MOODY'S ANALYTICS

MODELING METHODOLOGY

FEBRUARY 2021

Authors

Amnon Levy
Managing Director – Research

Kamal Kumar
Director – Research

Acknowledgements

The authors thank Andrew Kaplin, Stephanie Lee, Martha Sellers, Qiang Meng, Nihil Patel, Yaakov Tsaig, Yashan Wang, Tomer Yahalom, Jing Zhang, and Fanlin Zhu for their comments and suggestions.

Contact Us

Americas
+1.212.553.1658
clientservices@moodys.com

Europe
+44.20.7772.5454
clientservices.emea@moodys.com

Asia (Excluding Japan)
+85 2 2916 1121
clientservices.asia@moodys.com

Japan
+81 3 5408 4100
clientservices.japan@moodys.com

An Overview of Modeling Credit Portfolios

Abstract

This document provides a high-level overview of the modeling methodologies implemented in Moody's Analytics RiskFrontier™ and their business applications. To address the challenges faced by credit risk or credit portfolio managers, RiskFrontier models and calculates a credit investment's value at the analysis date, its value distribution at a user-specified investment horizon, as well as its marginal contribution to portfolio risk, i.e. the portfolio-referent risk, of every instrument in the portfolio. The approach is designed to explicitly analyze a wide range of credit investments and contingencies, including term loans with prepayment options and grid pricing, dynamic utilization in revolving lines of credit, bonds with put and call options, equities, credit default swaps, retail instruments, commercial real estate loans, and structured instruments. Institutions holding credit portfolios can utilize RiskFrontier and its outputs to help increase stakeholder value by measuring, managing, and increasing portfolio return/risk, while also ensuring capital adequacy and regulatory compliance.

Table of Contents

1. Introduction

All credit portfolios, whether global, regional, corporate, small and midsize enterprises (SME), commercial real estate (CRE), or retail, face uncertainty in losses. For example, Figure 1 depicts the delinquency rates for U.S. real estate, commercial and industrial (C&I), and consumer (retail) loans from 1987 to 2020. Looking at the figure, it is evident that variation in credit quality is substantial and default risk can increase quickly and in correlation with the macro-environment. Even in the aggregate, loss variability cannot be eliminated completely by diversification. Coupled with tight lending margins, even small miscalculations in risk and pricing can undermine profitability.[1]

With this in mind, credit risk can be substantially reduced through managed diversification. As investors adopt a diversification measurement, credit risk, and the rewards for bearing it, will ultimately be owned by those who can diversify it best. The challenge for every risk or portfolio manager is to measure and understand the economic risks in their portfolio, and ensure they are properly compensated. An active portfolio manager cannot do this without practical and conceptually sound methods for both measuring diversification and determining portfolio holdings to minimize risk concentrations and maximize return in credit portfolios.

Figure 1 Delinquency rates on U.S. loans.

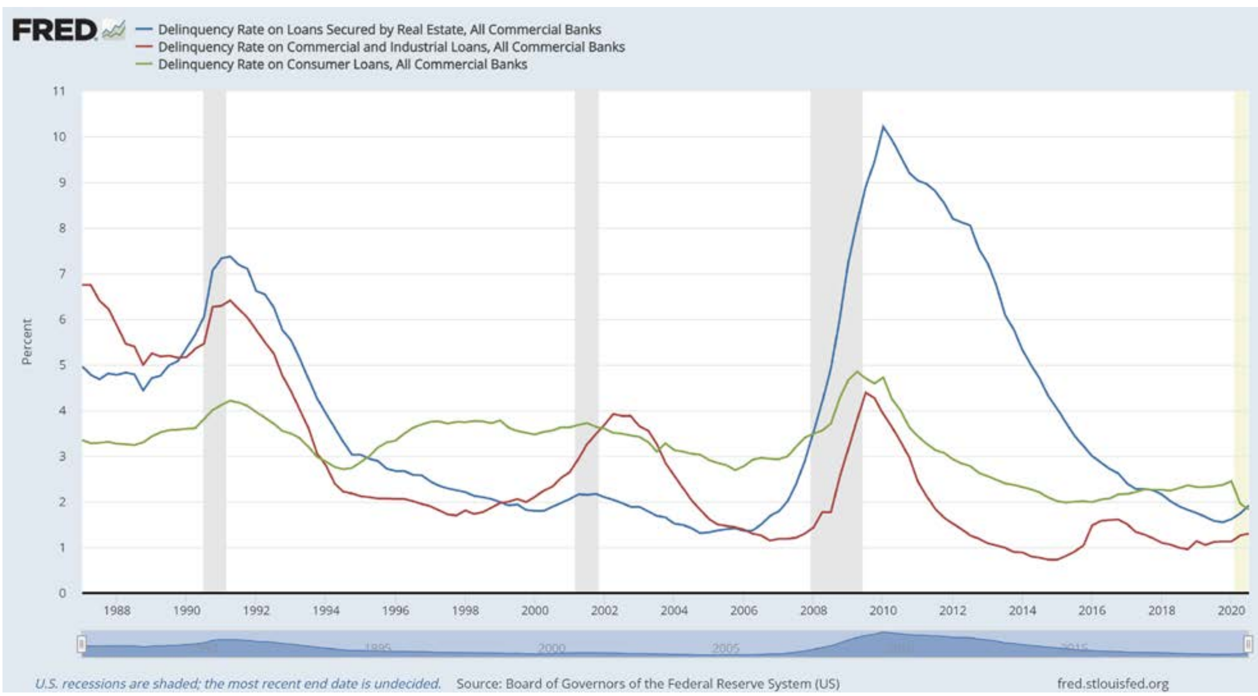

To address the challenges faced by credit risk or credit portfolio managers, RiskFrontier models each credit investment's value at the analysis date, its distribution of returns over an investment horizon, and the joint credit risk with all other instruments in the portfolio. The approach is designed to explicitly analyze a wide range of credit investments and contingencies, including term loans with prepayment options and grid pricing, dynamic utilization in revolving lines of credit, bonds with put and call options, equities, credit default swaps, and structured instruments. Moreover, RiskFrontier's open framework provides users with tremendous flexibility in specifying the valuation approach, migration model, and correlations. As such, RiskFrontier produces a quantitative set of actionable goals for portfolio management based on a granular model that can be tailored to the specific needs of an organization.

[1] For an additional discussion, see Kealhofer and Bohn (2001).

The portfolio model provides accurate economic estimates of portfolio risk measured as economic capital (i.e., Value-at-Risk, or VaR), unexpected loss (i.e., standard deviation, or UL), or expected shortfall (i.e., Conditional Value-at-Risk, or CVaR). It also provides accurate estimates of portfolio-referent risk (i.e. marginal contribution to portfolio risk) for each instrument in the portfolio, measured as risk contribution (i.e., marginal contribution to portfolio standard deviation) and tail risk contribution (i.e., marginal contribution to portfolio tail loss or VaR). The aforesaid portfolio and instrument risk measures account for concentration/diversification across dimensions such as country and industry, as well as reference name concentration. The model does so through a bottom-up approach that considers the following:

- Variation in instrument level inputs (e.g., holding amount, loss given default (LGD), tenure, contingencies, and fees).
- Variation in borrower level inputs (e.g. probability of default (PD), term structure of PD, size (sales or assets), industry, geography, and, if applicable, retail product or real estate property type).
- Variation in correlations across borrowers through a correlation model, such as the Moody's Analytics Global Correlation Model (GCorr™).[2] GCorr is a multi-factor correlation model — consisting of close to 1,000 geographical, sectoral, and national and regional macroeconomic factors — is updated and validated annually and is based on a long time series of empirical and granular data capturing intra- and inter-asset class correlations for a broad range of asset classes (public and private firms, CRE, retail, sovereign, and project finance).
- Using aforesaid instrument- and borrower-level inputs together with valuation, credit migration, GCorr, and Monte Carlo simulation models and framework, RiskFrontier calculates return, portfolio-referent risk, and return/risk (risk-adjusted performance) measures for each instrument in the portfolio, as well as for segments, sub-portfolios, and the overall portfolio. With this granular view, RiskFrontier supports portfolio- and instrument-level actions including pricing, approval, limit setting, position sizing, hedging, selling, and structuring. It facilitates the comparison of compensation for portfolio-referent credit risk against the amount of capital allocated to support a particular investment. It is important to note that return and portfolio-referent risk measures, such as risk-adjusted return on capital (RAROC), provide the portfolio manager with a rank ordering of the most beneficial and detrimental investments based on their marginal impact on portfolio return/risk. The analysis guides the manager in improving portfolio performance by recommending where to focus additional investments, and indicating where hedging is needed.
- In addition, RiskFrontier can help financial institutions assess and ensure capital adequacy from economic, accounting, and regulatory perspectives. Specifically, RiskFrontier is designed not only to measure required economic capital and quantify credit risk concentrations, but also to address requirements such as stress testing. More generally, RiskFrontier functionality addresses the Pillar 2 supervisory review processes.[3]

To summarize, the portfolio model behind RiskFrontier helps the risk manager or portfolio manager overcome the high- dimensional problem associated with measuring and managing portfolio-referent risk and return of the underlying investments in the portfolio. The remainder of this document provides an overview of how RiskFrontier addresses these challenges. This paper is organized as follows:

- Section 2 discusses the valuation methodologies.
- Section 3 explains how the risk and return of each individual investment is modeled within the portfolio.
- Section 4 provides an overview of the Global Correlation model (GCorr), as well as a discussion of other correlation models that can be utilized in RiskFrontier
- Section 5 explains the Monte Carlo methods that amalgamate instrument-level valuation and risk-and-

[2] For more information, see "Modeling Credit Correlations: An Overview of the Moody's Analytics GCorr Model" by Huang, Lanfranconi, Patel, and Pospisil (2012).
[3] For detailed discussions, see Gibbon (2007).

return analysis with correlations to build up the value/loss distribution for the portfolio, and provide actionable portfolio-referent risk measures for each instrument. This section also reviews several features associated with Monte Carlo analysis, including DealAnalyzer®, which allows for portfolio-referent risk-and-return calculations to be conducted rapidly for a new investment or hedge position without reanalyzing the entire credit portfolio.

» Section 6 describes RiskFrontier's key business applications.

2. Valuation

Valuation is fundamental to credit portfolio analysis. Given the lack of market prices for most credit instruments, an accurate model is essential. The valuation model used in marking a portfolio to market can have dramatic effects on the perceived portfolio values, as well as the rank-ordering of instruments' return and portfolio-referent risk. A portfolio or risk manager must ensure an economically consistent valuation approach that correctly handles the specifics of each instrument across the entire portfolio.

Within the context of RiskFrontier, there are two classes of valuation models: the valuation models used to analyze single-name instruments (e.g., loans or bonds), and models used for structured instruments (e.g., CDOs). We discuss horizon value distribution in Section 3.

2.1 Single-Name Credit Instruments

In the spirit of the open framework, RiskFrontier provides substantial modeling flexibility on the valuation front. Broadly speaking, users can value instruments using book-style methodologies, market-based measures (e.g., prices), or the Moody's Analytics lattice valuation method.[4] The approach that best fits an institution depends on its needs. In practice, institutions utilize differing approaches across their sub-portfolios if, for example, market prices are not consistently available for some sub-portfolio. Moreover, many institutions analyze their portfolios under multiple settings to get different views of their portfolio value, such as book value versus mark-to-market.

Under book value, the value of an instrument is par less amortized upfront fees. While book value provides a sense for the instrument value at origination (value is typically near par at origination), it does not adjust in any way for changes in the credit quality of the underlying reference entity. Book valuation can be very useful for evaluating the performance of an instrument's contractual pricing. It is also useful in that it can easily align with an institution's other systems that recognize only value changes resulting from cash losses.

If data are available, users can utilize market prices or credit curves (through spreads or matrix pricing) to value instruments. This option is useful for instruments that are either the same as or similar to instruments that are traded regularly. Although the application of prices is straightforward (value is equal to price), valuation using credit curves entails employing a reference rate to discount contractually promised cash flows. It is important to understand the caveat that applying credit curves from one market or instrument to another can lead to erroneous analysis. This is relevant when instruments have different terms and conditions that impact their respective spreads. As will be demonstrated below, this issue is addressed in the lattice structure, which explicitly models terms and conditions.

The lattice structure values an instrument using a bottom-up approach. Cash flows are modeled as a function of the contract as well as the state of a reference entity's credit quality. The structure lends itself to the explicit modeling of credit contingencies. Figure 2 provides a simplified visual illustration of the lattice structure. Credit states for the reference entity are represented on the vertical axis, and time is represented on the horizontal

[4] For a detailed discussion of the non-lattice valuation methodologies, see "Modeling Credit Portfolios." For additional details regarding the lattice models, see "Modeling Credit Portfolios," and Levy, Hu, and Li, (2007).

axis. A time-credit state combination is referred to as a "node." The probability of migrating from any one node to another at a subsequent time point is captured by the transition probabilities, which can be based on the Moody's Analytics Empirical Credit Migration model, or on a user-provided transition matrix (e.g., rating-based migration). The default state is an absorbing state, so all transition probabilities from the default state to non-default states are zero.

Figure 2 The lattice.

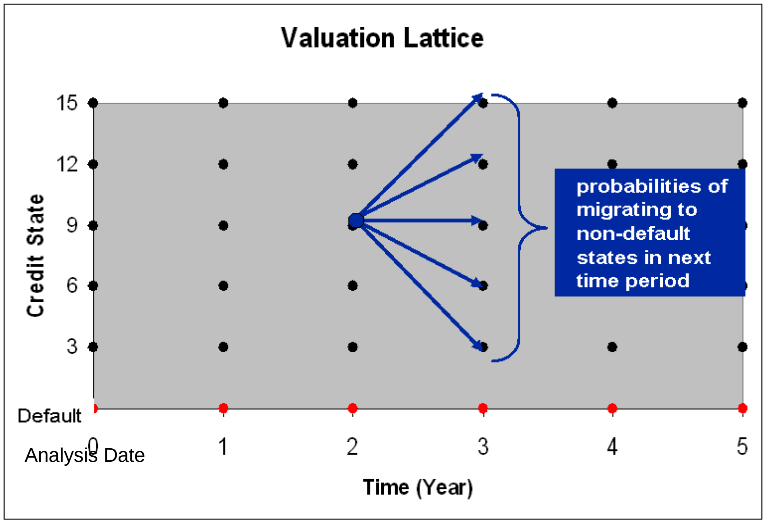

As Figure 3 demonstrates, the lattice allows for the modeling of credit qualities over time, credit migration probabilities, and credit-contingent cash flows. Given these, one can model contingencies that are a function of the reference entity's credit state. For example, the nodes for which a prepayment option on a loan is exercised can be determined by comparing the value of the instrument as an ongoing concern with the value to the borrower associated with prepayment.

Valuation within the lattice structure is conducted recursively backward in time. Starting at the second to last period, expected cash flows from the last period are discounted back to arrive at a continuation value at each node; for a plain vanilla loan there is no uncertainty in the value of the non-default nodes in the last period, so no calculation is necessary. Note that the default probability depends on the node; starting at a high node (i.e., a high credit quality) will result in a lower probability of default and higher value. The value will typically be lower for lower credit quality nodes. After adding any accumulated cash flows (e.g., interest payments, amortized principal, or prepayment penalties), the values for nodes in the preceding period are computed, taking into account contingencies, if any.

A few subtleties are worth pointing out. First, a separate lattice is constructed for each instrument. After all, contingencies and terms and conditions are instrument-specific, and so cash flows at each node are instrument-specific. Second, the probability of reaching a node is specific to the reference entity. This follows from the fact that migration probabilities depend on the default probability term structure for the reference entity; the default probability term structure can be different for each reference entity in the portfolio. Third, valuation requires discounting cash flows that face risks, such as default risk. Similar to the capital asset pricing model (CAPM), the lattice accounts for the systematic portion of risk in cash flows. However, instead of discounting cash flows at a beta-adjusted rate, the lattice discounts cash flows using risk-neutral valuation techniques.

The motivation for using risk-neutral valuation techniques is driven by their computational ease when applied to the problem at hand. These techniques were pioneered by Black and Scholes (1973), and Merton (1973) in their seminal work on option pricing. Merton (1974) applied the same analysis to the pricing of corporate liabilities, and created a structural framework that established the theoretical underpinning of subsequent developments in academic research and industrial application pioneered by Moody's Analytics. It is important to note that similar to the CAPM framework, each asset's covariance with the market determines its systematic risk adjustment; the risk-neutral valuation method includes an adjustment for each reference entity's systematic risk.

With the lattice in hand, one can analyze a wide range of dynamics associated with credit contingencies, since an arbitrary cash flow can be associated with each node. For example, one can model the common observation that deterioration in credit quality is associated with a drawdown of credit lines, as demonstrated in Figure 3. The lattice accurately accounts for differing usage and non-usage fees, and the impact of the payoff or drawdown on cash flows. In the same spirit, the RiskFrontier lattice models loans with prepayment options or grid pricing, callable or puttable bonds, CDSs (explicitly accounting for counterparty risk), custom instruments with described arbitrary cash flows, and exposure profiles that capture the credit component of instruments facing risks other than credit risk, such as counterparty risk of derivatives. The lattice also models equity, both public and private.

Figure 3 The lattice structure applied to modeling the dynamic usage pattern in a line of credit.

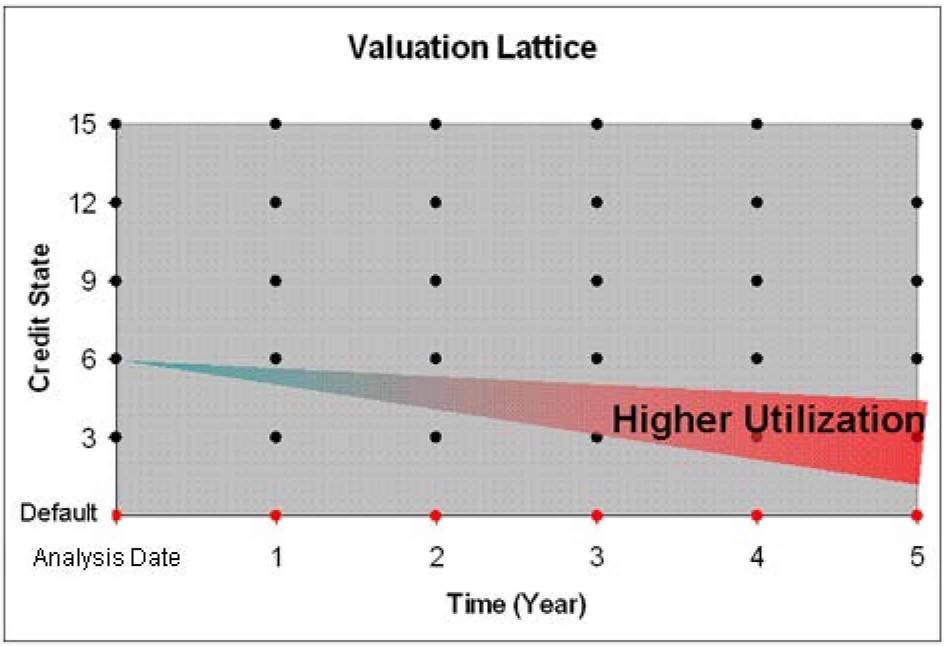

As mentioned, RiskFrontier allows for modeling of migration using the Moody's Analytics Empirical Credit Migration model, or through a user-specified credit migration matrix (e.g., a ratings-based migration matrix). The Moody's Analytics Empirical Credit Migration model uses information from the Moody's Analytics EDF model to determine the probability of migrating to different credit states. As demonstrated in Dwyer and Qu (2007), the EDF model provides the most accurate point-in-time measure of default probability. As such, an EDF-based migration model provides the most accurate description of how actual point-in-time credit quality (i.e., default probability) evolves. Alternatively, users can specify a Moody's rating transition matrix if the interest is to model migration using a through-the-cycle measure of credit quality.

To get a sense of the dynamics associated with different migration models, Figure 4 presents a Moody's one-year migration matrix (top left), as well as a Moody's Analytics EDF-based one-year migration matrix (bottom right). Although the Moody's Analytics matrix is not derived directly from the Moody's Analytics Empirical Credit Migration model, it does represent the kind of dynamics observed in the Moody's Analytics Empirical Credit Migration model. In particular, one can see that EDF credit measures are substantially more dynamic (or less sticky) than agency ratings; the probability of remaining within a Baa category is much higher under a ratings-based approach than the probability of remaining within an EDF-range that corresponds to a Baa rating.[5] This follows

[5] It is worth mentioning that the relationship between EDF measures and ratings can be non-monotonic. As such, this exercise is supposed to provide a sense for the dynamics; a formal comparison is much more difficult.

from the fact that EDF values measure credit quality at a point-in-time, while Moody's ratings are through-the-cycle measures, and do not change as frequently.[6]

Figure 4 A comparison of a Moody's rating transition matrix and a Moody's Analytics EDF-based transition matrix.

Moody's one-year migration matrix based on actual ratings changes

Initial Rating	Rating at End of One Year							
	Aaa	Aa	A	Baa	Ba	B	Caa	Default
Aaa	91.56	7.73	0.69	0.00	0.02	0.00	0.00	0.00
Aa	0.86	91.43	7.33	0.29	0.06	0.02	0.00	0.01
A	0.06	2.64	91.48	5.14	0.53	0.10	0.02	0.02
Baa	0.05	0.22	5.16	88.70	4.60	0.84	0.26	0.19
Ba	0.01	0.07	0.52	6.17	83.10	8.25	0.63	1.26
B	0.01	0.05	0.19	0.41	6.27	81.65	5.92	5.50
Caa	0.00	0.04	0.04	0.25	0.79	10.49	69.91	18.47

Source: Moody's Investor Services, 2006*

MKMV migration matrix:

- More granular (30 states)
- More dynamic
- Based on a very large data pool
- Multiple horizons
- Adjusted for borrower's PD term structure

MKMV one-year migration matrix based on non-overlapping EDF ranges

Equivalent Rating	Equivalent Rating at End of One Year							
	Aaa	Aa	A	Baa	Ba	B	Caa	Default
Aaa	65.50	20.10	8.50	3.90	1.20	0.80	0.10	0.00
Aa	25.53	36.57	25.80	9.20	2.17	0.50	0.10	0.07
A	4.10	18.03	41.03	27.17	7.77	1.63	0.17	0.07
Baa	0.43	2.87	19.17	46.27	24.87	5.47	0.67	0.30
Ba	0.10	0.17	2.33	21.27	49.77	22.97	2.70	0.60
B	0.00	0.00	0.20	2.00	18.50	55.70	21.40	2.17
Caa	0.00	0.00	0.00	0.30	1.80	20.60	63.60	13.60

*Withdrawn Ratings are redistributed proportionally.

Source: Moody's KMV

2.2 Structured Instruments

The challenges of modeling structured instruments are numerous. Much of the complexity stems from the dependence of a structure's cash flows on a collateral pool whose credit qualities are heterogeneous. Understanding the combination of events and associated probabilities that lead to possible cash flow realizations can be computationally costly. For example, using a lattice structure to value a CDO with 100 reference entities in the collateral pool would require constructing a grid with 100 dimensions (i.e., one for each reference entity). This would be impractical. To avoid these costs, RiskFrontier does not utilize the lattice structure to value structured instruments. Instead, RiskFrontier makes simplifying assumptions and applies an analytic valuation method.

The analytic valuation approach is similar to the approaches presented in Gregory and Laurent (2003), and Hull and White (2004). The methodology relies on two fundamental modeling assumptions. First, the collateral terms and conditions, and correlations are assumed to be uniform across the collateral pool.[7] Second, losses can be described through a pass-through waterfall structure. More specifically, losses impact a subordinated class (e.g., a tranche) when losses on the collateral pool exceed the attachment point. The tranche will be worthless after losses on the collateral pool exceed the detachment point. However, the structured instrument can be parameterized to model CDOs, BDSs, CLOs, and asset-backed securities of various types, with different waterfall structures.

Figure 5 provides a simplified visual representation of the analytic approach. The first step is to build up the loss distribution for the underlying collateral pool at each point in time, as depicted at the top of the figure. Cash flows

[6] This implies that default rates associated with an agency rating varies over time. For example, Baa-rated firms experienced a 10-fold increase in default rates from 1998 to 2001.

[7] It is worth noting that the simplified correlation structure is only applied to valuation. As we will discuss shortly, risk and return calculation incorporate the rich correlation structure that is available in RiskFrontier.

to the tranche can be constructed by mapping the distribution of losses on the collateral pool to losses on the tranche, which are based on the attachment point (a) and detachment point (d), as depicted in the lower right-hand part of the figure.

Tranche value is computed by discounting risk-neutral expected tranche cash flows back to the analysis date.

Figure 5 The analytic valuation approach to modeling structured instruments.

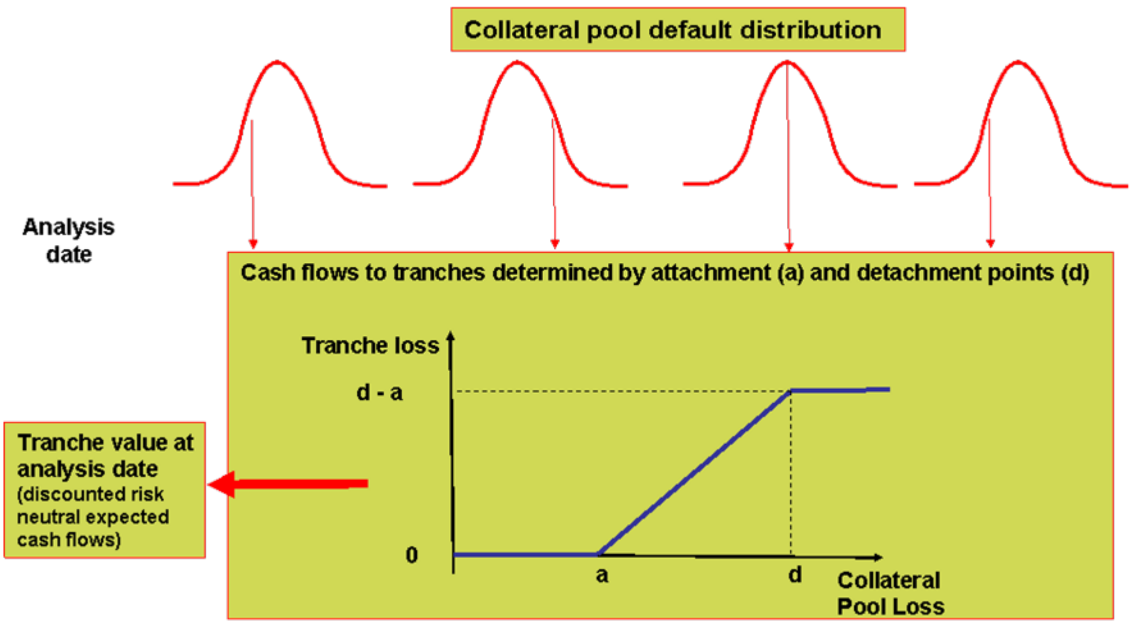

RiskFrontier allows users to offset distortions associated with the two fundamental modeling assumptions related to the analytic valuation method. First, distortions associated with collateral pool homogeneity are offset using information specific to the deal, as well as available market information. For example, the correlation parameter that is used for pricing can be calibrated to those implied by market prices or an external model.[8] It is important to note that the semi-analytic approach used for analyzing risk and return does not rely on homogeneity. For example, it leverages the chosen correlation model (e.g., GCorr) along with Monte Carlo techniques and does not have the same shortcomings. Second, users can also offset distortions associated with a pass-through waterfall structure. The attachment point can be calibrated so that the distress probability of the structured instrument corresponds to that produced by a model that accounts for the more complex waterfall structure. Similarly, the detachment point can be calibrated to the loss given distress on the structured instrument. This approach to setting attachment and detachment points allows users to leverage a richer stand-alone valuation tool in a computationally more efficient framework that accounts for portfolio concentration effects.

Moody's Structured Analytics (SA)-RiskFrontier (RF) data feed: Moody's Structured Analytics (SA) team maintains a database of publicly traded CLO, CMBS, RMBS, and ABS transaction data, a library of complex cash flow rules, and simulation models for assessing tranche level risk. SA-RF data feed automation (i) translates SA transaction data to RiskFrontier format, (ii) performs complex calculations, including the aforesaid calibration of attachment and detachment points, based on RiskFrontier analysis outputs to parameterize a CDO (or equivalent bond/loan) instrument to match tranche level risk based on SA simulations and models, and then (iii) loads relevant data and imports portfolios of synthetic CDO equivalents into RiskFrontier. This permits analysis of more complex structures,

[8] For a more detailed discussion, see "Modeling Credit Portfolios."

such as RMBS, CMBS, ABS, or cash CDO with complex waterfalls, in a portfolio context utilizing the RiskFrontier CDO module.

3. Measuring the Stand-Alone Instrument Risk and Return

We now discuss how the horizon value distribution is constructed and used, along with the valuation model, to measure an instrument's stand-alone risk and return. Similar to valuation, the model that defines the horizon value distribution is fundamental to credit portfolio management. The instrument's horizon value distribution, when combined with the analysis date value, defines the risk and return of the instrument. As in Section 2, "Valuation," the horizon valuation models used to analyze single-name instruments (e.g., loans or bonds) are different from those that are used for structured instruments (e.g., CDOs). As such, this section is divided into two parts.

3.1 Single-name Credit Instruments

RiskFrontier provides substantial flexibility when modeling the value distribution at horizon. As with analysis-date value, users can model the distribution incorporating book-style methodologies, default-no-default settings, market-based measures (e.g., prices or spreads), or utilizing the Moody's Analytics lattice structure.

Similar in spirit to book value, RiskFrontier allows users to define the expected horizon value of an instrument as an ongoing concern (i.e., conditional on not having defaulted) using a linear or exponential pull-to-par approach. As with book valuation, horizon linear and exponential valuation do not account for changes in the credit quality of the underlying reference entity.

As with analysis date valuation, market prices or credit curves (e.g., spreads or matrix pricing) can be used to describe the horizon value of an instrument as an ongoing concern. In the case of credit curves, the user also specifies the distribution of spreads at horizon. In other words, the Moody's Analytics Empirical Credit Migration model is not used in the analysis. Instead, each horizon credit state is associated with a forward credit curve, which, along with the reference rate, is used to discount the remaining promised cash flows. Here, the practitioner must be aware of the same caveats as in analysis date valuation: applying the distribution of spreads from one market to another can lead to erroneous analysis when instruments have different terms and conditions.

Finally, at horizon, the user can compute the expected value and value distribution of an instrument using the lattice methodology. Figure 6 provides a simplified visual representation of how the horizon value distribution is constructed. In the example, the analysis is conducted to a horizon of two years. The value of the instrument at each of the non-default nodes at horizon is a combination of post- and pre-horizon cash flows. These cash flows include not only coupons or fees, but also any cash flow-associated contingencies, such as prepayment.

Using the lattice, valuation of cash flows after horizon is conducted using risk-neutral valuation techniques as discussed in Section 2.1, "Single-name Credit Instruments," (i.e., backward induction). In this example, a higher credit state will typically be associated with a higher value, because the probability of default is lower. Cash flows before horizon are assumed to be reinvested at the credit-risk-free rate and aggregated at horizon (i.e., forward). It is worth pointing out that the different possible paths to a node are considered when aggregating these cash flows.[9] The single ***default*** node at horizon actually represents a set of states. That is, in the event of default, recovery value is uncertain and is modeled through a beta distribution with parameters from Moody's Analytics LossCalc™, or a user-specified mean and variance.[10]

[9] It is assumed that all cash flows realized before horizon are reinvested to horizon at the credit default risk-free rate.

[10] In fact, users can specify an entire term structure for the expected recovery and variance in RiskFrontier. This can be particularly useful when modeling project finance where recovery is typically lower at the initial phases of the project.

To arrive at the instrument value distribution in the non-default states, it is necessary to associate a value with each horizon node, as well as a probability of arriving at that node. The probability of arriving at each horizon node is computed by concatenating the credit migration probability distribution (i.e. the migration matrix) at each time-step before horizon. In Figure 6, there are two time steps before horizon. As such, the probability of arriving at any horizon node is computed by considering the different possible paths that lead to this node. Of course, the default node is an absorbing state. As discussed in Section 2.1, the probability of arriving at a node is determined by the chosen credit migration model (e.g., the Moody's Analytics Empirical Credit Migration model).

It is important to stress that for calculating value/loss distribution, return, and risk at the investment horizon, the relevant distribution of states at horizon is the distribution observed in the physical world (i.e., the distribution under the physical measure). For example, for the purpose of risk management, it is relevant to focus on actual probabilities of default and probabilities of the counterparty improving or deteriorating in credit quality. This physical distribution is distinctly different from the distribution under the risk-neutral measure, which is relevant when valuing risky future cash flows. In other words, the relevant probability of defaulting or migrating to a credit state at horizon is the physical probability. Once that credit state is realized, the value of the instrument can be computed using risk-neutral valuation techniques that combine future physical probabilities and the investors' risk aversion, as expressed in market prices.

With the horizon value distribution in the non-default and default states, and the analysis date value in hand, the computation of risk and return is straightforward. In particular, Unexpected Loss is computed as the standard deviation of the return, and the expected spread, which measures compensation for risk, is computed as the expected return less the credit-risk-free rate.

Figure 6 How an instrument's value distribution at horizon is generated.

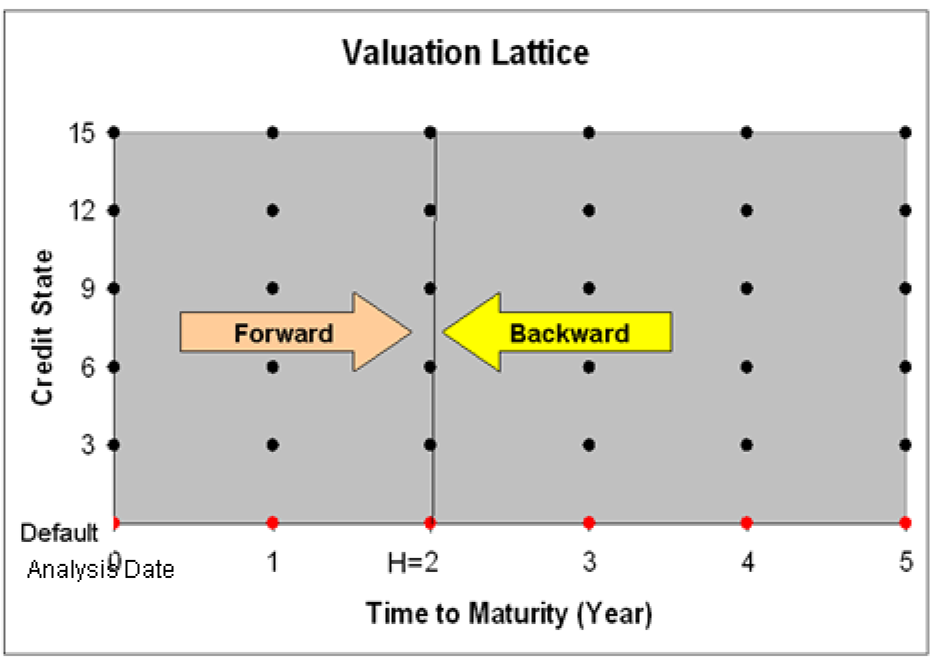

3.2 Structured Instruments

As discussed in Section 2.1, "Single-name Credit Instruments," *valuation* of a structured instrument is complex and time-consuming. When compared with single-name instruments, the complexity associated with modeling the *horizon value distribution* is much more for a structured instrument. One challenge is to keep track of the horizon credit states of all reference entities in the collateral pool, as they play a role in the horizon value of the structured instrument.

Another challenge is to associate a probability to each combination of reference entity credit states. This is coupled with the desired property that the model properly account for credit events associated with reference entities, which overlap the collateral pool of the structured instrument and other parts of the portfolio.

To overcome these challenges, RiskFrontier uses a semi-analytic approach. This approach combines a Monte Carlo simulation, which accurately represents cash losses prior to horizon, with the analytic valuation method (described in Section 2.2, "Structured Instruments"), which represents the horizon value of the structured instrument as a going concern. The analytic horizon value is pre-computed for each point of a two-dimensional grid: one dimension is the cumulative loss on the collateral pool to horizon; the other dimension is the average credit quality of the remaining principal (average forward 1-year default probability). The structured instrument's horizon value distribution is constructed by simulating each reference entity's horizon credit state in accordance with other parts of the portfolio. The value from cash flows prior to horizon is then combined with the forward-analytic value from the grid to arrive at a value at horizon.

It is important to stress that, within the context of the horizon value distribution, homogeneity is only assumed when computing the analytic value of the structured instrument as a going concern at each point on the grid. Meanwhile, the simulation accurately accounts for heterogeneity in default probability, the likelihood of multiple defaults as implied from the specified correlation model, loss given default, time to maturity, and notional. Thus, risk and return metrics incorporate the heterogeneous characteristics of the collateral pool. In addition, RiskFrontier is uniquely able to capture the risk concentration due to reference entities that are common to the collateral pool and other parts of the portfolio.

The last step of the analysis corrects for inconsistencies between assumptions (e.g., correlation) used in constructing the horizon value distribution, and assumptions used in valuation at the analysis date. To rectify this discrepancy, the simulated horizon values are shifted by a constant. This constant is set to guarantee that the discounted expected value (using risk-neutral valuation techniques) is equal to the value at the analysis date.

Finally, it is worth pointing out that stand-alone risk and return statistics for structured instruments are simulation-based, unlike those for single-name instruments. Unexpected loss is computed as the standard deviation of the return, and the expected spread is computed as the expected return less the credit-risk-free rate.

4. Correlation

A credit portfolio or structured product is composed of multiple instruments. Clearly, the overall risk of a portfolio depends not only on the risks of the individual instruments, but also on how the future values of these investments are correlated. Given the importance of correlation in determining the risk of a portfolio, Moody's Analytics has developed a rich global factor model (GCorr Corporate) that provides pairwise asset (not equity) correlations for approximately 33,000 publicly traded firms, the Private Firm R-squared Calculator that provides mid-market correlations, as well as correlation parameterizations for a wide range of asset classes, such as retail instruments, CRE, and asset-backed securities. In addition to the Moody's Analytics models and data, RiskFrontier provides substantial flexibility by allowing specification of an arbitrary factor model or pairwise correlation matrix for reference entities in the portfolio, where the joint distribution can be described through a Normal copula. To facilitate in parameterization, the following three structures are available:[11, 12]

» Flexible Factor Model. It is an expansion of the Moody's Analytics factor model (GCorr), in which additional factors (e.g., retail, SME, or CRE factors) are included. This allows users to refine the correlation model associated with the extended asset classes, while leveraging the richness of GCorr in a single unified correlation framework.

» General Factor Model. It allows users to specify an arbitrary linear factor model. The factors can be economic in nature, or represent latent factors. The general factor model framework requires users to specify the covariance matrix between the factors.

» Pairwise Correlation Model. It provides the most flexibility, because users specify the correlation between every pair of reference entities in the portfolio.

Regardless of the specified correlation structure, RiskFrontier utilizes a bottom-up approach where asset correlations describe the co-movement in borrowers' credit states, which, together with borrower's PD and instrument characteristics, impact the co-movement in instrument values that are part of the output. Parameterized correlations can represent the correlation of credit states to be used to describe the co-movement in forward default probabilities as represented in the lattice structure and the Moody's Analytics Empirical Credit Migration model, or, alternatively, the co-movement in credit curves. The details of how the joint movement of credit states is combined with the instrument horizon value distributions to describe the joint distribution of instrument values can be found in the Monte Carlo section that follows.

The remainder of this section provides an overview of the GCorr Corporate model, which uses a factor model rather than direct historical observations to measure asset return correlations between firms. The factor model approach leverages information in areas where data are plentiful and allows for inference in areas where data are scarce and noisy. Zeng and Zhang (2001) use extensive empirical data to show that historical correlations are subject to a large amount of sampling error, thereby limiting their usefulness in predicting future correlations. The predictive power of a factor model, in contrast, stems largely from its control over these errors. A good follow-up paper is Zhang, Zhu, and Lee (2008), which provides a link between the predicted default correlations from the GCorr Corporate model and correlations implied directly from realized corporate defaults.

Some practitioners have expressed valid concerns about measuring default correlations since they are difficult to estimate; defaults are infrequent and typically occur once, at most, for any entity. The approach described below cleverly addresses this dynamic aspect of default correlations by decomposing default (and value) correlations into two component pieces: asset (business value) correlation, which is relatively less dynamic, and credit quality (PD), which is enormously dynamic and requires forward-looking measures and constant monitoring. The description below is about the estimation of firm-level asset correlations. The RiskFrontier user can then put these together with highly predictive credit quality measures and get a useful, predictive, and actionable view of portfolio risk.

[11] For a detailed discussion of the correlation module, see "Modeling Credit Portfolios."

[12] The Private Firm R-squared Calculator can be used to estimate asset correlation between private firms and between private firms and other assets in the credit portfolio.

A factor model approach imposes a structure on the correlation of asset returns. The correlation between the asset returns of any pair of firms can be explained by the firms' relationships to a set of common factors. In particular, GCorr Corporate can be thought of as having a factor model within a factor model. To begin, a firm's credit risk can be decomposed into a systematic component (i.e., the firm's composite factor) and a firm-specific (i.e., idiosyncratic) component; this can be thought of as the first factor model. The systematic component can be described by a second factor structure. In particular, a firm's composite factor is defined by the firm's industry and country composition. The dynamics are driven by industry- and country-specific loadings on a set of global, regional, and industrial sector factors that are common to all countries and industries. Furthermore, each industry and country has a specific (i.e. idiosyncratic) risk component that is unrelated to other industry- and country-specific risks. Thus, the GCorr Corporate asset correlation between two firms is defined by the covariance between the two respective composite factors, along with the correlation between each firm and its own composite factor. More formally, the variables that define the GCorr Corporate correlation structure are the following:

- The relation between a firm and its composite factor (this is referred to as R-squared or RSQ).
- A firm's industry and country composition.
- The loadings on the common global, regional, and industrial sectors that define the correlation between industry-country combinations.

Figure 7 provides a visual representation of the Moody's Analytics GCorr Corporate factor structure. Starting from the top, a firm's risk is decomposed into systematic and firm-specific risk; this decomposition is determined by the firm's R- squared. The systematic risk is defined by the firm's industry and country composition (i.e., a composite factor). As demonstrated in the bottom row, dynamics for each industry and country are defined by loadings on the global, regional, and industrial sectors. The remaining industry and country risks are specific to the respective industry and country.

Figure 7 The GCorr Corporate Factor Structure.

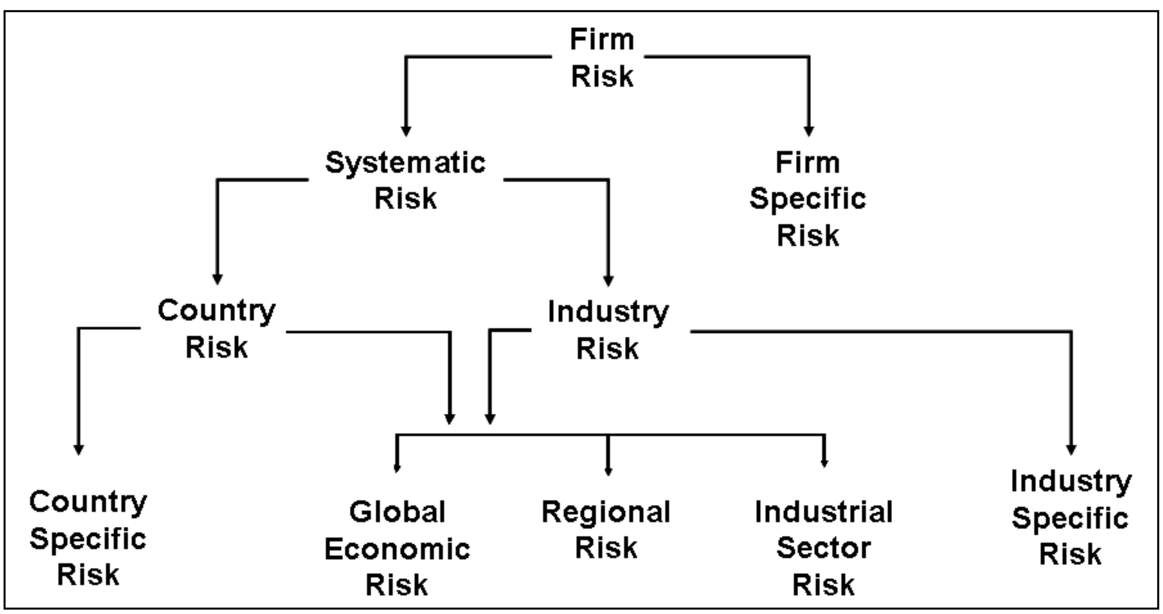

5. Bringing it All Together with Monte Carlo Simulation

So far, this document has provided an overview of valuation models, models describing horizon value distribution at the instrument level, and the correlation module (along with GCorr). This section describes how these components are used in the Monte Carlo engine to construct the portfolio value distribution and to calculate portfolio-referent risk statistics.

As discussed above, the problem of analytically computing the joint distribution of returns for a credit portfolio is a daunting task. Instead of describing the joint distribution analytically, the Monte Carlo engine samples from the joint distribution of credit states as defined by the correlation model. The engine then utilizes the mapping between credit state and instrument value (available through the instrument horizon value distribution) to sample from the joint distribution of instrument values. Said another way, the Monte Carlo engine simulates the credit states of the world for the underlying reference entities in the portfolio, and then maps credit states to corresponding values for each instrument. By simulating a large number of trials, where each trial represents one combination of credit states for all the reference entities, one can obtain a description of the joint distribution of instrument values, and, as a result, the portfolio value (or loss) distribution.

Figure 8 provides a visual representation of the simulation for two firms. Focusing on the graph in the top right, the x-axis represents the credit state for reference entity X, and the y-axis represents the credit state for reference entity Y. Ovals on the graph represent iso-probability curves for the joint distribution of credit states. Curves in the center have a higher probability mass and those further out have less, indicating that extreme credit states are less likely. The fact that the ovals are squeezed toward the top right and bottom left indicates that the credit states are correlated; reference entity X is more likely to realize a bad credit state if reference entity Y realizes a bad credit state (at horizon).

To aid with the visual representation, the marginal distribution of credit states for reference entities X and Y are plotted along the x-axis and y-axis, respectively. The shaded portion of the left x-axis and bottom y-axis represents the default states for reference entities X and Y, respectively. Meanwhile, the value of each instrument, as a function of credit state, is depicted on the top left and bottom right. In each trial, the Monte Carlo engine samples from the joint distribution of credit states. This is done by simulating the underlying factors as defined by the correlation model, associating the factor loadings for each reference entity to arrive at the systematic portion of the reference entity's risk. In addition, an idiosyncratic shock is simulated for each firm to arrive at a credit state for each firm. If the realized credit state falls below the default threshold, a random recovery amount is simulated. If the reference entity is not in default, the credit state is mapped onto a going-concern instrument value.

Computing portfolio and portfolio-referent risk statistics is straightforward after a large sample of instrument values is simulated. For portfolio-level statistics, instrument values can be aggregated for each trial to arrive at a sample of portfolio values. Unexpected Loss (i.e., standard deviation or UL) for the portfolio can be estimated by taking the standard deviation of the sample. Similarly, portfolio capital can be computed by looking at the loss threshold associated with a particular target probability. Meanwhile, portfolio-referent risk statistics, i.e. an instrument's marginal contribution to portfolio risk, can be computed by relating the instrument value with the portfolio value. The covariance between the instrument value and the portfolio value can be estimated for the purpose of computing instrument-level risk contribution i.e., instrument's marginal contribution to portfolio standard deviation (i.e. portfolio UL).

Meanwhile, Tail-Risk Contribution (i.e., marginal contribution to tail or extreme portfolio loss) is computed by keeping track of the loss on each instrument when the portfolio falls within the user-specified (tail) loss interval.

Having laid out the mechanics of how the portfolio and portfolio-referent risk statistics are computed, it is worth considering the economics of when it is relevant to focus on each of the measures. In particular, an institution's preferences should determine which statistics are most relevant. For example, an institution whose primary focus is to minimize the year-to-year changes in portfolio value should probably focus on portfolio UL and each instrument's marginal contribution to portfolio UL (i.e., risk contribution). If, instead, an institution's primary focus is to minimize the likelihood of incurring extreme losses (e.g., to protect their Aaa rating), then the focus should probably be on capital measured at a low target probability, and each instrument's marginal contribution to those extreme outcomes (i.e., Tail- Risk Contribution). In some cases, multiple stakeholders may have different preferences. For example, an institution's shareholders may care about portfolio volatility or UL, while regulators may be concerned with the likelihood of extreme loss. One approach is for the institution to focus on a low target probability capital measure at the portfolio level, and use risk contribution when allocating risk at the instrument level. Loosely speaking, this approach allows the institution to rank order instruments and optimize its portfolio based on instruments' marginal impact on portfolio volatility or UL, while assessing and ensuring capital adequacy based on regulatory or economic capital requirements.

Figure 8 The Monte Carlo engine.

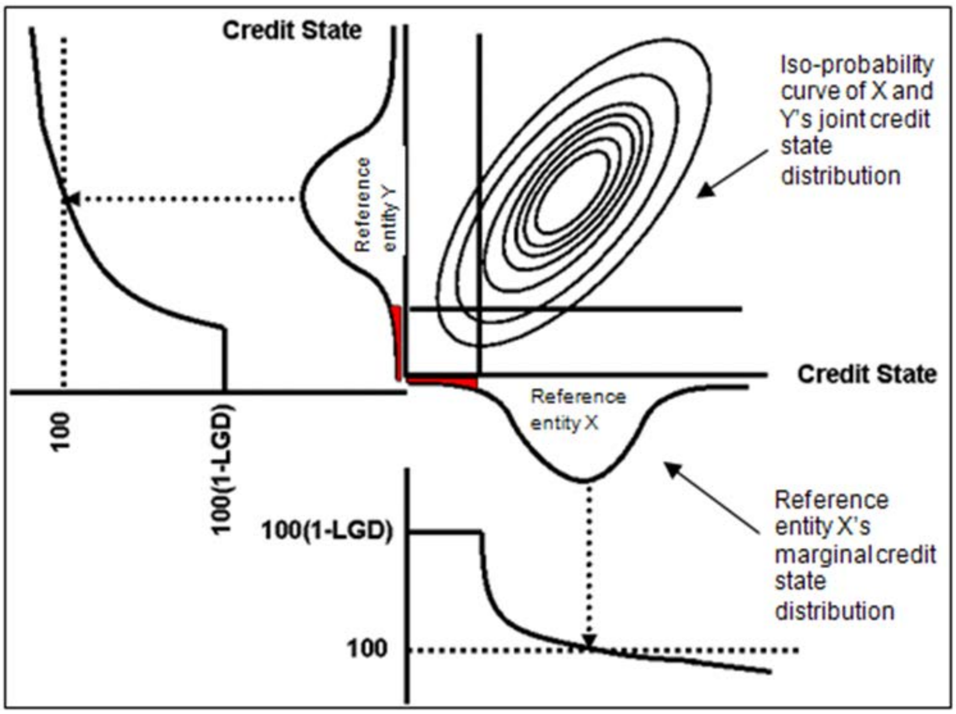

The following four features related to the Monte Carlo engine are worth highlighting:

- Simulation Services distributes the Monte Carlo calculations across multiple processors and machines. This allows for an improvement in computation speed, as well as an ability to analyze portfolios with a large number of instruments. The engine utilizes a Multi-Stream Monte Carlo approach that controls the sequence of pseudo-random numbers in such a way that ensures that results are not affected by whether the calculations are performed in a distributed fashion.
- An importance-sampling approach to simulation is available in RiskFrontier. The methodology improves upon the simulation speed without loss of accuracy. The basic idea behind importance sampling is to change the distribution from which the random samples are taken in a way that preserves the statistic; asymptotically, as the number of trials approaches infinity, importance sampling will produce the same simulated statistics as standard Monte Carlo. A greater concentration of the sample is chosen from the region with the greatest impact on the calculation. For computing risk statistics, this means sampling those scenarios that lead to large losses. Although performance is portfolio-dependent, the approach frequently reduces the number of simulations by a factor of ten or more (i.e., one-tenth the number of trials are needed to achieve the same level of accuracy).
- RiskFrontier facilitates the analysis of large homogeneous pools of instruments. The functionality allows for accurate analysis of a homogeneous pool of arbitrary size in a fraction of the time it would otherwise take to analyze each instrument individually; it takes about as much time to

analyze the pool as it does a stand-alone instrument. This functionality is useful when, for example, modeling retail pools (e.g., a pool of credit card loans), where the benefits of granularly analyzing characteristics of each borrower and loan may not be sufficient to offset the increased computation time and data requirements. When modeling instruments using the homogeneous pool functionality, the model is designed to compute simulated statistics with approximately the same level of accuracy as when the instruments were modeled individually.

» DealAnalyzer is a calculation engine that brings portfolio analytics to the desktops of loan originators and active managers. DealAnalyzer provides the ability to structure, value, and price transactions accurately in light of existing credit portfolio risk. DealAnalyzer leverages the Multi-Stream Monte Carlo methodology to analyze the return and portfolio-referent risk characteristics of a deal without re-analyzing (re-simulating) the entire portfolio. Intuitively, Multi-Stream Monte Carlo allows for the timely regeneration of random sequences simulated for the existing portfolio, so that simulated values from the deal can be appropriately associated with values from the existing portfolio's value distribution, allowing for the computation of risk statistics.

6. Business Applications

Credit portfolios, consisting of loans, bonds, and other credit assets, make up a large portion of all assets on most financial institutions' balance sheets. Hence, credit portfolio value, return, and risk are key drivers of shareholder/stakeholder value, return, and risk. Institutions can increase stakeholder value by increasing portfolio return/risk while also ensuring capital adequacy and regulatory compliance. This process, for the most part, requires five key steps:

Key Steps and Questions

1. Measuring portfolio risk and ensuring capital adequacy
   a. Loss distribution: How much capital is required to absorb possible portfolio losses?
   b. Stress testing: What are the expected losses and capital requirements under stress scenarios?
   c. Reverse stress testing: What scenarios, borrowers, sectors/regions contribute most to large losses?
   d. Capital planning: What actions will be taken to maintain capital adequacy as stress scenarios unfold?
   e. Risk appetite: What are the target levels (risk appetite) and limits (risk tolerances) for risk-taking?
2. Quantifying each exposure's marginal impact on portfolio risk
   a. Risk Contribution: What is the increase in portfolio risk for a $1 increase in exposure value?
   b. Risk concentrations: Which borrowers, sectors, and regions impact portfolio risk the most?
3. Setting limits to prevent excessive risk concentrations
   a. Segment limits: What is the notional limit given a limit on each segment's Risk Contribution?
   b. Borrower limits: What is the notional limit given a limit on borrower's Risk Contribution?
4. Increasing/reducing credit exposure to increase portfolio return/risk
   a. Asset selection: Which exposures should be increased/decreased to increase portfolio return/risk?
   b. Asset selection and capital requirements: How should assets be selected given capital requirements?
5. Pricing and sizing new deals to increase portfolio return/risk
   a. Risk-based pricing: What is the deal's Risk Contribution and Sharpe Ratio, i.e., deal's marginal impact on portfolio return/risk, taking account of deal's notional, standalone risk, and correlation with the portfolio?

Financial institutions can perform the five steps and their sub-steps and address questions listed above using RiskFrontier functionalities and outputs, as described next.

6.1 Measuring Portfolio Risk and Ensuring Capital Adequacy

Loss distribution: How much capital is required to absorb possible portfolio losses?

RiskFrontier's portfolio overview functionality outputs portfolio level risk metrics, including (i) loss quantile (aka Credit Value at Risk (CVaR) or Economic Capital ("EC")), (ii) unexpected loss, i.e., standard deviation or volatility of portfolio losses at a future horizon, and (iii) portfolio value and loss distributions — all of which inform and support an institution's capital adequacy assessment, capital planning and management, risk strategy, risk appetite, and risk tolerance, and are increasingly being reported and communicated both internally and externally.

Stress testing: What are the expected losses and capital requirements under stress scenarios?

RiskFrontier's conditional distribution functionality uses a Monte Carlo simulation engine, together with the GCorr framework, to capture systematic factors' (e.g. country and industry factors) and macroeconomic factors' (e.g. GDP, unemployment rate, etc.) impacts on borrower credit quality and the resulting exposure and portfolio losses, and produce outputs, including portfolio loss distribution, expected loss (EL), unexpected loss (UL), and economic capital (EC) conditional on user-specified stress scenarios. GCorr model coverage continues to expand and currently includes close to 1,000 total factors, including about 250 national, regional, and international macroeconomic variables.

Reverse stress testing: What scenarios, borrowers, sectors/regions contribute most to large losses?

In RiskFrontier with GCorr Macro module, macroeconomic variables do not replace GCorr systematic credit risk factors during simulation and in modeling portfolio risk, but are simulated in correlation and along with credit risk factors, thus facilitating reverse stress testing. Specifically, RiskFrontier's Monte Carlo simulation output file for an unconditional simulation can be used to identify scenarios, factors, borrowers, sectors, or regions associated with any target level of portfolio loss.

Capital planning: What actions will be taken to maintain capital adequacy as stress scenarios unfold?

Previously described loss distribution, stress testing, and reverse stress testing results from RiskFrontier can be used to:

- Project capital demand and capital supply and assess capital adequacy under stress scenarios
- Plan and determine management actions (e.g. capital increase or risk/exposure reduction) that will be taken to prevent and to recover from possible capital deficiencies, if and as stress scenarios unfold
- Use such capital planning as part of internal capital adequacy assessment process and discussion with regulators, supervisors, and stakeholders

Risk appetite: What are the target levels (risk appetite) and limits (risk tolerances) for risk-taking?

Previously described portfolio risk and capital adequacy assessment results and metrics from RiskFrontier can be used to estimate capital demand and supply and compare them with market opportunity and institution's risk strategy to inform and determine:

- Risk capacity, i.e., the maximum amount of risk an institution can support in pursuit of its business objectives.
- Risk appetite[13], i.e., the target amount or range of risk an institution is willing to accept in pursuit of its business objectives; risk appetite is lower than risk capacity to allow for a safety margin.
- Risk tolerance and risk limits, i.e., boundaries of risk-taking within which an institution is willing to operate; risk appetite is articulated more qualitatively and at aggregate level, while risk tolerances and limits are more quantitative and granular, e.g. borrower and segment limits.

6.2 Quantifying Each Exposure's Marginal Impact on Portfolio Risk

Risk contribution: What is the increase in portfolio risk for a $1 increase in exposure value?

Exposure's Risk Contribution ("RC_i"), defined as the increase in portfolio risk (i.e., standard deviation of portfolio value at horizon) for a $1 increase in exposure's value, can be shown to equal the exposure's standalone risk per unit exposure value, i.e., standard deviation of exposure value at horizon per unit exposure value (aka "unexpected loss" or "UL_i"), multiplied by the correlation between exposure value and portfolio value at horizon ("$CORR_{i,p}$").

$$Risk\ Contribution\ (RC_i) = Increase\ in\ portfolio\ risk\ for\ a\ \$1\ increase\ in\ exposure\ value = CORR_{i,p^*}UL_i$$

Risk concentrations: Which borrowers, sectors, and regions impact portfolio risk the most?

RiskFrontier calculates Risk Contribution, taking into account each exposure's notional, standalone risk per unit exposure value (driven mostly by PD, LGD, and maturity), as well as correlation with other exposures in the portfolio, regardless of the source of correlation or risk concentration, e.g. borrower, sector, or region, thus eliminating the need for any concentration add-on (relative to an assumed benchmark for a supposedly well-diversified portfolio). While there is no consensus on a definition of concentration, RiskFrontier's Risk Contribution output can be used to identify risk concentrations and exposures (or sectors) that could cause large losses to the portfolio, and decide where to add/reduce exposure to increase diversification and reduce risk concentration and portfolio risk.

[13] For more information, see "Quantifying Risk Appetite in Limit Setting" by Kaplin, Levy, Meng, and Pospisil (2015).

6.3 Setting Segment or Borrower Limits to Prevent Excessive Risk Concentrations

RiskFrontier's portfolio-referent instrument-level marginal return/risk measures, calculated as expected return divided by Risk Contribution (RC), can be used not only to increase portfolio return/risk via asset selection, as described in the next section, but also to monitor instrument-level marginal return/risk and prevent excessive risk concentrations by rejecting deals or reducing exposures that have high RC and low marginal return/risk and that would decrease portfolio return/risk. Doing so requires reliable tools and metrics and unflinching discipline.

As a practical, conservative alternative and/or supplementary approach, and as a backstop to the aforementioned marginal return/risk-based monitoring and limit setting, most institutions also calculate and enforce absolute notional limits in order to prevent excessive risk concentrations and to keep risk levels consistent with risk appetite and risk strategy. Limits based on notional exposure are a good start, but they have significant limitations:

- They do not account for borrower- or segment-specific risk.
- They do not account for the impact of borrower's or segment's correlation and concentration relative to the portfolio.
- They are more qualitative and subjective, and thus prone to subjective bias, and not easily defendable.
- They involve a manual process and are difficult to update frequently.
- They are not dynamic; they generally remain fixed despite changes to portfolio holdings and profile.

RiskFrontier users can overcome these limitations by using a quantitative and risk-sensitive approach. Specifically, RiskFrontier's limit setting utility uses analytic formulas, together with Monte Carlo simulation results, to calculate the relationship ("Risk vs. Notional Curve") between a borrower's or segment's "Notional Weight" and the corresponding "Risk Contribution Weight" for varying levels of borrower's or segment's Notional Weight, assuming no change to the rest of portfolio. Risk Contribution Weight is defined as the borrower's or segment's (analysis date) value-weighted Risk Contribution as a fraction of the sum of (analysis date) the value-weighted Risk Contribution of all exposures in the portfolio. Given a set of user-specified sector or borrower Risk Contribution Weight limits and using the aforementioned Risk vs. Notional Curve, RiskFrontier limit setting utility calculates the corresponding set of risk-adjusted notional limits for each segment or borrower, assuming no change to the rest of the portfolio. Users can specify the sector Risk Contribution Weight limits to be the same or different across segments.

6.4 Increasing/Reducing Credit Exposure to Increase Portfolio Return/Risk

Asset selection: Which exposures should be increased/decreased to increase portfolio return/risk?

It can be shown that an exposure's (standalone) expected return equals that exposure's marginal contribution to portfolio expected return, regardless of exposure notional and portfolio characteristics. Thus, an exposure's Sharpe Ratio, defined as "expected return / Risk Contribution," is a marginal measure quantifying exposure's marginal impact on portfolio return/risk and can be used to decide which exposures to increase/decrease in order to increase portfolio return/risk. Portfolio return/risk can be increased by increasing exposures with Sharpe Ratios higher than the portfolio Sharpe Ratio and decreasing exposures with Sharpe Ratios lower than portfolio Sharpe Ratio.

Portfolio Sharpe Ratio changes with the portfolio and can be used as a dynamic hurdle rate, without the need to compare exposure's Sharpe Ratio with an exogenously imposed hurdle rate, which often is static and not portfolio-sensitive, despite changes in the portfolio. Furthermore, using an exogenously imposed hurdle rate is superfluous (relative to using portfolio Sharpe Ratio as the hurdle rate) and may lead to asset selection that can actually decrease portfolio Sharpe Ratio. RiskFrontier calculates and reports each exposure's expected return and Sharpe Ratio, together with Risk Contribution, correlation with portfolio, and other risk measures discussed previously, and which can be used to selectively increase/decrease exposures and increase portfolio return/risk.

Asset selection and capital requirements: How should assets be selected given capital requirements?

If the financial institution's regulatory or economic capital requirements appear to constrain the aforesaid asset selection, the institution can overcome such apparent constraints by reducing portfolio size by scaling down (decreasing) all current exposures in the portfolio, or by raising additional capital at market rates. In either case and to the extent there is no gain/loss relative to market price, the institution can continue to increase portfolio return/risk by rebalancing portfolio using the previously marginal return/risk contributions, while simultaneously and independently meeting regulatory and/or economic capital requirements.

6.5 Pricing and Sizing New Deals to Increase Portfolio Return/Risk

Risk-based pricing: What is the deal's Risk Contribution and Sharpe Ratio, i.e., deal's impact on portfolio return/risk, taking account of deal's notional, standalone risk, and correlation with portfolio?

The prior section focused on how RiskFrontier outputs can be used to decide what to buy, sell, or hold to increase portfolio return/risk. What is equally important is to do so at the right price and with the correct size and terms. Sizing and structuring deals appropriately become even more important when the institution is a price-taker rather than a price-setter. RiskFrontier's Trade and DealAnalyzer modules enable institutions to quickly calculate return and risk measures for a new deal or a basket of new deals, considering the deal's standalone risk (driven mostly by PD, LGD, and maturity), as well as portfolio correlation and risk concentrations. Leveraging key metrics from the original portfolio analysis, RiskFrontier can quickly (within seconds and at the point of origination) calculate marginal return and risk statistics, including expected return (or credit earnings), Risk Contribution (RC), Sharpe Ratio, and incremental portfolio economic capital for new deals. Alternately, given RC for a deal, one can calculate the pricing required to meet or exceed portfolio Sharpe Ratio or any user-specified hurdle rate. Users can also calculate return/risk measures by substituting RiskFrontier's expected return (numerator) with a user-specified expected return, which should be a marginal measure, in order to remain consistent with Risk Contribution (denominator), which is marginal.

References

Black, Fischer, and Myron Scholes, "The Pricing of Options and Corporate Liabilities." *Journal of Political Economy*, Vol 81, No 3, 637– 659, 1973.

Dwyer, Douglas and Shisheng Qu, "EDF 8.0 Model Enhancements." Moody's KMV White Paper, 2007.

Gibbon, John, "RiskFrontier and Basel II, Pillar 2." Moody's KMV Whitepaper, 2007.

Huang, Jimmy, et al. "Modeling Credit Portfolios." Moody's Analytics Whitepaper, 2015.