Netspar

Network for Studies on Pensions, Aging and Retirement

Netspar PANEL PAPERS

Antoon Pelsser

Pricing in incomplete markets

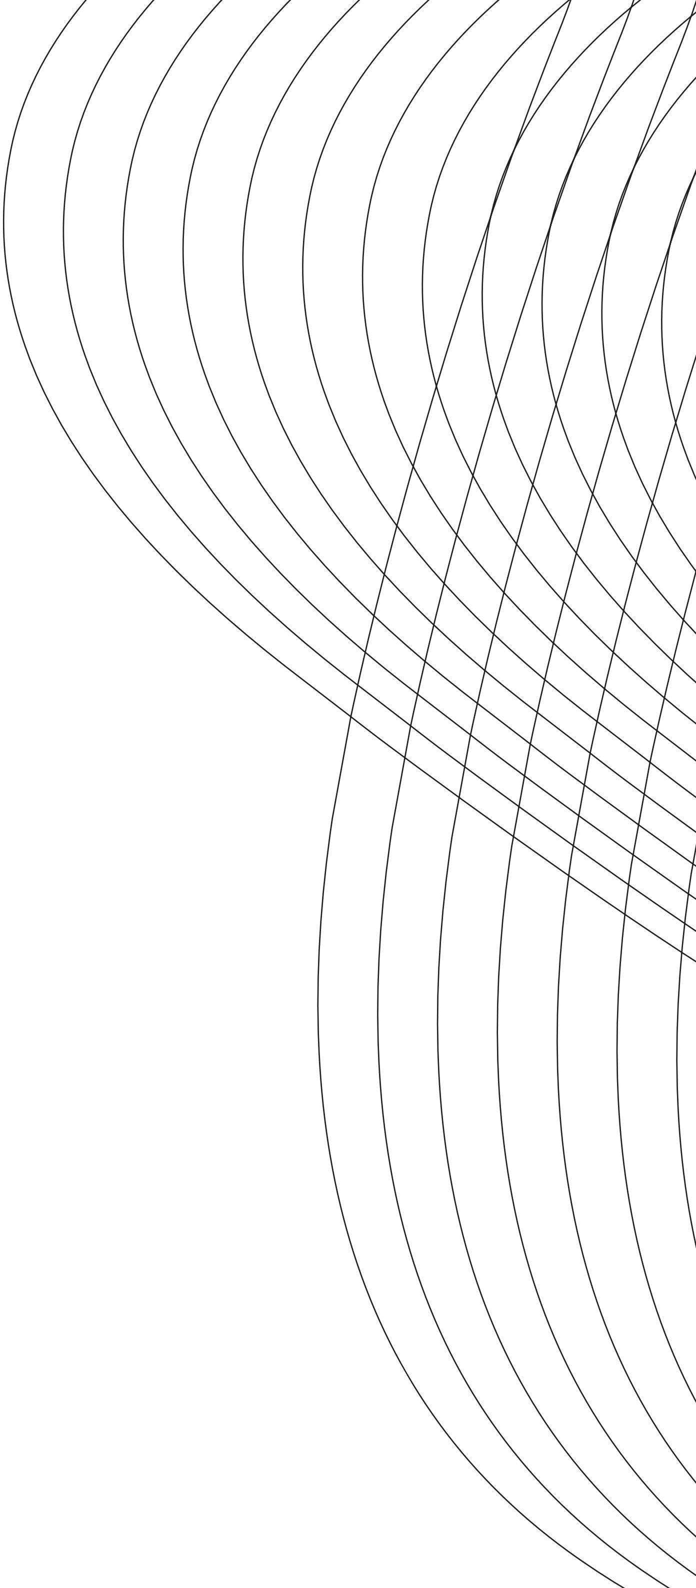

Antoon Pelsser

Pricing in incomplete markets

PANEL PAPER 25

Network for Studies on Pensions, Aging and Retirement

Colophon
Panel Papers is a publication of Netspar
June 2011

Design
B-more Design
Bladvulling, Tilburg

Printing
Printing Office Tilburg University

Editorial address
Netspar, Tilburg University
PO Box 90153, 5000 LE Tilburg
info@netspar.nl

CONTENTS

PREFACE

Netspar stimulates debate and fundamental research in the field of pensions, aging and retirement. The aging of the population is front-page news, as many baby boomers are now moving into retirement. More generally, people live longer and in better health while at the same time families choose to have fewer children. Although the aging of the population often gets negative attention, with bleak pictures painted of the doubling of the ratio of the number of people aged 65 and older to the number of the working population during the next decades, it must, at the same time, be a boon to society that so many people are living longer and healthier lives. Can the falling number of working young afford to pay the pensions for a growing number of pensioners? Do people have to work a longer working week and postpone retirement? Or should the pensions be cut or the premiums paid by the working population be raised to afford social security for a growing group of pensioners? Should people be encouraged to take more responsibility for their own pension? What is the changing role of employers associations and trade unions in the organization of pensions? Can and are people prepared to undertake investment for their own pension, or are they happy to leave this to the pension funds? Who takes responsibility for the pension funds? How can a transparent and level playing field for pension funds and insurance companies be ensured? How should an acceptable trade-off be struck between social goals such as solidarity between young and old, or rich and poor, and individual freedom? But most important of all: how can the benefits of living longer and healthier be harnessed for a happier and more prosperous society?

The Netspar Panel Papers aim to meet the demand for understanding the ever-expanding academic literature on the consequences of aging populations. They also aim to help give a better scientific underpinning of policy advice. They attempt to provide a survey of the latest and most relevant research, try to explain this in a non-technical manner and outline the implications for policy questions faced by Netspar's partners. Let there be no mistake. In many ways, formulating such a position paper is a tougher task than writing an academic paper or an op-ed piece. The authors have benefitted from the comments of the Editorial Board on various drafts and also from the discussions during the presentation of their paper at a Netspar Panel Meeting.

I hope the result helps reaching Netspar's aim to stimulate social innovation in addressing the challenges and opportunities raised by aging in an efficient and equitable manner and in an international setting.

Henk Don
Chairman of the Netspar Editorial Board

Affiliations

Antoon Pelsser – Maastricht University

PRICING IN INCOMPLETE MARKETS

Abstract
This Netspar Panel Paper discusses the pricing of contracts in an incomplete market setting. For life insurance companies and pension funds, it is always the case in practice that not all of the risks in their books can be hedged. Hence, the standard Black-Scholes methodology cannot be applied in this situation. The paper discusses and compares several methods that have been proposed in the literature in recent years: the Cost-of-Capital method (the current industry standard), Good Deal Bound pricing, and pricing under Model Ambiguity. Although each of these methods has a very different economic starting point, we show that all three converge for small time-steps to the same limit. This convergence provides a basis for comparing the different parameters used by the three methods. From this comparison we conclude that the current cost-of-capital of 6% used by the industry and CEIOPS is too low, since it is not in line with the values implied by the Good Deal Bound and Model Ambiguity methods. A cost-of-capital of 12% is needed to bring the method in line with the other two methods.

1. Management Summary & Policy Recommendations

Life insurance companies and pension funds have liabilities on their books with very long-dated maturities. The valuation and risk-management of these very long-dated contracts is therefore an important problem in practice.

The standard theory (based on replicating the cash flows) fails because there are simply no financial contracts that last this long. In well-developed economies (such as the euro-zone countries and the US) government bonds have maturities up to 30 years.

On the other hand, regulators in many countries (especially in Europe under the Solvency II project) are insisting that insurance companies (and in The Netherlands, also pension funds) value their liabilities on a "market-consistent" basis. Hence, to value these long-dated cash flows in a market-consistent way, one is forced to extend the term-structure of interest rates, which can be observed from financial markets, beyond the maturity of the longest dated instrument that can be observed in the market. In the current economic circumstances, with low long-term interest rates, pension funds are reporting low funding levels as a consequence of these valuation rules. A related issue is how to select financial instruments that give the best possible investment strategy (or hedge) for these very long-dated cash flows. In many cases this involves striking a balance between seeking assets with a higher return, at the expense of accepting a higher mismatch risk between the liabilities and the assets.

From a scientific point of view, the problem of pricing these very long-dated contracts boils down to the valuation of contracts in an incomplete markets setting. This means trying to price contracts where not all of the risks can be traded (and hedged) in financial markets. In the past ten years significant progress has been made regarding this subject. This Panel Paper discusses and compares several methods that have been proposed in the literature: the Cost-of-Capital method (the current industry standard), Good Deal Bound pricing, and pricing under Model Ambiguity. We show that each of these three methods converges for small time-steps to the same limit. This convergence provides a basis for comparing the different parameters used by the three methods.

The results presented in this paper allow us to provide the following policy recommendations:

- The "Cost-of-Capital" method proposed by the insurance industry and CEIOPS (i.e. the market-consistent price of an insurance contract, which is determined by the market value of the *replicating portfolio*, plus a mark-up for the unhedgeable risks: *the risk margin*; see EIOPA (2010) for further details) has qualitatively the right properties, but lacks a solid theoretical foundation. A pricing method with a rigourous theoretical foundation can be obtained by using the pricing methods put forward in this paper.

- In particular, the formulas for calculating market-consistent prices for *multi-year products*—as put forward by EIOPA in QIS5—lack a theoretical basis, and should be seen as a coarse approximation at best. The main problem is that the proposed QIS5-methodology is not time-consistent. We recommend that CEIOPS adopts a time-consistent pricing method based on backward-induction calculation techniques.

- The time-consistent pricing method proposed in this paper calculates prices under an "actuarially prudent" model, where (for each time-step) the best-estimate mean is adjusted by k times the standard deviation of the unhedgeable risk of the whole portfolio. The Good Deal Bound approach implies $k > 0.25$, the Model Ambiguity approach implies $k \approx 0.30$, and the Cost-of-Capital approach implies $k = 0.15$. The values k for the first two approaches are in line with each other, but the value implied by the Cost-of-Capital method seems too low. A value of $k = 0.30$ is needed to bring the Cost-of-Capital method in line with the other two methods, which corresponds to a cost-of-capital of 12% (instead of the 6% currently proposed by the industry and CEIOPS).

- Regulators are particularly vulnerable to model risk. When the regulator puts forward a very explicitly specified standard model (as is currently happening under Solvency II), then competitive market forces will ensure that most of the risk accumulates at the "weakest point" of the regulator's model. To guard against this model risk, we propose that the regulator adopts a robust approach to model risk. This can be achieved by putting forward several alternative models, and the industry should then calculate Solvency Capital on the basis of the worst outcome under the different models.

2. Introduction

Life insurance companies and pension funds have liabilities on their books with very long-dated maturities. Most people start saving for their pension from age 25, and people are expected to live to age 85, with the oldest people living to age 115. Hence, pension funds and life insurance companies are facing contractual obligations that can easily last 60 years—and sometimes even 80 or 90 years—into the future. The valuation and risk-management of these very long-dated contracts is therefore an important problem. To give a feel for the size of the problem: for life-insurance and pension products, a portion of roughly 20% of the net present value of the cash flows is located in the tail of 30+ years.

The standard theory (based on replicating the cash flows) fails because there are simply no financial contracts which last this long. In well-developed economies (such as countries in the euro-zone and the US), the longest government bonds have maturities up to 30 years. In developing countries (such as Eastern Europe, and Latin America and Asia), government bonds are issued with much shorter maturities (typically only up to ten years, and sometimes even much shorter).

On the other hand, regulators in many countries (especially in Europe under the Solvency II project) are insisting that insurance companies (and in The Netherlands, also pension funds) value their liabilities on a “market-consistent” basis. Hence, to value these long-dated cash flows in a market-consistent way, one is forced to extend the term-structure of interest rates, which can be observed from financial markets, beyond the maturity of longest dated instrument that can be observed in the market. In the current economic circumstances, with low long-term interest rates, pension funds are reporting low funding levels as a consequence of these valuation rules. A related issue is how to select financial instruments that give the best possible investment strategy (or hedge) for these very long-dated cash flows. In many cases this involves striking a balance between seeking assets with a higher return, at the expense of accepting a higher mismatch risk between the liabilities and the assets.

Pricing calculations serve multiple purposes. One of these is price-setting, which involves the calculation of the amount of money for which a contract can be sold to a customer. A second purpose has to do with pricing calculations used as a basis for corporate policy. This involves determining what the profit margin is for each contract sold. Alternatively, by determining for which price the profit is equal to zero, an institution can find the minimum price at which a product still can be sold profitably. These types of calculations are typically made when new products are being introduced by the institution. Third, pricing calculations are done for reporting and capital adequacy purposes. In this case, one uses the pricing calculations to (re)calculate the value of all assets and liabilities in the balance sheet based on current economic circumstances. As a result, one can then determine the surplus (or the coverage ratio) of assets versus liabilities. In practice, different calculation methods are often applied for the different pricing purposes. Ideally, one should use the same calculation methodology for all applications in order to ensure internal consistency.

From a scientific point of view, the problem of pricing very long-dated contracts boils down to the valuation of contracts in an incomplete markets setting. This means that we are trying to price contracts where not all of the risks can be traded (and hedged) in financial markets. In the past ten years significant progress has been made regarding this subject. Several approaches have been investigated with the common goal of trying to identify a pricing measure (or pricing kernel) that prices traded risks consistently with prices observed in the market and that also includes an extension for non-traded risks. The big problem is how to construct such an extension in a sensible way.

This panel paper first discusses the Cost-of-Capital (CoC) method proposed by the insurance industry. This method has become the *de facto* industry standard, which has also been adopted by the European Union for the Quantitative Impact Studies (QIS) in the Solvency II process. The idea behind the CoC method is that the insurance company has to hold a buffer for the non-hedgeable risks on top of the replicating portfolio. Hence, pricing consists of a "best-estimate" term plus a mark-up for the non-hedgeable risks. We discuss how to construct a time-consistent extension of the CoC methodology, and we derive an equation for how to calculate CoC prices.

A second approach discussed in this paper, is the Good Deal Bound (GDB) method. The GDB approach looks at the risk/return trade-off of non-hedgeable risks. This risk/return trade-off for the non-hedgeable risks is then compared to the risk/return trade-off that we can observe for traded assets (where it is called the *market price of risk*). The GDB method then calculates prices for non-hedgeable assets by making sure that the risk/return trade-off for any asset does not exceed a given upper bound. This upper bound is put on the prices, under the assumption that economic agents will exploit trading opportunities that are "too good" (i.e. have a risk/return trade-off that is too high).

The third approach discussed in this paper is based on model ambiguity and robustness. Although this methodology has been widely used in engineering for decades, it has attracted attention in economics only in recent years. (See, for example, the book *Robustness* by Hansen and Sargent (2007).) The fundamental premise in the robustness approach is that we are uncertain about the correct specification of our model. Therefore, when we try to make decisions (like pricing and hedging a liability) we explicitly want to take the model-uncertainty into account. This can be implemented mathematically as follows. First, specify a set of alternative models to the current base model. Then assume that we are playing against a "malevolent mother nature" that tries to pick the worst possible model out of the set of alternative models (given the decisions we have committed to). Since we are, however, aware of this, we try therefore to make decisions that are as resilient as possible given the worst-case actions of mother nature.

This Panel Paper shows that each of these three approaches converges for small time-steps in the limit to the same pricing equation. This is illustrated with several examples. Unfortunately, most of the academic literature discussed in this paper is written in rather abstract mathematical language, making the results very difficult to access for non-technical readers. One of the contributions that this Panel Paper hopes to make is to present the results in a more intuitive way.

The remainder of this paper is organised as follows. Section 3 briefly recalls the results of how to calculate prices in a complete market setting. Section 4 then analyses the other extreme, an incomplete market setting when we only have risks that are not traded in a market. In this setting we derive our main results about the mathematical equivalence of the three pricing methods under consideration. Section 5 considers the case in which we have both types of risks (traded and non-traded), and we show how the results from the previous section generalise in this case. Finally, Section 6 shows some applications of the pricing methods we have developed.

3. Pricing in Complete Markets

This section provides an overview of the theory of pricing payoffs in *complete markets*. In a complete market, every risk driver can be traded in a market—and every risk can thus be hedged. In the case of complete markets, every payoff can be priced explicitly using arbitrage-free pricing.

3.1 Binomial Tree

To illustrate the main ideas, we use a simple mathematical setting. We have a risk driver $W_x(t)$, which is a Brownian Motion. We also assume an asset price process $x(t)$, which is given by the diffusion equation

$$dx = m(t, x)\, dt + \sigma(t, x)\, dW_x, \tag{3.1}$$

where $m(t, x)$ and $\sigma(t, x)$ denote the drift and diffusion of the return process $x(t)$. We assume that $x(t)$ can be traded in a market.

For example, if we model a stock price $S(t)$ as $x(t)$, and we set $m(t, x) = \mu x$ and $\sigma(t, x) = \sigma x$, then we recover the famous Black and Scholes (1973) model.

We also assume a riskless asset B, which earns the risk-free interest rate r. The value of the riskless asset is given by

$$dB = rB\, dt. \tag{3.2}$$

We wish to consider a discretisation scheme for the return process $x(t)$ for the time period $[t, t + \Delta t]$ in the form of a binomial tree:

$$x(t+\Delta t) = x(t) + m\Delta t + \begin{cases} +\sigma\sqrt{\Delta t} & \text{with prob. } \frac{1}{2} \\ -\sigma\sqrt{\Delta t} & \text{with prob. } \frac{1}{2}, \end{cases} \tag{3.3}$$

where we have suppressed the dependence of $m(t, x)$ and $\sigma(t, x)$ on (t, x) for ease of notation.

3.2 Pricing by Replication

Suppose we have a derivative $f(t, x)$ that has a payoff that depends on x. Suppose that we know the price of the derivative at time $t + \Delta t$ for any value of $x(t + \Delta t)$; i.e. we know the function $f\big(t + \Delta t, x(t + \Delta t)\big)$. The question is: how do we determine the value for f one time-step earlier at time t?

The answer to this question was developed by Fisher Black, Myron Scholes and Robert Merton in the early 1970s. They used the notion of *pricing by replication*, which won Scholes and Merton the Nobel Prize in economics in 1997. The idea works as follows. Suppose we buy a portfolio of D units of the risky asset x and an amount B invested in the risk-free asset. Then at time t this portfolio has value $(Dx(t) + B)$. At time $t + \Delta t$ the portfolio has two possible values (using the binomial discretisation (3.3))

$$\begin{cases} Dx_+ + e^{r\Delta t}B & \text{with prob. } \frac{1}{2} \\ Dx_- + e^{r\Delta t}B & \text{with prob. } \frac{1}{2}, \end{cases} \tag{3.4}$$

where $x_\pm$ is shorthand notation for $x_\pm := x(t) + m\Delta t \pm \sigma\sqrt{\Delta t}$.

Given the binomial discretisation for $x(t + \Delta t)$, the derivative $f()$ has two possible values at time $t + \Delta t$: either $f_+ := f(t + \Delta t, x_+)$ or $f_- := f(t + \Delta t, x_-)$. If we want to match the values of our portfolio $(Dx(t) + B)$ with the value of our derivative f at time $t + \Delta t$, we have to solve the following system of equations:

$$\begin{cases} Dx_+ + e^{r\Delta t}B = f_+ \\ Dx_- + e^{r\Delta t}B = f_-. \end{cases} \tag{3.5}$$

The solution is given by $D = \frac{f_+ - f_-}{x_+ - x_-}$ and $B = e^{-r\Delta t}\frac{f_- x_+ - f_+ x_-}{x_+ - x_-}$.

We have now explicitly constructed the replicating portfolio for the derivative $f()$. The brilliant insight of Black, Scholes and

Merton was that the price of the derivative $f(t, x)$ at time t must be equal to the price $(Dx(t) + B)$ of the replicating portfolio. If this would not be the case, there would be an arbitrage opportunity: two different prices for two instruments that have exactly the same value at time $t + \Delta t$. Therefore, we calculate the value at time t of the derivative $f(t, x)$ by evaluating $(Dx(t) + B)$:[1]

$$f(t,x) = \tfrac{1}{2}\left(1 - r\Delta t - \left(\tfrac{m(t,x)-rx}{\sigma(t,x)}\right)\sqrt{\Delta t}\right) f_+ + \tfrac{1}{2}\left(1 - r\Delta t + \left(\tfrac{m(t,x)-rx}{\sigma(t,x)}\right)\sqrt{\Delta t}\right) f_-. \quad (3.6)$$

The term $\frac{m(t,x)-rx}{\sigma(t,x)}$ measures the excess return above the risk-free rate of the risky asset divided by the standard deviation of the risky asset. This ratio is known as the *market price of risk*, which will be denoted by

$$\lambda(t,x) := \frac{m(t,x) - rx}{\sigma(t,x)}. \quad (3.7)$$

The market price of risk is a positive quantity, as the return $m(t, x)$ on a risky asset is larger than the return rx on a risk-free asset. The market price of risk will turn out to be quite crucial in the rest of our story.

3.3 Deflator Pricing

The binomial pricing equation (3.6) admits several different interpretations. The first interpretation worth noting is the

[1] Please note that the equality is not exact, as we have omitted terms of higher order than Δt. On the other hand, the binomial approximation (3.3) of the process x is also not exact. However, when we consider the limit for $\Delta t \to 0$ then all of the approximations converge to the correct answer.

interpretation as a pricing operator with respect to a *deflator* or *pricing kernel*.

Interpret (3.6) as taking an expectation, using the original binomial probabilities $\frac{1}{2}$ and $\frac{1}{2}$, of the adjusted derivative values $\left(1 - r\Delta t - \lambda(t,x)\sqrt{\Delta t}\right) f_+$ and $\left(1 - r\Delta t + \lambda(t,x)\sqrt{\Delta t}\right) f_-$. The adjustment factor is different for the "plus" and the "minus" state of the world; hence, the adjustment factor is a random variable. In fact, we can interpret the adjustment factor as the binomial discretisation of the random variable $\xi(t)$, which is given by

$$d\xi = -r\xi\, dt - \lambda(t,x)\xi\, dW_x. \tag{3.8}$$

The random variable $\xi(t)$ is known as the *deflator* or the *pricing kernel*. Note that the volatility of the pricing kernel is equal to minus the market price of risk $-\lambda(t,x)$. The minus sign indicates than whenever $W_x(t)$ decreases, then $\xi(t)$ increases. This has the effect of putting more weight on "bad" outcomes of the process x (i.e. low values of W_x) than on "good" outcomes (high values of W_x).

Using the deflator interpretation, re-write the pricing equation (3.6) as

$$f(t,x) = \frac{\mathbb{E}_t[\xi(t+\Delta t) f(t+\Delta t)]}{\xi(t)}, \tag{3.9}$$

where $\mathbb{E}_t[\,]$ denotes the expectation operator conditional on the information available at time t, in particular the information that the process $x(t)$ at time t is equal to the value x.

3.4 Risk-Neutral Pricing

An alternative interpretation of the pricing equation (3.6) is as a discounted *risk-neutral* expectation. Instead of using the original binomial probabilities and adjusting the payoff (as was done in Section 3.3), we can adjust binomial probabilities and leave the payoff unchanged. When doing this, we must ensure that the new probabilities are created still sum to 1. The adjusted binomial probabilities are given by $\frac{1}{2}\left(1 - r\Delta t - \lambda(t,x)\sqrt{\Delta t}\right)$ and $\frac{1}{2}\left(1 - r\Delta t + \lambda(t,x)\sqrt{\Delta t}\right)$. However, when these two numbers are added together we get $(1 - r\Delta t)$, which is less than 1. An elegant way to adjust the weight-factors is by re-writing[2] them as $e^{-r\Delta t}\frac{1}{2}\left(1 - \lambda(t,x)\sqrt{\Delta t}\right)$ and $e^{-r\Delta t}\frac{1}{2}\left(1 + \lambda(t,x)\sqrt{\Delta t}\right)$. Now re-write the pricing equation (3.6) as

$$f(t,x) = \mathbb{E}_t^{\mathbb{Q}}\left[e^{-r\Delta t} f(t+\Delta t)\right], \tag{3.10}$$

where $\mathbb{E}_t^{\mathbb{Q}}[\,]$ denotes the conditional expectation operator with respect to the adjusted binomial probabilities

$$q = \tfrac{1}{2}\left(1 - \lambda(t,x)\sqrt{\Delta t}\right) \tag{3.11a}$$

$$1 - q = \tfrac{1}{2}\left(1 + \lambda(t,x)\sqrt{\Delta t}\right) \tag{3.11b}$$

for the “plus” and “minus” state, respectively.

Like in Section 3.3, the adjusted probabilities q and $(1-q)$ put more weight on the “minus” state compared to the original binomial probabilities of $\frac{1}{2}$. In fact, the expectation of $x(t+\Delta t)$ is calculated using the adjusted binomial probabilities, we find that $\mathbb{E}_t^{\mathbb{Q}}[x(t+\Delta t)] = x(t)e^{r\Delta t}$. Hence, under the adjusted probabilities, the process $x(t)$ grows with the risk-free rate r, which is lower than the true growth rate $m(t,x)$ of the process $x(t)$.

Wrapping up this section, we would like to stress that the pricing formulæ (3.9) and (3.10) will give exactly the same outcome. They

---

[2]Again, we are ignoring terms of higher order than Δt.

are nothing more than different representations of the same binomial pricing equation (3.6).

3.5 Partial Differential Equation

Up until now, we have focused extensively on the pricing of a derivative contract for one single time-step $[t, t + \Delta t]$. Obviously, we are ultimately interested in pricing contracts of the whole life $[0, T]$. One method for converting a "one-step" pricing formula into a "whole interval" pricing formula is to apply the one-step pricing formula using a backward-induction procedure. In other words, start at the end-date T and then move backward in time by repeatedly applying the one-step pricing formula for each time-step Δt. The backward-in-time nature of this algorithm ensures that at each time t during the calculation the valuation formula accounts for all of the remaining uncertainty until the maturity date T. Hence, use of backward-induction makes it possible to construct a pricing operator that is *time-consistent*.

This subsection considers the limit for $\Delta t \to 0$. Assume that $f(t + \Delta t, x)$ is sufficiently smooth in t and x, such that we can apply for all values of (t, x) the Taylor approximation

$$\begin{aligned} f\big(t + \Delta t, x + h\big) = f\big(t + \Delta t, x\big) + \\ f_x\big(t + \Delta t, x\big)h + \tfrac{1}{2} f_{xx}\big(t + \Delta t, x\big)h^2 + \mathcal{O}(h^3), \end{aligned} \quad (3.12)$$

where subscripts on f denote partial derivatives. If we apply the binomial approximation (3.3) for the process $x(t)$, this yields $f_+ = f\big(t + \Delta t, x + m\Delta t + \sigma\sqrt{\Delta t}\big)$ and $f_- = f\big(t + \Delta t, x + m\Delta t - \sigma\sqrt{\Delta t}\big)$.

If we substitute the Taylor approximation (3.12) into these expressions and then substitute into the one-step binomial pricing equation (3.6), this yields

$$0 = f\big(t + \Delta t, x\big) - f\big(t, x\big) + rxf_x\big(t + \Delta t, x\big)\Delta t + \\ \tfrac{1}{2}\sigma^2 f_{xx}\big(t + \Delta t, x\big)\Delta t - rf\big(t + \Delta t, x\big)\Delta t + \mathcal{O}(\Delta t^2). \tag{3.13}$$

Note that due to the adjustment factors in the binomial pricing equation (3.6), the true growth rate $m(t, x)$ has disappeared from (3.13), and has been replaced by the risk-free growth rate rx that multiplies the term $f_x\big(t + \Delta t, x\big)\Delta t$.

Now, divide by Δt and take the limit for $\Delta t \to 0$, which yields

$$f_t + rxf_x + \tfrac{1}{2}\sigma^2 f_{xx} - rf = 0, \tag{3.14}$$

where we have suppressed the dependence on (t, x) to lighten the notation. Equation (3.14) is a *partial differential equation* (pde) for the derivative price $f(t, x)$. The price of any derivative on the underlying process $x(t)$ is a solution to (3.14) with respect to a boundary condition $f\big(T, x(T)\big)$ that defines the payoff as a function of $x(T)$ at the maturity date T.

3.6 Literature Overview

The literature on pricing in complete markets has been developed and extended since the 1970s. It started with the seminal papers by Black and Scholes (1973) and Merton (1973). The binomial tree pricing model was developed by Cox et al. (1979). The connection to martingale measures was developed by Harrison and Kreps (1979) and Harrison and Pliska (1981). Significant generalisations were achieved for more general stochastic processes by Delbaen and Schachermayer (1994). For an introduction to pricing derivatives, see the textbook by Hull (2009).

4. Pricing non-hedgeable Risk

Section 3 considered the case of a complete market—one in which the underlying risk driver can be traded. This section considers the opposite case, where the underlying risk drivers *cannot* be traded. This makes it no longer possible to construct a replicating portfolio, which was the underlying basis for the pricing method in Section 3. Instead, we have to define a pricing operator to determine the value of a payoff.

This has been the subject of study of actuaries for a long time. The basic idea for a pricing operator is to use the expected value of the payoff minus a "penalty term" that depends on the risk of the payoff. Many different pricing operators have been proposed (for an overview, see Gerber (1979), Deprez and Gerber (1985), Young (2004a) and the textbook by Kaas et al. (2008)). Actuaries make a distinction between two main classes of pricing operators. One class uses standard deviation as a measure of risk, the other class uses variance as the measure of risk. This Panel Paper focusses on pricing operators of the first class. Section 4.6 provides a literature overview of alternative pricing methods that belong to the second class.

4.1 Binomial Tree

Let us introduce a new risk driver W_y, which is a Brownian Motion. We also assume a process $y(t)$, which is given by

$$dy = a(t, y)\, dt + b(t, y)\, dW_y. \tag{4.1}$$

Note that we assume that $y(t)$ *cannot* be traded in a market. Like in Section 3, here we also consider the binomial discretisation for the process y for a time-step Δt as

$$y(t+\Delta t) = y(t) + a\Delta t + \begin{cases} +b\sqrt{\Delta t} & \text{with prob. } \frac{1}{2} \\ -b\sqrt{\Delta t} & \text{with prob. } \frac{1}{2}. \end{cases} \tag{4.2}$$

Furthermore, we want to consider a derivative $g(t, y)$ which has a payoff that depends on y.

4.2 Cost-of-Capital Pricing

A pricing principle that is widely used in practice is the Cost-of-Capital (CoC) Principle. This was introduced by the Swiss insurance supervisor as a part of the method to calculate solvency capitals for insurance companies (see, e.g. Keller and Luder, 2004). In recent years, the CoC method has been widely adopted by the insurance industry in Europe, and has also been prescribed as the standard method by the European Insurance and Pensions Supervisor for the Quantitative Impact Studies (see EIOPA, 2010).

The CoC approach is based on the following economic reasoning. First consider the “expected loss” $\mathbb{E}[g(T, y)]$ of the insurance claim $g(T, y)$ as a basis for pricing. But this is not enough; the insurance company also has to hold a capital buffer against the “unexpected loss”. This buffer is calculated as a Value-at-Risk over a time horizon (typically one year) and a probability threshold q (usually 0.995, or even higher). The unexpected loss is then calculated as $\mathrm{VaR}_q[g(T, y)]$. The capital buffer is borrowed from the shareholders of the insurance company (i.e. the buffer is subtracted from the surplus in the balance sheet). Given the very high confidence level, in many cases the buffer can be returned to the shareholders—although there is a chance that the capital buffer is needed to cover an unexpected loss. Hence, the shareholders require compensation for this risk in the form of a *cost-of-capital* premium. This cost-of-capital premium needs to be included in the pricing of the insurance contract. If we denote the cost-of-capital by δ, then the CoC pricing equation is given by

$$g(t,y) = e^{-r(T-t)}\left(\mathbb{E}_t[g(T,y)] + \delta \mathrm{VaR}_{q,t}[g(T,y)]\right). \quad (4.3)$$

4.2.1 Time-Consistency

The pricing method defined in equation (4.3) has a methodological problem: it is defined for a one-year horizon (i.e. $t = T - 1$). An important practical question is, how to extend the pricing formula to longer horizons? The approach adopted by the industry is a simple rule-of-thumb; see Keller and Luder (2004). The idea is as follows: you first make a projection of the contract value along the "best-estimate path" of the risk driver given by $\mathbb{E}_t\left[g(T,y(T)) \mid y(t) = \mathbb{E}_0[y(t)]\right]$ for all $0 \le t \le T$. Then, at annual points ($t = 1, 2, 3, \ldots$) you approximate the Value-at-Risk (VaR) by considering the impact of a 99.5% shock for the risk driver from the best-estimate path. Finally, the present value of all shocks is added and multiplied by δ.

Let us consider an example. For ease of exposition assume $r = 0$. Suppose there is a two-year product with a payoff $e^{bW_y(2)}$. The best-estimate path is given by $\mathbb{E}_t[e^{bW_y(2)} | W_y(t) = 0] = e^{\frac{1}{2}b^2(2-t)}$ for $t = 0, 1, 2$. A one-year 99.5% worst-case shock on $W_y(t)$ is given by an increase in value to $W_y(t) + 2.58$. Hence, the Value-at-Risk in year t is approximated by applying the one-year shock to the best-estimate path as $e^{\frac{1}{2}b^2(2-t)}(e^{2.58b} - 1)$. Finally, the CoC price for this two-year product would be calculated as

$$e^{\frac{1}{2}b^2 2} + \delta(e^{\frac{1}{2}b^2} + 1)(e^{2.58b} - 1). \quad (4.4)$$

If $b = 50\%$ and $\delta = 6\%$ then we calculate a price of 1.62.

A disadvantage of the "best-estimate path method" is that the dynamics of the risk driver $y(t)$ are completely ignored for the VaR calculation. If we move one year ahead in time, then the risk driver will be at the value $y(1)$, which will differ from the best estimate value $\mathbb{E}_0[y(1)]$. Hence, the CoC price of the product at time $t = 1$ is based on a different best-estimate path than the calculation at $t = 0$. Therefore, the "best-estimate path method" used by the industry is *not* time-consistent.

How can we obtain a time-consistent version of the CoC pricing operator? One approach (similar to that taken for complete markets in Section 3) to use a backward-induction method. In fact, Jobert and Rogers (2008) prove that every time-consistent valuation operator can be obtained by backward-induction of a one-step pricing operator. Returning to the example, given the payoff at $T = 2$, we can calculate the price at time 1 conditional on the value of $W_y(1)$ as

$$e^{bW_y(1)+\frac{1}{2}b^2} + \delta\big(e^{b(W_y(1)+2.58)+\frac{1}{2}b^2} - e^{bW_y(1)+\frac{1}{2}b^2}\big).$$

This expression can be simplified to

$$e^{bW_y(1)+\frac{1}{2}b^2}\left(1 + \delta(e^{2.58b} - 1)\right).$$

Given the price at time 1, which is now an explicit function of $W_y(1)$, we can again calculate the CoC price at $t = 0$. This leads to the formula

$$e^{\frac{1}{2}b^2 2}\left(1 + \delta(e^{2.58b} - 1)\right)^2. \tag{4.5}$$

If we take again $b = 50\%$ and $\delta = 6\%$, we find a price of 1.72.

When we compare the price (4.4) of the "best-estimate path" method with the backward-induction price (4.5), then we immediately see the effect of the best-estimate path approximation. In (4.4) one adds the terms $\delta(e^{2.58b} - 1)$ and

$\delta e^{\frac{1}{2}b^2}(e^{2.58b}-1)$ to the price $e^{\frac{1}{2}b^2 2}$. Whereas the backward-induction method explicitly takes the “capital-on-capital” effect into account by multiplying the price $e^{\frac{1}{2}b^2 2}$ twice with the factor $(1+\delta(e^{2.58b}-1))$, the inclusion of the “capital-on-capital” effect leads to a time-consistent pricing operator.

4.2.2 Partial Differential Equation

As a final step in our argument, we change the length of the time-step in the Cost-of-Capital pricing operator from one year to Δt, and consider the limit for $\Delta t \to 0$. Note that when comparing Value-at-Risk quantities at different time-scales Δt, these have to be scaled back to a per annum basis; this is done by dividing the VaR term by[3] $\sqrt{\Delta t}$. Then, consider that the cost-of-capital δ behaves like an interest rate: it is the compensation the insurance company needs to pay to its shareholders for borrowing the buffer capital over a certain period. The cost-of-capital is expressed as a percentage per annum; hence, over a time-step Δt the insurance company has to pay a compensation of $\delta\Delta t$ per € of buffer capital. This yields a “net scaling” of $\delta\Delta t/\sqrt{\Delta t} = \delta\sqrt{\Delta t}$.

For a single time-step Δt, this yields the following expression for the CoC price:

$$g\big(t, y(t)\big) = e^{-r\Delta t}\Big(\mathbb{E}_t[g\big(t+\Delta t, y(t+\Delta t)\big)] + \delta\sqrt{\Delta t}\,\mathrm{VaR}_{q,t}[g\big(t+\Delta t, y(t+\Delta t)\big)]\Big). \quad (4.6)$$

[3]The scaling by $\sqrt{\Delta t}$ is a result of using Brownian Motion to describe the evolution of risk. More general stochastic process (such as Lévy processes) may require different scaling factors, but this is beyond the scope of the current paper.

With this pricing operator for a Δt-step we apply the backward-induction method to determine the time-consistent CoC price for a payoff $g(T, y)$ at time T, and take the limit $\Delta t \to 0$.

Note that for small Δt, the variance at time t of the process $g(t + \Delta t, y)$ is given by $b^2 g_y^2 \Delta t$. Furthermore, in a diffusion setting, for small Δt, all risks are very close to a normal distribution. Hence, the VaR at time t is closely approximated by k times the standard deviation $kb|g_y|\sqrt{\Delta t}$, where the constant k is given by the inverse cumulative normal distribution of the VaR confidence level q (i.e. $k = \Phi^{-1}(q)$). Given that $g(t + \Delta t, y)$ is sufficiently smooth to be twice continuously differentiable in y, we can then (similar to the manipulations in Section 3.5) substitute the Taylor approximation of the function $g(t + \Delta t, y)$ for the binomial approximation (4.2) into (4.6), divide by Δt and take the limit for $\Delta t \to 0$, which yields[4] the following partial differential equation (pde) for the price operator $g(t, y)$:

$$g_t + a g_y + \tfrac{1}{2} b^2 g_{yy} + \delta k b |g_y| - rg = 0. \tag{4.7}$$

A comparison of the pricing equation (4.7) and the complete market pricing equation (3.14) reveals two important differences. First, note the additional term $\delta k b |g_y|$. This is the “penalty term” that the Cost-of-Capital method adds for writing the non-hedgeable claim $g(T, y)$. Second, it seems that we have not changed the drift term a for the process y in the pricing equation. But this is not entirely true. In fact, whenever the payoff $g(T, y)$ is monotonous in $y(T)$, then the sign of g_y is unique, and the two terms depending on g_y can be added together to obtain $(a \pm \delta k b) g_y$. Therefore, the CoC price $g(t, y)$ can be

[4] For a more elaborate derivation, see Bayraktar and Young (2008) or Delong (2011).

represented with respect to the "risk-adjusted" process y:

$$dy = \big(a(t, y) \pm \delta k b(t, y)\big)\, dt + b(t, y)\, dW_y, \qquad (4.8)$$

where the sign of $\pm\delta k$ is determined by the sign of g_y. This allows for a very nice interpretation of the pricing equation (4.7). When pricing an non-hedgeable claim $g(T, y)$, we adjust the "best-estimate" drift of the process y in a "conservative" direction. In other words, the drift is adjusted upwards or downwards by $\delta k b(t, y)$ depending on the sign of g_y. Making the price more conservative by adjusting the drift is a time-honoured actuarial practice known as *prudence*.

Revisiting the example from Section 4.2.1, recall there was a payoff of $e^{y(2)}$ with $r = 0$, $a = 0$ and $b = 0.50$. This payoff is monotonically increasing in y and this claim can be priced by adjusting the drift of y upward to δkb. Hence, we calculate a price at time 0 of $e^{(\delta kb + \frac{1}{2}b^2)2}$. If we take $\delta = 6\%$, $k = 2.58$ and $b = 0.50$, then the price at time 0 is 1.50.

This section concludes by noting that that Cost-of-Capital method suffers from a weakness: there is relatively little economic justification for choosing the correct values of δ and k. The report CRO-Forum (2006) considers a wide variety of arguments that lead to a wide range of possible parameter values. In the end, the CRO-Forum recommends setting $\delta = 0.06$ and $k = \Phi^{-1}(0.995) = 2.58$, leading to a total factor of $\delta k = 0.15$.

4.3 Good Deal Bound Pricing

A very different approach on pricing in incomplete markets was introduced by Cochrane and Saá-Requejo (2000). It is based on the following idea. Suppose you are offered the opportunity to enter the following lottery: with a probability of $\frac{1}{2}$ you get a payoff of 1000, or 1. The initial price of the lottery is 2. In terms of the theory developed in Section 3, this is not an arbitrage opportunity. However, it does represent a "very good deal". We get something with an expected value of 500.50 for a price of 2. There is, however, risk involved: the expected value is $\frac{1}{2}1000 + \frac{1}{2}1 = 500.50$, and the standard deviation is $\sqrt{\frac{1}{2} * (1000 - 500.50)^2 + \frac{1}{2} * (1 - 500.50)^2} = 499.50$. But you have to be extremely risk-averse to bring the price you are willing to pay from 500 down to below 2, in order to not participate in this lottery.

The tools developed in Section 3 can be used to calculate the price for this lottery by adjusting the probabilities of the outcomes. In this example we have to solve[5] for the adjusted probability q the equation $2 = q * 1000 + (1 - q) * 1$, which leads to $q = 1/999$ and $(1 - q) = 998/999$. Comparing the ratio between the adjusted probabilities q and the original probabilities $\frac{1}{2}$, like in equation (3.11), we see that the ratio is extremely large—almost a factor 2000 in this example. Hence, "extremely good deals" imply very large probability ratios. Another way of looking at this is to look at the factor $\lambda()$ that was introduced in (3.8), which is the volatility of the deflator $\xi()$. Section 3.4 established that the deflator volatility $\lambda()$ can also be interpreted as the market price of risk, if it was in a complete market. Hence, for "extremely good deals" the market price of risk $\lambda(t, y)$ is very large.

This brings us to the idea of *Good Deal Bounds*. In an incomplete market setting, we cannot trade in the underlying risk driver y. Hence, we cannot calibrate the martingale measure to the prices of traded assets. On the other hand, it is unrealistic to assume that agents in the economy will leave "extremely good deals"

[5]For ease of exposition we ignore the effect of discounting in this example.

unexploited. Cochrane and Saá-Requejo (2000) introduced the idea of putting an upper bound on the deflator volatility $\lambda()$ to distinguish "normal deals" from "extremely good deals". Furthermore, Cochrane and Saá-Requejo (2000) proposed using market prices of risk that we can observe for traded risks as a benchmark for non-traded risks.

Suppose that we put an upper bound κ on the deflator volatility. This makes it possible to search for the upper and lower bounds on the price for a derivative $g(T, y)$ by considering all pricing deflators with a volatility less than or equal to κ. These upper and lower bounds for the price represent the "ask" and "bid" prices for an agent with good deal bound κ.

This may all sound quite complicated, but it allows us to exploit the structure between deflators, risk-neutral probabilities, and the drift of the risk driver already explored in Section 3. Let us return to the binomial discretisation (4.2) of the process y. By putting a bound κ on the deflator volatility, we infer from equation (3.11) that we are considering adjusted probabilities in the range

$$q \in \left[\tfrac{1}{2}\left(1 - \kappa\sqrt{\Delta t}\right), \tfrac{1}{2}\left(1 + \kappa\sqrt{\Delta t}\right)\right]. \tag{4.9}$$

But, changing the binomial probabilities is equivalent to changing the drift of the process y. Hence, alternatively, we can also say that we consider specifications for the stochastic process y where the adjusted drift $a^*(t, y)$ is somewhere in the range

$$a^*(t, y) \in \left[a(t, y) - \kappa b(t, y), a(t, y) + \kappa b(t, y)\right]. \tag{4.10}$$

Using the derivation from Section 3.5, we infer that any price of a derivative $g(t, y)$ that falls within the Good Deal Bounds is described by the partial differential equation (pde)

$$g_t + a^* g_y + \tfrac{1}{2} b^2 g_{yy} - rg = 0, \tag{4.11}$$

where a^* is taken from the interval (4.10).

When seeking the highest and lowest prices that are "on the edge" of the good deal bound interval, we have to find the drift a^* that minimises or maximises the price $g(t, y)$ for each time-step. For example, when we want to maximise the price $g(t, y)$, then we should either put a^* at the upper bound $a(t, y) + \kappa b(t, y)$ whenever $g_y(t, y) > 0$ or put a^* at the lower bound $a(t, y) - \kappa b(t, y)$ whenever $g_y(t, y) < 0$. Therefore, we can represent the good deal bound price $g(t, y)$ with respect to the "risk-adjusted" process y:

$$dy = \big(a(t, y) \pm \kappa b(t, y)\big)\, dt + b(t, y)\, dW_y, \qquad (4.12)$$

where the sign of $\pm\kappa$ is determined the sign of g_y.

Note that the structure of the risk-adjusted process (4.12) is exactly the same as the structure of the Cost-of-Capital pricing process (4.8), provided we take $\kappa = \delta k$.

On the other hand, given that we have the interpretation of κ as an upper bound for the deflator volatility, which in traded markets is equal to the market price of risk, this information can be used to get more guidance on setting κ. Considering equity markets , then we can calculate the market price of risk. For typical equity markets (see e.g. Dimson et al. (2002)) we see an excess return above the risk-free rate of around 4%, and a volatility of around 16%, leading to a market price of risk of approximately $4/16 = 0.25$. From this calculation we infer that the upper bound κ should be larger than 0.25. Note that in this light the value $\delta k = 0.15$ implied by the Cost-of-Capital method seems to be on the low side.

4.4 Model Ambiguity & Robustness

This subsection introduces a third perspective for pricing contracts in incomplete markets. This is the notion of *Model Ambiguity*. This means that we explicitly take into consideration that the mathematical models we use to describe the world are not exact, but may be misspecified.

Model ambiguity can be illustrated as follows. Suppose we try to estimate the expected return of investing in an equity index (say, the Standard & Poor's (S&P) index). Historical observations can then be used to estimate the expected return, but the estimate of the expected return will then be subject to estimation error. It turns out that since equity returns are relatively volatile, it is very difficult to obtain an accurate estimate for the expected return. Assume that the volatility of the S&P index is around 16%, and that we use 25 years of data. Then the standard error for the estimate of the expected return is $16\%/\sqrt{25} = 3.2\%$. Suppose that the estimate of the expected return is 8%; then the 95% confidence interval for this estimate is $[8\% - 1.96 * 3.2\%, 8\% + 1.96 * 3.2\%] = [1.7\%, 14\%]$. Even if we would use 100 years of historical data, our 95% confidence interval is still $[4.9\%, 11\%]$. Using more years of historical data will give us a more accurate answer only, if the data-generating process has remained the same during the entire period. It is highly questionable whether the economy of 100 years ago is representative of today's economy. Thus, it is clearly very difficult to obtain an accurate estimate of something as simple as the expected return of an equity index. The same observation is also true for the expected increase in human longevity: actuaries have been constantly revising their projections about forecasts of human longevity in the last 20 years.

Suppose we accept the impossibility of accurately "knowing" the correct model specification. How can we deal with this *model ambiguity*? One approach is to assume that economic agents are concerned about making bad decisions based on misspecified models. To deal with this problem, assume that agents resort to *robust optimisation* methods. This means that agents try to make their decisions in such a way that they explicitly incorporate the fact that the "true" model of Mother Nature may deviate from the mathematical model used by the agent for decision making. The notion of model ambiguity and robust optimisation in economics has been made popular in recent years by Hansen and Sargent (see, for example, their book *Robustness*, Hansen and Sargent (2007)).

How can the notion of robust optimisation be implemented in our setting? We start by making some strong simplifying assumptions. First, assume that the true model for the process $y(t)$ is of the form (4.1). The only uncertainty that we have concerns the correct specification of the drift term $a(t, y)$. Hence, we assume that we know the correct specification of the diffusion term $b(t, y)$. These are, of course, very strong assumptions indeed, but given the difficulties in estimating even a "simple" parameter as the expected return, this seems like a good starting point.[6]

Second, assume that the agent is able to specify a confidence interval of "reasonable" values for the drift $a^*(t, y)$. We will return later to the question of how to specify a confidence interval of reasonable values. In the one-dimensional case being considered in this section, a confidence interval for the drift has the form $a^*(t, y) \in [a_L, a_H]$. Another way of representing a confidence interval is to say we have a point estimate $a(t, y)$ that is located in the centre of the confidence interval, and a width of

[6]For a more elaborate justification of considering only uncertainty in the drift, see Hansen and Sargent (2007, Chapter 1).

the confidence interval given by 2ν times the standard deviation $b(t, y)$ of the process $y(t)$. This leads to the representation

$$a^*(t, y) \in [a(t, y) - \nu b(t, y), a(t, y) + \nu b(t, y)]. \quad (4.13)$$

Note that this confidence interval is of exactly the same form as the good deal bound equation (4.10). However, the parameter ν now has the interpretation as being the width of the confidence interval for $a^*(t, y)$.

In this setup we can achieve a robust notion of pricing a derivative $g(T, y)$ by calculating the expectation of $g(T, y)$ under the "worst" model specification for the non-hedgeable process y. In particular, for an insurance company that has written the claim $g(T, y)$ payable at time T, the "worst" model specification is that choice for the drift $a^*(t, y)$ in the interval (4.13) that maximises the value of the expectation. Using exactly the same argumentation as in Section 4.3, we find that the robust price $g(t, y)$ can be represented by taking the expectation with respect to the "worst case" process y:

$$dy = \big(a(t, y) \pm \nu b(t, y)\big)\, dt + b(t, y)\, dW_y, \quad (4.14)$$

where the sign of $\pm\nu$ is determined the sign of g_y.

Given our interpretation of ν as the width of the confidence interval, how can we determine ν? In other words: how can we establish an interval of "reasonable" values for $a^*(t, y)$? We can offer two (closely related) arguments. The first argument is that historical data can be used to estimate the parameter a. The confidence interval from the parameter estimate the becomes the measure for the interval of "reasonable" values for a^*. Using a 95% confidence interval based on 25 years of historical data yields $\nu = 1.96/\sqrt{25} = 0.39$. For 50 years of historical data we obtain

$\nu = 1.96/\sqrt{50} = 0.28$. The second argument is to consider the question: which alternative model specifications are statistically indistinguishable from our current model, given the available data? The answer (in the case of uncertainty in the mean) is given by all values for the drift *a* that are in the 95% confidence interval (4.13).

Comparing the values for ν of 0.39 or 0.28 to the lower bound of 0.25 found for κ reveals these values are nicely in agreement with each other. Furthermore, we arrive (once again) at the conclusion that the value $\delta k = 0.15$ used by the insurance industry is on the low side.

4.5 A New Value of the Cost-of-Capital

To summarise our discussion on the Cost-of-Capital: the conclusions drawn both in Section 4.3 and in the previous section indicate that the CoC parameter $\delta k = 0.15$ currently used by the insurance industry seems too low.

A value of $\delta k = 0.30$ seems much more appropriate when this is compared to the values implied by the Good Deal Bound and the Model Ambiguity methods. Setting $\delta k = 0.30$ and using $k = \Phi^{-1}(0.995) = 2.58$ implies a cost-of-capital parameter $\delta = 12\%$, which is basically doubling the current value of 6% proposed by the insurance industry.

Setting $\nu = 0.30$ in the Model Ambiguity method, this corresponds to using $(1.96/0.30)^2 = 43$ years of historical data to estimate the mean of the process $y(t)$. This also seems a reasonable tradeoff between using as much historical data as possible, without going back so far in time that it becomes hard to believe that the data is still representative for today's economy.

4.6 Alternative Approaches

This section presents various alternative approaches to the theory developed in this paper so far.

4.6.1 Variance Pricing & Utility Indifference Pricing

A large body of literature focuses on utility indifference pricing. The roots can be traced back to Hodges and Neuberger (1989). The idea is that the assumption is made that the behaviour of agents can be described by a utility function, then a *utility indifference price* for accepting an (non-hedgeable) claim can be found. For exponential utility functions, quite explicit results can be found (see Zariphopoulou (2001); Young and Zariphopoulou (2002); Musiela and Zariphopoulou (2004); Hugonnier et al. (2005); Hu et al. (2005); Musiela and Zariphopoulou (2009b); Henderson (2002, 2005) and Henderson and Hobson (2009)). Also for power utility functions, partial results are known, see Hobson (2004); Monoyios (2006). For a general overview, see the book *Indifference Pricing* by Carmona (2009).

A big disadvantage of utility-based pricing is that it depends on the specification of the utility function at a specific horizon T. This introduces an artificial dependency in the pricing on the horizon T. Attempts to resolve this issue were proposed by Henderson and Hobson (2007) and Musiela and Zariphopoulou (2007, 2009a).

4.6.2 Strong Time-Consistency

Our derivations have used conditional expectations that are sequentially evaluated using backward-induction arguments. This leads to the pricing pde's we have found in equation (4.11). Use of backward-induction techniques allows us to construct pricing methods that are *strongly time-consistent*; see Hardy and Wirch (2005) and Jobert and Rogers (2008). However, the concept of strong time-consistency for pricing methods is not uncontroversial. See Roorda et al. (2005); Roorda and Schumacher (2007) for a discussion.

4.6.3 Bayesian Approach

Finally, note that as an alternative to robust optimisation, it is possible to use Bayesian methods. In the Bayesian approach the uncertainty about the model specification is specified in the form of prior and posterior probability distributions on the parameter space. Portfolio optimisation and pricing is then carried out by "averaging" over the parameter space (i.e. averaging over the different alternative model specifications). For a discussion and examples, see Lutgens (2004); Lutgens and Schotman (2010).

4.7 Literature Overview

As mentioned in the text, the Cost-of-Capital (CoC) approach was originally proposed by the insurance industry (see CRO-Forum (2006)), based on ideas put forward by the Swiss insurance supervisor in the so-called Swiss Solvency Test (SST) (see Keller and Luder (2004)). For a critical discussion on the risk measure implied by the SST see Filipovic and Vogelpoth (2008). The CoC method was adopted by the European Union as the standard method for the calculations in the Quantitative Impact Studies of the Solvency II process; see CEIOPS (2008); EIOPA (2010).

Good Deal Bound pricing has been introduced by Cochrane and Saá-Requejo (2000). Their basic ideas were extended by Černý and Hodges (2002), Becherer (2009), Björk and Slinko (2006) andKlöppel and Schweizer (2007b). The connections between Good Deal Bound pricing and the numéraire portfolio (which we have called the stochastic discount factor, see eq. (3.8)) have been explored by Becherer (2001, 2009), Karatzas and Kardaras (2007), Christensen and Larsen (2007) and Delong (2011).

Jaschke and Küchler (2001) highlighted the connections between Good Deal Bound pricing and the rich theory of coherent risk measures. Coherent risk measures were introduced by Artzner et al. (1999, 2007). Later, this has been extended to the more general class of convex risk measures by Föllmer and Schied (2002) and Cheridito et al. (2005). The connection between convex risk measures and pricing can be found in Cvitanić and Karatzas (1999), Carr et al. (2001), Frittelli and Rosazza Gianin (2002), Detlefsen and Scandolo (2005), Klöppel and Schweizer (2007a), Stadje (2010), Delbaen et al. (2010), and the papers by Cherny (2007, 2009, 2010). In the actuarial literature, see Denuit et al. (2006) and Goovaerts and Laeven (2008).

Model Ambiguity and robustness were made popular in economics by Hansen and Scheinkman (1995), Hansen and Sargent (2001), Cagetti et al. (2002), Cont (2006), Hansen et al. (2006) and Hansen and Sargent (2007). However, ideas for robustness in statistics are much older, and date at least back to Huber (1981) (for a new edition, see Huber and Ronchetti, 2009). Also, several authors have applied robustness ideas to portfolio optimisation; see Kirch (2002), Goldfarb and Iyengar (2003), Maenhout (2004), Coleman et al. (2007), Gundel and Weber (2007), Rogers (2009), Föllmer et al. (2009), Iyengar and Ma (2010) and Kerkhof et al. (2010). Another branch of literature studies the notion of model ambiguity on decisions that economic agents make; see Duffie and Epstein (1992a), Duffie and Epstein (1992b), Chen and Epstein (2002), Maccheroni et al. (2006), Epstein and Schneider (2008) and Riedel (2009).

5. Combining Hedgeable and Non-Hedgeable Risk

This section pushes our analysis one step further. We investigate an environment in which we have both a financial risk process $x(t)$ that can traded and hedged in a market, and also an non-hedgeable insurance risk process $y(t)$. Basically, this section seeks to combine the results from Sections 3 and 4.

5.1 Model Ambiguity and Hedging

The process for the financial risk $x(t)$ is given in equation (3.1) and the process for the non-hedgeable risk $y(t)$ is given in (4.1). Similar to the setup in Section 4, an agent is considered that is uncertain about the true value of the drift parameters m and a of the financial and the insurance processes, respectively. Assume that the agent faces no uncertainty about the diffusion coefficients σ, b and the correlation parameter ρ between the the Brownian Motions W_x and W_y.

To help us describe the uncertainty set, we introduce some further notation. The vector of drift rates μ, and the covariance matrix Σ are defined as follows

$$\mu := \begin{pmatrix} m \\ a \end{pmatrix}, \quad \Sigma := \begin{pmatrix} \sigma^2 & \rho\sigma b \\ \rho\sigma b & b^2 \end{pmatrix}. \tag{5.1}$$

We now make the assumption that the joint uncertainty in the drift rates is described by the following set:

$$\mathcal{K} := \{\mu_0 + \varepsilon \mid \varepsilon' \Sigma^{-1} \varepsilon \leq k^2\}. \tag{5.2}$$

The specification of the uncertainty in this form is motivated (like in Section 4.4) by the fact that the economic agent can use econometric estimation techniques to estimate the drift rates. The estimation leads to the point estimate μ_0. However, there is uncertainty surrounding this estimate. This uncertainty typically is proportional to the covariance matrix Σ. In other words, the agent assumes that the true values of the drift parameters lie somewhere within the confidence interval given by the set $\mathcal{K}$. In the one-dimensional case we considered in Section 4.4, the uncertainty set $\mathcal{K}$ for the drift rate a simplifies to $a \in [a_0 - kb, a_0 + kb]$.

We want to investigate what price the agent will attribute to a derivative that depends both on financial and insurance risk and has payoff $g(t + \Delta t, x, y)$ at time $t + \Delta t$. We furthermore assume (like in Section 3.2) that the agent can hedge the financial risk by investing an amount D_t in the risky asset x at time t, but cannot trade in the insurance asset y. Hence, the robust rational agent solves the following optimisation problem for each time-step $[t, t + \Delta t]$:

$$\begin{aligned} \max_{D_t} \min_{\varepsilon_t} \quad & e^{-r\Delta t}\mathbb{E}_t^{[\varepsilon_t]}\left[g\left(t + \Delta t, x_{t+\Delta t}, y_{t+\Delta t}\right) - D_t x_{t+\Delta t}\right] \\ \text{s.t.} \quad & \varepsilon' \Sigma^{-1} \varepsilon \leq k^2, \end{aligned} \tag{5.3}$$

where $\mathbb{E}^{[\varepsilon_t]}[\,]$ denotes taking the expectation using the drift term $\mu_0 + \varepsilon_t$ for the processes x and y.

This "maxmin" optimisation problem can be interpreted as a two-player game. First, the agent chooses an amount D_t to invest in x, and then Mother Nature chooses the worst possible perturbation ε_t of the drift μ_0. But the agent is aware of the bad intentions of Mother Nature, and therefore chooses the amount D_t that maximises the minimal outcome of Mother Nature.

The optimal solution for D_t is given by

$$D_t^* := -\left(g_x + \rho\frac{b}{\sigma}g_y\right) + \frac{\lambda}{\sqrt{k^2 - \lambda^2}}\frac{b\sqrt{1-\rho^2}}{\sigma}|g_y|, \tag{5.4}$$

where λ is the market price of risk defined in equation (3.7).

Several interesting things about this solution are noteworthy. First, the solution D_t^* is only well-defined for $\lambda^2 < k^2$. This corresponds exactly to the condition encountered in Section 4.3: that the good deal bound κ should be larger than the market price of risk for the financial market. Stated differently, if the agent is confident that even in the worst case a positive excess return can be made by investing in the financial market (i.e. when $\lambda^2 > k^2$), then the agent will try to invest a massive amount in the financial market and has a confident expectation of getting very rich.

The second interesting thing to note is that the optimal hedge position consists of two parts: the hedge portfolio $-(f_x + \rho \frac{b}{\sigma} f_y)$ and a "speculative" portfolio that is determined by the product of the residual non-hedgeable risk $b\sqrt{1-\rho^2}/\sigma\, f_y$ and the "market confidence factor" $\lambda/\sqrt{k^2-\lambda^2}$. The market confidence factor shoots to infinity if λ approaches k. For small values of λ, the market confidence factor is approximately equal to λ/k, and the speculative investment is then approximately proportional to the market price of risk scaled down by a factor of k.

5.2 Agent's Valuation

If we substitute for each time-step $[t, t+\Delta t]$ the optimal solutions for (D_t^*, ε_t^*) into (5.3), and then take the limit for $\Delta t \to 0$ (as we did in Section 3.5), we find a partial differential equation for the price $g(t, x, y)$:

$$g_t + rg_x + a^* g_y + \tfrac{1}{2}\sigma^2 g_{xx} + \rho\sigma b g_{xy} + \tfrac{1}{2} b^2 g_{yy} - rg = 0, \quad (5.5)$$

where the drift term a^* for the insurance process is given by

$$a^* = a_0 - \rho\lambda b \pm b\sqrt{(1-\rho^2)(k^2-\lambda^2)}, \qquad (5.6)$$

where the sign of the last term depends on the sign of g_y. Although the expressions for a^* may look a bit complicated, some very nice interpretations can be given for the expressions.

The first interpretation for a^* is an economic interpretation. Recall that the formula for the optimal hedge given in equation (5.4) consists of two parts: a hedge portfolio and a "speculative" portfolio that depends on λ. Suppose that the agent would only choose the hedge portfolio $D_t := -(f_x + \rho b/\sigma f_y)$. Substitutign D_t into (5.3) yields an expression very similar to (5.5), except that the drift term for the insurance process would be given by $a = a_0 - \rho\lambda b \pm kb\sqrt{1-\rho^2}$. Therefore, by including the speculative portfolio into the optimal hedge, the agent can finance part of uncertainty in a^* by exploiting the expected excess return on equities. This then results in the optimal drift term a^*, where the residual non-hedgeable insurance risk $kb\sqrt{1-\rho^2}$ is shrunk by an additional factor $\sqrt{k^2-\lambda^2}$.

Note also the downward adjustment of the drift of the non-traded asset y by the term $-\rho b\lambda$. This downward adjustment compensates exactly the excess return of y due to the correlation ρ with the traded asset x. This effect makes sense if we think in terms of the Capital Asset Pricing Model (CAPM). In the CAPM, the expected return of any asset is given by the formula $\mathbb{E}[dy(t)] = \big(r + \beta(m-r)\big)dt$, where β is given by the formula $\beta := \rho b\sigma/\sigma^2$. Combining the expression for β with the definition of the market price of risk λ given in (3.7) yields $\mathbb{E}[dy(t)] = \big(r + \rho b\lambda\big)dt$. Hence, the drift term a^* in the agent's valuation adjusts the drift in two steps: the first step $-\rho\lambda b$ corrects the drift of y for the excess return injected by the correlation with the traded asset, and the second term $b\sqrt{(1-\rho^2)(k^2-\lambda^2)}$ adjusts the drift of y for the non-hedgeable risks.

The second interpretation of equation (5.6) is a geometric interpretation. Recall that the uncertainty set $\mathcal{K}$ is an ellipsoid centred around μ_0. Because the financial component x of the risk vector is perfectly replicated, this means that the uncertainty regarding the mean of the financial risk is eliminated, and is replaced by the risk-free return r. The uncertainty for the mean of the insurance process y is now confined to the intersection of uncertainty set $\mathcal{K}$ and the line $m = r$. The intersection of a line and an ellipsoid has two solutions: exactly the two solutions given in equation (5.6).[7]

This geometric interpretation is illustrated in Figure 1, using the following parameters: $m_0 = \{4\%, 7\%, 10\%\}$, $r = 4\%$, $a_0 = 0$, $\sigma = 0.15$, $b = 1$, $\rho = 0.75$ and $k = 2/\sqrt{25} = 0.4$. The three different values of m_0 lead to $\lambda = \{0, 0.2, 0.4\}$, and these three cases are illustrated in the sub-figures 1(a), 1(b) and 1(c). The uncertainty set $\mathcal{K}$ is given by the interior of the ellipse, and the point estimate μ_0 is given by the point in the centre. The line $m = r$ is the vertical dotted line, and the intersection with the ellipse gives the two solutions for a^*. Figure 1(a) illustrates the case $\lambda = 0$. In this case, $a^* = a_0 \pm kb\sqrt{1-\rho^2}$, which is the "naive" confidence interval for a equal to the point estimate a_0 plus/minus k times the non-hedgeable insurance risk $b\sqrt{1-\rho^2}$. The other sub-figures illustrate that for larger values of λ, the ellipse moves to the right. This is a reflection of the fact that higher values of λ correspond to higher point estimates of m_0. When the

[7]Note that this geometric interpretation is equivalent to the result found in Barrieu and El Karoui (2005), where the pricing measure is characterised as the "inf-convolution" of the set of test measures (in our case, the ellipsoid) and the set of martingale measures (in our case, the line $m = r$).

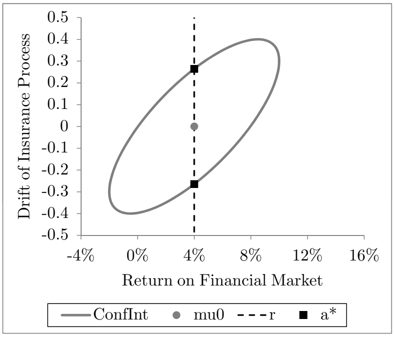

(a) $\lambda = 0$

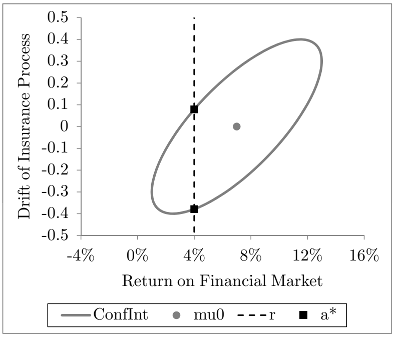

(b) $\lambda = 0.2$

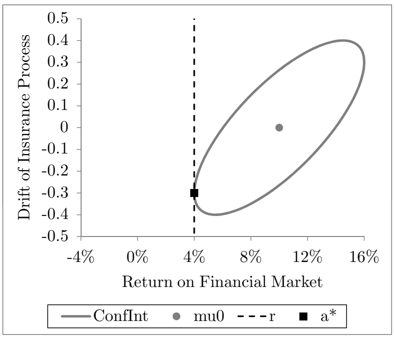

(c) $\lambda = 0.4$

Figure 1: Confidence interval for μ for different values of λ.

ellipse moves to the right, the points a^* move down (due to the correlation term $-\rho\lambda b$), and the points move closer together (due to the factor $\sqrt{k^2 - \lambda^2}$). Figure 1(c) illustrates the largest allowed value for λ. For larger values of λ, the ellipse no longer intersects the line $m = r$, and the optimisation (5.3) no longer has a finite solution.

5.3 On the choice of k

This section concludes by elaborating on the choice of k. In the one-dimensional case treated in Section 4.5 we took $k = 1.96/\sqrt{43} = 0.30$. This choice was motivated by considering the 95% confidence interval for the estimate of the mean of the process $y(t)$ using 43 years of historical data.

In the two-dimensional case we considered in this section, we have to modify this argument slightly. If we used 43 years of data to estimate the vector $\mu = (m, a)'$, then the confidence interval for μ is given by the set $\mathcal{K}$ defined in equation (5.2). The set $\mathcal{K}$ describes a 95% confidence interval if we take k^2 equal to the 95% critical value of a χ^2-distribution with 2 degrees of freedom, divided by the number of observations. This value is given by $k^2 = 5.99/43 = 0.139$, which leads to the value $k = 0.37$.

For the one-dimensional case, we would take the 95% critical value of a χ^2-distribution with 1 degree of freedom, divided by 43, which leads to $k = \sqrt{3.84/43} = 0.30$, which is the same value that was derived in Section 4.5.

6. Applications

This section uses several examples to illustrate the concepts developed in this paper.

6.1 Pricing Long-Dated Cash Flows
As mentioned in the introduction, the pricing of very long-dated cash flows is an important problem for insurance companies and pension funds. This section illustrates the application of the pricing approach outlined in this paper to this problem.

We start by assuming that the interest rates are stochastic, and can be described by a Vasicek (1977) model. In this model we take the instantaneous short rate $r(t)$ and model it as

$$dr = a\big(\theta - r\big)\,dt + \sigma\, dW_r. \tag{6.1}$$

In this equation, θ is the long-term average of the interest rates, and a is the speed of mean reversion.

When pricing cash flows, we have to distinguish two cases. For maturities up to 30 years, there are bonds traded in financial markets, and the market is complete. This means that a portfolio of cash flows is priced at the same price as the value of a *replicating portfolio* of bonds that matches the cash-flow pattern. This price can also be calculated by discounting the cash flows at the zero-curve implied by the market.

For cash flows beyond 30 years, there are no bonds available, which leads us in an incomplete market situation. This calls for application of the methods outlined in Section 4 to determine the value of the cash flows.

For pricing short cash flows (i.e. cash flows promised to policyholders), we have to adjust the interest rates downwards. Applying the methodologies of Section 4 to the interest rate

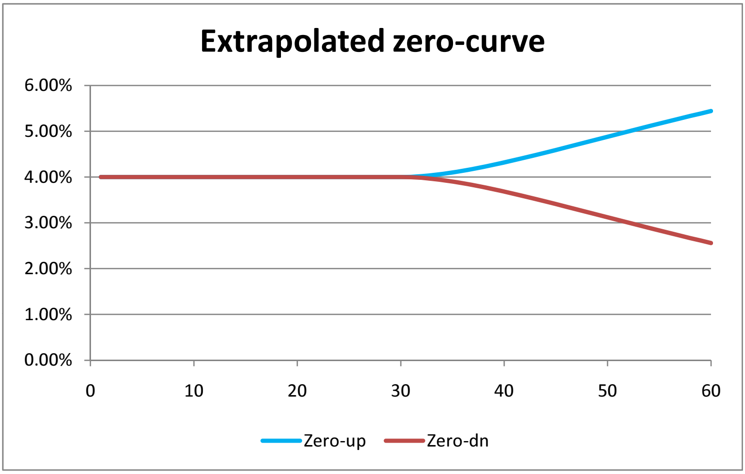

Figure 2: Extrapolated term-structure of interest rates

dynamics (6.1), we get

$$dr = a\left(\theta - k\tfrac{\sigma}{a} - r\right) dt + \sigma\, dW_r, \tag{6.2}$$

where we set $k = 0.30$. This formula implies that a zero rate with maturity $T > 30$ has to be adjusted up or down (for long or short cash flows, respectively) by the formula

$$\frac{k\sigma}{aT}\left((T-30) - \left(\frac{1-e^{-a(T-30)}}{a}\right)\right). \tag{6.3}$$

Assume that the long-term nominal interest rate is 4% (for example, composed of 2% inflation plus 2% real interest rates). If it is also assumed that $a = 0.05$ and $\sigma = 0.01$, then it is possible to explicitly calculate the extension of the zero-curve beyond the 30-year point using formula (6.3). This is illustrated in Figure 2. The left part of the curve between year 0 and 30 shows a market-curve that is flat at 4%. Beyond year 30 we extrapolate the curve using the up and down adjustment to the zero-rates given by equation (6.3). The bottom line shows the extrapolation of the curve for short cash flows (i.e. cash flows promised to policyholders). The upper line shows the extrapolation of the curve for long cash flows (i.e. cash flows to be received from policyholders). The figure reveals, for long maturities the incompleteness of the market becomes more and more pronounced in the extrapolated discount rates.

For further examples, see Cairns (2000), De Jong (2008a) and Iyengar and Ma (2010).

6.2 Pricing Longevity Risk

At the moment there is no well-developed market for longevity risk. Hence, it is reasonable (for the time being) to make the assumption that longevity cannot be hedged, which implies that we have to rely on the methods outlined in Section 4 for pricing longevity (and/or mortality) risk.

To provide a feel for the size of the numbers, this section focuses on one particular summary statistic: the life expectancy at birth of Dutch males. Figure 3 plots the historical development of the life expectancy at birth between 1950 and 2006, based on data downloaded from the Human Mortality Database (see `www.mortality.org`). Two things are immediately obvious from this graph: first, life expectancy has significantly increased in the past 50 years: from 70 years in 1950 to nearly 78 years in 2006. Second, the different trends are clear: between 1950 and 1960, life expectancy was increasing; then from 1960 until 1970, life expectancy was decreasing; since 1970, life expectancy has seen a steady increase. It is exactly these "trend breaks" that make

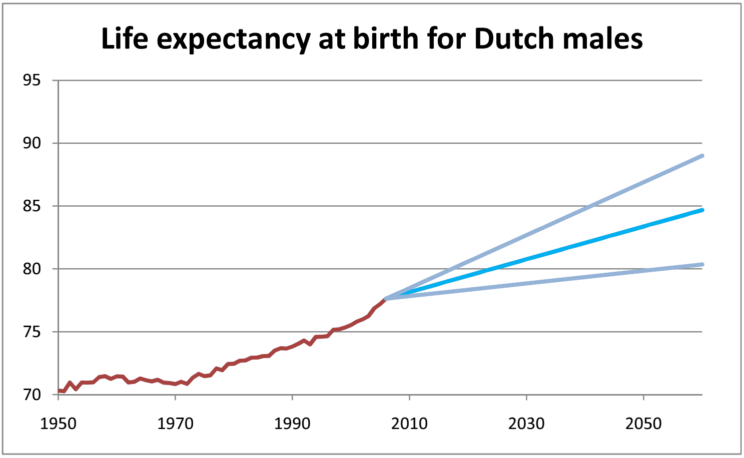

Figure 3: Life Expectancy at Birth for Dutch Males

forecasting of human mortality so difficult.

Based on almost 60 years of data, an average increase in life expectancy of 1.6 months per years is estimated, and a standard deviation of 3.2 months per year. The Dutch Actuarial Association has recently published new life tables, identifying a trend of 1.8 months increase in life expectancy for Dutch males per year, which is very close to the number found here.

However, when we want to price a portfolio of contracts where we worry about people living longer (as is typical for life insurance and pension portfolios). Using the methods outlined in Section 4, we should (for pricing purposes) therefore adjust the trend upwards (in the more conservative direction) by 0.30 standard deviation, leading to a “prudent” trend of 1.6+0.30*3.2 = 2.6 months per year increase in life expectancy. The “prudent” up and down trends are also illustrated in Figure 3 for the period from

2010 until 2060.

For further examples, see Wang et al. (1997), Young and Zariphopoulou (2002), Young (2004b), Milevsky et al. (2006), Bayraktar and Young (2007), Milevsky and Young (2007), Bayraktar and Young (2008), Hári et al. (2008), Ludkovski and Young (2008), Young (2008) and Bayraktar et al. (2009).

6.3 Pricing Non-Hedgeable Equity Risk

The final example examines the pricing of equity risk. A typical market-consistent setting assumes that equities can be freely traded in financial markets, and that equity risk is fully hedgeable. This means that equity risk is priced using the risk-neutral methods outlined in Section 3.

However, in the case of very large pension funds this assumption is questionable. If very large pension funds would buy or sell very large equity positions, they would move the market prices. Hence, large pension funds are not "price takers", but can move the market. In particular, in times of crisis, large pension funds could push the market down even further if they would try to sell equities in response to the drop in prices. Fortunately, this has not happened during the last two crises, thanks to adequate "relaxation" of the underfunding rules by the Dutch Central Bank.

If we take the fact that large pension funds cannot trade without moving the market to the extreme that large pension funds cannot trade at all, then we have again an incomplete market situation, and equity exposures should be priced using the methods of Section 4.

Another example of "equity incompleteness" is the case of insurance companies that give profit-sharing to their policyholders based on the performance of their own investment portfolio. In the traditional approach, such as that of Grosen and Jorgensen

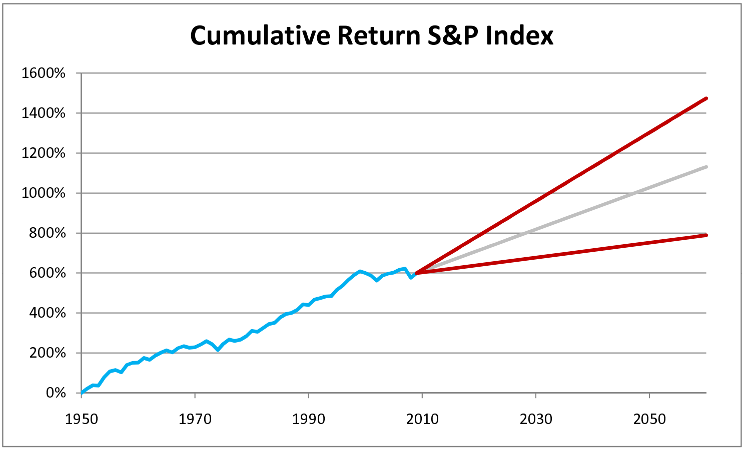

Figure 4: Cumulative Return in Standard and Poor's Stock Index

(2000), such profit-sharing options are priced with the risk-neutral methods of Section 3. However, there is a problem: risk-neutral pricing is based on the cost of the replicating portfolio of the contract. But in the case of profit-sharing on your own portfolio, it is impossible to hedge: at the moment you start buying instruments to hedge your own profit-sharing options, you start changing the composition of your asset portfolio, which then starts changing the nature of the profit-sharing. One approach to solve this problem was suggested by Kleinow (2009), where one tries to find the investment portfolio with profit-sharing that is "self-hedging". The solution to this approach is that the only portfolios that are "self-hedging" are those that have no investment risk (and thus have perfectly predictable profit-sharing). Although mathematically correct, this outcome does not reflect the behaviour of insurance companies in reality.

An alternative approach would be to assume that these profit-sharing options are "non-hedgeable" and should be valued using the methods of Section 4.

What kind of pricing do we get for the "non-hedgeable equity" approach? Figure 4 plots the cumulative return of the S&P equity index between 1950 and 2009. The average return over this period is 10.4% with a standard deviation of 16.9%. When we need to price an non-hedgeable equity position as an asset or as a written put-option, we adjust the return down by 0.30 standard deviation to 5.33%. This higher level of conservatism reflects the additional risk associated with holding an non-hedgeable position.

For further examples, see Davis (1997, 2006) and De Jong (2008b).

References

Artzner, P., Delbaen, F., Eber, J., and Heath, D. (1999). Coherent measures of risk. *Mathematical Finance*, 9(3):203–228.

Artzner, P., Delbaen, F., Eber, J., Heath, D., and Ku, H. (2007). Coherent multiperiod risk adjusted values and bellman's principle. *Annals of Operations Research*, 152(1):5–22.

Barrieu, P. and El Karoui, N. (2005). Inf-convolution of risk measures and optimal risk transfer. *Finance and Stochastics*, 9(2):269–298.

Bayraktar, E., Milevsky, M., David Promislow, S., and Young, V. (2009). Valuation of mortality risk via the instantaneous sharpe ratio: Applications to life annuities. *Journal of Economic Dynamics and Control*, 33(3):676–691.

Bayraktar, E. and Young, V. (2007). Hedging life insurance with pure endowments. *Insurance: Mathematics and Economics*, 40(3):435–444.

Bayraktar, E. and Young, V. (2008). Pricing options in incomplete equity markets via the instantaneous sharpe ratio. *Annals of Finance*, 4(4):399–429.

Becherer, D. (2001). The numeraire portfolio for unbounded semimartingales. *Finance and Stochastics*, 5(3):327–341.

Becherer, D. (2009). From bounds on optimal growth towards a theory of good-deal hedging. In Albrecher, H., Runggaldier, W., and Schachermayer, W., editors, *Advanced Financial Modelling*, Radon Series on Computational and Applied Mathematics (Vol. 8), pages 27–51. De Gruyter.

Björk, T. and Slinko, I. (2006). Towards a general theory of good-deal bounds. *Review of Finance*, 10(2):221–260.

Black, F. and Scholes, M. (1973). The pricing of options and corporate liabilities. *Journal of Political Economy*, 81:637–659.

Cagetti, M., Hansen, L., Sargent, T., and Williams, N. (2002). Robustness and pricing with uncertain growth. *Review of Financial Studies*, 15(2):363–404.

Cairns, A. (2000). Some notes on the dynamics and optimal control of stochastic pension fund models in continuous time. *ASTIN bulletin*, 30(1):19–56.

Carmona, R. (2009). *Indifference pricing: theory and applications*. Princeton University Press.

Carr, P., Geman, H., and Madan, D. (2001). Pricing and hedging in incomplete markets. *Journal of Financial Economics*, 62(1):131–167.

CEIOPS (2008). Qis4 technical specifications. Technical report, Committee for European Insurance and Occupational Pension Supervisors, MARKT/2505/08 [http://www.ceiops.eu].

Černý, A. and Hodges, S. (2002). The theory of good-deal pricing in financial markets. In Geman, H., Madan, D., Pliska, S. R., and Vorst, T., editors, *Mathematical Finance, Bachelier Congress 2000*, pages 172–202. Springer Verlag.

Chen, Z. and Epstein, L. (2002). Ambiguity, risk, and asset returns in continuous time. *Econometrica*, 70(4):1403–1443.

Cheridito, P., Delbaen, F., and Kupper, M. (2005). Coherent and convex monetary risk measures for unbounded cadlag processes. *Finance and Stochastics*, 9(3):369–387.

Cherny, A. (2007). Pricing and hedging European options with discrete-time coherent risk. *Finance and Stochastics*, 11(4):537–569.

Cherny, A. (2009). Capital allocation and risk contribution with discrete-time coherent risk. *Mathematical Finance*, 19(1):13–40.

Cherny, A. (2010). Risk-reward optimization with discrete-time coherent risk. *Mathematical Finance*, 20(4):571–595.

Christensen, M. and Larsen, K. (2007). No arbitrage and the growth optimal portfolio. *Stochastic Analysis and Applications*, 25(1):255–280.

Cochrane, J. H. and Saá-Requejo, J. (2000). Beyond arbitrage: good-deal asset price bounds in incomplete markets. *Journal of Political Economy*, 108(1):79–119.

Coleman, T., Kim, Y., Li, Y., and Patron, M. (2007). Robustly hedging variable annuities with guarantees under jump and volatility risks. *Journal of Risk and Insurance*, 74(2):347–376.

Cont, R. (2006). Model uncertainty and its impact on the pricing of derivative instruments. *Mathematical Finance*, 16(3):519–547.

Cox, J. C., Ross, S. A., and Rubinstein, M. (1979). Option pricing: A simplified approach. *Journal of Financial Economics*, 7:229–263.

CRO-Forum (2006). A market cost of capital approach to market value margins. Technical report, Chief Risk Officer Forum [http://www.croforum.org/publications.ecp].

Cvitanić, J. and Karatzas, I. (1999). On dynamic measures of risk. *Finance and Stochastics*, 3(4):451–482.

Davis, M. (1997). Option pricing in incomplete markets. In Dempster, M. and Pliska, S., editors, *Mathematics of Derivative Securities*, pages 216–226. Cambridge University Press.

Davis, M. (2006). Optimal hedging with basis risk. In Kabanov, Y., Liptser, R., Stoyanov, J., and Davis, M., editors, *From Stochastic Calculus to Mathematical Finance*, pages 169–187. Springer Berlin Heidelberg.

De Jong, F. (2008a). Pension fund investments and the valuation of liabilities under conditional indexation. *Insurance: Mathematics and Economics*, 42(1):1–13.

De Jong, F. (2008b). Valuation of pension liabilities in incomplete markets. *Journal of Pension Economics and Finance*, 7(03):277–294.

Delbaen, F., Peng, S., and Rosazza Gianin, E. (2010). Representation of the penalty term of dynamic concave utilities. *Finance and Stochastics*, 14(3):449–472.

Delbaen, F. and Schachermayer, W. (1994). A general version of the fundamental theorem of asset pricing. *Mathematische Annalen*, 300(3):463–520.

Delong, L. (2011). No-good-deal, local mean-variance and ambiguity risk pricing and hedging for an insurance payment process. Working paper, Warsaw School of Economics.

Denuit, M., Dhaene, J., Goovaerts, M., Kaas, R., and Laeven, R. (2006). Risk measurement with equivalent utility principles. *Statistics & Decisions*, 24(1/2006):1–25.

Deprez, O. and Gerber, H. (1985). On convex principles of premium calculation. *Insurance: Mathematics and Economics*, 4(3):179–189.

Detlefsen, K. and Scandolo, G. (2005). Conditional and dynamic convex risk measures. *Finance and Stochastics*, 9(4):539–561.

Dimson, E., Marsh, P., and Staunton, M. (2002). *Triumph of the Optimists*. Princeton University Press.

Duffie, D. and Epstein, L. (1992a). Asset pricing with stochastic differential utility. *Review of Financial Studies*, 5(3):411.

Duffie, D. and Epstein, L. (1992b). Stochastic differential utility. *Econometrica: Journal of the Econometric Society*, 60(2):353–394.

EIOPA (2010). QIS5 technical specifications. Technical report, European Insurance and Occupational Pensions Authority, [https://eiopa.europa.eu].

Epstein, L. and Schneider, M. (2008). Ambiguity, information quality, and asset pricing. *The Journal of Finance*, 63(1):197–228.

Filipovic, D. and Vogelpoth, N. (2008). A note on the Swiss solvency test risk measure. *Insurance: Mathematics and Economics*, 42(3):897–902.

Föllmer, H. and Schied, A. (2002). Convex measures of risk and trading constraints. *Finance and Stochastics*, 6(4):429–447.

Föllmer, H., Schied, A., and Weber, S. (2009). Robust preferences and robust portfolio choice. In Zhang, B. ., editor, *Handbook of Numerical Analysis, XV, Mathematical Modeling and Numerical Methods in Finance*, pages 29–89. Springer.

Frittelli, M. and Rosazza Gianin, E. (2002). Putting order in risk measures. *Journal of Banking & Finance*, 26(7):1473–1486.

Gerber, H. (1979). *An introduction to mathematical risk theory*. Irwin Homewood.

Goldfarb, D. and Iyengar, G. (2003). Robust portfolio selection problems. *Mathematics of Operations Research*, 28(1):1–38.

Goovaerts, M. and Laeven, R. (2008). Actuarial risk measures for financial derivative pricing. *Insurance: Mathematics and Economics*, 42(2):540–547.

Grosen, A. and Jorgensen, P. L. (2000). Fair valuation of life insurance liabilities: The impact of interest rate guarantees, surrender options, and bonus policies. *Insurance: Mathematics and Economics*, 26:37–57.

Gundel, A. and Weber, S. (2007). Robust utility maximization with limited downside risk in incomplete markets. *Stochastic Processes and Their Applications*, 117(11):1663–1688.

Hansen, L. and Sargent, T. (2001). Robust control and model uncertainty. *American Economic Review*, 91(2):60–66.

Hansen, L. and Sargent, T. (2007). *Robustness*. Princeton University Press.

Hansen, L., Sargent, T., Turmuhambetova, G., and Williams, N. (2006). Robust control and model misspecification. *Journal of Economic Theory*, 128(1):45–90.

Hansen, L. P. and Scheinkman, J. A. (1995). Back to the future: Generating moment implications for continuous-time markov processes. *Econometrica*, 63(4):767–804.

Hardy, M. and Wirch, J. (2005). The iterated cte–a dynamic risk measure. *North American Actuarial Journal*, 8(4):62–75.

Hári, N., De Waegenaere, A., Melenberg, B., and Nijman, T. (2008). Longevity risk in portfolios of pension annuities. *Insurance: Mathematics and Economics*, 42(2):505–519.

Harrison, J. and Kreps, D. (1979). Martingales and arbitrage in multiperiod securities markets. *Journal of Economic Theory*, 20(3):381–408.

Harrison, M. and Pliska, S. (1981). Martingales and stochastic integrals in the theory of continuous trading. *Stochastic Processes and Their Applications*, 11:215–260.

Henderson, V. (2002). Valuation of claims on nontraded assets using utility maximization. *Mathematical Finance*, 12(4):351–373.

Henderson, V. (2005). Explicit solutions to an optimal portfolio choice problem with stochastic income. *Journal of Economic Dynamics and Control*, 29(7):1237–1266.

Henderson, V. and Hobson, D. (2007). Horizon-unbiased utility functions. *Stochastic Processes and their Applications*, 117(11):1621–1641.

Henderson, V. and Hobson, D. (2009). Utility indifference pricing: An overview. In Carmona (2009).

Hobson, D. (2004). Stochastic volatility models, correlation, and the q-optimal measure. *Mathematical Finance*, 14(4):537–556.

Hodges, S. and Neuberger, A. (1989). Optimal replication of contingent claims under transaction costs. *Review of Futures Markets*, 8(2):222–239.

Hu, Y., Imkeller, P., and Müller, M. (2005). Utility maximization in incomplete markets. *Annals of Applied Probability*, 15(3):1691–1712.

Huber, P. (1981). *Robust statistics*. John Wiley & Sons Inc.

Huber, P. and Ronchetti, E. (2009). *Robust statistics (2nd Ed.)*. John Wiley & Sons Inc.

Hugonnier, J., Kramkov, D., and Schachermayer, W. (2005). On utility-based pricing of contingent claims in incomplete markets. *Mathematical Finance*, 15(2):203–212.

Hull, J. C. (2009). *Options, Futures, and Other Derivatives*. Prentice-Hall, Upper Saddle River, N.J., 7th edition.

Iyengar, G. and Ma, A. (2010). A robust optimization approach to pension fund management. *Journal of Asset Management*, 11(2):163–177.

Jaschke, S. and Küchler, U. (2001). Coherent risk measures and good-deal bounds. *Finance and Stochastics*, 5(2):181–200.

Jobert, A. and Rogers, L. (2008). Valuations and dynamic convex risk measures. *Mathematical Finance*, 18(1):1–22.

Kaas, R., Goovaerts, M., Dhaene, J., and Denuit, M. (2008). *Modern Actuarial Risk Theory: Using R*. Springer Verlag.

Karatzas, I. and Kardaras, C. (2007). The numéraire portfolio in semimartingale financial models. *Finance and Stochastics*, 11(4):447–493.

Keller, P. and Luder, T. (2004). White paper of the swiss solvency test. Technical report, Swiss Federal Office of Private Insurance (BPV), Bern,

Switzerland.[http://www.bpv.admin.ch/de/pdf/white_paper_sst.pdf].
Kerkhof, J., Melenberg, B., and Schumacher, H. (2010). Model risk and capital reserves. *Journal of Banking & Finance*, 34(1):267–279.
Kirch, M. (2002). *Efficient hedging in incomplete markets under model uncertainty*. PhD thesis, Humboldt University, Berlin.
Kleinow, T. (2009). Valuation and hedging of participating life-insurance policies under management discretion. *Insurance: Mathematics and Economics*, 44(1):78–87.
Klöppel, S. and Schweizer, M. (2007a). Dynamic indifference valuation via convex risk measures. *Mathematical Finance*, 17(4):599–627.
Klöppel, S. and Schweizer, M. (2007b). Dynamic utility-based good deal bounds. *Statistics & Decisions*, 25(4/2007):285–309.
Ludkovski, M. and Young, V. (2008). Indifference pricing of pure endowments and life annuities under stochastic hazard and interest rates. *Insurance: Mathematics and Economics*, 42(1):14–30.
Lutgens, F. (2004). *Robust Portfolio Optimization*. PhD thesis, Maastricht University.
Lutgens, F. and Schotman, P. (2010). Robust portfolio optimisation with multiple experts. *Review of Finance*, 14(2):343.
Maccheroni, F., Marinacci, M., and Rustichini, A. (2006). Ambiguity aversion, robustness, and the variational representation of preferences. *Econometrica*, 74(6):1447–1498.
Maenhout, P. (2004). Robust portfolio rules and asset pricing. *Review of Financial Studies*, 17(4):951.
Merton, R. (1973). The theory of rational option pricing. *Bell Journal of Economics and Management Science*, 4:141–183.
Milevsky, M., Promislow, S., and Young, V. (2006). Killing the law of large numbers: mortality risk premiums and the sharpe ratio. *Journal of Risk and Insurance*, 73(4):673–686.
Milevsky, M. and Young, V. (2007). Annuitization and asset allocation. *Journal of Economic Dynamics and Control*, 31(9):3138–3177.
Monoyios, M. (2006). Characterisation of optimal dual measures via distortion. *Decisions in Economics and Finance*, 29(2):95–119.
Musiela, M. and Zariphopoulou, T. (2004). An example of indifference prices under exponential preferences. *Finance and Stochastics*, 8(2):229–239.
Musiela, M. and Zariphopoulou, T. (2007). Investment and valuation under

backward and forward dynamic exponential utilities in a stochastic factor model. *Advances in Mathematical Finance*, pages 303–334.

Musiela, M. and Zariphopoulou, T. (2009a). Portfolio choice under dynamic investment performance criteria. *Quantitative Finance*, 9(2):161–170.

Musiela, M. and Zariphopoulou, T. (2009b). The single period binomial model. In Carmona (2009).

Riedel, F. (2009). Optimal stopping with multiple priors. *Econometrica*, 77(3):857–908.

Rogers, L. (2009). Optimal and robust contracts for a risk-constrained principal. *Mathematics and Financial Economics*, 2(3):151–171.

Roorda, B. and Schumacher, J. (2007). Time consistency conditions for acceptability measures, with an application to tail value at risk. *Insurance: Mathematics and Economics*, 40(2):209–230.

Roorda, B., Schumacher, J., and Engwerda, J. (2005). Coherent acceptability measures in multiperiod models. *Mathematical Finance*, 15(4):589–612.

Stadje, M. (2010). Extending dynamic convex risk measures from discrete time to continuous time: A convergence approach. *Insurance: Mathematics and Economics*.

Vasicek, O. (1977). An equilibrium characterization of the term structure. *Journal of Financial Economics*, 5:177–188.

Wang, S., Young, V., and Panjer, H. (1997). Axiomatic characterization of insurance prices. *Insurance: Mathematics and Economics*, 21(2):173–183.

Young, V. (2008). Pricing life insurance under stochastic mortality via the instantaneous sharpe ratio. *Insurance: Mathematics and Economics*, 42(2):691–703.

Young, V. R. (2004a). Premium principles. In Teugels, J. and Sundt, B., editors, *Encyclopedia of Actuarial Science*, volume 3, pages 1322–1331. Wiley.

Young, V. R. (2004b). Pricing in an incomplete market with an affine term structure. *Mathematical Finance*, 14(3):359–381.

Young, V. R. and Zariphopoulou, T. (2002). Pricing dynamic insurance risks using the principle of equivalent utility. *Scandinavian Actuarial Journal*, 2002(4):246–279.

Zariphopoulou, T. (2001). A solution approach to valuation with unhedgeable risks. *Finance and Stochastics*, 5(1):61–82.

SUMMARY OF THE DISCUSSION

By Marc de Graaf

Pricing in Incomplete Markets
By Antoon Pelsser (Maastricht University)
Discussants: Hans Schumacher (1st) and Dirk Broeders (2nd)

Hans Schumacher summarized the contribution by stating that the paper argues that three methods of computing prices in incomplete markets are mathematically equivalent. The liability is modeled in continuous time, and expectation is taken after modifying the drift by κ times the volatility in the adverse direction. For κ, the value 0.5 is suggested. This specific recommendation cuts through the academic discussion and provides a clear starting point.

Schumachers main comments focused on the purpose of pricing, the efficient boundary, the linearity assumption of the price, the 6% rule, the scaling, and the type of analysis. Schumacher mentioned that the purpose of price computations could be price setting, incremental market value computations, or determining the value of liabilities for regulatory purposes. He presumed that Pelssers paper is primarily focused on the last of these. This was confirmed by Pelsser.

The discussant continued by pointing out that regulation is used to prevent accidents from happening, without constraining activity too much. The ‘right’ balance presumably is a political decision, so that the contribution of science is perhaps just the provision of a suitable parameterization of the efficient boundary. From this perspective, the question is not: what is the right value of κ? But rather: do we trace out the efficient boundary by letting κ vary?

The pricing equation in the paper is nonlinear, so that in general the computed value of a portfolio of two liabilities is not equal to the sum of the computed values of the liabilties separately. To get a linear pricing equation, the correction to the drift should take place through a covariance (with the kernel process) rather than through a standard deviation. Such a correction would lead to an extra charge for liabilities that are large during bad economic times (and a reduced charge for liabilities that are small during bad economic times). Pelsser argued that nonlinearity is implicitly included in section 5. Here the phi's can turn negative if losses turn large. Schumacher suggested that this argument takes the covariance into account. Pelsser agreed, but explained that the mathematics in the paper had to be reduced, according to the editorial board.

The 6% rule for the cost of capital has acquired the status of a standard, but the origin is not very clear. When considering the value of 6%, the resulting Sharpe ratio is smaller than the typical Sharpe ratios found in equity markets. Does this mean that the market is more conservative than the regulator? Schumacher found it to be an interesting question, but asserted that it might be important to take into account that an investment decision is not the same as a regulatory decision.

The scaling of VaR by the square root of dt does not reflect a general property of stochastic processes. Presumably the scaling law could be justified at least for small dt's under suitable assumptions (i.e. jumps are not allowed).

Finally, Schumacher noted that the approach followed by Pelsser is closer to the actuarial tradition than to the standard literature on finance. The analysis focuses on a single risk and does not incorporate a pricing kernel.

The second discussant, **Dirk Broeders**, emphasized the relevance of this paper. Pension assets and liabilities stretch far into the future. Broeders found, however, that the paper neglects the derivative markets, which also offer prices beyond 30 years from now. Is the derivative market considered not to be a real market? Broeders also argued that if pension funds are price setters, then they can use the methodology in this paper, but then hedging would no longer be useful. Furthermore, Broeders illustrated that the techniques proposed in this paper have a significant impact. For example applying example 6.1 from the paper to pension funds results in a decrease in the funding ratio of 10 percentage points. He also suggested that Pelsser could discuss the proposals under Solvency 2 (SII) of an ultimate forward rate and an illiquidity premium.

Pelsser argued that in SII a normal distribution is assumed. This assumption has a problematic implication, which is that multiple worst-case scenarios will not happen at the same time. However, it appears that when there is an economic crisis, everything goes bad (and the worst-case scenarios do occur jointly at the same time). This is where the SII falls apart, and gives 40% of diversification benefits away for free.

Broeders argued that, to some extent, model risk is already taken into account in supervision, and that in the future partial internal models are being considered to buffer for specific risk factors.

Broeders wrapped up his comments by summarizing the key point of the paper; namely, that for unhedgeable risks a prudent approach is advocated. Pelsser replied that, to some extent, he is ‘pricing the impossible’. However, as soon as more prices are available, the model adapts to it. Broeders stated that Pelsser kept the market price of risk constant, but that it does react over time. When averaging, one could deal with the cyclical behavior of the market (this was a question from the public). Numerically, one could implement the crisis effect on the willingness to take risk (which is smaller just after a crisis occurred).

PUBLICATIONS IN THE PANEL PAPERS SERIES

1. *Saving and investing over the life cycle and the role of collective pension funds*
   Lans bovenberg , Ralph Koijen, Theo Nijman and Coen Teulings
2. *What does behavioural economics mean for policy? Challenges to savings and health policies in the Netherlands*
   Peter Kooreman and Henriëtte Prast
3. *Housing wealth and household portfolios in an aging society*
   Jan Rouwendal
4. *Birth is the sessenger of death – but policy may help to postpone the bad news*
   Gerard van den Berg and Maarten Lindeboom
5. *Phased and partial retirement: preferences and limitations*
   Arthur van Soest and Tunga Kantarci
6. *Retirement Patterns in Europe and the U.S.* (2008)
   Arie Kapteyn and Tatiana Andreyeva
7. *Compression of morbidity: A promising approach to alleviate the societal consequences of population aging?* (2008)
   Johan Mackenbach, Wilma Nusselder, Suzanne Polinder and Anton Kunst
8. *Strategic asset allocation* (2008)
   Frank de Jong, Peter Schotman and Bas Werker
9. *Pension Systems, Aging and the Stability and Growth Pact* (2008) Revised version
   Roel Beetsma and Heikki Oksanen
10. *Life course changes in income: An exploration of age- and stage effects in a 15-year panel in the Netherlands* (2008)
    Matthijs Kalmijn and Rob Alessie
11. *Market-Consistent Valuation of Pension Liabilities* (2009)
    Antoon Pelsser and Peter Vlaar
12. *Socioeconomic Differences in Health over the Life Cycle: Evidence and Explanations* (2009)
    Eddy van Doorslaer, Hans van Kippersluis, Owen O'Donnell and Tom Van Ourti
13. *Computable Stochastic Equilibrium Models and their Use in Pension- and Ageing Research (2009)*
    Hans Fehr
14. *Longevity risk* (2009)
    Anja De Waegenaere, Bertrand Melenberg and Ralph Stevens

15. *Population ageing and the international capital market* (2009)
Yvonne Adema, Bas van Groezen and Lex Meijdam
16. *Financial Literacy: Evidence and Implications for Consumer Education* (2009)
Annamaria Lusardi and Maarten van Rooij
17. *Health, Disability and Work: Patterns for the Working-age Population* (2009)
Pilar García-Gómez, Hans-Martin von Gaudecker and Maarten Lindeboom
18. *Retirement Expectations, Preferences, and Decisions* (2010)
Luc Bissonnette, Arthur van Soest
19. *Interactive Online Decision Aids for Complex Consumer Decisions: Opportunities and Challenges for Pension Decision Support* (2010)
Benedict Dellaert
20. *Preferences for Redistribution and Pensions. What Can We Learn from Experiments?* (2010)
Jan Potters, Arno Riedl and Franziska Tausch
21. *Risk Factors in Pension Returns* (2010)
Peter Broer, Thijs Knaap and Ed Westerhout
22. *Determinants of Health Care Expenditure in an Aging Society* (2010)
Marc Koopmanschap, Claudine de Meijer, Bram Wouterse and Johan Polder
23. *Illiquidity: implications for investors and pension funds* (2011)
Frank de Jong and Frans de Roon
24. *Annuity Markets: Welfare, Money's Worth and Policy Implications* (2011)
Edmund Cannon, Ian Tonks
25. *Pricing in incomplete markets* (2011)
Antoon Pelsser

Pricing in incomplete markets

This Netspar Panel Paper by Antoon Pelsser (Maastricht University) discusses the pricing of contracts in an incomplete market setting. For life insurance companies and pension funds, it is always the case in practice that not all of the risks in their books can be hedged. Hence, the standard Black-Scholes methodology cannot be applied in this situation. The paper discusses and compares several methods that have been proposed in the literature in recent years: the Cost-of-Capital method (the current industry standard), Good Deal Bound pricing, and pricing under Model Ambiguity.