Stochastic recovery rate: impact of pricing measure's choice and financial consequences on single-name products

Paolo Gambetti[1], Geneviève Gauthier[2] and Frédéric Vrins[1]*

[1] Louvain Finance Center & CORE, Université catholique de Louvain

[2] HEC Montréal and GERAD

Abstract

The ISDA CDS pricer is the market-standard model to value credit default swaps (CDS). Since the Big Bang protocol moreover, it became a central quotation tool: just like options prices are quoted as implied vols with the help of the Black-Scholes formula, CDSs are quoted as running (conventional) spreads. The ISDA model sets the procedure to convert the latter to an upfront amount that compensates for the fact that the actual premia are now based on a standardized coupon rate. Finally, it naturally offers an easy way to extract a risk-neutral default probability measure from market quotes. However, this model relies on unrealistic assumptions, in particular about the deterministic nature of the recovery rate. In this paper, we compare the default probability curve implied by the ISDA model to that obtained from a simple variant accounting for stochastic recovery rate. We show that the former typically leads to *underestimating* the reference entity's credit risk compared to the latter. We illustrate our views by assessing the gap in terms of implied default probabilities as well as on credit value adjustments (CVA) figures and pricing mismatches of financial products like deep in-/out-of-the-money standard CDSs and digital CDSs (main building block of credit linked notes, CLNs).

1 Introduction

Credit Default Swaps (CDSs) are credit derivatives allowing one party (*protection buyer*) to buy protection on a given *reference entity* from another party (*protection seller*) Fabozzi (2003). These instruments act as insurance contracts between the two parties, who trade the default risk of the reference entity for a given notional N up to a maturity date T. Before the crisis, the protection buyer had to make quarterly payments (defined by a specific calendar, called *IMM dates*) with amounts determined by the coupon rate c (called *running* spread), the contract's notional N and some day-count convention. This spread was the premium that makes the deal be worth zero at inception. Hence, the running spread agreed at inception with the prevailing *par* (also called *break-even*). In exchange of those premia, the protection seller agrees to make a payment (called *contingent flow*) to the protection buyer in case the reference entity effectively defaults prior to the contract's maturity. In the case of standard CDS contracts, the amount of the contingent flow depends on the contract notional as well as on the actual *recovery rate* of the firm that will be determined either by the residual value of issued Bonds or via an auction. The ISDA (International Swap and Derivatives Association) proposed a model to value CDS contracts based on a particular specification of the general no-arbitrage pricing equations. This model became the standard on the market. It is for example the default model in Bloomberg (the other alternatives being just variants of the latter), but also inspired most of the other pricing platforms like Markit or Summit, that differ from the ISDA model only by minor specificities.

But this way of trading CDSs is not really convenient from a Treasury management perspective. For instance, if a trader wants to close the position afterwards (i.e. at a different running spread) by entering in a reverse trade, she is left with a stream of residual payments. In order to facilitate back office operations, a standardization was needed. The fundamental reviewing of CDS trading conventions

*Contact information: Voie du Roman Pays 34, B-1348 Louvain-la-Neuve, Belgium. E-mail: frederic.vrins@uclouvain.be.

is called the Big Bang protocol Markit (March 13, 2009). While traders still quote CDS contracts in the form of a running spread (now called *conventional spread*), the mechanism of the quarterly payments made by the protection buyer has been standardized. The premium payments are decomposed in two parts: the first part consists in quarterly payments but, in contrast with the former convention, the coupon rate is no longer the quoted (*running*) spread. Instead, it is a standard coupon rate (k, say), being either 100 or 500 bps depending on the credit risk of the counterparty at inception.[1] As there is in general no reason that the credit risk premium of a reference entity (i.e. the quoted spread c) matches with the standard coupon rate k, a financial adjustment needs to take place. This correction takes the form of an upfront payment. For instance, if $c > k$, the quarterly payments made by the protection buyer to the protection seller do not properly compensate the value of the protection leg, and the trade would have a positive Mark-to-Market (MtM) to the buyer. To compensate the protection seller, the buyer has to pay (if positive, receive if negative) an additional (upfront) amount that precisely compensates the MtM. With this additional upfront amount (corresponding to the difference between (i) the present value of the payments made if the coupon rate were c and (ii) the present value of those made according to the standard rate k), the contract is *at par*. Therefore, the Big Bang protocol is just another way to schedule the premium flows. In such a setup, a trader willing to close her trade with another party only needs to make an upfront payment. In particular, all the quarterly flows will cancel each other, independently of the prevailing levels of quoted spreads. In this context, the ISDA model also plays the role of *converter*. Just like Black-Scholes formula allows to convert "quoted vols" to "cash amounts", the ISDA CDS pricer (also called *ISDA converter* since then) allows to quote deals on a running basis rather than on upfront amounts, which is much more intuitive and convenient for traders Markit (2004). The ISDA converter wipes out any ambiguity about how to convert quoted spreads to upfront amounts by making the present value of quarterly payments in the new setup (upfront + premia based on the standard coupon rate k) equal to the present value of the payments in the old setup assuming the quoted is the running spread c (i.e. without upfront, but replacing k by c in the premia).

On the top of being a pricer and a converter, it also became the central tool to select a pricing measure in (incomplete) credit markets. In the classical no-arbitrage setup, the value of a contract is the present value of the future cashflows. In the CDS case, this means that at inception, the value of the premium leg (upfront + present value of payments based on the standard coupon rate) agrees with the protection leg (present value of the contingent cashflow). The ISDA pricer consists in simplifying the pricing equations to have a simple form for both legs. Therefore, inverting the ISDA pricing equations, starting from quoted spreads of a set of CDS contracts, provides an easy way to pick up a pricing probability measure and, in particular, to back out a parametric risk-neutral default probability curve. This measure could then be used to compute other quantities, like credit valuation adjustments (CVA) or prices of non-standard credit-sensitive deals. Whereas most of these simplifications have little impact and are quite realistic, one of them may have substantial financial consequences: the recovery rate that determines the contingent payment in case of default is assumed to be known, and is typically set to 40% (or 20% for sovereigns). While 40% is indeed close to the average of observed recovery rates from past defaults, fixing the recovery rate to that level contradicts empirical evidence. Since the '90s in fact, researchers started to analyze recovery rates information and build evidence about their time-series and cross-sectional variation. It became clear that recovery rates exhibit significant differences across seniority levels of the defaulted bonds and sector to which issuers belong (Altman and Kishore (1996)): interactions among these two features can also play an important role. While it is straightforward to think that lack of collateralization and higher degrees of subordination lead to lower recoveries, industry effects have been justified by asset redeployability and/or anticipated government support considerations[2] (Shleifer and Vishny (1992), Acharya et al. (2007), Sarbu et al. (2013)). Moreover, there is clear evidence that recovery rates are negatively correlated with the business cycle and in particular with default probabilities: on average, the higher the default rate, the lower the recovery rate (Frye (2000), Hu and Perraudin (2002), Altman et al. (2005), Boudreault et al. (2013)). It is worth noting that the interest about recovery rates uncertainty is not only limited to academia, but it is also shared among the industry under the form of research reports, mainly by rating agencies (Hamilton et al. (2001), Cantor and Varma (2004), Keisman and Van de Castle (1999), Van de Castle and Keisman (2000)). The main message of all these studies is

[1] On the European contracts, a standard coupon rate of 25 bps is also considered for a couple of entities having a very high creditworthiness.

[2] This is particularly true in the banking sector, even though there are famous counter-examples.

again that recovery rates are unknown before default, and they should be considered as random variables. Surprisingly, none of the standard approaches account for this crucial point.

The goal of this paper is to show that in some circumstances, disregarding this reality may have important consequences in terms of pricing and risk-management of financial products. We first recall in Section 2 the no-arbitrage pricing equations associated to CDS contracts and show that the ISDA model essentially is an approximation where the recovery rate assumption plays a central role. We then introduce a simple CDS model in Section 3. The purpose of this model is to provide some intuition about the possible consequence of the recovery rate's uncertainty. The model is used in Section 4 as an external benchmark to identify the potential model risk embedded in the standard pricers. We start with an empirical analysis of recovery rates based on the Moody's Default and Recovery Database. This will provide guidelines to adjust the recovery rate parameters in our "in-house" CDS model. Then, the two models are compared from various points of views: implied default probabilities, CVA figures and mispricing of instruments like digital CDSs (customized CDSs that strike the recovery rate to 0% contractually, they are the cornerstone of credit linked notes, CLNs Fabozzi et al. (2007)) as well as on existing standard CDSs that are not explicitly quoted.

2 Pricing equations

We derive the pricing equations under the assumption that we are valuing deals at time t coming prior to maturity T and that the firm did not default yet. All contracts are valued from the standpoint of the protection buyer. In order to ease the notations the valuation date is used as a refence time ($t = 0$). Therefore, depending on the context, a same symbol s can be used to loosely denote either a date or the remaing time from the valuation date ($t = 0$) to date s. For instance, T represents both the contract maturity and the time-to-maturity. This is a common abuse of notation that drastically eases the mathematical exposition.

2.1 General no-arbitrage CDS pricing equations

The price of a CDS is given by the difference between the risk-neutral expectation of the discounted contingent cashflow with that of the discounted flows of the premium leg. Let us start with the protection leg. The default of the reference entity triggers the payment of the contingent flow at default's time, should the latter comes prior to the contract maturity T. The payment is a fraction $L = (1 - R)$ of the notional N where R stands for the recovery rate of the firm:

$$(1 - R)N\, \mathbb{1}_{\{\tau \leq T\}}\ ,$$

where $\mathbb{1}_{\{\omega\}}$ stands for the indicator function defined as 1 if ω is true and 0 otherwise. The corresponding present value in a non-arbitrage setup is

$$\mathrm{Prot}(T) := N\, \mathbb{E}\left[\mathbb{1}_{\{\tau \leq T\}} \frac{(1 - R)}{\beta_\tau}\right]\ ,$$

where $\mathbb{E}$ stands for the expectation operator under a chosen risk-neutral probability measure $\mathbb{Q}$ and β denotes the (risk-free) bank account numéraire.

Seen from the valuation time $t = 0$, the premium leg consists in a stream of (future) periodic flows paid at specific dates $t_1 < \ldots < t_m = T$. We denote by t_0 either the last payment date (prior to t) or the inception date (if none premium payment took place before t). The cashflows associated to the quarterly payments can be split in two parts. The first term deals with the protection between the valuation date and the next coupon date. The part of the coupon associated to the period (t, t_1) is dued if $\tau \geq t_1$. Otherwise, only a part of it (corresponding to the (t, τ) period) has to be paid. More precisely, let $\Delta(s, t)$ be the fraction of year between the dates (s, t) according to the specific day count convention. Then, the discounted flow associated to the credit protection for the period $(t, t_1]$ is[3]

$$kN\, \mathbb{1}_{\{\tau \geq t_1\}} \frac{\Delta(t, t_1)}{\beta_{t_1}} + kN\, \mathbb{1}_{\{\tau < t_1\}} \frac{\Delta(t, \tau)}{\beta_\tau}\ .$$

[3]Note that the period between the last payment date (or inception) t_0 and the valuation date does not enter the picture as it deals with past protection, and there is no reason to pay for it. The corresponding amount $kN\Delta(t_0, t)$ is called the *accrued* and explains the difference between clean and dirty prices as for Bonds.

The remaining terms deal with the following payments. As above, each of them can be split in two terms. A first term that deals with the case where there is no default up to the payment date (in which case the full coupon is due) and a second term (called *rebate*) accounting for the fact that whenever the default takes place between two payment dates, only a part of the coupon (corresponding to the period up to default) has to be paid. More specifically, the discounted cashflow is

$$kN\left(\sum_{i=2}^{m} \mathbb{I}_{\{\tau\geq t_i\}}\frac{\Delta(t_{i-1},t_i)}{\beta_{t_i}}+\sum_{i=2}^{m} \mathbb{I}_{\{t_{i-1}<\tau<t_i\}}\frac{\Delta(t_{i-1},\tau)}{\beta_\tau}\right) .$$

Hence, the risk-neutral present value of the premium flows is given by

$$\mathrm{Prem}(T) = \mathrm{up} + k\,\mathrm{PV01}(T) \;,$$

where $\mathrm{PV01}(T)$ is called the *risky duration* of the deal and is given by the present value of the premium payments based on a unitary coupon rate. Denoting $t_i^+ := \max(t, t_i)$ for conciseness [4], it comes

$$\mathrm{PV01}(T) \;:=\; N\sum_{i=1}^{m}\left(\Delta(t_{i-1}^+,t_i)\,\mathbb{E}\left[\frac{\mathbb{I}_{\{\tau\geq t_i\}}}{\beta_{t_i}}\right]+\mathbb{E}\left[\mathbb{I}_{\{t_{i-1}<\tau<t_i\}}\frac{\Delta(t_{i-1}^+,\tau)}{\beta_\tau}\right]\right) .$$

Consider the special case where the risk-free rate underlying the numéraire β is independent from both the default time τ and recovery rate R. Then,

$$\mathrm{Prot}(T) \;=\; -N\int_0^T (1-\mathbb{E}[R|\tau=u])P(u)dG(u) \tag{1}$$

$$\mathrm{PV01}(T) \;=\; N\sum_{i=1}^{m}\left(\Delta(t_{i-1}^+,t_i)P(t_i)G(t_i)-\int_{t_{i-1}^+}^{t_i}\Delta(t_{i-1}^+,u)P(u)dG(u)\right) \tag{2}$$

where

$$P(s) := \mathbb{E}\left[\frac{1}{\beta_s}\right] \;,\quad s\geq 0$$

is the price of a risk-free zero-coupon bond paying 1 unit of currency at time S and

$$G(s) := \mathbb{E}\left[\mathbb{I}_{\{\tau>s\}}\right] = \mathbb{Q}(\tau>s) = \mathbb{Q}(\tau\geq s) \;,\quad s\geq 0$$

is the risk-neutral survival probability of the reference entity[5]. Observe that the assumption that the reference entity did not default by the valuation time ($\tau > 0$) leads G to be non-increasing after t with $G(0) = 1$.

2.2 ISDA pricing equations

The ISDA model derives from JP Morgan routines. Except little subtleties that are negligible for our purposes, the protection leg and risky duration are given by [6]

$$\begin{aligned}
\overline{\mathrm{Prot}}(T) &:= (1-x)N\sum_{i=1}^{m}\bar{P}(t_i)\left(\bar{G}(t_{i-1}^+)-\bar{G}(t_i)\right)\\
\overline{\mathrm{PV01}}(T) &:= N\sum_{i=1}^{m}\Delta(t_{i-1}^+,t_i)\bar{P}(t_i)\frac{\bar{G}(t_{i-1}^+)+\bar{G}(t_i)}{2}\\
&= N\sum_{i=1}^{m}\Delta(t_{i-1}^+,t_i)\bar{P}(t_i)\left(\bar{G}(t_i)+\frac{\bar{G}(t_{i-1}^+)-\bar{G}(t_i)}{2}\right) .
\end{aligned}$$

[4] This notation is needed to deal with the term $i = 1$, to make sure that $t_{i-1}^+ = t_0^+ = t$ and not t_0

[5] In this paper, we assume that τ admits a density. In particular, $\mathbb{Q}(\tau = s) = 0$ for all $s \in \mathbb{R}^+$.

[6] In particular, we adopt a discretization scheme in line with the payment schedule (i.e. quarterly), and assume that all payments impacted by the occurence of the reference entity's default take place at the first payment date following the default event.

A key assumption is that it assumes a fixed loss-given-default, i.e. it requires the knowledge of the recovery rate x. In the standard ISDA pricer, the discount curve $\bar{P}(.)$ is the risk-free discount curve built from the prices of specific instruments taken from the previous day. The curve $\bar{G}$ represents the reference entity's survival probability, and is built in agreement with market quotes (see Section 2.3).

These equations resemble those obtained by no-arbitrage provided in (1) and (2). Indeed, replacing $\bar{P} \leftarrow P$ and $\bar{G} \leftarrow G$ and assuming independence between default time and recovery (setting $x \leftarrow \mathbb{E}[R|\tau = u] = \mathbb{E}[R]$) the protection leg expressions agree up to the discretization of the integral. Same applies to the risky duration, provided that one assumes that the payments always take place at payment date (even in case of default) and that if the default happens between two payment dates, half of the coupon is due. Most of these assumptions are known to have little impact. The impact of discretization is limited because the only effect is to discount from slightly different dates. For the same reason, the impact of choosing $\bar{P}$ instead of P is most of the time negligible. The parametric curve $\bar{G}$ allows to obtain a continuous survival probability curve from a limited number of quotes, but due to calibration constraints, its impact is minimal. All in all, the fundamental difference between the ISDA model and the no-arbitrage CDS equations is arguably the assumption regarding the dependence between variables and processes. The credit-rate independence assumption is known to be acceptable (Brigo and Alfonsi (2005)). Eventually, the fundamental specificity of the ISDA *versus* the general CDS equations has to be found in the protection leg, and with the treatment of the recovery rate in particular.

2.3 Calibration of the ISDA model to observed mark-to-market values

It is clear from above that $\bar{G}$ aims at representing the survival function under the pricing measure. Credit models are incomplete. Hence several risk-neutral measures exist. However, in order to avoid arbitrage opportunities, any chosen measure must be *calibrated to the market.* In other words, the chosen measure has to comply with market quotes. We now show how such quotes uniquely determine the parametric curve $\bar{G}$ in the ISDA model when a set of CDS calibration instruments is provided.

We start from a set of n (say) liquid CDS quotes of the reference entity with various maturities ($T_1 < T_2 < \ldots < T_n$, typically 1,3,5 and sometimes 10 years) called *calibration* instruments. As we have only n constraints, the curve $\bar{G}(.)$ will have n degrees of freedom. In the ISDA model, it is parametrized via a *hazard rate* function h as

$$\bar{G}(s) = \mathrm{e}^{-\int_0^s h(u)du} \ , \quad s \geq 0 \ .$$

The function h is piecewise constant between the maturities of the calibration CDSs with flat extrapolation beyond the last martutity. In order to *calibrate* the ISDA model, we make sure that the model and market prices of the calibration instruments agree. In particular, we make sure that the protection and premium legs of the ISDA model match when the coupon rate is set to the quoted spread $c(T_1), \ldots, c(T_n)$ for the assumed recovery rates $\vec{x} = (x_1, \ldots, x_n)$.[7] Hence, this curve (i.e. the piecewise constant levels of h) is constructed iteratively (in the order given by the products' maturities) so as to ensure

$$c(T_i) = \frac{\overline{\mathrm{Prot}}(T_i)}{\overline{\mathrm{PV01}}(T_i)} \ , \quad i \in \{1, 2, \ldots, n\} \ . \tag{3}$$

These calibration constraints simultaneously provide the function $\bar{G}$ as well as the legs $\overline{\mathrm{Prot}}$ and $\overline{\mathrm{PV01}}$. As stressed before, the coupon rate actually used in quarterly payments is not the quoted spread $c(T)$ but instead the standard coupon rate k, and the MtM difference is compensated via the upfront. The ISDA model is the market standard procedure to determine the actual upfront amounts from quoted (conventional) spreads:

$$\mathrm{up}(k, T_i) = \overline{\mathrm{Prot}}(T_i) - k\overline{\mathrm{PV01}}(T_i) = (c(T_i) - k)\overline{\mathrm{PV01}}(T_i) \ . \tag{4}$$

The way $\bar{G}$ is constructed combined with the way c is converted to an upfront amount implies that all calibration instruments have, by construction, an ISDA-model price that is inline with the market.

[7] Usually, the term-structure of recovery rate is flat, i.e. $x_1 = \ldots = x_n = x$.

2.4 Calibration of in-house model to observed mark-to-market values

Let us now consider an alternative (*in-house*) model controlled by a set $\vec{\theta}$ of exogenous parameters. As before, the model price of the calibration instruments must agree with those implied by the model. There is a fundamental difference, however. Whereas the ISDA model is calibrated from quoted (conventional) spreads, the calibration of the in-house model has to be done on the *actual prices*. In particular, the target that we have to meet is that, for all calibration instruments, the difference between the model-implied protection leg $\mathrm{Prot}(T)$ and the risky duration $\mathrm{PV01}(T)$ needs to agree with the upfront resulting from the application of the ISDA converter to the quoted spreads (as per eq. (4)):

$$\mathrm{up}(k, T_i) = \mathrm{Prot}(T_i) - k\,\mathrm{PV01}(T_i)\ . \tag{5}$$

This puts some constraints on the protection leg and the risky duration and so on the survival probability curve G. The procedure is illustrated on Figure 1. The key point here is to notice that the differences between the two models will change the selected risk-neutral measure. In particular, there is no reason that the survival probability in the ISDA model (computed with the help of the measure selected during the ISDA calibration procedure) agrees with that of the in-house model (that is based on another measure, the one picked up from the in-house model calibration step).

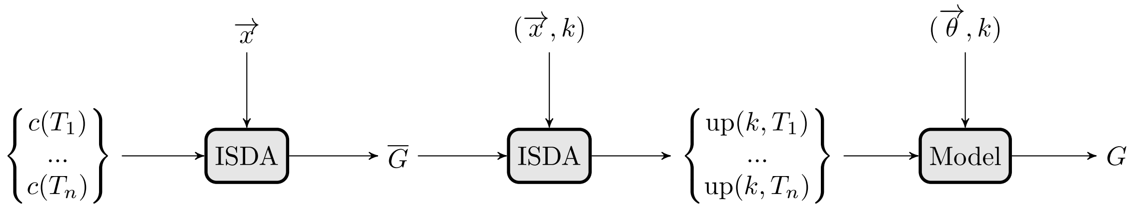

Figure 1: Methodology to extract the model implied default probability curve G from n quoted spreads. The procedure works as follows: (1) quoted spreads are first used to compute the ISDA's curve $\bar{G}$, (2) the ISDA model fed with $\bar{G}$ is then used to convert quoted spread to upfront amounts and (3) the survival probability curve G of the in-house model is tuned such that it implies the same prices (upfront) for the instruments used in the calibration procedure.

As explained above, the survival probability curve $\bar{G}$ obtained by inverting the ISDA CDS pricer equations provide a rather fair estimation of the risk-neutral default probability curve G *when the recovery rate is known in advance* (and agrees with the value plugged in the ISDA model). The other discrepancies have indeed very little impact. However, the curve $\bar{G}$ is obtained by calibrating the ISDA model (that assumes a fixed recovery) to prices of standard (i.e. floating recovery) CDS. This inconsistency suggests that one makes an error when assuming that the curve $\bar{G}$ is a valid proxy for G, depending on the stochastic properties of the recovery rate. In the next section, we provide a simple model to illustrate the potential impact of this way of working.

3 A simple model

The first models of stochastic recovery rates have been developed in the context of risk management. We refer the reader to Frye (2000), Jarrow (2001), Jokivuolle and Peura (2000), and Pykthin (2003), just to name a few. To the best of our knowledge, the pioneer work with regards to random recovery rates in the context of credit derivatives pricing is due to Andersen and Sidenius. In Andersen and Sidenius (2004), the authors extend the One Factor Gaussian Copula model for CDO pricing (also known as Li's model, Li (2016)), which was the market standard at that time. The credit crisis triggered a specific interest for stochastic recovery models. At some point, super senior CDO tranches (with attachment point higher than 60%) started to trade. However, such tranches are not worth anything under a fixed 40% recovery rate assumption. This is the best evidence that recovery rates have either to be decreased or, more realistically, have to be made stochastic. Since then, many stochastic recovery rate models have been introduced for the sake of pricing CDOs. We can mention for instance Gaspar and Slinko (2008), Ech-Chatbi (2008), Krekel (2010), Amraoui et al. (2012).

Surprisingly however, the impact of stochastic recovery rates on single-name CDS did not receive much attention. Yet, some authors tackled this point. For instance, Boudreault et al. (2013), Boudreault et al. (2015) and Bégin et al. (2017) propose a credit risk model in which the default intensity and the recovery rate are a non-linear function of the firm leverage ratio. This approach allows to capture the negative relationship between default probability and recovery rate observed in the empirical studies. The model is estimated using a filtering approach on a sample of CDS premiums. They show that the interrelation between recovery rate and default probability modifies the term structure of zero-coupon yield to maturity and impacts significantly the standard risk measures such as the VaR and the expected shortfall. While these methods can be adapted for this purpose, we prefer to introduce a simple model that is tailored to illustrate the potential impact of disregarding recovery rates uncertainty. Our goal here is indeed to put forward the impact of the uncertainty of recovery rate as well as its correlation with default rate on single-name credit derivatives. Therefore, it is more convenient to work with a model based on the standard approach, but adjusted for the effect we want to analyze.

In the sequel, we analyze "how far" can the default probability curve $\bar{G}$ extracted from the ISDA pricer be from the actual default probability curve G extracted from a CDS pricing model that would account for the stochastic nature of recovery rate. To concentrate on the effect we are effectively interested in, we disregard some technicalities that are known to have a minor impact on CDS valuation. First, we assume a unit notional ($N = 1$) and that the discount curve P and $\bar{P}$ agree and are parametrized by a constant and known risk-free rate r so that

$$P(s) = \bar{P}(s) = \mathrm{e}^{-rs}, \quad s \geq 0 \,.$$

Second, we make some simplifications in terms of payment schedule; the premiums are paid on a continuous-time basis. Finally, we assume that we have only one calibration instrument ($n = 1$). All these assumptions can be easily relaxed but significantly simplify the exposition. For the ISDA pricing equations, we assume that a single calibration CDS instrument is used, so that $h(s) = h$ and $\bar{G}(s) = \mathrm{e}^{-hs}$ for $s \geq 0$. Setting the ISDA recovery rate level to the value x, the corresponding protection leg and risky duration become

$$\begin{aligned}
\overline{\mathrm{Prot}}(T) &= -(1-x)\int_0^T P(u)d\bar{G}(u) = (1-x)h\int_0^T \mathrm{e}^{-(r+h)u}\,du = (1-x)\frac{h}{r+h}(1-\mathrm{e}^{-(r+h)T})\,, \\
\overline{\mathrm{PV01}}(T) &= \int_0^T P(u)\bar{G}(u)du = \int_0^T \mathrm{e}^{-(r+h)u}\,du = \frac{1-\mathrm{e}^{-(r+h)T}}{r+h}\,.
\end{aligned}$$

From (3), we find the *credit triangle* $h = c(T)/(1-x)$.

We now introduce our simple in-house model. A standard approach in credit risk modeling consists in representing the default time τ as the first jump of a Poisson process with intensity λ under $\mathbb{Q}$. By doing so, $G(T) = \mathbb{Q}(\tau > T) = \mathrm{e}^{-\lambda T}$. Accounting for the possible dependency of R and τ, one gets

$$\begin{aligned}
\mathrm{Prot}(T) &= \lambda\int_0^T (1-\mathbb{E}[R|\tau = u])\,\mathrm{e}^{-(r+\lambda)u}\,du\,, \\
\mathrm{PV01}(T) &= \int_0^T P(u)G(u)du = \int_0^T \mathrm{e}^{-(r+\lambda)u}\,du = \frac{1-\mathrm{e}^{-(r+\lambda)T}}{r+\lambda}\,.
\end{aligned}$$

Interestingly, we observe that the risky duration in both the ISDA and the in-house models reads $\mathrm{PV01}(h,T)$ and $\mathrm{PV01}(\lambda,T)$ where

$$\mathrm{PV01}(z,T) := \frac{1-\mathrm{e}^{-(r+z)T}}{r+z}\,.$$

As explained above, the protection leg needs to be modeled with care. One directly observes that when R and τ are independent, we have $h = \lambda$ (i.e. $G = \bar{G}$) whenever $\mathbb{E}[R] = x$. Generally speaking however, credit events and recovery rates are not independent. The most famous study supporting that claim is undoublty that of Altman et al (Altman et al. (2005)). This is why an alternative to ISDA, accounting for that effect, needs to be considered.

An easy choice consists in adopting a static copula setup. The idea is similar to the resampling approach used in counterparty risk application to compute the credit value adjustment (CVA) under wrong-way or right-way risk (Gregory (2010), Sokol (2011) or Vrins (2016)). One models the conditional expectation $\mathbb{E}\left[R_\tau|\tau=s\right]$ using a function $f(s)$ that collapses to $\mu_R := \mathbb{E}[R]$ when there is no dependency between τ and R. We start by fixing the marginal laws of R and τ. The survival probability function of τ is implied by our Poisson model and is denoted by $G(s)$. On the other hand, R is modeled with a Beta distribution with shape parameters α and β whose distribution function is denoted by $F_R(.;\alpha,\beta)$. Then, R and τ are coupled with a copula C with a fixed dependency parameter ρ, so that one can sample (R,τ) using independent uniform random variables (U,V):

$$(R,\tau) \sim C\left(F_R^{-1}(U;\alpha(\tau),\beta(\tau)), F_\tau^{-1}(V);\rho\right) .$$

Choosing the Gaussian copula with correlation ρ, the conditional distribution of R given $\tau \in ds$ is

$$f(s;\rho,Z)ds := F_R^{-1}\left(\Phi\left(\rho\Phi^{-1}(G(s)) + \sqrt{1-\rho^2}Z\right);\alpha,\beta\right) ds$$

where $Z \sim \mathcal{N}(0,1)$ and Φ its cumulative distribution function. Hence,

$$\mathbb{E}[R|\tau \in ds]/ds = \mathbb{E}\left[f(s;\rho,Z)\right] =: f(s;\rho)$$

so that

$$\mathbb{E}\left[\frac{R\,\mathbb{I}_{\{\tau\leq T\}}}{\beta_\tau}\right] = -\int_0^T f(u;\rho)P(u)dG(u) .$$

Eventually, the protection leg takes the following form:

$$\mathrm{Prot}(T) = \frac{\lambda\left(1-\mathrm{e}^{-(r+\lambda)T}\right)}{r+\lambda} - \lambda\int_0^T f(u;\rho)\,\mathrm{e}^{-(r+\lambda)u}\,du = \lambda\left(\mathrm{PV01}(\lambda,T) - \int_0^T f(u;\rho)\,\mathrm{e}^{-(r+\lambda)u}\,du\right) .$$

In order to correctly price the CDS contract with this model, the difference between the protection leg and the present value of the periodic premium payments must correspond to the upfront amount (the MtM of the deal without upfront); this is summarized in Figure 2. This in turns means that λ needs to satisfy eq. (5) which in this case reads as

$$\mathrm{up}(k,T) = (\lambda-k)\,\mathrm{PV01}(\lambda,T) - \lambda\int_0^T f(u;\rho)\,\mathrm{e}^{-(r+\lambda)u}\,du .$$

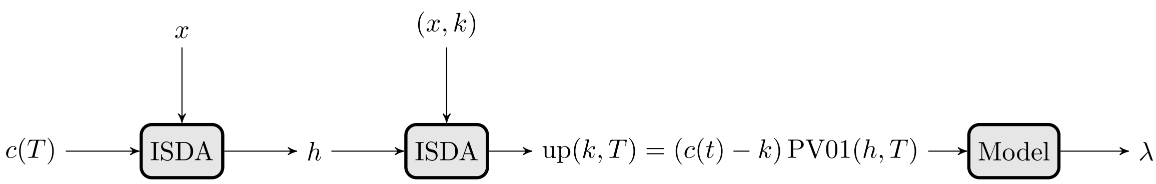

Figure 2: Methodology to calibrate the model's parameters to correctly price standard instruments. For simplicity, we assume here that only one CDS is quoted on the market so that only one conventional spread (with maturity T) is available for the reference entity. Hence, the hazard rate curve h collapses to a constant, and same for the default rate of the in-house model. The procedure works as follows: (1) quoted spreads are first used to compute the ISDA's parameters, (2) the ISDA model is used to convert quoted spread to upfront amounts and (3) the in-house model is calibrated such that it implies the same upfronts for the standard instruments.

Interestingly, $f(u;0) = \mu_R$ whenever $\rho = 0$ or when the variance of R vanishes. This means that if the recovery rate is deterministic (and equal to the constant used in the ISDA model) or if it is stochastic but independent from default's time (but whose risk-neutral expected value agrees with the constant used in the ISDA model), then $\lambda = h$. On the other hand, one can see that when $\rho > 0$, a high default probability (lower G) implies on average a low recovery rate. Hence, the observation in Altman et al.

(2005) suggests that one should use $\rho > 0$.[8] If we believe in our in-house model, the correct default rate λ can be obtained by looking at the hazard rate h implied by the ISDA model. In the other cases however, estimating the default rate λ (and equivalently the default probability curve G) from the hazard rate h leads to an error. We provide some order of magnitudes in the next section.

4 Numerical Experiments

We take the α and β shape coefficients so as to satisfy two constraints about the mean μ_R and variance σ_R^2 of the recovery rate. First, we want the conditional expectation of R given $\tau \in ds$ to be a given value μ_R when $\rho = 0$, i.e. $f(t;0) = \mu_R$. Second, we would like the variance σ_R^2 to be "valid". As the variance of a distribution whose support is $[0,1]$ and mean is π is upperbounded by that of a Bernoulli with parameter π, we choose:

$$\sigma_R = a\sqrt{\mu_R(1-\mu_R)} \quad , \quad a \in [0,1] \; .$$

In this expression, a controls the uncertainty we have about the recovery rate around the mean. These two constraints yield the shape and scale parameters

$$\alpha := \mu_R \left(\frac{\mu_R(1-\mu_R)}{\sigma_R^2} - 1 \right) \; , \quad \beta := \frac{1-\mu_R}{\mu_R}\alpha \; .$$

In the sequel, we consider the "ideal case" where the deterministic recovery rate value x used in the ISDA pricer is 40% does indeed correspond to $\mathbb{E}[R] = \mu_R$ (the alternatives will of course lead to more significant errors).

4.1 Maximum Likelihood estimation of the parameters

The shape parameters of the conditional distributions of recovery rate can be easily estimated either by moment matching techniques (plugging the sample mean and sample variance into the above formulas) or via maximum likelihood. The latter method works as follows. Given the probability density function of the Beta distribution parametrized by $\alpha, \beta > 0$

$$f(x;\alpha,\beta) = \frac{1}{B(\alpha,\beta)} x^{\alpha-1}(1-x)^{\beta-1} = \frac{\Gamma(\alpha+\beta)}{\Gamma(\alpha)\Gamma(\beta)} x^{\alpha-1}(1-x)^{\beta-1},$$

where $B(\cdot,\cdot)$ and $\Gamma(\cdot)$ denote the beta and gamma functions respectively, the likelihood function, given the sample of recovery rate observations $X = (x_1, ..., x_n)$, has the following form:

$$\mathcal{L}_n(\alpha,\beta|X) = \prod_{i=1}^{n} f(x_i;\alpha,\beta) = \prod_{i=1}^{n} \frac{\Gamma(\alpha+\beta)}{\Gamma(\alpha)\Gamma(\beta)} x_i^{\alpha-1}(1-x_i)^{\beta-1} \; .$$

The maximum likelihood estimator of α and β is then given by:

$$\frac{\partial \ell_n(\alpha,\beta|X)}{\partial \alpha} = n\Psi(\alpha+\beta) - n\Psi(\alpha) + \sum_{i=1}^{n} \log(x_i) = 0 \tag{6}$$

$$\frac{\partial \ell_n(\alpha,\beta|X)}{\partial \beta} = n\Psi(\alpha+\beta) - n\Psi(\beta) + \sum_{i=1}^{n} \log(1-x_i) = 0 \tag{7}$$

where $\Psi(\cdot)$ denotes the digamma function. In Table 1, we report summary statistics of recovery rate empirical distributions taken from Moody's Default and Recovery Database, together with the mean μ^{ML} and standard deviation σ^{ML} of the theoretical Beta conditional distributions obtained by maximum likelihood estimation with respect to the shape parameters α and β.[9] Conditioning is made with respect

[8]The reason why a negative relationship between recovery rate and default intensity corresponds to a positive ρ stems from the fact that the default-time is negatively correlated with the default probability: the higher the default intensity, the sooner the default time, on average.

[9]Note that, in the log-likelihood function we have to subtract a machine-precision quantity from full recoveries (i.e. when the recovery rate is equal to 100%).

	Min.	1st Qu.	Median	μ	3rd Qu.	Max.	σ	μ^{ML}	σ^{ML}
Junior Subordinated	0.63	8.00	13.99	20.34	33.00	74.00	18.98	20.68	17.74
Senior Subordinated	0.01	10.31	25.00	30.39	44.14	100.00	24.04	31.46	25.22
Senior Unsecured	0.01	15.00	31.94	36.63	56.06	100.00	26.03	37.85	26.80
Senior Secured	0.01	20.00	40.00	44.57	65.36	100.00	27.73	46.81	28.81
Banking	0.01	3.94	18.00	23.55	37.00	92.08	22.70	23.91	22.57
Capital Industries	0.01	14.43	30.14	36.35	57.00	100.00	26.18	37.05	26.81
Consumer Industries	0.01	15.00	31.67	36.95	55.00	100.00	25.73	38.60	26.48
Energy & Environment	0.01	19.00	36.75	38.86	52.82	100.00	25.70	39.89	27.93
FIRE	0.13	10.00	25.00	32.48	46.63	100.00	25.36	35.47	26.50
Media & Publishing	0.01	16.00	33.50	38.45	54.00	99.00	26.99	38.74	27.84
Retail & Distribution	0.50	13.62	29.00	33.52	48.62	99.50	25.63	35.97	26.98
Technology	0.25	10.00	23.75	30.32	45.00	100.00	25.36	32.91	26.18
Transportation	1.75	16.00	25.38	32.76	45.88	99.88	21.99	34.69	22.19
Utilities	13.99	43.63	67.18	62.87	84.65	100.00	26.46	65.27	27.26

Table 1: Summary statistics of empirical conditional distributions of recovery rates where recoveries are expressed as percentages of the bond face value. The last two columns display the mean and the standard deviation of the theoretical distribution with shape parameters α and β obtained via maximum likelihood. Data for the analysis are taken from Moody's Default and Recovery Database: the sample includes 2035 north-American bond defaults covering the period 1$^{\text{st}}$ January 1912 – 23$^{\text{rd}}$ January 2017. Regarding the optimization, we adopted a quasi-Newton method (L-BFGS algorithm with lower-bound constraints).

to the seniority of the defaulted Bond or the industry of the issuing company.
These results support the empirical findings discussed in Section 1 and confirm that the assumption of a constant recovery rate of 40% is indeed misleading. We observe important differences both in expected recovery rates μ^{ML} of different Bond seniorities (with mean recovery rates increasing with the level of seniority) and in their variability σ^{ML}. We also document a large intra-class variability (i.e. considerable values of σ^{ML} once we have conditioned for a specific seniority). In particular, we point out the higher variability of recoveries on Senior Secured Bonds: given the mean being around 50% and the standard deviation being nearly $1/\sqrt{12} \approx 28.87\%$, one should notice that this conditional distribution is in fact close to be Uniform in $[0, 1]$.[10] Similarly to Bond seniorities, we document sensible differences in mean and standard deviation of recovery rates conditional distributions when conditioning is made on the industrial sector. Also for this type of conditioning the recovery rates intra-class variability is high. In the sequel, we run our experiments as if the expected recovery rate μ^{ML} were exactly equal to the fixed one of 40% used as input in the ISDA pricer but with different levels of uncertainty. However, simulations could be run also by taking as inputs the last two columns of Table 1: this would lead to an amplification of the results. The time-evolution of recovery-rate using sliding window is depicted on Figure 3.

4.2 Impact on implied default probability

We look at the impact of the correlation parameter (ρ) for two uncertainty levels about the recovery rate value: a small uncertainty ($a = 10\%$ leading to $\sigma_R = 5\%$) and a large uncertainty ($a = 50\%$ leading to $\sigma_R = 25\%$, which is indeed close to the average σ^{ML} of Table 1). We consider a single calibration instrument which is a $T = 5Y$ CDS with running spread of 250 bps. The standard coupon is set to $k = 100$ bps. Figure 4 shows the mismatch between λ and h as well as a similar information but translated into survival probability at maturity to ease the interpretation. The ISDA model always returns the same constant value located at the intersection of the dotted lines. One can see that for $\rho > 0$ (suggested by empirical evidences), the survival probability implied by the "in-house" model is lower than that of the ISDA model. To put it differently, the ISDA model underestimates the (risk-neutral) default likelihood

[10]Uniform distribution is a particular case of the Beta distribution when both the shape parameters α and β are equal to one (indeed our estimation of this conditional distribution yields $\alpha = 0.9$ and $\beta = 1$). The flat distribution for Senior Secured bonds is in accordance with the findings of Schuermann (2004).

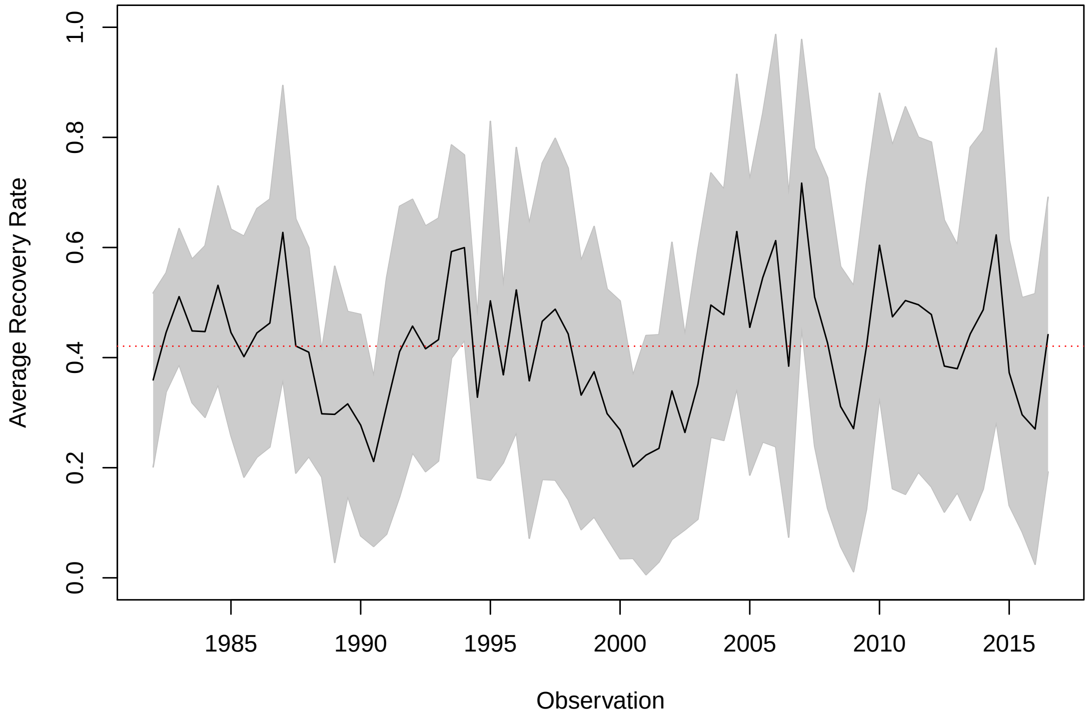

Figure 3: Half-yearly averages of recovery rates black solid line) with the ± 1 standard deviation envelope (grey area) for north-American bond defaults in the period 5^{th} January 1982-31^{st} December 2016. The red dotted line corresponds to the average recovery rate observed in this period.

in such cases.

4.3 Impact on deep in-/out-of-the-money standard CDSs

Suppose that a trader wants to price a CDS with outstanding maturity of 3Y and that the market only quotes a 5Y contract. The ISDA procedure would be to first extract the survival probability curve which, in the simplified setup depicted above, is given by the credit triangle: $h = c(5Y)/(1-x)$ (where x is the recovery rate assumed. This in turns indicates that the break-even spread for the 3Y contract is, $c(3Y) = c(5Y)$. In our simple in-house model, the procedure is slightly different. We still assume $\mu_R = x = 40\%$ but the implied default rate λ depends on ρ and a. From these assumptions, one can extract the break-even spread of the 3Y contract. Eventually the upfront is given by scaling the difference between the later spread and the standard coupon rate k with the risky duration of the deal computed according to the model. As above, one needs to convert the model spread $c(3Y)$ to a quoted spread $\hat{c}(3Y)$ in agreement to the quoting convention. A similar development as above yields

$$\hat{c}(3Y) = k + (c(3Y) - k)\frac{\text{PV01}(\lambda, 3Y)}{\text{PV01}(h, 3Y)} \, .$$

Figure 5 summarizes the methodology and Figure 6 illustrates the impact in terms of conventional spreads and MtM.

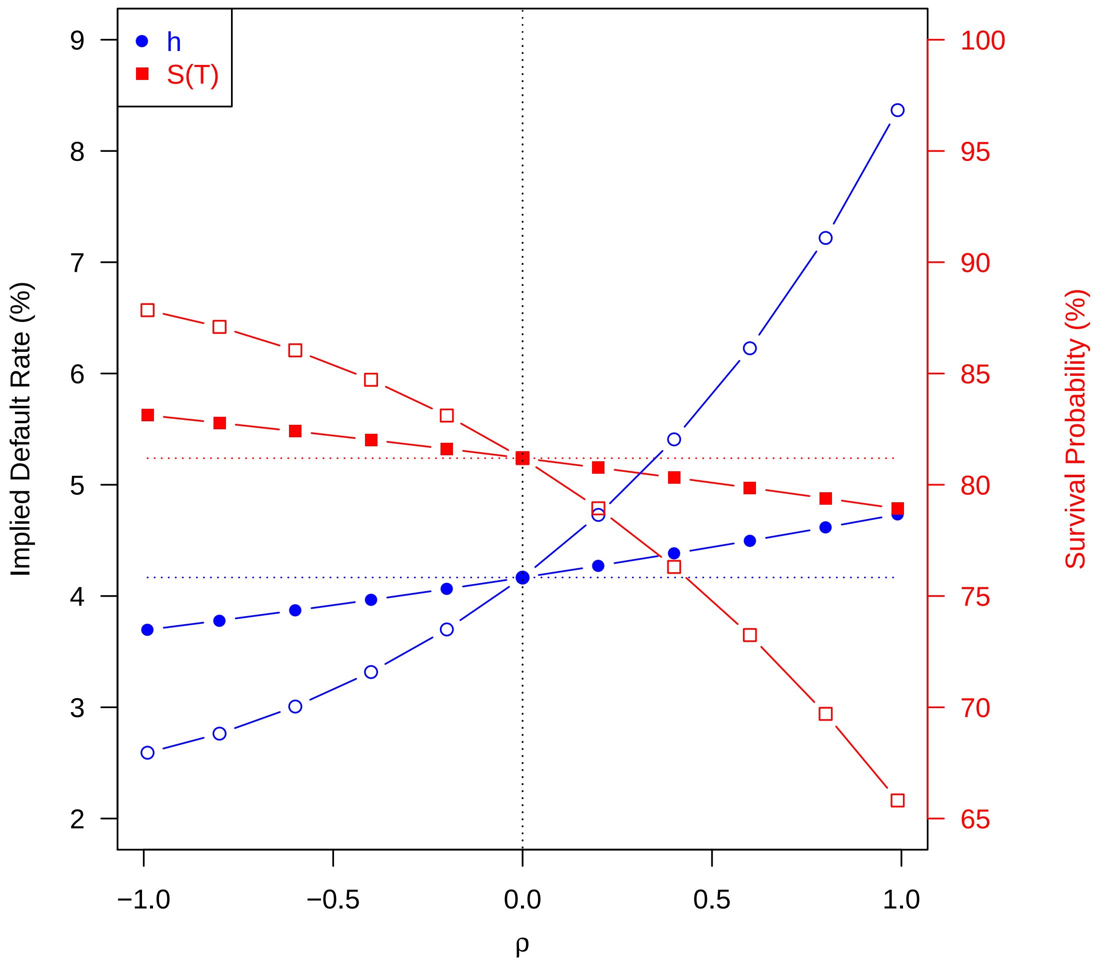

Figure 4: Impact of correlation on the implied default rate (blue, solid+dotted markes, left axis) and survival probability at maturity (red, solid+red squares, right axis) for two different values of recovery rate volatility: $a = 1/10$ (filled markers) and $a = 1/2$ (empty markers). Parameters: $x = \mu_R = 40\%$, calibration instrument: $T = 5Y$ maturity swap with $c(T) = 250$ bps, $k = 100$ bps.

4.4 Impact on CVA of a call

Credit value adjustment (CVA) is the current market price corresponding to the option – implicitly given to our counterparty– to default during the life of a trade. Let us consider a deal where we trade a call option with maturity T and strike K on a stock S (with volatility σ) with a counterparty whose default time τ is independent from S. Then, letting $\tilde{C}_s := C_s/\beta_s$ be the time-s discounted price of the call and G is the survival probability of the counterparty, Brigo and Vrins (2017)

$$\mathrm{CVA} = \mathbb{E}\left[\mathbb{I}_{\{\tau<T\}}(1-R)\frac{C_\tau^+}{\beta_\tau}\right] = -\int_0^T (1-\mathbb{E}[R|\tau=s])\,\mathbb{E}[\tilde{C}_s^+]dG(s)\;.$$

But the price of a call and the numéraire β are always non-negative, so that $\tilde{C}_s^+ = \tilde{C}_s$ and $\mathbb{E}[\tilde{C}_s] = \tilde{C}_0 = C_0$ as $\tilde{C}$ is a $\mathbb{Q}$-martingale. It becomes clear that in this context, CVA is nothing but the protection leg

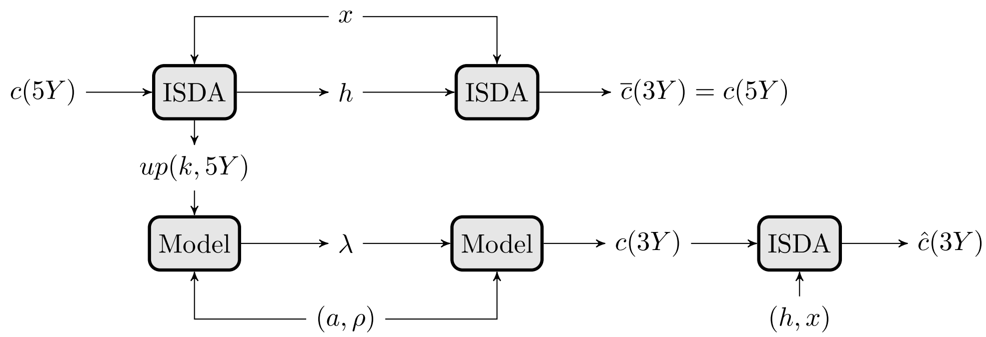

Figure 5: Methodology to analyze the model impact on the 3Y spread of a standard CDS starting from a 5Y spread of a standard CDS.

of a CDS whose reference entity is the counterparty, with zero risk-free rate and notional C_0:

$$\mathrm{CVA} = -C_0 \int_0^T (1 - \mathbb{E}[R|\tau = s])dG(s) \ .$$

In Figure 7(a) we compare two types of CVAs. The first one is computed by assuming an independent recovery rate whose expected value is equal to the ISDA level, $\mathbb{E}[R|\tau = s] = x = 40\%$ and used $\lambda \leftarrow h$ to be consistent:

$$\overline{\mathrm{CVA}} = h \int_0^T (1 - x)\, \mathbb{E}[\tilde{C}_s^+]\, \mathrm{e}^{-hs}\, du = (1 - x)C_0(1 - \mathrm{e}^{-hT}) \ ,$$

so that the CVA on a Call option with maturity T per unit of option premium is the ISDA protection leg $\overline{\mathrm{Prot}}(T)$ with $r = 0$. This is the horizontal black curve on Figure 7(a).

The second way to compute CVA is by extracting λ with our in-house model assuming a given pair (a, ρ) and then price CVA consistently. Hence, we are consistent with the ways recovery rate is considered in both the calibration and in the CVA pricing steps:

$$\mathrm{CVA} = \lambda \int_0^T (1 - f(s, \rho))\, \mathbb{E}[\tilde{C}_s^+]\, \mathrm{e}^{-\lambda s}\, ds = \lambda C_0 \int_0^T (1 - f(s, \rho))\, \mathrm{e}^{-\lambda s}\, ds \ .$$

As before, the “in-house” CVA, expressed per unit of option premium, is the “in-house” protection leg $\mathrm{Prot}(T)$ with $r = 0$. This is the blue curve on Figure 7(a). One can see that if we are consistent in the way we extract default probabilities and price CVA, the error is limited (the blue and black curves do not deviate too much from each other). If by contrast we lack consistency, i.e. use the internal model to compute the default probability curve but assume a fixed recovery in the CVA’s payoff as below,

$$\widetilde{\mathrm{CVA}} = \lambda(1 - x) \int_0^T \mathbb{E}[\tilde{C}_s^+]\, \mathrm{e}^{-\lambda s}\, ds = (1 - x)C_0(1 - \mathrm{e}^{-\lambda T}) \ .$$

This is similar to computing CVA assuming a deterministic recovery rate of x but assuming a stochastic recovery rate when extracting the default probability (by using λ instead of h). The inconsistency is striking by comparing this CVA (red curve) with the others on Figure 7(a).

4.5 Impact on par spread of digital CDSs

As a final experiment we focus on the price of digital CDS, i.e. a CDS contract where the recovery rate is *contractually* set to 0, so that the contingent flow is either 0 if $\tau > T$ or N otherwise. Applying the ISDA procedure, the hazard rate h is extracted from standard CDS, and the valuation of the digital CDS with the ISDA approach naturally leads to the spread $\bar{d}(5Y) = h = c(5Y)/(1 - x)$. Similarly, once the default rate λ of our in-house model has been extracted, it is straightforward to check that the break-even spread implied by the in-house model is $d(5Y) = \lambda$. Note that a proper comparison would require not to compare $\bar{d}$ and d but to compare $\bar{d}$ with the “ISDA-equivalent” spread $\hat{d}(5Y)$. Indeed,

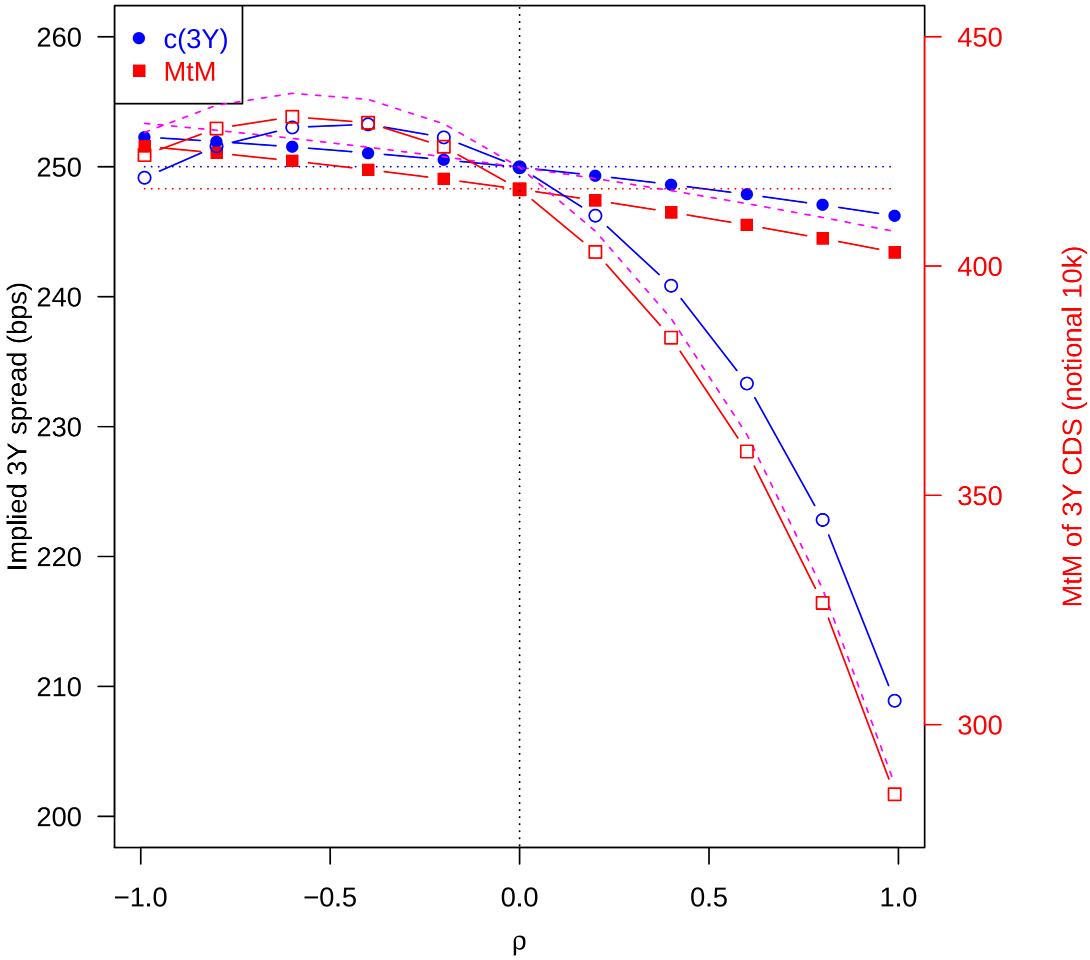

Figure 6: Impact of correlation on the implied break-even CDS spread with maturity 3Y (blue, dot markers, left axis) and corresponding MtM with 10k notional (red, square markers, right axis) for two different values of recovery rate volatility: $a = 1/10$ (filled markers) and $a = 1/2$ (empty markers). Parameters: $x = \mu_R = 40\%$, calibration instrument: 5Y maturity swap with $c = 250$ bps, $k = 100$ bps.

from the in-house model, the trader would not quote $d(5Y)$ but would quote a spread $\hat{d}(5Y)$ such that, when converting this spread according to the ISDA converter, she would find the same MtM (i.e. upfront amount) as the one implied by the in-house model when the actual (standard) coupon rate k is used. The spread $\hat{d}(5Y)$ is computed such that the "in-house model" price (correct to the trader) and "ISDA model" (conventional quotation) price agree:

$$(d(5Y) - k)\,\mathrm{PV01}(\lambda, T) = \mathrm{up} = (\hat{d}(5Y) - k)\,\mathrm{PV01}(h, T) \Rightarrow \hat{d}(5Y) = k + (d(5Y) - k)\frac{\mathrm{PV01}(\lambda, T)}{\mathrm{P\bar{V}01}(h, T)}\;.$$

Figure 8 summarizes the methodology and Figure 7(b) shows the impact. From the inputs given above, we find $\bar{d}(5Y) = c(5Y)/0.6 \approx 417$ bps (horizontal dotted blue line). By contrast, the break-even spread $d(5y)$ (blue) or its ISDA-equivalent $\hat{d}(5Y)$ (magenta) substantially depends on (a, ρ) but are relatively close from each other. Accounting for a positive dependency between ρ and λ leads to a digital CDS conventional spread larger than the one given by the simple rescaling of the conventional spread $c(5Y)/(1 - 0.4)$

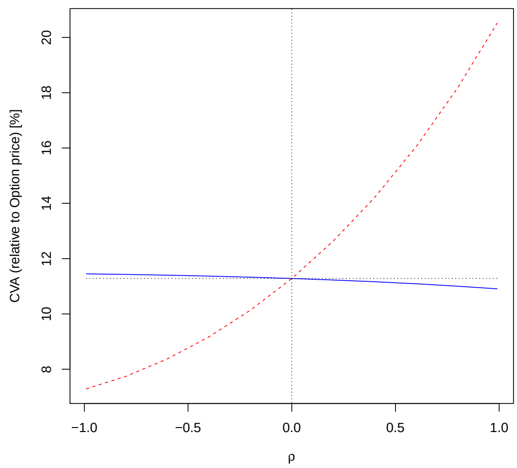

(a) The in-house model CVA ρ (blue), the "mixed-CVA" $\widetilde{\text{CVA}}$ (using G but $x = 40\%$ in the CVA payoff, red) and $\overline{\text{CVA}}$ obtained by assuming a fixed (here, $x = 40\%$) recovery rate everywhere (horizontal black). Parameters: $a = 0.5$, $S_0 = 50, K = 45, \sigma = 30\%$.

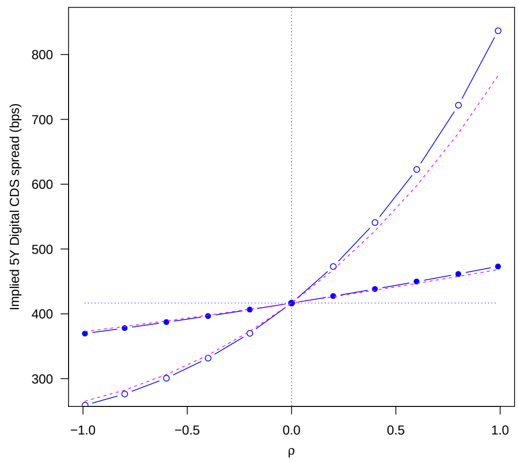

(b) Implied spread $\bar{d}$ (horizontal dots), d (blue) and $\hat{d}$ (magenta) of a digital CDS for two different values of recovery rate volatility: $a = 1/10$ (filled markers) and $a = 1/2$ (empty markers).

Figure 7: Impact of stochastic recovery rate on CVA (left) and implied par digital CDS spread (right) with respect to ρ. Calibration instrument: 5Y maturity swap with $c = 250$ bps and $k = 100$ bps.

suggested by the ISDA model.

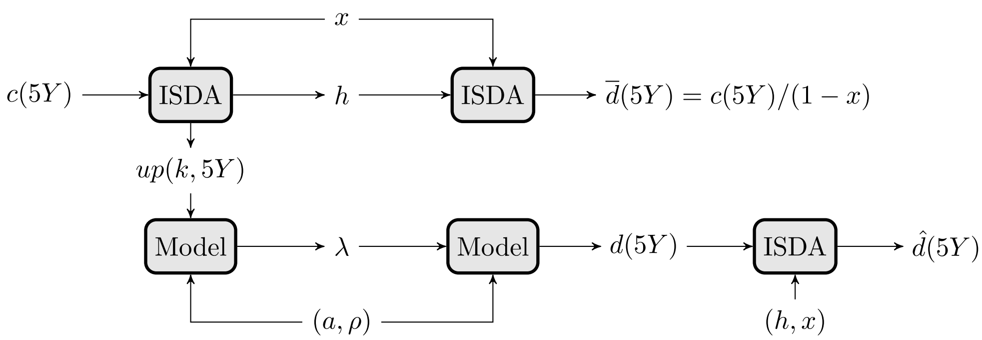

Figure 8: Methodology to analyze the model impact on the spread of a 5Y digital CDS starting from the spread of a 5Y standard CDS.

5 Conclusion

In this paper, we have stressed that the recovery rate uncertainty (which had been emphasized in many studies and again emphasized above) is a key driver of CDS prices. Moreover, quite a few empirical analyses suggest a negative dependency between recovery rate and default probability. Based on our simple in-house model, we have shown that the common way of extracting risk-neutral probabilities (i.e. inverting the equations derived from the ISDA CDS pricer, the standard pricer in the market) leads to *underestimate* the actual market-implied default probabilities in standard situations. This can have serious consequences when assessing credit risk of the firms. Moreover, it can also impact the valuation of less liquid instruments (like digital CDSs) or existing standard deals. This type of mispricing is of the same nature as those that led to introduce multi-curve pricing (*OIS discounting*): if one bootstraps the Libor curve from new trades, these trades will, by construction, price at par. By contrast, the impact of adopting the (correct) multi-curve approach will be clearly visible for swaps where the fixed rates substantially deviate from the prevailing swap rates. Therefore, just like Black-Scholes formula is a handy tool to communicate prices in terms of implied vols, the ISDA CDS model is a nice way to quote CDS prices as running premiums. Nevertheless, as pricer, it must be used with care as it neglects one of the major risks underlying such products. The figures produced with the simple model introduced above suggest that traders and risk managers may not make the economy of developing their own model when it comes to extract risk-neutral probabilities or price non standard deals. To the best of our knowledge, there is no standard alternative to the ISDA CDS pricer involving stochastic recovery rate. Developing such alternatives is the purpose of active research in the field.

6 Acknowledgments

This work was supported by the Fonds de la Recherche Scientifique - FNRS under Grant n J.0037.18.

References

F. Fabozzi, editor. *The Handbook of Fixed Income Securities*. Mc Graw-Hill (eighth ed.), 2003.

Markit. The CDS Big-Bang: understanding the changes to the Global CDS Contract and North American conventions. Technical report, March 13, 2009. http://www.markit.com/cds/announcements/resource/cds_big_bang.pdf.

Markit. ISDA CDS Standard Model. Technical report, 2004. http://www.cdsmodel.com/cdsmodel/.

E. Altman and V. Kishore. Almost everything you wanted to know about recoveries on defaulted bonds. *Financial Analysts Journal*, Nov.-Dec.:57–64, 1996.

A. Shleifer and R. Vishny. Liquidation values and debt capacity: A market equilibrium approach. *Journal of Finance*, 47(4):1343–1366, 1992. Reprinted in Michael J. Brennan, ed., The Theory of Corporate Finance, Edward Elgar Publishing Company, 1996.

V. Acharya, S. Bharath, and A. Srinivasan. Does industry-wide distress affect defaulted firms? evidence from creditor recoveries. *Journal of Financial Economics*, 85(3):787–821, 2007.

S. Sarbu, C. Schmitt, and M. Uhrig-Homburg. Market expectations of recovery rates. *Working Paper*, pages 1–42, 2013.

J. Frye. Collateral damage. *Risk*, page 9194, April 2000.

Y.-T. Hu and W. Perraudin. The dependence of recovery rates and defaults. Technical report, 2002.

E. Altman, B. Brady, A. Resti, and A. Sironi. The link between default and recovery rates: Theory, empirical evidence, and implications. *Journal of Business*, 78(6):22032227, 2005.

M. Boudreault, G. Gauthier, and T. Thomassin. Recovery rate risk and credit spreads in a hybrid credit risk model. *Journal of Credit Risk*, 9(3):3–39, 2013.

D. Hamilton, G. Gupton, and A. Berthault. Default and recovery rates of corporate bond issuers: 2000. Technical Report February, 2001.

R. Cantor and P. Varma. Determinants of recovery rates on defaulted bonds and loans for north american corporate issuers: 1983-2003. Technical Report December, 2004.

D. Keisman and K. Van de Castle. Recovering your money: insights into losses from defaults. Technical Report June, 1999.

K. Van de Castle and D. Keisman. Suddenly structure mattered: Insights into recoveries of defaulted debt. Technical Report May, 2000.

F. Fabozzi, H. Davis, and M. Choudhry. Credit-linked notes: a product primer. *Journal of Structured Finance*, 12(4):6777, 2007.

D. Brigo and A. Alfonsi. Credit default swaps calibration and option pricing with the SSRD stochastic intensity and interest rate model. *Finance and Stochastics*, 9:29–42, 2005.

R. Jarrow. Default parameter estimation using market prices. *Financial Analysts Journal*, 57(5):75–92, 2001.

E. Jokivuolle and S. Peura. A model for estimating recovery rates and collateral haircuts for bank loans. Technical report, Bank of Finland, 2000.

M. Pykthin. Unexpected recovery risk. *Risk*, pages 74–78, August 2003.

L. Andersen and J. Sidenius. Extensions to the gaussian copula: random recovery and random factor loadings. *Journal of Credit Risk*, 1(1):29–70, 2004.

D. Li. On default correlation: a copula function approach. Technical report, 2016.

R. Gaspar and I. Slinko. On recovery and intensity's correlation - a new class of credit models. *Journal of Credit Risk*, 4(2), 2008.

C. Ech-Chatbi. CDS and CDO pricing with stochastic recovery. Technical report, NA, 2008.

M. Krekel. Pricing distressed cdos with base correlation and stochastic recoveries. Technical report, 2010.

S. Amraoui, L. Cousot, S. Hitier, and J.-P. Laurent. Pricing CDOs with state-dependent stochastic recovery rates. *Quantitative Finance*, 12(8):1219–1240, 2012.

M. Boudreault, G. Gauthier, and T. Thomassin. Estimation of correlations in portfolio credit risk models based on noisy security prices. *Journal of Economic Dynamics and Control*, 61:334–349, 2015. URL https://doi.org/10.1016/j.jedc.2015.10.001.

J.-F. Bégin, M. Boudreault, and G. Gauthier. Firm-specific credit risk estimation in the presence of regimes and noisy prices. *Accepted in Finance Research Letter*, 2017.

J. Gregory. *Counterparty Credit Risk*. Wiley Finance, 2010.

A. Sokol. *A Practical Guide to Monte-Carlo CVA*. Risks Books, 2011.

F. Vrins. Wrong-way risk cva models with analytical epe profiles under gaussian exposure dynamics. *To appear in International Journal of Theoretical and Applied Finance*, 2016.

T. Schuermann. What do we know about loss given default? *Credit Risk Models and Management*, 2 (1), 2004.

D. Brigo and F. Vrins. CVA wrong way risk via change of measure: Actuarial vs risk-neutral pricing. In *Innovations in Insurance, Risk- and Asset-Management*, Munich (Germany), April 2017.