Collateral Valuation for Extreme Market Events

*Alejandro García and Ramazan Gençay**

Clearing and settlement systems are critical to the infrastructure of financial markets because of the large values of funds and securities that settle through them. For instance, in 2005, $49.9 trillion was settled through the Canadian securities clearing and settlement system (CDSX). Given the large values flowing through these systems, regulators and banking professionals have taken initiatives to make them safer.

A common factor in many of these initiatives is the use of collateral to manage financial risks. For example, participants in a clearing and settlement system may have to pledge collateral equivalent in value to the amount they owe. If a participant fails and is unable to pay the amount owing, the collateral can be sold to generate the needed funds. But collateral itself may consist of risky assets and thus can change in value over time. It is therefore necessary to require a pledge of collateral large enough to adequately cover all losses in the event of a failure.

To manage the risk created by the uncertainty surrounding the future value of collateral, the initial value of the 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."[1] The larger the haircut, the lower the risk, but the higher the costs incurred by participants using the system.

In this article, we propose a framework that can be used to compare different methods for calculating haircuts. Particular attention is paid to selecting an appropriate method for low-probability events (e.g., large, unexpected declines in asset prices) that might affect the stability of the financial system, and one that also takes into account the cost of pledging collateral.

Methods for Estimating Haircuts

Two components are needed to calculate a haircut for collateral. The first is a model of the distribution of losses (i.e., frequency with which the asset declines in value), since the distribution of returns is unknown. 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 collateral based on historical data for returns. These include:

- **Parametric approaches** that use historical data to obtain the parameters necessary to characterize a given distribution (e.g., Normal, *t*, etc.). These parameters are then used to approximate the return distribution, and the haircut is obtained from the resulting quantile, given a particular distribution and a confidence level.[2]
- **Non-parametric approaches**, such as historical-simulation techniques, that do not model the return distribution under some explicit parametric model, but instead use the empirical distribution of the data to estimate the quantiles, for a given confidence level.

1. The haircut represents the amount by which the security could decline in value subject to a confidence level and a holding period.

* This article summarizes García and Gençay (2006).

2. 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.

Along with choosing one of the above approaches, the estimation of haircuts requires a means of quantifying risk: a risk measure. Various risk measures can be used. One of the most common is the Value at Risk (VaR). We also use an alternative risk measure called Expected Shortfall (ES).[3]

The method for calculating a haircut can most easily be explained with an example. Consider an exposure of $100 in a system for clearing and settling securities. This exposure is collateralized by an asset that has a market price of $100. To estimate the haircut for such an asset, we use a parametric approach (e.g., a normal return distribution) and select a risk measure (e.g., VaR). Knowing that the asset has a daily percentange change in price with a mean of zero and a standard deviation of 3 per cent, we estimate the corresponding normal distribution. Next, we choose a confidence level for the haircut (e.g., 0.5 per cent)[4] and then select a holding period (e.g., 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.[5] This parametric approach, combined with VaR, yields a haircut of 7.72 per cent (quantile of the distribution), which is associated with a tail risk of 0.5 per cent (confidence level). With this haircut, the amount of collateral required to cover the exposure of $100, given the characteristics of the asset pledged, would be $108.36 which is (100/[1–haircut]).

Using Extreme Value Theory to Characterize the Distribution of Returns

A number of empirical observations generally hold for a wide range of financial time series.[6] One of these is that return series have fat tails. This means that compared with 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 the distribution of market returns. This is because the normal distribution cannot capture values at very low or high tails of the distribution. Extreme value theory (EVT) methods are more appropriate for modelling the tail behaviour of the distribution of returns for securities.[7]

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 parametrization known as the generalized extreme value distribution.[8]

The power of EVT methods to capture extreme events is illustrated in Gençay and Selçuk (2006), where the authors use data for Turkey's overnight interest rate prior to the crisis when the rate reached a level of 873 per cent on 1 December 2000 and 4,000 per cent on 21 February 2001. The authors find that estimation results from the pre-crisis data indicate that a day with overnight interest rates over 1,000 per cent (simple annual) could be expected every 4 years. In other words, the extraordinary levels observed during the crisis were in the nature of the economy before they actually materialized.

The Risk-Cost Frontier

Having suggested some alternative methods for estimating collateral haircuts, we now need a framework for comparing the methods. 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 method. Each method has its own trade-off between the risk that price fluctuations in collateral value are not covered by a haircut (tail risk), and the cost of pledging collateral, measured by the excess collateral above the exposure that corresponds

3. ES is a *coherent* alternative to VaR, where *coherence* is defined as axioms that capture the desired properties of a risk measure. This term is from Artzner et al. (1997, 1999).
4. This means that 1 day out of 200, the haircut would not be sufficient to cover the daily price fluctuations.
5. VaR is simply a quantile of the loss distribution of returns. This quantile represents the maximum loss that is not exceeded with a given high probability.
6. A good reference of the stylized facts for financial time series can be found in Mandelbrot (1963).
7. Embrechts, Klüppelberg, and Mikosch (1997) is a comprehensive source of theory and applications of extreme value theory to the finance and insurance literature.
8. This result is known as the Fisher-Tippett theorem.

Chart 1 Risk-Cost Frontier under Normality

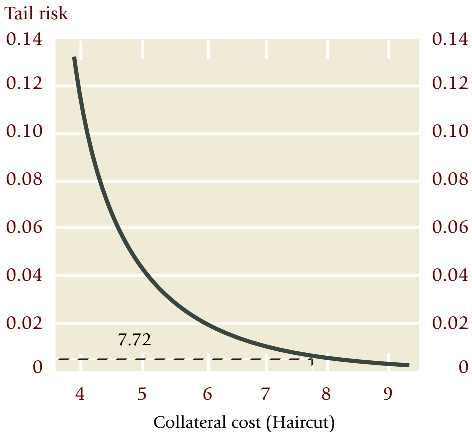

Note: Haircuts and tail risk values obtained from a simulation of a normal distribution with a zero mean and standard deviation of 3 per cent.

Chart 2 Comparison of Methods for Calculating Haircuts

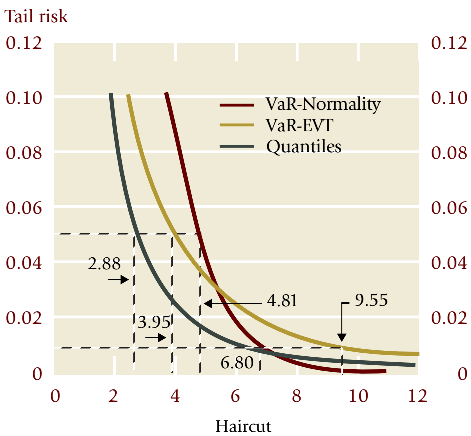

Note: Collateral returns are simulated from a *t* distribution with 2.2 degrees of freedom.

to the haircut (collateral cost). The trade-off exists because *larger haircuts* imply *lower tail risk* but *higher collateral cost*.

The risk-cost frontier can be constructed by calculating haircuts for different levels of tail risk but using the same method to model the return distribution. For example, the level of tail risk could start at 0.5 per cent and go up to 10 per cent. We can then calculate the associated haircuts. From these pairs of points, we can construct a risk-cost frontier. Chart 1 depicts the risk-cost frontier corresponding to the example given earlier (normal with mean zero and standard deviation of 3 per cent and a VaR risk measure).

Evaluating Haircut Estimation Methods

The risk-cost frontier can be used to compare different methods of calculating haircuts. Haircuts for the same levels of tail risk are calculated using different methods (i.e., combinations of (i) models for the loss distribution and (ii) risk measures).

The risk-cost frontier can then be used to determine the most appropriate method by selecting one whose frontier is closest to a benchmark frontier constructed from the data, but that does not cross it and, therefore, does not underestimate the haircuts. Consider the following example. First, the returns on a hypothetical asset are simulated using a *t*-distribution with 2.2 degrees of freedom. This specification shares similar statistical properties, such as fat tails, with those in financial time series. Two different methods are then used to estimate the haircuts. Knowing the underlying data-generating distribution allows us to determine that the best method for calculating the haircut is the one that has a risk-cost frontier closer to the risk-cost frontier calculated directly for the simulated data (using a non-parametric approach).

In this example, we compare two methods: both use a parametric approach, but one will assume a normal distribution and one an extreme value distribution. Both methods use VaR as the risk measure. Chart 2 shows the three risk-cost frontiers: the benchmark case with a green line (non-parametric approach for the empirical quantiles), the method based on the normal distribution with a red line, and the method that uses an extreme value theory distribution with a gold line.

Chart 2 illustrates the mismeasurement of risk when comparing the risk-cost frontier of the method that assumes a normal distribution, with the benchmark risk-cost frontier calculated from the simulated data (denoted by a green line). In Chart 2, we also observe that use of an extreme value distribution gives haircuts that are closer to benchmark given by the quantiles of the simulated *t* data (green line in Chart 2). Chart 2 suggests that the method that uses an extreme value distribution is the more appropriate one.

In our study, we also conduct the same analysis using real market data and find similar results. These results can be summarized as follows:

- Methods that use VaR on the assumption of normality overestimate (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 (green line in Chart 2). Thus, for the purpose of covering extreme risk, VaR with normality may not be adequate.
- VaR calculated with EVT methods provides a good fit in terms of slope to quantiles of the data. Nevertheless, VaR with EVT gives larger values for haircuts compared with the actual quantiles of the data. For the purpose of covering extreme risk, VaR with EVT is adequate. It should be kept in mind, however, that although they provide a cushion for extreme events, larger haircuts are costly to participants of the system.

Ultimately, the selection of the method for calculating haircuts depends on the weight placed on collateral costs versus coverage of extreme risk, and this depends on the objectives of the risk manager. Managers in critical financial infrastructures may choose to select a haircut that corresponds to a higher quantile than managers in organizations with greater tolerance for risk. 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 method for calculating haircuts.

Conclusions

We propose a framework that allows us to (i) characterize the risk-cost trade-off for a particular risk measure and method of haircut estimation, and (ii) compare different risk measures from alternative estimation methods, using the risk-cost frontier. The framework proposed is useful for understanding the risk-cost trade-off implied by the method used to calculate the collateral value (haircuts) that institutions must pledge 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.-M. Eber, and D. Heath. 1997. "Thinking Coherently." *Risk* 10: 68–71.

———. 1999. "Coherent Measures of Risk." *Mathematical Finance* 9: 203–28.

Embrechts, P., C. Klüppelberg, and T. Mikosch. 1997. *Modelling Extremal Events for Insurance and Finance.* New York: Springer.

García, A. and R. Gençay. 2006. "Risk-Cost Frontier and Collateral Valuation in Securities Settlement Systems for Extreme Market Events." Bank of Canada Working Paper No. 2006-17.

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

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