Market Timing with Option-Implied Distributions: A Forward-Looking Approach*

Alexandros Kostakis[a], Nikolaos Panigirtzoglou[b] and George Skiadopoulos[c]

First Draft: 1 April 2008 - This Draft: 15 February 2011
Longer Working Paper version of the *Management Science* publication

Abstract

We address the empirical implementation of the static asset allocation problem by developing a forward-looking approach that uses information from market option prices. To this end, constant maturity S&P 500 implied distributions are extracted and subsequently transformed to the corresponding risk-adjusted ones. Then, we form optimal portfolios consisting of a risky and a risk-free asset and evaluate their out-of-sample performance. We find that the use of risk-adjusted implied distributions times the market and makes the investor better off compared with the case where she uses historical returns' distributions to calculate her optimal strategy. The results hold under a number of evaluation metrics and utility functions and carry through even when transaction costs are taken into account. Not surprisingly, the reported market timing ability deteriorated during the recent subprime crisis. An extension of the approach to a dynamic asset allocation setting is also presented.

JEL Classification: C13, G10, G11, G13.

Keywords: Asset allocation, Option-implied distributions, Market timing, Performance evaluation, Portfolio Choice, Risk aversion.

* We are grateful to Wei Xiong (the Department Editor), an anonymous Associate Editor and two anonymous referees for their constructive, stimulating and thorough comments. We would also like to thank Carol Alexander, Timotheos Angelidis, Victor DeMiguel, Ozgur Demirtas, Daniel Giamouridis, Wolfgang Härdle, Jens Jackwerth, Chris Kind, Eirini Konstantinidi, Kostas Koufopoulos, Andrew Patton, Jacques Pezier, Michael Rockinger, Antonis Sangvinatsos, Pedro Santa-Clara, Georgios Skoulakis, Luis Viceira, Grigory Vilkov, Raman Uppal, Zvi Wiener and participants at the Aarhus Business School, Athens University of Economics and Business, Cass Business School, Cyprus University of Technology, ICMA University of Reading, ISCTE/Nova de Lisboa, Humboldt University (CASE), TU Munich University, VU University of Amsterdam, University of Piraeus, University of Warwick seminar series and the 2008 Bachelier World Congress (London), 2009 MathFinance Conference (Frankfurt), 2009 EFMA Conference (Milan), 2009 FMA European Conference (Turin), 2009 Inquire UK (Leeds), 2010 EFA Conference (Frankfurt) and 2011 HEC Lausanne Asset and Risk Management Conference for many helpful discussions and comments. Financial support from the Research Centre of the University of Piraeus and the Fundação Ciência e Tecnologia's grant number PTDC/EGE-ECO/099255/2008 is gratefully acknowledged. Previous versions of the paper were circulated under the title "Asset allocation with option-implied distributions: A forward-looking approach". Any remaining errors are our responsibility alone.

[a] Liverpool Management School, University of Liverpool, e-mail: A.Kostakis@liverpool.ac.uk
[b] Queen Mary, University of London, Department of Economics, Mile End Road, London E1 4NS. e-mail: n.panigirtzoglou@qmul.ac.uk
[c] Corresponding author. Department of Banking and Financial Management, University of Piraeus and Financial Options Research Centre, Warwick Business School, University of Warwick, gskiado@unipi.gr

The standard static optimal portfolio selection problem boils down to maximizing the expected utility derived by the one-period-ahead wealth. Maximisation of expected utility can be carried out in two alternative ways. The first obvious one is by performing a direct utility maximisation (e.g., Adler and Kritzman (2007) and Sharpe (2007)). The second is by maximising a Taylor series expansion up to a certain order that approximates expected utility (see e.g., Kroll, Levy and Markowitz (1984), Jondeau and Rockinger (2006), Guidolin and Timmermann (2008), Garlappi and Skoulakis (2009) and references therein). This approach results in portfolio choice based on some moments of the returns' distribution; the mean-variance optimization à la Markowitz is the most popular example. Implementation of the two routes requires estimation of the portfolio returns probability density function (PDF) and its moments, respectively. To this end, the literature has so far used historical data (*backward-looking approach*, see e.g., DeMiguel, Garlappi and Uppal (2009), for a review of various historical estimators). Inevitably, the issue of estimation error in the inputs of expected utility maximisation arises (e.g., Merton (1980), Chan, Karceski and Lakonishok (1999)) and the optimal portfolio may be mis-calculated (see e.g., Chopra and Ziemba (1993), Kan and Zhou (2007) and references therein). Mis-calculation of the optimal portfolio reduces investor's utility (see Siegel and Woodgate (2007)).

To avoid the use of historical PDFs, this paper takes a very different approach and develops an empirical procedure to using stock index implied distributions as inputs to calculate the optimal portfolio. By definition, implied distributions are extracted from the market option prices that reflect the market participants' expectations; they refer to the distribution of the asset price that serves as underlying to the option. The horizon of the distribution matches the expiry date of the option. Therefore, the appeal of the suggested approach is that implied distributions are inherently forward-looking and may serve as more accurate forecasts of the moments/distribution that will yield economic value to an investor in an asset allocation setting (this is known as moments timing or more generally distributional timing, see Jondeau and Rockinger, 2008, and the references therein). The suggested *forward-looking approach* can be viewed as a generalisation of the literature that suggests forecasting volatility by the implied volatility (the second moment of the implied distribution) rather than backward-looking measures of volatility that use historical data (see Poon and Granger (2003) for a review of this literature). It can also be viewed as part of the literature that suggests using information from option prices rather than historical data to estimate parameters that are of crucial importance to quantify risk and perform asset allocation such as the risk-premium (Duan and Zhang, 2010), beta (see e.g., Siegel, (1995), Chang, Christoffersen, Jacobs and Vainberg (2009)) and correlation coefficients (see e.g., Driessen, Maenhout and Vilkov (2009)), as well as to forecast future returns of financial assets and growth in real economic activity (Bakshi, Panayotov and Skoulakis, (2010), Cremers and Weinbaum, (2010), Golez, (2010), Xing, Zhang and Zhao, (2010)); see also Giamouridis and Skiadopoulos (2010) for a review.

We consider an asset universe that consists of a risky (the S&P 500 index) and a riskless asset. This setup has been commonly used in the literature (see e.g., Wachter (2002) and Chacko and Viceira (2005)) and is also encountered in practice.[1] First, we extract constant maturity one-month S&P 500 implied PDFs by applying the method of Bliss and Panigirtzoglou (2002) that has been found to be robust to the presence of measurement errors in the data. Then, we convert them to the corresponding risk-adjusted ones by employing the approach of Bliss and Panigirtzoglou (2004).[2] This transformation is dictated by financial theory because the implied distributions are measured under the risk-neutral probability measure while the calculation of optimal portfolios requires the real-world (also termed physical) PDF. Next, we use the risk-adjusted S&P 500 implied distributions to calculate the optimal portfolio. Finally, we compare the out-of-sample performance of the derived optimal

[1] Given that the risky asset under consideration is an index, this seemingly simple setup is in fact quite general and has a number of practical applications. In particular, there are three practical situations where an institutional investor invests in an asset universe that consists of a single risky and one risk-free asset. First, any index fund manager invests in cash (earning the risk-free rate) and the index; the cash is used to deal with redemptions, fixed costs and other operating expenses. Second, portfolio insurance strategies (e.g., constant proportional portfolio insurance) are a typical example. Third, and more generally, any management of market exposure (e.g., beta for equities, duration for bonds) can be understood as an asset allocation process between that risky asset and the risk-free asset. These strategies are known as “tactical asset allocation products” and have become quite popular recently. Examples for these types of products are index replicating portfolios with a cap for volatility or an outright target for volatility (e.g., the S&P 500 and EURO-STOXX 50 risk-control indices). In addition, this setting applies also to the case of individual investors due to the spectacular growth of exchange traded funds (ETFs). A number of popular ETFs are written on a market index, so an individual investor can allocate her funds between cash (earning the risk-free rate) and the ETF. These examples are in line with the spirit of Tobin’s two-fund separation theorem which is a building block in modern financial theory and still serves as a benchmark for investor’s behaviour and fund management practice despite the potential arguments against its validity.

[2] The term "risk-adjusted" is used to remind that risk-preferences are embedded and to distinguish it from the term "historical distribution"; the latter is used to define the PDF estimated solely from time series of asset prices.

strategies based on the risk-adjusted implied distributions/moments with that of the optimal strategies based on historical distributions/moments.

To check the robustness of the obtained results and shed light on whether implied distributions should be preferred to backward-looking ones for asset allocation purposes, we conduct a number of robustness tests. First, the risk-adjustment of implied distributions is performed by assuming alternative utility functions (exponential and power) for the representative (average) agent. Second, the optimal portfolios are calculated by maximising the expected utility per se and its truncated Taylor series expansion, separately. This is to check whether the use of a moment-based rule (e.g., the popular mean-variance analysis) will affect the properties of the derived optimal portfolios (see e.g., Jondeau and Rockinger (2006) for a comparison of the optimal portfolios derived by direct and Taylor series expansion maximisation in an in-sample historical estimators setting). Third, various utility/value functions and degrees of risk aversion that describe the preferences of the marginal (individual) investor are employed. The rationale justifying these partial-equilibrium exercises is that there exists an individual investor who is price-taker, i.e. takes these already extracted distributions as exogenously given and maximizes her own utility without affecting market prices because she only holds a small portion of the market wealth. In line with the existing asset allocation literature, the individual investor, whose portfolio choice we examine, is distinct from the representative agent. Standard and behavioral utility functions are used. In particular, exponential and power utility functions as well as the disappointment aversion setting introduced by Gul (1991) are employed. The latter has been used to explain investors' behavior with respect to their stock holdings (see e.g., Barberis, Huang and Santos (2001) and Ang, Bekaert and Liu (2005)) and option holdings (see Driessen and Maenhout (2007)). In particular, we employ a kinked value function to examine whether our results are robust in the presence of loss aversion.

Furthermore, we use a number of measures (Sharpe ratio, opportunity cost, portfolio turnover and risk-adjusted returns net of transaction costs) to assess the optimal forward-looking portfolio's performance with that of the backward-looking one. To this end, we also employ alternative ways of estimating the historical distribution of returns. Finally, we study the impact of the recent 2007-2009 subprime crisis on the reported findings. Interestingly, the suggested approach can be also extended to a multi-period asset allocation setting. We provide such an extension within Wachter's (2002) setting. The solution of the portfolio choice for a risk-averse, long-term investor who maximizes expected utility over terminal wealth is presented. Moreover, we analytically discuss how the extracted risk-adjusted option-implied distributions can be utilized in such a framework and present optimal myopic and hedging demands for various levels of risk aversion and investment horizons.

We conclude this introduction by discussing the relation of our work to the existing literature. There is already a significant literature on methods to extract implied PDFs as well as their potential applications to policy-making (see e.g., Söderlind and Svensson (1997)), option pricing and risk management (Ait-Sahalia and Lo (2000), Panigirtzoglou and Skiadopoulos (2004), Alentorn and Markose (2008)) and forecasting the future value of the underlying asset (Bliss and Panigirtzoglou (2004), Anagnou-Basioudis, Bedendo, Hodges and Tompkins (2005), Kang and Kim (2006), and Liu, Shackleton, Taylor and Xu (2007)). Jackwerth (2004) also provides an excellent review of the applications of implied distributions. However, their use for asset allocation purposes has not yet been considerably explored. There are two papers that are related to our study. Concurrently but independently, Ait-Sahalia and Brandt (2008) propose the use of implied PDFs to solve the intertemporal consumption and portfolio choice problem within the martingale approach setting of Cox and Huang (1989). However, theythe optimal portfolio choice problem is not addressed; only the properties of the derived optimal consumption paths are examined.[3] Very recently and subsequently to our paper, DeMiguel, Plyankha, Uppal and Vilkov (2010) have presented a study that is closer to ours. They consider an asset universe consisting of a large number of US stocks to assess whether the use of option implied moments can improve the out-of-sample performance of the formed optimal static portfolios. They find that the use of implied volatilities that are risk-adjusted by either the volatility risk premium or the option implied skewness yields optimal portfolios that earn greater Sharpe ratios (accompanied with higher portfolio turnover though) than the portfolios based on historical information. However, there are two important differences between this work and ours: first, the authors constrain their

[3]Jabbour, Pena, Vera and Zuluaga (2008) also use information from option prices to construct optimal portfolios. However, their definition of optimality is not in terms of maximising expected utility. Instead, the optimal portfolio is defined as the one that minimises the Conditional Value-at-Risk. This definition may be restrictive since it does not capture all the characteristics of the utility function of the investor. In addition, their study focuses on the properties of the suggested algorithm and does not provide further tests on its out-of-sample performance relative to a method that uses historical data to calculate the optimal portfolio.

analysis in a mean-variance setting, while we take into account the whole PDF to calculate the optimal portfolios and our approach can be extended to a multi-period setup. Second, their risk adjustment of implied volatilities, albeit innovative, is not grounded on financial theory.

The rest of the paper is structured as follows. Section I outlines the methodology to find the optimal portfolio by direct maximisation. Section II describes the data sets, the method to extract the implied distributions, and how their risk-adjusted analogues are derived. The following Section explains the implementation of the forward and backward-looking approach and discusses their relative performance under a number of metrics. Section IV presents some further robustness tests. Sections V and VI investigate the effect of loss aversion and sources for the discrepancy in the performance of the two approaches, respectively. Section VII repeats the analysis over the 2007-2009 period. Section VIII provides the extension of the proposed approach to a dynamic asset allocation setting. The last Section concludes and presents the implications of this study, as well as, suggestions for future research.

I. Calculating the Optimal Portfolio

Consider a risk-averse investor with utility function $U(W)$ where $U''(W) < 0 \; \forall \, W$. At any point in time t, the investor decides about her optimal allocation of wealth W_t between a risky and a riskless asset over the period [t, t+1] (static allocation problem). To fix ideas, let the return of the risky and the riskless asset from time t to t+1 be r_{t+1} and $r_{f,t+1}$ respectively. Let also the weights of wealth invested in the risky and the riskless asset at time t over the next period be α_t and α_t^f, respectively, where $\alpha_t + \alpha_t^f = 1$. Then, the optimal portfolio at time t is constructed by maximising the expected utility of wealth at time t+1 with respect to the portfolio weights, i.e.

$$\max_{\alpha_t} E[U(W_{t+1})] \tag{1}$$

where

$$W_{t+1} = W_t(1 + \alpha_t r_{t+1} + \alpha_t^f r_{f,t+1}) \tag{2}$$

Without loss of generality, initial wealth is normalised to one, i.e. W_t=1. Therefore,

$$W_{t+1} = 1 + \alpha_t r_{t+1} + \alpha_t^f r_{f,t+1} \tag{3}$$

At any point in time t, the problem of the direct maximization of the expected utility is defined as:

$$\max_{\alpha_t} E[U(W_{t+1})] = \max_{\alpha_t} E[U(1+\alpha_t r_{t+1} + \alpha_t^f r_{f,t+1})] \\ = \max_{\alpha_t} \int U(1+\alpha_t r_{t+1} + \alpha_t^f r_{f,t+1}) dF(r_{t+1}) \tag{4}$$

$$s.t.\ \alpha_t + \alpha_t^f = 1 \tag{5}$$

where $F(\bullet)$ is the cumulative real-world conditional distribution function (CDF) of the return of the risky asset r_{t+1} at time $t+1$; the CDF depends only on the return of the risky asset, since $r_{f,t+1}$ is known ex ante (at time t). Portfolio weights for the risky asset are constrained in the interval [-1, 2], i.e. leverage up to only 100% is allowed. This is a realistic assumption for the type of asset universe we consider.

II. The Dataset

The data set consists of S&P 500 futures options monthly closing prices (January 1986 through December 2009) traded on the Chicago Mercantile Exchange (CME). The analysis will be conducted in two stages. First, the period January 1986 to August 2007 shall be employed. This is a period that includes both bearish and bullish regimes as well as the dawn of the recent sub-prime crisis (see Brunnermeier, 2009). Next, the impact of the period September 2007 – December 2009 where the crisis became pronounced shall be studied. The CME S&P 500 options contract is an American style futures option; the underlying futures is the CME S&P 500 futures contract. The expiry dates of the S&P 500 options coincide with these of the futures contracts; these trade out to one year with expiries in March, June, September, and December. In addition, there are monthly serial options contracts out to one quarter; these were introduced in 1987. Options and futures expire on the third Friday of the expiry month. For serial months there is no corresponding futures expiry and the options settle on the closing price of the S&P 500 futures contract that expires next or just after the options expiry. The associated value of the underlying is the settlement price of the S&P 500 futures contract maturing on or just after the option expiry date. The risk-free rate used in this study is the one-month LIBOR rate taken from Bloomberg. The dividend yield is calculated as the twelve-month rolling dividends per share divided by the stock index price obtained from Datastream.

A. Extracting the Implied Distribution

We estimate the implied PDFs using the non-parametric method suggested by Bliss and Panigirtzoglou (2002) and currently used by the Bank of England. This method is chosen because they document that it generates PDFs that are robust to quite significant measurement errors in the quoted option prices. The technique uses the Breeden and Litzenberger's (1978) non-parametric result and employs a natural cubic spline to fit implied volatilities as a function of the deltas of the options in the sample.

In particular, Breeden and Litzenberger (1978) show that assuming that option prices are observed across a continuum of strikes, the second derivative of a European call price with respect to the strike price delivers the risk neutral PDF. However, in practice, available option quotes do not provide a continuous call price function. To construct such a function, a natural cubic spline is used to interpolate across implied volatilities (see also Jiang and Tian (2007) for a similar choice).[4] In addition, it is necessary to extrapolate the spline beyond the range of available implied volatilities so as to extract the tails of the PDF. To this end, we force the spline to extrapolate smoothly in a horizontal manner (see also Jiang and Tian, (2005), Carr and Wu (2009), for a similar choice).[5] We do this by introducing two pseudo-data points spaced three strike intervals above and below the range of strikes in the cross sections and set implied volatilities equal to the implied volatilities of the respective extreme-strike options. These pseudo-data points are added to the cross sections before spline-fitting takes place. Extrapolating the implied volatility function in this manner has the effect of smoothly pasting log-normal tails onto the implied density function beyond the range of traded strikes.

We calculate implied volatilities from option prices by using the analytical quadratic approximation of Barone-Adesi and Whaley (BAW, 1987) in order to capture the early exercise premium of the American-style S&P 500 futures options. In addition, the implied volatility calculated via the BAW formula can be inserted in Black's (1976) formula to calculate the European option prices (see BAW, 1987, for a discussion). Hence, Breeden and

[4]To fit the natural cubic spline to implied volatilities, a value for the smoothing parameter of the spline needs to be chosen. We choose a value of 0.99 that yields well-behaved PDFs and fits option prices well (see Bliss and Panigirtzoglou (2002, 2004) for an extensive discussion). Moreover, Bliss and Panigirtzoglou (2004) and Kang and Kim (2006) find that the forecasting performance of the implied PDF does not depend on the smoothing parameter for a wide range of values.

[5] We have also used the smooth pasting condition proposed by Jiang and Tian (2007) for extrapolation purposes. However, this delivered implausibly large implied volatilities.

Litzenberger's (1978) result can also be applied to our American option dataset despite the fact that it was derived for European options.[6] The delta metric is constructed by converting strikes into their corresponding call deltas by using the at-the-money implied volatility.[7] Hence, a set of implied volatilities and corresponding deltas is constructed for each available contract.

For the purposes of calculating the implied volatilities, we impose the standard filtering constraints. Only at-the-money and out-of-the-money options are used because they are more liquid than in-the-money. Hence, measurement errors in the calculation of implied volatilities due to bid-ask spreads and non-synchroneous trading (Harvey and Whaley (1991)) are less likely to occur. In addition, we discard option prices that violate Merton's (1973) arbitrage bounds and option prices with less than five working days to maturity since they are excessively volatile as market participants close their positions. Implied volatilities of deltas greater than 0.99 or less than 0.01 are also eliminated. These volatilities correspond to far out-of-the-money call and put prices, which have generally low liquidity. An implied volatility curve is constructed if there are at least three implied volatilities, with the lowest delta being less than or equal to 0.25 and the highest delta being greater than or equal to 0.75. This ensures that the available strikes cover a wide range of the PDF available outcomes. In the case that the range of strikes does not spread along the required interval, no PDF is

[6]Inserting the BAW implied volatilities in Black's (1976) rather than in BAW model does not affect the derived probabilities. This is because the size of the early exercise premium is very small in our case, since only short maturity (less than six months), out-of-the money options are used, and the cost of carry of the underlying asset is zero. BAW (1987) illustrate that out-of-the money options have very small early exercise premiums of the order of 0.01 (see Tables II and III in their paper, pages 313 and 314, respectively). This small size becomes even more insignificant when compared with the tick size error (0.05 for the S&P 500 futures options used in the paper). Moreover, in the case that the cost of carry is zero (Table III) the early exercise premium is smaller as compared to a 4% cost of carry case (Table II). They also show that the early exercise premium decreases as the time-to-maturity decreases. Therefore, the effect of the adjustment is very small on the option prices, and hence on the derived probabilities.

[7]The (call) delta metric is preferred to strike (or moneyness metric) because it takes values between zero and one irrespectively of the maturity of the contract; this is in contrast to the range of strikes that varies with the maturity widely. In addition, it is well known that the interpolated implied volatilities are more stable under a delta than a strike metric. A small delta corresponds to a high strike (i.e. out-of-the-money calls), while a large delta corresponds to a low strike (i.e. in-the-money calls). Black's (1976) model is used to calculate deltas. In line with Bliss and Panigirtzoglou (2002, 2004) and Liu, Shackleton, Taylor and Xu (2007), we use the at-the-money implied volatility so as the ordering of deltas is the same as that of the strikes. Using the implied volatilities that correspond to each strike could change the ordering in the delta space in cases where steep volatility skews are observed. This would result in generating volatility smiles with artificially created kinks.

extracted. Once the spline is fitted, 5,000 points along the function are converted back to option price/strike space using Black's (1976) model. The 5,000 call price/strike data points are used to differentiate twice the call price function numerically so as to obtain the estimated PDF.

For the purposes of forming optimal portfolios, we construct constant one-month maturity implied PDFs using the methodology described in Panigirtzoglou and Skiadopoulos (2004). This is done as follows. First, the implied volatility curve of a synthetic constant one-month maturity option contract is constructed. This is done in three steps. First, for each expiry contract, a spline interpolation is performed across implied volatilities as a function of delta. Implied volatilities corresponding to nine values of delta (ranging from 0.1 to 0.9) are retained. Next, spline interpolation is applied across the implied volatilities of contracts with different maturities for any one of the nine values of delta; the one-month maturity implied volatilities are picked. In the final step, once this discrete constant one-month maturity has been obtained (nine implied volatility points corresponding to nine deltas), a continuous implied volatility function is constructed by spline interpolating across these nine deltas. Finally, we back out the constant one-month maturity implied PDF by following the Bliss and Panigirtzoglou (2002) method described above. This exercise is repeated at the end of each month.

A final point to be taken into account is that in the case of the S&P 500 futures options, the extracted implied distributions are measured in the space of the variable

$$x = \frac{F_{T,T}}{F_{t,T}} - 1 = \frac{S_T}{F_{t,T}} - 1 \tag{6}$$

where $F_{t,T}$ is the price at time t of the futures contract on the S&P 500 that matures at T=1 month. However, for the purposes of our analysis, we are interested in measuring implied distributions in the space

$$y = \frac{S_T}{S_t} - 1 \tag{7}$$

To switch from the x to y consistently, we use the no-arbitrage formula

$$S_t = \frac{F_{t,T}}{1 + (r_{f,t+1} - d_t) \times 1/12} \tag{8}$$

where d_t is the dividend yield at time t. Plugging equation (8) in equation (7) yields

$$y = (1+x)\times[1+(r_{f,t+1}-d_t)\times 1/12]-1 \quad (9)$$

B. Risk-adjusting the Implied Distributions

There is a subtle point in the case where risk-neutral densities are used for asset allocation purposes. Option implied distributions are formed under the risk-neutral probability measure. Hence, the option risk-neutral densities need to be risk-adjusted so as to be converted to the corresponding actual probability measure distributions required to calculate optimal portfolios (equation (4)). The transformation uses the well-known link between the measured at time t risk-neutral distribution $q_t(S_T)$ and statistical distribution (also termed real-world, actual, or physical) $p_t(S_T)$ of the asset price S_T at time T ($t \leq T$). To fix ideas, assume that a representative agent with utility function $U(\cdot)$ exists. Then,

$$q_t(S_T) = \zeta_t(S_T)\times p_t(S_T) \quad (10)$$

where

$$\zeta_t(S_T) \equiv \exp[-r(T-t)]\frac{U'(S_T)}{U'(S_t)} \quad (11)$$

$\zeta_t(S_T)$ is the so-called pricing kernel. Equation (11) is derived by the first-order condition of the intertemporal expected utility maximisation problem of the representative agent (see also Ait-Sahalia and Lo (2000) for a detailed discussion). Equation (10) shows that given a utility function and the risk-neutral probabilities for the asset price returns, the corresponding risk-adjusted probabilities can be derived; the adjustment is non-linear and hence cannot be done by simply adding an econometrically estimated risk premium to every point of the implied PDF. The resulting risk-adjusted density function must be normalised to integrate to one. Hence, equations (10) and (11) yield

$$p_t(S_T) = \frac{\dfrac{q_t(S_T)}{\zeta_t(S_T)}}{\int \dfrac{q_t(x)}{\zeta_t(x)}dx} = \frac{\dfrac{q_t(S_T)}{U'(S_T)}}{\int \dfrac{q_t(x)}{U'(x)}dx} \quad (12)$$

To risk-adjust the risk-neutral densities [equation (12)] an assumption about the utility function of the representative agent needs to be made. We assume either one of the two most commonly used in the finance literature utility functions: (1) the negative exponential utility function, and (2) the power utility function.[8] The negative exponential utility function is defined as

$$U(W) = -\exp(-\eta W)/\eta, \quad \eta \neq 0 \tag{13}$$

where η is the coefficient of absolute risk aversion (ARA). The power utility function is defined as

$$U(W) = \frac{W^{1-\gamma} - 1}{1-\gamma}, \quad \gamma \neq 1 \tag{14}$$

where γ is the coefficient of constant relative risk aversion (RRA).

Both utility functions and thus the corresponding risk-adjusted densities depend on the value of the single parameter η (γ) that has an economic interpretation. We follow Bliss and Panigirtzoglou (2004) to determine this parameter in a three-step procedure. First, a sample of monthly fixed-expiry risk-neutral PDFs is extracted from the market option prices. Then, the extracted risk-neutral PDFs are converted to the corresponding subjective risk-adjusted PDFs for any given value of the single parameter η (γ). Finally, we find the value η^* (γ^*) of the risk aversion parameter that maximizes the forecasting ability of the risk-adjusted PDFs with respect to future realizations of the underlying index, i.e. the *p*-value of Berkowitz (2001) likelihood ratio statistic; the implicit assumption is that investors form rational expectations. This optimal value determines the (implied) risk aversion coefficient.[9] The coefficient η^* (γ^*) can be interpreted as the "average market" risk-aversion parameter for the sample time period considered.

For the purposes of our analysis, we derive a time series of η^* (γ^*). This is done by repeating the above three-step procedure on a monthly basis using a rolling window of K

[8]More flexible functional forms may be alternatively used for the utility function of the representative agent (see e.g. Kang and Kim (2006)). Equivalently, a more flexible specification for the pricing kernel may be adopted (see e.g., Rosenberg and Engle (2002)). However, these specifications have not been used in an asset allocation setting partly because the economic interpretation of their extra parameters is not obvious. Therefore, we employ the widely used power and exponential utility functions to risk-adjust option implied distributions as in Bliss and Panigirtzoglou (2004).

[9]In general, the risk-neutral PDF, the physical one, and the (differentiable) utility function of the representative agent are linked; the knowledge of any two of the three quantities delivers the third one. Therefore, the implied risk aversion can also be derived by knowledge of the risk-neutral PDF and the physical one (see e.g., Ait-Sahalia and Lo (2000) and Jackwerth (2004)). However, this approach is not applicable in our case since we are in search of the risk-adjusted physical PDF. Hence, we use the implied distribution and an assumed utility function in order to extract the corresponding risk-adjusted physical PDF.

monthly fixed-expiry risk-neutral PDFs and monthly realizations of the underlying index. That is, at each point in time t, we employ a time series of K monthly fixed-expiry risk-neutral PDFs (extracted on the dates from t - K to t -1) and their corresponding index realizations to estimate η^* (γ^*). Then, we use this estimated value η^* (γ^*) to risk-adjust the constant one month-maturity risk-neutral density extracted at time t; this derives the (risk-adjusted) subjective PDF over the t to t+1 horizon that will be used for the direct expected utility maximization [equation (4)].

Our methodology ensures that only information known to investors up to time t is employed to derive the risk-adjustment parameter η* (γ*), i.e. only the most recent t-K to t data on option prices and index realizations are used to adjust the constant one-month maturity risk-neutral density over the period between t and t +1. This will enable the subsequent evaluation of the suggested forward-looking asset allocation approach in an out-of-sample setting. The resulting time series of η^* (γ^*) is calculated by using alternative rolling windows of K=36,48,60, 72 monthly observations until we exhaust the whole sample. We consider alternative rolling windows of different sizes so as to check the robustness of our subsequent results to the choice of the rolling window that will be used to derive the risk-adjusted PDF.

III. Optimal Portfolios: Historical versus Implied Distributions

A. Implementation

In the case of direct maximisation [equation (4)], the CDF $F(r_{t+1})$ of the risky asset returns needs to be estimated to determine the optimal α_t at any point in time t. Two alternative "estimators" are compared: the empirical distribution estimated from monthly historical data up to time t (termed historical distribution), and the risk-adjusted implied distribution extracted from option prices at time t with expiry date at time t+1 -i.e. one month ahead expiry. Following Ait-Sahalia and Lo (2000), we estimate the historical distribution by means of a Gaussian kernel. To calculate the optimal portfolio, a grid search is performed.

Then, we follow a "rolling-window" procedure to compare the *out-of-sample* performance of the forward-looking approach to asset allocation with the historical backward-looking one. At any given point in time t, the optimal portfolio weights are determined by the forward and backward-looking estimators separately by maximising the expected utility; in the case of the backward-looking estimator, K=36,48,60, 72 monthly historical data up to time t are used. Next, we form the corresponding optimal portfolios and calculate the out-of-sample portfolio monthly return over the period [t, t+1]. This process is repeated (i.e. we rebalance the portfolio) until the end of the data set is reached; again, in the case of the historical estimator, a moving window of K monthly historical data is used so as to re-calculate the central moments of the updated dataset. Eventually, a time series of one-month out-of-sample portfolio returns is generated based on any given approach to estimating the required inputs to maximise expected utility.

Finally, an assumption about the utility function that describes the preferences of the individual investor needs to be made in order to find the optimal portfolio. We consider two alternative standard utility functions: the negative exponential utility function and the power utility function [equations (13) and (14), respectively]. In line with Jondeau and Rockinger (2006), we employ a grid search over possible values of the risky and risk-free asset weights to perform the direct maximisation [equation (4)].

B. Evaluation Metrics

We evaluate the alternative methodologies (i.e. option-implied distribution versus the historical one) in terms of certain characteristics of the respective optimal portfolios that are obtained out-of-sample. To this end, we use the Sharpe ratio (SR), the concept of opportunity cost, the portfolio turnover and a measure of the portfolio risk-adjusted returns net of transaction costs. The comparison of the backward and forward-looking approaches is carried out for any given expected utility function to be maximised when the risk-adjustment has been performed by the given utility function.

The SR is used to compare the risk-adjusted performance of the alternative investments during the whole time period (from t=1 to T) in line with the finance industry practice. The concept of opportunity cost has been introduced by Simaan (1993) to assess the economic significance of the difference in the performance of the two strategies (see also Jondeau and Rockinger (2006)). To fix ideas, let α^{imp} be the optimal portfolio choice derived by using the implied distribution approach. Similarly, let α^{hist} be the optimal portfolio choice that is obtained by employing the historical distribution. Denote by r_p^{imp} and r_p^{hist} the corresponding realized portfolio returns. The opportunity cost c is defined to be the return that needs to be added (or subtracted) to the one obtained by the strategy based on the historical distribution so as the investor becomes indifferent (in utility terms) between the two strategies, i.e.

$$E[U(1+r_p^{hist}+c)] = E[U(1+r_p^{imp})] \tag{15}$$

Therefore, in the case where the opportunity cost is positive (negative) the investor will be better (worse) off by adopting the risk-adjusted implied rather than the historical distribution as an input to calculate her optimal portfolio. Note that there is not necessarily a one-to-one correspondence between the SR and the opportunity cost. This is because the SR is a mean-variance measure while the opportunity cost is based on the assumed utility function and, hence, it takes into account the higher order moments of the portfolio returns distribution, too.

We compute the portfolio turnover (*PT*) to get a feel of the degree of rebalancing required to implement each one of the two strategies. In line with DeMiguel, Garlappi and Uppal (2009), for any portfolio strategy *k*, PT_k is defined as the average absolute change in the weights over the *T*-1 rebalancing points in time and across the *N* available assets (two in our case), i.e.

$$PT_k = \frac{1}{T-1}\sum_{t=1}^{T}\sum_{j=1}^{N}\left(\left|a_{k,j,t+1} - a_{k,j,t^+}\right|\right) \tag{16}$$

where $a_{k,j,t}$ is the portfolio weight in asset *j* at time *t* under strategy *k*, $a_{k,j,t+1}$ is the desired (based on the optimisation of expected utility) portfolio weight in asset *j* at time *t*+1 under strategy *k*, and a_{k,j,t^+} is the portfolio weight *before* rebalancing at *t*+1. For example, in the case of the 1/*N* strategy (i.e. 50% of the wealth invested in the risky asset and 50% of the wealth invested in the riskless asset), $a_{j,t} = a_{j,t+1} = 1/N$, but a_{k,j,t^+} may be different due to changes in asset prices between *t* and *t*+1. The *PT* quantity defined above can be interpreted as the average fraction (in percentage terms) of the portfolio value that has to be reallocated over the whole period.

Finally, we evaluate the historical and implied distributions strategies under the risk-adjusted, net of transaction costs, return-loss measure of DeMiguel, Garlappi and Uppal (2009). This measure provides an economic interpretation of the *PT* metric; it shows how the proportional transaction costs generated by the portfolio turnover affect the returns from any given strategy. To fix ideas, let *pc* be the proportional transaction cost. In the case where the portfolio is rebalanced, the total proportional cost is given by $pc \times \sum_{j=1}^{N}(\left|a_{k,\,j,\,t+1} - a_{k,\,j,\,t^+}\right|)]$. The evolution of the net of transaction costs wealth (NW_k) for strategy k is given by:

$$NW_{k,t+1} = NW_{k,t}(1 + r_{k,p,t+1})[1 - pc \times \sum_{j=1}^{N}\left|a_{k,j,t+1} - a_{k,j,t^+}\right|] \tag{17}$$

Then, the Return Net of Transaction Costs $RNTC_{k,t+1}$ for strategy k at time t+1 is given by:

$$RNTC_{k,t+1} = \frac{NW_{k,t+1}}{NW_{k,t}} - 1 \tag{18}$$

To calculate $NW_{k,t+1}$, we assume the proportional transaction cost *pc* for the S&P 500 (risky asset) to be equal to 50 basis points per transaction, as assumed in DeMiguel, Garlappi and Uppal (2009) and documented in references therein. On the other hand, *pc* is set equal to zero for the risk-free asset; this is a legitimate assumption since in practice no transaction fees are charged in the case where the investor deposits or withdraws an amount from the risk-free savings account.

The return-loss measure is calculated with respect to the implied distribution based strategy; it is defined as the additional return needed for the historical distribution based strategy to perform as well as the implied distribution based strategy. Let μ_{imp} and σ_{imp} be the monthly *out-of-sample* estimated mean and standard deviation of *RNTC* from the implied distribution based strategy, and μ_{hist} and σ_{hist} be the corresponding estimated quantities for the historical distribution based strategy. Then, the return-loss from the historical distribution based strategy is given by:

$$return - loss = \frac{\mu_{imp}}{\sigma_{imp}} \times \sigma_{hist} - \mu_{hist} \tag{19}$$

In the simplest case where σ_{imp}=σ_{hist} the return-loss measure amounts to the difference in the mean returns obtained under the two strategies.

C. Direct Maximisation: Results and Discussion

Table I shows the annualised SRs of the forward (Panels A and C) and backward-looking (Panel B and D) based strategies formed by direct maximisation of expected utility over the period 31/03/1992 to 31/08/2007. The maximisation of expected utility and the risk-adjustment of implied distributions have been implemented under the same assumed utility function for the individual and representative investor (i.e., exponential or power utility). The SRs are reported for different levels of absolute and relative risk aversion (ARA, RRA=2,4,6,8) and different sample sizes of the rolling window (36, 48, 60 and 72 observations, with corresponding SRs SR_36, SR_48, SR_60, and SR_72) used to risk-adjust the implied distribution and estimate the historical distribution by means of the Gaussian kernel.

-Table I about here-

We can see that in the case where either the exponential or power utility function is maximised, the optimal portfolios formed based on the forward-looking approach yield greater SRs than the corresponding portfolios based on historical distributions in most cases. This holds regardless of the degree of the investor's relative risk aversion and the employed window length. The greatest SR obtained by the risk-adjusted distribution is encountered in the case of $\eta=\gamma=2$ and $K=36$ months (SR=0.59 and 0.57, respectively), while the corresponding SR obtained by the historical estimators is 0.53. Notice that for any given level of risk aversion, the SRs decrease as the sample size of the rolling window increases. This implies that the recently arrived information should be weighted more heavily. Overall, the results suggest the superiority of the forward-looking approach and show that this does not depend on the choice of the utility function.

Table II shows the annualised opportunity cost over the period 31/03/1992 to 31/08/2007. Panels A and B show the results for the cases where the expected utility is maximised under an exponential and a power utility function, respectively. Results are reported for different sizes of the rolling window (36, 48, 60 and 72 observations) used to risk-adjust the implied distribution and estimate the historical distribution. The risk-adjustment has been performed by assuming that the utility function of the representative agent is exponential (Panel A) and power (Panel B).

-Table II about here-

We can see that the opportunity cost is positive in most cases regardless of the window of estimation and degree of risk aversion, i.e. the investor is better off by adopting the risk-adjusted implied rather than the historical distribution to obtain the optimal trading strategy. In particular, in the case where the individual investor uses a negative exponential function to calculate the optimal portfolio, the opportunity cost is positive for K=36, 48 months. This holds regardless of the level of his ARA; the opportunity cost becomes as high as 1.92% for the case of η=2 and K=48 months. In the case of the power utility investor, the magnitude of the opportunity costs is now even greater compared to the case of exponential utility, underlining the usefulness of option-implied distributions for the formation of optimal portfolios. In particular, the opportunity cost reported for the case of γ=8 and K=48 months is as high as 3.42%. The magnitude of the opportunity costs are of similar order to the ones reported by Jondeau and Rockinger (2006).

Nevertheless, there are some cases where the opportunity cost is negative. This occurs when the implied distributions are adjusted assuming an exponential utility function for the representative agent (for $\eta \geq 6$ and K=60, 72 months). This finding requires further explanation. It should be reminded that unlike Sharpe ratios that take into account only the mean and the standard deviation of excess portfolio returns, the opportunity cost metric takes also into account the higher-order moments, as well. In particular, a Taylor expansion of the exponential and power utility function illustrates that portfolio returns with negative skewness and excess kurtosis induce severe penalties in utility terms. In fact, the greater the degree of risk aversion, the greater this penalty becomes. Unreported results show that there are a series of cases, especially when the implied distributions are risk-adjusted by means of an exponential utility function, where the portfolio returns exhibit a greater degree of negative skewness and excess kurtosis as compared to the returns of portfolios formed on the basis of historical distributions. As a result, the mean-variance superiority of the portfolios' returns that make use of option-implied distribution is offset in some cases, due to the properties of their higher moments; this leads to the negative opportunity costs reported in Panel A of Table II.

Table III shows the portfolio turnover results. Panels A and B (D and E) show the portfolio turnover for the cases where the expected utility is maximised under an exponential (power) utility function. Results are reported for various levels of risk aversion for the individual investor and sizes of the rolling window (36, 48, 60 and 72 observations) used to risk-adjust the implied distribution and estimate the historical distribution. The risk-adjustment has been performed by assuming that the utility function of the representative agent is exponential (Panels A and B) and power (Panels D and E). Panels C and F show the portfolio turnovers' ratio of the risk-adjusted implied distributions to the historical distribution-based strategies under an exponential and a power utility function, respectively. We can see that the portfolio turnover decreases as the risk aversion increases, as expected. In addition, the ratio of the portfolio turnovers of the implied to the historical distribution-based strategies is slightly greater than one in most of the cases. This indicates that the portfolio turnover is slightly greater in the case where the investor uses the risk-adjusted implied distributions as an input in her asset allocation formation.

-Table III about here-

Table IV (Panels A and B) shows the annualised return-loss in the case where the expected utility is maximised directly under an exponential and power utility function, respectively. Results are reported for the different sizes of the rolling window (36, 48, 60 and 72 observations) used to risk-adjust the implied distribution and estimate the historical distribution. In general, the investor is roughly 0.2% to 3.7% per annum worse-off in risk-adjusted terms after deducting transaction costs, if she adopts the backward-looking approach. This implies that the greater transaction costs incurred by the forward-looking approach (arising from the fact that the portfolios based on the risk-adjusted implied distributions have greater turnover than the ones based on historical distributions) cannot offset the corresponding extra risk-adjusted returns of this approach. Therefore, the superiority of portfolios derived from the risk-adjusted implied distributions is confirmed, even after deducting the incurred transaction costs.

-Table IV about here-

IV. Further Robustness Tests

In this section we perform some further tests to assess the robustness of the superiority of the forward-looking approach demonstrated in subsection III.C. First, optimal portfolios are formed by means of a Taylor series approach for the case of the backward and forward looking approach, separately, and their performance is compared. Second, the backward-looking approach is re-assessed by estimating historical distributions via GARCH-type models.

A. Optimal Portfolio: Truncated Taylor Series Expansion

Let the mean value $\overline{W}_{t+1}$ of the future wealth defined by equation (3) be

$$\overline{W}_{t+1} = E_t(W_{t+1}) = 1 + \alpha_t \mu_{t+1} + \alpha_t^f r_{f,t+1} \tag{20}$$

where $\mu_{t+1} = E_t(r_{t+1})$. Then, at any point in time *t*, the expected utility approximated by an infinite order Taylor series expansion around $\overline{W}_{t+1}$ is given by

$$E[U(W_{t+1})] = E[\sum_{k=0}^{\infty} \frac{U^{(k)}(\overline{W}_{t+1})(W_{t+1} - \overline{W}_{t+1})^k}{k!}] \tag{21}$$

Equation (21) can be re-written, under certain assumptions (see Garlappi and Skoulakis (2009) and references therein) as:

$$E[U(W_{t+1})] = \sum_{k=0}^{\infty} \frac{U^{(k)}(\overline{W}_{t+1})}{k!} E[(W_{t+1} - \overline{W}_{t+1})^k] \tag{22}$$

For the purposes of our analysis, we calculate the optimal portfolios for *k*=2,4 and compare them with the ones derived from direct maximisation of expected utility. This will enable us to understand the features of the suggested forward-looking approach in a moments-based portfolio formation setting that is widely used. The case of *k*=2 corresponds to the familiar mean-variance Markowitz analysis while *k*=4 incorporates also the skewness and kurtosis of the returns distribution and has been extensively used in the literature. In particular, it can be shown that (see e.g., Jondeau and Rockinger (2006), Guidolin and Timmermann (2008) and references therein)

$$E[U(W_{t+1})] \approx U(\overline{W}_{t+1}) + \frac{U^{(2)}(\overline{W}_{t+1})}{2!}\sigma^2_{p,t+1} + \frac{U^{(3)}(\overline{W}_{t+1})}{3!}s^3_{p,t+1} + \frac{U^{(4)}(\overline{W}_{t+1})}{4!}k^4_{p,t+1} \tag{23}$$

where

$$\sigma^2_{p,t+1} \equiv E[(r_{p,t+1} - \mu_{p,t+1})^2] = \alpha_t^2 M_{2,t+1} \tag{24}$$

$$s^3_{p,t+1} \equiv E[(r_{p,t+1} - \mu_{p,t+1})^3] = \alpha_t^3 M_{3,t+1} \tag{25}$$

$$k_{p,t+1}^4 \equiv E[(r_{p,t+1} - \mu_{p,t+1})^4] = \alpha_t^4 M_{4,t+1} \quad (26)$$

and $M_{i,t+1}$ denotes the ith central moment at time t+1, i=1,2,3,4, i.e.

$$M_{i,t+1} \equiv E[(r_{t+1} - \mu_{t+1})^i],\ i = 2,3,4. \quad (27)$$

In the case of the negative exponential and power utility functions, the fourth order truncated Taylor series expansions [equation (23)] are given respectively by:

$$E[U(W_{t+1})] \approx -\frac{1}{\eta}\exp(-\eta \overline{W}_{t+1})(1 + \frac{\eta^2}{2}\sigma_{p,t+1}^2 - \frac{\eta^3}{6}s_{p,t+1}^3 + \frac{\eta^4}{24}k_{p,t+1}^4) \quad (28)$$

and

$$E[U(W_{t+1})] \approx \frac{\overline{W}_{t+1}^{1-\gamma} - 1}{1-\gamma} - \frac{\gamma}{2}\overline{W}_{t+1}^{(-\gamma-1)}\sigma_{p,t+1}^2 + \frac{\gamma(\gamma+1)}{6}\overline{W}_{t+1}^{(-\gamma-2)}s_{p,t+1}^3 - \frac{\gamma(\gamma+1)(\gamma+2)\overline{W}_{t+1}^{(-\gamma-3)}}{24}k_{p,t+1}^4 \quad (29)$$

Equations (28) and (29) are maximised with respect to α_t to obtain the optimal portfolio choice α_t^*; a grid search over possible values of the risky and risk-free asset weights is performed again. To implement the maximisation, the central moments M_t need to be estimated. These are alternatively extracted from the estimated historical distribution (sample historical moments, see also Jondeau and Rockinger (2006)) and the risk-adjusted implied distribution. We find similar results to the case of direct maximisation of the utility function, i.e. the optimal portfolios based on the forward-looking approach outperform those based on the historical approach; the detailed results are not reported due to space limitations.[10] Hence, the superiority of the proposed methodology is confirmed in the case of moments-based portfolio formation just as was the case with the optimal portfolios derived by direct maximisation. Interestingly, the case of k=4, delivers almost identical in value results with the full optimisation. This confirms in our setting the argument of Jondeau and Rockinger (2006) that the four-moment optimization strategy provides a very good approximation of the full scale utility optimization approach.

B. Optimal Portfolio: Alternative Estimators of the Historical Distribution

In contrast to Section III where the historical distribution of returns was estimated as a

[10] We have also computed optimal portfolios by using the sample moments estimated from the historical data as input in the Taylor series expansion instead of the moments extracted from the PDF estimated by the Gaussian kernel. We find that the forward looking based portfolios outperform again the historical ones. Therefore, the use of the Gaussian kernel does not have any negative effect on the performance of the historical portfolios.

(smoothed) histogram of past returns, we simulate GARCH-type models to provide the historical distributions (see also Liu, Shackleton, Taylor and Xu (2007) for a similar approach). We estimate the following specifications: a constant plus error term mean equation and a GARCH(1,1) model for the variance equation, an AR(1) model for the mean equation and a GARCH(1,1) model for the variance equation, an EGARCH(1,1) model (Nelson, 1991) for the variance equation to account for the asymmetric leverage effect, and an AR(1) model for the mean equation and a GARCH(1,1)-in-mean model for the variance equation to account for volatility-feedback effects in mean in the spirit of Engle, Lilien and Robins (1987). We estimate every specification by using a conditional normal and a *t*-student distribution for the residuals (Bollerslev, 1987) respectively, in order to capture the empirically documented fat-tailed unconditional returns distribution. Overall, we use eight alternative models to simulate the one-month horizon S&P 500 PDF; each model is estimated recursively from 31/03/1992 to 31/08/2007 by using a rolling window of 72 monthly observations and 100,000 simulation paths were generated at each time step to construct the PDF. We use the simulated PDFs derived by each model at every time step to calculate optimal portfolios. Then, we compare their performance with that of the optimal portfolios obtained under the forward-looking approach. Again, we find that the optimal portfolios formed under the forward-looking approach outperform those formed under the backward-looking approach (results are not reported due to space limitations).

V. The Effect of Loss Aversion

This section investigates whether forward-looking portfolios still outperform the historical ones in the case where the individual investor is loss averse; the investor is more sensitive to reductions in her financial wealth than to increases relative to a reference point and hence the value function that describes investors' preferences is steeper in the domain of losses than in the region of gains. This is a characteristic that cannot be captured by the standard utility functions that have been considered in the previous sections. To this end, investor's preferences are assumed to be described by a disappointment aversion (DA) setting, firstly introduced by Gul (1991).

The DA setting increases sensitivity to bad events (disappointments). It scales up the probabilities of all bad events by the same factor and scales down the probabilities of good events by a complementary factor, with good and bad defined as better and worse than a reference point, respectively. This framework has been employed in recent asset allocation studies so as to capture the presence of loss aversion. In addition, loss aversion may explain some stylised facts such as non-participation (i.e. zero investment in the risky asset) and the success of capital-guarantee products. For instance, Ang, Bekaert and Liu (2005) find that it can generate equity holdings that are consistent with the empirical evidence of non-participation. Driessen and Maenhout (2007) have also used it to address asset allocation questions for portfolios of stock and options. In addition, this setting is firmly grounded in decision theory and is very similar to expected utility; it retains all the axioms underlying expected utility but the independence axiom that is replaced by a weaker version so as to accommodate the Allais paradox (see Gul (1991) and Ang, Bekaert and Liu (2005) for a discussion). In line with Driessen and Maenhout (2007), a DA value function $V(W_T)$ based on a power utility function is employed, i.e.:

$$V(W_T) = \begin{cases} \frac{W_T^{1-\gamma}-1}{1-\gamma} & \text{if } W_T > \mu_W \\ \frac{W_T^{1-\gamma}-1}{1-\gamma} - (\frac{1}{A}-1)[\frac{\mu_W^{1-\gamma}-1}{1-\gamma} - \frac{W_T^{1-\gamma}-1}{1-\gamma}] & \text{if } W_T \leq \mu_W \end{cases} \quad (30)$$

where μ_W is the reference point relative to which gains or losses are measured, γ the RRA coefficient that controls the concavity of the value function in each region, and $A \leq 1$ is the coefficient of DA that controls the relative steepness of the value function in the region of gains versus the region of losses. The loss aversion decreases as A increases; A=1, corresponds to the case of the standard power utility function where there is no loss aversion. The main modelling advantage of this value function is that it is a one-parameter extension of the power utility function; hence, it nests the latter as a special case and inherits its attractive features. We follow Driessen and Maenhout (2007) and employ two values for A=0.6, 0.8 so as to consider the effect of DA; the weight of the risky asset will decrease as the loss aversion increases.

To maximise the expected value of the DA function [equation (37)], μ_W has to be defined first. We assume that μ_W equals the initial wealth invested at the risk-free rate, i.e. $\mu_W = W_t(1+r^f)$. This choice of the reference point is in line with Barberis, Huang and Santos (2001) and implies that the investor uses the risk-free rate as a benchmark to code a gain or a loss. For instance, if the riskless rate is 4 percent, the investor will be disappointed if her stock market investment returns only 3 percent. In fact, this is a realistic assumption. Veld and Veld-Merkoulova (2008) conduct a study on investors' behavior and find that a significant portion of investors use the risk-free rate as a reference point to distinguish between losses and gains.

Table V shows the annualised SRs of the forward and backward-looking strategies obtained by direct maximisation of the DA value function of the individual investor (maximisation of a Taylor series expansion is not possible since the DA value function is not globally differentiable). The SRs are reported for different levels of relative risk aversion and different values of DA (A=0.6 and A=0.8) for the individual investor.

-Table V about here-

We can see that the portfolios based on the risk-adjusted implied distributions yield greater SRs compared to the ones formed on the basis of historical distributions. This finding holds for any given level of RRA, any degree of DA, and any choice of the rolling window length. These results confirm the conclusions of the previous sections that the use of forward-looking option-implied distributions may prove beneficial. There are two additional observations to make. First, the SRs increase from A=0.6 to A=0.8. This is because the participation of the investor to the risky asset increases as the loss aversion of the investor decreases and this enables her to reap the realised risk premium. Second, the SRs decrease as the length of the rolling window increases, regardless of the employed methodology (forward or backward-looking). This is consistent with the findings of the previous sections.

Table VI shows the opportunity costs for the cases where the DA value function is maximised. The risk-adjustment has been performed by assuming that the utility function of the representative agent is exponential (Panel A) and power (Panel B).

-Table VI about here-

The results reported in Panel A are mixed. In particular, the opportunity cost tends to be positive (negative) in the case where a rolling window of 36 and 48 (60 and 72) observations is used. The explanation for this finding lies again in the higher moments of the portfolio returns' distributions. In particular, the portfolio returns derived by implied distributions that are risk-adjusted by an exponential utility function are, in some cases, characterized by a greater degree of negative skewness and excess kurtosis as compared to the corresponding returns of portfolios formed using historical distributions. On the other hand, in the case where the power utility function is employed to risk-adjust the implied distributions (Panel B), the opportunity cost is positive in all cases. Hence, the superiority of the optimal portfolios formed on the basis of implied distributions that are risk-adjusted by means of a power utility function is confirmed under the opportunity cost metric.

Panels A and B of Table VII show the ratio of the portfolio turnovers of the risk-adjusted implied distribution to the historical distribution-based strategies when the adjustment of the implied distributions has been performed by an exponential and a power utility function, respectively. The strategies are obtained by maximising a DA value function. We can see that the ratio is less than one in almost half of the cases. This is in contrast to the portfolio turnover obtained under the exponential and power utility functions that was greater than one. The results imply that the use of implied distributions is preferable to that of historical distributions in terms of the portfolio turnover. Panels C and D compare the out-of-sample performance of these two approaches under the annualised return-loss metric. The investor achieves an enhancement of up to 6.46% p.a. in terms of risk-adjusted, net of transaction costs, excess returns when she utilizes option-implied distributions for asset allocation purposes. Therefore, the superiority of the forward-looking approach in the presence of DA is even more pronounced when transaction costs are taken into account, compared to the case with the standard utility functions.

-Table VII about here-

VI. Sources of outperformance

Given that the forward-looking approach is found to be superior to the backward-looking one, we proceed to identify the source of its superiority. We use the following procedure to identify which one of the forward-looking risk-adjusted moments accounts for the reported outperformance. We calculate SRs of optimal strategies based on maximising the expected utility of the individual investor by means of a Taylor series expansion of order four by substituting repeatedly one central moment with the value of the corresponding risk-adjusted moment and the remaining three moments with the corresponding values of the central moments obtained from the historical PDF [see equations (28) and (29)]. This exercise is performed for the exponential and power utility function separately, and repeated four times so as to check whether the outperformance of the forward-looking approach stems from either the risk-adjusted mean, variance, skewness or kurtosis (i.e., in the first round, maximisation is implemented by using the risk-adjusted mean and the 'historical' variance, skewness and kurtosis. Then, in the second round, maximisation is implemented by using the risk-adjusted variance and the 'historical' mean, skewness and kurtosis, and so on). Then, we compare the obtained SRs with the corresponding ones obtained by maximising expected utility through a 4th-order Taylor series expansion using only historical moments as inputs (these are almost identical to the ones obtained by direct maximisation using as input the historical PDF, as mentioned in Section IV.A.).

Table VIII reports the annualised SRs obtained from the described above exercise in the case where an exponential utility function describes the preferences of the individual investor; Panels A to D tabulate the SRs using as input the first four forward-looking moments, respectively. A comparison with the SRs obtained by direct maximisation using as input the historical PDF (Panel B of Table I) shows that the outperformance of the forward-looking approach is due to the use of the forward-looking mean since this delivers the highest SRs; the use of the other three forward-looking moments leads to SRs of similar magnitude as compared to the "historical" ones. Similar results are also obtained in the case where a power utility function is assumed to describe the preferences of the individual investor. The results imply that the use of information from option markets allows the investor to time the market more effectively than historical information (see also Golez (2010) for a similar finding based on the ability of option-implied dividend-to-price ratios to forecast future returns within a mean-variance setting).[11] This is also confirmed by unreported results from application of the Treynor-Mazuy (1966) model to formally test the market timing ability of the proposed

[11] A by-product of the exercise is the implication that forecasting the mean accurately is of first order importance when compared separately to forecasting volatility, skewness, and kurtosis for asset allocation purposes. This result does not invalidate previous findings of the literature since results depend on the asset universe, time period under scrutiny, and the undertaken research method. For instance, Fleming, Kirby and Ostdiek (2001) and DeMiguel, Plyakha, Uppal and Vilkov (2010) investigate the economic value of volatility timing without investigating market timing, while Jondeau and Rockinger (2008) investigate the economic value of forecasting the first three moments simultaneously (distributional timing) versus that of volatility timing.

forward-looking approach.

-Table VIII about here-

VII. The effect of the 2007-2009 crisis

In this section, we explore whether the recent subprime/ liquidity crisis has an impact on the previously documented outperformance of the forward-looking approach. September 2007 is widely regarded as the "official" kick-off of the crisis since it marks the insolvency of the U.K. bank Northern Rock. The crisis may have been alleviated with the inception of a coordinated international bailout in October 2008, yet its effects "...might well drag on over the next few years" (Brunnermeier (2009, p.98)). Therefore, we repeat the analysis described in the previous sections over the September 2007-December 2009 period. Table IX reports the average excess return obtained under the forward (Panels A and C) and backward-looking (Panels B and D) based strategies. Results are tabulated for different sample sizes of the rolling window (36, 48, 60 and 72 observations, with corresponding mean excess returns mean_36, mean_48, mean_60, and mean_72) used to risk-adjust the implied distribution and estimate the historical distribution by means of the Gaussian kernel. Two remarks can be drawn. First, the average excess returns are negative under both approaches. This is expected given the bearish nature of the period under scrutiny. In this case, it is not meaningful to compare the performance of the two methods by means of risk-adjusted measures; the usual positive risk-return relationship may not hold for negative returns and paradoxes in the ranking of strategies may appear (see e.g., Israelsen, 2005). Second, based on the reported average excess returns, we can see that the forward-looking based strategy yields greater average losses than the backward-looking one. This comes as no entire surprise either; it is in line with the evidence that a number of well-known strategies (e.g., value and momentum) did not prove to be profitable over the recent crisis period even though they were consistently among investor's favourites before the crisis (see Jones, 2010). It is also in line with the findings of the literature that finds that option implied distributions do not anticipate stock market crashes (see e.g., Bates, 1991, Gemmill and Saflekos, 2000).

To understand the source of the previous results, we calculate the average optimal weight of the risky asset delivered by each strategy. Table X reports the average optimal weight of the risky asset obtained under the forward (Panels A and C) and backward-looking (Panels B and D) based strategies during this crisis period. Results are tabulated for different sample sizes of the rolling window (36, 48, 60 and 72 observations, with corresponding average weight weight_36, weight_48, weight_60, and weight_72). We can see that the historical based optimal portfolios assign, on average, smaller weights to the risky asset compared with the forward based portfolios. This is because the historical approach extrapolates past returns since it is based on a kernel estimator for a rolling window of observations. Hence, it proves to be relatively less unsuccessful during the crisis due to its lengthy duration. In particular, as the utilized rolling window starts incorporating the initial negative returns of the crisis period, this leads to lower portfolio weights assigned to the risky asset in the subsequent periods. Hence, the out-of-sample performance is less bad during the months of autumn/ winter 2008 when the disastrous S&P 500 returns come along. Given that the pattern and timing of the unprecedented shocks affecting the financial system could not have been anticipated or attributed to a particular fundamental economic process, the historical approach is "lucky" to exploit the negative momentum that unfolds during this prolonged crisis, especially when short rolling windows are used (e.g., 36 observations).[12]

VIII. Dynamic portfolio choice with option-implied distributions

This section shows how the risk-adjusted option-implied distributions can be utilized in a dynamic asset allocation setup, if one is ready to assume specific dynamics for the risky asset's returns and the underlying risk factor, as it is standard in the literature (e.g., Wachter, 2002, Sangvinatsos and Wachter, 2005, Liu, 2007). In particular, in what follows we adopt the setup of Wachter (2002) for an investor with power utility function defined over her terminal wealth. The rationale for adopting this setup is discussed below.

A. Solving the Portfolio Choice Problem

[12] We have also compared the out-of-sample performance of the forward-looking and historical portfolios over the entire March 1992 – December 2009 period. We find mixed results depending on the way that the historical PDF is estimated (i.e. Gaussian kernel or GARCH type models). This is not surprising though since the 2007-2009 period is an outlier given the clustering of unprecedented events that take place then; the probability of a collapse of an institution like the Lehman Brothers, the liquidation of numerous hedge funds, major liquidity problems in the interbank market, etc, is miniscule. Therefore, its inclusion is likely to distort findings reported over the" normal" 1992-2007 period. Interestingly, the outperformance of the Gaussian kernel based historical portfolios vanishes in the case where no bounds on portfolio weights are imposed.

Let the dynamics of the risky asset (the S&P 500) returns be given by the following stochastic differential equation (SDE):

$$\frac{dS_t}{S_t} = \mu_t dt + \sigma dw_t \tag{31}$$

where w_t denotes a standard Brownian motion defined on the probability space (Ω, F, P) with filtration F and time set $[0,T], 0 \le T \le \infty$. The underlying risk factor is the market price of risk X_t defined by:

$$X_t = \frac{\mu_t - r_f}{\sigma} \tag{32}$$

It is further assumed that X_t follows an Ornstein-Uhlenbeck process, given by the following SDE:

$$dX_t = -\lambda_X (X_t - \bar{X})dt - \sigma_X dw_t \tag{33}$$

The volatilities σ and σ_X are assumed to be constant and strictly positive and λ_X is assumed to be greater than or equal to zero. This setup implies a perfect negative correlation between shocks affecting stock returns and the risk factor, enabling us to solve the multi-period problem via the martingale method of Cox and Huang (1989). The investor seeks to maximize:

$$\max_{\alpha_t} E_t \left[\frac{W_T^{1-\gamma}}{1-\gamma} \right] \tag{34}$$

subject to the wealth budget constraint, represented by the process:

$$dW_t = \left[\alpha_t (\mu_t - r_f) + r_f \right] W_t dt + \alpha_t \sigma W_t dw_t \tag{35}$$

where α_t denotes the portion of wealth invested in the risky asset. It can be shown that (see Wachter, 2002):

$$\alpha_t^* = \frac{1}{\gamma} \frac{\mu_t - r_f}{\sigma^2} - \frac{1}{\gamma} \frac{\sigma_X}{\sigma} \left[A_1(\tau) X_t + A_2(\tau) \right] \tag{36}$$

where

$$A_1(\tau) = \frac{1-\gamma}{\gamma} \frac{2(1-e^{-\eta\tau})}{\left[2\eta - (b+\eta)(1-e^{-\eta\tau}) \right]} \tag{37}$$

$$A_2(\tau) = \frac{1-\gamma}{\gamma} \frac{4\lambda_X \bar{X}(1-e^{-\eta\tau/2})^2}{\eta\left[2\eta-(b+\eta)(1-e^{-\eta\tau})\right]} \tag{38}$$

and $b = 2\left(\frac{\gamma-1}{\gamma}\sigma_X - \lambda_X\right)$, $\eta = \sqrt{b^2-4ac}$, $a = \frac{1-\gamma}{\gamma}$, $c = \frac{1}{\gamma}\sigma_X^2$ where $\tau \equiv T-t$.

Three remarks about equation (36) are in order. First, the optimal portfolio choice for a multi-period investor can be calculated using information extracted from risk-adjusted option-implied distributions. Second, the multi-period optimal portfolio weight is determined by two components; the first one corresponds to the familiar static mean-variance demand, while the second one corresponds to the hedging demand that arises due to the desire of the intertemporal optimizer to smooth her wealth path by hedging away the shocks affecting the investment opportunity set. Notice that the hedging demand, and hence the optimal portfolio weight, depends on investor's horizon, as this is captured by τ. Third, the portion of wealth allocated to the risky asset changes with the prevailing market conditions since both the myopic and hedging demand components depend on the prevailing investment opportunities, reflected by the market price of risk X_t.

Finally, a note on the choice of Wachter's (2002) setup is in order. The aim of this section is to illustrate how, in principle, the proposed forward-looking methodology can be utilized in an intertemporal portfolio choice setup. Therefore, we have decided to keep the exposition as simple as possible without loss of generality. We regard Wachter's (2002) setup to be useful to this end because it utilizes the market price of risk as the underlying risk factor. This can be directly extracted from risk-adjusted option-implied distributions, the use of which we suggest in this study. This availability renders the estimation of the risk factor dynamics and the determination of the risky asset's weights straightforward. However, we acknowledge that Wachter's (2002) assumptions may be restrictive; the volatility parameter is assumed to be constant and the market of price of risk and returns' innovations are assumed to be perfectly negatively correlated. Yet, the proposed methodology could be also useful in a richer intertemporal portfolio choice setup. For example, Chacko and Viceira (2005) assume volatility to be stochastic; in that case, an extra hedging demand component would appear in the risky asset demand of the multi-period investor.

B. Results

In this subsection, we compute optimal portfolio weights for the intertemporal investor using information from the risk-adjusted option-implied distributions and show how these differ from the corresponding myopic weights. For implementation purposes, the parameters of the SDE describing the dynamics of the risk factor X [equation (33)] need to be estimated. To this end, we use the time series of the annualized market prices of risk for the period February 1992 to December 2009, extracted from option-implied distributions that have been risk-adjusted using a power utility function and a window of 60 observations. The maximum likelihood estimates of these coefficients are λ_X=1.11, $\bar{X}$=0.4815 and σ_X=1.0191. Then, the values for functions A_1 and A_2 [equations (37) and (38), respectively] are calculated. Finally, we determine the optimal portfolio weights assigned to the risky asset at any point in time over the period under scrutiny, for any degree of relative risk aversion ($\gamma > 1$) and investment horizon; to this end, we use the prevailing option-implied market price of risk on each date of the examined period.

Figure I shows the evolution of the myopic (T=0) and hedging demands [first and second terms of equation (36), respectively] for the risky asset over the period from February 1992 to December 2009; the hedging demand has been plotted for investor's horizons T=5, 10 years. We can see that there is a strong hedging motive for the intertemporal investor, since the hedging demand for the risky asset is much greater than the corresponding myopic one. This hedging motive becomes stronger for investors with longer horizons leading to even greater portfolio weights. This is due to the assumed perfect negative correlation between returns and the market price of risk; an adverse shock to returns increases X (i.e. future investment opportunity improve), and hence motivates the long-term investor to hold more of her wealth in the risky asset. The figure also demonstrates that both components in the portfolio choice fluctuate with market conditions, as expected.

The pronounced hedging motive arising for the intertemporal optimizer leads to particularly high hedging and total demands for the risky asset. To illustrate this finding, Table XI provides detailed information with respect to the median and maximum hedging and total demands that have been calculated for various investment horizons as well as for various levels of risk aversion during the whole sample period. We can see that the portfolio weights for the intertemporal investor become quite dramatic in particular periods. This comes as no surprise though since it is a well known feature of dynamic asset allocation that has been extensively documented in prior literature (see *inter alia* Campbell and Viceira (1999, 2001), Wachter (2002), Berkelaar, Kouwenberg and Post (2004), Campbell, Chacko, Rodriguez and Viceira (2004), Sangvinatsos and Wachter (2005)).

Figure II shows the indicative relationship between the hedging demand for the risky asset and investor's horizon; the portion of wealth assigned to the risky asset due to hedging is plotted as a function of investor's horizon for three different levels of RRA=2, 4, and 8, corresponding to the option-implied market price of risk prevailing on February 28, 1992 (i.e. X_t=0.34). In fact, the hedging demand increases with investor's horizon, confirming that the hedging motive is greater for long-term rather than short-term investors in this setup (see Wachter (2002) for a similar finding). The rate of increase in the hedging demand is also greater (for moderate investment horizon levels) as the investor becomes more risk averse. This is because the more risk averse the investor is, the more sensitive she is towards shocks affecting her wealth path, increasing her hedging motive. Equation (36) shows that the hedging motive depends, in the first place, on investor's exposure to this shock through her myopic demand. Therefore, the previous effect can be even better understood by plotting the corresponding ratio of hedging to myopic demands for this particular date. This is done in Figure III which shows that hedging demands dominate myopic ones, especially as the horizon and the degree of risk aversion increases.

Next, we evaluate the out-of-sample performance of the portfolio strategies derived from the myopic and multi-period setting by using the risk-adjusted option-implied information just as they have already been described in Section III.A. These are reported for the indicative case of RRA=4 and for the time period February 1992- August 2007. The multi-period strategies with fixed 5 and 10-year horizons that use option-implied information yield SR=0.51 and SR=0.52, respectively, and outperform the corresponding myopic one that yields SR=0.42. For comparability purposes we have also calculated the out-of-sample performance of multi-period portfolio strategies based on historical distributions. To determine the corresponding portfolio weights we need to calculate the values of functions A_1 and A_2 in equations (37) and (38). To this end, we estimate the coefficients of the SDE (33) describing the risk factor dynamics using historical distributions. To perform this estimation we utilize a time series of annualized market prices of risk implied from the historical distributions that have been extracted via a Gaussian kernel. The maximum likelihood estimates of the SDE coefficients for the historical market price of risk dynamics are λ_X=0.2492, $\bar{X}$=0.2782 and σ_X=0.3371. The multi-period strategies with fixed 5 and 10-year horizons using historical information yielded out-of-sample SR=0.35 and SR=0.38, respectively, for the period February 1992- August 2007. These are considerably lower than the SRs yielded by the corresponding multi-period strategies using option-implied information.[13]

It should be acknowledged, though, that mean-variance measures are not the most appropriate ones for evaluating multi-period strategies. Therefore, we also calculate the opportunity cost metric suggested by Sangvinatsos and Wachter (2005) to examine the amount a long-term investor would be willing to pay in order to follow the optimal multi-period strategy instead of the suboptimal myopic one. To calculate this cost one needs to resort to the indirect utility function of this long-term investor that is computable through the corresponding dynamic programming problem (see Cox and Huang (1989) for the equivalence of the two solution approaches). In particular, Wachter (2002, p. 72) shows that the indirect utility function $J(\cdot)$ for this problem is given by:

$$J(W,X,t)=\frac{W^{1-\gamma}}{1-\gamma}H(X,t)=\frac{W^{1-\gamma}}{1-\gamma}\exp\left\{A_1(\tau)X^2/2+A_2(\tau)X+A_3(\tau)\right\} \tag{39}$$

Let $\hat{H}(X,t)$ denote the function in equation (39) calculated under the optimal long-term portfolio strategy and $\tilde{H}(X,t)$ the function calculated under the suboptimal myopic strategy. Sangvinatsos and Wachter (2005, p. 214) define the opportunity cost as:

$$\wp\left(X_0,0\right)=1-\left[\frac{\tilde{H}(X_0,0)}{\hat{H}(X_0,0)}\right]^{\frac{1}{\gamma-1}} \tag{40}$$

Figure IV shows the opportunity cost (as a percentage of initial wealth) of following a suboptimal myopic strategy instead of the optimal multi-period strategy as a function of investor's horizon for three levels of RRA=2, 4 and 6, corresponding to the prevailing option-

[13] It should be noted that, in line with the results presented in Section VII, during the recent crisis period the dynamic strategies using option-implied information performed poorly and yielded more negative out-of-sample returns relative to the ones that utilized historical information only. The explanation for this poor performance is that the strong hedging motive that is present in multi-period strategies with option-implied information leads to very high allocations to the risky asset (S&P 500). Given the very poor record of S&P 500 returns during the recent crisis period, the poor performance of these multi-period strategies comes as no surprise.

implied market price of risk in February 28, 1992. This exercise reveals that the opportunity cost is extremely high. Furthermore, it increases as the investor's horizon increases and the degree of relative risk aversion becomes greater. These findings are due to the fact that the longer the investor's horizon and the more risk averse she becomes, the stronger the hedging motive becomes. In other words, the possibility to hedge away intertemporal shocks to the underlying opportunity set becomes of utmost importance for a long-term and highly risk-averse investor.

IX. Conclusions

This paper takes a forward-looking approach to asset allocation by suggesting a way of using information from market option prices (option-implied distributions) to calculate the optimal portfolio. The motivation for doing so is that by their nature, implied distributions are forward-looking. Therefore, they are expected to proxy the true unknown distribution/moments of asset returns that are required in any asset allocation problem more precisely than historical distributions do. Next, we test the validity of our hypothesis by comparing the *out-of-sample* performance of the forward-looking approach to that of a typical backward-looking one. Finally, we extend the suggested approach to a dynamic asset allocation setting.

The commonly used asset space of an index and a risk-free asset is considered. We extract implied distributions from the S&P 500 futures options and subsequently convert them to the corresponding risk-adjusted ones. We perform the risk-adjustment by backing out the coefficient of (absolute or relative) risk-aversion of the representative investor. We obtain optimal portfolios and compare them to the ones derived by historical distributions. To check the robustness of the results, we perform maximisation of the individual investor's utility function per se (direct maximisation) and its Taylor series approximation, separately. The effect of loss aversion, as well as that of the recent subprime crisis, is also investigated. Furthermore, we use a number of criteria to assess the out-of-sample performance of the optimal portfolios and estimate the benchmark optimal historical portfolios by alternative models.

The analysis over the period March 1992 to August 2007 reveals that using option-implied information increases the investor's out-of-sample risk-adjusted returns and makes her better off compared with the case where she uses only historical distributions. The results hold regardless of the performance measure, specification of the utility function, estimator of the historical portfolio, and objective value function to be maximised. Most importantly, the superiority of the forward-looking approach for asset allocation purposes is also confirmed in the case where transaction costs are taken into account. Not surprisingly, over the recent subprime crisis both approaches yielded negative average returns. In the case of multi-period portfolio choice, our results confirm the conclusion of prior studies that the risky asset's ability to hedge unfavourable shocks to the underlying investment opportunity set is of utmost importance for a long-term investor, leading to high hedging demands.

Our results have at least four implications. First, the use of information from option markets can be used for market timing purposes. In particular, we have found that the use of the forward-looking mean drives the outperformance of the suggested approach. Second, this finding confirms that expected returns is the most crucial input for asset allocation purposes endorsing the conclusions of previous studies (e.g., Merton (1980) and Chopra and Ziemba (1993)); higher order moments were not the source of outperformance within our setting. This in line with the findings of Jondeau and Rockinger (2006) who find that, for moderate values of risk aversion, the formation of static optimal portfolios is not considerably affected by departing from a mean-variance setting.[14] Third, the recent crisis has harmed the reported market timing ability of the proposed method, as it has also done with that of a number of commonly used investment strategies (see Fabozzi, Focardi and Jonas (2010)). This does not invalidate the application of the presented method, as well as that of other well known strategies though, given that unprecedented market conditions have been experienced over the recent crisis.[15] Fourth, our method offers an alternative way to estimate the (forward-looking) market risk premium (see Duan and Zhang (2010) and references therein for alternative methods).

The presented framework for asset allocation opens up at least three avenues for future research. First, the benefits from using risk-adjusted implied distributions to form optimal

[14] However, it should be noted that Jondeau and Rockinger (2008) find that higher moments do play a role over shorter time intervals considered within a dynamic setting. This may be attributed to the fact that bursts of higher moments tend to average out over the longer periods of time considered in their earlier static setting paper.

[15] For instance, the VIX volatility index, a measure of investor anxiety about the U.S. stock market (sometimes termed "fear index") soared to an all-time high of 80.06 on October 27, 2008 as fear reached record levels.

portfolios should be explored for alternative risky assets. This can be done by extracting implied distributions from other option markets too. Second, given the vast literature on alternative methods to extract implied distributions, these may be extracted by an alternative method to the one employed in this paper. The risk-adjustment of the implied distribution may also be performed by other methods than the Bliss and Panigirtzoglou (2004) one (see e.g., Liu, Shackleton, Taylor and Xu (2007)). Finally, the asset allocation problem should be investigated within the suggested approach in the case where there are more than one risky assets in the investor's portfolio.

References

Adler, T., and M. Kritzman, 2007, Mean-variance versus full-scale optimization: In and out of sample, *Journal of Asset Management* 7, 302-311.

Ait-Sahalia, Y., and A. W. Lo, 2000, Nonparametric risk management and implied risk aversion, *Journal of Econometrics* 94, 9-51.

Ait-Sahalia, Y., and M. W. Brandt, 2008, Consumption and portfolio choice with option-implied state prices, NBER Working Paper.

Alentorn, A., and S. Markose, 2008, Generalized extreme value distribution and extreme economic value at risk (EE-VaR), in E. J. Kontoghiorghes, B. Rustem, and P. Winker, eds: *Computational methods in Financial Engineering*: *Essays in Honor of Manfred Gilli*, (Springer Berlin Heidelberg).

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* 14, 41-66.

Ang, A., G. Bekaert, and J. Liu, 2005, Why stocks may disappoint, *Journal of Financial Economics* 76, 471-508.

Bakshi, G., Panayotov, G., and G. Skoulakis, 2010, Improving the predictability of real economic activity and asset returns with forward variances inferred from option portfolios, *Journal of Financial Economics*, forthcoming.

Barberis, N., M. Huang, and T. Santos, 2001, Prospect theory and asset prices, *Quarterly Journal of Economics* 116, 1-53.

Barone-Adesi, G., and R. E. Whaley, 1987, Efficient analytical approximation of American option values, *Journal of Finance* 42, 301-320.

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

Berkelaar, A. B., R. Kouwenberg and T. Post, 2004, Optimal portfolio choice under loss aversion, *Review of Economics and Statistics* 86, 973-987.

Berkowitz, J., 2001, Testing density forecasts with applications to risk management**,** *Journal of Business and Economics Statistics* 19, 465-474.

Black, F., 1976, The pricing of commodity contracts, *Journal of Financial Economics* 3, 167-179.

Bliss, R. R., and N. Panigirtzoglou, 2002, Testing the stability of implied probability density functions, *Journal of Banking and Finance* 26, 381-422.

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

Bollerslev, T., 1987, A conditionally heteroskedastic time series model for security prices and rates of return data. *Review of Economics and Statistics* 59, 542-547.

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

Brunnermeier, M., 2009, Deciphering the liquidity and credit crunch 2007-2008, *Journal of Economic Perspectives* 23, 77-100.

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

Campbell, J. Y., and L. M. Viceira, 2001, Who should buy long-term bonds?, *American Economic Review* 91, 99-127.

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

Carr, P., and L. Wu, 2009, Variance risk premiums, *Review of Financial Studies* 22, 1311-1341.

Chacko, G. and L. M. Viceira, 2005, Dynamic consumption and portfolio choice under stochastic volatility and incomplete markets, *Review of Financial Studies* 18, 1369-1402.

Chan, L. K.C., J. Karceski, and J. Lakonishok, 1999, On portfolio optimisation: Forecasting covariances and choosing the risk model, *Review of Financial Studies* 12, 937-974.

Chopra, V. K. and W. T. Ziemba, 1993, The effect of errors in means, variances and covariances on optimal portfolio choice, *Journal of Portfolio Management* 19, 6-11.

Chang, B.-Y., P. Christoffersen, K. Jacobs, and G. Vainberg, 2009, Option-implied measures of equity risk, Working Paper, McGill University.

Cox, J., C. and C.-F. Huang, 1989, Optimal consumption and portfolio policies when asset prices follow a diffusion process, *Journal of Economic Theory* 49, 33-83.

Cremers, M. and D. Weinbaum, 2010, Deviations from put-call parity and stock return predictability, *Journal of Financial and Quantitative Analysis* 45, 335-367.

DeMiguel, V., L. Garlappi, and R. Uppal, 2009, How inefficient are simple asset allocation strategies?, *Review of Financial Studies* 22, 1915-1953.

DeMiguel, V., Y. Plyakha, R. Uppal, and G. Vilkov, 2010, Improving portfolio selection using option-implied volatility and skewness, Working Paper, Goethe University.

Driessen, J., P. J. Maenhout, and G. Vilkov, 2009, The price of correlation risk: evidence from equity options, *Journal of Finance* 64, 1377-1406.

Driessen, J., and P. J. Maenhout, 2007, An empirical portfolio perspective on option pricing anomalies, *Review of Finance* 11, 561-603.

Duan, J-C., and W. Zhang, 2010, Forward-looking market risk premium, Working Paper, National University of Singapore.

Engle, R., F., Lilien, D. M. Robins, R.P., 1987. Estimating time varying risk premia in the term structure: The ARCH-M model. *Econometrica* 55, 391-407.

Fabozzi, F.J., Focardi, S.M., and C. Jonas, 2010, *Investment management after the global financial crisis*, Research Foundation of CFA Institute.

Fleming, J., C. Kirby, and B. Ostdiek, 2001, The economic value of volatility timing, *Journal of Finance* 56, 329-352.

Garlappi, L., and G. Skoulakis, 2009, Taylor series approximations to expected utility and optimal portfolio choice, Working Paper, University of British Columbia.

Gemmill, G., and A. Saflekos, 2000. How useful are implied distributions? Evidence from stock-index options, *Journal of Derivatives* 7, 83–98.

Giamouridis, D., and G. Skiadopoulos, 2010, The informational content of financial options for quantitative asset management: A review, in B. Scherer and K. Winston, eds: *Handbook of Quantitative Asset Management*, (Oxford University Press), forthcoming.

Golez, B., 2010, Expected returns and dividend growth rates implied in derivative markets, Working paper, Pompeu Fabra University.

Gul, F., 1991, A theory of disappointment aversion, *Econometrica* 59, 667-686.

Guidolin, M., and A. Timmermann, 2008, International asset allocation under regime switching, skew and kurtosis preferences**,** *Review of Financial Studies* 21, 889-935.

Harvey, C. R., and R. E. Whaley, 1991, S&P 100 Index option volatility, *Journal of Finance* 46, 1551-1561.

Israelsen, C.L., 2005, A refinement to the Sharpe ratio and information ratio, *Journal of Asset Management* 5, 423-427.

Jabbour, C., J. F. Peňa, J. C. Vera, and L. F. Zuluaga, 2008, An estimation-free, robust CVaR portfolio allocation model, *Journal of Risk* 11, 57-78.

Jackwerth, J., 2004, *Option-implied risk-neutral distributions and risk aversion,* Research Foundation of CFA Institute.

Jiang, G. and Y. Tian, 2005, The model-free implied volatility and its information content, *Review of Financial Studies* 18, 1305-1342.

Jiang, G. and Y. Tian, 2007, Extracting model-free volatility from option prices: An examination of the VIX index, *Journal of Derivatives*, 1-26.

Jondeau, E., and M. Rockinger, 2008, The economic value of distributional timing, Working paper, University of Lausanne.

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

Jondeau, E., and M. Rockinger, 2003, Conditional volatility, skewness, and kurtosis: Existence, persistence, and comovements, *Journal of Economic Dynamics and Control* 27, 1699-1737.

Jones, R.C., 2010, Maybe it really is different this time, *Journal of Portfolio Management* 36, 60-72.

Kan, R., and G., Zhou, 2007, Optimal portfolio choice with parameter uncertainty, *Journal of Financial and Quantitative Analysis* 42, 621-656.

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

Kroll, Y., H. Levy, and H. M. Markowitz, 1984, Mean-variance versus direct utility maximization, *Journal of Finance* 36, 47-61.

Liu, J., 2007, Portfolio selection in stochastic environments, *Review of Financial Studies* 20, 1-39.

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 and Finance* 31, 1501-1520.

Merton, R. C., 1973, A rational theory of option pricing, *Bell Journal of Economics and Management Science* 41, 141-183.

Merton, R. C., 1980, On estimating the expected return on the market: An exploratory investigation, *Journal of Financial Economics* 8, 323-361.

Nelson, D., B., 1991, Conditional heteroskedasticity in asset returns: A new approach. *Econometrica* 59, 347-370.

Panigirtzoglou, N., and G. Skiadopoulos, 2004, A new approach to modeling the dynamics of implied distributions: Theory and evidence from the S&P 500 options, *Journal of Banking and Finance* 28, 1499-1520.

Poon, S.-H., and C. W.J., Granger, 2003, Forecasting volatility in financial markets: A review, *Journal of Economic Literature* 26, 478-539.

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

Sangvinatsos, A., and J. A. Wachter, 2005, Does the failure of the expectations hypothesis matter for long-term investors? *Journal of Finance* 60, 179-230.

Sharpe, W. F., 2007, Expected utility asset allocation, *Financial Analysts Journal* 63, 18-30.

Siegel, A. F., 1995, Measuring systematic risk using implicit beta, *Management Science* 41, 124-128.

Siegel, A. F. and A. Woodgate, 2007, Performance of portfolios optimised with estimation error, *Management Science* 53, 1005-1015.

Simaan, Y., 1993, What is the opportunity cost of mean-variance investment strategies? *Management Science* 39, 578-587.

Söderlind, P., and L. E.O., Svensson, 1997, New techniques to extract market expectations from financial instruments, *Journal of Monetary Economics* 40, 383-429.

Treynor, J. L., and K. K. Mazuy, 1966, Can mutual funds outguess the market? *Harvard Business Review* 44, 131-136.

Veld, C., and Y. V. Veld-Merkoulova, 2008, The risk perceptions of individual investors, *Journal of Economic Psychology* 29, 226-252.

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, 63-91.

Xing, Y., X. Zhang, and R. Zhao, 2010, What does individual option volatility smirk tell us about future equity returns?, *Journal of Financial and Quantitative Analysis* 45, 641-662.

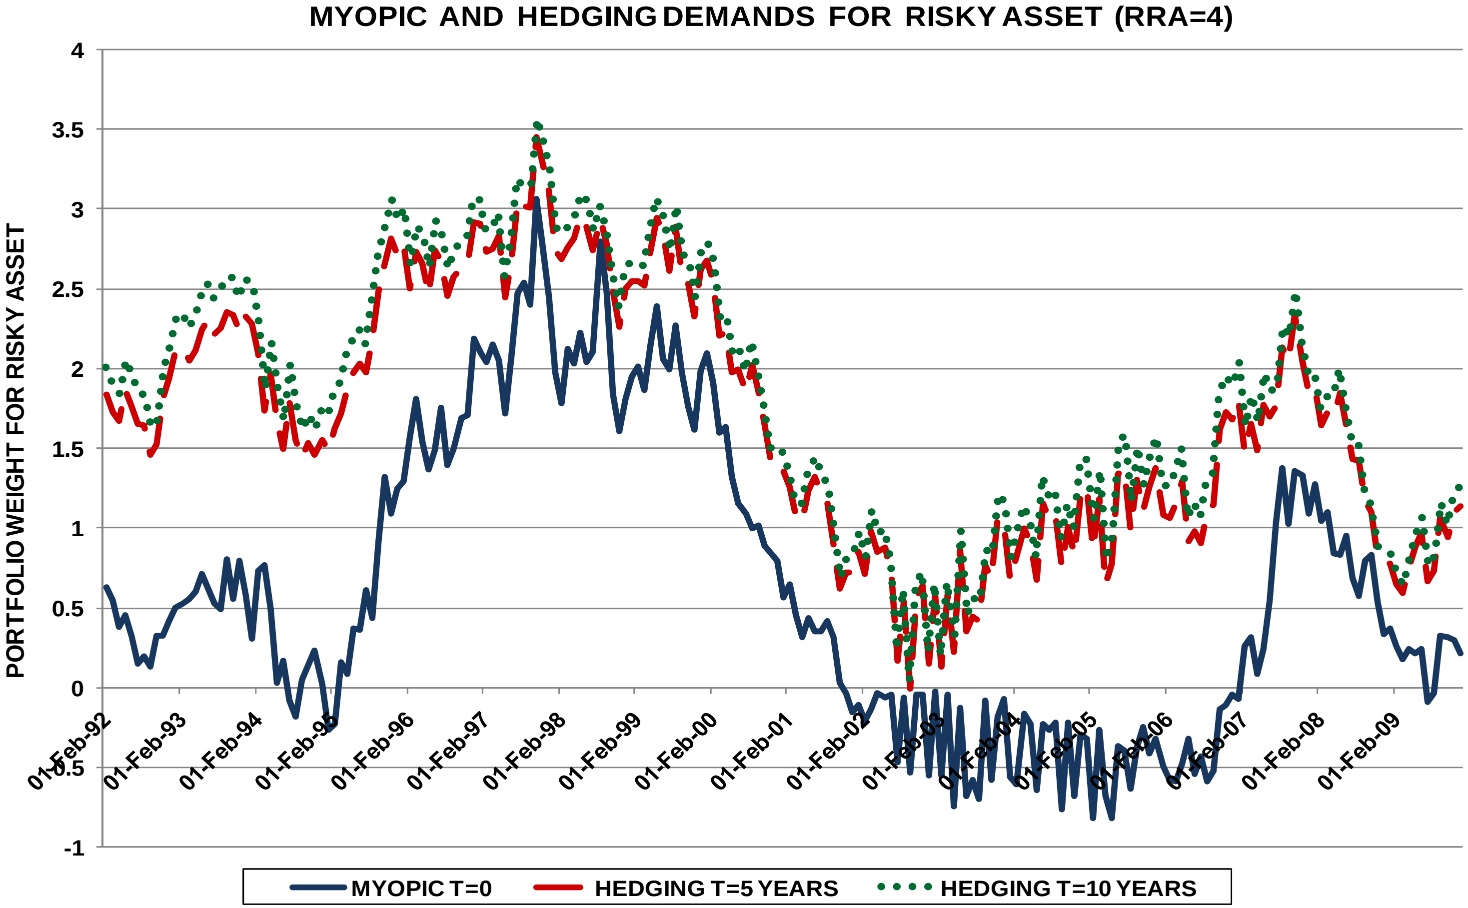

Figure I. The figure shows the unconstrained optimal risky asset demand for a myopic investor (solid curve) with coefficient of Relative Risk Aversion RRA=4, as well as the corresponding optimal hedging demands for an investor with a 5-year horizon (dashed curve) and a 10-year horizon (dotted curve). The period extends from February 1992 to December 2009, utilizing the market price of risk calculated by option-implied distributions that have been risk-adjusted using a rolling window of 60 observations and a power utility function. The myopic and hedging demands are given by the first component and second component, respectively, of equation (36).

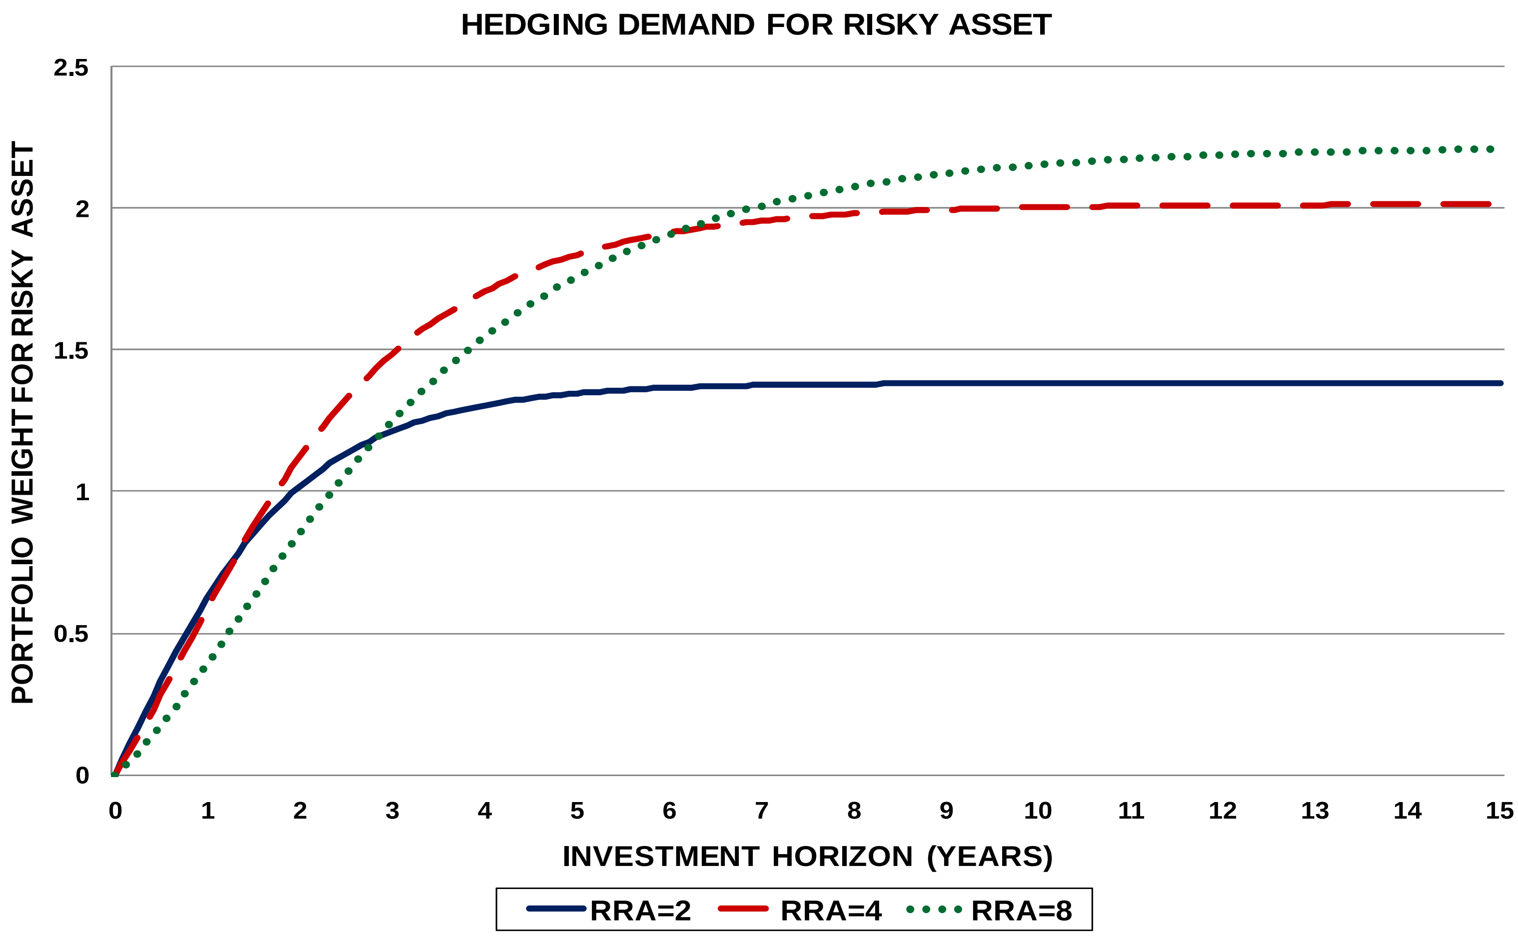

Figure II. This figure shows the hedging demand for the risky asset as a function of investor's horizon for three levels of Relative Risk Aversion (RRA). This hedging demand has been calculated for the risk-adjusted option-implied market price of risk prevailing on February 28, 1992 (i.e. the annualized market price of risk is set equal to X_t=0.34). The solid curve corresponds to the case of RRA=2, the dashed curve to the case of RRA=4 and the dotted curve to the case of RRA=8. The hedging demand is given by the second component of the portfolio choice expression in equation (36).

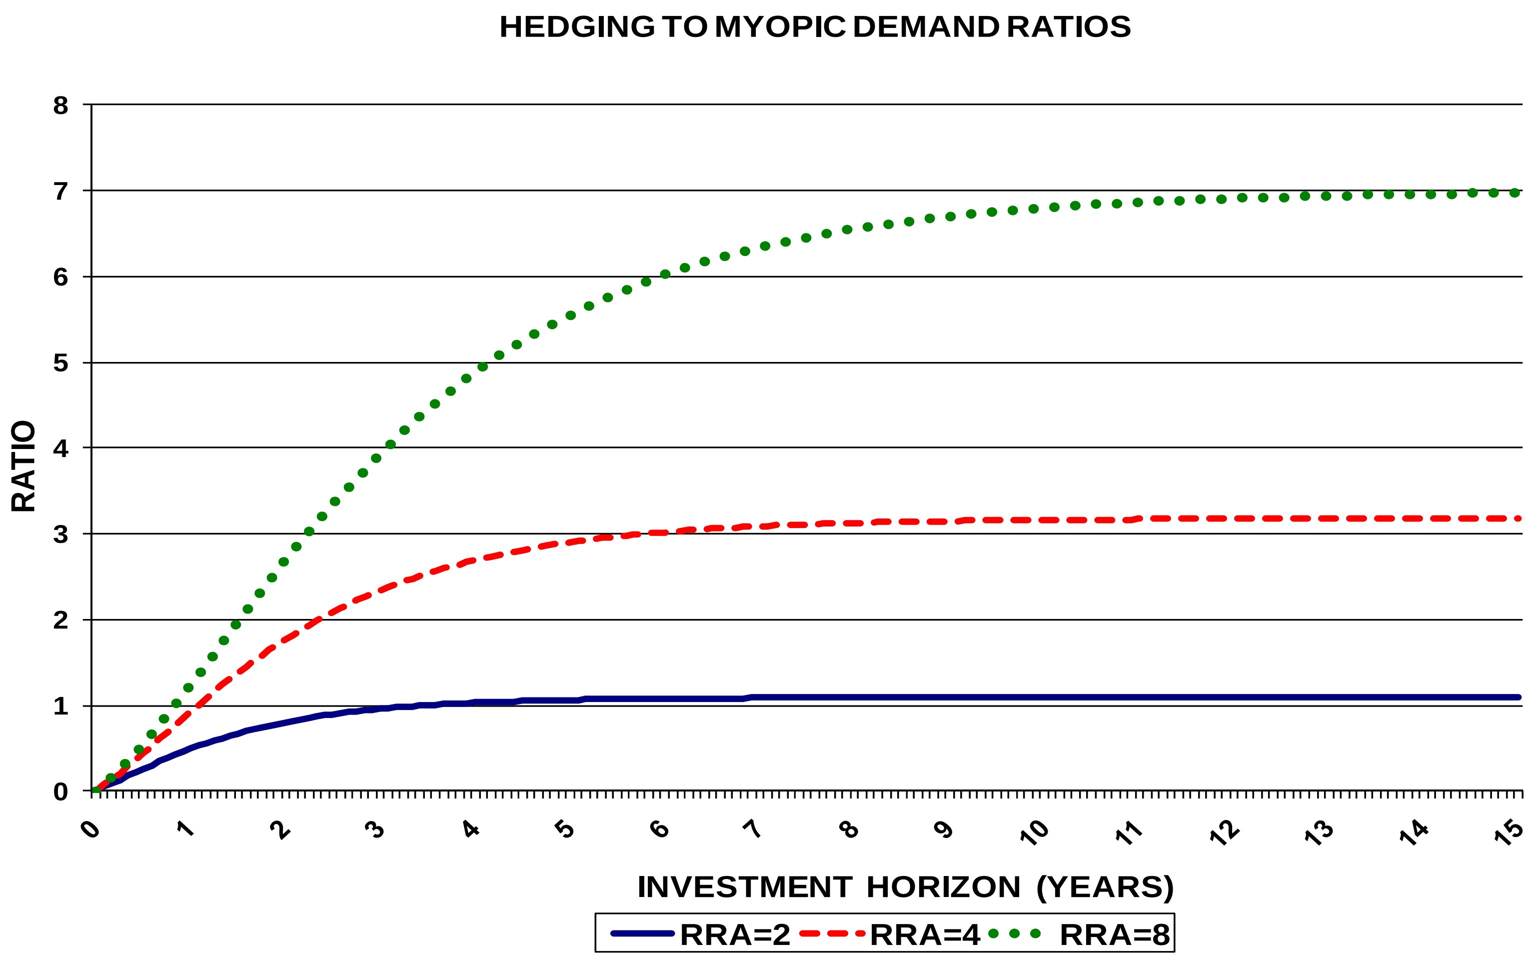

Figure III. This figure shows the ratio of hedging to myopic demand for the risky asset as a function of investor's horizon for three levels of Relative Risk Aversion (RRA). This ratio has been calculated for the risk-adjusted option-implied market price of risk prevailing on February 28, 1992 (i.e. the annualized market price of risk is set equal to X_t=0.34). The solid curve corresponds to the case of RRA=2, the dashed curve to the case of RRA=4 and the dotted curve to the case of RRA=8. The myopic and hedging demands are given by the first component and second component, respectively, of equation (36).

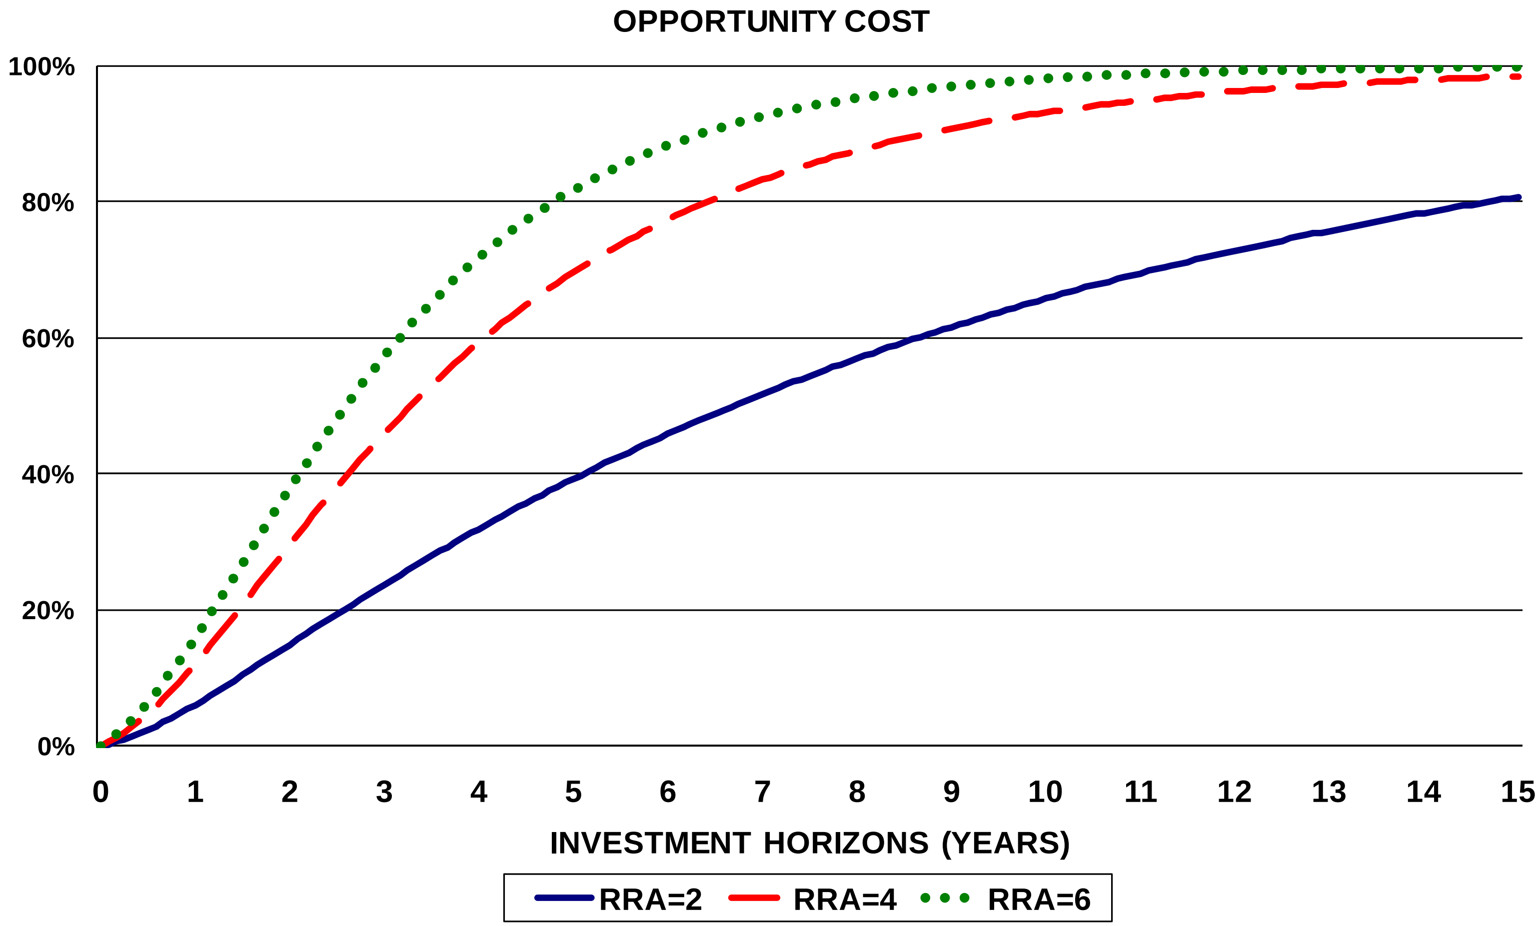

Figure IV. This figure shows the opportunity cost (as a percentage of initial wealth) of following a suboptimal myopic strategy instead of the optimal multi-period strategy as a function of investor's horizon for three levels of Relative Risk Aversion (RRA=2,4,6). This cost has been calculated for the risk-adjusted option-implied market price of risk prevailing in February 28, 1992 (i.e. the annualized market price of risk is set equal to X_t=0.34). The solid curve corresponds to the case of RRA=2, the dashed curve to the case of RRA=4 and the dotted curve to the case of RRA=6.

List of Tables

Table I

Sharpe Ratios obtained by Direct Maximisation of Expected Utility

Annualised Sharpe Ratios (SRs) for the period 31/03/1992 to 31/08/2007. Panels A and C report the SRs obtained by the optimal strategy based on the risk-adjusted implied distributions. The risk-adjustment (expected utility maximisation) has been performed assuming that the representative (individual) agent has an exponential and a power utility function, respectively. Panels B and D report the SRs obtained by the optimal strategy based on the historical distributions where maximisation of the exponential and power utility function has been performed, respectively. The SRs are reported for different levels of absolute and relative risk aversion (ARA, RRA=2, 4, 6 ,8) and different sizes of the rolling window (36, 48, 60 and 72 observations, with corresponding SRs SR_36, SR_48, SR_60, and SR_72) used to risk-adjust the implied distribution and estimate the historical distribution by means of the Gaussian kernel.

Panel A: Risk-Adjusted Implied Distributions & Exponential Utility function				
	ARA=2	ARA=4	ARA=6	ARA=8
Sharpe Ratio_36	0.59	0.51	0.46	0.46
Sharpe Ratio_48	0.50	0.49	0.45	0.45
Sharpe Ratio_60	0.33	0.38	0.35	0.35
Sharpe Ratio_72	0.28	0.33	0.32	0.32
Panel B: Historical Distributions & Exponential Utility function				
	ARA=2	ARA=4	ARA=6	ARA=8
Sharpe Ratio_36	0.53	0.48	0.45	0.45
Sharpe Ratio_48	0.41	0.44	0.35	0.35
Sharpe Ratio_60	0.32	0.35	0.31	0.31
Sharpe Ratio_72	0.26	0.27	0.26	0.26
Panel C: Risk-Adjusted Implied Distributions & Power Utility function				
	RRA=2	RRA=4	RRA=6	ARA=8
Sharpe Ratio_36	0.57	0.51	0.51	0.52
Sharpe Ratio_48	0.50	0.46	0.47	0.48
Sharpe Ratio_60	0.34	0.38	0.40	0.41
Sharpe Ratio_72	0.27	0.31	0.33	0.34
Panel D: Historical Distributions & Power Utility function				
	RRA=2	RRA=4	RRA=6	ARA=8
Sharpe Ratio_36	0.53	0.48	0.45	0.35
Sharpe Ratio_48	0.41	0.44	0.35	0.25
Sharpe Ratio_60	0.32	0.35	0.31	0.31
Sharpe Ratio_72	0.26	0.27	0.26	0.26

Table II

Direct Maximisation of Expected Utility: Annualised Opportunity Cost over the Period 31/03/1992 to 31/08/2007

Panels A and B show the opportunity cost (how much worse off the investor is in return terms by adopting the historical distribution rather than the risk-adjusted implied distribution to obtain the optimal trading strategy) for the cases where the expected utility is maximised under an exponential and power utility function, respectively. Results are reported for different sizes of the rolling window (36, 48, 60 and 72 observations) used to risk-adjust the implied distribution and estimate the historical distribution by means of a Gaussian kernel estimator. The risk-adjustment has been performed by assuming that the representative agent has an exponential (Panel A) and a power (Panel B) utility function.

Panel A: Exponential Utility function				
	ARA=2	ARA=4	ARA=6	ARA=8
36_Obs	1.08%	0.60%	0.24%	0.36%
48_Obs	1.92%	0.72%	1.44%	1.08%
60_Obs	0.24%	0.24%	-0.84%	-0.60%
72_Obs	0.60%	0.24%	-1.20%	-0.84%
Panel B: Power Utility function				
	RRA=2	RRA=4	RRA=6	RRA=8
36_Obs	0.60%	0.84%	1.56%	2.76%
48_Obs	1.92%	0.48%	3.24%	3.72%
60_Obs	0.48%	1.32%	1.92%	1.44%
72_Obs	0.36%	1.20%	1.20%	0.96%

Table III

Direct Maximisation of Expected Utility: Portfolio Turnover over the Period 31/03/1992 to 31/08/2007

Panels A and B (D and E) show the portfolio turnover for the cases where the expected utility is maximised under an exponential (power) utility function. Results are reported for different sizes of the rolling window (36, 48, 60 and 72 observations) used to risk-adjust the implied distribution and estimate the historical distribution by means of a Gaussian kernel estimator. The risk-adjustment has been performed by assuming that the representative agent has an exponential (Panels A and B) and a power (Panels D and E) utility function. Panels C and F show the ratio of the turnover generated by the strategy based on risk-adjusted implied distributions relative to that generated by the strategy based on historical distributions under an exponential and power utility function, respectively.

Panel A: Risk-adjusted Implied Distributions & Exponential Utility function				
	ARA=2	ARA=4	ARA=6	ARA=8
turnover_36	45.59%	33.20%	22.69%	17.02%
turnover_48	33.11%	25.38%	18.40%	13.71%
turnover_60	31.96%	23.47%	17.82%	13.38%
turnover_72	28.76%	21.07%	15.02%	11.20%
Panel B: Historical Distributions & Exponential Utility function				
	ARA=2	ARA=4	ARA=6	ARA=8
turnover_36	34.19%	34.03%	27.52%	20.60%
turnover_48	34.13%	27.64%	23.00%	16.49%
turnover_60	29.06%	23.09%	15.94%	11.32%
turnover_72	24.88%	20.83%	13.24%	9.40%
Panel C: Turnover Ratio: Risk-adjusted Implied Distributions/Historical distributions				
	ARA=2	ARA=4	ARA=6	ARA=8
ratio_36	1.33	0.98	0.82	0.83
ratio _48	0.97	0.92	0.80	0.83
ratio _60	1.10	1.02	1.12	1.18
ratio _72	1.16	1.01	1.13	1.19
Panel D: Risk-adjusted Implied Distributions & Power Utility function				
	RRA=2	RRA=4	RRA=6	RRA=8
turnover_36	52.16%	38.43%	29.04%	22.89%
turnover_48	42.42%	32.15%	22.88%	18.51%
turnover_60	51.43%	37.38%	26.74%	21.17%
turnover_72	41.53%	26.99%	19.48%	15.76%
Panel E: Historical Distributions & Power Utility function				
	RRA=2	RRA=4	RRA=6	RRA=8
turnover_36	34.23%	33.83%	27.48%	25.56%
turnover_48	34.25%	27.38%	22.88%	21.10%
turnover_60	28.88%	22.69%	15.84%	11.30%
turnover_72	24.59%	20.47%	13.15%	9.37%
Panel F: Turnover Ratio: Risk-adjusted Implied Distributions/Historical distributions				
	RRA=2	RRA=4	RRA=6	RRA=8
ratio _36	1.52	1.14	1.06	0.90
ratio _48	1.24	1.17	1.00	0.88
ratio _60	1.78	1.65	1.69	1.87
ratio_72	1.69	1.32	1.48	1.68

Table IV

Return-Loss for Direct Maximisation over the Period 31/03/1992 to 31/08/2007

Panels A and B show the annualised return-loss in the case where the expected utility is maximised directly under an exponential and power utility function, respectively. Results are reported for different sizes of the rolling window (36, 48, 60 and 72 observations) used to risk-adjust the implied distribution and estimate the historical distribution by means of a Gaussian kernel estimator. The risk-adjustment has been performed by assuming that the representative agent has an exponential (Panel A) and a power (Panel B) utility function.

Panel A: Return-Loss for direct maximization of Exponential Utility function				
	ARA=2	ARA=4	ARA=6	ARA=8
return-loss_36	0.84%	0.54%	0.22%	0.24%
return-loss_48	1.99%	0.83%	1.46%	1.08%
return-loss_60	0.16%	0.63%	0.47%	0.34%
return-loss_72	0.45%	1.06%	0.74%	0.55%
Panel B: Return-Loss for direct maximization of Power Utility function				
	RRA=2	RRA=4	RRA=6	RRA=8
return-loss_36	0.18%	0.35%	0.73%	1.99%
return-loss_48	1.67%	0.28%	1.57%	2.40%
return-loss_60	-0.43%	0.10%	0.61%	0.53%
return-loss_72	-0.22%	0.34%	0.54%	0.49%

Table V

Sharpe Ratios obtained by Direct Maximisation of the Disappointment Aversion Value Function

Entries report the annualised Sharpe Ratios (SRs) for the period 31/03/1992 to 31/08/2007. Panels A and B report the SRs obtained by the optimal strategy based on the risk-adjusted implied distributions derived by assuming that the representative agent has an exponential and a power utility function, respectively. Panel C reports the SRs obtained by the optimal strategy based on the historical distributions. The SRs are reported for different levels of relative risk aversion (RRA=2,4,6,8) and different sizes of the rolling window (36, 48, 60 and 72 observations with corresponding SRs SR_36, SR_48, SR_60, and SR_72) used to risk-adjust the implied distribution. Entries in each panel are reported for values of the parameter *A*=0.6,0.8 of the disappointment aversion utility function.

Panel A: Risk-Adjusted Implied Distributions by Exponential Utility function				
A=0.6	RRA=2	RRA=4	RRA=6	RRA=8
Sharpe Ratio_36	0.41	0.33	0.35	0.36
Sharpe Ratio_48	0.49	0.40	0.42	0.43
Sharpe Ratio_60	0.46	0.37	0.38	0.40
Sharpe Ratio_72	0.38	0.34	0.36	0.37
A=0.8	RRA=2	RRA=4	RRA=6	RRA=8
Sharpe Ratio_36	0.49	0.83	0.40	0.41
Sharpe Ratio_48	0.45	0.42	0.41	0.42
Sharpe Ratio_60	0.40	0.37	0.37	0.38
Sharpe Ratio_72	0.32	0.34	0.34	0.35
Panel B: Risk-Adjusted Implied Distributions by Power Utility function				
A=0.6	RRA=2	RRA=4	RRA=6	RRA=8
Sharpe Ratio_36	0.40	0.44	0.45	0.46
Sharpe Ratio_48	0.46	0.44	0.45	0.46
Sharpe Ratio_60	0.45	0.48	0.49	0.49
Sharpe Ratio_72	0.36	0.39	0.40	0.41
A=0.8	RRA=2	RRA=4	RRA=6	RRA=8
Sharpe Ratio_36	0.50	0.45	0.47	0.48
Sharpe Ratio_48	0.46	0.43	0.45	0.46
Sharpe Ratio_60	0.44	0.44	0.46	0.47
Sharpe Ratio_72	0.33	0.34	0.36	0.37
Panel C: Historical Distributions				
A=0.6	RRA=2	RRA=4	RRA=6	RRA=8
Sharpe Ratio_36	0.24	0.26	0.25	0.07
Sharpe Ratio_48	0.24	0.13	0.13	-0.06
Sharpe Ratio_60	0.10	0.10	0.10	0.10
Sharpe Ratio_72	0.05	0.05	0.05	0.05
A=0.8	RRA=2	RRA=4	RRA=6	RRA=8
Sharpe Ratio_36	0.45	0.38	0.38	0.25
Sharpe Ratio_48	0.46	0.33	0.30	0.17
Sharpe Ratio_60	0.40	0.30	0.30	0.30
Sharpe Ratio_72	0.26	0.21	0.21	0.21

Table VI

Direct Maximisation of the Disappointment Aversion Value Function: Annualised Opportunity Cost over the Period 31/03/1992 to 31/08/2007

Results are reported for different sizes of the rolling window (36, 48, 60 and 72 observations) used to risk-adjust the implied distribution and estimate the historical distribution by means of a Gaussian kernel estimator. The risk-adjustment has been performed by assuming that the representative agent has an exponential (Panel A) and a power (Panel B) utility function. Entries in each panel are reported for both values of the disappointment aversion parameter (*A*=0.6, 0.8) employed in this study.

Panel A: Risk-Adjusted Implied Distributions by Exponential Utility function

A=0.6	RRA=2	RRA=4	RRA=6	RRA=8
36_obs	0.15%	-0.05%	-0.02%	0.15%
48_obs	0.21%	-0.01%	0.00%	0.17%
60_obs	0.19%	-0.14%	-0.09%	-0.06%
72_obs	-0.07%	-0.21%	-0.14%	-0.10%
A=0.8	RRA=2	RRA=4	RRA=6	RRA=8
36_obs	0.09%	0.48%	0.01%	0.14%
48_obs	-0.01%	0.08%	0.05%	0.17%
60_obs	-0.01%	-0.06%	-0.09%	-0.07%
72_obs	0.06%	-0.09%	-0.11%	-0.08%

Panel B: Risk-Adjusted Implied Distributions by Power Utility function

A=0.6	RRA=2	RRA=4	RRA=6	RRA=8
36_obs	0.13%	0.06%	0.05%	0.20%
48_obs	0.17%	0.14%	0.10%	0.25%
60_obs	0.25%	0.13%	0.10%	0.08%
72_obs	0.02%	0.03%	0.03%	0.03%
A=0.8	RRA=2	RRA=4	RRA=6	RRA=8
36_obs	0.10%	0.12%	0.10%	0.21%
48_obs	0.01%	0.16%	0.16%	0.25%
60_obs	0.08%	0.17%	0.12%	0.10%
72_obs	0.11%	0.10%	0.07%	0.06%

Table VII

Direct Maximisation of the Disappointment Aversion Value Function: Portfolio Turnover and Return-Loss: 31/03/1992 - 31/08/2007

Panels A and B show the ratio of the portfolio turnovers of the risk-adjusted implied distribution to the historical distribution based strategies. Panels C and D show the annualised return-loss. The strategies are obtained by maximising a disappointment aversion value function. Results are reported for different sizes of the rolling window (36, 48, 60 and 72 observations) used to risk-adjust the implied distribution and estimate the historical distribution by means of a Gaussian kernel estimator. The risk-adjustment has been performed by assuming that the representative agent has an exponential (Panels A and C) and a power (Panels B and D) utility function. Entries in each panel are reported for both values of the disappointment aversion parameter (A=0.6,0.8) employed in this study.

Panel A: Turnover Ratio: Risk-adjusted Implied Distributions by Exponential Utility/Historical distributions				
A=0.6	RRA=2	RRA=4	RRA=6	RRA=8
ratio_36	0.99	0.92	1.10	0.82
ratio _48	0.84	1.12	1.34	0.91
ratio _60	0.97	1.95	2.39	2.78
ratio _72	0.92	1.67	2.04	2.41
A=0.8	RRA=2	RRA=4	RRA=6	RRA=8
ratio _36	0.83	0.80	0.77	0.67
ratio _48	0.92	0.77	0.84	0.71
ratio _60	1.07	1.12	1.35	1.59
ratio _72	0.94	1.07	1.22	1.44
Panel B: Turnover Ratio: Risk-adjusted Implied Distributions by Power Utility/Historical distributions				
A=0.6	RRA=2	RRA=4	RRA=6	RRA=8
ratio_36	1.01	1.12	1.36	0.99
ratio_48	0.84	1.18	1.43	0.96
ratio_60	1.29	1.99	2.42	2.77
ratio_72	1.34	1.77	2.15	2.47
A=0.8	RRA=2	RRA=4	RRA=6	RRA=8
ratio_36	0.98	0.96	1.00	0.85
ratio_48	0.94	0.85	0.96	0.81
ratio_60	1.37	1.32	1.61	1.83
ratio_72	0.76	0.86	1.11	1.31
Panel C: Return-Loss for the Disappointment Aversion value function (Exponential utility risk-adjustment)				
A=0.6	RRA=2	RRA=4	RRA=6	RRA=8
return-loss_36	2.31%	0.78%	0.62%	1.70%
return-loss_48	3.26%	2.10%	1.45%	2.45%
return-loss_60	3.79%	1.39%	0.96%	0.74%
return-loss_72	2.63%	1.17%	0.81%	0.62%
A=0.8	RRA=2	RRA=4	RRA=6	RRA=8
return-loss_36	0.87%	6.46%	0.39%	1.52%
return-loss_48	0.02%	1.48%	1.23%	2.11%
return-loss_60	-0.11%	0.94%	0.56%	0.45%
return-loss_72	1.21%	1.51%	0.99%	0.79%
Panel D: Return-Loss for the Disappointment Aversion value function (Power utility risk-adjustment)				
A=0.6	RRA=2	RRA=4	RRA=6	RRA=8
return-loss_36	2.18%	1.67%	1.22%	2.25%
return-loss_48	2.92%	2.29%	1.57%	2.56%
return-loss_60	3.58%	1.86%	1.26%	0.94%
return-loss_72	2.34%	1.26%	0.86%	0.65%
A=0.8	RRA=2	RRA=4	RRA=6	RRA=8
return-loss_36	0.74%	1.01%	1.01%	2.05%
return-loss_48	0.06%	1.54%	1.44%	2.29%
return-loss_60	0.45%	1.56%	1.11%	0.86%
return-loss_72	1.39%	1.46%	1.06%	0.84%

Table VIII

Sources of Outperformance of the Forward-looking Approach: The Exponential Utility Case

Panels A to D report the annualised Sharpe Ratios (SRs) obtained by maximisation of a fourth order Taylor series expansion of expected utility over the period 31/03/1992 to 31/08/2007 by substituting repeatedly one central moment with the value of the corresponding forward-looking risk-adjusted moment and the remaining three with the corresponding values of the central moments obtained from the historical PDF. For example Panel A reports the SRs obtained by using the risk-adjusted mean and the 'historical' variance, skewness and kurtosis. Panel B reports the SRs obtained by using the risk-adjusted variance and the 'historical' mean, skewness and kurtosis, and so on. The risk-adjustment (maximisation) has been performed assuming that the representative (individual) agent has an exponential utility function. The SRs are reported for different levels of absolute and relative risk aversion (ARA, RRA=2, 4, 6, 8) and different sizes of the rolling window (36, 48, 60, and 72 observations, with corresponding SRs SR_36, SR_48, SR_60, and SR_72) used to risk-adjust the implied distribution.

Panel A: Forward-looking mean and historical variance, skewness and kurtosis

	ARA=2	ARA=4	ARA=6	ARA=8
Sharpe Ratio_36	0.61	0.56	0.54	0.50
Sharpe Ratio_48	0.48	0.49	0.51	0.50
Sharpe Ratio_60	0.31	0.40	0.44	0.45
Sharpe Ratio_72	0.24	0.32	0.38	0.41

Panel B: Forward-looking variance and historical mean, skewness and kurtosis

	ARA=2	ARA=4	ARA=6	ARA=8
Sharpe Ratio_36	0.49	0.37	0.35	0.35
Sharpe Ratio_48	0.37	0.26	0.25	0.25
Sharpe Ratio_60	0.21	0.20	0.19	0.19
Sharpe Ratio_72	0.22	0.15	0.13	0.13

Panel C: Forward-looking skewness and historical mean, variance, and kurtosis

	RRA=2	RRA=4	RRA=6	RRA=8
Sharpe Ratio_36	0.53	0.45	0.43	0.43
Sharpe Ratio_48	0.40	0.39	0.34	0.34
Sharpe Ratio_60	0.32	0.30	0.30	0.30
Sharpe Ratio_72	0.26	0.24	0.24	0.24

Panel D: Forward-looking kurtosis and historical mean, variance, and skewness

	RRA=2	RRA=4	RRA=6	RRA=8
Sharpe Ratio_36	0.53	0.46	0.45	0.45
Sharpe Ratio_48	0.41	0.40	0.36	0.36
Sharpe Ratio_60	0.31	0.31	0.31	0.31
Sharpe Ratio_72	0.25	0.25	0.25	0.25

Table IX

Average Excess Return obtained by Direct Maximisation over the Period 28/09/2007-31/12/2009

Panels A and C report the monthly average excess returns obtained by the optimal strategy based on the risk-adjusted implied distributions. The risk-adjustment (expected utility maximisation) has been performed assuming that the representative (individual) agent has an exponential and a power utility function, respectively. Panels B and D report the monthly average excess returns obtained by the optimal strategy based on the historical distributions where maximisation of the exponential and power utility function has been performed, respectively. Results are reported for different sizes of the rolling window (36, 48, 60 and 72 observations with corresponding average excess returns mean_36, mean_48, mean_60, and mean_72) used to risk-adjust the implied distribution and estimate the historical distribution by means of a Gaussian kernel estimator.

Panel A: Risk-Adjusted Implied Distributions & Exponential Utility function				
	ARA=2	ARA=4	ARA=6	ARA=8
mean_36	-1.90%	-0.92%	-0.61%	-0.46%
mean_48	-2.00%	-1.01%	-0.68%	-0.51%
mean_60	-2.44%	-1.35%	-0.90%	-0.67%
mean_72	-1.62%	-0.81%	-0.54%	-0.41%
Panel B: Historical Distributions & Exponential Utility function				
	RRA=2	RRA=4	RRA=6	RRA=8
mean_36	-0.89%	-0.22%	-0.15%	-0.13%
mean_48	-1.19%	-0.61%	-0.46%	-0.35%
mean_60	-1.60%	-1.33%	-0.90%	-0.67%
mean_72	-1.44%	-0.72%	-0.48%	-0.36%
Panel C: Risk-Adjusted Implied Distributions & Power Utility function				
	RRA=2	RRA=4	RRA=6	RRA=8
mean_36	-2.33%	-1.49%	-1.18%	-1.03%
mean_48	-2.34%	-1.56%	-1.23%	-1.07%
mean_60	-2.42%	-1.56%	-1.24%	-1.07%
mean_72	-1.24%	-0.94%	-0.82%	-0.76%
Panel D: Historical Distributions & Power Utility function				
	RRA=2	RRA=4	RRA=6	RRA=8
mean_36	-0.89%	-0.22%	-0.14%	-0.13%
mean_48	-1.18%	-0.60%	-0.46%	-0.35%
mean_60	-1.60%	-1.33%	-0.90%	-0.68%
mean_72	-1.43%	-0.72%	-0.48%	-0.36%

Table X

Average optimal weight of the risky asset obtained by Direct Maximisation over the Period 28/09/2007-31/12/2009

Panels A and C report the average optimal weight of the risky asset obtained by the optimal strategy based on the risk-adjusted implied distributions. The risk-adjustment (expected utility maximisation) has been performed assuming that the representative (individual) agent has an exponential and a power utility function, respectively. Panels B and D report the average optimal weight of the risky asset obtained by the optimal strategy based on the historical distributions where maximisation of the exponential and power utility function has been performed, respectively. Results are reported for different sizes of the rolling window (36, 48, 60 and 72 observations with corresponding average excess returns weight_36, weight_48, weight_60, and weight_72) used to risk-adjust the implied distribution and estimate the historical distribution by means of a Gaussian kernel estimator.

Panel A: Risk-Adjusted Implied Distributions & Exponential Utility function				
	ARA=2	ARA=4	ARA=6	ARA=8
weight_36	68.07%	39.51%	26.33%	19.75%
weight_48	82.22%	46.69%	31.13%	23.35%
weight_60	107.25%	61.49%	40.99%	30.75%
weight_72	78.18%	39.09%	26.06%	19.55%
Panel B: Historical Distributions & Exponential Utility function				
	RRA=2	RRA=4	RRA=6	RRA=8
weight_36	-7.68%	-10.95%	-11.37%	-8.98%
weight_48	13.72%	5.86%	3.38%	2.54%
weight_60	47.94%	40.08%	26.58%	19.93%
weight_72	28.21%	14.11%	9.40%	7.05%
Panel C: Risk-Adjusted Implied Distributions & Power Utility function				
	RRA=2	RRA=4	RRA=6	RRA=8
weight_36	72.49%	60.67%	49.74%	44.18%
weight_48	86.53%	65.35%	52.83%	46.50%
weight_60	110.41%	76.07%	60.05%	51.94%
weight_72	80.88%	57.88%	47.73%	42.62%
Panel D: Historical Distributions & Power Utility function				
	RRA=2	RRA=4	RRA=6	RRA=8
weight_36	-8.07%	-11.16%	-11.69%	-9.29%
weight_48	13.30%	5.48%	3.13%	2.41%
weight_60	-7.10%	-3.50%	-2.35%	19.90%
weight_72	27.98%	14.08%	9.39%	7.05%

Table XI

Median and maximum portfolio allocations to the risky asset for a multi-period investor over the period 28/02/1992- 31/12/2009

Panel A reports the median and maximum optimal portfolio weights to the risky asset for a multi-period investor for different investment horizons (*T*=2, 5 and 10 years) and levels of Relative Risk Aversion (RRA=2, 4, 6, 8 and 10), given by equation (36). The case of *T*=0 years corresponds to the optimal myopic portfolio allocation (i.e. the first component of equation 36). These optimal portfolio weights have been calculated for the period 28/02/1992- 31/12/2009 using information on the market price of risk extracted by option-implied distributions that have been risk-adjusted via a power utility function and a window of 60 monthly observations. Panel B reports the corresponding median and maximum hedging demands for the multi-period investor (i.e. the second component of equation (36)) during the same time period.

Panel A: Total demand

	Median weights				Max weights			
	T=0	*T*=2	*T*=5	*T*=10	*T*=0	*T*=2	*T*=5	*T*=10
RRA=2	83.66%	163.12%	196.3%	199.91%	612.57%	865.09%	896.58%	899.01%
RRA=4	41.83%	134.55%	197.26%	214.23%	306.28%	575.55%	651.07%	663.28%
RRA=6	27.89%	111.05%	187.90%	220.42%	204.19%	437.04%	532.42%	554.83%
RRA=8	20.91%	93.60%	177.61%	222.36%	153.14%	353.43%	456.55%	487.47%
RRA=10	16.73%	79.73%	165.54%	215.92%	122.51%	297.02%	402.02%	439.58%

Panel B: Hedging demand

	Median weights				Max weights			
	T=0	*T*=2	*T*=5	*T*=10	*T*=0	*T*=2	*T*=5	*T*=10
RRA=2	0%	82.97%	116.93%	120.64%	0%	252.53%	284.01%	286.44%
RRA=4	0%	92.91%	165.66%	183.49%	0%	269.26%	344.79%	357.00%
RRA=6	0%	81.66%	169.11%	199.58%	0%	232.86%	328.23%	350.64%
RRA=8	0%	70.57%	162.70%	200.74%	0%	200.29%	303.41%	334.32%
RRA=10	0%	61.66%	153.31%	198.14%	0%	174.51%	279.51%	320.85%