European Journal of Operational Research 236 (2014) 657–663

Contents lists available at ScienceDirect

European Journal of Operational Research

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

Decision Support

No-arbitrage bounds for financial scenarios

Alois Geyer[a,b,*], Michael Hanke[c,1], Alex Weissensteiner[d,e,2,3]

[a] WU (Vienna University of Economics and Business), D4.4.144, Welthandelsplatz 1, A-1020 Vienna, Austria
[b] Vienna Graduate School of Finance, D4.4, Welthandelsplatz 1, A-1020 Vienna, Austria
[c] Institute for Financial Services, University of Liechtenstein, Liechtenstein
[d] School of Economics and Management, Free University of Bolzano, Italy
[e] Technical University of Denmark, Dept. of Management Engineering, Denmark

ARTICLE INFO

Article history:
Received 5 July 2013
Accepted 13 January 2014
Available online 23 January 2014

Keywords:
Finance
Scenarios
No-arbitrage bounds
Financial optimization

ABSTRACT

We derive no-arbitrage bounds for expected excess returns to generate scenarios used in financial applications. The bounds allow to distinguish three regions: one where arbitrage opportunities will never exist, a second where arbitrage may be present, and a third, where arbitrage opportunities will always exist. No-arbitrage bounds are derived in closed form for a given covariance matrix using the least possible number of scenarios. Empirical examples illustrate the practical potential of knowing these bounds.

1. Introduction

We are interested in constructing discrete scenarios and simulations of financial asset returns which are free from arbitrage. This requirement for financial optimization models has been pointed out by, among others, Klaassen (2002); Sodhi (2005), or Geyer, Hanke, and Weissensteiner (2010, 2013). To that end, we investigate the *theoretical* relation between expected excess returns and the associated arbitrage opportunities. Our innovation is to provide bounds for expected excess returns to determine whether or not arbitrage is possible theoretically, *before* simulations have been carried out. These bounds hold for the least possible number of scenarios irrespective of the particular algorithm to be used, provided the algorithm matches the given covariance matrix.

The present paper can be put into the context of many methods[4] which have been developed to obtain discrete approximations of continuous distributions. Such methods attempt to find a compromise between (statistical) accuracy and the curse of dimensionality, or focus on computational efficiency. For example, in moment matching (see, e.g., Høyland, Kaut, & Wallace, 2003; Høyland & Wallace, 2001) the pre-specified moments of asset returns are matched using a rather small number of scenarios (i.e. discrete mass points) to obtain *statistically* acceptable approximations of continuous distributions. As Klaassen (2002) has pointed out, the resulting scenarios (or trees) may allow for arbitrage opportunities, since this aspect is not controlled for in moment matching algorithms. Klaassen describes how arbitrage opportunities can be detected, and emphasizes the need to routinely check for arbitrage *after* each simulation. Alternatively, he suggests adding constraints to the nonlinear moment matching problem. While he points out that adding constraints "will complicate the numerical optimization" (p. 1516), he does not provide any (numerical) evidence on the severity of this complication, however. As a matter of fact, we are not aware of any moment matching or other scenario generation algorithm which includes such constraints to prevent arbitrage at the outset.

Klaassen (1998) considers the problem of reducing the size of an already arbitrage-free tree while maintaining the absence of arbitrage. He suggests aggregating trees across states and/or time to obtain smaller trees with comparable properties. However, since his approach only works under the risk-neutral measure, it is not suitable for financial applications like portfolio optimization as pointed out by Kouwenberg (2001) and Geyer et al. (2010, 2013).

Scenario reduction methods as proposed by, for example, Pflug (2001) or Heitsch and Römisch (2003, 2009) are driven by a similar objective. They retain only a few paths from a very large number of scenarios such that the approximate (i.e. "reduced") and the original distribution are close in terms of some probability metric. However, applying scenario reduction techniques entails the risk of arriving at scenario trees which admit arbitrage opportunities.

* Corresponding author at: WU (Vienna University of Economics and Business), D4.4.144, Welthandelsplatz 1, A-1020 Vienna, Austria. Tel.: +43 1 31336 4559.
E-mail addresses: alois.geyer@wu.ac.at (A. Geyer), michael.hanke@uni.li (M. Hanke), alex.weissensteiner@unibz.it, alwe@dtu.dk (A. Weissensteiner).
[1] Tel.: +423 265 1155.
[2] Tel.: +39 0471 013496.
[3] Tel.: +45 4525 4401.
[4] Kall and Mayer (2011, p. 366) provide an overview; the most important methods are reviewed below.

http://dx.doi.org/10.1016/j.ejor.2014.01.027
0377-2217/

A necessary condition for the absence of arbitrage is that the branching factor (i.e., the number of arcs emanating from a node) at each node of the tree must at least equal the number of non-redundant assets in the optimization problem (e.g. Harrison & Kreps, 1979). While it is straightforward to use that minimum branching factor, resulting trees may still admit arbitrage opportunities, since that condition is not sufficient. This is complicated by the fact that existing implementations of scenario reduction algorithms do not allow for controlling the branching factor for each node in the tree. Pflug (2001) proposes a scenario reduction method to optimally discretize a continuous distribution by minimizing the supremum of the distance between the objective function evaluated using the original probability distribution and its discrete approximation. However, Geyer et al. (2010) show that the equivalence of the original and approximated problems rests on the (implicit) assumption that the sup-distance is finite. This assumption is violated when there are arbitrage opportunities in the approximated problem. In principle, no-arbitrage constraints could be added to the discretization problem. However, deriving a constrained solution for that problem in closed form does not seem to be a trivial task. As it stands, trees obtained by scenario reduction must subsequently be checked for arbitrage.

King, Koivu, and Pennanen (2005) use a Gauss-Hermite process for numerically computing bounds for the arbitrage-free prices of an option. An attractive feature of these processes is that the discretized one-step conditional probabilities match a maximum number of moments of the normal distribution (i.e. with v branches, the Gauss-Hermite quadrature matches $2v-1$ moments). The no-arbitrage bounds for an option are based on convex duality properties and require some other options available for trading. In this respect their paper also belongs to those which require solving an optimization problem to make statements about no-arbitrage (bounds).

This review of the literature makes clear that existing simulation methods mainly focus on statistical properties. The aspect of no-arbitrage, which is key in financial optimization, usually only enters ex-post. The (occasionally) proposed inclusion of constraints to guarantee no-arbitrage has not been implemented in any of the well-known scenario generation procedures we are aware of. In this paper we take a point of view which *starts out* from the no-arbitrage requirement. We investigate how the pre-specified first and second moments *determine* arbitrage possibilities in the discrete state space of simulated returns. This point of view and the associated results distinguish the present paper from the literature. Before simulations have been done, we are able to answer the question whether a vector of expected excess returns does allow for arbitrage or not. As a main contribution of the paper, this replaces the usual tests for arbitrage which require solving a linear program (see Klaassen, 2002). For a given covariance matrix we derive bounds for expected excess returns in closed form. Using well-known results from linear algebra and standard techniques from convex optimization we are able to distinguish three possible cases on the basis of two concentric (hyper) ellipsoids. These separate the space of all possible expected excess returns into three regions: (a) Arbitrage opportunities always exist. In this case simulation or re-sampling need not even be attempted. For the application at hand the intended stochastic features of asset returns need to be reconsidered, i.e. assumptions about covariances and/or expected excess returns need to be altered, thereby taking into account that these parameters may have been estimated with large standard errors. (b) Arbitrage may or may not exist depending on the sample at hand. In this case arbitrage checks are still required to decide whether re-sampling is necessary or not. However, the distance of the vector of expected excess returns from the origin of the (hyper) ellipsoids indicates how likely the need for re-sampling is *ex-ante*. This distance can also be used in cases (a) and (b) to quantify the required change in expected excess returns to guarantee no-arbitrage. (c) Arbitrage opportunities will never exist. Knowing this has the advantage that ex-post arbitrage checks become redundant.

To derive these bounds and to identify these regions we make use of the Fundamental Theorem of Asset Pricing. We only assume that the assets' first and second moments exist, and make no assumptions about the distribution of returns. To account for the aspect of dimensionality we derive the bounds for the smallest possible state space. This implies that the associated trees have the least possible size but still match the first two moments of asset returns. We finally illustrate how information about bounds can be empirically derived and used.

In Section 2 we derive no-arbitrage bounds for expected excess returns. We also provide results from empirical illustrations in Section 3. Main features of the approach are summarized in Section 4.

2. No-arbitrage bounds

The results on no-arbitrage bounds derived below will hold *irrespective* of a particular realization of excess returns $\mathbf{R}$. Thus, the no-arbitrage properties of expected excess returns $\boldsymbol{\mu}$ can be judged ex-ante (i.e. before simulations are run). However, for *deriving* these bounds it is instructive to start by considering no-arbitrage conditions associated with a particular realization, and then generalize to any realization. In Section 2.1 we show how to simulate returns such that their covariance is exactly matched. In Section 2.2 we state necessary and sufficient conditions for no-arbitrage. In Section 2.3 we explore the discrete state space associated with simulated returns. We use the conditions established in Section 2.2 to derive general bounds for expected excess returns to rule out arbitrage opportunities in simulated returns (first for a particular realization and subsequently for any realization).

2.1. Return realizations

To simulate single-period returns we proceed in two steps. First, we generate realizations of *mean-zero* returns $\mathbf{Y}$ with the target covariance $\boldsymbol{\Phi}$. Second, we consider *excess* returns $\mathbf{R}$ with mean $\boldsymbol{\mu}$. All *numerical* results in the main part of the paper are based on normal random numbers. However, while any simulation requires distributional assumptions, we emphasize that all conclusions and closed-form results on no-arbitrage bounds derived in Section 2.3 *do not* depend on the distribution of returns. We only assume that first and second moments exist.

When defining the number of scenarios m (i.e. the number of realizations of random returns) we need to account for three aspects: (a) According to Harrison and Kreps (1979) and Harrison and Pliska (1981) a necessary condition for no-arbitrage is $m \geqslant n$ (i.e. at least as many discrete states as assets). (b) As we are going to show in the next paragraph, more realizations than assets are required to exactly match the covariance matrix $\boldsymbol{\Phi}$, which implies $m \geqslant n+1$. (c) Our objective is to match the covariance and to rule out arbitrage with the smallest possible scenarios (or trees) to safeguard against the curse of dimensionality. Therefore, for the rest of Section 2 we set $m \equiv n+1$; this leads to the minimal tree size, which satisfies the necessary condition for no-arbitrage and allows to match the covariance exactly. At the same time, this also corresponds to considering a complete market with n risky assets and a risk-free asset.

To construct $\mathbf{Y}$ we generate an $m \times n$ matrix of realizations $\mathbf{X}$ which are required to have the following properties: each row is equally likely, the mean of each column is zero, the standard deviation of each column is one, and the columns of $\mathbf{X}$ are orthogonal (i.e. $(\mathbf{X}'\mathbf{X})/m = \mathbf{I}$). In general, the columns of a simulated $m \times n$