Non perturbative key-rate contributions to bond returns

Marco Marchioro
www.statpro.com
Version 1.0

November 2016*

Abstract

We show a non-perturbative method to split the bond return into components coming from the different key-rates, the credit spread, the carry term, and another couple of minor components. We explicitly show the results for the case of a fixed-rate coupon bond with a price computed using cash-flow discounting with a yearly-compounded yield curve. We also explain how to extend the method to a much wider-range of securities depending on an interest-rate curve, the credit spread, and possibly some other financial variables. Among the minor return components we find a term coming from the compounding effect the of interest rates and the credit spread, as well as a contribution that summarises the effect of the other risk drivers. Finally we show how the bond carry term, in the presence of an accruing coupon, can be split into a coupon part and a convergence contribution.

Keywords: performance contribution, fixed-income attribution, key rates, bond pricing, interest-rate derivatives, fixed-income securities, pricing functions, interest-rate contribution, credit contribution, cash-flow discounting, spread/rate contribution, carry contribution, convergence contribution, coupon contribution

1 Introduction

In fixed-income attribution one is often interested in splitting the security return into contributions that have a precise financial meaning. For example when we hold a bond that has a certain return we would like to know how much of this return is due to interest-rate changes, credit-spread variations, and so on and so forth. Furthermore, since there usually are different rate changes for the different term-structure maturities, one may also be keen to split the overall interest-rate return according to contributions given by the different curve sections.

In this paper we show a non-perturbative method that can be used to split the security return into contributions coming from the different curve maturities, the credit spread, a calendar term, and possibly some other terms. We stress that the method is non-perturbative because it works regardless the displacement size of either the interest rates or the credit spreads.

Performance attribution In the asset-management industry *performance attribution* is normally used in determining which part of the return is created by the different actors in the investment process. Typically this type of analysis is performed on a portfolio relative to its benchmark, however total-return attribution is also a possibility, for example, for certain hedge funds. There are many methods available for performance at-

*Latest minor review November 9, 2016

tribution and a good overview is provided by reference [1].

Consider for example a mutual-fund house with a currency department, an allocation group, a selection division, and a trading desk, then the period return R could be split in contributions as

$$R^{\text{attr}} = C^{\text{fx}} + C^{\text{alloc}} + C^{\text{trad}} + R\,. \qquad (1)$$

In the above expression the different components may have the following interpretation:

1. C^{fx}, which only present for securities that are in a currency different from the portfolio currency, represents the return due to the exchange-rate variations

2. C^{alloc} represents the return due to choosing a portfolio sector allocation different from that in the benchmark

3. C^{trad} is the performance due to the trading timing as opposed to buy a security at the beginning of the day or sell it at the end

4. R is the price-variation return due to the market price changes from the beginning of the period to the end of it. (We denote this contribution by the letter R because it is sometimes called the security *simple return*.)

Equation (1) is just one example of many possible in performance attribution and, depending on the investment process, more or less terms might be involved.

Fixed-income attribution For investment portfolios with a strong presence of bond securities, the changes in the interest-rate curves and in the credit spreads affect multiple assets at the same time. A detailed description of fixed-income attribution can be found in reference [2], which however provides the traditional perturbational approach to the subject. In fixed-income attribution we are interested in determining the return contributions that are common to multiple bonds. For a single security we consider the price-variation return R, defined as

$$R = \frac{\Delta P}{P_b} = \frac{P_e + C - P_b}{P_b}\,, \qquad (2)$$

where P_b is the bond price at the beginning of the period, P_e that at the end, and C is the coupon paid during the period. To simplify the formulas that follows, from now on, we assume C=0.

In the reminder of this paper we show how to split the price-variation return, defined above, into contributions coming from the various key rates, the credit spread, the calendar term and few other terms. In order to simplify the description of the contribution computations, as seen in the next sections, we make a lot of assumptions, however it should be noticed that, the presupposition of small interest-rate variations is not one of them.

Simplified-model assumptions In order to precisely split the return into contributions one should take into account all the details and the market conventions connected to the bond trading and to the quantitative models behind the pricing functions. Therefore a complete treatment of the subject of security contributions would fill a small volume. The aim of this paper is not to describe all the details involved in the computation of fixed-income contributions, however that of outlining the basic concepts behind the method. Hence, in order to provide a method explanation as intuitive as possible, we make the following assumptions:

- We analyse the return over a single period and do not show how to compound together the contributions from different periods.
- We assume that no coupons are paid during the observation period.
- We do not consider accruals, i.e. we always assume to be at the beginning of the coupon. (As a consequence the dirty and the clean prices always match.)
- We presume that there is an immediate transaction settlement so that the exchange of cash for the security is instantaneous.

- We do not consider as relevant the difference between the quoted price and the bond evaluation, performed with our quantitative model and with the available market data.
- We consider a continuous time measured in years and do not bother to deal with the details of the different market conventions related to the computation of the actual cash flows from the coupon rates.
- We model the market-expected credit risk using a flat spread. In other words we consider the same credit spread for all cashflow maturities.
- In order to provide explicit formulas for the computation of contributions we consider only the pricing function of fixed-rate coupon bonds.

As it turns out all these assumption can be relaxed in order to obtain a proper quantitative computation of performance contributions, however in doing so we would turn our attention away from the simple concepts behind the method. Hence in this paper, unless stated otherwise, we work under the assumptions outlined above.

2 Sample fixed-rate coupon bond

To illustrate the computation of the return contributions for an elementary security, in this section, we consider a sample fixed-rate coupon bond. While the test bond we use is very simple, the results that we find can be applied to a much wider range of securities. In order to focus on the method itself we initially consider a flat interest rate term structure and a flat credit spread. The term flat in this contest means that we are not considering the maturity dependence of interest rates and credit spreads.

2.1 Return of a fixed-rate coupon bond

To obtain additive return contributions in a straightforward way, we consider a bond that jointly depends on the interest rate r and the credit spread s through the yield y=$r+s$. Also we consider a flat yield curve so that all cash flows can be discounted using the same yield y for all maturities. For simplicity we assume a bond on it's first day of trading so that there is no accrual and the clean price matches the dirty price. We measure time as a continuous variable so that the real number T=1.0 represents the maturity of one year, regardless the number of days in that year. Also, unless explicitly stated otherwise, we make all the assumptions oulined at the end of section 1.

We examine a bond expiring four years from the current date, paying at the end of each year coupons all with the same value C. At maturity the bond holder will receive, together with the usual coupon, also a redemption amount of 100.0. Assuming the settlement date to match the evaluation date, given a bond yield y we can compute the security price as

$$P(y) = \sum_{i=1}^{4} D(y, T_i) \cdot C + D(y, T_4) \cdot 100.0 \,,$$

where the discount factor $D(y, T)$ can be evaluated, for example, using annual interest compounding as

$$D(y, T) = \frac{1}{(1+y)^T} \,,$$

with T the cash-flow time to maturity.

For example consider the bond-price evolution from the period beginning at t_b and ending at t_e. When the initial yield is y^b=5.0% and the coupon amount is C=10, we can compute the bond cash-flow present values as the sum of the discounted coupons and redemption as shown in the following table,

T_i	y_i	$D(y_i, T_i)$	Cash	Disc.
1	5.0%	95.2%	10	9.524
2	5.0%	90.7%	10	9.070
3	5.0%	86.4%	10	8.638
4	5.0%	82.3%	110	90.497
		Sum	140	117.730

so that the corresponding bond price is

$$P_b = P(y^b) = P(5.0\%) = 117.730 \,.$$

In the previous table the column T_i denotes the maturity of a given cash flow, the column y_i the yield used for discounting that cash flow, the column *Cash* the cash-flow future value, and finally the column *Disc.* the value of the discounted cash flow (or it's present value). We keep the price precision to at least three decimal places because later we need to compute the difference between prices for different yields.

Consider now the same bond at time t_e, i.e. after a certain time has passed, however less than a day so that we can use the same cash-flow time to maturity as the previous table. In subsection 3.4 we relax this hypothesis by introducing the carry component. Let us suppose that during the period between t_b and t_e, the yield increases to a new value

$$y^e = 5.3\% ,$$

so that the new bond cash-flow present values can be computed as

T_i	y_i	$D(y_i, T_i)$	Cash	Disc.
1	5.3%	95.0%	10	9.497
2	5.3%	90.2%	10	9.019
3	5.3%	85.6%	10	8.565
4	5.3%	81.3%	110	89.470
		Sum	140	116.550

to give a new price:

$$P_e = P(y^e) = P(5.3\%) = 116.550 .$$

Suppose that we purchase the security at t_b, for a cash amount P_b. If we were to sell the same security at time t_e, at the model price P_e, we would make a cash profit of

$$\Delta P = P_e - P_b = -1.179 ,$$

which turns out to be a loss. The corresponding return, or more precisely the price-variation return, can be computed as

$$R = \frac{P_e - P_b}{P_b} = -1.002\% .$$

The goal of this paper is to describe a method to split the cash profit, and therefore also the return R, into additive components determined by the bond structure. Since the return R can simply be obtained by dividing the cash profit by the initial price P^b we limit ourselves to compute the fixed-income contributions to the cash profit (or loss)

$$\Delta P = P_e - P_b .$$

When the percentage contributions are needed, we can obtain them by simply dividing the cash contributions by P_b.

2.2 The cash-flow method for return contributions

Since both prices P_b and P_e are obtained as the sum of the discounted cash flows, in order to compute their difference we can sum the cash increment of each present value for all coupon maturities. We detail the computation of these differences in the following table:

T_i	Disc. y^b	Disc. y^e	C_i^y	c_i^y
1	9.524	9.497	-0.027	-0.023%
2	9.070	9.019	-0.052	-0.044%
3	8.638	8.565	-0.074	-0.063%
4	90.497	89.470	-1.027	-0.872%
Sum	117.730	116.550	-1.179	-1.002%

In the second column we write the discounted cash flow at T_i computed with a yield y^b, in the third column the present value of the i-th cash flow discounted using y^e, in the C_i^y column we compute the difference between the third and the second columns. Finally in the last column we show the percentage contribution c_i^y with respect to the initial price P_b.

At this point it is natural to split the cash profit ΔP as the sum of contributions coming from each coupon maturity (i.e. the fourth column of the above table):

$$\Delta P = C_1^y + C_2^y + C_3^y + C_4^y .$$

Given this result we can similarly split the price-variation return as

$$R = c_1^y + c_2^y + c_3^y + c_4^y , \qquad (3)$$

with $c_i^y = C_i^y / P_b$.

Observing the contribution numbers above we notice that, even though the bond coupons paid are all of the same size and the yield curve is flat, the cash flows paid later provide a higher contribution than those paid earlier. In figure 1 we give a graphical representation of equation (3) and show both the *single* coupon contributions to the return (blue arrows) together with the total return R (red arrow).

2.3 The simulated-price method for return contributions

The cash-flow method described above to compute the maturity contributions to the cash profit has the main advantage to be straightforward and very intuitive. The problem with that method is that it is limited to yield contributions and that we need the exact knowledge of the cash-flow structure. Also it is only applicable to securities that can be computed using cash-flow discounting and its generalisation to a generic quantitative model for bond pricing is questionable. We therefore provide here an alternative and equivalent method that only requires the computation of the security present value, however on simulated yield curves. For this reason we use the term *simulated-price method* for the procedure that we are about to describe.

First of all we notice that the price changes from P^b to P^e because all yields y_i's, for each maturity T_i, move from y^b to y^e simultaneously. However for any given i=1, 2, 3, 4, in order to obtain the contribution C_i^y we only need the yield y_i to change from y^b to y^e. Therefore to compute, for example, C_1^y we evaluate the bond price on a simulated yield curve with y_1=y^e and y_i=y^b for i=2, 3, 4. In other words we let only the yield used to compute the first cash flow to be y^e, while we keep all other yields to y^b and obtain the following discounted cash flows:

T_i	y_i	$D(y, T_i)$	Cash	Disc.
1	5.3%	95.0%	10	9.497
2	5.0%	90.7%	10	9.070
3	5.0%	86.4%	10	8.638
4	5.0%	82.3%	110	90.497
		Sum	140	117.703

We notice that the shortest-maturity contribution C_1^y can be computed as

$$\begin{aligned} C_1^y &= P(\{y^e, y^b, y^b, y^b\}) - P(y^b) \\ &= 117.703 - 117.730 = -0.027\,, \end{aligned}$$

where the notation $P(\{y^e, y^b, y^b, y^b\})$ means that the bond price has been computed by setting the first-maturity yield to y^e and all others to y^b.

Similarly, for the other contributions we have

$$\begin{aligned} C_2^y &= P(\{y^b, y^e, y^b, y^b\}) - P(y^b) = -0.052\,, \\ C_3^y &= P(\{y^b, y^b, y^e, y^b\}) - P(y^b) = -0.074\,, \\ C_4^y &= P(\{y^b, y^b, y^b, y^e\}) - P(y^b) = -1.027\,. \end{aligned}$$

The contribution results obtained in this way exactly match those obtained using the cash-flow method. The advantage of this method, however, is that it only requires the ability to compute security prices with a non-flat yield curve. Therefore the simulated-price method can also be used in those cases where we compute the bond price with a quantitative method that is not cash-flow discounting.

Non flat yield curve Let us consider the case of a non-flat yield curve and assume that the bond price P^b at t_b is computed using the yield vector $\mathbf{y}^b$, with $\mathbf{y}^b$=$\{y_1^b, y_2^b, y_3^b, y_4^b\}$, while the price P^e at the end of the period with $\mathbf{y}^e$=$\{y_1^e, y_2^e, y_3^e, y_4^e\}$. The cash profit in this case can be computed as

$$\Delta P = P(\{y_1^e, y_2^e, y_3^e, y_4^e\}) - P(\{y_1^b, y_2^b, y_3^b, y_4^b\})\,.$$

Using the simulated-price method, for example for the first maturity, we can compute the return contribution as

$$C_1^y = P(\{y_1^e, y_2^b, y_3^b, y_4^b\}) - P(\{y_1^b, y_2^b, y_3^b, y_4^b\})\,,$$

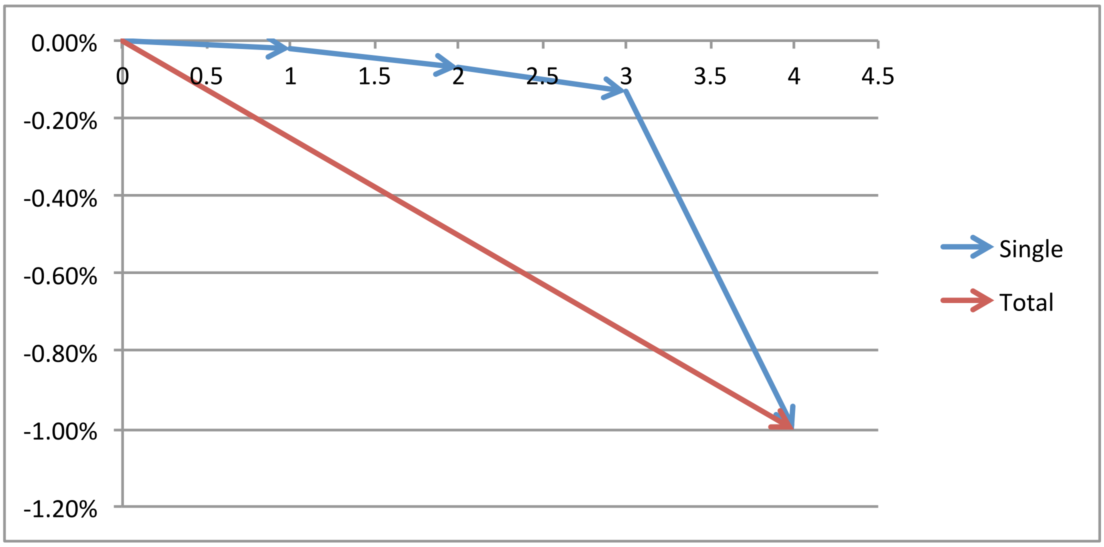

Figure 1: Graphical representation of equation (3): the blue arrows represents the single return contributions coming from each cash flow and the red arrow represents their sum.

while the other contributions can be computed similarly. In this way we extend the computation of the yield contributions to non-flat yield term structures.

Notice, again, how all computations have been performed without making the assumption of small yield-curve variations, hence the method is non perturbative.

2.4 Interest-rate and credit-spread contributions

In previous subsections we describe two methods to compute the maturity-dependent contributions to the return of a security, when priced in terms of the yield curve. In this section we assume the yield for each coupon maturity to be given by the sum of a risk-free interest rate r and a credit spread s. Hence for each maturity T we consider the discount factor as a function of both the interest rate and the credit spread:

$$D(r, s; T) = \frac{1}{(1 + r + s)^T}\,.$$

We use the same example discussed earlier for the fixed-rate coupon bond expiring in four years, however we assume that the yield at the beginning of the period y_b is obtained as the sum of r_b=4.0% and s_b=1.0%, i.e. y_b=r_b+s_b=5.0%. We compute the discounted cash flows to obtain P^b as

$$P_b = P(s^b, r^b) = P(4.0\%, 1.0\%) = 117.730\,,$$

which yield the same result as the earlier computation, however here we stress that the yield is now obtained as the sum of a risk-free interest rate and a credit spread. We summarise this result in the following table:

T_i	r_i	s_i	Cash	Disc.
1	4.0%	1.0%	10	9.524
2	4.0%	1.0%	10	9.070
3	4.0%	1.0%	10	8.638
4	4.0%	1.0%	110	90.497
		Sum	140	117.730

We proceeded similarly to compute the discounted cash flows for the bond price at the end of the period, here we assume y_e=r_e+s_e=5.3% with r_e=4.1% and s_e=1.2%. Notice that the yield increment

$$\Delta y = y_e - y_b = 5.3\% - 5.0\% = 0.3\%\,,$$

can be viewed as the sum of the increase in the interest rate of

$$\Delta r = r_e - r_b = 4.1\% - 4.0\% = 0.1\%\,,$$

and that in the credit-spread:

$$\Delta s = s_e - s_b = 1.2\% - 1.0\% = 0.2\%\,.$$

In the following we work out the discounted cash flows, computed at the end of the period,

T_i	r_i	s_i	Cash	Disc.
1	4.1%	1.2%	10	9.497
2	4.1%	1.2%	10	9.019
3	4.1%	1.2%	10	8.565
4	4.1%	1.2%	110	89.470
		Sum	140	116.550

so that we obtain an ending price

$$P_e = P(s^e, r^e) = P(4.1\%, 1.2\%) = 116.550\,.$$

Notice that if we had computed the return components using the cash-flow method, we would have obtained the yield contributions to the return but we could have not split these return components into an interest-rate and a credit-spread part. On the other hand, when we use the simulated-price method we choose to *affect* only one variable influencing the bond price at a time. For example to compute the credit component, we *simulate* the bond price that would occur if only the credit spread were to change from s_b to s_e, while leaving the interest rates to r_b:

$$\begin{aligned} C^s &= P(r_b, s_e) - P(r_b, s_b) \\ &= 116.942 - 117.730 = -0.788\,. \end{aligned}$$

We do not move the credit spread from s_b to s_e for each single maturity separately because we are just interested in the credit-spread contribution as a whole.

To compute the interest-rate components we can proceed similarly by displacing the interest rate used in computing the discounted cash flows separately for each maturity T_i, with i=1, 2, 3, 4. For the first maturity, for example, the simulated price $P(\{r^e, r^b, r^b, r^b\}, s_b)$ can be computed using the following table:

T_i	r_i	s_i	Cash	Disc.
1	4.1%	1.0%	10	9.515
2	4.0%	1.0%	10	9.070
3	4.0%	1.0%	10	8.638
4	4.0%	1.0%	110	90.497
		Sum	140	117.721

so that, in this way, we obtain

$$\begin{aligned} C^r_1 &= P(\{r^e, r^b, r^b, r^b\}, s^b) - P(r^b, s^b) \\ &= 117.721 - 117.730 = -0.009\,. \end{aligned}$$

Similarly for the other maturity components we have

$$\begin{aligned} C^r_2 &= P(\{r^b, r^e, r^b, r^b\}, s^b) - P(r^b, s^b) = -0.017\,, \\ C^r_3 &= P(\{r^b, r^b, r^e, r^b\}, s^b) - P(r^b, s^b) = -0.025\,, \\ C^r_4 &= P(\{r^b, r^b, r^b, r^e\}, s^b) - P(r^b, s^b) = -0.344\,. \end{aligned}$$

Notice that the sum of the credit-spread component C^s and the maturity-dependent interest-rate components C^r_i's is not exactly equal to the cash profit ΔP:

$$C^s + C^r_1 + C^r_2 + C^r_3 + C^r_4 = -1.183 \neq -1.179\,.$$

The difference between ΔP and the component sum, namely

$$C^{r*s} = \Delta P - [C^s + C^r_1 + C^r_2 + C^r_3 + C^r_4] = 0.004\,, \quad (4)$$

can be attributed to the compounding effects of moving simultaneously both interest rates and credit spreads (more on this cross term later).

Using the results obtained so far, we can write the cash profit in terms of its contributions as

$$\Delta P = C^s + C^r_1 + C^r_2 + C^r_3 + C^r_4 + C^{r*s}\,, \quad (5)$$

so that it is clearly expressed as the sum of the credit-spread part, the components coming from each coupon maturity, and the cross contribution born from the compounding of the interest rates and the credit spread.

3 Return contributions of fixed-income securities

3.1 Interest-rate nodes and key rates

In the numerical example of the previous session we compute the return contributions for a fixed-rate coupon bond depending on an interest-rate term structure and a credit spread. In this section we generalize those results to any fixed-income security depending on an interest-rate curve, the credit spread, and possibly on some other financial variables.

Interest-rate nodes According to the result described by equation (5) we can split the sample-bond profit into six components: four coming from the cash-flow maturities, one from the credit spread, and the last one from the interaction of the two. That result is only valid for the sample fixed-rate coupon bond examined in the previous section. Another financial security will have in general a different cash-flow structure and hence we might obtain a different number of contributions to the profit ΔP (and hence to the price-variation return R). In the limiting case of a security that pays, or receives, a cash flow each day until maturity, in principle we could compute a return contribution for each of those days.

The method described in the previous section, therefore, in practice is not applicable to a real-world portfolio where cash flows are observed almost every day. Indeed we would need to compute a simulated price for each day until the portfolio maturity. Such a task is usually neither feasible nor desirable. It is not feasible because, very often, the pricing function that allows us to obtain the security value is computationally expensive; it is not desirable because how would we analise all the contributions coming for each day as their number could possibly be in the hundreds or even in the thousands? Therefore instead of considering all the future cash-flow dates one by one, we define a fine grid of *interest-rate nodes*.

In table 1 we show as an example a possible choice for the interest-rate nodes. Given the choice of N nodes with maturities T^n_1, ..., T^n_N, in terms of notation, we denote with $\mathcal{T}=\{T^n_1, \ldots, T^n_N\}$ the ordered set of interest-rate nodes. The choice of the interest-rate nodes should be consistent with the choice of nodes used to bootstrap the interest-rate curve. In other words the bootstrap nodes should always be among the interest-rate nodes.

We consider now a generic fixed-income security with a price that depends on the N interest rates $r_1, \ldots, r_N$, the credit spread s, and, possibly, another vector of financial variables $\mathbf{x}$. The security price can then be computed formally as

$$P = P(\mathbf{r}, s, \mathbf{x})\,, \tag{6}$$

where we denote with $\mathbf{r}$ the vector of interest rates $r_1, \ldots, r_N$ at the nodes $\mathcal{T}$ and with $\mathbf{x} = x_1, x_2, \ldots$, the other financial variables, also known as the other risk drivers, affecting the price.

Key rates The choice of the interest-rate nodes should be made once for all and it should be the same for all currencies for which we have an interest-rate term structure. When dealing with different groups of securities, depending on the task to be accomplished, the node grid might be too fine grained. Indeed it is customary to analyse the fixed-income contributions on a coarser grid. We therefore introduce a number K of key rates by grouping together neighbouring nodes.

For example, using the nodes described in table 1, we could create five key rates by grouping together maturities that belong to the same row:

$$\begin{aligned}
K^R_1 &= \{T^n_{1d}, T^n_{2d}, T^n_{1w}, T^n_{1m}, T^n_{2m}, T^n_{3m}\}\,,\\
K^R_2 &= \{T^n_{6m}, T^n_{9m}, T^n_{1y}, T^n_{1y6m}, T^n_{2y}, T^n_{3y}\}\,,\\
K^R_3 &= \{T^n_{4y}, T^n_{5y}, T^n_{6y}, T^n_{7y}, T^n_{8y}, T^n_{9y}\}\,,\\
K^R_4 &= \{T^n_{10y}, T^n_{11y}, T^n_{12y}, T^n_{13y}, T^n_{14y}, T^n_{15y}\}\,,\\
K^R_5 &= \{T^n_{20y}, T^n_{25y}, T^n_{30y}, T^n_{40y}, T^n_{50y}, T^n_{60y}\}\,.
\end{aligned}$$

We can interpret K^R_1 as the very-short maturity rates, K^R_2 as the short-maturity rate, K^R_3 as the intermediate rates, K^R_4 as the long-term ones, and finally K^R_5 as the very-long-term rates.

Other grouping of interest-rate nodes are also possible, depending on the task at hand. For example when dealing with a fixed-income attribution in which there are two desks, one in charge of bonds with maturities

1 day	2 day	1 week	1 month	2 months	3 months
6 months	9 months	1 year	18 months	2 years	3 years
4 years	5 years	6 years	7 years	8 years	9 years
10 years	11 years	12 years	13 years	14 years	15 years
20 years	25 years	30 years	40 years	50 years	60 years

Table 1: Possible list of maturities for interest-rate nodes. Each node affects the interest rates at the given maturity and at all the maturities until the next node. The last node influences all maturities greater than 60 years. Note that in this case N=30.

less than, for example, five years and another for securities with longer maturities, we might just have two key rates:

$$\begin{aligned} \mathbb{K}^R_1 &= \{T^n_{1d}, T^n_{2d}, \ldots, T^n_{4y}, T^n_{5y}\}\,, \\ \mathbb{K}^R_2 &= \{T^n_{6y}, T^n_{7y}, \ldots, T^n_{50y}, T^n_{60y}\}\,. \end{aligned}$$

In this case K=2, $\mathbb{K}^R_1$ is the key rate associated to the short maturities, while the $\mathbb{K}^R_2$ with the long ones.

3.2 Computation of key-rate contributions

Let us apply the simulated-price method to a generic security with a pricing function described by equation (6). Let's assume the risk drivers, i.e. the financial variable affecting the security price, to be $(\mathbf{r}^b, s^b, \mathbf{x}^b)$ at the beginning of the period and $(\mathbf{r}^e, s^e, \mathbf{x}^e)$ at the end. We assume here that t_b and t_e belong to the same day so that we are not bothered with the effect of different evaluation times or coupon payments. In subsection 3.4 we relax this hypothesis by introducing the carry component.

Credit-spread component We compute the credit-spread component by simulating the price of a security in which only the spread varies from s^b to s^e, however leaving the interest rates and the other variables unchanged to $\mathbf{r}^b$ and $\mathbf{x}^b$:

$$C^s = P(\mathbf{r}^b, s^e, \mathbf{x}^b) - P(\mathbf{r}^b, s^b, \mathbf{x}^b)\,.$$

In subsection 3.3 we provide explicit results for the credit-spread contribution C^s in the case of a zero-coupon bond.

Interest-rate node contributions For each interest-rate node T^n_i, with i=1, . . . , N, we compute an interest-rate component C^r_i by simulating the price that we would obtain if all the interest rates r_j, with $j \neq i$, were observed at the beginning of the period, however with r_i evaluated at the end:

$$C^r_i = P(\{r^b_1, \ldots, r^e_i, \ldots, r^b_N\}, s^b, \mathbf{x}^b) - P(\mathbf{r}^b, s^b, \mathbf{x}^b)\,.$$

Note that in computing the node contributions C^r_i's we are assuming that both the credit spread and the other variables are observed at the beginning of the period. Again, the explicit analytical computation of the contributions C^r_i's is performed in subsection 3.3 for the case of a zero-coupon bond.

Spread/rate cross contribution As we noted earlier, there is a compound effect arising from varying at the same time the credit spread and the interest rates. This compounding effect can be computed by simulating the increase of the security price that is due to a concurrent change of both the interest rates and the credit spread but that it is not due by either effects alone. Hence, similarly to equation (4), we define

$$C^{r*s} = P(\mathbf{r}^e, s^e, \mathbf{x}^b) - P(\mathbf{r}^b, s^b, \mathbf{x}^b) - C^s - \sum_{i=1}^{N} C^r_i\,.$$

In principle this contribution could be further split into the components arising from the variation of each interest-rate node, however in practice such a detailed contribution is not needed. Once more, see subsection

3.3 for the analytical computation of this contribution in the case of a zero-coupon bond.

Other risk-driver components Notice that in the computation of the terms C^s, C^r_i, and C^{r*s} the other risk drivers $\mathbf{x}$ are always evaluated at the beginning of the period. For this reason the sum of the above contributions does not, in general, matches the price-variation return R. We therefore introduce the *other risk-driver component* as the contribution that is missing to obtain the return R:

$$C^x = R - C^s - \sum_{i=1}^{N} C^r_i - C^{r*s} .$$

This expression has the following interpretation: the return that is neither due to a variation of the credit spread alone, nor to the increase of each interest rate by itself, nor to the compounding of the two, must be due to the variation of the other risk drivers.

Gathering all the terms together we write the cash price variation as

$$\Delta P = C^s + \sum_{i=1}^{N} C^r_i + C^{r*s} + C^x . \tag{7}$$

This equation states that the cash profit can be split into N+3 basic risk-driver contributions that can be computed for any security with a price given by equation (6).

Key-rate contributions As mentioned earlier, one usually introduces K key rates in order to provide a practical fixed-income attribution to the interest-rate components. Therefore we compute the key-rate contributions as the sum, for each key rate, of the contributions of the corresponding nodes. For each key rate j=1, . . . , K we define

$$C^k_j = \sum_{i \in \mathcal{K}_j} C^r_i ,$$

so that the cash return decomposition can be written as

$$\Delta P = C^s + \sum_{j=1}^{K} C^k_j + C^{r*s} + C^x . \tag{8}$$

This equation states that the price-variation return can be written as the sum of the spread contribution C^s, the components coming from each key rate C^k_j, for j=1, . . . , K, the cross spread-interest term C^{r*s}, and finally the return C^x coming from all the other risk drivers. This equation is one of the main results of the current paper.

3.3 Explicit results for coupon bonds

In this subsection we provide explicit analytic results for some common bonds computed with a yearly-compounded discount factor.

To compute the return decomposition described by equation (8) we need the security pricing function $P(\mathbf{r}, s, \mathbf{x})$ which, most of the times, is obtained with a quantitative model and is computed on some numerical method calibrated on market data. In the case of a coupon bond priced using cash flow discounting, however, we can even compute explicitly the return contributions.

Hence consider a bond paying M coupons $C_i(\mathbf{x})$, for i=1, . . . , M, depending on a risk-driver vector $\mathbf{x}$ that contains neither the credit spread nor the interest rates. Given a discount factor $D(T; \mathbf{r}, s)$ for a cash flow paid at a certain future date T, we can write the bond price as

$$P(\mathbf{r}, s, \mathbf{x}) = \sum_{i=1}^{M} D(T_i; \mathbf{r}, s) C_i(\mathbf{x}) ,$$

where T_i , for i=1, . . . , M, is the maturity date of the i-th coupon.

Since in the above expression all the dependency on the interest rates and the credit spread is embedded in the discount factor, in computing the spread and the interest-rate contributions we can just consider the special case of a zero-coupon bond. It is then immediate to show that the return contributions are additive in the coupons. Therefore we can just focus on computing the components for a zero-coupon bond with a fixed maturity T, with an observed interest rate r and a credit spread s.

Contributions for zero-coupon bonds When annual compounding is used to compute the cash flows, assuming C=1, we can write the zero-coupon-bond pricing function $P_z(r, s; T)$ as

$$P_z(r, s; T) = \frac{1}{(1 + r + s)^T} .$$

Assuming that from the beginning of the period to its end the interest rate changes from r_b to r_e and the credit spread from s_b to s_e, we are interested in computing the contributions to the cash gain

$$\begin{aligned} \Delta P_z &= P_z(r_e, s_e; T) - P_z(r_b, s_b; T) \\ &= \frac{1}{(1 + y_e)^T} - \frac{1}{(1 + y_b)^T} , \end{aligned}$$

where $y_b = r_b + s_b$ is the yield at the beginning of the period and $y_e = r_e + s_e$ that at the end. In the following derivations it is useful to express the above price variation as

$$\Delta P_z = \frac{1}{(1 + y_b)^T} \left[\left(1 - \frac{\Delta y}{1 + y_b + \Delta y} \right)^T - 1 \right] ,$$

where Δy=$y_e - y_b$ is the yield displacement across the period.

Spread and interest-rate contributions Using the simulated-price method we compute the credit-spread component of the cash return for the zero-coupon bond as

$$\begin{aligned} C_z^s &= P_z(r_b, s_e) - P_z(r_b, s_b) \\ &= \frac{1}{(1 + y_b)^T} \left[\left(1 - \frac{\Delta s}{1 + y_b + \Delta s} \right)^T - 1 \right] , \end{aligned}$$

where $\Delta s = s_e - s_b$ is the credit-spread increase across the period. Similarly we proceed with the computation of the interest-rate contribution as

$$\begin{aligned} C_z^r &= P_z(r_e, s_b) - P_z(r_b, s_b) \\ &= \frac{1}{(1 + y_b)^T} \left[\left(1 - \frac{\Delta r}{1 + y_b + \Delta r} \right)^T - 1 \right] , \end{aligned}$$

where $\Delta r = r_e - r_b$ is the interest-rate displacement from the beginning of the period to its end.

Compound contribution of interests and credit As already noted earlier, the two contributions for C_z^r and C_z^s do not add up to ΔP_z because a compounding term is present. In the case of the zero-coupon bond we can compute this effect explicitly as

$$C_z^{s*r} = \Delta P_z - C_z^r - C_z^s ,$$

so that we have

$$\begin{aligned} C_z^{s*r} = \frac{1}{(1 + y_b)^T} \Bigg[&\left(1 - \frac{\Delta s + \Delta r}{1 + y_b + \Delta s + \Delta r} \right)^T + 1 \\ &- \left(1 - \frac{\Delta r}{1 + y_b + \Delta r} \right)^T - \left(1 - \frac{\Delta s}{1 + y_b + \Delta s} \right)^T \Bigg] . \end{aligned}$$

This term is typically small for short maturities but can be of significance when the displacements are important and for longer maturities.

In table 2 we provide a computation of the return components for a zero-coupon bond at different maturities for an interest-rate curve observed on the market at the end of year 2015. The displacements of the interest rate and that of the credit spread used to obtain the values of table 2 are compatible with values that could be observed on the market during relatively stressed market conditions.

We notice how for short maturities the credit-spread component is dominant, while for longer maturities the interest-rate component assumes a more significant role. Also observe how the cross contribution C_z^{r*s} is typically small and negligible unless the bond maturity is prominent, in which case it reaches 9 basis points. Also we notice that while the interest-rate and the credit spread contributions are negative, their compounding contribution is positive. One should be aware that, as small as the cross contribution C_z^{r*s} might be, when we need to explain every fraction of basis point of the price-variation return R, we need to include this term among the contributions. Furthermore, recall that these are contributions computed for a period of a single day and that, in general, they need to be compounded with the corresponding components coming from several other days. Therefore the compounding of the cross term over a long period of time, might generate significant contributions.

T	r_b	Δr	ΔP_z	C_z^r	C_z^s	C_z^{r*s}
1	0.30%	0.03%	-0.272%	-0.029%	-0.243%	0.000%
2	0.50%	0.05%	-0.571%	-0.095%	-0.476%	0.000%
5	0.70%	0.07%	-1.433%	-0.315%	-1.121%	0.003%
10	1.40%	0.09%	-2.572%	-0.688%	-1.899%	0.015%
15	1.90%	0.10%	-3.234%	-0.940%	-2.326%	0.031%
20	2.00%	0.10%	-3.632%	-1.062%	-2.618%	0.048%
25	2.10%	0.11%	-3.890%	-1.223%	-2.736%	0.069%
30	2.20%	0.12%	-3.957%	-1.329%	-2.718%	0.090%

Table 2: Exact contributions for a zero-coupon bond computed using yearly compounding at different maturities. In all cases we have set s_b=1% and Δs=0.25%.

3.4 The carry return contribution

So far we computed the security performance contributions due to changes in the interest rates and the credit spread, assuming that there was no variation of the cash-flow times to maturity. In practice however we are typically interested in periods that span one or more days, so that a change in the time to maturity is significant.

For example consider the zero-coupon bond price as a function of the current time t, the yield y, and the maturity T:

$$P_z(y;t,T) = \frac{1}{(1+y)^{(T-t)}}\,,$$

and notice the dependence from the $T-t$ term. Even when the yield does not change from the beginning of the period t_b to its end t_e, the bond price has a return of

$$\begin{aligned} R_c &= \frac{P_z(y;t_e,T) - P_z(y;t_b,T)}{P_z(y;t_b,T)} \\ &= \frac{(1+y)^{(T-t_b)}}{(1+y)^{(T-t_e)}} - 1 = (1+y)^{\Delta t} - 1\,, \end{aligned} \quad (9)$$

where Δt=$t_e - t_b$ is the elapsed time.

The contribution to the return due solely from the passage of time is sometimes called *carry return* or the *calendar component* of return. Notice that in the previous example, when y is infinitesimally small we have

$$R_c = y \cdot \Delta t\,,$$

so that the carry return is proportional to the yield. However this relationship is not true in general.

Carry return for a generic security In the pricing function for a generic security given by equation (6), we do not consider the explicit dependence on the current time t. In order to obtain the calendar return component we now study the effect of an explicit time dependence so that we write the bond price at time t as

$$P = P(\mathbf{r}, s, \mathbf{x}; t)\,. \quad (10)$$

In considering the price variation from the beginning of the period to its end we are interested in extracting the carry component from the cash return:

$$\Delta P^{(t)} = P(\mathbf{r}^e, s^e, \mathbf{x}^e; t_e) - P(\mathbf{r}^b, s^b, \mathbf{x}^b; t_b)\,. \quad (11)$$

We can easily accomplish this task by adding and subtracting the term $P(\mathbf{r}^b, s^b, \mathbf{x}^b; t_e)$ to the right-hand side of the previous equation to obtain:

$$\Delta P^{(t)} = \Delta P^{(r,s,x)} + C^{\text{carry}}\,,$$

where

$$\Delta P^{(r,s,x)} = P(\mathbf{r}^e, s^e, \mathbf{x}^e; t_e) - P(\mathbf{r}^b, s^b, \mathbf{x}^b; t_e)$$

and

$$C^{\text{carry}} = P(\mathbf{r}^b, s^b, \mathbf{x}^b; t_e) - P(\mathbf{r}^b, s^b, \mathbf{x}^b; t_b)\,.$$

The carry term C^{carry} has the meaning of the cash difference accumulated across the period that can be attributed to the passage of time alone. Indeed in computing this term we are not interested at all at the value of the risk drivers $\mathbf{r}$, s, and $\mathbf{x}$ at the end of the period, but only at the beginning. Also note that the computation of the carry term can be performed at time t_b, since no variables at t_e are needed.

The risk-driver component $\Delta P^{(r,s,x)}$, on the other hand, provides the increase in value at time t_e that is caused by a change of the risk drivers from $\mathbf{r}^b$, s^b, and $\mathbf{x}^b$, to, respectively, $\mathbf{r}^e$, s^e, and $\mathbf{x}^e$, without observing a change in the cash-flow time to maturity (we are keeping a constant t=t_e). Since in the description of the simulated-price method we always assumed the variation in maturity to be negligible, we can use those results, see equation (8), to split $\Delta P^{(r,s,x)}$ into additive components.

Coupon and convergence terms Very often, in trading securities that pay pre-determined coupons, we write the pricing function as the sum of a clean price $\hat{P}$ and a coupon accrual proportional to a constant A:

$$P = \hat{P}(\mathbf{r}, s, \mathbf{x}; t) + A \cdot (t - t_a) \ , \qquad (12)$$

where t_a is the beginning of the coupon accrual. In this case the carry return,

$$\begin{aligned} C^{\text{carry}} &= A \cdot (t_e - t_b) \\ &\quad + \hat{P}(\mathbf{r}^b, s^b, \mathbf{x}^b; t_e) - \hat{P}(\mathbf{r}^b, s^b, \mathbf{x}^b; t_b) \, , \end{aligned}$$

can be written as the sum of coupon term

$$C^{\text{cpn}} = A \cdot (t_e - t_b) \ ,$$

proportional to A, and a convergence term

$$C^{\text{cnv}} = \hat{P}(\mathbf{r}^b, s^b, \mathbf{x}^b; t_e) - \hat{P}(\mathbf{r}^b, s^b, \mathbf{x}^b; t_b) \, ,$$

depending only from the variation of the clean price across the period.

Observe how the coupon term does not contain any dependence on the risk drivers since its value was already set at time t_a. For a zero-coupon bond, by definition, A=0 so that C^{cpn}=0 and the convergence term can be explicitly computed as in equation (9). The convergence term provides the bond cash increase that is due to a variation of the risk drivers $\mathbf{r}$, s, and $\mathbf{x}$, however that is not already embedded into the coupon.

4 Model extensions

In the previous section we show that, as long as we can compute the security price using equation (10), we can split the cash profit as

$$\Delta P = C^s + \sum_{j=1}^{K} C_j^k + C^{r*s} + C^x + C^{\text{cnv}} + C^{\text{cpn}} \, ,$$

where

- C^s is credit-spread component,
- for each key rate $\mathrm{K}_j^{\mathrm{R}}$ with j=1, ..., K, C_j^k is the j-th key-rate contribution;
- C^{r*s} is the cross interest-rate/credit-spread component,
- C^x is the term due to the presence of other risk drivers,
- C^{cnv} is the carry convergence term,
- and, finally, C^{cpn} is the coupon term, present only when for securities that fix coupons.

Note that when a coupon is paid during the observation period, it's value should be included in the term C^{cpn}.

As it is usually the case in fixed-income attribution, we are interested in computing the contributions of the price-variation return R. In this case we have

$$R = c^s + \sum_{j=1}^{K} c_j^k + c^{r*s} + c^x + c^{\text{cnv}} + c^{\text{cpn}} \, , \qquad (13)$$

where the price-variation contributions and can be obtained from the cash one by dividing by the price at the beginning of the period P_b:

$$c^{\text{any}} = \frac{C^{\text{any}}}{P_b} \, .$$

The return contributions have a similar interpretation to their cash counterparts.

Results for multiple interest-rate curves or spreads The result described in equation (13) was obtained under the assumption that the security price would depend only on a single interest-rate curve and one credit spread. In practice however, sometimes we need to compute the contributions for securities depending on multiple interest-rate curves and/or more then one credit spread. Let us consider, for example, the case of a security that depends on two credit spreads s_1 and s_2 so that the price can be written as

$$P = P(\mathbf{r}, s_1, s_2, \mathbf{x}; t)\,.$$

In order to compute the different return contributions we now have to simulate separately the two credit spreads. It can be shown that, see reference [3] for more details, the return decomposition can be written as

$$\begin{aligned} R \,=\, & c_1^s + c_2^s + \sum_{j=1}^{K} c_j^k \\ & + c^{r*s_1} + c^{r*s_2} + c^{s_1*s_2} + c^{r*s_1*s_2} \\ & + c^x + c^{\mathrm{cnv}} + c^{\mathrm{cpn}}\,, \end{aligned}$$

where c_1^s and c_2^s are the contributions coming separately from each credit spread, the terms c_j^k's are the key-rate contributions, c^{r*s_1}, c^{r*s_2}, $c^{s_1*s_2}$ are the components obtained by displacing two risk drivers at the same time; finally, $c^{r*s_1*s_2}$ is the piece obtained by displacing at the same time both credit spreads and the interest rates but that was not already accounted for in the other terms. Very often the term $c^{r*s_1*s_2}$, being typically very small even for large maturities, is included in the c^x component. Also note that, in the above expression, c^x, c^{cnv}, and c^{cpn} denote the same contributions as shown in equation (13).

Multiple interest-rate term structures There are securities, such as a cross-currency swap with a fair value that depends on more than one interest-rate term structure. In this case, using the simulation method described earlier, we obtain a return decomposition with contributions from the key rates coming from all interest-rate curves involved, plus the appropriate cross terms.

Aggregation of results at portfolio level The results obtained in this paper are applicable to a single instrument held in a portfolio, or even in a benchmark. Depending on the chosen performance-attribution process the single-asset results are typically aggregated at different portfolio levels. The details of the aggregation procedure may vary from case to case, however since the single-period contributions are additive there is a great flexibility during this phase. In the aggregation process the key-rate contributions computed for different securities can be summed up to provide the portfolio-level key-rate components.

Model and quoted price The simulated-price method described in this paper is based on the assumption that we have available the security pricing function so that we can write the security quote as in equation (10). We call the value thus obtained the *model price* because it is obtained using a quantitative model. In practice, because of the supply and demand dynamics, the market purchase and sale of securities happens at a different price: the *quoted price*. More research is needed to develop a method to compute the performance contributions on quoted prices.

Further split of the risk-driver component As a final remark, we notice that, using the simulation method, the contribution from the other risk drivers c^x can itself be split into sub-components depending on what exactly are the variables $\mathbf{x}$. For example for a convertible bond we may split c^x into an equity component, a volatility component and, possibly, a bunch of other components coming from the interaction of multiple variables at the same time. See reference [3] for more details on the subject.

Conclusion In conclusion we are able to split of the price-variation return into components because we know how the security price depends from the different risk drivers affecting its value. We believe that as simple as it is, the price-simulation method described in this paper is very powerful and an helpful tool in fixed-income attribution.

References

[1] Carl Bacon. *Practical Portfolio Performance Measurement and Attribution*. Wiley Finance, second edition, 2008. 2

[2] Andrew Colin. *Fixed Income Attribution*. Wiley Finance, 2005. 2

[3] Marco Marchioro. Performance risk-driver contributions of non-linear portfolio. In *Risk and Performance Attribution*, Quantitative Research Series. StatPro website, June 2017. Previews soon to be available at marchioro.org/papers/. 14