XAIA INVESTMENT

Pricing single-name CDS options: a review of standard approaches

Jan-Frederik Mai
XAIA Investment GmbH
Sonnenstraße 19, 80331 München, Germany
jan-frederik.mai@xaia.com

Date: March 20, 2014

Abstract The derivations of the market standard formulas for the pricing of single-name CDS options are reviewed. Adaptations to the current market practice of trading CDS at standardized running coupons with an initial upfront payment are discussed.

1 Single-name CDS

A credit default swap (CDS) is an insurance contract between two parties. The protection buyer pays a periodic premium[1], called *CDS running coupon*, to the insurance seller, who in return compensates the insurance buyer for potential losses resulting from a *credit event* during the lifetime of the contract. A credit event is related to a third party, the so-called *reference entity*, and is typically triggered when the latter fails to make a coupon payment on an outstanding bond, files for bankruptcy, or restructures its debt in an unfavorable way for certain bond holders. The pricing of a CDS requires the computation of the expected value of the net present value of all payments to be made by protection buyer and seller. The value of the CDS for the protection buyer is then given by the expected net present value of compensation payments received minus the expected net present value of all coupon payments to be made. The analogous logic implies that the CDS value for the protection seller is the negative of the CDS value for the protection buyer. At inception, the value of the contract should be zero, which basically determines the CDS running coupon. However, it has become market standard that the running coupon is standardized, e.g. to 500 bps, so that in order for the CDS to be worth zero at inception an *upfront payment* from one party to the other is required. It is convention that the upfront payment is always quoted in percent of the notional and as if it is paid by the insurance buyer, i.e. it is negative if the protection seller has to make this initial payment. Let us now introduce some notation and have a closer look at the mathematics involved. To simplify notation, we assume unit notional for the CDS contract, i.e. upon occurrence of the credit event the protection seller compensates the protection buyer for potential losses incurred by some deliverable debt obligation with unit notional.

1.1 Conventions and notations

Formally, we work on a filtered probability space $(\Omega, \mathcal{F}, (\mathcal{F}_t), \mathbb{P})$ satisfying the usual hypotheses, where $\mathcal{F}_t$ denotes all information available to market participants at time t and $\mathbb{P}$ denotes a

[1] By market standard convention premium payments are made quarterly at the so-called *IMM dates* March 20, June 20, September 20, and December 20.

risk-neutral pricing measure (under which the discounted versions of all tradable assets are martingales). The future time point at which the credit event happens is modeled as an $(\mathcal{F}_t)$-stopping time and denoted by τ. We furthermore consider time points $t_0 \leq t_E < t_1 < \ldots < t_N$, where $t_1, \ldots, t_N$ denote the premium payment dates of the CDS. Moreover, t_0 denotes the last IMM date before the settlement date t_E of the CDS (required for the computation of the accrued interest payment to be made at inception), t_1 denotes the first IMM date after settlement and t_N denotes the maturity date of the CDS. Discount factors are denoted by $DF(t,T)$, and one may think of them as being given in terms of a deterministic reference short rate r_t via $DF(t,T) = \exp\left(-\int_t^T r_u\, du\right)$. The present value at time t – provided default has not yet occurred – of the sum over all (clean) cash flows to be made by the protection buyer (the *discounted premium leg*) is given by

$$
\begin{aligned}
DPL(t,t_E,t_N) := \sum_{i\,:\,t<t_i} &\underbrace{(t_i - \max\{t,t_{i-1}\})\, DF(t,t_i)\, 1_{\{\tau>t_i\}}}_{\text{coupon payment upon survival}} \\
&+ \underbrace{(\tau - \max\{t,t_{i-1}\})\, DF(t,\tau)\, 1_{\{\max\{t,t_{i-1}\}<\tau\leq t_i\}}}_{\text{accr. coupon upon default}}.
\end{aligned}
$$

Notice that this definition of the discounted premium leg assumes the running coupon to be equal to one. Denoting the running coupon by c, the insurance buyer actually has to pay $c\, DPL(t,t_E,t_N)$. The present value at time t – provided default has not yet occurred – of the potential default compensation payment to be made by the protection seller (the *discounted default leg*) is given by

$$
DDL(t,t_E,t_N) := (1-R)\, DF(t,\tau)\, 1_{\{t<\tau\leq t_N\}},
$$

where $R \in [0,1]$ denotes the random recovery rate in case of a default event. The value of the CDS from the point of view of the portection buyer at time t – provided default has not yet occurred – is hence

$$
\begin{aligned}
&CDS(t,t_E,t_N) \\
&\quad = \mathbb{E}[DDL(t,t_E,t_N)\,|\,\mathcal{F}_t] - c\,\mathbb{E}[DPL(t,t_E,t_N)\,|\,\mathcal{F}_t], \quad (1)
\end{aligned}
$$

and the value $CDS(t_E,t_E,t_N)$ gives precisely the upfront amount to be paid by the protection buyer to the seller at inception t_E. Even though the upfront payment determines the CDS value, it is convention to quote a CDS in terms of a so-called *CDS running spread*[2]. The market quotes running spreads for CDS contracts whose first coupon payment is the next IMM date, i.e. $t = t_E = 0$ and CDS insurance starts immediately. In this case the running spread of a CDS is defined as the unique value which, when plugged in for the running coupon c in (1), makes the CDS-value zero. This means that a fictitious CDS contract with running coupon equal to this value trades at zero upfront. In mathematical terms, the running spread is defined (only for $t = t_E = 0$) by

$$
s_0 = s_0(0,t_N) := \frac{\mathbb{E}[DDL(0,0,t_N)]}{\mathbb{E}[DPL(0,0,t_N)]}.
$$

[2]Also called *par CDS spread*.

The conversion of the quoted running spread s_0 into an actually tradable upfront amount $CDS(0, 0, t_N)$ at given running coupon c is a standard routine in the market, relying on a rearrangement of the terms in (1). It is accomplished by computing

$$CDS(0, 0, t_N) = (s_0 - c)\,\mathbb{E}[DPL(0, 0, t_N)], \tag{2}$$

where the computation of the required expectation value by convention relies on the so-called ISDA standard model, which we briefly explain in the next section.

1.2 Standard pricing methods

Regarding the pricing of CDS, the most minimal modeling approach is to assume that the recovery rate R is a deterministic constant and τ has an exponential distribution with rate parameter λ, which is also called the (deterministic and constant) *default intensity* of τ. In this case, it can be shown[3] that the approximation $\boxed{s_0 \approx \lambda\,(1 - R)}$, which is based on the assumption of continuous rather than discrete premium payments in DPL and a flat interest rate curve, is not too bad. Applying standard recovery assumptions, the default intensity λ can be matched perfectly to the market quoted running spread s_0, which constitutes the so-called ISDA standard pricing model to be used by market participants for the conversion into actually tradable upfronts via (2). More generally, in order to define a model which explains a battery of quoted CDS contracts with different maturities (a so-called *CDS-curve*), it is market standard to relax the assumption of τ being exponentially distributed to the case when the default intensity of τ is deterministic and piecewise constant. The piecewise levels of the default intensity can then be bootstrapped from the observed market prices in order to explain the whole CDS curve jointly, see, e.g., Hull, White (2000). It is important to observe, however, that such a model does **not** suffice in order to price CDS options. This is due to the fact that the deterministic nature of the default intensity implies that market sentiment, whose timely evolution is captured in $(\mathcal{F}_t)$, has no effect on the default intensity, and hence on the future evolution of the running CDS spread. Consequently, within such a simplistic model the latter is a smooth, deterministic function over time until default. In reality, a CDS option is a financial product whose intention is to trade the uncertainty about the running CDS spread at a future point in time. For this, more involved mathematical approaches are necessary.

2 Single-name CDS options

In contrast to regular CDS contracts, a CDS option settled today at $t = 0$ gives its holder the right to enter as protection buyer into a CDS contract which starts at the option expiration date $t_E > 0$ in the future. The option contract specifies a certain *strike running spread* $s^{(K)}$ at which the future CDS can be entered into. For the protection buyer this means that if the running spread at time t_E is higher than the strike running spread, it might make sense to exercise the option. If not, then protection might instead rather be bought in the regular CDS market at lower cost without exercising the option. The CDS option is said to trade *knockout* if

[3] See, e.g., (Hull, 2008, p. 500).

it becomes worthless in case a credit event occurs before expiry, i.e. if $\tau < t_E$. It is said to trade *no-knockout*, if the option provides its holder with an additional *front end protection*, i.e. if the event $\tau < t_E$ triggers a default compensation payment to the option holder – provided the latter decides to exercise the option. From an economic viewpoint, the volatile nature of the daily changing running spread constitutes the major risk which is traded via CDS options. From a mathematical viewpoint, this suggests the use of a stochastic model for the running spread at the future time point t_E. Using the notations introduced above, it is natural to simply define the forward running spread as[4]

$$s_t = s_t(t_E, t_N) := \frac{\mathbb{E}[DDL(t, t_E, t_N) \,|\, \mathcal{F}_t]}{\mathbb{E}[DPL(t, t_E, t_N) \,|\, \mathcal{F}_t]}, \quad 0 \leq t < \tau. \qquad (3)$$

Intuitively, this constitutes a stochastic process which at time $t = t_E$ equals the then prevailing running spread, as desired. Unfortunately, in order to derive useful pricing formulas from this intuitive idea the following technical difficulties have to be overcome:

(a) The definition (3) is not well-defined for scenarios in which the credit event is triggered before the option expiry, i.e. $\tau < t_E$, because it is only defined for $t < \tau$. For $t \geq \tau$, the value s_t as written down in (3) would equal zero divided by zero, which makes no sense.

(b) When defining τ via the so-called *canonical construction*[5] from an exogenously modeled stochastic process for its default intensity, and plugging this model into (3), the target random variable s_{t_E} is rather complicated. In order to derive "simple" pricing formulas, e.g. a Black-type formula, one rather would like to model the quantity s_{t_E} exogenously. The justification of such an approach relies on measure-change techniques similar as in the interest rate derivative pricing literature.

Both issues have been dealt with in the literature, e.g. in Jamshidian (2004); Schönbucher (2004); Brigo, Morini (2005); Armstrong, Rutkowski (2009); Martin (2012), and we briefly summarize and comment on these approaches in the upcoming subsections.

2.1 The "old" zero upfront case

In this subsection we assume that the underlying CDS, that might be entered into at t_E, is available in the market at zero upfront, i.e. the running coupon c is unknown today but fixed at t_E at the then prevailing level of the CDS running spread, and no upfront payment is made at t_E. This situation has been market practice before the Big Bang Protocol 2009 and is untypical in the marketplace nowadays, but simplifies the mathematical derivation massively. In particular, it suffices to treat knockout options, because the price of a no-knockout option is then simply the price of the respective knockout option plus the expected value of the discounted protection leg of a CDS contract with maturity

No-knockout case?

[4] We mostly omit to denote the dependence of s_t on the underlying CDS covenants, i.e. we only write s_t instead of $s_t(t_E, t_N)$ for the sake of brevity.

[5] See, e.g., Jeanblanc, Rutkowski (2000).

t_E (the so-called *front end protection*). This decomposition of a no-knockout option into two simpler parts becomes more difficult when the underlying CDS trades at a pre-determined running coupon, as will be explained in the next subsection.
The model-free price of a knockout CDS option is obviously given by the expectation value

$$\begin{aligned} CDSO(0) &= \mathbb{E}\Big[DF(0,t_E)\,\Big(\mathbb{E}[DDL(t_E,t_E,t_N)\,|\,\mathcal{F}_{t_E}] \\ &\qquad - s^{(K)}\,\mathbb{E}[DPL(t_E,t_E,t_N)\,|\mathcal{F}_{t_E}]\Big)_+\,1_{\{\tau>t_E\}}\Big] \\ &= \mathbb{E}\Big[DF(0,t_E)\,\mathbb{E}[DPL(t_E,t_E,t_N)\,|\,\mathcal{F}_{t_E}]\,1_{\{\tau>t_E\}}\,\big(s_{t_E}-s^{(K)}\big)_+\Big]. \end{aligned} \tag{4}$$

We review two different approaches to derive Black-type formulas for knockout CDS options from this general expression, one by Schönbucher (2004) and one by Brigo, Morini (2005); Armstrong, Rutkowski (2009) along the general ideas of Jeanblanc, Rutkowski (2000); Jamshidian (2004). For practical implications, both approaches are equivalent, because they ultimately lead to the same Black-type formula. Nevertheless, we think it is educational to point out the subtle differences between the two approaches.

Schönbucher's ansatz

Schönbucher (2004) introduces a so-called *survival measure* $\mathbb{P}_S$ associated with the process $t \mapsto \mathbb{E}[DPL(t,t_E,t_N)\,|\,\mathcal{F}_t]$, which appears in the denominator of (3). As mentioned earlier, it is possible that this denominator process becomes zero, hence it cannot serve as a numeraire and accordingly does not give rise to an equivalent numeraire measure, as is common practice in interest rate derivatives pricing. Nevertheless, this is precisely what's achieved in Schönbucher (2004) by the so-called *survival measure* – modulo some small technical differences compared with the standard change of numeraire technique. Without going into details[6] it is shown that the value of the CDS option can be written as

$$\boxed{CDSO(0) = \mathbb{E}[DPL(0,t_E,t_N)]\,\mathbb{E}_S[(s_{t_E}-s^{(K)})_+],} \tag{5}$$

where $\mathbb{E}_S$ denotes the expectation with respect to the survival measure $\mathbb{P}_S$. Moreover, the process s_t can be shown to be a martingale under $\mathbb{P}_S$. Let us collect a couple of remarks on this formula.

(i) The formula is well-defined because $\mathbb{P}_S$ is constructed in such a way that default becomes impossible[7] under $\mathbb{P}_S$, i.e. s_t is well-defined for all t under $\mathbb{P}_S$ even though it is not well-defined under $\mathbb{P}$.

(ii) The formula is useful for applications in the sense that it allows to work with two "disjoint" models: one for the distribution

[6] The interested reader is referred to Schönbucher (2004).

[7] It is shown in (Schönbucher, 2004, Proposition 6) that $\mathbb{P}_S$ must intuitively be thought of as a numeraire measure associated with the denominator process in (3) conditioned on survival until t_N, where the "conditioning upon survival" makes the denominator strictly positive and hence a numeraire.

of the default time τ under $\mathbb{P}$ in order to compute the first expectation value (e.g. piecewise default intensity), and one for the dynamics of s_t under $\mathbb{P}_S$ in order to compute the second expectation value (e.g. lognormal dynamics to obtain a Black formula).

(iii) Theoretically, the spread s_t is linked to the model for τ, because it is composed of conditional expectations involving default indicators, cf. (3). When separately modeling the dynamics of s_t under $\mathbb{P}_S$ and the law of τ under $\mathbb{P}$, as is market practice, compatibility of the two models is not assured, but implicitely assumed. Moreover, such an approach has the additional shortcoming that there is no interrelation between spread volatility and default probability at all.

(iv) The measure $\mathbb{P}_S$ is not equivalent to the risk-neutral measure $\mathbb{P}$ because default is impossible under $\mathbb{P}_S$, which is consistent with the separation of the pricing into a defaultable model (for the first expectation) and a non-defaultable model (for the second expectation) in the following sense: when defining the dynamics of s_t under $\mathbb{P}_S$, our intuition might lead us to apply volatility parameters as observed from historical time series of the running spread, and these observed values are conditioned on survival, so might be thought of as being observed under $\mathbb{P}_S$ rather than $\mathbb{P}$.

(v) One drawback of this approach becomes apparent when two CDS options must be explained jointly, where the two underlying CDS contracts start at the same date t_E, but have different maturities. This would require a joint model for two running spreads with different maturities, but the measure $\mathbb{P}_S$ depends on the maturity.

Subfiltration ansatz

The approach of Brigo, Morini (2005); Armstrong, Rutkowski (2009), which is basically along the lines of Jamshidian (2004) and relying on techniques based on the fundamental work of Jeanblanc, Rutkowski (2000), differs from Schönbucher (2004) in that the technical problem (a) above is circumvented by putting more structure on the market filtration $(\mathcal{F}_t)$. More precisely, the natural filtration of the default indicator process $t \mapsto 1_{\{\tau \leq t\}}$ is clearly contained in $(\mathcal{F}_t)$, since a default is assumed to be observable in the marketplace. Consequently, it makes sense to think of $(\mathcal{F}_t)$ as being obtained as the combination of the default indicator filtration and a disjoint “rest information”, denoted $(\mathcal{H}_t)$. Indeed, in standard default intensity modeling frameworks $(\mathcal{H}_t)$ denotes the natural filtration of the default intensity, hence this assumption is motivated by this prominent modeling example. It is assumed that $\mathbb{P}(\tau > t \,|\, \mathcal{H}_t)$ is positive almost surely, which again is satisfied by typical examples of default intensity models. With this subfiltration framework at hand, it is possible to replace the conditional expectations with respect to $\mathcal{F}_t$ in (3) by conditional expectations with respect to $\mathcal{H}_t$, see Brigo, Morini (2005); Armstrong, Rutkowski (2009) for details. More clearly, they consider a slightly alternative running CDS spread defined by

$$\bar{s}_t := \frac{\mathbb{E}[DDL(t, t_E, t_N) \,|\, \mathcal{H}_t]}{\mathbb{E}[DPL(t, t_E, t_N) \,|\, \mathcal{H}_t]}, \quad t \geq 0.$$

This results in a definition of the running spread which is defined for all $t \geq 0$, and can therefore be used further to write down a formula for the CDS option as an expectation value under the model pricing measure $\mathbb{P}$ getting rid of the default indicator, namely[8]

$$\begin{aligned} &CDSO(0) \\ &\quad = \mathbb{E}\Big[DF(0, t_E)\, \mathbb{E}\big[DPL(t_E, t_E, t_N) \,\big|\, \mathcal{H}_{t_E}\big]\, (\bar{s}_{t_E} - s^{(K)})_+\Big]. \end{aligned}$$

In order to get rid of the ugly conditional expectation term inside this expectation, Brigo, Morini (2005); Armstrong, Rutkowski (2009) now apply the usual change of numeraire trick[9] from interest rate modeling in order to derive the same formula (5) as Schönbucher (2004), only the measure $\mathbb{P}_S$ is replaced by a numeraire measure $\bar{\mathbb{P}}$ associated with the denominator in the definition of $\bar{s}_t$, which is strictly positive and hence a valid numeraire. Moreover, it is shown that $\bar{s}_t$ is a martingale under $\bar{\mathbb{P}}$. For practical applications, e.g. the derivation of a simple Black-type formula, this approach is just as useful Schönbucher's, yielding the identical result for practical applications. Nevertheless, it is further outlined in Brigo, Morini (2005) how their approach is able to allow for an extension to modeling several running spreads associated with different maturities jointly. This is accomplished along the idea of LIBOR market models in the interest rate world and sets the approach apart from Schönbucher (2004).

2.2 The "new" upfront case

It is now assumed that the CDS, which is entered into at expiration t_E of the option, trades upfront, i.e. the strike running coupon $c^{(K)}$ to be paid is pre-determined and there is a strike upfront payment $u^{(K)}$ to be made at t_E. In this situation, which is usual since the Big Bang Protocol 2009, the model-free price of a knockout CDS option is given by the expectation value

$$\begin{aligned} CDSO(0) = \mathbb{E}\Big[DF(0, t_E)\, \Big(&\mathbb{E}[DDL(t_E, t_E, t_N) \,|\, \mathcal{F}_{t_E}] - u^{(K)} \\ &- c^{(K)}\, \mathbb{E}[DPL(t_E, t_E, t_N) \,|\, \mathcal{F}_{t_E}]\Big)_+ 1_{\{\tau > t_E\}}\Big] \\ = \mathbb{E}\Big[&DF(0, t_E)\, 1_{\{\tau > t_E\}}\, \big(CDS(t_E, t_E, t_N) - u^{(K)}\big)_+\Big], \end{aligned}$$

where both $u^{(K)}$ and $c^{(K)}$ are pre-determined in the option contract. This formula causes technical difficulties, because it suggests that we should model the CDS value $CDS(t, t_E, t_N)$, i.e. the upfront payment, exogenously by some stochastic process. Instead, we are rather used to modeling the running spread s_t,

[8] This computation follows from the fundamental identity $\mathbb{E}[Z \,|\, \mathcal{F}_t] = \frac{1_{\{\tau > t\}}}{\mathbb{P}(\tau > t \,|\, \mathcal{H}_t)}\, \mathbb{E}[Z \,|\, \mathcal{H}_t]$ of Jeanblanc, Rutkowski (2000), applied with $t = t_E$ and Z being the expression under the expectation in (4). This is briefly outlined in the Appendix.

[9] Actually, Brigo, Morini (2005) apply some further technical assumptions on the subfiltration structure, but Armstrong, Rutkowski (2009) point out that this is not necessary.

which is not directly incorporated. There are at least two reasons why it is more natural to model the running spread rather than the CDS value exogenously: **Firstly**, its popularity. The running spread is a single quantity providing information about the creditworthiness of the reference entity, whereas the upfront must always be communicated along with its associated running coupon. The running spread allows to be compared with other spread measures of the reference entity, e.g. bond Z-spreads, and this is probably the reason why it is popular in the market and standard quotation. **Secondly**, it is more natural to define a mathematical model for the running spread. A running spread is a positive quantity and our mathematical tool box is rich of positive stochastic processes capturing many observed stylized facts. In contrast, the CDS value may be positive or negative, is bounded from above and below, and hence more difficult to model directly.

How do we tackle this pricing problem?

To the best of our knowledge, Martin (2012) is the only reference dealing with this issue. His idea is to express the quantity $\mathbb{E}[DPL(t_E, t_E, t_N) \,|\, \mathcal{F}_t]$ – conditioned on survival until t_E – as a function of s_{t_E}, say $f(s_{t_E})$. This implies, together with the idea of the formula (2), that we may rewrite the CDS option value as

$$CDSO(0) = \mathbb{E}\Big[DF(0,t_E)\, 1_{\{\tau > t_E\}}\, \big((s_{t_E} - c^{(K)})\, f(s_{t_E}) - u^{(K)}\big)_+\Big],$$

which appears to be a convenient starting point at least for numerical computations when specifying dynamics for s_t. The disturbing default indicator in the latter expectation can now be got rid of by either applying Schönbucher's survival measure idea (this is proposed in Martin (2012)), or by imposing a subfiltration structure. By precisely the same steps as in Brigo, Morini (2005) we obtain the formula

$$\boxed{CDSO(0) = \mathbb{E}[DPL(0,t_E,t_N)]\, \bar{\mathbb{E}}\Big[\Big(\bar{s}_{t_E} - c^{(K)} - \frac{u^{(K)}}{f(\bar{s}_{t_E})}\Big)_+\Big],} \tag{6}$$

where the second expectation is again taken with respect to the measure $\bar{\mathbb{P}}$, under which $\bar{s}_t$ is a martingale. From a theoretical point of view, the assumption of the premium leg being given as a function of the spread is not justified but intuitive. From a practical point of view, an implementation of formula (6) requires an approximation for the function f. Akin to the ISDA standard model, one might choose f as the concatenation of (i) the function which finds the exponential rate parameter in the standard ISDA model matching the prevailing running CDS spread and (ii) the function which computes the expected value of the premium leg with the obtained exponential rate[10]. For later reference, we call this choice for f the "exact" choice. The approach taken in Martin (2012) is even one step simpler. A parametric form for f is assumed, which is motivated by the approximation $s_0 \approx \lambda\,(1-R)$

[10] Notice that f should also depend on the interest rate term structure at t_E, which might be specified as the ISDA standard for the respective currency. Using determinsitic discount factors, forward discount factors at time t_E should be used in the calculation.

discussed earlier. More precisely, it is assumed that

$$f(s_{t_E}) \approx \frac{1 - e^{-(r+(1+\epsilon)\, s_{t_E}/(1-R))\,(t_N - t)}}{r + (1+\epsilon)\, s_{t_E}/(1-R)}, \tag{7}$$

where R is the constant recovery rate parameter, r denotes a flat interest rate parameter, and $\epsilon > 0$ is a small correction parameter which is introduced in order to make the approximative formula for f match the "exact" choice of f when fitted to certain market observables. Alternatively, it is not too costly to just stick with the "exact" choice for f, which, however, cannot be written down in algebraically closed form, but instead must be evaluated numerically. For $t_E = 0.6192$ and $t_N = 5.6192$, using a (forward) discounting curve derived from 3m-tenor based EUR swap rates as prescribed by the standard ISDA method for EUR denominated CDS, Figure 1 visualizes the "exact" choice for f and the approximation (7) with $\epsilon = 0$ and different values for r. This clearly shows that the approximation is not bad –indeed almost perfect – provided one picks the right r, which seems to act like a parallel shift. This method of picking r is similar to Martin's method of picking ϵ (and choosing r differently). From a mathematical viewpoint, the parameter ϵ is redundant because its contribution can be shifted into the second parameter r. However, it might make sense to work with two parameters (ϵ, r), e.g. by deriving r somehow from interest rate data as proposed in Martin (2012), in order to equip it with some economic meaning.

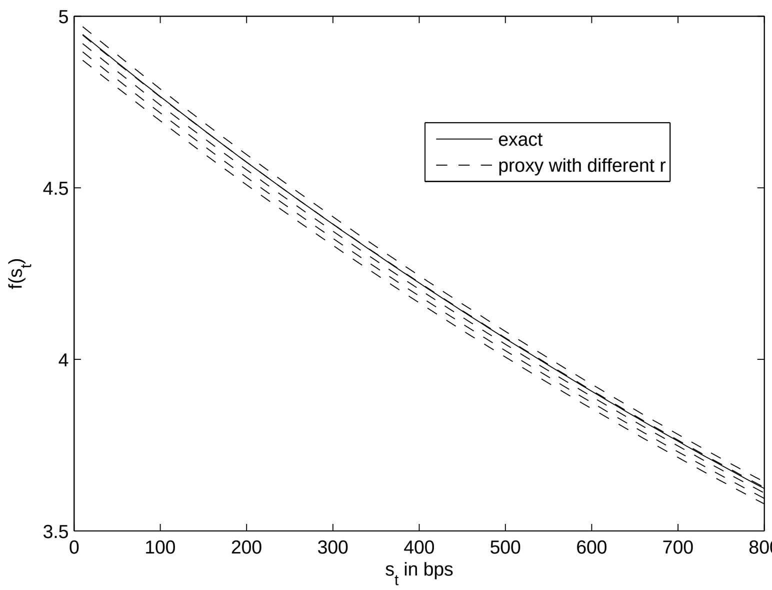

Fig. 1: Visualization of the "exact" choice of the function f, and its approximation (7) with different values of r and $\epsilon = 0$, for an exemplary forward-starting CDS with maturity 5y.

It is obvious that the formula (6) boils down to the formula (5) of the previous subsection in the special case $u^{(K)} = 0$ and $c^{(K)} = s^{(K)}$. This is the only case in which a Black-type formula is obtained when assuming a lognormal distribution for the running spread. In all other cases, the expectation value must be evaluated numerically by integration against the lognormal density, which is not a big issue in practice, however.

Practical implication

A natural, interesting question is to what extent the new upfront quotation has an effect on CDS option prices. An at-the-money CDS option should be one in which the pair $(u^{(K)}, c^{(K)})$ corresponds to the underlying CDS contract having the value zero from today's point of view. This means the values $c^{(K)}$ and $u^{(K)}$ are in a one-to-one relationship via the following algorithm.

Algorithm 1 (Standard upfront conversion)
Input are the strike coupon $c^{(K)}$ and the currently prevailing CDS running spread for a CDS that starts immediately and has the same maturity (i.e., $1y$, $5y$ etc., not the same maturity date!) as the underlying CDS of the option.

- Derive an exponential rate parameter λ from the given running spread by means of the standard ISDA model.
- Using the ISDA standard model from the first step, compute $\mathbb{E}[DPL(0, t_E, t_N)]$ and $\mathbb{E}[DDL(0, t_E, t_N)]$.
- Return $u^{(K)} = \mathbb{E}[DDL(0, t_E, t_N)] - c^{(K)}\,\mathbb{E}[DPL(0, t_E, t_N)]$.

This algorithm constitutes the market standard for computing $u^{(K)}$ from a given strike running coupon c_K. The value of the CDS option is not invariant under the choice of these pairs, see Figure 2. The higher the strike running coupon $c^{(K)}$, the smaller the associated upfront $u^{(K)}$, which is paid immediately upon exercise of the option. For instance, if the upfront is negative, then the option holder might have an incentive to exercise the option even though it is out of the money at t_E. This is a striking difference from call option contracts in equity markets.
Similarly, an out-of-the-money CDS option should be one in which the pair $(u^{(K)}, c^{(K)})$ corresponds to the underlying CDS contract having a negative value from today's point of view (of the protection buyer). Having set the strike coupon $c^{(K)}$ to a level above the currently prevailing running spread, it is reasonable to compute an associated upfront via Algorithm 1, where the currently prevailing running spread in the input parameters is replaced by $c^{(K)}$.

No-knockout case?

In contrast to the previous subsection, a no-knockout option is slightly more complex compared with a knockout option. The reason is that if the credit event occurs before expiry t_E, exercising the option (which is necessary in order to receive default compensation) provides the holder with the cash flow $1 - R - u^{(K)}$, which is only positive when the upfront is smaller than the default compensation. Hence, the front end protection is not received unconditionally, but instead contingent on the fact that $1 - R > u^{(K)}$. With a constant recovery assumption, the pricing is achieved similarly as in the previous subsection as the sum of a knockout option and the front end protection. The latter, however, is now multiplied with the indicator $1_{\{1-R>u^{(K)}\}}$. Since the front end protection obviously contains a significant amount of recovery risk, the use of a stochastic recovery model might be reasonable. The latter, however, complicates the pricing, especially when independence between recovery rate and spread level cannot be assumed.

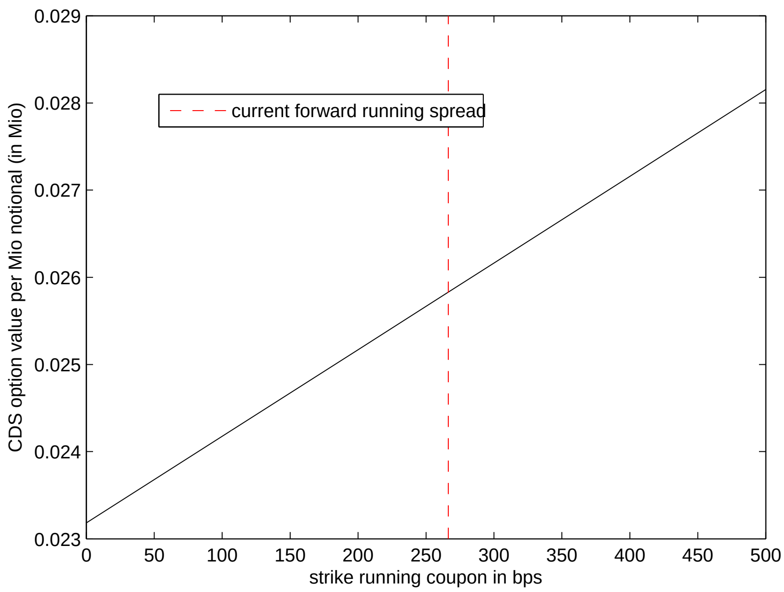

Fig. 2: At-the-money CDS options on an exemplary CDS with maturity 5y for different pairs $(u^{(K)}, c^{(K)})$. The spread $\bar{s}_{t_E}$ is assumed to have a lognormal law with volatility parameter $\sigma = 60\%$. The dotted red line indicates the current forward running spread, i.e. the pair $\big(0, s_0(t_E, t_N)\big)$. It is obvious that the option value increases in the strike coupon.

3 Conclusion

We reviewed the mathematical derivations of the market standard formulas for single-name CDS options.

Appendix

It is briefly explained how to get rid of the disturbing indicator in (4) via the subfiltration structure. To be precise, it holds that

$$\begin{aligned}
&\mathbb{E}\Big[DF(0,t_E)\,\Big(\mathbb{E}[DDL(t_E,t_N)\,|\,\mathcal{F}_{t_E}] \\
&\qquad\qquad - s^{(K)}\,\mathbb{E}[DPL(t_E,t_N)\,|\mathcal{F}_{t_E}]\Big)_+ 1_{\{\tau>t_E\}}\Big] \\
&= \mathbb{E}\Big[DF(0,t_E)\,\Big(\mathbb{E}[DDL(t_E,t_N)\,|\,\mathcal{H}_{t_E}] \\
&\qquad\qquad - s^{(K)}\,\mathbb{E}[DPL(t_E,t_N)\,|\mathcal{H}_{t_E}]\Big)_+ \frac{1_{\{\tau>t_E\}}}{\mathbb{P}(\tau>t_E\,|\,\mathcal{H}_{t_E})}\Big] \\
&= \mathbb{E}\Big[DF(0,t_E)\,\Big(\mathbb{E}[DDL(t_E,t_N)\,|\,\mathcal{H}_{t_E}] \\
&\qquad\qquad - s^{(K)}\,\mathbb{E}[DPL(t_E,t_N)\,|\mathcal{H}_{t_E}]\Big)_+\Big],
\end{aligned}$$

where the last equality is obtained from the tower property of conditional expectation when conditioning the expression under expectation on $\mathcal{H}_{t_E}$ and using the fact that all terms except for the indicator are measurable with respect to $\mathcal{H}_{t_E}$.

Acknowledgements

Helpful comments by German Bernhart on earlier versions of this manuscript are gratefully acknowledged.

References

A. Armstrong, M. Rutkowski, Valuation of credit default swaptions and credit default index swaptions, *International Journal of Theoretical and Applied Finance* **12** (2009) pp. 1027–1053.

D. Brigo, M. Morini, CDS market formulas and models, *Proceedings of the 18th Annual Warwick Option Conference, available at http://www.damianobrigo.it/cdsmktfor.pdf* (2005).

J. Hull, A. White, Valuing credit default swaps I: no counterparty default risk, *Journal of Derivatives* **8:1** (2000) pp. 29–40.

J. Hull, Options, futures and other derivatives, *Pearson Prentice Hall, 7th edition* (2008).

F. Jamshidian, Valuation of credit default swaps and swaptions, *Finance and Stochastics* **8** (2004) pp. 343–371.

M. Jeanblanc, M. Rutkowski, Default risk and hazard processes, *in Mathematical Finance – Bachelier Congress 2000, eds. H. Geman et. al., Springer* (2000) pp. 281–312.

R. Martin, A CDS option miscellany, *Working paper, available at http://arxiv.org/abs/1201.0111* (2012).

P. Schönbucher, A measure of survival, *Risk Magazine* **17** (2004) pp. 79–85.