Bank of Canada   Banque du Canada

Risk-Cost Frontier and Collateral Valuation in Securities Settlement Systems for Extreme Market Events

by

Alejandro García and Ramazan Gençay

Risk-Cost Frontier and Collateral Valuation in Securities Settlement Systems for Extreme Market Events

by

Alejandro García[1] and Ramazan Gençay[2]

[1]Monetary and Financial Analysis Department
Bank of Canada
Ottawa, Ontario, Canada K1A 0G9
agarcia@bankofcanada.ca

[2]Department of Economics
Simon Fraser University
Burnaby, British Columbia, Canada V5A 1S6
rgencay@sfu.ca

Contents

Acknowledgements

Abstract

The authors examine how the use of extreme value theory yields collateral requirements that are robust to extreme fluctuations in the market price of the asset used as collateral. In particular, they study the risk and cost attributes of market risk measures by constructing a risk-cost frontier for the collateral pledged to cover exposures in a securities settlement system. The frontier can be used as a diagnostic tool to understand the risk-cost trade-off of different methodologies to calculate collateral value (haircuts) and select the most efficient alternative in a variety of settings.

JEL classification: G0, G1, C1
Bank classification: Financial stability; Payment, clearing, and settlement systems; Econometric and statistical methods

Résumé

Les auteurs examinent comment la théorie des valeurs extrêmes permet d'obtenir une évaluation robuste du montant de la garantie nécessaire à la couverture des fluctuations extrêmes de la valeur de marché de l'actif remis en nantissement. Ils étudient en particulier les caractéristiques des risques et des coûts propres aux mesures du risque de marché en construisant une frontière risque-coût pour la garantie utilisée en couverture des risques liés à un système de règlement de titres. Cette frontière peut servir d'outil de diagnostic pour comprendre l'arbitrage risque-coût associé aux diverses méthodes de calcul des décotes et déterminer, pour des cadres différents, le choix le plus efficient.

Classification JEL : G0, G1, C1
Classification de la Banque : Stabilité financière; Système de paiement, de compensation et de règlement; Méthodes économétriques et statistiques

1. Introduction

Clearing and settlement systems play a critical role in the infrastructure of financial markets, and in recent years there has been increased attention on settlement risk associated with them. This has led to the development of international standards of risk control for different kinds of clearing and settlement arrangements, such as the recommendations for securities settlement systems (BIS 2001), and the recommendations for central counterparties (BIS 2004). A key element of risk management in these recommendations is the pledging of collateral by participants to cover risk that they bring to the system. Pledging collateral, however, is costly to participants and therefore it is important to balance risk and efficiency issues. A critical component that affects this balance is the valuation of the collateral. An accurate valuation of collateral is critical because there is a delay between the time that a participant pledges the collateral and the time at which the collateral may have to be used to cover money owing by a defaulting participant. During that time, the collateral may change in value. For this reason, the total value of collateral is discounted, or 'haircutted' to take account of the risk of future price declines. How these haircuts are calculated is critical to both good risk management (holding sufficient collateral) and achieving an adequate cost level to the participants of the system. The greater the haircut, the greater the total amount of collateral that must be pledged.

This paper proposes a framework to study and compare the risk and cost attributes of commonly used practices to calculate haircuts, such as parametric methods based on the normal distribution to calculate Value-at-Risk (VaR), and recent methods of capturing tail events based on extreme value theory (EVT). Furthermore we assess whether it is likely that EVT methods bring benefits in terms of risk coverage and efficiency for the participants in the system. We apply these EVT methods to VaR and Expected Shortfall (ES) (a coherent alternative to VaR), in the context of a securities settlement system.[1] We evaluate these methodologies and risk measures using the proposed framework for assessing the risk-cost trade-off.

In our study we focus on equities used as collateral. For equities a mismeasurement of risk can occur when estimation methods assume normal distributions which are used to estimate haircuts for securities whose prices exhibit fat-tailed distributions. Equities are more likely to exhibit fat tails than securities such as debt.

Such mismeasurement of risk is most likely when extreme market events occur (such

---

[1]A *coherent* risk measure is defined in section 4.

as a drastic price drop in equity markets) and the returns of the equity instruments take on values from the extreme tail of the return distribution. Moreover, if defaults in the system are also more likely during extreme market events and the liquidation of collateral may be more frequent, then it is important to reduce the risk of uncollateralized exposures by accurately measuring the tail (and therefore the corresponding risk) when calculating an adequate haircut for a collateral instrument. This is a necessary requirement for the robustness of the system.[2]

Reducing collateral cost to participants of the system is also important given the significant amount of resources that are devoted to allowing trades to settle. Although the cost of collateral may be less than the total values traded owing to collateral-sharing arrangements such as collateral pools, as well as to the use of netting positions, the cost of collateral remains large owing to the sheer size of securities transactions in a given year. Table 1 illustrates that 40.7 trillion Canadian dollars (approximately U.S. 29 trillion) were settled in 2003 by the Canadian Securities Settlement System (CDSX). This is 33 times the Canadian GDP for 2003 and thus justifies the importance of (i) accurately measuring risks,[3] and (ii) designing efficient settlement systems for improving the welfare of participants and the system as a whole.

	1999	2000	2001	2002	2003
Value	25.6	23.9	31.7	26.1	40.7
Value/GDP (times)	26	22	29	23	33

Table 1: VALUE OF SECURITIES TRANSACTIONS IN CANADA

This table shows the value of trades in trillions of Canadian dollars (first row) and as a ratio of GDP (second row). Source: BIS statistics on payment and settlement systems in selected countries, March 2005.

The contributions of this paper are (i) to show how *extreme value theory* leads to efficient measures of haircuts that adequately reflect the risk derived from the tail of the return distribution, (ii) to propose a framework to study the risk and cost trade-off for a given risk measure used to calculate haircuts in a securities settlement system, and (iii) to show the robustness of extreme value theory when calculating haircuts at high quantiles of the return distribution. This last point implies that a risk measure may provide adequate coverage to the market risk of the underlying asset used as collateral, especially at high quantiles of the

---

[2]A default is understood in this context as the failure of the buyer to supply the funds part associated with a given trade.

[3]There are a number of risks associated with securities trading. We will focus on two sources: *settlement risk* associated with the funds part of the trade, and *market risk* associated with the collateral used to support settlement risk.

return distribution, when EVT methods are employed. The contributions of our paper are not based on any particular securities settlement system. Our paper presents an overview of techniques, and sets up a methodology to study the valuation of collateral in a variety of settings.

This paper is organized as follows. In section 2, we study the nature of the risks that are being measured and describe the role of collateral in a securities settlement system. In section 3, we introduce a framework to assess risk and efficiency issues associated with the calculation of the value of collateral. In section 4, we provide a concise review of the literature related to risk measurement, where emphasis is placed on presenting a framework that defines the characteristics of a coherent risk measure, such as ES. Section 5 describes several methods of estimation associated with the different risk measures. Section 6 compares three methodologies to calculate haircuts; these are (i) VaR with a normal distribution, (ii) VaR with a distribution based on extreme value theory (EVT) methods, and (iii) ES with EVT methods. We compare these methodologies by constructing a frontier for the risk-cost trade-off in the system. The final section offers some conclusions.

This paper is part of a two-year research program. With this paper we propose a framework to study collateral valuation; we do this by linking the literature related to risk measurement, extreme value theory, and the infrastructure of the financial system. A future paper will extend the framework proposed in this paper to (i) address a *portfolio* of collateral rather than a single asset, (ii) expand the analysis to include *debt* instruments in the collateral portfolio, and (iii) test our findings using *actual data* of collateral instruments used in the Canadian securities settlement system.

2. The Economics of Securities Settlement Systems

In many countries, a central securities depository (CSD) or the central bank operates the system that facilitates the provision of securities settlement services. This system is an electronic platform governed by a set of rules and procedures that allows the exchange of funds for securities in a *secure* and *efficient* manner and with a desirable balance between these two characteristics. These systems have two main processes: *clearing* and *settlement*. Clearing refers to the process that determines the securities and funds obligations for each customer in the system (basically, "who owes what to whom"), and settlement refers to the process of transferring the security ownership and the corresponding funds to the respective parties.

The risks and costs of securities settlement systems affect the decisions of participants in the system to trade and settle securities. There are many different kinds of risks and costs associated with securities settlement systems:

- For the seller of securities, there is a risk of not receiving the funds after sending the securities.

- For the buyer of securities, there is a risk of not receiving the securities after sending the corresponding funds. There is also the cost of collateral to support the payment risk associated with the funds owed from trades that will be settled during the day.

- For all participants, there may be risks and costs associated with residual obligations resulting from the default of another participant. If large enough, such costs may lead to secondary defaults.

- For regulators, risk and cost for the entire system are important, since they are concerned with preserving the financial stability of the system as well as enhancing the welfare of all participants in the system.

Finding the right balance between risk and cost in securities settlement systems is very important for a well-functioning financial system. For example, a system that is highly secure but also inefficient (costly) could result in customers abandoning its use and perhaps switching to more risky practices to process their trades. Similarly, an efficient system (i.e., low transaction/collateral costs) that is also very risky may mean that a participant default could cause knock-on effects for other participants and lead to financial instability. From the perspective of a policy-maker, neither system is desirable for achieving a social optimum. However, at the margin there may be a trade-off between further improvements in risk-proofing and efficiency. In section 3, we provide a simple framework in the spirit of Berger, Hancock, and Marquardt (1996) to illustrate the trade-off between the risk and costs in the securities settlement system.

2.1 The role of collateral

In general, collateral is used as an instrument that mitigates the risk of financial losses. In securities settlement systems, buyers, sellers, and the CSD are all exposed in some form or another to financial losses when a participant in the system fails to pay/transfer either the funds or securities corresponding to the trade.

Consider the following example that illustrates the role of collateral. We assume a system where securities are settled on a *gross basis*, with final transfer of securities occurring *intraday*, and funds transfers settled on a *net basis*, with final transfer of funds occurring at the *end of the day*. In this system, a trade is entered into between Broker A (the seller of securities) and Broker B (the buyer of securities). Both parties to the trade have a securities account with the CSD; A has a balance of 500 Company X securities, and B has no Company X securities. A trade is created when A agrees to sell 200 Company X securities to B for a price of \$1,000. B may or may not have the funds to complete the trade; if it does not have them, the system allows it to enter into a negative (debit) funds position. For this example, we assume that B incurs such a debit funds position. This creates a risk to A, because once the securities are transferred to B, there is the risk that B may not have the funds to pay A. We refer to this as *payment risk*.

Definition 1 (Payment Risk) *Payment risk is the risk that a participant in the system defaults on its funds obligation at the end of the day.*

To eliminate payment risk for the seller, the CSD provides a guarantee to the seller that once the trade is approved for settlement and the securities are transferred, the funds will be transferred to the seller at the end of the day, even if the buyer defaults on its obligation. To provide this guarantee, the CSD manages payment risk by requiring B (the buyer) to pledge collateral ex ante to cover the payment risk that it brings to the system.

But collateral can change in value from the time it is pledged to the time it is realized to close out the funds position of a defaulting participant, as illustrated in Figure 1. Figure 1 shows the value of payment risk in the top panel, and two possible states of the world for the price of the collateral supporting such risk in the bottom panels. Let us refer to the bottom left panel as scenario 1, and the bottom right panel as scenario 2. For scenario 1, collateral value is constant from the time it is immobilized, $t + 1$, to the time the funds are due, $t + 3$. This assumption may be unrealistic when the collateral is composed of equity instruments, which leads us to scenario 2. In scenario 2, we observe that collateral value changes from the time it is immobilized, $t + 1$, to the time the system closes, $t + 3$. Such fluctuation results in uncollateralized payment risk at $t + 3$. One possible solution to avoid uncollateralized payment risk is to require more collateral ex ante. The question is how much more? The answer is an amount sufficient so that payment risk is collateralized (subject to a confidence level) despite the price volatility in the collateral value. One common approach is to discount

Figure 1: PAYMENT RISK AND COLLATERAL

TOP PANEL: Funds required from Broker B as a result of buying 200 Company X securities from Broker A. Payment risk = Funds required - Funds available = 1,000. BOTTOM LEFT PANEL: Collateral value supporting payment risk in scenario 1. BOTTOM RIGHT PANEL: Collateral supporting payment risk in scenario 2, when there is volatility in the price of collateral.

the market value of the collateral so that more collateral is pledged initially to cover payment risk. We refer to this "discount" as a *haircut.*

Definition 2 (Haircut) *A haircut represents the amount that the security used as collateral could decline in value from the time the participant fails to pay the funds it owes to the time the collateral is sold in the market to cover the funds obligation.*

From this example, we observe that the methodology used to calculate the haircut of a security is critical in determining the appropriate value of collateral. This is the focus of our paper. In particular, we focus on the calculation of haircuts for equity instruments, and where the only source of uncertainty in the future price is created by *market risk.*

3. Risk-Cost Frontier

To study different methodologies of calculating haircuts, we first need a framework that allows us to compare each methodology. We do so by comparing the risk-cost trade-off

implied by each. The trade-off arises where a higher haircut implies a higher collateral cost to participants, but a reduction in settlement risk to the system.

This framework enables us to (i) address the robustness (i.e., sufficient collateral in an extreme event) and accuracy of a given technology to measure risk, and thus to value collateral, and (ii) compare and rank different technologies to measure risk.

We focus on the trade-off between the risk of having price fluctuations in collateral value that are not covered by the haircut, which we call *tail risk*, and the cost of pledging collateral, measured by the excess collateral above payment risk that corresponds to the haircut, which we call *collateral cost.*

We construct a frontier by obtaining a sequence of cost-risk pairs of points (XY coordinates). In these pairs, *cost* is represented by the extra collateral value required by the application of the haircut to the market value of the security (or securities), and *risk* is represented by the size of the tail implied by the risk measure under study. This sequence is plotted in an XY plane, where the X coordinates correspond to *collateral cost* and the Y coordinates correspond to *tail risk*.[4]

This methodology can be illustrated with a simple example where the risk-cost frontier is constructed for *one* participant in the securities settlement system that uses *one* equity instrument as collateral. The haircut for such an instrument is calculated using Value-at-Risk as the risk measure.[5] The first step in constructing the risk-cost frontier is to calculate the VaR for different confidence levels. The result of the VaR calculations is represented in a vector, XY, defined as follows:

$$XY = \begin{cases} VaR_{tr_1} & tr_1 \\ \vdots & \vdots \\ VaR_{tr_n} & tr_n \end{cases}$$

where tr_i for $i = 1, 2, ..., n$, represents the size of tail risk (i.e., confidence level) for VaR, and n represents the number of points of the risk-cost frontier that we calculate.[6] The next step

---

[4]Cost can be represented in percentage terms as the *haircut* for the case of the frontier with one security used as collateral, and as the *dollar value* when there is more than one security used as collateral.

[5]A risk measure is a correspondence from a space of random variables (e.g., stock returns) to a scalar (e.g., Value-at-Risk). A common risk measure used to calculate haircuts is Value-at-Risk. In section 4, we provide an overview of two risk measures: Value-at-Risk and Expected Shortfall.

[6]For instance, for VaR with a confidence level of 1 per cent, VaR_1, the associated tail risk, tr, is 1 per

7

is to assign the VaR measures (first column of the XY vector) as the vector of haircuts, h. The vector h represents how the haircut varies as the size of the tail changes:

$$XY = \begin{cases} VaR_{tr_1} & tr_1 \\ \vdots & \vdots \\ VaR_{tr_n} & tr_n \end{cases} = \begin{cases} h_1 & tr_1 \\ \vdots & \vdots \\ h_n & tr_n \end{cases} .$$

The risk-cost frontier is constructed by mapping the first column of the XY vector corresponding to collateral cost (X-coordinate), and the second column of the XY vector corresponding to tail risk (Y-coordinate). This is shown in Figure 2.

VaR confidence level	Tail risk	Haircut	Monetary cost ($)
2 per cent	$tr_1 = 2$	3 per cent	1.5 million
1 per cent	$tr_2 = 1$	5 per cent	2.5 million

Table 2: RISK-COST TRADE-OFF

An example of the risk-cost trade-off for a participant when payment risk is $50 million and the risk measure used is VaR. Moving from the initial tail risk value, tr_1, to a lower value of tail risk, tr_2, costs the participant in the system an additional $1 million or, equivalently, a 2 per cent increase in the haircut.

Figure 2 demonstrates the trade-off between tail risk and collateral cost. For example, consider that $tr_1 = 2$ per cent, which would correspond to a VaR_2; $tr_2 = 1$ per cent, which would correspond to a VaR_1; $h_1 = 3.0$ per cent, and $h_2 = 5.0$ per cent. With these values, we observe that a 1 per cent reduction in tail risk leads to an additional 2 per cent increase in the haircut. In monetary terms, Table 2 indicates that if a payment risk of $50 million is to be collateralized, then a 1 per cent decrease in tail risk would result in an increase in marginal collateral cost of $1 million.

---

cent. Such tail risk implies that 1 per cent of the time the haircut will not cover the price fluctuation in the value of the collateral instrument. For this example, losses in the market price of collateral correspond to the left tail of the return distribution.

Figure 2: RISK-COST FRONTIER IN A SECURITIES SETTLEMENT SYSTEM

This figure illustrates the resulting risk-cost frontier for one participant, with one equity instrument used as collateral. The X coordinates represent collateral cost captured by the haircut; that is, the higher the haircut the more costly it is to use the instrument as collateral. The Y coordinates represent the size of the tail of the distribution that exceeds the VaR measure.

This framework can then be used to compare different risk measures. Consider a security used as collateral that exhibits high expected returns but with occasional large losses: one would expect it to exhibit fat tails on its return distribution. Assume that the returns of this security are realizations from a probability distribution that is unknown to the risk manager who is calculating the haircut. For now, let us assume that the true return distribution is known to the econometrician. This would allow us to compare the closeness of the frontier of a given risk measure used by the risk manager with the true frontier of quantiles of the data. For instance, we could compare the risk-cost frontier associated with two risk measures calculated by the risk manager: (i) VaR assuming a normal distribution to characterize the returns, and (ii) VaR assuming a generalized extreme value distribution to characterize the same returns. Figure 3 illustrates an example of a possible result that may be obtained when conducting such an experiment. In this particular case there is a clear mismeasurement of risk, which results from the assumption of normality. The risk measure that uses an extreme value distribution gives a haircut value that is closer to the quantiles of the data.

9

Figure 3: RISK-COST FRONTIER COMPARISON

This figure illustrates the resulting risk-cost frontier for one participant, with one equity instrument used as collateral, when using VaR with a normal distribution and VaR with an extreme value distribution. The figure shows the degree of risk mismeasurement that may result from using a thin tail distribution when the true distribution has fat tails.

This section has summarized the framework used to evaluate different risk measures. Sections 4 and 5 introduce the key concepts related to risk measurement and the associated estimation methods. In section 6, a case study is presented to illustrate the methodology for calculating haircuts, using the concepts reviewed in the previous sections.

4. Measurement of Risk

What is risk? And how should one go about measuring it? We understand risk as the uncertainty of observing an undesirable state of the world in a future time, and a measure of risk as the *correspondence* between a space of random variables and a scalar value. One aspect of measuring risk requires finding a distribution that is an accurate description of the probabilities associated with all states of the world.

When valuing collateral, the haircut is a reflection of the risk of a change in its market price. The calculation of such a haircut is a critical determinant of the value given to collateral. A haircut is a measure of risk that maps a distribution of returns into a scalar. When calculating haircuts for equity instruments used as collateral, we are interested in selecting a particular probability distribution for the returns of the asset that gives us the best possible estimate of the losses (tail of the unobserved distribution) with a given probability.

Specifically, we are interested in the left tail of the return distribution, because this area represents extreme negative returns. Given our focus on losses, we adopt the convention that a loss is represented by a positive number, and a profit with a negative number. Similarly, a risk measure indicates risk when it is positive, and no risk when it is zero or negative.

The definition of a risk measure is broad in the sense that there could be many *correspondences*/risk measures that may be used. However, only a subset of all correspondences/risk measures are appropriate indicators of risk. This brings us to the concept of *coherence*, which is studied below.

4.1 Axioms of coherent risk measures

Coherence is a term that captures the desired properties of a risk measure. This term is due to Artzner, Delbaen, Eber, and Heath (1997, 1999), who through an axiomatic formulation set the foundations for *coherent* risk measures.

Consider a set of x of real-valued random variables where the function ρ is a real-valued risk measure. We interpret these random variables as the negative returns on an equity instrument that is used as collateral.

The following four conditions are required for a risk measure to be considered coherent:

- *Positive homogeneity.* This property states that having λ times the security is equivalent to scaling the risk coming from the single security by a factor of λ. This is an intuitive property, since the equivalence implies that there should not be any diversification effects from having λ times of the same security. Mathematically, this property can be presented as follows:
$$\rho(\lambda x) = \lambda \rho(x). \tag{1}$$
  In terms of collateral, this property suggests that more collateral of the same type increases risk owing to concentration in one asset.

- *Subadditivity.* This property represents the benefits of diversification of a portfolio; that is, the risk derived from the portfolio $(x + y)$ is lower than (or equal to) the risk derived from the sum of the risk of the individual securities (x,y). Mathematically, this property can be presented as follows:

$$\rho(x + y) \leq \rho(x) + \rho(y). \tag{2}$$

A violation of subadditivity would imply that a collateral portfolio could be separated into smaller portfolios, repledged, and a higher value (lower haircut) obtained than before. If subadditivity does not hold, the risk measure could provide misleading information, which may lead to under-collateralization.

- *Monotonicity.* This property represents the notion that a higher return is associated with a higher risk. Mathematically, this property can be presented as follows:

$$x \leq y \Rightarrow \rho(x) \leq \rho(y). \tag{3}$$

For collateral valuation, this property implies that collateral instruments that have higher returns than others also have higher volatility, and thus receive a higher haircut.

- *Translational invariance.* This property implies that adding n units of the risk-free asset (or cash) with returns r_0 to a random return of a security leads to a decline in risk. Mathematically, this property can be presented as follows:

$$\rho(x + nr_0) = \rho(x) - n. \tag{4}$$

For a portfolio of collateral, this property implies that, in the absence of inflation, adding cash as collateral reduces the risk of negative changes in the value of collateral by the amount of cash introduced in the portfolio.

We next define two risk measures: Value-at-Risk and Expected Shortfall.
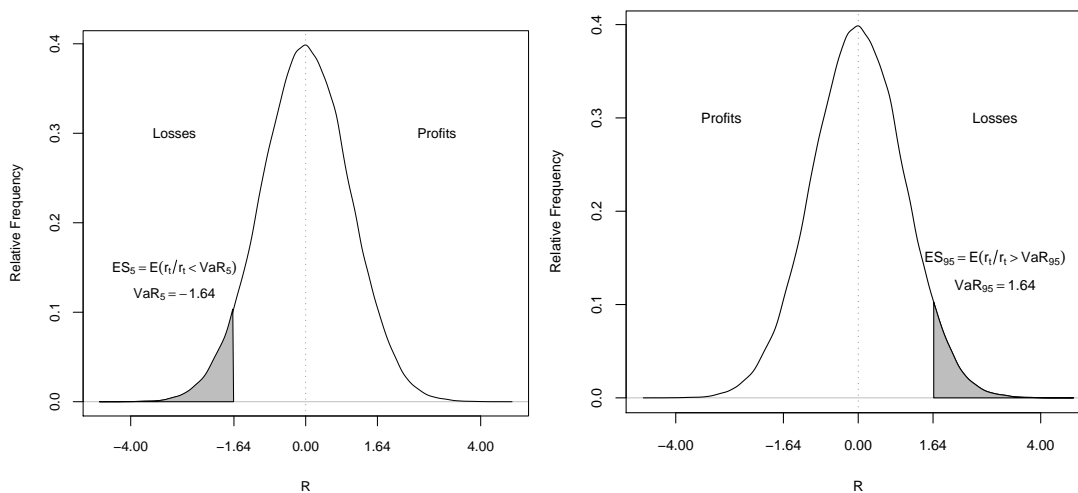
4.2 Different risk measures

Value-at-Risk and Expected Shortfall are two risk measures that are commonly used in finance to determine the value of a loss for an asset, with a given probability.

To put these measures into context, let us start by defining $r_t = \log(p_t/p_{t-1})$ to be the returns at time t, and p_t the price of an asset (or portfolio) at time t. Let the sample of observations be denoted by $r_t, t = 1, 2, \ldots, n$ where n is the sample size and r_t has a distribution function F with mean μ_t and variance σ_t^2.

4.2.1 Value-at-Risk

The VaR is the minimum potential loss in value of a portfolio given the specifications of market conditions, time horizon, and level of statistical confidence. The notation of the VaR

Panel (a):

Relative Frequency

Losses   Profits

$ES_5 = E(r_t / r_t < VaR_5)$

$VaR_5 = -1.64$

R

(a) VaR - Return distribution

Panel (b):

Relative Frequency

Profits   Losses

$ES_{95} = E(r_t / r_t > VaR_{95})$

$VaR_{95} = 1.64$

R

(b) VaR - Negative Return distribution

Figure 4: VAR AND ES RISK MEASURES

These figures illustrate normally distributed return distributions with mean zero and standard deviation of 1. The 5 per cent VaR is equal to -1.64 for the distribution of returns (panel a), and the 95 per cent VaR is equal to 1.64 for the distribution of losses (negative returns, panel b). The corresponding values for ES for both distributions are equal to the average value of returns that are less than -1.64 for panel a, and the average value of the returns that are greater than 1.64 for panel b.

depends on the way the distribution is represented. Generally, for risk management we adopt the convention of representing the distribution of losses or negative returns as shown in panel b of Figure 4. When this is the case, VaR represents a high quantile of the distribution.

VaR's popularity originates from the aggregation of several components of risk at firm and market levels into a single number. The popularity of VaR can be traced back to the seminal work of Markowitz (1952), who noted that one should be interested in risk as well as return and advocated the use of standard deviation as a measure of dispersion. The acceptance and use of VaR has been spreading rapidly since its inception in the early 1990s. Because of VaR's simplicity, computational easiness, and ready applicability, it has become a standard measure used in financial risk management. Many authors have claimed, however, that VaR has several conceptual problems. Artzner, Delbaen, Eber, and Heath (1997, 1999), for example, state the following problems: (i) VaR measures only percentiles of profit-loss distributions, and thus disregards any loss beyond the VaR level ("tail risk"), and (ii) VaR is not coherent since it is not subadditive. Szegö (2005) specifies the conditions under which VaR can be used and recommends other risk measures that are appropriate to investigate tail events. Furthermore, Szegö (2005) highlights that the use of VaR may give incentives to

13

stretch the tail exceeding VaR and thereby reduce VaR. In addition to these limitations, we consider the following:

- VaR may lead to a wide variety of results under a wide variety of assumptions and methods, and the selection of these assumptions and methods is critical to accurately measuring the risk. This limitation is common to other risk measures; however, given the broad use of VaR we consider that such use should be accompanied by a careful exploration of the data to aid in the selection of the estimated return distribution.

- VaR explicitly does not address exposure in extreme market conditions and it may violate coherence in certain settings. In terms of extreme risk, we show in section 5 how an estimation technique that uses extreme value theory can help VaR to better measure the tail of the distribution and thus obtain better estimates for high quantiles.

- As mentioned previously, VaR is not always a coherent risk measure, because it does not satisfy the assumption of subadditivity. For a subadditive measure, portfolio diversification always leads to a reduction in risk, while for risk measures that violate subadditivity the diversification may lead to an increase in the overall portfolio risk. Because VaR asks the question "What is the *minimum loss* incurred in the α per cent worst cases in a portfolio?", it does not satisfy the subadditivity condition of a coherent risk measure under certain conditions. A more robust question is "What is the expected loss incurred in the α per cent worst cases in a portfolio?", which is in line with the definition of Expected Shortfall.

- Another limitation of VaR is that it focuses on a single, somewhat arbitrary point. An alternative to selecting an arbitrary quantile is to use another risk measure that provides more information on the tail, such as expected shortfall, which calculates the average loss after the VaR quantile.

There are several methods for VaR calculations, such as the variance-covariance approach with normal distribution, this approach with Student's t distribution, historical simulation, and the generalized Pareto distribution (GPD) approach. The GPD approach involves extreme value methods of estimation. A brief review of some of these methods is presented in section 5.

4.2.2 Expected Shortfall

The Expected Shortfall (ES) of an asset or a portfolio is the average loss given that VaR has been exceeded. For example, the α per cent ES is the conditional mean of r_t given that $r_t > \text{VaR}_t(\alpha)$:

$$ES_t(\alpha) = E[r_t | r_t > \text{VaR}_t(\alpha)]. \tag{5}$$

Although ES is a coherent measure, it is subject to a similar limitation as VaR, in that it would underestimate the tails if the underlying distribution has thicker tails than the assumed return distribution that was used to calculate VaR. In such a setting, it is more desirable to use ES with extreme value methods to get a correct measure of the tail and the corresponding risk.

Similar to VaR, there are several methods for ES calculations, such as the variance-covariance approach, historical simulation, and the extreme value methods of estimation.

5. Methods of Estimation

There are various ways to model the return distribution of an asset. Generally, the models can be classified as those that use either a *parametric* or a *non-parametric* approach. The parametric approach uses a given distribution to model the return distribution, whereas the non-parametric approach uses historical data directly and does not make any distributional assumptions. Two examples of parametric approaches are those that use the normal distribution and those that use distributions based on extreme value theory to characterize the tail of the return distribution. After presenting these two methods, we briefly summarize a non-parametric approach called historical simulation.

5.1 Parametric approach: Normal distribution

This method uses a normal distribution to represent the sample return distribution, $r_t \sim N(\mu_t, \sigma_t^2)$. When this is the case, the calculation of $\text{VaR}_t(\alpha)$ reduces to

$$\text{VaR}_t(\alpha) = \mu_t + \sigma_t \cdot q(\alpha), \tag{6}$$

where $q(\alpha)$ is the α-quantile of the standard normal distribution.

If $r_t \sim N(\mu_t, \sigma_t^2)$, then $ES_t(\alpha)$ may be computed as the mean of a truncated normal random variable:

$$ES_t(\alpha) = \mu_t + \sigma_t \frac{\Phi}{1 - \Phi},$$

where Φ is the normal cumulative distribution function. The main benefit of using a normal distribution is the simplicity of the calculation of the risk measures. The main drawback, specifically when calculating haircuts for equity instruments, is that the normal distribution is not a very accurate representation of the true distribution of returns for equity instruments. In particular, fat tails observed for the returns of equity instruments are not captured by the normal distribution. To address this limitation, we require a different distributional assumption that better approximates the sample tail behaviour. One approach is to use EVT methods to select such distribution.

An enhancement to the use of a normal distribution consists of determining better estimates for the variance. This is important because the sample variance as an estimator of the standard deviation, although simple, has drawbacks at high quantiles of a fat-tailed empirical distribution. The quantile estimates for the right tail (left tail) are biased downwards (upwards) for high quantiles of a fat-tailed empirical distribution. Therefore, the risk is underestimated with a normality assumption. Another drawback of normality is that it is not appropriate for asymmetric distributions. Despite these drawbacks, this approach is commonly used for calculating the VaR from holding a certain portfolio, since the VaR is additive when it is based on sample variance under the normality assumption. Instead of the sample variance, the standard deviation can be estimated by a statistical model. Since financial time series exhibit volatility clustering, the ARCH (Engle 1982) and GARCH (Bollerslev 1986) are popular models for volatility modelling.[7] Although the conditional distribution of the GARCH process has normal tails, the unconditional distribution has some excess kurtosis. However, this may not be sufficient for modelling fat-tailed distributions, since the tails of the unconditional distribution decay exponentially fast. In these cases, the GARCH-t (GARCH with student-t innovations) model may be an alternative. A weakness of the GARCH models is that they generally produce highly volatile quantile estimates (see Gençay, Selçuk, and Ulugülyağcı 2003). Excessive volatility of quantile estimates is not desirable in risk manage-

---

[7]ARCH and GARCH refer to autoregressive conditional heteroscedasticity and generalized autoregressive conditional heteroscedasticity, respectively.

ment, since it is costly to adjust the required capital frequently and is difficult to regulate.

5.2 Parametric approach: Extreme value theory

The second method to model the return distribution is based on extreme value theory (EVT). EVT is a powerful and fairly robust framework in which to study the tail behaviour of a distribution.[8]

5.2.1 Fundamental concepts: EVT and extreme risk

EVT can be thought of as a theory that provides methods for modelling extremal events. Extremal events are those realizations of risk that take values from the tail of the probability distribution. EVT provides the tools to estimate a distribution of the tails through statistical analysis of the empirical data.

Within the EVT context, there are two approaches to model the extremal events (Figure 5). One of them is the direct modelling of the distribution of minimum or maximum realizations, *block maxima models*. The other one is modelling the exceedances of a particular threshold, *peak-over-threshold models*. To identify extremes, one approach considers dividing sample data in blocks. The maxima in these blocks are considered extreme events in the sample data. This approach is followed in block maxima models. Another way to identify extremes in the sample data consists of selecting the observations that exceed a given high threshold. This approach is followed in peak-over-threshold models.

We concentrate on the latter models since we consider them more useful for applications when the data on extreme events are rather limited. Before studying these models, we discuss the central result of extreme value theory: the Fisher-Tippett theorem.

5.2.2 Fundamental concepts: Distribution of the maxima and the Fisher-Tippett theorem

The normal distribution is the important limiting distribution for sample averages as summarized in a central limit theorem. Similarly, the family of extreme value distributions is used to study the limiting distributions of the sample extrema. This family can be presented

---

[8]Embrechts, Kluppelberg, and Mikosch (1997) is a comprehensive source of theory and applications of the extreme value theory to the finance and insurance literature. Recent applications of EVT can be found in Gençay, Selçuk, and Ulugülyağcı (2002), Gençay, Selçuk, and Ulugülyağcı (2003), and Gençay and Selçuk (2006).

Figure 5: APPROACHES TO MODELLING EXTREMAL EVENTS

On the left is the block maxima approach, and on the right the peak-over-threshold approach. Gilli and Këllezi (2005) motivate these concepts using a similar figure.

under a single parameterization known as the generalized extreme value (GEV) distribution. The theory deals with the convergence of maxima, that is, the limit law for the maxima. To illustrate this, consider $r_t, t = 1, 2, \ldots, n$, an uncorrelated sample of returns with a common distribution function $F(x) = \Pr\{r_t \le x\}$, which has mean (location parameter) μ and variance (scale parameter) σ^2.[9] Denote the sample maxima[10] of r_t by $M_1 = r_1$, $M_2 = max(r_1, r_2)$, and, in general, $M_n = max(r_1, ..., r_n)$, where $n \ge 2$, and let $\Re$ denote the real line. If there exists a sequence $c_n > 0$, $d_n \in \Re$ and some non-degenerate distribution function H such that

$$\frac{(M_n - d_n)}{c_n} \xrightarrow{d} H,$$

then H belongs to one of the following three families of distributions:

Gumbel:
$$\Lambda(x) = e^{-e^{-x}}, \quad x \in \Re,$$

Fréchet:
$$\Phi_\alpha(x) = \begin{cases} 0, & x \le 0 \\ e^{-x^{-\alpha}}, & x > 0 \quad \alpha > 0, \end{cases}$$

---

[9]For convenience, we will assume that $\mu = 0$ and $\sigma^2 = 1$ in this section.
[10]The sample maxima is $min(r_1, ..., r_n) = -max(-r_1, ..., -r_n)$.

18

Weibull:

$$\Psi_\alpha(x) = \begin{cases} e^{-(-x^\alpha)}, & x \le 0 \quad \alpha < 0 \\ 1, & x > 0. \end{cases}$$

The Fisher and Tippett (1928) theorem[11] suggests that the asymptotic distribution of the maxima belongs to one of the three distributions above,[12] regardless of the original distribution of the observed data.[13]

By taking the reparameterization $\xi = 1/\alpha$, due to von Mises (1936) and Jenkinson (1955), Fréchet, Weibull and Gumbel distributions can be represented in a unified model with a single parameter. This representation is known as the generalized extreme value distribution (GEV):

$$H_\xi(x) = \begin{cases} e^{-(1+\xi x)^{-\frac{1}{\xi}}} & \text{if } \xi \ne 0, \ 1+\xi x > 0 \\ e^{-e^{-x}} & \text{if } \xi = 0, \end{cases}$$

where $\xi = 1/\alpha$ is a shape parameter and α is the tail index.

The class of distributions of $F(x)$ where the Fisher-Tippett theorem holds is quite large.[14] One of the conditions is that $F(x)$ has to be in the domain of attraction for the Fréchet distribution[15] ($\xi > 0$), which in general holds for the financial time series. Gnedenko (1943) shows that if the tail of $F(x)$ decays like a power function, then it is in the domain of attraction for the Fréchet distribution. The class of distributions whose tails decay like a power function is large and includes the Pareto, Cauchy, Student-t, and mixture distributions. These distributions are the well-known heavy-tailed distributions.

---

[11]The first formal proof of the Fisher-Tippett theorem is given in Gnedenko (1943).

[12]In conventional statistics, a Weibull distribution function $F_\alpha(x)$ is defined as $F_\alpha(x) = 1 - e^{-x^\alpha}$ for $x > 0$. The Weibull distribution function $\Psi_\alpha(x)$ above is concentrated on $(-\infty, 0)$ and it is $\Psi_\alpha(x) = 1 - F_\alpha(-x)$ for $x < 0$. $F_\alpha(x)$ and $\Psi_\alpha(x)$ have completely different extremal behaviour. In the extreme value theory literature, $\Psi_\alpha(x)$ is referred to as the Weibull distribution. See Embrechts, Kluppelberg, and Mikosch (1997, Ch. 3).

[13]The interested reader will find the full development of the theory in Leadbetter, Lindgren, and Rootzén (1983) and de Haan (1990).

[14]McNeil (1997, 1999), Embrechts, Kluppelberg, and Mikosch (1997), Embrechts, Resnick, and Samorodnitsky (1998) and Embrechts (1999) have excellent discussions of the theory behind the extreme value distributions from the risk-management perspective.

[15]See Falk, Hüssler, and Reiss (1994).

5.2.3 Fundamental concepts: Distribution of exceedances over a threshold

In general, we are not only interested in the maxima of observations, but also in the behaviour of large observations that exceed a high threshold. One method of extracting extremes from a sample of observations, r_t, $t = 1, 2, \ldots, n$ with a distribution function $F(x) = \Pr\{r_t \le x\}$ is to take the exceedances over a predetermined, high-threshold u (Figure 6). Exceedances of a threshold u occur when $r_t > u$ for any t in $t = 1, 2, \ldots, n$. An excess over u is defined by $y = r_i - u$.[16]

Figure 6: DISTRIBUTION FUNCTION AND DISTRIBUTION OVER THRESHOLD

The distribution function of the returns is shown on the left panel, and the distribution function for the exceedances over the threshold u is shown on the right panel. Gilli and Këllezi (2005) motivate these concepts using a similar figure.

Given a high threshold u, the probability distribution of excess values of r over threshold u is defined by

$$F_u(y) = \Pr\{r - u \le y | r > u\}, \tag{7}$$

which represents the probability that the value of r exceeds the threshold u by at most an amount y given that r exceeds the threshold u. This conditional probability may be written as

$$F_u(y) = \frac{\Pr\{r - u \le y, r > u\}}{\Pr(r > u)} = \frac{F(y + u) - F(u)}{1 - F(u)}. \tag{8}$$

Since $x = y + u$ for $r > u$, we have the following representation:

$$F(x) = [1 - F(u)] F_u(y) + F(u). \tag{9}$$

Notice that this representation is valid only for $r > u$.

A theorem by Balkema and de Haan (1974) and Pickands (1975) shows that for sufficiently high threshold u, the distribution function of the excess may be approximated by

---

[16]This is also referred to as peaks-over-threshold (POT).

the generalized Pareto distribution (GPD), because as the threshold gets large, the excess distribution $F_u(y)$ converges to the GPD. The GPD in general is defined as

$$G_{\xi,\sigma,v}(x) = \begin{cases} 1 - \left(1 + \xi \frac{x-v}{\sigma}\right)^{-1/\xi} & \text{if } \xi \neq 0 \\ 1 - e^{-(x-v)/\sigma} & \text{if } \xi = 0, \end{cases} \tag{10}$$

with

$$x \in \begin{cases} [v, \infty], & \text{if } \xi \geq 0 \\ [v, v - \frac{\sigma}{\xi}], & \text{if } \xi < 0, \end{cases}$$

where $\xi = 1/\alpha$ is the shape parameter, α is the tail index, σ is the scale parameter, and v is the location parameter. When $v = 0$ and $\sigma = 1$, the representation is known as the standard GPD. There is a simple relationship between the standard GPD $G_\xi(x)$ and $H_\xi(x)$ such that $G_\xi(x) = 1 + \log H_\xi(x)$ if $\log H_\xi(x) > -1$.

The GPD embeds a number of other distributions. When $\xi > 0$, it takes the form of the ordinary *Pareto distribution*. This particular case is the most relevant for financial time-series analysis, since it is a *heavy-tailed* one. For $\xi > 0$, $E[X^k]$ is infinite for $k \geq 1/\xi$. For instance, the GPD has an infinite variance for $\xi = 0.5$ and, when $\xi = 0.25$, it has an infinite fourth moment. For the security returns or high-frequency foreign exchange returns, the estimates of ξ are usually less than 0.5, implying that the returns have finite variance Dacorogna, Gençay, Müller, Olsen, and Pictet (2001). When $\xi = 0$, the GPD corresponds to the *thin-tailed* distributions, and it corresponds to *finite-tailed* distributions for $\xi < 0$.

The importance of the Balkema and de Haan (1974) and Pickands (1975) results is that the distribution of excesses may be approximated by the GPD by choosing ξ and setting a high threshold u. The GPD model can be estimated with the maximum-likelihood method. For $\xi > -0.5$, Hosking and Wallis (1987) present evidence that maximum-likelihood regularity conditions are fulfilled and the maximum-likelihood estimates are asymptotically normally distributed. Therefore, the approximate standard errors for the estimator of ξ can be obtained through maximum-likelihood estimation.

For the tail estimation, recall from equation (9) that

$$F(x) = [1 - F(u)] F_u(y) + F(u).$$

Since $F_u(y)$ converges to the GPD for sufficiently large u, and since $x = y + u$ for $r > u$, we have

$$F(x) = [1 - F(u)]\, G_{\xi,\sigma,u}(x - u) + F(u). \tag{11}$$

After determining a high-threshold u, the last term on the right-hand side can be determined by $(n - n_u)/n$, where n_u is the number of exceedances and n is the sample size. As a result, we have the following estimator:

$$
\begin{aligned}
\hat{F}(x) &= \left(1 - \tfrac{n-n_u}{n}\right) G_{\hat{\xi},\hat{\sigma},u}(x - u) + \tfrac{n-n_u}{n} \\
&= \tfrac{n_u}{n} G_{\hat{\xi},\hat{\sigma},u}(x - u) + \tfrac{n-n_u}{n} \\
&= 1 + \tfrac{n_u}{n}\left(G_{\hat{\xi},\hat{\sigma},u}(x - u) - 1 \right).
\end{aligned}
$$

Therefore, the tail estimator becomes

$$\hat{F}(x) = 1 - \frac{n_u}{n}\left(1 + \hat{\xi}\frac{x - u}{\hat{\sigma}}\right)^{-1/\hat{\xi}} \quad \text{given that} \quad G_{\xi,\sigma,u}(x) = 1 - \left(1 + \xi\frac{x - u}{\sigma}\right)^{-1/\xi} \tag{12}$$

where $\hat{\xi}$ and $\hat{\sigma}$ are the maximum-likelihood estimators. Notice that the estimator in equation (12) is valid only for $r > u$.

5.2.4 Estimators: Value-at-Risk with EVT

For a given probability $q > F(u)$, an estimate of the VaR may be calculated by inverting the tail estimate in equation (12) to obtain[17]

$$\text{VaR}_t(\alpha) = u + \frac{\hat{\sigma}}{\hat{\xi}}\left[\left(\frac{n}{n_u}\alpha\right)^{-\hat{\xi}} - 1\right], \tag{13}$$

where u is a threshold, $\hat{\sigma}$ is the estimated scale parameter, $\hat{\xi}$ is the estimated shape parameter, n is the sample size, n_u is the number of exceedances, and $\alpha = 1 - q$.[18]

---

[17] Also, see Embrechts, Kluppelberg, and Mikosch (1997, p. 354) and McNeil (1999).

[18] As an example, suppose that in daily stock returns the threshold is determined as 6 per cent and estimated parameters are $\hat{\sigma} = 0.05$ and $\hat{\xi} = 0.50$. Further, suppose that $n = 1,000$ and $n_u = 50$. The VaR at 1 per cent is

$$
\begin{aligned}
\text{VaR}_t(0.01) &= 0.06 + \frac{0.05}{0.50}\left[\left(\frac{1000}{50}0.01\right)^{-0.50} - 1\right] \\
&= 0.184.
\end{aligned}
$$

That is, the stock return will not exceed 18.4 per cent in one day 99 per cent of the time.

5.2.5 Estimators: Expected shortfall with EVT

The expected shortfall (ES) is related to Value-at-Risk through

$$ES_t(\alpha) = \text{VaR}_t(\alpha) + E[r_t - \text{VaR}_t(\alpha)|r_t > \text{VaR}_t(\alpha)], \tag{14}$$

where the second term in equation (14) is simply the mean excess distribution $F_{VaR_t(\alpha)}(y)$ over the threshold $\text{VaR}_t(\alpha)$. The GPD approximation to $F_{VaR_t(\alpha)}(y)$, yields

$$E[r_t - \text{VaR}_t(\alpha)|r_t > \text{VaR}_t(\alpha)] = \frac{\sigma_t + \xi(\text{VaR}_t(\alpha) - u)}{1 - \xi} \tag{15}$$

provided that $\xi < 1$. Equation (13) together with equation (15) yields

$$ES_t(\alpha) = \frac{\text{VaR}_t(\alpha)}{1 - \hat{\xi}} + \frac{\hat{\sigma}_t - \hat{\xi}u}{1 - \hat{\xi}}. \tag{16}$$

ES is simple to calculate, and allows an easy switch from any risk-management system that works with VaR.

The use of ES with extreme value methods allows the risk measure to be coherent, as well as to capture the tails of the distribution more accurately. However, the implementation of this risk measure when calculating the haircut of an asset not only depends on the better measurement of risk but also on the collateral costs associated with it. In particular, it depends on how the risk and cost attributes compare with other risk measures. An example of this analysis is presented in section 6.

5.3 Non-parametric approach: Historical simulation

For completeness, we briefly present a non-parametric method to model the return distribution called historical simulation. This method estimates the quantiles of an underlying realized distribution. The problem with this approach is that the empirical distribution function is not one-to-one but constant between two realizations. That is, we may not have observations corresponding to certain quantiles of the underlying distribution. A simple solution may be rounding the probability level to the nearest empirical probability and then

taking the corresponding quantile as the desired quantile estimate. A more appropriate so-
lution is to smooth the empirical distribution function with piecewise linear interpolation or
kernel interpolation so that it is one-to-one.

The historical simulation method may fit the sample well, around the moderate quantiles,
since no parametric form for the distribution is assumed. The disadvantage of this method
is that the high quantile estimates are not reliable, since they are calculated from only a
few observations. Furthermore, it is not possible to obtain any quantile estimates above the
highest observed quantile.

6. Haircuts and Risk Measures

In this section we use the risk-cost frontier to select among different methodologies to calculate
haircuts. We do this for two cases, one based on simulated data, and the other based on
market price data.

- *Simulated data:* We simulate 10,000 observations from a Student-t distribution with
  2.2 degrees of freedom that represent daily returns of an asset used as collateral. We
  select this specification because it represents fat tails observed in financial return series.
  Empirical research that documents such stylized facts started during the 1960s. For
  example, Mandelbrot (1963) emphasizes that financial returns exhibit fat tails and
  proposes using the stable class of distributions. Commonly used distributions to model
  fat tails are the Student-t and Pareto distribution.

- *Market data:* We use the closing price of an equity instrument listed in the Toronto
  Stock Exchange to calculate daily returns. We do not pretend to assess the quality of
  this instrument as an investment and therefore we do not mention its name. We use it
  only as an example to illustrate the methodology proposed in this paper.

Before exploring these cases, let us introduce some context where these assets may be
used as collateral.

6.1 Context

Consider an operator/risk manager of a securities settlement system that accepts equity
instruments as collateral to provide a guarantee that all trades approved for settlement will

be completed by the end of the day. The risk faced by the risk manager is the future loss (negative returns) in the value of collateral, in particular when such collateral is required to close the trades of a defaulting participant. To mitigate such risk, collateral is discounted by a haircut that should reflect the future price volatility (market risk) for the time that collateral is used to support the corresponding payment risk. The risk manager is asked to take only *one* asset as collateral. In general, we refer to this asset as security A, and its continuously compounded returns as $r_{a,t}$. To determine the haircut, the operator starts by considering the returns of A as a random variable generated from an unknown distribution function, F_a. To calculate an adequate haircut, the risk manager needs to develop a model that accounts for the observed variations in the price of collateral. This is done by selecting a probability distribution, $\hat{F}_a$, that provides the best possible estimate of the tail area of F_a, and thus provides a good fit to the simulated data (or market data, depending on the case). Recall that F_a is not observed by the risk manager. Once $\hat{F}_a$ is determined, the risk manager can measure the degree of market risk that needs to be covered, given the $\hat{F}_a$ distribution, and the desired confidence level and holding period. Such risk is summarized in a single number known as the risk measure, which then is assigned as the value of the haircut for the collateral instrument.

6.2 Haircut calculation methodology

It is useful for the risk manager to conduct an exploration of the data, that is, to conduct a preliminary analysis of the historical time series of returns $r_{a,t}$ to determine the type of tail (heavy, thin, or finite) that corresponds to the negative returns. This analysis will give an indication of whether EVT methods should be used to estimate $\hat{F}_a$. The exploration of the sample data consists of using several tools such as *qq* plots, the mean excess function, and the Hill estimator. These techniques are explained next.

6.2.1 Preliminary data analysis: qq plots

In the extreme value theory and applications, the *qq* plot (quantile-quantile plot) is typically plotted against the exponential distribution (a distribution with a thin-sized tail) to measure the fat-tailness of a distribution. If the data are from an exponential distribution, the points on the graph would lie along a positively sloped straight line. If there is a concave presence, this would indicate a fat-tailed distribution, whereas a convex departure is an indication of a short-tailed distribution.[19]

---

[19]If the sample is a realization from a distribution that has the same form as the reference distribution but with different scale and/or location parameters, the *qq* plot is still linear. In this case, the intercept of the

6.2.2 Preliminary data analysis: Mean excess function

A second tool is the sample mean excess function (MEF), $e(u) = E[r - u | r > u]$, which can be estimated by

$$e_n(u) = \frac{\sum_{i=1}^{n}(r_i - u)}{\sum_{i=1}^{n} I_{\{r_i > u\}}} \tag{17}$$

where I is an indicator function. The MEF is the sum of the excesses over the threshold u divided by the number of data points that exceed the threshold u. An estimate of the mean excess function describes the expected overshoot of a threshold once an exceedance occurs. If the empirical MEF is a positively sloped straight line above a certain threshold u, this indicates that the data follows the GPD with a positive shape parameter, ξ. On the other hand, exponentially distributed data would show a horizontal MEF, while short-tailed data would have a negatively sloped line.

6.2.3 Preliminary data analysis: Hill plot

Another tool in threshold determination is the Hill plot.[20] Hill (1975) proposes an estimator of ξ when $\xi > 0$ (Fréchet case). By ordering the data with respect to their values as $r_{1,n}, r_{2,n}, r_{3,n}, \ldots, r_{n,n}$ where $r_{1,n} \geq r_{2,n} \geq r_{1,n} \geq \ldots \geq r_{n,n}$, the Hill estimator of the tail index ξ is

$$\hat{\xi} = \frac{1}{k-1} \sum_{i=1}^{k-1} \ln r_{i,n} - \ln r_{k,n} \text{ for } k \geq 2 \tag{18}$$

where k is the upper-order statistics (the number of exceedances),[21] n is the sample size, and $\alpha = 1/\xi$ is the tail index. A Hill plot is constructed such that estimated ξ is plotted as a function of k upper-order statistics or the threshold. A threshold is selected from the plot where the shape parameter ξ is fairly stable. The Hill estimator is proven to be a consistent estimator of $\xi = 1/\alpha$ for fat-tailed distributions.[22]

A difficulty of the Hill estimator is the ambiguity of the value of threshold parameter, k. In threshold determination, we face a trade-off between bias and variance. If we choose a low threshold, the number of observations (exceedances) increases and the estimation becomes

---

qq plot indicates the location parameter while the scale parameter determines the slope.

[20]See Embrechts, Kluppelberg, and Mikosch (1997) for a detailed discussion and several examples of the Hill plot.

[21]The ith element from the ordered sample, $r_{i,n}$ is called the ith upper-order statistic.

[22]The conditions on k and n for weak consistency of the Hill estimator are given in Mason (1982) and Rootzén, Leadbetter, and de Haan (1992). Deheuvels, Hausler, and Mason (1988) investigate the conditions for the strong consistency of the Hill estimator. From Hall (1982) and Goldie and Smith (1987), it follows that $(\hat{\xi} - \xi)k^{1/2}$ is asymptotically normally distributed with zero mean and variance ξ^2.

more precise. However, choosing a low threshold also introduces some observations from the centre of the distribution and the estimation becomes biased. While the estimates of ξ based on a few largest observations are highly sensitive to the number of observations used, the estimates based on many elements from the top of the ordering are biased.[23] Therefore, a careful combination of several techniques, such as the qq plot, the Hill plot, and the MEF should be considered in threshold determination.

6.2.4 Preliminary data analysis: Choosing a model for the tail

This section illustrates the exploration of the data for two cases. The first is where the operator uses the asset represented by the simulated $t(2.2)$ data. The second is where the operator uses market data as collateral.

Analysis for the simulated data: Figure 7 shows the properties of the data. The simulated $t(2.2)$ data exhibit heavier tails when compared with a normal distribution. The following points corroborate this finding:

- *Panel a* shows several large negative returns that exceed three standard deviations away from the mean, the most negative return corresponding to -78.78 per cent.

- *Panel b* compares the left tail of the simulated data with that of data from a normal distribution. We observe that the simulated $t(2.2)$ data have larger losses at lower quantiles than the normal distributed data.

- *Panel c* shows a quantile-quantile plot (qq plot) for the normal distributed returns and for the simulated $t(2.2)$ returns. The plot confirms that $r_{a,t}$ has a heavier tail in its negative returns when compared with a normal distribution with the same mean and standard deviation.

- *Panel d* corroborates the presence of heavy tails.

Analysis for the market data: In the same manner as for the simulated $t(2.2)$ data, Figure 8 also shows that returns calculated for the market data have heavier tails when compared with returns that follow a normal distribution. The main difference in the results is that the market data have negative returns that are not as far into the tail as the simulated data. The most negative return corresponds to -43.78 per cent.

---

[23]Danielsson, de Haan, Pend, and de Vries (2001) propose a standardized procedure for choosing an optimal k value.

(a) Returns

(b) Densities

(c) *qq* plot

(d) Mean excess plot

Figure 7: SIMULATED DATA

This figure illustrates: a) simulated $t(2.2)$ returns, b) comparison between normal and simulated $t(2.2)$ returns, c) *qq* plot for normal and simulated $t(2.2)$ returns, and d) mean excess function for simulated $t(2.2)$ returns.

(a) Returns

(b) Densities

(c) *qq* plot

(d) Mean excess plot

Figure 8: MARKET DATA

This figure illustrates: a) market data returns, b) comparison between normal and market data returns, c) *qq* plot for normal and market data returns, and d) mean excess function for the market data returns.

For both cases, the exploration of the data shows that $\hat{F}_a$ should be chosen so that it exhibits fat tails, and thus should point the risk manager to use the EVT methods. With this in mind, we proceed to select $\hat{F}_a$ using the peaks-over-threshold method. We use a threshold of 5, which comes from a stable region in the mean excess function. We also base the selection of the threshold on the behaviour of the risk-cost frontier. Ideally, we would like a threshold that yields a model with an associated risk-cost frontier that is close to a frontier constructed from the quantiles of the data. This is desirable since the closer the two frontiers are, the closer is the haircut derived from the model (GPD) to the haircut calculated from the data.

Another consideration in the selection of the threshold is the stability of the risk-cost frontier. The stability of the threshold is shown in Figure 9. Here the risk-cost frontier is plotted for VaR when different thresholds are used in the estimation. The frontier allows us to compare the robustness of the risk measure's estimation to different thresholds: namely, whether for different thresholds the risk-cost frontier changes dramatically, either by getting closer or further from the frontier obtained from the quantiles of the data. Figure 9 shows that at the 1 per cent tail risk, the frontier is robust to changes in the threshold. However, at the 5 per cent tail risk, changes in the threshold lead to significant changes in the corresponding haircut. We focus on the 1 per cent tail risk, since the objective is to calculate a haircut for an extremal event.

With a threshold of 5, the losses for the simulated $t(2.2)$ data are reduced from 4,994 losses to 168 threshold exceedances. On the basis of these data, the shape parameter, ξ, and the location parameter, β, are estimated to be 0.42 and 2.87, where the value of ξ shows the heavy tailedness of the data. Figure 10(a) shows the empirical excess distribution (points) and the estimated model (smooth curve), and Figure 10(b) shows the empirical tail distribution (points) and the estimated model (smooth curve). From Figure 10 we conclude that the GPD model fits well the empirical observations, and thus possibly F_a. With such a model for $\hat{F}_a$ we can estimate VaR and ES with EVT methods for the desired confidence level. Our analysis using EVT yields a VaR_1 of 9.55 per cent, and a corresponding ES_1 of 17.83 per cent.

For the case of market data, the same analysis was conducted and yields the following results. The shape parameter, ξ, and the location parameter, β, are estimated to be 0.42 and 0.019, where, here too, the value of ξ shows the heavy tailedness of the data. The dataset of losses is reduced from 2,402 losses to 155 when a threshold of 0.05 (5 per cent) is used to estimate $\hat{F}_a$. Our analysis here using EVT yields a VaR_1 of 0.11 (11 per cent) and a

VaR–EVT and Quantiles

Figure 9: SIMULATED $t(2.2)$ RETURNS - CHOOSING A THRESHOLD

This figure shows how different thresholds change in the risk-cost frontier for VaR when it is calculated using the peak-over-threshold method.

(a) Modelling the excess distribution


(b) Modelling the tail of the distribution

Figure 10: SIMULATED $t(2.2)$ RETURNS - EXCESS AND TAIL DISTRIBUTION

This figure compares the estimated models (represented by a smooth line) for the excess distribution and the tail of the distribution, with the actual observations. The figure shows that the model is a good fit to the data.

(a) Modelling the excess distribution

(b) Modelling the tail of the distribution

Figure 11: MARKET DATA RETURNS - EXCESS AND TAIL DISTRIBUTION

This figure compares the estimated models (represented by a smooth line) for the excess distribution and the tail of the distribution, with the actual observations. The figure shows that the model is a good fit to the data.

corresponding ES_1 of 0.18 (18 per cent). Figure 11 shows that the GPD model also fits the data well.

This section has shown how the exploration of the data for different datasets was useful to determine an appropriate $\hat{F}_a$.

6.3 Study of risk and cost attributes for different risk measures

The exploration of the data led to the conclusion that security A (for both datasets) exhibits a fat tail in its negative returns. This observation led to the use of EVT to estimate an appropriate distribution for the tail, and thus the corresponding haircut. In the absence of this analysis, a thin tail distribution such as the normal distribution may have been used in the calculation of the risk measures. In this section we examine the implications of selecting a distribution $\hat{F}_a$ that does not fit the sample data well (e.g., a thin-tailed distribution).

In Table 3 we display for the simulated $t(2.2)$ the different values obtained case for the risk measures and the quantiles of the data for different values of tail risk. The main points that can be taken from Table 3 are listed below:

Tail risk	VaR EVT	ES EVT	VaR Normal	Quantile
0.01	77.30	134.82	10.87	102.58
0.1	28.19	50.01	9.03	19.56
1	9.55	17.83	6.80	6.20
2	6.67	12.86	6.00	4.45
3	5.34	10.55	5.50	3.66
5	3.95	8.15	4.81	2.88
10	2.49	5.63	3.75	1.90

Table 3: SIMULATED $t(2.2)$ RETURNS - RISK MEASURES

This table shows the resulting risk measures at the 0.01, 0.1, 1, 2, 3, 5, and 10 per cent tail risk levels for the returns of security A, when the risk measures are calculated with a GPD or a normal distribution. Also included is the corresponding value associated with the quantile calculation from the data. The quantile value serves as the simulated $t(2.2)$ returns measurement. That is, the more accurate is the measurement of the risk measures, the closer it should be to the quantile values.

- VaR with EVT reports haircut values that are close to the quantiles of the data. However, VaR with EVT overestimates the haircuts when compared with the haircuts obtained from the quantiles of the data.

- VaR with normality further overestimates the haircuts for values greater than the 3 per cent tail risk, when compared with the haircuts corresponding to both the quantiles and VaR with EVT. From 1 to 3 per cent tail risk, VaR with normality provides smaller haircuts than VaR with EVT, although still greater than those implied by the quantiles of the data. For values lower than 1 per cent tail risk, VaR with normality provides lower haircuts than those of VaR with EVT *and* those corresponding to the quantiles of the data.

- As noted before, a critical point to determine the appropriate risk measure to use is the 1 per cent tail risk. For values lower than this critical value, VaR with EVT provides desirable haircuts (greater than or equal to the quantile), whereas VaR with normality provides unacceptable ones (lower than the quantiles).

Table 3 gives us an indication that the thin-tailed assumption, when used to calculate VaR for low tail risk levels (below 1 per cent), may result in *under collateralization*. We can confirm the points taken from Table 3 by comparing the risk-cost frontiers of the different risk measures. This is shown in Figure 12. The following points can be made for each panel in Figure 12.

(a) Panel 1

(b) Panel 2

(c) Panel 3

(d) Panel 4

Figure 12: SIMULATED $t(2.2)$ RETURNS ANALYSIS OF THE RISK-COST TRADE-OFF

Panel 1 shows the frontier constructed from the quantiles of the data, and the frontier resulting from VaR estimated with EVT methods. Panel 2 shows the frontier constructed from the quantiles of the data, and the frontier resulting from VaR estimated with a normal distribution. Panels 3 and 4 compare the frontiers for three models, VaR with normality, VaR with EVT methods, and ES with EVT methods.

Panel 1: We observe that there is a good fit (in terms of the slope) of the frontier of VaR with EVT and the frontier of the quantiles of the data. Nevertheless, VaR with EVT gives greater values for haircuts compared with the quantiles. The greater haircuts may be costly to participants, although they provide a cushion for extremal events.

Panel 2: VaR with normality *overestimates* or *underestimates* the values for the haircuts, depending on different locations in the curve. For the purpose of covering extremal risk, VaR with normality may not be adequate.

Panel 3: We observe that VaR with normality is preferable than VaR with EVT only for tail risk values that are between 3 to 1 per cent.

Panel 4: We observe that ES with EVT may be too costly, since the haircuts derived for all tail risks greatly exceed the haircuts derived from the quantiles of the data.

We also conducted backtesting for the haircuts obtained for VaR with the two parametric methods employed. Backtesting requires calculation of the number of times the returns exceed the haircut, and then comparison of this number with the number expected by the size of tail risk. For example, the haircut corresponding to VaR_1 implies that 1 per cent of the returns should be greater than the haircut (i.e. VaR_1). We observe in Table 4 that for VaR_1 using a normal distribution, the returns exceeded the haircut 0.96 per cent of the time, and when using the GPD only 0.46 per cent of the time, when we expected 1 per cent of the observations to be greater than the haircut. For the VaR_5 and VaR_{10}, VaR with EVT provides a number of failures closer to the expected values in comparison with VaR with normality. The backtesting analysis thus confirms the results previously obtained by the frontier.

Using simulated $t(2.2)$ data, we have shown the type of analysis to determine the haircuts that may be applied to the return distribution. The results of the analysis presented in this section are specific to the data-generating process that we have used, and thus are meant to highlight (rather than to make specific points) the methodology that can be used to select the appropriate risk measure. Similar results are obtained for the market data. These results are summarized in the following two points. First, VaR with EVT provides an estimate that does not lead to uncollateralized risk at most high quantiles of the distribution (Figure 13, panel 1). Second, VaR with normality provides lower haircuts than VaR with EVT, but underestimates risk at high quantiles.

	Normal distribution	GPD distribution
Number of observations	10,000	
VaR$_1$		
Expected violations	100	100
Actual violations	96	46
Actual violations (per cent)	0.96	0.46
VaR$_5$		
Expected violations	500	500
Actual violations	190	277
Actual violations (per cent)	1.90	2.77
VaR$_{10}$		
Expected violations	1000	1000
Actual violations	304	622
Actual violations (per cent)	3.04	6.22

Table 4: Simulated $t(2.2)$ Returns – Backtesting Risk Measures

The results of testing the confidence level of VaR for different thresholds of tail risk, when using a normal or GPD distribution.

(a) Panel 1   (b) Panel 2

(c) Panel 3   (d) Panel 4

Figure 13: MARKET DATA - ANALYSIS OF THE RISK-COST TRADE-OFF

This figure shows the risk-cost frontier analysis for the market data. Panel 1 shows the frontier constructed from the quantiles of the data and the frontier resulting from VaR estimated with EVT methods. Panel 2 shows the frontier constructed from the quantiles of the data and the frontier resulting from VaR estimated with a normal distribution. Panels 3 and 4 compare the frontiers for three models, VaR with normality, VaR with EVT methods, and ES with EVT methods.

There are two general points that we take from these case studies. First, the risk-cost frontier is a useful tool with which to compare risk measures. Second, the selection of the particular methodology (risk measure) can be based on three properties: *coherence*, *efficiency*, and *accuracy*. Recent research has shown that the breakdown of coherence does not seem to occur easily unless, (i) the return distribution is *excessively* fat tailed, or (ii) when derivative contracts are used. Regarding the latter two properties, we define efficiency as a measure of the distance from the risk-cost frontier given by the empirical (i.e., built from the data) distribution and the frontier given by the data-generating process, and accuracy as a measure of the goodness of fit of the risk measure at high quantiles of the return distribution. The weight placed on each property depends on the objective of the risk manager.

7. Conclusions

In this paper we investigated how estimation techniques based on extreme value theory affect the risk and cost attributes of different risk measures used to calculate haircuts for collateral instruments that exhibit fat tails. In doing so, we have proposed a framework that permits us to (i) characterize the risk-cost trade-off for a particular risk measure and estimation methodology, and (ii) contrast such a risk measure with other risk measures and estimation methodologies. This framework is a valuable diagnostic tool with which to understand the risk-cost trade-off implied by the internal methodology to calculate collateral value (haircuts) of institutions that use collateral to cover their exposures. These institutions may be clearing houses, central counterparties, payment system operators, central banks, or commercial banks determining their risk capital.

Most importantly, our results highlight the importance of carefully exploring the returns of the collateral instrument to determine the most appropriate model for the tails. There are at least three possible directions for future research. First, the proposed framework could be extended to a portfolio of collateral. Second, future work should consider the effects of liquidity shortages that prevent the rapid liquidation of collateral during extreme events. Third, debt instruments are a significant part of the portfolio of collateral, and thus a study of their valuation for extreme events is an important area to mitigate payment risk.

References

Artzner, P., F. Delbaen, J. Eber, and D. Heath. 1997. "Thinking Coherently." *Risk* 10: 33–49.

———. 1999. "Coherent Measures of Risk." *Mathematical Finance* 9: 203–228.

Balkema, A.A. and L. de Haan. 1974. "Residual Lifetime at Great Age." *Annals of Probability* 2: 792–804.

Berger, A.N., D. Hancock, and J.C. Marquardt. 1996. "A Framework for Analyzing Efficiency, Risks, Costs, and Innovations in the Payments System." *Journal of Money, Credit and Banking* 28: 696–732.

BIS. 2001. "Recommendations for Securities Settlement Systems." *Bank for International Settlements and International Organization of Securities Commissions Report* .

———. 2004. "Recommendations for Central Counterparties." *Bank for International Settlements and International Organization of Securities Commissions Report* .

Bollerslev, T. 1986. "Generalized Autoregressive Conditional Heteroscedasticity." *Journal of Econometrics* 31: 307–327.

Dacorogna, M.M., R. Gençay, U.A. Müller, R.B. Olsen, and O.V. Pictet. 2001. *An Introduction to High-Frequency Finance.* San Diego: Academic Press.

Danielsson, J., L. de Haan, L. Pend, and C.G. de Vries. 2001. "Using a Bootstrap Method to Choose the Sample Fraction in Tail Index Estimation." *Journal of Multivariate Analysis* 76: 226–248.

deHaan, L. 1990. "Fighting the ARCH-enemy with Mathematics." *Statistica Neerlandica* 44: 45–68.

Deheuvels, P., E. Hausler, and D.M. Mason. 1988. "Almost Sure Convergence of the Hill Estimator." *Mathematical Proceeding of Cambridge Philosophical Society* 104: 371–381.

Embrechts, P. 1999. "Extreme Value Theory in Finance and Insurance." Manuscript, Department of Mathematics, ETH, Swiss Federal Technical University.

Embrechts, P., C. Kluppelberg, and C. Mikosch. 1997. *Modeling Extremal Events for Insurance and Finance.* Berlin: Springer.

Embrechts, P., S. Resnick, and G. Samorodnitsky. 1998. "Extreme Value Theory as a Risk Management Tool." Manuscript, Department of Mathematics, ETH, Swiss Federal Technical University.

Engle, R.F. 1982. "Autoregressive Conditional Heteroscedastic Models with Estimates of the Variance of United Kingdom Inflation." *Econometrica* 50: 987–1007.

Falk, M., J. Hüssler, and R. Reiss. 1994. *Laws of Small Numbers: Extremes and Rare Events.* Basel: Birkhäuser.

Fisher, R.A. and L.H.C. Tippett. 1928. "Limiting Forms of the Frequency Distribution of the Largest or Smallest Member of a Sample." *Proceeding of Cambridge Philosophical Society* 24: 180–190.

Gençay, R. and F. Selçuk. 2006. "Overnight Borrowing, Interest Rates and Extreme Value Theory." *European Economic Review* 50: 547–63.

Gençay, R., F. Selçuk, and A. Ulugülyağcı. 2002. "EVIM: A Software Package for Extreme Value Analysis in Matlab." *Studies in Nonlinear Dynamics and Econometrics* 5: 213–239.

———. 2003. "High Volatility, Thick Tails and Extreme Value Theory in Value-at-Risk Estimations." *Insurance: Mathematics and Economics* 33: 337–356.

Gilli, M. and E. Këllezi. 2005. "An Application of Extreme Value Theory for Measuring Financial Risk." Mimeo.

Gnedenko, B.V. 1943. "Sur la Distribution limite du Terme d'une Serie Aleatoire." *Annals of Mathematics* 44: 423–453.

Goldie, C.M. and R.L. Smith. 1987. "Slow Variation with Remainder: Theory and Applications." *Quarterly Journal of Mathematics, Oxford 2nd Series* 38: 45–71.

Hall, P. 1982. "On Some Simple Estimates of an Exponent of Regular Variation." *Journal of the Royal Statistical Society, Series B* 44: 37–42.

Hill, B.M. 1975. "A Simple General Approach to Inference about the Tail of a Distribution." *Annals of Statistics* 3: 1163–1174.

Hosking, J.R.M. and J.R. Wallis. 1987. "Parameter and Quantile Estimation for the Generalized Pareto Distribution." *Technometrics* 29: 339–349.

Jenkinson, A.F. 1955. "Distribution of the Annual Maximum (or Minimum) Values of Meteorological Elements." *Quarterly Journal of Royal Meteorological Society* 81: 145–158.

Leadbetter, M., G. Lindgren, and H. Rootzén. 1983. *Extremes and Related Properties of Random Sequences and Processes.* Springer Series in Statistics. New York Berlin: Springer-Verlag.

Mandelbrot, B. 1963. "New Methods in Statistical Economics." *Journal of Political Economy* 71: 421–40.

Markowitz, H. 1952. "Portfolio Selection." *The Journal of Finance* 7: 77–91.

Mason, D.M. 1982. "Laws of Large Numbers for Sums of Extreme Values." *Annals of Probabability* 10: 756–764.

McNeil, A.J. 1997. "Estimating the Tails of Loss Severity Distributions using Extreme Value Theory." *ASTIN Bulletin* 27: 1117–137.

———. 1999. "Extreme Value Theory for Risk Managers." *Internal Modeling and CAD II, Risk Books* 93–113.

Pickands, J. 1975. "Statistical Inference using Extreme Order Statistics." *Annals of Statistics* 3: 119–131.

Rootzén, H., M.R. Leadbetter, and L. de Haan. 1992. "Tail and Quantile Estimators for Strongly Mixing Stationary Processes." Report, Department of Statistics, University of North Carolina.

Szegö, G. 2005. "Measures of Risk." *European Journal of Operational Research* 163: 5–19.

vonMises, R. 1936. "La distribution de la plus grande de n valeurs." *Rev. Math. Union Interbalcanique 1, 141-160. Reproduced in selected papers of Richard von Mises, American Mathematical Society, 1964* 2: 271–294.

Bank of Canada Working Papers
Documents de travail de la Banque du Canada

Working papers are generally published in the language of the author, with an abstract in both official languages, and are available on the Bank's website (see bottom of page). *Les documents de travail sont généralement publiés dans la langue utilisée par les auteurs; ils sont cependant précédés d'un résumé bilingue. On peut les consulter dans le site Web de la Banque du Canada, dont l'adresse est indiquée au bas de la page.*

2006