The Valuation of Corporate Coupon Bonds^{*}

Jens Hilscher[†], Robert A. Jarrow[‡], and Donald R. van Deventer[§]

January 3, 2022

Abstract

This paper shows that, for a sample of corporate bond prices, credit spreads and therefore discount rates of promised coupons and principal differ substantially. To better fit this stylized fact we propose and estimate a tractable, arbitrage-free valuation model for corporate bonds that includes a more realistic recovery rate process. Existing spread calculations assume that coupons and principal have the same seniority. Thus, these spread calculations are misspecified because they include recovery rates for coupons due after default. Misspecification errors resulting from such a coupon recovery assumption depend on recovery rates, coupons, maturity, and default probabilities and can be substantial in size. We provide two pieces of evidence in support of the no-coupon recovery model: (i) for a large sample of bond market transactions our model has lower pricing errors than one assuming recovery on all coupons, and (ii) the model's outperformance magnitude is closely linked to model misspecification errors.

1 Introduction

Credit spreads, the difference between yields to maturity on risky debt and safe government bonds, are commonly used both as measures of risk in bond prices and to price those bonds. In the corporate bond literature, for example Collin-Dufresne, Goldstein, and Martin (2001) identify drivers of variation in credit

^{*}Helpful comments from John Y. Campbell and seminar participants at the Bank of England, the Bank of Finland, the Board of Governors of the Federal Reserve, the European Bank for Reconstruction and Development, the International Monetary Fund, the World Bank, Brandeis University, the 2019 GEA meetings at the Frankfurt School of Finance and Management, and the University of Technology, Sydney are gratefully acknowledged. We thank Tengda Gong for providing excellent research assistance.

[†]University of California, Davis, California 95616. email: jhilscher@ucdavis.edu.

[‡]Samuel Curtis Johnson Graduate School of Management, Cornell University, Ithaca, N.Y. 14853 and Kamakura Corporation, Honolulu, Hawaii 96815. email: raj15@cornell.edu.

[§]Kamakura Corporation, Honolulu, Hawaii 96815. email: dvandeventer@kamakuraco.com.

spreads while Campbell and Taksler (2003) and Gilchrist-Zakrajšek (2012) explore determinants of credit spreads. Other work, e.g. Elton, Gruber, Agrawal, and Mann (2001) and Huang and Huang (2012), decompose a coupon bond's credit spread into its various components: the expected loss, a default risk premium, a illiquidity risk premium, and an adjustment for the deductibility of government bond income for state taxes.¹ A second stream of the literature prices bonds or related securities using a reduced form model (see Duffee (1999), Duffie, Pedersen, and Singleton (2003), Driessen (2005)). A careful reading of these papers shows that they all assume that a single credit spread or spread term structure can be used to value risky debt.

The underlying assumption (either implicit or explicit) is that a coupon bond is equivalent to a portfolio of risky zero-coupon bonds valued using a single spread or spread term structure. The number of zero-coupon bonds held in the portfolio corresponds to the promised coupons and principal with their maturities corresponding to the payment dates (see expression (4) in the text). Importantly, both promised coupons and principal payments are discounted using the same spread. For the credit spread estimation literature, this implicit assumption follows because all promised coupons and principal are included when computing a bond's credit spread. In the reduced form model literature, the recovery rate process utilized is the "recovery of market value (RMV)" due to Lando (1998) and Duffie and Singleton (1999), which implies this result. This pricing approach assumes that, when discounting, coupon and principal cash flows are treated the same, and, therefore, that both promised payments entitle the holder to a recovery in default. For subsequent discussion, we therefore call this approach the "full-coupon recovery" model.

In this paper we first compare bond valuations discounting all cash flows in the same way to one with different discount rates for coupons and principal. The first has one credit spread for all cash flows, while the second allows credit spreads for coupon and principal payments to differ. In our 120 thousand corporate coupon bond transactions data set, we find that coupon spreads are much larger than principal spreads. We also find that using different spreads results in much lower pricing errors – the average root mean squared pricing error is cut more than in half, dropping from 42 cents to 19 cents. These results provide initial evidence in support of valuing coupon and principal cash flows differently.

As shown by Jarrow (2004), a single term structure of risky zero-coupon bonds used for valuing coupon bonds is valid if and only if all of the risky zero-coupon bonds are of equal seniority and all have the same recovery rate in the event of default. However, this assumption is inconsistent with industry practice. After default, as evidenced by financial restructurings and default proceedings, only the

¹Other work explores determinants of risky debt yield spreads in the sovereign context, e.g. Duffie, Pedersen and Singleton (2003) or Hilscher and Nosbusch (2010).

bond's principal becomes due, and no additional coupon payments are made on or after the default date. This implies that coupon and principal payments cannot be valued using the same (single) credit spread or spread term structure and that basing a bond valuation model on this erroneous assumption of equal seniority will produce model prices that have misspecification errors. Our empirical results comparing using either one or two (principal and coupon) spreads imply that market participants understand the difference in seniority between coupon and principal payments and that this differences is reflected in market prices.

Consistent with our findings on coupon and principal spreads, industry practice has been confirmed in the recovery rate estimation literature where it has been shown that alternative recovery rate processes,² either the "recovery of face value (RFV)" or the "recovery of Treasuries (RTV)" formulations, provide a better approximation to realized recovery rates than does RMV (see Guha and Sbuelz (2005) and Guo, Jarrow and Lin (2008)).³ And, it is also well known that both the RFV and RTV recovery rate processes are consistent with a zero recovery on coupons promised after default. Therefore, these recovery rate processes do not imply the full-coupon recovery model. See Jarrow and Turnbull (2000), Longstaff, Mithal, and Neis (2005), Bielecki and Rutkowski (2002), chapter 13, and Collin-Dufresne and Goldstein (2001) for models with zero recovery on coupons promised after default. However, these papers do not consider the effect of varying seniority on single spread calculations or separate principal and coupon spreads; they do not provide a comprehensive analysis of the empirical implications of different pricing approaches for a large data set of coupon bond transaction prices, and they do not provide a clear and easy-to-calculate measure (the misspecification error's standard deviation) that identifies the effect of using a misspecified coupon recovery rate assumption on bond pricing errors.⁴

The purpose of this paper is to explore, both theoretically and empirically, the magnitude of the effect on bond prices of assuming zero recovery on coupons promised after default. In order to do so we derive an empirically tractable reduced form bond pricing model, the form of which is new to the literature. For subsequent discussion, we refer to our model as the "no-coupon recovery" model and the

 $^{^2 \}mathrm{See}$ Bielecki and Rutkowski (2002), Chapter 8 for a discussion of these different recovery rate processes.

 $^{^3 \}rm{See}$ also Guha, Sbuelz and Tarelli (2020), who provide evidence in support of RFV when studying high-yield bond duration.

⁴Bakshi, Madan and Zhang (2006) use the Lehman Bond price data set to compare different recovery assumptions for a sample of 25 BBB-rated bonds over a nine-year period. They find that pricing errors decline when choosing the RTV or RFV rather than RMV specification. Our paper uses a much larger data set, demonstrates the effect on pricing of using different spreads for principal and coupons, explores the drivers of model misspecification errors on pricing both theoretically and empirically, and estimates the effect of illiquidity on prices. We also provide direct evidence of prices reflecting no recovery on coupons by looking at prices of bonds immediately after default.

existing model as the "full-coupon recovery" model. We derive an intuitive and straightforward-to-implement pricing formula in which model prices depend only on the risk-free term structure, the term structure of default probabilities, and two parameters that we estimate – the recovery rate and a parameter capturing illiquidity. We also show that the bond can be valued using two "building block" securities: survival digitals, which pay off only if there is no default, and default digitals, which pay off only in the event of default. The decomposition is useful both to build pricing intuition and for the empirical implementation. Our model implies that, consistent with our approach of using two spreads, bond prices can be computed using two different discount rate functions – spread curves – one for coupons and one for principal.

We perform a calibration of the model to demonstrate that the spreads appropriate for discounting coupon and principal payments can be quite different. We also calculate pricing errors of the no-coupon recovery and the full-coupon recovery models. These are model misspecification errors, which result from the erroneous assumption of positive recovery of coupons after default. Misspecification errors are larger if recovery rates, default probabilities, maturity or coupon payments are larger. For example, for a 10-year bond with a recovery rate of 40%, a coupon of 2.61%, and an annual default probability of 1%, the full-coupon recovery model will assign a price that is \$0.50 too large. If it is a 30-year bond, the price is \$3.61 too large, a substantial difference relative to the correct price, which is equal to par. We calculate exact misspecification errors and also provide an approximate formula that can be used to estimate the misspecification error magnitudes. In this estimate, misspecification errors are proportional to the recovery rate and the coupon size, and they are approximately proportional to the default probability and the square of the number of coupon payments, which results in a close relation to maturity.

We then return to our empirical work and present direct evidence of a difference in seniority between principal and coupons. We provide three examples of issuers that have filed for bankruptcy: Lehman Brothers, Pacific Gas and Electric (PG&E), and Weatherford International. We use both the misspecified full-coupon recovery and our no-coupon recovery model to price the bonds. We find that pricing errors from using the full-coupon recovery model are between five and ten times as large as the no-coupon recovery model's pricing errors. Observed prices are thus consistent with market participants assuming zero recovery on coupons and they are inconsistent with the assumption of equal recovery. This analysis provides independent evidence in support of industry pricing practice implying zero recovery on coupons after default.

Expanding our analysis to the full data set, we perform a comparative analysis of the no-coupon and the full-coupon recovery models. Our sample consists of daily market prices for a collection of liquidly traded bonds over the period from September 2017 through January 2021. We fit both models to the data separately. If market prices reflect zero recovery of coupons after default, the no-coupon recovery model will outperform the full-coupon recovery model. Indeed, we find that, in the full sample, the average no-coupon recovery model's root mean squared pricing error is 3.2 cents smaller than the full-coupon recovery model pricing error, and that the difference is statistically significant.

Although it is useful to calculate the average outperformance, this is not particularly informative. Our model predicts that the size of the outperformance is directly related to the size of the misspecification error – the pricing effect of switching from assuming zero recovery of coupons to full recovery. A small average level of outperformance may therefore simply be the result of using a sample in which misspecification errors are small. Instead, the relevant test of our model is whether or not the misspecification error can explain variation in the magnitude of no-coupon recovery model outperformance. If, for example, a bond is of short maturity with only a few coupons and with a low default probability, the two models predict the same price (the misspecification error is close to zero) and the no-coupon recovery model will not outperform. But, if the maturity is long and the default probability is substantial (and zero recovery of coupons is reflected in the data) the no-coupon recovery model's outperformance will be large.

Our empirical approach thus proceeds in two steps. First, we calculate misspecification errors and then check that the outperformance of the no-coupon recovery model depends on them. Misspecification errors are large (in our data, 5% are larger than \$1.17). However, the median misspecification error is small; it is equal to 2.4 cents. There are thus some bonds for which zero recovery on coupons is highly relevant and ignoring it will result in large pricing errors, while for other bonds the distinction is less relevant.

Second, we compare pricing errors of the two models and relate them to misspecification errors. This amounts to a horse race between the models, but one that not only tests average outperformance but also tests whether our model's prediction regarding relative outperformance is present in the data. As pointed out, we find evidence of model outperformance in the full sample. More importantly, as we focus on higher and higher misspecification error observations, the no-coupon recovery model's outperformance becomes larger. That is, outperformance is larger when default probability and recovery rate are larger, when maturity is longer, and when coupons are larger. In other words, our model accurately forecasts when zero recovery of coupons is important for pricing.

Since we are pricing portfolios of bonds for each issuer day and since the full-coupon recovery model can adjust parameters, what actually matters is not the average level of the misspecification error but rather the within issuer day misspecification error's standard deviation. Regressing the difference in the full-coupon and the no-coupon recovery models root mean squared errors (i.e. model

outperformance) on the misspecification error standard deviation, the coefficient is highly significant and the regression has a R^2 of 57%. When focusing attention on the top quartile of misspecification error standard deviation observations, the pricing error difference increases from 3.2 cents for the full sample to 12 cents. For the top decile, outperformance increases to 24 cents. Model outperformance happens where our model predicts, providing evidence that market participants are pricing bonds according to the no-coupon rather than the full-coupon recovery model.

In the process of fitting the no-coupon recovery model to the data we estimate implied recovery rates and illiquidity discounts. In a final step we show that variation in both of these parameters track what we expect. Average recovery rates, though a little higher, are generally in line and consistent with previous work. Variation in the illiquidity effect we estimate has a strikingly close relationship to the Aaa-Treasury spread, although outside of the crisis periods the illiquidity discount is slightly smaller than the Aaa-Treasury spread. When the Covid-19 pandemic hit markets in March of 2020, both the Aaa-Treasury and our illiquidity measure spiked (also see Kargar et al. (2021)). The close relationship between our new measure of liquidity and the Aaa-Treasury spread, neither of which uses the same data nor methodology to compute, provides independent validation of our pricing model.

The outline of the paper is as follows. Section 2 presents a simple two period model to price bonds and calculates coupon- and principal-specific spreads. Section 3 presents the model for valuing risky coupon bonds. Section 4 discusses the estimation procedures, and Section 5 presents some illustrative pricing results for three companies that filed for bankruptcy. Section 6 presents a comparative analysis of two alternative pricing models, discusses variation in model fit and parameters over time, and shows robustness to differing model implementations. Section 7 concludes.

2 A Simple Credit Spread Model

Our objective is to price bonds of the same issuer but with different maturities and coupons. Short-maturity and low-coupon bonds derive value mainly from the bond's principal. For longer maturity and higher-coupon bonds, the present value of coupons plays a larger role. If coupons have a zero recovery after default, while the principal payment has a positive recovery, both cash flows will not have the same discount rate. Using the same credit spread for both will result in an inability to price bonds with different maturities and coupons. On the other hand if, as the standard literature on credit spreads implicitly assumes, both have the same seniority, one spread will suffice. A priori it is not clear if the effect we are focusing on is empirically large or small. Spreads appropriate for discounting coupons and principal may be similar and pricing errors using either one or two spreads may also be similar. We note that we are interested in pricing multiple bonds simultaneously. It is, of course, possible to calculate a bond-specific yield to maturity and therefore a bond specific credit spread. This, however, does not provide a pricing methodology, but it is simply a transformation of the price into another quantity.

2.1 A Comparison of One versus Two Credit Spreads

Before proceeding with our full model and estimation, we first examine the difference in the two pricing approaches empirically. We price bonds using either a single spread or two spreads - one to discount coupon payments and one to discount principal payments. When pricing a single bond, one can always find a credit spread that will price the bond perfectly. The model may be mis-specified, but it will nevertheless be able to fit the data. However, when pricing two or more bonds using the same credit spread, it is possible to compare the two pricing approaches.

Our data, which we describe in more detail in Section 4, consists of 120 thousand observations from September 2017 to January 2021. The data is for 181 issuers. Table 1 Panel A reports summary statistics for the full sample. Ninety percent of the data have ratings between BB+ and A^{-5} maturities between 0.3 and 8.1 years,⁶ and coupons between 1.7 and 5.7 percent. The average observation comes from an issuer day with seven observations and ninety percent of observations are from issuer-days with between two and 14 observations. The average bond is priced in a sample that has an average maturity range of five years, and ninety percent of observations have maturity ranges between 1.3 and 8.9 years. Our sample is therefore appropriate to study how default risky coupon bonds should be priced.

We price bonds assuming either a single spread or two spreads, one to discount coupons and one to discount principal payments. Specifically, we minimize the volume weighted squared pricing errors. When a single credit spread is used, the average spread is equal to 100 basis points, while the median is equal to 75 bps. If instead we use two spreads, the mean principal spread is equal to 54 bps, while the mean coupon spread is 746 bps. The market, therefore, uses substantially higher discount rates for coupons than it does for principal payments. This finding supports the idea that principal payments are safer because they deliver potentially large recovery values in the event of default. Coupon payments, in contrast, do

⁵The sample consists primarily of investment grade bonds since many high yield bonds have call features, all of which are excluded. An analysis of callable bonds goes beyond the scope of this paper.

⁶Our main model implementation is based on default probabilities that extend to a maturity of ten years, and so we restrict attention to observations with that maximum maturity.

Table 1: Summary Statistics, Credit Spreads, Implied Default Probabilities and Recovery Rates

This table reports detailed summary statistics for the full sample of bond prices. Panel A contains observationlevel statistics. Rating is the S&P issuer credit rating; maturity range is the difference for each issuer day between the maximum and minimum maturity; maturity SD is the issuer day maturity standard deviation; number of observations counts how many bond prices are in the data set each issuer day (we require a minimum of two). Below, we report the credit spread that minimizes the volume weighted squared pricing error for each issuer-day (Spread); Principal and Coupon spreads are also chosen to minimize the squared pricing error. Implied default probability and implied recovery rate are based on the coupon and principal spreads; details of the calculation are provided in the text. Panel B reports maturity range and model fit statistics at the issuer-day level. We report root volume-weighted mean squared errors and volume-weighted R-squared levels. R-squared is calculted relative to a model using the average bond price as the predictor and can therefore be negative. For those cases we report zero.

		Panel	A: Bond level st	tatistics		
	Coupon	Maturity	Maturity range	Maturity SD	Rating	Number of observations
Mean	3.2	3.1	5.0	2.0	A-	7
SD	1.2	2.3	2.3	0.8	two notches	4
p5	1.7	0.3	1.3	0.7	AA-	2
p50	2.9	2.5	5.0	1.9	A-	6
p95	5.7	8.1	8.9	3.4	BB+	14
	Spread	Principal Spread	Coupon Spread	Implied def prob	Implied recovery rate	
Mean	100	54	746	7.0%	85%	
SD	119	115	625	5.2%	25%	
p5	31	0	114	1.1%	0%	
p50	75	33	620	6.0%	94%	
p95	257	171	1864	17.1%	100%	
Number of is	suer days: 120.786	5				

Panel B: Issuer-day level statistics							
	Spread BMSE	C and D DMSE	Spread	C and P			
	Spread RIVISE	C and P RIVISE	R-squared	R-squared			
Mean	0.42	0.19	0.84	0.96			
SD	0.39	0.30	0.27	0.13			
p5	0.03	0.00	0.00	0.86			
p50	0.34	0.10	0.95	0.99			
p95	1.14	0.66	1.00	1.00			
Number of issuer days: 25,018							

not pay off in default and therefore need a larger discount rate.

Table 1 Panel B reports statistics measuring model fit. The average root mean squared error when using a single spread for principal and coupons is equal to 0.42. If we use different spreads for principal and coupons, the RMSE drops by more than one half to 0.19. The difference is even larger when looking at median pricing errors which drop from 0.34 to 0.1 when using two spreads. The pattern is also reflected in the volume weighted R^2 measures.⁷ The median R^2 for the single spread model is 95% while it is 99% for the two spread model. The mean R^2 is 84% for the single spread model and 96% for the model allowing for principal and coupon spreads. We check that this difference is not driven by some issuer days with two observations that may be able to be fit perfectly using the two spread model (the model is nonlinear and so a perfect fit is not guaranteed). When restricting attention to issuer days with at least five bonds, the model fit as measured by R^2 increases and is equal to 88% on average for the single spread model and 97% for the two spread model. The mean root mean squared error for the two approaches are slightly larger and equal to 0.53 and 0.26, respectively. As an additional check we estimated spreads at the monthly rather than the daily level. This allows for much larger groups of observations, making it less likely that there are only very few observations priced for a given spread. Our results are robust to this change.

2.2 Implied Default Probabilities and Recovery Rates in a Two Period Model

In this section we introduce a two period model that reflects the different seniority of principal and coupons. This model allows us to identify the default probability separately from the recovery value. The intuition underlying the more general model can be seen in this two period model, and it enables us to take a closer look at the principal and coupon spread data.

Assume that a bond pays back both the coupon and principal next period. If default occurs before maturity, the bond pays a fractional recovery value of principal and the coupon has a zero recovery. The appropriate discount rates and spreads for the two promised cash flows can be found by pricing two building block securities. The first is a survival digital that pays one dollar next period if no default has occurred, with value C. The second is a default digital that pays one dollar next period if default occurs, with value R. Since the results are mutually exclusive and exhaustive, a portfolio of both securities pays one dollar next period for sure.

 $^{^7\}mathrm{Using}$ two spreads allows the model one additional degree of freedom. To take this into account we have also

calculated adjusted R^{2} 's (unreported) which result in the same pattern.

For simplicity, and for the purposes of this example only, we assume a zero risk-free interest rate, frictionless markets, and no-arbitrage. No arbitrage implies the existence of risk neutral probabilities that can be used for valuation. Let Q be the risk neutral probability of default. Then C + R = 1 and C = (1 - Q), which is the risk neutral probability of survival. Note that the value only depends on the default probability.

Now, the bond's value V = C + R * (recovery rate) because this corresponds to its payoffs next period in terms of the two digital securities and a recovery rate. The appropriate discount rate to be used for principal payments, therefore, depends on the recovery rate. Using the two digitals, we can rewrite the bond's value as V = (1 - Q) + Q * (recovery rate), enabling us to calculate the default probabilities implied by the coupon spreads and the recovery values implied by principal spreads.

Table 1 Panel A reports these results. Median implied probabilities are equal to 6.0%. Median implied recovery rates are equal to 94%. The very large recovery rates result from the large difference in coupon and principal spreads (for a recovery rate of zero, the two spreads are the same). We note that default probabilities implied using this simple, stylized approach are very high. Later, when we estimate the full model we introduce an illiquidity discount which captures the effect of corporate bonds trading at lower prices due to trading frictions. Our simple model in this section does not allow for such an effect, likely resulting in artificially high implied default probabilities. Similarly, implied recovery rates are higher than what previous studies have found (Jankowitsch, Nagler, Subramanyam (2014)). Nevertheless, the spread-implied default probabilities and recovery values show that the market expects high recovery values and that bonds generally trade below their frictionless value, supporting the introduction of an illiquidity discount.

We next discuss our full model. When estimating the full model, instead of using bond price data to back out the implied default probabilities, we use independent estimates of the default probabilities. And, we also relax the assumption of frictionless markets and estimate an illiquidity discount.

3 The Pricing Model

This section presents the pricing model, which is based on the reduced form model of Jarrow and Turnbull (1995).⁸ To streamline the exposition, the details of the mathematical formulation are contained in the appendix. We assume that traded in the economy are default-free zero-coupon bonds of all maturities, a default-free money market account, and a risky coupon bond (to be described later). The market is assumed to be frictionless and competitive. Both the frictionless

⁸We note that Jarrow and Turnbull (2000) contains a reduced form model with zero recoveries paid on coupons due after default.

and competitive market assumptions are relaxed, subsequently, when we add a illiquidity discount to the valuation formula (see expression (11) below).

The default-free money market account earns interest continuously at the default-free spot rate of interest, r_t . The money market account's time t value is denoted by

$$B_t = e^{\int_0^t r_s ds} \tag{1}$$

with $B_0 = 1$. We let the time t value of a default-free zero-coupon bond paying a dollar at time T be strictly positive and denoted by p(t,T) > 0.

We consider a firm that issues a bond with a coupon of C dollars, a face value equal to L dollars, and a maturity date T. The bond pays the C dollar coupons at intermediate dates $\{t_1, ..., t_m = T\}$, but only up to the default time τ . For notational convenience, let the current time $t = t_0$. If default happens in the time interval $(t_{k-1}, t_k]$, then the bond pays a stochastic recovery rate of $\delta_{t_k} \in [0, 1]$ at time t_k on the notional of L dollars.⁹ It is important to note that default can happen anytime within this interval, but the payment only occurs at the end. If default does not happen, the face value of L dollars is repaid at time T.

3.1 Risk Neutral Valuation

To value the risky coupon bond, we assume (i) that the markets for both the default-free coupon bonds and the risky coupon bond are arbitrage-free and (ii) that enough credit derivatives trade on the risky firm so that the enlarged market is complete (see Jacod and Protter (2010) for a set of sufficient conditions on an incomplete market such that the expanded market is complete). Given the trading of credit default swaps, this is a reasonable approximation.

Under these assumptions, by the second fundamental theorem of asset pricing (see Jarrow and Protter (2008)), risk-neutral valuation applies. Denote the time $t \leq t_1$ value of the coupon bond as v_t , then

$$v_{t} = \sum_{k=1}^{m} CE^{\mathbb{Q}} \left[\mathbf{1}_{\{\tau > t_{k}\}} e^{-\int_{t}^{t_{k}} r_{u} du} |\mathcal{F}_{t} \right] + LE^{\mathbb{Q}} \left[\mathbf{1}_{\{\tau > T\}} e^{-\int_{t}^{T} r_{u} du} |\mathcal{F}_{t} \right]$$
$$+ \sum_{k=1}^{m} LE^{\mathbb{Q}} \left[\mathbf{1}_{\{t_{k-1} < \tau \le t_{k} \le T\}} \delta_{t_{k}} e^{-\int_{t}^{t_{k}} r_{u} du} |\mathcal{F}_{t} \right]$$
(2)

where \mathbb{Q} denotes the risk-neutral probabilities.

This formula reflects the expected discounted value of the random cash flows to the risky coupon bond using the risk-neutral probabilities \mathbb{Q} . The discount

⁹In practice, a portion of the next coupon payment after default represents some accrued interest, earned, but not yet paid. This accrued interest has a recovery rate associated with it. With a slight loss of generality we exclude this accrued interest payment in the stochastic recovery rate δ_{t_k} defined above. We appreciate the comments from a law firm, Morrison & Foerster, in this regard.

rate is the default-free spot rate r_t . The adjustment for risk is via the use of the risk-neutral probabilities, instead of the statistical probabilities.

The cash flows correspond to the coupon payments $C1_{\{\tau>t_k\}}$ for $k = 1, \ldots, m$ and the principal $L1_{\{\tau>T\}}$, but only if no default occurs prior to the payment date. A recovery payment on the principal is made in the event that default occurs within the time interval $(t_{k-1}, t_k]$, $1_{\{t_{k-1} < \tau \leq t_k \leq T\}}\delta_{t_k}L$, summed across all the intervals over the bond's life. Note that no recovery payments are included for any coupon payments occurring after the default date.

To facilitate analytic tractability, and with a minor loss of generality, we assume that the default-free spot rate r_t , the default time τ , and the recovery rate process δ_t are independent under \mathbb{Q} given the information at time t that doesn't depend on the histories of these three processes. This is a weak assumption on the evolutions of the default-free spot rate, the default time, and the recovery rate because it imposes very little structure on their evolutions under the statistical probabilities.

Under the statistical probabilities, these processes need not be independent. Hence, nonzero pairwise correlations between the observed default-free spot rate, the default time, and the recovery rate processes, which are realizations under the statistical probabilities, are not excluded by this assumption. And, it is well known that non-zero correlations across the default-free spot rate, default times, and recovery rates have been observed in historical data.

Using the independence assumption, equation (2) simplifies to

$$v_t = Lp(t,T) \left[1 - Q(t,T) \right] + (C - Ld_t) \sum_{k=1}^m p(t,t_k) \left[1 - Q(t,t_k) \right] + Ld_t \sum_{k=1}^m p(t,t_k) \left[1 - Q(t,t_{k-1}) \right]$$
(3)

where

$$d_t = E^{\mathbb{Q}}\left[\delta_\tau \left| \mathcal{F}_t \right]\right]$$

is the time t futures price for a contract receiving the recovery rate at time T^* (see the appendix for the details) and $Q(t,t_i) = Prob^{\mathbb{Q}} [\tau \leq t_i | \mathcal{F}_t]$ is the risk neutral probability of default before t_i . We refer to this price as the "no-coupon recovery" model to emphasize that it captures no recovery on coupons after default. We call d_t the recovery rate futures price. It is the relevant measure that results in value today derived from future fractional recovery of principal, and it is what we estimate in our empirical work. As shown in the appendix, it is expected to be slightly larger than the recovery rate if paid at time t, δ_t .

In this form it is easy to see that the value of this coupon bond is not equal to the sum of the coupons and principal times the value of a collection of zero-coupon bonds from a single risky term structure that pay the fractional recovery rate in the event of default. Let $D(t, t_k)$ denote the time t value of such a risky zerocoupon bond promising to pay a dollar at time t_k for k = 1, ..., m with recovery rate δ_t in default. Then,

$$v_t^{full\ coupon} = \sum_{k=1}^m CD(t,t_k) + LD(t,T) = Lp(t,T) [1 - Q(t,T)] + (C - Ld_t) \sum_{k=1}^m p(t,t_k) [1 - Q(t,t_k)] + Ld_t \sum_{k=1}^m p(t,t_k) [1 - Q(t,t_{k-1})] + \sum_{k=1}^m C(m+1-k) d_t p(t,t_k) [Q(t,t_k) - Q(t,t_{k-1})].$$
(4)

This expression is called the "full-coupon recovery model." As pointed out in the introduction, the RMV recovery rate assumption is a sufficient condition for this pricing model. The difference between this model and ours is that expression (4) contains terms omitted in expression (3). The additional terms in expression (4) are the recovery values for the coupons that would have been paid after the default date, i.e. $\sum_{k=1}^{m} C(m+1-k)d_tp(t,t_k)[Q(t,t_k) - Q(t,t_{k-1})]$.¹⁰ These terms are included because all cash flows are discounted using a single risky term structure of zero-coupon bonds. We note that, given the different seniorities of the principal and coupons, we anticipate the no-coupon recovery model to outperform the full-coupon recovery model when simultaneously pricing more than one bond. If one instead assigns different recovery rates to different bonds of the same firm traded on the same day, the pricing method extends. However such an assumption is in direct contradiction to industry practice, which is to pay the same recovery rate on equal seniority bonds (see Section 5).

3.2 Survival and Default Digital Securities

To simplify expression (4), we introduce two simple credit derivatives: a survival and default digital. These building block securities provide additional intuition and insight to our pricing approach and were first discussed in the literature by Madan and Udal (1998). These securities are generalized versions of those discussed in Section 2.

At time t, consider the time interval $[t_{k-1}, t_k]$ for k = 0, ..., m.¹¹ We define a survival and a default digital as follows

• (A Survival Digital) This security pays \$1 at time t_k only if default occurs after t_k . The value of this security at time $t < t_k$ is

$$z(t,t_k) = E^{\mathbb{Q}} \left[\mathbb{1}_{\{\tau > t_k\}} e^{-\int_t^{t_k} r_u du} \left| \mathcal{F}_t \right] \right].$$
(5)

• (A Default Digital) This security pays \$1 at time t_k if default occurs within $(t_{k-1}, t_k]$. The value of this security at time $t < t_k$ is

$$x(t,t_k) = E^{\mathbb{Q}} \left[\mathbb{1}_{\{t_{k-1} < \tau \le t_k\}} e^{-\int_s^{t_k} r_u du} \left| \mathcal{F}_t \right] \right].$$
(6)

¹⁰This term follows because if default occurs during the time interval $(t_{k-1}, t_k]$, the remaining future coupons are $\sum_{j=k}^{m} C = (m+1-k)C$. In the full coupon recovery model, one gets a recovery payment on all the remaining coupons.

¹¹Note that when $k = m, t_m = T$.

Using the conditional independence assumption, we can write these as

$$z(t, t_k) = p(t, t_k)[1 - Q(t, t_k)]$$
(7)

and

$$x(t,t_k) = p(t,t_k)[Q(t,t_{k+1}) - Q(t,t_k)].$$
(8)

The identity $1_{\{t_{k-1} < \tau\}} = 1_{\{t_{k-1} < \tau \le t_k\}} + 1_{\{\tau > t_k\}}$ implies that t^{12}

$$p(t,t_k)[1-Q(t,t_{k-1})] = x(t,t_k) + z(t,t_k).$$
(9)

The left side represents the present value of the t_k maturity zero-coupon bond at time t, which is received only if there is no default before or at time t_{k-1} . The right side is the present value of the default digital for the interval $(t_{k-1}, t_k]$ and the survival digital with maturity t_k . This valuation relationship corresponds to the discussion in Section 2.1, where we noted that a portfolio of a default and survival digital with maturities one period into the future pay \$1 for sure.

Using these two digitals, expression (3) can be written as

$$v_t = \sum_{k=1}^m Cz(t, t_k) + Lz(t, T) + Ld_t \sum_{k=1}^m x(t, t_k).$$
(10)

This expression shows that a risky coupon bond can always be decomposed into a portfolio of survival and default digitals.¹³

3.3 An Illiquidity Discount

Corporate bond markets are illiquid relative to Treasury bonds or exchange traded equities. This illiquidity implies that corporate bond prices may reflect an illiquidity discount (see Jarrow and Turnbull (1997), Duffie and Singleton (1999), Cherian, Jacquier, and Jarrow (2004)). Note also the preliminary evidence discussed in Section 2 supporting the presence of such a discount. An illiquidity discount modifies the previous valuation formula to implicitly incorporate the impact on pricing due to transaction costs and trading constraints.

¹²Indeed,

$$E^{\mathbb{Q}}\left[1_{\{t_{k-1}<\tau\}}e^{-\int_{s}^{t_{k}}r_{u}du}\left|\mathcal{F}_{t}\right.\right] = E^{\mathbb{Q}}\left[1_{\{\tau\leq t_{k}\}}e^{-\int_{s}^{t_{k}}r_{u}du}\left|\mathcal{F}_{t}\right.\right]$$
$$+E^{\mathbb{Q}}\left[1_{\{\tau>t_{k}\}}e^{-\int_{s}^{t_{k}}r_{u}du}\left|\mathcal{F}_{t}\right.\right],$$

which simplifies to the given expression.

¹³We note that the valuation of corporate coupon bonds is only one example of the advantages of pricing using these building block securities. Other securities that depend on default, e.g. CDS contracts, sovereign bonds, and "bail in" bonds can also be priced using the building block securities as well.

It is important to note that transactions costs (including bid/ask spreads) are a special case of an illiquidity cost paid when trading, which are implicitly included via an illiquidity discount (see Cetin, Jarrow and Protter (2004) for the theoretical justification of this statement). Similarly, taxes paid on coupons and capital gains can also be interpreted as a type of transaction cost, and hence they too are implicitly included in the illiquidity discount as well.¹⁴

We apply the illiquidity discount function $e^{\alpha_t(T-t)}$ symmetrically to all the cash flows promised to the coupon bond. This symmetry enables similar illiquidity discount impacts across different coupon bonds issued by the same credit entity. Given this, we can rewrite the coupon bond's value as

$$v_t^{liq} = \sum_{k=1}^m Cz(t, t_k) e^{\alpha_t(t_k - t)} + Lz(t, T) e^{\alpha_t(T - t)} + Ld_t \sum_{k=1}^m x(t, t_k) e^{\alpha_t(t_k - t)}.$$
(11)

It is important to emphasize that in the subsequent estimation, both the recovery rate d_t and the illiquidity parameter α_t are stochastic, hence, they can vary randomly across time due to changing market conditions. Our estimation procedure allows for these estimated parameter values to reflect this randomness.¹⁵ Expression (11) is the valuation model estimated in the next section.

3.4 Comparative Statics - Misspecification Errors

This section builds intuition for misspecification errors when using the full-coupon recovery model, expression (4) instead of the no-coupon recovery model, expression (3). Recall that the misspecification error, the difference between the full-coupon and no-coupon recovery model prices, is equal to $\sum_{k=1}^{m} C(m+1-k)d_t p(t,t_k)[Q(t,t_k)-Q(t,t_{k-1})]$. Note that these misspecification errors are always positive.

We next quantify the magnitudes of these misspecification errors and we provide a simple approximation that allows us to relate the misspecification errors to the model's inputs. Later, we relate the predicted misspecification errors to patterns in the data.

¹⁴The complication of explicitly including illiquidity costs (transaction, taxes) into the model is that different traders face different taxes and transaction costs based on their trading activities. Consequently, to determine a market price, an equilibrium model is needed. Equilibrium models are notoriously ladened with unrealistic assumptions. Furthermore, an argument can be made that the marginal trader, who determines the market price, is the lowest illiquidity cost trader. Here, we note that many institutions pay small transaction costs and there do exist non-taxable institutions that purchase corporate debt.

¹⁵We use implicit estimation at a fixed time t allowing α_t to depend on the information available at time t.

Misspecification Error Determinants

For illustrative purposes we make the following simplifying assumptions: (1) coupon bonds are priced on coupon dates, (2) the risk-free term structure of interest rates and the term structure of default probabilities are flat,¹⁶ (3) the coupon is set so that the no-coupon recovery model's bond price is equal to par, and (4) that there is no illiquidity discount ($\alpha_t = 0$), though we relax this last assumption when we consider the effect of model parameters on spreads later in this section. Combined, these imply that the misspecification error is fully determined by the bond's maturity, the issuer's default probability, the recovery rate, and the risk-free rate.

Table 2 reports misspecification errors for the different inputs, assuming that the risk-free term structure is flat at 2%. The par value of the bond is set to 100. As expected, misspecification errors increase with the bond's maturity, the issuer's default probability, and the recovery rate. For short maturity 2-year bonds, the misspecification errors are less than 0.10, while the misspecification errors are equal to or above 0.50 for the 10-year bonds. For 30-year bonds the misspecification error is equal to 12.32 for a 30-year bond with a default probability of 2% and a recovery rate of 80%.

The misspecification errors, which represents the present value of the recovery on all of the coupon payments after default, depend on the default probability, the recovery rate, the coupon payment, and the bond's maturity. For the first coupon payment the present value of the payoff in the event of default is equal to the discounted value of the product of the coupon rate, the recovery value, and the probability of default, i.e. $Cd_t p(t, t_1)Q(t, t_1)$. For the second coupon, default can occur either in the first or in the second period. The present value of the payoff in period one is the same as for the first coupon. The conditional expectation at time one of the payoff in period two is also the same as for the first coupon. However, this payment still needs to be discounted back to time zero and adjusted to take into account that the firm must have survived to period one. For longer time periods, a similar logic applies. This implies that the misspecification error is exactly proportional to the product of the coupon payment and the recovery value. In contrast, the total misspecification error is only approximately proportional to the error resulting from the first coupon. The reason is that the probability of survival depends on the default probability and because payments farther into the future will be discounted by a larger amount.

Next we consider the effect of the bond's maturity or, equivalently, the number of coupon payments. Ignoring discounting and the adjustment for the probability of survival, the expected recovery payment from the second coupon is twice as large as for the first coupon. The reason is that default on the second coupon can happen

¹⁶The assumption of a flat term structure is for illustrative purposes only. In the empirical implementation we use a term structure of default probabilities, which is not assumed to be flat.

Table 2: Misspecification Error Determinants

In this table we calculate prices based on the no-coupon recovery and the coupon recovery models. We assume a flat risk-free term structure of 2%, a flat default probability term structure and different maturities and recovery rates. Coupons are chosen so that bonds trade at par. We report the misspecification error resulting from using the coupon recovery instead of the reduced form model (column five, "missspec. error").

maturity	recovery	default probability	coupon	misspec. error
2	0.4	0.01	2.61%	0.03
2	0.4	0.02	3.23%	0.06
2	0.8	0.01	2.21%	0.04
2	0.8	0.02	2.42%	0.09
5	0.4	0.01	2.61%	0.14
5	0.4	0.02	3.23%	0.33
5	0.8	0.01	2.21%	0.23
5	0.8	0.02	2.42%	0.50
10	0.4	0.01	2.61%	0.50
10	0.4	0.02	3.23%	1.19
10	0.8	0.01	2.21%	0.84
10	0.8	0.02	2.42%	1.78
30	0.4	0.01	2.61%	3.61
30	0.4	0.02	3.23%	8.21
30	0.8	0.01	2.21%	6.10
30	0.8	0.02	2.42%	12.32

either in the first or the second period. For a bond receiving m coupon payments the factor is therefore m(m+1)/2. This means that the approximate total error is equal to $Cd_tp(t,t_1)Q(t,t_1)m(m+1)/2$. We later use this approximate error to identify portfolios of bonds that are likely to be mispriced by the full-coupon recovery model.

To summarize, the misspecification error is zero if the recovery rate, the default probability, or the coupon payment is zero. The error grows approximately with the square of the number of coupon payments and is exactly proportional to the product of the coupon payment and the recovery rate. Thus, bonds with significant recovery values, default probabilities, and with intermediate to long maturities will have significant misspecification errors.

Pricing with Two Credit Spread Curves

Pricing bonds with a single credit spread term structure is based on the incorrect assumption that coupons and principal have the same recovery. However, as pointed out in Section 2, this means that bonds should be priced using two spread curves, one for discounting coupons and one for discounting the principal. If there is a misspecification error using this full-coupon recovery model to price bonds, then the two curves will be different.

Table 3 provides some illustrative examples of credit spread curves. We use the same methodology as before to illustrate these spreads, varying recovery rates and default probabilities as in the previous table. The only difference is that here we introduce the effect of the illiquidity discount. Panel A reports principal spreads, Panel B reports coupon spreads. As long as there is a positive recovery, coupon spreads lie above principal spreads since the latter will be worth more and thus are discounted less. The difference between coupon and principal spreads is close to the product of the default probability and the recovery rate, which follows from the misspecification error relation given above, where, for the first coupon, the misspecification error is equal to $Cd_t p(t, t_1)Q(t, t_1)$. A larger default probability makes all spreads higher. If there is no illiquidity discount, coupon spreads are approximately equal to the default probability, and since differences relative to principal spreads depend on the default probability, frictionless spreads are approximately proportional to the default probability. The effect of the illiquidity discount is seen to be symmetric, affecting all cash flows equally. Indeed, both credit spreads increase by the amount of the illiquidity discount.

It is useful to note that the standard spread calculation, which assumes equal seniority of principal and coupons, will result in a spread that cannot be used to discount either cash flows with zero or positive recovery. For the former (the coupons), the spread will be too low and for the latter (the principal) it will be too high. Thus, using a single spread (or spread curve) to price a new bond with a different maturity or coupon will result in misspecification errors. In addition,

Table 3: Illust	rative Credit	Spread	Curves
-----------------	---------------	--------	--------

This table reports spreads appropriate for discounting coupons and principal (C and P) for various maturities, default probabilities, recovery rates, and liquidity values. As in the previous table, we assume a flat risk-free term structure of 2% and a flat default probability term structure. Panel A reports spreads appropriate for disounting Principal, Panel B reports spreads appropriate for discounting coupons. 0.01 0.02 0.02 0.01 0.01 0.02 0.02

Def prob	0.01	0.01	0.02	0.02	0.01	0.01	0.02	0.02
Recovery	0.4	0.4	0.4	0.4	0.8	0.8	0.8	0.8
Liquidity	0	-0.5	0	-0.5	0	-0.5	0	-0.5
Maturity								
			Panel A: I	Principal sp	reads			
1	0.61%	1.11%	1.21%	1.72%	0.20%	0.70%	0.40%	0.90%
2	0.60%	1.11%	1.20%	1.70%	0.19%	0.69%	0.38%	0.88%
3	0.59%	1.10%	1.19%	1.69%	0.18%	0.68%	0.36%	0.85%
4	0.59%	1.09%	1.17%	1.67%	0.17%	0.67%	0.34%	0.83%
5	0.58%	1.09%	1.16%	1.66%	0.16%	0.66%	0.32%	0.80%
6	0.58%	1.08%	1.14%	1.64%	0.15%	0.65%	0.30%	0.78%
7	0.57%	1.07%	1.13%	1.62%	0.14%	0.63%	0.28%	0.75%
8	0.57%	1.06%	1.12%	1.61%	0.13%	0.62%	0.26%	0.73%
9	0.56%	1.06%	1.10%	1.59%	0.12%	0.61%	0.24%	0.70%
10	0.55%	1.05%	1.09%	1.57%	0.11%	0.60%	0.22%	0.68%
			Panel B:	Coupon spi	reads			
1	1.02%	1.52%	2.04%	2.55%	1.02%	1.52%	2.04%	2.55%
2	1.02%	1.52%	2.04%	2.55%	1.02%	1.52%	2.04%	2.55%
3	1.02%	1.52%	2.04%	2.55%	1.02%	1.52%	2.04%	2.55%
4	1.02%	1.52%	2.04%	2.55%	1.02%	1.52%	2.04%	2.55%
5	1.02%	1.52%	2.04%	2.55%	1.02%	1.52%	2.04%	2.55%
6	1.02%	1.52%	2.04%	2.55%	1.02%	1.52%	2.04%	2.55%
7	1.02%	1.52%	2.04%	2.55%	1.02%	1.52%	2.04%	2.55%
8	1.02%	1.52%	2.04%	2.55%	1.02%	1.52%	2.04%	2.55%
9	1.02%	1.52%	2.04%	2.55%	1.02%	1.52%	2.04%	2.55%
10	1.02%	1.52%	2.04%	2.55%	1.02%	1.52%	2.04%	2.55%

using this 'standard' spread calculation to assess the market's implied risk pricing is not possible.

4 Data and Estimation

The details of the estimation procedures are as follows. To fit the valuation model to market prices, we obtained traded coupon bond prices for the 850 trading days between September 2017 and January 2021 using the TRACE system. The price in the TRACE system does not represent the full amount paid for the bond. The full amount paid is the price plus accrued interest. We compare the full amount paid (the present value of the bond purchase) with the valuation model in expression (11). For each firm, we eliminated from the sample any subordinated bonds, callable and putable bonds, structured bonds, bonds with "death puts" or a "survivor option," and floating rate bonds. Survivor option bonds distort bond prices both because they are issued in small amounts (typically \$20 million or less per tranche) and because the value of the embedded put option is significant. The survivor option feature has become more common in recent years.¹⁷ Finally, to be included in our sample, the bond issue's daily trade volume had to exceed \$50,000 (in almost every case, volume was much larger) and with at least two bonds traded.¹⁸ We also exclude some bonds of European issuers subject to a 2014 EU regulation allowing regulators to demand an exchange of senior debt securities into equity. Because data assembly and cleaning costs are substantial.¹⁹

¹⁷The largest issuers of survivor option bonds as of 2016 included General Electric, Goldman Sachs, Bank of America, Wells Fargo, Ford Motor, HSBC Holdings, National Rural Utilities Cooperative Finance Corporation, Dow Chemical, Prospect Capital, and Barclays PLC. A typical survivor option bond's terms are described as follows in a recent prospectus supplement from General Electric Capital Corporation: "Specific notes may contain a provision permitting the optional repayment of those notes prior to stated maturity, if requested by the authorized representative of the beneficial owner of those notes, following the death of the beneficial owner of the notes, so long as the notes were owned by the beneficial owner or his or her estate at least six months prior to the request. This feature is referred to as a 'Survivor's Option.' Your notes will not be repaid in this manner unless the pricing supplement for your notes provides for the Survivor's Option. The right to exercise the Survivor's Option is subject to limits set by us on (1) the permitted dollar amount of total exercises by all holders of notes in any calendar year, and (2) the permitted dollar amount of an individual exercise by a holder of a note in any calendar year."

¹⁸To ensure model convergence, for issuer-days with only two observations we also require that the maturities are more than 270 days apart.

¹⁹It is necessary to screen out callables and survivor options, data on which is only available in the pricing supplement. The SEC and FINRA do not maintain public access to prospectus data for more than about 5 years in easily accessible form. Thus including, for example, data from the financial crisis is not feasible. In addition, the TABB group finds a very high frequency of errors "TABB Group analysis shows reconciliation differences in more than 20% of new issues." (http://www.finregalert.com/an-sec-mandated-corporate-bond-data-monopolywill-not-help-quality/). There are also non-trivial computational costs.

we restrict our attention to this sample. After the restrictions discussed above our sample consists of 120 thousand observations for 172 issuers.

We use the U.S. Treasury yields reported daily by the U.S. Department of the Treasury²⁰ and derive the maximum smoothness Treasury forward rate curves from these data (see Adams and van Deventer (1994)). Using these historical forward rate curves, we compute the term structure of default-free zero coupon bond prices on all dates.

Lastly, to facilitate the estimation of the intensity process, we assume that the default time τ corresponds to the first jump time of a Cox process with intensity $\lambda_t = \lambda_t(\Gamma_t) \geq 0$ where $\Gamma_t = (\Gamma_1(t), \ldots, \Gamma_m(t))' \in \mathbb{R}^m$ are a collection of stochastic processes characterizing the state of the firm and the market at time t. In addition, we assume that default risk is diversifiable in the sense of Jarrow, Lando, and Yu (2005).²¹ This assumption enables the estimation of the default intensities using historical time series data, without the need to adjust the intensity process for a default jump risk premium.

In conjunction, these two assumptions imply that we can estimate the default probabilities using a proportional hazard rate model (see Fleming and Harrington (1991), p. 126), i.e.

 $\lambda_t(\Gamma_t) = \theta e^{\phi \Gamma_t}$

where θ is a constant and where ϕ is a vector of constants. For an application of such a hazard rate model applied to corporate default probabilities see Chava and Jarrow (2004).

As discussed in Jarrow, Lando, and Yu (2005), this assumption does not imply that risky coupon bonds earn no risk premium. Quite the contrary. If the state variables Γ_t driving the default process represent systematic risk, which is the most likely case, then risky coupon bond prices necessarily earn a risk premium due to the bond price's correlation to Γ_t . The diversifiable risk assumption just states that the timing of the default event itself, after conditioning on Γ_t , is diversifiable in a large portfolio. Alternatively stated, in a poor economy all firms are more likely to default. But, the timing of which firms actually default depends on the idiosyncratic risks of the firm's management and operations.

The default process parameters (θ, ϕ) from the proportional hazard rate model were provided by the Kamakura Risk Information Services division of Kamakura Corporation.²² Kamakura uses a refinement of the approach employed by Chava and Jarrow (2004) to estimate these parameters that are then used to construct the full term structure of cumulative default probabilities.²³ Specifically, for each issuer-day we calculate cumulative default probabilities from the 10-year term

²⁰https://www.treasury.gov/resource-center/data-chart-center/interest-

rates/Pages/TextView.aspx?data=yield

²¹For additional detail, see the appendix.

 $^{^{22}\}mathrm{See}$ www.kamakuraco.com.

²³The model underlying the default probability calculations is similar to the one used in Camp-

structure of monthly marginal default probabilities (the monthly probability of default conditional on no prior default) that Kamakura generates from separate maturity-specific models. The state variables used in the Kamakura hazard rate estimation include both firm specific and macroeconomic variables. Importantly, the default probabilities do not use traded bond or CDS prices as inputs. Default probabilities are therefore separate inputs relative to the observed bond prices that we fit using our model.²⁴

Given the zero coupon bond prices and the default process (above), we estimate the recovery rate futures price d_t and illiquidity discount factor α_t for each issuer on each day. We compare the model values (expression (11)) to the market prices using a non-linear least squares estimation, calculated on a volume-weighted basis, to solve for the best fitting values of (d_t, α_t) . We restrict the recovery rate futures price to lie between zero and 0.8 and the illiquidity discount to lie between zero and -5%. Doing so will reduce the influence of observed bond price errors on the estimates. We check that our results are robust to relaxing these restrictions (and discuss the results together with other robustness checks in Section 6.3).

5 Illustrations: Coupon and Principal Seniority in Default

Before moving to the full sample estimation, this section provides evidence that market participants are aware of the difference in seniority between principal and coupons in default. We consider three companies that filed for bankruptcy: Lehman, PG&E, and Weatherford International. Lehman is chosen because of the size and importance of its bankruptcy. The latter two firms are in our sample because each firm has a sufficient number of bonds traded. In each case we focus on senior bonds, including callable bonds, because on the day bankruptcy is announced the call option is worthless and can be ignored. We fit the no-coupon recovery and full-coupon recovery models to the data. The key reason for analyzing issuer bonds after they file for bankruptcy is that the default probability equals 100%. The recovery amount for the no-coupon recovery model is the recovery rate times the notional of \$100 (par value) for each of the bonds. The recovery amount for the full-coupon recovery model is \$100 plus the dollar coupon times the number of remaining payments on each bond, a different amount for each issue. For each issuer day and for both of the models, we run the regression NPV

bell, Hilscher and Szilagyi (2008, 2011), who extend Chava and Jarrow (2004) and Shumway (2001). Campbell et al. show that the default probability measure is a more accurate predictor of failure than Moody's EDF numbers, data that have been widely used in academic studies, e.g. Berndt, Douglas, Duffie and Ferguson (2018).

 $^{^{24}}$ It is possible to amend our model so that implied default probabilities can be estimated based on observed bond prices. We discuss the robustness of our results to such a change in Section 6.3.



Figure 1: Lehman Brothers Pricing Errors

= (delta)(notional recovery amount) to derive the recovery rate and the present values (price plus accrued interest) for each bond.

Figures 1-3 depict the pricing errors. We order the bonds by maturity. Pricing errors when using the no-coupon recovery model (in blue) are substantially lower than those resulting from the full-coupon recovery model (in red). Mean absolute errors are more than five times as large for Lehman (2.0 vs. 11.0) and almost ten times as large for PG&E (2.1 vs. 19.7) and Weatherford International (2.6 vs. 21.0). Running regressions of actual versus predicted prices, for all three firms the R^2 's are larger than 99% for the no-coupon recovery model. When using the full-coupon recovery model, the R^2 's are between 84% and 92%.

The full-coupon recovery model results in prices that are too large, especially for bonds of longer maturities that have more coupons, which, if they were of equal seniority, would entitle the bond holder to a recovery value. However, in default those coupons are worthless and so any coupon paying bond would have pricing errors that are positive as long as the model was using unbiased inputs. However, in an attempt to fit the data, the model tries to reduce the average pricing error resulting in bonds with short maturities being underpriced and bonds with long maturities being overpriced. The maximum errors lie between 19.6 and 37.4. It is worth noting that average market prices are equal to 32.4 (Lehman), 78.2 (PG&E), and 65.0 (Weatherford International) so that the maximum errors are around one half the market price. The (negative) minimum errors are similar in size lying between -37.2 and -14.2. In contrast, the no-coupon recovery model's maximum



Figure 3: Weatherford International Pricing Errors



and minimum pricing errors are much smaller. They lie between 3.9 and 9.1 and -5 and -2.7, and so are approximately one quarter of the full-coupon recovery model pricing errors. Importantly, and in direct support of the no-coupon recovery model, its pricing errors have no clear pattern relative to the bond's maturity.

To summarize, Lehman, PG&E and Weatherford International's bond prices provide direct evidence in support of the no-coupon relative to the full-coupon recovery model. Failing to take into account the different seniority of coupons and principal results in substantial pricing errors, which have a predictable pattern consistent with our model.

6 Comparing the Two Pricing Models

This section provides a comparative analysis of the no-coupon and full-coupon recovery models.²⁵ The full-coupon recovery model produces different prices only if misspecification errors (see section 3.4) are non-zero. We therefore investigate *both* if the no-coupon recovery model has better fit on average, but also if it has better fit when misspecification errors are larger. Both of these predictions are implications of our model. If bonds are priced according to the no-coupon recovery model, then a necessary condition for model outperformance is the presence of misspecification errors.

Table 4 reports summary statistics of model parameters and model fit. The median default probability used as an input to the model is 0.8% and the mean is equal to 1.4%. We estimate the mean recovery rate futures price as 47%, and the mean illiquidity parameter as -0.4%. It is useful to note that both the recovery rate futures price and illiquidity estimates are reasonable. This is, of course, not guaranteed given that our estimates are based purely on traded bond prices and our no-coupon recovery model. Jankowitsch, Nagler, and Subrahmanyam (2014) report an average recovery rate value of 0.38. Since our recovery rate is the recovery rate futures price, as shown in the appendix, it is expected that our estimate should be slightly larger than these estimates. At the same time, our estimated 0.26% (median) and 0.44% (mean) illiquidity discount, though somewhat lower, are broadly consistent with, for example, the spread between Aaa-rated corporate bonds and Treasury debt. In section 6.2 we discuss variation in estimated illiquidity over time and note that it also spikes in March of 2020, similar to the Aaa-Treasury spread.

²⁵In a previous version of the paper we also compared the no-coupon recovery model to one based on ratings. In that model coupons are assumed to have full recovery and the credit spread is assumed to depend only on the rating. The ratings-based valuation model is consistent with numerous pronouncements from the Basel Committee on Banking Supervision (2010, 2017). It performs poorly primarily because of the erroneous assumption that all firms that have the same rating have the same risk; analyzing it is therefore less relevant when comparing no-coupon and full-coupon recovery models.

Table 4: Model Parameters and Fit

This table reports summary statistics of model parameters and model fit. Panel A reports bond level full sample statistics. Default probability is the average annualized default probability over the maturity of the bond, as discussed in the text; Recovery rate and liquidity are the parameters estimated from the main no-coupon recovery model. Estimation is done minimizing the volume-weighted squared pricing error. Recovery rate is restricted to lie between zero and 0.8; liquidity is restricted to lie between zero and -5%. Panel B reports model fit statistics at the issuer-day level, as well as misspecification error standard deviation, also at the issuer-day level. We report root volume-weighted mean squared errors and R-square (calculated relative to a model using the average bond price as the predictor, replaced with zero when negative). Panel B also reports model fit statistics for a subsample of issuer-day observations for the subsample with the top 25 percent of within issuer-day day misspecification.

Panel A: Bond level statistics							
	Def Breh	Recovery	Liquidity	Misspecification			
	DerProb	rate	Liquidity	error			
Mean	1.4%	0.47	-0.44%	0.23			
SD	2.1%	0.31	0.66%	0.62			
p5	0.1%	0.00	-1.7%	0.00			
p50	0.8%	0.52	-0.26%	0.024			
p95	4.9%	0.80	0.00%	1.17			
Number of observ	vations: 120,786						
		Panel B: I	Model fit				
	No coupon	Full coupon	No coupon rec	Full coupon	Miss offer CD		
	recovery RMSE	recovery RMSE	R-square	rec R-square	wiss error 50		
Mean	0.354	0.385	0.885	0.872	0.20		
SD	0.78	0.82	0.24	0.25	0.41		
p5	0.00	0.00	0.12	0.00	0.00		
p50	0.186	0.192	0.984	0.983	0.05		
p95	1.21	1.34	1.00	1.00	0.97		
Number of issuer	days: 25,018						
	Тор	quartile of misspe	cificatio error varia	ance			
Mean	0.68	0.80	0.80	0.75	0.68		
SD	0.60	0.71	0.28	0.32	0.60		
p5	0.00	0.00	0.00	0.00	0.23		
p50	0.56	0.66	0.92	0.90	0.47		
p95	1.78	2.07	1.00	1.00	1.72		
Number of issuer days: 6,254							

Compared to the implied recovery rates from the simple model in Section 2, recovery rates in the no-coupon recovery model are lower. This is partly a result of imposing an upper bound on recovery values of 0.8 but it is mainly due to the inclusion of an illiquidity discount that can capture frictions; empirically we see an important role for such a discount.

Misspecification Errors

Before comparing the no-coupon and full-coupon recovery model pricing errors, we analyze the misspecification errors. Using the unbiased estimates from the nocoupon recovery model as inputs, the misspecification error is computed as the difference in prices using the full-coupon recovery instead of the no-coupon recovery model. The full-coupon recovery model can only exhibit a worse fit if misspecification errors are present. If they are zero, the two models are the same. Recall that the misspecification error is approximately equal to $Cd_tp(t,t_1)Q(t,t_1)m(m+1)/2$. This approximation is quite accurate in capturing any variation in misspecification errors. When we regress actual on predicted misspecification error approximation formula is multiplicative in coupon rate, recovery rate, default probability and maturity. From the summary statistics of those four (reported in Tables 1 and 4) we already know that there will be a substantial fraction of the sample for which the misspecification errors are small.

Table 4 Panel A shows that the median misspecification error is 2.4 cents. Thus, as expected, half of our data set is not greatly affected by the pricing differences between the two models. However, the mean misspecification error is almost ten times as large and equal to 23 cents, and so there are many bonds for which there are large differences in the prices implied by the two models. The 95th percentile of the misspecification error distribution is 1.17, which is substantial, and the 75th percentile (unreported) is 0.17, which also quite large. The median observed bond price is equal to 102.32. As a result, these numbers are directly comparable to those in Table 1 of Section 3.4.²⁶

6.1 Model Performance

As noted previously, if bonds are priced according to the no-coupon recovery model and not the full-coupon recovery model, this has two implications. First, in a sample of bonds that have large default risk, we should detect that the no-coupon recovery model has a better fit. Second, the outperformance will be larger when

 $^{^{26} {\}rm Of}$ course, when estimating the models separately, which is what we do next, the full-coupon recovery model may adjust parameters. As we saw in Section 5 when discussing the bond prices of defaulted companies, an incorrect model not only produce biased parameter estimates, but its fit is also worse.

the two models disagree by more (i.e. when the misspecification errors are larger and more variable). We provide evidence supporting both of these implications below. We note that there is nothing mechanical about the relation between misspecification errors and model outperformance. If the data were priced according to the full-coupon recovery model, we would find that that model outperforms the no-coupon recovery model, and that it does so by more when the misspecification errors are larger.

We compare the two models after fitting them independently to the data. For each issuer-day, we estimate both the no-coupon and full-coupon recovery models and calculate the volume-weighted root mean squared errors and R^2 . These average error statistics are reported in Table 4 Panel B.

As expected, we find that the pricing errors are smaller for the no-coupon recovery model (35.4 cents) as compared to the full-coupon recovery model (38.5 cents). However, also as expected, there are many observations for which the error difference is not that large. We see this by comparing the median error, equal to 18.6 cents (no-coupon recovery) and 19.2 cents (full-coupon recovery). In fact, when calculating the error difference, the median is almost equal to zero (0.1 cents). This is a direct result of analyzing the full sample of transactions, one that includes many observations for which we know a priori that model differences are small (misspecification errors are close to zero) and therefore pricing differences will be near zero. The same pattern is present when analyzing the R^2 for the two models. The mean R^2 s are equal to 88.5% (no-coupon recovery) and 87.2% (full-coupon recovery).²⁷ The median R^2 are almost the same, equal to 98.4% and 98.3% respectively, and the median R^2 difference is close to zero.

What is more relevant for the performance comparison is how the models perform when their prices differ, resulting in misspecification errors. In fact, having a large misspecification error standard deviation within the data sample is crucial to validating the no-coupon recovery model. When predicted misspecification errors are small, for example because the default probabilities are low, there is only a small difference between the two models and so the pricing differences will also be small. But, even if predicted misspecification errors are large, the fullcoupon recovery model may still generate lower pricing errors, for example, if the bonds in the specific firm-day sample have very similar predicted misspecification errors. In this case the misspecified model will be able to adjust, for example by choosing a lower recovery rate, and the resulting pricing error differences relative to the no-coupon recovery model will be small, albeit at the cost of biased model parameter estimates. But, if there is a high variability of misspecification errors, the incorrect model will fail to fit prices as well as the no-coupon recovery model.

²⁷Both pricing models are non-linear and do not have a constant term. R^2 is calculated relative to a model with only a constant and it is therefore possible to end up with negative R^2 , if the non-linear model is outperformed by the constant model. In those cases we set R^2 equal to zero.

Reflecting the large portion of the data with low misspecification errors, we see that a large fraction of the data also has a low standard deviation of those errors. The mean standard deviation is 20 cents (Table 4 Panel B) but the median is only 5 cents.

As predicted, we find that the no-coupon recovery model outperformance is large when the misspecification error standard deviation is large. The second set of results reported in Panel B is for the subsample of issuer days in the top quartile of the misspecification error standard deviation distribution. In that group, the minimum misspecification error standard deviation is 21 cents and the mean is 68 cents. Correspondingly, the differences in root mean squared error and R^2 s are much larger, exactly as our model predicts. The mean RMSE is 68 cents (no-coupon recovery) and 80 cents (full-coupon recovery) and the medians are 56 cents (no-coupon recovery) and 66 cents (full coupon recovery). R^2 follows the same pattern – the averages are equal to 80% (no-coupon recovery) and 75% (fullcoupon recovery). Even though we are considering only 25% of issuer days, this sample reflects prices from 36,528 observations, which is 30.2% of observations, or 5.8 observations per issuer-day compared to 4.8 for the full sample.

The Determinants of No-coupon Recovery Model Outperformance

We next explore the determinants of the no-coupon recovery model's outperformance. In Table 5 we first measure outperformance for the full sample and several subsamples, and then regress outperformance on different sets of explanatory variables. On average, the no-coupon recovery model provides a better fit (also see Table 4 Panel B). Pricing errors are 3.2 cents larger when using the full-coupon recovery model and the difference is statistically significant. We have already seen that the (in)ability of the full-coupon recovery model to fit the data reflects its misspecified assumption. Thus, we expect a strong relationship between the nocoupon recovery model outperformance and the misspecification error's standard deviation.²⁸

We first focus on the subsample with the top 25% of default probabilities, an independent input to the model. For that subsample, the average root mean squared errors are 10.6 cents larger for the full coupon recovery model; the difference is three times the full sample difference. This outperformance is even larger when considering the subsample with the highest 25% average misspecification error issuer days. Here the outperformance is 11.6 cents on average. As expected, the best variable to find large outperformance is the misspecification error's standard

 $^{^{28}}$ Recall that when estimating this model, the algorithm searches for those parameters to maximize its fit. The examples of the issuers in bankruptcy nicely show that considering only the average pricing error of the full-coupon recovery model misses the strong association of the pricing errors with the bond's maturity, a pattern that is a direct consequence of the model's misspecification.

deviation. It may be large because of a large dispersion in maturities combined with large default probabilities. The misspecified model does not have sufficient degrees of freedom to match the data well and the no-coupon recovery model has, on average, a 12.1 cent lower pricing error than the full-coupon recovery model. The pattern is the same, only stronger, when examining the pricing errors in the top deciles. For large default probability issuer days the outperformance is equal to 16 cents, for the misspecification error it is 23.6 cents and for the standard deviation it is equal to 24.3 cents. Pricing error differences are large when the model predicts them to be large.

We next explore in more detail what determines the size of the no-coupon recovery model outperformance. We regress pricing error differences on the four variables that are the main determinants of the misspecification error (recovery rate, maturity, coupon, and default probability). Recovery rate, maturity, and default probability are all highly significant and come in with a positive sign; the coupon rate also has a positive sign but is insignificant. All four variables enter the misspecification error, but they do so in a specific way. When we include the actual misspecification error in the regression, the coefficients of all four variables become indistinguishable from zero. Meanwhile, the model fit increases from 26.2% to 47.6%. Dropping the four variables results in almost the same fit of 47.3%.

Of course, these variables do not measure the dispersion of the misspecification errors, though the misspecification error level and standard deviation are correlated to each other and to the four variables. When we include the misspecification error standard deviation together with the four variables, dispersion is the most significant variable and the regression fit increases to 57.3%. Removing the four variables results in almost the same fit of 56.6%. The coefficients on the four inputs are now negative, most likely capturing non-linearities in the relationship rather than their independent importance.

We also find that the misspecification error standard deviation is what is important for model outperformance. When both the level and the standard deviation are included together, the level becomes insignificant and the model fit is practically the same as using only the standard deviation. The ability of the misspecification error standard deviation to explain variation in model outperformance is direct evidence supporting our hypothesis that the market prices bonds according to the no-coupon recovery model instead of the full-coupon recovery model.

To summarize, our analysis provides evidence that the pricing error differences between the two models are statistically significant for the full sample. Importantly, variation in model outperformance occurs exactly at those times when the model predicts it. This evidence is for a sample consisting of 94.6% investment grade debt (98.9% with a rating of BB+ or above) and thus one where market participants perceive default is not imminent. As we discussed earlier, when default

Table 5: Determinants of Model Performance

This table reports results from regressions of root mean squared errors for different samples (Panel A) and on different sets of explanatory variables (Panel B). Root mean squared error is calcualted at the issuer-day level and is volume weighted (same as the actual estimation). "PD" is default probability, "Miss error" is the average within issuer-day misspecification error and "Miss error SD" is its standard deviation. Recovery rate is from the no-coupon recovery model; coupon, maturity, and default probability are averaged at the issuer-day level. Standard errors, reported below coefficients, are clustered by issuer and date; *** denotes significant at the 1% level, ** at the 5% level and * at the 10% level.

	Panel A: Root mean squared error differences: subsamples						
Constant	0.032***	0.106***	0.116***	0.121***	0.160***	0.236***	0.243***
	(0.010)	(0.028)	(0.028)	(0.027)	(0.047)	(0.035)	(0.035)
	full	PD	Miss error	Miss err SD	PD	Miss error	Miss err SD
Sample	Tun	top quartile	top quartile	top quartile	top decile	top decile	top decile
Issuer days	25,018	6,254	6,254	6,254	2,501	2,501	2,501
	Panel P	Determinant	s of variation in	n root mean so	wared error di	fferences	
Recovery	0.035***	-0.004		-0.022**			
	(0.013)	(0.009)		(0.008)			
Maturity	0.019***	-0.000		-0.004*			
	(0.006)	(0.002)		(0.002)			
Coupon	0.002	-0.002		-0.001			
	(0.006)	(0.002)		(0.002)			
Def Prob	2.466***	0.463		-0.457			
	(0.823)	(1.009)		(0.879)			
Miss error		0.144***	0.151***			0.022	
		(0.019)	(0.018)			(0.018)	
Miss err SD				0.232***	0.207***	0.182***	
				(0.021)	(0.021)	(0.038)	
Constant	-0.080***	0.005	-0.001	0.014**	-0.011***	-0.010***	
	(0.026)	(0.008)	(0.003)	(0.006)	(0.002)	(0.002)	
P. coupro	0.262	0.476	0.472	0 572	0.566	0.569	
n-square	0.262	0.476	0.473	0.573	0.566	0.568	
issuer days	25,018	25,018	25,018	25,018	25,018	25,018	

risk is larger, differences across the models are even larger.

Errors and Biased Parameter Estimates in the Full-Coupon Recovery model: A Quasi-Simulation

One way to check that our results are due to the misspecified assumptions underlying the full-coupon recovery model is to use a simulation to calculate pricing errors. We do this using our actual data as the basis for a quasi-simulation. Specifically, we assume that the no-coupon recovery model's parameters are unbiased and calculate market prices based on that model. We then use the full-coupon recovery model to price these bonds. This enables us to trace out the effects of using these misspecified assumptions on the prices.

We find that the expected biases arise. The misspecification error's standard deviation is closely related to the pricing error and there is a downward bias on the recovery rate. When regressing issuer-date level root mean squared errors on the misspecification error's standard deviation, the R^2 is equal to 75% for the full sample. Thus, whenever the full-coupon recovery model has more difficulty using the parameters to fit the data, the actual mean-squared errors are larger. We also find that the full-coupon recovery parameter estimates are indeed biased relative to the no-coupon recovery model's parameters, which in this exercise are assumed to be the actual underlying parameters. The effect is more pronounced for recovery rates. The full-coupon recovery model's estimated recovery rate is 3% lower for the full sample and 6.7% lower for the subsample with the highest quartile of the misspecification error's standard deviation. For the liquidity parameter the effect is much smaller. For the full sample, the bias is equal to 0.02% (full-coupon-implied liquidity is more negative), while it is 0.13% for the top 10% of the misspecification error's standard deviation subsample. In short, the patterns we find in the data are directly driven by the full-coupon model's assumptions.

6.2 Time Variation in Model Performance and Parameter Estimates

We have documented the no-coupon recovery model's outperformance in the full sample. We now turn to studying the time variation in its outperformance. Each week we average all the issuer days with above-median misspecification error variance and then average the difference in the full-coupon and no-coupon recovery model RMSE. Figure 4 plots the time series of the no-coupon recovery model's outperformance. There is some noticeable time variation. Toward the end of 2018 and in the beginning of 2019 the model's outperformance increases, reaching a peak of about 14 cents. This episode happened contemporaneously with a stock market downturn and a corresponding increase in volatility and default probabilities. Then, at the beginning of the pandemic in early 2020 we also see an increase



Figure 4: Error Difference Over Time

in model outperformance reaching a weekly average of 26 cents. Note that, with the exception of two averages at the end of 2020, all of the remaining 175 weeks exhibit the average outperformance of the no-coupon recovery model.

As reported in Table 5, the misspecification error standard deviation explains 57% of the variation in no-coupon recovery model outperformance as measured by the RMSE difference. Figure 5 shows this pattern graphically. We plot the same weekly averages of the above-median misspecification error standard deviation observations as in Figure 4. The figure shows a clear linear pattern. Easily identified are the pandemic observations as those with higher model outperformance. This provides evidence in support of our pricing model, because these observations are the ones where the misspecification error's standard deviation is large. However, it is not immediately apparent how close the relationship is between misspecification error and model outperformance outside of 2020. To check this, we ran a regression of all issuer day RMSE differences, and the R^2 actually increases to 65%.

Figure 6 graphs the weekly averages, confirming the clear relationship present in the full-sample results. The relationship continues to be linear. The only



Figure 5: Error Difference RMSE

difference is that all of the weekly observations of average misspecification error standard deviations between 0.6 and 1.3 are no longer present.

We next turn to the illiquidity and recovery rate futures price (recovery rate, for short) parameters over time. Figure 7 plots the weekly average of the illiquidity parameters. We again notice a slight increase in late 2018, early 2019. The much more striking and larger increase in the illiquidity parameters occur at the beginning of the pandemic. The average illiquidity parameter increases to just above 3%. This increase is matched by the increase in the Aaa-Treasury spread, which we also plot in the figure. We interpret the Aaa-Treasury spread as a related measure of illiquidity. Indeed, as is evident from the figure, the two series move together; the correlation is equal to 77%. The close relationship between the Aaa-Treasury spread provides independent validation of our approach modeling corporate bond illiquidity.

Finally, we plot the average recovery rate over time in Figure 8. As pointed out in Section 2, the fact that there is no recovery of coupons in the event of default means that recovery rates can be directly inferred from bond prices. This is an important empirical novelty compared to an inability to disentangle recovery rates



Figure 6: Error Differences not 2020



Figure 7: Liquidity Parameter versus Aaa-Treasury spread

and default probabilities if equal seniority is assumed and there is a single spread used to discount all cash flows. The most noticeable pattern is an increase in recovery rates in late 2018 and early 2020. Bond prices can decline either because of lower recovery rates or because of more illiquidity. Our estimates suggest that, at least during part of the sample period, the two effects moved in opposite directions.

6.3 Robustness

This section provides several robustness tests of the model's performance.

Daily Observation Cutoffs

In our estimation, we fit the recovery and illiquidity parameters for each issuer day, and we require a minimum of two bond price observations for each issuer day. Another possibility is to choose minimum observation cutoffs, for example requiring at least five or ten observations for each issuer day. Such a restriction has the benefit of reducing estimation noise but comes at the cost of shrinking the sample size. We have checked that our results are robust to increasing the



Figure 8: Recovery Rate Futures Price Parameters

minimum required number of observations.

Monthly Observations

Another way in which noise in parameter estimates can be reduced is to estimate the recovery and illiquidity parameters for all the issuer observations in one month rather than for one day. This approach significantly increases the number of observations used to fit the recovery and illiquidity parameters. It also has the benefit of not reducing the overall sample size. Our results are robust to this change too. In particular, we see a similar level of outperformance of the no-coupon recovery model, though both models fit less well since there is now less parameter flexibility. Our findings suggests that outperformance is not directly linked to sample size. This result is consistent with our findings that the misspecification error's variance explains well the variation in model outperformance.

Implied Default Probabilities

In our main results we use independently estimated default probabilities from a proportional hazard model and then implicitly estimate the illiquidity and recovery rate parameters. We now check that our results do not depend on these probabilities. We can also fit our model (11) assuming a flat term structure of default probabilities and estimate the default intensity implicitly. In the model, a change in the default probability affects the present value of all cash flows, both coupons and principal. The same is true for a change in the illiquidity parameter. From expression (11) and the spread curve examples in Table 3 we know that the effect is not exactly the same however, and so it is possible to estimate both separately. Nevertheless, to guard against unstable estimates, we fit two different models. One where, in addition to the illiquidity and the recovery parameter, we estimate a default intensity from the bond data. The second, a model where we set the illiquidity discount equal to zero and estimate only a recovery rate and a default intensity. As before, we do this both at the issuer-day and the issuer-month levels. Our main result of outperformance of the no-coupon recovery model is robust to this approach.

Alternative Bounds for the Illiquidity and Recovery Parameters

In Section 2 we saw that spread-implied recovery values can be quite large, and larger than the upper bound of 80% on the recovery rates imposed when estimating our model. We relax this restriction and instead fit our model assuming that recovery rates can reach 100%. We do this both for the issuer-day estimation approach which forms the basis for our main results as well as for the monthly observation methodology. We also allow the illiquidity parameter to reach 10%.

Our results are robust to this change. There continues to be substantial outperformance of the no-coupon recovery model relative to the full-coupon recovery model. We actually see an improvement in the overall fit of the no-coupon recovery model's performance.

7 Conclusion

This paper makes three contributions to the literature on corporate bond pricing. First, we provide evidence that the common assumption of equal seniority of principal and coupon payments is not supported by market transaction prices. Coupon spreads are much higher than principal spreads, and assuming that a single spread can be used to discount all of the bond's cash flows results in pricing errors that are more than twice as large compared to using different, seniority-specific spreads.

Second, we propose a new and tractable coupon bond valuation model, which includes a more realistic recovery rate process that distinguishes between coupon payments received before and after default. Our approach has important benefits: (1) we show how to price bonds using building block securities – survival digitals, which pay off in the event of no default, and default digitals, which pay off in the event of default. In this way we show that seniority-specific discount rates arise when taking into account different recovery rates of principal and coupons. Using these building block securities for pricing makes the model intuitive and straightforward to implement. (2) The model has a clear prediction about the importance of modeling the market practice of zero recovery on coupons paid after default. We calculate misspecification errors – those resulting from using the fullcoupon recovery model rather than the no-coupon recovery model. When these errors are large, the no-coupon and the full-coupon recovery model's predictions differ by more. Misspecification errors depend directly on the coupon, recovery rate, default probability, and time to maturity, and they can be substantial in size.

Third, using a large sample of bond transaction prices we demonstrate that our no-coupon recovery model's predictions are reflected in the data and that, therefore, market participants are pricing default risky coupon bonds taking zero coupon recovery after default into account. The model has a clear prediction when no recovery on coupons after default is quantitatively relevant for pricing and when it is less important – if default probabilities, coupons, or maturity are small, the effect of differing recovery assumptions for coupons has only a very small impact. There is clear evidence supporting this prediction in the data. We can detect our model's pricing effects in the full sample. In addition, and more importantly, we find that model outperformance is closely related to the misspecification error's standard deviation. When that standard deviation is large, our model predicts that bond prices will be the most affected by the erroneous assumption of full-coupon recovery after default. The fact that the misspecification error's standard deviation explains model outperformance so well is thus direct evidence supporting the no-coupon recovery model. Separately, we find that the no-coupon recovery outperformance is evident when considering bond prices of companies in bankruptcy.

Finally, our model allows for direct estimation of implied recovery rates and illiquidity effects on bond prices. When the global pandemic hit in 2020, default probabilities increased substantially and bond prices dropped. The no-coupon recovery model outperformance increased, as predicted by our model. At the same time, illiquidity increased markedly, consistent with standard corporate bond illiquidity measures such as the Aaa-Treasury spread.

References

- K. Adams and D. van Deventer, 1994, "Fitting Yield Curves and Forward Rate Curves with Maximum Smoothness," Journal of Fixed Income, June, 52 - 62.
- [2] G. Bakshi, D. Madan, and F. Zhang, 2006, "Understanding the Role of Recovery in Default Risk Models: Empirical Comparisons and Implied Recovery Rates," September, FDIC CFR Working Paper No. 06.
- [3] Basel Committee on Banking Supervision, 2010, Basel III: A global regulatory framework for more resilient banks and banking systems, December, Bank for International Settlements.
- [4] Basel Committee on Banking Supervision, 2017, Basel III: Finalizing the post-crisis reforms, December, Bank for International Settlements.
- [5] A. Berndt, R. Douglas, D. Duffie, and M. Ferguson, 2018, "Corporate Credit Risk Premia," Review of Finance, 22(2), 419–454.
- [6] T. Bielecki and M. Rutkowski, 2002, Credit Risk: Modeling, Valuation and Hedging, Springer Verlag, Berlin.
- [7] P. Bremaud, 1980, Point Processes and Queues: Martingale Dynamics, Springer Verlag, Berlin.
- [8] J. Y. Campbell, J. Hilscher, and J Szilagyi, 2008, "In Search of Distress Risk," Journal of Finance, 63(6), 2899–2939.
- J. Y. Campbell, J. Hilscher, and J. Szilagyi, 2011, "Predicting Financial Distress and the Performance of Distressed Stocks," Journal of Investment Management, 9(2), 14–34.

- [10] J. Y. Campbell, and G. B. Taksler, 2003, "Equity Volatility and Corporate Bond Yields," Journal of Finance, 58, 2321-2350.
- [11] U. Cetin, R. Jarrow and P. Protter, 2004, "illiquidity Risk and Arbitrage Pricing Theory," Finance and Stochastics, 8 (3), (August), 311 - 341.
- [12] S. Chava and R. Jarrow, 2004, "Bankruptcy Prediction with Industry Effects," Review of Finance, 8 (4), 537 - 569.
- [13] J. Cherian, E. Jacquier, and R. Jarrow, 2004, "A Model of the Convenience Yields in On the Run Treasuries," Review of Derivatives Research 7, 79 - 97.
- [14] P. Collin-Dufresne and R. Goldstein, 2001, "Do Credit Spreads Reflect Stationary Leverage Ratios," Journal of Finance, 56 (6), 1929 - 1957.
- [15] P. Collin-Dufresne, R. Goldstein, and J. Martin, 2001, "The Determinants of Credit Spread Changes," Journal of Finance, 56 (5), 2177 - 2207.
- [16] J. Driessen, 2005, "Is Default Event Risk Priced in Corporate Bonds?," The Review of Financial Studies, 18 (1), 165 - 195.
- [17] G. Duffee, 1999, "Estimating the Price of Default Risk," The Review of Financial Studies," 12 (1), 197 - 226.
- [18] D. Duffie, L. Pedersen, and K. Singleton, 2003, "Modeling Sovereign Yield Spreads: A Case Study of Russian Debt," Journal of Finance, 63 (1), 119 -159.
- [19] D. Duffie and K. Singleton, 1999, "Modeling Term Structures of Defaultable Bonds," Review of Financial Studies, 12 (4), 687 - 720.
- [20] E. Elton, M. Gruber, D. Agrawal, and C. Mann, 2001, "Explaining the Rate Spread on Corporate Bonds," Journal of Finance, 56 (1), 247 - 277.
- [21] T. Fleming and D. Harrington, 1991, Counting Processes and Survival Analysis, John Wiley & Sons, Inc., New York.
- [22] S. Gilchrist and E. Zakrajšek, 2012, "Credit Spreads and Business Cycle Fluctuations," American Economic Review, 102 (4), 1692-1720.
- [23] R. Guha and A. Sbuelz, 2005, "Structural Recovery of Face Value at Default," unpublished working paper, University of Verona, Italy.
- [24] X. Guo, R. Jarrow and H. Lin, 2008, "Distressed Debt Prices and Recovery Rate Estimation," Review of Derivatives Research, 11, 171 - 204.

- [25] D. Heath, R. Jarrow and A. Morton, 1992, "Bond Pricing and the Term Structure of Interest Rates: A New Methodology for Contingent Claims Valuation," Econometrica, 60 (1), 77 - 105.
- [26] J. Hilscher, Y. Nosbusch, 2010, "Determinants of Sovereign Risk: Macroeconomic Fundamentals and the Pricing of Sovereign Debt," Review of Finance, 14 (2), April, 235–262.
- [27] J. Huang and M. Huang, 2012, "How Much of the Corporate-Treasury Yield Spread Is Due to Credit Risk?," Review of Asset Pricing Studies, 2 (2), December, 153–202.
- [28] J. Jacod and P. Protter, 2010, "Risk Neutral Compatibility with Option Prices," Finance and Stochastics, 14, 285 - 315.
- [29] R. Jankowitsch, F. Nagler, and M. G. Subrahmanyam, 2014, "The determinants of recovery rates in the US corporate bond market," *Journal of Financial Economics*, 114 (1), 155-177.
- [30] R. Jarrow, 2004, "Risky Coupon Bonds as a Portfolio of Zero-coupon Bonds," Finance Research Letters, 1, 100 - 1005.
- [31] R. Jarrow, 2009, "Credit Risk Models," Annual Review of Financial Economics, 1, 37 - 68.
- [32] R. Jarrow, 2021, Continuous-Time Asset Pricing Theory: A Martingale-Based Approach, 2nd Edition, Springer.
- [33] R. Jarrow, D. Lando and F. Yu, 2005, "Default Risk and Diversification: Theory and Empirical Applications," Mathematical Finance, 15 (1), (January), 1 - 26.
- [34] R. Jarrow and M. Larsson, 2012, "The Meaning of Market Efficiency," Mathematical Finance, 22 (1), 1-30.
- [35] R. Jarrow and P. Protter, 2008, "An Introduction to Financial Asset Pricing," Handbooks in OR&MS, Chap. 1 (J. R. Birge & V. Linetsky eds.), Vol. 15. Philadelphia: Elsevier B.V.
- [36] R. Jarrow and S. Turnbull, 1995, "Pricing Derivatives on Financial Securities Subject to Credit Risk," Journal of Finance, 50 (1), 53 - 85.
- [37] R. Jarrow and S. Turnbull, 1997, "An Integrated Approach to the Hedging and Pricing of Eurodollar Derivatives," Journal of Risk and Insurance, 64 (2), 271 - 299.

- [38] R. Jarrow and S. Turnbull, 2000, "The Intersection of Market and Credit Risk," Journal of Banking and Finance, 224, 271 - 299.
- [39] M. Kargar, B. Lester, D. Lindsay, S. Liu, P. O. Weill, and D. Zúñiga, 2021, "Corporate Bond illiquidity during the COVID-19 Crisis," The Review of Financial Studies, 34 (11), November, 5352–5401.
- [40] D. Lando, 1998, "On Cox Processes and Credit Risky Securities," Review of Derivatives Research, 2, 99 - 120.
- [41] F. Longstaff, S. Mithal, and E. Neis, 2005, "Corporate Yield Spreads: Default Risk or illiquidity? New Evidence from the Credit Default Swap Market," Journal of Finance, 60 (5), 2213 - 2253.
- [42] D. Madan and H. Unal, 1998, "Pricing the Risks of Default," Review of Derivatives Research, 2, 121 - 160.
- [43] P. Protter, 2005, Stochastic Integration and Differential Equations, 2nd edition, Springer-Verlag: New York.
- [44] T. Shumway, 2001, "Forecasting Bankruptcy More Accurately: A Simple Hazard Model," Journal of Business, 74(1), 101–124.

Appendix.

The Bond Valuation Formula

This appendix formalizes the pricing model in Section 3. We consider a continuous trading model on a finite horizon $[0, T^*]$. The uncertainty in the model is characterized by a complete filtered probability space $(\Omega, \mathcal{F}, (\mathcal{F}_t)_{t \in [0, T^*]}, \mathbb{P})$ where the filtration $(\mathcal{F}_t)_{t \in [0, T^*]}$ satisfies the usual hypotheses and $\mathcal{F} = \mathcal{F}_{T^*}$.²⁹ Here \mathbb{P} is the statistical probability measure.

The default-free money market account earns interest continuously at the default-free spot rate of interest, r_t , which is adapted to $(\mathcal{F}_t)_{t \in [0,T^*]}$. As in the text, money market account's time t value is³⁰

$$B_t = e^{\int_0^t r_s ds}.\tag{12}$$

The default-free zero-coupon bond, denoted by p(t,T) > 0, is adapted to $(\mathcal{F}_t)_{t \in [0,T^*]}$.

Let $\Gamma_t = (\Gamma_1(t), \ldots, \Gamma_n(t))' \in \mathbb{R}^n$ be a collection of stochastic processes characterizing the state of the firm and the market at time t with \mathcal{F}_t^{Γ} representing the filtration generated by the state variables Γ_t up to and including time $t \ge 0$. We assume that these state variables are adapted to \mathcal{F}_t , which implies that $\mathcal{F}_t^{\Gamma} \subset \mathcal{F}_t$. We assume that r_t is \mathcal{F}_t^{Γ} - measurable.

We assume that these state variables are adapted to t_t , then TWe assume that r_t is \mathcal{F}_t^{Γ} - measurable. Let $\lambda : [0, T^*] \times \mathbb{R}^n \longrightarrow [0, \infty)$, denoted $\lambda_t = \lambda_t(\Gamma_t) \ge 0$, be jointly Borel measurable with $\int_0^{T^*} \lambda_t(\Gamma_t) dt < \infty$ a.s. \mathbb{P} , and let $N_t \in \{0, 1, 2, \cdots\}$ with $N_0 = 0$ be a Cox process conditioned on \mathcal{F}_T^{Γ} with $\lambda_t(\Gamma_t)$ its intensity process (see Lando (1998)). Finally, let the default time $\tau \in [0, T^*]$ be the stopping time adapted to the filtration \mathcal{F}_t defined by

$$\tau \equiv inf \{t > 0 : N_t = 1\}.$$

The function $\lambda_t(\Gamma_t)$ is the firm's default intensity. A Cox process is a point process which, conditional upon the information set generated by the state variables process Γ_t over the entire trading horizon $[0, T^*]$ behaves like a standard Poisson process. In particular,

$$\mathbb{P}(\tau > \mathcal{T} \left| \mathcal{F}_{T^*}^{\Gamma} \lor \mathcal{F}_t \right) = e^{-\int_t^{\mathcal{T}} \lambda_u du}$$

and

$$\mathbb{P}(\tau > \mathcal{T} | \mathcal{F}_t) = E^{\mathbb{P}} \left[e^{-\int_t^{\mathcal{T}} \lambda_u du} | \mathcal{F}_t \right].$$

Letting the time $t \leq t_1$ value of the coupon bond be v_t , we add the following assumption.

²⁹See Protter (2005) for the definitions of these various terms.

 $^{^{30}{\}rm Of}$ course, we assume the necessary measurability and integrability such that the following expression is well-defined.

Assumption. (Existence of an Equivalent Martingale Measure) There exists an equivalent probability measure \mathbb{Q} such that

$$\frac{p(t,\mathcal{T})}{B_t} \text{ for all } \mathcal{T} \in [0,T^*] \quad and \quad \frac{v_t}{B_t}$$

are \mathbb{Q} martingales.

Equivalent means that both the probability measures \mathbb{P} and \mathbb{Q} agree on zero probability events. It is well known that this assumption implies that the market is arbitrage-free. See Jarrow and Protter (2008). Given an equivalent martingale measure \mathbb{Q} , define $\lambda_t \equiv \lambda_t \kappa_t$ to be the intensity process of the Cox process under \mathbb{Q} where $\kappa_t(\omega) \geq 0$ is a predictable process with $\int_0^{T^*} \lambda_t(\Gamma_t) \kappa_t dt < \infty$ a.s. \mathbb{P} (see Bremaud (1980), p. 167). The process $\kappa_t(\omega)$ represents a default jump risk premium.

Let \mathcal{F}_t^r denote the filtration generated by r_t , \mathcal{F}_τ the σ -algebra generated by the default time τ (see Protter (2005), p. 5), and \mathcal{F}_t^{δ} the filtration generated by the recovery rate process δ_t , which is the hypothetical recovery rate obtained if the assets were liquidated at time t. Finally, let \mathcal{G}_t be a filtration independent of $\mathcal{F}_t^r \vee \mathcal{F}_t^\delta \vee \mathcal{F}_\tau$ where $\mathcal{F}_t \equiv \mathcal{F}_t^r \vee \mathcal{F}_t^\delta \vee \mathcal{F}_\tau \vee \mathcal{G}_t$. We add the following assumption.

Assumption. (Conditional Independence)

The default-free spot rate r_t , the default time τ , and the recovery rate process δ_t are independent under \mathbb{Q} given the filtration \mathcal{G}_t for all t.

Proof. (The Bond Valuation Formula)

The only difficult term is $E^{\mathbb{Q}}\left[1_{\{t_{k-1} < \tau \leq t_k \leq T\}} \delta_{t_k} e^{-\int_t^{t_k} r_u du} |\mathcal{F}_t\right]$. We have by conditional independence that

$$E^{\mathbb{Q}}\left[1_{\{t_{k-1}<\tau\leq t_k\leq T\}}\delta_{t_k}e^{-\int_t^{t_k}r_udu} |\mathcal{F}_t\right]$$
$$= E^{\mathbb{Q}}\left[1_{\{t_{k-1}<\tau\leq t_k\leq T\}} |\mathcal{F}_t\right]E^{\mathbb{Q}}\left[\delta_{t_k} |\mathcal{F}_t\right]E^{\mathbb{Q}}\left[e^{-\int_t^{t_k}r_udu} |\mathcal{F}_t\right].$$

Using the definition of d_t and $p(t, t_k)$ gives

$$v_t = \sum_{k=1}^m Cp(t, t_k) E^{\mathbb{Q}} \left[\mathbf{1}_{\{\tau > t_k\}} | \mathcal{F}_t \right] + Lp(t, T) E^{\mathbb{Q}} \left[\mathbf{1}_{\{\tau > T\}} | \mathcal{F}_t \right] \\ + \sum_{k=1}^m Lp(t, t_k) d_t E^{\mathbb{Q}} \left[\mathbf{1}_{\{t_{k-1} < \tau \le t_k \le T\}} | \mathcal{F}_t \right].$$

The Recovery Rate Futures Price

This section explains why d_t can be interpreted as a futures price. Define

$$d_t \coloneqq E^{\mathbb{Q}}\left[\delta_{T^*} \mid \mathcal{F}_t\right]$$

for $t \in [0, T^*]$ where δ_{T^*} corresponds to the recovery rate if default happens at time T^* . That this can be interpreted as a futures price follows from the commodities pricing literature (for example see Jarrow (2021), chapter 6). Since d_t is a \mathbb{Q} -martingale, the recovery rate futures price is equal to

$$d_t = E^{\mathbb{Q}} \left[\delta_\tau \left| \mathcal{F}_t \right] \right],$$

which is the risk adjusted expected value of the recovery rate at the default time τ . The relation between the recovery rate process δ_t and the futures price d_t is given by the following expression

$$\delta_t = d_t E^{\mathbb{Q}} \left[\frac{B_t}{B_\tau} \, | \mathcal{F}_t \right]. \tag{13}$$

Proof (Expression (13)).

Define the recovery rate process δ_t on $[0, T^*]$ as the liquidation value of the firm at time t. Using the notation in the text, the following time line documents relevant dates.

Both τ and τ^* are stopping times with $\tau^* > \tau$. τ is the firm's default time. Due to cross-defaulting provisions, if one liability defaults, all default at the same time. τ^* corresponds to the date the payments on all of the firm's liabilities are paid, after liquidation or financial restructuring. It occurs, as indicated, after the default date. Time T^* is when the model ends. The arbitrage free value of the recovery rate payment at time $t \in [0, T^*]$ is

$$\delta_t = E^{\mathbb{Q}} \left[\frac{\delta_{T^*}}{B_{T^*}} \, |\mathcal{F}_t \right] B_t.$$

Using iterated expectations, we obtain

$$\delta_t = E^{\mathbb{Q}} \left[\frac{\delta_\tau}{B_\tau} \, |\mathcal{F}_t \right] B_t. \tag{14}$$

Under the conditional independence of δ_t and r_t , we can rewrite this as

$$\delta_{t} = E^{\mathbb{Q}} \left[E^{\mathbb{Q}} \left[\delta_{\tau} \left| \mathcal{F}_{t} \lor \tau \right] E^{\mathbb{Q}} \left[\frac{B_{t}}{B_{\tau}} \left| \mathcal{F}_{t} \lor \tau \right] \right| \mathcal{F}_{t} \right]$$
$$= \sum_{i=1}^{M} E^{\mathbb{Q}} \left[\delta_{t_{i}} \left| \mathcal{F}_{t}; t_{i-1} < \tau \le t_{i} \right] E^{\mathbb{Q}} \left[\frac{B_{t_{i}}}{B_{\tau}} \left| \mathcal{F}_{t}; t_{i-1} < \tau \le t_{i} \right] Prob^{\mathbb{Q}} \left[t_{i-1} < \tau \le t_{i} \left| \mathcal{F}_{t} \right] \right]$$

$$= \sum_{i=1}^{M} E^{\mathbb{Q}} \left[\delta_{t_{i}} \left| \mathcal{F}_{t}; t_{i-1} < \tau \leq t_{i} \right] E^{\mathbb{Q}} \left[\frac{B_{t_{i}}}{B_{\tau}} \left| \mathcal{F}_{t}; t_{i-1} < \tau \leq t_{i} \right] \left[Q(t, t_{i}) - Q(t, t_{i-1}) \right] \right]$$

where $t_M = T^*$ and $Q(t, t_i) = E^{\mathbb{Q}} \left[\mathbb{1}_{\tau \leq t_i} | \mathcal{F}_t \right] = Prob^{\mathbb{Q}} \left[\tau \leq t_i | \mathcal{F}_t \right].$

Recall that all payments occur at the next payment date after default.

Using the conditional independence of δ_t and r_t from τ , we can rewrite this as

$$\delta_t = \sum_{i=1}^M E^{\mathbb{Q}} \left[\delta_{t_i} \left| \mathcal{F}_t \right] E^{\mathbb{Q}} \left[\frac{B_t}{B_{t_i}} \left| \mathcal{F}_t \right] \left[Q(t, t_i) - Q(t, t_{i-1}) \right] \right].$$

Last, using the definition of d_t and noting that $E^{\mathbb{Q}}\left[\frac{B_t}{B_{t_i}} | \mathcal{F}_t\right] = p(t, t_i)$ we get

$$\delta_t = d_t \sum_{i=1}^M p(t, t_i) \left[Q(t, t_i) - Q(t, t_{i-1}) \right].$$

We note that this implies

$$\delta_t = d_t E^{\mathbb{Q}} \left[\frac{B_t}{B_\tau} \, |\mathcal{F}_t \right].$$

End of proof.

Given this expression, it is easily seen that the recovery rate futures price is strictly greater than the recovery rate δ_t since $E^{\mathbb{Q}}\left[\frac{B_t}{B_{\tau}} | \mathcal{F}_t\right] < 1$. This difference is expected to be small since interest rates are small over our sample period.