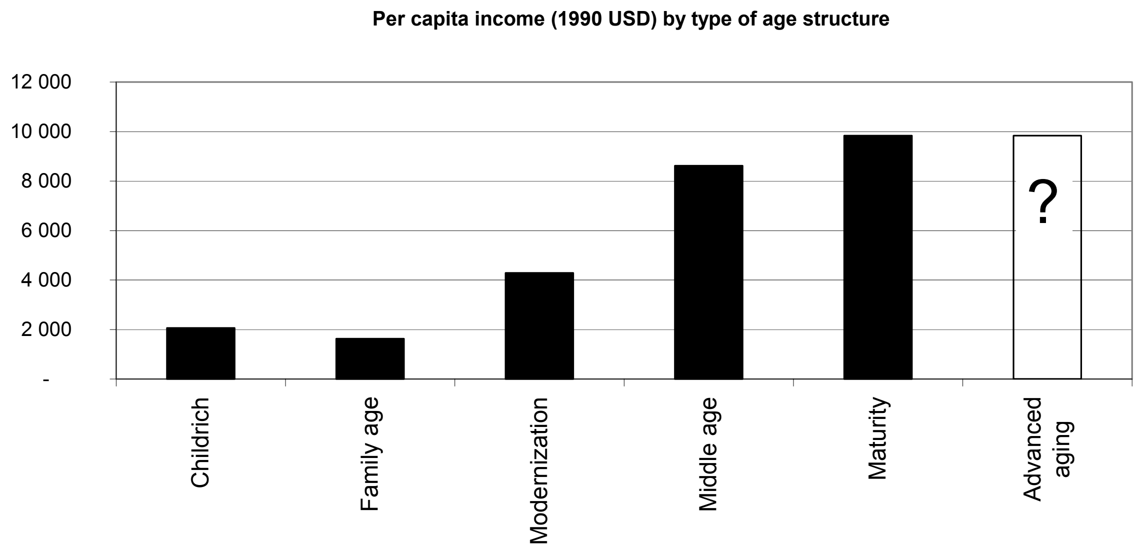

Figure 4: Average income per capita in 1990 USD in different age structure clusters. For the advanced aging stage we have no data.

as obviously the home bias in the relation between domestic saving and investment, should make capital accumulation easier when a large part of the population are in the middle age phase. There are indirect effects through the public sector finances since revenue and public consumption and transfers are asymmetrically distributed with respect to age. There are also more long-run mechanisms where for example increasing longevity provides incentives both for retirement saving and increased human capital accumulation. Several other factors having to do with age-specific demand and income patterns could also be invoked as parts of the explanation. Here we do not go further into this discussion. but are satisfied that there are good solid reasons why we should expect a correlation of GDP growth and age structure.

The rest of this paper is devoted to study whether the basic idea that there is a correspondence between age structure and income level can be put to practical use for long-run forecasting.

3 Forecasting problems

There are a number of problems to deal with once the basic specification of an acceptable regression equation of GDP on demographic variables has been accomplished. We briefly comment on the following six problems.

1. The uncertainty inherent in the projected demographic variables.
2. Collinearity of the explanatory variables.
3. Model heterogeneity across countries.
4. Structural breaks.
5. Model heterogeneity over time.
6. Non-stationarity of time series.

Here we only mention the first point since it is, of course, of crucial importance for validating the forecasts. But that uncertainty is to a first-order approximation mainly a question of the assumptions made on fertility, migration and mortality in the demographic projections, and here we wish to concentrate on the other points. Ideally probabilistic demographic forecasts should be plugged into the model to deal with this issue in an explicit way.

The second point is not necessarily a problem. Provided that the correlation patterns between demographic variables are stable, the group of variables should forecast well even if there are collinearity distortions of the individual coefficients. However, we know that the projections imply that these correlation patterns actually will change, and thus it becomes necessary to trade off collinearity problems against the precision of the demographic information. Thus in the specification part neither the fit of the equations nor the plausibility of individual coefficients are central concerns. The main concern is that the coefficient estimates yield stable forecasts for a reasonably long period.

The third point is a common worry when working with cross-country panels. Is it legitimate to assume a homogeneous model for such a variety of countries, different in size, location, history, institutions and natural resources? In some sense, every country is a unique economic system related to its neighbours by a multitude of different relations. There are countries like Luxembourg, which cannot really be pictured as an independent national economy. There are vast differences between a country like Sweden and a country like Lesotho which would seem to invalidate any comparison. Still the aim of this exercise is exactly to try to extract some general pattern of how demographic change relates to economic change. Using the panel approach allows us to control for unobservables that are constant over the estimation period, but we may also need to account for differences in the model across countries. In the extreme we might even consider individual models for every country. That would, however, not be an efficient use of the available information, not to mention that there are a substantial number of countries where our information is scant and/or of questionable quality. To obtain reliable forecasts for the individual countries as well as for the global aggregate it may, however, be necessary to control for some differences.

The fourth point about structural breaks in the data is in a sense the corresponding problem of heterogeneity in the model across time. Identifying and controlling for structural breaks within countries thus also becomes a concern. Some events liable to cause structural breaks may be worldwide influences like oil crises, some of more regional importance as the opening of free markets within the EU, some only affecting specific countries like the genocide in Rwanda. Breaks with a very temporary influence may be controlled for by dummies. Others like the breakdown of the Bretton-Woods system are likely to have much more long-lasting effects, and others like the breakdown of the Soviet Union also have important spin-off effects on neighbouring countries. Our aim can hardly be to model all possible peculiarities that can arise, but only to try and identify those of significance for the forecasting model.

The fifth point has already been treated above as far as it concerns individual countries, but there are some difficulties at the global scale that we wish to draw attention to. We start from the assumption that our cross-country observations are drawn from a data generating process that is at least partly common to all countries, viz. the demographic transition and the concurrent industrialization and aging of the population. However, our observations are drawn only from a limited period of this transition for each country. Some countries are just in the beginning of this process while others are entering a second transition into an ageing society. While the observations from more developed countries provide information to forecast (to some extent at least) the evolution of the less developed countries our sample does not contain information regarding the ageing society and how it will adapt to a rising dependency burden. There is thus an asymmetry in the information available which in some sense is analogous to the well-known difficulty of identifying cohort, period and age effects in microeconometric studies due to the exact linearity of such effects.

The sixth point is a standard problem of all time series analysis, since the regression of non-stationary series on each other can cause spurious regression. However, recent work demonstrates that this problem is substantially ameliorated in a panel context. Moreover, recent work shows that GDP and age structure are likely to be cointegrated implying that OLS estimates of the relation are in fact superconsistent.

3.1 Two estimation strategies

There is always a trade-off of drawbacks and advantages in different approaches to forecasting. Our target is to devise a forecasting model that only relies on variables which can be independently projected. First of all, this makes structural modelling less interesting since in the end we will have to rely on a reduced model anyhow. Second, degrees of freedom are limited in spite of using annual data on a world panel and overfitting can easily become a problem, especially since high degrees of collinearities prevent the use of too finely divided age groups. Our choice here has been to pursue two different paths. One approach emphasizes the time series modelling using only five different age groups and disregard the heterogeneity issues across countries except for fixed effects. The other approach emphasizes the heterogeneity issue and takes the time series issues less seriously.

Thus we have chosen not to pursue the first two problems in this paper. Points 3-5 will be adressed in a manner conformal to the demographic transition theory that has inspired our attempt to use demographic information for long-term forecasting but we will here abstain from modelling structural breaks at the individual country level.

3.2 Time series approach

The basic model for this approach is a regression in levels of the logarithm of per capita GDP, y, on the logarithms of age shares, a

$$y_{it} = \alpha t + \sum_{k=0-14}^{65+} \beta_k a_{kit} + \eta_i + \varepsilon_{it} \qquad k = 0-14, 15-29, 30-49, 50-64, 65+ \tag{1}$$

Interpreting this we assume that GDP per capita can be described by a Cobb-Douglas index of age shares and an exponential trend that is intended to capture technological change. The logarithmic form ameliorates problems with heteroskedasticity and also makes it possible to include the whole distribution of age shares in the fixed effect estimation since the exact linear dependence of the variables is broken. Based on previous work an aggregation of age groups to children 0-14, young adults 15-29, mature adults 30-49, middle aged 50-64, and old age 65+ is known to work well in growth equations for the OECD. This corresponds roughly to the age intervals in which humans are first dependent on parents, second finding their place in adult life and forming a family, third raising their family, fourth preparing for retirement and fifth retiring. The exact limits, of course vary both with time and culture, as well as the institutions that govern the economic effects of each age group. As a rough approximation making a trade-off against collinearity this is a reasonable starting point.

3.3 Heterogeneity approach

There are numerous reasons why we should expect the coefficients of age shares to vary over time and countries, for one example consider length of education. This variable in turn is dependent on expected returns which are at least partly determined by life expectancy. Therefore, interacting the age shares with life expectancy at birth, $e0$, could correct for some of the heterogeneity across countries that depends on stage of development.

$$y_{it} = \alpha t + \sum_{k=0-14}^{65+} (\beta_k + \gamma_k e0_{it}) a_{kit} + \eta_i + \varepsilon_{it} \tag{2}$$

without sacrificing the basic requirement of using only demographic variables. We develop this further by also attempting to take account of population density and urbanization.

4 Fixed age share parameters

In this section we investigate the behavior of an age model with fixed age share parameters. Our aim here is to study the properties of a homogeneous model (apart from intercepts) in predicting income levels. We proceed in three steps. First, we investigate the specification using the whole sample of 111 countries with data at least 1961-1994 and in some cases for the full period 1950-1996. Second, we examine the behavior of the forecasting model out-of-sample. Third, the 2050 forecasts on UN projections are presented and commented.

4.1 Specification tests

Our aim of long-run forecasting prevents us from using other economic variables that are hard to forecast as explanatory variables, but we are at liberty to use deterministic trends, lags of the dependent variable and different transformations of demographic variables. In this sub-section we evaluate some different possible specifications. First, in Table 1 we present a benchmark pooled model in column 1. The residual tests clearly indicate that we have unaccounted time and country effects. Although the Q-test reported is not fully appropriate in this context it gives some indication of the strength of the auto-correlation in the residuals which is monotonously tapering off, indicating an AR(1) model for these residuals. However, scrutinizing the individual country residuals it becomes very obvious that the autocorrelation is far from similar over countries and thus fitting a common autoregressive errors model will be misleading.

Using a quadratic trend in the specification does not eliminate time effects in the residuals but makes them considerably smaller. Individual effects are, however, not much affected. Thus, the next two columns compare a fixed and random effects specification of this variant. The differences are small but the efficiency gain that one should get from using random effects seems to be very small or non-existent.

The serial correlation of errors is, however, still highly significant indicating that omitted variables bias may be substantial, so after some experimentation we decided to include three lags of the income variable to correct for this in column 4 and 5. The efficiency gain from using random effects still seems negligible, and we conclude that the simpler and more robust fixed effects estimator is preferable.

The pattern of lag coefficients suggests that the hypothesis of a common unit root cannot be dismissed, but the residuals nevertheless are reasonably stationary. However, this is clearly an important point to investigate further.

Dep var	OLS	Country effects		Including lags	
log(GDP/cap)		Fixed	Random	Fixed	Random
Constant	**10.30**		**5.952**		**0.396**
	(10.4)		(17.0)		(3.94)
$\log a_{0-14}$	**-1.427**	**-1.697**	**-1.691**	-0.050	-0.012
	(5.80)	(20.5)	(20.5)	(1.89)	(0.45)
$\log a_{15-29}$	**0.972**	**-0.178**	**-0.176**	0.011	0.024
	(4.53)	(2.79)	(2.77)	(0.54)	(1.28)
$\log a_{30-49}$	**1.029**	0.127	0.131	**0.055**	0.055
	(4.53)	(1.88)	(1.95)	(2.69)	(2.81)
$\log a_{50-64}$	**-1.094**	**0.153**	**0.166**	0.023	0.023
	(4.54)	(3.06)	(3.33)	(1.47)	(1.62)
$\log a_{65+}$	**1.199**	-0.072	**-0.052**	-0.009	0.009
	(10.5)	(2.01)	(15.0)	(0.77)	(0.89)
Trend	**0.039**	**0.048**	**0.048**	**0.002**	**0.001**
	(6.82)	(42.7)	(42.1)	(4.80)	(2.44)
Trend squared	**-0.00073**	**-0.00068**	**-0.00069**	**-0.00004**	**-0.00002**
	(6.87)	(33.0)	(33.5)	(4.96)	(3.14)
y_{t-1}				**1.054**	**1.068**
				(70.3)	(71.8)
y_{t-2}				**-0.116**	**-0.118**
				(5.38)	(5.48)
y_{t-3}				0.018	0.026
				(1.20)	(1.76)
$\bar{R}^2$	0.668	0.961	0.961	0.997	0.997
F time effects	**2.08**	**2.29**	**2.25**	**4.06**	**3.98**
F country eff	**186.6**	–	–	–	–
Ljung-Box $Q\,(6)$	**22672**	**13404**	**13428**	6.38	7.02

Table 1: Absolute t-values in parentheses. Bold face indicates that the estimates are significant on the 5 percent level. Newey-West robust errors allowing for three lags are reported for the pooled model. Standard errors for the panel model are adjusted for the unbalanced panel.

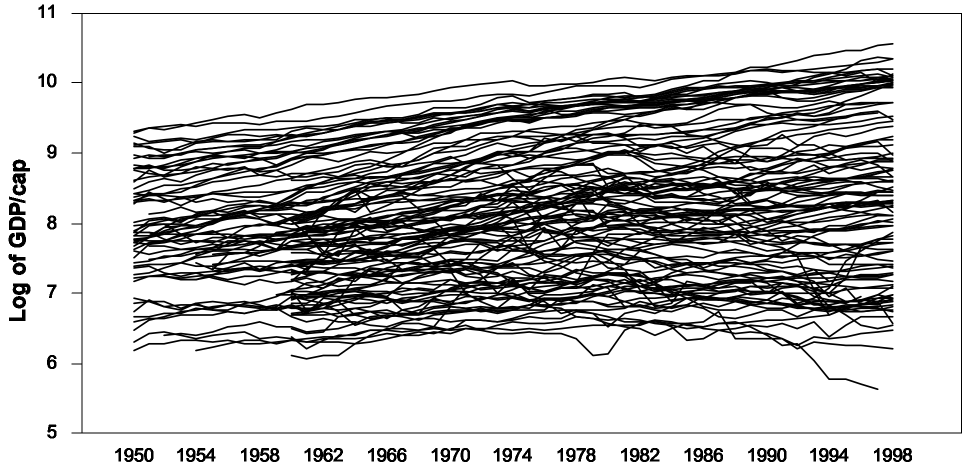

Figure 5: Time series for 111 countries of the dependent variable the log of GDP per capita. In the scale 6 corresponds to about \$400 and 10 to about \$22 000 in 1996 USD.

The magnitude of the estimates decrease as expected when we include the lags but the decrease is uneven over the age coefficients, such that the basic hump-shaped pattern we would expect with positive effects from the working age groups actually becomes steeper when we include the lags, and more so in the fixed effect case than in the random effects case.

Plotting the log of GDP per capita for all the time series (in Figure 5) it becomes apparent that although most of the individual country series are trended, a common linear trend is not obviously present corroborating the previous conclusion that a non-linear trend is needed to capture the common deterministic trend. Most of the series also seem far too smooth to be generated by a common stochastic trend (with asymptotically infinite variance).

The question then must be raised whether it is meaningful to pursue a common homogeneous model or turn to an attempt to model the country heterogeneity. The residuals of the models in Table 1 exhibit rather strong non-normality, mainly due to high kurtosis in the distribution. This might be generated by outliers or structural breaks but it could also be a sign of a non-standard error distribution. Considering the purpose of this paper it would be an advantage if a parsimonious homogeneous forecast equation could be made to work.

Including time effects instead of trends stabilizes the age coefficient estimates somewhat (not reported). But neither time effects nor quadratic trends are very desirable to have in a forecast specification. The former because time effects can be predicted only by ad hoc assumptions, the latter because a quadratic trend is likely to depress income too much in a long-run exercise since the influence of the quadratic term will dominate the whole equation as we project more than fifty years ahead. The fat error distribution tails also raise a suspicion of structural breaks that may distort both deterministic and stochastic trend estimates. Some experimentation showed that this suspicion was valid, generating unreasonably gloomy forecasts into the negative region for many countries.

We therefore turned to a demographic variable that could be projected but still catch some of the non-linear trend. The scatterplot of logged GDP/capita versus the log of life expectancy at birth in Figure 6 points to a kinked linear relationship that becomes more well determined and steeper at higher GDP levels. This suggests that life expectancy (which is also demographically projected) can be used as a substitute for the trend. The picture also suggests that this relationship might hold only for high incomes. In Table 2 column 1 we report a model where a linear relationship with life expectancy is imposed for observations with per capita incomes above $\exp(8.5) = 4914$ USD per capita.

Theoretically a model with life expectancy at birth entering in this way could be a consequence of longer life expectancy leading to on average longer education necessary to maintain the rate of technological change in production. The chosen break point would imply that only when life expectancy at birth reach about $\exp(4.0) = 55$ or higher education becomes profitable enough for mass education to make convergence to technologically leading nations viable. This is, however, a side track in the current paper where it only need be noted that this specification was clearly statistically superior to alternatives like including linear and quadratic trends. However, the precision of the age coefficients become compromised because of the high correlation to life expectancy, although the basic hump shape pattern in the point estimates is preserved. This indicates that age structure per se actually picks up much of this variation in any case. Moreover, a forecast model would have to predict the switch point for all countries, conceivably also reversion back to the previous regime if war or epidemics lower life expectancy and income.

It remains to analyze the risk of spurious regression if we are regressing variables integrated of order one on each other. Theoretically the left hand variable may be unlimited and contain a unit root (since GDP/capita roughly could be modelled as a geometric Brownian motion with drift). Thanks to

Figure 6: Scatterplot of the logarithm of GDP per capita versus the logarithm of life expectancy of birth. Points connected by lines for each country. On the horizontal scale 3.5 corresponds to 33 years of life and 4.25 to 70 years of life expected at birth.

the log specification this also hold for the age shares (all negative values are theoretically possible) and life expectancy (although limited from above and in practice also from below since life expectancies below 1 will mean almost immediate extinction). Anyway, we are regressing rather smooth trended variables on other low-volatility trended variables making inference uncertain even if we would not have any exact random walks. In column 2 we therefore report the differenced ECM counterpart to column 1 with the lagged residuals from the long-run regression measuring the deviation from the long-run relation. According to Zovitt the t-statistic for the residual coefficient should have standard critical values and we would conclude that the hypothesis that the demographic variables and GDP per capita are cointegrated cannot be rejected. However, since this is a panel with country-specific effects we have a bias in the estimated coefficient making that conclusion questionable. We therefore tested the stationarity of the residuals finding no evidence of non-stationary behavior. Since we below also test the out-of-sample properties of forecasts we did not at this stage find it necessary to pursue the issue further with panel cointegration tests. It should be noted that the differenced coefficients does not leave much of an impact for age structure in the short-run dynamics except for the long-run error correction mechanism. That is hardly surprising since those slow-moving variables should not determine much of the business cycle behavior.

This cannot be taken to imply that we actually have estimated a long-run cointegrating vector reflecting the true relation between age structure and economic development. As will be adamantly clear below the relationship certainly shifts over time and stage of development and our general argument does not allow the conclusion that we actually observe any kind of steady state process but only a limited stretch of a secular transition path. But it should assure us that the statistical problem of spurious regression with random coefficients, diverging t-ratios and so on, does not seem any serious problem. If it were, long-run forecasts would behave erratically, which we show below is not the case.

Beyond the collinearity issue in using life expectancy to catch the long-term trend it could also introduce some simultaneity, since it is a viable hypothesis that contemporary correlation between life expectancy and GDP level is due to life expectancy being increased by higher standards of living and health care rather than life expectancy causing higher GDP. Although there is evidence that decreasing mortality (better health) in the working population increases GDP, life expectancy at birth is also highly influenced by GDP levels, as is survival at old age. However, we were not able to find valid instruments using lags of available variables, so further investigation of this issue must either use a systems approach or find a set of valid instruments.

This section should also be compared to a balanced sample excluding countries with wars and other anomalies.

4.2 Out-of-sample tests

In Table 3 are the cross-section means of forecast errors from a pure age model with life expectancy as a linear independent variable using data up to 1990. This is compared to a demographic model with lags and two other comparisons; one pure lag model and the naive model naive model with forecast $\hat{y}_{i,t+k} = y_{it}$ at k steps ahead outperforms the other models in terms of accuracy over the 8-year horizon reported here. However, while errors tend to increase rather fast both for the naive model and the two models including lags that is not the case for the pure demographic model. Moreover, the latter three models have a tendency towards a positive mean error $y_{it} - \hat{y}_{it}$ while the pure demographic model tends to have negative mean errors indicating a tendency to overshoot actual outcomes. Looking at different horizons of forecasting this general pattern appears consistent.

The demographic model including lags is much better at short horizons, but the errors are increasing fast and especially the mean error. According

Dep var	Fixed country effects	
log(GDP/cap)		Difference
Constant		
$\log a_{0-14}$	**-0.052**	-0.102
	(2.08)	(0.40)
$\log a_{15-29}$	0.003	0.218
	(0.16)	(1.19)
$\log a_{30-49}$	0.008	-0.023
	(0.41)	(0.13)
$\log a_{50-64}$	0.011	-0.028
	(0.76)	(0.24)
$\log a_{65+}$	-0.007	0.116
	(0.61)	(1.39)
loge0 high income	**0.004**	0.429
	(4.34)	(1.30)
y_{t-1}	**1.058**	**0.935**
	(70.9)	(9.99)
y_{t-2}	**-0.118**	**-0.103**
	(5.51)	(6.14)
y_{t-3}	0.017	0.024
	(1.15)	(1.61)
Long-run residual		**-0.870**
		(9.16)
$\bar{R}^2$	0.997	0.0821
F time effects	**4.39**	**3.87**
F country eff	–	–
Ljung-Box $Q\,(6)$	7.82	8.46

Table 2: Absolute t-values in parentheses. Bold face indicates that the estimates are significant at the 5 percent level. Newey-West robust errors allowing for three lags are reported.

Forecast	Mean error				Mean Abs Error				RMSE			
1990	Demograph		Comparison		Demograph		Comparison		Demograph		Comparison	
of year	Pure	Lags	3 lags	Naive	Pure	Lags	3 lags	Naive	Pure	Lags	3 lags	Naive
1991	-0.10	0.01	0.02	0.01*	0.25	0.04*	0.04*	0.04*	0.31	0.06	0.06	0.05*
1992	-0.11	0.02	0.03	0.01*	0.26	0.07	0.08	0.06*	0.33	0.09	0.10	0.08*
1993	-0.12	0.03	0.04	0.01*	0.28	0.10	0.11	0.09*	0.37	0.13	0.14	0.12*
1994	-0.13	0.03	0.05	0.01*	0.30	0.14	0.15	0.12*	0.39	0.18	0.19	0.16*
1995	-0.11	0.06	0.08	0.04*	0.31	0.16	0.18	0.14*	0.39	0.21	0.22	0.18*
1996	-0.11	0.09	0.11	0.06*	0.32	0.18	0.21	0.16*	0.40	0.24	0.26	0.21*
1997	-0.09*	0.11	0.15	0.09*	0.32	0.21	0.24	0.18*	0.41	0.27	0.29	0.23*
1998	-0.09*	0.13	0.17	0.10	0.31	0.21	0.24	0.17*	0.39	0.26	0.30	0.23*

Table 3: Comparison of forecast errors between a naive model, the best pure lag model with three lags, a pure demographic model and a demographic model including three lags of the dependent variable. Asterisks mark the smallest errors at each horizon.

to a RMSE criterion, which penalizes large errors we would definitely prefer the naive model although the demographic model with three lags is only marginally worse. for short term forecasting. The same evaluation holds for the mean absolute error but for the mean error the pure demographic model becomes competitive at the longest horizons.

One important reason for the naive model to do so well is that it starts at the current level and then only accumulates the deviations from this level for each country as the horizon widens. The other models are fixed effects models where the deviation between the model and the actual values in the last year in sample may be quite substantial giving rise to substantial errors at short horizons. The lagged models condition on recent values and thus have less of this problem, which is well known in the forecasting literature (and often corrected by an add-in factor that adjusts the level of the forecast to current values). The price for this is, however, that the truly independent variables get less influence on the forecasts with rather quickly rising errors.

Recursive estimation of the demographic model shows that coefficients are quite unstable up to the beginning of the 1990s. This may be an indication of a problem that is inherent in estimations where the variables are changing only slowly. Since most of the variance is at low frequencies we need long periods of data in order to get reliable estimates. In a sense we do not have enough degrees of freedom at those frequencies. That is why we have chosen to report out-of-sample only from 1991, in spite of the general pattern being rather similar also for longer horizons.

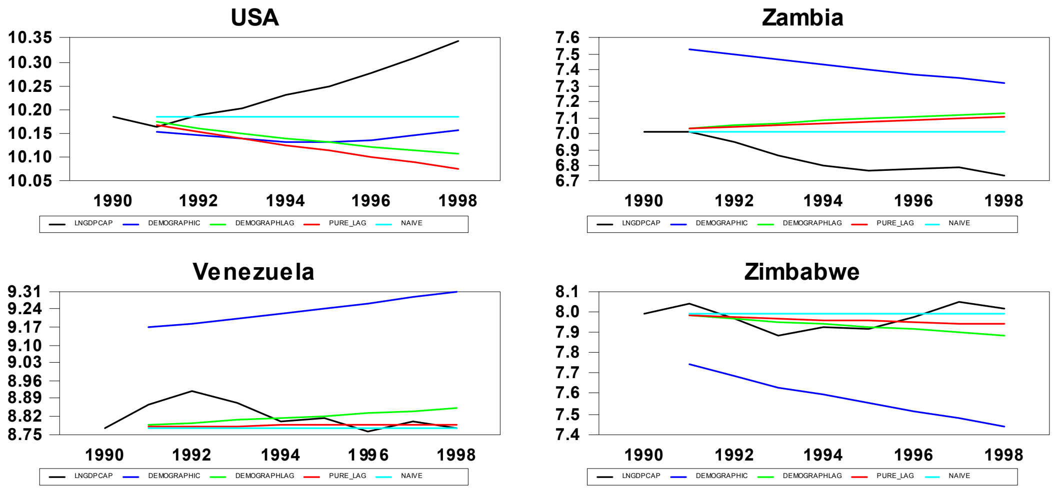

Figure 7: Four examples of forecast models compared to the actual outcome.

It remains to evaluate the importance of outliers for these results. A comparison with excluded countries can yield further information.

At this point the forecasting performance of the demographic model seems rather disappointing although the relative stability of the pure demographic model does hint at the possibility that better results can be achieved by taking account of heterogeneity. Figure 7 shows the respective forecasts for four countries and compares to actual outcomes. While the naive model is best in all cases we do get a more nuanced picture of why that is so. The lagged models behave rather similar and essentially forecasts the previous trend movement this leads completely wrong in the case USA and Zambia where the previous trends happen to be reversed in the 1990s but works pretty well in Venezuela and Zimbabwe where the economies stagnated during the period. However, in the case of USA the demographic model predicts a trend break although later and weaker than actually was the case. In Zambia the downward trend would actually have been captured by the demographic model if we had inserted an add-in factor to adjust the level. In Venezuela and Zimbabwe the demographic model would have done much better with an add-in factor but still worse than the other models. Clearly by using an add-in factor (i.e. adjusting the country-specific level at which the forecasts start) the performance of the demographic model would be greatly improved. More importantly the demographic model at least has some ability to forecast trend breaks which the other models seems to lack.

In the next subsection we investigate whether long-term predictions seem sensible or not.

4.3 Forecasts using a non-lag model

The figures below report the regression and forecasts up to 2050 of the model where η_i is a country-specific intercept

$$y_{it} = \sum_{k=0-14}^{65+} \beta_k a_{kit} + \gamma \log e0 + \eta_i + \varepsilon_{it}$$

This is the same model for which out-of-sample errors are reported in Table 3 above but using data up to 1996 to generate the forecast. Performing the same experiment with the lagged model the forecasts tends to yield similar but damped results once the influence of initial conditions in the lag model wears off. As the out-of-sample errors above indicates the demographic model without lags tends to overestimate while the lag model tends to underestimate. In the very long run up to 2050 the lag model therefore implies considerably lower growth rates. The mean growth rate in the pure age model is forecasted to be 2.6 percent with a minimum of -7.6 and a maximum of 15.9. The mean growth rate of the lag model on the other hand is only 0.9 percent with a minimum of -27 percent and a maximum of 42 percent, thus also a much larger spread.

The figures have been adjusted to the same scale and contain the whole time series for each country both outcomes and predictions from the model. Although the logarithmic scale makes it less obvious we can still discern how most developed countries show a tendency to flatten out in the forecast period as they reach their advanced aging phase. Developing countries mostly take off and start to close in on the developed countries. Some Sub-Saharan countries tend to lag behind, however, but there is still improvement rather than the negative trend many of them have actually experienced.

However, the demographic model may have a tendency for overestimation as was demonstrated above. It bears emphasizing that there is no tendency for the model to explode or behaving erratically even at very long horizons. In the next section we study to what extent the conclusions so far are modified by taking account of heterogeneity due to different stages of the demographic transition.

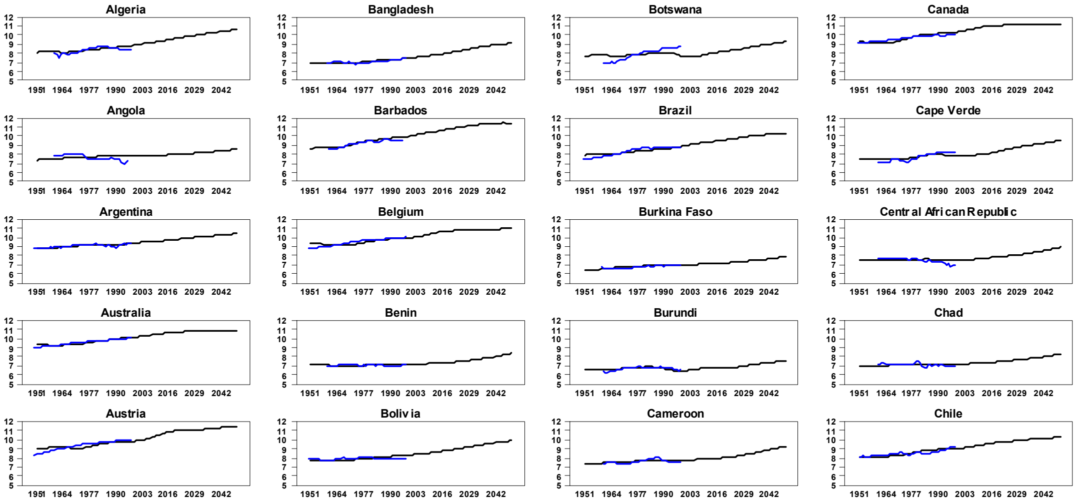

Figure 8:

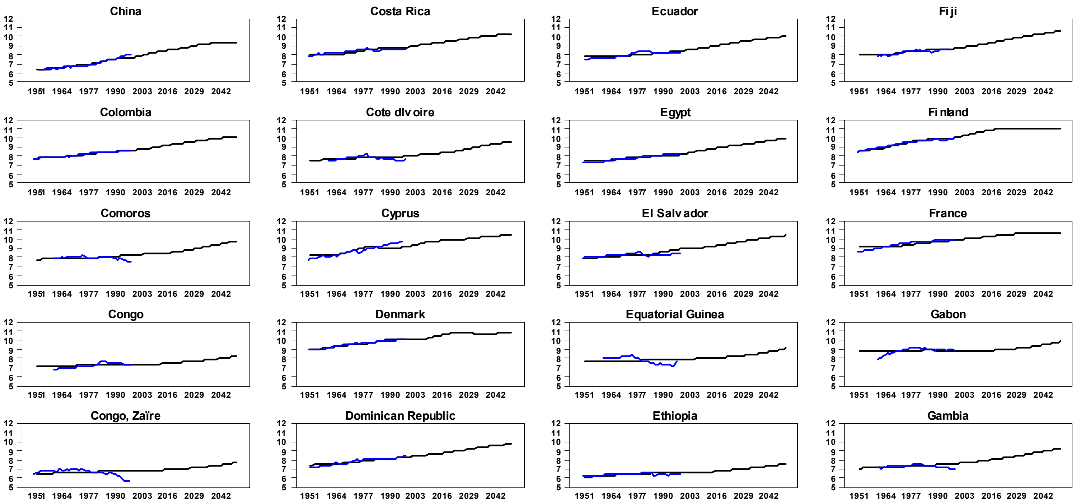

Figure 9:

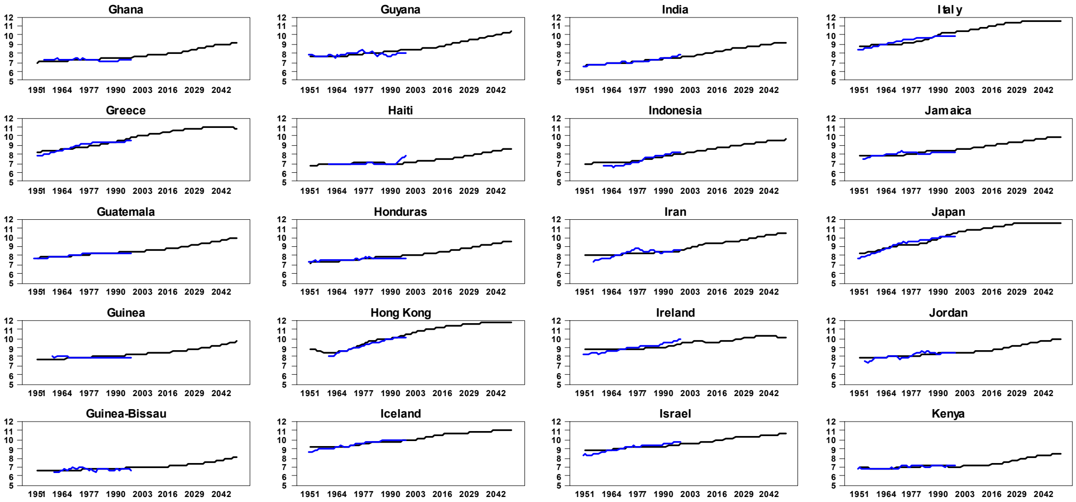

Figure 10:

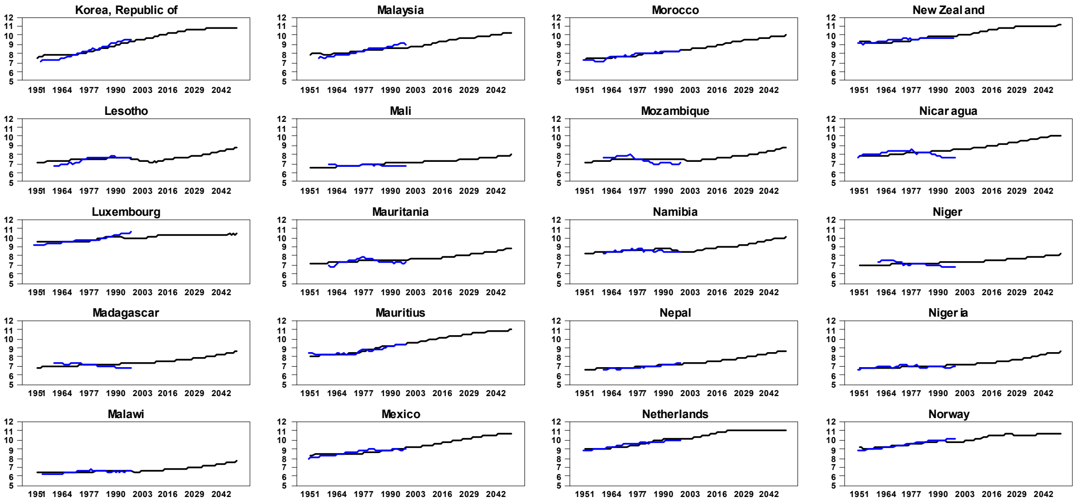

Figure 11:

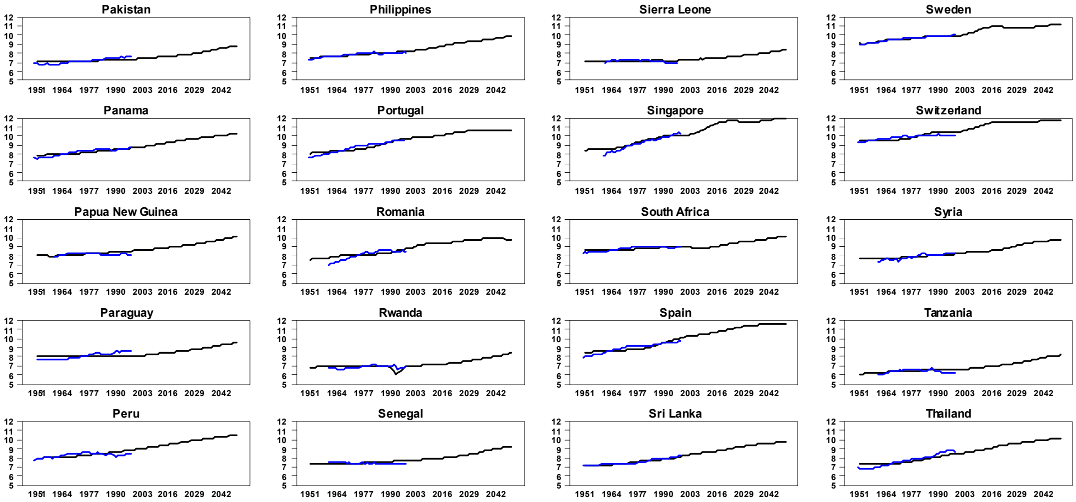

Figure 12:

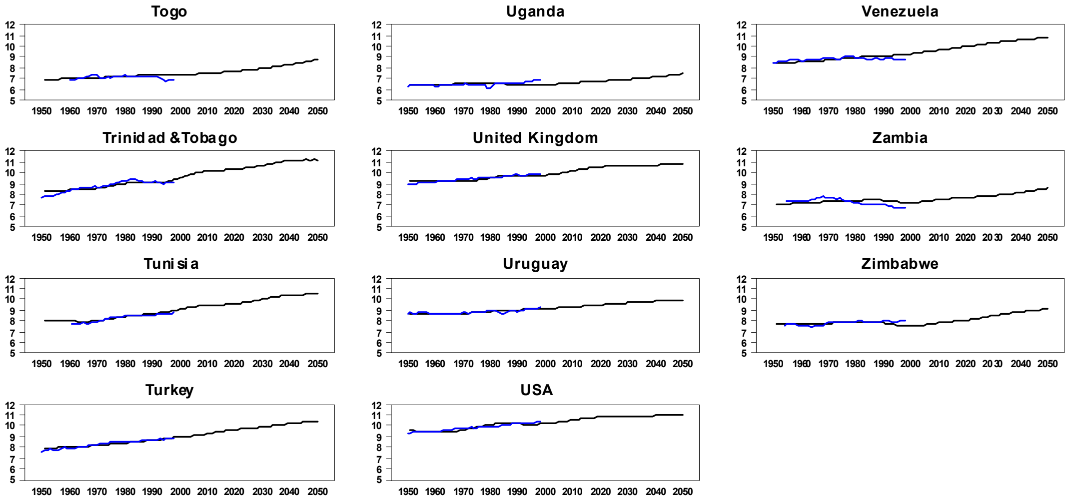

Figure 13:

5 Time-varying age share parameters: Using demographic transition theory to forecast global income growth

The previous section indicates that the homogenous demographic model is rather inexact as a forecasting tool but seems to have an unusually stable error distribution as the horizon widens. In this section we extend the model to allow for more heterogeneity across countries and time. There are, of course, numerous ways to do that, but we take our clues from demographic transition theory as we modify the model in order to improve on its forecasting properties.

5.1 Extended model

The extended model incorporates three demographic factors that influence economic development: Mortality, age structure and urban structure. Each of these factors are closely connected to the demographic transition. Declining mortality is what triggers the demographic transition. Age structure change is, as argued above, a process that necessarily accompanies the transition. Urbanization, finally is a spatial redistribution of the population that also has characterized all countries undergoing demographic transition. All these three factors have also been focused in the discussion of the last decade concerning the determinants of national economic growth.

Our aim is to analyze if we in historic data can find such stable correlations between economic development, mortality, age structure and urban structure that allow us to forecast future trends in income growth.

The model we have chosen to estimate deviates from the basic model in equation 2 in the following ways:

1. *Log per capita income is assumed to be a quadratic function of life expectancy at birth.* As was established above the relation between development and trend as well as life expectancy is non-linear. Although the best fit seems to be a kinked linear relation to life expectancy this is hard to use for forecasting but a quadratic function should capture at least some of the same features.

2. *Log per capita income is assumed to be positively influenced by a well-developed urban system as measured by the population size of the largest city (in logs).* However, urbanization effects are allowed to differ between countries at different levels of mortality. The positive effect of

urban size on economic development has long been a staple of regional development theory. During the 1990s it has also received new interest from leading economists. To measure urbanization by the size of the largest city is, however, to our knowledge a new approach, as well as the idea that city system effects can differ according to the level of mortality.

3. *Log per capita income is influenced by shifts in the age structure of the population.* The relations between age structure and per capita income are allowed to change at different levels of mortality. Depending on whether developed or developing countries have been focused there has been some differences in what age groups that seem to be the most important for influencing economic development. Moreover, theoretical research on the effect of life expectancy on the optimal length of education, fertility, saving and retirement decision indicate that the economic life cycle should change in response to changing mortality. This implies that it can be misleading to assume that the profile of age structure effects on per capita income will be the same in countries at different mortality levels. Rather we would expect the pattern to shift with life expectancy. To capture this our specification allows for an interaction between age effects and a quadratic function of life expectancy.

4. *Recent observations have been given a larger weight than early observations.* Through the use of this weighting procedure the estimate of the country specific intercept will be more strongly influenced by conditions in the later part of the estimation period. Given that we have structural change in the relations this is important when the purpose is forecasting.

Thus we arrive at a model which can be written

$$w_t y_{it} = \sum_{k=0-14}^{65+} (\beta_k + \gamma_k e0_{it} + \delta_k e0_{it}^2) w_t a_{kit} + \beta_u w_t \pi_{it} + \gamma_u w_t \pi_{it} e0_{it} + \eta_i + \varepsilon_{it} \quad (3)$$

where π is the population in the largest city and $w_t = \log(t - 1959)$.

5.2 Results

Table 4 presents the results obtained by estimating the model presented above. All the parameters refer to the same model. Column 1 gives the results for the direct effects. Column 2 gives the parameter estimates obtained for interactions with life expectancy at birth. Column 3 gives the results from interactions with life expectancy squared. Of the 20 parameter estimates presented in the table all except 5 are significantly different from zero at the 5 %-level.

The estimates show that life expectancy, as expected, is directly correlated with per capita income. However, the estimates also confirm the hypothesis that life expectancy modifies the effect of demographic structure. Thus, a well developed urban system has a positive effect on per capita income, but only when life expectancy rises above 62 years. Life expectancy also modifies the effects of age structure.

The estimates show that, as expected, there is a strong negative effect on per capita income from high levels of child dependency. The strongest negative effect of child dependency is found when life expectancy is around 60 years. With further increases in life expectancy the effect becomes somewhat weaker but it is still negative when life expectancy has reached 75 years.

The effect of young adults is much less negative. This implies that per capita income will improve when child dependency declines and the share of young adults increases. However, as life expectancy increases above 70 years the effect of the young adult share becomes negative. This might indicate that increasing length of education in low mortality populations reduces the income boosting effect of young adults.

A surprising result is that in high and medium mortality populations (life expectancy < 65 years) there is an effect on per capita income of young middle-aged adults (30-49 years) almost as negative as the effect of children. When life-expectancy increases above 70 years, however, a high share of people aged 30-49 years becomes a strong booster of per capita income. A negative effect of prime aged adults seems counter-intuitive but there are, at least, two possible substantive explanations. One is that adults aged 30-49 are the primary care takers of dependent children. Since most high and medium mortality populations still have high fertility rates this age group, therefore, may face severe time restrictions with respect to their participation in market activities. A second possibility is given by the fact that a high share of young middle aged adults in populations with moderate and low life expectancy is associated with increased risk for severe disturbances to the social order like revolts, riots and civil war. It remains far from clear what lies behind this association but the existence of such a pattern would explain why a high share of people aged 30-49 does not give rise to the expected positive effects on per capita income. However, we cannot at this point ignore the possibility that it is simply a statistical artefact due to the high negative correlation between this group and the age share of children.

The 50-64 group follows a similar but less pronounced pattern. A weak negative effect when E0 is around 60 turns positive when life expectancy increases above 75 years.

For the old age share the effect is negative but the estimated parameters are small. To some extent this reflects the use of log shares. In populations with few old age people small percentage point increases in the old age share translates into large increases in the log share.

One way to summarize these findings is to note that at low and medium levels of life expectancy the age effects on per capita income are dominated by the balance between children and young adults. Child-rich populations tend to be poor whereas countries with declining child dependency and an expanding young adult population enjoy rising per capita income. At higher levels of life expectancy it is instead a high share of middle age adults (30-49 and 50-64 years old) that ensure good economic prospects. The hypothesis that changing life expectancy can alter the profile of age structure effects has, thus, been corroborated.

Table 4 does not give the estimated country effects. These are displayed in Figure 14 which shows that a substantial part of the variation in per capita income is captured by the country effects. These strong country effects underscore that demography is not the only factor that determines per capita income. However, the success of demographic transition theory gives us strong reasons to believe that population forecast based on the transition hypothesis are relatively reliable. Therefore, a forecast of future changes in per capita income based on the model presented here can provide a picture of possible global trends that are consistent with historical correlations between demographic and economic change.Performing the same out-of-sample experiment as above, i.e. estimating the model on data up to 1990 and then comparing forecasts to actual outcomes we get the results in Table 5. Comparing to Table 3 there is a substantial improvement on the demographic model but we retain the feature of a very slow increase of errors with the horizon. However, the naive model is still better.

Judging from the mean errors we have substantially less of a problem with systematic bias in the forecasts when we take better account of heterogeneity. In the next section we present the results.

5.3 Forecasts

Below the 2000-2050 forecasts per capita income in different countries are presented. These forecasts are all based on the medium variant of World Population Prospects: The 2000 Revision. The fertility assumptions applied by the Population Division in the medium variant are reproduced in appen-

	Direct effect	Interaction effects $(e0 - 61.8)$	$(e0 - 61.8)^2$
Intercept	-0.3580 (-0.31)		
$e0$	**0.0276** (5.24)		
$e0^2$	**-0.0009** (-3.75)		
ln maxpop	-0.0004 (-0.02)	**0.0095** (10.62)	.
lna014	**-2.5880** (-6.75)	0.0392 (0.88)	**0.0058** (3.45)
lna1529	0.2121 (0.83)	0.0079 (0.29)	**-0.0025** (-2.34)
lna3049	**-1.4792** (-7.25)	**0.1012** (4.78)	**0.0041** (4.39)
lna5064	**-0.6598** (-5.03)	**0.0211** (2.28)	**0.0025** (4.54)
lna65	**-0.1782** (-2.62)	**0.0194** (3.00)	**-0.0008** (-2.14)
Country effects	F-ratio	204.31	
RMSE	1980-1998	0.146	

Table 4: Effects of demographic structure and life expectancy on log per capita income. Interaction model, (t-values in parentheses). Weighted with log (year-1959).

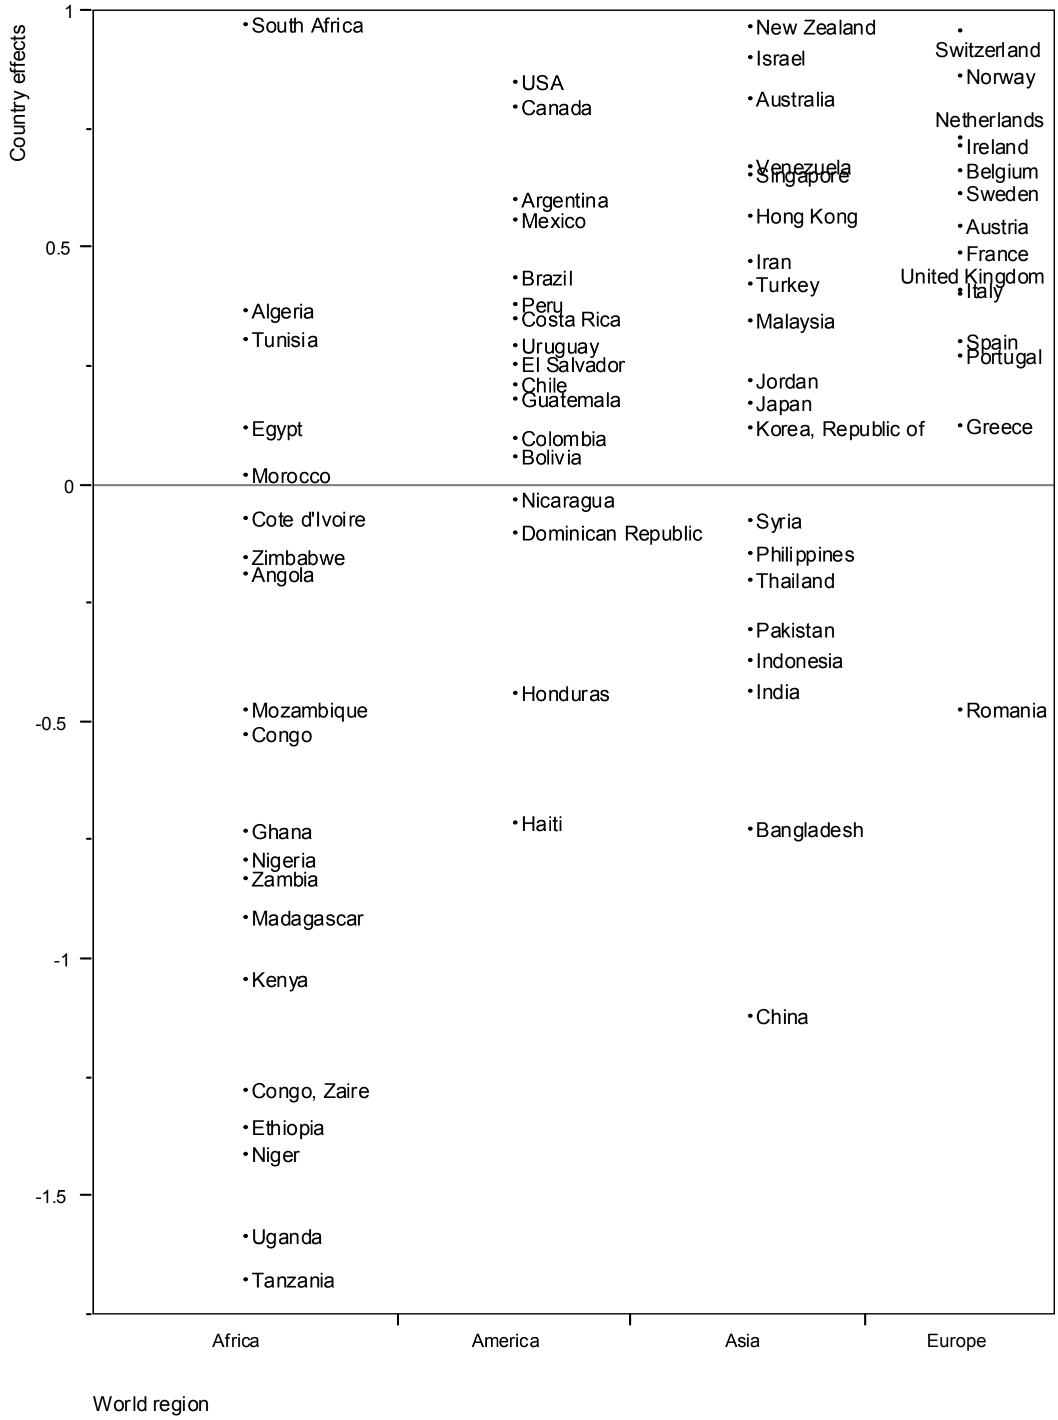

Country effects
1
0.5
0
-0.5
-1
-1.5
South Africa
Algeria
Tunisia
Egypt
Morocco
Cote d'Ivoire
Zimbabwe
Angola
Mozambique
Congo
Ghana
Nigeria
Zambia
Madagascar
Kenya
Congo, Zaire
Ethiopia
Niger
Uganda
Tanzania
USA
Canada
Argentina
Mexico
Brazil
Peru
Costa Rica
Uruguay
El Salvador
Chile
Guatemala
Colombia
Bolivia
Nicaragua
Dominican Republic
Honduras
Haiti
New Zealand
Israel
Australia
Venezuela
Singapore
Hong Kong
Iran
Turkey
Malaysia
Jordan
Japan
Korea, Republic of
Syria
Philippines
Thailand
Pakistan
Indonesia
India
Bangladesh
China
Switzerland
Norway
Netherlands
Ireland
Belgium
Sweden
Austria
France
United Kingdom
Italy
Spain
Portugal
Greece
Romania
Africa
America
Asia
Europe
World region

Figure 14:

	Mean error	Mean absolute error	Root mean square error
1991	-0.033	0.148	0.194
1992	-0.041	0.162	0.208
1993	-0.050	0.180	0.238
1994	-0.050	0.195	0.265
1995	-0.033	0.214	0.282
1996	-0.022	0.226	0.302
1997	-0.005	0.231	0.314
1998	0.009	0.211	0.272

Table 5: Forecast errors for the model estimated at the sample 1950-1990. Averages over countries.

dix A. According to the Population Division mortality in the medium version is projected on the basis of the models of change of life expectancy produced by the United Nations. In countries highly affected by the HIV/AIDS epidemic, estimates of the impact of the disease are made explicitly through assumptions about the future course of the epidemic, that is, by projecting the yearly incidence of HIV infection.

The Population Division provides predicted values for the variable "population in the largest city" only up to 2015. Because of this the forecast we present below is based on a separate projection. This projection is based on an interpolation of observed country-specific growth rates for the population in the largest city. This observed growth rate has been adjusted for effects on the growth rate from shifts in the growth rate of the 0-14 age group, the 50-64 age group and the share of the 0-14 age group as well as allowing for a time-trend in the growth rates for the population of the largest city.

These forecasts have been produced for 82 countries. Countries for which no forecast have been made are either not included in our original sample or do not have values available for the size of their largest city in UN's Urbanization prospects.

The parameter values used in the forecast are those that are presented in Table 4 and Figure 14. One important modification, however, is that in the forecast the maximum value of life expectancy has been fixed to 76 years. The motivation for this is that in the sample used for estimating the model only a few countries have experienced life expectancies above 76 years for an extended period. As a consequence, there is no information about how increased life expectancies beyond 76 years will affect economic growth. In the forecast we present, this lack of knowledge is represented by an assumption that countries with a life expectancy above 76 years will

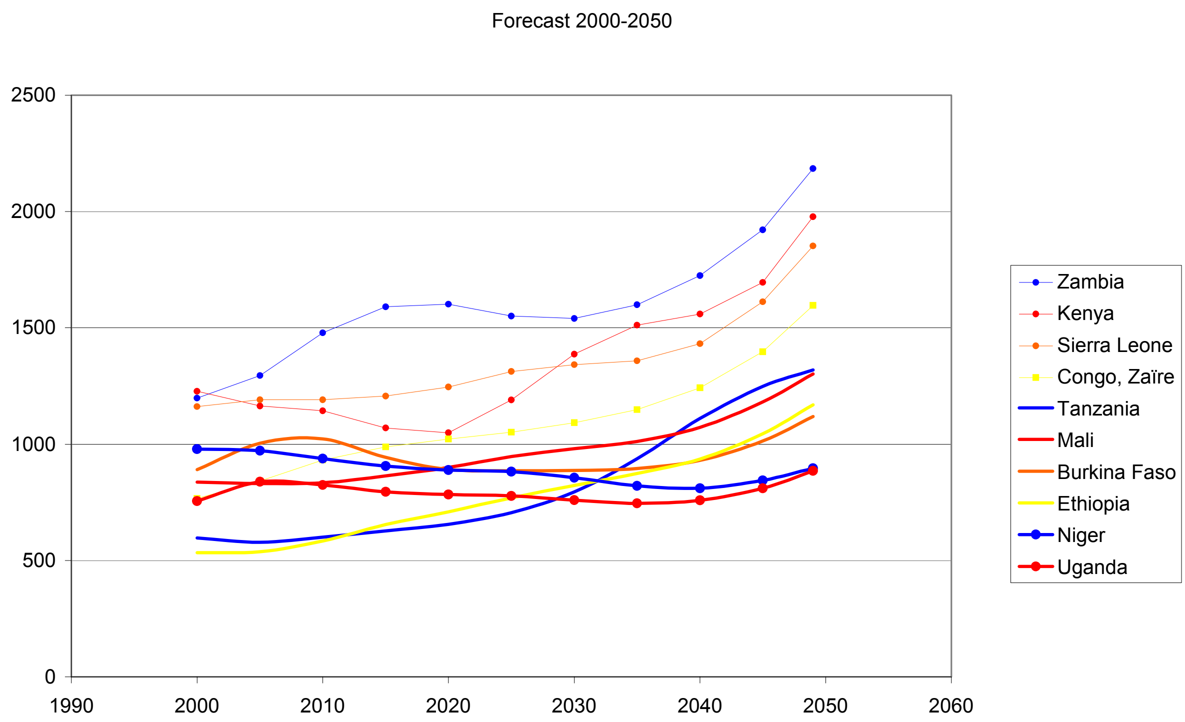

Figure 15:

behave in a way similar to countries that have only a life expectancy of 76 years.

The maximum income level in these forecasts are closer to those of the homogeneous lag model than the pure demographic model, but the general pattern that developed countries will tend to stagnate while developing countries will tend to take off on a growth path is confirmed. The exceptions are mainly in Sub-Saharan Africa and due to the AIDS epidemic, but even there we predict a take-off eventually.

6 Conclusion: A better world

Using demographic projections we evaluate forecasting models for GDP per capita based on demographic information. We pursue two different strategies. The first approach attempts to estimate a globally homogeneous model. The most important conclusion from this approach is that a forecast model that only uses age distribution information seems to have a rather narrow forecast

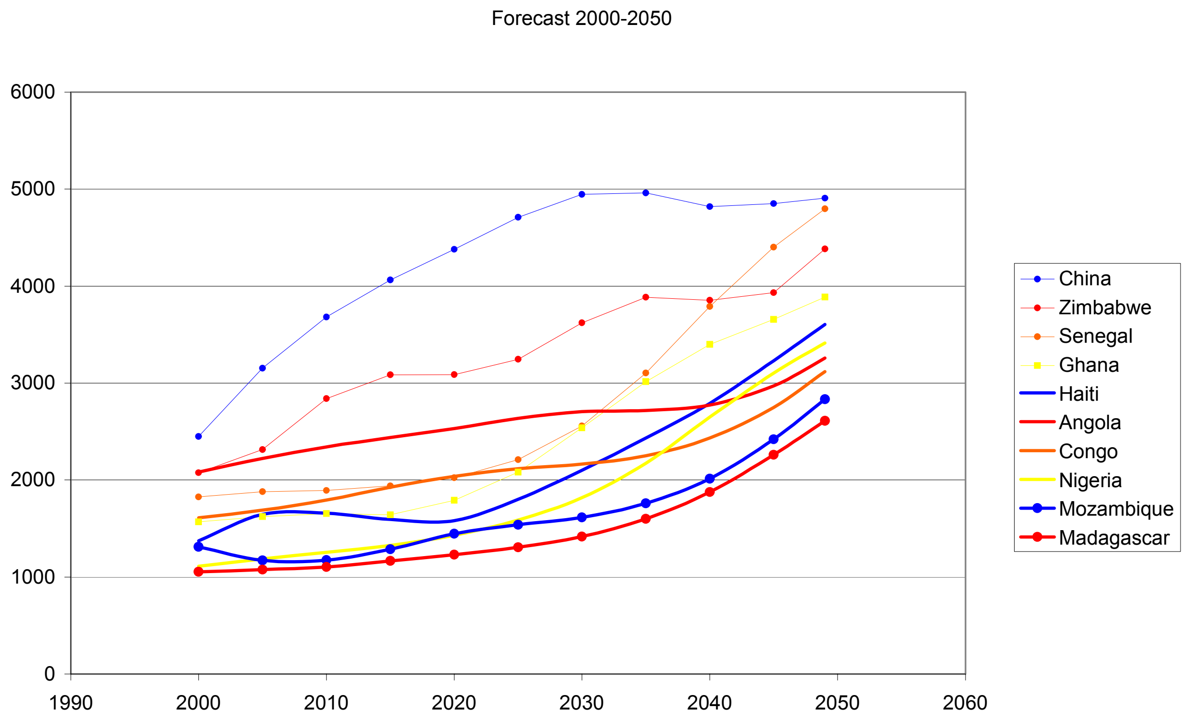

Forecast 2000-2050
6000
5000
4000
3000
2000
1000
0
1990
2000
2010
2020
2030
2040
2050
2060
China
Zimbabwe
Senegal
Ghana
Haiti
Angola
Congo
Nigeria
Mozambique
Madagascar

Figure 16:

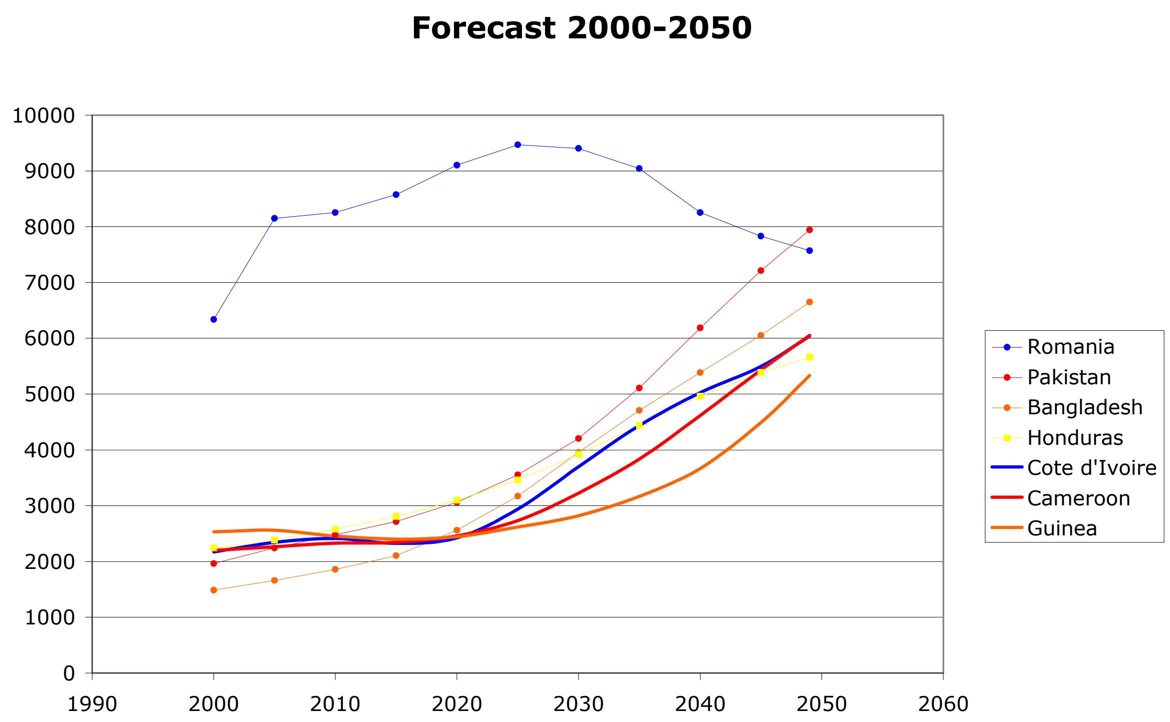
Forecast 2000-2050
10000
9000
8000
7000
6000
5000
4000
3000
2000
1000
0
1990
2000
2010
2020
2030
2040
2050
2060
Romania
Pakistan
Bangladesh
Honduras
Cote d'Ivoire
Cameroon
Guinea

Figure 17:

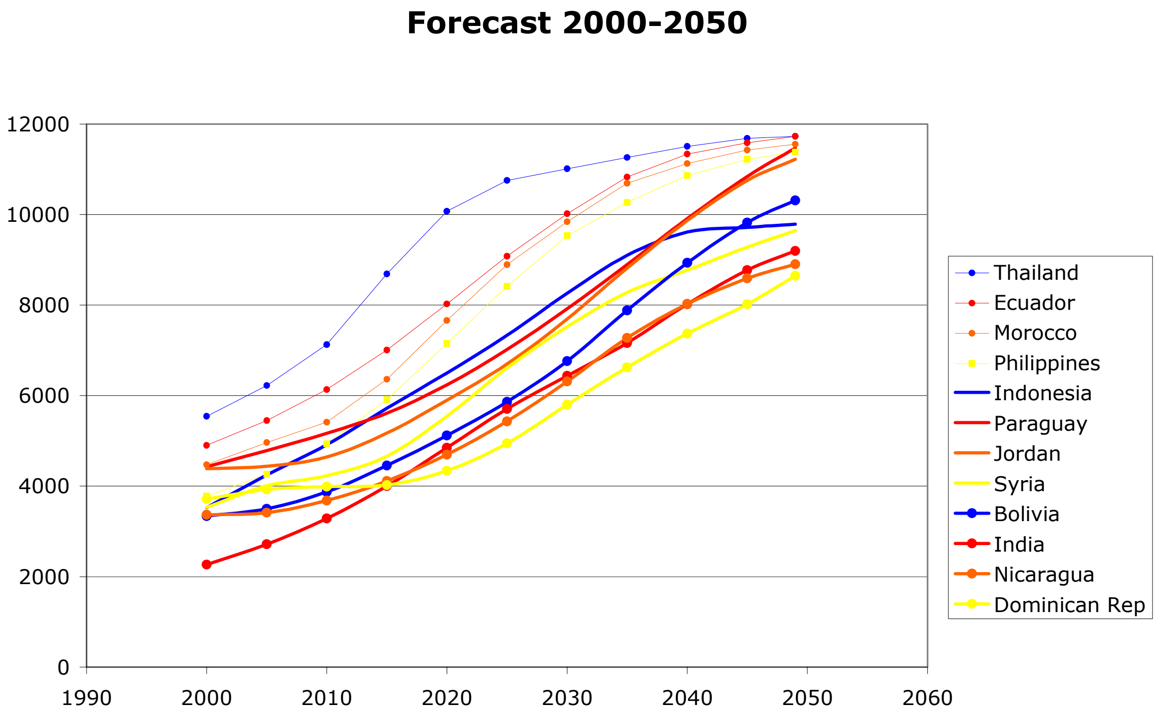
Forecast 2000-2050
12000
10000
8000
6000
4000
2000
0
1990
2000
2010
2020
2030
2040
2050
2060
Thailand
Ecuador
Morocco
Philippines
Indonesia
Paraguay
Jordan
Syria
Bolivia
India
Nicaragua
Dominican Rep

Figure 18:

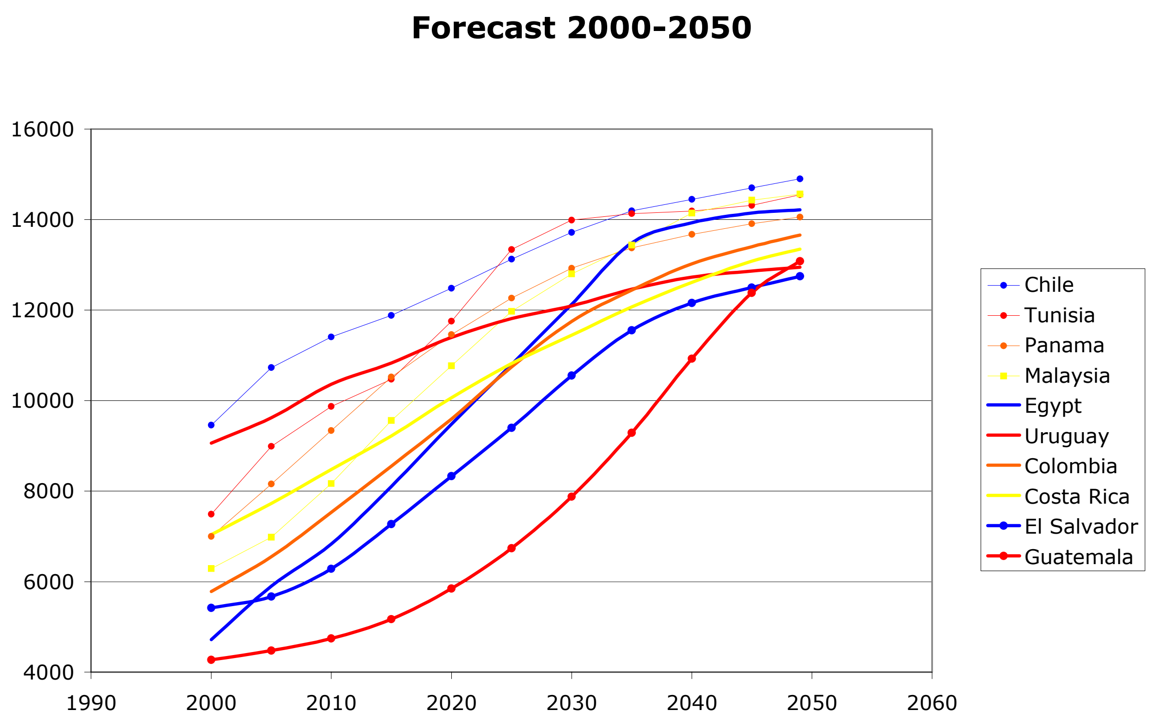
Forecast 2000-2050
16000
14000
12000
10000
8000
6000
4000
1990
2000
2010
2020
2030
2040
2050
2060
Chile
Tunisia
Panama
Malaysia
Egypt
Uruguay
Colombia
Costa Rica
El Salvador
Guatemala

Figure 19:

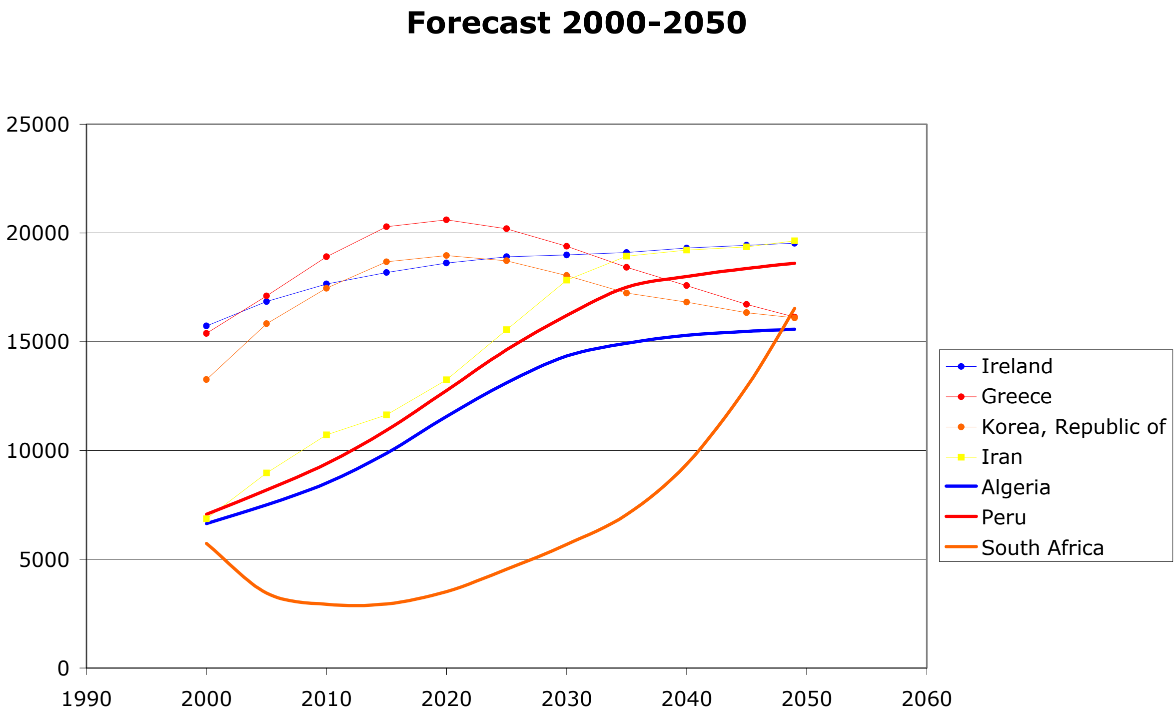
Forecast 2000-2050
25000
20000
15000
10000
5000
0
1990
2000
2010
2020
2030
2040
2050
2060
Ireland
Greece
Korea, Republic of
Iran
Algeria
Peru
South Africa

Figure 20:

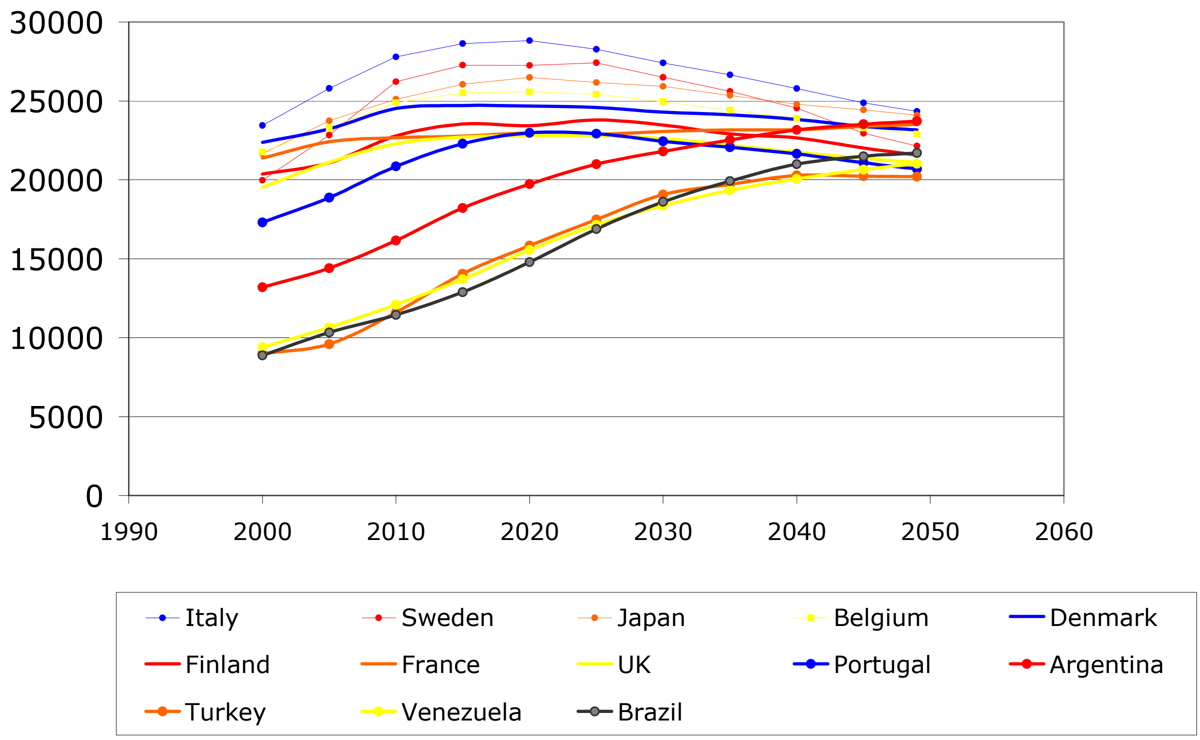
30000
25000
20000
15000
10000
5000
0
1990
2000
2010
2020
2030
2040
2050
2060
Italy
Sweden
Japan
Belgium
Denmark
Finland
France
UK
Portugal
Argentina
Turkey
Venezuela
Brazil

Figure 21:

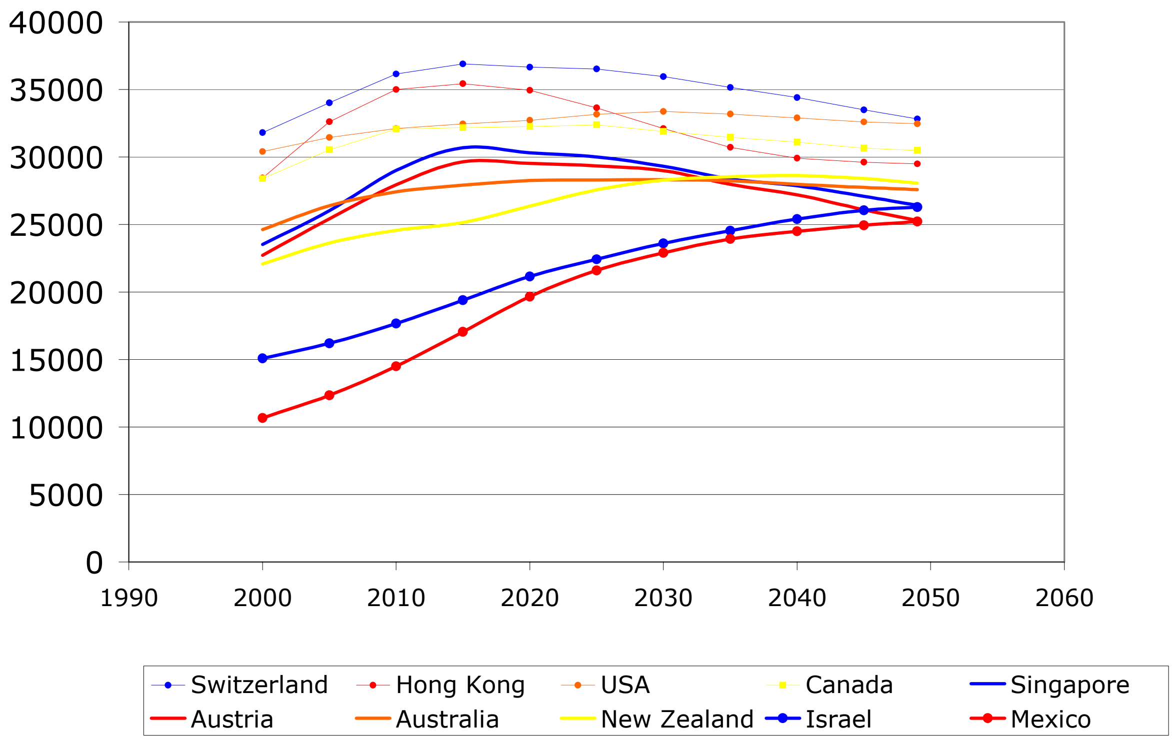

40000
35000
30000
25000
20000
15000
10000
5000
0
1990
2000
2010
2020
2030
2040
2050
2060
Switzerland
Hong Kong
USA
Canada
Singapore
Austria
Australia
New Zealand
Israel
Mexico

Figure 22:

error distribution that increases only at a very slow rate. The precision of forecast for individual countries is, however, not very impressive.

The other approach allows for heterogeneity by interaction with life expectancy variables and also includes urbanization trends in the model. This model narrows down the forecast errors. The picture emerging from both models is, however, similar in its general trends. The currently aging developed countries will experience a stagnating or even negative growth trend in income. Most developing countries will, however, experience accelerating growth and converge towards the income levels of the developed world. The main exceptions to this are to be found in sub-Saharan Africa where the impact of AIDS on the age distribution is expected to postpone any growth take-off. However, even in these countries the UN assumptions that the AIDS epidemic will be brought to an end results in increasing growth rates towards the end of the period.

Thus we expect that demographic change and the demographic gift following from decreased fertility rates will substantially decrease the share of people living in extreme poverty in the world. The negative aging effects on developed countries do not imply any catastrophic decrease of living standards. The futures scenario that appears from this forecasting exercise is thus a rather bright one, where economic prosperity and equality increases and fewer people will live in extreme poverty.

The reliability of our forecasts, of course, must be subjected to further tests and studies. There are, however, at least three circumstances pointing in the direction that the method compares favorably to any alternatives. First, demographic structure is the most fundamental determinant of human resources that we can easily measure. Second, this structure can be reliably projected for a long period of time thanks to the inertia inherent in demographic momentum. Third, the slow increase of the forecast errors that we find in our out-of-sample tests makes it improbable that even large stochastic disturbances make long-term forecasts meningless, as is often the case with other income forecasts. Although our point forecasts for individual countries in 2050 may be highly uncertain, the statistical patterns generated seem reasonable enough.

In the absence of contrary evidence we therefore find that demographic projections indicate a better world within the next half century.

BIBTEX

References to be added.

Appendix A: Fertility assumptions in the medium variant of of World Population Prospects: The 2000 Revision

A. FERTILITY ASSUMPTIONS

Fertility assumptions are described in terms of the following groups of countries:

1. High-fertility countries: Countries that until 2000 have had no fertility reduction or only an incipient decline;

2. Medium-fertility countries: Countries where fertility has been declining but whose level is still above replacement level (2.1 children per woman);

3. Low-fertility countries: Those countries with fertility at or below replacement level (2.1 children per woman) plus a few with levels very close to replacement levels that are expected to fall below replacement level in the near future.

Medium-fertility assumptions:

1. Fertility in high-fertility countries is generally assumed to decline at an average pace of nearly one child per decade starting in 2005 or later. Consequently, some of these countries do not reach replacement level by 2050.

2. Fertility in medium-fertility countries is assumed to reach replacement level before 2050.

3. Fertility in low-fertility countries is generally assumed to remain below replacement level during most of the projection period, reaching by 2045-2050 the fertility of the cohort of women born in the early 1960s or, if that information is lacking, reaching 1.7 children per woman if current fertility is below 1.5 children per woman or 1.9 children per woman if current fertility is equal to or higher than 1.5 children per woman.