

Risk Management in Emerging Markets: Practical Methodologies and Empirical Tests

Marios Nerouppos

Cyprus International Institute of Management, Cyprus

David Saunders

University of Waterloo, Canada

Costas Xiouros

University of Southern California, U.S.A.

Stavros A. Zenios

University of Cyprus, Cyprus

Risk management has undergone a remarkable transformation over the past fifteen years, with most new methods having been designed for the concerns of large institutions operating in well-developed financial markets. This paper addresses a problem faced by smaller institutions operating in emerging markets, namely the significant lack of data. As many risk management techniques are data intensive, this problem may seem insurmountable. This paper introduces a new method, enriched historical simulation, which supplements the data in an emerging market with data from other markets. The principle behind this methodology is that when many markets are considered, the essence of emerging market economies comes to the fore, with local idiosyncrasies being washed out. This principle is illustrated on the problem of estimating Value-at-Risk on the Cyprus and Athens Stock Exchanges. Numerical tests show that standard models underestimate risks, but that estimates are improved significantly with the use of external data (JEL: C10, C80, G10, G15).

Keywords: risk management, historical simulation, value-at-risk, emerging markets

* During this research, Marios Nerouppos was a research associate at RiskLab Cyprus at the Cyprus International Institute of Management. David Saunders is an assistant professor in the Department of Statistics and Actuarial Science, University of Waterloo, Canada. He gratefully acknowledges support from CLR stockbrokers in the form of a sponsored chair at the Cyprus International Institute of Management, from the University of Pittsburgh in the form of an internal grant, and from NSF grant: DMS-0310656. Costas Xiouros is a Ph.D. student at the Marshall School of Business at the University of Southern California. Stavros A. Zenios is professor and director of the HERMES European Center of Finance and Economics, University of Cyprus and Senior Fellow at the Financial Institutions Center, the Wharton School in Philadelphia, PA. This research was funded in part by the HERMES European Center of Excellence on Computational Finance and Economics under EU contract ICA1-CT-2000-70015.

(Multinational Finance Journal, 2006, vol. 10, no. 3/4, pp. 179–221)

© *Multinational Finance Society*, a nonprofit corporation. All rights reserved.

DOI: 10.17578/10-3/4-2

I. Introduction

Emerging markets present unique challenges for the design and implementation of risk management systems. There are many reasons why the time is ripe to study these challenges. The first is globalization and the fact that local economies are gradually integrating with more developed, highly competitive markets. The second is that the methodologies for risk management (and in particular market risk management) are at a stage where they are well developed and understood in advanced markets; the peculiarities of emerging markets can perhaps be understood as perturbations of these common models. The third is the adoption of the internal models approach for measuring market risk by the Basel Committee on Risk Management of the Bank for International Settlements. Emerging market banks face a substantial competitive disadvantage if they are forced to continue using the standardized approach. Finally, in emerging markets risk management is being developed concurrently with the financial system as a whole. This is in contrast to the situation in more developed economies where financial markets developed over time, and were quite advanced before risk management became a hot topic. It makes the study of risk management in such environments vital, as these economies try to “hit the ground running”.

A. The State of Risk Management in Emerging Markets

Risk managers in emerging markets face a number of challenges that do not present themselves to their colleagues in more developed economies. The first and most apparent is the often chaotic state of the local economy. The second is the short history of these markets. This has a number of significant consequences. One is the relative novelty of financial markets (both to institutions and households; this can be a major cause of speculative bubbles). Another problem, equally important from a risk management point of view, is that there is a startling scarcity of available data. Often, the institutional mechanisms that lead to the plethora of data in advanced markets do not exist (e.g. derivatives exchanges, secondary markets, and even regular auctions of a standard set of government bonds). Furthermore, those data that are available are contaminated for many reasons. Since many emerging markets have gone through some period of crisis, the history of local financial variables is of questionable value in calibrating mathematical

models for assessing future risks. Any current price data that are available must be viewed in light of the low volumes and liquidity of local markets. All of these factors lead to tremendous difficulties for risk management.

Emerging markets often bear significant liquidity risk. During periods of business as usual, volumes on the exchange are often extremely low, while during unusual periods volumes are extremely high. The monthly volumes and the level of the Cyprus Stock Exchange are presented in figure 1.¹ Institutional restrictions frequently prohibit short-selling, and it is unlikely to have a liquid derivatives market, or even a secondary market for instruments such as government bonds. Further evidence of the impressive illiquidity in some emerging markets is provided by the presence of large transactions costs. Figure 2 shows the average bid-ask spreads for both the Cypriot FTSE 20 and the German DAX 30 stock indices during the period 4/1/1999 to 29/12/2000. This figure illustrates that transactions costs in emerging markets can be dramatically greater than in developed markets, even during periods of business as usual. This situation is only exacerbated during times of crisis.

It is impossible to arrive at a completely satisfactory solution to the problems posed by the absence of data in emerging markets. There is no way to develop and calibrate a theoretically consistent model using the limited resources available while taking into account all of the (often significant) idiosyncrasies of the local market. In such an environment there is no choice but to attempt to develop methods that are intuitively plausible and effective in practice, both in historical tests and in the current market. This is the approach taken in this paper.

The main novelty in this paper is the use of data from other emerging markets, referred to as enriched historical simulation (EHS), to supplement risk management calculations in the local market. This approach is illustrated on the problem of calculating Value-at-Risk on the Cypriot and Greek stock exchanges. The justification for this method is that emerging markets share many common properties, as discussed above, and that when many of them are considered, local idiosyncrasies are dwarfed by the influence of these common factors. Therefore, data from other markets, while seemingly unrelated to the local market, may carry important information that is relevant to risk management. It should be noted that using foreign markets' experiences

1. In this paper, all stock exchange values are plotted in log-scale.




FIGURE 1.— Level and Monthly Volumes of the CSE Index

as possible scenarios for the local market in a stress testing framework has already been suggested by Dembo et al. (2000).

This paper discusses two methodological issues in emerging markets. The first is the best method for estimating the standard risk measure Value-at-Risk (VaR) (see, for example, Jorion ([2000])).² The second is how to address problems arising from the lack of data. The proposed method is to use data from other emerging markets to supplement the data in the local market for the purpose of scenario generation. The inclusion of data from many other emerging markets highlights the commonalities that exist between these markets, while local idiosyncrasies have less effect. Ultimately, the test of any scenario generation methodology is how well it performs in practice. This paper presents extensive tests of the proposed methodology on the stock exchanges of both Cyprus (a less developed market) and Athens (a more developed market).

In the present study, the focus is on the perspective of a local investor, with substantial (perhaps all) of its investment in the emerging

2. This should in no way be interpreted as an endorsement of VaR as an ideal (or even adequate) risk measure on the part of the authors. Rather, this measure is used in the current work due to its standing as a benchmark in financial risk management, and because of its practical importance in light of current regulations. Some of the (many) shortcomings of VaR along with some alternatives are discussed briefly later in the paper.




FIGURE 2. — Bid-ask Spreads for Cypriot and German Capital Markets

market. Consequently, considerations such as currency risk, which would play an important role in the risk management of, for example, a global emerging markets fund, are ignored in this study. In many developing markets (as was the case with Cyprus for the period of this study), there are strict controls on the amount of a portfolio that can be held outside the country. Investors are therefore forced to concentrate on the local market and its constituent risks. This motivates the focus of the current paper.

The remainder of the paper is structured as follows. The second section outlines the statistical models of returns that are used throughout the paper. The third section presents a statistical analysis of the history of the Cyprus and Athens Stock Exchanges. The fourth section discusses risk measurement in general, and in particular the industry standard measure Value-at-Risk. The fifth section discusses methodologies for scenario generation and VaR estimation, focusing on the new method of enriched historical simulation (i.e. on the use of data available from other markets for the purpose of risk measurement). The sixth section presents the results of back-testing our methodology on the Cyprus and Athens Stock Exchanges. The seventh section reviews the results and presents general conclusions.

II. Statistical Methodology

This section presents the basic statistical models of returns that will be employed throughout the rest of the paper. The return between time $t - 1$ and time t is defined to be:

$$r_t = \ln \left(\frac{S_t}{S_{t-1}} \right),$$

where S_t is the closing level of the index at time t . The methodologies that will be used can be divided into two categories: parametric and non-parametric. Effectively, these methodologies correspond to the two general models considered below.

A. Parametric Models

A general specification for the returns is given by:

$$r_t = \mu_t + \sigma_t \varepsilon_t$$

where, conditional on the information up to time t , μ_t and σ_t are the expected return and standard deviation of returns respectively, and ε_t are *i.i.d.* random shocks.

The Exponentially Weighted Moving Average (EWMA) and the Generalized Autoregressive Conditionally Heteroskedastic (GARCH(1,1)) models are the most common models of this class.

$$\text{EWMA: } \sigma_t^2 = \lambda \sigma_{t-1}^2 + (1 - \lambda)(r_{t-1} - \mu_{t-1})^2$$

$$\text{GARCH(1,1): } \sigma_t^2 = \omega + \alpha (r_{t-1} - \mu_{t-1})^2 + \beta \sigma_{t-1}^2$$

EWMA is just a special case of GARCH and is also known as Integrated-GARCH or IGARCH.³ Statistical tests and parameter

3. The EWMA model is a GARCH(1,1) with $\omega = 0$, $\beta = \lambda$ and $\alpha = 1 - \beta$. It has the advantage that there is only one parameter which can be easily estimated using Ordinary Least Squares (OLS) by minimizing the squared deviation of the model variance from the unexpected squared returns, i.e:

$$\hat{\lambda} = \arg \min \sum_{t=1}^n \left[\sigma_t^2 - (r_t - \mu_t) \right]^2$$

where r_t are the observed returns and μ_t is the expected mean. This implicitly assumes that

estimates for the above models are based on the assumption that the shocks ε_t are *i.i.d.* standard normal random variables.

For the purposes of this paper, it is assumed that the expected return at all times is zero, i.e. $\mu_t = 0$ for all t . This assumption is common in many market risk calculations and over small time horizons, such as those considered in this paper, is not significant.

B. Non-Parametric Models

In these models, the returns are assumed to be independently and identically distributed (*i.i.d.*).⁴ At time t , returns are simulated based on a discrete probability distribution P_t on a set of possible outcomes Ω_t . The method for determining the possible returns Ω_t and the corresponding probabilities P_t is what distinguishes each method in this class. Examples of models in this class include straightforward historical simulation, the weighted historical simulation algorithm of Boudoukh et al. (1998) and the enriched historical simulation method introduced in this paper.

III. Statistical Analysis of Emerging Market Exchanges

By a “bubble” we mean a particular pattern that often appears in emerging markets, when financial variables deviate for a prolonged time from their equilibrium values. The following sequence of market regimes is typical: The Calm Before the Storm: An initial period of low volumes and low volatility, when very few investors have entered the market; The Upswing: A period of rapid growth where many investors, both individual and institutional, enter the market. This period is marked

current volatility is given by the absolute value of the spot unexpected return. Moreover, this estimation technique, being a non-parametric one, does not require the shocks ε_t to be normally distributed. As a result, the normality test of the standardized returns is not an appropriate diagnostic to test the validity of the model. A better way to estimate the parameters of both of the models is to use Maximum Likelihood Estimation (MLE), which is consistent with the GARCH model and can be easily tested. Assuming the ε_t are *i.i.d.* standard normal random variables, the distribution of the return r_t conditional on the information up to and including time $t - 1$ is normal with mean μ_t and standard deviation σ_t . Therefore, using the conditional densities we can construct the likelihood function and estimate the parameters. The GARCH method is estimated using maximum likelihood and setting $\mu_t = 0$. For more details see the technical report by Nerouppos et al. (2002).

4. This can be achieved by taking $\mu = 0$ and $\sigma = 0$ so that $r_t = \varepsilon_t$.




FIGURE 3. — The Periods of the Speculative Bubble of the CSE

by high volumes and a rapid growth in the level of the general market index; The Crash: The euphoria subsides, as investors realize that the securities traded on the market are overvalued. Panic selling ensues, and the market plummets back towards its initial level; The Calm After the Storm: Volumes and volatility reduce as investors are once bitten, twice shy. Unfortunately, this stage is often the calm before the next storm.

The second and third of these stages will sometimes be grouped under the heading “the storm”. The typical stages are illustrated in figure 3. Speculative bubbles are a remarkable phenomenon of mass psychology. While it is not the purpose of this paper to address this directly, we point out the interesting work of MacKay (1995—first published in 1852) and Shiller (2000).

The Cyprus Stock Exchange (CSE) started operating on 29/3/1996 and its history can be best explained when divided into three periods. The first period was characterized by low investor interest, thus low volumes, particularly low volatility and persistence of the CSE General Index around the initial level of 100. The first period was followed by a typical speculative bubble, when interest in investment became substantial. It took less than one and a half years for the bubble to burst. The after-bubble period dropped the General Index back to the initial level. It was substantially less volatile than the second period, while at the same time significantly more volatile than the first. The complete

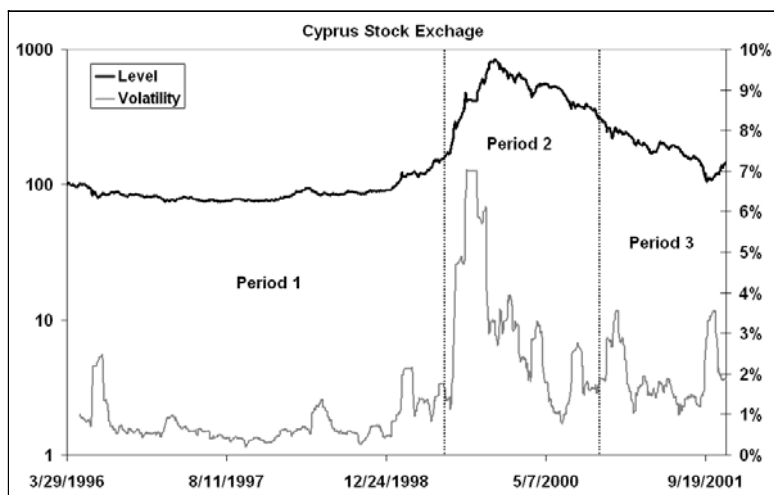


FIGURE 4. — Level and Volatility of the CSE Index

history of the CSE General Index is shown in figure 4, together with the period division and the monthly volatilities.⁵ The three periods are defined as: Period One: 29/3/1996 – 30/6/1999 (The Calm Before the Storm), Period Two: 1/7/1999 – 31/10/2000 (The Storm), Period Three: 1/11/2000 – 23/11/2001 (The Calm After the Storm). Of course, this division is based on a posteriori knowledge about the presence of the bubble. This knowledge is used only to test the models under different stock market regimes. It is not used in any way in developing the models.

The descriptive statistics of the daily returns on the CSE index both for the entire history as well as the three periods separately are presented in table 1. The assumption of unconditional normality is soundly rejected by statistical tests.⁶ Note particularly the positive skewness and high excess kurtosis; these will make it more difficult to efficiently measure risk. Furthermore, note that the behavior of the CSE

5. In the graphs displaying histories of indices, monthly volatilities are calculated using the historical estimate $\sigma = \sigma_{25,t} = \sqrt{\frac{1}{24} \sum_{i=0}^{24} (r_{t-i} - \bar{r}_{25,t})^2}$ where $\bar{r}_{25,t} = \frac{1}{25} \sum_{i=0}^{24} r_{t-i}$.

6. Normality was also tested visually using QQ-plots and histograms with fitted normal distributions, but these figures are not shown here.

TABLE 1. Descriptive Statistics of CSE General Index Daily Returns

	Mean (%)	St. Dev. (%)	Skew.	Kurt.	Min (%)	Max (%)	A**	D**	N
All	0.0277	1.9450	1.85	24.41	-10.08	23.68	54.09 (0.0%)	0.014 (<1.0%)	1371
Period 1	0.0627	0.9520	-0.24	9.81	-7.68	5.07	15.95 (0.0%)	0.092 (<1.0%)	801
Period 2	0.1878	3.2670	1.81	11.22	-10.08	23.68	7.737 (0.0%)	0.113 (<1.0%)	306
Period 3	-0.2638	2.1110	-0.38	2.22	-9.17	5.59	1.327 (0.2%)	0.060 (2.8%)	264

Note: * Anderson-Darling normality test: Normality is rejected when the p -value is below 5% ** Kolmogorov-Smirnov normality test: Normality is rejected when the p -value is below 5%.

seems to be significantly different during the different periods. This highlights the importance of including periods of financial turmoil in the testing period of any VaR methodology.

Next, conditional normality of the returns is tested through the estimation of the two parametric models, namely EWMA and GARCH(1,1). These have been estimated for the CSE General Index using the entire history of returns as well as the series of the three periods separately. Parameter estimates are shown in table 2. Beginning with the GARCH(1,1) parameters, note that the volatility is highly persistent (indicated by high β) during period two while in periods one and three persistence is lower, but still relatively high. The long-run volatility is given by:

$$\sigma_{\infty} = \sqrt{\frac{\omega}{1 - \alpha - \beta}}$$

For period two, σ_{∞} is very high, around 4.75%, while the volatilities for all the other periods (one, three and the entire history) imply a much lower long-run volatility, less than 2.5%. This observation signifies the importance of modeling volatility changes for proper risk management in emerging markets.

The decay factor (λ) of the EWMA model was estimated using both the method of maximum likelihood (MLE) and ordinary least squares (OLS). OLS was used in RiskMetrics' specification and it involves no parametric assumptions on the distribution of the residuals. However, observe that the maximum likelihood method is consistent with RiskMetrics' methodology, where VaR is estimated by assuming conditional normality. Comparing the two sets of estimates, observe that the difference between the two is larger the greater the departure from the normality assumption, as should be expected. For example, the returns in period three are relatively close to being normal and the two estimates are correspondingly close. For the entire series, the OLS estimate is very close to 0.94, which is the value that RiskMetrics considers typical for emerging markets. The maximum likelihood estimate converges to a value slightly less than 0.9 and in this way gives much more weight to the realized returns and assumes less persistence in the volatility. The qualitative difference in the estimates of volatility persistence in the GARCH and EWMA models is also significant. For GARCH, persistence is highest in period two, while for EWMA this is when it is lowest.

TABLE 2. GARCH(1,1) and EWMA Estimated Parameters for the CSE General Index

	ω	α	β	$\sigma^2(\%)$	$\lambda(\text{ML})$	$\lambda(\text{OLS})$
All	0.0237 (0.0050)	0.1798 (0.0177)	0.8137 (0.0177)	1.906	0.8816	0.9479
Period 1	0.0493 (0.0117)	0.2434 (0.0412)	0.7066 (0.0405)	0.994	0.9012	0.7596
Period 2	0.0688 (0.0516)	0.1259 (0.0257)	0.8711 (0.0268)	4.752	0.8667	0.9567
Period 3	0.5207 (0.2590)	0.2591 (0.0907)	0.6259 (0.1165)	2.128	0.7879	0.8032

TABLE 3. Descriptive Statistics of EWMA Standardized Returns of the CSE General Index

	Mean (%)	St. Dev. (%)	Skew.	Kurt.	Min (%)	Max (%)	A**	D**	N
All	-0.0106	1.1105	0.392	8.06	-5.06	9.36	6.445 (0.0%)	0.053 (<1.0%)	1371
Period 1	0.0675	1.1009	0.488	9.57	-5.37	9.02	4.750 (0.0%)	0.059 (<1.0%)	801
Period 2	-0.0641	1.1232	0.566	6.19	-3.47	6.07	2.377 (0.0%)	0.071 (<1.0%)	306
Period 3	-0.2152	1.1304	-0.542	4.98	-5.71	2.60	0.461 (25.8%)	0.047 (>15%)	264

Note: * Anderson-Darling normality test: Normality is rejected when the p -value is below 5%. ** Kolmogorov-Smirnov normality test: Normality is rejected when the p -value is below 5%.

TABLE 4. Descriptive Statistics of GARCH(1,1) Standardized Returns of CSE General Index

	Mean (%)	St. Dev. (%)	Skew.	Kurt.	Min (%)	Max (%)	A**	D**	N
All	-0.0146	1.0416	0.261	6.86	-4.86	7.20	5.843 (0.0%)	0.054 (<1.0%)	1371
Period 1	0.0659	0.9992	0.296	6.73	-3.9	7.01	3.858 (0.0%)	0.049 (<1.0%)	801
Period 2	-0.0556	1.0516	0.709	3.98	-3.39	5.97	2.624 (0.0%)	0.068 (<1.0%)	306
Period 3	-0.1583	0.9981	-0.16	3.66	-4.05	2.33	0.331 (52.0%)	0.030 (>15.0%)	264

Note: * Anderson-Darling normality test: Normality is rejected when the p -value is below 5%. ** Kolmogorov-Smirnov normality test: Normality is rejected when the p -value is below 5%.

TABLE 5. Ljung-Box Statistic (lag 30) for CSE General Index

	LB 1*	LB 2**	LB 3***
All	260.9 (0.0%)	455.21 (0.0%)	34.966 (24.33%)
Period 1	104.6 (0.0%)	387.56 (0.0%)	32.774 (33.24%)
Period 2	100.3 (0.0%)	68.698 (0.07%)	31.163 (40.74%)
Period 3	48.98 (1.57%)	80.944 (0.0%)	27.209 (84.11%)

Note: * Standardized Returns: r_t/σ_t . ** Squared Returns: r_t^2 . *** Squared Standardized Returns: r_t^2/σ_t^2 .

The time series of the parameter estimates are shown in figure 6 in appendix A. This figure shows how long it takes for the estimates to converge to their equilibrium values. After one year the decay factor of the EWMA model converges to equilibrium (around 0.9); the GARCH(1,1) parameters converge to their long-run values after three years.

Table 3 shows the descriptive statistics and normality tests of the residuals for the entire history and the three period estimations for the EWMA model. While the residuals of the first and the second period as well as the residuals of the entire history easily reject the assumption of conditional normality, the EWMA model seems to be able to fit the period three data relatively well. During periods of business as usual, this model is probably able to function effectively in capturing the risk. However, it dramatically fails to do so during periods of financial turmoil.

Considering the residuals of the GARCH(1,1) model for the entire history as well as the three periods separately leads to the conclusion that this model is only able to explain the price movements during normal periods, as was the case for the EWMA model. However, due to the more general specification of the GARCH(1,1) model it can cope better than the EWMA model during crisis periods.

Finally, we apply Ljung-Box tests in order to see how much of the autocorrelation is explained by the GARCH(1,1) model. LB1 and LB2 test whether there is any autocorrelation within the standardized and the squared returns respectively, while LB3 tests whether the GARCH model manages to adequately describe the volatility process. From the results of table 5 it follows that the GARCH(1,1) model manages to

describe the volatility of each subperiod. However, the models do not eliminate the autocorrelation within the standardized returns.

A. Statistical Analysis of the Athens Stock Exchange

In this section the statistical analysis is repeated for the case of the more developed Athens Stock Exchange (ASE). The ASE started operating on 3/1/1988 and its history can be separated into four periods, with two bubbles taking place. The first period was characterized by low volatility and a persistence of the index around the initial level, followed by a small positive trend in the level. The first period was followed by a rise of the index (first bubble), reaching the level of 715 (5/7/1990) and then slowly decaying to the level of 400. The following four years were a relatively quiet period when the ASE General Index remained around the level of 400 with rather small volatilities. During the following period another bubble took place and the General Index reached the level of 3,360 (17/9/1999). The complete history of the ASE General Index, together with monthly volatilities, is shown in figure 5. The four periods are defined as: Period One: 3/1/1988 – 31/8/1998; Period Two: 1/9/1989 – 3/1/1993; Period Three: 4/1/1993 – 3/1/1997, Period Four: 4/1/1997 – 18/12/2001.

The ASE returns were analyzed with the same tests as were the CSE General Index returns. The results are presented in appendix B. The descriptive statistics of the daily returns for both the entire history and the periods are shown in table 8. The results show that in no period can the returns of the index be taken to be normally distributed. The returns distribution is peaked and has particularly fat tails. Moreover, the distribution of the returns is not symmetric since the skewness is significantly different than zero. The normality assumption is tested and rejected through the empirical distribution and the QQ-plot (not shown here) as well as the statistics of the empirical distribution, see table 8.

For the EWMA model, maximum likelihood estimation was used and the results are shown in table 9. The daily estimates of the decay parameter tend towards the value $\lambda = 0.94$, which is also the value originally used by RiskMetrics, see figure 7. For the GARCH(1,1) model the estimated parameters are given in table 9. The estimated parameters for the entire history are converging to values similar to those obtained for the CSE ($\omega \approx 0.01$, $\alpha \approx 0.2$, $\beta \approx 0.8$). The long-run volatility for the entire history is relatively high, around 4%. This is notably higher than the CSE long-run volatility, which is around 2%.

GARCH parameters converge to equilibrium much more slowly for Athens (approximately six years) than for Cyprus (approximately three years). For both indices, the estimate of the EWMA decay factor converges in about a year. This observation may be relevant for risk management applications in emerging markets with particularly short histories.

The normality tests for standardized returns indicate that none of the periods has conditionally normal returns, see table 10. The Ljung-Box Statistics imply that the GARCH(1,1) model manages to describe the volatility for every period. On the other hand, the models do not eliminate the autocorrelation within the standardized returns, except for period one, see table 11.

IV. Risk Measurement

This section recalls the definition of Value-at-Risk (VaR), and briefly discusses its properties as a risk measure. As VaR's many shortcomings are notorious, some alternative measures are also presented and discussed.

VaR is generally defined for a portfolio as the maximum possible loss over a given time horizon for a prescribed confidence level. Let $\Pi_{t+\tau}$ denote the value of a portfolio at the horizon $T = t + \tau$, assuming that today is time t . Then the portfolio Value-at-Risk at the horizon T and the confidence level α is given by:

$$\text{VaR}_{t,\tau}(\Pi; \alpha) = \sup\{\xi \mid P[\Pi_{t+\tau} \geq \xi] \leq 1 - \alpha\}.$$

One of the main functions of any financial risk management operation is to determine the amount of capital that the institution must hold in order to be reasonably sure of covering its potential losses. Methods for determining this capital amount are specified by the local regulators, often following the lead of the Basel Committee on Risk Management of the Bank for International Settlements. Institutions following the Basel guidelines may either calculate capital allocation for market risk using standard formulae provided by the committee, or by calculating their Value-at-Risk based on internal models (if local regulators accept the use of internal models, and the bank's models are tested and approved). In a regulatory environment where the use of internal models

is approved (as is the case in all developed markets, but not the case in many emerging markets, due to the factors discussed in the introduction), required regulatory capital is determined as a multiple of the bank's Value-at-Risk. More precisely, the bank is required to hold an amount C given by the equation:

$$C = \max\left(\text{VaR}_{t-1,\tau}(\Pi;\alpha), \kappa \cdot \overline{\text{VaR}}_{t,\tau}(\Pi;\alpha)\right),$$

where Π is the bank's portfolio, $\alpha = 0.99$, τ is ten days, κ is determined by the performance of the bank's models in a back-testing experiment, and:

$$\overline{\text{VaR}}_{t,\tau}(\Pi;\alpha) = \frac{1}{60} \sum_{i=t-60}^{t-1} \text{VaR}_{i,\tau}(\Pi;\alpha).$$

The multiplicative factor κ in the above equation is determined by the number of times that portfolio losses have exceeded the corresponding 99% VaR (using a two week time horizon) over the past two years. It usually takes the value of three, but it can increase if the number of exceptions is greater than five, and can rise up to four if the number of exceptions reaches ten or more during the period. This multiplier can be viewed as an insurance against model risk or imperfect assessment of specific risks. Another view of this multiplier is as a safety factor against non-normal market moves. Regulatory realities, together with the ease of its use and implementation, have made VaR an industry standard for risk measurement.

Even though VaR is the industry standard, it suffers from many serious drawbacks. Of tremendous practical and theoretical importance is the fact that VaR fails to be subadditive; one can define portfolios Π^1 and Π^2 such that $\text{VaR}_{t,\tau}(\Pi^1 + \Pi^2;\alpha) > \text{VaR}_{t,\tau}(\Pi^1;\alpha) + \text{VaR}_{t,\tau}(\Pi^2;\alpha)$. In such a case, VaR fails to reflect the reduction in risk due to diversification (see Acerbi and Tasche [2002], and Artzner et al. [1999]). In the language of Artzner et al. (1999) this means that VaR is not a coherent risk measure. Additionally, (and this is of particular importance in emerging markets), VaR ignores the magnitude of possible losses beyond a given threshold. For example, it is easy to see that a reduction in VaR may simply entail a stretching of a portion of the tail exceeding VaR. Daniélsón (2002) argues that the use of VaR could actually exacerbate financial crises. A more thorough critique of the use of VaR as a risk measure is given by Szegö (2002) (see also the references

therein).

One of the alternatives to VaR is *expected shortfall*, which (assuming that the probability distribution of the portfolio value is continuous) is simply:

$$ES_{t,\tau}(\Pi; \alpha) = E\left[\Pi_t - \Pi_{t+\tau} \mid (\Pi_t - \Pi_{t+\tau}) \geq \text{VaR}_{t,\tau}(r; \alpha)\right].$$

For a definition in the general case (i.e. with possible probability “atoms”), and a proof that this risk measure is coherent in the sense of Artzner et al. (1999), see Acerbi and Tasche (2002). It is important to note that expected shortfall captures the risk of loss where it is most dramatic, i.e. in the tail of the distribution. The entire tail, not only a threshold, determines the risk measure. It is evident that expected shortfall can be estimated easily from an empirical distribution. As a final note, we mention that portfolio selection problems using expected shortfall are generally easier to solve than those using VaR (see Rockafellar and Uryasev [2002]). *Spectral risk measures* resemble expected shortfall, but introduce an auxiliary weight function representing an individual’s risk aversion, see Acerbi (2002) and Tasche (2002) for details.

All the methods for calculating VaR in this paper can be adapted easily to the computation of expected shortfall. For the scenario based methods, the estimation of expected shortfall is straightforward, as it is just the average loss over the scenarios where loss exceeds VaR; for variance-covariance methods (see below) which assume conditional normality, expected shortfall is a simple function of the mean and standard deviation. The motivation for the exclusive use of VaR is twofold. Firstly, the major aim of this study is to perform a comparative analysis, which is easier when considering VaR by assessing the performance of each method with the number of exceedances. Secondly, the method for calculating capital allocation proposed by the Basel Committee is based on VaR, and this measure therefore still must be taken to be the benchmark.

V. VaR Methodologies

This section presents the methods for scenario generation and VaR calculation that are employed in the back-testing in the following

sections. Particular emphasis is paid to the use of data from other emerging markets to enhance risk management in the local environment.

There are numerous methods for computing Value-at-Risk. These include methods based on expansion and normal approximation (such as RiskMetrics' method presented below and its various higher order generalizations, and the Cornish-Fisher expansion), Monte-Carlo methods (exemplified by the other methods described below), and extreme value theory. Many studies have compared the various methods. See for example Allen (1994), Beder (1995), Hendricks (1996), Dave and Stahl (1997), Duffie and Pan (1997), and Manganello and Engle (2001). The methods that are employed in the present study can be divided into the following categories: 1) Variance-Covariance Methods (VCM); 2) Monte-Carlo Simulation Methods (MC); 3) Historical Simulation Methods (HS).

The first two are parametric while the historical simulation methods are non-parametric. The most significant difference between the parametric and non-parametric methods in the context of this paper is that the non-parametric models allow the use of data from external markets. Other stock market indices with longer histories may carry very useful information even though there is no correlation or direct relevance to the local market. The data may be able to reflect stock market characteristics that are very difficult to incorporate into a model, like investors' psychology during 'unusual' periods. This kind of information may be universal and not stock market specific. Consistent with this reasoning is the use of time-series of stock market indices that had the experience of unusual periods. The following stock markets were taken for this study: Greece: ASE Index; Italy: MIB30 Index; Mexico: MEXBOL Index; NASDAQ: NASDAQ 100; Portugal: BVLX Index; Thailand: SET Index.

The above indices and the monthly volatilities of daily returns are plotted in figures 8 (Greece, Italy, Mexico) and 9 (NASDAQ, Portugal, Thailand) in appendix C. The descriptive statistics of the daily returns are shown in table 12 in the appendix, together with the CSE General Index descriptive statistics.⁷ The normality tests are also shown in table 12. Finally, the decay parameter λ of the EWMA model was estimated for all the time series. The results are shown in table 13 (see appendix C).

7. These statistics were not used directly in selecting the external market indices for enriched historical simulation.

The following observations are important when comparing the CSE with the other indices: the CSE General Index has significantly greater skewness (1.85) and excess kurtosis (24.40) than any of the other indices. Even though it has a maximum return (23.68%) that is the highest maximum among all indices, one would expect skewness and excess kurtosis to be closer to the values for the other indices. However, the recent bubble has been a significant part of the CSE history and therefore it has a large effect on the entire history's statistics. This effect should fade out through time.

It is difficult to provide an exact recipe for selecting which external markets to use as data. The markets listed above were chosen using a number of criteria. Each market has undergone periods of relative tranquility and volatility, without the presence of a complete system-wide collapse (the "crash" of the Cyprus Stock Exchange is larger than that of any of the other markets considered). The series are long enough to be useful for sampling techniques. Finally, they are taken from widely varying markets and it is therefore hoped that they span a wide range of possible market regimes. In the historical sampling methods described below, the probabilities assigned to the local and external markets were set arbitrarily. A more advanced version of the technique could attempt to fit these weights based, for example, on a regression of returns, or, more in line with the motivation of the current paper, by selecting the weights that would have performed optimally in the back-test up to the current date. Improved methods for selecting the data series and the corresponding weights are important topics for future study.

An important point is that in the enriched historical simulation framework, the theoretical relationship that may exist between the local market and a particular foreign market is inconsequential. For example, in the case of the Athenian and Cypriot markets, there may exist some predictive power in the lag returns of the Athens Stock Exchange index in forecasting the Cyprus Stock Exchange index. However, if the previous day's return is used in our EHS, it is just one more return in the pool that is used to generate the historical scenarios and carries the same weight as any other return from any other market. Therefore, problems like cointegration should not affect the results of EHS.

The Monte-Carlo and Historical simulation methods are scenario based. For this reason, for each pair (α, τ) of confidence level and time horizon, respectively, two VaR estimates will be considered. These are *nVaR*: Non-parametric VaR with $\alpha\%$ confidence, i.e. the lower $1 - \alpha\%$

percentile of the empirical distribution for the scenario set; *pVaR*: Parametric VaR with $\alpha\%$ confidence in which the normal distribution is fitted to the empirical distribution of the returns of the scenario set. The VaR estimate is then the lower $1 - \alpha\%$ percentile of the fitted distribution (computed using the inverse of the cumulative normal distribution function, and the scenario set mean and standard deviation). Two confidence levels (95% and 99%) will be considered, along with three time horizons (one week, two weeks and one month). Therefore, each scenario-based method generates in total twelve VaR estimates (six for each confidence level), while the variance-covariance methods will give six VaR estimates (three for each confidence level, corresponding to the three time horizons).

A. Variance-Covariance (VCV) Methods

This approach to VaR estimation is based on the standard methodology pioneered at J.P. Morgan. The RiskMetrics model assumes that the distribution of the returns is normal with mean zero and variances σ_t^2 . It follows that $\Pi_{t+\tau} + \Pi_t$ is a mean zero normal random variable with single-period variance $\sqrt{\tau}\sigma_t^2$ and the equation:

$$P[\Pi_t - \Pi_{t+\tau} \geq \text{VaR}_{t,\tau}(\Pi; \alpha)] = 1 - \alpha,$$

implies that: $\text{VaR}_{t,\tau}(\Pi; \alpha) = z_{[\alpha]} \sigma_t \sqrt{\tau}$.

For instance, for a ten day horizon at $\alpha = 0.95$ confidence level, $\text{VaR}_{t,10}(\Pi; \alpha) = 1.65\sigma_t \sqrt{10}$. The volatility (σ_t) is based on the EWMA model of daily returns and is estimated each day using the newly available information. This is referred to as the “RiskMetrics” method or “VCV-EWMA”. Alternatively, the volatility estimate can be made using the GARCH(1,1) model. This is referred to as the “VCV-GARCH(1,1)” method.

B. Monte-Carlo Simulation Methods

The methods described above are based on fitting models for *daily* returns. VaR estimates for longer horizons are generated by scaling the

volatility accordingly. This is something of an approximation. In order to assess the errors and determine whether the results can be improved, the EWMA and the GARCH(1,1) models are also used in a Monte-Carlo context. Once again, the model for returns (assuming $\mu = 0$) is:

$$r_t = \sigma_t \varepsilon_t,$$

where ε_t are *i.i.d.* with a standard normal distribution. The volatility parameter can be estimated using either the EWMA or the GARCH(1,1) model. These are referred to in the tables as “MC-EWMA” and “MC-GARCH(1,1)” respectively. The parameters used for each simulation are estimated using all the available local information. The difference between this method and the previous one lies in the concurrent simulation of the volatility and the return. The future volatility is estimated through the endogenous parameters σ_t and r_t , and the exogenous parameter ε_t , in contrast to the previous method where the future volatility is constant until the time horizon.⁸

C. Historical Simulation Methods

The historical simulation method is widely used for VaR calculation due to its simplicity and the fact that no particular parametric assumption is made regarding the distribution of the returns. Even though the original method suffers from drawbacks such as theoretical inconsistency and under-responsiveness to conditional risk, it enjoys wide popularity. The enriched historical simulation method tries to cope with the fact that the history of an emerging market does not carry enough information for effective risk management.

The original Historical Simulation (HS) method estimates VaR using a window of fixed length T of recent historical returns.⁹ The window length T is chosen to strike a balance between relevance (returns from the distant past may not be as relevant as the previous day’s return) and accuracy (many returns are required for an accurate estimate). Once T is chosen, the VaR of the next period return is simply estimated as the quantile of the empirical distribution $\{r_{t-T+1}, \dots, r_t\}$.

This method implicitly assumes that the returns are identically and

8. The initial value of σ_t^2 is $\sigma_0^2 = \omega/1 - \alpha - \beta$, which is the expected volatility.

9. Returns can be of any horizon, i.e. daily, weekly, etc.

independently distributed (*i.i.d.*) and for this reason the weights of all the T returns are the same in the empirical distribution. The selection of any window of returns T is inconsistent with the *i.i.d.* assumption in the sense that in estimating the VaR for the return r_{t+1} , the return r_{t-T+1} holds the same weight as r_t while r_{t-T} has no weight at all.¹⁰

Bootstrapping Historical Data

The historical methods used in this paper apply the statistical technique of bootstrapping to the problem of estimating VaR. In the language of Hall (1992) the (uncorrected) bootstrap estimator of VaR is simply the corresponding quantile of the empirical distribution of returns, described in the previous section. Results for this estimator are presented in the rows in the tables labeled “HS”.

The simple bootstrap estimate can be “improved” by correcting its bias based on a resampling technique. The data are repeatedly resampled (with replacement) and quantiles of the generated samples are computed. The difference between the average of these “bootstrapped” quantiles and the quantile of the empirical distribution gives an estimate of the bias of the quantile estimator. The new estimate for VaR subtracts this bias from the empirical quantile.¹¹ This technique is applied by generating 1000 samples of size 500 from the original returns series and estimating the quantile of each sample. For each time horizon, samples are generated in two different ways. In the first method (labeled “BHS” in the tables), weekly, biweekly, and monthly returns are generated by sampling five, ten and twenty-one daily returns respectively from the original series in an independent manner. This is consistent with the assumption that daily returns are *i.i.d.* The second method (labeled “BBHS”, for “block bootstrap” in the tables) generates samples based on the series of weekly, biweekly and monthly returns that actually occurred (this is consistent with the existence of autocorrelation in the daily returns).

10. A more theoretically appealing method, introduced by Boudoukh, Richardson and Whitelaw (1998), avoids this inconsistency by applying exponentially decaying weights to the returns. Thus, the Weighted Historical Simulation Method essentially eschews the *i.i.d.* assumption and imposes autocorrelations on the returns. It tries to build on the stylized fact of volatility clustering; the decay factor is usually above 0.97.

11. It should be noted that this method converges slowly, and other methods based on first smoothing the distribution have better performance. This additional sophistication would be unlikely to change the ordinal rankings of the methods in this paper.

D. Enriched Historical Simulation Method

The Enriched Historical Simulation Method (EHS) is based on the original bootstrapping HS method. It is an attempt to enrich the information set using data from other markets. This method is non-parametric, and assumes that all returns are independently and identically distributed. For this reason the entire historical set is used, instead of a window of returns T , thus avoiding the inconsistency with other historical methods discussed earlier.

In order to specify a non-parametric method, the set of possible returns Ω_t and the probabilities P_t must be specified. In the enriched historical simulation method, the set Ω_t contains all the returns previously observed on the local index as well as prior returns from a selected set of other indices. The external indices used in this study were described above. As already discussed, the probability weights for the external series in this study were set arbitrarily.¹²

Enriched Historical simulation 1: Sampling from CSE and other Indices

With this historical simulation method the scenario sets are generated by sampling from the CSE historical returns distribution and the distributions of returns of other indices. Arbitrarily, a weight of 82% is assigned to the local market (CSE) information set and 3% to each of the other six foreign market information sets.¹³ For a particular date t , historical scenario sets were generated by sampling from the following periods: CSE: 29/3/1996 – t ; ASE: 2/3/1998 – t ; Italy: 3/11/1997 – t ; Mexico: 1/3/1995 – 31/12/1998; Nasdaq: 1/4/1998 – t ; Portugal: 1/4/1987 – 1/3/1989; Thailand: 1/10/1991 – 30/10/1998. In the tables, results for this method are contained in the rows labeled “EHS1”.¹⁴

Enriched historical simulation 2: Sampling from CSE and proxy scenarios

12. More sophisticated models, based on fitting these probabilities, will be the focus of future work.

13. In this case of the ASE being the local market, the CSE history is neglected. The weights are 82.5% to ASE and 3.5% to each of the other 5 remaining foreign market information sets.

14. The dates for the Mexican, Portuguese and Thai exchanges were selected as the beginning and end of local bubbles.

This variant of the enriched historical simulation method makes use of the local information set in the same way as in the previous historical simulation method and uses the foreign information set by creating what-if scenarios. These what-if scenarios use the actual evolutions of the other indices as proxies for possible future evolutions of the local market index. More precisely, a weight of 94% is given to the historical scenarios that are generated by sampling from the local market returns and 1% to each of the six what-if scenarios.¹⁵ The periods that were used for these historical scenarios are the same as the periods of the previous historical simulation method. The results for this method are reported in the rows of the tables labeled “EHS2”.

VI. Empirical Study

This section discusses results of tests of the VaR methodologies described in the previous section. The methods are assessed using the back-testing algorithm recommended by the Basel Committee of the Bank for International Settlements (2001).

The back-test compares VaR measures for one week, two week and one month changes of the portfolio at a 95% or 99% confidence level, against the actual portfolio loss for the corresponding time period. Assuming that the risk factors are correctly modeled and that markets behave accordingly, we expect, on average, that the absolute value of the actual loss will exceed the 99% VaR only five days out of the last five hundred.

A VaR estimate is generated using the methods described above for a particular time horizon and confidence level. For each method, an overall score is defined as follows. On a test date t , the violation score is given by:

$$V_{p,\tau,m,t} = \begin{cases} 1 & \text{if the actual loss is less than the VaR figure,} \\ 0 & \text{otherwise} \end{cases}$$

where p is the one-sided confidence level, t is the time-horizon and m is the method. If the null hypothesis (H_0) is that the methodology is

15. The probability weights are lower because EHS2 makes stronger assumptions, adopting the autocorrelations as well as the returns of the external market. In the case of the ASE being the local market, the weights are 95% for the ASE and 1% for each of the other five remaining foreign market information sets.

correct, then for every $V_{p,\tau,m,t}$

$$H_0 = \left\{ \begin{array}{l} 1 \text{ with probability } p \\ 0 \text{ with probability } 1-p \end{array} \right\},$$

for all p , τ , m , and $t = 1, \dots, T$. The overall score, henceforth called the *violation ratio*, for the entire testing period $(1, \dots, T)$ is defined as:¹⁶

$$V_{p,\tau,m} = \frac{1}{T} \sum_{t=1}^T V_{p,\tau,m,t}.$$

Assuming that the violation scores are independent of each other and that T is large, the Central Limit Theorem implies that under the null hypothesis $V_{p,\tau,m}$ is approximately normally distributed with mean p and variance $\frac{p(1-p)}{n}$. Therefore the null hypothesis for a particular method m , confidence level p , and time horizon τ is defined as:¹⁷

$$H_0 : V_{p,\tau,m} \sim N \left(p, \sqrt{\frac{p(1-p)}{T}} \right).$$

In order to compare VaR figures one should consider the adjusted violation ratios, defined as:

$$\tilde{V}_{p,\tau,m} = \frac{1}{p} V_{p,\tau,m},$$

with the corresponding null hypothesis:

$$H_0^* : \tilde{V}_{p,\tau,m} \sim N \left(1, \sqrt{\frac{(1-p)}{pT}} \right).$$

16. The term violation ratio is somewhat misleading. We have a violation when the actual loss is great than the VaR. However, the violation ratio is greater the greater the number of non-violations that we have, since non-violations should occur exactly 100- α % of the time.

17. Normally distributed random variables are presented as $V \sim N(\mu, \sigma)$.

The above distributions can be used to test the null hypothesis by checking whether the estimated violation ratio $V_{p,\tau,m}$ is within the 95% confidence interval specified by the distribution.

A. Cyprus Stock Exchange VaR Test Results

The testing period for the CSE General Index begins before the bubble, on 9/2/1999 and lasts until the end of the times series, 26/10/2001. Scenario sets using all simulation methods were generated for each of the dates of the testing period. The VaR figures generated were compared with the actual losses of the CSE General Index for the dates of the scenario sets. However, for some dates the comparisons were not made. These were the dates on which the forecasting period (one month ahead) had many days for which there was no trading. More precisely, no testing was made for a date on which there were more than two days of no trading in the one week ahead, or five days of no trading in the next two weeks ahead, or ten days of no trading in the one month ahead. From the 646 trading days in the testing period, twenty days were lost because of these restrictions. It should be noted here that, in principle, it does not matter how many trading days there are within a particular time-horizon as what matters is the flow of information, which drives the prices. However, the above constraints were applied to avoid any thin-trading effects.

The results are presented in table 6. The adjusted violation ratios for each method and time horizon are shown. The left part of the table is for the 95% confidence level, while the right part is for the 99% confidence level. The methods almost uniformly underestimate the risk of the index, as could be expected given the chaotic nature of the local market.

The null hypothesis as defined above states that the 95% VaR adjusted violation ratios for a sample of 626 are normally distributed:

$$\tilde{V}_{p,\tau,m} \sim N(1, 0.0087),$$

resulting in a 95% symmetric confidence interval [0.9820,1.0180]. For the 99% VaR violation ratios, the distribution is:

$$\tilde{V}_{p,\tau,m} \sim N(1, 0.0040).$$

A 95% symmetric confidence interval in this case is [0.9921,1.0078].

One way to test each method is to compare its adjusted violation ratio against the appropriate confidence interval. Only BHS-pVaR at the 95% confidence level falls within the confidence interval specified by the null hypothesis.

Examining the results in the table more closely, we see that Enriched Historical Simulation outperforms straightforward historical simulation on short time horizons and for higher confidence levels. EHS1 outperforms all other nVaR methods on the one week and two week time horizons at the 99% confidence level. For the longer time horizons, the quality of enriched historical simulation deteriorates, and historical simulation using only the local data is competitive. The MC methods using EWMA and GARCH(1,1) outperform their VCV counterparts since the VCV methods do not take into consideration that the volatility changes with time. The nVaR adjusted violation ratios of the Monte-Carlo methods are close to the pVaR ones. This is natural as these simulation methods assume a conditionally normal distribution for the returns.

One striking feature of the results is that for the historical methods the pVaR estimates uniformly outperform their nVaR counterparts. This is because they assume a symmetric distribution for the returns (not due to normality). As is evident from the statistical analyses of section 3, the CSE General Index returns are far from normally distributed. The scenario sets generated by the historical methods were positively skewed during the period when the index was going up and it took these methods a while after the turning point to adjust in producing more symmetric distributions. This is particularly noticeable with the bootstrap historical simulation methods that use only local information (i.e. HS, BHS, and BBHS). It is quite reasonable to assume that the downside potential of the returns is of the same magnitude as the upside potential; the parametric method fills in the downside of the distribution, even when no downside has yet been observed. Readers should be cautioned against generalizing these observations to other markets, as they derive partially from the distinctive “single bubble” nature of the history of the CSE. In general, the pVaR approach entails the danger of seriously mis-estimating the risk and therefore undermining the profitability of the institution. Some evidence

TABLE 6. Adjusted VaR Test Results for CSE General Index

Method	95% VaR			99% VaR		
	1-week	2-week	1-month	1-week	2-week	1-month
HS-n VaR	0.8643	0.8492	0.8139	0.9472	0.9149	0.8633
HS-p VaR	0.9265	0.9366	0.9198	0.9698	0.9714	0.9843
BHS-n VaR	0.8794	0.8593	0.8660	0.9278	0.8972	0.8826
BHS-p VaR	0.9568	0.9854	1.0123	0.9778	0.9907	1.0101
BBHS-n VaR	0.8609	0.8307	0.7685	0.9326	0.8778	0.8003
BBHS-p VaR	0.9282	0.9383	0.9232	0.9698	0.9714	0.9907
EHS1-n VaR	0.8979	0.8508	0.7735	0.9746	0.9552	0.8778
EHS1-p VaR	0.9534	0.9265	0.8424	0.9778	0.9714	0.9423
EHS2-n VaR	0.8593	0.8408	0.7516	0.9375	0.9068	0.8245
EHS2-p VaR	0.9383	0.9080	0.8340	0.9762	0.9601	0.9294
VCV-GARCH(1,1)	0.8897	0.8612	0.8102	0.9159	0.8869	0.8651
MC-GARCH(1,1)-n VaR	0.9047	0.8576	0.7836	0.9633	0.9439	0.9230
MC-GARCH(1,1)-p VaR	0.9164	0.8677	0.8189	0.9536	0.9310	0.9101
VCV-EWMA(RiskMetrics)	0.8744	0.8273	0.7970	0.8972	0.8729	0.8487
MC-EWMA-n VaR	0.8996	0.8626	0.8054	0.9310	0.8859	0.8584
MC-EWMA-p VaR	0.9080	0.8593	0.8088	0.9262	0.8988	0.8633

TABLE 7. Adjusted VaR Test Results for ASE General Index

Method	95% VaR			99% VaR		
	1-week	2-week	1-month	1-week	2-week	1-month
HS-n VaR	0.9551	0.9529	0.9366	0.9903	0.9758	0.9789
HS-p VaR	0.9833	0.9865	0.9908	0.9831	0.9799	0.9955
BHS-n VaR	0.9496	0.9399	0.9171	0.9955	0.9737	0.9622
BHS-p VaR	0.9952	1.0006	1.0136	0.9862	0.9883	1.0018
BBHS-n VaR	0.9529	0.9496	0.9193	0.9893	0.9695	0.9456
BBHS-p VaR	0.9843	0.9876	0.9908	0.9831	0.9799	0.9955
EHS1-n VaR	0.9648	0.9713	0.9681	0.9935	0.9883	0.9883
EHS1-p VaR	0.9908	0.9908	0.9898	0.9883	0.9883	0.9966
EHS2-n VaR	0.9438	0.9481	0.9492	0.9735	0.9725	0.9725
EHS2-p VaR	0.9775	0.9786	0.9862	0.9798	0.9798	0.9892
VCV-GARCH(1,1)	0.9525	0.9503	0.9623	0.9641	0.9662	0.9725
MC-GARCH(1,1)-n VaR	1.0461	1.0472	1.0515	1.0091	1.0101	1.0101
MC-GARCH(1,1)-p VaR	1.0472	1.0505	1.0526	1.0101	1.0101	1.0101
VCV-EWMA(RiskMetrics)	0.9833	0.9811	0.9887	0.9841	0.9810	0.9862
MC-EWMA-n VaR	0.9666	0.9764	0.9699	0.9819	0.9725	0.9735
MC-EWMA-p VaR	0.9753	0.9732	0.9721	0.9788	0.9756	0.9767

supporting this view is given by the results of the back-test for the ASE.

B. Athens Stock Exchange Test Results

The testing for the ASE General Index starts before the second bubble, on 2/2/1998 and ends on 18/12/2001. In this testing period, 971 trading days were included. Under the null hypothesis as defined above, the 95% VaR adjusted violation ratios for a sample of 971 are normally distributed:

$$\tilde{V}_{p,\tau,m} \sim N(1, 0.0070),$$

resulting in the 95% symmetric confidence interval [0.9856, 1.0143]. For the 99% VaR violation ratios the distribution is:

$$\tilde{V}_{p,\tau,m} \sim N(1, 0.0032),$$

and the 95% symmetric confidence interval is [0.9936, 1.0064]. Table 7 gives the adjusted violation ratios for each method and time horizon.

The results in this case are more mixed. Almost all methods perform better on this more advanced market. It is interesting to note that now some methods overestimate the risk (indicated by adjusted violation ratios greater than one). In this case the EHS1-pVaR method is competitive even with BHS1-pVaR; both methods produce estimates within the confidence interval for the 95% confidence level. At this confidence level, the enriched historical simulation EHS1-nVaR method outperforms all other nVaR methods. Once again, RiskMetrics performs poorly compared to the other methods.

Conclusions for some methods in the Athens results are difficult. For example, observe that the non-parametric Monte-Carlo GARCH does quite poorly at the 95% confidence level, but produces estimates satisfying the confidence interval test at the 99% confidence level. The general conclusions that may be drawn from this analysis are as follows. For the Cyprus Stock Exchange, enriched historical simulation generally improves risk assessment (although not by enough to achieve the number of exceptions that would guarantee the lowest capital multiplier defined by BIS). For the Athens Stock Exchange, EHS methods satisfy the BIS back-testing criteria, while methods that use local data exclusively fail to pass the test. Finally, in all cases, the use of data from other markets provides substantially more accurate risk assessment than

the standard RiskMetrics methodology.

VII. Conclusion

The challenges facing risk managers in emerging markets are both numerous and complex. This paper addresses two of the challenges, the estimation of the industry standard risk measure Value-at-Risk and the lack of data for scenario generation. The paper proposes the use of data from other emerging markets in order to supplement these calculations. The methodology is illustrated on both the Cyprus and Athens stock exchanges. Numerical tests reveal that in both cases methods using data from other markets improve risk measurement in comparison to techniques using only data from the local market. It is posited that these results are due to the fact that data from other markets carry information that is common to all emerging markets and that the inclusion of data from many markets means that the local idiosyncrasies of these markets will have little effect.

There are many avenues for future research and possible improvement of the results presented herein. There is a plethora of different distributions and mathematical models that have been developed in advanced markets and should be tested before being applied in the local market (e.g. Lévy processes, stochastic volatility, jump-diffusions, and bootstrapping historical residuals as in Barone-Adesi et al. [1998,2000]). For enriched historical simulation in particular, an important direction of future study is to develop algorithms for selecting which historical series should be included in a simulation, and with what probabilities. One possibility would be to choose the weights at each time to be those that would have performed optimally in the past.

Enterprise risk management for institutions operating in emerging markets is still in its infancy. While there exist many well-developed techniques in advanced markets, these techniques do not translate effortlessly into the emerging market context. However, if the innovation, creativity and persistence that have been applied to risk management in developed markets are employed in the emerging market context, the associated difficulties should not prove insurmountable.

References

- Acerbi, C. 2002. Spectral measures of risk: A coherent representation of subjective risk aversion. *Journal of Banking and Finance* 27:1505–1518.
- Acerbi, C., and Tasche, D. 2002. On the coherence of expected shortfall. *Journal of Banking and Finance* 27:1487–1503.
- Allen, M. 1994. Building a Role Model. *Risk* 7 (July): 73–80.
- Artzner, P.; Delbaen, F.; Eber, J.-M.; and Heath, D. 1999. Coherent measures of risk. *Mathematical Finance* 9:203–228.
- Barone-Adesi, G.; Giannopoulos, K.; and Vosper, L. 2000. Backtesting Derivative Portfolios with FHS. Working paper.
- Barone-Adesi, G.; Giannopoulos, K.; and Vosper, L. 1998. VaR without correlations for portfolios of derivative securities. Working paper.
- Barone-Adesi, G.; Giannopoulos, K.; and Vosper, L. 2000. Filtering historical simulation. Backtest analysis. Working paper, 2000.
- Basel Committee on Banking Supervision. 1996. Overview of the amendment to the capital accord to incorporate market risks. Basel Committee Publications no. 23, Basel: Bank for International Settlements.
- Basel Committee on Banking Supervision. 1997. Amendment of the capital accord to incorporate market risks. Basel Committee Publications no. 24, Basel: Bank for International Settlements.
- Basel Committee on Banking Supervision. 2001. The new Basel capital accord. Consultative document. Basel: Bank for International Settlements.
- Beder, T. 1995. VaR: Seductive but dangerous. *Financial Analysts Journal* 51:12–24.
- Boudoukh, J.; Richardson, M.; and Whitelaw, R. 1998. The best of both words. *Risk* 11 (November): 64–67.
- Crouhy, M.; Galai, D.; and Mark, R. 2001. *Risk Management*. New York: McGraw-Hill.
- Daniélsson, J. 2002. The emperor has no clothes: Limits to risk modeling. *Journal of Banking and Finance*. 27:1273–1296.
- Dave, R., and Stahl, G. 1997. On the accuracy of VaR estimates based on the variance-covariance approach. Working paper.
- Dembo, R.; Aziz, A.; Rosen, D.; and Zerbs, M. 2000. Mark to future: A framework for measuring risk and reward. Toronto: Algorithmics Inc.
- Duffie, D., and Pan, J. 1997. An overview of Value-at-Risk. *Journal of Derivatives* 4:7–49.
- Gentle, J. 2002. *Elements of Computational Statistics*. New York: Springer.
- Hall, P. 1992. *The Bootstrap and Edgeworth Expansion*. New York: Springer.
- Hendricks, D. 1996. *Evaluation of Value-at-Risk models using historical data*. Federal Reserve Bank of New York Policy Review, New York: Federal Reserve Bank.
- Jorion, P. 2000. *Value at Risk: The New Science of Risk Management*. New York: John Wiley & Sons.

- MacKay, C. 1995. *Extraordinary Popular Delusions and the Madness of Crowds*. New York: Crown Publishers.
- Manganelli, S., and Engle, R.F. 2001. Value at Risk Models in Finance. Working paper 75, European Central Bank.
- Nerouppos, M.; Saunders, D.; Xiouros, C.; and Zenios, S.A. 2002. The risks of the Cyprus and Athens stock exchanges. HERMES European Center of Excellence on Computational Finance and Economics Working Paper 02–05. Nicosia: University of Cyprus.
- Pritsker, M. 2002. The hidden dangers of historical simulation. Federal Reserve Board Working Paper 27, New York: Federal Reserve Bank.
- Rockafellar, R.T., and Uryasev, S. 2002. Conditional Value-at-Risk for general loss distributions. *Journal of Banking and Finance*. 27:1443–1471.
- Shiller, R. 2000. *Irrational Exuberance*. Princeton, N.J.: Princeton University Press.
- Szegö, G. 2002. Measures of risk. *Journal of Banking and Finance*. 27:1253–1272.
- Tasche, D. 2002. Expected shortfall and beyond. *Journal of Banking and Finance*. 27:1519–1533.

APPENDIX A. Additional Figures and Tables for the Cyprus Stock Exchange

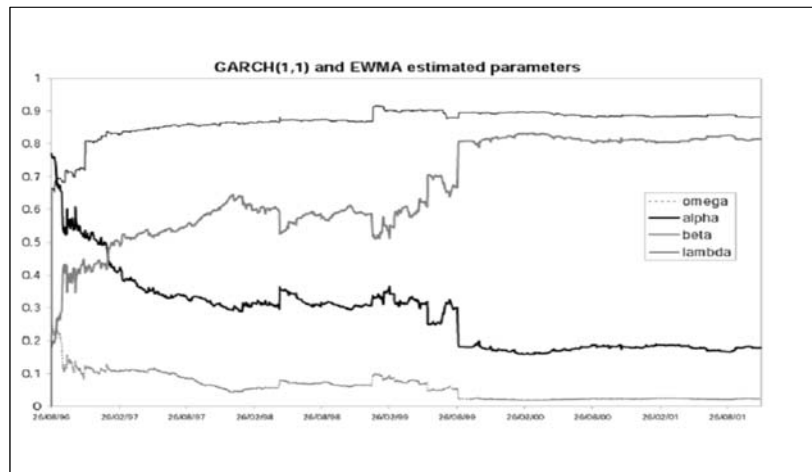


FIGURE 6. — GARCH(1,1) and EWMA Parameter Estimates for the CSE Daily Returns

APPENDIX B. Additional Figures and Tables for the Athens Stock Exchange

TABLE 8. Descriptive Statistics of ASE General Index Daily Returns

	Mean (%)	St. Dev. (%)	Skew.	Kurt.	Min (%)	Max (%)	A**	D**	N
All	0.0749	1.908	0.019	6.39	-14.63	15.31	59.2 (0.0%)	1.084 (<1.0%)	3484
Period 1	0.0587	1.567	-2.22	22.5	-14.63	6.093	14.365 (0.0%)	0.127 (<1.0%)	409
Period 2	0.1177	2.431	0.367	5.37	-12.42	15.31	12.111 (0.0%)	0.077 (<1.0%)	817
Period 3	0.0501	1.194	0.103	4.52	-6.764	5.866	11.331 (0.0%)	0.074 (<1.0%)	1000

Note: * Anderson-Darling normality test: Normality is rejected when the p -value is below 5%. ** Kolmogorov-Smirnov normality test: Normality is rejected when the p -value is below 5%.

TABLE 9. GARCH(1,1) and EWMA Estimated Parameters for the ASE General Index

	ω	α	β	$\sigma^2(\%)$	$\lambda(\text{ML})$	$\lambda(\text{OLS})$
All	0.0107 (0.0050)	0.1873 (0.0177)	0.8063 (0.0173)	4.086	0.9371	0.9361
Period 1	0.6049 (0.1727)	0.2082 (0.0707)	0.5808 (0.1012)	1.692	0.9368	0.9373
Period 2	0.0382 (0.0149)	0.0603 (0.0100)	0.9368 (0.0106)	3.720	0.9725	0.9700
Period 3	0.1038 (0.2590)	0.2181 (0.0907)	0.7115 (0.1165)	1.215	0.8574	0.8502

TABLE 10. Descriptive Statistics of GARCH(1,1) Standardized Returns of ASE General Index

	Mean (%)	St. Dev. (%)	Skew.	Kurt.	Min (%)	Max (%)	A ² *	D**	N
All	0.0306	1.004	-0.405	11.93	-12.65	8.08	26.87 (0.0%)	0.040 (<1.0%)	3484
Period 1	0.0492	0.995	-3.935	39.85	-10.83	2.89	13.100 (0.0%)	0.052 (<1.0%)	409
Period 2	0.0378	0.991	0.366	6.85	-6.53	6.99	7.644 (0.0%)	0.127 (<1.0%)	817
Period 3	0.0455	1.002	0.393	1.06	-3.33	4.26	3.111 (0.0%)	0.070 (<1.0%)	1000

Note: * Anderson-Darling normality test: Normality is rejected when the *p*-value is below 5%. ** Kolmogorov-Smirnov normality test: Normality is rejected when the *p*-value is below 5%.

TABLE 11. Ljung-Box Statistic (lag 30) for ASE General Index

	LB 1*	LB 2**	LB 3***
All	22.97 (81.0%)	3.134 (100.0%)	3.036 (100.0%)
Period 1	55.36 (<1.0%)	176.6 (0.0%)	14.904 (99.0%)
Period 2	62.52 (<1.0%)	330.1 (0.0%)	31.072 (41.0%)
Period 3	58.36 (<1.0%)	152.73 (0.0%)	24.471 (75.0%)

Note: * Standardized Returns: r_t/σ_t . ** Squared Returns: r_t^2 . *** Squared Standardized Returns: r_t^2/σ_t^2 .




FIGURE 7. — GARCH(1,1) and EWMA Parameter Estimates for the ASE Daily Returns

APPENDIX C. Additional Figures and Tables for other Exchanges




FIGURE 8. — Level and Volatilities of General Indices of Greece, Italy and Mexico




FIGURE 9. — Level and Volatilities of the General Indices of Nasdaq, Portugal and Thailand

TABLE 12. Descriptive Statistics of Daily Returns of all Indices

Index	Mean (%)	St. Dev. (%)	Skew.	Kurt.	Min (%)	Max (%)	A ² *	D**	N
CSE	0.0277	1.945	1.85	24.4	-10.08	23.68	54.09 (0.0%)	0.014 (<1.0%)	1371
ASE	0.0749	1.908	0.019	6.39	-14.63	15.31	59.21 (0.0%)	0.084 (<1.0%)	3484
Italy	0.0433	1.426	-0.255	4.05	-9.843	9.185	48.79 (0.0%)	0.058 (<1.0%)	7296
Mexico	0.1197	1.808	0.03	5.96	-14.31	12.15	36.05 (0.0%)	0.069 (<1.0%)	3509
Nasdaq	0.0683	1.932	0.068	5.43	-10.38	17.20	40.44 (0.0%)	0.072 (<1.0%)	3260
Portugal	0.0649	1.287	-0.26	14.9	-13.53	10.75	134.7 (0.0%)	0.122 (<1.0%)	3622
Thailand	0.0088	1.898	-0.058	5.87	-17.39	11.35	45.82 (0.0%)	0.068 (<1.0%)	3585

Note: * Anderson-Darling normality test: Normality is rejected when the *p*-value is below 5%. ** Kolmogorov-Smirnov normality test: Normality is rejected when the *p*-value is below 5%.

TABLE 13. EWMA lambda Estimates for all Indices

Index	Maximum Likelihood			OLS
	$\mu_t = 0$	$\mu_t = \bar{\mu}$	$\mu_t = \bar{\mu}_{10,t}$	$\mu_t = 0$
Italy	0.9316	0.9314	0.9141	0.8916
Mexico	0.9500	0.9509	0.9254	0.9252
Nasdaq	0.9566	0.9563	0.9354	0.9352
Portugal	0.9050	0.9138	0.8712	0.7965
Thailand	0.9040	0.9041	0.8686	0.9071