Optimal portfolio allocation using option implied information

Maria Kyriacou * Jose Olmo †,‡ Marius Strittmatter§

October 28, 2020

Abstract

This paper explores option-implied information measures for optimal portfolio allocation. We introduce two state variables constructed from option prices. The first state variable is the risk premium on the risky asset and the second variable is the market price of risk. We also explore a lognormal distribution, a mixture of lognormal distributions and a binomial tree for constructing the implied risk neutral density function. Using a combination of statistical and economic measures applied to a portfolio given by the one-month US Treasury bill and the S&P 500 Index we show the good performance of option-implied information measures for optimal portfolio allocation.

Keywords: strategic asset allocation, option implied information, risk-neutral density function, ARCH models, performance evaluation

*Economics Department, University of Southampton, University Rd., Southampton, SO17 1BJ, United Kingdom. E-mail: M.Kyriacou@soton.ac.uk

†Corresponding author: Department of Economic Analysis, Universidad de Zaragoza, Gran Vía 2, 50005. Zaragoza, Spain. Phone: +34 876 55 4682. E-mail: joseolmo@unizar.es

‡Department of Economics, University of Southampton. University Rd., Southampton, SO17 1BJ, United Kingdom. Jose Olmo acknowledges financial support from Fundación Agencia Aragonesa para la Investigación y el Desarrollo (ARAID).

§Economics Department, University of Southampton, University Rd., Southampton, SO17 1BJ, United Kingdom. E-mail: ms2g15@soton.ac.uk

1 Introduction

Option prices contain market expectations about the distribution of future asset prices. Investors can take advantage of these expectations as valuable input for the purposes of volatility modelling as in Bollerslev et al. (2011), risk management as in Cao et al. (2010), and, more recently, for optimal asset allocation, see the recent contributions by Kostakis et al. (2011), DeMiguel et al. (2013) and Kempf et al. (2015). Market expectations are derived from the distribution function of asset returns. Breeden and Litzenberger (1978) propose a convenient closed-form expression for obtaining such distribution function implied from option prices under the assumption of risk-neutral investors. This risk-neutral distribution function can be retrieved from option prices on the underlying asset. Importantly, the literature on financial derivatives has developed novel methods for obtaining the counterpart objective distribution function that incorporates individuals' risk aversion, see Bakshi et al. (2003), Bliss and Panigirtzoglou (2004) and Anagnou-Basioudis et al. (2005).

An alternative approach for retrieving the objective distribution function describing the random behavior of stock prices is to use historical asset prices. Under this modelling framework, time series of stock returns are used as inputs of parametric and nonparametric procedures for predicting the conditional distribution of asset prices. Prominent examples of this include location-scale time series models for the parametric case or the empirical distribution of returns for the nonparametric case. A popular example within the parametric case is that of the family of autoregressive models for conditional mean and ARCH, GARCH type processes for conditional volatility, see Engle (1982) and Bollerslev (1986). Other notable examples using the historical approach in option pricing are Jackwerth (2000), Aït-Sahalia and Lo (2000), Prignon and Villa (2002) and Rosenberg and Engle (2002). An additional advantage of this approach is that it allows the use of Monte-Carlo simulation methods for forecasting asset prices, once the model parameters and the distribution error are estimated. In particular, simulation methods can be straightforwardly used to obtain estimates of the objective distribution function of asset prices.

A strand of empirical studies comparing option-implied methods against historical time series models for predicting expected asset returns has revealed the outperformance of option prices methods, see, for example, Anagnou-Basioudis et al. (2005), Liu et al. (2007) and Shackleton et al. (2010). This result can be attributed to the forward-looking nature of option prices. Indeed, when markets are efficient, the option price should reflect market expectations of the underlying asset for the time until option expiration in a way that any information arriving into the market until expiry date should be priced into the distribution. In contrast, the estimates of future asset prices obtained from historical data are merely based on the past dynamics of asset prices. In this paper, we explore these findings from an asset allocation perspective.

To the best of our knowledge, empirical work on this area is still very limited, with seminal references being the work of Kostakis et al. (2011), DeMiguel et al. (2013) and Kempf et al. (2015). The main contribution of our study is to explore alternative state variables based on option-implied information to construct optimal portfolios as in Campbell and Viceira (1999). To do this, we use a parametric portfolio policy rule for infinitively-lived investors and investigate the performance of two different state variables that track the excess return on the underlying risky asset. We entertain an investment portfolio comprised by two assets (a risky and a risk-free asset). The dynamic weights on the risky asset determining the optimal allocation are driven by a state variable that tracks the excess return on the risky asset.

Our first investment portfolio considers the risk premium on the stock return as state variable replacing the log dividend-price ratio proposed in Campbell and Viceira (1999). The state variable defining the second investment portfolio is the market price of risk given by the ratio of the risk premium over the standard deviation of the asset return. The conditional mean and variance of the excess return on the risky asset are then estimated using forward-looking information from option prices written on the underlying risky asset. For this purpose, we need to convert the risk-neutral conditional mean and variance of the asset price into the conditional mean and variance of the asset return evaluated under the objective distribution function. Therefore, our second contribution is to propose a Taylor expansion to approximate the latter two statistical moments of the return on the risky asset. To the best of our knowledge this expansion has not been developed in this context yet, see Jondeau and Rockinger (2006) for the use of higher moments in a related context. The third methodological contribution of this study is to explore the robustness of the portfolio allocations obtained from the above state variables to different modelling choices of the implied risk-neutral distribution of stock returns.

We consider a lognormal distribution for asset prices 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).

The suitability of the different state variables and methods to approximate the implied distribution of asset prices is assessed in-sample and out-of-sample using statistical performance measures such as the Sharpe and Sortino ratios, as well as economic performance measures such as the value function measuring investor's utility and a raw measure of revenue in percent terms. We also use these measures to compare the statistical and economic performance of our parametric portfolios against the benchmark portfolio introduced by Campbell and Viceira (1999) and characterized by the log dividend-price ratio. Our optimal portfolios obtained from option-implied information are also compared against portfolios constructed from historical information. The latter portfolios are obtained using the risk premium and the market price of risk, as above, but estimated using parametric ARCH type models. These methods are applied to a portfolio given by the US one-month Treasury bill and the S&P 500 Index.

Three main findings arise from the empirical study. First, the option-implied approach provides superior portfolio performance compared to the historical approach using economic performance measures, but not in statistical terms. In particular, Sharpe and Sortino ratios are comparable and sometimes higher for the historical approach than for the option-implied approach. In contrast, the value function and revenue measures are superior for the portfolios constructed from option-implied information. These findings are obtained for those portfolios using as state variables the risk premium and the market price of risk. Nonetheless, we find that the long-term portfolio proposed by Campbell and Viceira (1999) based on the log dividend-price ratio is difficult to outperform. Second, we show that the functional form of the implied risk-neutral distribution function does not play a major role on performance for those portfolios exploiting implied information from option prices. In particular, our empirical study 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 optimal portfolio weights that use the option-implied risk-premium as state variable yield superior portfolio performance compared to the weights constructed from the market price of risk. Furthermore, restricting the portfolio rules for market timing highlights the timing abilities of the implied risk premium compared to the market price of risk.

These results are in line with the findings observed in the empirical literature that show that option-implied information improve market timing and yield superior portfolio allocation. Related studies focus on solving option-implied mean-variance portfolios and compare them against the historical approach, equally-weighted strategies and investment on the index. DeMiguel et al. (2013) and Kempf et al. (2015) show that the use of option-implied information is preferable when finding an optimal allocation for multiple assets compared to different benchmarks. 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. Another influential study very related to our method is that of 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 (2004). This investment strategy is compared against a benchmark constructed from historical data. The empirical findings of Kostakis et al. (2011) 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 empirical findings of our present study.

The rest of the paper is organised as follows. Section 2 outlines the investor's asset allocation problem for a lifetime investor with a parametric portfolio policy rule. Section 3 presents different methods to retrieve the option-implied distribution of the underlying risky asset and discusses a transformation to obtain the corresponding objective density function. Section 4 introduces statistical and economic performance measures to assess portfolio performance and compare it against two benchmark portfolios. Section 5 applies these methods to a portfolio given by the US one-month Treasury bill and the S&P 500 Index. Section 6 summarizes the main findings of the study and concludes. A mathematical appendix details the approximation of the first two conditional moments of the log-return on the risky asset using a Taylor expansion. Finally, tables are collected in a second appendix.

2 Investor's asset allocation problem

This section outlines the lifetime portfolio allocation problem solved in Campbell and Viceira (1999) which serves as a useful benchmark for our study. The authors provide a simple tractable analytical solution to an optimization problem with an infinitely-lived investor who chooses both consumption and portfolio allocation in a discrete time setting. The infinitely-lived investor is assumed to have Epstein-Zin-Weil preferences (Epstein and Zin (1989) and Weil (1989)) and maximizes the expected utility of future wealth as in (1) below.

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

where $u(c_t)$ is a period utility function, c_t denotes consumption and $E_t[\cdot]$ is the conditional expectation at time t; $\gamma > 0$ is the coefficient of risk aversion; $\theta = (1-\gamma)/(1-\psi^{-1})$, δ is the discount factor and, finally, ψ is the elasticity of intertemporal substitution.

Campbell and Viceira (1999) consider a two-asset portfolio. The investor seeks an optimal allocation between a single risk-free asset with constant log return r_f and a single risky asset with log return r_{t+1}. The dynamics of the log-return on the risky asset are defined as:

$$r_{t+1} = E_t[r_{t+1}] + u_{t+1}, \tag{2}$$

with $E_t[r_{t+1}]$ being the conditional expected return on the risky asset and u_{t+1} the innovation process. This process is normally distributed with zero mean and variance σ_u^2. The corresponding excess return on the risky asset is $r_{t+1}^e = r_{t+1} - r_f$. An important feature of this framework is that the risk premium on the risky asset is state-dependent and driven by a single state variable x_t, which satisfies

$$E_t[r_{t+1}^e] = f(x_t), \tag{3}$$

with $f(x_t)$ a deterministic function of the state variable. Following Campbell and Viceira (1999), the state variable x_t is modelled as a mean-reverting AR(1) process with mean μ and persistence parameter $|\phi| < 1$. Mathematically,

$$x_{t+1} = \mu + \phi(x_t - \mu) + \eta_{t+1}. \tag{4}$$

The associated innovation process η_{t+1} is conditionally homoskedastic and normally distributed with zero mean and variance σ_η^2. The two innovations η_{t+1} and u_{t+1} are correlated with each other, with the correlation coefficient denoted by $\sigma_{\eta u}$.

The log-return on the optimal investment portfolio is defined as $r_{t+1}^P = r_f + \alpha_t r_{t+1}^e$, with α_t the parametric portfolio weight that is defined as a deterministic function of the state variable. Campbell and Viceira (1999) derive the optimal consumption and portfolio policies under the following assumptions on the functional form of the portfolio weight α_t and the log consumption-wealth ratio. More specifically,

$$\alpha_t = a_0 + a_1 x_t \tag{5}$$

$$c_t - w_t = b_0 + b_1 x_t + b_2 x_t^2 \tag{6}$$

The fixed parameters a_0, a_1, b_0, b_1 defining the linear portfolio policy are determined by the following equations (see Proposition 1 in Campbell and Viceira (1999)).

$$a_0 = \frac{1}{2\gamma} - \frac{b_1}{1-\psi}\frac{\gamma-1}{\gamma}\frac{\sigma_{\eta u}}{\sigma_u^2} - \frac{b_2}{1-\psi}\frac{\gamma-1}{\gamma}\frac{\sigma_{\eta u}}{\sigma_u^2}2\mu(1-\phi) \tag{7}$$

$$a_1 = \frac{1}{\gamma\sigma_u^2} - \frac{b_2}{1-\psi}\frac{\gamma-1}{\gamma}\frac{\sigma_{\eta u}}{\sigma_u^2}2\phi, \tag{8}$$

As anticipated, these expressions depend on both the risk-aversion parameter γ and the elasticity of intertemporal substitution,ψ, which emerges from the Epstein-Zin-Weil utility function. The remaining parameters of the log consumption-wealth ratio are retrieved by a recursive non-linear system and the exact definition of the associated formulae terms and the recursive procedure to obtain $[b_0, b_1, b_2]$ are summarized in Proposition 2 of Campbell and Viceira (1999).

The expression for the optimal portfolio weight α_t consists of two components. The first term defines the myopic asset demand. The myopic component is proportional to the risk-premium and inversely proportional to the volatility on the risky asset and the relative risk aversion coefficient. The second component captures the intertemporal hedging demand. This term captures the excess allocation to the risky asset used to compensate against changes in the investment opportunity set over time. The presence of this term in long-term portfolio allocation problems goes back to Merton (1969, 1971, 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 changes in the investment opportunity set is. Two special cases also become apparent from (7) and (8). Firstly, the hedging demand becomes zero when $\sigma_{\eta u} = 0$ and secondly, when individuals are risk-neutral, that is $\gamma = 1$, and the optimal allocation to the risky asset α_t is only driven by the myopic component.

The optimal portfolio allocation is characterized by the parameters $[\mu, \phi, \sigma_u^2, \sigma_\eta^2, \sigma_{\eta u}]$. Campbell and Viceira (1999) estimate a restricted VAR(1) model using ordinary least square methods with the state variable x_t as single predictor:

$$\begin{pmatrix} r_{t+1}^e \\ x_{t+1} \end{pmatrix} = \begin{pmatrix} \theta_0 \\ \beta_0 \end{pmatrix} + \begin{pmatrix} \theta_1 \\ \beta_1 \end{pmatrix} x_t + \begin{pmatrix} \varepsilon_{1,t+1} \\ \varepsilon_{2,t+1} \end{pmatrix} \tag{9}$$

where $(\varepsilon_{1,t+1}, \varepsilon_{2,t+1})' \backsim N(0, \Omega)$ and

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

The parameters defining the optimal consumption and portfolio allocation problem are identified from the parameters of the restricted VAR(1) model 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 normalized in a way so that the intercept a_0 of the optimal policy function determines the optimal allocation to stocks and the optimal consumption-wealth ratio when the risk premium on the risky asset is zero. Under this scenario a myopic investor would not allocate any wealth to it. Therefore, any asset demand for a lifetime portfolio allocation would reflect intertemporal hedging demand. This normalization 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 $b_1^* = b_1 - b_2\sigma_u^2$, while the parameters a_1 and b_2 are not affected by this transformation.

A crucial element in defining the parametric portfolio policy lies in the choice of state variable x_t. Campbell and Viceira (1999) propose the log dividend-price ratio, which is, in turn, motivated by findings in Campbell and Shiller (1988), Fama and French (1988), Hodrick (1992), and others, which found this variable to be a good predictor of stock returns. In this paper we take an alternative approach and propose two state variables x_t constructed from option-implied information. Our first state variable is the risk premium on the risky asset. Using the above formulation in expression (3), we consider $f(x_t) = x_t$ such that $x_t = E_t[r_{t+1}^e]$. Our second state variable is the market price of risk initially proposed by Wachter (2002) and applied to option prices by Kostakis et al. (2011). This variable resembles the Sharpe ratio and is defined as $x_t = \frac{E_t[r_{t+1}^e]}{\sqrt{Var_t[r_{t+1}]}}$, with $Var_t[r_{t+1}]$ denoting he conditional variance of the risky asset one-period ahead.

Rather than exploiting the historical information contained in financial ratios as in Campbell and Viceira (1999) and, in particular, its predictive ability, we focus on exploiting the content of forward-looking information extracted from option prices. The information obtained from option prices enters the state variables via the first two moments of the log-returns on the underlying asset $E_t[r_{t+1}]$ and $Var_t[r_{t+1}]$. Thereby, our next objective is to approximate these moments from the implied distribution of the risky asset under risk-neutrality conditions.

3 Obtaining Option Implied Moments

Our option-implied approach requires the use of information from option prices transmitted by the first two moments. In this section we explore option-implied moments obtained from three different density types; namely, a lognormal density, a mixture-lognormal and a binomial tree. In a model with lognormal density it is assumed that the volatility is constant along the range of strike prices. This allows one to estimate the implied density function using the two closest option prices above and below the current forward price on the underlying risky asset. More specifically, the implied volatility between both options can be interpolated based on the current forward price to match the at-the-money volatility. In contrast, employing a mixture-lognormal density allows for more complex shapes. Estimation takes into account the entire cross-section of option prices to capture the non-normal shape contained in option prices. The use of a binomial tree is, nevertheless, the most flexible density type owing to its nonparametric nature.

We now review the mixture of two lognormal density functions. The lognormal case can be viewed as a particular case of the mixture model.

3.1 Mixture of two Lognormal Densities

Under the absence of 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.[1] The no-arbitrage call option price is defined as

$$c_{t,m}(K) = e^{-r_{f,t}(T-t)} \int_K^{\infty} (s-K) f_{Q,t}(s) ds, \tag{10}$$

with $r_{f,t}$ the risk-free rate, T the time to maturity and $f_{Q,t}(\cdot)$ the risk-neutral density. From this expression, we can obtain the risk-neutral density function evaluated at the different strike prices K as

$$f_{Q,t}(K) = e^{r_{f,t}(T-t)} \frac{\partial^2 c_{t,m}}{\partial K^2}. \tag{11}$$

In this setting, the risk-neutral density function $f_{Q,t}(K)$ is assumed to be a mixture of two lognormal distribution functions, see Ritchey (1990), Melick and Thomas (1997), Brigo and Mercurio (2002) and Liu et al. (2007). This literature shows that the call option price is derived from two weighted models, see Black (1976):

$$c_{t,m}(K \mid \theta, r_f, T) = w_t c_{t,B}(\mu_1, T, K, r_{f,t}, \sigma_1) + (1-w_t) c_{t,B}(\mu_2, T, K, r_{f,t}, \sigma_2), \tag{12}$$

where $c_{t,m}(K \mid \theta, r_{f,t}, T)$ denotes the option price as a function of the strike price and parametrized by a set of model parameters denoted as θ, the risk-free rate r_f, and the time to maturity T;

[1] c and K denote here vectors of prices for the entire cross-sections of options in the market.

$c_B(\mu_1, T, K, r_{f,t}, \sigma_1)$ and $c_B(\mu_2, T, K, r_{f,t}, \sigma_2)$ are the option prices obtained from Black (1976)'s model. The statistical moments of order n are obtained from the following expression:

$$E_{t,rn}[S_T^n] = w_t \mu_1^n exp(0.5(n^2 - n)\sigma_1^2(T - t)) + (1 - w_t)\mu_2^n exp(0.5(n^2 - n)\sigma_2^2(T - t)), \qquad (13)$$

with $E_{t,rn}[\cdot]$ denotes the expected value evaluated under the risk-neutral implied distribution function; $\theta = [\mu_1, \mu_2, \sigma_1, \sigma_2, w_t]$, with $0 \leq w_t \leq 1$. Each of the weighted models in Black (1976)'s formulation has its own mean $[\mu_1, \mu_2]$ and variance $[\sigma_1, \sigma_2]$. This results in five parameters that need to be determined.

The vector of parameters θ is obtained by minimizing the squared error between the vector of observed prices, $c_{t,o}$, and model generated option prices, $c_{t,m}$, as

$$\min_{\theta} \sum_{i=1}^{N} \left[c_{t,o}(K_i) - c_{t,m}(K_i \mid \theta, r_f, T)\right]^2, \qquad (14)$$

where K_i denotes the cross-section of N option prices available for the underlying risky asset. The optimization 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 F_t, has to equal the expected value of the stock price at maturity, S_T, under the risk-neutral density function. More formally,

$$F_t = E_{t,rn}[S_T] = \int_0^{\infty} s f_{Q,t}(s) ds = w_t \mu_1 + (1 - w_t)\mu_2. \qquad (15)$$

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. Bliss and Panigirtzoglou (2004) provide a framework to transform the risk-neutral density into its objective counterpart, denoted as $f_{P,t}(S_T)$. These authors relate both density functions through the marginal utility of the stock price for the investor. For investors with a power utility function, $u(y) = \frac{y^{1-\gamma}}{1-\gamma}$, where γ denotes the coefficient of relative risk aversion, the objective density function is obtained as

$$f_{P,t}(S_T) = \frac{\frac{f_{Q,t}(S_T)}{u'(S_T)}}{\int_0^\infty \frac{f_{Q,t}(y)}{u'(y)}dy} = \frac{S_T^\gamma f_{Q,t}(S_T)}{\int_0^\infty y^\gamma f_{Q,t}(y)dy}. \tag{16}$$

The integral in the denominator of (16) ensures the objective density integrates to 1. The power utility function adds one more parameter to the set θ inherited by the risk-neutral density. Importantly, Liu et al. (2007) show that applying the result in (16) for the mixture lognormal density function and a power utility function results in another mixture lognormal density function. The transformed parameter values of the objective mixture lognormal density function are

$$\mu_i^\star = \mu_i exp(\gamma\sigma_i^2(T-t)) \text{ for } i = 1, 2 \text{ and} \tag{17}$$

$$\frac{1}{w^\star} = 1 + \frac{1-w}{w}\left(\frac{\mu_2}{\mu_1}\right)^\gamma exp(0.5(\gamma^2-\gamma)(\sigma_2^2-\sigma_1^2)(T-t)). \tag{18}$$

The transformation from the risk-neutral to the objective density function only changes the values for the mean of the distributions and the weights in the mixture of distributions from w to $w^\star$. The volatilities σ_1 and σ_2 are not affected by the transformation. Then, the first n statistical moments evaluated under the objective density function are

$$E_t[S_T^n] = w^\star\mu_1^{\star n} exp(0.5(n^2-n)\sigma_1^2(T-t)) + (1-w^\star)\mu_2^{\star n} exp(0.5(n^2-n)\sigma_2^2(T-t)), \tag{19}$$

with $E_t[\cdot]$ denoting the conditional expected value under the density function $f_{P,t}(\cdot)$.

It is worth noting that in order to align this approach to retrieve the objective density function with Campbell and Viceira (1999)'s framework, that considers an Epstein-Zin utility function for describing individual's preferences, we assume $\psi = \frac{1}{\gamma}$ in (1). In doing so, we reduce the recursive utility function to a power utility function and the framework proposed in Campbell and Viceira (1999) can be naturally applied in our context. Therefore, we can use the long-term portfolio optimization framework developed and solved by these authors with option-implied information, and obtain the objective density function from the transformation (16). The restriction $\psi = \frac{1}{\gamma}$ imposes an elasticity of intertemporal substitution equal to the inverse of the individual's risk aversion coefficient. This assumption offers a closed-form solution to obtain the objective density function of stock prices. In general, Epstein-Zin utility functions do not have a closed-form solution as in (16) for the transformation of the risk-neutral density function to the objective density function of stock prices.

In our empirical exercise we also consider the case of a lognormal density. This distribution corresponds to the case $w_t = 1$ in (12) for all t. The right hand side of this expression vanishes and only the simple Black (1976) model remains with $\theta = [\mu_1, \sigma_1]$. Since the results for the lognormal density are embedded in the mixture case, we do not discuss them separately.

3.2 Binomial Tree

This method allows for greater flexibility by obtaining risk-neutral probabilities without imposing a specific functional form to the risk-neutral density function. This method goes back to Rubinstein (1994), who proposes the use of a binomial tree for option pricing. In this model, the risk-neutral density function is discrete with $n+1$ possible values for the stock price S_T at expiry date T, with corresponding probabilities $q_t = (q_{1t}, ..., q_{mt})$, such that $q_{1t}, ..., q_{mt} \geqslant 0$ and $\sum_{j=1}^{m} q_{jt} = 1$. As in the previous section, the density function is priced risk-neutral and the expected value of the density needs to equal the forward price F_t. More formally, $\sum_{j=1}^{m} q_{jt} S_{jT} = F_t$ characterizes the risk-neutrality constraint. Under those conditions the option pricing formula in (10) becomes:

$$c_{t,m}(K; q_t, m) = e^{-r_{f,t}(T-t)} \sum_{j=1}^{m} \max(S_{jT} - K, 0) q_{jt}, \tag{20}$$

where the first n statistical moments associated to the risk-neutral density function are obtained as

$$E_{t,rn}[S_T^n] = \sum_{j=1}^{m} q_{jt} S_{jT}^n, \tag{21}$$

with $E_{t,rn}$ the conditional expected value evaluated at time t under the risk-neutral set of probabilities.

To estimate the vector q_t, Jackwerth and Rubinstein (1996) propose a range of different objective functions. However, they recommend to minimise a penalty function given by a combination of a smoothness function g of the risk-neutral density function and a squared loss function G measuring the distance between the fit of observed and model option prices similar to (14). We follow this approach, which leads to the following minimisation problem: $\min_{q_t} g(q_t) + \alpha G(q_t)$, where

$$g(q_t) = \sum_{i=2}^{m-1} [q_{i-1,t} - 2q_{i,t} + q_{i+1,t}]^2 \tag{22}$$

$$G(q_t) = \sum_{j=1}^{N} [c_m(K_j; q_t, m) - c_o(K_j; q_t, m)]^2 + \left[\sum_{i=1}^{m} max[0, -q_{i,t}]\right]^2 + \left[\sum_{i=1}^{m} q_{it} - 1\right]^2 \tag{23}$$

$$+ \left[\sum_{i=1}^{m} q_{it} S_{iT} - F_t\right]^2, \tag{24}$$

with $\alpha > 0$ a penalty term that regularizes the optimization problem; N is the number of strike prices for the underlying stock.

Using an analogous expression to (16) adapted to the discrete case, we transform the probability distribution defined by the vector $q_t = (q_{1t}, \ldots, q_{mt})$ in the binomial model into a set of probabilities $p_t = (p_{1t}, \ldots, p_{mt})$ defined as

$$p_{it} = \frac{\frac{q_{it}}{u'(S_{iT})}}{\sum_{j=1}^{m} \frac{q_{jt}}{u'(S_{jT})}} = \frac{S_{iT}^{\gamma} q_{it}}{\sum_{j=1}^{m} S_{jT}^{\gamma} q_{jt}}, \tag{25}$$

for $i = 1, \ldots, m$, that define the objective probability distribution of S_T and S_{iT} is the value closest to S_T. Thereby, the first n statistical moments associated to the objective probability distribution are obtained as

$$E_t[S_T^n] = \sum_{j=1}^{m} S_{jT}^n p_{jt}, \tag{26}$$

with E_t the conditional expected value evaluated under the objective set of probabilities p_t.

3.3 Moment conversion through Taylor expansions

The above methods allow one to obtain the statistical moments of asset prices evaluated under the risk-neutral and objective probability distributions. In contrast, the state variables driving the parametric portfolio policy rule α_t are defined as a function of the mean and variance of asset returns. Therefore, we need to transform the expected values of asset prices into the expected values of asset returns.

To do this, we use a Taylor expansion of the stock price S_T. In what follows, we derive closed-form approximations for $E_t[r_{t+1}]$ and $V_t[r_{t+1}]$ as a function of the first moments of S_{t+1}. Note that, in this setting, we assume that $T = t+1$. Let $r_{t+1} = \ln S_{t+1} - \ln S_t$, we focus on computing $E_t[\ln S_{t+1}]$ first. To do this, we use a second order Taylor expansion of $\ln S_{t+1}$ about $E_t[S_{t+1}]$, and obtain

$$\ln S_{t+1} = \ln E_t[S_{t+1}] + \frac{1}{E_t[S_{t+1}]}(S_{t+1} - E_t[S_{t+1}]) - \frac{1}{2E_t[S_{t+1}]^2}(S_{t+1} - E_t[S_{t+1}])^2 + o_P((S_{t+1} - E_t[S_{t+1}])^2). \tag{27}$$

Appendix A shows the following relationship between the conditional expected return and the conditional expected value of the stock price:

$$E_t[r_{t+1}] \approx \ln E_t[S_{t+1}] - \ln S_t - \frac{1}{2E_t[S_{t+1}]^2} V_t[S_{t+1}], \tag{28}$$

and

$$V_t[r_{t+1}] \approx \ln^2 E_t[S_{t+1}] + \frac{1 - \ln E_t[S_{t+1}]}{E_t[S_{t+1}]^2} V_t[S_{t+1}] - \frac{V_t[S_{t+1}]^{3/2}}{E_t[S_{t+1}]^3} Skew_t(S_{t+1}) + \frac{V_t[S_{t+1}]^2}{4E_t[S_{t+1}]^4} Kurt_t(S_{t+1}) \tag{29}$$

$$+ \ln^2 S_t - 2 \underbrace{\left(\ln E_t[S_{t+1}] - \frac{1}{2E_t[S_{t+1}]^2} V_t[S_{t+1}] \right)}_{E_t[\ln S_{t+1}]} \ln S_t \tag{30}$$

$$- \underbrace{\left(\ln E_t[S_{t+1}] - \ln S_t - \frac{1}{2E_t[S_{t+1}]^2} V_t[S_{t+1}] \right)^2}_{E_t[r_{t+1}]}, \tag{31}$$

with $Skew_t(S_{t+1})$ and $Kurt_t(S_{t+1})$ denoting the conditional skewness and kurtosis of the stock price S_{t+1}. The next section discusses estimation of the model parameters.

4 Evaluation Criteria

This section presents the performance measures used to assess the suitability of the state variables obtained from option-implied information, and introduces two benchmark competitors that will be discussed below.

4.1 Performance measures

We evaluate the performance of the different investment portfolios using out-of-sample monthly returns $r_{P,t}$ on the investment portfolio. The optimal portfolio allocation α_t is performed for the whole sample of 243 periods. For the out-of-sample evaluation period, we consider a fixed window with the first 180 periods for estimating the model parameters. The remaining periods are used to evaluate the out-of-sample performance. We rely on Exchange Traded Funds (ETF) to calculate the performance of the risky asset. The advantage of using an ETF to evaluate the performance of the portfolio allocation compared to an index value is the possibility of including trading costs and administration fees in the portfolio. The returns of the suggested portfolio allocation are calculated on the basis of a theoretical fund 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 each period is adjusted for the difference in weight between $t-1$ and t. Furthermore, we restrict the portfolio weight to be in the interval $[-1, 2]$ to avoid unrealistic leverage.

The resulting returns are first standardised to 4-week horizons before computing the performance measures. This is necessary since option expiry dates are not necessarily connected to each other. If the estimation horizon is not connected, we assume that the investor holds the portfolio until the next rebalancing date. Based on the standardised out-of-sample returns, we evaluate the performance of the proposed state variables using two statistical criteria given by the I) Sharpe ratio and II) Sortino ratio; and two economic measures given by the III) Value function; IV) Portfolio revenue. These criteria are calculated over a series of h realisations.

The annualised Sharpe ratio (SR) is computed using the unconditional average excess return on the portfolio over the out-of-sample period, denoted as $\hat{\mu}_P = \frac{1}{h}\sum_{t=1}^{h}(r_{P,t} - r_{f,t})$, and corresponding unconditional sample variance, denoted as $\hat{\sigma}_P^2 = \frac{1}{h-1}\sum_{t=1}^{h}(r_{P,t} - \hat{\mu}_P)^2$. The annualized Sharpe ratio is defined as

$$\widehat{SR}_P = \frac{365}{28}\frac{\hat{\mu}_P}{\sqrt{\hat{\sigma}_P^2}}. \tag{32}$$

To add statistical significance to the comparison of Sharpe ratios across investment strategies, we apply the Sharpe ratio test introduced by Ledoit and Wolf (2008). We follow the implementation in DeMiguel et al. (2014) and consider $1,000$ bootstrap resamples and an expected block size equal to 5. To gain further insights on the statistical performance of the different methods, we also compute the Sortino ratio. This measure extends the Sharpe ratio by only considering the downside volatility defined as $\hat{\sigma}_{adj}^2 = \frac{1}{h-1}\sum_{t=1}^{h}(\min(0, r_{P,t} - r_{f,t}))^2$, and, therefore, putting the emphasis on returns falling below the risk-free rate. The annualized Sortino ratio is

$$\widehat{SoR} = \frac{365}{28}\frac{\hat{\mu}}{\sqrt{\hat{\sigma}_{adj}^2}} \tag{33}$$

Our first economic measure is the value function. This function is obtained from the realized values of the utility function (1) under the restriction that the elasticity of intertemporal substitution is the inverse of the risk aversion coefficient. In this case the Epstein-Zin-Weil utility function reduces to a standard power utility function. Lastly, the average portfolio revenue is evaluated. In this case we compare the aggregate revenue obtained from each portfolio over the life of the investment.

As a further exercise, we analyze the portfolio allocation itself to gain insight on the effectiveness of the myopic and hedging demand components. Therefore, we restrict the portfolio rule while adjusting optimally for consumption. This procedure is outlined in detail in Campbell and Viceira (1999). We consider three different restricted portfolio rules: First, restriction of the timing component. This is achieved by setting the state variable x_t in (5) to the unconditional expected log excess return $E[r_{P,t}]$. Second, restriction of the hedging demand by ignoring the hedging demand on the portfolio allocation. This is accomplished by setting the covariance $\sigma_{\eta u}$ to zero. Third, both components are restricted by imposing $x_t = E[r_t^P]$ and $\sigma_{\eta u} = 0$.

4.2 Benchmark portfolios

In the empirical application we consider two benchmark portfolios competing against the portfolios constructed from state variables that use option-implied information. These competitors are *i*) a parametric portfolio with portfolio weights driven by the log dividend-price ratio and *ii*) a historical approach in which the risk premium and market price of risk state variables are estimated from historical data.

The first benchmark portfolio corresponds to the long-term portfolio initially proposed by Campbell and Viceira (1999) that considers the log dividend-price ratio as single state variable. The log dividend-price ratio enters the portfolio allocation directly as predictor in (9). Second, we implement a historical approach as benchmark. More specifically, we model the conditional mean and variance of asset returns using a time series model. The conditional mean process is modelled using an autoregressive process of order one and the volatility component is modelled using a threshold-GARCH model. This proccess is defined as

$$r_{t+1} = \mu + \phi r_t + \varepsilon_{t+1} \tag{34}$$

$$\varepsilon_{t+1} = u_{t+1}\sqrt{h_{t+1}}. \tag{35}$$

$$h_{t+1} = w + (\alpha + \beta I_t)\varepsilon_t^2 + \delta h_t \tag{36}$$

with $I_t = 0$ if $r_t \geq 0$ and $I_t = 1$ if $r_t < 0$. This process dates back to Glosten et al. (1993) and Zakoian (1994). In contrast to standard ARCH and GARCH processes, see Engle (1982) and Bollerslev (1986), this model accommodates asymmetries on the volatility process due to the leverage effect (negative shocks to asset prices entail higher volatility in the following period compared to positive shocks). To guarantee the stationarity of the process, the parameters of this process satisfy that $w, \alpha, \beta, \delta > 0$ and $\alpha + \frac{\beta}{2} + \delta < 1$. The conditional mean and volatility processes are $E_t[r_{t+1}] = \mu + \phi r_t$ and $V_t[r_{t+1}] = h_{t+1}$.

5 Empirical application

5.1 Data

The empirical application presented in this section is based on the S&P 500 Index (SPX). The observation period ranges from 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. 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, the risk-neutral density is estimated using the closing price on the option observation date, which results in 243 periods.

In this empirical study we only consider European options which are always written on the underlying spot market index. The structure of European options allows us to assume that 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 Mercantile Exchange. 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-the-money are taken into account. In-the-money options are less frequently traded and, therefore, prices might not fully 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 1. If necessary, options 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 $F_t = S_t e^{(r_{f,t}-d_t)(T-t)}$. The risk-free rate r_f is retrieved via bootstrapping from zero-coupon bonds. The dividend yield d_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 d_t provides the log dividend price ratio $p_t - d_t$ in the style of Campbell and Viceira (1999). Lastly, we obtain ETF bid and ask prices for the the SPDR S&P 500 ETF to evaluate the portfolio allocation. The ETF was selected based on the fund inception date and size.

5.2 Parameter estimation

This section presents the estimation results of the risk-neutral parameters, the transformation into the objective density function and the estimation results of the historical approach. Following this, we explore the implications of the parameter values obtained from the portfolio allocation across state variables and methods. Results presented in this section are obtained from estimates using the entire sample available.

Table 2 presents the risk-neutral density estimation error for the mixture of lognormal density functions and the binomial tree. To ensure that the squared error converges to a global minimum, multiple initial values have been applied into the nonlinear 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 function fits precisely the data. Table 3 presents the parameter estimates for the AR(1)-TGARCH(1,1) model introduced above obtained from the whole evaluation period.

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. Expression (16) shows that the transformation relies on an assumption about the utility of the representative investor. A conventional standpoint is to assume a power utility function. To avoid the excessive parametrization of the model we consider the estimation of the risk aversion coefficient obtained by Bliss and Panigirtzoglou (2004) to characterize our representative investor. These authors provide an extensive study to estimate the risk aversion coefficient from option prices and evaluate the results with a sequence of different statistical measures. For the following empirical study we consider $\gamma = 4.02$ for a time horizon of 4 weeks to match the time to expiry of our option sample. Furthermore, in order to align our model that considers a representative investor endowed with a power utility function with Campbell and Viceira (1999)'s framework, that considers an individual with Epstein-Zin utility function, we assume $\psi = \frac{1}{\gamma}$ in (1). In this scenario the recursive utility function reduces to the power utility case and the framework proposed by Campbell and Viceira (1999) can be naturally applied in our context. The reader should note that this assumption has implications on the portfolio policy rule α_t as the coefficients a_0 and a_1 in (7) and (8) simplify.

Table 4 reports the estimation results for the restricted VAR(1) model for the SPX dataset using the entire sample. 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. Similar results are found for the historical approach. Nevertheless, in both cases, the restricted VAR(1) model given by the log dividend-price ratio as state variable provides a higher R^2. Comparing the dynamics of the two equations within the VAR(1) model 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. 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 viceversa. Therefore, the ratio between both factors is relatively stable compared to the use of the risk premium as single state variable.

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 applied model here there is no direct link between the risk premium and volatility, so both factors are less related to each other, which helps 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 lines up with the alternative methods. Overall, values are similar to the seminal paper by Campbell and Viceira (1999).

5.3 Portfolio weight characteristics

As briefly discussed in Section 2, there are two components of the portfolio weight, namely myopic and hedging demand. While the myopic demand entirely derives from the expected return and variance of the innovation process σ_u^2, the hedging demand component depends on the covariance $\sigma_{\eta u}$ between the innovations to the state variable η_{t+1} and to the risky asset 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 on 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), cases where $\sigma_{\eta u} < 0$ are 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. The values for $\sigma_{\eta u}$ in Table 4 are positive in the historical approach. This might be a result of the slow adjustment of this method. Since the historical returns are simply extrapolated into the future, this method lags anticipating changes in the market. However, to further analyze 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.

Table 5 displays the average portfolio weights and contribution of the hedging demand. In this 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. In contrast, the comparison of the risk premium and market price of risk as state variables with the log dividend-price ratio reveal important differences. Overall, the average portfolio weight for the log dividend-price ratio exceeds the weights corresponding to the alternative methods considerably. The contribution of the hedging demand to the portfolio weight for the log dividend-price ratio based portfolio is also much larger than for the other portfolio strategies. This empirical finding can be partly explained by the magnitude of the correlation between the innovations to the log dividend-price ratio state variable and to the stock returns. As mentioned earlier, this correlation gives an indication of how effectively the state variable can be used to hedge against changes in investment opportunities.

5.4 Performance results

Table 6 reports the results of the different performance measures covered in the study. We consider two separate exercises given by two diferent evaluation periods: a first exercise that spans the full dataset and a second exercise that divides the sample set into in-sample and out-of-sample periods. As discussed in Section 4.1, the in-sample analysis considers the full sample whereas for the out-of-sample analysis the sample is split in two subsamples. The first subsample is used to estimate the parameters in the VAR(1) model in (9) and the remaining periods are then used to evaluate the performance using the obtained parameter values.

In addition to the Sharpe and Sortino ratios, we also compute the value function (1) under the assumption that the elasticity of intertemporal substitution is equal to the inverse of the coefficient of relative risk aversion. In this case the value function corresponds to that of a constant relative risk aversion (CRRA) utility function. For completeness, we also report the revenue of the investment strategy in percent terms. Table 6 is divided in three panels. The top panel reports the performance measures for the risk premium obtained from option-implied information under different choices of the option-implied density function. The middle panel reports the same empirical exercise for the market price of risk obtained from option-implied information, and the bottom panel focuses on the performance measures obtained from Campbell and Viceira (1999)'s approach based on the log dividend-price ratio.

First, we compare the option-implied state variables given by the risk premium and the market price of risk for different choices of the distribution function implied from option prices. The risk premium state variable seems to perform considerably better than the market price of risk. This result suggests that the risk premium obtained from option-implied information is a good predictor of the actual future excess returns on the risky asset. More specifically, large values of this state variable predict large excess returns and hence, a larger allocation to the risky asset. In contrast, the portfolio allocation driven by the market price of risk gets less exposure to the risky asset under large expected returns. The reason for this is that the risk premium is adjusted by the underlying risk, proxied by the conditional variance. The choice of method for modelling the implied distribution of the risky return does not seem to be instrumental for obtaining these findings. These results hold for both the full sample and the out-of-sample periods.

Second, we compare the portfolios obtained from using option-implied information against the corresponding portfolios in which the state variables are estimated from time series models (historical approach). The results are mixed; statistical measures such as the Sharpe and Sortino ratios show the superiority of the historical approach. In contrast, economic measures such as the value function and portfolio revenue suggest the opposite entailing a superior performance of the option-implied methods compared to the historical approach.

Finally, we compare the option-implied methods against the portfolio policy using the log dividend-price ratio as state variable. In this case the results show a clear outperformance of this method compared to the state variables proposed in this paper. This finding suggests that the log dividend-price ratio is a better predictor of the excess return on the risky asset that state variables based on forward-looking information. A potential explanation for these findings is the ability of the log dividend-price ratio and, more generally, of financial ratios for predicting excess returns. In our case the persistence of the log dividend-price ratio entails persistent portfolio weights that take full advantage of periods of large returns on the risky asset. In contrast, our approach based on information obtained from option prices is based on state variables that are stationary and revert quickly to the mean implying a set of portfolio weights that revert more quickly to the intercept a_0 in (5) than under the log dividend-price ratio.

Table 7 adds statistical significance to the performance measures discussed above by reporting the p-values of the Sharpe ratio test developed by Ledoit and Wolf (2008). The null hypothesis corresponds to equality of Sharpe ratios across strategies and the alternative hypothesis implies that one strategy outperforms the other. The results in this table are very inconclusive and do not allow to draw any meaningful conclusion about the superiority of one strategy over the others in terms of Sharpe ratios.

5.5 Effect of hedging and timing the portfolio

This section examines the influence of the components comprising the portfolio weights, namely hedging demand and timing of the portfolio, on portfolio performance. To do so, the portfolio rule is set fixed while allowing the investor to optimally adjust their consumption. Table 8 reports changes in the value function (1) when restricting the portfolio rules compared to the unrestricted policy given in (5) for the estimated parameters. This exercise is divided into a full sample exercise and an out-of-sample evaluation period. If the additional component (hedging demand or portfolio timing) is beneficial, then the value function should display a negative value indicating a reduction compared to the unrestricted rule.

We analyze first the effect of the timing component to the portfolio using the value function as performance measure. The timing component is set fixed to equal to the long run log excess return. The results reported in Table 8 show a large loss in value function with respect to the unrestricted portfolio for those portfolios constructed from option-implied information and using the risk premium and the market price of risk as state variables. For the historical approach, the effects of market timing seem to be much smaller than for the other strategies. In contrast, for the log dividend-price ratio the effect of timing the portfolio is even larger than for option-implied portfolios.

In the second exercise we restrict the ability of the portfolio to hedge. This is done by neglecting the hedging demand component from the portfolio. To assess the performance of this restricted portfolio we compare the value functions across strategies and portfolios. The results are qualitatively similar to the previous exercise analyzing the timing component. The results are less sizeable for the out-of-sample evaluation period but still reflect in all cases a loss in performance (value function) compared to the unrestricted portfolio.

Unsurprisingly, the combination of no hedging demand and no timing yields the worst performance across investment strategies and portfolios.

6 Conclusion

This paper explores the role of option-implied information from a long-term optimal asset allocation perspective. To do this, we propose two state variables constructed from option-implied information and build parametric portfolios based on the long-term asset allocation problem presented in Campbell and Viceira (1999). The first variable is given by the risk premium on the risky asset and the second variable is the market price of risk. We also explore different choices proposed in the literature to extract the information from option prices for constructing the predictive density function of the underlying risky asset. In particular, we consider a lognormal distribution, a mixture of lognormal distributions and a binomial tree. A third methodological contribution is to approximate the conditional first two moments of the expected log return on the risky asset using a Taylor expansion of the log of the asset price. This approximation is necessary for constructing our state variables based on the risk premium and the market price of risk.

We assess these strategies to construct optimal parametric portfolios using statistical measures such as the Sharpe and Sortino ratios, and economic measures such as the value function and the revenue in percent terms. The empirical results to a portfolio given by the one-month US Treasury bill and the S&P 500 Index are mixed. The risk premium estimated from forward-looking information outperforms the market price of risk as a suitable state variable for constructing optimal portfolios. The reason for this outperformance is the superior ability of the risk premium compared to the market price of risk for predicting the excess return on the risky asset. This result is robust to different specifications of the implied distribution function of the risky asset obtained from option prices. We also compare these strategies based on implied information with the Campbell and Viceira (1999) parametric long-term portfolio given by the log dividend-price ratio as state variable and an historical approach in which the estimates of the risk premium and the conditional variance are obtained from time series of historical prices. The results show the outperformance of the historical approach for statistical performance measures such as the Sharpe and Sortino ratios and the outperformance of the option-implied information for economic measures given by the value function obtained from the investor's lifetime utility function and revenue in percent terms.

Overall, these results highlight the importance of considering option-implied information as an alternative and valuable option for constructing investment portfolios. Suitable state variables are obtained from the first moments of the distribution of asset returns that are obtained from forward-looking measures. Despite the appeal of these measures, our results show that it is difficult to beat financial ratios such as the log dividend-price ratio. In fact, whereas we show that the use of the implied risk premium and the market price of risk state variables results in better value function performance than the use of historical methods, the performance improvement is small compared to the improvement resulting from Campbell and Viceira approach using the log dividend-price ratio – as seen in our empirical exercise. The rationale behind this empirical result is the strong ability of financial ratios for predicting stock returns.

Data availability statement

The data that support the findings of this study are available from the corresponding author upon reasonable request.

Appendix A: Taylor expansion of $\ln S_{t+1}$

In this section we derive closed-form approximations for $E_t[r_{t+1}]$ and $V_t[r_{t+1}]$ as a function of the first moments of S_{t+1}. Let $r_{t+1} = \ln S_{t+1} - \ln S_t$, we focus on computing $E_t[\ln S_{t+1}]$ first. To do this, we use a second order Taylor expansion of $\ln S_{t+1}$ about $E_t[S_{t+1}]$, and obtain

$$\ln S_{t+1} = \ln E_t[S_{t+1}] + \frac{1}{E_t[S_{t+1}]}(S_{t+1} - E_t[S_{t+1}]) - \frac{1}{2E_t[S_{t+1}]^2}(S_{t+1} - E_t[S_{t+1}])^2 + o_P((S_{t+1} - E_t[S_{t+1}])^2). \tag{37}$$

Taking the conditional expectation in the above expression, we obtain

$$E_t[\ln S_{t+1}] = \ln E_t[S_{t+1}] - \frac{1}{2E_t[S_{t+1}]^2} V_t[S_{t+1}] + o_P(V_t[S_{t+1}]), \tag{38}$$

that can be approximated as

$$E_t[\ln S_{t+1}] \approx \ln E_t[S_{t+1}] - \frac{1}{2E_t[S_{t+1}]^2} V_t[S_{t+1}]. \tag{39}$$

Now, by construction, $E_t[r_{t+1}] = E_t[\ln S_{t+1}] - \ln S_t$. Then,

$$E_t[r_{t+1}] \approx \ln \frac{E_t[S_{t+1}]}{S_t} - \frac{1}{2E_t[S_{t+1}]^2} V_t[S_{t+1}] = \ln E_t\left[\frac{S_{t+1}}{S_t}\right] - \frac{1}{2E_t[S_{t+1}]^2} V_t[S_{t+1}] \tag{40}$$

Similarly, we have $V_t[r_{t+1}] = E_t[r_{t+1}^2] - E^2[r_{t+1}]$, and such that

$$E_t[r_{t+1}^2] = E_t[\ln^2 S_{t+1}] + \ln^2 S_t - 2E_t[\ln S_{t+1}] \ln S_t. \tag{41}$$

We only need to derive $E_t[\ln^2 S_{t+1}]$ to obtain the conditional variance as the remaining expressions are already known. From expression (37), we have

$$\ln^2 S_{t+1} = \ln^2 E_t[S_{t+1}] + \frac{1}{E_t[S_{t+1}]^2}(S_{t+1} - E_t[S_{t+1}])^2 + \frac{1}{4E_t[S_{t+1}]^4}(S_{t+1} - E_t[S_{t+1}])^4 \tag{42}$$

$$+ \frac{2\ln E_t[S_{t+1}]}{E_t[S_{t+1}]}(S_{t+1} - E_t[S_{t+1}]) - \frac{\ln E_t[S_{t+1}]}{E_t[S_{t+1}]^2}(S_{t+1} - E_t[S_{t+1}])^2 \tag{43}$$

$$- \frac{1}{E_t[S_{t+1}]^3}(S_{t+1} - E_t[S_{t+1}])^3 + o_P((S_{t+1} - E_t[S_{t+1}])^4). \tag{44}$$

Taking conditional expectations,

$$E_t[\ln^2 S_{t+1}] \approx \ln^2 E_t[S_{t+1}] + \frac{1}{E_t[S_{t+1}]^2} V_t[S_{t+1}] + \frac{1}{4E_t[S_{t+1}]^4} E_t(S_{t+1} - E_t[S_{t+1}])^4 \tag{45}$$

$$- \frac{\ln E_t[S_{t+1}]}{E_t[S_{t+1}]^2} V_t[S_{t+1}] \tag{46}$$

$$- \frac{1}{E_t[S_{t+1}]^3} E_t(S_{t+1} - E_t[S_{t+1}])^3. \tag{47}$$

Rearranging the terms, we obtain

$$E_t[\ln^2 S_{t+1}] \approx \ln^2 E_t[S_{t+1}] + \frac{1 - \ln E_t[S_{t+1}]}{E_t[S_{t+1}]^2} V_t[S_{t+1}] - \frac{1}{E_t[S_{t+1}]^3} \mu_{3t} + \frac{1}{4E_t[S_{t+1}]^4} \mu_{4t}, \tag{48}$$

with $\mu_{3t} = E_t(S_{t+1} - E_t[S_{t+1}])^3$ and $\mu_{4t} = E_t(S_{t+1} - E_t[S_{t+1}])^4$.

We can write this expression as a function of the first four standardized central conditional moments (mean, variance, skewness and kurtosis) of S_{t+1}, that is,

$$E_t[\ln^2 S_{t+1}] \approx \ln^2 E_t[S_{t+1}] + \frac{1 - \ln E_t[S_{t+1}]}{E_t[S_{t+1}]^2} V_t[S_{t+1}] - \frac{V_t[S_{t+1}]^{3/2}}{E_t[S_{t+1}]^3} Skew_t(S_{t+1}) + \frac{V_t[S_{t+1}]^2}{4E_t[S_{t+1}]^4} Kurt_t(S_{t+1}), \tag{49}$$

with $Skew_t(S_{t+1})$ and $Kurt_t(S_{t+1})$ denoting the skewness and kurtosis parameters of the conditional distribution of the stock price S_{t+1}.

Now, using the above expressions, and (39) and (41), we can derive the expression for $E_t[r_{t+1}]$ and $V_t[r_{t+1}]$. More formally,

$$E_t[r_{t+1}] \approx \ln E_t[S_{t+1}] - \ln S_t - \frac{1}{2E_t[S_{t+1}]^2} V_t[S_{t+1}], \tag{50}$$

and

$$V_t[r_{t+1}] \approx \ln^2 E_t[S_{t+1}] + \frac{1 - \ln E_t[S_{t+1}]}{E_t[S_{t+1}]^2} V_t[S_{t+1}] - \frac{V_t[S_{t+1}]^{3/2}}{E_t[S_{t+1}]^3} Skew_t(S_{t+1}) + \frac{V_t[S_{t+1}]^2}{4E_t[S_{t+1}]^4} Kurt_t(S_{t+1}) \tag{51}$$

$$+ \ln^2 S_t - 2 \underbrace{\left(\ln E_t[S_{t+1}] - \frac{1}{2E_t[S_{t+1}]^2} V_t[S_{t+1}] \right)}_{E_t[\ln S_{t+1}]} \ln S_t \tag{52}$$

$$- \underbrace{\left(\ln E_t[S_{t+1}] - \ln S_t - \frac{1}{2E_t[S_{t+1}]^2} V_t[S_{t+1}] \right)}_{E_t[r_{t+1}]}^2. \tag{53}$$

Appendix B: Tables

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

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

Table 1: Summary statistics of the amount of options and by moneyness in % on each observation date for the period from January 1996 until April 2016. The amount of options vary quite strongly due to the increasing amount of available option prices over the past years.

Method	G in %			g		
	Mean	Median	Standard deviation	Mean	Median	Standard deviation
Mix-log	16.70	11.65	19.19			
Bin-Tree	2.13	0.93	5.27	17.93	5.33	88.05

Table 2: Risk neutral density estimation error for the mixture-lognormal density (mix-log) and binomial tree (bin-tree) for the period from January 1996 until April 2016.

μ	ϕ	ω	δ	α	β	d.o.f.
0.043	-0.031	0.018	0.903	0.000	0.164	8.785
0.006***	0.078*	0.000***	0.000***	0.000***	0.000***	0.000***
(0.012)	(0.155)	(0.003)	(0.009)	(0.000)	(0.014)	(0.986)

Table 3: Parameter estimates of the AR-TGARCH model in (34) representing the historical approach. P-values are labelled with *, **, and *** which indicate statistical significance at the 10%, 5% and 1% levels, respectively. Standard errors are reported in parentheses.

		Restricted VAR(1)				**Derived Model**			
x_t	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	$\begin{pmatrix} 0.000 \\ 0.003 \end{pmatrix}$	$\begin{pmatrix} 0.234 \\ 0.691 \end{pmatrix}$	$\begin{bmatrix} 2.456E-03 & -3.405E-04 \\ -3.405E-04 & 9.344E-05 \end{bmatrix}$	$\begin{pmatrix} 0.008 \\ 0.478 \end{pmatrix}$	0.002	0.691	$\begin{bmatrix} 2.456E-03 & -7.952E-05 \\ -7.952E-05 & 5.097E-06 \end{bmatrix}$	-0.711
	Mix-Log	$\begin{pmatrix} -0.002 \\ 0.003 \end{pmatrix}$	$\begin{pmatrix} 0.414 \\ 0.726 \end{pmatrix}$	$\begin{bmatrix} 2.449E-03 & -2.531E-04 \\ -2.531E-04 & 4.730E-05 \end{bmatrix}$	$\begin{pmatrix} 0.011 \\ 0.527 \end{pmatrix}$	0.002	0.726	$\begin{bmatrix} 2.449E-03 & -1.047E-04 \\ -1.047E-04 & 8.098E-06 \end{bmatrix}$	-0.744
	Bin-Tree	$\begin{pmatrix} -0.003 \\ 0.003 \end{pmatrix}$	$\begin{pmatrix} 0.399 \\ 0.730 \end{pmatrix}$	$\begin{bmatrix} 2.447E-03 & -2.765E-04 \\ -2.765E-04 & 5.566E-05 \end{bmatrix}$	$\begin{pmatrix} 0.012 \\ 0.532 \end{pmatrix}$	0.002	0.730	$\begin{bmatrix} 2.447E-03 & -1.102E-04 \\ -1.102E-04 & 8.845E-06 \end{bmatrix}$	-0.749
	Hist	$\begin{pmatrix} 0.002 \\ 0.000 \end{pmatrix}$	$\begin{pmatrix} 0.101 \\ 0.938 \end{pmatrix}$	$\begin{bmatrix} 2.466E-03 & 3.960E-06 \\ 3.960E-06 & 4.347E-07 \end{bmatrix}$	$\begin{pmatrix} 0.004 \\ 0.881 \end{pmatrix}$	0.002	0.938	$\begin{bmatrix} 2.466E-03 & 4.002E-07 \\ 4.002E-07 & 4.442E-09 \end{bmatrix}$	0.121
MPR	Log	$\begin{pmatrix} -0.005 \\ 0.047 \end{pmatrix}$	$\begin{pmatrix} 0.042 \\ 0.732 \end{pmatrix}$	$\begin{bmatrix} 2.454E-03 & -1.990E-03 \\ -1.990E-03 & 3.233E-03 \end{bmatrix}$	$\begin{pmatrix} 0.009 \\ 0.535 \end{pmatrix}$	0.002	0.732	$\begin{bmatrix} 2.454E-03 & -8.325E-05 \\ -8.325E-05 & 5.660E-06 \end{bmatrix}$	-0.706
	Mix-Log	$\begin{pmatrix} -0.008 \\ 0.047 \end{pmatrix}$	$\begin{pmatrix} 0.055 \\ 0.736 \end{pmatrix}$	$\begin{bmatrix} 2.450E-03 & -1.622E-03 \\ -1.622E-03 & 2.351E-03 \end{bmatrix}$	$\begin{pmatrix} 0.011 \\ 0.540 \end{pmatrix}$	0.002	0.736	$\begin{bmatrix} 2.450E-03 & -8.985E-05 \\ -8.985E-05 & 7.220E-06 \end{bmatrix}$	-0.676
	Bin-Tree	$\begin{pmatrix} -0.010 \\ 0.039 \end{pmatrix}$	$\begin{pmatrix} 0.071 \\ 0.773 \end{pmatrix}$	$\begin{bmatrix} 2.447E-03 & -1.261E-03 \\ -1.261E-03 & 1.478E-03 \end{bmatrix}$	$\begin{pmatrix} 0.012 \\ 0.596 \end{pmatrix}$	0.002	0.773	$\begin{bmatrix} 2.447E-03 & -9.001E-05 \\ -9.001E-05 & 7.530E-06 \end{bmatrix}$	-0.663
	Hist	$\begin{pmatrix} 0.002 \\ 0.024 \end{pmatrix}$	$\begin{pmatrix} 0.002 \\ 0.834 \end{pmatrix}$	$\begin{bmatrix} 2.466E-03 & 1.048E-03 \\ 1.048E-03 & 1.280E-03 \end{bmatrix}$	$\begin{pmatrix} 0.004 \\ 0.692 \end{pmatrix}$	0.002	0.834	$\begin{bmatrix} 2.466E-03 & 2.252E-06 \\ 2.252E-06 & 5.910E-09 \end{bmatrix}$	0.000
$d_t - p_t$		$\begin{pmatrix} 0.121 \\ -0.100 \end{pmatrix}$	$\begin{pmatrix} 0.030 \\ 0.975 \end{pmatrix}$	$\begin{bmatrix} 2.419E-03 & -2.419E-03 \\ -2.419E-03 & 2.721E-03 \end{bmatrix}$	$\begin{pmatrix} 0.023 \\ 0.950 \end{pmatrix}$	0.002	0.975	$\begin{bmatrix} 2.419E-03 & -7.141E-05 \\ -7.141E-05 & 2.371E-06 \end{bmatrix}$	-0.943

Table 4: Estimation results for the period from January 1996 until April 2016 for the SPX index using the risk-premium (RP), market price of risk (MPR) and log dividend price ratio ($d_t - p_t$) as state variable. The state variables are derived assuming the lognormal density (log), mixture-lognormal density (mix-log), binomial-tree (Bin-tree) and historical approach (hist).

x_t	Method	Mean optimal allocation in % $\alpha_t = [a_0 + a_1(\mu + \sigma_u^2/2)] \times 100$	Fraction hedging demand in % $[\alpha_{t,h}(\mu;\gamma,\phi)/\alpha_t(\mu;\gamma,\phi)] \times 100$
RP	Log	0.38	8.1
	Mix-Log	0.39	12.1
	Bin-Tree	0.39	12.2
	Hist	0.35	-0.2
MPR	Log	0.38	9.8
	Mix-Log	0.38	10.7
	Bin-Tree	0.37	9.5
	Hist	0.35	0.3
$d_t - p_t$		1.14	67.9

Table 5: The left column displays the mean optimal percentage allocation to the risky asset computed over the period January 1996 to April 2016. The right column panel displays the mean hedging demand relative to the mean total demand ($\alpha_{t,hedging}(\mu;\gamma_P,\Psi) = \alpha_t(\mu;\gamma_P,\Psi) - \alpha_{t,myopic}(\mu;\gamma_P,\Psi)$). All values are based on the full-sample estimates of the return processes for the risk premium (RP), market price of risk (MPR) and log dividend price ratio ($d_t - p_t$) state variables. The state variables are derived assuming the lognormal density (log), mixture-lognormal density (mix-log), binomial-tree (Bin-tree) and historical approach (hist).

x_t	Method	Inner-Sample				Out-of-Sample			
		Sharpe Ratio	Sortino Ratio	Value func. in %	Revenue in %	Sharpe Ratio	Sortino Ratio	Value func. in %	Revenue in %
RP	Log	0.148	0.135	0.191	9.293	0.757	1.055	0.127	4.461
	Mix-Log	0.134	0.124	0.229	13.770	0.666	1.999	0.155	8.172
	Bin-Tree	0.136	0.124	0.220	13.405	0.610	2.684	0.152	8.048
	Hist	0.233	0.198	0.145	0.635	0.953	1.043	0.111	2.083
MPR	Log	0.102	0.096	0.203	16.148	0.607	1.243	0.135	10.477
	Mix-Log	0.089	0.085	0.219	19.940	0.486	1.194	0.155	15.446
	Bin-Tree	0.117	0.110	0.191	14.129	0.291	0.559	0.148	12.634
	Hist	0.210	0.180	0.145	0.307	0.744	0.686	0.111	1.209
$d_t - p_t$		0.294	0.226	0.836	13.607	0.853	0.836	0.518	13.299

Table 6: The inner-sample results refers to the entire sample horizon from January 1996 to April 2016 over 243 periods. The out-of-sample evaluation is performed from January 2011 until April 2016 over 63 periods. The state variables are denoted as risk-premium (RP), market price of risk (MPR) and log dividend price ratio ($d_t - p_t$). The state variables are derived assuming the lognormal density (log), mixture-lognormal density (mix-log), binomial-tree (Bin-tree) and historical approach (hist).

		RP				MPR				$d_t - p_t$
		Log	Mix-Log	Bin-Tree	Hist	Log	Mix-Log	Bin-Tree	Hist	
RP	Log	-	0.79	0.86	0.58	0.37	0.38	0.41	0.65	0.85
	Mix-Log		-	0.98	0.79	0.79	0.27	0.32	0.80	0.96
	Bin-Tree			-	0.79	0.43	0.27	0.40	0.85	0.97
	Hist				-	0.98	0.90	0.98	0.11	0.50
MPR	Log					-	0.36	0.84	0.99	0.83
	Mix-Log						-	0.54	0.87	0.73
	Bin-Tree							-	0.97	0.83
	Hist								-	0.53
$d_t - p_t$										-

Table 7: The table shows the results of the Sharpe Ratio Test testing $H_0 : SR_1 = SR_2$ for the entire sample horizon (January 1996 to April 2016). The state variables are denoted as risk-premium (RP), market price of risk (MPR) and log dividend price ratio ($d_t - p_t$). The state variables are derived assuming the lognormal density (Log), mixture-lognormal density (Mix-Log), binomial-tree (Bin-Tree) and historical approach (Hist).

x_t	Method	Inner-Sample				Out-of-Sample			
		Hedging		No-Hedging		Hedging		No-Hedging	
		Timing	No-Timing	Timing	No-Timing	Timing	No-Timing	Timing	No-Timing
RP	Log	0.191	-33.4	-5.7	-33.9	0.127	-15.3	-2.0	-15.6
	Mix-Log	0.229	-44.4	-9.2	-44.8	0.155	-30.9	-4.7	-31.2
	Bin-Tree	0.220	-42.1	-9.2	-42.6	0.152	-29.4	-4.7	-29.7
	Hist	0.145	-12.7	0.1	-12.7	0.111	-4.9	-0.3	-4.9
MPR	Log	0.203	-37.4	-7.1	-37.8	0.135	-20.6	-3.0	-20.9
	Mix-Log	0.219	-42.1	-8.0	-42.5	0.155	-30.6	-4.6	-30.9
	Bin-Tree	0.191	-33.7	-6.8	-34.1	0.148	-28.1	-4.3	-28.3
	Hist	0.145	-12.5	0.0	-12.6	0.111	-3.0	-0.1	-3.1
$d_t - p_t$		0.836	-84.6	-66.0	-84.7	0.518	-80.4	-60.1	-80.5

Table 8: The panel displays results of the value function for restricted and unrestricted portfolio rules. For the restricted portfolio rules the relative change of the value function compared to the unrestricted rule is displayed. The inner-sample evaluation refers to the entire sample horizon and the unconditional mean implied by the risk-premium process. The out-of-sample evaluation uses the average realised return to calculate the value function of the investor. The state variables are denoted as risk-premium (RP), market price of risk (MPR) and log dividend price ratio ($d_t - p_t$). The state variables are derived assuming the lognormal density (log), mixture-lognormal density (mix-log), binomial-tree (Bin-tree) and historical approach (hist).

References

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

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

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

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

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

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

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

Bollerslev, T., M. Gibson, and H. Zhou (2011). Dynamic estimation of volatility risk premia and investor risk aversion from option-implied and realized volatilities. *Journal of Econometrics 160*(1), 235 – 245.

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

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

Campbell, J. Y. and R. Shiller (1988). Stock prices, earnings, and expected dividends. *Mathematical Finance* (XLIII), 661–676.

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

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

Cao, C., F. Yu, and Z. Zhong (2010). The information content of option-implied volatility for credit default swap valuation. *Journal of Financial Markets 13*(3), 321 – 343.

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

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

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

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

Fama, E. and K. French (1988). Dividend yields and expected stock returns. *Journal of Financial Economics* (XXII), 3–27.

Glosten, L. R., R. Jagannathan, and D. E. Runkle (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.

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

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

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

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

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

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

Liu, X., M. B. Shackleton, S. J. Taylor, and X. Xu (2007). Closed-form transformations from risk-neutral to real-world distributions. *Journal of Banking* & *Finance 31*(5), 1501–1520.

Melick, W. R. and C. P. Thomas (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.

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

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

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

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

Shackleton, M. B., S. J. Taylor, and P. Yu (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.

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), 6391.

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

Zakoian, J.-M. (1994). Threshold heteroskedastic models. *Journal of Economic Dynamics and Control* (18), 931–955.

Data citation: [Chicago Mercantile Exchange options and futures historical data]; 2020; Options and futures historical data; `https://optionmetrics.com/`