Working Paper No. 471

Population Forecasts, Fiscal Policy, and Risk

by

Shripad Tuljapurkar
Stanford University

August 2006

Final paper for the conference, "Government Spending on the Elderly" at The Levy Economics Institute of Bard College, April 28-29, 2006.

The Levy Economics Institute
P.O. Box 5000
Annandale-on-Hudson, NY 12504-5000
http://www.levy.org

ABSTRACT

This paper describes how stochastic population forecasts are used to inform and analyze policies related to government spending on the elderly, mainly in the context of the industrialized nations. The paper first presents methods for making probabilistic forecasts of demographic rates, mortality, fertility, and immigration, and shows how these are combined to make stochastic forecasts of population number and composition, using forecasts of the U.S. population by way of illustration. Next, the paper discusses how demographic models and economic models can be combined into an integrated projection model of transfer systems such as social security. Finally, the paper shows how these integrated models describe various dimensions of policy-relevant risk, and discusses the nature and implications of risk in evaluating policy alternatives.

JEL Classifications: J11, J18, H55, H60

Keywords: Demography, stochastic population forecast, social security, fiscal policy, probability of solvency

1. INTRODUCTION

This paper describes stochastic population forecasts as they relate to the development of policies related to government spending on the elderly, mainly in the context of the industrialized nations. I begin by discussing methods of using probabilistic forecasts for projecting demographic rates, mortality, fertility, and immigration. I show how these are combined to make stochastic forecasts of population number and composition, illustrating with forecasts of the United States population. Next, I discuss how demographic models and economic models can be combined into an integrated projection model of transfer systems such as Social Security. I show how these integrated models describe various dimensions of policy-relevant risk and discuss the nature and implications of risk in evaluating policy alternatives. Finally, I consider briefly the distinct issues in the projection of health care spending.

The background of this paper will be familiar to many readers. United States demography in the early 21st century will be shaped first by the aging of the baby boomers and then by the sustained increase in human life span that began over a century ago. These factors will result in an unprecedented increase in the proportion of people who are over 65 years old, starting in about 2010. Since existing public programs for pensions and health care are largely paid for by transfers from taxpayers (in the prime working ages of 20 to 65) to people over 65, the fiscal cost of these transfers will likely have to go up unless there are changes in the benefits paid by these public programs. Much the same phenomenon is taking place contemporaneously across the industrialized world and will also occur after some decades in most other countries of the world, most notably India and China. Research on this topic addresses the possible individual and policy responses to these forthcoming demographic changes. At the aggregate level, the central quantities of interest are the taxes collected to pay transfers and the size of the total transfer paid to the older population. These have two components, per-capita tax and transfer rates, and the numbers of taxpayers and recipients. If we focus on the rates, as Richard Disney said in his 1996 book, projected fiscal changes have nothing to do with demography. If we focus on the second, demography matters. Of course, in the real world both matter, as does the interaction between them. In this paper, my primary focus is on demography, but the integrated models discussed here provide a tool for the exploration of interactions, as I will illustrate.

The next section of this paper discusses the determinants of population change and the nature of stochastic forecasting models, then considers, in turn, mortality, fertility, and migration. Section 3 discusses stochastic population projections for the United States and some of their features. Section 4 considers integrated projection models that combine economic, policy, and demographic variables, and presents in outline such a model for the United States Social Security system. Results from this model are used to illustrate and discuss policy analysis in the context of the substantial uncertainty in our projections of the future. In particular, I discuss ways in which one constructs and analyzes measures of risk. In Section 5, I indicate how the integrated model of Section 4 can be adapted to address health care costs, and conclude with some reflections on the value of these methods.

Much of the work underlying this paper has been done in collaboration with Ronald Lee over the past 15 years. Significant contributions to parts of this effort have also been made by Mike Anderson, Carl Boe, and Ryan Edwards. My goal here is to present the current state of this work, so the citations are not exhaustive.

2. DEMOGRAPHIC FORECASTS: VITAL RATES

Demographers study and forecast the numbers and composition of populations. It is possible to subdivide populations quite finely, for instance by age, sex, ethnicity, education, marital status, and household status. But our ability to accurately model the dynamics involved decreases as disaggregation increases, so most forecasts only project population composition by age and sex. To do this, we require forecasts of age-specific and sex-specific vital rates, fertility, mortality, and immigration. Given a launch year, say 2005, and the population composition in that year, plus projections of all of rates for every year in the forecast horizon, the population forecast is simply a book-keeping exercise called the cohort-component method. The key to the stochastic forecasting approach is that it incorporates the temporal volatility of vital rates and projects these into the future to produce a range of forecasts with associated probabilities. The hard part of this is to project the rates.

Mortality

Mortality rates for large national populations have a characteristic age pattern for ages through 85 and the age pattern of mortality at ages between 85 and 110 is now also fairly well understood (Thatcher, Kannisto, and Vaupel 1998). To make projections, we want to know how the level of mortality falls at all ages, and reliable forecasts are made possible by Lee and Carter's (1992) discovery that mortality rates in the United States have fallen at roughly the same age-specific exponential rate over many decades. Lee and Carter's conclusion has been found to apply to virtually every industrial country, starting with the G7 countries (Tuljapurkar, Li, and Boe 2000), and represents the most striking discovery about mortality change in the industrial age. Letting M(x,t) be the central death rate for age x in year t, the Lee-Carter model can be written in the form,

$$\log M(x,t) = \log M(x,t-1) - z\, b(x) + E(t). \qquad \text{(Equation 1)}$$

Here, z is a rate of decline common to every age and b(x) is an age-specific response factor that tells us the relative rate of decline specific to age x, and E(t) is a residual that is typically small in magnitude. The constancy of the rate of decline in z means that mortality falls at a roughly constant exponential rate over time.

For a forecaster, the Lee-Carter model has the equally important feature that deviations around the long-run trend can be effectively modeled using stochastic processes. Fitting the model (1) yields residuals E(t) that have no trend but can be described by a relatively simple stochastic process. Lee and Carter (1992), and many subsequent studies, find that E(t) is well-described by a random walk; some variants of this model use low-order autoregressive processes (UK Government Actuary 2001). The model in Equation 1 and a stochastic model for E(t) are fitted to some span of historical data. Starting with known mortality rates in a launch year, one can then generate an infinity of forecasts for subsequent years: first project a sample path w of the residual process, say E(t,w) for each forecast year t, and then use (1) to compute a corresponding forecast M(x,t,w). The probability distribution of the sample paths E(t,w) can be derived from the fitted model (analytically or by simulation) and this yields the probability distribution of M(x,t,w). Finally, these forecasts of central death rates are converted to age-specific probabilities of death using standard demographic methods.

Several recent studies have tested and extended the Lee-Carter method. The original method has been simplified slightly (Tuljapurkar, Li, and Boe 2000), tested using historical data and found accurate (Lee and Miller 2000), modified slightly for particular countries (Booth, Maindonald, and Smith 2002), better algorithms have been developed to guarantee smoothness and to generate optimal forecast models for the E(t) process (Hyndman and Booth 2006), and alternative models for E(t) have been estimated and tested (UK Government Actuary 2001). The largest systematic application of the Lee-Carter method has been done by the European DEMWEL (demographic uncertainty and the sustainability of social welfare systems) project, covering 19 countries in and neighboring the EU; these results are soon to be published. A recently proposed new method to forecast mortality applies the extrapolative insight of the Lee-Carter method to the age distribution of deaths (Bongaarts 2005); the latter has not been tested against the Lee-Carter method. The power of the Lee-Carter result is that the linearity of the mortality trend in Equation 1 is revealed by an analysis of the variance in the data, and is not imposed. Most alternatives, including the Bongaarts approach, assert a model and fit its parameters; there has not been much systematic comparison of competing models.

Figure 1 shows forecasts of life expectancy at birth (derived by the Lee-Carter method) in the United States for both sexes combined, with a launch date in 2001. These are 100-year forecasts—it is a brave assumption that the model will hold for a century, but the model does accurately describe the past seven decades of mortality experience. The forecast is displayed as a set of three lines—the lowest indicates the 2.5 percentile level of the forecast in each year, meaning that the probability that the life expectancy falls below that lowest line in any year is 2.5 percent. The middle and upper lines are the median (50th percentile) and the 97.5 percentile. Thus, this forecast provides a range of outcomes that are predicted to contain the future life expectancy with high probability. In addition, this method provides any desired number of sample paths of forecast mortality rates, and in a simulation, these sample paths will have a probability distribution that is reflected in the prediction intervals shown in Figure 1. Note the 95 percent prediction interval for life expectancy is 2.5 years wide after 25 years and five years wide after 50 years. The median life expectancy forecast 100 years out is about 87 years; by way of comparison, the life expectancies in both Japan and Sweden today are over 85 years.

Fertility

Fertility is described by age-specific birth rates per person per year for women between (usually) ages 15 and 50. The sum of the age-specific fertility rates over the reproductive ages is the total fertility rate (TFR). The most noticeable empirical fact about fertility in the past half-century or so is that it has varied enormously over the years. From the 1950s to the 1970s many industrialized countries experienced first a boom and then a bust in TFR, leading to large baby booms. Fertility change in recent decades in many industrialized countries has put a dent in a long-held belief that TFR should stabilize at or close to the population replacement value of 2.05. Among the rich industrialized countries, the United States now stands out with a relatively high (though below replacement) fertility that has held stable for over a decade; France is not far different. Virtually all other rich industrialized countries have experienced declines in fertility to TFR levels well below replacement. Although demographers have argued that this decline is a temporary consequence of much-delayed childbearing, TFR levels have stayed low for several years, and there seem to be no reasons to expect a sizeable rebound (Lesthaege and Willems 1999). Economists have advanced endogenous theories of fertility that predict swings in fertility in response to changing population composition, most famously Easterlin's (1976) cohort size story for the United States, but these have found little empirical support and do not seem useful in a serious forecasting model.

The volatility of fertility over time means that, in contrast to what we know about mortality, we do not have a strong trend to exploit in making forecasts. However, the age pattern of fertility (i.e., the relative, as opposed to absolute, fertility of different ages) changes relatively slowly over time, even though the levels (TFR) can change relatively rapidly over time. In making a forecast, it seems sensible to expect that past volatility in fertility is likely to continue in the coming decades. Given this state of understanding, Lee and Tuljapurkar (1994) proposed a stochastic model of fertility that is conceptually similar to the Lee-Carter mortality model but has very different dynamics. The model involves two equations,

$$F(x,t) = A(x) + B(x)\,F(t), \qquad \text{(Equation 2)}$$

$$F(t) = F^* + c(F(t-1) - F^*) + u(t) + d\,u(t-1) \qquad \text{(Equation 3)}$$

Equation 3 is a stochastic model for fertility F(x,t) of age x in year t, A(x) and B(x) are constant age-specific schedules; the level of fertility is governed by the time-factor F(t) and the constant F*, and the u(t) are stochastic innovations. Equation 2 takes the output from Equation 3 and translates it into fertility change. A key assumption in this model is that the expected value of the TFR will converge over time to a specified value F*. This is a subjective assumption made in all long-run fertility forecasts that I am aware of; Tuljapurkar and Boe (1997) examine the consequences of assuming that the long-run mean, F*, is drawn from a prior distribution that reflects our uncertainty about long-run trends. We estimate the fertility model starting with an a priori specification of the long run constraint, F*, and historical data. Just as with mortality, this estimation process yields a forecast engine that generates sample paths, F(x,t,w), of age-specific fertility over time. An alternative approach to forecasting fertility (Alho and Spencer 1997) uses a more flexible parameterization of the stochastic process that generates uncertainty over time, but yields similar qualitative predictions.

Figure 2 displays probabilistic forecasts of TFR for the United States with a launch date of 2003. As with the life expectancy forecasts in Figure 1, the plot displays several forecast lines: shown here are five annual percentile bounds ranging from a 2.5 percent low to a 97.5 percent high. In contrast with Figure 1, in which mortality forecasts display a modest rate of increase in uncertainty with forecast horizon, the uncertainty range for TFR in Figure 2 rises extremely rapidly with forecast horizon. It is important to note that any actual forecast will not follow the percentile lines, but will fluctuate over time between the highest and lowest percentiles. It is also worth pointing out that the 2.5 percentile of projected TFR approaches about 1.0, a value close to that currently observed in South Korea, Italy, and several other countries, whereas the 97.5 percentile of projected TFR approaches about 3, a value substantially below United States TFR at the peak of the baby boom. Thus, the uncertainty shown in Figure 3 may be high, but is certainly consistent with historical experience.

Migration

This is the final component of vital rates and is driven largely by policy so that it does not make much sense to model the historical dynamics. For United States forecasts, we adopt a scenario method and use (with equal probability) high-medium-low immigration alternatives in line with either the United States Census Bureau or the Social Security Administration. The medium scenario is simply a continuation of recent past patterns and includes both legal and illegal migrants.

3. POPULATION PROJECTIONS

Given stochastic models for mortality and fertility, as well as migration scenarios, we start with the population in a launch year and generate a large number of sample paths of the vital rates and then corresponding sample paths of population forecasted by age and sex over time. Using a large number of sample paths (several thousand) allows us to estimate prediction intervals that contain any desired probability for population numbers and ratios. It is worth noting that, in general, the forecast uncertainty obtained using stochastic methods is much wider than the uncertainty indicated by typical scenario forecasts.

The effects of combining stochastic projections of mortality, fertility, and immigration into a population forecast for the United States are illustrated by considering several aspects of the projected population. A much broader set of forecasts for the G-7 countries is presented and discussed in Tuljapurkar, Anderson, and Li (2001). Figure 3 displays percentile prediction intervals for United States total population starting in 2003. Note that the uncertainty in the forecast increases with the forecast horizon, as indicated by the growing width of the 95 percent prediction interval relative to the median forecast—for total population the 95 percent prediction interval is about 20 percent of the median forecast after 25 years and increases to 40 percent after 50 years.

Figure 4 displays probability percentiles for projections of the old-age dependency ratio (population over 65 divided by population 20 to 64). The size of this old age dependency ratio is a key determinant of the balance of flows in any transfer system for old age populations, whether for pensions or health care. As is well known, the net balance of payments in any such transfer system depends on a product of two ratios: the dependency ratio and the ratio of per-capita transfer payments to per-capita receipts. The projected dependency ratio has small uncertainty until 2030 because most of the change in this ratio until then is due to the aging of the baby boom. Note that the aging of the baby boom will, with near-certainty, increase the dependency ratio from its 2003 value of about 0.2 to a value of 0.35 by 2030. The prediction intervals for the dependency ratio are not as wide as for total population; the 95 percent prediction interval in Figure 3 is about 0.1 of the median projection after 25 years and just over 0.2 after 50 years. It is striking that even after the baby boom has largely passed, from 2050 onwards, the probability that the dependency ratio will fall below 0.35 is only about 1/6—a striking result given the large uncertainty about the predictions.

Figures 5, 6, and 7 display the uncertainty differently, in terms of population pyramids shown at 5, 25, and 50 year forecast horizons. Each figure shows a standard population pyramid with age groups arranged vertically for ages up to 105 years—for simplicity, I show the total population at each age and do not distinguish the two sexes. For each pyramid, the colors indicate percentiles at 2.5 percent (yellow), 50 percent (blue), and 97.5 percent (red) for the forecasted population in each age interval. Observe that the largest numerical uncertainty is at the younger ages and is driven by the substantial uncertainty about fertility that is illustrated in Figure 2. As we move further out in time, from 5 to 25 to 50 years, the uncertainty band due to fertility at the younger ages moves up the pyramid. The uncertainty in mortality projections, shown by the multicolored bands at the highest ages, is modest by comparison to the uncertainty in fertility, although it is not in fact trivial in absolute terms. The pattern of growth of uncertainty in these pyramids corresponds directly with the growth of uncertainty shown for the old-age dependency ratio in Figure 4—that ratio becomes increasingly uncertain as the number of working age people becomes uncertain, i.e., at a time horizon greater than 20 years.

These stochastic forecasts give us a complete set of probabilistic projections of population number and composition over time. Since the method generates a probability distribution over sample paths (i.e., possible futures), there is complete consistency of the prediction probabilities for numbers in various age segments (e.g., children or retirees) with prediction probabilities for ratios such as the old-age dependency ratio.

4. PENSION SYSTEMS AND RISK

The approach we take to model pension systems and other fiscal expenditures (Lee and Tuljapurkar 2001; Tuljapurkar, Anderson, and Lee 2001) builds on the work of the Social Security Administration (SSA) actuaries in the United States, but is simpler in that we do not disaggregate as much as they are required to do. We follow SSA in modeling both disability and old-age support to produce a model of the United States OASDI system. The analytical structure is straightforward—we project labor force participation rates, productivity growth, real interest rates, and returns on the stock market (the S&P 500) using multivariate autoregressive models that test for correlations between the series. The economic models are estimated using historical data in conjunction with expert opinion where necessary—for example, we constrain the long-run behavior of real interest rate and labor force participation rates. These models are used to generate stochastic sample paths for the economic variables by Monte Carlo simulation. To complete the model, we need tax schedules that determine payments into the system and benefit schedules that determine payments out of the system. We take tax and benefit schedules to be given by existing law, including the indexation of wages for the computation of benefits, and the indexation of benefits. We follow SSA in using historical data to estimate retirement patterns, in particular the fractions of each cohort that opt to receive benefits at different ages in the retirement window. We also follow SSA assumptions concerning behavioral responses to planned changes in retirement age and incentives.

The core of our model tracks the balance B(t,w) in the Social Security trust fund on every sample path using our forecast of economic and demographic variables,

$$B(t,w) = (1 + r(t,w))*B(t-1,w) + \{\text{Tax receipts } t-1,w\} - \{\text{Benefits paid } t-1,w\}$$

(Equation 4)

Given a launch year in which we know the holdings of the fund, and some forecast horizon, we employ the stochastic population projections to generate a sample path of future population by age and sex in each year, and the economic models to generate a sample path for each of interest rate, labor force participation, wages, taxes, beneficiary status, and benefit payment schedules by age in each year. The latter schedules are applied to the projected population to generate the flows on the right side of Equation (4). Thus, the final output is a sample path of the flows into and out of the trust fund, as well as the level of the fund itself. The resulting set of forecast sample paths can be used to assess the state of the system. To explore the effects of privatization, we use our joint stochastic model of interest rates and the rate of return on the S&P 500 index of equities. For any particular privatization scheme, we modify Equation (4) appropriately to incorporate both risky and risk-free components of the fund.

The way in which we use these projections is to ask: given some fiscal quantity, what is its predicted probability distribution over time? An illustration is given in Figure 8, which describes the probability distribution of the projected actuarial balance under current law. This is a summary measure that can be interpreted as the additional tax—as a percentage of income that would balance the flows in Equation 4 over a specified time horizon. The histogram shown in Figure 8 indicates the distribution of projected actuarial balances for the United States OASDI over a 75-year horizon. The vertical lines indicate the SSA low, medium, and high actuarial balances. What is striking is that even with the substantial uncertainty in economic and demographic variables, there is a projected probability of over 50 percent that the additional tax would lie between 1 percent and 3 percent. In addition, the probability that fortuitous combinations of future conditions (e.g., higher than expected wage growth, high interest rates, high fertility) will keep the actuarial balance at or above zero is only about 0.05. The probabilistic approach gives us a nuanced understanding of how economic and demographic changes affect the system over time, and also gives us the odds are that they will generate particular outcomes.

Another, and possibly more important, application of this model is to explore the impact of policy changes on particular outcomes of interest. To model a policy change, say an increase in the normal retirement age (NRA), requires that we incorporate not simply a shift in benefit schedules, but also any resulting shift in behavior. Thus, an increase in the NRA may change the probabilities with which people opt to claim benefits at different ages, their labor force participation rates, and also payouts from the disability program. We proceed by incorporating the responses or mechanisms predicted by other models (behavioral, endogenous economic, empirical) to provide information on such shifts. Our modeling approach would also allow us to incorporate uncertainty about these responses although we have not done so.

I provide two illustrations of the analysis of policy options based on Anderson, Tuljapurkar, and Lee (2003). In these examples, I focus on the solvency of the trust fund as a criterion of interest; this is not the most important criterion but is simply convenient for illustrative purposes. Solvency simply tells us whether in any given year, t, the balance in Equation 4 is positive. The probability of solvency in any future year, t, is found directly by counting the fraction of sample paths, w for which $\{B(t,w) > 0\}$, and the answer will clearly depend on the time horizon, here indicated by t. In the illustration that follows, I use a time horizon of 50 years and examine the probability that the fund is solvent in each of those 50 years. Under current law, it is well known that the fund will almost certainly be insolvent before these 50 years are up, so we will use the models to ask what happens to the probability of solvency if we change policy in specified ways.

As a first example, we ask how solvency is affected by a combination of two policy changes to be implemented immediately—one is to raise the taxes collected for Social Security, the other is to invest some fraction of the trust fund (collectively) in the S&P 500 index. Figure 9 shows how the probability of solvency changes in response. The current situation is indicated by the origin where we have neither additional taxes nor investment, and shows the probability of solvency to be essentially zero. If we move along either of the horizontal axes, we are tracking either an increase in tax by the specified percentage, or the immediate investment of a specified proportion of the trust fund surplus into equities. As we move into the interior, we are tracking the joint effect of both kinds of change. An interesting feature of the plot is the synergy between the two policy variables. A tax increase of 1 percent by itself raises solvency probability from near zero to about 40 percent. If we also invested 20 percent of the trust fund in the stock market, we get up to a solvency probability of about 60 percent.

As a second illustration, we again consider the probability of solvency over 75 years, but examine the effect of a different pair of policies. One policy, as before, is to raise the taxes going into Social Security, the other is to increase the normal retirement age (NRA) to 67, 68, or 69 on a more accelerated schedule. The final NRA values and the target dates for their achievement are shown in the figure. In the baseline case at the origin we have current law taking us to an NRA of 67 in 2022. Other possibilities shown are achieved by a steady increase over the periods indicated starting immediately. The results of these changes on solvency over 50 years are shown in Figure 10. Here again we observe synergies between the two policy options and they are even bigger than those in Figure 9. A tax increase of 1 percent by itself raises solvency probability from near zero to about 40 percent. If we also increase the NRA steadily up to age 69 by 2024, we get up to a solvency probability of about 80 percent.

The analyses illustrated in Figures 8 and 9 can be conducted for other objective functions under other policy scenarios. These analyses reveal unexpected synergies between policy options and also the importance of the timing of policy change. The analyses also highlight the fact that many policy options can improve the fiscal situation of the OASDI program but still leave us with a substantial probability that some target, such as solvency, will not be achieved. Thus, we can rank alternative policies in terms of the probabilities that they will achieve some set of targets.

5. MODELING HEALTH EXPENDITURE AND MODELING POLICY

Public health care spending can be modeled in much the same way as we model the OASDI system. The key step is to determine age-schedules of expenditure by sex and to describe how these schedules will evolve over time. As the population ages, we expect a long-run shift in the age profiles themselves—the large expenditures that characterize the last two years of life will progressively occur at later ages. In addition, and this is critical, we must project changes in the absolute level of expenditure at all ages over time. The latter is well known to be driven by cost inflation in the short run, but may also change over time in response to trends in health. It is easy to see that if costs increase faster than wages then the transfers needed to pay for these costs must also increase, and eventually any such system will be unsustainable. Forecasts usually analyze alternative long run scenarios of change while incorporating historical volatility as we did with our models of fertility. Lee and Tuljapurkar (2000) discuss and illustrate these and more general aspects of the forecasting of fiscal expenditures for a variety of government programs.

I argue that probabilistic methods provide a rich and necessary tool for the analysis of public policy, especially in settings that involve intergenerational transfers and, thus, require analyses over long time horizons. In such settings, uncertainty will usually grow over time and policies must be evaluated in terms of their probability of achieving targets of various kinds (Tuljapurkar 1997; Anderson, Tuljapurkar, and Lee 2003). The methods described here provide a natural setting in which to evaluate multiple policy objectives and also to identify policy combinations that are optimal in the sense that they maximize the probability of achieving particular goals.

REFERENCES

Alho, J. M. and B. D. Spencer. 1997. "The practical specification of the expected error of population forecasts." *Journal of Official Statistics* 13(3): 203–225.

Anderson, M., S. Tuljapurkar, and R.D. Lee. 2003. "Chances are … Stochastic forecasts of the social security trust fund and attempts to save it." Presented at the 1999 Conference on Retirement Research, Center for Retirement Research, Boston College. Available online from the Michigan Retirement Research Center.

Anderson, M., S. Tuljapurkar, and N. Li. 2001. "How accurate are demographic projections used in forecasting pension expenditure?" in T. Boeri, A. Borsch-Supan, A. Brugiavini, R. Disney, A. Kapteyn, and F. Peracchi (eds.) *Pensions: More Information, Less Ideology*. Dordrecht, The Netherlands: Kluwer Academic Publishers

Bongaarts, John. 2005. "Long-Range Trends in Adult Mortality: Models and Projection Methods." *Demography* 42(1): 23–49.

Booth, H., J. Maindonald, and L. Smith. 2002. "Applying Lee-Carter under conditions of variable mortality decline." *Population Studies* 56(3): 325–336.

Disney, R. 1996. "Can we afford to grow older? A perspective on the economics of ageing." Cambridge, MA: MIT Press.

Easterlin, R. A. 1976. "The conflict between aspirations and resources." *Population and Development Review* 2(3/4): 417–425.

Hyndman, R. J. and H. Booth. 2006. "Stochastic Population Forecasts Using. Functional Data Models For Mortality, Fertility And Migration." Working Papers in Demography 99. Australian National University. Available at: demography.anu.edu.au/Publications/WorkingPapers/99.pdf .

Lee, R.D. and L. Carter. 1992. "Modeling and Forecasting U.S. Mortality," *Journal of the American Statistical Association* 87(419): 659–671

Lee, R. D. and T. Miller. 2001. "Evaluating the Performance of the Lee-Carter method for Forecasting Mortality." *Demography* 38(4): 537–49.

Lee, R. D. and S. Tuljapurkar. 2000. "Population forecasting for fiscal planning: Issues and innovations." in A. Auerbach and R. Lee (eds.) *Demography and Fiscal Policy*. Cambridge, U.K.: Cambridge University Press.

Lesthaeghe, R. and P. Willems. 1999. "Is Low Fertility a Temporary Phenomenon in the European Union?" *Population and Development Review* 25(2): 211–228.

Oeppen, J. and J. W. Vaupel. 2002. "Broken limits to life expectancy." *Science* 296(5570): 1029–31.

Thatcher, A.R., V. Kannisto, and J.W. Vaupel. 1998. *The Force of Mortality at Ages 80 to 120*. Odense: Odense University Press.

Tuljapurkar, S., N. Li, and C. Boe. 2000. "A Universal Pattern of Mortality change in the G7 Countries." *Nature* 405(6788): 789–92.

Tuljapurkar, S. 1997. "Uncertainty In Demographic Forecasts: Methods and Meanings." Proceedings of the 1996 Conference on Applied Demography, Bowling Green State University, Bowling Green, Ohio.

Tuljapurkar, S. and C. Boe. 1999. "Validation, information and probabilistic priors in stochastic forecasts." *International Journal of Forecasting* 15(4): 259–271.

Tuljapurkar, S., M. Anderson, and R.D. Lee. 2003. "Effects of Changing the Social Security Retirement Age on the Dynamics of the OASDI Program." Presented at the October 1998 meeting of the Social Security Advisory Board, Capitol Hill, Washington D.C. Available online from the Michigan Retirement Research Center.

UK Government Actuary. 2001. "National Population Projections: Review of Methodology for Projecting Mortality." Government Actuary's Department, *National Statistics Quality Review Series* Report No. 8.

White, K. M. 2002. "Longevity Advances in High-Income Countries." *Population and Development Review* 28(1): 59–76.

FIGURE LEGENDS

Figure 1. Life expectancy at birth (sexes combined) in the United States forecast using the Lee-Carter method. The three lines shown are the 2.5, 50, and 97.5 percentile lines of the probability forecast.

Figure 2. Total fertility rate in the United States forecast using the stochastic model described in the text. The lines shown are the 2.5, 16.7, 50, 83.3, and 97.5 percentile lines of the probability forecast.

Figure 3. Probabilistic forecast of the total population of the United States. The lines shown are the 2.5, 16.7, 50, 83.3, and 97.5 percentile lines of the probability forecast.

Figure 4. Probabilistic forecast of the old-age dependency ratio of the United States. The lines shown are the 2.5, 16.7, 50, 83.3, and 97.5 percentile lines of the probability forecast.

Figure 5. Age pyramid drawn from a probabilistic forecast of United States population at a time horizon of ten years. The bars show total population by age, the inner and outer pyramids contain 95 percent of the probability projected.

Figure 6. Age pyramid drawn from a probabilistic forecast of United States population at a time horizon of twenty-five years. The bars show total population by age, the inner and outer pyramids contain 95 percent of the probability projected.

Figure 7. Age pyramid drawn from a probabilistic forecast of United States population at a time horizon of fifty years. The bars show total population by age, the inner and outer pyramids contain 95 percent of the probability projected.

Figure 8. Histogram of summarized actuarial balance in a stochastic projection of the United States OASDI system. The balance is computed over a 75-year time horizon.

Figure 9. Probability of solvency of the OASDI trust fund over a 50-year period in response to tax increases and investment of some fraction of the trust fund balance in equities.

Figure 10. Probability of solvency of the OASDI trust fund over a 50-year period in response to tax increases and more rapid increases in the normal retirement age.

FIGURE 1: Life Expectancy at Birth (sexes combined) in the United States

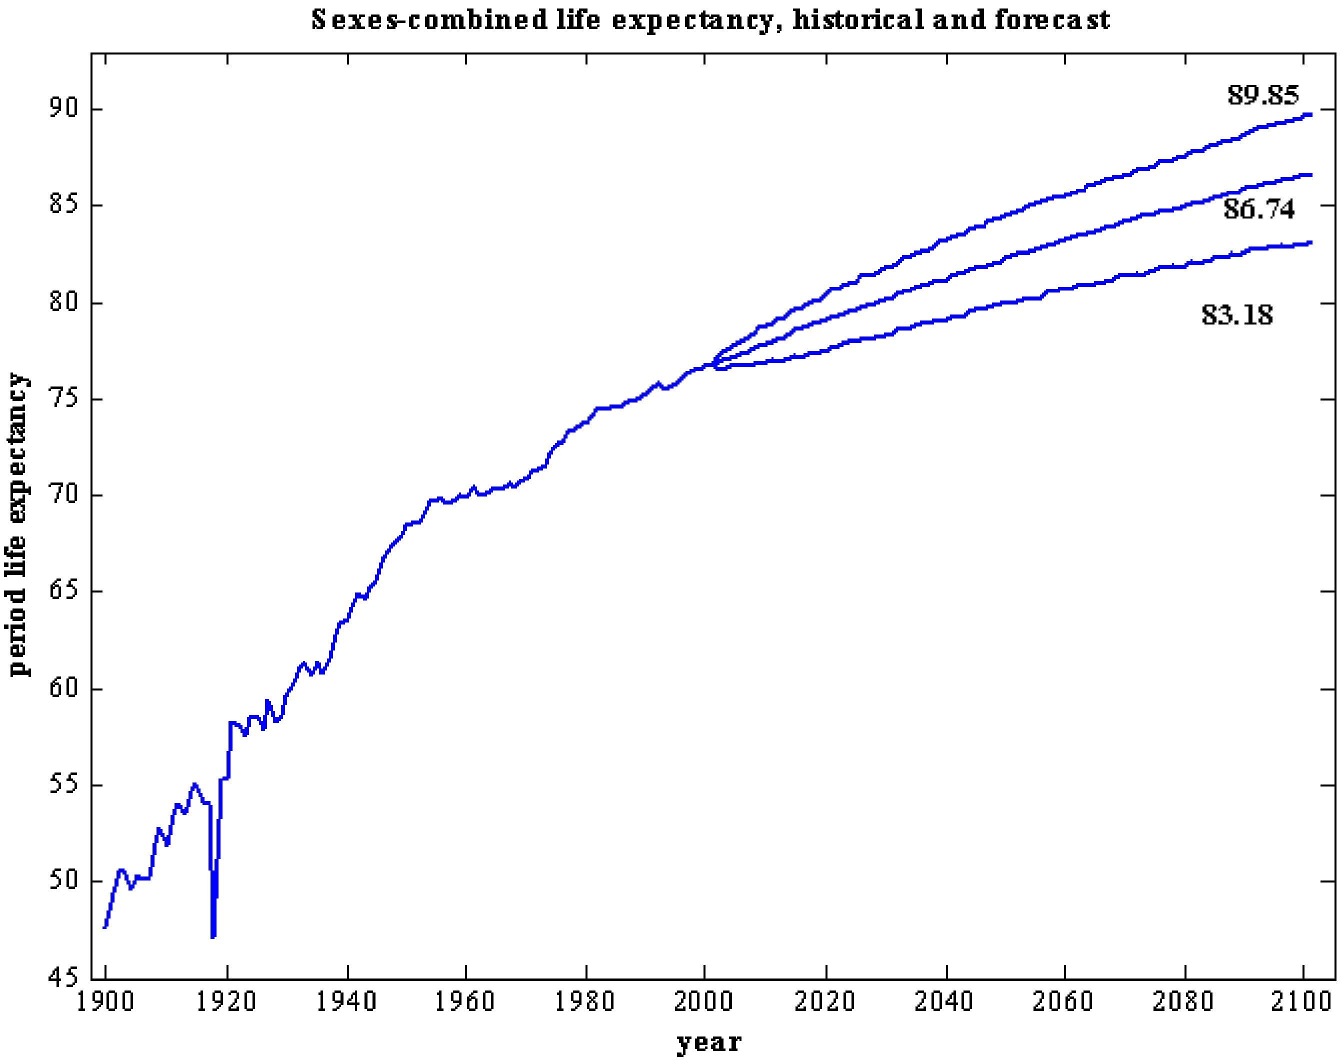

FIGURE 2: Total Fertility Rate in the United States

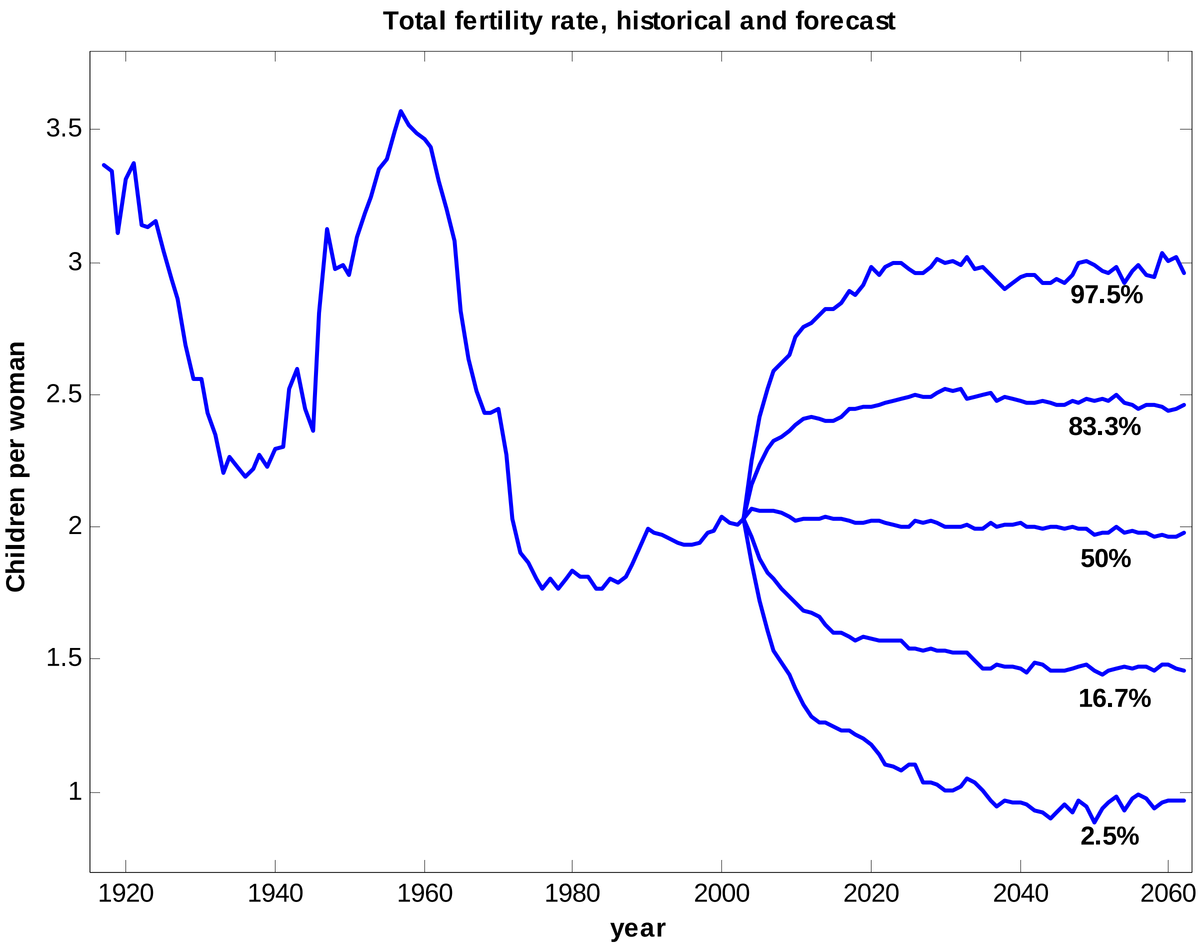

FIGURE 3: Probabilistic Forecast of the Total Population of the United States

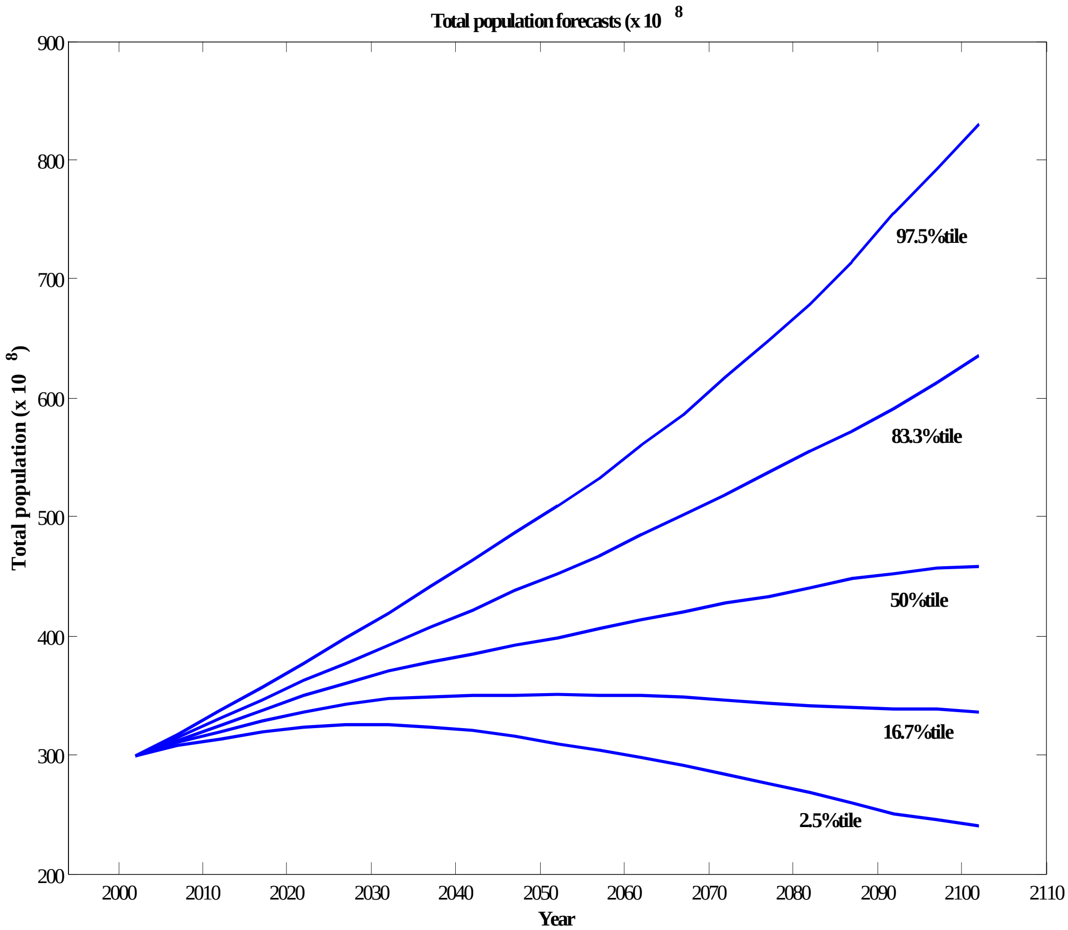

FIGURE 4: Probabilistic Forecast of the Old-Age Dependency Ratio of the United States

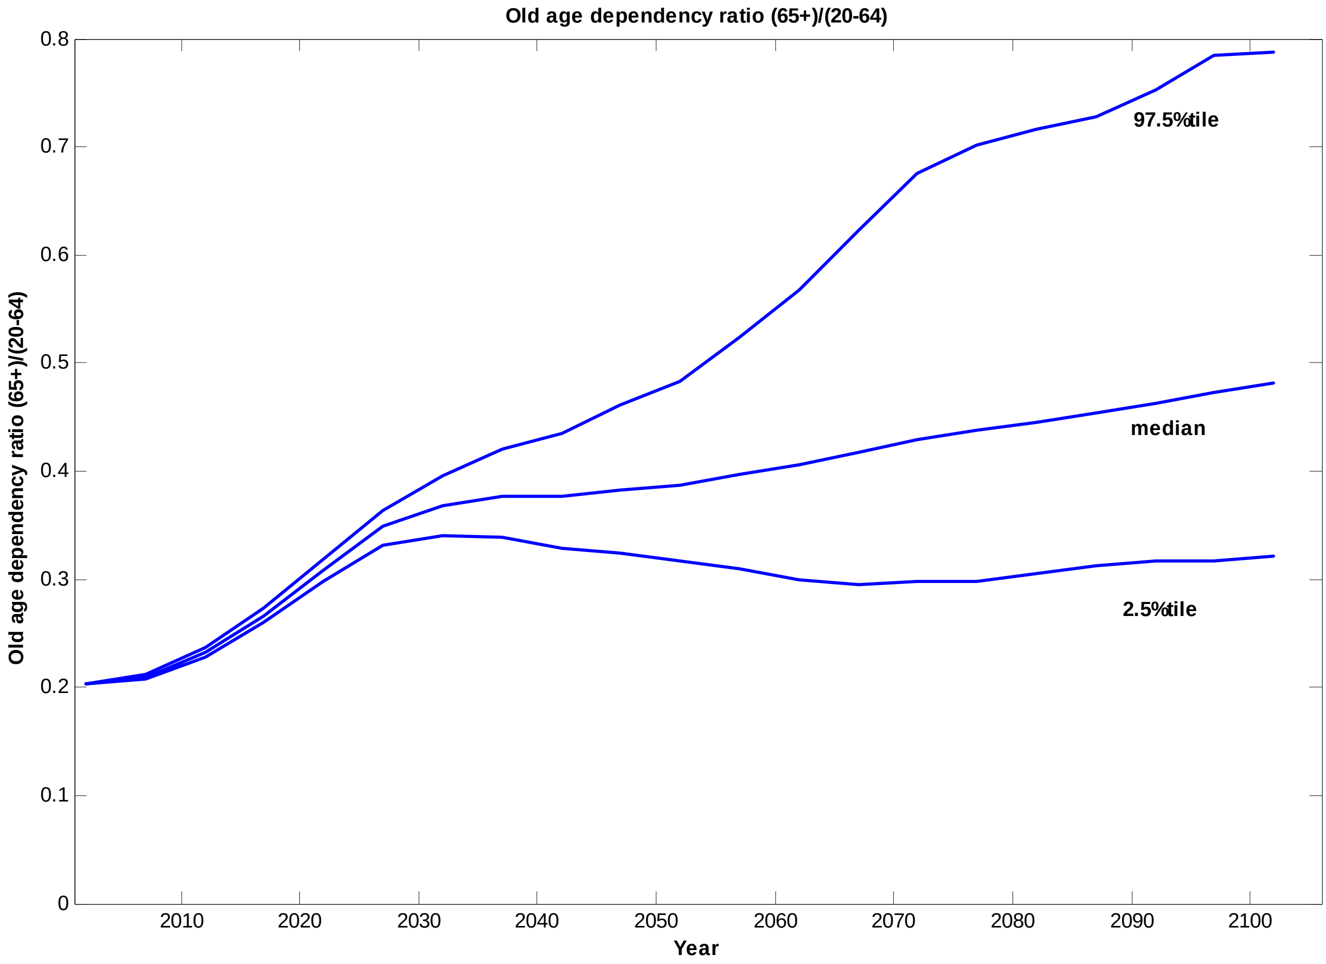

FIGURE 5: Age Pyramid Drawn from a Probabilistic Forecast of United States Population at a Time Horizon of Ten Years

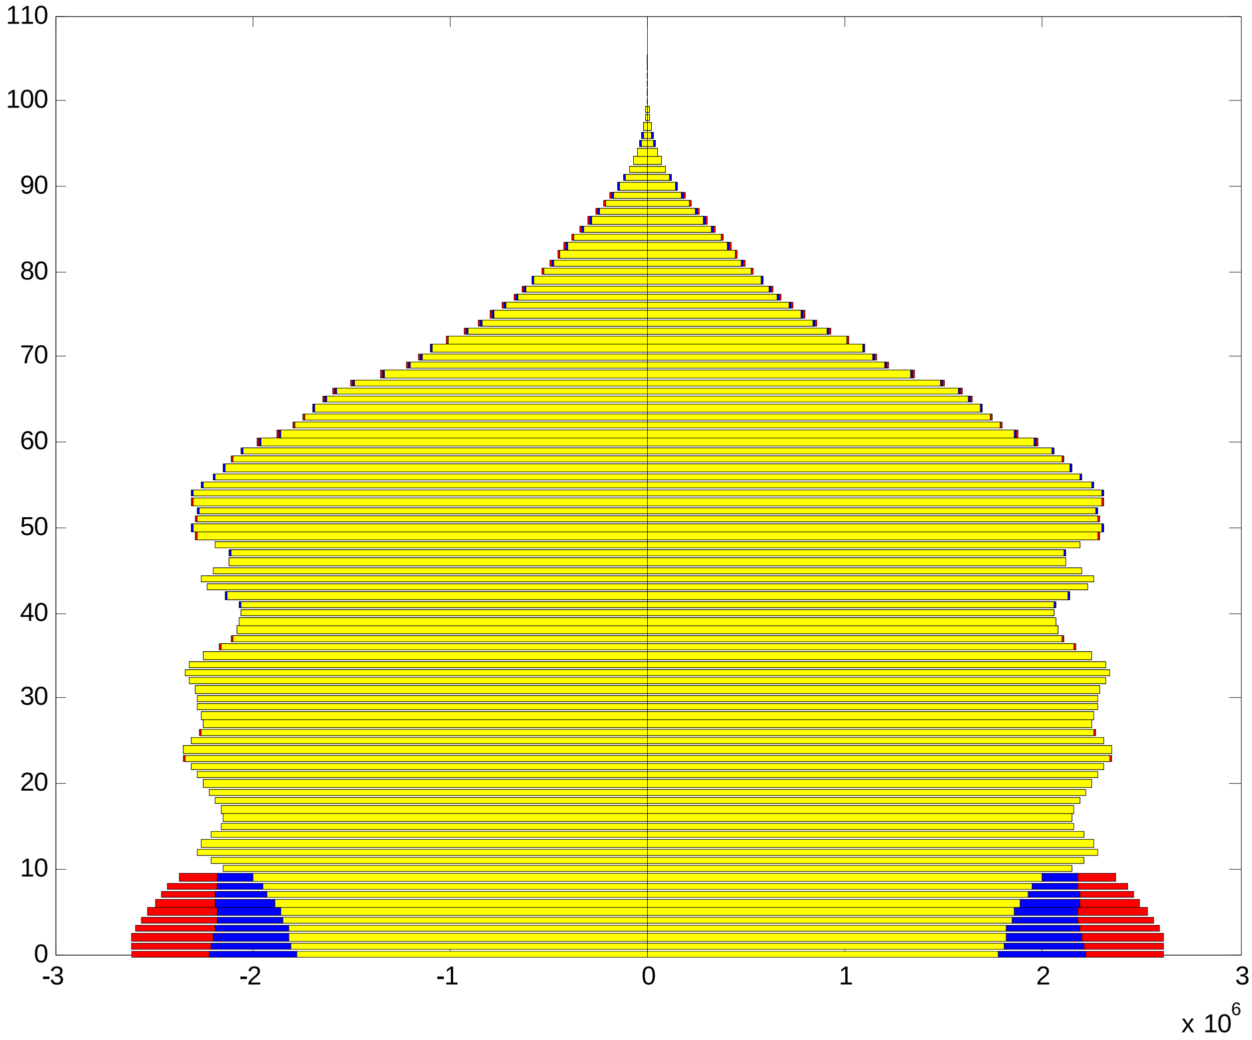

FIGURE 6: Age Pyramid Drawn from a Probabilistic Forecast of United States Population at a Time Horizon of Twenty-Five Years

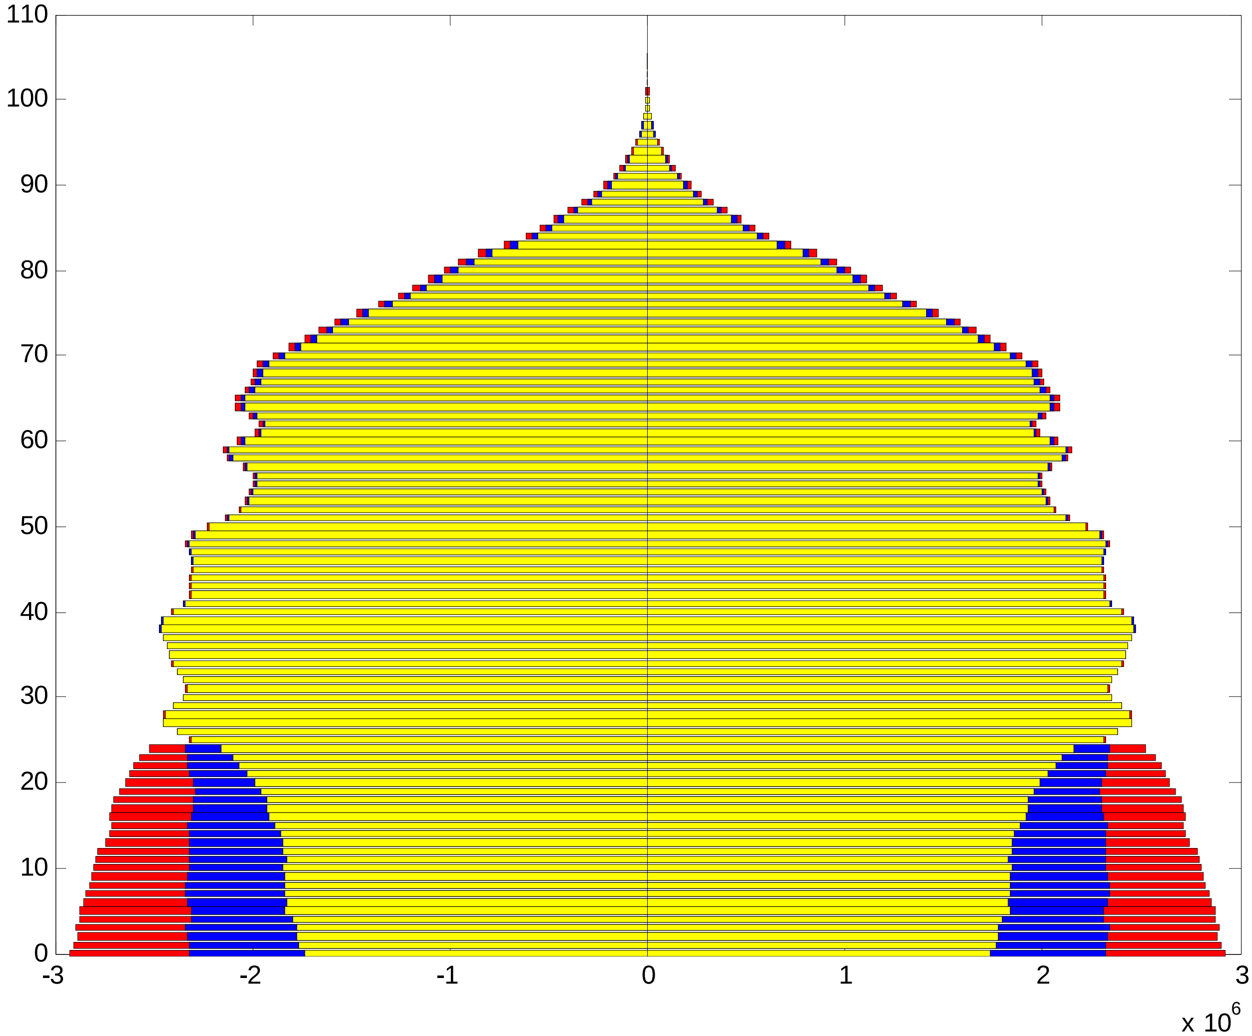

FIGURE 7: Age Pyramid Drawn from a Probabilistic Forecast of United States Population at a Time Horizon of Fifty Years

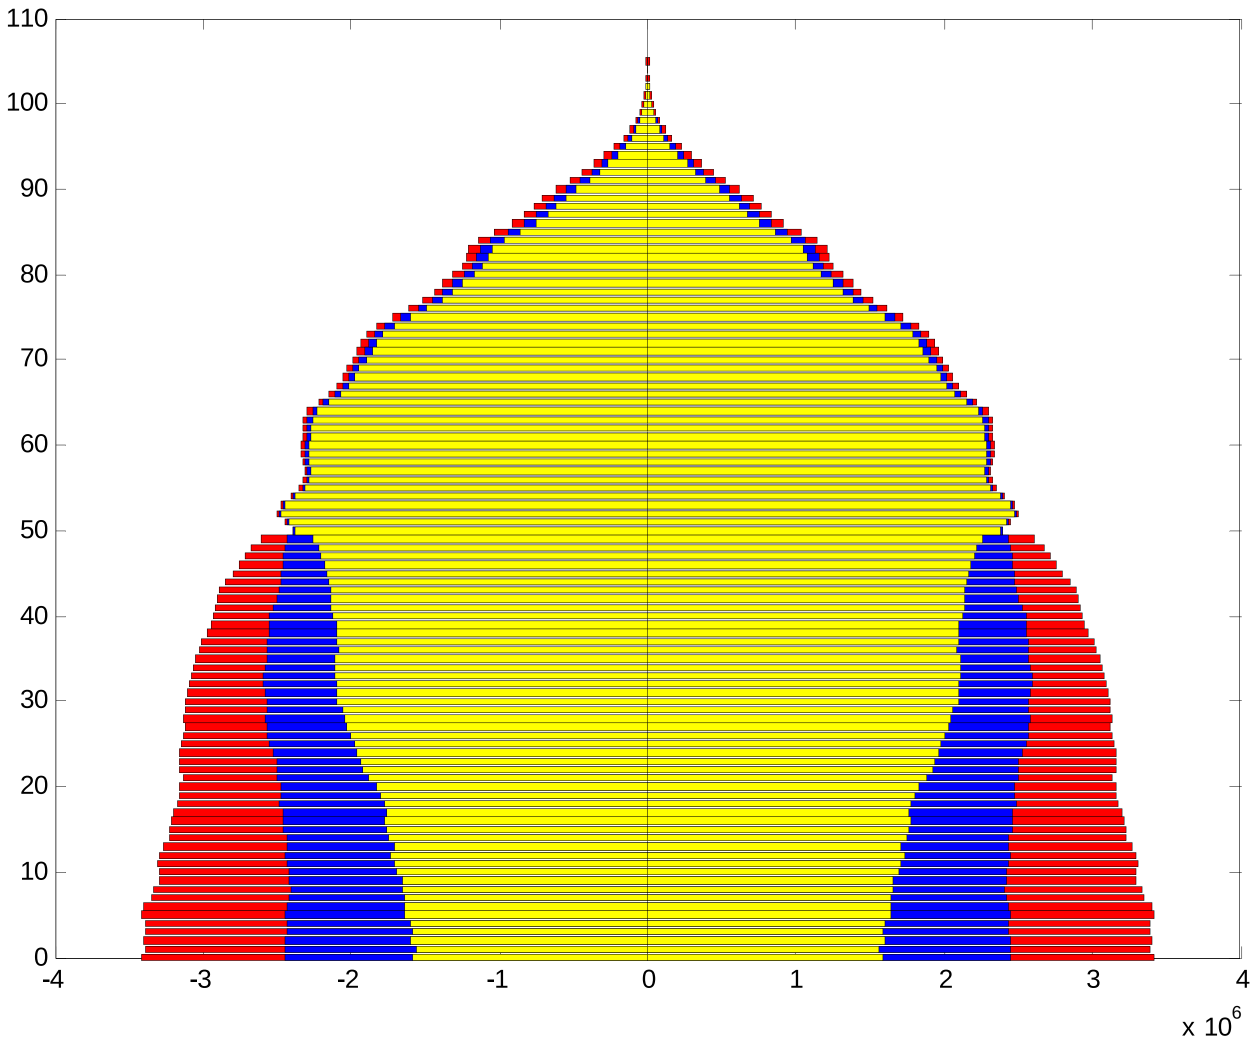

FIGURE 8: Histogram of Summarized Actuarial Balance in a Stochastic Projection of the United States OASDI System

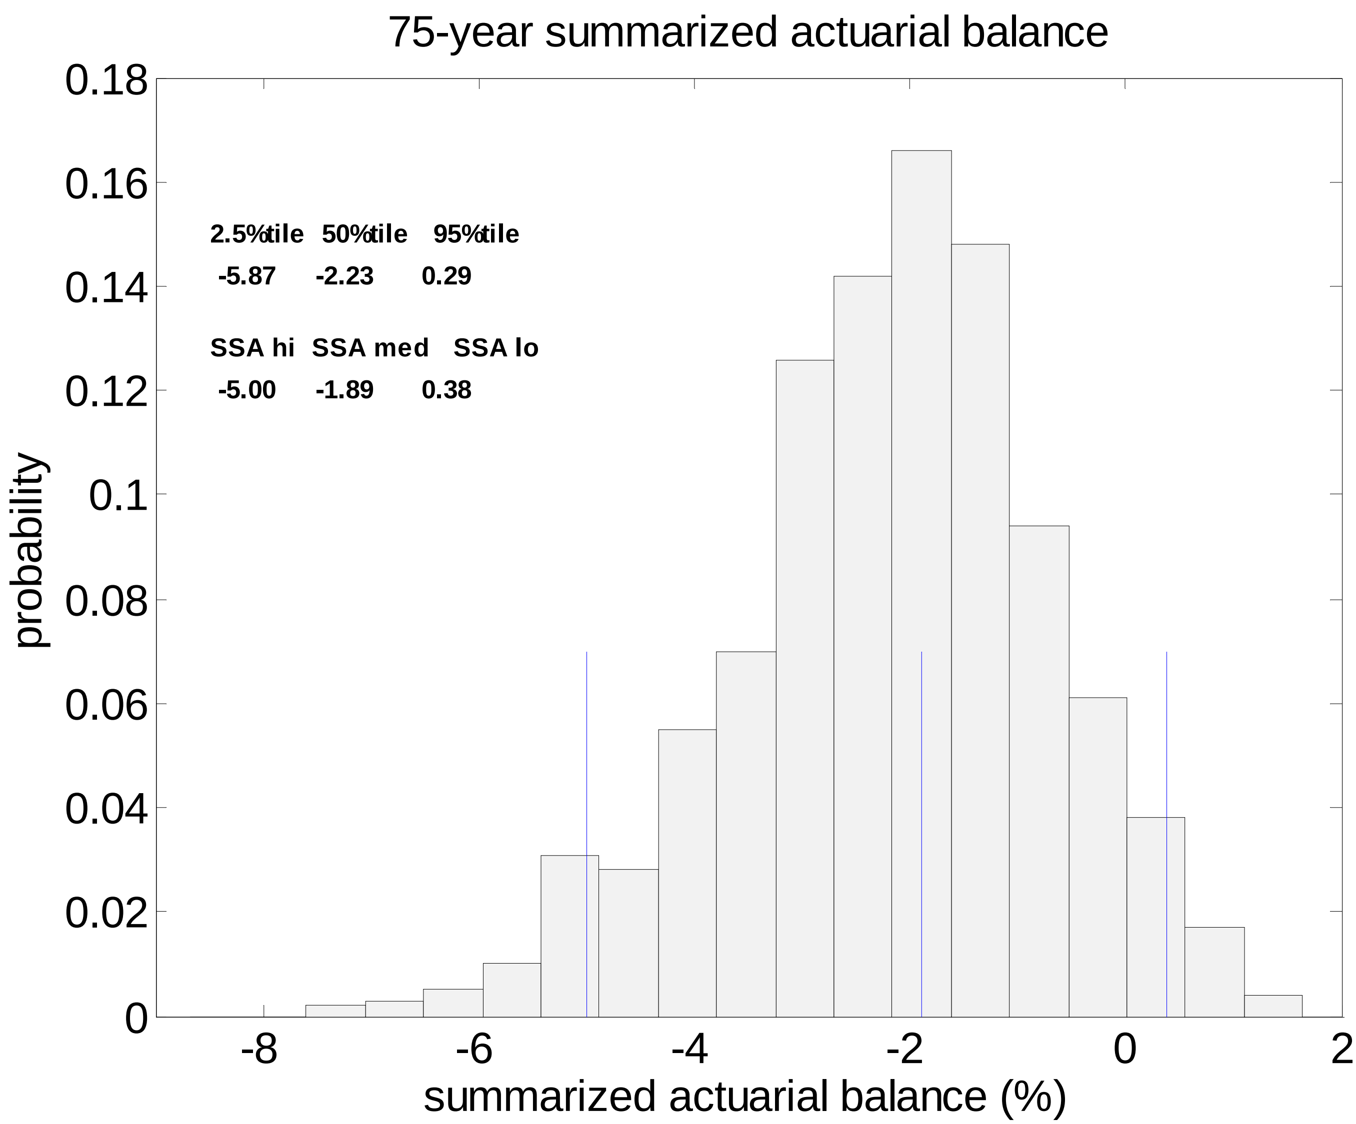

FIGURE 9: Probability of Solvency of the OASDI Trust Fund Over a 50-Year Period in Response to Tax Increases and Investment of Some Fraction of the Trust Fund Balance in Equities

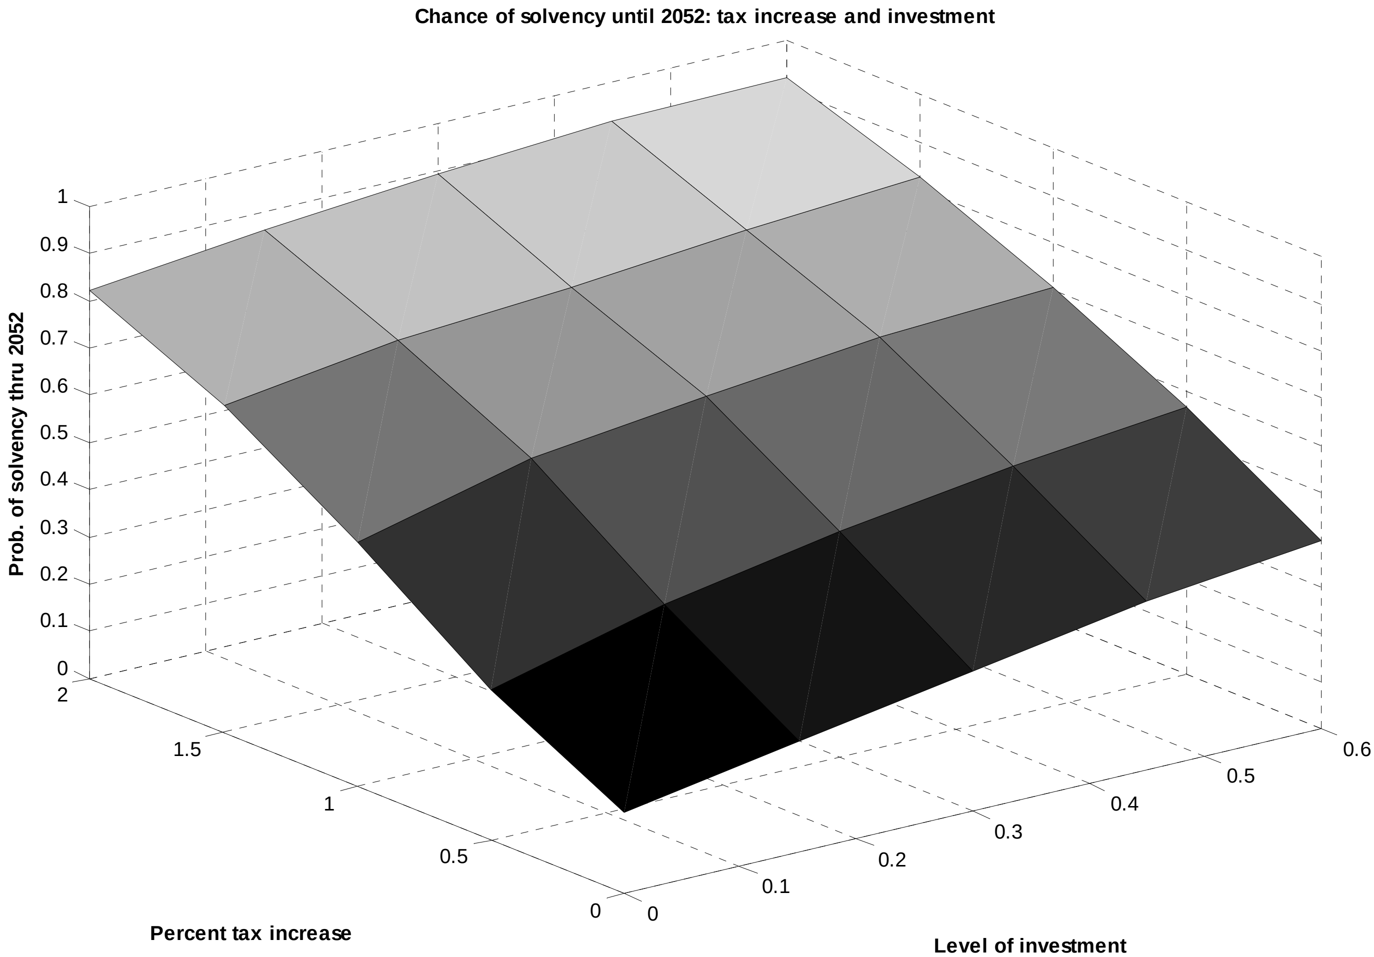

FIGURE 10: Probability of Solvency of the OASDI Trust Fund over a 50-Year Period in Response to Tax Increases and More Rapid Increases in the Normal Retirement Age