

1 – What are callable bonds?

When you take out a fixed rate mortgage to buy a house, you usually have the option of pre-paying the mortgage. The common term to use is “to refinance”. And people would refinance their mortgages when it is cheap to do so – when interest rates are low.

Callable bonds are very similar except that now companies are the borrowers. They issue callable bonds to borrow money for whatever reason (not necessarily to buy houses). Being callable, such bonds give them the right to “call home” the bonds – prepay their borrowings – when they see fit, which usually means when interest rates are low.

To pay off the bonds, the issuers usually have to pay the holder the face value of the bonds. For many callable bonds, however, the issuers will need to pay some premium on top of the face value. This premium acts as some compensation for the lenders who upon being prepaid, have to find new borrowers at generally lower interest rates. The price that the issuers have to pay is the **call price**.

Since callable bonds are attractive to borrowers, they are disliked by lenders. Although lenders get compensated through higher coupon rates with callable bonds, to tone down call risks with callable bonds, many issuers introduce a call protection period during which a callable bond cannot be called. A typical callable bond structure will look like 10 NC 5: which means the bond has 10 years till maturity and only callable after year 5.

2 – Why are callable bonds?

It is obvious that callable bonds give borrowers the option to refinance when interest rates are low. In other words, it is one way companies hedge against possible decreases in future interest rates. For this reason, callable bonds are very popular before 1990. In fact, before 1970 almost all corporate bonds were issued with call features. Between 1970 and 1990, about 80% of fixed rate corporate bonds were callable. Due to the development of the interest rate derivatives markets in the late eighties and early nineties, there has been a big drop in callable bonds issuance – now accounting for only 30% of the total. This is understandable since with derivatives, it becomes ever easier to hedge against interest rate risks. With callable bonds, providers of capital (lenders) also act as insurance providers. This may not be necessarily optimal - the same way a person may not be good at both tennis and finance. Due to cost savings from specialization, companies may find it more cost-effective to borrow by issuing straight bonds and buy insurance against interest rate risks from specialized insurance providers.

However, another reason why firms may still find callable bonds desirable is that by issuing callable bonds they can send a strong positive signal to the markets about the quality of their business. The reasoning goes as follows: If a firm is confident about their business and believes that their credit quality will improve in the future (which will lower their borrowing costs), it makes sense for them to issue a callable bond. As soon as the market realizes their better values, they can simply call the old expensive bond and replace it with a bond which pays lower coupons. On the other hand, if a firm knows that they are not doing particularly well and their credit quality is very likely to deteriorate, it makes sense for them to issue a non-callable bond to lock in a borrowing rate.

3 – Yields to call and yields to worst

Traders like to think in terms of yield to maturity simply because it is seemingly easier to understand. “A bond is trading at a yield of 5%” seems more straightforward as compared to a bond trading at 95.24% of face value. For this reason, markets have come up with yields measures for callable bonds as well. We will talk about these measures in this section.

Strictly speaking, yield to maturity is out of question for callable bonds. The simple reason is that callable bonds don't have fixed maturities. Take for an example, the 10 NC 5 bond (10 year stated maturity, only callable after 5 years). If the issuer, for some reasons, decides to call the bond at end of year 5/beginning of year 6, the maturity of the bond is 5 years. However, it is also possible that the issuer may let the bond live until its usual maturity of 10 years. Without a fixed maturity, we are not certain about the cash flows either and as such a yield to maturity cannot be computed.

However, traders are in love with yields measures and thus they have come up with at least 2 ways of computing yields for callable bonds.

- First, they assume that the bond, though callable, will not be called at all during its entire life. In our 10 NC 5 bond example, this means the bond's maturity will be 10 years. Yield computed with this assumption is still called yield to maturity.
- Second, they assume that the bond will be called with certainty. In our 10 NC 5 bond example, this means that the bond will mature at year 5. Yield computing with this assumption is yield to call. Many callable bonds however have multiple call dates. For example, our 10 NC 5 bond can be called anytime after year 5 until year 10. In this case, we need to be very specific about the call assumption. If we assume that the bond will be called at the end of year 5 with certainty, strictly speaking, the resulting yield will be called **yield to first call**.

To avoid possible confusion, let me give a simple example for us to quickly grasp the concept. To make it simple, let's work with a 2-year bond that can only be called at end of year 1 for a call price of \$100. This bond, currently selling for \$99, has a face value of \$100 and is paying a semi-annual coupon rate of 8% p.a.

Yield to maturity

To compute yield to maturity of this callable bond, we will make the assumption that the bond will be held to maturity regardless. Therefore, the cash flows from the bond will simply be:

- At time 0.5: \$4
- At time 1.0: \$4
- At time 1.5: \$4
- At time 2.0: \$104

The yield to maturity of the bond will then be y such that:

$$99 = \frac{4}{1 + \frac{y}{2}} + \frac{4}{\left(1 + \frac{y}{2}\right)^2} + \frac{4}{\left(1 + \frac{y}{2}\right)^3} + \frac{104}{\left(1 + \frac{y}{2}\right)^4}$$

Solve this for y , we have: $y = 8.55\%$.

Yield to call

To compute yield to call of this callable bond, we will make the assumption that the bond will be called with certainty. Therefore, the cash flows from the bond will be:

- At time 0.5: \$4
- At time 1.0: \$4 + \$100 = \$104

The yield to call of the bond will then be y such that:

$$99 = \frac{4}{1 + \frac{y}{2}} + \frac{104}{\left(1 + \frac{y}{2}\right)^2}$$

Solve this for y , we have: $y = 9.07\%$.

Yield to worst

For a callable bond, yield to worst is simply the minimum between the yield-to-maturity and the yield to call. In the above example, yield to worst is simply minimum of (8.55%, 9.07%) = 8.55%.

A word of caution

Let's assume that we go out to a Bloomberg terminal to check out prices of bonds of comparable credit quality to the callable bond above and find out the following:

- A 2-year non-callable bond is trading at a yield of 8.5% (or a price of \$99.10)
- A 1-year non-callable bond is trading at a yield of 8.4% (or a price of \$99.62)

Comparing the pricing information here to that of the callable bond, it seems really weird. From our calculations,

- the callable bond offers a yield of 8.55% if it is held to maturity. In this case, its cash flows are exactly the same as a 2-year non-callable bond which offers a yield of only 8.5%.
- the callable bond offers a yield of 9.07% if it is called regardless. In this case, its cash flows are exactly the same as a 1-year non-callable bond which offers a yield of only 8.4%.
- in other words, worst comes to worst, the bond earns a yield to worst of 8.55% which is still better than either of the yields offered by the 1-year or 2-year non-callable bond.

It seems that the callable offers higher (than necessary) yields when compared to the non-callable. Put it differently, the callable is selling for \$99 which is cheaper than both of the 1-year and 2-year non-callable. What is going on? Is the market not functioning well? Or are we missing something?

It turns out that **if the market is functioning well**, the callable ought to be cheaper than both the 1-year **as well as** the 2-year non-callable. To see why the callable should be cheaper than the non-callables, let's compare their cash flows:

time	1-year Non-callable	2-year Non-callable	callable
0.5	\$4	\$4	\$4
1.0	\$104	\$4 + Market price at time 1.0	= \$4 + minimum of (\$100, market price at time 1.0 of the 2-year non-callable)

The second column contains cash flows of the 1-year non-callable which is quite straightforward. The third column contains the cash flows of the 2-year non-callable up to time 1.0. The first cash flow is of course the \$4 coupon at time 0.5. To come up with a cash flow for the 2-year non-callable at time 1.0, I assume we collect the coupon of \$4 and sell this bond at time 1.0. The cash flow at time 1 for this bond will be \$4 + its market price at time 1.0.

The last column of the table contains cash flows to the callable. Nothing is special about the first cash flow – simply a coupon of \$4. The cash flow at time 1.0, however, is crucial since this is where the bond issuer can exercise their call right. Let's see how the issuer of the callable makes his/her decision at time 1.0. It turns out to be quite simple. Since the call gives the issuer the right to buy back the 2-year non-callable bond at time 1.0 for a price of \$100, it only makes sense for the issuer to buy the 2-year non-callable back if it is selling for more than \$100 at time 1.0. Therefore:


- If the market value of the 2-year non-callable is less than \$100, the issuer shall not call the bond → the callable will be the same as the 2-year non-callable. In this case, the cash flow of the callable at time 1.0 will be the same as that of the 2-year non-callable, **which is less than** the cash flow of the 1-year non-callable.
- If the market value of the 2-year non-callable is greater than \$100, the issuer will call the bond → the callable will be the same as the 1-year non-callable. In this case, the cash flow of the callable at time 1.0 will be the same as that of the 1-year non-callable, **which is less than** the cash flow of the 2-year non-callable.

As can be seen, the issuer's objective is to minimize his/her cash flow obligations of the callable bond. Therefore, by exercising the call feature optimally, the issuer makes sure that the cash flow of the callable at time 1.0 will be the minimum between those of the 1-year non-callable and 2-year non-callable. In other words, compared to either of the non-callable, the callable entails strictly-less-than-or-equal cash flows. As such it is obvious that the callable has to be cheaper than both the 1-year and 2-year non-callable.

4 – Valuation of a callable bond

You agree that the callable bond above should sell for less than the 1-year non-callable as well as the 2-year non-callable, but exactly how much less? Of course it is easy if you know the market price. But what

if the bond doesn't trade that frequent? So that you know, 80% of bonds never trade more than once a year. In such instances, to value a callable bond, our modeling knowledge becomes handy. This is because it is quite straightforward to value a callable bond if we have an interest rate tree. Let's assume we have the following tree of semi-annual interest rates.



As you may notice, there are some crazy interest rates (like 15.44%) in the tree, but that's all right, every tree, if extended long enough, would give that. Importantly, probabilities for extreme occurrences are quite small. And also, our current focus is on how to use the tree for pricing, not how reasonable the tree is.

Pricing of the 2-year non-callable

Let's start to see how we can use this tree to price the 2-year non-callable. And I promise that it is a smooth transition from pricing non-callable bonds to pricing callable bonds. Remember that the cash flows from this bond are as follows:

- At time 0.5: \$4
- At time 1.0: \$4
- At time 1.5: \$4
- At time 2.0: \$104

To price this bond, we will of course do the usual thing by walking backwards along the tree, starting at time 1.5.

- At time 1.5, we are not sure what the price of the bond will be, but we know that there are 4 scenarios. If the 6-month interest rate is 15.44%, the price (excluding the \$4 coupon at time 1.5) of the bond will simply be $\frac{104}{1 + \frac{15.44\%}{2}} = 96.55$. Similarly, if the 6-month interest rate at time 1.5 is 11.44% or 8.48% or 6.28%, the value of the bond at time 1.5 would be \$98.37, \$99.77, \$100.83 respectively.
- Let's take a step back to time 1. At time 1, the price of the bond can be computed using the risk-neutral pricing equation. For example, if we are in the highest node at time 1, the price of the bond will be $\frac{0.5 \times (96.55 + 98.37) + 4}{1 + \frac{11.91\%}{2}} = \95.75 . The \$4 in the numerator is, of course, the coupon that we will receive at time 1.5 regardless of where we are (going up or down). If we are in the

middle node (or the lowest node), by similar calculations, the price of the bond would be 98.72 (or 101.00).

- Now let’s take a step back to time 0.5. If you are in the upper node (or lower node), by very similar calculations, the price of the bond would be \$96.79 (or \$100.44).
- Finally, take a step back to the current time – time 0. Applying the risk-neutral pricing equation one last time, the price of the bond is $\frac{0.5 \times (96.79 + 100.44) + 4}{1 + \frac{7.09\%}{2}} = \99.10 .

Putting the bond’s values at every node of the tree together, we will have the following price tree. This price tree corresponds to the interest tree that we start with. The way we interpret this tree is the same as how we interpret the interest rate tree. For example, we know the price of the bond is \$99.10 now, but we are not sure what the price of the bond would be at time 0.5. But we **know** that it could only be either 96.79 or 100.43 with equal **risk-neutral** probabilities of 50%. The 96.79 (100.43) price corresponds to the scenario when interest rate is 9.19% (6.81%) at time 0.5. And so on, each of the price/value we see here corresponds to one interest rate node we see on the interest rate tree.

0	0.5	1	1.5
			96.5451
		95.75492	
	96.7878		98.3727
99.10046		98.716	
	100.4394		99.7719
		101.0012	
			100.8344

OK, you understand why we need an interest rate tree – simply because it allows us to price the bond above. But once you have the price (at time 0), isn’t that the end? Why bother putting all the prices together to build the above price tree? What purpose does it serve? It turns out that, from the above tree, the price of the 2-year callable bond is only a few calculations away. Let’s now turn to how we can use the tree to price the 2-year callable bond with a call price of \$100.

Pricing of the 2-year callable

In pricing the 2-year callable, the key is just to remember that at time 1 the issuer of the callable will optimally use his/her call right.

- At time 1.5, if the callable has not been called, it will be the same as a 2-year non-callable bond. As such, the values of the callable and the non-callable must be identical at every node of the tree at time 1.5.
- At time 1.0, the issuer may call the bond. However, the issuer will call the bond only if the value of the bond is higher than what he needs to pay in calling it: \$100. Checking the 3 scenarios at time 1, it only makes sense for the issuer to buy back the bond if the value of the bond is \$101.00. Paying \$100 for the bond, effectively the issuer nets \$1.00 thanks to the call feature of

the bond. On the contrary, it will not make any sense for the issuer to call back the bond if its total value is either 95.75 or 98.72. In such instances, it is better for the issuer to leave the bond uncalled. To price the callable, therefore, requires one modification in the bond value at time 1 on the lowest node. Instead of 101.00, since the issuer would optimally call the bond here, the value of the callable should really be \$100 at this node. This modification, in turn, will lower the value of the callable at time 0.5 (lower node only) and ultimately the price of the callable at time 0.


- Stepping back to time 0.5, the total value of the bond at the upper node remains unchanged.

The total value of the bond at the lower node however will change to $\frac{0.5 \times (98.72 + 100) + 4}{1 + \frac{6.81\%}{2}} = \99.96 .

- Finally, let's take the step back to time 0. The price of the callable would be:

$$\frac{0.5 \times (96.79 + 99.96) + 4}{1 + \frac{7.09\%}{2}} = \$98.87.$$

Putting all the values that we just calculated above in a tree, we have:



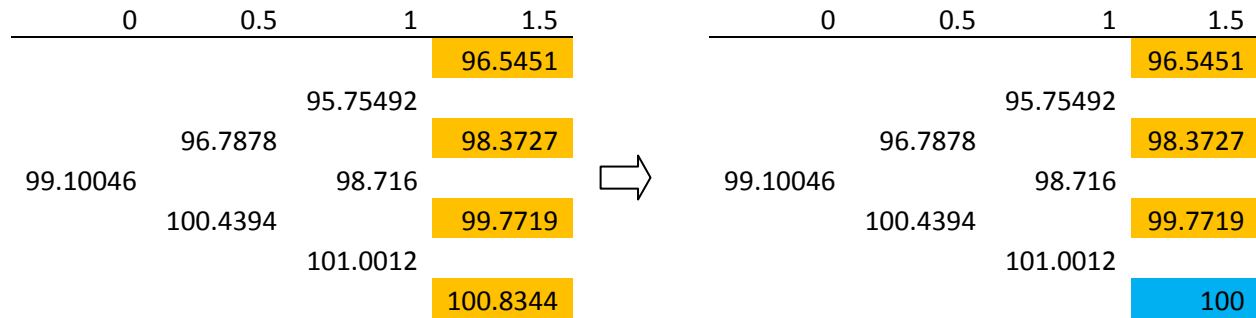
Note that the differences to the tree of non-callable bond prices are highlighted with the blue color. These are the nodes that are affected by the bonds being called at time 1. Note that these nodes correspond to the lowest branch of the tree where interest rates are low. This makes sense because, as we know, bonds are best called/refinanced when interest rates are low.

Pricing of the 2-year callable – What if the bond is callable at time 1.5 as well?

I always like to start things out simple. That's why I've illustrated how to price the callable bond assuming that we can only call the bond at time 1. You can ask the question of what if the bond can be called at time 1.5 as well. In fact many bonds allow for multiple call dates. Further, what if at different call dates, we have different call prices? Fortunately, though more complicated, all these concerns can be addressed using the same framework that we've just gone through.

Let's consider the same callable bond: face \$100, semi-annual coupon of 8%, but now callable either at time 1 with a call price of \$102 or at time 1.5 with a call price of \$100. The key to pricing this bond is simply to start with the price tree of the **non-callable**, and then check at time 1.5 and 1, whether it makes sense to call the bond at any of the node.

First, let's check if it's optimal to call the bond at time 1.5. Remember that the call price at time 1.5 is \$100. If he/she doesn't call the bond and decides to let the bond live until maturity, the bond's value will be the same as that of its non-callable counter-part. As such, starting with the price tree for the non-callable bond and focus on its value at time 1.5, we will be able to tell when the issuer would call the bond: only when the value of the bond exceeds 100.



The tree on the left is the tree of non-callable bond price. If you are wondering where I got the tree from, I just copied and pasted from above when we were pricing the **non-callable** bond. If you already forgot how we got these numbers, please go back and see my detailed calculations up there. For now, I only paint orange the nodes at time 1.5 to highlight the fact that we are only checking whether it is callable at time 1.5.


Examining the 4 possible scenarios at time 1.5, it is easy to see that the issuer will only call the bond at the lowest node where the value of the bond if let alive (until maturity) is \$100.8344. As such for the tree on the right, I replace the value of the bond at the lowest node by 100 and paint the node blue to show that the bond would be called if we get to this node.

So that is done at time 1.5. The second thing we would like to do is to go back to time 1 and check if it is optimal to call the bond anywhere at time 1. You may be thinking: easy stuff – we'd do the same thing again, checking for all possible scenarios at time 1 and comparing the value of the non-callable bond to the call price. If you think so, that would be too fast! Before we get to that stage, we need to **adjust** the bond values at time 1 to reflect the change(s) we have made to the tree at time 1.5.

Specifically, remember how we get the value of 101.0012 at the lowest node of the tree at time 1? It is simply the risk-neutral expected cash flows discounted at the risk-free rate of 6.54%. At time 1.0, if you are in the lowest node, we know that the value of the bond would be either 99.7719 or 100.8344. As such, the value of the bond including the coupon would be: $\frac{0.5 \times (99.77 + 100.83) + 4}{1 + \frac{6.54\%}{2}} = \101.0012 . That is,

for the non-callable bond. Now for the callable, we already work out that if the value of the bond is 100.83 at time 1.5, the issuer will call the bond. As such, the value of the bond the bond on the lowest

node at time 1.0 would be: $\frac{0.5 \times (99.77 + 100) + 4}{1 + \frac{6.54\%}{2}} = \100.60 .

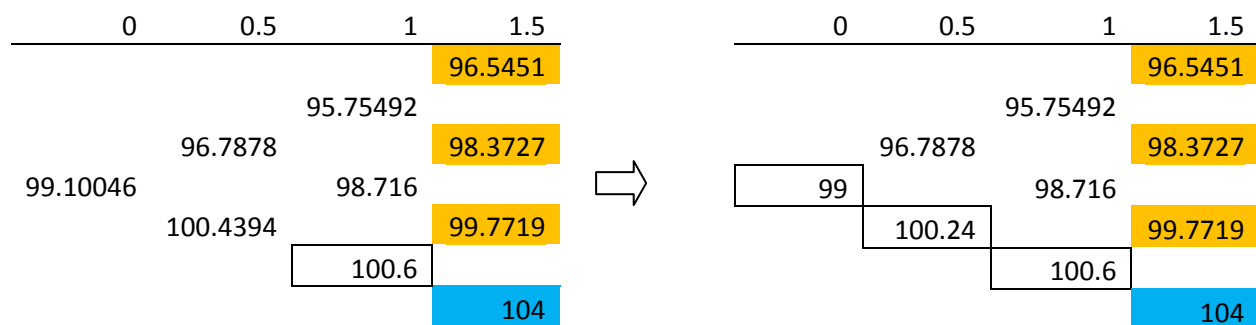


Only after adjusting the nodes of the tree at time 1 as shown above, we can proceed and check whether it is optimal for the issuer to call the bond at time 1. Remember that the call price at time 1 is different. It is \$102. Examining the upmost node of the tree at time 1, the total value of the bond is only \$95.75. It is therefore not worth to pay \$102 for the bond at that node. Similarly, for the other two nodes at time 1, it turns out that it is not optimal for the issuer to call the bond either. You can see that this is simply due to the high call price (\$102) at time 1. If the call price at time 1 were still \$100, the issuer would optimally call the bond at the lowest node where he/she would pay \$100 for the bond that is worth \$100.6.

Now, going back to time 0.5, we need to adjust the value of the bond at the lowest node at time 0.5 as well. This modification is necessary due to the changes we made to the values of the bonds at time 1.

The total value of the bond at the lowest node at time 0.5 should be: $\frac{0.5 \times (98.72 + 100.6) + 4}{1 + \frac{6.81\%}{2}} = \100.24 .

Similarly, going back to time 0, the value of the bond should be: $\frac{0.5 \times (96.78 + 100.24) + 4}{1 + \frac{7.09\%}{2}} = \99




5 – Spreads due to optionality

Let’s take a step back and think about all that we have done: seemingly, *all of a sudden*, we price bonds using trees of interest rates! A lot of calculations make us miss dearly the simple discounting exercises that we used to do: all we need is just a yield curve and then we will just discount 1-year cash flows using the 1-year discount rate, 2-year cash flows using the 2-year discount rate and so on, all we need to

care about is the consistency between the timing of the cash flows and the horizon of the interest rates. All we know is: all of a sudden, life gets so complicated! Can we get back to the simple discounting calculations?

Alright, let's do that. From the tree that we start with (which I put below to save you time flipping back the pages), I can compute the interest rates for different horizons.



I assume that all of you know how to calculate interest rates of different horizons from an interest rate tree. And as such, I won't show the details of my calculations here. However, if you are wondering how I got the above interest rates, you should go back to my notes on Interest Rate Models.

Pricing of the non-callable bond:

Given the above interest rates, in order to price the 2-year non-callable bond, all we need to do is to discount its cash flows using the appropriate interest rates:

$$P = \frac{4}{1 + \frac{7.09\%}{2}} + \frac{4}{\left(1 + \frac{7.54\%}{2}\right)^2} + \frac{4}{\left(1 + \frac{8.03\%}{2}\right)^3} + \frac{104}{\left(1 + \frac{8.55\%}{2}\right)^4} = \$99.10$$

You can see that the price, \$99.10, matches what we got before from the interest tree.

Pricing of the callable bond:

However, when it comes to pricing the callable bond, without the tree, we are stuck! From our earlier calculations, we understand that the price of the callable is lower than the price of the non-callable and should be \$98.87. But it seems that, without having the tree to determine when it is optimal to call the bond, we won't be able to arrive at this price.

Some traders, however, don't like to carry a bulky tree around. They prefer the easiness of the familiar discounting exercises. As such, they decide to do the following: since they know the callable (due to its callability) will be cheaper than its non-callable counter-part, in order to price the callable bond, they will **add a spread** to the discount rates used to price the non-callable. Let's say the spread is 13 basis points. The price of the callable bond will be:

$$P = \frac{4}{1 + \frac{7.09\% + 0.13\%}{2}} + \frac{4}{\left(1 + \frac{7.54\% + 0.13\%}{2}\right)^2} + \frac{4}{\left(1 + \frac{8.03\% + 0.13\%}{2}\right)^3} + \frac{104}{\left(1 + \frac{8.55\% + 0.13\%}{2}\right)^4} = \$98.87$$

which exactly matches the price of the callable bond we have above. I am sure you know how I come up with such a spread that gives the exact price of \$98.87 – Solver, what else?

Once we have this spread, it is **seemingly** convenient because we can then carry the spread around and price other callable bonds by adding the same spread to their discount rates. This practice is **dangerous**, however, since the value of a call option is different from bond to bond depending on their coupon rates, their call price etc. As such, please be careful if you ever do this at work. If you treasure safety, I would recommend using an interest rate tree.

6 – Zero-volatility spread (or Z-spread or static spread)

We have been using the risk-free interest rate tree to price these two bonds with the implicit assumption that they come without default risk. This is **not reasonable**. In fact, accounting for default risk, liquidity risk etc., prices of the bonds would be lower than what we had previously. Let's assume that because of these factors (default risk, liquidity risk), the callable is only selling for \$97.33 instead of \$98.87. To look for a spread for this bond, we again choose a number s that when added to the risk-free discount rates will recover the market price of \$97.33.

$$\$97.33 = \frac{4}{1 + \frac{7.09\% + s}{2}} + \frac{4}{\left(1 + \frac{7.54\% + s}{2}\right)^2} + \frac{4}{\left(1 + \frac{8.03\% + s}{2}\right)^3} + \frac{104}{\left(1 + \frac{8.55\% + s}{2}\right)^4}$$

Using Solver, I have $s = 100$ basis points. Not surprising! From the last section, even without credit and liquidity risk, and just due to optionality, we already have a positive spread of 13 basis points. Now the bond has more risks attached to it, the price is reduced to reflect the extra risks, as such the spread should be larger to account for not only the optionality of the bond but also the credit and liquidity risks associated with it.

This spread is called the zero-volatility spread or the zero spread or the z-spread or the static spread of the bond. Now z-spread or zero spread is just a short form for zero-volatility spread. Why is it called zero-volatility or static spread? Well, it is static relative to other spreads that come off the interest rate tree that we would consider in the next section. Loosely speaking, a spread that comes off a bulky tree seems more dynamic. Likewise, with a tree, we, sort of, see the volatility of interest rates. If the tree is fat, interest rates are volatile, if it is thin, interest rates are stable. As such, a spread that comes from just the risk-free interest rates (like what we have here) as opposed to one that comes off a tree lacks the volatility element, hence is called zero-volatility spread.

Naming business aside, two things are important about static spreads:

- It takes the shape of the yield curve into account (since it is a constant spread added to each of the discount rate for each maturity → we need a yield curve to compute the spread)
- It is some sort of a total spread since it includes everything: some element of optionality, some element of credit risk, some element of liquidity risk etc.

7 – Option-adjusted spread

Option-adjusted spread is an important (though potentially confusing) concept often used in contexts of callable bond, mortgage, mortgage-back-securities pricing.

To understand option-adjusted spread as well as why it has such a name, think about the following situation: **We observe the price of the callable bond to be \$97.33 and we would like to use this information to make some inferences regarding the price of an identical bond issued by the same issuer except that it is non-callable.** I make the previous sentence bold to show that it is important to bear this context in mind in better understanding the concept of option-adjusted spread.


Alright, from the calculations in the previous section, we know that the z-spread of the callable bond is 100 basis points. But, of course, we can't use this spread to price the non-callable because the z-spread of the callable bond includes **everything**. It includes not only credit risk, liquidity risk but also the optionality of the callable bond. While the part of the spread that accounts for credit risk and liquidity risk should be the same for both the callable and non-callable bonds, the non-callable bond has no option in it. As such, it would be unfair to price the non-callable bond using a spread that includes some optionality component in it.

Naturally, therefore, we would like to take away the part that is due to optionality of the callable and use the remaining part to price the non-callable bond. This makes sense because if you take away the optionality component from the callable's z-spread, the remaining spread must be due to credit risks and liquidity risks which are the same for both the callable and non-callable. **This spread is called the option-adjusted spread.** The name derives from the fact that we start from the static spread of a callable and in order to price the non-callable, we need to **adjust** the spread for the **optionality** component in it. To tie everything in an equation, we have:

$$\text{Static spread} = \text{option-adjusted spread} + \text{spread due to optionality of the bond}$$


But how would we do that? How could we disentangle the option and non-option components of the static spread of the callable bond? **The answer: We need an interest rate tree.** You may have a "why a tree" question right now, but let me defer answering that question later, let me show you how we find the option-adjusted spread from a tree first and then explain why later.

First of all, the interest tree that we used before is no longer appropriate! That tree was default-free. Now that our bonds are subject to default risks and liquidity risks, we need to discount their cash flows heavier. And we need to do that at every node of the tree. This suggests that we need to add a spread s to each of the interest rates in the binomial tree. That way, we would discount the bond's cash flows heavier at every node.



In looking for the option-adjusted spread of the callable bond which is selling at t\$97.33, I will choose a spread s in a way that when I use the resulting tree (on the right above) to price the callable, it would recover the market value of the callable bond (at \$97.33). As usual, this process can only be done by trial and error which can be automated by the Solver function in Excel.

I will leave the Solver part to you. For now, to further illustrate how the process works, let's try a random value of $s = 99$ basis points. If it $s = 99$ basis points, we will have a new tree of interest rates that accounts for default/liquidity risks of the bond. The tree will be as follows:



Upon having the tree, we can use the tree to price our callable bond following the usual process. To save space and time, I will not show the details of the pricing process. Rather, I just include here the final tree of bond values. If you like, you can do all the pricing calculations yourself and check your answers against mine. If you are not sure how to price a callable bond using an interest rate tree, please refer to section 4 of this node. In pricing the bond, remember that this bond has a face value of \$100, paying semi-annual coupon rate of 8% and can be called for a call price of \$100 at time 1.0 and time 1.0 only. Note that I paint blue all the nodes that need adjustments due to the call feature of the callable bond.

Amazing enough, with a spread of 99 basis points, we indeed recover the market price of the callable bond which is \$97.33. Ok, so I cheated. I said "let's try a random value of $s = 99$ basis points". The value of 99 basis points I chose to try was not random. I used Solver in Excel and worked out that $s = 99$ basis points would give me the price of the callable bond that I want (\$97.33).

Hopefully, by now you understand at least in a technical, mechanical sense how to compute the option-adjusted spread from the interest rate tree. (And as I mention earlier, compared to the static spread, this option-adjusted spread seems more dynamic, less static flavor since it comes off a tree.) Still, you

may be wondering why such a procedure would give us the exact spread that we want – the option-adjusted spread – the spread that has no option component in it.

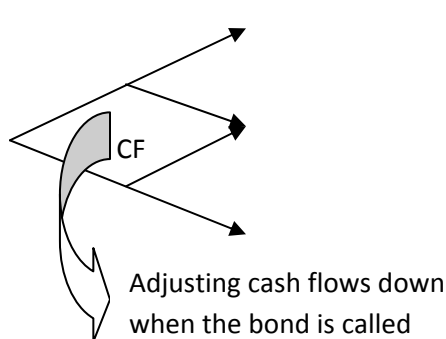
Fair enough – let me explain it.

In explaining it, I find it helps to look back and see how things move along. First, we price the non-callable default-free bond to be \$99.10. Second, we show that if the bond becomes callable, the callable default-free bond should be priced at a lower value of \$98.87, a reduction of \$0.23. Finally, if we allow for the fact that the bond is defaultable, the price of the callable defaultable bond should be even lower at \$97.33, an additional reduction of \$1.54.


It is **important** to recognize that, in using the tree, the two reductions in bond price occur in **different manners**.

- To account for the callability of the bond, we adjust downwards the values of the bonds at time 1.0 at nodes where it is optimal for the issuer to call the bond. This is because the issuer of the callable bond will optimally call the bond whenever the value of the bond is **greater** than the call price he/she has to pay in calling the bond. It is important to understand that all we do here is to adjust the cash flows downwards. We never have to modify our interest rates at each node of the tree to account for the callability of the bond because it is unnecessary.
- To account for default/liquidity risks, unlike how we allow for callability, we don't forcibly modify the cash flows. Instead, we simply discount the cashflows heavier. This involves pushing up our interest rates at each node of the tree by a positive spread.

If you can think of pricing as generally dividing expected future cash flows by (1+the discount rate), or $price = \frac{Expected\ Future\ Cash\ flows}{1+discount\ rate}$, the loosely speaking, the first price reduction (to account for the bond's callability) occurs through a reduction of future cash flows (a reduction in the numerator). On the other hand, the second price reduction (to account for the bond's default/liquidity risks) occurs through an increase in the discount rate (an increase in the denominator).



How to adjust for callability




How to adjust for credit risks/liquidity risks

Since we already account for the callability of the bond by adjusting the cash flows down whenever the bond is called, the spread (99 basis points in the above example) we add to the risk-free interest rate

tree **has nothing to do with the callability of the bond**. In other words, such a spread by which we push the whole tree up in pricing the callable **only accounts for the credit risks and liquidity risks of the callable bond**. Therefore, the spread of 99 basis points that we found above is the option-adjusted spread that we need – a spread without the optionality component!

Once we find the option-adjusted spread, we can use it to price the non-callable bond since we would have a new interest rate tree that allows for the credit/liquidity risks of the bond issuer.



Using the tree to price the non-callable bond should be straight forward and even simpler than pricing the callable bond because we don't even have to check for when it is optimal for the issuer to call the bond. Again, I won't go through the details of the pricing process here. Rather, I would put here the resulting price tree for you to compare your calculations to.

According to my calculations, the final price of the non-callable bond is \$97.35, just slightly above the price of the callable bond at \$97.33. This means that the value of the callable feature is only \$0.02? Or in other words, if we go back to our equation:

$$\text{Static spread} = \text{option-adjusted spread} + \text{spread due to optionality}$$

Since the option-adjusted spread is 99 basis points and the static spread is 100 basis points, the spread due to optionality is really small: $100 - 99 = 1$ basis point! This is crazy! Due to our calculations earlier, the spread due to optionality is 13 basis points (remember?). What happens to it that reduces it by 13-fold? Answer: the extra credit/liquidity risks. But why?

If the firm has credit risks/liquidity risks, its bond should generally sell for less compared to the case when it has no credit or liquidity problems. This should be true at every node of the tree because to account for credit/liquidity risks, we need to use higher discount rates at every node of the tree. To better illustrate this, I will put the price tree of the non-callable bond with and without credit/liquidity risks together and hopefully you can have a sense of what I mean. Just comparing any pair of corresponding nodes, you would see that the value of the defaultable bond is always less than its default-free counter- part.

Default-free non-callable bond				Defaultable non-callable bond			
0	0.5	1	1.5	0	0.5	1	1.5
			96.5451				96.1034
		95.75492				94.8869	
	96.7878		98.3727		95.48639		97.9142
99.10046		98.716		97.34568		97.807	
	100.4394		99.7719		99.0705		99.3003
		101.0012				100.0601	
			100.8344				100.3528

As a consequence, the value of the call option to the issuer of the defaultable callable bond will be smaller. This should be clear by looking at the lowest node at time 1 where the issuer of the two bonds will call (because the value of the bond here is greater than \$100) and comparing the value of the call to each issuer. The issuer of the default-free bond pockets the difference of \$1.00012 between the value of the bond (if let alive) and what he has to pay in calling (\$100). On the other hand, the issuer of the defaultable bond earns only \$0.0601.

Another way of thinking about this is: relative to the default-free case, if you have the extra credit/liquidity risks and thus have to face relatively higher borrowing costs, you will be less likely to call the bond – the same way you will be less likely to refinance your mortgage if the current interest rates are high. If you are less likely to call the bond, its value should be smaller.

And if the value of the callable feature becomes smaller, naturally the spread due to optionality becomes smaller as well. This explains why the spread due to optionality has decreased from 13 basis points in case of no default/liquidity risks to 1 basis point when we allow for credit/liquidity risks.

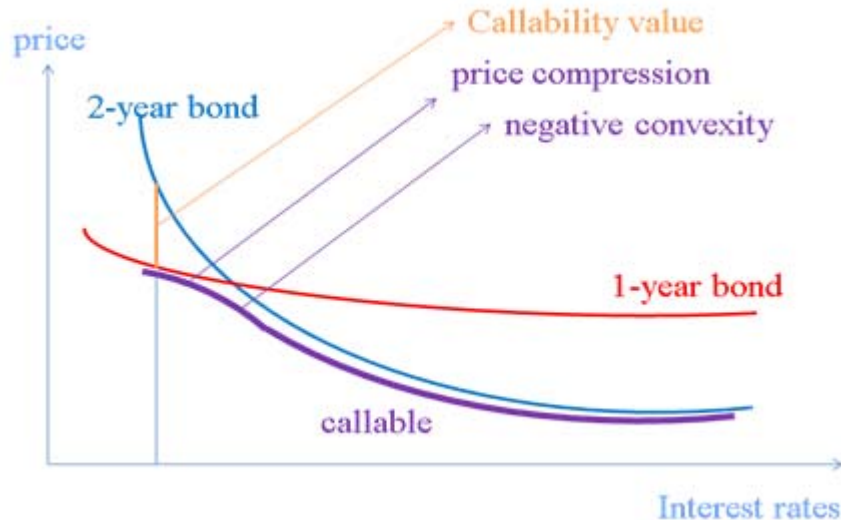
As some final words of caution, the option-adjusted spread measures that we learn so far are model-dependent. In other words, we need some interest rate tree – some model of interest rate – to calculate this measure. Whenever we talk about model, there is one extra dimension of risk, namely model risk. As such to be precise, the option-adjusted spread contains credit/liquidity risks as well as model risks. The parts related to credit/liquidity risks should be positive! However, the part related to the model risks, it could be positive or negative depending on how we construct the model.

8 – Callable bond prices and interest rates

As I already mentioned in the intro class, bond prices and interest rates are like two ends of a see-saw. When interest rates go up, bond prices go down and vice versa. The same analogy applies to the relationship between callable bond prices and interest rates. If you plot the price of the 2-year callable bond considered earlier against interest rates, you will have a negatively sloped graph as you would for other bonds.

However, as I showed above, since the 2-year callable is bounded from above by prices of the 1-year non-callable as well as the 2-year non-callable, the pricing function of the 2-year callable bond will take a

somewhat special shape. In the graph below, I plot the pricing function of the 2-year callable bond (in purple color) together with that of the 1-year (in red color) and 2-year (in blue color) non-callable bonds for comparison purposes.



When interest rates are really high, borrowers are less likely to refinance their borrowings. In other words, there are high chances that the issuer will let the bond live until maturity with a very high probability. In these cases, the cash flows from the callable are very much similar to those coming from the 2-year non-callable bond. This explains why the shape of the purple graph (callable pricing function) comes very close to that of the blue curve (2-year non-callable pricing function) as interest rates rise. In the other extreme, when interest rates are low, borrowers are highly likely to refinance their borrowings. If the issuer calls the 2-year callable with a very high chance, the cash flows from the callable are very much like those coming from the 1-year non-callable. This explains why the purple graph comes very close to the red graph (1-year non-callable pricing function) as interest rates go really low.

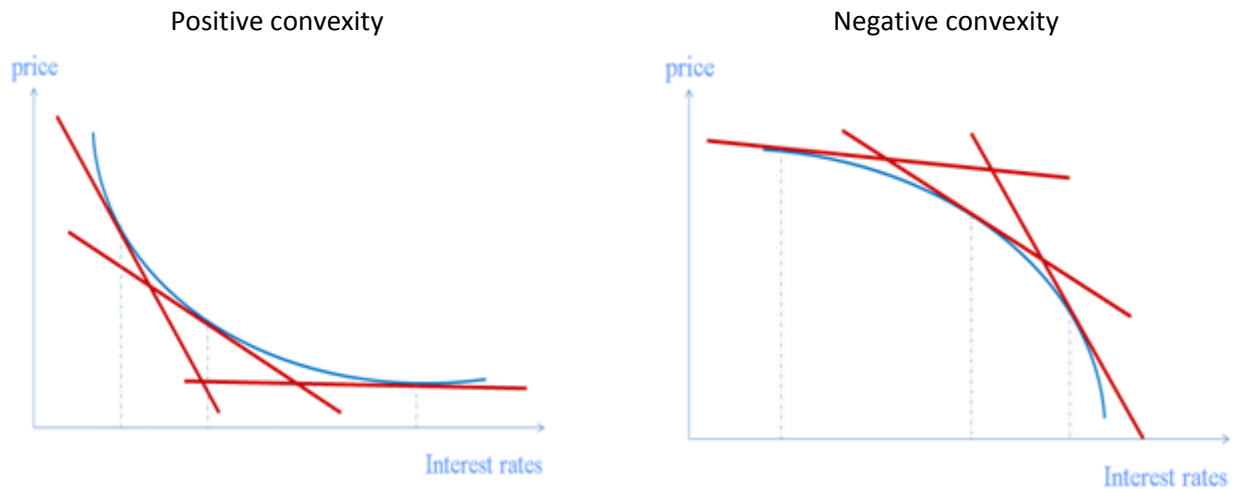
We can see that as interest rates decrease, the pricing wedge between the 2-year callable and 2-year non-callable becomes more and more pronounced. This difference is precisely due to callable feature attached to the callable bond. This difference is the value of the callability of the bond. In addition, we can see from the graph that given a reduction in interest rates, bonds would generally appreciate in values. However, due to the callable feature, the extent to which callable bonds' prices increase is much less than that for the 2-year non-callable. This phenomenon is referred to as price compression.

Finally, in the intermediate range of interest rates, the pricing function of the callable displays a negatively convex shape. If we are worried since you are not sure what negative convexity means, don't worry we will tackle that with more details shortly in the next section.

9 – Duration/dollar duration and convexity/dollar convexity of callable bonds

Before talking about either duration or dollar duration of callable bonds, I would like to remind you of how I showed, in our intro class, that the dollar duration of bonds is simply the slope of the tangent line of the pricing function. After all, (dollar) duration measures the sensitivity of bond price with respect to changes in interest rates. If the slope of the tangent line is steep, a given change in interest rate will lead to a large change in bond price. On the other hand, when the tangent line is rather flat, a given change in interest rates will only cause a small deviation in bond prices.

It turns out that dollar duration as well as duration of regular bonds decrease as interest rates increase. Why is that? Just think of how you compute duration of a zero coupon bond: $\frac{T}{1+\frac{y}{2}}$ where T is maturity and y is the semi-annual interest rate. Obviously, as interest rates (y) increase, duration goes down. Graphically speaking, as we move along the pricing function of regular bonds (like what we have below on the left hand side), the slope of the tangent line gradually decreases – the bond becomes less and less sensitive to changes in interest rates.

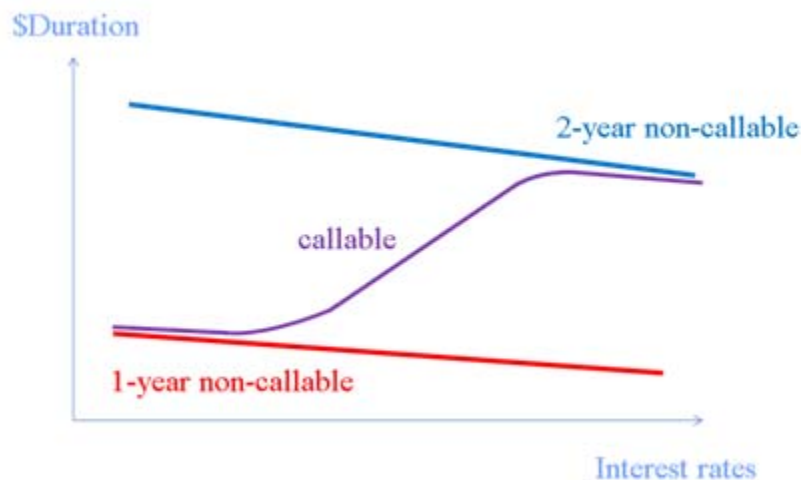


It is precisely this decrease in duration as interest rates increase that gives rise to positive convexity for the regular bonds. To contrast between positive convexity and negative convexity I also include on the right hand side an example of negative convexity. There you see that as interest rates increase the tangent lines become steeper and steeper.

Another interesting observation is that:


- With positive convexity, the blue curve always lies above the tangent lines. If you remember well the duration and convexity approximation, the difference between the blue curve and the red tangent line is precisely what our convexity adjustments go after. If the blue curve always lies above the red line, this explains why our convexity adjustment is always positive – hence the name positive convexity.
- With negative convexity, the blue curve instead always lies below the tangent lines. This time, the difference between the blue curve and the red line is always negative. As such the convexity adjustment for this case will always be negative – hence the name negative convexity.

Let's now get back to our 2-year callable bond and think about its dollar duration and how it will change as interest rates increase. For illustration purposes, I plot on a graph here the \$duration of the 2-year non-callable (in blue) and the \$duration of the 1-year non-callable (in red) together with the \$duration of the callable bond (in purple). You can see that the \$duration of the 2-year non-callable and 1-year non-callable decrease as interest rates increase as I explained above. In addition, \$duration of the 2-year non-callable is always higher than \$duration of the 1-year non-callable. This makes sense since for a longer maturity, the 2-year non-callable should be more sensitive to interest rate changes than the 1-year non-callable. Let me know explain the shape of the purple graph which tells us how the \$duration of the callable changes as interest rates change.



- For really low interest rates, borrowers are likely to refinance their borrowings –equivalently, there are high chances the 2-year callable bond will be called. Therefore, for very low interest rates, the 2-year callable bond behaves very similarly to the 1-year non-callable. As such, for the low range of interest rates, the duration/dollar duration of the 2-year callable will look very much like that of the 1-year non-callable. In other words, the purple graph should come really close to the red line when interest rates are really low.
- For really high interest rates, borrowers are almost certain not to refinance – or in other words, there are high chances that the 2-year callable will live until maturity. As such, for the high range of interest rates, the duration/dollar duration of the 2-year callable will look very much like that of the 2-year non-callable. In other words, the purple graph should come really close to the blue line when interest rates are really high.
- For the intermediate range of interest rates, we don't know exactly how the purple graph will turn out. However, one thing we know for sure is that the purple graph should be continuous. To get from where it is when interest rates are low (close to the red line) to where it is when interest rates are high (close to the blue line), therefore, means that \$duration of the callable bond will increase as interest rates increase **at least for some intermediate range of interest rates.**

As explained above, this behavior of duration (increase when interest rates increase) creates negative convexity. This is exactly what we see from the figure I include in section 8, which I will put here again to save you time flipping back:



One way to think about how this negative convexity comes about is to imagine that you are driving along the red curve and approaching whether the blue curve and the red curve interests. Then you want to make a right turn into the blue curve. When you are making a right turn, as long as you are not going deadly slow, you will make a bending shape similar to the purple graph that we have, which is negative convexity.

After all, what is the deal about negative convexity? Why is it so important that we have wasted quite some time talking about it? It turns out that it has quite important implications in hedging, especially for those companies that invest in fixed rate mortgages that, as you will see later on, also display negative convexity. In minimizing their exposure to interest rate risks, naturally these firms would like to balance/match the durations and convexities of their assets and liabilities. Therefore, if they have negative-convexity assets, they would like to have negative-convexity liabilities that give them a natural hedge. And one way to have negative-convexity liabilities is to issue callable bonds.

10 – Computation of Duration/dollar duration and convexity/dollar convexity of callable bonds


As you already see, unfortunately, duration and convexity measures of callable bonds are quite different from those of the more regular bonds. Due to this, all the techniques that we learn in computing duration and convexity for the regular non-callable bonds are no longer applicable here. I will first show you the process by which we can compute the (dollar) duration of the 2-year callable bond. Next, I will show you how the convexity can also be computed, using a similar process. These duration and convexity measures are called effective (dollar) duration and effective (dollar) convexity to differentiate them from the other measures that we have learnt. For the sake of brevity, in what follows, I will omit the word effective in front of duration and convexity with the implicit assumption that you understand that I refer to effective (dollar) durations and (dollar) convexity.

(Dollar) Duration:

First, I mentioned earlier, dollar duration is the slope of the tangent line to the pricing function of any bond, callable or non-callable. For non-callable regular bonds, fortunately, we have formulas. With callable bonds, however, we don't have that luxury. To compute (dollar) duration for a callable bond, we need to go through a set of steps that turn out to be applicable to every bond:

1. First, we compute the current value of the bond at the current level of interest rates. Let's say the current value of the bond is V_0 .
2. Second, we increase the interest rate level by a small amount Δy . How small? Something smaller than 10 basis points would be small. We then compute the value of the bond at this new level of interest rates and call it V_+ .
3. Third, we decrease the interest rate level by the same small amount Δy and compute the value of the bond at this new level of interest rates and call it V_- .
4. Dollar duration of the bond would be $\frac{V_+ - V_-}{2 \times \Delta y}$.
5. Duration of the bond would be $\frac{V_+ - V_-}{2 \times V_0 \times \Delta y}$.

I will now explain why these steps make sense. My explanation (though short and intuitive) only serves the purpose of satisfying those curious about the reasoning behind these steps. It won't be on the test. As such, those of you who think that it is too much to read already, you can skip this section and go straight to the example of how we can actually implement these steps.




Basically, if we have a nice and neat equation that gives us the slope of the tangent line to the pricing function or dollar duration, that would be nice. Otherwise, the slope of the tangent line would be similar to the slope of the brown dotted line on the graph above. This dotted line connects 2 points on the blue curve: the first point is when we decrease interest rates by a small amount Δy and the second point is

when we increase interest rates by the same small amount Δy . The slope of this line is simply $\frac{V_+ - V_-}{2 \times \Delta y}$ which is the formula we use for the dollar duration. Since duration = dollar duration/price, the duration of the bond is simply $\frac{V_+ - V_-}{2 \times V_0 \times \Delta y}$.

Importantly, none of the steps is particular to callable bonds, which means: the process is applicable to any kind of bonds or interest rate sensitive securities.


To provide a concrete example of how to carry out the above steps, let's consider computing the dollar duration and duration of our 2-year callable bond.

- The good news is step 1 is already done because we already price the callable bond in the preceding sections. To remind you of what we did, I include here the interest rate tree we used as well as the resulting price tree. Again, nodes painted blue are those affected by the bond being called at time 1.




So our $V_0 = 97.33$.

- In step 2, we need to increase interest rates by a small amount Δy . Let's assume $\Delta y = 10$ basis points. We will talk a little bit about this later, but for now we will add 10 basis points to each of the node of the interest rate tree and then use the tree to price the callable bond again. Again, I will not give the details of the pricing calculations but rather post here the resulting price tree for you to compare your calculations against.



As you can see, as interest rates increase by 10 basis points, bond price decrease to $V_+ = \$97.17$. It turns out that at nowhere along the tree that it is optimal to call the bond given this increase in interest rates. For this reason, we don't have any blue nodes this time.

- Step 3 is similar to step 2, except that we now subtract 10 basis points from the original interest rate tree. The resulting rate tree and price tree are as follows:



As interest rates decrease, bond price increases to $V_- = \$97.49$

Given V_+ , V_- , V_0 and $\Delta y = 10$ basis points, we can compute the bond's dollar duration and duration as follows:

- Dollar duration = $\frac{V_+ - V_-}{2 \times \Delta y} = \frac{97.17 - 97.49}{2 \times 0.0010} = -160$.
- Duration = $\frac{\text{Dollar duration}}{\text{price}} = \frac{-160}{97.33} = -1.6439$.

The negative signs in front of dollar duration and duration are just to indicate that bond prices and interest rates move in opposite directions: as interest rates increase, bond prices decrease and vice versa.

(Dollar) Convexity:

It turns out that it is quite straightforward to compute convexity and dollar convexity once you have V_+ , V_- and V_0 all ready. I will first show you the formulas to perform the needed calculations. Next, for those curious, I briefly and graphically explain the ideas behind the formulas. Again, this part will not be on the test. Therefore, if you are not interested in reasoning, you could safely skip it.

- To compute dollar convexity, we use the following formula: $\frac{V_+ + V_- - 2V_0}{2(\Delta y)^2}$.

Plugging the values for V_+ , V_- , V_0 in the formula, dollar convexity of the callable bond is:

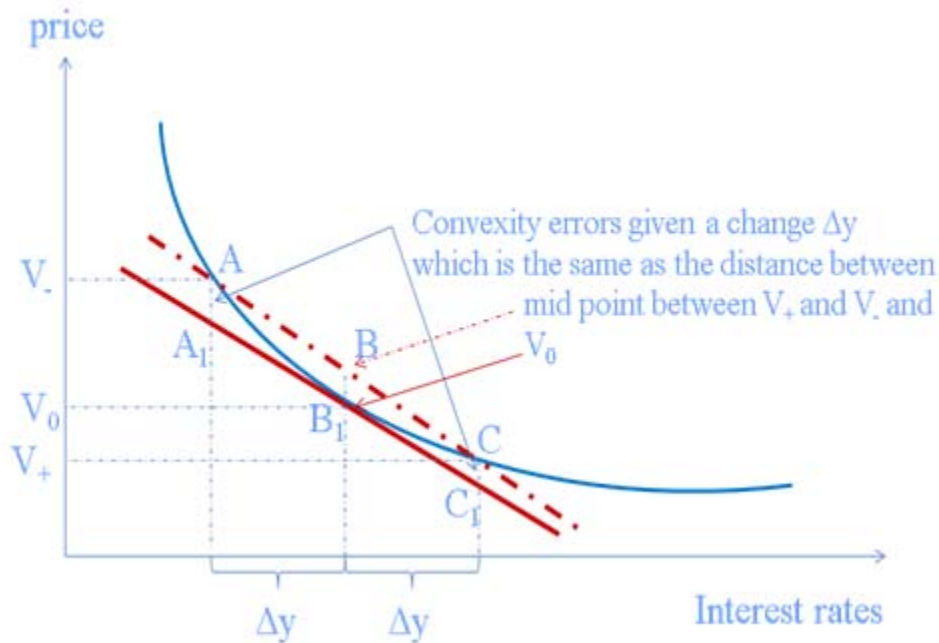
$$\frac{97.1707 + 97.4853 - 2 \times 97.3318}{2 \times (0.0010)^2} = -3800$$

- To compute convexity, we simply divide dollar convexity by the price, which will give:

$$\frac{-3800}{97.33} = -39.04.$$

As you can see, the (dollar) convexity for the callable is negative.

Let me now briefly provide the intuition behind the formula : $\frac{V_+ + V_- - 2V_0}{2(\Delta y)^2}$ for dollar convexity.



As you already know, the whole idea of convexity adjustment is to correct for the deviations between the blue curve and its tangent line. Given an increase of Δy in interest rates, the difference between the blue curve and the red tangent line is the distance CC_1 . Given a decrease of Δy in interest rates, the difference between the blue curve and the red tangent line is the distance AA_1 . Although we don't know exactly what CC_1 and AA_1 are, we know their average which is the distance BB_1 . The height of B is simply the average of V_+ and V_- . the height of B_1 is V_0 . Therefore, the length of BB_1 is simply: $\frac{V_+ + V_-}{2} - V_0 = \frac{V_+ + V_- - 2 \times V_0}{2}$. This explains why the convexity formula, $\frac{V_+ + V_- - 2V_0}{2(\Delta y)^2}$ is proportional to $\frac{V_+ + V_- - 2 \times V_0}{2}$.