Applications of extreme value theory to collateral valuation

Alejandro García
Senior Analyst, Financial Risk Office, Bank of Canada

Ramazan Gençay
Professor of Economics, Simon Fraser University[1]

Clearing and settlement systems play a critical role in the infrastructure of financial markets because of the large values of funds and securities that settle through these systems in a given year. For instance, U.S.$49.9 trillion were settled in 2005 by the Canadian securities clearing and settlement system (CDSX). Given these large values flowing through clearing and settlement systems, regulators and banking professionals have undertaken initiatives to make them safer.

A common factor in many of these initiatives is the use of collateral to manage financial risks. For example, when participants owe money in a clearing and settlement system they may need to pledge collateral equivalent in value to the amount owing. If a participant fails, and is unable to pay the amount owed, the collateral can then be sold to generate the needed funds. However, collateral itself may consist of risky assets and thus can change in value over time. An important concern, therefore, is to require a large enough pledge of collateral, so that if a failure occurs all losses can be adequately covered. To manage the risk created by the uncertainty surrounding the future value of collateral, the initial value of collateral is discounted. In other words, participants must pledge a greater amount of collateral than the amount owing. This discount is often referred to as the 'haircut.' The higher the haircut the lower the risk that the collateral does not cover the exposure, but the higher the costs incurred by participants using the system. The methodology to calculate a haircut can most easily be explained with an example.

Consider an exposure of $100 in a securities clearing and settlement system. This exposure is collateralized by an asset that has a market price of $100. One way of estimating a haircut for such an asset is to assume a distribution for the returns of the underlying asset — for example, a normal return distribution — and select a risk measure — for example, Value-at-Risk (VaR). Supposing that the asset's daily percentage change in price has a mean of zero and a standard deviation of 3 percent we can estimate the corresponding normal distribution. Next we choose a confidence level for the haircut — for example, 0.5 per cent[2] — and then we select a holding period — for example, 1 day. Finally, we calculate the corresponding VaR obtained from a normal distribution with the mean and standard deviation of the data and assign this value as the haircut[3]. Using this approach yields a haircut of 7.72 percent, which is associated with a tail risk of 0.5 percent (confidence level). With this haircut the amount of collateral that would be required to cover the exposure of $100, given the characteristics of the asset pledged would be $108.36 which is [100/(1-haircut)].

The objective of the paper is to propose a framework that can be used to compare different methodologies to calculate haircuts. We pay particular attention to selecting a methodology that is appropriate to cover low probability events (i.e., large unexpected declines in asset prices) which might have a high impact on the stability of the financial system, and which also takes into account considerations related to the cost of pledging collateral.

Methodology for estimating haircuts

Two components are necessary to calculate a haircut for collateral. The first is a model of the distribution of losses (i.e., the frequency with which the asset declines in value), since the distribution of returns is unknown, and the second is a risk measure, which can be thought of as a way of mapping the loss distribution into a single number (the haircut). There are several ways to model the loss distribution for the collateral based on historical return data. These include:

- **Parametric approaches –** which are those that use historical return data to obtain the necessary parameters to characterize a given distribution (i.e., Normal, t, Pareto, etc.). These parameters are then used to approximate the return distribution, and the haircut is obtained from the resulting quantile given a distribution and a confidence level[4].
- **Non-parametric approaches –** such as historical simulation techniques do not model the return distribution under

1 The views expressed in this paper are those of the authors. No responsibility for them should be attributed to the Bank of Canada or Simon Fraser University. We would like to thank Dinah Maclean for her comments and suggestions with this paper. We gratefully acknowledge the financial support from the Natural Sciences and Engineering Research Council of Canada, the Social Sciences and Humanities Research Council of Canada, and the Bank of Canada.

2 This means that 1 day out of 200 the haircut would not be sufficient to cover the daily price fluctuations.

3 VaR is simply a quantile of the loss distribution of returns. This quantile represents the maximum loss which is not exceeded with a given high probability.

4 Quantiles are points taken at regular intervals from the cumulative distribution function. Dividing the ordered data into q equal sized data subsets is the motivation for q-quantiles. The quantiles are the corresponding data values marking the boundaries between consecutive subsets.

some explicit parametric model, but instead use the empirical distribution of the data to estimate the quantiles given a confidence level.

Along with choosing one of the above mentioned approaches, the estimation of haircuts requires a means of quantifying risk; that is a risk measure. There are multiple risk measures that may be used. In this paper we use two measures Value-at-Risk (VaR) and Expected Shortfall (ES).

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. 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 et al. (1997, 1999), for example, state the following problems: VaR measures only percentiles of profit-loss distributions and thus disregards any loss beyond the VaR level (tail risk), and that VaR is not coherent since it is not sub-additive. Because of these problems, we also use Expected Shortfall (ES), which addresses some of the problems of VaR.

The ES of an asset is the average loss given that VaR has been exceeded. For example, the α percent ES is the conditional mean of r_t given that $r_t > VaR_t(\alpha)$: $ES_t(\alpha) = E[r_t \mid r_t > VaR_t(\alpha)]$, where r_t represents the time series of returns for the asset pledged as collateral. Although ES is a coherent measure (sub-additive) it is still 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. This is why choosing the appropriate distribution is of critical importance to obtain an accurate characterization of risk.

Using extreme value theory to characterize the distribution of returns

There are a number of empirical observations that generally hold for a wide range of financial time series [Mandelbrot (1963)]. One of these is that return series have fat tails. This means that compared to a normal distribution there are fewer observations around the mean and more in the tails or extremes of the distribution. This is true for many equities and certain fixed income instruments that may be pledged as collateral. For such assets, it is not appropriate to use a normal distribution to estimate market return distributions. This is because the normal distribution cannot capture values at very low or high tails of the distribution. An alternative is to use extreme value theory (EVT) methods to model the tail behavior of the return distribution [Embrechts et al. (1997)]. The intuition of EVT is as follows. While the normal distribution is the important limiting distribution for sample averages (central limit theorem), the family of extreme value distributions is used as the limiting distribution of the sample extremes. Thus, it is more relevant when we are interested in the extremes of the distribution. This family can be presented under a single parameterization known as the generalized extreme value distribution (GEV)[5]. The GEV distribution has two parameters, one is called the shape parameter, and the other is called the tail index; both of these parameters describe the distribution.

Another useful result from extreme value theory is related to the behavior of large observations that exceed a high threshold. For example, if one asks the question what is the distribution of losses that exceed a 5% loss, a result from EVT tells us that given a high threshold u (i.e., 5% in this case), the proba-

5 This result is known as the Fisher-Tippett theorem.

bility distribution of excess values of the returns over threshold u may be approximated by the generalized Pareto distribution (GPD)[6]. Similarly to the GEV, the GPD embeds a number of other distributions. In particular, when its shape parameter ξ is greater than zero 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. When $\xi = 0$ the GPD corresponds to the thin-tailed distributions, and it corresponds to finite-tailed distributions for $\xi < 0$.

The risk-cost frontier

Having suggested some alternative methodologies for collateral haircuts, that is, combinations for modeling the loss distribution and the risk measure, we now need a framework for comparing and selecting methodologies. We propose the 'risk-cost frontier' as such a framework. The frontier is a way of summarizing the risk-cost trade-off implied by each methodology one is comparing. Each methodology has its own trade-off between the risks that price fluctuations in collateral value are not covered by a haircut, which we call tail risk, and the cost of pledging collateral, measured by the excess collateral above the exposure that corresponds to the haircut, which we call collateral cost. The trade-off exists because larger haircuts imply lower tail risk but higher collateral cost.

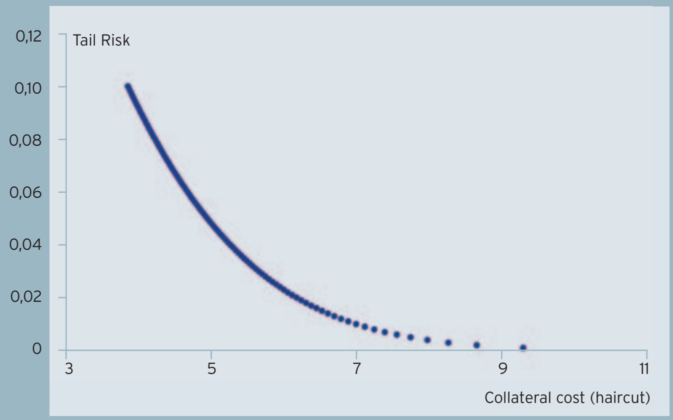

Figure 1 – Risk-cost frontier under normality (Haircuts and tail risk values obtained from a simulation of a normal distribution with a zero mean and standard deviation of 3 percent.)

The risk-cost frontier can be constructed by calculating haircuts for different levels of tail risk but using the same methodology to model the return distribution. To do this the level of tail risk could start, for example, at the 0.5 percent level and go up to 10 percent[7]. For these levels of tail risk we can calculate the associated haircuts. Having these estimations of haircuts and tail risks we can construct a risk-cost frontier from each pair of points of haircut and tail risk. Figure 1 depicts the risk-cost frontier corresponding to the example given earlier (normal with mean zero and standard deviation of 3 percent and a VaR risk measure).

This framework can then be used to compare different methodologies to calculate haircuts for the same asset. This can be done by comparing different risk-cost frontiers emanating from different methodologies to a benchmark frontier constructed from the observed return data. Having this information can help organizations understand their potential for model risk as well as establish confidence intervals for the haircuts.

Evaluating haircut methodologies with the frontier

Using the risk-cost frontier we can compare different methodologies to calculate haircuts. This can be done by calculating haircuts for the same levels of tail risk but using different methodologies (i.e., combinations of models for the loss distribution and risk measures). The risk-cost frontier can be used to determine the most appropriate methodology by selecting the one that has a frontier that is closest to a benchmark frontier, constructed from the data, but does not cross it, and therefore, does not underestimate the haircuts. To illustrate this, consider the following example using simulated data. We simulate data for the return series of a hypothetical asset using a t-distribution with 2.2 degrees of freedom. This specification shares similar statistical properties, such as fat tails, with those found in financial time series. We then use two different methodologies to estimate the haircuts. Knowing the underlying data generating distribution (t with 2.2 d.f.) allows us to determine that the best methodology to calculate the haircut is the one that has a risk-cost frontier closer to the benchmark risk-cost frontier calculated

6 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 the generalized Pareto distribution (GPD), because as the threshold gets larger, the excess distribution converges to the GPD.

7 This is equivalent to a haircut where there is a 0.5 (or 10) percent probability that the haircut does not cover the decline in price for the security used as collateral.

directly for the simulated data (using a non-parametric approach).

In this example (t with 2.2 d.f.) we compare two methodologies. Both use a parametric approach, but one assumes a normal distribution and the other an extreme value distribution. Both methodologies use VaR as the risk measure. This comparison allows us to isolate the benefits of using the extreme value distribution. Figure 2 shows the three risk-cost frontiers: the benchmark case with a red line (non-parametric approach for the empirical quantiles), the methodology based on the normal distribution with a green line, and the methodology that uses an extreme value theory distribution with a blue line. Figure 2 illustrates the mismeasurement of risk when comparing the risk-cost frontier of the methodology that assumes a normal distribution with the benchmark risk-cost frontier calculated from the simulated data (denoted by a red line in Figure 2). Also in Figure 2 we observe that the methodology that uses an extreme value distribution gives haircuts that are closer to benchmark (especially at low levels of tail risk), given by the quantiles of the simulated t data (red line in Figure 2). Figure 2 suggests that the methodology that uses an extreme value distribution is the more appropriate one.

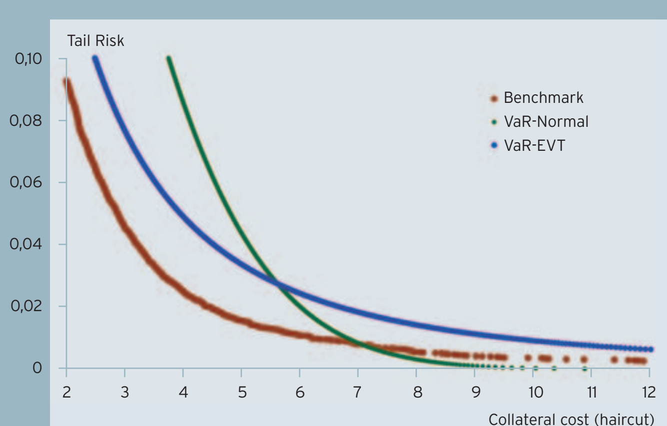

Figure 2 – Methodologies to calculate haircuts for simulated data (VaR-Normal and VaR-EVT)

We also conduct the same analysis using real market data. We use the closing price of the Dow Jones Industrials Average Index from the 3rd of January, 1930 to the 23rd of May, 2006. Using the time series of prices we construct the daily returns. Using the returns for the Dow Jones Index we compare three methodologies: VaR with normal distribution, VaR with an extreme value distribution, and ES with an extreme value distribution. The result of the risk-cost frontier analysis is shown in Figure 3[8].

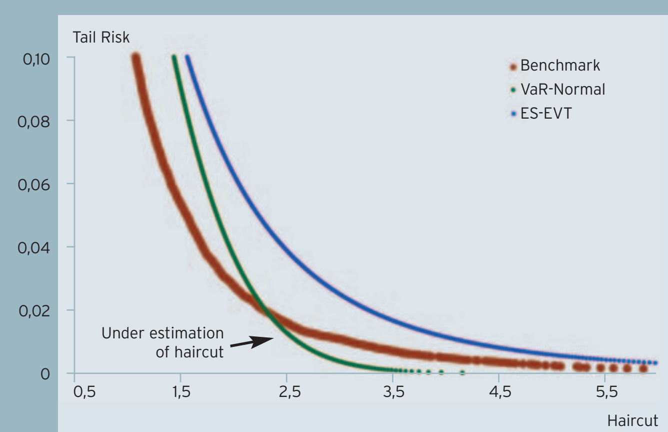

Figure 3a - Methodologies to calculate haircuts for real data

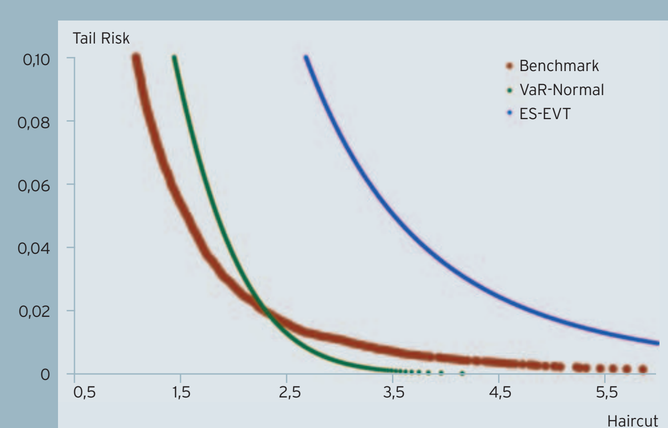

Figure 3b - Methodologies to calculate haircuts for real data

8 Figure 3 shows the risk-cost frontier analysis for the Dow Jones Industrial Average data. The top panel shows the benchmark frontier constructed from the quantiles of the data, the frontier resulting from VaR estimated with a normal distribution, and the frontier resulting from VaR estimated with EVT methods. The bottom panel shows the benchmark frontier constructed from the quantiles of the data, the frontier resulting from VaR estimated with a normal distribution, and the frontier resulting from ES estimated with EVT methods.

The results from our analysis using the risk-cost frontier for simulated and real data can be summarized as follows:

1. Methods that use VaR on the assumption of normality over- (at high levels of tail risk) and underestimate (at low levels of tail risk) the values for the haircuts. This happens because the risk-cost frontier that uses the normality assumption crosses the benchmark frontier constructed from the empirical quantiles (red line in Figures 2 and 3). Thus, for the purpose of covering extreme risk, VaR with normality may not be adequate.
2. VaR calculated with EVT methods provides a good fit in terms of slope to quantiles of the data. Nevertheless, VaR with EVT gives greater values for haircuts compared to the actual quantiles of the data. For the purpose of covering extreme risk, VaR with EVT is adequate. However, it should be kept in mind that although they provide a cushion for extreme events, greater haircuts are costly to participants of the system pledging collateral.
3. Although ES is a coherent risk measure, it represents a large cost in terms of the extra collateral required to cover the exposure. The use of ES is a more appropriate risk measure when calculating haircuts for portfolios, and thus, it may not be necessary in this case since we are focusing on individual securities.

Ultimately, the selection of the methodology to calculate haircuts depends on the weight placed on collateral costs versus coverage of extreme risk and this depends on the objectives of the risk manager. For example, those risk managers in critical financial infrastructures may choose to select a haircut that corresponds to higher quantile, compared to those in other organizations that may select a haircut that corresponds to a lower quantile since they may have greater risk tolerance. No matter what the weights placed on risk and cost may be a careful examination of the statistical properties of the return distribution is always recommended in order to select the most appropriate methodology to calculate haircuts.

Conclusion

In this paper we propose a framework that allows us to characterize the risk-cost trade-off for a particular risk measure and haircut estimation methodology, and to compare different risk measures from alternative estimation methodologies, using the risk-cost frontier. The framework proposed is useful to understand the risk-cost trade-off implied by the methodology used 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.

References

- Artzner, P., F. Delbaen, J. Eber, and D. Heath, 1997, "Thinking coherently," Risk, 10, 33–49
- Artzner, P., F. Delbaen, J. Eber, and D. Heath, 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
- BIS, 2001, "Recommendations for securities settlement systems," Bank for International Settlements and International Organization of Securities Commissions Report
- BIS, 2004, "Recommendations for central counterparties," Bank for International Settlements and International Organization of Securities Commissions Report
- Dacorogna, M. M., R. Gençay, U. A. Müller, R. B. Olsen, and O. V. Pictet, 2001, An introduction to high-frequency finance, Academic Press, San Diego
- 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
- de Haan, 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 Proceedings of Cambridge Philosophical Society, 104, 371-381
- Embrechts, P., C. Kluppelberg, and C. Mikosch, 1997, "Modeling extremal events for insurance and finance," Springer, Berlin
- Falk, M., J. Hüssler, and R. Reiss, 1994, "Laws of small numbers: extremes and rare events," Birkhäuser, Basel
- Fisher, R. A. and L. H. C. Tippett, 1928, "Limiting forms of the frequency distribution of the largest or smallest member of a sample," Proceedings 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. Ulugülyagcı, 2002, "EVIM: A software package for extreme value analysis in Matlab," Studies in Nonlinear Dynamics and Econometrics, 5, 213-239
- Gençay, R., F. Selçuk, and A. Ulugülyagcı, 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, Springer-Verlag, New York Berlin
- 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
- McNeil, A. J., 1999, "Extreme value theory for risk managers," Internal Modeling and CAD II, Risk Books, pages 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
- von Mises, R., 1936, "La distribution de la plus grande de n valeurs," Rev. Math. Union Interbalcanique1, 141-160. Reproduced in selected papers of von Mises, R., 1964, American Mathematical Society, 2, 271–294