CPB Memorandum

CPB Netherlands Bureau for Economic Policy Analysis

Sector(s)	:	Arbeidsmarkt en Welvaartsstaat
Unit(s)/Project	:	Pensioenen: Risicowaardering
Author(s)	:	D. Peter Broer
Number	:	241
Date	:	March 16, 2010

Macroeconomic Risks and Pension Returns

This paper assesses the risk factors associated with the return to both funded and unfunded social security. The paper distinguishes productivity risk (as a proxy for wage risk), demographic risk, and asset return risk in its various forms. It provides a survey of the literature that discusses these risk factors, and it quantifies these risks by means of a VAR models that generates a *joint* distribution of these risk factors.

It appears that there is only limited mean reversion in both productivity growth and asset returns, so that risks continue to grow with the length of the horizon. As a result, the level of pension benefits becomes progressively more uncertain over longer time horizons. For PAYG systems, the main source of uncertainty in benefits is productivity risk. The long-run annual benefit growth risk is 2%, about the same size as the long-run expected growth rate. In a PAYG DC scheme expected benefit growth rates are slightly negative over the next few decades.

Equity-funded pensions are substantially more risky than PAYG pensions, with a long-term annual standard deviation of returns of 14%, but they offer substantially larger expected returns as well. The contribution of demographic risk factors to pension return risk appears to be limited, because the effect of demographic changes on dependency ratios can be predicted several decades ahead.

Contents

1 Introduction

The return to pension plans is uncertain. Returns depend on productivity growth and changes in dependency ratios for a paygo scheme, on capital market returns for funded DC plans, and on a combination of these determinants for a funded DB scheme. All these determinants are intrinsically uncertain and, furthermore, a case can be made that the uncertainty increases with the time horizon (see Section 2.2). This affects the value of pension plans. The claims that arise from participation in a pension scheme can be seen as quasi-assets that are characterized by their risk-return profiles .[1] Hence, the risk inherent in a pension scheme is an important characteristic, since, together with the risks of private assets, it determines how well a household is diversified. This scope for diversification affects the value of the pension arrangement to the household.

This paper concentrates on the long-term development of the determinants of pension returns. It tries to quantify both the expected development and the associated uncertainties. For this, I use a state space approach that generates a *joint* distribution of these determinants. This distribution is used to evaluate the returns characteristics of a paygo pension system.

The structure of this document is as follows. In Section 2 I present the long-run determinants of pension returns and I review the literature on each of them. Section 3 discusses the structure of the statistical model to be used and in Section 4 I describe the implied joint distribution of pension return determinants. Section 5 concludes.

[1] See Matsen and Thøgersen (2004). Another characteristic of a pension scheme is its degree of actuarial fairness, see Lindbeck and Persson (2003). That aspect plays no role in this paper.

2 Risk Factors

The risk factors to be considered are productivity risk (as a proxy for wage risk), demographic risk, and asset return risk.[2]

Pay-as-you-go (paygo) pension plans levy contributions from the young, working-age, population and pays transfer income to the elder, retired, population. In principle, total transfers equal total contributions period by period. This makes the dependency ratio a potentially important risk factor. A paygo scheme may be either a defined benefit (DB) scheme, in which the transfer is guaranteed either unconditionally, or conditional on some other economic variable like a wage index or a price index, or a defined contribution (DC) scheme, in which the contribution rate is fixed.[3] These characteristics matter for the risk-sharing properties of the scheme. In a DB scheme, workers incur both productivity risk, through their wage income, and demographic risk, via their contribution rate, whereas the paygo benefits of retirees are only subject to productivity risk. In a DC scheme, workers incur only productivity risk, while paygo benefits are subject to both productivity risk and demographic risk. In addition the size of the risks differs with the length of the time period till retirement.

Funded pension systems primarily face asset return risk. Stock and bond price risks are the main source of variation in the market value of pension fund assets. In a purely private defined contribution scheme these market values directly affect the value of future pension benefits of the participants. In a pure defined benefit (DB) scheme, liabilities are also relevant, and liability risk depends primarily on the volatility in the risk-free rate (nominal or real, depending on the type of liability). For example, during the stock market crash of 2008, the funding ratio of most Dutch DB pension funds dropped by some 40%, and about half of the fall was related to the decline in the nominal risk-free rate. In either type of scheme, the participants are interested in the real value of their benefits, even though most DB schemes only guarantee a nominal amount. Hence, from the point of view of the participants, inflation risk is an important ingredient of the risk profile of the pension system, even if the liabilities of the pension fund are not subject to inflation risk.

Pension funds are typically long-run investors. As such, they are not primarily interested in fluctuations in returns at business cycle frequencies. Still, solvability restrictions imply that changes in volatility affect their value at risk. Since volatility is correlated over time, it stands as a separate risk factor. Finally, *valuation risk* of assets can be regarded as a separate risk factor as

[2] These risk factors also determine most of the resulting consumption risk of retirees. An important omission is health care costs.

[3] In addition, the scheme may be Bismarckian, with a transfer that is proportional to previous contributions, or Beveridgian, with an unconditional transfer. This is primarily relevant in terms of the actuarial fairness of the scheme.

well. Valuation risk occurs during periods of crashes and panics, which cause asset markets to thin out. It can either be treated as a separate risk factor or as a cause of a "fat tail," or non-normality, in the asset return distribution.

In summary, I distinguish the following risk factors in the return to pension plans

$\phi_{t,\tau}$	fertility of age group τ at time t
$\lambda_{t,\tau}$	mortality rate of age group τ at time t
ζ_t	total factor productivity
π_t	inflation
r_{f_t}	short-term (risk-free) rate[4]
r_{b_t}	mid-term bond rate [4]
r_{k_t}	capital returns
σ_r	valuation risk ("disaster risk")

Within the present framework, these risk factors are considered to be jointly exogenous.[5] A proper characterization of these risk factors therefore requires the specification of their *distribution*. As the risks may be correlated over time as well as cross-correlated, a multivariate autoregressive model of the risk factors offers a natural framework to describe this distribution. As a guideline in the selection of an appropriate model, I first survey the literature.

2.1 Demographic risk

Population growth rates are difficult to predict over long periods. Keilman et al. (2008), Table 2.1, present official forecasts of old-age dependency ratios in 2050 that shift by as much as 10%-points over the period 1994-2004. This large revision is an indication of non-stationarity of the underlying demographic process. The standard model for both mortality and fertility is Lee and Carter (1992), which is of the form

$$\ln \phi_{t+1,\tau} = \alpha_{\phi,\tau} + \beta_{\phi,\tau}\, \mu_{\phi,t} + \varepsilon_{\phi,t+1,\tau} \tag{2.1a}$$

$$\ln \lambda_{t+1,\tau} = \alpha_{\lambda,\tau} + \beta_{\lambda,\tau}\, \mu_{\lambda,t} + \varepsilon_{\lambda,t+1,\tau} \tag{2.1b}$$

Here (μ_ϕ, μ_λ) are the hidden state variables for fertility and mortality, respectively, and $(\phi_{t,\tau}, \lambda_{t,\tau})$ are the observed realizations of the rates of fertility and mortality by age τ. The $\varepsilon_{t,\tau}$ are i.i.d. random variables with mean zero. The hidden state develops according to a stochastic difference

[4] The risk-free rate is the discount on one-year Treasury bills. *Future* bill rates are uncertain so that the short-term risk-free rate is also a risk factor. The mid-term bond rate is measured as the return on bonds with a remaining maturity of five years.

[5] This assumption is justified in the small open economy case. In a closed-economy setting more fundamental risk factors, like entrepreneurial risk, need to be considered.

equation:

$$\mu_{i,t+1} = \delta_{i,1}\,\mu_{i,t} + \delta_{i,0} + \eta_{i,t+1} \quad i \in \{\lambda, \phi\} \tag{2.1c}$$

For mortality, the development of the hidden state variable appears to be described well by a random walk with drift ($\delta_{\lambda,1} = 1$). Estimates for the U.S. show that $\delta_{\lambda,0} \approx -0.365$ and $\mathrm{E}\left[\eta_\lambda^2\right] \approx 0.65$. For persons of age $\tau \gtrapprox 60$, $\beta_{\lambda,\tau} \approx 0.03$, so that the rate of mortality of the elderly falls on average by about 1% per year. Hári et al. (2008) estimate the Lee-Carter model using a Kalman filter approach. They show that the estimate of mortality trend δ_0 is rather uncertain if one allows it to be time-varying.

The fertility process is more difficult to model. μ_ϕ varies over time, but it is not a random walk. Lee and Tuljapurkar (1994) use an ARIMA(1,0,1) model for μ_ϕ with a restriction on mean fertility of 2.1 child per woman. The parameter estimates are then $\delta_{\phi,1} = 0.968$ and $\delta_{\phi,0} = 2.1 \times (1 - 0.968) = 0.067$. This results in a 95% long-term confidence interval for fertility from 1 to 3 children per woman. Obviously, this implies huge differences in possible population growth rates. In the "Uncertain Population of Europe" (UPE) project (Keilman et al., 2008), a similar approach is chosen, but Northern and Western European countries are assumed to have a long-run average fertility rate of 1.8 and Southern European countries an average fertility rate of 1.4. Again, this is indicative of a large amount of uncertainty in population growth rates.

In Appendix A, I show that the development of fertility of Dutch women can be explained reasonably well using a version of the Lee-Carter model. Compared to (2.1) above, the changes are that fertility is modelled per cohort, instead of by age, and that *two* states are used, one describing the total fertility of a cohort, and the other one the shift in the age distribution of fertility for that cohort. Modelling fertility shifts by cohort should agree better with changes in fundamental determinants of fertility than modelling fertility shifts by age. However, compared to Lee and Tuljapurkar (1994), the uncertainty about future fertility is just as large, basically because the current value of the fertility state vector is imperfectly observed (see Figure A.1 in Appendix A).

Another source of demographic uncertainty is migration. The UPE forecasts assume that the change in migration is normally distributed, $\Delta M_{t,\tau} = S_{t,\tau}\,(\eta_\tau + \delta_{t,\tau})$, where the scales $S_{t,\tau}$ are country-specific and η and δ are random variables. For the Netherlands $S_{t,\tau} \approx 2.0$, and $\eta \propto N(0, 0.3)$ and $\delta \propto N(0, 0.7)$.

Assessment

Demographic change is a non-stationary process. Current estimates of total fertility in Western Europe are below the sustainable reproduction rate, but the uncertainty is large enough that future population growth cannot be excluded. Dependency ratios too may be unstable, depending on the interplay between fertility and mortality.

2.1.1 Spill-overs from Productivity Risk to Demographic Risk

Becker et al. (1990) analyse the relation between intergenerational utility smoothing and fertility. In their model a higher rate of technical progress leads to a lower fertility rate. Galor and Weil (1996, 2000) argue that shifts in the rate of technological progress have raised the relative value of market time of women, and led to a decline in fertility, together with an increased investment in human capital. Alders and Broer (2005) show that persistent shocks in the *level* of productivity may lead to lower fertility. However, the size of the effect depends strongly on the size of the substitution elasticity between children and goods, about which little is known.

There is evidence that labour market conditions affect migration. In line with the Harris and Todaro (1970) model, wage levels and unemployment levels affect the inflow of immigrants (see e.g. Karemera et al. (2000); Clark et al. (2002); Mayda (2005)). As a full labour market model is outside the scope of this paper, I try to capture labour market effects through the productivity risk factor. This variable contains both a cyclical component to proxy unemployment variations and a deviation-from-trend effect to proxy for wage level effects.[6] It appears that the effects of productivity on migration are primarily cyclical; the long run elasticity of the share of immigrants in the Dutch population with respect to productivity is about 0.1%. Appendix B contains a description of the migration equations used in the VAR model.

Assessment

Labour market conditions have a clear effect on migration. These effects can be proxied by labour productivity. In the Netherlands, the long-term effect of productivity differences on migration are limited. The size, or indeed the sign, of any direct effect of productivity on fertility is rather uncertain. There is some evidence that sectoral shifts have facilitated the boost in the participation rate of women, which may have contributed to the fertility decline. As sectoral shifts are not included among the relevant risk factors, this spill-over will not be taken into account either.

[6] Effects of changes in conditions in the rest of the world on the net migration flow are neglected, even though they do contribute to total migration risk.

2.2 Productivity

In a standard growth accounting framework, output growth can be decomposed in the growth of inputs and total factor productivity (TFP) growth,

$$\dot{y}/y = \frac{\dot{L}}{L} + \frac{F_K K}{y}\left(\frac{\dot{K}}{K} - \frac{\dot{L}}{L}\right) + \frac{F_t}{y}$$
$$\frac{\dot{L}}{L} = \frac{\dot{H}}{H} + \frac{\dot{N}}{N}$$

where F_t/y represents TFP growth, H represents human capital per worker, and N denotes working-age population. Demographic risk affects production growth, as ageing has a direct effect on N, and possibly labour productivity, via an effect on H (see e.g. Fougère and Mérette, 1999). Comparing productivity differences between countries over the period 1820–1985, Prescott (1998) argues that TFP is the primary determinant of labour productivity, not capital intensity or human capital per worker.

TFP growth is not directly observable. In growth accounting, it is computed as a residual (the "Solow residual") from a decomposition of output growth (see Hulten, 2000, for a survey). In RBC models, TFP growth is estimated from the Solow residual by treating it as a latent variable. Let

$$y_t = \zeta_t X_t F[K_t, L_t] \tag{2.2a}$$
$$\ln \zeta_t = \rho \ln \zeta_{t-1} + \gamma t + \varepsilon_t \tag{2.2b}$$

where ζ_t denotes the stochastic component of productivity and X_t the deterministic component. The Solow residual is $\frac{\dot{S}}{S} = \frac{\dot{\zeta}}{\zeta} + \frac{\dot{X}}{X}$. Filtering out ζ_t from S_t, this approach should lead to an unbiased estimate of both components of productivity.[7] Empirical research for the U.S. shows that technology shocks are highly correlated over time, with $\rho \approx 0.98$ and $\sigma_\varepsilon \approx 0.0072$ on a quarterly basis (see King and Rebelo, 1999).[8]

The large autocorrelation of the stochastic component of technology shocks makes it difficult to statistically distinguish between transitory shocks and persistent shocks, which lead to a permanent change in the *level* of TFP. The long-run implications of these two possibilities are very different, because with persistent shocks the income effect of a shock is much larger, whereas the intertemporal substitution effect is smaller. It is therefore important to be able to

[7] There are however substantial problems with the correct measurement of inputs, e.g. with respect to quality adjustments and spill-overs from R&D. See Carlaw and Lipsey (2003) for a discussion. Any systematic measurement error in the Solow residual will also affect the TFP measure.

[8] On an annual basis, the standard deviation of output shocks is 0.028. The implausibly large size of these shocks, which are difficult to associate with observable events, is one of the main criticisms of RBC models. Indeed, the high correlation of shocks triggers intertemporal substitution effects and is one of the main driving mechanisms of RBC models.

discriminate between these possibilities, and a large body of empirical research has been created, which deals mostly with the existence of a unit root in GDP per capita, which is basically the same issue as a unit root in TFP, while GDP is less prone to measurement error.

2.2.1 A Unit Root in GDP?

The issue whether there is a unit root in GDP was put on the agenda by Nelson and Plosser (1982), who concluded that many macroeconomic variables appear to be difference stationary over a 60-year span,[9] implying that an unexpected shock has a permanent effect on the level of the variable. In particular, a 1% shock in GNP should lead to a revision of the level of future GNP of *more* than 1%, because of the lags in the data generating process.[10] In fact, the unemployment rate is the only variable for which a unit root is clearly rejected, suggesting that variables like real per capita GNP, the GNP deflator, real wages, and the nominal interest rate may all possess a unit root. Nelson and Plosser's result was reconfirmed e.g. by Campbell and Mankiw (1987), using a more elaborate testing framework. Thus, during the eighties a consensus emerged about the presence of a unit root in many economic time series.

The unit root consensus was broken in the nineties. Rappoport and Reichlin (1989) and Perron (1989) showed that if structural breaks in the regression coefficients are allowed, most variables in the Nelson-Plosser database are stationary, with the possible exception of price series and the (nominal) interest rate. Zivot and Andrews (1992) show that the Great Depression and the oil price shock of 1974 are the most likely candidates for structural breaks. The relevance of the Nelson-Plosser findings for the post-war period has also been questioned. Christiano and Eichenbaum (1990) and Rudebusch (1993) argue that if uncertainty wrt. lag length is taken into account, unit root tests based on post-war U.S. data have insufficient power to reliably distinguish between both hypotheses. The choice of sample period appeared to be vital to the conclusions anyway. Using a *longer* sample period, starting in 1870, Diebold and Senhadji (1996) were able to reject the null hypothesis of a unit root in favour of stationarity also in the absence of structural breaks. Sen (2004) and Vougas (2007) take a similar approach as Perron and Rappoport and Reichlin, using an extension of the original Nelson-Plosser data set. They both conclude that the data are trend-stationary, provided that one allows for breaks or nonlinearities in the constant term or the time trend.

[9] The sample period in the Nelson-Plosser data set ranges from 1860-1990 to 1909-1970.

[10] Nelson and Plosser (1982) estimate an equation of the form

$$z_t = \mu + \gamma t + \rho_1 z_{t-1} + \sum_{i=2}^{k} \rho_i \Delta z_{t-i+1} + u_t$$

for different data series. For real GNP per capita they obtain $\hat{\gamma} = 0.004$ and $\hat{\rho}_1 = 0.818$, which is statistically not significantly different from unity under the null hypothesis of a unit root.

One side effect of using long sample periods is that the homogeneity of the data can be violated more easily, as argued by Murray and Nelson (2000). In Murray and Nelson (2002), the authors use a Markov switching model to show that the Great Depression was characterized by more volatile shocks. Taking this conditional heteroskedasticity into account, they conclude that GDP *does* have a unit root. The argument that the trend in GDP may show breaks is used on a different level by Gaffeo et al. (2005) to argue that the lack of temporal homogeneity in GDP data over long periods makes it effectively impossible to draw reliable inferences about unit roots in the data generating process.

Figure 2.1 Annualised variance ratio of log GNP for the U.S.(source: Cochrane (1988))

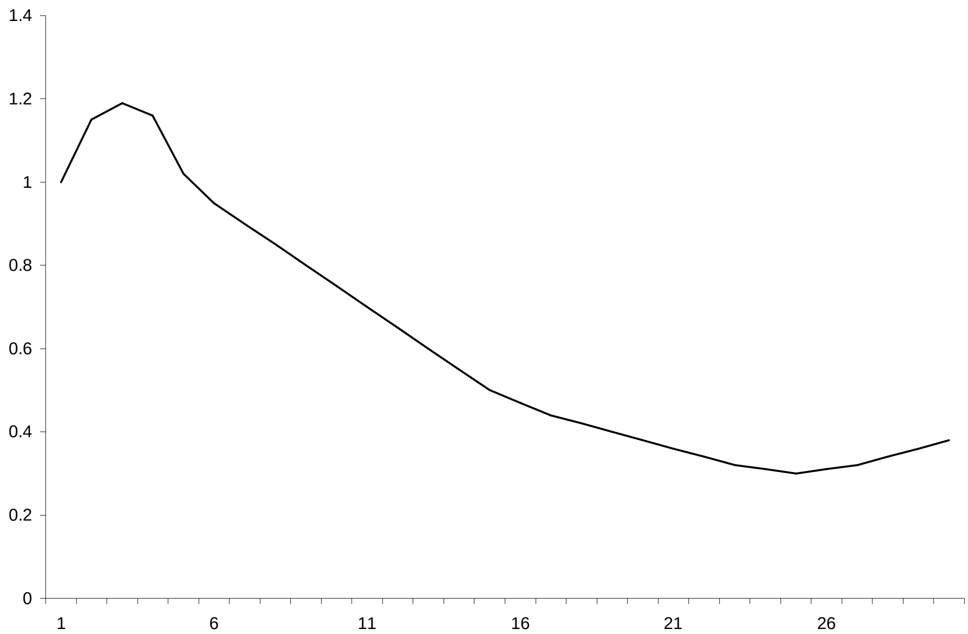

An intermediary position between trend stationarity and random walk has been proposed by Cochrane (1988). Cochrane finds that the annualised variance ratio of GNP, $\frac{1}{k}\operatorname{var}\left(\ln \frac{\mathrm{GNP}_t}{\mathrm{GNP}_{t-k}}\right)$, falls with the length of the time period considered, but stays positive (see Figure 2.1). If GNP_t is a random walk, the annualised variance should be constant, whereas if it is trend stationary, the annualised variance should fall to zero. Cochrane concludes that the non-stationary component in GNP is fairly small relative to the annual variation in GNP. In fact, the observed gradual decline in the annualised variance, at a rate smaller than $1/k$, suggests that GNP is fractionally integrated with parameter $\frac{1}{2} < d < 1$, see Diebold and Lindner (1996). This indicates that GNP may indeed be non-stationary, but with a variance that grows less rapidly than that of a random walk. This result is corroborated by Diebold and Rudebusch (1989), who obtain an estimate of $d \approx 0.7$ for per capita GNP.

Figure 2.2 **Annualised variance ratio of log GDP for the Netherlands**

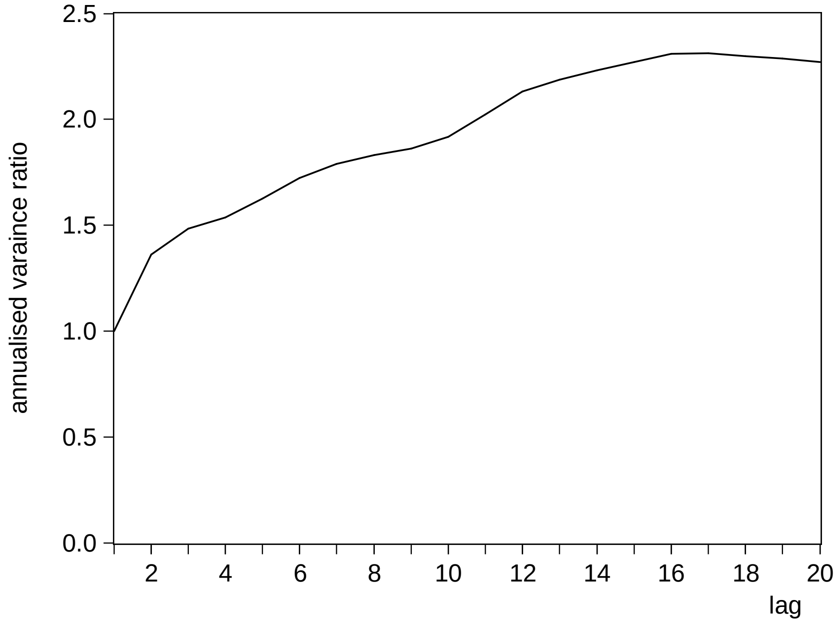

The variance profile in Figure 2.1 is specific to the United States. Cogley (1990) applied the method of Cochrane (1988) to a number of OECD countries. It appears that the mean reversion found for the United States is not present to the same degree in other countries. Several countries (e.g. Italy, France) display increasing annualised variances. Figure 2.2 shows that for the Netherlands too the annualised variance ratio increases for periods up to twenty years. Indeed, the variance ratio grows more strongly than that of any country included in Table 2 of Cogley, with a 15-20 year ratio of about 2.3.

The annualised variance profiles can also be used to assess the degree of *mean reversion* in GNP growth (i.e. the degree to which positive shocks are followed by 'compensating' negative shocks). A variance ratio that falls with the length of the time interval signals mean reversion. Using the results of Diebold and Rudebusch (1996), several cases may be distinguished:

1. A trend-stationary process with fractional integration parameter $d \leq \frac{1}{2}$ is strongly mean-reverting in growth rates, and shows an annualised variance ratio that declines asymptotically at rate k^{-1}, where k denotes the lag length.
2. An integrated process with $\frac{1}{2} < d < 1$ is weakly mean-reverting in the long run with an annualised variance ratio that declines asymptotically at rate $k^{2(d-1)}$
3. Nonstationary processes with fractional integration parameter $d \geq 1$ have asymptotically nondecreasing variance ratio's. Any tendency towards mean reversion in variance ratios is confined to the short- to medium term.

The only country for which asymptotic mean reversion is statistically significant appears to be the United States. Results for Canada and the U.K. also suggest mean reversion. Results for continental Europe do not support mean reversion. GDP growth in the Netherlands in particular shows no sign of mean reversion for periods of up to thirty years.[11]

The estimates of GNP variance growth for different countries do not take into account possible spill-overs between countries. As the U.S. was the technology leader over the sample period, the apparent nonstationarity of GNP for the other countries might be due to catching up. In addition, to the extent that convergence between countries occurs (Barro, 1996; Temple, 1999), GNP processes must be cointegrated, which precludes a fundamental (i.e., asymptotic) difference in variance profiles.[12]

The debate continues. Lima and De Jesus Filho (2008) and Cook (2008) argue that a nonlinear alternative to the standard ADF test leads to the conclusion that U.S. GDP *is* trend stationary. The nonlinear correction serves mostly to deal with the outliers during the Great Depression, like in the paper by Murray and Nelson. However, the conclusion is different. On the other hand, Benati (2007) concludes that labour productivity growth should be regarded as time-varying, which makes labour productivity an I(1) process and virtually implies that GDP per capita is also nonstationary (provided that labour participation rates are stationary).

Assessment

The currently dominant view in the unit root discussion is that quantity series may generally be characterized as trend stationary, *provided that* breaks or nonlinearities in the level or the trend parameter are included. The problem with this point of view is that the possibility of future breaks in the parameters of the data-generating process (DGP) is left open. In a sense, conditioning the stationarity of the data on exogenous structural breaks begs the unit root question. If the breaks are *not* one-time events, they must be part of the DGP.[13] As such, the *unconditional* development of e.g. GNP may be nonstationary.

In addition, there is evidence that the stochastic process for GNP in Anglo-Saxon countries is best described as an intermediary case between stationarity and non-stationarity, in the sense that forecast uncertainty increases with the horizon, but not as fast as in case of a random walk.

[11] In Appendix C, I present two productivity equations for the Netherlands, a trend-stationary autoregressive equation in log productivity, and a state space formulation with variable productivity growth, which has a unit root. It appears that the trend stationary model implies *nonstationary* forecast intervals, once parameter uncertainty is taken into account. Furthermore, both models are statistically equivalent, in the sense that a non-nested model specification test remains inconclusive. The main difference between both specifications is in the implied expected future productivity growth.

[12] See Appendix C for an example.

[13] In the papers by Barro (2006); Barro and Ursúa (2008); Barro et al. (2009), the breaks are treated as “rare disasters,” see Section 2.3.7.

However, for continental Europe the available evidence firmly points in the direction of a unit root for GNP. [14]

Another point, originally raised by Christiano and Eichenbaum (1990), is whether the stationarity issue is all that important. Since there is considerable uncertainty about the correct DGP anyway, for most series it would seem restrictive to a priori impose $\rho = 1$ or $\rho < 1$.[15] What matters from a practical point of view is mostly the degree of mean reversion in GNP growth rates over various horizons. As strong mean reversion does not find support in the data, it is probably best not to impose this.

2.2.2 Spill-overs from demographic risk to productivity risk

Population ageing may affect productivity through various channels. Productivity is age-dependent, so that the age composition of the labour force affects aggregate productivity. In addition, the ability to handle product innovations or process innovations may deteriorate with age. I discuss both channels in turn.

Direct measurement of the relation between age and productivity is difficult. Implicit contracts and tenure effects obscure the effect of productivity on wages (see the survey of De Hek and Van Vuuren (2008)). A few studies exists that attempt to directly measure the effect of age on productivity. OECD (1998) surveys psychological literature to find that cognitive skills deteriorate only to a small extent between the ages of 40 and 65. Kotlikoff and Gokhale (1992) use a panel data set on age-earnings profiles within a single large U.S. firm. They find that productivity does indeed fall with age for elder workers for all job types considered. They tentatively conclude that productivity at age 65 is only one-third of peak productivity (at age 45). However, Hellerstein et al. (1999) , in a study that uses data from different plants, do not find that productivity of elder workers is lower than that of the reference group. Dostie (2006) accounts for unobserved heterogeneity between workers and plants and also concludes that productivity of elder workers does not fall.[16] Prskawetz and Lindh (2006) reach a different conclusion. They present a survey of piece-rate studies from which they conclude that individual productivity has an inverted U-shaped profile. They attribute the productivity decline at later age to a reduction in cognitive abilities. A limitation of their survey is that it contains mainly studies

[14] For the Netherlands, the augmented Dickey-Fuller statistic for postwar log GDP per capita is -1.66, and for the growth rate of GDP per capita it is -5.26. This points to a unit root in the level of GDP. For log labour productivity per hour the D-F statistic is -1.07, and for productivity growth it is -2.8. It follows that a random walk in productivity *growth* cannot be rejected on statistical grounds at the 5% level.

[15] Stock (1991) computes asymptotic confidence intervals for the largest autoregressive root of the series in the original Nelson-Plosser data set. He finds that the unit root $(\rho = 1)$ is included in all 90% confidence intervals, except for unemployment, which is below one, and the interest rate, where it is above one.

[16] However, elder workers with at least a college degree are paid more than their productivity.

on the performance of artists and scientists, rather special groups.

A few papers look specifically at the effect of mortality risk on economic growth. There is substantial evidence of a positive relation between health status and productivity. Weil (2007) estimates that differences in health status between developing countries and developed countries account for a productivity difference of about 10%. However, this relation need not hold in general equilibrium. Young (2005) considers the effect of high mortality risk among young Africans, due to AIDS, on economic growth and concludes that the general equilibrium effects of this risk, via a higher capital intensity, dominate the outcome. That is, a higher mortality risk among young adults boosts aggregate productivity. Acemoglu and Johnson (2007) also finds a negative relation between longevity and GDP per working- age population, with an elasticity of around -2. Since most of the improvement in life expectancy in developed countries has occurred in post-retirement age groups, it is unclear whether these results have any bearing on the effects of demographic risk on productivity risk.

Boersch-Supan (2003) has estimated the expected effects of ageing on German productivity growth, assuming that the productivity profiles found by Kotlikoff and Wise (1989) may be applied to Germany. These profiles deviate from the observed wage profiles, as older workers obtain a wage in excess of their marginal product. He finds that the expected drop in productivity will probably not exceed 3% over the next two decades, i.e. at most 0.15% per year, or 0.1% per percent increase in the dependency ratio. This suggests that the effects of demographic shifts on productivity growth is minor.

A different question is whether age composition has an effect on productivity *growth*. The existence of a relation between age composition and TFP growth is plausible on theoretical grounds, as laid out e.g. in Greenwood and Yorukoglu (1997) and Violante (2002). It also finds support in the observed productivity decline of elder artists and scientists, surveyed in Prskawetz and Lindh (2006). Lindh and Malmberg (1999) and Nahuis et al. (2000) find that the growth rate of GDP per worker depends positively on the share of workers in the age group 50-64 and negatively on the share of the age group 65+. Lindh and Malmberg interpret this finding in terms of age-specific human capital, which may be largest for workers over 50. Nahuis et al. interpret the finding in terms of a resistance to the adoption of new ideas, which increases with the age of the worker. A problem with both studies is that the age composition is based on the population, rather than the workforce. Feyrer (2007) and Werding (2008) repeat the analysis using workforce composition data. They find that the age group 40-49 contributes most to aggregate labour productivity.

The effect of age composition on TFP growth cannot be inferred from the studies cited above. These studies mainly recover the macro effects of the age-productivity profile that forms the basis of Mincer (1974). The productivity index used is measured in terms of an unweighted aggregate of workers, instead of a Divisia index that uses the cost share of age groups in the total wage bill as weights. As a consequence of this inconsistent aggregation, these results mainly show composition effects. Indeed, from the estimates in Nahuis et al. the effect of the dependency ratio on log labour productivity is about -0.09, in good agreement with the calculations of Boersch-Supan (2003). The quantitative significance of an age composition effect for productivity growth at the macro level remains to be discovered.

A different transmission channel is provided by scale effects in the production of new ideas. Lee (1988) and Kremer (1993) argue that population growth leads to technical progress. Combining this scale effect with a negative feedback from productivity to fertility, Jones (2001) shows that this mechanism can be used to explain both the take-off of productivity growth during the industrial revolution *and* the subsequent demographic shift. Jones uses an elasticity of 0.75 for the effect of population size on productivity growth, which implies a semi-elasticity of $0.75\dot{A}/A \approx 0.01$. In the long run, this is not negligible.

Assessment

There is empirical evidence of a relation between the age structure of the workforce and aggregate labour productivity. This effect may possibly be interpreted as a composition effect of age-dependent labour productivity, even though the micro-econometric evidence for this interpretation is contradictory. The effect of the dependency ratio on *unweighted* log labour productivity is about -0.1. There is no firm evidence of a spill-over from age composition on TFP. There is some weak evidence of a spill-over from population size on productivity, which implies that fertility risk could have a bearing on long-run productivity risk.

2.3 Financial Asset Returns

The basic trade-off in asset portfolio choice is between risk and return. In the classic papers of Merton (1969) and Samuelson (1969), the distribution of returns is time-invariant. With constant relative risk aversion, the optimal portfolio composition is then also time-invariant and does not depend on the time horizon of the investor. The time invariance of asset return risk is empirically untenable, however. Already in the sixties, Mandelbrot (1963) showed that asset return risk is correlated over time. Research in the past three decades has reconfirmed Mandelbrot's conclusions and as a result, we have moved a considerable distance from the classic random walk hypothesis.

Campbell (2003) and Lettau and Ludvigson (2009) provide an overview of some stylised facts on asset market returns. For the purposes of this paper, the following observations seem relevant

1. For the U.S., the average (geometric) real return on stocks over the postwar period was 8%, while the real return on (risk-free) Treasury bills was 0.9%. This compares with a return on stock in the Netherlands of 14% and on short-term government bonds of 3.4%. The EU average returns are about 9% for stock and 2% for government bonds. Hence the post-war excess return of equity over bonds is about 7%, both in the US and the EU.[17]
2. Stock return has an annual standard deviation of 15% in US data, whereas the return on T-bills has a 1.7% standard deviation
3. The US excess return of equity over bonds is in substantial measure predictable over longer horizons.
4. The Sharpe ratio (the "market price of risk") for the U.S. stock market is countercyclical and highly volatile.
5. The volatility in the stock market is persistent, with a mean lag of about a year.

From this list, we see that the assumption of the Merton-Samuelson model are untenable. There is substantial variation in the volatility of returns over time. In addition, the market price of risk (the Sharpe ratio) also fluctuates over the business cycle. These stylised facts are discussed below in terms of a number of related issues: the size of the equity risk premium, the predictability of equity returns, volatility, the price of risk, and the term structure of asset returns.

2.3.1 The Equity Premium

Stylised fact 1 in Section 2.3 states that, historically, stock has earned an excess return over risk-free assets of about 7%. Mehra and Prescott (1985) were the first to point out that this observation is not compatible with standard consumption-based asset pricing in complete markets,[18] in view of the relatively smooth development of consumption over time.

In a consumption-based asset pricing model, the equity premium depends on the covariance of the marginal utility of consumption with the return of the asset. It is convenient to express the equity premium as a fraction of the return risk, which gives the *price of risk*, χ, as

$$\chi_t \equiv \frac{\mathrm{E}[R_{t+1}] - R^f_{t+1}}{\sigma_t(R_{t+1})} = -R^f_{t+1}\,\sigma_t(m_{t+1})\,\rho_t(m_{t+1}, R_{t+1}) \tag{2.3}$$

where R^f represents the risk-free return, m is the stochastic discount rate of the investor, and $\rho(m,R)$ denotes the correlation between the stochastic discount rate and the return to the asset. For a given price of risk, the equity premium varies in proportion with the return risk on the asset. The price of risk depends on general consumption risk exposure of the investor ($\sigma(m)$),

[17] Campbell's data for Europe cover the period 1970-1999. Taking the stock market crashes of 2002 and 2008 into account lowers the excess return by about 2% points. For the U.S., the crash causes the excess return over the period 1891-2008 to fall by about 0.5%.

[18] Cochrane (2008) notes that the problem was already clearly identified by Shiller (1982)

and on the correlation between the marginal utility of consumption (or wealth) and the asset return. The correlation between consumption growth and the return on equity varies between 0.25 and slightly negative for the G-7 countries (Sarkar and Zhang, 2004, quarterly data). For the risk-free return, $R_f \approx 1$ is a reasonable approximation. With a standard deviation of consumption growth of about 2%, $\chi \approx 0.25 \times 0.02\gamma = 0.005\gamma$, where γ is the rate of relative risk aversion. Empirically, the Sharpe ratio, which proxies the price of risk, is about 0.4 on average for equity, so that we need $\gamma \approx 80$. However, from micro experiments we obtain values of the relative risk tolerance $(1/\gamma)$ of around 0.25 (see e.g. Barsky et al., 1997, Table IIa). Hence the equity premium puzzle is a *quantitative* puzzle about explaining the size of the average observed excess return on equity in relation to the return risk incurred. In addition, a high coefficient of risk aversion implies a high risk-free rate (the risk-free rate puzzle, Weil (1989)), so that setting γ at a high value only avoids the Scylla at the cost of falling victim to the Charybdis.

The equity premium puzzle has generated a huge literature. Recent surveys are given by Salomons (2008), Donaldson and Mehra (2008), and Mehra and Prescott (2008). From the literature the following types of explanation can be distilled:

- preference related: non-expected utility, habit formation, loss aversion
- incomplete markets: liquidity constraints, uninsurable micro-risks, transactions costs, limited participation
- finite sample bias: parameter uncertainty, catastrophic events, survivor bias

For practical purposes the important question is what can be said about future (expected) excess returns on equity. Do excess returns vary systematically with economic conditions? Are excess returns to some extent predictable? These issues will be addressed in the sections below.

Here I discuss the issue of finite sample bias. Were the observed high excess returns to equity of the past an indication of high *expected* excess returns over that same period? Did investors ex ante have any reason to expect these high returns? If the observed high excess returns in the U.S. and other long-lasting stock markets are a consequence of survivor bias, as argued by Jorion and Goetzmann (1999), the expected excess return, i.e. the equity premium, was probably much lower. However, Dimson et al. (2008), using a database of seventeen countries over a 106-year interval, estimate the historical world equity premium at 4.1%, a bit below the U.S. figure of 4.5%.[19]. After adjustment for survivorship bias, this figure drops to 4.0%, a negligible correction.

A related point is that conditions in the past century were not representative for current conditions so that *future* returns may turn out to be much lower. Fama and French (2002) argue

[19] Geometric means (i.e. annualised returns). The higher figure of stylised fact 1 in Section 2.3 above refers to the post-war period.

that the high average return over the period 1950-2000 is caused by a decline in discount rates.[20] Using dividend/price ratios and dividend growth, Fama and French estimate an equity premium of between 2.5% to 4.3% over that period, depending on whether one uses dividend or earnings forecasts. Dimson et al. (2008) use a similar decomposition into fundamentals to argue that the current equity premium should be in the range $3\% - 3.5\%$. The estimate of Dimson et al. is compounded as given in Table 2.1 below. The expected annualised return on equity is between 4.3% and 4.8% and with a real risk-free rate of 1% this yields an equity premium of between 3.3% and 3.8%. The tricky part in these estimates is the expected dividend yield. The stock

Table 2.1 Decomposition of the estimated equity premium (Dimson et al. (2008))

	$\dot{d}/d$ [A]	$\Delta p/d$ [B]	p/d [C]	$\dot{e}/e$ [D]	i [E]	equity premium
scenario						
high	0.01	0	0.038	0	0.008	0.038
low	0.0	0	0.032	0	0.008	0.032

[A] real dividend growth rate.
[B] change in dividend yield.
[C] dividend yield.
[D] change in real exchange rate.
[E] U.S. real interest rate.

market crashes of the 21st century may very well result in higher dividend yields in the near future. More generally, the use of fundamentals to estimate equity premia neglects predictable variations in equity risk and the price of risk. I address these questions below.

Assessment

A "reasonable" lower bound for the equity premium, based on fundamentals, is in the order of 3%. However, it is plausible that the current equity premium is higher than that, given that equity return risk has risen. An upper bound may be set at the historically observed excess return of 7%.

2.3.2 Asset Return Predictability

Estimating the current equity premium is a special case of the more general problem of efficient prediction of asset prices from current information. Under the random walk hypothesis of asset markets (see Fama, 1965), such an effort would be futile. However, research in the seventies showed that excess asset returns are to some extent predictable, see e.g. Bodie (1976) for the real return to equity and Fama and Schwert (1977) for the real return on short-term bonds. Predictability has also been established wrt. medium-term excess returns of equity, see e.g. Fama

[20] This argument finds support in the results of Lettau and Van Nieuwerburgh (2008) (Section 2.3.2).

and French (1988). It appears from the survey in Lettau and Ludvigson (2009) that a whole range of variables is able to predict excess stock returns, including price-dividend ratios, term spreads, and proxies for the consumption-wealth ratio. Generally, the predictability of stock returns appears to improve with the horizon. A widely used variable is the price-dividend ratio (see Campbell and Viceira, 1999; Campbell and Shiller, 2001), which is able to explain 10% of the variation in excess returns in the U.S. stock market over a 10-year horizon.[21]

The dividend-smoothing model implies that price-dividend ratios contain information about discount rates of investors but not about future dividends, see Cochrane (1994) and Appendix D. If discount rates are stationary, so are price-dividend ratios, and any deviation from the unconditional mean must signal a future corrective price adjustment. A fall in the price-dividend ratio of 1% therefore generates a predicted excess return of 1%. The price-dividend ratio does not forecast well over short horizons, however. According to the survey in Table 2 of Lettau and Ludvigson (2009) it is outperformed by the consumption-wealth ratio proposed by Lettau and Ludvigson (2001). Moreover, Lettau et al. (2008) and Lettau and Van Nieuwerburgh (2008) show that the ability of dividend yields to forecast returns has declined after 1995. There appears to be a structural break in dividend yields from 1995 on, with substantially lower dividend yields after this date. Lettau and Van Nieuwerburgh show that conditional on this break, the predictability of returns from dividend yield remains intact. In line with the dividend-smoothing hypothesis, they argue that the break in dividend yields signals a break in discount rates, with substantially lower risk premia being applied after 1994.

Business cycle information also helps predicting stock returns. Fama and French (1989) show that expected excess returns for bonds and equity are countercyclical (with as predictors dividend yield, the term spread, and the default spread). Risk premia rise at the start of a recession and fall again during recovery. Campbell and Diebold (2009) show that business expectations as measured by the Livingstone survey predict excess returns. Expected excess returns are found to be countercyclical. Lettau and Ludvigson (2001) provide a link between return predictability and 'real' macroeconomic variables by showing that the consumption-wealth ratio forecasts both stock returns and dividend growth at horizons as short as one quarter.[22,23] The consumption-wealth ratio therefore contains information about stock prices that is independent of the price-dividend ratio. In addition, the consumption-wealth ratio

[21] However, this finding does not hold for continental Europe. For the Netherlands in particular, Campbell and Shiller find no relation between dividend-price ratios and stock prices growth.

[22] Wealth is taken to include human wealth, i.e. the present value of future wage-related income. Lettau and Ludvigson construct a proxy for human wealth based on current labour income, *assuming* that the return on human wealth is a stationary process.

[23] The estimates suggest that a 1% increase in the consumption-wealth ratio boosts predicted excess returns by 5% over a 1-year horizon, c.q. 7% over a two-year interval.

outperforms the price-dividend ratio at periods of up to four years, and the latter variable only contributes at horizons of five to six years. Moreover, the relation between stock returns and the consumption-wealth ratio appears to be more stable than that between returns and dividend yields. Lettau and Ludvigson (2005) show that the consumption-wealth ratio does not display a structural break after 1994, and maintains a stable relation to stock returns after 1994.

The predictability of the equity premium also has implications for investment in productive assets. If there is a term structure in the equity premium, investment in long-lived fixed assets should depend more on future equity premia than the current expected excess returns. Lettau and Ludvigson (2002) find some confirmation for this theoretical relation in that they show that the consumption-wealth ratio also predicts long-term investment. Hence, a hike in the equity risk premium will curtail investment in the short run, because of a higher cost-of-capital. However, in the long run, the equity premium will mean revert, and investment will recover.

The results from studies of return predictability suggest that predictability improves with the length of the horizon over which the return is measured. These multi-period returns are computed by summing single-period returns, so that adjacent data points contain overlapping data. Valkanov (2003) shows that estimators using overlapping data are not consistent and their t-statistics do not have standard distributions. In addition, the R^2 statistic cannot be used as an indicator of goodness-of-fit, so that the apparent improvement in predictability with the length of the horizon may be spurious. Valkanov proposes a correction factor of $1/\sqrt{T}$ for the t-statistics, in combination with an adjusted asymptotic distribution. With these corrections, the case for long horizon predictability appears to be considerably weaker.

Another approach to long-run predictability is to estimate a VAR model of one-period returns and derive the multi-period returns from the VAR model. Campbell et al. (2003) and Campbell and Viceira (2005) investigate excess return predictability of a number of financial assets in the context of such a model.[24] The estimated model implies declining standard deviations of the real return to equity and 5-year bonds, as a result of mean reversion in the returns. In contrast, the standard deviation of the real return on short-term bonds quadruples, due to inflation risk, making short-term bonds just as risky as equity in the long run.

Campbell and Viceira do not present standard errors of multiperiod return forecast errors, so that it is difficult to judge the reliability of the model for long-run forecasts. Potentially, parameter uncertainty could render their long-term results statistically insignificant.[25] In addition, Goetzmann and Jorion (1993) and Nelson and Kim (1993) argue that small sample bias

[24] The structure of the Campbell and Viceira (2005) model is discussed in Appendix E.

[25] Moreover, the variance profile estimates in Figure 1 of Campbell and Viceira (2005) do not follow from the published VAR coefficients and error covariances in their Table 2. On the basis of these parameter values, the variance profiles are almost completely flat.

affects estimated coefficients. Dividend yields are highly autocorrelated, so that explanatory variables are correlated with lagged values of the dependent variables, which violates the independence assumption of OLS. They use bootstrapping methods to approximate the finite-sample distributions of the test statistics. Based on their results, the statistical significance of the results in Fama and French (1988) is much reduced.

Hoevenaars et al. (2008) use an extended version of the Campbell and Viceira model, adding more assets, the credit spread, and inflation-linked liabilities. The additional assets are assumed not to generate dynamic feedback on the primary assets, which separates the model into a "core" VAR model and a supplementary VAR. Their results for the core VAR confirm those of Campbell and Viceira. Mean reversion in stock returns is attributed both to dividend yield behaviour and the credit spread.

The finding that excess returns are in substantial measure predictable has been challenged by Goyal and Welch (2003, 2008), who use out-of-sample forecasting to test the predictive ability of dividend yields and other factors. They find unstable parameter estimates and poor forecasting performance, compared to forecasts based on the sample mean. Cochrane (2007, 2008) and Lettau and Ludvigson (2009) argue that the results in Goyal and Welch primarily show the effect of finite-sample bias on forecasting errors, and do not deny the existence of a structural relation between dividend yield and future excess returns. However, the results of Goyal and Welch (2008) do cast doubt on the possibility to successfully use dividend yield information in a portfolio investment strategy, as advocated in Campbell and Viceira (2002).

Assessment

Asset excess returns show predictable variation, i.e. the equity premium is not constant. Financial variables like price-dividend ratios and term spreads have some forecasting ability over medium-term horizons (5-6 years). A fall in the price-dividend ratio of 1% raises predicted excess returns by about 1% over this horizon. Another relevant indicator appears to be the consumption-wealth ratio, which predicts about 5% of excess returns per percent increase in the consumption-wealth ratio over horizons of up to six years.

The case for *out-of-sample* asset predictability is fairly weak, due to parameter uncertainty in regression equations and statistical problems with overlapping data in long-term returns. It is therefore doubtful whether a successful portfolio investment strategy can be devised that is based on the predictability of asset returns.

2.3.3 Volatility and Excess Returns

In the Merton-Samuelson model (Merton, 1969; Samuelson, 1969) investors assume that the distribution of asset returns is time invariant. However, Mandelbrot (1963) already noted that volatility is clustered over time. Changes in market risk may induce changes in equity premia. Provided that the price of risk in (2.3) remains constant, a doubling of return volatility results in a doubling of risk premia. This relation is not evident from the data. Black (1976b) showed that changes in volatility are *negatively* correlated with excess returns, so that returns below expectations boost volatility and vice versa. The variation in return volatility raises two questions: what causes volatility changes, and how do volatility changes affect the price of risk?

Merton (1980) was the first to systematically investigate the empirical relation between volatility and risk premium. He used observed volatility as a measure of risk, which obscures most of the relation between risk premium and volatility, due to the large variance of stock returns, and the corresponding low signal-noise ratio. Christie (1982) investigates the relation between stock return variances and a number of economic variables. He finds a strong effect of leverage and interest rates on volatility. This finding supports Black's "leverage effect" hypothesis: a fall in stock prices boosts the leverage of the firm, which increases the volatility of the firm's market value.

Schwert (1989) also investigates the relation of stock market volatility to other economic variables, such as a recession period, financial leverage, predicted industrial production volatility, and predicted inflation volatility. The recession dummy is the only variable which is consistently significant over all subperiods distinguished. The predicted volatility variables are only intermittently significant. The statistical significance of financial leverage disappears when it is tested together with other variables. In any case, it can explain only a small fraction of stock market volatility. Conversely, the evidence for the hypothesis that stock market volatility affects macroeconomic volatility is somewhat stronger. Schwert concludes that there exists a "volatility puzzle."

French et al. (1987) point out that the observed correlation between volatility and return consists of two components, a positive relation between the expected excess return and *expected* volatility, and a negative relation between the actual excess return and *realised* volatility. An unexpected, persistent, increase in volatility will initially lower excess returns, due to the effect of a higher risk premium on asset prices, and afterwards, on average, lead to an increase in excess returns. French et al. (1987) and Bollerslev et al. (1988) model this relation in terms of a GARCH-in-mean process, where the conditional expected return is a function of volatility. They find a positive relation between predicted volatility and excess returns. However, this effect is dominated quantitatively by the far larger effect of unexpected changes in volatility on holding period yields, as noted by Black (1976b). Generally, the relative size of these effects depends on the persistence in volatility. French et al. obtain a mean lag in volatility of about 12 months and Bollerslev et al. of about 2 quarters. This implies that an unexpected shock in volatility may push up risk premia for more than a year, so that current stock prices have to fall by a substantial amount.

Nelson (1991) and Bollerslev et al. (1994) use an exponential GARCH (EGARCH) model in which variances depend both on the size and sign of the residuals, to capture the asymmetric effects of shocks on volatility through the "leverage effect" discussed above. They find that the volatility process is a composite of two AR(1) processes, one with a half-life of a few days, and the other with a half-life ranging from about one year (Bollerslev et al.) to seven years (Nelson). In addition, the asymmetry effect of the EGARCH model is strongly significant, negative return shocks create considerably more volatility than positive shocks.

Bekaert and Wu (2000) investigate the asymmetric volatility of stock returns in terms of two competing hypotheses: the leverage effect and the volatility feedback effect. They test their model both at the firm level and at the level of the market. At the firm level, they document an extra volatility feedback effect via the covariance asymmetry between stock returns and market return in response to positive and negative shocks. They find that asymmetric volatility is mostly driven by volatility feedback and time-varying risk premia. Leverage effects play only a minor role.

Parametric specifications of volatility in terms of ARCH models condition volatility in terms of a small number of observable variables. Instead, in *stochastic volatility models* the variance is a latent stochastic process (see Andersen et al., 2006, section 4). Estimated volatility then depends on the whole history of the stochastic process. Brandt and Kang (2004) model both conditional returns and conditional volatility as latent state variables. The specification includes both a volatility-in-mean and a mean-in-volatility effect, so that conditional expected volatility is affected both by persistence, and by deviations of conditional expected returns from their long-term mean. Their response functions indicate that a one standard deviation (3%) shock in volatility is fairly persistent, and takes more than twelve months to revert to normal. The positive volatility-in-mean effect generates a temporary increase in expected excess returns, with a peak of 0.5% after about 8 months. A one standard deviation innovation in mean returns (0.66%) causes a fall in expected volatility of about 1.5%, which causes expected excess returns to first undershoot the long-run level before reverting to normal after two years. The Sharpe ratio falls in response to an innovation in volatility and increases in response to an innovation in expected returns.

A new development in volatility modelling is *realized volatility* estimation (see Andersen et al., 2006, section 5). In this approach, high frequency data are used to construct a non-parametric estimate of the development of volatility over time. The estimated volatility series can be used as a benchmark to evaluate GARCH volatility forecasts, or as direct input in a standard ARFIMA forecasting scheme. Andersen et al. (2001, 2003) find that the log of daily volatility is a stationary Gaussian process with long memory, with fractional integration parameter $d \approx 0.4$. Volatility is also asymmetric, stock market downturns are associated with larger volatility than upturns. However, this relation is stronger at the market level than at the individual stock level, so that a leverage effect does not find direct support in the data.

Assessment

Asset return volatility shows substantial correlation over time, with mean lags of about one year. Volatility is strongly countercyclical, the main economic factor that boosts volatility is the start of a recession. Another factor that influences volatility is financial leverage, but empirical evidence on the size of this effect is conflicting. Changes in predicted volatility boost risk premia, but unexpected changes in volatility cause a fall in realised excess returns, so that volatility and excess returns are negatively correlated in the short run. Negative excess return shocks boost volatility more than positive shocks.

2.3.4 The Price of Risk

The trade-off between expected volatility and expected excess returns can be summarised in the market price of risk (see (2.3)). The realized Sharpe ratio is not a good measure of the market price of risk, as both the numerator and the denominator are measured with error. Lettau and Ludvigson (2009) estimate the price of risk by using *predicted* excess returns and *predicted* volatility to construct a conditional Sharpe ratio for the period 1953.Q1-2001.Q1, based on the CRSP VW index. Their estimates point to a rather volatile price of risk, with values between -0.5 and 1.75,[26] that is strongly countercyclical. From their figure 2, the price of risk rises substantially during most recessions, and on average falls again during expansions, although the pattern during expansions is less clear than during recessions.[27]

Estimates for the price of risk can also be inferred from stochastic volatility models. In the model of Brandt and Kang (2004), the conditional Sharpe ratio is countercyclical. Volatility increases near the start of a recession, and raises the Sharpe ratio. Near the trough of the cycle, positive return innovations boost the Sharpe ratio again, and this effect is reinforced by the mean-in-volatility feedback effect, which dampens volatility.

Most consumption-based asset pricing models fail to capture the cyclical properties of the price of risk, as the conditional distribution of the stochastic discount rate ($\sigma(m)$ in (2.3)) in these models does not vary much. Campbell and Cochrane (1999) succeed in introducing non-trivial variation in the price of risk within a consumption-based model by introducing habit

[26] Negative estimates arise because the predicted excess return is occasionally negative.

[27] Their Figure 1 shows that the results are partly driven by a procyclical conditional volatility, which is at variance with some of the studies cited in Section 2.3.3.

formation. With habit formation, consumers are more risk-averse during a downturn, because the habit term in their utility function is slow to adjust, which boosts their marginal utility of consumption. The stochastic discount factor is now

$$m_{t+1} = \mathrm{e}^{-\delta} \left(\frac{S_{t+1}}{S_t} \frac{C_{t+1}}{C_t} \right)^{-\gamma} \tag{2.4}$$

where S denotes the *surplus consumption ratio*, $S_t = (C_t - X_t)/C_t$ and X is the habit. The dynamics of X are "proudly reverse engineered" (Cochrane, 2001) to generate precautionary saving in bad states and dampen the intertemporal substitution effect that would otherwise drive up the risk-free rate.[28] This is done by making the habit adjust faster in bad states than in good states. So, in a bad state S_{t+1}/S_t is expected to be lower than in a good state and the stochastic discount factor is higher. Still, the risk premium is high, because $S^{-\gamma}$ is much more volatile than $C^{-\gamma}$ and excess stock returns correlate negatively with $S^{-\gamma}$. With reference to (2.3), the value added of the habit formation assumption is that, both $\sigma(m_{t+1})$ and $\rho(m_{t+1}, R_{t+1})$ are boosted, while both the variance of consumption and the correlation between consumption and stock returns remains low.

The Campbell/Cochrane model does a good job of explaining both the equity-premium puzzle and the risk-free rate puzzle, at the cost of a somewhat contrived specification of habit formation. The model generates, by calibration, the average Sharpe ratio (0.4). It is also able to produce a countercyclical conditional Sharpe ratio, and, consequently, rising equity premia at the start of a recession. However, the model is not able to match the volatility of the conditional Sharpe ratio estimated in Lettau and Ludvigson (2009). This suggests that the model may need other risk factors with cyclical properties.

Bansal and Yaron (2004) specify a model with two risk factors, consumption growth and dividend yields. Consumption growth has time-varying volatility and both risk factors have persistence in growth rates. Utility is of the Epstein and Zin (1989) variety, with an IES of 1.5, and a relative risk aversion parameter of 10. Because of the long-run impact of news on expected volatility and growth, the asset price implications of small shocks are large. The model is able to reproduce both the equity premium and the risk-free rate, and to capture the volatility feedback effect on asset prices.

The Campbell/Cochrane model uses only one risk factor, consumption growth. Extending the structural model to include more risk factors complicates the analysis, as it implies an expansion of the state space. In the finance literature, these complications are evaded by using a reduced-form approach, in which the stochastic discount factor is specified directly in terms of several underlying risk factors. A common specification is the affine yield model of Duffie and

[28] A bad state is defined as a state with low S.

Kan (1996) (see Appendix F for a description of the model in discrete time). Affine models have some limitations, however. They have difficulty explaining the cross-sectional distribution of interest rates, and they are of necessity linear, even though the price of risk often behaves nonlinear. Finally, the factors in an affine model often have little economic content.

The cyclical properties of the price of risk are also relevant for investment in fixed assets. Lettau and Ludvigson (2002) find that long-term investment is affected by changes in the risk premium, via the consumption-wealth ratio. The implication is that variations in the cost of capital are not completely captured by variations in interest rates. In a general equilibrium context, the relevant variable in (2.3) is $\sigma(m_{t+e})$, the volatility of the stochastic discount rate. If the stochastic discount rate of investors is more volatile, investors are behaving in a more risk-averse way. This induces a downward slope in the term structure of the risk premium and leads to a countercyclical shadow cost of capital. At the trough of the cycle, forward-looking investors expect that the cost of capital will go down, and they therefore invest in long-lived project in anticipation of the recovery.

Assessment

The price of risk is countercyclical, and rises sharply during a recession. In combination with countercyclical volatility, this implies strongly countercyclical *expected* excess returns on equity. The cyclical properties of the price of risk also may help to explain the behaviour of investment in fixed assets over the cycle, in particular the observation that investment leads the cycle. As the real risk-free rate is also countercyclical (King and Watson, 1996), expected equity returns are even more countercyclical.

2.3.5 The Term Structure of Asset Returns

For some assets, forward rates are available over long horizons, e.g. government bonds offer maturities of 30 years and more.[29] From the yield curves for these assets, information can be distilled about the perception of investors about future risks.[30] In particular, the expected development of the stochastic discount rate can be read from the term structure of the risk-free rate.

Empirical evidence on the term structure of the real risk-free rate is not always available. In most countries an indexed bond market does not exist, so that the real rate can only be inferred by estimating expected inflation. A risk-averse investor who only has access to nominal bonds

[29] The Dutch Central Bank publishes a 60-year yield curve, to be used by pension funds to discount their future pension obligations, see www.statistics.dnb.nl/index.cgi?lang=uk&todo=Rentes.

[30] However, asset markets thin out into the distant future, which makes the measurement of risk perceptions increasingly unreliable.

will however apply an inflation risk premium to these bonds, so that the nominal bond rate contains a risk premium that may vary with the inflationary regime (see below), which blocks any direct inferences about the real risk-free rate. For these cases, it may be better to resort to a priori reasoning.

The standard Ramsey pricing formula arises if utility is CRRA and consumption follows a random walk with drift, $\ln c_{t+1} = \ln c_t + \mu + \varepsilon_{t+1}$. The stochastic discount rate is now

$$m_{t+1} = \exp\left[-\delta - \gamma\mu - \gamma\varepsilon_{t+1}\right] \tag{2.5}$$

where δ is the rate of time preference and γ the coefficient of relative risk aversion. Let $R_{t,T} = \prod_{i=1}^{T}(1 + r_{t+i})$ be the T-period safe rate of return. It follows from the arbitrage equation ((F.1e) in Appendix F) that the yield curve is given by

$$\left(R_{t,T}\right)^{1/T} = \exp\left[\delta + \gamma\mu - \tfrac{1}{2}\gamma^2\sigma^2\right] \tag{2.6}$$

The yield curve is perfectly flat. The safe rate is higher in a fast-growing economy, and lower if future consumption is more uncertain.[31]

The assumption that consumption follows a random walk with drift may overstate the amount of uncertainty in future consumption. The literature discussed in Section 2.2.1 indicates that there is probably *some* mean reversion in GDP, so that long-term annualised variances are lower than short-term variances. Gollier (2007) shows that a declining annualised variance in consumption should result in an upward sloping yield curve. As consumption and GDP are cointegrated, long-term risk-free discount rates should be *larger* than short-term rates, because mean reversion in consumption growth implies that long-term uncertainty per unit of time in consumption is smaller than short-term uncertainty. Using Cochrane (1988), Gollier shows that a value of $\gamma = 4$ results in a short-term discount rate of 4%, which gradually rises to 6% at horizons of 20 years and longer.

On the other hand, the *growth rate* of consumption may also be uncertain. Bansal and Yaron (2004) show that there is persistence in both the rate and the volatility of consumption growth. This implies consumption variance ratios that increase for periods of up to ten years. Gollier (2008) considers a stronger case in which the distribution of the growth rate of consumption is *unknown*, so that variance ratios keep increasing.[32] Increasing variance ratios cause the annual

[31] Stern (2006) considers the parameter values $\delta = 0.1\%$, $\mu = 1.3\%$, $\gamma = 1$ and $\sigma^2 = 0$, which results in a safe rate of 1.4%. This estimate neglects consumption risk and uses an estimate of γ at the lower end of the range. Putting $\mu = 0.018$ and $\sigma = 3.56\%$ (Kocherlakota, 1996, Table 1) raises the risk-free rate to 1.9%, while setting $1/\gamma = 0.25$, (Barsky et al., 1997) raises the risk-free rate to 6.5%.

[32] Presumably, the growth rate also follows a random walk, otherwise people would gradually learn its value, anticipation of which should be reflected in the optimal discount rate (cf. Weitzman, 2007).

discount rate to fall gradually with the length of the horizon. Gollier (2008) gives an example in which the yield curve gradually falls from 3% in the near future to 1.5% in the distant future.

Ang et al. (2008) use a term structure model of nominal rates, with as state variables inflation and two latent factors, to estimate the real term structure for the US. They incorporate four different inflationary regimes in their model, which follow a Markov chain. The *unconditional* term structure of the real rate that results is fairly flat, at around 1.3% up to 5 year maturities. The regimes with relatively volatile and high real rates have a downward sloping yield curve, that starts at around 2% and falls off to 1.5% after five years.

For the UK, a market for price-indexed bonds exists since 1983. Apart from some indexation lag issues, this market allows for a direct measurement of real rates that are, presumably, risk-free (apart from liquidity risk). From (Evans, 2003, Table 1) it appears that the real yield curve in the UK was on average downward sloping over the period 1983-1995, from 5% to 4% over a 10-year horizon. Evans estimates a term structure model with regime switching, to capture changes in inflation regimes. He finds that real yield curves are upward sloping in the regimes where inflation is slowly rising or falling, and downward sloping in the regime with quickly rising inflation. The long-run real yield is 5% in all regimes, but short-term yields vary between -0.5% for the slowly falling inflation regime and 8% for the quickly rising inflation regime. In fact, the current (as of 09/03/09) UK real yield curve starts at 1.5% (2.5 years) and gradually falls to zero at 25 years, which appears to contradict one of the maintained assumptions of the model.[33]

For the US, a price-indexed bond market exists since 1997, a comparatively short time span. Furthermore, the market was not very liquid in its initial stages. As a result, there is a lack of empirical studies that address this market. The current (09/03/09) real yield curve rises from 1.1% for five-year bonds to 2.2% for 20-year bonds.[34] The difference in long-term real risk-free yields between the UK and US points to an expected appreciation of the pound sterling.

In the Euro area, the French treasury issues inflation-linked bonds since 1998.[35] Combining information about break-even inflation rates and real yields for 2015 with the nominal yield curve published by the European Central Bank,[36] results in three points on the real yield curve, 0.8% at 6 years, 2.4% at 10 years, and 1.9% at 30 years.

[33] See www.bankofengland.co.uk/statistics/yieldcurve/ As the nominal yield curve rises from 0.4% at 0.5 years to 4.5% at 20 years, this implies that inflation is expected to rise to 4.5%.

[34] See www.ustreas.gov/offices/domestic-finance/debt-management/interest-rate/. The nominal yield curve rises from 0.1% at one month to 4.2% at 30 years.

[35] Inflation-linked bonds are also being issued by Germany, Italy, and Greece.

[36] See www.aft.gouv.fr/aft_en_21/debt_management_51/products_248/oateurois_and_btaneurois_257/index and www.ecb.int/stats/money/yc/html/index.en. There is a small difference between the French break-even inflation rate and the rate for the Euro zone.

In policy analyses, it may be useful to evaluate the consumption stream itself. A claim to consumption is a relatively risky asset, since it correlates strongly with the marginal utility of consumption. This raises the price of risk (see (2.3)). Let $p_{c,t}$ be the price of a (permanent) claim to this flow. The return to this claim is $R_{c,t} = (p_{c,t+1} + c_{t+1})/p_{c,t}$. E.g., with CRRA utility pricing, it follows from $\mathrm{E}[m_{t+1}R_{c,t}] = 1$ that the value of the claim is

$$\frac{p_{c,t}}{c_t} = \mathrm{E}\left[m_{t+1}\frac{c_{t+1}}{c_t}\left(1+\frac{p_{c,t+1}}{c_{t+1}}\right)\right] \quad \Rightarrow \tag{2.7a}$$

$$\frac{p_{c,t}}{c_t} = \left\{\exp\left[\delta + (\gamma-1)\mu - \tfrac{1}{2}(\gamma-1)^2\sigma^2\right] - 1\right\}^{-1} \tag{2.7b}$$

Using the parameter values in Kocherlakota (1996) (see Footnote 31), for $\gamma = 1$, $p_c/c = 1000$, and for $\gamma = 4$, $p_c/c = 20$. Correcting for consumption growth, the discount rate is between 1.9% and 7.2%, slightly above the respective safe rates. With high consumption risk the gap widens, e.g. for $\sigma = 10\%$ and $\gamma = 4$, the safe rate is -0.7% and the consumption discount rate is 2.8%.

As CRRA asset pricing is empirically rejected, it may be better to use the habit formation model of Campbell and Cochrane (1999) to evaluate risky claims. In that case, the stochastic discount rate is state-dependent and time-varying (see (2.4)). Campbell and Cochrane find that the price of a consumption claim as defined in (2.7a) varies between $p_c/c = 6$ and $p_c/c = 24$, depending on the initial state.[37]. The implied annualised discount rates are 22% and 6.3% respectively, which points to large fluctuations in risk premia. The risk premium is substantial (3.8%) even in good states, because these states may not last (i.e. the yield curve is upward sloping). The large annualised risk premium in bad states is a consequence of an even larger short-term risk premium, that falls again over longer horizons as habits adapt (see (2.4) above). As there is mean reversion, the price-consumption ratio forecasts future expected returns.

The working-paper version (Campbell and Cochrane, 1995) contains more information about the implied term structure of discount rates. Figure 2.3 presents the elasticities of the term structure of excess stock returns with respect to the price-consumption ratio. A one percent higher price/consumption ratio results in a one percent lower expected return in ten year's time. The long run excess return is constant, but the term structure displays slow mean reversion.

Bansal and Yaron (2004) also derive a term structure of returns from their model, but they provide little quantitative detail. As in Campbell and Cochrane (1999), the term structure depends on the initial state of the economy, which in their model consists of the latent growth rate factor and consumption volatility. Table 6 in Bansal and Yaron (2004) provides partial information, which agrees with Figure 2.3.

The term structure of interest rates in the Campbell/Cochrane model does not fit the observations very well. To match the risk premia of bonds and equity on the one hand and the

[37] Assuming $\gamma = 2$, $\mu = 1.89\%$, and a risk-free rate of 2.5%, see their table 1.

Figure 2.3 Elasticities of log excess returns with respect to log price/consumption in the Campbell/Cochrane model. Source: Campbell and Cochrane (1995), Table 3.4

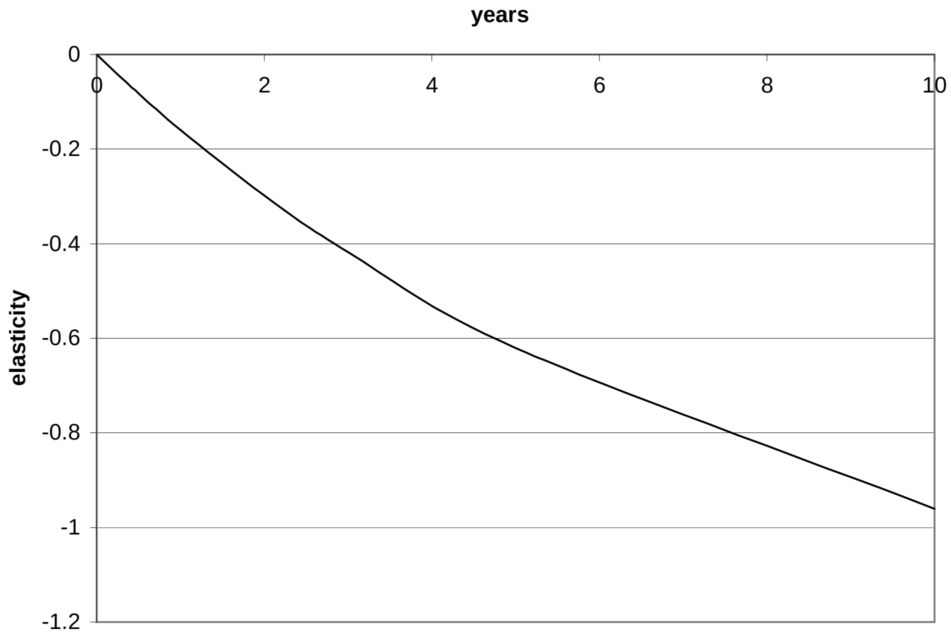

yield curve of interest rates at the other hand, it appears necessary to include more than one risk factor. Lemke and Werner (2009) use an affine risk factor model for this purpose. In their model, the state space consists of inflation, the dividend yield, and two latent factors that define the real interest rate. The innovations in these variables constitute the risk factors that drive the model. The stochastic discount factor is specified as an affine function of the risk factors, see (F.1a) in Appendix F. The model is estimated on U.S. data for inflation, ex dividend stock returns, dividend yields, and zero-coupon bond yields of different maturities, spanning the period 1983 till December 2008. By construction, the estimated equity premium is a linear function of the state variables, two of which are latent. The results show, in general, an upward-sloping term structure, with short-run risk premia over the safe rate that vary between 0.5% and 4%, and long-run premia between 2% and 4%. That is, short-term risk premia are more volatile than long-run premia. The estimated term structure in January 2008 is different. The equity premium over 2009 is 4.2%, rising to 5% per year over a five-year period, and falling again to an annualised rate of 4% in the long run.

The term structure of portfolio returns has implications for portfolio management. The Samuelson/Merton rule that a constant fraction of wealth should be invested in risky assets no longer holds if the price of risk varies with the investment horizon. Campbell and Viceira (1999, 2002) argue that the optimal portfolio of a long-term investor should contain more stock than that of a short-term investor, because annualised return risk is falling with the length of the horizon (i.e. the term structure of the risk premium is downward sloping). This general prescription is at variance with the state-dependent term structure derived by Campbell and Cochrane (1999), and it is not supported either by the empirical term structures obtained by Lemke and Werner (2009). However, the results of Goyal and Welch (2003, 2008) suggest that the reliability of the estimated term structure of equity returns may be low. In fact, the standard errors for the estimated equity premium term structure provided by Lemke and Werner (2009) suggest that the reliability of the estimated equity risk premia is highest for long horizons. For short horizons, up to about eight years, the standard error is larger than one percent.

Assessment

Empirically, the yield curve for the *real* risk-free rate appears to be flat or somewhat upward sloping during times of low or moderate inflation. In a high inflation regime, the real yield curve is falling. A flat real yield curve is consistent with the random walk hypothesis of consumption. Apparently, in high inflation regimes consumption uncertainty over long horizons is larger than what is implied by the random walk hypothesis.

The shape of the yield curve of risky assets depends on the state of the economy. In bad states, the yield curve is falling, and in good states, rising. Hence, the value of a claim on consumption differs substantially, depending on the initial state. Currently, the US equity premium yield curve is hump-shaped, with a short-run equity premium of 4% that rising in the medium run to an annualised return of 5%, and falling again in the long-run to about 4%. As the real risk-free yield curve is 1% at short maturities and 2% at 30 years maturity, this implies a short-term expected real return to equity of 5% and a long-term annualised return of 6%.

2.3.6 Mean Reversion?

The review of the literature in Sections 2.3.2-2.3.5 lead to the following conclusions

1. There is fairly strong evidence for mean reversion in equity returns
2. Mean reversion in asset returns does not imply reduced risk exposure for long-term investors

There is both direct and indirect evidence for mean reversion. First, asset returns show systematic dependence on *stationary* economic variables, notably the price-dividend ratio over medium-term horizons, and the consumption-wealth ratio over short- to medium-term horizons (Section 2.3.2). Second, there is substantial indirect evidence in terms of autocorrelation in volatility and the price of risk. Volatility is clearly correlated with the onset of a recession, and has a half time of about one year. ARCH models and stochastic volatility models show that there is a volatility-in-mean effect, implying that equity premia go up at the start of a recession, i.e. at the same time that realized excess returns fall (Section 2.3.3). Third, mean reversion is reinforced by the finding that the price of risk varies countercyclically, and usually rises during a recession (Section 2.3.4).

The existence of mean reversion in asset returns does not mean that it can be usefully exploited in terms of an investment strategy, however. First, out-of-sample prediction of asset returns is much more difficult than in-sample prediction, due to parameter uncertainty. Second, high expected excess returns coincide with, and are triggered by, high and persistent volatility. The above-normal expected return must be balanced against the above-normal return risk.[38] Third, the price of risk rises during recessions, which signals that the risk-absorbing capacity of investors is limited. Private investors worry about their consumption possibilities, along the lines sketched by Campbell and Cochrane (1999), and start to sell. Institutional long-term investors may be less risk-averse, but they need to meet specific solvency conditions, and with falling stock prices, these conditions become more stringent as well. Indeed, in a standard Value-at-Risk framework, the increase in volatility may boost the price of risk of institutional investors just as much as that of private investors.

2.3.7 Rare Disasters

Aggregate equity returns have an annual standard deviation of about 16%. A fall of the stock market of 50%, as happened in 2008, is therefore at a probability level of 4σ and should occur once every 30,000 years if returns are lognormally distributed. The actual frequency of such crashes is however substantially higher. The trigger may be a financial crisis, but also war.

Events like financial crises are commonly interpreted as a system failure, with nonlinear effects of the original shock. As a result, these ‘rare disasters’ cannot conveniently be modelled as a volatility spike, of the type described in Section 2.3.3 above. First, the effects of the shocks are invariably negative. Second, these effects appear to have much greater persistence than the effects of ‘common’ volatility shocks, which have a half-life of one year (Cerra and Saxena (2008) and Reinhart and Rogoff (2009) show that the aftermath of a financial crises takes more than 5 years). Third, the size and duration of their effect is endogenous.[39]

In terms of the goal of the present study, the issue is how rare disasters can be modelled as a

[38] Time-varying volatility is missing from the VAR analyses of Campbell and Viceira (2002) and Campbell and Viceira (2005) on the argument that it is short-lived. However, a half life of one year of volatility corresponds very well with the average duration of a recession.

[39] Bryant (1980) and Diamond and Dybvig (1983) specify models of financial crises in which random shifts in consumer behaviour may generate liquidity problems. In case of a sudden surge of liquidity demand, banks need to liquidate long-run assets prematurely to satisfy consumer withdrawals, which can only be done at a loss. Allen and Gale (1994, 2004) extend the models of Bryant and Diamond and Dybvig by introducing an asset market. In their model, a small shock in the liquidity preference of consumers may induce large shocks in asset prices. The size of volatility depends primarily on the level of liquidity in the market. If the amount of liquidity held by banks is small, forced selling may lead to a sharp fall in asset prices, and bank default. The volatility risk arises endogenously as a result of incomplete insurance markets: the liquidity preference shock is private information, and cannot be insured against.

risk factor. Barro (2006) and Barro and Ursúa (2008) model rare disasters in terms of a random event with low probability. In their dataset of series of consumption and GDP disasters, due to financial crises, war, etc., disasters occur with a probability of about 3.5% per year, with a mean effect on GDP of 20%, and a mean duration of 3½ years. In a large majority of cases, equity prices fall during a disaster.[40] Conditional on a fall, equity prices drop by 30% on average. By including the possibility of consumption disasters in the stochastic discount rate, Barro (2006) obtains a realistic equity premium.[41,42]

The notion of rare disasters is related to the discussion of unit roots in GDP in Section 2.2.1. The analyses of Reinhart and Rogoff (2008), Cerra and Saxena (2008), and Reinhart and Rogoff (2009) suggest that financial crises are a source of structural breaks in GDP growth. Barro et al. (2009) generalise the analysis in Barro and Ursúa (2008) by allowing for both permanent and transitory effects of disasters and for a spread of disasters over several periods. A disaster causes a permanent fall in consumption, plus 'undershooting' of the long-run outcome. They find that allowing for partial recovery after a disaster substantially raises the level of risk aversion needed to explain the equity premium.

Assessment

Rare disasters occur with a probability of between 1% and 3% per year. They have an impact on consumption and GDP that is several times larger than that of standard business cycle risks. As a result, these risks have a large impact on risk premia and rates of return. Rare disasters *usually* also have a large negative impact on stock returns. As a result, stock returns also have a 'fat lower tail'.

2.3.8 Spill-overs from Demographic Risk on Real Rates of Return

Worldwide ageing affects both supply and demand on capital markets. According to the standard life-cycle hypothesis, consumption smoothing leads to wealth accumulation during the working life and dissaving in retirement. This implies that, *ceteris paribus*, population ageing will lead to a relative abundance of capital, rising real wages and falling rates of return on capital, see e.g. Miles (1999). However, this prediction depends strongly on the assumptions wrt. public spending and social security. Kotlikoff et al. (2001) generate a baseline scenario with a three

[40] Chile and Argentina in the seventies, respectively the eighties, are the exceptions.

[41] Note that in this "reduced-form" approach the number of risk factors does not increase. Instead, the rare disasters affects the *shape* of the risk distribution.

[42] The notion of fat tails in asset returns also arises in Weitzman (2007), who points out that investors do not known the exact distribution of returns, and must estimate the relevant parameters from a finite sample. This typically results in t-distributions for the predicted returns, rather than normally distributed prediction errors. As a result, the equity premium may become arbitrarily large (cf. Geweke, 2001).

percentage-point increase in the real interest rate, due to the crowding out of capital. This contrasts with a fall in the interest rate of one percentage point that is predicted to occur if the social security system were to be gradually privatised.[43]

The pace of ageing is spread unevenly over countries, with OECD countries, Eastern Europe, and China generally ageing faster than most non-OECD countries.[44] This difference in age composition creates an opportunity for capital flows to the slow-ageing countries that may level most of the effect of ageing on factor prices. The effect of ageing on capital flows and world interest rates has been investigated in a number of papers. Brooks (2003b) obtains a fall of the world interest rate of 1%-point in a model without social security. This omission may not be crucial on a global scale. Domeij and Flodén (2006) do include a paygo system in their model, and they conclude to a fall in interest rates of 1.5%-points over the next century and a half. However, Fehr et al. (2004), in a study that includes capital flows among the US, the EU, and Japan, reach essentially the same conclusion as Kotlikoff et al. (2001).

These differences in predictions depend substantially on the assumptions wrt. saving behaviour in developing countries. Fehr et al. (2005) extend the analysis of Fehr et al. (2004) by taking the impact of China on international capital flows into account. This reverses the interest rate scenarios of their previous paper. As China is ageing rapidly, and does not currently support a large social security system, Chinese savings are sufficiently large to compensate for the crowding out effects that dominate factor price changes in the OECD. Without the effect of China on international capital flows, interest rates would increase by 1%,[45] while including China changes the baseline scenario to a predicted *fall* of 1%-point. In addition, the fall in interest rates is predicted to be considerably larger if the OECD countries change their pension system, either by cutting replacement rates of by privatising pensions.

Boersch-Supan et al. (2006) study the impact of ageing on international capital flows, using different capital mobility scenarios. In these scenarios the rate of return on capital falls by about 1%-point.[46] Krueger and Ludwig (2007) and Ludwig et al. (2007) extend the model of Boersch-Supan et al. with idiosyncratic productivity risk and mortality risk. This does not fundamentally affect saving behaviour, and the interest rate predictions are very similar to those in Boersch-Supan et al..

[43] This scenario is indeed an important element of the 2006 Pension Protection Act, which provides for a massive shift to 401(k) plans.

[44] Turkey and Mexico are OECD countries that are ageing slowly

[45] Another change wrt. the previous paper is that government investment is taken into account, which raises the saving ratio

[46] The model used has a calibration problem in that the rate of return in 2000 differs by capital mobility scenario. The long-term rate of return is 6% in all scenarios. Hence, the fall in the rate of return is *larger* in the scenario where capital is mobile worldwide.

The studies cited above do not consider macroeconomic risk. Hence they cannot distinguish between bond returns and equity returns. However, the fundamental reason for a fall in interest rates due to population ageing runs in terms of the ratio between accumulated net wealth and labour supply, i.e. the capital-labour ratio. The basic prediction is therefore that the rate of return on productive capital will fall. Although all other rates of return are linked with the marginal product of capital, they may be affected differently by ageing in a risky environment. Old households have a different portfolio composition from young households because their human capital hedge is much lower. As a result, population ageing *ceteris paribus* should result in lower demand for equity and a higher demand for low risk assets. This shift in demand raises the equity premium (Brooks, 2000). The size of this effect depends on the size of the equity premium itself, however. In Brooks (2002) the effect is rather small.

One explanation of the equity premium that has a bearing on future expectations of the excess return on stock is the junior-can't-borrow argument of Constantinides et al. (2002). This states that the size of the equity premium is determined by the inability of the young to participate in the stock market. As a result, stock is held mainly by older workers, who demand a relatively high risk premium, because their human capital hedge is lower. If the number of elder households increases relative to the number of young and middle-aged households, the reluctance of the elderly to hold equity will push up the equity premium. There are two versions of this story, the asset market meltdown hypothesis, which predicts a fall in stock prices (Poterba, 2001; Abel, 2001), and a predicted plunge of the risk free rate (Brooks, 2003a). Brooks shows that the size of the change in the risk premium is dampened by the presence of a DB social security system, which creates a hedge for retired households and allows them to take on more risk. However, in his stylised model the risk-free rate still falls 5%-points below its steady-state value around 2020 if households are borrowing-constrained, even with a DB social security system in place.

The empirical evidence for an effect of age structure on rates of return is mixed. Bakshi and Chen (1994) find an effect of about 60 basis points on the U.S. risk premium following a one percent change in the mean age. However, Poterba (2004) finds little empirical evidence of demographic structure on asset returns in the United States. Goyal (2004) finds a positive effect of *changes* in the size of the 45-64 cohort on excess stock returns and a negative effect of changes in the size of the 65+ cohort. Ang and Maddaloni (2005) confirm the result of Bakshi and Chen for the U.S., but they find that for most other countries the fraction of retired household has a negative effect on the equity premium. The effect is especially strong for countries with a large DB social security system.

The age composition of the workforce may also affect macroeconomic volatility. Ríos-Rull (1996), Gomme et al. (2004), and Jaimovich and Siu (2009) find that labour market volatility depends in a U-shaped fashion on the age distribution of the work force. Jaimovich and Siu argue that the ageing of the workforce explains about one quarter of the reduction in volatility during the Great Moderation. In view of the strong relation between volatility and asset returns, this provides another channel for an effect of ageing on the equity premium.

Assessment

The rate of return on productive capital will fall as a result of population ageing. The "consensus" estimate is a fall of 1-1.5%-points, which may be compared with an expected rise in the dependency ratio of 30%-points. There is some uncertainty as to the size of the fall, which is related to uncertainty about the fiscal policy pursued by the major OECD countries and the resulting crowding out effects.

There is also evidence of a relation between the age composition of the workforce and the volatility of GDP. Plausibly, this may imply an effect from demographic risk on asset returns as well, but this issue remains open for research.

There is considerably uncertainty as to the development of the risk-free rate. There are theoretical reasons to expect an increase in the equity premium as a result of ageing. If so, the predicted fall in the rate of return on capital will lead to an even larger fall in the risk free rate. However, the empirical evidence for a relation between age structure and risk premium is both weak and conflicting.

2.3.9 Spill-Overs from Demographic Risk to Inflation risk

Spill-overs from inflation risk to real returns are implicit in the VAR models of Campbell and Viceira and Hoevenaars et al.. In these models, the nominal yield has a strong negative feedback on the excess stock return, and a positive feedback on the real short term yield. However, in terms of the inflation forecast, the model implies a simple regression on past inflation rates, long term interest rates, and the yield spread. Other macroeconomic variables are not considered. One indicator that comes to mind is government debt or, more generally, the government intergenerational balance. The lack of sustainability of current pension schemes in most western countries may create inflationary pressure.

Lindh and Malmberg (2000) argue that demographic shifts may create inflationary pressure through changes in saving rates. In a linear regression of CPI inflation on population age shares for a pooled regression of OECD countries, they find a significant negative effect on inflation of the size of young and middle-aged cohorts. The effects are fairly substantial, a with a cohort size half-elasticity of about -5 for the middle-aged cohort. The evidence of a positive effect on inflation of old cohorts is rather mixed, however. Only young retirees appear to affect inflation positively, whereas old retirees have a significant negative effect on inflation. On the basis of their estimates, the short-run effect of a change in the old-age dependency ratio on the rate of inflation would be between zero and one percent, and the long-run effect at most 3%. Lindh (2004) finds essentially the same result using a richer set of explanatory variables. He also shows that age structure information helps in out-of-sample forecasts of medium-term inflation.

Bettendorf and De Wachter (2007) investigate the effect of age composition on the relative price of non-tradeables. They find a statistically significant effect of the old age dependency ratio on the relative productivity of the nontradeable goods sector. Based on this productivity effect, they forecast an increase of the relative price of non-tradables in the EU of about 25% over the next five decades, corresponding to an equal increase in the dependency ratio. Given world market prices, this may lead to an increase in the aggregate price level of about 10%, or a rise in inflation of about 0.2% over the period. The rise in the price level of non-tradeables would affect elder cohorts in particular, in view of their higher expenditure share of non-tradeables.

Assessment

The empirical evidence for an effect of age structure on inflation is fairly weak. In terms of the effects of ageing on the general price level the results are inconclusive. Ageing may affect the price of non-tradeables. The estimated size of this effect, if it exists, on the price level is in the order of a one percent increase in the price level for each point increase in the dependency ratio. As the dependency ratio is expected to rise by one percent per year over the next 25 years, this may push up inflation by between 0.25%-point and 1%-point over this same period.

2.3.10 Spill-overs from productivity to asset market returns

In the long run all factor-augmenting technical progress must be labour augmenting. Labour-augmenting technology shocks θ_L have a positive short-run effect on marginal capital productivity given by $\frac{\partial \ln \partial y/\partial K}{\partial \ln \theta_L} = \frac{s_L}{\sigma}$ where s_L denotes the cost share of labour and σ the substitution elasticity. For $\sigma \approx 0.5$ and $s_L \approx 0.7$, the elasticity is about 1.4. However, the net effect of TFP shocks on rates of return also depends on the labour response. King and Watson (1996) and Beaudry and Guay (1996) show that output shocks have a negative effect on real interest rates, with a half-elasticity of about -0.3. Galí (1999) claims that TFP shocks in the short run lead to a reduction in hours worked that severs any short-term link between technology shocks and capital returns. This argument is taken further by Beaudry and Portier (2006, 2005), who demonstrate that stock prices lead TFP growth by a few years, suggesting that news of TFP innovations precedes the actual increase in productivity. Their results are supported by Avouyi-Dovi and Matheron (2006), who show that both for the US and the Euro area stock prices are negatively correlated with productivity at high frequencies, and positively at periods of between six quarters and eight years. For the US, these results are statistically significant, less so for the Euro area. Jaimovich and Rebelo (2009) and Lorenzoni (2009) construct calibrated business cycle models to show that noisy information about productivity growth can generate expectations-driven business cycles.

Indirect support for the announcement effects interpretation of the results in Beaudry and Portier is offered by Campbell and Shiller (2001), who report that dividend-price ratios do *not* help to predict future productivity growth. This is in agreement with the dividend smoothing hypothesis (see Appendix D), which predicts that dividends and market value react in equal proportion to future productivity shocks.

The finding that stock prices lead TFP growth suggests the existence of an unobserved state variable in TFP growth, about which current productivity shocks offer only limited information. This effect may be modelled by generalising the stochastic process for TFP in (2.2b). Let (ζ, μ_ζ) be the hidden state variables for the TFP process,

$$\ln y/L = \zeta_t + \varepsilon_t \tag{2.8a}$$

$$\zeta_t = \zeta_{t-1} + \mu_{\zeta,t-1} \tag{2.8b}$$

$$\mu_{\zeta,t} = \mu_{\zeta,t-1} + \rho\,(\psi - \mu_{\zeta,t-1}) + \eta_t \tag{2.8c}$$

A shock in TFP can be the result of either a realization of ε or η. Given that there is substantial autocorrelation in productivity growth, what matters for future productivity forecasts is the estimate of μ . If investors have information about the origin of the shock, they will be able to forecast future TFP based on knowledge of the underlying state variable. However, if the state variable μ_ζ cannot be directly observed, there will be a lag between TFP shocks and revisions in estimates of μ_ζ. To the outside observer it will therefore appear as if stock prices lead TFP growth. Using the estimates in Beaudry and Portier (2006), a 15%-point increase in stock prices is associated with a 1% increase in TFP. It follows that the effect of a one percent shock in TFP *growth* (η) on stock returns is about $\exp[0.15/\rho] - 1\%$.

It can be argued that the observed positive link at business cycle frequencies between productivity innnovations and equity returns must also hold in the long run. Indeed, if we accept the constancy of labour's share in GDP as a stylized fact, the long-run development of capital income and wage income must occur at the same pace. Along this line, Benzoni et al. (2007) argue that wage income and dividend income are cointegrated. Using an Augmented Dickey-Fuller test, they obtain a significant coefficient of -0.26 for the cointegrating vector over the period 1929-2004. This suggests that wage income and dividend income are strongly positively correlated at period lengths of four years and over. However, over the period 1947-2004 the ADF statistic is not significant at 5%.[47]

[47] Benzoni et al. assume that the cointegration vector consists of log dividends and log wages only. However, if the rate of

Bohn (2009) uses a VAR model to estimate 30-year correlations between productivity and capital returns. He reports a positive correlation of between 30% and 66%, depending on the specification of the VAR and the cointegrating vector. In addition, the residual volatility in capital returns, conditional on productivity, is only a bit higher than that of production or productivity itself, after correction for a growth trend.

In Section 2.3.7 it is noted that large negative shocks in GDP, or TFP, most often coincide with large negative shocks in stock returns. While there is no causal link from TFP to stock returns in this case, it is still useful to utilize the observed correlation between these two types of events in stochastic simulation of risk factors. The inventory made by Barro and Ursúa (2008), table C2, suggests that stock prices fall on average more than GDP.

Assessment

Productivity shocks have a positive effect on contemporaneous *realized* excess capital returns, with a half-elasticity elasticity of around unity. They have a negative short-term effect on real interest rates, with a half-elasticity of around -0.3. Shocks to (unobserved) structural TFP growth have a strongly positive effect on excess stock returns. However, this correlation between stock returns and productivity takes two years and more to materialize. There is also empirical evidence that dividends and productivity *levels* are cointegrated. This implies that the relation between stock returns and productivity growth should also hold in the long run.

return to capital shifts as a result of ageing, this assumption is not valid. In a perfectly competitive Cobb-Douglas world without taxes, steady-state dividends are given by

$$\frac{D}{w\,L} = \frac{y - w\,L - I}{w\,L} = \frac{p_k\,K}{w\,L}\left(1 - \frac{\delta + \psi}{p_k}\right) = \frac{1-\alpha}{\alpha}\,\frac{r_k - \psi}{r_k + \delta}$$

where δ is the depreciation rate of capital, ψ is the growth rate of capital, and α is labour's share in production. It follows that the cointegrating relation is

$$\ln D - \ln(w\,L) - \ln\left((r_k - \psi)/(r_k + \delta)\right) + \ln\left(\alpha/(1-\alpha)\right) \tag{2.9}$$

If the r_k term in (2.9) is stationary, it may be dropped. However, the point is that the ageing process is expected to lead to a *permanent* shift in the return to capital. Consequently, the r_k term is *not* stationary. E.g. a decline of 1%-point in capital returns leads to a fall of $\frac{r_k-\psi}{r_k+\delta}$ in (2.9) of more than ten percent, depending on the exact value of δ (≈ 0.06) and ψ (≈ 0.02). The failure to include the return to capital in their cointegrating vector may explain why Benzoni et al. do not find a significant ADF statistic over the post-war period.

3 A State Space Model of Social Security Risks

As argued in Section 1, the return on social security generally depends on three factors, the old-age dependency ratio, (labour) productivity, and capital market returns. This section presents a quantitative model to forecast these factors, and their joint distribution, based on the literature review in Section 2. The risk factors listed in Section 2 include demographic risks, productivity, inflation, and rates of return on bonds and stock. To complete the asset market description, we add the risk-free rate. As in Campbell and Viceira (2005), it is convenient to model the dynamics of the excess returns over the risk-free rate. The state vector is

$$\mathbf{x}_t' = \left(\mathbf{x}_{1,t}', \mathbf{x}_{2,t}'\right) \tag{3.1a}$$

$$\mathbf{x}_{1,t}' = \left(z_{1,t}, z_{2,t}, \mu_{\lambda,t}, \zeta_t^{NL}, \zeta_t^{US}, \mu_{\zeta,t}^{NL}, \mu_{\zeta,t}^{US}, \ln R_{f_t}, \ln R_{b_t}^e, \ln R_{k_t}^e, \ln R_{f_t}^e, \ln D/V, \ln V, \delta_t\right) \tag{3.1b}$$

$$\mathbf{x}_{2,t}' = \left(\mathbf{n}_t', \mathbf{m}_t'\right) \tag{3.1c}$$

where the z_i are the state variables of the fertility equation. $R_f = 1 + r_f$ denotes the risk-free rate, and δ is the disaster indicator, $\delta = 1$ in case of a systemic crisis, $\delta = 0$ otherwise. The excess returns are defined as $R_b^e \equiv \frac{1+r_b}{1+r_f}$, $R_k^e \equiv \frac{1+r_k}{1+r_f}$, $R_f^e \equiv \frac{1+r_f}{1+\pi}$. D/V is the dividend-price ratio, and V the (real) market value of equity. $\mathbf{n}_t' = (n_{t,1}, \dots, n_{t,T})$ denotes the vector of age cohort sizes and $\mathbf{m}_t' = (m_{t,1}, \dots, m_{t,T})$ denotes net migration.

The survey of the literature in Section 2 above identifies a number of empirically relevant interactions between the different risk factors, viz. from demographic risk (age composition) to rates of return, from demographic risk to productivity risk, from productivity to rates of return, and interactions between the rates of return in general. Interactions between demographic risk and other risk factors are somewhat special because the relevant link is the age composition of the population, which depends on lagged realizations of the fundamental risk factors fertility and mortality. Consequently demographic shifts are the result of a high-dimensional process that is characterized by long lags

$$\mathbf{n}_{t+1} = \mathbf{H}_t \, \mathbf{n}_t + \mathbf{m}_t \tag{3.2a}$$

$$\mathbf{H}_t = \begin{pmatrix} \phi_{t-1,1} & \phi_{t-2,2} & \cdots & \phi_{t-T+1,T-1} & \phi_{t-T,T} \\ 1-\lambda_{t,1} & 0 & \cdots & 0 & 0 \\ 0 & \ddots & \ddots & \vdots & \vdots \\ \vdots & \ddots & \ddots & 0 & 0 \\ 0 & \cdots & 0 & 1-\lambda_{t,T-1} & 1-\lambda_{t,T} \end{pmatrix} \tag{3.2b}$$

$\mathbf{H}$ is the demographic transition matrix.[48] I assume that the age composition of the population

[48] Fertility $\phi_{t,\tau}$ is indexed by *birth year t* and age τ.

can be adequately captured by a single variable, the dependency ratio d_t

$$d_t = \sum_{\tau=65}^{T} n_{t,\tau} / \sum_{\tau=18}^{64} n_{t,\tau} \tag{3.2c}$$

d_t is a nonlinear transformation of past realizations of the demographic risk factors.

The complete state space model is given by dynamic equations for the state variables (3.3d-3.3f), observation equations for the hidden state variables (3.3a-3.3c), and definition equations for the dependency ratio, total population, and total migration (3.3i). The interactions between these variables may be modelled as a VAR, as in the Campbell/Viceira model discussed in Appendix E. However, it is more convenient to use an error-correction formulation, so as to include the long-run impact of demographic factors and productivity on rates of return.

$$\ln \phi_{t,\tau} = \alpha_{\phi,\tau} + \beta_{\phi,\tau} z_{1,t} + \gamma_{\phi,\tau} z_{2,t} + \varepsilon_{\phi,t+1,\tau}, \qquad \tau = 1,\ldots,T \tag{3.3a}$$

$$\ln \lambda_{t,\tau} = \alpha_{\lambda,\tau} + \beta_{\lambda,\tau} \mu_{\lambda,t} + \varepsilon_{\lambda,t+1,\tau}, \qquad \tau = 1,\ldots,T \tag{3.3b}$$

$$\ln \frac{y_t}{L_{h_t}} = \zeta_t + \varepsilon_t \tag{3.3c}$$

$$\Delta \mathbf{x}_{1,t+1} = \mathbf{\Psi} \Delta \mathbf{x}_{1,t} + \boldsymbol{\psi}_0 + (\mathbf{I} - \mathbf{\Xi}) \left(\mathbf{x}^*_{1,t} - \mathbf{x}_{1,t} \right) + \boldsymbol{\varepsilon}_{x,t+1} + \delta_{t+1} \boldsymbol{\varepsilon}_{\delta,t+1} \tag{3.3d}$$

$$\mathbf{x}^*_{1,t} = \mathbf{c}_0 + \mathbf{C} \mathbf{x}_{1,t} \tag{3.3e}$$

$$\mathbf{n}_t = \mathbf{H}_{t-1} \mathbf{n}_{t-1} + \mathbf{m}_t \tag{3.3f}$$

$$\frac{\mathbf{m}_t}{\bar{n}_t} = \mathbf{a}_0 + \mathbf{a}_1 \ln t + \mathbf{A}_2 \frac{\mathbf{m}_{t-1}}{\bar{n}_{t-1}} + \mathbf{a}_3 \frac{\bar{m}_t}{\bar{n}_t} + \mathbf{a}_4 \frac{\bar{m}_{t-1}}{\bar{n}_{t-1}} + \boldsymbol{\varepsilon}_{m,t} \tag{3.3g}$$

$$\frac{\bar{m}_t}{\bar{n}_t} = \alpha_0 + \alpha_1 \frac{\bar{m}_{t-1}}{\bar{n}_{t-1}} + \sum_{i=0}^{2} \alpha_{i+2} \ln \zeta_{t-i} + \varepsilon_{m,t} \tag{3.3h}$$

$$d_t \equiv \sum_{\tau=65}^{T} n_{t,\tau} / \sum_{\tau=18}^{64} n_{t,\tau}; \qquad \bar{n}_t \equiv \sum_{\tau=1}^{T} n_{t,\tau}; \qquad \bar{m}_t \equiv \sum_{\tau=1}^{T} m_{t,\tau} \tag{3.3i}$$

(3.3a is the fertility equation for the cohort of birth year t and age τ. The fertility distribution of the cohort depends on two unobserved state variables, z_1 and z_2, that describe the development of total (completed) fertility of the cohort and the timing of fertility, respectively. This specification generalises the Lee-Carter model discussed in Section 2.1. Appendix A provides more detail. (3.3b) describes the mortality at age τ at time t, using a standard Lee-Carter model with a single state variable $\mu_{\lambda,t}$ that may be taken to represent the state of medical technology. In (3.3c), productivity fluctuates around a structural level ζ_t, that grows according to a difference stationary process along the lines of the system discussed in (2.8). See Appendix C for details.

Equations 3.3f are accounting equations that keep track of the demographic structure.

The system of equations in (3.3d) represents the core of the VAR system that describes the development of the state variables. Short-term dynamics are contained in $bm\Psi\Delta\mathbf{x}_{1,t}$ and feedback effects are given by the term $(\mathbf{I} - \mathbf{\Xi}) \left(\mathbf{x}^*_{1,t} - \mathbf{x}_{1,t} \right)$. The $\delta_{t+1} \boldsymbol{\varepsilon}_{\delta,t+1}$ term in (3.3d) represents the effects of a disaster in period $t+1$ on asset returns. δ is a binary random variable, that takes the value 1 ("disaster") with probability 1.38%, in accordance with the recommendation by Barro et al. (2009). The vector $\boldsymbol{\varepsilon}_\delta$ contains the size of the disaster, if it occurs. In $\boldsymbol{\varepsilon}_\delta$ the only non-zero entries are those for the excess equity return and the nominal and real risk-free rate. It is assumed that the real and nominal risk-free rate are affected in equal size, so that the impact of the disaster on inflation is zero. The equity return shock is uniformly distributed on $[-0.7, 0]$ and the real rate shock is a non-random -2% (see Barro and Ursúa, 2008, figure 4 and table 13). That is, a disaster usually causes a stock market crash, and it always lowers the risk-free rate, as a consequence of increased precautionary saving, but it does not have a systematic effect on inflation.

3.1 Parameter Values

The model has a large number of parameters. Only a subset of parameters has been estimated for this paper. These estimates are reported in Appendices A-E. The remaining parameters have been taken from the literature, based on the survey given in the preceding sections. Most parameters are specified in Tables 3.1-3.3, and the remaining ones in Tables A.2 in Appendix A and B.1 - B.2 in Appendix B.

Table 3.1 reports the adjustment coefficients for the state variables, except for the fertility and mortality states z_1, z_2, and μ_λ, which have no interactions with the other state variables.[49] The adjustment coefficients for the productivity states and the return states are block-diagonal. As explained in Appendix C, there is a spill-over term from U.S. productivity to Dutch productivity, that represents a catching-up effect.

Table 3.1 Adjustment matrix (Ξ)

	ζ^{NL}	μ_ζ^{NL}	ζ^{US}	μ_ζ^{US}	$\ln R_f$ [a]	$\ln R_f^e$ [b]	$\ln R_b^e$ [b]	$\ln R_k^e$ [b]	$\ln \frac{D}{V}$
ζ^{NL}	1	1							
μ_ζ^{NL}	– 0.046	0.423	0.046						
ζ^{US}			1	1					
μ_ζ^{US}				0.716					
$\ln R_f$					0.807	0.001	0.003	0.004	0.002
$\ln R_f^e$					0.366	0.031	0.002	0.002	– 0.001
$\ln R_b^e$					0.314	– 0.001	0.004	0.002	– 0.005
$\ln R_k^e$		0.07			– 1.154	– 0.037	– 0.021	– 0.006	0.044
$\ln D/V$					3.239	– 1.788	– 0.302	0.016	0.866

[a] $R_f \equiv 1 + r_f$
[b] $R_x^e \equiv 1 + r_x^e,\ x \in \{f, b, e\}$

[49] The adjustment speed matrix is diagonal in the fertility and mortality states, $\Xi_{z_1,z_1} = 0.952$, $\Xi_{z_2,z_2} = 0.98$, and $\Xi_{\mu_\lambda,\mu_\lambda} = 1$.

The impact effects specified in Table 3.2 are mostly zero. Based on the survey of productivity spill-overs in Section 2.3.10 the impact effect of productivity shocks on returns has been set at -0.3. By assumption, there is no effect on excess returns. Based on the survey in Section 2.3.8, the impact elasticity of shifts in the dependency ratio on productivity is set at -0.1. This corresponds with an assumed long-term effect of equal size in Table 3.3. The coefficient ψ_0 represents the mortality drift speed (identical to $\delta_{\lambda,0}$ in (2.1c)).

Table 3.2 Impact matrix (Ψ)

	$ln\,y/L-\zeta$	Δd	$\boldsymbol{\psi}_0$
Δz_1			
Δz_2			
$\Delta\mu_\lambda$			– 0.36
$\Delta\zeta$			
$\Delta\mu_\zeta$		– 0.1	
$\Delta\ln R_f$			
$\Delta\ln R_f^e$	– 0.3		
$\Delta\ln R_b^e$			
$\Delta\ln R_k^e$			
$\Delta\ln D/V$			

The long-term effects are specified in Table 3.3. The effect of the dependency ratio on rates of return has been discussed in Section 2.3.8. A fall in rates of return of about 1% in relation to an increase of the dependency ratio of 25% leads to an elasticity of -3 in terms of the risk-free rate. The constant term in the fertility risk factor corresponds to a fertility of 1.8 child per women (see Section 2.1). Benzoni et al. (2007) find that wages (productivity) are cointegrated with dividends. A weak form of this cointegration relationship has been included by imposing a restriction on the relation between dividend yields and equity returns:

$$\ln(1-D/V_t^*)+\ln R_k^{e*}=\ln R_f^{e*}+\mu_\zeta^{US}+\ln\bar{n}_t/\bar{n}_{t-1}+\mathrm{E}[\delta\,\varepsilon_\delta] \tag{3.4}$$

The market value of equity must grow in proportion with the size of the economy, and therefore with productivity and the size of the workforce. The growth rate of the market value of firms is $\ln V_t/V_{t-1}=\ln R_k+\ln(1-D/V)$, from which (3.4) follows. A boost in U.S. structural productivity growth thus leads to a fall in the equilibrium dividend-price ratio, temporarily above-normal dividend yields, and higher excess returns on equity during the transition period.

The constant terms of the rates of return have been calculated from Table 1 in Campbell and Viceira (2005), with the exception of the constant in the D/V row, which equals $0.0152-0.7571+0.0138\times0.35$. Note that the real and nominal risk-free rate are effectively a

Table 3.3 **Cointegration matrix ($\mathbf{C}$)**

	μ_ζ^{US}	$\ln \bar{n}/\bar{n}_{-1}$	$d - d_{2006}$	$\mathbf{c}_0$
z_1				– 0.4735
z_2				– 0.0551
μ_ζ				0.0264
$\ln R_f$			– 0.03	0.0337
$\ln R_f^e$			– 0.03	0.0152
$\ln R_b^e$				0.0137
$\ln R_k^e$				0.7571
$\ln 1 - D/V$	1	1		– 0.7366

submodel within the VAR model.[50] This submodel can also be written in terms of the nominal risk-free rate and the rate of inflation, but this would imply a more complicated lag structure.

[50] in Campbell and Viceira (2005) the other parameters of these equations are not significant, with the exception of the spread, which is not included here

4 Simulation Results

Figures 4.1-4.17 provide estimates of the distribution of a number of variables over time, based on 10000 Monte-Carlo runs. The graphs show the development over time of the median (solid lines), the mean (dashed lines), the 2.5% tails, and the 0.5% tails. In addition, there are a few figures that show estimated autocorrelation functions. The graphs are arranged by return and risk factor, first demographic risk (Figures productivity risk,

4.1 Demographics

Figures 4.1-4.3 provide an overview of the uncertainty in demographics. According to Figure 4.1, the dependency ratio will increase almost surely in the medium run. The long run development is less certain. Dependency ratio uncertainty increases slowly over the next thirty years, due to uncertain mortality and migration, as the future workers over that period have already been born. There is a boost in dependency ratio uncertainty when cohorts of currently unborn workers start to enter. A second increase in uncertainty occurs as these workers retire.

The larger part of dependency ratio uncertainty is due to fertility. Figure 4.2 shows that total fertility may end up anywhere between 1.5 and 2.7 children per woman.[51] In addition, uncertainty is boosted by the effect of net migration. If productivity growth is strong, this leads to a larger inflow of young age cohorts, which lowers dependency ratios somewhat. According to the estimates in Appendix B, migration is expected to depress the dependency ratio by about five percentage points.

The effect of future mortality rates on the life expectancy of a 65 year old person is given in Figure 4.3. Life expectancy increases due to the decline in expected mortality. However, uncertainty is sufficiently large that, in the medium run, an increase in mortality is not inconceivable. In fifty years, life expectancy of a 65-year old may have risen to 90, but it may also have remained constant.

[51] Total fertility is an artificial characteristic, as it is computed as a cross-section of fertility rates at a given point in time. However, in the long run this rate coincides with the completed fertility of a cohort, if the latter is stationary. The underlying uncertainty is depicted in Figure A.1 in Appendix A, which shows the confidence intervals of the two state variables that determine fertility developments. The point is that there is substantial uncertainty in the *current* value of these state variables, because fertility behaviour of current cohorts has so far been only partially observed.

Figure 4.1 Dependency ratio

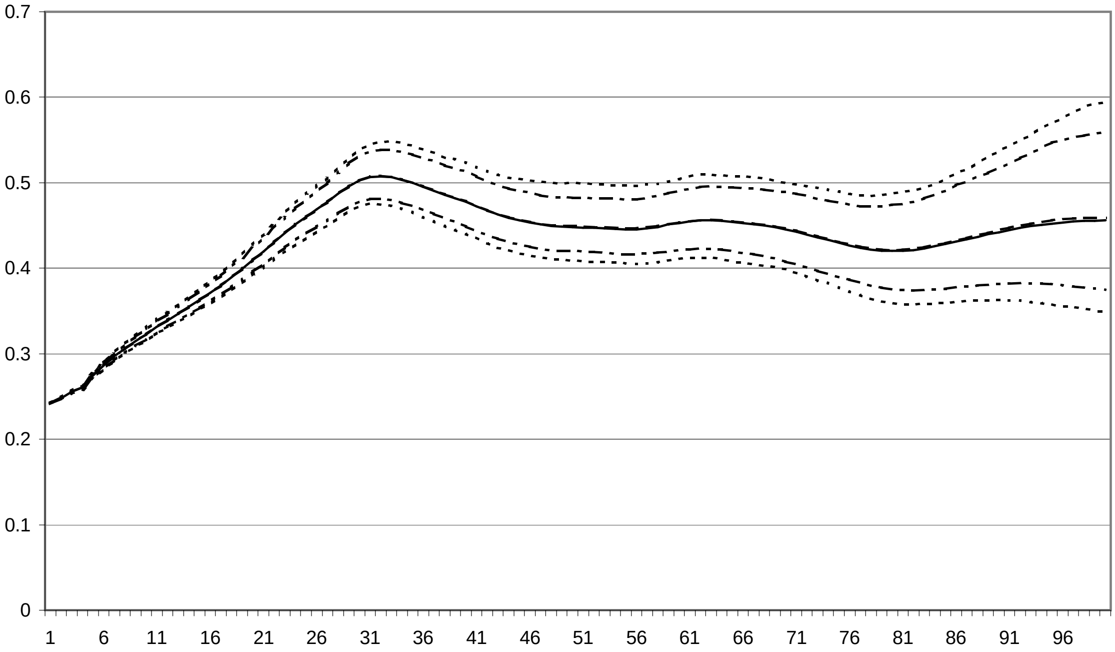

Figure 4.2 Total fertility

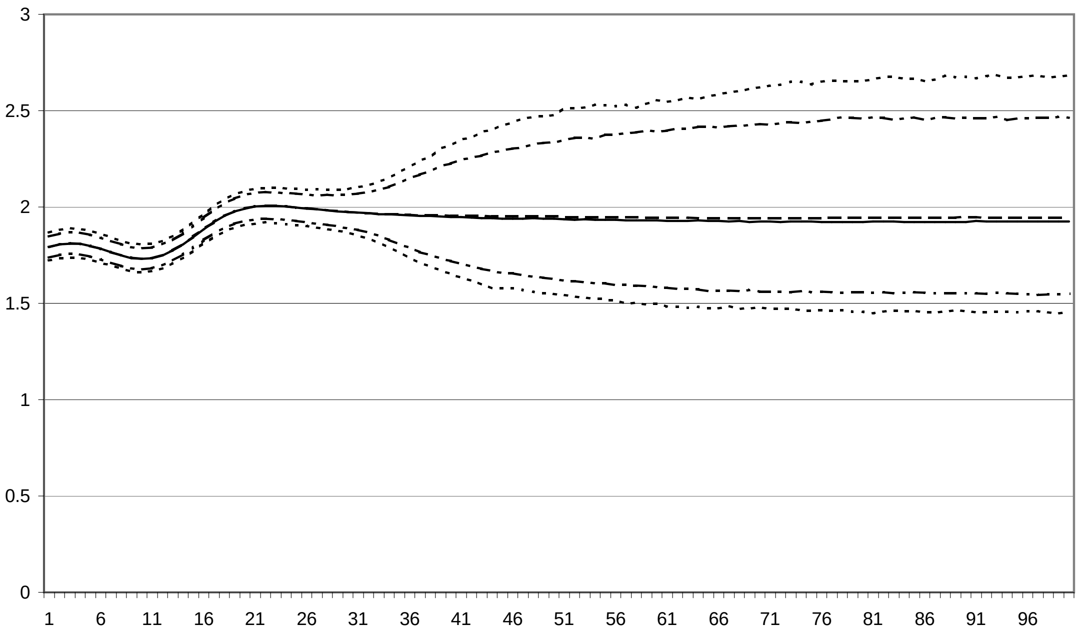

Figure 4.3 Expected life of people of age 65

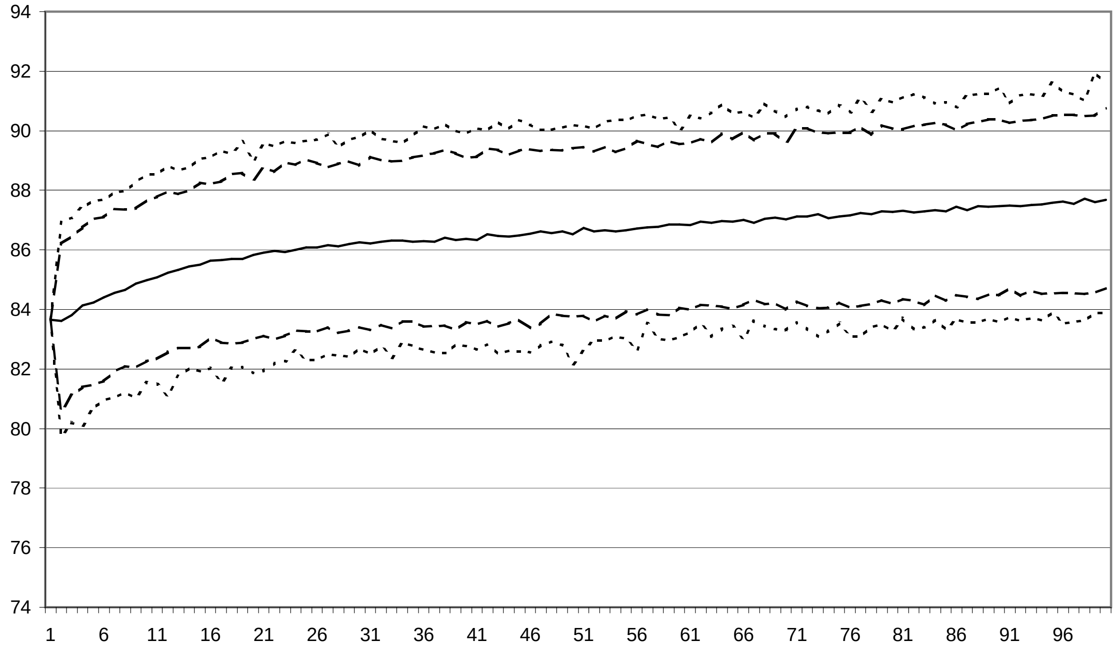

4.2 Productivity

Figures 4.4-4.6 present the outcomes for the labour productivity process. In the long run, labour productivity growth is on average 1.8%. However, the recovery of productivity growth from the historically low rates around the turn of the century takes about two decades. First, during the next two decades the demographic shift lowers expected productivity growth by about 0.1%. Moreover, catching up with productivity growth in the United States takes time. Figure 4.4 shows that there is considerable uncertainty about the size of productivity growth in any given year. In addition, the lower tail of productivity growth is stretched out considerably as a result of the occurrence of rare disasters, as discussed in Section 2.3.7. With small probability, shocks of up to -10% may occur, and recovery takes time.[52] Shocks in productivity growth are persistent, as displayed in Figure 4.5. This characteristic of the model corresponds well with the increasing variance profile of productivity levels for the Netherlands as presented in Section 2.2.1. As a result of the unit root in productivity, long-run productivity levels are rather uncertain (Figure 4.6). In a century, productivity may be anywhere between the same and eight times the present

[52] These shocks may occur either in the Netherlands or, independently, in the United States. In the latter case, the flow of innovations from the U.S. to the Netherlands peters out for a few years.

level (implying average annual growth rates between 0% and 2.1%).[53]

Figure 4.4 Labour productivity growth

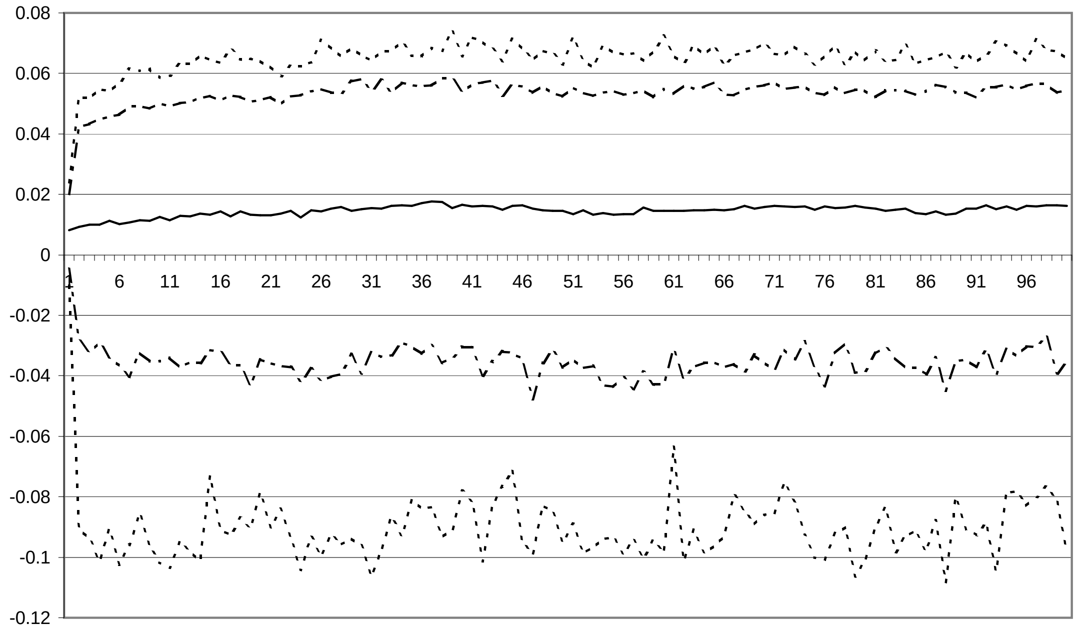

4.3 Asset Returns

Figures 4.7 and 4.8 give the distribution of the real return on 5-year bonds and short-term bills. Median returns and expected returns on both assets fall by about 1.2%, predominantly on account of population ageing (see Section 2.3.8). There is hardly any autocorrelation in real bond returns, but the return on short-term bills is correlated over a period of about 5-10 years (Figure 4.9).

The fall in real rates of return is not related to inflationary pressure (Figure 4.10), as the projection shows only a small upward drift. However, the margin of uncertainty of future inflation is considerable. This is partly caused by the substantial persistence in inflation, which boosts inflation risk, see Figure 4.11. As a result, the return on nominal assets is rather uncertain over long time spans.

Figure 4.12 shows that the real rate of return to equity shows considerable variation, at a standard deviation of 21.5%, a bit higher than in Campbell and Viceira (2005) because of the

[53] The distribution is skewed towards lower productivity as a result of the rare disasters as well as the log-linearity of the productivity distribution.

Figure 4.5 Autocorrelations in labour productivity growth

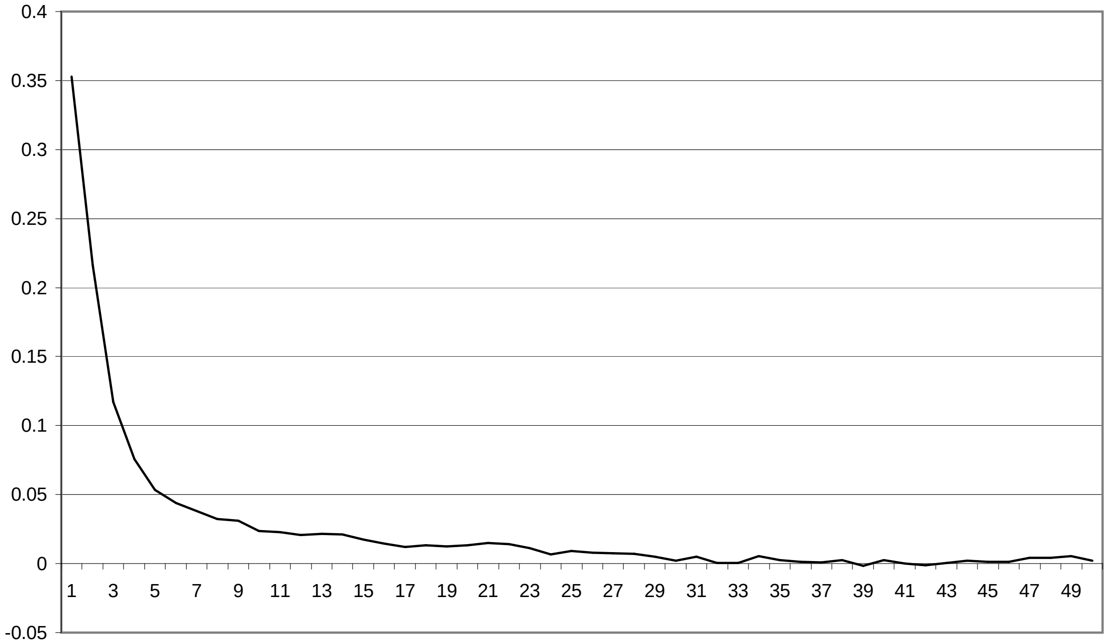

Figure 4.6 Labour productivity

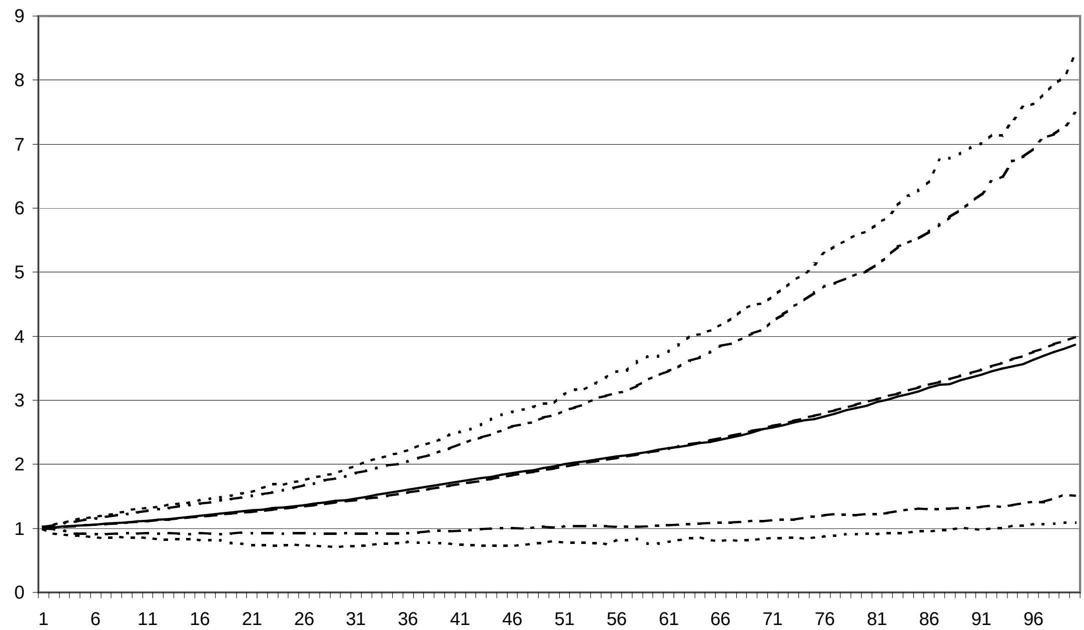

Figure 4.7 Real rate of return to medium-term bonds

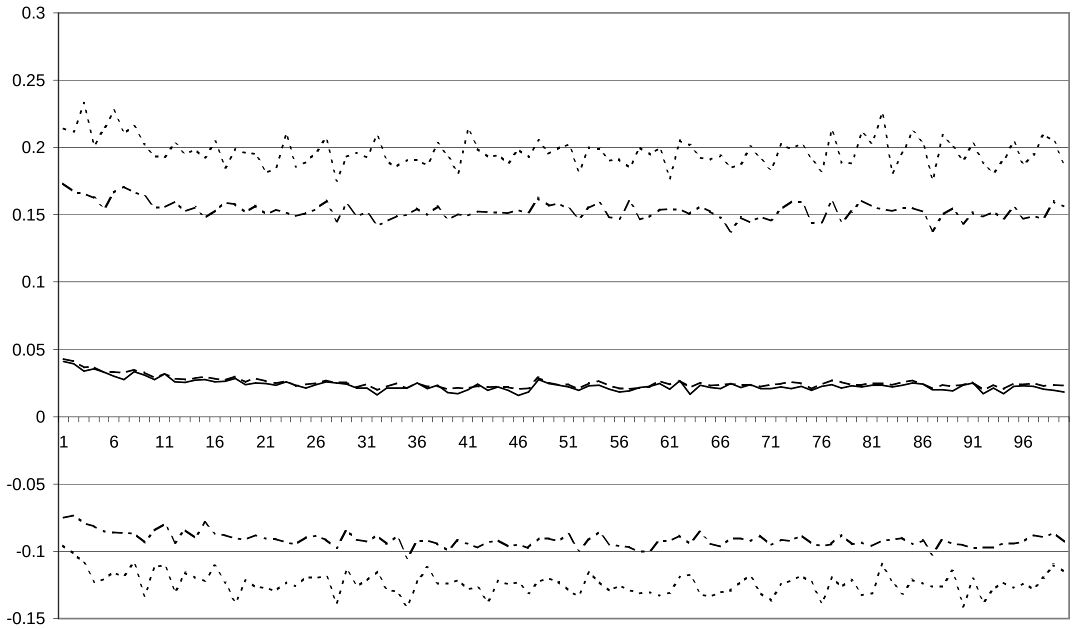

Figure 4.8 Real rate of return to one-year deposits

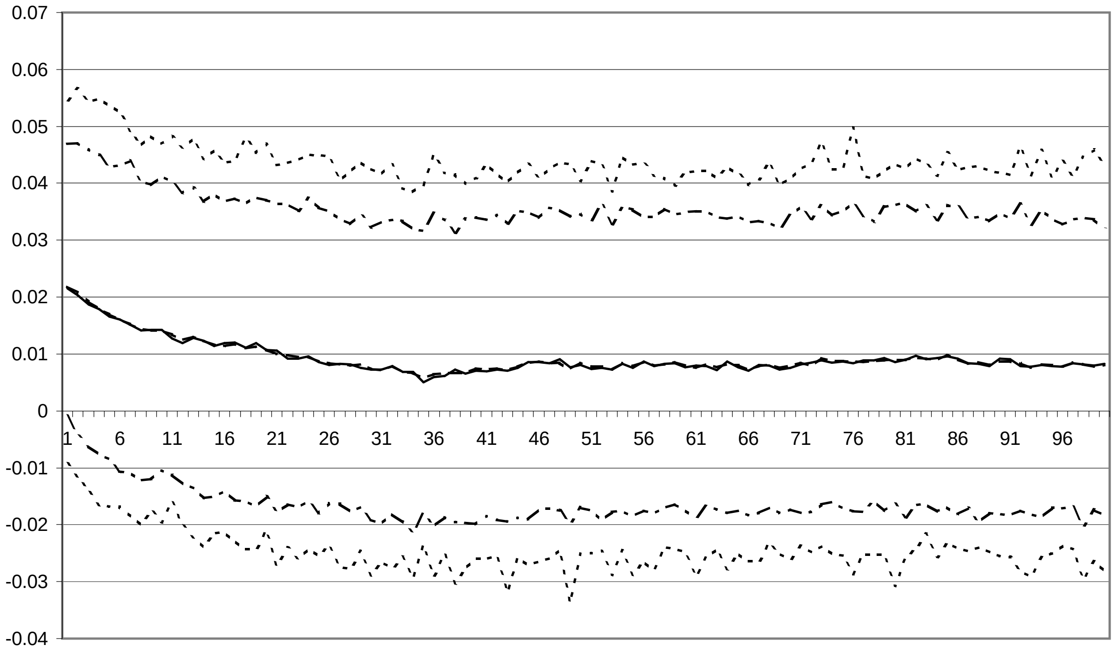

Figure 4.9 Autocorrelations in the real return to one-year deposits

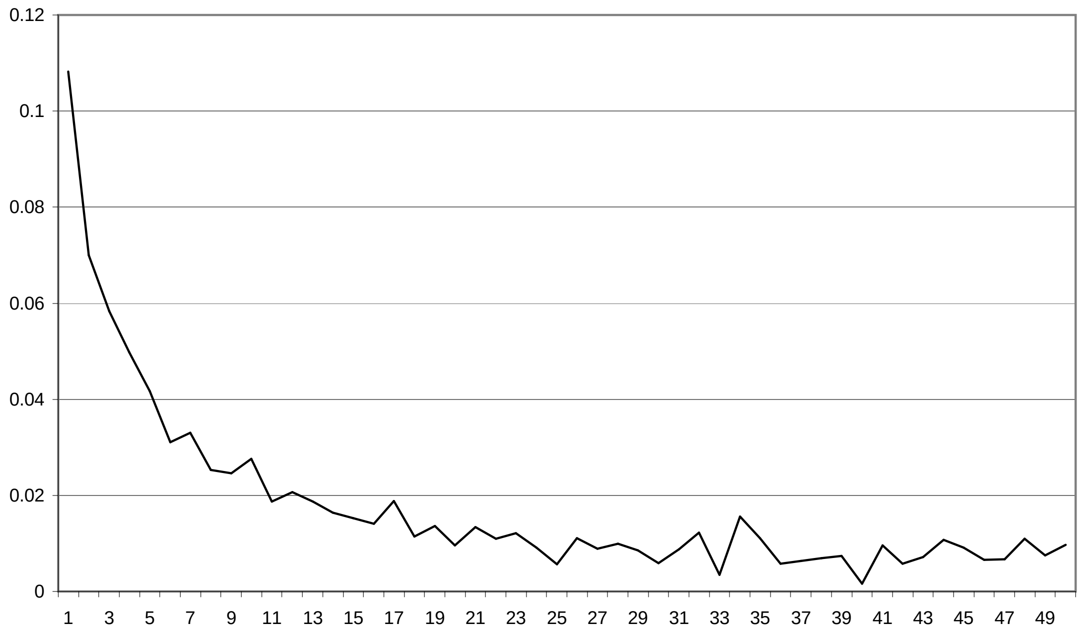

Figure 4.10 Consumer price inflation

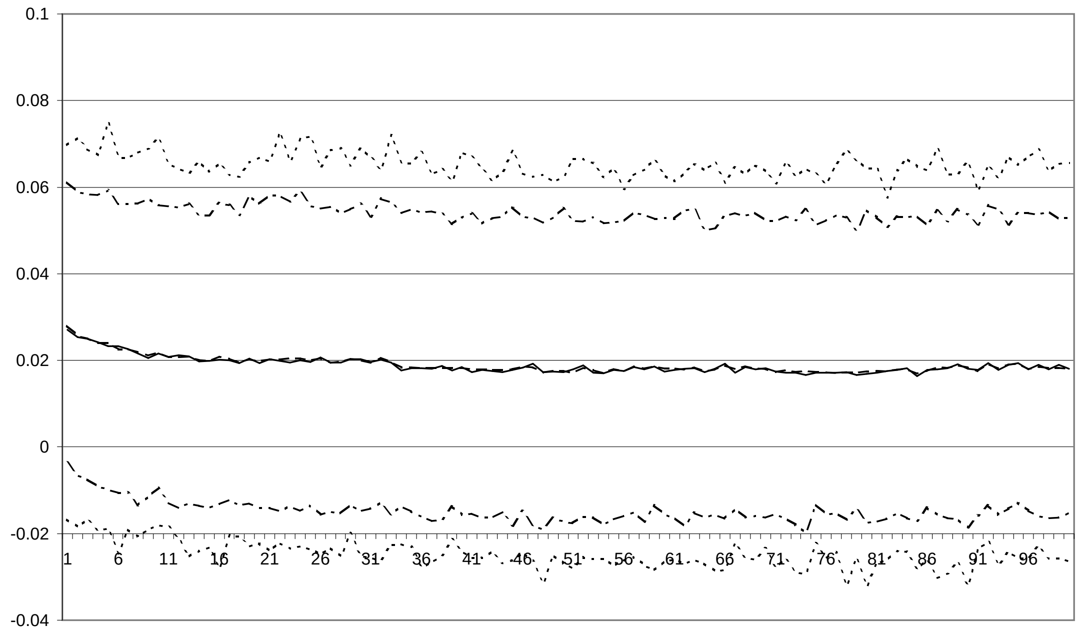

Figure 4.11 Autocorrelations in consumer price inflation

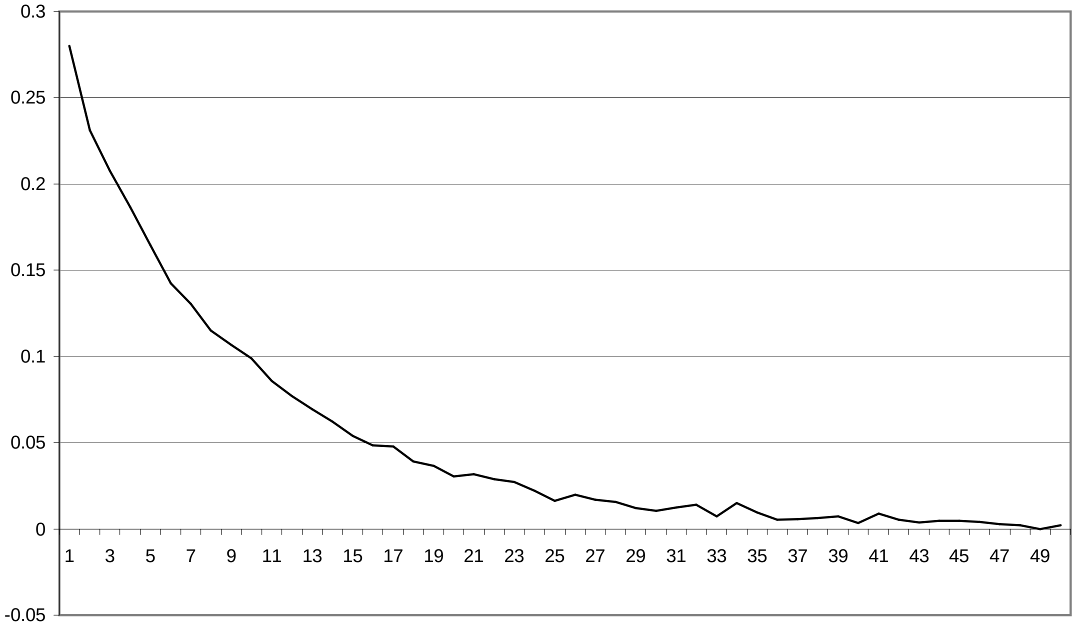

spill-over from other risk factors, including the rare disasters. The large standard deviation of returns makes it difficult to discern from the figure that the median return on equity falls from 7.9% to 7.1% during the first 25 years. The mean return falls from 9.1% to 8.4% during this period, after which it stabilizes around 8.6%.[54] The rare disasters spread out the lower tail of the return distribution, but because the lognormal distribution is positively skewed, the effect of this on the graph is not conspicuous. The autocorrelation function shows only small negative correlations, i.e. little mean reversion (Figure 4.13). This finding is also apparent from the annualised standard deviations of equity returns in Figure 4.14, which indicates only a modest reduction in risk over longer horizons.[55] For the sake of comparison, figure 4.14 also includes the annualized risk involved in the growth rate of PAYG benefits in a pure DC system.

Figure 4.15 presents the distribution over time of the growth rate of benefits in a defined-contribution PAYG scheme. In a DC scheme, benefits grow with labour productivity minus the growth in the dependency ratio. The expected real growth rate of benefits is negative for the next three decades in a DC scheme, on account of the rising dependency ratio and the below-par productivity growth. The case where the demographic risk is shifted fully to workers

[54] The deviation between the mean and median originates from the (approximate) log-normality of the distribution of r_k, combined with the large standard deviation.

[55] This finding contrasts with the results presented in Campbell and Viceira (2005), who find a substantial fall in equity risk with increasing horizon. A numerical check of their variance profiles using their published parameter estimates reveals however that their annualised equity variance profile is completely flat.

Figure 4.12 Real rate of return to equity

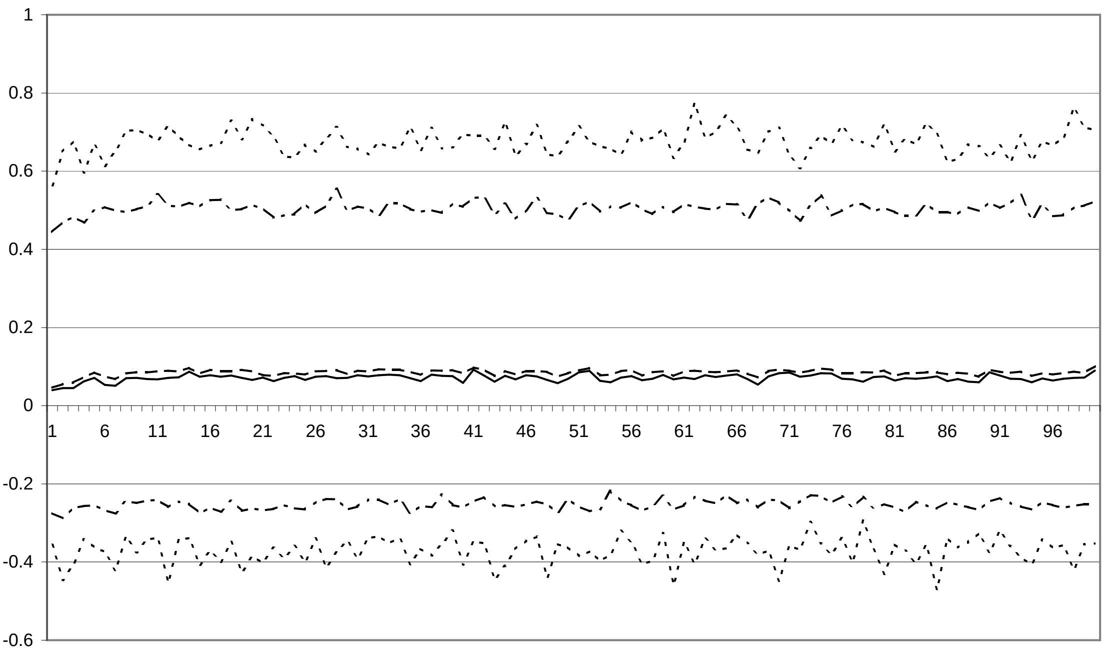

Figure 4.13 Autocorrelations in equity returns

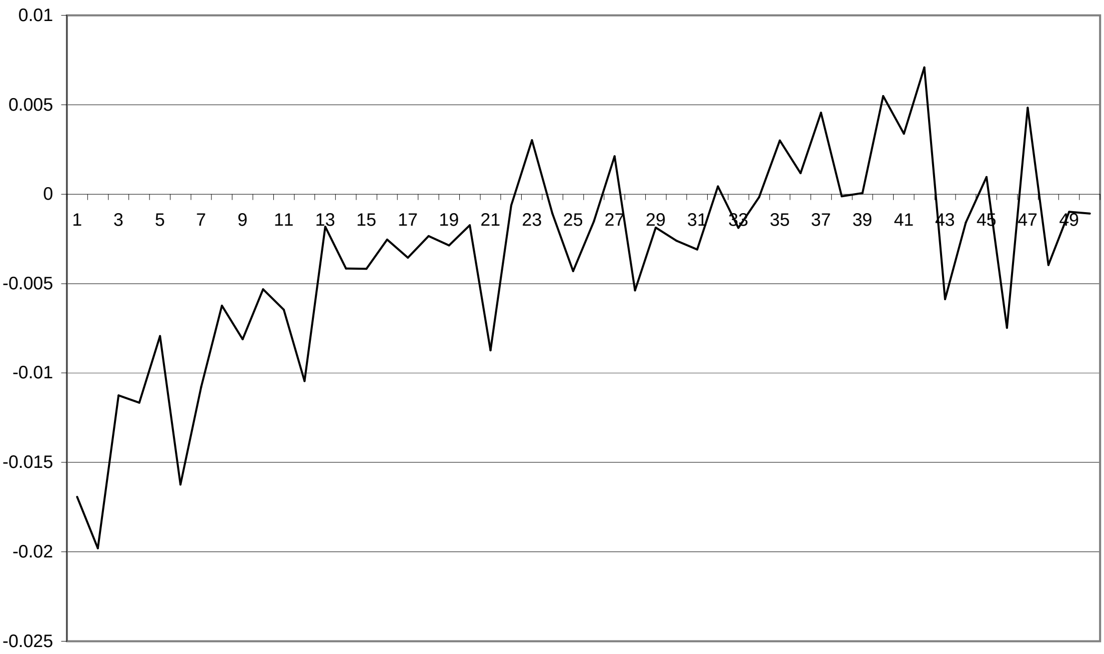

Figure 4.14 Annualised standard deviations of asset returns

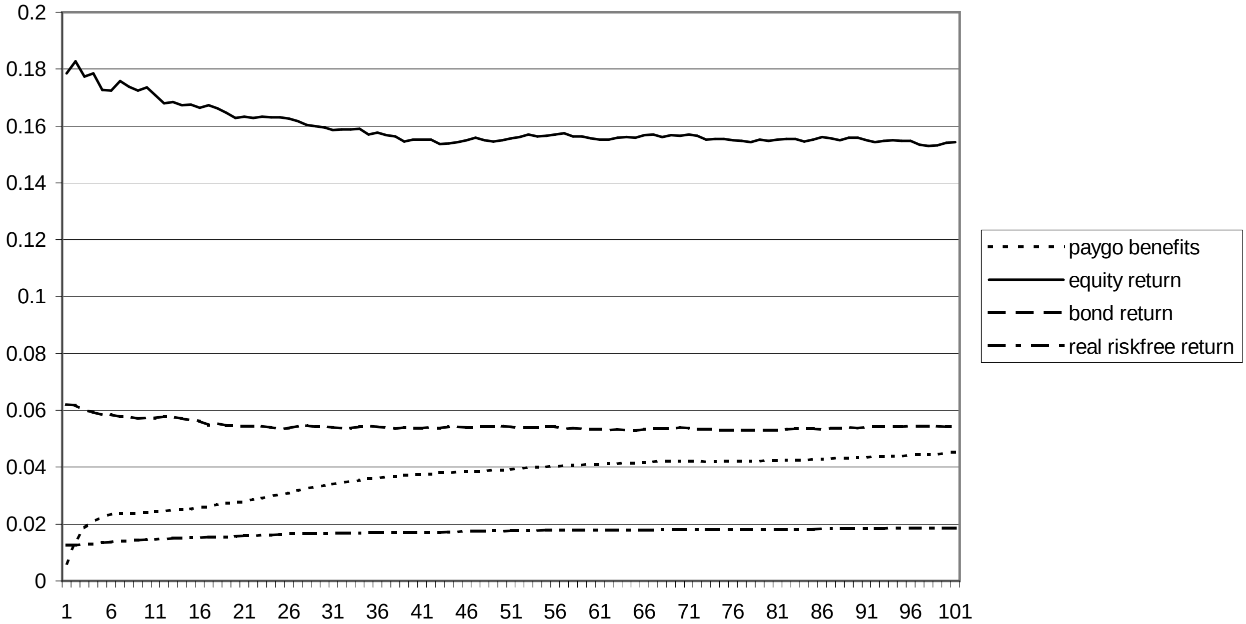

by maintaining a Defined Benefit scheme is depicted in Figure 4.16. In this case the expected doubling of the PAYG contribution rate lowers the growth of net real wages. Comparing Figure 4.16 with Figure 4.15 shows that the expected effect of demographic shifts on net wages in a DB scheme is considerably less than the effect on benefits in a DC scheme. However, the amount of *risk* is very similar in both schemes, as the demographic change is well predictable. The differences between Figures 4.15 and 4.16 thus mainly reflect differences in intergenerational redistribution.

The issue of how to address the implicit threat to the stability of the scheme is the subject of a policy debate in the Netherlands. In this respect, it is relevant that there is considerable uncertainty about the growth rate of benefits even over the next decade. In the long run, when the demographic transition has completed, the expected PAYG benefit growth is in line with expected productivity growth. Still, demographic changes induce emphpredictable fluctuations in PAYG benefits and/or contributions even a century ahead.

The uncertainty about future PAYG benefit levels mainly derives from the uncertainty in productivity growth. Because of the possibility of production disasters, the downward risk is substantially larger than the upward risk. Figure 4.17 shows that the correlation profile of PAYG benefit growth resembles that of productivity growth. As a result, annualised PAYG benefit risk increases with the length of the horizon. In the long run PAYG is almost as risky as medium-term bonds, as shown in Figure 4.14.

Figure 4.15 Growth of real per capita benefits in a DC PAYG pension system

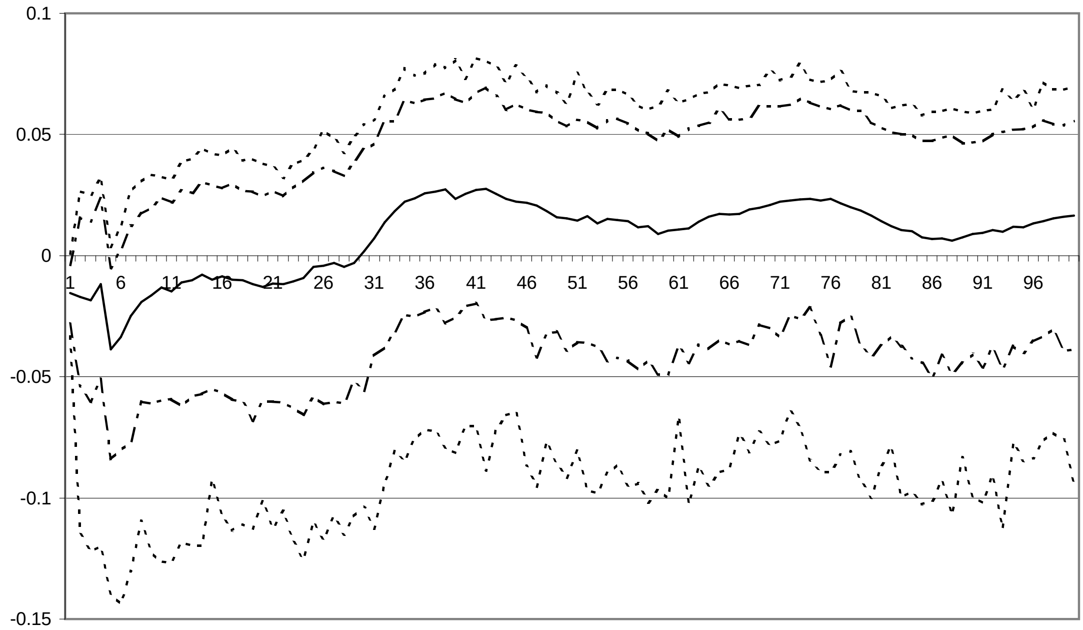

Figure 4.16 Growth of net real wages in a DB PAYG pension system

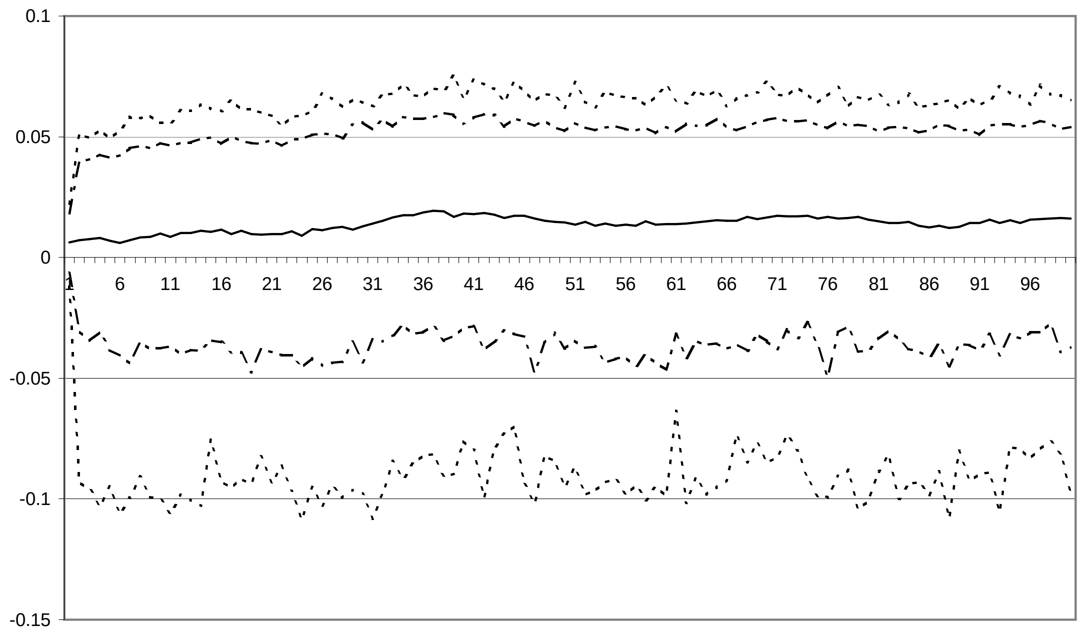

Figure 4.17 Autocorrelations in real benefit growth in a DC PAYG scheme

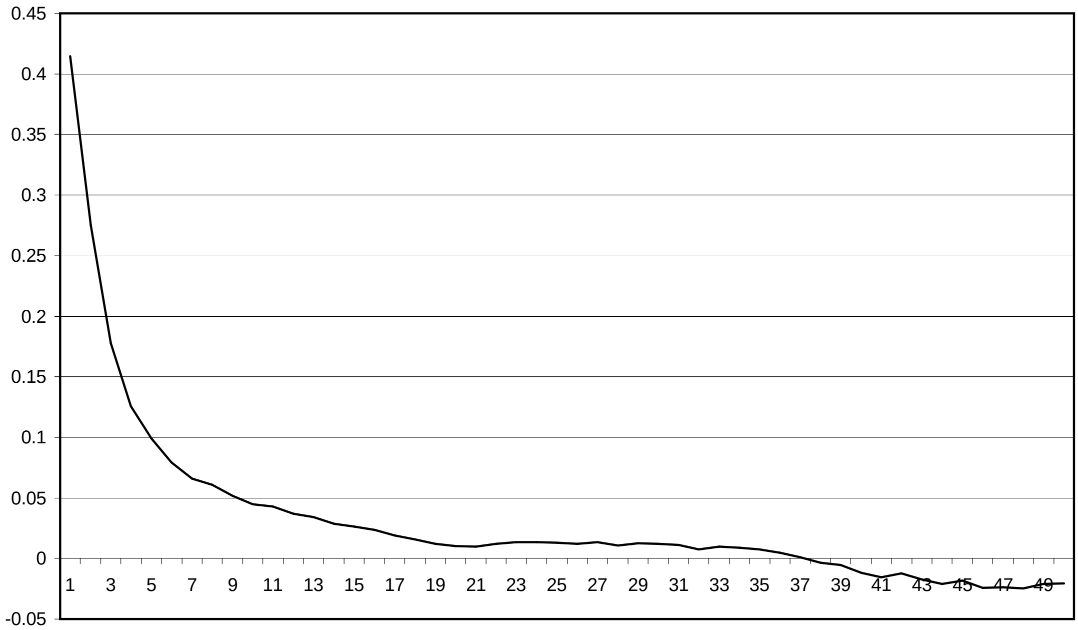

5 Conclusion

This paper takes stock of the available empirical information about the size of the main risk factor associated with the return to social security, viz. demographic risk (fertility, mortality, migration), productivity risk, and asset return risk (inflation, bond returns, and equity return).

In terms of future consumption levels, productivity is the most important risk factor. As a rule-of-thumb, the risk increases by about 2% per year. As a result, the benefit level risk in a defined-contribution PAYG scheme is of the same order of magnitude as the return risk on mid-term bonds. There is *some* evidence of a deceleration of the rate of growth of productivity risk over longer horizons, but this mainly applies to Anglo-Saxon countries. Annualised variances for per capita GNP decline only slowly for the U.S., and appear to increase over time for continental Europe.

The risk involved in the development of the old-age dependency ratio is limited. The dependency ratio will change substantially in the next 70 years, but the larger part of this fluctuation is predictable, so that, by definition, it does not constitute a risk factor. It follows that the effect of demographic risk on PAYG pension benefits is second-order over the relevant planning period.

Asset return risk declines with the length of the investment period as a result of mean reversion, but the effect is limited and extends only for about 5-6 years. Exploiting mean reversion in a pension fund portfolio strategy would seem to be problematical, because the effect appears to be mainly due to persistence in volatility. Expected returns are above average if volatility is above average. Before the latest financial crisis, there was some agreement in the literature that the equity premium should be in the range 3-4% (annualised returns over long horizons). Current estimates from term structure models are about one percent higher, in line with the increase in volatility of stock markets. Given a current U.S. long-term real risk-free rate of two percent, the market appears to expect U.S. long-term annualised real equity returns of about 6-7%, a bit below the returns over the past century. Returns in Western Europe may deviate, depending on the development of the real exchange rate.

Demographic risk spills over to financial markets. Most of the literature argues that ageing lowers rates of return. The consensus estimate is 1%-point. However, if tax burdens increase steeply, a fall in rates of return need not happen. There is also evidence of a relation between migration and differential TFP growth. Regions with relatively strong TFP growth receive a net inflow of young workers, which lowers the dependency ratio and thereby demographic risk.

GDP disasters like wars and financial crises are infrequent but have a strong impact on the lower tail of the productivity risk distribution. As a result, a zero-growth scenario for the next century cannot be completely excluded. In a large majority of cases, GDP disasters also have a strong negative impact on asset returns. As a result, equity returns possess a fat lower tail which makes equity much more risky than what would follow from lognormally distributed returns.

Possible extensions of the present analysis are

1. The present analysis only considers fundamental risk factors. It does not discuss how these risk factors are distributed over generations. It would be informative to investigate the effect of different institutional arrangements on the distribution of risk over generations;
2. The VAR risk model is only partly data-based. The empirical content of the VAR model may be improved by estimating a complete system;
3. The asset return model is based on U.S. data solely. It would be useful to include data for other regions as well, to obtain a better estimate of portfolio risk;
4. The productivity analysis for the Netherlands is based only on post-war data. This is too short for reliable conclusions about unit roots. A useful extension of the analysis would involve collecting long historic series on GDP and employment;
5. It would be useful to construct consumption-wealth ratios for the Netherlands, or the Euro area, to verify the usefulness of this statistic for predicting excess returns and/or investment in productive capital.

Appendices

Appendix A Fertility

Compared to the Lee-Tuljapurkar model described in Section 2.1, the present specification models fertility by cohort, instead of by age, and uses *two* states, one describing the total fertility of a cohort, and the other one the shift in the age distribution of fertility for that cohort. The specification is

$$\ln \phi_{t+1,\tau} = \alpha_{\phi,\tau} + \beta_{\phi,\tau}\, z_{1,t} + \gamma_{\phi,\tau}\, z_{2,t} + \varepsilon_{\phi,t+1,\tau} \tag{A.1a}$$

$$z_{i,t+1} = \sum_{j=1}^{2} \delta_{ij,1}\, z_{j,t} + \delta_{i,0} + \eta_{z_i,t+1} \qquad i \in \{1,2\} \tag{A.1b}$$

The $z_{i,t}$ are the hidden state variables of the cohort with birth year t, and the $\phi_{t,\tau}$ are the observed realizations of the fertility rate of this cohort at age τ. The $\varepsilon_{t,\tau}$ are i.i.d. random variables with mean zero. Fertility of age group τ at time t is then $\phi_{t-\tau,\tau}$ and total fertility in period t is given by $\sum_{\tau=0}^{\infty} \phi_{t-\tau,\tau}$. The restrictions on the coefficients are $\sum_\tau \beta_{\phi,\tau} = 1$ and $\sum_\tau \gamma_{\phi,\tau} = 0$. This allows us to interpret z_1 as the completed fertility state and z_2 as the fertility timing state. From the estimates of the state variables in Figure A.1, it appears that completed fertility was already steeply declining from the cohorts born from the mid-thirties on. Initially this was masked by a shift towards early childbirth, but starting with the post-war cohort, fertility also shifted to later ages.

The coefficients of A.1 are given in Tables A.1 and A.2. The coefficient $\delta_{22,1}$ has been fixed at 0.98 to avoid an estimate slightly above one. Both state variables move to a long-run value. The fertility coefficients in Table A.2 show that peak fertility for the 1935 cohort is estimated to have been around age 28. The shift in the state variables z_1 and z_2 then gradually pushed fertility to a lower level overall and with peak fertility at a higher age.

Table A.1 Coefficients of the fertility state equations (A.1b)

	$\delta_{i,0}$	$\delta_{i1,1}$	$\delta_{i2,1}$
state			
1	– 0.4735	0.9524	0
2	– 0.0551	0	0.9800

Figure A.1 Fertility states for Dutch women by year of birth

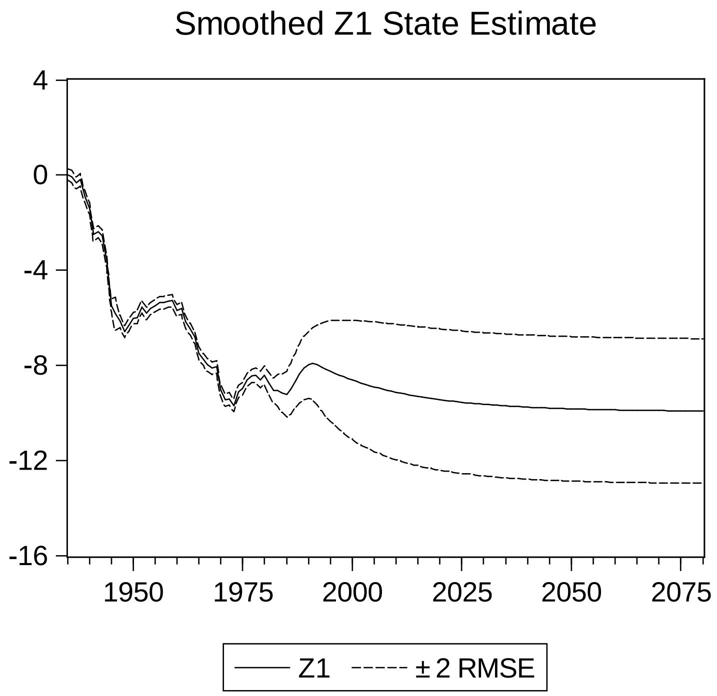

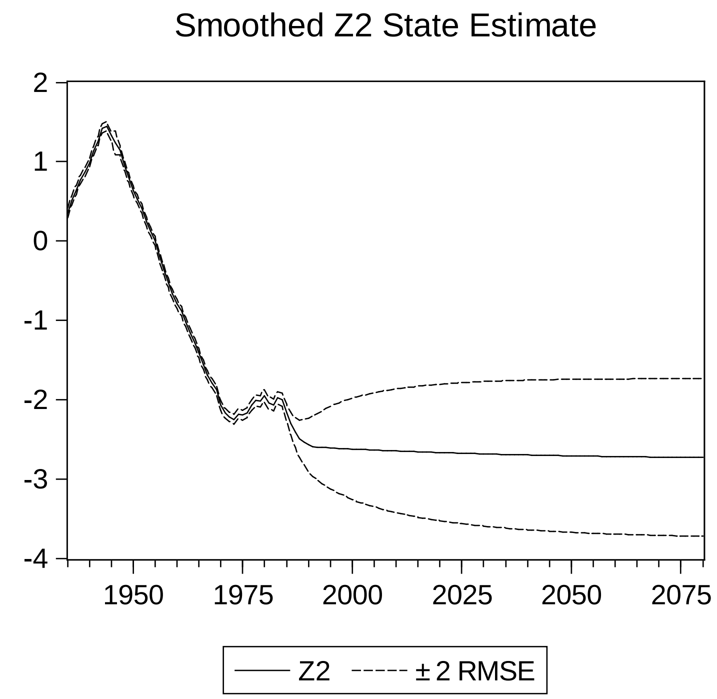

Table A.2 **Coefficients of the fertility observation equations (A.1a)**

τ	$\alpha_{\phi,\tau}$	$\beta_{\phi,\tau}$	$\gamma_{\phi,\tau}$
15	– 8.1990	– 0.1097	0.2450
16	– 6.8980	– 0.1191	0.3533
17	– 5.5221	– 0.0926	0.4436
18	– 4.5526	– 0.0662	0.4693
19	– 3.8573	– 0.0450	0.4412
20	– 3.3422	– 0.0276	0.4191
21	– 2.9086	– 0.0127	0.4157
22	– 2.5151	0.0045	0.4032
23	– 2.1639	0.0259	0.3666
24	– 1.8977	0.0431	0.3144
25	– 1.7091	0.0558	0.2463
26	– 1.5823	0.0660	0.1677
27	– 1.5384	0.0698	0.0859
28	– 1.5383	0.0736	– 0.0047
29	– 1.5803	0.0760	– 0.0920
30	– 1.6610	0.0765	– 0.1736
31	– 1.7841	0.0771	– 0.2510
32	– 1.9470	0.0759	– 0.3161
33	– 2.1429	0.0718	– 0.3633
34	– 2.3661	0.0671	– 0.4064
35	– 2.6616	0.0512	– 0.4246
36	– 3.0279	0.0272	– 0.4178
37	– 3.3567	0.0153	– 0.4222
38	– 3.7498	– 0.0056	– 0.3999
39	– 4.1244	– 0.0197	– 0.3779
40	– 4.4690	– 0.0203	– 0.3819
41	– 4.9046	– 0.0271	– 0.3506
42	– 5.2996	– 0.0117	– 0.3436
43	– 5.8252	– 0.0138	– 0.2625
44	– 6.3155	0.0073	– 0.2147
45	– 6.8597	0.0297	– 0.1413
46	– 7.3980	0.0660	– 0.0625
47	– 7.7656	0.1595	0.1358
48	– 8.1831	0.1926	0.1434
49	– 7.8073	0.2388	0.7564

Appendix B Migration

The migration model describes *net* migration to the Netherlands by age, using a two-stage approach. Total migration relative to the Dutch population size is assumed to be labour market driven. In times of a strong demand for labour, migration into the Netherlands is relatively high. The cyclical indicator used in the VAR model is labour productivity. Migration per age group is assumed to be a fixed fraction of total migration. The specification is

$$\frac{m_t}{n_t} = \alpha_0 + \alpha_1 \frac{m_{t-1}}{n_{t-1}} + \sum_{i=0}^{2} \alpha_{i+2} \ln \zeta_{t-i} + \varepsilon_t \tag{B.1a}$$

$$\frac{m_{t,i}}{n_t} = \alpha_{0,i} + \alpha_{1,i} \ln t + \alpha_{2,i} \frac{m_{t-1,i}}{n_{t-1}} + \alpha_{3,i} \frac{m_t}{n_t} + \alpha_{4,i} \frac{m_{t-1}}{n_{t-1}} + \varepsilon_{t,i} \tag{B.1b}$$

where m is migration, n denotes population size, and ζ is labour productivity. The coefficients of eqs. (B.1) are given in tables B.1-B.2. Table B.1 shows that total net migration into the Netherlands depends on current and lagged productivity and population size. The long-run elasticity of productivity is 0.1%, virtually negligible. On average, migration is about 0.3% of the population, without a clear trend.

Table B.1 Coefficients of the aggregate migration equation (B.1a)

	α_0	α_1	α_2	α_3	α_4
value	– 0.0004	0.6020	– 0.0223	0.0649	– 0.0422
standard dev.	(0.004)	(0.129)	(0.013)	(0.022)	(0.012)

Table B.2 gives the coefficients of the migration equation per age group. For young age groups, immigration depends on total migration, as a representation of family migration. For older age groups, this effects weakens somewhat, and total migration can be seen to represent labour market conditions.[56] The time trend shows a shift in the age composition over time towards persons in their late twenties. It has been specified as a logarithmic trend to weaken its effect in the long run.

Table B.2: Coefficients of the age-specific migration equations (B.1b)

τ	$\alpha_{0,\tau}$	$\alpha_{1,\tau}$	$\alpha_{2,\tau}$	$\alpha_{3,\tau}$	$\alpha_{4,\tau}$
1	0.000051	– 0.000012	0.718482	0.008824	– 0.006731
2	0.000108	– 0.000030	0.674855	0.018685	– 0.011424

continued on next page ...

[56] It appeared infeasible to estimate labour productivity effects for each age group separately. The coefficients vary substantially over age groups and are frequently insignificant.

Table B.2: Coefficients of the age-specific migration equations (continued)

τ	$\alpha_{0,\tau}$	$\alpha_{1,\tau}$	$\alpha_{2,\tau}$	$\alpha_{3,\tau}$	$\alpha_{4,\tau}$
3	0.000147	– 0.000040	0.497530	0.020314	– 0.010906
4	0.000164	– 0.000045	0.533240	0.019257	– 0.010875
5	0.000147	– 0.000041	0.618307	0.021188	– 0.013314
6	0.000133	– 0.000037	0.643473	0.019752	– 0.012988
7	0.000113	– 0.000033	0.590686	0.020756	– 0.011819
8	0.000082	– 0.000025	0.599613	0.019639	– 0.011309
9	0.000089	– 0.000026	0.588728	0.019130	– 0.011787
10	0.000101	– 0.000029	0.535633	0.018156	– 0.009262
11	0.000107	– 0.000030	0.556901	0.018676	– 0.010355
12	0.000085	– 0.000025	0.623136	0.019705	– 0.011897
13	0.000069	– 0.000020	0.684998	0.017228	– 0.010925
14	0.000100	– 0.000028	0.522190	0.016285	– 0.008043
15	0.000106	– 0.000029	0.445143	0.016572	– 0.007342
16	0.000090	– 0.000025	0.516049	0.018953	– 0.008937
17	0.000050	– 0.000014	0.661339	0.021658	– 0.012580
18	0.000081	– 0.000022	0.569849	0.022065	– 0.009642
19	0.000038	– 0.000009	0.626570	0.024846	– 0.013295
20	– 0.000007	0.000006	0.731088	0.027500	– 0.020590
21	– 0.000010	0.000007	0.753872	0.028713	– 0.023184
22	– 0.000063	0.000021	0.770555	0.027437	– 0.022165
23	– 0.000080	0.000026	0.752174	0.027055	– 0.020909
24	– 0.000098	0.000031	0.747848	0.031915	– 0.024665
25	– 0.000120	0.000038	0.664675	0.029589	– 0.019616
26	– 0.000153	0.000048	0.575412	0.033578	– 0.021793
27	– 0.000184	0.000055	0.556147	0.030250	– 0.017996
28	– 0.000342	0.000097	0.201720	0.030343	– 0.007057
29	– 0.000256	0.000072	0.371479	0.026974	– 0.010216
30	– 0.000137	0.000039	0.656343	0.024754	– 0.016787
31	– 0.000190	0.000052	0.433177	0.025332	– 0.012144
32	– 0.000151	0.000042	0.499215	0.022246	– 0.011792
33	– 0.000083	0.000024	0.626541	0.017895	– 0.011883

continued on next page ...

Table B.2: Coefficients of the age-specific migration equations (continued)

τ	$\alpha_{0,\tau}$	$\alpha_{1,\tau}$	$\alpha_{2,\tau}$	$\alpha_{3,\tau}$	$\alpha_{4,\tau}$
34	– 0.000083	0.000022	0.598750	0.018739	– 0.011346
35	– 0.000038	0.000011	0.765083	0.015612	– 0.012898
36	– 0.000021	0.000006	0.824792	0.016363	– 0.013831
37	– 0.000055	0.000013	0.571591	0.015776	– 0.008465
38	– 0.000023	0.000005	0.670542	0.013985	– 0.008791
39	– 0.000027	0.000005	0.604730	0.012482	– 0.006077
40	– 0.000017	0.000003	0.643921	0.010568	– 0.005937
41	– 0.000022	0.000004	0.658494	0.011532	– 0.006820
42	– 0.000028	0.000005	0.485556	0.010824	– 0.004455
43	– 0.000022	0.000003	0.391518	0.010379	– 0.003533
44	– 0.000025	0.000004	0.245904	0.009770	– 0.002555
45	0.000000	– 0.000002	0.579455	0.008677	– 0.004720
46	0.000000	– 0.000003	0.288896	0.008438	– 0.002806
47	0.000008	– 0.000005	0.319706	0.007477	– 0.002272
48	0.000006	– 0.000004	0.323070	0.006724	– 0.001995
49	– 0.000002	– 0.000001	0.427341	0.005640	– 0.002669
50	0.000017	– 0.000006	0.434228	0.005468	– 0.002572
51	0.000021	– 0.000008	0.248798	0.005944	– 0.002732
52	0.000012	– 0.000004	0.651969	0.005203	– 0.004005
53	0.000021	– 0.000008	0.460269	0.005031	– 0.002872
54	0.000031	– 0.000010	0.422688	0.003833	– 0.001816
55	0.000025	– 0.000009	0.432829	0.004368	– 0.002106
56	0.000026	– 0.000009	0.402804	0.003989	– 0.001840
57	0.000025	– 0.000008	0.510373	0.003198	– 0.001809
58	0.000022	– 0.000008	0.522649	0.003394	– 0.001855
59	0.000034	– 0.000011	0.389475	0.002431	– 0.000586
60	0.000043	– 0.000014	0.288126	0.002957	– 0.000915
61	0.000028	– 0.000009	0.536255	0.002218	– 0.000913
62	0.000024	– 0.000008	0.509892	0.002756	– 0.001204
63	0.000010	– 0.000004	0.677907	0.002266	– 0.001455
64	0.000010	– 0.000004	0.495546	0.002125	– 0.000556

continued on next page ...

Table B.2: Coefficients of the age-specific migration equations (continued)

τ	$\alpha_{0,\tau}$	$\alpha_{1,\tau}$	$\alpha_{2,\tau}$	$\alpha_{3,\tau}$	$\alpha_{4,\tau}$
65	0.000007	– 0.000002	0.564329	0.002001	– 0.001220
66	0.000020	– 0.000007	0.706504	0.002145	– 0.001199
67	0.000014	– 0.000005	0.721139	0.001874	– 0.001203
68	0.000008	– 0.000003	0.495385	0.001200	– 0.000325
69	0.000018	– 0.000006	0.432182	0.001628	– 0.000614
70	0.000008	– 0.000002	0.615707	0.000986	– 0.000466
71	0.000005	– 0.000002	0.571500	0.001343	– 0.000767
72	0.000002	– 0.000001	0.436022	0.001260	– 0.000635
73	0.000000	0.000000	0.509395	0.001120	– 0.000713
74	0.000003	– 0.000001	0.339115	0.000824	– 0.000438
75	– 0.000001	0.000000	0.628249	0.000603	– 0.000499
76	– 0.000002	0.000000	0.114440	0.000836	– 0.000182
77	– 0.000001	0.000000	0.477442	0.000364	– 0.000188
78	– 0.000003	0.000001	0.163725	0.000562	– 0.000097
79	– 0.000005	0.000001	0.338565	0.000119	0.000150
80	– 0.000003	0.000001	0.453816	0.000392	– 0.000202
81	– 0.000004	0.000001	0.195005	0.000356	– 0.000192
82	– 0.000004	0.000001	0.392312	0.000111	0.000037
83	– 0.000005	0.000001	0.125671	0.000225	0.000036
84	– 0.000003	0.000001	0.407340	0.000100	– 0.000020
85	– 0.000004	0.000001	– 0.118040	0.000031	0.000175
86	– 0.000003	0.000001	0.052530	0.000089	0.000075
87	– 0.000001	0.000000	0.089676	0.000024	0.000025
88	– 0.000001	0.000000	0.111058	0.000031	– 0.000012
89	– 0.000002	0.000001	– 0.159731	0.000050	0.000011
90	– 0.000001	0.000000	– 0.162866	0.000065	0.000023
91	0.000001	0.000000	– 0.222624	0.000095	– 0.000013
92	0.000000	0.000000	– 0.104047	0.000012	0.000049
93	0.000000	0.000000	0.124329	0.000031	– 0.000005
94	– 0.000001	0.000000	– 0.199483	– 0.000022	0.000018
95	0.000000	0.000000	– 0.181421	0.000015	– 0.000003

continued on next page ...

Table B.2: Coefficients of the age-specific migration equations (continued)

τ	$\alpha_{0,\tau}$	$\alpha_{1,\tau}$	$\alpha_{2,\tau}$	$\alpha_{3,\tau}$	$\alpha_{4,\tau}$
96	0.000000	0.000000	− 0.189988	− 0.000010	− 0.000001
97	0.000000	0.000000	− 0.294416	0.000028	− 0.000028
98	0.000000	0.000000	− 0.057140	0.000001	0.000006
99	0.000000	0.000000	− 0.263704	0.000005	0.000000
100	0.000000	0.000000	− 0.143369	− 0.000005	0.000005

Appendix C Productivity Growth

The productivity equation describes GDP per hour for the Netherlands, as given in the databank of The Conference Board (2009). Figure C.1 displays hourly productivity growth over the past

Figure C.1 Growth rates of GDP per hour in the Netherlands

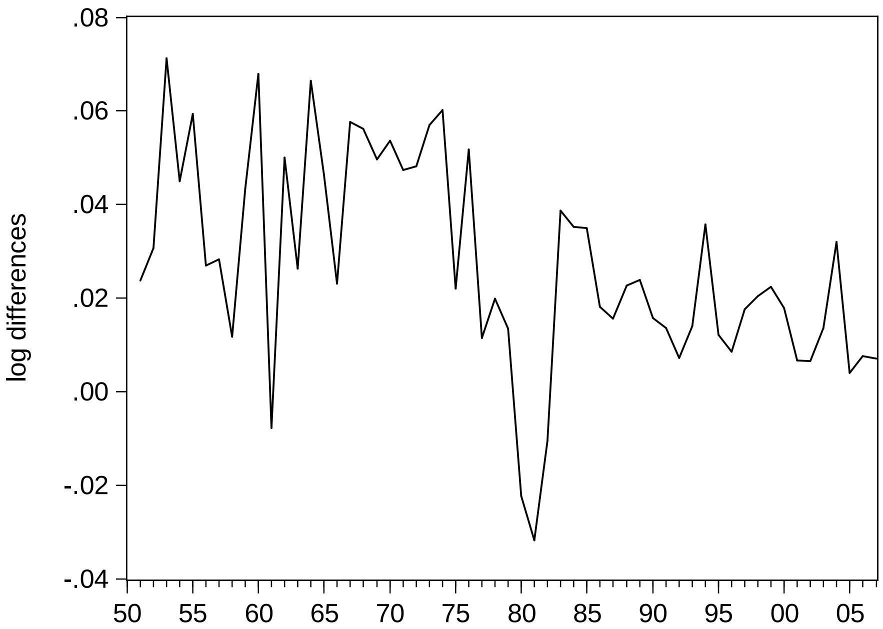

decades. The slowdown in productivity growth from the mid-seventies is apparent. Consequently, trend stationarity of labour productivity is firmly rejected, using a standard Dickey-Fuller test, while even difference stationarity is in doubt, with a marginal significance level of 5.6%. I present two different specifications, the first being a straightforward AR(2) formulation, and the second a state space formulation with a hidden state variable. The AR(2) specification gives the following results (standard errors in parentheses)

$$\Delta \ln \frac{y_t}{L_{h_t}} = \underset{(1.37)}{0.3689} - \underset{(0.02)}{0.025} \left(\ln \frac{y_{t-1}}{L_{h_{t-1}}} - \underset{(0.02)}{0.0057}\, t \right) + \underset{(0.13)}{0.276} \Delta \ln \frac{y_{t-1}}{L_{h_{t-1}}} + \varepsilon_t \tag{C.1}$$

$$\hat{\sigma}_\varepsilon = 0.019; \qquad \text{D-W} = 2.11$$

There is a considerable probability that the equation possesses a unit root.[57] As a result, confidence intervals for forecasted log productivity become unbounded in the long run. Figure

[57] The marginal significance level under the null hypothesis of a unit root is 92%.

Figure C.2 Forecasts for Dutch log productivity per hour, using equation (C.1)

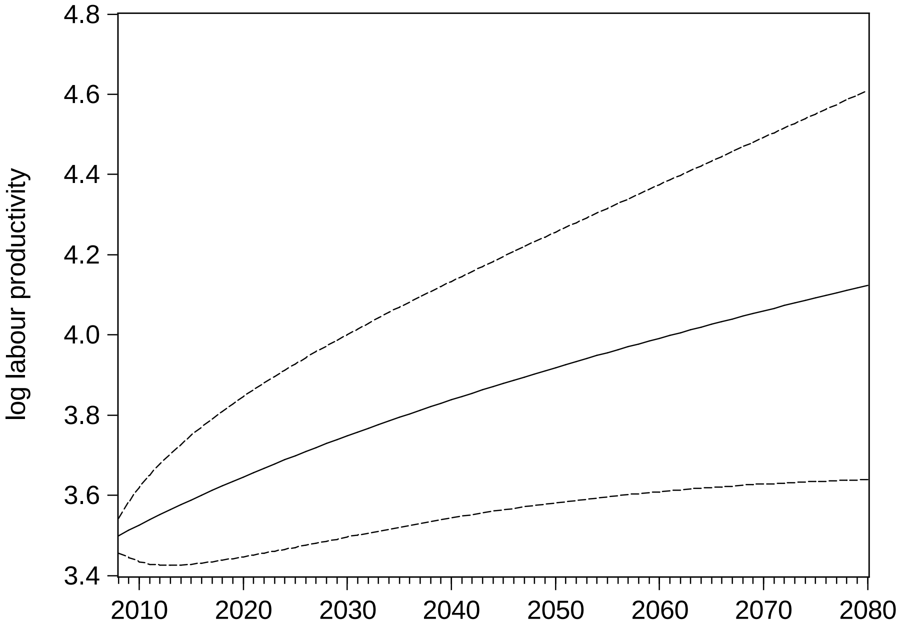

Figure C.3 Difference of log hourly productivity in the U.S. and the Netherlands

C.2 presents the forecast intervals of equation (C.1), including parameter uncertainty, using standard distribution theory. [58] A striking feature of equation (C.1) is the low estimated value of annual productivity growth (0.6%), substantially below observed growth rates in the past decade. The model in fact extrapolates the slowdown in productivity growth that started in the eighties. This result is more or less implied by the assumed trend-stationarity of productivity, which leaves little room in the specification for a recovery of productivity growth. [59]

A possible explanation for the slowdown in productivity growth is the observation that the gap with U.S. productivity had been substantially closed by the early seventies, so that catching-up disappeared as a source of productivity growth. Figure C.3 shows the development of the productivity gap between the U.S. and the Netherlands. Inclusion of the productivity gap in (C.1) reduces the probability of a unit root in Dutch productivity, but creates an implausibly large long-run effect of U.S. productivity on Dutch productivity, with an elasticity of 2.5.

As labour productivity is probably non-stationary anyway (see Section 2.2.1), it may be better to assume nonstationarity from the outset. The model specified below assumes that productivity growth may vary over time. $\ln y/L_h$ depends on a hidden state variable ζ_t, which grows at a rate μ_ζ. μ_ζ in turn follows an AR(1) process. Because catching up may have contributed to productivity growth in the fifties and sixties, the model uses the gap between *structural* U.S. and Dutch productivity as an explanatory variable for structural Dutch productivity growth. This implies that U.S. productivity growth is a risk factor for Dutch productivity growth. Hence U.S. productivity is also modelled, using the same type of model.[60]

$$\ln y_t^i/L_{h_t}^i = \zeta_t^i + \varepsilon_t^i \qquad i \in \{NL, US\} \tag{C.4a}$$

$$\zeta_t^i = \zeta_{t-1}^i + \mu_{\zeta_{t-1}}^i \qquad i \in \{NL, US\} \tag{C.4b}$$

$$\mu_{\zeta_t}^{NL} = \mu_{\zeta_{t-1}}^{NL} - \rho_{NL}\left(\mu_{\zeta_{t-1}}^{NL} - \psi_{NL} - \omega\left(\zeta_{t-1}^{US} - \zeta_{t-1}^{NL}\right)\right) + \eta_t^{NL} \tag{C.4c}$$

$$\mu_{\zeta_t}^{US} = \mu_{\zeta_{t-1}}^{US} - \rho_{US}\left(\mu_{\zeta_{t-1}}^{US} - \psi_{US}\right) + \eta_t^{US} \tag{C.4d}$$

For both countries, the point of this specification is that productivity growth is stationary but can deviate from its long-run average ψ for substantial periods of time. The *level* of productivity on the other hand is not stationary, but *may* change permanently, which it does in response to

[58] These intervals are therefore to be regarded as lower bounds of the correct forecast intervals.

[59] In addition to the apparent lack of mean reversion in productivity levels, there is also substantial autocorrelation in the variance of productivity growth. Re-estimating (C.1) using a GARCH(1,2) specification for the error term gives

$$\ln \frac{y_t}{L_{h_t}} = \underset{(1.0)}{-0.202} - \underset{(0.0005)}{0.00013}\,t + \underset{(0.11)}{1.362} \ln \frac{y_{t-1}}{L_{h_{t-1}}} - \underset{(0.11)}{0.374} \ln \frac{y_{t-2}}{L_{h_{t-2}}} + \varepsilon_t \tag{C.2}$$

$$\hat{\sigma}_\varepsilon^2 = \underset{(0.000012)}{0.000013} + \underset{(0.11)}{0.331}\,\hat{\varepsilon}_{t-1}^2 - \underset{(0.051)}{0.192}\,\hat{\sigma}_{\varepsilon,t-1}^2 + \underset{(0.06)}{0.736}\,\hat{\sigma}_{\varepsilon,t-2}^2 \tag{C.3}$$

This equation has a long run expected error standard deviation of 1.0%, substantially smaller than (C.1). However, the estimated rate of technical progress in (C.2) is – 1.0%.

[60] Over the post-war period, unit root tests strongly suggest the presence of a unit root in U.S. productivity

random shocks in η. Shocks in ε on the other hand have only a temporary effect on productivity. In addition, there is a catching-up effect for Dutch productivity, the size of which is determined by ω. This implies that, in the long run, Dutch productivity growth is the same as U.S. productivity growth, but the level may differ.[61]

The model (C.4) has been estimated using a Kalman filter. Estimation results are presented in Table C.1. For the U.S., the long-run rate of technical progress is $\psi_{US} = 1.9\%$, and for the

Table C.1 Estimation Results for Equation (C.4)

	value	standard error
parameter		
ρ_{NL}	0.5775	0.189
ρ_{US}	0.2837	0.172
ψ_{NL}	0.0118	0.006
ψ_{US}	0.0188	0.002
ω	0.0792	0.021
$\hat{\sigma}_{\varepsilon}^{NL}$	0.0061	0.025
$\hat{\sigma}_{\varepsilon}^{US}$	0.0053	0.016
$\hat{\sigma}_{\eta}^{NL}$	0.0144	0.029
$\hat{\sigma}_{\eta}^{US}$	0.0053	0.021

Netherlands it is 1.2%, substantially lower. However, the model implies that, due to catching-up, the Dutch rate converges to the U.S. rate, while the level of labour productivity in the Netherlands will on average lag behind by about 9%.[62] The rate of technology diffusion from the U.S. is estimated to be $\omega = 8\%$ per year The estimates of σ_{ε} and σ_{η} suggest that for the Netherlands shocks in structural labour productivity growth (η) contribute more to variations in labour productivity than variations in labour productivity over the business cycle (ε). For the U.S., both type of shocks are of equal magnitude. These results are in good agreement with the variance ratio diagrams discussed in Section 2.2.1.

In addition to the parameters presented in Table C.1, the hidden state variables are also part of the estimation results. Figure C.4 presents the estimated (smoothed) values of the hidden state variables,[63] and their forecasts till 2080. Because of the small error variance in the signal equation (C.4a), the estimates of ζ_t have narrow confidence bands over the sample period. In addition, forecast confidence intervals for ζ and $\ln y/L_h$ virtually coincide. The productivity growth state variables show considerable variation, in particular for the Netherlands. Because of

[61] This catching-up effect only makes sense if Dutch productivity levels are below U.S. levels, which is true for the sample period.

[62] The expected long-run difference in productivity equals $(\psi_{US} - \psi_{NL})/\omega \approx 9\%$

[63] Smoothed estimates are computed using the information in the *full* sample, so they use more information than available for prediction at the date of the state variable.

Figure C.4 Estimated state variables of the productivity equations (C.4)

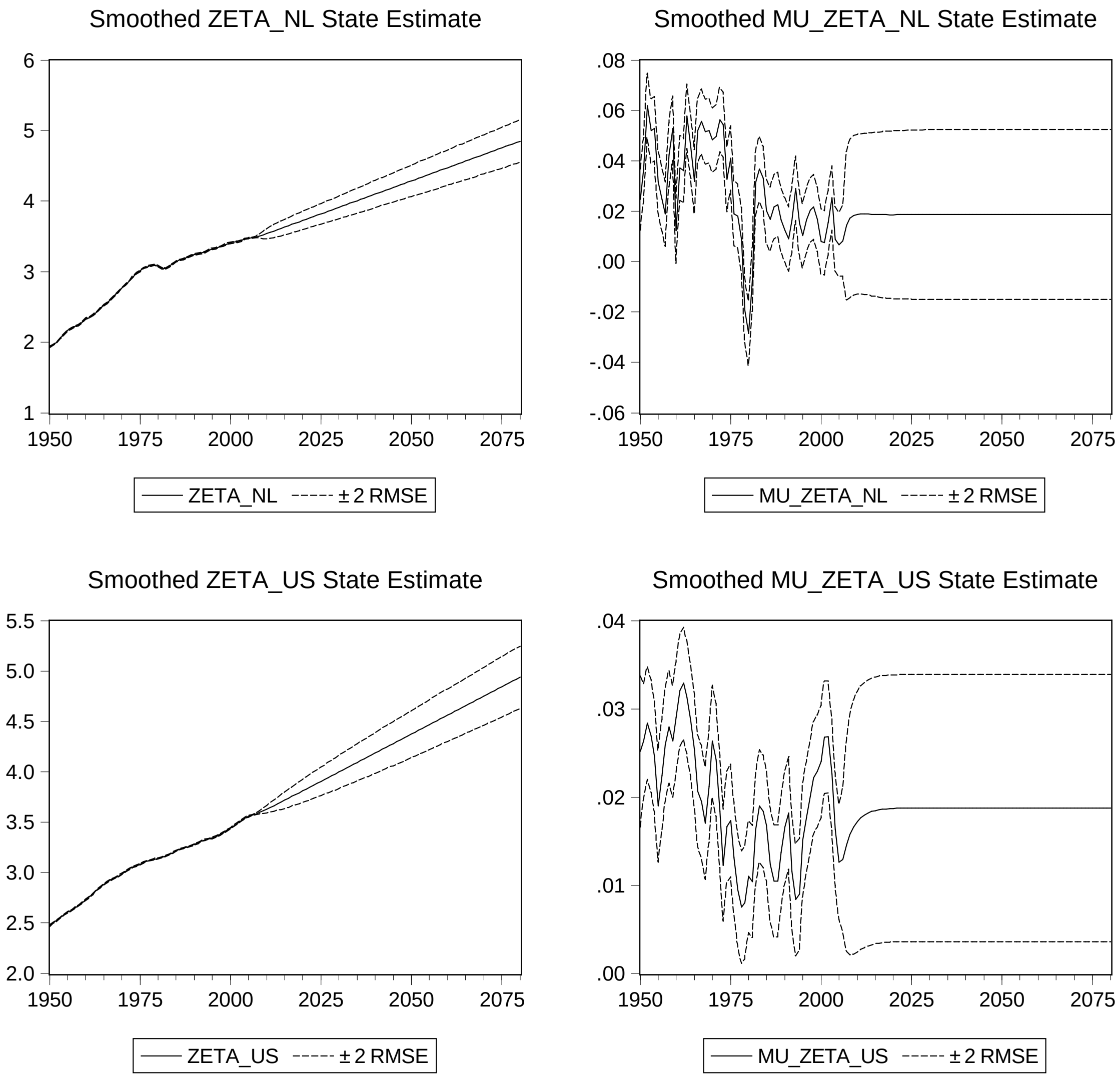

the uncertainty in the growth rate, the forecast for the level of productivity is fairly uncertain over long horizons. Over a 75-year horizon, the 95% confidence interval is about 50% of the expected level. However, *in expectation*, both growth rates converge to 1.9%, the long-run estimate for the U.S.

Forecasting of a state space formulation is done via the filter gain matrix. Let the general form of the state space model be as

$$\mathbf{y}_t = \mathbf{Z}\mathbf{x}_t + \boldsymbol{\varepsilon}_t \tag{C.5a}$$

$$\mathbf{x}_{t+1} = \mathbf{A}\mathbf{x}_t + \boldsymbol{\eta}_{t+1} \tag{C.5b}$$

Forecasts are given as $\hat{y}_t = \mathbf{Z}\hat{\mathbf{x}}_t$ where $\hat{x}$ is the estimate of the state vector. The state is updated as

$$\hat{\mathbf{x}}_{t+1} = \mathbf{A}\,\hat{\mathbf{x}}_t + \mathbf{G}\,(\mathbf{y}_t - \hat{\mathbf{y}}_t) \tag{C.6a}$$

$$\mathbf{G} = \mathbf{P}\,\mathbf{Z}'\,\left(\mathbf{Z}\mathbf{P}\mathbf{Z}' + \mathbf{\Omega}\right)^{-1} \tag{C.6b}$$

The coefficients of the filter gain matrix $\mathbf{G}$ are presented in Table C.2. A 1% forecast error in

Table C.2 Filter gain coefficients

	$\hat{\varepsilon}^{NL}$	$\hat{\varepsilon}^{US}$
ζ^{NL}	1.192	0.016
μ_{ζ}^{NL}	0.104	0.043
ζ^{US}	– 0.0004	1.027
μ_{ζ}^{US}	– 0.0005	0.217

Dutch productivity leads to an upward revision of structural (unobserved) productivity of 1.2% and an upward revision of Dutch structural productivity growth of 0.1%. For the U.S., the revision of the productivity level is approximately one-to-one at 1.02%, whereas estimated structural productivity growth goes up by 0.2%. U.S. productivity innovations also carry information about future Dutch productivity, as indicated by the cross-effect of U.S. forecast errors on Dutch productivity estimates. After a shock, productivity growth gradually returns to its long-run level of 1.9%. The permanent effect of a shock in productivity levels is therefore larger than the immediate effect, and given by $1.02 + 0.217/(1-\rho^{US}) = 1.32\%$.

Appendix D Dividends, Stock Prices, and Equity Return

This appendix discusses the relation between dividends, stock prices, and stock returns (see also Campbell et al. (1997), Ch. 7). The return to equity is composed of dividends and capital gains, in the form of increases in share prices. At the start of period t, the firm has $n_{v_{t-1}}$ shares outstanding. The market price per share is denoted p_{v_t}, and the market value of the firm is $V_t = p_{v_t} n_{v_{t-1}}$. The firm then issues $n_{v_t} - n_{v_{t-1}}$ new shares.[64] These n_{v_t} shares are traded *cum dividend*, i.e. with the dividend falling to the buyer.[65] The return r_k to equity n_{v_t} is therefore given by

$$1 + r_{k_{t+1}} = \frac{p_{v_{t+1}}}{p_{v_t} - d_t} \qquad \Leftrightarrow \tag{D.1a}$$

$$1 + r_{k_{t+1}} = \frac{V_{t+1}}{V_t - D_t - VN_t} \tag{D.1b}$$

where D_t denotes dividend payments, $d_t = D_t / n_{v_t}$ is dividends per share, and $VN_t = p_{v_t} \left(n_{v_t} - n_{v_{t-1}}\right)$ denotes the value of new share issues by the firm. For simplicity of notation, I assume that dividend payments are not taxed. Generally, $r_{k_{t+1}}$ is a random variable at time t.

From (D.1a), it follows that

$$\begin{aligned} \ln 1 + r_{k_{t+1}} &= -\ln\left(\frac{p_{v_t}}{p_{v_{t+1}}}\left(1 - \frac{d_t}{p_{v_t}}\right)\right) \\ &\approx \ln \frac{p_{v_{t+1}}}{p_{v_t}} + \frac{d_t}{p_{v_t}} \end{aligned}$$

The log return to equity consists of capital gains and dividend return. However, with perfect and complete markets, the only relevant information is the distribution of r_k. The division between dividends and retained earnings is irrelevant. The relevant market equilibrium condition is the alignment of returns with the stochastic discount factor of investors, m_t

$$\mathrm{E}_t\left[m_{t+1}\left(1 + r_{k_{t+1}}\right)\right] = 1 \tag{D.2}$$

Using (D.1a), we obtain from (D.2) $p_{v_t} = d_t + \mathrm{E}_t\left[m_{t+1} p_{v_{t+1}}\right]$. Expanding the recursion yields

$$\frac{p_{v_t}}{d_t} = 1 + \mathrm{E}_t\left[\sum_{i=1}^{\infty} \prod_{j=1}^{i} m_{t+j} \frac{d_{t+j}}{d_{t+j-1}}\right] \tag{D.3}$$

The price-dividend ratio is seen to depend on future discount factors m_{t+j} and growth rates of dividends d_{t+j}/d_{t+j-1}.

[64] The number of new shares issued is known at the start of period t, when the price p_{v_t} of shares is determined.

[65] Alternatively, if trades are ex dividend, the original owner decides about production and investment in the current period. This must imply that trade occurs at the end of the period, rather than the beginning, as is assumed here.

Suppose the discount factor is based on the marginal utility of consumption

$$m_{t+1} = \left(\frac{c_{t+1}}{c_t}\right)^{-\gamma} \mathrm{e}^{-\rho}$$

and that both log consumption and log dividends follow a random walk (Hall (1978), Cochrane (1994))

$$\ln c_{t+1} = \ln c_t + \psi + \varepsilon_{t+1}$$
$$\ln d_{t+1} = \ln d_t + \delta + \eta_{t+1}$$

Then

$$\begin{aligned}\frac{p_{v_t}}{d_t} &= 1 + \mathrm{E}_t\left[\sum_{i=1}^{\infty} \exp\left[\sum_{j=1}^{i} -\gamma\psi - \rho - \gamma\varepsilon_{t+j} + \delta + \eta_{t+j}\right]\right] \\ &= 1 + \sum_{i=1}^{\infty} \exp\left[-(\gamma\psi + \rho - \delta)\, i + \tfrac{1}{2} i \left(\gamma^2\sigma_\varepsilon^2 + \sigma_\eta^2\right)\right]\end{aligned}$$

It follows that

$$\frac{p_{v_t}}{d_t} \approx 1 + \left(\gamma\psi + \rho - \delta - \tfrac{1}{2}\left(\gamma^2\sigma_\varepsilon^2 + \sigma_\eta^2\right)\right)^{-1} \tag{D.4}$$

In this case stock prices move one-to-one with dividends, so that the price-dividend ratio is constant. Stock prices depend strongly on the growth rate and variance of dividends, however.

The issue is whether dividends can be described as a random walk. Thirty years after Black (1976a) first raised the issue, the reasons why firms pay dividends are still not completely understood. One of these reasons may be that, in a world of asymmetric information, dividends are a way for firms to signal their financial viability (see e.g. Baker et al. (2002)). In this respect, a reasonable assumption is that firms set dividends equal to a sustainable fraction of "permanent earnings" (Cochrane, 1994).

Assume that permanent earnings are evaluated at the risk-free rate r_f. That is, in analogy to (D.1b), define permanent earnings W by

$$\frac{\mathrm{E}_t[W_{t+1}]}{1 + r_f} = W_t - D_t - VN_t \tag{D.5}$$

and assume that dividends are determined by the rule[66]

$$D_t = r_f W_t \tag{D.8}$$

It follows from (D.5) and (D.8) that $\mathrm{E}[D_{t+1}] = (1+r_f)(1-r_f)D_t$, or

$$\mathrm{E}_t[D_{t+1}] \approx D_t \tag{D.9}$$

It follows that a shock in earnings leads to a simultaneous change in dividends. Prices adjust immediately in reaction to the change in market value, and the dividend-price ratio remains unchanged. A shock in prices without a corresponding movement in dividends on the other hand must reflect a change in the discount rate (the risk premium). Provided that discount rates are stationary processes, the discount rate must eventually return to its unconditional mean and so must the price of stock. Hence the dividend smoothing model implies that the dividend price ratio can be used to forecast equity returns because, in essence, the dividend price ratio conveys information about the deviation of the discount rate from its unconditional mean. Even so, predictability of stock returns does not imply arbitrage opportunities on the stock market. When the dividend price ratio is high, the price of risk is high as well.

[66] An important aspect of the model is that firms do *not* try to incorporate information about the price of risk in their dividend policy. Suppose alternatively that dividends equal the safe rate of return on the firm's *market value*

$$D_t = r_{f_t} V_t \tag{D.6}$$

Assume that new share issues are zero, $VN_t = 0$, then (D.1b) gives $\mathrm{E}_t\left[m_{t+1}D_{t+1}\right] = \left(1 - r_{f_t}\right) D_t \approx \frac{D_t}{1+r_{f_t}}$, so

$$\mathrm{E}_t\left[m_{t+1}D_{t+1}\right] \approx \mathrm{E}_t\left[m_{t+1}\right] D_t \tag{D.7}$$

It now follows that D_t is a martingale in risk-adjusted probabilities (see Ljungqvist and Sargent (2004), Chapter 13.4), but not in standard probabilities. Indeed, in this case dividends would also reflect changes in the price of risk, and the dividend-price ratio would only vary in response to changes in the risk-free rate.

Appendix E The Campbell-Viceira model

The structure of the Campbell and Viceira (2005) model is as follows. Asset returns are denoted as $(r_{0,t}, x_{1,t}, \ldots, x_{n,t})$, where $r_{0,t}$ is the real return on the benchmark asset and $x_{i,t} = r_{i,t} - r_{0,t}$ is the excess return on the other assets. Let $\mathbf{s}_t \in \mathbb{R}^{m-n-1}$ denote the state vector, then $\mathbf{z}_t' \equiv (r_{0,t}, \mathbf{x}_t, \mathbf{s}_t)$, where $\mathbf{z}_t \in \mathbb{R}^m$. The VAR is specified as

$$\mathbf{z}_{t+1} = \mathbf{\Phi}_0 + \mathbf{\Phi}_1 \mathbf{z}_t + \mathbf{v}_{t+1} \tag{E.1a}$$

$$\mathbf{v}_{t+1} \propto \mathcal{N}(\mathbf{0}, \mathbf{\Sigma}_v) \tag{E.1b}$$

From (E.1a) the conditional variance of $\mathbf{z}$ is $\mathrm{var}(\mathbf{z}_t|\mathbf{z}_0) = \sum_{i=0}^{t-1} \mathbf{\Phi}_1^i \mathbf{\Sigma}_v \mathbf{\Phi}_1^{i\,\prime}$ The long-run variance of $\mathbf{z}$, $\mathbf{\Sigma}_z$, is defined by letting $t \to \infty$ in this expression. The long-run (unconditional) mean and variance of $\mathbf{z}$ are given by[67]

$$\boldsymbol{\mu}_z = (\mathbf{I}_m - \mathbf{\Phi}_1)^{-1}\, \mathbf{\Phi}_0 \tag{E.2a}$$

$$\mathrm{vec}(\mathbf{\Sigma}_z) = (\mathbf{I}_m \otimes \mathbf{I}_m - \mathbf{\Phi}_1 \otimes \mathbf{\Phi}_1)^{-1}\, \mathrm{vec}(\mathbf{\Sigma}_v) \tag{E.2b}$$

For the long-term mean and variance to be defined, it is necessary that $|\det(\mathbf{\Phi}_1)| < 1$. The point of using a VAR specification is that the long-run expected returns and variances may deviate from the conditional short-run returns. E.g., with mean reversion the variances in $\mathbf{\Sigma}_z$ will be smaller than those in $\mathbf{\Sigma}_v$. Generally, the average return risk of holding a portfolio for k periods depends on k as

$$\begin{aligned}
&\frac{1}{k}\,\mathrm{var}\left(\sum_{t=1}^{k} \mathbf{z}_t \,|\, \mathbf{z}_0\right) = \frac{1}{k}\,\mathrm{var} \sum_{t=1}^{k} \sum_{i=0}^{t-1} \mathbf{\Phi}_1^i v_{t-i} \quad \Rightarrow \\
&\frac{1}{k}\,\mathrm{var}\left(\sum_{t=1}^{k} \mathbf{z}_t \,|\, \mathbf{z}_0\right) = \sum_{i=0}^{k-1} \frac{k-i}{k}\, \mathbf{\Phi}_1^i\, \mathbf{\Sigma}_v\, \mathbf{\Phi}_1^{i\,\prime}
\end{aligned} \tag{E.3a}$$

It follows that

$$\lim_{k\to\infty} \frac{1}{k}\,\mathrm{var}\left(\sum_{t=1}^{k} \mathbf{z}_t \,|\, \mathbf{z}_0\right) = \sum_{i=0}^{\infty} \mathbf{\Phi}_1^i\, \mathbf{\Sigma}_v\, \mathbf{\Phi}_1^{i\,\prime} = \mathbf{\Sigma}_z \lesseqgtr \mathbf{\Sigma}_v \tag{E.3b}$$

The estimated model is presented in Table 2 in Campbell and Viceira (2005). It contains the following variables: real return on T-bills, excess stock return, excess bond return, nominal bill rate, log dividend-price ratio, and yield spread. The latter three variables are state variables that appear to be well described in terms of an AR(1) process. The interaction among the return variables is also weak, except for a feedback from the bond return on the stock return. Most of

[67] The variance can be solved from $\mathbf{\Sigma}_z = \mathbf{\Phi}_1 \mathbf{\Sigma}_z \mathbf{\Phi}_1{}' + \mathbf{\Sigma}_v$. Applying the vec operator and noting that $\mathrm{vec}(\mathbf{\Phi}_1 \mathbf{\Sigma}_z \mathbf{\Phi}_1{}') = \mathbf{\Phi}_1 \otimes \mathbf{\Phi}_1 \mathrm{vec}(\mathbf{\Sigma}_z)$ gives the desired result.

the long-term effects come from the gradual spill-over from the state variables to the returns. As the model is stable, the estimated long-run returns must be approximately equal to the sample averages, i.e. a 1.5% return for T-bills, 6% excess return for stock, and a 1.4% excess return for Treasury bonds.

Appendix F The Affine Yield Model

The exposition follows Ang and Piazzesi (2003). In the affine yield model there are n independent risk factors, $\boldsymbol{\eta}' = (\eta_1, \ldots, \eta_n)$, with $\eta_i \propto \mathrm{N}(0,1)$. The stochastic discount factor m of investors is specified as

$$-\ln m_{t+1} = \ln(1 + r_{f_t}) + \boldsymbol{\chi}_t' \boldsymbol{\eta}_{t+1} + \tfrac{1}{2} \boldsymbol{\chi}' \boldsymbol{\chi} \tag{F.1a}$$

Each risk factor has its own 'price' χ_{i_t}.[68] The prices of risk are affine in the state of the economy,

$$\boldsymbol{\chi}_t = \boldsymbol{\upsilon}_0 + \boldsymbol{\Upsilon} \mathbf{x}_t \tag{F.1b}$$

where $\mathbf{x}_t$ is a vector of state variables that determine the returns

$$\mathbf{x}_{t+1} = \boldsymbol{\xi}_0 + \boldsymbol{\Xi}_0 \mathbf{x}_t + \boldsymbol{\Sigma} \boldsymbol{\eta}_{t+1} \tag{F.1c}$$

$$\ln(1 + \mathbf{r}_{t+1}) = \boldsymbol{\rho}_0 + \mathbf{P} \mathbf{x}_{t+1} \tag{F.1d}$$

Elements of $\mathbf{x}$ may comprise both observed and latent variables. m_{t+1} is not observed directly. However, asset returns obey the no-arbitrage conditions

$$\mathrm{E}\left[m_{t+1}\left(1 + r_{i,t+1}\right)\right] = 1, \quad i \in \{1, \ldots, n\} \tag{F.1e}$$

(F.1e) and the observable part of (F.1c) can be used to estimate the parameters of the model. The model is a generalisation of the Campbell/Viceira model of Appendix E, with as extra element a more elaborate and arbitrage-free term structure of asset returns, that can be obtained recursively from the stochastic discount factor as

$$p_{t,\tau} = \mathrm{E}\left[m_{t+1}\left(p_{t+1,\tau-1} + d_{t+1}\right]\right., \quad \tau \in \{1, \ldots, \infty\} \tag{F.2}$$

where d_t is the coupon or the dividend of the asset. The boundary conditions are $p_{t,0} = 1$, and, for equity, $\lim_{T \to \infty} \mathrm{E}\left[\prod_{i=1}^{T} m_{t+i} p_{t+T,\infty}\right] = 0$.

[68] Note that $\mathrm{E}[m_{t+1}] = \left(1 + r_{f_t}\right)^{-1}$ and $\mathrm{E}[m_{t+1} \mathrm{e}^{\eta_{i_t}}] = \mathrm{e}^{\frac{1}{2} - \chi_{i_t}} / \left(1 + r_{f_t}\right)$. As $\mathrm{E}[\mathrm{e}^{\eta_i}] = \mathrm{e}^{\frac{1}{2}}$, $\mathrm{e}^{\chi_i} \approx 1 + \chi_i$ is the return needed to persuade the investor to hold one unit of risk factor i.

References

Abel, A.B., 2001, Will bequests attenuate the predicted meltdown in stock prices when baby boomers retire?, *The Review of Economics and Statistics*, vol. 83, pp. 589–595.

Acemoglu, D. and S. Johnson, 2007, Disease and development: The effect of life expectancy on economic growth, *Journal of Political Economy*, vol. 115, pp. 925–985.

Alders, P.A. and D.P. Broer, 2005, Ageing, fertility, and growth, *Journal of Public Economics*, vol. 89, pp. 1075–1095.

Allen, F. and D. Gale, 1994, Limited market participation and volatility of asset prices, *The American Economic Review*, vol. 84, no. 4, pp. 933–955.

Allen, F. and D. Gale, 2004, Financial fragility, liquidity, and asset prices, *Journal of the European Economic Association*, vol. 2, no. 6, pp. 1015–1048.

Andersen, T.G., T. Bollerslev, P.F. Christoffersen and F.X. Diebold, 2006, Volatility and correlation forecasting, in G. Elliot, C.W.J. Granger and A. Timmerman, eds., *Handbook of Economic Forecasting*, vol. 1, chap. 15, pp. 777–878, North-Holland.

Andersen, T.G., T. Bollerslev, F.X. Diebold and H. Ebens, 2001, The distribution of realized stock return volatility, *Journal of Financial Economics*, vol. 61, no. 1, pp. 43–76.

Andersen, T.G., T. Bollerslev, F.X. Diebold and P. Labys, 2003, Modeling and forecasting realized volatility, *Econometrica*, vol. 71, no. 2, pp. 579–625.

Ang, A., G. Bekaert and M. Wei, 2008, The term structure of real rates and expected inflation, *Journal of Finance*, vol. 63, no. 2, pp. 797–849.

Ang, A. and A. Maddaloni, 2005, Do demographic changes affect risk premiums? Evidence from international data, *The Journal of Business*, vol. 78, pp. 341–379.

Ang, A. and M. Piazzesi, 2003, A no-arbitrage vector autoregression of term structure dynamics with macroeconomic and latent variables, *Journal of Monetary Economics*, vol. 50, pp. 745–787.

Avouyi-Dovi, S. and J. Matheron, 2006, Productivity and stock prices, *Financial Stability Review*, vol. 8, pp. 81–94.

Baker, H.K., G.E. Powell and E.T. Veit, 2002, Revisiting the dividend puzzle: Do all of the pieces now fit?, *Review of Financial Economics*, vol. 11, pp. 241–261.

Bakshi, G.S. and Z. Chen, 1994, Baby boom, population aging, and capital markets, *The Journal of Business*, vol. 67, pp. 165–202.

Bansal, R. and A. Yaron, 2004, Risks for the long run: A potential resolution of asset pricing puzzles, *The Journal of Finance*, vol. 59, no. 4, pp. 1481–1509.

Barro, R.J., 1996, Determinants of economic growth: A cross-country empirical study, NBER Working Paper 5698, NBER.

Barro, R.J., 2006, Rare disasters and asset markets in the twentieth century, *Quarterly Journal of Economics*, vol. 121, no. 3, pp. 823–866.

Barro, R.J., E. Nakamura, J. Steinsson and J. Ursúa, 2009, Crises and recoveries in an empirical model of consumption disasters, mimeo, Columbia University.

Barro, R.J. and J.F. Ursúa, 2008, Macroeconomic crises since 1870, *Brookings Papers on Economic Activity*, vol. 1, pp. 255–350.

Barsky, R.B., F.T. Juster, M.S. Kimball and M.D. Shapiro, 1997, Preference parameters and behavioral heterogeneity: An experimental approach in the health and retirement study, *Quarterly Journal of Economics*, vol. 112, no. 2, pp. 537–579.

Beaudry, P. and A. Guay, 1996, What do interest rates reveal about the functioning of real business cycle models?, *Journal of Economic Dynamics and Control*, vol. 20, no. 9-10, pp. 1661–1682.

Beaudry, P. and F. Portier, 2005, The "news" view of economic fluctuations: Evidence from aggregate Japanese data and sectoral U.S. data, *Journal of the Japanese and International Economies*, vol. 19, no. 4, pp. 635–652.

Beaudry, P. and F. Portier, 2006, Stock prices, news, and economic fluctuations, *American Economic Review*, vol. 96, no. 4, pp. 1293–1307.

Becker, G.S., K.M. Murphy and R. Tamura, 1990, Human capital, fertility, and economic growth, *Journal of Political Economy*, vol. 98, pp. S12–S37.

Bekaert, G. and G. Wu, 2000, Asymmetric volatility and risk in equity markets, *Review of Financial Studies*, vol. 13, pp. 1–42.

Benati, L., 2007, Drift and breaks in labor productivity, *Journal of Economic Dynamics and Control*, vol. 31, no. 8, pp. 2847–2877.

Benzoni, L., P. Collin-Dufresne and R.S. Goldstein, 2007, Portfolio choice over the life-cycle when the stock and labor markets are cointegrated, *The Journal of Finance*, vol. 62, no. 5, pp. 2123–2167.

Bettendorf, L.J.H. and H. de Wachter, 2007, Ageing and the relative price of nontradeables, TI Discussion Paper 07-064, Tinbergen Institute.

Black, F., 1976a, The dividend puzzle, *Journal of Portfolio Management*, vol. 2, pp. 5–8.

Black, F., 1976b, Studies of stock price volatility changes, in *Proceedings of the 1976 Meetings of the Business and Economic Statistics Section*, pp. 177–181, American Statistical Association.

Bodie, Z., 1976, Common stocks as a hedge against inflation, *Journal of Finance*, vol. 31, pp. 459–470.

Boersch-Supan, A., 2003, Labor market effects of population aging, *Labour*, vol. 17, pp. 5–44.

Boersch-Supan, A., A. Ludwig and J. Winter, 2006, Aging, pension reform, and capital flows: A multi-country simulation model, *Economica*, vol. 73, pp. 625–658.

Bohn, H., 2009, Intergenerational risk sharing and fiscal policy, mimeo, University of California at Santa Barbara.

Bollerslev, T., R.F. Engle and D.B. Nelson, 1994, ARCH models, in R.F. Engle and D.L. McFadden, eds., *Handbook of Econometrics*, vol. 4, chap. 49, pp. 2959–3038, Elsevier.

Bollerslev, T., R.F. Engle and J.M. Wooldridge, 1988, A capital asset pricing model with time-varying covariances, *Journal of Policital Economy*, vol. 96, no. 1, pp. 116–131.

Brandt, M.W. and Q. Kang, 2004, On the relationship between the conditional mean and volatility of stock returns: A latent VAR approach, *Journal of Financial Econometrics*, vol. 72, no. 2, pp. 217–257.

Brooks, R., 2000, What will happen to financial markets when the baby boomers retire?, IMF Working Paper 00/18, IMF.

Brooks, R., 2002, Asset-market effects of the baby boom and social-security reform, *American Economic Review, Papers and Proceedings*, vol. 92, pp. 402–406.

Brooks, R., 2003a, The equity premium and the baby boom, Econometric Society North American Winter Meeting 155, Econometric Society.

Brooks, R., 2003b, Population aging and global capital flows in a parallel universe, *IMF Staff Papers*, vol. 50, pp. 200–221.

Bryant, J., 1980, A model of reserves, bank runs, and deposit insurance, *Journal of Banking and Finance*, vol. 4, no. 4, pp. 335–344.

Campbell, J.Y., 2003, Consumption-based asset pricing, in G.M. Constantinides, M. Harris and R.M. Stulz, eds., *Handbook of the Economics of Finance*, vol. 1B, chap. 13, pp. 803–888, North-Holland.

Campbell, J.Y., Y.L. Chan and L.M. Viceira, 2003, A multivariate model of strategic asset allocation, *Journal of Financial Economics*, vol. 67, no. 1, pp. 41–80.

Campbell, J.Y. and J.H. Cochrane, 1995, By force of habit: A consumption-based explanation of aggregate stock market behavior, NBER Working Paper 4995, NBER.

Campbell, J.Y. and J.H. Cochrane, 1999, By force of habit: A consumption-based explanation of aggregate stock market behavior, *Journal of Polical Economy*, vol. 107, no. 2, pp. 205–251.

Campbell, J.Y., A.W. Lo and A.C. MacKinlay, 1997, *The Econometrics of Financial Markets*, Princeton University Press.

Campbell, J.Y. and N.G. Mankiw, 1987, Are output fluctuations transitory?, *The Quarterly Journal of Economics*, vol. 102, pp. 857–880.

Campbell, J.Y. and R.J. Shiller, 2001, Valuation ratios and the long-run stock market outlook: An update, NBER Working Paper 8221, NBER.

Campbell, J.Y. and L.M. Viceira, 1999, Consumption and portfolio decisions when expected returns are time varying, *The Quarterly Journal of Economics*, vol. 114, pp. 433–495.

Campbell, J.Y. and L.M. Viceira, 2002, *Strategic Asset Allocation: Portfolio Choice for Long-Term Investors*, Oxford University Press, Oxford.

Campbell, J.Y. and L.M. Viceira, 2005, The term structure of the risk-return tradeoff, CEPR Discussion Paper 4914, CEPR.

Campbell, S.D. and F.X. Diebold, 2009, Stock returns and expected business conditions: Half a century of direct evidence, *Journal of Business and Economic Statistics*, vol. 27, no. 2, pp. 266–278.

Carlaw, K.I. and R.G. Lipsey, 2003, Productivity, technology and economic growth: What is the relationship?, *Journal of Economic Surveys*, vol. 17, pp. 457–495.

Cerra, V. and S.C. Saxena, 2008, Growth dynamics: The myth of economic recovery, *American Economic Review*, vol. 98, no. 1, pp. 439–457.

Christiano, L.J. and M. Eichenbaum, 1990, Unit roots in real GNP: Do we know and do we care?, *Carnegie-Rochester Conference Series on Public Policy*, vol. 32, pp. 7–62.

Christie, A.A., 1982, The stochastic behavior of common stock variances: Value, leverage and interest rate effects, *Journal of Financial Economics*, vol. 10, no. 4, pp. 407–432.

Clark, X., T.J. Hatton and J.G. Williamson, 2002, Where do U.S. immigrants come from, and why?, NBER Working Paper 8998, NBER.

Cochrane, J., 2008, Financial markets and the real economy, in R. Mehra, ed., *Handbook of the Equity Risk Premium*, chap. 7, pp. 237–325, Elsevier.

Cochrane, J.H., 1988, How big is the random walk in GNP?, *Journal of Political Economy*, vol. 96, pp. 893–920.

Cochrane, J.H., 1994, Permanent and transitory components of GNP and stock prices, *Quarterly Journal of Economics*, vol. 109, pp. 241–265.

Cochrane, J.H., 2001, *Asset Pricing*, Princeton University Press, Princeton.

Cochrane, J.H., 2007, The dog that did not bark: A defence of return predictability, *Review of Financial Studies*, vol. 21, no. 4, pp. 1533–1575.

Cogley, T., 1990, International evidence on the size of the random walk in output, *Journal of Political Economy*, vol. 98, no. 3, pp. 501–518.

Constantinides, G.M., J.B. Donaldson and R. Mehra, 2002, Junior can't borrow: a new perspective on the equity premium puzzle, *Quarterly Journal of Economics*, vol. 117, pp. 269–296.

Cook, S., 2008, More uncertainty: on the trending nature of real GDP in the US and UK, *Applied Economics Letters*, vol. 15, no. 9, pp. 667–670.

Diamond, D.W. and P.H. Dybvig, 1983, Bank runs, deposit insurance, and liquidity, *Journal of Political Economy*, vol. 91, no. 3, pp. 401–419.

Diebold, F.X. and P. Lindner, 1996, Fractional integration and interval prediction, *Economics Letters*, vol. 50, pp. 305–313.

Diebold, F.X. and G.D. Rudebusch, 1989, Long memory and persistence in aggregate output, *Journal of Monetary Economics*, vol. 24, no. 2, pp. 189–209.

Diebold, F.X. and G.D. Rudebusch, 1996, Measuring business cycles: A modern perspective, *Review of Economics and Statistics*, vol. 78, pp. 67–77.

Diebold, F.X. and A.S. Senhadji, 1996, The uncertain unit root in real GNP: Comment, *American Economic Review*, vol. 86, pp. 1291–1298.

Dimson, E., P. Marsh and M. Staunton, 2008, The worldwide equity premium: A smaller puzzle, in R. Mehra, ed., *Handbook of the Equity Risk Premium*, chap. 11, pp. 467–514, Elsevier.

Domeij, D. and M. Flodén, 2006, Population aging and international capital flows, *International Economic Review*, vol. 47, pp. 1013–1032.

Donaldson, J. and R. Mehra, 2008, Risk-based explanations of the equity premium, in R. Mehra, ed., *Handbook of the Equity Risk Premium*, chap. 2, pp. 37–99, Elsevier.

Dostie, B., 2006, Wages, productivity and aging, IZA Discussion Paper No. 2496, IZA.

Duffie, D. and R. Kan, 1996, A yield-factor model of interest rates, *Mathematical Finance*, vol. 6, no. 4, pp. 379–406.

Epstein, L.G. and S. Zin, 1989, Substitution, risk aversion, and the temporal behavior of consumption and asset returns: A theoretical framework, *Econometrica*, vol. 57, pp. 937–969.

Evans, M.D., 2003, Real risk, inflation risk, and the term structure, *The Economic Journal*, vol. 113, no. 487, pp. 345–389.

Fama, E.F., 1965, The behavior of stock-market prices, *The Journal of Business*, vol. 38, no. 1, pp. 34–105.

Fama, E.F. and K. French, 1989, Business conditions and expected returns on stocks and bonds, *Journal of Financial Economics*, vol. 25, no. 1, pp. 23–49.

Fama, E.F. and K.R. French, 1988, Permanent and temporary components of stock prices, *Journal of Political Economy*, vol. 96, pp. 246–273.

Fama, E.F. and K.R. French, 2002, The equity premium, *The Journal of Finance*, vol. 57, no. 2, pp. 637–659.

Fama, E.F. and G.W. Schwert, 1977, Asset returns and inflation, *Journal of Financial Economics*, vol. 5, no. 2, pp. 115–146.

Fehr, H., S. Jokisch and L.J. Kotlikoff, 2004, Fertility, mortality, and the developed world's demographic transition, Cesifo Working Paper 1326, Cesifo.

Fehr, H., S. Jokisch and L.J. Kotlikoff, 2005, Will China eat our lunch or take us out to dinner? Simulating the transition paths of the U.S., EU, Japan, and China, NBER Working Paper 11668, NBER.

Feyrer, J., 2007, Demographics and productivity, *The Review of Economics and Statistics*, vol. 89, no. 1, pp. 100–109.

Fougère, M. and M. Mérette, 1999, Population ageing and economic growth in seven OECD countries, *Economic Modelling*, vol. 16, pp. 411–427.

French, K.R., G.W. Schwert and R.F. Stambaugh, 1987, Expected stock returns and volatility, *Journal of Financial Economics*, vol. 19, no. 1, pp. 3–29.

Gaffeo, E., M. Gallegati and M. Gallegati, 2005, Requiem for the unit root in per capita real GDP? Additional evidence from historical data, *Empirical Economics*, vol. 30, pp. 37–63.

Galí, J., 1999, Technology, employment, and the business cycle: Do technology shocks explain aggregate fluctuations?, *American Economic Review*, vol. 89, pp. 249–271.

Galor, O. and D.N. Weil, 1996, The gender gap, fertility, and growth, *American Economic Review*, vol. 86, pp. 374–387.

Galor, O. and D.N. Weil, 2000, Population, technology, and growth: From Malthusian stagnation to the demographic transition and beyond, *American Economic Review*, vol. 90, pp. 806–828.

Geweke, J., 2001, A note on some limitations of CRRA utility, *Economics Letters*, vol. 71, no. 3, pp. 341–345.

Goetzmann, W.N. and P. Jorion, 1993, Testing the predictive power of dividend yields, *The Journal of Finance*, vol. 48, pp. 663–679.

Gollier, C., 2007, The consumption-based determinants of the term structure of discount rates, *Mathematics and Financial Economics*, vol. 1, pp. 81–101.

Gollier, C., 2008, Discounting with fat-tailed economic growth, *Journal of Risk and Uncertainty*, vol. 37, pp. 171–186.

Gomme, P., R. Rogerson, P. Rupert and R. Wright, 2004, The business cycle and the life cycle, Federal Reserve Bank of Cleveland Working Paper 04-04, Federal.

Goyal, A., 2004, Demographics, stock market flows, and stock returns, *Journal of Financial and Quantitative Analysis*, vol. 39, no. 1, pp. 115–142.

Goyal, A. and I. Welch, 2003, Predicting the equity premium with dividend ratios, *Management Science*, vol. 49, pp. 639–654.

Goyal, A. and I. Welch, 2008, A comprehensive look at the empirical performance of equity premium prediction, *Review of Financial Studies*, vol. 21, pp. 1455–1508.

Greenwood, J. and M. Yorukoglu, 1997, 1974, *Carnegie-Rochester Conference Series on Public Policy*, vol. 46, pp. 49–95.

Hall, R.E., 1978, Stochastic implications of the life cycle-permanent income hypothesis: Theory and evidence, *Journal of Political Economy*, vol. 86, pp. 971–987.

Hári, N., A. de Waegenaere, B. Melenberg and T.E. Nijman, 2008, Estimating the term structure of mortality, *Insurance: Mathematics and Economics*, vol. 42, pp. 492–504.

Harris, J.R. and M.P. Todaro, 1970, Migration, unemployment and development: A two-sector analysis, *American Economic Review*, vol. 60, pp. 126–142.

Hek, P. de and D. van Vuuren, 2008, Wage and productivity over the worker's life cycle, CPB Document forthcoming, CPB Netherlands Bureau for Economic Policy Analysis.

Hellerstein, J.K., D. Neumark and K.R. Troske, 1999, Wages, productivity, and worker characteristics: Evidence from plant-level production functions and wage equations, *Journal of Labor Economics*, vol. 17, pp. 409–446.

Hoevenaars, R.P.M.M., R.D.J. Molenaar, P.C. Schotman and T.B.M. Steenkamp, 2008, Strategic asset allocation with liabilities: Beyond stocks and bonds, *Journal of Economic Dynamics and Control*, vol. 32, no. 9, pp. 2939–2970.

Hulten, C.R., 2000, Total factor productivity: A short biography, NBER Working Paper 7471, NBER.

Jaimovich, N. and S. Rebelo, 2009, Can news about the future drive the business cycle?, *The American Economic Review*, vol. 99, no. 4, pp. 1097–1118.

Jaimovich, N. and H.E. Siu, 2009, The young, the old, and the restless: Demographics and busisess cycle volatility, *The American Economic Review*, vol. 99, no. 3, pp. 804–826.

Jones, C.I., 2001, Was an industrial revolution inevitable? Economic growth over the very long run, *Advances in Macroeconomics*, vol. 1, no. 2.

Jorion, P. and W.N. Goetzmann, 1999, Global stock markets in the twentieth century, *The Journal of Finance*, vol. 54, no. 3, pp. 953–980.

Karemera, D., V.I. Oguledo and B. Davis, 2000, A gravity model analysis of international migration to North America, *Applied Economics*, vol. 32, pp. 1745–1755.

Keilman, N., H. Cruijsen and J.M. Alho, 2008, Changing views of future demographic trends, in J.M. Alho, S.E.H. Jensen and J. Lassila, eds., *Uncertain Demographics and Fiscal Sustainability*, chap. 2, pp. 11–54, Cambridge University Press.

King, R.G. and S.T. Rebelo, 1999, Resuscitating real business cycles, in J.B. Taylor and M. Woodford, eds., *Handbook of Macroeconomics*, vol. 1B, chap. 14, pp. 927–1007, Elsevier.

King, R.G. and M.W. Watson, 1996, Money, prices, interest rates and the business cycle, *Review of Economics and Statistics*, vol. 78, pp. 35–53.

Kocherlakota, N.R., 1996, The equity premium puzzle: It's still a puzzle, *Journal of Economic Literature*, vol. 34, pp. 42–71.

Kotlikoff, L. and D. Wise, 1989, Employee retirement and a firm's pension plan, in D.A. Wisse, ed., *The Economics of Aging*, pp. 279–334, University of Chicago Press, Chicago.

Kotlikoff, L.J. and J. Gokhale, 1992, Estimating a firm's age-productivity profile using the present value of workers' earnings, *Quarterly Journal of Economics*, vol. 107, pp. 1215–1242.

Kotlikoff, L.J., K. Smetters and J. Walliser, 2001, Finding a way out of America's demographic dilemma, NBER working paper 8258, NBER.

Kremer, M., 1993, Population growth and technological change: One million B.C. to 1990, *The Quarterly Journal of Economics*, vol. 108, pp. 681–716.

Krueger, D. and A. Ludwig, 2007, On the consequences of demographic change for rates of return to capital, and the distribution of wealth and welfare, *Journal of Monetary Economics*, vol. 54, pp. 49–87.

Lee, R.D., 1988, Induced population growth and induced technological progress: Their interaction in the accelerating stage, *Mathematical Population Studies*, vol. 1, no. 3, pp. 265–288.

Lee, R.D. and L.R. Carter, 1992, Modeling and forecasting U.S. mortality, *Journal of the American Statistical Association*, vol. 87, pp. 659–671.

Lee, R.D. and S. Tuljapurkar, 1994, Stochastic population forecasts for the United States: Beyond high, medium, and low, *Journal of the American Statistical Association*, vol. 89, pp. 1175–1189.

Lemke, W. and T. Werner, 2009, The term structure of equity premia in an affine arbitrage-free model of bond market and stock market dynamics, ECB Working Paper 1045, European Central Bank.

Lettau, M. and S.C. Ludvigson, 2001, Consumption, aggregate wealth and expected stock returns, *Journal of Finance*, vol. 56, pp. 815–849.

Lettau, M. and S.C. Ludvigson, 2002, Time-varying risk premia and the cost of capital: An alternative implication of the Q theory of investment, *Journal of Monetary Economics*, vol. 49, pp. 31–66.

Lettau, M. and S.C. Ludvigson, 2005, Expected returns and expected dividend growth, *Journal of Financial Economics*, vol. 76, no. 3, pp. 583–626.

Lettau, M. and S.C. Ludvigson, 2009, Measuring and modeling variation in the risk-return tradeoff, in Y. Ait-Sahalia and L. Hansen, eds., *Handbook of Financial Econometrics*, vol. 1, chap. 12, p. forthcoming, Elsevier.

Lettau, M., S.C. Ludvigson and J. Wachter, 2008, The declining equity premium: What role does macroeconomic risk play?, *The Review of Financial Studies*, vol. 21, no. 4, pp. 1653–1687.

Lettau, M. and S. van Nieuwerburgh, 2008, Reconciling the return predictability evidence, *Review of Financial Studies*, vol. 21, pp. 1607–1652.

Lima, L.R. and J. de Jesus Filho, 2008, Further investigation of the uncertain trend in US GDP, *Applied Economics*, vol. 40, no. 9, pp. 1207–1216.

Lindbeck, A. and M. Persson, 2003, The gains from pension reform, *Journal of Economic Literature*, vol. 41, no. 1, pp. 74–112.

Lindh, T., 2004, Medium-term forecasts of potential GDP and inflation using age structure information, *Journal of Forecasting*, vol. 23, pp. 19–49.

Lindh, T. and B. Malmberg, 1999, Age structure effects and growth in the OECD, 1950-1990, *Journal of Population Economics*, vol. 12, pp. 431–449.

Lindh, T. and B. Malmberg, 2000, Can age structure forecast inflation trends?, *Journal of Economics and Business*, vol. 52, no. 1-2, pp. 31–49.

Ljungqvist, L. and T.J. Sargent, 2004, *Recursive Macroeconomic Theory*, MIT Press, 2nd ed.

Lorenzoni, G., 2009, A theory of demand shocks, *The American Economic Review*, vol. 99, no. 5, pp. 2050–2084.

Ludwig, A., D. Krueger and A.H. Boersch-Supan, 2007, Demographic change, relative factor prices, international capital flows, and their differential effects on the welfare of generations, NBER Working Paper 13185, NBER.

Mandelbrot, B., 1963, The variation of certain speculative prices, *The Journal of Business*, vol. 36, no. 4, pp. 394–419.

Matsen, E. and Ø. Thøgersen, 2004, Designing social security - a portfolio choice approach, *European Economic Review*, vol. 48, pp. 883–904.

Mayda, A.M., 2005, International migration: A panel data analysis of economic and non-economic determinants, IZA Discussion Paper 1590, IZA.

Mehra, R. and E.C. Prescott, 1985, The equity premium: a puzzle, *Journal of Monetary Economics*, vol. 15, no. 2, pp. 145–161.

Mehra, R. and E.C. Prescott, 2008, Non-risk-based explanations of the equity premium, in R. Mehra, ed., *Handbook of the Equity Risk Premium*, chap. 3, pp. 101–115, Elsevier.

Merton, R.C., 1969, Lifetime portfolio selection under uncertainty: The continuous-time case, *The Review of Economics and Statistics*, vol. 51, no. 3, pp. 247–257.

Merton, R.C., 1980, On estimating the expected return on the market: An exploratory investigation, *Journal of Financial Economics*, vol. 8, no. 4, pp. 323–361.

Miles, D., 1999, Modelling the impact of demographic change upon the economy, *The Economic Journal*, vol. 109, pp. 1–36.

Mincer, J., 1974, *Schooling, Experience and Earnings*, National Bureau of Economic Research.

Murray, C.J. and C.R. Nelson, 2000, The uncertain trend in U.S. GDP, *Journal of Monetary Economics*, vol. 46, pp. 79–95.

Murray, C.J. and C.R. Nelson, 2002, The great depression and output persistence, *Journal of Money, Credit and Banking*, vol. 34, pp. 1090–1098.

Nahuis, R., E. Canton and H. de Groot, 2000, Vested interests, ageing, and growth, Tech. Rep. 3, CPB Netherlands Bureau for Economic Policy Analysis.

Nelson, C.R. and M.J. Kim, 1993, Predictable stock returns: The role of small sample bias, *The Journal of Finance*, vol. 48, no. 2, pp. 641–661.

Nelson, C.R. and C.I. Plosser, 1982, Trends and random walks in macroeconomic time series. some evidence and implications, *Journal of Monetary Economics*, vol. 10, no. 2, pp. 139–162.

Nelson, D.B., 1991, Conditional heteroskedasticity in asset returns: A new approach, *Econometrica*, vol. 59, no. 2, pp. 347–370.

OECD, 1998, Work force ageing: Consequences and policy responses, Working Paper AWP 4.1, OECD.

Perron, P., 1989, The great crash, the oil price shock, and the unit root hypothesis, *Econometrica*, vol. 57, pp. 1361–1401.

Poterba, J.M., 2001, Demographic structure and asset returns, *Review of Economics and Statistics*, vol. 83, pp. 565–584.

Poterba, J.M., 2004, The impact of population aging on financial markets, NBER Working Paper 10851, NBER.

Prescott, E.C., 1998, Needed: A theory of total factor productivity, *International Economic Review*, vol. 39, pp. 525–551.

Prskawetz, A. and T. Lindh, 2006, The impact of population ageing on innovation and productivity growth in Europe, Research Report 28, Vienna Institute of Demography.

Rappoport, P. and L. Reichlin, 1989, Segmented trends and non-stationary time series, *The Economic Journal*, vol. 99, pp. 168–177.

Reinhart, C.M. and K. Rogoff, 2008, This time is different: A panoramic view of eight centuries of financial crises, NBER Working Paper 13882, NBER.

Reinhart, C.M. and K.S. Rogoff, 2009, The aftermath of financial crises, NBER Working Paper 14656, NBER.

Ríos-Rull, J.V., 1996, Life-cycle economies and aggregate fluctuations, *The Review of Economic Studies*, vol. 63, pp. 465–489.

Rudebusch, G.D., 1993, The uncertain unit root in real GNP, *American Economic Review*, vol. 83, no. 1, pp. 264–272.

Salomons, R., 2008, A theoretical and practical perspective on the equity risk premium, *Journal of Economic Surveys*, vol. 22, no. 2, pp. 299–329.

Samuelson, P.A., 1969, Lifetime portfolio selection by dynamic programming, *The Review of Economics and Statistics*, vol. 51, pp. 239–246.

Sarkar, A. and L. Zhang, 2004, Time-varying consumption correlation and the dynamics of the equity premium: Evidence from the G-7 countries, Federal Reserve Bank of New York Staff Report 181, Federal Reserve Bank of New York.

Schwert, G.W., 1989, Why does stock market volatility change over time?, *Journal of Finance*, vol. 44, no. 5, pp. 1115–1153.

Sen, A., 2004, Are US macroeconomic series difference stationary or trend-break stationary?, *Applied Economics*, vol. 36, pp. 2025–2029.

Shiller, R.J., 1982, Consumption, asset markets, and economic fluctuations, *Carnegie-Rochester Conference Series on Public Policy*, vol. 17, pp. 203–238.

Stern, N., 2006, *The Economics of Climate Change: The Stern Review*, Cambridge University Press.

Stock, J.H., 1991, Confidence intervals for the largest autoregressive root in U.S. macroeconomic time series, *Journal of Monetary Economics*, vol. 28, pp. 435–459.

Temple, J., 1999, The new growth evidence, *Journal of Economic Literature*, vol. 37, pp. 112–156.

The Conference Board, 2009, Total economy database, http://www.conference-board.org/economics.

Valkanov, R., 2003, Long-horizon regressions: Theoretical results and applications, *Journal of Financial Economics*, vol. 68, no. 2, pp. 201–232.

Violante, G.L., 2002, Technological acceleration, skill transferability, and the rise in residual inequality, *Quarterly Journal of Economics*, vol. 117, pp. 297–338.

Vougas, D.V., 2007, Is the trend in post-WW II US real GDP uncertain or non-linear?, *Economics Letters*, vol. 94, pp. 348–355.

Weil, D.A., 2007, Accounting for the effect of health on economic growth, *Quarterly Journal of Economics*, vol. 122, pp. 1265–1306.

Weil, P., 1989, Overlapping families of infinitely-lived agents, *Journal of Public Economics*, vol. 38, no. 2, pp. 183–198.

Weitzman, M.L., 2007, Subjective expectations and asset-return puzzles, *American Economic Review*, vol. 97, pp. 1102–1130.

Werding, M., 2008, Ageing and productivity growth: Are there macro-level cohort effects of human capital?, CESifo Working Paper 2207, CESifo.

Young, A., 2005, The gift of the dying: The tragedy of AIDS and the welfare of future african generations, *Quarterly Journal of Economics*, vol. 120, pp. 423–466.

Zivot, E. and D.W.K. Andrews, 1992, Further evidence on the great crash, the oil-price shock, and the unit-root hypothesis, *Journal of Business & Economic Statistics*, vol. 10, pp. 25–44.