LSF Research Working Paper Series

N°. 12 - 02

Date: April 10, 2012

Title: Conditioned Higher Moment Portfolio Optimisation Using Optimal Control

Author(s)*: Marc Boissaux, Jang Schiltz

Abstract: Within a traditional context of myopic discrete-time mean-variance portfolio investments, the problem of conditioned optimisation, in which predictive information about returns contained in a signal is used to inform the choice of portfolio weights, was first expressed and solved in concrete terms by [1]. An optimal control formulation of conditioned portfolio problems was proposed and justified by [2]. This opens up the possibility of solving variants of the basic problem that do not allow for closed-form solutions through the use of standard numerical algorithms used for the discretisation of optimal control problems. The present paper applies this formulation to set and solve variants of the conditioned portfolio problem which use the third and fourth moments as well as the variance. Using backtests over a realistic data set, the performance of strategies resulting from conditioned optimisation is then compared to that obtained using analogous optimisation strategies which do not exploit conditioning information. In particular, we report on both ex ante improvements to the accessible expected return-risk boundaries and the ex post results obtained.

Keywords: Skewness, Kurtosis, Optimal Control, Portfolio Optimization

JEL Classification : C02, C61, G11

***Corresponding Author's Address:** Tel: +352 46 66 44 5716; Fax : 352 46 66 44 6834
E-mail address: marc.boissaux@uni.lu

The LSF Research Working Paper Series is available online: http://www.lsf.lu/eng/Research/Working-Papers/2011
For editorial correspondence, please contact: martine.zenner@uni.lu

University of Luxembourg
Faculty of Law, Economics and Finance
Luxembourg School of Finance
4 Rue Albert Borschette
L-1246 Luxembourg

Conditioned Higher Moment Portfolio Optimisation Using Optimal Control

Marc Boissaux and Jang Schiltz
marc.boissaux@uni.lu, jang.schiltz@uni.lu

Luxembourg School of Finance
University of Luxembourg, Luxembourg

April 10, 2012

Abstract

Within a traditional context of myopic discrete-time mean-variance portfolio investments, the problem of conditioned optimisation, in which predictive information about returns contained in a signal is used to inform the choice of portfolio weights, was first expressed and solved in concrete terms by [1]. An optimal control formulation of conditioned portfolio problems was proposed and justified by [2]. This opens up the possibility of solving variants of the basic problem that do not allow for closed-form solutions through the use of standard numerical algorithms used for the discretisation of optimal control problems.

The present paper applies this formulation to set and solve variants of the conditioned portfolio problem which use the third and fourth moments as well as the variance. Using backtests over a realistic data set, the performance of strategies resulting from conditioned optimisation is then compared to that obtained using analogous optimisation strategies which do not exploit conditioning information. In particular, we report on both ex ante improvements to the accessible expected return-risk boundaries and the ex post results obtained.

Keywords: Skewness, Kurtosis, Optimal Control, Portfolio Optimization
JEL Classification: C02, C61, G11

1 Introduction

The present paper contributes to two strains of portfolio optimisation literature. The first is conditioned portfolio optimisation, which discusses the mathematically correct treatment of information external to the investment assets themselves within what is otherwise a classical portfolio optimisation context. The second is portfolio optimisation involving higher moments of returns, which attempts to optimise for expected levels of portfolio returns moments beyond mean and variance. The optimal control formulation of conditioned portfolio problems introduced in [2] allows for generic numerical solution methods to be applied in the context of conditioned optimisation if single signal series are used, and was applied to obtain constrained-weight solutions to the basic conditioned mean-variance problem in [3]. In this paper, the approach is applied to the higher-moment problem context. We formulate and backtest two constrained-weight higher-moment problem variants which avoid nonconvex objective functions. In both cases, the use of conditioning information significantly improves observed strategy performance with respect to all metrics optimised by each problem formulation. We also briefly discuss and give results for the full four-moment problem using quartic polynomial utility functions, and find that results provide evidence that the full problem can be worked in practice even though its potentially nonconvex objective function may cause numerical issues.
The structure of the paper is as follows. Section 2 introduces the two previously separate problems of conditioned mean-variance portfolio optimisation and of optimisation involving higher moments of returns, and then moves on to describe the data set used for the present study. Section 3 reports on the backtesting setup used and defines optimal control formulations for two optimisation problems involving the fourth moment of returns. The various results obtained are then presented and discussed. Four-moment results are also given for confirmation purposes. Section 4 concludes.

2 Context

2.1 The conditioned portfolio problem

Following a related theoretical argument in [4], the type of portfolio problem nowadays often referred to as conditioned was formulated and solved in [1]. The authors consider the problem of mean-variance portfolio optimisation within a discrete-time myopic investment world, such that only two time instants are considered - an initial time t, at which the investment choice is made, and a final time $t+1$, at which the investment returns are examined. A vector of signals s is considered; what makes these signals interesting to the portfolio manager is that there is assumed to exist some measurable relationship $\mu(s)$ between the signal as observed at the initial time and the return as revealed at the final time. The fundamental signal-return relationship is then

$$r_{t+1} = \mu(s_t) + \epsilon, \tag{1}$$

where the time indices on r and on s will be suppressed in what follows. Here ϵ is a noise term whose conditional mean given s is assumed to be zero, and there is no specific a priori requirement on the functional form of $\mu(s)$.

The Ferson and Siegel paper derives expressions for mean-variance optimal portfolio weight functions given the signal-return relation (1) holds and whether a risk-free asset is available or not. Further closed-form solutions covering the benchmark tracking error minimisation versions of the basic problem are derived in [5]. Empirical illustrations respectively studies of the solutions to these problems are provided in papers such as [5], [6] or [7]. [8] report on an optimisation approach comparable, but not identical, to that proposed by Ferson and Siegel, in which quadratic utility is directly maximised, and the index value is used in a quadratic term that implements a penalty for leverage.

Except for [8], all of the above papers cover the basic version of the conditioned portfolio problem for which admissible weights are not constrained. As in the case of the classic Markowitz problem, constrained variants of the standard Ferson and Siegel formulation can generally only be solved by introducing a numerical scheme. In the case where a single signal series is used, [2] describes how to express conditioned optimisation problems in optimal control terms and proves extended versions of the relevant classical necessity and sufficiency results which are the Pontryagin Minimum Principle and the Mangasarian sufficiency theorem. This implies that generic variants of the basic problem, for which the signal support may in general correspond to the full real axis, can be analysed like any other optimal control problem, as well as solved numerically using any of the standard available discretisation schemes. We note in this context that the typical conditioned problem involves optimisation of the *unconditional* mean with respect to the *unconditional* variance even though the investment manager by assumption has access to the signalling information. Indeed the optimality of that approach has been confirmed by a generic function theoretical argument (see [4]), by financial intuition (the manager is evaluated by generally uninformed investors, who judge manager performance based on their observation of the unconditional moments, see [1]) and by empirical comparisons with the alternative strategy where conditional moments are optimised (see [5]). We thus follow this problem setting.

2.2 Portfolio optimisation involving higher moments of returns

A second direction of research in portfolio optimisation has concerned itself with the different ways in which it is possible to consider moments of portfolio returns beyond mean and variance when specifiying investor preferences and formulating the optimisation problem. The two higher moments of particular interest have been the third moment, or skewness, and the fourth moment, or kurtosis. Here, skewness allows for the representation of non-symmetric preferences, which mean and variance do not. A positive skewness implies a possibility of large positive outcomes greater than would be available for a symmetrical returns distribution with identical mean and variance. Consequently, investors are expected to prefer positive skewness and portfolio research will aim to maximise skewness for a certain level of risk, see e.g. [9] or [10]. Kurtosis, on the other hand, is a risk metric which may be considered instead of, or in addition to, variance. A leptokurtic distribution is 'peaked' with respect to the standard normal. As such more probability mass is found in the tails compared to the benchmark, and extreme outcomes are proportionally more likely. Unlike the case of variance, there is some slight ambiguity in the use of kurtosis as a risk measure: in particular, as pointed out in [11], one possible infinite-kurtosis limit corresponds to the returns distribution converging to a Dirac measure, which is hence risk-free. Also, as is the case with all even moments, kurtosis is unable to distinguish between the two tails of the distribution and thus penalises long-tailed positive outcomes just as much as

negative ones. In common with most of the optimisation literature involving higher moments (see e.g. [12], [13] or [14]), we will nevertheless minimise kurtosis, rather than maximise it as done in [11][1]. Overall, we feel that the minimisation of the probability of extreme losses is definitely a desirable feature of optimisation problems and ties in with the use of risk management metrics such as VaR or CVaR in portfolio optimisation (see e.g. [15] or [16]): this is not the case to the same extent with classical mean-variance optimisation.
The specific type of conditioned optimisation problem this paper focuses on starts with the basic problem as analysed in [1] and mentioned in the above. Use of the optimal control formulation justified in [2] allows for the addition of portfolio weights constraints. A particularly interesting constraint for real-world applications is obtained when we exclude the possibility of negative investment amounts i.e. short positions. For legal or risk management reasons, these may not be available to investors. The resulting mean-variance constrained-weight conditioned optimisation problem (with allowable weight values in $[0, 1]$ given the sum of portfolio weights must always give 1) is discussed and evaluated, using the same data set as in the present article, in [3]. In the present paper, we merge that problem with higher-order optimisation specifically by considering the fourth moment of returns. While optimisation of the skewness is conceptually another attractive proposition, it is only covered briefly in this paper given that objective functions dominated by skewness terms are nonconvex. It is true that all optimisation problems defined in this paper contain nonlinear equality constraints and are thus nonconvex, such that any optima found may only be local in any case. However, we have found in practice that the use of an appropriate solver[2] yields solutions that consistently appear global if the objective function remains convex, but may fail convergence tests or clearly correspond to local optima if convexity of the objective function is violated. Hence we initially prefer to retain a convex objective function, which is the case for the problem based on either absolute kurtosis or the uncentred fourth moment of returns. Having thus obtained convincing three-moment results, we then list the results achieved when taking skewness into account and observe that, for the specific data set used at least, these do not point to any significant numerical issues.

2.3 Data

The data set used collects eleven years of daily returns data chosen to represent a market relevant to investors with domestic currency EUR. This market is made up of ten different funds[3] chosen across both equity and fixed income markets as well as Morningstar style classifications. All funds involved provide EUR return quotes and manage at most a proportion of 30% in non-EUR assets such that the impact of currency risk on the choice of investments remains reasonable. The data covers business days from January 1999 to February 2010: in total, each series contains 2891 returns. Funds rather than individual assets were chosen given they provide a level of built-in diversification and a ten-asset universe composed of funds is thus seen as more attractive than a universe of similar size composed of individual equity assets; additionally, interest-rate exposure is easily achieved through funds. Investment strategies involving funds and requiring frequent portfolio rebalancings, such as the ones being examined, have become realistically achievable even for small investors with the advent of exchange traded funds (ETF). Although the funds listed above are not ETFs, this choice was made purely because of the need to obtain sufficiently exhaustive historical data series: actively managed funds comparable to those used are nowadays accessible in an ETF format. Summed log-returns for each asset vary from -27.91% to 40.43% over the backtesting period: summary statistics for the asset market are given in table 1. These show, in particular, that returns cannot realistically be modelled as normal - a stylised fact that is generally assumed for daily return series, see e.g. [18].
Of particular interest is the fact that the data covers two periods of crisis. The first of these is defined by the bursting of the dot-com bubble: stock market implications of this crisis are strongly visible in the data set used over a period spanning spring 2000 to spring 2003, with an average log return of -74.28% for the ten assets under consideration. The second corresponds to the initial bear market linked to the financial crisis

[1] We also disagree with the contention in [11] that kurtosis minimisation coincides with variance minimisation. Both problems are different given that the kurtosis formulation involves the fourth comoment matrix, which expresses asset return dependencies not present in the covariance matrix.

[2] We have used the interior-point solver integrated within the `fmincon` function found in the Optimization Toolbox of MATLAB. This incorporates a sequential quadratic programming (SQP) loop to linearise the problem equality constraints if necessary: see [17] for details.

[3] AXA L Fund Equity Europe (AXA), Credit Suisse Bond Fund Management Company Luxembourg Small (CSU), Dekalux Midcap TF (DEK), Dexia Luxpart C (DEX), DWS Euro Bonds Long (DWS), Fidelity Funds Euro Bond Fund A Global Certificate (FIB), Fortis L Fund Equity Socially Responsible Europe (FOB), Invesco Pan European Small Cap Equity Fund Lux (INV), KBC Money Euro Medium Cap (KBC) and Morgan Stanley European Currencies High Yield Bond (MSE). In every case the reinvesting variant of the fund was picked.

Asset	Cumulative log return	Standard deviation	Skewness	Kurtosis	Jarque-Bera statistic
AXA	-27.91%	1.343	-0.092	8.989	4306
CSU	22.56%	1.386	-0.235	9.031	4390
DEK	3.55%	1.338	-0.670	6.744	1897
DEX	12.91%	1.236	-0.155	13.369	12914
DWS	40.43%	0.208	-0.400	6.235	1332
FIB	29.86%	0.146	-0.489	4.884	540
FOB	-27.54%	1.308	-0.104	9.294	4757
INV	-18.07%	1.238	-0.957	9.159	4991
KBC	30.01%	0.081	-1.645	23.959	54024
MSE	-6.62%	0.481	-1.506	20.719	38771

Table 1: Cumulative log returns, return standard deviation, (relative) skewness, (relative) kurtosis and Jarque-Bera statistic of individual assets over the entire data set.

still ongoing at the time of writing; its impact on the data set is seen on the interval from summer 2007 to spring 2009, over which the average log return observed across the ten asset market is -59.66%.

For reasons of space, the market variant discussed is that in which no risk-free asset is available. We concentrate on that case as it seems more realistic given truly risk-free assets do not exist and the usual proxies cannot (as is generally assumed) be entered without penalty in any required position size. Additionally, the market chosen incorporates several large cap bond funds and one money market fund whose risk level is rather low in any case, such that they constitute some approximation of risk-free assets in themselves. The case with risk-free asset does not add any fundamental complexity to the problem, and backtests executed using the one week EURIBOR rate as a risk-free proxy were carried out to confirm that all results obtained are compatible with the case reported.

Attractiveness and performance of the conditioning approach strongly depends on the quality of the signal used. In [3], several different signals are exercised, with the conclusion that the 'pure' equity risk signal represented by the VDAX index (i.e. the DAX index equivalent of the CBOE volatility index VIX) performs best, at least for the data set under consideration. Accordingly, the VDAX signal is also used in the present study.

3 Results

In line with most of the literature on portfolio optimisation involving kurtosis (see e.g. [13] or [14]), we concentrate on *absolute kurtosis*, i.e. the fourth central moment of portfolio returns. The equivalent quantity normalised with respect to portfolio standard deviation is usually known as the *relative* kurtosis. While the latter has the advantage of scaling kurtosis to have a comparable order of magnitude to standard deviation, it leads to numerically complex optimisation problems given the quotient of convex functions need not be convex itself. This is not an issue with the absolute form of the moment, which has the additional appeal of directly generalising the expression used for portfolio variance. The fourth moment is introduced in two separate ways. Initially (see subsection 3.1), the classical mean-variance problem is adapted by simply replacing the pertinent risk metric by absolute kurtosis. In this mean-kurtosis (MK) problem, the investor then aims to minimise expected absolute kurtosis for a given level of expected return. Subsection 3.2 defines a number of three-term polynomial utility functions to carry out optimisation strategies using three uncentred moments at once - the results may be called mean-variance-kurtosis (MVK) optima even though, strictly speaking, variance and kurtosis are not used in that case. Finally, subsection 3.3 reports on results obtained using quartic polynomial utility functions with four nonzero terms: these suggest that, for the present data context, the objective function nonconvexities induced by the third-order term do not cause any numerical issues, although the universality of that conclusion cannot be guaranteed.

3.1 Mean-kurtosis (MK) optimisation

Even though the only other paper we are aware of that considers the simple mean-kurtosis problem is [11], the tail risk connotations held by the fourth moment, along with its numerical tractability, nevertheless make it an interesting object for study. The possibility of improving the kurtosis-expected return tradeoff through the use of conditioning information within an investment context which prohibits short positions is then attractive. To set up that problem, a formulation for the unconditional absolute kurtosis of the portfolio returns in the presence of the conditioning information relationship (1) is initially obtained by replacing and applying the law of iterated expectations.

We reuse the Kronecker product notation used in e.g. [12], with which portfolio expected skewness and kurtosis in the absence of conditioning information, and given coskewness and cokurtosis matrices S^3 and κ^4 respectively, can be expressed as

$$E\Big[(P-\mu_P)^3\Big] = u'S^3(u \otimes u) \tag{2}$$

and

$$E\Big[(P-\mu_P)^4\Big] = u'\kappa^4(u \otimes u \otimes u) \tag{3}$$

given scalar portfolio weights u, an observed portfolio return P and a required portfolio return μ_P. In these formulations, $\otimes$ denotes the Kronecker product operator: thus the coskewness matrix has order $n \times n^2$ and the cokurtosis matrix has order $n \times n^3$.

We now consider the case where conditioning information is present. Introduce portfolio weight functions of the signal value $u(s)$ and define $\Gamma(s) = u'(s)\mu(s)$. Using, in particular, the law of iterated expectations with the relationship $E[\epsilon|s] = 0$, the expression for the unconditional absolute kurtosis of the expected portfolio return P in the presence of conditioning information is then found as

$$\begin{aligned} E\Big[(P-\mu_P)^4\Big] = E\Bigg[&(\Gamma(s))^4 + 6(\Gamma(s))^2u'(s)\Sigma_\epsilon^2u(s) + \\ &4\Gamma(s)u'(s)S_\epsilon^3(u(s)\otimes u(s)) + u'(s)\kappa_\epsilon^4(u(s)\otimes u(s)\otimes u(s)) - \\ &4(\Gamma(s))^3E[\Gamma(s)] - 12\Gamma(s)u'(s)\Sigma_\epsilon^2u(s)E[\Gamma(s)] - \\ &4u'(s)S_\epsilon^3(u(s)\otimes u(s))E[\Gamma(s)] + 6(\Gamma(s))^2\Big(E[\Gamma(s)]\Big)^2 + \\ &6u'(s)\Sigma_\epsilon^2u(s)\Big(E[\Gamma(s)]\Big)^2 - 3\Big(E[\Gamma(s)]\Big)^4\Bigg]. \end{aligned} \tag{4}$$

Here the $n \times n$ matrix $\Sigma_\epsilon^2 = E[\epsilon\epsilon'|s]$ is the conditional covariance matrix implied by the fitted relationship (1), the $n \times n^2$ matrix $S_\epsilon^3 = E[\epsilon(\epsilon' \otimes \epsilon')|s]$ is the conditional coskewness matrix and the $n \times n^3$ matrix $\kappa_\epsilon^4 = E[\epsilon(\epsilon' \otimes \epsilon' \otimes \epsilon')|s]$ is the conditional cokurtosis matrix. Now denote $K(s)$ the quantity within the expectation operator in (4). Then the optimal control formulation of the MK problem involves minimising the expectation integral cost function

$$J_{D_S}(x(s), u(s)) = \int_{s^-}^{s^+} K(s)p_S(s)ds.$$

Here D_S is the support of s spanning the interval from s^- to s^+, where neither boundary needs to be finite given the developments of [2], and $p_S(s)$ is the density of s. Additionally, a problem state variable described by the differential equation $\dot{x}_1 = u'(s)\mu(s)p_S(s)$ enforces the expected returns constraint using the terminal condition $x_1(s^+) = \mu_P$.

The backtesting procedure then computes the efficient frontier at each rebalancing point (using a standard *direct collocation* optimal control discretisation algorithm, see e.g. [20] for details) and obtains ex post observed returns for hypothetical investors targeting different expected returns. Table 2 gives the sample mean values generated by the three strategies under consideration[4] and for different performance metrics as obtained both ex ante and ex post. Results are shown for every fourth point on the 21-point discretised efficient frontiers, such that the portfolio expected return, and thus the anticipated risk the investor accepts, increases with each point. We initially consider the ex ante figures obtained. While the normal Sharpe ratio (SR) is indicated, it clearly ignores kurtosis and thus best fits the intent of the mean-variance solution. Plausibly, the ex ante Sharpe ratios obtained for the unconditioned mean-kurtosis strategy are lower than the unconditioned mean-variance ones across the entire range of points. For higher risk points, however, the conditioned mean-kurtosis strategy SRs are highest: this can be seen as interesting evidence of how the use of signalling information can fundamentally add to the power of the resulting solution.

A straightforward extension of the classical Sharpe ratio for contexts involving higher-order moments is suggested in [19]. Using a Taylor series expansion of an exponential utility function truncated after the first

[4] Unconditioned mean-variance (Markowitz) optimisation (MV UNC), unconditioned and conditioned mean-kurtosis optimisation (MK UNC, MK CON)

Point #	1	2	3	4	5	6
			Ex ante			
MV UNC std	0.056	0.099	0.175	0.279	0.412	0.640
MV UNC SR	0.199	0.407	0.397	0.354	0.311	0.246
MV UNC rel. skewness	-20.632	-7.241	-6.160	-3.753	-2.010	-0.489
MV UNC rel. kurtosis	18.793	17.260	29.265	30.668	28.998	27.470
MV UNC abs. kurtosis	0.0002	0.0017	0.0277	0.1871	0.8328	4.6008
MV UNC ASR	0.057	0.159	0.159	0.219	0.242	0.224
MK UNC std	0.057	0.102	0.179	0.284	0.416	0.640
MK UNC SR	0.193	0.397	0.390	0.348	0.308	0.246
MK UNC rel. skewness	-5.787	-3.864	-4.750	-3.198	-1.762	-0.478
MK UNC rel. kurtosis	9.889	12.747	23.518	25.667	25.880	27.344
MK UNC abs. kurtosis	0.0001	0.0014	0.0240	0.1668	0.7765	4.5928
MK UNC ASR	0.154	0.262	0.211	0.238	0.249	0.224
MK CON std	0.128	0.141	0.167	0.211	0.266	0.334
MK CON SR	0.087	0.286	0.416	0.470	0.482	0.472
MK CON rel. skewness	-1.928	-0.284	0.638	1.005	1.300	1.563
MK CON rel. kurtosis	6.372	5.534	9.233	17.460	25.207	35.008
MK CON abs. kurtosis	0.002	0.002	0.007	0.034	0.126	0.436
MK CON ASR	0.084	0.277	0.407	0.431	0.415	0.377
			Ex post			
MV UNC return	0.009	0.023	0.035	0.043	0.047	0.037
MV UNC std	0.073	0.138	0.256	0.402	0.556	0.805
MV UNC SR	0.129	0.165	0.137	0.107	0.084	0.045
MV UNC rel. skewness	-2.147	-0.650	-0.463	-0.569	-0.699	-0.709
MV UNC rel. kurtosis	27.904	14.623	15.419	14.403	11.242	9.452
MV UNC abs. kurtosis	0.001	0.005	0.066	0.376	1.074	3.964
MV UNC ASR	0.120	0.160	0.134	0.105	0.083	0.045
MK UNC return	0.009	0.023	0.034	0.041	0.046	0.037
MK UNC std	0.078	0.143	0.262	0.408	0.567	0.807
MK UNC SR	0.112	0.159	0.130	0.101	0.080	0.045
MK UNC rel. skewness	-2.109	-0.790	-0.753	-0.805	-0.748	-0.709
MK UNC rel. kurtosis	30.211	13.905	15.115	13.987	11.811	9.480
MK UNC abs. kurtosis	0.001	0.006	0.071	0.388	1.219	4.000
MK UNC ASR	0.106	0.153	0.127	0.099	0.079	0.045
MK CON return	0.008	0.018	0.027	0.039	0.052	0.060
MK CON std	0.154	0.171	0.218	0.293	0.371	0.461
MK CON SR	0.054	0.107	0.124	0.134	0.140	0.129
MK CON rel. skewness	-1.996	-0.995	-0.653	-0.340	0.190	0.088
MK CON rel. kurtosis	17.075	10.779	16.010	18.145	15.545	13.806
MK CON abs. kurtosis	0.009	0.009	0.036	0.133	0.293	0.620
MK CON ASR	0.053	0.104	0.121	0.132	0.139	0.128

Table 2: Mean (ex ante and ex post) metrics of portfolio returns obtained for unconditioned (UNC) and conditioned (CON) mean-variance (MV) respectively mean-kurtosis (MK) optimisation. In particular, SR is the (unmodified) Sharpe ratio and ASR the adjusted Sharpe ratio suggested by [19].

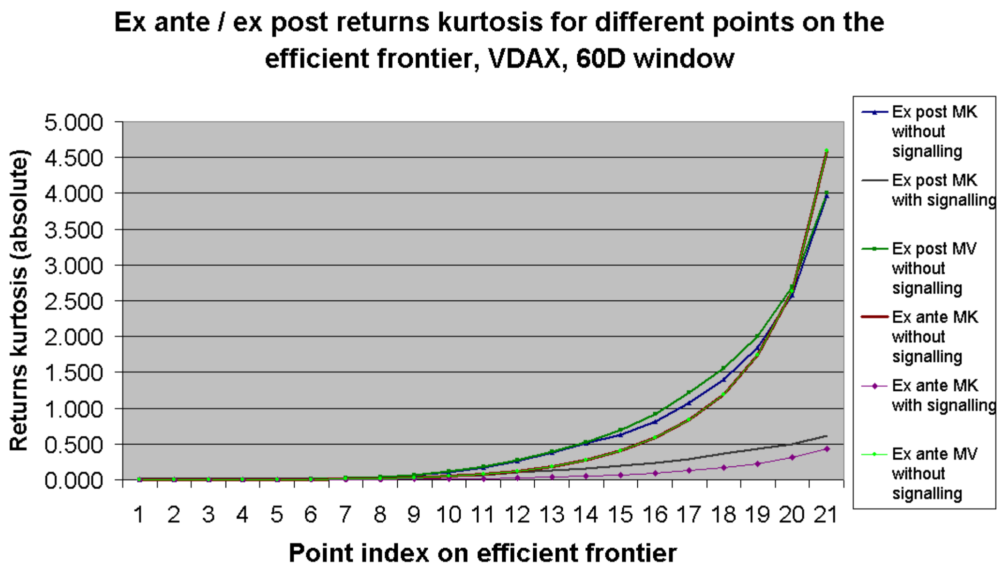

Figure 1: MK ex ante and ex post mean absolute returns kurtoses for all 21 points on the discretised efficient frontier.

four terms, the paper formulates an adjusted Sharpe ratio (ASR) as

$$ASR = SR\left[1 + \frac{1}{6}s_r^3 - \frac{1}{24}\kappa_r^4 SR^2\right],$$

where s_r^3 and κ_r^4 are the *relative* versions of skewness and kurtosis discussed in what precedes. The ASR is given in table 2 for all cases covered, along with the *relative* skewness and kurtosis values, which are not, as discussed, optimised by the given strategy but necessary for the calculation of the ASR. Differences between SR and ASR are most noticeable in those cases where high absolute values of relative skewness and kurtosis are observed, such as the first point for the mean-variance strategy. For smaller values of skewness and kurtosis, ASR values approach SR values. Compared across the three strategies, the ASR values are consistent. In particular, the unconditioned mean-kurtosis figures are higher than the unconditioned mean-variance ones. For all but the least risky points covered, the conditioned values are substantially more attractive still. Note that this is not a necessary given as the absolute rather than relative kurtosis is used in the optimisation, and the skewness is not optimised at all. ASR improvements are not observed for the two minimum risk points, where the use of conditioning information is seen to entail decreases in both SR and ASR. As can be seen in table 2, this is not because of relative skewness or kurtosis, for which the most preferable values are obtained for the conditioned mean-kurtosis procedure across the entire range of points. Instead, the standard deviation is lower for unconditioned than for conditioned mean-kurtosis, leading to a larger SR - an observation that does not reflect on the MK optimisation strategy discussed as it does not take standard deviation into account.

The same observation holds for absolute kurtosis, which is of course the quantity the mean-kurtosis algorithms are aiming to minimise. As with standard deviation, this is highest for the conditioned optimiser at very low levels of risk. However, it is observed that all strategies yield very low kurtosis values at these levels, to the point where this difference may be due to numerical errors and is seen as largely immaterial. As investor risk appetite is increased, ex ante kurtosis levels obtained for the conditioned strategy remain at a low level while the unconditioned benchmark levels rise much more strongly. Figure 1 graphically depicts this observation. Finally, skewness is observed to be significantly more attractive using the conditioned optimiser even though it is not being considered in the problem formulation, such that this may be a coincidence.

Overall, the figures just discussed suggest that the use of signalling information thus generates significant ex ante benefits over traditional approaches. This general tendency remains true in the ex post case. While the use of conditioning information leads to lower values of SR and ASR for low levels of risk, the results observed without signalling degrade significantly from about the middle of the range of expected returns considered whereas conditioned performance remains stable or even slightly increases, leading to a large differential at the high-risk end of the range. Similar observations hold for relative skewness and absolute kurtosis (the latter of which, in particular, increases exponentially with risk using both unconditioned strategies - see figure 1), while no pronounced tendency is visible for relative kurtosis.

It is, however, suggested that the low-risk points near the minimum-variance point correspond to a level of risk aversion so high that it is not relevant in practice. The ex post standard deviations obtained for the unconditioned mean-variance strategy at the first few points on the 21-point efficient frontier are $0.073, 0.078$ and 0.092, and the three first levels of additive cumulative returns obtained by the present backtest are 27.22%, 38.72% and 48.43%. By comparing to those obtained for the utility function based optimisation in the corresponding backtest of [3], it can be seen that all three of these values show less risk and lower returns than were realised for the highest quadratic risk aversion coefficient of $\lambda = 10$. This suggests that the part of the investment range of practical interest in the present is given by the riskier half of the range, i.e. precisely the interval over which conditioned optimisation outperforms unconditioned optimisation by a larger and larger margin.[5]

Next considering the standard deviation of returns, the relevant entries in table 2 confirm that the decreases in kurtosis are not achieved through increases in the standard deviation for both the ex ante and ex post values. Again, the mean-variance and mean-kurtosis unconditioned optimisers achieve results very close to one another. The standard deviations of returns yielded by the conditioned procedure are higher than for the unconditioned methods at the low expected return points. However, their levels are more and more significantly lower than those obtained using the mean-variance optimiser as the investor aims to achieve higher expected returns while accepting that the resulting risk may be higher. This is a very strong result as the conditioned mean-kurtosis optimal portfolios thus exhibit significantly less risk over the more interesting part of the risk-return function domain, *according to either risk metric definition used*, than the dedicated mean-variance optimal portfolio.

Finally, the mean additive returns obtained are also given in table 2. For the lower expected returns half of the efficient frontier considered, conditioned mean-kurtosis optimisation offers the sort of trade-off that it would be plausible to expect in advance when choosing the kurtosis risk metric: a significant reduction in portfolio kurtosis is achieved in exchange for an increase in standard deviation and a decrease in the ex post returns figure. However, the conditioned mean-kurtosis portfolio becomes very advantageous for higher target returns. As the unconditioned strategies show decreasing returns, the ex post return curve in the presence of signalling remains almost linear in the expected returns index, such that both unconditioned optima are seen to be dominated in expected returns as well as both risk metrics considered.

In particular, the diminishing returns shown by the unconditioned series may be conjectured an alternative manifestation of the exponentially increasing ex post risk metric levels: clearly, the high risk of the resulting portfolios is likely to severely affect their crisis resilience in particular. To check this suggestion for the extreme case of the riskiest portfolio used, figure 2 plots the corresponding time path of additive returns as delivered by both mean-kurtosis strategies. The unconditioned path is visibly more volatile than the conditioned path. An examination of the figure furthermore shows that the unconditioned strategy suffers from two protracted periods of severe drawdowns corresponding to the two crises contained in the data set used, while the conditioned returns experience a milder drawdown period during the collapse of the internet bubble, and traverse the crisis starting in 2008 with comparatively small drawdowns. This graphical observation is suggested by the drawdown figures obtained as well: while the maximum drawdown (MD) observed for the unconditioned strategy is as much as 70.38%, with a maximum drawdown duration (MDD) of 1026 days, the conditioned MD is much lower at 23.91% with an MDD of 865 days.

In conclusion, the given results consistently show that, for the data set used, introduction of conditioning information visibly and measurably adds to the crisis resilience of the resulting portfolios, and results in investment strategies that dominate their classical equivalents for every metric considered as long as a minimum level of risk is accepted by the investor.

3.2 Mean-variance-kurtosis (MVK) optimisation

Clearly, any extension of the optimisation problem formulation to span more than two moments at the same time requires the investor to specify their respective preferences across these moments. One way of achieving this has been the specification of a moment-dependent goal function in the context of polynomial goal

[5] The minimum-variance point of course reflects the presence of, in particular, the KBC money market fund, which approximates a risk-free asset.

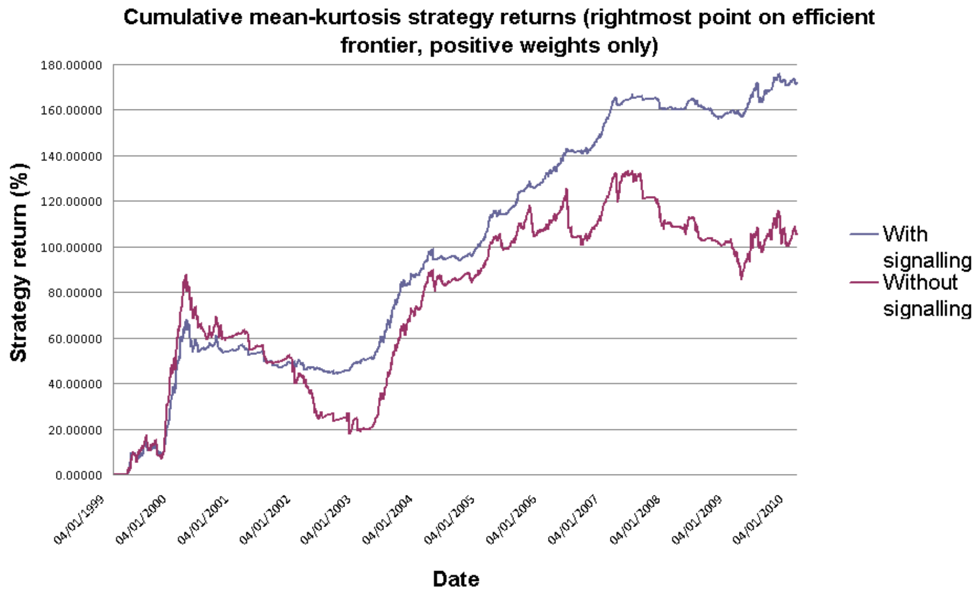

Figure 2: MK ex post returns time path for highest attainable expected return point.

programming, or PGP (see e.g. [9] or [21]). These PGP goal functions involve distances from the different portfolio moments involved to their 'aspired levels' as found through initial unconstrained optimisation passes for each moment. Accordingly, they intuitively, although not formally, correspond to polynomial utility functions, which provide for the moment-ordering device that will be preferred in what follows. The use of polynomial utility functions constitutes a formally correct way of describing investor preferences that consider only the initial p moments within the von Neumann-Morgenstern utility framework (see [22] for the original exposition of utility axioms): the equivalence between such a set of investor preferences and consistent use of a p-th degree polynomial as the investor's utility function is shown in [23].

A quartic polynomial utility function used to express preferences over the first four moments of investment returns has the form

$$U(x) = a_1 x + a_2 x^2 + a_3 x^3 + a_4 x^4, \tag{5}$$

where the present subsection's discussion will consider the subset of functions for which $a_3 = 0$ so as to avoid objective function nonconvexity associated with third moment preferences.[6] The polynomial utility functions thus use uncentred moments, which can be seen as proxies to the centred moments of interest. Classical results (see e.g. [24]) show that any optimisation carried out with respect to uncentred moments will necessarily involve the pertinent comoment matrices as well, leading to what is at least a related problem. From their point of definition in [22], various desirable properties have been associated with utility functions. While the main point of focus of the present paper is given by the moment optimisation problems proper, we should nevertheless attempt to use utility functions that verify as many of those properties as possible, over as large a domain as possible. This necessarily requires a compromise. For one thing, while expected positive daily returns in the data set used span the range from 0 to 1.0125%, $E[X^p] \leq (E[X])^p$ by Jensen's inequality in general for any integer $p \geq 2$ and any random variable X, such that this inevitably only gives a rough indication of the range which would need to be covered. Additionally, it turns out that the above range does not allow for sufficient flexibility in the coefficients of quartic polynomial utility functions that verify the properties discussed in what follows, such that an exclusive focus on these properties would not allow for the resulting functions to reflect significantly varying preferences with respect to the different expected moments of returns. Noting that most daily returns in the data set used are comprised between -0.5% and 0.5%, we will therefore introduce two sets of utility functions that are well-behaved over two intervals which are at least pertinent to the context.

Two desirable higher-order properties of utility functions were introduced by Kimball in the articles [25] and [26] respectively. [25] names *prudence* the property that marginal utilities are convex, while [26] defines *temperance* as the equivalent fourth-order property, corresponding to convex second derivatives of the utility

[6] The constant term a_0 can be absorbed into the remaining coefficients with no loss in generality.

	a_1	a_2	a_3	a_4
Function #1	0.1001	-0.0201	0	-0.0101
Function #2	0.4901	-0.0001	0	-0.0901
Function #3	0.9600	-0.4301	0	-0.0001
Function #4	1.0000	-0.0201	0	-0.1701
Function #5	1.0000	-0.9601	0	-0.0401
Function #6	1.0000	-0.5601	0	-0.8401
Function #7	1.0000	-0.0401	0	-1.8801
Function #8	1.0000	-0.0401	0	-0.0401
Function #9	1.0000	-0.5000	0	0.0000

Table 3: Polynomial coefficients for the MVK utility functions (indexed from #1 to #9) used.

functions. Jurczenko and Maillet ([27], theorem 4) obtain a set of sufficient conditions on the coefficients vector of (5) to guarantee that the properties of non-satiability, risk averseness, prudence and temperance are all jointly verified over the range of returns for which the given conditions hold.

While the present section's requirement that $a_3 = 0$ means that the resulting utility functions cannot implement investor prudence[7], the other three properties turn out to be feasible for functions thus constrained over the range of returns reflected in the data set used. Utility functions in the first set of four given in table 3.2 verify the theorem's conditions, with the exception of that required for investor prudence, over the range from 0 to 1.0125%: clearly, portfolios with a negative expected return do not yield positive utility values i.e. are dominated by a null position and can thus be excluded a priori. However, only a very limited amount of flexibility is available as regards possible values for the fourth moment coefficient a_4 if the resulting utility function is still to verify the remaining conditions. To visualise the impact of stronger aversions to the fourth expected moment of returns, the second set of functions, from #5 to #8, accepts that the same set of conditions is only verified over the interval from 0 to 0.5. Finally, function #9 is a quadratic function included as a comparative benchmark.

As with the previous variant of the conditioned optimisation problem, the strategies were checked using a backtesting setup. In this case, no efficient frontiers were constructed as these require expected return inputs. Instead, each utility function was directly maximised within the given portfolio weight constraint, and the resulting weights function interpolated at the observation point of the VDAX signal. Since

$$E[U(r)] = E\Big(a_1 r + a_2 r^2 + a_4 r^4\Big)$$

for the functions under consideration which verify $a_3 = 0$, the optimal control problem thus involves minimising the cost function

$$J_{D_S}(x(s), u(s)) = -\int_{s^-}^{s^+} (a_1 \dot{x}_1(s) + a_2 \dot{x}_2(s) + a_4 \dot{x}_4(s))ds.$$

Here we have used the same notation as in the previous subsection. Additionally, the problem state variables corresponding to the unconditioned expected uncentred first, second and fourth moments of returns in the presence of conditioning information are described by the differential equations

$$\dot{x}_1 = u'(s)\mu(s)p_S(s),$$

$$\dot{x}_2 = \Big((u'(s)\mu(s))^2 + u'(s)\Sigma_\epsilon^2 u(s)\Big)p_S(s)$$

[7] This is also the reason why the stronger conditions given by theorem 5 in [27], which specify quartic polynomial utility functions that exhibit decreasing absolute risk aversion, decreasing absolute prudence and constant or increasing relative risk aversion, are not met by the given set: nonzero preferences over all four moments are necessary for this.

and

$$\dot{x}_4 = \Big((u'(s)\mu(s))^4 + 6(u'(s)\mu(s))^2 u'(s)\Sigma_\epsilon^2 u(s) + 4u'(s)\mu(s)u'(s)S_\epsilon^3(u(s)\otimes u(s)) + u'(s)\kappa_\epsilon^4(u(s)\otimes u(s)\otimes u(s)) \Big) p_S(s),$$

where each expression is as before obtained through replacement of the assumed signal relationship (1) and use of the law of iterated expectations.

Utility function #	1	2	3	4	5	6	7	8	9 (MV)
Ex ante									
MVK UNC SR	0.408	0.369	0.451	0.372	0.486	0.466	0.432	0.345	0.457
MVK UNC rel. skewness	-0.559	-0.507	-0.603	-0.512	-0.268	-0.300	-0.328	-0.561	-0.543
MVK UNC rel. kurtosis	11.308	8.620	10.259	8.895	8.560	7.012	5.673	11.852	15.119
MVK UNC abs. kurtosis	0.049	0.073	0.030	0.072	0.002	0.003	0.004	0.253	0.021
MVK UNC ASR	0.361	0.339	0.367	0.341	0.435	0.426	0.403	0.314	0.378
MVK UNC utility	0.008	0.047	0.057	0.094	0.037	0.046	0.059	0.110	0.056
MVK CON SR	0.487	0.459	0.518	0.461	0.577	0.515	0.456	0.427	0.528
MVK CON rel. skewness	0.471	0.281	0.819	0.299	0.533	0.312	0.179	0.413	0.796
MVK CON rel. kurtosis	6.828	5.913	8.954	6.004	6.607	6.089	5.639	7.348	8.583
MVK CON abs. kurtosis	0.081	0.099	0.066	0.101	0.005	0.006	0.008	0.372	0.047
MVK CON ASR	0.473	0.445	0.503	0.447	0.553	0.494	0.440	0.415	0.512
MVK CON utility	0.012	0.067	0.090	0.136	0.054	0.057	0.064	0.168	0.087
Rel. utility improvement	50.781%	44.377%	56.054%	45.327%	46.998%	26.097%	8.745%	52.287%	54.888%
Ex post									
MVK UNC SR	0.137	0.119	0.154	0.121	0.180	0.173	0.154	0.109	0.157
MVK UNC rel. skewness	-1.692	-1.715	-1.621	-1.732	-3.063	-2.674	-2.574	-0.969	-1.818
MVK UNC rel. kurtosis	19.792	18.159	23.259	18.452	46.732	34.951	29.363	12.557	25.104
MVK UNC abs. kurtosis	0.350	0.545	0.212	0.545	0.064	0.068	0.093	0.898	0.177
MVK UNC ASR	0.130	0.114	0.144	0.115	0.152	0.152	0.139	0.107	0.145
MVK UNC utility	-0.001	-0.023	0.004	-0.042	-0.005	-0.044	-0.133	0.011	0.002
MVK CON SR	0.186	0.174	0.202	0.176	0.226	0.191	0.171	0.153	0.206
MVK CON rel. skewness	-0.368	-0.432	-0.218	-0.432	-1.834	-1.537	-1.530	0.025	-0.447
MVK CON rel. kurtosis	13.166	12.075	16.276	12.132	32.263	22.623	22.603	9.431	17.997
MVK CON abs. kurtosis	0.293	0.334	0.278	0.336	0.064	0.053	0.061	0.725	0.231
MVK CON ASR	0.180	0.170	0.194	0.171	0.195	0.175	0.159	0.152	0.196
MVK CON utility	0.001	0.005	0.011	0.012	0.000	-0.028	-0.074	0.040	0.010

Table 4: Mean (ex ante and ex post) metrics of portfolio returns obtained for unconditioned (UNC) and conditioned (CON) polynomial utility function based mean-variance-kurtosis (MVK) optimisation. In particular, SR is the (unmodified) Sharpe ratio and ASR the adjusted Sharpe ratio suggested by [19].

The backtesting result set obtained for this problem consists of eleven different series, i.e. one for each utility function applied. Results are shown in table 3.2. Initially, note that both Sharpe ratios and adjusted Sharpe ratios are improved through conditioning for all of the utility functions used, ex ante as well as ex post. This observation can be taken as further support of the suggestion, made in the preceding discussion of mean-kurtosis optimisation where the use of signalling information worsened results close to the minimum-variance point, that the minimum-risk points considered there were not relevant for most practical optimisation contexts: in these MVK problems, by construction, only points that encode interesting compromises between expected return and risk are considered. Even so, given the MVK strategy optimises purely a function of uncentered moments, it is not by definition indicative of outperformance, although positive, that SR and ASR results are improved by conditioning. Similarly, it is interesting to observe that relative skewness is increased, and relative kurtosis decreased, for the conditioned strategy, across the entire set of utility functions and ex post as well as ex ante. Still, visible outperformance corresponds to increases in investment utilities: ex ante (expected) and ex post (realised) figures are shown both in table 3.2 and in figure 3.

By looking at the ex ante utility scores, the added value of conditioning information is apparent, with most scores showing an increase around 50%.[8] The only exceptions to this are utility function #6, which penalises both second and fourth moments as much as was feasible given the conditions from [27], and function #7, which models the largest aversion to the fourth moment in the set used. It seems plausible to suggest these again represent limit cases, for which the risk minimisation component of the optimisation tradeoff becomes predominant, such that there remains little scope for manoeuvre. This limit status is in a way confirmed by the ex post utilities, where #6 and #7 yield the only negative utilities for the conditioned strategy. However, the central observation here is that the conditioned approach again outperforms the unconditioned approach ex post for all utility functions. Relative improvements are seen to be substantial, but cannot

[8] While the specific utility scores are clearly a reflection of the decision to use a utility function constant term of zero, this increase will always be observed in the moment dependent parts of the utility scores obtained and it thus seems a meaningful measure of outperformance.

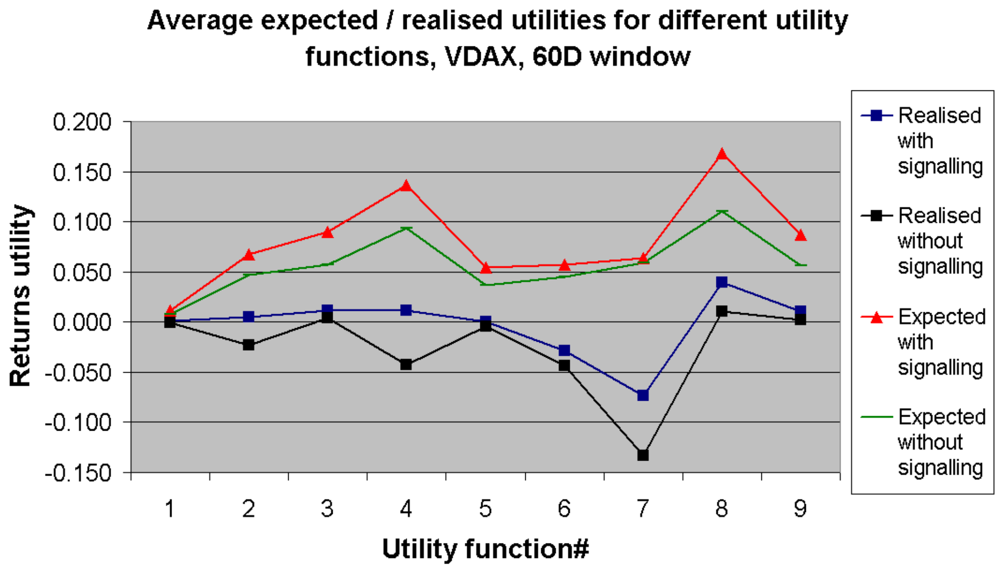

Figure 3: Expected (ex ante) and observed (ex post) utility values for both unconditioned and conditioned optimisers with the given set of MVK utility functions.

be quantified using simple ratios since most utilities resulting from the unconditioned strategy backtest are negative. Overall, then, it is clear that the use of conditioning information strongly and consistently improves both ex ante and ex post investment utilities for the data set used.

3.3 Mean-variance-skewness-kurtosis (MVSK) optimisation

The full four-moment problem, obtained by using nonzero values for a_3 in the utility function form (5), can be designated mean-variance-skewness-kurtosis, or MVSK, even if, as in this case, uncentred moments are used. Although the inclusion of a nonzero third-moment term potentially yields a nonconvex optimisation problem objective function, which makes the numerical solution process more delicate as discussed in section 3.2, formal extension of the problem examined in the previous subsection is straightforward. Indeed, it is sufficient to add a new state to the optimal control formulation whose final value then corresponds to the uncentred third expected moment of returns. Using iterated expectations as before, this is easily obtained as

$$\dot{x}_3(s) = \left((u'(s)\mu(s))^3 + 3u'(s)\mu(s)u'(s)\Sigma_\epsilon^2 u(s) + u'(s)S_\epsilon^3(u(s) \otimes u(s)) \right) p_S(s).$$

We note that this is then the only typically nonconvex term in the resulting utility functions, and that we would not normally expect it to dominate the sum of the remaining three terms in most cases, such that it seems reasonable to expect the problem to solve correctly in the majority of cases. This configuration can be contrasted to that obtained using the PGP approach (see e.g. [14]), which requires the solution of an exclusively third-moment auxiliary problem: that problem in practice proves nontrivial to solve in the unconditioned case (where a simple grid search approach is required but still realistic) and very difficult in the conditioned case, where the number of variables resulting from the problem discretisation is in general too large to allow for a grid search loop with sufficient coverage of the feasible domain. Whether the results obtained are reasonable can best be checked by verifying that they are compatible with those for the MVK problem, which we know to be correct.
Given the greater freedom afforded by considering utility functions with four nonzero coefficients, the eight new functions introduced, given in table 5, meet the stronger conditions of theorem 5 in [27] over the two subintervals $[0, 0.5]$ and $[0, 0.25]$ comprised within the range of observed data returns. The theorem then implies that they verify the standard properties of decreasing absolute risk aversion (DARA), constant or increasing relative risk aversion (CRRA) and decreasing absolute prudence for the intervals concerned, which as in the MVK case allows us to specify functions that model strong preferences with respect to the third and fourth moments while matching standard desirable characteristics over at least part of the domain.

	a_1	a_2	a_3	a_4
Function #1	0.5600	-0.2001	0.1001	-0.0401
Function #2	0.6600	-0.2801	0.1201	-0.0401
Function #3	0.9800	-0.3201	0.1601	-0.0601
Function #4	1.0000	-0.3801	0.2001	-0.0801
Function #5	0.4401	-0.2401	0.2601	-0.2201
Function #6	1.0000	-0.4001	0.4401	-0.4001
Function #7	1.0000	-0.8400	0.7000	-0.4401
Function #8	1.0000	-0.9000	0.8600	-0.6200
Function #9	1.0000	-0.5000	0	0.0000

Table 5: Polynomial coefficients for the MVSK utility functions (indexed from #1 to #9) used.

Utility function #	1	2	3	4	5	6	7	8	9(MV)
Ex ante									
MVSK UNC SR	0.430	0.442	0.424	0.434	0.457	0.441	0.479	0.481	0.457
MVSK UNC rel skewness	-0.263	-0.248	-0.258	-0.242	-0.147	-0.212	-0.058	-0.047	-0.543
MVSK UNC rel kurtosis	15.763	16.409	16.339	15.609	9.823	10.000	10.039	9.270	15.119
MVSK UNC abs kurtosis	0.055	0.043	0.070	0.049	0.010	0.016	0.005	0.004	0.021
MVSK UNC ASR	0.370	0.375	0.364	0.373	0.413	0.399	0.431	0.436	0.378
MVSK UNC utility	0.040	0.043	0.073	0.069	0.025	0.063	0.045	0.043	0.056
MVSK CON SR	0.487	0.497	0.480	0.491	0.515	0.507	0.535	0.532	0.528
MVSK CON rel skewness	1.034	1.155	1.047	1.053	0.763	0.703	0.959	0.885	0.796
MVSK CON rel kurtosis	9.011	9.677	9.263	8.998	6.821	6.660	7.684	7.389	8.583
MVSK CON abs kurtosis	0.145	0.127	0.182	0.132	0.025	0.035	0.017	0.013	0.047
MVSK CON ASR	0.485	0.495	0.477	0.489	0.510	0.501	0.532	0.528	0.512
MVSK CON utility	0.066	0.071	0.121	0.114	0.037	0.096	0.068	0.063	0.087
Rel. utility improvement	64.881%	65.272%	65.829%	65.074%	49.749%	51.727%	51.092%	45.875%	54.888%
Ex post									
MVSK UNC SR	0.144	0.148	0.141	0.145	0.162	0.154	0.173	0.177	0.157
MVSK UNC rel skewness	-1.337	-1.451	-1.226	-1.401	-2.028	-1.894	-2.345	-2.458	-1.818
MVSK UNC rel kurtosis	19.539	21.093	18.603	20.212	24.640	22.860	31.256	33.294	25.104
MVSK UNC abs kurtosis	0.303	0.257	0.346	0.284	0.115	0.160	0.081	0.072	0.177
MVSK UNC ASR	0.137	0.140	0.135	0.138	0.149	0.143	0.155	0.156	0.145
MVSK UNC utility	-0.013	-0.013	-0.020	-0.026	-0.029	-0.066	-0.053	-0.064	0.002
MVSK CON SR	0.183	0.189	0.178	0.185	0.199	0.190	0.208	0.209	0.206
MVSK CON rel skewness	0.044	0.053	0.113	0.010	-0.650	-0.805	-0.759	-1.133	-0.447
MVSK CON rel kurtosis	11.911	12.632	11.406	12.112	14.930	16.454	17.024	20.803	18.010
MVSK CON abs kurtosis	0.405	0.372	0.470	0.379	0.116	0.185	0.089	0.086	0.231
MVSK CON ASR	0.181	0.186	0.176	0.182	0.189	0.180	0.196	0.193	0.196
MVSK CON utility	-0.006	-0.008	-0.008	-0.014	-0.021	-0.058	-0.047	-0.065	0.010

Table 6: Mean (ex ante and ex post) metrics of portfolio returns obtained for unconditioned (UNC) and conditioned (CON) polynomial utility function based mean-variance-skewness-kurtosis (MVSK) optimisation. In particular, SR is the (unmodified) Sharpe ratio and ASR the adjusted Sharpe ratio suggested by [19].

The backtest was then carried out similarly to that executed in the MVK case. While a grid search step was implemented in the case of the unconditioned problem, this was not done for the conditioned problem as mentioned in the above. Results are then given in table 6. Compared to those obtained in the MVK case, they appear consistent with respect to the performance improvements shown by the use of information. Although the degradation in metrics between ex ante and ex post figures does seem to be slightly larger in general than was the case in the preceding subsection, this would not be the result of numerical issues, which would affect both ex ante and ex post figures. Rather, this particular development may plausibly result from the greater estimation errors expected when working with both coskewness and cokurtosis matrices at the same time. The backtest corresponding to utility function #8 does yield an ex post utility figure that is minimally smaller in the conditioned case: as this is not observed for the remaining utility functions and remaining metrics still show improvements through the use of information, it is unlikely that any numerical issues would be involved in that case either.

4 Conclusion

The present paper has formulated conditioned portfolio optimisation strategies involving the fourth moment of returns in two different ways, and reported the results of an empirical investigation, carried out using backtesting over a realistic data set, into their constrained-weight performance. To integrate the fourth moment into the optimisation problem, two approaches were considered: a mean-kurtosis formulation equivalent to mean-variance but using kurtosis as the risk metric, and a utility function formulation which allows for an optimisation strategy that simultaneously considers the first, second and fourth moments of returns. After introducing both conditioned optimisation and optimisation involving higher moments, the optimal control formulation was presented in each case, and the backtesting methodology described. Results were given and discussed, specifically with respect to the classical mean-variance case and the unconditioned equivalents of the conditioned problems introduced. The performance improvements obtained by using signal information were found to be very substantial for both problem variations, both ex ante and ex post and across the entire set of metrics considered, for any levels of ex ante risk beyond the lowest ones.

One interesting point not linked to the conditioning issue per se concerns the choice of problem variant by an individual investor. While the MK problem contains no difficulty of interpretation, it does not allow for variance to be taken into account by the optimiser. This point is fixed by the MVK formulation, but we emphasise that the mapping between utilities and moment preferences is not entirely transparent in that case. Indeed, the polynomial utility functions used involve three moments at a time, and the relevant moments in this case are additionally non-centred moments, whose maximisation or minimisation may or may not directly correspond to optimisation of the more frequently used centred (absolute) moments, or even their normalised (relative) versions. For the investor purely interested in minimising kurtosis, the MK strategy may thus be preferable. However, knowledge of their three-term utility function will, as has been seen, result in systematic outperformance by the conditioned strategy, in a way that coherently takes into account three, rather than just two, returns moments at the same time. In this sense both approaches are seen to be interesting to their respective target sets of investors, and have been experimentally confirmed to be amenable to strong performance improvements by the conditioned optimisation strategies discussed.

Finally, the more general MVSK version of the utility function formulation was backtested. This introduces preferences with respect to the third expected moment of returns and the solutions obtained are more likely to be affected by numerical issues because of the objective function nonconvexities introduced. However, the results obtained do not contradict those obtained for the simpler MVK problem: it has been suggested that, given the nonconvex third-order term does not appear by itself in the given polynomial formulation of the optimisation problem, it is likely that this full MVSK problem can also be solved without issues in most practical cases.

References

[1] Wayne E. Ferson and Andrew F. Siegel. The efficient use of conditioning information in portfolios. *The Journal of Finance*, 56(3):967–982, 2001. ISSN 00221082.

[2] Marc Boissaux and Jang Schiltz. An optimal control approach to portfolio optimisation with conditioning information. Luxembourg School of Finance Working Paper 10-09, Luxembourg School of Finance, 2010.

[3] Marc Boissaux and Jang Schiltz. Practical weight-constrained conditioned portfolio optimisation using risk aversion indicator signals. Luxembourg School of Finance Working Paper 11-12, Luxembourg School of Finance, 2011.

[4] Lars Peter Hansen and Scott F. Richard. The role of conditioning information in deducing testable restrictions implied by dynamic asset pricing models. *Econometrica*, 55(3):587–613, 1987. ISSN 0012-9682.

[5] I-Hsuan Ethan Chiang. Modern Portfolio Management with Conditioning Information. *SSRN eLibrary*, 2008.

[6] Devraj Basu, Chi-Hsiou Hung, Roel C. Oomen, and Alexander Stremme. When to Pick the Losers: Do Sentiment Indicators Improve Dynamic Asset Allocation? *SSRN eLibrary*, 2006.

[7] Devraj Basu, Roel C. Oomen, and Alexander Stremme. Exploiting the Informational Content of the Linkages between Spot and Derivatives Markets? *Warwick Business School Working Paper Series*, 2006.

[8] Jinhui Luo, Philip Saks, and Steve Satchell. Implementing risk appetite in the management of currency portfolios. *Journal of Asset Management*, 9:380–397(18), 2009.

[9] Tsong-Yue Lai. Portfolio selection with skewness: A multiple-objective approach. *Review of Quantitative Finance and Accounting*, 1:293–305, 1991. ISSN 0924-865X. 10.1007/BF02408382.

[10] Gustavo Athayde and Renato Flores. Finding a maximum skewness portfolio. Computing in economics and finance 2001, Society for Computational Economics, April 2001.

[11] James Stacey. Multi-Dimensional Risk and Mean-Kurtosis Portfolio Optimization. *Journal of Financial Management and Analysis, Vol. 21, No. 2, December 2008*, 2008.

[12] Gustavo Athayde and Renato Flores. Incorporating skewness and kurtosis in portfolio optimization: a multidimensional efficient set. In *Advances in Portfolio Construction and Implementation*, pages 243–257. Butterworth-Heinemann, Boston, EUA, 2003.

[13] Eric Jondeau and Michael Rockinger. Optimal portfolio allocation under higher moments. *European Financial Management*, 12(1):29–55, 2006.

[14] Kin Keung Lai, Lean Yu, and Shouyang Wang. Mean-variance-skewness-kurtosis-based portfolio optimization. *Computer and Computational Sciences, International Multi-Symposiums on*, 2:292–297, 2006.

[15] R. Tyrrell Rockafellar and Stanislav Uryasev. Optimization of conditional value-at-risk. *Journal of Risk*, 2:21–41, 2000.

[16] Enrico De Giorgi. A note on portfolio selection under various risk measures. IEW - Working Papers iewwp122, Institute for Empirical Research in Economics - University of Zurich, 2002.

[17] Richard H. Byrd, Mary E. Hribar, and Jorge Nocedal. An interior point algorithm for large-scale nonlinear programming. *SIAM J. Optim.*, 9(4):877–900, 1999. ISSN 1052-6234. Dedicated to John E. Dennis, Jr., on his 60th birthday.

[18] Rama Cont. Empirical properties of asset returns: stylized facts and statistical issues. *Quantitative Finance*, 1: 223–236, 2001.

[19] Jacques Pézier and Anthony White. The relative merits of investable hedge fund indices and of funds of hedge funds in optimal passive portfolios. ICMA Centre Discussion Papers in Finance icma-dp2006-10, Henley Business School, Reading University, 2006.

[20] John T. Betts. *Practical methods for optimal control and estimation using nonlinear programming*, volume 19 of *Advances in Design and Control*. Society for Industrial and Applied Mathematics (SIAM), Philadelphia, PA, first edition, 2001. ISBN 0-89871-488-5.

[21] Ryan J. Davies, Harry M. Kat, and Sa Lu. Fund of hedge funds portfolio selection: A multiple-objective approach. City University Working Paper, City University London, 2003.

[22] John von Neumann and Oskar Morgenstern. *Theory of games and economic behavior*. Princeton University Press, Princeton, N.J., third edition, 1953. ISBN 0-691-00362-9.

[23] Sigrid M. Muller and Mark Machina. Moment preferences and polynomial utility. *Economics Letters*, 23(4): 349–353, 1987.

[24] Maurice G. Kendall. *The Advanced Theory of Statistics. Volume I, 2nd ed.* Charles Griffin & Company Limited, 42 Drury Lane, London, W.C.2, 1945.

[25] Miles S Kimball. Precautionary saving in the small and in the large. *Econometrica*, 58(1):53–73, January 1990.

[26] Miles S. Kimball. Precautionary motives for holding assets. Working Paper 3586, National Bureau of Economic Research, January 1991.

[27] Emmanuel Jurczenko and Bertrand Maillet. Theoretical foundations of asset allocations and pricing models with higher-order moments. In *Multi-moment Asset Allocation and Pricing Models*, pages 1–36. John Wiley and Sons, 2006.