

**The Emperor has no Clothes:**

**Limits to Risk Modelling**

**By**

**Jon Danielsson**

**SPECIAL PAPER 126**

**October 2000**

---

**FINANCIAL MARKETS GROUP**  
AN ESRC RESEARCH CENTRE

---

**LONDON SCHOOL OF ECONOMICS**



**Any opinions expressed are those of the author and not necessarily those of the Financial Markets Group.**

ISSN 1359-9151-126

# The Emperor has no Clothes: Limits to Risk Modelling

Jón Daniélsson\*  
Financial Markets Group  
London School of Economics  
*www.RiskResearch.org*

June 2000

## Abstract

This paper considers the properties of risk measures, primarily Value-at-Risk (VaR), from both internal and external (regulatory) points of view. It is argued that since market data is endogenous to market behavior, statistical analysis made in times of stability does not provide much guidance in times of crisis. In an extensive survey across data classes and risk models, the empirical properties of current risk forecasting models are found to be lacking in robustness while being excessively volatile. For regulatory use, the VaR measure is lacking in the ability to fulfill its intended task, it gives misleading information about risk, and in some cases may actually increase both idiosyncratic and systemic risk. Finally, it is hypothesized that risk modelling is not an appropriate foundation for regulatory design, and alternative mechanisms are discussed.

---

\*I have benefited from comments by Richard Payne, Casper de Vries, Charles Goodhart, Kevin James, and Bob Nobay; the mistakes and analysis is of course my responsibility alone. For correspondence [j.danielsson@lse.ac.uk](mailto:j.danielsson@lse.ac.uk). My papers can be downloaded from [www.RiskResearch.org](http://www.RiskResearch.org) or [fmq.lse.ac.uk/~jond](http://fmq.lse.ac.uk/~jond).

# 1 Introduction

Recent years have witnessed an enormous growth in financial risk modelling both for regulatory and internal risk management purposes. There are many reasons for this; stronger perception of the importance of risk management, deregulation enabling more risk taking, and technical advances encouraging both risk taking and facilitating the estimation and forecasting of risk. This impact has been felt in regulatory design: market risk regulations are now model based, and we see increasing clamoring for regulatory credit risk modelling. The motivations for market risk modelling are obvious. The data is widely available, a large number of models analyzing financial data exist, rapid advances in computer technology enable the estimation of the most complicated models, and the ever-increasing supply of well educated graduates, all led to a sense of “can do” within the technical modelling community. Empirical modelling has been of enormous use in applications such as derivatives pricing and risk management, and is being applied successfully to the new fields of credit and operational risk.

There is however precious little evidence that risk modelling actually works. Has risk management delivered? We don’t know for sure, and probably never will. If regulatory risk management fulfills its objectives, we will not observe any systemic failures, so only the absence of crisis can prove the system works. There is however an increasing body of evidence that inherent limitations in risk modelling technology, coupled with imperfect regulatory design, is more like a placebo rather than the scientifically proven preventer of crashes it is sometimes made out to be. Below I survey some of this evidence; the general inaccuracy and limitations of current risk models, the impact of (externally imposed uniform) risk constraints on firm behavior, the relevance of statistical risk measures, and the feedback between market data, risk models, and the beliefs of market participants. I finally relate this evidence to the current debate on regulatory design where I argue against the notion of model based regulations, be it for market, credit, or operational risk.

An explicit assumption in most risk models is that market price data follows a stochastic process which only depends on past observations of itself and other market data, not on outside information. While this assumption is made to facilitate modelling<sup>1</sup>, it relies on the hypothesis that there are so many market participants, and they are so different that in the aggregate their actions are essentially random and can not influence the market. This implies that the role of the risk forecaster is akin to a meteorologists job, who can forecast the weather, but not influence it. This approach to modelling has a number of shortcomings from the point of view of financial markets. If risk measurements influence people’s behavior, it is inappropriate to assume market prices follow an independent stochastic process. This becomes especially relevant in times

---

<sup>1</sup>Incorporating outside information in statistical risk models is very hard

of crisis when market participants hedge related risks leading to the execution of similar trading strategies. The basic statistical properties of market data are not the same in crisis as they are during stable periods; therefore, most risk models provide very little guidance during crisis periods. In other words, the risk properties of market data change with observation. If, in addition, identical external regulatory risk constraints are imposed, regulatory demands may perversely lead to the amplification of the crisis by reducing liquidity. There is some evidence that this happened during the 1998 crisis. Indeed, the past 3 years have not only been the most volatile in the 2nd half of the 20th century but also the era of intensive risk modelling.

In order to forecast risk, it is necessary to assume a model which in turn is estimated with market price data. This requires a number of assumptions regarding both model design and the statistical properties of the data. It is not possible to create a perfect risk model, and the risk forecaster needs to weigh the pros and cons of the various models and data choices to create what inevitably can only be an imperfect model. I present results from an extensive survey of model risk forecast properties, employing a representative cross-section of data and models across various estimation horizons and risk levels. The results are less than encouraging. All the models have serious problems with lack of robustness and high risk volatility, implying that unless a risk model is chosen with considerable skill and care, the model outcomes will be as accurate as predictions of the outcomes of a roulette wheel. Off-the-shelf models can not be recommended.

Current market risk regulations are based on the 99% Value-at-Risk (VaR) measure obtained from a risk model. The VaR number can, under some assumptions, provide an adequate representation of risk, however such assumptions are often unrealistic, and result in the misrepresentation of risk. There are (at least) four problems with the regulatory VaR measure. First, it does not indicate potential losses, and as a result is flawed, even on its own terms. Second, it is not a coherent measure of risk. Third, its dependence on a single quantile of the profit and loss distribution implies it is easy to manipulate reported VaR with specially crafted trading strategies. Finally, it is only concerned with the 99% loss level, or a loss which happens 2.5 times a year, implying that VaR violations have very little relevance to the probability of bankruptcy, financial crashes, or systemic failures.

The role of risk modelling in regulatory design is hotly debated. I argue that the inherent flaws in risk modelling imply that neither model based risk regulations nor the risk weighing of capital can be recommended. If financial regulations are deemed necessary, alternative approaches such as *state contingent* capital levels and/or *cross-insurance* systems are needed.

## 2 Risk Modelling and Endogenous Response

The fundamental assumption in most statistical risk modelling is that the *basic* statistical properties of financial data during stable periods remain (almost) the same as during crisis. The functional form of risk models is usually not updated frequently, and model parameters get updated slowly.<sup>2</sup> This implies that risk models can not work well in crisis. The presumed inability of risk models to work as advertised has not gone unnoticed by commentators and the popular press;

“Financial firms employed the best and brightest geeks to quantify and diversify their risks. But they have all – commercial banks, investment banks and hedge funds – been mauled by the financial crisis. Now they and the worlds regulators are trying to find out what went wrong and to stop it happening again. ... The boss of one big firm calls super-sophisticated risk managers ‘high-IQ morons’ ”

The Economist, November 18, 1998, pp. 140–145.

There is a grain of truth in this quote: Statistical financial models do break down in crisis. This happens because the statistical properties of data during crisis is different than the statistical properties in stable times. Hence, a model created in normal times may not be of much guidance in times of crisis. Morris and Shin (1999) suggest that most statistical risk modelling is based on a fundamental misunderstanding of the properties of risk. They suggest that (most) risk modelling is based on the incorrect assumption of a single person (the risk manager) solving decision problems with a natural process (risk). The risk manager in essence treats financial risk like the weather, where the risk manager assumes a role akin to a meteorologist. We can forecasts the weather but can not change it, hence risk management is like a “game against nature”. Fundamental to this is the assumption that markets are affected by a very large number of heterogeneous market participants, where in the aggregate their actions become a randomized process, and no individual market participant can move the markets. This is a relatively innocuous assumption during stable periods, or in all periods if risk modelling is not in widespread use. However, the statistical process of risk is different from the statistical process of the weather in one important sense: forecasting the weather does not (yet) change the statistical properties of the weather, but forecasting risk does change the nature of risk. In fact, this is related to Goodharts Law:

**Law 1 (Goodhart (1974))** *Any statistical relationship will break down when used for policy purposes.*

---

<sup>2</sup>Models, such as GARCH, do of course pick up volatility shocks, but if the parameterization and estimation horizon remain the same, the basic stochastic process is the same, and the model has the same steady state volatility, which is only updated very slowly.

We can state a corollary to this

**Corollary 1** *A risk model breaks down when used for its intended purpose.*

Current risk modelling practices are similar to pre-rational expectations Keynesian economics in that risk is modelled with behavioral equations that are invariant under observation. However, just as the economic crisis of the 1970's illustrated the folly of the old style Keynesian models, so have events in financial history demonstrated the limitations of risk models. Two examples serve to illustrate this. The Russia crisis of 1998 and the stock market crash of 1987.

Consider events during the 1998 Russia crisis. (See e.g. Dunbar (1999)). At the time risk had been modelled with relatively stable financial data. In Figure 3 on page 26 we see that the world had been in a low volatility state for the preceding half a decade, volatility had somewhat picked up during the Asian crisis of 1997, but those volatility shocks were mostly confined to the Far East, and were leveling off in any case. In mid year 1998 most financial institutions employed similar risk model techniques and often similar risk constraints because of regulatory considerations. When the crisis hit, volatility for some assets went from 16 to 40, causing a breach in many risk limits. The response was decidedly one-sided, with a general flight from volatile to stable assets. This of course amplified price movements and led to a sharp decrease in liquidity. In other words, the presence of VaR based risk limits led to the execution of similar trading strategies, escalating the crisis.

This is indeed similar to events surrounding a previous crisis, the 1987 crash when a method called portfolio insurance was very much in vogue. (See e.g. Jacobs (1999)). An integral component in portfolio insurance is that complicated hedging strategies with futures contracts are used to dynamically replicate options in order to contain downside risk. These dynamic trading strategies worked well in the stable pre-crisis periods since they depended on the presence of functioning futures markets. However, one characteristic of the '87 crash was that the futures markets ceased to function properly because the institutions who used portfolio insurance were trying to execute identical trading strategies, which only served to escalate the crisis.

If every financial institution has its own trading strategy, no individual technique can lead to liquidity crisis. However, each institution's behavior does move the market, implying that the distribution of profit and loss is endogenous to the banks decision-making process. In other words, risk is not the separate stochastic variable assumed by most risk models, instead, risk modelling affects the distribution of risk. A risk model is a model of the aggregate actions of market participants, and if many of these market participants need to execute the same trading strategies during crisis, they will change the distributional properties of risk. As a result, the distribution of risk is different during crisis than in other periods, and risk modelling is not only useless but may exasperate the crisis, by leading to large price swings and lack of liquidity.

The role played by regulation during these scenarios is complex. It is rational for most banks to reduce exposure in the event of a risk shock, independent of any regulatory requirements, and if banks have similar incentives and employ similar risk models, that alone can lead to a snowball effect in trading strategies during crisis. Indeed this is what happened during the 1987 crisis when regulation played no direct role. However, if regulation restricts the banks scope for pursuing individually optimal strategies, causing banks to act in a more uniform manner during crisis it may lead to an escalation of the crisis. Whether this actually happens is yet unknown. In the current regulatory environment, the banks themselves model risk, hence risk models are individual to the bank and as a results not all institutions follow the same risk control trading strategy. Hence, regulation may not be the straight-jacket one might think. However, the burden of proof is on the regulators to demonstrate that the regulations do not escalate the crisis. They have not yet done so.

The analysis presented here does not imply that risk modelling is inherently pointless. Internally, banks do benefit from hedging as demonstrated by e.g. Froot, Scharfstein, and Stein (1993), and risk models are reliable in stable periods, especially in dealing with idiosyncratic shocks. However, even though (most) current risk models are useless during crisis, this is only a reflection of the current state of risk technology. Risk models grounded in the microeconomic theory of financial crisis, have considerable promise.

### 3 Empirical Properties of Risk Models

Risk forecasting depends on a statistical model and historical market price data. The modeller makes a number of assumptions about the statistical properties of the data, and from that specifies the actual model. This model will always be based on objective observations and subjective opinions, and therefore the quality of the model depends crucially on the modeller's skill. As a result, it is not possible to create a perfect model. Each model has flaws, where the modeller weighs the pros and cons of each technique and data set, juggling issues like the choice of the actual econometric model, the length of the estimation horizon, the forecast horizon, and the significance level of the forecasts. In fact, due to these limitations, the resulting model is endogenous to its intended use. Two different users, who have different preferences but identical positions and views of what constitutes risk, require different risk forecasts. This happens because risk modelling is conditional on the user's loss functions. The weighing of the pros and cons is different for different users, resulting in different risk models and hence different risk forecasts, even for identical positions. In order to address some of these issues I refer below to results from a survey I made (Danielsson 2000) of the forecast properties of various models and data sets.

The data is a sample from the major asset classes, equities, bonds, foreign ex-

change, and commodities. Each data set consists of daily observations and spans at least fifteen years. The discussion below is based on forecasts of day-by-day Value-at-Risk during the 1990s (2500 forecasts per dataset on average) The risk level is mostly the regulatory 99% but I also consider lower and higher risk levels. The survey was done with a single asset return. While a portfolio approach would be more appropriate, this raises a number of issues which I thought best avoided. As such, my survey only presents a best case scenario, most favorable to risk models.

A large number of risk models exist and it is not possible to examine each and every one. However, most models are closely related to each other, and by using a carefully selected subsample I am confident that I cover the basic properties of most models in use. The models studied are:

- Conditional volatility models (normal<sup>3</sup> and student-t GARCH)
- Unconditional models (historical simulation (HS) and extreme value theory (EVT))

(More details on the survey can be found in Appendix A.)

The risk modeller is faced with many challenges, some of the most important are:

- Robustness of risk forecasts (Section 3.1)
- Volatility of risk forecasts (Section 3.2 on page 9)
- Determination of the appropriate measuring horizon (Section 3.3 on page 10)
- Determination of the holding period (Section 3.4 on page 11)
- Underestimation of downside risk due to asymmetries in correlation structures (Section 3.5 on page 13)

### 3.1 Robustness of Risk Forecasts

For a risk model to be considered reliable, it should provide accurate risk forecasts across different assets, time horizons, and risk levels within the same asset class. The robustness of risk models has been extensively documented, and there is not much reason to report detailed analysis here, my results correspond broadly to those from prior studies. I use violation ratios<sup>4</sup> to measure the accuracy of risk forecasts. If the violation ratio is larger than 1, the model is underforecasting risk (it is thin tailed relative to the data), and if the violation ratio is lower than 1, the model is overforecasting risk (the model is thick tailed relative to the

---

<sup>3</sup>The RiskMetrics™ model is a restricted normal GARCH model (IGARCH) and hence its performance can only be worse than for the normal GARCH.

<sup>4</sup>The realized number of VaR violations over expected number of violations. By violation I mean that realized loss was larger than the VaR forecast



data). Violation ratios are the most common method for ranking models, since they directly address the issue of forecast accuracy. The risk level used is the regulatory 99%, see Table 1 on page 24.

An ideal model has violation ratios close to 1 across asset classes and significance levels. While what constitutes “close to 1” is subjective, the range of 0.8 to 1.2 is a useful compromise. Based on this criteria, the results are depressing. For example, the normal GARCH model produces violation ratios ranging from 0.37 to 2.18, and even for the same data set, e.g. S&P-500, the violation ratios range from 0.91 to 1.46. The other estimation methods have similar problems but not on the same scale. Every method overestimates the bond risk, and underestimates the risk in Microsoft stock. The normal GARCH model (and by extension RiskMetrics<sup>TM</sup>) has the overall worst performance, the results for the other models are mixed. Danielsson (2000) reports that a similar picture emerges from all considered risk levels, but the ranking among models changes. At the 95% risk level, the normal GARCH model performs generally best, while at 99.9% EVT is best.

It is also interesting to consider the importance of the estimation horizon. Conventional wisdom seems to be that short horizons are preferred, say one year or 250 days. This is indeed a regulatory recommendation in some cases. It is however not supported by my results. For example, for oil prices and the Student-t GARCH model, the violation ratio is 1.38 when the model is estimated with 300 days, and only 1.04 when the model is estimated with 2,000 days. Similar results have been obtained in some other cases. The reason has to do with the fact that a conditional volatility reverts to the steady state volatility, which is dependent on the estimation horizon. If the estimation horizon is too short, the model steady state volatility reflects past high/low volatility states which are less relevant than the long run average.

These results show that no model is a clear winner. The forecasts, cover a very wide range, and the lack of robustness is disconcerting. Furthermore, the estimation horizon has considerable impact on the forecast accuracy. One conclusion is that none of these models can be recommended, but since these models form the basis of almost every other model, this recommendation is too strong. My approach here is to use off-the-shelf models. A skilled risk manager considering specific situations is able to specify much more accurate models. This is indeed the situation internally in many banks. For reporting purposes, where the VaR number is an aggregate of all the banks risky positions, the use of an accurate specially designed model is much harder and off-the-shelf models are more likely to be used. This coupled with ad hoc aggregation methods for risk across operations, and the lack of coherence in VaR, can only have an adverse effect on model accuracy.

## 3.2 Risk Volatility

Fluctuations in risk forecasts have serious implications for the usefulness of a risk model; however, risk forecast fluctuations have not been well documented. The reason for this is unclear, but the importance of this issue is real. If a financial institution has a choice between two risk models both of which forecast equally well, but one model produces much less volatile forecasts, it will be chosen. And if risk forecasts are judged to be excessively volatile, it may hinder the use of risk forecasting within a bank. If a VaR number routinely changes by a factor of 50% from one day to the next, and factor 2 changes are occasionally realized, it may be hard to sell risk modelling within the firm. Traders are likely to be unhappy with widely fluctuating risk limits, and management does not like to change market risk capital levels too often. This is due to many reasons, one of them phantom price volatility. Furthermore Andersen and Bollerslev (1998) argue that there is an built-in upper limit on the quality of volatility forecasts (around 47%).

In my survey I use two measures of fluctuations in risk forecasts;

- The volatility of the VaR, i.e. the standard deviation of VaR forecasts over the sample period
- The VaR forecast range, i.e. the maximum and minimum VaR forecast over the sample period

Both measures are necessary. The VaR volatility addresses the issue of day-to-day fluctuations in risk limits, while the VaR range demonstrates the worst-case scenarios.

A representative sample of the results using the S&P-500 index is presented below in Table 2 on page 25. Consider e.g. the regulatory 99% level and the 300 day estimation horizon. The return volatility is 0.9%, and the volatility of the VaR estimates is 0.7%. It is almost like we need a risk model to access the risk in the risk forecasts! The largest drop in returns is -7.1% (in 2527 observations), while the lowest normal GARCH model forecast is -7.5% at the 99% level or an event once every 100 days. With longer estimation horizons both the volatility and the range decrease, suggesting that longer estimation horizons are preferred. The same results are obtained from the other data sets. Another interesting result is that the least volatile methods are historical simulation (HS) and extreme value theory (EVT). The reason is that conditional volatility models are based on a combination of long estimation horizons (more than 250 days) along with very short run VaR updating horizons (perhaps five days). In contrast, the HS and EVT methods are unconditional and as a result produce less volatile risk forecasts. Note that a hybrid conditional volatility and EVT method, such as the methods proposed by McNeil and Frey (1999) produce VaR forecasts which are necessarily more volatile than the condition volatility methods.

The contrast between GARCH and EVT volatility can be seen in Figure 4 on page 27 which shows Hang Seng index returns during the last quarter of 1997. Set in the middle of the Asian crisis, the Hang Seng index is very volatile, with the largest one day drop of more than 15%. Both models have an excessive amount of violations, but while the EVT forecast is relatively stable throughout the quarter, the GARCH forecast is very volatile. The lowest GARCH VaR is 19%, and the model takes more than a month to stabilize after the main crash. In addition, the main contributor to the GARCH VaR volatility is the positive return of 18% following the main crash. Since conditional volatility models like GARCH have a symmetric response to market movements, a positive and negative market movement has the same impact on the VaR.

Because the Value-at-Risk numbers are quantiles of the profit and loss distribution it is not surprising that they are volatile. However, I find them surprisingly volatile. It is not uncommon for VaR numbers to double from one day to the next, and then revert back. If VaR limits were strictly adhered to, the costs of portfolio rebalancing would be large. This has not gone unnoticed by the financial industry and regulators. Anecdotal evidence indicates that many firms employ ad hoc procedures to smooth risk forecasts. For example, a bank might only update its covariance matrix every three months, or treat risk forecasts from conditional volatility models as an ad hoc upper limit for daily risk limits. Alternatively, covariance matrices are sometimes smoothed over time using non-optimal procedures. If Value-at-Risk is used to set risk limits for a trading desk, strict adherence to a VaR limit which changes by a factor of two from one day to the next is indeed costly. The same applies to portfolio managers who need to follow their mandate, but would rather not rebalance their portfolios too often. In addition, since regulatory VaR is used to determine market risk capital, a volatile VaR leads to costly fluctuations in capital if the financial institution keeps its capital at the minimum level predicted by the model. This may turn cause lack of confidence in risk models and hinder their adoption within a firm. Anecdotal evidence indicates that (some) regulators consider bank capital as a constant to be allocated to the three categories of risk, market, credit, and operational, and not the widely fluctuating quantity predicted by the models.

### 3.3 Model Estimation Horizon

The estimation of a risk model depends on sufficiently long historical data series being available. The regulatory suggestion is (at least) 250 days, and anecdotal evidence indicates that short estimation horizons are very much preferred. This must be based on one of two assumptions;

- Older data is not available, or is irrelevant due to structural breaks
- Long run risk dynamics are so complicated that they can't be modelled

While the first assumption is true in special cases, e.g. immediately after a new instrument is introduced such as the Euro, and in emerging markets, in general, it is not correct. The second assumption is partially correct: long run risk dynamics are complicated and often impossible to model explicitly; however, long run patterns can be incorporated, it just depends on the model.

Long run risk dynamics are not a well understood and documented phenomena, but it is easy to demonstrate the existence of long cycles in volatility. Consider Figure 3 on page 26 which demonstrates changes in average daily volatility for the second half of the 20th century. Daily volatility ranges from 0.5% to almost 2% in a span of few years. The 1990s demonstrate the well-known U-shaped pattern in volatility.

Observing these patterns in volatility is one thing, modelling them is another. Although existing risk models do not yet incorporate this type of volatility dynamics, conceivably this is possible. Most conditional volatility models, e.g. GARCH, incorporate both long run dynamics (through parameter estimates) and very short-term dynamics (perhaps less than one week). Long memory volatility models may provide the answers; however, their risk forecasting properties are still largely unexplored.

The empirical results presented in Table 1 on page 24 and Table 2 on page 25 show that shorter estimation horizons do not appear to contribute to more accurate forecasting, but longer estimation horizons do lead to lower risk volatility. For that reason alone longer estimation horizons are preferred.

### 3.4 Holding Periods and Loss Horizons

Regulatory Value-at-Risk requires the reporting of VaR for a 10 day holding period. This is motivated by a fear of liquidity crisis where a financial institution might not be able to liquidate its holdings for 10 days straight. While this may be theoretically relevant, two practical issues arise;

- The contradiction in requiring the reporting of a 10 day 99% Value-at-Risk, i.e. a two week event which happens 25 times per decade, in order to catch a potential loss due to a liquidity crisis which is unlikely to happen even once a decade. Hence the probability and problem are mismatched.
- There are only two different methods of doing 10 day Value-at-Risk in practice:
  - Use non-overlapping<sup>5</sup> 10 day returns to produce the 10 day Value-at-Risk forecast
  - Use a scaling law to convert one day VaRs to 10 day VaRs (recommended by the Basel Committee on Banking Supervision (1996))

---

<sup>5</sup>Overlapping returns cannot be used for obvious reasons

Both of these methods are problematic

If 10 day returns are used to produce the Value-at-Risk number, the data requirements obviously also increase by a factor of 10. For example, if 250 days (one year) are used to produce a daily Value-at-Risk number, ten years of data are required to produce a 10 day Value-at-Risk number with the same statistical accuracy. If however 250 days are used to produce 10 day Value-at-Risk numbers, as sometimes is recommended, only 25 observations are available for the calculation of something which happens in one observation out of every hundred, clearly an impossible task. Indeed, at least 3,000 days are required to directly estimate a 99% 10 day VaR, without using a scaling law.

To bypass this problem most users tend to follow the recommendation in the Basel regulations (Basel Committee on Banking Supervision 1996) and use the so-called ‘*square-root-of-time*’ rule, where a one day VaR number is multiplied by the square root of 10 to get a 10 day VaR number. However, this depends on surprisingly strong distribution assumptions, i.e. that returns are normally iid:

- Returns are normally distributed
- Volatility is independent over time
- The volatility is identical across all time periods

Needless to say, all three assumptions are violated. However, creating a scaling law which incorporates violations of these assumptions is not trivial. (For a technical discussion on volatility and risk time scaling rules see appendix B.) For example, it is almost impossible to scale a one day VaR produced by a normal GARCH model to a 10 day VaR (see Drost and Nijman (1993) or Christoffersen and Diebold (2000)). Using *square-root-of-time* in conjunction with conditional volatility models implies an almost total lack of understanding of statistical risk modelling. The problem of time scaling for a single security, e.g. for option pricing, is much easier than the scaling of an institution wide VaR number, which currently is impossible. When I have asked risk managers why they use the *square-root-of-time* rule, they reply that they do understand the issues (these problems have been widely documented), but they are required to do this anyway because of the demand for 10 day VaRs for regulatory purposes. In other words, regulatory demands require the risk manager to do the impossible! I have encountered risk managers who use proper scaling laws for individual assets, but then usually in the context of option pricing, where the pricing of path dependent options depends crucially on using the correct scaling method and accurate pricing has considerable value. The pricing of path dependent options with fat tailed data is discussed in e.g. Caserta, Daniélsson, and de Vries (1998).

In fact, one can make a plausible case for the *square-root-of-time* rule to be twice as high, or alternatively half the magnitude of the real scaling factor. In other words, if a daily VaR is one million, a 10 day VaR equal to 1.5 million or 6 million is as plausible. Indeed, given current technology, it is not possible

to come up with a reliable scaling rule, except in special cases. The market risk capital regulatory multiplier of 3 is sometimes justified by the uncertainty in the scaling laws, e.g. by Stahl (1997), however as suggested by Danielsson, Hartmann, and de Vries (1998), it is arbitrary.

Estimating VaR for shorter holding horizons (intra day VaR, e.g. one hour) is also very challenging due to intraday seasonal patterns in trading volume, as frequently documented, e.g. by Danielsson and Payne (2000). Any intraday VaR model needs to incorporate these intraday patterns explicitly for the forecast to be reliable, a non-trivial task.

Considering the difficulty given current technology of creating reliable 10 day VaR forecasts, regulatory market risk capital should not be based on the 10 day horizon. If there is a need to measure liquidity risk, other techniques than VaR need to be employed. If the regulators demand the impossible, it can only lead to a lack of faith in the regulatory process.

### 3.5 Asymmetry and Correlations

One aspect of risk modelling which does not seem to get much attention is the serious issue of changing correlations. It is well known that measured correlations are much lower when the markets are increasing in value, compared to market conditions when some assets increase and other decrease, and especially when the markets are falling. Indeed, the worse market conditions are, the higher the correlation is: in a market crash, all assets collapse in value and correlations are close to hundred percent. However, most risk models do not take this into consideration. For example, conditional volatility models, e.g. GARCH or Risk-Metrics, produce correlation estimates based on normal market conditions, hence these models will tend to underestimate portfolio risk. Furthermore, since the correlations increase with higher risk levels, a conditional variance-covariance volatility model which performs well at the 95% level, will not perform as well at the 99%

This problem is bypassed in methods which depend on historical portfolio returns, since they preserve the correlation structure. An example of methods which use historical portfolio returns are historical simulation, extreme value theory, and Student-t GARCH<sup>6</sup>

One study which demonstrates this is by Erb, Harvey, and Viskanta (1994) who consider monthly correlations in a wide cross-section of assets in three different market conditions (bull markets, bear markets, mixed). They rank data according to the market conditions, and report correlations for each subsample. A

---

<sup>6</sup>A non-ambiguous representation of the multivariate Student-t does not exist, hence in practice a Student-t conditional volatility model can only be used with a single asset, not a portfolio. This applies to most other non-normal distribution as well. In addition, the normal GARCH model can be used in this manner as well.

small selection of their results is reported below:

Asset Pair		Up-Up	Down-Down	Mixed	Total
US	Germany	8.6	52	-61	35
	Japan	21	41	-54	26
	UK	32	58	-60	50
Germany	Japan	4.6	24	-47	40
	UK	22	40	-62	42
Japan	UK	12	21	-54	37

We see that correlations are low when both markets are increasing in value, for example for the U.S. and Germany the correlation is only 8.6%. When both of these markets are dropping in value, the correlation increases to 52%. Similar results have been obtained by many other authors using a variety of data samples.

These problems are caused because of the non-normal nature of financial return data. The only way to measure tail correlations is by using bi-variate extreme value theory where under strict assumptions it is possible to measure tail correlations across probability levels, see e.g. Longin (1998) or Hartmann, Straetmans, and de Vries (2000). However, this research is still in an early stage.

Another problem in correlation analysis relates to international linkages. Since not all markets are open at the exact same time, volatility spillovers may be spread over many days. This happened during the 1987 crisis where the main crash day was spread over two days in Europe and the Far East. A naïve analysis would indicate that the US markets experienced larger shocks than other markets, however this is only an artifact of the data. This also implies that it is very hard to measure correlations across timezones. Any cross-country analysis is complicated by market opening hours, and is an additional layer of complexity.

## 4 The Concept of (Regulatory) Risk

A textbook definition of risk is volatility<sup>7</sup>, however, volatility is a highly misleading concept of risk. It depends on the notion of returns being normal iid<sup>8</sup>, but since they are not, volatility only gives a partial picture. Consider Figure 5 on page 27 which shows 500 realizations of two different return processes,  $A$  which is normally distributed and  $B$  which is not normal. For the purpose of risk, returns  $B$  are clearly more risky, for example the regulatory 99% Value-at-Risk for asset  $A$  is 2, while the VaR for  $B$  is 7. However, the volatility of  $A$  is 1, while the volatility of  $B$  is 0.7. If volatility was used to choose the less risky asset,

<sup>7</sup>The standard deviation of returns.

<sup>8</sup>See Section 3.4 on page 12 for more on iid normality

the choice would be  $B$ , but if Value-at-Risk was used correctly<sup>9</sup> to make the choice it would be  $A$ . This demonstrates the advantages of using a distribution independent measure, like VaR, for risk.

Value-at-Risk (VaR) is a fundamental component of the current regulatory environment<sup>10</sup>, and financial institutions in most countries are now expected to report VaR to their supervisory authorities. (See the Basel Committee on Banking Supervision (1996) for more information on regulatory VaR.) Value-at-Risk as a theoretical definition of risk has some advantages, primarily;

- Ease of exposition<sup>11</sup>
- Distributional independence

There are some disadvantages in the Value-at-Risk concept as well;

- It is only one point on the distribution of profit and loss
- It is easy to manipulate, leading to moral hazard, hence potentially increasing risk, while reporting lower risk

We address each of these issues in turn.

## 4.1 Lower Tail and Alternative Risk Measures

Regulatory Value-at-Risk measures one point on the profit and loss (P/L) distribution of a portfolio of a financial institution, i.e. the 1% lower quantile. This is demonstrated in Figure 1 on the next page which shows the cumulative distribution function (CDF) for returns.

Value-at-Risk does not take into account the entire lower tail of the profit and loss (P/L) distribution, only the quantile. However this may not be a relevant measure in many cases. What matters is how much money a bank loses when a disaster strikes, not the minimum amount of money it loses on a bad day. If the VaR is 1 million, one has no way of knowing whether the maximum possible loss is 1.1 million or 100 million. While users may implicitly map the VaR number into a more useful measure, perhaps relying on dubious distributional assumptions, this can not be recommended. If a different measure is needed, it should be modelled explicitly.

---

<sup>9</sup>If a volatility model (e.g. GARCH) was used to produce the VaR, it would incorrectly choose  $B$


<sup>10</sup>Mathematically, regulatory VaR is defined as:

$$0.01 = \text{Probability}[Loss^{10\text{Day}} \geq VaR].$$

<sup>11</sup>It is however surprising how many authors confuse Value-at-Risk with other risk measures, c.f. expected shortfall (see Section 4.1)



Figure 1: Value-at-Risk and the CDF of Profit and Loss



In addition, Artzner, Delbaen, Eber, and Heath (1999) note that VaR is not a coherent measure of risk because it fails to be subadditive.<sup>12</sup> They propose to use instead the expected shortfall measure which measures the expected loss conditional on reaching the VaR level. A related measure is the first lower partial moment which attempts to map the entire lower tail into one number.<sup>13</sup>

However, this need not to be a serious criticism. As discussed by Cumperayot, Danielsson, Jorgensen, and de Vries (2000) and Danielsson, Jorgensen, and de Vries (2000), these three risk measures provide the same ranking of risky projects under second order stochastic dominance, implying that Value-at-Risk is a sufficient measure of risk. This however only happens sufficiently far the tails. Whether the regulatory 99% is sufficiently far out, remains an open question.<sup>14</sup>

<sup>12</sup>A function  $f$  is subadditive if  $f(x_1 + \dots + x_N) \leq f(x_1) + \dots + f(x_N)$ .

<sup>13</sup>The formal definition of expected shortfall is:


$$\int_{-\infty}^t x \frac{f(x)}{F(t)} dx.$$

First lower partial moments:

$$\int_{-\infty}^t (t - x) f(x) dx = \int_{-\infty}^t F(x) dx.$$

<sup>14</sup>Answering this question ought not to be difficult, it only requires a comprehensive empirical study, as the theoretical tools do exist.

Figure 2: Impact of options on the C.D.F. of Profit and Loss



## 4.2 Moral Hazard

The reliance on a single quantile of the P/L distribution as in VaR is conducive to the manipulation of reported risk. Consider Figure 1 on the preceding page from the perspective of a bank which is faced with externally imposed regulations targeting the 99% point on the curve. In addition, suppose the bank considers the regulatory VaR unacceptably high. Ahn, Boudukh, Richardson, and Whitelaw (1999) consider the optimal response by the bank to this by the use of options. Daniélsson, Jorgensen, and de Vries (2000) consider a richer example where the bank sets up trading strategies to manipulate the VaR. Assume that the VaR before any response is  $VaR_0$  and that the bank really would like the VaR to be  $VaR^D$  where the desired VaR is  $VaR^D > VaR_0$ . One way to achieve this is to write a call with a strike price right below  $VaR_0$  and buy a put with a strike right above  $VaR^D$ , i.e.  $X^c = VaR_0 - \epsilon$  and  $X^p = VaR^D + \epsilon$ . The effect of this will be to lower expected profit and increase downside risk, see Figure 2. This is possible because the regulatory control is only on a single quantile, and the bank is perfectly within its rights to execute such a trading strategy. A measure like expected shortfall or first lower partial moment render this impossible. The regulatory focus on a simple measure like VaR may thus perversely increase risk and lower profit, while the intention is probably the opposite.

This example is very stylistic since it assumes that a bank knows the tail distribution, but as argued above, it is unlikely to do so. However, the example demonstrates how trading strategies can be used to mis-represent risk, and it is not hard to create a trading strategy which works with uncertain tail distributions.

### 4.3 The Regulatory 99% Risk Level

The regulatory risk level is 99%. In other words, we expect to realize a violation of the Value-at-Risk model once every hundred days, or 2.5 times a year on average. Some banks report even lower risk levels, JP Morgan (the creator of VaR and RiskMetrics) states in its annual report that in 1996, its average daily 95% VaR was \$36 million. Two questions immediately spring to mind: why was the 99% level chosen, and how relevant is it?

The first question may have an easy answer. Most models only have desirable properties in a certain probability range and the risk level and risk horizon govern the choice of model. For example, at the 95% risk level, conditional normal models such as normal GARCH or RiskMetrics are the best choice. However, the accuracy of these models diminishes rapidly with lower risk levels, and at the 99% risk level they cannot be recommended, and other, harder to use techniques must be employed. In general, the higher the risk level, the harder it is to forecast risk. So perhaps the 99% level was chosen because it was felt that more extreme risk levels were too difficult to model?

The question of the relevance of a 99% VaR is harder to answer because it depends on the underlying motivations of the regulators. The general perception seems to be that market risk capital is required in order to prevent systemic failures. However, systemic failures are very rare events, indeed so rare that one has never been observed in modern economies. We have observed near-systemic collapses, e.g. in the Scandinavian banking crisis, but in that case even a meticulously measured VaR would not have been of much help.

The regulatory risk level is clearly mismatched with the event it is supposed to be relevant for, i.e. systemic collapse. In other words, the fact that a financial institution violates its VaR says nothing about the probability of the firm's probability of bankruptcy, indeed, one expects the violation to happen 2.5 times a year. There is no obvious mapping from the regulatory risk level to systemic risk, however defined. Whether there is no link between regulatory risk and systemic risk is still an open question.

## 5 Implications for Regulatory Design

The arguments voiced above suggest that modelling as a regulatory tool cannot be recommended. It does not imply anything about the need to regulate. Bank regulation is a contentious issue which is beyond the scope of this paper. Assuming that regulations are here to stay, the important question must be whether it is possible to create a regulatory mechanism that is successful in reducing systemic risk, but not too costly.

For most of the 1990s the answer to this question seemed to be risk modelling,

and it is only after the Asian and Russian crisis that modelling as a regulatory tool has come under serious criticism. Risk modelling is simply too unreliable, it is too hard to define what constitutes a risk, and the moral hazard issues are too complicated for risk modelling to be an integral part of regulatory design, whether for market, credit, or operational risk. This is a reflection of the current state of technology.

My analysis focuses on market risk models. Market risk modelling is much easier than credit risk modelling due to the abundance of accurate market risk data and more established methodology, compared to the lack of reliable credit risk data. My criticism applies equally to credit and operational risk models, hence the case for model based credit and operational regulations is even weaker than for market risk.

If the authorities want banks to hold minimum capital, crude capital adequacy ratios are the only feasible way. Risk weighing of capital will not work for the same reason as regulatory risk modelling does not work. However, it is unrealistic to expect banks to meet minimum capital ratios in times of crisis, therefore, such capital would have to be *state contingent*, i.e., a bank would be allowed to run capital down during crisis. The question what constitutes a crisis, of course still remains to be answered.

Another possible way is to follow the lead of New Zealand and do away with minimum capital, but require banks instead to purchase insurance, in effect require financial institutions to cross insure each other. This market solution has the advantage that much more flexibility is built into the system while at the same time sifting the burden of risk modelling back to the private sector. While such a system may work for small country like New Zealand which can insure in larger markets, it is still an open question whether this would work for larger economies.

## 6 Conclusion

Empirical risk modelling forms the basis of the market risk regulatory environment as well as internal risk control. Market risk regulations are based on using the lower 1% quantile of the distribution of profit and loss ( $VaR^{99\%}$ ) as the statistic reported as risk. This paper identifies a number of shortcomings with regulatory Value-at-Risk (VaR), where both theoretic and empirical aspects of VaR are analyzed.

I argue that most existing risk models break down in times of crisis because the stochastic process of market prices is endogenous to the actions of market participants. If the risk process becomes the target of risk control, it changes its dynamics, and hence risk forecasting becomes unreliable. This is especially prevalent in times of crisis, such as events surrounding the Russia default of 1998.

In practice, Value-at-Risk (VaR) is forecasted using an empirical model in conjunction with historical market data. However, current risk modelling technology still in the early stages of development, is shown in the paper to be lacking in the robustness of risk forecasts, and to produce excessively volatile risk forecasts. If risk modelling is not done with great skill and care, the risk forecast will be unreliable to the point of being useless. Or even worse, it may impose significant but unnecessary costs on the financial institution, due to the misallocation of capital and excessive portfolio rebalancing.

This, however, is only a reflection on the current state of technology. A risk model which incorporates insights from economic and financial theory, in conjunction with financial data during crisis, has the potential to provide much more accurate answers by directly addressing issues such as liquidity dynamics. There is a need for a joint market and liquidity risk model, covering both stable and crisis periods.

The theoretic properties of the VaR measure, conceptually result in VaR providing misleading information about a financial institution's risk level. The very simplicity of the VaR measure, so attractive when risk is reported, leaves the VaR measure wide open to manipulation. This in turn implies that founding market risk regulations on VaR, not only can impose considerable costs on the financial institution, it may act as a barrier to entry, and perversely increase both bank and systemic risk.

The problems with risk modelling have not gone unnoticed. Anecdotal evidence indicates that many firms employ ad hoc procedures to smooth risk forecasts, and that (some) regulators consider capital as a constant rather than the widely fluctuating variable suggested by the models. Risk modelling does, however, serve a function when implemented correctly internally within a firm, but its usefulness for regulatory purposes is very much in doubt.

## A Empirical Study

The empirical results are a subset of results in Daniélsson (2000). I used there 4 common estimation methods

- Normal GARCH
- Student t GARCH
- Historical simulation
- Extreme value theory

as well as representative foreign exchange, commodity, and equity datasets containing daily observations obtained from DATASTREAM, from the first recorded observation until the end of 1999.

- S&P 500 index
- Hang Seng index
- Microsoft stock prices
- Amazon stock prices
- Ringgit pound exchange rates
- Pound dollar exchange rates
- Clean U.S. government bond price index
- Gold prices
- Oil prices

I estimated each model and dataset with a moving 300, 1,000, and 2,000 day estimation window, and forecast risk one day ahead. Then I record the actual return, move the window and reestimate. This is repeated until the end of the sample.

## B Scaling Laws

This discussion is partially based on de Vries (1998) and Dacorogna, Muller, Pictet, and de Vries (1999), and it draws on insight from extreme value theory.

The following holds for all iid distributions for which the second moment is defined.

- The variance of sum is sum of variances  $V(X + Y) = 2 \times V(X)$  if  $V(X) = V(Y)$  and  $COV(X, Y) = 0$ . This is called self additivity.

If in addition,  $X$  is normally distributed, the self additivity property extends to the tails:

- Implication for the quantile:  $\Pr[X \leq x] = p$ 
  - for the sum over two days, the probability of an outcome,

$$\Pr[X_1 + X_2 \leq 2^{\frac{1}{2}}x] = p$$

- For the sum over  $T$  days

$$\Pr[X_1 + X_2 + \dots + X_T \leq T^{\frac{1}{2}}x] = p$$

- i.e.  $VaR_T = T^{\frac{1}{2}}VaR_{\text{one day}}$

If however  $X$  is iid but *not* normal, the self additivity property does not apply to the tails, however, even if heavy tailed distributions are typically not self additive, the tails are self additive in a special way:

- Consider i.i.d. fat tailed daily returns  $X_t$  where  $\Pr[X_t \leq x] = p$
- Then for a sum of the returns
  - $\Pr[X_1 + X_2 \leq 2^{\frac{1}{\alpha}}x] = p$
- where  $\alpha$  is the tail index
- $\alpha$  is also the number of finite bounded moments
- therefore  $VaR_T = T^{\frac{1}{\alpha}}VaR_{\text{one day}}$

It is known that for financial data  $\alpha > 2$  (if  $\alpha \leq 2$  the variance is not defined, with serious consequences for much financial analysis). Since we can assume that the tail index  $\alpha > 2$ , then

$$T^{\frac{1}{2}} > T^{\frac{1}{\alpha}},$$

which has a number of interesting consequences, e.g.

- The use of the *square-root-of-time* rule to obtain multi-day Value-at-Risk estimates eventually **overestimates** the risk
- The  $\sqrt{T}$  rule usually used for multi day returns
- For fat tailed data, a  $T$  day VaR extrapolated from a one day VaR

$$VaR_T = T^{\frac{1}{\alpha}}VaR$$

Therefore, the  $\sqrt{T}$  rule will eventually lead to an overestimation of the VaR as  $T$  increases.

This has implications in other areas besides risk, e.g. in the pricing of path dependent options, see e.g. Caserta, Danielsson, and de Vries (1998).

## C tables



Table 1: 99% VaR Violation Ratios 1990–1999

Data	Estimation horizon	GARCH Normal	GARCH Student-t	Historical simulation	Extreme value theory
S&P-500	300	1.46	1.07	0.79	0.79
S&P-500	1,000	1.27	0.83	0.95	0.99
S&P-500	2,000	0.91	0.67	1.07	1.07
US bond	300	0.94	0.66	0.49	0.49
US bond	1,000	0.53	0.57	0.66	0.41
US bond	2,000	0.37	0.49	0.67	0.37
Oil	300	1.76	1.38	1.17	1.17
Oil	1,000	1.67	1.30	0.92	1.00
Oil	2,000	1.64	1.04	0.35	0.35
Hang Seng	300	2.18	1.41	0.69	0.69
Hang Seng	1,000	1.90	1.29	1.17	1.21
Hang Seng	2,000	2.02	1.21	1.09	1.09
Microsoft	300	2.24	1.78	1.58	1.58
Microsoft	,1000	1.98	1.60	1.84	1.74
Microsoft	2,000	2.25	1.69	2.06	1.87
GBP/USD	300	2.13	1.42	0.79	0.79
GBP/USD	1,000	1.85	1.18	0.63	0.63
GBP/USD	2,000	1.62	1.14	0.47	0.47

**Notes:** Each model was estimated with three different estimation horizons, 300, 1000, and 2000 days. The expected value for the violation ratio is one. A value larger than one indicates underestimation of risk, and a value less than one indicates overestimation. These results are a part of results in Danielsson (2000)

Table 2: S&amp;P-500 Index 1990-1999. VaR Volatility


Risk Level	Statistic	Estimation Horizon	Returns	GARCH Normal	GARCH Student-t	Historical simulation	Extreme value theory
5%	SE	300	0.89	0.47	0.44	0.41	0.41
5%	Min	300	-7.11	-5.32	-4.17	-2.19	-2.19
5%	Max	300	4.99	-0.74	-0.66	-0.73	-0.73
5%	SE	2,000	0.89	0.42	0.41	0.12	0.12
5%	Min	2,000	-7.11	-3.70	-3.31	-1.55	-1.55
5%	Max	2,000	4.99	-0.81	-0.74	-1.15	-1.15
1%	SE	300	0.89	0.66	0.68	0.71	0.71
1%	Min	300	-7.11	-7.52	-6.68	-3.91	-3.91
1%	Max	300	4.99	-1.05	-1.09	-1.26	-1.26
1%	SE	2,000	0.89	0.60	0.64	0.29	0.33
1%	Min	2,000	-7.11	-5.23	-5.45	-2.72	-2.84
1%	Max	2,000	4.99	-1.14	-1.26	-1.90	-1.90
0.4%	SE	300	0.89	0.76	0.82	1.94	1.94
0.4%	Min	300	-7.11	-8.58	-8.25	-7.11	-7.11
0.4%	Max	300	4.99	-1.19	-1.36	-1.81	-1.81
0.4%	SE	2,000	0.89	0.68	0.80	0.74	0.63
0.4%	Min	2,000	-7.11	-5.96	-6.84	-4.27	-4.18
0.4%	Max	2,000	4.99	-1.30	-1.61	-2.50	-2.49
0.1%	SE	300	0.89	0.88	1.06	1.94	1.94
0.1%	Min	300	-7.11	-9.99	-10.92	-7.11	-7.11
0.1%	Max	300	4.99	-1.39	-1.82	-1.81	-1.81
0.1%	SE	2,000	0.89	0.80	1.06	2.10	1.42
0.1%	Min	2,000	-7.11	-6.94	-9.31	-8.64	-7.49
0.1%	Max	2,000	4.99	-1.52	-2.26	-3.13	-3.15

**Notes:** For each model and the risk level, the table presents the standard error (SE) of the VaR forecasts, and the maximum and minimum forecast throughout the sample period. These results are a part of results in Daniélsson (2000)

Figure 3: Estimated Daily Unconditional Volatility (Smoothed) S&P-500



(a) 1950-1999



(b) 1990's

Figure 4: Daily Hang Seng Index 1997 and 99% VaR

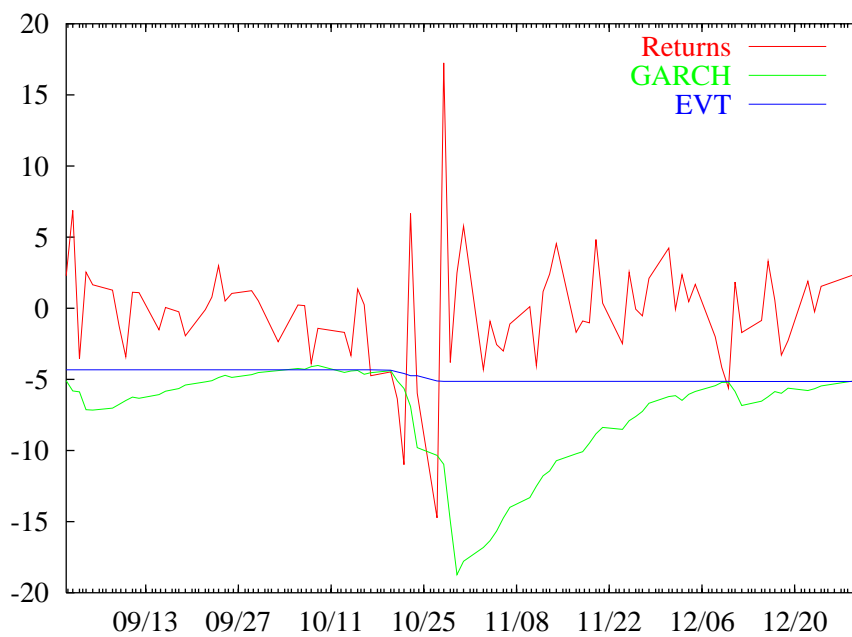




Figure 5: Which is More Volatile and which is more Risky?



(a) Returns  $A$



(b) Returns  $B$

## References

- AHN, D. H., J. BOUDUKH, M. RICHARDSON, AND R. F. WHITELAW (1999): “Optimal Risk Management Using Options,” *J. Finance*, 54 (1), 359–375.
- ANDERSEN, T., AND T. BOLLERSLEV (1998): “Answering the Skeptics: Yes, Standard Volatility Models Do Provide Accurate Forecasts,” *International Economic Review*, 38, 885–905.
- ARTZNER, P., F. DELBAEN, J. EBER, AND D. HEATH (1999): “Coherent Measure of Risk,” *Mathematical Finance*.
- BASEL COMMITTEE ON BANKING SUPERVISION (1996): *Overview of the amendment to the capital accord to incorporate market risk*.
- CASERTA, S., J. DANÍELSSON, AND C. G. DE VRIES (1998): “Abnormal Return, Risk, and Options in Large Data Sets,” *Statistica Neerlandica*, November.
- CHRISTOFFERSEN, P. F., AND F. X. DIEBOLD (2000): “How Relevant is Volatility Forecasting for Financial Risk Management?,” *Review of Economics and Statistics*, 82, 1–11.
- CUMPERAYOT, P. J., J. DANÍELSSON, B. N. JORGENSEN, AND C. G. DE VRIES (2000): “On the (Ir)Relevancy of Value-at-Risk Regulation,” Forthcoming book chapter: Springer Verlag.
- DACOROGNA, M. M., U. A. MULLER, O. V. PICTET, AND C. G. DE VRIES (1999): “Extremal foreign exchange rate returns in extremely large data sets,” [www.few.eur.nl/few/people/cdevries/workingpapers/workingpapers.htm](http://www.few.eur.nl/few/people/cdevries/workingpapers/workingpapers.htm).
- DANÍELSSON, J. (2000): “(Un)Conditionality and Risk Forecasting,” Unfinished working paper, London School of Economics [www.RiskResearch.org](http://www.RiskResearch.org).
- DANÍELSSON, J., P. HARTMANN, AND C. G. DE VRIES (1998): “The cost of conservatism: Extreme returns, Value-at-Risk, and the Basle ‘Multiplication Factor’,” *Risk*, January 1998, [www.RiskResearch.org](http://www.RiskResearch.org).
- DANÍELSSON, J., B. N. JORGENSEN, AND C. G. DE VRIES (2000): “Risk Managements and Complete Markets,” Unfinished working paper, [www.RiskResearch.org](http://www.RiskResearch.org).
- DANÍELSSON, J., AND R. PAYNE (2000): “Real Trading Patterns and Prices in Spot Foreign Exchange Markets,” working paper [www.RiskResearch.org](http://www.RiskResearch.org).
- DE VRIES, C. G. (1998): “Second Order Diversification Effects,” [www.few.eur.nl/few/people/cdevries/workingpapers/workingpapers.htm](http://www.few.eur.nl/few/people/cdevries/workingpapers/workingpapers.htm).
- DROST, F. C., AND T. E. NIJMAN (1993): “Temporal Aggregation of GARCH Processes,” *Econometrica*, 61, 909–927.
- DUNBAR, N. (1999): *Inventing Money: The Story of Long-Term Capital Management and the Legends Behind It*. John Wiley.

- ECONOMIST (1998): "Risk management. Too clever by half," *The Economist*, pp. 140–145.
- ERB, HARVEY, AND VISKANTA (1994): "Forecasting International Correlations," *Financial Analysts Journal*.
- FROOT, K. A., D. S. SCHARFSTEIN, AND J. C. STEIN (1993): "Risk Management: Coordinating Corporate Investments and Financing Policies," *J. Finance*, XLVIII, no 4, 1629–1658.
- GOODHART, C. A. E. (1974): "Public lecture at the Reserve bank of Australia," reprinted.
- HARTMANN, P., S. STRAETMANS, AND C. G. DE VRIES (2000): "Asset Market Linkages in Crisis Periods," working paper.
- JACOBS, B. (1999): *Capital Ideas and Market Realities*. Blackwell.
- LONGIN, F. (1998): "Correlation structure of international equity markets during extremely volatile periods," [babel.essec.fr:8008/domsite/cv.nsf/WebCv/Francois+Longin](http://babel.essec.fr:8008/domsite/cv.nsf/WebCv/Francois+Longin).
- MCNEIL, A. J., AND R. FREY (1999): "Estimation of Tail-Related Risk Measures for Heteroskedastic Financial Time Series: An Extreme Value Approach," *Journal of Empirical Finance*, forthcoming.
- MORRIS, S., AND H. S. SHIN (1999): "Risk Management with Interdependent Choice," Mimeo Nuffield College, Oxford <http://www.nuff.ox.ac.uk/users/Shin>.
- STAHL, G. (1997): "Three cheers," *Risk*, 10, 67–69.