European Journal of Operational Research 234 (2014) 469–480

Contents lists available at SciVerse ScienceDirect

European Journal of Operational Research

journal homepage: www.elsevier.com/locate/ejor

Mean–variance optimal portfolios in the presence of a benchmark with applications to fraud detection

C. Bernard, S. Vanduffel *

[a] University of Waterloo, Department of Statistics and Actuarial Sciences, 200 University Avenue West, Waterloo, ON, N2L3G1, Canada
[b] Vrije Universiteit Brussel, Faculty of Economics, Pleinlaan 2, 1050 Brussels, Belgium

ARTICLE INFO

Article history:
Available online 27 June 2013

Keywords:
Mean–variance
Fraud detection
Optimal portfolio
Correlation constraints

ABSTRACT

We first study mean–variance efficient portfolios when there are no trading constraints and show that optimal strategies perform poorly in bear markets. We then assume that investors use a stochastic benchmark (linked to the market) as a reference portfolio. We derive mean–variance efficient portfolios when investors aim to achieve a given correlation (or a given dependence structure) with this benchmark. We also provide upper bounds on Sharpe ratios and show how these bounds can be useful for fraud detection. For example, it is shown that under some conditions it is not possible for investment funds to display a negative correlation with the financial market and to have a positive Sharpe ratio. All the results are illustrated in a Black–Scholes market.

1. Introduction

Markowitz (1952) and Roy (1952) were first in proposing a quantitative approach to determine the optimal trade-off between mean (return) and variance (risk). Their framework is nowadays known as mean–variance analysis and has become very influential as it combines algebraic simplicity with practical applicability. Markowitz pursued the study of optimal investment portfolios and his seminal works initiated a tremendous amount of research heading in several directions, ranging from the study of other notions for measuring risk, multi-period models (Cui et al., 2013; Mossin, 1968), non-negative final wealth (Korn & Trautmann, 1995) and imperfect markets (Lim, 2004; Xia & Yan, 2006) to the inclusion of ambiguity on the returns (Goldfarb & Iyengar, 2003) or an uncertain horizon (Martellini & Urosevič, 2006). See also Leung, Wong, and Ng (2012) and Zhang, Zhang, and Xiao (2009). More recently, several authors have been working on quadratic hedging or mean–variance hedging, which corresponds to the problem of approximating, with minimal mean squared error, a given payoff by the final value of a self-financing trading strategy in a financial market; see e.g., Lim (2006), Pham (2000) and Schweizer (1992, 2010), to cite only a few.

In an important contribution, Basak and Chabakauri (2010) have fully characterized time-consistent dynamic mean–variance optimal strategies. At any date prior to maturity, a time-consistent optimal strategy is the best possible mean–variance efficient allocation of wealth, assuming that an optimal mean variance strategy is also selected at each later instant in time. Mean–variance optimal strategies that are derived in a static setting[1] violate time-consistency in the sense that it may become optimal for a mean–variance investor to deviate away from this optimal mean–variance strategy during the investment horizon. However, these optimal strategies (derived in the static setting) can still be justified by assuming that the investor is pre-committed at time *0* and thus executes the dynamic investment strategy that has been decided at time $t = 0$. While time-consistency is a natural requirement, the assumption of pre-commitment is compatible with an investment practice in which a (retail) investor purchases a financial contract (from a financial institution) and does not trade (herself) afterwards.[2] It also fits with the behavior of an investment manager who revisits (optimizes) his portfolio periodically and sticks to his strategy between two dates. In practice, managers and other investors may also have additional constraints when optimizing their portfolio. One motivation for having constraints is that optimal (unconstrained) strategies are typically long with the market index and perform poorly in poor economic situations (Bernard, Boyle, & Vanduffel, 2013). In this paper, we show that the static setting is well suited to deal with a certain type of constraints that we motivate economically. See also Wang and Forsyth (2011) for a numerical approach of mean variance efficiency in a time-consistent framework

[1] By a "static setting", we mean that the strategy is derived at the initial time $t = 0$ as the mean–variance efficient optimum with respect to the terminal wealth W_T without consideration for its properties at intermediate dates.

[2] This is also consistent with the work of Goldstein, Johnson, and Sharpe (2008) who propose a tool that allows consumers to specify their desired probability distribution of terminal wealth at maturity.

* Corresponding author. Tel.: +32 2 6291046; fax: +32 2 6292060.
E-mail addresses: c3bernar@uwaterloo.ca (C. Bernard), steven.vanduffel@vub.ac.be (S. Vanduffel).

0377-2217/$ - see front matter © 2013 Elsevier B.V. All rights reserved.
http://dx.doi.org/10.1016/j.ejor.2013.06.023

and numerical comparisons of pre-committed strategies and time-consistent strategies.

Traditional mean–variance optimization consists in finding the best pre-committed allocation of assets assuming a buy-and-hold strategy or a constant-mix strategy (which requires a dynamic rebalancing to ensure a constant percentage invested in each asset). The question raises then how pre-committed mean–variance efficient portfolios can be derived when all strategies are allowed and available. Of course, allowing for more trading strategies and thus more degrees of freedom will further enhance optimality. The *first contribution* in this paper is to derive optimal mean–variance strategies in this setting. We show that the optimal portfolio consists of a short position in the stochastic discount factor used for pricing derivatives and a long position in cash. We are also able to compute the maximum possible Sharpe ratio (Sharpe, 1967) of an optimal (mean–variance efficient) strategy. Bounds on the Sharpe ratio can be useful to regulators or other market participants for fraud detection, i.e. to assess whether the reported performance of a strategy is feasible or not. Recall for example that the Sharpe ratio of Madoff's strategy lied far above the maximum Sharpe ratio for plausible strategies (Bernard & Boyle, 2009).

In the second part of the paper we extend our study to the case when there is additional information on the strategy, for example on the way it interacts with the financial market or any other benchmark asset as a source of background risk. Our *second contribution* is then to derive tighter bounds on the Sharpe ratio. This is useful for improved fraud detection or abnormal performance reporting. For example, it is shown that under some conditions it is not possible for investment funds to display negative correlation with the financial market and to have a positive Sharpe ratio.

Considering the interaction with a benchmark asset is also a natural way to make mean–variance efficient strategies more resilient against declining markets. Indeed, the mean–variance efficient portfolios derived in the first part of the paper provide no protection against bear markets. In practice, many investors reward strategies that offer protection or, more generally, that exhibit some desired dependence with any other source of background risk (which we refer to as a benchmark). Our *third contribution* is to derive mean–variance optimal allocation schemes for investors who exhibit state-dependent preferences in the sense that they care about the first two moments of the strategy's distribution *and* additionally aim at obtaining a desired correlation or dependence with a benchmark asset.

The rest of the paper is organized as follows. The optimal portfolio problem and the assumptions on the financial market are presented in Section 2. Section 3 provides explicit expressions for mean–variance efficient portfolios when there are no trading constraints as well as a first application to fraud detection. Sections 4 and 5 extend these preliminary results to the case when there are constraints on the correlation (respectively the dependence) with a benchmark and illustrate how these results are particularly useful to improve fraud detection tools. Final remarks are presented in Section 6.

2. Market setting

In this section, we provide our main assumptions and definitions. Let $(\Omega, \mathcal{F}, \mathbb{P})$ be a probability space that describes a financial market. Assume that all market participants agree to use a (non-negative) stochastic discount process $(\xi_t)_t$ for pricing, i.e. the price at time 0 for a strategy with terminal payoff X_T (paid at time $T > 0$) writes as

$$c(X_T) = \mathbb{E}[\xi_T X_T]. \tag{1}$$

Note that the price of the unit cash-flow at time T is given by $c(1) = \mathbb{E}[\xi_T]$ and we define the risk-free rate r such that $e^{-rT} = \mathbb{E}[\xi_T]$. All payoffs X_T are assumed to be square integrable ensuring that $c(X_T) < +\infty$. In particular $\text{var}(\xi_T) < +\infty$. We remark that this practice is usually motivated by assuming a frictionless and arbitrage-free financial market where the usual definition of absence of arbitrage is employed.[3] In particular, we do not take into account transaction costs (Pelsser & Vorst, 1996). When the market is complete (all payoffs can be replicated) the stochastic discount factor ξ_T is uniquely given, but in general an infinite number of choices is possible. However, using a milder notion of arbitrage, Platen and Heath (2009) argue that under some conditions, the stochastic discount factor ξ_T corresponds to the inverse of the so-called Growth Optimal Portfolio (GOP)[4] and also that the latter can be proxied by a market index. This motivates why in the remainder of the paper we refer to $1/\xi_t$ as *"the market index"* and we denote it by S_t^*. The pricing formula (1) can then be interpreted as the arithmetic average of the possible outcomes all expressed in units of the market index. Note how low values for the market index S_t^* correspond to high values for the discount factor ξ_t. This is consistent with economic theory in the sense that the states of a downturn are usually the most expensive states to insure and therefore correspond to the states ω where the highest values for the discount factor $\xi_T(\omega)$ are observed.

In the remainder of the paper, we consider an investor with a fixed horizon $T > 0$ without intermediate consumption. We denote by $W_0 > 0$ her initial wealth. For convenience, we assume that all ξ_t ($t > 0$) are continuously distributed.

3. Unconstrained mean–variance optimal portfolios

3.1. Mean–variance efficiency

Roy (1952) and Markowitz (1952) propose a quantitative approach to find mean–variance efficient allocation among risky assets assuming a buy-and-hold strategy. Their technique can also be applied in the context of constant-mix strategies. In this section we study mean–variance efficient portfolios when there are no restrictions on the possible strategies. Finding optimal policies turns out to be surprisingly simple. Indeed, let us first observe that an optimal mean–variance efficient final payoff X_T^* must necessarily be the cheapest possibility to generate a maximum mean for the given variance level.[5] Otherwise it is easy to contradict the optimality of this payoff. Indeed, if the optimum is not the cheapest strategy then there is thus another strategy that is cheaper and also has maximum mean. The cost benefit can be invested in the risk free account and one obtains a strategy that has a higher mean for the same variance, which contradicts the mean–variance efficiency of the strategy.

Constructing a cheapest strategy amounts to minimizing the price (1). Observe that since $\text{std}(X_T^*)$ and $\text{std}(\xi_T)$ are both fixed and finite, minimizing the price (1) is equivalent to minimizing the correlation[6] between X_T^* and the discount factor ξ_T. It is then a standard result in statistics that this occurs if and only if the optimal

[3] The no-free-lunch with vanishing risk (NFLVR) concept is the prevalent way to describe absence of arbitrage. In their fundamental theorem of asset pricing, Delbaen and Schachermayer (1994) essentially state that NFLVR is equivalent to the existence of a stochastic discount factor (also called state-price density process).

[4] The Growth Optimal Portfolio is a diversified strategy which ultimately outperforms all other strategies with probability one. In the literature it also appears as the *Kelly portfolio*.

[5] In other words, the optimal mean–variance portfolio must be cost-efficient in the sense defined by Bernard, Boyle, et al. (2013).

[6] $\text{corr}(X_T, \xi_T) = \frac{\mathbb{E}[\xi_T X_T] - \mathbb{E}[\xi_T]\mathbb{E}[X_T]}{\text{std}(\xi_T)\,\text{std}(X_T)}$. As the moments of ξ_T are given, it follows that for given moments $\mathbb{E}[X_T]$ and $\text{std}(X_T)$ minimizing $\mathbb{E}[\xi_T X_T]$ is equivalent to minimizing correlation $\text{corr}(X_T, \xi_T)$.