INDEX CDS OPTIONS: A REVIEW OF PRICING APPROACHES

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

Date: February 2, 2017

Abstract This is a survey of methods proposed in the literature and the marketplace regarding the pricing of index CDS options. The challenges of the topic are highlighted, and the heavy assumptions on which common formulas rely are pointed out.

1 Introduction and notation

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 risk-neutral pricing measure (under which the discounted versions of all tradable assets are $(\mathcal{F}_t)$ martingales). For simplicity, discount factors are assumed to be deterministic and denoted by $DF(t,T)$ throughout[1]. An index credit default swap (index CDS) is an insurance contract between two parties. The insurance buyer makes periodic premium payments to the protection seller. In return, the insurance buyer is compensated by the protection seller for losses occuring in a reference basket of credit-risky assets during the lifetime of the contract. More precisely, denoting the number of assets in the reference basket by d and the default time of asset k in the basket by τ_k, which is assumed to be an $(\mathcal{F}_t)$-stopping time, the relative number L_t of defaulted assets in the basket at time t, and the remaining nominal N_t at time t, are given by

$$L_t = \frac{1}{d} \sum_{k=1}^{d} 1_{\{\tau_k \leq t\}} \in [0,1], \quad N_t = 1 - L_t, \quad t \geq 0.$$

The protection buyer pays a quarterly coupon c at each IMM date during the contract life time on the remaining nominal. We denote the last IMM date before settlement of the contract by t_0, and the IMM dates during the contract lifetime by $t_1, \ldots, t_N$, with t_N being at the same time the maturity of the contract. If the index CDS is settled at t_E this means that $t_0 \leq t_E < t_1 < \ldots < t_N$. Denoting the index CDS running coupon by c, the present value at time t of the sum over all (clean) cash flows to be made by the

[1] One should think of them as being given in terms of a reference short rate r_t via $DF(t,T) = \exp\Big(-\int_t^T r_u \,\mathrm{d}u\Big)$.

protection buyer (the *discounted premium leg*) is given by

$$c \cdot DPL(t,t_N) := c \cdot \sum_{i\,:\,t<t_i} \underbrace{(t_i - \max\{t,t_{i-1}\})\,DF(t,t_i)\,N_{t_i}}_{\text{coupon payment}}$$
$$+ \underbrace{\frac{t_i - \max\{t,t_{i-1}\}}{2}\,DF(t,t_i)\,(N_{\max\{t,t_{i-1}\}} - N_{t_i})}_{\approx \text{ accrued coupon upon default(s)}}.$$

The term in the second line above is just an approximation of the aggregated accrued coupon payments to be made in case of observed defaults during t_{i-1} and t_i. It vanishes if no default is observed in the respective period and is based on (i) the assumption of all default times occuring in the middle of the interval $[\max\{t,t_{i-1}\},t_i]$ (so-called *midpoint rule*) and (ii) the assumption of all accrued coupon payments being made at the end of each period. Upon default of one asset in between t_{i-1} and t_i the protection buyer receives at the end of the period t_i the compensation payment $(1-R_i)/d$ per unit of nominal, where $R_i \in [0,1]$ denotes the recovery rate of the asset i. For the sake of simplicity it is assumed that $R_1 = \ldots = R_d =: R$ for a constant R in the sequel, i.e. all recovery rates are assumed to be non-random and identical to all names in the equally weighted basket. Consequently, the present value at time t of the default compensation payments to be made by the protection seller (the *discounted default leg*) is given by

$$DDL(t,t_N) := (1-R) \sum_{i\,:\,t<t_i} DF(t,t_i)\,(L_{t_i} - L_{\max\{t,t_{i-1}\}}).$$

The value of the index CDS contract for the protection buyer at time t is hence given by

$$ICDS(t) = \mathbb{E}[DDL(t,t_N)\,|\,\mathcal{F}_t] - c\,\mathbb{E}[DPL(t,t_N)\,|\,\mathcal{F}_t].$$

Assume for a moment that $t_E = 0$, i.e. the CDS settles immediately. At inception of the contract the value $ICDS(0)$ has to be paid by the protection buyer to the protection seller (if it is negative, the protection buyer receives money), and it is called the *upfront payment* of the index CDS. The coupon rate c satisfying $ICDS(0) = 0$ is called *index CDS running spread* and denoted by s_0. Furthermore, for later reference we define s_t as the unique root of the equation $ICDS(t) = 0$ for all $t \geq 0$, which is well-defined for all $t < \tau_{[d]} := \max\{t_1,\ldots,t_d\}$, i.e. before all names in the basket default. It is given by

$$s_t = \frac{\mathbb{E}[DDL(t,t_N)\,|\,\mathcal{F}_t]}{\mathbb{E}[DPL(t,t_N)\,|\,\mathcal{F}_t]}, \quad t_E = 0 \leq t < \tau_{[d]}. \tag{1}$$

If $t_E > 0$, i.e. the index CDS settles in the future, it is also called a *forward index CDS*, and the quantity

$$s_{0,t_E} := \frac{\mathbb{E}[DDL(t_E,t_N)\,|\,\mathcal{F}_0]}{\mathbb{E}[DPL(t_E,t_N)\,|\,\mathcal{F}_0]}.$$

is called the *forward index CDS running spread*, which equals today's market-expected index CDS running spread for a forward

index CDS settled at t_E. Similarly, we may define a stochastic process s_{t,t_E} for all $0 \leq t < \tau_{[d]}$ as the unique root of the equation $ICDS(t) = 0$, explicitely denoting the dependence on the settlement date. Clearly, $s_{t,0} = s_t$.
It is market convention that index CDS with immediate settlement $t_E = 0$ are quoted[2] in terms of their running spread s_0, while the coupon c is standardized, e.g. to 100 bps. In order to compute the corresponding index CDS value/ upfront payment $ICDS(0)$, one needs to compute

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

Since the involved expectation value is model-dependent in general, but a price agreement of both contractual parties requires a common basis, it is market convention to compute $\mathbb{E}[DPL(0, t_N)]$ in the conversion formula (2) as a function of s_0, i.e. $f_0(s_0) := \mathbb{E}[DPL(0, t_N)]$, as:

(a) Assume that $\tau_1 = \ldots = \tau_d =: \tau$ and that τ has an exponential distribution with rate $\lambda > 0$.

(b) Choose λ such that $ICDS(0) = 0$ under assumption (a), when the coupon is given by the input spread s_0, i.e. $c = s_0$, and discount factors are obtained from a battery of ISDA-defined swap rates by standard bootstrapping routines.

(c) With λ obtained from (b), $f_0(s_0) = \mathbb{E}[DPL(0, t_N)]$ is computed under the assumption (a).

For later reference we also define the functions f_t for $t > 0$ via $f_t(s_0) := \mathbb{E}[DPL(t, t_N)]$, which are defined in precisely the same way as above, only replacing $t = 0$ by arbitrary $t \geq 0$. An efficient approximation for the function f_t is given by

$$f_t(s_0) \approx \frac{1 - e^{-\left(r + \frac{s_0}{1-R}\right)\,(t_N - t)}}{r + \frac{s_0}{1-R}}, \tag{3}$$

which relies on the assumption of continuous CDS coupon payments and a flat short rate $r_t \equiv r$ used for discounting all cash flows via $DF(t, T) = \exp(-r\,(T - t))$.

2 Index CDS options

An index CDS option settled at $t = 0$ is a financial contract providing its holder the right, but not the obligation, to enter as protection buyer into a (forward) index CDS at option expiry $t_E > 0$, which is at the same time the settlement date of the underlying index CDS. The option holder specifies a *strike index CDS running spread* $s^{(K)}$ at which the underlying index CDS can be settled, even though the then prevailing index CDS running spread s_{t_E,t_E} might be higher than $s^{(K)}$. More clearly, in view of the market conversion formula (2), upon exercise of the option at time t_E the option holder has to pay the upfront $(s^{(K)} - c)\, f_{t_E}(s^{(K)})$ on the full nominal[3] $N_0 = 1$, rather than $(s_{t_E,t_E} - c)\, f_{t_E}(s_{t_E,t_E})$ on

[2] Similarly, forward index CDS are quoted in terms of their forward index CDS running spread s_{0,t_E}.

[3] Notice that N_0 is the remaining nominal in the basket at option settlement $t = 0$. It is assumed here that all names in the basket are still alive at $t = 0$. If not, the nominal of the contract simply has to be adjusted in all formulas.

the then prevailing remaining nominal N_{t_E}, which would be the upfront of a regular (non-forward) index CDS at time t_E – provided not all assets in the basket have defaulted until t_E, in which case s_{t_E,t_E} is not defined. By convention, an index CDS option always trades *no-knockout* meaning that its holder is compensated for defaults prior to t_E, i.e. receives the payment $(1-R)\,L_{t_E}$, if the option is exercised[4]. Consequently, the model-free value of the option equals

$$\boxed{\begin{aligned} ICDSO(0) := \mathbb{E}\Bigg[DF(0,t_E)\,\bigg(&(1-R)\,L_{t_E} \\ &+ 1_{\{\tau_{[d]}>t_E\}}\,\Big\{(s_{t_E,t_E}-c)\,f_{t_E}(s_{t_E,t_E})\,N_{t_E} \\ &\qquad - (s^{(K)}-c)\,f_{t_E}(s^{(K)})\Big\}\bigg)_+\Bigg]. \end{aligned}} \tag{4}$$

Notice that all pricing formulas presented in the next section are simplifications from the general formula (4) which rely on more or less (un)realistic assumptions. The latter are required, because a direct evaluation of formula (4) within a model capturing all desired risks involved in an index CDS option contract is challenging without resorting to time-consuming Monte Carlo engines.

2.1 Pricing

Presupposing the standard assumption of deterministic interest rates, which makes the discount factor $DF(0,t_E)$ non-random, the remaining challenges in the evaluation of the general formula (4) are as follows.

(i) Noticing that $\{\tau_{[d]} > t_E\} = \{L_{t_E} < 1\}$, two random objects appear under the expectation in (4), and in theory they are not stochastically independent. These are L_{t_E} and s_{t_E,t_E}. Unfortunately, it is not straightforward to single out one of these two objects in a separate expectation value, because they are not linearly combined.

(ii) A bottom-up model for both quantities based on d correlated default intensity processes, which would be the canonical and most intuitive ansatz, is not very feasible, since it is almost impossible to evaluate the expectation value (4) without time-consuming Monte Carlo algorithms due to the high dimensionality.

(iii) Appealing to the well-known single name CDS option case, the market is used to think of the major driver of randomness s_{t_E,t_E} as a lognormal random variable. Similarly, the stochastic object L_{t_E} is the main driver of randomness in so-called CDO tranche contracts, for which numerous efficient pricing models exist. However, these two structurally different model types cannot easily be combined in a joint model for both, without giving up a lot of practical viability.

There exists some literature tackling this challenging pricing issue. Armstrong, Rutkowski (2009) provides a very reader-friendly

[4] It might happen that a default occurs before expiry but $s_{t_E,t_E} < s^{(K)}$. In this case, exercise of the option only makes sense if the default compensation payment exceeds the upfront payoff loss.

and rigorous derivation of standard pricing formulas and market conventions, and the approach is identical to the one in Brigo, Morini (2009, 2011). In particular, an explanation for the simpler formula of Pedersen (2003), relying on additional simplifying assumptions, is included, and it is justified how a market-conventional, simple Black-type formula is obtained as an approximation of the more rigorous formulas. The approaches of Jackson (2005) and Martin (2012) are different and have to be discussed separately. All mentioned references have in common that their goal is to impose a lognormality assumption for the spread s_{t_E,t_E} in order to derive Black-type formulas, while the "disturbing quantity" L_{t_E} is treated with different levels of carefulness by imposing convenient assumptions or introducing tricky measure-changes during the derivation.

2.2 Derivation of Black-type formulas

We distinguish between simple approaches which ignore the default compensation term L_{t_E} completely, and such approaches which take this term into account – at least in a simplified way.

2.2.1 Approaches ignoring default compensation

Assumption 1: $L_{t_E} = 0$

The simplest way to get rid of the disturbing random variable L_{t_E} is to assume $L_{t_E} = 0$, i.e. to assume that there will be no default until expiry of the option. This simplifies pricing massively, since we are left with essentially only one random object being the spread s_{t_E,t_E}. More precisely, the pricing formula simplifies to

$$ICDSO(0) = DF(0, t_E) \times \\ \times \, \mathbb{E}\Big[\Big((s_{t_E,t_E} - c)\, f_{t_E}(s_{t_E,t_E}) - (s^{(K)} - c)\, f_{t_E}(s^{(K)})\Big)_+\Big].$$

Using standard measure-changing techniques inherited from the single-name CDS option case, it is further possible to rewrite the latter formula in terms of two expectation values as

$$\boxed{ICDSO(0) = DF(0, t_E)\, \mathbb{E}[DPL(t_E, t_N)] \times \\ \times \, \bar{\mathbb{E}}\Big[\Big((s_{t_E,t_E} - c) - \frac{(s^{(K)} - c)\, f_{t_E}(s^{(K)})}{f_{t_E}(s_{t_E,t_E})}\Big)_+\Big],} \quad (5)$$

where the second expectation is taken with respect to a measure $\bar{\mathbb{P}}$ under which $\{s_{t,t_E}\}_{t\in[0,t_E]}$ is a martingale. For mathematical details, the interested reader is referred to Mai (2014). Formula (5) may be evaluated numerically under the assumption of a lognormal distribution for the random variable s_{t_E,t_E}. However, this is not yet a Black formula because of the function f_{t_E}. Nevertheless, formula (5) allows to evaluate the two involved expected values separately. For the first expectation, required are only default probabilities, which are extracted by the standard ISDA-model from the forward index CDS running spread s_{0,t_E}. For the second expectation, required is only a model for the martingale process $\{s_{t,t_E}\}_{t\in[0,t_E]}$ for which a geometric Brownian motion, hence a lognormal law for s_{t_E,t_E}, is canonical.

Assumption 2: no-upfront trading

The market typically quotes prices for index CDS options in terms of implied Black volatilities, which are computed by an even simpler formula than (5), which relies on the assumption that the underlying forward index CDS trades no-upfront (which is not the case in the marketplace, of course). This means that the forward index protection buyer has to pay $s_{t_E,t_E}\, f_{t_E}(s_{t_E,t_E})$, which has to be compared with the payments $s^{(K)}\, f_{t_E}(s_{t_E,t_E})$ that would have to be made in case of an index CDS option exercise. Consequently, it can be shown that formula (5) is replaced by

$$\boxed{\begin{aligned} ICDSO(0) = DF(0,t_E)\,\mathbb{E}[DPL(t_E,t_N)] \times \\ \times\, \bar{\mathbb{E}}\Big[(s_{t_E,t_E} - s^{(K)})_+\Big]. \end{aligned}} \tag{6}$$

Similar to formula (5), formula (6) allows to evaluate the expressions $\mathbb{E}[DPL(t_E,t_N)]$ and $\hat{\mathbb{E}}\Big[(s_{t_E,t_E} - s^{(K)})_+\Big]$ separately. For the first expectation, required are only default probabilities, which are extracted by the standard ISDA-model from the observed forward index CDS running spread s_{0,t_E}. For the second expectation, required is only a model for the martingale process $\{s_{t,t_E}\}_{t\in[0,t_E]}$ for which a geometric Brownian motion, hence a lognormal law for s_{t_E,t_E}, is canonical, leading to a Black-type formula. The latter formula is used by the market in order to convert prices to implied Black volatilities.

2.2.2 Approaches with default compensation

We briefly sketch the derivation of Armstrong, Rutkowski (2009), which basically coincides with the one of Brigo, Morini (2009, 2011). These approaches avoid the assumption $L_{t_E} = 0$ of the previous paragraph. Although the mathematical derivation is essentially analogous to the single-name CDS option case, the economic interpretation appears a lot more awkward in the multi-name case, so we sketch the logic in the sequel, pointing out the numerous necessary assumptions made.

Assumption 1: subfiltration structure

First of all, it is natural to assume that the market filtration $(\mathcal{F}_t)$ includes the natural filtration of the indicator $t \mapsto 1_{\{\tau_{[d]} \leq t\}}$, because the market participants observe defaults. Consequently, we may think of $\mathcal{F}_t$ as being generated by this information and a disjoint "rest information", which we denote by $\mathcal{H}_t$. We take further for granted the assumption $\mathbb{P}(\tau_{[d]} > t \,|\, \mathcal{H}_t) > 0$ almost surely for all t. Since the event $\{\tau_{[d]} > t\}$ is not contained in $\mathcal{H}_t$ by construction, the latter assumption is rather natural and seems to be not too restrictive. It is required later on in order to be able to guarantee positivity of the process $t \mapsto \mathbb{E}[DPL(t,t_N) \,|\, \mathcal{H}_t]$, so that it can be used as a numeraire. The resulting change of numeraire will help us to get rid of the disturbing indicator $1_{\{\tau_{[d]} > t_E\}}$ in formula (4).

Assumption 2: simplified payoff

The second assumption[5], which is imposed in order to simplify the payoff in formula (4), is

$$f_{t_E}(s^{(K)}) \approx f_{t_E}(s_{t_E,t_E})\, N_{t_E} \approx \mathbb{E}[DPL(t_E,t_N) \,|\, \mathcal{F}_{t_E}],$$

on the event $\{\tau_{[d]} > t_E\}$, where these quantities are well-defined. Notice that the second approximation is relevant because the

[5] It is precisely this assumption, which makes technical problems due to an upfront trading of index CDS disappear. In the single-name case we do not make this assumption and, consequently, end up with a technical difference between zero-upfront traded CDS and upfront-traded CDS, see Mai (2014) for details.

function f_{t_E} computes the premium leg in a simplified way based on the single-name standard ISDA model, which needs not be consistent with the computation of $\mathbb{E}[DPL(t_E, t_N)\,|\,\mathcal{F}_{t_E}]$ within a given multivariate model for the default times $(\tau_1, \ldots, \tau_d)$. Furthermore, this assumption is clearly critical since it equates the constant $f_{t_E}(s^{(K)})$ with a random quantity that is most likely a decreasing function of the index CDS running spread s_{t_E,t_E}.
Having these two assumptions at hand, the crucial step in the derivation is to handle the accumulated loss variable L_{t_E}. To this end, it is useful to define a modified running spread quantity, denoted *loss-adjusted running index CDS spread*, and given via

$$\hat{s}_{t,t_E} := \frac{\mathbb{E}[1_{\{\tau_{[d]} > t_E\}}\, DF(t, t_E)\,(1-R)\,L_{t_E} + DDL(t, t_N)\,|\,\mathcal{H}_t]}{\mathbb{E}[DPL(t, t_N)\,|\,\mathcal{H}_t]}.$$

With this artificial spread definition, formula (4) allows to be rewritten as the sum of two simpler expectation values, namely

$$\begin{aligned} ICDSO(0) &= \mathbb{E}\Big[DF(0, t_E)\,(1-R)\,L_{t_E}\,1_{\{\tau_{[d]} \le t_E\}}\Big] \qquad (7)\\ &+ \mathbb{E}\Big[DF(0, t_E)\,\mathbb{E}[DPL(t_E, t_N)\,|\,\mathcal{H}_{t_E}]\,(\hat{s}_{t_E,t_E} - s^{(K)})_+\Big], \end{aligned}$$

a justification is provided in the Appendix for the sake of clarity. The second expectation value might now be attacked by a standard change of numeraire technique with the numeraire process $t \mapsto \mathbb{E}[DPL(t, t_N)\,|\,\mathcal{H}_t]$, yielding ultimately the useful pricing formula

$$\boxed{\begin{aligned} ICDSO(0) = DF(0, t_E)\Big(&(1-R)\,\mathbb{P}(\tau_{[d]} \le t_E)\\ &+ \mathbb{E}[DPL(0, t_N)]\,\hat{\mathbb{E}}[(\hat{s}_{t_E,t_E} - s^{(K)})_+]\Big). \end{aligned}} \qquad (8)$$

In the latter formula $\hat{\mathbb{E}}$ denotes the expectation with respect to a measure $\hat{\mathbb{P}}$ related to the aforementioned numeraire. Moreover, it can be shown that $\{\hat{s}_{t,t_E}\}_{t \in [0,t_E]}$ is a martingale under $\hat{\mathbb{P}}$. Similar as in the single-name case, see Mai (2014), formula (8) achieves a convenient separation between a model for $(\tau_1, \ldots, \tau_d)$ under $\mathbb{P}$ and a second model for the evolution of $\hat{s}_{t,t_E}$ under $\hat{\mathbb{P}}$. The conclusions from this observation are almost identical to the conclusions in the single-name case. However, let us point out one problem for practical applications that is not present in the single-name case: when imposing lognormal dynamics on $\hat{s}_{t,t_E}$ under $\hat{\mathbb{P}}$, required is the start value $\hat{s}_{0,t_E}$, but the market participants rather like to have a formula involving the observable quantity s_{0,t_E}. The difference between the two quantities depends only on the expected number of defaults in the basket until t_E.

Assumption 3: the spreads s_t and $\hat{s}_t$ differ only by a computable constant

The further simplifying assumption

$$\begin{aligned} \hat{s}_{t,t_E} - \underbrace{\frac{\mathbb{E}[DDL(t, t_N)\,|\,\mathcal{H}_t]}{\mathbb{E}[DPL(t, t_N)\,|\,\mathcal{H}_t]}}_{=:\hat{s}_{t,t_E}} &= \frac{\mathbb{E}[1_{\{\tau_{[d]} > t_E\}}\, DF(t, t_E)\,(1-R)\,L_{t_E}\,|\,\mathcal{H}_t]}{\mathbb{E}[DPL(t, t_N)\,|\,\mathcal{H}_t]}\\ &\approx \frac{\mathbb{E}[(1-R)\,DF(t, t_E)\,L_{t_E}]}{\mathbb{E}[DPL(t, t_N)]} =: \epsilon(t) \end{aligned}$$

can help to circumvent this issue. Define $\bar{s}_{t,t_E}$ like s_{t,t_E} only with expectations with respect to $\mathcal{F}_t$ in numerator and denominator replaced by expactations with respect to $\mathcal{H}_t$, so that $\bar{s}_{0,t_E} = s_{0,t_E}$, as desired. It can be shown that $\{\bar{s}_{t,t_E}\}_{t\in[0,t_E]}$ is also a martingale under $\hat{\mathbb{P}}$, and that the pricing formula (8) further simplifies to

$$\boxed{\begin{aligned} ICDSO(0) &= DF(0,t_E)\,\Big((1-R)\,\mathbb{P}(\tau_{[d]}\leq t_E) \\ &\quad + \mathbb{E}[DPL(0,t_N)]\,\hat{\mathbb{E}}\Big[\Big(\bar{s}_{t_E,t_E}-\big(s^{(K)}-\epsilon(0)\big)\Big)_+\Big]\Big). \end{aligned}} \tag{9}$$

It is now clear how the remaining expectation with respect to $\hat{\mathbb{P}}$ in (9) gives rise to a Black-formula which, in contrast to (8), relies on an adjusted strike spread and the observable spread s_{0,t_E} as input. It is further pointed out by Armstrong, Rutkowski (2009) that some market participants completely neglect the first term, because $\mathbb{P}(\tau_{[d]}\leq t_E)\approx 0$. With a similar far-fetched argument, one could also drop the strike-adjustment (i.e. set $\epsilon(0)=0$), which then ultimately leaves one with precisely the standard market-typical Black formula (6).

3 Summary

The existing methods for pricing index CDS options have been surveyed, with a focus on the mathematical techniques behind them. In particular, it has been pointed out that the formulas used in the marketplace rely on quite heavy and unrealistic assumptions. It has also been explained why the derivation of more realistic pricing formulas is quite challenging.

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, Last option before the armageddon, *Risk Magazine* (2009) pp. 78–83.

D. Brigo, M. Morini, No-armageddon arbitrage-free equivalent measure for index options in a credit crisis, *Mathematical Finance* **21** (2011) pp. 573–593.

A. Jackson, A new method for pricing index default swaptions, *Working Paper, Global Credit Derivatives Research, Citigroup, London* (2005).

J.-F. Mai, Pricing single-name CDS options: an overview, *XAIA homepage article* (2014).

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

C.M. Pedersen, Valuation of portfolio credit default swaptions, *Quantitative Credit Research, Lehman Brothers* (2003).

Appendix

We briefly justify the derivation of formula (7) from formula (4) under the two simplifying assumptions mentioned. The second assumption readily implies

$$\begin{aligned} ICDSO(0) &= \mathbb{E}\Big[DF(0,t_E)\,\Big((1-R)\,L_{t_E} \\ &\quad + 1_{\{\tau_{[d]}>t_E\}}\,\mathbb{E}[DPL(t_E,t_N)\,|\,\mathcal{F}_{t_E}]\,(s_{t_E,t_E}-s^{(K)})\Big)_+\Big]. \end{aligned}$$

XAIA INVESTMENT

Considering the two disjoint events $\{\tau_{[d]} > t_E\}$ and $\{\tau_{[d]} \le t_E\}$ separately, the first expectation value in (7) simply corresponds to the armageddon event $\{\tau_{[d]} \le t_E\}$. On the event $\{\tau_{[d]} > t_E\}$ we compute

$$\begin{aligned}
&\mathbb{E}\Big[1_{\{\tau_{[d]}>t_E\}}\, DF(0,t_E)\,\Big((1-R)\,L_{t_E} \\
&\quad + 1_{\{\tau_{[d]}>t_E\}}\,\mathbb{E}[DPL(t_E,t_N)\,|\,\mathcal{F}_{t_E}]\,(s_{t_E,t_E} - s^{(K)})\Big)_+\Big] \\
&= \mathbb{E}\Big[1_{\{\tau_{[d]}>t_E\}}\, DF(0,t_E)\,\Big((1-R)\,\mathbb{E}[L_{t_E}\,|\,\mathcal{F}_{t_E}] \\
&\quad + \mathbb{E}[DDL(t_E,t_N)\,|\,\mathcal{F}_{t_E}] - s^{(K)}\,\mathbb{E}[DPL(t_E,t_N)\,|\,\mathcal{F}_{t_E}]\Big)_+\Big] \\
&= \mathbb{E}\Big[\frac{1_{\{\tau_{[d]}>t_E\}}}{\mathbb{P}(\tau_{[d]} > t_E\,|\,\mathcal{H}_{t_E})}\, DF(0,t_E)\,\Big((1-R)\,\mathbb{E}[L_{t_E}\,|\,\mathcal{H}_{t_E}] \\
&\quad + \mathbb{E}[DDL(t_E,t_N)\,|\,\mathcal{H}_{t_E}] - s^{(K)}\,\mathbb{E}[DPL(t_E,t_N)\,|\,\mathcal{H}_{t_E}]\Big)_+\Big] \\
&= \mathbb{E}\Big[\frac{1_{\{\tau_{[d]}>t_E\}}}{\mathbb{P}(\tau_{[d]} > t_E\,|\,\mathcal{H}_{t_E})}\, DF(0,t_E) \times \\
&\qquad \times \mathbb{E}[DPL(t_E,t_N)\,|\,\mathcal{H}_{t_E}]\,(\hat{s}_{t_E} - s^{(K)})_+\Big] \\
&= \mathbb{E}\Big[DF(0,t_E)\,\mathbb{E}[DPL(t_E,t_N)\,|\,\mathcal{H}_{t_E}]\,(\hat{s}_{t_E,t_E} - s^{(K)})_+\Big].
\end{aligned}$$