Journal of Banking & Finance 40 (2014) 257–270

Contents lists available at ScienceDirect

Journal of Banking & Finance

journal homepage: www.elsevier.com/locate/jbf

Is recovery risk priced?

Timo Schläfer [a], Marliese Uhrig-Homburg [b,*]

[a] Goldman Sachs International, EC4A 2BB London, UK
[b] Karlsruhe Institute of Technology (KIT), 76131 Karlsruhe, Germany

ARTICLE INFO

Article history:
Received 27 June 2013
Accepted 23 November 2013
Available online 5 December 2013

JEL classification:
G13

Keywords:
Credit risk
Recovery risk
Implied recovery rates
Risk premia

ABSTRACT

Recovery risk to explain corporate debt premia has not received much attention so far, most likely due to the difficulties around decomposing the expected loss. We exploit the fact that differently-ranking debt instruments of the same issuer face identical default risk but different default-conditional recovery rates. This allows us to isolate implied recovery under the T-forward measure without any of the rigid assumptions employed by prior studies. We find a pronounced systematic component in recovery rates for which investors should receive a premium. Comparisons to physical realizations show that the premium is quite time-stable and similar for different debt seniorities.

1. Introduction

Empirical research on corporate default risk shows that expected losses are insufficient to explain the corporate bond spread. Elton et al. (2001) therefore decompose the spread additionally into a tax premium and a risk premium arising from systematic factors related to corporate bond returns. The magnitude of this risk premium has long been a puzzle to economists. Recent research has looked at liquidity effects (Longstaff et al., 2005) and premia on default risk such as spread risk, jump-to-default risk (Driessen, 2005 or Berndt et al., 2008), and contagion risk (Collin-Dufresne et al., 2010). A risk premium associated with the uncertainty about the recovery rate is ignored in all of these studies, although a systematic component in recovery risk is likely to exist (Altman et al., 2005), and even diversifying the idiosyncratic component could be unachievable in practice. While the economic effects of uncertainty in default rates and recovery rates are clearly distinct, empirically disentangling default risk premia from recovery risk premia is a difficult task because default and loss rates are essentially multiplicatively linked in most approaches to credit risk modeling (Jarrow and Turnbull, 1995; Duffie and Singleton, 1999). For example, the very same corporate bond price can be due to a high risk-neutral default rate and a low risk-neutral loss given default or vice versa.

If there were no recovery risk, corporate credit risk would be largely unaffected by this identification problem. Research on the determinants of physical recovery in default, however, shows that recovery rates are stochastic and that they tend to be particularly low in times of financial distress and rising default rates. Most empirical evidence on recovery risk (Altman et al., 2005; Cantor and Varma, 2005; Acharya et al., 2007) stems from the physical measure. There is essentially no reliable data bearing on the stochastic nature of risk-neutral recovery rates. Our research is motivated by this lack of empirical evidence. While it extends earlier findings on the determinants of physical recovery rates to the risk-neutral world, our paper is concerned with much broader and more fundamental questions: Do investors price recovery risk? How does the recovery risk premium compare to the default risk premium?

To shed light on these questions, we devise an approach to isolating interest rate risk, default risk, and recovery risk under the T-forward measure without assuming stochastic independence between any of these factors. This is achieved by exploiting the fact that differently-ranking debt instruments of the same issuer face identical default risk but different default-conditional recovery rates. We show that the ratio of premia of two CDSs referencing such instruments is a function of recovery risk only and use this measure to estimate firms' implied probability distribution of recovery given default at a particular point in time. Our methodology therefore differs crucially from the standard CDS valuation approach that treats the recovery rate simply as a constant parameter.

For our sample of 37 mainly sub-investment grade corporates, we find that the mean of this distribution is much lower than that of its physical counterpart, suggesting the presence of a recovery risk premium. Further, we find that implied expected recovery

* Corresponding author. Tel.: +49 721 60848183; fax: +49 721 60848190.
E-mail address: uhrig@kit.edu (M. Uhrig-Homburg).

0378-4266/$ - see front matter
http://dx.doi.org/10.1016/j.jbankfin.2013.11.033

rates are strongly affected by the economic environment, declining substantially in times of distress and exhibiting a strong negative relation to implied default rates, especially so if these are high. This confirms that there is a systematic component in recovery risk that cannot be diversified away and justifies why the market should price such risk. By relating our estimates of implied recovery to physical realizations, we make the recovery risk premium and its evolution over time explicit. We find that coefficients of risk aversion are quite time-stable and similar for different debt seniorities. Based on our estimates of implied recovery rates, we can infer firms' implied probability of default at a particular point in time and find default risk premia to be somewhat lower than recovery risk premia. Robustness tests indicate that neither the assumed functional form of the implied probability distribution of recovery nor liquidity premia crucially drive our main results.

2. Related literature

Our study is related to three strands of the finance literature. First, it adds to the bourgeoning empirical research on risk premia in corporate debt markets such as the recent studies of Driessen (2005), DeJong and Driessen (2007), Berndt et al. (2008), Elkamhi and Ericsson (2008), and Collin-Dufresne et al. (2010). To the best of our knowledge, this study is the first to estimate recovery risk premia over time and for different seniorities. Our results imply a negative covariance between the pricing kernel and the recovery rate. As noted by Chen et al. (2009) this is one channel to reconcile the well-known discrepancies between predictions of structural models (documented, for instance, in Huang and Huang (2012)) and (much higher) observed market spreads.

Second, our work relates to the empirical literature on recovery rates under the physical measure. Studies such as Altman and Kishore (1996), Renault and Scaillet (2003), Altman et al. (2004), Cantor and Varma (2005), or Emery and Ou (2009) have shown that recovery in default is strongly affected by the ranking of an obligation and the assets securing it, with senior and well-collateralized instruments typically recovering significantly more than junior and/or unsecured obligations. A number of firm- and industry-specific factors are found to exert influence on recovery rates, as well. Acharya et al. (2007) show that recovery rates are lower if the issuer's industry is in distress or illiquid, particularly so if that industry's assets are specific, i.e. of limited use to other industries. Finally, the relevance of macroeconomic factors is well-documented, too: Recovery rates relate positively to GDP growth and S&P 500 returns and negatively to speculative-grade credit spreads and equity volatility (Altman et al., 2005; Cantor and Varma, 2005; Trück et al., 2005; Chava et al., 2011). This implies that there is a systematic component in recovery risk that cannot be diversified away. In addition, recovery rates are found to be negatively related to default rates, and Acharya et al. (2007) and Bruche and González-Aguado (2009) find that this is particularly so for senior, well-collateralized obligations. Our analysis shows that implied recovery rates, too, are related to proxies for firm- and industry-specific financial health. In particular, they tend to be higher for issuers with low leverage, a high share of tangible assets, strong liquidity, and more so if an issuer's industry is in a robust condition. Implied instrument-specific recovery rates are first and foremost driven by the reference obligation's seniority and the issuer's capital structure: Our estimates are on average more than twice as high for senior secured loans as for senior unsecured bonds and more than four times as high for senior unsecured bonds as for senior subordinated bonds. This extends earlier findings on the determinants of physical recovery rates to the risk-neutral world.

Third, our approach is related to the literature concerned with the estimation of implied recovery rates. The methods put forth by Bakshi et al. (2006), Gaspar and Slinko (2008), and Das and Hanouna (2009) have in common that they define implied default and recovery rates as functions of one (or more) state variable(s) such that a separation of both factors is feasible, provided that their relation is defined explicitly. A second procedure, suggested by Zhang (2003), Pan and Singleton (2008), and Schneider et al. (2010) employs CDS term structure data and relies on the assumption that implied recovery rates are constant across maturities. Zhang arrives at economically meaningful results, assuming additionally that recovery rates are constant over time. Another approach, suggested by Madan and Unal (1998), Unal et al. (2003), Le (2007), and Song (2008), derives equations that are entirely free of either default or recovery risk. In the first case (Madan and Unal, Unal et al.), this is achieved by relating the priority of junior and senior debt holders' claims to proxies of issuers' capital structure. This allows estimating the entire implied probability distribution of recovery given default, provided that the functional form of this distribution is specified. In the second case (Le, Song), implied probabilities of default are estimated from equity/equity option data and premia of spot and forward CDSs, respectively, and are then used to calculate implied recovery rates. In all four studies it is assumed that implied default and recovery rates are stochastically independent, or else a separation would not be feasible. In summary, prior approaches to estimating implied recovery suppose either (i) constant implied recovery rates (over time, over firms, or both), (ii) an explicit relation to the implied probability of default or (iii) independence between the two. Our method differs critically from these approaches in that we isolate interest rate risk, default risk, and recovery risk under the T-forward measure without assuming independence between any of these factors.

3. Methodology

CDSs allow trading the credit risk associated with certain debt-related events. The CDS buyer pays a periodic premium ("premium leg") in exchange for a default-contingent compensation ("protection leg"). At inception of the CDS, the premium is commonly chosen such that the value of both legs is identical.[1] The time t value of the premium leg $P_t(s_t, T) = s_t A_t(T)$ is the product of the CDS premium s_t and the price of an annuity $A_t(T)$ paying one until the reference entity defaults in τ or the CDS expires in T, whichever happens first. To evaluate the protection leg, assume that recovery is defined as a fraction of a present value of face (recovery of treasury).[2] If the reference entity defaults in $\tau \leqslant T$, the protection leg is associated with an insurance payment $PR_\tau(\rho, T) = (1-\rho)B_\tau(T)1_{\{\tau \leqslant T\}}$ at default τ, where $B_\tau(T)$ denotes the time τ-value of a default-free zero coupon bond with maturity T, ρ is the reference obligation's recovery rate,[3] and $1_{\{\cdot\}}$ is the indicator function that equals one if the event described in $\{\cdot\}$ occurs and zero otherwise.

The normalized process $PR_t(\rho, T)/B_t(T)$ is then a martingale under the T-forward measure $\widetilde{Q}$ and because of $B_T(T) = 1$ we can write:

$$PR_t(\rho, T) = E_t^{\widetilde{Q}}[(1-\rho)1_{\{\tau \leqslant T\}}]B_t(T). \tag{1}$$

Rearranging the right hand-side of Eq. (1) shows that the value of the protection leg is equal to the product of the default-condi-

[1] At the time of this writing, single-name CDSs have moved to fixed coupons of 100 BPs and 500 BPs, depending on the reference entity's credit quality. The present value of both legs then generally differs such that either party needs to make an upfront payment.

[2] Bakshi et al. (2006) test several recovery assumptions and find that market participants do not anticipate immediate recovery of face.

[3] In practice, more than one obligation may be deliverable under the CDS, resulting in a cheapest-to-deliver (CTD) option for the protection buyer. In this case, ρ should be interpreted with respect to the CTD obligation.