

OpenGamma Quantitative Research

**The Pricing and Risk
Management of Credit
Default Swaps, with a Focus
on the ISDA Model**

Richard White
Richard@opengamma.com

Abstract

In the paper we detail the reduced form or hazard rate method of pricing credit default swaps, which is a market standard. We then show exactly how the ISDA standard CDS model works, and how it can be independently implemented. Particular attention is paid to the accrual on default formula: We show that the original formula in the standard model is slightly wrong, but more importantly the proposed fix by Markit is also incorrect and gives a larger error than the original formula.

We finish by discussing the common risk factors used by CDS traders, and how these numbers can be calculated analytically from the ISDA model.

Contents

1	Introduction	1
2	The Standard CDS Contract	1
2.1	Dates	2
2.2	The Coupons	3
2.3	Pricing	3
2.4	Differences before and after ‘Big Bang’	4
3	The Basic Tools	4
3.1	Day Count Conventions	4
3.1.1	ACT/365 and ACT/365 Fixed	5
3.2	Interest Rate Curves	5
3.3	Credit Curve	6
4	Pricing a CDS	7
4.1	The Protection leg	7
4.2	The Premium Leg	8
4.3	CDS Valuation	9
4.3.1	Points Upfront	9
4.3.2	Market Value	9
4.3.3	The sign of the Accrued Premium	9
4.4	Approximation to the PV	10
4.5	The Par Spread	10
4.5.1	Par Spread Approximation	11
4.6	The Credit Triangle	11
4.6.1	PUF approximation	11
5	The ISDA Model	12
5.1	The ISDA Model Curve	12
5.2	Combined Nodes	13
5.3	Protection Leg	13
5.3.1	Limiting Case	13
5.4	Premium Leg	14
5.4.1	Formulae given in ISDA Standard CDS Model C Code	15
5.4.2	Limit Case	16
6	Calibration of the Credit Curve	16
6.1	The Single CDS Case	16
6.1.1	Conversion between PUF and Quoted Spreads	17
6.2	Multiple CDSs	18
6.2.1	Calibration from Quoted Spreads	19
6.2.2	Conversion between PUF and Par Spreads	19
7	Sensitivity to the Credit Curve	19
7.1	Hedging the Credit Curve	21
7.1.1	Hedging the ISDA Model Credit Curve	22
7.2	Rebalancing a Portfolio	23

8	Spread Sensitivity or Credit DV01	25
8.1	Par Spread Sensitivity	25
8.1.1	Hedging against Par Spread Moves	26
8.1.2	Rebalancing a Portfolio using Spread Sensitivities	27
8.2	Quoted Spread Sensitivity	28
8.2.1	Parallel CS01	28
8.2.2	Bucketed CS01	28
8.3	Sensitivity to Arbitrary Spreads	31
8.4	Credit Gamma or Convexity	31
9	Yield Curve Sensitivities	32
9.1	Hedging Yield Curve Movements	32
9.2	Money Market and Swap Sensitivities	32
10	Other Risk Factors	33
10.1	Recovery Rate Risk	33
10.2	Value on Default or Jump to Default	33
11	Beyond the ISDA Model	34
11.1	The Discount Curve	34
A	ISDA Model Dates	34
A.1	The Premium Leg	35
A.1.1	The Standard CDS Contract	36
A.1.2	Protection at Start of Day	37
A.2	The Protection Leg	37
B	Curve Sensitivity for the ISDA Model	37
B.1	Protection Leg Sensitivity	37
B.2	Premium Leg Sensitivity	38
B.2.1	Premiums Only	38
B.2.2	Accrual Paid On Default	38
C	Yield Curve Sensitivity for the ISDA Model	39
C.1	Protection Leg Sensitivity	39
C.2	Premium Leg Sensitivity	39
C.2.1	Premiums Only	39
C.2.2	Accrual Paid On Default	39
D	The ISDA Model Yield Curve Bootstrap	40
D.1	Swap Pricing	40
E	Market Data Used in Examples	41

1 Introduction

A Credit Default Swap (CDS) is a form of insurance against the default of a debt issuing entity.¹ This can be a corporation, a municipality or sovereign state. The protection lasts for a specified period (e.g. five years), and if the reference entity defaults in this period, the protection buyer receives a payment from the protection seller. In return, the buyer of protection makes regular (e.g. every three months) premium payments to the protection seller. These payments cancel in the event of a default, otherwise the contract cannot be cancelled before maturity, but it can be sold or unwound (at cost).

2 The Standard CDS Contract

Here we describe the new (post ‘Big Bang’) CDS contract. These are often referred to as vanilla CDS, standard CDS, Standard North American Contract (SNAC) or Standard European Contract (STEC). The differences from old, or legacy contracts are given at the end of the section.

What constitutes a default is contract specific and legally very technical. We will use the term *default* to mean any credit event that will trigger the payment of the protection leg of the CDS. More details can be found in The ISDA Credit Derivatives Definitions [ISD], or in more digestible form here [O’k08]. What follows is a description of a CDS contract suitable to build a mathematical pricing model, rather than a rigorous legal description.

A CDS contract is made between two parties - the protection buyer and the protection seller - and references a particular obligor (i.e. a corporation, municipality or sovereign state). The buyer of protection makes regular (quarterly) premium payments to the protection seller, until the expiry of the CDS contract or the reference obligor suffers a default (if this occurs first). For a particular period, there is an accrual start date, s_i , an accrual end date, e_i , and a payment date, p_i . The amount paid on the payment date is

$$\text{Notional} \times DCC(s_i, e_i) \times C$$

where C is the fixed coupon amount, DCC is the day count convention² used for premium payments. The commonly quoted currencies are USD, EUR and JPY and the notional is usually quite large - \$1MM to \$10MM (or equivalent in other currencies). In most CDS contracts (including SNAC), on default the protection buyer must pay the accrued premium from the last accrual date to the default date, τ , i.e. an amount $\text{Notional} \times DCC(s_i, \tau) \times C$, where $s_i < \tau < e_i$.

In the event of a default, the protection buyer would deliver one of the reference obligor’s defaulted bonds³ to the protection seller in return for par (physical settlement). In the event that there is not a ready supply of defaulted bonds in the market, an auction is conducted to establish the bond’s recovery rate and the contract is cash settled. The details of this are beyond the scope of this paper (again see [ISD] or [O’k08]). It is enough for us to assume that the protection seller pays the protection buyer an amount of $\text{Notional} \times (1 - RR)$, where RR is the established recovery rate of the obligor’s bonds. A less liquid type of CDS contract is a

¹The entity issues debt in the form of bonds, which it may default on. Unlike normal insurance, there is no obligation to hold the insured asset (i.e. the bonds of the reference entity), so CDS can be used to speculate on default or credit downgrades.

²Except for bespoke trades this is ACT/360 - see section 3.1 for a short description of day count conventions.

³Generally any maturity of bond is acceptable, so the protection buy will deliver the cheapest bond available.

Digital Default Swap (DDS) - this contractually specifies a recovery rate and cash settles. From a modelling perspective, it can be treated as an ordinary CDS.

Regardless of the trade date, the first coupon is paid in full. If the trade date is not an accrual start date, the protection buyer would receive from the protection seller the accrued premium to that point as a rebate.

2.1 Dates

There are several dates and date sets are defined in CDS contracts. These are:

- *Trade Date*. The date when the trade is executed. This is denoted as T , with $T+n$ meaning n days after the trade date.
- *Step-in or Protection Effective Date*. This is usually $T+1$. This is the date from which the issuer is deemed to be risky. Note, this is sometimes just called the *Effective Date*, however this can cause confusion with the legal effective date which is $T-60$ or $T-90$.
- *Valuation Date*. This is the date that future expected cash flows are discounted to. If this is the trade date, then we will report the *market value*, while if it is the *Cash-settle Date* then we will report the *cash settlement*.
- *Cash-settle Date*. This is the date that any upfront fees are paid. It is usually three working dates after the trade date.
- *Start or Accrual Begin Date*. This is when the CDS nominally starts in terms of premium payments, i.e. the number of days in the first period (and thus the amount of the first premium payment) is counted from this date. It is also known as the *prior coupon date* and is the previous IMM date before the trade date.
- *End or Maturity Date*. This is when the contract expires and protection ends - any default after this date does not trigger a payment. This is an (unadjusted) IMM date.
- The *Payment Dates*. These are the dates when premium payments are made. They are IMM dates adjusted to the next good business day.
- The *Accrual Start and End Dates*. These are dates used to calculate the exact value of each premium payment. Apart from the final accrual end date, these are adjusted⁴ IMM dates.

The ISDA model takes the trade, step-in, cash-settle, start and end dates as inputs and calculates payment and accrual start/end dates based on a set of rules. All inputs to the ISDA model and date generation rules are discussed in appendix A.

Prior to 2003, there were no fixed maturity dates, so a six-month (6M) CDS would have a maturity exactly six months after the step-in date [O'k08]. Since then, maturity dates for standard contracts have been fixed to the IMM dates of the 20th of March, June, September and December. The meaning of a 6M CDS is six months after the next IMM date from the trade date. Table 2.1 below gives some examples.

⁴using the *following* business day adjustment convention.

Trade date	6M	1Y
18-Jun-2013	20-Dec-2013	20-June-2014
19-Jun-2013	20-Dec-2013	20-June-2014
20-Jun-2013	20-March-2014	20-Sep-2014
21-Jun-2013	20-March-2014	20-Sep-2014

Table 1: Six month and one year CDS maturity dates for a set of trade dates crossing an IMM date.

2.2 The Coupons

Like the Interest Rate Swaps (IRS) on which they were modelled, CDSs originally traded at par - i.e. they were constructed to have zero cost of entry. The market view on the credit quality of the reference obligor was reflected in the coupon or *par spread*. Since the coupon was then fixed, the CDS could have positive or negative mark-to-market (MtM) throughout its lifetime, depending on the market's updated view of the credit quality.

Highly distressed issuers will either default in the near-term or recover and have a much smaller chance of default over the medium-term. The par-spread of a CDSs on non-investment grade obligors would typically be very high (sometimes exceeding 10,000bps) to cover the seller of protection from the (high) chance of a quick default. However, this does not suit the buy of protection, since they may be left paying huge premiums on an obligor who's credit quality has improved markedly. To better reflect the risk profile of these type of CDS, they generally traded with an upfront fee (paid by the protection buyer) and a standard (much lower) coupon.

Following the credit crisis, in 2009 ISDA issued the 'Big Bang' protocol in an attempt to restart the market by standardising CDS contracts [Roz09]. All CDS contracts would now have a standard coupon⁵ and an up-front charge, quoted as a percentage of the notional - *Points Up-Front* (PUF), which reflected the market's view of the reference obligor's credit quality. This amount is quoted as if it is paid by the protection buyer. However, it can be negative, in which case it is paid to the protection buyer.

The standardised coupons, together with maturity only on IMM dates, makes CDS contracts extremely fungible, which in turn makes market prices readily available. For example, a 1Y CDS on M&S issued on 20-Jun-2013 (and with a coupon of 100bps), will be identical to every other 1Y CDS on M&S (with 100bps coupon) issued up to the 19-Sep-2013, i.e. they will have identical premium payments and a maturity of 20-Sep-2014 - this is known as *on-the-run*. After the 19-Sep-2013, new 1Y CDSs will have a maturity of 20-Dec-2014, so the 'old' 1Y CDSs (which are now effectively 9M CDSs) will be less liquid - this is known as *off-the-run*.

2.3 Pricing

Since their inception, CDSs have been priced via a survival probability curve (and yield or discount curve for computing the PV of future cash flows). Section 3 describes the interest rate and credit curves and section 4 shows how to price a CDS given these curves.

In practice on-the-run CDSs for standard maturities will be highly liquid, and therefore have market quoted prices. These quotes (on the same reference entity) are used to construct a credit curve (see section 6), which in turn can be used to price off-the-run CDSs, which are less liquid.

⁵100 or 500bps in North America, and 25, 100, 500 or 1000bps in Europe.

2.4 Differences before and after ‘Big Bang’

Below we list the main differences between standard CDSs issued before and after the ISDA ‘Big Bang’ of April 2009, that affect the pricing model. Other differences mainly concern the determination of credit events and the recovery rate. This is taken from [Mar09].

- **Legal Effective Date.** This was T+1 (i.e. the same as the step-in date), which could cause risk to offsetting trades (as there may be a delay in credit event becoming known). It is now T-60 (credit events) or T-90 (succession events).
- **Accrued Interest.** For legacy trades only accrued interest from the step-in date (T+1). If the trade date was more than 30 days before the first coupon date, the first coupon is reduced to just be the accrued interest over the shortened period (*short stub*); if there are less than 30 days, nothing is paid on the (nominal) first coupon and on the second coupon (effectively the first) the full accrued over the extended period (*long stub*) is paid - i.e. the normal premium for this period plus the portion not paid on the first coupon date. Following this *front stub* period, normal coupons are paid.

For standard CDS, interest is accrued from the previous IMM date (the prior coupon date), so all coupons are paid in full. This accrued interest (from the prior coupon date to the step-in date) is rebated to the protection buyer, on the cash settlement date.

- **coupons.** Previously CDSs were issued with zero cost of entry, so the coupon was such that the fair value of the CDS was zero (par spread). Now CDSs are issued with standard coupons, and an upfront fee is paid.

3 The Basic Tools

3.1 Day Count Conventions

The start and end of accrual periods, payment dates and default time are all calendar dates. Cash-flows, such as premium payments, are calculated by converting a date interval into a real number representing a fraction of a year. This is achieved by determining the number of days between two events, and dividing by a denominator - the details of these Day Count Conventions (DCC) are beyond the scope of this paper; some of the more widely used ones are detailed here [Res12]. The simplest are of the form $\text{ACT}/\text{FixedNumber}$ - where ACT is the actual number of days between two events and the fixed denominator is 360 or 365 (or even 365.25). For our purposes we will denote the year-fraction, Δt , between two dates, d_1 and d_2 as

$$\Delta t = \text{DCC}(d_1, d_2) \tag{1}$$

The DCC used to calculate the premium payments for single-name CDS is almost always $\text{ACT}/360$.

When it comes to defining continuous time discount and survival curves, we need to convert intervals between the trade date (i.e. ‘now’) and payment times (or default times) to year fractions. In principle we can choose any DCC; however, generally the year fractions are not additive, so for three dates $d_1 < d_2 < d_3$ we have

$$\text{DCC}(d_1, d_2) + \text{DCC}(d_2, d_3) \neq \text{DCC}(d_1, d_3).$$

This is not a desirable feature for a continuous time curve. Actual over fixed denominator DCCs do not have this problem. Therefore we prefer to use ACT/365 (fixed) for the curves.⁶

3.1.1 ACT/365 and ACT/365 Fixed

What we have called ACT/365 above is often referred to as ACT/365 Fixed or ACT/365F. Conversely within the standard ISDA model, ACT/365 means ACT/ACT ISDA - this is used for accrual fractions in some currencies.⁷

The rest of this paper involves calculations which use real numbers. It is understood that all dates have been converted to year fractions from the trade date (using ACT/365F), and all payments (fixed) have been calculated using the correct DCC (usually ACT/360). Hence we will sometimes refer to a date, when strictly we mean a year fraction.

3.2 Interest Rate Curves

In the interest rate world, the price of a zero-coupon bond⁸ plays a critical role. The price at some time, t , for expiry at $T \geq t$, is written as $P(t, T)$, where $P(t, T) \leq 1.0$ and $P(T, T) = 1.0$. This quantity is used to discount future cash flows, so is known as the *discount factor*.

If the instantaneous short (interest) rate, $r(t)$ is deterministic, we may write

$$P(t, T) = e^{-\int_t^T r(s)ds} \quad (2)$$

If $r(t)$ follows some stochastic process, the price of a zero coupon bond may be written as

$$P(t, T) = \mathbb{E}^{\mathbb{P}} \left[e^{-\int_t^T r(s)ds} | \mathcal{F}_t \right] \quad (3)$$

where the expectation is taken in the risk neutral measure. We may also define the instantaneous forward rates, $f(t, s)$ $s \geq t$ and $f(t, t) = r(t)$, such that

$$P(t, T) = e^{-\int_t^T f(t, s)ds} \iff f(t, s) = -\frac{\partial \ln[P(t, s)]}{\partial s} = -\frac{1}{P(t, s)} \frac{\partial P(t, s)}{\partial s} \quad (4)$$

Finally, we define the continuously compounded yield on the zero coupon bond (or zero rate), $R(t, T)$, as

$$P(t, T) = e^{-(T-t)R(t, T)} \iff R(t, T) = -\frac{1}{T-t} \ln[P(t, T)] \quad (5)$$

When $t = 0$ we may drop the first argument and simply write $P(T)$, $f(T)$ and $R(T)$ for the discount factor, forward rate and zero rate to time T .

The graphs of $P(T)$ and $R(T)$ against T for $0 \leq T \leq T^*$ (where T^* is some upper cutoff) are known as the discount and yield curves respectively. Since they are linked by a simply monotonic transform, we can treat these curves as interchangeable,⁹ and refer to them generically as yield curves.

⁶This is the default in the ISDA model. Note, ACT/ACT which is also popular is not additive.

⁷Specifically for some Libor curves.

⁸This is largely a hypothetical bond which pays no coupons and returns one unit of currency at expiry - its fair value before expiry must be less than or equal to one.

⁹The link to the forward curve is more complex as it involves differentiation of the yield/discount curve to obtain the forward curve, or integration of the forward curve to obtain the yield/discount curve.

Each currency will have its own collection of curves, which are implied from various instruments in the market. ‘Risk’ free curves in the US are built from Overnight Index Swaps (OIS). The ISDA standard model uses the Libor curve for discounting, with a tenor (i.e. 3M, 6M etc) that is currency specific.

The ISDA model uses a simple bootstrap approach to construct a Libor curve with piecewise constant (instantaneous) forward rates from market quotes for money market and swap rates. This is discussed further in appendix D.

3.3 Credit Curve

In the reduced form used for CDS pricing, it is assumed that default is a Poisson process, with an intensity (or hazard rate) $\lambda(t)$. If the default time is τ , then the probability of default over an infinitesimal time period dt , given no default to time t is

$$\mathbb{P}(t < \tau \leq t + dt | \tau > t) = \lambda(t)dt \quad (6)$$

where $\mathbb{P}(A|B)$ denotes a conditional probability of A given B . The probability of surviving to at least time, $T > t$, (assuming no default has occurred up to time t) is given by

$$Q(t, T) = \mathbb{P}(\tau > T | \tau > t) = \mathbb{E}(\mathbb{I}_{\tau > T} | \mathcal{F}_t) = e^{-\int_t^T \lambda(s)ds} \quad (7)$$

where $\mathbb{I}_{\tau > T}$ is the indicator function which is 1 if $\tau > T$ and 0 otherwise. It should be noted that the intensity here does not represent a real-world probability of default.

Up until this point we have assumed that the intensity is deterministic¹⁰ - if it is extended to be a stochastic process, then the survival probability is given by

$$Q(t, T) = \mathbb{E}^{\mathbb{P}} \left[e^{-\int_t^T \lambda(s)ds} | \mathcal{F}_t \right] \quad (8)$$

It is quite clear that the survival probability $Q(t, T)$ is playing the same role as the discount factor $P(t, T)$, as is the intensity $\lambda(t)$ and the instantaneous short rate $r(t)$. We may extend this analogy and define the (forward) hazard rate, $h(t, T)$ as

$$Q(t, T) = e^{-\int_t^T h(t, s)ds} \iff h(t, s) = -\frac{\partial \ln[Q(t, s)]}{\partial s} = -\frac{1}{Q(t, s)} \frac{\partial Q(t, s)}{\partial s} \quad (9)$$

and the zero hazard rate,¹¹ $\Lambda(t, T)$ as

$$Q(t, T) = e^{-(T-t)\Lambda(t, T)} \iff \Lambda(t, T) = -\frac{1}{T-t} \ln[Q(t, T)] \quad (10)$$

The survival probability curve, $Q(t, T)$, the forward hazard rate curve, $h(t, T)$, and the zero hazard rate curve, $\Lambda(t, T)$, are equivalent, and we refer to them generically as credit curves.

The forward hazard rate represents the (infinitesimal) probability of a default between times T and $T + dT$, conditional on survival to time T , as seen from time $t < T$. The unconditional probability of default between times T and $T + dT$ is given by

$$\mathbb{P}(T < \tau \leq T + dT | \tau > t) = Q(t, T)h(t, T) = -\frac{\partial Q(t, T)}{\partial T} \quad (11)$$

¹⁰the default time is random even when the intensity is deterministic.

¹¹since it is the analogy of the zero rate. In a bond context, this has been called the ZZ-spread (zero-recovery, zero-coupon) [Ber11].

4 Pricing a CDS

The pricing mechanism we present below is quite standard and can be found (in various forms) in any textbook covering the subject [Hul06, O’k08, Cha10, LR11]. We assume that a continuous time interest rate and credit curve is extraneously given, and are defined from the trade date ($t = 0$) to at least the expiry of the CDS. All dates have been converted into year fractions (from the trade-date) using the curve DCC.¹²

Below we consider prices at the cash-settlement date rather than the trade date. The *cash settlement* value (or dirty price from bond lexicon) of a CDS is the discounted value of expected future cash flows, and ignores the accrued premium. What is normally quoted is the *upfront fee* or *cash amount* (the clean price), which is simply the dirty price with any accrued interest added. For a newly issued legacy CDS, there is no accrued interest (recall, interest accrues from $T+1$), so dirty and clean price are the same.

Clean and dirty price are used in the ISDA model, but are not standard terms in the CDS world. We chose to use the term to aid clarity: dirty means without accrued premium and clean means with accrued premium.

We take the trade date as $t = 0$ and the maturity as T . The step-in (effective protection date) is t_e and the valuation (cash-settle date) is t_v . All cash flows are discounted to the valuation date - this is what we call the present value (or PV). This pricing is valid for new contracts (both spot starting and forward starting) as well as seasoned trades (e.g. for MtM calculations), and for a broad set of contract specifications.

4.1 The Protection leg

The protection leg of a CDS consists of a (random) payment of $N(1 - RR(\tau))$ at default time τ if this is before the expiry (end of protection) of the CDS (time T) and nothing otherwise. The present value of this leg can be written as

$$PV_{\text{Protection Leg}} = N\mathbb{E}[e^{-\int_{t_v}^{\tau} r(s)ds}(1 - RR(\tau))\mathbb{I}_{\tau < T}] \quad (12)$$

Under the assumptions that recovery rates are independent of interest or hazard rates¹³, and independent of the default time, this can be rewritten as

$$PV_{\text{Protection Leg}} = N\mathbb{E}[1 - RR(\tau)]\mathbb{E}[e^{-\int_{t_v}^{\tau} r(s)ds}\mathbb{I}_{\tau < T}] = N(1 - RR)\mathbb{E}[e^{-\int_{t_v}^{\tau} r(s)ds}\mathbb{I}_{\tau < T}] \quad (13)$$

where $RR = \mathbb{E}[RR(\tau)]$ is the expected recovery rate.

Under the further assumption that interest rates and hazard rates are independent, this becomes

$$PV_{\text{Protection Leg}} = -\frac{N(1 - RR)}{P(t_v)} \int_0^T P(s) \frac{dQ(s)}{ds} ds = -\frac{N(1 - RR)}{P(t_v)} \int_0^T P(s) dQ(s) \quad (14)$$

With no other information about the nature of the yield and credit curves, the PV of the protection leg must be computed by numerical integration.

¹²We use ACT/365F for the discount and credit curves. The trade (in terms of calculation of premiums) uses ACT/360.

¹³A negative correlation between spreads and recovery rates has been observed[Cha10, Ber11], but this is generally ignored.

4.2 The Premium Leg

The premium leg consists of two parts: Regular premium (or coupon) payments (e.g. every three months) up to the expiry of the CDS, which cease if a default occurs; and a single payment of accrued premium in the event of a default (this is not included in all CDS contracts but is for SNAC).

If there are M remaining payments, with payment times $t_1, t_2, \dots, t_i, \dots, t_M$, period end times $e_1, e_2, \dots, e_i, \dots, e_M$ ¹⁴ and year fractions of $\Delta_1, \dots, \Delta_i, \dots, \Delta_M$, then the present value of the premiums only is

$$PV_{\text{Premiums only}} = NC \mathbb{E} \left[\sum_{i=1}^M \Delta_i e^{-\int_{t_v}^{t_i} r(s) ds} \mathbb{I}_{e_i < \tau} \right] = \frac{NC}{P(t_v)} \sum_{i=1}^M \Delta_i P(t_i) Q(e_i) \quad (15)$$

where the second equation follows from the independence assumption. The quantity $P(T)Q(T) \equiv B(T)$ is known as the risky discount factor - the PV of the premium payments is just the risky discounted value of the cash flows.

The second part of the premium leg is the accrued interest paid on default. If the accrual start and end times are $(s_1, e_1), \dots, (s_i, e_i), \dots, (s_M, e_M)$, its PV is given by

$$\begin{aligned} PV_{\text{accrued interest}} &= NC \mathbb{E} \left[\sum_{i=1}^M \text{DCC}(s_i, \tau) e^{-\int_{t_v}^{\tau} r(s) ds} \mathbb{I}_{s_i < \tau < e_i} \right] \\ &= -\frac{NC}{P(t_v)} \sum_{i=1}^M \left[\int_{s_i}^{e_i} \text{DCC}(s_i, t) P(t) \frac{dQ(t)}{dt} dt \right] \end{aligned} \quad (16)$$

We use ACT/365 to measure year fractions for the curves. The accrual fraction may use a different day count convention.¹⁵ To account for this we may write

$$PV_{\text{accrued interest}} = -\frac{NC}{P(t_v)} \sum_{i=1}^M \left[\eta_i \int_{s_i}^{e_i} (t - s_i) P(t) \frac{dQ(t)}{dt} dt \right] \quad (17)$$

where $\eta_i = \Delta_i / \text{DCC}_{\text{curve}}(s_i, e_i)$ - i.e. it is the ratio of the year fraction of the interval measured using the accrual DCC, to the interval measured using the curve DCC. When both DCC use actual days for the numerator and a fixed denominator, this will be a fixed number.¹⁶ This is now amenable to numerical integration.

The full PV of the premium leg is

$$\begin{aligned} PV_{\text{premium}} &= PV_{\text{Premiums only}} + PV_{\text{accrued interest}} \\ &= \frac{NC}{P(t_v)} \sum_{i=1}^M \left[\Delta_i P(t_i) Q(e_i) - \eta_i \int_{s_i}^{e_i} (t - s_i) P(t) \frac{dQ(t)}{dt} dt \right] \end{aligned} \quad (18)$$

This of course scales linearly with the coupon, C .

¹⁴The period end time is the final time that a default can occur to count as a default in that period. It usually equals the payment time, but can be before it; notably for the ISDA model.

¹⁵it is usually ACT/360

¹⁶In most cases it will be simply $365/360 \approx 1.0139$.

The *Risky PV01* (RPV01) (aka *spread PV01*, *risky duration* or *PVBP*) is usually defined as the value of the premium leg per basis point of spread (coupon). To avoid scaling factors, we define it as the *value of the premium leg per unit of coupon*,¹⁷ so:

$$PV_{\text{premium}} = C \times RPV01 \quad (19)$$

4.3 CDS Valuation

The premium leg we discussed above is the *dirty* PV. For clarity we label the Risky PV01 given above, $RPV01_{\text{dirty}}$. The cash settlement (or dirty PV) for the buyer of protection, is given by

$$PV_{\text{dirty}} = PV_{\text{Protection Leg}} - C \times RPV01_{\text{dirty}} \quad (20)$$

As with bonds, this value has a sawtooth pattern against time, driven by the discrete premium payments. To smooth out this pattern, the accrued interest between the accrued start date (immediately before the step-in date) and the step-in date is subtracted from the premium payments. So

$$RPV01_{\text{clean}} = (RPV01_{\text{dirty}} - N \times DCC(s_1, t_e)) \quad (21)$$

It is this clean RPV01 that is normally quoted, so when we refer to RPV01 without a subscript we mean the clean RPV01. The two PV can then be written as

$$\begin{aligned} PV_{\text{clean}} &= PV_{\text{Protection Leg}} - C \times RPV01 \\ PV_{\text{dirty}} &= PV_{\text{Protection Leg}} - C \times RPV01 - NC \times DCC(s_1, t_e) \end{aligned} \quad (22)$$

This PV_{clean} is the *upfront amount* - it is the net amount paid by the protection buyer, on the cash-settlement date to enter a new CDS contract. Recall the current (post ‘Big-Bang’) treatment of coupons: The protection buyer pays the next coupon in full on the coupon date (even if this is the next day); in return the buyer receives (from the protection seller) the accrued interest[Cha10], which is paid on the cash-settle date.

4.3.1 Points Upfront

As already mentioned, the upfront amount is just the PV_{clean} in our notation. Points Upfront (PUF) is just this value per unit of notional (expressed as a percentage). The *clean price* is defined as 1 - PUF (again expressed as a percentage). It is now market standard to quote CDSs in terms of PUF and a standard coupon.

4.3.2 Market Value

The market value is simply PV_{dirty} discounted to the trade date. That is, $P(t_v) = 1$ in the above equations. It will be slightly smaller in magnitude than the cash settlement.

4.3.3 The sign of the Accrued Premium

The quantity $NC \times DCC(s_1, t_e)$ is positive and is called the *accrued amount*. By market convention the sign of the accrued for the buyer of protection is negative, so accrued is defined as

$$\text{accrued} = -NC \times DCC(s_1, t_e) \quad (23)$$

For the seller of protection, the accrued is positive.

¹⁷i.e. our value is 10,000 times the ‘standard’ value.

4.4 Approximation to the PV

A common approximation is to assume that any payment resulting from a default is made on the next scheduled payment date [Ber11]. We can then approximate the protection leg as

$$PV_{\text{Protection Leg}} \approx \frac{N(1 - RR)}{P(t_v)} \sum_{i=1}^M P(t_i) [(Q(t_{i-1}) - Q(t_i))] \quad (24)$$

In addition, if we assume that on average default happens mid-way through the period, we can approximate the accrued interest paid on default as

$$PV_{\text{accrued interest}} \approx \frac{NC}{2P(t_v)} \sum_{i=1}^M \Delta_i P(t_i) (Q(t_{i-1}) - Q(t_i)) \quad (25)$$

and

$$\begin{aligned} RPV01_{\text{dirty}} &\approx \frac{N}{P(t_v)} \sum_{i=1}^M \Delta_i P(t_i) \left[Q(t_i) + \frac{(Q(t_{i-1}) - Q(t_i))}{2} \right] \\ &= \frac{N}{P(t_v)} \sum_{i=1}^M \Delta_i P(t_i) \left[\frac{(Q(t_{i-1}) + Q(t_i))}{2} \right] \end{aligned} \quad (26)$$

So the PV of a CDS can be calculated reasonably accurately without numerical integration. This in turn means one can use more complex representations for the curves, without making calibration slow. We discuss this further in section 11.

4.5 The Par Spread

Prior to the new ISDA ‘Big Bang’ protocol of 2009, CDS contracts were entered into at zero cost. The coupon that makes the *clean* PV of the CDS zero is known as the par spread, S_p . It is given simply as

$$S_p(T) = \frac{PV_{\text{Protection Leg}}}{RPV01_{\text{clean}}} = \frac{-(1 - RR) \int_0^T P(s) \frac{dQ(s)}{ds} ds}{\sum_{i=1}^M \left[\Delta_i P(t_i) Q(t_i) + \eta_i \int_{s_i}^{e_i} (t - s_i) P(t) \frac{dQ(t)}{dt} dt \right] - N \times DCC(s_1, t_e)} \quad (27)$$

If a CDS is issued with a coupon C (which may be the initial par spread), and some time later the par spread for a CDS with the same maturity¹⁸ is S_p , then the MtM value of the CDS is

$$MtM = (S_p - C) \times RPV01(P, Q) \quad (28)$$

The mark-to-market will be clean or dirty depending on which RPV01 we use. We have emphasised that the RPV01 depends on the discount and credit curves, which will have changed since the CDS was issued. Hence the equation is not as straightforward as it appears, since one must recalculate the value of RPV01.

¹⁸and of course with all other details except the coupon identical.

4.5.1 Par Spread Approximation

Using the approximation of section 4.4 we have

$$S_p(T) = \frac{(1 - RR) \sum_{i=1}^M P(t_i) [(Q(t_{i-1}) - Q(t_i))] }{\sum_{i=1}^M \Delta_i P(t_i) \left[\frac{(Q(t_{i-1}) + Q(t_i))}{2} \right]} \quad (29)$$

4.6 The Credit Triangle

If the premium leg were paid continuously, its value would be given by

$$PV_{\text{premium, continuous}} = \frac{NC}{P(t_v)} \int_0^T P(t) Q(t) dt \quad (30)$$

and the par spread (ignoring any accrued interest) would be given by

$$S_p = -(1 - RR) \frac{\int_0^T P(t) Q(t) dt}{\int_0^T P(t) dQ(t)} \quad (31)$$

Under the further assumption that the hazard rate¹⁹ is constant at a level, λ , this becomes the *credit triangle* relationship

$$S_p = (1 - RR) \lambda \frac{\int_0^T P(t) e^{-\lambda t} dt}{\int_0^T P(t) e^{-\lambda t} dt} = (1 - RR) \lambda \quad \leftrightarrow \quad \lambda = \frac{S_p}{1 - RR} \quad (32)$$

This serves as a *rule of thumb* for switching between spreads and hazard rates, and as a first guess at the hazard rate for calibration. Some subtleties, such as what recovery rate should be used, are discussed in [Ber11].

4.6.1 PUF approximation

For a fixed premium and PUF we have

$$[(1 - RR)\lambda - C] \int_0^T P(t) e^{-\lambda t} dt = PUF \quad (33)$$

To generate another rule of thumb, we first assume that interest rates are flat at a level, r , and that $(r + \lambda)T \ll 1$. The integral then becomes

$$\int_0^T e^{-(r+\lambda)t} dt = \frac{1 - e^{-(r+\lambda)T}}{r + \lambda} \approx T - \frac{(r + \lambda)T^2}{2} + \dots$$

Keeping just the first term, we may write

$$\lambda \approx \frac{C + PUF/T}{1 - RR} \quad (34)$$

¹⁹In this case, instantaneous, forward and zero hazard rates all have the same value, so we don't make a distinction.

5 The ISDA Model

The ISDA CDS Standard Model²⁰ is maintained by Markit. It is written in the C language and is the evolution of the JP Morgan CDS pricing routines. As well as the CDS date logic²¹ the model makes the assumption that both the yield curve and credit curve are piecewise constant in their respective forward rates. This reduces the integrals of equations 14 and 17 to sums of simple analytic parts, and hence eliminates the need for numerical integration in the pricing. The curve assumptions also greatly simplify calibration, which we discuss in section 6.

The assumption of piecewise constant forward rates between specified nodes *is the ISDA model* - the official C code is an implementation of this, but any implementation using the same curve assumptions will produce the same prices to machine precision.²² OpenGamma has produced its own implementation in Java.

One further point to note, is that the ISDA model is quite general about the contract specification. It can be used to price CDSs with any maturity date (it knows nothing about IMM dates), start date and payment interval. So the contract specifics are inputs to the model - for standard contracts, this would be a maturity date on the relevant IMM date, a start date of the IMM date immediately before the trade date, and quarterly premium payments.

5.1 The ISDA Model Curve

We assume that the yield curve has n_y nodes at times $\mathcal{T}^y = \{t_1^y, t_2^y, \dots, t_i^y, \dots, t_{n_y}^y\}$ and that the credit curve has n_c nodes at $\mathcal{T}^c = \{t_1^c, \dots, t_i^c, \dots, t_{n_c}^c\}$. At the i^{th} node of the yield curve the discount factor is given by $P_i = \exp(-t_i^y R_i)$, and at the i^{th} node of the credit curve the survival probability is given by $Q_i = \exp(-t_i^c \Lambda_i)$.

We now just consider the credit curve (the equivalent values for the yield curve follow by analogy). The survival probability at non-node times is given by

$$Q(t) = \exp\left(-\left[\frac{t_i^c \Lambda_i (t_{i+1}^c - t) + t_{i+1}^c \Lambda_{i+1} (t - t_i^c)}{\Delta t_i^c}\right]\right) \quad \text{for } t_i < t < t_{i+1} \quad (35)$$

where $\Delta t_i^c = t_{i+1}^c - t_i^c$. The interpolation is linear in the log of the survival probability. Using equation 9 we have

$$h(t) = \frac{-t_i^c \Lambda_i + t_{i+1}^c \Lambda_{i+1}}{\Delta t_i^c} \equiv h_{i+1} \quad \text{for } t_i < t < t_{i+1} \quad (36)$$

which justifies our previous statement that the forward rates are piecewise constant.²³ This allows us to write

$$\begin{aligned} Q(t) &= Q_i e^{-h_{i+1}(t-t_i^c)} \\ \Lambda(t) &= \frac{\Lambda_i t_i^c + h_{i+1}(t - t_i^c)}{t} \quad \text{for } t_i < t < t_{i+1} \end{aligned} \quad (37)$$

For times before the first node we have simply

$$Q(t) = e^{-h_1 t} \quad \text{for } 0 \leq t < t_1 \quad (38)$$

²⁰latest version 1.8.2 at the time of writing.

²¹This is not strictly part of the *model* as it is specified in the CDS contract.

²²This is 1 part in 10^{15} using double precision floating point operations.

²³Note: the forward rate is undefined at the nodes.

and after the last node

$$Q(t) = Q_{n_c} e^{-h_{n_c}(t-t_{n_c}^c)} \quad \text{for } t > t_{n_c}^c \quad (39)$$

That is, the forward hazard rate is constant after the penultimate node.

5.2 Combined Nodes

We let the chronologically ordered combined set of unique nodes²⁴ be

$$\mathcal{T} = \mathcal{T}^y \cup \mathcal{T}^c = \{t_1, \dots, t_i, \dots, t_n\}$$

where $n \leq n_y + n_c$. Between any two consecutive nodes we are guaranteed that both the forward interest rate and forward hazard rate are constant.

We distinguish values (e.g. discount factors, survival probabilities) at these nodes by $\bar{Q}_i \equiv Q(t_i)$, $\bar{P}_i \equiv P(t_i)$ etc.

5.3 Protection Leg

For the protection leg we modify the set of nodes so that final node, t_n , is taken as the maturity of the CDS (so nodes after this time are not included), and the *effective protection start* is added as the zeroth node, t_0 .

The value of the protection leg (equation 14) may be written as

$$PV_{\text{Protection Leg}} = \frac{N(1-RR)}{P(t_v)} \sum_{i=1}^n I_i \quad \text{where } I_i = - \int_{t_{i-1}}^{t_i} P(t) dQ(t) \quad (40)$$

The individual integral elements are given by

$$\begin{aligned} I_i &= - \int_{t_{i-1}}^{t_i} P(t) dQ(t) = \bar{P}_{i-1} \bar{Q}_{i-1} h_i \int_{t_{i-1}}^{t_i} e^{-(f_i+h_i)(t-t_{i-1})} dt \\ &= \frac{h_i}{f_i + h_i} (\bar{B}_{i-1} - \bar{B}_i) = \frac{\hat{h}_i}{\hat{f}_i + \hat{h}_i} (\bar{B}_{i-1} - \bar{B}_i) \end{aligned} \quad (41)$$

where $\bar{B}_i = \bar{P}_i \bar{Q}_i$ and we have defined

$$\begin{aligned} \hat{f}_i &\equiv f_i \Delta t_{i-1} = \ln(\bar{P}_{i-1}) - \ln(\bar{P}_i) \\ \hat{h}_i &\equiv h_i \Delta t_{i-1} = \ln(\bar{Q}_{i-1}) - \ln(\bar{Q}_i) \end{aligned} \quad (42)$$

- the point of this last step is to not divide by the time step since it cancels out anyway.

5.3.1 Limiting Case

Care must be taken when $\hat{f}_i + \hat{h}_i \rightarrow 0$. We start by defining the function $\epsilon(x)$, its derivatives and their Taylor expansions

$$\begin{aligned} \epsilon(x) &\equiv \frac{e^x - 1}{x} \approx 1 + \frac{1}{2}x + \frac{1}{6}x^2 + \frac{1}{24}x^3 + \dots \\ \epsilon'(x) &\equiv \frac{(x-1)(e^x - 1) + x}{x^2} \approx \frac{1}{2} + \frac{1}{3}x + \frac{1}{8}x^2 + \frac{1}{30}x^3 + \dots \\ \epsilon''(x) &\equiv \frac{(e^x - 1)(x^2 - 2x + 2) + x^2 - 2x}{x^3} \approx \frac{1}{3} + \frac{1}{4}x + \frac{1}{10}x^2 + \frac{1}{36}x^3 + \dots \end{aligned} \quad (43)$$

²⁴that is, if the yield and credit curve have a common node, it is only counted once.

With this, equation 41 can be written as

$$I_i = B_{i-1} \hat{h}_i \epsilon(-(\hat{f}_i + \hat{h}_i)) \quad (44)$$

For small values of $\hat{f}_i + \hat{h}_i$ we may use the Taylor expansion form of $\epsilon(x)$ which is well defined for $x \rightarrow 0$.

5.4 Premium Leg

The premium leg consists of M risky discounted premium payments and M integrals which represent the accrued interest paid on default. The premium payments pose no problem, since we may use equation 15 directly. Therefore we just focus on the integral for default in the k^{th} payment period. From equation 17 this is

$$I^k = - \int_{s_k}^{e_k} (t - s_k) P(t) \frac{dQ(t)}{dt} dt \quad (45)$$

where s_k and e_k and the accrual start and end for the k^{th} period.²⁵ If we now truncate our set of combined nodes, so it only contains the n_k nodes between s_k and e_k exclusively, then add s_k and e_k as the first and last node (i.e. $t_0 = s_k$ and $t_{n_k} = e_k$) we have

$$\begin{aligned} I^k &= - \sum_{i=1}^{n_k} \int_{t_{i-1}}^{t_i} (t - s_k) P(t) \frac{dQ(t)}{dt} dt \\ &= \sum_{i=1}^{n_k} \left[\bar{P}_{i-1} \bar{Q}_{i-1} h_i \int_{t_{i-1}}^{t_i} (t - s_k) e^{-(f_i + h_i)(t - t_{i-1})} dt \right] \\ &= \sum_{i=1}^{n_k} \left[\frac{h_i}{f_i + h_i} \left(\frac{\bar{B}_{i-1} - \bar{B}_i}{f_i + h_i} + \bar{B}_{i-1}(t_{i-1} - s_k) - \bar{B}_i(t_i - s_k) \right) \right] \\ &= \sum_{i=1}^{n_k} \left[\frac{\hat{h}_i}{\hat{f}_i + \hat{h}_i} \left(\Delta t_{i-1} \left(\frac{\bar{B}_{i-1} - \bar{B}_i}{\hat{f}_i + \hat{h}_i} - \bar{B}_i \right) + (t_{i-1} - s_k)(\bar{B}_{i-1} - \bar{B}_i) \right) \right] \end{aligned} \quad (46)$$

where we have made use of the identity $\bar{B}_i = \bar{B}_{i-1} e^{-(f_i + h_i)(t_i - t_{i-1})}$ in simplifying the expression. Since the nodes are generally further apart than the three months between payments, we will usually have $n_k = 1$ and the integral for each period reduces to a single simple formula

$$I^k = \frac{\Delta t_0 \hat{h}_1}{\hat{f}_1 + \hat{h}_1} \left(\frac{\bar{B}_0 - \bar{B}_1}{\hat{f}_1 + \hat{h}_1} - \bar{B}_1 \right) \quad (47)$$

Putting this all together we have

$$PV_{\text{premium}} = \frac{NC}{P(t_v)} \sum_{k=1}^M [\Delta_k P(t_k) Q(e_k) - \eta_i I^k] \quad (48)$$

for the value of the premium leg when the forward rates (yield curve) and forward hazard rates (credit curve) are piecewise constant, i.e. the ISDA model.

²⁵In the ISDA model, if protection is from start of day (i.e standard CDS) these are both rolled back one day. The exception to this is the maturity date, so the overall effect is to give one extra day of protection.

5.4.1 Formulae given in ISDA Standard CDS Model C Code

This implementation (versions 1.8.2 and below) gives this formula as (using our notation and rearranging)

$$I_{\text{ISDA}}^k = \sum_{i=1}^{n_k} \left[\frac{\hat{h}_i}{\hat{f}_i + \hat{h}_i} \left(\Delta t_{i-1} \left(\frac{\bar{B}_{i-1} - \bar{B}_i}{\hat{f}_i + \hat{h}_i} - \bar{B}_i \right) + (t_{i-1} - \tilde{s}_k)(\bar{B}_{i-1} - \bar{B}_i) \right) \right] \quad (49)$$

where $\tilde{s}_k = s_k - 1/730$ - i.e. the accrual start is reduced by half a day (on a 365-based day count). The reason for this is unclear²⁶ and the formula is identical to equation 46 except for this half-day adjustment to the accrual start. In fact we may write

$$I_{\text{ISDA}}^k = I^k + \frac{1}{730} \sum_{i=1}^{n_k} \frac{\hat{h}_i}{\hat{f}_i + \hat{h}_i} (\bar{B}_{i-1} - \bar{B}_i) \quad (50)$$

with the second term seen as an error term.

In a note in December 2012 [Mar12] Markit states that the formula (eqn. 49) is incorrect. They give a correction (for the $n_k = 1$ case) in the paper, and also added it as a comment to the 1.8.2 C code. The corrected equation is (in our notation)

$$I_{\text{Markit fix}}^k = \sum_{i=1}^{n_k} \left[\frac{\Delta t_{i-1} \hat{h}_i}{\hat{f}_i + \hat{h}_i} \left(\frac{\bar{B}_{i-1} - \bar{B}_i}{\hat{f}_i + \hat{h}_i} - \bar{B}_i \right) \right] \quad (51)$$

This is the same as our equation 47 for $n_k = 1$, i.e. if the forward rate and forward hazard rate are both constant throughout the coupon period (i.e. there are no nodes within the coupon period). However, this is not true in general, so the correction is still wrong. We may write

$$I_{\text{Markit fix}}^k = I^k - \sum_{i=1}^{n_k} \left[\frac{\hat{h}_i}{\hat{f}_i + \hat{h}_i} ((t_{i-1} - s_k)(\bar{B}_{i-1} - \bar{B}_i)) \right] \quad (52)$$

where again the second part can be seen as an error term. Clearly for $n_k = 1$ the error term vanishes (since $t_0 = s_k$), but for $n_k > 1$ the error is larger than for the original ISDA model formula.

Both Markit and Bloomberg still use the code with equation 49 and are yet to switch to equation 51.²⁷ Nevertheless some groups have started using the Markit fix, even though it is wrong in some circumstances.

The correct formula is that given in equation 46, which is the original formula in the Standard ISDA model without the half day adjustment. However, to be pragmatic OpenGamma has implemented all three formulae given above. The price of five CDS (tenors of 3M, 6M, 1Y, 5Y and 10Y) with a notional of 10MM using all three formulae is given in table 2, along with the percentage difference from what we believe to be the correct formula (labeled OpenGamma fix). All the CDSs have been priced from the same credit curve.²⁸ The original ISDA formula is out by less than 0.01%, while the Markit fix is out by more than 0.6% for the short dated CDS - this is because the formula mis-prices periods that have a curve node fall within then, and these tend to be the early premium periods; the effect is diluted for long dated CDS.

²⁶An amount of 10^{-50} is also added to the denominators to prevent divide by zero problems, rather than taking the limits correctly (e.g. by a Taylor expansion). We can ignore this.

²⁷presumably for backwards compatibility as the change in price is non-negligible.

²⁸If the ‘wrong’ formula is used to calibrate the curve and the same (wrong) formula is the used to price CDS, the cancelation of errors will lessen the effect.

	3M	6M	1Y	5Y	10Y
OpenGamma Fix	-5,698.56	-10,868.26	-21,118.52	317,678.48	837,013.78
ISDA Original	-5,699.06	-10,869.22	-21,120.39	317,660.59	836,976.56
Markit Fix	-5,662.91	-10,829.15	-21,071.65	317,761.01	837,135.37
Error: ISDA Original	0.009%	0.009%	0.009%	-0.006%	-0.004%
Error: Markit Fix	-0.626%	-0.360%	-0.222%	0.026%	0.015%

Table 2: Comparison of the effect of the three different accrual-on-default formulae discussed in the text, on the price (for 10MM notional) of five CDSs. The error is given versus the OpenGamma fix.

5.4.2 Limit Case

We must again consider the case of $\hat{f}_i + \hat{h}_i \rightarrow 0$. Using $\epsilon(x)$ and $\epsilon'(x)$ (eqn. 43) we may write

$$I^k = \sum_{i=1}^n \left[\hat{h}_i \bar{B}_{i-1} \left((t_{i-1} - s_k) \epsilon(-(\hat{f}_i + \hat{h}_i)) + \Delta t_{i-1} \epsilon'(-(\hat{f}_i + \hat{h}_i)) \right) \right] \quad (53)$$

and once again small values of $\hat{f}_i + \hat{h}_i$ can be handled using the Taylor expansion form of $\epsilon(x)$ and $\epsilon'(x)$. The Markit fix formula (eqn. 51) can also be handled this way, giving

$$I_{\text{Markit fix}}^k = \sum_{i=1}^n \left[\Delta t_{i-1} \hat{h}_i \bar{B}_{i-1} \epsilon'(-(\hat{f}_i + \hat{h}_i)) \right] \quad (54)$$

6 Calibration of the Credit Curve

In this section we will concentrate on the calibration of the ISDA model. In section 11 we will discuss calibration using other representations of the credit curve.

The ISDA CDS Standard Model calibrates a yield curve from money market instruments (deposits) and swaps. We discuss the details of this in appendix D, but for now we assume an appropriate yield curve exists.

6.1 The Single CDS Case

If for a particular reference entity, there is only one CDS tenor available with reliable pricing information (i.e is liquid), the best we can do is to find the single constant hazard rate, λ , that will reprice the CDS. For a fixed coupon, C , and recovery rate, the price is a (non-linear) function of λ . This is shown in figure 1 below for a range of recovery rates. The function is monotonically increasing, starting at the (negative) discounted value of the premium payments when $\lambda = 0$ (since there is no value in the protection leg) and asymptotically reaching the Loss Given Default ($\text{LGD} = 1 - \text{RR}$) as $\lambda \rightarrow \infty$ - here default is immediate.

For (market observed) clean price, PV_{clean} we can define the price residual functions as

$$G(\lambda) = (PV_{\text{Protection Leg}}(\lambda) - C \times RPV01_{\text{clean}}(\lambda) - PV_{\text{clean}})/N \quad (55)$$

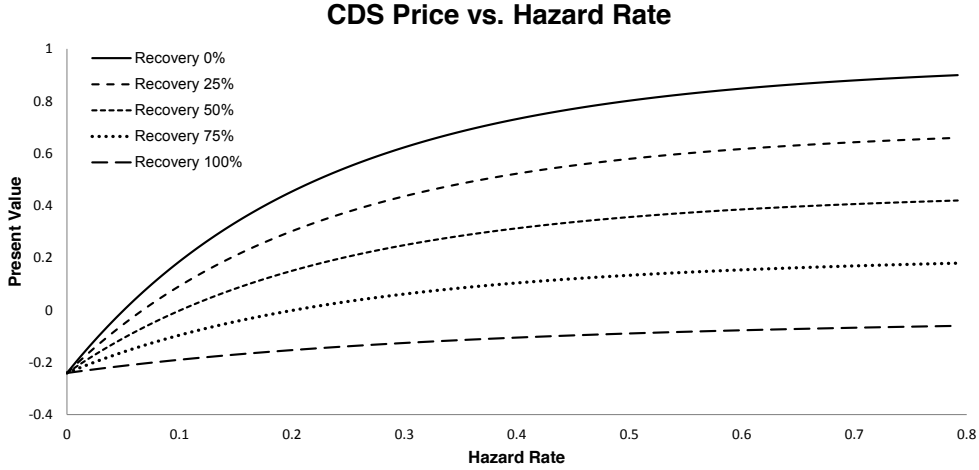


Figure 1: The normalised price (i.e. unit notional) for a standard 5Y CDS against hazard rate for a range of recovery rates. The CDS premium is 500bps.

This is also monotonic in λ and is guaranteed to have a root $G(\lambda^*) = 0$ provided that

$$-CN \sum_{i=1}^M \Delta_i P(t_i) \leq P(t_v) PV_{\text{clean}} < (1 - RR)N \quad (56)$$

A price outside this range is arbitrageable, so should not exist in an efficient economy. Without the first derivative, $G'(\lambda)$, one can use a gradient-free root finder such as the *bisection method*, *secant method* or *Van Wijngaarden-Dekker-Brent Method* (see [PTVF07] for details of these) to find the implied hazard rate. We find last of these to be the most effective. If the derivative available analytically the *Newton-Raphson method* is preferred - appendix B gives formulae for the first derivative for the ISDA model.

For a zero upfront CDS, the credit triangle $\lambda = S_p / (1 - RR)$ can be used as (a very good) initial guess. With PUF, we use the modified form $\lambda = (C + PUF/T) / (1 - RR)$. In either case, the root finder will converge in only a few iterations.

6.1.1 Conversion between PUF and Quoted Spreads

When CDS trades with a fixed premium and PUF (as all standard contracts now do), they can be given an equivalent *quoted spread*.²⁹ This is calculated by first finding the constant hazard rate that reprices the CDS, as above. With this constant hazard rate, the par spread of the CDS is calculated - this is the *quoted spread*.

To reverse the procedure, one first finds the constant hazard rate that reprices the CDS to zero given its premium is set to its quoted spread. One then prices the CDS from the constant hazard rate with its *actual premium*.

²⁹This is also called flat spread - see figure 4 and the accompanying text for a discussion on the (mis)use of this term.

There is a one-to-one mapping between price (PUF) and quoted spread for a CDS, and the link can be seen as analogous to the link between price and implied volatility of an option. Figure 2 shows PUF against quoted spread for fixed coupons of 0bps (i.e. no coupons), 100bps and 500bps, for a 10Y CDS with a recovery rate of 40%. When quoted spread equals the coupon, the PUF is zero.

Whether the price is given as PUF or a quoted spread, the actual trade is the same - an upfront fee is paid and then premiums, with standard coupons, are paid on the premium dates.

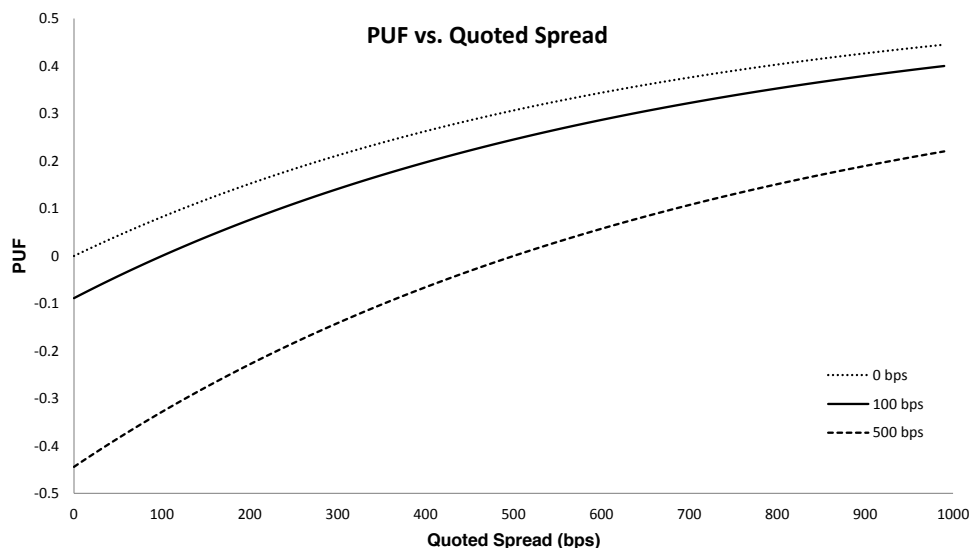


Figure 2: The Points Up-Front (PUF) of a 10Y CDS against quoted spread for three levels of fixed coupons. The recovery rate is 40%.

6.2 Multiple CDSs

Usually CDS quotes are available for several tenors on the same name, and one is interested in building a credit curve that will exactly reprice all of the quoted instruments.

A common approach for curve construction is to specify an interpolating scheme and a set of nodes (or knots) of equal number to the number of instruments, and at times matching the maturities.³⁰ We work with the zero hazard rate curve $\Lambda(t)$. Let $\mathbf{\Lambda} = (\Lambda_1, \dots, \Lambda_n)^T$ be the vector of curve nodes and $\mathbf{V} = (V_1, \dots, V_n)^T$ be the vector of market prices, then we must solve the vector equation

$$\mathbf{G}(\mathbf{\Lambda}) = \mathbf{V} \quad (57)$$

where $\mathbf{G}(\cdot)$ is the function that prices all the CDSs given the credit curve that results from the nodes $\mathbf{\Lambda}$. We need in effect to invert this equation, and have

$$\mathbf{\Lambda} = \mathbf{G}^{-1}(\mathbf{V}) \quad (58)$$

³⁰This isn't required, but generally results in smoother curves.

This can be handled by a multi-dimensional root finding technique such as the Newton and Quasi-Newton methods [PTVF07]. We have used this approach to good effect for the construction of multi-curves (i.e. separate discount and index curves) in the interest rate setting [Whi12].

Since the interpolator used in the ISDA model is *local* - meaning that the survival probability (or equivalently the zero hazard rate) at a particular time only depends on the value of the two adjacent nodes - the vector equation may be broken down as

$$\begin{aligned}
 G_1(\Lambda_1) &= V_1 \\
 G_2(\Lambda_1, \Lambda_2) &= V_2 \\
 &\vdots \\
 G_i(\Lambda_1, \Lambda_2, \dots, \Lambda_i) &= V_i \\
 &\vdots \\
 G_n(\Lambda_1, \dots, \Lambda_i, \dots, \Lambda_n) &= V_n
 \end{aligned} \tag{59}$$

The first equation may be solved by one-dimensional root finding for Λ_1 ; with this the second equation is also solved by one-dimensional root finding for Λ_2 . In this way the whole system is solved by n one-dimensional root findings, using the values found at previous steps: This process is known as *bootstrapping* a curve.³¹

The framework we have discussed included both the case where the market quotes are par spreads (i.e. the CDSs have zero cost of entry and the market view is reflected entirely through the premium or spread), in which case $\mathbf{V} = \mathbf{0}$, and when CDSs trade with a fixed premium and quotes are given as PUF.

6.2.1 Calibration from Quoted Spreads

Since by convention quoted spreads are computed from a constant hazard rate, they are different from par spread (although numerically close) - it is entirely inconsistent to treat these quotes as if they were par spreads to build a credit curve.

If CDSs with fixed premiums have quoted spreads (rather than PUF), one must first convert these to PUF before proceeding to bootstrap the credit curve.

6.2.2 Conversion between PUF and Par Spreads

It is consistent to build a credit curve from PUF quotes, then use this curve to compute the equivalent par spreads of CDSs with the same maturities, or indeed go the other way. Whether viewed in terms of PUF (or quoted spreads) or par spreads, the underlying credit curve is, by construction, identical. One use of this is to view credit risk in terms of shocks to the par spreads.

7 Sensitivity to the Credit Curve

We start by considering an arbitrary CDS, with present value, V and (with no loss of generality) a unit notional. We assume an interpolated credit curve with n nodes has been constructed, but make no further assumptions at this stage.

³¹The term bootstrap is used loosely to mean implying curves from market data, even when the technique used may not actually be a bootstrap.

Under this framework, quantities of fundamental importance are the sensitivity of the CDS value to the nodes of the credit curve. Formally these are:

$$\frac{\partial V(\mathbf{\Lambda})}{\partial \Lambda_k} \tag{60}$$

which form a vector of sensitivities. For any model, an element can be calculated by central difference, i.e.

$$\frac{\partial V(\mathbf{\Lambda})}{\partial \Lambda_k} \approx \frac{V(\Lambda_1, \dots, \Lambda_k + \eta, \dots, \Lambda_n) - V(\Lambda_1, \dots, \Lambda_k - \eta, \dots, \Lambda_n)}{2\eta} \tag{61}$$

where η is some small number.³² For a typical credit curve with eight nodes, the computation of the sensitivity vector takes 16 calls to the pricing code.

In most cases, the sensitivity can be calculated analytically. The general technique is often called *Algorithmic Differentiation*, which works by repeatedly applying the chain rule to elementary operations and functions [GW08, Hen12]. This is more accurate (i.e. machine precision) and faster. In the case of the ISDA Model, this becomes relatively simple. Appendix B gives explicit formulae for the protection and premium leg sensitivities to the credit curve.

Table 3 shows the PV sensitivity vectors for six CDSs with maturities corresponding to the nodes of the credit curve. As expected, there is no sensitivity to nodes beyond the maturity of the CDS. The value along the diagonal increases with maturity, as there is more scope for large price changes for longer-tenor CDSs. The numbers are the sensitivity of a unit notional CDS to the hazard rate. This is often scaled to be the sensitivity on a \$10MM notional to a 1bps change of hazard rate - in that case the first entry in the table would be \$308.69.

CDS Maturity	Curve Nodes					
	P6M	P1Y	P3Y	P5Y	P7Y	P10Y
20-Dec-11	0.30903	0.00000	0.00000	0.00000	0.00000	0.00000
20-Jun-12	0.00632	0.59684	0.00000	0.00000	0.00000	0.00000
20-Jun-14	0.00633	0.02487	1.66846	0.00000	0.00000	0.00000
20-Jun-16	0.00633	0.02487	0.14165	2.41260	0.00000	0.00000
20-Jun-18	0.00633	0.02487	0.14166	0.23604	2.87787	0.00000
20-Jun-21	0.00633	0.02487	0.14166	0.23605	0.35704	3.31445

Table 3: The PV sensitivity to zero hazard rates at the credit curve nodes for a set of standard CDSs with a coupon of 100bps and a recovery rate of 40%. The trade date is 13-Jun-2011. This can be viewed as a Jacobian matrix for the PV sensitivity to credit curve nodes.

A second table (4) shows the par-spread sensitivity to the zero hazard rate at standard maturity points for a set of CDSs. The trade date is 13-Jun-2011, the CDS maturities are on the June and December IMM dates out to ten years and the recovery rate is 40%. The two maturities (20-Jun-11 and 20-Dec-11) only have sensitivities to the six-month node (20-Dec-11), with a value that is consistent with the credit triangle (section 4.6). Even for CDS maturities that do not correspond to credit curve nodes, there is no sensitivity beyond the next node after that maturity.

³²Typically $\eta \approx 10^{-4} - 10^{-8}$. The optimal choice is beyond the scope of this paper - see any text book on numerical modelling.

CDS Maturity	Curve Nodes					
	P6M	P1Y	P3Y	P5Y	P7Y	P10Y
20-Jun-11	0.5921	0.0000	0.0000	0.0000	0.0000	0.0000
20-Dec-11	0.5932	0.0000	0.0000	0.0000	0.0000	0.0000
20-Jun-12	0.0055	0.5879	0.0000	0.0000	0.0000	0.0000
20-Dec-12	0.0041	0.2990	0.2890	0.0000	0.0000	0.0000
20-Jun-13	0.0033	0.1544	0.4337	0.0000	0.0000	0.0000
20-Dec-13	0.0027	0.0672	0.5212	0.0000	0.0000	0.0000
20-Jun-14	0.0023	0.0095	0.5791	0.0000	0.0000	0.0000
20-Dec-14	0.0022	0.0089	0.3821	0.1956	0.0000	0.0000
20-Jun-15	0.0020	0.0083	0.2379	0.3388	0.0000	0.0000
20-Dec-15	0.0019	0.0078	0.1269	0.4491	0.0000	0.0000
20-Jun-16	0.0018	0.0073	0.0394	0.5361	0.0000	0.0000
20-Dec-16	0.0017	0.0070	0.0371	0.3786	0.1590	0.0000
20-Jun-17	0.0016	0.0066	0.0352	0.2500	0.2889	0.0000
20-Dec-17	0.0015	0.0063	0.0336	0.1422	0.3978	0.0000
20-Jun-18	0.0014	0.0061	0.0321	0.0515	0.4895	0.0000
20-Dec-18	0.0014	0.0058	0.0307	0.0493	0.3918	0.1009
20-Jun-19	0.0013	0.0056	0.0296	0.0474	0.3078	0.1876
20-Dec-19	0.0013	0.0054	0.0285	0.0457	0.2343	0.2636
20-Jun-20	0.0012	0.0052	0.0276	0.0441	0.1697	0.3304
20-Dec-20	0.0012	0.0051	0.0267	0.0427	0.1126	0.3895
20-Jun-21	0.0012	0.0049	0.0259	0.0415	0.0621	0.4417

Table 4: The par-spread sensitivity to zero hazard rates at the credit curve nodes for a set of CDSs with recovery rate 40%.

7.1 Hedging the Credit Curve

With no loss of generality we take our CDS to have unit notional. We will hedge this with a collection of m other CDSs (on the same name) with present values V_i and (relative) notionals w_i - these are hedge ratios. The total value of the portfolio is

$$\Pi(\mathbf{\Lambda}) = V(\mathbf{\Lambda}) - \sum_{i=1}^m w_i V_i(\mathbf{\Lambda}) \quad (62)$$

We have emphasised the dependence on (the nodes of the) credit curve; the dependence on the yield curve we suppress. The sensitivity of the portfolio to the k^{th} node is given by

$$\frac{\partial \Pi(\mathbf{\Lambda})}{\partial \Lambda_k} = \frac{\partial V(\mathbf{\Lambda})}{\partial \Lambda_k} - \sum_{i=1}^m w_i \frac{\partial V_i(\mathbf{\Lambda})}{\partial \Lambda_k} \quad (63)$$

To ensure the portfolio is insensitive to any of the curve nodes, we must have

$$\mathbf{J}^T \mathbf{w} = \mathbf{v}_{\Lambda} \quad (64)$$

where

$$\begin{aligned} \mathbf{v}_\Lambda &= \left(\frac{\partial \Pi(\boldsymbol{\Lambda})}{\partial \Lambda_1}, \dots, \frac{\partial \Pi(\boldsymbol{\Lambda})}{\partial \Lambda_n} \right)^T, \\ \mathbf{w} &= (w_1, \dots, w_m)^T \\ \text{and } J_{i,k} &= \frac{\partial V_i(\boldsymbol{\Lambda})}{\partial \Lambda_k} \end{aligned}$$

In general the price Jacobian, \mathbf{J} , is not square,³³ so we must solve the system in a least squares sense. Formally

$$\mathbf{w} = (\mathbf{J}\mathbf{J}^T)^{-1}(\mathbf{J}\mathbf{v}_\Lambda) \quad (65)$$

will minimise $\|\mathbf{v}_\Lambda - \mathbf{J}^T \mathbf{w}\|_2$, i.e. give the lowest sum-of-squares of sensitivity of the portfolio to the credit curve nodes.

7.1.1 Hedging the ISDA Model Credit Curve

If we build an ISDA model credit curve (i.e. with piecewise constant forward hazard rates) and the hedging instruments are the same instruments used to construct the credit curve, then the Jacobian is square and is of lower diagonal form (i.e. it looks like table 3), i.e. $J_{i,j} = 0$ for $j > i$.

This means equation 64 can be solved by back-substitution

$$\begin{aligned} w_n &= \frac{\partial V}{\partial \Lambda_n} / \frac{\partial V_n}{\partial \Lambda_n} \\ w_{n-1} &= \left(\frac{\partial V}{\partial \Lambda_{n-1}} - \frac{\partial V_{n-1}}{\partial \Lambda_n} w_n \right) / \frac{\partial V_{n-1}}{\partial \Lambda_{n-1}} \\ &\vdots \\ w_1 &= \left(\frac{\partial V}{\partial \Lambda_1} - \sum_{i=2}^n \frac{\partial V_i}{\partial \Lambda_1} w_i \right) / \frac{\partial V_1}{\partial \Lambda_1} \end{aligned} \quad (66)$$

If the next node after the expiry of our CDS is k , then

$$\frac{\partial V}{\partial \Lambda_i} = 0 \quad \text{for } i > k$$

and

$$w_i = 0 \quad \text{for } i > k$$

In table 5 we show the hedge ratios for a set of CDSs. The trade date is (again) 13-Jun-2011 so the first CDS is only a week from maturity, so it is hedged with a small amount of the six-month CDS (maturity 20-Dec-11). All the CDSs that have a maturity corresponding to a hedging instrument, have a hedge ratio of exactly one to that instrument only, i.e. the portfolio cancels to contain no trades - this is a sanity check. For all other maturities, most of the significant hedge ratios are for the instruments with maturities immediately before and after the CDSs, i.e. all the CDSs could be hedged with just two standard CDSs. Furthermore, as the hedge ratios (almost) sum to one, this hedge makes our portfolio insensitive to default (the exception is the first, one week CDS).

³³usually $m < n$, that is, we hedge with fewer instruments than were used to build the curve.

CDS Maturity	Hedge Instrument Maturity					
	P6M	P1Y	P3Y	P5Y	P7Y	P10Y
20-Jun-11	0.03204	0.00000	0.00000	0.00000	0.00000	0.00000
20-Dec-11	1.00000	0.00000	0.00000	0.00000	0.00000	0.00000
20-Jun-12	0.00000	1.00000	0.00000	0.00000	0.00000	0.00000
20-Dec-12	0.00002	0.73999	0.26002	0.00000	0.00000	0.00000
20-Jun-13	0.00001	0.48703	0.51300	0.00000	0.00000	0.00000
20-Dec-13	0.00000	0.23923	0.76079	0.00000	0.00000	0.00000
20-Jun-14	0.00000	0.00000	1.00000	0.00000	0.00000	0.00000
20-Dec-14	-0.00000	-0.00001	0.73060	0.26963	0.00000	0.00000
20-Jun-15	-0.00001	-0.00001	0.47556	0.52474	0.00000	0.00000
20-Dec-15	-0.00000	-0.00001	0.23171	0.76851	0.00000	0.00000
20-Jun-16	0.00000	0.00000	0.00000	1.00000	0.00000	0.00000
20-Dec-16	-0.00000	-0.00001	-0.00002	0.72870	0.27157	0.00000
20-Jun-17	-0.00001	-0.00001	-0.00003	0.47298	0.52737	0.00000
20-Dec-17	-0.00000	-0.00001	-0.00002	0.22942	0.77083	0.00000
20-Jun-18	0.00000	0.00000	0.00000	0.00000	1.00000	0.00000
20-Dec-18	-0.00001	-0.00001	-0.00003	-0.00004	0.80969	0.19075
20-Jun-19	-0.00001	-0.00002	-0.00005	-0.00006	0.63029	0.37038
20-Dec-19	-0.00001	-0.00003	-0.00006	-0.00007	0.45940	0.54134
20-Jun-20	-0.00001	-0.00002	-0.00005	-0.00006	0.29752	0.70312
20-Dec-20	-0.00001	-0.00001	-0.00003	-0.00004	0.14427	0.85611
20-Jun-21	0.00000	0.00000	0.00000	0.00000	0.00000	1.00000

Table 5: The hedge ratios to make different maturity CDSs insensitive to infinitesimal changes to the credit curve. All the CDSs have a coupon of 100bps and a recovery rate of 40%. The trade date is 13-Jun-2011.

What we have shown is the delta hedge of the credit curve; the portfolio is still sensitive to large movements of the curve. As an example we consider hedging the 4Y CDS with the 3Y and 5Y standard instruments (with hedge ratios of 0.47556 and 0.52474). Figure 3 shows the change in the portfolio PV (for a notional of \$10MM of the 4Y CDS) against parallel and tilt³⁴ shifts to the zero hazard rates. Even for a large (100bps) shift in either direction the change is only \$250 for parallel shifts; this compares to a loss of \$220,000 for a 100bps down shift and a gain of \$212,000 for the same up shift, on the unhedged position. The shapes suggest that our hedged portfolio will gain (a small amount) from any shift in the credit curve (i.e. we have positive gamma). However, this will be offset by the portfolio losing value over time (negative theta).

7.2 Rebalancing a Portfolio

A trader may have a portfolio of illiquid CDSs (on the same name) and wish to swap these for a new portfolio of liquid CDSs, such that the new portfolio maintains the same sensitivities to the credit curve. A particular case is when the CDSs were issued pre Big-Bang (and thus have coupons set to whatever the par spreads were on the original trade date), and the trader would

³⁴the first node is lowered and the last node raised by the given amount - the other nodes have proportionate shifts, so the curve is tilted about the central node.

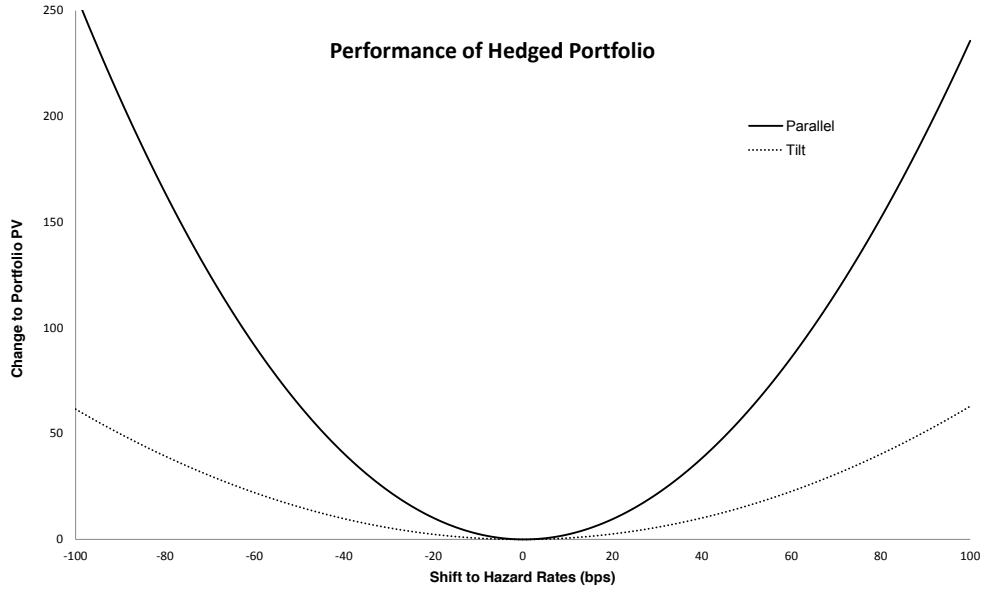


Figure 3: The performance of a hedged CDS portfolio to parallel and tilt shifts to the zero hazard rates.

prefer to hold new CDS with standard coupons. Such a process is called *portfolio compression*.

One way to compute the new portfolio is as follows: our original portfolio has p positions with notionals w_1, \dots, w_p and CDS PVs V_1, \dots, V_p ; the target portfolio has m positions with notionals $\hat{w}_1, \dots, \hat{w}_m$ and CDS PVs $\hat{V}_1, \dots, \hat{V}_m$. So that both portfolios have the same sensitivity to (infinitesimal) changes to the zero hazard rates, we require

$$\sum_{i=1}^p w_i \frac{\partial V_i(\boldsymbol{\Lambda})}{\partial \Lambda_k} = \sum_{i=1}^m \hat{w}_i \frac{\partial \hat{V}_i(\boldsymbol{\Lambda})}{\partial \Lambda_k} \quad \forall k \quad (67)$$

In matrix notation this becomes

$$\mathbf{J}^T \mathbf{w} = \hat{\mathbf{J}}^T \hat{\mathbf{w}} \quad (68)$$

which can be solved exactly provided that the Jacobian $\hat{\mathbf{J}}$ is square - i.e. there are as many instruments in the new portfolio as there are nodes in the credit curve. Generally one may wish to solve this problem in least-squares sense. Formally

$$\hat{\mathbf{w}} = (\hat{\mathbf{J}}\hat{\mathbf{J}}^T)^{-1} \hat{\mathbf{J}}\mathbf{J}^T \mathbf{w} \quad (69)$$

The trader may also desire that the new portfolio have the same value on default as the old, and that there is no cost (other than fees) in switching to the new portfolio. These two constraints

are met by

$$\begin{aligned}\sum_{i=1}^p w_i &= \sum_{i=1}^m \hat{w}_i \\ \sum_{i=1}^p w_i V_i &= \sum_{i=1}^m \hat{w}_i \hat{V}_i\end{aligned}\tag{70}$$

These can be incorporated by prepending two columns to both Jacobians, i.e.

$$\mathbf{A} = \begin{pmatrix} 1 & V_1 & & \\ 1 & V_2 & & \\ \vdots & \vdots & \mathbf{J} & \\ 1 & V_{p-1} & & \\ 1 & V_p & & \end{pmatrix} \quad \mathbf{B} = \begin{pmatrix} 1 & \hat{V}_1 & & \\ 1 & \hat{V}_2 & & \\ \vdots & \vdots & \hat{\mathbf{J}} & \\ 1 & \hat{V}_{m-1} & & \\ 1 & \hat{V}_m & & \end{pmatrix}\tag{71}$$

then solving

$$\hat{\mathbf{w}} = (\mathbf{B}\mathbf{B}^T)^{-1}\mathbf{B}\mathbf{A}^T\mathbf{w}\tag{72}$$

Finally, relative importance to the difference constraints³⁵ can be set by multiplying \mathbf{A} and \mathbf{B} by a diagonal matrix of weights.

This procedure does not tell the trader which CDSs to choose for the new portfolio, just what the optimal (in some sense) notionals are. A range of different portfolios could be tested, choosing the one with the smallest value of $|\mathbf{A}^T\mathbf{w} - \mathbf{B}^T\hat{\mathbf{w}}|_2$. This could be automated in some optimisation scheme, however this is outside the scope of this paper.

8 Spread Sensitivity or Credit DV01

As the market traditionally traded on par spread, traders like to think in terms of the effect on their portfolio due changes in spread rather than the credit curve directly.

The credit DV01 (CDV01) or CS01 (Credit Spread 01) is defined as the change in price of a CDS for a one basis point increase in spread. This simple definition belies a great deal of confusion as to which spread we are talking about.

8.1 Par Spread Sensitivity

If a credit curve is calibrated from a set of par spreads, we may bump up a single spread by 1bps and calibrate a new *bumped* credit curve. The difference in price of a reference CDS priced off the bumped versus the normal curve is the CS01 (of the CDS) to that par spread. This set of numbers (from individual bumped par spreads) is generally known as *bucketed CS01*. If all the par spreads are simultaneously bumped by 1bps (i.e. in parallel), the resultant single number is known as *parallel CS01*. Since the bump is finite (albeit small), the sum of the bucketed CS01 does not exactly equal the parallel CS01.

³⁵for example, you may care little about the sensitivity to the 10Y node, but a lot about both portfolios having the same value.

The CS01 defined above is just the (scaled) forward difference approximation to the sensitivity, $\frac{\partial V}{\partial S_k}$ when we bump the k^{th} par spread. To compute this exactly, we first consider the par spread Jacobian,³⁶ \mathbf{J}_S , where

$$(\mathbf{J}_S)_{i,j} = \frac{\partial S_i}{\partial \Lambda_j} \quad (73)$$

is a square matrix. Let $\mathbf{v}_s = \left(\frac{\partial V}{\partial S_1}, \dots, \frac{\partial V}{\partial S_n} \right)^T$, then

$$\mathbf{J}_S^T \mathbf{v}_s = \mathbf{v}_\Lambda \quad (74)$$

and the system can be solved for \mathbf{v}_s . In the ISDA model case, the par spread Jacobian is lower triangular, so the system can be solved by back-substitution. The elements of the vector \mathbf{v}_s are the analytical (exact) versions of the bucketed CS01; the sum of the elements is the equivalent of the parallel CS01.

Table 6 shows this sensitivity for a set of CDSs. The pattern is as expected in that there is no sensitivity to spreads beyond the one immediately after the CDS's maturity. An important point to note is that all the CDSs we price have a coupon of 100bps; if they had coupons equal to the par rates, then for a CDS with a maturity corresponding to one of the six pillar³⁷ maturities, the sensitivity would be only that spread, and the value would be just the CDS's RPV01.

Since it is more normal in the credit world to compute these numbers by a *bump and reprice* (i.e. a scaled forward finite difference), in table 7 below we compare the sensitivity numbers for a 8Y CDS (maturity 20-Jun-2019) computed analytically and by forward finite difference.³⁸ The numbers start to differ in the fourth significant figure,³⁹ and as previously mentioned the sum of the buckets does not equal the parallel in the finite difference case.

Other than matching the *accepted numbers* from Bloomberg, Markit etc., there is no compelling reason to compute CS01 by forward finite difference rather than analytically when using the ISDA model; the latter will be much quicker as it does not require (multiple) recalibration of the credit curve for each bump of spread.

8.1.1 Hedging against Par Spread Moves

To hedge against spread changes, we again consider a portfolio as in equation 62, but now write its value in terms of the spreads, \mathbf{S} and use instruments to hedge with the same maturities as those used to construct the credit curve.

$$\Pi(\mathbf{S}) = V(\mathbf{S}) - \sum_{i=1}^n w_i V_i(\mathbf{S}) \quad (75)$$

To hedge against the k^{th} par spread we require

$$\frac{\partial \Pi(\mathbf{S})}{\partial S_k} = \frac{\partial V(\mathbf{S})}{\partial S_k} - \sum_{i=1}^n w_i \frac{\partial V_i(\mathbf{S})}{\partial S_k} = 0 \rightarrow \mathbf{v}_s = \Theta^T \mathbf{w} \quad (76)$$

³⁶sensitivity of par spreads to the zero hazard rates - this is what was shown in table 4

³⁷pillar as they "hold up the curve".

³⁸In the normal way of quoting CS01, the PV difference is not divided by the bump size (1bps), and they are given on full notional (i.e. 10MM), so the numbers quoted will be 1000 ($10^7 \times 10^{-4}$) larger than here.

³⁹this is to be expected as we approximate a derivative with forward finite difference with a bump of 10^{-4} .

CDS Maturities	Tenors						Total
	P6M	P1Y	P3Y	P5Y	P7Y	P10Y	
20-Jun-11	0.0194	0.0000	0.0000	0.0000	0.0000	0.0000	0.0194
20-Dec-11	0.5229	0.0000	0.0000	0.0000	0.0000	0.0000	0.5229
20-Jun-12	0.0010	1.0173	0.0000	0.0000	0.0000	0.0000	1.0183
20-Dec-12	0.0005	0.7516	0.7495	0.0000	0.0000	0.0000	1.5016
20-Jun-13	-0.0000	0.4931	1.4787	0.0000	0.0000	0.0000	1.9718
20-Dec-13	-0.0005	0.2399	2.1930	0.0000	0.0000	0.0000	2.4324
20-Jun-14	-0.0010	-0.0045	2.8825	0.0000	0.0000	0.0000	2.8770
20-Dec-14	-0.0015	-0.0068	2.0895	1.2134	0.0000	0.0000	3.2946
20-Jun-15	-0.0020	-0.0090	1.3388	2.3615	0.0000	0.0000	3.6893
20-Dec-15	-0.0024	-0.0110	0.6210	3.4585	0.0000	0.0000	4.0661
20-Jun-16	-0.0028	-0.0130	-0.0610	4.5003	0.0000	0.0000	4.4235
20-Dec-16	-0.0030	-0.0138	-0.0641	3.2455	1.5964	0.0000	4.7610
20-Jun-17	-0.0031	-0.0145	-0.0670	2.0628	3.1001	0.0000	5.0782
20-Dec-17	-0.0033	-0.0152	-0.0697	0.9364	4.5312	0.0000	5.3794
20-Jun-18	-0.0034	-0.0158	-0.0722	-0.1247	5.8784	0.0000	5.6622
20-Dec-18	-0.0034	-0.0159	-0.0721	-0.1239	4.7172	1.4310	5.9330
20-Jun-19	-0.0034	-0.0159	-0.0719	-0.1230	3.6228	2.7787	6.1872
20-Dec-19	-0.0034	-0.0160	-0.0717	-0.1221	2.5802	4.0613	6.4283
20-Jun-20	-0.0034	-0.0160	-0.0715	-0.1211	1.5926	5.2751	6.6556
20-Dec-20	-0.0034	-0.0161	-0.0712	-0.1202	0.6578	6.4229	6.8698
20-Jun-21	-0.0034	-0.0161	-0.0709	-0.1191	-0.2223	7.5023	7.0705

Table 6: The sensitivity of the PV of a set of CDSs to the par spreads of the CDSs used to construct the credit curve. The last column shows the sensitivity of all the spreads moving in parallel. The (priced) CDSs all have a coupon of 100bps. All CDSs have a recovery rate of 40% and the trade date is 13-Jun-2011.

The quantity \mathbf{v}_s we met in equation 74; Θ is another Jacobian - it is the sensitivity of the hedging CDSs to the par spreads (it corresponds to selected rows of table 6). We take our hedging CDSs to be instruments one can readily buy, therefore they have a fixed coupon of 100bps. If instead we used par-CDSs, the Jacobian would be diagonal, with values equal to the CDS's RPV01.

Since Θ is lower triangular, we can solve equation 76 by back-substitution. We performed this exercise for our standard set of CDS; the results we obtained are identical to these in table 5. This should come as no surprise - one is hedging against infinitesimal movements of the hazard rates, and the other is hedging against infinitesimal movements of the par spreads, but these are just two ways of showing exactly the same information. The latter calculation involves eight matrix back-substitutions, while the former takes just one to produce the same result.

8.1.2 Rebalancing a Portfolio using Spread Sensitivities

In section 7.2 we discussed swapping an old portfolio (of CDSs on a single name) for a new one that had the same sensitivity to the zero hazard rates. The same exercise can be carried out for sensitivity to par spreads. Formally we seek to minimise

$$|\Theta^T \mathbf{w} - \hat{\Theta}^T \hat{\mathbf{w}}|_2 \quad (77)$$

Calculation Method	P6M	P1Y	P3Y	P5Y	P7Y	P10Y	Sum	Parallel
Analytic	-0.00343	-0.01593	-0.07193	-0.12300	3.62276	2.77873	6.18718	6.18718
Forward FD	-0.00343	-0.01593	-0.07193	-0.12300	3.62328	2.78004	6.18903	6.18322

Table 7: The sensitivity of the PV of 8Y CDSs to the par spreads of the CDSs used to construct the credit curve. The calculation methods are analytic and forward finite difference (or bump and reprice). The bump in the forward difference is 1bps.

If the Jacobians Θ^{40} are computed analytically, the same weights for the new portfolio will emerge as for computing them via zero hazard rate sensitivities - except more work is involved in the calculation.

Most often Θ and $\hat{\Theta}$ are computed by finite difference (i.e. bump and reprice), which involves refitting the credit curve several times. While this will ultimately produce very similar results to the analytic method of section 7.2, this method is considerably slower.

8.2 Quoted Spread Sensitivity

Any standard CDS issued after April 2009 will have a standard coupon and trade with an upfront fee (Points Up-Front - PUF). Section 6.1.1 details conversion between PUF and quoted spreads. Given a set of CDSs with PUF (or equivalently quoted spreads), one can build the credit curve and manage the market credit risk by considering movements to the underlying zero hazard rates as in section 7.1. Practitioners seem to prefer to see risk to the quoted spread (as this is a direct market observable), which results in some rather peculiar methodologies.

8.2.1 Parallel CS01

This is calculated by pricing the CDS using its quoted spread (see section 6.1.1), then bumping the quoted spread by 1bps and recalculating the price; the difference is the (parallel) CS01. This number is reported by Markit, Bloomberg etc. It can be regarded as the (forward) finite difference approximation for

$$\frac{\partial V}{\partial S} = \left(\frac{\partial V}{\partial \lambda} \right) / \left(\frac{\partial S}{\partial \lambda} \right) \quad (78)$$

The difference can be quite large; for a 10Y CDS with a 100bps coupon and a 130bps quoted spread (trade level), the (scaled) CS01 is 6.1885 while the exact derivative is 6.5880.

8.2.2 Bucketed CS01

There is no universally accepted definition for *bucketed CS01*. We will present a few methodologies.

Bumping Flat Spreads

- For a target CDS, choose a set of maturities (pillars) that you want to measure sensitivity to (these could be the standard liquid points of 6M, 1Y, 3Y, 5Y, 7Y and 10Y).
- Set the spreads at these points equal to the quoted spread of the target CDS.

⁴⁰recall these are the sensitivity of the PV of CDSs to the par rates of the CDSs used to build the credit curve.

- Build a credit curve from the pillar CDSs assuming the spreads are par spreads.
- Price the target CDS from this curve.
- Bump each spread in turn by 1bps, build a credit curve and price the target CDS from this new curve.
- The differences from the original price are the bucketed CS01s.

There are a number of odd features of this method. Firstly flat spreads are not equivalent to a constant hazard rate. This is shown clearly in figure 4 where we plot the par spread of different maturity CDSs for a constant hazard rate of 100bps (as usual the recovery rate is 40%). The constant hazard rate has given rise to a term structure of spreads. The range is not large (around 0.1bps), however since the reverse is also true (a flat term structure of spreads will not give a constant hazard rate), the price of the target CDS will be (slightly) wrong before we bump the spreads. Secondly, absolutely no account is made for the actual credit curve, just the fictitious one that is constructed from the flat spreads.

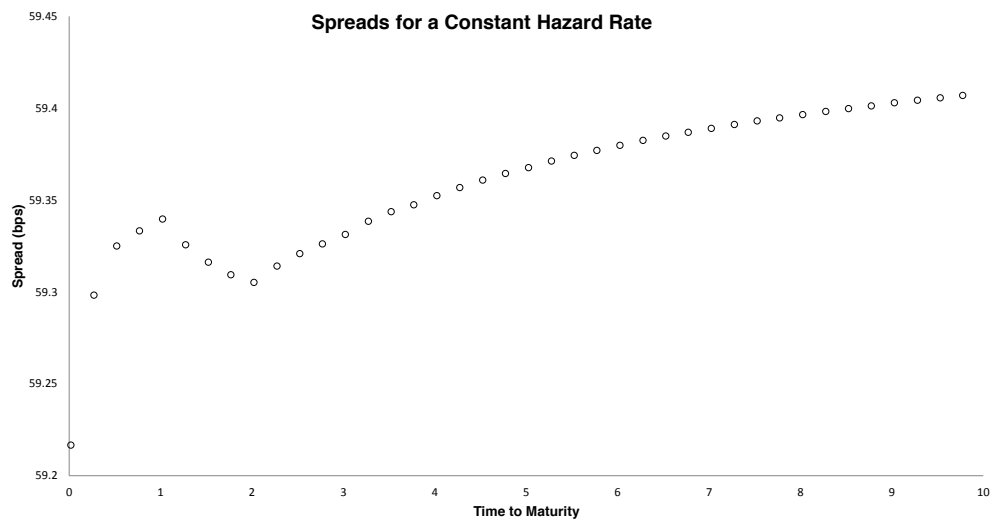


Figure 4: The par spread for a set of CDSs with a constant hazard rate of 100bps and a recovery rate of 40%. Clearly, a constant hazard rate induces a (small) term structure of spreads.

This method is, however, widely used. If the (flat) spreads are bumped in parallel, the result is called parallel CS01, although it will differ from the result of the method in the last section. Table 8 shows the results of this calculation for a set of CDS with maturities one year apart. The pattern is similar to what we saw for par spread sensitivity - that should not be surprising, since after setting all the spreads to the same value, the methods are identical, and hence an equivalent analytic flat spread CS01 can be calculated.

CDS Maturities	Bucketed						sum	Parallel
	P6M	P1Y	P3Y	P5Y	P7Y	P10Y		
20-Jun-11	0.0194	0.0000	0.0000	0.0000	0.0000	0.0000	0.0194	0.0194
20-Jun-12	0.0009	1.0177	0.0000	0.0000	0.0000	0.0000	1.0186	1.0185
20-Jun-13	-0.0005	0.4954	1.4720	0.0000	0.0000	0.0000	1.9669	1.9667
20-Jun-14	-0.0009	-0.0045	2.8723	0.0000	0.0000	0.0000	2.8669	2.8667
20-Jun-15	-0.0020	-0.0100	1.3489	2.3155	0.0000	0.0000	3.6525	3.6511
20-Jun-16	-0.0025	-0.0127	-0.0608	4.4386	0.0000	0.0000	4.3627	4.3616
20-Jun-17	-0.0028	-0.0144	-0.0681	2.0550	3.0233	0.0000	4.9931	4.9899
20-Jun-18	-0.0030	-0.0152	-0.0711	-0.1237	5.7673	0.0000	5.5544	5.5523
20-Jun-19	-0.0030	-0.0154	-0.0714	-0.1234	3.5840	2.6922	6.0629	6.0574
20-Jun-20	-0.0030	-0.0154	-0.0704	-0.1207	1.5873	5.1380	6.5158	6.5100
20-Jun-21	-0.0029	-0.0152	-0.0686	-0.1165	-0.2208	7.3404	6.9164	6.9129

Table 8: The *flat spread* method for parallel and bucketed CS01.

Bumping Pillar Quoted Spreads

- For a target CDS, choose a set of pillars that you want to measure sensitivity to (these should correspond to liquid CDS with available market prices).
- From the quoted spreads at the pillars, build a credit curve as detailed in section 6.2.1.
- Price the target CDS from this curve.
- Bump each quoted spread in turn by 1bps, build a credit curve and price the target CDS from this new curve (see comment below).
- The differences from the original price are the bucketed CS01s.

Table 9 shows the results of this method. The pattern is very different: CDSs corresponding to a pillar (e.g. 10Y CDS, maturity 20-Jun-2021) only have sensitivity to that pillar, and all the others only have any significant sensitivity to the adjacent pillars. This is just a reflection of the fact that quoted spreads have a one-to-one relationship with CDS prices.

The constructed credit curve will imply a price for the target CDS; this price may not be the same as that quoted in the market (via PUF or quoted spread). The only way to ensure that the target CDS is priced correctly is to include it as an instrument used to build the credit curve.

Bumping Implied Par Spreads

- Build a credit curve as for the *bumping pillar quoted spreads* method above.
- Imply equivalent par spreads at the pillars as in section 6.2.2.
- Use these implied par spreads as in section 8.1 to compute parallel and bucketed CS01.

This method obviously produces results constant with section 8.1, i.e. a pattern like table 6. It also gives sensitivity to (implied) par spreads rather than quoted spreads.

CDS Maturities	Bucketed						Sum	Parallel
	P6M	P1Y	P3Y	P5Y	P7Y	P10Y		
20-Jun-11	0.0194	0.0000	0.0000	0.0000	0.0000	0.0000	0.0194	0.0194
20-Jun-12	0.0000	1.0185	0.0000	0.0000	0.0000	0.0000	1.0185	1.0185
20-Jun-13	0.0000	0.4961	1.4708	0.0000	0.0000	0.0000	1.9669	1.9675
20-Jun-14	0.0000	0.0000	2.8667	0.0000	0.0000	0.0000	2.8667	2.8667
20-Jun-15	0.0000	0.0000	1.3641	2.2901	0.0000	0.0000	3.6542	3.6513
20-Jun-16	0.0000	0.0000	0.0000	4.3651	0.0000	0.0000	4.3651	4.3618
20-Jun-17	0.0000	0.0000	-0.0001	2.0666	2.9329	0.0000	4.9995	4.9901
20-Jun-18	0.0000	0.0000	0.0000	0.0000	5.5634	0.0000	5.5634	5.5526
20-Jun-19	0.0000	0.0000	-0.0001	-0.0003	3.5101	2.5680	6.0777	6.0579
20-Jun-20	0.0000	0.0000	-0.0001	-0.0003	1.6585	4.8768	6.5349	6.5106
20-Jun-21	0.0000	0.0000	0.0000	0.0000	0.0000	6.9382	6.9382	6.9136

Table 9: The *bumping pillar quoted spreads* method for parallel and bucketed CS01.

8.3 Sensitivity to Arbitrary Spreads

Traders often want to see bucketed CDS01 to more (or different) maturities than the CDSs used to construct the credit curve. One way to do this is to build the credit curve in the normal way from the pillar CDSs, then imply par spreads at the specified bucket maturities. These par spreads (and the corresponding fictitious CDSs) can be used to compute CS01 as in section 8.1. The normal caveat that the target CDS may not be priced correctly applies.

8.4 Credit Gamma or Convexity

This is the second order sensitivity to the credit curve or spread. The general curve gamma is

$$\Gamma_{i,j} = \frac{\partial^2 V(\mathbf{\Lambda})}{\partial \Lambda_i \Lambda_j} \quad (79)$$

although one rarely considers the cross terms. This can be found by finite difference, or for the ISDA model, analytically.⁴¹

What is more often quoted is the spread gamma or convexity (or gamma CS01). This is normally calculated by finite difference as

$$\frac{\partial^2 V(\mathbf{\Lambda})}{\partial S_i^2} \approx \frac{V(\mathbf{\Lambda}[S_1, \dots, S_i + \eta, \dots, S_n]) - 2V(\mathbf{\Lambda}[\mathbf{S}]) + V(\mathbf{\Lambda}[S_1, \dots, S_i - \eta, \dots, S_n])}{\eta^2} \quad (80)$$

where η is one basis point (1bps) and the notation $V(\mathbf{\Lambda}[\mathbf{S}])$ means build the credit curve $\mathbf{\Lambda}$ from the spreads, then use the credit curve to price the CDS.

The change in value of a CDS due to movements of the par spreads (of the instruments used to construct the credit curve) can be written as

$$\Delta V \approx \sum_{i=1}^n \frac{\partial V}{\partial S_i} \Delta S_i + \sum_{i=1, j=1}^n \frac{\partial^2 V}{\partial S_i \partial S_j} \Delta S_i \Delta S_j + \dots \quad (81)$$

⁴¹This second derivative is more involved, and we are yet to implement it analytically.

9 Yield Curve Sensitivities

Similarly to the credit curve sensitivity, for an interpolated yield curve we take this to be the sensitivity of the CDS's PV to the zero rates at the curve nodes.

$$\frac{\partial V}{\partial R_i}$$

We call the vector of these values \mathbf{v}_R . In the general case this can be found by finite difference. For the ISDA model exact formulae can be derived, and these are given in appendix C. We also give the sensitivity of money market instruments and swaps to the yield curve - since these are used to construct the yield curve, and form natural hedging instruments against movements to the yield curve. Table 10 shows the sensitivity of a set of CDS to the yield curve nodes (zero rates).⁴² This should be compared with table 3, where the sensitivities are at least an order of magnitude larger.

CDS Maturity	Yield Curve Nodes									
	1M	3M	6M	1Y	3Y	5Y	7Y	10Y	11Y	12Y
20-Jun-11	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
20-Jun-12	-0.0001	0.0002	0.0004	0.0018	0.0001	0.0000	0.0000	0.0000	0.0000	0.0000
20-Jun-13	0.0000	0.0002	0.0004	-0.0021	-0.0015	0.0000	0.0000	0.0000	0.0000	0.0000
20-Jun-14	0.0000	0.0002	0.0004	-0.0037	-0.0093	0.0001	0.0000	0.0000	0.0000	0.0000
20-Jun-15	0.0001	0.0002	0.0004	-0.0037	-0.0406	-0.0138	0.0000	0.0000	0.0000	0.0000
20-Jun-16	0.0002	0.0002	0.0004	-0.0037	-0.0513	-0.0581	0.0001	0.0000	0.0000	0.0000
20-Jun-17	0.0003	0.0002	0.0004	-0.0037	-0.0513	-0.1080	-0.0192	0.0000	0.0000	0.0000
20-Jun-18	0.0004	0.0002	0.0004	-0.0037	-0.0513	-0.1251	-0.0800	0.0001	0.0000	0.0000
20-Jun-19	0.0005	0.0002	0.0004	-0.0037	-0.0513	-0.1251	-0.1470	-0.0159	0.0000	0.0000
20-Jun-20	0.0006	0.0002	0.0004	-0.0037	-0.0513	-0.1251	-0.1852	-0.0655	0.0000	0.0000
20-Jun-21	0.0007	0.0002	0.0004	-0.0037	-0.0513	-0.1251	-0.1980	-0.1439	0.0002	0.0000

Table 10: The PV sensitivity to zero rates at the yield curve nodes for a set of standard CDSs with a coupon of 100bps and a recovery rate 40%. The trade date is 13-Jun-2011.

9.1 Hedging Yield Curve Movements

A CDS position (or a portfolio of CDSs) can be hedged against infinitesimal movements of the yield curve by adding money market contracts and swaps as hedging instruments. The calculation proceeds exactly as in section 7.1 except that we are now considering the sensitivity to the nodes of the yield curve.

9.2 Money Market and Swap Sensitivities

Mirroring what happens with the credit curve, traders often want to see the sensitivity of a CDS to the money market and swap rates of the instruments used to construct the yield curve rather than to the derived zero rates.

⁴²To fit the table in we have reduced the number of nodes on the yield curve.

Traditionally, these numbers, like other sensitivities in the credit world, have been calculated by forward finite difference: Each market rate is bumped in turn by one basis point; a new yield curve is bootstrapped; and finally the CDS is repriced from the new yield curve (the credit curve remains unchanged). The difference in price is the *bucketed IR01*. If the rates are bumped in parallel this is the *parallel IR01*.

The exact derivative can be calculated exactly as the spread sensitivity. Calling S_i the i^{th} market rate (money market or swap)⁴³, the market rate Jacobian is

$$(\mathbf{J})_{i,j} = \frac{\partial S_i}{\partial R_j} \quad (82)$$

which is a square matrix. Let $\mathbf{v}_s = \left(\frac{\partial V}{\partial S_1}, \dots, \frac{\partial V}{\partial S_n} \right)^T$, then

$$\mathbf{J}_S^T \mathbf{v}_s = \mathbf{v}_R \quad (83)$$

As \mathbf{J}_S is lower triangular, we can solve for \mathbf{v}_s by back-substitution.

We can hedge a CDS portfolio's sensitivity to market interest rates quotes exactly like we hedged a CDS to market CDS spread - but again, this will give the same results as hedging directly against the yield curve.

10 Other Risk Factors

10.1 Recovery Rate Risk

The sensitivity of the PV to the recovery rate (assuming the credit curve is immutable) can be read straight from equation 14:

$$\frac{\partial PV}{\partial RR} = \frac{N}{P(t_v)} \int_0^T P(s) dQ(s) \quad (84)$$

This is of course just the (negative of) the value of the protection leg with zero recovery rate.

A different type of recovery risk can be computed by refitting the credit curve with a bumped recovery risk, but leaving the trade recovery untouched.

10.2 Value on Default or Jump to Default

Immediately prior to default, the CDS has some value V (to the protection buyer). After default, the contract is cancelled, so there is an immediate loss of $-V$ ⁴⁴[O'k08]. The protection buyer pays the accrued interest, A and receives $N(1 - RR)$, so the full Value on Default (VoD) for the protection buyer is

$$\text{VoD} = -V + N(1 - RR) - A \quad (85)$$

The VoD for the protection seller is the negative of this (where V is still the PV prior to default for the protection buyer). In this formula, V is the dirty price. Since the clean price is $V_{\text{clean}} = V_{\text{dirty}} + A$, we can absorb A and use the clean price instead.

If default happens after a large widening of spreads (i.e. the default is somewhat expected), then VoD will be small.

⁴³ S has been previously used for spread, we hope this does not cause confusion.

⁴⁴If V was negative, this is a gain.

11 Beyond the ISDA Model

The ISDA model is a market standard and can be made robust and quick to calibrate. However, the curves (yield and credit) it produces are not realistic since they imply sudden jumps in forward rates. This in turn means that prices and risk factors implied from the curves should be viewed with caution.

For interest rate derivatives, we have used a range of interpolators in a multi-curve setting [Whi12]. These interpolators can all be expressed as piecewise polynomials [dB78, Iwa13]. Since the integrals involved in CDS pricing take exponents of the zero curves, anything beyond linear interpolation will not have a closed form solution for the integrals. The integrals could be computed numerically, with a tradeoff between accuracy and computation time. Alternatively, one could completely remove the need to carry out the integrals by using the approximation discussed in section 4.4.

Whether we compute the integrals numerically or remove them entirely, we are no longer in a position to bootstrap the credit curve as the Jacobian (of CDS PVs to credit curve nodes) is no longer triangular. Instead we must solve for all the nodes simultaneously using a multi-dimensional root finder - this is now standard in the interest rate world.

Depending on the interpolator used,⁴⁵ the resultant credit curve will be smoother (no jumps in the forward rate). Whether the prices and risk factors are ‘better’ than those produced from the ISDA model can only be determined with substantial back-testing, which is outside the scope of this paper.

11.1 The Discount Curve

Since CDS are over-the-counter (OTC) derivatives, it is assumed that discounting should be at the (risky) inter-bank rate (i.e. Libor or its equivalents). Interest rate swaps are now almost always collateralised according to a credit support annex (CSA), and discounting is performed using the risk free curve (which is normally build from overnight index swaps (OIS). With CDS trades moving onto exchanges, the correct discount curve should be the OIS rather than Libor.

A ISDA Model Dates

The ISDA model is formulated with an integer date representation; this is taken to mean that the trade and all payments are assumed to be made at 11:59pm in the relevant time zone. For completeness we list all inputs to the ISDA model. The exact meaning of certain terms will be clarified when they are used in subsequent sections.

- *Trade Date.* The date when the trade is executed. This is denoted as T, with T+n meaning n days after the trade date
- *Step-in or Protection Effective Date.* This is usually T+1. This is when protection (and risk) starts in terms of the model. Note the legal effective date is T-60 or T-90 for standard contracts.
- *Valuation or Cash-settle date* This is the date for which the present value (PV) of the CDS is calculated. It is usually three working dates after the trade date.

⁴⁵some interpolators, notably cubic spline, can produce widely oscillating curves, which reprice all the instruments but are hardly believable.

- *Start or Accrual Begin Date* . This is when the CDS nominally starts in terms of premium payments, i.e. the number of days in the first period (and thus the amount of the first premium payment) is counted from this date.
- *End or Maturity Date* This is when the contract expires and protection ends - any default after this date does not trigger a payment.
- *Pay Accrual on Default* Is a partial payment of the premium made to the protection seller in the event of a default. This is normally true.
- *Payment Interval* The interval between premium payments. This is three months for a standard CDS contract.
- *stub-type* This can be front-short, front-long, back-short or back-long. Front-short is normal, although somewhat moot when maturity and nominal payments are all on IMM dates.
- *Protection at Start of Day* Does protection start at the beginning of the day.
- *Recovery Rate* The recovery rate.
- *Business-day Adjustment* How are adjustments for non-business days made. This is normally *following*, i.e. move forward in steps of one day until a good business day is found.
- *Calendar* Calendar defining what is a non-business day. Except for CDS denominated in JPY (which use the Tokyo holiday calendar), normally a weekend only calendar is used.
- *Accrual Day Count Convention* DDC used to calculate the amount of the premium payments. This is normally ACT/360.

A.1 The Premium Leg

The payments on the premium leg depend on three sets of dates:

- The payment dates - actual cash payments are made on these dates
- The accrual start dates
- The accrual end dates

The premium payment is given by the year fraction between the accrual start and end dates multiplied by the coupon and notional. As is normal in day counts, the first day is included and the second is excluded.

Example: For an accrual start of 18-Feb-2009 and an accrual end of 20-May-2009, the number of days is 91; using the usual day count convention of ACT/360, the year fraction is $91/360 \approx 0.2528$.

Given the specification of a *stub-type* and the *payment interval*, nominal dates are generated by either rolling back from the maturity date in integer multiples of the payment interval, until the start date is passed (for front stubs) or forward from the start date until the maturity date is passed (for back stubs). For front-stub the first date is set to the start date, and for front-long the second date is removed. For back-stub the last date is set to the maturity, and for back-long the penultimate date is removed.

Example: if the start date is 30-Jul-2012, the maturity is 29-May-2013, the stub is front-short and the payment interval is 3M, the nominal dates will be: 29-May-2013, 01-Mar-2013, 29-Nov-2012, 29-Aug-2012, and 30-Jul-2012.

Accrual dates are generated by:

- The first accrual start date is set to the first nominal date (which will be the start date), i.e. it is unadjusted
- The i^{th} payment date is set to the $(i + 1)^{th}$ nominal date business-day adjusted according to a specified holiday calendar and business day convention.
- All subsequent accrual start dates are set to the previous payment date.
- All but the last accrual end date are set to the payment date.
- The final accrual end date is set to the final nominal date (i.e. the maturity), plus one day,⁴⁶ i.e. it is unadjusted.

A.1.1 The Standard CDS Contract

For a standard contract the maturity is always an unadjusted IMM date⁴⁷ and the coupon interval is 3M (quarterly), so the nominal payment dates will all be IMM dates. The business day convention is *following*, the day count convention is ACT/360 and the holiday calendar is determined by currency. For cash-settled amounts and non-JPY currencies the calendar is taken as a weekday only [ISD09].

The accrual begin date is taken as the latest adjusted IMM on or before the step-in date (T+1). So while the ISDA model does not adjust the accrual start date, it should be adjusted before being input to the model.

Example: A 2Y CDS with a trade date of 30-Jul-13 will have a maturity of 20-Sep-15 (two years on from the next IMM date after the trade date). With a weekday calendar, the dates and cash-flows for a 10MM notional 100bps CDS are shown in table A.1.1 below.

Accrual Start (Inclusive)	Accrual End (Exclusive)	Payment Date	Days in Period	Amount
20-Jun-13	20-Sep-13	20-Sep-13	92	25555.56
20-Sep-13	20-Dec-13	20-Dec-13	91	25277.78
20-Dec-13	20-Mar-14	20-Mar-14	90	25000.00
20-Mar-14	20-Jun-14	20-Jun-14	92	25555.56
20-Jun-14	22-Sep-14	22-Sep-14	94	26111.11
22-Sep-14	22-Dec-14	22-Dec-14	91	25277.78
22-Dec-14	20-Mar-15	20-Mar-15	88	24444.44
20-Mar-15	22-Jun-15	22-Jun-15	94	26111.11
22-Jun-15	21-Sep-15	21-Sep-15	91	25277.78

Table 11: dates and cash-flows for a 2Y CDS with a trade date of 30-Jul-13.

⁴⁶This is for the normal case that protection is from the start of the day; if it is not, the extra day is not added.

⁴⁷20th March, June, September and December.

A.1.2 Protection at Start of Day

This is handled in the ISDA model by subtracting one day from *credit observation dates*. This means that the period end date (which first appears in equation 15) is taken as the accrual end date minus one day. The dates for the calculation of the accrual on default also have this treatment, so are the accrual start and ends dates minus one day.

A.2 The Protection Leg

The effective start date is taken as the later of the step-in date and the start date. If protection is at start of day, one day is subtracted from this. The end date is unadjusted. These two dates (where converted to year fractions from the trade date) are used in equation 40 to compute the protection leg.

B Curve Sensitivity for the ISDA Model

A quantity that will be used in subsequent calculations is the sensitivity of the survival probability at time t , $Q(t)$, to the value of the zero hazard rate at the i^{th} node. This is given simply by

$$\frac{\partial Q(t)}{\partial \Lambda_i} = \begin{cases} -t_i Q_i & \text{if } t = t_i \\ -t_i Q(t) \frac{(t-t_{i-1})}{\Delta t_{i-1}} & \text{if } t_{i-1} < t < t_i \\ -t_i Q(t) \frac{(t_{i+1}-t)}{\Delta t_i} & \text{if } t_i < t < t_{i+1} \\ 0 & \text{otherwise} \end{cases} \quad (86)$$

Of course, exactly analogous equations hold for the yield curve (i.e. discount factor sensitivity to zero rates at nodes).

B.1 Protection Leg Sensitivity

Using the notional of equation 41, we define the i^{th} integral element sensitivity to the survival probability at t_j as

$$\frac{\partial I_i}{\partial \bar{Q}_j} = \begin{cases} \frac{1}{\hat{f}_i + \hat{h}_i} \left[\frac{\hat{f}_i (\bar{B}_{i-1} - \bar{B}_i)}{\bar{Q}_{i-1} (\hat{f}_i + \hat{h}_i)} + \hat{h}_i \bar{P}_{i-1} \right] & \text{if } j = i-1 \\ -\frac{1}{\hat{f}_i + \hat{h}_i} \left[\frac{\hat{f}_i (\bar{B}_{i-1} - \bar{B}_i)}{\bar{Q}_i (\hat{f}_i + \hat{h}_i)} + \hat{h}_i \bar{P}_i \right] & \text{if } j = i \\ 0 & \text{otherwise} \end{cases} \quad (87)$$

From the chain rule, the i^{th} integral element sensitivity to (zero hazard rate at the) j^{th} node is:

$$\frac{\partial I_i}{\partial \Lambda_j} = \sum_{k=1}^n \frac{\partial I_i}{\partial \bar{Q}_k} \frac{\partial \bar{Q}_k}{\partial \Lambda_j} = \frac{\partial I_i}{\partial \bar{Q}_{i-1}} \frac{\partial \bar{Q}_{i-1}}{\partial \Lambda_j} + \frac{\partial I_i}{\partial \bar{Q}_i} \frac{\partial \bar{Q}_i}{\partial \Lambda_j} \quad (88)$$

Finally, the protection leg sensitivity to the credit curve is

$$\frac{\partial PV_{\text{Protection Leg}}}{\partial \Lambda_j} = \frac{N(1-RR)}{P(t_v)} \sum_{i=1}^n \frac{\partial I_i}{\partial \Lambda_j} \quad (89)$$

B.2 Premium Leg Sensitivity

B.2.1 Premiums Only

The sensitivity of the premiums only is very straightforward; from equation 15 we can derive

$$\frac{\partial PV_{\text{Premiums only}}}{\partial \Lambda_j} = \frac{NC}{P(t_v)} \sum_{i=1}^M \Delta_i P(t_i) \frac{\partial Q(e_i)}{\partial \Lambda_j} \quad (90)$$

B.2.2 Accrual Paid On Default

We focus on the solution of the integral for the k^{th} premium period given in equation 53. The i^{th} element of that sum we call I_i^k , and for ease of reference we reproduce it here:

$$I_i^k = \hat{h}_i \bar{B}_{i-1} \left((t_{i-1} - s_k) \epsilon(-(\hat{f}_i + \hat{h}_i)) + \Delta t_{i-1} \epsilon'(-(\hat{f}_i + \hat{h}_i)) \right)$$

The sensitivity of this to the survival probability is:

$$\frac{\partial I_i^k}{\partial Q_j} = \begin{cases} \bar{P}_{i-1} \left(\phi_1 (1 + \hat{h}_i) - \phi_2 \hat{h}_i \right) & \text{if } j = i - 1 \\ \frac{\bar{B}_{i-1}}{Q_i} \left(-\phi_1 + \hat{h}_i \phi_2 \right) & \text{if } j = i \\ 0 & \text{otherwise} \end{cases} \quad (91)$$

where we have defined the following variables

$$\begin{aligned} \phi_1 &= (t_{i-1} - s_k) \epsilon(-(\hat{f}_i + \hat{h}_i)) + \Delta t_{i-1} \epsilon'(-(\hat{f}_i + \hat{h}_i)) \\ \phi_2 &= (t_{i-1} - s_k) \epsilon'(-(\hat{f}_i + \hat{h}_i)) + \Delta t_{i-1} \epsilon''(-(\hat{f}_i + \hat{h}_i)) \end{aligned} \quad (92)$$

The sensitivity to the zero hazard rates at the nodes is:

$$\frac{\partial I_i^k}{\partial \Lambda_j} = \frac{\partial I_i^k}{\partial \bar{Q}_{i-1}} \frac{\partial \bar{Q}_{i-1}}{\partial \Lambda_j} + \frac{\partial I_i^k}{\partial \bar{Q}_i} \frac{\partial \bar{Q}_i}{\partial \Lambda_j} \quad (93)$$

The full sensitivity of the premium leg is

$$\frac{\partial PV_{\text{premium}}}{\partial \Lambda_j} = \frac{NC}{P(t_v)} \sum_{i=1}^M \left[\Delta_i P(t_i) \frac{\partial Q(e_i)}{\partial \Lambda_j} - \eta_i \sum_{k=1}^{n_i} \frac{\partial I_k^i}{\partial \Lambda_j} \right] \quad (94)$$

Since we have used the $\epsilon()$ form, the limit $\hat{f}_i + \hat{h}_i \rightarrow 0$ is handled correctly.

To get the sensitivity using the original ISDA formula, simply replace s_k with $\tilde{s}_k = s_k - 1/730$. To use the Markit fix formula, drop the first terms from equation 92 (i.e. set $t_{i-1} - s_k = 0$).

C Yield Curve Sensitivity for the ISDA Model

The formulae in this section are very similar to those in appendix B and we follow closely the notation used there.

C.1 Protection Leg Sensitivity

The sensitivity of the i^{th} integral element to the discount factor at t_j is

$$\frac{\partial I_i}{\partial \bar{P}_j} = \begin{cases} \frac{\hat{h}_i}{\hat{f}_i + \hat{h}_i} \left[\frac{-(\bar{B}_{i-1} - \bar{B}_i)}{\bar{P}_{i-1}(\hat{f}_i + \hat{h}_i)} + \bar{Q}_{i-1} \right] & \text{if } j = i - 1 \\ \frac{\hat{h}_i}{\hat{f}_i + \hat{h}_i} \left[\frac{(\bar{B}_{i-1} - \bar{B}_i)}{\bar{P}_i(\hat{f}_i + \hat{h}_i)} - \bar{Q}_i \right] & \text{if } j = i \\ 0 & \text{otherwise} \end{cases} \quad (95)$$

From the chain rule, the i^{th} integral element sensitivity to (zero rate at the) j^{th} node is:

$$\frac{\partial I_i}{\partial P_j} = \sum_{k=1}^n \frac{\partial I_i}{\partial \bar{P}_k} \frac{\partial \bar{P}_k}{\partial R_j} = \frac{\partial I_i}{\partial \bar{P}_{i-1}} \frac{\partial \bar{P}_{i-1}}{\partial R_j} + \frac{\partial I_i}{\partial \bar{P}_i} \frac{\partial \bar{P}_i}{\partial R_j} \quad (96)$$

Finally, the protection leg sensitivity to the yield curve is

$$\frac{\partial PV_{\text{Protection Leg}}}{\partial R_j} = \frac{N(1 - RR)}{P(t_v)} \sum_{i=1}^n \frac{\partial I_i}{\partial R_j} \quad (97)$$

C.2 Premium Leg Sensitivity

C.2.1 Premiums Only

Again, the sensitivity of the premiums only is very straightforward; from equation 15 we can derive

$$\frac{\partial PV_{\text{Premiums only}}}{\partial R_j} = \frac{NC}{P(t_v)} \sum_{i=1}^M \Delta_i \frac{\partial P(t_i)}{\partial R_j} Q(e_i) \quad (98)$$

C.2.2 Accrual Paid On Default

This can be written most compactly using $\epsilon(x)$ and its derivatives from equation 43. The sensitivity of this to the discount factor is:

$$\frac{\partial I_i^k}{\partial \bar{P}_j} = \begin{cases} \hat{h}_i \bar{Q}_{i-1} (\phi_1 - \phi_2) & \text{if } j = i - 1 \\ \frac{\bar{B}_{i-1}}{\bar{P}_i} \hat{h}_i \phi_2 & \text{if } j = i \\ 0 & \text{otherwise} \end{cases} \quad (99)$$

where the variables ϕ_1 and ϕ_2 were defined in equation 92. The sensitivity to the zero rates at the nodes is:

$$\frac{\partial I_i^k}{\partial R_j} = \frac{\partial I_i^k}{\partial \bar{P}_{i-1}} \frac{\partial \bar{P}_{i-1}}{\partial R_j} + \frac{\partial I_i^k}{\partial \bar{P}_i} \frac{\partial \bar{P}_i}{\partial R_j} \quad (100)$$

Again, to use the original ISDA formula, simply replace s_k with $\tilde{s}_k = s_k - 1/730$, and to use the Markit fix formula, drop the first terms from equation 92 (i.e. set $t_{i-1} - s_k = 0$).

D The ISDA Model Yield Curve Bootstrap

The yield curve is constructed from money market rates (spot (L)ibor rates) with maturities out to 1Y (typically these are 1M, 2M, 3M, 6M, 9M and 12M), and swap rates with maturities out to 30Y (typically the swap rates are 2Y-10Y, 15Y, 20Y, 25Y and 30Y). The combined maturities (chronologically ordered) form the nodes of the yield curve. For n_M money market rates and n_S swap rates, we have a total of $n = n_M + n_S$ nodes.

The curve is defined from a particular *spot date*. Nominal maturity dates are measured from the spot date and (actual) maturity dates are adjusted to good business days (normally using modified following).

For curve nodes corresponding to money market rates, the discount factor at that node is given by

$$P_i = (1 + m_i t_i)^{-1}$$

where m_i is the relevant money market rate and t_i is the year fraction measured from the spot date using the correct DCC for that currency. This has no dependence on any other node in the yield curve.

D.1 Swap Pricing

The ISDA model makes the standard single curve textbook assumption that the present value (on a unit notional) of the floating leg of an IRS can be written as

$$1 - P(T)$$

where $P(T)$ is the discount factor from the last payment (the maturity of the swap). The nominal payments on the floating leg are calculated stepping backwards from the maturity in integer multiples of the payment interval; these are then adjusted to be good business days. The PV of the fixed leg is given by

$$C \sum_{i=1}^M \Delta_i P(t_i)$$

where C is the fixed leg coupon (swap rate), $\Delta_i = DCC_{\text{fixed leg}}(t_{i-1}, t_i)$, and $P(t_i)$ are the discount factors for the fixed payments, of which there are M . Since the swap must have zero PV, we require

$$(1 + C \Delta_M) P(T) + C \sum_{i=1}^{M-1} \Delta_i P(t_i) - 1 = 0$$

where $P(T) \equiv P(t_M)$. The value $P(T)$ will correspond to a curve node, while the other discount factors are found by interpolation (given the ‘current’ node and all the previous ones) - there is no dependence on nodes forward of the current. We can find the value of $P(T)$ that solves this equation using one-dimensional root finding.⁴⁸

In summary, for nodes corresponding to money market rates we simply read out the value of the node, while all other nodes (corresponding to swaps) are found by bootstrapping along the curve in a very similar way to what is done for the credit curve.

⁴⁸We actually solve for the zero rate rather than the discount factor.

E Market Data Used in Examples

Unless stated otherwise, the examples in this paper use the following market data.

The trade date is 13-Jun-2011. The instruments used to build the yield curve (deposit and swap rates) are snapped the day before (12-Jun-2011), and the spot date is three days forward from this (i.e. 15-June-2011). The table 12 below shows the market data to build a Euro yield curve. The deposits (money market) use a ACT/360 DCC, while the fixed leg of the swap uses 30/360 DCC with annual payments. In the table we also show the year fractions (from ACT/365F DCC) and fitted zero rates measured from the CDS trade date.

Type	Tenor	Rate	Year Fraction	Zero Rates
Depo	1M	0.445%	0.088	0.451%
Depo	2M	0.949%	0.173	0.945%
Depo	3M	1.234%	0.258	1.232%
Depo	6M	1.776%	0.507	1.778%
Depo	9M	1.935%	0.756	1.937%
Depo	1Y	2.084%	1.008	2.082%
Swap	2Y	1.652%	2.014	1.629%
Swap	3Y	2.018%	3.011	1.998%
Swap	4Y	2.303%	4.008	2.287%
Swap	5Y	2.525%	5.011	2.512%
Swap	6Y	2.696%	6.011	2.688%
Swap	7Y	2.825%	7.011	2.822%
Swap	8Y	2.931%	8.016	2.934%
Swap	9Y	3.017%	9.014	3.024%
Swap	10Y	3.092%	10.014	3.104%
Swap	11Y	3.160%	11.014	3.178%
Swap	12Y	3.231%	12.014	3.256%
Swap	15Y	3.367%	15.016	3.407%
Swap	20Y	3.419%	20.022	3.451%
Swap	25Y	3.411%	25.027	3.421%
Swap	30Y	3.412%	30.033	3.411%

Table 12: The date use to build a Euro curve with a spot date of 15-Jun-2011. The fitted curve is show measured from the CDS trade date (13-Jun-2011).

The credit curve is built from six liquid CDS quotes (6M, 1Y, 3Y, 5Y, 7Y and 10Y) given as par spreads. Apart from quoting as par spreads, they are assumed to be standard CDSs - i.e. quarterly premium payments with a DCC of ACT/360; a following business-day convention (with weekend only calendar); and a prior coupon (accrual start date) of 21-Mar-2011 (so there are 85 days of accrued interest).⁴⁹ Table 13 shows the data along with the survival probabilities at the curve knots (pillars).

The generic CDS we price off these curves is a standard ISDA CDS with a coupon of 100bps.

⁴⁹In a previous version we had incorrectly stated the accrual start date as 20-Mar-2011 (which is a Sunday); thanks to Christian Fenger for pointing this out.

Tenor	Maturity	Spread (bps)	Year Fraction	Survival Probability
6M	20-Dec-11	79.27	0.521	0.99307
1Y	20-Jun-12	79.27	1.022	0.98645
3Y	20-Jun-14	122.39	3.022	0.93915
5Y	20-Jun-16	169.79	5.025	0.86259
7Y	20-Jun-18	192.71	7.025	0.78866
10Y	20-Jun-21	208.60	10.027	0.69051

Table 13: Data used to build the credit curve with a trade date of 13-Jun-2011. The year fractions (measured from the trade date) and survival probability of the knots (pillars) of the fitted curve are shown.

References

- [Ber11] Arthur M. Berd. A guide to modelling Credit Term Structure. In *The Oxford Handbook of Credit Derivatives*. Oxford, 2011.
- [Cha10] Geoff Chaplin. *Credit Derivatives*. Wiley Finance, 2010.
- [dB78] Carl de Boor. *A Practical Guide to Splines*. Springer, 1978.
- [GW08] A. Griewank and A. Walther. *Evaluating derivatives: principles and techniques of algorithmic differentiation*. SIAM, second edition, 2008.
- [Hen12] Marc Henrard. Adjoint Algorithmic Differentiation: Calibration and Implicit Function Theorem. *Journal of Computational Finance*, to appear., 2012.
- [Hul06] John C. Hull. *Options, futures, and other derivatives*. Prentice Hall, sixth edition, 2006.
- [ISD] 2003 ISDA Credit Derivatives Definitions.
- [ISD09] ISDA. ISDA Standard CDS Converter Specifications. Technical report, ISDA, 2009.
- [Iwa13] Yukinori Iwashita. Piecewise Polynomial Interpolations. Technical report, OpenGamma, 2013.
- [LR11] Alexander Lipton and Andrew Rennie, editors. *The Oxford Handbook of Credit Derivatives*. Oxford, 2011.
- [Mar09] Markit. The CDS Big Bang: Understanding the Changes to the Global CDS Contract and North American Conventions. Technical report, Markit, March 13, 2009.
- [Mar12] Markit. ISDA Model Accrual on Default fix. Technical report, Markit Group Limited, 2012.
- [O’k08] Dominic O’kane. *Modelling single-name and multi-name Credit Derivatives*. Wiley F, 2008.

- [PTVF07] William H. Press, Saul A. Teukolsky, William T. Vetterling, and Brian P. Flannery. *Numerical Recipes*. Cambridge University Press, 2007.
- [Res12] Quantitative Research. *Interest Rate Instruments and Market Conventions Guide*. OpenGamma, 1.0 edition, April 2012. Available at <http://docs.opengamma.com/display/DOC/Analytics>.
- [Roz09] Saha Rozenberg. A ‘Big Bang’ in the Credit Derivatives Universe. *Derivatives Week*, 2009.
- [Whi12] Richard White. Multiple Curve Construction. Technical report, OpenGamma, 2012. Available at docs.opengamma.com.

OpenGamma Quantitative Research

1. Marc Henrard. Adjoint Algorithmic Differentiation: Calibration and implicit function theorem. November 2011.
2. Richard White. Local Volatility. January 2012.
3. Marc Henrard. My future is not convex. May 2012.
4. Richard White. Equity Variance Swap with Dividends. May 2012.
5. Marc Henrard. Deliverable Interest Rate Swap Futures: Pricing in Gaussian HJM Model. September 2012.
6. Marc Henrard. Multi-Curves: Variations on a Theme. October 2012.
7. Richard White. Option pricing with Fourier Methods. April 2012.
8. Richard White. Equity Variance Swap Greeks. August 2012.
9. Richard White. Mixed Log-Normal Volatility Model. August 2012.
10. Richard White. Numerical Solutions to PDEs with Financial Applications. February 2013.
11. Marc Henrard. Multi-curves Framework with Stochastic Spread: A Coherent Approach to STIR Futures and Their Options. March 2013.
12. Marc Henrard. Algorithmic Differentiation in Finance: Root Finding and Least Square Calibration. January 2013.
13. Marc Henrard. Multi-curve Framework with Collateral. May 2013.
14. Yukinori Iwashita. Mixed Bivariate Log-Normal Model for Forex Cross. January 2013.
15. Yukinori Iwashita. Piecewise Polynomial Interpolations. May 2013
16. Richard White. The Pricing and Risk Management of Credit Default Swaps, with a Focus on the ISDA Model. September 2013
17. Yukinori Iwashita. Conventions for Single-Name Credit Default Swaps. December 2013
18. Arroub Zine-eddine. Brazilian Swaps. December 2013.
19. Arroub Zine-eddine. Inflation: Instruments and curve construction. January 2014.
20. Arroub Zine-eddine. Inflation caps and floors. January 2014.
21. Arroub Zine-eddine. Convexity adjustment for inflation derivatives. January 2014.
22. Richard White. Forward CDS, Indices and Options. March 2014
23. Yukinori Iwashita. Tree Option Pricing Models. April 2014

About OpenGamma

OpenGamma helps financial services firms unify their calculation of analytics across the traditional trading and risk management boundaries.

The company's flagship product, the OpenGamma Platform, is a transparent system for front-office and risk calculations for financial services firms. It combines data management, a declarative calculation engine, and analytics in one comprehensive solution. OpenGamma also develops a modern, independently-written quantitative finance library that can be used either as part of the Platform, or separately in its own right.

Released under the open source Apache License 2.0, the OpenGamma Platform covers a range of asset classes and provides a comprehensive set of analytic measures and numerical techniques.

Find out more about OpenGamma

www.opengamma.com

Download the OpenGamma Platform

developers.opengamma.com/downloads

Europe

OpenGamma
185 Park Street
London SE1 9BL
United Kingdom

North America

OpenGamma
125 Park Avenue
25th Floor, Suite 2525
New York, NY 10017
United States of America

