The Tail that Wags the Dog: Integrating Credit Risk in Asset Portfolios

NORBERT J. JOBST AND STAVROS A. ZENIOS

NORBERT J. JOBST
is a Fellow at the
HERMES Center on
Computational Finance and
Economics, University of
Cyprus, in Nicosia, Cyprus,
and Ph.D. candidate, De-
partment of Mathematical
Sciences, Brunel Univer-
sity, West London, UK.
He may be contacted at
norbert.jobst@brunel.ac.uk

STAVROS A. ZENIOS
is professor and director of
the HERMES Center on
Computational Finance
and Economics, Univer-
sity of Cyprus, Director of
RiskLab, Cyprus Interna-
tional Institute of Manage-
ment and Senior Fellow
at the Financial Institu-
tions Center, the Wharton
School in Philadelphia,
PA. He may be contacted
at zenioss@ucy.ac.cy

redit risky securities and most of their derivatives are characterized by a large chance of positive returns and a very small probability of large investment losses. The distribution of price returns of these instruments is asymmetric and highly skewed, exhibiting very flat tails on the downside. That investors are compensated for assuming the low-probability risk of losses is evidenced by the total return of corporate bond indices versus the U.S. Treasury market, as summarized in Exhibit 1 for the 11.5-year period January 1990–June 2001.

On a duration-adjusted basis the Merrill Lynch Eurodollar index, for instance, realized an annualized return of 9.35% with a standard deviation of 4.48% over the period December 1985 to March 2001. The Eurodollar index outperformed the U.S. Treasury index during this period, which realized duration adjusted returns of 8.74% with a standard deviation of 4.73%.

We might well ask whether the return of the credit risky portfolio adequately compensates for the low-probability large losses events. Given the proliferation of corporate bond issues, the constant stream of innovations in credit derivatives, and their increased use in the asset portfolios of institutions, it is appropriate that credit risk pricing models receive widespread attention; see, for instance, Saunders [1999] and Schönbucher [2000]. Costly lower-tail outcomes also have an impact on the practice of enterprise risk management as pointed out by Stulz [1996]. Models for integrated risk management in the context of credit risky securities are few, however.

We demonstrate the significance of the tails in shaping the risk profile of credit asset portfolios. We start with the obvious: namely, that properly simulated credit events result in risk profiles with flat tails on downside risk (i.e., losses) and limited upside potential (i.e., gains). We develop from this more subtle and important observations:

- Losses are probabilistic events, and without adequate accuracy the low-probability events may be missed. We demonstrate this in the section entitled *"Where is the Tail?."*
- Recognizing that these low-probability events can lead to different optimal portfolios is shown in the section entitled *"Tail Effects on Efficient Frontiers."*
- Even so, standard risk measures do not adequately penalize the low-probability events. The resulting portfolios might be efficient with respect to the standard metric, yet unacceptable to portfolio managers because of the probability of achieving substantial losses. Different risk measures such as conditional VaR should therefore be used. Relevant models are introduced in the section entitled *"Optimizing the Right Risk Metric."*
- With appropriate modeling, long-term performance goals can be met without suffering catastrophic blows from the tails in the short run. This is demonstrated in the section, *"Long-Term Performance with Short-Term Tails."*

EXHIBIT 1

Total Return of Broad Market Indexes
January 1990–June 2001

Index	Return (%)
Merrill Lynch Eurodollar	139.68
U.S. Treasury	137.37
U.S. Agency Master (AAA)	142.39
U.S. Corporate Domestic Bonds	156.34

These observations are supported by empirical analysis carried out using the simulation models developed recently by Jobst and Zenios [2001]. In our numerical experiments we simulate 500 economic scenarios and 5,000 credit events, for a total of 2.5 million simulation runs. Portfolio models are optimized on samples of 5,000 to 10,000 scenarios, drawn from the 2.5 million simulations.[1]

WHERE IS THE TAIL?

The causes of loss due to credit assets are many and complex. Credit risk can be described as the changing expectations of an obligator's ability or willingness to fulfill its obligations on a certain date or at any time beyond. Losses may result from a default or a change in market value due to credit quality migration. In general, credit risk for a single instrument may be decomposed into default risk, migration risk, and security-specific risks

that cause idiosyncratic spread changes. Default is the low-probability, large-impact, event.

Tools such as CREDITRISK+ from CSFB, Credit-Metrics from J.P. Morgan, Credit Portfolio View from McKinsey & Co, or KMV's Portfolio Manager allow us to gain important insights into credit risks, but a number of important aspects are missing. CREDITRISK+, for example, assesses credit risk due to default losses only without taking into consideration the term structure of credit spreads. CreditMetrics allows calculation of the present value of a portfolio of credit risk-sensitive assets depending on credit risk only. Market risk is not incorporated explicitly. As a result, no other risks apart from credit risk can be assessed for their impact on the valuation of the portfolio.

Jobst and Zenios [2001] show how some popular pricing models can be extended to the valuation and simulation of portfolios of credit risky securities and their derivatives. These extensions allow us to estimate the risk profile of portfolios taking into account market and spread risk, and the risks of rating migration, defaults, and recovery. The simulations reveal—and quantify—the flat tails due to credit events.

For instance, Exhibit 2 shows a flat lower tail when credit rating migrations are simulated under current economic conditions. This tail is absent when simulating only market and spread changes under constant volatility.

The tail in this example is quite pronounced as it was simulated assuming that the term structure of risk-free rates

EXHIBIT 2

Distribution of Returns of a Baa Bond Portfolio
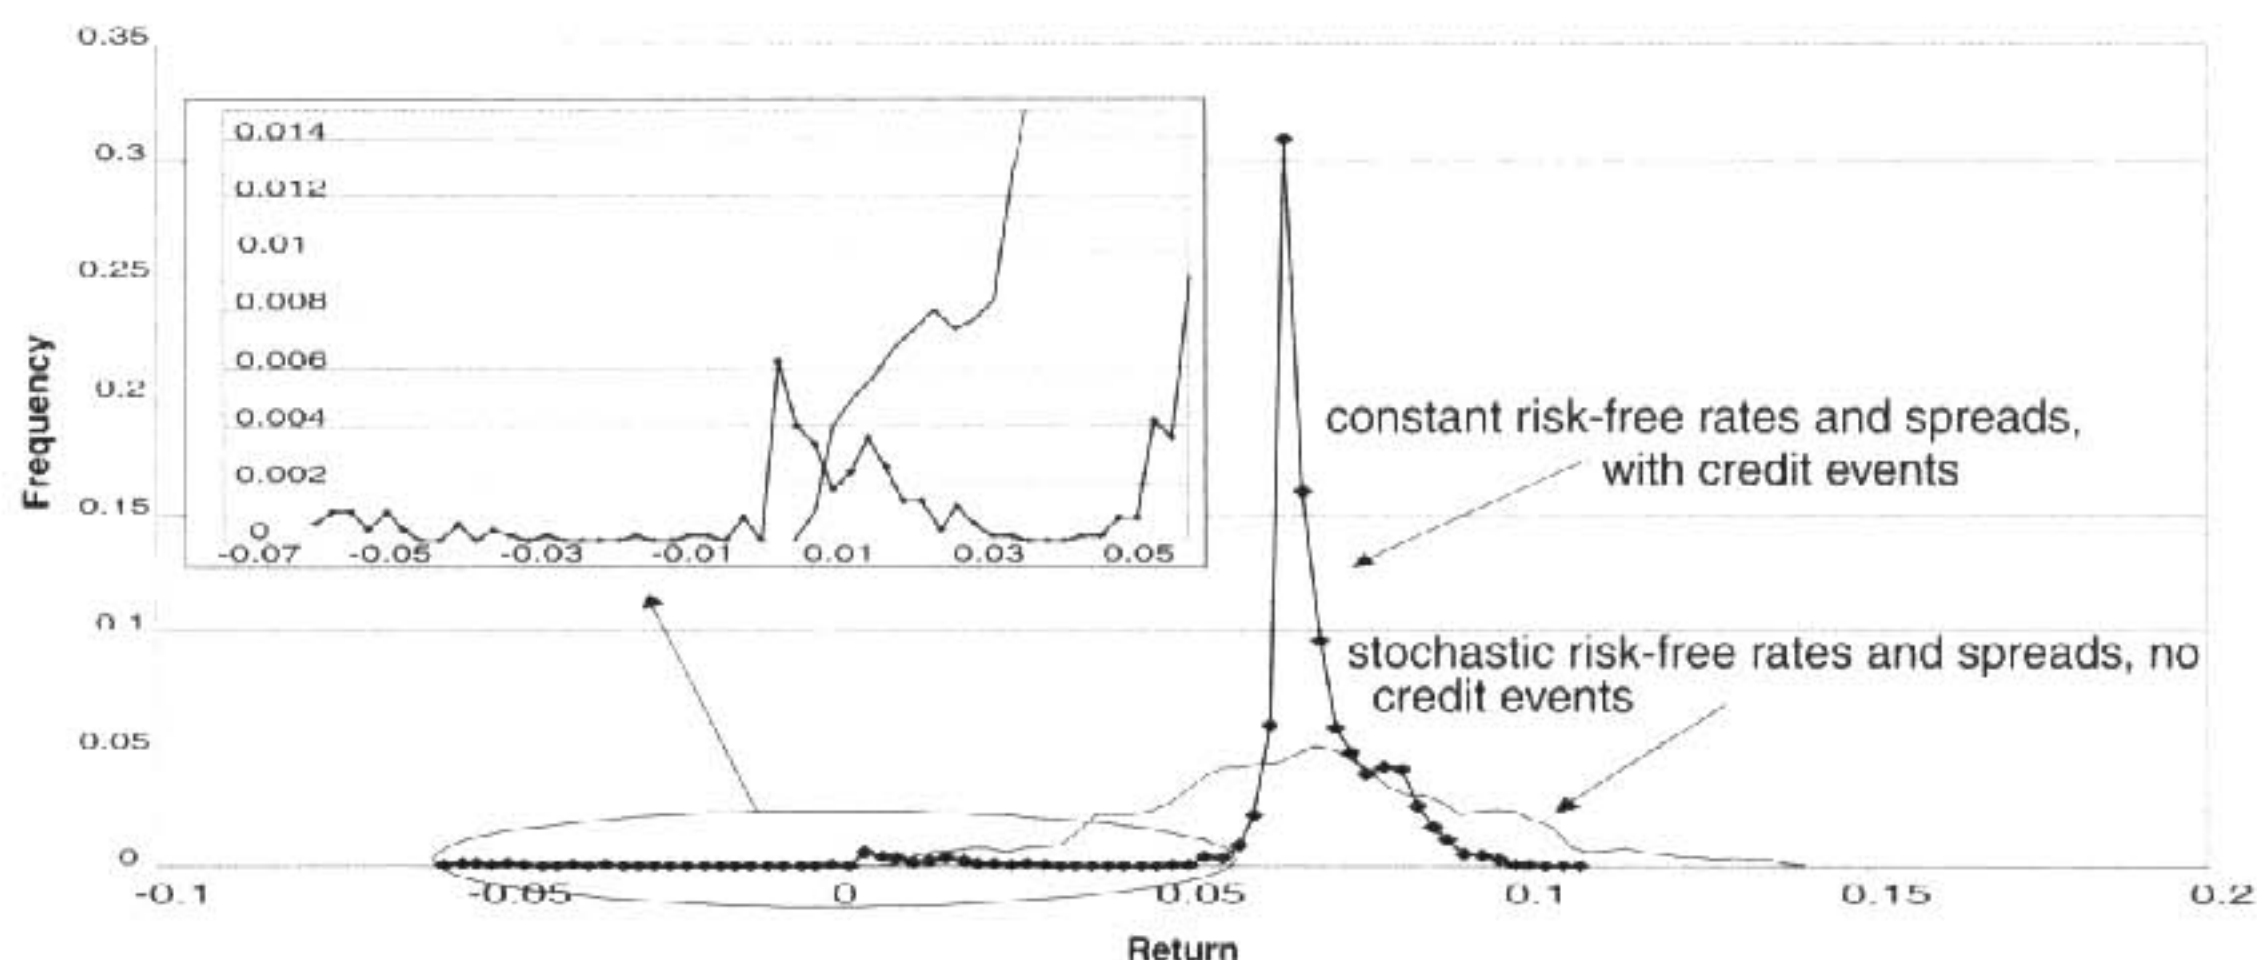
Heavy line: Constant risk-free rates and credit spreads with credit rating migrations and default. Fine line: Stochastic risk-free rates and credit spreads with constant volatility and without credit rating migrations or default.

and credit spreads remains unchanged. Exhibit 3 develops the simulation of the same portfolio integrating market and credit spread risk, and then adding the credit events, i.e., credit rating migrations and defaults. The tail is once more evident but not when we look at the 0.95 quantile of the distributions.

The 0.95 quantiles of the two distributions shown in Exhibit 3 are positive and very close to each other, but the 0.99 quantiles have the opposite signs and differ by an order of magnitude. There are no losses at the 0.95 probability level, even when credit events are properly simulated. At the 0.99 probability level, however, we observe losses of 1.7% when credit events are simulated.

The lesson to take away is that tails are probabilistic events, and without an accurate simulation method they may be missed.

Improved simulation may take the form of accurate numerical methods for tail resolution or the use of confidence intervals for the extremes of the loss distribution, or any other technique that allows us to observe the tails that drive low-probability events. We will see later that this observation has ramifications for the choice of an appropriate risk metric for portfolio optimization.

TAIL EFFECTS ON EFFICIENT FRONTIERS

Ignoring the tails has a significant effect on the efficient frontiers. We simulate first the distribution of returns of 17 bonds rated Baa, *without* credit rating migration and defaults, and apply a mean absolute deviation portfolio optimization model (the MAD of Konno and Yamazaki [1991]) to these simulated data.

The efficient frontier of portfolio expected return against its mean absolute deviation is shown by the fine line in Exhibit 4. The dashed line represents the frontier calculated according to the expected return and the mean absolute deviation of the optimized portfolios and using the distribution of returns *with* credit rating migrations and

---

EXHIBIT 3
Distribution of Returns of a Baa Bond Portfolio
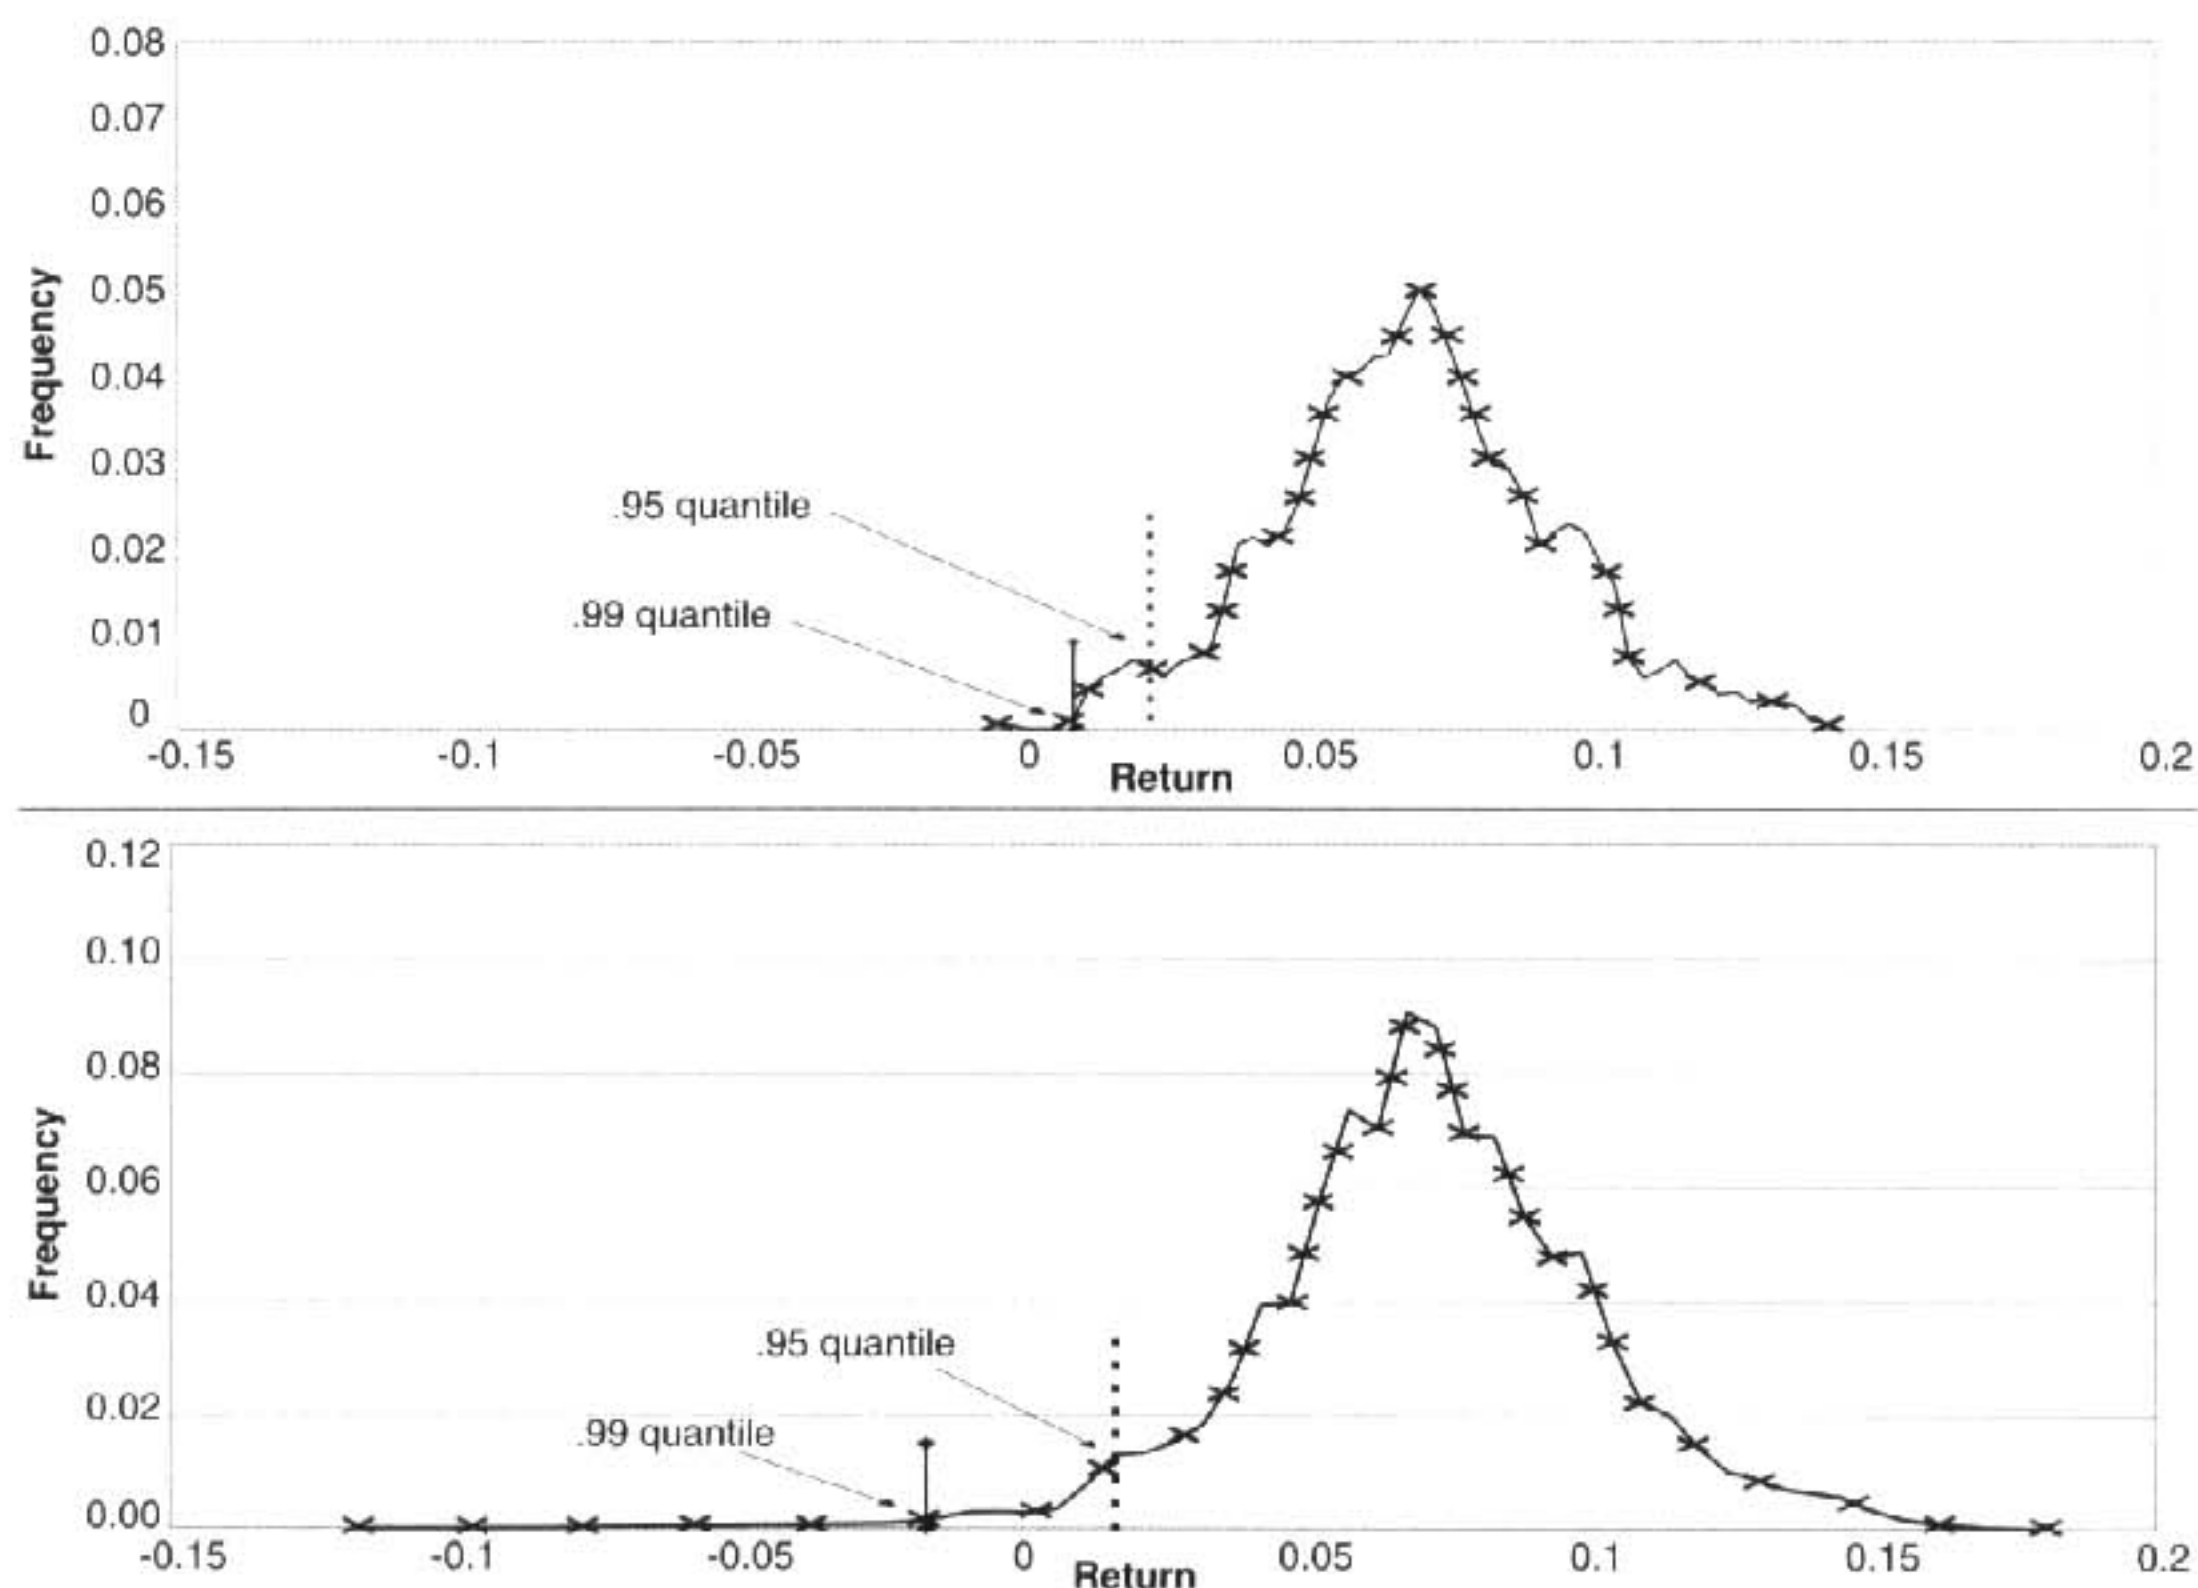
First graph: Without credit rating migrations or defaults. Second graph: With credit rating migrations and defaults.

defaults. Thus we perform an out-of-sample sensitivity analysis of the frontier using tail scenarios that are not included in the scenario sample of the optimization model.

There is nothing efficient about the optimized portfolios obtained by ignoring the tails, once the tails are properly accounted for. This is our second observation: Tails distort the risk-return frontiers, making seemingly efficient portfolios inefficient ones.

Running the mean absolute deviation portfolio optimization using the distribution *with* credit events, we obtain a frontier that is very close to the out-of-sample frontier and eliminates the inefficient portfolios; this frontier is shown by the thick solid line in Exhibit 4. Does this imply that it is sufficient to accurately simulate the tails, and then develop portfolios that optimally trade off expected return against risk? The answer is of course affirmative, although the mean absolute deviation risk measure does not properly account for the tails.

The distribution of returns of the minimum-risk portfolio obtained using the MAD model is shown in Exhibit 5. Note there is a small probability of losses in excess of 80% of the portfolio value. These losses are likely to be catastrophic, and when they occur they will most likely—due to bankruptcy—block the prospects of the long-term expected return. The long-term expected return of the minimum-risk portfolio obtained using MAD is 5.5%, a return that will be realized only if the portfolio is not ruined in the short term.

This is the same observation Stulz [1996] makes in explaining the discrepancy between the corporate use of derivative securities advocated by theory, and their actual use in practice as revealed by the Wharton surveys (see Bodnar, Hayt, and Marston [1998] for the latest survey results). It is a common limitation when applying mean-variance or MAD models to optimize non-normal distributions; Leland [1999] shows similar problems with Sharpe ratios.

What then should be done in order to properly account for the tail effects? The answer is to select a risk metric that penalizes appropriately extreme events, and then optimize the portfolio composition with respect to this metric of risk.

OPTIMIZING THE RIGHT RISK METRIC

Value at Risk (VaR) has become an industry standard for measuring extreme events and integrating disparate sources of risk. VaR answers a particular question: What is the maximum loss with a given confidence level (say, $\alpha \times 100\%$) over the target horizon? Its calculation also reveals that with probability $(1 - \alpha)100\%$ the losses will exceed VaR.

Consider a portfolio with value $V(x, \tilde{p})$. This is a function of the holdings $x = (x_i)_{i=1}^m$ of assets in the portfolio, and of the random asset prices $\tilde{p}$. If the current value of the portfolio is V_0, the losses in portfolio value are given by the *loss function:*


EXHIBIT 4
Frontier Generated Using MAD Portfolio Optimization
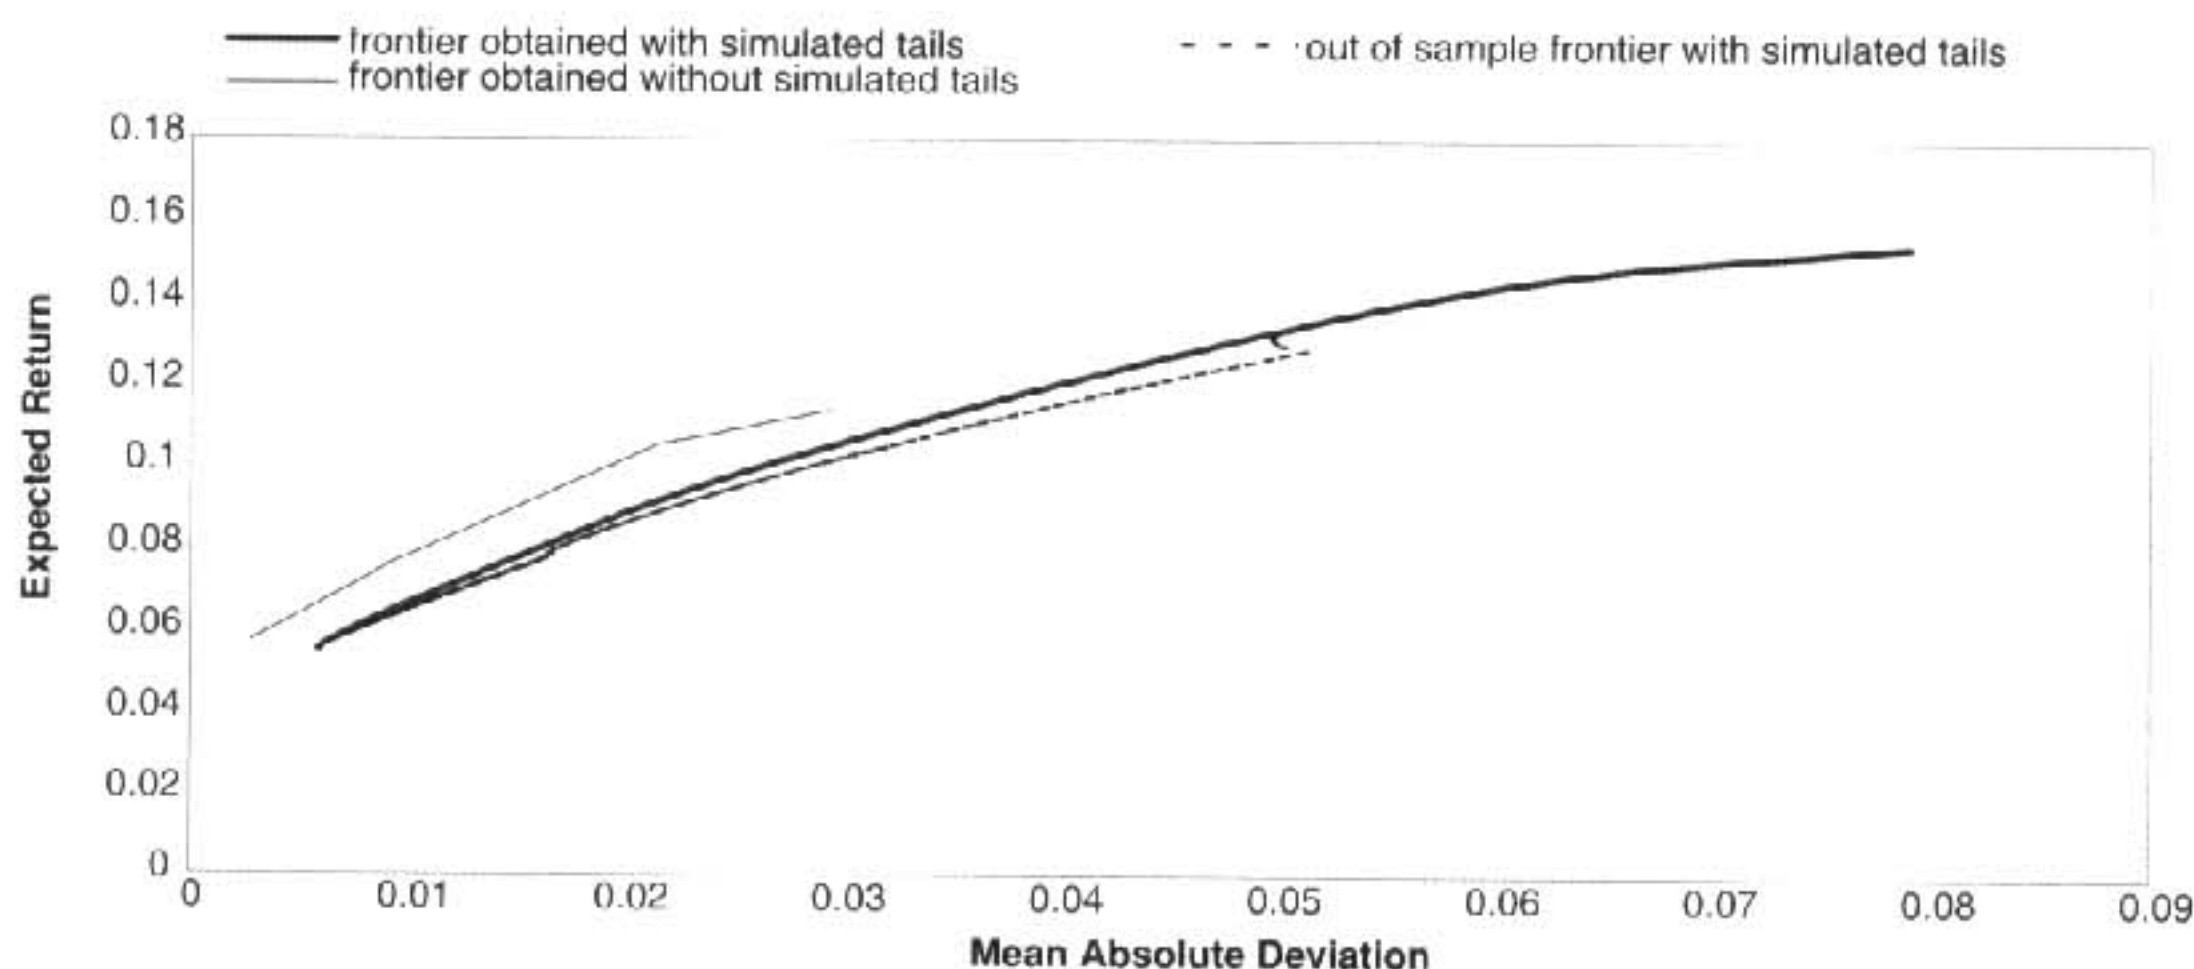
Fine line: Without simulating the tails due to credit events. Dashed line: Out-of-sample performance when credit events are included. Heavy line: Frontier traced using a MAD model with simulated data of credit events.

EXHIBIT 5
Distribution of Returns for Minimum-Risk Portfolio Obtained Using MAD Model
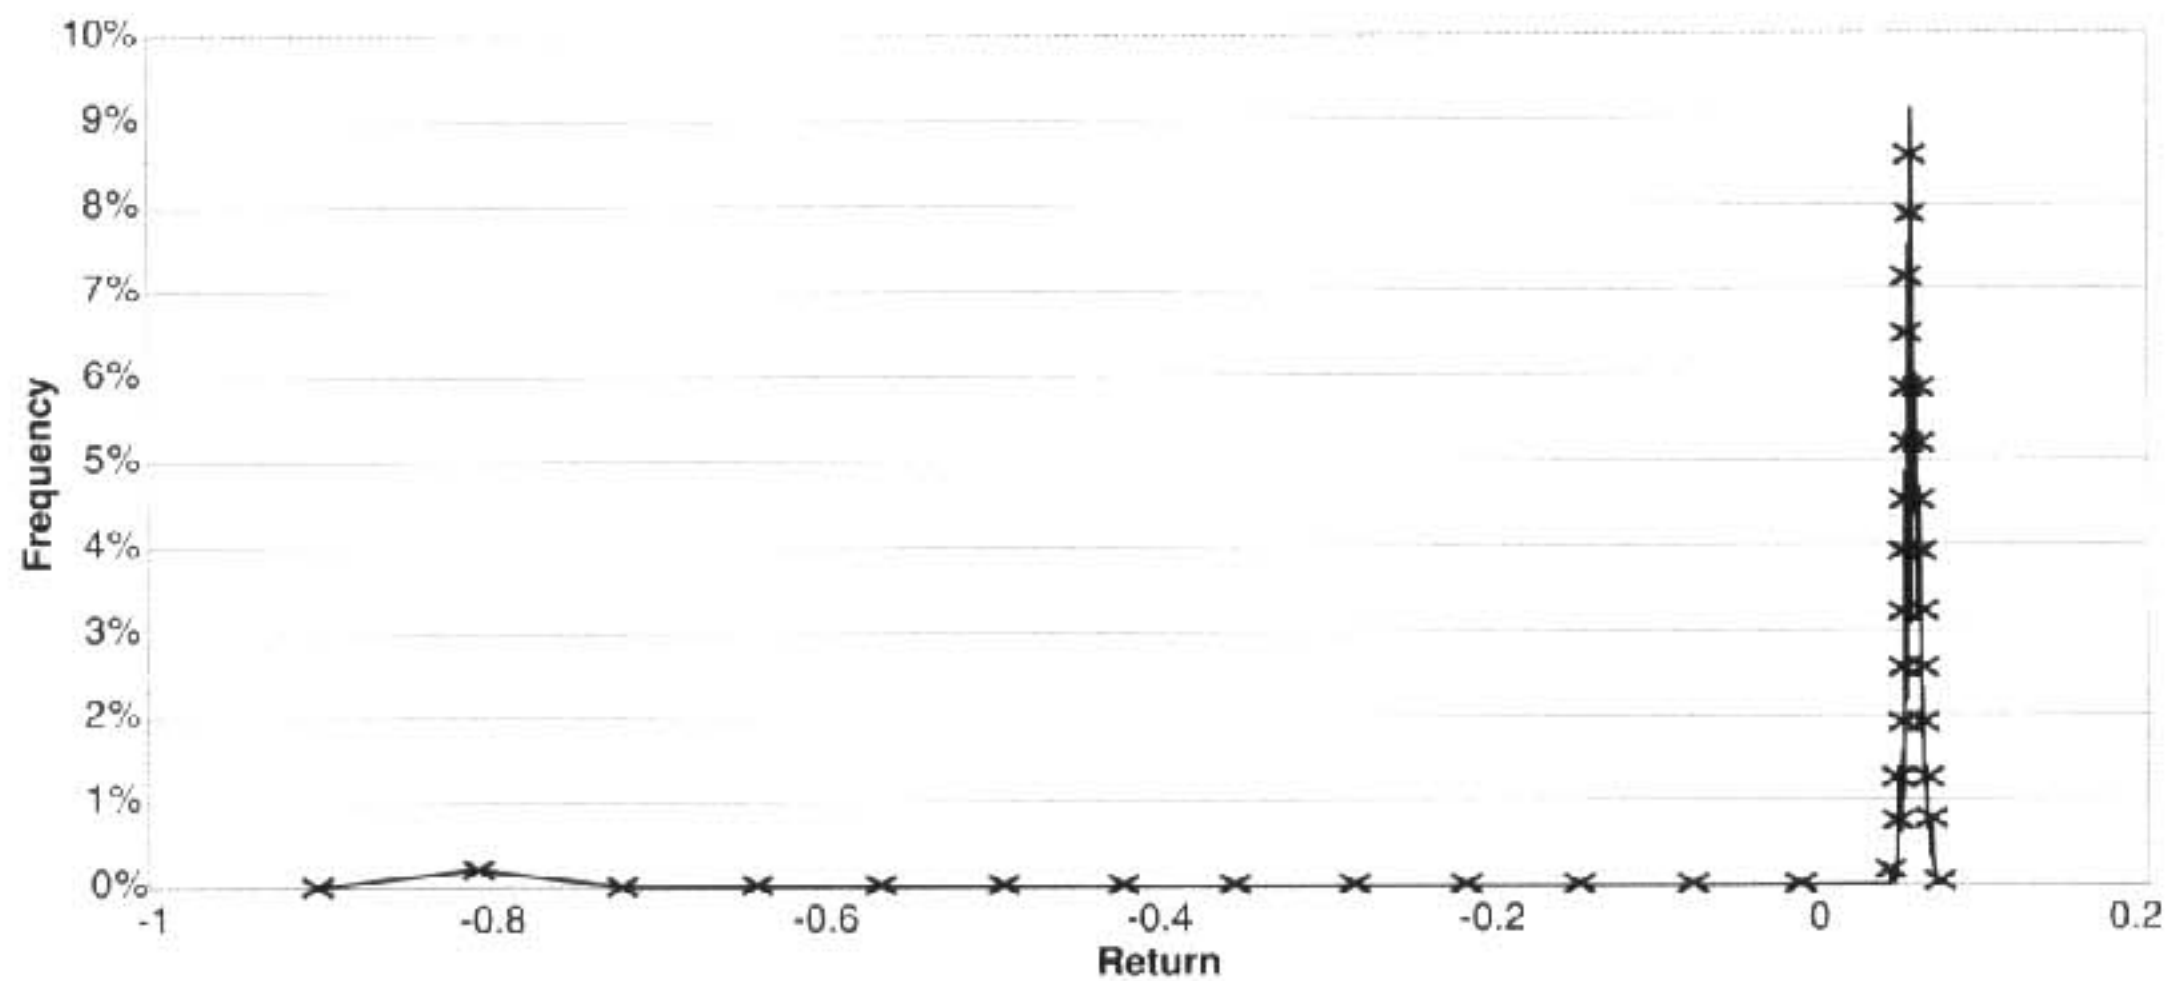
$$L(x, \widetilde{P}) = V_0 - V(x, \widetilde{P}) \qquad (1)$$

The relation between the loss function and portfolio returns is given by

$$L(x, \widetilde{r}) = -R(x, \widetilde{r})V_0 \qquad (2)$$

We assume a discrete scenario setting in which all random quantities take values from a finite and discrete scenario set indexed by members of a set Ω. That is, $\widetilde{P} \in \{P^l\}_{l \in \Omega}$, and the objective probability associated with each scenario $l \in \Omega$ is given by p^l. Under this assumption, the probability that the loss function does not exceed some threshold value ζ is given by the probability function:

$$\psi(x, \zeta) = \sum_{\{l \in \Omega | L(x, P^l) \leq \zeta\}} p^l \qquad (3)$$

The value at risk of the portfolios is then defined as follows:

Definition 1. *The value at risk (VaR) of a portfolio at the α probability level is the left α quantile of the losses of the portfolio, i.e., the lowest possible value so that the probability of losses less than VaR exceeds $\alpha \times 100\%$. It is given as:*

$$\text{VaR}(x, \alpha) \doteq \min \{\zeta \in \mathbb{R} \mid \psi(x, \zeta) > \alpha\} \qquad (4)$$

The quantile ζ is the left endpoint of the non-empty interval consisting of the values ζ so that $\Psi(x, \zeta) = \alpha$. The dependence of VaR on the confidence level α is sometimes made explicit by referring to α–VaR.

Exhibit 3 illustrates the VaR of the return for a credit risky portfolio of Baa bonds as 1.7% at the 0.95 probability level, and −1.7% at the 0.99 probability. There is a 1% chance of losses in excess of −1.7% and a 5% chance that returns will be less than 1.7%.

The VaR measure reveals nothing about the extent of the losses beyond the given confidence level. Such losses can be catastrophic. Long-Term Capital Management is a case in point. LTCM was estimated to have a VaR of only −5% at the 0.95 probability level, but a return of −80% in September 1998 wiped out a position of $1.85 trillion and threatened a global meltdown of the financial markets.

A measure of risk that goes beyond the information revealed by VaR is the expected value of the losses that exceed VaR. This quantity is called *expected shortfall, conditional loss,* or *conditional VaR*; see, e.g., Embrechts, Klüppelberg, and Mikosch [2000].

For general distributions, the conditional VaR is defined as a weighted average of VaR and the expected losses that are strictly greater than VaR. For discrete distributions, and under a mild technical condition that the probability of scenarios with losses strictly greater than VaR is exactly equal to $1 - \alpha$, i.e., $\psi(x, \zeta) = \alpha$, the following definition applies:

Definition 2. *The conditional value at risk (CVaR) of the losses of the portfolio is the expected value of the losses, conditioned on the losses being in excess of VaR:*

$$\text{CVaR}(x, \alpha) = \varepsilon[L(x, P^l) \mid L(x, P^l) > \zeta] \tag{5}$$

$$= \frac{\sum_{\{l \in \Omega \mid L(x,P^l) > \zeta\}} p^l L(x, P^l)}{\sum_{\{l \in \Omega \mid L(x,P^l) > \zeta\}} p^l} \tag{6}$$

$$= \frac{\sum_{\{l \in \Omega \mid L(x,P^l) > \zeta\}} p^l L(x, P^l)}{1 - \alpha} \tag{7}$$

where the last equality follows from the condition $\psi(x, \zeta) = \alpha$. The dependence of CVaR on the confidence level α is made explicit by referring to α–CVaR.

It follows from the definitions that CVaR is always greater than or equal to VaR. Both VaR and CVaR are functions of the asset allocation vector x and the percentile parameter α. It is natural to seek to minimize these measures by judiciously specifying the composition of the asset portfolio.

VaR is difficult to optimize when it is calculated using discrete scenarios. The VaR function is non-convex and non-smooth, and it has multiple local minima. CVaR, however, can be minimized using linear programming formulations; see Rockafellar and Uryasev [2000].

Consider the minimization of conditional value at risk given above by

$$\text{CVaR}(x, \alpha) = \frac{\sum_{\{l \in \Omega \mid L(x,P^l) > \zeta\}} p^l L(x, P^l)}{1 - \alpha} \tag{8}$$

This function can be expressed as a linear model with the use of auxiliary variables. Let:

$$y_+^l = \max [0, L(x, P^l) - \zeta], \text{ for all } l \in \Omega \tag{9}$$

y_+^l is equal to zero when the losses are less than or equal to the value at risk, ζ, and it is equal to the excess loss when the losses exceed ζ.

With this definition of y_+^l, we write:

$$\sum_{l \in \Omega} p^l y_+^l = \sum_{\{l \in \Omega \mid L(x,P^l) \le \zeta\}} p^l y_+^l + \sum_{\{l \in \Omega \mid L(x,P^l) > \zeta\}} p^l y_+^l \tag{9.1}$$

$$= 0 + \sum_{\{l \in \Omega \mid L(x,P^l) > \zeta\}} p^l \left(L(x, P^l) - \zeta \right)$$

$$= \sum_{\{l \in \Omega \mid L(x,P^l) > \zeta\}} p^l L(x, P^l) - \zeta \sum_{\{l \in \Omega \mid L(x,P^l) > \zeta\}} p^l$$

$$= \sum_{\{l \in \Omega \mid L(x,P^l) > \zeta\}} p^l L(x, P^l) - \zeta(1 - \alpha)$$

Dividing both sides by $(1 - \alpha)$ and rearranging terms, we get

$$\zeta + \frac{\sum_{l \in \Omega} p^l y_+^l}{1 - \alpha} = \frac{\sum_{\{l \in \Omega \mid L(x,P^l) > \zeta\}} p^l L(x, P^l)}{1 - \alpha} \tag{10}$$

The term on the right-hand side is $\text{CVaR}(x, \alpha)$ of Equation (8). It can be optimized using linear programming to minimize the term on the left-hand side.

We minimize CVaR subject to constraints on the asset allocation of the form $x \in X$, where X denotes the set of feasible solutions, and the condition that the expected value of the portfolio exceeds some target μ. Using the equivalent definition of CVaR from (10), we write the model as follows:

$$\underset{x \in X}{\text{Minimize}} \quad \zeta + \frac{\sum_{l \in \Omega} p^l y_+^l}{1 - \alpha} \tag{11}$$

$$\text{subject to} \quad \sum_{i=1}^{m} \bar{P}_i x_i \ge \mu \tag{12}$$

$$y_+^l \ge L(x, P^l) - \zeta \text{ for all } l \in \Omega \tag{13}$$

$$y_+^l \ge 0 \qquad \text{for all } l \in \Omega \tag{14}$$

$(\bar{P} = \sum_{l \in \Omega} p^l P_i^l$ is the expected value of the price of asset i.) Since the loss function $L(x, P^l)$ is linear [see Equation (1)], the model is linear (see the appendix for further details). The solution to this model gives us the minimum CVaR^* for a given target expected value μ, and the VaR value ζ^* corresponding to the minimum CVaR portfolio. (Recall that $\text{CVaR} \ge \text{VaR}$ and, hence, $\text{CVaR}^* \ge \zeta^*$.)

An efficient frontier trading off expected shortfall against expected portfolio value is traced by varying the parameter μ. We develop the CVaR efficient frontier of portfolios of Baa bonds at the 0.99 and 0.95 probability levels (see Exhibit 6). Exhibit 6 also plots the trade-offs

between CVaR and the expected return of the optimal portfolios obtained using a MAD model. That is, we take the portfolios of the efficient frontier of Exhibit 4 and calculate their CVaR.

Note from the two frontiers in Exhibit 6 (first graph) that there is nothing efficient about the MAD optimized portfolios when using a risk metric that properly accounts for the tails of the optimized portfolios. It is not sufficient to capture the tails in the simulation phase, as we do in Exhibit 4. We must also optimize the appropriate risk metric, as in Exhibit 6.

In other words, to avoid distortions of the efficient frontier due to the tail events, we need to optimize a risk metric that appropriately penalizes the tails. CVaR provides such a risk metric. Note, however, from the second graph in Exhibit 6 that the distortions of the frontiers calculated at the 0.95 probability level are barely noticeable for a wide range of target portfolio values, while the differences are distinct for a wider range of target values at the 0.99

probability level. We reemphasize that tail effects can be captured only with adequate accuracy of the models.

The Tail that Wags the Dog

The risk profile of a portfolio is shaped by the attention the risk manager pays to the tails. Taking a CVaR perspective on risk management substantially reduces the tails. Exhibit 7 shows the distribution of returns of the minimum risk portfolio obtained when minimizing CVaR at probability levels 0.95 and 0.99. The tail extends up to −35% when minimizing CVaR at the 0.95 probability level, but it shrinks to −10% when minimizing CVaR at the 0.99 probability level. Both of these losses are substantially smaller than the losses in excess of 80% realized when minimizing the mean absolute deviation measure of risk as demonstrated in Exhibit 5.

Of course the choice of a risk metric has an effect on the upside potential of the portfolio. We can see in the

E X H I B I T 6
Frontier Generated Using CVaR Portfolio Optimization
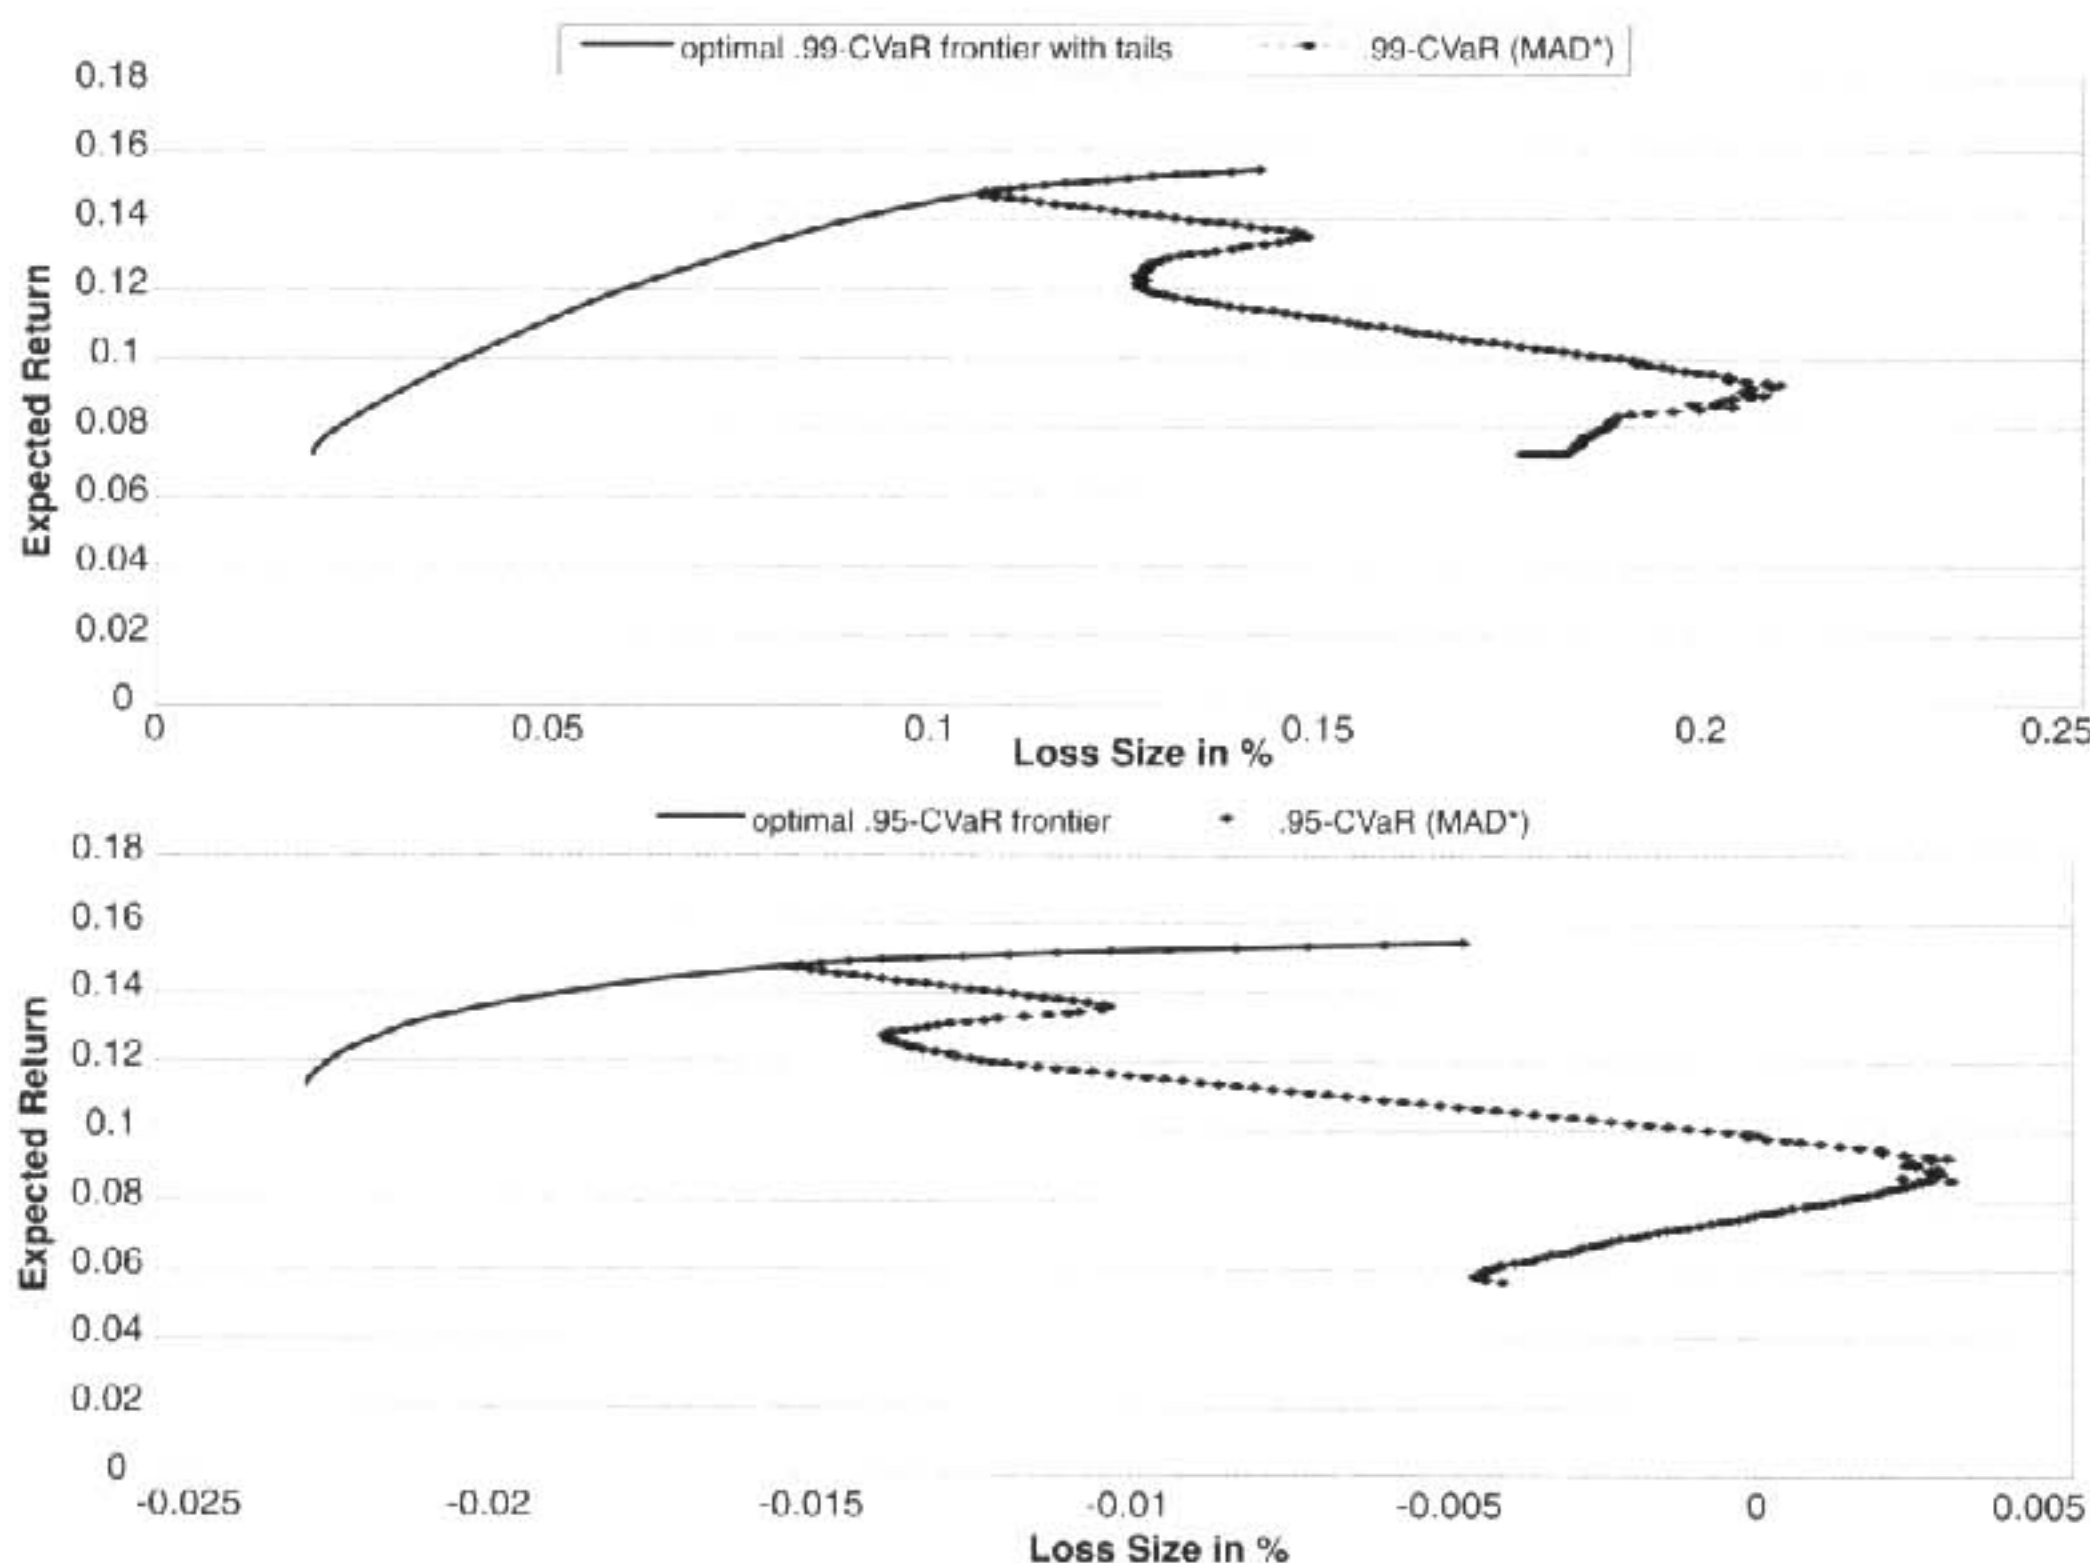
Solid line: With credit event simulations. Dashed line: CVaR of efficient portfolios optimized with MAD model.
First graph: CVaR estimated at 0.99 probability level. Second graph: CVaR estimated at 0.95 probability level.

distributions of Exhibits 5 and 7 that the upside potential is reduced as the tails are shrunk. There are the usual trade-offs between upside potential and downside risk, but in the context of credit risky securities the downside risk is hidden in the tail and not in the variance or the mean absolute deviation. In this respect CVaR has an important role to play in tracing efficient frontiers for the management of credit risk.

To VaR or to CVaR?

There is some debate among both academicians and practitioners as to whether VaR or CVaR is the appropriate metric for risk management applications. VaR clearly has an advantage in the practice of risk measurement, where it is considered the industry standard; see, e.g., Jorion [1996]. CVaR, on the other hand, appears to be the metric of choice in the insurance industry; see, e.g., Embrechts, Klüppelberg, and Mikosch [2000]. Axiomatic characterizations of risk metrics—the notion of *coherence* suggested by Artzner et al. [1999]—favor CVaR, which is coherent, over VaR, which is not coherent.

Notions of coherence notwithstanding, the fact that CVaR provides a bound for VaR, as well as the increasing acceptance of VaR estimates by regulators, has somewhat shadowed the debate. Exhibit 8 shows the estimated VaR of CVaR-optimized portfolios obtained with and without simulations of the tails. As expected, CVaR provides an upper bound for VaR, although the bound need not be tight, especially when the tails are properly simulated. Furthermore, the frontier of VaR against expected returns for the CVaR-optimized portfolios need not be efficient. The distortion of the VaR frontier is pronounced when the tails are properly simulated.

The results in Exhibit 8 clearly make the point that the choice between VaR and CVaR has significant ramifications in the risk management of credit risky portfolios. Given the flat tails witnessed in our simulations, and

---

EXHIBIT 7
Distribution of Returns for Minimum-Risk Portfolio
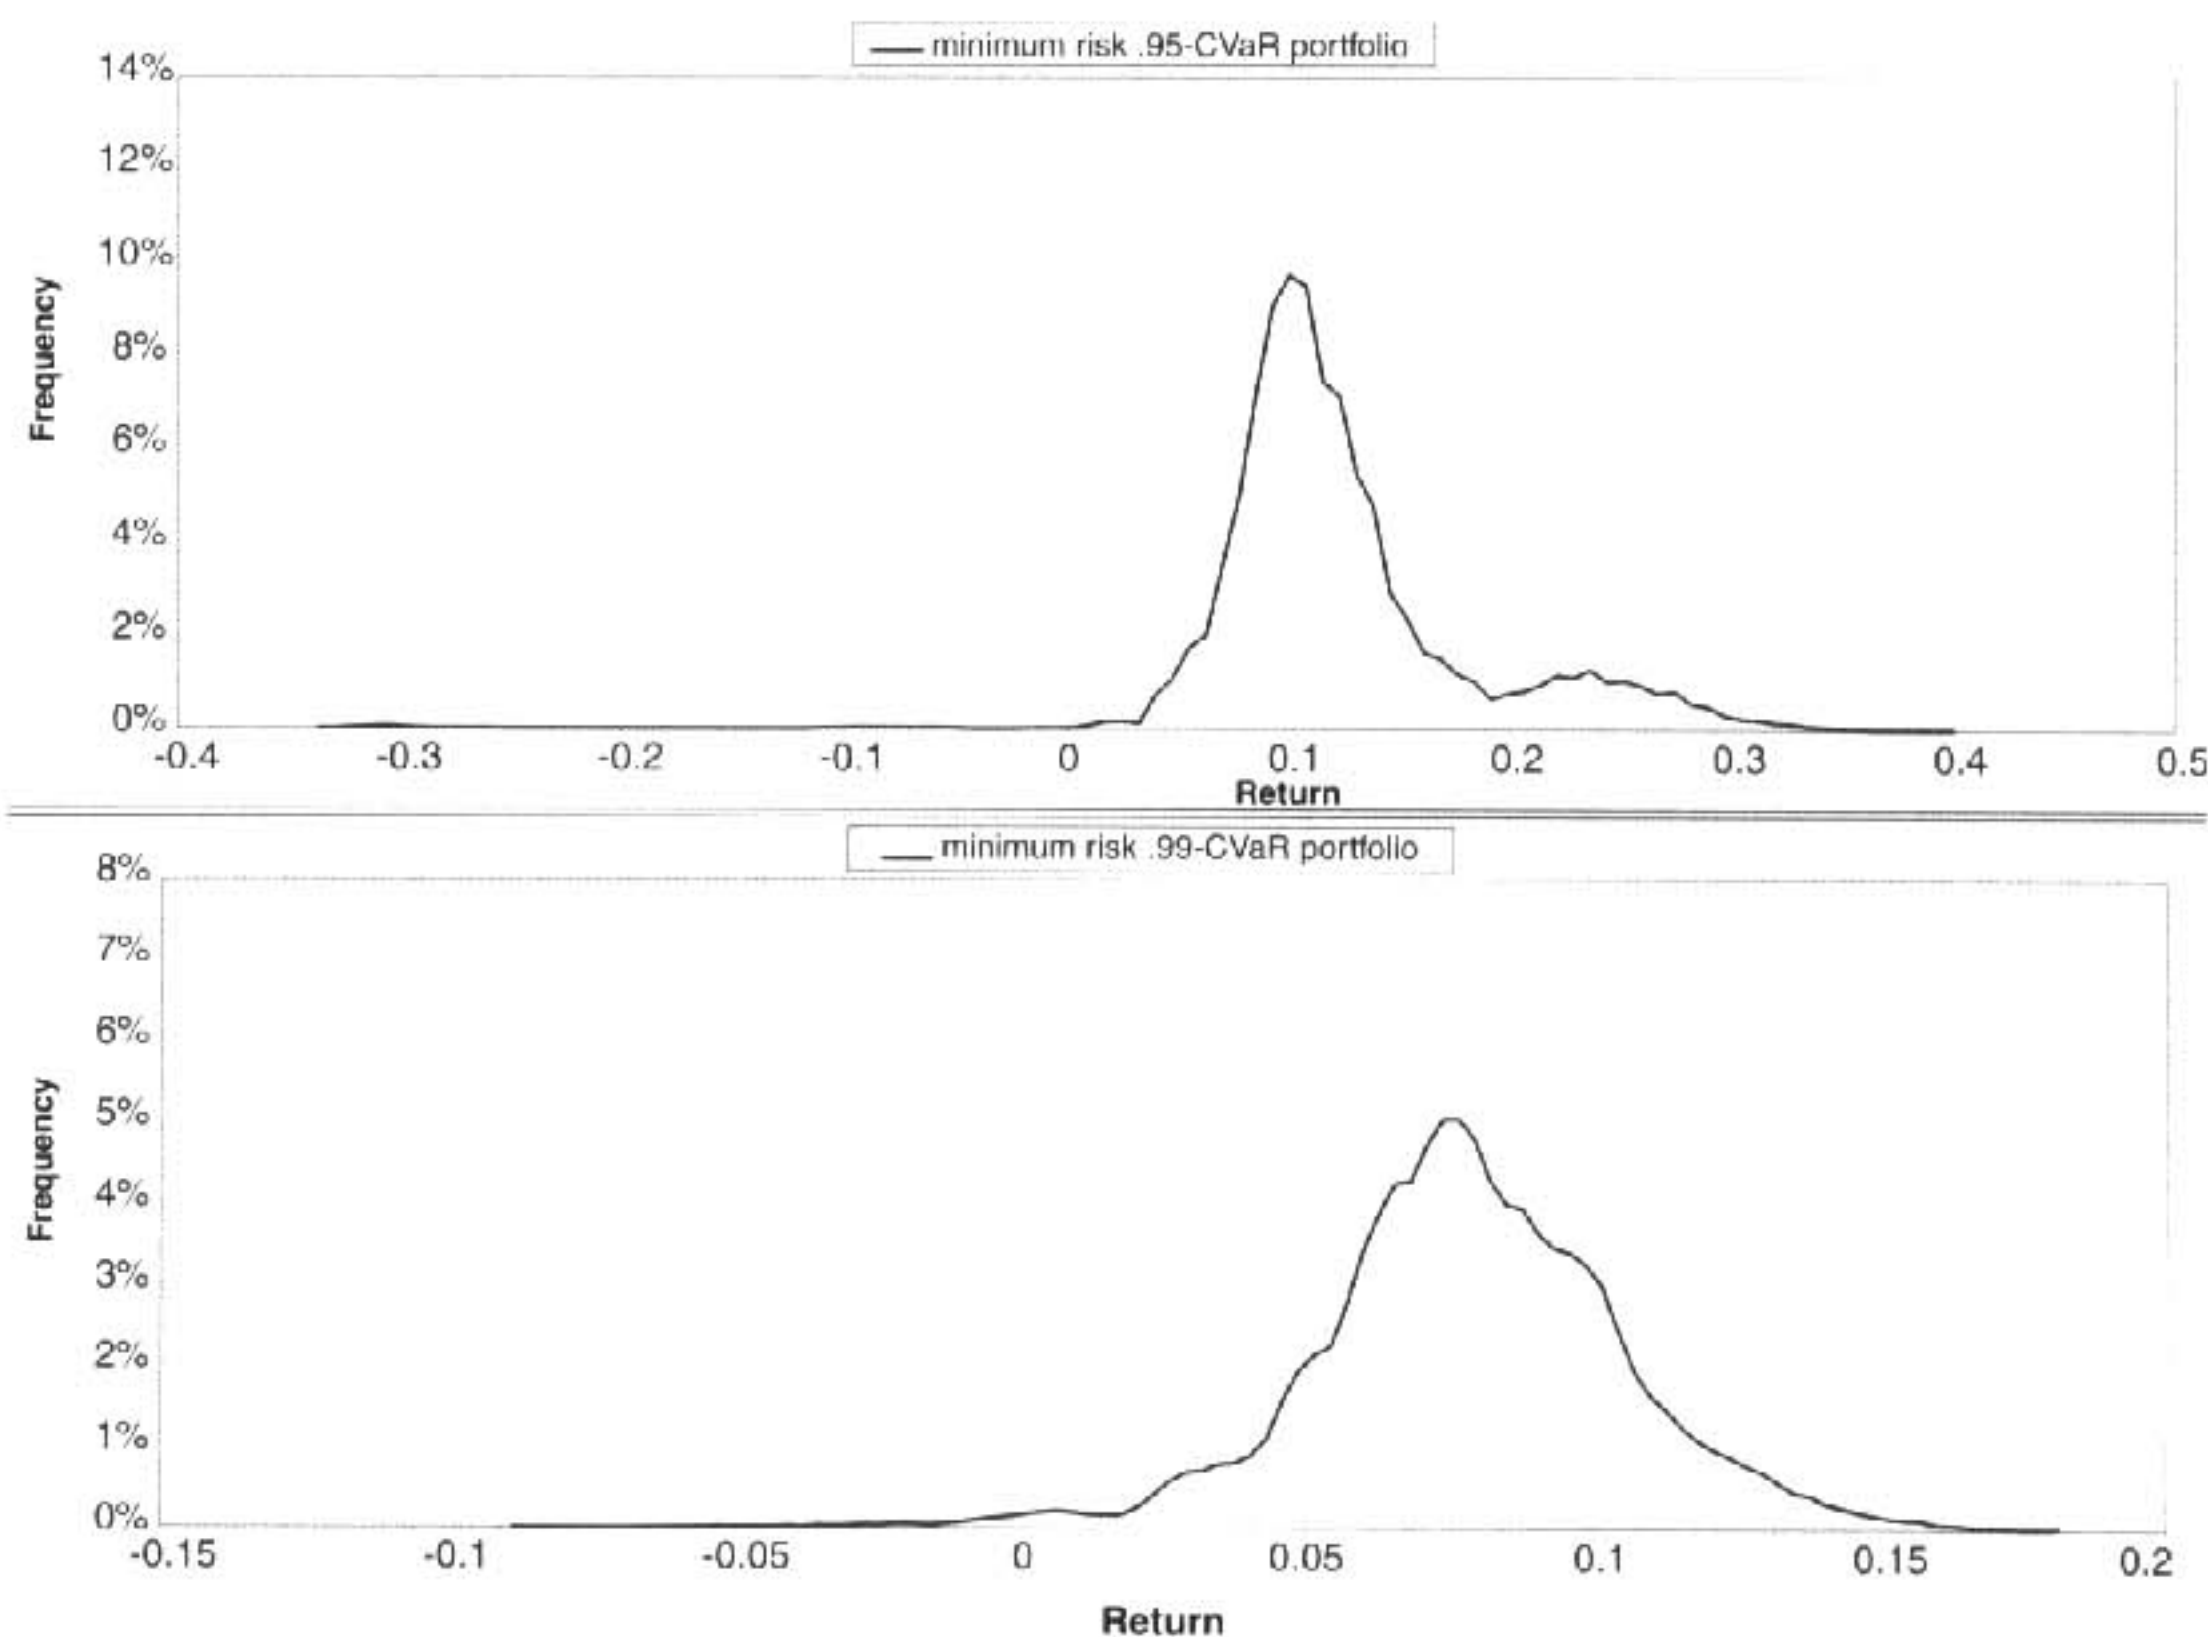
First graph: Using CVaR model at 0.95 probability level.
Second graph: Using CVaR model at 0.99 probability level.

the coherence properties of CVaR, we argue that CVaR optimization provides the appropriate risk management framework for credit risky portfolios. The adoption of CVaR criteria for credit risk management by Andersson et al. [2001] is well justified, although their model does not include all the sources of risk incorporated in the simulations of Jobst and Zenios [2001].

LONG-TERM PERFORMANCE WITH SHORT-TERM TAILS

Optimization of portfolio performance for the long run ignores the short-term effects. This has been the tradition in myopic single-period optimization models. Ignoring the short-term effects can be catastrophic in the presence of tails. In particular, the long-term (expected) potential of a portfolio strategy may never be realized if an extreme event in the short run results in bankruptcy.

Long-Term Capital Management is a case in point. When LTCM suffered losses of 80% in September 1998, the New York Federal Reserve orchestrated a bailout. Fourteen banks invested $3.6 billion in return for a 90% stake in the firm. The fund eventually recovered its losses and posted positive returns, but the original stakeholders were not there any more.

We look at the short-term effects of the tails on the portfolios obtained with the MAD and the 0.99-CVaR optimization models. We take the minimum-risk portfolio from the efficient frontiers of both models at a 12-month risk horizon, and simulate with out-of-sample scenarios the distributions of returns at months 3, 6, and 9. The results for the portfolios obtained by the MAD models are shown in Exhibit 9. Results for the CVaR model are in Exhibit 10.

The expected return of the MAD optimized portfolio over the 12-month period is 5.5%. The worst case losses are on the order of 2% in the first three months, but

EXHIBIT 8
CVaR Efficient Portfolios and Estimated VaR
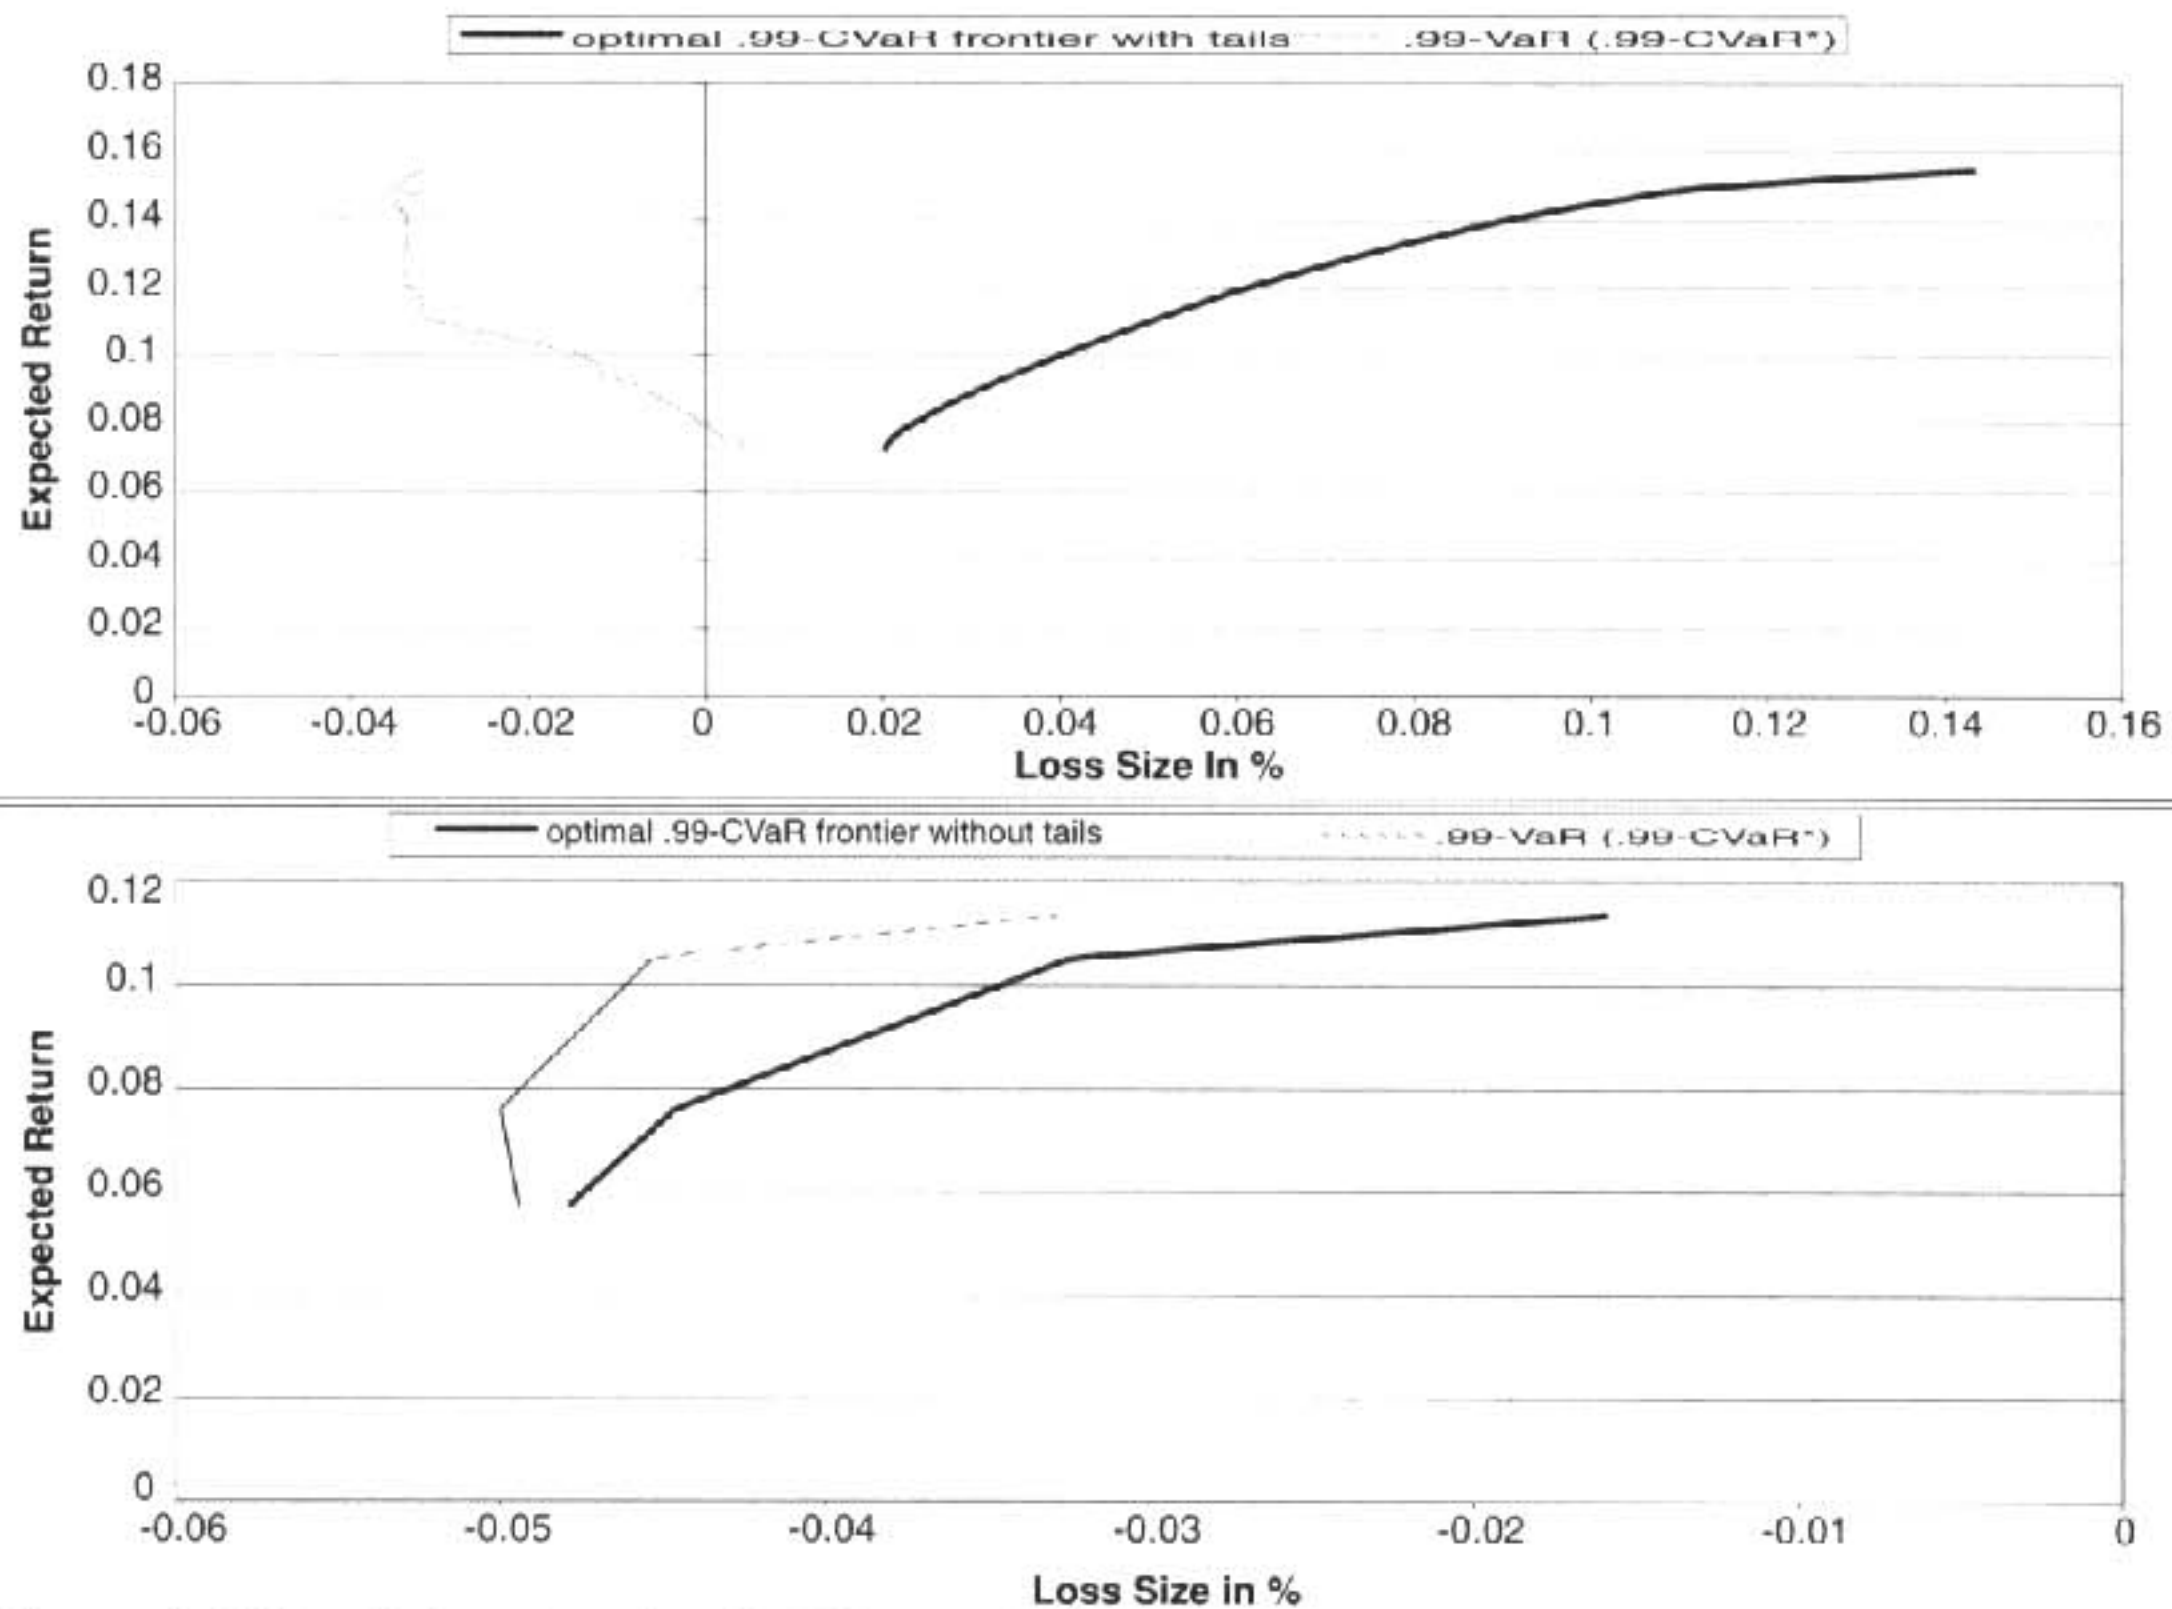
First graph: With tail effects. Second graph: Without tail effects.

Solid lines: CVAR efficient portfolios. Dashed lines: Estimated VaR.

CVaR estimated at 0.99 probability level.

they jump to 80% at months 6, 9, and 12. The probability of these losses also increases with time—from 0.17% at month 6 to 0.34% at month 12. Although the probabilities are small, these losses are potentially catastrophic. While the expected return of the MAD optimized portfolio increases with time, so does the probability of a catastrophic event that will prevent the portfolio path from obtaining the long-run terminal value.

The expected return of the CVaR optimized port-folio over the 12-month period is 7.2% (Exhibit 10). The worst case losses of this portfolio remain at around 10% throughout this time period. The probability of losses increases marginally from month 3 to month 12. In any event, these losses are not catastrophic, and the long-term expected return can be achieved. The CVaR optimized portfolio not only has higher expected return than the MAD optimized portfolio in the long run, but it also has better downside risk profile in the short run.

EXHIBIT 9
### Distribution of Returns of Minimum-Risk MAD Optimized Portfolio
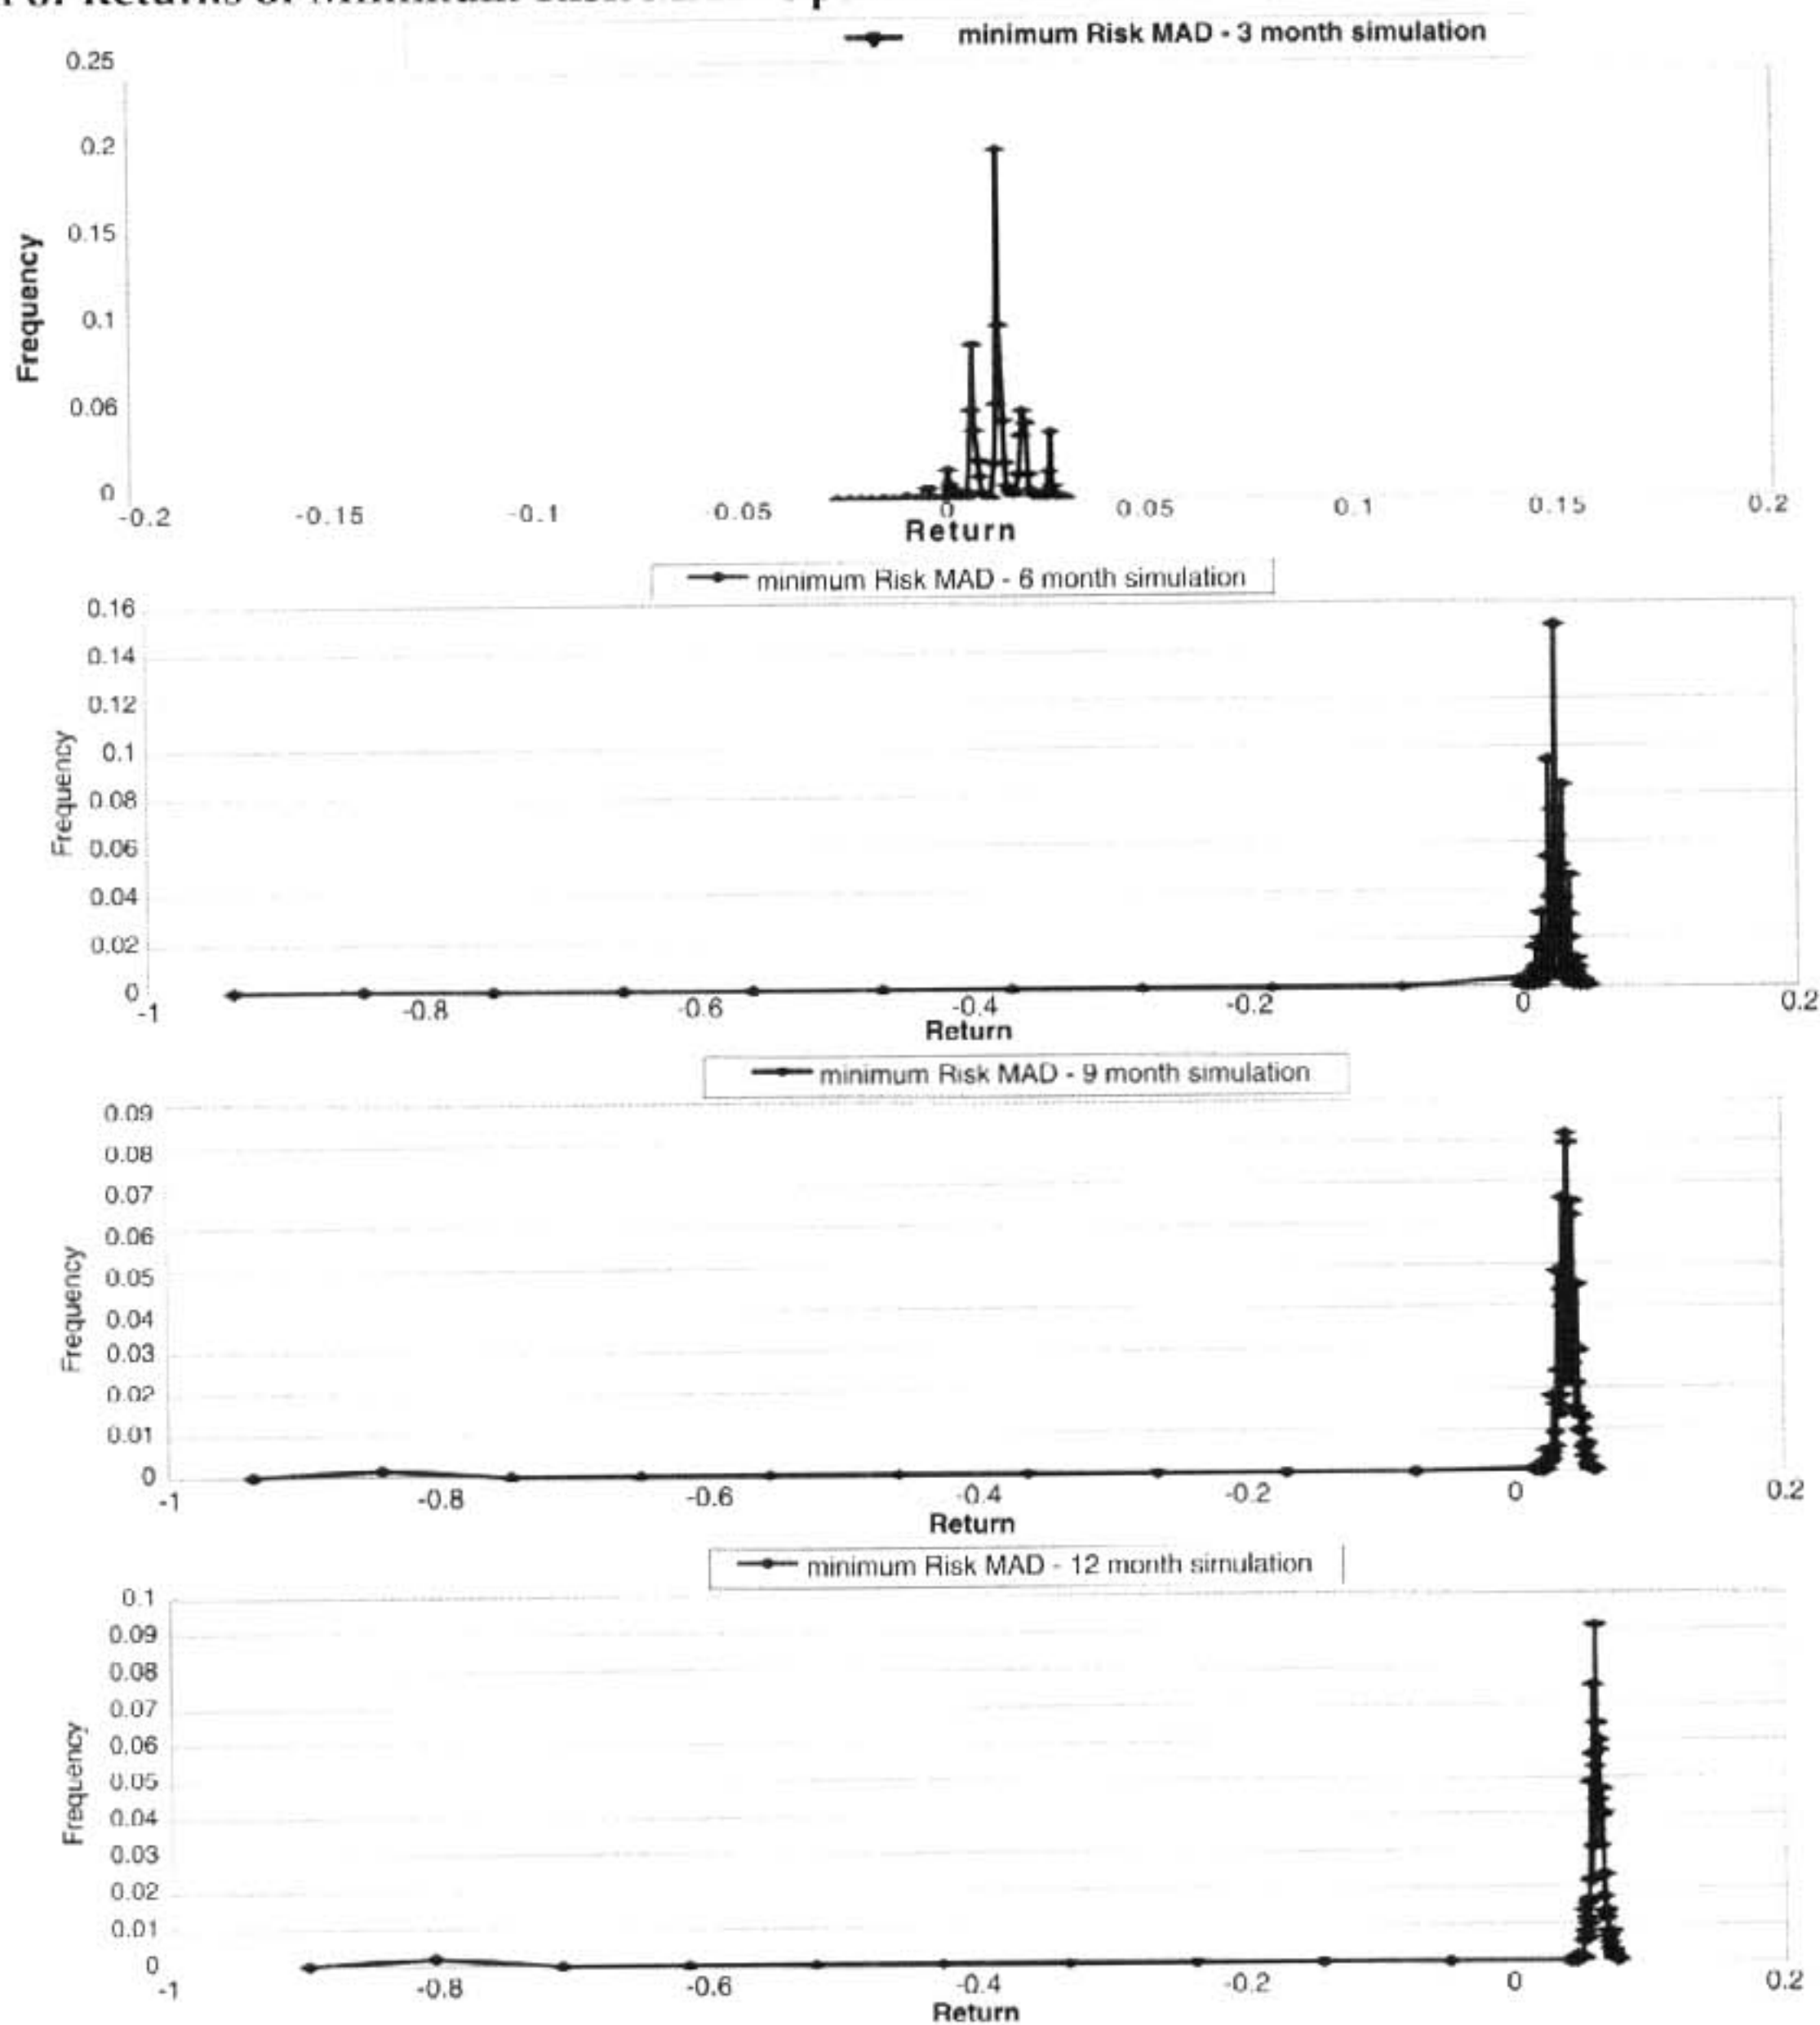
The impact of the tails when compounded through time appears to be devastating, especially in the MAD portfolio (assuming that investors will not rebalance). If one uses reduced-form models of default such as those discussed in Jobst and Zenios [2001], credit events are typically unpredictable. This implies that investors will not be able to rebalance prior to a credit event. In fact, there is strong evidence that credit events may in part be predictable (for example, from equity returns), and such information can be included in the reduced-form framework (see Schön-bucher and Schubert [2001]). This implies that investors would use these signals and reduce their exposure.

Multiperiod optimization models that incorporate new information are currently available in the framework of stochastic programming and have been successfully applied in financial planning (see Kouwenberg and Zenios [2001]). Using stochastic programming to improve further the performance of the models introduced here is a promising direction for further research.

EXHIBIT 10
Distribution of Returns of Minimum-Risk CVaR Optimized Portfolio
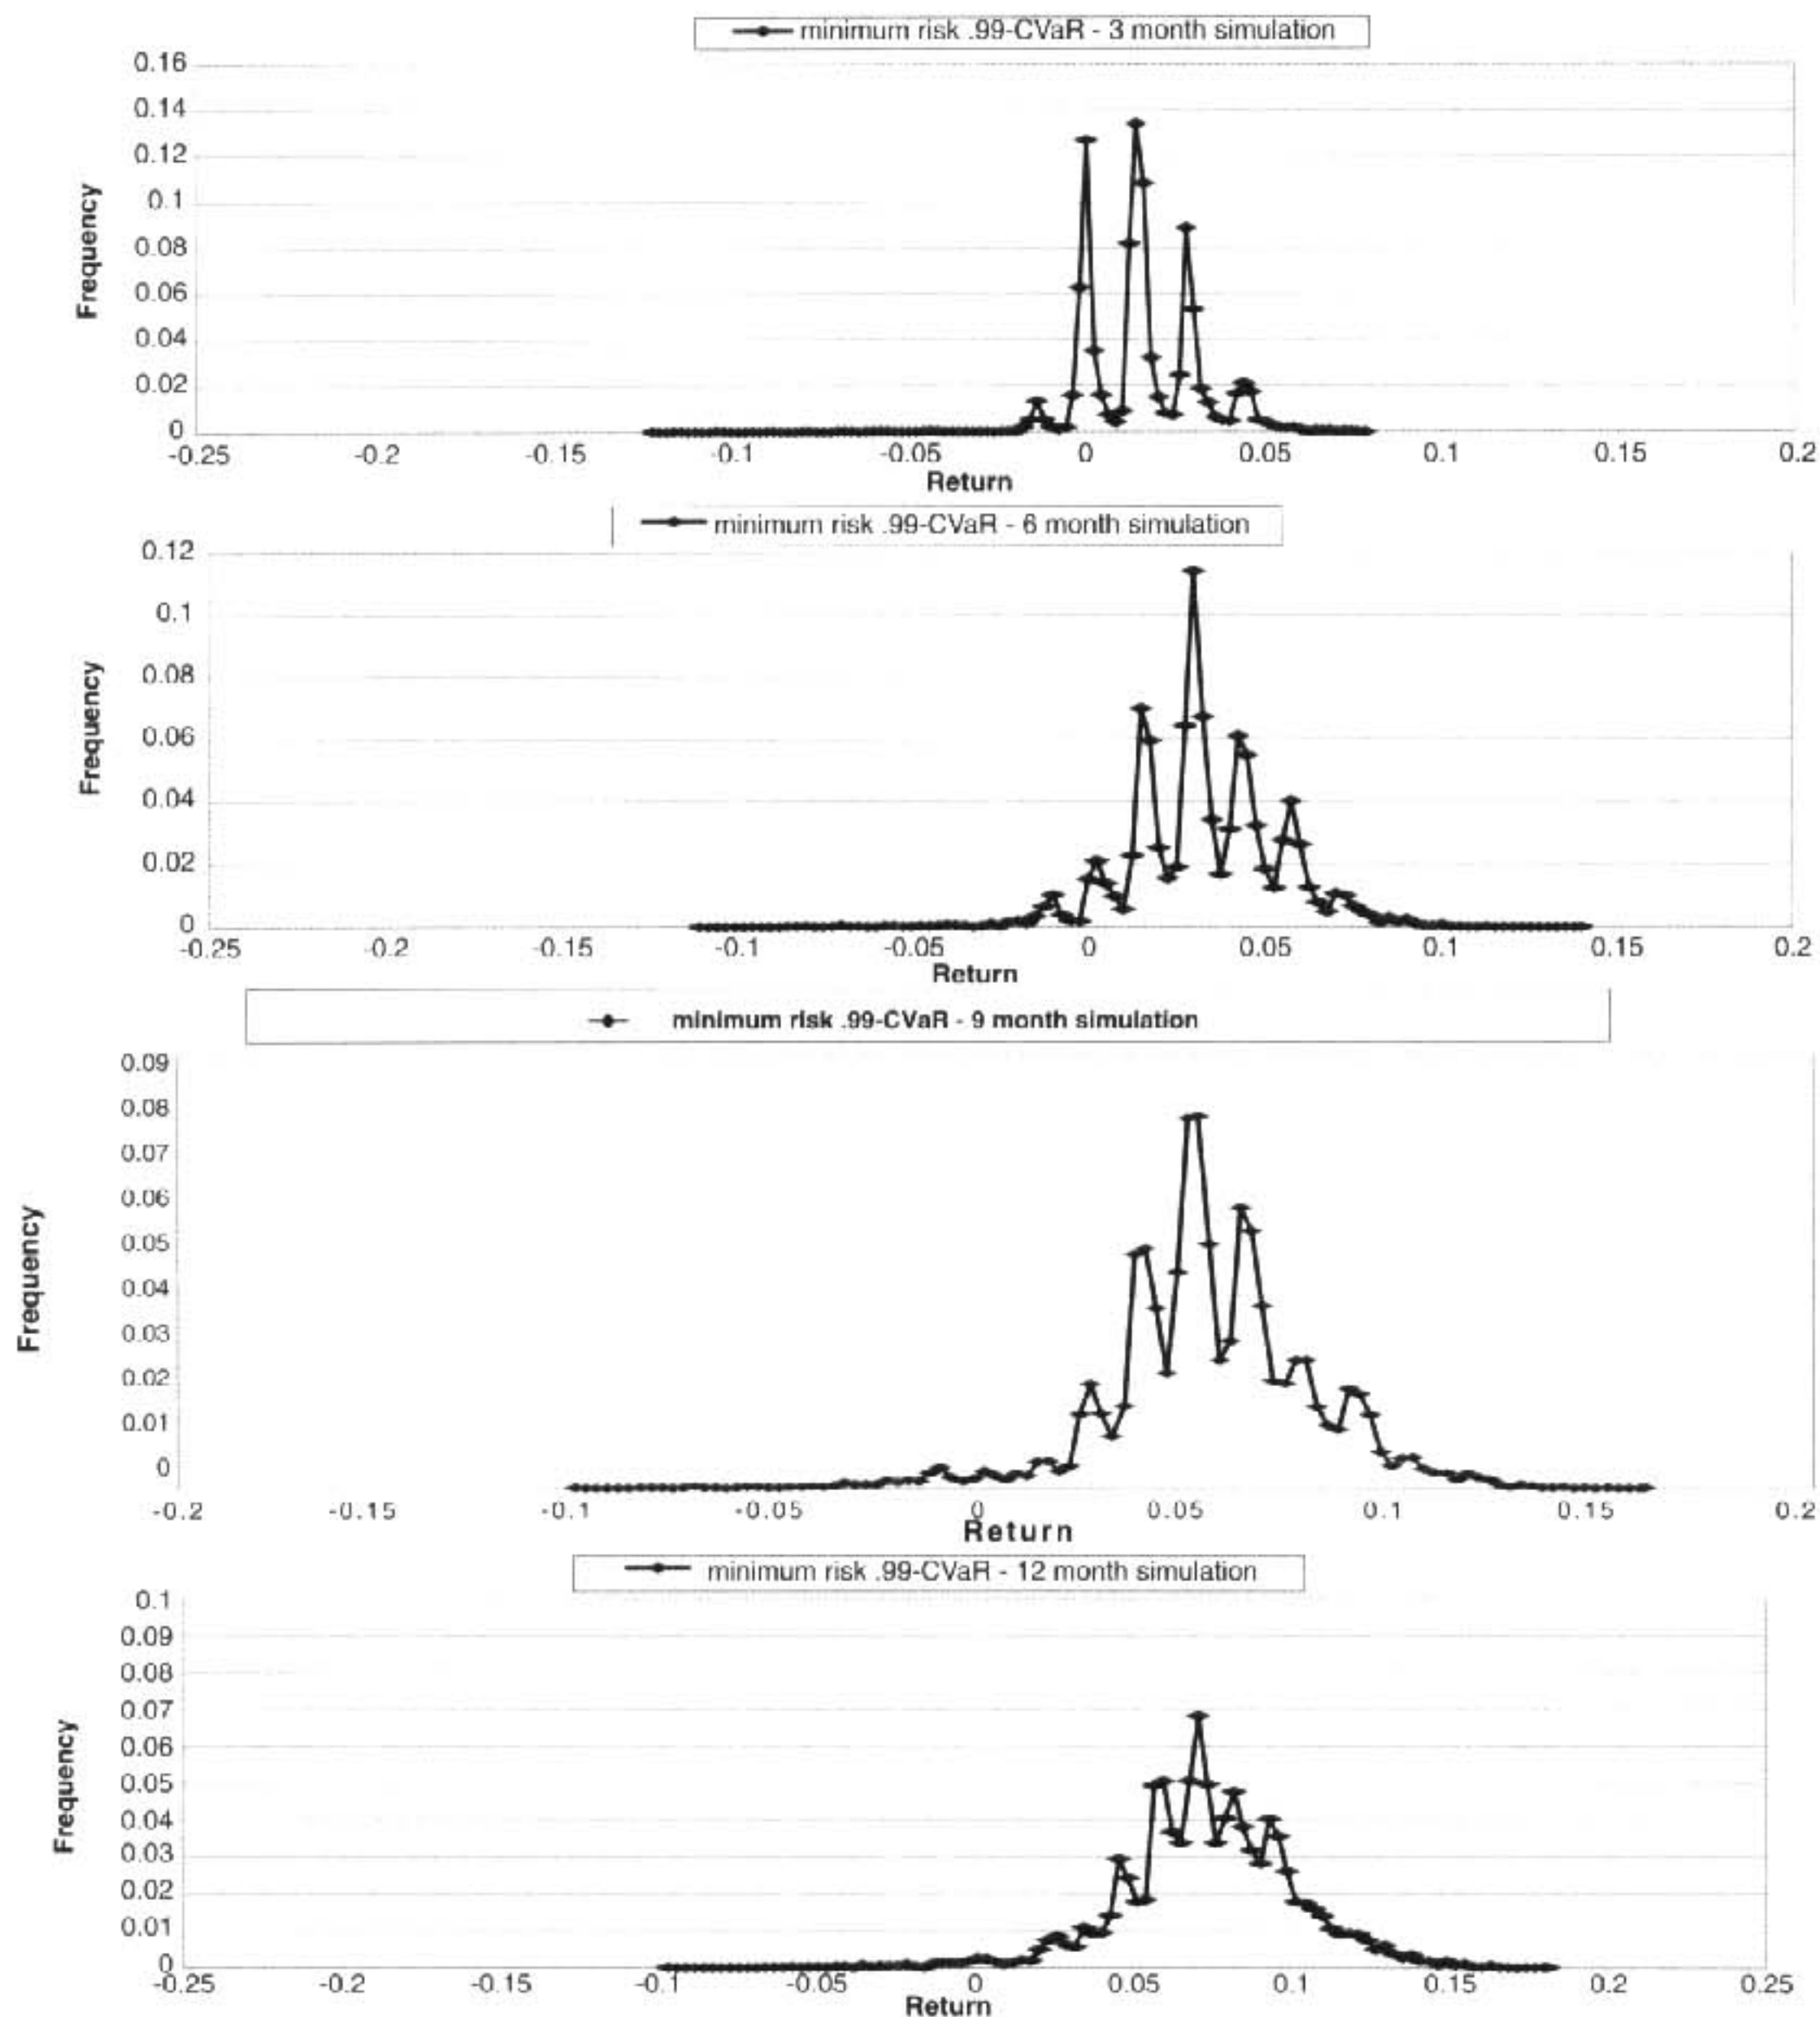
CONCLUDING REMARKS

We have highlighted the pitfalls of ignoring the low-probability costly events of default while integrating credit risky assets into portfolios. The shape of the risk profile of the portfolio is affected substantially by the tails caused by credit events. In order to efficiently tradeoff long-term expected return with risks, we must recognize that the risks are in the tails. The extreme credit events must be properly simulated with sufficient accuracy.

Furthermore, a risk metric must be chosen that accurately captures the impact of the tails. Conditional value-at-risk provides such a metric. Its use ensures that long-term goals can be met, without suffering catastrophic blows from the tails in the short run. The models we develop for limiting tail effects should also be applicable to portfolios of assets with correlated defaults, such as collateralized loan or debt obligations.

A P P E N D I X
Linear Programming Model

We can reformulate Equations (11)–(14) as a standard linear program suitable for implementation using any linear programming environment available such as the linear programming solver in EXCEL or any other mathematical programming system.[2]

Variables:

x_i: Holdings of asset i;

y_+^l: Auxiliary variable under scenario l of the index set of scenarios Ω; and

ζ: Value at Risk at a given confidence level α.

Parameters:

V_0: Initial portfolio value (set equal to one in our experiments);

P_i^0: Initial price of asset i;

$\bar{r}_i$: Expected return of asset i;

r_i^l: Return of asset i under scenario l;

p^l: Probability of scenario l;

μ: Target portfolio return;

α: Confidence level for CVaR calculation.

$$\underset{x \in X}{\text{Minimize}} \quad \zeta + \frac{1}{1-\alpha}\sum_{l \in \Omega} p^l y_+^l \tag{A-1}$$

$$\text{subject to} \quad \sum_{i=1}^{m} P_i^0 x_i = V_0 \tag{A-2}$$

$$\sum_{i=1}^{m} P_i^0 (1 + \bar{r}_i) x_i \geq (1 + \mu) V_0 \tag{A-3}$$

$$y_+^l \geq - \sum_{i=1}^{m} x_i \, r_i^l \, P_i^0 - \zeta \quad \text{for all } l \in \Omega \tag{A-4}$$

$$y_+^l \geq 0 \qquad \text{for all } l \in \Omega \tag{A-5}$$

$$x_i \geq 0 \qquad \text{for all } i = 1... \, m \tag{A-6}$$

ENDNOTES

[1]To carry out the 2.5 million simulation runs on a portfolio of approximately 600 securities takes on the order of two hours on a Pentium 800Mhz machine. The application could be significantly sped up by using closed-form solutions for bond pricing instead of the more general tree implementation used in Jobst and Zenios [2001]. Additional efficiency gains could be achieved using distributed networks of workstations (see Cagan, Carriero, and Zenios [1993]), which would improve real time performance.

[2]Optimization models of the size reported here would be solved in one to two minutes of computing time on a Pentium 800Mhz.

REFERENCES

Andersson, F., H. Mauser, D. Rosen, and S. Uryasev. "Credit Risk Optimization with Conditional Value-at-Risk Criterion." *Mathematical Programming*, Series B, 89 (2001), pp. 273–291.

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

Bodnar, G.M., G.S. Hayt, and R.C. Marston. "1998 Wharton Survey of Financial Risk Management by U.S. Non-Financial Firms." *Financial Management*, 27 (4) (1998), pp. 70–91.

Cagan, L.D., J. Carriero, and S.A. Zenios. "A Computer Network Approach to Pricing Mortgage-Backed Securities." *Financial Analysts Journal*, March/April, 1993.

Embrechts, P., C. Klüppelberg, and T. Mikosch. "Modeling Extremal Events for Insurance and Finance." Volume 33 of *Applications of Mathematics, Stochastic Modeling and Applied Probability*. Heidelberg: Springer-Verlag, 2000.

Jobst, N.J., and S.A. Zenios. "Extending Credit Risk (Pricing) Models for Simulation and Valuation of Portfolios of Interest Rate and Credit Risk Sensitive Securities." Working paper 01-03, HERMES Center on Computational Finance and Economics, University of Cyprus, Nicosia, CY, 2001.

Jorion, P. "Value at Risk: The New Benchmark for Controlling Market Risk." Burr Ridge, IL: Irwin Professional Publishing, (1996).

Konno, H., and H. Yamazaki. "Mean-Absolute Deviation Portfolio Optimization Model and Its Applications to Tokyo Stock Market." *Management Science*, 37(5) (1991), pp. 519–531.

Kouwenberg, R., and S.A. Zenios. "Stochastic Programming Models for Asset Liability Management." Working paper 01-01, HERMES Center on Computational Finance and Economics, University of Cyprus, Nicosia, CY, 2001.

Leland, H.E. "Beyond Mean Variance: Performance Measurement in a Nonsymmetrical World." *Financial Analysts Journal*, 55(1), January/February 1999, pp. 27–36.

Rockafellar, R.T., and S. Uryasev. "Optimization of Conditional Value-at-Risk." *The Journal of Risk*, 2(3) (2000), pp. 21–41.

Saunders, A. *Credit Risk Measurement: New Approaches to Value at Risk and Other Paradigms*. New York: John Wiley and Sons, 1999.

Schönbucher, P. *Credit Risk Modeling and Credit Derivatives*. Ph.D. thesis, Rheinische Friedrich-Wilhelms-Universität, Bonn, 2000.

Schönbucher, P., and D. Schubert. "Copula-Dependent Default Risk in Intensity Models." Working paper, Department of Statistics, Bonn University, 2001.

Stulz, R. "Rethinking Risk Management." *Journal of Applied Corporate Finance*, 9(8) (1996), p. 24.