Contact*	Phone
New York	
David T. Hamilton	1.212.553.1653
Richard Cantor	

Measuring Corporate Default Rates

Summary

Measurement of the probability of default for a corporate exposure over a given investment horizon is often the first step in credit risk modeling, management, and pricing. Many market practitioners base their parameter estimates on results reported in rating agency default studies. Although the comparability of default rates reported by the agencies has increased in recent years, many differences in default rate calculation methodologies remain and care should be taken to understand how these differences may limit their comparability.

One important and poorly understood methodological difference is whether default rate estimates are statistically adjusted for issuer rating withdrawals, which occur when borrowers shift from rated public to unrated private debt finance or when all their debts are extinguished outright. Unadjusted default rates report the share of rated issuers that were observed to have experienced a default over a particular measurement horizon. Withdrawal-adjusted default rates, however, are estimates of the share of rated issuers *expected* to default under the assumption that withdrawn issuers would have faced the same risk of default as other similarly rated issuers if they had stayed in the data sample. Both methods have legitimate uses under appropriate assumptions, but each method makes a different statement about default risk for a given historical data set.

In addition to their being statements of historical fact, unadjusted default rates may be useful benchmarks for the expected likelihood of default for obligations that have *fixed maximum potential tenors* and expected rating withdrawal rates similar to those exhibited by issuers in the empirical sample on which default rates were estimated. These requirements severely limit their applicability as proxies for expected default rates in practice, however. Expected rating withdrawal patterns for specific credit exposures are unlikely to be closely related to the historical average rating withdrawal pattern for corporate issuer rating histories. Furthermore, unadjusted default rate estimates are likely to be downwardly biased because rating agencies have incomplete knowledge of subsequent defaults once firms are no longer rated.

In contrast, withdrawal-adjusted default rates are the appropriate estimates of expected default rates for obligations with specific *expected realized tenors*. Adjusted default rates provide common yardsticks for default risk for credit exposures across all sectors, regardless of differences in rating withdrawal rates. Moreover, marginal default rates calculated using the withdrawal-adjusted method may be interpreted as default intensities, which are critical inputs to most credit pricing models. Moody's typically reports its default rates on a withdrawal-adjusted basis, although it also provides unadjusted default statistics as well.

In this *Special Comment* we review the mechanics and rationale behind Moody's corporate default rate calculation methodology. We discuss the relative merits of adjusting or not adjusting for rating withdrawals and the importance of the assumption that firms whose ratings are withdrawn would have faced similar default risk as firms that did not withdraw if had they remained in the data sample. We demonstrate that the available evidence suggests this is a reasonable assumption.

Moody's Investors Service
Global Credit Research

Table of Content

1. Introduction

The measurement of the probability of default for a corporate exposure is often the first step in credit risk modeling, management, and pricing. Rating agency default studies are widely-used sources for estimates of these important parameter values. The default statistics reported in rating agency studies are based on rich source data sets, containing a large number of corporate rating histories and credit events.[1] It is frequently assumed that the default statistics reported by the rating agencies are calculated using more or less the same methodology and may, therefore, be used interchangeably, compared, and interpreted more or less consistently. Furthermore, it is often taken for granted that the default statistics reported by the rating agencies are equally appropriate measures of risk for a given purpose.

In the past decade there has indeed been a convergence in the methodologies used by the agencies to calculate cumulative default rates, and their methodologies currently share many similarities. Most rating agencies emphasize issuer-based default statistics rather than dollar-volume based statistics; average default rate estimates for an historical time period are calculated using a cohort-based approach; and, long-term multi-year default rates are derived using a discrete-time hazard rate method.[2] Despite these similarities, the default rates for corresponding rating categories reported by the rating agencies often differ significantly. While variations in default rates by rating category across agencies are to be expected due to differences in rating methodologies, discongruities in the distributions of the underlying rated populations, variations in the agencies' definitions of default, the historical time periods under study and the periodicity of observation, an additional reason to expect differences is that each rating agency's default rate calculation methodology differs in its statistical treatment of issuer rating withdrawals.[3]

Default rate calculation methodologies generally take one of two approaches to dealing with rating withdrawals when calculating default rates: ignore them and make no adjustment; or adjust for rating withdrawals by treating them as randomly censored data. Under the *no adjustment for withdrawals* method, issuers whose ratings are withdrawn are treated as if they remained in the data sample over the entire measurement horizon. An attempt is made to monitor their subsequent default status. If no default is observed, the firm is assumed not to have defaulted. Hence, the no adjustment method takes a relatively simple view of the evolution of credit risk in that there are only two possible outcomes, default or non-default. Under the *withdrawal-adjusted* method, issuers whose ratings are withdrawn are treated as randomly censored data, meaning that it is assumed that firms whose ratings are withdrawn would have faced the same risk of default as other similarly rated issuers if they had stayed in the sample. The adjusted-for-withdrawals method recognizes that there are three possible end-of-period outcomes: default, survival, and rating withdrawal. Rating withdrawals represent losses from the data sample before the final outcome of interest (default or survival) is observed. Moody's default statistics are most often reported using the adjusted-for-withdrawals method, although Moody's also reports unadjusted default rates.

Both calculation methods have legitimate uses under the appropriate assumptions, but each method makes a different statement about default risk. As they are derived from historical corporate rating histories and default data, the default rate estimates generated by each method represent a view of the "actual" default experience of a given data sample. However, empirical default rates are frequently used as proxies for *expected* default probabilities, and it is for this purpose that the treatment of rating withdrawals becomes an important concern. Unadjusted default rates may be useful benchmarks for the expected likelihood of default for obligations that have fixed *maximum potential* tenors and expected rating withdrawal rates similar to those exhibited by issuers in the empirical sample on which default rates were estimated. In contrast, withdrawal-adjusted default rates are the appropriate estimates of expected default rates for obligations with specific *expected realized* tenors. Withdrawal-adjusted default rates therefore provide common yardsticks for comparing default risk for credit exposures across *all* sectors, regardless of difference in rating withdrawal rates.

In many respects, the issue is similar to that studied in Altman (1989). Altman (1989) maintained that prevailing methods for calculating multi-year bond default rates were unsuitable as estimates of expected default risk because they failed to account for maturities, calls, and other early redemptions that occur prior to the end of a given measurement horizon. Altman's mortality rate estimator recognized that calculating default rates based on the surviving population was the relevant measure of *expected* default risk. Coming to a similar conclusion, Asquith, et. al. (1989) showed that default rates estimates are materially affected by early bond redemptions, as nearly two-thirds of high yield bonds in their data sample had been called, defaulted, or exchanged within 10 years of issuance.

---

1. *Moody's database records the rating histories and defaults of over 19,000 Moody's-rated corporate and sovereign bond issuers since 1919. See Hamilton and Varma (2006).*
2. *Moody's (and other rating agencies) also reports default rates derived by calculating multi-period rating transition matrices. Although we do not discuss this method in this Special Comment, transition matrix-derived default rates – which generally report rating withdrawals as a distinct state – are very close to those derived using the unadjusted method discussed later in this study.*
3. *Differences in default rate calculation methods aside, it is important to keep in mind that disparities in default rates across agency rating scales are likely to result from differences across agencies' fundamental rating practices. Moreover, Moody's ratings, for example, are relative rankings, and Moody's does not attempt to hit particular default rate targets when assigning corporate ratings. See Fons, Cantor, and Mahoney (2002).*

The adjustments advocated by Altman (1989) and Asquith et. al. (1989) therefore amounted to adjusting for survival bias. But because rating agency default rates are typically *issuer* (or corporate family) based, adjusting for withdrawals depends critically on the assumption of random censoring. An issuer's rating may be withdrawn for a variety of reasons. One common reason is that a company has extinguished all of its rated public debt due to scheduled maturities, company-initiated calls, investor-initiated puts, or mergers and acquisitions. In many cases, the issuer is no longer at risk of default after a rating withdrawal because the withdrawal event corresponds to the extinguishment all of its debt obligations. However, in many other cases an issuer remains at risk of default after its rating has been withdrawn because it has replaced all of its public, rated debt with unrated, typically private, debt. The relevant question is whether issuer rating withdrawals are uninformative events or are correlated with changes in credit quality.

The remainder of this *Special Comment* is organized into four sections. In the first section we review the general cumulative default rate calculation methodology followed by Moody's and other rating agencies. We also identify certain features of Moody's default rates that distinguish them from other approaches. In Section 3 we explain the mechanics of Moody's adjustment for rating withdrawals and discuss the rationale underlying the unadjusted and withdrawal-adjusted methods. Following a long line of academic research, we argue that withdrawal-adjusted default rates have the most general use for applications requiring estimates of *expected* future default risk for a stated investment horizon. In Section 4 we analyze the hypothesis of the neutrality of issuer rating withdrawals. We demonstrate that the available evidence suggests the assumption of random censoring is reasonable.

2. Cumulative Default Rate Methodology

The cumulative default rate calculation methodology used by Moody's (and other agencies) is a discrete-time approximation of the nonparametric continuous-time hazard rate approach.[4] A pool of issuers, called a *cohort*, is formed on the basis of the rating held on a given calendar date (or set of dates), and the default/survival status of the members of the cohort is tracked over some stated time horizon. The time horizon T for which we desire to measure a default rate is divided into evenly spaced time intervals (e.g. months, years) of length t. Hence, the data is discrete in that the time to default is not measured continuously. In each time interval, some fraction of the cohort that has survived up to that time may default. The *marginal default rate* is the probability that an issuer that has survived in the cohort up to the beginning of a particular interval t will default by the end of the time interval. The T-horizon *cumulative default rate* is defined as the probability of default from the time of cohort formation up to and including time horizon T.

Cohorts of issuers can be formed on the basis of their original ratings or on the ratings held as of the cohort formation date. The original rating method, studied by Altman (1989), groups issuers into pools based on the first rating that was assigned to the issuer (or one of its obligations); such pools consist only of first-time issuers that were rated as of the cohort formation date(s).[5] In contrast, the cohort rating method on which Moody's and other agencies' corporate default studies often rely are based on pools of issuers holding a given rating on the cohort date regardless of original rating or time since issuance. Because Moody's long-term corporate ratings address the likelihood of default over multiple time horizons, regardless of age or time to maturity, Moody's usually reports corporate default rates based on the rating held on the cohort date rather than on original ratings.[6]

Mathematically, the marginal default rate in time interval t, $d(t)$, for a cohort of issuers formed on date y holding rating z is defined as the number of defaults $x(t)$ from the cohort that occur in the time interval t divided by the effective size of the cohort, $n(t)$, at the start of time t:

(2.1)
$$d_y^z(t) = \frac{x_y^z(t)}{n_y^z(t)}$$

Initially, $n(t)$ is equal to the number of issuers in the pool holding rating z on the cohort formation date. As time from the initial cohort date passes the size of the denominator falls because some issuers in the cohort fail to survive to the next time interval. As we discuss in detail in the next section, differences in the default rates reported by the rating agencies arise to a large extent because each rating agency models the default/survival process differently.

---

4. *The method is essentially that of Cutler and Ederer (1958). This approach is sometimes referred to as the life-table or actuarial method.*
5. *The original rating method captures the impact of the now well-known aging or seasoning effect (i.e. the term structure of default risk for a given issuance year and rating category). Marginal default (hazard) rates exhibit more pronounced "humps" relative to the cohort rating method.*
6. *Bond level ratings are statements about expected loss severity, which incorporates loss-given-default as well as default probability. The ratings referenced in Moody's default studies are senior unsecured (or estimated senior unsecured) issuer-level ratings, which control for loss severity (see Hamilton (2005)).*

Cumulative default rates for investment horizons of length *T*, denoted *D(T)*, are built up from the marginal default rates, and are found by subtracting the product of the fraction of surviving cohort members in each of the *t* time intervals from unity:

(2.2)
$$D_y^z(T) = 1 - \prod_{t=1}^{T}[1 - d_y^z(t)]$$

Or, expanding equation 2.2 (and dropping indices for brevity):

(2.3)
$$D(T) = d(1) + d(2)[1 - d(1)] + d(3)[(1 - d(1))(1 - d(2))] + ... + d(T)(\prod_{t=1}^{T-1}[1 - d(t)])$$

Equation 2.3 highlights the fact that a cumulative default rate is a conditional probability. In the first time period, a fraction of the credit exposures in the cohort either defaults or survives. The credit exposures that survive period one may then go on to default or survive in period two; those that survive period two may go on to default or survive in period three, etc. Because the time periods are non-overlapping and the probability of default in each period is assumed to be independent, the *T*-period cumulative default rate is defined as one minus the product of the *T* marginal survival rates.

Issuer-based default rates receive particular emphasis in the rating process because the expected *likelihood* of default of a bond issuer holding a given rating is expected be the same regardless of differences in the nominal sizes of the exposures.[7] For example, the expected likelihood of default for a B-rated corporate issuer should be the same whether the size of the exposure is $200 million or $2 billion, everything else equal. Issuer-based default rates give equal weight to all issuers in the default rate calculation. Dollar volume based default rates, which weight each exposure by the total face (or market) value of its outstanding bonds, are useful statistics for portfolio benchmarking, but they are less useful for forming expectations about future default *probabilities*.[8]

The frequency with which cohorts are formed also impacts the accuracy of the average default probability estimates for a given rating category. The higher the sampling frequency – equivalently, the shorter the time interval between cohorts – the more accurate the estimates of expected default rates for a given rating category become. Closer cohort spacing captures rating changes and default events that occur in small time intervals, an important consideration when an issuer's rating is undergoing rapid change. The effect of cohort spacing on default rate estimates becomes clear in the following example. Consider the senior unsecured rating history for LTV Steel Company up to its default on July 17, 1986:

Table 2.1 – LTV Steel Company Rating History

Rating Date	Rating	Event
11/18/1970	A	First rating assigned
4/26/1982	A3	Alphanumeric rating assigned
5/5/1982	Baa2	Downgraded
10/18/1982	Baa3	Downgraded
11/18/1983	Ba1	Downgraded
3/20/1985	Ba3	Downgraded
8/9/1985	B3	Downgraded
7/17/1986	Caa	Defaulted

Using annual cohort spacing, LTV Steel Company's default is recorded for the A-rated cohorts from 1971-1982, the Baa3-rated 1983 cohort, the Ba1-rated cohorts in 1984 and 1985, and the B3 1986 cohort. If one instead formed cohorts at monthly intervals, the default event gets captured at the appropriate time horizon for every rating in its rating history, including its A3, Baa2, Ba3 and Caa ratings that are ignored under annual cohort spacing. Moody's has traditionally reported its average cumulative default rates calculated using annual cohort spacing (cohorts of issuers formed on January 1 of each year). In Moody's 2005 default study, Moody's moved to monthly cohort spacing in calculating its average cumulative default rates. Moody's believes that monthly cohort spacing strikes a reasonable balance between the competing goals of informational efficiency and tractability.[9]

---

7. When a firm defaults on one bond it usually defaults on all its bonds due to cross-default clauses in bond indentures. Additionally, in some bankruptcy codes (e.g. U.S. Chapter 11 and France's "sauvegarde" procedure) an automatic stay provision triggered upon a bankruptcy filing creates perfect cross default, causing all debt to default at the same time (unless the bankruptcy judge grants a waiver). This approach is also consistent with the structural view of credit risk (e.g. Merton (1974)) which regards default as an issuer-level phenomenon that is primarily a function of firm-level characteristics, such as its operating performance and liability structure.
8. Fridson (1991) is an interesting discussion of the many different ways to measure default rates that addresses this and other topics.

While investors may be interested in the cumulative default experience of a particular cohort, averages taken over many cohort periods (which capture the effects of several macroeconomic and credit cycle peaks and troughs) are required to estimate *expected* cumulative default probabilities. The average cumulative default rate for a given historical time period is calculated by first averaging the period t marginal default rates across all available cohort dates y in the historical data set Y, then calculating the cumulative rates using equation 2.2 or 2.3. Moody's average cumulative default rates are weighted averages, where each period's marginal default rate is weighted by the relative size of the cohort (proportion of issuers) in each time interval t.[10]

Equation 2.4 shows that the calculation of the average cumulative default rate for rating class z, $\overline{D}^z(T)$, is derived from the weighted average marginal default rates, $\overline{d}^z(t)$, calculated from all the available cohort marginal default rates in the historical data set Y:

(2.4)
$$\overline{D}^z(T) = 1 - \prod_{t=1}^{T} [1 - \overline{d}^z(t)]$$

where

(2.5)
$$\overline{d}^z(t) = \frac{\sum_{y \in Y} x_y^z(t)}{\sum_{y \in Y} n_y^z(t)}$$

If, for example, one were calculating the average three-year cumulative default rate for the 2003-2005 time period (with annual cohort spacing), one would first take the weighted average of the $d(1)$ from each of the three cohort years. The second year's average marginal default rate would consist of the weighted average of the two cohort years (2003 and 2004) with two years of exposure available, $d(2)$. The third year average marginal default rate would simply consist of the 2003 cohort's third year marginal default rate since it is the only cohort with three years of history available (i.e. it would receive 100% weight). The average cumulative default rate would then simply be calculated according to equation 2.4 using the weighted average marginal default rates derived using equation 2.5.

Note that this procedure for calculating average cumulative default rates maximizes the existing historical information by using *all* the available rating and marginal default rate data, not just issuers with rating histories that endure for a period of at least length T. While the third year's *marginal* default rate is calculated from just one cohort year (the 2005 cohort), the three-year *cumulative* default rate reflects information on conditional default/survival derived from *all three* cohort years. What may seem like the simplest method – deriving average cumulative default rates directly from the cohort cumulative default rates – limits the estimated average to the set of cohorts with at least T full periods of data. For long-horizon default rate averages, this requirement throws away much useful data (as well as raising the noise of the estimate). More importantly, however, deriving average cumulative default rates directly from cohort cumulative default rates may result in seriously biased and inconsistent estimates of expected cumulative default risk. For example, average cumulative default rates could possibly be *decreasing* if the historical data sample consists of default rates that have been very high in recent cohorts but very low in past cohorts.

---

9. There is a tradeoff between informational efficiency and tractability when calculating default rates using duration methods. Default/survival times are precisely measured using continuous time methods, but the resulting output may be quite unwieldy. Making default event times discrete by arbitrarily choosing the width of the marginal time intervals the distance between cohort formation dates results in some loss of information, but greatly facilitates the presentation and interpretation of cohort cumulative default rates. For example, investors are often interested in default rates for certain discrete time horizons (e.g. one, five, ten years). As the time interval t is allowed to shrink so as to be so small that at most one default occurs within an interval, the derived default rates approach the continuous time estimate (Kaplan and Meier (1958)).

10. Weighted averages place greater weight on more recent cohorts as both the number and total dollar volume of bond issuance has experienced secular growth over time. This is appealing from a statistical sampling point of view, but also because defaults tend to be correlated with periods of active bond issuance. Simple averaging may be appropriate in some circumstances; e.g. the impact of macroeconomic fluctuations on multi-year default rates.

3. Adjusting For Rating Withdrawals

The calculation methodologies for cohort and average cumulative default rates described in the previous section are generally followed by all the major rating agencies (again, with minor variations). Default rate calculation methodologies diverge, however, on their assumptions about the default/survival process. Whereas most rating agencies' marginal default rates are incrementally adjusted for defaults that occurred in the past, Moody's default rates also account for rating withdrawals that occur prior to the end of the measurement period. Rating withdrawals complicate the calculation of default rates because there will be some issuers initially included in a cohort that will be lost from the data sample before the final outcome of interest (default or survival) is observed. In practice, a rating agency's approach to modeling the survival process is reflected in its calculation of the effective cohort size in each time interval; i.e. the denominator of equation 2.1.

Under the unadjusted method, the effective size of a cohort of issuers rated z formed on date y in time interval t is equal to the initial size of the cohort less the total number of issuers that have defaulted prior to the current time interval (and therefore cannot default in the future).[11] The denominator of the cohort marginal default rate (equation 2.1) for the unadjusted method is therefore calculated:

$$(3.1) \qquad n_y^z(t) = n_y^z(0) - \sum_{i=2}^{t-1} x_y^z(i)$$

In contrast, the withdrawal-adjusted method recognizes that there are three possible end-of-period outcomes: default, survival, and rating withdrawal.[12] The cohort size at time t is calculated as in equation 3.1, but with an additional adjustment for the number of issuers that have had their ratings withdrawn in periods prior to the current time interval. Additionally, a small adjustment is made for rating withdrawals that occur *within* the current time interval. Withdrawn ratings that occur within an interval are treated as if they were censored at the midpoint of the interval; i.e. were at risk for half the time.[13] Equation 3.2 shows the calculation of the denominator for the adjusted for withdrawals method.

$$(3.2) \qquad n_y^z(t) = n_y^z(0) - \sum_{i=2}^{t-1} x_y^z(i) - \sum_{i=2}^{t-1} w_y^z(i) - \tfrac{1}{2} w_y^z(t)$$

Unadjusted default rates are highly intuitive. They report the share of issuers that were observed to have experienced a default over a particular time horizon. Unadjusted default rates are clearly useful benchmarks for the likelihood of default for obligations that have fixed maximum tenors and expected rating withdrawal patterns similar to those of the empirical sample from which the default rate estimate was derived. [14] In contrast, withdrawal-adjusted default rates are more complex in both calculation and interpretation. Withdrawal-adjusted default rates are based in part on "hypothetical" data whose accuracy depends on the assumption that issuers whose ratings are withdrawn would have defaulted at the same average rates as other similarly-rated issuers. One might reasonably ask, therefore, why bother adjusting default rates for rating withdrawals? Unadjusted default rates, it turns out, have three shortcomings – not shared by withdrawal-adjusted default rates – that limit their usefulness as measures of expected default risk for similarly rated obligations.

Firstly, unadjusted cumulative default rates are downwardly biased measures of default risk because one cannot observe all defaults experienced by issuers after their ratings are withdrawn.[15] Figure 3.1 gives an indication of the magnitude of the problem. Of the 1,201 corporate bond issuers since 1980 that have defaulted and had their ratings withdrawn, the percentage of defaults observed after the rating withdrawal date (5%) is relatively low compared to that before the withdrawal date (95%). Many rating withdrawals occur when issuers retire their public debt with proceeds raised through private bank borrowings. These firms remain at risk of default but rating agencies cannot easily track

---

11. The unadjusted method is often referred to as the "static pool" method. However, the term has been subject to some confusion. Sometimes, the term static pool is meant to imply no adjustment for withdrawals. Other times, the term static pool has been used to refer to what we have defined (in Section 2) as the cohort approach to calculating multi-year default rates. Using our terminology, the static pool method can be defined as a cohort-based method that does not adjust for rating withdrawals.

12. Moody's considers the three possible end-of-period outcomes mutually exclusive. Issuers that default and have their rating withdrawn in the same time interval are categorized as defaults, not withdrawals.

13. The within-period adjustment is valid only if one is confident that defaults are observable after a rating withdrawal within the current time interval. The time interval must, therefore, be reasonably short (such as one year or less). Of course, for small enough time intervals the effect of the within period adjustment is immaterial.

14. A prime example is static synthetic corporate CDOs, which reference the debt obligations of a large number of corporations over a common and fixed maturity. In the event that all of the public debts and syndicated loans of a corporation are paid off, the risk in the CDO associated with that entity disappears. In such a structure, the historical average rating withdrawal pattern of the typical corporate issuer may be very relevant. For cash CDOs, on the other hand, the pattern of issuer rating withdrawals is less relevant, since it is unlikely to be closely related to the realized maturity patterns loans and bonds that comprise the structure's collateral pool.

their subsequent default experience. Hence, without adjustments for rating withdrawals, measured default rates are likely to understate true long-term issuer default rates. Under the withdrawal-adjusted method, rating withdrawals are viewed as events that randomly censor rating histories and events of default that occur after rating withdrawals need not be tracked. Rather, it is assumed that issuers whose ratings are withdrawn, had they remained in the sample, would have experienced the same rate of default as that experienced by similarly-rated issuers that did remain in the sample.

Figure 3.1 – Distribution of Distance between Observed Withdrawal and Default Times
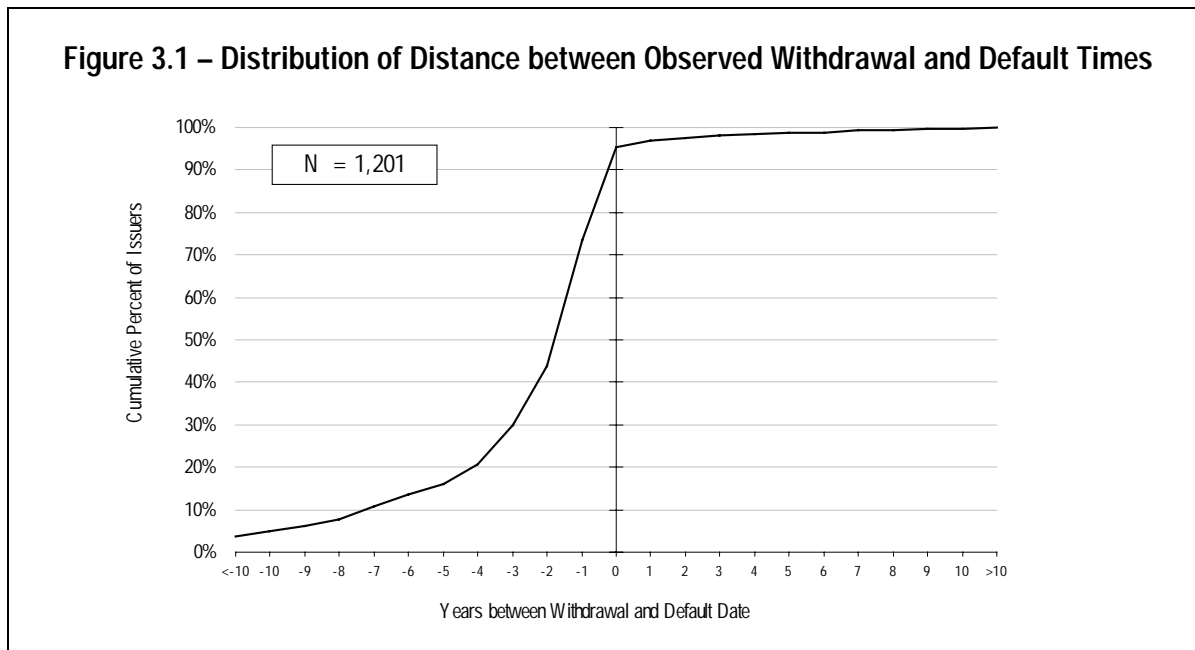
Secondly, the relevance of unadjusted default statistics as guides of future expected default experience is limited to sets of issuers with similar rating withdrawal patterns. While all issuers that carry the same rating can reasonably be expected to have roughly the same withdrawal-adjusted default rates, their unadjusted default rates are likely to vary significantly. Borrowers from different industries within the non-financial corporate sector and across different sectors altogether (non-financial corporates, financial corporates, and sovereigns) have very different rating withdrawal patterns and should therefore be expected to have markedly different unadjusted default rates by rating category. Moreover, structured finance securities have rating withdrawal patterns that differ considerably from one another and from those of the broad universe of corporate issuers.[16] As a result, it is hard to imagine how one can realistically use unadjusted corporate issuer default rates to benchmark and compare the risks of such diverse obligations. Withdrawal-adjusted default rates, in contrast, facilitate comparisons of default rates across asset classes with markedly different rating withdrawal patterns.

Lastly, credit spreads are unlikely to be closely related with unadjusted default rates. For a risk neutral investor, the appropriate discount rate for debt obligations subject to default risk is, in theory, equivalent to the risk free rate $r(t)$ plus a spread to account for expected loss in default.[17] Hence, $R(t) = r(t) + d(t)L$, where $d(t)$ is the expected marginal default rate and L is the expected (and assumed constant) rate of loss given default. If we assume that investors recovery nothing in the event of default (i.e. $L=1$), then the required spread $s(t)$ for a risk-neutral investor is simply the expected marginal default rate: $s(t) = R(t)-r(t) = d(t)$. The appropriate measure of $d(t)$ is the withdrawal-adjusted marginal default rate, not the unadjusted marginal default rate. If an investor's expected investment horizon were, say, 10 years, then (s)he would only require compensation for default risk only on exposures expected to survive for at least 10 years. In

---

15. It has been asserted (e.g. DeRosa-Farag, et. al. (1999)) that withdrawal-adjusted cumulative default rates are "biased" too high due to the correction for data censoring. However, this perspective confuses concerns about sample size and statistical significance with issues of bias. Consider the following oft-used hypothetical example. Suppose that there were 10 bond issuers in a cohort, nine of which had their ratings withdrawn over a 10-year time span for benign reasons such as mergers or retirement of debt. If, in the 10th year, the one company that was still rated were to default, the withdrawal-adjusted marginal default rate would be 100%. This sample statistic is not biased. However, because only one issuer was at risk of default in its tenth year, the statistical reliability of the 100% point estimate is virtually nil. In order for the unadjusted default rate to reach 100%, all nine of the censored issuers would need to default after their ratings were withdrawn. The unadjusted method assumes that there were ten issuers at risk of default in the tenth year, lowering the empirical marginal default rate estimate to 10%. The small sample problem does not go away by changing the definition of the default rate.

16. See Hu (2004).

17. In addition to compensation for expected credit losses, credit spreads may also be influenced by tax effects, interest rate risk premia, and a variety of other potential sources of risk premia. The academic literature in this area is large. See Duffie and Singleton (2003) for an overview.

fact, as long as withdrawal-adjusted marginal default rates are used in pricing, no further adjustments for realized rating withdrawals are required. In Appendix B we demonstrate this argument with an example.

As proxies for expected default probabilities, the advantages of withdrawal-adjusted default rates are consequently threefold: they avoid the downward bias that can arise from incomplete knowledge of defaults for firms whose ratings are withdrawn; they provide a common yardstick for measuring default risk for issuers and obligations across different sectors regardless of differences in rating withdrawal rates and, thus, associate a single time profile of default rates for each rating category for all types of credit exposures. Lastly, they provide useful and relevant data for pricing a wide variety of debt obligations.

Adjusting cumulative default rates for rating withdrawals, however, changes the meaning of the reported default rates from the simple concept of a share of a pool of issuers that were observed to default over a specific horizon. Rather, withdrawal-adjusted default rates report the historical frequency of default over a particular horizon conditional on having a rating outstanding for that length of time. For example, the withdrawal-adjusted 10-year cumulative default rate provides the answer to the following question: if one expects to have exposure to an issuer for T years, what is its probability of default? As long as rating withdrawals are randomly censored, then withdrawal-adjusted rates are appropriate estimates. The unadjusted default rate provides the answer to a simpler question: if one has a *maximum potential* maturity of T years, what is the expected probability of default? The answer provided, however, is potentially downwardly biased (as discussed above) and is only accurate if the issuer withdrawal experience of the underlying sample matches the expected rating withdrawal experience of the obligation in question.

An example should make the differences between the two methods clear. Table 3.1 shows detailed calculation of the 1- through 10-year cumulative default rates for the January 1, 1996 cohort of B-rated corporate bond issuers. The table shows the number of defaults, $x(t)$ and rating withdrawals, $w(t)$, in each year after the cohort formation date, as well as the effective size of the denominator in each time interval, $n(t)$, for each of the two methods. The table also shows the marginal default rates and the resulting cumulative default rates for each method (calculated using equation 2.2).

Table 3.1 – 10-Year Cumulative Default Rates: Adjusted vs. Unadjusted Methods

January 1, 1996 Cohort of B-Rated Corporate Issuers

t	x(t)	w(t)	Withdrawal-Adjusted Method			Unadjusted Method		
			n(t)	d(t)	D(t)	n(t)	d(t)	D(t)
0	0	0	519	0.00%	0.00%	519	0.00%	0.00%
1	7	55	491.5	1.42%	1.42%	519	1.35%	1.35%
2	13	51	431.5	3.01%	4.39%	512	2.54%	3.85%
3	19	61	362.5	5.24%	9.41%	499	3.81%	7.51%
4	12	42	292	4.11%	13.13%	480	2.50%	9.83%
5	17	23	247.5	6.87%	19.10%	468	3.63%	13.10%
6	21	12	213	9.86%	27.07%	451	4.66%	17.15%
7	19	28	172	11.05%	35.13%	430	4.42%	20.81%
8	8	22	128	6.25%	39.18%	411	1.95%	22.35%
9	4	14	102	3.92%	41.57%	403	0.99%	23.12%
10	1	14	84	1.19%	42.26%	399	0.25%	23.31%

Table 3.1 shows that, at any given measurement horizon, the withdrawal-adjusted method results in higher default rate estimates than the unadjusted method, with the difference growing larger as the time horizon lengthens. The unadjusted 10-year cumulative default rate shows that 23.31% of issuers originally in the cohort defaulted by the tenth year. The unadjusted method is calculated *as if* the 322 issuers whose ratings were withdrawn had remained in the cohort and did not default over the entire 10 year measurement period. In contrast, the withdrawal-adjusted 10-year cumulative default rate method yields an estimate of 42.26%. Default rates calculated using the withdrawal-adjusted method are based on the number of issuers that remain at risk (i.e. have not previous defaulted nor had their ratings withdrawn) of default in each time interval. For example, the marginal default rate in the tenth year is 25 basis points under the unadjusted method; under the withdrawal-adjusted method, it is nearly five times higher, 1.19%. The withdrawal-adjusted approach recognizes that at the start of the tenth year only 84 issuers actually remained at risk of default. Appendix A shows average cumulative unadjusted and withdrawal-adjusted default rates for a 20 year time horizon.

In addition to the three advantages of the withdrawal-adjusted method discussed above, the method also generates default probability estimates with intuitive and appealing statistical characteristics. The data in Table 3.1 illustrates that cumulative default rates calculated using the withdrawal-adjusted method will, at sufficiently long time horizons, approach 100% much more quickly than unadjusted default rates. This has a natural statistical interpretation:

conditional on survival, all firms will likely eventually default. Cumulative default rates for a given cohort calculated using the unadjusted method, on the other hand, may never approach 100% over *any* measurement horizon. In order for the cumulative default rate to approach 100%, all the issuers whose ratings were withdrawn would need to be observed to ultimately default.

As we have mentioned several times in the preceding sections, the accuracy of the withdrawal-adjusted default rate measure depends critically upon the validity of the assumption of random data censoring. We analyze the validity of this assumption in the next section.

4. Assessing The Neutrality Of Issuer Rating Withdrawals

An issuer rating withdrawal indicates that Moody's has ceased to rate all the publicly rated bonds of an issuer. At the *bond* level, rating withdrawals are overwhelmingly correlated with scheduled or anticipated redemptions. Table 5.1 shows the reasons for rating withdrawals organized into five categories. The table shows that of the 137, 414 bond[18] rating withdrawals between 1980 and 2005, 97% were associated with maturity, calls, puts, conversions, or mergers. The business reasons category includes instances where the issuer chose to stop paying for a rating or the size of the bond issue was increased or decreased (and re-rated), or Moody's removed the rating because of lack of information from the issuer. The defaulted category includes cases where the bond rating was withdrawn on or shortly after the date of default.

Table 5.1 – Reasons for Rating Withdrawals, 1980-2005

Reason	%Bonds	%Issuers
Matured	68.84%	32.00%
Called, put, converted, etc.	27.85%	43.97%
Reason unknown	3.15%	21.31%
Business reasons	0.11%	2.02%
Defaulted	0.05%	0.70%
N	137,414	3,275

Table 5.1 also shows that 76% of *issuer* rating withdrawals corresponded to the final maturity, call, etc. of its bonds. However, even if it were known that all the bonds of an issuer were withdrawn due to, say, maturity (therefore making default on those particular bonds impossible), the rationale behind the firm's decision exit the rated bond market might reveal information about its default risk. Bond rating withdrawals, which are closely associated with maturities and redemptions of specific bonds, and issuer rating withdrawals, which are related to a firm's decision to exit or issue in the rated public bond market, reflect different corporate finance choices of a firm. It is impossible to know with complete certainty whether an issuer defaults (or would have defaulted) after all its ratings are withdrawn. It is in this sense that issuer rating withdrawals are censoring events.

An issuer rating withdrawal might signal heightened future credit risk. For example, a bond issuer experiencing financial distress may be forced into the private or short-term debt market. Contrarily, an issuer rating withdrawal might be negatively correlated with default risk if issuers experiencing improving credit quality choose to pay off their rated debt obligations or replace debt with equity. As a matter of rating policy, the rating outstanding immediately prior to a withdrawal is intended to reflect Moody's view of the credit at the time of the withdrawal.[19] Hence, rating withdrawals are supposed to be neutral events that are not systematically correlated with changes in default risk.

The neutrality of issuer rating withdrawals is, ultimately, an empirical question, yet there is almost no published research on the subject. Carty (1997) is the only study that has attempted to assess whether treating rating withdrawals as randomly censored data is justified. Carty's analysis using Moody's data was, however, indirect because he analyzed the reasons for *bond* rating withdrawals and their correlation with *issuer* rating withdrawals. In this section we attempt to asses whether the assumption that firms whose ratings are withdrawn would have faced similar default risk as firms that did not withdraw if had they remained in the data sample is valid. We seek to answer two specific questions. Firstly, does default risk increase or decrease leading up to or shortly following the rating withdrawal date? Secondly, is the *level* of default risk correlated with rating withdrawal events? Establishing the neutrality of issuer rating withdrawals is not straightforward as there is no direct statistical test for random versus informative censoring. We therefore attempt to infer the neutrality of withdrawals by examining several indications of default risk near rating

---

18. *Includes coupon, discount, and convertible bonds.*
19. *See Moody's (2004).*

withdrawal dates. Collectively, the results we present in this section provide evidence that the assumption of random censoring is reasonable.

A large body of literature[20] has shown that default rates vary with past rating actions: default rates are relatively higher conditional on a past downgrade and relatively lower conditional on a past upgrade. If rating withdrawal rates exhibit the same dependence on past rating actions, then we might have cause to doubt the random censoring hypothesis. For example, if rating withdrawal rates are higher conditional on a past downgrade, then rating withdrawals might represent "hidden" defaults. Table 4.1 shows average one-year issuer rating withdrawal rates conditional on rating upgrades, downgrades, and no changes in the prior year based on monthly cohorts of corporate issuers between 1983 and 2005. (A rating action represents a change of one alphanumeric rating notch or more; e.g. Ba2 to Ba3.) For investment-grade rated issuers in the aggregate, issuer rating withdrawal rates do not appear to exhibit strong systematic differences when conditioned on rating changes in the year prior to the rating withdrawal. Among speculative-grade rated issuers, those whose ratings were unchanged in the past year exhibit somewhat higher rates of rating withdrawal. Overall, the likelihood of a rating withdrawal does not appear to be influenced by rating changes in the past year.

Similar results are evident when issuer rating withdrawal frequencies are conditioned on outlook and review (Watchlist) status. Moody's rating outlooks and reviews provide indications of the likely direction and timing of future credit rating changes. Cantor and Hamilton (2004) and Cantor and Hamilton (2005) showed that default rates for similarly rated issuers differ when conditioned on rating outlook and Watchlist status. Similar to the analysis of rating actions, a positive correlation between outlooks/review and rating withdrawal rates would cast a doubt on the validity of the assumption of uninformative rating withdrawals. Table 4.2 shows average one-year issuer rating withdrawal rates between 1995 and 2005 conditional on outlook and Watchlist status held at the start of each monthly period.[21] Issuers with outlooks other than stable have exhibited higher rating withdrawal rates, but there does not appear to be a systematic bias in one direction: one-year rating withdrawal rates are roughly symmetric for positive and negative outlooks and Watchlist assignments.

Table 5.2 – One-Year Rating Withdrawal Rates Conditional on Prior Rating Action, 1983-2005

	Rating Action in Prior Year		
	Upgraded	Unchanged	Downgraded
Investment-Grade	5.39%	4.20%	4.30%
Speculative-Grade	9.62%	12.64%	9.88%
All Rated	6.79%	6.76%	7.02%

### Table 5.3 – One-Year WR Rates Conditional on Outlook Status, 1995-2005

	Outlook on Cohort Date				
	Watch Up	Positive	Stable	Negative	Watch Down
Investment-Grade	6.25%	4.79%	3.70%	3.62%	5.78%
Speculative-Grade	7.04%	7.99%	6.56%	7.51%	5.54%
All Rated	6.54%	6.41%	4.92%	5.51%	5.70%

Moody'sKMV EDFs offer a way to analyze changes in default risk both before and after an issuer rating withdrawal occurs. We analyzed monthly Moody'sKMV EDFs in the 12 months before and the 6 months after withdrawal dates for 5,577 issuers between 1998 and 2005. The data was divided into two samples: 206 issuers with EDF data available that actually experienced a rating withdrawal, as well as a control group consisting of 5,371 issuers that did not experience a rating withdrawal. For the non-withdrawn subset, we measured EDFs around the same rating withdrawal dates as for the withdrawn subset. Because EDFs can vary between 0 and 0.20, and because we are interested in changes in EDFs around withdrawal dates, we normalized the EDFs so that they are equal to 1 in the twelfth month prior to the rating withdrawal date. Figure 4.1 shows the average EDF indices around rating withdrawal dates for the two groups of issuers. The graph shows that the changes in the average EDFs of the two subsets exhibit very similar patterns in the months leading up to and after the rating withdrawal date. Both series exhibit a slight upward trend leading up to the rating withdrawal date, but the magnitude of the change is relatively small.

---

20. See, for example, Altman (1991), Carty (1997), and Cantor and Hamilton (2004).
21. Moody's introduced the Watchlist in 1991 and rating outlooks in 1995.

Figure 5.1 – Average EDF Index around Withdrawal Dates

EDF Index

Time to Rating Withdrawal (Months)

——Not Withdrawn   - - - - Withdrawn
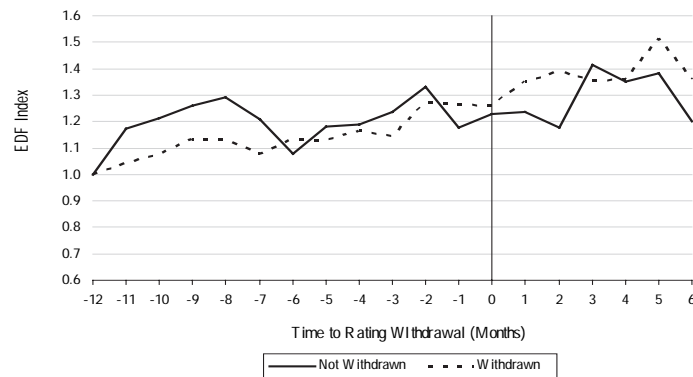
As a final assessment of the neutrality hypothesis we compared the weighted average *levels* of credit ratings for issuers that experienced a rating withdrawal to those that did not. We formed annual cohorts for 4,833 corporate bond issuers over the 1983 to 2005 time period, and recorded their estimated senior unsecured ratings from the cohort year to cohort year plus 5 years. To calculate the average rating level at each point in time, we weighted ratings using Moody's CDO rating factors.[22] To the extent that rating levels are correlated with future default risk (as clearly demonstrated by Moody's default studies), any significant differences observed in the rating levels in the stratified samples would cast doubt on the neutrality hypothesis.

Table 4.3 presents the results of this exercise. The table shows the number of notches of difference between the ratings of the withdrawn sub-sample and the non-withdrawn sub-sample. By construction, the average ratings for the two sub-samples on the cohort dates are the same; as time elapses rating changes may occur. Negative values in the table indicate the number of rating notches lower the withdrawn sub-sample was relative to the non-withdrawn sub-sample. The relevant question is whether the sub-sample that experienced a rating withdrawal exhibited significant rating improvement or degradation relative to the sub-sample that did not experience a rating withdrawal. The data shows that 47% of the time the average ratings of the sub-samples is the same at year 5; 42% of the time the average ratings of the withdrawn sub-sample is lower than the not withdrawn sample by one rating notch at year 5. Hence, it appears that the average ratings of a sub-set of issuers whose ratings are ultimately withdrawn exhibit little difference compared to the average ratings of issuers whose ratings are not withdrawn over the same time period.

Table 5.4 – Withdrawn vs. Non-Withdrawn Factor-Weighted Rating Notch Differences

Cohort Rating	+1 year	+2 years	+3 years	+4 years	+5 years
Aaa	0	0	0	-1	0
Aa1	0	0	-1	-1	0
Aa2	0	0	0	-1	0
Aa3	-1	-1	0	0	-1
A1	-1	-1	0	-1	-1
A2	0	0	0	-1	-1
A3	-1	-1	-1	-1	-1
Baa1	0	0	-1	-1	-1
Baa2	0	0	1	0	0
Baa3	0	-1	-1	-1	-2
Ba1	0	0	-1	-1	-1
Ba2	0	0	0	0	0
Ba3	0	0	0	-1	-1
B1	0	0	0	-1	-1
B2	0	0	0	0	0
B3	0	0	0	0	0
Caa1	0	0	-1	-1	-1
Caa2	0	0	0	0	0
Caa3	-1	0	0	0	0

Sample period: 1983-2005

---

22. *Moody's CDO rating factors are described in Yoshizawa and Witt (2003).*

Appendix A. 20-Year cumulative default rate tables

Table A1 – Average Cumulative Default Rates by Whole Letter Rating, Unadjusted vs. Withdrawal-Adjusted

Unadjusted

	Years After Cohort Formation Date																			
	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20
Aaa	0.00	0.00	0.00	0.02	0.08	0.14	0.19	0.25	0.30	0.36	0.41	0.47	0.52	0.55	0.57	0.60	0.63	0.64	0.64	0.64
Aa	0.01	0.02	0.04	0.09	0.15	0.20	0.25	0.29	0.32	0.34	0.37	0.40	0.45	0.49	0.51	0.54	0.57	0.60	0.62	0.65
A	0.02	0.10	0.21	0.31	0.41	0.51	0.60	0.70	0.79	0.87	0.95	1.02	1.08	1.13	1.21	1.28	1.36	1.43	1.50	1.55
Baa	0.18	0.49	0.86	1.25	1.60	1.91	2.18	2.41	2.64	2.87	3.08	3.28	3.47	3.64	3.80	3.92	4.03	4.12	4.20	4.27
Ba	1.18	2.98	4.85	6.52	7.86	8.90	9.68	10.34	10.90	11.40	11.84	12.25	12.59	12.89	13.12	13.34	13.52	13.64	13.75	13.84
B	5.41	11.04	15.41	18.46	20.66	22.17	23.21	23.85	24.29	24.59	24.78	24.91	25.00	25.08	25.14	25.18	25.19	25.20	25.20	25.20
Caa-C	19.90	29.37	34.86	37.81	39.32	40.16	40.58	40.88	41.06	41.18	41.23	41.23	41.23	41.23	41.23	41.23	41.23	41.23	41.23	41.23

Withdrawal-Adjusted

	Years After Cohort Formation Date																			
	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20
Aaa	0.00	0.00	0.00	0.03	0.11	0.18	0.27	0.36	0.45	0.56	0.66	0.78	0.90	0.97	1.04	1.12	1.22	1.25	1.25	1.25
Aa	0.01	0.02	0.05	0.12	0.19	0.29	0.38	0.46	0.51	0.58	0.65	0.76	0.92	1.06	1.16	1.29	1.44	1.58	1.74	1.91
A	0.02	0.10	0.24	0.37	0.51	0.67	0.83	1.01	1.22	1.42	1.63	1.82	2.02	2.21	2.49	2.81	3.17	3.53	3.88	4.20
Baa	0.18	0.53	0.98	1.52	2.06	2.60	3.13	3.65	4.23	4.89	5.50	6.17	6.85	7.56	8.24	8.84	9.41	9.97	10.44	10.91
Ba	1.23	3.31	5.75	8.26	10.57	12.65	14.48	16.28	18.05	19.86	21.62	23.41	25.15	26.82	28.29	29.78	31.14	32.17	33.15	33.97
B	5.65	12.35	18.65	24.09	29.06	33.50	37.47	40.71	43.59	46.12	47.56	48.77	49.65	50.51	51.26	51.77	51.96	52.12	52.12	52.12
Caa-C	21.12	33.53	43.47	51.01	56.52	61.05	64.58	68.50	71.98	74.72	75.16	75.16	75.16	75.16	75.16	75.16	75.16	75.16	75.16	75.16

Sample Period: 1970-2005, monthly cohort spacing

Appendix B. Corporate bond pricing with rating withdrawals

Using a reduced form model similar to that of Fons (1994), we show that the credit spread required by a risk neutral investor is independent of the rate withdrawal; i.e., the frequency that a bond prepays prior to maturity. The equation below represents the price a risk neutral investor would pay at time zero for a risky bond with face value of 1 payable at maturity date T. The bond pays an annual coupon $c(t)$. The bond is subject to three possible end-of-year outcomes: it may default, it may withdraw, or it may survive. In each year, the issuer may default with probability $d(t)$; if the issuer does not default, it may then pay off the bond early and withdraw with probability $w(t)$. For simplicity, we assume that loss-given-default for the bond is 100%. $r(t)$ denotes the risk-free rate.

$$P(0,T) = \sum_{i=1}^{T} \left[\left[\prod_{i=1}^{T}[1-d(i)][1-w(i-1)]w(i) \right] \left[\sum_{i=1}^{T} \left(\frac{c(i-1)}{\prod_{t=1}^{T}(1+r(i-1))} \right) + \frac{1+c(i)}{\prod_{t=1}^{T}(1+r(i))} \right] \right]$$

$$+ \sum_{i=1}^{T} \left[\left[\prod_{i=1}^{T}[1-d(i-1)][1-w(i-1)]d(i) \right] \sum_{i=1}^{T} \left(\frac{c(i-1)}{\prod_{t=1}^{T}(1+r(i-1))} \right) \right]$$

$$+ \left[\prod_{i=1}^{T}[1-d(i)][1-w(i)] \right] \left[\sum_{i=1}^{T} \left(\frac{c(i)}{\prod_{t=1}^{T}(1+r(i))} \right) + \frac{1}{\prod_{t=1}^{T}(1+r(i))} \right]$$

The three terms on the right-hand side of the equation are the probability weighted payoffs in the event of withdrawal, default, and survival to maturity, respectively. In the event of a withdrawal prior to maturity, an investor receives all the coupons and principal up to the withdrawal date. If default occurs in year t, an investor only receives coupons paid prior to t and loses the entire face value upon default. The last term in the equation is the payoff if the bond survives – i.e. does not withdraw and does not default – to maturity.

Suppose the coupon paid in each year is equal to the sum of the risk free rate and the marginal default probability adjusted for survival,

$$c(t) = \frac{r(t)+d(t)}{1-d(t)}$$

Upon substituting this value of c(t) into the formula, it can easily be seen that the price of the bond at time 0 should be 1 and all the terms involving $w(t)$ drop out. That is, the bond will price at par if the coupon rate is set equal to the risk free rate plus an adjustment for default risk; however, the withdrawal rate has no impact on the required spread. Since credit spreads are compensation for bearing risk, only marginal default rates matter for pricing, not marginal withdrawal rates.

References

Altman, E. I., 1989. "Measuring Corporate Bond Mortality and Performance," *Journal of Finance* 44: 909-22.

Asquith, P., Mullins, D. W., Wolff, E. D, 1989. "Original Issue High Yield Bonds: Aging Analyses of Defaults, Exchanges, and Calls," *Journal of Finance* 44: 923-52.

Carty, L. V., 1997. "Measuring Changes in Corporate Credit Quality," *Moody's Global Credit Research*, July.

Cutler, J. S., Ederer, F., 1958, "Maximum Utilization of the Life Table Method in Analyzing Survival," *J. Chronic Dis.* 8:699-712.

DeRosa-Farag, S., Blau, J., Matousek, P., Chandra, I., 1999. "Default Rates in the High Yield Market," *Journal of Fixed Income*, June: 7-31.

Duffie, D. and K. Singleton, 2003, *Credit Risk: Pricing, Measurement, and Management*. Princeton University Press.

Fons, J., 1994. "Using Default Rates to Model the Term Structure of Credit Risk," *Financial Analysts Journal*, September/October: 25-32.

Fons, J., Cantor, R., and Mahoney, C., 2002. "Understanding Moody's Corporate Bond Ratings and Rating Process," *Moody's Global Credit Research*, June.

Fridson, M. S., 1991. "Everything You Ever Wanted to Know about Default Rates," *Extra Credit*, Merrill Lynch High Yield Strategy, July/August: 4-14.

Hamilton, D. T., 2005. "Moody's Senior Ratings Algorithm and Estimated Senior Ratings," *Moody's Global Credit Research*, July.

Hamilton, D. T., Varma, P., 2005. "Default and Recovery Rates of Corporate Bond Issuers," *Moody's Global Credit Research*, January.

Hu, J., 2004. "Structured Finance Rating Transitions: 1983-2003," *Moody's Global Credit Research*, February.

Kaplan, E. L. and P. Meier, 1958, "Non-Parametric Estimation from Incomplete Observations," *Journal of the American Statistical Association*, 53: 437-481.

Merton, R. C., 1974. "On the Pricing of Corporate Debt: The Risk Structure of Interest Rates," *Journal of Finance*, 29: 449-470.

Moody's Investors Service, 2004. "Moody's Guidelines for the Withdrawal of Ratings," *Moody's Global Credit Research*, January.

Yoshizawa, Y. and G. Witt, 2003. "Moody's Approach to Rating Synthetic CDOs," *Moody's Structured Finance Rating Methodology*, July.

To order reprints of this report (100 copies minimum), please call 1.212.553.1658.
Report Number: 100779

Authors	**Production Associate**	**Senior Production Associate**
David T. Hamilton	*Tara Cheparev*	*Wing Chan*
Richard Cantor		