The inflation column

November 2007

Christopher C. Finger, Fabien Couderc
chris.finger@riskmetrics.com, fabien.couderc@riskmetrics.com

The most common report template used in our risk application is called the Customizable Table Report. (Our marketing team surely disavows the less than enticing name.) In a standard report layout, positions are represented in rows, while types of risk are represented in columns. For research and development, adding rows to the report, typically through a pricing model for a new type of security, is, while not quite routine, also not profound. It is a matter of working within our framework. Adding columns, on the other hand, is needed when we encounter an altogether new type of risk, necessitating not just new pricing models but the more fundamental work of defining a new family of risk factors.

Of course, pension funds and individuals have recognized inflation risk in their liabilities for longer than any conversation about inflation as a proper risk factor. At the same time, governments hold inflation-linked assets (primarily in the form of future tax revenues), and so a market in inflation-linked bonds was natural. The UK was first to issue these bonds, in the early 1980s, and were followed by most major sovereigns by the 1990s. Trading volume has risen significantly in the last several years, with monthly turnover in the Euro inflation-linked market now at about 60% of the total inflation-linked inventory.[1]

Though this is hardly a new market, inflation exposure through bonds is mostly inextricable from interest rate exposure and the technicals of the bond market, and thus the convention has been to lump inflation risk into the overall fixed income risk on inflation-linked bonds.

Inflation swaps are thus a crucial innovation for two reasons: first, they isolate exposure to inflation, decoupling interest rate and inflation risk in much the same way that credit default swaps decouple credit from interest rate risk; second, they provide for synthetic exposure to inflation, meaning that in theory, inflation can be traded at any term between any two counterparties, rather than only through government auctions. Though inflation swaps have been traded for longer, it was only in 2003 and 2004 that the market truly established itself. By 2005, in the UK, total demand for swaps actually exceeded that for bonds.[2]

There are thus three distinct ways to gain (and possibly manage) exposure to inflation: naturally (for example, by a pension fund granting inflation-indexed benefits in the normal course of its business), through inflation-linked bonds and through inflation swaps. Natural exposures are for the most part liabilities, and therefore short inflation, while bond positions are for the most part long; swaps afford exposure in either direction. The time has come then for a common language to express the inflation risk consistently across these three sources, and for an inflation column in our risk report.

[1] See Amblard (2005).

[2] See Barclay's Capital (2006).

Realized and implied

To begin our definitions, it is important to distinguish two notions of inflation: realized and implied. Realized inflation is what central banks report—the actual inflation experienced over one month in the recent past. It gives us something real to reference but (obviously) can only reference historical periods, and can only be observed on a monthly (or possibly quarterly) basis, typically with a lag.

To be specific, realized inflation is observed on a well-defined index. A government entity defines a basket of goods and services, and reports the price of this basket, possibly averaged over different locations, during a given month. This price is the inflation index, and changes in this index from month to month the realized inflation. Multiple indices do exist for some regions,[3] and so to describe a bond precisely, it is necessary to specify not only the issuer and currency, but also the inflation index referenced.

Gathering the price information takes time, and so the publication of the index for one month does not occur until several weeks after month end.[4] We may think of the inflation index as an exchange rate—the number of units of our nominal currency (for instance, dollars, euros or pounds) needed to buy one unit of our inflation currency (the basket) at a specific point in time.

With financial products that reference future realized inflation, we have a notion of implied inflation whenever we observe prices on these products. Implied inflation, strictly speaking, is thus tied to a particular product or set of products. It can reference any future period that the financial products do, and can be observed at a daily frequency, or as often as the relevant financial products trade.

For the standard inflation swap, the notion of implied inflation is natural. In the standard swap contract, the inflation payer contracts to pay a notional amount times the realized inflation over some future period, while the fixed payer contracts to pay a fixed rate times the same notional amount. The quoted inflation swap rate is the fixed rate at which both parties agree to enter the contract with no upfront payment. This quoted rate can then be thought of as the market's expectation of the inflation to be realized over the referenced period. There are further details to sort out to make this a useful risk factor, but the interpretation of the quoted swap rates is clear.

In a sense, the inflation swap contracts play a similar role to credit default swaps, isolating a single source of risk such that their prices have a clear interpretation. Also as with credit, inflation-linked bonds mingle inflation risk with other risk sources, particularly interest rates and liquidity, making the interpretation more difficult. In fact, the choice of risk factors for an inflation-linked bond is not at all obvious.

Break-even defined

Inflation-linked bonds for the most part now all follow the Canadian model, implemented first by (yes) Canada in 1990 and adopted by most every issuer since. Under this model, a bond pays a fixed coupon

[3]For instance, France issues bonds linked to both a French and a harmonized European inflation index.

[4]Though the index by construction does not reflect prices on any specific day, it is market convention to consider the inflation index for a specific month to be effective as of the first day of that month.

Figure 1: Investing in real and nominal bonds

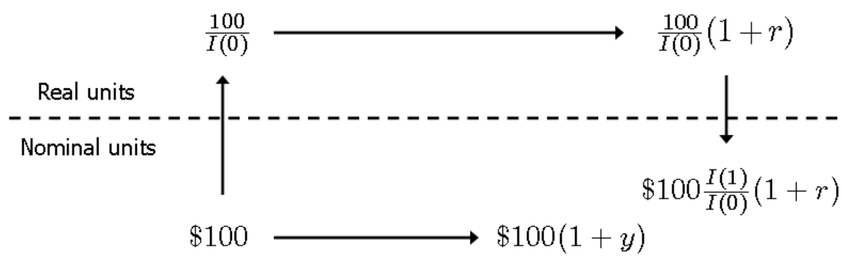

rate on a face value that inflates with realized inflation. For instance, say an inflation-linked bond pays a two percent annual coupon and that the initial face value is \$100. Suppose realized inflation in the first year is three percent. Then the face value inflates to \$103 and the first coupon is two percent of this amount. For the next period, realized inflation is applied to the \$103 face value, and the coupon is paid on this new amount. At maturity, the bond pays its final inflated face value.

To make the sources of risk more clear, we describe the same investment in a slightly different way. (See Figure 1.) We begin with the face value of \$100, then convert this amount into a number of units—obtained by dividing \$100 by today's inflation index, $I(0)$—of the inflation currency. We then invest in the inflation currency in a bond yielding two percent. After one period, our investment has grown to $102/I(0)$ units of the inflation currency. At that time, we convert back into the nominal currency at the prevailing inflation index value $I(1)$, leaving us with a final value of $102 \cdot I(1)/I(0)$. The final yield on our investment can thus be expressed as $(1+r)(1+i)-1$, where $i = I(1)/I(0) - 1$ is the realized rate of inflation, and r is the real yield, that is, the yield we earn in the inflation currency. Note that only i is unknown at the time we make our investment.

At the initial date, we could have alternately invested in a nominal (that is, conventional) bond. Let y be the nominal yield. We define the break-even rate of inflation (BEI) as the rate of inflation that would need to occur such that our two strategies (investing in inflation-linked or nominal bonds) produce the same total yield. Equating the two yields, we see that

$$BEI = \frac{1+y}{1+r} - 1.$$ [5]

The BEI is for the most part an expression of the market's expectation of future inflation, but other effects, including the relative liquidity premia for the nominal and real bonds, are also embedded. For inflation swaps, as discussed earlier, the quoted rates effectively are break-even inflation themselves.

Thus from two market observables—the prices of the nominal and inflation-linked bonds—we may calculate three potentially interesting quantities: the nominal yield (y), the real yield (r), and the break-even inflation (BEI). The relationship above linking the three implies that the three quantities cannot move independently, and that only two may be considered as sources of risk. Which two is for us to choose.

[5] For small values of y and r, this approximates the familiar Fisher equation, $y = r + i$.

Historically, bond investors chose to express risk in terms of real and nominal yields. Thus, the nominal value of an inflation-linked bond was driven by the real yield and the current inflation index. This was sensible given the analogy with foreign-denominated bonds, whose value is linked to the yield in the foreign currency and the spot exchange rate.

This view has three major drawbacks, however. First, while spot foreign exchange rates change daily, the inflation index is published only monthly, and with a lag, and cannot be treated with standard risk models. Consequently, the risk due to this factor is typically not modeled statistically, but rather only shocked through stress tests.

Second, real interest rates are strongly correlated with nominal rates, and thus it is difficult to argue that the two represent distinct sources of risk. We plot nominal and real rates and BEI for the US and France[6] in Figure 2. Over this period, the correlations of weekly moves in the nominal and real rates are 87% for the US and 91% for France. In contrast, the correlations between nominal rates and BEI are 47% for the US and 40% for France. Interestingly, these relationships do change in time: in particular, BEI has moved in synch with nominal rates since early summer in the US.

Third, while real interest rates do behave like standard risk factors, they bear no relation to either the natural inflation exposures in pension liabilities or to inflation swaps. Thus using the real rate framework affords no mechanism to isolate the offset in inflation risk in pension liabilities from buying inflation-linked bonds, nor to compare the relative market prices of inflation in the swap and bond markets.

A good factor

We have argued that BEI is the most appropriate choice for an inflation risk factor, as it is only moderately correlated with nominal rates and has meaning within the contexts of all three possible sources of inflation risk. There are, however, refinements needed in order to satisfy two more criteria for good risk factors: homogeneity and portability.

Homogeneity is a somewhat elusive concept. Essentially, what we want is that our risk factor is well suited to the time series models we apply for forecasting. As we have discussed in previous issues, we seek to apply a similar forecasting model to all risk factors. This is not to assert that all risk factors are the same. Rather, it is to say that since the forecasting model does not account for peculiarities of individual risk factors, we have to do our work up front, removing anything that is special about a risk factor before applying the forecasting model. That a risk factor is homogeneous is that the risk factor has the same meaning for us whether we observe it ten years ago or today, and in February or August.

One obvious necessity for us to produce homogeneous inflation risk factors is to work with BEI curves. The BEI on a ten-year bond at issuance is a reflection of inflation expectations over ten years, while as the bond ages, its break-even inflation pertains to an ever shorter period. Just as with nominal bonds, where we choose not to use the yields to maturity of individual bonds as risk factors, we build a curve of break-even inflation. The points on our curve are constant maturity zero-coupon points:

[6] The OATe bonds, linked to harmonized European inflation, ex-tobacco

Figure 2: Five-year zero-coupon nominal rate, real rate and break-even inflation. US inflation-linked bonds and French OATe

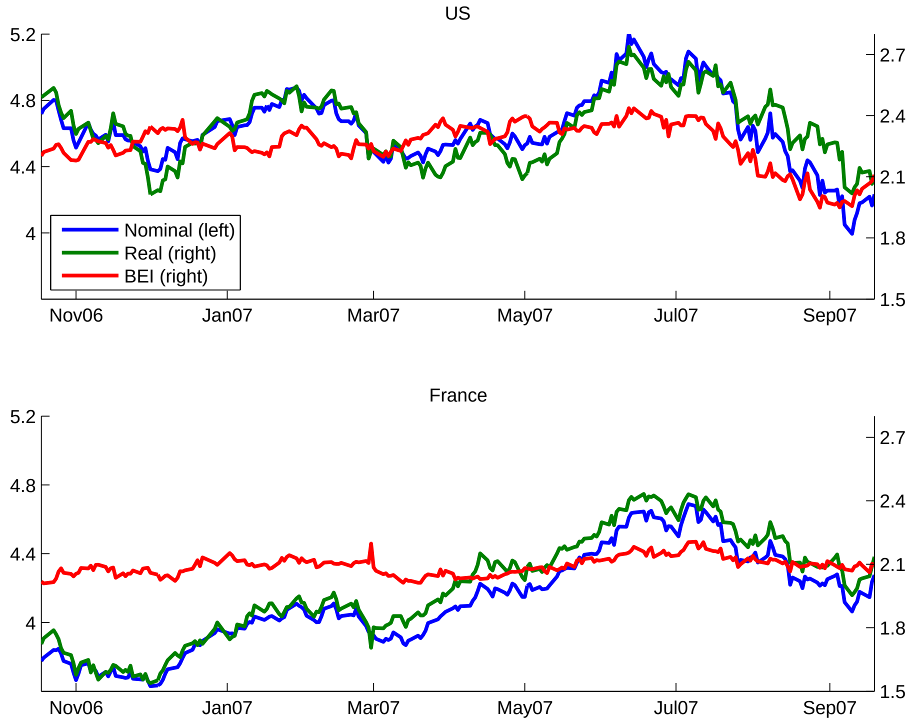

the market-implied inflation from today for specific lengths of time into the future.

More specific to inflation is the issue of seasonality. Generically, seasonality is a predictable, repeating pattern that, while contributing to the amount of fluctuation in a factor, should not be treated as a random quantity. We do expect seasonality in inflation, as certain components of the basket of goods and services, notably energy and commodity prices, will systematically be highest at the same time every year.

We illustrate the magnitude of seasonality effects in inflation with the US Consumer Price Index. In Figure 3, we first plot the overall monthly realized inflation, along with a horizontal bar indicating the mean inflation rate (corresponding to 2.6% annually) over the period. Notably, we see a significant number of months where realized inflation is negative. Though this is at first counterintuitive, we note that the same months each year tend to produce negative inflation, corresponding to the notion of seasonality. Twelve-month inflation, that typically reported, ranged between 1.1% and 4.6% over our sample period.

The second and third plots in Figure 3 show the decomposition of the overall inflation into a seasonal pattern[7] and a residual (or unexplained) component.

[7]Extracted using a multiplicative X-11 method. In its simplest form, the method extracts seasonality using a twelve month centered moving average, then smooths the seasonality for a given month through a centered four observation moving average. See Shiskin et al (1967).

Figure 3: Inflation decomposition, US Consumer Price Index, 2003-2007

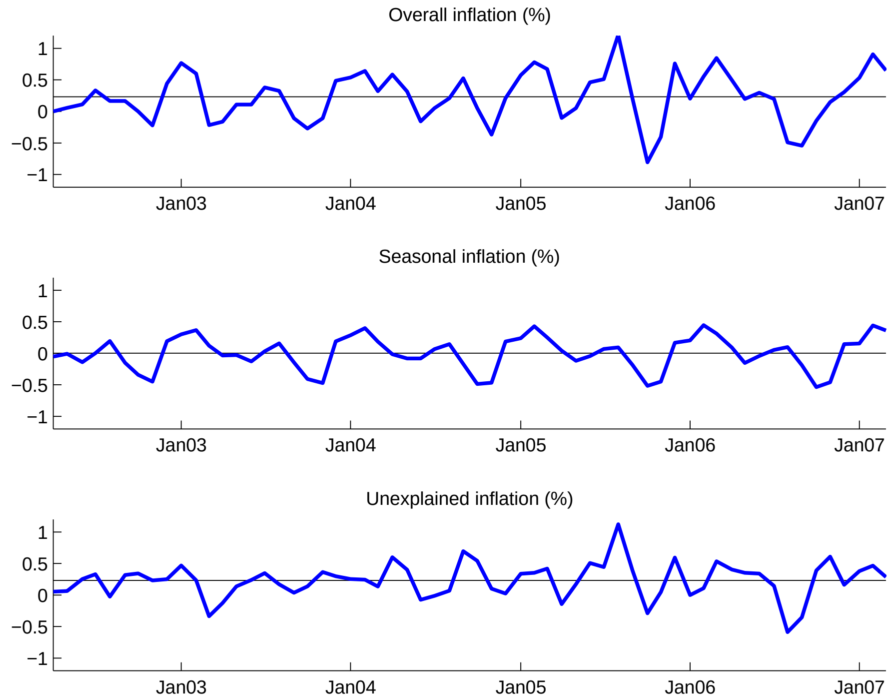

The seasonal piece is remarkably stable, with a peak each February of roughly 0.5% and a trough each October of roughly -0.5%. The unexplained component looks not predictable, but random, and is of comparable magnitude to the seasonal component. For forecasting, then, it does not make sense to treat the overall inflation as an unexplainable, random quantity. Rather, we should take advantage of the predictability of the seasonal component, and treat only the unexplained component as random. Since the seasonal and unexplained portions are of comparable magnitude, ignoring this decomposition could result in a twofold overestimation of risk.

Admittedly, one-month realized inflation is the quantity with greatest seasonality effects. Naturally, if we look at inflation over a full year, or an integer number of years, we observe no seasonality. Moreover, for inflation over longer periods, seasonality is less crucial: eighteen-month inflation is a combination of one-year inflation (with no seasonality) and six-month inflation (which clearly depends on which six months are considered); 20.5 year inflation is dominated by the twenty-year component, and the seasonal six-month component has little practical effect.

All of these arguments apply equally to market-implied as well as realized inflation. Thus, for our purposes, we define our risk factors as the unexplained (or seasonally adjusted) BEI. Following the logic above, this has material impacts on our risk forecasts for shorter, non-integer maturities.

Portability

Besides its impact on forecasting, our seasonality adjustment plays a role in our ability to compare inflation through time. That November 2006 inflation was lower than February 2006 inflation is not surprising, since February inflation is seemingly always higher than November inflation. What is of interest is which month's inflation was higher relative to the typical seasonal inflation we would expect. With BEI, we may want to ask whether the 2.5-year inflation implied by the market looks high relative to where it was six months ago, and whether today's rates represent an attractive opportunity to purchase an inflation hedge. Clearly, this is a question we must ask in the context of seasonally adjusted inflation.

This brings us to the last criterion for our risk factor: portability. If homogeneity refers to factors having the same meaning through time, portability refers to their having the same meaning across assets.

Thus beyond comparing through time, we would also like to make comparisons between break-even inflation implied by different markets, particularly bonds and swaps, and even across different inflation indices within the same market. This facilitates judgments on relative value, and gives us a common language to describe our different sources of inflation risk. These tasks are not possible using risk factors that embed the specific technicalities of the different markets.

What we require is a bridge, not unlike what we have proposed for credit spread risk.[8] For spread risk, though each market (bonds and credit default swaps) does express risk in terms of spread, the spreads used conventionally by each market mean subtly different things. Thus, it is impossible to say whether a yield spread on a bond of 50bp is expensive relative to a fair spread of 60bp on a default swap. Only by converting all spreads into a common language can we make relative value (or risk) assessments.

Consider an example. Suppose on 25 October 2007, we enter a five-year inflation swap on the harmonized European inflation index (HICP). By convention, this swap will work with a three-month lag, meaning that it matures on 25 October 2012, but will cover the increase on the HICP from 25 July 2007 to 25 July 2012. To obtain the HICP level associated with July 25, another convention—the interpolation rule—applies. For swaps, the convention is that the July index level applies to any effective date in that month. (For bonds, convention dictates that we interpolate between the July and August index values.)

In fact, as of October 25, August and September inflation had already been published. Thus, viewed from the day of swap inception, the covered period is composed of three parts:

1. From 25 July 2007 to 1 September 2007, for which we know the realized inflation (0.43%) with certainty,
2. From 1 September 2007 to 25 October 2007, for which we do not know the inflation with certainty (October inflation is not yet available), but which has already occurred, meaning we should have some notion of what inflation was experienced, and
3. From 25 October 2007 to 25 July 2012, which occurs entirely in the future.

[8] See Finger (2005).

The market quote we see for our five-year inflation swap incorporates the three periods, and yet only for the third period is inflation truly unknown. It is really only the inflation expectation over this third period that we should consider at risk. Moreover, the inflation we attribute to the first period depends on the lag and the interpolation rule, both of which may vary across the bond and swap markets. Finally, note that although we started with a standard swap, with an integer number of years to maturity, the portion of the swap that is at risk covers a period of four years and nine months, meaning that it is subject to seasonal effects; a similar swap initiated in March 2007, covering a period from March 2007 to December 2011 and thus not including the winter 2012 months when inflation should be higher, would presumably command a lower quoted rate, at least adjusted for the first two periods.

In the end, we strip out the known (or well estimated) inflation pertaining to the first two periods, leaving us only with the market's view on future inflation. With the seasonality adjustment applied as well, we arrive at a truly portable risk factor—the adjusted break-even inflation—that is free of market conventions and is comparable across products and time periods. What distinguishes different break-even curves is simply the market from where they are sourced, but not their definition.

Risk and beyond

Clearly, our first application is as we stated at the outset: a consistent, intuitive treatment of inflation risk derived from natural positions, bond positions and swap positions which we can place in the new column on our risk report. For the natural inflation exposures, we are free to choose the break-even curves that best suit our needs for valuation and risk analysis, based on the quality and relevance of the data rather than on which market conventions happen to line up best with our liabilities. For all exposures, the risk factors fluctuate daily, with the predictable parts of the fluctuation extracted, and so are suitable for our risk forecasting models.

From here, where do we go? By defining a clean inflation risk factor, we have the appropriate quantity to model for inflation caps and floors, and for structured investment products referencing inflation. We also have all of the pieces to decompose corporate-issued inflation-linked bonds into interest rate, inflation and credit risks.

All this is work ahead of us, but it is back to working within the framework, that is, adding rows. Our new inflation column should serve us well for some time.

Further reading

- Amblard, G. (ed.) (2005). Inflation-linked products in the Euro area. Association des Marchés de Taux en Euro.
- Barclay's Capital (2006). Global inflation linked products: A user's guide.
- Finger, C. (2005). Spread values, Research Monthly, November.
- Shishkin, J. et al (1967). The X-11 variant of the census method II seasonal adjustment program. Technical Paper No. 15. U.S. Department of Commerce.