CARISMA

A multi-disciplinary research centre focussed on understanding, modelling, quantification, management and control of RISK

- School of Information Systems, Computing and Mathematics
- Brunel Business School
- School of Engineering
- School of Social Sciences

TECHNICAL REPORT

CTR/78-08

HMM Based Scenario Generation for an Investment Optimisation Problem

Christine Gregory, Diana Roman, Gautam Mitra,

CARISMA

Centre for the Analysis of Risk and Optimisation Modelling Applications

Brunel UNIVERSITY WEST LONDON

Noname manuscript No.
(will be inserted by the editor)

HMM based scenario generation for an investment optimisation problem

Christina Erlwein · Gautam Mitra · Diana Roman

the date of receipt and acceptance should be inserted later

Abstract The Geometric Brownian motion (GBM) is a standard method for modeling financial time series. An important criticism of this method is that the parameters of the GBM are assumed to be constants; due to this fact, GBM has been considered unable to properly capture important features, like extreme behaviour or volatility clustering. We propose an approach by which, the parameters of the GBM follow a regime switching model, more precisely a hidden Markov model (HMM). Thus, financial time series are modeled via a hidden Markov model (HMM) with a GBM in each state. Using this approach, we generate scenarios for a financial portfolio optimisation problem in which the portfolio CVaR is minimised. Numerical results are presented.

Keywords scenario generation, hidden Markov model, Geometric Brownian motion, asset allocation, optimal parameter estimation

1 Introduction

Stochastic programming is an important tool for portfolio optimisation. Since the future returns of most assets are uncertain, the outcome of an investment is uncertain. Usually, the decision maker tries to achieve an optimal combination of risk and expected return. This is done by solving stochastic programs: optimisation problems in which some of the parameters (the future returns) are not certain, but described by distributions. In order for the stochastic programs to be numerically solved, the distributions involved are approximated by discrete distributions with a finite number of outcomes (scenarios). An overview of financial optimisation problems can be found in (Mulvey (2001)).

Thus, a crucial issue for financial optimisation problems (with portfolio selection as a particular case) is the scenario generation for the future realisations of the time series involved. This is closely related to a proper modeling of the financial time series.

Among the most popular models for financial time series are the Geometric Brownian motion (GBM) (see for example Ross (2002)) and ARCH-GARCH models (Engle (1982), Bollerslev (1986)). With GBM, stock prices are approximated by continuous time stochastic processes. ARCH-GARCH models were designed for modelling time-dependent variance. However, both methods are accepted only with reservation; one of the main criticism is that they underestimate the lower tails of distributions, thus the severity of unfavourable

Christina Erlwein
Fraunhofer Institute, ITWM, Fraunhofer Platz 1, 67663 Kaiserslautern, Germany
Tel: +49 631 31600-4266 Fax: +49 631 316005352 E-mail: christina.erlwein@itwm.fraunhofer.de

Gautam Mitra · Diana Roman
CARISMA, Brunel University, Uxbridge, UB8 3PH, Great Britain *and*
OptiRisk Systems Ltd, Uxbridge, UB9 4DA, Great Britain
E-mail: gautam.mitra@brunel.ac.uk E-mail: diana.roman@brunel.ac.uk

events. Both models may be used for scenario generation by sampling from the assumed distributions. A more recent method of scenario generation for general stochastic programs is the "moment-matching" approach (Hoyland et al (2003)) in which it is assumed that the moments of the "true" distribution are known; a discrete distributions with the same moments is then constructed. An overview of scenario generation for general stochastic programs is given in (Kaut and Stein (2003)).

Over the last years, hidden Markov models (HMM) became increasingly popular for modelling financial time series. One main idea behind these models is that some information on the market is hidden in noisy observations, e.g. price movements. The underlying hidden information can symbolise different stages of a business cycle, like expansion, peak, recession, trough and recovery, which influence the price movements. The market may switch from time to time from a state to another; each state is characterised by different parameters of the model for the observable stochastic process. In this case, one single model is valid only for short periods; regime switching models offer a better way of modelling.

A hidden Markov model (HMM) is a particular regime-switching model, in which there are two stochastic processes involved: apart from the one of interest, that is observable (e.g. asset prices), there is an underlying stochastic process describing the system's state over time, that is not observable, i.e. "hidden". HMMs were applied to a variety of fields, in particular to speech recognition (Rabiner (1989)); a review of HMMs and their applications is given in (Rabiner (2002)).

A review of the applications of HMM for financial time series modelling and forecasting is given in (Zhang (2004)). In (Hamilton and Susmel (1994)) weekly time series were modelled using regime switching with ARCH-type models within regimes. A simpler regime switching model for monthly returns with log-normal distributions is proposed in (Hardy (2001)). (Messina and Toscani (2007)) use autoregressive HMM for modeling univariate financial time series and for generating scenarios to describe their possible future evolution. In order to validate the HMM based scenario generation method, they show that the resulting Monte Carlo-sampled distributions can replicate with good approximation the empirical distribution of the observed data. In (Roman et al (2008)) multi-variate financial time series are modelled by using HMMs and mixtures of multi-variate normal distributions; based on this approach, scenarios are generated for a mean-downside risk optimisation problem. It is shown that the scenario generator based on this method has a reasonable stability and can capture extreme movements of asset prices. However, the resulting model is computationally demanding and can only handle a limited number of time series.

In this paper, we use HMM and Geometric Brownian motion for modelling financial time series. The motivation behind is that, although GBM is a traditional and popular method, it has recognised limitations. It fails to capture some important observed aspects, like volatility clustering and extreme behaviour. The limitations of GBM are mainly due to the fact that the parameters of the GBM (the "drift" and the volatility) are assumed to be constant. In our approach we aim to address this shortcoming by allowing the GBM parameters to switch between different states.
These parameters are estimated using a parameter estimation filtering approach (Elliott et al (1995)). Information of the Markov chain is filtered out of the observation process and optimal parameter estimates are derived with a reference probability measure technique. Once optimal parameters are estimated using historical information of the prices, scenario paths can be generated according to the parameter estimates and the filtered transition probability of the underlying Markov chain.

In our numerical experiments, we generate scenarios for the future returns of the FTSE100 index and of gold spots. The quality of the scenario generator is tested in a portfolio optimisation problem, in which the mean-CVaR approach is used, namely, the CVaR of the portfolio is minimised while a minimal portfolio expected return is imposed. We conclude that the scenario generation method proposed is reasonably stable.

The rest of the paper is organised as follows. In Section 2 the proposed approach for modelling financial time series is described. The basics of GBM and of HMM are presented. Once the parameter estimation and the future Markov chain estimation are made, the model can be used as a scenario generator. In Section 3 we describe the portfolio selection problem and the mean-CVaR optimisation model. Section 4 presents the numerical results. The scenarios generated with the approach presented in Section 2 are used as input in a mean-CVaR model; the stability of this scenario generator is investigated. Conclusions are presented in Section 5.

2 Modelling financial time series using HMM and GBM

2.1 The Geometric Brownian motion

The Geometric Brownian motion (GBM) is a continuous time stochastic process that is widely used as a reasonable approximation of stock price dynamics. Although stock prices are stochastic processes in discrete time, this approximation is generally accepted by practitioners for short - medium time periods. Formally, a stochastic process S_t follows a GBM if it satisfies the following stochastic differential equation:

$$dS_t = \mu S_t dt + \sigma S_t dW_t. \tag{1}$$

where W_t is a Wiener process, and μ (the "drift"), σ (the volatility) are parameters of the model, assumed to be constant.

The equation (1) has the following analytic solution, where S_0 is the initial value (at time $t = 0$):

$$S_t = S_0 exp\{(\mu - \frac{\sigma^2}{2})t + \sigma W_t\}. \tag{2}$$

One shortcoming of the model is that, in this form, it cannot account for observable phenomena in financial markets, like extreme behaviour and volatility clustering.
It is known that a Wiener process W_t can be simulated as $\sqrt{t}z_t$, where z_t is N(0,1) distributed. Thus, the GBM can be used for generating possible scenarios for S_t by sampling from the standard normal distribution and using (2).

2.2 Hidden Markov models with GBM as observation processes

With the HMM approach, there are two stochastic processes involved (both in discrete time): in addition to the one of interest, which is observable (e.g. stock prices), there is another one describing the "state of the system", which is not directly observable. At each point in time, the system is in one of possible n states and it may switch from the current state to another (or it may stay into the same state) according to some transition probabilities. Only an observation from the first stochastic process can be made. The underlying, hidden stochastic process describing the system's state is a Markov chain, i.e. the system's state at one point in time depends only on the system's state at the previous point in time and not on the entire history. Another assumption is that the transition probabilities are not time dependent (they only depend on the two involved states).
Each state generates observations for the stochastic process of interest according to its own probability density. In our approach, we assume that the observations of the stock prices are generated according to discretised Geometric Brownian motions; each state in the Markov chain has different parameters of the GBM. This approach may address the shortcomings of the GBM, described at the previous section.

Thus, the log-return process $y_t = ln\frac{S_t}{S_{t-1}}$ has the following dynamic in discrete time:

$$y_{t+1} = f(\mathbf{x}_t) + \sigma(\mathbf{x}_t) z_{t+1}. \tag{3}$$

where $f(\mathbf{x}_t) = (\mu(\mathbf{x}_t) - \frac{\sigma^2(\mathbf{x}_t)}{2})$. The parameters f and σ are governed by the Markov chain $\mathbf{x}$ in discrete time and are therefore able to switch between different regimes. The $z_t's$ are a sequence of independent, identically distributed (IID) standard normal random variables independent of $\mathbf{x}$.

We denote the scalar product with $< \cdot, \cdot >$, and with n the number of states of the Markov chain. The parameters associated with GBMs are of the form $\mathbf{f} = (f_1, f_2, \ldots, f_n)^\top$ and $\boldsymbol{\sigma} = (\sigma_1, \sigma_2, \ldots, \sigma_n)^\top$ such that $f(\mathbf{x}_t) = \langle \mathbf{f}, \mathbf{x}_t \rangle$ and $\sigma(\mathbf{x}_t) = \langle \boldsymbol{\sigma}, \mathbf{x}_k \rangle$ where σ_i's are all positive for every $1 \leq i \leq n$. The observation at time $t+1$ depends on the state of $\mathbf{x}$ at time t. This is a one-step delay model and is reasonable as y may not react to $\mathbf{x}$ immediately.

Under the real world probability measure P, the Markov chain $\mathbf{x}$ has the dynamics

$$\mathbf{x}_{t+1} = \boldsymbol{\Pi}\mathbf{x}_t + \mathbf{v}_{t+1} \tag{4}$$

where $\mathbf{v}_{t+1}$ is a martingale increment and $\Pi = (\pi_{ji})_{,j,i=1\dots n}$ is the transition probability matrix with $\pi_{ji} = P(\mathbf{x}_{k+1} = \mathbf{e}_j | \mathbf{x}_k = \mathbf{e}_i)$, i.e. the probability of transiting from state j to state i (as stated before, the assumption is that these transition probabilities are independent of time).

Thus, the parameters to be estimated are: the number of Markov states n, the transition probabilities $\pi_{ji},_{j,i=1\dots n}$ and the parameters of the discretised GBM's $\mathbf{f}$ and $\boldsymbol{\sigma}$.

While optimal values can be derived recursively for $\mathbf{f}$, $\boldsymbol{\sigma}$ and Π, the number of states is supposed to be determined in advance; they are estimated separately from the rest of the parameters. In practice, the state dimension of the hidden Markov chain may be dictated by the actual application or may be determined in an empirical manner by visual inspection of plots of the data (Messina and Toscani (2007), Geyer and Ziemba (2007)). In (Roman et al (2008)) the AIC criterion is used (Akaike (1974)) in order to determine the number of Markov states.

Regarding the rest of the parameters, filters are developed for the Markov chain and related quantities and recursive optimal parameter estimation are derived. Denote by $\mathcal{Y}_k$ the complete filtration generated by the sequence of $y_1, y_2, \dots, y_k$.
The jump process of the Markov chain, namely the number of times the Markov chain jumps from state i to state j up to a given time k is denoted by $J_k^{(ji)}$.

$$J_k^{(ji)} = \sum_{l=1}^{k} \langle \mathbf{x}_{l-1}, \mathbf{e}_i \rangle \langle \mathbf{x}_l, \mathbf{e}_j \rangle \tag{5}$$

The occupation time process which indicates how long the Markov chain stays in state i up to time k is denoted by $O_k^{(i)}$.

$$O_k^{(i)} = \sum_{l=1}^{k} \langle \mathbf{x}_{l-1}, \mathbf{e}_i \rangle \tag{6}$$

Furthermore we need an auxiliary process $T_k^{(i)}$ for filtering the Markov chain.

$$T_k^{(i)}(g) = \sum_{l=1}^{k} \langle \mathbf{x}_{l-1}, \mathbf{e}_i \rangle \, g(y_l) \tag{7}$$

where g is a function that will take the form $g(y) = y$ or $g(y) = y^2$. (8)

These processes are filtered out of our observation process y under a reference probability measure $\widetilde{P}$ following an HMM filtering method (Elliott et al (1995)). A reverse measure change back to the real world measure P then gives us the filter equation for the processes of the Markov chain under the physical measure. To construct the real world measure P from $\widetilde{P}$, we define the processes λ_l and Λ_l by

$$\lambda_l = \frac{\phi\left(\sigma\left(\mathbf{x}_{l-1}\right)^{-1}\left(y_l - f(\mathbf{x}_{l-1})\right)\right)}{\sigma(\mathbf{x}_{l-1})\phi(y_l)}$$

$$\Lambda_k = \prod_{l=1}^{k} \lambda_l, \quad k \geq 1, \quad \Lambda_0 = 1 \tag{9}$$

where $\phi(z)$ is the probability density function of the standard normal distribution. We refer to Λ_k as the Radon-Nikodŷm derivative of P with respect to $\widetilde{P}$, $\left.\frac{dP}{d\widetilde{P}}\right|_{\mathcal{H}_k} = \Lambda_k$.
Write $\mathbf{D}$ for the diagonal matrix whose ith entry on the diagonal is

$$\frac{\phi\left(\frac{y_{k+1} - f_i}{\sigma_i}\right)}{\sigma_i \phi(y_{k+1})}. \tag{10}$$

First we would like to estimate $\mathbf{x}$, given the observations under P, the real world probability. Let $\widehat{\mathbf{x}}^i = P(\mathbf{x_k} = \mathbf{e}_i|\mathcal{Y}_k) = E\left[\langle \mathbf{x_k}, \mathbf{e}_i\rangle | \mathcal{Y}_k\right]$ be the conditional distribution of $\mathbf{x}_k$ given $\mathcal{Y}_k$ under P. From Bayes' theorem, $\widetilde{E}[\Lambda_k \mathbf{x}_k|\mathcal{Y}_k] = \widetilde{E}[\Lambda_k|\mathcal{Y}_k]E[\mathbf{x}_k|\mathcal{Y}_k]$ We define $\boldsymbol{\xi}_k := \widetilde{E}[\Lambda_k \mathbf{x}_k|\mathcal{Y}_k]$. Noting that $\sum_{i=1}^n \langle \mathbf{x}_k, \mathbf{e}_i\rangle = 1$, we have

$$\sum_{i=1}^n \widetilde{E}\left[\langle \Lambda_k \mathbf{x}_k, \mathbf{e}_i\rangle\right] = \widetilde{E}\left[\Lambda_k \sum_{i=1}^n \langle \mathbf{x}_k, \mathbf{e}_i\rangle \middle| \mathcal{Y}_k\right] = \widetilde{E}\left[\Lambda_k|\mathcal{F}_k\right] = \sum_{i=1}^n \langle \boldsymbol{\xi}_k, \mathbf{e}_i\rangle .$$

Again noting that $\widehat{\mathbf{x}}_k = E[\mathbf{x}_k|\mathcal{Y}_k]$, we get an explicit form for the conditional distribution, $\widehat{\mathbf{x}}_k = \frac{\boldsymbol{\xi}_k}{\sum_{i=1}^n \langle \boldsymbol{\xi}_k, \mathbf{e}_i\rangle}$.

The filters for the processes are also derived by applying the change of measure technique, this leads to

$$\begin{aligned}
\gamma(J^{(ji)}\mathbf{x})_l &= \boldsymbol{\Pi}\mathbf{D}(y_l)\gamma(J^{(ji)}\mathbf{x})_{l-1} \\
&\quad + \left\langle \boldsymbol{\xi}_{l-1}, \mathbf{e}_i\right\rangle \frac{\phi\left(\sigma_i^{-1}(y_l - f_i)\right)}{\sigma_i \phi(y_l)} \pi_{ji}\mathbf{e}_j, && (11)\\
\gamma(O^{(i)}\mathbf{x})_l &= \boldsymbol{\Pi}\mathbf{D}(y_l)\gamma(O^{(i)}\mathbf{x})_{l-1} \\
&\quad + \left\langle \boldsymbol{\xi}_{l-1}, \mathbf{e}_i\right\rangle \frac{\phi\left(\sigma_i^{-1}(y_l - f_i)\right)}{\sigma_i \phi(y_l)} \boldsymbol{\Pi}\mathbf{e}_i && (12)\\
\gamma(T^{(i)}(g)\mathbf{x})_l &= \boldsymbol{\Pi}\mathbf{D}(y_l)\gamma(T^{(i)}(g)\mathbf{x})_{l-1} \\
&\quad + \left\langle \boldsymbol{\xi}_{l-1}, \mathbf{e}_i\right\rangle \frac{\phi(\sigma_i^{-1}(y_l - f_i))}{\sigma_i \phi(y_l)} g(y_l)\boldsymbol{\Pi}\mathbf{e}_i. && (13)
\end{aligned}$$

with $\boldsymbol{\xi}_{k+1} = \boldsymbol{\Pi}\mathbf{D}\boldsymbol{\xi}_k$. More details and proofs for these filter derivations can be found in (Mamon et al (2008)).

The derived filters for the processes are used for the recursive parameter estimation formulas. The recursive formulas for the model parameters are calculated with the Expectation-Maximisation (EM)- algorithm. The maximum likelihood estimation (MLE) for the parameters within the EM-algorithm makes use of the derived adaptive filters of the observation process. In the calculation of the MLE, the filters substitute terms involving the observation process. The following recursive optimal parameter estimates for the transition probabilities π_{ji}, the mean f_i and the variance σ_i of the observation process are derived (see (Mamon et al (2008)) for a proof).

$$\widehat{\pi}_{ji} = \frac{\widehat{J}_k^{(ji)}}{\widehat{O}_k^{(i)}} = \frac{\gamma(J^{(ji)})_k}{\gamma(O^{(i)})_k}, \tag{14}$$

$$\widehat{f}_i = \frac{\widehat{T}_k^{(i)}}{\widehat{O}_k^{(i)}} = \frac{\gamma(T^{(i)}(y))_k}{\gamma(O^{(i)})_k} \tag{15}$$

and

$$\widehat{\sigma}_i = \sqrt{\frac{\widehat{T}_k^{(i)}(y^2) - 2\widehat{f}_i \widehat{T}_k^{(i)}(y) + \widehat{f}_i^2 \widehat{O}_k^{(i)}}{\widehat{O}_k^{(i)}}}. \tag{16}$$

Once the parameters are estimated, this method can be used as a scenario generator for the future prices (returns), after the Markov chain for the next time period is estimated. As derived above, the conditional distribution of the Markov chain for the next time step is given by $E[\mathbf{x}_{k+1}|\mathcal{Y}_k] = \boldsymbol{\Pi}\hat{\mathbf{x}_k}$ with $\hat{\mathbf{x}_k} = E[\mathbf{x}_k|\mathcal{Y}_k] = \frac{\boldsymbol{\xi}_k}{\sum_{i=1}^n \langle \boldsymbol{\xi}_k, \mathbf{e}_i\rangle}$. Scenario generation is then made by sampling from the corresponding distributions, according to (3).

In our numerical experiments we consider two different frameworks for the scenario generations. In the first setting both time series are assumed to be independent, the optimal parameters for each process are derived in separate algorithms. The second setting assumes a dependency between the time series. Both observation processes are governed by the same Markov chain and are therefore estimated as vector observations. A thorough discussion on the filter derivation and optimal parameter estimation for vector observations can be found in (Elliott et al (1995)).

3 The mean-CVaR model for portfolio optimisation

A crucial issue in financial optimisation is the choice criterion used. The result of a financial decision (say a portfolio return) is uncertain and regarded as a random variable (with a discrete distribution, using scenario generation). Thus a natural question is what criterion to use for making choices among random returns. Various models for optimal portfolio selection have been proposed, starting with the popular Mean-Variance approach (Markowitz (1952)), in which decisions (i.e., portfolio weights) are taken such that the portfolio's expected return is maximised while a risk measure (the variance) is minimised. More recently portfolio optimisation problems include more sophisticated risk measures. A popular risk measure in industry at the moment is Value-at-Risk (VaR), which indicates the maximum amount to be lost at a particular confidence level. One drawback of this measure in the application to stochastic optimisation problems is that it is non-smooth and non-convex and can therefore lead to multiple local extrema (Artzner et al (1999)). On the other hand, the Conditional Value-at-Risk (CVaR), also known as Mean Excess Loss or Mean shortfall, which is approximately equal to the conditional expectation of losses beyond VaR, is a risk measure with good theoretical and computational properties (Pflug (2000)). CVaR models were developed for portfolio optimisation, where in one framework (Rockafellar and Uryasev (2000)) risk is minimised with a given minimum level of portfolio return, while in other frameworks (Krokhmal et al (2002)) the return is maximised with a given maximum level of risk, or a weighted combination of risk and return is optimised. All formulations lead to the same efficient frontier (Krokhmal et al (2002)).

Let R_x be a random variable representing the return of a portfolio x over a given holding period and $A\% = \alpha \in (0,1)$ a percentage which represents a sample of "worst cases" for the outcomes of R_x (usually, $\alpha = 0.01 = 1\%$ or $\alpha = 0.05 = 5\%$).

The definition of CVaR at the specified level α is the mathematical transcription of the concept "average of losses in the worst $A\%$ of cases" (Acerbi and Tasche (2002)), where the "loss" associated with R_x is usually described by the random variable $-R_x$, meaning a negative return is a loss, a positive return is a "gain" (however, the definition of CVaR allows for a general loss function.)
For a detailed definition of CVaR, see (Rockafellar and Uryasev (2002)). CVaR is approximately equal to the average of losses greater than or equal to VaR (at the same confidence level α); in some cases, the equality is exact.

An important result, proved in (Rockafellar and Uryasev (2002)), is that the CVaR of a portfolio x can be calculated by solving a convex optimisation problem. Moreover and very importantly, CVaR can be optimised over the set of feasible decision vectors (feasible portfolios) and this is also a convex optimisation problem. In (Rockafellar and Uryasev (2002)), an auxiliary function is used, $F : X \times \mathbb{R} \to \mathbb{R}$,

$$F_\alpha(x, v) = \frac{1}{\alpha} E[-R_x + v]^+ - v,$$

where $[u]^+ = u$ if $u \geq 0$ and 0 otherwise. They proved that minimising CVaR over X can be done by minimising F_α over $X \times \mathbb{R}$.

When the random returns under consideration are represented as discrete random variables (via scenario generation), the CVaR optimisation problem is a LP.

The algebraic formulation of the mean-CVaR model is given below:

$$\min \quad v + \frac{1}{\alpha S} \sum_{i=1}^{S} y_i \tag{M-CVaR}$$

Subject to:

$$\sum_{j=1}^{N} -r_{ij} x_j - v \leq y_i, \quad \forall i \in \{1 \ldots S\}$$

$$y_i \geq 0, \quad \forall i \in \{1 \ldots S\}$$

$$\sum_{j=1}^{N} \mu_j x_j \geq d; \quad x \in X.$$

where the parameters of the model are:

- $S =$ the number of scenarios
- $N =$ the number of assets
- $r_{ij} =$ return of asset j under scenario i, $i = 1 \ldots S$, $j = 1 \ldots N$
- $\mu_j =$ the expected return of asset j, $j = 1 \ldots N$ ($\mu_j = \frac{1}{S}\sum_{i=1}^{S} r_{ij}$))
- $d =$ the desired expected return of the portfolio

The decision variables of the model are:

- $x_j =$ the fraction of portfolio wealth invested in asset j, $j = 1 \ldots N$
- $v =$ the negative of an α-quantile of the portfolio return distribution (it is an approximation for the α-VaR of the portfolio)
- $y_i =$ the magnitude of the negative deviations of the portfolio return from the α-quantile, for every scenario $i \in \{1 \ldots S\}$ (they are 0 if the portfolio return is higher than the α-quantile)

4 Numerical experiments: A case study for an investment problem

We consider the following portfolio optimisation problem. An investor faces the problem of finding the optimal ratio between investing in gold spots, FTSE100 stocks or putting his money into a bank account. We consider a one time-step problem: decisions have to be taken now in order to give the best return at the next time period.
An investment in gold is traditionally seen as a long-term investment that can play a role in hedging against inflation and political or economic problems. Like other commodities, gold is traded on spot and future markets, but it has specific characteristics which are not common within commodity markets. Because of its role as a global currency, prices on spot markets are global. Gold supply and demand does not depend on seasonality, it has a low risk of supply interruption and low storage and insurance costs. Furthermore gold has no risk of spoilage and the consumption level relative to inventory is low. Unlike price processes for other commodities, a convenience yield is not included in the price process due to the distinctive gold features.
We generate scenarios for the future prices of gold and FTSE100 index according to the scenario generation method proposed in section 2. These scenarios are further used in a mean-CVaR optimisation model (section 3) in order to take the required decisions.
The dataset used for generating the scenarios are historical prices of gold spots and FTSE 100 index, monitored daily over a 1-year period.
We considered a number of $n = 3$ Markov states. This is consistent with Roman et al (2008), Messina and Toscani (2007), Geyer and Ziemba (2007), who, either by inspecting data or by using statistical techniques, estimated the number of Markov states to 3. It is in fact more or less agreed that a number of 3 Markov states usually gives a good representation, without overfitting the model.

We generate scenarios under two different frameworks. First, we consider that the two time series of interest are independent and thus governed by different Markov chains. Secondly, we consider them governed by the same Markov chain; observations are generated as vectors.
The algorithm for estimating the rest of the parameters runs 24 times on the data sets. For each data set the parameters are updated when new information arrives after batches of 10 data points. Scenarios are generated starting from the last data point.

Figures 1 and 2 depict the optimal parameter estimates for the gold spot prices and for FTSE 100, respectively (estimated independently). Scenarios for the next time period are generated, following the parameter estimation.

Under the assumption of correlated time series, we consider that both time series are governed by the same hidden Markov chain, which filtered out from both processes. Under this framework we estimate the optimal parameters of both observation processes simultaneously and we generate vector observations.

Figure 3 shows the parameter estimation of the observation processes for the vector observation case. Possible scenarios for both time series are then generated and depicted in figure 4.

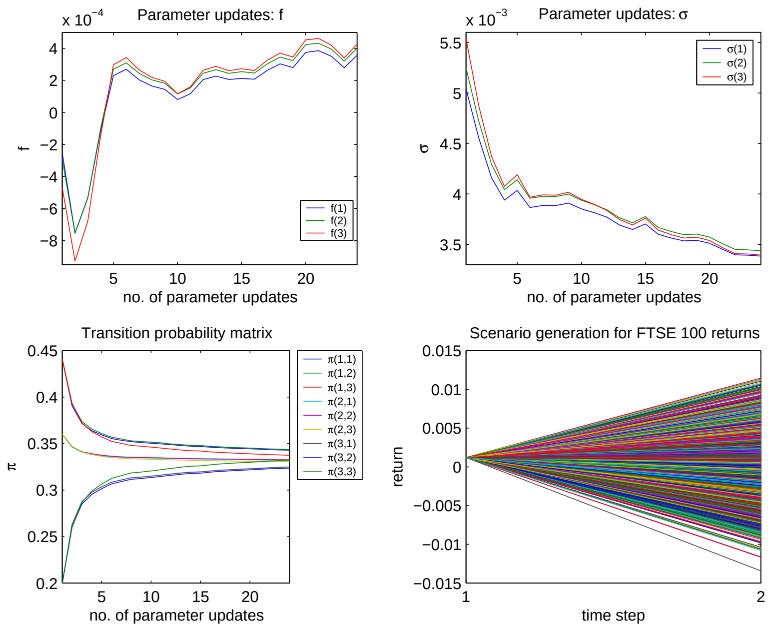

Fig. 1 Parameter estimation and scenario generation for FTSE 100. Different lines correspond to the different Markov states

4.1 Stability of the scenario generators

Stability is one of the most important requirements from a scenario generator. Since any scenario generator has an element of randomness, the final outcome depends on the values drawn from the corresponding distributions (in our case, the standard normal distribution). Thus, every time a scenario generator is run, different sets of scenarios are obtained. Stability guarantees that the optimal solution of the optimisation problem of interest does not depend (but, possibly, only to a small extent) on the specific scenario set chosen.

For each framework (independent / correlated time series) we use the same methodology. We generate 10 sets, each with 5000 scenarios for the future returns of FTSE100 and gold. The scenarios obtained are used as input in the mean-CVaR optimisation problem (the scenarios are considered to have equal probability). We consider CVaR at confidence level $\alpha = 1\%$. The daily return on the bank account is assumed to be 0.0001. The minimal required level of daily portfolio return is set to $d = 0.0005$. The optimisation problems are formulated in AMPL (Fourer et al (1989)) and solved with the FortMP solver (Ellison et al (1999)).

The optimal solutions and objectives, for the case of independent observations, are displayed in Table 1. All values of the objective function lie between 0.002844 and 0.003325 with a mean of 0.0031 and variance of $2.7070e-008$.

Scenario set		1	2	3	4	5	6	7	8	9	10
Objective function		0.003325	0.00321	0.002955	0.003199	0.003222	0.003035	0.003004	0.002844	0.002927	0.002898
Percentages invested in:	FTSE	9.77	17.64	9.14	16.88	15.27	12.03	15.84	10.38	13.63	8.81
	Gold	24.37	20.71	22.26	21.8	22.76	23.08	21.08	21.44	21.8	22.38
	Bond	65.96	61.65	68.6	61.32	61.97	64.89	63.08	68.18	64.57	68.81

Table 1 Stability analysis for 10 scenario sets with independent observations

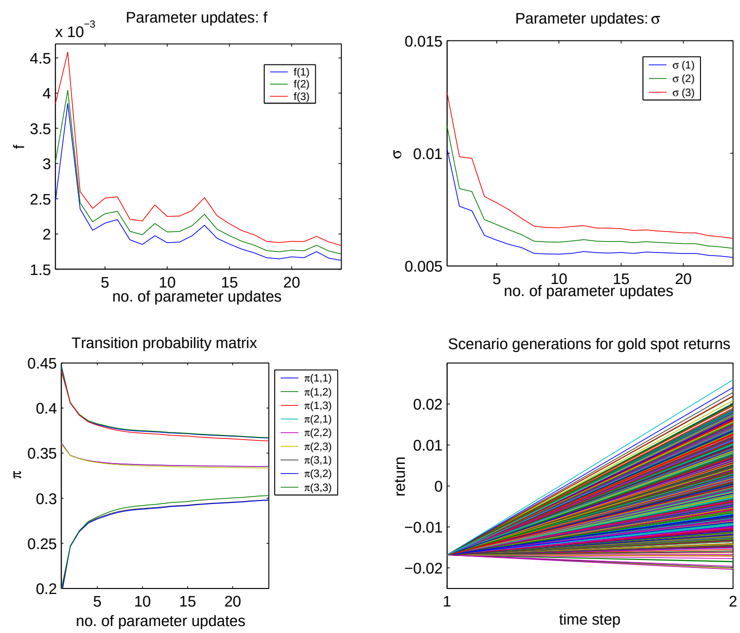

Fig. 2 Parameter estimation and scenario generation for gold spots

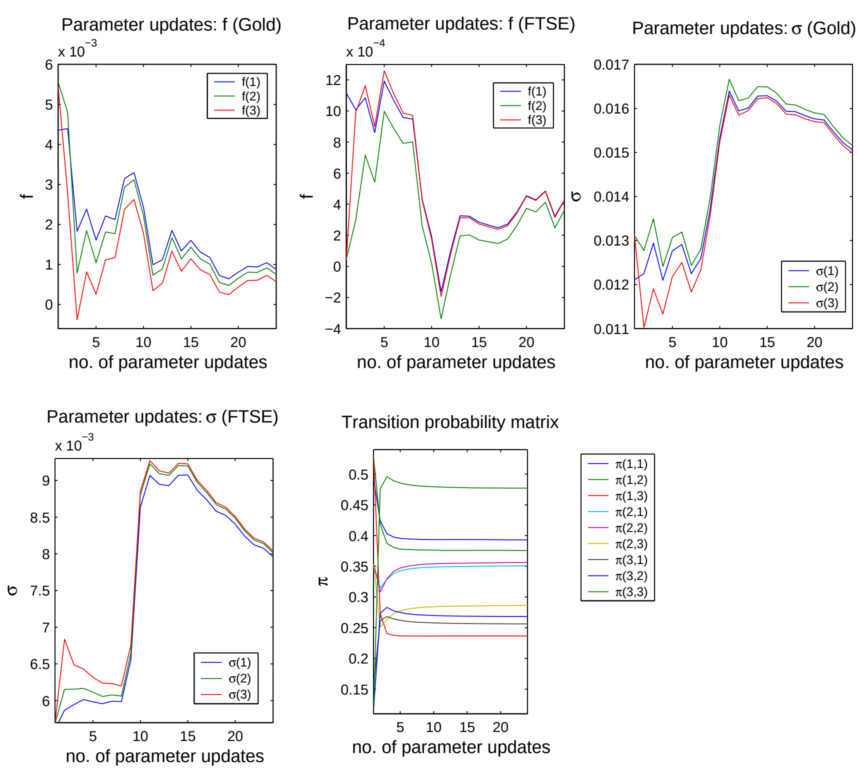

Fig. 3 Simultaneous parameter estimation for observation process of gold spots and FTSE 100 returns

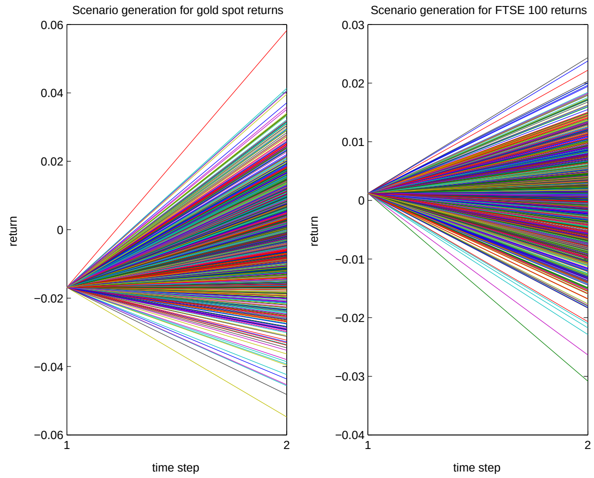

Fig. 4 Scenario generation for returns on gold spots and FTSE 100

The optimal solutions and objectives, for the case of vector observations, are displayed in Table 2. All values of the objective function lie between 0.015058 and 0.027259 with a mean of 0.0193 and variance of $1.4093e - 005$.

Scenario set		1	2	3	4	5	6	7	8	9	10
Objective function		0.018185	0.015058	0.015768	0.021392	0.027259	0.016276	0.023035	0.019834	0.017423	0.0185903
Percentages invested in:	FTSE	42.47	46.13	40.58	99.47	30.78	58.77	43.06	63.75	47.06	63.1
	Gold	39.15	28.83	35.16	0.53	69.22	27.22	56.94	36.25	37.52	36.9
	Bond	18.38	25.04	24.26	0	0	14.01	0	0	15.42	0

Table 2 Stability analysis for 10 scenario sets with vector observations

It is evident from the tables above that the scenario generator is much more stable under the framework of independent observations. The stability under this framework is good, with respect to both the optimal values and the optimum. The similarity of optimal solutions guarantees a good out-of-sample stability of this scenario generator (i.e. stability with respect the "true", unknown distribution of the random returns), which is the most important aspect.
When generating vector observations, there is somewhat high variability of both optimums and optimal solutions.

Economically this can be explained by the largely uncorrelated price of gold to any other stock. Due to its characteristics as a stable long-term investment and a global currency, gold prices are highly likely to be independent of stock prices or indices. This is supported by stable optimal solutions calculated with independent generated scenarios.

4.2 The investment problem with weight constraints

In this section we modify our original investment problem and add constraints on the weights (fractions of capital) invested in the different assets. All weights now have an upper bound of 50%, i.e. no more than half of the budget is allowed to be invested in either gold spots, FTSE 100 stocks or bonds. Taking into account the results from the previous section, we only use the scenario generator that assumes independent observations. We generate five sets containing 5,000 scenarios and five sets containing 20,000 scenarios; we

solve the constrained mean-CVaR optimisation problem with this input.
As it may be seen from Table 3, the stability of the solutions and objective is very good. The average objective function over all generated scenario sets is 0.003538409 with a very low variance of $3.77072e-09$. As already specified in the previous section, the similarity of the optimal solutions implies a good out-of-

Scenario set		5,000 generated scenarios					20,000 generated scenarios				
		1	2	3	4	5	1	2	3	4	5
Objective function		0.003676	0.003541	0.003575	0.003459	0.00354	0.003522	0.003569	0.003517	0.003465	0.00352
Percentages invested in:	FTSE	29.38	32.67	32.6	31.57	30	31.84	30.55	31.87	32.04	30.05
	Gold	20.62	17.33	17.4	18.43	20	18.16	19.45	18.13	17.96	19.95
	Bond	50	50	50	50	50	50	50	50	50	50

Table 3 Investment problem with weight constraints, results for 10 scenario sets

sample stability.

5 Conclusions

In this paper, we modeled financial time series using a hidden Markov model approach. This enables the parameters of the observation process (which we assumed to be a discretised Geometric Brownian motion) to switch between different economic regimes (states) that are not directly observable. Thus our method, while not departing fundamentally from traditional approaches, tries to address a major shortcoming (the assumption that the parameters of the GBM are constant over time).
The parameters of the model are estimated using a filtering approach. Once the parameter estimation and the estimation of the Markov chain for the next time period are done, the model can be used for generating scenarios of future realisations for financial time series.
We tested the quality of the scenario generator in a financial decision making problem: a single period portfolio optimisation problem, in which the basket of assets is formed by FTSE 100, gold spots and a bank account. Scenarios are generated for the future returns of FTSE 100 and of gold spots under two frameworks: once, assuming that these time series are independent and follow different underlying Markov chains; secondly, assuming that the time series are correlated and governed by the same Markov chain. The scenarios generated were used as an input in a mean-CVaR optimisation problem. We tested the stability of the scenario generator by producing several different sets of scenario of the same size and solving the corresponding mean-CVaR model with each of these sets. We concluded that the scenario generator based on independent observations gives a good stability, both in-sample and, more importantly, out-of-sample. The focus of future research is to extend and test the scenario generator for multiple time step optimisation problems.

Acknowledgements This work was carried out while Christina Erlwein was a NET ACE sponsored Intern at OptiRisk Systems; and Diana Roman was a KTP associate of Brunel Unversity assigned to OptiRisk systems. The support of these organisations are gratefully acknowledged.

References

Acerbi C, Tasche D (2002) On the coherence of Expected Shortfall. Journal of Banking and Finance 26:1487–1503

Akaike H (1974) A new look at the statistical model identification. IEEE Transactions on Automatic Control AC 19:716–723

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

Bollerslev T (1986) Generalized Autoregressive Conditional Heteroskedasticity. Journal of Econometrics 31(2):307–327

Elliott R, Aggoun L, Moore J (1995) Hidden Markov models: estimation and control. Springer

Ellison E, Hajian M, Levkovitz R, Maros I, Mitra G (1999) A Fortran based mathematical programming system FortMP. Brunel University, Uxbridge, UK, NAG Ltd, Oxford UK

Engle RF (1982) Autoregressive Conditional Heteroskedasticity with Estimates of the Variance of U.K. Inflation. Econometrica 50(2):987–1007

Fourer R, Gay DM, Kernighan B (1989) AMPL: A Mathematical Programming Language. AT&T Bell Laboratories

Geyer A, Ziemba W (2007) The innovest austrian pension fund financial planning model innoalm. To appear in Operations Research, available at: http://edoc.hu-berlin.de/docviews/abstract.php?lang=ger&id=26731

Hamilton JD, Susmel R (1994) Autoregressive Conditional Heteroskedasticity and Changes in Regime. Journal of Econometrics 64(2):307–333

Hardy M (2001) A Regime-Switching Model of Long-Term Stock Returns. North American Actuarial Journal 5(2):41–53

Hoyland K, Kaut M, Stein W (2003) A Heuristic for Moment-Matching Scenario Generation. Computational Optimization and Applications 24(2-3):169 – 185

Kaut M, Stein W (2003) Evaluation of scenario-generation methods for stochastic programming. Stochastic Programming E-Print Series (SPEPS) 14

Krokhmal P, Palmquist J, Uryasev S (2002) Portfolio optimization with conditional value-at-risk objective and constraints. The Journal of Risk 4(2):21–41

Mamon R, Erlwein C, Gopaluni R (2008) Adaptive signal processing of asset price dynamics with predictability analysis. Information Sciences 178(1):203–219

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

Messina E, Toscani D (2007) Hidden markov models for scenario generation. IMA Journal of Management Maths Advance Access published online on October 8, 2007

Mulvey J (2001) Introduction to financial optimization: mathematical programming special issue. Mathematical Programming 89(2):205–216

Pflug G (2000) Some remarks on the value-at-risk and the conditional value-at-risk, probabilistic constrained optimization: methodology and applications. Kluwer Academic Publishers 38:272–281

Rabiner LR (1989) A tutorial on Hidden Markov Models and selected applications in speech recognition. Proc of the IEEE 77(2):257–286

Rabiner LR (2002) Hidden Markov Processes. IEEE Transactions on Information Theory 48(6):1518– 1569

Rockafellar R, Uryasev S (2000) Optimization of Conditional Value-at-Risk. Journal of Risk 2(3):21–41

Rockafellar R, Uryasev S (2002) Conditional value-at-risk for general loss distributions. Journal of Banking and Finance 26:1443–1471

Roman D, Mitra G, Spagnolo N (2008) Hidden markov models for financial optimisation problems. CARISMA Technical report CTR/77/08, available on-line at www.carisma.brunel.ac.uk, to appear in: IMA Journal of Management Maths

Ross S (2002) An Elementary Introduction to Mathematical Finance. Cambridge University Press

Zhang Y (2004) Prediction of financial time series with hidden markov models. M.Sc. Thesis, Simon Fraser University, China