Appearing in the *Journal of Risk Finance*, vol. 6, #5, pp.388-403.

Value at Risk with Info-gap Uncertainty

Yakov Ben-Haim
Yitzhak Moda'i Chair in Technology and Economics
Faculty of Mechanical Engineering
Technion — Israel Institute of Technology
Haifa 32000 Israel
yakov@technion.ac.il
http://www.technion.ac.il/yakov
Tel:+972-4-829-3262, Fax: +972-4-829-5711

Contents

Keywords: Value at Risk, Knightian uncertainty, information-gaps, robustness, stress testing.

Abstract

The value at risk of a financial investment is assessed as the quantile of an estimated probability distribution of the returns. Estimating a VaR from historical data entails two distinct sorts of uncertainty: probabilistic uncertainty in the estimation of a pdf from historical data, and non-probabilistic Knightian info-gaps in the future size and shape of the lower tail of the pdf. A pdf is estimated from historical data, while a VaR is used to predict future risk. Knightian uncertainty arises from the structural changes, surprises, etc., which occur in the future and therefore are not manifested in historical data. In this paper we concentrate entirely on Knightian uncertainty and we do not consider the statistical problem of estimating a pdf. We use info-gap decision theory to study the robustness of a VaR to Knightian uncertainty in the distribution. We show that VaRs, based on estimated pdfs, have no robustness to Knightian errors in the pdf. We derive an info-gap safety factor which multiplies the estimated VaR in order to obtain a revised VaR with specified robustness to Knightian error in the pdf. We define a robustness premium as a supplement to the incremental VaR for comparing portfolios.

1 Introduction

The physicist Nils Bohr was fond of the Danish aphorism that "Prediction is always difficult, especially of the future" (Moore, 1989). Bohr lived in tumultuous times and he knew what he was talking about. The aphorism applies to financial risk whether or not Bohr had that in mind. For instance, Hendricks (1996) finds extensive empirical evidence for

> two well-known characteristics of daily financial market data. First, extreme outcomes occur more often and are larger than predicted by the normal distribution (fat tails). Second, the size of market movements is not constant over time (conditional volatility).

Just as Bohr said.

The dispute about predictability of market outcomes has raged for a long time and is here to stay. Fama (1965, p.34) argued that "[T]he series of price changes has no memory, that is, the past cannot be used to predict the future in any meaningful way." (Bohr again). While more recently (Fama and French, 1988, p.247) we read about "the mounting evidence that stock returns are predictable." (Bohr, applied to econometrics).

Of course, 'predictable' does not mean 'infinitely reliable':

> Virtually all of the approaches [to assessing value at risk] produce accurate 95th percentile risk measures. The 99th percentile risk measures, however, are somewhat less reliable and generally cover only between 98.2 percent and 98.5 percent of the outcomes (Hendricks, 1996).

VaRs are evaluated from estimated probability density functions (pdfs); estimates are based on history. This limits the accuracy of estimated VaRs as predictions of *future* risk for reasons which we will divide into two categories. The first category is estimation uncertainty, while the second category we will refer to broadly as Knightian uncertainty.

The evaluation and management of estimation uncertainty is in the province of statistical analysis. Numerous powerful methods are available for estimating quantiles and for evaluating the error of those estimates. Standard errors of quantile estimates can be evaluated (Kendall, Stuart and Ord, 1987, §10.9). Confidence intervals can be constructed for quantile values (Degroot, 1986, p.563). Sign tests use order statistics for testing hypotheses on the value of an estimated quantile (Kendall, Stuart and Ord, 1979, §32.2–10). Kernel smoothing methods provide a rich array of methods for estimating a pdf (Härdle, 1990). The starting point of the current paper is the assumption that a pdf of returns has been competently estimated, and the reliability of this estimate, vis á vis the historical data, has been established using appropriate statistical tools. We are not concerned with the statistical task of estimating a pdf.

This paper is concerned with the second class of factors which limit the accuracy of VaRs: Knightian uncertainty about the future. The systematic errors in VaRs mentioned earlier suggest that something more than random estimation error is involved. The most troublesome source of error in predicting *future* VaRs from *historical* data is that things can change. The true pdf in the future can differ substantially from the true realization in the past. Changes in a market or in its economic, social and political environment can cause significant changes in the actual shape of a pdf. Environmental change arising well outside the specific market in question is virtually impossible to predict, and experience has shown many examples of profound and sometimes sudden alterations. This source of error can be neither assessed nor rectified by statistical methods since future innovation has no manifestation in historical data. Unpredictable changes are ones which cannot be modelled probabilistically, and against which one can in no way insure in any actuarial sense. This uncertainty is what Frank Knight called a "true uncertainty" for which "there is no objective measure of the probability", as opposed to risk which is probabilistically measurable (Knight, 1921, pp.46, 120, 231–232). In this paper we focus on the unpredictable and non-stochastic Knightian uncertainties which accompany a VaR estimate.

We will use information-gap decision theory (Ben-Haim, 2001) to assess the robustness of the estimated VaR to Knightian uncertainty in the estimated pdf. The info-gap analysis focusses on the **robustness question:** how wrong, in the Knightian sense, can the estimated pdf be without jeopardizing the VaR estimate at a specified level of statistical confidence? A VaR which is highly robust to Knightian uncertainty in the pdf is a reliable assessment of risk, while a VaR whose robustness to Knightian error in the pdf is low is not a useful risk measure.

In section 2 we formulate a specific info-gap model for uncertainty in the pdf with which the VaR is estimated. The info-gap robustness function is formulated in section 3 and applied to an example of a normal VaR in section 4. The info-gap analysis leads to a robustness premium associated with an incremental VaR in section 5. Summary and discussion appears in section 6.

2 Info-gap Uncertainty

We begin with the definition of some basic quantities.

W is the initial portfolio value, in monetary units. V is the revenue from the portfolio in a single period, in monetary units. R is the rate of return $= V/W$.

$f(R)$ is a normalized pdf for R. $\widetilde{f}(R)$ is our best estimate of the pdf of R.

$R_\star$ is the lowest acceptable rate of return. This is a cutoff value, typically negative, so that $R_\star W$ equals the greatest tolerable monetary loss in a single period.

c is a confidence level, typically 0.01 or 0.05, expressing the greatest acceptable probability at which the return will be less than $R_\star$.

VaR $= R_\star W$ is the absolute Value at Risk. When $R_\star$ is chosen as the cth quantile of the estimated pdf $\widetilde{f}(R)$, we have confidence $1-c$ that the monetary loss will not be more negative than VaR (Jorion, 2001, p.22).

Our best estimate of the pdf of the rate of return is $\widetilde{f}(R)$, and it is with this pdf that we evaluate a VaR in order to assess risk. But an estimate is something short of the truth. There are both estimation errors and Knightian uncertainties in the pdf. It is with the Knightian uncertainties — non-stochastic and unpredictable errors — that we concern ourselves in this paper.

An **info-gap model of uncertainty** quantifies the gap between what we *do know* about the probability distribution of the rate of return — the estimated pdf $\widetilde{f}(R)$ — and what we *need to know* in order to make a truly responsible risk assessment — the true pdf $f(R)$. In this section we will present an info-gap model for doing this. Background discussion of the relation between Knight's conception of non-probabilistic uncertainty and info-gap theory is found in (Ben-Haim, 2001, section 12.5). Similarly, Shackle's "non-distributional uncertainty variable" bears some similarity to info-gap analysis (Shackle, 1972, p.23). Likewise, Kyburg recognized the possibility of a "decision theory that is based on some non-probabilistic measure of uncertainty." (Kyburg, 1990, p.1094).

An info-gap model is not a probabilistic quantification of uncertainty. An info-gap model depends on much less information than is needed to formulate and verify a probability model. An info-gap model is thus less informative than a probabilistic model. On the other hand an info-gap model can be implemented in situations with severe gaps in our knowledge of rare events: very low future returns.

An info-gap model for uncertainty in the pdf of the returns is a family of nested sets of pdfs. Each level of nesting corresponds to an **horizon of uncertainty** h**.** Each set contains all pdfs deemed feasible at that horizon of uncertainty. More pdfs are included as the horizon of uncertainty increases.

Info-gap models can be realized in many different forms (Ben-Haim, 2001). For example, the horizon of uncertainty may express the fractional error of the estimated pdf. In this case the info-gap model is formulated as follows.

Let $\mathcal{P}$ be the set of all normalized pdfs on the real numbers. Any pdf $f(R)$ must belong to $\mathcal{P}$. The **fractional error** info-gap model for Knightian uncertainty in the pdf of the returns is the following family of nested sets of pdfs:

$$\mathcal{F}(h, \widetilde{f}) = \left\{ f(R): \ f(R) \in \mathcal{P}, \ \left| f(R) - \widetilde{f}(R) \right| \le h\widetilde{f}(R) \text{ for all } R \right\}, \quad h \ge 0 \tag{1}$$

At horizon of uncertainty h, the set $\mathcal{F}(h, \widetilde{f})$ contains all pdfs $f(R)$ which differ from the best estimate $\widetilde{f}(R)$ by no more than a fraction h. As h gets larger, the uncertainty set $\mathcal{F}(h, \widetilde{f})$ becomes more inclusive, which is why h is the 'horizon of uncertainty'. The estimated pdf, $\widetilde{f}(F)$, belongs to $\mathcal{F}(h, \widetilde{f})$ at all horizons of uncertainty. In the absence of uncertainty, when $h = 0$, $\mathcal{F}(0, \widetilde{f})$ is the singleton set $\{\widetilde{f}\}$. Because the horizon of uncertainty h is not known, the info-gap model is not a single set but rather an unbounded family of nested sets of pdfs.

3 Info-gap Robustness

When we use the estimated pdf, $\widetilde{f}(R)$, to evaluate the VaR, we must ask ourselves: how confident are we in this estimate, in light of the info-gaps — Knightian uncertainties — in the estimated pdf? We will use an info-gap model to assess our confidence. To do this we first define the quantile function.

$q(c, f)$ is the cth quantile of the pdf $f(R)$. That is:

$$\int_{-\infty}^{q(c,f)} f(R)\,\mathrm{d}R = c \tag{2}$$

How do we evaluate the lowest acceptable return, $R_\star$? If we use the estimated pdf $\widetilde{f}(R)$, then $R_\star$, at level of confidence $1 - c$, is the cth quantile of $\widetilde{f}(R)$:

$$R_\star = q(c, \widetilde{f}) \tag{3}$$

Consequently the VaR which is calculated from the estimated pdf is:

$$\mathrm{VaR}(c, \widetilde{f}) = q(c, \widetilde{f})W \tag{4}$$

However, since the estimated pdf has unknown non-stochastic Knightian errors, we cannot be confident that the probability of returns less than $R_\star$ is in fact c. Nor can we be confident that eq.(4) is the true VaR. If the actual distribution of returns is given by $f(R)$, then the corresponding VaR at confidence $1 - c$ is:

$$\mathrm{VaR}(c, f) = q(c, f)W \tag{5}$$

We now return to the fundamental **robustness question** mentioned at the end of section 1. Let $R_\star$ be a specified value of the lowest acceptable rate of return. Typically, $R_\star$ will be dictated by management. The question of robustness to info-gaps in the pdf of R is: how wrong — in the Knightian sense — can $\widetilde{f}(R)$ be and still guarantee a VaR, at confidence $1 - c$, no less than $R_\star W$? The answer is the info-gap robustness function which is defined as:

$$\widehat{h}(R_\star, c) = \max\left\{h : \min_{f \in \mathcal{F}(h, \widetilde{f})} q(c, f) \geq R_\star\right\} \tag{6}$$

We can 'read' this equation from left to right as follows: the robustness $\widehat{h}(R_\star, c)$ of cutoff return $R_\star$ with required confidence $1 - c$ is the maximum horizon of uncertainty h up to which all pdfs $f(R)$ in $\mathcal{F}(h, \widetilde{f})$ have cutoff returns no less than $R_\star$. Stated differently, the robustness $\widehat{h}(R_\star, c)$ is the greatest Knightian uncertainty — as distinct from estimation error — up to which the value at risk, at confidence $1 - c$, is no less than $R_\star W$.

It is important to emphasize that the robustness $\widehat{h}(R_\star, c)$ is *not* a minimax algorithm. In minimax robustness analysis one *min*imizes the *max*imal adversity. This is not what the info-gap robustness function does. There is no maximal adversity in an info-gap model of uncertainty: the worst case at any horizon of uncertainty h is less damaging than some realization at a greater horizon of uncertainty. Since the horizon of uncertainty is unbounded, there is no worst case and the info-gap analysis cannot and does not purport to ameliorate a worst case. The info-gap robustness $\widehat{h}(R_\star, c)$ is the greatest horizon of uncertainty at which the worst return is no less than $R_\star$. This is not the same as identifying a worst possible return and ameliorating that case.

It is also important to emphasize that the info-gap robustness addresses Knightian uncertainty and not estimation error. Even if there were no estimation error (an idyllic situation) there would still be a need for the robustness function, which evaluates the vulnerability of the VaR estimate, based on historical data, to changes in future behavior.

It is readily shown that the robustness to info-gaps in the pdf decreases as the cutoff return $R_\star$ increases:

$$R_{\star 2} > R_{\star 1} \quad \text{implies} \quad \widehat{h}(R_{\star 2}, c) \leq \widehat{h}(R_{\star 1}, c) \tag{7}$$

This means that greater aspirations are more vulnerable than lesser ones: $R_{\star 2}$ is a less negative and hence more ambitious cutoff return than $R_{\star 1}$, and $R_{\star 2}$ entails lower immunity to uncertainty than $R_{\star 1}$.

Furthermore, the robustness becomes zero at the cutoff value corresponding to the best estimate of the pdf:

$$\widehat{h}(R_\star^o, c) = 0 \quad \text{if} \quad R_\star^o = q(c, \widetilde{f}) \tag{8}$$

$R_\star^o$ is the cutoff value calculated from the estimated pdf $\widetilde{f}(R)$. Relation (8) asserts that we have no immunity to Knightian uncertainty in the estimated pdf, when aspiring to a cutoff value as favorable as $R_\star^o$. Only less ambitious (more negative) cutoff values can be relied on with confidence $1 - c$ in light of the unknown Knightian errors in the estimated pdf. In terms of VaR, relation (8) means that the VaR calculated as the cth quantile of the best-estimated pdf, has no immunity to Knightian error in this pdf.

4 Example: Robustness of a Normal VaR

In the appendix, section 7, we derive an expression for the info-gap robustness of a calculated VaR, defined in eq.(6). The VaR is evaluated from a distribution of returns $\widetilde{f}(R)$ which is estimated to be normal with mean μ and variance σ^2. The robustness function can be derived analogously for other forms of the estimated pdf. We use the fractional-error info-gap model of eq.(1).

The robustness to uncertainty in the pdf, for cutoff return $R_\star$ at confidence $1 - c$, is:

$$\widehat{h}(R_\star, c) = \begin{cases} \dfrac{c}{\Phi\left(\frac{R_\star - \mu}{\sigma}\right)} - 1 & \text{if } \Phi\left(\dfrac{R_\star - \mu}{\sigma}\right) \leq c \\ 0 & \text{else} \end{cases} \tag{9}$$

$\Phi(\cdot)$ is the cumulative distribution function of the standard normal distribution. The 'else' condition holds if $R_\star \geq q(c, \widetilde{f})$. In eq.(8) we denoted this cutoff value $R_\star^o = q(c, \widetilde{f})$.

4.1 Numerical Results

The robustness of eq.(9) is evaluated numerically in fig. 1. The estimated pdf is normal with mean and variance $\mu = 0.05$ and $\sigma^2 = 0.01$. The robustness is plotted for three different values of the statistical confidence c: 0.01, 0.03 and 0.05.

The monotonic trade-off between immunity to info-gaps in the pdf, $\widehat{h}(R_\star, c)$, and cutoff return $R_\star$, is demonstrated by the strongly negative slopes of the curves. This is the graphical expression of eq.(7). In each case, as expected from eq.(8), the robustness equals zero when the cutoff value equals the cth quantile of the estimated distribution $\widetilde{f}$: $R_\star^o = -0.115$, -0.138 and -0.183 for $c = 0.05$, 0.03 and 0.01 respectively.

The negative slope of the robustness curves in fig. 1 can be thought of as a robustness cost of increasing the cutoff frequency. The slope of the essentially linear part of these curves is about -200. This means that the robustness must be reduced by 2 units in order to increase the cutoff value by 0.01. For instance, on the $c = 0.03$ curve, increasing the cutoff value from -0.22 to -0.21 reduces the robustness from 7.6 to 5.4. Robustness of 5.4 is still substantial: this VaR, $-0.21W$ at 97% confidence, is immune to 540% Knightian error in the estimated pdf. However, this robustness is lower by 220 percentage points than the VaR at -0.22 which, also with 97% confidence, is immune

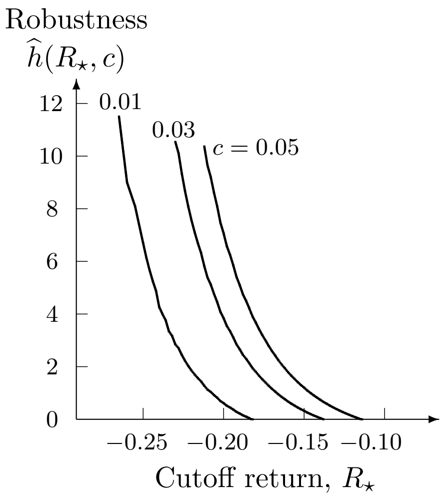

Figure 1: Robustness vs. cutoff return for three levels of confidence. $\widetilde{f} = \mathcal{N}(0.05, 0.01)$.

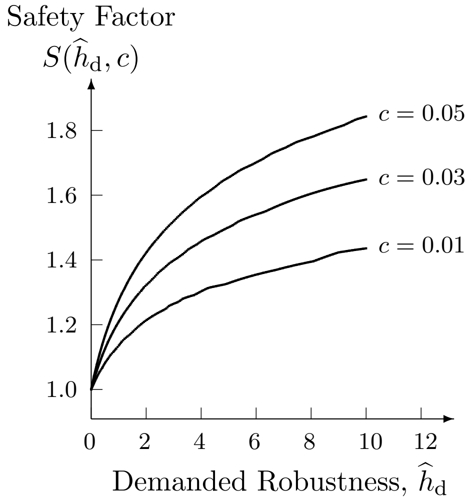

Figure 2: Safety factor vs. demanded robustness for three levels of confidence. $\widetilde{f} = \mathcal{N}(0.05, 0.01)$.

to 760% error in the estimated pdf. The VaR anticipated from the estimated pdf, $-0.138W$ at 97% confidence, has zero robustness to Knightian error.

4.2 Safety Factor

Our best-estimated pdf, $\widetilde{f}(R)$, implies a VaR value of $R_\star^o W$ at confidence $1-c$, where $R_\star^o = q(c, \widetilde{f})$. However, the robustness to info-gaps in the pdf is zero for this VaR, as asserted in eq.(8). We must 'migrate' up the robustness curve, relinquishing performance (accepting more negative cutoff $R_\star$) in order to gain positive robustness to Knightian uncertainty in the pdf. In this section we derive a "safety factor" for expressing how much performance to relinquish.

The analyst must choose a probability value c to express the demanded level of confidence with respect to the random variation of R. In like manner the analyst chooses a value $\widehat{h}_{\rm d}$ of robustness to express the demanded confidence with respect to Knightian info-gaps in the pdf of R. For instance, choosing $\widehat{h}_{\rm d} = 3$ means that the analyst seeks a cutoff return $R_\star$ which, with probabilistic confidence $1-c$, is the worst (least) rate of return for all pdfs whose fractional deviation from $\widetilde{f}(R)$ is no greater than 300%.

Since $\widehat{h}(R_\star, c)$ is a monotonically decreasing function of $R_\star$, any positive choice of $\widehat{h}_{\rm d}$ translates into a unique choice of $R_\star$. Let us denote this choice by $R_\star(\widehat{h}_{\rm d})$. From eq.(9) one can show that this value is in fact a quantile of $\widetilde{f}$:

$$R_\star(\widehat{h}_{\rm d}) = q\left(\frac{c}{\widehat{h}_{\rm d}+1},\ \widetilde{f}\right) \tag{10}$$

$R_\star(\widehat{h}_{\rm d})$ is the cutoff value which, at level of confidence $1-c$, has info-gap robustness $\widehat{h}_{\rm d}$. In contrast, $R_\star^o$ is the cutoff value based on the estimated pdf, and has zero info-gap robustness. The comparison of $R_\star(\widehat{h}_{\rm d})$ against $R_\star^o$ is a safety factor:

$$S(\widehat{h}_{\rm d}, c) = \frac{R_\star(\widehat{h}_{\rm d})}{R_\star^o} = \frac{q\left(\dfrac{c}{\widehat{h}_{\rm d}+1},\ \widetilde{f}\right)}{q(c, \widetilde{f})} \tag{11}$$

$S(\widehat{h}_{\rm d}, c)$ is the factor by which the estimated cutoff value $R_\star^o$ (based on the estimated pdf $\widetilde{f}$) should be multiplied in order to obtain a VaR, at probabilistic confidence $1-c$, which has info-gap robustness equal to $\widehat{h}_{\rm d}$. In other words,

$$\text{VaR} = R_\star^o W \tag{12}$$

is a VaR at $1-c$ statistical confidence but with zero robustness to Knightian uncertainty in the pdf, while

$$\mathrm{VaR} = S(\widehat{h}_{\rm d}, c) R_\star^o W \tag{13}$$

is a VaR at $1-c$ confidence with robustness $\widehat{h}_{\rm d}$.

For computational purposes it is more convenient to express the safety factor in terms of quantiles of the standard normal distribution, whose density is denoted ϕ. Eq.(11) becomes:

$$S(\widehat{h}_{\rm d}, c) = \frac{\mu + \sigma q\left(\dfrac{c}{\widehat{h}_{\rm d}+1},\ \phi\right)}{\mu + \sigma q(c, \phi)} \tag{14}$$

where μ and σ^2 are the mean and variance of the estimated normal distribution $\widetilde{f}(R)$.

The safety factor is evaluated numerically in fig. 2. Consider the 97% confidence curve, $c = 0.03$. The safety factor for a demanded robustness of 2 is $S(2, 0.03) = 1.32$. In order to guarantee immunity to 200% errors in the estimated pdf, the cutoff for the estimated VaR at 97% confidence, $R_\star^o$, should be multiplied by 1.32. Thus $R_\star = 1.32 \times (-0.138) = -0.182$ is the cutoff value for a VaR at 97% with robustness to 200% Knightian error in the pdf. The safety factors for 600% and 1000% robustness to pdf error, $\widehat{h}_{\rm d} = 6$ and 10 respectively, are 1.54 and 1.65 (at $c = 0.03$).

Perhaps the most striking aspect of the safety factors in fig. 2 is that $S(\widehat{h}_{\rm d}, c)$ *decreases* as the level of confidence $1-c$ *increases.* The curve for $c = 0.05$ lies above the curve for $c = 0.03$ which lies above the curve for $c = 0.01$. In the current example, for any demanded robustness $\widehat{h}_{\rm d}$ one needs a smaller safety factor for 99% confidence than for 97% which is smaller than for 95% confidence. We can understand this from eq.(11), where the numerator is more negative than the denominator. An increase in c causes the numerator to become less negative proportionately less than the denominator. As c increases, both quantiles, $R_\star(\widehat{h}_{\rm d})$ and $R_\star^o$, become less negative, the former less so than the latter.

This can be explained more intuitively as follows. The estimated cutoff at $c = 0.01$ is $R_\star^o = -0.183$. The robustness to Knightian error in $\widetilde{f}$ is zero for this value, but since it is fairly negative it needs a relatively small correction to reach a cutoff with given positive robustness $\widehat{h}_{\rm d}$. The estimated cutoff at $c = 0.03$ is $R_\star^o = -0.138$, and a greater correction is needed to reach a cutoff with the same positive robustness.

5 Incremental VaR

The incremental VaR is a valuable tool for comparing portfolios, for instance in assessing the risk-implications of including or removing a particular asset or set of assets (Dowd, 1998, p.48). A portfolio change which results in less negative VaR is desirable; a change which makes the VaR more negative is not.

However, VaRs are evaluated from estimated pdfs. In sections 3 and 4 we have seen that a VaR calculated as the cth quantile of an estimated pdf has no robustness against Knightian error in that pdf. Only a more negative VaR will have positive robustness. The analyst faces an irrevocable trade-off between performance and robustness when estimating VaRs from distributions with Knightian uncertainty.

In this section we will use the info-gap robustness analysis to evaluate incremental VaRs. We will find that, under certain circumstances, the robustness curves can cross, which implies a reversal of the preferences between portfolios indicated by the estimated VaRs.

5.1 Robustness Premium

Consider two portfolios whose rates of return have different probability distributions. Let $\widetilde{f}_1(R)$ and $\widetilde{f}_2(R)$ represent the estimated pdfs of the returns for these portfolios. As before, let $q(c, f)$ denote the cth quantile of the pdf $f(R)$. The VaRs of the estimated distributions are $q(c, \widetilde{f}_1)W$ and $q(c, \widetilde{f}_2)W$

where W is the value of each portfolio. The incremental VaR is the difference between these two values, and portfolio i is preferred if $q(c, \widetilde{f}_i) > q(c, \widetilde{f}_j)$ where $i = 1$ or 2 and $j = 3 - i$.

Let $\mathcal{F}_i(h, \widetilde{f}_i)$ denote an **info-gap model** for the uncertainty in the estimated pdf $\widetilde{f}_i(R)$.

Let $R_\star$ be a cutoff value of the rate of return, and $1-c$ a statistical confidence level for that cutoff value. As in eq.(6), the **robustness** to Knightian uncertainty in the pdf, of portfolio i with cutoff $R_\star$ and confidence $1 - c$, is the greatest horizon of uncertainty in info-gap model $\mathcal{F}_i(h, \widetilde{f}_i)$ up to which all pdfs have quantiles no less than $R_\star$:

$$\widehat{h}_i(R_\star, c) = \max \left\{ h : \min_{f \in \mathcal{F}_i(h, \widetilde{f}_i)} q(c, f) \geq R_\star \right\} \tag{15}$$

From eq.(7) we know that $\widehat{h}_i(R_\star, c)$ decreases as $R_\star$ becomes less negative: robustness trades off against performance. From eq.(8) we know that the robustness is zero for a VaR based on the estimated pdf: $\widehat{h}_i(R_\star, c) = 0$ if $R_\star = q(c, \widetilde{f}_i)$.

The total lack of robustness to Knightian uncertainty in the pdf, of the estimate $q(c, \widetilde{f}_i)$, suggests that ranking the portfolios according to these quantiles may be unreliable. If portfolio i is more robust than portfolio j, over a substantial and relevant range of cutoff values $R_\star$, then i should be preferred over j. This leads us to define the **robustness premium:**

$$\pi_{ij}(R_\star, c) = \widehat{h}_i(R_\star, c) - \widehat{h}_j(R_\star, c) \tag{16}$$

The use of the robustness premium in comparing portfolios is best illustrated with an example, to which we now proceed.

5.2 Example: Incremental Normal VaR

Suppose that the estimated pdf of the rate of return of the ith portfolio, $\widetilde{f}_i(R)$, is normal with mean μ_i and variance σ_i^2, for $i = 1$ and 2. Furthermore let the info-gap model for uncertainty in this pdf be the fractional error model of eq.(1). Thus the robustness function for portfolio i, $\widehat{h}_i(R_\star, c)$, is given by eq.(9). In the subsequent discussion we will allow i to be either 1 or 2 and $j = 3 - i$.

The relative disposition of the robustness curves $\widehat{h}_i(R_\star, c)$ and $\widehat{h}_j(R_\star, c)$ — which curve lies above the other and whether they cross at some value of $R_\star$ — determines the robustness premium $\pi_{ij}(R_\star, c)$ of portfolio i over j. This in turn depends on the standardized cutoff values, $(R_\star - \mu_i)/\sigma_i$. From eq.(9) we learn that:

$$\frac{R_\star - \mu_i}{\sigma_i} < \frac{R_\star - \mu_j}{\sigma_j} \iff \Phi\left(\frac{R_\star - \mu_i}{\sigma_i}\right) < \Phi\left(\frac{R_\star - \mu_j}{\sigma_j}\right) \iff \widehat{h}_i(R_\star, c) \geq \widehat{h}_j(R_\star, c) \tag{17}$$

where '$\iff$' means 'is equivalent to' or 'if and only if'. The inequality between the robustnesses, $\widehat{h}_i$ and $\widehat{h}_j$, is strict, that is '$>$' rather than '$\geq$', if $\Phi\left(\frac{R_\star - \mu_i}{\sigma_i}\right) < c$ because then $\widehat{h}_i(R_\star, c) > 0$.

The portent of relations (17) is illustrated in figs. 3 and 4. In both figures portfolio i is less volatile than portfolio j: $\sigma_i < \sigma_j$. However, in fig. 3 the mean estimated return of portfolio i is less than the mean estimated return of portfolio j, $\mu_i < \mu_j$, while the means are reversed in fig. 4.

The curves of standardized cutoff in figs. 3 and 4 cross when $R_\star = \mu_i - \frac{\delta}{\varepsilon}\sigma_i$ (δ and ε are defined in the figure captions). Let R_c denote this value of $R_\star$. For values of $R_\star$ less than R_c, relations (17) imply that portfolio i is more robust than portfolio j, and less robust for $R_\star$ greater than R_c.

Consider fig. 3. Portfolio i has lower estimated mean return and, for $R_\star > R_\mathrm{c}$, also has less robust VaR than portfolio j, together suggesting that j is preferable over i. However, the robustness comparison is reversed for $R_\star < R_\mathrm{c}$, suggesting the need to balance more reliable VaR (in favor of i) against lower mean return (mitigating against i) in choosing between the portfolios.

The situation is less ambiguous in fig. 4. The only relevant cutoff values $R_\star$ are less than R_c, and portfolio i is more robust than portfolio j throughout this range, as well as having greater estimated mean return. Portfolio i is the clear choice.

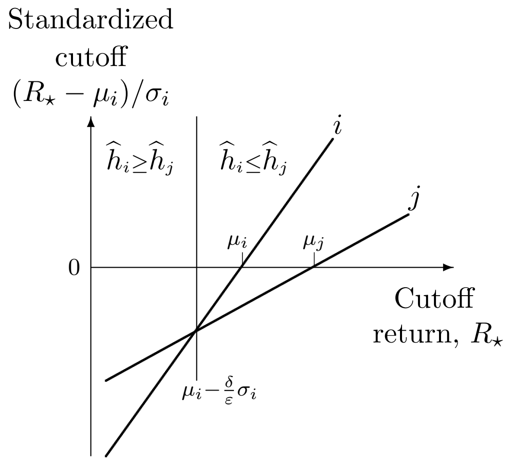

Figure 3: Standardized cutoff vs. cutoff. $\sigma_j = \sigma_i + \varepsilon$, $\varepsilon > 0$. $\mu_j = \mu_i + \delta$, $\delta > 0$.

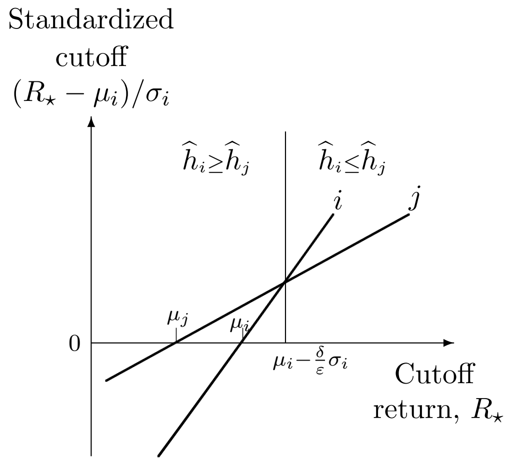

Figure 4: Standardized cutoff vs. cutoff. $\sigma_j = \sigma_i + \varepsilon$, $\varepsilon > 0$. $\mu_j = \mu_i + \delta$, $\delta < 0$.

Robustness curves for two portfolios are shown in fig. 5 and again in fig. 6 for a greater range of cutoff values. Portfolio i is estimated to be less volatile than portfolio j, but i is also estimated to have a lower mean return than j which corresponds to the situation in fig. 3. The VaR which is based on the estimated pdf is more negative for portfolio i than for portfolio j: the estimated 95% quantiles are $q(0.05, \widetilde{f}_i) = -0.118$ while $q(0.05, \widetilde{f}_j) = -0.114$. Based on these VaR estimates one should prefer j over i.

However, the robustnesses of these estimated VaRs are zero (eq.(8) again). Also, we see from fig. 5 that the robustness curves of these portfolios cross at a cutoff of $R_c = \mu_i - \frac{\delta}{\varepsilon}\sigma_i = -0.15$. Portfolio j is more robust than portfolio i for $R_\star > -0.15$, but the robustness is rather low, and the robustness premium for j over i is small. However, for $R_\star < -0.15$ the robustness premium for i over j gets increasingly large. As we see in fig. 6, at $R_\star = -0.2$ the robustnesses are $\widehat{h}_i = 8.47$ and $\widehat{h}_j = 7.06$ so the robustness premium of portfolio i over j is $\pi_{ij} = 1.41$. That is, for a VaR of $-0.2W$ at 95% confidence, portfolio i has 141 percentage points more robustness to Knightian uncertainty in the pdf than portfolio j.

From these considerations the cautious analyst may choose portfolio i over portfolio j. Even though i is estimated to have lower mean, its estimated volatility is lower and the risk estimates at low VaRs are substantially more robust for i than for j.

6 Summary and Discussion

VaR assesses financial risk by evaluating the probability of loss resulting from stochastic variation of the rate of return. VaR is based, in one way or another, on historical data reflecting this variation, usually as an estimated pdf. This paper is devoted to supplementing VaR with tools for evaluating the impact of Knightian uncertainty in the future realizations of this estimated pdf. As such, the robustness and robustness premium developed here can be used as stress-testing methods. In addition, the robustness together with the safety factor can assist in evaluating and regulating investment.

6.1 Summary

VaR is a prediction of the future based on evidence from the past. We often use this evidence to estimate a pdf. *Future* innovations, surprises and structural or environmental changes of all sorts cannot, by definition, be manifested in *historical* data. This Knightian uncertainty is non-

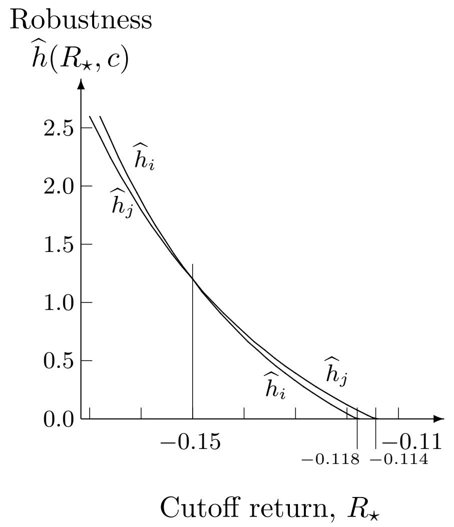

Figure 5: Robustness vs. cutoff return for two portfolios. $\mu_i = 0.03$, $\sigma_i = 0.09$, $\mu_j = 0.05$, $\sigma_j = 0.10$. $c = 0.05$.

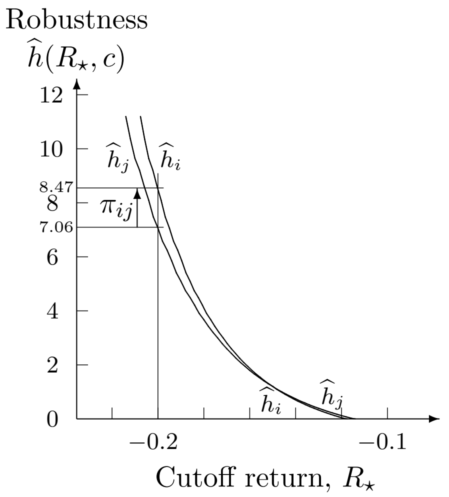

Figure 6: Robustness vs. cutoff return for two portfolios. $\mu_i = 0.03$, $\sigma_i = 0.09$, $\mu_j = 0.05$, $\sigma_j = 0.10$. $c = 0.05$.

probabilistic and is particularly pertinent to the far lower tail of the distribution. We have used **info-gap models** to quantify this highly unstructured Knightian uncertainty. Our examples have been based on the fractional-error info-gap model of eq.(1). An info-gap model is an unbounded family of nested sets of pdfs, $\mathcal{F}(h, \widetilde{f})$, $h \geq 0$. The uncertainty sets become more inclusive as the horizon of uncertainty h grows. There is no greatest horizon of uncertainty so there is no worst case in an info-gap model.

The **info-gap robustness function** $\widehat{h}(R_\star, c)$ is defined in eq.(6). The robustness is the greatest horizon of uncertainty in the pdf, up to which all pdfs have VaRs no less than a specified value $R_\star W$, at a specified level of confidence $1 - c$. A large robustness $\widehat{h}(R_\star, c)$ means that VaR $= R_\star W$ is highly immune to Knightian uncertainty in the pdf and is reliable as a guide to investment. Low robustness implies vulnerability of this VaR to Knightian uncertainty and indicates that it is an unreliable assessment of risk. The info-gap robustness function evaluates the greatest tolerable Knightian uncertainty. However, the info-gap robustness is not a minimax analysis (we do not ameliorate a worst case) since, in an info-gap model of uncertainty, there is no worst case: any given horizon of uncertainty is less severe than all greater horizons of uncertainty.

The robustness function has two critical properties. First, robustness trades-off against the value of $R_\star$: as $R_\star$ becomes less negative the robustness becomes smaller, eq.(7). Second, a VaR evaluated as a quantile of an estimated distribution has zero robustness to error in that pdf, eq.(8). Combining these two properties we see that a VaR based on a 'best estimate' of the pdf is completely unreliable vis á vis Knightian uncertainty. Only more negative and less optimistic VaRs have positive immunity to Knightian error in the pdf. These properties are illustrated in fig. 1.

Let $R_\star^o$ be the cutoff value at confidence $1 - c$ evaluated as the cth quantile of an estimated pdf. The corresponding VaR is $R_\star^o W$, which has zero robustness to error in the pdf. The **info-gap safety factor** $S(\widehat{h}_{\rm d}, c)$ is the factor by which $R_\star^o$ should be multiplied in order to obtain a VaR at statistical confidence $1 - c$ which has info-gap robustness equal to $\widehat{h}_{\rm d}$, eq.(11). The safety factor leads to a VaR estimate which deals with both the probabilistic risk and the Knightian uncertainty in the pdf. Safety factors are illustrated in fig. 2.

The upshot of the info-gap analysis of VaR is that one should consider an entire robustness curve, not only a single estimate based on the estimated pdf; the usual single estimate is the endpoint of the robustness curve and has zero robustness. This has practical implications when using the VaR to compare portfolios. Specifically, robustness curves of different portfolios can cross, as in fig. 5. Intersection of robustness curves implies that the preference between the portfolios depends on the demanded robustness to uncertainty in the future realizations of the pdf. Equivalently, the preference depends on the desired cutoff $R_\star$. In fig. 5, the **robustness premium** favors portfolio i over j for $R_\star < -0.15$, and favors j otherwise.

6.2 Stress Testing

Dowd adopts a healthy attitude of pluralism in characterizing stress testing as "a variety of different procedures that attempt to gauge the vulnerability of our portfolio to hypothetical events." (Dowd 1998, p.121). Dowd discusses stress testing by analysis of selected scenarios as well as by various extreme-value and worst-case analyses. When trying to anticipate the unknown it is necessary to implement a range of conceptually distinct methods because each method captures a different aspect of the situation.

Info-gap robustness analysis is a stress testing tool with similarities to the maximum-loss and worst-case methods. However, it is important to point out two fundamental differences between info-gap analysis and these methods, differences which make the info-gap approach a useful supplement. First, the robustness function is *not* a worst-case or minimax assessment. There is no worst case in an info-gap model of uncertainty: as the horizon of uncertainty h grows, the uncertainty sets $\mathcal{F}(h, \widetilde{f})$ become more inclusive. The robustness function $\widehat{h}(R_\star, c)$ does not identify a worst case. What is evaluated is the greatest horizon of uncertainty up to which the performance is acceptable. This in no way asserts that the real variation is acceptable. The utility of the robustness function is in comparing alternative investments in order to determine which is more immune and which is less, and in assessing capital requirements in terms of Knightian uncertainty in the estimated pdf.

The second basic difference between info-gap and extreme-value methods is that info-gap analysis deals non-probabilistically with severe uncertainty. An info-gap model quantifies the Knightian uncertainty — the lack of information and understanding of unmeasured future changes or surprises — which accompanies an estimate of the pdf. The info-gap robustness assesses the impact of Knightian uncertainty without introducing measure functions or probabilistic assumptions and requirements such as normality or large samples. For instance, an info-gap analysis has no requirement analogous to the asymptotic character of extreme value distributions. Likewise, info-gap analysis is "non-parametric" in a stronger sense than "non-parametric statistics": an info-gap model quantifies uncertainty without measure functions at all. Info-gap modelling and analysis is a sparse, stark and skeptical approach to representation and management of ignorance (Ben-Haim, 2004).

7 Appendix: Derivation of the Robustness Function, eq.(9)

First note that the following two inequalities are equivalent:

$$q(c, f) \geq R_\star \quad \Longleftrightarrow \quad c \geq \int_{-\infty}^{R_\star} f(R)\,\mathrm{d}R \tag{18}$$

Provided that c is small, which in practice will always be the case, it is evident that the minimum of the cth quantile of f, $\min_{f \in \mathcal{F}(h, \widetilde{f})} q(c, f)$ in the definition of the robustness in eq.(6), occurs when the far lower tail is as 'fat' as possible at horizon of uncertainty h. That is, the minimum quantile occurs when the far lower tail of the unknown distribution is:

$$f(R) = (1 + h)\widetilde{f}(R) \tag{19}$$

Because c is small this holds on the far lower tail without violating the normalization requirement of a pdf.

The robustness, $\widehat{h}(R_\star, c)$, is the greatest horizon of uncertainty, h, up to which the cth quantile is no less than $R_\star$. That is, using eqs.(18) and (19), $\widehat{h}(R_\star, c)$ is the least upper bound of the set of h-values for which:

$$c \geq \int_{-\infty}^{R_\star} (1 + h)\widetilde{f}(R)\,\mathrm{d}R \tag{20}$$

If $q(c, \widetilde{f})$, the cth quantile of the estimated distribution $\widetilde{f}$, is less than (more negative than) $R_\star$, then the integral in eq.(20) exceeds c when $h = 0$. In this case the robustness is zero. On the other hand, if $q(c, \widetilde{f}) \geq R_\star$, then the robustness $\widehat{h}(R_\star, c)$ satisfies:

$$c = (1+\widehat{h}) \int_{-\infty}^{R_\star} \widetilde{f}(R)\, \mathrm{d}R \quad (21)$$

$$= (1+\widehat{h}) \Phi \left(\frac{R_\star - \mu}{\sigma} \right) \quad (22)$$

where $\Phi(\cdot)$ is the cumulative distribution function of the standard normal distribution. Solving eq.(22) for $\widehat{h}$ yields eq.(9).

Acknowledgement

The author is indebted to Prof. Jacob Boudoukh for valuable comments and criticisms.

8 References

Ben-Haim, Yakov (2001), *Information-gap Decision Theory: Decisions Under Severe Uncertainty,* Academic Press, San Diego.

Ben-Haim, Yakov (2004), Uncertainty, probability and information-gaps, *Reliability Engineering and System Safety,* 85: 249–266.

DeGroot, Morris H. (1986), *Probability and Statistics,* 2nd ed., Addison-Wesley, Reading, MA.

Dowd, Kevin (1998), *Beyond Value at Risk: The New Science of Risk Management,* Wiley, New York.

Fama, Eugene F. (1965), The behavior of stock-market prices, *The Journal of Business,* vol. 38, no. 1, pp.34–105.

Fama, Eugene F. and Kenneth R. French (1988), Permanent and temporary components of stock prices, *Journal of Political Economy,* vol 96, no. 2, pp.246–273.

Härdle, Wolfgang (1990), *Applied Nonparametric Regression,* Cambridge University Press, Cambridge, U.K.

Hendricks, Darryll (1996), Evaluation of value-at-risk models using historical data, *Economic Policy Review,* Federal Reserve Bank of New York, April 1996, Vol. 2, no. 1.

Jorion, Philippe (2001), *Value at Risk: The New Benchmark for Managing Financial Risk,* McGraw-Hill, New York.

Knight, Frank H. (1921), *Risk, Uncertainty and Profit.* Houghton Mifflin Co. Re-issued by University of Chicago Press, 1971.

Kendall, Maurice, Alan Stuart and J. Keith Ord (1987), *The Advanced Theory of Statistics.* Vol. 1: *Distribution Theory,* 5th ed., Oxford University Press, New York.

Kendall, Maurice, Alan Stuart and J. Keith Ord (1979), *The Advanced Theory of Statistics.* Vol. 2: *Inference and Relationship,* 4th ed., Charles Griffin & Co., London.

Kyburg, Henry E. jr. (1990), Getting fancy with probability, *Synthese,* 90: 189–203.

Moore, Walter (1989), *Schrödinger: Life and Thought,* Cambridge University Press, Cambridge.

Shackle, G.L.S. (1972), *Epistemics and Economics: A Critique of Economic Doctrines,* Cambridge University Press, re-issued by Transaction Publishers, New Brunswick, 1992.