Mathematical Finance, Vol. 12, No. 3 (July 2002), 239–269

PORTFOLIO VALUE-AT-RISK WITH HEAVY-TAILED RISK FACTORS

Paul Glasserman

Graduate School of Business, Columbia University

Philip Heidelberger

IBM Research Division, Yorktown Heights, NY

Perwez Shahabuddin

IEOR Department, Columbia University

This paper develops efficient methods for computing portfolio value-at-risk (VAR) when the underlying risk factors have a heavy-tailed distribution. In modeling heavy tails, we focus on multivariate t distributions and some extensions thereof. We develop two methods for VAR calculation that exploit a quadratic approximation to the portfolio loss, such as the delta-gamma approximation. In the first method, we derive the characteristic function of the quadratic approximation and then use numerical transform inversion to approximate the portfolio loss distribution. Because the quadratic approximation may not always yield accurate VAR estimates, we also develop a low variance Monte Carlo method. This method uses the quadratic approximation to guide the selection of an effective importance sampling distribution that samples risk factors so that large losses occur more often. Variance is further reduced by combining the importance sampling with stratified sampling. Numerical results on a variety of test portfolios indicate that large variance reductions are typically obtained. Both methods developed in this paper overcome difficulties associated with VAR calculation with heavy-tailed risk factors. The Monte Carlo method also extends to the problem of estimating the conditional excess, sometimes known as the conditional VAR.

Key Words: value-at-risk, delta-gamma approximation, Monte Carlo, simulation, variance reduction, importance sampling, stratified sampling, conditional excess, conditional value-at-risk

1. INTRODUCTION

A central problem in market risk management is estimation of the profit-and-loss distribution of a portfolio over a specified horizon. Given this distribution, the calculation of specific risk measures is relatively straightforward. Value-at-risk (VAR), for example, is a quantile of this distribution. The expected loss and the expected excess loss beyond some threshold are integrals with respect to this distribution. The difficulty in estimating these types of risk measures lies primarily in estimating the profit-and-loss distribution itself, especially the tail of this distribution associated with large losses.

P. Glasserman and P. Shahabuddin were partially supported by NSF NYI Award DMI9457189 and NSF Career Award DMI9625297, respectively.

Manuscript received August 2000; final revision received January 2001.

Address correspondence to Paul Glasserman, 403 Uris Hall, Columbia University, New York, NY 10027; pg20@columbia.edu.

All methods for estimating or approximating the distribution of changes in portfolio value rely (at least implicitly) on two types of modeling considerations: assumptions about the changes in the underlying risk factors to which a portfolio is exposed, and a mechanism for translating these changes in risk factors to changes in portfolio value. Examples of relevant risk factors are equity prices, interest rates, exchange rates, and commodity prices. For portfolios consisting of positions in equities, currencies, commodities, or government bonds, mapping changes in the risk factors to changes in portfolio value is straightforward. But for portfolios containing complex derivative securities this mapping relies on a pricing model.

The simplest and perhaps most widely used approach to modeling changes in portfolio value is the variance-covariance method popularized by RiskMetrics (1996). This approach is based on assuming (i) that changes in risk factors are conditionally multivariate normal over a horizon of, say, one day, two weeks, or a month, and (ii) that portfolio value changes linearly with changes in the risk factors. ("Conditionally" here means conditional on information available at the start of the horizon; the unconditional distribution need not be normal.) Under these assumptions, the portfolio profit-and-loss distribution is conditionally normal; its standard deviation can be calculated from the covariance matrix of the underlying risk factors and the sensitivities of the portfolio instruments to these risk factors. The attraction of this approach lies in its simplicity. But each of the assumptions on which it relies is open to criticism, and research in the area has tried to address the shortcomings of these assumptions.

One line of work has focused on relaxing the assumption that portfolio value changes linearly with changes in risk factors while preserving computational tractability. This includes, in particular, the "delta-gamma" methods developed in Britten-Jones and Schaefer (1999), Duffie and Pan (2001), Rouvinez (1997), and Wilson (1999). These methods refine the relation between risk factors and portfolio value to include quadratic as well as linear terms. Methods that combine interpolation approximations to portfolio value with Monte Carlo sampling of risk factors are considered in Jamshidian and Zhu (1997), Picoult (1999), and Shaw (1999). Low variance Monte Carlo methods based on exact calculation of changes in portfolio value are proposed in Cardenas et al. (1999), Glasserman, Heidelberger, and Shahabuddin (2000), and Owen and Tavella (1999).

Another line of work has focused on developing more realistic models of changes in risk factors. It has long been observed that market returns exhibit systematic deviation from normality: across virtually all liquid markets, empirical returns show higher peaks and heavier tails than would be predicted by a normal distribution, especially over short horizons. Early studies along these lines include Mandelbrot (1963), Fama (1965), Praetz (1972), and Blattberg and Gonedes (1974). More recent investigations, some motivated by value-at-risk, include Bouchaud, Sornette, and Potters (1997), Danielsson and de Vries (1997), Eberlein and Keller (1995), Eberlein, Keller, and Prause (1998), Embrechts, McNeil, and Straumann (2001), Hosking, Bonti, and Siegel (2000), Huisman et al. (1998), Koedijk, Huisman, and Pownall (1998), McNeil and Frey (1999), Heyde (1999). Using different approaches to the problem and different sets of data, these studies consistently find high kurtosis and heavy tails. Moreover, most studies find that the tails in financial data are not so heavy as to produce infinite variance (as would be implied by a nonnormal stable distribution), though higher order moments (e.g., fifth and higher) may be infinite.

This paper contributes to both lines of investigation by developing methods for calculating portfolio loss probabilities when the underlying risk factors are heavy