



ELSEVIER

Journal of Empirical Finance 7 (2000) 247–269

Journal of
EMPIRICAL
FINANCE

www.elsevier.com/locate/econbase

Portfolio selection with limited downside risk

Dennis W. Jansen^{a,*}, Kees G. Koedijk^b, Casper G. de Vries^c

^a Department of Economics, Texas A&M University, College Station, TX 77843-4228, USA

^b Maastricht University and L.I.F.E., Maastricht, Netherlands

^c Erasmus University Rotterdam and Tinbergen Institute, Rotterdam, Netherlands

Abstract

A safety-first investor maximizes expected return subject to a downside risk constraint. Arzac and Bawa [Arzac, E.R., Bawa, V.S., 1977. Portfolio choice and equilibrium in capital markets with safety-first investors. *Journal of Financial Economics* 4, 277–288.] use the Value at Risk as the downside risk measure. The paper by Gouriéroux, Laurent and Scaillet estimates the optimal safety-first portfolio by a kernel-based method, we exploit the fact that returns are fat-tailed, and propose a semi-parametric method for modeling tail events. We also analyze a portfolio containing the two stocks used by Gouriéroux et al. and discuss the merits of the safety-first approach. © 2000 Elsevier Science B.V. All rights reserved.

JEL classification: G0; G1

Keywords: Value at Risk; Extreme value theory

1. Introduction

The portfolio choice of a safety-first investor is to maximize expected return subject to a downside risk constraint. Roy's (1952) and Arzac and Bawa's (1977) safety-first investor uses Value at Risk (VaR) as the measure for downside risk. Empirical applications of the safety-first principle have tended to use the Chebyshev bound. But this estimator can be rather imprecise as an estimate of the downside risk. The preceding paper by Gouriéroux et al. (2000) (henceforth, G–L–S) makes a number of contributions in addition to deriving analytical

* Corresponding author. Tel.: +1-979-845-7358; fax: +1-979-847-8757.

E-mail address: d-jansen@tamu.edu (D.W. Jansen).

expressions for the first and second derivatives of the VaR. In their empirical application, they introduce a kernel-based approach for estimating VaR, as well as a discussion of how to check the convexity of the estimated VaR-efficient portfolio set. In this paper we exploit the fact that asset returns are fat-tailed in order to estimate the downside risk and to calculate optimal safety-first portfolios.

We first show that the portfolio choice problem facing the G–L–S investor is the same as the portfolio choice problem faced by a safety-first investor as introduced by Roy (1952) and developed by Arzac and Bawa (1977). Then, we develop an alternative procedure for estimating VaR and constructing VaR-efficient portfolios using methods from the statistical theory of extremes. While G–L–S estimate their VaR-efficient portfolios using a kernel estimator to replace the unknown distribution of the portfolio by a smooth approximation, a nonparametric approach that attempts to estimate the entire distribution, we use a semiparametric approach that utilizes characteristics of the tails of the distribution. We are able to characterize, for an unknown distribution of returns, the behavior in the tails of the distribution, both in sample and for extreme values beyond those realized in our sample. Subsequently, we apply the method to several data sets, including the data used by G–L–S, and conclude by discussing the merits of the safety-first criterion.

We have several motives for this research program. First, we believe portfolio selection with limited downside risk to be a practical problem. Even if agents are endowed with standard concave utility functions such that to a first order approximation they would be mean-variance optimizers, practical circumstances often impose constraints that elicit asymmetric treatment of upside potential and downside risk. Regulatory concerns require commercial banks to report a single number, the so-called VaR, which gives the expected loss on their trading portfolio if the lowest 1% quantile return would materialize. Capital adequacy is judged on the basis of the size of this expected loss. Likewise, pension funds are often required by law to structure their investment portfolio such that the risk of underfunding is kept very low, e.g. equity investment may be capped.

Second, there is a wealth of experimental evidence for loss aversion (see, e.g. Kahneman et al., 1990). Other evidence is provided through consumption behavior. As Deaton (1991, 1992) shows, consumption responds asymmetrically to good and bad states. Similarly, within the mean-variance setup there is a range of returns such that consumption is too low for survival. Over this range, modeling risk by the expectation of squared returns may not be useful. Better may be to collapse all returns below the survival threshold as being equally risky.

Modeling limited downside risk has often been identified with the safety-first criterion developed by Roy (1952), although we have already commented that the G–L–S portfolio choice problem for VaR-efficient portfolios is essentially the

same problem as faced by Arzac and Bawa's safety-first investor.¹ In the past, safety-first was regarded as impractical for two reasons. The safety-first criterion was originally stated in terms of securing a minimal return level with a high probability. The criterion was made operational through the use of the Chebyshev bound on the probability of failure. But for minimal return levels below the riskless rate, the portfolio problem under the safety-first criterion is degenerate (see, e.g. Levy and Sarnat, 1972). Arzac and Bawa (1977) also noted that the original criterion fails to order risky assets, which are unambiguously ordered by the principle of absolute preference. By considering a lexicographic form of the safety-first principle, to the extent that the investor is only concerned about safety when the failure probability is above a critical level and otherwise maximizes expected return, and by allowing for borrowing and lending, Arzac and Bawa were able to resolve the theoretical shortcomings of the original criterion.

The other concern with the safety-first choice criterion has been the calculation of the failure probability. In practice it is not precisely known how returns are distributed. The literature, therefore, has proceeded by using the Chebyshev bound. The problem is that while this bound is robust, it is also highly inaccurate.

The question, then, is whether one can improve upon the estimate of the failure probability. If possible, this would also provide a genuine motive for using the safety-first principle. In particular, a method is called for which takes into account the fact that asset returns are heavy tailed, generating failure probabilities that are considerably different from normal probabilities. G–L–S provides one technique for doing so, a technique which involves nonparametric estimation of the entire distribution. We provide an alternative, semi-parametric method for modeling the tail of the distribution. For problems involving downside risks that are not in the tails of the distribution, the G–L–S technique should be superior, since our technique has nothing to say about the central part of the distribution. But for problems involving downside risks that are far in the tails of the distribution, our technique should be superior, since it is built upon knowledge of the tail behavior that must be followed by any unbounded fat-tailed distribution.²

The setup of this paper is as follows. In Section 2, we discuss the two inputs for portfolio selection with limited downside risk: safety-first theory as developed by Roy (1952) and Arzac and Bawa (1977), and statistical extreme value theory. In Section 3, we use the latter theory to estimate tail indices for stocks and bonds. In

¹ Roy developed the safety-first criterion at the same time as Markowitz introduced mean-variance analysis. See Bernstein (1992) for an interesting account of this coincidence. Jorion (1997) gives a good account of the development of VaR.

² There is one single feature of asset return data which is almost undisputed: Returns display high kurtosis. That is, in comparison to the bell-shaped normal density, there are more returns in the center and more returns in the tails of the distribution. This feature is often referred to as the heavy-tail property. In fact, the assumption of Frechet tails is commonplace for financial data (see Jansen and de Vries, 1991 or Longin, 1996).

Section 3, these estimates for tail fatness are used to construct portfolios with limited downside risk. In Section 4 we show how the analysis can be extended to longer horizons. Section 5 contains our conclusion.

2. Safety-first theory and extreme value theory

This presentation closely follows Arzac and Bawa (1977) but adopts the G–L–S notation for comparison to their paper. Investors’ preference ordering is represented as (π, μ) , where $\pi = 1$ if the probability (P) of a large negative return is less than a specified critical value δ , $P < \delta$, and otherwise $\pi = 1 - P$. The critical loss value δ will be called the risk probability (RP). For a given π , the investor then maximizes the expected return, μ .

We can state the portfolio problem as:

$$\max_{a, a_0}(\pi, \mu) \text{ subject to } \alpha' P_t + \alpha_0 = W_t \tag{1}$$

where

$$\begin{aligned} \pi &= 1 \text{ if } P = Pr(a' P_{t+1} + a_0 r \leq s) \leq \delta \\ \pi &= 1 - P \text{ otherwise} \end{aligned} \tag{2}$$

and

$$\mu = E a' P_{t+1} + a_0 r \tag{3}$$

where P_t denotes the vector of initial market values of assets at time t , P_{t+1} the vector of market value of assets at time $t + 1$, W_t the initial wealth of the investor, a_0 the amount of lending (if $a_0 > 0$) or borrowing (if $a_0 < 0$), r the risk-free gross rate of return, and a' the vector giving the amount of risky assets in the portfolio. Note that a safety-first investor must specify both the disaster level of wealth, s , and the maximal acceptable probability of this disaster, δ , which we refer to as the risk level.

The above problem can be restated in terms of gross returns R , $R = a' P_{t+1} / a' P_t$, by rewriting the safety-first condition as

$$\begin{aligned} Pr(a' P_{t+1} \leq s - a_0 r) &= Pr\left(\frac{a' P_{t+1}}{a' P_t} < \frac{s - a_0 r}{W_t - a_0}\right) \\ &= Pr\left(R \leq r + \frac{s - W_t r}{W_t - a_0}\right) \leq \delta. \end{aligned} \tag{4}$$

It is useful in this case to define a value $q_\delta(R)$ that such that there is a $\delta\%$ chance of returns less than or equal to this value.

$$Pr(R \leq q_\delta(R)) = \delta. \tag{5}$$

The negative of this quantile q_δ will be referred to as the VaR. The safety-first criterion is violated whenever

$$q_\delta(R) < r + \frac{s - W_t r}{W_t - a_0} \tag{6}$$

Note that a safety-first investor will exhibit risk aversion if the critical wealth level s is smaller than his secure final wealth Wr , an assumption that seems reasonable in application. For example, s could be stated as a fraction of wealth, say $0.8W$ or $0.9W$, in which case Wr most certainly exceeds s . In this case, a safety-first investor will decline a fair risk that violates the safety-first criterion in favor of pure lending at the risk-free rate r . But a risk averse safety-first investor will always buy some part of a divisible favorable risk, and the amount bought will be the maximum satisfying $(W_t - a_0)q_\delta(R) + a_0r = s$, or

$$W_t - a_0 = \frac{s - W_t r}{q_\delta(R) - r}. \quad (7)$$

Thus, as long as some favorable assets are available, the portfolio problem can be rewritten as the maximization of the excess returns:

$$\max_{a, a_0} \mu = W_t r + (W_t - a_0)(\bar{R} - r), \quad \bar{R} = E(R), \quad (8)$$

where the maximization is carried out among those portfolios that satisfy the safety-first criterion; that is, among those portfolios that have a probability of $1 - \delta$ or greater of maintaining a value in excess of s . By combining (7) and (8), this problem reduces to

$$\max_{a, a_0} \mu = Wr - (s - Wr)(\bar{R} - r)/(r - q_\delta(R)) \quad (9)$$

This implies that the risk averse safety-first investor can first maximize the ratio of the risk premium to the return opportunity loss that he is willing to incur with probability δ , i.e.

$$\max_a (\bar{R} - r)/(r - q_\delta(R)) \quad (10)$$

and then the investor can pick the scale of the risky part of his portfolio.

In the paper by G–L–S, they set up this same problem as maximizing an expected return subject to a VaR constraint. Thus, Arzac and Bawa's problem and G–L–S's problem are isomorphic. Other than some notational matters, the two approaches are the same. Thus, the portfolio selection problem in the G–L–S paper can be seen as an application of their results applying a nonparametric estimation procedure to the Arzac and Bawa portfolio selection problem. More generally, the similarity between the Arzac and Bawa problem and the G–L–S problem points out the similarity of VaR problems and Roy's (1952) safety-first criterion.

A natural question to ask is when and how the portfolio selection under safety-first will differ from such traditional methods as mean variance. In safety-first, the parameter δ and the actual portfolio distribution determines the VaR,

$q_\delta(R)$. This measure of risk may be preferred to the variance in some cases, because it is only based on large negative returns and it may apply in situations when the variance is not sufficient to measure the probability of a large negative shock.

There are some cases in which the safety-first problem can be restated as an equivalent mean-variance problem. Arzac and Bawa (1977) indicated that if $q_\delta(R)$ can be written as $\bar{R} - g(\delta)h(\gamma, \Theta)$, where Θ are the parameters of the distribution, then safety-first reverts to mean-variance. This happens, for example, when the distribution of returns is normal, Student's t , or stable with common characteristic exponent in the interval (1,2). It also happens when $q_\delta(R)$ is derived using Chebyshev's inequality.

Further, if the distribution of the returns is known, then the two-part optimization can be relatively straightforward (if the convolution that is needed for calculating $q_\delta(R)$ is not too complicated). But in most practical situations the return distribution is not known and $q_\delta(R)$ has to be estimated.

Roy himself initially proposed using the Chebyshev inequality, and the subsequent literature has largely followed up on this suggestion.³ But there are strong reasons to doubt that the Chebyshev inequality gives a sufficiently tight bound in practice (see de Haan et al., 1994). For one thing, it is well known that asset returns are not normal distributed but instead have fat tails. Below, we are able to exploit this fact to obtain a sharper bound than the Chebyshev bound.

How do we exploit the fact that asset distributions are heavy tailed? First, we note that in order to operationalize the safety-first measure of risk, we need to calculate tail probabilities. Hence, a semi-parametric approach that only models the tail of the distribution parametrically suffices. Second, such a semi-parametric approach may have several advantages. Vis-a-vis a fully parametric approach, tail estimates are less likely to be biased than parametric estimates which also need to reflect the center characteristics of the distribution.⁴ In comparison with the nonparametric approach, the semi-parametric approach can easily deal with out-of-sample VaR levels q . Third, the strength of the semi-parametric approach resides in the fact that all heavy-tailed distributions exhibit, to a first order approximation, identical behavior far out in the tails. This feature will be made precise in a moment. Fourth, the tail self-similarity permits a simple extension from the single period problem to the multiperiod portfolio investment problem. While the rule is different, it is as simple as the square root of time rule for the normal distribution (or for the Chebyshev, for that matter). That is discussed in Section 4.

³ A textbook discussion of safety-first and the use of the Chebyshev bound is provided in Elton and Gruber (1995).

⁴ See Mittnik et al. (1998) for further discussions of the choice between fully parametric and semi-parametric estimates for the tail index.

The unifying feature of heavy-tailed distributions is the regular variation property.⁵ A distribution $F(x)$ is said to be regularly varying at negative infinity if

$$\lim_{t \rightarrow \infty} \frac{F(-tx)}{F(-t)} = x^{-\alpha} \quad (11)$$

where $x > 0$ and $\alpha > 0$. The parameter α is called the tail index for obvious reasons. A similar condition applies to the positive tail. By l'Hopital's rule, it is immediate that regular variation implies that the tails of the density decline by a power. This is different from such distributions as the normal, which exhibit exponential decline, and exponential decline is eventually always faster than any power decline. From this one can understand why all the moments can be bounded for the normal distribution, while distributions which are regularly varying only have bounded moments less than α . Hence, the label 'heavy tailed.' It can be checked that, e.g. Student's t and the non-normal stable Paretian (or sum-stable) distributions satisfy (11), as well as the unconditional stationary distribution of certain ARCH processes with conditionally normal innovations.⁶

As we indicated, the regular variation property can be exploited to model (δ, q_δ) . Regular variation (11) implies that, asymptotically,

$$F(-x) = ax^{-\alpha} \quad (12)$$

as $x \rightarrow \infty$, for some $\alpha > 0$. The a and α parameters are specific to the distribution at hand, but can be estimated.

Given (12), we turn to estimation of the VaR quantile, q_δ . Let n be the sample size. Consider two failure probabilities p and t , such that one is in-sample and close to $1/n$, say $p > 1/n$, and one is out-of-sample, say $t < 1/n$. From expansion (12), we have

$$p = a(-q_p)^{-\alpha} \text{ and } t = a(-q_t)^{-\alpha} \quad (13)$$

for $q_t < q_p < 0$ and large n . Division of p by t and rearrangement gives

$$q_t \approx q_p (p/t)^{1/\alpha}. \quad (14)$$

Let $X_{(m)}$ be the m -th ascending order statistic from the sample $\{X_1, \dots, X_n\}$, i.e. $X_{(1)} \leq X_{(2)} \leq \dots \leq X_{(m)} \leq \dots \leq X_{(n)}$. Since the empirical distribution function

⁵ More detailed account of the regular variation property can be found in Resnick (1987) and Geluk and de Haan (1987).

⁶ For the Student's t distribution, α equals the degrees of freedom; for the stable distributions α equals the characteristic exponent.

is a mean squared error consistent estimator for $F(q)$, we will replace p by m/n and use $X_{(m)}$ for q_p . Conditional on having an estimate for α , our estimator for the unit VaR $-q_t$ is

$$\hat{q}_t = X_{(m)} \left(\frac{m}{nt} \right)^{1/\hat{\alpha}} \quad (15)$$

To estimate the tail index α , we use the well-known Hill (1975) moment estimator. This estimator can be interpreted as the maximum likelihood estimator of α in (12), conditioned on a low threshold value $X_{(m+1)}$ such that the RHS of (12) is a good approximation to $F(-x)$.⁷

The Hill estimator reads:

$$1/\hat{\alpha} = \frac{1}{m} \sum_{i=1}^m \left[\log(X_{(n+1-i)}/X_{(n-m)}) \right], \quad (16)$$

where the $X_{(i)}$ are again the lowest (negative) order statistics. Goldie and Smith (1987) provide a relatively simple proof of the fact that $(1/\alpha - 1/\hat{\alpha})m^{0.5}$ is asymptotically normal $N(0, 1/\alpha^2)$ if m increases suitably rapidly with n . de Haan et al. (1994) subsequently show that

$$\frac{\sqrt{m}}{X_{m+1} \log(m/np)} (\hat{q}_t - q_t) \sim N(0, 1/\alpha^2) \quad (17)$$

asymptotically.

3. Data analysis — tail estimates for stocks and bonds

Here, we consider the problem of choosing between investing in a mutual fund of bonds or a mutual fund of stocks. We use 67 years of monthly data on a US bond index and a US stock index, 1926.01–1992.12, which we obtained from the CRSP database. We also present, separately, an analysis of the two stocks considered by G–L–S, Thomson-CSF and L’Oreal. On these, we have 546 daily observations.

In Table 1, we present summary statistics. The top portion presents statistics for our monthly index of US stock and bond returns. We include both corporate bonds and government bonds, as well as stocks. Note that the returns on corporate and government bonds are highly positively correlated. Also, the mean return on stocks is about 0.008 per month, almost twice as high as the mean return on

⁷ See Berliant et al. (1996).

Table 1
Summary statistics

Returns on monthly US corporate bonds, government bonds, and stocks			
	Corporate bonds	Government bonds	Stocks
Mean	0.004445	0.003938	0.007943
Standard deviation	0.019782	0.021674	0.055702
Range	+0.133, -0.093	+0.142, -0.088	+0.320, -0.340
Skewness	0.746	0.746	-0.488
Kurtosis	10.027	8.496	9.888
J–B normality test	1728.6	1086.6	1621.4
No. of observations	804	804	804
Correlation of returns			
	Corporate bonds, government bonds: 0.838		
	Corporate bonds, stocks: 0.231		
	Government bonds, stocks: 0.183		
Returns on daily French stocks			
	Stock 1 (Thomson-CSF)	Stock 2 (L'Oreal)	
Mean	0.0000495	0.0005861	
Standard deviation	0.01261	0.01129	
Range	+0.0399, -0.0452	+0.0401, -0.0434	
Skewness	-0.239	0.061	
Kurtosis	4.114	4.311	
J–B normality test	33.46	39.44	
No. of observations	546	546	
Correlation of returns			
	Stock 1, stock 2: 0.385		

bonds. Corporate bonds have a mean return of 0.0044 per month, which exceeds the 0.0039 mean return on US government bonds and reflects the risk premium paid on corporate bonds. The standard deviation of stock returns is more than twice that on either of the bond series, and the range of sample values of stock returns is much wider than the range of values for the two bond returns. Table 1 also reports kurtosis calculations, which indicate that these three series have fat tails, and the Bera–Jarque test statistics, which decisively reject the normality of these returns. Our conclusion from the summary statistics is that there is little reason to think that either stock or bond returns are normally distributed. This is important to our application, since we will assume that the limiting shape of the tail of the return distribution is fat-tailed.

The bottom half of Table 1 presents summary statistics on the two French stocks. The mean returns are widely different, 0.00005 and 0.0006 per day, respectively, which translates into something like 1.25% and 16.2% annually, respectively. As we will see, these widely differing mean returns over the period of the G–L–S sample makes an application of our safety-first procedure somewhat problematic, and it would make an application to mean variance equally problem-

atic. The issue is that, with such a short sample and a resulting wide variation in the sample-specific mean returns (over the past 10-year period, the mean returns are much closer), the portfolio choice problem is swamped by the difference in means, especially for two stocks that show such similar risk characteristics whether measured by standard deviation or, as we shall see, by risk quantiles. Indeed, these two stocks show a great similarity in standard deviations, ranges, and excess kurtosis. For both stocks, the Bera–Jarque test statistic would lead us to reject normality.

In Table 2, the top portion of the table reports estimates of the tail index for both the upper and the lower tails of the distribution of our monthly stock and bond returns. For stocks, the point estimate for the lower tail index is 2.60, while the point estimate for the upper tail index is estimated to be 2.66. These results are comparable to those of Jansen and de Vries (1991), Loretan and Phillips (1994) and Longin (1996). We also report the two standard error confidence bands beneath the point estimates, and these can be quite wide. For the lower tail this confidence band stretches from about 1.2 to 4.0, while for the upper tail it stretches from about 2.1 to 3.3. Thus, for the lower tail the confidence band includes a values for the tail index for which the variance is not defined. Finally, in the column labeled m , we report the number of order statistics chosen by Hall's (1990) bootstrap method for estimating the tail index.⁸ Choosing the 'correct' value for m is important for the properties of the estimate of the tail index, yet choosing m appropriately has proven quite difficult and depends on unknown parameters of the underlying distribution. Hall's bootstrap is one attempt at dealing with this issue, but problems remain, and his bootstrap itself requires some assumption about unknown parameters of the underlying distribution. But we use his bootstrap procedure here, and for the lower tail of stock returns it indicates $m = 13$, while for the upper tail index it indicates $m = 79$.

If we look at our three return series over the common sample period 1926–1992, we see that while the point estimates for α in the upper and lower tails of the stock return series are very close, this is not as true for the bond series. For US government bonds, the point estimates of the tail indices appear to differ substantially, with the tail index in the lower tail estimated as 5.39 and the upper tail index estimated as 2.98. On the basis of a χ^2 test for equality, however, not even in this case could we reject equality of the upper and lower tail indices.⁹ For corporate bonds, the lower tail estimate is 2.93 and the upper tail value is 3.53. In these cases the two standard error confidence bands are quite wide, however, and include the point estimates from the other tail.

⁸ Hall (1990) describes a bootstrap procedure for calculating the number of order statistics to use in getting an estimate for the tail index. His procedure requires an initial starting estimate for m , which we take as \sqrt{n} . Unfortunately, Hall's procedure is not a panacea.

⁹ This test is described and used in Jansen and de Vries (1991).

Table 2
Estimates of tail indices and exceedence values

	m	X_{n-m}	$\alpha (\pm 2 \text{ S.E.})$	X_p : Exceedences		
				$p = 1/n$	$p = 1/(1.5n)$	$p = 1/(2n)$
<i>Monthly stock and bond returns, 1926.01–1992.12 (n = 804)</i>						
Stocks 1926–1992 lower tail	13	-0.13150	2.601 (± 1.414)	-0.352 (-0.230, -0.760)	-0.412 (-0.254, -1.086)	-0.460 (-0.274, -1.442)
Stocks 1926–1992 upper tail	79	0.06009	2.660 (± 0.587)	0.311 (0.228, 0.487)	0.362 (0.259, 0.599)	0.403 (0.284, 0.694)
US bonds 1926–1992 lower tail	9	-0.05055	5.390 (± 3.521)	-0.076 (-0.060, -0.104)	-0.082 (-0.062, -0.120)	-0.086 (-0.064, -0.133)
US bonds 1926–1992 upper tail	17	0.06042	2.983 (± 1.763)	0.135 (0.092, 0.257)	0.155 (0.099, 0.348)	0.170 (0.106, 0.439)
Corporate bonds 1926–1992 lower tail	16	-0.03843	2.932 (± 1.437)	-0.099 (-0.068, -0.184)	-0.114 (-0.074, -0.241)	-0.125 (-0.079, -0.296)
Corporate bonds 1926–1992 upper tail	16	0.05591	3.537 (± 1.733)	0.122 (0.088, 0.199)	0.137 (0.095, 0.245)	0.149 (0.106, 0.287)
<i>Daily data on French stocks (n = 546)</i>						
Stock 1	21	0.0275	4.37 (± 1.941)	-0.054 (-0.042, -0.076)	-0.058 (-0.044, -0.084)	-0.063 (-0.046, -0.097)
Stock 2	13	0.0285	4.829 (± 2.625)	-0.049 (-0.038, -0.069)	-0.052 (-0.039, -0.076)	-0.056 (-0.041, -0.089)

We also tested whether the bounded variance assumption, i.e. $\alpha > 2$, was supported by the data. Given the failure to reject similarity of the point estimate for the tail index using the upper and lower tail, we combined these tails to estimate the tail index and test for a bounded variance. In the case of stock returns, we could reject the null of $\alpha < 2$ at a significance level of 8%, while for bonds the marginal significance level for this test was below 1%. Note that for the lower tail of stock returns, this null hypothesis was not rejected. This seems to be largely due to the low number of order statistics selected by the Hall bootstrap procedure for the lower tail, and the effect of this low value for m in the formula for the variance of the tail index estimate.¹⁰

In the rightmost three columns of Table 2, we calculate the unit-investment VaR values corresponding to various probabilities. The first of these columns is headed $p = 1/n$, to indicate that this calculation applies when the probability is such that we would expect one occurrence in our sample. Looking at stock returns, we see that the estimated exceedence value in the lower tail is -0.35 or a -35% return. Thus, we estimate that there is a 1 in 804 chance that we will lose US\$0.35 or more on a US\$1 investment. In the upper tail, we estimate that there is a 1 in 804 chance of a monthly return of $+31\%$ or greater. Notice that we estimate much lower VaR values for the two bond series. For corporate bonds, there is a 1 in 804 chance that we will experience a monthly return of -9.9% or less, and a monthly return of $+12.2\%$ or higher. We also notice that even though the tail indices for the lower tails of the two bond series differ considerably the VaR estimates are quite similar. This can be explained by the values of $X_{(n-m)}$ in the middle column, which is a measure of the scale, see formula (15). Thus, even though the government bond series is thinner tailed it has a large scale. Over a large range, these two effects more or less cancel each other.

The most interesting calculations are those for probabilities much lower than $1/n$, since we could use the empirical distribution for probabilities of $1/n$ or larger. Based on our tail-behavior assumptions, we can estimate the VaR values of events that occur with much less frequency than once in our sample. Hence, we report a column headed $p = 1/(2n)$, a probability of 1 in 1608. For stock returns, we estimate that there is a 1 in 1608 chance of a -46% or smaller return, and the

¹⁰ The estimation of the tail index parameter, the properties of the estimators, and the importance of the choice of m , are all well-known issues in the literature on extremes. Our solution to this issue — the Hill estimator with the Hall bootstrap — is but one of large number of possibilities. For instance, Danielsson et al. (1996) proposed an estimator that was the ratio of second to first moments, showing that it had lower bias than the Hill estimator in certain cases. Huisman et al. (2000) developed a small-sample estimator for the tail index. If we had used their estimator, we would have estimated the tail index parameters for the two French stocks as 7.438 and 4.475 instead of 4.37 and 4.829 as indicated in Table 2. Using the Huisman et al. estimator would have changed little in our stock selection exercise, as we use the lower estimate for the tail index for all portfolios involving both stocks, so that for all but one portfolio the tail index value we used, 4.37, differed little from that we would have used following Huisman et al., 4.475.

same chance of a +40% or more gain. In contrast, for corporate bonds there is the same 1 in 1608 chance of a –12.5% or lower return on or of a +14.9% or greater gain. These estimates accord well with a priori notions that bonds are less subject to such large negative or positive returns as are stocks.

The bottom portion of Table 2 reports similar statistics for our sample of daily returns on our two French stocks. We only report the statistics that are calculated on the combined data from the upper and lower tails. Note that the estimated tail indices are very similar, either 4.5 or 4.8. The exceedences are calculated for three probabilities, either 1 in 546 or 0.0018, or 0.75 in 546 or 0.00135, or 0.5 in 546 or 0.0009. In all cases, the exceedence quantiles are similar for our two stocks. In fact, for the probability 0.0009, an event that only occurs on average once in a period twice the size of our sample, we estimate the daily exceedence quantile to be either 0.055 or 0.064 for our two stocks, less than 10% in 1 day.

We next turn to illustrate portfolio selection with limited downside risk. We begin with our monthly US stock and bond data, and proceed to construct hypothetical portfolios consisting of linear combinations of our stock index and our corporate bond index. (Government bond returns are highly correlated with corporate bonds and hence portfolios consisting of government bonds and stocks are similar to the portfolios we investigate here.) We vary the fraction of stocks (and bonds) from 0% to 100% by steps of size 10%. This gives us 11 portfolios. To construct the extreme quantiles and corresponding probabilities from Eq. (18), we need the scale measure $X_{(m)}$ and the tail index α . The $X_{(m)}$ are readily found for each portfolio by ordering the returns from high to low. Note that due to the mixing of stock and bond returns, the ordering may change as the proportions change. The α values, however, are unaffected by the mixing as long as there is some positive fraction of stocks in the portfolio. This follows from proposition 1.7 of Geluk and de Haan (1987), which states that a convolution of two regularly varying variables produces a random variable which has the same tail properties as the fattest tail of the two convoluting variables, i.e. the fattest tail dominates. Thus, for all portfolios except the one which involves 100% bonds, we take $\alpha = 2.601$, cf. Table 2, and for the portfolio with 100% bonds, we take $\alpha = 2.932$.

We first calculate for each hypothetical portfolio the VaR quantiles that go with various probabilities. These are reported in Table 3 without further comment. Recall that the safety-first investor must specify the critical level of wealth, s , and the probability δ that wealth will not decline to the critical level or lower. Of course, the critical wealth level, s , can be mapped into a critical return level, as is implicit in Eq. (7) above. As an example, assume that $\delta = 0.000625$, and further assume that our investor states that the critical wealth level s is a fraction of wealth, $0.70W$. This maps into a gross return of 0.70 (a net return of –30%), and if there were no risk-free asset available then our safety-first investor would choose a portfolio of 70% stocks, where $q_\delta = 1 - 0.2889 = 0.7111$. However, our risk averse safety-first investor can also borrow or lend at the risk-free rate r , and in the first stage of his portfolio choice problem maximizes $(\bar{R} - r)/(r - q_\delta(R))$.

Table 3
Estimated VaR levels corresponding to the stated probabilities

Portfolio of monthly US stocks and corporate bonds	Probabilities (expected number of occurrences in parentheses)			
	0.0025 (2 in 804)	0.00125 (1 in 804)	0.000625 (0.5 in 804)	0.0001 (0.08 in 804)
100% Stock	-0.2696	-0.3519	-0.4594	-0.9293
90% Stock	-0.2449	-0.3197	-0.4173	-0.8441
80% Stock	-0.2158	-0.2816	-0.3677	-0.7437
70% Stock	-0.1877	-0.2450	-0.3198	-0.6469
60% Stock	-0.1695	-0.2213	-0.2889	-0.5844
50% Stock	-0.1427	-0.1863	-0.2431	-0.4918
40% Stock	-0.1125	-0.1468	-0.1917	-0.3877
30% Stock	-0.0986	-0.1287	-0.1680	-0.3398
20% Stock	-0.0953	-0.1244	-0.1624	-0.3286
10% Stock	-0.0793	-0.1035	-0.1351	-0.2734
0% Stock	-0.0780	-0.0988	-0.1251	-0.2337
<hr/>				
Portfolio of daily French stocks	0.0018 (1 in 546)	0.00135 (0.75 in 804)	0.0009 (0.5 in 546)	
100% Stock 2	-0.0486	-0.0515	-0.0561	
90% Stock 2	-0.0472	-0.0503	-0.0550	
80% Stock 2	-0.0438	-0.0467	-0.0511	
70% Stock 2	-0.0422	-0.0449	-0.0491	
60% Stock 2	-0.0420	-0.0447	-0.0489	
50% Stock 2	-0.0421	-0.0449	-0.0491	
40% Stock 2	-0.0424	-0.0452	-0.0494	
30% Stock 2	-0.0439	-0.0468	-0.0512	
20% Stock 2	-0.0473	-0.0504	-0.0551	
10% Stock 2	-0.0499	-0.0531	-0.0581	
0% Stock 2	-0.0539	-0.0575	-0.0628	

We calculate this quantity for two values of r in Table 4, and for two choices of δ , $\delta = 0.0025$ and $\delta = 0.00625$. Note that we calculate values for q_δ using Table 3.

For $\delta = 0.0025$, we consider the case $r = 1$, so that the risk-free asset pays a zero net return, and $r = 1.00303$ (which translates into an annual interest rate of 3.7%), the average return on US Treasury Bills over 1926–1992. Optimal portfolios in Table 4 are marked with double asterisks (**). We see that for $r = 1$, the optimal portfolio is one that contains 10% stock and 90% bonds, while for $r = 1.00303$ the optimal portfolio is the one containing 40% stocks and 60% bonds.

How much will our investor place in the risk-free asset? This depends on s , the safety-first wealth level. For $s = 0.70W$, we can calculate the amount borrowed or lent at the risk-free rate from $(W + b)q_\delta(R) - br = s$. For $r = 1$ and the optimal portfolio of 10% stocks, we have that $b = (0.70 - q_d(R))W / (q_\delta(R) - r) =$

Table 4
Portfolio selection for monthly US stocks and bonds

Portfolio	$q_\delta(R)$	$(R-r)/(r-q_\delta),$ $r=1$	$(R-r)/(r-q_\delta),$ $r=1.00303$
<i>Portfolio selection with $\delta = 0.0025$</i>			
100% Stock	1 - 0.2696	0.02946	0.01802
90% Stock	1 - 0.2449	0.03101	0.01841
80% Stock	1 - 0.2158	0.03357	0.01926
70% Stock	1 - 0.1877	0.03674	0.02026
60% Stock	1 - 0.1695	0.03860	0.02036
50% Stock	1 - 0.1427	0.04341	0.02172
40% Stock	1 - 0.1125	0.05196	0.02437 ^a
30% Stock	1 - 0.0986	0.05574	0.02426
20% Stock	1 - 0.0953	0.05397	0.02150
10% Stock	1 - 0.0793	0.06046 ^a	0.02143
0% Stock	1 - 0.0780	0.05701	0.01747
<i>Portfolio selection with $\delta = 0.000625$</i>			
100% Stock	1 - 0.4594	0.01729	0.01063
90% Stock	1 - 0.4173	0.01820	0.01086
80% Stock	1 - 0.3677	0.01970	0.01137
70% Stock	1 - 0.3198	0.02156	0.01197
60% Stock	1 - 0.2889	0.02265	0.01204
50% Stock	1 - 0.2431	0.02548	0.01285
40% Stock	1 - 0.1917	0.03049	0.01446 ^a
30% Stock	1 - 0.1680	0.03271	0.01441
20% Stock	1 - 0.1624	0.03167	0.01278
10% Stock	1 - 0.1351	0.03548	0.01277
0% Stock	1 - 0.1251	0.03553 ^a	0.01104

^a Indicates optimal portfolio among available choices.

$-0.2207W/(0.9207 - 1) = 2.7831W$. For an investor with wealth normalized to $W = 1$, we have borrowing of an additional US\$2.7831 to invest in a portfolio consisting of 10% stocks and 90% bonds. The mean return on this portfolio is 1.00479, so this investor earns a mean return of $(3.7831 \times 1.00479) - 2.7831 = 1.0181$. If disaster strikes, his wealth ends up as $(3.7831 \times 0.9207) - 2.7831$, or 0.70, just as the safety-first criterion demands.

What is wrong with investing 100% in stocks? Clearly the mean return is higher, 1.00794. But the safety-first criterion severely limits the amount of wealth our investor can place in such a risky investment. With $s = 0.7W$ (and $r = 1$), we calculate that our investor will borrow $b = (0.70 - q_d(R))W/(q_\delta(R) - r) = (0.70 - 0.7304)W/(0.7304 - 1) = 0.1128W$, i.e. he will borrow 0.1128 for every dollar of wealth. The expected return on this portfolio is $(1.11280 \times 1.00794) - 0.1128 = 1.00884$, less than the return on the leveraged portfolio containing only 10% stocks. If disaster strikes, the 100% stock portfolio will decline in value to $(1.1128 \times 0.7304) - 0.1128 = 0.70$, just as required by the safety-first constraint.




Fig. 1. US stock and bond index.

The conclusion, then, is that a leveraged portfolio containing mostly bonds generates a superior rate of return given the constraint of holding the risk of a gross return of 0.70 to a 0.0025 probability. Thus, a safety-first investor will not want to go too far into stocks. Of course, this conclusion depends strongly on the quantile estimates, which in turn depend on the choice of δ , the probability that the portfolio selection criterion will be violated. Also playing a role is the choice of s , the safety-first portfolio value.

(Fig. 1) illustrates the portfolio choice problem, plotting the mean return versus VaR for a portfolio of stock and bond indices when $r = 1.003$. If we turn to the daily data on the two French stocks analyzed by G–L–S, we see in Table 3, bottom half, the calculated VaR levels corresponding to the stated probabilities

Table 5
Portfolio selection for daily French stocks

Portfolio	$q_\delta(R)$	$(R - r)/(r - q_\delta)$, $r = 1$
100% Stock 2	1 - 0.0486	0.0121 ^a
90% Stock 2	1 - 0.0472	0.0113
80% Stock 2	1 - 0.0438	0.0109
70% Stock 2	1 - 0.0422	0.0101
60% Stock 2	1 - 0.0420	0.0089
50% Stock 2	1 - 0.0421	0.0075
40% Stock 2	1 - 0.0424	0.0062
30% Stock 2	1 - 0.0439	0.0048
20% Stock 2	1 - 0.0473	0.0033
10% Stock 2	1 - 0.0499	0.0021
0% Stock 2	1 - 0.0539	0.0009

Portfolio selection with $\delta = 0.0018$.

^aOptimal portfolio.

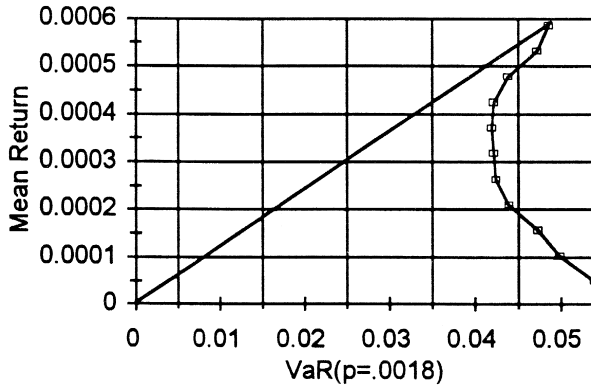


Fig. 2. G–L–S stocks.

(0.0018, 0.00135, 0.0009) for the various portfolio proportions. These are much more similar across the various portfolio combinations of the two French stocks than they were for the various combinations of the US stock and bond indices in the top half of Table 3.

Using these calculated VaR levels and the historical return data as an estimate of the future mean return, we report in Table 5 the values of the VaR and the portfolio selection criterion for the various portfolios of the two stocks. Here, we can see quite clearly the effect of the large difference in mean returns over the sample period. The large mean return on the second stock swamps any changes in the VaR estimates, which are quite similar across portfolios, and leads to a corner solution where only one stock is held in the portfolio.

(Fig. 2) illustrates this feature. The limited downside risk portfolio selection criterion chooses a corner solution where the entire portfolio is invested in one stock. This is not a unique feature of our portfolio selection criterion. Fig. 3 illustrates mean-variance portfolio selection for these two stocks, using the in-sam-




Fig. 3. G–L–S stocks, mean variance.

ple variance and covariance values to calculate the portfolio standard deviation. In a mean-variance framework the sample period data for the two French stocks would also lead to a corner solution.

This application shows that the safety-first criterion is not without its problems. If the list of assets is unrestricted, safety-first portfolios can be very unbalanced. As Dert and Oldenkamp (1997) and Vorst (1999) noted, the safety-first constraint may be met by investing just enough into the riskless bond, while the rest of the wealth is allocated to buying for out of the money call options (since expected returns are increasing in the exercise price).

We leave for future research the amended safety-first criterion whereby investors maximize expected utility of the traditional form subject to their wealth constraint, but also subject to the side constraint that the VaR level must be maintained. Whatever the choice theoretic aspects of this reformulated safety-first criterion, it is the practical problem faced by the financial industry that operates under such an additional constraint under the Basle rules.

4. Extension to longer investment horizons

Suppose that the investor has a longer horizon than the monthly implicit in our use of monthly returns. There is nothing in the safety-first theory that requires one being wedded to a particular horizon. Extreme value theory allows one to easily change the time horizon of the analysis, without changing the frequency of the observations. This is analogous to the square root of time rule if normality prevails; but the multiplication factor is different if the distributions are heavy tailed. To show this, consider the expressions for the cdf's of two return variables, X and Y , where X is a one-period return, and Y is a k -period return or k -convolution defined as:

$$Y = \sum_{i=1}^k X_i.$$

It is assumed that the X_i are i.i.d.¹¹ In that case, from Feller's theorem (Feller, 1972, VIII.8), if

$$F(-x) \approx ax^{-\alpha}, \quad (18)$$

as $x \rightarrow \infty$, then

$$P\{Y \leq -x\} \approx kax^{-\alpha}. \quad (19)$$

That is, the cdf for the extremes of the k -period return Y and the cdf for the extremes of the single period return X are closely related, differing only by the

¹¹ Since GARCH and other models of conditional heteroskedasticity are frequently applied to financial data, it bears emphasis that this i.i.d. assumption is perhaps a strong one, one which may be violated and whose violation may have important consequences for certain results in this section.

factor k . Note that for the Chebyshev-based approach, the same multiplicative factor k enters the RHS if one considers a k -period return. Hence, the asymptotic difference between the two approaches remains equally important.

From Eq. (19), it is straightforward to obtain the asymptotic change in the quantile that is needed to keep the probability level (risk level) invariant under the k -period convolution. For x , sufficiently large $P\{Y \leq -k^{1/\alpha}x\} \approx p\{X \leq -x\}$. Hence, there is a simple method for calculating the k -period unit-VaR from the estimated single period VaR measure \hat{q}_t that was given in (15). The k -period unit-VaR is

$$k^{1/\alpha} \hat{q}_t. \quad (20)$$

Recall that under normality, or if the Chebyshev approach is used, then the rule is to multiply the one-period quantile by the square root of time. If the distribution is fat-tailed, the rule for the extreme quantiles is to apply an α -root of time multiplication factor. This is further discussed in Dacarogna et al. (1995).

Thus, we can easily extend the lexicographic safety-first criterion by incorporating choices over different time horizons, in the sense that the safety requirement which is most binding has to be met before return maximization becomes an issue. No new parameter estimates are necessary for the multiple horizon case, since (20) states that the k -horizon problem is just a scaled up version of the single period horizon problem. Also, as is shown by Dacarogna et al. (1995), one can benefit from the highest frequency at which the data are available, since no reestimation of the tail index $\hat{\alpha}$ is necessary for the low frequency data. Extremes from heavy-tailed distributions have a fractal nature. Thus, efficiency of $\hat{\alpha}$ is generally highest on the highest available frequency of the data. Hence, there are good reasons for using an estimator of the tail index over high-frequency data even when the estimated tail index is to be used to calculate exceedences over lower frequency or time-aggregated data.

We conducted a small simulation to illustrate the above arguments. We simulated 10,000 observations from a Student's t with 3 df , and calculated loss levels for three different risk levels for both the single-period and two multiple-period returns. These are reported in Table 6. The true loss-level values are reported in the square brackets. The q estimates are reasonably close to the true values, but the results deteriorate if one moves to the p value, which is below the inverse of the sample size (in the last row of the table). The precision and bias of course also increase if one moves to the 2-period and 4-period horizons. The combination of extremely low risk levels and multiple horizons in the lower right-hand corner of the table yields a 95% confidence band as wide as (45–144). One might ask whether reestimation instead of calculating the multiple horizon results from the single period results would not improve the situation. In fact, if one used the reestimation procedure the confidence band becomes even wider, i.e. it ranges from -57 to $+135$, which is even worse. The reason is that the reduction in sample size which one has to bear with if one reestimates the

Table 6
Simulated Student's t loss levels and time aggregation

Risk level	Single period	Multiple period	
	x_p	$2^{1/\alpha} x_p$	$4^{1/\alpha} x_p$
0.0005	13.60 (11.07–16.14) [12.92]	17.49 (13.52–24.47) [16.64]	22.51 (16.44–28.59) [21.40]
0.0001	24.42 (17.49–31.35) [22.20]	31.45 (21.15–41.75) [16.64]	40.53 (25.46–55.59) [35.79]
0.00001	56.71 (32.28–81.13) [47.93]	73.16 (38.35–107.96) [60.50]	94.43 (45.21–143.64) [76.35]

The numbers presented are averages over 250 replications. The 95% bootstrap confidence intervals are presented in parentheses. In each replication, a sample of 10,000 Student's t (3 df) distributed observations is drawn. Within square brackets, the theoretical quantile values are reported.

parameters on the aggregated return data outweighs the potential gain from using the aggregated data. (It can be shown that direct estimation for the single period data should give upwards biased quantiles, as is the case in Table 6, while for the multiperiod convoluted Student's t data the bias should go in the other direction. This latter fact is not present in the Table 8, since we used the single period estimates to calculate the multiperiod quantiles. But the bias reduction that direct estimation would yield, does not outweigh the loss in efficiency.) Tables 7 and 8 contain the estimated VAR levels corresponding to the stated probabilities for various portfolios and the optimal portfolios for the annual horizon. With a net risk-free rate of zero, the optimal portfolio contains 10% stocks, for a net risk-free rate of 0.003, the optimal portfolio is one that contains 40% stocks. Note that these results are obtained for a risk level equal to the inverse of the sample size. The

Table 7
Estimated VaR levels corresponding to stated probabilities

Portfolios	Probabilities (expected number of occurrences in parentheses)	
	0.015 (1 in 67)	0.01 (1 in 100)
100% Stock	–0.3518	–0.4112
90% Stock	–0.3197	–0.3736
80% Stock	–0.2817	–0.3292
70% Stock	–0.2526	–0.2952
60% Stock	–0.2214	–0.2587
50% Stock	–0.1863	–0.2178
40% Stock	–0.1469	–0.1716
30% Stock	–0.1287	–0.1504
20% Stock	–0.1243	–0.1453
10% Stock	–0.1035	–0.1209
0% Stock	–0.0987	–0.1133

US stock and bond portfolios, annual (67 years of annual data).

Table 8
Portfolio selection with $\delta = 0.01$

Portfolio	$q_\delta(R)$	$(R - r)/(r - q_\delta),$ $r = 1$	$(R - r)/(r - q_\delta),$ $r = 1.0370$
100% Stock	1 - 0.4112	0.2422	0.1397
90% Stock	1 - 0.3736	0.2529	0.1400
80% Stock	1 - 0.3292	0.2714	0.1429
70% Stock	1 - 0.2952	0.2853	0.1421
60% Stock	1 - 0.2587	0.3057	0.1423
50% Stock	1 - 0.2178	0.3395	0.1450
40% Stock	1 - 0.1716	0.4010	0.1525 ^a
30% Stock	1 - 0.1504	0.4235	0.1424
20% Stock	1 - 0.1453	0.4030	0.1183
10% Stock	1 - 0.1209	0.4419 ^a	0.1041
0% Stock	1 - 0.1133	0.4263	0.0752

^a Indicates optimal portfolio among choices listed.

resulting portfolio allocation is very similar to the one under the monthly investment horizon. For the $r = 1$ case and $s = 0.7W$ we find that the average yearly net return on the optimal portfolio is 13%; while if $r = 1.037$ and $s = 0.8W$, the average return on the optimal portfolio is 7%.

5. Conclusion

In this paper, we deal with the trade-off between expected return and risk when investors desire limited downside risk. We show that portfolio selection with limited downside risk includes both the safety-first investor of Roy and, more specifically, Arzac and Bawa (1977), as well as the VaR-constrained investor of G–L–S. We calculate the optimal VaR-efficient portfolios for both a portfolio of US stocks and bonds indices observed monthly over the period 1926–1992 and for a portfolio of two French stocks observed daily for a period in the 1990s. These applications show how extreme value theory proves to be a useful procedure for estimating VaR-efficient portfolios and for describing portfolio risk for events far out in the tails of the distribution. But we also discussed the limitations of the current safety-first criterion, and argued it may be worthwhile to investigate its merits as an additional side constraint to regular utility maximization.

Acknowledgements

The authors thank Hans Dewachter, Andre Lucas, Ronald Mahieu, Francois Nissen, Adrian Pagan, and Peter Schotman, two anonymous referees and Franz

Palm (the editor) for helpful comments, and Laurens de Haan for inspiring and encouraging these applications. They also thank seminar participants at Maastricht University, the Federal Reserve Bank of Dallas, the University of Notre Dame, and Texas A&M University. Jansen acknowledges research assistance from the Maastricht University, the Private Enterprise Research Center, and the Office of International Coordination at Texas A&M University.

References

- Arzac, E.R., Bawa, V.S., 1977. Portfolio choice and equilibrium in capital markets with safety-first investors. *Journal of Financial Economics* 4, 277–288.
- Berliant, J., Teugels, J.L., Vynckier, P., 1996. *Practical Analysis of Extreme Values*. Leuven Univ. Press, Belgium.
- Bernstein, P., 1992. *Capital Ideas*. Free Press, New York.
- Dacarogna, M.M., Muller, U.A., Pictet, O.V., de Vries, C.G., 1995. The distribution of extremal foreign exchange returns in extremely large data sets. Tinbergen Institute Working Paper TI 95-70.
- Danielsson, J., Jansen, D.W., de Vries, C.G., 1996. The method of moments ratio estimator for the tail shape parameter. *Communications in Statistics — Theory and Methods* 25 (4), 711–720.
- Deaton, A., 1991. Saving and liquidity constraints. *Econometrica* 59, 1221–1248.
- Deaton, A., 1992. *Understanding Consumption*. Oxford Univ. Press, Oxford.
- de Haan, L., Jansen, D.W., Koedijk, K., deVries, C.G., 1994. Safety first portfolio selection, extreme value theory and long run asset risks. In: Galombos (Ed.), *Proceedings from a Conference on Extreme Value Theory and Applications*. Kluwer Academic Publishing, Dordrecht, pp. 471–487.
- Dert, C., Oldenkamp, B., 1997. Optimal guaranteed return portfolios and the casino effect, Technical report 9704, Erasmus Center for Financial Research.
- Elton, E., Gruber, M., 1995. *Modern Portfolio Theory and Investment Analysis*. Wiley, New York.
- Feller, W., 1972. 2nd edn. *An Introduction to Probability Theory and its Applications*, vol. 2, Wiley, New York.
- Geluk, J.L., de Haan, L., 1987. Regular variation, extension, and Tauberian theorems. *CWI Tract*, vol. 40, CWI, Amsterdam.
- Goldie, C.M., Smith, R.L., 1987. Slow variation with remainder: theory and applications. *Quarterly Journal of Mathematics* 38, 45–71.
- Gourieroux, C., Laurent, J.P., Scaillet, O., 2000. Sensitivity analysis of values at risk, manuscript.
- Hall, P., 1990. Using the bootstrap to estimate mean squared error and select smoothing parameter in nonparametric problems. *Journal of Multivariate Analysis* 32, 177–203.
- Hill, B.M., 1975. A simple general approach to inference about the tail of a distribution. *The Annals of Statistics* 3, 1163–1174.
- Huisman, R., Koedijk, K., Kool, C., Palm, F., 2000. Tail-index estimates in small samples. *Journal of Business and Economic Statistics*, forthcoming.
- Jansen, D.W., de Vries, C.G., 1991. On the frequency of large stock returns: putting booms and busts into perspective. *The Review of Economics and Statistics* 73, 18–24.
- Jorion, P., 1997. *Value at Risk*. McGraw-Hill, New York.
- Kahneman, D., Knetsch, J.L., Thaler, R.H., 1990. Experimental tests of the endowment effect and the Coase theorem. *Journal of Political Economy* 99, 1325–1350.
- Levy, H., Sarnat, M., 1972. Safety first — an expected utility principle. *Journal of Financial and Quantitative Analysis* 7, 1829–1834.
- Longin, F., 1996. The asymptotic distribution of extreme stock market returns. *Journal of Business* 69, 383–408.

- Loretan, M., Phillips, P.C.B., 1994. Testing the covariance stationarity of heavy-tailed time series: an overview of the theory with applications to several financial datasets. *Journal of Empirical Finance* 1, 211–248.
- Mittnik, S., Paolella, M., Rachev, S., 1998. A tail estimator for the index of the stable Paretian distribution. *Communication in Statistics — Theory and Methods* 27, 1239–1262.
- Resnick, S., 1987. *Extreme Values, Regular Variation and Point Processes*. Springer-Verlag, New York.
- Roy, A.D., 1952. Safety first and the holding of assets. *Econometrica* 20, 431–449.
- Vorst, T., 1999. Portfolio optimization subject to VaR, mimeo, Erasmus University Rotterdam.