1 Introduction

Research on the determinants of recovery in default shows that there is a systematic component in recovery risk and that in times of financial distress and rising default rates, recovery rates tend to be particularly low. This has important ramifications for credit risk management and stress testing: Altman, Brady, Resti, and Sironi (2005) estimate that assuming constant recovery rates or independence from systematic factors underestimates value at risk, and hence economic capital, by approximately 30%. In its framework documentation on capital measurement and capital standards (Basel II), the Basel Committee on Banking Supervision accordingly demands that recovery estimates "reflect economic downturn conditions where necessary to capture the relevant risks"[1].

Market-implied recovery rates, whilst not directly conveying real-world expectations, concern the subject of risk management as well, albeit in a more subtle manner: In addition to the above, the provisions of Basel II require that "if recovery rates are negatively related to default rates, loss given default parameters must embed forecasts of future recovery rates that are lower than expected during more neutral conditions"[2]. For instance, in summer 2008 it was unambiguous that a major financial crisis was under way and that the months to come might see a considerable number of default events and possibly lower-than-average recovery rates. In such an environment the risk manager is thus obliged to reduce estimates of future physical recovery rates accordingly, but deciding just how substantial such an adjustment should reasonably be is at her own discretion. Historical recovery rates are of little help in answering this question: They are backward-looking by definition and only if data for comparable situations is available can they serve as a basis to infer reasonable assumptions. Implied recovery rates, on the other hand, carry information as to the market's expectation of future economic conditions. Analyzing how changes in these expectations have affected implied figures can thus provide more precise an indication of how model parameters need be altered to adequately account for current risks.

Reliable estimates of implied recovery rates are furthermore indispensable for the pricing of many credit-risky assets, although practitioners frequently use comparable actual realizations as a substitute. For instance, senior unsecured bonds have historically recovered around 40% and this figure is oftentimes used as a proxy for the implied expected recovery rate of credit

1 Basel Committee on Banking Supervision (2006), §468.
2 Basel Committee on Banking Supervision (2005), p2.

default swaps (CDSs) on these bonds. This practice, however, is problematic for at least three reasons. First, if the dynamics of implied recovery rates are similar to those observed under the physical measure, i.e. if figures are lower in downturns and vice versa, using a constant input factor will generally under-price the CDS in times of distress and generally over-price it if the economy fares well. Second, Le (2007) notes that employing historical observations does not adequately reflect the market's expectation of the future economic development. As mentioned earlier, implied estimates are forward-looking and thus more apt for this purpose. Third and maybe most importantly, if recovery risk is systematic and hence not diversifiable, investors will require a premium for bearing this risk such that implied expected recovery rates will on average be lower than comparable physical realizations. Güntay, Madan, and Unal (2003) and Pan and Singleton (2008) remark that using historical recovery rates for asset pricing therefore assumes risk-neutrality and disregards the risks involved.

Understanding the characteristics of implied recovery rates should therefore be of interest to a variety of market participants, be it risk managers, traders, or developers of forward-looking credit risk models. Efforts to the end of enhancing this understanding are, however, exacerbated importantly by an identification problem known since Jarrow and Turnbull (1995) and Duffie and Singleton (1999): In most approaches to credit risk modeling, default and loss rates are essentially multiplicatively linked, complicating a separate identification of either factor. Prior studies concerned with the extraction of implied recovery rates negotiate this hurdle only by employing rather implausible assumptions. Specifically, these are either i) constant implied recovery rates, ii) an explicit, arbitrarily imposed relationship between (stochastic) implied default and recovery rates or iii) independence between the two. These characteristics are at odds with what is known about recovery under the physical measure, and it thus is questionable whether implied dynamics are accounted for appropriately.

It is the objective of this thesis to devise a robust, assumption-light approach to estimating implied recovery rates from CDS premia that delivers reliable, economically meaningful estimation outcomes. To this end, two peculiarities in the cash flows of CDSs are exploited: First, the protection seller receives a payment only in the event the reference entity defaults but not otherwise. This is different from the structure of other credit-risky assets, such as bonds, under which payments can occur in either case. Second, the present value of the CDS premium payments is a function only of the implied probability of default but not of the implied expected recovery rate. If the CDS pricing equation is formulated under the T-forward measure, these two properties allow it to isolate interest rate risk, default risk, and recovery risk without assuming independence between any of these factors. As a result, it can be shown that the ratio of premia of two CDSs referencing a junior and a senior debt instrument of a given issuer, respectively, is a function only of the implied expected recovery rates of these two instruments but not of the implied probability of default. The identification problem is thus overcome in an elegant manner, and rigid presuppositions such as those employed by earlier studies are eschewed.

In a second step, issuers' capital structure is analyzed such that, based on the priority of claims, a relation between the ratio of firm value to liabilities at default and instrument-specific recovery rates can be established. Calibrating model parameters to market data then permits estimating the entire probability distribution of recovery conditional on default for a particular firm at a particular point in time. The mean and the standard deviation of this distribution are allowed to vary stochastically and no parametric relationship to the implied probability of default is imposed.

The practical implementation of this method is illustrated using CDSs referencing senior unsecured and senior subordinated bonds as well as loan-only credit default swaps (LCDSs), an altogether new asset class. LCDSs share the purpose of "traditional" CDSs in that they allow trading the credit risk associated with some debt obligation but are intended for use with leveraged loans as opposed to bonds. Leveraged loans are senior secured loans of sub-investment grade issuers and usually rank senior to all other debt of a borrower. Thus, debt of three distinct seniorities is employed in the estimation procedure, setting the analysis on a broad and stable basis.

The results of this thesis add manifold to the existing literature on implied recovery and shed light on several aspects of the topic that have hitherto not been examined at all or only in considerable less detail. Particular emphasis is on finding answers to the following questions:

- How are implied recovery rates related to firm-specific and industry-specific factors and what is their reaction to changes in the economic environment?
- How do debt cushion effects influence implied instrument-specific recovery rates in theory and in practice?
- How do average historical recovery rates compare to implied estimates and what does this relation reveal about the premium for taking recovery risk?
- Can estimates of implied recovery be used to deduce the implied probability of default and what is the relation of both factors under the pricing measure?

To prepare the ground for this program, Chapter 2 commences with a summary of what is known about recovery under the physical measure. This brings first insights into the general characteristics of recovery rates and later allows setting estimation results into perspective. Next, the difficulties associated with a separation of implied default and recovery rates are illustrated using the example of a defaultable zero coupon bond, followed by a survey of how earlier research tackled this issue. Particular attention is being paid to the methods and assumptions employed therein, and, based on the foregone discussion of physical realizations, the plausibility of results is assessed.

Chapter 3 then develops the approach pursued in this thesis, showing how recovery information can be extracted from pairs of CDSs on corporate debt. To the end of identifying suitable instruments for a practical implementation, Chapter 4 gives an overview of common types of CDSs and discusses potential combinations. This makes it necessary to compare in detail the properties of CDSs and LCDSs such that divergent contract specification can be properly accounted for.

Chapter 5 first describes the data and details the empirical specification of the model as well as the calibration procedure. Thereafter, estimation results are presented and the characteristics of implied recovery rates are examined from various vantage points. Robustness is assessed through two alternative model parameterizations, a different approach to calibrating the model, and by comparing implied estimates to historical realizations. Next, estimates of implied recovery rates are used to deduce implied probabilities of default and the relation of both factors under the pricing measure is examined. Again, historical estimates are employed to set implied estimates into perspective.

Finally, Chapter 6 summarizes the principal findings and concludes with suggestions for future research. The ideas and results presented in Chapters 3 and 5 are taken in large part from Schläfer and Uhrig-Homburg (2010b). Chapter 4 is a synopsis of Schläfer and Uhrig-Homburg (2010a).