UNIVERSITY OF SOUTHAMPTON

FACULTY OF SOCIAL SCIENCES
SCHOOL OF ECONOMIC, SOCIAL & POLITICAL SCIENCES
DEPARTMENT OF ECONOMICS

Option Implied Information and Portfolio Allocation

by

Marius Strittmatter

Thesis for the degree of Doctor of Philosophy

June 2019

UNIVERSITY OF SOUTHAMPTON

ABSTRACT

FACULTY OF SOCIAL SCIENCES
SCHOOL OF ECONOMIC, SOCIAL & POLITICAL SCIENCES
DEPARTMENT OF ECONOMICS

Doctor of Philosophy

Option Implied Information and Portfolio Allocation

by Marius Strittmatter

A key concern in the financial literature is if asset prices can be predicted. In the past, studies mainly investigated on the predictability of past asset prices or financial ratios. This thesis takes an alternative route and centres on information derived from options. Utilising options is appealing for two major reasons: First, the structure enables to obtain a forward looking or implied estimate. Second, the cross-section gives detailed insight about the future states of the underlying asset. The starting point of this thesis is to investigate on option implied risk-aversion. A simulation study is developed to analyse risk-aversion estimates from option prices. The simulation suggests that risk-aversion varies drastically over time and the volatility of the underlying has a strong impact. Subsequently, the thesis addresses the use of option implied information within portfolio allocation. The considered investor is infinitely lived with a time varying risk-premium. The dynamics of the risk-premium are derived using various implied predictor variables. They are extracted considering different assumptions about the underlying asset price density. The findings suggest that the option implied risk-premium is the preferable predictor and flexible densities do not benefit notably. Furthermore, performance measures suggest that option implied predictors are superior to historical predictors and similar to financial ratios. The final chapter elaborates further on the previously considered portfolio allocation. It proposes a novel estimation method to obtain consistent estimates of the risk-premium process. The method relies on the observable option implied risk-premium. It avoids the use of a multivariate framework when working with financial ratios. Only with the proposed method, the dynamics of the state variable are economically reasonable. Further, the applied performance measures improve.

Contents

Declaration of Authorship ix

Acknowledgements xi

1 Introduction 1

1.1 Background 2

1.2 Literature Review 4

1.2.1 Option Pricing and Density Forecasting 4

1.2.2 Forward Looking Information and Portfolio Allocation 7

1.3 Contribution 10

1.3.1 Paper 1 11

1.3.2 Paper 2 12

1.3.3 Paper 3 14

2 Uncovering the Distribution of Option Implied Risk Aversion 17

2.1 Introduction 18

2.2 Option implied densities 23
2.2.1 Mixtures of lognormal densities 24
2.2.2 Generalized Beta Densities 25
2.3 Risk Aversion Evaluation 25
2.3.1 Simulation framework 26
2.3.2 Likelihood criteria 27
2.3.3 Market assumptions and estimation 28
2.3.3.1 Geometric Brownian Motion 29
2.3.3.2 Stochastic Volatility 29
2.3.3.3 Expected risk-aversion 30
2.4 Simulations 30
2.4.1 Geometric Brownian Motion 31
2.4.2 Stochastic volatility 34
2.5 Empirical application 40
2.5.1 Data and Estimation 40
2.5.2 Results 42
2.6 Conclusion 48
A Appendix 50

3 Strategic Asset Allocation and Multiperiod Investors using Option-Implied State Variables 51

3.1 Introduction 52
3.2 Methodology 57
3.2.1 Option Implied State Variables 58
3.2.2 Obtaining Option Implied Moments 58
3.2.2.1 Mixture of two Lognormal Densities 59
3.2.2.2 Binomial Tree 62
3.2.2.3 Moment Conversion 63
3.2.3 Portfolio allocation 64
3.3 Evaluation Criteria 67
3.3.1 Data 67

3.3.2 Performance Measures 69
3.3.3 Benchmark Approaches 71
3.4 Estimation 73
3.4.1 State variables 73
3.4.2 State variable Dynamics 75
3.4.3 Portfolio weight characteristics 78
3.5 Empirical findings 81
3.5.1 Performance results 81
3.5.2 Tests 84
3.5.3 Suboptimal portfolio choice results 85
3.5.4 Trade revenue 89
3.6 Conclusion 90
B Appendix 93

4 Consumption and Portfolio Decisions when Expected Returns are Time Varying: An Option Implied Approach 101

4.1 Introduction 102
4.2 Econometric framework 106
4.3 Benchmark approach 108
4.3.1 Classic estimation method 109
4.3.2 Alternative state variables 109
4.4 Application to portfolio allocation 111
4.5 Evaluation criteria 113
4.5.1 Data 113
4.5.2 Suboptimal portfolio choice 116
4.5.3 Out-of-sample performance 116
4.6 Empirical findings 118
4.6.1 Dynamics state variable 119
4.6.2 Estimation portfolio allocation 121
4.6.3 Portfolio weight and consumption 125

4.6.4 Utility and out-of-sample performance 128
4.7 Extension for option implied variance 132
4.8 Conclusion 136
C Appendix 138

5 Conclusion 149
5.1 Summary 150
5.2 Future Work 153

References 156

List of Tables

2.1 Mean Squared Error 32
2.2 Likelihood 33
2.3 Option Statistics 41
2.4 Estimation Results RND 41
2.5 Excess Log Likelihood 43
2.6 AG Test 47

3.1 Option statistics 67
3.2 RND estimation errors 73
3.3 Risk-aversion estimates 74
3.4 Historical approach estimation 75
3.5 Estimation restricted VAR 76
3.6 Mean allocation to risky asset and fraction due to hedging demand 80
3.7 Out-of-sample performance 82
3.8 Statistical tests 86
3.9 Restricted portfolio policy 87

3.10 Trade revenue in % 91

4.1 Option statistics 114
4.2 Options by moneyness 114
4.3 Stochastic process 122
4.4 Parameter estimates portfolio rule 124
4.5 Mean weight and fraction of hedging demand 127
4.6 Average utility and restricted portfolio rules - classic method 130
4.7 Average utility and restricted portfolio rules - direct method 131
4.8 Out-of-sample performance 133
4.9 Restricted portfolio rules 139
4.10 Parameter statistics of the historical approach. 140

List of Figures

2.1 Distribution of Risk-Aversion lognormal density 37
2.2 Distribution of Risk-Aversion mixture-lognormal 38
2.3 Distribution of Risk-Aversion GB2 density 39
2.4 Time Series Risk Aversion assuming Geometric Brownian Motion 44
2.5 Time Series Risk Aversion assuming Stochastic Volatility 45

3.1 Time series myopic and hedging demand - market price of risk 97
3.2 Time series myopic and hedging demand - risk premium 98
3.3 Time series myopic and hedging demand - historical and log dividend price ratio 99

4.1 Time series state variables 120
4.2 Time series portfolio weight direct method 141
4.3 Time series portfolio weight classic method 142
4.4 Time series hedging demand direct method 143
4.5 Time series hedging demand classic method 144
4.6 Time series consumption direct method 145

4.7 Time series consumption classic method 146
4.8 Time series portfolio weight extension 147
4.9 Time series consumption extension 148

Declaration of Authorship

I, Marius Strittmatter, declare that the thesis entitled *Option Implied Information and Portfolio Allocation* and the work presented in the thesis are both my own, and have been generated by me as the result of my own original research.

I confirm that:

1. This work was done wholly or mainly while in candidature for a research degree at this University;
2. Where any part of this thesis has previously been submitted for a degree or any other qualification at this University or any other institution, this has been clearly stated;
3. Where I have consulted the published work of others, this is always clearly attributed;
4. Where I have quoted from the work of others, the source is always given. With the exception of such quotations, this thesis is entirely my own work;

5. I have acknowledged all main sources of help;
6. Where the thesis is based on work done by myself jointly with others, I have made clear exactly what was done by others and what I have contributed myself;
7. Parts of this work have been published as:
   - Kyriacou, M., Olmo, J. and Strittmatter, M. (2019). Uncovering the Distribution of Option Implied Risk Aversion. *Journal of Mathematical Finance*, 9:81-104.

Signed: ....

Date: ....

Acknowledgements

I am especially grateful for the support and suggestions of my supervisors Jose Olmo and Maria Kyriacou who guided me through my time as PhD student. Further, I would like to express my gratitude to Alexandros Kostakis for his discussion at the economics seminar series, and Ana-Maria Fuertes for the comments at the INFINITI Conference 2018. I would like to thank the economics department for the feedback on various occasions, notably Héctor Calvo Pardo, Jean-Yves Pitarakis, Arnau Valladares-Esteban, Michael Hatcher, Thomas Gall and Anastasios Magdalinos. Also thanks to my friends and fellow PhD colleagues including Flavia Bellotto, Abu Siddique, Larissa Marioni, Armine Ghazaryan, Andrea Giovannetti, Chi Wan Cheang and Li Lunzheng who accompanied me through the time at the University of Southampton. Lastly, I want to acknowledge my family and friends for their encouragement and support.

CHAPTER 1

Introduction

The introduction provides an overview of the literature covered within this thesis and outlines the performed research. It is structured into three parts. First, the major milestones of derivative markets are outlined with a focus on options. Further, the key breakthroughs in the academic literature are highlighted. Second, a comprehensive literature review is provided covering the main areas of this thesis: option pricing, density forecasting and portfolio allocation. Third, the methodology and contribution to the literature of each chapter is summarised.

1.1 Background

Over the course of the last 40 years, derivative instruments play an increasingly important role in financial markets. The class of derivatives compromises a range of different products including forwards, swaps, options and many other types. In today's world, they have a broad range of applications including the creation of structured investment solutions with non-linear payoffs, speculation, hedging, but also as tool for executive compensations. Thereby, the value of a derivative instrument roots in the underlying security. In many cases, this is a price of a traded asset. For example, this thesis will be looking into index options, which depend on the price of an underlying equity index.

However, what were the key factors in the last decades that helped derivatives to today's popularity? Probably the starting point of modern derivatives trading was in 1848 with the introduction of the Chicago Board of Trade (CBOT). The initial objective was to provide farmers and merchants with standardised qualities and quantities of grains. However, soon after the successful introduction of the CBOT, the first future contracts were developed. Speculators quickly figured out that trading a future is an appealing alternative to the grain itself. One of the driving success factors was the high standardisation of tradable future contracts and the ability to hedge against the variation in the underlying asset. In contrast to futures, entering option contracts was far more difficult. Interested parties needed to be matched individually. This resulted in a general lack of liquidity and made options an unattractive instrument. Just in 1973, the Chicago Board of Options Exchange (CBOE) was established and the cornerstone for today's option markets was laid. The CBOE was the first exchange that provided a central market place to trade standardised options. This lead to sufficient market participants and liquidity resulting in the success of options until today (Hull, 2012).

Next to the CBOE, another important milestone links to the year 1973. The academic literature marks a fundamental result in option pricing theory. Black and Scholes (1973) and Merton (1973) published their work on the celebrated Black-Scholes-Merton (or Black-Scholes) model, which won the Nobel Prize for economics in 1997. Even though option pricing was subject of academic research before, scientists had difficulties to obtain correct discount rates. Black-Scholes-Merton provide two different approaches to handle this issue. Black and Scholes derive a connection between the required market return on the option and stock using the capital asset pricing model. Merton chooses a different route by constructing a riskless portfolio using the underlying asset and option. The advantage of this procedure is the greater generality since it does not rely on the assumptions of the capital asset pricing model.

Even though the Black-Scholes-Merton model provides a fundamental result in option pricing theory it is not without flaws. The Black Monday in 1987 marks an important event that lead to further developments in option pricing theory. Up to this point there was no evidence that the assumption of constant volatilities through strike prices had been violated (Rubinstein, 1985). After the Black Monday, option prices exhibit the characteristic volatility “smile” or “skew” (Rubinstein, 1994). The Black-Scholes-Merton model cannot account for these characteristics due to restrictive assumptions. However, the volatility priced into the option has strong influence on the option’s price. The higher the volatility the higher the price. Most prominent are out-of-the-money put options in equity markets. They tend to be valued with higher volatilities compared to at-the-money or in-the-money options. An attempt to explain this phenomena argues with the fear of option sellers towards large losses in the underlying asset. Large losses tend to decrease liquidity and increase the likelihood of not being able to hedge the position. Therefore, the option seller seeks for a higher compensation to account for the additional risk. This fundamental change in the structure of volatilities has strong implications on the probability distribution that is used to price the option. Generally, it can be observed that the shape of the distribution has excess kurtosis and negative skew. Therefore, the assumption of a normal distribution is insufficient and attempting to accurately price options using the Black-Scholes-Merton model is not successful. This key issue motivates a large body of literature discussed in the following section.

1.2 Literature Review

In the following, the main body of literature related to this thesis is reviewed. The beginning elaborates on different distributional assumptions of the underlying asset price density to account for the observed non-constant volatilities. Further, the commonly applied transformation methods from risk-neutral to subjective density are discussed. Lastly, the portfolio allocation literature is reviewed. It covers the classic case of the myopic investor and its extension to long-term oriented allocation problems. Moreover, the available literature on option implied information and portfolio allocation is reviewed.

1.2.1 Option Pricing and Density Forecasting

The fundamental changes in the option market due to the events in 1987 motivate a range of different methods in option pricing theory. The main criticism is that a satisfactory underlying asset price distribution needs to be more general than provided by the lognormal family. Particularly, distributional properties should be able to accommodate advanced shapes in regards of skewness and kurtosis. Jackwerth (1999) provides a comprehensive literature review on potential distributional assumptions. The described methods rely on the theoretical result of complete markets by Breeden and Litzenberger (1978). In the following, the most common methods in the literature are reviewed to model risk-neutral distributions. Thereby, the classification relies on Jackwerth (1999):

A popular approach is the application of parametric distribution types. This stream can be categorised into three groups. First, expansion methods. They assume a simple underlying distribution and add correction terms to provide greater flexibility in the style of a Taylor series expansion. Prominent examples for an underlying normal distribution are hypergeometric (or Kummers) function (Rockinger and Abadir, 1997), four parameter Hermite polynomials (Abken et al., 1996a,b), Gram-Charlier expansions (Jondeau and Rockinger, 2001; Corrado and Su, 1997); for an underlying gamma distribution are Laguere polynomials (Brenner and Eom, 1997); for an underlying lognormal, normal and binomial are cumulants (Jarrow and Rudd, 1982; Longstaff, 1995; Corrado and Su, 1996; Rubinstein, 1998).

Second, generalised distribution methods. They provide more flexibility due to additional parameters compared to the (log-) normal distribution. A common feature of this type of distribution is that in the limit it contains commonly known distributions. Examples include generalised beta functions of the second kind (Aparicio and Hodges, 1998) going back to Bookstaber and McDonald (1987), Johnson family of densities (Posner and Milevsky, 1998) going back to Johnson (1949), Burr III distribution (Sherrick et al., 1996a) and Burr XII distribution (Sherrick et al., 1992, 1996b).

Third, mixture methods. This class of density achieves greater flexibility by combining several simple distributions with each other. An issue using this type of specification is that the amount of parameters increases quickly, which might over fit the data. Examples of successful applications are mixtures of two and three lognormal distributions (Ritchey, 1990; Melick and Thomas, 1997; Liu et al., 2007).

A natural alternative to parametric specifications are non-parametric methods. They aim to achieve a higher flexibility by allowing for more general functions. These can be divided into three groups. First, kernel methods. This type of method treats each data point, as for example the implied volatility for a specific strike price, as centre of a region in which the true function could be laying. The region is modelled by a kernel, which is normally a normal distribution. It allows to model the drop in likelihood when moving away from the observed data point. Examples of this method include kernel estimators across five dimensions (stock price, strike price, time to expiration, interest rate, dividend yield) (Aït-Sahalia and Lo, 1998; Pritsker, 1998) and bivariate kernel across time to expiration and moneyness (Rookley, 1997).

Second, maximum entropy methods. This type of method is to some extend similar to a Bayesian framework. They aim to impose a minimum level on unknown or missing information. Therefore, a prior distribution is chosen to find the posterior distribution by maximising the cross-entropy. Examples include using the uniform and lognormal distribution as priors (Buchen and Kelly, 1996) or alternatively the historical distribution as prior (Stutzer, 1996).

Third, non-parametric curve fitting methods. The objective of this method is to find a function that minimises the squared difference to the underlying data while accounting for a reasonable curvature of the function. Examples for this type of method include approaches that fit the volatility smile first and then the risk-neutral distribution (Shimko, 1993; Malz, 1997; Brown and Toft, 1999; Campa et al., 1998; Aparicio and Hodges, 1998; Rosenberg and Engle, 1997; Rosenberg, 1998; Jackwerth, 2000). Alternatively, the risk-neutral density can be approximated directly (Rubinstein, 1994; Jackwerth and Rubinstein, 1996).

Altering the risk-neutral density is a popular and relatively straightforward adjustment to deal with non-constant volatility. Another route to bring more flexibility into the Black-Scholes-Merton framework is to extend the underlying price process. The inital assumption of a Brownian Motion can achieve greater flexibility when adding additional components to it. This can include a jump diffusion (Merton, 1976; Cox and Ross, 1976; Malz, 1996), stochastic volatility (Heston, 1993), or both in combination (Bates, 1991, 1996, 2000). An alternative to stochastic continuous time processes are binomial trees. They are a special case of discretised, one-dimensional diffusion processes. Cox et al. (1979) developed a binomial tree that matches the Black-Scholes-Merton framework. Later this approach is further extended to allow for greater flexibility (Rubinstein, 1994; Jackwerth and Rubinstein, 1996; Jackwerth, 1997), forward solved binomial trees (Barle and Cakici, 1998; Derman and Kani, 1994), or an extension to trinomial trees (Dupire, 1994).

Until this point the presented literature focused on accurate distributional assumptions within the option pricing formula to reflect observed option prices. To gain insight on the underlying asset the obtained distribution needs to be transformed. This transformation is necessary since options are priced risk-neutral. They do not reflect a compensation for the risk in the underlying asset as it is commonly observed in financial markets. The transformed density has multiple names in the literature including physical, objective, real-world or subjective density. Here the density will be referred to as subjective density. The validity of this transformation is underlined in forecasting exercises (Liu et al., 2007;

Shackleton et al., 2010). The predictability of the subjective density is superior to the risk-neutral density. Moreover, it beats alternative benchmarks that rely on a history of asset prices.

In finance, a commonly used approach to transform the density is to use asset pricing theory. Aït-Sahalia and Lo (2000) and Rosenberg and Engle (2002) can show that the option pricing formula can be connected to the stochastic discount factor or pricing kernel. With this result, the stochastic discount factor can be linked to the utility function of the representative agent. Making assumptions about the utility allows to transform the risk-neutral density into the subjective counterpart. Empirical work relies on a range of different utility functions to perform the transformation. This includes a power utility (Bakshi et al., 2003; Bliss and Panigirtzoglou, 2004; Liu et al., 2007) and exponential utility (Bliss and Panigirtzoglou, 2004). An alternative to this procedure is the parametric recalibration method by Fackler and King (1990). It does not rely on any assumptions regarding the utility function. This procedure has been successfully implemented by Liu et al. (2007); Wang (2009) and Shackleton et al. (2010). Further, Shackleton et al. (2010) extends this approach for a non-parametric recalibration method. They also propose a risk-premia transformation within the Heston (1993) and Bates (1996) model.

The calibration of the subjective parameter values relies on two different approaches. First, the maximum likelihood criteria. Based on the risk-neutral density the subjective parameters are obtained by maximising the out-of-sample log-likelihood of realised asset price returns (Bao et al., 2007; Shackleton et al., 2010). Alternatively, the p-value of the LR_3 statistic of Berkowitz (2001) is maximised (Bliss and Panigirtzoglou, 2004). Liu et al. (2007) compares both approaches and concludes that none is superior to the other.

1.2.2 Forward Looking Information and Portfolio Allocation

Most likely the starting point of modern portfolio theory is made by Markowitz (1952). In his analysis he showed that all investors independent of risk-preferences invest only in one optimal portfolio containing multiple assets. Thereby, the risk-preferences affect how strongly investors are exposed to the optimal portfolio. More aggressive investors will allocate a higher proportion or even use leverage, whereas defensive investors will allocate a smaller fraction. The finding that the composition of the portfolio weight is not adjusted is also known as mutual fund theorem by Tobin (1958).

In the real world investors deviate from the optimal portfolio suggested by Markowitz. There are multiple reasons why this is the case. Investors might alter their individual portfolio compositions due to individual preferences. A major impact are risk-preferences. For example, more aggressive individuals tend to allocate a higher fraction into stocks than bonds. Further, different objectives to invest money could influence the composition. For example, saving for retirement might result in other investment decisions than maximising profits.

The literature deals with these problems by introducing the concept of a long-term investor. Key conclusion is that long-term investors do not necessary allocate their wealth in the same way as short-term investors. Long-term investors might be more interested in how wealth affects their well-being over time rather than simply maximising wealth itself. With these thoughts in mind different asset classes that are not attractive in the short run might become attractive in the long run. For example, cash in the short run is attractive, whereas in the long-run it is risky due to changing interest rates and associated reinvestment risk. Stocks on the other side have a high variation in the short run but there is support that in the long-run income streams are more stable. Therefore, a long run investor might be more willing to hold a larger fraction in stocks than in cash. Further, the concept of a long-term investor allows to give an explanation why more risk-averse investors allocate a higher fraction into bonds than equities since bonds provide a more certain payoff profile.

The first developments in this field study long-term portfolio choice under restrictive conditions. Results show that in these special cases, the long-term allocation narrows down to the same optimal weights as short-term investors (Samuelson, 1963, 1969; Mossin, 1968; Merton, 1969; Fama, 1970). Based on these fundamental findings research has been performed for specific asset classes. Modigliani and Sutch (1966) argue that long-term bonds are safe assets when having a long-term horizon. Stiglitz (1970) and Rubinstein (1976a,b) build a theoretical framework to highlight this argument. The influence of human capital is studied by Mayers (1976) and Fama and Schwert (1977b). Merton (1971, 1973) provides the basis of working with time varying investment opportunities. Rubinstein (1976b) and Breeden (1979) study this aspect in further detail. They illustrate the results in terms of consumption risk. This way of interpreting portfolio choice influenced important work as for example Lucas (1978), Grossman and Shiller (1981), Shiller (1982), Hansen and Singleton (1983), Mehra and Prescott (1985) and many others.

Even though the idea of a long-term investor can be well tracked in theory its implementation was confronted with issues for a long time. A major problem is the development of closed form solutions. Many results need to be approximated and required advanced numerical methods. With the development of these methods the assumption of a long-term investor can be well tracked. Under the assumption that asset price returns are to some extend predictable (Campbell, 1987; Campbell and Shiller, 1988a,b; Fama and French, 1988, 1989; Hodrick, 1992) the first models were developed. Kim and Omberg (1996) mark the starting point of allowing for time varying investment opportunities. They provide a closed form solution of an investor that chooses optimal portfolio policies with no transaction costs and HARA utility. Brennan et al. (1997) further extend the model to allow for frequent rebalancing. Balduzzi and Lynch (1999) and Balduzzi and Lynch (2000) consider transaction costs, which reduces the optimal amount of rebalancing. Barberis (2000) provides a long-term investor within a Bayesian framework to account for parameter uncertainty and shows that ignoring estimation risk may overallocate to the risky asset. Liu (2007) further extends the assumptions about underlying market dynamics by providing a closed form solution allowing for quadratic term structure models for bond returns and stochastic volatility for stock returns.

Of central interest in the subsequent sections is the study of Campbell and Viceira (1999). As one of the first, they solve analytically the portfolio allocation of a long-term investor that determines his optimal portfolio and consumption policies. The investor allocates his wealth by maximising Epstein and Zin (1989) and Weil (1989) preferences, which are derived from consumption rather than wealth. Further, the equity premium is time varying. Later, this approach was extended by allowing for changing interest rates and inflation (Campbell et al., 2003) and the risk of changing volatility (Chacko and Viceira, 2005).

So why are option prices interesting in the context of portfolio allocation? The previous section discussed that option prices are implied estimates and reflect current market expectations. These information are beneficial in forecasting exercises. Therefore, it is a natural step to use option implied information to solve a portfolio allocation. If option prices really contain valuable information about the future they should improve the performance of the portfolio allocation compared to alternative methods. This area of research is relatively new. The starting point is the work of Kostakis et al. (2011) who successfully applies option implied information assuming a myopic investor. They analyse the effect when introducing higher moments into the portfolio allocation. Furthermore, they employ a long-term investor. They conclude that option implied information are beneficial when trying to time the market and that a long-term investor is beneficial compared to a myopic one. Another aspect is the question on finding an optimal mean variance portfolio similar to the approach by Markowitz. DeMiguel et al. (2013) and Kempf et al. (2015) construct portfolios using option implied expected returns and volatility. A key contribution of their work is to derive option implied covariance estimates between their considered assets. They support the use of option implied information and highlight the effectiveness of option implied covariance.

1.3 Contribution

This section outlines the methodology and contribution of this three paper thesis. Each subsection summarises paper 1 to 3, which form the core sections 2 to 4 of the thesis.

1.3.1 Paper 1

The first paper of this thesis contributes to the field of risk-aversion estimation from option prices. Estimation of risk-aversion within a representative agent model is known to be a difficult problem as highlighted by Mehra and Prescott (1985). In contrast, the density forecasting literature could show that with the use of option prices risk-aversion estimates are within more reasonable bounds. Nevertheless, a major issue in the density forecasting literature are the limited sample sizes, which heavily impact the risk-aversion estimate. The obtained estimates tend to vary depending on the considered sample period and length. Therefore, the first paper investigates the variation of risk-aversion over time. It aims to obtain a "current" estimate of risk-aversion using a simulation rather than an average over a past estimation period.

Risk-aversion estimates using option prices are obtained in a two step procedure. First, the risk-neutral dynamics are obtained from observed option prices using a more flexible assumption of the risk-neutral density within the option pricing formula. The underlying asset price density is assumed to be a lognormal, mixture-lognormal and generalized beta distribution of the second kind to analyse the impact of different degrees of flexibility. The density parameters are obtained using a cross-section of option prices with 4-weeks until expiry. The estimation is performed by minimising the squared error between observed and model option prices. Second, the obtained risk-neutral parameters need to be transformed in order to reflect risk-premia in the market. This transformation relies on an assumption about the preferences of the representative investor. In this case, the power utility is chosen. Based on the previously obtained risk-neutral parameter estimates the additional parameter from the power utility is derived using final asset price outcomes. The estimation of the relative risk-aversion parameter is conducted by maximising the log-likelihood of final asset price outcomes. The maximisation is performed over multiple past periods using a rolling window. Therefore, the obtained risk-aversion estimate is an average of the considered transformation horizon. It can be observed that risk-aversion estimates vary quite drastically over time when moving the rolling window forward.

The proposed simulation aims to retrieve a "current" estimate of risk-aversion and provides insight into the variation of risk-aversion. For the simulations only the risk-neutral density at observation date t is relevant to retrieve the risk-aversion estimate. Making assumptions about the underlying asset price process allows to obtain a range of asset price outcomes. The considered asset price process relies on a Geometric Brownian motion and extension for stochastic volatility. It enables to simulate multiple asset price outcomes for a single risk-neutral density. Therefore, the estimation of the risk-aversion parameter can be performed multiple times. The advantage of this procedure is that only the most recent risk-neutral density is used in combination with the dynamics of the underlying asset. Consequently, the risk-aversion estimate is referred to as “current” since it does not rely on a past transformation horizon. Further, it allows to obtain a distribution of risk-aversion estimates because the simulation can be repeated arbitrarily.

The empirical findings have three major implications. First, the assumption of the underlying risk-neutral density has strong impact on the variation of risk-aversion. More flexible density types reduce the variation in risk-aversion. Second, overall the findings suggest a high variation in risk-aversion. Stochastic volatility in the underlying asset price process has a strong impact on the uncertainty of risk-aversion. This is in contrast to other studies as for example Bliss and Panigirtzoglou (2004) who only use realised asset price outcomes and report lower variations of risk-aversion. Third, applying the log-likelihood of asset price outcomes as evaluation criterion shows that the out-of-sample forecasting abilities using the proposed simulation method improves density forecasts compared to a rolling window methodology.

1.3.2 Paper 2

The second paper studies the impact of option implied information within portfolio allocation. Options are an interesting source of information due to their forward looking characteristics and the detailed insight into the future state of the underlying asset. The main contribution in this chapter is to study the impact of option implied state variables on an infinitely lived investor. Traditionally, portfolio allocation problems rely on state variables derived from financial ratios. These ratios are replaced by two option implied alternatives: 1. The risk-premium, 2. The market price of risk. Furthermore, different distributional assumptions of the underlying asset price are applied to study the impact of different degrees of flexibility.

The considered portfolio allocation is the infinitely lived investor by Campbell and Viceira (1999). The investor allocates his funds between a risky and risk-free asset by maximising his utility of future consumption. His preferences are derived from an Epstein and Zin (1989) and Weil (1989) utility. Furthermore, he considers a single state variable, which is the risk-premium. In their set-up it follows an AR(1) process. This state variable is responsible for the myopic and hedging demand of the risky asset. In their initial paper they use the log dividend price ratio to derive the dynamics of the risk-premium process. In this chapter the log dividend price ratio is replaced by the option implied risk-premium and market price of risk. The risk-premium and market price of risk are defined as the excess return towards the risk-free asset, and the ratio of the risk premium and volatility, respectively.

The option implied moments, risk-premium and volatility, are derived from a range of different methods. Thereby, the previously described procedure to obtain the option implied density is followed. The considered density types include lognormal, mixture-lognormal and binomial tree. The transformation from risk-neutral to subjective density is performed for 36, 48, 60, 72, 84, 96, 108 and 120 months. Subsequently, the moments can be derived and transformed so they relate to the return distribution. The log dividend price ratio and an alternative historical approach is used as benchmark approach for the option implied state variables. The historical approach derives the state variable from a past series of asset price returns. It relies on assumptions about the return and volatility process to obtain the risk-premium and volatility.

The empirical findings underline the timing abilities of option implied information compared to the historical approach. Only the use of the log dividend price ratio makes the option implied approaches less outstanding. Comparing the different methods to derive the option-implied state variables suggests that there are no major differences between different distributional assumptions. This might be a consequence of only considering the first two moments to calculate the state variables. Analysing the state variables suggests that within the chosen set-up the option implied risk-premium is superior towards the alternative market price of risk. All findings are measured out-of-sample taking into account trading costs.

1.3.3 Paper 3

The second paper suggests that the risk-premium is the favourable predictor variable. The third paper picks up this finding and continues to work directly with the risk-premium implied by option prices. The analysis is again undertaken within the framework of Campbell and Viceira (1999). Thereby, the risk-premium process in their portfolio allocation is of central interest. It follows an autoregressive process with one lag, which also acts as state variable. Campbell and Viceira (1999) propose to derive the dynamics of the risk-premium process from a VAR(1) model. It contains the risk-premium and predictor variable. Typical predictors are financial ratios, which is in their case the log dividend price ratio. Using a VAR(1) model is necessary since financial ratios do not give direct insight on the risk-premium.

Central advantage of option prices is that the risk-premium can be observed directly. Therefore, this chapter proposes an estimation framework that directly and consistently obtains the dynamics of the risk-premium process from option prices. This approach makes the estimation method described by Campbell and Viceira (1999) redundant since a direct proxy of the risk-premium is available. This proxy is retrieved from the model-free approach by Martin (2017). This method is chosen since it diminishes the need to transform risk-neutral to subjective dynamics. However, it requires to approximate two integrals that rely on the cross-section of option prices. These integrals need to be approximated since option prices cannot be observed with continuous strikes. Therefore, the obtained option implied risk-premium is a proxy of the objective risk-premium. This paper develops a framework that uses this proxy of the risk-premium to obtain consistent estimates of the objective risk-premium process.

The proposed direct estimation method comes with multiple advantages. Most striking is that the option implied risk-premium is not conflicted with non-stationarity issues as it is commonly detected for financial ratios. Furthermore, when working with the option implied risk-premium only the direct estimation method obtains economically reasonable parameter values compared to the approach by Campbell and Viceira (1999). The main impact is on the hedging demand for the risky asset. Just in case of the direct estimation method, the average hedging demand is positive. It implies that the infinitely lived investor demands on average a higher fraction of the risky asset. Only in this case the investor maintains a positive share in the risky asset when expected returns are zero. These findings are underpinned by the applied evaluation criteria. The suggested method provides strong improvements in the value function and out-of-sample performance compared to the VAR(1) methodology. Furthermore, the option implied approach is compared against a historical benchmark, simple investment into the index and the use of the log dividend price ratio as state variable. As in paper 2, the option implied approach is clearly in excess to the historical benchmark and index investment. Furthermore, the results are of equal magnitude compared to the log dividend price ratio.

CHAPTER 2

Uncovering the Distribution of Option Implied Risk Aversion

This chapter explores the dynamics of risk aversion of a representative agent with an iso-elastic utility function. In contrast to most of the existing literature, this chapter estimates the coefficient of relative risk aversion using option prices. To do this, the risk-neutral density function obtained from a cross-section of option prices is transformed to an objective distribution function that reflects individuals' risk aversion through a CRRA utility function. The dynamics of the relative risk-aversion coefficient are obtained by repeating the same estimation procedure over rolling windows. This procedure uncovers strong variation in risk aversion over time. Alternatively, a simulation procedure is proposed to construct confidence intervals for the risk-aversion coefficient in each period. The robustness of these confidence intervals is assessed under different assumptions on the data generating process of stock prices. The results imply a strong influence of volatility on the variation of risk aversion. In an empirical application, the forecasting performance of the simulation approach using the retrieved risk-aversion estimates is compared against the rolling window methodology. Overall, the simulation based approach obtains better forecasting results.

2.1 Introduction

Stock options are priced using risk-neutral expectations of the payoff of the underlying asset. These risk-neutral expectations are usually incorporated in the option pricing formula within a flexible specification of the density function to describe the price of the underlying asset. In contrast, stock prices are priced by discounting the expected value of future cash flows of the asset under an objective distribution function. Both approaches incorporate the individuals' attitude towards risk. The risk-neutral approach embeds individuals' risk aversion in the risk-neutral probability distribution of the risky payoffs, that are discounted using the risk-free rate. Opposing to this procedure, the objective approach uses objective probabilities, inferred from historical prices or alternative parametric methods, and discounts future payoffs using a stochastic discount factor. This discount factor incorporates a risk premium over the risk-free rate.

There are different approaches to connect the risk-neutral valuation with the objective valuation. A convenient formulation that relies on an utilitarian approach assumes a parametric relationship between the distribution functions by a utility function. The analytical tractability of power utility functions offers a convenient modelling device for describing individuals' preferences. In this paper the parametric relationship between the risk-neutral and objective distribution functions, and the power utility function is exploited to extract consistent estimates of the coefficient of relative risk aversion of a representative agent. There are two reasons to use option prices for extracting risk aversion from risk-neutral valuations. First, risk-aversion estimates are economically important and provide relevant information for understanding the performance and dynamics of financial markets. Second, the empirical finance literature confirms the significant improvement in the ability of parametric models for forecasting stock prices that incorporate reliable estimates of risk aversion.

The link between expected returns and investors' risk-preferences in financial markets is the foundation of research concerning the risk-return relationship. Several methods were developed over time in order to determine the risk-premia underlying asset markets. Popular approaches are, for example, the standard and consumption-based capital asset pricing model. These pioneering asset pricing methods are usually estimated based on asset price returns, treasury yields and consumption data. Although these models are still the building blocks of current asset pricing models and the foundation of today's financial economics, their success in estimating the underlying risk premia in financial markets have been limited. A prominent example is the seminal article by Mehra and Prescott (1985) on the equity premium puzzle. These authors retrieve extremely high values of the risk aversion coefficient in order for the asset pricing model to be compatible with observed asset prices. These values are not economically justified. Other studies have tried to correct these estimates by imposing different assumptions on individuals' preferences. For example, Arrow (1971) makes strong assumptions on the flexibility of the underlying utility function. This author finds estimates of the risk aversion coefficient equal to one and consistent with the presence of a log-utility function for modelling the preferences of a representative agent. Friend and Blume (1975) find values of the risk aversion coefficient in the range of two. Hansen and Singleton (1982) in their seminal study propose a novel GMM econometric framework for estimating the model parameters. However, their parameter estimates are consistent with previous estimates and oscillate between 0 and 1. Mehra and Prescott (1985), discussed above, find values of the risk aversion coefficient around 55 that are difficult to reconcile with practitioners' and financial economists' expectations on the underlying risk aversion in financial markets. Epstein and Zin (1991) propose a recursive utility function for modelling individuals' preferences and find under this novel approach a risk aversion coefficient that oscillates between 0.4 and 1.4. It is important to remark that the Epstein-Zin utility function is able to disentangle individuals' risk aversion from the rate of intertemporal substitution. Ferson and Constantinides (1991) propose habit formation models in which individuals' preferences depend on a reference point under which perceived utility of the asset is zero. This approach provides more variability in the estimates of the risk aversion coefficient. Thus, these authors find a risk aversion coefficient between 0 and 12. Cochrane and Hansen (1992) in an asset pricing context find values in the range of $40 - 50$ consistent with those found in Mehra and Prescott (1985). Jorion and Giovannini (1993) and Normandin and St-Amour (1998) find values between 5.4 and 11.9 and smaller than 3, respectively. More recently, Aït-Sahalia and Lo (2000) derives a coefficient of risk aversion of 12.7 and Guo and Whitelaw (2001) using option prices, asset price returns and consumption data find a value of the relative risk aversion coefficient equal to 3.52. These authors exploit option prices to extract implied volatility measures for estimating the risk aversion coefficient.

The conclusion of this literature on risk aversion is that although it well acknowledges the existence of a premium in financial markets to compensate individuals for the presence of risk in asset prices under individuals' risk aversion, there is much controversy on how to measure it and what values are consistent with the realizations of stock prices. There are three major approaches for retrieving risk-aversion coefficients based on option prices. They all try to combine risk-neutral dynamics from option prices with the realized dynamics of asset prices. The first approach estimates the risk-neutral density based on option prices separately from the density of asset prices obtained from historical asset price data. Afterwards, both densities are compared against each other to retrieve the risk preferences of market participants (Aït-Sahalia and Lo, 2000; Jackwerth, 2000; Aït-Sahalia et al., 2001; Weinberg, 2001; Pérignon and Villa, 2002; Rosenberg and Engle, 2002; Kliger and Levy, 2002; Anagnou-Basioudis et al., 2005; Ziegler, 2007). The second approach jointly estimates the continuous-time dynamics of the risk-neutral and the subjective process. Option and asset prices are used to jointly estimate the pricing kernel that incorporates the premia for price, volatility and jump risks. (Bates, 2000; Chernov and Ghysels, 2000; Pan, 2002; Shackleton et al., 2010). The third approach assumes a parametric form of the risk-aversion function, which is combined with the risk-neutral density in order to derive the subjective density. (Fackler and King, 1990; Bliss and Panigirtzoglou, 2004; Anagnou-Basioudis et al., 2005; Kang and Kim, 2006; Liu et al., 2007; Shackleton et al., 2010).

The aim is to follow a combination of the first and third approach. The focus of the work is on the transformation proposed by Bliss and Panigirtzoglou (2004). One of the objectives is to gain insight on the variation of risk-aversion estimates over time and, hence, through the business cycle. Bliss and Panigirtzoglou (2004) provide quantitative methods to evaluate risk-aversion estimates. These authors introduce two different approaches. First, a Monte-Carlo simulation method to retrieve the standard error of risk-aversion estimates. Second, a Bootstrapping method to capture the influence of actual data and potential model misspecification on the reliability of risk-aversion estimates. Despite the appealing of both approaches, these methods face challenges due to the limited availability of asset prices outcomes and option prices. To overcome this issue, different parametric data generating processes are considered to construct the risk-neutral distribution function for describing the dynamics of option and asset prices. The proposed method distinguishes from other studies by retrieving the risk aversion estimate taking into account the current risk-neutral density in combination with a data generating process to model the subjective dynamics. Thereby, the process is applied to simulate final asset price outcomes, which are then used to transform the risk-neutral density. Thereby, only the currently observed density is of relevance within the simulation set-up. This enables to only consider current market expectations. By choosing this approach the work in this section distinguishes from studies as for example Aït-Sahalia and Lo (2000); Audrino and Meier (2012); Rosenberg and Engle (2002); Christoffersen et al. (2013); Bakshi et al. (2010) aiming to gain more current insights of current risk-aversion. Central difference is that they consider the entirety of observations through time of risk-neutral densities in combination with physical densities derived from realised asset price returns. Linn et al. (2017) criticises the approach chosen by the previously mentioned authors in regards of their estimated pricing kernels and propose an alternative estimation method which is more robust. This chapter cannot avoid this criticism. However, it needs to be highlighted that the aim of the proposed method is in obtaining a current estimate of risk-aversion that is not an average over a past time horizon.

Variation in the estimates of risk aversion is obtained by constructing rolling windows of 120 monthly observations and obtaining estimates of risk aversion for each window. This number is sufficiently large to avoid too much variation in risk-aversion estimates and at the same time is small enough to capture changes in market conditions that trigger shifts in risk aversion. Confidence intervals of the risk-aversion estimates can be constructed by simulating the prices from a parametric data generating process. The simulated distribution of risk-aversion estimates provides valuable insights on the variation of the risk-aversion coefficient. Furthermore, in the proposed set-up the underlying asset price process allows for different characteristics starting with the base case of a Geometric Brownian Motion and its extension to accommodate stochastic volatility processes. The simulation results are compared with the classic approach followed by Bliss and Panigirtzoglou (2004) and Liu et al. (2007). These authors derive their risk-aversion estimates based on past series of risk-neutral densities and realised asset prices.

The extensive simulation exercise provides the following insights. First, the results suggest that there is more variation in the risk-aversion coefficient than implied by Bliss and Panigirtzoglou (2004). Not surprisingly, the volatility component has a major impact on the coefficient of relative risk aversion. The choice of parametric model for describing asset prices also has a major effect on the estimates of risk aversion for a representative agent. Thus, it needs to be noted that imposing a lognormal density on asset prices produces greater variation in the coefficient of risk aversion compared to the mixture-lognormal and generalised beta distribution of the second kind. Lastly, a forecasting exercise is performed to assess the relative out-of-sample ability of the different data generating processes for estimating the risk-aversion coefficient to forecast stock prices. The forecasting performance is evaluated using the subjective density derived from the average simulated risk-aversion estimate. The findings are compared against the classic approach obtained from Bliss and Panigirtzoglou (2004) and Liu et al. (2007). Overall, the simulation based approach obtains better forecasting results than using a rolling window methodology. This fact is particularly apparent in the period after the sub-prime crisis.

The paper is structured in six sections. Section 2.2 introduces the considered density types and transformation from risk-neutral to subjective density. Section 2.3 describes the proposed simulation method. Section 2.4 elaborates on the simulation results. Section 2.5 provides the summary statistics of the observed S&P 500 option and index prices. Further, estimation errors and risk-neutral characteristics are reported, and empirical results are discussed. Lastly, section 2.6 concludes the main findings and gives suggestions for further research.

2.2 Option implied densities

The risk-aversion estimates obtained in this section are retrieved from option prices. This approach goes back to the density forecasting literature for option pricing as, for example, Liu et al. (2007). The foundations to this literature on option pricing is due to Breeden and Litzenberger (1978). These authors show that when there are no-arbitrage opportunities and options are available for all strikes K, a call-option c can be evaluated by:

$$c(K) = e^{-rT} \int_K^{\infty} (S - K) f_Q(S) dS, \tag{2.2.1}$$

where S is the stock price, r the risk-free rate, T the time to maturity and f_Q a unique risk-neutral density. All parameters can be observed in the market except for the risk-neutral density. Assuming a parametric form of f_Q, its parameters θ can be obtained by minimising the squared error between observed c_o and model-based c_m call prices:

$$\arg\min_{\theta} G(\theta) = \sum_{i=1}^{N} \left(c_o(K_i) - c_m(K_i \mid \theta)\right)^2, \tag{2.2.2}$$

with N the number of options in the cross-section. The risk-neutral valuation involving the density function f_Q is restricted by the risk-neutrality constraint $F = E^Q[S_T]$, stating that the expected value of f_Q equals the forward price F. To reflect the dynamics of the underlying asset, the risk-neutral density needs to be transformed into a subjective counterpart density function that incorporates individuals' risk-aversion profile. Here the transformation in Bliss and Panigirtzoglou (2004) is followed, who relate the risk-neutral f_Q to a subjective density function f_P by making an assumption about the preferences of a representative investor. This parametric relationship goes back to Aït-Sahalia and Lo (2000). In this study, a power utility function u is assumed with $u' = x^{-\gamma}$:

$$f_P(S_T) = \frac{\frac{f_Q(S_T)}{u'(S_T)}}{\int_0^{\infty} \frac{f_Q(y)}{u'(y)} dy} = \frac{S_T^{-\gamma} f_Q(S_T)}{\int_0^{\infty} y^{-\gamma} f_Q(y) dy}. \tag{2.2.3}$$

The constant relative risk-aversion parameter γ, inherited from the power utility and responsible for the transformation into the subjective density, is the central aspect of the subsequent analysis. As in Liu et al. (2007), its value is determined by a series of M densities, maximising the log-likelihood of the realization of asset prices:

$$log(L(S_{T,1}, S_{T,2}, ..., S_{T,M} \mid \gamma)) = \sum_{i=1}^{M} log(f_{P,i}(S_{T,i} \mid \hat{\theta}_i, \gamma)) \qquad (2.2.4)$$

In the empirical analysis, three different density types are considered with different degrees of flexibility: Lognormal density (LN), mixture-lognormal density (MLN) and generalized beta density of the second kind (GB2). Their characteristics and transformation are outlined in the subsequent subsections.

2.2.1 Mixtures of lognormal densities

Ritchey (1990) introduces a mixture of lognormal densities, MLN herafter, to evaluate option prices. Following Liu et al. (2007), this distribution is defined as a weighted combination of two lognormal densities. A European call price is evaluated as the sum of two Black (1976) models:

$$c(K \mid \theta) = wc_B(F_1, T, K, r, \sigma_1) + (1 - w)c_B(F_2, T, K, r, \sigma_2). \qquad (2.2.5)$$

with $\theta = (F_1, \sigma_1, F_2, \sigma_2, w)$. The risk-neutrality constraint for the MLN reduces the amount of free parameters to four and by stating that:

$$F = wF_1 + (1 - w)F_2, \qquad (2.2.6)$$

with F the forward price. Transforming the density into its subjective counterpart as in (2.2.3) leads to another mixture-lognormal density. The subjective parameters become:

$$\begin{aligned} &\tilde{\theta} = (\tilde{F}_1, \tilde{F}_2, \sigma_1, \sigma_2, \tilde{w}), \qquad (2.2.7) \\ &\tilde{F}_i = F_i e^{\gamma\sigma^2 T}, \text{ for } i = 1, 2, \text{ and} \\ &\frac{1}{\tilde{w}} = 1 + \frac{1 - w}{w} \left(\frac{F_2}{F_1}\right)^{\gamma} exp\left(0.5(\gamma^2 - \gamma)(\sigma_2^2 - \sigma_1^2)T\right). \end{aligned}$$

The lognormal density (LN) is a special case of the above for $w = 1$. As a result, this distribution is not described separately.

2.2.2 Generalized Beta Densities

The generalized beta density of the second kind, GB2 hereafter, was introducted by Bookstaber and McDonald (1987) and was applied to option prices by Liu et al. (2007). They show that a European call option price can be evaluated by:

$$c(K \mid \theta) = Fe^{-rT}[1 - F_\beta(v(K,a,b) \mid p + 1/a, q - 1/a) - Ke^{-rT}[1 - F_\beta(v(K,a,b) \mid p, q], \tag{2.2.8}$$

with $\theta = (a, b, p, q)$, all greater than zero, $v(K, a, b) = (x/b)^a/[1 + (x/b)^a]$ and F_β the distribution function of a random variable following a beta distribution. The risk-neutrality constraint reduces the amount of free parameters from four to three:

$$F = \frac{bB(p + 1/a, q - 1/a)}{B(p, q)}\ ,\ aq > 1, \tag{2.2.9}$$

where B is defined in terms of gamma functions $B(p, q) = \Gamma(p)\Gamma(q)/\Gamma(p + q)$. The transformation of the GB2 risk-neutral density as in (2.2.3) leads a different type of GB2 density function. The subjective parameters are then $\tilde{\theta} = \left(a, b, p + \frac{\gamma}{a}, q - \frac{\gamma}{a}\right)$ with $aq > \gamma$.

2.3 Risk Aversion Evaluation

One of the biggest challenges in the evaluation of risk-aversion estimates is the limited availability of data. While option prices provide an accurate description of risk-neutral market expectations, subjective dynamics can only be observed as single price outcomes. Linking a single risk-neutral density to realized asset prices to estimate the degree of risk aversion of a representative agent produces high variation in the parameter estimates. One way of overcoming this issue is to obtain risk aversion over several periods, as presented in (2.2.4). However, this approach requires access to a sufficiently long dataset. In situations where the time horizon is too short, estimates might be misleading and will not accurately reflect the risk preferences of the representative agent.

Bliss and Panigirtzoglou (2004) introduce a Monte-Carlo simulation method and a Bootstrapping procedure to gain better insight into the variation of risk-aversion estimates. However, their approach does not overcome the problem of single price realisations. More importantly, risk aversion is only estimated from past data, which may accurately reflect individuals' past preferences but not necessarily those relevant for the forthcoming period.

This study proposes an alternative Monte-Carlo simulation exercise that addresses all these problems. The simulation setting is structured in a way that only considers the most recent risk-neutral dynamics and simulates potential asset price outcomes based on different stochastic processes. This allows to estimate risk aversion by taking into account current market dynamics. Importantly, it is possible to simulate several asset price outcomes for the same time period providing a more detailed outlook of the anticipated variation of the underlying risk aversion.

2.3.1 Simulation framework

In the simulation framework current market dynamics are captured by a set of assumptions on risk-neutral and subjective processes. Given these, a whole range of potential asset price outcomes can be simulated and used to estimate risk aversion for a specific period. Simulation-based risk-aversion estimates provide better insight into the full spectrum of potential risk-aversion estimates. Accordingly, the utility function can be validated. The Monte-Carlo simulation framework consists of the following five steps[1]:

1. Make assumptions about the subjective process that includes the utility function of the representative agent and the risk-neutral process.
2. Obtain options from the risk-neutral process and fit a risk-neutral density into option prices. Since, the risk-neutral process does not change, estimated risk-neutral parameter values do not change for the time horizon of M periods.
3. Generate M asset price outcomes based on the subjective process for M risk-neutral

[1]Note that for the empirical application, Step 2 does not draw option prices from the risk-neutral process since options can be observed directly in the market.

densities.

4. Estimate the risk-aversion parameter γ over the estimation horizon of M periods based on (2.2.4) .

5. Repeat Steps 3 and 4 H times.

2.3.2 Likelihood criteria

Log-likelihood based tests on the asset price outcomes can be useful in providing insights on the quality of the retrieved risk-aversion values, as used in the work of Bao et al. (2007), Liu et al. (2007) and Shackleton et al. (2010) to compare the quality of their density methods on equity indices. The log-likelihood function is defined as:

$$log\ L_m = \sum_{t=1}^{i} log(f_{m,t}(S_{t+1})) \tag{2.3.1}$$

where f is the density function, m the density type and i the number of densities. The log-likelihood function serves two different testing purposes. First, a log-likelihood ratio test is constructed to assess statistically whether the additional risk-aversion parameter γ significantly improves the forecasting performance of the densities:

$$H_0\colon \gamma = 0 \tag{2.3.2}$$

$$H_1\colon \gamma \neq 0. \tag{2.3.3}$$

The corresponding test statistic is defined as:

$$2^{*}LR = -2(L_{m,\gamma=0} - L_{m,\gamma\neq 0}). \tag{2.3.4}$$

Under the null $2^{*}LR \sim \chi^2$ distribution with 1 degree of freedom.
Second, it is tested whether the two competing methods, denoted as m and n, provide equal expected log-likelihoods as described:

$$H_0\colon E[log(f_{m,t}(S_{t+1}))] = E[log(f_{n,t}(S_{t+1}))] \tag{2.3.5}$$

$$H_1\colon E[log(f_{m,t}(S_{t+1}))] \neq E[log(f_{n,t}(S_{t+1}))] \tag{2.3.6}$$

This test is proposed by Amisano and Giacomini (2007), AG hereafter, and was applied to option based densities by Shackleton et al. (2010). The AG test is an out-of sample test, which uses a rolling window estimation scheme to rank density forecasts. Its loss function is based on a weighted Likelihood ratio test and offers the advantage that it can be robust to heterogeneous data sets, while it can be applied to densities obtained from both parametric and non-parametric models. The test is based on the difference of log-likelihood of each period, denoted by d_t. The test statistic, AG, is retrieved on the basis of the average difference $\hat{d}$ and standard deviation of s_d of d with integer times $i, \ldots, j$, as:

$$AG = (j - i + 1)^{0.5} \frac{\hat{d}}{s_d}. \tag{2.3.7}$$

Under the null hypothesis H_0, it follows that $AG \sim N(0, 1)$.

2.3.3 Market assumptions and estimation

Two different parametric forms are imposed to model market dynamics. First, market dynamics are assumed to follow a Geometric Brownian Motion as in Black and Scholes (1973). Second, stochastic volatility is incorporated into the model by implementing the model by Heston (1993). In both cases, it is assumed that the volatility parameters are derived from option prices including also the subjective dynamics. Therefore, within the proposed simulation framework there is no adjustment for the volatility premium as it is commonly observed for option prices (Bekaert and Hoerova, 2014). Only the drift parameter is changed between specifications in order to account for the risk-premium of the subjective approach. This transformation is similar to the risk-premium transformation in Shackleton et al. (2010). The drift term of the subjective process is estimated using historical asset price returns under the assumption of lognormality of stock prices. An extension of the proposed methodology would be to account for volatiltiy premia. Shackleton et al. (2010) provides a procedure that adds a linear drift term to the variance process of the Heston (1993) model to account for volatility premia. Finally, it is illustrated how expected risk aversion can be obtained by taking into account the risk-neutral and subjective dynamics.

2.3.3.1 Geometric Brownian Motion

In the first specification, it is assumed that markets follow a Geometric Brownian Motion as in Black and Scholes (1973):

$$dS_t = \mu S_t dt + \sigma S_t dW_t \quad \text{with } W_t \sim \mathcal{N}(0,1) \tag{2.3.8}$$

with μ the drift term and σ the constant volatility. Option prices can be evaluated by the well-known Black and Scholes formula as follows.

$$c(X) = S_t e^{-qT} N(d_1) - X e^{-rT}(d_2) \tag{2.3.9}$$

$$d_1 = \frac{ln\left(\frac{S_t}{X}\right) + \left(r - q + \frac{\sigma^2}{2}\right)T}{\sigma\sqrt{T}} \tag{2.3.10}$$

$$d_2 = d_1 - \sigma\sqrt{T}. \tag{2.3.11}$$

The parameter estimation is performed as in (2.2.2) by only considering a single at-the-money option.

2.3.3.2 Stochastic Volatility

An alternative to the geometric Brownian motion is the stochastic volatility model of Heston (1993). This specification allows for more flexibility and captures more accurately observed market dynamics. As in the previous model, the change in asset prices follows a Geometric Brownian Motion as in (2.3.8), whereas the volatility parameter σ follows an Ornstein-Uhlenbeck process.

$$d\sigma_t = \kappa(\theta - v_t)dt + \xi\sqrt{v_t}dZ_t \quad \text{with } Z_t \sim \mathcal{N}(0,1) \tag{2.3.12}$$

where μ denotes the drift term, σ the volatility of stock price, κ the mean reversion speed, θ the long-run variance, v_t the current variance and ξ the volatility of volatility parameter. Heston (1993) provides a closed form solution to price a European call option in the form of the Black and Scholes formula:

$$c(X) = Se^{-qT}P_1 - Xe^{-rT}P_2 \tag{2.3.13}$$

In contrast to the estimation of the risk-neutral density, the risk aversion parameters are not obtained directly from option prices as in (2.2.2). The Heston model is estimated by minimising the squared error between market implied volatility and model implied volatility[2].

2.3.3.3 Expected risk-aversion

In line with Jackwerth (2000), Aït-Sahalia and Lo (2000) and Rosenberg and Engle (2002), the risk-aversion coefficient can be determined by the relation of the utility function to the risk-neutral and objective density function. The absolute risk-aversion, RA, can be then defined as:

$$RA(x) = -\frac{u''(x)}{u'(x)} = \frac{f_P'(x)}{f_P(x)} - \frac{f_Q'(x)}{f_Q(x)} \tag{2.3.14}$$

Using this result, the relative risk-aversion, γ, can be retrieved by:

$$\gamma = xRA(x). \tag{2.3.15}$$

Since it is assumed that the transformation from risk-neutral to subjective dynamics only impacts the mean return, the risk-aversion estimate is the same at every point x.

2.4 Simulations

This simulation exercise aims to unveil the influence of different characteristics of the underlying asset price process on risk aversion. As a starting point, the conventional Geometric Brownian Motion is used to model asset price dynamics and then extended to accommodate stochastic volatility, as discussed in section 2.3. The simulations provide useful insights about how and to what extent risk-aversion estimates are influenced by different factors within the considered models. The main point of interest is the influence of stochastic volatility on the risk-aversion estimates.

[2] The required calculations for this method are computationally extensive and the weighting approach by Christoffersen et al. (2009) helps to reduce the amount of computations (see appendix A.1)

Therefore, the simulations below include six different scenarios that examine the impact on risk aversion, as depicted in Panels A-F in the tables 2.1 and 2.2, and figure 2.1, 2.2 and 2.3. Simulation A evaluates the change in the mean return, while Simulation B evaluates the impact of volatility, under a Geometric Brownian Motion. The remainder panels focus on the presence of stochastic volatility. The volatility process is modelled by an Ornstein-Uhlenbeck process. Simulation C investigates on the effect of the mean reversion κ of volatility. Simulation D evaluates the impact of fat-tailed distributions via the volatility of volatility ξ. Lastly, Simulation E and F evaluate the impact of skewness in conjunction with a fat-tailed distribution. The latter might be the empirically most relevant case which can provide guidance for practical purposes.

Each simulation scheme derives risk aversion over an estimation horizon of $H = 120$ monthly periods. As described in section 2.3, the risk-neutral and subjective parameters are kept constant over the entire estimation horizon to ensure only the current market dynamics are considered. The simulated option prices are obtained one month prior to expiry. One month consists of 21 trading days, which results in 252 trading days per year. Each day is discretised by ten sub-steps to assure a sufficient precision of the simulated asset prices. The entire simulation is repeated $N = 5,000$ times.

2.4.1 Geometric Brownian Motion

Within the Geometric Brownian Motion, both the mean, μ, and the variance,σ^2, influence the mean and variation of the simulated risk-aversion estimates, respectively. Figure 2.1, 2.2 and 2.3 A displays the case of an stepwise increase of μ from -0.5% to 4%. It is apparent that risk-aversion increases when the expected return rises since the compensation for risk becomes higher. Consequently, the mean of the subjective density shifts to reflect the premium in the underlying asset price process.

Figure 2.1, 2.2 and 2.3 B displays the results of a change in variance. Compared to the previous case an increase in the variance does not result in a major increase of risk-aversion. The variance mainly influences the variation of risk-aversion. However, this

	median			average			MSE			
E[γ]	log	mix-log	GB2	log	mix-log	GB2	log	mix-log	GB2	simulation
0.00	0.00	0.07	0.02	0.01	0.03	0.03	10.17	**8.48**	9.58	A
1.50	1.50	1.25	1.47	1.48	1.43	1.46	9.73	**7.11**	9.18	A
2.50	2.60	2.59	2.54	2.59	2.58	2.53	9.91	9.78	**9.29**	A
3.50	3.46	3.42	3.36	3.49	3.45	3.40	10.18	9.91	**9.50**	A
4.50	4.51	3.83	4.38	4.48	4.15	4.37	9.88	6.10	**9.22**	A
0.00	0.08	0.13	0.09	0.09	0.11	0.11	9.86	**8.74**	9.27	B
0.33	0.35	0.35	0.35	0.37	0.37	0.38	3.40	3.39	**3.21**	B
0.40	0.37	0.37	0.38	0.38	0.38	0.38	1.93	1.93	**1.84**	B
0.43	0.47	0.47	0.47	0.47	0.47	0.47	1.45	1.45	**1.38**	B
0.44	0.42	0.42	0.43	0.42	0.42	0.42	1.13	1.13	**1.07**	B
0.00	-0.08	0.04	-0.11	-0.05	0.03	-0.10	41.77	**26.22**	29.63	C
0.00	-0.12	-0.01	-0.03	-0.15	-0.05	0.03	20.95	**14.07**	16.58	C
0.00	-0.06	0.00	-0.01	-0.07	-0.01	-0.04	17.23	13.73	**12.83**	C
0.00	-0.05	0.01	0.00	-0.10	-0.04	0.02	14.53	11.44	**11.21**	C
0.00	-0.05	0.00	-0.05	-0.06	-0.01	-0.03	13.24	11.01	**10.92**	C
0.00	0.03	0.05	0.11	-0.01	0.01	0.12	10.12	**9.23**	9.48	D
0.00	0.02	0.03	0.04	0.02	0.03	-0.01	11.38	10.74	**10.31**	D
0.00	-0.07	0.11	0.05	-0.05	0.01	-0.03	15.75	**9.85**	13.58	D
0.00	-0.10	-0.04	-0.07	-0.02	0.02	-0.07	26.32	20.87	**20.39**	D
0.00	-0.13	-0.03	0.00	-0.14	-0.06	0.01	43.43	**30.26**	30.43	D
0.00	-0.07	0.14	-0.04	-0.05	-0.02	-0.04	9.97	**7.17**	9.22	E
0.00	-0.05	-0.03	-0.02	0.05	0.05	0.06	9.82	9.61	**9.11**	E
0.00	0.01	0.01	0.05	0.17	0.17	0.18	10.62	10.51	**9.79**	E
0.00	0.04	0.08	0.08	0.21	0.22	0.22	10.11	9.31	**9.29**	E
0.00	-0.04	-0.01	0.02	0.19	0.20	0.20	10.21	9.67	**9.24**	E
0.00	-0.04	0.13	0.15	0.15	0.20	0.19	43.50	**29.65**	30.79	F
0.00	-0.02	0.22	0.29	0.36	0.38	0.38	42.20	**28.61**	30.06	F
0.00	-0.30	0.10	0.17	0.46	0.46	0.46	45.14	**28.50**	30.14	F
0.00	-0.02	0.41	0.53	0.82	0.77	0.78	45.08	**27.45**	29.34	F
0.00	-0.27	0.31	0.45	0.73	0.69	0.70	48.53	**27.89**	30.13	F

Table 2.1: The table provides the results for all 6 simulations. The expected value of gamma $E[\gamma]$ is obtained from the theoretical result in section 2.3. The average, median and mean squared error for all density types are the results of the simulation. The bold values for the mean squared error mark the lowest values in each row.

log		mix-log		GB2		
RND	RWD	RND	RWD	RND	RWD	simulation
0.57	**1.08**	0.48	0.96	0.00	0.49	A
0.47	**1.07**	0.42	0.96	0.00	0.58	A
0.46	**1.29**	0.46	1.28	0.00	0.81	A
0.45	**1.56**	0.40	1.50	0.00	1.08	A
0.43	**1.94**	0.37	1.72	0.00	1.47	A
0.45	**0.95**	0.34	0.81	0.00	0.48	B
0.50	**1.03**	0.50	1.03	0.00	0.52	B
0.42	**0.94**	0.42	0.94	0.00	0.51	B
0.43	**1.02**	0.43	1.02	0.00	0.57	B
0.46	1.05	0.46	**1.05**	0.00	0.58	B
0.00	1.81	111.32	112.80	119.66	**121.21**	C
0.00	0.92	42.47	43.23	43.48	**44.31**	C
0.00	0.77	24.70	25.39	28.18	**28.83**	C
0.00	0.66	19.04	**19.62**	18.81	19.38	C
0.00	0.61	13.26	13.81	14.93	**15.48**	C
0.00	0.50	0.48	0.97	1.97	**2.45**	D
0.00	0.56	8.84	9.38	10.75	**11.27**	D
0.00	0.75	22.30	22.92	29.28	**29.97**	D
0.00	1.20	56.31	57.39	68.88	**69.96**	D
0.00	1.88	102.91	104.52	124.70	**126.30**	D
0.00	0.50	0.27	0.70	2.14	**2.61**	E
0.00	0.49	0.81	1.30	2.18	**2.65**	E
0.00	0.53	0.52	1.04	2.21	**2.72**	E
0.00	0.50	0.74	1.22	2.23	**2.72**	E
0.00	0.51	0.98	1.48	2.38	**2.86**	E
0.00	1.88	99.49	101.07	119.19	**120.82**	F
0.00	1.81	95.83	97.34	117.18	**118.73**	F
0.00	1.91	104.51	106.06	117.58	**119.18**	F
0.00	1.89	99.33	100.82	112.75	**114.29**	F
0.00	2.00	104.96	106.50	113.51	**115.09**	F

Table 2.2: The table contains the average maximum log-likelihood estimate based on 2.2.4 for the risk-neutral density (RND) and subjective density (RWD). The bold values mark the highest average value of log-likelihood in each row.

relation is inverse, meaning when variance is high (low) the variation of risk aversion is low (high). This feature can be confirmed by (2.2.7), which clearly illustrates that the risk-premium is linked to the product of volatility and risk-aversion. Thus, an increase in volatility results in lower variation of risk-aversion since the risk-aversion has to change less to achieve a specific premium. Overall, the different density types do not exhibit major differences between each other.

This is also the case for the mean squared error in table 2.1. However, results change across specifications of the simulation. In simulation A, the mean squared error is similar independently of the values for μ. In contrast, the mean squared error varies heavily in simulation B due to the changes in variance. Generally, increases in variance reduce the mean squared error. For both simulations, A and B, mean and median infer that risk aversion is distributed symmetrically. This is due to the restrictive assumptions of the Geometric Brownian motion which does not allow for asymmetries in the asset price distribution. The log-likelihood, in table 2.2, supports this finding by obtaining the highest values for the lognormal distribution. This is not surprising because the underlying asset price assumptions infer a lognormal density. Quite similar are the values of the mixture-lognormal and GB2 density. Both densities can resemble a lognormal density as a special case and therefore values can be close to each other. The observed differences can be most likely traced back to the estimation error. Generally, the subjective density obtains the highest values since the assumed asset price dynamics infer positive risk-aversion.

2.4.2 Stochastic volatility

Extending the Geometric Brownian motion to accommodate stochastic volatility allows to further analyse the impact of volatility in the underlying asset. The focus is on the effect of the following parameters: the speed of mean reversion to the long run volatility, κ, the volatility of volatility, ξ, and the correlation between the two Wiener processes of the stock price and volatility, ρ. Figure 2.1, 2.2 and 2.3 C displays the results across different values of κ. The results suggest that an increase in κ has a similar influence as an increase in variance which can straightforwardly be explained by the mechanics of κ. A high value of κ results in a faster reversion to the long run mean of the variance.

In contrast, the parameter ξ has an opposite effect on risk aversion, as shown in figure 2.1, 2.2 and 2.3 D. Increasing ξ results in more variation of risk-aversion. This can be attributed to the fact that ξ models the tail thickness of the asset price distribution. The simulation unveils that a higher frequency of extreme outcomes outweighs the impact of higher volatility, resulting in higher variation of risk-aversion.

Simulations E and F elaborate on the remaining parameter ρ, which captures asymmetry in the asset price distribution. Figure 2.1, 2.2 and 2.3 E and F exhibit the influence of ρ on the risk aversion and its interaction with ξ. Figure 2.1, 2.2 and 2.3 E assumes a low value of ξ, whereas in figure 2.1, 2.2 and 2.3 F ξ is relatively high. Both plots show that with an increase in ρ the asymmetry of the relative risk-aversion coefficient increases. This can be seen by the diverging mean and median parameters but also in the shrinking upper percentiles. The higher ξ is the greater is the absolute difference between median and average.

Overall, the more flexible mixture-lognormal and GB2 density are more capable to deal with the asymmetric shape of the underlying distribution. In all of the considered simulations assuming stochastic volatility the distribution of risk-aversion is more symmetric and displays a lower variation. This is also reflected in the mean squared error. It clearly favours the more flexible mixture-lognormal and GB2 density. However, between both density types it is not clear which one is favourable. This might be rooted in the relatively restricted assumptions about the underlying asset price process, which could be extended to allow for even greater flexibility.

The simulation also provides insights on the maximum likelihood estimates and compares the subjective to the risk-neutral likelihoods. The results are tabulated in table 2.2 and draw a more complete picture along with results of the Mean Squared Error in table 2.1. In simulation A and B likelihood values are fairly close to each other. However, this changes from simulation C and ongoing. It shows that the GB2 density clearly fits the data best, enforcing that a flexible density type is preferable when the underlying asset price distribution moves further away from lognormal. Furthermore, when the subjective density is compared to the risk-neutral density the likelihood is in all cases higher suggesting an improved fit. This comes as a surprise since risk-aversion is zero. The additional risk-aversion parameter provides more flexibility when transforming the data, which results on average in a better fit of the simulated asset price outcomes. Therefore, the log-likelihood measure of the subjective density exceeds the values of the risk-neutral density. However, those differences are of minor magnitude.

The mean squared error underlines the findings that more flexible density types help to obtain more stable estimates of risk-aversion. As for the boxplot the variation of risk-aversion is lower. However, comparing the mixture-lognormal and GB2 density with each other does not lead to clear results. Depending on the simulation, one method is preferable over the other.

Collectively, simulation results indicate that the main influence on the certainty of risk-aversion may be attributed to the volatility and the volatility of volatility. The mean mainly shifts the risk-aversion estimate in one or the other direction. In addition, flexible class of densities enable to obtain more accurate estimates for risk-aversion when the asset price distribution is skewed and has excess kurtosis. The fact that more flexible density types are favourable could also be shown by the likelihood comparison of the simulation. Finally, variation of risk-aversion strongly depends on the volatility component. Beyond that it could be shown that asymmetry in the asset price distribution also skews risk-aversion estimates. All these results provide useful insights to practitioners and highlight the importance of accounting for volatility when evaluating risk aversion estimates.

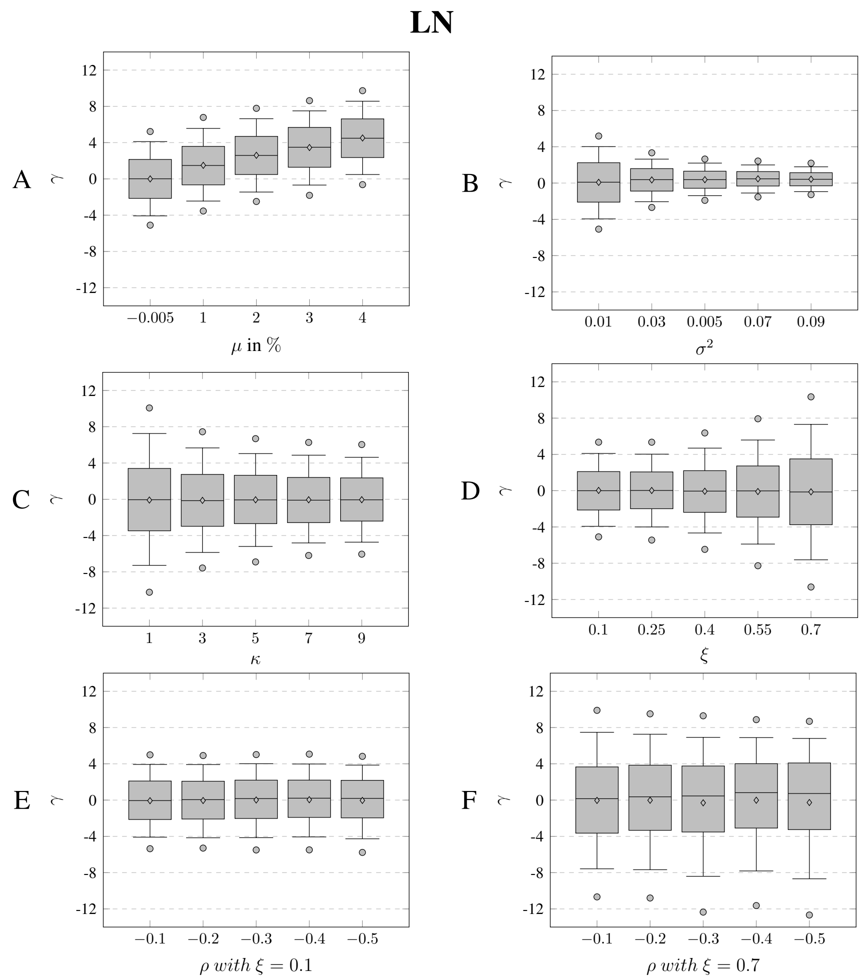

Figure 2.1: Distribution of estimated γs applying the lognormal density. A and B assume geometric Brownian motion as underlying process. The remaining plots are simulated applying the stochastic volatility model. The body of box plot is defined by the interquartile range of γ estimates. The horizontal line dividing the box and diamond mark the median and average, respectively. The "whiskers" define the 10^{th} and 90^{th} percentile and the dots the 5^{th} and 95^{th} percentile. If not separately specified in the graph the parameter values for $\mu = -0.005$, $r = 0$, $q = 0$ and $\sigma^2 = 0.01$. For the case of the stochastic volatility model $\kappa = 1$, $\theta = 0.01$, $\xi = 0.7$ and $\rho = 0$.

MLN

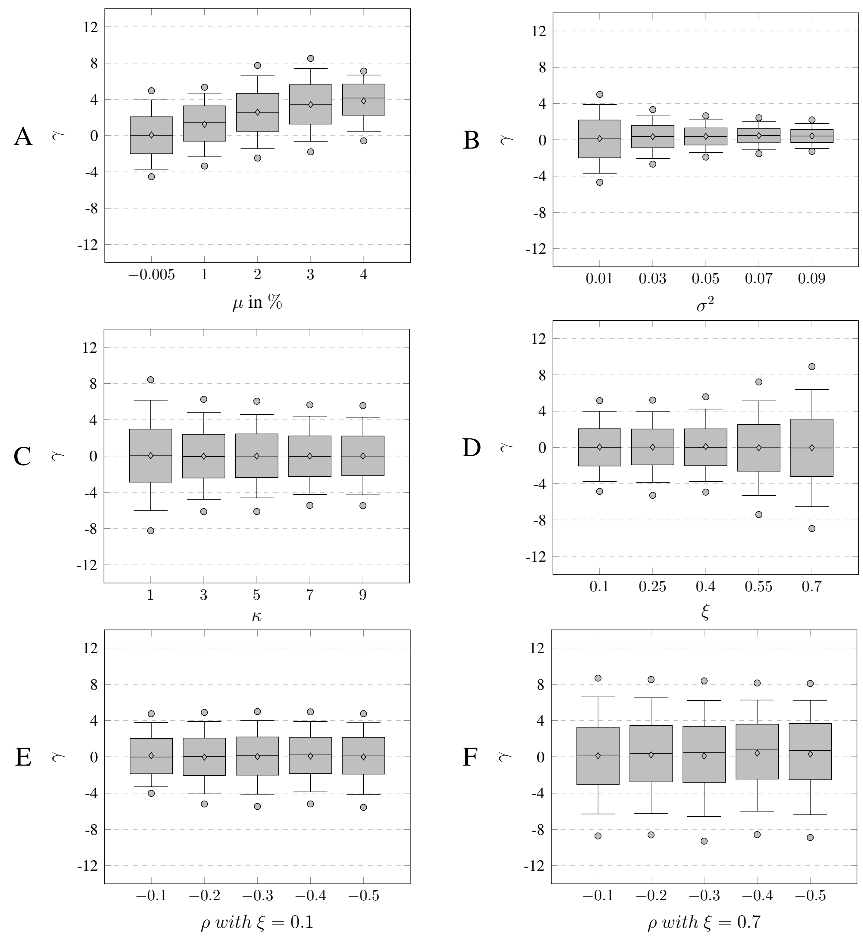

Figure 2.2: Distribution of estimated γs applying the mixture-lognormal density. A and B assume geometric Brownian motion as underlying process. The remaining plots are simulated applying the stochastic volatility model. The body of box plot is defined by the interquartile range of γ estimates. The horizontal line dividing the box and diamond mark the median and average, respectively. The "whiskers" define the 10^{th} and 90^{th} percentile and the dots the 5^{th} and 95^{th} percentile. If not separately specified in the graph the parameter values for $\mu = -0.005$, $r = 0$, $q = 0$ and $\sigma^2 = 0.01$. For the case of the stochastic volatility model $\kappa = 1$, $\theta = 0.01$, $\xi = 0.7$ and $\rho = 0$.

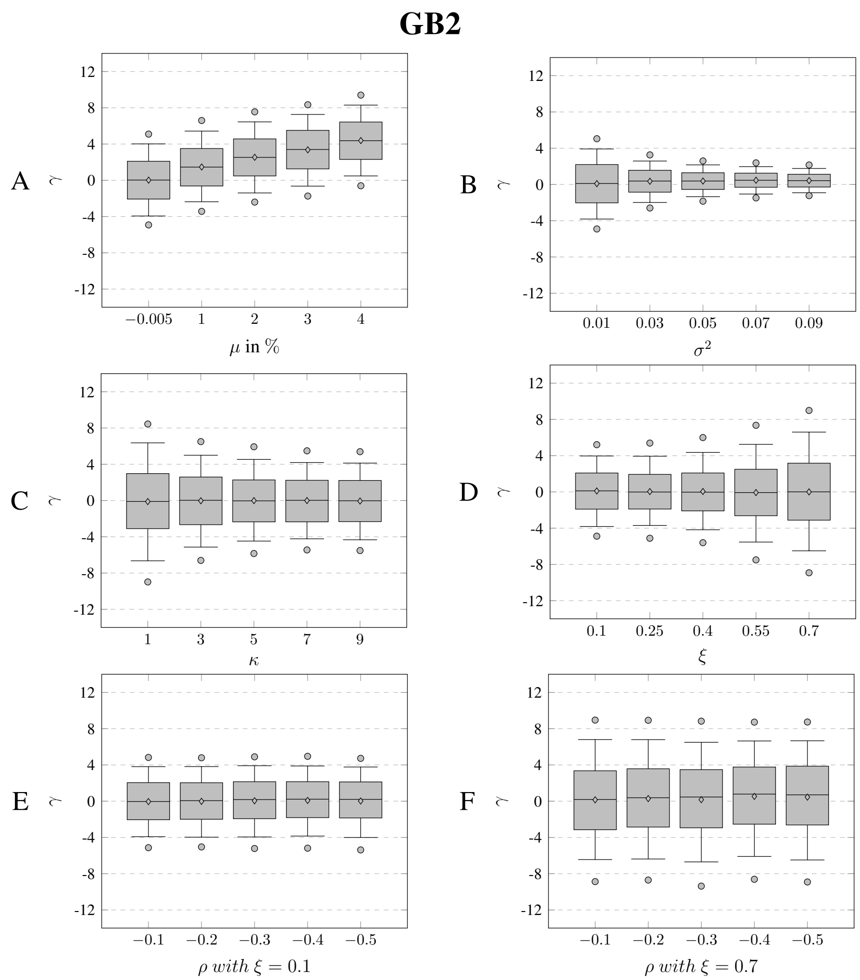

Figure 2.3: Distribution of estimated γs applying the GB2 density. A and B assume geometric Brownian motion as underlying process. The remaining plots are simulated applying the stochastic volatility model. The body of box plot is defined by the interquartile range of γ estimates. The horizontal line dividing the box and diamond mark the median and average, respectively. The "whiskers" define the 10^{th} and 90^{th} percentile and the dots the 5^{th} and 95^{th} percentile. If not separately specified in the graph the parameter values for $\mu = -0.005$, $r = 0$, $q = 0$ and $\sigma^2 = 0.01$. For the case of the stochastic volatility model $\kappa = 1$, $\theta = 0.01$, $\xi = 0.7$ and $\rho = 0$.

2.5 Empirical application

2.5.1 Data and Estimation

This study considers European call and put options, which can only be exercised at the maturity date. Option prices are obtained from Optionmetrics for the period from January 1996 until May 2016 on the S&P 500 index. The considered expiry dates are 28 days before the third Friday of each month. Only options 28 days before expiry are considered to ensure that estimation horizons are non-overlapping. In case that there are no option prices available for the 28 days before expiry, the observation date is shifted by one day. If there are still no options available the observation date is dropped from the sample. This results in 242 observation dates for the time period under study. The relevant option prices are selected based on their moneyness. All call (put) options with a moneyness greater (smaller) than 1.03 (0.97) are removed from the sample since they are less frequently traded. Moneyness is defined as the ratio between strike price and underlying stock price. Furthermore, if there is no traded volume the option is removed as well. The estimation of the risk-neutral density is then performed taking into account the mid option price. Put options are transformed into call-options by the put-call parity relationship. In case there are two options available for the same strike price, the average mid price is used. In case dividends and zero-coupon rates do not match the maturity date, rates are interpolated between the two closest dates. Future prices F are calculated synthetically by the pricing equation $F_t = S_t e^{(r_t - q_t)T}$.

The risk-neutral density parameters are estimated minimising (2.2.2) at each observation date. Multiple start values were supplied into the minimisation problem to ensure the minimum is robust. Table 2.4 provides the summary statistics of the option pricing errors. The reported values are in a similar range as, for example, Shackleton et al. (2010). There are remarkable differences between the more flexible densities and the lognormal approach. The pricing errors are attributed to the greater flexibility of the MLN and GB2, which better replicate the non-normal distribution contained in option prices.

	Total number	Average	Max	Min
Calls	7,411	31	104	7
Puts	12,857	53	264	11
Overall	20,268	84	350	20

Moneyness	S/K	Total number	In %
Deep OTM put	>1.10	6,629	32.71
OTM put	1.03-1.10	3,106	15.32
Near the money	0.97-1.03	6,342	31.29
OTM call	0.90-0.97	3,159	15.59
Deep OTM call	<0.90	1,032	5.09
Overall		20,268	100.00

Table 2.3: Total number of options and options by moneyness.

Density	Median	Average	Standard deviation
LN	3.21	4.89	6.04
MLN	0.10	0.16	0.17
GB2	0.09	0.17	0.22

Table 2.4: Mean squared error of 2.2.2 between model and observed option prices.

2.5.2 Results

The empirical results, applying the simulation method to observable S&P 500 index options, are presented in figure 2.4 and 2.5. The graphs display the mean and median simulated risk-aversion estimate enclosed by confidence intervals. Furthermore, the estimation using a sequence of risk-neutral densities and asset price outcomes as in Bliss and Panigirtzoglou (2004) and Liu et al. (2007) is displayed. This procedure is denoted as "historical risk-aversion".

For the case of the Geometric Brownian Motion, the results of the different density types are similar with no major discrepancies. Overall, the simulation method suggests a symmetric distribution of risk aversion since the mean and median are close to each other and the percentiles are evenly distributed. Furthermore, during times of high volatility, as for example at the subprime crisis, the variation in risk aversion lowers. These observations confirm the results from the simulation.

More complex are the findings assuming stochastic volatility for the price process. The greater flexibility of the underlying process allows for an asymmetric shape. While in the previous case the mean and median for all density types are close together, they diverge for the lognormal case. The median tends to be in excess towards the median, imposing a negative skew. This observation is a result of the negative skew induced by the stochastic volatility process. The lognormal distribution is not able to accommodate the more complex shape of the underlying probability distribution. Therefore, the estimation of the risk-aversion parameter compensates for this inflexibility, resulting in a skewed distribution of risk aversion. This observation highlights the importance of flexible density types when trying to estimate risk aversion. This finding lines up with the characteristics of the simulation, that more flexible density types yield a lower mean square error. Furthermore, the confidence intervals have a stronger variation compared to the Geometric Brownian Motion. This is not a surprise due to the greater complexity of stochastic volatility. However, remarkable are the differences in the lower confidence intervals. The MLN and GB2 swing far less into negative values of risk aversion, which further highlights the importance of flexible density types.

	Excess log-likelihood			2*LR		
Method	LN	MLN	GB2	LN	MLN	GB2
RND	0.00	17.04	96.44			
Hist.	1.04	18.03	97.43	2.09	1.69	1.98
GBM	1.20	18.17	97.58	2.40	2.26	2.28
SV	1.18	18.17	97.58	2.36	2.25	2.27

Table 2.5: The excess log-likelihood is defined as excess value compared to the weakest method. In this case the worst performing method is the risk-neutral LN. The remaining methods denote transformed subjective densities.

Overall, it is remarkable how uncertain the estimates are even within basic assumptions as displayed in figure 2.4 and 2.5 . Generally, the simulation-based methods tend to exceed the risk-aversion estimates of the historical approach. Nevertheless, the confidence intervals mostly encompass the historical estimates, particularly when assuming stochastic volatility.

To gain further insight, the obtained mean risk-aversion estimate is used to transform the risk-neutral density into its subjective counterpart. Afterwards,the predictability is evaluated based on the log-likelihood of the final asset price outcomes. This results in 122 evaluation periods since the historical method retrieves the estimates using a 10 year rolling window. The log-likelihood is reported in table 2.5 in excess to the risk-neutral LN. The reported values can clearly show that the more flexible density types are preferable towards the LN. Furthermore, the GB2 is clearly in excess towards the MLN. The additional parameter γ is further tested using the log-likelihood ratio test. All methods are significant at the 20% level. The reason for the low significance level could be rooted in the short sample size. Liu et al. (2007) encounter similar difficulties and argue the insignificance is a type II error. These authors base this claim on the findings of Merton (1980) and others since the risk-premium is small relative to volatility and, therefore, difficult to estimate.

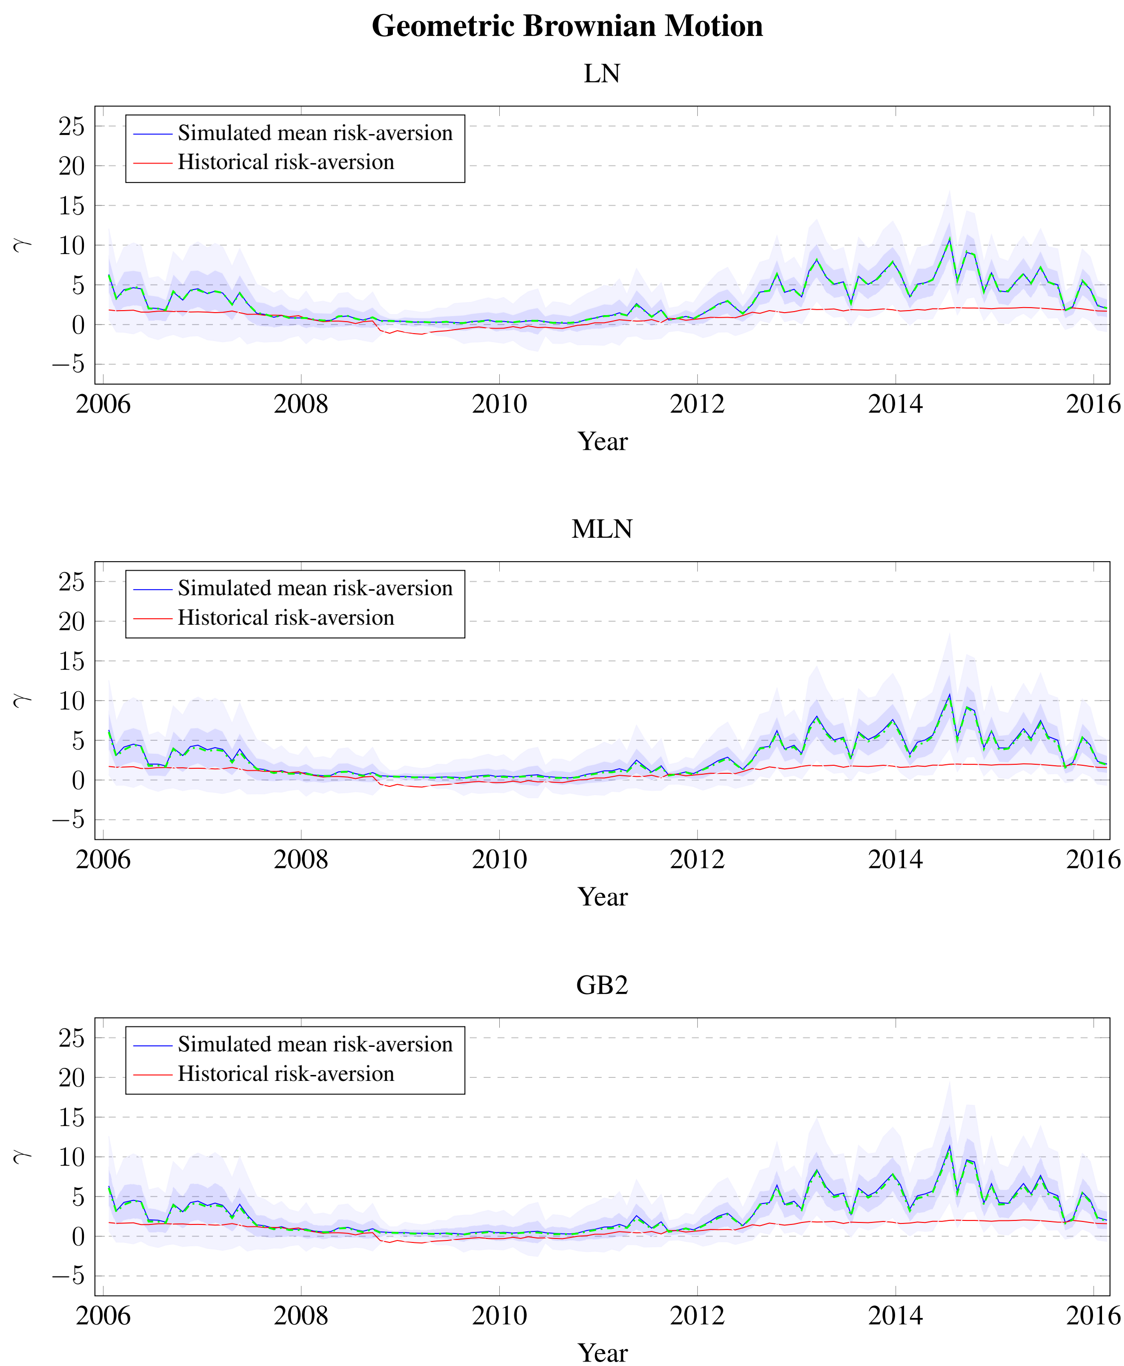

Figure 2.4: Risk-aversion simulation results and historical estimates. The simulation assumes the asset price follows a geometric Brownian motion. The dark blue line denotes the average estimate of risk-aversion whereas the green dashed-dotted line denotes the median. The inner confidence interval (dark blue) defined by the 25^{th} and 75^{th} percentile. The outer confidence interval (light blue) is defined by the 5^{th} and 95^{th} percentile. The red line defines the risk-aversion obtained by the classic method.

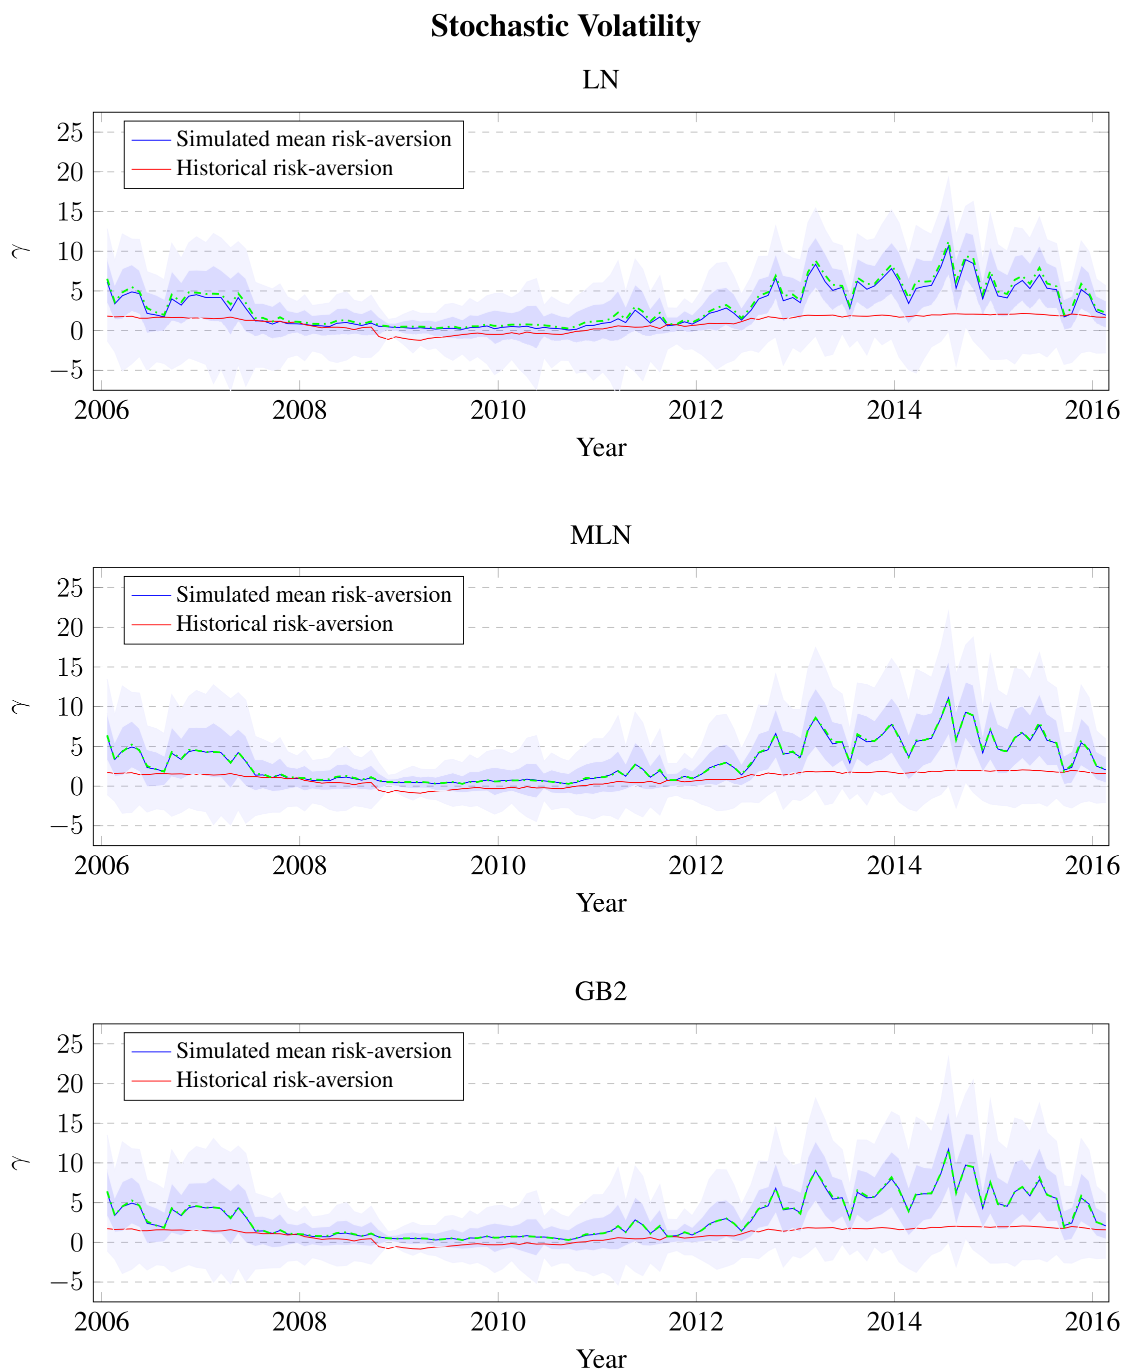

Figure 2.5: Risk-aversion simulation results and historical estimates. The simulation assumes the asset price follows stochastic volatility. The dark blue line denotes the average estimate of risk-aversion whereas the green dashed-dotted line denotes the median. The inner confidence interval (dark blue) defined by the 25^{th} and 75^{th} percentile. The outer confidence interval (light blue) is defined by the 5^{th} and 95^{th} percentile. The red line defines the risk-aversion obtained by the classic method.

Finally, two competing forecasts are tested for equal expected log-likelihoods following Amisano and Giacomini (2007). The results are reported in table 2.6. Comparing the different density types with each other clearly shows that the GB2 is the preferable method. The test rejects the null hypothesis at the 1% level in all cases. Within each density type findings are less clear. Only for the historical risk aversion, the subjective density is significantly different from the risk-neutral density. In none of the cases, the proposed simulation method is significantly different to the risk-neutral density. A possible explanation for this observation could be that the volatility dynamics are assumed to be the same in the risk-neutral and subjective model dynamics. It is well known that option prices contain a volatility premium and the volatility in the subjective process tends to be lower. Therefore, when assuming the risk-neutral and subjective volatility dynamics to be the same, it implicitly assumes similar expected log-likelihoods.

Analysing the obtained risk-aversion estimates from an economic perspective gives some interesting insights into financial markets but also raises some questions. First, the representative investor is almost risk-neutral during periods of distress but clearly risk-averse during periods of low volatility. Second, the variation in risk-aversion is the lowest during crisis periods and increases heavily during times of low volatility. This observation is easily explained from a technical perspective. However, the economic intuition is far more interesting since it implies that the variety of investors during crisis periods is far lower than during less risky times. Third, the confidence intervals show that risk-aversion can become negative imposing a risk-seeking representative investors. This problem is less critical when working with more flexible distribution types but it does not fully vanish. Therefore, this issue cannot be fully ruled out and is a shortcoming of the employed approach. Potential ways of tackling this issue can include the estimation approach, as in Linn et al. (2017), or underlying utility assumptions.

		Log				mix log				GB2			
		RND	hist.	sim. 1	sim. 2	RND	hist.	sim. 1	sim. 2	RND	hist.	sim. 1	sim. 2
log	RND												
	hist.	1.91*											
	sim. 1	1.13	0.18										
	sim. 2	1.12	0.16	0.31									
mix log	RND	2.09**	2.01**	1.91*	1.91*								
	hist.	2.16**	2.09**	2**	2**	2.08**							
	sim. 1	2.23**	2.16**	2.08**	2.08**	1.07	0.18						
	sim. 2	2.23**	2.17**	2.09**	2.09**	1	0.16	0.06					
GB2	RND	11.58***	11.76***	11.26***	11.27***	30.92***	30.71***	28.25***	27.89***				
	hist.	11.44***	11.65***	11.19***	11.19***	30.15***	30.93***	28.86***	28.55***	2.11**			
	sim. 1	11.73***	11.97***	11.59***	11.59***	28.82***	29.91***	30.89***	30.79***	1.05	0.17		
	sim. 2	11.74***	11.99***	11.62***	11.62***	28.58***	29.69***	30.94***	30.88***	0.98	0.15	0.04	

Table 2.6: The table contains the test-statistics of the Amisano and Giacomini (2007) test. The critical values are denoted with (*) at the 10% level, (**) at the 5% level and (***) at the 1% level.

2.6 Conclusion

This paper provides a novel methodology for uncovering the presence of risk aversion in asset prices. Risk aversion is interpreted as the coefficient of relative risk aversion characterizing a power utility function. In contrast to most of the literature, risk aversion is extracted from the risk-neutral distribution underlying the risk-neutral valuation of option prices. More specifically, a parametric expression is exploited relating the risk-neutral distribution function with an objective distribution describing the true probability distribution of asset prices and a power utility function describing individuals' risk attitude. This expression provides a closed-form solution that links risk-neutral expectations with a subjective valuation that incorporates the preferences of a representative agent. This approach allows to estimate consistently the coefficient of relative risk aversion as the sample size increases.

The study also proposed a simulation scheme to introduce uncertainty into the pricing models. By doing so, confidence intervals were constructed that reflect the uncertainty around the parameter estimates of relative risk aversion. It also accommodated the presence of dynamics in risk aversion by proposing a parametric model for describing the dynamics of asset prices. Finally, as an additional contribution, different types of density functions were exploited for describing the probability law of asset prices. By doing so, the predictive ability of the model was explored and assessed, out of sample, the reliability of the different estimates of the relative risk-aversion coefficient for prediction purposes.

The results, obtained from an extensive simulation exercise, suggest that there is more variation in the risk-aversion coefficient than implied by Bliss and Panigirtzoglou (2004). Not surprisingly, the volatility component has a major impact on the coefficient of relative risk aversion. This is due to the direct link between volatility, risk aversion and the risk premium. A way to make the risk-aversion estimate less varying and more stable across time is to assume more flexible density types. Especially, when stochastic volatility is present more flexible density types obtain more stable values. The choice of a parametric model for describing asset prices also has a major effect on the estimates of risk aversion. Thus, it is noted that imposing a lognormal density on asset prices produces greater variation in the coefficient of risk aversion compared to the mixture-lognormal and generalised beta distribution of the second kind. Lastly, a forecasting exercise was performed to assess the relative predictive ability of each estimation method out of sample. For all density types, there is no statistically significant difference between the risk-aversion coefficient based on the Monte-Carlo simulation and the classic approach by Bliss and Panigirtzoglou (2004) and Liu et al. (2007). In contrast, comparing both approaches against the risk-neutral density reveals statistically significant differences.

Within this paper the attention centred on the classic case of a power utility function. Therefore, it can be extended in different directions to obtain further insights into the degree of risk aversion in financial markets and its variation over time. First, an interesting avenue for further research is to evaluate the influence of the volatility premium on risk-aversion estimates. Second, exploring alternative utility functions that provide a different characterization of individuals' risk aversion can be a fruitful strategy to uncover different patterns of risk aversion under different assumptions on individuals' risk profile. Finally, the simulation exercise clearly shows that individuals' risk aversion changes over time. It might be of much interest to applied economists and financial practitioners to explore suitable time series models for forecasting asset prices constructed from reliable estimates of the risk-aversion coefficient.

A Appendix

A.1 Weighting approach by Christoffersen et al. (2009)

Christoffersen et al. (2009) suggests to weight the mean squared error by the Black and Scholes Vega. Vega is defined as the first derivative after volatility:

$$Vega_{BS} = \frac{\partial c(X)}{\partial \sigma_{BS}} = Se^{-qT}\Theta(d) \tag{A.1}$$

with

$$d = \frac{log(\frac{S}{K}) + (r - q + \frac{\sigma_{BS}^2}{2})T}{\sigma_{BS}\sqrt{T}}, \tag{A.2}$$

and Θ the cumulative standard normal distribution. The loss function is defined as:

$$\arg\min_{\theta} G(\theta) = \frac{1}{N}\sum_{i=1}^{N} \frac{(c_{obs}(X_i) - c_{theo}(X_i \mid \theta))^2}{Vega_{BS}^2}, \tag{A.3}$$

with $\theta = (\kappa, \theta, \sigma, \rho, \xi)$.

CHAPTER 3

Strategic Asset Allocation and Multiperiod Investors using Option-Implied State Variables

This chapter studies the impact of option-implied information on an infinitely lived investor that solves an allocation problem between a risky and risk-free asset. The investor determines his optimal weight taking into account a single state variable, which is chosen as option implied risk-premium and market price of risk. The option implied state variables are obtained by making assumptions about the risk-neutral density and subsequent transformation to obtain the subjective probabilities. Multiple densities are employed to analyse the impact of different degrees of flexibility. The results are compared against a single investment into the risky asset, historical approach relying on a GARCH type model, and the log dividend price ratio acting as state variable. Applying this methodology to the S&P 500 acting as risky asset, shows that the use of option implied information is superior towards the historical approach and full investment into the risky asset but it is less clear towards the results of the log dividend price ratio. Further, the option-implied risk-premium is a more suitable state variable compared to the market price of risk. The chosen densities to retrieve the option-implied information provide similar results and are not superior to each other within the chosen set-up. All results are measured out-of-sample taking into account trading costs.

3.1 Introduction

Option prices contain market expectations about the distribution of future asset prices. Breeden and Litzenberger (1978) propose a convenient closed-form expression for obtaining such distribution functions. This pricing equation assumes a risk-neutral distribution function for the future payoffs of the underlying asset. Hence, a simple procedure to calibrate the proposed option pricing formula is by fitting the model option prices into a cross-section of observed option prices. Jackwerth (1999) provides a survey of different risk-neutral density types applied within option pricing. This distribution function does not contain any premia and reflects the risk-neutral market expectations towards the underlying asset. In order to incorporate investors' views on risk aversion, the risk-neutral probabilities need to be transformed into a subjective probability distribution that explicitly considers the preferences of the representative investor. The resulting subjective density, also called "real-world density", "objective density", incorporates the market expectations of the underlying asset price as outlined by Bakshi et al. (2003), Bliss and Panigirtzoglou (2004) and Anagnou-Basioudis et al. (2005). The procedures involving option prices are referred to as "option implied approach".

An alternative approach to retrieve the subjective density describing the random behavior of stock prices is to use the history of asset prices and, in particular, statistical procedures such as the empirical distribution of returns or location-scale time series models for predicting the conditional distribution of asset prices. A natural example of these models is the family of autoregressive models for the conditional mean and ARCH type processes for the conditional volatility, see Engle (1982). Other examples of the historical approach are Jackwerth (2000), Aït-Sahalia and Lo (2000), Pérignon and Villa (2002) and Rosenberg and Engle (2002). The strength of this approach is the possibility of simulating different scenarios to forecast asset prices once the model parameters and the distribution error are estimated. In particular, simulation methods are applied to obtain estimates of the subjective distribution function of asset prices.

Both approaches are compared in the literature from different angles, see for example

Anagnou-Basioudis et al. (2005), Liu et al. (2007) and Shackleton et al. (2010). The density forecasting literature emphasises the comparison between the option implied and historical approach and consistently shows that the option implied approach can forecast more accurately future asset prices compared to the historical approach. The main reason for this is the forward-looking nature of option prices. When markets are efficient, the option price should reflect market expectations of the underlying asset for the time until expiration. This means that any newly arriving information from observation until expiry date are priced into the distribution. A different route is taken by the historical approach. Its estimates rely on past asset price outcomes. Therefore, when forecasting asset prices the historical approach imposes that past dynamics also apply for the future. This means that time series models slowly adjust their dynamics in a forecasting exercise when market conditions change rapidly. In contrast, methods based on implied information obtained from option markets are faster in incorporating new information and dynamics adjust more timely to changing market expectations.

In this chapter, these findings motivate the use of option implied information in an asset allocation perspective. To the best knowledge, the amount of empirical work in this area is still very limited, with seminal work by Kostakis et al. (2011), DeMiguel et al. (2013) and Kempf et al. (2015). To evaluate if option implied information offers statistically and economically significant advantages in an asset allocation context, an investment portfolio is employed comprising two assets (a risky and a risk-free asset). The optimal portfolio allocation is obtained as the result of maximizing the expected utility of an infinitely-lived representative agent.

Following the asset allocation literature, more specifically the seminal works of Campbell and Viceira (1999) and Aït-Sahalia and Brandt (2001), a linear parametric portfolio policy rule is considered for modelling the dynamics of the optimal portfolio weights. In this setting, the individuals' portfolio allocation is dictated by a set of state variables with power to predict the future dynamics of asset returns. In this chapter, the implied risk premium as a novel state variable is proposed. It is compared against the log dividend price ratio proposed by Campbell and Viceira (1999) that has acted as benchmark in the long-term asset allocation literature, and also against the market price of risk, recently proposed by Kostakis et al. (2011) in a similar setting. The proposed state variable takes advantage of the forward-looking information implied from option prices with the risky asset as underlying. In fact, this variable is constructed from the expected value of the underlying risk-neutral distribution function used for pricing options. The comparison between optimal portfolio allocation strategies is based on a range of standard performance measures such as the mean return, Sortino ratio, CAPM and portfolio turnover, see DeMiguel et al. (2013) for a formal definition of these performance measures in an out-of-sample context.

This chapter also contributes to the empirical literature on long-term asset allocation by exploring the sensitivity of the optimal asset allocation exercise to different choices of the risk-neutral distribution function. In contrast to standard parametric portfolio policies, the proposed option implied approach considers the choice of different risk-neutral distributions as an additional aspect to take into account when constructing optimal portfolio weights. The channel through which the risk-neutral distribution function affects the investment portfolio is by the construction of the state variable – the implied risk premium in this case – that is obtained from the entirety of available option prices under the risk-neutral density function. This chapter applies a lognormal density as in Black and Scholes (1973), a mixture lognormal density as in Ritchey (1990), and a non-parametric binomial tree as in Jackwerth and Rubinstein (1996).

The comparison of investment strategies characterized by different portfolio weights obtained under different choices of state variables and risk-neutral density functions is applied to a portfolio given by the US three-month Treasury bill and the S&P 500 Index. All of the results are measured out of sample taking into account trading costs. Rather than using index prices, prices of exchange traded funds (ETF) are taken into account that capture the trading features of the risky asset – ETFs have the advantage that the price derives from a real traded portfolio whereas the alternative index value is just a theoretical construct. The obtained portfolio returns of the option implied and historical approach are evaluated on the basis of the Sharpe ratio, Sortino Ratio, Capital Asset Pricing Model (CAPM) and trading costs.

The empirical study implies three main findings. First, the option implied approach is superior to the historical approach. This finding is greatly supported by empirical forecasting exercises reported in the literature, as for example Liu et al. (2007) or Shackleton et al. (2010). These studies reveal strong timing abilities of the option implied approach relative to a historical benchmark. In this chapter, it can be shown that this finding also applies in a portfolio allocation context even when trading costs are subtracted. Interestingly, the portfolio allocation under the proposed state variable cannot outperform the Campbell and Viceira (1999) optimal portfolio constructed using the log dividend price ratio. As a byproduct, the poor performance of the option-implied approach during the subprime crisis falls behind the benchmark investment strategies including the simple investment into the S&P 500 Index. The results are inconclusive about which approach provides superior results during the subprime crisis. Second, the influence of risk-neutral density function for the option implied approach reveals no significant differences between the lognormal density with a single at-the-money option compared to the alternatives, which exploit the entire cross-section of option prices. Third, the option implied risk premium acting as state variable is more suitable compared to the market price of risk. The performance results of the risk premium are more consistent across periods[1]. Furthermore, restricting the portfolio rules for market timing highlights the timing abilities of the implied risk premium. The timing abilities of the market price of risk tend to have a positive impact as well but the magnitude is lower.

The results confirm the findings in the literature that the use of option implied information improves market timing and yields superior portfolio allocation. Related studies focussed on solving option implied mean-variance portfolios, as for example DeMiguel et al. (2013) and Kempf et al. (2015). They show that the use of option implied information is preferable compared to benchmark approaches such as the historical approach, equally-weighted strategies and full investment into the index. Explicitly, DeMiguel et al. (2013) compares implied variances and historical correlations against historical variances and implied correlations whereas Kempf et al. (2015) considers a fully implied approach. Both papers analyse crisis periods separately. DeMiguel et al. (2013) mentions that the

[1] The results can be confirmed applying the Dow Jones Industrial Average and NASDAQ 100 as risky asset.

subprime crisis period does not effect the overall results. Kempf et al. (2015) confirms this observation and makes an even stronger statement for the period of the dotcom bubble and subprime crisis period by claiming that the crisis period contributes strongly to the overall performance of the portfolio. The here presented findings cannot confirm that the subprime crisis leads to performance gains. Reason for this might be the different structure of the portfolio allocation problem, which focuses on the optimal portfolio decision of an infinitely-lived investor maximizing the expected value of a recursive Epstein-Zin utility function. Another influential study that is strongly related to the proposed method is Kostakis et al. (2011). These authors exploit higher moments in a portfolio allocation between a risky and a risk-free asset. As in Jondeau and Rockinger (2006), Kostakis et al. (2011) expand the expected utility of future wealth via a Taylor series expansion. This procedure allows them to introduce mean, variance, skewness and kurtosis of the option implied density into the portfolio allocation. The implied density extracted from option prices is obtained via splines as in Bliss and Panigirtzoglou (2002). This investment strategy is compared against a benchmark constructed from historical data. The empirical findings of the study illustrate the outperformance of the option implied approach compared to the historical approach. These findings are also confirmed by DeMiguel et al. (2013) and are similar in spirit to the presented empirical findings.

The here proposed study is concerned with a dynamic investment portfolio characterized by a set of portfolio weights driven by a state variable. Kostakis et al. (2011) also applies their methodology to the dynamic portfolio allocation of Wachter (2002) for a finite-lived investor. In their setting, innovations to the state variable and stock returns are perfectly negatively correlated. In this study, the estimated correlation for the market price of risk is negative but not close to -1. Beyond this, the findings line up with Kostakis et al. (2011). Both studies find that the performance of portfolios comprised by a risk-free and a risky asset deteriorates during the subprime crisis. Outside this period, option implied information yields superior results compared to the historical approach and simple investment into the index. However, the comparison of the option implied approach against the method characterized by the log dividend price ratio reveals a mixed performance.

The rest of the chapter is organised as follows. Section 3.2 outlines the procedures to obtain the two state variables, the risk-premium and market price of risk, considering the different methods within the option implied and historical approach and portfolio allocation exercise. Section 3.3 describes the data, evaluation methods and benchmark approaches. Section 3.4 summarises the in-sample estimation results of the portfolio allocation. Section 3.5 presents the out-of-sample results applying the suggested performance measures. Lastly, section 3.6 reviews the main findings of the study.

3.2 Methodology

The selection of suitable state variables is a crucial aspect for long-term asset allocation. These variables have a strong impact on the structure of the optimal weights and the effectiveness of the allocation problem. Historically, many studies (Campbell and Viceira, 1999; Campbell et al., 2003; Chacko and Viceira, 2005) use financial ratios as state variables. However, their structure relies on past information and their implications on the allocation in the future are not always clear. As an alternative, this chapter proposes to derive state variables using options as a source of information. They are a tempting alternative due to their predictive power documented in forecasting exercises (Bliss and Panigirtzoglou, 2004; Liu et al., 2007; Shackleton et al., 2010). The reason for the success of options is generally attributed to the cross-section of option prices containing detailed information about the future state of the underlying asset. Furthermore, estimates are forward-looking and theoretically reflect current market expectations. Three contributions are made with the chosen approach. First, two different option implied state variables are evaluated including the market price of risk, as in Kostakis et al. (2011), and the here proposed option implied risk-premium. Second, the impact of different assumptions about the underlying asset price density is investigated. Different degrees of flexibility are taken into account to find any trades between methods. Third, the state variables are applied to evaluate the out-of-sample performance of the optimal allocation problem of an infinitely-lived investor with time varying expected returns, see Campbell and Viceira (1999).

3.2.1 Option Implied State Variables

To describe future asset price returns two different option implied state variables x_t^I are considered. First, the option implied risk premium. It is defined as the expected return of a risky asset $E_t[r_{t+1}]$ in excess to the risk-free rate $r_{f,t}$. Choosing the risk premium as state variable is an intuitive approach since it directly reflects current premia in the market:

$$x_{RP,t}^I = E_t[r_{t+1}] - r_{f,t}. \tag{3.2.1}$$

As an alternative, the market price of risk x_{MPR}^I is employed, as proposed by Wachter (2002) and applied for option prices in Kostakis et al. (2011). This state variable can be interpreted as an extension of the risk premium since it scales the risk premium by the conditional volatility $\sqrt{Var_t[r_{t+1}]}$. It is formally defined as

$$x_{MPR,t}^I = \frac{E_t[r_{t+1}] - r_{f,t}}{\sqrt{Var_t[r_{t+1}]}}. \tag{3.2.2}$$

The information obtained from option prices enters the state variables via the first two moments: $E_t[r_{t+1}]$ and $Var_t[r_{t+1}]$. Thereby, the necessary moments are retrieved relying on different assumptions about the risk-neutral density and assumptions about the representative investor.

3.2.2 Obtaining Option Implied Moments

To obtain option implied moments, three density types are applied: First, a lognormal density (Log). Due to the assumption of constant volatility across strike prices, the density is estimated using the two closest option prices above and below the current forward price. The implied volatility between both options is interpolated based on the current forward price to match the at-the-money volatility. Second, a mixture-lognormal density (Mix-log). This density allows for more complex shapes. Its estimation takes into account the entire cross-section of option prices to capture the non-normal shape contained in option prices. Third, a binomial tree (Bin-Tree). This method is the most flexible density type due to its non-parametric characteristics. Applying the different density types aims to analyse if the entire cross-section benefits the proposed state variables and if there are any gains when applying more flexible density types.

Since the densities are retrieved from option prices they are priced risk-neutral. Therefore, they need to be transformed to reflect the objective probabilities of the underlying asset. The transformation relies on an assumption about the utility of the representative investor. Throughout the study, it will be assumed that the representative investor follows a power utility. The transformation adds a risk-aversion parameter γ_D to the density, which needs to be estimated. This risk-aversion coefficient should not be mistaken with the risk-aversion parameter embedded in the portfolio allocation of the individual investor that will be denoted hereafter by γ_P. The latter risk-aversion coefficient refers to the relative risk-aversion coefficient of an individual with a recursive Epstein-Zin utility function. Since the risk-aversion of individuals cannot be simply observed, γ_P is freely chosen. Consequently, the values for the relative risk-aversion coefficient in the portfolio allocation problem can be different from the values obtained in the transformation between risk-neutral and objective density functions[2].

3.2.2.1 Mixture of two Lognormal Densities

When there are no arbitrage opportunities Breeden and Litzenberger (1978) show that there is a unique risk-neutral density f_Q for all possible values of the underlying asset price S. The density can be inferred when there are European call prices c available for all strike prices K with the same time to maturity T. The risk-neutral density is then defined as:

$$c_m(K) = e^{-r_f T} \int_K^\infty (S-K) f_Q(S) dS, \qquad (3.2.3)$$

[2]First, a power utility is assumed for retrieving the density of stock returns from the risk-neutral density of option prices. This risk aversion coefficient is considered to characterize the relative risk aversion of a representative investor. Second, the obtained subjective density function is utilised to solve the sophisticated asset allocation problem of an individual investor with recursive preferences modelled by an Epstein-Zin utility function.

with r_f the risk-free rate, T the time to maturity and f_Q is the risk-neutral density. Going back to Ritchey (1990), Melick and Thomas (1997), Brigo and Mercurio (2002) and Liu et al. (2007), the risk-neutral density is assumed to be a mixture of two lognormal distributions. These authors show that the call price is derived from two weighted Black (1976) models such that:

$$c_m(K \mid \theta, r_f, T) = wc_B(F_1, T, K, r_f, \sigma_1) + (1-w)c_B(F_2, T, K, r_f, \sigma_2), \tag{3.2.4}$$

where the density moments are obtained as:

$$E[S_T^n] = wF_1^n exp(0.5(n^2-n)\sigma_1^2 T) + (1-w)F_2^n exp(0.5(n^2-n)\sigma_2^2 T). \tag{3.2.5}$$

The parameters characterising this model are $\theta = [F_1, F_2, \sigma_1, \sigma_2, w]$, with $0 \leq w \leq 1$. Each of the weighted Black (1976) models have its own mean $[F_1, F_2]$ and variance $[\sigma_1, \sigma_2]$. This results in five parameters θ that need to be determined. They are calibrated using observed option prices. The objective is to minimise the squared error between observed c_o and model-based c_m option prices:

$$\min_{\theta} \sum_{i=1}^{N} \left[c_o(K_i) - c_m(K_i \mid \theta)\right]^2 \tag{3.2.6}$$

The optimisation is restricted by the risk-neutrality constraint. It requires that the risk-neutral expectation of the underlying asset price, represented by the current forward price, has to equal the expected value of the risk-neutral density:

$$F = \int_0^{\infty} Sf_Q(S)dS = wF_1 + (1-w)F_2. \tag{3.2.7}$$

The constraint reduces the amount of free parameters θ by 1.

It is well known that the risk-neutral density does not reflect the objective dynamics of the underlying asset. Going back to Aït-Sahalia and Lo (2000), Bliss and Panigirtzoglou (2004) provide a framework to transform the risk-neutral density into its objective counterpart. They link the risk-neutral density to a parametric form of the utility function of the representative investor that allows one to obtain the objective density function.

This density function, f_P, is linked to the risk-neutral density f_Q and utility function u as follows:

$$f_P(S_T) = \frac{\frac{f_Q(S_T)}{u'(S_T)}}{\int_0^\infty \frac{f_Q(y)}{u'(y)} dy} = \frac{S_T^{\gamma_D} f_Q(S_T)}{\int_0^\infty y^{\gamma_D} f_Q(y) dy}. \tag{3.2.8}$$

The integral in the denominator of (3.2.8) ensures the objective density integrates to 1. In the ongoing analysis, it is assumed that the preferences of the representative investor are derived from a power utility function with $u'(S) = S^{-\gamma_D}$, where $u(S) = \frac{S^{1-\gamma_D}-1}{1-\gamma_D}$. The power utility specification adds an additional parameter, the relative risk-aversion coefficient γ_D, to the parameters θ inherited from the risk-neutral density. At this stage, all parameters are known except γ_D. To estimate its value, the approach by Liu et al. (2007) is followed by maximising the log-likelihood of asset price outcomes:

$$\max_{\gamma_D} \sum_{i=1}^{m} log(f_{P,i}(S_{T,i} \mid \hat{\theta}_i, \gamma_D)) \tag{3.2.9}$$

Liu et al. (2007) shows that the transformation (3.2.8) applied to the mixture-lognormal density function under the assumption of a power utility for modelling investors' preferences results in another mixture-lognormal density. The transformed parameter values of the new mixture-lognormal density are as follows:

$$\theta = (F_1, F_2, \sigma_1, \sigma_2, w), \tag{3.2.10}$$

$$F_i^\star = F_i exp(\gamma_D \sigma_i^2 T) \text{ for } i = 1, 2 \text{ and} \tag{3.2.11}$$

$$\frac{1}{w^\star} = 1 + \frac{1-w}{w} \left(\frac{F_2}{F_1}\right)^{\gamma_D} exp(0.5(\gamma_D^2 - \gamma_D)(\sigma_2^2 - \sigma_1^2)T). \tag{3.2.12}$$

The transformation from risk-neutral to objective density only changes the values for the mean and the portfolio weight of the process but leaves the volatility of the process unaffected by the transformation.

Beyond the mixture-lognormal density, the simple case of a lognormal density is considered. It is a special case of the mixture-lognormal density when $w = 1$. In this case, the right hand side of (3.2.4) vanishes and only the simple Black (1976) model remains with $\theta = [F_1, \sigma_1]$. The risk-neutrality constraint reduces to $F = F_1$, which leaves σ_1 as single free parameter. Since the results for the lognormal density are embedded in the mixture case, they are not discussed separately.

3.2.2.2 Binomial Tree

To allow for greater flexibility, a non-parametric density is employed. This choice allows to obtain risk-neutral probabilities without making a priori assumptions on the functional form of the risk-neutral density. It is also better in adapting to possible complex expectations of market participants. The here considered approach goes back to Rubinstein (1994), who proposes the use of a binomial tree. The resulting risk-neutral density is discrete with $n+1$ possible values for the stock price S_j at expiry date T with corresponding probabilities $p = (p_0, p_1, ..., p_n)$. The entirety of values define the risk-neutral density. The density is constrained by $p_0, p_1, ..., p_n \geqslant 0$ and $\sum_{j=0}^{n} p_j = 1$. As in the previous section, the density is priced risk-neutral and the expected value of the density needs to equal the future price F at estimation $\sum_{j=0}^{n} p_j S_j = F$ (risk-neutrality constraint). Under those conditions the option pricing formula in (3.2.3) becomes:

$$c_m(K) = e^{-rT} \sum_{j=0}^{n} \text{max}(S_j - K, 0) p_j. \tag{3.2.13}$$

where the density moments are obtained by:

$$E[S_T^n] = \sum_{j=0}^{n} S_j^n p_j. \tag{3.2.14}$$

The initial values for S_j and p are chosen by the standard binomial tree in Cox et al. (1979) using the obtained implied volatility from the lognormal density in section 3.2.2.1. To estimate the vector p, Jackwerth and Rubinstein (1996) propose a range of different objective functions. However, they recommend to minimise a smoothness function g of the risk-neutral density in combination with the fit of observed and model-based option prices G as in (3.2.6). Here the recommended approach is followed, which requires to minimise:

$$\min_{p} g(p) + mG(p), \text{ where} \tag{3.2.15}$$

$$g(p) = \sum_{j=0}^{n} (p_{j-1} - 2p_j + p_{j+1})^2. \tag{3.2.16}$$

with $p_{-1} = p_{n+1} = 0$. After visually checking the densities, a trade-off parameter of $m = 0.0005$ is chosen. In the application it provides the best trade-off between smoothness and fit of option prices without overfitting the observed option prices. Subsequently, the density is transformed as in (3.2.8). The expression changes to the discrete case:

$$f_P(S_T) = \frac{\frac{p_T}{u'(S_T)}}{\sum_{j=0}^{n} \frac{p_j}{u'(S_j)}} = \frac{S_T^{\gamma_D} p_T}{\sum_{j=0}^{n} S_j^{\gamma_D} p_j}. \tag{3.2.17}$$

where p_T is the probability closest to the final asset price outcome S_T. Performing the transformation results in another binomial tree. There is no closed form solution as it is the case for the (mixture-) lognormal density due to the non-parametric shape.

3.2.2.3 Moment Conversion

The state variables are referring to relative changes in the underlying asset ("returns"). Using (3.2.5) and (3.2.14) the central moments of the asset price densities are obtained. These asset price moments need to be transformed into asset return moments. The transformation is performed from the perspective of the investor solving the portfolio allocation. The current stock price is chosen as as reference point for computing the moments of the returns since the final profit of the investor is the difference between the price today and next rebalancing date. It should also be noted that the stock dividend accrued over the time interval also needs to be added to the expected return as part of the investor's earnings over the holding period. This procedure is not straightforward as the dividend is not incorporated in the expected value of the objective density. The forward price lowers when anticipating dividend payments. Therefore, the expected value needs to be adjusted. The variance of the return density is obtained applying the same procedure but related this time to the squared stock price. This leads to the following expressions for the conditional expected return and variance:

$$E_t[r_{t+1}] = log\left(\frac{E_t[S_{t+1}]}{S_t}\right) + q_t \tag{3.2.18}$$

$$Var_t[r_{t+1}] = \frac{Var_t[S_{t+1}]}{S_t^2} \tag{3.2.19}$$

3.2.3 Portfolio allocation

The proposed state variables are evaluated within the framework of Campbell and Viceira (1999). These authors assume an infinitely-lived investor with Epstein and Zin (1989) and Weil (1989) preferences who maximises the expected utility of future wealth. Thereby, the investor seeks an optimal allocation between a single risk-free asset with constant log return r_f, a single risky asset with log return r_{t+1} and innovations u_{t+1} in the risky asset's return:

$$r_{t+1} - E_t[r_{t+1}] = u_{t+1}. \tag{3.2.20}$$

The innovations are normally distributed with zero mean and variance σ_u^2. The expected excess log return on the risky asset is state dependent and is determined by a single state variable x_t:

$$E_t[r_{t+1}] - r_f = x_t \tag{3.2.21}$$

The state variable x_t or risk-premium follows a mean-reverting AR(1) process with mean μ and persistence ϕ. The innovations η_{t+1} are conditionally homoskedastic and normally distributed with zero mean and variance σ_η^2:

$$x_{t+1} = \mu + \phi(x_t - \mu) + \eta_{t+1} \tag{3.2.22}$$

The two innovations η_{t+1} and u_{t+1} are correlated with each other by $\sigma_{\eta u}$. The choice to model the risk-premium process with an AR(1) is rooted in the characteristics of the underlying data. The empirical work in section 3.4 and 3.5 relies on monthly observations since option prices cannot be observed for any expiry date. Therefore, it is practical to assume that the underlying stochastic process is in discrete time. The lag of one for the autoregressive process has been confirmed taking into account appropriate testing procedures. An alternative to the assumed process would be to work in continuous time applying an Ornstein-Uhlenbeck process in the style of Chacko and Viceira (2005). It allows to model the risk-premium process continuously as it is the case in financial markets. However, it would complicate the here intended analysis without a clear benefit for the proposed research agenda. Therefore, this path is not followed.

Based on these assumptions, Campbell and Viceira (1999) guess a form of the optimal consumption and portfolio policies. They suggest that the state variable x_t is linear in the optimal portfolio weight and quadratic in the log-consumption-wealth ratio:

$$\alpha_t = a_0 + a_1 x_t \tag{3.2.23}$$

$$c_t - w_t = b_0 + b_1 x_t + b_2 x_t^2 \tag{3.2.24}$$

This results in unknown parameters $[a_0, a_1, b_0, b_1, b_2]$. The parameters defining the linear portfolio policy are determined by the following expression:

$$a_0 = \frac{1}{2\gamma_P} - \frac{b_1}{1-\Psi}\frac{\gamma_P - 1}{\gamma_P}\frac{\sigma_{\eta u}}{\sigma_u^2} - \frac{b_2}{1-\Psi}\frac{\gamma_P - 1}{\gamma_P}\frac{\sigma_{\eta u}}{\sigma_u^2}2\mu(1-\phi) \tag{3.2.25}$$

$$a_1 = \frac{1}{\gamma_P \sigma_u^2} - \frac{b_2}{1-\Psi}\frac{\gamma_P - 1}{\gamma_P}\frac{\sigma_{\eta u}}{\sigma_u^2}2\phi, \tag{3.2.26}$$

where γ_P and Ψ are the risk-aversion and elasticity intertemporal substitution from the Epstein-Zin utility. The remaining parameters of the log consumption wealth ratio are retrieved by a recursive non-linear system. The exact definition of the terms and the recursive procedure to obtain $[b_0, b_1, b_2]$ is outlined in Campbell and Viceira (1999) and would exceed the scope of this chapter.

The expression for the optimal portfolio weight α_t consists of two components to capture the asset demand. The first term defines the myopic asset demand. The myopic component is proportional to the risk-premium and inversly proportional to the volatility and risk-aversion. The second component captures the intertemporal hedging demand going back to Merton (1969), Merton (1971) and Merton (1973). The covariance term between the two innovations has a crucial impact on the intertemporal hedging demand. The higher the covariance the better the ability of the risky asset to hedge against changing investment opportunities over time. The weights have two special cases: 1. When returns are unpredictable the hedging demand is zero and $\sigma_{\eta u} = 0$; 2. When $\gamma_p = 1$ the portfolio weight reduces to the myopic component.

To solve the allocation problem the parameters $[\mu, \phi, \sigma_u^2, \sigma_\eta^2, \sigma_{\eta u}]$ are of central interest. They are inferred from the suggested option implied state variables. As in Campbell and Viceira (1999), they are estimated by restricted VAR(1) via OLS with the state variable as predictor:

$$\begin{pmatrix} r_{t+1} - r_{f,t+1} \\ x_{t+1}^I \end{pmatrix} = \begin{pmatrix} \theta_0 \\ \beta_0 \end{pmatrix} + \begin{pmatrix} \theta_1 \\ \beta_1 \end{pmatrix} x_t^I + \begin{pmatrix} \varepsilon_{1,t+1} \\ \varepsilon_{2,t+1} \end{pmatrix} \tag{3.2.27}$$

where $(\varepsilon_{1,t+1}, \varepsilon_{2,t+1}) \backsim N(0, \Omega)$ and

$$\Omega = \begin{bmatrix} \Omega_{11} & \Omega_{12} \\ \Omega_{21} & \Omega_{22} \end{bmatrix}.$$

Subsequently, the required parameters are retrieved as following: $\mu = \theta_0 + \theta_1\beta_0/(1-\beta_1)$, $\phi = \beta_1$, $\sigma_\eta^2 = \theta_1^2\Omega_{22}$, $\sigma_u^2 = \Omega_{11}$ and $\sigma_{\eta u} = \theta_1\Omega_{12}$.

In a last step, the parameters are normalised defining the optimal portfolio and consumption rules as in Campbell and Viceira (1999). After the normalisation, the intercepts of the optimal policy functions are the optimal allocation to stocks and the optimal consumption-wealth ratio when the expected excess return is zero. In this case the risky asset has a risk-premium of zero and a myopic investor would not allocate any wealth to it. Therefore, any asset demand would be generated by the intertemporal hedging demand. The normalisation is achieved by setting $a_0^* = a_0 - a_1(\sigma_u^2/2)$, $b_0^* = b_0 - b_1(\sigma_u^2/2) + b_2(\sigma_u^4/4)$ and $b_1^* = b_1 - b_2\sigma_u^2$. The parameters a_1 and b_2 are not affected by this transformation.

3.3 Evaluation Criteria

3.3.1 Data

The empirical application is based on the S&P 500 Index (SPX). The observation period ranges from 20th January 1996 until the 15th April 2016. In the empirical application the time series is split into the subprime crisis period from 18th August 2007 until 19th December 2009 and the period after the crisis from 19th December until 15th April 2016. Within these periods the option expiry dates schedule the observation dates and allocation periods. Relevant option expiry dates are the third Friday each month. In order to avoid autocorrelation problems, as discussed in Bliss and Panigirtzoglou (2004), estimation lengths are fixed at 4 weeks prior to expiry. This ensures that periods are non-overlapping. Following the procedures outlined in section 3.2.2, the risk-neutral density is estimated using the closing price on the option observation date. This results in 243 periods.

	Mean	Median	Standard deviation	Minimum	Maximum
Amount	58.94	44	34.27	16	185

	Deep OTM Call	OTM Call	ATM Options	OTM Put	Deep OTM Put
F/K	< 0.90	$0.90 - 0.97$	$0.97 - 1.03$	$1.03 - 1.10$	$1.10 <$
Amount	3.23	15.75	22.9	24.21	33.91

Table 3.1: Summary statistics of at-the-money and out-of-the-money options fulfilling the conditions outlined in this section. The option statistics are displayed by moneyness in % based on each observation date. The amount of options vary quite strongly due to the increasing amount of available option prices over the past years.

In the empirical application only European options are considered that are always written on the underlying spot market index. The structure of European options allows to assume that they are written on the corresponding forward, as discussed in Liu et al. (2007). Options as well as future prices are obtained end of day from the Chicago Mercentile Exchange. The obtained option prices are checked against no arbitrage constraints. The options passing the no arbitrage constraints are then selected based on their moneyness. Moneyness is defined as the ratio between current forward and strike price. To ensure prices are accurate only options that are at-the-money and out-of-the-money are taken into account. For in-the-money options it can be observed that they are less frequently traded and therefore prices might not reflect market expectations. To qualify as eligible option the moneyness for a call-option (put-option) needs to exceed (be below) 0.97 (1.03). Lastly, the option price itself is checked. If there is no bid price quoted the option is dropped from the sample since it cannot be actively traded any more. Furthermore, any options with a price of less than $3/8$ are excluded. Too low prices might not reflect the true value of the option due to the proximity of tick sizes. The eligible option prices are summarised in table 3.1. If necessary, option prices are transformed via the put-call parity relationship. In case there exists a call- and put-option for the same strike the average call price is used.

Corresponding future prices only exist for quarterly expiry dates, namely March, June, September and December for every year. Since options mature monthly, the closest future price is interpolated to match the monthly expiry date. This is performed using the following definition of the future price $F_t = S_t e^{(r_{f,t} - q_t)T}$. The risk-free rate r_f is retrieved via bootstrapping from zero-coupon bonds. The dividend yield q_t is calculated using the total dividend on the underlying index over the previous year divided by the current stock price. Taking the logs of q_t provides the log dividend price ratio $p_t - q_t$ in the style of Campbell and Viceira (1999). Lastly, ETF bid and ask prices for the the SPDR S&P 500 ETF are obtained to evaluate the portfolio allocation. The ETF was selected based on the fund inception date and size.

3.3.2 Performance Measures

The performance of the proposed state variables is evaluated based on monthly out-of-sample returns $r_{P,t}$ of the portfolio allocation. To estimate the parameter values for the portfolio allocation in section 3.2.3 and risk-aversion γ_D in section 3.2.2, a rolling and recursive window estimation is applied. The transformation using a rolling window takes into account monthly periods of 36, 48, 60, 72, 84, 96, 108 and 120. The empirical application relies on the transformation using 120 periods in order to accurately reflect the long-term dynamics within the portfolio allocation. Shorter interval tend to lack reasonable parameter estimates.

The performance of the risky asset is based on Exchange Traded Funds (ETF). Using an ETF to evaluate the performance of the portfolio allocation comes with multiple advantages compared to using index value. Most striking is that it replicates the underlying asset in an actually traded portfolio with no fractions of shares. Furthermore, paid dividends and its taxation are directly reflected in the price. Trading costs and management fees are also incorporated into the ETF's price. Beyond the trading costs, bid and ask prices are used to implement the suggested portfolio allocation.

The out-of-sample returns of the suggested portfolio allocation are calculated on the basis of a theoretical fund with a value of one. On the basis of the optimal weights, shares are bought using the ask price and sold using the bid price. Depending on whether the portfolio is long (short) in the risky asset, the funds' value is determined by multiplying the shares with the bid (ask) price. The remaining share in the risk-free asset is interest-bearing using the corresponding risk-free rate. In case the portfolio weight in the risky asset exceeds one, interest is charged to finance the leverage. The portfolio weight in every period is adjusted for the difference in weight between $t-1$ and t. Furthermore, the portfolio weight is restricted in the interval of $[-1, 2]$ to avoid unrealistic leverage.

The resulting out-of-sample returns are first standardised to 4-week horizons before any performance measures are calculated. This is necessary since option expiry dates are not necessarily connected to each other. If the estimation horizon is not connected, it is assumed that the investor holds the portfolio until the next rebalancing date. Based on the standardised out-of-sample returns, the performance of the proposed state variables is evaluated using four different criteria: I) Sharpe ratio; II) Sortino ratio; III) CAPM; and IV) portfolio turnover. They are calculated over a series of N realisations. The annualised Sharpe ratio (SR) using the mean return μ and variance σ^2 is determined as following:

$$\hat{\mu} = \frac{1}{N}\sum_{t=1}^{N} r_{P,t} - r_{f,t} \tag{3.3.1}$$

$$\hat{\sigma}^2 = \frac{1}{N-1}\sum_{t=1}^{N}(r_{P,t} - \mu)^2 \tag{3.3.2}$$

$$\widehat{SR} = \frac{365}{28}\frac{\hat{\mu}}{\sqrt{\hat{\sigma}^2}} \tag{3.3.3}$$

The obtained out-of-sample Sharpe ratios are futher tested. The Sharpe ratio are compared by the proposed test of Ledoit and Wolf (2008). They provide a studentized time series bootstrap test for equality of Sharpe ratios. For the analysis, $1,000$ bootstrap resamples and an expected block size equal to 5 is applied as in DeMiguel et al. (2014). All approaches are tested against the simple investment into the underlying index.

To gain further insight into the relative performance of the investment strategies, the Sortino ratio is computed. This ratio extends the Sharpe ratio by only considering the downside volatility σ^2_{adj} and therefore puts emphasis on returns falling below the risk-free rate:

$$\hat{\sigma}^2_{adj} = \frac{1}{N-1}\sum_{t=1}^{N}(\min(0, r_{P,t} - r_{f,t}))^2 \tag{3.3.4}$$

$$\widehat{SoR} = \frac{365}{28}\frac{\hat{\mu}}{\sqrt{\hat{\sigma}^2_{adj}}} \tag{3.3.5}$$

The CAPM is evaluated assuming that the returns of the underlying risky asset are the market returns r_m. The values for $alpha$ are annualised by multiplying them by $\frac{365}{28}$:

$$r_{P,t} - r_{f,t} = \widehat{alpha} + \widehat{beta}(r_{m,t} - r_{f,t}) \tag{3.3.6}$$

Furthermore, the obtained values for $alpha$ are tested applying a simple t-test to check whether the constant term $alpha$ is significantly different from zero.

To analyse to what extent the portfolio weights need to be adjusted when moving from t to $t+1$ the average portfolio turnover per holding period is calculated. It is the average absolute difference in the risky asset's weight α between two consecutive rebalancing dates:

$$Turnover = \frac{1}{N-1}\sum_{t=1}^{N-1}|\alpha_{t+1} - \alpha_t| \tag{3.3.7}$$

Lastly, the portfolio allocation itself is evaluated to gain insight on the effectiveness of the myopic and hedging demand component. This is performed by restricting the portfolio rule while adjusting optimally for consumption. This procedure is outlined in detail in Campbell and Viceira (1999). Three different restricted portfolio rules are considered: First, restriction of the timing component. This is achieved by setting the state variable x_t^I in (3.2.23) to the unconditional expected log excess return μ. Second, restriction of the hedging demand. Ignoring the hedging demand leads to a myopic portfolio rule. This is accomplished by setting the covariance $\sigma_{\eta u}$ to zero. Third, both components are restricted with $x_t^I = \mu$ and $\sigma_{\eta u} = 0$.

3.3.3 Benchmark Approaches

The suggested option implied state variables are benchmarked against three general methods. I) full investment into the risky asset; II) the log dividend price ratio ($d_t - p_t$) as in Campbell and Viceira (1999); and III) the historical approach (hist).

The first benchmark approach allocates all available funds into the risky asset and does not make any adjustments over time. This simple strategy is chosen since it is far more easy to implement compared to the here proposed methodology. The second benchmark refers to financial ratios. They are commonly applied as state variables in a portfolio allocation context. However, they are rarely compared against the option implied approach in the density forecasting literature. Therefore, the option implied state variables are benchmarked against the log dividend price ratio as initially proposed by Campbell and Viceira (1999). The log dividend price ratio enters the portfolio allocation directly as a predictor in (3.2.27). Lastly, a historical approach is applied as benchmark. It is a classic choice from the density forecasting literature and relies on a series of asset price returns. Making assumptions about the return and volatility dynamics allows to obtain conditional moments. Thereby, the historical information enters the portfolio allocation in the same way as for the option implied approach. The conditional mean and variance is retrieved from the previously specified processes, which then enter the allocation problem via the two proposed state variables in (3.2.1) and (3.2.2). The reported values in the empirical application rely on a rolling estimation window of ten years, which is matched with the option observation dates. The return process is assumed to follow an autoregressive of order one - AR(1) process. For the conditional volatility the GJR asymmetric GARCH process is proposed going back to Glosten et al. (1993) with t-distributed white noise u_{t+1}. This model reacts differently to positive and negative shocks to the return process and extends very naturally standard GARCH models. The proposed location-scale model for the historical returns is:

$$r_{t+1} = \mu + \phi r_t + u_{t+1}\sqrt{h_{t+1}}. \tag{3.3.8}$$

$$h_{t+1} = w + (\alpha + \beta I_t)\varepsilon_t^2 + \delta h_t \tag{3.3.9}$$

$$I_t = \begin{cases} 0, & \text{if } r_t \geq 0 \\ 1, & \text{if } r_t < 0 \end{cases} \tag{3.3.10}$$

where $w, \alpha, \beta, \delta > 0$ and $\alpha + \frac{\beta}{2} + \delta < 1$. It is important to mention that the dividend yield needs to be added to the conditional mean since it reduces the index value. Further, alternative specifications of the historical approach are employed including normally distributed white noise process, constant return process with $\phi = 0$ and a symmetric specification of the volatility process. However, the impact on the results is of minor magnitude, which is why only the outlined model above is reported.

3.4 Estimation

This section presents and discusses the estimation results of the risk-neutral parameters, the transformation into the subjective density and the estimation results of the historical approach. Following this, the behaviour of parameter values of portfolio allocation is explored across different indices, state variables and methods. Results presented in this section take into account the entire sample for retrieving the necessary parameter values of the portfolio allocation.

3.4.1 State variables

Table 3.2 presents the key summary statistics of the estimated values. To ensure that the squared error converges to its minimum, multiple initial values have been applied into the non-linear optimisation problem. Note that the squared errors of the lognormal density are not displayed since the free parameter σ is obtained directly from the interpolation of the implied volatility of the two closest at-the-money options. This guarantees that the density precisely fits the data.

Method	G $\times$ 100			g $\times$ 1000		
	Mean	Median	Standard deviation	Mean	Median	Standard deviation
Mix-log	16.698	11.646	19.193			
Bin-Tree	2.601	1.305	4.419	0.131	0.077	0.167

Table 3.2: Risk-neutral density estimation errors

Once the risk-neutral parameter values are estimated, the relative risk-aversion parameter γ_D of the representative investor can be derived. Unlike the estimation of the risk-neutral density, the relative risk-aversion is estimated using all observations within the rolling window. The relative risk-aversion is updated on each new observation date before the state variable at time point t is determined. The statistics are reported in table 3.3 and refer to the 10-year rolling window estimation. Applying the lognormal density obtains average estimates for the relative risk-aversion of 1.163, for the mixture-lognormal density 1.115, while for the binomial tree of 1.053, which are consistent across methods.

Method	Mean	Median	Standard deviation	5th Perc.	95th Perc.	Min	Max
Log	1.163	1.276	0.959	-0.468	2.262	-0.823	2.388
Mix-log	1.115	1.161	0.861	-0.307	2.131	-0.586	2.258
Bin-Tree	1.053	1.102	0.772	-0.223	1.951	-0.458	2.052

Table 3.3: Statistics of risk-aversion estimates for the 10-year rolling window.

Further, it can be observed that the estimates of the SPX are consistent with results from other related studies. Specifically, Shackleton et al. (2010) report an average MSE for the risk-neutral density of 0.11 with median 0.03 and standard deviation of 0.21. Bliss and Panigirtzoglou (2004) obtain a relative risk-aversion of 4.08 for the period from 1992 until 2001, whereas Liu et al. (2007) retrieve a value of 1.85 for the period from 1993 until 2003 for the FTSE 100. The discrepancies in these statistics can be attributed to the different time periods covered in each study. In Bliss and Panigirtzoglou (2004) the authors exclude most of the bearish market period starting in 2000, which can explain why their risk-aversion estimates are found to be so high. Similarly, Liu et al. (2007) utilize a sample, which covers most of the bearish period resulting in a lower risk-aversion. Even though the obtained average risk-aversion coefficients fall below these reported values, they are reasonable as the obtained estimates cover in large parts both the dotcom bubble and the financial subprime crisis. Losses in the underlying indices have a strong impact on the risk-aversion estimates, which results in average values below the values reported in the existing literature.

Parameter	Mean	Median	Standard deviation
μ	0.043	0.047	0.015
ϕ	-0.017	-0.007	0.039
ω	0.011	0.010	0.004
δ	0.927	0.929	0.017
α	0.005	0.000	0.010
β	0.115	0.124	0.045
d.o.f.	8.504	7.889	2.580

Table 3.4: Parameter statistics for the historical approach.

3.4.2 State variable Dynamics

Table 3.5 reports the estimation results for the restricted VAR(1) and derived model for the SPX. The obtained results utilize the entire sample rather than using a rolling windows to ensure that results can be displayed in a coherent manner.

The R^2 of the first equation in the restricted VAR(1), which models the excess return in the next period, indicates similar results (up to two decimal places) between the option implied and the historical approach. This may come as a surprise since the density forecasting literature emphasises the strength of option implied forecasts compared to the historical approach. Therefore, one may expect a distinct higher R^2 compared to the historical approach. One characteristic that may affect these values may be the different levels of persistence as will be discussed in the next paragraph. Hence, comparing the models merely on values of R^2 may not help to obtain a clear picture. Another possible attribute is that the strength of the option implied approach lays in the asymmetric shape of the underlying probability distribution. The asymmetry of the underlying distribution is not incorporated in the state variable, which might negatively effect the predictability.

		Restricted VAR(1)				Derived Model			
x_t^I	Method	$\begin{pmatrix}\theta_0\\\beta_0\end{pmatrix}$	$\begin{pmatrix}\theta_1\\\beta_1\end{pmatrix}$	$\begin{bmatrix}\Omega_{11} & \Omega_{12}\\\Omega_{21} & \Omega_{22}\end{bmatrix}$	R^2	μ	ϕ	$\begin{bmatrix}\sigma_u^2 & \sigma_{u\eta}\\\sigma_{u\eta} & \sigma_\eta^2\end{bmatrix}$	ρ
RP	Log	$\begin{pmatrix}0.000\\0.002\end{pmatrix}$	$\begin{pmatrix}0.487\\0.679\end{pmatrix}$	$\begin{bmatrix}2.455E-03 & -1.719E-05\\-1.719E-05 & 2.462E-05\end{bmatrix}$	$\begin{pmatrix}0.009\\0.679\end{pmatrix}$	0.002	0.679	$\begin{bmatrix}2.455E-03 & -8.378E-05\\-8.378E-05 & 5.844E-06\end{bmatrix}$	-0.699
	Mix	$\begin{pmatrix}-0.001\\0.002\end{pmatrix}$	$\begin{pmatrix}0.702\\0.691\end{pmatrix}$	$\begin{bmatrix}2.450E-03 & -1.430E-04\\-1.430E-04 & 1.646E-05\end{bmatrix}$	$\begin{pmatrix}0.011\\0.477\end{pmatrix}$	0.002	0.691	$\begin{bmatrix}2.450E-03 & -1.004E-04\\-1.004E-04 & 8.107E-06\end{bmatrix}$	-0.712
	Bin-tree	$\begin{pmatrix}0.000\\0.002\end{pmatrix}$	$\begin{pmatrix}0.403\\0.652\end{pmatrix}$	$\begin{bmatrix}2.458E-03 & -1.780E-04\\-1.780E-04 & 2.786E-05\end{bmatrix}$	$\begin{pmatrix}0.007\\0.424\end{pmatrix}$	0.002	0.652	$\begin{bmatrix}2.458E-03 & -7.196E-05\\-7.196E-05 & 4.520E-06\end{bmatrix}$	-0.680
	Hist	$\begin{pmatrix}-0.005\\0.000\end{pmatrix}$	$\begin{pmatrix}1.069\\0.972\end{pmatrix}$	$\begin{bmatrix}2.454E-03 & 5.856E-06\\5.856E-06 & 6.336E-07\end{bmatrix}$	$\begin{pmatrix}0.009\\0.940\end{pmatrix}$	0.003	0.972	$\begin{bmatrix}2.454E-03 & 6.259E-06\\6.259E-06 & 7.238E-07\end{bmatrix}$	0.149
MPR	Log	$\begin{pmatrix}0.010\\0.007\end{pmatrix}$	$\begin{pmatrix}-0.193\\0.818\end{pmatrix}$	$\begin{bmatrix}2.459E-03 & 2.318E-04\\2.318E-04 & 5.833E-05\end{bmatrix}$	$\begin{pmatrix}0.007\\0.668\end{pmatrix}$	0.002	0.818	$\begin{bmatrix}2.459E-03 & -4.478E-05\\-4.478E-05 & 2.176E-06\end{bmatrix}$	-0.612
	Mix	$\begin{pmatrix}0.011\\0.007\end{pmatrix}$	$\begin{pmatrix}-0.223\\0.823\end{pmatrix}$	$\begin{bmatrix}2.459E-03 & 2.086E-04\\2.086E-04 & 4.648E-05\end{bmatrix}$	$\begin{pmatrix}0.007\\0.676\end{pmatrix}$	0.002	0.823	$\begin{bmatrix}2.459E-03 & -4.653E-05\\-4.653E-05 & 2.313E-06\end{bmatrix}$	-0.617
	Bin-tree	$\begin{pmatrix}0.011\\0.006\end{pmatrix}$	$\begin{pmatrix}-0.241\\0.827\end{pmatrix}$	$\begin{bmatrix}2.457E-03 & 2.237E-04\\2.237E-04 & 4.822E-05\end{bmatrix}$	$\begin{pmatrix}0.008\\0.682\end{pmatrix}$	0.002	0.809	$\begin{bmatrix}2.457E-03 & -5.386E-05\\-5.386E-05 & 2.795E-06\end{bmatrix}$	-0.650
	Hist	$\begin{pmatrix}0.008\\-0.023\end{pmatrix}$	$\begin{pmatrix}0.076\\0.688\end{pmatrix}$	$\begin{bmatrix}2.462E-03 & 1.783E-03\\1.783E-03 & 3.300E-03\end{bmatrix}$	$\begin{pmatrix}0.006\\0.476\end{pmatrix}$	0.002	0.688	$\begin{bmatrix}2.462E-03 & -4.148E-05\\-4.148E-05 & 1.786E-06\end{bmatrix}$	-0.625
$d_t - p_t$		$\begin{pmatrix}0.121\\-0.100\end{pmatrix}$	$\begin{pmatrix}0.030\\0.975\end{pmatrix}$	$\begin{bmatrix}2.419E-03 & -2.419E-03\\-2.419E-03 & 2.721E-03\end{bmatrix}$	$\begin{pmatrix}0.023\\0.950\end{pmatrix}$	0.002	0.975	$\begin{bmatrix}2.419E-03 & -7.141E-05\\-7.141E-05 & 2.371E-06\end{bmatrix}$	-0.943

Table 3.5: Estimation results for the entire sample period for the restricted VAR and resulting dynamics of the risk-premium process.

Analysing the state variables unveils that there are state variable specific characteristics that become visible. The option implied approach attains higher values of the R^2 using the risk-premium as state variable compared to the market price of risk. In line with this observation is the historical approach. Nevertheless, compared to the log dividend price ratio both state variables fall behind independently of the underlying data type.

Comparing the dynamics of the two equations within the VAR(1) suggests that the bottom equation, describing the dynamics of the state variable, obtains far lower values of the R^2. This does not come as a surprise since the first equation aims to predict the risk-premium in the next period, which is known to be a difficult task.

Apart from the R^2, the persistence of ϕ in the derived AR(1) model provides useful insights about the dynamics of the applied state variables. The option implied approach using the risk-premium as state variable obtains overall the lowest values of persistence. This can be attributed to the sole dependence on the observed option prices at observation date, which make the option implied risk-premium more reactive.

Compared to the risk-premium, the market price of risk obtains higher values of persistency. The higher level of persistence is most likely a consequence of the relation between the risk-premium and volatility. High volatility normally implies a high risk-premium and vice versa. Therefore, the ratio between both factors is relative stable compared to the single use of the risk-premium.

The historical approach diverges from these observations. The risk-premium is highly persistent, which can be traced back to the estimation method using a historical time series of asset price returns. The model simply extrapolates the past returns into the future resulting in a low variation of the risk-premium. Opposite to the risk-premium is the market price of risk. Here the persistence is lower and discrepancies in the results may be attributed to the different model specifications. In the option implied approach the risk-premium is strongly connected to the variance priced in the options. For the historical approach this is not directly the case. In the here applied model there is no direct link between the risk-premium and volatility. Both factors are less related to each other, which promotes to reduce persistence.

Lastly, the estimation results of the log dividend price ratio are in line with the reported values in Campbell and Viceira (2000) Erratum Table 1. The R^2 is the highest compared to the alternative methods and state variables. Furthermore, the log dividend price ratio has a high persistence in line with empirical observations and the expected log excess return μ lines up with the alternative methods. Overall, values are similar to the initial paper of Campbell and Viceira (1999).

3.4.3 Portfolio weight characteristics

As briefly discussed in section 3.2.3 there are two components of the portfolio weight, namely myopic and hedging demand. While myopic demand derives entirely from expected returns and variance σ_u^2, hedging demand component values depend on the covariance $\sigma_{\eta u}$ between the innovations in the state variable η_{t+1} and innovations in the risky asset's return u_{t+1}. If $\sigma_{\eta u}$ equals zero the demand for the risky asset is purely derived from the myopic component. Therefore, to effectively hedge against changes in the investment opportunity set, $\sigma_{\eta u}$ should be of high magnitude.

In addition, the sign of $\sigma_{\eta u}$ also plays a key role how the hedging demand contributes to the portfolio weight. In cases where $\sigma_{\eta u}$ is negative (positive) the hedging demand increases (decreases) the weight in the risky asset. As discussed in Campbell and Viceira (1999), the case where $\sigma_{\eta u} < 0$ is of empirical relevance. Only when the latter holds, the long term investor maintains a positive proportion in the risky asset when expected returns are zero. This characteristic holds true in the performed analysis except for the historical approach using the risk-premium as state variable. This might be a result of the slow adjustment of the historical approach. Since the historical returns are simply extrapolated into the future, this method lags anticipating changes in the market. This is ultimately reflected in a low magnitude and positive covariance $\sigma_{\eta u}$.

However, to further analyse how the covariance $\sigma_{\eta u}$ impacts the portfolio allocation the average portfolio weight is studied by setting the state variable x_t to its long run log excess return $\mu + \frac{\sigma_u^2}{2}$. Table 3.6 displays the average portfolio weights and contribution of the hedging demand. In the here considered application the covariance $\sigma_{\eta u}$ is mostly negative, increasing the average share in the risky asset.

Comparing the state variables with each other shows that the risk-premium and market price of risk obtain similar results. Even though the share of hedging demand varies, the resulting average weight is quite stable. However, comparing the risk-premium and market price of risk with the log dividend price ratio shows some tremendous differences. Overall the average portfolio weight exceeds the alternative methods by more than 30%. Furthermore, the share of hedging demand on the overall portfolio weight is far higher. To explain this finding the correlation ρ between innovations in the state variable and returns provides good insight. It gives an indication how effectively the state variable can be used to hedge against changes in investment opportunities. Clearly, the log dividend price ratio is ahead of the alternative approaches displaying the highest values of correlation.

The risk-aversion parameter γ_P impacts the average weight in the risky asset in the same way for all methods. The higher the risk-aversion the lower the average weight in the risky asset. This finding is not a surprise since a more risk-averse investor dislikes uncertainty and therefore allocates a lower fraction into the risky asset. The influence of the elasticity of inter-temporal substitution ψ is negligible for the here proposed state variables. The weight does not change in any meaningful way. Therefore, the empirical findings in section 3.5 only report the case when $\psi = 0.5$. In contrast, the log dividend price ratio behaves as outlined in Campbell and Viceira (1999). An increase in ψ increases the portfolio weight in the risky asset. However, the impact is considerably smaller compared to the risk-aversion γ_P.

These characteristics are reflected in the portfolio weights over time. Appendix B.2 displays the variation of the hedging demand and myopic component. For the option implied and historical approach the majority of variation in the weight of the risky asset is a consequence of the myopic demand. Thereby, the use of the risk-premium leads to stronger spikes in the risky portfolio weight compared to the market price of risk. The hedging demand is of small magnitude. Comparing the different approaches with each other unveils that the highest variation in the hedging demand component is achieved

x_t^I	Method	γ_P	Mean allocation (in %): E.I.S			Hedging demand (in %): E.I.S		
			0.25	0.5	0.75	0.25	0.5	0.75
RP	Log	2	74.51	74.50	74.49	5.35	5.34	5.33
		4	38.38	38.37	38.36	8.12	8.09	8.07
		6	25.85	25.84	25.83	9.05	9.03	9.00
	Mix-log	2	75.50	75.49	75.48	6.67	6.66	6.65
		4	39.21	39.20	39.19	10.16	10.13	10.10
		6	26.49	26.48	26.47	11.34	11.31	11.27
	Bin-Tree	2	73.58	73.58	73.57	4.20	4.19	4.18
		4	37.64	37.63	37.63	6.36	6.34	6.32
		6	25.29	25.28	25.28	7.08	7.06	7.04
	Hist	2	75.51	75.55	75.60	-3.75	-3.68	-3.62
		4	37.11	37.16	37.21	-5.55	-5.41	-5.27
		6	24.60	24.64	24.68	-6.13	-5.96	-5.80
MPR	Log	2	72.75	72.75	72.74	3.89	3.88	3.87
		4	37.15	37.14	37.14	5.88	5.86	5.85
		6	24.94	24.94	24.93	6.55	6.53	6.51
	Mix-log	2	72.74	72.74	72.73	3.95	3.94	3.94
		4	37.15	37.15	37.14	5.97	5.95	5.94
		6	24.95	24.94	24.94	6.65	6.63	6.62
	Bin-Tree	2	72.22	72.22	72.22	3.44	3.44	3.43
		4	36.78	36.78	36.77	5.20	5.19	5.18
		6	24.67	24.67	24.67	5.79	5.78	5.76
	Hist	2	70.24	70.24	70.25	-0.77	-0.76	-0.76
		4	34.98	34.99	34.99	-1.16	-1.15	-1.14
		6	23.29	23.30	23.30	-1.28	-1.27	-1.26
Classic	Log(DivY)	2	133.32	135.69	138.42	44.26	45.23	46.31
		4	116.16	118.38	120.85	68.01	68.61	69.25
		6	104.16	105.45	106.80	76.22	76.51	76.81

Table 3.6: The left column displays the mean optimal percentage allocation to the risky asset. The right column panel displays the mean hedging demand relative to the mean total demand. All values are based processes reported in table 3.5 assuming different values of risk aversion γ and elasticity of inter-temporal substitution Ψ.

by the option implied methods using the risk-premium. Underlying the log dividend price ratio, the variation in the weights changes fundamentally. The hedging demand component has a far stronger influence on the variation of the weights. It can be well observed that the hedging demand has on average positive fraction in the risky asset, which varies notably depending on the state variable. Overall, the weights reflect the findings of the average portfolio weights in table 3.6.

3.5 Empirical findings

The presented results are based on the 10-year rolling window,[3] which is the most sensible choice to maintain a sufficient long estimation window, while keeping an adequate period length to analyse the out-of-sample performance. Two different evaluation periods are considered: 1. Subprime crisis, from 18th August 2007 until 19th December 2009; 2. Post crisis, from 19th December 2009 until 15th April 2016. The subprime crisis period is evaluated separately due to the exceptional market conditions. Following this, the results of the statistical tests, impact of restricted portfolio rules and trade revenue on the performance measures are presented.

3.5.1 Performance results

The out-of-sample Sharpe ratio, Sortino ratio and CAPM are reported in table 3.7. The Sharpe and Sortino Ratios are stated in excess to the full investment into the risky asset. The CAPM assumes that the risky asset is the market portfolio. The risky asset is in all cases the SPX and is denoted as "Index".

[3] Other choices of window lengths went down to 3 years in yearly increments starting at the 10-year rolling window. Beyond that a recursive window is considered with a fixed starting point but extending end. Results are in most cases similar. The shorter the interval the more inconsistent and unstable are the results. Furthermore, in case of short estimation windows, portfolio allocation of the long-term investor struggles to pick up the market's long-term dynamics.

x_t^I	Method	γ_P	Sharpe Ratio (%)			Sortino Ratio (%)			Alpha (%)			Beta (%)		
			Crisis	After	All	Crisis	After	All	Crisis	After	All	Crisis	After	All
	Index		-53.3	66.0	15.0	-41.8	53.4	11.4						
RP	Log	2	**37.5**	**12.5**	**19.9**	**21.8**	**41.0**	**28.4**	-11.8	**4.6**	**5.2**	-54.5	27.9	-19.0
		4	**57.2**	**6.7**	**21.1**	**47.0**	**41.9**	**32.0**	-5.5	**3.0**	**3.9**	-37.8	14.2	-15.6
		6	**59.6**	**6.3**	**15.7**	**50.1**	**40.9**	**24.3**	-4.1	**2.0**	**2.9**	-31.2	9.6	-13.9
	Mix-log	2	**40.5**	**14.1**	**20.3**	**26.5**	**43.0**	**27.2**	-11.8	**4.8**	**5.4**	-57.8	25.2	-22.0
		4	**60.1**	**7.4**	**22.4**	**50.3**	**46.9**	**32.3**	-5.3	**3.1**	**4.1**	-39.6	14.3	-16.6
		6	**61.9**	**8.1**	**17.4**	**52.5**	**49.3**	**25.2**	-4.0	**2.1**	**3.1**	-32.4	9.7	-14.5
	Bin-Tree	2	**34.8**	**18.6**	**24.0**	**21.4**	**46.8**	**30.0**	-11.8	**5.1**	**5.7**	-51.8	31.6	-15.8
		4	**56.4**	**13.9**	**23.1**	**45.5**	**49.8**	**32.1**	-5.4	**2.9**	**4.0**	-36.4	16.9	-13.7
		6	**60.3**	**14.5**	**17.9**	**50.4**	**52.0**	**25.5**	-3.9	**2.0**	**3.0**	-30.4	11.4	-12.6
	Hist	2	**38.0**	-7.8	**20.5**	**22.9**	**5.3**	**26.7**	-14.7	**0.8**	**7.2**	-70.3	99.5	1.0
		4	**42.4**	-11.9	**16.4**	**28.9**	**0.1**	**21.2**	-9.3	**0.2**	**4.3**	-47.0	58.6	-2.7
		6	**30.5**	-9.7	**11.1**	**17.4**	**3.0**	**14.4**	-8.1	**0.3**	**2.5**	-33.3	39.0	-2.8
MPR	Log	2	-46.1	-0.3	-28.3	-46.9	**24.3**	-23.3	-20.5	**1.0**	-2.3	-1.5	41.2	18.2
		4	-52.5	**2.4**	-33.0	-49.5	**27.7**	-26.7	-12.8	**0.7**	-1.5	-4.4	21.2	7.4
		6	-51.6	**3.4**	-31.4	-49.2	**29.3**	-25.5	-8.4	**0.5**	-1.0	-2.9	14.2	5.0
	Mix-log	2	-51.7	-0.9	-32.1	-48.0	**23.8**	-26.1	-21.1	**1.0**	-3.0	4.7	40.6	21.7
		4	-56.9	**1.9**	-34.9	-50.8	**28.4**	-28.2	-12.6	**0.6**	-1.7	-0.8	20.9	9.4
		6	-56.2	**2.9**	-33.4	-50.5	**30.0**	-27.1	-8.2	**0.4**	-1.1	-0.6	14.1	6.3
	Bin-Tree	2	-55.8	**1.1**	-35.6	-49.0	**25.7**	-28.4	-22.1	**1.2**	-3.7	11.4	39.2	25.1
		4	-60.6	**4.0**	-37.1	-52.1	**30.5**	-29.5	-12.7	**0.7**	-2.0	2.8	20.2	11.3
		6	-60.2	**5.0**	-35.6	-52.0	**32.2**	-28.4	-8.3	**0.5**	-1.2	1.8	13.6	7.5
	Hist	2	-46.1	-32.4	-48.8	-43.6	-22.0	-37.9	-25.3	-2.2	-7.1	23.1	49.3	36.2
		4	-39.5	-29.4	-47.9	-41.0	-18.7	-37.7	-16.2	-1.0	-3.7	-0.2	24.6	11.5
		6	-38.7	-28.4	-47.2	-40.6	-17.6	-37.2	-10.8	-0.6	-2.4	-0.7	16.4	7.3
$d_t - p_t$		2	-9.5	**34.6**	-10.7	-0.2	**46.6**	-8.5	-16.4	**5.5**	-1.7	94.4	79.3	90.4
		4	-17.3	**32.5**	-27.1	-4.7	**42.8**	-19.3	-19.1	**3.2**	-6.0	98.4	53.7	82.3
		6	-23.4	**31.7**	-36.2	-8.2	**41.2**	-25.0	-21.3	**2.3**	-8.2	100.8	41.2	78.7

Table 3.7: The panel displays the proposed out-of-sample performance measures of the SPX acting as risky asset. The elasticity of inter-temporal substitution parameter is $\Psi = 0.5$ for all cases.

First, the option implied state variables are analysed. Using option implied state variables leads to consistent excess performance relative to the underlying index for the period after the crisis. This finding is consistent across evaluation metrics. Comparing both state variables with each other favours the risk-premium. It exceeds consistently the performance measures of the market price of risk. For the period during the crisis, the risk-premium obtains large Sharpe and Sortino Ratios in excess to the index. In contrast, this finding is not in line with the values of alpha, which are all negative. However, taking into account the entire sample positive Sharpe ratios as well as positive alphas are obtained for the option implied risk-premium

Second, different methods of obtaining the option implied state variables are explored. To analyse if the flexibility of the underlying density has an impact on the results, three different density types are employed where each density has a different degree of flexibility. The performance results do not strongly favour a particular method. For the risk-premium, the more flexible density types tend to obtain slightly higher values as well as for the market price of risk. However, these differences are minor in magnitude. This might not be expected since the existing literature, for example Liu et al. (2007), favours more flexible density types when trying to predict asset prices. However, the proposed methodology only considers the first two moments. The advantage of option implied densities when trying to forecast asset prices mainly stems from the non-normal shape. Therefore, higher moments seem to be a main influence, which are not considered in the applied state variables shown here.

Third, the option implied approach is compared against the benchmark approaches. The historical approach clearly lacks behind the option implied state variables taking into account the sub-periods. Taking into account the entire sample the findings are less obvious for the case of the risk-premium but the option implied approach is still in favour. A fundamental issue is the way the risk-premium is obtained. Since the estimation relies on a past time series, expected returns are simply extrapolated into the future. Thus, the historical approach lacks on reaction speed and any changes in market conditions cannot be anticipated timely. Contrary to this observation is the log dividend price ratio. Compared to the alternative benchmark approaches it obtains more convincing out-of-sample results.

Its performance after the crisis is in excess to the underlying index and similar to the option implied approach. However, during the crisis results distinguish more strongly from option implied approach. The alphas are on average far more negative. An interesting insight is also provided by the values for Beta. The log dividend price ratio is on average more exposed to market risks compared to the option implied approach. Particularly, during the crisis the beta of the log dividend price ratio is close to one, whereas for the option implied approach Betas are negative or close to zero. When evaluating the entire sample period suggest that the option implied risk-premium is the most promising state variable. However, these findings need to be interpreted with care since the option implied approach is dependent on the transformation from risk-neutral to subjective density. Particularly, the excess performance during the crisis is difficult to justify, which leads to the high performance for the entire sample. Overall, it is not clear whether the option implied approach is universally preferable relative to the log dividend price ratio.

3.5.2 Tests

In this section the differences in the out-of-sample Sharpe Ratio and CAPM-alphas are tested. The Sharpe Ratio is evaluated for the difference between two Sharpe Ratios following the robust procedure based on bootstrap Confidence Intervals proposed by Ledoit and Wolf (2008). This procedure is used to test the difference between the Sharpe ratio of different applied portfolio allocations SR_{strat} against the simple investment into the index SR_{index}. Therefore, the null hypothesis can be defined as $H_0 : SR_{strat} = SR_{index}$. The index is chosen as benchmark since it is a simple strategy that does not require a complex method approach. Furthermore, beating the simple investment into the index with an alternative strategy seems to be a tough challenge. The alternative test for the CAPM-alpha is a standard t-test with $H_0 : alpha = 0$. The underlying index is chosen as market portfolio in order to obtain the CAPM and corresponding test statistics.

The test results are presented in table 3.8. Overall, the statistical tests applied to the Sharpe Ratio and CAPM-alpha do not suggest any statistical evident results. These may be attributed to two sources of problems, which arise when evaluating the two performance measurements. First, the sample size is relatively small due to monthly observations. Second, the variation in the portfolio allocation tends to be similar to the benchmark since both are strongly dependent on the movements of the SPX. Therefore, statistical differences are harder to detect. Nevertheless, some tendencies can be observed in the reported p-values that support the reported Sharpe ratios and CAPM results. The Sharpe ratio test does not detect any differences for the option implied and historical approach. In contrast, the log dividend price ratio attains the lowest p-values for the period after crisis taking into account the methods producing an excess Sharpe ratio relative to the index. Next to this finding the historical approach in combination with the market price of risk is close to the 10% significance level. It underlines the low Sharpe ratio falling behind the index. The t-test of the CAPM-alpha provides slightly more transparent results. For the option implied approach using the risk-premium, the relevant p-values vary around the 20% level for the period after the crisis. The same applies for the market price of risk during the crisis period. However, the p-values for the log dividend price ratio are at the 10% level. It confirms the high values for alpha in the period after the crisis. Taking into account the entire sample tends to lead to no statistically evident results even though the sample horizon is longer. However, the p-values are less varying through different specifications of state variables and values of risk-aversion.

All in all, the test statistics confirm the strength of the log dividend price ratio. They also somewhat confirm the poor performance of the historical approach. The test statistics for the option implied approach just provide a tendency but no statistically evident results.

3.5.3 Suboptimal portfolio choice results

In order to analyse the influence of the components of the portfolio weights of the long-term investor the portfolio rule is restricted while allowing the investor to optimally adjust their consumption. Table 3.9 displays the performance changes in the Sharpe Ratio when restricting the portfolio rules relative to the unrestricted policy. If the component is beneficial the Sharpe Ratio should display a negative value indicating a reduction in the Sharpe Ratio relative to the unrestricted rule.

x_t^I	Method	γ_P	Sharpe Ratio test			Alpha t-test		
			Crisis	After	All	Crisis	After	All
RP	Log	2	0.821	0.868	0.771	0.133	0.223	0.250
		4	0.691	0.948	0.733	0.439	0.235	0.222
		6	0.668	0.970	0.788	0.549	0.239	0.279
	Mix-log	2	0.799	0.842	0.803	0.122	0.195	0.230
		4	0.703	0.950	0.751	0.445	0.230	0.202
		6	0.664	0.941	0.799	0.556	0.224	0.253
	Bin-Tree	2	0.824	0.773	0.727	0.140	0.182	0.219
		4	0.709	0.829	0.701	0.458	0.207	0.212
		6	0.674	0.794	0.762	0.580	0.201	0.262
	Hist	2	0.823	0.695	0.783	0.021	0.888	0.327
		4	0.772	0.607	0.802	0.120	0.960	0.373
		6	0.865	0.631	0.872	0.114	0.899	0.450
MPR	Log	2	0.689	0.895	0.531	0.156	0.658	0.605
		4	0.639	0.960	0.465	0.114	0.590	0.544
		6	0.653	0.980	0.495	0.117	0.567	0.571
	Mix-log	2	0.652	0.881	0.470	0.152	0.671	0.508
		4	0.609	0.957	0.450	0.118	0.603	0.484
		6	0.620	0.970	0.480	0.120	0.580	0.509
	Bin-Tree	2	0.574	0.940	0.387	0.151	0.611	0.426
		4	0.558	0.992	0.399	0.121	0.545	0.425
		6	0.574	0.986	0.410	0.123	0.522	0.449
	Hist	2	0.678	0.137	0.249	0.207	0.321	0.215
		4	0.719	0.164	0.349	0.189	0.389	0.296
		6	0.726	0.225	0.322	0.189	0.414	0.308
$d_t - p_t$		2	0.910	0.299	0.646	0.572	0.068	0.836
		4	0.844	0.274	0.342	0.491	0.064	0.441
		6	0.771	0.241	0.219	0.430	0.059	0.284

Table 3.8: The left column displays the Sharpe Ratio Test testing $H_0 : SR_{strat} = SR_{index}$ and the t-Test testing $H_0 : alpha = 0$. The elasticity of inter-temporal substitution parameter is $\Psi = 0.5$ for all cases.

x_t^I	Method	γ_P	Hedging + no timing			No hedging + timing			No hedging + no timing		
			Crisis	After	All	Crisis	After	All	Crisis	After	All
RP	Log	2	**-64.54**	**-20.37**	**-28.78**	0.45	**-0.54**	**-0.50**	**-60.86**	**-21.12**	**-29.20**
		4	**-89.27**	**-11.87**	**-40.12**	0.14	**-1.50**	**-1.44**	**-85.13**	**-13.15**	**-41.17**
		6	**-91.53**	**-10.44**	**-33.01**	0.12	**-1.57**	**-1.47**	**-86.76**	**-11.93**	**-43.28**
	Mix-log	2	**-66.21**	**-22.26**	**-28.07**	**-0.18**	**-0.70**	**-1.27**	**-63.14**	**-23.12**	**-29.47**
		4	**-85.41**	**-12.55**	**-27.15**	**-0.60**	**-1.76**	**-2.26**	**-84.66**	**-14.09**	**-34.49**
		6	**-86.54**	**-12.22**	**-20.55**	**-0.50**	**-2.02**	**-2.31**	**-85.57**	**-14.03**	**-28.49**
	Bin-Tree	2	**-60.65**	**-26.88**	**-33.29**	**-0.10**	**-0.38**	**-0.35**	**-57.67**	**-27.29**	**-32.97**
		4	**-81.90**	**-19.69**	**-30.61**	0.22	**-0.28**	**-0.70**	**-81.14**	**-20.34**	**-35.06**
		6	**-85.18**	**-19.46**	**-24.24**	**-0.54**	**-0.29**	**-1.03**	**-84.18**	**-20.20**	**-28.86**
	Hist	2	26.60	**-22.59**	**-11.44**	**-0.66**	**-0.08**	**-0.26**	27.44	**-22.42**	**-10.92**
		4	30.16	**-15.24**	**-6.67**	1.66	**-0.08**	0.32	28.88	**-14.92**	**-7.04**
		6	43.82	**-16.35**	**-3.30**	3.12	**-0.03**	0.91	42.88	**-15.97**	**-3.47**
MPR	Log	2	15.19	**-7.56**	23.04	2.45	**-1.27**	**-0.60**	16.28	**-8.92**	20.11
		4	23.20	**-7.16**	31.82	2.22	**-2.12**	**-2.77**	24.99	**-9.42**	27.37
		6	22.88	**-7.02**	31.68	2.52	**-2.43**	**-3.14**	24.92	**-9.63**	26.71
	Mix-log	2	21.45	**-7.15**	27.11	1.69	**-1.34**	**-0.78**	22.34	**-8.53**	24.29
		4	28.39	**-6.74**	34.00	1.64	**-2.24**	**-2.77**	29.84	**-9.06**	29.66
		6	28.33	**-6.60**	33.90	1.88	**-2.57**	**-3.16**	29.99	**-9.27**	29.03
	Bin-Tree	2	25.30	**-9.05**	30.56	1.29	**-1.54**	**-0.79**	26.38	**-10.61**	27.74
		4	31.78	**-8.58**	36.12	1.31	**-2.57**	**-2.68**	33.53	**-11.21**	31.76
		6	31.89	**-8.43**	36.06	1.55	**-2.96**	**-3.07**	33.90	**-11.46**	31.17
	Hist	2	14.94	23.85	41.64	0.42	**-1.18**	0.19	16.82	22.89	40.20
		4	9.67	23.60	43.89	2.96	**-1.75**	0.74	12.75	22.17	41.96
		6	9.27	23.51	44.17	2.41	**-1.94**	0.16	12.80	21.93	42.12
Classic	Log(DivY)	2	**-10.06**	**-32.69**	**-11.85**	**-11.19**	7.36	**-12.52**	**-10.52**	**-28.27**	**-16.29**
		4	**-3.49**	**-29.34**	**-5.52**	**-22.83**	11.79	**-17.34**	6.23	**-23.16**	8.22
		6	1.35	**-27.84**	**-1.98**	**-18.51**	13.45	**-9.11**	13.93	**-21.28**	19.23

Table 3.9: The panel displays the difference in Sharpe Ratio between the restricted and unrestricted portfolio rule. The elasticity of inter-temporal substitution parameter is $\Psi = 0.5$ for all cases.

First, the timing component is restricted by setting it to the long run log excess return μ. For the option implied risk-premium this results in performance losses independently of the time period. It suggests that the option implied risk-premium benefits the timing abilities of the portfolio allocation. Similar findings cannot be confirmed for the option implied market price of risk. During the crisis the timing component has a negative impact whereas after the crisis this is found to be positive. Thereby, the magnitude is far lower compared to the option implied risk-premium. For the case of the historical approach market timing seems to benefit only when considering the risk-premium as state variable for the period after the crisis. For the log dividend price ratio the results are not as clear as for the option implied risk premium but the performance tends to be negatively affected when restricting the timing component.

Second, the hedging demand component is restricted. The option implied risk-premium tends to obtain higher Sharpe Ratios values during the crisis and lower Sharpe Ratios values after the crisis. This result can be easily justified since the hedging demand is mainly a fixed fraction in the risky asset. During the crisis restricting the portfolio rule reduces the losses resulting in Sharpe Ratio gains. However, after the crisis the opposite is the case and the Sharpe Ratio reduces. Using the option implied market price of risk as state variable lines up with the option implied risk-premium. Remarkably, for the entire sample the performance drops consistently for the option implied approach. Similar observations can be made for the historical approach except for the entire sample. The results of the log dividend price ratio are the opposite. When restricting the hedging demand the Sharpe ratio declines, whereas after the crisis it increases. These findings might be an outcome of the hedging demand on the fixed fraction a_0 and the weight restrictions on the risky asset. As shown in section 3.4.3 the share of the hedging demand is the highest for the log dividend price ratio. Excluding the hedging demand might on average reduce the weights to more moderate levels, which decreases the risk of large losses. All in all, the hedging demand has a smaller impact on the Sharpe Ratio compared to the timing component. This is due to its smaller share on the overall portfolio weight as shown in section 3.4.3.

The combination of no hedging demand and no timing yields roughly the net of performance losses or gains for the previous two cases. Overall, the findings in this section strengthen the results of the option implied risk-premium. Particularly, the timing abilities are consistent through periods and indices. However, in terms of magnitude, the log dividend price ratio is ahead of the option implied risk-premium.

3.5.4 Trade revenue

In order to get an understanding how actively a strategy needs to be managed the trade revenue is analysed. The trade revenue is defined as the average of the absolute difference in weight from two consecutive holding periods. Table 3.10 summarises the trade revenue for the considered evaluation periods and methods. Also the impact of trading costs is studied. In order to get insights on how trading costs impact performance, the Sharpe Ratio is analysed with and without trading costs. When trading costs are ignored the bid and ask price is simply the mid price of the ETF. This does not avoid the management fees incorporated in the ETF price itself but should still give an intuition of the influence on the Sharpe Ratio.

Overall, the results suggest that during the crisis period portfolio revenues are higher compared to the period after the crisis. This can be associated with the higher volatility and quickly changing market conditions during the crisis. Also, there is a relation between portfolio revenue and risk-aversion γ_P in the portfolio allocation. The higher the risk-aversion is the lower the trade revenue. This link can be explained by the preferences of the investor. If risk-aversion increases the willingness to invest into the risky asset lowers and the weight in the risky asset reduces. As a consequence, adjustments in the portfolio weight are of smaller magnitude and the portfolio revenue reduces.

Results on trade revenue indicate that the option implied approach tends to obtain higher trade revenues compared to the historical approach. This finding is a consequence of the estimation procedure. The historical approach provides more stable estimates, which reduces the need to adjust big portions of the risky portfolio weight. In contrast, the option implied approach is more reactive and therefore trade revenues are higher. Focusing solely on state variables suggests that for the option implied approach, the trade revenue tends to be higher using the risk-premium as state variable. For the historical approach the opposite occurs having higher revenues using the market price of risk. In contrast to these findings, the log dividend price ratio tends to have a smaller trade revenue, which is also more balanced.

In some cases the trade revenue increases with an increase in risk-aversion γ_P, which contradicts the relation stated above. This is a by-product of the portfolio weight restriction in the interval $[-1, 2]$. When the portfolio weight exceeds the limits in subsequent periods no trade revenue is recorded. Since this is more likely for lower values of γ_P, the trade revenue can be smaller compared to the higher values of γ_P.

Lastly, the impact of the trade revenue on the Sharpe Ratio is stronger during the crisis than after the crisis. This is a consequence of the higher revenues during the crisis period. The option implied and historical approach are in some cases heavily affected by trading costs resulting in peak reductions in the Sharpe Ratio of -10.58 during the crisis. After the crisis the impact reduces but is in a few cases still higher than the log dividend price ratio. The more consistent trade revenues of the log dividend price ratio positively affect the losses in the Sharpe Ratio. They have a maximum magnitude during the crisis of -3.92. Overall the impact is not as strong as for the alternative methods.

3.6 Conclusion

This chapter explored the role of option implied information from a long-term optimal asset allocation perspective. The main contribution is to propose the implied risk premium based on the risk-neutral distribution of stock prices as the state variable that drives the optimal portfolio weights. The investment strategy obtained from this parametric portfolio policy is compared against alternative investment strategies. Two types of strategies were considered. The first type constructs the optimal portfolio allocation from past information obtained from historical data and is represented by location-scale time series models such as the AR-ARCH family of models. The second type constructs the investment portfolio using a linear parametric portfolio policy in which the dynamics of the portfolio weights

x_t^I	Method	γ_P	Trade Revenue (%)			Sharpe Ratio (%) diff.		
			Crisis	After	All	Crisis	After	All
RP	Log	2	41.20	23.93	28.79	-6.90	-0.90	-3.79
		4	27.12	15.59	18.83	-4.55	-0.35	-2.52
		6	20.40	10.54	13.32	-3.34	-0.36	-2.01
	Mix-log	2	38.47	29.79	32.24	-8.99	-2.24	-5.43
		4	26.08	18.61	20.71	-5.92	-1.75	-3.80
		6	19.70	12.45	14.49	-4.36	-1.73	-3.02
	Bin-Tree	2	39.58	27.17	30.66	-10.58	-2.62	-6.22
		4	26.67	15.94	18.96	-6.71	-2.21	-4.40
		6	20.46	10.65	13.41	-4.86	-2.18	-3.43
	Hist	2	15.57	8.92	10.80	-1.99	-0.93	-1.19
		4	11.65	10.67	10.94	-1.64	-0.78	-0.99
		6	10.02	7.09	7.91	-1.62	-0.78	-0.96
MPR	Log	2	21.66	18.25	19.21	-2.70	-0.71	-1.33
		4	13.74	9.25	10.51	-4.46	-0.72	-1.95
		6	9.15	6.20	7.03	-4.42	-0.73	-1.93
	Mix-log	2	25.50	18.68	20.60	-3.00	-0.87	-1.55
		4	15.10	9.45	11.04	-4.57	-0.87	-2.09
		6	10.06	6.32	7.37	-4.55	-0.88	-2.07
	Bin-Tree	2	28.91	18.03	21.09	-3.04	-0.96	-1.64
		4	16.48	9.13	11.20	-4.43	-0.96	-2.11
		6	11.00	6.11	7.49	-4.43	-0.96	-2.10
	Hist	2	26.47	14.68	18.00	-2.28	-2.71	-1.94
		4	18.91	7.30	10.57	-3.77	-2.70	-2.44
		6	12.95	4.86	7.14	-3.72	-2.69	-2.42
$d_t - p_t$		2	30.33	16.35	20.28	-3.63	-2.65	-2.68
		4	31.23	9.03	15.28	-3.92	-2.61	-2.53
		6	33.46	6.29	13.94	-3.91	-2.60	-2.40

Table 3.10: The left column displays the trade revenues of the risky asset. The right column displays the difference between the Sharpe Ratio when the returns are calculated using bid and ask prices, and mid prices. The elasticity of inter-temporal substitution parameter is $\Psi = 0.5$ for all cases.

are driven by different state variables, particularly as competitor state variables the log dividend price ratio and the market price of risk are considered.

The performance of these different investment portfolios was assessed within the context of long-lived investors using different economic and statistical performance measures in an out-of-sample context. The results highlight the outperformance of the implied approach over historical methods. This finding is consistent with the literature that emphasizes the role of the derivatives market for extracting forward-looking information from market option prices. The second finding of the empirical study is to note that the implied risk premium is a good predictor of future returns on the risky asset. This is observed implicitly by noting the good performance of the investment portfolio based on the parametric portfolio policy rule that uses the implied risk premium as state variable. This portfolio is superior to the investment that uses the market price of risk as state variable. However, it does not beat the investment strategy proposed by Campbell and Viceira (1999) based on the log of the dividend price ratio.

Finally, it has been investigated on the performance of the investment portfolios depending on specific choices of the risk-neutral distribution function to price options. For that reason parametric and nonparametric distribution functions were contemplated. In principle, the choice of risk-neutral distribution function is an additional parameter that needs to be considered in the optimal asset allocation problem that uses implied information. However, the application to a portfolio given by the risk-free rate and the S&P 500 Index shows that the out-of-sample performance of the portfolio is quite insensitive to the specific choice of risk-neutral distribution function among the set of candidates usually proposed in the literature (lognormal density as in Black and Scholes (1973), a mixture lognormal density as in Ritchey (1990), and a nonparametric binomial tree as in Jackwerth and Rubinstein (1996)). A potential reason for this finding is that the state variables under study characterizing each investment portfolio only consider the risk-premium and volatility (market price of risk). Hence, higher moments of the risk-neutral distribution do not play any role in driving the optimal portfolio allocation.

B Appendix

B.1 Consumption and portfolio decisions

This appendix replicates the theoretical results of the optimal consumption and portfolio decisions. These findings need to be attributed to the novel study of Campbell and Viceira (1999). They are not part of the here conducted research and are purely added for completeness of this thesis.

Campbell and Viceira (1999) solve a discrete time portfolio allocation of an infinitely lived investor with Epstein and Zin (1989) and Weil (1989) preferences:

$$U(C_t, E_t[U_{t+1}]) = \left\{(1-\delta)C_t^{(1-\gamma)/\theta} + \delta(E_t[U_{t+1}^{1-\gamma}])^{1/\theta}\right\}^{\theta/(1-\gamma)} \tag{B.1}$$

with $\theta = (1-\gamma)/(1-\psi^{-1})$, δ the discount factor and ψ the elasticity of intertemporal substitution. The investor seeks an optimal allocation between a single risk-free asset with constant log return r_f, a single risky asset with log return r_{t+1} and innovations u_{t+1} in the risky asset's return:

$$r_{t+1} - E_t[r_{t+1}] = u_{t+1} \tag{B.2}$$

The innovations are normally distributed with zero mean and variance σ_u^2. After the optimal weights are determined the gross portfolio return R_p can be calculated by:

$$R_{p,t+1} = \alpha_t(R_{t+1} - R_f) + R_f \tag{B.3}$$

where $R_{t+1} = exp\{r_{t+1}\}$ and $R_f = exp\{r_f\}$. The expected excess log return on the risky asset is state dependent and derived by a single state variable x_t:

$$E_t[r_{t+1}] - r_f = x_t \tag{B.4}$$

The state variable x_t or risk-premium follows a mean-reverting AR(1) process with mean μ and persistence ϕ. The innovations η_{t+1} are conditionally homoskedastic and normally distributed with zero mean and variance σ_η^2:

$$x_{t+1} = \mu + \phi(x_t - \mu) + \eta_{t+1} \tag{B.5}$$

The two innovations η_{t+1} and u_{t+1} are correlated with each other. The covariance $\sigma_{\eta u}$ between the two innovations creates the intertemporal hedging demand for the risky asset of the long-term investor. The higher the covariance the better the ability of the risky asset to hedge against changes in the investment opportunity set over time.

In order to solve the portfolio allocation problem, the individual investor chooses his consumption and portfolio policies to maximise (B.1) subject to the budget constraint:

$$W_{t+1} = R_{p,t+1}(W_t - C_t) \tag{B.6}$$

Using this type of budget constraint, Epstein and Zin (1989) and Weil (1989) show that the optimal consumption and portfolio policy must satisfy the following Euler equation for any asset i:

$$1 = E_t\left[\left\{\delta\left(\frac{C_{t+1}}{C_t}\right)^{-1/\psi}\right\}^{\theta} R_{p,t+1}^{-(1-\theta)} R_{i,t+1}\right] \tag{B.7}$$

In the here considered case, i represents the single risky and risk-free asset. Therefore, $i = p$ and (B.7) reduces to:

$$1 = E_t\left[\left\{\delta\left(\frac{C_{t+1}}{C_t}\right)^{-1/\psi} R_{p,t+1}\right\}^{\theta}\right]. \tag{B.8}$$

After dividing (B.1) by W_t and inserting the budget constraint, Campbell and Viceira (1999) obtain the following expression for the utility per unit of wealth:

$$V_t = \left\{(1-\delta)\left(\frac{C_t}{W_t}\right)^{1-1/\psi} + \delta\left(1-\frac{C_t}{W_t}\right)^{1-1/\psi}\left(E_t[V_{t+1}^{1-\gamma}R_{p,t+1}^{1-\gamma}]\right)^{1/\theta}\right\}^{1/(1-\psi)}. \tag{B.9}$$

where $V_t \equiv U_t/W_t$. Using this form of value function enables to rewrite the expression as a combination of a power function of $(1-\delta)$ and the consumption-wealth ratio as shown by Epstein and Zin (1989) and Epstein and Zin (1991):

$$V_t = (1-\delta)^{-\psi/(1-\psi)}\left(\frac{C_t}{W_t}\right)^{1/(1-\psi)}. \tag{B.10}$$

To obtain optimal consumption and portfolio policies, Campbell and Viceira (1999) use the log-linear approximation of the Euler equation and intertemporal budget constraint of Campbell (1993). Following this approach the Euler equation (B.7) becomes:

$$0 = \theta log\delta - \frac{\theta}{\psi} E_t \Delta c_{t+1} + \theta E_t r_{p,t+1} + \frac{1}{2} var_t \left(\frac{\theta}{\psi} \Delta c_{t+1} - \theta r_{p,t+1} \right), \tag{B.11}$$

and the budget constraint (B.6):

$$\Delta w_{t+1} \approx r_{p,t+1} + \left(1 - \frac{1}{\rho}\right)(c_t - w_t) + k, \tag{B.12}$$

where $k = log(\rho) + (1-\rho)log(1-\rho)/\rho$ and $\rho = 1 - exp\{E(c_t - w_t)\}$. Based on expression (B.11) and (B.12), Campbell and Viceira (1999) characterise the optimal portfolio rule:

$$\alpha_t = \frac{1}{\gamma} \frac{E_t r_{1,t+1} - r_f + 0.5\sigma_{1,1,t}}{\sigma_{1,1,t}} - \left(\frac{1}{1-\psi}\right)\left(\frac{\gamma - 1}{\gamma}\right) \frac{\sigma_{1,c-w,t}}{\sigma_{1,1,t}}. \tag{B.13}$$

The expression for the optimal portfolio weight α_t consists of two components that capture the demand for the risky asset. The first term defines the myopic asset demand. The myopic component is proportional to the risk-premium and inversly proportional to the volatility and risk-aversion. The second component captures the intertemporal hedging demand going back to Merton (1969), Merton (1971) and Merton (1973). It describes the demand of the investor to hedge against changes in the investment opportunity set. Within this framework there are two special cases: 1. When returns are unpredictable the hedging demand is zero, so $\sigma_{1,c_t - w_t,t} = 0$; 2. When $\gamma = 1$ the portfolio weight reduces to the myopic component.

Lastly, Campbell and Viceira (1999) guess a form of the optimal consumption and portfolio policies. This is necessary since $\sigma_{1,c-w,t}$ depends on future decisions about portfolio weights and consumption as shown by Campbell (1993). Therefore, to finally solve the allocation problem Campbell and Viceira (1999) guess that the optimal portfolio weight of the risky asset α_t is linear and the log-consumption-wealth ratio $c_t - w_t$ is quadratic in the state variable x_t:

$$\alpha_t = a_0 + a_1 x_t \tag{B.14}$$

$$c_t - w_t = b_0 + b_1 x_t + b_2 x_t^2 \tag{B.15}$$

This results in unknown parameters $[a_0, a_1, b_0, b_1, b_2]$. The parameters defining the linear portfolio policy can be solved analytically:

$$a_0 = \frac{1}{2\gamma} - \frac{b_1}{1-\Psi}\frac{\gamma-1}{\gamma}\frac{\sigma_{\eta u}}{\sigma_u^2} - \frac{b_2}{1-\Psi}\frac{\gamma-1}{\gamma}\frac{\sigma_{\eta u}}{\sigma_u^2}2\mu(1-\phi) \tag{B.16}$$

$$a_1 = \frac{1}{\gamma\sigma_u^2} - \frac{b_2}{1-\Psi}\frac{\gamma-1}{\gamma}\frac{\sigma_{\eta u}}{\sigma_u^2}2\phi. \tag{B.17}$$

The remaining parameters of the log consumption wealth ratio are retrieved by a recursive non-linear system. The exact definiton of the terms and the recursive procedure to obtain $[b_0, b_1, b_2]$ is outlined in Campbell and Viceira (1999) proposition 2. The results imply that the approximate value function per unit of wealth is given by:

$$V_t = exp\left\{\frac{b_0 - \psi log(1-\delta)}{1-\psi} + \frac{b_1}{1-\psi}x_t + \frac{b_2}{1-\psi}x_t^2\right\}, \tag{B.18}$$

and $b_2/(1-\psi) > 0$. As final step we normalise the parameters defining the optimal portfolio and consumption rules as in Campbell and Viceira (1999). After the normalisation, the intercept of the optimal policy functions are 1. the optimal allocation to stocks; 2. the optimal consumption-wealth ratio when the expected excess return is zero. In this case, the risky asset has a risk-premium of zero and a myopic investor would not allocate any wealth to it. Therefore any asset demand would be generated by the intertemporal hedging demand. We achieve the normalisation by setting $a_0^* = a_0 - a_1(\sigma_u^2/2)$, $b_0^* = b_0 - b_1(\sigma_u^2/2) + b_2(\sigma_u^4/4)$ and $b_1^* = b_1 - b_2\sigma_u^2$. The parameters a_1 and b_2 are not affected by this transformation.

B.2 Portfolio weights

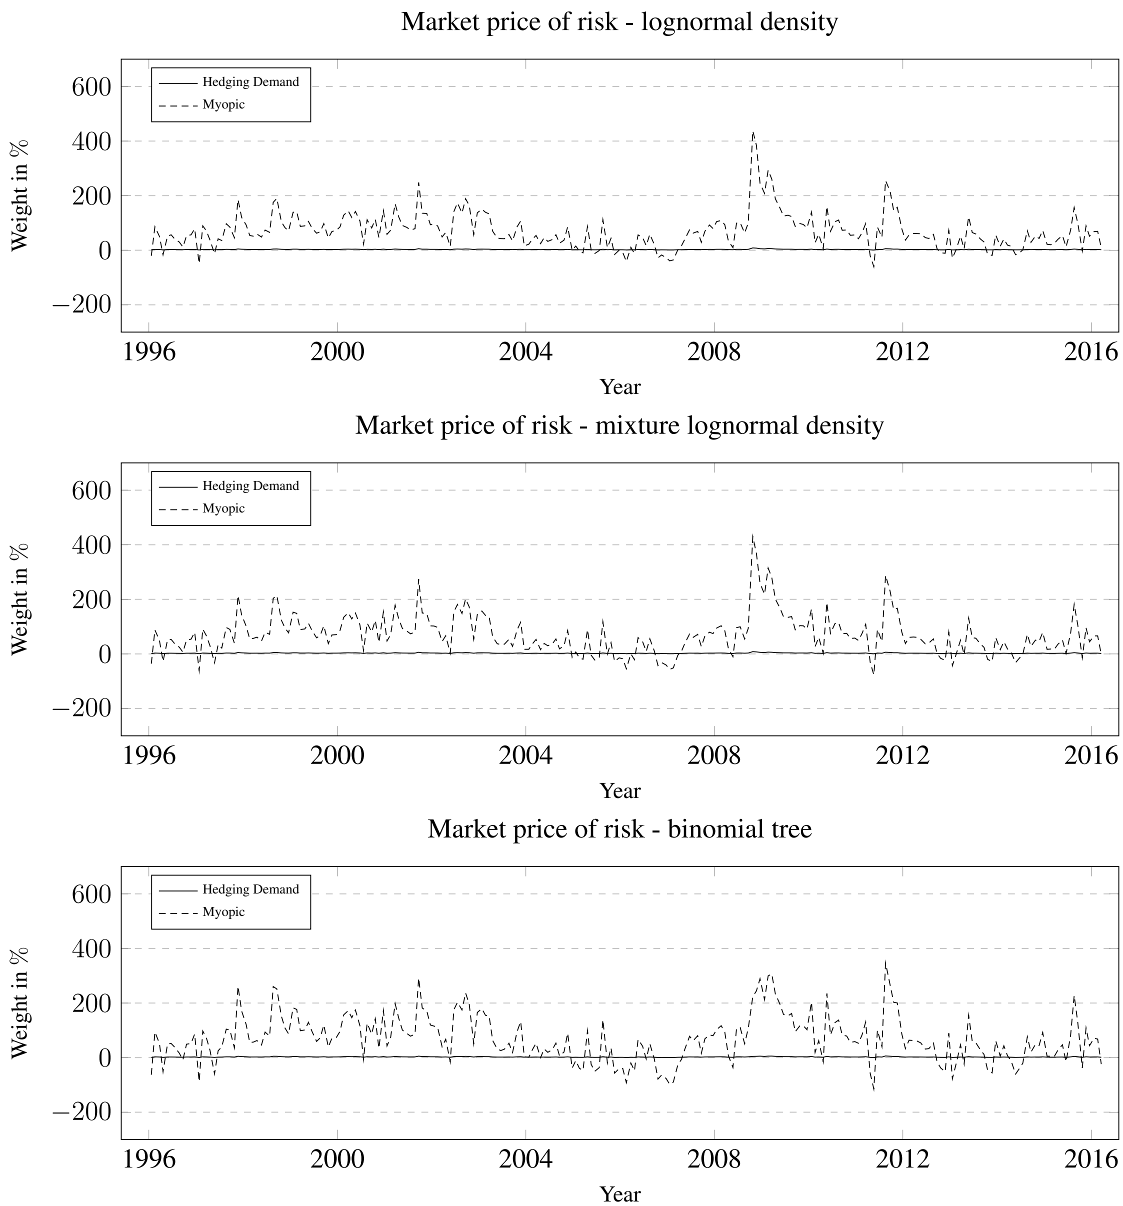

Figure 3.1: The figure displays the myopic and hedging demand following the optimal portfolio rule for the case $\gamma = 2$ and $\psi = 0.5$.

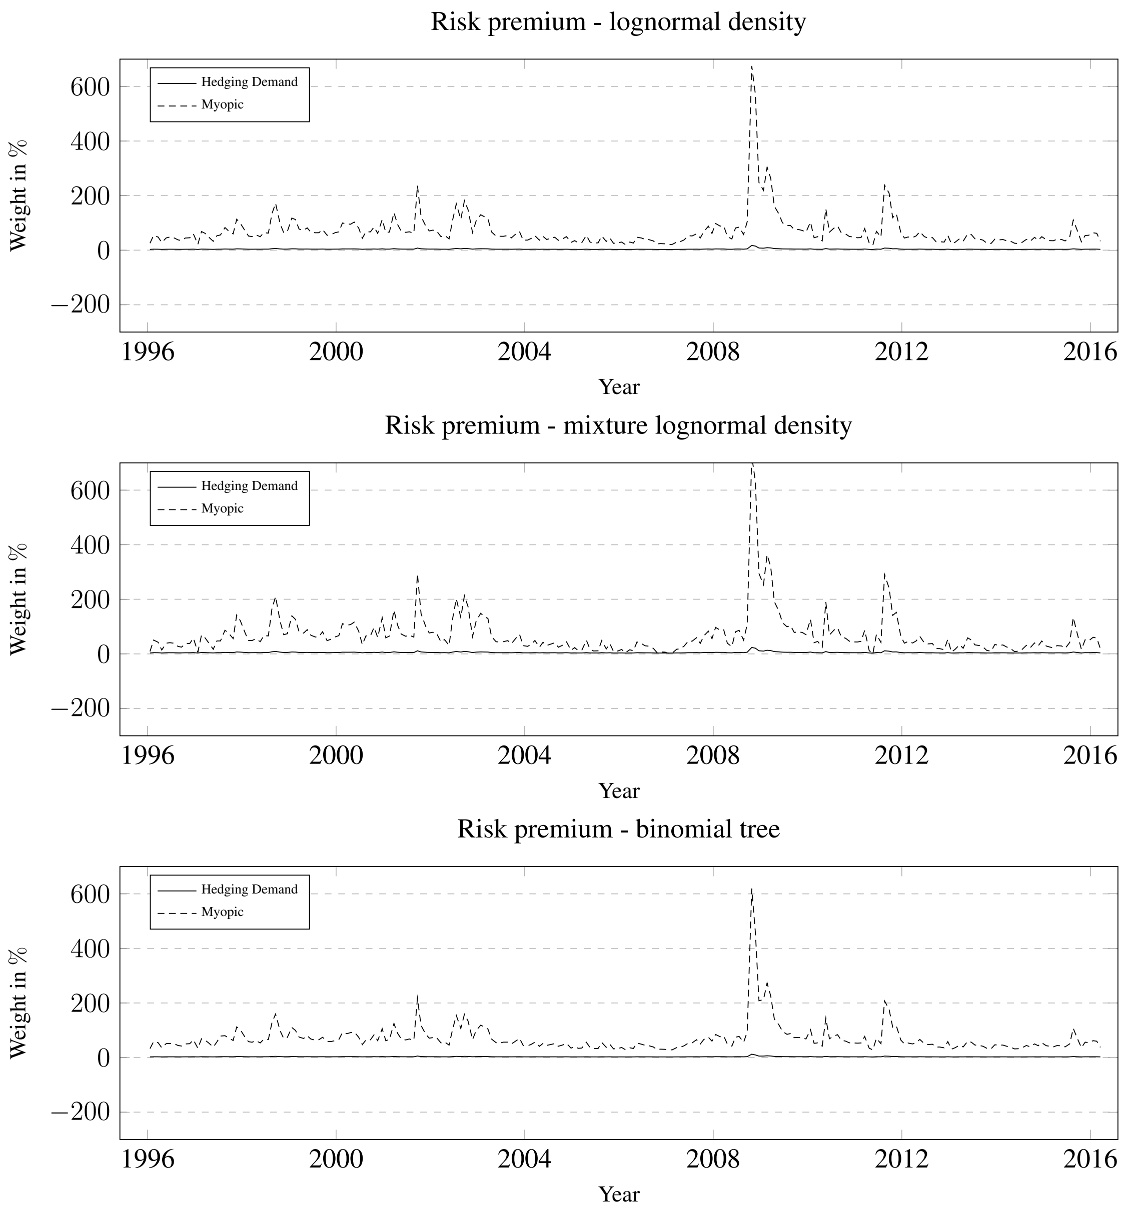

Figure 3.2: The figure displays the myopic and hedging demand following the optimal portfolio rule for the case $\gamma = 2$ and $\psi = 0.5$.

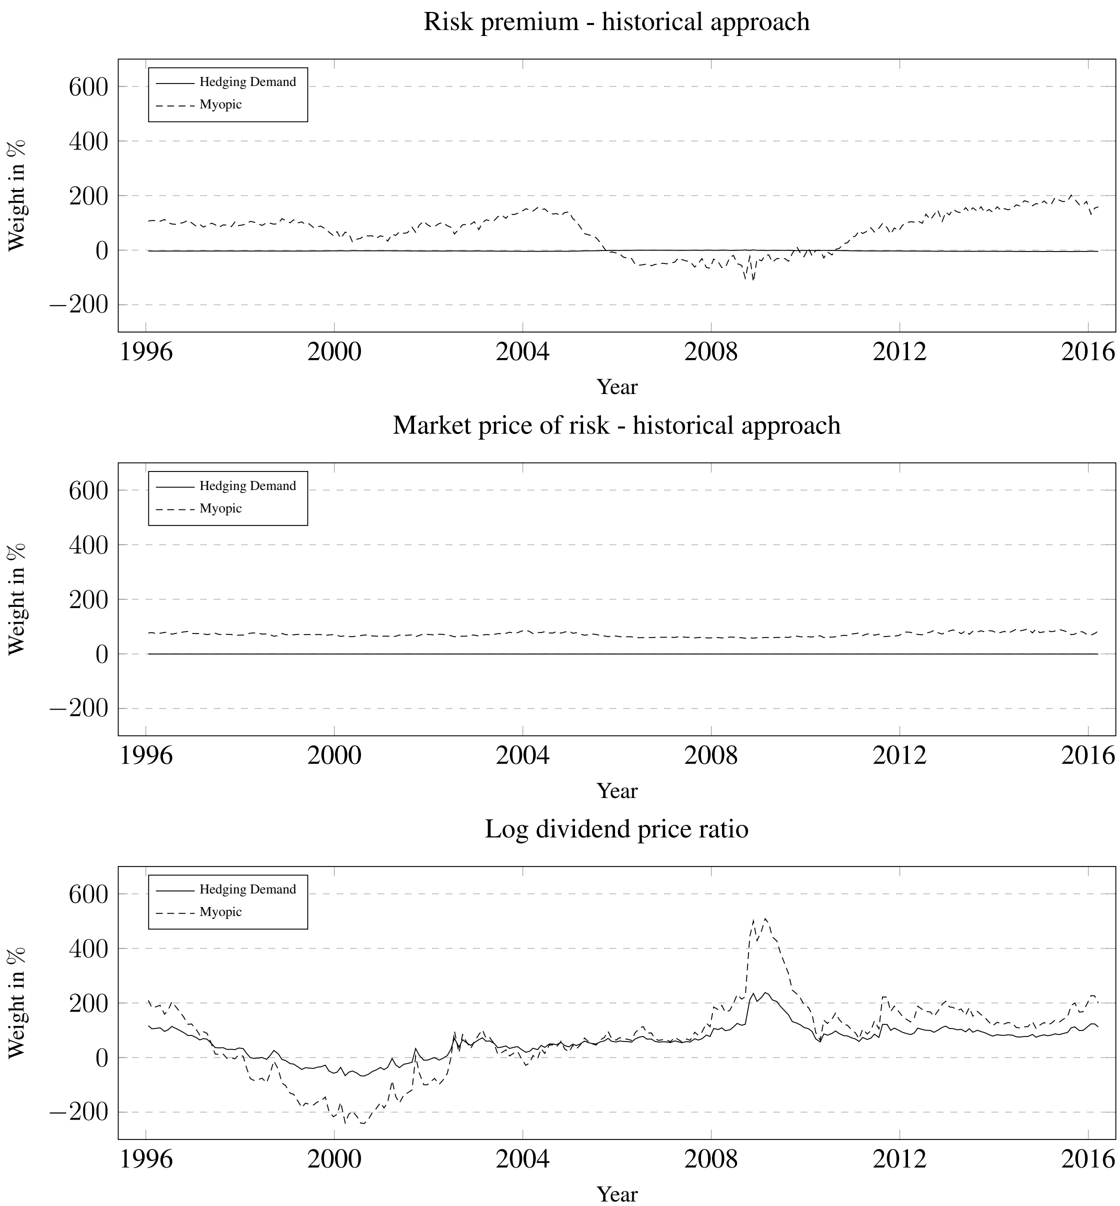

Figure 3.3: The figure displays the myopic and hedging demand following the optimal portfolio rule for the case $\gamma = 2$ and $\psi = 0.5$.

CHAPTER 4

Consumption and Portfolio Decisions when Expected Returns are Time Varying: An Option Implied Approach

This paper proposes a novel estimation method to directly obtain the dynamics of the risk-premium process from option prices. Contrary to the use of financial ratios, using a cross-section of option prices allows to directly retrieve implied estimates of the risk-premium. Using the option implied risk-premium, the estimation method obtains consistent estimates of the risk-premium process, which is then used to solve the portfolio allocation by Campbell and Viceira (1999). The findings suggest that the proposed estimation method is superior to the use of times series models as proposed in the related literature. The investor obtains high utility increases and strong improvements in the out-of-sample performance. Furthermore, the use of option prices beats a simple investment into the index and historical approach.

4.1 Introduction

The optimal portfolio allocation of a multiperiod investor is strongly dependent on the assumptions of market dynamics. A prominent observation is that asset returns vary through time and investment opportunities are not constant (Campbell, 1987; Fama and French, 1988, 1989; Campbell and Shiller, 1988a,b; Hodrick, 1992). Based on those empirical findings optimal portfolio and consumption policies can be derived within a multiperiod set-up. The most obvious impact is on the portfolio weight. Alongside to the myopic component, it accommodates a hedging demand component. This additional component is used by the investor to hedge against changes in the investment opportunity set. However, a core issue is that the dynamics and expectations, as for example expected returns, of the underlying asset cannot be directly observed. They need to be captured by suitable predictors to obtain an accurate picture about the current state. Typical predictor variables are financial ratios. Campbell and Viceira (1999); Campbell et al. (2003); Campbell et al. (2004); Chacko and Viceira (2005) incorporate such ratios within their portfolio allocation. They solve a portfolio allocation assuming an infinitely lived investor and time varying investment opportunities. To capture the changes in the opportunity set, they apply different financial ratios as predictors for the risk-premium, which are associated with some predictive power. Those financial ratios include the dividend price ratio, earnings price ratio and book to market ratio (Campbell and Shiller, 1988a; Fama and French, 1988; Cochrane, 1992). Next to financial ratios, other predictors as the short term interest rate (Fama and Schwert, 1977a; Campbell, 1987; Glosten et al., 1993) or yield spread between long- and short-term bonds (Fama, 1984; Fama and French, 1989; Campbell and Shiller, 1991) are used.

However, the use of those ratios is associated with some flaws. For example, dividends and earnings itself are more difficult to forecast than returns. Therefore, most of the variation in financial ratios is due to the variation in expected returns. This is concerning in a forecasting exercise since it implies that the use of financial ratios is conflicted in several aspects. Most prominent is the discussion on persistency. Typical tests have difficulties to reject the hypothesis of a unit root. The testing procedures evaluating the forecastability signal weak statistical evidence for a good predictive power after adjusting the test for high persistency (Nelson and Kim, 1993; Stambaugh, 1999; Ang and Bekaert, 2007; Ferson et al., 2003; Valkanov, 2003). Next to this econometric issue is the little evidence for the out-of-sample forecasting abilities as discussed in Bossaerts and Hillion (1999); Goyal and Welch (2003). Closely related are the findings that the forecasting quality of financial ratios varies over time. Paye and Timmermann (2006), Gonzalo and Pitarakis (2017) reject the hypothesis of constant regression coefficients in their study. Also in the more complex framework of Kostakis et al. (2015) to capture the dynamics of financial ratios there is no strong evidence of predictability in the post 1952 period.

The here outlined methodology aims to overcome the econometric issues associated with financial ratios by introducing an option-implied state variable. The analysis relies on the assumptions of Campbell and Viceira (1999). They assume that the risk-premium acts as state variable and follows an AR(1) process. The dynamics of the risk-premium are derived from the log dividend price ratio within a multivariate set-up. This procedure is the benchmark approach. It is referred to as "classic estimation method". The chapter contributes to this stream of literature by proposing an estimation framework that retrieves consistent estimators from risk-premium estimates inferred by option prices. It is referred to as "direct estimation method". The advantage of this method is that: First, the risk-premium is an implied estimate for the period until option expiry. Second, the risk-premium is obtained from a cross-section of option prices and does not require an estimation based on a past time series. The cross-section is in relation to the strike price and provides detailed information about the states of the underlying asset. Third, The option-implied estimate is model-free and directly obtains the risk-premium without making any assumptions about the underlying distribution, as in Martin (2017). Fourth, The dynamics of the AR(1) risk-premium process can be directly inferred using the option-implied estimates.

In the empirical analysis, the proposed approach is applied to the Dow Jones, NASDAQ-100 and S&P 500 acting as risky asset. There are three key empirical findings. First, the estimation results suggest that the option-implied risk-premium is persistent but not close to a unit root. This is validated by a Dickey-Fuller test. The results clearly reject the null-hypothesis of a unit root and therefore the econometric difficulties associated with the financial ratios are avoided.

Second, only the proposed direct estimation method delivers parameters that are empirically reasonable using the option-implied risk-premium. The key role plays the covariance between innovations in unexpected stock returns and state variable. Only for the direct estimation method the covariance is coherent negative, implying low expected returns today are followed by high returns in the future and vice versa. The negative covariance has strong implications on portfolio weights, particularly on the hedging demand. If the covariance is negative the hedging demand increases the average portfolio weight that is allocated to the risky asset. Furthermore, the direct estimation method makes the portfolio and consumption policies more reactive to changes in the investment opportunity set. Particularly, during extreme market events like the subprime crisis. The improved abilities are ultimately evaluated by the out-of-sample performance which strongly increases when comparing the classic with the direct estimation method.

Third, using the option-implied risk-premium confirms the findings in the density forecasting literature, as for example Liu et al. (2007) or Shackleton et al. (2010). The option-implied risk-premium is compared relative to a historical benchmark that relies on a past series of asset price returns. The obtained out-of-sample performance using the option-implied approach generally exceeds the values of the historical benchmark. Furthermore, the results are compared against the initially chosen approach by Campbell and Viceira (1999) using the log dividend price ratio. Unfortunately, the proposed estimation method using the option-implied risk-premium can only close up to this benchmark approach. To further analyse why there is no major difference between both methods the average utility implied by the model dynamics is taken into account. Thereby, the portfolio rules for market timing and hedging demand are restricted while solving for optimal consumption policies. For the case of the log dividend price ratio it can be seen that restricting the portfolio rule results in high losses of utility whereas for the case of the option-implied risk-premium the losses are only of minor magnitude. This is a concerning finding since it imposes that option-implied information within the chosen set-up do not benefit the investor in regards of timing the market and hedging against changes in the investment opportunity set. An attempt to overcome this issue is to adjust the volatility of the return process by an option-implied estimate to obtain a richer set of information from option prices. It can be observed that the decisions on portfolio weights and consumption become more flexible over time and higher fractions in the risky asset shift towards times of lower volatility. Even though this might improve the applied measures, the findings stand in conflict with mathematical issues. A key problem is that the correlation is now a combination of historical and option-implied information. As a consequence, the correlation exceeds in some cases -1 and the portfolio allocation cannot be solved. Therefore, this aspect needs to be studied to greater extend in the future since it is not possible within the chosen framework.

This chapter contributes to the literature by allowing for richer assumptions within the portfolio allocation using option prices and a novel direct estimation method. More recent, developments in this area focus on mean-variance portfolio allocation with multiple assets. DeMiguel et al. (2013) and Kempf et al. (2015) centre their attention on the gains of option-implied covariance. At this stage, only Kostakis et al. (2011) provides insight on multiperiod investors and the use of option-implied information. The studied allocation problem goes back to Wachter (2002). The investor follows a power utility and derives his investment decisions from the market price of risk that acts as state variable. They conclude that the out-of-sample performance exceeds the values of the myopic investor and therefore the multiperiod investor is preferable.

This study is structurred as follows. Section 4.2 describes the suggested estimation framework. Section 4.3 presents the initial estimation procedure in Campbell and Viceira (1999) and alternative state variables. Section 4.4 outlines the portfolio allocation. Section 4.5 describes the data and evaluation methods. Section 4.6 summarises the estimation and presents the out-of-sample performance. Section 4.7 presents the results using option implied volatility within the portfolio allocation. Lastly, section 4.8 reviews the main findings of this chapter.

4.2 Econometric framework

The suggested estimation framework is based on the assumptions made in Campbell and Viceira (1999). They propose that the log excess return on a risky asset is state dependent. The state of the risky is determined by a single state variable x_t:

$$E_t[r_{t+1}] - r_f = x_t. \tag{4.2.1}$$

The dynamics of the state variable are described by an AR(1) process:

$$x_{t+1} = \mu + \phi(x_t - \mu) + \eta_{t+1}, \tag{4.2.2}$$

where the innovations η_{t+1} are conditionally homoskedastic, normally distributed white noise error, $\eta_{t+1} \sim \mathcal{N}(0,\, \sigma_\eta^2)$. The innovations in the unexpected log return u_{t+1} are then defined as:

$$r_{t+1} - E_t[r_{t+1}] = u_{t+1}, \tag{4.2.3}$$

where u_t is also conditionally homoskedastic, normally distributed $u_{t+1} \sim \mathcal{N}(0,\, \sigma_u^2)$. Further, they are correlated with the innovations in the state variable. Therefore:

$$var_t(u_{t+1}) = \sigma_u^2 \tag{4.2.4}$$

$$cov_t(u_{t+1}, \eta_{t+1}) = \sigma_{\eta u} \tag{4.2.5}$$

Campbell and Viceira (1999) propose to estimate the necessary parameter values within a VAR framework. They infer the dynamics of the risk-premium process in (4.2.1) by using the log dividend price ratio as suggested in Campbell and Shiller (1988a), Fama and French (1988) and Hodrick (1992). However, using financial ratios, in particular the log dividend price ratio, is heavily discussed in the literature due to multiple econometric issues. Most prominent and concerning is the debate on stationarity and the little evidence of predictability.

Therefore, an alternative framework is proposed that does not rely on financial ratios. This section makes two contributions: First, the latent process describing the risk-premium of the risky asset is replaced by an observable proxy that is obtained from information implied by option prices; second, an autoregressive process is proposed to estimate consistent estimators. It is of particular interest to obtain consistent estimators for σ_u^2 and $\sigma_{\eta u}$.

Let $x_t^I = \widehat{E}_t[r_{t+1}^e]$, where $\widehat{E}_t[r_{t+1}^e] = \widehat{E}_t[r_{t+1}] - r_{f,t}$, be the estimated risk-premium obtained from option prices. Under some regularity conditions that will be discussed below, this process provides consistent estimates of the objective risk premium $x_t = E_t[r_{t+1}^e]$. To obtain estimates for x_t^I from option prices the approach by Martin (2017) is employed. The estimator x_t^I is constructed using a set of cross-sectional option prices. Martin (2017) shows if the negative correlation condition $cov_t(M_t R_t, R_t)$ holds, with gross return R_t and stochastic discount factor M_t, the lower bound of the risk-premium can be derived from option prices. He provides several examples that this condition holds for the here considered case of equity indices. The lower bound is defined as:

$$\widehat{E}_t[r_{t+1}^e] \geq \frac{2}{S_t^2}\left[\int_0^{F_t} p(K)dK + \int_{F_t}^{\infty} c(K)dK\right] + q_t \tag{4.2.6}$$

where S_t is the current stock price, c the call prices and p the put prices. In his analysis he shows that the lower bound is relatively tight and can be used as proxy for the risk-premium. However, the two integrals in (4.2.6) need to be approximated since strike prices cannot be observed continuously[1]. Hence, it can be safely assumed that $x_t^I = x_t + \frac{\varepsilon_t}{\sqrt{N_t}}$, with ε_t a zero-mean serially uncorrelated random variable and N_t the number of option prices on each observation date t. In this framework, the parameters of the autoregressive process (4.2.2) can be estimated using standard OLS methods. More formally, it proposes to estimate the following observable processes:

$$x_{t+1}^I = \mu + \phi(x_t^I - \mu) + \eta_{t+1}^I, \tag{4.2.7}$$

$$r_{t+1}^e = x_t^I + u_{t+1}^I, \tag{4.2.8}$$

with η_{t+1}^I and u_{t+1}^I zero-mean error terms with variance $\sigma_{\eta^I}^2$ and $\sigma_{u^I}^2$, respectively. After

[1] In Appendix C.1 the approximation of the continuous integral in (4.2.6) is outlined

some simple algebra: $\eta_t = \eta_t^I - \frac{\varepsilon_t}{\sqrt{N_t}} + \phi\frac{\varepsilon_{t-1}}{\sqrt{N_t}}$ and $u_t = u_t^I + \frac{\varepsilon_t}{\sqrt{N_t}}$. Then:

$$\begin{aligned} cov(u_{t+1},\eta_{t+1}) &= cov(u_{t+1}^I,\eta_{t+1}^I) \qquad (4.2.9)\\ &+ \frac{1}{\sqrt{N_t}}\left[cov(\eta_{t+1}^I,\varepsilon_{t+1}) - cov(u_{t+1}^I,\varepsilon_{t+1}) - \frac{1}{\sqrt{N_t}}var(\varepsilon_{t+1})\right], \text{ and}\\ var(u_{t+1}) &= var(u_{t+1}^I) + \frac{1}{\sqrt{N_t}}\left[cov(u_{t+1}^I,\varepsilon_{t+1}) + \frac{1}{\sqrt{N_t}}var(\varepsilon_{t+1})\right]. \qquad (4.2.10)\end{aligned}$$

Since η_{t+1}^I and u_{t+1}^I are not correlated with η_{t+1}, $cov(\varepsilon_{t+1},\eta_{t+1}^I) = 0$ and $cov(\varepsilon_{t+1},u_{t+1}^I) = 0$. When $N_t \to \infty$, the second term on the right hand side in (4.2.9) and (4.2.10) tends to zero. The resulting consistent estimators of $\sigma_{\eta u}$ and σ_u^2 can then be obtained via OLS:

$$\widehat{\sigma}_{\eta u} = \frac{1}{T}\sum_{t=1}^{T}(r_{t+1}^e - x_t^I)(x_{t+1}^I - \widehat{\mu} - \widehat{\phi}(x_t^I - \widehat{\mu})), \qquad (4.2.11)$$

with $\widehat{\mu}$ and $\widehat{\phi}$ the OLS estimators of the autoregressive process (4.2.8), and

$$\widehat{\sigma}_u^2 = \frac{1}{T}\sum_{t=1}^{T}(r_{t+1}^e - x_t^I)^2. \qquad (4.2.12)$$

The estimators $\widehat{\sigma}_{\eta u}$ and $\widehat{\sigma}_u^2$ allow to characterise the dynamics within the framework of Campbell and Viceira (1999). The effectiveness of the proposed direct estimation method is evaluated by solving their suggested portfolio allocation. The performance is measured out-of-sample and compared to the benchmark approaches in section 4.3. The details of the portfolio allocation are outlined in section 4.4.

4.3 Benchmark approach

This section outlines the initial estimation procedure of Campbell and Viceira (1999) using a restricted VAR model, which will be the benchmark approach for the suggested estimation method. Furthermore, the originally suggested log dividend price ratio is applied as state variable. This type of state variable is not suitable for the suggested estimation method since it does not directly reflect the risk-premium. Next to the log dividend price ratio, a state variable is employed that estimates the risk-premium from a history of asset prices. This historical approach has been commonly used in the density forecasting literature and is also applied in Kostakis et al. (2011).

4.3.1 Classic estimation method

Campbell and Viceira (1999) provide a framework to derive the dynamics of the risk-premium process in (4.2.8) from a predictor variable z_t. Typical predictors are financial ratios, which do not directly give insight into the risk-premium. Therefore, they propose to estimate a restricted VAR(1) model, which allows to derive the underlying characteristics of the risk-premium. The VAR(1) is restricted by inferring the risk-premium process solely by the lagged predictor variable.

$$\begin{pmatrix} r_{t+1} - r_{f,t+1} \\ z_{t+1} \end{pmatrix} = \begin{pmatrix} \theta_0 \\ \beta_0 \end{pmatrix} + \begin{pmatrix} \theta_1 \\ \beta_1 \end{pmatrix} z_t + \begin{pmatrix} \varepsilon_{1,t+1} \\ \varepsilon_{2,t+1} \end{pmatrix} \tag{4.3.1}$$

where $(\varepsilon_{1,t+1}, \varepsilon_{2,t+1}) \sim \mathcal{N}(0, \Omega)$ and

$$\Omega = \begin{bmatrix} \Omega_{11} & \Omega_{12} \\ \Omega_{21} & \Omega_{22} \end{bmatrix}.$$

The structure of (4.3.1) is identical to a multivariate regression model using the same explanatory variables in all equations. Therefore, under the assumption of normality there is no difference between maximum likelihood estimation and OLS regression for each of the equations. From the VAR system in (4.3.1) the values for the AR(1) risk-premium process (4.2.1) are obtained: $\mu = \theta_0 + \theta_1\beta_0/(1-\beta_1)$, $\phi = \beta_1$, $\sigma_\eta^2 = \theta_1^2\Omega_{22}$, $\sigma_u^2 = \Omega_{11}$ and $\sigma_{\eta u} = \theta_1\Omega_{12}$. As predictor variable z_t the log dividend price ratio, the historical approach in subsequent section and the option-implied risk-premium is used.

4.3.2 Alternative state variables

The results are compared against two alternative state variables. First, the log dividend price ratio as proposed by Campbell and Viceira (1999). It is defined as the log of the total dividend paid over the past year divided by the current stock price. Second, a historical benchmark. The historical benchmark relies on information contained in a past series of asset price returns. It is well known that asset price behaviour can be well described by ARCH type models going back to Engle (1982). To describe the returns an AR(1) process is used with constant μ, persistence parameter ϕ and residuals ε:

$$r_{t+1} = \mu + \phi r_t + \varepsilon_{t+1}. \tag{4.3.2}$$

The residuals consist of a conditional volatility component h and white noise process u. h is modelled by an asymmetric GARCH type model as in Glosten et al. (1993). Their model imposes that the impact of negative returns on volatility is stronger compared to positive returns. This aspect is accounted for by a dummy variable I_t which depends on the last return. This structure of the GARCH model provides more flexility and results in a more flexible shape of the return distribution. To further account for a non-normal shape the white noise component is assumed to be t-distributed as suggest by Bollerslev (1986) to allow for fat tails:

$$\varepsilon_{t+1} = u_{t+1}\sqrt{h_{t+1}} \tag{4.3.3}$$

$$h_{t+1} = w + (\alpha + \beta I_t)\varepsilon_t^2 + \delta h_t \tag{4.3.4}$$

$$I_t = \begin{cases} 0, & \text{if } r_t \geq 0 \\ 1, & \text{if } r_t < 0 \end{cases} \tag{4.3.5}$$

where $w, \alpha, \beta, \delta > 0$ and $\alpha + \frac{\beta}{2} + \delta < 1$. Beyond this approach three alternative specifications are applied. 1. The t-distribution is replaced by a normally distributed white noise process, the return process in (4.3.2) is replaced by a constant mean ($\phi = 0$) using a 2. normally and 3. a t-distributed white noise process.

The parameter estimation is matched with the option implied approach. The GARCH model is estimated on each option observation date using the last 10 years of daily historical asset prices. In Liu et al. (2007) and Kostakis et al. (2011) the historical approach is implemented via Monte-Carlo simulation to obtain the distribution of asset price outcomes. For the here considered allocation problem this is not necessary since only the expected return is of interest. Therefore, the exact solution for the conditional mean of the return process in (4.3.2) is retrieved:

$$\widehat{E}_t[r^e_{t+1}] = \hat{\mu} + \hat{\theta} r_t + q_t - r_{f,t} \tag{4.3.6}$$

The expected dividend q_t needs to be added on top of the expected return since any dividend payments negatively effect the underlying index value and are not reflected in the conditional mean. Due to the similar results of the applied ARCH model specifications, section 4.6 only reports the values of the model in (4.3.2) with t-distributed white noise process.

4.4 Application to portfolio allocation

The outlined framework in section 4.2 relies on the assumptions made by Campbell and Viceira (1999). Consequently, it is natural to use their suggested portfolio allocation to evaluate the proposed estimation method. They solve a discrete time portfolio allocation between a risk-free and risky asset of an infinitely lived investor with Epstein and Zin (1989) and Weil (1989) preferences:

$$U(C_t, E_t[U_{t+1}]) = \left\{(1-\delta)C_t^{(1-\gamma)/\theta} + \delta(E_t[U_{t+1}^{1-\gamma}])^{1/\theta}\right\}^{\theta/(1-\gamma)} \tag{4.4.1}$$

with $\theta = (1-\gamma)/(1-\psi^{-1})$, δ the discount factor and ψ the elasticity of intertemporal substitution. The gross portfolio return R_p is calculated by:

$$R_{p,t+1} = \alpha_t(R_{t+1} - R_f) + R_f \tag{4.4.2}$$

where $R_{t+1} = exp\{r_{t+1}\}$ and $R_f = exp\{r_f\}$. In order to solve the portfolio allocation problem, the individual investor chooses his consumption and portfolio policies to maximise (4.4.1) subject to the budget constraint:

$$W_{t+1} = R_{p,t+1}(W_t - C_t) \tag{4.4.3}$$

Using this type of budget constraint, Epstein and Zin (1989) and Weil (1989) provide the optimality conditions for the consumption and portfolio policies. To finally solve the portfolio allocation, Campbell and Viceira (1999) guess a form of the optimal consumption and portfolio policies. This is necessary since the optimal policies depend on future decisions about portfolio weights and consumption. They guess that the optimal portfolio weight of the risky asset α_t is linear and the log-consumption-wealth ratio $c_t - w_t$ is quadratic in the state variable x_t:

$$\alpha_t = a_0 + a_1 x_t \tag{4.4.4}$$

$$c_t - w_t = b_0 + b_1 x_t + b_2 x_t^2 \tag{4.4.5}$$

In the empirical application the weight of the risky asset in (4.4.4) is restricted in the range of $[-1, 2]$ to avoid unrealistic leverage. The expression of the portfolio weight and consumption-wealth ratio results in the unknown parameters $[a_0, a_1, b_0, b_1, b_2]$. The parameters defining the linear portfolio policy can be solved analytically:

$$a_0 = \frac{1}{2\gamma} - \frac{b_1}{1-\Psi}\frac{\gamma-1}{\gamma}\frac{\sigma_{\eta u}}{\sigma_u^2} - \frac{b_2}{1-\Psi}\frac{\gamma-1}{\gamma}\frac{\sigma_{\eta u}}{\sigma_u^2}2\mu(1-\phi) \tag{4.4.6}$$

$$a_1 = \frac{1}{\gamma\sigma_u^2} - \frac{b_2}{1-\Psi}\frac{\gamma-1}{\gamma}\frac{\sigma_{\eta u}}{\sigma_u^2}2\phi. \tag{4.4.7}$$

The remaining parameters of the log consumption wealth ratio are retrieved by a recursive non-linear system. The exact definiton of the terms and the recursive procedure to obtain $[b_0, b_1, b_2]$ is outlined in Campbell and Viceira (1999) proposition 2. The results imply that the approximate value function per unit of wealth is given by:

$$V_t = exp\left\{\frac{b_0 - \psi log(1-\delta)}{1-\psi} + \frac{b_1}{1-\psi}x_t + \frac{b_2}{1-\psi}x_t^2\right\}, \tag{4.4.8}$$

and $b_2/(1-\psi) > 0$.

The expression for the optimal portfolio weight (4.4.6) and (4.4.7) consists of two components that capture the demand for the risky asset. The first term in (4.4.7) defines the myopic asset demand:

$$\text{Myopic Demand} = \frac{x_t}{\gamma_P \sigma_u^2} \tag{4.4.9}$$

The myopic component is proportional to the risk-premium and inversely proportional to the volatility and risk-aversion. The remaining parts capture the intertemporal hedging demand going back to Merton (1969, 1971, 1973):

$$\text{Hedging Demand} = \frac{1}{2\gamma} - \frac{1}{1-\psi}\frac{\gamma-1}{\gamma}\frac{\sigma_{\eta u}}{\sigma_u^2}\left[b_1 + 2b_2\left(\mu(1-\phi) + \phi x_t\right)\right] \tag{4.4.10}$$

It describes the demand of the investor to hedge against changes in the investment opportunity set.

Lastly, the parameters defining the optimal portfolio and consumption rules are normalised as in Campbell and Viceira (1999). After the normalisation, the intercept of the optimal policy functions are the optimal allocation to stocks and the optimal consumption-wealth ratio when the expected excess return is zero. In this case, the risky asset has a risk-premium of zero and a myopic investor would not allocate any wealth to it. Therefore any asset demand would be generated by the intertemporal hedging demand. The normalisation is achieved by setting $a_0^* = a_0 - a_1(\sigma_u^2/2)$, $b_0^* = b_0 - b_1(\sigma_u^2/2) + b_2(\sigma_u^4/4)$ and $b_1^* = b_1 - b_2\sigma_u^2$. The parameters a_1 and b_2 are not affected by this transformation.

Now there are two special cases concerning the demand for the risky asset: 1. When returns are unpredictable the hedging demand is zero, so $\sigma_{\eta u} = 0$; 2. When $\gamma = 1$ the portfolio weight reduces to the myopic component.

4.5 Evaluation criteria

This section outlines the considered data and adjustments to it. Further, describes the applied evaluation criteria including the suboptimal portfolio choice, out-of-sample return calculation using ETF prices and performance measures.

4.5.1 Data

In the empirical study three major US indices are considered: Dow Jones Industrial Index (DJX), NASDAQ 100 Index (NDX) and S&P 500 Index (SPX). The respective observation periods range from 18th October 1997, 20th April 1996 and 20th January 1996 until the 15th April 2016. Within these periods the option expiry dates schedule the observation dates and allocation periods. Relevant option expiry dates are the third Friday each month. The estimation length is fixed at 4-weeks before expiry to avoid overlapping intervals. The option implied risk-premium is then estimated using the closing price on the option observation date. This results in 222 periods for the DJX, 240 periods for the NDX and 243 periods for the SPX.

Index	Mean	Median	Standard deviation	Minimum	Maximum
DJX	13.34	12.00	6.17	5.00	48.00
NDX	49.68	34.00	37.36	16.00	240.00
SPX	58.94	44.00	34.27	16.00	185.00

Table 4.1: Summary statistics of the amount of options on each observation date. The values vary quite strongly due to the increasing amount of available option prices over the past years.

	Deep OTM Call	OTM Call	ATM Options	OTM Put	Deep OTM Put
F/K	<0.9	0.9 - 0.97	0.97 - 1.03	1.03 - 1.10	1.10 <
DJX	0.95	10.44	51.77	30.43	6.42
NDX	8.67	15.43	21.62	22.42	31.85
SPX	3.23	15.75	22.90	24.21	33.91

Table 4.2: Summary statistics of option prices by moneyness in $\%$ for each observation date.

The summary statistics of the eligible option prices are displayed in table 4.1 and 4.2. Only European options are considered in this study and are always written on the underlying spot market index. Due to the structure of European options it allows to assume they are written on the corresponding forward as discussed in Liu et al. (2007). Option as well as future prices are obtained end of day from the Chicago Mercentile Exchange for the NDX and SPX, and from the Chicago Board of Trade for the DJX. The obtained option prices are checked against no arbitrage constraints. They must fulfil:

$$c(K) \geq (F_t - K)e^{-r_{f,t}T}, \tag{4.5.1}$$

$$p(K) \geq (K - F_t)e^{-r_{f,t}T}, \tag{4.5.2}$$

where c is the call and p the put price, F the forward and K the strike price. The options passing (4.5.1) and (4.5.2) are then selected based on their moneyness. Moneyness is defined as the ratio between current forward and strike price. To ensure prices are accurate only options that are at-the-money and out-the-money are taken into account. For in-the-money options it can be observed that they are less frequently traded and therefore prices might not be accurate. To qualify as eligible option the moneyness for a call-option (put-option) needs to exceed (be below) 0.97 (1.03). Lastly, the option price itself is checked. If there is no bid price quoted, the option is dropped from the sample since it cannot be actively traded any more. Furthermore, any options with a price of less than $3/8$ are excluded. Too low prices might not reflect the true value of the option due to the proximity of tick sizes. If necessary options are transformed via put-call parity $c(K) + Ke^{-r_{f,t}T} = p(K) + F_te^{-r_{f,t}T}$. In case there are options for the same strike the average price is used.

Corresponding future prices only exist for quarterly expiry dates, namely March, June, September and December for every year. Since options mature monthly, the closest future price is interpolated to match the monthly expiry date. This is performed using the following definition of the future price $F_t = S_te^{(r_{f,t}-q_t)T}$. The risk-free rate r_f is retrieved via bootstrapping from zero-coupon bonds. The dividend yield q_t is calculated using the total dividend on the underlying index over the previous year divided by the current stock price. Taking the logs of q_t provides the log dividend price ratio as in Campbell and Viceira (1999).

4.5.2 Suboptimal portfolio choice

In Campbell and Viceira (1999), they restrict the portfolio rules to highlight the benefits of timing the stock market and using stocks to hedge against deteriorations in the investment opportunity set. In their analysis, they focus how optimal timing and hedging provides large utility gains in the value function. This procedure is employed to: 1. Compare the utility of the different estimation procedures with each other; 2. How the suggested estimation methods benefit the portfolio components of the long-term investor.

The restricted portfolio weights are determined in the same fashion as in the case of the unrestricted rule but the consumption rule is still optimally adjusted. Appendix C.2 displays the optimal consumption and corresponding restricted portfolio rules. First, the impact of the timing component is analysed. This is performed by setting the state variable x_t to the unconditional expected log excess return μ of the derived VAR(1) system. Second, the impact of the hedging demand is analysed. Ignoring hedging demand leads to a myopic portfolio rule. This is performed by setting the covariance $\sigma_{\eta u}$ to zero. Lastly, both components are restricted by setting the equity allocation to the average allocation under the myopic portfolio rule and set the covariance to zero. In this case market timing as well as hedging demand are ignored.

4.5.3 Out-of-sample performance

The predictability of option prices is heavily discussed in the density forecasting literature. Within this stream the quality of option implied forecasts is evaluated based on the final index value at option expiry date. Index values have the disadvantage that they are not traded assets and just reflect a theoretical value. For example, indices mostly do not contain trading costs, management fees and consist of fractions of shares. To base the results of the portfolio allocation on more realistic grounds the performance of the risky portfolio allocation is evaluated using Exchange Traded Funds (ETF). The ETFs were selected based on the fund inception date and size. For the DJX the SPDR Dow Jones Industrial Average ETF trust, the NDX the Power Shares QQQ Trust series 1 and the SPX the SPDR

S&P 500 ETF Trust are chosen. Using an ETF to evaluate the performance of the portfolio allocation comes with multiple advantages. Most striking is that it replicates the risky asset in an actually traded portfolio with no fractions of shares. Furthermore, paid dividends and its taxation are directly reflected in the price. Lastly, trading costs and management fees are incorporated in the ETF's price. Beyond the trading costs within the ETF, bid and ask prices for the ETF at the listed exchange are used to implement the suggested portfolio allocation.

The out-of-sample returns $r_{t,t+1}^{strategy}$ of the suggested portfolio allocation are calculated on the basis of a theoretical fund value of one. The funds are allocated using the optimal weights from the considered portfolio allocation. The optimal weights are determined on the basis of a rolling-window estimation. The window length is 10 years in order to accurately capture the long-term dynamics in the market within the portfolio allocation. In this case, the optimal weights are determined taking into account all available information from ten years ago until t. The out-of-sample return of the portfolio allocation is then determined using the funds value FV. It is defined as the sum of the fair value of the risky asset and the risk-free asset. The return is then calculated using the FV from t until $t+1$:

$$r_{t,t+1}^{strategy} = \frac{1}{T}\frac{28}{365}log\left(\frac{FV_{t+1}}{FV_t}\right), \tag{4.5.3}$$

where T is the number of days between t and $t+1$. Dividend payments are not added since the used ETFs are reinvesting funds. Therefore, any paid dividends are reflected in the ETF price. Before applying the suggested performance measures the calculated returns are standardised to 4-week horizons. This is necessary since the forecasting horizons are not always connected to each other and it is assumed that the investor holds the portfolio until the next rebalancing date. The return calculation takes into account that shares are bought using the ask price and sold using the bid price. Depending if the portfolio is long or short in the risky asset, the funds value is determined by multiplying the shares with the bid or ask price, respectively. The remaining share in the risk-free asset is interest-bearing using the corresponding risk-free rate. In case the weight in the risky asset exceeds one, interest is charged to finance the leverage. The portfolio weight in every period is adjusted for the difference in weight between $t-1$ and t.

Using the obtained realised, out-of-sample returns $r_{t,t+1}^{strategy}$ for the portfolio allocation, the funds performance is evaluated using four criteria. First, the Sharpe Ratio. It is defined as the quotient between annualised average excess return μ to the risk-free rate r_f and annualised volatility σ of portfolio returns:

$$\widehat{\mu}^{strategy} = \frac{1}{N}\frac{365}{28}\sum_{t=1}^{N}(r_{t,t+1}^{strategy} - r_{f,t}), \tag{4.5.4}$$

$$\widehat{\sigma}^{strategy} = \left[\frac{1}{N-1}\frac{365}{28}\sum_{t=1}^{N}(r_{t,t+1}^{strategy} - \mu)^2\right]^{0.5}, \tag{4.5.5}$$

$$\widehat{\text{SR}}^{strategy} = \frac{\widehat{\mu}^{strategy}}{\widehat{\sigma}^{strategy}}. \tag{4.5.6}$$

It is a measure that describes the excess return in terms of risk. Generally, the higher the Sharpe Ratio the better the trade-off between return and risk. Second, the Sortino Ratio. It is an extension of the Sharpe Ratio by using the downside deviation σ_{adj} instead of the volatility. In the calculations of the volatility, positive as well as negative returns are considered equally. However, investors are normally more concerned about negative returns than positive. Therefore, the downside deviation takes only into account returns that fall below a specified threshold. The threshold is in this case the risk-free rate:

$$\widehat{\sigma}_{adj}^{strategy} = \left[\frac{1}{N-1}\frac{365}{28}\sum_{t=1}^{N}(\min(0, r_t^{strategy} - r_{f,t}))^2\right]^{0.5}, \tag{4.5.7}$$

$$\widehat{\text{SoR}}^{strategy} = \frac{\widehat{\mu}^{strategy}}{\widehat{\sigma}_{adj}^{strategy}}. \tag{4.5.8}$$

4.6 Empirical findings

This section presents the empirical findings applying the previously outlined evaluation criteria. Section 4.6.1 discusses the estimation results of the state variable. Section 4.6.2 elaborates on the estimation results of the portfolio allocation. Section 4.6.3 evaluates the portfolio weight and consumption decisions over the sample horizon. Lastly, section

4.6.4 analyses the utility gains when using the direct estimation method and impact on the out-of-sample performance.

4.6.1 Dynamics state variable

The results of the outlined state variables are displayed in table 4.1. Among the three methods there are great differences in the characteristics over time. The option implied risk-premium displays the highest variation. Particularly, during times of distress risk-premia increase sharply. In contrast to this observation is the historical approach. The historical approach displays less variation and is more constant over time. Further, during times of distress, as for example the subprime crisis in 2008, the risk-premium lowers.

Comparing the option implied risk-premium with the historical risk-premium suggests an inverse relation. This can be easily justified. The historical risk-premium is an extrapolation of past returns. Therefore, if markets fall in value it results in a lower risk-premium. Contrary, to this observation are the dynamics of the option implied risk-premium. When markets are in distress, the implied volatility contained in the option price tends to increase. Therefore, the integral in (4.2.6) grows, which results in an increase in the option implied risk-premium. Furthermore, the option implied approach is more reactive. It does not rely on a past series of asset price returns which makes the obtained parameter values less persistent.

The log dividend price ratio tends to react in a similar fashion as the option implied risk-premium. During times of distress the ratio increases. This is a consequence of the construction of the ratio. It considers the dividends over the past year divided by the current stock price. Due to the lag of one year for the dividend payments, changes in this factor happen relatively slow compared to the asset price. Therefore, when markets are in distress the ratio increases since the dividend payments do not adjust immediately but the asset price does.

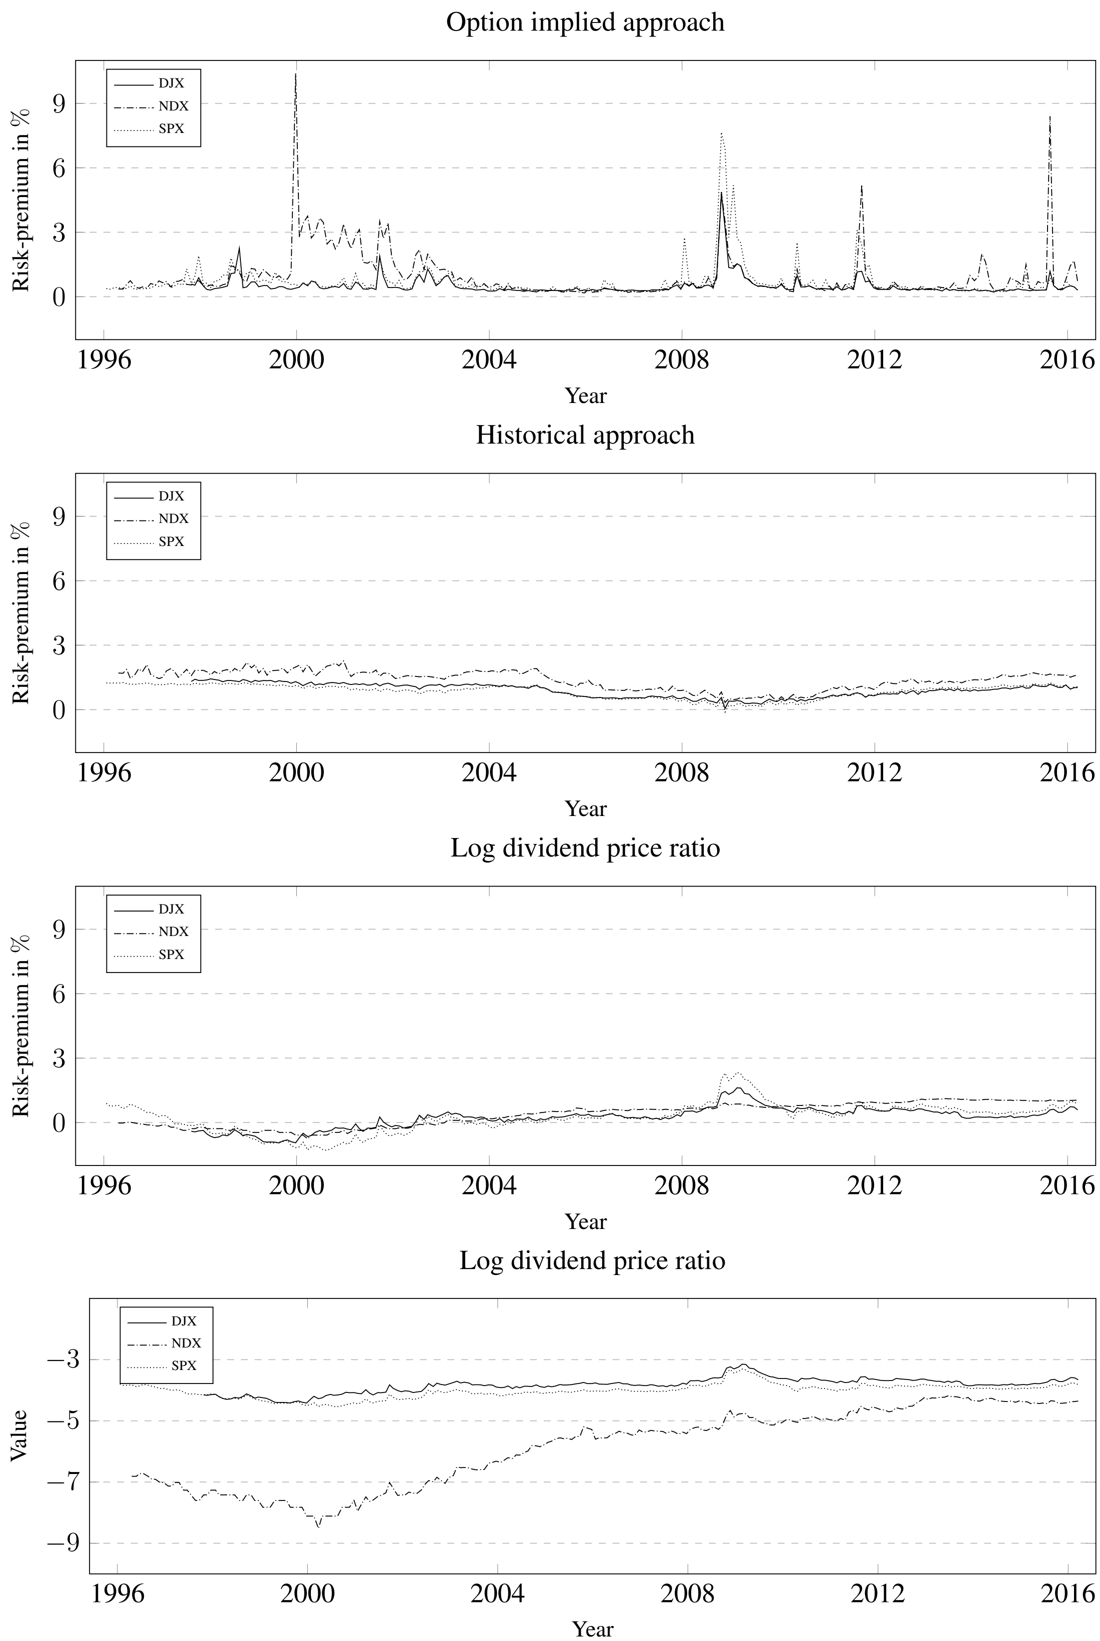

Figure 4.1: State variables over sample horizon.

The state variables of the NDX pose an exception compared to the other indices. For the option implied approach there are high risk-premia during the early 2000's. Further, the log dividend price ratio is relatively low and only comes closer to the DJX and SPX moving towards the end of the sample horizon. Only for the historical approach values are similar even though the risk-premium of the NDX mostly exceeds the values of the SPX and DJX. These observations have two reasons. First, the NDX at the beginning of the sample horizon consisted of mainly young companies. Young companies tend to pay no dividend or only low dividends. Therefore, the log dividend price ratio is relatively small. The more the companies mature over time, the higher the dividend payments grow. Therefore, the log dividend price ratio of the NDX closes up to the other two indices. Second, the dotcom bubble in the early 2000's. Due to the focus of the NDX in the technology sector, it was heavily affected by the bubble. Therefore, option implied risk-premia strongly exceed the values for the alternative indices. Further, the risk-premium of the historical approach is higher due to the strong recovery in the subsequent years. Those characteristics are later reflected in the estimation results and empirical findings.

4.6.2 Estimation portfolio allocation

The reported estimation results are on the basis of the entire sample horizon. Table 4.3 reports the results using the classic estimation method and the results using the here suggested direct estimation method.

Most striking are different estimation results for the covariance $\sigma_{\eta u}$ between the two different estimation procedures. Applying the classic estimation method for the option implied and historical approach results in inconsistent parameter values. Depending on the index, the sign of $\sigma_{\eta u}$ changes. This has strong implications on the portfolio allocation since $\sigma_{\eta u}$ dictates how great the share of hedging demand is on the overall portfolio weight, and if the weight in the risky asset is increased or lowered. Campbell and Viceira (1999) conclude that the case $\sigma_{\eta u} < 0$ is relevant since it implies mean-reversion in stock returns. This means that unexpected high returns today imply lower returns in the future. Furthermore, the magnitude of $\sigma_{\eta u}$ plays a crucial role. The higher the magnitude, the better the asset can be used to hedge against changing investment opportunities. Particularly, the positive values for $\sigma_{\eta u}$ cast doubt if the classic estimation method is suitable to capture the model dynamics.

				Classic Estimation Method				
Method	Index	μ	ϕ	σ_u^2	σ_η^2	$\sigma_{\eta u}$	ρ	DF p-value
Option	DJX	0.002	0.594	2.418E-03	1.023E-05	-9.646E-05	-0.613	0.000***
	NDX	0.004	0.452	5.746E-03	3.673E-07	8.437E-06	0.184	0.000***
	SPX	0.002	0.677	2.466E-03	1.034E-07	9.512E-06	0.596	0.000***
Hist	DJX	0.002	0.970	2.432E-03	8.884E-08	2.334E-06	0.159	0.427
	NDX	0.004	0.949	5.737E-03	9.982E-07	-1.960E-05	-0.259	0.406
	SPX	0.003	0.972	2.454E-03	7.238E-07	6.259E-06	0.149	0.423
Log(DivY)	DJX	0.004	0.973	2.409E-03	1.157E-06	-4.926E-05	-0.933	0.383
	NDX	0.029	0.998	5.716E-03	2.208E-07	-2.556E-05	-0.719	0.126
	SPX	0.002	0.975	2.419E-03	2.371E-06	-7.141E-05	-0.943	0.621
				Direct Estimation Method				
Method	Index	μ	ϕ	σ_u^2	σ_η^2	$\sigma_{\eta u}$	ρ	DF p-value
Option	DJX	0.002	0.594	2.416E-03	1.277E-05	-1.108E-04	-0.631	0.000***
	NDX	0.005	0.452	5.934E-03	1.074E-04	-1.584E-04	-0.198	0.000***
	SPX	0.002	0.677	2.563E-03	3.687E-05	-1.864E-04	-0.606	0.000***
Hist	DJX	0.000	0.971	2.448E-03	5.966E-07	5.120E-06	0.134	0.427
	NDX	0.001	0.948	5.846E-03	1.776E-06	2.172E-05	0.213	0.406
	SPX	0.000	0.972	2.469E-03	6.311E-07	5.158E-06	0.131	0.423

Table 4.3: Parameter estimates for the risk-premium process $x_{t+1} = \mu + \phi(x_t - \mu) + \eta_{t+1}$. The p-values refer to the Dickey-Fuller-Test, H_0: $\phi = 1$, H_1: $\phi < 1$.

Applying the direct estimation method overcomes these issues. For the option implied approach $\sigma_{\eta u}$ is consistently negative implying a positive hedging demand. It underpins the suitability of the direct estimation method in combination with the option implied approach. Contrary, to this observation are the results of the historical approach. Here $\sigma_{\eta u}$ displays positive values for all indices. This finding can be linked to the nature of the historical approach. Past returns are simply extrapolated into the future. Therefore, after a shock in the returns arrived, it is extrapolated into the future. Since it is well known that high returns today imply lower returns tomorrow, the extrapolation of shocks is inaccurate and has to result in $\sigma_{\eta u} > 0$. Lastly, the estimation results for the log dividend price ratio are only available for the classic estimation method since it does not directly reflect the risk-premium. Therefore, the direct estimation method cannot be applied. The estimation results line up with the findings in Campbell and Viceira (1999). Remarkably are the values of ρ compared to the alternative approaches. They are of the highest magnitude which will later on drastically affect portfolio weights. Beyond the covariance, the unconditional excess return μ and variance σ_u^2 are affected by the direct estimation method. For the option implied approach the differences between the estimation methods are of small magnitude. In contrast, the use of the direct estimation method using the historical approach drastically reduces μ.

The different estimation methods change the parameter values $[a_0^\star, a_1]$ of the portfolio rule as presented in table 4.4. Thereby, both parameters are not affected equally. The changes in $a_0^\star$ are stronger compared to a_1. This is a consequence of the changes in $\sigma_{\eta u}$ for the direct estimation method. As discussed before, the major impact of the covariance is on the hedging demand. Thereby, the hedging demand effects the constant fraction in the risky asset stronger compared to varying fraction since it ensures a sufficient share when returns are low. Furthermore, the consistently negative covariance keeps the share of hedging demand positive. The historical approach lines up with these observations. However, the positive covariance obtains a negative value for $a_0^\star$, which reduces on average the portfolio weight. Overall, the historical and option implied approach do not exceed the values of $a_0^\star$ of the log dividend price ratio. This is rooted in the higher value of the correlation which implies better hedging abilities.

Method	γ_P	Classic estimation			Direct estimation		
		DJX	NDX	SPX	DJX	NDX	SPX
				Intercept $a_0^\star \times 100$			
ModFree	2	2.34	-0.05	-0.25	2.58	1.19	5.35
	4	1.84	-0.04	-0.19	2.03	0.91	4.46
	6	1.38	-0.03	-0.14	1.53	0.68	3.43
Historical	2	-0.47	0.90	-1.25	-0.39	-0.46	-0.40
	4	-0.35	0.70	-0.89	-0.28	-0.33	-0.28
	6	-0.26	0.52	-0.65	-0.20	-0.24	-0.20
Log(DivY)	2	26.84	1.41	31.08	26.84	1.41	31.08
	4	33.65	2.96	46.01	33.65	2.96	46.01
	6	31.83	3.76	48.04	31.83	3.76	48.04
				Slope a_1			
ModFree	2	210.67	86.98	202.30	211.49	84.90	204.60
	4	106.37	43.48	101.03	106.95	42.62	105.05
	6	71.15	28.99	67.32	71.58	28.45	70.69
Historical	2	204.06	88.54	199.78	201.08	84.13	199.22
	4	101.66	44.64	98.97	99.80	41.74	98.84
	6	67.69	29.84	65.78	66.37	27.75	65.73
Log(DivY)	2	256.40	94.41	291.02	256.40	94.41	291.02
	4	152.26	52.81	201.31	152.26	52.81	201.31
	6	110.04	38.24	159.71	110.04	38.24	159.71

Table 4.4: The table displays the parameter estimates for $a_0^\star$ and a_1. All values are based on the estimates of the return processes reported in table 4.3. The elasticity of inter-temporal substitution parameter is $\Psi = 0.5$ for all cases.

The impact on the slope coefficient a_1 is rather small. There are no major differences between the two estimation methods. The covariance mainly impacts the constant fraction, as can be seen in (4.4.6) and (4.4.7). Further, other parameters determining a_1 changed less drastically resulting in similar values. The comparison of a_1 with the log dividend price ratio is not possible due to the different type of data.

Lastly, the persistency using the option implied risk-premium falls strongly below the unit root while still being persistent. Testing for a unit root using a Dickey-Fuller underlines that the option implied state variable does not encounter stationarity issues compared to the alternative state variables.

4.6.3 Portfolio weight and consumption

This section analyses the impact of the different estimation procedures on the portfolio and consumption decisions. Table 4.5 displays the average portfolio weight and the fraction due to hedging demand. Different values for the elasticity of intertemporal substitution were not considered since the impact on the results are only of minor magnitude as shown in chapter 3. An increase in elasticity of intertemporal substitution results in an increase in hedging demand for the log dividend price ratio. For the alternative methods this change has almost no impact on the hedging demand even though a new estimation technique is introduced.

As discussed in the previous section the main impact of the direct estimation method is a change in the covariance $\sigma_{\eta u}$, which increases the constant term $a_0^\star$ in the portfolio rule. Further, it is well known that $\sigma_{\eta u}$ is one of the main drivers in determining the hedging demand. For the option implied approach, it can be observed that for the direct estimation method the fraction due to hedging demand increases. Further, the mean allocation in the risky asset tends to be higher as well. Both observations are rooted in $\sigma_{\eta u}$ and associated increase in the constant term $a_0^\star$. Contrary to this observation are the findings of the historical approach. The direct estimation method results in lower average portfolio weights compared to the classic estimation method. On one hand, this is a consequence of the reduced value in μ. On the other hand, the fraction due to hedging demand is negative implying a reduction of the average portfolio weight. For both state variables, the results are consistent across indices using the direct estimation method. The classic estimation method does not provide consistent results and the share of hedging demand tends to be lower for the option implied approach. Comparing the log dividend price ratio relative to the two alternative state variables suggests that average portfolio weights and the fraction due to hedging demand is substantially higher.

To further analyse the portfolio weight appendix C.4 and C.5 display the weights over the sample horizon using the classic and direct estimation method. For the option implied approach it is apparent that for the NDX and SPX results are not reasonable using the classic estimation method. The weight in the risky asset and consumption-wealth ratio exhibit almost no changes. This is a surprise since the sample period covers the dotcom and subprime crisis which should strongly impact portfolio policies. Furthermore, the positive values of $\sigma_{\eta u}$ result in reductions in the portfolio weight. Especially, when risk-premia are high, reductions in the weight of the risky asset are strong. This stands in contradiction with the previously explained characteristics of asset price returns.

Applying the direct estimation method overcomes these issues. Consistently through the applied risky assets, the portfolio weights are more reactive. Furthermore, when there are high risk-premia in the market the portfolio weight increases. Therefore, the weight in the risky asset behaves in a similar way as the state variable. Ultimately, the consumption wealth ratio is more reactive to changes in investment opportunities. Generally, in times of financial distress the ratio increases whereas during more calm periods the ratio lowers. Compared to the historical approach, the option implied approach is more reactive. Further, the historical approach displays similar issues following the classic estimation method. The time series of the portfolio weight and consumption-wealth ratio is inconsistent across indices. When applying the direct estimation method, this problem is mostly overcome. However, due to the negative hedging demand, the portfolio weight is confronted with the issues explained before.

The log dividend price ratio is closer to the option implied approach, displaying similar

Method	γ_P	Classic estimation			Direct estimation		
		DJX	NDX	SPX	DJX	NDX	SPX
		Mean optimal percentage allocation to risky asset: $\alpha_t = [a_0^\star + a_1(\mu + \sigma_u^2/2)] \times 100$					
ModFree	2	74.63	57.95	70.18	71.10	71.52	77.48
	4	38.33	28.95	34.99	36.69	36.21	41.49
	6	25.79	19.30	23.30	24.73	24.24	28.35
Historical	2	70.57	58.14	75.55	28.45	29.28	27.89
	4	35.04	29.55	37.16	14.04	14.42	13.75
	6	23.31	19.82	24.64	9.32	9.57	9.13
Log(DivY)	2	160.07	300.41	135.69			
	4	112.77	170.21	118.38			
	6	89.01	124.88	105.45			
		Fraction due to hedging demand (in %): $[\alpha_{t,h}(\mu;\gamma,\phi)/\alpha_t(\mu;\gamma,\phi)] \times 100$					
ModFree	2	4.93	-0.13	-0.60	5.68	2.41	11.25
	4	7.46	-0.20	-0.90	8.60	3.63	17.14
	6	8.31	-0.22	-1.00	9.59	4.04	19.15
Historical	2	-1.41	3.09	-3.68	-2.99	-3.25	-3.11
	4	-2.12	4.68	-5.41	-4.38	-4.81	-4.55
	6	-2.36	5.22	-5.96	-4.83	-5.33	-5.01
Log(DivY)	2	32.63	7.79	45.23			
	4	52.18	18.63	68.61			
	6	59.61	26.06	76.51			

Table 4.5: The left panel displays the mean optimal percentage allocation to the risky asset. The lower panel displays the mean hedging demand relative to the mean total demand ($\alpha_{t,hedging}(\mu;\gamma,\Psi) = \alpha_t(\mu;\gamma,\Psi) - \alpha_{t,myopic}(\mu;\gamma,\Psi)$). All values are based on the estimates of the return processes reported in table 4.3. The elasticity of inter-temporal substitution parameter is $\Psi = 0.5$ for all cases.

characteristics over time. Differences are mainly in the slower adjustment of the portfolio weight. This is rooted in the calculation of the log dividend rice ratio as discussed in section 4.6.1. Additionally, the variation of the portfolio weight is stronger reflected in changes of the consumption-wealth ratio. It is more reactive and peaks during the financial crisis. For the NDX the log dividend price ratio poses and exception. Due to the strong growth of dividends, values are not reasonable.

Lastly, the hedging demand for the risky asset is reported in appendix C.6 and C.7. The variation over time in the hedging demand confirms the positive improvements when taking into account the entire demand for the risky asset. For the classic estimation method the hedging demand hardly changes over time and is close to zero for the option implied approach. The historical approach obtains similar characteristics but also displays greater inconsistencies between different risky assets. Applying the direct estimation method leads to more consistent movements comparing the different risky assets to each other. Further, the variation in hedging demand and magnitude increase tremendously particularly for the SPX. For the historical approach the variation is not as high as for the option implied approach but the hedging demand behaves more consistent comparing the different risky assets with each other. However, comparing the hedging demand with the log dividend price ratio still suggests tremendous differences in magnitude as well as variation over time. Major reason for the differences is the covariance $\sigma_{\eta u}$ that is responsible for the hedging demand.

4.6.4 Utility and out-of-sample performance

To quantitatively evaluate the direct estimation procedure two measures are taken into account: utility and out-of-sample performance. The results for the unconditional mean value function per unit of wealth are reported in table 4.6 and 4.7. For the unrestricted portfolio rule the absolute values are reported whereas the restricted portfolio rules display the relative change in the value function. For the option implied approach the unrestricted rule obtains high utility gains when using the direct estimation method. When using the classic estimation procedure the utility of the option implied approach cannot even exceed the historical approach. However, using the direct estimation method, the utility closes mostly up to the log dividend price ratio and clearly exceed the results of the historical approach. Further, this finding is consistent through indices, making the direct estimation method preferable.

Furthermore, restricting the hedging demand and timing component results in reductions in the value function. These reductions have a stronger impact using the direct estimation method. This is an important finding since it emphasizes that the direct estimation method better captures the information contained in option prices. Furthermore, it strengthens to allow for the more complex features of the portfolio allocation. Nevertheless, compared to the log dividend price ratio, the timing and hedging component have a rather small impact. Therefore, it is questionable if model assumption are suitable for the use of option implied information and if different assumptions about the model dynamics should be considered in order to fully accelerate information contained in option prices.

The out-of-sample performance results underline the findings of the value function. The out-of-sample performance is evaluated using 10-year rolling windows to capture the dynamics of the state variable. The window length of 10 years is selected as it provides a sufficiently long estimation horizon to capture long-run dynamics without loosing too much of the out-of-sample evaluation period. The performance is evaluated separately for the subprime crisis from 18th August 2007 until 19th December 2009 and after the crisis from 19th December 2009 until 15th April 2016 taking into account the Sharpe and Sortino Ratio. The evaluation period is split into two parts due to the exceptional market environment during the financial crisis. The results are reported in table 4.8 in excess to the simple investment into the index.

For the option implied approach the performance measures consistently improve during both market periods using the direct estimation method. It confirms the higher utilities reported in the previous paragraph. Furthermore, after the crisis the performance measures are in all cases in excess relative to the simple investment into the index. This is an important finding since the results using the classic estimation method were mostly behind the full investment into the risky asset.

	Method	γ_P	Classic estimation: Timing DJX	NDX	SPX	No timing DJX	NDX	SPX
Hedging	Mod-free	2	0.35	0.27	0.22	-0.45	-0.24	-0.23
		4	0.24	0.20	0.18	-0.23	-0.12	-0.11
		6	0.21	0.18	0.17	-0.16	-0.08	-0.08
	Hist	2	0.23	0.30	0.31	-2.17	-2.19	-4.41
		4	0.18	0.22	0.21	-1.09	-1.12	-2.17
		6	0.17	0.19	0.19	-0.72	-0.75	-1.44
	Log(DivY)	2	0.86	16.45	1.22	-11.33	-88.66	-16.44
		4	0.60	6.72	0.88	-7.67	-81.36	-13.01
		6	0.48	4.14	0.71	-5.89	-76.73	-10.98
No hedging	Mod-free	2	-4.38	0.14	0.44	-4.80	-0.11	0.21
		4	-4.11	0.12	0.36	-4.32	0.00	0.25
		6	-3.32	0.09	0.28	-3.47	0.01	0.20
	Hist	2	1.07	-3.18	3.45	-1.17	-5.26	-1.32
		4	0.86	-2.85	2.73	-0.26	-3.90	0.36
		6	0.66	-2.26	2.06	-0.09	-2.96	0.49
	Log(DivY)	2	-44.80	-23.04	-54.13	-49.28	-92.47	-58.60
		4	-53.95	-47.89	-65.55	-55.82	-91.45	-67.21
		6	-52.93	-59.59	-65.91	-54.20	-90.58	-67.00

Table 4.6: This table displays the expected value function in (4.4.8) implied by the risk-premium process in table 4.3 when setting the state variable to its long-run mean $\mu + \frac{\sigma_u^2}{2}$ for the classic estimation method. The upper left hand side of each table reports the utility of the unrestricted portfolio rule. The remaining parts display the utility losses when following a suboptimal portfolio rule. The elasticity of inter-temporal substitution parameter is $\Psi = 0.5$ for all cases.

	Method	γ_P	Direct estimation					
			Timing			No timing		
			DJX	NDX	SPX	DJX	NDX	SPX
Hedging	Mod-free	2	0.38	0.97	0.96	-0.50	-1.08	-1.55
		4	0.25	0.47	0.49	-0.26	-0.55	-0.82
		6	0.21	0.34	0.36	-0.17	-0.36	-0.55
	Hist	2	0.22	0.22	0.22	-1.94	-1.18	-2.20
		4	0.18	0.18	0.18	-0.95	-0.58	-1.08
		6	0.17	0.17	0.17	-0.63	-0.39	-0.71
No hedging	Mod-free	2	-4.94	-3.14	-12.84	-5.39	-4.17	-14.06
		4	-4.71	-3.40	-14.22	-4.93	-3.91	-14.83
		6	-3.84	-2.96	-12.62	-3.99	-3.30	-13.03
	Hist	2	1.24	1.64	1.36	-0.78	0.39	-0.95
		4	0.95	1.28	1.04	-0.05	0.66	-0.10
		6	0.71	0.96	0.78	0.04	0.55	0.02

Table 4.7: This table displays the expected value function in (4.4.8) implied by the risk-premium process in table 4.3 when setting the state variable to its long-run mean $\mu + \frac{\sigma_u^2}{2}$ for the direct estimation method. The upper left hand side of each table reports the utility of the unrestricted portfolio rule. The remaining parts display the utility losses when following a suboptimal portfolio rule. The elasticity of inter-temporal substitution parameter is $\Psi = 0.5$ for all cases.

Compared to the historical approach, the option implied approach benefits more heavily by the direct estimation method. For the historical approach the Sharpe and Sortino ratio increase in many cases but only of small magnitude. It confirms the weak timing abilities as discussed in the density forecasting literature. Furthermore, it underlines the adverse estimation results. In contrast, the log dividend price ratio obtains high values for Sharpe and Sortino ratio. The option implied approach is only able to close up to the log dividend price ratio but results do not exceed the performance measures. It lines up with the findings taking into account the utility.

4.7 Extension for option implied variance

Following the proposed direct estimation method leads to high weights in the risky asset during times of distress. Even though the portfolio weight reacts in the intended way, the long-term investor misses out during periods of high gains as for example from 2004 until 2007 or 2012 until 2016. The weight in the risky asset during these time periods obtains its lowest values which is not reasonable. A potential issue associated with the estimation procedure is the estimate of the return volatility σ_u^2. It is an estimate based on past returns. However, it is questionable if this estimate correctly reflects the volatility in the risky asset. From volatility forecasting it is well known that option implied volatility has advanced forecasting abilities compared to alternative methods, e.g. Blair et al. (2001), and varies over time. Therefore, the estimation procedure for σ_u^2 is amended by obtaining estimates from option prices. Even though this procedure does not allow for stochastic volatility it should help to improve the portfolio allocation. The option implied estimates of σ_u^2 are constantly updated and should more accurately reflect uncertainties in the market. Therefore, the implied estimates should deliver a more accurate picture about current market expectations.

To follow a similar methodology to retrieve the volatility as the risk-premium, the model-free approach by Britten-Jones and Neuberger (2000) is employed. Jiang and Tian (2005) extend this procedure to account for interest and dividend payments. To determine

Risky asset	Method	γ_P	Classic estimation				Direct estimation			
			Sharpe Ratio		Sortino Ratio		Sharpe Ratio		Sortino Ratio	
			Crisis	After	Crisis	After	Crisis	After	Crisis	After
DJX	Index		-0.549	0.605	-0.440	0.484	-0.549	0.605	-0.440	0.484
	Mod-free	2	-0.438	-0.068	-0.390	-0.003	-0.089	**0.162**	-0.048	**0.186**
		4	-0.393	-0.034	-0.350	**0.036**	-0.024	**0.208**	-0.016	**0.246**
		6	-0.375	-0.023	-0.335	**0.050**	**0.038**	**0.220**	**0.024**	**0.262**
	Hist	2	-0.567	-0.141	-0.536	-0.072	**0.052**	-0.085	**0.013**	-0.066
		4	-0.656	-0.241	-0.581	-0.176	**0.101**	-0.119	**0.048**	-0.089
		6	-0.654	-0.221	-0.580	-0.155	**0.116**	-0.092	**0.059**	-0.062
	Log(DivY)	2	-0.370	**0.177**	-0.224	**0.225**				
		4	-0.470	**0.181**	-0.278	**0.205**				
		6	-0.514	**0.180**	-0.300	**0.199**				
NDX	Index		-0.200	0.787	-0.157	0.609	-0.200	0.787	-0.157	0.609
	Mod-free	2	-1.117	-0.035	-0.964	**0.101**	-0.063	**0.082**	-0.040	**0.147**
		4	-1.108	**0.038**	-0.960	**0.208**	-0.128	**0.143**	-0.078	**0.233**
		6	-1.105	**0.064**	-0.959	**0.270**	-0.097	**0.189**	-0.059	**0.311**
	Hist	2	-0.545	-0.205	-0.371	-0.185	**0.003**	-0.081	**0.000**	-0.060
		4	-0.519	-0.153	-0.356	-0.138	**0.039**	-0.050	**0.027**	-0.023
		6	-0.511	-0.136	-0.351	-0.122	**0.051**	-0.021	**0.036**	**0.008**
	Log(DivY)	2	-0.028	-0.190	-0.019	-0.173				
		4	**0.093**	-0.104	**0.068**	-0.084				
		6	**0.135**	-0.078	**0.103**	-0.051				
SPX	Index		-0.533	0.660	-0.418	0.534	-0.533	0.660	-0.418	0.534
	Mod-free	2	-0.043	-0.317	-0.015	-0.202	-0.109	**0.093**	-0.060	**0.119**
		4	-0.132	-0.288	-0.070	-0.169	-0.104	**0.165**	-0.051	**0.259**
		6	-0.158	-0.279	-0.084	-0.159	-0.163	**0.181**	-0.082	**0.341**
	Hist	2	**0.380**	-0.078	**0.229**	**0.053**	**0.155**	-0.027	**0.093**	-0.001
		4	**0.424**	-0.119	**0.289**	**0.001**	**0.187**	-0.021	**0.116**	**0.013**
		6	**0.305**	-0.097	**0.174**	**0.030**	**0.197**	0.002	**0.124**	**0.040**
	Log(DivY)	2	-0.095	**0.346**	-0.002	**0.466**				
		4	-0.173	**0.325**	-0.047	**0.428**				
		6	-0.234	**0.317**	-0.082	**0.412**				

Table 4.8: Annualized out-of-sample Sharpe and Sortino ratios. The elasticity of inter-temporal substitution parameter is $\Psi = 0.5$ for all cases. Positive excess values are marked bold.

the volatility estimate, they consider the forward price F under the forward probability measure. The variance can then be obtained by the following expression:

$$E_0\left[\int_0^T \left(\frac{dF_t}{F_t}\right)^2\right] = 2\int_0^\infty \frac{c^F(T,K) - max(F_0 - K, 0)}{K^2} dK \tag{4.7.1}$$

The numerator of the integral in (4.7.1) becomes the difference between forward option price c^F and its inner value. The forward option price is defined as $c^F = ce^{r_{f,t}T}$. The integral requires a continuous sequence of option prices which are not observable in financial markets. Therefore, the integral needs to be approximated. It is assumed that option prices are sorted by the corresponding strike price from low to high. N_t is the amount of observed option prices:

$$Var_t[r_{t+1}] = E_0\left[\int_0^T \left(\frac{dF_t}{F_t}\right)^2\right] \approx \sum_{i=1}^{N_t}[g(T,K_i) + g(T,K_{i-1})](K_i - K_{i-1}) \tag{4.7.2}$$

$$\text{where } g(T,K) = [c_o^F(T,K) - max(F_0 - K, 0)]/K^2 \tag{4.7.3}$$

This expression results in a downwards bias of the implied volatility as discussed in Jiang and Tian (2005) since volatilities in the extreme upper and lower tails are missing. The approximation of $Var_t[r_{r+1}]$ is then introduced into the optimisation procedure of the portfolio allocation. For each observation date, the observed option implied variance is used to estimate σ_u^2 and the portfolio optimisation is performed to obtain optimal portfolio and consumption policies. The results for the portfolio weights and consumption-wealth ratio are displayed in appendix C.10 and C.11.

The plots show that the portfolio weights change drastically and avoid some of the issues when estimating σ_u^2 from a past time series. Most apparent for the option implied approach is that the weight in the risky asset increases in periods of low volatility and lowers during times of distress. This finding clearly differentiates from not using an option implied estimate for σ_u^2. The updated option implied estimates for σ_u^2 seem to put the expected returns better into perspective and scale the share in the risky asset depending on the market environment. These findings are quite intuitive since the weight in the risky asset is high during times of gains whereas during times of distress the portfolio weight is reduced. Furthermore, the weight in the risky asset exhibits a stronger variation over time. The historical approach obtains similar results. However, the variation in weight is stronger compared to the option implied approach. Overall, the weight in the risky asset for the option implied and historical approach becomes similar volatile compared to the log dividend price ratio. Nevertheless, they diverge quite drastically as for example before and during the subprime crisis. Within this time periods the weight in the risky asset is more or less the opposite compared to the option implied and historical approach.

Next to the weight in the risky asset, the consumption wealth ratio exhibits great changes for the option implied approach. The variation over time greatly increased peaking in periods of low volatility but maintaining low levels when markets are in distress. Compared to the alternative approaches the option implied approach obtains the highest variation in the consumption wealth ratio. The historical approach increases its variation too. However, it is on a far lower level compared to the option implied approach. This is a remarkable observation since the weight in the risky asset is of similar magnitude for some sub-periods.

Even though the observations can be explained intuitively there are some shortcomings following this procedure. First, within the portfolio allocation σ_u^2 is assumed to be constant and therefore accounts for the entire estimation horizon. This is not reasonable since estimates for the volatility change over time. Second, the estimation of the parameters $[\sigma_u^2, \sigma_\eta^2, \sigma_{\eta u}]$ is performed based on two different approaches: σ_η^2 using a past time series and σ_u^2 using option implied information. The resulting covariance $\sigma_{\eta u}$ is a combination of both. This way of estimating the parameters of the portfolio allocation is not reasonable from a mathematical perspective. Using the obtained values to calculate the correlation coefficient can result in values that exceed in some cases -1. If this applies the approximation method normally cannot converge or obtains unreasonable values for optimal weights and consumption. This is of course by no means reasonable. Nevertheless, this section should highlight that the introduction of option implied volatility impacts the portfolio allocation drastically. In order to accommodate this feature in a consistent and appropriate way the model assumptions need to be extended.

4.8 Conclusion

This study presented a novel estimation method to directly obtain the dynamics of the state variable within the portfolio allocation of a long-term investor. The portfolio allocation goes back to Campbell and Viceira (1999) and considers a single state variable. In the initial paper, the dynamics of the state variable or risk-premium process were obtained indirectly via a VAR(1) system using the log dividend price ratio. However, the research concerning with option implied information advanced over the past years and could show that the risk-premium can be estimated robustly from a cross-section of option prices. Therefore, this chapter proposed a direct estimation method to obtain consistent estimates of the AR(1) process that reflects the dynamics of the risk-premium within the considered portfolio allocation.

When Applying the direct estimation method, the results become more consistent across considered risky assets and portfolio weights as well as consumption policies. The main impact is on the covariance $\sigma_{\eta u}$ between the innovations in the state variable and log excess return. This component is responsible for the hedging demand of the long-term investor. Only the direct estimation method, obtains estimates of $\sigma_{\eta u} < 0$ that are empirically reasonable as discussed in Campbell and Viceira (1999). It ensures that the hedging demand has a positive share in the overall portfolio weight of the risky asset. The positive share of hedging demand is mainly reflected in an increase of $a_0^\star$ within the portfolio rule. To evaluate the two different estimation methods, the utility of the long-term investor is compared to each other. The utility using the direct estimation method clearly exceeds the classic estimation method when using option implied information. This is also confirmed by the out-of-sample results. The direct estimation method obtains high gains in the Sharpe and Sortino ratio. Comparing the findings of the option implied to the historical approach confirms the findings of the density forecasting literature. Utility as well as out-of-sample performance are in excess to the historical approach. However, relative to the log dividend price ratio, results only close up.

The empirical findings are most likely a consequence of the model assumptions, which are tailored for the log dividend price ratio. To overcome parts of these issues the option implied variance is incorporated into the portfolio allocation. As a result, the portfolio and consumption decisions change quite drastically and seem to be more suitable taking into account the market environment. However, the model comes to its limits by obtaining unreasonable values of the correlation between innovations in the state variable and excess returns. Future work should be centred around the model assumptions to incorporate the option implied risk-premium and variance in a more advanced framework.

C Appendix

C.1 Approximation Risk-premium

Since we cannot observe a continuous series of strike prices the two integrals need to be approximated. We approximate the first integral containing only put options where N^p is the amount of observed put prices as following:

$$\begin{aligned}\int_0^{F_t} p_o(K)dK \approx \sum_{i=2}^{N^p-1} p_o(K_i)\frac{K_{i+1}-K_{i-1}}{2} \\ + p_o(K_1)(K_2-K_1) \\ + p_o(K_{N^p})\left[\frac{(K_{N^p}-K_{N^p-1})}{2}+\frac{(F_t-K_{N^p})}{2}\right].\end{aligned} \tag{C.1}$$

The second and third term refer to the put option with the lowest and highest strike price. For the put option with the lowest strike the difference in the strike price towards K_2 is mirrored to the left hand side. For the put option with the highest strike K_{N^p} the area on the right hand side of the option is only considered until the future price. The second integral containing only call options is approximated in the same fashion with the amount of call option prices N^c:

$$\begin{aligned}\int_{F_t}^{\infty} c_o(K)dK \approx \sum_{i=2}^{N^c-1} c_o(K_i)\frac{K_{i+1}-K_{i-1}}{2} \\ + c_o(K_{N^c})(K_{N^c}-K_{N^c-1}) \\ + c_o(K_{N^c})\left[\frac{(K_2-K_1)}{2}+\frac{(K_1-F_t)}{2}\right]\end{aligned} \tag{C.2}$$

C.2 Restricted portfolio policies

	Portfolio rule	Optimal consumption rule given portfolio rule
Hedging		
Timing	$\alpha_t = a_0^* + a_1(x_t + \sigma_u^2/2)$	$c_t - w_t = b_0^* + b_1^*(x_t + \sigma_u^2/2) + b_2(x_t + \sigma_u^2/2)^2$
No-timing	$\alpha_t = a_0^* + a_1(\mu + \sigma_u^2/2)$	$c_t - w_t = b_0^{h,nt} + b_1^{h,nt}(x_t + \sigma_u^2/2)$
No-hedging		
Timing	$\alpha_t = \frac{x_t+\sigma_u^2/2}{\gamma\sigma_u^2}$	$c_t - w_t = b_0^{nh,t} + b_1^{nh,t}(x_t + \sigma_u^2/2) + b_2^{nh,t}(x_t + \sigma_u^2/2)^2$
No-timing	$\alpha_t = \frac{\mu+\sigma_u^2/2}{\gamma\sigma_u^2}$	$c_t - w_t = b_0^{nh,nt} + b_1^{nh,nt}(x_t + \sigma_u^2/2)$

Table 4.9: This table is from Campbell and Viceira (1999). From left to right, the first equation describes the unrestricted intertemporal optimisation problem. The second equation allocates a fixed fraction to the risky asset that is equal to the average allocation of the risky asset imposed by the optimal portfolio rule. Market timing is ignored in combination with imperfect hedging. The third case refers to a myopic investor ignoring intertemporal hedging but trying to time the market. Lastly, the fourth investor ignores hedging and timing and allocates a fixed fraction equal to the average myopic demand of the risky asset.

C.3 Estimation errors

Parameter	Mean			Median			Standard deviation		
	DJX	NDX	SPX	DJX	NDX	SPX	DJX	NDX	SPX
μ	0.047	0.069	0.043	0.053	0.075	0.047	0.016	0.020	0.015
ϕ	-0.016	0.001	-0.017	-0.010	-0.028	-0.007	0.037	0.058	0.039
ω	0.012	0.026	0.011	0.011	0.023	0.010	0.003	0.018	0.004
δ	0.925	0.924	0.927	0.927	0.928	0.929	0.018	0.027	0.017
α	0.009	0.018	0.005	0.002	0.021	0.000	0.010	0.014	0.010
β	0.106	0.089	0.115	0.114	0.078	0.124	0.049	0.037	0.045
d.o.f.	8.236	14.618	8.504	8.303	12.631	7.889	2.001	6.280	2.580

Table 4.10: Parameter statistics of the historical approach.

C.4 Portfolio weight - direct estimation method

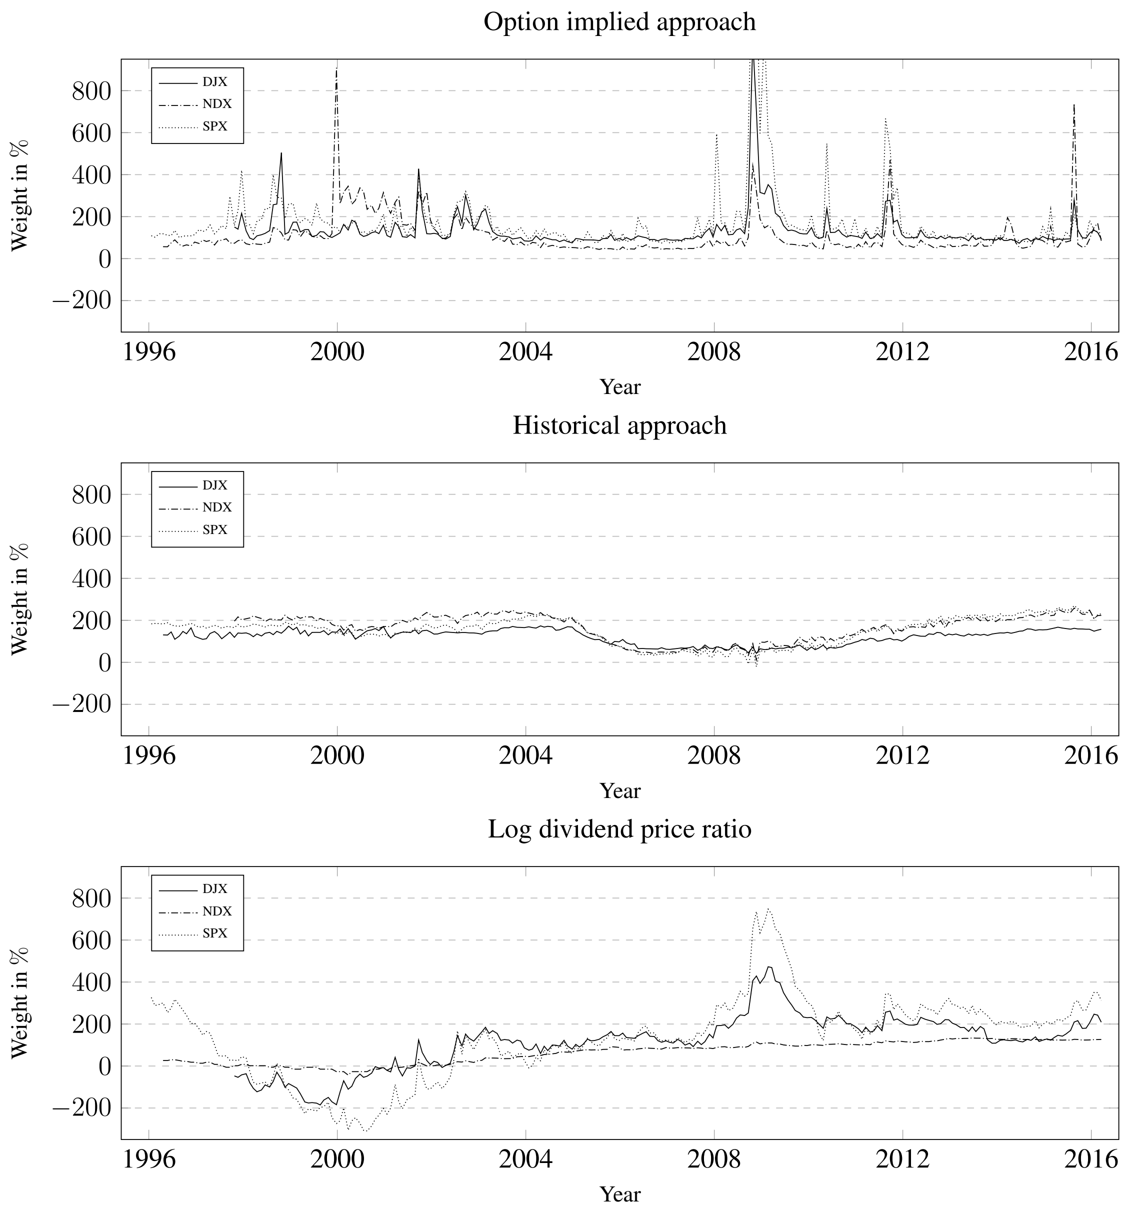

Figure 4.2: The figure displays the risky portfolio weight following the optimal portfolio rule for the case $\gamma = 2$ and $\psi = 0.5$.

C.5 Portfolio weight - classic estimation method

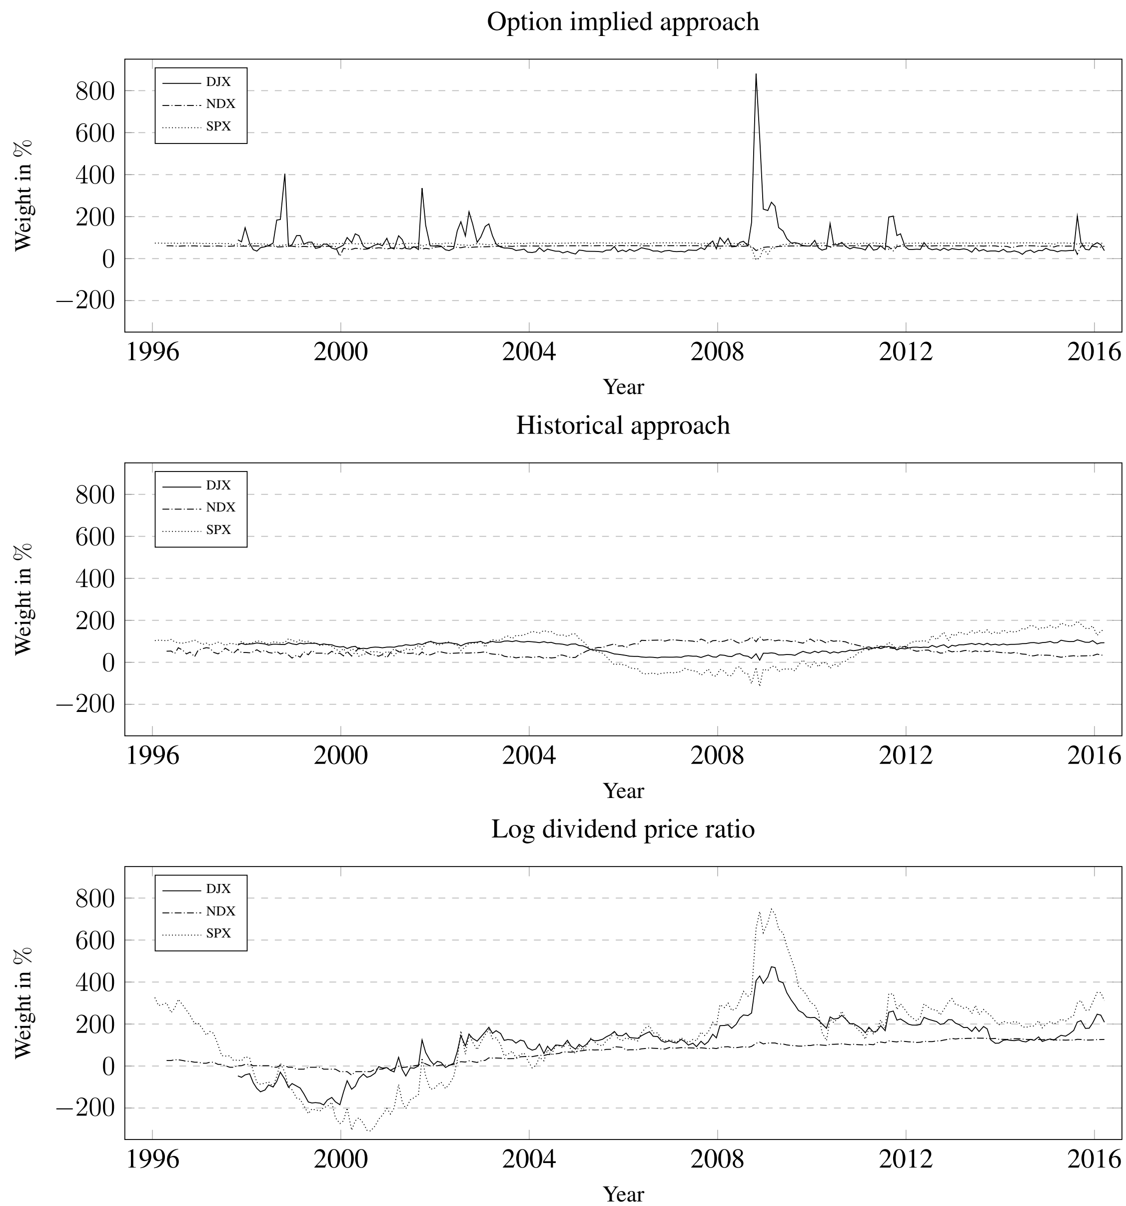

Figure 4.3: The figure displays the risky portfolio weight following the optimal portfolio rule for the case $\gamma = 2$ and $\psi = 0.5$.

C.6 Hedging demand - direct estimation method

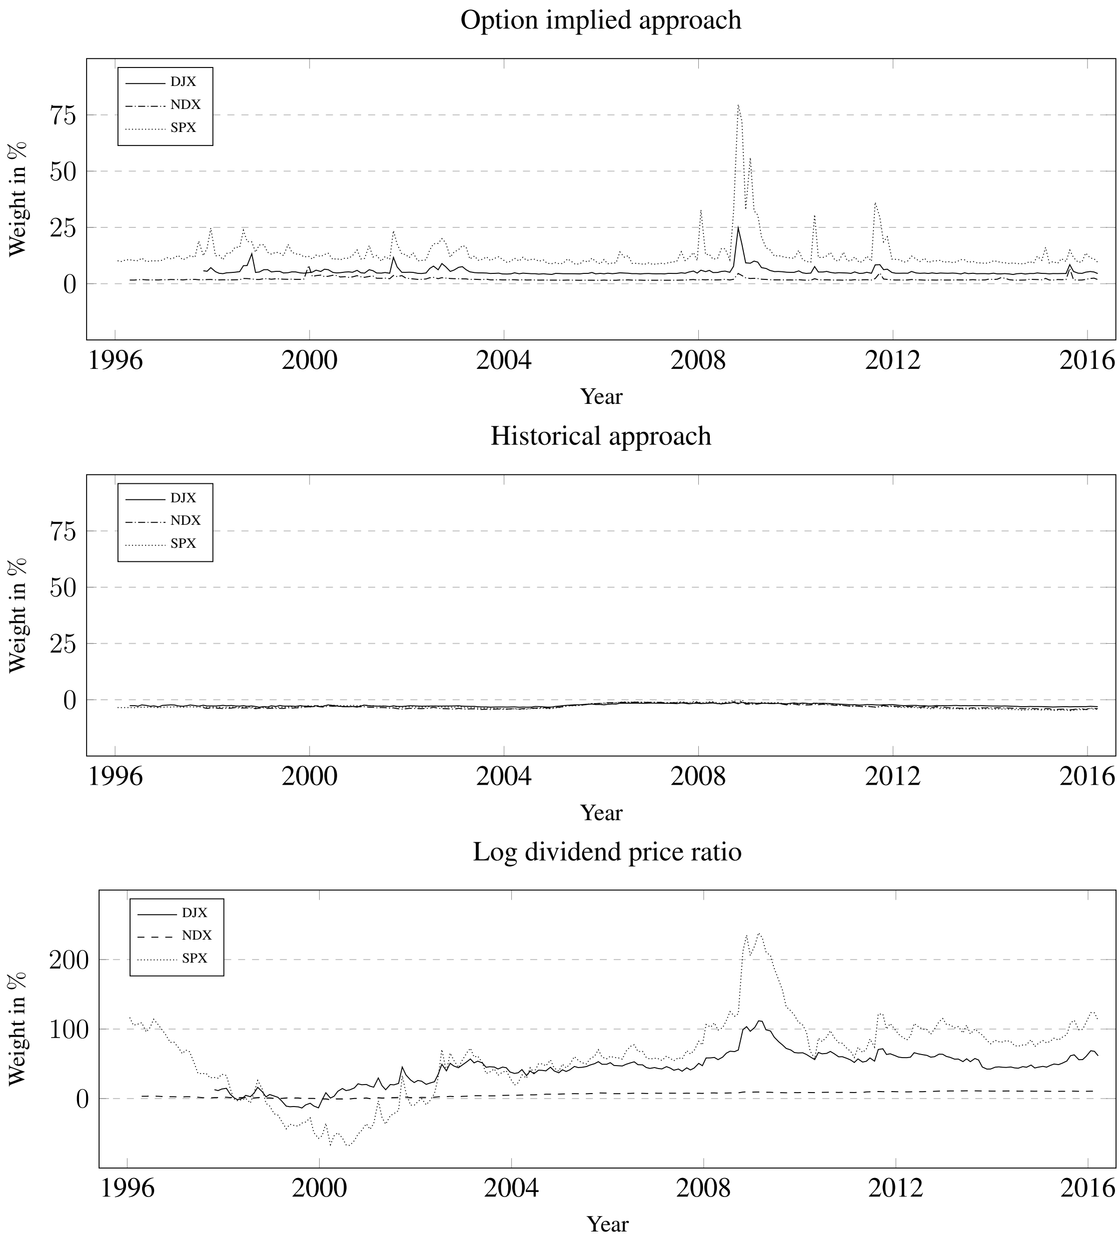

Figure 4.4: The figure displays the hedging demand following the optimal portfolio rule for the case $\gamma = 2$ and $\psi = 0.5$.

C.7 Hedging demand - classic estimation method

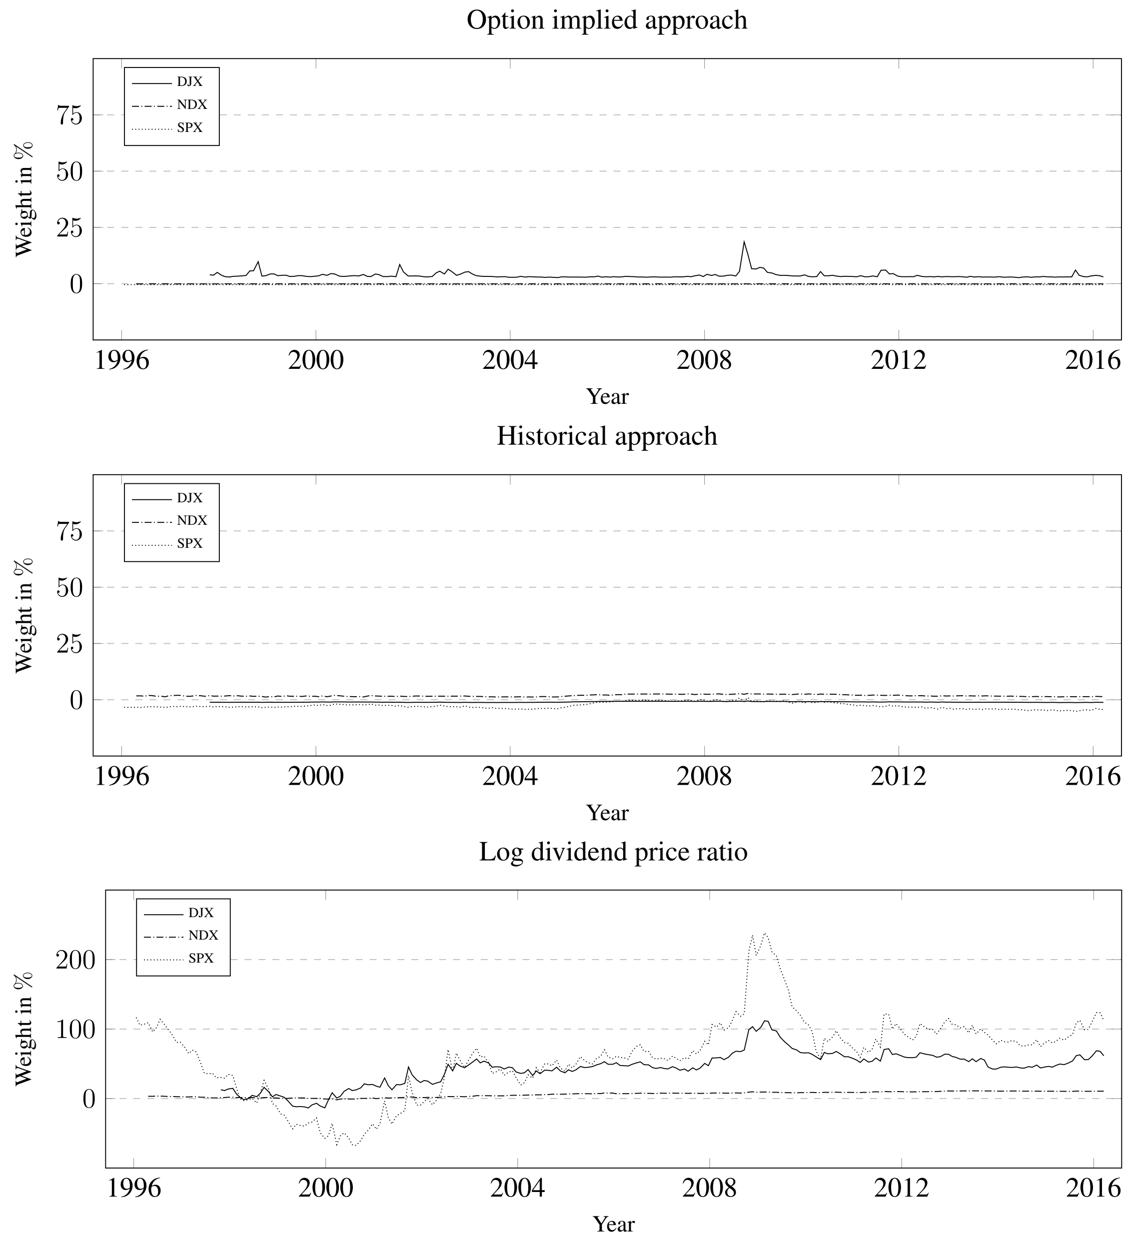

Figure 4.5: The figure displays the hedging demand following the optimal portfolio rule for the case $\gamma = 2$ and $\psi = 0.5$.

C.8 Consumption wealth ratio - direct estimation method

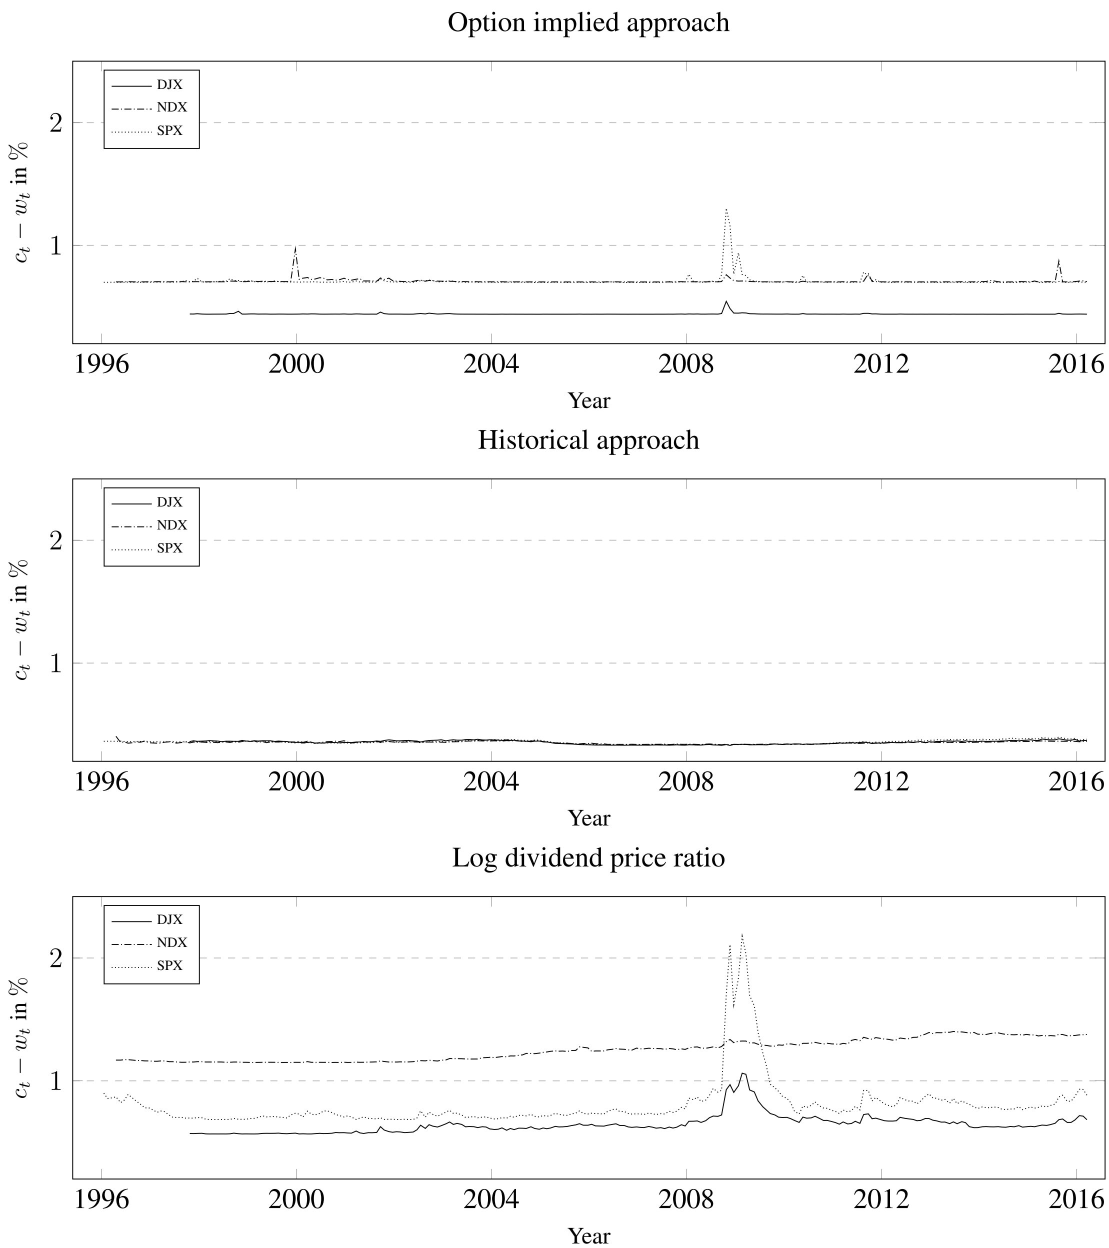

Figure 4.6: The figure displays the consumption wealth ratio following the optimal portfolio rule for the case $\gamma = 2$ and $\psi = 0.5$.

C.9 Consumption wealth ratio - classic estimation method

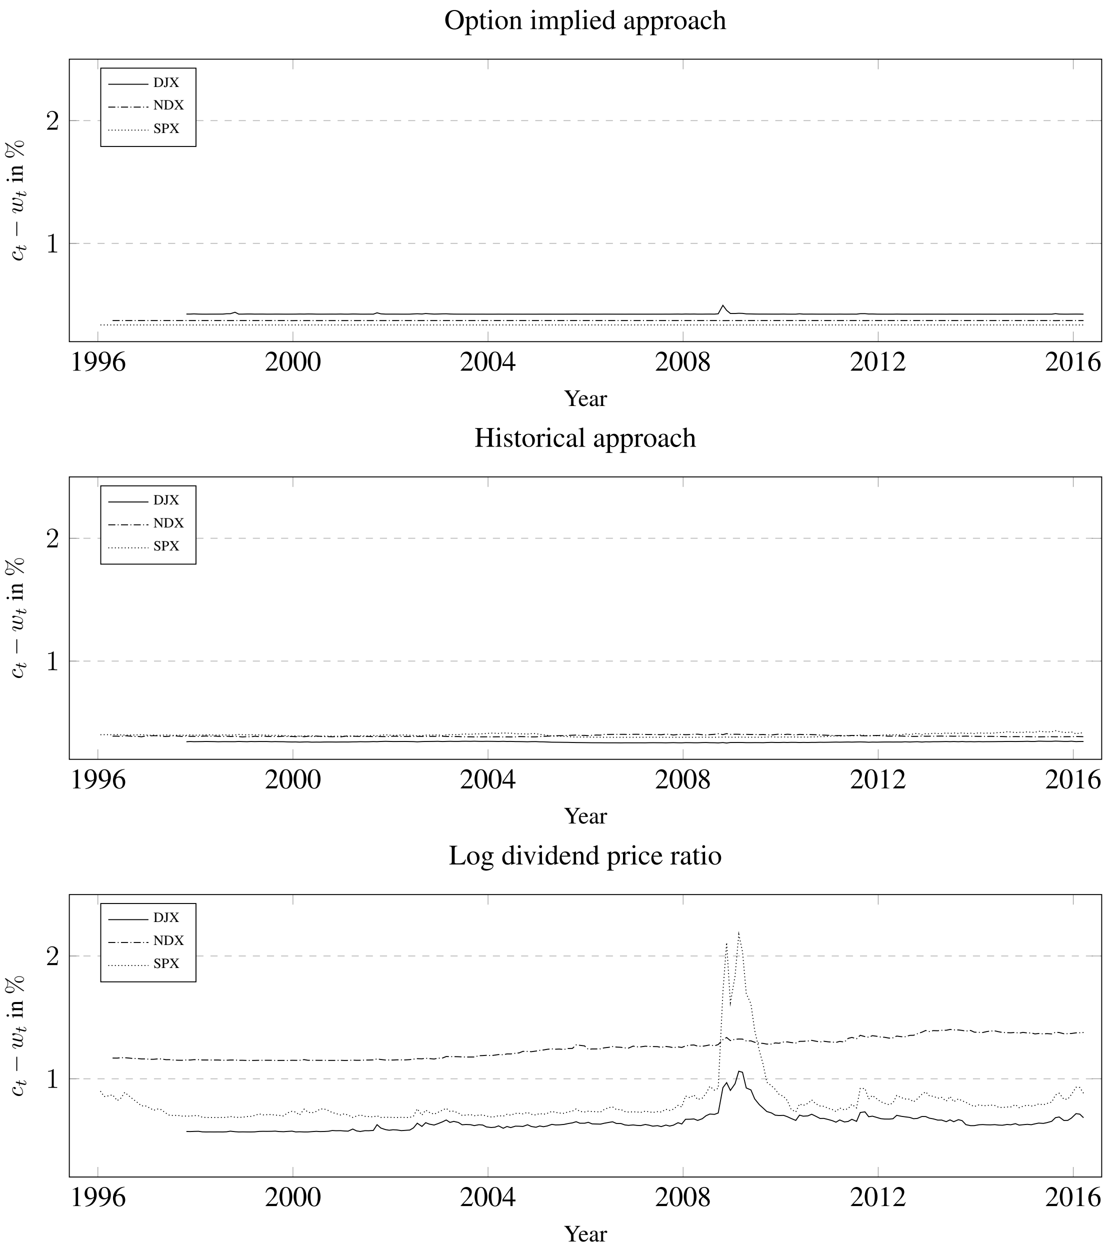

Figure 4.7: The figure displays the consumption wealth ratio following the optimal portfolio rule for the case $\gamma = 2$ and $\psi = 0.5$.

C.10 Portfolio weight - excursus

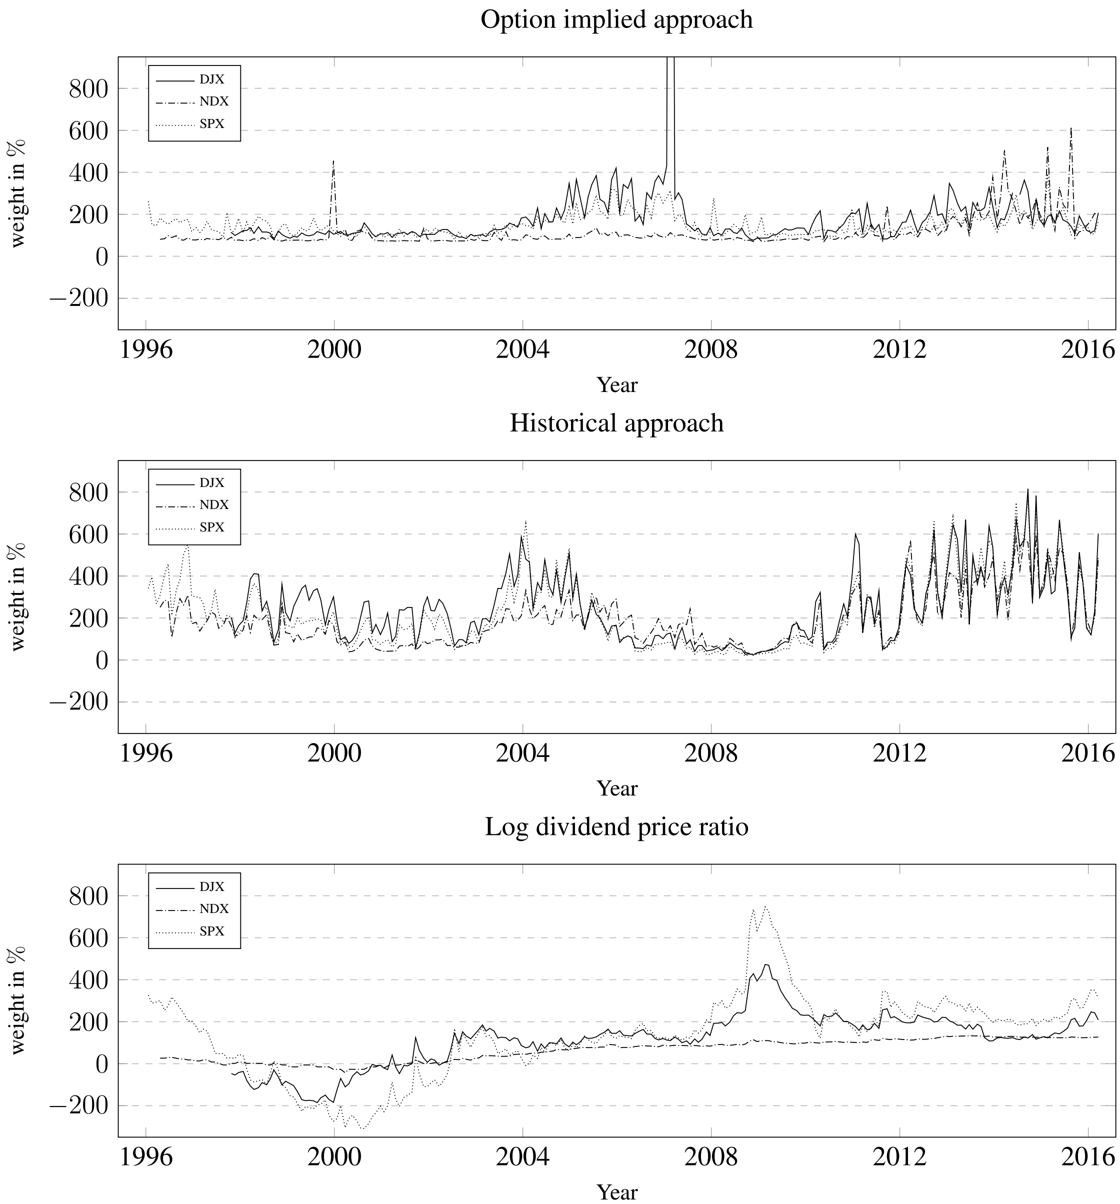

Figure 4.8: The figure displays the risky portfolio weight following the optimal portfolio rule with updated volatility σ_u for the case $\gamma = 2$ and $\psi = 0.5$.

C.11 Consumption wealth ratio - excursus

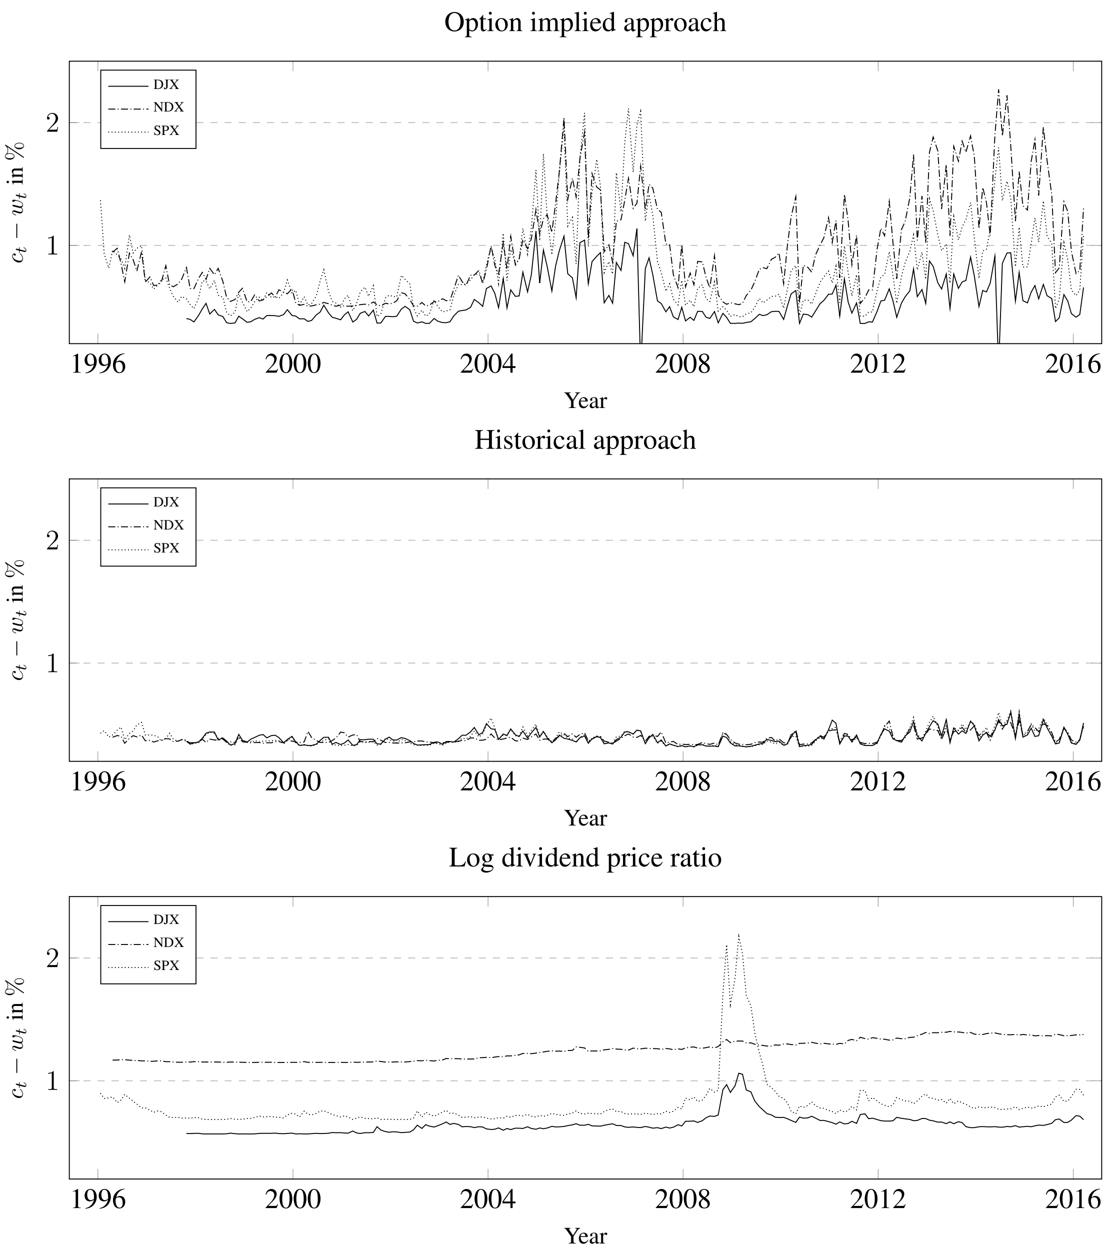

Figure 4.9: The figure displays the consumption wealth ratio following the optimal portfolio rule with updated volatility σ_u for the case $\gamma = 2$ and $\psi = 0.5$.

CHAPTER 5

Conclusion

This chapter summarises the main findings outlined in this thesis including chapter 2 dealing with risk-aversion measurement, chapter 3 on portfolio allocation and option implied state variables and chapter 4 on the developed direct estimation method. Further, an outlook of prospective research in the area of portfolio allocation using option prices and density forecasting is provided.

5.1 Summary

This thesis dealt with option implied information in economic applications. Thereby, the focus was on the use of option prices in the field of risk-aversion measurement and portfolio allocation. The performed research is motivated by the findings in the option pricing and density forecasting literature. These streams suggest that observed option prices imply an underlying asset price distribution that is non-normal. Further and more importantly, relying on a non-normal asset price distribution improves out-of-sample forecasting abilities when predicting the underlying asset price. Therefore, option prices are an appealing alternative to other sources of information (e.g. financial ratios) when trying to predict asset prices. Based on this finding, this thesis investigates within three chapters on the use of option implied information within risk-aversion measurement and portfolio allocation.

The first chapter contributes to the area of risk-aversion measurement using option prices. This field is rooted in the density forecasting literature, which relies on a transformation from risk-neutral to subjective dynamics. The transformation requires an assumption about the preferences of the representative investor. In this chapter, the investor follows a power-utility function. To transform the risk-neutral density into its subjective counterpart the risk-aversion parameter contained in the power-utility needs to be estimated. This transformation is normally performed taking into account a past series of risk-neutral densities that are linked to corresponding asset price outcomes to obtain an estimate of risk-aversion. In the density forecasting literature the risk-aversion estimate is many times a point estimate. Studies trying to evaluate those estimates suggest that they are relatively stable and within reasonable bounds. However, comparing the risk-aversion estimate over time using a sufficiently long rolling window unveils that risk-aversion varies strongly over time. To analyse and justify this observation, the first chapter proposes a simulation method to validate the obtained risk-aversion estimates and examines its variation. Contrary to other studies, the simulation aims to obtain current insights about risk-aversion rather than an average over a past time horizon. Therefore, only the current risk-neutral density is relevant to obtain risk-aversion estimates. The simulation makes assumptions about the underlying subjective asset price process to obtain risk-aversion estimates. Thereby, the market dynamics are assumed to follow a geometric Brownian motion and are further extended for stochastic volatility. The entirety of simulation outcomes allows to obtain a distribution of risk-aversion estimates and insight into its variation. The findings imply that the assumptions about the underlying risk-neutral distribution has a great impact on the variation of risk-aversion. More flexible distribution types reduce the variation. Nevertheless, the variation reported in this chapter is far higher than suggested by other studies (Bliss and Panigirtzoglou, 2004). Noteworthy, is that the structure of stochastic volatility has a strong impact on the distribution of risk-aversion and magnitude of variation. Lastly, the simulation is evaluated taking into account the out-of-sample performance. This is performed by taking into account the average risk-aversion from the simulation to transform the risk-neutral density. The obtained subjective density is then evaluated based on the log-likelihood of asset price outcomes. Comparing the proposed simulation with the classic method of estimating risk-aversion over a past time horizon suggest that forecasting abilities improve when using the proposed simulation method.

Central aspect within the third and fourth chapter is portfolio allocation. Both chapters investigate on how option implied information benefit strategic asset allocation between a risky and risk-free asset. The studied investor is infinitely lived with Epstein-Zin preferences going back to Campbell and Viceira (1999). The long-term investor makes his investment decisions based on a single state variable, the risk-premium. Central issue with the risk-premium is that it cannot be observed directly. Therefore, the literature derives the risk-premium process from state variables that are associated with predictive power towards the risk-premium. Typical examples for state variables are financial ratios (e.g. log dividend price ratio, price-earnings ratio).

The third chapter studies the impact when the financial ratio is replaced with information derived from option prices. The use of option prices in portfolio allocation is motivated by three aspects: First, option prices are forward looking since their pricing contains in theory the current market expectations until expiry. Second, the cross-section in the strike price allows to gain detailed insight about different states of the underlying asset. Third, option prices rely on currently observed prices and do not require a history of options prices as it is the case for financial ratios[1]. This chapter proposes to exploit the option implied risk-premium as state variable. To derive the option implied risk-premium a battery of different distributional assumptions are employed, including a lognormal density (Black, 1976), a mixture-lognormal density (Ritchey, 1990) and a non-parametric binomial tree (Jackwerth and Rubinstein, 1996). As benchmark state variables, the option implied market price of risk as in Kostakis et al. (2011) and the log dividend price ratio as proposed in Campbell and Viceira (1999) are applied. Furthermore, as alternative method to derive the proposed state variables a historical benchmark is implemented. It makes assumptions about the return and volatility process using a GARCH type model and is estimated using a series of asset price returns. Next to this benchmark the results are also compared against a simple investment into the risky asset. The findings suggest that there are no major differences among the proposed density types. This is most likely a consequence of only using the first two moments within the suggested state variables. Comparing the option-implied approach to the alternative benchmarks unveils that option implied information provide improved timing abilities towards the historical approach. Comparing it towards the simple investment shows that during the period after the crisis the option implied approach is preferable. However, during the crisis this is not the case. Analysing the market price of risk shows that the risk-premium is preferable within the considered portfolio allocation. Nonetheless, comparing the risk-premium to the log dividend price ratio it is not clear which state variable is in favour.

The fourth chapter is motivated by the findings in chapter three and develops and estimation framework. Starting point are the empirical findings of the option implied risk-premium. Chapter three suggests that the option implied risk-premium is a good state variable in regards of the applied performance measures to describe the risk-premium process of the portfolio allocation. However, the applied estimation framework of the risk-premium process goes back to Campbell and Viceira (1999) and derives its parameters using a restricted VAR(1). This is necessary since they originally worked with financial

[1] Indeed the transformation from risk-neutral to subjective density relies on a past time horizon but within the assumption of a power utility risk-aversion is constant. Therefore only the observed cross-section of option prices, which essentially dictates the characteristics of the risk-neutral density, is of relevance.

ratios which do not give direct insight into the risk-premium. Therefore, this chapter proposes a direct estimation method to obtain consistent estimates of the risk-premium process using the option implied risk-premium. The risk-premium is obtained as in Martin (2017). This procedure is model free. Main benefit is that it avoids to appoint a fixed horizon to transform the risk-neutral to subjective density. The empirical findings suggest that the proposed direct estimation method obtains consistent parameters across different risky assets which is not the case for the estimation method using the restricted VAR(1). The parameters are in line with the empirically relevant case described in Campbell and Viceira (1999) and exhibit a positive hedging demand. These observations are supported by the improvements in the value function and out-of-sample performance. Furthermore, a striking argument to use the option-implied risk-premium is that it is not conflicted with non-stationarity issues as it is the case for financial ratios.

5.2 Future Work

In the near future, the work on portfolio allocation and option implied information should be continued. The infinitely lived investor by Campbell and Viceira (1999) operates within a quite restrictive set-up. Most apparent is the assumption of constant volatility. A challenging task is to incorporate option-implied volatility. The starting point would be to design a suitable framework to capture the dynamics of option-implied volatility following Chacko and Viceira (2005). Furthermore, some work needs to be dedicated to design a portfolio allocation that is more tailored towards the use of option implied information rather than financial ratios. It can be expected that the portfolio and consumption decisions in this richer set-up will strongly diverge from the findings. Some preliminary work already indicates these expectations as suggested in section 4.7. Furthermore, allowing for a wider set of available securities, particularly to hedge positions within the portfolio, is an appealing route. Thereby, the focus should not be on stock selection but rather on the availability of derivative securities like options and futures or short selling.

The second area of work is to introduce machine learning algorithms to option pricing and density forecasting. Even though machine learning is on a more early stage within economics it has potentially interesting applications within option pricing and density forecasting. An interesting fact is the great flexibility within the machine learning algorithms which generally is of advantage when working with option implied distributions. Thereby, an appealing route are short term or high frequency predictions using option implied densities since they tend to perform worse compared to historical densities. Machine learning could draw another perspective into short term predictions and help to understand the weaker forecasting abilities of option implied densities. Particularly, interesting within this context is the analysis of high frequency option prices.

References

Abken, P. A., Madan, D. B., and Ramamurtie, S. (1996a). Estimation of risk-neutral and statistical densities by hermite polynomial approximation: With an application to eurodollar futures options. *Working Paper*, Federal Reserve Bank of Atlanta.

Abken, P. A., Madan, D. B., and Ramamurtie, S. (1996b). Pricing s&p 500 index options using a hilbert space basis. *Working Paper*, Federal Reserve Bank of Atlanta.

Aït-Sahalia, Y. and Brandt, M. W. (2001). Variable selection for portfolio choice. *The Journal of Finance*, 56(4):1297–1351.

Aït-Sahalia, Y. and Lo, A. (2000). Nonparametric risk management and implied risk aversion. *Journal of Econometrics*, 94(1):9–51.

Aït-Sahalia, Y., Wang, Y., and Yared, F. (2001). Do option markets correctly price the probabilities of movement of the underlying asset? *Journal of Econometrics*, 102(1):67–110.

Amisano, G. and Giacomini, R. (2007). Comparing density forecasts via weighted likelihood ratio tests. *Journal of Business & Economic Statistics*, 25(2):177–190.

Anagnou-Basioudis, I., Bedendo, M., Hodges, S. D., and Tompkins, R. (2005). Forecasting accuracy of implied and garch-based probability density functions. *Review of Futures Markets*, 11(1):41–66.

Ang, A. and Bekaert, G. (2007). Stock return predictability: Is it there? *The Review of Financial Studies*, 20(3):651–707.

Aparicio, S. D. and Hodges, S. (1998). Implied risk-neutral distribution: A comparison of estimation methods. *Working Paper*, Warwick University.

Arrow, K. J. (1971). Essays in the theory of risk-bearing. *North Holland, Amsterdam*.

Aït-Sahalia, Y. and Lo, A. W. (1998). Nonparametric estimation of state-price densities implicit in financial asset prices. *The Journal of Finance*, 53(2):499–547.

Audrino, F. and Meier, P. (2012). Empirical pricing kernel estimation using a functional gradient descent algorithm based on splines. *Working Paper No. 2012-10, University of*

St. Gallen.

Bakshi, G., Kapadia, N., and Madan, D. (2003). Stock return characteristics, skew laws, and the differential pricing of individual equity options. *The Review of Financial Studies*, 16(1):101–143.

Bakshi, G., Madan, D., and Panayotov, G. (2010). Returns of claims on the upside and the viability of u-shaped pricing kernels. *Journal of Financial Economics*, 97(1):130–154.

Balduzzi, P. and Lynch, A. W. (1999). Transaction costs and predictability: some utility cost calculations. *Journal of Financial Economics*, 52(1):47–78.

Balduzzi, P. and Lynch, A. W. (2000). Predictability and transaction costs: The impact on rebalancing rules and behavior. *The Journal of Finance*, 55(5):2285–2309.

Bao, Y., Lee, T.-H., and Saltoğlu, B. (2007). Comparing density forecast models. *Journal of Forecasting*, 26(3):203–225.

Barberis, N. (2000). Investing for the long run when returns are predictable. *The Journal of Finance*, 55(1):225–264.

Barle, S. and Cakici, N. (1998). How to grow a smiling tree. *Journal of Finanical Engineering*, 7(2):127–146.

Bates, D. S. (1991). The crash of '87: Was it expected? the evidence from options markets. *The Journal of Finance*, 46(3):1009–1044.

Bates, D. S. (1996). Jumps and stochastic volatility: Exchange rate processes implicit in deutsche mark options. *The Review of Financial Studies*, 9(1):69–107.

Bates, D. S. (2000). Post-'87 crash fears in the s&p 500 futures option market. *Journal of Econometrics*, 94(1):181–238.

Bekaert, G. and Hoerova, M. (2014). The vix, the variance premium and stock market volatility. *Journal of Econometrics*, 183(2):181–192.

Berkowitz, J. (2001). Testing density forecasts, with applications to risk management. *Journal of Business & Economic Statistics*, 19(4):465–474.

Black, F. (1976). The pricing of commodity contracts. *Journal of Financial Economics*, 3(1):167–179.

Black, F. and Scholes, M. (1973). The pricing of options and corporate liabilities. *Journal of Political Economy*, 81(3):637–654.

Blair, B. J., Poon, S.-H., and Taylor, S. J. (2001). Modelling S&P 100 volatility: The information content of stock returns. *Journal of Banking & Finance*, 25(9):1665–1679.

Bliss, R. R. and Panigirtzoglou, N. (2002). Testing the stability of implied probability density functions. *Journal of Banking & Finance*, 26(2):381–422.

Bliss, R. R. and Panigirtzoglou, N. (2004). Option-implied risk aversion estimates. *Journal of Finance*, 59(1):407–446.

Bollerslev, T. (1986). Generalized autoregressive conditional heteroskedasticity. *Journal of Econometrics*, 31(3):307–327.

Bookstaber, R. M. and McDonald, J. B. (1987). A general distribution for describing security price returns. *Journal of Business*, 60(3):401–424.

Bossaerts, P. and Hillion, P. (1999). Implementing statistical criteria to select return forecasting models: What do we learn? *The Review of Financial Studies*, 12(2):405–428.

Breeden, D. T. (1979). An intertemporal asset pricing model with stochastic consumption and investment opportunities. *Journal of Financial Economics*, 7(3):265–296.

Breeden, D. T. and Litzenberger, R. H. (1978). Prices of state-contingent claims implicit in option prices. *The Journal of Business*, 51(4):621–651.

Brennan, M. J., Schwartz, E. S., and Lagnado, R. (1997). Strategic asset allocation. *Journal of Economic Dynamics and Control*, 21(8):1377–1403.

Brenner, M. and Eom, Y. (1997). No-arbitrage option pricing: New evidence on the validity of the martingale property. *Working Paper*, New York University.

Brigo, D. and Mercurio, F. (2002). Lognormal-mixture dynamics and calibration to market volatility smiles. *International Journal of Theoretical and Applied Finance*, 5(4):427–446.

Britten-Jones, M. and Neuberger, A. (2000). Option prices, implied price processes, and stochastic volatility. *The Journal of Finance*, 55(2):839–866.

Brown, G. and Toft, K. B. (1999). Constructing binomial trees from multiple implied probability distributions. *The Journal of Derivatives*, 7(2):83–100.

Buchen, P. W. and Kelly, M. (1996). The maximum entropy distribution of an asset inferred from option prices. *The Journal of Financial and Quantitative Analysis*, 31(1):143–159.

Campa, J. M., Chang, P., and Reider, R. L. (1998). Implied exchange rate distributions: evidence from otc option markets. *Journal of International Money and Finance*, 17(1):117–160.

Campbell, J. Y. (1987). Stock returns and the term structure. *Journal of Financial Economics*, 18(2):373–399.

Campbell, J. Y. (1993). Intertemporal asset pricing without consumption data. *The American Economic Review*, 83(3):487–512.

Campbell, J. Y., Chacko, G., Rodriguez, J., and Viceira, L. M. (2004). Strategic asset allocation in a continuous-time var model. *Journal of Economic Dynamics and Control*, 28(11):2195–2214.

Campbell, J. Y., Chan, Y. L., and Viceira, L. M. (2003). A multivariate model of strategic asset allocation. *Journal of Financial Economics*, 67(1):41–80.

Campbell, J. Y. and Shiller, R. J. (1988a). The dividend-price ratio and expectations of future dividends and discount factors. *The Review of Financial Studies*, 1(3):195–228.

Campbell, J. Y. and Shiller, R. J. (1988b). Stock prices, earnings, and expected dividends. *The Journal of Finance*, 43(3):661–676.

Campbell, J. Y. and Shiller, R. J. (1991). Yield spreads and interest rate movements: A bird's eye view. *The Review of Economic Studies*, 58(3):495–514.

Campbell, J. Y. and Viceira, L. M. (1999). Consumption and portfolio decisions when expected returns are time varying. *The Quarterly Journal of Economics*, 114(2):433–495.

Campbell, J. Y. and Viceira, L. M. (2000). Consumption and portfolio decisions when expected returns are time varying: Erratum. *Manuscript Harvard University*.

Chacko, G. and Viceira, L. M. (2005). Dynamic consumption and portfolio choice with stochastic volatility in incomplete markets. *The Review of Financial Studies*, 18(4):1369–1402.

Chernov, M. and Ghysels, E. (2000). A study towards a unified approach to the joint estimation of objective and risk neutral measures for the purpose of options valuation. *Journal of Financial Economics*, 56(3):407–458.

Christoffersen, P., Heston, S., and Jacobs, K. (2009). The shape and term structure of the index option smirk: Why multifactor stochastic volatility models work so well. *Management Science*, 55(12):1914–1932.

Christoffersen, P., Heston, S., and Jacobs, K. (2013). Capturing option anomalies with a variance-dependent pricing kernel. *The Review of Financial Studies*, 26(8):1963–2006.

Cochrane, J. H. (1992). Explaining the variance of price–dividend ratios. *The Review of Financial Studies*, 5(2):243–280.

Cochrane, J. H. and Hansen, L. P. (1992). Asset pricing explorations for macroeconomics. Working Paper 4088, National Bureau of Economic Research.

Corrado, C. J. and Su, T. (1996). Skewness and kurtosis in s&p 500 index returns implied by option prices. *Journal of Financial Research*, 19(2):175–192.

Corrado, C. J. and Su, T. (1997). Implied volatility skews and stock index skewness and kurtosis implied by s&p 500 index option prices. *The Journal of Derivatives*, 4(4):8–19.

Cox, J. C. and Ross, S. A. (1976). The valuation of options for alternative stochastic processes. *Journal of Financial Economics*, 3(1):145–166.

Cox, J. C., Ross, S. A., and Rubinstein, M. (1979). Option pricing: A simplified approach. *Journal of Financial Economics*, 7(3):229–263.

DeMiguel, V., Nogales, F. J., and Uppal, R. (2014). Stock return serial dependence and out-of-sample portfolio performance. *The Review of Financial Studies*, 27(4):1031–1073.

DeMiguel, V., Plyakha, Y., Uppal, R., and Vilkov, G. (2013). Improving portfolio selection using option-implied volatility and skewness. *Journal of Financial and Quantitative Analysis*, 48(6):1813–1845.

Derman, E. and Kani, I. (1994). Riding on a smile. *Risk*, 7(2):127–146.

Dupire, B. (1994). Pricing with a smile. *Risk*, 7(1):18–20.

Engle, R. (1982). Autoregressive conditional heteroscedasticity with estimates of the variance of united kingdom inflation. *Econometrica*, 50(4):987–1007.

Epstein, L. G. and Zin, S. E. (1989). Substitution, risk aversion, and the temporal behavior of consumption and asset returns: A theoretical framework. *Econometrica*, 57(4):937–969.

Epstein, L. G. and Zin, S. E. (1991). Substitution, risk aversion, and the temporal behavior of consumption and asset returns: An empirical analysis. *Journal of Political Economy*, 99(2):263–286.

Fackler, P. L. and King, R. P. (1990). Calibration of option-based probability assessments in agricultural commodity markets. *American Journal of Agricultural Economics*, 72(1):73–83.

Fama, E. (1970). Multiperiod consumption-investment decisions. *American Economic Review*, 60(1):163–174.

Fama, E. F. (1984). The information in the term structure. *Journal of Financial Economics*, 13(4):509–528.

Fama, E. F. and French, K. R. (1988). Dividend yields and expected stock returns. *Jour-*

nal of Financial Economics, 22(1):3–25.

Fama, E. F. and French, K. R. (1989). Business conditions and expected returns on stocks and bonds. *Journal of Financial Economics*, 25(1):23–49.

Fama, E. F. and Schwert, G. (1977a). Asset returns and inflation. *Journal of Financial Economics*, 5(2):115–146.

Fama, E. F. and Schwert, G. (1977b). Human capital and capital market equilibrium. *Journal of Financial Economics*, 4(1):95–125.

Ferson, W. E. and Constantinides, G. (1991). Habit persistence and durability in aggregate consumption: Empirical tests. *Journal of Financial Economics*, 29(2):199–240.

Ferson, W. E., Sarkissian, S., and Simin, T. (2003). Spurious regressions in financial economics? *Journal of Finance*, 58(4):1393–1413.

Friend, I. and Blume, M. E. (1975). The demand for risky assets. *The American Economic Review*, 65(5):900–922.

Glosten, L. R., Jagannathan, R., and Runkle, D. E. (1993). On the relation between the expected value and the volatility of the nominal excess return on stocks. *Journal of Finance*, 48(5):1779–1801.

Gonzalo, J. and Pitarakis, J.-Y. (2017). Inferring the predictability induced by a persistent regressor in a predictive threshold model. *Journal of Business & Economic Statistics*, 35(2):202–217.

Goyal, A. and Welch, I. (2003). Predicting the equity premium with dividend ratios. *Management Science*, 49(5):639–654.

Grossman, S. J. and Shiller, R. J. (1981). The determinants of the variability of stock market prices. *The American Economic Review*, 71(2):222–227.

Guo, H. and Whitelaw, R. (2001). Risk and return: Some new evidence. *Working Paper No. 2001-001A, Federal Reserve Bank of St. Louis*.

Hansen, L. P. and Singleton, K. J. (1982). Generalized instrumental variables estimation of nonlinear rational expectations models. *Econometrica*, 50(5):1269–1286.

Hansen, L. P. and Singleton, K. J. (1983). Stochastic consumption, risk aversion, and the temporal behavior of asset returns. *Journal of Political Economy*, 91(2):249–265.

Heston, S. L. (1993). A closed-form solution for options with stochastic volatility with applications to bond and currency options. *The Review of Financial Studies*, 6(2):327–343.

Hodrick, R. J. (1992). Dividend yields and expected stock returns: Alternative procedures for inference and measurement. *The Review of Financial Studies*, 5(3):357–386.

Hull, J. (2012). *Options, futures, and other derivatives*. Upper Saddle River, N.J: Pearson/Prentice Hall, 8. edition.

Jackwerth, J. C. (1997). Generalized binomial trees. *The Journal of Derivatives*, 5(2):7–17.

Jackwerth, J. C. (1999). Option implied risk-neutral distributions and implied binomial trees: A literature review. *Journal of Derivatives*, 7(2):66–82.

Jackwerth, J. C. (2000). Recovering risk aversion from option prices and realized returns. *The Review of Financial Studies*, 13(2):433–451.

Jackwerth, J. C. and Rubinstein, M. (1996). Recovering probability distributions from option prices. *The Journal of Finance*, 51(5):1611–1631.

Jarrow, R. and Rudd, A. (1982). Approximate option valuation for arbitrary stochastic processes. *Journal of Financial Economics*, 10(3):347–369.

Jiang, G. J. and Tian, Y. S. (2005). The model-free implied volatility and its information content. *The Review of Financial Studies*, 18(4):1305–1342.

Johnson, N. L. (1949). Systems of frequency curves generated by methods of translation. *Biometrika*, 36(1):149–176.

Jondeau, E. and Rockinger, M. (2001). Gram–charlier densities. *Journal of Economic Dynamics and Control*, 25(9):1457–1483.

Jondeau, E. and Rockinger, M. (2006). Optimal portfolio allocation under higher moments. *European Financial Management*, 12(1):29–55.

Jorion, P. and Giovannini, A. (1993). Time-series tests of a non-expected-utility model of asset pricing. *European Economic Review*, 37(5):1083–1100.

Kang, B. J. and Kim, T. S. (2006). Option-implied risk preferences: An extension to wider classes of utility functions. *Journal of Financial Markets*, 9(2):180–198.

Kempf, A., Korn, O., and Saßning, S. (2015). Portfolio optimization using forward-looking information. *Review of Finance*, 19(1):467–490.

Kim, T. S. and Omberg, E. (1996). Dynamic nonmyopic portfolio behavior. *Review of Financial Studies*, 9(1):141–161.

Kliger, D. and Levy, O. (2002). Risk preferences heterogeneity: Evidence from asset markets. *Review of Finance*, 6(3):277–290.

Kostakis, A., Magdalinos, T., and Stamatogiannis, M. P. (2015). Robust econometric inference for stock return predictability. *The Review of Financial Studies*, 28(5):1506–1553.

Kostakis, A., Panigirtzoglou, N., and Skiadopoulos, G. (2011). Market timing with option-implied distributions: A forward-looking approach. *Management Science*, 57(7):1231–1249.

Ledoit, O. and Wolf, M. (2008). Robust performance hypothesis testing with the sharpe ratio. *Journal of Empirical Finance*, 15(5):850–859.

Linn, M., Shive, S., and Shumway, T. (2017). Pricing Kernel Monotonicity and Conditional Information. *The Review of Financial Studies*, 31(2):493–531.

Liu, J. (2007). Portfolio Selection in Stochastic Environments. *The Review of Financial Studies*, 20(1):1–39.

Liu, X., Shackleton, M. B., Taylor, S. J., and Xu, X. (2007). Closed-form transformations from risk-neutral to real-world distributions. *Journal of Banking & Finance*, 31(5):1501–

1520.

Longstaff, F. A. (1995). Option pricing and the martingale restriction. *The Review of Financial Studies*, 8(4):1091–1124.

Lucas, R. E. (1978). Asset prices in an exchange economy. *Econometrica*, 46(6):1429–1445.

Malz, A. M. (1996). Using option prices to estimate realignment probabilities in the european monetary system: the case of sterling-mark. *Journal of International Money and Finance*, 15(5):717–748.

Malz, A. M. (1997). Estimating the probability distribution of the future exchange rate from option prices. *The Journal of Derivatives*, 5(2):18–36.

Markowitz, H. (1952). Portfolio selection. *The Journal of Finance*, 7(1):77–91.

Martin, I. (2017). What is the expected return on the market? *The Quarterly Journal of Economics*, 132(1):367–433.

Mayers, D. (1976). Nonmarketable assets, market segmentation, and the level of asset prices. *The Journal of Financial and Quantitative Analysis*, 11(1):1–12.

Mehra, R. and Prescott, E. C. (1985). The equity premium: A puzzle. *Journal of Monetary Economics*, 15(2):145–161.

Melick, W. R. and Thomas, C. P. (1997). Recovering an asset's implied pdf from option prices: An application to crude oil during the gulf crisis. *Journal of Financial & Quantitative Analysis*, 32(1):91–115.

Merton, R. C. (1969). Lifetime portfolio selection under uncertainty: The continuous-time case. *The Review of Economics and Statistics*, 51(3):247–257.

Merton, R. C. (1971). Optimum consumption and portfolio rules in a continuous-time model. *Journal of Economic Theory*, 3(4):373–413.

Merton, R. C. (1973). An intertemporal capital asset pricing model. *Econometrica*, 41(5):867–887.

Merton, R. C. (1976). Option pricing when underlying stock returns are discontinuous. *Journal of Financial Economics*, 3(1):125–144.

Merton, R. C. (1980). On estimating the expected return on the market. *Journal of Financial Economics*, 8(4):323–361.

Modigliani, F. and Sutch, R. (1966). Innovations in interest rate policy. *The American Economic Review*, 56(1):178–197.

Mossin, J. (1968). Optimal multiperiod portfolio policies. *The Journal of Business*, 41(2):215–229.

Nelson, C. R. and Kim, M. J. (1993). Predictable stock returns: The role of small sample bias. *The Journal of Finance*, 48(2):641–661.

Normandin, M. and St-Amour, P. (1998). Substitution, risk aversion, taste shocks and equity premia. *Journal of Applied Econometrics*, 13(3):265–281.

Pan, J. (2002). The jump-risk premia implicit in options: Evidence from an integrated time-series study. *Journal of Financial Economics*, 63(1):3–50.

Paye, B. S. and Timmermann, A. (2006). Instability of return prediction models. *Journal of Empirical Finance*, 13(3):274–315.

Posner, S. E. and Milevsky, M. A. (1998). Valuing exotic options by approximating the spd with higher moments. *The Journal of Financial Engineering*, 7(2):109–125.

Pérignon, C. and Villa, C. (2002). Extracting information from options markets: Smiles, state–price densities and risk aversion. *European Financial Management*, 8(4):495–513.

Pritsker, M. (1998). Nonparametric density estimation and tests of continuous time interest rate models. *The Review of Financial Studies*, 11(3):449–487.

Ritchey, R. J. (1990). Call option valuation for discrete normal mixtures. *Journal of Financial Research*, 13(4):285–296.

Rockinger, M. and Abadir, K. (1997). Density-embedding functions. *Working Paper*, HEC.

Rookley, C. (1997). Fully exploiting the information content of intra day option quotes: Applications in option pricing and risk management. *Working Paper*, University of Arizona.

Rosenberg, J. V. (1998). Pricing multivariate contingent claims using estimated risk–neutral density functions. *Journal of International Money and Finance*, 17(2):229–247.

Rosenberg, J. V. and Engle, R. F. (1997). Option hedging using empirical pricing kernels. *Working Paper*, NBER.

Rosenberg, J. V. and Engle, R. F. (2002). Empirical pricing kernels. *Journal of Financial Economics*, 64(3):341–372.

Rubinstein, M. (1976a). The strong case for the generalized logarithmic utility model as the premier model of financial markets. *The Journal of Finance*, 31(2):551–571.

Rubinstein, M. (1976b). The valuation of uncertain income streams and the pricing of options. *The Bell Journal of Economics*, 7(2):407–425.

Rubinstein, M. (1985). Nonparametric tests of alternative option pricing models using all reported trades and quotes on the 30 most active cboe option classes from august 23, 1976 through august 31, 1978. *Journal of Finance*, 40(2):455–480.

Rubinstein, M. (1994). Implied binomial trees. *The Journal of Finance*, 49(3):771–818.

Rubinstein, M. (1998). Edgeworth binomial trees. *The Journal of Derivatives*, 5(3):20–27.

Samuelson, P. A. (1963). Risk and uncertainty: A fallacy of large numbers. *Scientia*, 57(98):108.

Samuelson, P. A. (1969). Lifetime portfolio selection by dynamic stochastic programming. *The Review of Economics and Statistics*, 51(3):239–246.

Shackleton, M. B., Taylor, S. J., and Yu, P. (2010). A multi-horizon comparison of density forecasts for the s&p 500 using index returns and option prices. *Journal of Banking &*

Finance, 34(11):2678–2693.

Sherrick, B. J., Garcia, P., and Tirupattur, V. (1996a). Recovering probabilistic information from option markets: Tests of distributional assumptions. *Journal of Futures Markets*, 16(5):545–560.

Sherrick, B. J., Irwin, S. H., and Forster, D. L. (1992). Option-based evidence of the nonstationarity of expected s&p 500 futures price distributions. *Journal of Futures Markets*, 12(3):275–290.

Sherrick, B. J., Irwin, S. H., and Forster, D. L. (1996b). An examination of option-implied s&p 500 futures price distributions. *Financial Review*, 31(3):667–694.

Shiller, R. J. (1982). Consumption, asset markets and macroeconomic fluctuations. *Carnegie-Rochester Conference Series on Public Policy*, 17:203–238.

Shimko, D. (1993). Bounds of probability. *Risk*, 6(4):33–36.

Stambaugh, R. F. (1999). Predictive regressions. *Journal of Financial Economics*, 54(3):375–421.

Stiglitz, J. E. (1970). A consumption-oriented theory of the demand for financial assets and the term structure of interest rates1. *The Review of Economic Studies*, 37(3):321–351.

Stutzer, M. (1996). A simple nonparametric approach to derivative security valuation. *The Journal of Finance*, 51(5):1633–1652.

Tobin, J. (1958). Liquidity preference as behavior towards risk. *The Review of Economic Studies*, 25(2):65–86.

Valkanov, R. (2003). Long-horizon regressions: theoretical results and applications. *Journal of Financial Economics*, 68(2):201–232.

Wachter, J. A. (2002). Portfolio and consumption decisions under mean-reverting returns: An exact solution for complete markets. *Journal of Financial and Quantitative Analysis*, 37(1):63–91.

Wang, Y.-H. (2009). The impact of jump dynamics on the predictive power of option-implied densities. *The Journal of Derivatives*, 16(3):9–22.

Weil, P. (1989). The equity premium puzzle and the risk-free rate puzzle. *Journal of Monetary Economics*, 24(3):401–421.

Weinberg, S. A. (2001). Interpreting the volatility smile: An examination of the information content of option prices. *Discussion Paper*, FRB International Finance.

Ziegler, A. (2007). Why does implied risk aversion smile? *The Review of Financial Studies*, 20(3):859–904.