OpenGamma Quantitative Research

Forward CDS, Indices and Options

Richard White
Richard@opengamma.com

OpenGamma Quantitative Research n. 22 | First version: March 13, 2014; this version October 2, 2014

Abstract

This paper is a followup to *The Pricing and Risk Management of Credit Default Swaps, with a Focus on the ISDA Model* [Whi13]. Here we show how to price indices (portfolios of CDSs) from the calibrated credit curves of the constituent names, and how to adjust those curves to match the market price of a index (basis adjustment). We then show how to price forward starting single-name CDSs and indices, since these are the underlying instruments for options on single-name CDSs and indices. The pricing of these options is the main focus of this paper. The model we implement for index options was first described by Pedersen [Ped03], and we give full implementation details and examples. We discuss the common risk factors (the Greeks) that are calculated for these options, given various ways that they may be calculated and show results for some example options. Finally we show some comparisons between our numbers and those displayed on Bloomberg's CDSO screens. All the code used to generate the results in this paper is available as part of the open source release of the OpenGamma Platform.

Contents

1 Introduction

In a previous paper [Whi13] we discussed how to price a single-name CDS, and gave particular attention to the ISDA model. In this paper we assume the reader is familiar with the mechanics of standard single-name CDS, and only clarify rather than repeat what is in other sources. We show how to build on a CDS pricing model[1] to price forward CDS, CDS indices, options on single-name CDS (CDS options or default swaptions) and options on CDS indices (CDS index options).

In section 2 we briefly review standard single-name CDS and how they are priced from a known yield curve and credit curve. We do not discuss the calibration of credit curves from market quotes of CDS prices, as this is covered in our previous paper [Whi13] (in the context of the ISDA model). Section 3 introduces the forward starting CDS, and shows how (with minor modifications) these can be priced in the same way as a spot starting CDS. Section 4 discussed European options on single-name CDS (default swaptions), and shows that for certain contract specifications these can be priced within a Black framework. In section 5 we introduce portfolios of single-name CDS - *CDS indices*, and show how to adjust individual credit curves so that the intrinsic price of the index matches the market. Section 6 then deals with European options on CDS indices, and in particular discusses Pedersen's one factor model [Ped03] to price these options, as well as an approximation using the Black formula. Section 7 discusses numerical implementation and approximations and gives performance metrics, while section 8 introduces risk measures associated with index options - the Greeks. Finally section 9 discusses how options are priced from the Bloomberg CDSO screen, with examples comparing the results from those screens with our own calculations.

Appendix A gives a very brief overview of option pricing theory and the description of yield (discount) and credit curves, while appendix B gives a full list of terms used in this paper.

2 Review of Standard CDS

We use *today* and *trade date* interchangeably to mean the date on which we are valuing a trade,[2] and assign this a value, t. A spot CDS has some cost of entry (which can be positive or negative), which is paid on the *cash settlement date*,[3] which we call t_{cs}.

Standard CDS contracts have fixed coupons[4] which are nominally paid every three months on the IMM dates (the 20^{th} of March, June, September and December). They are price quoted[5] with a *clean price*, and the actual cash payment (made on the cash settlement date) is this clean price (i.e. PUF $\times$ notional) adjusted for the *accrued amount*. The accrued amount is the year fraction[6] between the *accrual start date* (the previous IMM date) and the *step-in date* (one day after the trade-date) multiplied by the coupon; the convention is that this is quoted as a negative number for the buyer of protection and positive for the seller. Let Δ be the year fraction and C

[1] We do not restrict ourselves to the ISDA model.

[2] Trade date is sometimes taken as the inception date of a derivative (a CDS or an option in this context), but we wish to value the derivative at any date up to expiry.

[3] Usually three working days after the trade date.

[4] These are 100 or 500 bps in North America and 25, 100, 500 or 100 bps in Europe.

[5] Single name CDS are quoted with Points-Up-Front (PUF), which is the clean value as a percentage of notional; Indices are quoted as (100-PUF).

[6] On a ACT/360 day count.

the coupon, then the *cash settlement amount* (for the buyer of protection) is

$$\text{cash settlement amount} = \text{quoted price} + \text{accrued} = \text{quoted price} - C\Delta$$

The buyer of protection pays this amount at t_{cs} to the seller of protection.[7] The *market value* is the cash settlement amount, risk free discounted (the few days) from the cash settlement date back to today.

To clarify, *clean price* is the market quoted price, and *dirty price* is the cash settlement amount (clean price plus accrued), both of which are amounts 'seen' on the cash settlement date rather than the trade date.

2.1 Pricing from Yield and Credit Curves

We assume for the rest of this paper that there exists appropriate[8] yield (discount) and credit curves. See appendix A for a brief overview of the curves and our previous paper [Whi13] for details of calibration (of the ISDA model). Where terms are common between this and the previous paper, we use the same notation - a full list of terms is given in appendix B.

2.1.1 The Protection Leg

The value at time $t_{cs}(\geq t)$ for protection between t and T where no default has occurred by t is

$$\begin{aligned} V_{\text{Prot}}(t,T) &= \mathbb{E}_t[e^{-\int_{t_{cs}}^{\tau} r_s ds}(1-R(\tau))\mathbb{I}_{\tau \leq T}|\tau > t] \\ &= \frac{(1-\hat{R}(T))}{P(t,t_{cs})}\int_t^T P(t,u)[-\dot{Q}(t,u)]du \end{aligned} \tag{1}$$

where $\hat{R}(T) = \mathbb{E}_t[R(\tau)]$ is the expected recovery rate (for a CDS with maturity T), $\dot{Q}(t,u) \equiv \frac{dQ(t,u)}{du}$ and we have implicitly set the notional to unity. Where there is no term structure of recovery rates, we just use R. The division by the discount factor $P(t,t_{cs})$ occurs because protection is from t but we value it at t_{cs}.

2.1.2 The Premium Leg

The full value (or dirty price) of the premium at t_{cs} is

$$V_{\substack{\text{Prem}\\\text{dirty}}}(t,T) = \frac{C}{P(t,t_{cs})}\sum_{i=1}^{M}\left[\Delta_i P(t,t_i)Q(t,e_i) + \eta_i \int_{\max(s_i,T_1)}^{e_i}(t-s_i)P(t,u)[-\dot{Q}(t,u)]du\right] \tag{2}$$

where C is the CDS coupon, s_i, e_i and t_i are the accrual start, end and payment time for the i^{th} accrual period,[9] and η_i is the day count adjustment factor.[10] The sum is over all the accrual periods between t and T, i.e. $s_1 \leq t$, $t_1 > t$ and $e_M = T$.

The clean price is then simply

$$V_{\substack{\text{Prem}\\\text{clean}}}(t,T) = V_{\substack{\text{Prem}\\\text{dirty}}}(t,T) + C\Delta \tag{3}$$

where Δ is the year fraction between s_1 and $t+1$ (the step-in date).

---

[7] Of course, if the amount is negative, the seller of protection actually pays the buyer.

[8] The yield curve is for the currency in which the CDS is priced and the credit curve is for the reference entity.

[9] Usually have $e_i = t_i$, except for the last accrual period for which the accrual end is not business-day adjusted.

[10] Ratio of the year fraction for the i^{th} accrual period measured with the accrual day count convention to the same period measured with the curve day count convention. In almost all cases this is just 365/360.

From a mathematical sense, the dirty value (of the premium leg) is the principle quantity, since it is the actual economic value of future liabilities, and is computed directly from the yield and credit curves; the clean value is then found by adding the accrued amount.[11] However the market treats the clean value as if it is the principle quantity, with the dirty value 'calculated' by subtracting the accrued amount. This also holds for other quantities that depend on the value of the premium leg, such as the price, annuity and spread, which are usually quoted as clean. This is perfectly sensible, since in reality the credit curve is not exogenously given, but calibrated from the market clean prices of CDS.

The *risky annuity* or *risky PV01* (RPV01) is the value of the premium leg per unit of spread, i.e. $C = 1$ in the above equations.[12] We will refer to it as simply the annuity. The relationship between clean and dirty annuity is clearly

$$A(t,T) = A_D(t,T) + \Delta \tag{4}$$

where we have dropped to subscript for the clean annuity.

2.1.3 Price and Spread

The clean price is given by

$$V(t,T) = V_{\text{Prot}}(t,T) - V_{\underset{\text{clean}}{\text{Prem}}}(t,T) = V_{\text{Prot}}(t,T) - CA(t,T) \tag{5}$$

The spread is defined as the value of the coupon, C that makes the clean price zero. So

$$S(t,T) = \frac{V_{\text{Prot}}(t,T)}{A(t,T)} \tag{6}$$

This allows us to write the price as

$$V(t,T) = [S(t,T) - C]\, A(t,T) \tag{7}$$

where both the spread and the annuity are computed from the known yield and credit curves.

2.2 Constant Hazard Rate Pricing

The convention for CDS post-'Big-Bang'[13] is to quote the spread by pricing with a constant hazard rate; this has always been the convention for CDS indices. This spread is known as the *quoted spread* or *flat spread*, and differs from the par spread computed from the full credit curve (see White (2013) or Rozenberg (2009) [Whi13, Roz09] for further discussion).

If we set $Q(t,T) = \exp(-(T-t)\lambda)$, the CDS price can be written explicitly as a function of the (constant) hazard rate, λ,

$$V(t,T|\lambda) = \bar{V}_{\text{Prot}}(t,T|\lambda) - C\bar{A}(t,T|\lambda) \tag{8}$$

[11] The accrued amount is $C\Delta$ so is always positive (or zero).

[12] Strictly the RPV01 is the value per basis point of spread, so is one ten-thousandth of the value per unit of spread.

[13] In 2009 ISDA issued the 'Big Bang' protocol in an attempt to restart the market by standardising CDS contracts [Roz09].

where the 'bar' indicates quantities that have been calculated from a constant hazard rate. This is a monotonic function, so we may also find the inverse function;[14] given the clean price we can find λ - this is the implied hazard rate for a given market price of a CDS. Similarly, we may write the spread as a function of λ

$$\bar{S}(t,T|\lambda) \equiv \frac{\bar{V}_{\text{Prot}}(t,T|\lambda)}{\bar{A}(t,T|\lambda)} \tag{9}$$

or think of λ as a function of $\bar{S}$. This means we may write the annuity as

$$\bar{A}(t,T|\bar{S}) \equiv \bar{A}(t,T|\bar{S}(t,T)) = \bar{A}(t,T|\lambda(\bar{S}(t,T))) \tag{10}$$

which means: find λ (by root finding) that gives this spread, then compute the annuity for that λ. The CDS price can now be expressed as an explicit function of spread

$$V(t,T|\bar{S}) = (\bar{S} - C) \times \bar{A}(t,T|\bar{S}). \tag{11}$$

Figure 1 shows the annuity as a function of spread for a short (1Y) and long (10Y) expiry CDS. The annuity can be approximated extremely well by a low order polynomial.

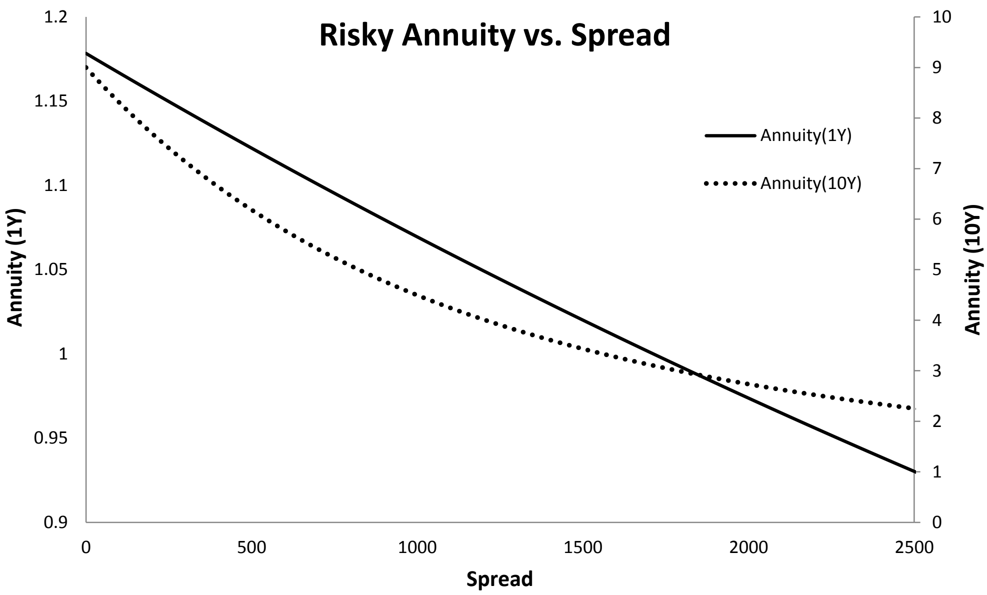

Figure 1: The annuity as function of spread for a 1Y and 10Y CDS. All calculations are preformed using a constant hazard rate.

It must be kept in mind that this a quoting convention only - $V(t,T) \leftrightarrow \bar{S}(t,T)$ - the constant hazard rate 'curve' cannot be used to price any other maturity CDS. Also the annuity, protection

[14]If the recovery rate, R is too high, there may be no solution for a given V.

leg and spread calculated this way will generally not equal their equivalent values calculated from the full credit curve, i.e.

$$\begin{aligned}\bar{V}_{\text{Prot}}(t,T) &\neq V_{\text{Prot}}(t,T)\\ \bar{A}(t,T|\bar{S}) &\neq A(t,T)\\ \bar{S}(t,T) &\neq S(t,T).\end{aligned}$$

However, since the price must be the same whether a full or flat credit curve is used, it does follow that

$$\begin{aligned}\bar{V}_{\text{Prot}}(t,T) - C\bar{A}(t,T|\bar{S}) &= V_{\text{Prot}}(t,T) - CA(t,T)\\ (\bar{S}(t,T) - C)\bar{A}(t,T|\bar{S}) &= (S(t,T) - C)A(t,T).\end{aligned} \tag{12}$$

3 Forward Starting CDS

A forward starting payer CDS entered into at time t will give protection against the default of an obligor for the period $T_e > t$ to T_m, in return for premium payments in that period. If the obligor defaults before the start of protection (i.e. $\tau < T_e$), the contract cancels worthless. This can easily be replicated by entering a long protection CDS with maturity T_m and a short protection position with maturity T_e for the same obligor. Since for standard CDS, premiums are all paid on the same dates (adjusted IMM dates) with standard coupons, these will exactly cancel up to T_e,[15] leaving only the coupons between T_e and T_m to pay. Furthermore, if a default occurs before T_e the protection payments will exactly cancel.[16] So the value of our forward CDS is given by

$$V(t,T_e,T_m) = V(t,T_m) - V(t,T_e). \tag{13}$$

Since this value depends on the difference in price of a long and short dated CDS, its price will be sensitive to the shape of the credit curve.

In practice, CDS prices may not be available for a particular forward protection span,[17] so we need to be able to price these instruments from the calibrated yield and (appropriate) credit curve.

3.1 Pricing from Curves

Assume that today, t, I enter a forward contract where I agree that at some future *expiration date*, T_e, I will enter a (payer) CDS contract with maturity T_m ($> T_e$) in exchange for a payment of $F - C\Delta$ made at the *exercise settlement date* $t_{es} \geq T_e$ provided that the reference entity has not defaulted by T_e. The exercise settlement date should correspond to the cash settlement date of a spot CDS entered at T_e.[18]

At T_e, a spot starting CDS with maturity T_m will have a (clean) price for the cash settlement date t_{es} given by $V(T_e,T_m)$. Therefore the mark-to-market value of the forward contract at T_e is

$$P(T_e,t_{es})\left[V(T_e,T_m) - F\right]\mathbb{I}_{\tau > T_e}$$

[15]This is not entirely true, since the final coupon period (for a standard CDS) has an extra day of accrued interest, so will not exactly match the coupon of the longer CDS.

[16]This may need some legal specifications, so that the defaulted bond is delivered from the short position before it is delivered to the long position.

[17]In particular if T_e does not correspond to an IMM date.

[18]Normally three working days after T_e.

and its expected value at t is

$$\mathrm{Fwd}(t,T_e,T_m)=\mathbb{E}_t\left[\frac{\mathbb{I}_{\tau>T_e}}{\beta_t(t_{es})}(V(T_e,T_m)-F)\right]$$

where $\beta_t(s)=\exp(\int_t^s r_u du)$ is the money market numeraire.[19] The value of F that makes this zero is the forward (clean) price; it is given by

$$\begin{aligned}F(t,T_e,T_m)&=\frac{1}{P(t,t_{es})Q(t,T_e)}\mathbb{E}_t\left[\frac{\mathbb{I}_{\tau>T_e}}{\beta_t(t_{es})}V(T_e,T_m)\right]\\&=\frac{1}{P(t,t_{es})Q(t,T_e)}V(t,T_e,T_m)\end{aligned}\tag{14}$$

where $V(t,T_e,T_m)$ is the expected present value of the forward CDS. This may be written as

$$\begin{aligned}V(t,T_e,T_m)&=\mathbb{E}_t\left[\frac{\mathbb{I}_{\tau>T_e}}{\beta_t(t_{es})}(V_{\mathrm{Prot}}(T_e,T_m)-CA(T_e,T_m))\right]\\&=V_{\mathrm{Prot}}(t,T_e,T_m)-CA(t,T_e,T_m).\end{aligned}\tag{15}$$

We will show below how to calculate the expected present value of the protection leg and annuity from the calibrated yield and credit curves. In general we define

$$\begin{aligned}Y(t,T_e)&=\mathbb{E}_t\left[\frac{\mathbb{I}_{\tau>T_e}}{\beta_t(t_{es})}Y(T_e)\right]\\&=\mathbb{E}_t\left[\frac{P(T_e,t_{es})\mathbb{I}_{\tau>T_e}}{\beta_t(T_e)}Y(T_e)\right]\end{aligned}\tag{16}$$

to be the fair value today of a contract to pay $Y(T_e)$ at $t_{es}\geq T_e$ provided the reference entity has not defaulted by T_e. Implicit in the definition is the fact that the reference entity has not defaulted by t. From the definition we have

$$Y(T_e,T_e)=P(T_e,t_{es})Y(T_e).\tag{17}$$

A useful result is when $Y(T_e)=1$, i.e. 1 unit of currency is paid at T_{es} provided no default occurred before T_e - this a defaultable bond with a payment delay. We argue in appendix A that this is given by

$$\mathbb{E}_t\left[\frac{P(T_e,t_{es})}{\beta_t(T_e)}\mathbb{I}_{\tau>T_e}\right]=P(t,t_{es})Q(t,T_e).\tag{18}$$

For a suitable choice of Y, the quantity $Y(t,T_e)$ is an asset which is positive provided that $t<\tau$, so it can be used as a numeraire in what is known as a *survival measure* [Sch04].

3.2 Valuing the Forward Protection

Protection between T_e and T_m can be statically replicated by buying protection to T_m and selling protection to T_e. Therefore

$$\begin{aligned}V_{\mathrm{Prot}}(t,T_e,T_m)&=P(t,t_{cs})\left(V_{\mathrm{Prot}}(t,T_m)-V_{\mathrm{Prot}}(t,T_e)\right)\\&=(1-\hat{R}(T_m))\int_t^{T_m}P(t,u)[-\dot{Q}(t,u)]du\\&\quad-(1-\hat{R}(T_e))\int_t^{T_e}P(t,u)[-\dot{Q}(t,u)]du.\end{aligned}\tag{19}$$

[19] It is the amount of money in a 'risk-free' bank account at time $s\geq t$ if one unit of currency in invested at t.

If there is no term structure of expected recovery rates (i.e. $\hat{R}(T_e) = \hat{R}(T_m) = R$), then this simplifies to

$$V_{\text{Prot}}(t, T_e, T_m) = (1 - R) \int_{T_e}^{T_m} P(t, u)[-\dot{Q}(t, u)]du, \tag{20}$$

so the present value of the forward starting protection can be computed directly from the known yield and credit curve (which have been calibrated to spot instruments).

3.3 Valuing the Forward Premium Leg

The present value of the dirty forward starting annuity is

$$\begin{aligned} A_D(t, T_e, T_m) &\equiv \mathbb{E}\left[e^{-\int_t^{t_{es}} r_s ds} \mathbb{I}_{\tau > T_e} A_D(T_e, T_m))\right] \\ &= \mathbb{E}\left[e^{-\int_t^{t_{es}} r_s ds} \mathbb{I}_{\tau > T_e} (A(T_e, T_e) - \Delta)\right] \\ &= A(t, T_e, T_m) - P(t, t_{es}) Q(t, T_e)\Delta. \end{aligned} \tag{21}$$

By writing this as $A_D(t, T_e, T_m) = A_D(t, T_m) - A_D(t, T_e)$, or by evaluating the terms in the expectation directly, one can show that the present value of the forward starting annuity is given by

$$\begin{aligned} A(t, T_e, T_m) = &\sum_{i=1}^{M} \left[\Delta_i P(t, t_i) Q(t, e_i) + \eta_i \int_{\max(s_i, T_e)}^{e_i} (t - s_i) P(t, u)[-\dot{Q}(t, u)]du\right] \\ &+ P(t, t_{es}) Q(t, T_e)\Delta \end{aligned} \tag{22}$$

where s_i, e_i and t_i are the accrual start and end times and payment time for the i^{th} accrual period.[20] The sum is over all the accrual periods between T_e and T_m, i.e. $s_1 \leq T_e$, $t_1 > T_e$ and $e_M = T_m$.

3.4 The Forward Spread

The forward spread is the coupon that makes the expected value of the forward CDS zero, that is

$$S(t, T_e, T_m) = \frac{V_{\text{Prot}}(t, T_e, T_m)}{A(t, T_e, T_m)}. \tag{23}$$

Since the discounting from t_{es} to T_e cancels,

$$S(T_e, T_e, T_m) = S(T_e, T_m) \tag{24}$$

so the forward spread observed T_e is just the spot spread. The present value of a forward starting CDS can be written as

$$\begin{aligned} V(t, T_e, T_m) &= \mathbb{E}_t\left[\frac{\mathbb{I}_{\tau > T_e}}{\beta_t(t_{es})} (S(T_e, T_m) - C) A(T_e, T_m)\right] \\ &= A(t, T_e, T_m) \mathbb{E}_t^{\mathbb{A}} \left[S(T_e, T_m) - C\right]. \end{aligned} \tag{25}$$

[20] As we mentioned earlier, usually have $e_i = t_i$, except for the last accrual period for which the accrual end is not business-day adjusted.

In the last line we have switched to the annuity measure (see appendix A for details). Since this can also be written as

$$V(t,T_e,T_m) = [S(t,T_e,T_m) - C]\,A(t,T_e,T_m) \tag{26}$$

we have that $\mathbb{E}_t^{\mathbb{A}}[S(T_e,T_m)] = S(t,T_e,T_m)$, that is, the forward spread is a Martingale in the annuity measure.[21]

3.4.1 Different Expressions for Forward Spread

The forward spread can also be expressed as

$$\begin{aligned} S(t,T_e,T_m) &= \frac{V_{\text{Prot}}(t,T_m) - V_{\text{Prot}}(t,T_e)}{A(t,T_m) - A(t,T_e)} \\ &= \frac{S(t,T_m)A(t,T_m) - S(t,T_e)A(t,T_e)}{A(t,T_m) - A(t,T_e)} \\ &= S(t,T_m) + \frac{A(t,T_e)}{A(t,T_m) - A(t,T_e)}(S(t,T_m) - S(t,T_e)) \end{aligned} \tag{27}$$

The final expression shows that, without explicitly calculating the annuities, if the term structure of spreads is upwards sloping (i.e. $S(t,T_m) > S(t,T_e)$), then the forward spread is greater than the spot spread (both with maturity T_m); the converse is true for a downwards sloping term structure of spreads.[22]

3.5 Forward Flat Spread

In the previous section we showed how to calculate a forward par spread, which depends on the full credit curve. The concept of a forward flat spread is less well defined.

We may write the present value of a forward starting CDS in terms of a flat spread (and annuity) as

$$V(t,T_e,T_m) = \mathbb{E}_t\left[\frac{P(T_e,t_{es})\mathbb{I}_{\tau>T_e}}{\beta_t(T_e)}\left(\bar{S}(T_e,T_m) - C\right)\bar{A}(T_e,T_m|\bar{S})\right]. \tag{28}$$

The present value of the (forward) flat annuity is

$$\bar{A}(t,T_e,T_m) = \mathbb{E}_t\left[\frac{P(T_e,t_{es})\mathbb{I}_{\tau>T_e}}{\beta_t(T_e)}\bar{A}(T_e,T_m|\bar{S})\right], \tag{29}$$

however, calculation of this value depends on the distribution of $\bar{S}$, so it cannot be calculated from the calibrated credit curve at t. Additionally, the flat annuity is not an asset in its own right and we cannot technically use it as a numeraire (even if we could calculate its value). We may however switch to the risky bond measure to write

$$F(t,T_e,T_m) = \mathbb{E}_t^{\mathbb{B}}\left[\left(\bar{S}(T_e,T_m) - C\right)\bar{A}(T_e,T_m|\bar{S})\right] \tag{30}$$

[21] This also follows from the fact that the spread is the ratio of an asset (the protection leg) to the numeraire (the annuity).

[22] Regardless of the shape of the credit curve, we always have $A(t,T_m) > A(t,T_e) \iff T_m > T_e$.

where the forward is defined in equation 14. Whatever the terminal distribution of $\bar{S}$ is in the risky bond measure, $\mathbb{B}$, it must satisfy the above equation. The standard approach is to assume it is log-normal, so

$$\bar{S}(T_e, T_m) = \bar{S}_0 \exp(Z\sigma\sqrt{T_e - t} - \sigma^2(T_e - t)/2) \tag{31}$$

where σ is the spread volatility, Z is a standard Gaussian random variable and $\bar{S}_0$ is a constant we must compute. The value of $\bar{S}_0$ is such that

$$\begin{aligned} F(t, T_e, T_m) &= \int_{-\infty}^{\infty} (\bar{S} - C)\bar{A}(T_e, T_m|\bar{S})e^{-z^2/2}dz \\ \text{where} \quad \bar{S} &= \bar{S}_0 \exp(z\sigma\sqrt{T_e - t} - \sigma^2(T_e - t)/2) \end{aligned} \tag{32}$$

and $\bar{A}(T_e, T_m|\bar{S})$ is calculated at t assuming deteministic interest rates. The integral can be commuted via Gauss-Hermite quadrature [PTVF07], and the value of $\bar{S}_0$ found by one-dimensional root finding (again [PTVF07]). Clearly

$$\bar{S}_0 = \mathbb{E}_t^{\mathbb{B}}\left[\bar{S}(T_e, T_m)\right] \tag{33}$$

so $\bar{S}_0$ is the expected value (under $\mathbb{B}$) of the forward flat spread, but it does depend on the extraneous parameter σ (spread volatility).

3.5.1 Flat ATM Forward Spread

A useful quantity for the options we will meet later, is the value of $\bar{S}$ that solves

$$F(t, T_e, T_m) = \left(\bar{S} - C\right)\bar{A}(T_e, T_m|\bar{S}). \tag{34}$$

This is of course just $\bar{S}_0$ for $\sigma = 0$. We call this the flat ATM forward spread, $\bar{K}_{ATM}$. This is simply a number that put in a formula[23] gives the correct forward price. It is not the expected value of the forward flat spread.

4 Default Swaptions

A default swaption is a European option that can be exercised at time $T_e > t$ to enter into a CDS for protection between T_e and T_m. A payer option exercises into a payer CDS and of course a receiver option exercises into a receiver CDS.

In a knockout option, if the reference entity defaults before the option expiry, the option is cancelled and there are no further cash flows. For a non-knockout payer default swaption, in the case of a default before expiry, the option holder would receive the default settlement amount.[24] It should be noted that since the crisis, options on single-name CDS are not a widely traded instrument, and the material below really serves as a precursor to the discussion of options on CDS indices (which are traded).

[23] The ISDA quoted spread to upfront calculation.

[24] In practice they would deliver a suitable defaulted bond at expiry in exchange for par. In some circumstances this would be cash settled instead.

4.1 Spread based Default Swaptions

In this case, if an option strike at a level K is exercised at T_e, one entered a bespoke CDS paying a coupon of K with accrual starting from the day after the exercise date (i.e T+1 so $\Delta = 0$), that is, the underlying is a pre-'Big-Bang' style CDS. When there was a market for options on single-name CDS, this was the standard contract specification.

The payoff in this case is

$$V_{\text{option}}(T_e, T_m) = P(T_e, t_{es})\mathbb{I}_{\tau > T_e} A(T_e, T_m)(\chi \left[S(T_e, T_m) - K\right])^+ \tag{35}$$

where $\chi = 1$ for a payer and $\chi = -1$ for a receiver option. The value of the option at time t is given by

$$V_{\substack{\text{option}\\\text{knockout}}}(t, T_e, T_m) = \mathbb{E}_t\left[\frac{P(T_e, t_{es})A(T_e, T_m)\mathbb{I}_{\tau > T_e}}{\beta_t(T_e)}(\chi \left[S(T_e, T_m) - K\right])^+\right]. \tag{36}$$

We may then use the annuity, $A(t, T_e, T_m)$, as a numeraire (see appendix A) and perform a change of measure, to write the option value as

$$V_{\substack{\text{option}\\\text{knockout}}}(t, T_e, T_m) = A(t, T_e, T_m)\mathbb{E}_t^{\mathbb{A}}\left[(\chi \left[S(T_e, T_m) - K\right])^+\right] \tag{37}$$

where $\mathbb{A}$ indicates we are in the risky annuity measure. As we have already discussed, the forward spread is a martingale in this measure.

The above argument is mathematically identical to that made for pricing interest rate swaptions. Any model that preserves the martingale property of the forward spread will produce arbitrage free option prices. The usual approach is to assume log-normal dynamics for S, i.e.

$$\frac{dS(u, T_e, T_m)}{S(u, T_e, T_m)} = \sigma(u)dW_u \tag{38}$$

which gives

$$V_{\substack{\text{option}\\\text{knockout}}}(t, T_e, T_m) = A(t, T_e, T_m)\text{Black}(S(t, T_e, T_m), K, T_e - t, \hat{\sigma}, \chi) \tag{39}$$

where

$$\begin{aligned}\hat{\sigma} &= \sqrt{\frac{1}{T_e - t}\int_t^{T_e} \sigma(u)^2 du}\\ \text{Black}(F, K, T, \sigma, \chi) &= \chi\left(F\Phi(\chi d_1) - K\Phi(\chi d_2)\right)\\ d_1 &= \frac{\ln\left(\frac{F}{K}\right) + \frac{\sigma^2 T}{2}}{\sigma\sqrt{T}} \quad d_2 = d_1 - \sigma\sqrt{T}\end{aligned} \tag{40}$$

is Black's formula [Bla76]. The ATM strike is given by $K_{ATM} = S(t, T_e, T_m)$.

While it is attractive to end up pricing with the familiar Black formula, there is no market for the underlying instrument, which makes hedging non-trivial, and therefore it is unlikely that this style of default swaption will trade in the future.

4.2 Price Based Default Swaptions

In order to have a liquid instrument as the underlying, the option must be price-based, which means, at T_e the option holder chooses whether to excise the option and pay the excise price, $G(K)$, in return for entering a standard CDS. The excise decision is made at T_e but the excise price is paid on the excise settlement date, t_{es} which corresponds to the settlement date of the underlying CDS. The excise price, $G(K)$, is clean, but what is actually paid at t_{es} is the dirty price, $G(K) - C\Delta$. This means the payoff at T_e is

$$\begin{aligned} V_{\substack{\text{payer}\\\text{knockout}}}(T_e) &= P(T_e, t_{es})\left[(S(T_e, T_m) - C)A(T_e, T_m) - C\Delta) - (G(K) - C\Delta)\right]^+ \\ &= P(T_e, t_{es})\left[(S(T_e, T_m) - C)A(T_e, T_m) - G(K)\right]^+ \end{aligned} \tag{41}$$

The accrual terms have cancelled in the payoff, but one must remember that the actual amount paid to exercise the option does include the accrued - this seems to be overlooked in the literature. The choice of $G(K)$ affects how simple this type of option is to price.

4.2.1 The Simple Case

The simplest case is when $G(K) = (K - C)A(T_e, T_m)$. That is, the payment exactly accounts for the difference in value of the premium leg for paying the actual coupon, C, rather than the nominal coupon, K. This gives a payoff that is identical to equation 35, and thus is priced in the same way. However, it is highly unlikely to have a exercise price of this form since it is not known until expiry, and depends on a subjective calculation of the annuity at T_e. Furthermore, as the exercise price is itself stochastic, one cannot simply hedge with the standard CDS of the same maturity as the underlying.

4.2.2 The Fixed Excise Price Case

The convention for index options is to set

$$G(K) = (K - C)\bar{A}(T_e, T_m|K)$$

that is, the annuity is computed using a spread level of K, which gives a payoff that is zero for $\bar{S}(T_e, T_m) \le K$ and is positive for $\bar{S}(T_e, T_m) > K$, so K is still a strike level. Technically $\bar{A}(T_e, T_m|K)$ depends on the yield curve observed at T_e, but in practice it is calculated on the trade date and treated as a known fixed payment. For some options (indices that trade on price) the excise price is given directly rather than quoted as a strike (which is then converted to a price). For the general case of an excise price G that is know at t, the option price is

$$\begin{aligned} V_{\substack{\text{option}\\\text{knockout}}}(t, T_e, T_m) &= \mathbb{E}_t\left[\frac{P(T_e, t_{es})\mathbb{I}_{\tau > T_e}}{\beta_t(T_e)}\left(\chi[(\bar{S}(T_e, T_m) - C\right)\bar{A}(T_e, T_m|\bar{S}) - G])^+\right] \\ &= P(t, t_{es})Q(t, T_e)\mathbb{E}_t^{\mathbb{B}}\left[\left(\chi[(\bar{S}(T_e, T_m) - C\right)\bar{A}(T_e, T_m|\bar{S}) - G])^+\right] \end{aligned} \tag{42}$$

where we have expressed the CDS price in terms of flat spread and annuity, and switched to the risky bond measure (see appendix A). Note, the survival probability, $Q(t, T_e)$, in the above expression is obtained from the full credit curve.

To compute the option price it is necessary to know the terminal distribution of $\bar{S}$ in this measure. If we assume log-normal dynamics as we did in section 3.5, then once $\bar{S}_0$ is calibrated the option price is easily calculated as

$$\begin{aligned} V_{\substack{\text{payer}\\\text{knockout}}}(t,T_e,T_m) &= P(t,t_{es})Q(t,T_e)\int_{-\infty}^{\infty}((\bar{S}-C)\bar{A}(\bar{S})-G)^{+}e^{-z^2/2}dz \\ &= P(t,t_{es})Q(t,T_e)\int_{z^*}^{\infty}(\bar{S}-C)\bar{A}(\bar{S})-G)e^{-z^2/2}dz \\ \text{with}\quad \bar{S} &= \bar{S}_0\exp(z\sigma\sqrt{T_e-t}-\sigma^2(T_e-t)/2) \end{aligned} \tag{43}$$

where z^* solves $(\bar{S}-C)\bar{A}(\bar{S})=G$. This integral must also be computed numerically, but since the payoff is not smooth, it is unwise to use Gauss-Hermite quadrature. It is better to use the second form with Simpson's rule or a higher order scheme. One further point to note is that while function $\bar{A}(\bar{S})$, which appears in the integrand, is smooth, it is computed via a root finding routine, therefore in implementations it is often approximated. This is discussed further in section 7.

4.2.3 Put-call Parity

The value of payer minus a receiver option is given by

$$\begin{aligned} V_{\substack{\text{payer}\\\text{knockout}}}(t,T_e,T_m)-V_{\substack{\text{receiver}\\\text{knockout}}}(t,T_e,T_m) &= \mathbb{E}_t\left[\frac{P(T_e,t_{es})\mathbb{I}_{\tau>T_e}}{\beta_t(T_e)}(V(T_e,T_m)-G(K))\right] \\ &= P(t,t_{es})Q(t,T_e)(F(t,T_e,T_m)-G(K)), \end{aligned} \tag{44}$$

so the ATM exercise price is given by $G(K_{ATM})=F(t,T_e,T_m)$, and the flat spread that solves this, is what we called the flat ATM forward spread.

4.3 No-knockout Options

A payer option can be structured without the knockout feature: In the event of a default before expiry, the option holder may deliver an appropriate defaulted bond at the expiry date in return for par. The value of this *front-end protection* is

$$V_{FEP}(t,T_e)=(1-R)(1-Q(t,T_e))P(t,t_{es}) \tag{45}$$

where the default settlement occurs at t_{es}. If the default is settled immediately rather than waiting to the option expiry, the front end protection is just worth the same as the protection leg of a (spot) CDS with maturity of T_e.[25]

If there is a default before expiry, then the front-end protection pays out and the rest of the option is worthless. If there is no default before expiry, the option is worth the same as the no knockout. The front-end protection is worthless for the holder of a receiver option (who would never choose to pay out in the event of a default). So we have

$$V_{\substack{\text{payer}\\\text{no-knockout}}}(t,T_e,T_m)=V_{FEP}(t,T_e)+V_{\substack{\text{payer}\\\text{knockout}}}(t,T_e,T_m). \tag{46}$$

[25] In a low interest rate environment and with expiries typically less than six months, the difference is tiny.

A complication to this picture is if the contract specifies that the excise price G is paid in exchange for default settlement. In this case the payoff is[26]

$$\left(\chi\left(V(T_e,T_m)\mathbb{I}_{\tau>T_e}+(1-R(\tau))\mathbb{I}_{\tau\leq T_e}-G\right)\right)^+.$$

It is not possible to decompose this into a knock-out option plus front-end protection. Also, the no-knockout receiver is worth more than the knockout receiver since the option holder will excise following a default provided that $G>(1-R(\tau))$, where $R(\tau)$ is the realised recovery rate. The option price can be decomposed as

$$V_{\text{no-knockout}}(t,T_e,T_m)=V_{\text{knockout}}(t,T_e,T_m)+\mathbb{E}_t\left[\frac{P(T_e,t_{es})\mathbb{I}_{\tau\leq T_e}}{\beta_t(T_e)}\left(\chi(1-R(\tau)-G)\right)^+\right]. \quad (47)$$

Computation of the second term requires a model for the realised recovery rate. A treatment is given in Martin(2012) [Mar12].

5 CDS Indices

CDS Indices are portfolios of single-name CDS, with typically between 40 and 125 entries. They pay coupons every three months (i.e. on adjusted IMM dates) to the buyer of the index (who is the seller of protection), and in return, the index buyer must receive any (qualifying) defaulted bonds in return for par. For constancy with single-name CDS, we use *payer* to mean the payer of coupons (i.e. the seller of the index/buyer of protection), and *receiver* to mean the recipient of coupon payments (i..e the buyer of the index/seller of protection).

The names in the index are given a weight, with w_j the weight of the j^{th} entry. If the notional of a index is N, then the notional of the j^{th} entry is just $N_j=Nw_j$, and

$$\sum_{j=1}^{J}w_j=1 \quad (48)$$

where J is number of names in the index. Most indexes are equally weighted, so $w=\frac{1}{j}$.

Let $J_D(t)$ be the number of names that have expired by some time t, then

$$\sum_{j=1}^{J}\mathbb{I}_{\tau_j>t}w_j=f(t)<1 \quad \Longleftrightarrow \quad J_D(t)>0 \quad (49)$$

where $f(t)$ is the *index factor* (at time t), and τ_j is the default time of the j^{th} name; so the index factor is less than one iff any names in the index have defaulted.

The amount paid on the i^{th} coupon payment date is

$$NC\Delta_i f(t_i)$$

where C is the index coupon and Δ_i is the year fraction for the i^{th} accrual period.

If the j^{th} name defaults at τ_j, the nominal amount paid out at τ_j is

$$Nw_j(1-R_j)$$

[26]The accrual still cancels provided accrual is paid on default.

where R_j is the realised recovery rate for that name. If this default occurs in the i^{th} accrual period, the buyer of protection must pay the accrual-on-default of

$$Nw_j\Delta(s_i,\tau_j)C$$

where $\Delta(s_i,\tau_j)$ is the year fraction between the accrual start date[27] and the default. So the cash settlement for the default is

$$Nw_j\left((1-R_j)-\Delta(s_i,\tau_j)C\right).$$

Provided that the definition of a default is the same, the payments above will exactly match those from owning the CDS (for the same maturity) on all the individual constituent of the index, each with notionals of $N_j = Nw_j$. So the index can, in principle, be statically replicated.

5.1 Index Value

The price of an index is given as a percentage in a form such that the cash settlement amount at $t_{cs} \geq t$[28] is

$$\text{cash settlement amount} = N \times f \times (1 - \text{quoted price} - C\Delta).$$

So the index quoted price and the points upfront (PUF) are related by

$$1 - \text{quoted price} = PUF$$

For the remainder of this paper we will assume that the index quoted price is expressed as PUF.

CDS indices are conventionally priced using a constant hazard rate, λ. This means the index annuity at time t for an index that expires at T may be written as $\bar{A}_I(t,T|\lambda)$ and computed exactly as in the single-name case. The value of the protection leg (for a unit notional) is given by

$$\bar{V}_{I,\text{Prot}}(t,T|\lambda) = (1-R_I)\lambda\int_t^T P(t,s)e^{-\lambda(s-t)}ds \tag{50}$$

where R_I is the index recovery rate, which is a fixed number with little relation to the recovery rates of the constituent names.[29] The index PUF is

$$PUF_I(t,T) = \bar{V}_{I,\text{Prot}}(t,T|\lambda) - C\bar{A}_I(t,T|\lambda) \tag{51}$$

and its value (for a payer) is

$$V_{I,\text{payer}}(t,T) = Nf(t)PUF_{I,\text{payer}}(t,T) \tag{52}$$

where $f(t)$ is the index factor at time t, and the protection and annuity are computed for unit notional. The index value will change sign depending on whether one is considering the payer or receiver, but the PUF is always given as shown above.

For a given index PUF, we can find the implied hazard rate λ. The maximum value of the protection leg is $(1-R_I)$ (i.e. immediate default) which is also the maximum value of the index PUF, so if the PUF of the index is above $(1-R_I)$ it is not possible to find a hazard rate and thus an equivalent spread.[30]

[27] which is also the previous payment, except for the first accrual period.
[28] The standard is three working days.
[29] It is usually 40%, but is 30% for CDX.NA.HY and other indices made of lower credit quality entities. Basically it is a number that lets you convert between PUF and a quoted spread.
[30] This is just an effect of using a recovery rate that is too high.

5.2 Index Spread

The index spread is defined as

$$\bar{S}_I(t,T|\lambda) \equiv \frac{\bar{V}_{I,\text{Prot}}(t,T|\lambda)}{\bar{A}_I(t,T|\lambda)} \tag{53}$$

where we have made it explicit that it is a function of the hazard rate (which in turn is implied from the index price). For a given λ this is independent of the index factor, $f(t)$.

Given a spread, one can also imply the hazard rate that satisfies the above equation. So, like the single-name case, the annuity can be expressed as a function of the spread, $\bar{A}_I(t,T|\bar{S}_I)$, and the index PUF and value may be written as

$$\begin{aligned} PUF_I(t,T|\bar{S}_I) &= (\bar{S}_I(t,T) - C)\bar{A}_I(t,T|\bar{S}_I) \\ V_{I,\text{payer}}(t,T|\bar{S}_I) &= Nf(t)(\bar{S}_I(t,T) - C)\bar{A}_I(t,T|\bar{S}_I). \end{aligned} \tag{54}$$

It is worth mentioning again that the (flat) annuity function $\bar{A}_I(t,T|\bar{S}_I)$ as we have defined it, is always for a unit notional index: This is so it can be regarded as a function of the index (flat) spread, $\bar{S}_I$, only, and not other state variables.[31]

An index value may be quoted as a PUF or a (flat) spread. Just as in the single-name CDS case, the conversion is via a constant hazard rate (i.e. the ISDA upfront model).

5.3 Intrinsic Value

Since an index is a portfolio of single-names, we can value it using the individual credit curves of the constituent names; this is the *intrinsic value*. At some time, t, the clean value of the index (for the protection buyer/seller of the index) is

$$\begin{aligned} V_{I,\text{intrinsic}}(t,T) &= N\sum_{j=1}^{J} \mathbb{I}_{\tau_j>t} w_j \left[V_{\text{Prot},j}(t,T) - CA_j(t,T)\right] \\ &= N\sum_{j=1}^{J} \mathbb{I}_{\tau_j>t} w_j \left(S_j(t,T) - C\right) A_j(t,T) \\ &= N\sum_{j=1}^{J} \mathbb{I}_{\tau_j>t} w_j \left(\bar{S}_j(t,T) - C\right) \bar{A}_j(t,T|\bar{S}_j). \end{aligned} \tag{55}$$

The last form shows that we can compute this quantity using the flat spreads of the constituent names. Of course spreads (of either form) are not defined for defaulted names, but the default indicator means they are not counted, and in practice the sum is taken over just the undefaulted names.

It is unlikely that this value will equal the market quoted value - this is known as the *index basis*. The biggest effect will come from a difference between what constitutes a credit event in the index and for a single-name CDS: In North America restructuring is not a credit event for indices while it is for standard CDS [O'k08]. The second effect is liquidity: An index will be more liquid than its constituent parts, so it is more expensive to buy protection on all the individual names than to sell the index (for the same total notional) [Ped03].

[31] Obviously it also depends on the payment schedule of the premium leg, but that is fixed.

5.4 Intrinsic Annuity and Spread

The intrinsic value of the index (eqn. 55) may be written in the form

$$V_{I,\text{intrinsic}}(t,T) = N\left(\frac{\sum_{j=1}^{J}\mathbb{I}_{\tau_j>t}w_jS_j(t,T)A_j(t,T)}{\sum_{j=1}^{J}\mathbb{I}_{\tau_j>t}w_jA_j(t,T)} - C\right)\sum_{j=1}^{J}\mathbb{I}_{\tau_j>t}w_jA_j(t,T)$$

Comparing this with equation 54, we see that we can define the intrinsic annuity and intrinsic spread as

$$A_{I,\text{intrinsic}}(t,T) \quad = \sum_{j=1}^{J}\mathbb{I}_{\tau_j>t}w_jA_j(t,T) \tag{56}$$

$$S_{I,\text{intrinsic}}(t,T) \quad = \frac{\sum_{j=1}^{J}\mathbb{I}_{\tau_j>t}w_jS_j(t,T)A_j(t,T)}{\sum_{j=1}^{J}\mathbb{I}_{\tau_j>t}w_jA_j(t,T)} \tag{57}$$

The difference between the index spread and the intrinsic spread is another way of seeing the index basis. These definitions give the intrinsic index value as

$$V_{I,\text{intrinsic}}(t,T) = N(S_{I,\text{intrinsic}}(t,T) - C)A_{I,\text{intrinsic}}(t,T). \tag{58}$$

The form is identical to equation 54 except there is no explicit index factor - this is because it is implicit in the definition of the intrinsic annuity, $A_{I,\text{intrinsic}}(t,T)$.[32] These intrinsic values could alternatively be calculated using flat spreads and annuities, which would result in sightly different values. However, since we are ultimately interested in computing forward index values, we will always work with values computed from the full credit curves of the constituent names.

5.5 Credit Curve Adjustment

The individual credit curves will be used to calculate a forward price for the index, and in turn price options on the index. As with all models to price derivatives, it is essential that the model prices the underlaying exactly. To this end we need to adjust the individual credit curves in some way so that the intrinsic value matches the market value. There are several ways to do this (methods are discussed in Pedersen [Ped03] and O'Kane [O'k08]).

5.5.1 Matching a Single Index

Since there is only one number, the market price, to match, we can only have a single adjustment factor. We choose to adjust the forward hazard rates directly, and apply a constant multiplication factor to the rate.[33] Each credit curve is transformed thus:

$$\begin{aligned}\tilde{h}_j(t,T) &= \alpha h_j(t,T)\\ \tilde{Q}_j(t,T) &= (Q_j(t,T))^{\alpha}\\ \tilde{\Lambda}_j(t,T) &= \alpha\Lambda_j(t,T)\end{aligned} \tag{59}$$

where $\alpha = \mathcal{O}(1)$. Since the ISDA standard model uses log-linear interpolation of survival probabilities,[34] these curve properties are preserved by applying the multiplication factor to the value of the knots, i.e. $\tilde{\Lambda}_j^i = \alpha\Lambda_j^i$, where Λ_j^i is the i^{th} knot of the j^{th} curve.

[32] we could normalise the intrinsic annuity (i.e. divide by $f(t)$) so that the index factor is explicit in the equal for an intrinsic value, but we find this more cumbersome.

[33] A multiplicative factor is preferred over an additive one, since it ensures that forward hazard rates remain positive, and is more in keeping with the dynamics of hazard rates - volatility tends to scale with the rate.

[34] Equivalently, piecewise constant hazard rates.

We may write the adjusted intrinsic value as a function of α, $V_{I,\text{intrinsic}}(t,T|\alpha)$, and find value of α that satisfies

$$V_{I,\text{intrinsic}}(t,T|\alpha) = V_{I,\text{market}}(t,T)$$

using a one-dimensional root finder.

We could, instead, make the adjustment so that the intrinsic spread matches the index spread. However the index values would not (exactly) match, and since the spread is just a quoting convention - it is the index value that matters - this is not desirable.

5.5.2 Matching Multiple Indices

What we mean here, is matching several indices that have the same constituents but different expiries (e.g. iTraxx Europe 3Y, 5Y, 7Y and 10Y). In general we could transform the forward hazard rates of each credit curve thus:

$$\begin{aligned}
\tilde{h}_j(t,T) &= \alpha(T)h_j(t,T) \\
\tilde{Q}_j(t,T) &= \exp\left(-\int_t^T \alpha(u)h_j(t,u)du\right) \\
\tilde{\Lambda}_j(t,T) &= \frac{1}{T-t}\int_t^T \alpha(u)h_j(t,u)du
\end{aligned} \tag{60}$$

and find the single curve $\alpha(T)$ such that the intrinsic price for each index equals the market value. The simplest way to do this is to make $\alpha(u)$ piecewise constant, i.e.

$$\alpha(u) = \alpha_i \text{ for } T_{i-1} \leq u < T_i$$

where T_i is the maturity of the i^{th} index. We can then *bootstrap* α by solving for each index maturity in turn.

For the ISDA model this turns out to be quite easy. The first step is to add additional knots to the credit curves at the index expiries (if these knots are not already present),[35] at a level interpolated from the existing points - this has no effect on any survival probabilities computed off the curves.

The next step is to multiply all the knots less than or equal to T_1 on all the curves by α_1, then all the knots greater that T_{i-1} and less than or equal to T_i by α_i and finally all the knots greater than T_{n-1} by α_n, where n is the number of index maturities. The values of α are found in turn by solving for the index value.

We finish this section with an example of this adjustment in practice. We choose the CDX.NA.HY.21-V1 index on 13-Feb-2014. The rates used to build the yield curve and the index prices are given in table 10. Data on par spreads at standard pillars and recovery rates for the 100 constituents are provided by Markit - we do not reproduce them here. With this data we can build the 100 individual credit curves and calculate the intrinsic price of the index at four different terms. We show the market and intrinsic prices in table 1 together with the equivalent quoted spreads. For the most liquid 5Y index the difference is less than one basis point between the market and intrinsic quoted spread.

Once we adjust the curves as described above, of course the adjusted intrinsic prices will match the market prices. It is instructive to see what effect this has on the par spreads of the

[35] An x year CDS has maturity x years from the next IMM date, while a x year index has maturity x years plus three months from the previous roll date, so the two maturities may not match.

	PUF		Quoted Spread	
Term	Intrinsic	Market	Intrinsic	Market
3Y	-7.248	-7.56	234.770	223.955
5Y	-7.591	-7.62	323.427	322.789
7Y	-6.576	-5.71	381.069	396.038
10Y	-7.433	-6.52	394.141	406.448

Table 1: The intrinsic and market prices for the CDX.NA.HY.21-V1 index on 13-Feb-2014.

constituents. In table 2 below we show original and adjusted par spreads[36] for four names in the index with very different credit qualities and spread term structures.

	Original Par Spreads			Adjusted Par Spreads		
	6M	5Y	10Y	6M	5Y	10Y
The AES Corp	19.0	187.0	262.0	17.9	195.9	271.6
Bombardier Inc.	29.0	319.0	412.0	27.3	333.1	426.2
J C Penney Co Inc	1594	1633	1519	1501	1647	1543
TX Competitive Elec Hldgs Co LLC	47657	31568	27110	44880	29129	25927

Table 2: The original par spreads, and the implied par spreads after credit curve adjustment for three names in the CDX.NA.HY.21-V1 index. Note, we only show three pillar points (6M, 5Y and 10Y) for space.

In this example, the spreads at the long end (5 years and more) increase while the spreads at the short end decrease - this would be expected from what we saw in table 1. The exception to this is the highly distressed *TX Competitive Elec Hldgs Co*, which has a one year survival probability of just 3.6%; in this case the 10Y spread is actually driven by the short end of the credit curve, which is adjusted down.

For the remainder of this note, unless explicitly stated, it is assumed that the intrinsic credit curves have been adjusted to match index prices before any values are computed from them. With that in mind we will refer to the intrinsic index spread and annuity as just $S_I(t,T)$ and $A_I(t,T)$ - these should not be confused with the flat index spread and annuity, $\bar{S}_I(t,T)$ and $\bar{A}_I(t,T|\bar{S})$.

6 Options on CDS Indices

Without loss of generality, we set the (initial) index notional to one (i.e. $N = 1$) and drop the term.

[36]Recall we adjust the credit curves directly , so these are the par spreads implied from the new credit curves.

6.1 Forward CDS Index

The (clean) value of the index observed at some time $T_e > t$ and valued at t_{es} is

$$\begin{aligned} V_I(T_e, T_m) &= f(T_e)(\bar{S}_I(T_e, T_m) - C)\bar{A}_I(T_e, T_m|\bar{S}_I) \\ &= (S_I(T_e, T_m) - C)A_I(T_e, T_m). \end{aligned} \tag{61}$$

We now consider the present value of this index using the (adjusted) credit curves of the constituent names. This gives

$$\begin{aligned} V_I(t, T_e, T_m) &= \mathbb{E}_t\left[\frac{P(T_e, t_{es})}{\beta(T_e)}\sum_{j=1}^{J}\mathbb{I}_{\tau_j > T_e} w_j A_j(T_e, T_m)(S_j(T_e, T_m) - C)\right] \\ &= \sum_{j=1}^{J}\left(w_j \mathbb{E}_t\left[\frac{\mathbb{I}_{\tau_j > T_e} P(T_e, t_{es}) A_j(T_e, T_m)}{\beta(T_e)}(S_j(T_e, T_m) - C)\right]\right) \\ &= \sum_{j=1}^{J} w_j A_j(t, T_e, T_m)\mathbb{E}_t^{\mathbb{A}_j}\left[S_j(T_e, T_m) - C\right] \\ &= \sum_{j=1}^{J} w_j A_j(t, T_e, T_m)(S_j(t, T_e, T_m) - C) \\ &= \sum_{j=1}^{J} w_j V_{\text{prot},j}(t, T_e, T_m) - C\sum_{j=1}^{J} w_j A_j(t, T_e, T_m). \end{aligned} \tag{62}$$

Either of the last two forms are readily calculated from the (adjusted) calibrated credit curves of the undefaulted constituents at t.

6.2 Forward Spread

We may define the forward intrinsic protection leg and annuity as

$$\begin{aligned} V_{\text{prot},I}(t, T_e, T_m) &= \sum_{j=1}^{J} w_j V_{\text{prot},j}(t, T_e, T_m) \\ A_I(t, T_e, T_m) &= \sum_{j=1}^{J} w_j A_j(t, T_e, T_m) \end{aligned} \tag{63}$$

and naturally the forward (intrinsic) index spread as

$$S_I(t, T_e, T_m) = \frac{\sum_{j=1}^{J} w_j V_{\text{prot},j}(t, T_e, T_m)}{\sum_{j=1}^{J} w_j A_j(t, T_e, T_m)} \tag{64}$$

which allows us to write the present value of the forward index as

$$\begin{aligned} V_I(t, T_e, T_m) &= V_{\text{prot},I}(t, T_e, T_m) - C A_I(t, T_e, T_m) \\ &= (S_I(t, T_e, T_m) - C)A_I(t, T_e, T_m). \end{aligned} \tag{65}$$

6.3 Default Adjusted Index Value

If the holder of a payer option on an index exercises that option at T_e, they enter a payer CDS index at time T_e, paying the coupon C, until expiry at T_m. In addition, any defaults between the start of the option and T_e are settled at t_{es}. The clean value of the index plus the default settlement at t_{es} is

$$V_I^D(T_e, T_m) = f(T_e)\left[\bar{S}_I(T_e, T_m) - C\right]\bar{A}_I(T_e, T_m|\bar{S}_I) + \sum_{j=1}^{J}\mathbb{I}_{\tau_j \leq T_e} w_j(1 - R_j) \tag{66}$$

This is known as the *default-adjusted forward portfolio swap price* [Ped03] or just the default-adjusted price. The full value of the index plus default settlement at T_e includes the accrued (on the remaining index) and the accrual-on-default (on the defaulted entries) so is given by

$$V_I^D(T_e, T_m) - \Delta C$$

6.3.1 Intrinsic Value

The default-adjusted price also has an intrinsic value given by

$$V_I^D(T_e, T_m) = \sum_{j=1}^{J} w_j\left[(1 - R_j)\mathbb{I}_{\tau_j \leq T_e} + \mathbb{I}_{\tau_j > T_e}\left(S_j(T_e, T_m) - C\right)A_j(T_e, T_m)\right] \tag{67}$$

and its expectation under the $\mathbb{T}$-forward measure is

$$\begin{aligned} F_I^D(t, T_e, T_m) &\equiv \mathbb{E}_t^{\mathbb{T}}\left[V_I^D(T_e, T_m)\right] \\ &= \sum_{j=1}^{J} w_j\mathbb{I}_{\tau_j \leq t}(1 - R_j) + \sum_{j=1}^{J} w_j\mathbb{I}_{\tau_j > t}(1 - R_j)(1 - Q_j(t, T_e)) \\ &\quad + \frac{V_I(t, T_e, T_m)}{P(t, t_{es})}. \end{aligned} \tag{68}$$

We call this the *ATM forward price.* The first term in the expression is just value from the names in index that have already defaulted by t; the second term is the expected value from defaults between t and T_e; and the final term is expected value of the index (equation 62). All these terms can be calculated from the (adjusted) credit curves of the constituent single-names.

At expiry, the ATM forward price is equal to the default-adjusted (index) price, i.e.

$$F_I^D(T_e, T_e, T_m) = V_I^D(T_e, T_m).$$

6.3.2 Homogeneous Pool Approximation

If the pool is completely homogeneous, meaning the recovery rates and credit curves of the constituents are all identical, then the expected default-adjusted price reduces to

$$\begin{aligned} F_I^D(t, T_e, T_m) &= (1 - f(t))(1 - R) + f(t)\left[(1 - R)(1 - Q(t, T_e))\right] \\ &\quad + \frac{f(t)}{P(t, t_{es})}\left(V_{\text{prot}}(t, T_e, T_m) - CA(t, T_e, T_m)\right) \end{aligned} \tag{69}$$

The protection leg and annuity are calculated for a full index from the single credit curve. The single credit curve, the *index curve*, can be built from index quotes[37] for different terms, and hence the value of $F_I^D(t, T_e, T_m)$ estimated without knowing anything about the constituents names. Even if the pool is highly heterogeneous, approximating it this way and ignoring the underlying names, can still produce fairly accurate results. We show examples of this when we look at Greeks and for comparisons with Bloomberg.

6.4 Option Payoff

The exercise price of the option is given (contractually) by

$$G(K) = (K - C)\bar{A}(T_e, T_m|K)) \tag{70}$$

where $\bar{A}(T_e, T_m|K)$ is the flat annuity for a spread level of K. This is quoted clean; the actual amount paid at the exercise settlement date, t_{es}, to exercise the option is $G(K) - \Delta C$.

The option payoff at T_e may be written as

$$\begin{aligned} \text{option payoff} &= P(T_e, t_{es})(\chi[V_I^D(T_e, T_m) - \Delta C - (G(K) - \Delta C)])^+ \\ &= P(T_e, t_{es})(\chi[V_I^D(T_e, T_m) - G(K)])^+. \end{aligned} \tag{71}$$

We see that the accrual terms cancel, but this is purely because the accrual is included in the (full) exercise price.[38]

If there have been no defaults by the exercise date, then the payoff is exactly zero if the index spread equals the strike. Figure 2 shows the payoff against index spread for a payer and receiver option with different numbers of defaults at expiry.

Formally, the option value is

$$V_{I,\text{option}}(t, T_e, T_m) = \mathbb{E}_t\left[\frac{P(T_e, t_{es})}{\beta_t(T_e)}(\chi[V_I^D(T_e, T_m) - G(K)])^+\right]. \tag{72}$$

The payoff depends both on the index spread and the number of defaults at expiry, so in principle we need to jointly model all the constituent names of the index. Following the normal change of measure rules, we also have

$$V_{I,\text{option}}(t, T_e) = P(t, t_{es})\mathbb{E}_t^{\mathbb{T}}\left[(\chi[V_I^D(T_e, T_m) - G(K)])^+\right] \tag{73}$$

where $\mathbb{T}$ is the measure associated with the numeraire $P(\cdot, t_{es})$. The excise price, $G(K)$ is strictly not known advance as it depends (weakly) on the yield curve seen at T_e. However it is usual to set $G(K)$ to its expected value at t since the variance from this is negligible next to the variance of V_I^D.

What is required is a model for the distribution of $V_I^D(T_e, T_m)$ under the $\mathbb{T}$ forward measure, which preserves its expectation given in equation 68. If the distribution of $V_I^D(T_e, T_m)$ is given by $\rho^{\mathbb{T}}(F)$ such that

$$F_I^D(t, T_e, T_m) = \int_{-\infty}^{\infty} F\rho^{\mathbb{T}}(F)dF \tag{74}$$

[37] For an index quoted on spread, these should first be converted to upfront amounts before bootstrapping the credit curve.

[38] This is something that is simply not mentioned in the literature we have reviewed.

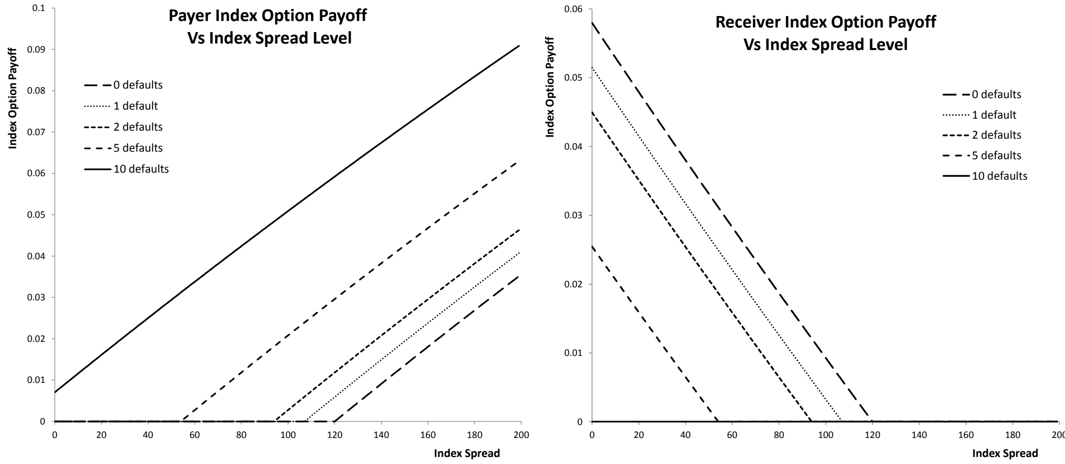

Figure 2: Index option payoff against index spread for various numbers of defaults by option expiry. The option is to enter an on-the-run 5Y index. The initial size of the index is 100 and the recovery rate of all constituents is 40%. If 10 or more defaults have occurred, then one would always exercise the payer option (never exercise the receiver option) regardless of the level of the index spread at expiry.

then the option price is

$$V_{I,\text{option}}(t,T_e) = P(t,t_{es}) \int_{-\infty}^{\infty} \left(\chi\left[F - G(K)\right]\right)^{+} \rho^{\mathbb{T}}(F)dF \tag{75}$$

Up until this point the analysis has been exact, and any distribution that satisfies equation 74 will produce arbitrage free prices.

6.5 Put-Call relation

The value of a receiver option is

$$V_{I,\text{Receiver}}(t,T_e,T_m) = P(t,t_{es})\mathbb{E}_t^{\mathbb{T}}\left[(G(K) - V_I^D(T_e,T_m))^{+}\right] \tag{76}$$

So the put-call relation is

$$\begin{aligned} V_{I,\text{Payer}} - V_{I,\text{Receiver}} &= P(t,t_{es})\mathbb{E}_t^{\mathbb{T}}\left[V_I^D(T_e,T_m) - G(K)\right] \\ &= P(t,t_{es})\left(F_I^D(t,T_e,T_m) - G(K)\right) \end{aligned} \tag{77}$$

So once the payer (receiver) price is calculated, the receiver (payer) price is available (almost) for free.

The ATM strike is defined as the strike level that makes the payer and receiver options have the same value; in this case this is given by

$$G(K_{ATM}) = F_I^D(t,T_e,T_m) \tag{78}$$

which justified us calling $F_I^D(t,T_e,T_m)$ the ATM forward value. This also allows us to define an ATM forward spread as the value of K_{ATM} that solves the above equation.

6.6 Default Adjusted Forward Index Spread

This is often referred to as just the forward index spread, however it should not be confused with either the forward index spread, $S_I(t, T_e, T_m)$ or the ATM forward spread $K_{ATM}(t, T_e, T_m)$ we defined earlier. It is defined as the ratio of the protection (including default settlement) to the forward annuity, so

$$\begin{aligned} S_I^D(t, T_e, T_m) &= \frac{P(t, t_{es}) \left(\sum_{j=1}^J w_j (1 - R_j)(1 - Q_j(t, T_e)) \right) + \sum_{j=1}^J w_j V_{\text{prot},j}(t, T_e, T_m)}{\sum_{j=1}^J w_j A_j(t, T_e, T_m)} \\ &= \frac{P(t, t_{es}) \left(\sum_{j=1}^J w_j (1 - R_j)(1 - Q_j(t, T_e)) \right) + V_{\text{prot},I}(t, T_e, T_m)}{A_I(t, T_e, T_m)} \end{aligned} \tag{79}$$

This of course is undefined in the armageddon scenario that the entire pool defaults by t. The default-adjusted forward spread is a martingale in the measure with the intrinsic annuity, $A_I(t, T_e, T_m)$, as the numeraire.

The present value of the default-adjusted index may be written as

$$\begin{aligned} V_I^D(t, T_e, T_m) &= \mathbb{E}_t \left[\frac{P(T_e, t_{es})}{\beta_t(T_e)} (S_I^D(T_e, T_m) - C) A_I(T_e, T_m) \right] \\ &= A_I(t, T_e, T_m) \mathbb{E}_t^{\mathbb{A}} \left[(S_I^D(T_e, T_m) - C) \right] \\ &= A_I(t, T_e, T_m)(S_I^D(t, T_e, T_m) - C) \end{aligned} \tag{80}$$

This is a useful result for the approximate option pricing formula we discuss in section 6.7.2.

6.6.1 Homogenous Pool

For a homogenous pool, the default-adjusted forward index spread may be written as

$$S_I^D(t, T_e, T_m) = \frac{P(t, t_{es}) \left(\frac{1 - f(t)}{f(t)} + (1 - R)(1 - Q(t, T_e)) \right) + V_{\text{prot}}(t, T_e, T_m)}{A(t, T_e, T_m)} \tag{81}$$

6.7 Option Pricing Models

6.7.1 The Pedersen Model

The following strategy is due to Pedersen [Ped03] - it is very similar to what was done in section 4.2.2. Let the terminal distribution of $V_I^D(T_e, T_m)$ under $\mathbb{T}$ be a function of a pseudo index spread, $\bar{X}$ (call default-adjusted forward portfolio spread in [Ped03]), which absorbs the payments from the defaulted names.

$$\begin{aligned} V_I^D(T_e, T_m) &= (\bar{X} - C)\bar{A}(T_e, T_m | \bar{X}) \\ \bar{X} &= \bar{X}_0 \exp\left(-\frac{\sigma^2}{2}(T_e - t) + \sigma Z \sqrt{T_e - t} \right) \end{aligned} \tag{82}$$

where Z is a standard Gaussian random variable, the annuity is calculated on the full index (i.e. no defaults, index factor, $f = 1$), and the value of $\bar{X}_0$ must be chosen so that equation 74 is satisfied.[39]

[39] A process for the default-adjusted spread, X_t, must by definition jump at defaults, so a log-normal distribution for X_{T_e} may be lacking enough kurtosis.

Writing $V_I^D(T_e,T_m)$ as a function of $\bar{X}_0$ and Z, $V_I^D(T_e,T_m|\bar{X}_0,Z)$, we have

$$\int_{-\infty}^{\infty} V_I^D(T_e,T_m|\bar{X}_0,Z)\phi(Z)dZ = F_I^D(t,T_e,T_m) \tag{83}$$

where $\phi(Z)$ is the Gaussian PDF. Since $V_I^D(T_e,T_m|\bar{X}_0,Z)$ is a smooth function of Z, the integral can be performed by Gauss-Hermite quadrature [PTVF07], with a small number of nodes.[40]

We root find for the value of $\bar{X}_0$ - the function that we are finding the root of involves a numerical integral, and the integrand is computed via a second root find (i.e. find the annuity for a given spread level). So we have a three level nest of numerical routines. Fortunately the roots can be found in a few iterations and (as already discussed) the integral can be computed with only a small number of function calls, so the overall cost of finding $\bar{X}_0$ is not too onerous. Of course, $\bar{X}_0$ depends on the volatility, σ, so must be recomputed for each value it takes.

The option value is now given by

$$\begin{aligned}
V_{\text{option}}(t,T_e,T_m) &= P(t,t_{es})\int_{-\infty}^{\infty}\left(\chi\left[V_I^D(T_e,T_m|\bar{X}_0,Z)-G(K)\right]\right)^{+}\phi(Z)dZ \\
&= P(t,t_{es})\mathbb{I}_{\chi=1}\int_{Z^*}^{\infty}\left(V_I^D(T_e,T_m|\bar{X}_0,Z)-G(K)\right)\phi(Z)dZ \\
&+ P(t,t_{es})\mathbb{I}_{\chi=-1}\int_{-\infty}^{Z^*}\left(G(K)-V_I^D(T_e,T_m|\bar{X}_0,Z)\right)\phi(Z)dZ
\end{aligned} \tag{84}$$

where Z^* solves $V_I^D(T_e,T_m|\bar{X}_0,Z)=G(K)$. This integral must also be computed numerically, however here we cannot use Gauss-Hermite quadrature since the function multiplying the Gaussian density is not well represented by a polynomial (it is discontinuous in its first derivative). It is better to use the second form (after solving for Z^*) and compute it using Simpsons rule[41] (or some higher order scheme) [PTVF07].

As an example of this in practice, we show in figure 3 the option price (for a spread based strike) where between 0 and 3 names in the index have already defaulted. The option has one month to expiry, and the index (which is a randomly generated example) initially had 100 (equally weighted) names.

6.7.2 A Modified Black Formula

We may write the option price in terms of the default-adjusted forward index spread:

$$\begin{aligned}
V_{\text{option}}(t,T_e) &= \mathbb{E}_t\left[\frac{P(T_e,t_{es})}{\beta_t(T_e)}(\chi[(S_I^D(T_e,T_m)-C)A_I(T_e,T_m)-G(K)])^{+}\right] \\
&= A_I(t,T_e,T_m)\mathbb{E}_t^{\mathbb{A}}\left[(\chi[S_I^D(T_e,T_m)-\tilde{K}])^{+}\right] \\
\text{where}\quad \tilde{K} &= C+\frac{G(K)}{A_I(T_e,T_m)}
\end{aligned} \tag{85}$$

This expression is exact, but since the modified strike, $\tilde{K}$ depends on the intrinsic index annuity at T_e, we cannot use this directly. However, if we make an approximation and set the modified

[40] We have found that as little as seven is sufficient.
[41] An upper limit must be chosen; Z between 6 and 8 would normally suffice.

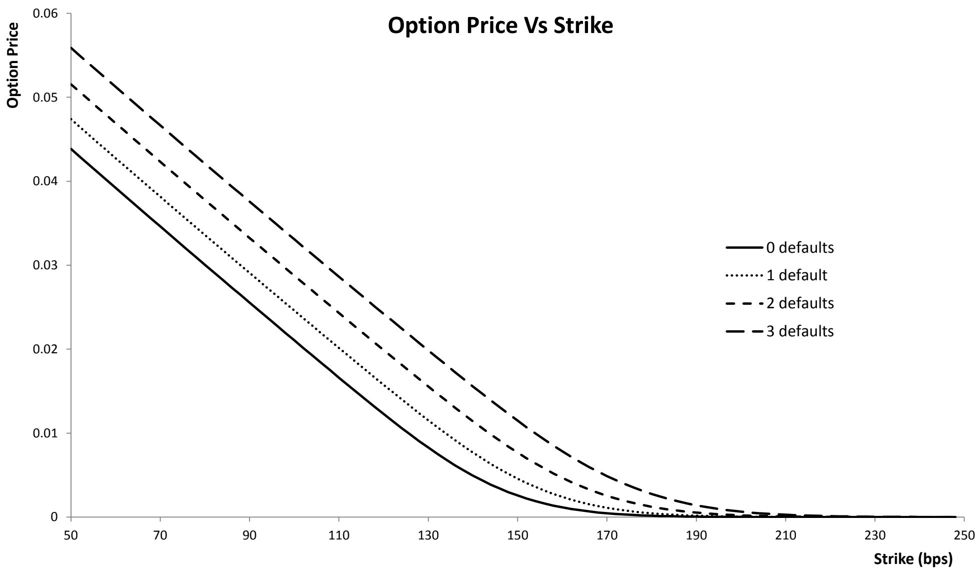

Figure 3: Payer option price when a number of names have defaulted from the index. The option is one month from expiry and the index initially had 100 names.

strike to be

$$\tilde{K} = C + \frac{G(K)P(t, t_{es})}{A_I(t, T_e, T_m)} \tag{86}$$

which can be calculated at time t, then since S_I^D is a Martingale in the annuity measure, the option can be priced with the Black formula. That is

$$V_{\text{option}}(t, T_e) \approx A_I(t, T_e, T_m)\text{Black}(S_I^D(t, T_e, T_m), \tilde{K}, T_e - t, \hat{\sigma}, \chi). \tag{87}$$

The put-call relationship is given by

$$\begin{aligned} V_{I,\text{Payer}} - V_{I,\text{Receiver}} &= A_I(t, T_e, T_m)\left(S_I^D(t, T_e, T_m) - \tilde{K}\right) \\ &= A_I(t, T_e, T_m)\left(S_I^D(t, T_e, T_m) - C\right) - P(t, t_{es})G(K) \\ &= V_I^D(t, T_e, T_m) - P(t, t_{es})G(K) \\ &= P(t, t_{es})\left(F_I^D(t, T_e, T_m) - G(K)\right) \end{aligned} \tag{88}$$

so this modified Black formula does preserve the put-call relationship, which is critical for any option pricing model.

We should not expect that for the same volatility level, σ, we obtain the same option prices from the Pedersen model as this modified Black formula. There are two reasons for this; firstly and most obviously, the modified Black formula is an approximation since we ‘freeze’ the modified strike at today’s value to make the expectation tractable; secondly, and more subtly, the volatility in Pedersen’s model is the log-normal volatility of a flat pseudo spread, while the volatility here is the log-normal volatility of the defaulted adjusted forward spread - these are not the same thing, so given them equal volatility will not correspond to equal option prices. The clearest way to demonstrate this is to produce option prices for a range of strikes with the Pedersen model using a flat volatility, then find the implied volatility for the modified Black model - that is, the volatility parameter that must be input into equation 87 to produce the same option price. Figure 4 shows a plot of this implied volatility when the volatility of the flat pseudo spread is 40%. There are two options: one with one week and the other with three months to expiry. The two resultant volatility 'smiles' are extremely close.

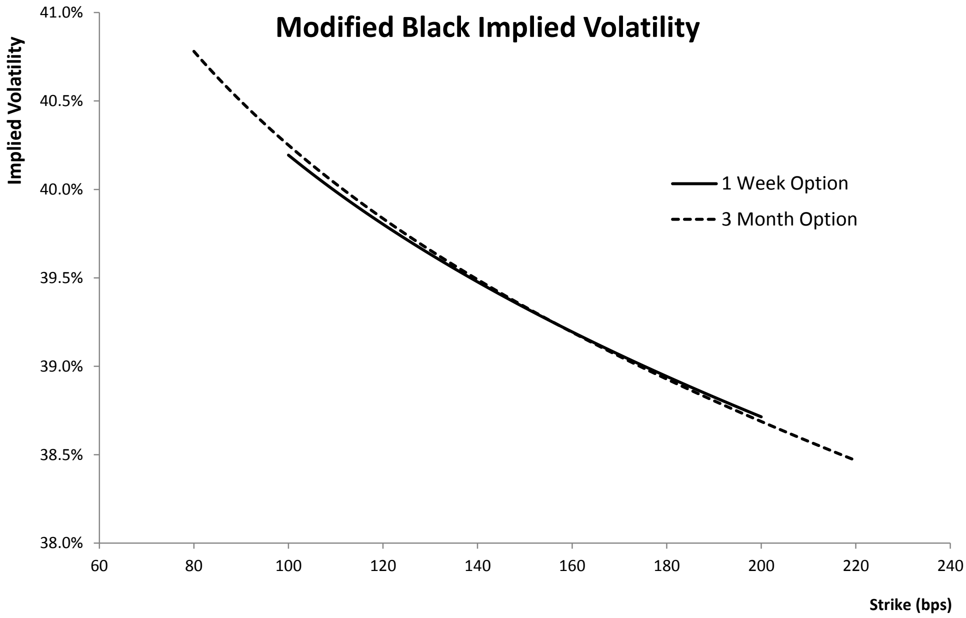

Figure 4: The implied volatility using the modified Black formula when option prices have been computed using the Pedersen model. The volatility used to generate the prices is 40%.

In this example the index coupon is 100bps, and when the strike is at this level, the exercise price, $G(K)$, is zero and the modified strike, $\tilde{K}$, is just C, so the modified Black formula is exact. We find the implied volatility for a strike of 100bps to be close to 40% (40.19% and 40.24% respectively), so the residual difference is because the two models use different spreads.

6.8 The Armageddon Scenario

The armageddon scenario is the theoretically possible, but unlikely, event that every name in an index defaults before the expiry of the option. It is claimed that this is incorrectly ignored in index option pricing. If the index annuity is used as a numeraire (as in the modified Black formula), then this can become zero while the option payoff is positive, which makes the option price theoretically unsound. However, while the default of every name means that forward spread given above is undefined, this has no effect on the pricing of index options we have presented using Pedersen's model.

The value of $V_I^D(T_e, T_m)$ in the event of a total default of the index by T_e is given by

$$V_I^D(T_e, T_m) = \sum_{j=1}^{J} w_j(1 - R_j)$$

which comes purely from the default settlement. The option payoff is

$$\text{armageddon option payoff} = P(T_e, t_{es})(\chi[\sum_{j=1}^{J} w_j(1 - R_j) - G(K)])^+ \tag{89}$$

which is also well defined since the calculation of the exercise price, $G(K)$, does not depend on the number of names remaining in the index (even if this is zero). We price the option by taking the expectation under the $\mathbb{T}$ forward measure which does not disappear in armageddon scenario. In fact the armageddon scenario simply represents the largest value $V_I^D(T_e, T_m)$ can take. Under Pedersen's model, we have

$$\lim_{X \to \infty} (X - C)\bar{A}(T_e, T_m|X) = 1 - R_I$$

so provided that $R_I \approx \sum_{j=1}^{J} w_j R_j$, the model included the armageddon scenarios in its states of the world.[42]

Some attempts have been made to estimate the probability of armageddon from tranche prices [BM09], but these do not exist for many indices, and where they do the results are somewhat suspect.

6.8.1 Extreme Strike Value

The limit of the exercise price as $K \to \infty$ is $1 - R_I$. It is often assumed that the payer option price will go to zero in this limit. However the realised value of $V_I^D(T_e, T_m)$ can exceed this limit (even without full default of the index) if there are enough defaults with realised recovery rates lower than expected (so that the default settlement is higher than expected), so the limit of the payer option price as $K \to \infty$ should be a (very) small positive number. Pedersen's model does not capture this effect. In the modified Black case, the default adjusted forward spread, S_I^D, can go arbitrarily high, so the option will have some (very low) value even for $K \to \infty$.

7 Numerical Implementation

Key to a fast implementation of Pedersen's model is the calculation of the flat annuity function $\bar{A}(t, T|\bar{S})$. The market standard way to do this is the ISDA upfront model, in which an intermediary constant hazard rate is solved for (with a root finding algorithm), then this is used to compute the annuity. An often used approximation is to assume that the coupons on the premium are paid continually. For constant hazard rate pricing this gives the annuity as

$$A(t, T|\lambda) = \frac{\eta}{P(t, t_c)} \int_t^T P(t, u)e^{-\lambda u} dt \tag{90}$$

where η is the ratio of the year fraction of the interval measured using the accrual day count convention, to the interval measured using the curve day count convention.[43] With the annuity

[42] We are not suggesting that the model gets the (risk neutral) probability of armageddon correct, just that it is an available state.

[43] In most cases it will be simply $365/360 \approx 1.0139$.

calculated this way, the spread is given by

$$S = \frac{(1-R)\lambda}{\eta}. \tag{91}$$

This gives us the hazard rate for a given spread level without root finding. The η term is normally missing in the *credit triangle*, however it greatly improves the estimate of λ at no real additional computational cost. The approximated hazard rate can then be used in equation 90 to give the annuity. Since the ISDA model uses log-linear interpolation for the discount factor [Whi13], the integral can be split up at the knots of the yield curve and computed as the sum of simple expressions; we call this approximation 1. We may also use the approximated hazard rate with a full calculation of the annuity (i.e. coupon paid every three months with accrual on default); we call this approximation 2.

Figure 5 shows the annuity for a 10Y CDS computed exactly and using the two credit triangle approximations. Approximation 2 is better for low spreads, while approximation 1 becomes better for high spreads. We prefer approximation 2, although the choice is not critical,[44] as the effect of the calculation of the annuity on the final option price is small since much of the variation is absorbed in the calibration of $\bar{X}_0$.

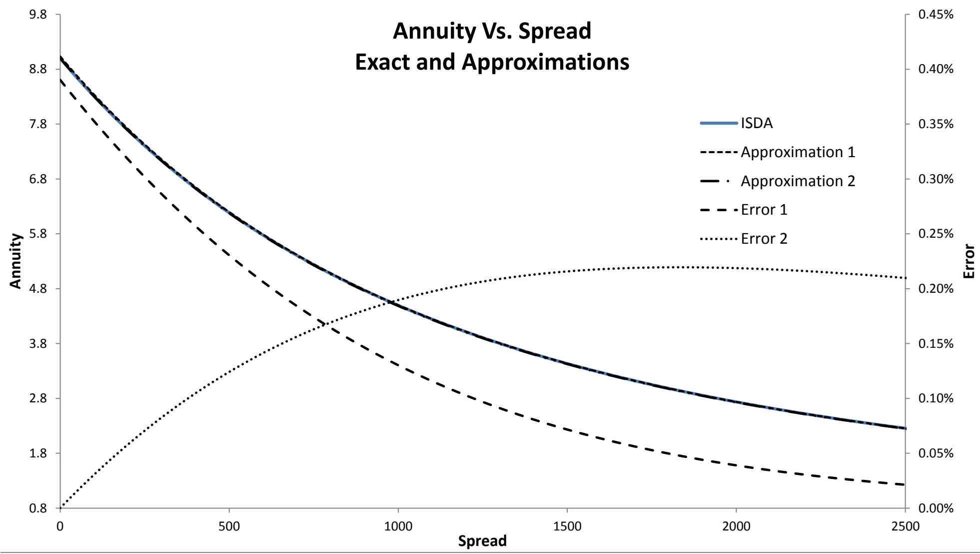

Figure 5: The annuity for 10Y CDS computed exactly and using the two credit triangle approximations. The three curves cannot be distinguished visually, and the relative errors are shown on the right hand scale.

Figure 6 shows out-the-money (OTM) option prices against strike for a one month option. The index has 125 names (none defaulted) and the individual recovery rates and credit curves have been randomly generated. Shown is the option price using both the ISDA model and the credit triangle approximation to compute the annuity. The error is lowest for ATM options at

[44]We could switch between the two approximations, but since this would mean a jump in annuity at the boundary we would need to split the integral in two to maintain high accuracy.

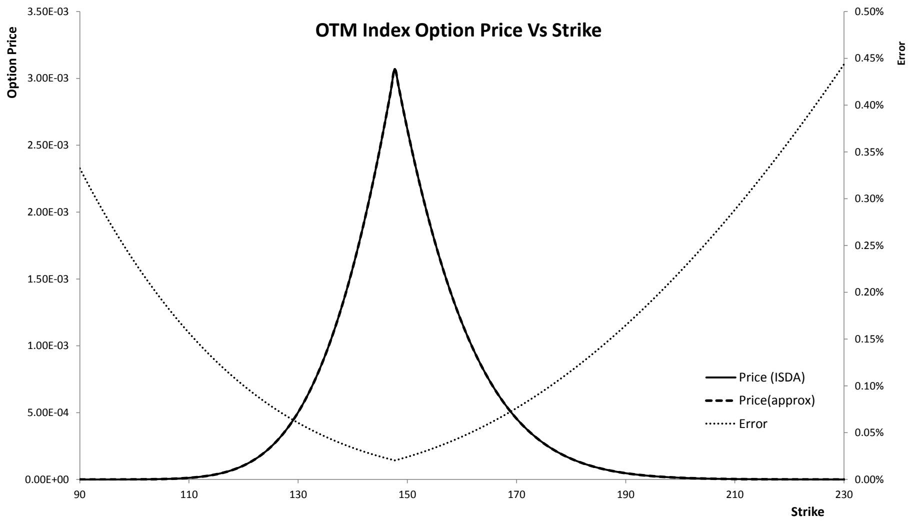

Figure 6: The price of OTM index options against strike using two ways to compute the annuity. The two prices cannot be distinguished visually, and the relative error from the approximate method is shown on the right hand axis. The option is one month from expiry, and the underlying index is 5Y with 125 names (non defaulted).

only 0.02%.

We find the time to price a single option (once the ATM forward price can been computed) to be 5.8ms using the full annuity calculation (ISDA upfront) and 1.4ms using the approximation.[45]

Given that the error is so small, and time to price a single option is about four times faster, there is no compelling reason not to use the approximation. However, as the exercise prices are actual cash payments, these should be calculated exactly - this has no material effect on the time to price an option as it is a one-off calculation.

It is worth pointing out that the timing difference is as small as it is, because we have used our own, cleanroom, implementation of the ISDA model, which is extremely fast. This and all the code used to produce the results in this paper is available in the OpenGamma open source release.

8 Greeks

In this section we discuss the sensitivity of the price of an option (to enter a CDS index), to both market observables (state variables) and model parameters - these are collectively known as the Greeks.

We showed in [Whi13] how to compute analytic Greeks for CDS prices, and this can be extended to compute analytic Greeks for options on CDS. However, here we restrict ourselves to the calculation of the Greeks by a *bump and reprice method* - while this can be seen as just a forward finite difference approximation to the (mathematical) derivative $\frac{\partial V}{\partial X}$, where V is the

[45]These numbers apply to specific hardware so only the relative times are meaningful.

option price and X is the state variable or model parameter, it is these finite difference Greeks rather than the exact mathematical derivatives (analytic Greeks) that are usually given.

When we discuss individual Greeks below, we will present values for options to enter the CDS.NA.HY.21 5Y index. The trade date is 13-Feb-2014, the expiry is 19-Mar-2014 and the maturity (of the index) is 20-Dec-2018. The market date is given in table 10. We also use par spread data (from Markit) on the individual names in the index to build their credit curve. The examples we give use the Pedersen model to price the index options with inputs calculated using the ISDA model, however the technique is quite general and can be applied to any pricing model.

8.1 General Bump and Reprice

The Greek corresponding to the variable X[46] is usually calculated as

$$G_X = \frac{V(X + \Delta X) - V(X)}{\Delta X}$$

where ΔX is the bump size - for some Greeks the difference is not divided by ΔX. As we already mentioned, this is a poor estimate for the derivative $\frac{\partial V(X)}{\partial X}$, but for historical reasons it is this number and not the derivative that is given.[47]

Often we have computed the option price V through some full pricing method, but wish, for speed, to calculate the bumped value $V(X + \Delta X)$ through some quicker (but less accurate) method.[48] In this case it is important to also calculate the base price, $V(X)$, with this quick method rather than using the full computed value - this will lead to cancellation of errors that will give a value of the Greek closer to the value compute using the full model for the bumped price.

8.2 CS01 or Spread DV01

This is normally defined as the change in value of option due to a 1bps increase in spread. There are several ways this could be computed.

The simplest (and quickest to compute) method is to assume a homogenous pool with a single constant hazard rate calibrated to the price or spread of the underlying index. The index spread is then bumped up by 1bps and a new (bumped) hazard rate computed. The difference in the option price with these two hazard rates is the CS01 - we call this *flat spread method*. It is very similar to what is done to calculate the CS01 for an index or single-name CDS (priced with a constant hazard rate) [Whi13].

Another method which also assumes an homogenous pool, involves building an index credit curve from the market quotes of the index at several terms. One then simply bumps the spread quotes[49] up by 1bps, rebuilds a bumped index credit curve and uses this to compute a bumped option price. The difference from the base price (i.e. the price calculated with the original index curve) is the CS01. We call this the *index curve method*.

If we are pricing considering the individual credit curves (adjusted to match index prices as in section 5.5), then one approach is to bump the (implied) spreads of all the curves by 1 bps, then

[46] This can be a scalar or vector.

[47] If one were estimating the derivative, one would use the central difference $\frac{\partial V(X)}{\partial X} \approx \frac{V(X+\Delta X) - V(X-\Delta X)}{2\Delta X}$ which has an accuracy of $\mathcal{O}(\Delta X^2)$ rather than the one sided difference which only has an accuracy of $\mathcal{O}(\Delta X^)$ - see any text on numerical analysis for details.

[48] In practice there may be hundreds of 'Greeks' to calculate.

[49] If the quotes are given as PUF, they must first be converted to quoted spreads.

reprice the option and take the difference. The trouble with this approach is that a 1bps shift is applied to all curves regardless of their current level, so high credit quality entries will have a relatively large shift compared to the lower credit quality enters in the index. An alternative is to first bump the quoted spread of the underlying index by 1bps, then readjust the individual curves to match this price. The option price with these readjusted curves is the bumped value. We call this the *intrinsic method*.

There are of course other ways to bump spreads that will give some measure of CS01. However, we restrict ourselves to the three methods presented above. The results are shown in table 3. The greatest variation between the methods is for ATM options, but even there the largest difference is less than $10 (on a notional of $100MM). In section 9 we compare the CS01 numbers from Bloomberg for the same example.

	Flat Index Spread		Index Curve		Intrinsic	
Strike Price	Payer	Receiver	Payer	Receiver	Payer	Receiver
103	138.95	-45379.86	139.05	-45379.76	138.86	-45378.90
104	781.47	-44737.34	781.92	-44736.88	781.09	-44736.68
105	3269.04	-42249.77	3270.49	-42248.31	3267.79	-42249.98
106	9939.58	-35579.22	9942.71	-35576.09	9936.83	-35580.94
107.144443	23642.68	-21876.12	23646.91	-21871.89	23638.73	-21879.04
108	34507.50	-11011.31	34510.81	-11007.99	34504.03	-11013.73
109	42497.16	-3021.64	42498.52	-3020.28	42495.09	-3022.68
110	45100.02	-418.79	45100.28	-418.52	45098.78	-418.99
111	45494.39	-24.42	45494.40	-24.40	45493.33	-24.43

Table 3: The CS01 for options to enter the CDX.NA.HY.21 5Y index calculated using the three methods discussed in the main text.

8.2.1 Bucketed CS01

A full bucketed CS01 would involve bumping each par spread[50] of each individual underlying credit curve in turn and computing an option price. This will produce several hundred numbers (a CS01 vector). One use for this very granular approach is a delta based VaR calculation, where scenarios of spread changes are known (e.g. from historical data). These delta spreads can be multiplied by the CS01 vector to produce a change in option price for that scenario. Further discussion is outside the scope of this paper.

Another form of bucketed CS01 is sensitivity to points along the index credit curve. One way to proceed is to build an index curve from market quotes of the index at different terms. This curve can then be used to imply index par spreads for different times (the buckets), not necessarily corresponding to the index terms. By bumping these (implied) spreads in turn we can build new index curves, then use these directly to price the option (homogenous pool approximation) or first adjust the underlying credit curves to the index prices implied by the new (bumped) index curve. Either way we will obtain a set of several CS01 corresponding to the buckets we specified.

[50]Since we usually first adjust the curve to match index prices, there are the adjusted spreads implied from the adjusted credit curves.

8.3 Delta

An option Delta is usually defined as the sensitivity of the option price to the price of the underlying - i.e. the first derivative of the option price with respect to the price of the underlying. If the underlying is taken as the spot index price, this gives a spot delta, while if is it taken as the default adjusted forward index, this gives a forward delta.

8.3.1 Forward Delta

We define this as the sensitivity to the ATM forward price

$$\Delta_F(t, T_e, T_m) = \frac{\partial V_{\text{option}}(t, T_e, T_m)}{\partial F_I^D(t, T_e, T_m)} \tag{92}$$

As usual, this can approximated by finite difference.

8.3.2 Spot Delta

Delta in the context of index options should be the sensitivity of the option price to the price of the underlying index. If the price of the index is bumped up by a small amount, then we can readjust the individual credit curves to match this new index price, then price the option from these curves. This is remarkably similar to the calculation of CS01 where we bumped the index spread rather than the index price. This leads to a natural definition of the delta as

$$\Delta(t, T_e, T_m) = \frac{\text{CS01}_{\text{option}}(t, T_e, T_m)}{\text{CS01}_I(t, T_m)} \tag{93}$$

This is a hedge ratio, since if I have a portfolio of one option and $-\Delta$ of the index, the portfolio's CS01 will be zero (it is important that both CS01s are measuring the same spread sensitivity). It is this CS01 ratio that Bloomberg use to compute the option delta. The measure is consistent with the analytic delta since,

$$\Delta(t, T_e, T_m) = \frac{\partial V_{\text{option}}(t, T_e, T_m)}{\partial V_I(t, T_m)} = \frac{\partial V_{\text{option}}(t, T_e, T_m)}{\partial S} \Big/ \frac{\partial S}{\partial V_I(t, T_m)} \tag{94}$$

To show this in practice we compute the delta for our example options. The two methods are central finite difference (where the underlying credit curves are readjusted to the bumped index prices) and the CS01 ratio. In the second case the index CS01 is computed by bumping the index quoted spread by one basis point (this is the standard way to compute an index CS01), while the option CS01 is computed by our *flat index spread* method. Table 4 shows the computed delta for payer and receiver options with these two methods. Again the numbers are close, but they are not expected to be exact as they are different things: one is an approximation to the mathematical delta, while the other is the hedge ratio to make a portfolio insensitive to a one basis point rise in the quoted spread of the index.

Finally in figure 7 we show a plot of the delta (computed with central finite difference), which shows a pretty typical shape.

8.4 Gamma

Gamma is the only second order Greek we will consider. Formally (spot) Gamma is defined as

$$\Gamma(t, T_e, T_m) = \frac{\partial^2 V_{\text{option}}(t, T_e, T_m)}{\partial V_I^2(t, T_m)} = \frac{\partial \Delta(t, T_e, T_m)}{\partial V_I(t, T_m)} \tag{95}$$

	Central Finite Difference		CS01 Ratio	
Strike Price	Payer	Receiver	Payer	Receiver
103	0.29%	-99.70%	0.31%	-99.69%
104	1.64%	-98.34%	1.72%	-98.27%
105	6.95%	-93.04%	7.18%	-92.81%
106	21.33%	-78.66%	21.83%	-78.16%
107.14444	51.25%	-48.74%	51.94%	-48.06%
108	75.26%	-24.72%	75.80%	-24.19%
109	93.13%	-6.86%	93.35%	-6.64%
110	99.03%	-0.96%	99.07%	-0.92%
111	99.93%	-0.06%	99.94%	-0.05%

Table 4: The spot delta for options on CDX.NA.HY using the two methods discussed in the main text.

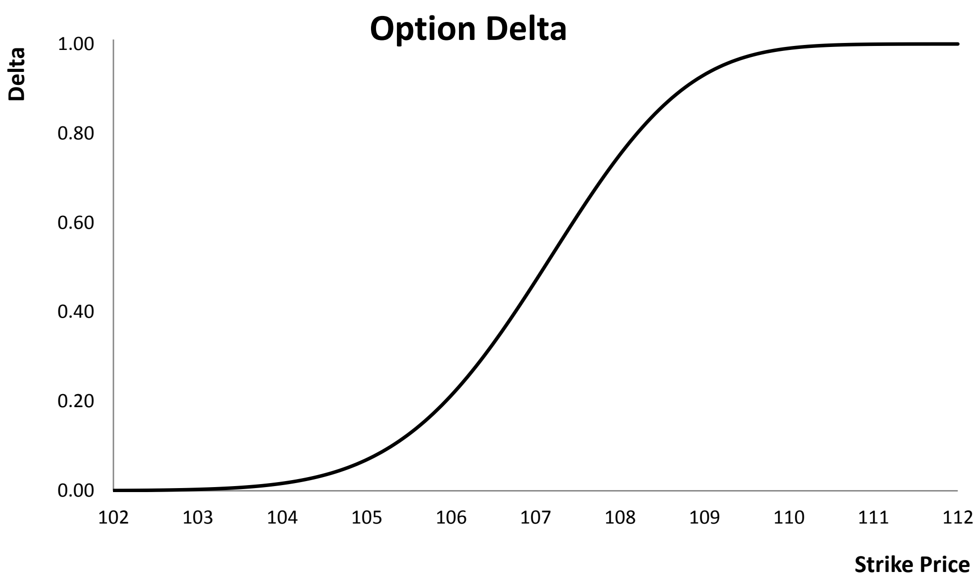

Figure 7: The delta for a payer option on the CDX.NA.HY.21 5Y index. The index is priced based, so the strike is given as an index price.

The standard finite difference way to approximate this is

$$\Gamma(t, T_e, T_m) \approx \frac{V_{\text{option}}(V_I + \Delta V_I) + V_{\text{option}}(V_I - \Delta V_I) - 2V_{\text{option}}(V_I)}{\Delta V_I^2} \tag{96}$$

where $V_{\text{option}}(X)$ means the option price computed with an index value of X using one of the methods we have already discussed. This can also be expressed (less accurately) as a forward finite difference using the Delta, i..e.

$$\Gamma(t, T_e, T_m) \approx \frac{\Delta(V_I + \Delta V_I) - \Delta(V_I)}{\Delta V_I} \tag{97}$$

8.4.1 Alternative Definition of Gamma

Bloomberg defines Gamma as the change in Delta for a 10bp rise in spread. Since Gamma may be computed as

$$\Gamma(t, T_e, T_m) = \frac{\partial \Delta}{\partial S}\frac{\partial S}{\partial V_I} \approx (\Delta(S + 10bp) - \Delta(S)) \times 1000\frac{\partial S}{\partial V_I} \tag{98}$$

it can be viewed as a scaled version of Gamma. In terms of CS01 ratios this can be written as

$$\Gamma_{\text{Bloomberg}} = \frac{\text{CS01}_{\text{option}}(S + 10bp)}{\text{CS01}_{\text{index}}(S + 10bp)} - \frac{\text{CS01}_{\text{option}}(S)}{\text{CS01}_{\text{index}}(S)}. \tag{99}$$

In table 5 we show the Gamma calculated by finite difference and as the change in delta for a 10bp rise in the spread.

	Finite Difference		CS01 Ratio	
Strike Price	Payer	Receiver	Payer	Receiver
103	0.6589	0.6588	0.4941%	0.4941%
104	3.0471	3.0471	2.0107%	2.0107%
105	9.9261	9.9260	5.7525%	5.7526%
106	21.6243	21.6242	10.9853%	10.9854%
107.14444	29.6437	29.6437	12.9199%	12.9200%
108	23.4869	23.4868	9.1124%	9.1124%
109	9.8287	9.8286	3.3217%	3.3218%
110	1.9173	1.9173	0.5631%	0.5632%
111	0.1491	0.1490	0.0379%	0.0380%

Table 5: Two different definitions of Gamma: The mathematical definition (computed by finite difference) and the definition as the change in delta for a 10bps rise in the spread. The second definition is used by Bloomberg.

8.5 Vega

For Pedersen model, this has a straight forward interpretation. It is simply

$$\nu(t, T_e, T_m) = \frac{\partial V_{\text{option}}(t, T_e, T_m)}{\partial \sigma} \tag{100}$$

This is often calculated as the change in option price for a 1% point rise in spread volatility, [51] i.e.

$$\nu(t, T_e, T_m) \approx V_{\text{option}}(\sigma + 1\%) - V_{\text{option}}(\sigma). \tag{101}$$

The Greeks we have calculated so far have involved bumping credit curves (either the intrinsic curves or an index curve), which affects the value of the ATM forward price, which in turn affects the option price. For Vega, the ATM forward price remains unchanged. This means that once an ATM forward price is established, the details of that calculation are irrelevant in the calculation of Vega.

In figure 8 we show a plot of Vega calculated using central finite difference (with a bump of 1bp which has been scaled to a 1% Vega). The value of Vega should be identical for payer and receiver options, and we find this is so (up to some numerical noise).

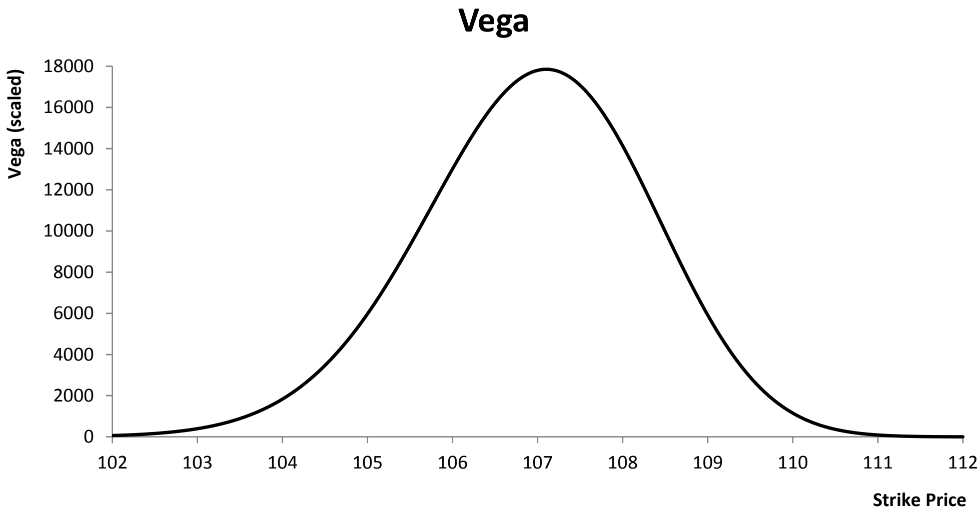

Figure 8: The Vega for a option on the CDX.NA.HY.21 5Y index. The index is priced based, so the strike is given as an index price.

8.6 Theta

Theta is the sensitivity of the option price to the time-to-expiry, so measures the *time decay* in the value of an option. It is often calculated as a one-day-theta - the difference in the option price one day ahead from the price now. Depending on what is held constant with respect to the time-to-expiry, different Theta can be calculated.

Since the ATM Forward value, $F_I^D(t, T_e, T_m)$ is a Martingale (under the $\mathbb{T}$ forward measure), its expected value one day ahead should be just today's value - so, just like for Vega, all we need to consider is the sensitivity of the option price to a change of time-to-expiry for a fixed ATM forward price. For a payer (receiver) option the potential for upside (downside) gains from taking large (small) values of the default-adjusted index value is reduced as the option duration

[51] Of course this is just the forward finite difference approximation scaled down by 10,000.

reduces, so Theta is negative. This calculation implicitly assumes defaults can occur over the one day period (that is what makes $F_I^D(t, T_e, T_m)$ a Martingale).

An alternative (and unconventional) measure of Theta is the value conditional the credit curves remaining the same and no defaults occuring (over the one day period). Viewed this way, Theta consists of two main parts; the spread diffusion part and the default part. For payer options the potential for upside gains from taking large values of the index spread is reduced as the option duration reduces; and the potential for gains from defaults from the index (and their subsequent settlement) is reduced as the time over which defaults can occur is reduced - so Theta is always negative. For receiver options the potential for downside gains from taking small values of the index spread is reduced as the option duration reduces; but the potential for losses from defaults from the index is also reduced as the time over which defaults can occur is reduced - this can give a positive Theta for deeply in the money receiver options.

The ATM forward value is calculated one day ahead assuming the realised value of the curves (credit and yield) one day forward are just their expected value (and of course the index factor remains the same). Once the new ATM forward is calculated Theta is computed as above. These two definitions of Theta are shown for our example option in table 6.

For payer options the Theta is larger in magnitude (more negative) for the no defaults version of Theta - some of the option value comes from the default settlement, so not considering defaults over the one day period lowers the option value more; this is most striking for deeply in the money options. The opposite is true for receiver options, where deeply in the money options gain value (positive Theta).

	With Default		Without Default	
Strike Price	Payer	Receiver	Payer	Receiver
103	-166.32	-148.17	-171.78	1998.55
104	-787.08	-773.30	-819.59	1346.36
105	-2610.19	-2600.78	-2752.95	-591.37
106	-5754.79	-5749.76	-6205.34	-4048.13
107.1444434	-7928.54	-7928.51	-9030.51	-6878.31
108	-6259.34	-6263.06	-7885.03	-5736.57
109	-2580.19	-2588.28	-4590.44	-2446.36
110	-480.55	-493.02	-2613.78	-474.07
111	-20.09	-36.93	-2171.23	-35.89

Table 6: Theta calculated allowing and not allowing defaults over the one day period.

8.7 Interest Rate DV01

This is the sensitivity of the option price to the interest rate. It is normally calculated by calibrating a bumped yield curve with instruments whose rates have been bumped up by 1bp.[52] The IR DV01 of an index is defined as the price calculated with this bumped yield curve when the index spread is kept constant, minus the price with the original yield curve (and of course the same spread).

[52]In practice you will obtain a similar result by bumping the knots of the calibrated yield curve by 1bp.

We may calculate IR DV01 of an option by first adjusting the underlying credit curves to match the index spreads using the bumped yield curve. If we price the option with these curves, this gives us a yield curve bumped price. We will obtain very similar results using the homogenous pool approximation. Table 7 shows the IR DV01 for our example options. The two calculation methods show remarkable agreement (less than a $1 difference for ATM options on a notional of $100MM).

	Flat Index Spread		Intrinsic	
Strike Price	Payer	Receiver	Payer	Receiver
103	1.53	-1800.47	1.53	-1799.45
104	12.02	-1779.14	12.00	-1778.14
105	65.56	-1714.78	65.46	-1713.86
106	247.55	-1521.96	247.26	-1521.23
107.1444	716.91	-1040.21	716.25	-1039.84
108	1166.39	-581.46	1165.50	-581.33
109	1553.41	-183.61	1552.39	-183.60
110	1696.57	-29.63	1695.54	-29.63
111	1713.34	-2.03	1712.31	-2.03

Table 7: The IR DV01 for options on CDX.NA.HY calculated from a flat index curve and using the intrinsic credit curve.

It is well noted that the IR01 for CDS is one or two orders of magnitude smaller than the equivalent CS01 [O'k08, Whi13], and we see the same thing for these risk factors on options on CDS indices.

8.8 Default Sensitivity

This is the change in the option premium due to the default of a single-name from the underlying index. Unlike other Greeks, which are (mathematical) derivatives, this is a jump caused by a sudden unpredictable event. To calculate this value for a given name, we remove that name from the index and recalculate the ATM forward price assuming the remaining intrinsic credit curves are unaffected by the default.[53] This 'bumped' ATM forward price is then used to recalculate the option premium and the difference is the default sensitivity (to that name).

We consider the same four names we looked at in section 5.5.2 (table 2), since they represent very difference credit qualities in the index. Figure 9 shows the result of the calculation for a payer option. For deep in the money options (which are certain to be exercised), the default of *AES Corp* causes a jump in premium of 7.67×10^{-3} - this is larger than the default settlement of 6×10^{-3} (the recovery rate of AES is 40%). Since the credit quality of AES is higher than the index average, its removal raises the value of the index (to a payer of coupons), which accounts for the extra premium gain.[54] At the opposite end, the default of *TX Competitive Elec Hldgs* causes a jump of only 2.72×10^{-4}, despite the fact that its default settlement is 9.4×10^{-3} (it has a recovery rate of only 6%) - this is because the default is already figured in the index value (the spreads are massive), so the default settlement is balanced by a drop in the index value.

[53]In reality, a default will lead to the widening of spreads of other names in the sector.

[54]The index has a negative value, so removal of an 'average' name will increase its value (move it closer to zero), but removal of a better than average name will cause a larger change.

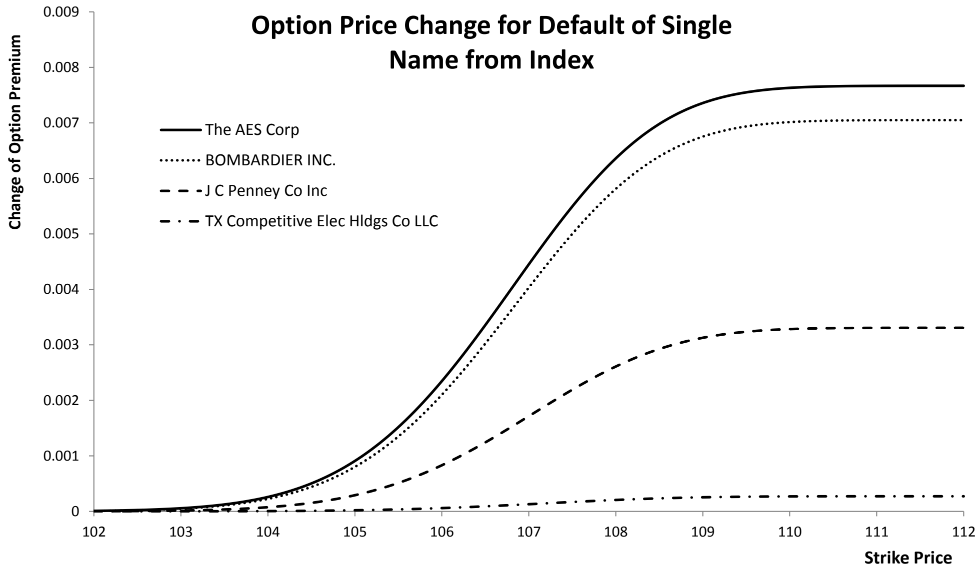

Figure 9: The default sensitivity for a option on the CDX.NA.HY.21 5Y index, for the default of specified names. The index is priced based, so the strike is given as an index price.

One could carry out this calculation for every name in the index. An alternative is to use the homogenous pool approximation and consider the loss of a single-name. After performing this calculation we obtain a curve that plateaus at 7.95×10^{-3}; ignoring discounting, this consists of 7.0×10^{-3} from the default settlement (recall the index recovery rate is 30%) and the remaining from the change to the forward index value.[55]

9 Comparison With Bloomberg

In this section we give a few examples of single-name and index options taken from Bloomberg's CDSO screen,[56] and show how their prices and Greeks compare with our own analytics.

9.1 Single Name Default Swaptions

Our first example is the option to enter the 5Y CDS referencing the republic of Italy. The trade date is 5-Feb-2014, the option expiry date is 20-Mar-2014 and the maturity of the underlying CDS is 20-Jun-2019. If the holder of a payer swaption exercises the option (on 20-Mar-2014) they enter a (bespoke) CDS with a running coupon equal to the strike and accrual starting from 21-Mar-2014 (i.e. T+1). This type of option was discussed in section 4.2.1 and can be priced using the Black formula.

[55]The PUF of the index will remain unchanged, so its value (per unit notional) will increase by 7.62×10^{-4}, i.e. -1% of the PUF.

[56]This allows you to see calculated values of the option premium and Greeks for a user input volatility level. It does not give actual traded option prices.

The yield curve is the ISDA Standard USD snapped at 1600 (New York time) on the 4-Feb-2014 using instruments (deposits and swaps) with spot date of 7-Feb-2014. The credit curve is built using spreads[57] and standard CDS. This market data is given in table 8.

Period	Type	Rate
1M	Deposit	0.16%
2M	Deposit	0.20%
3M	Deposit	0.24%
6M	Deposit	0.33%
1Y	Deposit	0.56%
2Y	Swap	0.44%
3Y	Swap	0.78%
4Y	Swap	1.19%
5Y	Swap	1.58%
6Y	Swap	1.92%
7Y	Swap	2.19%
8Y	Swap	2.42%
9Y	Swap	2.61%
10Y	Swap	2.76%
12Y	Swap	3.01%
15Y	Swap	3.25%
20Y	Swap	3.46%
25Y	Swap	3.55%
30Y	Swap	3.59%
Interest Rate Conventions		
Spot Date		07-Feb-2014
Swap DCC		30/360
MM DCC		A/360
Swap Interval		6M
Floating DCC		A/360
Holidays		none
Floating Interval		3M
Bad Day Conv		MF

Tenor	Spread (bps)
6M	57.4300
1Y	74.9700
2Y	111.3200
3Y	139.3200
4Y	157.6400
5Y	173.6600
7Y	209.2800
10Y	228.3500
CDS Conventions	
Frequency	Quarterly IMM
Day Count	ACT/360
Recovery Rate	40%

Table 8: Market data used to build the Standard ISDA USD yield curve (left), and Spreads on Republic of Italy CDSs for 5-Feb-2014 (right).

Having calibrated the yield and credit curve we compute the forward spread using equation 23, where the (forward) protection leg and annuity are computed using the ISDA model. We obtain a value of 182.764 bps compared to Bloomberg's value (ATM Fwd) of 182.767 bps. The annuity we calculate as 4.8354[58], which compares to the Bloomberg implied value[59] of 4.8364.

[57]These are treated as par spreads, even though the numbers given are actually the quoted spreads for those maturities.

[58]For unit notional and per unit of spread.

[59]This value is not given on the CDSO screen, we have implied it by looking at call-put values across a range of strikes.

We consider a range of strikes from 100 to 300 bps with a spread volatility of 40%. The time to expiry is taken as 0.11781 years.[60] In table 9 we show the calculated option premium and the implied volatility from the premiums given on the CDSO screen. Black1 uses our calculated values for the forward spread and annuity, while Black2 uses the Bloomberg values for these qualities. The implied volatility is the number that must be put into the Black formula to recover the premium given by CSDO (again Imp Vol1 and Imp Vol2 use our and Bloomberg's forward spread and annuity). In both cases the implied volatility is around 39% rather than the 40% we would expect for options priced with the Black formula. The discrepancy is much large than any error from calculating the forward spread and annuity, or from the calculation of the time-to-expiry. The only explanation we can offer is than Bloomberg is using some other model to price these type of options.

	Option Premium			Implied Volatility	
Strike	Bloomberg	Black1	Black2	Imp Vol1	Imp Vol2
100.000	4.0029595%	4.0020013%	4.0029599%	62.116%	39.201%
140.000	2.0777336%	2.0783381%	2.0788947%	39.616%	39.250%
150.000	1.6186780%	1.6211885%	1.6216446%	39.342%	39.220%
160.000	1.1956130%	1.2009104%	1.2012686%	39.243%	39.192%
170.000	0.8304912%	0.8386714%	0.8389382%	39.190%	39.164%
180.000	0.5396918%	0.5497805%	0.5499678%	39.153%	39.137%
182.767	0.4732540%	0.4835945%	0.4837624%	39.144%	39.130%
190.000	0.3274686%	0.3378529%	0.3379762%	39.121%	39.111%
200.000	0.1856680%	0.1948579%	0.1949340%	39.093%	39.086%
210.000	0.0986346%	0.1057882%	0.1058323%	39.066%	39.061%
220.000	0.0492900%	0.0542785%	0.0543026%	39.041%	39.037%
230.000	0.0232763%	0.0264402%	0.0264526%	39.017%	39.013%
250.000	0.0044651%	0.0054703%	0.0054732%	38.971%	38.968%
300.000	0.0000367%	0.0000583%	0.0000584%	38.865%	38.864%

Table 9: Comparison of option premiums given on CDSO to that calculated from the Black formula. The spread volatility is 40%, and implied volatility is for the premium quoted on CDSO. The sets of data 1 & 2 refer to what value is used for the forward and numeraire in the Black formula. See main text for details.

9.2 Index Options - CDX High Yield

Our first index option example is an option to enter the 5Y CDX.NA.HY.21-V1 - this is the same as was used in the Greeks section. The reason we have chosen to start with this index is because it trades on price rather than spread, so there are no issues with converting from a strike given as a spread to an exercise price. The notional amount is $100MM, the trade date is 13-Feb-2014, the option expiry is 19-Feb-2014,[61] and the maturity of the underlying index is 20-Dec-2018. The index has a coupon of 500bps and a recovery rate of 30%.

The current price is 107.62 (so PUF of -7.62) which corresponds to a quoted spread of 322.5621

[60] Using ACT/ACT ISDA day count between 5-Feb-2104 and 20-Mar-2014.

[61] This is chosen so there is no accrued premium.

bps.[62] The market data used to build the yield curve (ISDA Standard USD) and the prices for other terms of CDS.NA.HY.21 are given in table 10. It is worth noting that the CDSO screen gives prices at the standard tenors of 6M, 1Y, 2Y, 3Y, 4Y, 5Y, 7Y and 10Y even though only 3Y, 5Y, 7Y and 10Y trade on-the-run.[63] The prices for tenors less than 3Y are set to the 3Y value, but only the 3Y, 5Y, 7Y and 10Y appear to be used by Bloomberg for curve construction. We also only use these points, and that is what we show in the table. Of course only the 3Y and 5Y points can affect the value of options on the 5Y index.

Period	Type	Rate
1M	Deposit	0.15%
2M	Deposit	0.20%
3M	Deposit	0.24%
6M	Deposit	0.33%
1Y	Deposit	0.55%
2Y	Swap	0.47%
3Y	Swap	0.83%
4Y	Swap	1.27%
5Y	Swap	1.67%
6Y	Swap	2.02%
7Y	Swap	2.30%
8Y	Swap	2.53%
9Y	Swap	2.72%
10Y	Swap	2.89%
12Y	Swap	3.14%
15Y	Swap	3.38%
20Y	Swap	3.58%
25Y	Swap	3.67%
30Y	Swap	3.71%
Interest Rate Conventions		
Spot Date		17-Feb-2014
Swap DCC		30/360
MM DCC		A/360
Swap Interval		6M
Floating DCC		A/360
Holidays		none
Floating Interval		3M
Bad Day Conv		MF

Tenor	Price
3Y	107.5600
5Y	107.6200
7Y	105.7100
10Y	106.5200
CDS Conventions	
Frequency	Quarterly IMM
Day Count	ACT/360
Recovery Rate	30%

Table 10: Market data used to build the Standard ISDA USD yield curve (left), and prices of for CDX.NA.HY.21-V1 on 13-Feb-2014 (right).

[62] Recall the conversion is via the ISDA model.

[63] Seasoned indices (off-the-run) exist that may have, say, 6 months to maturity, but these are highly illiquid and generally will have different constituents to the on-the-run indices.

9.2.1 ATM Forward

What Bloomberg calls the *ATM Forward* for price based options, when expressed as an exercise price, gives the same price for payer and receiver options - so this is in line with what we defined as the ATM Forward in equation 68, i.e. it is the default adjusted forward index value. On the CDSO screen it is given as 107.1444 which corresponds to a forward default adjusted value of -$7,144,443.[64]

We calculate this value a few different ways: The first is to use the prices given in table 10 to build a pseudo credit credit for the index.[65] We may then use equation 69 to compute the ATM forward price. The second method is to compute the intrinsic value from the credit curves of the constituent entries of the index using equation 68. The third method is to first adjust the constituent credit curves to match the index prices (as described in section 5.5.2), then compute the intrinsic value. We consider the last of these to be the *correct* way to compute the ATM forward. Table 11 shows Bloomberg's value together with the three values we calculate. The unadjusted intrinsic value shows the largest difference (0.316%), while the homogenous pool and the adjusted intrinsic value show very close agreement with each other (a difference of less than $600 on a $100MM notional), and close agreement with Bloomberg (a difference of less than $6000 or 0.08%).

	ATM Forward	Error
Bloomberg	-$7,144,443.00	
Curve Built from Index Prices	-$7,149,840.40	-0.076%
Intrinsic Value	-$7,121,885.71	0.316%
Adjusted Intrinsic Value	-$7,150,407.60	-0.083%

Table 11: Value of the ATM Forward (derived) from CDSO together with the calculated values from three methods described in the main text. The error shows percentage difference from the Bloomberg value.

There is nothing in either the documentation [FNS11] or on the CDSO screens to indicate that Bloomberg considers the underlying credit curves in the computation of index options. Most likely they assume a homogenous pool and build a single credit curve from the prices of different terms of the index. That our version of this calculation does not exactly match is most likely down to the fact they use their own model to build the index curve in this case.

9.3 Option Prices

We price a set of these index options using our implementation of the Pedersen model presented in section 6.7.1. Firstly we use the ATM forward value $F_I^D(t, T_e, T_m)$ shown on the CDSO screen - we call this calculation 1. Secondly we use our own calculated value of $F_I^D(t, T_e, T_m)$ from building an index curve - we call this calculation 2. We could have used the adjusted intrinsic value, but as we showed above, the difference in this case is tiny. Table 12 shows the results of these calculations.

Even when we use the same forward as Bloomberg (so there is no issue with curve construction or whether the constituent curves are considered) our option prices differ from Bloomberg's. The

[64]We have backed out the extra accuracy from the put-call relationship.

[65]We use our implementation of the ISDA model that allows this. See [Whi13] for details.

	Bloomberg		Calculation 1		Calculation 2	
strike	Payer	Receiver	Payer	Receiver	Payer	Receiver
103.000	1050.57	4144889.03	1166.48	4144837.71	1150.74	4150217.96
104.000	7281.74	3151266.17	7845.01	3151702.55	7755.71	3157009.24
105.000	38630.22	2182760.61	40556.79	2184600.64	40180.03	2189619.88
106.000	154451.94	1298728.30	158851.69	1303081.86	157696.67	1307322.84
107.144	529900.91	529900.91	536240.00	536240.00	533468.84	538864.84
108.000	1053705.97	198274.25	1058884.29	203487.10	1054818.41	204817.21
109.000	1894442.50	39156.75	1896630.82	41419.94	1891603.29	41788.41
110.000	2859098.57	3958.78	2859437.34	4412.77	2854092.97	4464.40
111.000	3855162.16	168.34	3855042.43	204.19	3849649.48	207.23

Table 12: Prices of options on CDX.NA.HY index. See main text for details.

variation due to which method is used to convert between the (pseudo) spread and the forward price (i.e. the ISDA model or an approximation involving the credit triangle) is too small to explain this difference - when we change from the approximation discussed in section 7 to using the full ISDA model, the price of a receiver with an exercise price of 111 changes from 204.19 to 204.64, which is one-eightieth of the discrepancy we have with Bloomberg's number.

Furthermore, the option prices we compute with our own value of the ATM forward (calculation 2) differ far less from calculation 1 than from the Bloomberg numbers - it is a theme of this section, that while we are confident in our implementation, as we have set out in this paper, we cannot hope to exactly reproduce a black box such as Bloomberg.

An alternative way of looking at the discrepancy is shown in figure 10. There we show the implied volatility corresponding to the option prices given on CDSO - that is, the volatility we must put in our model to obtain Bloomberg prices (recall the volatility is set to 30%). Again a 40bps difference in implied volatility is large, given that the underlying models should be the same up to some implementation details.

For completeness we should mention the role of the discount factor, $P(t, t_{es})$. We compute this value as 0.999814 (from our bootstrapped yield curve), however what we back-out from Bloomberg's option prices is 0.99985 - in terms of zero rates[66] this is 0.174% versus 0.137%. Since the 1M and 2M deposit rates are 0.15% and 0.2% respectively (see table 10), the second number seems a little low. What this means is that the put-call relationship (eqn. 77) holds for all the sets of prices provided the 'correct' discount factor is used. The difference in discount factors does not account for the difference seen in the option prices.

9.3.1 Comparison of Greeks

In our Greeks section we used options on the 5Y CDX.NA.HY.21-V1 as examples. Here we compare our calculations with the values obtained from CDSO. Table 13 shows the full set of Greeks available on CDSO together with our calculations (most of which we show in section 8). For CS01 we use the flat spread method; delta and Gamma use CS01 ratio; Vega is for a 1% bump in (log-normal) volatility; Theta is the (normal) with defaults Theta over one day; and IR DV01 was computed with a flat index spread.

[66] These are given as continuously compounded.

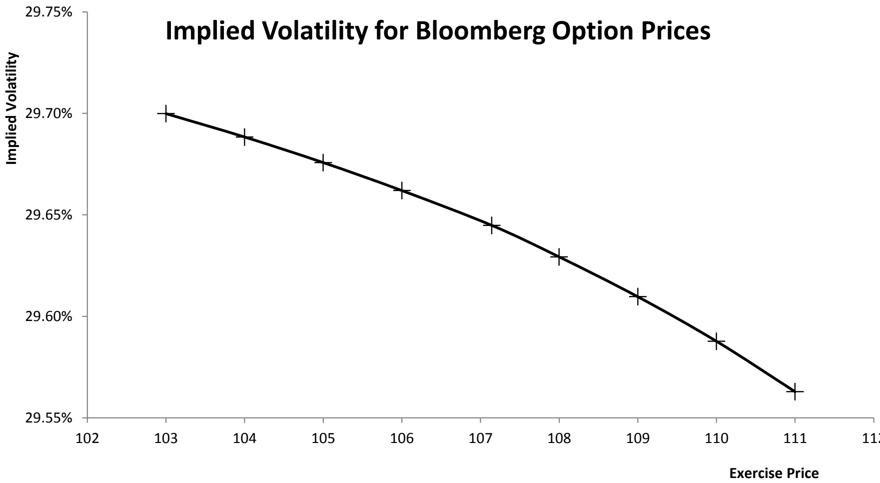

Figure 10: The implied volatility for the prices given on CDSO for options on CDX.NA.HY. The volatility is fixed at 30%, and this implied volatility is the number that must go into our implementation to recover the given price.

We have already seen that our options prices do not agree, so we would not expect an exact agreement on the Greeks. With that caveat, the numbers are close. The notable exception is the Theta of payer options: The values from CDSO for deeply in and out the money options (where there is little optionality) is around -$8000 rather than zero - in fact the entire Theta curve looks like it is pulled down by $8000. This is not just for this index; we find the same thing for options on ITraxx Xover (with the same magnitude) and options on CDX.NA.IG. The options we have considered have more than a month to expiry, yet on the CDSO screen it is easy to have a Theta (the one day change in the option price) that is greater than the current option price - this makes no sense to us.

9.4 Index Options - iTraxx Xover

Our second index option example is on 5Y iTraxx Xover.[67] Like most indices, this is spread based, so the strikes are given as spreads which must be converted to actual exercise prices. The notional is €100MM, the trade date is the 6-Feb-2014, the option expiry is 19-Feb-2014, and the maturity of the underlying index is 20-Dec-2018. The index has a coupon of 500 bps and a recovery rate of 40%. The current spread is 318.25 bps which corresponds to PUF of -7.768. The market data used to build the yield curve (ISDA Standard EUR) and the spreads for other terms of iTraxx Xover are given in table 14.

9.4.1 Exercise Price

There are zero accrued days on 19-Mar-2014, so the full exercise amount is just the clean price of the index at T_e (with index factor 1, valued at t_{es}) for a quoted spread of K. This can be calcu-

[67]iTraxx Europe Crossover Series 20 version 1 5Y.

		Bloomberg						OpenGamma					
	Strike Price	CS01	Delta	Gamma	Vega	Theta	IRDV01	CS01	Delta	Gamma	Vega	Theta	IRDV01
Payer	103	132.00	0.28%	0.2800%	419.00	-8413.00	2.00	138.95	0.31%	0.4941%	450.84	-166.32	1.53
	104	763.54	1.64%	1.9262%	1885.95	-8994.79	14.99	781.47	1.72%	2.0107%	1980.10	-787.08	12.02
	105	3264.81	7.03%	5.6222%	6026.67	-10738.22	77.63	3269.04	7.18%	5.7525%	6200.47	-2610.19	65.56
	106	10086.19	21.72%	10.8710%	12962.36	-13776.41	283.33	9939.58	21.83%	10.9853%	13160.42	-5754.79	247.55
	ATM	24225.91	52.17%	12.8024%	17633.09	-15816.23	797.40	23642.68	51.94%	12.9199%	17841.31	-7928.54	716.91
	108	35412.11	76.26%	8.9212%	13946.80	-14069.72	1275.26	34507.50	75.80%	9.1124%	14231.49	-6259.34	1166.39
	109	43518.12	93.72%	3.1267%	5843.59	-10424.58	1673.27	42497.16	93.35%	3.3217%	6120.34	-2580.19	1553.41
	110	46073.15	99.22%	0.4589%	1143.52	-8454.46	1814.60	45100.02	99.07%	0.5631%	1258.43	-480.55	1696.57
	111	46438.91	100.01%	0.0207%	89.13	-8086.11	1831.47	45494.39	99.94%	0.0379%	106.44	-20.09	1713.34
Receiver	103	-46327.39	-99.77%	0.5136%	418.59	-395.36	-1895.46	-45379.86	-99.69%	0.4941%	396.57	-148.17	-1800.47
	104	-45696.03	-98.41%	1.9784%	1885.95	-932.64	-1874.50	-44737.34	-98.27%	2.0107%	1834.10	-773.30	-1779.14
	105	-43194.76	-93.02%	5.6744%	6026.67	-2631.68	-1803.81	-42249.77	-92.81%	5.7526%	5974.56	-2600.78	-1714.78
	106	-36373.38	-78.33%	10.9233%	12962.36	-5625.48	-1590.06	-35579.22	-78.16%	10.9854%	13015.79	-5749.76	-1521.96
	ATM	-22233.66	-47.88%	12.8546%	17633.09	-7614.50	-1066.77	-21876.12	-48.06%	12.9200%	17842.76	-7928.51	-1040.21
	108	-11047.46	-23.79%	8.9735%	13946.80	-5830.01	-582.02	-11011.31	-24.19%	9.1124%	14136.90	-6263.06	-581.46
	109	-2941.45	-6.33%	3.1790%	5843.59	-2140.48	-175.96	-3021.64	-6.64%	3.3218%	5915.92	-2588.28	-183.61
	110	-386.42	-0.83%	0.5112%	1143.52	-125.97	-0.27	-418.79	-0.92%	0.5632%	1154.03	-493.02	-29.63
	111	-20.66	-0.04%	0.0316%	89.13	286.78	-1.65	-24.42	-0.05%	0.0380%	89.70	-36.93	-2.03

Table 13: Comparison of Greeks from options on CDX.NA.HY

Period	Type	Rate
1M	Deposit	0.22%
2M	Deposit	0.25%
3M	Deposit	0.29%
6M	Deposit	0.39%
9M	Deposit	0.47%
1Y	Deposit	0.55%
2Y	Swap	0.44%
3Y	Swap	0.58%
4Y	Swap	0.78%
5Y	Swap	1.00%
6Y	Swap	1.21%
7Y	Swap	1.40%
8Y	Swap	1.57%
9Y	Swap	1.73%
10Y	Swap	1.87%
12Y	Swap	2.09%
15Y	Swap	2.31%
20Y	Swap	2.46%
30Y	Swap	2.50%
Interest Rate Conventions		
Spot Date		10-Feb-2014
Swap DCC		30/360
MM DCC		A/360
Swap Interval		1Y
Floating DCC		A/360
Holidays		none
Floating Interval		6M
Bad Day Conv		MF

Tenor	Spread (bps)
6M	204.8700
1Y	204.8700
2Y	204.8700
3Y	204.8700
4Y	261.5600
5Y	318.2500
7Y	377.9800
10Y	401.3900
CDS Conventions	
Frequency	Quarterly IMM
Day Count	ACT/360
Recovery Rate	40%

Table 14: Market data used to build the Standard ISDA EUR yield curve (left), and Spreads on iTraxx Europe Crossover Series 20 for 6-Feb-2014 (right).

lated at T_e using the ISDA model[68] with the yield curve calibrated at that time. The dependance on the yield curve is weak[69], so as usual, we calculate $G(K)$ at t assuming deterministic interest rates.

For indices quoted on spread, the CDSO screen does not give the exercise price, so it must be backed out from the option prices. From equation 72 we have that

$$V_{\text{payer}}(t, T_e, T_m) - V_{\text{receiver}}(t, T_e, T_m) = P(t, t_{es})(F_I^D(t, T_e, T_m) - G(K)) \tag{102}$$

The first term in brackets is the ATM forward value, while the second is the exercise price. Since $G(K) = 0$ if $K = C$ (the index coupon), we may extract the first term from the payer and receiver prices at a strike equal to the index coupon. We find this value to be -€7,195,599, which as expected, is close to the current (clean) value of the index. As the first term does not depend on K, we can then easily extract $G(K)$ for a range of K.

In Figure 11 we show the exercise price computed from the ISDA model along with the values we have backed out from options prices given on the CDSO screen. Given that all parties must agree on the exercise price, that there are significant differences[70] between the values we compute from the ISDA model and those we extract from the option prices, is a concern.

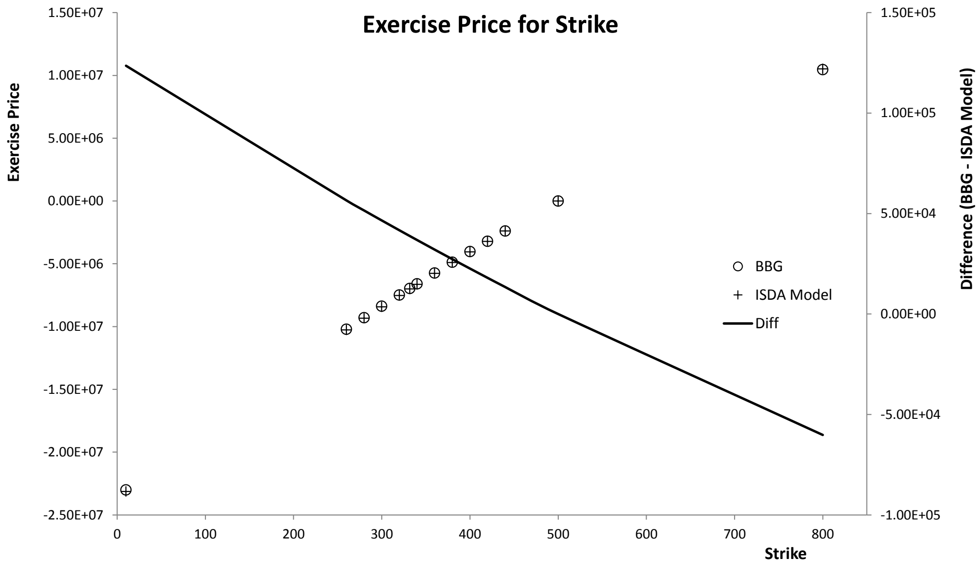

Figure 11: The exercise price for a index option for a range of strikes (given as basis point spreads). The crosses use the ISDA model to convert a spread to an upfront price, while the circles are the values extracted from Bloomberg option prices (see main text). The difference is shown as a line using the righthand scale.

[68] ISDA Standard Upfront Model.

[69] The IR DV01 is much smaller than the Credit DV01.

[70] By construction the two exercise prices must agree for a strike of 500 bps (and have zero value).

9.4.2 ATM Forward Value

As in the previous example, the index is considered homogeneous and the index curve built from the quoted spreads at 3Y, 5Y, 7Y and 10Y. To do this properly, one should first convert the spreads to PUF and then bootstrap a credit curve that recovers all the upfront amounts.

Once this curve is built we may compute the default adjusted index value from equation 69. We obtain a value of -€7,201,983, while, as we saw above, the value derived from CDSO is -€7,195,599, so there is a discrepancy of €6,385 (about 0.09%), which again we cannot reconcile without further details of how the value is obtained, which depends, among other things, on exactly how the credit curve is build.

9.4.3 ATM Forward Spread

The ATM forward given on CDSO is 331.79bps. However this is not the value of the strike that makes the price of the payer and receiver equal - that is a strike of 326.61bps.[71] The value we calculate using equation 78 (which we call the ATM forward spread, K_{ATM}) is 327.39bps - a discrepancy of about 0.8bps from the value of 326.61bps we extracted from CSDO. A calculation of the default-adjusted forward spread, S_I^D, using equation 81 gives a value of 331.65bps - this is within 0.15bps of the Bloomberg value for the ATM forward.

It seems reasonable to assume that Bloomberg uses (a variant of) equation 81 to calculate what it calls ATM Forward [spread]. However this number is not the ATM forward spread, rather it is the defaulted-adjusted forward spread.

9.4.4 Option prices

We price these index options using our implementation of the Pedersen model presented in section 6.7.1. Firstly we use the value of $F_I^D(t, T_e, T_m)$ and the exercise prices $G(K)$ derived from the option prices on the CDSO screen - we call this calculation 1. Secondly we use our own calculated values of $F_I^D(t, T_e, T_m)$ and $G(K)$ - we call this calculation 2. Table 15 shows the results of these calculations. The first set (calculation 1), unsurprisingly shows close agreement with Bloomberg, while the second set (calculation 2) show a larger discrepancy. In this second case we have used a different value of the ATM forward value (which is used to calibrate the model parameter $\bar{X}_0$) and different values for the exercise prices, so it would be surprising to obtain the same option values.

The pertinent question is which values are correct. The calculation of default adjusted index value depends on whether one builds a single index curve or adjusts all the consistent curves, and in either case exactly how this is done. The value is therefore somewhat subjective. The exercise price on the other hand is an amount that is actually paid if the option is exercised, and thus there must be an agreed procedure to convert between a strike given as a spread and this amount. Since for spot indices this procedure is the ISDA standard model, it would seem reasonable to use the same procedure to calculate the excise amount. However Bloomberg does something different which results in different exercise prices.

[71] We obtained this value by manually root finding.

	Bloomberg		Calculation 1		Calculation 2	
strike	Payer	Receiver	Payer	Receiver	Payer	Receiver
10	15,806,483	0	15,806,483	0	15,925,985	0
260	3,026,637	5,587	3,026,698	5,648	3,077,145	4,979
280	2,136,488	38,353	2,136,691	38,556	2,179,572	35,416
300	1,347,217	158,280	1,347,589	158,652	1,379,940	150,021
320	736,790	443,543	737,202	443,955	757,534	428,289
326.6135	582,865	582,865	583,251	583,251	599,880	565,518
331.7902	479,143	707,672	479,498	708,028	493,463	688,904
340	343,256	932,402	343,551	932,696	353,804	911,879
360	135,570	1,594,013	135,711	1,594,154	139,806	1,572,051
380	45,626	2,360,473	45,670	2,360,517	46,964	2,340,431
400	13,229	3,171,785	13,236	3,171,792	13,558	3,155,497
420	3,350	3,993,114	3,349	3,993,113	3,411	3,981,267
440	751.64	4,809,417	750.31	4,809,415	758.36	4,802,169
500	4.69	7,193,399	4.65	7,193,399	4.56	7,199,782
800	0.00	17,646,109	0.00	17,646,109	0.00	17,713,585

Table 15: Prices of options on ITraxx Xover. See main text for details.

A Mathematical Preliminaries

A.1 Optional Pricing Theory

We state here without proof a couple of useful results from Martingale pricing theory that we use in the main text. For details see one of the many books that cover this material in a financial context (for example [HK04, Shr04, Reb02, RMW09]).

Let $X(u)$ be the value of some asset at time u, that depends only on the value of some (finance) observables at time u. Let $N(u)$ be the value of some other asset at time u, such that $N(u) > 0\ \forall\, u < T^*$, where T^* is some horizon time. We may measure the value of X in terms of N, i.e.

$$\tilde{X}(u) = \frac{X(u)}{N(u)}.$$

The equivalent Martingale theory states that there exists a probability measure $\mathbb{P}^{\mathbb{N}}$ such that $\tilde{X}(u)$ is a Martingale. That is

$$\mathbb{E}_t^{\mathbb{N}}\left[\tilde{X}(u)\right] = \mathbb{E}_t^{\mathbb{N}}\left[\frac{X(u)}{N(u)}\right] = \tilde{X}(t) = \frac{X(t)}{N(t)} \tag{103}$$

for $u \geq t$. $\mathbb{E}_t^{\mathbb{N}}[\cdot]$ is shorthand for $\mathbb{E}_t^{\mathbb{N}}[\cdot|\mathcal{F}_t]$, which is the expectation under the $\mathbb{P}^{\mathbb{N}}$ measure conditional on the filtration $\mathcal{F}_t$, which loosely means that the expectation only depends on information available at t. $N(\cdot)$ is known as the numeraire associated with the measure $\mathbb{P}^{\mathbb{N}}$.

The value today (time t) of a derivative with payoff at T of $X(T)$ is given by

$$V(t) = \mathbb{E}_t\left[e^{-\int_t^T r_s ds} X(T)\right] = \beta_t(t)\mathbb{E}_t\left[\frac{X(T)}{\beta_t(T)}\right] \tag{104}$$

where r_s is the risk-fee (instantaneous) rate at s and $\beta_t(s) = e^{\int_t^s r_s ds}$ is the value of a money market account started at t (so $\beta_t(t) = 1$), i.e. it is the numeraire associated with the *risk-neutral measure*.[72] We can use the result of equation 103 to write the option price as

$$V(t) = N(t)\mathbb{E}_t^{\mathbb{N}}\left[\frac{X(T)}{N(T)}\right] \tag{105}$$

This is useful in the case that the payoff can be written in the form $X(T) = N(T)Y(T)$. Then

$$V(t) = N(t)\mathbb{E}_t^{\mathbb{N}}\left[Y(T)\right] \tag{106}$$

where $Y(\cdot)$ is a Martingale under $\mathbb{P}^{\mathbb{N}}$. Then any model form the dynamics of $Y(\cdot)$ that is a Martingale under $\mathbb{P}^{\mathbb{N}}$, will produce arbitrage free prices.

A.2 Curves

For a given t, $P(t,T)$ and $Q(t,T)$ are functions (or curves) of the independent variable T, and are know as the discount (or yield) curve and the credit curve. They form the bases of CDS pricing [Whi13, O'k08].

A.2.1 Discount Curve

We assume that at an observation time, t, the price of a risk-free[73] zero-coupon bond that pays 1 at time $T \geq t$ is $P(t,T)$. This is also known as the (risk-free) discount factor, since we can discount a known cash flow of c at time T to today (time t) by selling c zero-coupon bonds for $cP(t,T)$. Clearly $P(T,T) = 1$.

Define $P(t,S,T)$, $T > S \geq t$, as the forward discount factor observed at t for the period S to T. That is, I agree today (t) to pay $P(t,S,T)$ at time S to receive 1 at time T. By a standard no arbitrage argument, we have

$$P(t,S,T) = \frac{P(t,T)}{P(t,S)} \tag{107}$$

and $P(S,S,T) = P(S,T)$. Since interest rates evolve over time, generally $P(t,S,T) \neq P(S,T)$. However in the $\mathbb{S}$ forward measure which uses $P(\cdot,S)$ as the numeraire, we have

$$\mathbb{E}_t^{\mathbb{S}}\left[P(u,S,T)\right] = P(t,S,T) \quad \forall\, u \geq t \tag{108}$$

since the $P(u,S,T)$ is the ratio of an asset to the numeraire. So $P(u,S,T)$ is a Martingale in the $\mathbb{S}$ forward measure.

A.2.2 Credit Curve

Let τ be the default time of a particular obligor. If it has not defaulted by t (i.e. $\tau > t$), then the probability that it has not defaulted by $T \geq t$ is

$$Q(t,T) = \mathbb{E}_t[\mathbb{I}_{\tau > T}|\tau > t] = \mathbb{P}(\tau > T|\tau > t) \tag{109}$$

in the risk-neutral measure. If default has occurred (i.e. $\tau \leq t$), then $Q(t,T) = 0$.

[72] Often this is written as $\mathbb{E}_t^{\mathbb{P}}[\cdot]$ - we suppress the superscript when dealing with the risk-neutral measure.
[73] There is no chance that the counterpart will fail the pay.

Define $Q(t, S, T)$ as the probability observed at t that the obligor will not have defaulted by T conditional on it not defaulting by S. It is a standard conditional probability result that

$$Q(t, S, T) = \frac{Q(t, T)}{Q(t, S)}. \tag{110}$$

As the credit quality of the obligor will evolve over time, have $Q(t, S, T) \neq Q(S, T)$.

A.3 Survival Measure Pricing

Consider a derivative that pays an amount $Z(T)$ (which is fixed at T) at some time $t_{es} \geq T$ conditional on a reference obligor not defaulting by T - if default occurs after T, but before t_{es}, the payment is still made. The payoff at T is

$$P(T, t_{es})Z(T)\mathbb{I}_{\tau>T}$$

and the present value is given by

$$Z(t, T) = \mathbb{E}_t \left[\frac{P(T, t_{es})}{\beta_t(T)} Z(T)\mathbb{I}_{\tau>T} \right]$$

Assume that we can calculate the value $Z(t, T)$ by some means. Now consider a second derivative that pays an amount X(T) at t_{es} conditional on the **same** reference obligor not defaulting by T. Its present value is give by

$$X(t, T) = \mathbb{E}_t \left[\frac{P(T, t_{es})}{\beta_t(T)} X(T)\mathbb{I}_{\tau>T} \right]$$

We would like to price this by using the first asset, $Z(t, T)$, as a numeraire, but since its value becomes zero in the event of a default, it fails the technical condition that $N(u) > 0 \,\forall\, u < T^*$. However, we are saved by the fact that the derivative value, $X(t, T)$ also becomes zero in the event of a default, so there is no technical restriction. This means we are free to change to the $\mathbb{Z}$ survival measure [Sch04] with $Z(t, T)$ as the numeraire and write

$$X(t, T) = Z(t, T)\mathbb{E}_t^{\mathbb{Z}} \left[\frac{P(T, t_{es})X(T)\mathbb{I}_{\tau>T}}{P(T, t_{es})Z(T)\mathbb{I}_{\tau>T}} \right] = Z(t, T)\mathbb{E}_t^{\mathbb{Z}} \left[\frac{X(T)}{Z(T)} \right] \tag{111}$$

The survival indicator has disappeared from the expectation. This is particularly useful if we can write $X(T) = Z(T)Y(T)$, which gives

$$X(t, T) = Z(t, T)\mathbb{E}_t^{\mathbb{Z}} \left[Y(T) \right]$$

so the task is to compute the expectation of $Y(T)$ under $\mathbb{Z}$.

A.3.1 Option Payoff

Imagine an option payoff at T of the form

$$p(T, t_{es})\mathbb{I}_{\tau>T}(\chi[X(T) - KZ(T)])^+ \; = p(T, t_{es})\mathbb{I}_{\tau>T}(\chi Z(T)[Y(T) - K])^+$$

We may compute its present value thus

$$V_{\text{option}}(t, T) = Z(t, T)\mathbb{E}_t^{\mathbb{Z}} \left[(\chi[Y(T) - K])^+ \right]$$

Since Y is a Martingale in the $\mathbb{Z}$ survival measure, any model that preserves this property will produce arbitrage free prices. In particular the Black model can be used in this case.

A.3.2 The Risky Bond Measure

A risky (zero coupon) bond issued by some obligor will pay 1 at time t_{es} if the obligor has not defaulted by T.[74] Its present value is

$$B(t,T) = \mathbb{E}_t \left[e^{-\int_t^{t_{es}} r_s ds} \mathbb{I}_{\tau > T} \right] \tag{112}$$

If we assume that credit events are independent of interest rates,[75] we may write

$$B(t,T) = \mathbb{E}_t \left[e^{-\int_t^{t_{es}} r_s ds} \right] \mathbb{E}_t \left[\mathbb{I}_{\tau > T} \right] = P(t, t_{es}) Q(t,T) \tag{113}$$

However, it is not necessary to assume the independence of credit events and interest rates. We may change to the $\mathbb{T}$ forward measure (with $P(\cdot, T)$ as the numeraire), which gives

$$B(t,T) = P(t,T) \mathbb{E}_t^{\mathbb{T}} \left[P(T, t_{es}) \mathbb{I}_{\tau > T} \right] \tag{114}$$

If $T = t_{es}$ (i.e. there is no payment delay) and we redefine the survival probability as

$$Q(t,T) = \mathbb{E}_t^{\mathbb{T}} \left[\mathbb{I}_{\tau > T} \right] \tag{115}$$

then $B(t,T) = P(t,T)Q(t,T)$ with no independence assumption - all we have done is define $Q(t,T)$ as the survival probability in the $\mathbb{T}$-forward measure, rather than the risk-neutral measure.[76]

In practice when a payment delay exists, it will be only a few days (i.e. three working days from the exercise of an option to settlement of payments), so we may write

$$\mathbb{E}_t^{\mathbb{T}} \left[P(T, t_{es}) \mathbb{I}_{\tau > T} \right] \approx \mathbb{E}_t^{\mathbb{T}} \left[P(T, t_{es}) \right] \mathbb{E}_t^{\mathbb{T}} \left[\mathbb{I}_{\tau > T} \right] = \frac{P(t, t_{es})}{P(t,T)} \left[\mathbb{I}_{\tau > T} \right]$$

which gives

$$B(t,T) = P(t, t_{es}) Q(t,T)$$

The function $B(t, \cdot)$ is known as the risky discount curve (for a particular obligor and currency).

This risky bond is equivalent to setting $Z(T) = 1$ above. This means a derivative with payoff $X(T)$ (at t_{es}) conditional on the bond not defaulting, has a present value given by

$$V(t) = B(t,T) \mathbb{E}_t^{\mathbb{B}} \left[\frac{P(T, t_{es})}{B(T,T)} X(T) \mathbb{I}_{\tau > T} \right] = B(t,T) \mathbb{E}_t^{\mathbb{B}} \left[X(T) \right] \tag{116}$$

since $B(T,T) = P(T, t_{es}) \mathbb{I}_{\tau > T}$. To price the derivative we would need to compute the expectation of $X(T)$ under the risky bond measure.

A.4 The Annuity Measure

Consider a derivative price that can be written as

$$\begin{aligned} V(t,T) &= \mathbb{E}_t \left[e^{-\int_t^{t_{es}} r_s ds} A(T,S) X(T) \mathbb{I}_{\tau > T} \right] \\ &= \mathbb{E}_t \left[\frac{P(T, t_{es})}{\beta_t(T)} A(T,S) X(T) \mathbb{I}_{\tau > T} \right] \end{aligned} \tag{117}$$

[74] Assume the money is held in escrow between T and t_{es}, so it is still paid if there is a default in this period.
[75] Since an increase in debt cost is a contributing factor in bankruptcy, they are not really independent.
[76] It follows that if rates and credit events are independent, then $\mathbb{E}_t^{\mathbb{T}} \left[\mathbb{I}_{\tau > T} \right] = \mathbb{E}_t \left[\mathbb{I}_{\tau > T} \right]$.

where $A(T, S)$ is the annuity at T (for some CDS maturity $S > T$), with cash settlement at t_{es}. Since the annuity is an asset (the premium leg of a unit coupon), we may set $Z(T) = A(T, S)$ and change to the *annuity measure* $\mathbb{A}$. This gives the derivative price as

$$\begin{aligned} V(t,T) &= A(t,T,S)\mathbb{E}_t^{\mathbb{A}}\left[\frac{P(T,t_{es})A(T,S)\mathbb{I}_{\tau>T}}{P(T,t_{es})A(T,S)\mathbb{I}_{\tau>T}}\;X(T)\right] \\ &= A(t,T,S)\mathbb{E}_t^{\mathbb{A}}\left[X(T)\right] \end{aligned} \tag{118}$$

where

$$A(t,T,S) = \mathbb{E}_t\left[\frac{P(T,t_{es})}{\beta_t(T)}A(T,S)\mathbb{I}_{\tau>T}\right]$$

Any quantity defined as

$$Y(t,T) = \frac{X(t,T)}{A(t,T,S)}$$

will be a Martingale in the annuity measure, i.e. $\mathbb{E}_t\left[Y(u,T)\right] = Y(t,T)\,\forall u \geq t$. This includes the forward spread.

B List of Terms

B.1 General Terms

- t today or the trade date.
- t_{cs} Cash settlement time for a spot CDS/index. For standard CDS and CDS indices this is three working days after the trade date.
- T_e The exercise date of a CDS/index option or the forward start date of a forward starting CDS/index.
- t_{es} The exercise settlement date. This would normally correspond to the cash settlement date of a spot CDS with trade date T_e (i.e. t_{es} is three workings days after T_e).
- T_m The maturity (protection end) of a CDS.
- s_i The start of the i^{th} accrual period.
- e_i The end of the i^{th} accrual period.
- t_i The i^{th} payment date.
- Δ_i The year fraction of the i^{th} accrual period.
- Δ The year fraction in the current accrual period, i.e. the year fraction between s_1 and $t{+}1$ 9the spin-in date). Multiplied by the day count convention base (almost always 360) this gives the accrued days. Multiplied by the coupon and the notional, it gives the (magnitude) of the accrued. Also used for the year fraction at T_e for a forward starting CDS (i.e. one starting at T_e).
- C The coupon of a CDS.

- PUF Points Up-Front. The clean price of a CDS or index quoted as a percentage of notional.
- N The (initial) notional of a CDS or index.
- r_s Risk-free instantaneous interest rate at time s.
- $P(t,s)$ The price at t of a risk free zero coupon bond that pays 1 at $s \geq t$. $P(t,\cdot)$ is the discount curve observed at t.
- $Q(t,s)$ The survival probability - the probability that a reference entity that has not defaulted by t will not have defaulted (i.e. survived) by $s \geq t$. $Q(t,\cdot)$ is the survival or credit curve observed at t.
- $h(t,s)$ The forward instantaneous hazard rate. This is related to the survival probability by $Q(t,s) = \exp(-\int_t^s h(t,s')s'ds')$
- $\dot{Q}(t,s) \equiv \frac{dQ(t,s)}{ds} = -h(t,s)Q(t,s)$. This is the (negative of) the probability density for default at s given no default by $t \leq s$.
- $\Lambda(t,s)$ The zero hazard rate. This is related to the survival probability by $Q(t,s) = \exp(-\Lambda(t,s)(s-t))$
- $\mathbb{E}_t[\cdot]$ Risk neutral expectation given the information available up to time t.
- $\mathbb{E}_t^{\mathbb{X}}[\cdot]$ Expectation under some equivalent measure $\mathbb{X}$ given the information available up to time t.
- $\beta_t(s) = \exp(\int_t^s r_u du)$ is the money market numeraire; It is the amount of money in a 'risk-free' bank account at time $s \geq t$ if one unit of currency is invested at t.
- ΔX A small change in the quantity X. This should not be confused with the use of Δ for a year fraction.

B.2 Single Name CDS Terms

- τ The default time. The random time when a reference entity defaults.
- $\mathbb{I}_{\tau>t}$ Default indicator. Takes the value 1 if the subscript is true, i.e. the default (τ) is after t, and 0 if it is false.
- $R(\tau)$ Realised recovery rate for a default at τ.
- $\hat{R}(T)$ Expected recovery rate for protection to T.
- R Expected recovery rate when there is no term structure of recovery (i.e. no dependence on length of protection).
- η_i Ratio of the year fraction for the i^{th} accrual period measured with the accrual day count convention to the same period measured with the curve day count convention. In almost all cases this is just 365/360.
- $V_{\text{perm}}(t,T)$ Value of the premium leg of a CDS with protection end at T observed at t. By convention the valuation is rolled forward to $t_{cs} \geq t$.

- $V_{\text{prot}}(t, T)$ As above but for the protection leg.
- $A(t, T)$ The (risky) annuity, duration or RPV01 of a CDS with protection end at T. This is related to the premium leg by $V_{\text{perm}}(t, T) = C \times A(t, T)$.
- $S(t, T)$ The par spread of a CDS with maturity of T - the value of the coupon that would make the fair value of the CDS zero. A set of these (corresponding to different maturities of CDS on the same name) can be used to construct a full credit curve; conversely, to computed this requires a full credit curve. By construction the following relationship holds: $V_{\text{prot}}(t, T) = S(t, T)A(t, T)$
- λ Constant hazard rate. A (often hidden) parameter used to convert between PUF and a quoted spread for a CDS or index.
- $\bar{S}(t, T)$ The quoted or flat spread of a CDS maturity T. If a CDS is priced assuming a constant hazard rate, then all quantities we calculate from that 'flat' credit curve, we prefix with *flat* and use an over-bar in the maths. The flat spread is the spread that is calculated from a flat credit curve.
- $\bar{V}_{\text{prot}}(t, T)$ or $\bar{V}_{\text{prot}}(t, T|\bar{S})$ The value of the flat protection leg of a CDS. The value of the protection leg calculated using a flat (i.e. constant hazard rate) credit curve. See above.
- $\bar{A}(t, T)$ or $\bar{A}(t, T|\bar{S})$ The flat annuity. The value of the annuity calculated using a flat (i.e. constant hazard rate) credit curve (see above). In the second case we mean the annuity corresponding to the flat spread $\bar{S}$.
- $V(t, T)$ Value of a CDS with protection end (maturity) at T observed at t. By convention the valuation is rolled forward to $t_{cs} \geq t$.
- $V(t, T_e, T_m)$ Present value of a forward starting CDS. The CDS is entered at T_e and protection ends at $T_m > T_e$.This value includes the possibility of default before T_e (in which case the contract is worth zero - there is no frontend protection)
- $V_{\text{port}}(t, T_e, T_m)$ Present value of the protection leg of a forward starting CDS. See above.
- $A(t, T_e, T_m)$ Present value of the annuity of a forward starting CDS. See above.
- $F(t, T_e, T_m)$ Forward value of CDS. The expected value of a CDS at T_e conditional on the reference not having defaulted by T_e. The CDS maturity is T_m and the expecting is calculated with all the information available at t.
- $S(t, T_e, T_m)$ The forward spread observed at t for protection starting at $T_e > t$ and ending at $T_m > T_e$.
- $\bar{S}_0$ Expected value of the flat spread $\bar{S}(T_e, T_m)$ at t under the risky bond measure.
- $\bar{K}_{ATM}(t, T_e, T_m)$ The ATM forward flat spread defined in section 3.5.1.

B.3 CDS Index Terms

- J The initial number of entries in an index.
- J_D The number of defaults from an index. $J_D(t)$ the number of defaults at time t.
- w_i The weight of the i^{th} name in an index. For an equally weighted index this is just $1/J$.
- f The index factor - the fraction of names remaining in an index. $f(t)$ The index factor at time t.
- R_I Recovery rate of an index. This is purely used to convert between the index spread and an upfront amount.
- $V_I(t,T)$ (Clean) value of an index with maturity T.
- $\bar{V}_{I,\text{prot}}(t,T)$ Nominal value of the protection leg of an index using a constant hazard rate.
- $\bar{A}_I(t,T|\bar{S})$ Flat annuity of an index, calculated from the (flat) spread, $\bar{S}$.
- $\bar{S}_I(t,T)$ Index flat spread. An alternative way of expressing the price of an index.
- $V_{I,\text{prot}}(t,T)$ Value of the protection leg of an index, calculated using adjusted intrinsic credit curves.
- $A_I(t,T)$ Annuity of an index, calculated using adjusted intrinsic credit curves.
- $S_I(t,T)$ Index spread, calculated using adjusted intrinsic credit curves.
- $V_I^D(T_e,T_m)$ The clean value of an index plus the settlement value of any defaults from the index before T_e. This is known as the default-adjusted forward portfolio swap.
- $F_I^D(t,T_e,T_m)$ The ATM forward value or default-adjusted forward index value. Given by $F_I^D(t,T_e,T_m) \equiv \mathbb{E}_t^{\mathbb{T}}\left[V_I^D(T_e,T_m)\right]$
- $S_I(t,T_e,T_m)$ The forward index spread - this does not include default settlement.
- $S_I^D(t,T_e,T_m)$ The default-adjusted forward index spread.
- K_{ATM} The ATM forward spread. It is the strike (spread) that makes the exercise price equal to the The ATM forward value.

Additionally, any single-name term, X when considered as the j^{th} entry of an index is given the notation X_j.

B.4 Option Terms

- $V_{\substack{\text{option}\\\text{knockout}}}(t,T_e,T_m)$ The price of an option (payer or receiver) at time t on a single-name CDS with or without a knockout feature. The expiry is at T_e and the maturity of the underlying CDS is T_m.
- $V_{I,\text{option}}(t,T_e,T_m)$ The price of an option (payer or receiver) at time t on a CDS index. The expiry is at T_e and the maturity of the underlying index is T_m.

- χ Takes the value +1 for a payer option and −1 for a receiver option.
- $\sigma(t)$ Time dependent, deterministic volatility. $\hat{\sigma}$ or σ is root-mean-square (RMS) value.
- K Strike level expressed as a spread.
- $G(K)$ The exercise price of an option. This is often a function of the strike (spread) K.
- Black(F, K, T, σ, χ) The Black formula (see eqn 40).
- $\bar{X}$ A flat pseudo spread that absorbs defaults from the index. Used in the Pedersen model.
- $\bar{X}_0$ The calibrated mean of $\bar{X}$.
- $\tilde{K}$ The modified strike used the modified Black index option pricing formula (eqn. 87).

References

[Bla76] Fischer Black. The pricing of commodity contracts. *Journal of Financial Economics*, 3:167–179, 1976. 10

[BM09] Damino Brigo and Massimo Morini. Last Option before the Armageddon. *Risk*, 2009. 27

[FNS11] Bjorn Flesaker, Madhu Nayakkankuppam, and Igor Shkurko. Credit Index Options in CDSO. Technical report, Bloomberg, 2011. 42

[HK04] P. J. Hunt and J. E. Kennedy. *Financial Derivatives in Theory and Practice.* Wiley series in probability and statistics. Wiley, second edition, 2004. 49

[Mar12] R. J Martin. A CDS Option Miscellany. *ArXiv e-prints*, December 2012. 13

[O'k08] Dominic O'kane. *Modelling single-name and multi-name Credit Derivatives.* Wiley F, 2008. 15, 16, 37, 50

[Ped03] Claus M. Pedersen. Valuation of Portfolio Credit Default Swaptions. Technical report, Lehman Brothers, 2003. 1, 15, 16, 20, 23

[PTVF07] William H. Press, Saul A. Teukolsky, William T. Vetterling, and Brain P. Flannery. *Numerical Recipes.* Cambridge University Press, 2007. 9, 24

[Reb02] Riccardo Rebonato. *Modern pricing of interest-rate derivatives: the LIBOR Market Model and Beyond.* Princeton University Press, Princeton and Oxford, 2002. 49

[RMW09] Riccardo Rebonato, Kenneth Mckay, and Richard White. *The SABR Libor Market Model.* Wiley, 2009. 49

[Roz09] Saha Rozenberg. A 'Big Bang' in the Credit Derivatives Universe. *Derivatives Week*, 2009. 3

[Sch04] Philipp Schonbucher. A Measure of Survival. *Risk*, 2004. 6, 51

[Shr04] Steven E. Shreve. *Stochastic calculus for finance II: Continuous-Time Models.* Springer, 2004. 49

[Whi13] Richard White. The Pricing and Risk Management of Credit Default Swaps, with a Focus on the ISDA Model. Technical report, OpenGamma, 2013. 1, 2, 3, 28, 29, 30, 37, 42, 50

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

OpenGamma