

A New Approach to Comparing VaR Estimation Methods

CHRISTOPHE PÉRIGNON AND DANIEL R. SMITH

CHRISTOPHE PÉRIGNON is an associate professor at HEC in Paris, France. perignon@hec.fr

DANIEL R. SMITH is an associate professor in the faculty of business at Simon Fraser University in Vancouver, Canada, and a visiting fellow in the faculty of business in the Queensland University of Technology in Brisbane, Australia. dsmith@sfu.ca

We develop a novel backtesting framework based on multidimensional value-at-risk (VaR) that focuses on the left tail of the distribution of bank trading revenues. Our coverage test is a multivariate generalization of Kupiec's unconditional test. Applying our method to actual daily bank trading revenues, we find that nonparametric VaR methods, such as GARCH-based methods or filtered historical simulation, work best for bank trading revenues.

Value-at-risk (VaR) is the standard measure for market risk used by financial institutions and banking regulators. VaR quantifies the loss that a bank can face on its trading portfolio within a given period and for a given confidence interval. Since the first widely publicized appearance of the term VaR in a G-30 report in July 1993, numerous statistical methods have been proposed to compute this market risk measure, such as RiskMetrics and historical simulation (HS). Furthermore, over the past 10 years, an increasing number of banks have been setting their regulatory capital using in-house VaR measures. The quest for the most accurate VaR method is of great interest to regulators and risk managers in charge of developing banks' proprietary risk models. Unfortunately, current backtesting procedures for VaR models are known to lack power.

The key methodological contribution of our article is the development of a new

backtesting framework for VaR. It is based on multidimensional VaR, which is a vector of VaRs measured with the same horizon but different coverage probabilities or confidence levels. Our basic idea is that the accuracy of a given VaR method should not be assessed using only a single coverage probability or, in other words, a single observation on the trading revenue distribution. Instead, we focus on the left tail of the distribution of the trading revenues and consider K different confidence levels. Our coverage test is a multivariate generalization of the unconditional coverage test of Kupiec [1995], which remains the standard backtesting procedure. Furthermore, our test can be implemented with any combination of coverage probabilities or any time horizon—for example, [1% 2% ... 10%], [0.1% 0.2% 0.3%] or 1 day, 10 days.

To date, almost all empirical comparisons of VaR methods have been based on hypothetical positions in individual assets, interest rates, exchange rates, or stock indices. For instance, Ferreira and Lopez [2005] consider an equally weighted portfolio of short-term fixed-income positions in the U.S. dollar, German deutschemark, and Japanese yen using daily interest rate and foreign exchange data from 1979 through 2000.¹ The reason for this practice is because actual trading positions and daily trading revenues of banks are typically unknown to the public. Unlike previous studies, we use actual daily trading revenues