XAIA
INVESTMENT

Z-SPREADS FOR BONDS WITH OPTIONAL SINKING FEATURE: A BELLMAN EXERCISE

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

Date: May 2, 2013

Abstract It is explained how to compute a Z-spread for bonds with optional sinking feature. Such instruments equip their issuer with an option (but not an obligation) to redeem parts of the nominal before maturity; therefore the future cash flows generated by the bond are random. The proposed method coincides with the so-called “worst-ansatz“ in the special case of a callable bond. In the general case it relies on a dynamic programming technique based on the Bellman principle.

1 What is a Z-spread?

On fixed income markets investors as well as risk managers often face the task of assessing a return measure to a bond. On the one hand, an investor might base her investment decision on this number, together with a measure for the risk associated with the instrument. On the other hand, return measures for bonds are commonly defined as “risk factors“ in a financial institution's risk management process.

On the very first, but admittedly naive, glimpse this task is simple. One would expect that a straight coupon bond generates an annualized return which coincides precisely with its coupon rate. However, this is of course only true when the bond trades at par. For instance, in the case the bond trades below par it generates a return additional to the coupon, which is due to the so-called *pull-to-par effect*, i.e. the fact that the bond's full nominal is redeemed at maturity. Conversely, if the bond trades above par the bond's annualized return measure should be smaller than its coupon rate. Therefore, one has to implement methods that “annualize“ the pull-to-par effect.

One of the most common return measurements applied in the markets is the computation of a Z-spread. The fundamental idea of this method relies on the assumption that a bond carries a certain risk of default, i.e. the possibility that the bond issuer is not able to pay back the borrowed money at maturity. In order to compensate the bond holder for this risk the bond is expected to generate a return higher than the “risk-free“ interest rate. By risk-free rate we mean the interest rate which can be earned without exposure to default risk. Intuitively, the higher the default risk, the lower is the bond's market price and the higher its expected return. Consequently, the Z-spread is defined as a spread on top of the risk-free interest rate which is chosen such that the resulting new interest rate can explain the bond's market price when used for discounting the bond's cash flows. Mathematically, we assume the bond has unit notional and denote its coupon dates by $0 =: t_0 < t_1 < \ldots < t_m$, where t_m is the redemption date, which coincides with the last coupon date. Further assuming continuous compounding and denoting the risk-free, deterministic short rate curve by $\{r_t\}_{t\geq 0}$, the risk-free discount factor for time t is given by $DF(t) := \exp\Big(-\int_0^t r_s\,ds\Big)$. Consequently, if the bond was assumed to be free of credit risk its market price should be given by $B_0 := \sum_{i=1}^m C_i\,DF(t_i)$, where C_i denotes the actual cash flow which the bond generates at time t_i. For instance, if the bond is a plain vanilla coupon bond with coupon rate C, then[1] $C_i = C\,(t_i - t_{i-1})$ for $i = 1, \ldots, m-1$ and $C_m = C\,(t_m - t_{m-1}) + 1$. However, the observed market price B_0 of the bond typically deviates from its discounted cash flows. With given market price B_0, the bond's Z-spread is defined as follows.

Definition 1.1 (Z-Spread)
The Z-Spread z is the unique root of the equation

$$B_0 = \sum_{i=1}^m C_i \exp\Big(-\int_0^{t_i}(r_s + z)\,ds\Big).$$

In words, the Z-spread is a parallel shift of the risk-free interest rate curve $\{r_t\}_{t\geq 0}$ such that the bond's market price matches the sum of its cash flows when discounted with the shifted interest rate curve. The function on the right-hand side of the defining equation for the Z-spread is strictly decreasing in z, so that the root is unique. The root is typically positive, which means that the bond price is smaller than its default-free discounted cash flows. Otherwise, a negative Z-spread means that the bond is "less risky than risk-free", raising doubts about the initial pick of the risk-free short rate curve $\{r_t\}_{t\geq 0}$. Henceforth, we therefore restrict our analysis to bonds with non-negative Z-spreads.

2 What is an optional sinking feature?

Unlike a straight bond, a *sinkable* bond does not redeem the full nominal at once at maturity but instead pays back the nominal in several steps. For instance, there might be three equal installments of redemption payments within the last three years of a 10-year bond's lifetime. Such a sinkable feature does not make the pricing, and hence the Z-spread computation, for the bond more complicated than for straight bonds. Using the notation from the previous paragraph, it simply requires the cash flow amounts C_i to be computed differently. For instance, if the nominal of a 10-year bond sinks in three equal installments within the last three years before maturity, then $C_i = C\,(t_i - t_{i-1})$ for the first seven cash flows $i = 1, \ldots, 7$, and

$$C_8 = C\,(t_8 - t_7) + \frac{1}{3}, \quad C_9 = C\,(t_9 - t_8)\,\frac{2}{3} + \frac{1}{3},$$
$$C_{10} = C\,(t_{10} - t_9)\,\frac{1}{3} + \frac{1}{3}.$$

However, some bonds traded in the market provide its issuer with an option to redeem parts of the nominal at pre-specified time

[1] In practice, to compute the actual coupon payment amount at time t_i the coupon rate has to be multiplied with the year fraction between the time points t_{i-1} and t_i, according to some contract-specific daycount convention. For the sake of notational simplicity, we already interpret the time points t_i as year fractions with respect to the appropriate daycount convention.

points before maturity. The difference to a regular, mandatory sinking feature is that the issuer might opt for an early redemption of some of her nominal but need not do so. Also the sinking amount might be optional to the issuer. Therefore, the cash flows generated by the bond are not determined yet and depend on the issuer's decisions in the future[2]. Let us provide a tiny example for such a structure.

Example 2.1 (Bond with optional sinking feature)
Consider a 2-year bond with annual coupon rate C and unit nominal. The issuer is assumed to have the option to redeem half of the bond's nominal after one year but need not do so. Consequently, there are two different cash flow szenarios that might be generated by this bond: Either the issuer does not make use of her option, in which case the bond holder receives C after one year and $C+1$ at maturity. Or the issuer makes use of her option, in which case the bond holder receives $C+0.5$ after one year and $C \cdot 0.5 + 0.5$ at maturity.

Clearly, an extension of the Z-spread definition to also include bonds with an optional sinking feature is no longer trivial at all. In Example 2.1, should we compute two Z-spreads and include probability weightings to the two szenarios? How would we choose the respective probabilities? In the next paragraph we aim at answering precisely these questions. It provides a very basic, but therefore simple and intuitive method which truly extends the definition of a Z-spread to include optional sinkable features, i.e. the proposed method coincides with the regular Z-spread in case there is no optional sinking feature. But before we get there let us make the final remark that *callable bonds* arise as the special case when the sinkable option of the issuer is such that the full nominal has to be redeemed at once, only the timing is optional. We will see that the proposed algorithm for the Z-spread computation of a bond with optional sinking feature in case of a callable bond incidentally boils down to what is called the "worst-ansatz" on Bloomberg, so it might be considered a proper extension thereof.

Remark 2.2 (Decomposition into callable bonds)
The bond from Example 2.1 can be decomposed into a straight bond with maturity 2 years and notional 0.5, and a 2-year callable bond with notional 0.5. This implies that the pricing for this bond is equivalent to the pricing of two simpler instruments, and therefore the dynamic programming approach presented in the present article is not needed but can be resolved by applying the classical worst-ansatz to both parts. However, such a decomposition into callable bonds is not always possible, or at least not always obvious to find in general. Furthermore, within one's booking software system it might not be desired to treat one bond (with optional sinking feature) as a whole portfolio of numerous (callable) bonds.

[2] For example, Westvaco Corporation, a US packaging company, has issued a $150\$$ billion bond with optional sinking feature in March 1997, whose ultimate maturity is June 2027. On an annual basis the issuing company is allowed to redeem either 5% or 10% of the outstanding nominal.

3 Computation of a Z-spread using the Bellman principle

So far, the definition of a Z-spread involves no probability theory at all. In order to account for the uncertainty of cash flows due to the issuer's sinkable option we have to introduce some stochastic evolution of the bond price. Introducing only the most minimal amount of probability theory, we can reformulate the definition of a Z-spread such that it allows for a stochastic interpretation. To this end, we consider an exponential random variable τ on a probability space $(\Omega, \mathcal{F}, \mathbb{P})$, which we interpret as the issuer's default time, i.e. the future time point at which the issuer becomes bankrupt and cannot repay her debt. We assume that the exponential decay parameter of this random variable equals the bond's Z-spread z. In this case, the present value of the bond's cash flows is a random variable given by $\sum_{i=1}^{m} C_i \, DF(t_i) \, 1_{\{\tau > t_i\}}$ and the arbitrage-free market price of the bond in this model is the expectation of this random variable. Since $\mathbb{P}(\tau > t_i) = \exp(-z \, t_i)$, the definition of the Z-spread implies that this tiny model for the issuer's default time, henceforth called *simple default model*, (a) can explain the bond's observed market price and (b) provides a probabilistic interpretation for the Z-spread as an exponential rate parameter, also called *default intensity*. Consequently, the Z-spread can alternatively be defined as follows.

Definition 3.1 (Z-Spread)
The Z-Spread z is the unique default intensity of the exponential default time τ such that the simple default model price of the bond matches the bond's observed market price.

This stochastic interpretation of the Z-spread is already explained in Pedersen (2006). It is quite intuitive because the default intensity z can be considered as a credit spread which the issuer has to pay on top of the risk-free interest rate in order to compensate the bond holder for the intrinsic default risk. For our purpose, the simple default model is the most minimal stochastic modeling approach in the following sense: On the one hand, it allows us to derive an arbitrage-consistent[3] price for the bond with optional sinking feature, which is a fundamental requirement. On the other hand the introduced stochastic objects (which is only the exponential default time τ) are so minimal that they do not have an influence on the decision process of the issuer. The only information flow we consider is the evolution of the indicator process $1_{\{\tau > t\}}$ over time, which provides the issuer with no relevant information regarding her optimal sinking schedule. Mathematically speaking, this is because the so-called lack-of-memory property of the exponential distribution implies that the default indicator process $\{1_{\{\tau > t\}}\}_{t \geq 0}$ is Markovian. In other words, the issuer only observes whether she is already bankrupt or not, but her expectations regarding this default event do not fluctuate over time. All possible time-dependent information is modeled deterministically, namely the evolution of interest rates $\{r_t\}_{t \geq 0}$, as well as any additional information that might influence the creditworthiness of the issuer. The default intensity z is modeled as a constant and does not fluctuate over time, so the issuer knows how creditworthy she is and that this creditworthiness remains constant

[3] We mean arbitrage-consistent with respect to other bonds issued by the same company and/or possibly also other credit derivatives.

during the lifetime of the bond. Consequently, the optimal redemption schedule, which the issuer is going to choose within the simple default model, is already known now at time $t = 0$. However, the optimal schedule does depend on the default intensity. Let us provide a tiny example to illustrate this fact.

Example 3.2 (Bond with optional sinkable feature, cont.) Consider the 2-year bond with annual coupon rate $C = 0.04$ and unit nominal from Example 2.1. Furthermore, assume a flat interest rate $r_t \equiv 0.01$. Temporarily denote the default intensity in the simple default model by the unknown x, because we do not know the Z-spread yet. In case the issuer makes use of her early redemption option, the bond price is given by

$$B_0^a(x) := (0.04 + 0.5)\,\exp(-(0.01 + x)) \\ + (0.04 \cdot 0.5 + 0.5)\,\exp(-(0.01 + x)\,2).$$

In case she doesn't opt for early redemption it is given by

$$B_0^b(x) := 0.04\,\exp(-(0.01 + x)) \\ + (0.04 + 1)\,\exp(-(0.01 + x)\,2).$$

Within the tiny setup of the simple default model the issuer is going to choose the schedule which minimizes her expected payments, so that the bond's model price equals $\min\{B_0^a(x), B_0^b(x)\}$. As a minimum of two decreasing functions, this is a decreasing function in the default intensity x, which is visualized in Figure 1. For small x, i.e. high creditworthiness, the issuer decides to make use of her option and redeems early, whereas for large x, i.e. low creditworthiness, the issuer decides to redeem the full notional after two years. This is intuitive, since an issuer with low creditworthiness rather defers her due payments to later time points, hoping to get around them in case of a default event. The Z-spread according to Definition 3.1 is now determined as the unique root z of the equation $B_0 = \min\{B_0^a(z), B_0^b(z)\}$, where B_0 denotes the bond's observed market price.

How do we compute the price of a bond with optional sinkable feature in the general case? To this end, we make the following two simplifying discretization assumptions:

(a) There are only finitely many time points before maturity at which the issuer may possibly redeem parts of the bond nominal. Only for the sake of notational convenience, we additionally assume that these time points coincide with the coupon payment dates t_i of the bond. The latter assumption can clearly be relaxed without further theory by including more coupon dates with coupon payment amounts equal to zero.

(b) The unit bond nominal is partitioned into K equal parts of $1/K$ and at each possible redemption payment date t_i the issuer is forced to redeem a certain number of parts. The amount the issuer is allowed to redeem at time t_i is dependent on the outstanding notional $s \in \{0, 1/K, 2/K, \dots, 1\}$ and must be chosen from a pre-defined subset $D_i(s)$ of the set $\{0, 1/K, 2/K, \dots, s\}$, where $D_m(s) = \{s\}$ (implying that the remaining outstanding nominal has to be redeemed at time t_m).

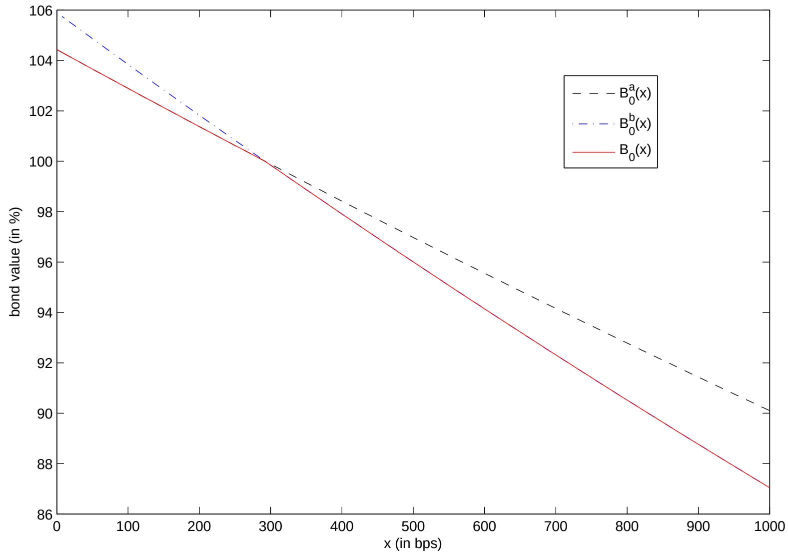

Fig. 1: Illustration of the model price for the bond described in Example 3.2, in dependence on the default intensity x in the simple default model.

For example, for the bond in Example 2.1 we have $m = 2$, and $D_1(1) = \{0, 1/2\}$. On the one hand, these discretization assumptions are not severe for practical purposes because both the number of parts K as well as the number of time points m can be made arbitrarily large, so that also continuous sinking intervals and amounts may be approximated arbitrarily close. When increasing the number of redemption (respectively coupon) time points m, one has to introduce additional cash flows C_i which are zero, i.e. $C_i = 0$, at redemption time points which are not coupon time points.
On the other hand, these discretization assumptions imply that there are only finitely many possible options of redemption schedules for the issuer to choose from. Like in Example 3.2, for each possible redemption schedule we obtain a possible bond price as a decreasing function in the default intensity, and the sinkable bond price is the minimum over all these functions, which itself is a decreasing function. In order to compute the Z-spread we have to find the default intensity z such that this function matches the bond's observed market quote B_0. To accomplish this root search a bisection routine is very efficient because of the monotonicity of the bond price function. However, the bijection routine itself requires multiple evaluations of the sinkable bond for different default intensities. Under our discretization assumptions each evaluation requires finding the minimum over up to K^m possible redemption schedules, which can be a very tedious exercise. Clearly, for large values of K and/or m this computational effort is unacceptable. Luckily, there is a way to solve this minimization problem more efficiently with computational complexity of the order $\mathcal{O}(K^2\, m)$, which is a substantial improvement and allows to compute the Z-spread within fractions of a second. The problem can be formulated in terms of a dynamic programming exercise and is solved by means of a backwardation technique

based on the Bellman principle. In particular, the involved algorithm for finding the optimal strategy is quite a compact code.
The formulation of the optimization problem as a deterministic dynamic program is as follows:

- At each time point t_i, $i = 1, \ldots, m$, the remaining notional takes a value in the state space $S := \{0, 1/K, \ldots, (K-1)/K, 1\}$.
- At each time point t_i, $i = 1, \ldots, m$, the issuer may redeem a certain amount $a_i \in A := S$ of the remaining nominal. A is typically called the action space.
- The set of admissible actions (i.e. redemptions) at time t_i with remaining nominal $s_i \in S$ is denoted by $D_i(s_i)$ and given as an arbitrary subset

$$D_i(s_i) \subset \{0, 1/K, 2/K, \ldots, s_i\}, \quad D_m(s_m) = \{s_m\}.$$

- The transition function from time t_i to time t_{i+1}, when the remaining nominal at t_i is $s_i \in S$ and the action taken is $a_i \in D_i(s_i)$, is given by $T_i(s_i, a_i) = s_i - a_i$.
- The so-called one-step cost functional gives the discounted cash flow the issuer has to pay at time t_i, when the remaining nominal is $s_i \in S$ and the action taken is $a_i \in D_i(s_i)$. It is given by $r_i(s_i, a_i) = DF(t_i)\,(C_i\, s_i + a_i)$.

Any admissible sequence of actions $(a_1, \ldots, a_m)$ gives a possible redemption schedule and induces a sequence of remaining nominals $(s_1, \ldots, s_m)$ with $s_1 = 1$ and $s_{i+1} = T_i(s_i, a_i)$, $i = 2, \ldots, m$. Given this terminology, the issuer's optimization problem can be formulated as follows: find an optimal sequence $(a_1, \ldots, a_m)$ of actions with $a_i \in D_i(s_i)$ for all $i = 1, \ldots, m$ such that the cost functional $\sum_{i=1}^{m} r_i(s_i, a_i)$, which is precisely the bond price under the redemption schedule $(a_1, \ldots, a_m)$, is minimized.
The benefit from reformulating the cost minimization problem as a dynamic program is that it can now be solved sequentially, i.e. by iteratively solving the m simpler optimization problems of minimizing the one-step cost functionals $r_1, \ldots, r_m$. This dynamic programming algorithm relies on the so-called Bellman principle, named after Bellman (1957). Denoting the optimal redemption schedule by $(a_1^*, \ldots, a_m^*)$, it basically states that the truncated action sequence $(a_k^*, \ldots, a_m^*)$ is an optimal redemption schedule for the truncated dynamic program restricted to the time points $t_k, \ldots, t_m$, for all $k = 1, \ldots, m$. This implies that we can solve the optimization problem by the following generic algorithm, inductively working backwards in time from t_m to t_1:

- Time t_m: before the final redemption is made, the remaining nominal s_m can have any value in S and the final redemption a_m^* is forced to equal s_m. This means that for all possible states s_m we obtain a minimizer $a_m^* = a_m^*(s_m) = s_m$ and a minimal one-step cash flow value of $r_m(s_m, a_m^*) =: V_m(s_m)$. We store both $a_m^*(s_m)$ and $V_m(s_m)$ for all possible states s_m.

- Time t_{m-1}: the remaining nominal s_{m-1} can have any value in S. For each of these values, we have to find the optimal redemption amount a^*_{m-1}, which depends on s_{m-1}. Using the Bellman principle, $a^*_{m-1} = a^*_{m-1}(s_{m-1})$ minimizes the two-step cash flow value

  $$a \mapsto r_{m-1}(s_{m-1}, a) + V_m\big(T_{m-1}(s_{m-1}, a)\big),\ a \in D_{m-1}(s_{m-1}).$$

  Denote the minimal value by

  $$\begin{aligned} V_{m-1}(s_{m-1}) := r_{m-1}(s_{m-1}, a^*_{m-1}(s_{m-1})) \\ + V_m\big(T_{m-1}(s_{m-1}, a^*_{m-1}(s_{m-1}))\big). \end{aligned}$$

  It is important to note that the values $V_m\big(T_{m-1}(s_{m-1}, a)\big)$ have already been computed in the first step for all required actions a. We store both $a^*_{m-1}(s_{m-1})$ and $V_{m-1}(s_{m-1})$ for all possible states s_{m-1}.

- Arbitrary time t_k, $k = 1, \dots, m-2$: the remaining nominal s_k can have any value in S. For each of these values, we have to find the optimal redemption amount a^*_k, which depends on s_k. Using the Bellman principle, $a^*_k = a^*_k(s_k)$ minimizes the $(m-k+1)$-step cash flow value

  $$a \mapsto r_k(s_k, a) + V_{k+1}\big(T_k(s_k, a)\big), \quad a \in D_k(s_k).$$

  Denote the minimal value by

  $$V_k(s_k) := r_k(s_k, a^*_k(s_k)) + V_{k+1}\big(T_k(s_k, a^*_k(s_k))\big).$$

  It is important to note that the values $V_{k+1}\big(T_k(s_k, a)\big)$ have already been computed in the previous step for all required actions a. We store both $a^*_k(s_k)$ and $V_k(s_k)$ for all possible states s_k.

- Final step: the value $V_1(s_1)$ for $s_1 = 1$ is the value of the sinkable bond, and the optimal redemption schedule is computed iteratively by

  $$\begin{aligned} & a^*_1 = a^*_1(1), \\ s_2 = T_1(1, a^*_1), \quad & a^*_2 = a^*_2(T_2(s_2, a^*_2(s_2))), \\ s_3 = T_2(s_2, a^*_2), \quad & a^*_3 = a^*_3(T_3(s_3, a^*_3(s_3))), \\ & \vdots \\ s_m = T_{m-1}(s_{m-1}, a^*_{m-1}), \quad & a^*_m = a^*_m(T_m(s_m, a^*_m(s_m))) = s_m. \end{aligned}$$

When analyzing the runtime of this algorithm, we see that it is much quicker than testing all possibly admissible redemption schedules $(a_1, \dots, a_m)$. In each step $1, \dots, m$, we have to solve at most K simple minimization problems (namely one for each possible state), each of which requires at most K comparisons (namely one for each admissible action). This leaves us with a runtime of at most $\mathcal{O}(K^2\, m)$. In contrast, the brute force approach of comparing all possible redemption strategies has computational effort in $\mathcal{O}(K^m\, m)$ in general, i.e. the number of possible redemption payment dates enters exponentially, which can be dramatic.

Finally, let us remark that the special case of a callable bond is included in the presented setup. In this case $K = 1$ and $D_i(s_i) \subset \{0, s_i\}$ for all i, meaning that the issuer may decide when to redeem the nominal, but has to redeem the full nominal in this case. Analyzing the dynamic programming algorithm above, this implies that the bond value $V_1(s_1)$ equals precisely the minimum over all possible m redemption schedules. This simple approach to evaluate a callable bond is called "worst-ansatz" on Bloomberg. Clearly, in this special case pricing with the presented dynamic programming approach is cracking a nut with a sledgehammer, and one would rather resort to brute force minimization. In the general case, as outlined above, the dynamic programming approach is a lot quicker.

References

R. Bellman, Dynamic programming, *Princeton University Press, Princeton, New Jersey* (1957).

C.M. Pedersen, Explaining the Lehman Brothers option adjusted spread of a corporate bond, *Fixed Income Quantitative Credit Research, Lehman Brothers* (2006).