Life after VaR

PHELIM BOYLE, MARY HARDY, AND TON VORST

PHELIM BOYLE
is a professor of finance at the University of Waterloo in Ontario, Canada.
pboyle@uwaterloo.ca

MARY HARDY
is an associate professor of actuarial science at the University of Waterloo.
mrhardy@uwaterloo.ca

TON VORST
is a professor of finance at Erasmus University in Rotterdam, The Netherlands.
vorst@few.eur.nl

The value at risk measure can induce perverse incentives. In fact, a trader can make $1 million in seven days with over 99% probability and still meet a VaR constraint. The trader achieves this objective with no initial funds by taking positions in standard call and put options, specifically by using very large short positions in deep out-of-the-money put options.

The use of conditional tail expectations (tail VaR) instead of VaR can curtail risk-taking tendencies and reinstate more prudent incentives. A new dynamic risk measure known as the iterated conditional tail expectation has attractive risk management properties.

Risk management has become of paramount importance in the financial industry. Value at risk (VaR), a quantile risk measure designed to summarize the risk exposure of an entire firm, is a ubiquitous and simple tool for measuring the risk of a portfolio. It is also used to measure and monitor the risk exposure of trading desks.

The traditional VaR measure violates two of the axioms that Artzner et al. [1997] specify for a *coherent* risk measure. The coherence axioms, a set of characteristics considered desirable for a risk measure, have framed most of the research in risk measures in the past decade:

- The risk measure should not be greater than the maximum possible loss.
- The risk measure should be greater than the mean loss. In other words, on average, the measure should provide sufficient capital to meet losses.
- A proportional change in the loss (e.g., a currency change) should induce a pro rata change in the risk measure, and the risk measure for a certain loss should be equal to the loss.
- The risk measure should be subadditive—that is, the risk measure applied to two separate losses should not be less than the risk measure applied to the aggregated loss. There should be no incentive to break up a portfolio to reduce the capital requirement.

The VaR measure does not, in general, satisfy the requirement that the measure should exceed the mean loss, or the requirement that the measure should be subadditive. As a result of this lack of coherence, investment decisions made with VaR as a binding constraint may introduce some perverse incentives, leading to strange behavior.

Basak and Shapiro [2001] analyze the portfolio investment problem in continuous time under a VaR restriction. They find that the VaR constraint induces an agent to invest more in risky assets than he or she would in the absence of this constraint. The presence of the VaR constraint causes agents to lose more (when large losses occur) than they would without the VaR constraint. Embedding the VaR constraint in an equilibrium model of asset pricing can thus amplify asset price volatility in down markets.

EXHIBIT 1
Benchmark Parameter Values

Parameter	Symbol	Numerical value
Initial Stock Price	S_0	100
Drift of Stock Price Process	μ	0.1
Volatility of Stock Price Process	σ	0.31
Risk free rate	r	0.05
Number of trading days in year	n_y	256
Number of trading days in horizon	n_h	7
Time to horizon in years	$\frac{n_h}{n_y}$	$\frac{7}{256}$

Danielsson, Shin, and Zigrand [2001] reach similar conclusions. They show that the widespread adoption of VaR may serve to increase rather than reduce financial instability in the system. Danielsson [2001] notes that VaR has other deep-rooted drawbacks, arguing not only that it can give misleading information about risk, but also that it may increase both idiosyncratic and systematic risk.

We first illustrate the perverse incentives of VaR using a simple example of standard options. We use the example to motivate alternative risk measures that overcome two of the main drawbacks of VaR. The first drawback is that VaR deals only with the probability of loss and not the severity of loss. The second drawback relates to the fact that VaR is a single-period measure and realistic risk management takes place in a multiperiod world.

Our example involves a trader who wants to make $1 million in seven days. The trader can take long and short positions in standard call and put options, all the while meeting the value at risk constraint. Banks expect their traders to find and exploit profit opportunities in the market, so it does not seem too farfetched to assume that some traders will attempt to game the risk management and monitoring systems in place. Indeed there are a number of well-publicized examples.

A case in point occurred in 2002. John Rusnak, a trader with Allfirst Financial, Inc., a subsidiary of the Allied Irish Banks, managed to chalk up losses of US$691 million through trading activities. One of his techniques was to manipulate the bank's VaR system. Boyle and Boyle [2001] describe other examples of where traders' activities, together with poor risk management and lax oversight, have produced excessive losses.

We modify an example due to Vorst [2000] to use the familiar standard call and put options. We illustrate that the conditional tail expectation does a much better job than VaR in identifying the risk in this example.

Both VaR and the CTE are essentially static one-period risk measures. There is now considerable interest in developing multiperiod or dynamic risk measures (see, for example, Artzner et al. [2003] and Riedel [2003]). We introduce an intuitive multiperiod risk measure.

I. GAMING VAR

We use a continuous-time model to illustrate how VaR can be manipulated. This example is the basis for multiperiod risk measures.

Assume that the stock price dynamics under the real world measure P are given by

$$dS_t = S_t(\mu dt + \sigma dW_t)$$

where W_t is a standard Brownian motion under P. If the initial stock price at time zero is S_0, the solution for the stock price at time t, S_t is given by

$$S_t = S_0 e^{\left((\mu - \frac{1}{2}\sigma^2)t + \sigma W_t\right.} \tag{1}$$

Since W_t has a normal distribution with mean zero and variance t we can readily derive the various quantiles of the stock price distribution at time t, assuming we are now at time zero.

The benchmark parameters are shown in Exhibit 1. Given these parameter values, the 1% quantile of the stock price distribution after seven days is 88.8846, under the real world measure P. Since it is convenient to work in integer stock prices, we record the probabilities as follows.

- The probability of the stock price being 87 or lower is 0.0030.
- The probability of the stock price being 88 or lower is 0.0058.
- The probability of the stock price being 89 or lower is 0.0107.

A trader who wishes to make $1 million and stay above the VaR threshold can:

- Purchase 1 million seven-day call options with a strike price of 87.
- Sell short 1 million seven-day call options with a strike price of 88.
- Finance the cost of the call options by selling short enough seven-day put options with a strike price of 87.

EXHIBIT 2

Market Value of Portfolio After Seven Days for Different Stock Prices

Realized stock price after seven days	Corresponding market value of portfolio for this stock price millions of dollars
90	1
89	1
88	1
87	0
86	-226.39
85	-452.78
84	-679.17

This strategy will provide the trader with \$1 million in seven days as long as the terminal stock price is at least 88. The probability of this occurring is 99.42%. The trader loses money only if the terminal stock price is 87 or lower, and the probability that this will occur is 0.30%, which lies below the VaR threshold.[1]

Thus, the trade does not contribute to the VaR. The VaR calculation relies only on the probability of loss, and makes no allowance for low-probability, high-severity, risks—of which this trade is an example. This trade has the potential to generate very severe losses, yet is not recognized as dangerous under the conventional VaR measure.

From the Black-Scholes-Merton model, the market price of the 87 call is 13.1233 and the market price of the 88 call is 12.1293. Hence the net cost of purchasing 1 million 87 calls and selling short 1 million 88 calls is 993,988.2. As the market price of the 87 put is 0.00439, this means the total number of puts to be sold short is 226,390,369.[2]

Exhibit 2 displays the profit (loss) at maturity against the final stock price. Note that for stock prices of 88 and higher the profit is \$1 million, but that the losses explode as the stock price goes below 87.

In practice there will be other controls in place to prevent this sort of gaming. There is normally a limit on the notional amount that can be traded by each trader. In addition, there is often daily monitoring of the delta and gamma of traders' positions. Risk managers do not live by VaR alone.

The example highlights the dangers of VaR when we assume it is the primary risk management tool. VaR's fatal flaw in this context is that it is concerned only with the probability of a loss and not with the severity of the loss. A trader has a powerful incentive to structure a portfolio so that there is a small probability of a huge loss. This suggest that a risk measure that takes into account both the probability of the loss and the severity may be more appealing.

It turns out that the conditional tail expectation (CTE), which is a measure of the expected losses in the left-hand tail of the profit and loss distribution, is a much more useful risk measure than value at risk. We provide first a general description of the CTE that does not cover every case, and then use a more precise definition that covers all the cases. The CTE is described for some parameter α, $0 \leq \alpha < 1$—say $\alpha = 0.9$—as the average loss, given that the loss falls in the worst $(1 - \alpha)$ or 1% of the distribution.

Our starting point is the value at risk measure, $V_{99\%}$, say, for a continuous loss distribution (or, more strictly, if $V_{\alpha+\varepsilon} > V_\alpha$ for any $\varepsilon > 0$). Then the CTE with parameter α is:

$$CTE_\alpha = E[\text{loss} \mid \text{loss} > V_\alpha] \tag{2}$$

where V_α is the α-value at risk for the loss.

Note that this definition, although appealing, does not give the right result when the value at risk number falls in a probability mass (that is, where there is some $\varepsilon > 0$ such that $V_{\alpha+\varepsilon} = V_\alpha$). In this case we use a more general formulation, calculating the CTE with parameter α as follows. Find $\beta^1 = \max\{\beta : V_\alpha = V_\beta\}$. Then:

$$CTE_\alpha = \frac{(1 - \beta')E[\text{loss} \mid \text{loss} > V_\alpha] + (\beta' - \alpha}{1 - \alpha} \tag{3}$$

Hence, while the 99% VaR measure tells us the level of losses that has a 1% chance of being exceeded, the 99% CTE tells us the average loss if the loss falls in the worst 1% of the distribution.

The CTE measure is coherent in the sense of Artzner et al. [1997]. This important property of the CTE has been discovered independently by Acerbi and Tasche [2001], Artzner et al. [1999], Longin [2001], and Wirch and Hardy [1999]. The risk measure appears under a variety of different names in the literature, including conditional tail expectation, tail VaR, BVaR, and CVaR.

In our numerical example, the conditional tail expectation can be computed in closed form. For notational ease, let $n_1 = 226{,}390{,}369$ denote the number of put options sold short and $n_2 = 1{,}000{,}000$ denote the number of call options.

Exhibit 3 displays the trader's profit and loss after seven days depending on different stock price ranges. We denote the stock price in seven days' time by S_7 and let $k_1 = 87$, $k_2 = 88$, and $k_3 = 88.8846$ (k_3 is the 1% quan-

EXHIBIT 3
Trader's Profit and Loss After Seven Days

Stock price range	Trader's profit(> 0) or loss(< 0)	Profit or loss
$0 < S_7 \leq k_1$	$-n_1(k_1 - S_7)$	Loss
$k_1 < S_7 \leq k_2$	$+ n_2(S_7 - k_1)$	Profit
$k_2 < S_7 \leq k_3$	$+n_2$	Profit
$k_3 < S_7$	$+n_2$	Profit

tile of the distribution of S_7).

Exhibit 3 provides the profit and loss distribution of the maturity payoff for stock prices that lie below the 1% quantile, $k_3 = 88.8846$. By taking the expected value of losses in the tail of the distribution, we can obtain the CTE. In this case, the losses occur when the put option is in the money at maturity. Hence the CTE is given by:

$$CTE = E_P[L \mid 0 < S_7 \leq k_3)]$$
$$= \frac{\int_0^{k_1} n_1(k_1 - s)f(s)ds - \int_{k_1}^{k_2} n_2(s - k_1)f(s)ds - \int_{k_2}^{k_3} n_2 f}{Prob[0 < S_7 \leq k_3]}$$

where $f(s)$ is the lognormal density function for S_7.

By working out the various expectations using the properties of the lognormal distribution, we can write the expression for the CTE:

$$CTE = \frac{1}{N(-d_5)}\Big[n_1\big(k_1N(-d_2) - S_0e^{\mu h}N(-d$$
$$-n_2\big\{(S_0e^{\mu h}\big(N(-d_3) - N(-d_1)\big) - k_1(N(-d_4) - N(-$$
$$-n_2(N(-d_5) - N(-d_4 \qquad (4)$$

where $h = 7/256$

$$d_1 = \frac{\left(ln(\frac{S_0}{k_1}) + (\mu + \frac{\sigma^2}{2})h\right)}{\sigma\sqrt{h}} \qquad d_2 = d_1 -$$

$$d_3 = \frac{\left(ln(\frac{S_0}{k_2}) + (\mu + \frac{\sigma^2}{2})h\right)}{\sigma\sqrt{h}} \qquad d_4 = d_3 -$$

$$d_5 = \frac{\left(ln(\frac{S_0}{k_3}) + (\mu - \frac{\sigma^2}{2})h\right)}{\sqrt{\ }}$$

Using this formula, we find that the 99% CTE for the example is 90.594 million, compared with a 99% VaR of zero.[3] The sheer size of this number indicates that the conditional tail expectation reflects the riskiness of the trader's strategy. Thus this measure eliminates the very perverse incentives associated with VaR.

The market value of the trader's position is zero at inception, but this is a highly levered position, so the market value can easily become negative very quickly. To see how this happens, consider the situation one day later. The options will now have only six days left to maturity.

The details are summarized in Exhibit 4. For a range of possible stock prices, the first column shows the probabilities of attaining a stock price of this level or lower after a single day, assuming that the initial stock price is 100. The last column shows the market value of the trader's portfolio for each stock price.

Note there is a 14.6% probability that the market value of the portfolio will be at least 800,000 in the red and a 5.7% probability that it will be 2.449 million in the red. If these market values occur, the alarm bells should sound—but this will be too late if the VaR measure is used to quantify risk at issue; a 99% VaR of zero contrasts with more than a 25% probability that the market value of the portfolio will be negative after one day. On the other hand, at 99% CTE, 90.594 million easily absorbs the risk over a single day.

Exhibit 5 indicates that the 99% value at risk of the one-day loss would not be zero, but instead $7.05 million, corresponding to a stock price after one day of $95.6122. The one-day 99% CTE would be $11.29 million. The difference between the two risk measures is much less than in the seven-day case, as the range of possible outcomes after one day is not nearly so great as after seven days.

We can repeat the calculations for any horizon, from one day to seven days, to show how the CTE and VaR differ as the horizon changes. These results are given in Exhibit 5, together with the 1% quantile of the stock price, S_t, at each horizon.

II. MULTIPERIOD RISK MEASURES

The CTE has been shown to have advantages over the VaR approach, but still suffers from the disadvantage that it is a single-period risk measure. That is, both CTE and VaR are calculated at time zero (say), by looking at the portfolio value at maturity, without any considera-

EXHIBIT 4

Market Value of Portfolio One Day Later

Stock price one day afterwards	Probability of stock price at this level or lower	Market value of trader's portfolio for this stock price
100	0.4958	+554,475
99	0.2983	+90,176
98	0.1461	-800,338
97	0.0568	-2,448,830
96	0.0171	-5,391,985
95	0.0039	-10,456,263

EXHIBIT 5

99% VaR and 99% CTE for Different Horizons
($ millions)

Horizon t	99% VaR	99% CTE	1%ile of S_t
1 day	7.05	11.15	95.612
2 days	12.52	23.88	93.863
3 days	16.73	38.66	92.546
4 days	18.72	55.86	91.454
5 days	16.95	71.40	90.504
6 days	8.94	85.77	89.656
7 days	0	90.59	88.885

tion of what could happen in between these two dates.

Suppose now that a CTE of $90.592 million is held at the start of the contract, and (for simplicity) ignore interest. Assume also that the CTE is recalculated daily, at the same 99% standard. The CTE required each day would depend on the underlying stock process. What is the probability that additional risk capital would be required?

Consider first the position after one day. Exhibit 6 repeats the information from Exhibit 4 but shows the 99% CTE, assuming the portfolio remains in force, rather than the market value of the position.[4]

The CTE brought forward, of 90.6 million, will be adequate if the stock price remains above 99. In fact, with six days to go, the capital of $90.59 million is the CTE for a stock price of S_1 = 98.874; this means that if the stock price falls below $98.874 on the second day, additional capital will be required if the contract is allowed to continue.

The probability of this is easily calculated, from the lognormal distribution, at 27.6%. So, even though the CTE is a measure from the 1% worst outcomes, there is more than a one-in-four chance that it will not be enough to meet the capital requirement the following day. Of course, the trader may decide to liquidate the portfolio (it will have positive value if the stock price is higher than 98.874), but to continue with the contract, holding capital determined by the 99% CTE, an injection of funds will be required.

Even if the stock price remains above 98.874 after the first day, if it falls below a threshold of 97.672 on the second day, the initial CTE will be insufficient. The probability that the initial CTE is adequate after one day, but inadequate after two days, is 5.7%.

When we consider all seven days of the contract, the probability that the initial CTE will be inadequate on one or more subsequent days is around 36%.

To calculate this, we first calculate the threshold point for each day. This is the price to which the stock must fall on that day for which the shortfall is exactly equal to the initial capital available, $90.59 million. We have already computed this for the first day—the threshold is 98.874, and for the second day it is 97.672.

The figures for all seven days are shown in the first column of Exhibit 7. The last column shows the probability that the stock price falls below the threshold for the first time on the t-th day.

To determine a measure of risk that takes into consideration intermediate capital requirements, it is necessary to use a multiperiod approach to measuring risk. Artzner et al. [2003] and Riedel [2003] propose a framework for a multiperiod risk measure. Hardy and Wirch [2004] propose a specific risk measure that satisfies the criteria established in that research, the *iterated CTE* or ICTE.

The iterated CTE is calculated by working backward from the final position. We illustrate this risk measure using our numerical example, and assuming that the risk manager recalculates the CTE at the end of the third day. We thus have a two-period situation, the first running from the start of the contract until the end of the third day, and the second running from the end of the third day to the end of the contract.

At the start of the contract, we can calculate the distribution of the CTE that will be required at the end of the third day. The CTE on day 3, $CTE_3(S_3)$, say, depends on the stock price S_3, which is random, with a lognormal distribution. So, at the start of the contract $CTE_3(S_3)$ is a random variable, and may be considered a liability at time 3. We can therefore calculate the CTE of this liability. That is, for this two-period model, the iterated CTE is the time zero CTE of the random variable CTE_3.

Since low values of the stock price S_3 give high

EXHIBIT 6

CTE of Portfolio One Day Later for Different Stock Prices

Stock price one day afterwards	Probability of a stock price of this level or lower	99% CTE of Trader's portfolio for this stock price $ millions
100	0.4958	39.8
99	0.2983	82.9
98	0.1461	165.8
97	0.0568	319.6
96	0.0171	519.3
95	0.0039	719.1

values for the CTE_3, the worst 100(1 – α)% of the distribution of CTE_3 corresponds to the lower 100(1 – α)% of the stock price distribution S_3:

$$ICTE = \int_0^{Q_\alpha} \mathrm{CTE}_3(s) f(s)\,ds$$

where Q_α is the lower 100(1 – α)% quantile of the distribution of S_3, $\mathrm{CTE}_3(s)$ is the CTE at time 3 for the final (time 7) liability, given that the stock price at time 3 is *s*, and *f*(*s*) is the density function of the appropriate lognormal distribution for the stock price S_3. The CTE standard is given by α, which may or may not be the same as the one-period standard used for $\mathrm{CTE}_3(S_3)$—that is, 99% in the example above. In fact, the use of 99% for the second iteration would produce a very conservative value for the iterated CTE. The ICTE would then represent the mean of the worst 1% values at time 3 of the worst 1% values at time 7—clearly we are very deep into the tail of the distribution.

Even using a relatively weak standard for the second CTE, say, $\alpha = 0.75$—that is, we hold at time zero the average of the lower quartile of all the possible values for the time 3 CTE—the iterated CTE at time zero would be $206 million; this compares with the single-period CTE of around $91 million.

Exhibit 8 shows the effect of changing the standard used for the second iteration on the risk measure. For $\alpha > 0.75$, the single-period CTE may be quite inadequate for demonstrating required capital after three days. If we split the term into more subperiods, the iterated CTE would be even higher.

While the CTE alerts the risk manager to the dangers of the low-probability, high-severity, outcome, the ICTE further demonstrates the ongoing risk of calls on capital implied by this trade. Clearly these capital requirements are too high for a practical capital requirement, however. The ICTE can be useful when one party is locked into a contract (this can be effectively true for some financial insurance contracts). In our example, since the trader can unwind the deal at any time, it would not be necessary to assume that the contract is maintained until the options mature, regardless of the underlying stock price. The iterated CTE might also be used as a measure to assess capital requirement risk when comparing contracts.

We might assume a shorter horizon and still use the multiperiod risk measure to measure the risk of a future capital draw. Suppose we assume a two-day horizon. From Exhibit 5, we know that the 99% CTE for the two-day horizon is approximately $24 million. We can compare this with the iterated CTE, where we take a two-step iteration, with a 99% standard at each step. The ICTE is then the 99% CTE at the start of the contract of the 99% CTE after one day of the portfolio value after two days. The ICTE is approximately $90 million.

What this number indicates is that if things go badly on the first day, action needs to be taken quickly. If things continue to go badly, the cost by the second day could be very significant.

These results suggest a possible risk management strategy consisting of two parts that are analogous to initial and maintenance margins in futures contracts.[5] That is, the VaR based on the terminal distribution for the position value would be used in setting the initial capital requirement, while the CTE (based on a less stringent probability criterion) would be used to set a minimum

EXHIBIT 7

Threshold Values for Daily Stock Price to Exactly Require Initial Capital of $90.592 Million

Days from issue t	Threshold Stock Price	Probability that the stock price first falls below the threshold on t-th day
1	98.874	0.276
2	97.672	0.057
3	96.370	0.019
4	94.927	0.009
5	93.268	0.003
6	91.198	0.001
7	86.596	0.000
Total		0.365

EXHIBIT 8

Iterated CTE at Time Zero for Different Levels of α Over Second Period

Value of α	ICTE at outset ($ millions)
99%	999
95%	615
90%	420
75%	206
50%	107

position value that would be monitored daily and would trigger a margin call (i.e., required capital infusion) if it were violated on any given day.

III. CONCLUSIONS

Our simple example exposes some of the inherent flaws in value at risk as a risk measure. The difficulty is that VaR takes into consideration only the probability of loss and not the potential size of the loss, lending itself to manipulation. Coherent risk measures are superior to VaR in this respect. The conditional tail expectation, for example, measures the expected losses in the tail, given that a loss has occurred. We show that the CTE is a much more effective risk measure in this regard. The CTE is also intuitive and easy to compute.

VaR and the CTE are traditionally implemented in a single-period framework, but there is now considerable interest in the construction of multiperiod risk measures. The single-period risk measure calculated to cover losses over, say, a seven-day period ignores the possibility that additional capital infusions may be required before the end of the period. We show in a relatively simple example that there is a high probability that additional risk capital will be required. It is reasonable to anticipate the additional capital requirement, and the multiperiod, or dynamic, risk measure is designed to do this.

The particular example of a multiperiod risk measure that we have described, the iterated CTE, treats the intermediate CTE as the loss, and takes the CTE of that loss. This number, interpreted as a capital requirement, protects the portfolio (to the extent of the solvency standard selected through the parameter) both from ultimate losses and from intermediate additional capital requirements.

The intermediate capital draw may be considered less critical than the ultimate losses because it may still be feasible to liquidate the portfolio without significant loss. It might be reasonable to use a lower capital standard for that part of the risk measure, maintaining the higher standard for the protection against ultimate losses.

The iterated CTE may also be used as an additional risk measure for comparing investments. A lower ICTE will indicate less of a chance that additional capital will be required as time passes.

ENDNOTES

The authors thank Weidong Tian for useful comments and Lin Yuan and Jessica Ling-Wai Lam for technical assistance. They are grateful to Raghu Sundaram for his comments on an earlier draft. Boyle and Hardy acknowledge research support from the Natural Sciences and Engineering Research Council of Canada.

[1]Note that these two probabilities do not add up to one since 0.9942 + 0.00030 = 0.9972. This is because the probability of the stock price being between 87 and 88 is 0.0028.

[2]We compute the option prices to nine decimal places to eliminate rounding errors.

[3]Some authors prefer to restrict profits so that only losses are taken into consideration. In this case, the second and third terms in the formula for the CTE would be replaced by zero, and the CTE would be 91.163 million.

[4]In other words, we do not mark the portfolio to the market values shown in the final column of Exhibit 4.

[5]We thank the editor for this suggestion.

REFERENCES

Acerbi, C., and D. Tasche. "On the Coherence of Expected Shortfall." Working paper, Abaxbank, Italy, 2001.

Artzner, P., F. Delbaen, J.-M. Eber, and D. Heath. "Coherent Measures of Risk." *Mathematical Finance*, 9(3) (1999), pp. 203-228.

——. "Thinking Coherently." *Risk*, 10(11) (1997).

Artzner, P., F. Delbaen, J.-M. Eber, D. Heath, and H. Ku. "Coherent Multiperiod Risk Adjusted Values and Bellman's Principle." Working paper, ETH Zurich, Switzerland, 2003.

Basak, S., and A. Shapiro. "Value-at-Risk-Based Risk Management: Optimal Policies and Asset Prices." *The Review of Financial Studies,* 14(2) (2001), pp. 371-405.

Boyle, P., and F. Boyle. *Derivatives: The Tools that Changed Finance*. London: Risk Books, 2001.

Danielsson, J. "The Emperor Has No Clothes: Limits to Risk Modelling." Working paper, Financial Markets Group, London

School of Economics, 2001.

Danielsson, J., H.S. Shin, and J.-P. Zigrand. "Asset Price Dynamics with Value at Risk Constrained Traders." Working paper, Financial Markets Group, London School of Economics, 2001.

Hardy, M.R., and J.L. Wirch. "The Iterated CTE: A Dynamic Risk Measure." *North American Actuarial Journal*, 8 (2004), pp. 62-75.

Longin, F.M. "Beyond the VaR." *The Journal of Derivatives,* 4 (2001), pp. 7-49.

Riedel, F. "Dynamic Coherent Risk Measures." Working paper, Stanford University Department of Economics, 2003.

Vorst, T. "Optimal Portfolios Under a Value at Risk Constraint." Working paper, Erasmus University Rotterdam, 2000.

Wirch, J.L., and M. R. Hardy. "A Synthesis of Risk Measures for Capital Adequacy." *Insurance: Mathematics and Economics*, 25 (1999), pp. 337-347.