Moody's ANALYTICS

MODELING METHODOLOGY

Modeling the Joint Credit-Interest Rate Dynamics on a Multi-Dimensional Lattice Platform: Model Validation and Applications in Risk Integration

Authors
Sunny Kanugo
Yanping Pan
Rama Sankisa
Yashan Wang

Contact Us

Americas
+1.212.553.1653
clientservices@moodys.com

Europe
+44.20.7772.5454
clientservices.emea@moodys.com

Asia-Pacific (Excluding Japan)
+85 2 3551 3077
clientservices.asia@moodys.com

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

Abstract

This document presents validation results for the credit-interest lattice or the multi-dimensional lattice (MDL) valuation model within Moody's Analytics RiskFrontier™. We focus on valuations of a large sample of corporate bonds, January 2006 – July 2013. We also produce valuations using the credit-only lattice and compare performance of the two lattice models.

We find that model valuations compare extremely well with market transaction prices. Further, the MDL model produces better valuations for high credit quality bonds, especially in higher interest rate regimes. These findings validate the model's ability to accurately value risky assets while accounting for both credit and interest rate risks.

We also compare the bottom-up approach for risk integration implemented in the joint lattice model with a traditional top-down approach. Our analysis shows how the joint lattice model provides more accurate risk measures for portfolios sensitive to both credit and interest rate risks. Validation results are robust across different interest rate environments as well as across instrument characteristics.

Table of Contents

1. Introduction

Since the late 1990s, the corporate bond market has seen tremendous growth, with outstanding debt increasing almost five-fold during the past 15 years.[1] Market transparency has also improved a great deal, especially after the introduction of TRACE by the NASD.[2] Figure 1 shows the outstanding debt and the average trading volume of corporate bonds. Outstanding debt has increased at an average annual rate of approximately 10% from 1996 – 2013. Trading volume has increased gradually since 2011, though it fluctuated around the 2007 – 2008 financial crisis and dropped in 2008 and 2011.

Figure 1   Outstanding Balance and Average Daily Trading Volume for the U.S. Corporate Bond Market
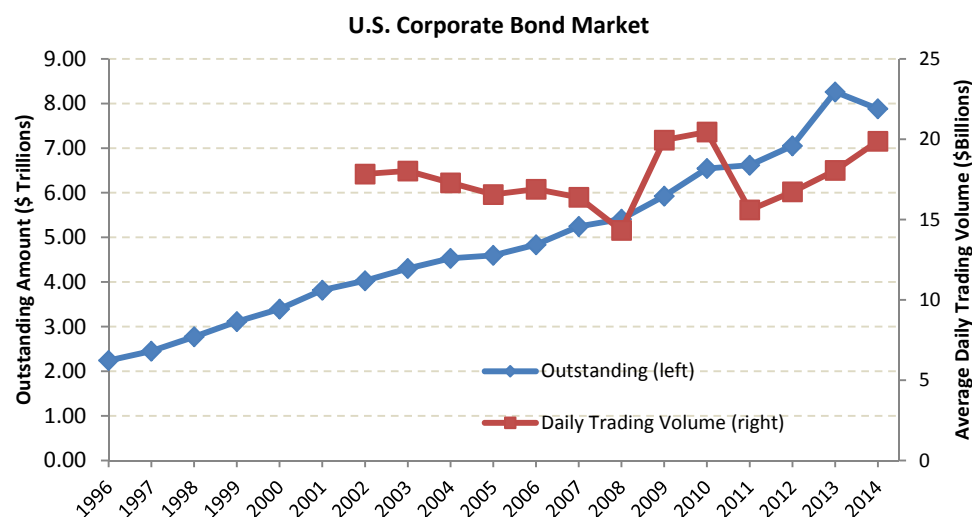

Figure 2 shows a steady growth in the issuance of both vanilla and callable bonds. Almost 90% of callable bonds have fixed coupon, whereas, for the vanilla bonds, the same proportion is approximately 60%. Further comparing callable bond and vanilla bond issuance, we observe an interesting phenomenon. Prior to the crisis, there were approximately 60% vanilla and 40% callable bonds issued. However, after the crisis, the issuance of callable bonds jumped to almost 60%, mainly driven by a large increase in the issuance of fixed rate callable bonds and a decrease in floating rate vanilla bonds.

---

[1] http://www.sifma.org/research/statistics.aspx

[2] Trade Reporting and Compliance Engine (TRACE) was introduced by NASD (presently FINRA) in 2002 to increase price transparency in the U.S. corporate debt market. All transactions executed in the OTC corporate bond market and any secondary market transactions of a TRACE eligible security have to be reported through the TRACE system.

Figure 2   New Issuance for the U.S. Corporate Bond Market

Fixed Rate Coupon


Floating Rate Coupon
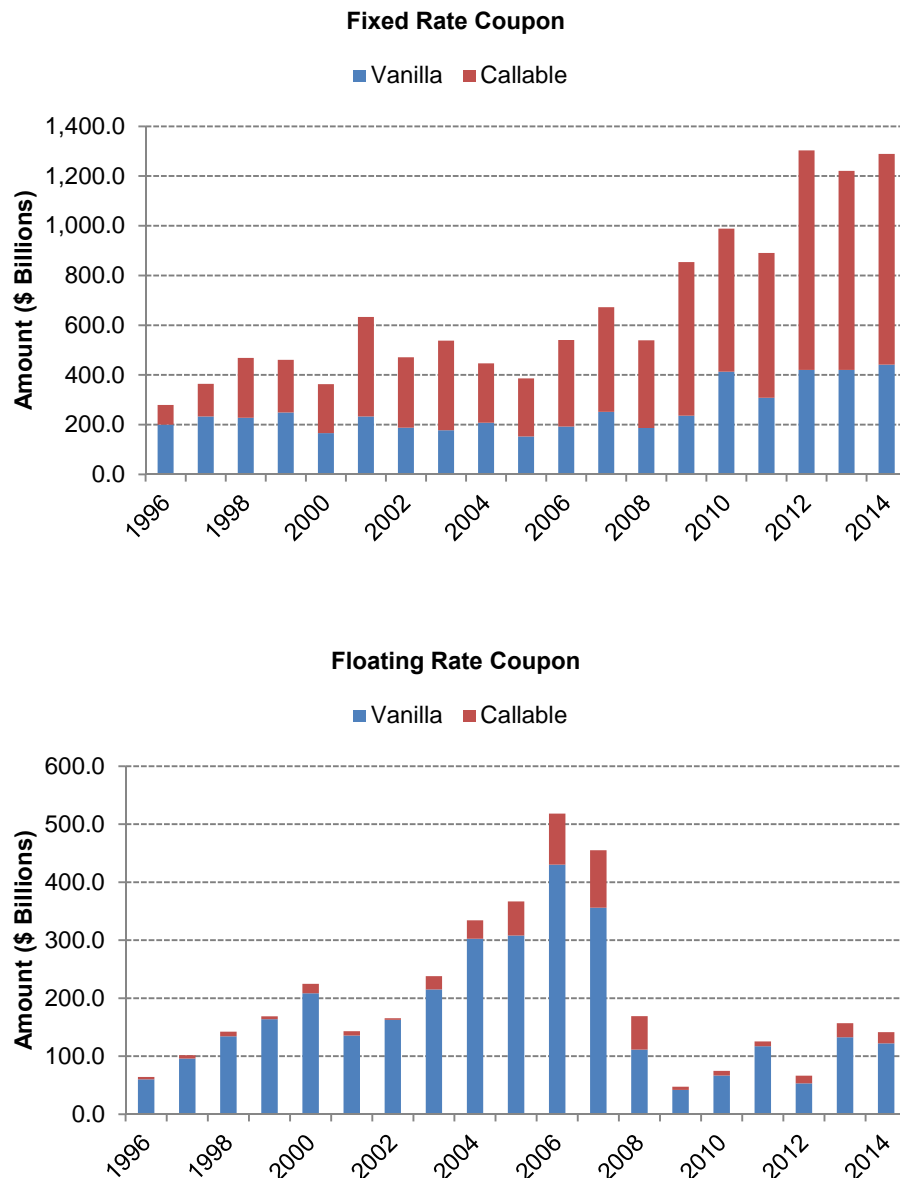
In terms of risk characteristics, bonds differ from other instruments such as loans. First, most bonds are fixed-rate instruments. Consequently, they are sensitive to changes in both credit quality and interest rates. Second, in addition to call options, equivalent to prepayment options for loans, bonds may include put options or options that convert them into equity positions. For loans, an improvement in issuer credit quality is the principal driver of prepayment decisions, whereas, for bonds, the joint movement of credit and interest is required to properly capture the inherent risks. For this reason, the valuation and risk assessment of bonds, either using credit migration alone or using option-adjusted spreads based on a stochastic interest rates model, would likely yield inaccurate results.

More generally, credit and interest rate risks are among the most important risks for a financial institution driving its losses. While various advanced models have been developed to measure these two risks in isolation, practices and techniques in risk aggregation are usually much less sophisticated. Financial institutions generally take the approach of assessing individual risk components first and then proceed to aggregate these components up to the level. Typical aggregation approaches, often called top-down approaches, include simple summation, with some fixed diversification percentages and a copula-based approach where, for

example, a variance-covariance matrix is used to aggregate broad risk types.[3] While these top-down approaches may be sufficient to produce reasonable integrated risk measures for certain types of portfolios where one risk clearly dominates, they lack the granularity needed to provide a clear picture of the scope and depth of the risks and the interaction between different types of risks when the portfolio is exposed to more than one source of risk.

In contrast to these top-down approaches, the credit-interest lattice or the multi-dimensional lattice (MDL) in RiskFrontier 4.0 implements a bottom-up approach to evaluate losses accounting for both credit and interest rate risk at the instrument level. In particular, the impact of the options and contingencies common in most assets and liabilities can be accounted for properly using this bottom-up approach. For example, consider the case of a fixed-rate callable bond. The decision to exercise the call option depends on the interest rate environment as well as the credit quality of the issuer. The issuer of a callable bond exercises the call option when its credit quality improves or when the market interest rate drops, or some combination of these two factors. A system such as that shown in Figure 3 is needed to determine when the bond is called i.e., at low interest rates, or good credit states, or a combination of the two. A top-down approach would generally have difficulty modeling this type of behavior in a consistent way.

Figure 3   Credit-Interest Lattice
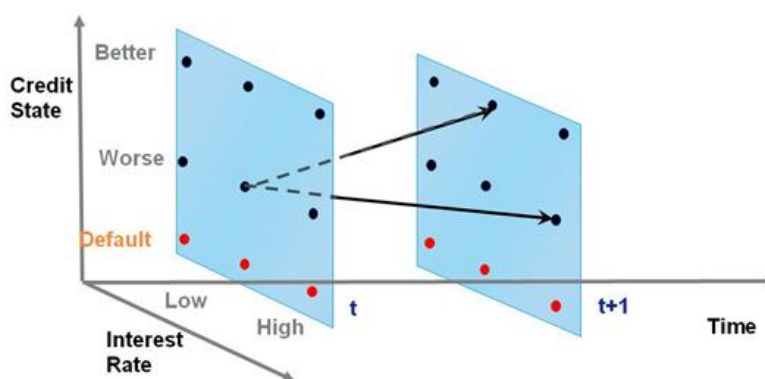
The purpose of this paper is to provide evidence on the valuation performance of the MDL model, as well as demonstrate how the bottom-up approach for risk integration implemented in the MDL outperforms a traditional top-down approach in accurately assessing portfolio risks.

We use the joint lattice model to value a comprehensive sample of corporate bonds and then compare model prices with actual transaction prices. The data used for this purpose consists of roughly 5,500 fixed rate corporate bonds observed from 1Q 2006 through 2Q 2013. Here we don't include any loans since majority of them tend to be of floating type thus insensitive to interest rate movements. Our analysis shows that the joint lattice model produces valuations that match reasonably well with prices observed in the market. The results are significantly better when CDS spread-implied risk measures are used and the sample consists of fairly liquid bonds.

Next, we use the joint lattice model to compute risks measures for portfolios of vanilla and callable bonds and then compare them with those of a typical top-down approach. Results show that, depending on the portfolio, the top-down approach can either overstate or understate the overall portfolio risk. Further, the error magnitude depends on interest rate levels and the correlation between credit and interest rate risks.

The remainder of this paper is organized as follows:

» Section 2 provides an overview of the theoretical framework for the joint modeling of credit and interest rates within RiskFrontier and the validation methodology

» Section 3 describes validation and different input parameters data.

» Section 4 discusses the valuation and risk integration results for different data segments.

» Section 5 summarizes the study and provides concluding remarks.

[3] For more information, see "Range of practices and issues in economic capital frameworks" and "Developments in Modeling Risk Aggregation" by the Basel Committee on Banking Supervision, March 2009 and October 2010. Papers by Chen, et al. (2010) and Pospisil, et al. (2013) discuss top-down approaches and provide examples.
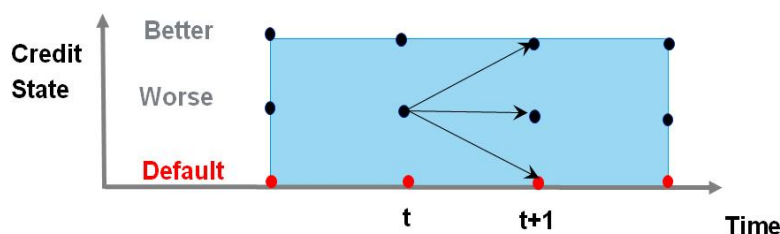
2. Validation Methodology

In this section, we provide an overview of the lattice framework within RiskFrontier and the methodology we use to validate the MDL model.

2.1 Credit-Only Lattice

Figure 4 illustrates the credit lattice structure used in RiskFrontier.

Figure 4   RiskFrontier Credit Lattice Structure

The credit-only lattice assumes that interest rates are deterministic. The vertical axis represents credit states for the reference entity, and the horizontal axis represents time. Each node on the lattice represents the credit quality of the reference entity at the corresponding time. The probability of migrating from one node to another node at a subsequent time point is captured by a set of transition probabilities, which can be derived from the Moody's Analytics Distance-to-Default (DD) dynamics model or from a user-provided migration model. The instrument value at the analysis date is calculated using a risk-neutral backward valuation process that begins at maturity and traverses the lattice, one step back at a time, computing the value at each node on the lattice as it progresses. In addition, the lattice model calculates the value grids that characterize the distribution of an instrument's value at the specified horizon using a forward-valuation process.

The forward process considers possible paths leading from the initial credit quality of the reference entity at the date of analysis to some credit state at the horizon date. The expected value of all possible paths leading to a credit state at horizon is combined with the risk-neutral value of cash flows beyond horizon to associate a value with each credit state at the horizon date. The horizon value distribution may then be used to conduct a risk-return analysis of a credit portfolio. The lattice model can explicitly model the options and credit contingencies, such as call and put options of floating rate corporate bonds, prepayment options of bank loans, and dynamic usage schedules of revolvers.

2.2 Credit-Interest Lattice

The credit-interest lattice is a multi-dimensional lattice (MDL) designed to jointly model the dynamics of credit risk and interest rate risk. As described in the white paper "Estimating Parameters of the Joint Interest Rate and Credit Model in RiskFrontier," the joint lattice model is constructed by combining the credit-only lattice with the interest rate lattice through a copula for which the correlation parameter is estimated in the Moody's GCorr Macro framework. On the joint lattice, movements of short rate level and credit quality are monitored and serve as the state of the world in assessing the economic value of future cash flows, along with decisions on contingencies such as call or put options. Instrument values at the analysis date and value of distributions at horizon are calculated for each instrument through the backward- and forward-valuation process similar to those on the credit-only lattice. At the portfolio level, interest rate shocks are simulated along with credit factors to produce instrument- and portfolio- value distributions, as well as risk-return statistics.
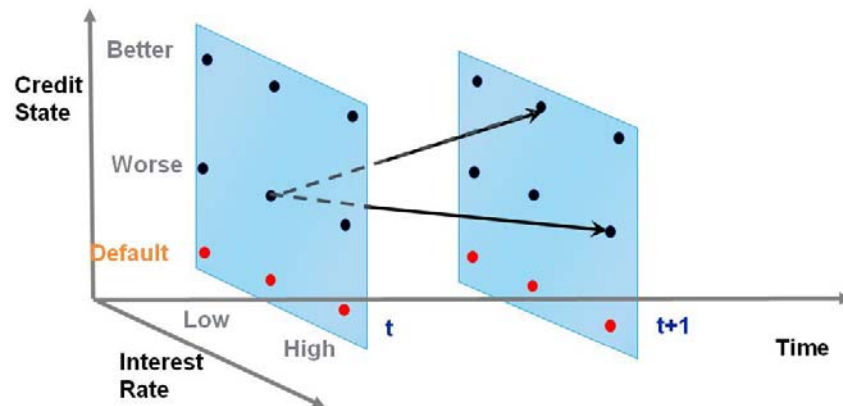
Figure 5   Credit-Interest Lattice

Figure 5 shows the joint lattice structure used in RiskFrontier. At a given time t, the lattice nodes are represented by $S_t = (CQ_t, IR_t)$, where CQ_t and IR_t correspond to the credit and interest rate lattice nodes at time t. We must also find the period-to-period transition probability $P(S_{t_1}, S_{t_2})$. Given the marginal transition probabilities along the credit dimension and along the interest rate dimension, and the correlation between credit and interest rate movements, the joint transition probabilities can be calculated on the joint lattice via a Gaussian copula.

MODELING EMBEDDED OPTIONALITY IN FIXED RATE INSTRUMENTS

Fixed rate instruments such as corporate bonds are often issued with flexible funding arrangements and may include one or more options, such as call options, put options, or an option to convert the bond into an equity position. In the case of call options, the debt issuer retains the right to repay the debt prematurely at a predetermined date and price, whereas, for put options, the lender reserves the right to terminate the contract prematurely. Once we have the joint credit-interest lattice with contractual principal and coupon cash flows laid out at each lattice node as per instrument's terms and conditions, we can use backward induction to calculate the instrument value at the analysis date and evaluate the embedded options at each step. We can also calculate the value distribution of the instrument at horizon on the joint lattice by combining the cash flows received before horizon and the future cash flow values at horizon.

2.3 An Overview of the Validation Methodology

We compare the two lattice models to quantify the additional value the credit-interest lattice provides over the credit-only lattice. In our approach, we compare the two lattice models by studying how closely the prices produced by these models track the market price benchmark of these bonds. Since the value of vanilla bonds is not affected by stochastic interest rates, we analyze the pricing errors only for callable bonds using the two lattice models. Further, we use only the high credit quality callable bond dataset for this analysis. The primary reason for this approach is that the transaction prices of the high credit quality callable bonds are relatively less volatile when compared to those of the low credit quality callable bonds. Also, the illiquidity costs on the higher credit quality callable bonds are much lower than those on the lower credit quality bonds.

We construct the market price benchmarks using the TRACE data. For most part of the analysis, we use the closing price benchmark. In Section 4.1, we describe how we construct this benchmark, and, in Section 6.1, we elucidate the motivation for constructing various alternative market price benchmarks.

3. Data

In this section, we describe the validation data and input parameters.

3.1 Bond Data

Bond data comes from FINRA's TRACE (Trade Reporting And Compliance Engine) and Reuter's EJV databases. The TRACE database provides end-of-the-day prices and other bond attributes such as coupon rate, issue date, maturity date, etc. We use the EJV database to download EJV model prices as well as call schedules for callable bonds. We apply the following filters to the data:

- » Include corporate bonds issued in the U.S. and denominated in USD
- » Include fixed-coupon only
- » Include callable bonds with "Cash Call" provision only. Other call features such as "Make Whole Call," or "Special Event Call" are not supported by RiskFrontier yet.
- » Include bonds with Agency Ratings
- » Exclude bonds with convertible, equity claw back, or private placement features

In addition to the aforementioned filters, we also exclude bonds with missing TRACE transaction prices or missing EDF value information. The data sample covers January 2006 – August 2013. We use approximately 6,500 unique bonds issued by 412 issuers, contributing to approximately 49,000 vanilla and 32,000 callable bonds observations at a monthly frequency. As discussed in Section 4.1, we subsequently apply liquidity filters to the full sample to test model performance on a liquid sample.

Figure 6 shows the distribution of the sample by origination year. The majority of the bonds in our sample are issued after 2001. However, there are a few bonds issued in early the 1990's as well. For the vanilla bonds sample, the maximum number of bonds come from 2004, whereas, for the callable bonds sample, most are from 2005. It is worth mentioning that there is a significant drop in the number of callable bonds issued after 2008, despite the fact that there is a significant increase in callable bonds issuance post-crisis. Our investigation shows that, after the crisis, there are more callable bonds that are either not "Cash Callable" or that have additional optionalities, such as equity claw back, convertible, etc.[4] Since we filter out such bonds from the sample, we notice a drop in the number of callable bonds after the crisis.
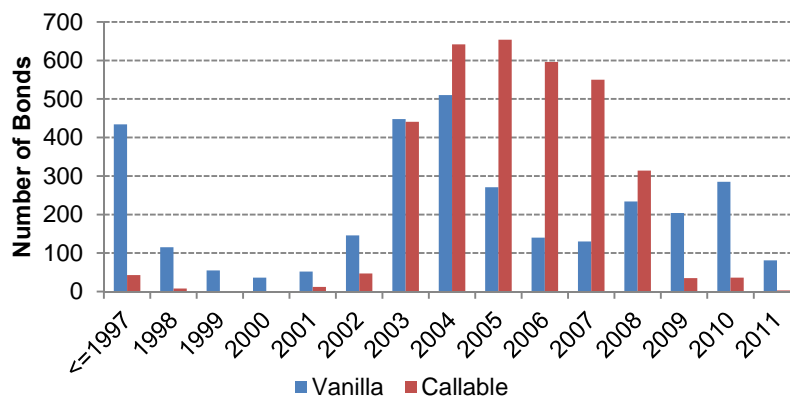
Figure 6   Sample Breakdown By Origination Year

Figure 7 shows the sample distribution by valuation year. Although we have a comparable number of callable and vanilla bonds in our sample, there are fewer callable bonds observations. This trait is due to lower trading activity for callable bonds and, hence, missing TRACE prices. That being said, the sample includes observations from both high and low interest rate periods.

---

[4] May be due to very low interest rates, which would diminish the benefit of bonds issued with cash-call feature only.

Figure 7   Sample Breakdown By Valuation Year
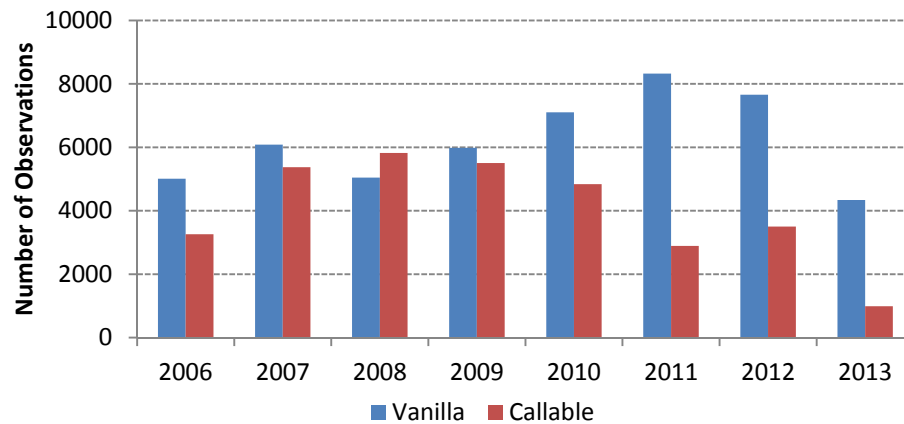
*Includes observations only up to June 2013

Figure 8 and Figure 9 display data distribution by industry and by debt type, respectively. The most common industries include Other Financials, followed by Bank and S&Ls and Consumer Goods and Durables. There are only a trace of bonds issued by firms in the Financials and Utilities industries. In terms of debt type, approximately 90% of the bonds fall into either the senior unsecured or the senior secured categories.

Figure 8   Sample Breakdown by Industry Sector
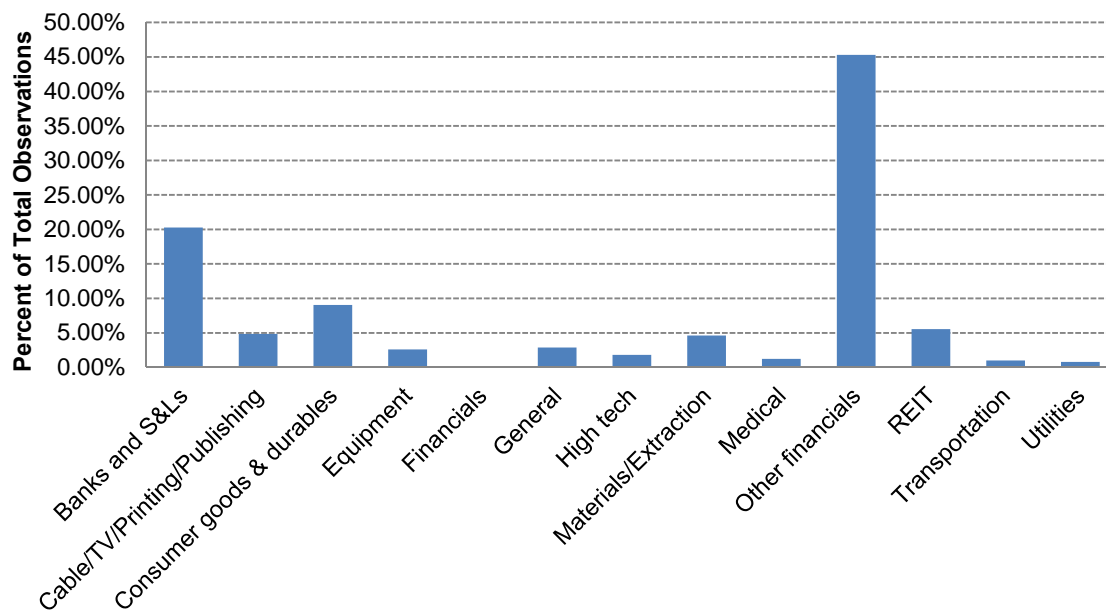
Figure 9   Sample Breakdown by Debt Type
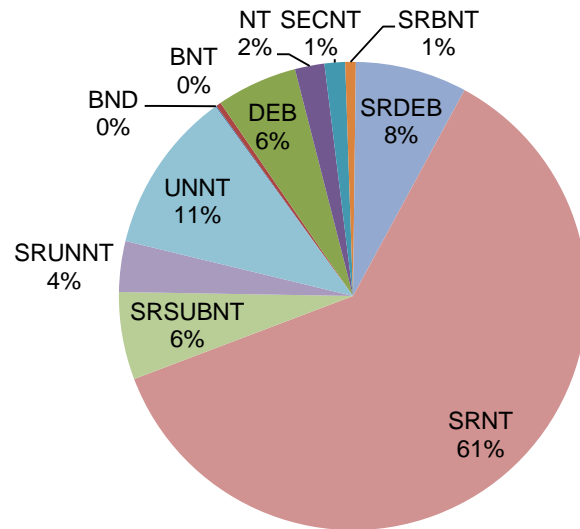
3.2 Input Parameters

To value the bonds using the lattice model, we require the following model parameters:

» **Credit Risk Parameters:** Credit risk parameters include the probability of default (PD), the loss given default (LGD), and the market Sharpe ratio (λ). We use both EDF- and CDS-implied parameters to value the bonds and to analyze the price differences between the model and the market. We use the market Sharpe ratio to transform between a physical and a risk-neutral measure using the following equation:

$$Q_t = N\left(N^{-1}(P_t) + \lambda\rho\sqrt{t}\right)$$

Where
t : time horizon
Q_t and P_t : cumulative risk-neutral and physical default probabilities
λ : the market Sharpe ratio or the market price of risk
ρ : the correlation between the asset return of the issuer and of the market
N and N^{-1} : the cumulative Normal distribution function and its inverse function

For CDS-implied parameters, we estimate separate market Sharpe ratios for investment grade and non-investment grade issuers, whereas, in the case of EDF-implied parameters, there is only one Sharpe ratio for the two risk categories.

Figure 10 displays the 90th percentile, median, and the 10th percentile of CDS Spread Implied EDF (SI EDF) credit measures for the sample period. As it can be noticed, SI EDF measures increased dramatically during the 2007 – 2008 crisis and reached their peak levels in 2009. The SI EDF measures have been decreasing since then, though the 2013 levels are still higher than that of 2006. Figure 11 shows the market Sharpe ratio variation over time. The graph plots the CDS-implied risk premiums for investment grade and non-investment grade categories. It is not surprising that the market Sharpe ratio also spiked in 2008 for both risk categories. However, the two Sharpe ratios have come down since then and have remained very close to one another.

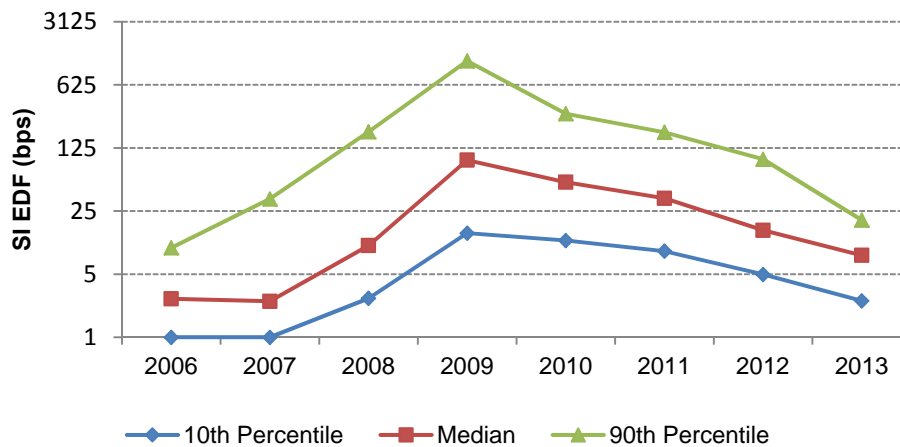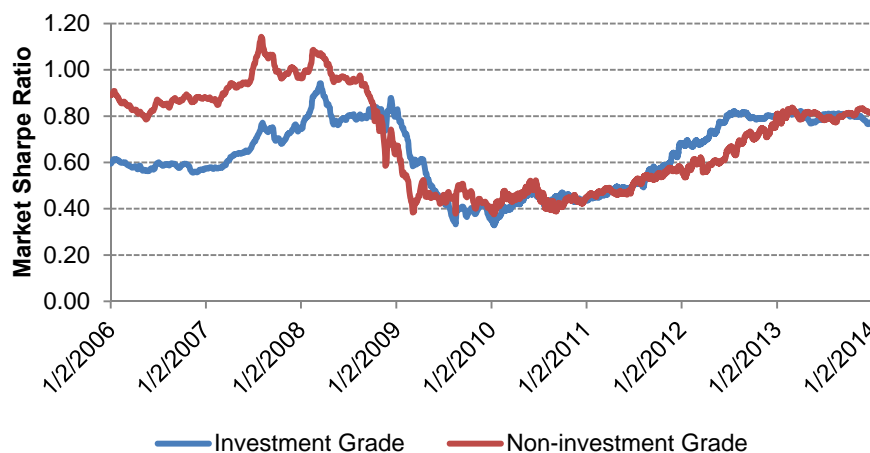Figure 10   EDF Credit Measure During the Sample Period

Figure 11   Market Sharpe Ratio During the Sample Period

» **The Hull-White Interest Rate Model Parameters:** The two parameters of the Hull-White one factor interest rate model — mean reversion (κ) and volatility (σ) — are estimated from the historical bond volatility using the "Unannualized Optimization" approach described in the white paper "Estimating Parameters in the Single-Factor Hull-White Model Using Historical Data" by Meng, et al. (2013). This approach is the same as the one used for the parameters provided in RiskFrontier.

Figure 12 and Figure 13 show the time series of κ and σ estimated using the "Unannualized Optimization" approach and the "Swaption" approach, respectively. The "Unannualized Optimization" approach-based parameters are more stable and capture the long-term trend. As mentioned in the above paper by Meng, et al., these parameters are better suited for the purpose of risk management. In this study, we briefly discuss how the valuation results change when using the "Swaption" approach-based parameters.

Figure 12   Hull-White Model Parameters Estimated from Bond Volatility
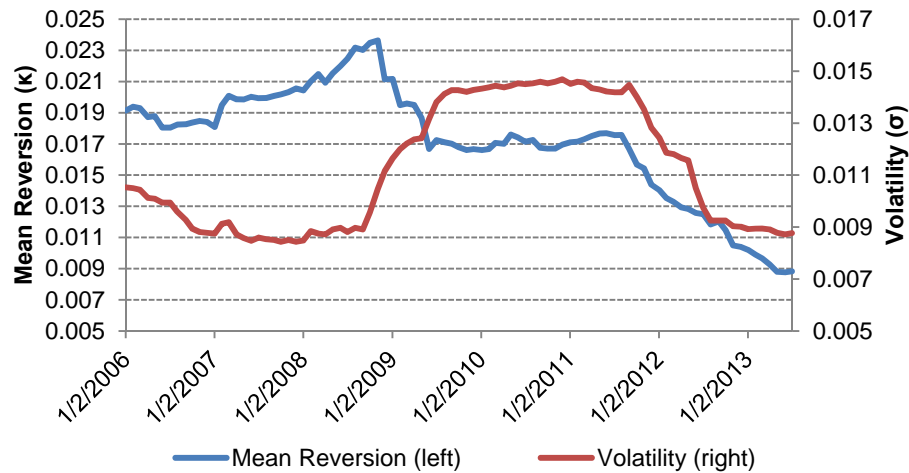
Figure 13   Hull-White Model Parameters Estimated from Swaption Volatility
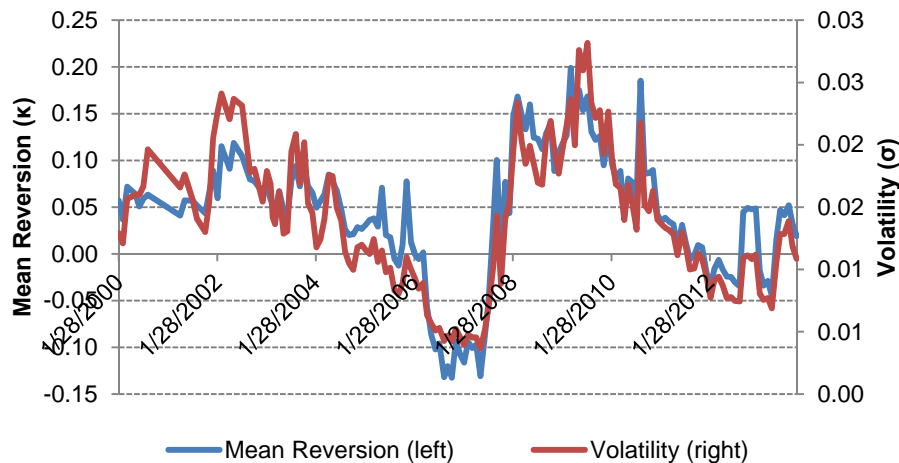
» **Correlation parameters:** We estimate the two correlation parameters — a) correlation between the asset returns of the issuer and the market (a measure of systematic risk), and b) correlation between the asset returns of the issuer and the interest rate — using the GCorr Macro framework.

4. Results

In this section, we first present the valuation results and then compare the performance of the MDL with the credit-only lattice. We show results for the full sample based on both CDS- and Equity-implied risk parameters and illustrate through an example why the three markets — the bond, the equity, and the CDS market — may not move together. We then focus on liquid sample results based on CDS-implied parameters and compare the performance of the two lattice models.

Next, we discuss the risk integration results using the bottom-up approach implemented in the MDL and a traditional top-down approach. We compare the risk measures based on these two approaches for vanilla and callable bond portfolios under different interest rate environments.

4.1 Valuation Results

FULL SAMPLE

We value the bonds in our sample using the MDL model and calculate the pricing error as the difference between the TRACE's transaction price and the model price. Figure 14 shows the pricing error distribution for vanilla bonds using both CDS- and equity-implied risk parameters. In both cases, the distribution is centered at zero and is nearly symmetric around it. This finding shows that there is no systematic bias in model prices, and, on average, model prices are close to market prices. Further, the CDS-based prices are closer to the market prices than the equity-based prices. Our analysis shows that higher pricing errors using the equity-implied parameters are driven primarily by Other Financials and REIT sectors, for which, the PD estimates for the 2009 – 2012 period were somewhat conservative. For the remaining sectors, the equity-based results are comparable with the CDS-based results.

Figure 14   Pricing Errors Distribution for Vanilla Bonds
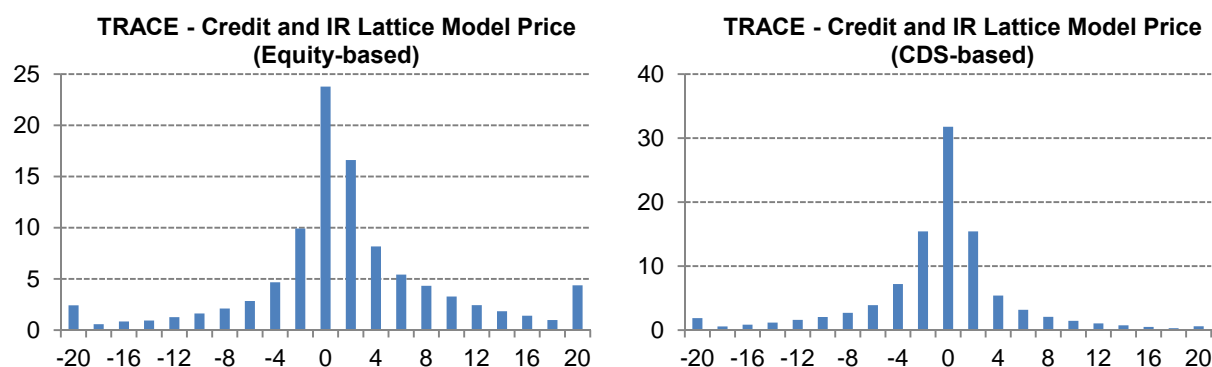
Figure 15 shows the error distribution for callable bonds using CDS-implied risk parameters. Table 1 shows the summary statistics of pricing error distributions. On average, the errors for callable bonds are higher than vanilla bonds. This result could be due to relatively lower trading activity for callable bonds, which would lead to higher bid-ask spreads and less accurate transaction prices. It is also possible that the price of the embedded option in the market reflects additional information not incorporated in our lattice model.
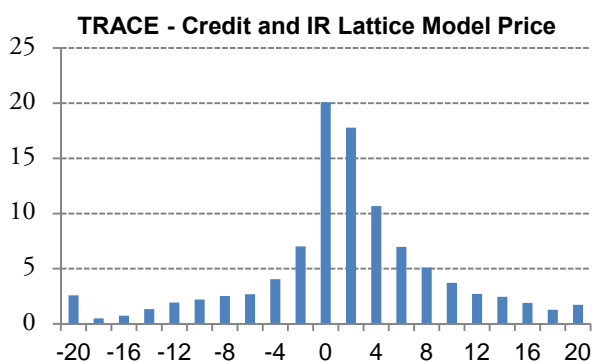
Figure 15   Pricing Errors Distribution for Callable Bonds

Table 1

Summary Statistics of Pricing Errors on Full Sample

BOND TYPE	STATISTIC	EQUITY-IMPLIED	CDS-IMPLIED
	Mean Error	0.80	-1.93
Vanilla	Mean Absolute Error	5.84	4.17
	Std Dev	9.52	6.54
	Mean Error	9.70	0.35
Callable	Mean Absolute Error	11.51	5.60
	Std Dev	12.39	8.45

We investigate the pricing errors further by studying the general pattern in the co-movements of equity, CDS, and bond markets. Our analysis shows that the three markets may not always move in the same direction. Figure 16 shows an example in which the higher credit riskiness reflected in the CDS and the Equity market is not observed in the bond market: the CDS spread increased from 133 bps to 283 bps, and the stock price dropped by 30%, but the bond price did not fluctuate much. Such behavior could also contribute to price differences between the model and the market.

Figure 16   Prudential Financial: Movements in the Equity and CDS Market and Their Impacts on Bond Price
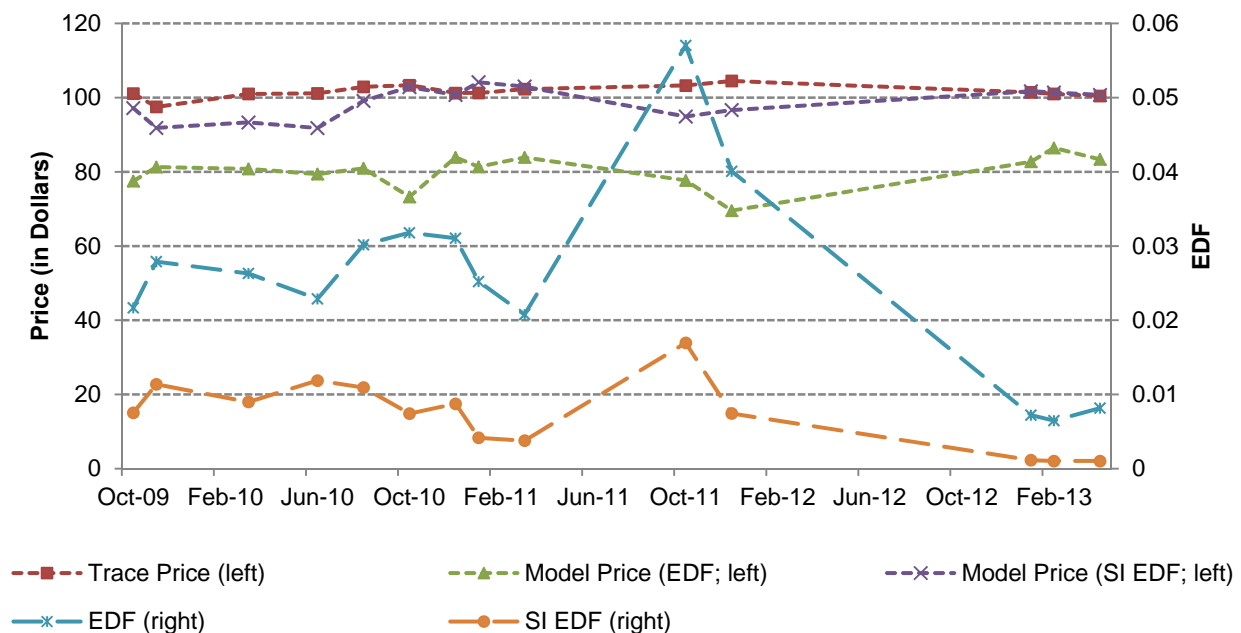
LIQUID SAMPLE

In our full sample analysis, we observe some large pricing errors (> $20 or < -$20) for both vanilla and callable bonds. Our further and detailed examination of these errors indicates that it may be driven by illiquidity in the bond market and, thus, motivates us to focus on the liquid sample.

To filter the bonds based on their liquidity, we estimate the following two commonly used measures of liquidity using end-of-day trading information. We use this method as we do not have the historical bid-ask spread information readily available.

$$Amihud_t = \frac{1}{N_t}\sum_{i=1}^{N_t} \frac{|r_i|}{Q_i}$$

$$Roll_t = 2\sqrt{-covariance(r_i, r_{i-1})}$$

Where

$Amihud_t$ = Amihud measure of illiquidity

$Roll_t$ = Roll's measure of illiquidity

r_i = return for day i for month t

Q = traded quantity for day i

N_t = number of days with trading information for month t

Figure 17 shows the time-series of the bond market illiquidity, measured using these two measures. Results are intuitive, in the sense that the market liquidity was high before and after the 2008 – 2009 crisis, but dropped dramatically during the crisis.[5]

Figure 17   Bond Market Illiquidity Over Time
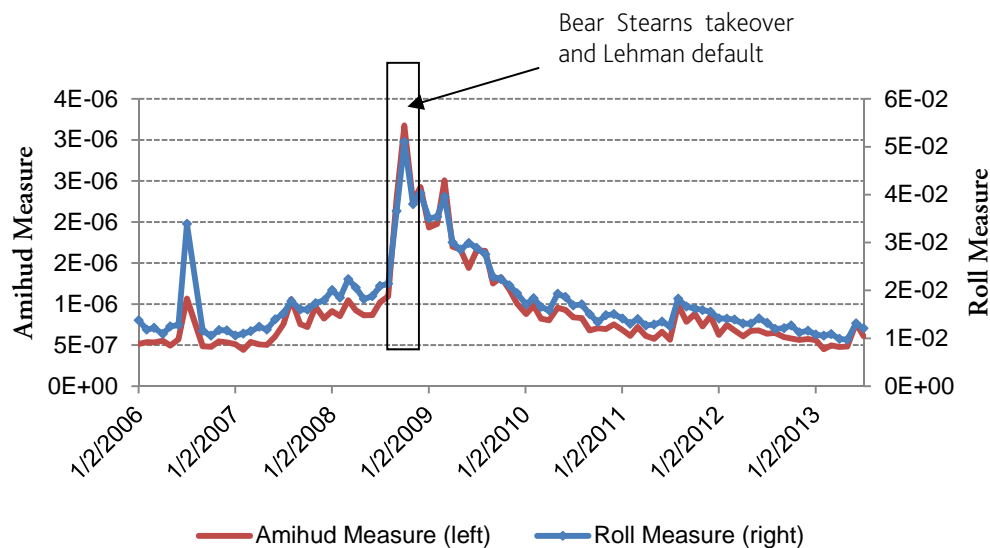
Using these liquidity measures, we group bonds into liquid and illiquid samples by applying the liquidity filter to each year separately. The liquid sample selected consists of bonds with Amihud measure in the 0-25[th] percentile from each year, whereas, the illiquid sample comprises the remaining bonds. We then compare the pricing errors for these two sub-samples, and find that the errors for the liquid sample are significantly smaller than those of the illiquid sample.

Figure 18 and Figure 19 show the pricing error distributions for vanilla and callable bonds, respectively. For the liquid sample, 88% of vanilla bond and 63% of callable bond pricing errors are within a $4 difference. This result compares with 63% of vanilla bonds and 45% of callable bonds with pricing errors for the illiquid sample. These results highlight that model performance is significantly enhanced when we filter out illiquid bonds.

---

[5]We require a bond to be traded for at least 10 days during a month to reliably estimate its liquidity measure.

Figure 18 Pricing Errors Distribution on Liquid and Illiquid Samples: Vanilla Bonds
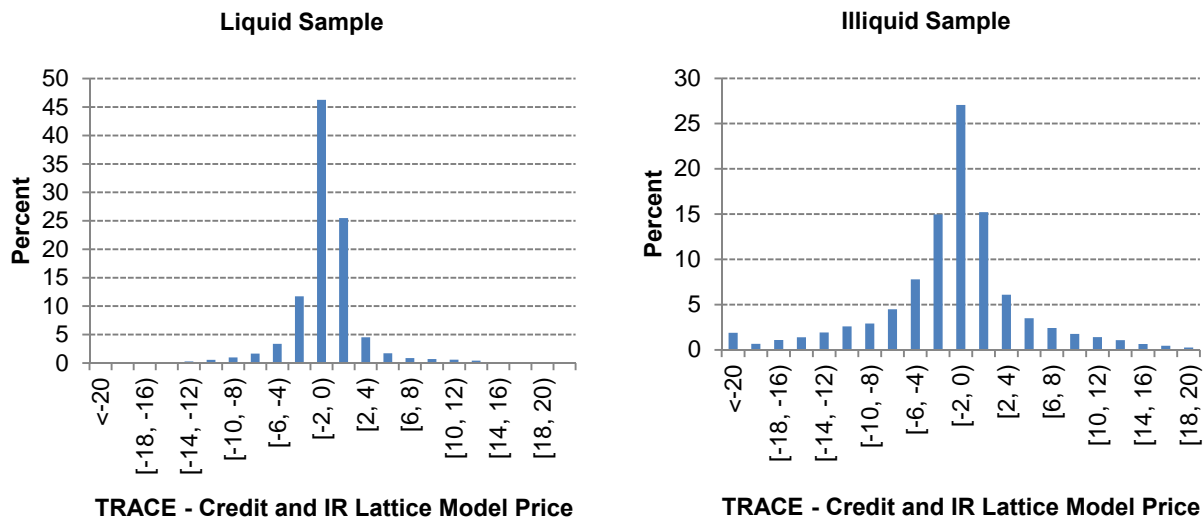
Figure 19 Pricing Errors Distribution on Liquid and Illiquid Samples: Callable Bonds
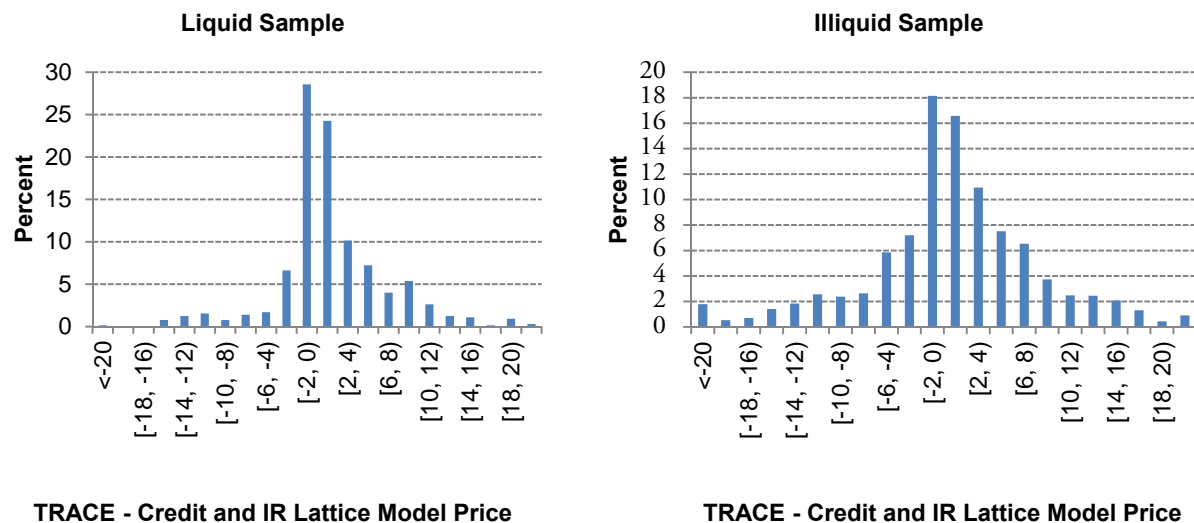

Table 2 and Table 3 show the summary statistics of vanilla bond and callable bond error distributions, respectively. The liquid sample pricing errors are roughly 30-50% smaller than those of the illiquid sample. It is worth mentioning that the errors for the lattice model are somewhat larger than the EJV errors. This finding is not surprising, as the EJV model uses TRACE pricing history to calibrate the option-adjusted spread that, in turn, is used in pricing. Hence, it is expected to better match trace prices.

Table 2

Summary Statistics of Pricing Errors on Liquid and Illiquid Sample: Vanilla Bonds

	LIQUID SAMPLE		ILLIQUID SAMPLE	
STATISTIC	CREDIT AND IR LATTICE	EJV	CREDIT AND IR LATTICE	EJV
Mean Error	-0.52	0.08	-1.74	-0.32
Mean Absolute Error	1.90	0.66	4.60	2.06
Standard Deviation	3.44	1.54	6.94	4.55
Number of Observations	7,481		22,193	
Number of Bonds	807		1,299	
Number of Issuers	146		188	

Table 3

Summary Statistics of Pricing Errors on Liquid and Illiquid Sample: Callable Bonds

	LIQUID SAMPLE		ILLIQUID SAMPLE	
STATISTIC	CREDIT AND IR LATTICE	EJV	CREDIT AND IR LATTICE	EJV
Mean Error	2.23	1.43	1.95	1.95
Mean Absolute Error	4.58	2.62	6.51	6.51
Standard Deviation	6.64	4.29	9.04	9.04
Number of Observations	995		5,736	
Number of Bonds	331		654	
Number of Issuers	52		55	

VALUATION RESULTS USING SWAPTION MARKET BASED HULL-WHITE MODEL PARAMETERS

The Hull-White interest rate model parameters used in our study are estimated from historical bond volatility. In this section, we discuss how results would compare if we use the parameters estimated from at-the-money swaptions' volatility. As noted in Section 3.2, the parameters from the swaption approach are more varying and resemble a point-in-time measure of interest rate dynamics.

Table 4 presents the average pricing errors for callable bonds,[6] based on mean reversion (κ) and volatility (σ) parameters estimated using the two approaches. The two sets of parameters produce fairly similar results. The mean absolute error using the swaption volatility-based parameters is about five cents lower than the bond volatility-based parameters.

Table 4

Pricing Errors Using Bond Volatility-Based and Swaption Volatility-Based Hull-White Model Parameters

|	LIQUID SAMPLE		ILLIQUID SAMPLE	
BOND TYPE	BOND VOLATILITY BASED PARAMETERS	SWAPTION VOLATILITY BASED PARAMETERS	BOND VOLATILITY BASED PARAMETERS	SWAPTION VOLATILITY BASED PARAMETERS
	1.35	1.15	0.53	0.36
Callable	3.63	3.57	5.55	5.49
	5.55	5.51	8.06	8.04

[6]For this comparison study we use only a sub-sample of our data. Hence, the bond volatility parameters based results slightly differ from those of in Table 3.

4.2 Comparison of Credit-Interest Lattice and Credit-Only Lattice

As mentioned in Section 2.3, we compare the two lattice models by studying how closely the model prices track the market price benchmark. In the following section, we describe some of the TRACE data features, and we discuss the construction of our "close price" benchmark used in Section 4.1, and how we refine it.

TRACE data includes the bond closing price, as well as additional information on trading activity, such as the total number of transactions and traded quantities on any given day. The number of transactions on a given day can vary depending on bond liquidity. In our dataset, the median number of transactions per day is two, whereas, the maximum number of transactions per day is 355 for Principal Life Insurance Company bond (CUSIP 74254PYL0) on July 7, 2008. We construct our "close price" benchmark using the closing or last transaction price on a given day (similar to the construction of equity benchmarks). If there are no transactions on that particular day, we look forward one business day and use the last traded price of that day. In case there are no transactions on the subsequent business day as well, we drop that bond observation from our dataset. Thus, we only use the price of the last transaction (sorted by the time stamp) and ignore other transaction prices on a given day. Although, this construction appears to be straightforward, in terms of computation and rationale, we notice that bonds with almost identical covenants (such as maturity, coupon frequency, call dates, coupon type, bond issuer, and seniority, but differ only in origination date), henceforth, "similar bonds," are priced differently by the market on a same day. Often times, the difference in benchmark TRACE prices between similar bonds is often greater than $5.00 on the same day. Figure 20 shows the close benchmark for couple of similar General Electric bonds issued one week apart. There are nine instances (about 5.6% of the total observations) where the prices of these two bonds differ by more than $5.00 between October 2004 and August 2013. We can also see that more than 50% of the time, the "close price" benchmarks of these bonds differ by about one dollar.

Figure 20   Close Benchmark Prices for Similar Bond Pair
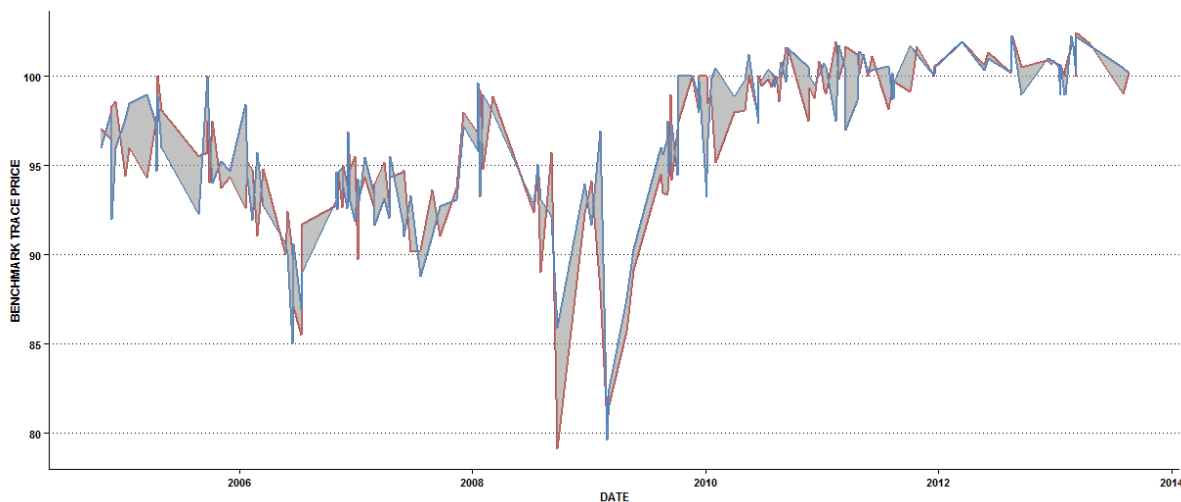
It is possible to make the conjecture that these price differences between similar bonds arise because of the sub-optimal construction benchmark. To that end, we test several approaches to constructing a robust benchmark that enables us to conduct a fair comparison of the two lattice models. The Appendix provides details different benchmark construction using the transaction level information and compares their performance. In the end, we use a fixed-benchmark with lag of three. It should be noted that the new benchmark merely allows us to capture the subtle differences between the two lattice models. As shown in Figure 21, there is no systematic bias between the close price benchmark used in Section 4.1 for comparing the model with the market and the fixed benchmark with lag of three. In other words, we do not find any evidence of one benchmark consistently outperforming or underperforming the other. So, the new benchmark would produce results almost identical to those based on the close price benchmark.

Figure 21 Monthly Mean Differences between Close Price and Fixed Benchmarks with Lag of Three

MONTHLY MEAN DIFFERENCE BETWEEN CLOSE PRICE AND FIXED BENCHMARK WITH LAG 3
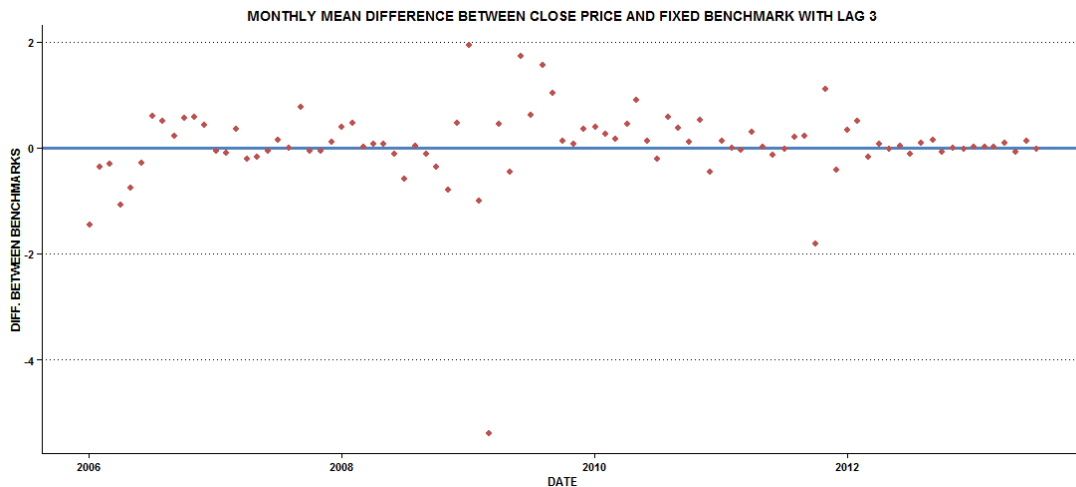
Next, we compare the credit-only and credit-interest rate model prices against the newly constructed benchmark with lag of three. Figure 22 presents the average monthly model prices and the corresponding benchmark price. During the high interest rate period (2006 – 2008), the prices generated by the Credit-interest lattice are closer to the benchmark price. This finding is in-line with our expectations, as the impact of interest rates on embedded optionality is more pronounced during times of high interest rates. In other words, there is a higher expectation that future interest rates would be lower, and the bond is more likely to be called by the issuer. This results in lower bond prices in the case of Credit-Interest lattice, such that the difference between the two lattice models accounts for the risk of reinvesting cash flows at lower interest rates (i.e., increased prepayment or reinvestment risk). However, when the interest rates drop significantly (2009 onwards), the value of embedded options increases and the impact of interest rates on embedded option prices is dampened (the sensitivity of option price to interest rate changes decreases as the option money-ness increases). Therefore, the average price from the two models converge in recent years and stay close to strike price (i.e., par).

Figure 22 Monthly Mean Model and Benchmark Prices Across Time

MONTHLY MEAN PRICES ACROSS TIME
BENCHMARK : FLOAT LAG - 3

— CREDIT PRICE — FLOATLAG 3 — IRR PRICE — NO-OPTION PRICE
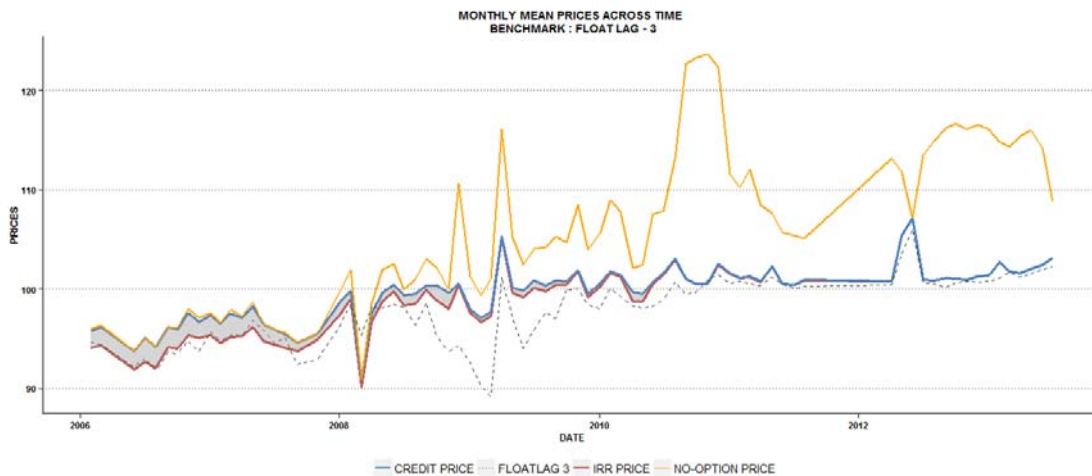
Figure 23 presents the error comparison results for high credit quality bonds across all years. Here we count the number of times one model outperforms the other in matching the trace price.[7] The study shows that the MDL model provides better valuations for bonds with high credit quality. In 2006, the MDL outperforms the credit-only lattice 68% times.[8] This finding suggests that the

---

[7] We use a threshold of $0.50 and count only those occurrences when the two models prices differ by more than the threshold value.

[8] The comparison of average pricing errors also showed similar pattern.

market perceives high credit quality bonds to be more sensitive to interest rate movements and as the credit quality deteriorates, the credit risk starts dominating the interest rate risk.

Figure 23   Model Comparisons Through Time
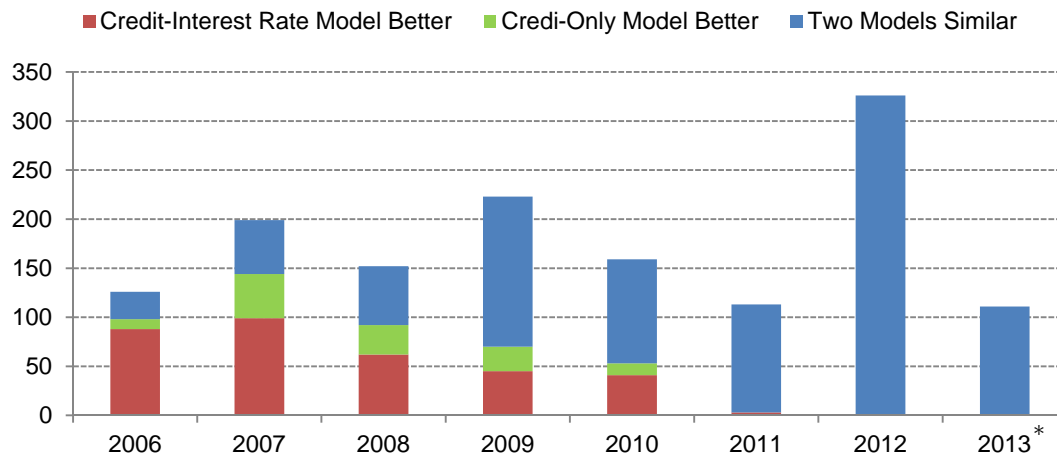
Figure 24 shows comparison of absolute pricing errors across various years. Clearly, MDL produces smaller errors when compared to those of credit-only lattice. We also show the market illiquidity proxy.

Figure 24   Comparison of Pricing Errors of Credit-Interest Lattice and Credit-Only Lattice
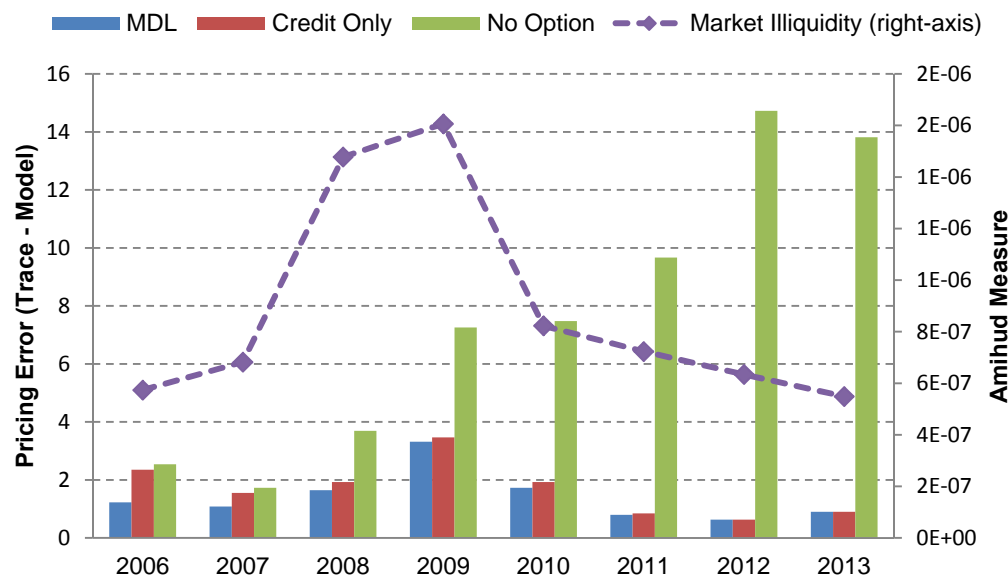
* Includes observations only up to June 2013

EXTRA OPTION VALUE DUE TO STOCHASTIC INTEREST RATES

Next, we illustrate how stochastic interest rates impact the value of the embedded option. For exposition, we construct a hypothetical callable bond with semi-annual payment frequency, five years to maturity, and 50% loss given default (LGD). The embedded call option is Bermudan, with strike price equal to par. We use an upward sloping yield curve and set the mean reversion and volatility parameters (σ) equal to 5% and 1%, respectively. The level of yield curve is such that a bond with default probability (PD) equal to 20 bps is priced around par. Given this information, we use the two lattice models to value bonds by varying coupon and default probability. Table 5 shows these valuation results, from which we can see that:

» the extra option value is maximum when the underlying bond is around par, and

» the extra option value is higher for bonds with lower default probability.

Table 5

Extra Option Value Due to Stochastic Interest Rates in Credit-Interest Lattice

		CREDIT-ONLY LATTICE				CREDIT-INTEREST LATTICE				EXTRA OPTION VALUE DUE TO STOCHASTIC INTEREST RATES			
						1 YEAR PD							
		5 BPS	10 BPS	20 BPS	100 BPS	5 BPS	10 BPS	20 BPS	100 BPS	5 BPS	10 BPS	20 BPS	100 BPS
Coupon = 3%	Bond Value	92.19	92.08	91.88	90.24	92.02	91.92	91.72	90.13				
	Option Value	0.00	0.00	0.00	0.00	0.17	0.17	0.16	0.11	0.17	0.17	0.16	0.11
Coupon = 4%	Bond Value	96.61	96.50	96.28	94.55	96.03	95.94	95.74	94.15				
	Option Value	0.00	0.00	0.00	0.00	0.57	0.56	0.53	0.40	0.57	0.56	0.53	0.40
Coupon = 5%	Bond Value	100.73	100.69	100.56	98.85	99.38	99.31	99.15	97.73				
	Option Value	0.30	0.22	0.12	0.01	1.64	1.60	1.53	1.13	1.34	1.38	1.41	1.12
Coupon = 6%	Bond Value	101.22	101.21	101.18	100.72	101.17	101.15	101.09	100.27				
	Option Value	4.22	4.11	3.90	2.45	4.27	4.17	3.99	2.90	0.05	0.06	0.09	0.45

Figure 25 displays the extra option value for the callable bonds in our sample. Similar to what we stated above, the value peaks when the underlying bond price is around par. These results show that the impact of stochastic interest rates on the option value depends on the credit riskiness of the bond, as well as the level of interest rates relative to the coupon.

Figure 25   Extra Option Value vs. Underlying Bond Price (Model Price)
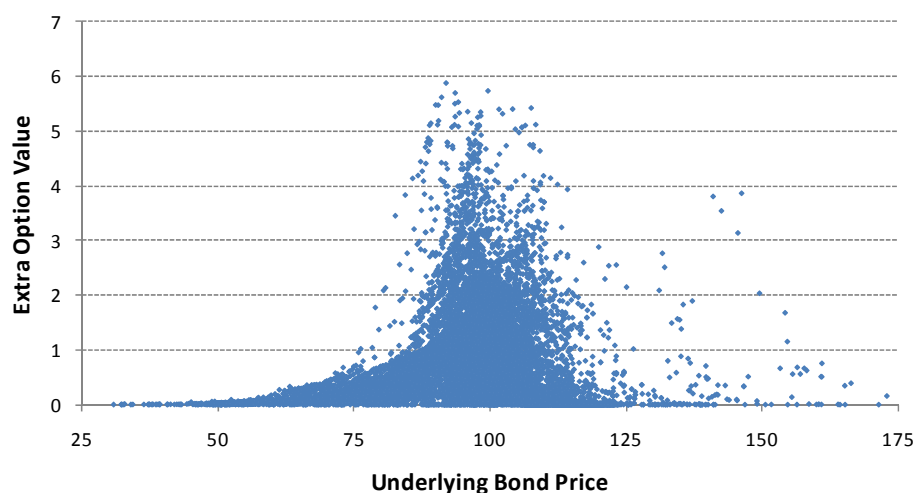
4.3 Risk Integration Results

We perform the risk integration study to compare the bottom-up approach implemented in the MDL with a traditional top-down approach. In the following section, we describe the construction of test portfolios and the computation of their risks using the top-down and bottom-up risk integration approaches.

TEST PORTFOLIOS' CONSTRUCTION

Starting with the full sample described in Section 3.1, we construct vanilla and callable bond portfolios separately for the following two periods:

» June 2006: High interest rate environment
» June 2011: Low interest rate environment

As the interest rate environments before and after the financial crisis are significantly different, we perform validation exercises for both periods to help us better understand the impact of interest rate levels on risk integration results. Table 6 presents the description of the test portfolios used in the study and Figure 26 displays the zero-EDF yield curves as of the analysis dates. It should be noted that the bonds constituting these portfolios are actual world bonds trading in the market. Also, the analysis is done using CDS-implied credit risk parameters.

Table 6

Summary Statistics of Test Portfolios for Risk Integration

ANALYSIS DATE	INSTRUMENT TYPE	NUMBER OF BONDS	REMAINING MATURITY (YRS)	COUPON	PD	LGD	RSQ
June 2006	Vanilla	1,898	7.32 yrs	5.71%	10 bp	59%	39%
	Callable	323	8.95 yrs	5.24%	9 bp	56%	42%
June 2011	Vanilla	2,099	6.62 yrs	5.72%	47 bp	57%	32%
	Callable	699	17.38 yrs	5.88%	86 bp	50%	32%

Figure 26   Yield Curve as of Analysis Date
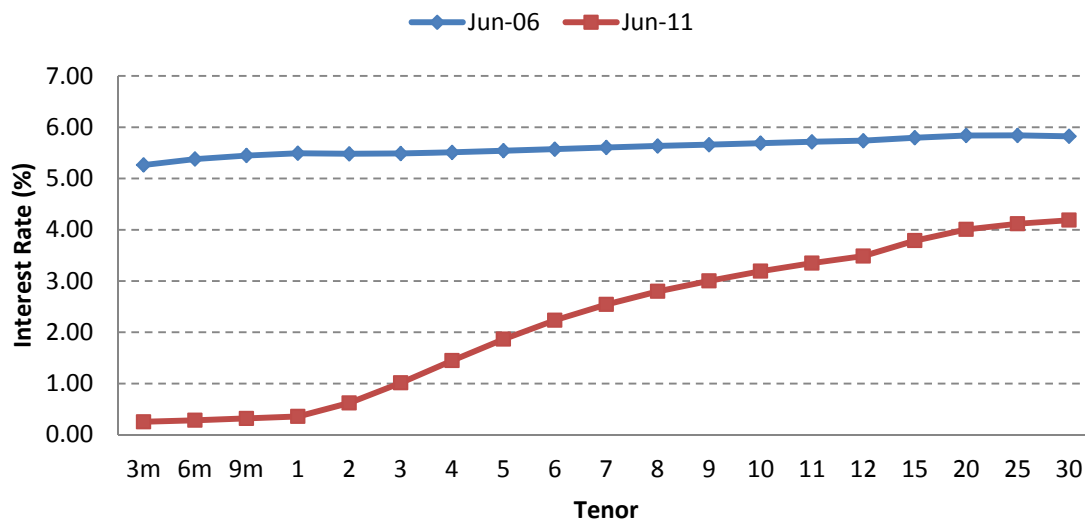
TOP-DOWN VERSUS BOTTOM-UP APPROACHES TO RISK INTEGRATION

Once we select the test portfolios, we run them in RiskFrontier to obtain their Unexpected Loss (UL) estimates using the credit-only lattice, the interest rate-only lattice, and the MDL, which implements the bottom-up approach.[9] The credit-only lattice model measures the credit-only component of the portfolio risk, assuming no interest rate risk (the interest rates are held deterministic). Similarly, the interest rate-only lattice model measures the interest rate risk-only component of the portfolio risk, assuming no credit risk. Since RF does not allow analyzing instruments with zero credit risk, we set the default probability (PD) equal to 0.1 bps (the lower bound in RF) to ensure that there is virtually zero credit risk and the interest rate-only lattice truly captures the interest rate-only risk component. The standalone credit and interest rate risk measures are then combined using the portfolio variance formula to calculate the UL for the top-down approach as:

$$UL_{topdown} = \sqrt{UL^2_{Creditonly} + UL^2_{IRonly} + 2 * \rho_{credit,IR} * UL_{Creditonly} * UL_{IRonly}}$$

Here $\rho_{credit,IR}$ is the correlation parameter between credit and interest rate risks. It is estimated within the Moody's GCorr Macro framework.

Table 7 shows the standalone ULs as well as the integrated ULs using the top-down and bottom-up approaches under the assumption of zero correlation between credit and interest rate risks. Table 8 shows similar results for 2011 portfolios when using the GCorr Correlation parameter, which is positive during our test period.

Table 7

Risk Integration Results Assuming Zero Correlation ($\rho_{credit,IR} = 0$)

ANALYSIS DATE	INSTRUMENT TYPE	CREDIT-ONLY UL	IR-ONLY UL	TOP-DOWN UL	BOTTOM-UP UL	% DIFFERENCE IN UL
June 2006	Vanilla	1.04%	4.48%	4.78%	4.08%	17.18%
	Callable	1.30%	3.59%	3.92%	4.08%	-3.82%
June 2011	Vanilla	2.15%	0.43%	2.20%	2.12%	3.58%
	Callable	2.49%	0.03%	2.50%	2.96%	-15.37%

Table 8

Risk Integration Results With GCorr Correlation ($\rho_{credit,IR} = 0.15$)

ANALYSIS DATE	INSTRUMENT TYPE	CREDIT-ONLY UL	IR-ONLY UL	TOP-DOWN UL	BOTTOM-UP UL	% DIFFERENCE IN UL
June 2011	Vanilla	2.15%	0.43%	2.27%	2.14%	6.30%
	Callable	2.49%	0.03%	2.51%	3.40%	-25.94%

Notice, the top-down approach overstates risks for vanilla bond portfolios, while understating risk for callable bond portfolios. The reason for overstating vanilla bond portfolio risk in the top-down approach is due to inappropriately accounting for interest rate risk in the default state. On the other hand, imprecise accounting for call options exercised at low credit quality or default states leads to understating the risks for callable bond portfolios.

With correlation between credit and interest rate risks, the error in UL is more pronounced: the vanilla bond portfolio UL error increases from 3.6% to 6.3%, whereas, for the callable bond portfolio, the error jumps from -15% to -26%. The increase in error is due to higher probability of the states where the top-down approach incorrectly accounts for the interest rate risks for vanilla bonds or call decisions for callable bonds. Further, the magnitude of the error depends on the interest rate environment. These results demonstrate that, for a portfolio sensitive to both credit and interest rate movements, the bottom-up approach implemented in the joint lattice is definitely better in assessing portfolio risks than a top-down approach.

---

[9] Portfolios are run with 1-year horizon and 10^6 simulation trials.

5. Conclusion

A financial institution's portfolio typically includes instruments such as corporate bonds, which are subject to both credit and interest rate risks. Since these two risks are generally the most important risk types driving loss, their accurate assessment is integral for the stability and profitability of the institution. When top-down approaches to risk integration are used for these portfolios, the computed risk measures are not reliable, as they do not account for the interaction between these two risks. In such cases, a bottom-up approach, as implemented in the MDL, is better suited for accurately modeling portfolio risks.

In this paper, we demonstrate the effectiveness of the MDL in modeling credit and interest rate risk together. Our study employs a simple empirical validation focusing on the valuation of a large sample of corporate bonds. We show that the joint lattice model within RiskFrontier produces valuations that compare well with market transaction prices. The price differences between the model and the market are within an acceptable range for bonds with different ratings and industries.

Next, we show how the bottom-up approach implemented in the joint lattice is better-suited for assessing portfolio risks. The findings are robust across different interest rate environments, instrument types, and the level of correlation between credit and interest rate risks. Through our analysis, we show how a top-down approach to risk integration may not be appropriate for credit risk portfolios also sensitive to interest rate movements.

Finally, we provide important insights into the variation of the extra option value due to stochastic interest rates. We illustrate that the extra option value is highest when the underlying bond price is near par and is lower for prices above or below par.

Appendix

Benchmark Construction

Our deeper analysis of the "close price" benchmark reveals that, while the benchmark is reasonable to assess the model performance, it may not be ideal for assessing the impact of interest rate risk on credit risky instruments. In particular, we find that "close price" benchmarks of similar bonds differ significantly across time. Given that these bond pairs have exactly same the characteristics, we believe that observed price differences could potentially arise under the following circumstances:

» The corporate bond market is not deeply liquid. As a consequence, the prices of securities traded in this market reflect an illiquidity discount. Depending on the open interest of these bonds in the market, the illiquidity discount varies. Also, considering the possibility that dealers transact over-the-counter with interested counterparties, there is the possibility of a lack of real-time price transparency, since counterparties can only reach out to a limited number of liquidity providing dealers. This market inefficiency could potentially result in significant deviation between traded prices and the fair values.

» After analyzing the transaction prices for an extended period of time, it appears that there are periods when one of the bonds in a similar bond pair exhibits relatively higher trading activity than the other. However, there is no systematic bias with respect to the volume of these bonds traded across time. It could be simply that these differences are due to some bond-specific information in the market, as the issuer-level information affects both bonds (for bonds with similar covenants and seniority). The bond-specific information could relate to the supply and demand of the bonds, caused by factors such as hedging needs and portfolio rebalancing. If the market participants trade bonds to change their portfolio composition, consistent with their mandate and, given that all market participants do not update their portfolios at the same time, there will be differences in the trading activity of such bonds. Due to these factors, the supply of similar bonds may not be the same, resulting in dealers applying different illiquidity discounts and, hence, different traded prices. We would reach similar conclusions if we assume there is no inherent (bond-specific) information causing the difference in prices. In that case, we could attribute the differences to the supply-demand imbalances caused by changes in the portfolio compositions of various market participants.

To mitigate the impact of these issues on our comparison of the two lattice models, we explore constructing benchmarks using additional transaction price data. We believe using the collective information from a host of market participants can produce a benchmark that is less sensitive to short-term, demand-supply fluctuations. We constructed our first set of benchmarks as the average of prices for all transaction on a given day and a specific number of transactions in the past. We refer to these benchmarks as "float" benchmarks. We specify a lag term to denote the number of previous transactions used for estimating the benchmark. So, for example, float benchmark price with lag value of three on September 1, 2013 would be the average of all transaction prices reported on September 1, 2013 and the last three transaction prices in the past. Given limited transaction data for some bonds, using an average price helps in establishing a reasonable price estimate for a given day. However, specifying an appropriate lag is important, as a higher lag value may result in too much smoothing of prices and average out any price variation due to changes in market risk parameters. Our careful analyses of similar bond pairs show that using the lag value of three is most appropriate for minimizing the deviation in prices of similar bonds while retaining systematic price fluctuations.

We construct another set of benchmarks, considering that the benchmark price is an average of a fixed number of transactions counting from the last transaction on the subject date. We refer to these benchmarks as "fixed" benchmarks. As before, this strategy also uses information from past transactions. However, the important difference is that we use only a fixed number of previous transactions, as opposed to an arbitrary number of transactions, as in the case of float benchmarks (float benchmarks use all the transactions available on the subject date. i.e., for a lag of three transactions, if there are 10 transactions on the subject date, float benchmark would end up using 13 transactions to estimate the benchmark price. But, a fixed benchmark with lag of three uses only the last three transactions on the subject date). This approach has one notable advantage, it is less likely to use price transactions from a different regime, especially for bonds with extremely limited transactions. Based on our analysis, a lag value of three provides the most optimal results in terms of lowering the price deviations of similar bonds, while retaining systematic price fluctuations. We also notice that longer lags result in a lower deviation of prices for similar bonds, but they inevitably smoothen the price time series and discard most of its features

Last, we construct another set of benchmarks where we use the filtered bond transaction data and retain only transactions with larger trade sizes. We refer to this as the "quantile filtered" approach. The rationale here is that the prices of the transactions with smaller trade size would carry a higher illiquidity discount, since the trades are usually executed between dealers and counterparties. Large trades are typically executed between dealers, and, hence, these trades tend to reflect the true fair value of bonds. We construct float and fixed benchmarks using this filtered dataset. Although the rationale for filtering appears conceptually sound, these benchmarks produced sub-optimal results, primarily because of the limited transaction data. For example, a float benchmark with lag of three on the filtered transaction dataset might potentially result in using transaction prices months apart.

Table 9 shows the mean absolute deviation of the benchmark prices constructed using various approaches.

Table 9

Mean Absolute Deviation of Benchmark Prices of Similar Bond Pair with CUSIPS 36966RCS3 and 36966RCY0

LAG	FIXED	QUANTILE FILTERING - FIXED	FLOAT	QUANTILE FILTERING – FLOAT	CLOSE
1	1.5665	1.3419	1.1612	0.9705	
3	1.0993	1.0201	0.9890	0.7833	
5	1.0274	0.8535	0.9194	0.7681	1.6651
10	0.8315	0.6454	0.7360	0.6252	

Figure 27 shows the benchmark price dynamics of the similar bond pair for varying levels of lag under four different benchmark construction settings. It is evident from Figure 27 that there is information loss with higher lag variables. Also, as mentioned earlier, there is a clear smoothing of price time series under the quantile filtering setting. The relative benefit of using the newly constructed benchmarks is, at best, marginal. Using the bond prices in the subject dataset, the average absolute difference between the close price benchmark and the fixed benchmark with lag of three is only $0.42. Based on these results, we use the float benchmark with lag of three to compare against our model prices generated using credit-only and credit-interest lattice.

Figure 27   Fixed and Floating Benchmarks with and without Applying Quantile Filtering
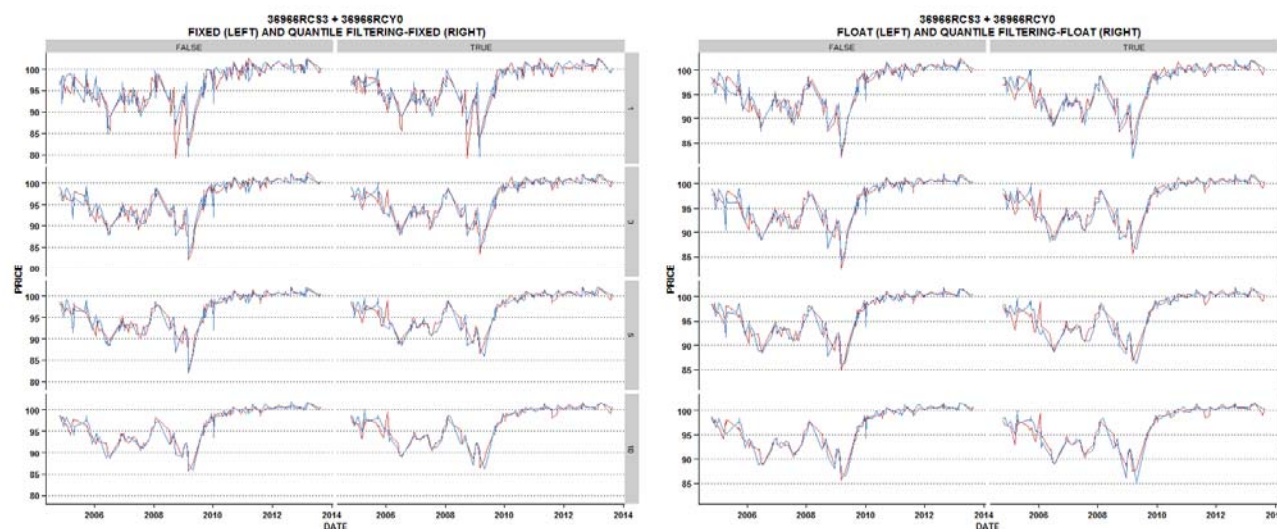
References

King, T. D., "An Empirical Examination of Call Option Values Implicit in U.S. Corporate Bonds." University of Wisconsin-Milwaukee, 2000.

Agarwal, D., I. Korablev, and D. W. Dwyer, "Valuation of Corporate Loans: A Credit Migration Approach." Moody's Analytics White Paper, 2008.

Kalotay, A., and L. Abreo, "Ratchet Bonds: Maximum Funding Efficiency at Minimum Transaction Costs." Andrew Kalotay Associates, Inc., 1999.

Chen, N., A. Kaplin, A. Levy, and Y. Wang, "Risk Integration: New Top-down Approaches and Correlation Calibration." Moody's Analytics White Paper, 2010.

Bliss, R. R., and E. I. Ronn, "Callable U.S. Treasury Bonds: Optimal Calls, Anomalies, and Implied Volatilities." *Journal of Business*, Vol. 71, 1998.

Longstaff, F. A., "Are Negative Option Prices Possible? The Callable U.S. Treasury-Bond Puzzle." *Journal of Business*, Vol. 65, 1992.

Hull, J., and A. White, "The General Hull-White Model and Super Calibration." *Financial Analysts Journal*, Vol. 57, 2001.

King, T. D., and C. Mauer, "Corporate Call Policy for Nonconvertible Bonds." *Journal of Business*, Vol. 73, 2010.

Meng, Q., A. Kaplin, and A. Levy, "Estimating Parameters in the Single-Factor Hull-White Model Using Historical Data." Moody's Analytics White Paper, 2013.

Meng, Q., Y. Pan, N. Patel, L. Pospisil, Y. Wang, and K. Yang, "Estimating Parameters of the Joint Interest Rate and Credit Model in RiskFrontier." Moody's Analytics White Paper, 2013.