Haircutting non-cash collateral

Value-at-risk-based, data-driven haircut models are subject to data quality issues and lack flexibility for further analysis. Wujiang Lou develops a complementary parametric haircut model to conduct sensitivity tests, capture market liquidity risk, allow idiosyncratic risk adjustments and incorporate relevant market information. Computational results show potential uses in designing collateral haircuts for collateral agreements, such as credit support annexes, and in capital calculations

Haircuts – discounts on the market value of securities taken in as collateral – draw their intuition from earlier stock loan brokers' desire to withstand stock market meltdowns without losses. That intuition remains largely intact, although when statisticians got involved and historical data became abundant, a confidence interval was used to qualify the haircut. For example, a 15% haircut would give a 99% confidence interval of no loss within 10 days, in typical value-at-risk language. Naturally, simple intuition such as this does not call for sophisticated methodologies or models. Today, haircuts appear in standard financial transaction documents, including the Master Repurchase Agreement (MRA) for repos, the credit support annex (CSA) (to the International Swaps and Derivatives Association Master Agreement) for swaps and derivatives, and exchange or central counterparty (CCP) clearing agreements. Basel's risk capital framework is the only place where the term 'haircut model' is found. For advanced banks adopting Basel's market risk capital rules, VAR methodology with at least two years of historical data and internal haircut models are allowed, although no technical specifics are given.[1]

Basel's haircut models and the Financial Stability Board's (FSB's) enhanced haircut framework (Financial Stability Board 2015) exemplify a data-driven approach to haircuts. Although prudentially guarded with qualitative and quantitative standards, a data-centric approach is only as good as the data used, and it carries the usual caveat that history may or may not repeat itself. Except for some on-the-run government securities, debt instruments do not possess market liquidity anywhere close to the equity market. It is customary in practice for historical data to be sourced from an untradeable proxy index or a representative portfolio that is similar in key product design features and risk characteristics, such as credit rating and maturity. Such a proxy is obviously subject to data accuracy issues, as underlying bonds trade sparsely. Another problem is that it eradicates progressively diverging idiosyncrasies as bonds age. When a bond is on a downgrade watch, or it has seen large spread widening relative to its peers, one may suspect it is riskier than other bonds and that its haircut will deviate away and move higher.

By relying exclusively and directly on historical data, a data-driven approach also lacks the flexibility of incorporating useful information when it becomes available. The unprecedented price behaviour of US residential mortgage-backed securities (RMBSs) in 2007–8, for instance, rendered their prior pricing history meaningless for setting haircuts via VAR. In particular, investment grade (IG) RMBSs on the subprime mortgage loans of 2005–7 vintages priced close to par before plunging into the teens. The VAR estimate at the time would have predicted single-digit haircuts, in line with Basel II's 8% haircut. Some banks, however, promptly hiked bilateral repo haircuts up by multiples of themselves (Gorton & Metrick 2012), taking into account future price volatilities that were already exhibited in their synthetic market kin: the asset-backed credit default swap index.

To address these limitations, this article constitutes a first effort to develop a parametric haircut model from asset pricing and credit risk perspectives, capturing asset volatility, jumps and market liquidity risk. It contributes to the literature by introducing credit risk measures to define haircuts such that these measures can satisfy certain predetermined criteria, eg, an AAA rating. Essentially, the original intuition of loss aversion is transformed into the credit enhancement language typical of credit derivatives and structured products. This article aims to develop a counterparty-independent haircut model, leaving counterparty-dependent haircuts to a separate effort (Lou 2016b).

Expanded haircut definitions

In a repo-style securities financing transaction, the lender is exposed to the borrower's default risk with a market-contingent exposure framed on a short window for default settlement. The margin period of risk (MPR) covers the time period from the last date when margin was met to the date when the defaulting counterparty is closed out following the completion of collateral asset disposal. The lender's exposure in a repo during the MPR is flat, as it is simply principal plus accrued and unpaid interest. A flat exposure could apply to over-the-counter derivatives netting sets under the CSA, if we assume that the derivatives exposure is hedged during an MPR with its primary market risks. So, in an idealised setting, we consider a counterparty (or borrower) C's default time at t, when the margin is last met; an MPR of u, during which there is no margin posting; and collateral assets that are sold at time $t+u$ instantaneously at the market, with a possible liquidation discount g, to account for market liquidity risk.

We denote the collateral market value as $B(t)$ and exposure to the defaulting counterparty C as $E(t)$. At time t, one share of the asset is margined properly, ie, $E(t) = (1-h)B(t)$, where h is a constant haircut, $1 > h \geqslant 0$. The margin agreement is assumed to have a zero minimum transfer amount. The lender would have a residual exposure $(E(t) - B(t+u)(1-g))^+$, where g is a constant, $1 > g \geqslant 0$. Exposure to C is assumed to be flat after t. We can write the loss function from holding the collateral as follows:

$$L(t+u) = E_t\left(1 - \frac{B_{t+u}}{B_t}\frac{1-g}{1-h}\right)^+ = (1-g)B_t\left(1 - \frac{B_{t+u}}{B_t} - \frac{h-g}{1-g}\right)^+ \tag{1}$$

[1] *The Financial Stability Board's new framework establishes a haircut floor for non-centrally cleared securities financing transactions, requiring the use of at least five years' historical data that includes at least one stress period (Financial Stability Board 2015). Banks' internal models are basically methods of selecting and justifying a proxy index or portfolio.*

Conditional on default happening at time t, the above determines a one-period loss distribution driven by the asset price return $B(t+u)/B(t)$.

Let $y = 1 - B_{t+u}/B_t$ be the price decline. If $g = 0$, $\Pr(y > h)$ is equal to $\Pr(L(u) > 0)$. If the price decline is less than or equal to h, there is no loss. A first dollar loss will occur only if $y > h$. Thus, h provides a cushion to protect against a loss being incurred. Given a target rating class's default probability p, the first loss haircut can be written as:

$$h_p = \inf\{h > 0\colon\ \Pr(L(u) > 0) \leqslant p\} \tag{2}$$

Let VAR_q denote the VAR of holding the asset; this is an amount the price decline will not exceed, given a confidence interval of q, say, 99%. In light of the adoption of expected shortfall (ES) in Basel IV, we can define a haircut as the ES under the q-quantile:

$$\begin{aligned} h_{\mathrm{ES}} &= \mathrm{ES}_q = \mathrm{E}[y \mid y > \mathrm{VAR}_q] \\ \mathrm{VAR}_q &= \inf\{y_0 > 0\colon\ \Pr(y > y_0) \leqslant 1 - q\} \end{aligned} \tag{3}$$

Without the liquidity discount, h_p is the same as VAR_q. If haircuts are set to VAR_q or h_{ES}, the market risk capital for holding the asset for a given MPR, defined as a multiple of VAR or ES, is zero. This implies that we can define a haircut to meet a minimum economic capital (EC) requirement C_0:

$$h_{\mathrm{EC}} = \inf\{h \in R^+\colon \mathrm{EC}[L \mid h] \leqslant C_0\} \tag{4}$$

where EC is measured as either VAR or ES minus expected loss (EL). For rating criteria employing the EL-based target for each rating class, we introduce one more definition of haircuts based on the EL target L_0:

$$h_{\mathrm{EL}} = \inf\{h \in R^+\colon E[L \mid h] \leqslant L_0\} \tag{5}$$

The EL target L_0 can be set based on the EL criteria of certain designated high credit ratings, whether internal or external. With an external rating such as Moody's, for example, a firm can set the haircut to a level such that the expected (cumulative) loss satisfies the EL tolerance L_0 of a predetermined rating target, eg, Aaa or Aa1. In (4) and (5), loss L's holding period does not have to be an MPR.

Unlike VAR_q, the definitions h_p, h_{EL} and h_{EC} are based on a loss distribution solely generated by collateral market risk exposure. As such, we no longer apply the usual wholesale credit risk terminology of probability of default (PD) and loss given default (LGD) to determine EL. Here, EL is computed directly from a loss distribution originated from market risk, where the haircut should be wholesale counterparty independent.

Collateral price dynamics

The loss function (1) is an out-of-the-money put on the collateral asset, which is predominantly decided by the asset return's skewness and tail characteristics. A study of haircuts is necessarily a study of tail behaviour. Asset price models with stochastic volatility and jumps in both their return and volatility are shown to improve empirical studies of stock indexes (Eraker *et al* 2003). The double exponential jump-diffusion model (DEJD) of Kou (2002) is popular in exotic and path-dependent options pricing due to its appealing asymmetric jump specification and the ease with which it performs transform analytics. Its extension, the mixed-exponential jump-diffusion model (MEM) by Cai & Kou (2011), is capable of producing a wide variety of skewed tail distributions. As the risk exposure window (MPR) covers a short period of time, just days or weeks, and stochastic volatility models are expected to have limited impact, particularly on haircut designs that depend on a negative tail, we choose the DEJD and MEM. We do so with a view that the MEM could be valuable in coping with the excessive skewness and fat tails of securitisation debts.

The log return of a jump-diffusion asset price B_t has the form:

$$X_t = \log\left(\frac{B_t}{B_0}\right) = \mu t + \sigma_{\mathrm{a}} W_t + \sum_{j=1}^{N_t} Y_j \tag{6}$$

where μ is the asset return, σ_{a} is the asset volatility, $W(t)$ is a Brownian motion, $N(t)$ is a Poisson process with intensity λ and Y_j is a random variable denoting the magnitude of the jth jump. With the MEM, Y_j, $j = 1, 2, \dots$, are a sequence of independent and identically distributed mixed-exponential random variables, with a probability density function $f_Y(x)$ given by:

$$f_Y(x) = p_{\mathrm{u}} \sum_{l=1}^{m} p_l \eta_l \mathrm{e}^{-\eta_l x} I\{x \geqslant 0\} + q_{\mathrm{d}} \sum_{j=1}^{n} q_j \theta_j \mathrm{e}^{\theta_j x} I\{x < 0\} \tag{7}$$

where p_{u} and q_{d} are up-jump and down-jump switching probabilities, $p_{\mathrm{u}} + q_{\mathrm{d}} = 1$. In addition, p_l is the weight (not necessarily in a probabilistic sense) of the lth up-jump mixture exponentially distributed at a rate of $\eta_l > 1$, $\sum p_l = 1$. Similarly, $\theta_j > 0$ is the jth down-jump mixture's rate, and q_j are weights that sum to 1, $\sum q_j = 1$. Obviously, this reduces to the DEJD when $m = n = 1$.

The probability of the cumulative loss exceeding an amount b, $L(u) \geqslant b$, ie, the tail cumulative density function (cdf), can be mapped to X_{u}'s cdf:

$$P_b \mid h = E[I\{L(u) \geqslant b\}] = \Pr\left(X_{\mathrm{u}} \leqslant \log\left(\frac{1 - h - b/B_0}{1 - g}\right)\right) \tag{8}$$

Fixing a haircut h, this gives the loss distribution P_b as a function of b. Fixing b, P_b becomes a function of h, which can be inverted to solve for h given a target level of P_b. VAR can be solved by setting $P_b = 1 - q$. Obviously, setting b to zero leads to (3). It is useful for implementations to note that (8) is translational in h and b, ie, $P_b \mid h = P_{b*} \mid h^*$, where $b^* = b + (h - h^*)B_0$.

The EL relates to the undiscounted European put option fair value:

$$E[L(u)] = (1 - g)E[(K - B_{\mathrm{u}})^+] = (1 - g)P(K) \tag{9}$$

where $K = ((1 - h)/(1 - g))B_0$ and $P(K)$ is the undiscounted put fair value. Fixing h, and thus strike K, the put fair value can be obtained by means of an inverse Laplace transform. Finding h given an EL target is a simple numerical inversion.

The intuition that a haircut is a cushion against investment loss is captured in (9) and can be seen easily when the $L = (K - B_{\mathrm{u}})^+$ term is reformatted as $((B_0 - B_{\mathrm{u}}) - hB_0)^+$ with $g = 0$. If we consider an investment of amount B_0, the haircut h effectively cuts the investment into two tranches, a senior tranche of $(1 - h)B_0$ and a subordinated tranche of hB_0. In standard collateralised debt obligation (CDO) terminology, h is the attachment point of the senior piece whose loss function is L. Supposing that L_{a} and L_{b} correspond to h_{a} and h_{b}, $h_{\mathrm{a}} < h_{\mathrm{b}}$, so $L_{\mathrm{a}} \geqslant L_{\mathrm{b}}$. $L_{\mathrm{a,b}} = L_{\mathrm{a}} - L_{\mathrm{b}}$ is the loss of a mezzanine tranche with an attachment point of h_{a} and a detachment point of h_{b}. Such a viewpoint is relevant for securitisation products, where different tranches are traded in the same market while underlying assets can be traded separately; but a data-driven haircut approach would determine haircuts irrespective of their structural linkage.

Jump-diffusion models such as the DEJD have been studied primarily for stock returns and bond yields. In our haircut model, it is also used for bond price returns. Modelling the price rather than the yield term structure for fixed-income securities is justified because the MPR is very short. Bond options, for example, are priced using the Black-Scholes option pricing model with lognormal bond price dynamics, as their terms are typically of three months or less.

The conditioned loss distribution (8) and EL (9) do not enlist the derivatives counterparty or repo borrower's credit quality. Repo-style transactions have recourse to the borrower's general credit, and, as such, repo haircuts are counterparty dependent and procyclic (Gorton & Metrick 2012). Haircuts in CSAs and regulatory capital contexts are counterparty independent, as collateral is deliberately used to mitigate counterparty credit risk. The extended definitions (2)–(5) allow counterparty-independent haircuts and counterparty-dependent haircuts to be modelled within the same analytical framework. The diffusion component of the DEJD model can be made to correlate with the dynamic spread of the counterparty or repo borrower to capture wrong-way risk, as is shown in a companion paper (Lou 2016b).

Numerical techniques

The computation of loss distribution and EL reduces to the cdf of X_t and the undiscounted European put valuation. For the mixed exponential jump-diffusion process $X(t)$ specified in (6) and (7), Cai *et al* (2014) develops a two-sided Laplace transform analysis. The Laplacian for the probability density function $f_{X(t)}$ is denoted by $L_{f_{X(t)}}$:

$$L_{f_{X(t)}}(s) = E[\mathrm{e}^{-sX_t}] = \int_{-\infty}^{\infty} \mathrm{e}^{-sx} f_{X(t)}(x)\,\mathrm{d}x = \mathrm{e}^{G(-s)t} \tag{10}$$

where the Lévy exponent function is:

$$G(x) = \tfrac{1}{2}\sigma_{\mathrm{a}}^2 x^2 + \mu x + \lambda\bigg(p_{\mathrm{u}} \sum_{l=1}^{m} \frac{p_l \eta_l}{\eta_l - x} + q_{\mathrm{d}} \sum_{j=1}^{n} \frac{q_j \theta_j}{\theta_j + x} - 1\bigg) \tag{11}$$

with a range of absolute convergence (ROAC) of $(-\min(\eta_l), \min(\theta_j))$ for $\mathrm{Re}(s)$, the real part of the complex number s. The Laplace transform for the cdf $F_{X(t)}$ of $X(t)$ is then simply $L_F(s) = L_f(s)/s$, with its ROAC $\mathrm{Re}(s) \in (0, \min(\theta_j))$. For the undiscounted European put option with strike K, if we denote its fair value P as a function of the normalised log strike k, such that $K = B_0\mathrm{e}^{-k}$, then $P(k) = E[(B_0\mathrm{e}^{-k} - B_t)^+]$. We apply the two-sided Laplace transform to $P(k)$, $k \in (-\infty, \infty)$:

$$L_{P(k)}(s) = \int_{-\infty}^{\infty} \mathrm{e}^{-sk} E[B_0(\mathrm{e}^{-k} - \mathrm{e}^{X(t)})^+]\,\mathrm{d}k = \frac{B_0}{s(s+1)} L_f(-s-1) \tag{12}$$

with a ROAC of $(-\min(\theta_j) - 1, -1)$. The Laplacian for a European call option is exactly the same as that of the put option, although the ROAC will be $\mathrm{Re}(s) \in (0, \min(\eta_j) - 1)$. Note the formula above is different from Cai *et al* (2014), where the Laplacian of the European call option price is conducted on the log of the strike K and $L_{\mathrm{EuC}}(s) = (B_0^{s+1}/(s(s+1)))L_f(-s-1)$. Our strike normalisation with moneyness is necessary, as it allows us to use the same error control over the put payoff (loss) and the tail cdf under the same parameters.

Having computed the two-sided Laplace transform of the pdf $L_{f_{X(t)}}$, the cdf $L_{F_{X(t)}}$ and the loss or put $L_{P(k)}$, the two-sided Laplacian inversion algorithm shown in the appendix is used to solve for loss distribution (8) and EL (9).

The DEJD model has six parameters $(\mu, \sigma_{\mathrm{a}}, \lambda, p, \eta, \theta)$ that need to be determined either through model estimation from historical data or calibration to traded instruments. In the context of regulatory haircuts, a historical estimation is generally required. When historical asset price data, either directly or as a proxy, is satisfactorily available and reliable, the model can be estimated via a maximum likelihood estimation (MLE). The likelihood function is written as $H(\mu, \sigma_{\mathrm{a}}, \lambda, p, \eta, \theta \mid x) = f_{X(t)}(x \mid \mu, \sigma_{\mathrm{a}}, \lambda, p, \eta, \theta)$ for a return data point x. The two-sided Laplacian of the pdf (10) can be inverted to arrive at $H(\cdot \mid x)$. Ramezani & Zeng (2007) provide an explicit likelihood function formula for the DEJD model involving the sum of double infinite series as a result of conditioning on up- and down-jump counts; they find the DEJD generally produces a better fit for stock indexes than for individual stocks.

Given a log return time series x_i, $i = 1, 2, \dots, N$, our estimation becomes an optimisation problem:

$$\max \prod_i H(\mu, \sigma_{\mathrm{a}}, \lambda, p, \eta, \theta \mid x_i)$$
$$\text{such that} \quad \sigma_{\mathrm{a}} > 0,\ \lambda > 0,\ 0 \leqslant p \leqslant 1,\ \eta > 1,\ \theta > 0 \tag{13}$$

There are no constraints other than these simple lower or upper bounds. The likelihood function can also be written as $(\mu, \sigma_{\mathrm{a}}, \lambda_{\mathrm{u}}, \lambda_{\mathrm{d}}, \eta, \theta \mid x)$ by a change of variable $\lambda_{\mathrm{u}} = \lambda p$ and $\lambda_{\mathrm{d}} = \lambda(1 - p)$. For the results presented below, we have taken optimisation routines off shelves from Matlab's median-sized constrained nonlinear optimiser `fmincon`, utilising an active-set algorithm, with sequential quadratic programming (SQP) and quasi-Newton line search.

In this article, we follow a simple estimation procedure. To begin, calculate the sample's mean μ_0 and standard deviation σ_0, and then estimate the remaining four parameters of the DEJD model with $(\mu, \sigma_{\mathrm{a}})$ fixed at (μ_0, σ_0). Now, re-estimate the model with μ fixed at μ_0, and σ_{a} relaxed so that five parameters are sought with initial values taken from the previous estimation's outcomes. Last, repeat the estimation with μ also relaxed in order to estimate the full set of six parameters. This procedure's intention is to find a stable local minimum close to the lognormal model, recognising throughout that finding the global minimum is difficult and time consuming for this type of highly nonlinear optimisation problem.

Applications and results

The counterparty-independent haircut model outlined above can be used to determine repo haircuts with or without recourse to a hindered or extremely weak repo counterparty. It could be a candidate for internal haircut models, as per Basel and FSB requirements, or used to design the non-cash collateral haircuts increasingly seen in collateral and margining agreements.[2] Below we demonstrate how the model is applied to three main non-cash

[2] *Collateral haircuts or CSA haircuts are necessarily counterparty independent. Consider overnight indexed swaps (OIS) discounting under a full cash CSA, where trades are indifferent to counterparties or free of bearers. Now, suppose that a party posts corporate bonds: the haircut applied should be such that it reproduces the counterparty-free bearer form, or the OIS discounting will no longer apply.*

1 Predicted main equity haircuts (MPR: 10 days) targeting hypothetical Moody's one-year loss rates, eg, 0.00003%, 0.00031%, 0.00075% and 0.00166%, for Aaa/Aa1/Aa2/Aa3, as per Bielecki (2008), and S&P's average one-year default rates, eg, 0.0005%, 0.001%, 0.01% and 0.02%, for AAA/AA+/AA/AA−, as per Standard & Poor's (2015)

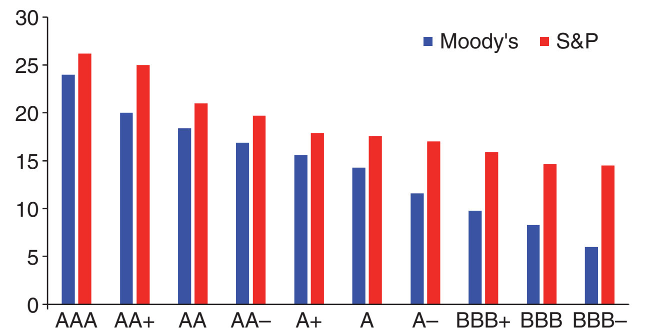

collateral classes – equity, corporate bonds and securitisation – intertwined with considerations of liquidity risk, haircut sensitivities and model risk.

Equity main index The Standard & Poor's 500 index (SPX) is commonly used as a proxy for US main equities. To satisfy the FSB's requirement of a five-year history consisting of at least one stress period, we choose the period from February 1, 2008 to February 1, 2013, as SPX experienced significant stress during the second half of 2008 and early 2009, at the height of the financial crisis. The estimated DEJD has:

$$(\mu, \sigma_{\mathrm{a}}, \lambda_{\mathrm{u}}, \lambda_{\mathrm{d}}, \eta_{\mathrm{u}}, \eta_{\mathrm{d}}) = (0.1984, 0.1512, 37.53, 40.24, 71.51, 60.56)$$

and produces −0.5136 in skewness and 10.50 in kurtosis, compared with a sample skewness of −0.2443 and kurtosis of 9.95.

Figure 1 shows predicted haircuts targeting Moody's and Standard & Poor's (S&P) IG credit ratings, ie, h_{EL} and h_p. The hypothetical S&P 'A' and above rating-targeted haircuts are about three to five points higher than Moody's corresponding rating-targeted haircuts. Caution should be taken, however, as these default rates and loss rates are examples and not directly comparable with each other. Here, our intention is to show that the methodology works for both PD and EL approaches.

Securitisation debts and liquidity risk Non-cash collateral is dominated by fixed-income instruments. Unlike equities, these instruments' market liquidity depth is very limited, except for on-the-run government debts. In practice, a data-driven haircut approach has to rely on non-tradable market indexes as proxies to the unobserved data. Because of its fixed maturity, a bond's historical price series (if reliably available) suffers from a progressive shortening of the remaining maturity. As a result, debt securities haircuts are estimated by putting them into residual maturity buckets.

Repo counterparties wishing to adopt regulatory haircuts often find the supervisory haircuts table is set too broadly, not distinguishing among many types or subclasses of assets and risk characteristics of securitised assets. US commercial mortgage-backed securities (CMBSs), one of the main securitisation products seen in repos, are not even found in the table. In fact, there is only one securitisation category. Typical triparty haircut schedules have much finer granularities, with close to 100 line items offered. With a parametric haircut model in place, more granular haircuts can be computed per asset type, subclass and maturity bucket.

A. Predicted CMBS 10-day haircuts per rating subclass and residual maturity bucket, targeting Aa1 with comparisons to raw haircuts and liquidity-adjusted haircuts

Rating	Maturity	Aa1 HC	2% LP dHC	5% LP dHC	LP HC	Raw HC (VAR)
AAA	1–5Y	5.08	1.88	4.70	6.96	4.80
AAA	5–10Y	12.27	1.68	4.22	16.49	11.52
AA	1–5Y	6.11	1.82	4.56	7.93	6.08
AA	5–10Y	20.35	1.63	4.08	24.43	23.49
A	1–5Y	9.14	1.84	4.60	10.98	6.93
A	5–10Y	20.22	1.65	4.12	24.34	26.44

Table A shows CMBS haircuts predicted using DEJD models of Bank of America Merrill Lynch CMBS price indexes estimated on daily time series from January 2008 to January 2013. The maturity up-sloping effect is evident as 5–10Y bucket haircuts are higher than those of 1–5Y bucket haircuts for all ratings. AAA rated CMBSs have much lower haircuts than those of AA and A rated bonds, eg, 5.08% for an AAA 1–5Y bucket compared with an AA at 6.11% and an A at 9.14%, when targeting Moody's Aa1 rating. The last column of table A lists raw data haircuts as the standard 10-day, 99-percentile VAR. The estimated models are able to produce haircuts that closely match the raw haircuts for 1–5Y maturity AAA and AA.

Greater differences are seen in 5–10Y maturity AA and A, eg, an AA haircut at 20.35% under an Aa1 target rating versus a 23.49% haircut from VAR. These can be attributed to market liquidity risk, intensified with private securitisation and lower-grade corporates, which trade infrequently and often rely on valuation services. Table A illustrates the additional haircuts (in percent) when 2% and 5% liquidity premiums (LPs) are applied. In this case, 2% can be seen as business-as-usual for non-super senior IG securitised products, while 5% is indicative of stress. Roughly speaking, these liquidity shocks translate into the same magnitude of additional haircuts.

Column 'LP HC' adds '2% LP dHC' to 1–5Y maturity and '5% LP dHC' to 5–10Y maturity. For AA 5–10Y CMBSs, the Aa1 haircut with 5% liquidity premium is 24.43%, slightly higher than the raw haircut of 23.49%. Using precrisis historical data (January 2002 to January 2007) shows a 3.44% haircut, while Gorton & Metrick (2012) show a bilateral haircut at 27.5% for AA rated CMBSs during the crisis. This elevated haircut is explained by the liquidity draught that faced asset-backed securities following the subprime crisis as well as sharp increases in experienced volatility and anticipated future volatility. Both liquidity and asset volatility are incorporated into our model.

Corporate debts and idiosyncratic adjustments The last major asset class we will demonstrate is US IG corporate bonds, using Bank of America Merrill Lynch US corporate bond price indexes from January 2008 to January 2013. Table B shows single A rated bonds with residual maturities of 1–5 and 5–10 years. Compared with Basel III supervisory haircuts for 20% risk-weighted wholesale issuers at 4% for a residual maturity of 1–5 years, and at 8% for 5+ years, Aaa is the target rating that gets closest to the supervisory haircuts. The differences between Aa1 targeted haircuts and the raw haircuts are of the same scale as bond bid/ask spreads.

Once a proxy historical price series is chosen and justified, the collateral haircut is trivially given via the VAR approach. An estimated parametric model will not add value if its sole use is to reproduce collateral haircuts. The parametric model is useful, however, in that it facilitates a meaningful

B. Haircuts for A rated corporate bonds. 1–5Y and 5–10Y maturities receive raw haircuts of 4.25% and 6.43%, respectively

Shifts	Aaa	Aa1	Aa2
HC (1–5Y)	3.87	2.98	2.64
HC (5–10Y)	6.49	5.19	4.68
$\mu + 1\%$	−0.03	−0.04	−0.04
$\sigma_a + 1\%$	0.37	0.34	0.32
$\lambda_u - 1$	0.01	0.01	0
$\lambda_d + 1$	0.07	0.04	0.04
$\eta_u + 10$	0.01	0	0
$\eta_d - 10$	0.26	0.2	0.18

Sensitivities shown for 5–10Y maturities with shifts based on $(\mu, \sigma_a, \lambda_u, \lambda_d, \eta_u, \eta_d) = (0.0729, 0.0525, 13.82, 31.90, 212.6, 225.6)$

sensitivity analysis of the collateral haircut, as shown in table B. Sensitivity to the up-jump rate and down-jump rate (η_u, η_d) is asymmetric as expected, as haircuts measure one-sided loss and depend on down jumps rather than up jumps. A roughly 3% shift in volatility will lead to a haircut increase of 1%. Given increased market volatility, haircut deltas or adjustments can then be computed and added to applicable haircuts; scenario analysis and stress tests can be conducted as desired.

Moreover, the parametric model allows idiosyncratic factors to be incorporated as adjustments. When model estimation is performed based on an index return, index constituents are averaged out. For a long-running index, constituents' credit quality and market performances will diverge as the index ages. So far, practitioners do not have a tool to adjust this ageing effect. With a parametric model at hand, one could create a delta model for the deteriorating credits. For instance, the down-jump size distribution parameter η_d can be adjusted down to reflect larger jump sizes, the down-jump intensity can be increased or the volatility σ_a can be hiked. On a first-order basis, the sensitivity table could thus be used to incorporate idiosyncratic risk characteristics into a proxy index or portfolio.

Parameterisation stability and model risk Compared with a data-driven haircut approach, a parametric model of haircuts introduces potential model risk. In our context, this is related to the model parameter estimation errors or stability. We conduct quarterly rolling SPX model estimations covering 2005 Q1 to 2009 Q4. For each quarter, a five-year sample is taken, eg, a 2005 Q1 sample from January 1, 2005 to January 1, 2010. The stress period is taken from January 7, 2008 to January 7, 2009, so the last two estimations (2009 Q3 and 2009 Q4) are not part of it. Figure 2 shows that model-computed standard deviations and haircuts drop for 2008 Q3/Q4 samples, because the remainder of the five-year series lies mostly outside of the stress period.

The jump parameters are shown in figure 3. The most volatile period occurs in 2008, deep in the financial crisis. As we approach the end of the sample series, the intensities for both up and down jumps become weaker, and jump sizes also decrease (the average up-jump size is the reciprocal of η_u), as the index has a much less volatile period following Q3 2009. The estimated models' kurtoses are generally in line with those of the data samples, indicating a good fit of the tail. The overall behaviour is expected and relatively stable. As appropriate, once model risk is determined and quantified to be substantial, a haircut add-on could be levied to compensate for it. The results shown above do not suggest an immediate need for this, but it is good to remain vigilant.

As discussed earlier, being overly dependent on historical data is a source of model risk, eg, the subprime mortgage bonds. An alternative to estimated models is a model calibrated to options or other relevant markets. Table C shows haircuts under each of two calibrated MEM models for the SPX: the Cai-Kou model (Cai & Kou 2011) is calibrated to SPX European options, while CKL-2014 (Cai *et al* 2014) is weighted more towards SPX option smiles. The Cai-Kou model is in line with a 15% supervisory haircut, but the smile-leaning CKL-2014 model overshoots haircuts. This reflects the danger of exclusively calibrating to deep out-of-the-money put options, which are known to consist of significant liquidity premiums.

2 Estimated sample standard deviation and Aa2 haircuts change as the five-year SPX series rolls quarterly from January 1, 2005 to January 10, 2009

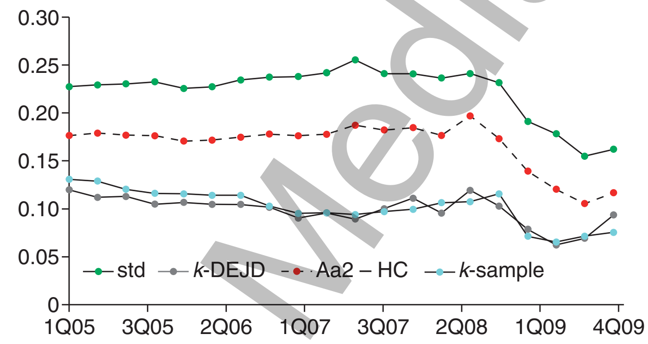

Also shown are sample kurtosis (k-sample) and model kurtosis (k-DEJD), scaled by 0.01

3 Estimated S&P 500 DEJD jump parameters between January 1, 2005 and January 10, 2009

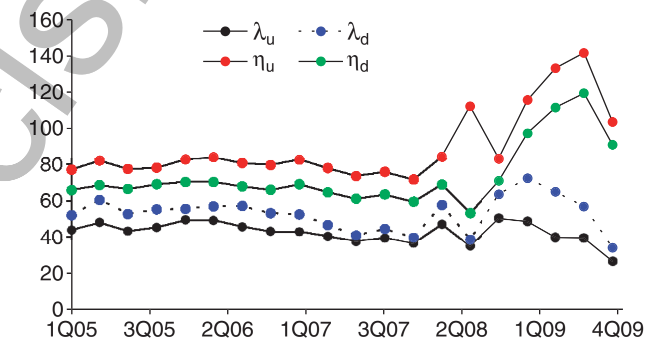

C. SPX haircuts under two implied risk-neutral DEJD models, compared with a raw haircut of 14.44% and a regulatory haircut of 15%

	Aa1	Aa2	Aa3
CKL-2014	34.48%	31.48%	28.66%
Cai-Kou	17.86%	15.93%	14.15%

Conclusion

A data-driven haircut approach such as VAR is subject to limitations in data availability, reliability and flexibility. In this article, we introduce a complementary parametric haircut model employing asset volatility and asymmetric jumps to conduct haircut sensitivity analysis, capture liquidity risk, adjust for idiosyncrasy and incorporate useful information from related markets. For example, a proxy data series fails to capture a specific bond's recent price behaviour and credit deterioration when it becomes evidently idiosyncratic. Our model allows us to stress volatility, downward jump magnitude or probability to quantify the impact.

The new model may find applications in haircut design for non-cash collateral, which is becoming increasingly accepted in CSAs and exchanges, or CCPs' clearing and margining agreements. It can be used as a candidate regulatory internal haircut model, following on from the FSB's strengthened regulatory haircut framework, which is expected to be implemented by the end of 2018. Preliminary results show that estimated DEJD models with a stress period are able to produce haircut levels consistent with collateral haircuts for equities, corporate bonds and securitised products, as typically seen in CSAs and Basel. Since there is no default history to study top (AAA) rated corporates' default frequency or loss rates, we recommend haircut designs targeting AA+ (Aa1) or AA (Aa2).

The model could be extended to study counterparty-dependent haircuts (eg, repo haircuts) with wrong-way risk taken into consideration. Application of the MEM as an extension of the DEJD model could be explored, with the view that its enhanced skew- and tail-capturing abilities might be needed for securitised debts where structural linkages exist. It would also be interesting to see stochastic volatility jump diffusion models' performances for haircut purposes.

Appendix: Laplace inversion procedures

Cai *et al* (2014) propose a two-sided Laplacian inversion algorithm. To solve for $f_{X(t)}$:

$$f(t) = f_{\mathrm{A}}(t,\sigma,C,N) + e_{\mathrm{T}}(t,\sigma,C,N) - e_{\mathrm{D}}(t,\sigma,C)$$

$$\begin{aligned} f_{\mathrm{A}}(t,\sigma,C,N) = {} & \frac{\exp(\sigma t)L_f(\sigma)}{2(|t|+C)} + \frac{\exp(\sigma t)}{|t|+C} \\ & \times \sum_{k=1}^{N}\left[(-1)^k \operatorname{Re}\left(\exp\left(-\frac{k\pi \mathrm{i} C\operatorname{sgn}(t)}{t+C\operatorname{sgn}(t)}\right)\right.\right. \\ & \left.\left.\times L_f\left(\sigma + \frac{k\pi\mathrm{i}}{t+C\operatorname{sgn}(t)}\right)\right)\right] \end{aligned} \tag{14}$$

where f_{A} provides an accurate approximation of $f(t)$ when the truncation error e_{T} and the discretisation error e_{D} are small. σ is a number in the ROAC, $C > 0$ is a shift constant to control e_{D} and $N > 0$ is the number of truncation terms to control e_{T}.

The truncation error e_{T} is bounded as:

$$|e_{\mathrm{T}}(t,\sigma,C,N)| \leqslant \frac{\varsigma(\sigma)\mathrm{e}^{\sigma t}}{\pi\xi\rho^{(1-\beta)/\xi}}\,\Gamma\left(\frac{(1-\beta)}{\xi}, \rho\left(\frac{\pi}{|t|+C}N\right)^{\xi}\right) \tag{15}$$

where $\Gamma(a,b)$ is the upper incomplete gamma function of order a and lower bound b. For the inversion of pdf Laplacian L_f, these parameters are listed below:

$$\beta = 0,\quad \rho = \tfrac{1}{2}\sigma_{\mathrm{a}}^2 t,\quad \xi = 2,\quad \zeta(\sigma) = \mathrm{e}^{tG(-\sigma)} \tag{16}$$

For cdf L_F, use $\beta = 1$. For European put options L_P, $\beta = 2$ and $\zeta(-\sigma-1)$ can be easily shown, following the derivation of Cai *et al* (2014). We propose a different set of parameters so the relative errors can be determined for all three inversions: (17), (18) and (19) for pdf, cdf and put, respectively:

$$\beta = 0,\quad \rho = \tfrac{1}{2}\sigma_{\mathrm{a}}^2 t,\quad \xi = 2,\quad \zeta(\sigma) = \mathrm{e}^{tG(-\sigma)} \tag{17}$$

$$\beta = 0,\quad \rho = \tfrac{1}{2}\sigma_{\mathrm{a}}^2 t,\quad \xi = 2,\quad \zeta(\sigma) = \frac{\mathrm{e}^{tG(-\sigma)}}{|\sigma|} \tag{18}$$

$$\beta = 0,\quad \rho = \tfrac{1}{2}\sigma_{\mathrm{a}}^2 t,\quad \xi = 2,\quad \zeta(\sigma) = \frac{\mathrm{e}^{tG(\sigma+1)}}{(\sigma+1)^2} \tag{19}$$

Noting that the $\varsigma(\sigma)\mathrm{e}^{\sigma t}$ term in (15) also appears in f_{A}'s leading term, while the rest of the variables in (15) do not depend on which Laplacian is being converted, the relative errors of e_{T} to f_{A} are therefore independent, so the same N can be used for all three inversions. ■

Wujiang Lou is a director in global fixed-income trading at HSBC in New York. The views and opinions expressed herein are exclusively the author's, not those of his employer or any of its affiliates.
The author wishes to thank Simon Juen and the reviewers for their helpful comments.
Email: wujiang.lou@us.hsbc.com.

REFERENCES

Bielecki T, 2008
Rating SME transactions
Report, May, Moody's Investors Service

Cai N and SG Kou, 2011
Option pricing under a mixed-exponential jump diffusion model
Management Science 57(11), pages 2067–2081

Cai N, SG Kou and ZJ Liu, 2014
A two-sided Laplace inversion algorithm with computable error bounds and its applications in financing engineering
Advances in Applied Probability 46, pages 766–789

Eraker B, M Johannes and N Polson, 2003
The impact of jumps in returns and volatility
Journal of Finance 53, pages 1269–1300

Financial Stability Board, 2015
Transforming shadow banking to resilient market-based finance: regulatory framework for haircuts on non-centrally cleared securities financing transactions
Report, FSB

Gorton G and A Metrick, 2012
Securitized banking and the run on repo
Journal of Financial Economics 104(3), pages 425–451

Kou SG, 2002
A jump-diffusion model for option pricing
Management Science 48(8), pages 1086–1101

Lou W, 2016a
Gap risk KVA and repo pricing
Risk November, pages 70–75

Lou W, 2016b
Repo haircuts and economic capital
Working Paper, February 8, SSRN, available at http://ssrn.com/abstract=2725633

Ramezani CA and Y Zeng, 2007
Maximum likelihood estimation of the double exponential jump-diffusion process
Annals of Finance 3(4), pages 487–507

Vazza D and NW Kraemer, 2015
Default, transition, and recovery: 2015 annual global corporate default study and rating transitions
Report, Standard & Poor's