MOODY'S ANALYTICS

MODELING METHODOLOGY

FROM MOODY'S KMV

Analyzing the Impact of Credit Migration in a Portfolio Setting

(A version of this has been published in the *Journal of Banking & Finance*, 2011, vol. 35, issue 12, pages 3145-3157)

Authors

Yaakov Tsaig

Amnon Levy

Yashan Wang

Contact Us

Americas
+1-212-553-1653
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

Abstract

Credit migration is an essential component of credit portfolio modeling. In this paper, we outline a framework for gauging the effects of credit migration on portfolio risk measurements.

For a typical loan portfolio, we find credit migration can explain as much as 51% of volatility and 35% of economic capital. We compare through-the-cycle migration effects, implied by agency rating transitions, with point-in-time migration, implied by EDF™ (Expected Default Frequency) transitions, and find that migration of point-in-time credit quality accounts for a greater fraction of total portfolio risk when compared with through-the-cycle dynamics.

In a stylized analytic setting, we show that, when controlling for PD term structure effects, higher likelihood of moving away from the current credit state does not necessarily imply greater risk.

Finally, we review methods for generating high-frequency transition matrices, which are needed to analyze instruments with cash flows or contingencies whose frequencies are asynchronous to an available transition matrix. We further demonstrate that the naïve application of such methods can result in material deviations to portfolio analytics.

Table of Contents

1 Introduction

Recent turmoil in the capital markets has led to a sharp rise in the number of negative rating actions taken by the leading rating agencies, signaling a deterioration in the credit quality of firms affected by adverse economic conditions. These credit quality dynamics highlight the importance of credit migration modeling as an integral part of modern credit risk solutions. Credit migration models play a vital role in pricing credit-risky instruments, as well as in assessing the risk of credit portfolios. Many credit instruments exhibit cash flows contingent on the credit quality of the reference entity, either implicitly (e.g., as in a prepayment option on a loan, which may be exercised in the event of improved credit quality of the borrower), or explicitly (e.g., a loan with grid pricing, which ties the contractual loan spread to the credit quality of the borrower). Pricing such credit-contingent cash flows requires a credit migration model. Further, to estimate the risk associated with a credit instrument or a portfolio, a credit migration model is needed to build a distribution of values at the analysis horizon.

Most practical credit risk models typically employ a discretized representation of credit qualities and utilize a finite-state Markov model, which posits that the probability of migrating to another credit state in the future does not depend on the past.[1] In a discrete-time framework, parameterizing a credit migration model amounts to estimating a Markov transition matrix, populated with migration probabilities over a specified time period. In a continuous-time framework, a generator matrix can be estimated, allowing for forecasts over any time horizon.

Parameterization of a Markovian credit migration model with an appropriate transition matrix is an important step in the specification of a credit risk model. Within the context of a credit portfolio management (CPM) framework, common considerations affecting the choice of a transition matrix include the following.

- Asset class/industry/domicile: A number of studies document the fact that rating transitions vary according to the industry or regional classification of the obligor.[2] Depending on the granularity of the model, portfolio managers may choose to employ a custom transition matrix that is industry- or region-specific or estimated for a particular asset class.
- Through-the-cycle vs. point-in-time: Since agency ratings are considered through-the-cycle measures of credit quality (Cantor and Mann, 2003, Altman and Rijken, 2004), portfolio managers whom want to gain a long-term view of portfolio risk often employ a rating transition matrix as the driver of credit migration. In addition, different agencies' credit ratings may imply different migration patterns, suggesting that managers take into consideration the particular agency data used.[3] We can also use point-in-time measures of credit quality, such as Moody's Analytics EDF credit measures, to estimate a model that provides a more dynamic view of credit migration.

Because credit migration is a fundamental component of a CPM model, it is important to understand the extent to which different migration matrices impact measures of portfolio risk. In particular, if portfolio risk is sensitive to the choice of a specific migration model, we must take care to ensure that portfolio statistics are driven by economically meaningful data, rather than by numerical artifacts or estimation noise.

While the literature on credit migration is plentiful, relatively little work addresses the parameterization of a credit migration model or the impact of credit migration on portfolio risk. Bangia et al. (2002) estimate transition matrices during economic expansion and contraction periods. They show that the loss distribution and economic capital of a synthetic bond portfolio vary significantly in different economic environments. Trück and Rachev (2005a) draw similar conclusions, finding that migration matrices estimated at different points of the business cycle produce significantly different VaR measures for a test portfolio. Further, they find a more dramatic impact on the confidence sets of estimated default probabilities. The experimental results documented in these papers suggest that the migration matrix choice can have a significant impact on the risk analysis of the portfolio.

Realizing the need for a quantitative measure of difference between credit migration matrices, several authors propose distance measures that aim to capture economically significant differences. Jafry and Schuermann (2004) develop a

[1] Recently, in light of empirical evidence suggesting a non-Markovian behavior of credit rating migration, some researchers have considered alternative models; see, e.g., Frydman and Schuermann (2008).

[2] See, for example, Nickell et al. (2000), or Kadam and Lenk (2008).

[3] Livingston et al. (2008) and Hill et al. (2010).

metric for comparing migration matrices based on singular values (closely related to the induced matrix 2-norm) and use it to compare different methods of estimating migration probabilities. They demonstrate that under their proposed metric, statistically significant differences between transition matrices result in material differences in the economic capital of a fictitious bond portfolio. In related work, Trück and Rachev (2005b) suggest a class of so-called directed difference indices as a measure of distance between transition matrices. They demonstrate that their proposed measures exhibit high correlation with differences in credit VaR, indicating that these measures capture economically significant differences in migration dynamics.

This paper complements existing work in three main areas.

- First, we construct a class of low-dimensional migration matrices parameterized by a varying diagonal element, and show that when the default probability term structure remains unchanged, increasing the probability mass assigned to the diagonal results in non-monotonic behavior of the volatility of the value distribution at the analysis horizon. This example supports observations made in Jafry and Schuermann (2004) and Trück and Rachev (2005b), that the diagonal dominance of a matrix, often a used as a visual indicator for the inherent "riskiness" of a migration matrix, is not a sufficient characteristic of migration risk. This example also highlights the interplay between default risk and migration risk implied by a transition matrix.
- Second, we propose a framework for measuring the portion of portfolio risk attributable to credit migration, which facilitates comparison of the impact of different transition matrices on portfolio risk. In contrast to the distance measures outlined in Jafry and Schuermann (2004) or Trück and Rachev (2005b), our framework yields a portfolio-referent distance measure. Further, our framework allows for decoupling of default risk and migration risk. We employ this framework to study the differences between two migration models, one based on agency ratings, and the other based on the EDF credit measure. We find that credit migration explains a significant portion of the risk attributed to a test portfolio, and that migration of point-in-time credit quality, measured by EDF transition rates, accounts for a greater fraction of total portfolio risk when compared with through-the-cycle dynamics reflected by agency rating migrations.
- The third contribution lies in studying the performance of numerical algorithms for generation of high-frequency transition matrices from an annual matrix. We demonstrate that in some cases, application of such methods results in distortions to the implied annual transition probabilities and, consequently, to the calculated risk measures as well.

The remainder of this paper is organized in the following way.

- Section 2 offers an overview of credit migration in the context of credit portfolio analysis.
- Section 3 describes a stylized analytic example that illustrates the behavior of an instrument's volatility in value as a function of migration probabilities.
- Section 4 details a framework to measure migration impact on portfolio risk, which can be applied to compare the migration effects of two different migration models. We use this framework to contrast EDF-based migration with rating-based migration.
- Section 5 discusses the problem of generating high-frequency transition matrices from an annual transition matrix. We describe two numerical algorithms for this purpose and demonstrate their performance when applied to a rating transition matrix and an EDF credit measure transition matrix.
- Section 6 provides concluding remarks.

2 An Overview of Credit Migration and Portfolio Risk

In this section, we review the foundations of a discrete-time, finite-state, homogenous Markov model of credit migration, and relate it to the estimation of a value distribution of a vanilla credit instrument. We present two commonly used transition matrices and highlight the challenges in comparing different views of migration in analyzing portfolio risk. The discussion lays a foundation for the analysis in subsequent sections, where we propose two approaches to studying the impact of credit migration on portfolio risk.

To begin, assume that the universe of credit qualities is represented by a set of K states, where the first state corresponds to the highest credit quality, and the K-th state corresponds to the default state. In a homogeneous Markov chain model, credit migration dynamics are fully specified through a transition matrix containing probabilities of moving from one credit state to another. Let T denote this migration matrix of dimensions $K \times K$ so that $T_{i,j}$ is the probability of moving from state i to state j in one period. Assume the K-th state, corresponding to default of the obligor, is absorbing; once an obligor is in default, it remains in default. Thus, the K-th row of the matrix contains zeros in all entries, except at the K-th entry, which equals one.

A credit migration model facilitates estimating a portfolio value distribution at the analysis horizon. For simplicity, we concentrate on a single vanilla instrument.[4] A discretized value distribution at horizon associates each credit state with a future value, as well as a probability of realizing it. We can calculate the latter directly from the migration matrix; by the Markovian property, if p_t is a vector containing probabilities of realizing each credit state at time t, then $p_{t+1} = T \cdot p_t$

We can calculate the former using risk-neutral valuation techniques; the value of the instrument at a horizon credit state is the expected value of future cash flows, calculated in the risk-neutral probability measure and discounted to the horizon date by the risk-free rate.[5] Figure 1 provides a schematic illustration.

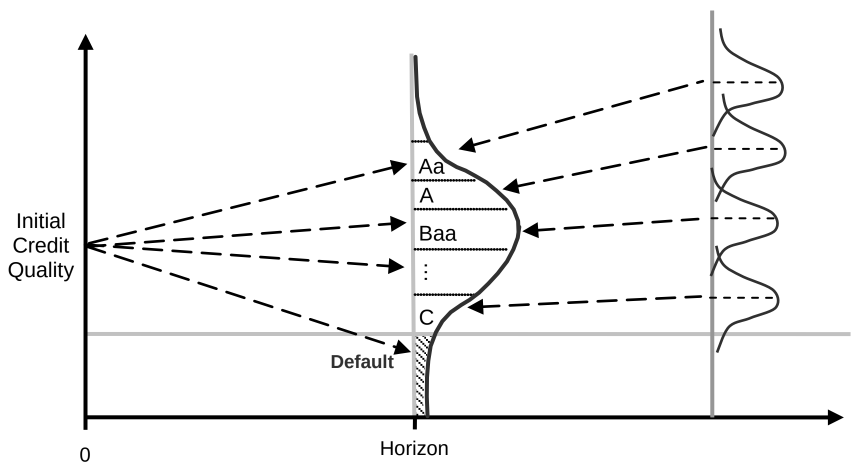

Figure 1 Migration probabilities determine the distribution of values at horizon.

Figure 1 and the preceding discussion hint at the relationship between transition probabilities and the horizon distribution of a vanilla instrument. We can express a discretized value distribution at horizon as two sets of values: a vector of probabilities, p_H, of realizing each credit state at horizon, and a vector of future values, V_H, at each credit state. The vector of horizon values, V_H, is calculated using the PD term structure implied by the transition matrix for each credit state.[6] The probability vector p_H is determined directly from the matrix. For example, in the special case where the time to horizon equals the transition period, p_H equals the transition matrix row that corresponds with the starting credit quality.

A distribution of values at the analysis horizon facilitates the calculation of return and risk measures for the instrument at hand. We calculate the instrument's return to horizon as the expected value of the distribution, divided by its present value. The volatility of the instrument at horizon, often called the unexpected loss (UL), equals the standard deviation of the distribution. Note, that for vanilla instruments, knowledge of the PD term structure implied from the transition matrix is sufficient to determine the expected value of the distribution and, hence, the return to horizon.

[4] Credit migration plays an important role in the valuation of instruments with embedded options. However, the dynamics of option values warrant special consideration, outside the scope of this paper.

[5] See Blöchlinger (2010), for a discussion on the transformation of physical probabilities to risk-neutral probabilities.

[6] Cash flows between the analysis date and the horizon date would also contribute to the value of the instrument at each credit state at horizon.

A credit migration model is often parameterized with transition rates estimated from historical data of credit qualities, measured in a variety of ways, such as using credit agency ratings, institution-specific credit ratings, or a measure of default probability. As mentioned in the introduction, the data source choice depends on the objective of the analysis as well as on the availability of data. We present two examples.

Table 1 displays a transition matrix based on Moody's annual broad rating migrations, published in Tennant et al. (2008). We estimate the matrix using a dataset consisting of U.S. financial and non-financial corporate debt issuers in the period 1970–2007. Using methodology outlined in Metz et al., (2008), we adjust this transition matrix to account for withdrawn ratings and to ensure that default probabilities strictly increase with decreasing credit quality.

Table 1 Adjusted Moody's annual broad rating migration rates, 1970–2007

	Aaa	Aa	A	Baa	Ba	B	Caa-C	Default
Aaa	91.62%	7.70%	0.66%	0.00%	0.02%	0.00%	0.00%	0.01%
Aa	1.13%	91.31%	7.21%	0.27%	0.06%	0.02%	0.00%	0.02%
A	0.07%	2.84%	91.29%	5.14%	0.51%	0.09%	0.02%	0.03%
Baa	0.05%	0.20%	5.15%	88.83%	4.54%	0.81%	0.24%	0.18%
Ba	0.01%	0.06%	0.42%	6.25%	82.94%	8.48%	0.63%	1.20%
B	0.01%	0.05%	0.18%	0.39%	6.21%	81.93%	6.23%	5.00%
Caa-C	0.00%	0.03%	0.03%	0.19%	0.73%	11.22%	68.56%	19.23%
Default	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	100.00%

Table 2 shows the second example, a transition matrix estimated from Moody's Analytics EDF™ (Expected Default Frequency) credit measure.[7] The data set we use for estimating the transition rates listed in Table 2 encompasses EDF data from U.S. based non-financial firms in the years 1990–2007. For the estimation, the space of admissible EDF values was discretized in eight bins, in a manner that ensures historic default probabilities in the estimated matrix match the default probabilities of the corresponding rating categories in the rating migration matrix listed in Table 1 (compare the right column of the matrices listed in Table 1 and Table 2).

Table 2 EDF migration rates, 1990–2007

	Aaa	Aa	A	Baa	Ba	B	Caa-C	Default
Aaa	67.45%	17.15%	10.80%	3.24%	0.86%	0.35%	0.14%	0.01%
Aa	23.19%	35.69%	32.80%	7.07%	0.94%	0.19%	0.09%	0.02%
A	1.94%	11.00%	54.97%	28.15%	3.04%	0.60%	0.26%	0.03%
Baa	0.04%	0.32%	13.38%	61.01%	20.09%	3.73%	1.25%	0.18%
Ba	0.01%	0.01%	0.40%	24.01%	48.16%	19.63%	6.57%	1.20%
B	0.00%	0.00%	0.07%	3.25%	26.69%	38.99%	25.99%	5.00%
Caa-C	0.00%	0.00%	0.05%	0.61%	4.88%	16.58%	58.64%	19.23%
Default	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	100.00%

Visual inspection of the two matrices reveals significantly more dispersion of the EDF-based matrix—the likelihood of a firm remaining in the same EDF bin is much smaller than the likelihood of a firm remaining in the corresponding rating category. We glean further insight from inspecting the PD term structures implied by each matrix.

[7] See Crosbie and Bohn (2003), and Dwyer and Qu (2007).

Figure 2 depicts the 10-year PD term structure implied by the two matrices for each of the seven rating categories. We notice that the EDF matrix implies higher long-term risk of default in the high quality credit categories. This relationship reverses in the lower credit qualities; the EDF-based matrix implies lower long-term default risk than the rating-based matrix. We attribute this latter effect to the mean-reversion behavior of EDF credit measures.[8]

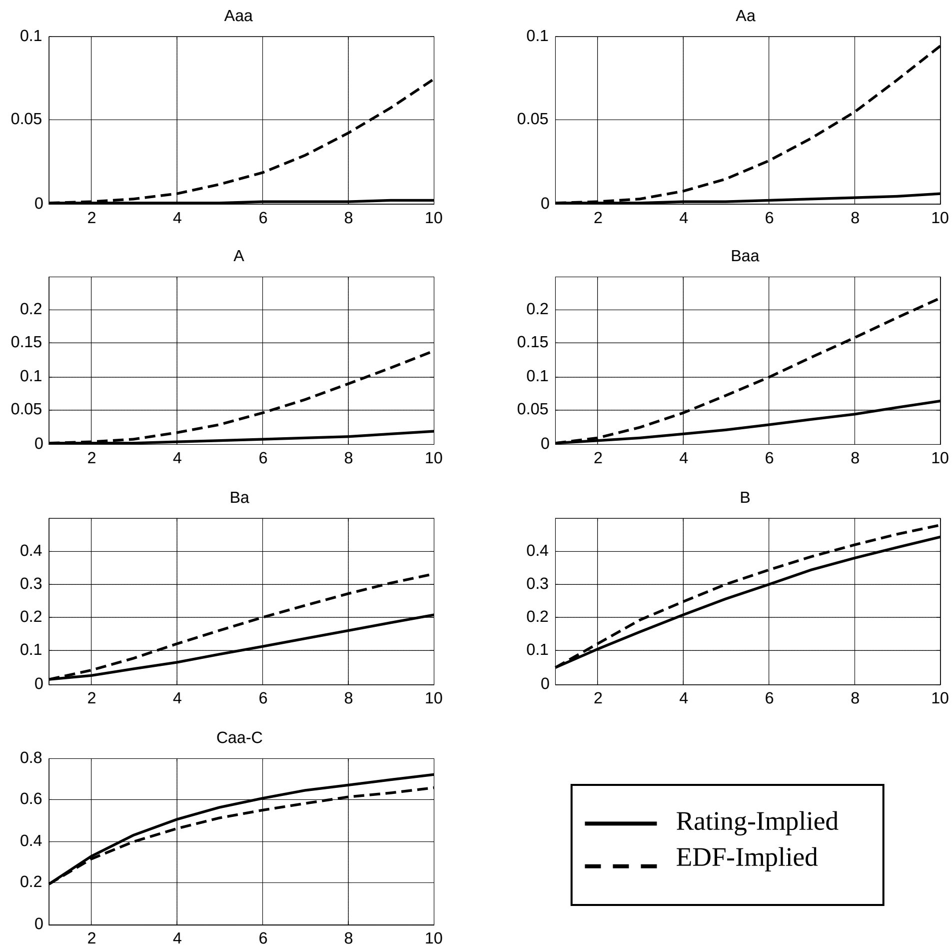

Figure 2 Comparison of implied PD term structures

The differences between a rating-based migration model and an EDF-based model stem from the different economic methodologies driving these two credit quality measures. The EDF model provides a point-in-time measure of default probability, which moves with the credit cycle, as demonstrated in Dwyer and Qu (2007). As such, an EDF-based migration model provides a description of how actual point-in-time credit quality (i.e., default probability) evolves over time. Similar credit cycle movement patterns have also been observed in empirical default rates.[9] In contrast, rating agencies strive to provide investors with a through-the-cycle measure of credit quality, by separating long-term drivers of credit risk from transient ones.[10] This methodology results in stable ratings, with relatively low incidence of migration between ratings. It also suggests that default rates associated with an agency's rating vary over time. For example, Baa-rated firms experienced a 10-fold increase in default rates from 1998 to 2001.[11]

---

[8] For an example, see Figure 5 in Dwyer and Qu, 2007.

[9] See recent works by Bruche and González-Aguado (2009), as well as Bonfim (2009).

[10] Cantor and Mann (2003), Altman and Rijken (2004).

[11] Cantor and Mann, 2003.

In light of the different characteristics exhibited by these two migration matrices, a natural question arises: How do we compare these two views of migration in the context of credit portfolio analysis? One straightforward approach is to analyze the portfolio in question using the valuation methodology illustrated in Figure 1, parameterized with each of the migration matrices, and then compare the analysis results. This approach is popular among credit practitioners because of its simplicity. However, this technique poses a significant challenge when comparing risk measures obtained in the different analyses. Recall that the return of a vanilla credit instrument can be inferred from the implied term structure of default probabilities for its starting credit quality. Different migration matrices imply, in general, different sets of PD term structures. Thus, analyzing a portfolio using disparate views of migration generally yields different portfolio returns. Comparing risk measures without fixing the portfolio's return is undesirable, since it is difficult to determine whether or not elevated risk levels occur because of elevated returns or because of other effects. An alternative approach that circumvents this challenge modifies each of the transition matrices to imply a PD term structure specified exogenously.[12] A drawback to this approach is that modifying the PD term structure implied by a matrix may distort the migration dynamics it implies.

In the following two sections, we present two approaches addressing this challenge. First, we construct a class of low-dimensional transition matrices with the PD term structure implied by each matrix fixed. This setup allows us to study how varying migration dynamics affect risk in a simplified setting. Second, we present an approach for isolating the component attributed to migration risk from the total risk of a credit portfolio. We demonstrate the applicability of this approach by comparing the migration risk implied by the rating transition matrix listed in Table 1 with the EDF transition matrix listed in Table 2 when applied to a test portfolio.

3 Migration Probabilities and Unexpected Loss

When considering migration probabilities' effect on instrument or portfolio risk, common reasoning suggests that a more dynamic transition matrix (i.e., a matrix that implies greater likelihood of moving away from the current credit state) results in greater horizon values dispersion. Similarly, a more dynamic matrix implies a higher probability of realizing large credit quality changes. Combined, these effects imply greater uncertainty about the instrument's future value, which translates to increased risk measures (e.g., standard deviation). As it turns out, the interplay between values and probabilities, which collectively determine the shape of the value distribution, can be more subtle. To demonstrate this behavior, we construct a class of transition matrices parameterized by a diagonal element. Above, we mentioned a key property of this class of matrices; starting from a particular credit quality, the PD term structure implied by any matrix in this family of matrices is fixed and independent of the parameterization, which allows for a meaningful comparison of risk levels implied by different matrices in this class. By linking transition probabilities to an instrument's risk, we can characterize the behavior of the unexpected loss of the distribution at horizon as the magnitude of the diagonal varies.

In detail, we consider a zero-coupon bond, with a face value of \$1, maturing in five periods. We set $K = 4$, so that the universe of credit qualities consists of three no-default states, labeled A, B, and C, and a default state, labeled D. At the time of analysis, the initial credit state of the bond is B, and the recovery in the event of default equals zero (there is no uncertainty around the recovery value).

[12] For an example, see Bluhm and Overbeck (2007).

The following annual transition matrix governs credit migration:

$$T = \begin{bmatrix} x & (1-q-x)/2+\Delta & (1-q-x)/2 & q-\Delta \\ (1-q-x)/2 & x & (1-q-x)/2 & q \\ (1-q-x)/2 & (1-q-x)/2-\Delta & x & q+\Delta \\ 0 & 0 & 0 & 1 \end{bmatrix}, \tag{1}$$

where $q > 0$, $\Delta \in (0, q)$ are fixed, and *x* varies in the interval $(0, 1-q-\Delta)$. The variable *x*, which determines the probability of remaining in the same no-default state during the next period, controls the dynamism of the matrix *T*; a high x value results in a "sticky" matrix, with low likelihood of migration from the current credit state, whereas, a low x value results in dynamic migration, with a high likelihood of migration from the current credit state to a different state. Below, we express the variance of the value distribution of the bond at a future date as a function of *x*, which enables us to study the behavior of the uncertainty around the horizon value as we move from a very "sticky" transition matrix to a very dynamic transition matrix, while maintaining a fixed return level.

Figure 3 displays the PD term structures implied by a matrix in this class in the first 10 periods, when $q = 10\%, \Delta = 5\%,$ and $x = 75\%$.

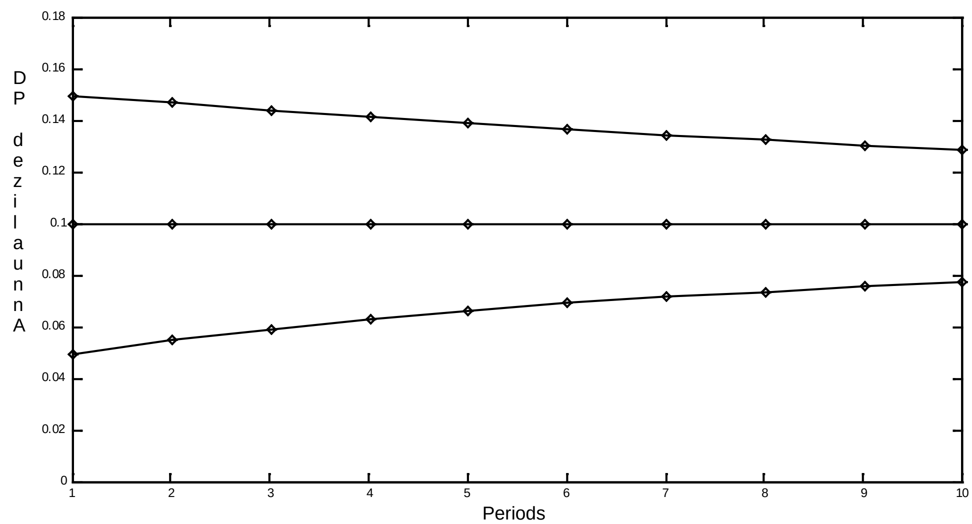

Figure 3 Implied PD term structures for each no-default credit state, when $q = 10\%$, $\Delta = 5\%$, and $x = 75\%$

The figure highlights two important properties of the transition matrix specified in Equation (1). First, the annualized PD term structure corresponding to starting credit state B is flat and equal to 10%. In Appendix A, we show that this property holds in a more general setting as well, namely, for all admissible *q* and Δ , the PD term structure corresponding to credit state B is flat and independent of *x*. As we discuss in the previous section, since the present value (PV) of the bond and its return to horizon are functions of this PD term structure, these quantities do not depend on x either, allowing for a comparison of risk levels obtained with different values of x, while maintaining a fixed return. A second property, evident in the figure, is the mean-reverting behavior of the credit quality associated with states A and C. Specifically, the PD term structure associated with state A slopes upward, and the PD term structure of state C slopes downward. Mean-reversion of credit quality is a property often observed in empirical data.[13]

---

[13] For an example see, Dwyer and Qu, 2007.

To arrive at an expression for the unexpected loss at the analysis horizon, set at the end of the first period, we employ a valuation framework in the spirit of the schematic shown in Figure 1, simplified to maintain analytical tractability. Specifically, we set the risk-free rate to zero and assume that the risk-neutral measure equals the physical probability measure (i.e., the market risk premium equals zero). In this setup, the zero-coupon bond's value at each credit state at horizon equals the sum of expected cash flows past horizon. Formally, the vector of horizon values at each credit state is given by

$$V_H = T^{n-1} \cdot V_n \tag{2}$$

where n denotes the number of periods to maturity and $V_n = \begin{bmatrix} 1 & 1 & 1 & 0 \end{bmatrix}'$ is the vector of bond values at maturity (the value at maturity equals one if the bond has not defaulted and zero if it has defaulted). Since we set the horizon one period into the future, the probability of realizing credit state j at horizon is given by the $(2, j)$-th element of the transition matrix T (corresponding to initial credit state B). Therefore, the variance of the value distribution at horizon is given by

$$\sigma_H^2 = \sum_{j=1}^{4} T_{2,j} \cdot \left(V_H(j)\right)^2. \tag{3}$$

For fixed q, Δ, the expression in Equation (3) yields a polynomial in *x* of degree $2(n-1)-1$. For the case $n=5$, illustrated in Figure 4, the expression for the variance as a function of *x* is

$$\sigma_H^2 = 10^{-2} \cdot \left(3.92 + 0.135x - 0.228x^2 + 0.3x^3 + 1.31x^4 - 3.33x^5 + 4.27x^6 - 2.85x^7\right). \tag{4}$$

Figure 4 displays the variance of the horizon value distribution as a function of *x* when $q = 10\%$ and $\Delta = 5\%$.

Inspection of the figure reveals a non-monotonic relationship between the unexpected loss and the likelihood of moving away from the current state. As we discuss previously, this behavior may seem surprising; one expects the graph to exhibit a downward sloping trend. However, careful examination of the dynamics at play reveals two countervailing effects contributing to the variance at horizon. First, as *x* increases, the probability of realizing state A or C at horizon falls, that is, the probability mass becomes more concentrated around state B. At the same time, the horizon value at state A increases, and the horizon value at state C decreases, while the value at state B remains unchanged. Therefore, values at horizon become more dispersed as *x* increases. Since these two effects pull the unexpected loss in opposite directions, the net effect is the non-monotonic behavior exhibited in Figure 4.

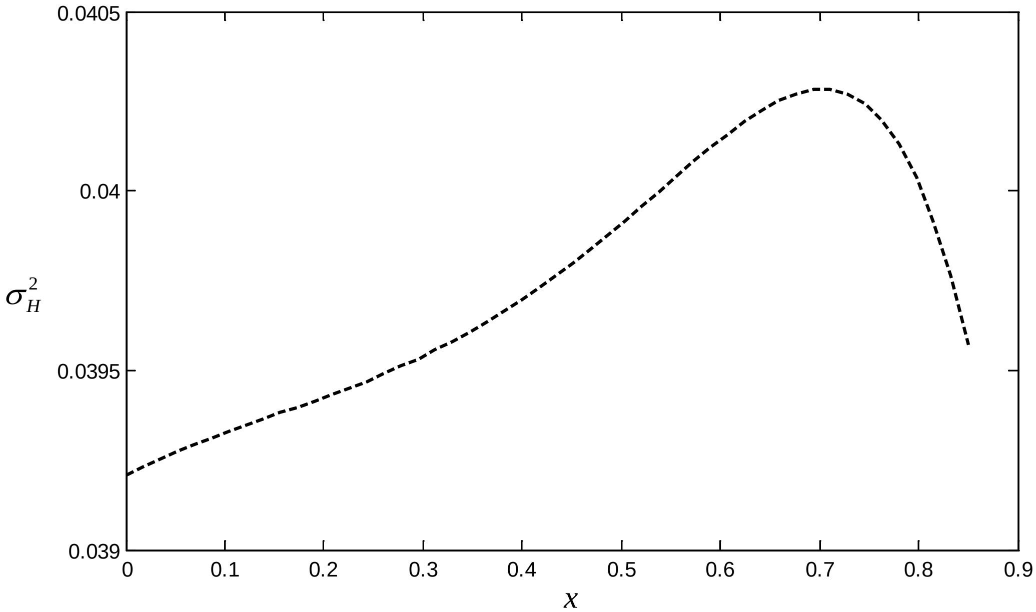

Figure 4 Variance of the horizon value distribution of a zero-coupon bond maturing in five periods as a function of x, when $q = 10\%$, $\Delta = 5\%$.

To summarize, we present counter evidence to the assertion that more dynamic migration implies greater risk. In the stylized setting we consider, modifying the dynamism of the matrix, while maintaining a fixed PD term structure, affects the values at horizon and the probabilities of realizing these values in opposing directions. Consequently, matrices at the extreme ends of the dynamic spectrum (i.e., very dynamic or very "sticky" matrices) imply lower variance at horizon than a matrix in the mid-range of the spectrum. While we construct this example in a very controlled setting, it nevertheless highlights important behavior that should be considered when examining the effect of different migration matrices on the risk of an individual instrument or a portfolio.

4 Measuring the Impact of Different Migration Matrices on Portfolio Risk

In the previous section, we demonstrate how controlled variation of the transition probabilities affects the variance of the value distribution of a zero-coupon bond at horizon in a stylized setting. A key construct in the analysis is the fact that the PD term structure associated with the instrument and, consequently, the expected value of the instrument at horizon, does not change as we modify transition probabilities, allowing for a comparison of risk levels with the return level fixed. When comparing the impact of different migration matrices in a general setting, it is difficult to control for each instrument's associated PD term structure and, thereby, maintain a constant level of return, without introducing distortions to the transition probabilities. In this section, we describe a framework for measuring the impact of a credit migration matrix on portfolio risk in a manner that accounts for PD term structure effects.

We utilize the Moody's Analytics RiskFrontier® credit portfolio model for the analysis. This model employs a bottom-up approach that associates a horizon value distribution for each instrument, utilizing a valuation framework similar to the approach described in Section 2. The model then constructs a portfolio value distribution using Monte Carlo simulation, taking into account correlations between portfolio instruments.[14]

[14] Levy (2008) provides a detailed exposition.

Within the context of a credit portfolio model, we propose the following approach for measuring the impact of credit migration on portfolio risk. We analyze the portfolio in question twice, once using a credit migration model and a second time using a default/no-default (D/ND) model. We can attribute the difference between the two resulting sets of risk measures to migration effects. The intuition is clear: an analysis of the portfolio without credit migration is the base case, that is, only a single no-default state exists at horizon. We then introduce credit migration into the model, keeping everything else fixed. Thus, we expect credit migration to drive risk measure changes. An important attribute of this approach is the fact that the present value and return to horizon of the portfolio are equal in both analyses, since the PD term structure associated with each instrument does not change. Consequently, this approach achieves two goals. First, it isolates the contribution of migration to the portfolio's risk measures. Second, it allows for a comparison of the impact of different migration matrices on portfolio risk.

Details of the default/no-default analysis follow: The model carries out the valuation of each instrument using the period-by-period default probability term structure implied from the migration matrix, yielding two values: an expected value conditional on default and an expected value given no-default (recall from the previous discussion that for vanilla debt instruments, we can infer the return to horizon from the counterparty's PD term structure). The Monte Carlo simulation procedure then associates one of these two values to the instrument at each trial, depending on whether the asset return has fallen above or below the default threshold. [15] This procedure mutes migration effects on portfolio risk. Migration probabilities still play a role in the resulting values at horizon. The risk-neutral PD term structure, used to calculate an expected value given no-default, is determined by transition probabilities in the matrix.

The portfolio used in this analysis (as well as all other analyses carried out in the paper) is described in a report by IACPM/ISDA (2006). This synthetic portfolio mimics a typical bank portfolio, designed for the purpose of comparing different models of economic capital. It consists of two loans to each of 3,000 obligors, yielding 6,000 floating rate term loans with annual coupons, for a total commitment of \$100 billion. Each obligor is assigned a credit rating on a seven-category scale, which determines its stand-alone credit risk. Loans are assumed non-prepayable.

We conduct the above analysis on the IACPM test portfolio using the Moody's rating migration matrix in Table 1. Results are displayed in Table 3. As indicated above, the present value of the portfolio and its return to horizon are the same, whether the analysis is carried out using the D/ND model or via the migration model. The change in the last two measures in Table 4, which offers two views of portfolio risk, is significant. Unexpected Loss (UL) equal to the standard deviation risk of the portfolio expressed as a percentage of its present value, is 65% greater in the migration analysis when compared with the D/ND analysis. Portfolio Economic Capital (EC), formally defined as the loss level in excess of Expected Loss that will not be exceeded with probability $1-\alpha$, has similarly increased by 37%.[16] Alternatively, we can say that migration accounts for 39% of portfolio volatility and 27% of economic capital. In other words, a sizeable portion of portfolio risk (measured in terms of UL or EC) can be attributed to migration effects.

Table 3 Impact of migration on portfolio risk and return, using a rating-based transition matrix

	D/ND Mode	Migration Mode	% Increase
MTM Value	\$100,340,446,638	\$100,340,446,638	0%
Expected Spread	0.506%	0.506%	0%
Expected Loss	0.658%	0.658%	0%
Unexpected Loss	0.697%	1.150%	65.11%
Economic Capital	4.587%	6.290%	37.11%

This approach facilitates comparison of different migration models, of interest to portfolio managers wishing to understand how different views of migration translate into different views of portfolio risk. In this spirit, we utilize this procedure to compare the Moody's ratings transition matrix in Table 1 with the EDF transition matrix in Table 2, when applied to the IACPM test portfolio. To quantify the impact of EDF-based migration on the risk of the IACPM test portfolio, we repeat the analysis of the portfolio using the EDF-based transition matrix. Results are displayed in Table 4.

[15] In this exercise, uncertainty in recovery is taken into account in the default/no-default analysis (as well as in the analysis with migration).

[16] In this analysis and later analyses presented in this paper, $\alpha = 0.1\%$.

Table 4 Impact of migration on portfolio risk and return, using EDF-based transition matrix

	D/ND Mode	Migration Mode	% Increase
MTM Value	$99,162,223,694	$99,162,223,694	0%
Expected Spread	0.753%	0.753%	0%
Expected Loss	0.670%	0.670%	0%
Unexpected Loss	0.702%	1.434%	104.13%
Economic Capital	4.586%	7.141%	55.71%

Inspecting portfolio economic capital and unexpected loss suggests that the impact on portfolio risk is quite dramatic; unexpected loss increases by 104% and capital by 56%. Put differently, migration accounts for 51% of portfolio volatility and 35% of economic capital. We note that direct comparison of the risk measures obtained from this analysis with the results listed in Table 3 is not meaningful, since portfolio returns are different. However, comparison of the relative increase in risk is valid, since we keep the return fixed when moving from D/ND analysis to migration analysis. Performing the latter comparative analysis, we see that when using the rating migration matrix, UL increases by 65% and EC increases by 37%, suggesting that the impact of EDF-based migration on risk measures is significantly larger than the increase observed in the rating-based analysis.

We summarize this section by reiterating the two important findings documented above. First, credit migration explains a large portion of the overall risk of the IACPM test portfolio (measured in terms of standard deviation as well as capital). Second, EDF-based migration, which describes the evolution of point-in-time credit quality, accounts for a greater fraction of total portfolio risk, when compared with ratings-based migration.

5 High-frequency Transition Matrices

In the discussion so far, we have focused on a portfolio containing instruments with annual cash flows, analyzed at a one-year horizon. This setting allows for an analysis employing solely an annual credit migration matrix. In practice, most credit instruments have cash flow streams specified at varying frequencies, which may not correspond with the one-year period typically used to estimate the migration matrix. A simple approach that addresses this issue involves modifying an instrument's cash flow structure to match the frequency of the migration matrix. For instance, with only an annual transition matrix available, we can modify instruments to have annual principal and interest cash flows. While this approach may be acceptable when the modification to the cash flows is minor, it will degrade the granularity and accuracy of the subsequent analysis when the change to the cash flow structure is substantial. Thus, in a general setting, this solution does not circumvent the need for an estimate of higher frequency transition probabilities.

In the context of a Markovian modeling framework, two formulations can be utilized to compute transition probabilities to arbitrary time horizons (Jarrow et al., 1997). The discrete-time formulation (the focus of this discussion so far) calls for a transition matrix for each time frequency needed (or one transition matrix for the finest frequency needed, assuming that other frequencies are integer multiples of the finest frequency). The continuous-time formulation specifies a time-homogeneous Markov process in terms of a generator matrix Q. It allows for a transition matrix $T(t)$ to horizon t to be expressed as the exponential of the generator matrix, $T(t) = \exp(Qt)$.

When sufficient historical credit transition data are available, it may be possible to estimate transition probabilities from the data. A transition matrix can be estimated for each time horizon needed in the analysis using the cohort method, which directly estimates empirical transition rates. Alternatively, a generator matrix for the continuous Markov process can be estimated from the data (Albert, 1962, Andersen et al., 1985). Lando and Skødeberg (2002) and Jafry and Schuermann (2004) argue that estimation of a generator matrix leads to superior estimation results, compared with the cohort approach. Regardless of the estimation method employed, if the available transition data is not sufficiently rich, estimation noise can lead to spurious transition probabilities and, consequently, unreliable analysis results.

Often, historical transition data is not readily available to credit practitioners, who obtain transition matrices of fixed-time periods (typically annual) from third parties (e.g., a credit rating agency).[17] In this case, estimates of transition probabilities over shorter horizons must be obtained from the specified transition matrix T. In the discrete-time setting, to obtain a transition matrix for a fraction $1/n$ of the period of *T*, we seek a stochastic matrix *X* that satisfies $X^n = T$. In he continuous-time setting, we seek a generator matrix Q with zero row sums and non-negative off-diagonal entries, such that $T = \exp(Q)$ (known as the embedding problem). Often, an additional requirement posed is that the PD column in the resulting short-horizon transition matrix be monotonically increasing, in line with the economic reasoning that the PD of a higher-quality state be lower than the PD of lower-quality state. This constraint relates closely to the stochastic monotonicity condition Jarrow et al. (1997) impose. Note that both the discrete formulation (finding the n^{th} -root) as well as the continuous formulation (the embedding problem) face issues of existence and uniqueness of solutions.[18]

The remainder of this section focuses on the discrete-time formulation and reviews two common algorithms for generating high-frequency matrices. We study the performance of the algorithms and analyze the implications on portfolio analysis results. In the following, we assume that an annual transition matrix is available, the most common case in practice. The discussion generalizes easily to the case where a matrix of a different tenor is available.

5.1 Series Approximation

Direct computation of the n^{th} -root of the annual transition matrix *T*, by computing the eigenvalue decomposition $T = VDV^{-1}$ (assuming *T* is diagonalizable) and letting $X = VD^{1/n}V^{-1}$, generally fails, since the eigenvalues of T may be complex valued. In addition, X is not guaranteed to be stochastic. A common approach that avoids the issue of a complex solution and constrains the solution as stochastic is based on a Taylor series expansion of the matrix root function (Israel et al., 2001). In detail, the m^{th} -order Taylor expansion of $T^{1/n}$ around the identity matrix is given by

$$S_m = I + \sum_{i=1}^{m} a_i (I - T)^i \tag{5}$$

where I denoted the $K \times K$ identity matrix. The series coefficients are given by

$$a_i = \frac{(-1)^i \cdot \frac{1}{n} \cdot (\frac{1}{n} - 1) \cdots (\frac{1}{n} - i + 1)}{i!} \tag{6}$$

Of course, the partial sum S_m is not guaranteed to yield a stochastic matrix. To make the resulting matrix stochastic, we set negative entries in the matrix to zero, and each row is normalized to sum to one. Note that it is possible that the resulting matrix S_m does not have a monotonically increasing PD column.

We demonstrate this method's performance by applying it to the rating transition matrix in Table 1 and the EDF migration matrix in Table 2 to generate monthly transition matrices (i.e., $n = 12$). Results are depicted in Figure 5 and Figure 6, respectively (element-wise differences in numeric form are listed in Appendix C, Table 7 and Table 8, respectively). Each plot shows the annual transition probabilities from a starting credit rating to any other credit rating on a log scale, obtained by raising the series approximation to the 12th power. The original annual transition probabilities used to generate the high-frequency matrix are overlaid. Inspecting the results reveals that, for the rating-based transition matrix, the Taylor series method produces a small average absolute difference between elements of the two annual transition matrices of 6.8×10^{-6}. This approach does not produce satisfactory results when applied to the EDF-based matrix, as evident in Figure 6. The average absolute difference between elements of the two annual matrices

[17] Credit quality measures are based on observable data, which may be available at discrete points in time. For example, data used for measuring credit quality of private firms is typically reported at an annual (or sometimes quarterly) frequency.
[18] See Israel et al., 2001.

equals 0.45%. Moreover, in multiple instances, the difference between the implied annual transition probabilities and the original probabilities is significant. For example, the approximation implies a one-year PD of about 2bp for Aaa-rated borrowers, roughly a twofold increase when compared with the PD in the original matrix, approximately 1bp.

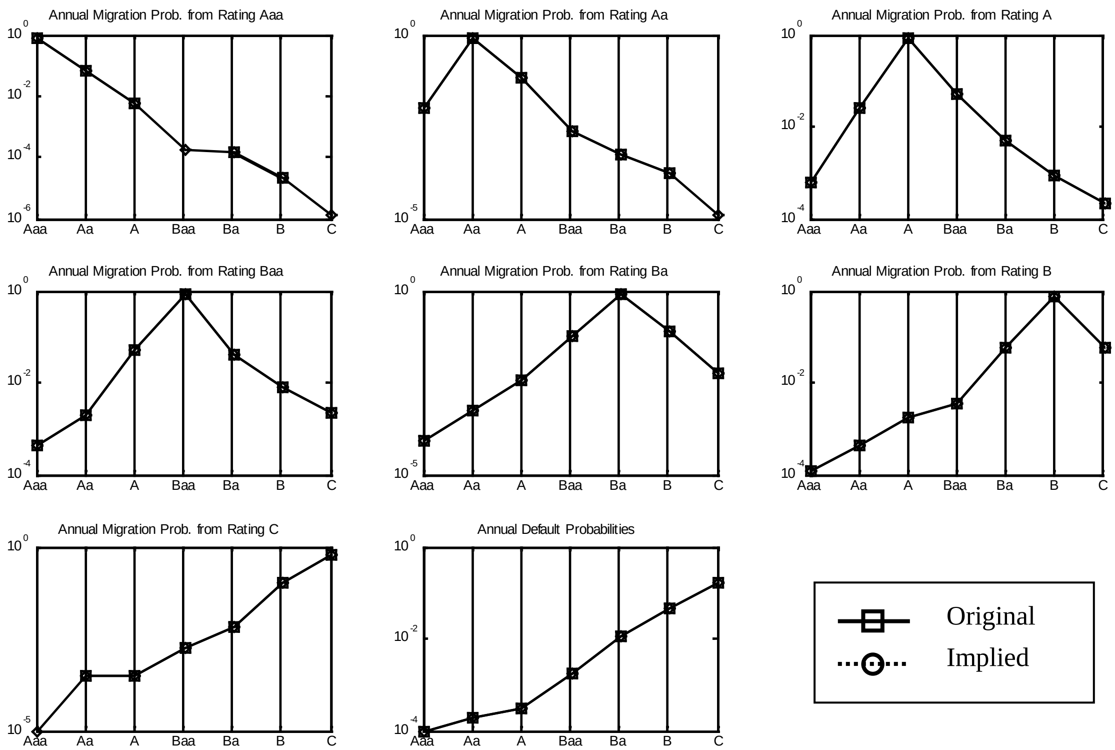

Figure 5 Comparison of Moody's annual rating transition probabilities with the annual transition probabilities implied from the monthly Taylor series approximation.

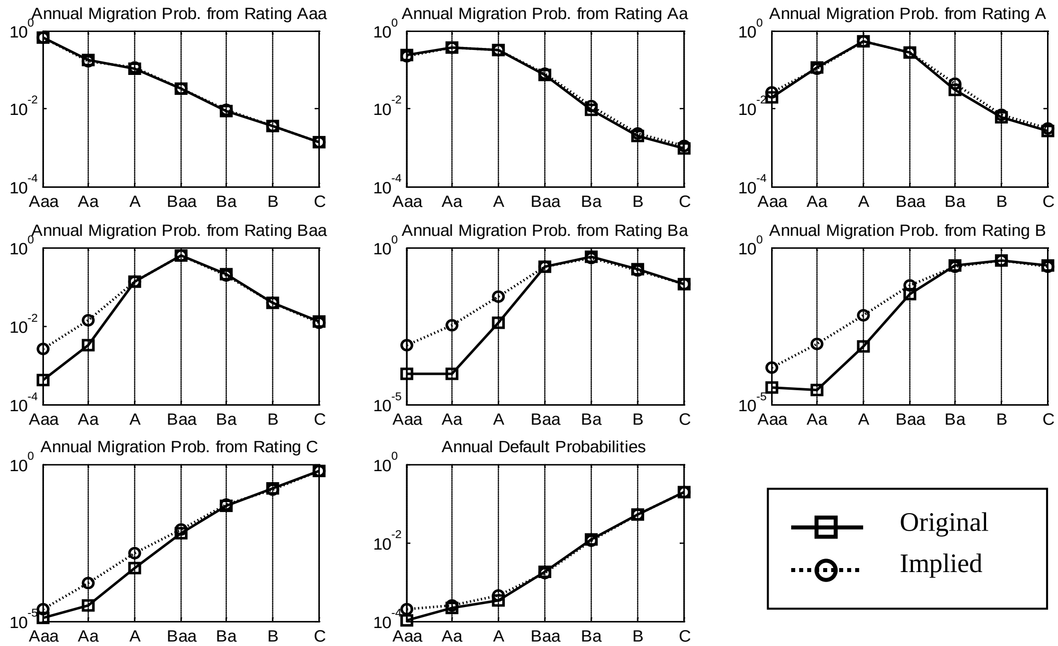

Figure 6 Comparison of EDF-based annual transition probabilities with the annual transition probabilities implied from the monthly Taylor series approximation.

5.2 Eigenspace Optimization

An alternative approach to the problem of generating a high-frequency transition matrix is to use optimization techniques to look for the matrix that, when raised to the n^{th} power, is closest to the original annual matrix in some distance measure (an analog formulation can be posed in the continuous-time setting). This approach guarantees that, when an exact solution does not exist, the resulting solution is the best available approximation (with respect to some specified criteria). An additional benefit to formulating the problem as an optimization problem is the allowance for explicit constraints on the solution, e.g., that the resulting matrix be stochastic.

Direct formulation of the matrix root problem, namely, minimizing $\left\| X^n - T \right\|$, results in a high-order nonlinear optimization problem that is difficult to solve. Instead, Crommelin and Vanden-Eijnden (2006) suggest a method of finding a matrix X that yields the closest eigenspace representation of the n^{th} root. To that end, recall that we can express the n^{th} root of the matrix T as $Q\Lambda^{1/n}Q^{-1}$, where $T = Q\Lambda Q^{-1}$ is the eigenvalue decomposition of T. Thus, we can write the optimization problem to find X as

$$
\begin{aligned}
&\min_{X \in \mathbb{R}^{K\times K}} \sum_{k=1}^{K} \left\| X v_k - \lambda_k^{1/n} v_k \right\|_F^2 + \left\| v_k^{-T} X - \lambda_k^{1/n} v_k^{-T} \right\|_F^2 \\
&\text{subject to} \qquad x_{ij} \ge 0, \quad i, j = 1,\dots,K \\
&\qquad\qquad\qquad \sum_{j=1}^{K} x_{ij} = 1, \quad i = 1,\dots,K \\
&\qquad\qquad\qquad x_{i+1,K} \ge x_{i,K}, \quad i = 1,\dots,K-1
\end{aligned} \tag{7}
$$

where (λ_k, v_k) are eigenpairs of the matrix T. In other words, the optimization solves for the matrix X whose eigenspectrum is closest to that of the matrix root of T, subject to the matrix X being stochastic with monotonically increasing default probabilities. Note that, while the terms inside the norm may well be complex, the overall value of the objective function will always be real, due to the use of the quadratic Frobenius norm. In fact, this optimization problem can be written as a quadratic program, and solved using a general-purpose quadratic programming solver.

We apply this technique to the rating transition matrix in Table 1 and the EDF migration matrix in Table 2 to generate monthly transition matrices. Overall, we see results similar to those obtained with the Taylor series method, with slightly lower error estimates. When applied to the rating transition matrix, the optimization approach produces a good approximation, with the average error between elements of the annual transition matrix equal to 6.76×10^{-6}. When applied to the EDF transition matrix, we see the optimization approach does not succeed in finding a satisfactory approximation; the average absolute difference between the implied annual matrix and the original matrix is 0.42%, slightly lower than the error produced by the Taylor approximation. Figure 7 gives a visual comparison of the transition probabilities at each credit quality, revealing that, in many cases, the error level is significant. For instance, the annual default probabilities implied by the optimization solution are lower than the corresponding probabilities PDs in the annual matrix by non-trivial amounts. Element-wise differences in numeric form are listed in Table 9 in Appendix C.

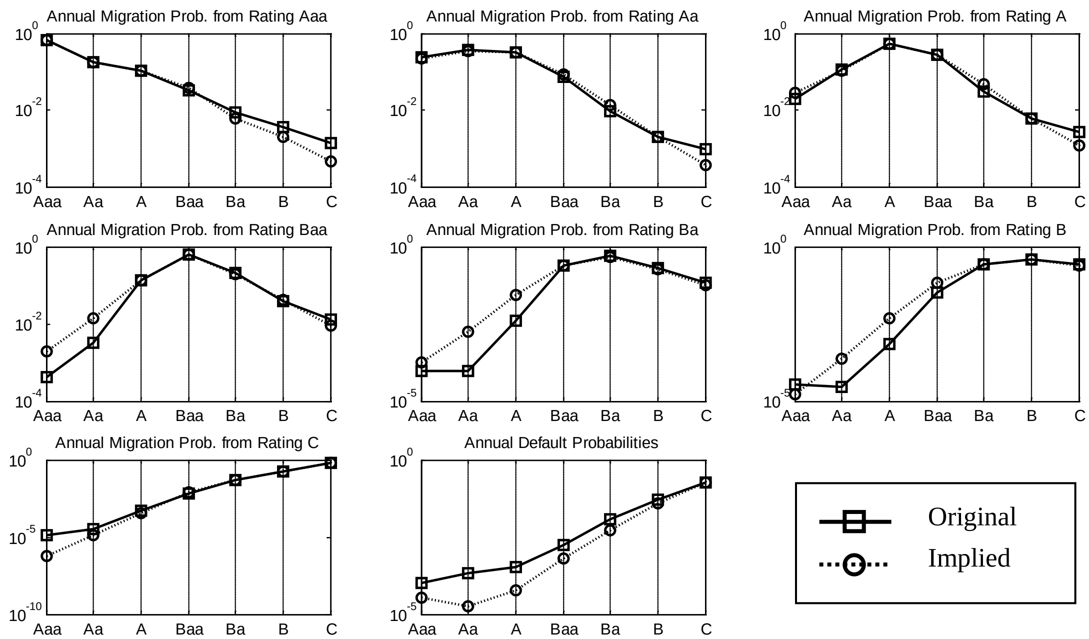

Figure 7 Comparison of EDF-based annual transition probabilities with the annual transition probabilities implied from the monthly eigenspace optimization solution.

To understand the impact of the approximation error on portfolio analysis results, we analyze the IACPM test portfolio with the annual transition matrices implied from the monthly eigenspace solution for the rating-based and EDF-based transition matrices. The summary statistics appear in Table 5 and Table 6, respectively, each contrasted with the risk and return measures obtained using the original annual transition matrix. The results of the rating-based analysis are nearly identical to the original results, since the implied annual transition rates conform very closely to the annual matrix. However, the results of the EDF-based analysis suggest a significant reduction in portfolio risk caused by the errors in the approximation. We attribute this finding to the fact that the eigenspace solution implies higher annual probabilities of migrating from low credit states to high credit states and, consequently, lower annual default probabilities for most credit qualities. The magnitude of these differences, attributed to approximation error and do not carry any economic significance, places the usability of the monthly solution in question.

Table 5 Comparison of portfolio-level output obtained using the original annual Moody's rating transition matrix and the annual transition matrix implied from the monthly eigenspace solution

	Original	Monthly-implied	Pctg. Difference
MTM Value	$100,340,446,638	$100,340,286,525	0.00%
Expected Spread	0.506%	0.507%	0.01%
Expected Loss	0.658%	0.658%	0.00%
Unexpected Loss	1.150%	1.150%	0.00%
Economic Capital	6.290%	6.291%	0.01%

Table 6 Comparison of portfolio-level output obtained using the original annual EDF transition matrix and the annual transition matrix implied from the monthly eigenspace solution

	Original	Monthly-implied	Pctg. Difference
MTM Value	$99,162,223,694	$99,901,869,139	0.75%
Expected Spread	0.753%	0.647%	-14.14%
Expected Loss	0.670%	0.546%	-18.52%
Unexpected Loss	1.434%	1.248%	-12.93%
Economic Capital	7.141%	5.890%	-17.52%

The failure of both algorithms to produce a monthly matrix that agrees with the EDF-based annual transition probabilities suggests that, quite likely, a good approximation does not exist for this particular EDF transition matrix. We stress that this does not imply that EDF migration dynamics cannot be captured using a time-homogeneous Markov process. Recall that the annual EDF transition rates displayed in Table 2 are estimated using a particular discretization of the space of EDF credit measures into eight buckets, which correspond to the eight rating categories (and annual PDs) of the rating transition matrix listed in Table 1. While this discretization yields an annual transition matrix that does not lend itself to a high-frequency decomposition, other discretizations exist that do not suffer from this issue.

To demonstrate this latter assertion, we apply an alternative discretization scheme to estimate annual EDF transition rates, and use the eigenspace optimization technique described in Section 5.2 to obtain monthly transition rates. Figure 8 provides a visual comparison of the annual EDF transition rates with annual probabilities implied from the generated monthly transition matrix (Table 10 in Appendix C lists the element-wise differences in numeric form). As evident in Figure 8, the monthly-implied probabilities conform closely with the estimated annual probabilities. The average absolute error between elements of the two matrices is 1.66×10^{-4}. Appendix B provides details on the numerical optimization procedure that yielded this alternative discretization. We note that in case the chosen discretization does not yield a good, high-frequency approximation, one can avoid the embedding problem by directly estimating a generator matrix for the continuous-time Markov process from historical transition rates ().[19] In general, the annual transition probabilities implied by the estimated generator matrix will not agree with the annual transition rates estimated directly using the cohort method.

[19] Albert (1962), Andersen et al. (1985), Lando and Skødeberg (2002), Jafry and Schuermann (2004).

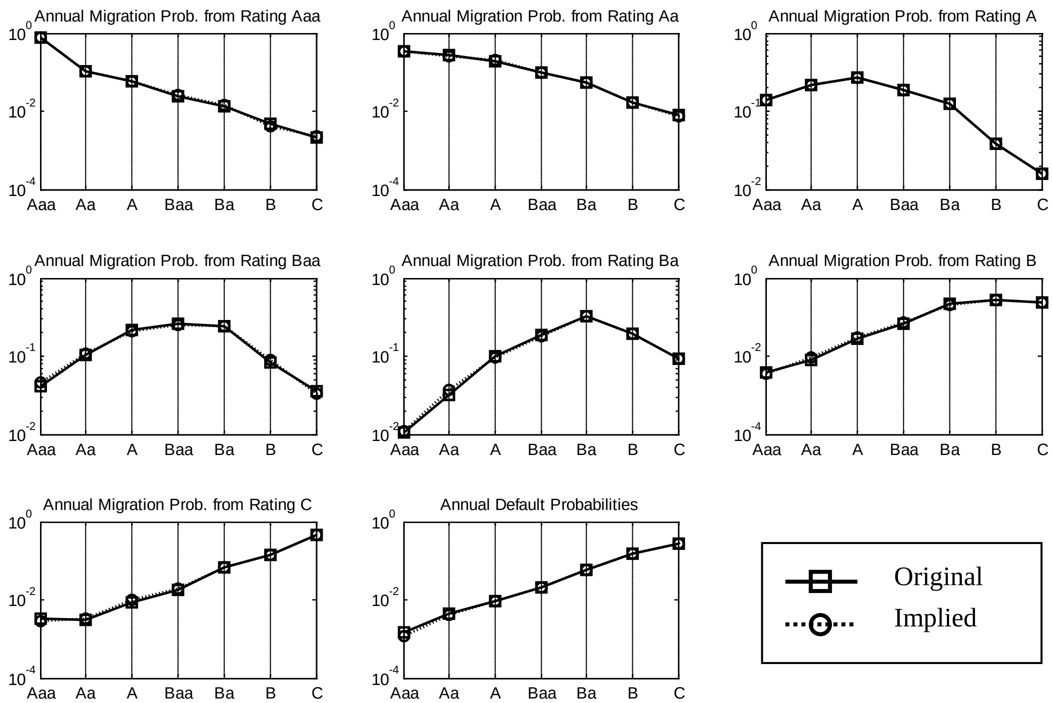

Figure 8 Comparison of EDF-based annual transition probabilities with the annual transition probabilities implied from the monthly eigenspace optimization solution, using an alternative EDF discretization scheme (see Appendix B for details).

6 Conclusion

In this paper, we study the impact of credit migration on credit portfolio risk and address some of the challenges a risk manager may face when specifying a transition matrix for portfolio analysis. The discussion highlights several important points to take into consideration in the parameterization process.

First, we demonstrate that examining a transition matrix's diagonal weights is insufficient to assess the migration risk it implies. In particular, we employ a stylized analytic setup to show that, when controlling for PD term structure effects, higher likelihood of moving away from the current credit state does not necessarily imply greater risk (as measured by the volatility of the value at horizon).

In the context of a portfolio analysis, we outline a framework that allows for assessing the portion of risk attributed to migration. Applying this framework to a rating transition matrix and an EDF transition matrix, we demonstrate that for a synthetic test portfolio, a point-in-time view of credit migration accounts for a greater fraction of total portfolio risk, when compared with a through-the-cycle view of migration.

Finally, we discuss the problem of estimating transition matrices to arbitrary time horizons from a given annual matrix. We review two methods for generating high-frequency transition matrices and document a case where the numerical artifacts caused by these methods results in material deviations of portfolio analytics.

In conclusion, our findings suggest that migration explains a large portion of the risk of a credit portfolio, and that the choice of a migration matrix has a significant impact on analysis results. These findings shed new light on the importance of proper specification of migration matrices in addition to the identification of other risk inputs such as PD and lost given default. Further, these findings suggest that care must be taken to ensure that the analysis results are not driven, in part, by numerical artifacts or estimation noise.

Appendix A

In this appendix, we verify the assertion made in Section 2 regarding the PD term structure implied by the transition matrix as defined in Equation (1). In the following, we denote by $P\left(S^i \to T^j\right)$ the probability of migrating from state S at period i to state T at period j, where $S,T \in \{A,B,C,D\}$. The following Lemma demonstrates that, when starting from credit state B, the PD term structure implied by the transition matrix in Equation (1) is flat.

Lemma 1

Suppose that at the beginning of the first period, the obligor is in credit state B. Then, for period $n > 0$,

a. $P\left(B^0 \to A^n\right) = P\left(B^0 \to C^n\right)$;

b. The annualized n-period default probability is equal to q.

Proof

The proof of the lemma is by induction. For $n=1$, we verify (a) and (b) by inspection. For $n=k$, assume that $P\left(B^0 \to A^k\right) = P\left(B^0 \to C^k\right)$, and that the annualized k-period PD is equal to q. The latter implies that

$$P\left(B^0 \to D^k\right) = 1-(1-q)^k . \tag{8}$$

For $n=k+1$, we have

$$P\left(B^0 \to A^{k+1}\right) = P\left(B^0 \to A^k\right)\cdot x + \left[P\left(B^0 \to B^k\right) + P\left(B^0 \to C^k\right)\right]\frac{1-q-x}{2}, \tag{9}$$

And

$$P\left(B^0 \to A^{k+1}\right) = P\left(B^0 \to A^k\right)\cdot x + \left[P\left(B^0 \to B^k\right) + P\left(B^0 \to C^k\right)\right]\frac{1-q-x}{2}, \tag{10}$$

Invoking the induction hypothesis yields (a). In addition,

$$\begin{aligned} P\left(B^0 \to D^{k+1}\right) &= P\left(B^0 \to A^k\right)(q-\Delta) + P\left(B^0 \to B^k\right)\cdot q + P\left(B^0 \to C^k\right)(q+\Delta) + P\left(B^0 \to D^k\right) = \\ &= q\cdot\left[P\left(B^0 \to A^k\right) + P\left(B^0 \to B^k\right) + P\left(B^0 \to C^k\right)\right] + P\left(B^0 \to D^k\right) = \\ &= q\cdot(1-q)^k + 1-(1-q)^k = \\ &= 1-(1-q)^{k+1}, \end{aligned} \tag{11}$$

where the second equality holds because $P\left(B^0 \to A^k\right) = P\left(B^0 \to C^k\right)$ by the induction hypothesis on (a), and the third equality holds since the sum in square parentheses is the probability of no-default up to period k. Annualizing the expression on the right-hand side of Equation (11) yields (b).

Appendix B

In Figure 8 we show the results of the eigenspace optimization algorithm described in Section 5.2 applied to an EDF transition matrix estimated using an alternative discretization of the EDF state space. As evident in the figure, the algorithm generates a monthly transition matrix that closely conforms to the annual EDF transition matrix. In contrast, the eigenspace solution obtained for the original discretization, depicted in Figure 7, shows significant deviations from the annual transition rates. This appendix describes a numerical optimization scheme for the discretization of the EDF space that results in the annual transition rates portrayed in Figure 8.

In detail, the problem involves finding a segmentation of the interval $[0,1)$ into K bins, $[0,b_1),[b_1,b_2),[b_2,b_3),\ldots,[b_{K-2},b_{K-1}),[b_{K-1},1)$, so that each bin defines a credit category consisting of a range of EDF values. We wish to find a segmentation that, when used to estimate an EDF transition matrix, results in minimal error between the annual transition rates and the annual rates implied from the monthly transition. Formally, we express the optimization problem as follows:

$$
\begin{aligned}
\min_{\mathbf{b}} & \left\|\log(T)-\log(X^{12})\right\|_1 \Big/ \left\|\log(T)\right\|_1 \\
& + \lambda \left\|\left(P(P^T P)^{-1}P^T - I\right)\log\left(T_K\right)\right\|_1 \Big/ \left\|\log\left(T_K\right)\right\|_1 \\
& + \gamma \left\|(P^T P)^{-1}P^T\left(\log\left(T_K\right)-\log\left(T_{\mathrm{Ref,K}}\right)\right)\right\|_1 \Big/ \left\|\log\left(T_{\mathrm{Ref,K}}\right)\right\|_1
\end{aligned}
\tag{12}
$$

subject to

$$
\begin{aligned}
& b_1 > 0, \\
& b_i > b_{i-1}, \quad i = 2,\ldots,K-1, \\
& b_{K-1} < 1,
\end{aligned}
$$

Where

- T is the annual transition matrix estimated from the segmentation vector **b**, and T_{K} denotes its K-th column;
- X is the monthly transition matrix calculated using the eigenspace optimization method;
- T_{Ref} is the EDF transition matrix listed in Table 2;
- $\|\ \|_1$ denotes the element-wise vector 1-norm;
- log ( ) for a vector or matrix is applied element-wise;
- $P = \begin{bmatrix} 1 & 1 & \cdots & 1 \\ 1 & 2 & \cdots & K-1 \end{bmatrix}^T$
- λ and γ are regularization parameters.

The optimization problem (Equation 11) has a non-linear objective with linear inequality constraints, and can be solved using an active-set method. The objective function in (Equation 11) consists of three terms. The first term is the sum of absolute differences between the annual transition probabilities in T and the implied annual transition probabilities from the monthly matrix X. We calculate differences on the logarithms of elements so that large probabilities do not dominate the error term. The dependence of T and X on the segmentation vector **b** is implicit; given a vector **b**, the transition matrix T is the result of estimating migration probabilities between bins defined in **b** using the available EDF data, and the matrix X is the result of applying a high-frequency generation algorithm to the matrix T. Thus, while an analytic form of the objective function is not readily available, it can be evaluated numerically for different realizations of the vector **b**.

To avoid optimal solutions that yield impractical segmentations of the probability mass, e.g., solutions in which most of the probability mass is concentrated in a single bucket, we introduce two regularization terms in the objective function.

Intuitively, we base these regularization terms on the common practice of segmenting the probability space in a log-linear fashion, so that buckets are more granular for low PDs than for high PDs. This property holds approximately in the Moody's rating transition matrix listed in Table 1. The first regularization term is the residual of the projection of the logarithm of the PD column of matrix T onto a two-dimensional subspace. A large residual term indicates that the PD column cannot be well represented by a log-linear function. The second regularization term penalizes segmentations whose log-linear projection differs in slope or intercept from the log-linear projection of the original segmentation.

The regularization parameters λ and γ can be tuned to control the magnitude of the penalizing terms. In the example shown in Figure 8, we used the values $\lambda = 0.04$, $\gamma = 0.06$. Finally, the constraints on $\mathbf{b}$ ensure that it constitutes an admissible segmentation of the interval $[0,1)$.

Note that we can tailor this optimization scheme to other problem formulations by tweaking the objective function or the constraints (e.g., the objective function can be modified to treat a different frequency than monthly). Yet, we recognize that the proposed procedure may not be applicable in a general setting. We view the choice of an appropriate segmentation of the space of PDs for the purposes of estimating a transition matrix from historical transition rates (either in a continuous time setting or in a discrete time setting) as an important topic that merits further treatment.

Appendix C

In this appendix, we list in numeric form the element-wise absolute error in the annual transition probabilities implied from the high-frequency solution, as depicted in Figure 5, Figure 6, Figure 7, and Figure 8. We list errors in matrix form, with each element of the error matrix equal to the absolute difference between the two elements in the corresponding matrices. We report errors at 1bp accuracy.

Table 7 Absolute differences between Moody's annual rating transition probabilities and annual transition probabilities implied from the monthly Taylor series approximation (cf. Figure 5).

	Aaa	Aa	A	Baa	Ba	B	Caa-C	Default
Aaa	91.62%	7.70%	0.66%	0.00%	0.02%	0.00%	0.00%	0.01%
Aa	1.13%	91.31%	7.21%	0.27%	0.06%	0.02%	0.00%	0.02%
A	0.07%	2.84%	91.29%	5.14%	0.51%	0.09%	0.02%	0.03%
Baa	0.05%	0.20%	5.15%	88.83%	4.54%	0.81%	0.24%	0.18%
Ba	0.01%	0.06%	0.42%	6.25%	82.94%	8.48%	0.63%	1.20%
B	0.01%	0.05%	0.18%	0.39%	6.21%	81.93%	6.23%	5.00%
Caa-C	0.00%	0.03%	0.03%	0.19%	0.73%	11.22%	68.56%	19.23%
Default	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	100.00%

Table 8 Absolute differences between EDF-based annual transition probabilities and annual transition probabilities implied from the monthly Taylor series approximation (cf. Figure 6).

	Aaa	Aa	A	Baa	Ba	B	Caa-C	Default
Aaa	91.62%	7.70%	0.66%	0.00%	0.02%	0.00%	0.00%	0.01%
Aa	1.13%	91.31%	7.21%	0.27%	0.06%	0.02%	0.00%	0.02%
A	0.07%	2.84%	91.29%	5.14%	0.51%	0.09%	0.02%	0.03%
Baa	0.05%	0.20%	5.15%	88.83%	4.54%	0.81%	0.24%	0.18%
Ba	0.01%	0.06%	0.42%	6.25%	82.94%	8.48%	0.63%	1.20%
B	0.01%	0.05%	0.18%	0.39%	6.21%	81.93%	6.23%	5.00%
Caa-C	0.00%	0.03%	0.03%	0.19%	0.73%	11.22%	68.56%	19.23%
Default	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	100.00%

Table 9 Absolute differences between EDF-based annual transition probabilities and annual transition probabilities implied from the monthly eigenspace optimization solution (cf. Figure 7).

	Aaa	Aa	A	Baa	Ba	B	Caa-C	Default
Aaa	91.62%	7.70%	0.66%	0.00%	0.02%	0.00%	0.00%	0.01%
Aa	1.13%	91.31%	7.21%	0.27%	0.06%	0.02%	0.00%	0.02%
A	0.07%	2.84%	91.29%	5.14%	0.51%	0.09%	0.02%	0.03%
Baa	0.05%	0.20%	5.15%	88.83%	4.54%	0.81%	0.24%	0.18%
Ba	0.01%	0.06%	0.42%	6.25%	82.94%	8.48%	0.63%	1.20%
B	0.01%	0.05%	0.18%	0.39%	6.21%	81.93%	6.23%	5.00%
Caa-C	0.00%	0.03%	0.03%	0.19%	0.73%	11.22%	68.56%	19.23%
Default	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	100.00%

Table 10 Absolute differences between EDF-based annual transition probabilities and annual transition probabilities implied from the monthly eigenspace optimization solution, using an alternative EDF discretization scheme (cf. Figure 8).

	Aaa	Aa	A	Baa	Ba	B	Caa-C	Default
Aaa	91.62%	7.70%	0.66%	0.00%	0.02%	0.00%	0.00%	0.01%
Aa	1.13%	91.31%	7.21%	0.27%	0.06%	0.02%	0.00%	0.02%
A	0.07%	2.84%	91.29%	5.14%	0.51%	0.09%	0.02%	0.03%
Baa	0.05%	0.20%	5.15%	88.83%	4.54%	0.81%	0.24%	0.18%
Ba	0.01%	0.06%	0.42%	6.25%	82.94%	8.48%	0.63%	1.20%
B	0.01%	0.05%	0.18%	0.39%	6.21%	81.93%	6.23%	5.00%
Caa-C	0.00%	0.03%	0.03%	0.19%	0.73%	11.22%	68.56%	19.23%
Default	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	100.00%

Acknowledgements

We would like to acknowledge helpful suggestions and comments by Kamyar Moud, Shisheng Qu, and an anonymous referee.

References

Albert, A., 1962. Estimating the infinitesimal generator of a continuous time, finite state Markov process. *The Annals of Mathematical Statistics*, 33(2), 727-753.

Altman, E.I., Rijken, H.A, 2004. How rating agencies achieve rating stability. *Journal of Banking & Finance* 28(11), 2679-2714.

Andersen, P., Borgan, Ø., Hjort, N., Arjas, E., Stene, J., Aalen, O., 1985. Counting process models for life history data: A review. *Scandinavian Journal of Statistics*, 12(2), 97-158.

Bangia, A., Diebold, F.X., Kronimus, A., Schagen, C., Schuermann, T., 2002. Ratings migration and the business cycle, with application to credit portfolio stress testing. *Journal of Banking & Finance* 26(2-3), 445-474.

Blöchlinger, A., 2010. Arbitrage-free credit pricing using default probabilities and risk sensitivities. *Journal of Banking & Finance*, In Press.

Bluhm, C., Overbeck, L., 2007. Calibration of PD term structures: To be Markov or not to be. *Risk*, 20(11), 98-103.

Bonfim, D., 2009. Credit risk drivers: Evaluating the contribution of firm level information and of macroeconomic dynamics. *Journal of Banking & Finance*, 33(2), 281-299.

Bruche, M., González-Aguado, C., 2009. Recovery rates, default probabilities, and the credit cycle. *Journal of Banking & Finance* 34(4), 754-764.

Cantor, R., Mann, C., 2003. Are corporate bond ratings procyclical? Moody's Investor Services Special Comment.

Crommelin, D.T., Vanden-Eijnden, E., 2006. Fitting timeseries by continuous-time Markov chains: A quadratic programming approach. *Journal of Computational Physics* 217(2), 782-805.

Crosbie, P., Bohn, J.R., 2003. Modeling default risk. Moody's KMV Whitepaper.

Dwyer, D.W., Qu, S., 2007. EDF™ 8.0 model enhancements. Moody's KMV Whitepaper.

Frydman, H., Schuermann, T., 2008. Credit rating dynamics and Markov mixture models. *Journal of Banking & Finance* 32(6), 1062-1075.

Hill, P., Brooks, R., Faff, R., 2010. Variations in sovereign credit quality assessments across rating agencies. *Journal of Banking & Finance,* 34(6), 1327-1343.

IACPM / ISDA, 2006. Convergence of credit capital models.

Israel R.B., Rosenthal, J.S., Wei, J.Z., 2001. Finding generators for Markov chains via empirical transition matrices, with applications to credit ratings. *Mathematical Finance* 11(2), 245-265.

Jafry, Y., Schuermann, T., 2004. Measurement, estimation and comparison of credit migration matrices. *Journal of Banking & Finance* 28(11), 2603-2639.

Jarrow, R.A., Lando, D., Turnbull, S.M., 1997. A Markov model for the term structure of credit risk spreads. *Review of Financial Studies* 10, 481-523.

References

Kadam, A., Lenk, P., 2008. Bayesian inference for issuer heterogeneity in credit ratings migration. *Journal of Banking & Finance*, 32(10), 2267-2274.

Lando, D., Skødeberg, T., 2002. Analyzing ratings transitions and rating drift with continuous observations. *Journal of Banking & Finance* 26(2/3), 423–444.

Levy, A., 2008. An overview of modeling credit portfolios. Moody's KMV Whitepaper.

Livingston, M., Naranjo, A., Zhou, L., 2008. Split bond ratings and rating migration. *Journal of Banking & Finance*, 32(8), 1613-1624.

Metz, A., Donmez, N., Cantor, R., 2008. Comparing withdrawal adjustment methods: An application of Moody's credit transition model. Moody's Investor Services Special Comment.

Nickell, P., Perraudin, W., Varotto, S., 2000. Stability of rating transitions. *Journal of Banking & Finance* 24(1-2), 203-227.

Tennant, J., Emery, K., Cantor, R., 2008. Corporate one-to-five-year rating transition rates. Moody's Investor Services Special Comment.

Trück, S., Rachev, T., 2005a. Credit portfolio risk and PD confidence sets through the business cycle. *Journal of Credit Risk*, 1(4).

Trück, S., Rachev, T., 2005b. Changes in migration matrices and credit VAR - a new class of difference indices. Working paper.