Pricing Covered Bonds

Chris Kenyon*

19 June 2009

1 Introduction

Covered bonds[1] have emerged as a potentially preferred funding vehicle from the credit crisis, and they were already a major part of many European financial systems. Although some prices are directly available on, for example Bloomberg[2], there is no detailed examination of how covered bonds should be priced taking into account the features that make them attractive to investors: i.e. over-collateralization of the bankruptcy-remote reference pool of assets, and covenants by the pool manager (issuer) on asset replacement. This study provides the first pricing methods for covered bonds based on a Triggered Refreshed CDO with Issuer Risk model that we introduce here.

Since covered bonds are relatively new to the pricing literature we include a generic description in Section 2.

We define a Refreshed CDO as one where the first m assets that *default* are replaced. A Triggered Refreshed CDO is one where replacement is triggered by a credit event other than default, e.g. downgrade, restructuring, etc. m expresses the liquidity available from the issuer to replace assets. A further step is to include issuer default not caused by the coverpool directly. We define this as a Triggered Refreshed CDO with Issuer Risk. Issuer Risk is important because, whilst the assets of the coverpool are bankruptcy-remote, the covenants on the manager are not, and the manager may default for reasons not directly related to the coverpool.

We provide analytic pricing for Triggered Refreshed CDOs with or without Issuer Risk, i.e. covered bonds, based on the factor Copula model. The analytic pricing is based both on direct calculation of the underlying asset default distribution, and an alternate recursive characterization. Since both of these methods are numerically intensive, involving integration over high

*DEPFA BANK Plc., chris.kenyon{at}depfa.com. The views expressed are those of the author only, no other representation should be attributed.

[1]Aka: Pfandbriefe (Germany); Obligations Foncières (France); Lettres de Gage (Lux); Cédulas Hipotecarias (Spain).

[2]For example Bloomberg pages JPBP, MSEA, BXPF, BPXF, locked but can be opened on request.

dimensions, we propose an accurate analytic approximation: "new-on-last refresh". We also provide exact and approximate simulation algorithms. The approximate algorithms allow addition of coverpool features to standard CDO algorithms with no change to their internals.

The contributions of this paper are threefold: first the introduction of the Triggered Refreshed CDO with Issuer Risk as a model for covered bonds; second exact pricing methods (analytic and Monte Carlo); and third an accurate approximation approach (analytic and Monte Carlo versions) that enables coverpool features to be added to previous standard CDO analyses or code. To the authors' knowledge these are the first pricing algorithms for covered bonds in the literature.

2 Covered Bonds

Covered bonds are bonds that are issued on an over-collateralized pool of assets that are bankruptcy-remote from the issuer. This coverpool of assets is also typically protected by legal covenants on the issuer, and potentially by issuer covenants agreed with rating agencies.. Thus covered bonds must be supported by legislation or their bankruptcy-remoteness will be questioned by the rating agencies [Hof09]. Specific details are embodied in the relevant legislation and rating agency methods [3]. A good introduction to the European context is [ECB08]. The reader is warned not to assume that any detail here is correct; the authors explicitly repeat the usual caveats and that this work is their opinion only and should not be used as a basis for any action or inaction.

Covered bonds are issued *pari passu* and *pro rata* on a reference pool of assets that are legally ring-fenced, i.e. bankruptcy-remote from the pool manager (the issuing bank). Typically the pool is rated by rating agencies (taking into account the notes issued on the pool) and the manager aims to keep the rating above a certain level. Whilst the manager might legally be able to allow the rating to drop there would be consequences: firstly the manager (i.e. the bank) might have difficulty issuing further covered bonds; secondly, if some of the notes from the pool were being used to provide liquidity (e.g. via the Fed or ECB), then those notes would cease to be eligible leaving the bank with a liquidity issue.

Legally a coverpool must typically be over-collateralized in nominal terms. The degree of over-collateralization is considered by the rating agencies with fairly clear criteria. Lately, the agencies have increased the over-collateralization ratio for banks to keep the rating of their covered bonds. Generally, assets are either related to soverign entities or related to mortgages (Pfandbriefe in Germany). Here we focus on the soverign case. Extension to mortgage related assets introduces complexity on the assets but not

[3] See http://ecbc.hypo.org/Content/Default.asp for a list.

a fundamental change in the coverpool mechanism. Sovereign-related assets are for example, sovereigns, local authorities, or fully-owned public sector enterprises (e.g. hospitals or toll roads). Coverpools must be diverse in terms of the percentage of each issuer type and geography. Coverpools are limited to a group of approved geographies, e.g. Canada, USA, Japan, Switzerland and the European Economic Area (which includes Iceland). There may be legal restrictions on rating for some or all geographical components.

3 Modelling

We use the stylized setting of vanilla CDOs, i.e. hazard rates are deterministic (default is of course stochastic), as are interest rates. We use a factor-based Copula pricing approach that we extend to deal with covered bond features.

We ignore requirements on asset-liability matches on issuances, e.g. duration, cashflows, etc and assume the pool is well-managed. Since the match may not be perfect the degree of liquidity of the pool is included by rating agencies to understand the potential for forced sales of illiquid assets. However, we do not include this level of detail. We also assume that if one bond is replaced it is replaced by exactly one bond so the number of underlyings in the pool is constant.

Notes issued on the coverpool are essentially a single $x\%$, to 100% tranche of a CDO where x is the degree of over-collateralization. Clearly if the over-collateralization drops below the legal minimum the manager would change from the bank to the regulating government. At that point we do not expect changes in the coverpool beyond an orderly runoff with the notes being paid *pari passu* and *pro rata*. We must now include the protection features of the covered bonds.

If an asset drops below a given rating, which we call the replacement rating, then it must be replaced to maintain coverpool rating. We call this refreshing the CDO. Thus an asset can only drop below the replacement rating if the supporting organization, which we call here the manager, has defaulted. Note that for "rating" we can substitute any other credit event that is significant for the coverpool manager relative to the asset in question (e.g. restructuring, or a changes in liquidity of the issuer of the underlying asset, or a change in the credit spread of underlying asset, etc.). We apply the same hazard rate setting for all event types.

In a factor-based approach there are common and idiosyncratic risks that lead to default of the underlyings. Here, replacing coverpool assets may also lead to default of the manager who may have insufficient liquidity to replace them. We term this replacement risk. Note that if replacement is free then the covered bonds are only sensitive to issuer risk (i.e. if the manager of the coverpool defaults, even free replacements are not available).

Below we first review factor-based pricing and then show how to price a refreshed CDO with replacement risk. A simple form of this would be to treat the liquidity available for replacement as subordination for the coverpool. However, this ignores the fact that replacement is required on a credit event not default. Hence we can view the replacement capability of the manager as bonds that will default on downgrade. Although bonds can recover from some credit events (that are not default events) we are only interested in the cost caused by the credit event (generally the "replacement trigger event"). Taking an extra-subordination approach would also ignore the fact that (some) liquidity is held in default-free instruments (up to default of the manager). We tackle this issues below.

We thus have the following models:

Refreshed CDO : the first m defaults by the n underlyings don' t matter, only the last n matter, m can be greater than n. This is the essential difference w.r.t. subordination — the liquidity cannot default. However the m is limited precisely by the fact that replacement reduces capital causing exhaustion of liquidity.

Triggered Refreshed CDO : the first m replacements have a different probability to the subsequent n, i.e. the refresh trigger is not default.

Triggered Refreshed CDO with Issuer Risk : m depends on $n + 1$ of the underlying event generators (i.e. the issuer).

We can include assets outside the coverpool within the modeling framework. We term this model a Triggered Stochastically Refreshed CDO because the replacement capacity (capital and replacement assets) is stochastically dependent on defaults (mostly) outside the coverpool. This is covered in Appendix I so as to keep the flow of the main text as clear as possible.

4 Pricing

We first recall factor-based CDO pricing, then build up to pricing a Refreshed CDO exactly using a series of Lemmas. We provide both direct (Lemma 5) and recursive (Lemma 6) calculations. The key distribution, in the general case, is n dimensional (even with factor conditioning) requiring integration over n dimensions for pricing. Hence we introduce an approximation, new-on-last-refresh. This means that *all* the assets are taken to start anew at the time of the last refresh. This is an approximation in the factor setting because, even with constant hazard rates and memoryless marginal distributions, the factor-conditional distributions are not memoryless. We show in the Example Results section that, for the coverpool considered, the pricing error of this approximation is bounded by ± 1 basis point for compound correlation up to 60%.

Following the analytic section we give simulation algorithms for exact and approximate pricing.

Next we consider a Triggered Refreshed CDO and then in Section 4.4 we include Issuer Risk, i.e. the possibility that the manager defaults independently to the coverpool. Issuer default moves m to zero when it happens because there is no longer any external capital to support the coverpool.

Proofs are in Appendix II.

4.1 Factor-based Pricing

We follow the usual setup, see [HW04, LG05, Sch05]. We give a Gaussian copula setting but any other copula can be directly substituted here. We assume that the marginal default probabilities are available, e.g. from CDS quotes. Let:

t_i: default time of i-th issuer.

$Q_i(t)$: cumulative risk-neutral default probability of the i-th issuer, hence the survival probability is given by $S_i(t) = 1 - Q_i(t)$.

$x_i = a_i H + \sqrt{1 - a_i^2} Z_i$, for a 1-factor model.

H, Z_i are independent zero-mean, unit-variance random variables, and $-1 \leqslant a_i < 1$; H represents the common risk and Z_i the idiosyncratic risks of the issuers.

F_i, G: cumulative distribution of x_i, Z_i respectively.

Under the copula model x_i maps to t_i under percentile to percentile. Hence we have:

Prob$(x_i < x|H) = G\left[\frac{x - a_i H}{\sqrt{1-a_i^2}}\right]$, so Prob$(t_i < t|H) = G\left[\frac{F^{-1}[Q_i(t)] - a_i c}{\sqrt{1-a_i^2}}\right]$

So conditional survival probability of i-th asset beyond T is just:

$$S_i(T|M) = 1 - G[].$$

4.2 Refreshed CDO

We continue within the factor setting. A Refreshed CDO is one where asset replacement is triggered by default of the asset. Each replacement has some cost (from the cost of default) but does not change the coverpool until this forces manager default (because the manager runs out of capital). Each of the n assets in the coverpool can require replacement and the replacement can also require replacement etc. Thus there are n generators of replacement events in the coverpool up to the point when no more replacements occur, i.e. the mth replacement.

We provide four pricing methods: exact analytic; approximate analytic; exact simulation; and approximate simulation. The approximation and simulation approaches are necessitated by our demonstration that an exact analytic approach requires integration over a large number of dimensions.

4.2.1 Valuation: Exact Analytic

The first m events do not affect the pool and the pool has a fixed number, n, assets. In a factor model, conditional on the conditioning factor c (possibly vector) all n underlyings default independently. In the refreshed case each underlying can generate multiple replacement events (default, downgrade, etc.). Prior to the mth replacement event we call the underlyings event generators.

Pricing relies on the following observations:

Obs0 Replacements are chosen at future points in time, not at the start.

Obs1 The n refreshed underlyings (generators) are conditionally independent and the successive survival times (to events) from each generator are also conditionally independent.

Obs2 Successive defaults of each generator occur at strictly increasing times (because they are sums of survival times).

Obs3 The mth default uses at most one default time from each generator. For example if $m = 5$ and $n = 3$ one way to get the 5th default is with the 2nd default from the first generator and the 3rd default from the second generator (assuming that these occur before the first default from the third generator). Two different default times from the same generator are never used.

Obs4 The expectation of a sum of random variables does not depend on the correlation between them.

Obs0 and Obs1 are key to understanding time-to-mth replacement in high correlation scenarios. For ease of exposition consider constant hazard rates, so the forward hazard rate is the same as the spot hazard rate. Suppose all the coverpool assets have a credit event at almost the same time (because of high correlation), then provided that assets in the general market exist that are OK (w.r.t. not having credit events), then these assets basically start anew when they enter the coverpool (because they are picked in the future and, we assume, the manager can always pick assets that are only now starting, i.e. genuinely new). If the manager had to designate potential replacements at the start (time zero) then, of course, the statistics would be completely different.

Obs3 is important because it means that although the conditional default times of the kth and $(k + 1)$th conditional default times from an individual generator are correlated it does not matter for this calculation. They are correlated because they are sums and their first k terms are the same.

Obs4 is important because it means that we can add contributions from different paths of default without caring that the distributions of the different paths are correlated.

We work up to the price of a Refreshed CDO using a couple of Lemmas.

Lemma 1. *If there are n independent generators of independent identically distributed events (per generator) with probability distribution functions of time between events $f_i(t)$, $i = 1, \ldots, n$ (and cumulative distribution functions $F_i(t)$), then the probability distribution function of the time of the mth event, e_m is:*

$$P[e_m = t] = \sum_{\alpha \in Comp(m,n)} \sum_{i=1}^{n} P[e_m = t \ \cap \ e[i]_{\alpha[i]} = t \ \cap_{j \neq i} e[j]_{\alpha[j]} \leqslant t]$$

where:

$$\begin{aligned} P[e_m = t \ \cap \ e[i]_{\alpha[i]} = t \ \cap_{j \neq i} e[j]_{\alpha[j]} \leqslant t] = 1_{\{\alpha[i]>0\}} \prod_{j=1}^{j=n} \Bigg(& 1_{\{i=j\}} f_{i*\alpha[i]}(t) \\ & + 1_{\{i \neq j\}} \Bigg(1_{\{\alpha[j]=0\}} (1 - F_j(t)) \\ + 1_{\{\alpha[j] \neq 0\}} \Bigg(F_{j*\alpha[j]}(t) - \int_{s=0}^{s=t} & f_{j*\alpha[j]}(s) F_j(t-s) \ ds \Bigg) \Bigg) \Bigg) \end{aligned}$$

and:

e_m is the random variable giving the time of the mth event (considering all generators).

$\cap$ means the intersection of the events on either side.

$e[i]_\beta$ is the random variable giving time of the βth event for ith generator. Since each of the generators generates iid events the pdf of $e[i]_\beta$ is given by $f_{i*\beta}$.

$\cap_{j \neq i} \xi_j$ means, for say $i = 2$ and $n = 4$, $\cap \ \xi_1 \cap \xi_3 \cap \xi_4$, i.e. intersection over all possible values of j except i.

1_a is the indicator function on the statement "a" (e.g. "a" could be $\{i = j\}$).

$\vec{f}$ is the vector of pdfs $f_i(t)$, $i = 1, \ldots, n$.

$f_{i*\beta}$ is the βth convolution of the probability density function f_i.

$F_{i*\beta}$ is the cumulative distribution function of $f_{i*\beta}$.

$Comp(m, n)$ is the set of integer compositions of m into n parts.

α is an element of $Comp(m, n)$, it is a vector of length n with non-negative integer entries that sum to m.

To give an example of $\alpha \in Comp(m, n)$ consider $Comp(5, 3)$, elements of this are $(5, 0, 0)$, $(0, 5, 0)$, $(3, 1, 1)$, etc. In fact there are $\binom{m+n-1}{n-1}$ elements of $Comp(m, n)$. If there are restrictions on the coverpool composition and hence on the number of events a generator can generate (imagine defaults of Länder in Germany — there are only a given number of Länder available) then we can use the restricted composition $Comp(m, n, b)$ where b is the upper bound, see [Opd08] for more details. However, we do not go into this further.

Proofs are in the Appendix II but we will give an informal description here of how the lemma works. We want the time of the mth event. When the mth event happens the number of events from all of the generators add up to m. This can happen in many ways, so we must sum over all possible ways, i.e. over all $\alpha \in Comp(m,n)$. If the mth event happens at time t this means that one of the n generators must have an event exactly at time t and all the other generators must have their events at or before t. Thus we must sum over all of the generators, letting each have their last, $d[i]$th event at time t. Of course if a generator has zero events then that gives zero contribution, hence the indicator function on $\alpha[i], 1_{\{\alpha[i]=0\}}$. The rest of the equations follow from the probability of having exactly q events before t.

Note also that by assuming that the generators generate identically distributed events we are implying that forward hazard curves for replacements are identical to spot hazard curves. At a cost of some notational complexity this could be addressed. However, see Obs0, replacements are chosen at future times, and the asset chosen in the future may not even exist at the spot date. We ignore these complexities here.

As well as the distribution of the time of the mth event e_m, we will also need the joint distribution of the times of the events $e[j]_{\alpha[j]} = v[j]$ conditional on the time of the mth event. Here $\vec{v}$ are the times of the $\alpha[j]$th events from each of the j events where j runs over $1, \ldots, n$. Of course one of the $v[j]$ is equal to the time of the mth event. This is contained in the following lemma.

Lemma 2. *If there are n independent generators of independent identically distributed events (per generator) with probability distribution functions of time between events $f_i(t)$, $i = 1, \ldots, n$ (and cumulative distribution functions $F_i(t)$), then given of the time of the mth event, $e_m = u$, and the number of events of each generator $\vec{\alpha}$, the joint distribution of the last event times of the other generators before u is:*

$$P[e_m = u \;\cap_i\; e[i]_{\alpha[i]} = \vec{v}[i] \cap \max \vec{v} = u] = \sum_{j=1}^{n} P[e_m = u \;\cap_i\; e[i]_{\alpha[i]} = \vec{v}[i] \cap \vec{v}[j] = u]$$

and:

$$P[e_m = u \;\cap_i\; e[i]_{\alpha[i]} = \vec{v}[i] \cap \vec{v}[j] = u] = 1_{\{\alpha[j]>0\}} \prod_{i=1}^{i=n} \Bigg(1_{\{i=j\}} f_{i*\alpha[i]}(u) + 1_{\{i \neq j\}} \Big(1_{\{\alpha[i]=0\}} (1 - F_j(u)) + 1_{\{\alpha[i]\neq 0\}}\, f_{i*\alpha[i]}(v[i])(1 - F_i(u - v[i])) \Big) \Bigg)$$

This is somewhat of a simplification of the previous Lemma. Now we need the form of conditional distributions, these are given by the next Lemma.

Lemma 3. *Conditional on the conditioning factor c the probability density function and cumulative probability density of the default time in a 1-factor Gaussian copula model are respectively $q_{i,|c}(t)$ and $Q_{i,|c}(t)$ and are given by:*

$$
\begin{aligned}
Q_{i,|c}(t) &= G\left[\frac{F^{-1}[Q_i(t)] - \sqrt{\rho_i}c}{\sqrt{1-\rho_i}}\right] \\
q_{i,|c}(t) &= \frac{d}{dt}Q_{i,|c}(t)
\end{aligned}
$$

where the marginal (unconditional) cumulative probability density of default time is $Q_i(t)$.

This is simply a restatement of the standard copula factor formulae (from the Factor-based Pricing section above) in more appropriate language for the analysis here.

If we assume that after m refresh events the Refreshed CDO (R-CDO) probability of u further defaults is as for the conventional CDO, then we only need one simple lemma to calculate the R-CDO value (Lemma 7).

However, this is an approximation because, conditional on the conditioning factor c, the distributions are not memoryless. That is, the distributions $Q_{i,|c}(t)$ in Lemma 3 are not in general Exponential even when the unconditional distributions are. (Note that this is different from the ability of the manager to pick new replacement assets that can genuinely be starting at the replacement time).

In the general case, the probability that an asset has defaulted depends on the time of the last replacement of that asset. Thus, even if the assets are homogeneous, the view that the coverpool has of them will not be homogeneous when replacements are taken into account (although the distribution will be homogeneous with respect to asset exchange). Hence we need to calculate the vector of default probabilities at time t given that m replacements are available. We use the following results.

N.B. although direct calculation must take into account the pattern of replacements over the assets, asset homogeneity is a symmetry property so the output joint distribution will be symmetric with respect to asset exchange (of course it will not be the product of independent distributions). Although some results for distributions of exchangeable variables are available (following B. de Finetti [Fel70]) we do not follow that route here.

Lemma 4. *If a generator generates independent identically distributed events with probability distribution function $f(t)$, (and cumulative distribution function $F(t)$), then the probability distribution function of the $k+1$st event by* *time t given that k events have taken place by u, $EvK(f,k,t,u)$ is:*

$$EvK(f,k,t,u) = \begin{cases} \int_{s=0}^{s=u} f_{*k}(s) \frac{F(t-s)-F(u-s)}{1-F(u-s)} ds, & k > 0 \\ \frac{F(t)-F(u)}{1-F(u)}, & k = 0 \end{cases}$$

The probability of an event by t if the kth event was at u is F(t-u).

We can now put all this together since the value of a CDO coupon given a set of independent default probabilities is simple to calculate (e.g. via the bucketing algorithm in [HW04]).

Lemma 5. *If the value of a CDO coupon given a vector of independent default probabilities $\vec{p}(t)$ at t is* $\mathrm{CDO}(\vec{p}(t))$*, then the value of a refreshed CDO coupon,* R-CDO*, given that the generators have iid events with the time between events distributed as $\vec{f}$ is:*

$$\begin{aligned} \text{R-CDO} &= P[e_m > t]\mathrm{CDO}(\vec{0}) + \\ & \int_{u=0}^{u=t} \sum_{\alpha \in Comp(m,n)} \sum_{i=1}^{i=n} \\ & \int_{\max(\vec{v}) \leqslant u,\ \vec{v}[i]=u} \mathrm{CDO}(\vec{p}\ |\xi(i,u,\vec{v},\alpha))\ P[\xi(i,u,\vec{v},\alpha)] d\vec{v}\ du \end{aligned}$$

where:

$\xi(i,u,\vec{v},\alpha)$ *are the events s.t.* $\{e_m = u \ \cap \ e[i]_{\alpha[i]} = u \ \cap_{j \neq i} e[j]_{\alpha[j]} = \vec{v}[j]\}$*;*

$P[e_m > t]$ *is given by Lemma 1;*

$P[\xi(i,u,\vec{v},\alpha)]$ *is given by Lemma 2;*

$$p[j]|\xi(i,u,\vec{v},\alpha) = \begin{cases} F[j](t-u) & \textit{if } j = i \\ EvK(f[j],d[j],t,u) & \textit{otherwise} \end{cases}$$

from Lemma 4.

It is clear from this Lemma that even if $f[j]$ are identical the elements of $\vec{p}$ will not be.

Valuation using the analytic expressions in Lemma 5 directly is numerically intensive since it involves, amongst other things, an $n+1$ dimensional integral. We can address this in two ways: either approximate; or use simulation. An obvious approximation suggests itself: assume that all assets start anew when the last refresh happens. This changes the $n+1$ dimensional integral to a 1-dimensional integral in the problem. We explore this below in Section 4.2.2 and report results in the Results section.

The structure of the problem also suggests an alternate characterization of the next-default-time for each of the underlying assets via recursion. This is demonstrated in the Lemma below.

Lemma 6. *Given the independent distributions of first default times of n assets $\vec{f}$, then if the first m have been replaced, the distribution of the next default times of the assets $next(\vec{f}, m)$ is:*

$$\begin{aligned} next(\vec{f}, m) &= \sum_{i=1}^{i,n} 1_{\{next(\vec{f},m-1)=\min(next(\vec{f},m-1))\}} next(\vec{f}, m-1) * f[i] \\ &\quad + 1_{\{next(\vec{f},m-1)\neq\min(next(\vec{f},m-1))\forall i\}} next(\vec{f}, m-1) \end{aligned}$$

where:

$next(\vec{f}, 0)$ is the distribution the first default times of the n assets (i.e. after zero replacements).

$*$ means convolution.

If more than one point is equal to the minimum then the convolution is applied to all with weighing proportional to $next(\vec{f}, m-1)$.

The characterization in Lemma 6 is direct and also illustrates how an exact simulation works.

Whilst Lemma 6 is simple and could be used in Lemma 5 it still requires integration over an n-dimensional domain. This is because the n, originally-independent, distributions are mixed together by the min applied over all of them.

4.2.2 Valuation: Approximate Analytic: New-On-Last-Refresh

We now assume that all assets are as new on last refresh, i.e. memoryless at that point (only)). From Lemmas 1 and 3 above we have the conditional probability distribution function that the m refreshes are exhausted at any time s from zero up to t_i. Thus we arrive at:

Lemma 7. *The price of a coupon of a Refreshed CDO_i paying at time t_i is, in terms of the conventional CDO_i coupon is:*

$$Refreshed\ CDO_i = df(t_i) \int_c \int_{s=0}^{s=t_i} CDO_i(\vec{g}(s, t_i), t_i)\ pdf(c)\ ds\ dc$$

where:

$df(t_i)$ is the discount factor.

$CDO_i(\vec{f}, t_i)$ is the payoff of a conventional CDO coupon i at t_i given a vector of (conditionally independent) default probabilities of its underlyings at time $\vec{g}$.

$\vec{g}(s, t_i) = Ev(\vec{f}_{re|c}, n, m, s) \times \vec{f}_{def|c}(s, t_i)$

$Ev(\vec{f}|c, n, m, s)$ is defined in Lemma 1.

$\vec{f}_{re}|c$ are the probability distribution functions of replacement events (labelled "re") conditional on the conditioning factor c.

$\vec{f}_{def|c}(s, t_i)$ are the probabilities of default events (labelled "def") conditional on the conditioning factor c for underlyings between times s and t_i.

$pdf(c)$ is the probability distribution function of the conditioning factor.

c is the conditioning factor (potentially a vector).

4.2.3 Valuation: Exact Simulation

Exact simulation is shown in Algorithm 1. This is simple compared to analytic calculation. It has, however, the disadvantage that it does the whole pricing itself. This cannot be used to adapt a pre-existing factor-based CDO pricing algorithm where the user does not have access to the full internals of the coupon pricing code. The new-on-last-refresh approximation requires less access to (or change of) the internals of previous pricing algorithm codes and thus enables more reuse of pre-existing codes.

input : Numbers of: samples to generate g; underlyings n (equal weight); replacements m; coupon dates $\vec{t}$.
output: list of vectors of default fractions at each coupon date $\vec{h}(g)$

h← {};
for $i \leftarrow 1$ **to** g **do**
    c ←conditioningFactor;
    get the first event from each of the n underlyings;
    $\vec{e} \leftarrow$ `event(c)` ;
    replace the first m events;
    for $j \leftarrow 1$ **to** m **do**
        $e(\min^{-1}(\vec{e})) \leftarrow$ `event(c)` $+ \min(\vec{e})$;
    end
    get the fraction defaulted on each coupon date;
    temp← {};
    for $k \leftarrow 1$ **to** `Length(`$\vec{t}$`)` **do**
        temp←temp $\cup \sum_{l=1}^{l=n} 1_{\{e(l)<t(k)\}}/n$
    end
    h←h ∪ temp;
end

Algorithm 1: exact simulation for R-CDO (∪ means "append to"). The conditioning factor is drawn randomly at each iteration i, as are the events. An arrow over a symbol indicates a vector.

4.2.4 Valuation: New-On-Last-Refresh Simulation

In this case we need to find the time of the mth event from n generators (and also the first event for the $n + 1$st generator). A direct method would be to simulate m events from each generator (except the $n+1$st of course, if Issuer Risk is included), sort them, and take the time of the mth. This would be of complexity $O(m \times n)$, assuming that event generation is significantly more costly than sorting. A more efficient algorithm is given in Algorithm 2.

input : Numbers of: underlyings n; replacements m; conditioningFactor.
output: time of last refresh t_l

if $m > 0$ **then**
  get the first replacement event from each of the n underlyings;
  $\vec{e} \leftarrow$ `event(conditioningFactor)` ;
  replace the first m underlyings that have replacement events;
  for $j \leftarrow 1$ **to** m **do**
    $e(\min^{-1}(\vec{e})) \leftarrow$ `event(conditioningFactor)` $+ \min(\vec{e})$;
  end
  $t_l \leftarrow \min(\vec{e})$;
else
  $t_l \leftarrow 0$;
end
return t_l;

Algorithm 2: new-on-last-refresh simulation of last refresh time for R-CDO. The conditioning factor is an input from the overall simulation of the CDO; it must be the same at that used for the coupon calculations. An arrow over a symbol indicates a vector. Only the time of the last replacement is output; this is the start time for all the underlyings with the new-on-last-refresh approximation.

The complexity of this algorithm is an initial $O(n)$ event generations, one minimum finding per iteration, and $O(m)$ iterations. Worst case minimum finding costs $n - 1$ comparisons but this is a very cheap operation. The complexity will be dominated by the $O(m + n)$ event generations.

The simulation algorithm works in the context of factor copula models, i.e. it is within an outer loop that sets the factor. The *same* factor is used for the conditional calculations for the CDO. Obviously if the time of the mth replacement is after a coupon date for the CDO then there are no losses to any tranche at that coupon date. Note that this approach can directly use forward hazard rates for replacements.

4.3 Triggered Refreshed CDO

Here the first m refresh events use different marginal densities Q_i^{tr} for the generators as compared with subsequent default events Q_i. No other change is required.

4.4 Triggered Refreshed CDO with Issuer Risk

The effect of issuer default for covered bonds is different to a standard CDO because the coverpool is bankruptcy-remote from the issuer. Issuer bankruptcy does however still have an effect, it means that no further replacements occur.

We model issuer default within the factor Copula approach by including the issuer as the $n+1$st event generator. The effect of an event from this generator is to reduce the further possible replacements to zero. That is, at any time the "conventional" CDO starts if, either m replacements have occurred, or if the $n+1$st generator has generated a single event.

Pricing just requires a simple modification to $Prob[e_m = t]$ of Lemma 1, i.e.:

Lemma 8. *If there are $n+1$ independent generators of independent identically distributed events (per generator) with probability distribution functions $f_i(t)$, $i = 1, \ldots, n+1$ (and cumulative distribution functions $F_i(t)$) for events from the $n+1$ generators respectively, then the probability distribution function of* **either** *the mth event from the first n generators by time t* **or** *the first event from generator $n+1$, $Prob[e_{last} = t]$ (considering all the $n+1$ generators) is:*

$$Prob[e_{last} = t] = Prob[e_m = t] + F_{n+1}(t) - Prob[e_m = t]F_{n+1}(t)$$

This can be used directly with the new-on-last-refresh approach. Modifying the n-dimensional joint distribution of later lemmas is possible but we leave it out in interest of conserving space. It is direct to include this in the exact and approximate simulation algorithms.

4.5 Cash vs. Synthetic Instruments

Clearly all covered bonds are cash instruments. Note however that typically there is no provision for any speed-up in the return of principal as the coverpool defaults. We follow [HW04] in considering the difference between synthetic valuation and cash valuation. [HW04] argued that the value of a cash CDO tranche is the value of the corresponding synthetic CDO tranche plus the remaining principal of the tranche. Since the covenants only affect when the actual CDO mechanism starts, and nothing after that, the same logic holds for Refreshed / Triggered / with Issuer Risk / CDO variants too.

Thus we can follow [HW04] in stating that the spread above the risk-free rate (breakeven rate in Hull & White's terminology) on a covered bond is the same as that on the breakeven rate on the synthetic equivalent. Clearly this does not hold when there are mismatches in duration, cashflows etc. However, we are not considering this level of detail here.

5 Example Results

As an example, we analyze a hypothetical stylized coverpool at the concentration level of "public sector enterprises" (PSE), e.g. cities, Länder, etc., with $n = 100$. The level determines the order of magnitude of replacement m and the number of underlyings (assets) in the coverpool n. An alternative would be to address the concentration level of sovereigns. For sovereigns the n and m of interest would be an order of magnitude less. However, it is not clear that if a sovereign defaults all PSE within that country also default.

To take one example in the case $(m, n) = (20, 100)$ we assume that the issuer has liquidity equal to the cost of replacing 20% of the coverpool. Typically there would be hundreds to thousands of assets in a coverpool but they are suject to the concentration, i.e. several bonds from the same PSE or several hundred depending on the same sovereign. Although a coverpool issuer may only have spare capital of 4% to 8% the may also be assets outside the coverpool that can add to this. Hence we consider a range up to 30%, of course if the issuer defaults this spare capital vanishes. In Appendix I we consider how assets outside the coverpool may be included explicitly.

Any parameters we pick as typical are likely to be overtaken by events given the recent volatility of most financial markets. For illustration we pick a spread of $100bps$ for our homogeneous PSEs and $500bps$ for our issuer. Correlation is the remaining parameter, note that since we have only one tranche of interest we use compound correlation (NOT base correlation). For illustration we use compound correlation in the range from 0% through 100%.

Finally, we consider that the covered bonds pay quarterly, have a maturity of 5 years, and match the assets in the coverpool. As stated earlier we look at the level of bonds and do not keep track of cashflow mismatches, durations, accruals, etc.

5.1 Protection Afforded by Covenants

We consider replacement with different levels of issuer default risk because covenants are *not* bankruptcy-remote. Figure 1 shows the effect of different levels of issuer risk to asset risk for varying compound correlation levels. The two plots show the median time to the mth (or last) event (right plot) and the early confidence interval (CI), 10% (left plot). Increasing issuer risk flattens the early part of the curve w.r.t. compound correlation (left plot) and drastically brings in the median of the distribution (right plot). That is, for high levels of issuer risk this even dominates the effect of compound correlation between the assets and the issuer.

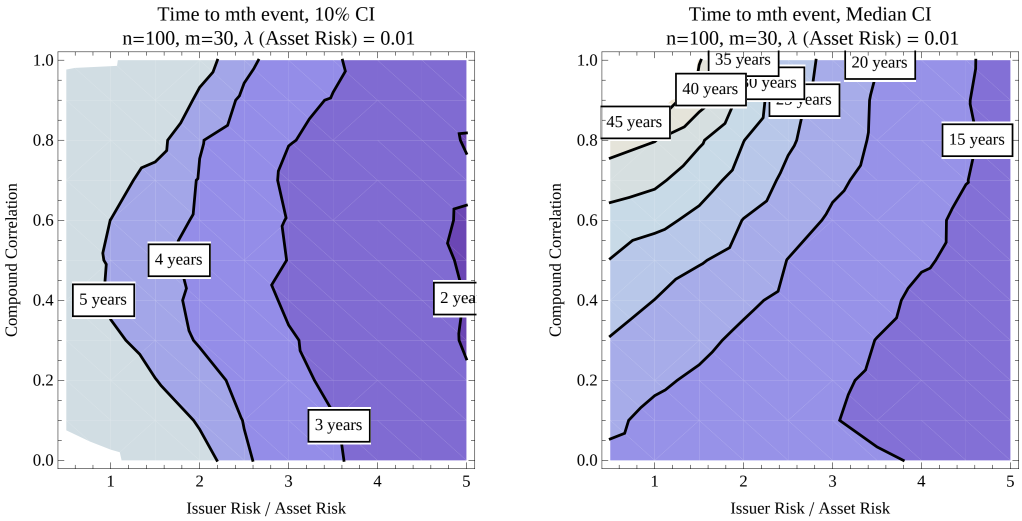

Figure 1: Time to mth event with PSE concentration (i.e. $n = 100$) risk: median on the right and 10% on the left. Analysis by compound correlation and ratio of Issuer hazard rate to Asset (PSE) hazard rate. Note that the contours have different spacings and levels between the two plots. For a 5-year covered bond the ideal situation is where the mth event (i.e. loss of protection) occurs beyond 5 years. See text for details.

5.2 Fair Spreads with Covenants and Over-Collateralization

Figure 2, shows the combined effect of covenants (replacement) and over-collateralization on fair spreads for covered bonds with and without issuer risk.

As expected from the previous results, because the issuer is more risky than the assets in the coverpool, issuer risk dominates available liquidity. I.e. defaults from the coverpool do not usually exhaust issuer liquidity before the issuer defaults and removes the possibility of further replacements.

5.3 Exact vs New-on-last-Refresh

The difference between exact calculation (whether via simulation or analytic methods) is important to see at what point approximation can be introduced into the pricing. Figure 3 shows the difference in basis points. Over the range of scenarios considered, m from 0% of the coverpool to 30%, and compound

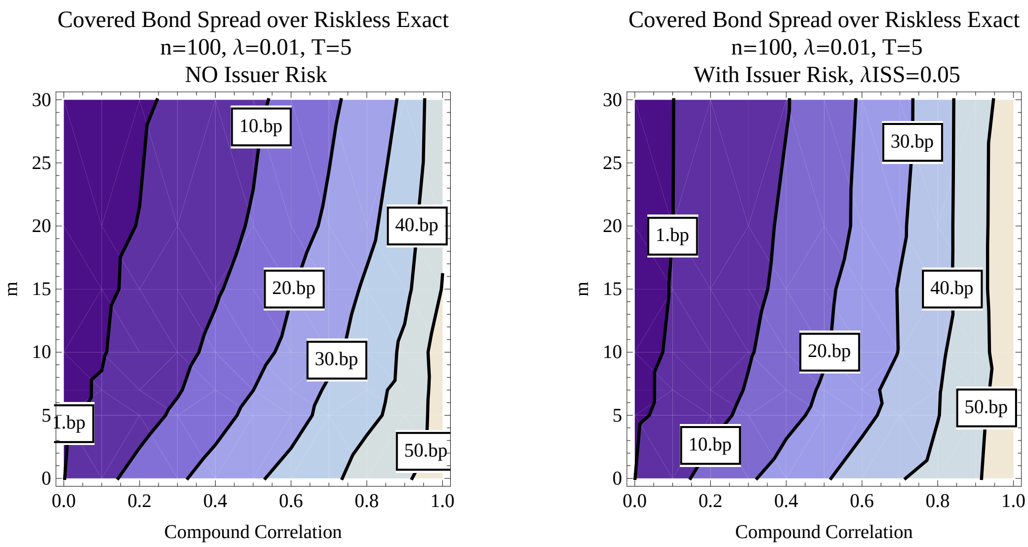

Figure 2: Fair spreads for covered bonds over riskless without (left) and with (right) issuer risk. $m/100$ is the fraction of the coverpool that the issuer can replace given available liquidity. Over-collateralization is 5% but because of the concentration risk the first sovereign default to affect the CDO goes to 10% loss immediately (for a 5% effect). Contour lines (except the lowest at 1bp) are at 10bp intervals.

correlation from 0 to 100%, the difference between exact calculation and the New-on-last-Refresh approximation is bounded by ±1bp for compound correlations below 60%..

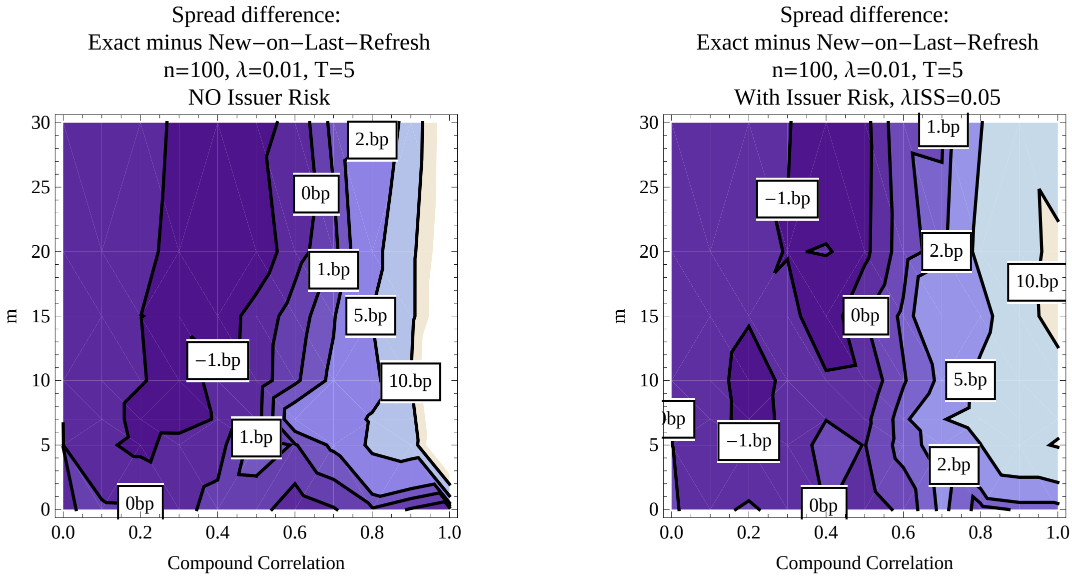

Figure 3: Difference in fair spreads for exact minus new-on-last refresh approximation. $m/100$ is the fraction of the coverpool that the issuer can replace given available liquidity. Over-collateralization is 5%. Contours are in basis points. The difference is bounded by ±1bp for compound correlations below 60%.

6 Discussion and Conclusions

We have introduced the Triggered Refreshed CDO with Issuer Risk as a model for covered bonds and provided analytic pricing methods. The analytic pricing was based both on direct calculation of the underlying distribution and an alternate recursive characterization. Since both of these methods are numerically intensive, involving integration over high dimensions, we propose an accurate analytic approximation: ”new-on-last refresh”. We also provide exact and approximate simulation algorithms. The approximate algorithms allow addition of coverpool features to standard CDO algorithms with no change to their internals. To the authors’ knowledge these are the first pricing algorithms for covered bonds in the literature.

Although the assets in the coverpool are bankruptcy-remote from the issuer the *covenants* on the coverpool are not. This is significant when the coverpool does require replacement and either this drives the issuer into bankruptcy, or when the issuer independently becomes bankrupt and is unable to carry out the replacement. Thus the creditworthiness and the liquidity of the issuer are significant for the value of the notes on the coverpool. These are key features of our pricing model. The figures in the Results section make clear this interplay. Note that if the assets were worse than the issuer then the figures would look quite different.

Covered bonds backed by public sector bonds have a variety of concentration risks that affect problem scaling. Considering sovereign risks only, a coverpool may only have an effective n (number of generators) of 5 to 10 and an available m (available replacements) of up to 3 considering the EEA together with USA/Canada/Japan/Switzerland and ignoring very small countries (such as Malta). At the level of local authorities, cites, etc, the effective n and m may be up to a couple of orders of magnitude larger. We considered this level choosing $n = 100$ and a range of m up to 30.

Adding a hierarchical default/replacment structure to mirror sovereigns and PSEs might form a useful extension but this is beyond the scope of this paper. However, this hierarchy would probably be better dealt with using exactly the current framework, i.e. using different factors to cover different levels. This is because default does not typically have a hierarchical structure, e.g. cities may have higher ratings and smaller spreads that the country they are in.

Covered bonds backed by mortgage assets, as stated at the outset, are outside of the scope. The models introduced here are directly applicable but would require different modeling of the underlying assets. Mortgages are significantly different from bonds and could benefit from inclusion of stochastic interest rates given the interaction with prepayment risk.

We accept that our setting is somewhat stylized. Arbitrary additional details can be added using our simulation approach without changing the basic structure of the problem.

Technically the main complication with respect to standard CDO factor models is that replacement introduces dependence between originally independent conditional distributions. This is exacerbated by the fact that conditional distributions are needed not just conditional default probabilities. Furthermore, the conditional distributions, even if marginally Exponential (i.e. simplest case), can be very different. The authors found (not reported here) that typical conditional distributions for the Exponential in the Gaussian 1-factor copula case are very similar to Gamma distributions.

Even for homogeneous assets the conditional joint distribution after replacement will not be the product of independent distributions, although it will be symmetric w.r.t. asset exchange and potentially reducible to a very simple form (following B. de Finetti [Fel70]). We leave this for future study.

The technical complications result in exact pricing being dependent on a full joint n-dimensional probability distribution. We provided analytic methods to calculate this distribution, both directly, and by a recursive characterization. However, an approximation, which we term new-on-last-refresh, greatly simplifies calculation and, in the examples considered, has an error bounded by ±1bp The difference is bounded by ±1bp for compound correlations below 60%. The new-on-last-refresh approximation also allow decoupling between the additional element of covered bonds, i.e. replacements, and the previous CDO coupon calculation. Alternatively, we provide exact and new-on-last-refresh simulation methods that are direct. Simulation also enables additional details to be considered that are beyond the scope of effective analytic methods.

In summary, we have introduced the Triggered Refreshed CDO with Issuer Risk as a model for covered bonds and provided analytic, effective approximate, and simulation pricing methods.

Appendix I: Triggered Stochastically Refreshed CDO

We can include assets outside the coverpool within the modeling framework. We present this in the new-on-last-refresh approximation. We take the view that these are available as free replacements for coverpool assets. However, they can also generate defaults are costly to the manager in terms of capital. When the manager's capital goes below a legally defined level the manager loses its banking license and the coverpool is then transferred to an administrator. We term this model a Triggered Stochastically Refreshed CDO (TSR-CDO) because the replacement capacity is stochastically dependent on defaults outside the coverpool.

The covenant protection goes through two phases, taking our usual new-on-last-refresh setting:

substitution phase In the first phase the "spare" assets outside the coverpool are substituted in as assets in the coverpool are downgraded. Only assets outside the coverpool can default because as soon as an asset inside the coverpool is downgraded it is substituted. Spares are available until either all of them have been used, or until default amongst the spares has exhausted available capital (represented by m).

replacement phase If all the assets outside the coverpool are substituted in, then, if there is capital remaining (represented by m', i.e. m reduced by the cost of any defaults) as assets in the coverpool default they are replaced. When m' is exhausted the coverpool is just a conventional CDO with the attachment point defined by the over-collateralization.

The previous non-stochastic Refreshed CDOs have only the replacement phase, but their m is not reduced by any defaults outside the coverpool as this is not part of the model. Note that we assume that after the substitution phase downgraded assets in the coverpool are not replaced until they default. Thus we only, additionally, need the probability distribution function of the substitution phase to price this model.

Analytic pricing depends on the following lemma, for notational simplicity we assume a homogeneous coverpool and that defaultable assets outside the coverpool are also homogeneous with the coverpool assets.

Lemma 9. *The conditional probability distribution function of the last substitution time t given l defaults, $LS_{|l}(\vec{f}, \vec{g}, n, k, t)$, where $l \leqslant \min(m, k)$ is:*

$$\begin{aligned} LS_{|l}(\vec{f}, \vec{g}, n, k, m, t) \;=\; & \binom{k+n}{n \quad l \quad (k-l)} (1 - G(t))^n \\ & \Big(F^l(t) \; (k - (l+1)) g(t) G^{k-(l+1)}(t) \\ & \quad + l f(t) F^{l-1}(t) \; G^{k-l}(t) \Big) \end{aligned}$$

where $f(t)$ is the pdf of default and $g(t)$ is the pdf of downgrade. Cumulative distribution functions are given by their respective capital letters.

Proof. The last substitution occurs when there are no more substitutes available, i.e. after k downgrade events (which may be within the coverpool or within the assets that are initially outside the coverpool). If a default occurs a downgrade must (in our model) have already occurred, thus we can condition on the number of defaults. When m defaults have occurred no more substitutions occur because the capital is exhausted and the manager defaults. Thus the number of defaults l is limited by the minimum of m and k.

We have assumed homogeneity, hence the multinomial coefficient on the ways to choose n non-downgraded assets, l defaulted assets, and $k-l$ downgraded assets.

We have n non-downgrades, l defaults, and $k-l$ downgrades occurring before t. The last event may be a default or a downgrade, hence the two components for downgrade and default. □

To price TSR-CDOs we integrate over the conditional possible last substitution times and sum over the conditioning factor (number of defaults) which specifies $m' = m - ql$ where q is the cost of default relative to the cost of replacement. That is, we assume that up to when there are no more substitutes, any defaulting asset simply causes a loss but does not result in a replacement asset. We also assume that the cost of default is different to the cost of a replacement asset for the coverpool after default (because at that point a new asset is bought). There are numerous accounting issues that may interact with this analysis, users should check for their particular situation.

Appendix II: Proofs

Proofs of Lemmas in the main text.

Lemma 1

Proof. We must consider every possible way to have exactly m events from n generators. These are precisely the set of integer compositions of m into n parts $Comp(m,n)$. Note that part sizes range between zero and m inclusive and that compositions are ordered. We call each composition $\alpha \in Comp(m,n)$ a case, e.g. from $m = 5, n = 3$ one case is $(3,2,0)$. We call the numbers in the case (i.e. $3,2,0$ in the example) the entries of the case.

We are interested in the probability distribution of the time of the mth event, i.e. what is the probability that this equals t? Thus one event is exactly at time t. Each case will give a contribution from each of the n entries that is non-zero because each non-zero entry could contain the mth event. So the example $(3,2,0)$ would generate two contributions.

The time of the last event in each entry in the case are independent because each of the generators are independent by assumption. Each of the events generated by each generator is also independent by assumption.

Thus for the generator i that generates the mth event, say the $\alpha[i]$th event from that generator, its pdf is just the $\alpha[i]$th convolution of f_i, i.e. $f_{i*\alpha[i]}(t)$.

If generator i gives the mth event, all of the other generator events occur before t. Furthermore, because we state that exactly m events occur before t, we know that each subsequent event for the other entries must occur after t.

- If the entry $\alpha[j]$ for the event generator is zero, that means that all the events for this generator occur after t, so its first event occurs after t, i.e. with probability $1 - F_j(t)$.

- If the entry $\alpha[j]$ is non zero, then the we know the $\alpha[j]$th event occurs before t and the next one afterwards, i.e. with probability:

$$\begin{aligned} P[\alpha[j] \leqslant t \leqslant \alpha[j]+1] &= \int_{s=0}^{s=t} f_{*\alpha[j]}(s)(1 - F_j(t-s))ds \\ &= F_{j*\alpha[j]}(t) - \int_{s=0}^{s=t} f_{*\alpha[j]}(s)F_j(t-s)\ ds \end{aligned}$$

  because the distribution of the $\alpha[j]$th event is the $\alpha[j]$th convolution of the pdf because the events are independent (by assumption).

Thus we arrive at the formula given. □

Lemma 2.

This is basically a small rearrangement of the previous Lemma and the proof follow directly.

Lemma 3. No proof, this is a restatement of known formulae.

Lemma 4

Proof. If the kth event was at u then the probability is just the cumulative distribution for the time remaining, $t - u$.

If $k = 0$ then we want the conditional probability of an event in (u, t) given no events in $(0, u)$, this is direct by Bayes Theorem, $P(A|B) = P(A)/P(A \cap B)$, where if an event happened by u it certainly happened by $t > u$.

If $k \geqslant 1$ we repeat the calculation for $k = 0$ conditioning on when the kth event happened, and then integrate. That is, we calculate $P(A) = \int P(A|B)P(B)dB$. The distribution of the time of the kth event is the k-fold convolution of one event because events are independent by assumption. □

Lemma 5

Proof. CDO depends only on the vector of default probabilities of its underlyings $\vec{p}$ provided that the default events are independent (which they are by construction).

If less than m replacements have occurred the default probabilities are all zero (if there was a default then a replacement would be available).

If m replacements have occurred then the probabilities of default are non-zero. The probability that an underlying defaults depends on its last replacement time. The last replacement time depends on how many replacements have occurred for that underlying. $Comp(m, n)$ contains all the possible combinations of m underlying replacements on n underlyings.

The probability that a particular combination $\alpha \in Comp(m, n)$ will occur depends on which underlying was the last to be replaced and when that occurred. This is described by ξ. Since we do not know when the last replacement occurred or which underlying was replaced we integrate over the possible times $u = 0$ to $u = t$, and consider the contribution from each underlying being the last one to be replaced at time u. If there are n underlyings this means that we have an integration over n dimensions.

Given a setup, ξ, if the event probabilities of each underlying are independent then the probabilities of the next event for each underlying are also independent. Hence the joint distribution of $P[\xi]$ is just the product of the individual distributions from earlier Lemmas.

The pdf of the next event time for a particular underlying j, given a setup, ξ, depends firstly on whether or not $i = j$. If $i = j$ we know when the previous event occurred, u, so the probability of default of the replacement is given by the cdf. The rest follows exactly from Lemma 4. □

Lemma 6

Proof. Whenever a default time is less than any other it is replaced by the sum of itself and a replacement. This constitutes convolution with the respective replacement distribution. If a point is never a minimum it is the same at the next step. □

Lemma 7

Proof. We have made the approximation that the probability that a coverpool asset has defaulted, is given by its conditional survival probability from the last refresh up to t_i. This is an approximation in that, whilst the unconditional time-to-default pdfs may (if Exponential) be memoryless, the distributions conditioned on the conditioning factor c, are not (as shown by Lemma 3 above). With this approximation, the survival probability from the last refresh up to t_i depends only on the time of the mth refresh and t_i. For constant hazard rates this is just a function of the difference between these times. □

Lemma 8

Proof. It is elementary that P(either) = P(one) + P(other) − P(both). Events from the first n generators and the $n+1$st generator are independent by assumption so P(both) = P(one) × P(other). □

References

[ECB08] ECBC. *European Covered Bond Fact Book 2008.* ECBC, Brussels, 2008.

[Fel70] W. Feller. *An introduction to probability theory and its applications. Second Edition, Ch VII.4.* Wiley, New York, 1970.

[Hof09] K.F. Hoffelder. Covered bond legislation seen necessary to jump-start alternative program to TLGP. *Financial Times*, May 7th, 2009.

[HW04] J.C. Hull and A. White. Valuation of a CDO and an n-th to default CDS without Monte Carlo simulation. *The Journal of Derivatives*, Winter, 2004.

[LG05] J-P. Laurent and J. Gregory. Basket Default Swaps, CDO's and Factor Copulas. *J. Risk*, 7(4):103–122, 2005.

[Opd08] J.D. Opdyke. A unified approach to algorithms generating unrestricted and restricted integer compositions and integer partitions. *SSRN eLibrary*, 2008.

[Sch05] Schonbucher. *Credit derivatives pricing methods.* Wiley, London, 2005.