OpenGamma Quantitative Research

Inflation: Instruments and curve construction

Arroub Zine-eddine
arroub@opengamma.com

OpenGamma Quantitative Research n. 19 | First version: April 2013, this version: January 2014.

Abstract

The most common inflation linear instruments and the curve construction are described in this note.

Contents

1 Introduction

In order to price and value any kind of inflation-linked cash flows, we first need to have an accurate inflation price index curve. Firstly, we will describe inflation indices and how they are used, then we will take a look at the most common inflation linear instruments and finally we will show how we can take observable market inputs and combine them with a seasonality vector to construct an appropriate inflation curve.
Inflation-linked options are not treated in this note; for a description and a pricing methodology of the most common inflation options see Zine-eddine (2013b).

2 Price index

In economics, inflation can be define in different ways but the most common definition is a rise of the general level of prices of goods and services in an economy over a period of time. For derivatives purposes, we use a slightly different definition: inflation is defined in terms of the percentage increments of a reference price index, which is the value of a representative basket of goods and services. The second definition is obviously supposed to define a proxy for the first definition.
Inflation derivatives pay cash-flows linked to widely available price indices. The standard price indices are the European HICP x, the US CPI-U and the UK RPI.

Price indices are published by national and/or international statistics services on a monthly basis. Usually they are published two weeks after the month end. The price index value used for the swaps (or any other inflation derivatives) must be known at each trading date, so price indices are used with a certain lag. The most commonly used lag is three months.

All inflation-linked instruments use price index values. Given a date, there are, by convention, two ways to interpolate into a price index curve. We provide a description of both methods hereunder.
We call the reference index the value effectively used for the cash flow determination.

2.1 Monthly index

In some cases the reference index is the price index of a month, i.e.

$$\text{ReferenceIndex} = \text{PriceIndex}(m)$$

with m the reference month (first of month) linked to the payment date. Some inflation-linked instruments use reference index values linked to other dates than the payment date; the calculation is done in the same way.

For example, a financial instrument paying an inflation-linked cash flow the 17th of June (namely the payment date) with a lag of three months will use the index of March.

2.2 Interpolated index

In other cases the reference index is linearly interpolated between two months. The interpolation is done with the number of days of the *payment* month (not the reference months). The reference

[0] First version: April 2013; this version: June 2013.

index is given by

$$\begin{aligned} \text{ReferenceIndex} &= \text{PriceIndex}(m_1) + \left(\frac{d-1}{D}\right)(\text{PriceIndex}(m_2) - \text{PriceIndex}(m_1)) \\ &= \alpha\text{PriceIndex}_1 + (1-\alpha)\text{PriceIndex}_2. \end{aligned}$$

where d is the day of the month of the payment date and D is the number of calendar days in the payment month. The reference dates (m_1 and m_2) are the first of two consecutive months.

For example, a financial instrument paying an inflation-linked cash flow the 17th of June (namely the payment date) with a lag of three months will use an interpolated value between the index of March and the index of April.

2.3 Notation for index

For the sake of clarity and to avoid repetition, we will use the following notation $I(t)$ for the reference price index of an inflation-linked flow paid at date t using the natural lag (by natural we mean the one used by market convention). For example if t is the 3rd of December 2013, the natural lag is 3 months and the inflation-linked flow pay a monthly (non interpolated) price index then

$$I(t) = \text{PriceIndex}(September 2013)$$

We will still call $I(t)$ the price index. This is an abuse of notation, this kind of notation is widely (implicitly) used in articles about inflation derivatives.

We give hereunder some examples of lag and interpolation conventions for swaps and bonds.

Area	Index	lag	Interpolation	Liquidity
Eurozone	HICPxT	Three months	Monthly	Yes
France	French CPIxT	Three months	Linear	Yes
UK	UK RPI	Two months	Monthly	Yes
US	HICPxT	Three months	Linear	Moderate
Canada	HICPxT	Three months	Monthly	Moderate
Australia	HICPxT	Three months	Monthly	Moderate
Italy	HICPxT	Three months	Monthly	No
Spain	HICPxT	Three months	Monthly	No

Table 1: Conventions for zero-coupon inflation swaps

Issuer	Index	lag	Interpolation	Liquidity
France OATi	French CPIxT	Three months	Linear	Yes
France OATei	HICPxT	Three months	Linear	Yes
Germany BUndei	HICPxT	Three months	Linear	Yes
Italy BTPSei	HICPxT	Three months	Linear	Yes
UK old	UK RPI	Height months	Monthly	Yes
UK new(since 2005)	UK RPI	Three months	Linear	Yes

Table 2: Conventions for inflation-linked bonds

3 Swaps description

3.1 Zero-coupon swap

A zero-coupon swap is the exchange of two flows on one given date t_p which is a certain number of years n after the start date t_s on a reference notional N. One flow is a fixed amount. The amount is quoted through a compounded annual rate. For a rate R the fix amount paid is

$$N((1+R)^n - X)$$

with $X = 0$ if the notional is exchanged and 1 if it is not exchanged. No exchange of notional is the standard in exchange markets.

The inflation flow is given by the change of the price index. The amount paid is

$$N\left(\frac{I(t_s)}{I(t_p)} - X\right).$$

The reference indices are linked to the coupon dates t_s and t_p. The usual linked is a *fixing lag* which is several months.

3.2 Year-on-year swap

Given a set of dates $(t_i)_{i=1\dots n}$, a year-on-year inflation swap is the exchange of two flows on each date t_i on a reference notional N. One flow is a fixed amount. The fix amount paid at time t_i is

$$NR_i$$

Where $(R_i)_{i=1\dots n}$ are rates fixed at the inception of the swap. Usually for standard contracts this rate is constant with respect to i. The inflation flow is given by the change of the price index. The amount paid at time t_i is

$$N\left(\frac{I(t_i)}{I(t_{i-1})}\right).$$

The reference indices are linked to the coupon dates t_i and t_{i-1}. The usual linked is a *fixing lag* which is several months. A period of 1 year between each payment date t_i is the standard in the derivative market; there are also some swaps with a period of 6 months or 3 months but this practice is marginal.

3.3 Revenue swap

A revenue swap is a series of zero-coupon swaps, a period of 1 year between each payment is a standard in the derivative market. This swap could also be called annuity inflation swap.

3.4 OATI swap

The purpose of OATI is to reproduce the pay-off of the French governmental bonds (namely the OATI). So the inflation-linked leg of this swap is the same as the one of the revenue swap. However, the other leg could pay the following fixed amount at each time t_i

$$NR_i$$

or it could also pay a libor leg. This kind of swap is used in asset swap structure.

3.5 Limited price indexation inflation swap

In standard definition this swap is only used linked to the UK RPI (of course, in OTC markets the same kind of structure can exist linked to other indices). The structure of this is the same as the revenue swap but each inflation indexed flow is capped and floored (usually a floor at 0% and a cap at 5%).

So, in the case of floor at 0% and a cap at 5% the inflation-linked amount paid at time t_i is

$$N\left(\min\left(\max\left(\frac{I(t_i)}{I(t_0)} - 1, 0\%\right), 5\%\right)\right).$$

Usually, the revenue swap part and the optional part are priced separately; in this document we only focus on the linear part.

4 Bonds description

The two most common forms of inflation bonds are the *Capital Indexed Bond* and the *Interest Indexed Bond*.

Let's denote t_s the start date, ie the date when the bond is issued, t_p the last payment date, ie when the nominal is reimbursed and $(t_i)_{i=1\ldots n}$ the payment dates of each coupon.

4.1 Capital index bond

These bonds have a fixed real coupon and real notional. The payments are the real coupons and notional multiplied by the ratio of reference indices.

The amounts paid for a notional N and a real coupon C are, at each coupon date t_i,

$$CN\frac{I(t_i)}{I(t_s)}$$

and at maturity

$$N\max\left(\frac{I(t_p)}{I(t_s)}, 1\right).$$

In practice, the floor at strike 0 is often omitted for pricing purposes but the actual payoff did contain a floor. The inflation bonds issued by France, UK and US all are in this category.

Each coupon and the notional are (in pricing terms) like the inflation leg of a zero-coupon swap with notional exchange.

4.2 Interest Indexed Bond

At each coupon date, the amount paid is the real coupon plus the indexation of the notional over the coupon period, i.e.

$$N\left(C+\frac{I(t_i)}{I(t_{i-1})}\right).$$

The notional payment at maturity is not adjusted for inflation.

Each coupon is (in pricing terms) similar to the inflation leg of a year-on-year swap. The final coupon is a fixed amount.

5 Curves and Discounting

The pricing of simple inflation instruments is done in a similar way to the techniques used in interest rates. A standard discounting curve is used. The inflated cash-flows are priced using discounted estimated cash-flow with the estimation done with a *price index* curve. Like in the case interest rates, the price index curve is defined as the function for which the *estimate and discount* formulas hold.

The price index is a monthly index. So, only the monthly points make sense. The data for the month are stored on the first of month date.

5.1 Definition of the forward price index

We denote $\mathcal{I}(t,t_p)$ the forward price index value as seen from the date t, which pays at the date t_p the price index $I(t_p)$. Let $(\Omega,\mathcal{F},\mathbb{P})$ be a probability space.

Now let's assume there is a discount curve in the nominal economy denoted by $t\longmapsto P^D(t_0,t)$ (which is the "standard" discount curve), and a discount curve in the real economy denoted by $t\longmapsto P^D_r(t_0,t)$. So an arbitrage argument give us the following relationship

$$\mathcal{I}(t,t_p)=I(t)\frac{P^D_r(t,t_p)}{P^D(t,t_p)} \tag{1}$$

As $t_0\longmapsto I(t_0)P^D_r(t_0,t_p)$ is a tradable security (its value in $t=t_p$ is the pay off of the inflation-linked leg of a zero-coupon swap) then $\mathcal{I}(t,t_p)$ is a martingale under the t_p-forward measure $\mathbb{E}^{t_p}$, which mathematically means

$$\mathcal{I}(t,t_p)=\mathbb{E}^{t_p}\left[\mathcal{I}(t_p,t_p)|\mathcal{F}_t\right]=\mathbb{E}^{t_p}\left[I(t_p)|\mathcal{F}_t\right] \tag{2}$$

where $\mathcal{F}_t$ is the σ-algebra representing the information until time t (which is the known information at t). This last formulae 2 give us the relation between the index price and the forward index curve, it will be useful to price inflation-linked securities with zero-coupon kind of cash flows . The price in t of an inflated flow paying $I(t_p)/I(t_0) - X$ in t_p (where X is a constant) satisfies

$$P^D(t_0, t_p)\left(\frac{\mathcal{I}(t, t_p)}{I(t_0)} - X\right). \tag{3}$$

In order to price every kind of inflation-linked flows, we still have to price year-on-year cash flows.

5.2 Year-on-Year cash flow pricing

Using previous notation, a year-on-year cash flow is linearly linked to the following quantity

$$\frac{I(t_i)}{I(t_{i-1})}$$

So, in order to price this cash flow, we want to evaluate the following quantity $\mathbb{E}^{t_i}\left[\frac{I(t_i)}{I(t_{i-1})}|\mathcal{F}_t\right] = \mathbb{E}^{t_i}\left[\frac{\mathcal{I}(t_i,t_i)}{\mathcal{I}(t_i,t_{i-1})}|\mathcal{F}_t\right]$. Unfortunately $\frac{\mathcal{I}(t,t_i)}{\mathcal{I}(t,t_{i-1})}$ is not a martingale under the natural payment measure (the t_i-forward measure). One way to solve this problem is to express $\mathbb{E}^{t_i}\left[\frac{I(t_i)}{I(t_{i-1})}|\mathcal{F}_t\right]$ in function of $\frac{\mathcal{I}(t,t_i)}{\mathcal{I}(t,t_{i-1})}$ using a convexity adjustment, typically we use this kind of formula

$$\mathbb{E}^{t_i}\left[\frac{I(t_i)}{I(t_{i-1})}|\mathcal{F}_t\right] = e^{\texttt{Adjustment}}\frac{\mathcal{I}(t,t_i)}{\mathcal{I}(t,t_{i-1})}$$

This `Adjustment` is model dependent, in particular there is some closed formula in the (price index) market model. This closed formula uses several different market data such as price index volatility, interest rates bond forward volatility, price index cross correlation and correlation between rates and price index.
For a description and analytical formulas of this convexity adjustment see Zine-eddine (2013a).

5.3 Interpolation

So a forward price index curve is computed through a calibration with a set of inflation-linked linear instruments (usually zero-coupons). The calculation of forward index price values is useful but it is not complete since it does not provide us with values for intermediate months because market quotes are annual. That's why we use an interpolation method. When it comes to interpolation, there are numerous possible approaches; there is not a strong market consensus on this interpolation so there is no definitive right answer.
Some methods seem to be more used by the market and also perhaps preferable to others. We will see (in a following part of this note) a method which combines interpolation and seasonality.
One of the most used methods is perhaps the log linear interpolation, for example, we know the index price forward for a date m (which is a first of the month, so we use the notation m) and also for the same month one year later $m + 1year$, we can define the annual growth r

$$r = log\left(\mathcal{I}(t, m + 1year)\right) - log\left(\mathcal{I}(t, m)\right)$$

then for $j \in [0, 12]$,

$$\mathcal{I}(t, m + jmonth) = \mathcal{I}(t, m) * exp(\frac{j * r}{12})$$

This interpolation method can easily be generalised when the gap between forward price is more than a year. This interpolation method is widely used because the underlying assumption is that the growth rate between each month is constant.

6 Seasonal adjustment

If we want to construct a realistic price index curve, we need to overlay our curve with seasonality. Seasonality is a change of price patterns that occurs at a given time of the year. For example, if the raise of a price index is 3% over a year, that means, for example, that the price index is worth 100 in January 2012 and 103 in January 2013. But this increase of the price index is not regular and is subject to monthly variations. Reasons for this are social and cultural factors (holidays, sales periods...), legal measures (tax regime change...), weather, and so on.
Seasonality is measured using statistical methods The market consensus is to use the X-12 ARIMA software developed by the U.S. Census Bureau or an equivalent method also based on the ARIMA method, X-12 ARIMA is an autoregressive integrated moving average seasonal adjustment software for only univariate time series models. For a full description of this methodology and this software see `http://www.census.gov/srd/www/x12a/`.
This method will decompose the initial time series into three components which are the trend, the seasonal factor and the white noise. It's based on Moving Average filters used to extract the trend and the seasonality from the time series. The implicit assumption of this method is that these three components are stochastic. We will not describe the details of this statistical method but how we use the output of this method as input for our model.
This statistical method can generate inputs for different methods; we describe them hereunder.

6.1 Multiplicative case

The seasonal adjustment we describe here is multiplicative in two ways. The adjustment is a multiplicative factor on the non-adjusted price and the adjustments themselves are stored as a multiplicative factor from one month to the next.

A factor f_i is associated to each month. The factor represents the multiplicative seasonal adjustment from that month to the next. The first month is January. The seasonal adjustment is

$$\mathrm{Adj}(t, \text{Month } i+1) = f_i \mathrm{Adj}(t, \text{Month } i).$$

We notice that the seasonality adjustment is dependent of t, which is typically today. The reason for this is that the seasonality adjustment is a number (or a vector of numbers to be precise) computed using historical. The cumulative seasonal adjustment over the year is 1, i.e.

$$\prod_{i=1}^{12} f_i = 1.$$

The price indices are given by

$$\mathcal{I}(t, m) = \mathrm{Adj}(m) \mathcal{I}_{\text{NoAdj}}(t, m).$$

Where $\mathcal{I}_{\text{NoAdj}}$ is the forward price index without seasonality adjustment (obtained with the market data and the interpolation).

6.2 Additive case

The seasonal adjustment we describe here is additive in two ways. The adjustment is an additive factor on the non-adjusted price and the adjustments themselves are stored as an additive factor from one month to the next.

A factor f_i is associated to each month. The factor represents the additive seasonal adjustment from that month to the next. The first month is January. The seasonal adjustment is

$$\text{Adj}(t, \text{Month } i+1) = f_i + \text{Adj}(t, \text{Month } i).$$

The cumulative seasonal adjustment over the year is 1, i.e.

$$\sum_{i=1}^{12} f_i = 0.$$

The price indices are given by

$$\mathcal{I}(t, m) = \text{Adj}(m) + \mathcal{I}_{\text{NoAdj}}(t, m).$$

6.3 Seasonality and interpolation

This method is also known as pseudo additive. We will describe here a mixture method which combines interpolation and seasonality. Let $(f_i)_{(1 \leq i \leq 12)}$ be some pseudo additive seasonal factors. As before, we assume we know the index price forward for a date m (which is a first of the month, so we use the notation m) and also for the same month one year later $m + 1year$. We can therefore define the annual growth r by the solution of the following polynomial equation of degree 12

$$\mathcal{I}(t, m + 1year) = \mathcal{I}(t, m) \prod_{i=1}^{12} (1 + r + f_i)$$

then for $j \in [1, 12]$,

$$\mathcal{I}(t, m + jmonth) = \mathcal{I}(t, m) \prod_{i=1}^{j} (1 + r + f_i)$$

Notice that we are working here with non-seasonally adjusted forward price index, and the seasonal factors should satisfy $\sum_{i=1}^{12} f_i = 0$.

7 Example of the curve construction

To construct a price index curve, we firstly have to define a basket of calibration instruments. Any kind of inflation linear can be chosen but usually, the market practice is to use zero-coupon swaps because those instruments are more liquid. Usually available maturities on exchange market are 1Y, 2Y, 3y, 4Y, 5Y, 6Y, 7Y, 8Y, 9Y, 10Y, 12Y, 15Y, 20Y, 25Y, 30Y and sometimes there is market quotes for 40Y and 50Y. We define the kind of index. We then choose an interpolation method and a seasonal method or the mixture method (seasonal and interpolation). Then we use a solver to construct the curve. The figure 1 is a zoom on the period between 1Y and 7Y.

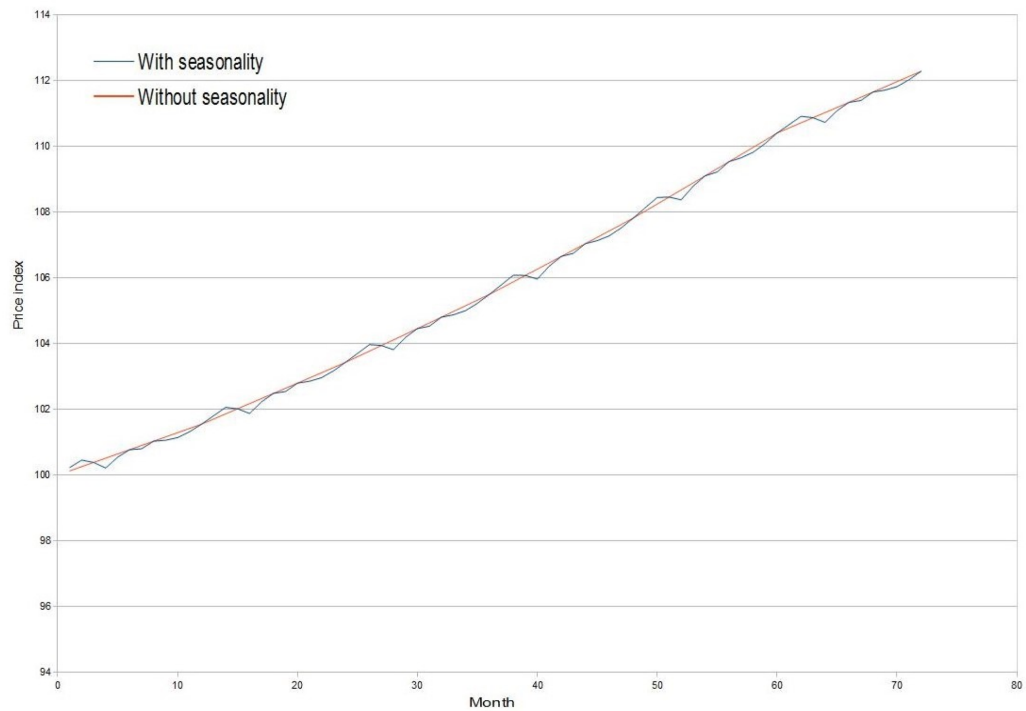

Figure 1: Price index forward curve until 7 years. With and without seasonality.

The choice of zero-coupon swaps to construct is due to several reasons

· it is the most liquid linear product,

· historically, zero-coupon came first,

· their structure is very simple which makes it easier to construct the curve with.

However, the implementation in the OpenGamma Analytics Library allows the user to use any linear inflation-linked instruments to construct his calibration basket.

One final word concerning performance, build 1000 price index curves (with seasonality and interpolation) takes 2.816 seconds, building 1000 price index curves and 1000 discount curves simultaneously takes 5.869 seconds Those tests have been done using 3.5 GHz Quad-Core Intel Xeon.

References

Zine-eddine, A. (2013a). Convexity adjustment for inflation derivatives. Technical report, OpenGamma. 6

Zine-eddine, A. (2013b). Inflation caps and floors. Technical report, OpenGamma. 1

OpenGamma Quantitative Research

1. Marc Henrard. Adjoint Algorithmic Differentiation: Calibration and implicit function theorem. November 2011.
2. Richard White. Local Volatility. January 2012.
3. Marc Henrard. My future is not convex. May 2012.
4. Richard White. Equity Variance Swap with Dividends. May 2012.
5. Marc Henrard. Deliverable Interest Rate Swap Futures: Pricing in Gaussian HJM Model. September 2012.
6. Marc Henrard. Multi-Curves: Variations on a Theme. October 2012.
7. Richard White. Option pricing with Fourier Methods. April 2012.
8. Richard White. Equity Variance Swap Greeks. August 2012.
9. Richard White. Mixed Log-Normal Volatility Model. August 2012.
10. Richard White. Numerical Solutions to PDEs with Financial Applications. February 2013.
11. Marc Henrard. Multi-curves Framework with Stochastic Spread: A Coherent Approach to STIR Futures and Their Options. March 2013.
12. Marc Henrard. Algorithmic Differentiation in Finance: Root Finding and Least Square Calibration. January 2013.
13. Marc Henrard. Multi-curve Framework with Collateral.

    May 2013.
14. Yukinori Iwashita. Mixed Bivariate Log-Normal Model for Forex Cross. January 2013.
15. Yukinori Iwashita. Piecewise Polynomial Interpolations. May 2013.
16. Richard White. The Pricing and Risk Management of Credit Default Swaps, with a Focus on the ISDA Model. September 2013.
17. Yukinori Iwashita. Conventions for Single-Name Credit Default Swaps. December 2013.
18. Arroub Zine-eddine. Brazilian Swaps. December 2013.
19. Arroub Zine-eddine. Inflation: Instruments and curve construction. January 2014.
20. Arroub Zine-eddine. Inflation caps and floors. January 2014.
21. Arroub Zine-eddine. Convexity adjustment for inflation derivatives. January 2014.

About OpenGamma

OpenGamma helps financial services firms unify their calculation of analytics across the traditional trading and risk management boundaries.

The company's flagship product, the OpenGamma Platform, is a transparent system for front-office and risk calculations for financial services firms. It combines data management, a declarative calculation engine, and analytics in one comprehensive solution. OpenGamma also develops a modern, independently-written quantitative finance library that can be used either as part of the Platform, or separately in its own right.

Released under the open source Apache License 2.0, the OpenGamma Platform covers a range of asset classes and provides a comprehensive set of analytic measures and numerical techniques.

Find out more about OpenGamma
`www.opengamma.com`

Download the OpenGamma Platform
`developers.opengamma.com/downloads`

Europe
OpenGamma
185 Park Street
London SE1 9BL
United Kingdom

North America
OpenGamma
280 Park Avenue
27th Floor West
New York, NY 10017
United States of America

OpenGamma