Journal of Banking & Finance 37 (2013) 3843–3854

Contents lists available at SciVerse ScienceDirect

Journal of Banking & Finance

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

The Risk Map: A new tool for validating risk models

Gilbert Colletaz [a], Christophe Hurlin [a], Christophe Pérignon [b,*]

[a] University of Orléans, Laboratoire d'Economie d'Orléans, France
[b] Finance Department, HEC Paris, France

ARTICLE INFO

Article history:
Received 30 May 2012
Accepted 26 June 2013
Available online 9 July 2013

JEL classification:
G21
G28
G32

Keywords:
Financial risk management
Tail risk
Basel III

ABSTRACT

This paper presents a new method to validate risk models: the *Risk Map*. This method jointly accounts for the number and the magnitude of extreme losses and graphically summarizes all information about the performance of a risk model. It relies on the concept of a *super exception*, which is defined as a situation in which the loss exceeds both the standard Value-at-Risk (VaR) and a VaR defined at an extremely low probability. We then formally test whether the sequences of exceptions and super exceptions are rejected by standard model validation tests. We show that the Risk Map can be used to validate market, credit, operational, or systemic risk estimates (VaR, stressed VaR, expected shortfall, and CoVaR) or to assess the performance of the margin system of a clearing house.

1. Introduction

The need for sound risk management has never been more essential than in today's financial environment. Of paramount importance for risk managers and regulators is the ability to detect misspecified risk models as they lead to misrepresentation of actual risk exposures (Blöchlinger, 2012). In this paper, we focus on a popular family of risk models, namely the tail risk models. These models can generate a variety of risk measures including the Value-at-Risk (VaR), which is defined as an extreme quantile of a return or profit-and-loss (P&L) distribution, as well as stressed VaR, expected shortfall, and CoVaR. In practice, tail risk models are used to measure downside risk (Bali et al., 2009), construct portfolios (Basak and Shapiro, 2001), quantify financial institutions' exposures to, and capital requirements for, market, credit, and operational risks (Jorion, 2007), set margin requirements for derivatives users (Booth et al., 1997), and measure the systemic-risk contribution of a financial institution (Adrian and Brunnermeier, 2011).

In this paper, we present a new tool, called the *Risk Map*, for validating risk models. To grasp the intuition of our approach, consider two banks that both have a one-day Value-at-Risk (VaR) of \$100 million at the 1% probability level. This means that each bank has a one percent chance of losing more than \$100 million over the next day. Assume that, over the past year, each bank has reported three VaR exceptions, or days when the trading loss exceeds its VaR, but the average VaR exceedance is \$1 million for bank A and \$999 million for bank B. In this case, standard backtesting methodologies would indicate that the performance of both models is equal (since both models lead to the same number of exceptions) and acceptable (since the annual number of exceptions is close enough to its target value of 2.5 = 1% of 250 trading days). The reason is that current backtesting methodologies only focus on the number of VaR exceptions and totally disregard the magnitude of these exceptions (Berkowitz, 2001; Stulz, 2008).

However, in practice, market participants and regulators do care about the magnitude of their losses.[1] It is indeed the severity of the trading losses, and not the exceptions per se, that jeopardize the solvency of financial institutions. For instance, banking regulators may want to penalize more heavily – in terms of capital requirements – a bank that experiences extremely large exceptions than a bank that experiences moderate exceptions. Furthermore, it makes a big difference whether the margin of a given derivative market participant is exceeded by a small or by a large amount as it creates a shortfall of the same amount that the clearing house must cover. To the best of our knowledge, there is no general hypothesis-testing

* Corresponding author. Address: Finance Department, HEC Paris, 1 Rue de la Libération, 78351 Jouy-en-Josas, France. Tel.: +33 139 67 94 11.

E-mail addresses: gilbert.colletaz@univ-orleans.fr (G. Colletaz), christophe.hurlin@univ-orleans.fr (C. Hurlin), perignon@hec.fr (C. Pérignon).

[1] In a recent survey on trading risk, members of the Basel Committee state that "account for the severity of losses beyond the confidence threshold [...] is especially important for regulators, who are [...] concerned about exactly these losses" (Basel Committee on Banking Supervision, 2011b).

0378-4266/$ - see front matter
http://dx.doi.org/10.1016/j.jbankfin.2013.06.006

framework available in the literature that accounts for both the number and the magnitude of extreme losses. The objective of this paper is to fill this gap.

The Risk Map approach jointly accounts for the number and the magnitude of the VaR exceptions. The basic intuition is that a large loss not only exceeds the regular VaR defined with a probability α (e.g. 1%) but is also likely to exceed a VaR defined with a much lower probability α' (e.g. 0.2%). On this ground, we define a VaR *exception* as $r_t < -VaR_t(\alpha)$, where r_t denotes the P&L, and a VaR *super exception* as $r_t < -VaR_t(\alpha')$, with α' much smaller than α. As an illustration, we show in Fig. 1 the joint evolution of the daily P&L, $VaR(\alpha)$, and $VaR(\alpha')$ for a hypothetical portfolio. We see that, as expected, $-VaR(\alpha')$ is systematically more negative than $-VaR(\alpha)$ as $VaR(\alpha')$ is measured further left in the tail of the P&L distribution. For this portfolio, there are five exceptions and four super exceptions.

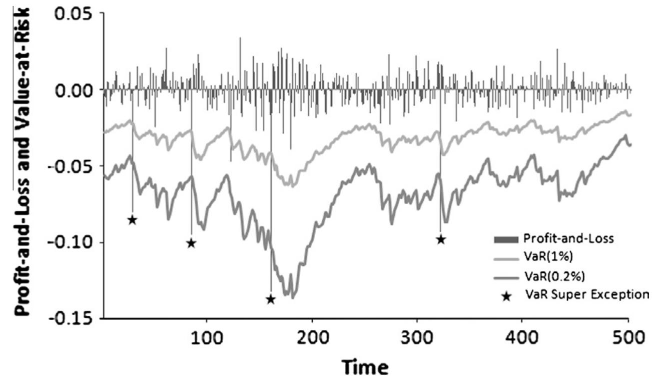

Fig. 1. VaR exception vs. VaR super exception. *Notes*: This figure displays the daily P&L, $VaR(\alpha)$, and $VaR(\alpha')$ for a hypothetical portfolio, with $\alpha = 1\%$ and $\alpha' = 0.2\%$. Both the P&L and VaR series are simulated using a t-Garch model. A VaR exception is defined as $r_t < -VaR_{t|t-1}(\alpha)$ whereas a super exception is defined as $r_t < -VaR_{t|t-1}(\alpha')$. Over this 500-day sample period, there are five exceptions and four super exceptions, which are highlighted with stars.

In practice, the choice of the probability α' is key. We show in this paper that this parameter can be either chosen freely by the user or defined endogenously. In the former case, the α' parameter must reflect the loss aversion of the investor whereas in the latter case, α' is set such that $VaR(\alpha')$ corresponds either to the stressed VaR or to the expected shortfall.

In order to validate the risk model, we formally test whether the sequences of exceptions and super exceptions satisfy standard backtesting conditions. Formally, we need to test the joint null hypothesis that the probability of having an exception is α and that the probability of having a super exception is α'. Several statistical strategies can be used to test this null hypothesis (e.g. log-likelihood ratio, hit regression test) and some statistical tests are readily available in the literature. We then report the p-value of the statistical test in a three-dimensional graph, which we call the Risk Map, that graphically summarizes all information about the performance of a risk model. To sum up, what the Risk Map approach allows us to do is to simplify an intricate problem defined over a {loss frequency, loss severity} domain to an easier problem defined over a {loss frequency, loss frequency} domain.

There are several advantages to the Risk Map approach. First, it preserves the *simplicity* of the standard validation techniques (Kupiec, 1995), while still accounting for the magnitude of the losses. Thus, the Risk Map approach is a three-dimensional generalization of the "Traffic Light" system (Basel Committee on Banking Supervision, 2006, 2011a) which remains the reference backtest methodology for banking regulators.[2] Second, it is a *formal* hypothesis testing framework that provides p-values and rejection ranges. Indeed, it allows us to jointly test the null hypothesis that both the numbers of VaR exceptions and super exceptions are accurate. Third, the Risk Map approach is *general* and can be used with any tail risk model as it only relies on the sequence of exceptions and super exceptions. In particular, no assumptions need to be made regarding the distribution of the P&L. For instance, it can be used to backtest the market VaR of a single asset, portfolio, trading desk, business line, bank, insurance company, mutual fund, or hedge fund (Berkowitz et al., 2011). It also permits to jointly validate the standard and stressed VaRs that banks must compute under Basel III, as well as expected shortfalls. Furthermore, the Risk Map can be used to backtest credit-risk VaRs (Lopez and Saidenberg, 2000), operational-risk VaRs (Dahen and Dionne, 2010), or VaR-based margins for derivative users (Cruz Lopez et al., 2013). Finally, we show that the Risk Map can be used to validate the systemic risk measure recently proposed by Adrian and Brunnermeier (CoVaR) as it is defined as the conditional quantile of a bank asset return. The Risk Map is, to the best of our knowledge, the first method allowing one to backtest a systemic risk measure.

The outline of the paper is as follows. In the next section, we describe our model validation framework. In Section 3, we present several applications of the Risk Map methodology that sequentially deal with market risk, systemic risk, and margin requirements. We summarize and conclude our paper in Section 4.

2. Validation framework

2.1. Background

Let r_t denote the return or P&L of a portfolio at time t and $VaR_{t|t-1}(\alpha)$ the ex-ante one-day ahead VaR forecast for an α coverage rate conditionally on an information set $\mathcal{F}_{t-1}$. In practice, VaR is computed using either non-parametric techniques, such as Historical Simulation (Pritsker, 2006), or parametric techniques, such as Monte Carlo simulation (Broadie et al., 2011). If the VaR model is adequate, then the following relation must hold:

$$\Pr[r_t < -VaR_{t|t-1}(\alpha)] = \alpha. \tag{1}$$

Let $I_t(\alpha)$ be the hit variable associated with the ex-post observation of a VaR(α) violation at time t:

$$I_t(\alpha) = \begin{cases} 1 & \text{if } r_t < -VaR_{t|t-1}(\alpha), \\ 0 & \text{otherwise.} \end{cases} \tag{2}$$

Backtesting procedures are typically only based on the violation process $\{I_t(\alpha)\}_{t=1}^{T}$. As stressed by Christoffersen (1998), VaR forecasts are valid if and only if this violation sequence satisfies the Unconditional Coverage (UC) hypothesis.[3] Under the UC hypothesis,

[2] Accounting for VaR exceedances is particularly important in the context of the Basel III regulation (Basel Committee on Banking Supervision, 2011a). Indeed, under Basel III, capital requirements for market risk depend on both the VaR and stressed VaR (calibrated to historical data from a continuous 12-month period of significant financial stress) of the bank:

$$c = \max\{VaR_{t|t-1}; m \cdot VaR_{avg}\} + \max\{sVaR_{t|t-1}; m_s \cdot sVaR_{avg}\}.$$

As the value of the multiplicative factors m and m_s only depends on the backtesting results of the VaR, but not of the stressed VaR, it is particularly important to account for the magnitude of the VaR exceedances.

[3] Validation tests are also based on the independence hypothesis (IND), under which VaR violations observed at two different dates for the same coverage rate must be distributed independently. Formally, the variable $I_t(\alpha)$ associated with a VaR violation at time t for a coverage rate α should be independent of the variable $I_{t-k}(\alpha)$, $\forall k \neq 0$. In other words, past VaR violations should not be informative about current and future violations. When the UC and IND hypotheses are simultaneously valid, VaR forecasts are said to have a correct Conditional Coverage (CC), and the VaR violation process is a martingale difference, with $\mathbb{E}[I_t(\alpha) - \alpha \mid \mathcal{F}_{t-1}] = 0$. For a test of the CC hypothesis, see Christoffersen (1998), Christoffersen and Pelletier (2004), Engle and Manganelli (2004), and Berkowitz et al. (2011).