Available online at www.sciencedirect.com
SCIENCE DIRECT®
Journal of Banking & Finance 28 (2004) 2789–2811
Journal of BANKING & FINANCE
www.elsevier.com/locate/econbase

The relationship between credit default swap spreads, bond yields, and credit rating announcements

John Hull *, Mirela Predescu, Alan White

Joseph L. Rotman School of Management, University of Toronto, 105 St. George Street, Toronto, ON, Canada M5S 3E6

Available online 4 August 2004

Abstract

A company's credit default swap spread is the cost per annum for protection against a default by the company. In this paper we analyze data on credit default swap spreads collected by a credit derivatives broker. We first examine the relationship between credit default spreads and bond yields and reach conclusions on the benchmark risk-free rate used by participants in the credit derivatives market. We then carry out a series of tests to explore the extent to which credit rating announcements by Moody's are anticipated by participants in the credit default swap market.

JEL subject classification: G13; G14
Keywords: Credit default swaps; Credit risk; Bonds; Credit ratings

1. Introduction

Credit derivatives are an exciting innovation in financial markets. They have the potential to allow companies to trade and manage credit risks in much the same way

* Corresponding author. Tel.: +1 416 978 8615; fax: +1 416 971 3048.
E-mail addresses: hull@rotman.utoronto.ca (J. Hull), mirela.predescu01@rotman.utoronto.ca (M. Predescu), awhite@rotman.utoronto.ca (A. White).

0378-4266/$ - see front matter © 2004 Elsevier B.V. All rights reserved.
doi:10.1016/j.jbankfin.2004.06.010

as market risks. The most popular credit derivative is a credit default swap (CDS). This contract provides insurance against a default by a particular company or sovereign entity. The company is known as the reference entity and a default by the company is known as a credit event. The buyer of the insurance makes periodic payments to the seller and in return obtains the right to sell a bond issued by the reference entity for its face value if a credit event occurs.

The rate of payments made per year by the buyer is known as the CDS spread. Suppose that the CDS spread for a 5-year contract on Ford Motor Credit with a principal of $10 million is 300 basis points. This means that the buyer pays $300,000 per year and obtains the right to sell bonds with a face value of $10 million issued by Ford for the face value in the event of a default by Ford. [1] The credit default swap market has grown rapidly since the International Swaps and Derivatives Association produced its first version of a standardized contract in 1998.

Credit ratings for sovereign and corporate bond issues have been produced in the United States by rating agencies such as Moody's and Standard and Poor's (S&P) for many years. In the case of Moody's the best rating is Aaa. Bonds with this rating are considered to have almost no chance of defaulting in the near future. The next best rating is Aa. After that come A, Baa, Ba, B and Caa. The S&P ratings corresponding to Moody's Aaa, Aa, A, Baa, Ba, B, and Caa are AAA, AA, A, BBB, BB, B, and CCC, respectively. To create finer rating categories Moody's divides its Aa category into Aa1, Aa2, and Aa3; it divides A into A1, A2, and A3; and so on. Similarly S&P divides its AA category into AA+, AA, and AA−; it divides its A category into A+, A, and A−; etc. Only the Moody's Aaa and S&P AAA categories are not subdivided. Ratings below Baa3 (Moody's) and BBB− (S&P) are referred to as "below investment grade".

Analysts and commentators often use ratings as descriptors of the creditworthiness of bond issuers rather than descriptors of the quality of the bonds themselves. This is reasonable because it is rare for two different bonds issued by the same company to have different ratings. Indeed, when rating agencies announce rating changes they often refer to companies, not individual bond issues. In this paper we will similarly assume that ratings are attributes of companies rather than bonds.

The paper has two objectives. The first is to examine the relationship between credit default swap spreads and bond yields. The second is to examine the relationship between credit default swap spreads and announcements by rating agencies. The analyses are based on over 200,000 CDS spread bids and offers collected by a credit derivatives broker over a 5-year period.

In the first part of the paper we point out that in theory the N-year CDS spread should be close to the excess of the yield on an N-year bond issued by the reference entity over the risk-free rate. This is because a portfolio consisting of a CDS and a

[1] In a standard contract, payments by the buyer are made quarterly or semi-annually in arrears. If the reference entity defaults, there is a final accrual payment and payments then stop. Contracts are sometimes settled in cash rather than by the delivery of bonds. In this case there is a calculation agent who has the responsibility of determining the market price, x, of a bond issued by the reference entity a specified number of days after the credit event. The payment by the seller is then $100 - x$ per $100 of principal.

par yield bond issued by the reference entity is very similar to a par yield risk-free bond. We examine how well the theoretical relationship between CDS spreads and bond yield spreads holds. A number of other researchers have independently carried out related research. Longstaff et al. (2003), using the Treasury rate as the benchmark risk-free rate, and find significant differences between credit default swap spreads and bond yield spreads. Blanco et al. (2003) use the swap rate as the risk-free rate and find credit default swap spreads to be quite close to bond yield spreads. They also find that the credit default swap market leads the bond market so that most price discovery occurs in the credit default swap market. Houweling and Vorst (2002) confirm that the credit default swap market appears to use the swap rate rather than the Treasury rate as the risk-free rate. Our research is consistent with these findings. We adjust CDS spreads to allow for the fact that the payoff does not reimburse the buyer of protection for accrued interest on bonds. We estimate that the market is using a risk-free rate about 10 basis points less than the swap rate.

The second part of the paper looks at the relationship between credit default swap spreads and credit ratings. Some previous research has looked at the relationship between stock returns and credit ratings. Hand et al. (1992) find negative abnormal stock returns immediately after a review for downgrade or a downgrade announcement, but no effects for upgrades or positive reviews. Goh and Ederington (1993) find negative stock market reaction only to downgrades associated with a deterioration of firm's financial prospects but not to those attributed to an increase in leverage or reorganization. Cross-sectional variation in stock market reaction is documented by Goh and Ederington (1999) who find a stronger negative reaction to downgrades to and within non-investment grade than to downgrades within the investment grade category. Cornell et al. (1989) relates the impact of rating announcements to the firm's net intangible assets. Pinches and Singleton (1978) and Holthausen and Leftwich (1986) find that equity returns anticipate both upgrades and downgrades.

Other previous research has considered bond price reactions to rating changes. Katz (1974) and Grier and Katz (1976) look at monthly changes in bond yields and bond prices, respectively. They conclude that in the industrial bond market there was some anticipation before decreases, but not increases. Using daily bond prices, Hand et al. (1992) find significant abnormal bond returns associated with reviews and rating changes. [2] Wansley et al. (1992) confirm the strong negative effect of downgrades (but not upgrades) on bond returns during the period just before and just after the announcement. Their study concludes that negative excess returns are positively correlated with the number of rating notches changed and with prior excess negative returns. [3] This effect is not related to whether the rating change caused the firm to become non-investment grade. By contrast, Hite and Warga (1997) find that the strongest bond price reaction is associated with downgrades to and within the non-investment grade class. Their findings are confirmed by Dynkin et al. (2002) who report significant underperformance during the period leading up to

[2] An exception was a "non-contaminated" subsample, where there were no other stories about the firm other that the rating announcement.

[3] An example of a one-notch change is a change from Baa1 to Baa2.

downgrades with the largest underperformance being observed before downgrades to below investment grade. A recent study by Steiner and Heinke (2001) uses Eurobond data and detects that negative reviews and downgrades cause abnormal negative bond returns on the announcement day and the following trading days but no significant price changes are observed for upgrades and positive review announcements. This asymmetry in the bond market's reaction to positive and negative announcements was also documented by Wansley et al. (1992) and Hite and Warga (1997).

Credit default swap spreads are an interesting alternative to bond prices in empirical research on credit ratings for two reasons. [4] The first is that the CDS spread data provided by a broker consists of firm bid and offer quotes from dealers. Once a quote has been made, the dealer is committed to trading a minimum principal (usually \$10 million) at the quoted price. By contrast the bond yield data available to researchers usually consist of indications from dealers. There is no commitment from the dealer to trade at the specified price. The second attraction of CDS spreads is that no adjustment is required: they are already credit spreads. Bond yields require an assumption about the appropriate benchmark risk-free rate before they can be converted into credit spreads. As Section 3 will discuss, the usual practice of calculating the credit spread as the excess of the bond yield over a similar Treasury yield is questionable.

As one would expect, the CDS spread for a company is negatively related to its credit rating: the worse the credit rating, the higher the CDS spread. However, there is quite a variation in the CDS spreads that are observed for companies with a given credit rating. In Section 4 of the paper we consider a number of questions such as: To what extent do CDS spreads increase (decrease) before and after downgrade (upgrade) announcements? Are companies with relatively high (low) CDS spreads more likely to be downgraded (upgraded)? Does the length of time that a company has been in a rating category before a rating announcement influence the extent to which the rating change is anticipated by CDS spreads?

In addition to the credit rating change announcements, we consider other information produced by Moody's that may influence, or be influenced by, credit default swap spreads. These are reviews (also called Watchlists), and outlook reports. A review is typically either a review for upgrade or a review for downgrade. [5] It is a statement by the rating agency that it has concerns about the current rating of the entity and is carrying out an active analysis to determine whether or not the indicated change should be made. The third type of rating event is an outlook report

[4] Other empirical research on credit default swaps that has a different focus from ours is Cossin et al. (2002) and Skinner and Townend (2002). Cossin et al. examine how much of the variation in credit default swap spreads can be explained by a company's credit rating and other factors such as the level of interest rates, the slope of the yield curve, and the time to maturity. Skinner and Townend argue that a credit default swap can be viewed as a put option on the value of the underlying reference bond. Using a sample of sovereign CDS contracts, they investigate the influence of factors important in pricing put options on default swap spreads.

[5] Occasionally a firm is put on review with no indication as to whether it is for an upgrade or a owngrade. We ignore those events in our analysis.

from a rating agency analyst. These reports are similar to the types of reports that an equity analyst with an investment bank might provide. They are distributed via a press release (available on the Moody's website) and indicate the analyst's forecast of the future rating of the firm. Outlooks fall into three categories: rating predicted to improve, rating predicted to decline, and no change in rating expected. [6] To the best of our knowledge, ours is the first research to consider Moody's outlook reports. [7]

The rest of this paper is organized as follows. Section 2 describes our data. Section 3 examines the relationship between CDS spreads and bond yields and reaches conclusions on the benchmark risk-free rate used in the credit derivatives market. Section 4 presents our empirical tests on credit rating announcements. Conclusions are in Section 5.

2. The CDS data set

Our credit default swap data consist of a set of CDS spread quotes provided by GFI, a broker specializing in the trading of credit derivatives. The data covers the period from January 5, 1998 to May 24, 2002 and contains 233,620 individual CDS quotes. Each quote contains the following information:

1. The date on which the quote was made. [8]
2. The name of the reference entity.
3. The maturity of the CDS.
4. Whether the quote is a bid (wanting to buy protection) or an offer (wanting to sell protection), and
5. The CDS spread quote in basis points.

A quote is a firm commitment to trade a minimum notional of 10 million USD. [9] In some cases there are simultaneous bid and offer quotes on the same reference entity. When a trade took place the bid quote equals the offer quote.

The reference entity may be a corporation such as Blockbuster Inc., a sovereign such as Japan, or a quasi-sovereign such as the Federal Home Loan Mortgage Corporation. During the period covered by the data CDS quotes are provided on 1599 named entities: 1502 corporations, 60 sovereigns and 37 quasi-sovereigns. Of the reference entities 798 are North American, 451 are European, and 330 are Asian and Australian. The remaining reference entities are African or South American.

[6] In our analysis we ignore outlooks where no change is expected.

[7] Standard and Poor's (2001) considers the outlook reports produced by S&P.

[8] The quotes in our data set are not time stamped.

[9] The vast majority of the quotations are for CDSs denominated in USD. However, there is increasing activity in EUR and JPY. The proportion of the quotes denominated in USD from 1998 to 2002 is: 100%, 99.9%, 97.7%, 92.2%, and 71.4%.

The maturities of the contracts have evolved over the last 5 years. Initially, very short-term (less than 3 months) and rather longer-term (more than 5 years) contracts were relatively common. As trading has developed, the 5-year term has become by far the most popular. Approximately 85% of the quotes in 2001 and 2002 are for contracts with this term. [10]

The number of GFI quotations per unit of time has risen steadily from 4,759 in 1998 to an effective rate of over 125,000 quotes per year in 2002. The number of cases of simultaneous bid/offer quotes has risen from 1401 per year in 1998 to an effective rate of 54,252 per year in 2002. The number of named entities on which credit protection is available has also increased from 234 in 1998 to 1152 in 2001, the last year for which a full year of data is available.

The CDS rate quoted for any particular CDS depends on the term of the CDS and the credit quality of the underlying asset. The vast majority of quotes lie between 0 and 300 basis points. However, quotes occasionally exceed 3000 basis points. [11] The typical quote has evolved over the life of the market. In the first two years the prices quoted tended to decline which is consistent with a developing market in which competition is lowering the prices. However in the last 3 years it appears that the typical quote has been increasing. This is consistent with our observation that the average quality of the assets being protected is declining.

3. CDS spreads and bond yields

In theory CDS spreads should be closely related to bond yield spreads. Define y as the yield on an n-year par yield bond issued by a reference entity, r as the yield on an n-year par yield riskless bond, and s as the n-year CDS spread. The cash flows from a portfolio consisting of the n-year par yield bond issued by the reference entity and the n-year credit default swap are very close to those from the n-year par yield riskless bond in all states of the world. The relationship

$$s = y - r \tag{1}$$

should therefore hold approximately. If s is greater than $y - r$, an arbitrageur will find it profitable to buy a riskless bond, short a corporate bond and sell the credit default swap. If s is less than $y - r$, the arbitrageur will find it profitable to buy a corporate bond, buy the credit default swap and short a riskless bond.

There are a number of assumptions and approximations made in this arbitrage argument. In particular:

[10] At the end of 2002 the market began to standardize contract maturity dates. This means that the most popular maturity is approximately 5 years rather than exactly 5 years.

[11] Such high spreads may seem surprising but are not unreasonable. Suppose it was known with certainty that an entity would default in 1 year and that there would be no recovery. The loss 1 year from now would be 100% and to cover this cost it would be necessary to charge a CDS spread of about 10,000 basis points per year. If it were known that the entity would default in 1 month's time the spread would be 120,000 basis points per year, but it would be collected for only 1 month.

1. The argument assumes that market participants can short corporate bonds. Alternatively, it assumes that holders of these bonds are prepared to sell the bonds, buy riskless bonds, and sell default protection when $s > y - r$.
2. The argument assumes that market participants can short riskless bonds. This is equivalent to assuming that market participants can borrow at the riskless rate.
3. The argument ignores the "cheapest-to-deliver bond" option in a credit default swap. Typically a protection seller can choose to deliver any of a number of different bonds in the event of a default. [12]
4. The arbitrage assumes that interest rates are constant so that par yield bonds stay par yield bonds. By defining the corporate bond used in the arbitrage as a par corporate floating bond and the riskless bond as a par floating riskless bond we can avoid the constant interest rate assumption. Unfortunately, in practice par corporate floating bonds rarely trade.
5. There is counterparty default risk in a credit default swap. (We discuss this later.)
6. The circumstances under which the CDS pays off is carefully defined in ISDA documentation. The aim of the documentation is to match payoffs as closely as possible to situations under which a company fails to make payments as promised on bonds, but the matching is not perfect. In particular, it can happen that there is a credit event, but promised payments are made.
7. There may be tax and liquidity reasons that cause investors to prefer a riskless bond to a corporate bond plus a CDS or vice versa.
8. The arbitrage assumes that the CDS gives the holder the right to sell the par bond issued by the reference entity for its face value plus accrued interest. In practice it gives the holder the right to sell a bond for its face value.

As discussed by Duffie (1999) and Hull and White (2000) it is possible to adjust for the last point. Define A^* as the expected accrued interest on the par yield bond at the time of the default. The expected payoff from a CDS that gives the holder the right to sell a par yield bond for its face value plus accrued interest is $1 + A^*$ times the expected payoff on a regular CDS. To adjust for this we can replace Eq. (1) by

$$s = \frac{y - r}{1 + A^*}. \tag{2}$$

3.1. Alternative risk-free rates

The main problem in using Eq. (2) lies in choosing the risk-free rate, r. Bond traders tend to regard the Treasury zero curve as the risk-free zero curve and measure a corporate bond yield spread as the spread of the corporate bond yield over the yield

[12] The claim made by bondholders on the assets of the company in the event of a default is the bond's face value plus accrued interest. All else equal, bonds with low accrued interest are therefore likely to be cheapest to deliver. Also, in the event of a restructuring, the market may not expect all bonds to be treated similarly. This increases the value of the cheapest-to-deliver bond option.

on a similar government bond. By contrast, derivatives traders working for large financial institutions tend to use the swap zero curve (sometimes also called the LIBOR zero curve) as the risk-free zero curve in their pricing models because they consider LIBOR/swap rates to correspond closely to their opportunity cost of capital.

The choice of the Treasury zero curve as the risk-free zero curve is based on the argument that the yields on bonds reflect their credit risk. A bond issued by a government in its own currency has no credit risk so that its yield should equal the risk-free rate of interest. However, there are many other factors such as liquidity, taxation, and regulation that can affect the yield on a bond. For example, the yields on US Treasury bonds tend to be much lower than the yields on other instruments that have very low credit risk. One reason for this is that Treasury bonds have to be used by financial institutions to fulfill a variety of regulatory requirements. A second reason is that the amount of capital a financial institution is required to hold to support an investment in Treasury bonds is substantially smaller than the capital required to support a similar investment in low risk corporate bonds. A third reason is that the interest on Treasury bonds is not taxed at the state level whereas the interest on other fixed income investments is taxed at this level. For all of these non-credit-risk reasons, the yields on US Treasury bonds tend to be depressed relative to the yields on other low risk bonds. [13]

The swap zero curve is normally calculated from LIBOR deposit rates, Eurodollar futures, and swap rates. The credit risk associated with the swap zero curve is somewhat deceptive. The rates for maturities less than 1 year in the swap zero curve are LIBOR deposit rates and are relatively easy to understand. They are the short-term rates at which one financial institution is willing to lend funds to another financial institution in the inter-bank market. The borrowing financial institution must have an acceptable credit rating (usually Aa). From this it might be assumed that longer rates are also the rates at which Aa-rated companies can borrow. This is not the case. The n-year swap rate is lower than the n-year rate at which an Aa-rated financial institution borrows when $n > 1$. It represents the credit risk in a series of short-term loans to Aa borrowers rather than the credit risk in one long-term loan to Aa borrowers. Consider for example the 5-year swap rate when LIBOR is swapped for a fixed rate of interest and payments are made semi-annually. This is the rate of interest earned when a bank (a) enters into the 5-year swap and (b) makes a series of 10 six-month loans to companies with each of companies being sufficiently creditworthy that it qualifies for LIBOR funding at the beginning of its 6-month borrowing period. From this it is evident that rates calculated from the swap zero curve are very low risk rates, but are not totally risk-free. They are also liquid rates that are not subject to any special tax treatment.

[13] See Duffee (1996) and Reinhart and Sack (2001) for a further discussion of the market for Treasury instruments.

3.2. Test of Eq. (2)

To test Eq. (2) we chose a sample of 31 reference entities that were very actively quoted in our CDS data set. These are listed in Table 1. The reference entities were chosen to span the rating categories and to represent a range of different industries. We used only CDS quotes on these reference entities that corresponded to trades (that is, the bid quote equaled the offer quote).

For each of the reference entities we determined the CUSIPs of all the outstanding bond issues. The total number of issues considered was 964. The characteristics of each issue were downloaded from Bloomberg and the bonds to be included were selected using the following major criteria:

1. Bonds must not be puttable, callable, convertible, or reverse convertible.
2. Bonds must be single currency (USD) bonds with fixed rate, semi-annual coupons that are not indexed.

Table 1
List of the reference entities included in the analysis in Section 3

AT & T Corp
Bank of America Corporation
BT Group Plc
BHP Billiton Limited
Computer Associates Internatonal Inc
Deutsche Telekom AG
Enron Corp
Federal National Mortgage Association
Ford Motor Credit Company
France Telecom
General Electric Capital Corporation
General Motors Acceptance Corp
Goodyear Tire and Rubber Co
Hilton Hotels Corporation
Hutchison Whampoa Ltd
International Paper Co
Lehman Brothers Holdings Inc
Merrill Lynch & Co Inc
Morgan Stanley Dean Witter & Co
Pearson Plc
Philip Morris Companies Inc
Qwest Communications International Inc
Raytheon Company
Sears Roebuck Acceptance Corp
Sprint Corp
Telstra Corporation Limited
TRW Inc
Tyco International Ltd
Vodafone Group Plc
Williams Companies Inc
WorldCom Inc

3. Bonds must not be subordinated or structured.
4. The issue must not be a private placement.

We also filtered the bonds on their time to maturity to eliminate long maturity issues. After applying these criteria there were 183 issues remaining. Indicative yields for these issues for the period from January 1, 1998 to July 15, 2002 were downloaded from Bloomberg.

The CDS quotes were merged with the bond data in the following way. For each CDS transaction a corresponding 5-year bond par yield, y, was estimated by regressing yield against maturity for all the bonds of the reference entity on that date. [14] The time to maturity of the bonds used in the regression had to be between 2 and 10 years, and there had to be at least one bond with more than 5 years to maturity and one with less than 5 years to maturity. The regression model was then used to estimate the 5-year yield. This resulted in a total of 370 CDS quotes with matching 5-year bond yields. Of these 111 of the quotes were for reference entities in the Aaa and Aa rating categories, 215 for reference entities in the A rating category, and 44 for reference entities in the Baa rating category. Since all bonds paid interest semi-annually we assume that $A^* = y/4$ in Eq. (2) so that

$$y - s(1 + y/4) = r. \tag{3}$$

To test this equation we considered two alternative models:

$$y - s(1 + y/4) = a + br_{\mathrm{T}} + \varepsilon \tag{4}$$

and

$$y - s(1 + y/4) = a + br_{\mathrm{S}} + \varepsilon, \tag{5}$$

where r_{T} is the 5-year Treasury par yield, r_{S} is the 5-year swap rate, and ε is a normally distributed error term. [15] The regression results are shown in Table 2.

The model in Eq. (5), where the risk-free rate is the swap rate, provides a better fit to the data than the model in Eq. (4), where the risk-free rate is the Treasury rate. The ratio of sums of squared errors is 1.513. Under the hypothesis that the models are equally good this statistic should be distributed $F(368, 368)$. As a result we are able to reject the hypothesis that the models are equally good with a very high degree of confidence.

The model in Eq. (3) predicts that $a = 0$ and $b = 1$. We are unable to reject the hypothesis that $a = 0$ for both versions of the model. The value of b is significantly greater than 1.0 at the 1% confidence level when the Treasury rate is used as the risk-

[14] We tried other schemes to estimate the 5-year par yield. One of them was the interpolation method used by Blanco et al. (2003) where a synthetic 5-year bond yield is created from one large bond issue below and one above the five year maturity. However, none of these schemes proved to be better than the procedure we used. We also carried out the tests using mid-market CDS quotes where the bid/offer spread was less than 10 basis points. The results were similar but the standard errors were larger.

[15] The five-year swap rate is the par yield that would be calculated from the swap zero curve and was downloaded from Bloomberg. The 5-year Treasury par yield was estimated as the yield on the constant maturity 5-year Treasury bond taken from the Federal Reserve database.

Table 2
Test of the relationship between 5-year CDS spreads and 5-year bond yields using the assumption that the risk-free rate is (i) the Treasury rate, r_T, and (ii) the swap rate, r_S

	a	b	Standard error of residuals	Adjusted R^2
Eq. (4): risk-free rate is the Treasury rate	0.12 (0.070)	1.10 (0.014)	0.250	0.941
Eq. (5): risk-free rate is the swap rate	0.09 (0.059)	0.972 (0.010)	0.203	0.961

The CDS spread and all rates are measured in percentage points. Standard errors are shown in parentheses.

Table 3
Comparison of implied risk-free rate, r, with the corresponding Treasury rate, r_T, and the corresponding swap rate, r_S

	$r - r_T$ (bp)		$r - r_S$ (bp)		Q	
	Mean	SE (of mean)	Mean	SE (of mean)	Mean	SE (of mean)
Aaa/Aa	51.30	1.97	−9.55	1.31	0.834	0.0250
A	64.33	1.82	−5.83	1.59	0.927	0.0229
Baa	84.93	3.63	−2.21	2.79	0.967	0.0364
All Categories	62.87	1.38	−6.51	1.06	0.904	0.0160

The implied risk-free rate is about 63 basis points higher than the Treasury rate and about 6 basis points lower than the swap rate. The variable Q is the fraction of the distance from the Treasury rate to the swap rate at which the implied risk-free rate is found.

free rate and significantly less than 1.0 at the 1% confidence level when the swap rate is used as the risk-free rate. This suggests that the benchmark risk-free rate used by CDS market participants is between the Treasury rate and the swap rate. [16]

3.3. The benchmark risk-free rate

To investigate the benchmark risk-free rate further we examined the statistics of $r - r_T$ and $r - r_S$ where r is the implied risk-free interest rate calculated using Eq. (3). These statistics are summarized in Table 3. The table also shows statistics on a variable, Q, which is defined as

$$Q = \frac{r - r_T}{r_S - r_T}.$$

This is a measure of the fraction of the distance from the Treasury rate to the swap rate at which the implied risk-free rate is found. The results show that on average the implied risk-free rate lies 90.4% of the distance from the Treasury rate to the swap rate, 62.87 basis points higher than the Treasury rate and 6.51 basis points lower than the swap rate. Our results are consistent with those of Houweling and

[16] We tried alternative tests adjusting for heteroskedasticity. The results were very similar.

Vorst (2002) who use the CDS market to argue that market participants no longer see the Treasury curve as the risk-free curve and instead use the swap curve and/or the repo curve. Houweling and Vorst use Eq. (1) rather than Eq. (2) in their tests.

Table 3 shows that, as the credit quality of the reference entity declines, the implied risk-free rate rises. A possible explanation for this is that there is counterparty default risk in a CDS (that is, there is some possibility that the seller of the CDS will default). Hull and White (2001) provide an analytic approximation for the impact of counterparty default risk on CDS spreads. Using their formula with reasonable estimates of the parameters we were able to provide only a partial explanation of the differences between the results for rating categories in Table 3. We conclude that the results may be influenced by other factors such as differences in the liquidities of the bonds issued by reference entities in different rating categories.

The estimates made for the Aaa and Aa reference entities are probably most indicative of the benchmark risk-free rate applicable to liquid instruments. The impact of counterparty default risk on CDS spreads for these reference entities is extremely small and market participants have indicated to us that bonds issued by these reference entities tend to be fairly liquid. Our best estimate is therefore that the benchmark 5-year risk-free rate is on average about 10 basis points less than the swap rate or about 83% of the way from the Treasury rate to the swap rate.

4. CDS spreads and rating changes

Both the credit default swap for a company and the company's credit rating are driven by credit quality, which is an unobservable attribute of the company. Credit spreads change more or less continuously whereas credit ratings change discretely. If both were based on the same information we would expect rating changes to lag credit spread changes. As explained by Cantor and Mann (2003) rating agencies have stability as one of their objectives. (They try and avoid getting into a position where a rating change is made and has to be reversed a short time later.) This stability objective is also likely to cause rating changes to lag credit spread changes. However, rating agencies base their ratings on many different sources of information, some of which are not in the public domain. The possibility of rating changes leading credit spreads cannot therefore be ruled out.

In this section we carry two sorts of tests. We first condition on rating events and test whether credit spreads widen before and after rating events. We then condition on credit spread changes and test whether the probability of a rating event depends on credit spread changes. Our tests use the GFI database described in Section 2 and databases from Moody's that contain lists of rating events during the period covered by the GFI data.

We used the quotes in the GFI database between October 1, 1998 and May 24, 2002. We restricted our analysis to 5-year quotes on reference entities that were corporations rated by Moody's. We would have liked to proceed as in Section 3 and retain as observations only data where an actual trade was reported (that is, the bid quote equals the offer quote). However, this would have been led to insufficient observations for our empirical tests. We therefore chose to search for situations where there are both bid quotes and offer quotes for a reference entity on a particular day and they are reasonably close together. When there were both bid and offer quotes for a reference entity on a day we calculated U, the maximum of the bid quotes and V, the minimum of the offer quotes. If U and V were less than 30 basis points apart, we calculated a "spread observation" for the reference entity for the day as $0.5(U + V)$. The total number of spread observations obtained in this way for the period considered was 29,032. [17]

Macroeconomic effects cause the average level of CDS spreads to vary through time. For example, all CDS spreads increased sharply after September 11, 2001. To allow for this in our empirical tests, we calculated an index of CDS spreads for companies in each of the following three categories: Aaa and Aa, A, and Baa. (The Aaa and Aa categories were combined because there were relatively few reference entities in each category. We did not consider below investment grade categories because it is relatively rare for a CDS to trade on a reference entity in these categories.) This enabled us to convert each spread observation into an "adjusted spread observation" by subtracting the appropriate spread index. [18] An implicit assumption in our adjustment procedure is that all companies in a rating category have the same sensitivity to the index. We repeated all our tests without subtracting the spread index. It is reassuring that the results, including the level of significance, were similar to those we report here.

We considered six types of Moody's rating announcements: downgrades, upgrades, review for downgrade, review for upgrade, positive outlook, and negative outlook. We will refer to downgrades, reviews for downgrade and negative outlooks as "negative events" and upgrades, reviews for upgrade, and positive outlooks as "positive events".

4.1. Spread changes conditional on rating events

Our first test considered the changes in adjusted CDS spreads that occur before and after a Moody's rating event. [19] This is similar to a traditional event study. In our analysis we eliminated all Moody's events that were preceded by another event in the previous 90 business days. This controls for contamination. We define the time interval $[n_1, n_2]$ as the time interval lasting from n_1 business days after the event to n_2 business days after the event where n_1 and n_2 can be positive or negative. Thus $[-90, -61]$ is the time interval from 90 days before the event to 61 days before the

[17] If we had defined observations as situations where a trade was indicated we would have had a total of 5056 observations.

[18] We adjust the CDS spreads after a rating change by the index corresponding to the "old" rating category (the rating before the event). In this way we avoid any discontinuities at the time of the event that would have contaminated the announcement day effect.

[19] Using announcements from Standard and Poor's or Fitch as well as Moody's would have had the advantage of capturing more rating events, but would have had the disadvantage of leading to some double counting of events.

event; $[1, 10]$ is the time interval from 1 day after the event to 10 days after the event day; and so on. We calculated the "adjusted spread change" for interval $[n_1, n_2]$ as the adjusted spread observation for day n_2 minus the adjusted spread for day n_1. When there was no observation on the adjusted spread available for a day we estimated an adjusted spread observation by interpolating between adjacent observations. [20]

We considered whether the mean adjusted spread change for a rating event is significantly greater than (less than) zero for negative (positive) events. The distribution of the adjusted spread change often had a pronounced positive skew and the sample size (i.e., number of rating events for which the spread change could be calculated) was sometimes quite low so that a standard t-test was inappropriate. This led us to use the bootstrap technique described by Efron and Tibshirani (1993). Suppose that the values sampled for the adjusted spread change are $s_1, s_2, \ldots, s_n$, the mean adjusted spread change is $\bar{s}$, and the standard deviation of the spread change is $\hat{\sigma}$. The bootstrap test of whether the mean adjusted spread change is greater than zero is based on the distribution of the t-statistic: $t = \sqrt{n}(\bar{s}/\hat{\sigma})$. Define $\tilde{s}_i = s_i - \bar{s}$ for $i = 1, \ldots, n$. Our null hypothesis is that the distribution of the adjusted spread change corresponds to the distribution where $\tilde{s}_1, \tilde{s}_2, \ldots, \tilde{s}_n$ are equally likely. We will refer to this distribution, which has a mean of zero, as the null distribution. We repeat the following a large number of times: sample n times with replacement from the null distribution and calculate $t^{\mathrm{B}} = \sqrt{n}(\bar{s}^{\mathrm{B}}/\hat{\sigma}^{\mathrm{B}})$, where $\bar{s}^{\mathrm{B}}$, $\hat{\sigma}^{\mathrm{B}}$ are the sample mean and standard deviation. This provides an empirical distribution for t under the null hypothesis. By comparing t with the appropriate percentile of this distribution we are able to test whether the null hypothesis can be rejected at a particular confidence level.

Our results for negative rating events are shown in Table 4. [21] By pooling all observations we find a significant (at the 1% level) increase in the CDS spread well in advance of a downgrade event. In the case of reviews for downgrade and negative outlooks there is a significant (at the 1% level) increase in the CDS spread during the 30 days preceding the event. CDS spreads increase by approximately 38 bps in the 90 days before a downgrade, by 24 bps before a review for downgrade, and by 29 basis points before a negative outlook. When observations are pooled there are no significant changes in CDS spread during the 10 business days after any type of negative event.

Announcement day effects are captured by the $[-1, +1]$ interval. The announcement day effect for reviews for downgrade are significant at the 1% level when all companies are pooled (as well as for A and Baa companies considered separately). The average increase in the CDS spread at the time of a review for downgrade is almost 10 basis points. For downgrades and negative outlooks the average CDS in-

[20] An exception is that we never interpolated across day zero. If after applying our interpolation rules there was no observation on day n_1 but there were at least two observations between day n_1 and n_2 we used the next observation after day n_1 as a substitute for the observation on day n_1. The other rules we used were analogous. One implication of the rules is that our day 1 and day -1 results are produced only from spread observations on those days, not from interpolated spread observations.

[21] The sample size for an entry in Table 4 may be less than the corresponding number of events because there was sometimes insufficient data to calculate the spread change for a rating event.

Table 4
The mean change in the adjusted CDS spread during an interval that is prior to or after a negative rating event

	No. of events	Time interval				
		[−90, −61]	[−60, −31]	[−30, −1]	[−1, 1]	[1, 10]
Downgrade						
Aaa/Aa	17	−0.349	2.731*	2.524	0.440	−4.054
A	39	7.496	10.246**	16.347*	6.806*	5.929
Baa	27	37.069**	9.259	23.186	−9.610	35.304
All	83	14.076**	8.356**	15.001**	3.769	8.163
Review for downgrade						
Aaa/Aa	18	2.251	1.443	5.273*	0.663	−0.924
A	57	6.681*	−1.778	12.237**	12.175**	−0.190
Baa	39	6.044	10.694	21.186*	11.376**	−2.725
All	114	5.979*	3.157	14.573**	9.883**	−1.000
Negative outlook						
Aaa/Aa	7	1.597	9.543	−2.364	2.945	−2.663
A	39	4.567*	9.278**	14.587**	1.425	7.045*
Baa	23	3.361	2.097	29.214*	2.243	−8.899
All	69	3.979	7.032*	17.739**	1.961	0.575

The time interval [−90, −61] is from 90 business days before the event to 61 business days before the event. Other time intervals are defined similarly. The adjusted CDS spread on a day is the actual CDS spread less an index of the average CDS for the rating category.

* Indicates that the adjusted CDS spread change is greater than zero at the 5% confidence level.

** Indicates that it is greater than zero at the 1% confidence level.

creases for all companies, although positive, are not significant at the 5% level. This suggests that there is significant information in a review for downgrade, but perhaps not in a downgrade or negative outlook.

To summarize, there is evidence that the CDS market anticipates all three types of negative credit events. There is evidence of announcement day effects at the time of a review for downgrade. We did not find significant post-announcement day effects and conclude that CDS spreads fully adjust to the information in rating changes by day +1.

We carried out a similar test to that in Table 4 for positive events (upgrades, reviews for upgrades, and positive outlooks). We found virtually no significance, although we were reassured by the results that the average changes in adjusted CDS spreads were mostly negative. There are two possible reasons for our results. The first is that positive rating events are anticipated much less than negative rating events. (This would be consistent with the conclusions of other researchers, mentioned in the introduction, who looked at bond yields and equity prices.) The second is that the number of positive rating events is not large enough to get significance. The total number of positive rating events in our sample was 59 (22 upgrades, 24 positive reviews, and 13 positive outlooks) whereas the number of negative rating events was 266.

4.2. Estimating the probability of rating events

In our next set of tests we examine whether CDS spreads are useful in estimating the probability of a rating event. The test in Section 4.1 considers the adjusted spread change conditional on a rating event. Here we consider the probability of a rating event conditional on the adjusted spread change.

To carry out the analysis we constructed a set of non-overlapping 30-day time intervals for each reference entity and observed whether a particular rating event occurred in the 30 days following the end of the interval. Those intervals that contained a rating event of any kind were eliminated, thus controlling for contamination. We also eliminated intervals that did not include at least two spread observations on the reference entity.

In our first test we used the logistic model:

$$P = \frac{1}{1 + \mathrm{e}^{-a-bx}}, \tag{6}$$

where x is the adjusted spread change in a 30-day interval, P is the probability of a rating event during the 30 days following the end of the interval, and a and b are constants. We determined a and b from a maximum likelihood analysis. The adjusted spread change is defined as the last spread observed in the interval less the first spread observed in the interval. Our sample consisted of observations for all combinations of intervals and reference entities, except when they were eliminated for one of the reasons mentioned above.

The results are shown in Table 5. When companies in all rating categories are considered, the coefficient of the adjusted spread change is significant at the 1% level for the probability of a downgrade or a negative outlook, and is significant at the 5% level for the probability of a review for downgrade. In the case of downgrades, the coefficient of the adjusted spread change is significant at 1% for each rating category.

To provide an intuitive measure of the impact of x on P we calculated

$$\left.\frac{\mathrm{d}P}{\mathrm{d}x}\right|_{x=\bar{x}}$$

for the best fit values of a and b where $\bar{x}$ is the mean value of x. This measures the increase in the probability of a rating event for a one basis point increase in the adjusted spread change. We refer to this as the "probability sensitivity measure" or PSM in the table.

A natural alternative to looking at the adjusted spread change is to look at the adjusted spread level. We therefore took our sample of observations from the previous experiment and set x equal to the average adjusted spread level in an interval. The logistic model is the same as before and the PSM measure is defined as before. The results are shown in Table 6. They are similar to those in Table 5. Adjusted spread levels are significant at 1% level for downgrades and review for downgrades when all rating categories are pooled, but overall the results for outlooks are not significant.

Table 5
Results of the logistic regression

	All	Aaa/Aa	A	Baa
Downgrade				
a	-3.4783^{**}	-3.512^{**}	-3.3305^{**}	-3.9054^{**}
	(0.1104)	(0.2764)	(0.1433)	(0.247)
b	0.0183^{**}	0.071^{**}	0.0163^{**}	0.0225^{**}
	(0.00285)	(0.0245)	(0.00353)	(0.00515)
McFadden's LRI	0.0490	0.0465	0.0392	0.1018
PSM	0.000546	0.002093	0.000562	0.000445
Review for downgrade				
a	-3.3754^{**}	-3.9189^{**}	-3.3592^{**}	-3.1827^{**}
	(0.1047)	(0.3251)	(0.1456)	(0.1711)
b	0.00693^{*}	0.0132	0.00713	0.00592
	(0.00296)	(0.0389)	(0.00441)	(0.00392)
McFadden's LRI	0.0051	0.0012	0.0046	0.0062
PSM	0.000224	0.000254	0.000235	0.000228
Negative outlook				
a	-3.9933^{**}	-4.7016^{**}	-3.793^{**}	-4.1111^{**}
	(0.1402)	(0.4858)	(0.1783)	(0.2654)
b	0.00968^{**}	0.0434	0.012^{**}	0.00573
	(0.00316)	(0.049)	(0.00396)	(0.00562)
McFadden's LRI	0.0123	0.0124	0.0191	0.0052
PSM	0.000175	0.000399	0.000265	0.000091

$$P = \frac{1}{1 + e^{-a-bx}}$$

where x is the change in the adjusted spread during a 30-day period and P is the probability of a rating event during the subsequent 30 days. LRI is the McFadden's likelihood ratio index. PSM is the probability sensitivity measure, defined as the increase in the probability of a rating event for a 1bp increase in the adjusted spread change. Standard errors are shown in parentheses. Significance of coefficients is calculated using the Wald Chi-square test.
* Indicates significance at the 5% level.
** Indicates significance at the 1% level.

Analysts often look at what are termed cumulative accuracy profile curves (CAP curves) when comparing alternative models for predicting rating events. [22] (In other contexts these are called Lorenz curves.) Suppose that a variable x is proposed as an indicator of the probability of a particular rating event. A CAP curve is a plot of quantiles of the rating event against quantiles of x. It provides a visual qualitative guide to the predictive power of x. For example, it might show that the highest 10% of observations on x accounted for 30% of the rating events; the highest 25% of observations on x account for 50% of the rating events; and so on.

The logistic model is open to the criticism that it relies on a particular functional form for the relationship between the probability of a rating event and the

[22] See for example Moody's (2003).

Table 6
Results of the logistic regression

	All	Aaa/Aa	A	Baa
Downgrade				
a	−3.6336**	−3.6692**	−3.5833**	−3.9682**
	(0.1069)	(0.2638)	(0.148)	(0.2265)
b	0.00943**	0.0403**	0.0131**	0.00899**
	(0.0011)	(0.0152)	(0.002)	(0.0013)
McFadden's LRI	0.0765	0.0356	0.0647	0.1663
PSM	0.000249	0.001041	0.000378	0.000172
Review for downgrade				
a	−3.4994**	−3.8301**	−3.5401**	−3.3426**
	(0.1)	(0.2716)	(0.1451)	(0.1676)
b	0.00593**	−0.00213	0.00939**	0.00479**
	(0.001)	(0.0215)	(0.0022)	(0.0012)
McFadden's LRI	0.0262	0.0001	0.0283	0.0375
PSM	0.000174	−0.000044	0.000275	0.000162
Negative outlook				
a	−4.1246**	−4.6862**	−3.9917**	−4.183**
	(0.1347)	(0.1347)	(0.1801)	(0.2475)
b	0.00303	−0.00546	0.00692*	0.00127
	(0.0017)	(0.0332)	(0.0029)	(0.0025)
McFadden's LRI	0.0043	0.0004	0.0123	0.0013
PSM	0.000048	−0.000049	0.000130	0.000019

$$P = \frac{1}{1 + e^{-a-bx}}$$

where x is the average level of the adjusted spread during a 30-day period and P is the probability of a rating event during the subsequent 30 days. LRI is the McFadden's likelihood ratio index. PSM is the probability sensitivity measure, defined as the increase in the probability of a rating event for a 1bp increase in the average spread level. Standard errors are shown in parentheses. Significance of coefficients is calculated using the Wald Chi-square test.

* Indicates significance at the 5% level.

** Indicates significance at the 1% level.

explanatory variable. We therefore decided to develop a non-parametric test based on the idea underlying CAP curves. Consider again our first test where we calculate the change in the adjusted spread during a 30-day interval and observe whether a particular rating event occurs during the following 30 days. For each observation we assign a score of 1 if the rating event does occur and a score of zero if it does not occur.

We divided the observations into two categories: a high spread change category, H, and a low spread change category, L. The categories are defined as:

H: The set of observations in which the adjusted CDS spread change is greater than the $(100 - p)$th percentile of the distribution of all changes.

L: The set of observations in which the adjusted CDS spread change is less than the $(100 - p)$th percentile of the distribution of all changes.

We then counted the total score (i.e., the total number of rating events) for all the observations in each category.

Suppose that there are a total of N rating events of the type being considered, with n being from category H and $N-n$ being from category L. Our null hypothesis is that there is a probability p of any one of these events being from category H and $1-p$ of it being from category L. The probability of observing exactly n events from category H under the null hypothesis is

$$\pi(n) = \frac{N!}{n!(N-n)!} p^n (1-p)^{N-n}.$$

In a one-tailed test for negative events when the confidence level is q the critical value of n is the smallest value of n for which

$$\sum_{i=n}^{N} \pi(i) < q.$$

In a one-tailed test for positive rating events the critical value of n is the largest value for which

$$\sum_{i=0}^{n} \pi(i) < q.$$

The results for the negative rating events and three different values of p are shown in Table 7. The results are similar to those in Table 5. When all rating categories are considered together we get significant results for all rating events, except for review for downgrades and $p = 50\%$. This indicates that the adjusted spread change does contain useful information for estimating the probability of rating events. The results for the Aaa/Aa category show less significance than for other rating categories.

We proceeded similarly when looking at adjusted spread levels. The observations were divided into two categories, a high spread level category, H, and a low spread level category, L:

H: The set of observations for which the adjusted spread level is greater than the $(100-p)$th percentile of the distribution of all adjusted spread levels.

L: The set of observations for which the adjusted spread level is less than the $(100-p)$th percentile of the distribution of all adjusted spread levels.

We then counted the total score for all the observations in each category. The test of the significance of the results is the same as that given above for the Table 7.

The results for negative events are shown in Table 8. Again we find that adjusted spread levels have about the same explanatory power as adjusted spread changes in estimating the probability of rating events.

We carried out similar tests to those in Tables 5–8 for positive events. We found very little significance. As in the case of the results in Section 4.1, this may be because positive events are not anticipated or it may be because the number of positive events is quite small.

Table 7
The ability of adjusted spread changes during a 30-day interval to predict rating events during the 30 days following the interval

p	All	Aaa/Aa	A	Baa
Downgrade				
50	62.4**	66.7	62.1*	72.0*
25	42.6**	50.0*	41.4**	56.0**
10	28.7**	33.3**	24.1**	48.0**
Review for downgrade				
50	58.2	70.0	60.8	56.8
25	39.8**	30.0	43.1**	37.8
10	23.5**	10.0	21.6*	24.3**
Negative outlook				
50	76.4**	80.0	77.1**	73.3
25	50.9**	60.0	48.6**	40.0
10	30.9**	40.0	34.3**	20.0

The table shows the percentage of the events that occur when the adjusted spread change is above the $100 - p$ percentile of the distribution of adjusted spread changes. The null hypothesis is that the probability of an event occurring is p. The adjusted CDS spread on a day is the actual CDS spread less an index of the average CDS for the rating category.
* Indicates significance at the 5% level.
** Indicates significance at the 1% level.

Table 8
The ability of adjusted spread levels to predict rating events 30 days ahead

p	All	Aaa/Aa	A	Baa
Downgrades				
50	79.7**	65.0	79.7**	85.3**
25	58.5**	45.0*	53.1**	73.5**
10	37.3**	25.0*	31.3**	55.9**
Review for downgrade				
50	72.4**	50.0	71.2**	79.1**
25	45.7**	7.1	45.8**	60.5**
10	28.4**	0.0	27.1**	27.9**
Negative outlook				
50	67.8**	50.0	72.2**	70.6
25	47.5**	16.7	44.4**	47.1*
10	15.3	16.7	16.7	5.9

The table shows the percentage of the events that occur when the adjusted spread level is above the $100 - p$ percentile of the distribution of adjusted spread levels. The null hypothesis is that the probability of an event occurring is p. The adjusted CDS spread on a day is the actual CDS spread less an index of the average CDS for the rating category.
* Indicates significance at the 5% level.
** Indicates significance at the 1% level.

Researchers such as Altman and Kao (1992) and Lando and Skodeberg (2002) find that the probability of the credit rating change for a company depends on how long the company has been in its current rating category. The more recently a company has changed its credit rating the more likely it is to do so again in the next short period of time. This phenomenon is sometimes referred to as ratings momentum.

To test whether the length of time a company has been in its current rating category is a useful explanatory variable we modified the logistic model in Eq. (6) to

$$P = \frac{1}{1 + \mathrm{e}^{-a-bx-cu}},$$

where u is the length of time since the company's rating has changed, x is as before either the adjusted spread change or the adjusted spread level, and a, b and c are constants. Although the sign of c was almost invariably negative (indicating that the longer a company has been in its rating category the less likely a rating event is), it was not significant for any of the rating events we consider. This may be because CDS spreads reflect the information in the u.

5. Conclusions

Credit default swaps are a recent innovation in capital markets. There is a theoretical relationship between credit default swap spreads and bond yield spreads. We find that the theoretical relationship holds fairly well and that we are able to use it to estimate the benchmark 5-year risk-free rate used by participants in the credit default swap market. Our conclusion is that the risk-free rate used by market participants is about 10 basis points less than the 5-year swap rate on average. Alternatively it can be characterized as above the Treasury rate by about 83% of the spread between the 5-year swap rate and the 5-year Treasury rate.

We have conducted two types of analyses exploring the relationship between the credit default swap market and ratings announcements. In the first type of analysis we examine credit default swap changes conditional on a ratings announcement. We find that reviews for downgrade contain significant information, but downgrades and negative outlooks do not. There is anticipation of all three types of ratings announcements by the credit default swap market. In the second type of analysis we examine ratings announcements conditional on credit spread levels and credit spread changes. Either credit spread changes or credit spread levels provide helpful information in estimating the probability of negative credit rating changes. We find that 42.6% of downgrades, 39.8% of all reviews for downgrade and 50.9% of negative outlooks come from the top quartile of credit default swap changes.

Our results for positive rating events were much less significant than our results for negative rating events. This is consistent with the work of researchers who have looked at the relationship between rating events and bond yields, but may be influenced by the fact that there were far fewer positive rating events in our sample.

Acknowledgments

Joseph L. Rotman School of Management, University of Toronto. We are grateful to Moody's Investors Service for financial support and for making their historical data on company ratings available to us. We are grateful to GFI for making their data on CDS spreads available to us. We are also grateful to Jeff Bohn, Richard Cantor, Yu Du, Darrell Duffie, Jerry Fons, Louis Gagnon, Jay Hyman, Hui Hao, Lew Johnson, Chris Mann, Roger Stein, and participants at a Fields Institute seminar, meetings of the Moody's Academic Advisory Committee, a Queens University workshop, and an ICBI Risk Management conference for useful comments on earlier drafts of this paper. Matthew Merkley and Huafen (Florence) Wu provided excellent research assistance. Needless to say, we are fully responsible for the content of the paper.

References

Altman, E., Kao, D., 1992. The implications of corporate bond rating drift. Financial Analysts Journal (May–June), 64–75.

Blanco, R., Brennan, S., Marsh, I.W., 2003. An empirical analysis of the dynamic relationship between investment grade bonds and credit default swaps. Working Paper, Bank of England, May.

Cantor, R., Mann, C., 2003. Measuring the Performance of Corporate Bond Ratings. Special Comment, Moody's Investor's Services, April.

Cornell, B., Landsman, W., Shapiro, A.C., 1989. Cross-sectional regularities in the response of stock prices to bond rating changes. Journal of Accounting, Auditing and Finance 4, 460–479.

Cossin, C., Hricko, T., Aunon-Nerin, D., Huang, Z., 2002. Exploring the determinants of credit risk in credit default transaction data: Is fixed income markets' information sufficient to evaluate credit risk. Working Paper, University of Lausanne, May.

Duffee, G., 1996. Idiosyncratic variation in Treasury bill yields. Journal of Finance 51, 527–551.

Duffie, D., 1999. Credit swap valuation. Financial Analysts Journal (January/February), 73–87.

Dynkin, L., Hyman, J., Konstantinovsky, V., 2002. Sufficient diversification in credit portfolios. Journal of Portfolio Management (Fall), 89–114.

Efron, B., Tibshirani, R.J., 1993. An Introduction to the Bootstrap. Chapman & Hall.

Goh, J.C., Ederington, L.H., 1993. Is a bond rating downgrade bad news, good news, or no news for stockholders? Journal of Finance 48 (December), 2001–2008.

Goh, J.C., Ederington, L.H., 1999. Cross-sectional variation in the stock market reaction to bond rating changes. Quarterly Review of Economics and Finance 39 (1), 101–112.

Grier, P., Katz, S., 1976. The differential effects of bond rating changes among industrial and public utility bonds by maturity. Journal of Business 49 (April), 226–239.

Hand, J., Holthausen, R.W., Leftwich, R.W., 1992. The effect of bond rating agency announcements on bond and stock prices. Journal of Finance 47 (June), 733–752.

Hite, G., Warga, A., 1997. The effect of bond-rating changes on bond price performance. Financial Analysts Journal (May/June), 35–51.

Holthausen, R.W., Leftwich, R.W., 1986. The effect of bond rating changes on common stock prices. Journal of Financial Economics 17 (1), 57–89.

Houweling, P., Vorst, T., 2002. An empirical comparison of default swap pricing models. Working Paper, Erasmus University Rotterdam, June.

Hull, J.C., White, A., 2000. Valuing credit default swaps I: No counterparty default risk. Journal of Derivatives 8 (1), 29–40 Fall.

Hull, J.C., White, A., 2001. Valuing credit default swaps II: Modeling default correlations. Journal of Derivatives 8 (3), 12–21.

Katz, S., 1974. The price and adjustment process of bonds to rating reclassifications: A test of bond market efficiency. Journal of Finance 29 (May), 551–559.
Lando, D., Skodeberg, T., 2002. Analyzing rating transitions and rating drift with continuous observations. Journal of Banking and Finance 26, 423–444.
Longstaff, F.A., Mithal, S., Neis, E., 2003. The credit default swap market: Is credit protection priced correctly. Working Paper, Anderson School, UCLA, August.
Moody's, 2003. Default and Recovery Rates of Corporate Bond Issuers: A Statistical Review of Moody's Ratings Performance, 1920–2002. Moody's Investor's Services, February.
Pinches, G.E., Singleton, J.C., 1978. The adjustment of stock prices to bond rating changes. Journal of Finance 33 (1), 29–44.
Reinhart, V., Sack, B., 2001. The changing information content of market interest rates. BIS Working Paper, prepared for Autumn 2001, Central Bank Economist Meeting at the Bank for International Settlements, Basel, Switzerland.
Skinner, F.S., Townend, T.G., 2002. An empirical comparison of default swaps. International Review of Financial Analysis 11, 297–309.
Standard and Poor's, 2001. Outlook proves a reliable indicator of company default probability. Working Paper.
Steiner, M., Heinke, V.G., 2001. Event study concerning international bond price effects of credit rating actions. International Journal of Finance and Economics 6, 139–157.
Wansley, J.W., Glascock, J.L., Clauretie, T.M., 1992. Institutional bond pricing and information arrival: The case of bond rating changes. Journal of Business Finance and Accounting 19 (5), 733–750.