Imperial College London
Department of Mathematics

A Quadratic Gaussian Year-on-Year Inflation Model

Hringur Gretarsson

Submitted in part fulfilment of the requirements for the degree of
Doctor of Philosophy in Mathematical Finance of Imperial College London
and the Diploma of Imperial College
Thursday 31 January 2013

Declaration

I herewith certify that all material in this dissertation is my own work. Other sources of information that have been used are duly acknowledged.

Hringur Gretarsson

Abstract

We introduce a new approach to model the market smile for inflation-linked derivatives by defining the Quadratic Gaussian Year-on-Year inflation model—the QGY model. We directly define the model in terms of a year-on-year ratio of the inflation index on a discrete tenor structure, which, along with the nominal discount bond, is driven by a log-quadratic function of a multi-factor Gaussian Markov process.

We find closed-form expressions for the drift of the inflation index and for inflation-linked swaps. We get a Black-Scholes-type pricing formula for year-on-year inflation caplets in semi-analytical form. The formula contains an integral of a multivariate Gaussian density over a quadratic domain. In a two-dimensional case, we show how this integral reduces to a one-dimensional integration along the boundary of a conic section.

In the case where the year-on-year inflation ratio is driven by two factors, we specify a spherical parameterisation. This gives an intuitive control over the curvature and the skew of the year-on-year inflation smile and shows the maximum curvature and skew obtainable with a particular three-factor version of the QGY model. Within this three-factor model, we identify a parameterisation to control the autocorrelation structure of the inflation index.

We calibrate the model to year-on-year inflation options on the UK's Retail Prices Index (RPI) and the eurozone's Harmonised Index of Consumer Prices Excluding Tobacco (HICPxT) and get a good fit to the smile of implied volatilities. We use the calibrated model to price HICPxT zero-coupon inflation options and RPI limited price indices (LPIs). Furthermore, we provide methods to interpolate the process for the inflation index and the year-on-year inflation ratio between dates on the tenor structure.

Keywords: Year-on-year inflation modelling; multi-factor log-quadratic Gaussian model; stochastic-volatility parameterisation; inflation autocorrelation; year-on-year inflation calibration; LPI pricing.

Acknowledgments

In this thesis, I cover results that I have obtained as part of postgraduate studies in Mathematical Finance at Imperial College London. While undertaking these studies, I have been funded by Lloyds Banking Group. This has made it possible for me to learn from highly knowledgeable people both in academia and in the financial industry. This has also given me the opportunity to attend interesting conferences around the world. For all this I am forever thankful.

I have had the good fortune to work closely with the members of the Inflation Derivatives Quantitative Research team at Lloyds. I am very grateful to Manlio Trovato and Diana Ribeiro for their great support throughout my research. They have helped me take the necessary steps for the research, given me motivation throughout my studies and helped me improve my knowledge in mathematical finance and programming. I would like to thank Diana for verifying formulas and calculation methods presented in the thesis and helping me improve the structure and arguments in the thesis. I would also like to thank the inflation traders at Lloyds, Charles Harris and Evan Macdonald, for giving me useful insights into inflation markets.

At Imperial College, I would like to thank my supervisor, Prof. Mark H.A. Davis, for his guidance over the years and for making valuable comments regarding various parts of the thesis. Finally, I would like to thank Björgvin Sigurðsson and Chris Kenyon for making useful suggestions for the thesis and Russell Moxham for proofreading the thesis in detail.

'You can do it!' MÓ (2006).

This thesis is dedicated to my father, Gretar Reynisson,
and to the memory of my mother, Margrét Ólafsdóttir.

Contents

List of Tables

List of Figures

Introduction

Bank of England Governor Sir Mervyn King (2012) said that 'Low and stable inflation is a prerequisite for economic success'. This is supported by examples of hyperinflation, such as in Germany after the First World War, and deflation, as experienced in Japan in the 1990s and in Europe in 2008 and 2009, which have proved to have detrimental consequences for economic growth and have often gone hand-in-hand with recession. Most governments therefore carefully monitor inflation, it being one of the key indicators of economic welfare. Central banks also look at inflation when deciding whether to change their interest rates to stimulate or dampen economic activity.

Inflation indices track price changes in the economy using well-established calculation methods. The annual percentage increase in the inflation index gives the year-on-year inflation rate—an important measure of current inflation in the economy. Inflation-linked derivatives are financial contracts with payoffs that are linked to the values of inflation indices and year-on-year inflation rates. There is a wide variety of inflation-linked derivatives traded in financial markets around the world. The main buyers of these derivatives are investors and companies sensitive to movements in inflation, such as pension funds, utility companies and retail companies.

Rather than considering macroeconomic models used for forecasting inflation by central banks and investment funds, we focus on inflation models that are suitable for pricing inflation-linked derivatives in investment banks. In this thesis, we introduce a new inflation model that we believe fills a gap in the current literature on inflation modelling. This thesis is aimed at explaining the complex properties of inflation markets and we hope that it will inspire further research into inflation modelling.

We now briefly describe the main inflation-linked derivatives and inflation models and give an overview of the contribution and structure of this thesis.

Inflation-linked derivatives

The three main types of inflation-linked derivative are bonds, swaps and options. They are linked to a particular inflation index, either directly or through the year-on-year inflation rate.

As inflation can reduce the return on investments, many governments have issued inflation-linked bonds. They guarantee a given real return because a bond's repayments compensate for the increase in an inflation index over a given period. The prices of inflation-linked bonds and nominal bonds determine the forward inflation index, which gives an estimate of future inflation and can therefore help shape government policies.

Investors can protect their investments against inflation by swapping the value of the forward inflation index with the realised value of the inflation index using a zero-coupon inflation swap. It is important to note that these swaps can incorporate a payment delay convexity correction, which depends on the correlation between inflation and nominal interest rates. Another type of swap occurs when the realised year-on-year rate is swapped with the year-on-year inflation forward rate. These swaps incorporate a year-on-year convexity correction, which depends on the autocorrelation of the inflation index.

Zero-coupon options on the inflation index are an alternative protection against inflation. The buyer of a zero-coupon inflation cap has the option to receive the relative increase in the value of the inflation index in exchange for a strike (given amount) at maturity. The buyer of a zero-coupon inflation floor has the option to receive a strike (given amount) in exchange for the relative increase in the value of the inflation index at maturity. Therefore, a cap is only exercised if the realised inflation index is above the strike, while a floor is only exercised if the realised inflation index is below the strike. Similar options exist on the year-on-year inflation rate.

The prices of inflation-linked options suggest that the distributions of the inflation index and the year-on-year inflation rate have excess kurtosis and negative skewness. These distributional properties are equivalent to positive curvature and negative skew in the so-called volatility smiles that are implied by the prices of inflation-linked options.

Inflation models

Different inflation models have been proposed in the literature to price inflation-linked derivatives. Three main types of inflation model exist: foreign-exchange analogy models, index models and year-on-year models.

The first type uses the foreign-exchange analogy by interpreting the inflation index as the exchange rate from a hypothetical real currency to the nominal currency. Jarrow and Yıldırım (2003) define this type of model using the Heath-Jarrow-Morton framework, with a multi-factor extension provided by Trovato et al. (2009). These models have no curvature or skew in the volatility smiles. To incorporate curvature and skew, a multi-currency quadratic Gaussian model by McCloud (2008) could be applied to model inflation. However, for foreign-exchange analogy models, it is difficult to specify values for parameters of the real currency because of illiquidity of derivatives on the real rate.

The second type of inflation model directly defines a process for the forward inflation index. Belgrade et al. (2004) use this approach in a market model framework. This model has good flexibility in defining the correlation structure of the inflation index. However, the model has no curvature or skew in the inflation volatility smiles. Mercurio and Moreni (2006, 2009) extend the approach of Belgrade et al. and control the smile by using the SABR model by Hagan et al. (2002) and the Heston model (1993). These stochastic-volatility extensions only give approximate option-pricing formulas and make it difficult to simulate the models' processes. Furthermore, market models are high-dimensional and therefore computationally intensive for calibration and simulation.

The year-on-year inflation derivatives have become more liquid in the last few years, giving rise to the third type of inflation model, which models the year-on-year inflation rate directly. Year-on-year inflation models have intuitive parameters and good analytical properties for the year-on-year inflation rate. Kenyon (2008) assumes that the year-on-year rate is Gaussian and provides mixture methods and stochastic-volatility extensions to control the year-on-year inflation volatility smile. On the other hand, it is difficult to use these models to price zero-coupon swaps and options. A log-normal model by Jäckel and Bonneton (2010), with multi-factor extension by Trovato et al. (2009), specifies a process for the year-on-year ratio of the inflation index, as well as modelling the nominal interest rate. However, the model does not have smile.

Objective and contribution

Inflation models should account for the main distributional properties of the inflation index and of the year-on-year inflation rate. These include the curvature and skew of the inflation volatility smiles, as well as the year-on-year and payment delay convexity corrections. Furthermore, an inflation model should have good analytical tractability, give fast calibration to inflation-linked options and show intuitive behaviour in the model parameters. It should also be easy to simulate the modelling processes so one can calculate the prices of exotic inflation-linked derivatives.

Existing inflation models do not incorporate all these features. We therefore introduce a new inflation model, which we call the Quadratic Gaussian Year-on-Year inflation model—the QGY model. The model combines good analytical properties of quadratic Gaussian processes and the intuitive parameterisation of year-on-year inflation models.

We directly define the QGY model in terms of the year-on-year ratio of the inflation index, on a discrete tenor structure. Using this ratio, we can get processes for both the year-on-year inflation rate and the inflation index. The year-on-year inflation ratio, along with the nominal discount bond, is driven by a log-quadratic function of a multi-factor Gaussian process. The year-on-year inflation ratio is a Markovian function of this Gaussian process. The log-quadratic form and the Markovianity make it possible to calculate both zero-coupon and year-on-year inflation swaps analytically, to get semi-analytical pricing formulas for year-on-year inflation options and to incorporate curvature and skew in the year-on-year inflation volatility smiles.

In a three-factor set-up, the model offers intuitive parameters and fast pricing and calibration to inflation-linked swaps and year-on-year inflation options. Furthermore, we can control the payment delay and year-on-year inflation convexity corrections. Because the modelling processes are functions of a multi-factor Gaussian process, it is simple to perform a Monte Carlo simulation to price zero-coupon inflation options and exotic inflation-linked derivatives. We also interpolate the modelling processes to be able to price inflation-linked derivatives that depend on the inflation index for dates not on the model's original tenor structure.

As well as deriving formulas in the QGY model and analysing its behaviour, we have implemented all the parts of the three-factor model in the quantitative library at

Lloyds Bank. The implementation involved the coding of the modelling processes, pricing formulas, reparameterisations, interpolations and a part of the calibration, using an object-oriented approach in C++. These functions were then linked to the Microsoft Excel spreadsheet application.

A paper on the QGY model by M. Trovato, D. Ribeiro and H. Gretarsson was published in Risk Magazine in September 2012. The paper covered the main pricing formulas, provided a three-factor reparameterisation of the model and showed pricing results from a calibration to year-on-year inflation options.

Structure

This thesis is divided into three parts. Part I consists of four chapters, where we cover the preliminaries to the QGY model by looking at inflation markets, derivatives, pricing framework and models. In Part II, we look at the QGY model in detail in five chapters. This involves specifying the modelling processes, listing the necessary pricing methods, deriving pricing formulas, defining parameterisations, calibrating the model, pricing exotic options and providing interpolation methods for the year-on-year inflation process and its parameters. Part III contains five appendices, in which we explain calculation methods that are used in the thesis.

In Chapter 1, we discuss economic theories on the causes and consequences of inflation, look at the history and features of consumer price indices in the UK and the eurozone and discuss the inflation-linked bond markets for the UK's Retail Prices Index (RPI) and the eurozone's Harmonised Index of Consumer Prices Excluding Tobacco (HICPxT). By using inflation-linked bonds, we define the forward inflation index and introduce the notion of a real currency.

In Chapter 2, we define inflation-linked derivatives that are traded in the market. We look at the payment delay convexity correction and see how it changes with the correlation between inflation and nominal interest rates. We look at the year-on-year convexity correction and see how it depends on the autocorrelation of the inflation index. We define particular inflation volatilities implied by the prices of zero-coupon and year-on-year inflation options and show examples from RPI and HICPxT inflation option markets.

In Chapter 3, we explain a pricing framework for inflation by assuming the existence of a so-called pricing kernel, which is driven by a multi-dimensional (and continuous) Brownian motion. In this framework, we show how to change between measures and how to define the foreign-exchange analogy for inflation.

We use this framework in Chapter 4, where we review inflation models that exist in the literature. These include foreign-exchange analogy inflation models, inflation index models and year-on-year inflation rate models. We also review the class of quadratic Gaussian models that have been applied to the short rate and used in a multi-currency framework.

In Chapter 5, we specify the modelling processes of the QGY model and provide formulas from McCloud (2008) for the expectations of log-quadratic Gaussian processes.

In Chapter 6, we calculate pricing formulas for inflation-linked swaps and options. We derive closed-form expressions for the drift of the inflation index and for inflation-linked swaps. These swaps include the forward inflation index with a payment delay and the year-on-year inflation forward rate. We derive a Black-Scholes-type semi-analytical pricing formula for year-on-year inflation options.

In Chapter 7, we look at a special version of the three-factor QGY model so that we can look in more detail at the distribution of the year-on-year inflation ratio. We show how to control the autocorrelation structure of the inflation index and the correlation between inflation rates and nominal interest rates. We define a spherical reparameterisation of the inflation ratio in order to get an intuitive control over the year-on-year inflation volatility smile.

In Chapter 8, we calibrate the model to market prices of year-on-year inflation options and forward rates on RPI and HICPxT. We also calibrate the model to market prices on other dates to look at the stability of the model parameters. We simulate the QGY model using Monte Carlo methods to calculate the prices of zero-coupon options on HICPxT. We also simulate the QGY model to price RPI limited price indices (LPIs), which are path-dependent exotic derivatives on the RPI year-on-year rates.

In Chapter 9, we provide two interpolation methods, which make it possible to price inflation-linked derivatives that depend on the value of the inflation index on dates that are not part of the discrete tenor structure of the QGY model. The first interpolation method is well suited to defining a process for the inflation index, while the second interpolation method is ideal for defining a process for the year-on-year inflation rate. We price inflation-linked swaps and options to verify the stability of these interpolation methods.

In the appendices, we explain calculation methods that are used in the thesis. In Appendix A, we show that the Hull-White short-rate equation is equivalent to the log-linear Gaussian discount bond. In Appendix B, we derive the variance, skewness and excess kurtosis of a quadratic Gaussian distribution. In Appendices C and D, we detail calculations by McCloud (2008) on the expectations of log-quadratic Gaussian processes. In Appendix E, we show how to reduce a two-dimensional integral of a Gaussian density over a conic domain to a finite sum of one-dimensional integrals. We use this method to calculate the prices of year-on-year inflation options in the QGY model.

Part I

Preliminaries

1 Inflation market

Since the invention of money, economies have experienced variations in its purchasing power—the quantity of goods and services it can buy. Inflation occurs when the purchasing power of money decreases and deflation occurs when the purchasing power of money increases. In §1.1, we discuss in more detail the purchasing power of money and the causes and consequences of inflation.

To measure the purchasing power of money, most countries calculate and publish consumer price indices once a month.[1] In §1.2, we look at UK's Retail Prices Index (RPI) and the eurozone's Harmonised Index of Consumer Prices (HICP). Although they differ slightly in their calculation methods, they are composed of a combination of goods and services, which is updated annually to reflect changes in the consumption pattern and expenditure of a typical household within the given economy.

One way to reduce the effect that inflation has on wealth is to use consumer price indices to adjust nominal amounts, so that these amounts keep their purchasing power over time. Governments have issued inflation-linked bonds, for which predetermined nominal amounts are adjusted for changes in prices and paid at certain dates. The prices of these bonds show the expected inflation and can help governments set their economic policies. Inflation-linked bonds and derived instruments are covered in §1.3. There we will also introduce the real currency and the real forward rate.

A detailed discussion of the conventions, behaviour and derivatives in the inflation market can be found in Deacon et al. (2004) and Kerkhof (2005) and in documents by the UK's Office for National Statistics (2012b) and the eurozone's Eurostat (2001).

[1] The Australian Consumer Price Index is calculated and published once every three months.

1.1 Purchasing power of money

In the description of inflation, an important term is the purchasing power of money. It represents goods and services that can be purchased with a unit of money. Inflation occurs when one needs more money over time to buy a fixed quantity of goods and services. On the other hand, deflation happens when one needs less money over time to buy a fixed quantity of goods and services.

Money

It is clear that purchasing power depends on which goods and services are purchased and in relation to what money. Friedman and Friedman (1980) mention several forms of money. Stones were used on the island of Yap in the western Pacific Ocean until the beginning of the twentieth century, shells were traded between Indians and the early European settlers of America, tobacco was used as currency in Virginia in the seventeenth and eighteenth centuries and cigarettes were exchanged for goods in post-war Germany.

Gold has been used as money for centuries, with the gold standard having been introduced in the UK and the United States in the nineteenth century and used until 1971. Using the gold standard meant that paper money could be exchanged for a fixed weight of gold. As an example, the US dollar was fixed at $20.67 per troy ounce (31.10 grams) of gold at the beginning of the twentieth century. As Fisher notes (1928), although the currency was exchangeable with a fixed number of grams of gold, it was not fixed in regard of the number of goods and services it could buy—that is, its purchasing power changed over time, as with all other types of money.

Today, the most accepted mean of exchange is paper money and coins—i.e. fiat money. It is issued by central banks, although commercial banks were once allowed to print it. Checking deposits, which are money held in transactional accounts, are another form of money. They can be used to pay for goods and services, through the writing of cheques and the use of debit cards. As described by Samuelson and Nordhaus (2010), the introduction of fractional-reserve banking allowed commercial banks to own only a fraction (e.g. one-tenth) of the amount they lent out. This fraction is reserved in paper money or with the central bank. In turn, this makes the supply of money in deposit accounts much higher than paper money. Paper money and chequing deposits form the most narrowly defined part of the money supply, with the other part composed of less liquid monetary assets, such as savings deposits and money market funds.

Goods and services

To measure the purchasing power of money, one needs to use a combination—a basket—of goods and services. This basket should be representative of the typical expenditure of households at a given point in time and in the relevant region. It was not until the twentieth century that a systematic approach was introduced to estimate the weights of the items of this basket and to check their prices.

The inverse of the purchasing power of money is the price level experienced in the economy at any given time. Consumer price indices (CPIs) are used to approximate the price level. From a given index, one can calculate the relative changes in the price of the basket over time. Fisher (1922) gives an overview of research that has been carried out to define the basket and to calculate a consumer price index in the UK for the last thousand years. A more recent study by O'Donoghue et al. (2004), gives estimates of the index since 1750. We will look in more detail at CPIs in the UK and the eurozone in §1.2.

Economics of inflation

Fisher (1928) explains changes in price levels by looking at the relative changes between the circulation of money and the circulation of goods in the economy. Inflation occurs when the circulation of money increases relative to the circulation of goods, and opposite for deflation. Modern economic theory uses the concept of demand-pull inflation and cost-push inflation. Demand-pull inflation occurs when the demand for goods and services increases relatively more than production. Cost-push inflation is when production costs increase, e.g. from an increase in oil prices.

In the 1920s, Germany had to pay reparations, necessitating in a colossal increase in the supply of German marks. This caused a period of hyperinflation, with price levels increasing 100-billionfold from 1922 to 1923. As discussed in Samuelson and Nordhaus (2010), unstable and high levels of inflation—even if nowhere close to those in 1920s Germany—have various undesirable effects on the economy. An unexpected inflation shock can redistribute the wealth[2] from the lender to the borrower of a nominal fixed rate loan. In this case, the fixed nominal repayments that the borrower pays to the lender end up having less purchasing power than expected.

Because of the many undesirable effects of unstable and high inflation, many governments have taken up an inflation target policy to keep inflation stable at a reasonable level. This level is usually around 2%-4% annual inflation. Central banks can adjust the nominal interest rates of their reserves, to control the money supply in the economy. In a recession, lowering the interest rate increases the money supply and helps stimulate economic activity. As the interest rate cannot go below 0%, a liquidity trap can occur, where the money supply cannot be increased by further reducing the interest rate. In a recession, prices often fall, creating deflation. Therefore, an inflation target of 0% would not be optimal, as it would increase the probability of entering into a liquidity trap, prolonging the recession.

1.2 Consumer price indices in the UK and the eurozone

Here we look at the history of consumer price indices in the UK and the eurozone and we look at the calculation methods behind them. In 1914, to protect workers during the First World War, the UK Government began to construct an inflation index that measured the increase in the prices of goods and services. The items and their corresponding weights were chosen according to a 1904 survey of households' expenditure but it was not until 1936 that a new survey was conducted to update the items and weights. After an update of the index in 1955, the Retail Prices Index (RPI) was introduced in January 1956.

Since 1962, the items and weights that comprise RPI have been updated annually to reflect changes in household consumption patterns. The main source is the Living Costs and Food Survey, which excludes households in the top 4% for income, as well as pensioner households that derive more than 75% of their income from pensions and benefits. In Table 1.1 we show the main groups of RPI and their corresponding weights for different years. The prices of the items in the index are updated every month, through the checking of about $180,000$ prices for around 800 items all over the UK, and the corresponding value

[2] Here, wealth means the quantity of goods and services a given nominal amount can buy—that is, the purchasing power of that amount—equating to the amount adjusted for the price level at any given time.

RPI group	**1987**	**2007**	**2008**	**2009**	**2010**	**2011**	**2012**
Housing	157	238	254	236	237	238	237
Motoring expenditure	127	133	133	121	144	137	131
Food	167	105	111	118	112	118	114
Leisure services	30	68	65	67	65	64	71
Household goods	73	66	66	70	67	65	67
Household services	44	65	64	61	59	63	62
Alcoholic drink	76	66	59	63	64	60	56
Catering	46	47	47	50	47	47	47
Clothing and footwear	74	44	42	39	40	44	46
Fuel and light	61	39	33	49	40	42	45
Personal goods/services	38	39	41	41	41	38	39
Leisure goods	47	41	38	38	37	36	33
Tobacco	38	29	27	27	27	28	29
Fares/travel costs	22	20	20	20	20	20	23

Table 1.1: High level RPI weights since 1987. Weights are specified as parts per 1000 of the all item RPI. (Source: Office for National Statistics)

is calculated and then published in the middle of the following month.

RPI uses arithmetic averages to calculate the value of the subcategories of the index. As the weights are changed in February every year, January is chosen as the base month for the next twelve months. The items are stratified by product, location, item, class and group. Up to and including the location layer, elementary aggregates are calculated in a given month t with either of the two formulas:

$$\begin{array}{ll} \text{Average of price relatives} & \frac{1}{N}\sum_{i=1}^{N}\frac{p_{i,t}}{p_{i,0}} \\ \text{Ratio of average prices} & \sum_{i=1}^{N}\frac{p_{i,t}}{N}\Big/\sum_{i}\frac{p_{i,0}}{N}. \end{array}$$

Here $p_{i,t}$ is the price of the i-th item in a given subcategory in month t ($t = 0$ represents January). The first formula puts equal weight on each price quotation, while the second formula is used when the weights are assumed to be proportional to the price.

Aggregating items, classes, and groups is done with a weighted arithmetic average:

$$\sum_j w_j I_{j,t}\Big/\sum_j w_j.$$

Here $I_{j,t}$ is the value of the j-th subindex in month t (with previous January the base month), and w_j is its corresponding weight, which is fixed for the whole year.

To calculate the index over a period of more than a year, a method called chain-linking is applied for the class and group indices. As an example, we calculate the change in the index between January 2010 and July 2012. Using the weights for the year 2010, we calculate the change in the index between January 2010 and January 2011 using the corresponding prices in these months. In the same way, using the weights for the year 2011, we calculate the change in the index between January 2011 and January 2012. Finally, using the weights for the year 2012, we calculate the change in the index between January 2012 and July 2012. Multiplying these changes together gives the change in the index from January 2010 to July 2012:

$$I_{\text{Jul12/Jan10}} = I_{\text{Jul12/Jan12}} \times I_{\text{Jan12/Jan11}} \times I_{\text{Jan11/Jan10}}$$

In 1993, the eurozone created the Harmonised Index of Consumer Prices (HICP), which every member state calculates using the same methodology and structure. First it was used to see whether each member state fulfilled the necessary price stability and convergence criteria to join the European Union. Now it is used to assess the price stability of the eurozone as a whole. The eurozone's HICP is calculated from the average HICP in each member state, weighted for the state's relative share in the euro area's overall consumption. The difference between HICP and RPI is that HICP excludes mortgage interest payments, includes charges for financial services and uses a broader population to calculate its weights. The weights in HICP in the UK (called CPI) are also different from RPI, as the UK's CPI uses household final monetary consumption expenditure calculated annually from the National Accounts.

The main difference is that, when calculating the elementary aggregates, HICP uses the geometric average, which is always lower than the arithmetic averages used in RPI. The geometric average is given by

$$\sqrt[N]{\prod_{i=1}^{N} \frac{p_{i,t}}{p_{i,0}}}.$$

Here $p_{i,t}$ is the price of the i-th item in a given subcategory in month t ($t = 0$ represents the base month). The geometric average incorporates some substitution effect. When one product becomes more expensive relative to another similar product, people tend to buy less of the relatively more expensive product and more of the relatively cheaper product.

As discussed in a document by the Office for National Statistics (2003a), the substitution effect is related to utility theory, where people maximise their utility by spending a given budget on the right combination of goods and services. Likewise, people tend to buy sufficient quantities of goods and services to maintain the same standard of living over time, which gives a cost of living index. This type of index is highly subjective, because of the assumptions concerning utility and the substitution effect. HICP captures some of this effect between months, while RPI does not count for any substitution, except when the weights are changed annually. Estimates show that this formula effect makes the annual rate of RPI around 0.6% higher than the annual rate of UK HICP.

1.3 Inflation-linked bonds and forwards

In terms of a consumer price index, the term indexation refers to using the index to adjust an amount (denoted in some currency) to maintain its purchasing power over time. As an example, RPI was set to 100 in January 1987, indicating that £1,000 in January 1987 was equivalent to ten units ('baskets of goods and services') of the index. In December 2011, the index had increased to 239.4, as a result of which £1,000 only bought 4.177 (= £1,000/239.4) units of the index. For £1,000 to maintain its purchasing power from January 1987 to December 2011, indexation with RPI would adjust it to £2,394 (= £1,000 × 239.4/100).

In the UK, RPI has been used for wage bargaining and for price adjustment in private contracts. RPI excluding mortgage interest payments was used to target inflation by the government from 1992 until December 2003, when the UK HICP was renamed CPI and came into use instead. The index that was used for the indexation of benefits, tax credits and public service pensions was also changed from RPI to CPI in April 2011.

Inflation-linked bond issuance

An inflation-linked bond is a contract, in which the bond issuer makes predefined payments to bond holders at given dates. Each of the payments is adjusted through indexation of a given inflation index. Bond indexation dates back as far as the eighteenth century when the state of Massachusetts issued bonds linked to silver in 1742, as well as to a basket of corn, beef, wool and leather during the American Revolution in 1780. Deacon et al. (2004) mention reasons why governments issue inflation-linked bonds. After the Second World War, governments that experienced high and volatile inflation issued inflation-linked bonds as these gave investors more stable returns. These governments included those of Argentina, Brazil, Mexico and Iceland.

For countries with stable inflation, inflation-linked bonds can still provide a lower funding cost as they remove uncertainty about future inflation—the inflation risk premium. Inflation-linked debt can also serve as a diversification tool against a portfolio of non-indexed debt, and help governments assess the future inflation that is expected by market participants. Overall, this can help a government meet its inflation targets and improve the overall credibility of its monetary policy.

In 1981, the UK issued its first inflation-linked bond. Its half-yearly coupons and the payment of principal at maturity were indexed to RPI. Canada and the United States issued inflation-linked bonds in 1993 and 1997 respectively, followed by a bond issued by France in September 1998. The latter bond was linked to the French CPI excluding tobacco, as laws did not allow tobacco to be included. In October 2001, France became the first country to issue an inflation-linked bond linked to the eurozone's HICP Excluding Tobacco (HICPxT). Since then, Germany and Italy have also issued HICPxT inflation-linked bonds.

The RPI inflation-linked government bonds issued before September 2005 have an eight-month payment lag. For example, the value of the inflation index referring to mid-November 1995 was published in mid-December 1995 and used to index a coupon paid in July 1996. For the other half-yearly coupon, the value of the inflation index referring to mid-May 1996 was published in mid-June 1996 and used to index a coupon paid in January 1997. This means that when the previous coupon was paid in July 1996 the value of the next coupon, to be paid in January 1997, was known.

On the other hand, all HICPxT inflation-linked bonds, as well as all RPI inflation-linked bonds issued since September 2005, have a three-month payment lag. For example, the values of the index referring to mid-April and to mid-May 2001 were published in mid-May and in mid-June 2001 respectively and were used to index coupons paid on 1 July and 1 August 2001 respectively. A coupon paid on other days in July 2001 was then indexed through linear interpolation between the coupons on 1 July and 1 August 2001. If we assume that the value of the inflation index cannot be inferred in the period between the reference date and the publication date, there is only around a six-week lag from the

	Issue date	Maturity	Coupon	Coupon months	Market value (£ million)
Gilts	08-Feb-2006	22-Nov-2017	1.25%	May / Nov	17,329.24
	11-Jul-2007	22-Nov-2022	1.875%	May / Nov	23,910.76
	26-Apr-2006	22-Nov-2027	1.25%	May / Nov	24,059.37
	23-Nov-2011	22-Mar-2029	0.125%	Mar / Sep	8,036.13
	29-Oct-2008	22-Nov-2032	1.25%	May / Nov	19,716.95
	25-May-2011	22-Mar-2034	0.75%	Mar / Sep	11,717.34
	21-Feb-2007	22-Nov-2037	1.125%	May / Nov	19,597.44
	28-Jan-2010	22-Mar-2040	0.625%	Mar / Sep	14,792.90
	24-Jul-2009	22-Nov-2042	0.625%	May / Nov	14,881.91
	21-Nov-2007	22-Nov-2047	0.75%	May / Nov	14,234.68
	25-Sep-2009	22-Mar-2050	0.5%	Mar / Sep	14,653.47
	23-Sep-2005	22-Nov-2055	1.25%	May / Nov	19,486.51
	26-Oct-2011	22-Mar-2062	0.375%	Mar / Sep	14,100.77
Stocks	21-Feb-1985	16-Aug-2013	2.5%	Feb / Aug	21,618.00
	19-Jan-1983	26-Jul-2016	2.5%	Jan / July	27,471.42
	12-Oct-1983	16-Apr-2020	2.5%	Apr / Oct	24,474.82
	30-Dec-1986	17-Jul-2024	2.5%	Jan / July	22,874.41
	12-Jun-1992	22-Jul-2030	4.125%	Jan / July	16,408.31
	11-Jul-2002	26-Jan-2035	2%	Jan / July	19,359.62

Table 1.2: RPI-linked UK Treasury bonds in issue on 23 July 2012. Treasury gilts (top half) have a three-month indexation lag, while Treasury stocks (bottom half) have an eight-month indexation lag. (Source: UK Debt Management Office)

time the index is published to the time it is used to index the relevant coupon.

In Table 1.2 we list the RPI-linked UK Treasury bonds that were in issue on 23 July 2012. The bonds with a three-month indexation lag are called Treasury gilts, while the the bonds with an eight-month indexation lag are called Treasury stocks. The half-yearly coupons represent per cent notional value, which is appropriately indexed by RPI. The market value is as of 23 July 2012.

Inflation-linked zero-coupon bond

Stripping a bond means separating its coupon payments from the principal so that each coupon can be sold separately as a zero-coupon bond. This is mainly carried out by investment banks and dealers, often on government bonds. By stripping nominal bonds we get nominal zero-coupon bonds for different maturities T with present values at time t given by P^n_{tT}. These bonds are also called nominal discount bonds. Buying one unit of the bond at time t for P^n_{tT} units of the nominal currency pays out one unit of the nominal currency at maturity T.

By stripping inflation-linked bonds we get *inflation-linked zero-coupon bonds* for different maturities T with present values at time t given by P^i_{tT}. Buying one unit of the bond at time t for P^i_{tT} units of the nominal currency pays out I_T units of the nominal currency at maturity T. To define I_T, we account for the payment lag that is incorporated into the original bond. Therefore, I_T refers to the value of the inflation index (possibly interpolated) that was published at an earlier publication date, which for most inflation-linked bonds on RPI and HICPxT is around six weeks before the payment date T. In order to ease notation, we assume that the publication date falls on the corresponding reference payment date T. This approximation is acceptable for the short lags of six weeks, however longer time delays must be properly taken into account within a given inflation model.

Forward inflation index and swap market

An inflation index forward contract—commonly known as an *inflation-linked zero-coupon swap*—is an agreement between two parties made at time t, to exchange I_T units of currency for a fixed amount I_{tT} at time T. I_{tT} is called the *forward inflation index* and is determined at time t and quoted in the market by the annual compounded rate k_{tT} defined through the relation $I_{tT} = (1 + k_{tT})^{T-t}$. Over the last few years, the inflation-linked zero-coupon swap market has become independent from the inflation-linked bond market, with its own rules and conventions. For RPI, the main difference is that inflation-linked swaps usually trade with only a two-month payment lag instead of the three- and eight-month lags for inflation-linked bonds. This means that the value of the index referring to mid-April is published in mid-May and used to index a swap that is paid in June.

We note that the value of I_{tT} paid at time T and discounted to time t is $P^n_{tT} I_{tT}$, while the value of I_T paid at T and discounted to time t is P^i_{tT}. From the definition of I_{tT}, these present values should match to prevent arbitrage. We therefore have

$$I_{tT} = P^i_{tT} / P^n_{tT}.$$

In Figure 1.1, we look at the historical and forward inflation indices for RPI and HICPxT. We plot the historical inflation levels I_T / I_t and year-on-year inflation rates I_T / I_{T-1},

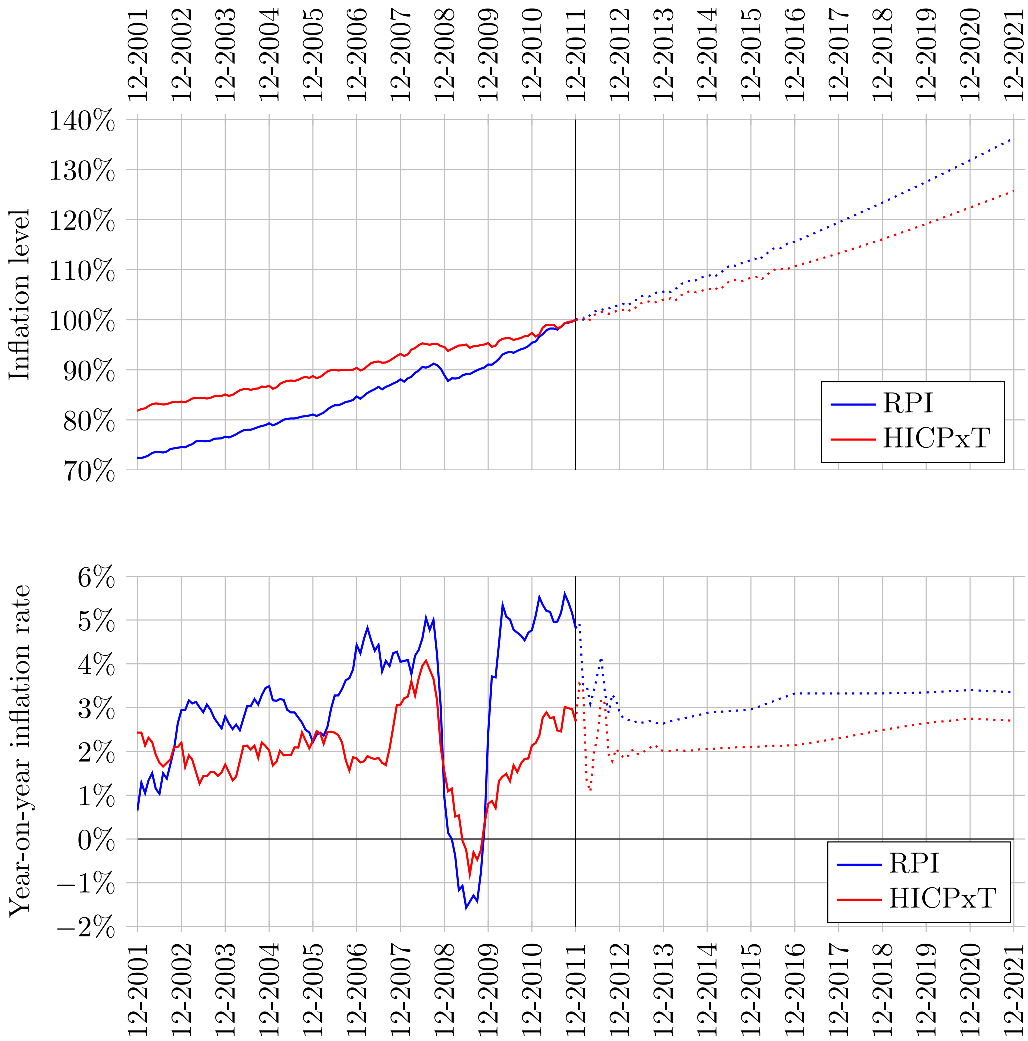

Figure 1.1: The historical and forward inflation for RPI and HICPxT. The solid lines show the historical inflation levels I_T/I_t (top) and year-on-year inflation rates $I_T/I_{T-1} - 1$ (bottom), with monthly reference dates T from December 2001 to December 2011, where the base reference date t is December 2011. The dotted lines show the forward inflation levels I_{tT}/I_t (top) and the unadjusted year-on-year inflation forward rates $I_{tT}/I_{t,T-1} - 1$ (bottom), with monthly reference dates T from December 2011 to December 2021, where the base reference date t is December 2011. (Source: UK Debt Management Office and Markit)

with monthly reference dates T from December 2001 to December 2011, where the base reference date t is December 2011. We also plot the forward inflation levels I_{tT}/I_t and the year-on-year rate of the forward inflation index $I_{tT}/I_{t,T-1} - 1$, with monthly reference dates T from December 2011 to December 2021, where the base reference date t is December 2011.

Real currency, real bond and Fisher equation

Hypothetically, we can think of the goods and services comprising the inflation index as a currency—the real currency. One unit of the real currency can then be exchanged for I_T units of the nominal currency at time T. Related to the real currency is the notion of a real discount bond. Buying one unit of the real bond at time t for P^r_{tT} units of the real currency pays out one unit of the real currency at maturity T. The nominal value of this payment at time T is I_T units of nominal currency, which at time t is worth P^i_{tT}. On the other hand, P^r_{tT} can be converted to the nominal currency at time t, giving $I_t P^r_{tT}$ units of nominal currency. These cash flows should match, giving

$$P^r_{tT} = P^i_{tT} \big/ I_t. \tag{1.1}$$

From these definitions, we can now define the instantaneous nominal forward rate f^n_{tT}, the instantaneous real forward rate f^r_{tT} and the instantaneous inflation forward rate f^i_{tT}:

$$\begin{aligned} f^n_{tT} &= -\frac{\partial}{\partial T} \ln P^n_{tT} \;, \quad P^n_{tt} = 1 && \iff && P^n_{tT} = \exp\left[-\int_t^T f^n_{ts} ds\right] , \\ f^r_{tT} &= -\frac{\partial}{\partial T} \ln P^r_{tT} \;, \quad P^r_{tt} = 1 && \iff && P^r_{tT} = \exp\left[-\int_t^T f^r_{ts} ds\right] , \\ f^i_{tT} &= \frac{\partial}{\partial T} \ln I_{tT} \;, \quad I_{tt} = I_t && \iff && I_{tT} = I_t \exp\left[\int_t^T f^i_{ts} ds\right] . \end{aligned}$$

From these relationships, we get the Fisher equation for instantaneous rates:

$$P^r_{tT} I_t = P^n_{tT} I_{tT} \qquad \Longrightarrow \qquad f^r_{tT} = f^n_{tT} - f^i_{tT}.$$

As an example we set $I_t = 100$, $T - t = 1$ year, and flat nominal rate $f^n_{tT} = 2\%$. With flat inflation rate $f^i_{tT} = -1\%$ (deflation), the real rate will be $f^r_{tT} = 2\% - (-1\%) = 3\%$ and the real discount bond will be $P^r_{tT} = \exp[-3\%] \approx 0.97$.

On the other hand if the inflation rate is $f^i_{tT} = 3\%$, we have a negative real rate $f^r_{tT} = 2\% - 3\% = -1\%$, indicating that the nominal bond depreciates in value in real terms. This gives $P^r_{tT} = \exp[-(-1\%)] \approx 1.01$, showing that one needs to buy more than one unit of the inflation basket at time t to be sure to get at least one unit of the inflation basket at a later time T.

2 Inflation-linked derivatives

In this chapter, we will look at the different inflation-linked derivatives that are traded in the market. In §2.1, we look at zero-coupon swaps that are settled after the inflation index is published. We define the payment delay convexity correction and derive a formula that shows that this correction decreases in the correlation between inflation rates and nominal interest rates. Derivatives that depend on the year-on-year inflation rate are considered in §2.2. We define the year-on-year convexity correction and show that it decreases in the autocorrelation of the inflation index. We will also look at the limited price index (LPI), a more exotic derivative that depends on the year-on-year inflation rate.

Both zero-coupon and year-on-year inflation caps and floors are regularly traded. In §2.3, we define and calculate absolute log-normal volatilities that are implied by the prices of options on the UK's Retail Prices Index (RPI) and the eurozone's Harmonised Index of Consumer Prices Excluding Tobacco (HICPxT). We find that the volatility of the inflation index increases in approximately linearly with time, while the volatility of the year-on-year rate is stable over time. For a given maturity, we observe the volatility smile and see that both the zero-coupon index and the year-on-year inflation rate have a positive curvature and a negative skew in their smiles. This indicates excess kurtosis (fat tails) and a negative skewness in their distribution.

2.1 Zero-coupon derivatives

A particular generalisation of the inflation-linked zero-coupon swap is to add a payment delay so that the payment of I_T happens at time T', later than T. The other leg payment of this swap, which is denoted by $I_{tTT'}$, is the forward value at time t of paying I_T at T'. It is important to note that $I_{tTT'}$ will not be the same as the forward inflation index I_{tT}. The difference $I_{tTT'} - I_{tT}$ is called a *payment delay convexity correction*. Changing from the T'-forward measure to the T-forward measure, we get[1]

$$\begin{aligned}I_{tTT'} &= \mathbb{E}_t^{T'}\left[I_T\right] = \mathbb{E}_t^{T}\left[I_T \frac{P^n_{TT'} / P^n_{tT'}}{P^n_{TT} / P^n_{tT}}\right] = \frac{\mathbb{E}_t^{T}\left[I_T P^n_{TT'}\right]}{P^n_{tT'} / P^n_{tT}} \\ &= \frac{\mathbb{E}_t^{T}\left[I_T\right]\mathbb{E}_t^{T}\left[P^n_{TT'}\right] + \mathrm{Cov}_t^{T}\left[I_T, P^n_{TT'}\right]}{P^n_{tT'} / P^n_{tT}} = I_{tT}\left(1 + \frac{\mathrm{Cov}_t^{T}\left[I_T, P^n_{TT'}\right]}{I_{tT}\, P^n_{tT'} / P^n_{tT}}\right).\end{aligned}$$

We see that the payment delay convexity correction decreases in the correlation between inflation rates and nominal interest rates. A positive correlation gives a negative correlation (and covariance) between the nominal discount bond $P^n_{TT'}$ and the inflation index I_T, resulting in $I_{tTT'} \leq I_{tT}$. In this case, the pay-delayed forward has a lower value than the regular forward.

An example of a swap with a payment delay is when a utility company enters into a particular swap with a bank at time T_0. In this swap, the company pays I_{T_1} / I_{T_0} at time T_1 and then, for each T_i, $i = 2, 3, \ldots, N$ (T_N being the maturity), it pays

$$\frac{I_{T_i} - I_{T_{i-1}}}{I_{T_0}}.$$

Here $T_i - T_{i-1}$ can be five years, creating a payment delay because $I_{T_{i-1}}$ fixes at T_{i-1} but is received later at time T_i. The motivation for this swap is that, instead of there being one payment I_{T_N} at maturity T_N, the payments are spread out, effectively reducing the credit risk of the company towards the bank.

Zero-coupon caps and floors

Options on the compounded inflation index—*zero-coupon inflation options*—are also regularly traded. At maturity T, a *zero-coupon inflation cap* bought at time t with strike k (an annualised rate) pays out the positive cash flow

$$\left(\frac{I_T}{I_t} - (1+k)^{T-t}\right)^+ := \max\left[\frac{I_T}{I_t} - (1+k)^{T-t}, 0\right].$$

The corresponding *zero-coupon inflation floor* has the payoff

$$\left(\frac{I_T}{I_t} - (1+k)^{T-t}\right)^- := \max\left[(1+k)^{T-t} - \frac{I_T}{I_t}, 0\right].$$

[1] Denoting $\mathbb{E}_t^{T'}[\cdot]$ the expectations in the $\mathbb{Q}^{T'}$ forward measure conditional on time t (see §3.1), we have

$$I_{tTT'} = \mathbb{E}_t^{T'}\left[I_T\right], \text{ while } I_{tT} = \mathbb{E}_t^{T}\left[I_T\right].$$

Furthermore, the covariance between some random processes A and B is defined by

$$\mathrm{Cov}_t^{T'}\left[A, B\right] := \mathbb{E}_t^{T'}\left[\left(A - \mathbb{E}_t^{T'}\left[A\right]\right)\left(B - \mathbb{E}_t^{T'}\left[B\right]\right)\right].$$

In §2.3 we will define the volatility smile corresponding to the prices of these options. It is possible that zero-coupon inflation options are paid at a later time T', giving a payment delay of $T' - T$, which creates a payment delay convexity correction as for the inflation-linked zero-coupon swap.

2.2 Year-on-year derivatives

The year-on-year inflation rate is an important financial quantity, which is calculated as the per cent change in the inflation index over a one-year period, $I_T / I_{T-1} - 1$. Governments use it as a part of their inflation target policies, where they require that the rate is stable around a certain value.

A *year-on-year inflation swaplet* is an agreement made at time t between two parties, to exchange the realised value of the year-on-year rate for a fixed amount y_{tT} at time T, with $t < T$. y_{tT} is called the *(adjusted) year-on-year inflation forward rate* and is determined at time t. On the other hand, the year-on-year increase in the forward inflation index I_{tT} is given by $I_{tT}/ I_{t,T-1} - 1$ and called the *unadjusted (or naïve) year-on-year inflation forward rate.* The difference $y_{tT} - (I_{tT}/ I_{t,T-1} - 1)$, is called the *year-on-year convexity correction* indicating that the value of year-on-year inflation forward rate is usually different from the value of the unadjusted year-on-year inflation forward rate for $t < T - 1$. To approximate this correction, we use a first-order Taylor expansion on $1 / I_{T-1}$ around $I_{t,T-1,T} = \mathbb{E}_t^T [I_{T-1}]$, and the result on the payment delay convexity correction:

$$
\begin{aligned}
y_{tT} + 1 &= \mathbb{E}_t^T \left[\frac{I_T}{I_{T-1}} \right] \approx \mathbb{E}_t^T \left[I_T \left(\frac{1}{I_{t,T-1,T'}} - \frac{1}{I_{t,T-1,T}^2} \left(I_{T-1} - I_{t,T-1,T} \right) \right) \right] \\
&= \frac{I_{tT}}{I_{t,T-1,T}} \left(2 - \frac{\mathbb{E}_t^T [I_T I_{T-1}]}{I_{tT} I_{t,T-1,T}} \right) = \frac{I_{tT}}{I_{t,T-1,T}} \left(2 - \frac{\mathbb{E}_t^T [I_T]}{I_{tT}} \frac{\mathbb{E}_t^T [I_{T-1}]}{I_{t,T-1,T}} - \frac{\mathrm{Cov}_t^T [I_T, I_{T-1}]}{I_{tT} I_{t,T-1,T}} \right) \\
&= \frac{I_{tT}}{I_{t,T-1,T}} \left(1 - \frac{\mathrm{Cov}_t^T [I_T, I_{T-1}]}{I_{tT} I_{t,T-1,T}} \right) \\
&= \frac{I_{tT}}{I_{t,T-1}} \frac{1 - \mathrm{Cov}_t^T [I_T, I_{T-1}] / (I_{tT} I_{t,T-1,T})}{1 + \mathrm{Cov}_t^{T-1} \left[I_{T-1}, P_{T-1,T}^n \right] \Big/ \left(I_{t,T-1} \, P_{tT}^n / P_{t,T-1}^n \right)}.
\end{aligned}
$$

We see that increasing the inflation index autocorrelation lowers the correction, while increasing the correlation between inflation rates and nominal interest rates increases the correction. This is because $I_T / I_{T-1} - 1$ is paid at time T, so there is a payment delay of one year from when the value of I_{T-1} is known at time $T-1$, until I_{T-1} is paid at time T, which creates a payment delay convexity correction.

The inflation index autocorrelation $\mathrm{Corr}_t^{T'}[I_{T'}, I_T]$ is typically positive, which decreases the year-on-year convexity correction. On the other hand, a positive correlation between inflation rates and nominal interest rates increases the year-on-year convexity correction. The effect from the positive autocorrelation is usually larger than the effect from the positive nominal-inflation correlation, causing the year-on-year convexity correction to be negative, which means that the adjusted year-on-year inflation forward rate is lower than the unadjusted year-on-year inflation forward rate.

A *year-on-year inflation swap* starting at time t with maturity T is a series of year-on-year inflation swaplets. At each time $s = t+1, t+2, \ldots, T$, the actual value of the year-on-year inflation rate $I_s / I_{s-1} - 1$, is swapped for a fixed rate X. This rate is determined at time t from the weighted average of the adjusted year-on-year inflation forward rates y_{ts}:

$$X = \sum_{s=t+1}^{T} y_{ts} P_{ts}^n \Bigg/ \sum_{s=t+1}^{T} P_{ts}^n.$$

The year-on-year convexity correction is measurable from the market prices of these swaps, which are quoted as a weighted average convexity correction of the difference between the adjusted and unadjusted year-on-year inflation forward rates. This value, α_T, is called the year-on-year convexity margin and is computed as follows:

$$\sum_{s=t+1}^{T} \left(\frac{I_{ts}}{I_{t,s-1}} - 1 + \alpha_T \right) P_{ts}^n = \sum_{s=t+1}^{T} y_{ts} P_{ts}^n$$

$$\Longrightarrow \quad \alpha_T = \sum_{s=t+1}^{T} \left(y_{ts} - \left(\frac{I_{ts}}{I_{t,s-1}} - 1 \right) \right) P_{ts}^n \Bigg/ \sum_{s=t+1}^{T} P_{ts}^n.$$

In Figure 2.1, we plot this convexity margin for RPI and HICPxT, as well as the adjusted and unadjusted year-on-year inflation forward rates for each index. For RPI, the gap between the two rates gradually increases, causing the year-on-year convexity margin to decrease with maturity to around minus four basis points (-0.04%) for the thirty-year maturity. This is considerably lower than for HICPxT, where there is very little difference between the adjusted and unadjusted rates. From the previous discussion on convexity corrections, one can partly explained these different characteristics of RPI and HICPxT by looking into the inflation index autocorrelation and the correlation between inflation rates and nominal interest rates.

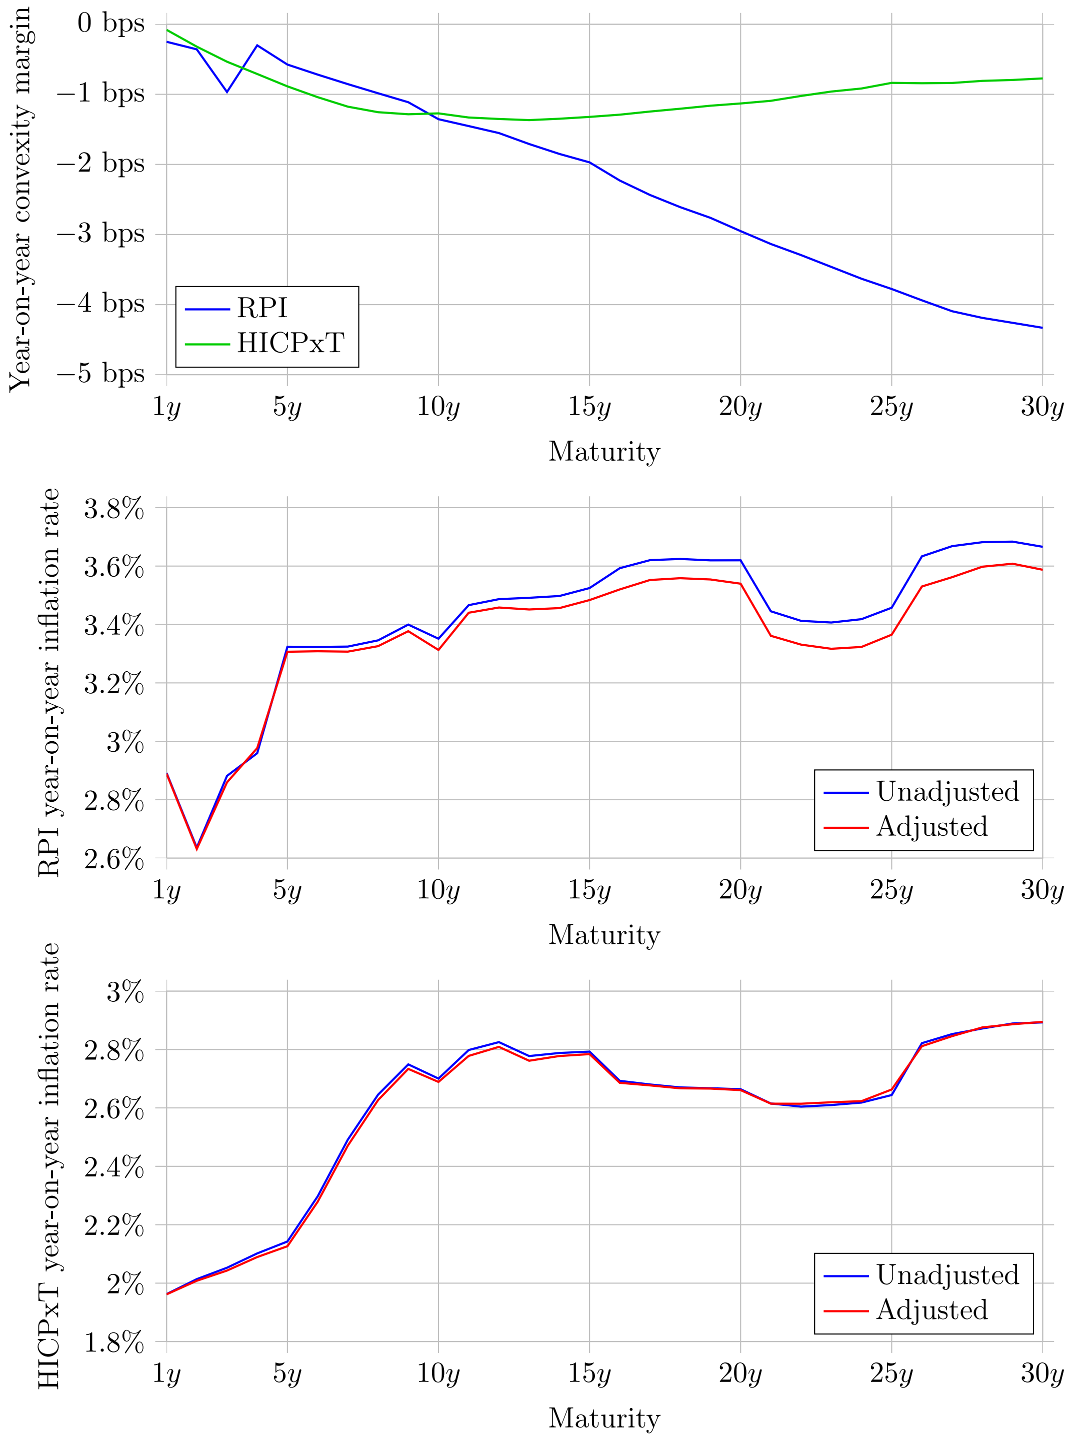

Figure 2.1: The graphs show the year-on-year convexity margins (top), and the unadjusted and the adjusted year-on-year inflation forward rates for RPI (middle) and HICPxT (bottom) on 31 January 2012. See §2.2 for definitions. (Source: Markit)

Year-on-year caps and floors

A *year-on-year inflation cap* with a strike K starting at time t with maturity T is a series of *year-on-year inflation caplets* paid out at $s = t+1, t+2, \dots, T$. For each time s, the caplet pays out

$$\left(\frac{I_s}{I_{s-1}} - 1 - K\right)^+ .$$

A *year-on-year inflation floor* consists of a series of *year-on-year inflation floorlets*

$$\left(\frac{I_s}{I_{s-1}} - 1 - K\right)^- .$$

In §2.3, we will define the volatility smile corresponding to the prices of these options.

Limited price indices

A more exotic year-on-year product is the *limited price index* (LPI), which is obtained by compounding year-on-year collar options. Each collar has a floor, f, and a cap, c, on the year-on-year inflation rate:

$$\max\left[f, \min\left[c, \frac{I_T}{I_{T-1}} - 1\right]\right].$$

Usually we have $f = 0\%$ and $c = 3\%, 5\%, +\infty$ (no cap). We denote by $P_T^{f,c}$ the value of the LPI at time T, floored at f and capped at c. For a base time t, we set $P_t^{f,c} := I_t$ and iteratively define

$$P_T^{f,c} = P_{T-1}^{f,c}\left(1 + \max\left[f, \min\left[c, \frac{I_T}{I_{T-1}} - 1\right]\right]\right),$$

for all $T = t+1, t+2, \dots$. The LPI is path-dependent on the inflation index, as $P_T^{f,c}$ is a function of I_s for $s = t, t+1, t+2, \dots, T$. We let $P_{tT}^{f,c}$ be the time t value of the LPI index at time T. The market quotes the LPI swap rate $s_{tT}^{f,c}$, which is the excess annual rate of return of LPI over the forward inflation index:

$$s_{tT}^{f,c} := \left(\frac{P_{tT}^{f,c}}{P_t^{f,c}}\right)^{(T-t)^{-1}} - \left(\frac{I_{tT}}{I_t}\right)^{(T-t)^{-1}}.$$

In Figure 2.2, we plot the annual rates of return for the forward inflation index and three different LPIs for RPI on 31 January 2012. All three LPIs have a floor at 0%, with no cap, a 5% cap, and a 3% cap respectively.

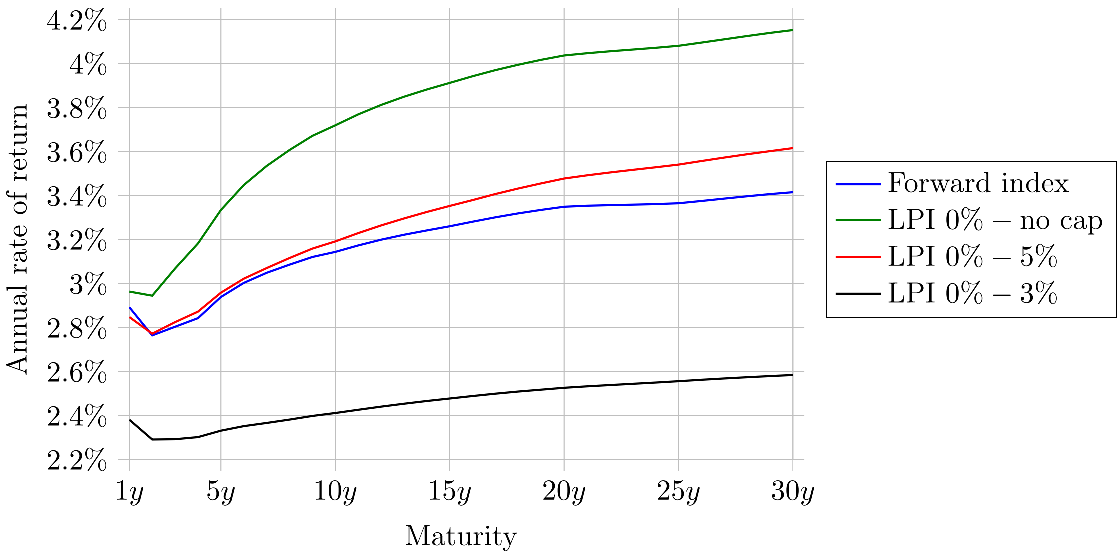

Figure 2.2: Forward inflation index and LPI annual rates of return for RPI on 31 January 2012. The LPIs have no cap, a 5% cap, and a 3% cap respectively, with all of them having a 0% floor. See §2.2 for a definition of the limited price index. (Source: Markit)

2.3 Inflation smile

Here we look at the prices of inflation-linked options in more detail—in particular, zero-coupon and year-on-year inflation caps and floors. Given an option on some financial variable with a given maturity and strike, rather than looking directly at the option's price, it is a common practice to use the implied volatility. The implied volatility arises if one models the financial variable using a Black-Scholes model (1973). The method assumes that the financial variable at a given time has a log-normal distribution with a mean equal to the corresponding forward and some log-normal volatility. For each option, this volatility is adjusted so that the expectations of the option's payoff match the forward value of the option's price.

We define the general inflation ratio $Y_{ST} := I_T/I_S$ for times S, T with $S < T$. We denote the forward value of Y_{ST} at time t by Y_{tST}. We now assume that Y_{ST} follows a log-normal distribution with a log-normal volatility σ_{tST} at time t:

$$Y_{ST} = Y_{tST} \exp\left[\sigma_{tST} Z - \tfrac{1}{2}\sigma_{tST}^2\right],$$

where Z has a standard normal distribution. A call option on Y_{ST} with a strike K pays out at maturity T the positive cash flow

$$\left(\frac{I_T}{I_S} - K\right)^+ .$$

From the Black-Scholes option formula, the forward price V^c_{tST} of this call option at time t is given by

$$\begin{aligned} V^c_{tST}[K] &= Y_{tST}\mathcal{N}\left[d + \tfrac{1}{2}\sigma_{tST}\right] - K\mathcal{N}\left[d - \tfrac{1}{2}\sigma_{tST}\right], \\ d &:= \ln\left[\frac{Y_{tST}}{K}\right] \Big/ \sigma_{tST}, \end{aligned}$$

where $\mathcal{N}[\cdot]$ is the standard normal cumulative distribution function. Given the forward price $V^c_{tST}[K]$ of a call,[2] we can find a value $\sigma_{tST}[K]$ so that, if we set $\sigma_{tST} = \sigma_{tST}[K]$ into the option formula, we recover $V^c_{tST}[K]$. We call $\sigma_{tST}[K]$ the *absolute log-normal volatility* of the option. The relationship between the traditional implied volatility $\sigma^{iv}_{tST}[K]$ and $\sigma_{tST}[K]$ is given by $\sigma_{tST}[K] = \sigma^{iv}_{tST}[K]\sqrt{T-t}$.

For zero-coupon inflation options we have $S = t$, so $Y_{ST} = Y_{tT} = I_T/I_t$. When the maturity T increases we expect the absolute volatility $\sigma_{tST}[K]$ to increase, so to normalise it we define the *zero-coupon inflation volatility*, $\sigma^{ZC}_{tT}[k]$, as

$$\sigma^{ZC}_{tT}[k] := \frac{1}{T-t}\sigma_{ttT}\left[(1+k)^{T-t}\right].$$

For year-on-year inflation options, we have $S = T-1$, so $Y_{ST} = Y_{T-1,T} = I_T/I_{T-1}$. The absolute volatility $\sigma_{t,T-1,T}[K]$ is the log-normal volatility over a one-year period from $T-1$ to T, so we do not normalise $\sigma_{t,T-1,T}[K]$ depending on the maturity T. We therefore define the *year-on-year inflation volatility*, $\sigma^{YoY}_{tT}[k]$, as

$$\sigma^{YoY}_{tT}[k] := \sigma_{t,T-1,T}[1+k].$$

Market data

We now want to see what the values of the zero-coupon and year-on-year inflation volatilities are in inflation markets by looking at the prices of zero-coupon and year-on-year inflation options. Table 2.1 shows a summary of the inflation-linked derivatives covered in this chapter and the previous chapter. In the table, we have tried to give an indication of how liquid each derivative is for RPI and HICPxT inflation markets. Liquidity of a

[2] If we have the price of a put, V^p_{tST}, with payoff $(I_T/I_S - K)^-$, the corresponding call price is obtained from the put-call parity, $V^c_{tST} = Y_{tST} - K + V^p_{tST}$.

	Zero-coupon	**Year-on-year**
Bonds	Very liquid	Undefined
Swaps	Liquid	Rare
w/payment delay	Occasional/rare	Rare
Options	Liquid on 0% floors $< 10y$ for RPI	Semiliquid for RPI
	Liquid on $\sim 0\%$ floors for HICPxT	Liquid for HICPxT
LPIs	Not defined	Occasional for RPI

Table 2.1: Summary of inflation-linked derivatives

derivative can be estimated from how frequently the derivative is traded and how narrow the bid-offer spread of its price is.

By looking at broker screens and talking to the inflation traders at Lloyds Banking Group, we got an idea on the liquidity of inflation-linked derivatives. Compared to the nominal interest rate markets, inflation markets are very illiquid. Therefore, Table 2.1 shows the relative liquidity between the inflation-linked derivatives that are traded in inflation markets.

Typically, inflation-linked swaps and options are only quoted for a small number of maturities and strikes. When looking for representative market data, we therefore use inflation-linked derivatives' consensus prices that are published at the the end of each month by Markit's Totem Service. For each quoted product, Markit provides the average price from the main market makers, as well as the range, standard deviation, kurtosis and skew of the prices' distribution.

For both RPI and HICPxT, we have market quotes for zero-coupon and year-on-year inflation swaps, floors for -2%, -1%, 0%, 1%, 2% and 3% strikes and caps for 2%, 3%, 4% and 5% strikes (6% and 7% strikes only for RPI), with annual maturities from one to thirty years. In Figure 2.3 we show the zero-coupon inflation volatilities $\sigma_{tT}^{ZC}[k]$ from forward prices of zero-coupon inflation caps and floors on 31 January 2012. In Figure 2.4 we show the year-on-year inflation volatilities $\sigma_{tT}^{YoY}[k]$ from forward prices of year-on-year inflation caplets and floorlets on 31 January 2012.

For a fixed strike, the zero-coupon and year-on-year inflation volatilities are low for maturities of up to ten years. This may be because, in the short term, the market participant's ability to make better estimations of how the inflation index will move results in lower volatilities. For higher maturities, the volatilities are high and stable, possibly because it is more difficult to estimate the movements of the inflation index further into the future. The stability of the zero-coupon inflation volatility $\sigma_{tT}^{ZC}[k]$, indicates that the volatility of I_T/I_t increases linearly with T.

For a fixed maturity T, the volatilities vary with the strike, giving the so-called *smile* of implied volatilities, which indicate positive excess kurtosis (giving the positive *curvature* of the smile) and non-zero skewness (giving the *skew* of the smile) in the distribution of the logarithm of I_T/I_t and I_T/I_{T-1}. The curvature is strong for both RPI and HICPxT, with the lowest volatilities for strikes around the forward, and increasing volatilities as the strikes move away from the forward. Low or high inflation often indicates economical uncertainty and recession, which people want protection against. This uncertainty increases the volatilities at lower and higher strikes and produces the positive curvature.

Let f be the value of the forward at a given maturity. A negative skew means that the volatility at a strike $f - m$ is higher than the volatility at the strike $f + m$ for most (if not all) values $m > 0$. A floor at the zero-percent inflation rate is implicit in many inflation-linked contracts, making the lower strikes more expensive than the higher ones and giving a negative skew. The skew is more negative for RPI than for HICPxT, because of a demand for LPIs on RPI. The demand comes from a requirement for pension funds in the UK, where a portion of pensioners' contributions needs to increase at the LPI rate. Pension funds therefore take a long position in LPIs, which pushes up the prices of year-on-year floors with a 0% strike (long position) and pushes down the prices of year-on-year caps with 3% and 5% strikes (short position), hence giving a large negative skew.

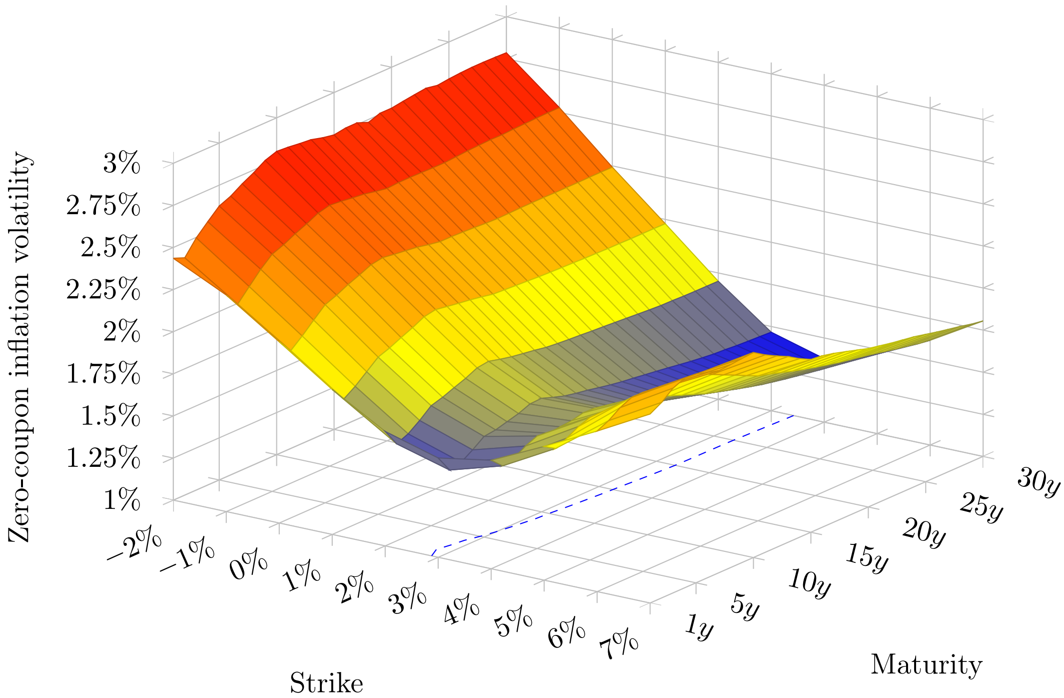

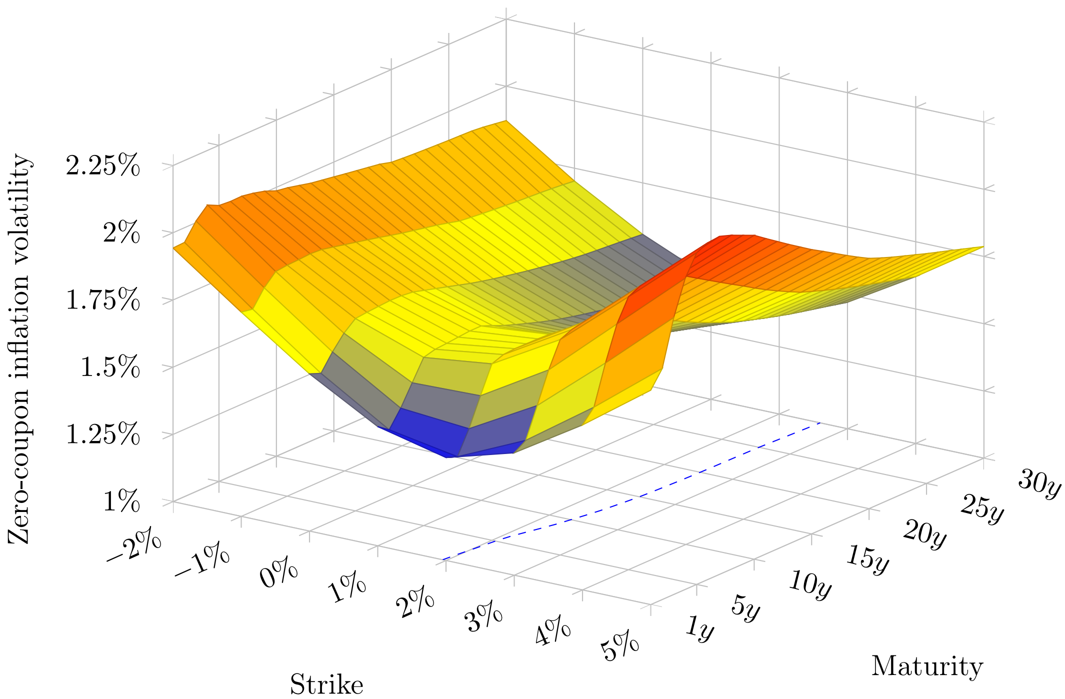

Figure 2.3: As defined in §2.3, the graphs show the zero-coupon inflation volatility surfaces for RPI (top) and HICPxT (bottom) on 31 January 2012. The blue and broken curve represents the strikes that are equal to the annual rate of the forward inflation index. (Source: Markit)

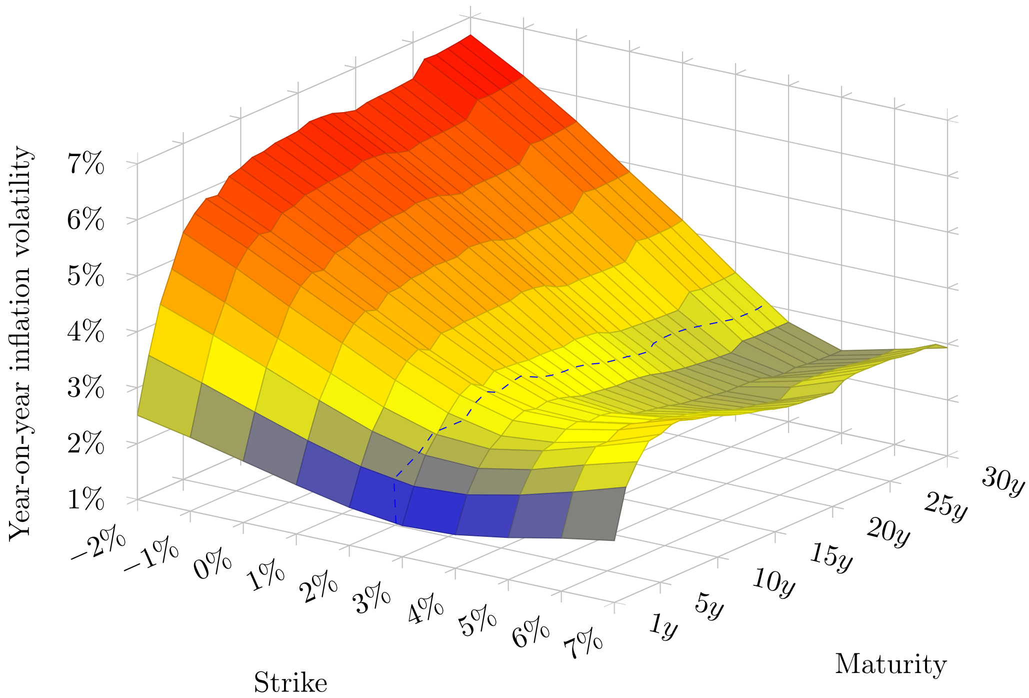

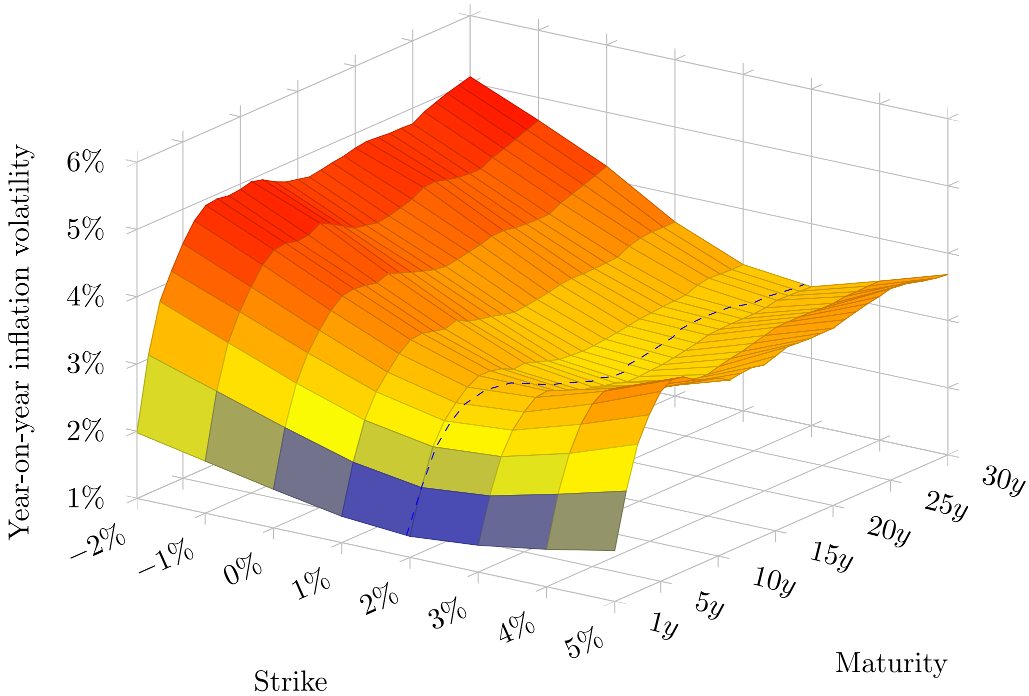

Figure 2.4: As defined in §2.3, the graphs show the year-on-year inflation volatility surfaces for RPI (top) and HICPxT (bottom) on 31 January 2012. The blue and broken curve represents the volatility at strikes that are equal to the year-on-year inflation forward rate. (Source: Markit)

3 Pricing framework

In this chapter, we set up a framework for pricing inflation-linked derivatives. We will use it to specify the different inflation models in the next chapter. In §3.1, we assume the existence of a pricing kernel in a given currency. In the framework, we can represent any price process in a simple way, making it easy to change between equivalent martingale measures. In §3.2, we detail the foreign-exchange analogy for inflation. This analogy builds on the notion of a real currency and can be used to model inflation.

3.1 Martingale pricing

We specify a real-world probability space $(\Omega, \mathcal{F}, \mathbb{P})$ with a filtration $\mathcal{F}_t$ that is generated by a multidimensional $\mathbb{P}$-Brownian motion W_t.[1] Given some currency c, we assume the existence of a strictly positive pricing kernel π_t^c as described in Duffie (2001). We have $\pi_0^c = 1$. The pricing kernel has the property that if S_t^c is the price process of some non-dividend paying asset S denoted in units of the currency c, then the process M_t^{Sc} defined by

$$M_t^{Sc} := \pi_t^c S_t^c, \tag{3.1}$$

is a $\mathbb{P}$-martingale. From the martingale representation theorem—see Karatzas and Shreve (1991)— we can write the stochastic differential equation—the SDE—for M_t^{Sc} as

$$dM_t^{Sc} = \theta_t^{Sc} dW_t, \qquad M^{Sc} = S_0^c, \tag{3.2}$$

where θ_t^{Sc} is a multidimensional $\mathcal{F}_t$-measurable process. We assume the existence of a strictly positive risk-free bank account B_t^c. Then there exists a one-dimensional strictly

[1] Here $\mathbb{P}$ is the real-world (historical) measure and not any risk-neutral (equivalent martingale) measure. We also note that, as the filtration $\mathcal{F}_t$ is generated by a continuous Brownian motion, any $\mathcal{F}_t$-adapted martingale will be continuous. In particular, it will not have any jumps.

positive risk-free short-rate process r_t^c so that

$$\frac{dB_t^c}{B_t^c} = r_t^c dt. \tag{3.3}$$

From (3.1), it holds that $\pi_t^c B_t^c$ is a strictly positive $\mathbb{P}$-martingale, so, from (3.2), there exists a multidimensional process λ_t^c so that

$$\frac{d\left(\pi_t^c B_t^c\right)}{\pi_t^c B_t^c} = -\lambda_t^c dW_t. \tag{3.4}$$

Using Itō calculus[2] on (3.3) and (3.4), the SDE for the pricing kernel π_t^c is given by

$$\frac{d\pi_t^c}{\pi_t^c} = \frac{d\left(\pi_t^c B_t^c / B_t^c\right)}{\pi_t^c B_t^c / B_t^c} = -r_t^c dt - \lambda_t^c dW_t. \tag{3.5}$$

From (3.2) and (3.5), we get in a similar way the SDE for the price process S_t^c:

$$dS_t^c = d\left(\frac{M_t^{S^c}}{\pi_t^c}\right) = \left(r_t^c S_t^c + \lambda_t^c \psi_t^{S^c}\right) dt + \psi_t^{S^c} dW_t.$$

Here $\psi_t^{S^c} = \theta_t^{S^c}/\pi_t^c + \lambda_t^c S_t^c$ is the absolute volatility of S_t^c. If S_t^c is strictly positive we let $\sigma_t^{S^c}$ be the strictly positive relative volatility of S_t^c with $\psi_t^{S^c} = \sigma_t^{S^c} S_t^c$ and we write

$$\frac{dS_t^c}{S_t^c} = \left(r_t^c + \lambda_t^c \sigma_t^{S^c}\right) dt + \sigma_t^{S^c} dW_t. \tag{3.6}$$

We can interpret λ_t^c as the market risk premium for the currency c.

Changing the measure

Let $\mathbb{Q}^S$ be the measure with a numeraire S^c—the price process of S in currency c. This means that V_t^c / S_t^c is a $\mathbb{Q}^S$-martingale for all strictly positive non-dividend paying assets V, with its price process in units of the currency c denoted by V_t^c. Examples of a particular measure is the bank-account-risk-neutral measure, where we set $S_t^c = B_t^c$, and the U-forward measure, where we set $S_t^c = P_{tU}^c$ for some time U.

[2] For two Itō processes X and Y, with Y strictly positive, we have $d\left(1/Y_t\right) = -dY_t/Y_t^2 + d\left\langle Y\right\rangle_t / Y_t^3$ where $\left\langle Y\right\rangle_t$ is the quadratic variation process of Y—see Karatzas and Shreve (1991). The SDE for the ratio of X and Y, is then given by

$$\begin{aligned} d\frac{X_t}{Y_t} = \frac{1}{Y_t}dX_t + X_t d\frac{1}{Y_t} + d\left\langle X, \frac{1}{Y_t}\right\rangle_t &= \frac{1}{Y_t}\left(dX_t - X_t\frac{dY_t}{Y_t} + X_t\frac{d\left\langle Y\right\rangle_t}{Y_t^2} - \frac{d\left\langle X,Y\right\rangle_t}{Y_t}\right) \\ &= \frac{X_t}{Y_t}\left(\frac{dX_t}{X_t} - \frac{dY_t}{Y_t} + \frac{d\left\langle Y\right\rangle_t}{Y_t^2} - \frac{d\left\langle X,Y\right\rangle_t}{X_t Y_t}\right), \end{aligned}$$

where in the last step we assumed that X is strictly positive. Here $\left\langle X,Y\right\rangle_t$ is the cross-variation process of X and Y.

From (3.1), we have that $\pi_t^c S_t^c$ is a $\mathbb{P}$-martingale. Using (3.5) and (3.6), the change-of-measure process from $\mathbb{P}$ to $\mathbb{Q}^S$, conditional on the filtration $\mathcal{F}_t$ at time t, is given by

$$\left.\frac{d\mathbb{Q}^S}{d\mathbb{P}}\right|_{\mathcal{F}_t} = \frac{S_t^c/S_0^c}{\pi_0^c/\pi_t^c} = \frac{\pi_t^c S_t^c}{\pi_0^c S_0^c} = \exp\left[\int_0^t \left(\sigma_s^{S^c} - \lambda_s^c\right) dW_s - \tfrac{1}{2}\int_0^t \left(\sigma_s^{S^c} - \lambda_s^c\right)^2 ds\right].$$

From the Girsanov theorem, a $\mathbb{Q}^S$-Brownian motion W_t^S is then specified by

$$W_t^S := W_t - \int_0^t \left(\sigma_s^{S^c} - \lambda_s^c\right) ds. \tag{3.7}$$

We can express V_t^c in the measure $\mathbb{P}$ in the same way as the asset S in (3.6), with

$$\frac{dV_t^c}{V_t^c} = \left(r_t^c + \lambda_t^c \sigma_t^{V^c}\right) dt + \sigma_t^{V^c} dW_t,$$

with the relative volatility of V_t^c given by a multidimensional process $\sigma_t^{V^c}$. Using the change-of-measure from $\mathbb{P}$ to $\mathbb{Q}^S$ in (3.7), under $\mathbb{Q}^S$ we have

$$\begin{aligned}
\frac{dS_t^c}{S_t^c} &= \left(r_t^c + \left(\sigma_t^{S^c}\right)^2\right) dt + \sigma_t^{S^c} dW_t^S,\\
\frac{dV_t^c}{V_t^c} &= \left(r_t^c + \sigma_t^{S^c}\sigma_t^{V^c}\right) dt + \sigma_t^{V^c} dW_t^S.
\end{aligned}$$

Now we define the inverse numeraire process $D_t^{S^c} := S_0^c/S_t^c$. Using Itō calculus, we get

$$\frac{d\left(D_t^{S^c} V_t^{V^c}\right)}{D_t^{S^c} V_t^{V^c}} = \left(\sigma_t^{V^c} - \sigma_t^{S^c}\right) dW_t^S,$$

confirming the fact that $D_t^{S^c} V_t^c = V_t^c/\, S_t^c \cdot S_0^c$ is indeed a $\mathbb{Q}^S$-martingale.

3.2 Nominal currency and inflation

Now we look at a particular nominal currency n. Everything in §3.1 now holds for the nominal currency if we put $c = n$. For example, the nominal pricing kernel π_t^n is given by

$$\frac{d\pi_t^n}{\pi_t^n} = -r_t^n dt - \lambda_t^n dW_t, \tag{3.8}$$

where r_t^n is the nominal short-rate process and λ_t^n is the nominal market risk premium process. We introduced various inflation-linked financial quantities in §1.3. We would like to express these quantities in terms of the measure $\mathbb{P}$ by using the martingale property in (3.1). To do this, we will use the expectations under $\mathbb{P}$ conditional on $\mathcal{F}_t$, $\mathbb{E}_t\left[\cdot\right] := \mathbb{E}\left[\cdot|\mathcal{F}_t\right]$.

The value of a nominal discount bond, which pays one unit of nominal currency at time T, is P_{tT} at time t. From (3.1), we know that $\pi^n_t P^n_{tT}$ is a $\mathbb{P}^n$-martingale, so we get

$$\pi^n_t P^n_{tT} = \mathbb{E}_t\left[\pi^n_T P^n_{TT}\right] = \mathbb{E}_t\left[\pi^n_T\right].$$

I_T is the value of the inflation index observed at time T. The inflation-linked bond P^i_{tT} describes the value at time t of a cash flow paying I_T units of nominal currency at time T. We therefore have

$$\pi^n_t P^i_{tT} = \mathbb{E}_t\left[\pi^n_T P^i_{TT}\right] = \mathbb{E}_t\left[\pi^n_T I_T\right]. \tag{3.9}$$

The value of the forward inflation index, I_{tT}, is determined at time t (it is $\mathcal{F}_t$-measurable) and is exchanged for I_T units of currency at time T. The values of these two cash flows match at time t, i.e. $\mathbb{E}_t\left[\pi^n_T I_{tT}\right] = \mathbb{E}_t\left[\pi^n_T I_T\right]$, which gives us the relation

$$I_{tT} = \frac{\mathbb{E}_t\left[\pi^n_T I_T\right]}{\mathbb{E}_t\left[\pi^n_T\right]} = \frac{P^i_{tT}}{P^n_{tT}}.$$

Foreign-exchange analogy and real currency

As in §1.3, the real discount bond P^r_{tT} can be thought of as the value at time t of one unit of goods and services comprising the inflation index at T—that is, one unit of the real currency r. From (1.1), we have the relation

$$P^i_{tT} = I_t P^r_{tT}. \tag{3.10}$$

We can interpret the inflation index I_t as the exchange rate from the real currency to the nominal currency. The forward inflation index I_{tT} can then be interpreted as the number of units of nominal currency agreed upon at time t, for the exchange of one unit of real currency at time T. Inserting (3.10) into (3.9) gives

$$\pi^n_t I_t P^r_{tT} = \mathbb{E}_t\left[\pi^n_T I_T\right]. \tag{3.11}$$

Following Hughston (1998), we define π^r_t with

$$\pi^r_t := \pi^n_t I_t. \tag{3.12}$$

π^r_t is strictly positive as both π^n_t and I_t are strictly positive. For some one-dimensional process r^r_t and for some multidimensional process λ^r_t, we let π^r_t be given by

$$\frac{d\pi^r_t}{\pi^r_t} := -r^r_t dt - \lambda^r_t dW_t, \qquad \pi^r_0 = I_0 \tag{3.13}$$

$$\implies \quad \pi^r_t = I_0 \exp\left[\int_0^t \left(-r^r_s - (\lambda^r_s)^2\right) ds - \int_0^t \lambda^r_s dW_s\right].$$

From (3.11) and (3.12), we can write $\pi_t^r P_{tT}^r = \mathbb{E}_t[\pi_T^r]$ so $\pi_t^r P_{tT}^r$ is a $\mathbb{P}$-martingale. We then have

$$P_{tT}^r = \mathbb{E}_t\left[\frac{\pi_T^r}{\pi_t^r}\right] = \mathbb{E}_t\left[\exp\left[-\int_t^T r_s^r ds\right]\exp\left[-\int_t^T \lambda_s^r dW_s - \int_t^T (\lambda_s^r)^2\, ds\right]\right].$$

Taking the logarithm and differentiating with respect to T, we have[3]

$$r_t^r = -\left.\frac{\partial \ln P_{tT}^r}{\partial T}\right|_{T=t} = -\frac{1}{P_{tT}^r}\left.\frac{\partial P_{tT}^r}{\partial T}\right|_{T=t},$$

so we interpret the process r_t^r as the risk-free real short rate.

Given some strictly positive non-dividend paying continuous asset S, we can express its price process in units of nominal currency with S_t^n and in units of real currency with $S_t^r := S_t^n/I_t$. Now it holds that $\pi_t^r S_t^r = \pi_t^n I_t \cdot S_t^n/I_t = \pi_t^n S_t^n$ and $\pi_t^n S_t^n$ is a $\mathbb{P}$-martingale, so it follows that $\pi_t^r S_t^r$ is a $\mathbb{P}$-martingale. As in (3.6), it therefore follows that there exists a multidimensional volatility process $\sigma_t^{S^r}$ so that we have

$$\frac{dS_t^r}{S_t^r} = \left[r_t^r + \lambda_t^r \sigma_t^{S^r}\right] dt + \sigma_t^{S^r} dW_t,$$

so λ_t^r can be interpreted as the real market risk premium. We have therefore established that everything in §3.1 holds for the real currency by setting $c = r$, where the process π_t^r satisfies the properties of the real pricing kernel.

Writing (3.12) as $I_t = \pi_t^r/\pi_t^n$, we deduce from (3.8) and (3.13) that

$$\frac{dI_t}{I_t} = \left[r_t^n - r_t^r + \lambda_t^n\left(\lambda_t^n - \lambda_t^r\right)\right] dt + \left(\lambda_t^n - \lambda_t^r\right) dW_t = \left[r_t^i + \lambda_t^n \sigma_t^i\right] dt + \sigma_t^i dW_t,$$

where the inflation short rate r_t^i is the difference of the nominal and real short rates, $r_t^i := r_t^n - r_t^r$, and the multidimensional inflation index volatility process σ_t^i is the difference of the nominal and real risk-premium processes, $\sigma_t^i := \lambda_t^n - \lambda_t^r$.

[3] Let the measure $\mathbb{Q}$ be defined by the following change-of-measure process from $\mathbb{P}$ to $\mathbb{Q}$:

$$\left.\frac{d\mathbb{Q}}{d\mathbb{P}}\right|_{\mathcal{F}_t} = \exp\left[-\int_0^t \lambda_s^r dW_s - \tfrac{1}{2}\int_0^t (\lambda_s^r)^2\, ds\right], \qquad \text{which gives} \qquad P_{tT}^r = \mathbb{E}_t^{\mathbb{Q}}\left[\exp\left[-\int_t^T r_s^r ds\right]\right],$$

where we denote the expectations in the measure $\mathbb{Q}$ conditional on $\mathcal{F}_t$ as $\mathbb{E}_t^{\mathbb{Q}}[\cdot]$. Assuming that $r_T^r\, \mathrm{e}^{-\int_t^T r_s^r ds}$ is integrable and using the dominated convergence theorem, we get

$$\begin{aligned}
\frac{\partial \ln P_{tT}^r}{\partial T} &= \frac{1}{P_{tT}^r}\frac{\partial P_{tT}^r}{\partial T} = \frac{1}{P_{tT}^r}\lim_{\epsilon\searrow 0}\frac{P_{t,T+\epsilon}^r - P_{tT}^r}{\epsilon}\\
&= \frac{1}{P_{tT}^r}\lim_{\epsilon\searrow 0}\frac{1}{\epsilon}\mathbb{E}_t^{\mathbb{Q}}\left[\exp\left[-\int_t^T r_s^r ds\right]\left(\exp\left[-\int_T^{T+\epsilon} r_s^r ds\right] - 1\right)\right]\\
&= \frac{1}{P_{tT}^r}\lim_{\epsilon\searrow 0}\frac{1}{\epsilon}\mathbb{E}_t^{\mathbb{Q}}\left[\exp\left[-\int_t^T r_s^r ds\right]\left(-\epsilon r_T^r + \mathcal{O}\left[\epsilon^2\right]\right)\right]\\
&= \frac{-1}{P_{tT}^r}\mathbb{E}_t^{\mathbb{Q}}\left[\exp\left[-\int_t^T r_s^r ds\right] r_T^r\right] \xrightarrow[T\searrow t]{} -r_t^r.
\end{aligned}$$

From real to nominal

As was detailed in §3.1, given some strictly positive continuous asset X with X_t^n its asset price process in units of nominal currency, we construct a measure $\mathbb{Q}^X$ (with corresponding expectations noted by $\mathbb{E}^X$) associated with the inverse numeraire $D_t^{X^n} = X_0^n / X_t^n$. In the same way, given some strictly positive continuous asset Y with Y_t^r its asset price process in units of real currency, we construct a measure $\mathbb{Q}^Y$ (with corresponding expectations noted by $\mathbb{E}^Y$) associated with the inverse numeraire $D_t^{Y^r} = Y_0^r / Y_t^r$.

Now let V be some strictly positive continuous asset with V_t^r its asset price process in units of the real currency. Much as with the reasoning of Trovato et al. (2009), from no-arbitrage arguments we know that the value of the price process V_T^r discounted under the measure $\mathbb{Q}^Y$ back to time t and then converted back to units of the nominal currency at I_t is the same as the value of converting V_T^r at time T at the exchange rate I_T to units of nominal currency and then discounting it back to time t under the measure $\mathbb{Q}^X$. This also follows from the observation that $D_t^{Y^r} V_t^r$ is a $\mathbb{Q}^Y$-martingale and $D_t^{X^n} I_t V_t^r$ is a $\mathbb{Q}^X$-martingale:

$$\left.\begin{aligned} \mathbb{E}_t^Y \left[D_T^{Y^r} V_T^r\right] &= D_t^{Y^r} V_t^r \\ \mathbb{E}_t^X \left[D_T^{X^n} I_T V_T^r\right] &= D_t^{X^n} I_t V_t^r \end{aligned}\right\} \quad \Longrightarrow \quad \frac{\mathbb{E}_t^Y \left[D_T^{Y^r} V_T^r\right]}{D_t^{Y^r}} = V_t^r = \frac{\mathbb{E}_t^X \left[D_T^{X^n} I_T V_T^r\right]}{I_t D_t^{X^n}}.$$

Hence we get

$$\begin{aligned} \mathbb{E}_t^Y \left[D_T^{Y^r} V_T^r\right] &= \frac{D_t^{Y^r}}{I_t D_t^{X^n}} \mathbb{E}_t^X \left[D_T^{X^n} I_T V_T^r\right] \\ &= \frac{D_t^{Y^r}}{I_t D_t^{X^n}} \mathbb{E}_t^X \left[\frac{I_T D_T^{X^n}}{D_T^{Y^r}} D_T^{Y^r} V_T^r\right] = \mathbb{E}_t^X \left[\frac{m_T}{m_t} D_T^{Y^r} V_T^r\right], \end{aligned} \tag{3.14}$$

where we have defined the strictly positive process

$$m_t := I_t \frac{D_t^{X^n}}{D_t^{Y^r}}. \tag{3.15}$$

Choosing $V_T^r = 1 \big/ D_T^{Y^r}$ we get from (3.14) that

$$1 = \mathbb{E}_t^Y [1] = \mathbb{E}_t^X \left[\frac{m_T}{m_t}\right] \quad \Longrightarrow \quad m_t = \mathbb{E}_t^X [m_T],$$

so m_t is a $\mathbb{Q}^X$-martingale. Therefore m_t defines the change-of-measure process from the measure $\mathbb{Q}^X$ to the measure $\mathbb{Q}^Y$. In particular, the process for the inflation index I_t is given by

$$I_t = m_t \frac{D_t^{Y^r}}{D_t^{X^n}}.$$

4 Literature review

In this chapter, we review models that are used for pricing inflation-linked derivatives. These inflation models can be split into three types: foreign-exchange analogy inflation models, inflation index models and year-on-year inflation models.

In §4.1, we look at inflation models that use this foreign-exchange analogy—the Jarrow-Yıldırım inflation model—and a particular multidimensional extension of that model. In §4.2, we look at market models that directly specify the forward inflation index. These include extensions to stochastic volatility models. In §4.3, we look at inflation models, which specify a process for the year-on-year rate. In §4.4, we review the class of quadratic Gaussian models, which incorporate curvature and skew in the volatility smile.

In §4.5, we provide a summary of the models discussed in the chapter.

4.1 Jarrow-Yıldırım inflation model and extensions

In previous chapters, we have seen that, from looking at inflation-linked bonds, we can define a real currency. The inflation index can then be seen as the exchange rate from the real currency to the nominal currency. It is therefore possible to adapt any foreign-exchange model and use it to model inflation. Jarrow and Yıldırım (2003) use the modelling framework by Heath et al. (1992) and define the nominal and real forward rate processes and the inflation index in the real measure $\mathbb{P}$ with

$$\begin{aligned} df_{tT}^{n} &:= \alpha_{tT}^{n}dt + \xi_{tT}^{n}dW_t^{\mathbb{P},n}, \\ df_{tT}^{r} &:= \alpha_{tT}^{r}dt + \xi_{tT}^{r}dW_t^{\mathbb{P},r}, \\ dI_t/\, I_t &:= \mu_t^{I}dt + \sigma^{I}dW_t^{\mathbb{P},I}. \end{aligned}$$

Here $\left(W_t^{\mathbb{P},n}, W_t^{\mathbb{P},r}, W_t^{\mathbb{P},I}\right)$ are $\mathbb{P}$-Brownian motions with correlations ρ^{nr}, ρ^{nI} and ρ^{rI}, the drifts α_{tT}^{r}, α_{tT}^{r} and μ_t^{I} are t-adapted processes, the volatilities ξ_{tT}^{n} and ξ_{tT}^{r} are deterministic functions and the volatility σ^I is a positive constant. Mercurio (2005) summarises the study by Jarrow and Yıldırım (2003). Jarrow and Yıldırım assume that the forward rate volatilities are time-homogeneous:

$$\xi^n_{tT} := \sigma^n \exp\left[-a^n (T-t)\right],$$
$$\xi^r_{tT} := \sigma^r \exp\left[-a^r (T-t)\right],$$

where σ^n, σ^r, a^n and a^r are positive constants. Denoting the nominal and real instantaneous short rates with $r^n_t = f^n_{tt}$ and $r^r_t = f^r_{tt}$ respectively, they then prove that, under the nominal risk-neutral measure $\mathbb{Q}^*$ the following holds,

$$dr^n_t = \left(\nu^n_t - a^n r^n_t\right) dt + \sigma^n dW^{*,n}_t,$$
$$dr^r_t = \left(\nu^r_t - \rho^{rI}\sigma^r\sigma^I - a^r r^r_t\right) dt + \sigma^r dW^{*,r}_t,$$
$$dI_t/\, I_t = \left(r^n_t - r^r_t\right) dt + \sigma^I dW^{*,I}_t.$$

Here $\left(W^{*,n}_t, W^{*,r}_t, W^{*,I}_t\right)$ are $\mathbb{Q}^*$-Brownian motions with correlations ρ^{nr}, ρ^{nI} and ρ^{rI}, and ν^n_t and ν^r_t are deterministic functions that fit the initial term structure of nominal and real interest rates respectively. The processes r^n_t and r^r_t follow the Vašíček short-rate equation (1977). An extension of the model is to make the parameters time-dependent. This gives the Hull-White equation (1990, 1993), which, for the nominal short-rate process, can be written as

$$dr^n_t := \left(\nu^n_t - a^n_t r^n_t\right) dt + \sigma^n_t dW^{*,n}_t. \tag{4.1}$$

Let $\mathbb{Q}^n$ be the T-forward measure in the nominal currency and with an inverse numeraire $D^n_t := P^n_{0T}/\, P^n_{tT}$ where P^n_{tT} is the value at time t of the nominal discount bond with maturity T. Using the Hull-White equation for r^n_t, in Appendix A, we follow calculations by Hughston (1994), which show that D^n_t is a log-linear Gaussian process under $\mathbb{Q}^n$:

$$D^n_t = P^n_{0t} \exp\left[-\phi^n_t x_t - \tfrac{1}{2}\left(\phi^n_t\right)^2 G_t\right]. \tag{4.2}$$

Here x_t is a Gaussian process with variance G_t and driven by a $\mathbb{Q}^n$-Brownian motion W_t:

$$x_t = \int_0^t \sigma_s dW_s, \qquad G_t = \int_0^t \sigma^2_s ds, \qquad W_t = W^*_t + \int_0^t \sigma_s \left(\phi_T - \phi_s\right) ds.$$

Furthermore, we have the relations

$$\sigma_s = \sigma^n_s \exp\left[\int_0^s a^n_u du\right],$$
$$\phi^n_t = \int_t^U \exp\left[-\int_0^s a^n_u du\right] ds.$$

We note that D_t^n only depends on x at time t—that is, D_t^n is a Markovian function of x_t.

A multidimensional version of the Gaussian driving process x_t is given by

$$x_t := \int_0^t \operatorname{diag}[\sigma_s]\, dW_s. \tag{4.3}$$

Here σ_s is a multidimensional deterministic process and W_t is a multidimensional $\mathbb{Q}^n$-Brownian motion with correlation matrix ρ_t. Therefore, x_t has a zero mean under $\mathbb{Q}^n$ and a covariance matrix G_t with entries $G_t^{ij} = \int_0^t \rho_s^{ij} \sigma_s^i \sigma_s^j ds$. Gretarsson (2009) and Trovato et al. (2009) describe a multidimensional extension on (4.2) by defining an inflation model using an inverse numeraire D_t^n in the nominal currency, an inverse numeraire $D_t^r := P_{0T}^r / P_{tT}^r$ in the real currency, and a change-of-measure process m_t between the nominal and the real currency as in (3.15). The model is now defined in the following way:[1]

$$\begin{aligned}
D_t^n &:= P_{0t}^n \exp\left[-\phi_t^n \cdot x_t - \tfrac{1}{2}\phi_t^n \cdot G_t \phi_t^n\right], \\
D_t^r &:= P_{0t}^r \exp\left[-\phi_t^r \cdot \widetilde{x}_t - \tfrac{1}{2}\phi_t^r \cdot G_t \phi_t^r\right], \\
m_t &:= I_0 \exp\left[-\phi_t^m \cdot x_t - \tfrac{1}{2}\phi_t^m \cdot G_t \phi_t^m\right].
\end{aligned}$$

Here ϕ_t^n, ϕ_t^r and ϕ_t^m are multidimensional deterministic processes, compared to a one-dimensional process as in (4.2). $\widetilde{x}_t$ is a multidimensional Gaussian process with mean zero in $\mathbb{Q}^r$. By changing the measure from $\mathbb{Q}^r$ to $\mathbb{Q}^n$, we get the relation $\widetilde{x}_t = x_t + G_t \phi_t^m$, and the forward inflation index I_{tT} is then given by

$$I_{tT} = I_{0T} \exp\left[\phi_T^i \cdot x_t - \tfrac{1}{2}\phi_T^i \cdot G_t \phi_T^i + \phi_T^i \cdot G_t \phi_T^n\right], \tag{4.4}$$

where we have defined the term volatility for the inflation index with $\phi_t^i := \phi_t^n - \phi_t^r - \phi_t^m$. This model gives simple, closed-form pricing formulas for zero-coupon and year-on-year inflation caps and floors, as well as having the advantage of being easily extended to higher dimensions.

The rates in the models in this sections are Gaussian, so the models do not give curvature and skew in the inflation volatility smiles. Furthermore, as noted by Jäckel and Bonneton (2010), there are no liquid market prices for derivatives on the real rate. With the inflation models that build on the foreign-exchange analogy it can therefore prove difficult to obtain values for the volatility parameters of the real rate and its correlation

[1] For a vector x and a vector or a matrix A, the notation $A \cdot x$ is the same as $A^\intercal x$ where $A^\intercal$ is the transpose of A.

to inflation rates and to nominal interest rates. The inflation models in the next two sections get around this problem by modelling the dynamics of the inflation index without modelling a process for the real rate. These models have more intuitive parameterisation, which makes calibration to market data easier. Extensions of these models can also give a better fit to the market inflation volatility smiles.

4.2 Inflation index market model with stochastic volatility

Belgrade et al. (2004) use the market model framework by Brace et al. (1997) to model inflation without specifying a process for the real rate. The equation for the nominal discount bond is given by

$$\frac{dP^n_{tT}}{P^n_{tT}} := r^n_t dt + \Gamma^n_{tT} dW^n_t$$

and the forward inflation index is modelled on a discrete tenor structure $\{T_i\}_{i\in\mathbb{N}}$ by

$$\frac{dI_{tT_i}}{I_{tT_i}} := \mu^I_{tT_i} dt + \sigma^I_{tT_i} dW^{T_i}_t,$$

with Γ^n_{tT}, $\mu^I_{tT_i}$ and $\sigma^I_{tT_i}$ some deterministic functions. Here $\left(W^n, \{W^{T_i}\}_{i\in\mathbb{N}}\right)$ are Brownian motions in the risk-neutral measure $\mathbb{Q}^n$ with correlations ρ^{nT_i} and $\rho^{T_iT_j}$.

Belgrade et al. (2004) note that inflation-linked options in this framework can be priced with the Black-Scholes formula. For a general form of σ_{tT_i}, they provide the variance of the inflation index. They also show the relation between the year-on-year inflation volatilities and the zero-coupon inflation volatilities. They then look at a particular form of σ_{tT_i} including a constant volatility and a time-homogeneous volatility (Hull-White).

The model by Belgrade et al. does not give skew and curvature in the inflation volatility smile. Mercurio and Moreni (2006) therefore extend the inflation market model model so that the inflation index has a stochastic-volatility factor that is of the same form as in the Heston model (1993). For a given tenor structure $\{T_i\}_{i\in\mathbb{N}}$, the nominal forward rates are defined with

$$F^n_{tT_i} := \frac{P^n_{tT_{i-1}}/P^n_{tT_i} - 1}{T_i - T_{i-1}}. \tag{4.5}$$

Mercurio and Moreni specify the model in the following way for some $i \in \mathbb{N}$:

$$\begin{aligned} dF^n_{tT_i}/F^n_{tT_i} &:= (\ldots)\, dt + \sigma^n_{T_i} dW^{n,T_i}_t, \\ dI_{tT_i}/I_{tT_i} &:= (\ldots)\, dt + \sigma^I_{T_i}\sqrt{V_t} dW^{I,T_i}_t, \\ dV_t &:= \alpha\,(\theta - V_t)\, dt + \epsilon\sqrt{V_t} dW_t, \end{aligned}$$

Here $\sigma^n_{T_i}$, $\sigma^I_{T_i}$, α, θ and ϵ are positive constants with $2\alpha\theta > \epsilon$ to ensure that $V_t \geq 0$. Furthermore, W^{n,T_i}_t, W^{I,T_i}_t and W_t are correlated Brownian motions in a given measure. Mercurio and Moreni then give semi-analytical formulas for the prices of year-on-year inflation caplets by using Fourier transform methods. The model adds smile to the implied volatilities of inflation-linked options but it can be difficult to calibrate the model to market prices for all maturities as the parameters of the volatility process V_t are not time dependent.

In a similar way, Mercurio and Moreni (2009) use the SABR model by Hagan et al. (2002) to add smile to the inflation market model. For a given tenor structure $\{T_i\}_{i\in\mathbb{N}}$, Mercurio and Moreni specify the model in the following way for some $i \in \mathbb{N}$:

$$\begin{aligned} dF^n_{tT_i}/F^n_{tT_i} &:= \sigma^n_{T_i} dW^{n,T_i}_t, \\ dI_{tT_i}/I_{tT_i} &:= \sum_{j=\beta_t}^{i} V^{T_j}_t dW^{I,T_i,T_j}_t, \\ dV^{T_i}_t/V^{T_i}_t &:= \nu^{T_i} dW^{V,T_i}_t, \end{aligned}$$

with $\sigma^n_{T_i}$ and ν^{T_i} positive constants, $\beta_t = \{j \in \mathbb{N} : T_{j-1} \leq t < T_j\}$, and where W^{n,T_i}_t, W^{I,T_i,T_j}_t and W^{V,T_i}_t are correlated Brownian motions in the T_i-forward measure. They then work out the dynamics of the year-on-year inflation ratio so that the corresponding caplet can be valued with a SABR log-normal formula as given by Hagan et al. (2002). Because each forward inflation index is driven by a separate volatility process, the model calibrates very well to the inflation smile for a given maturity. Nevertheless, to price exotic options, for example limited price indices (LPIs), we would have to simulate each of the volatility processes, as well as each forward inflation index. The difficulty in simulating the SABR processes, in addition to the high-dimensionality, makes it very computationally intensive to price exotic options.

Finally, we note that the inflation market models covered here are very flexible when it comes to obtaining a particular correlation structure of the inflation index. However, because of the many Brownian motions, it can still prove difficult to determine this structure from the market prices of inflation-linked derivatives. The high dimensionality also make it computationally expensive to calibrate and price inflation-linked derivatives with the models.

4.3 Year-on-year inflation models

Year-on-year inflation options have become more liquid in recent years, giving rise to inflation models that directly model either the year-on-year inflation ratio $I_T/\,I_{T-1}$ or the year-on-year inflation rate $I_T/\,I_{T-1}-1$. Jäckel and Bonneton (2010) outline a mean-reverting year-on-year model, which models the year-on-year inflation ratio. A multi-dimensional extension is the year-on-year inflation model introduced by Trovato et al. (2009) (see also Gretarsson, 2009). For a given measure $\mathbb{Q}^n$ in the nominal currency, the corresponding inverse numeraire is defined by a log-linear Gaussian process

$$D_t^n := P_{0t}^n \exp\left[-\phi_t^n \cdot x_t - \tfrac{1}{2}\phi_t^n \cdot G_t \phi_t^n\right], \tag{4.6}$$

where x_t is a multidimensional Gaussian process as in (4.3). Let $\{T_k\}_{k\in\mathbb{N}}$ be a given tenor structure with $0 = T_0 < T_1 < T_2 < \ldots$. The forward inflation ratio $Y_{tT_k} = I_{tT_k}\,/I_{tT_{k-1}}$ for $k \geq 1$ is defined as a multi-factor log-linear Gaussian process

$$Y_{tT_k} := \frac{I_{0T_k}}{I_{0T_{k-1}}} \exp\left[A_{tT_k} + \phi_{T_k}^y \cdot x_{t\wedge T_k} - \tfrac{1}{2}\phi_{T_k}^y \cdot G_{t\wedge T_k}\phi_{T_k}^y\right], \tag{4.7}$$

where $\phi_{T_k}^y$ is a multidimensional deterministic process and $t \wedge T_k := \min(t, T_k)$. The drift correction A_{tT_k} is a one-dimensional deterministic process, which is fitted so that the model matches the time zero value I_{0T_k} of the forward inflation index I_{tT_k}. To retrieve I_{tT_k}, we multiply the year-on-year ratios together:

$$I_{tT_k} = I_0 \prod_{i=1}^{k} Y_{tT_i} = I_{0T_k} \exp\left[\sum_{i=1}^{k}\left(A_{tT_i} + \phi_{T_i}^y \cdot x_{t\wedge T_i} - \tfrac{1}{2}\phi_{T_i}^y \cdot G_{t\wedge T_i}\phi_{T_i}^y\right)\right]. \tag{4.8}$$

To price zero-coupon and year-on-year inflation options, it is not necessary to introduce a real currency as done for the foreign-exchange analogy inflation models in §4.1. The pricing formulas are given in closed form and it is simple to calibrate the model to market prices by adjusting ϕ^y and the variance G_t of x_t. However, this model does not give smile to the implied volatilities of inflation-linked options.

Kenyon (2008) models the year-on-year inflation rate $I_T/\,I_{T-1} - 1$ as a Gaussian process. He lists different methods to get smile, such as using mixture modelling or replacing the normal distribution with a normal-gamma distribution. For the finite Gaussian mixture model the density $p_t[z]$ of the year-on-year inflation rate is given by

$$p_t[z] = \sum_{i=1}^{N} \lambda_i p_t^i[z], \quad \text{with} \quad \sum_{i=1}^{N} \lambda_i = 1.$$

Here $p_t^i[z]$ are normal densities with a given mean and variance. This model gives simple pricing formulas for year-on-year inflation options and a good fit to the year-on-year inflation smiles. We note that it is difficult to use these modelling methods to price zero-coupon options, because the inflation index can take negative values in the model but not in reality.

4.4 Quadratic Gaussian models

In §4.1, we looked at models that used the foreign-exchange analogy for modelling inflation. There the short rate was a linear function of a Gaussian process, so the models had no curvature or skew in the inflation volatility smile. Here, we therefore review the quadratic Gaussian framework, where the short rate is a quadratic function of a Gaussian process, giving us control over the volatility smile. The quadratic Gaussian model was introduced by Beaglehole and Tenney (1991) and has been widely studied by academics for modelling the short rate. Ahn et al. (2002) and Assefa (2007) give a good overview of the literature and provide detailed pricing results. Furthermore, papers by Piterbarg (2009) and McCloud (2010) illustrate good properties and flexibility on the part of the model in terms of calculations, parameterisation, pricing and calibration.

As in (4.3), we let x_t be a Gaussian process in the bank-account-risk-neutral measure $\mathbb{Q}^*$:

$$x_t := \int_0^t \operatorname{diag}[\sigma_s]\, dW_s.$$

Here σ_s is a multidimensional deterministic process and W_t is a multidimensional $\mathbb{Q}^*$-Brownian motion with correlation matrix ρ_t. Therefore, x_t has a zero mean under $\mathbb{Q}^*$ and a covariance matrix G_t with entries $G_t^{ij} = \int_0^t \rho_s^{ij} \sigma_s^i \sigma_s^j ds$. Following Piterbarg (2009), the short rate in a multi-factor quadratic Gaussian model is given by

$$r_t = \alpha_t + \phi_t \cdot x_t + \frac{1}{2} x_t \cdot \Psi_t x_t,$$

where the scalar α_t, the vector ϕ_t, and the matrix Ψ_t, are all deterministic functions. The discount bond is then given by

$$P_{tT} = \frac{P_{0T}}{P_{0t}} \exp\left[-\phi_{tT} \cdot x_t - \frac{1}{2} x_t \cdot \Psi_{tT} x_t - \alpha_{tT}\right].$$

Here the matrix Ψ_{tT} and the vector ϕ_{tT} are deterministic functions, which satisfy the Riccati equations

$$\begin{aligned}\frac{d}{dt}\Psi_{tT} &= -\Psi_t + \frac{1}{2}\Psi_{tT} \cdot G_t \Psi_{tT} \\ \frac{d}{dt}\phi_{tT} &= -\phi_t + \Psi_{tT} \cdot G_t \phi_{tT},\end{aligned}$$

with terminal conditions $\Psi_{tT} = 0$ and $\phi_{tT} = 0$. We require that the model fits the initial condition of the discount bond, $\mathbb{E}_0^t \left[P_{tT}\right] = P_{0T}/\, P_{0t}$, which determines the scalar function α_{tT}. Here $\mathbb{E}_0^t$ represents the expectations under the t-forward measure conditional on time zero. The discount bond can then be written as

$$P_{tT} = \frac{P_{0T}}{P_{0t}} \frac{\exp\left[-\phi_{tT} \cdot x_t - \frac{1}{2}x_t \cdot \Psi_{tT} x_t\right]}{\mathbb{E}_0^t \left[\exp\left[-\phi_{tT} \cdot x_t - \frac{1}{2}x_t \cdot \Psi_{tT} x_t\right]\right]}$$

McCloud (2008, 2010) introduces a multi-currency quadratic Gaussian modelling framework to model interest rates. Instead of specifying a process for the short rate, he uses the potential approach of Rogers (1997). Given some universal measure $\mathbb{Q}$ (with 0-conditional expectations noted by $\mathbb{E}_0$), for each currency c, we define a measure process m_t^c, and an inverse numeraire D_t^c corresponding to a measure $\mathbb{Q}^c$ defined by the Radon-Nikodym derivative $d\mathbb{Q}_t^c/\, d\mathbb{Q}_t = m_t^c$. The quantities m_t^c and D_t^c are log-quadratic functions[2] of a Gaussian process x_t in $\mathbb{Q}$,

$$m_t^c\left[\phi_t^{mc}, \Psi_t^{mc}\right] := s_{0t}^c \frac{\exp\left[-\phi_t^{mc} \cdot x_t - \frac{1}{2}x_t \cdot \Psi_t^{mc} x_t\right]}{\mathbb{E}_0\left[\exp\left[-\phi_t^{mc} \cdot x_t - \frac{1}{2}x_t \cdot \Psi_t^{mc} x_t\right]\right]}, \tag{4.9a}$$

$$D_t^c\left[\phi_t^c, \Psi_t^c\right] := P_{0t}^c \frac{\exp\left[-\phi_t^c \cdot x_t - \frac{1}{2}x_t \cdot \Psi_t^c x_t\right]}{\mathbb{E}_0\left[\exp\left[-\phi_t^c \cdot x_t - \frac{1}{2}x_t \cdot \Psi_t^c x_t\right]\right]}. \tag{4.9b}$$

Here ϕ_t^{mc} and ϕ_t^c are deterministic vector-valued processes and Ψ_t^{mc} and Ψ_t^c are deterministic symmetric matrix-valued processes. McCloud (2008) then gives expressions for the discount factor P_{tT}^c in currency c and for the foreign-exchange index process s^{cd} from some currency d to currency c, and shows a method to calculate a linear multi-currency option under the forward measure of each currency.

We note that one could apply this model for inflation through the application of the foreign-exchange analogy. This would extend the linear Gaussian inflation model in §4.1 to a quadratic Gaussian inflation model. Consequently, we would have similar problems estimating the parameters of the illiquid real rate.

[2] In his paper, McCloud (2008) calls the functions exponential-quadratic, indicating that a quadratic function is in the exponent. Instead, to underline the similarities to log-linear functions, we call the functions log-quadratic, indicating that the logarithms of the functions are quadratic functions.

4.5 Summary

In this chapter, we have reviewed a range of inflation models. Table 4.1 shows an overview of the three types of inflation model. Each type contains models that incorporate curvature and skew in the inflation volatility smiles.

The models based on the foreign-exchange analogy specify the real discount bond. One could do the same by adapting the multi-currency quadratic Gaussian model from McCloud (2008), to model the inflation volatility smiles. Nevertheless, for these models it is complicated to estimate the parameters of the real rate.

The inflation market model and its stochastic-volatility extensions specify the inflation index directly. However, these market models are computationally intensive because of their high dimensionality. Furthermore, the stochastic-volatility processes only give approximate pricing formulas and are difficult to simulate.

Alternatively, year-on-year models specify the year-on-year inflation rate, $I_T/\,I_{T-1}-1$, or ratio, $I_T/\,I_{T-1}$. The models by Kenyon (2008) can control the year-on-year inflation volatility smile but are not well-suited to pricing zero-coupon inflation swaps and options. The log-normal year-on-year inflation ratio model from Trovato et al. (2009) has simple pricing formulas and intuitive parameters to describe the year-on-year inflation volatilities and can also be used to model the inflation index. However, the model has no curvature or skew in the inflation volatility smiles.

To add curvature and skew to the smile, we define a new model—the Quadratic Gaussian Year-on-Year inflation model—abbreviated to the QGY model. We specify processes for the nominal inverse numeraire and the year-on-year inflation ratio. Each process is a Markovian log-quadratic function of a multidimensional Gaussian process. This makes it possible to get control over the volatility smile, both for nominal interest rate options and for inflation-linked options. Furthermore, by defining the year-on-year inflation ratio, rather than the year-on-year inflation rate, on a discrete tenor structure, we get a well-defined process for the inflation index.

The QGY model lends itself to low dimensionality, intuitive parameterisation, analytical formulas, fast calibration and ease of simulation. We will look at the QGY model in detail in Part II.

	Without smile	**With smile**
FX analogy	Jarrow and Yıldırım (2003) (multi-factor log-linear Gaussian extension by Trovato et al., 2009)	Cross-currency log-quadratic Gaussian by McCloud (2008)
Inflation index	Belgrade et al. (2004)	Mercurio and Moreni (2006, 2009)
Year-on-year inflation	Kenyon (2008); Jäckel and Bonneton (2010) (multi-factor log-linear Gaussian extension by Trovato et al., 2009)	Kenyon (2008);

Table 4.1: Overview of inflation models. With the QGY model we fill the gap in the table by providing a year-on-year inflation ratio model with smile.

Part II

The QGY Model

5 Specifications

In this chapter, we define the Quadratic Gaussian Year-on-Year inflation model—the QGY model—and list the necessary pricing methods for the model. As is seen in Table 5.1, the QGY model is a year-on-year inflation model with a smile.

In the model, we specify a strictly positive process for the year-on-year inflation ratio in order to get a well-defined process for the inflation index. All the year-on-year inflation ratios are Markovian functions of the same (low-dimensional) Gaussian process, which distinguishes the model from market models. The year-on-year ratios are log-quadratic functions of this Gaussian process, which makes it easy to simulate the modelling processes. The quadratic term enables the model to incorporate curvature and skew in the inflation volatility smiles while still having good analytical properties.

In §5.1, we define the modelling processes in the QGY model. In §5.2, we provide a formula for the expectations of a function of a log-quadratic Gaussian process. To derive this formula, we follow McCloud (2008), as detailed in Appendix C. In §5.3, we show the log-linearity of log-quadratic Gaussian processes.

5.1 Modelling processes

We let $\mathbb{Q}^n$ be some measure in the nominal currency. As in Chapter 3, we specify a probability space $(\Omega, \mathcal{F}, \mathbb{Q}^n)$. The filtration $\mathcal{F}_t$ is generated by a multidimensional $\mathbb{Q}^n$-Brownian motion W_t that has a correlation matrix ρ_t. The t-conditional expectations are denoted by $\mathbb{E}_t^n[\cdot] := \mathbb{E}^n[\cdot|\mathcal{F}_t]$, where $\mathbb{E}^n$ are the expectations under $\mathbb{Q}^n$. We define an $\mathcal{F}_t$-adapted multidimensional Gaussian process x_t, $t \geq 0$, with

$$x_t := \int_0^t \operatorname{diag}[\sigma_s]\, dW_s. \tag{5.1}$$

	Without smile	**With smile**
FX analogy	Jarrow and Yıldırım (2003) (multi-factor log-linear Gaussian extension by Trovato et al., 2009)	cross-currency log-quadratic Gaussian by McCloud (2008)
Inflation index	Belgrade et al. (2004)	Mercurio and Moreni (2006, 2009)
Year-on-year inflation	Kenyon (2008); Jäckel and Bonneton (2010) (multi-factor log-linear Gaussian extension by Trovato et al., 2009)	Kenyon (2008); **QGY model**

Table 5.1: Overview of inflation models. The QGY model is a year-on-year inflation ratio model with inflation volatility smile.

Here σ_s is a strictly positive multidimensional deterministic process with the same dimension as W_t. We deduce that x_t has a mean zero under $\mathbb{Q}^n$ and a covariance matrix G_t with entries $G_t^{ij} = \int_0^t \rho_s^{ij} \sigma_s^i \sigma_s^j ds$. In the QGY model, we model the nominal inverse numeraire in the following way:

Definition 5.1.1. *In the QGY model, the inverse numeraire of the measure $\mathbb{Q}^n$ is denoted by D_t^n and defined by a normalised log-quadratic process:*

$$D_t^n \left[\phi_t^n, \Psi_t^n\right] := P_{0t}^n \frac{\exp\left[-\phi_t^n \cdot x_t - \frac{1}{2} x_t \cdot \Psi_t^n x_t\right]}{\mathbb{E}_0^n \left[\exp\left[-\phi_t^n \cdot x_t - \frac{1}{2} x_t \cdot \Psi_t^n x_t\right]\right]}. \tag{5.2}$$

Here ϕ_t^n is a deterministic vector-valued process and Ψ_t^n is a deterministic symmetric matrix-valued process.

The expectations in the denominator of (5.2) normalise the inverse numeraire so that we match the initial condition of the nominal discount bond, P_{0t}^n—that is, $\mathbb{E}_0^n \left[D_t^n \left[\phi_t^n, \Psi_t^n\right]\right] = P_{0t}^n$. In the QGY model, we model inflation in the following way.

Definition 5.1.2. *Let $\{T_k\}_{k \in \mathbb{N}}$ be a strictly increasing tenor structure, $0 = T_0 < T_1 < \dots$. In the QGY model, the inflation ratio $Y_{T_k} = I_{T_k} / I_{T_{k-1}}$ is a log-quadratic process for all integers $k \geq 1$:*

$$Y_{T_k} \left[\phi_{T_k}^y, \Psi_{T_k}^y\right] := \frac{I_{0T_k}}{I_{0T_{k-1}}} \exp\left[A_{T_k} - \phi_{T_k}^y \cdot x_{T_k} - \tfrac{1}{2} x_{T_k} \cdot \Psi_{T_k}^y x_{T_k}\right]. \tag{5.3}$$

The vector-valued process $\phi_{T_k}^y$, the symmetric-matrix-valued process $\Psi_{T_k}^y$ and the one-dimensional drift process A_{T_k} are deterministic.

We note that both the process for the inverse numeraire D_t^n and the process for the year-on-year inflation ratio Y_{T_k} are functions of the same Gaussian process x. By adjusting the dimension of x, we can get a rich correlation structure within and between the two processes. We also note that the year-on-year rate is Markovian in x because Y_{T_k} only depends on x_{T_k}.[1] The matrix $\Psi^y_{T_k}$ makes it possible to get non-zero skewness and excess kurtosis, which give us a convenient way to control the skew and the curvature of the volatility smile of year-on-year options. This can be seen in the following proposition. The derivation is given in Appendix B (we note that $x_{T_k} | x_t$ has a Gaussian distribution with mean x_t and variance $G_{tT_k} := G_{T_k} - G_t$).

Proposition 5.1.1. *Conditional on $\mathcal{F}_t$, the variance κ_2, the skewness $\kappa_3 \big/ \kappa_2^{3/2}$ and the excess kurtosis $\kappa_4 \big/ \kappa_2^2$ of the quadratic Gaussian process $\ln Y_{T_k}$, $T_k \geq t$, are deduced from*

$$\begin{aligned}
\kappa_2 &= \frac{1}{2}\operatorname{tr}\left[\left(\Psi^y_{T_k} G_{tT_k}\right)^2\right] + \left(\phi^y_{T_k} + \Psi^y_{T_k} x_t\right) \cdot G_{tT_k}\left(\phi^y_{T_k} + \Psi^y_{T_k} x_t\right), \\
\kappa_3 &= -\operatorname{tr}\left[\left(\Psi^y_{T_k} G_{tT_k}\right)^3\right] - 3\left(\phi^y_{T_k} + \Psi^y_{T_k} x_t\right) \cdot G_{tT_k} \Psi^y_{T_k} G_{tT_k}\left(\phi^y_{T_k} + \Psi^y_{T_k} x_t\right), \\
\kappa_4 &= 3\operatorname{tr}\left[\left(\Psi^y_{T_k} G_{tT_k}\right)^4\right] + 12\left(\phi^y_{T_k} + \Psi^y_{T_k} x_t\right) \cdot G_{tT_k}\left(\Psi^y_{T_k} G_{tT_k}\right)^2\left(\phi^y_{T_k} + \Psi^y_{T_k} x_t\right).
\end{aligned}$$

We have defined the year-on-year ratios $Y_{T_k} = I_{T_k} / I_{T_{k-1}}$, $k = 1, 2, \ldots$, on a discrete tenor structure. The ratios are strictly positive, so we get the process for the inflation index by multiplying together the processes for the year-on-year inflation ratios:

$$I_{T_k} = I_0 \prod_{i=1}^{k} Y_{T_i}\left[\phi^y_{T_i}, \Psi^y_{T_i}\right] = I_{0T_k} \exp\left[\sum_{i=1}^{k} A_{T_i}\right] \prod_{i=1}^{k} \exp\left[-\phi^y_{T_i} \cdot x_{T_i} - \tfrac{1}{2} x_{T_i} \cdot \Psi^y_{T_i} x_{T_i}\right]. \tag{5.4}$$

I_{T_k} is not Markovian in x, as it depends on x_t for $t = T_1, T_2, \ldots, T_k$.[2]

5.2 Expectations of log-quadratic Gaussian processes

To calculate prices of inflation-linked derivatives in the QGY model, we need to calculate the expectations of the inverse numeraire and the year-on-year inflation ratio, which are both quadratic Gaussian processes. We use calculations by McCloud (2008) to calculate these expectations.

[1] Here we note that we could model the year-on-year rate $I_{T_k} / I_{T_{k-1}} - 1$ and make it a quadratic function of x_{T_k}. This is not practical because the year-on-year rate would then not be strictly positive. This would make it difficult to define a process for the inflation index I_{T_k}.

[2] We will see that it is still possible to get closed-form pricing formulas for zero-coupon inflation swaps. Prices of zero-coupon inflation options can be calculated with Monte Carlo methods that are simple as they only require the simulation of the Gaussian process x_{T_i} for $i = 1, 2, \ldots, k$.

Let f be some European payoff function that depends on the Gaussian process x in (5.1). At maturity T it pays out $f[x_T]$. The process x_T conditional on $\mathcal{F}_t$ for $t < T$ has a multi-factor Gaussian distribution with mean x_t and variance $G_{tT} := G_T - G_t$. Detailed derivation of the following proposition is given in Appendix C.[3]

Proposition 5.2.1. *(McCloud, 2008) The expectations of a log-quadratic Gaussian process multiplied by the function $f[x_T]$ are given by*

$$\begin{aligned} &\mathbb{E}_t^n\left[\exp\left[-\phi \cdot x_T - \tfrac{1}{2} x_T \cdot \Psi x_T\right] f[x_T]\right] \\ &= \mathbb{E}_t^n\left[\exp\left[-\phi \cdot x_T - \tfrac{1}{2} x_T \cdot \Psi x_T\right]\right] \mathbb{E}_t^n\left[f\left[\widetilde{x}_T\right]\right], \end{aligned} \tag{5.5}$$

with
$$\begin{aligned} &\mathbb{E}_t^n\left[\exp\left[-\phi \cdot x_T - \tfrac{1}{2} x_T \cdot \Psi x_T\right]\right] \\ &= \det\left[1 + \Psi G_{tT}\right]^{-\frac{1}{2}} \exp\left[\tfrac{1}{2}\phi \cdot G_{tT}\left(1 + \Psi G_{tT}\right)^{-1}\phi\right] \\ &\quad \times \exp\left[-\left(1 + \Psi G_{tT}\right)^{-1}\phi \cdot x_t - \tfrac{1}{2} x_t \cdot \left(1 + \Psi G_{tT}\right)^{-1}\Psi x_t\right]. \end{aligned} \tag{5.6}$$

Here ϕ is a deterministic vector and Ψ is a deterministic symmetric matrix. $\widetilde{x}_T$ is a Gaussian process that, conditional on time t, has mean $\widetilde{x}_t$ and variance $\widetilde{G}_{tT}$:

$$\begin{aligned} \widetilde{x}_T &:= \widetilde{x}_t + \sqrt{\widetilde{G}_{tT}}\sqrt{G_{tT}}^{-1}\left(x_T - x_t\right), \\ \widetilde{x}_t &:= \widetilde{G}_{tT}\left(G_{tT}^{-1} x_t - \phi\right), \\ \sqrt{\widetilde{G}_{tT}} &:= \sqrt{G_{tT}}\left(1 + \sqrt{G_{tT}}^{\intercal}\Psi\sqrt{G_{tT}}\right)^{-\frac{1}{2}}, \\ \widetilde{G}_{tT} &= \sqrt{\widetilde{G}_{tT}}\sqrt{\widetilde{G}_{tT}}^{\intercal} = G_{tT}\left(1 + \Psi G_{tT}\right)^{-1}. \end{aligned}$$

Corollary 5.2.1. *The expectations in* (5.6) *can be written as*

$$\mathbb{E}_t^n\left[X_T[\phi, \Psi]\right] = E_{tT}[\phi, \Psi]\, X_{tT}[\phi, \Psi], \tag{5.7}$$

where $X_{tT}[\phi, \Psi] := X_t\left[\Theta_{tT}[\phi, \Psi]\right] = \exp\left[-M_{tT}[\Psi]\phi \cdot x_t - \tfrac{1}{2} x_t \cdot M_{tT}[\Psi]\Psi x_t\right],$

$$\begin{aligned} X_t[\phi, \Psi] &:= \exp\left[-\phi \cdot x_t - \tfrac{1}{2} x_t \cdot \Psi x_t\right], \\ \Theta_{tT}[\phi, \Psi] &:= \left(M_{tT}[\Psi]\phi, M_{tT}[\Psi]\Psi\right), \\ M_{tT}[\Psi] &:= \left(1 + \Psi G_{tT}\right)^{-1}, \\ E_{tT}[\phi, \Psi] &:= \det\left[M_{tT}[\Psi]\right]^{-\frac{1}{2}} \exp\left[\tfrac{1}{2}\phi \cdot G_{tT} M_{tT}[\Psi]\phi\right]. \end{aligned}$$

[3] The derivation is based on the fact that both the modelling processes and the density of the Gaussian distribution are log-quadratic functions. This makes it possible to complete the squares and get at least semi-analytical formulas for the expectations. To be able to use this formula, we have therefore not included higher-order polynomial terms, such as cubic terms, in the modelling processes.

The $\mathcal{F}_t$-conditional expectations of a log-quadratic Gaussian process $X_T[\phi, \Psi]$ are therefore given by multiplying the deterministic correction term $E_{tT}[\phi, \Psi]$ and the log-quadratic Gaussian process $X_{tT}[\phi, \Psi]$.[4] Here we assume that the matrix $1 + \Psi G_{tT}$ is positive-definite. This condition is equivalent to the matrix's having only strictly positive eigenvalues. It ensures that the matrix has a strictly positive determinant and is invertible. Therefore, the transformation matrix $M_{tT}[\Psi]$ and the transformation correction $E_{tT}[\phi, \Psi]$ are well defined.

The modelling quantities defined in §5.1 can now be written as follows:

$$D_t^n[\phi_t^n, \Psi_t^n] = P_{0t}^n \frac{X_t[\phi_t^n, \Psi_t^n]}{E_{0t}^n[\phi_t^n, \Psi_t^n]}, \tag{5.8a}$$

$$Y_{T_k}\left[\phi_{T_k}^y, \Psi_{T_k}^y\right] = \frac{I_{0T_k}}{I_{0T_{k-1}}} \exp[A_{T_k}] X_{T_k}\left[\phi_{T_k}^y, \Psi_{T_k}^y\right], \tag{5.8b}$$

$$I_{T_k} = I_{0T_k} \exp\left[\sum_{i=1}^k A_{T_i}\right] \prod_{i=1}^k X_{T_i}\left[\phi_{T_i}^y, \Psi_{T_i}^y\right]. \tag{5.8c}$$

5.3 Log-linearity of log-quadratic Gaussian processes

A useful property of the log-quadratic Gaussian process is the log-linearity of its parameters.

Lemma 5.3.1. *For some vectors ϕ and ϕ' and symmetric matrices Ψ and Ψ', we have*

$$\begin{aligned} &\exp\left[-\phi \cdot x_t - \tfrac{1}{2} x_t \cdot \Psi x_t\right] \exp\left[-\phi' \cdot x_t - \tfrac{1}{2} x_t \cdot \Psi' x_t\right] \\ &= \exp\left[-\left(\phi + \phi'\right) \cdot x_t - \tfrac{1}{2} x_t \cdot \left(\Psi + \Psi'\right) x_t\right]. \end{aligned}$$

We therefore get

$$\begin{aligned} X_t[\phi, \Psi] X_t\left[\phi', \Psi'\right] &= X_t\left[\phi + \phi', \Psi + \Psi'\right] =: X_t\left[(\phi, \Psi) + \left(\phi', \Psi'\right)\right], \\ X_t[\phi, \Psi] X_{tT}\left[\phi', \Psi'\right] &= X_t\left[(\phi, \Psi) + \Theta_{tT}\left[\phi', \Psi'\right]\right]. \end{aligned} \tag{5.9}$$

[4] For $\Psi = 0$ the log-quadratic process $X_{tT}[\phi, \Psi]$ reduces to a log-linear Gaussian process $\exp[-\phi \cdot x_T]$. Conditional on time t, $-\phi \cdot x_T$ has a normal distribution with mean $-\phi \cdot x_t$ and variance $\phi \cdot G_{tT}\phi$. In this case, the conditional expectations then take the following form:

$$\mathbb{E}_t^n[\exp[-\phi \cdot x_T]] = \exp\left[\tfrac{1}{2}\phi \cdot G_{tT}\phi\right] \exp[-\phi \cdot x_t].$$

We have $E_{tT}[\phi, 0] = \exp\left[\frac{1}{2}\phi \cdot G_{tT}\phi\right]$ and $X_{tT}[\phi, 0] = X_t[\phi, 0] = \exp[-\phi \cdot x_t]$ and the parameter ϕ is not transformed.

6 Pricing formulas

In this chapter, we will calculate pricing formulas in the QGY model for inflation-linked derivatives. In §6.1, we show a formula for the nominal discount bond process. In §6.2, we derive a closed-form expression for the drift of the year-on-year process in the QGY model, so that the QGY model matches a given forward inflation index. Using similar calculations, in §6.3 we get the price of a general period-on-period inflation swaplet in closed form. This makes it possible to price both zero-coupon swaps with payment delay and year-on-year swaplets.

In §6.4 we derive the pricing formula for a year-on-year inflation caplet. The formula is of the same form as the Black-Scholes formula, except that we need to integrate a standard Gaussian density over a domain bounded by a quadratic equation. In two dimensions, using polar coordinates, we transform this integral to a finite sum of one-dimensional integrals. This makes pricing both faster and more accurate. This method is detailed in Appendix E.

6.1 Nominal bond

To calculate the process for the nominal bond, we use calculations given by McCloud (2008). We detail his derivation in Appendix D.

Proposition 6.1.1. *(McCloud, 2008) In the QGY model, the expectations of the nominal inverse numeraire D_t^n in* (5.2) *are given by*

$$\begin{aligned}\mathbb{E}_t\left[D_T^n\left[\phi_T^n,\Psi_T^n\right]\right] &= P_{0T}^n \mathbb{E}_t\left[\frac{X_T\left[\phi_T^n,\Psi_T^n\right]}{E_{0T}^n\left[\phi_T^n,\Psi_T^n\right]}\right] \\ &= P_{0T}^n \frac{X_t\left[\Theta_{tT}\left[\phi_T^n,\Psi_T^n\right]\right]}{E_{0t}^n\left[\Theta_{tT}\left[\phi_T^n,\Psi_T^n\right]\right]} = \frac{P_{0T}^n}{P_{0t}^n} D_t^n\left[\Theta_{tT}\left[\phi_T^n,\Psi_T^n\right]\right]. \qquad (6.1)\end{aligned}$$

The expression for the nominal discount bond price process is now given by

$$P_{tT}^n = \frac{\mathbb{E}_t\left[D_T^n\left[\phi_T^n,\Psi_T^n\right]\right]}{D_t^n\left[\phi_T^n,\Psi_T^n\right]} = \frac{P_{0T}^n}{P_{0t}^n}\frac{D_t^n\left[\Theta_{tT}\left[\phi_T^n,\Psi_T^n\right]\right]}{D_t^n\left[\phi_T^n,\Psi_T^n\right]}. \qquad (6.2)$$

We see from (6.1) that the conditional expectations of the inverse numeraire are obtained by transforming the parameters using the transformation Θ_{tT} defined in §5.2. As we are mainly interested in pricing inflation-linked options, we refer to McCloud (2008) for pricing formulas for options on nominal interest rates.

6.2 Inflation index drift

The initial value I_{0T_k} of the forward inflation index I_{tT_k} is given by market prices of zero-coupon inflation swaps. For the QGY model to match I_{0T_k} for $k = 1, 2, \ldots$ we require that the following expression holds:

$$\begin{aligned}I_{0T_k} &= \frac{1}{P_{0T_k}^n}\mathbb{E}_0^n\left[D_{T_k}^n I_{T_k}\right] \\ &= \frac{I_{0T_k}}{E_{0T_k}^n\left[\phi_{T_k}^n,\Psi_{T_k}^n\right]}\exp\left[\sum_{i=1}^k A_{T_i}\right]\mathbb{E}_0\left[X_{T_k}\left[\phi_{T_k}^n,\Psi_{T_k}^n\right]\prod_{i=1}^k X_{T_i}\left[\phi_{T_i}^y,\Psi_{T_i}^y\right]\right].\end{aligned}$$

As I_{T_k} depends on x_{T_i} for $i = 1, 2, \ldots, k$, we iteratively apply the tower property of conditional expectations,[1] along with the expressions for the expectations of a log-quadratic

[1]The tower property states that for some $s \le t \le T$ and a function f, we have

$$\mathbb{E}_s^n\left[f\left[x_T\right]\right] = \mathbb{E}_s^n\left[\mathbb{E}_t^n\left[f\left[x_T\right]\right]\right].$$

process detailed in §5.2 and its log-linear property in §5.3. We get

$$
\begin{aligned}
I_{0T_k} &= \frac{1}{P_{0T_k}} \mathbb{E}_0^n \left[D_{T_k}^n I_{T_k} \right] \\
&\overset{(5.8a,5.8c)}{=} I_{0T_k} \exp\left[\sum_{i=1}^{k} A_{T_i} \right] \left(E_{0T_k} \left[\phi_{T_k}^n, \Psi_{T_k}^n \right] \right)^{-1} \\
&\quad \times \mathbb{E}_0^n \left[\left(\prod_{i=1}^{k-1} X_{T_i} \left[\phi_{T_i}^y, \Psi_{T_i}^y \right] \right) \mathbb{E}_{T_{k-1}}^n \left[X_{T_k} \left[\phi_{T_k}^y + \phi_{T_k}^n, \Psi_{T_k}^y + \Psi_{T_k}^n \right] \right] \right] \\
&\overset{(5.7)}{=} I_{0T_k} \exp\left[\sum_{i=1}^{k} A_{T_i} \right] \left(E_{0T_k} \left[\phi_{T_k}^n, \Psi_{T_k}^n \right] \right)^{-1} E_{T_{k-1}T_k} \left[\phi_{T_k}^y + \phi_{T_k}^n, \Psi_{T_k}^y + \Psi_{T_k}^n \right] \\
&\quad \times \mathbb{E}_0^n \left[\left(\prod_{i=1}^{k-1} X_{T_i} \left[\phi_{T_i}^y, \Psi_{T_i}^y \right] \right) X_{T_{k-1}} \left[\Theta_{T_{k-1}T_k} \left[\phi_{T_k}^y + \phi_{T_k}^n, \Psi_{T_k}^y + \Psi_{T_k}^n \right] \right] \right] \\
&\overset{(5.9)}{=} I_{0T_k} \exp\left[\sum_{i=1}^{k} A_{T_i} \right] \left(E_{0T_k} \left[\phi_{T_k}^n, \Psi_{T_k}^n \right] \right)^{-1} E_{T_{k-1}T_k} \left[\phi_{T_k}^y + \phi_{T_k}^n, \Psi_{T_k}^y + \Psi_{T_k}^n \right] \\
&\quad \times \mathbb{E}_0^n \left[\left(\prod_{i=1}^{k-2} X_{T_i} \left[\phi_{T_i}^y, \Psi_{T_i}^y \right] \right) \mathbb{E}_{T_{k-2}}^n \left[X_{T_{k-1}} \left[\begin{bmatrix} \phi_{T_{k-1}}^y \\ \Psi_{T_{k-1}}^y \end{bmatrix} + \Theta_{T_{k-1}T_k} \begin{bmatrix} \phi_{T_k}^y + \phi_{T_k}^n \\ \Psi_{T_k}^y + \Psi_{T_k}^n \end{bmatrix} \right] \right] \right].
\end{aligned}
$$

By continuing in the same manner, iteratively taking expectations conditional on times $T_{k-2}, T_{k-3}, \ldots, T_2, T_1, 0$, cancelling I_{0T_k} and rearranging terms, we get a closed-form formula for the drift term of the inflation index I_{T_k}.

Proposition 6.2.1. *In the QGY model, the drift term of the inflation index I_{T_k} is given by*

$$
\exp\left[\sum_{i=1}^{k} A_{T_i} \right] = \frac{E_{0T_k} \left[\phi_{T_k}^n, \Psi_{T_k}^n \right]}{\prod_{i=1}^{k} E_{T_{i-1}T_i} \left[\Xi_{T_i}^{T_k} \right]},
$$

where the parameter transformation $\Xi_{T_i}^{T_k}$, $k = 1, 2, \ldots, k$, is defined in the following way:

$$
\begin{aligned}
\Xi_{T_k}^{T_k} &:= \left(\phi_{T_k}^y + \phi_{T_k}^n, \Psi_{T_k}^y + \Psi_{T_k}^n \right), \\
\Xi_{T_i}^{T_k} &:= \Theta_{T_i T_{i+1}} \left[\Xi_{T_{i+1}}^k \right] + \left(\phi_{T_i}^y, \Psi_{T_i}^y \right), \quad \textit{for } i = k-1, k-2, \ldots, 1.
\end{aligned}
$$

In the calculations of the drift we apply this property k times:

$$
\mathbb{E}_0^n \left[\prod_{i=1}^{k} X_{T_i} \right] = \mathbb{E}_0^n \left[X_{T_1} \mathbb{E}_{T_1}^n \left[X_{T_2} \mathbb{E}_{T_2}^n \left[\cdots X_{T_{k-1}} \mathbb{E}_{T_{k-1}}^n \left[X_{T_k} \right] \cdots \right] \right] \right].
$$

The drift term of the inflation ratio $I_{T_k}/\, I_{T_h}$, for some $h < k$, can be retrieved from the relation

$$\exp\left[\sum_{i=h+1}^{k} A_{T_i}\right] = \exp\left[\sum_{i=1}^{k} A_{T_i}\right] \Big/ \exp\left[\sum_{i=1}^{k-1} A_{T_i}\right].$$

In particular, the drift term A_{T_k} of the year-on-year inflation ratio $Y_{T_k} = I_{T_k}/\, I_{T_{k-1}}$ is given by

$$\exp\left[A_{T_k}\right] = \exp\left[\sum_{i=1}^{k} A_{T_i}\right] \Big/ \exp\left[\sum_{i=1}^{k-1} A_{T_i}\right].$$

6.3 Inflation swaplet

In Chapter 2, we saw two types of inflation-linked swap. For the zero-coupon inflation swap agreed upon at time t, the value of the inflation index, I_u, is swapped for a predetermined cash amount at time $T \geq u$. For $T = u$ this amount is the forward inflation index I_{tu}. For $T > u$ we get a payment delay convexity correction as discussed in §2.1. For the year-on-year inflation swaplet agreed upon at time t, the year-on-year inflation rate, $I_u/\, I_{u-1} - 1$, is swapped for the year-on-year inflation forward at time u. This forward incorporates a year-on-year convexity correction as described in §2.2.

Both types of swap can be described by the general period-on-period inflation swaplet, where, at time T, $I_u/\, I_s$ is swapped for the value of this payment at time t. Here we assume that $t \leq s < u \leq T$. For the zero-coupon inflation swap, we have $s = t$, while for the year-on-year inflation swap, we have $s = u - 1$.[2] To calculate the value of the general inflation swaplet in the QGY model, we will assume that the reference dates of the inflation ratio $I_u/\, I_s$ fall on the model tenor structure $\{T_i\}_{i\in\mathbb{N}}$—that is $u = T_k$ and $s = T_h$ for some $k > h \geq 0$. We use a method similar to the one we used in §6.2, where we calculated the value of the drift A_{T_k}. We use the modelling processes in (5.8a) and (5.8c) and apply the log-linear property of (5.9), along with the expectation formula in (5.7) and the tower property of conditional expectations.

[2] If I_u is swapped for an amount X that is determined at time $t < u$, then $I_u/\, I_t$ is swapped for $X\,/I_t$. If $I_u/\, I_{u-1}$ is swapped for an amount Y, then $I_u/\, I_{u-1} - 1$ is swapped for $Y - 1$.

Proposition 6.3.1. *Let $0 \leq t < T_{h+1} < T_k \leq T$. The value of a swaplet on the general inflation ratio I_{T_k} / I_{T_h} with a payment delay $T - T_k$ is given at time t by*

$$
\begin{aligned}
&\frac{1}{D_t^n}\mathbb{E}_t^n\left[D_T^n \frac{I_{T_k}}{I_{T_h}}\right] \\
&= \frac{P_{0T}^n}{D_t^n}\frac{I_{0T_k}}{I_{0T_h}} \exp\left[\sum_{i=h+1}^{k} A_{T_i}\right] \mathbb{E}_t^n\left[\frac{X_T\left[\phi_T^n, \Psi_T^n\right]}{E_{0T}\left[\phi_T^n, \Psi_T^n\right]} \prod_{i=h+1}^{k} X_{T_i}\left[\phi_{T_i}^y, \Psi_{T_i}^y\right]\right] \\
&= \frac{P_{0T}^n}{D_t^n}\frac{I_{0T_k}}{I_{0T_h}} \frac{1}{E_{0T}\left[\Theta_{T_kT}\left[\phi_T^n, \Psi_T^n\right]\right]} \exp\left[\sum_{i=h+1}^{k} A_{T_i}\right] X_{tT_{h+1}}\left[\Xi_{T_{h+1}}^{T_kT}\right] \\
&\quad \times E_{tT_{h+1}}\left[\Xi_{T_{h+1}}^{T_kT}\right]\left(\prod_{i=h+2}^{k} E_{T_{i-1}T_i}\left[\Xi_{T_i}^{T_kT}\right]\right). \qquad (6.3)
\end{aligned}
$$

The expression for the drift term $\exp\left[\sum_{i=h+1}^{k} A_{T_i}\right]$ was given in Proposition 6.2.1, and the parameter transformation Ξ^{T_kT} is defined iteratively in the following way:

$$
\begin{aligned}
\Xi_{T_k}^{T_kT} &:= \Theta_{T_kT}\left[\phi_T^n, \Psi_T^n\right] + \left(\phi_{T_k}^y, \Psi_{T_k}^y\right) \\
\Xi_{T_i}^{T_kT} &:= \Theta_{T_iT_{i+1}}\left[\Xi_{T_{i+1}}^{T_kT}\right] + \left(\phi_{T_i}^y, \Psi_{T_i}^y\right), \quad \textit{for } i = h+1, h+2, \ldots, k-1.
\end{aligned}
$$

The transformation $\Xi_{T_k}^{T_kT}$ for the swaplet is slightly different from the transformation $\Xi_{T_i}^{T_k}$ for the drift in Proposition 6.2.1. This is because of the payment delay $T - T_k$ and because we only need to iterate the expectations down to T_{h+1}.

6.4 Inflation caplet

An inflation ratio I_{T_k} / I_{T_h}, $h < k$ that covers multiple periods of the tenor structure $\{T_i\}_{i\in\mathbb{N}}$ is driven in the QGY model by the driving process x_{T_i} for all i so that $h < i \leq k$. To simplify calculations, we therefore look at the case where $h = k - 1$ so that $I_{T_k} / I_{T_h} = I_{T_k} / I_{T_{k-1}} = Y_{T_k}$ is only driven by x_{T_k}. Let $0 \leq t < T_k \leq T$. The value of a one-period inflation caplet with a payment delay $T - T_k$ is given at time t by

$$
V_t := \frac{1}{D_t^n}\mathbb{E}_t^n\left[D_T^n\left(\frac{I_{T_k}}{I_{T_{k-1}}} - K\right)^+\right] = \frac{1}{D_t^n}\mathbb{E}_t^n\left[D_T^n\left[\phi_T^n, \Psi_T^n\right]\left(Y_{T_k}\left[\phi_{T_k}^y, \Psi_{T_k}^y\right] - K\right)^+\right].
$$

Using the tower property of conditional expectations, along with the formula for the nominal inverse numeraire in (6.1), we get

$$
V_t = \frac{1}{D_t^n}\frac{P_{0T}^n}{P_{0T_k}^n}\mathbb{E}_t^n\left[D_{T_k}^n\left[\Theta_{T_kT}\left[\phi_T^n, \Psi_T^n\right]\right]\left(Y_{T_k}\left[\phi_{T_k}^y, \Psi_{T_k}^y\right] - K\right)^+\right].
$$

Now $D^n_{T_k}$ is a log-quadratic Gaussian process, so combining (5.5) and (5.6) we get

$$\begin{aligned}
V_t &= \frac{1}{D^n_t}\frac{P^n_{0T}}{P^n_{0T_k}}\ \mathbb{E}^n_t\left[D^n_{T_k}\left[M_{T_kT}\left[\Psi^n_T\right]\phi^n_T, M_{T_kT}\left[\Psi^n_T\right]\Psi^n_T\right]\right]\ \mathbb{E}^n_t\left[\left(\widetilde{Y}_{T_k}\left[\phi^y_{T_k},\Psi^y_{T_k}\right]-K\right)^+\right] \\
&= \frac{1}{D^n_t}\ \mathbb{E}^n_t\left[D^n_T\left[\phi^n_T,\Psi^n_T\right]\right]\ \mathbb{E}^n_t\left[\left(\widetilde{Y}_{T_k}\left[\phi^y_{T_k},\Psi^y_{T_k}\right]-K\right)^+\right] \\
&= P^n_{tT}\ \mathbb{E}^n_t\left[\left(\widetilde{Y}_{T_k}\left[\phi^y_{T_k},\Psi^y_{T_k}\right]-K\right)^+\right]. \qquad (6.4)
\end{aligned}$$

Here, $\widetilde{Y}_{T_k}\left[\phi^y_{T_k},\Psi^y_{T_k}\right]$ is obtained from $Y_{T_k}\left[\phi^y_{T_k},\Psi^y_{T_k}\right]$ by replacing x_{T_k} with $\widetilde{x}_{T_k}$:

$$\widetilde{x}_{T_k} := \widetilde{x}_t + \sqrt{\widetilde{G}_{tT_k}}\sqrt{G_{tT}}^{-1}\left(x_{T_k}-x_t\right).$$

It is a Gaussian process, which, conditional on t, has mean $\widetilde{x}_t$ and variance $\widetilde{G}_{tT_k}$ under $\mathbb{Q}^n$:

$$\widetilde{x}_t := \widetilde{G}_{tT_k}\left(G^{-1}_{tT_k}x_t - M_{T_kT}\left[\Psi^n_T\right]\phi^n_T\right), \qquad (6.5a)$$

$$\sqrt{\widetilde{G}_{tT_k}} := \sqrt{G_{tT_k}}\left(1+\sqrt{G_{tT_k}}^{\intercal}M_{T_kT}\left[\Psi^n_T\right]\Psi^n_T\sqrt{G_{tT_k}}\right)^{-\frac{1}{2}}, \qquad (6.5b)$$

$$\widetilde{G}_{tT_k} := \sqrt{\widetilde{G}_{tT_k}}\sqrt{\widetilde{G}_{tT_k}}^{\intercal} = G_{tT_k}\left(1+M_{T_kT}\left[\Psi^n_T\right]\Psi^n_T G_{tT_k}\right)^{-1}. \qquad (6.5c)$$

Now $\widetilde{Y}_{T_k}$ is a log-quadratic Gaussian process,

$$\widetilde{Y}_{T_k}\left[\phi^y_{T_k},\Psi^y_{T_k}\right] = \frac{I_{0T_k}}{I_{0T_{k-1}}}\exp\left[A_{T_k}-\phi^y_{T_k}\cdot\widetilde{x}_{T_k}-\tfrac{1}{2}\widetilde{x}_{T_k}\cdot\Psi^y_{T_k}\widetilde{x}_{T_k}\right],$$

so by applying (5.5) to (6.4), we get

$$\begin{aligned}
\frac{V_t}{P^n_{tT}} &= \mathbb{E}^n_t\left[\widetilde{Y}_{T_k}\left[\phi^y_{T_k},\Psi^y_{T_k}\right]\mathbb{1}_{\left\{\widetilde{Y}_{T_k}\left[\phi^y_{T_k},\Psi^y_{T_k}\right]>K\right\}}\right] - K\ \mathbb{E}^n_t\left[\mathbb{1}_{\left\{\widetilde{Y}_{T_k}\left[\phi^y_{T_k},\Psi^y_{T_k}\right]>K\right\}}\right] \\
&= \mathbb{E}^n_t\left[\widetilde{Y}_{T_k}\left[\phi^y_{T_k},\Psi^y_{T_k}\right]\right]\ \mathbb{P}^n_t\left[\widetilde{\widetilde{Y}}_{T_k}\left[\phi^y_{T_k},\Psi^y_{T_k}\right]>K\right] - K\ \mathbb{P}^n_t\left[\widetilde{Y}_{T_k}\left[\phi^y_{T_k},\Psi^y_{T_k}\right]>K\right]. \qquad (6.6)
\end{aligned}$$

Here, $\widetilde{\widetilde{Y}}_{T_k}\left[\phi^y_{T_k},\Psi^y_{T_k}\right]$ is obtained from $\widetilde{Y}_{T_k}\left[\phi^y_{T_k},\Psi^y_{T_k}\right]$ by replacing $\widetilde{x}_{T_k}$ with $\widetilde{\widetilde{x}}_{T_k}$:

$$\widetilde{Y}_{T_k}\left[\phi^y_{T_k},\Psi^y_{T_k}\right] = \frac{I_{0T_k}}{I_{0T_{k-1}}}\exp\left[A_{T_k}-\phi^y_{T_k}\cdot\widetilde{\widetilde{x}}_{T_k}-\tfrac{1}{2}\widetilde{\widetilde{x}}_{T_k}\cdot\Psi^y_{T_k}\widetilde{\widetilde{x}}_{T_k}\right],$$

$$\widetilde{\widetilde{x}}_{T_k} := \widetilde{\widetilde{x}}_t + \sqrt{\widetilde{\widetilde{G}}_{tT_k}}\sqrt{\widetilde{G}_{tT_K}}^{-1}\left(\widetilde{x}_{T_k}-\widetilde{x}_t\right).$$

Conditional on t, the Gaussian process $\widetilde{\widetilde{x}}_{T_k}$ has mean $\widetilde{\widetilde{x}}_t$ and variance $\widetilde{\widetilde{G}}_{tT_k}$ under $\mathbb{Q}^n$:

$$\widetilde{\widetilde{x}}_t := \widetilde{\widetilde{G}}_{tT_k}\left(\widetilde{G}^{-1}_{tT_k}\widetilde{x}_t - \phi^y_{T_k}\right), \qquad (6.7a)$$

$$\sqrt{\widetilde{\widetilde{G}}_{tT_k}} := \sqrt{\widetilde{G}_{tT_k}}\left(1+\sqrt{\widetilde{G}_{tT_k}}^{\intercal}\Psi^y_{T_k}\sqrt{\widetilde{G}_{tT_k}}\right)^{-\frac{1}{2}}, \qquad (6.7b)$$

$$\widetilde{\widetilde{G}}_{tT_k} := \sqrt{\widetilde{\widetilde{G}}_{tT_k}}\sqrt{\widetilde{\widetilde{G}}_{tT_k}}^{\intercal} = \widetilde{G}_{tT_k}\left(1+\Psi^y_{T_k}\widetilde{G}_{tT_k}\right)^{-1}. \qquad (6.7c)$$

We note that the following variable, z, has a standard multi-factor Gaussian distribution:

$$z := \sqrt{G_{tT}}^{-1}\left(x_{T_k} - x_t\right) = \sqrt{\widetilde{G}_{tT_K}}^{-1}\left(\widetilde{x}_{T_k} - \widetilde{x}_t\right) = \sqrt{\widetilde{\widetilde{G}}_{tT_K}}^{-1}\left(\widetilde{\widetilde{x}}_{T_k} - \widetilde{\widetilde{x}}_t\right).$$

Writing $\widetilde{x}_{T_k} = \widetilde{x}_t + \sqrt{\widetilde{G}_{tT_k}}z$, the probability term in (6.6) is given by

$$\begin{aligned}
&\mathbb{P}_t^n\left[\widetilde{Y}_{T_k}\left[\phi_{T_k}^y, \Psi_{T_k}^y\right] > K\right] \\
&= \mathbb{P}_t^n\left[\phi_{T_k}^y \cdot \widetilde{x}_{T_k} + \tfrac{1}{2}\widetilde{x}_{T_k} \cdot \Psi_{T_k}^y \widetilde{x}_{T_k} < \ln\left[\frac{I_{0T_k}}{I_{0T_{k-1}}}\frac{\exp\left[A_{T_k}\right]}{K}\right]\right] \\
&= \mathbb{P}_t^n\Bigg[\phi_{T_k}^y \cdot \left(\widetilde{x}_t + \sqrt{\widetilde{G}_{tT_k}}z\right) + \tfrac{1}{2}\left(\widetilde{x}_t + \sqrt{\widetilde{G}_{tT_k}}z\right) \cdot \Psi_{T_k}^y\left(\widetilde{x}_t + \sqrt{\widetilde{G}_{tT_k}}z\right) \\
&\qquad < \ln\left[\frac{I_{0T_k}}{I_{0T_{k-1}}}\frac{\exp\left[A_{T_k}\right]}{K}\right]\Bigg] \\
&= \mathcal{N}\left[\widetilde{\mathcal{D}}\right],
\end{aligned}$$

where we have denoted the standard Gaussian distribution function over some domain $\mathcal{D} \in \mathbb{R}^{dim}$ with

$$\mathcal{N}\left[\mathcal{D}\right] := (2\pi)^{-\frac{dim}{2}}\int_{\mathcal{D}}\exp\left[-\tfrac{1}{2}z \cdot z\right]dz, \tag{6.8}$$

and $\widetilde{\mathcal{D}}$ is a quadratic domain:

$$\begin{aligned}
\widetilde{\mathcal{D}} = \Bigg\{z \in \mathbb{R}^{dim} : &\left(\sqrt{\widetilde{G}_{tT_k}}^{\intercal}\left(\phi_{T_k}^y + \Psi_{T_k}^y\widetilde{x}_t\right)\right) \cdot z + \tfrac{1}{2}z \cdot \left(\sqrt{\widetilde{G}_{tT_k}}^{\intercal}\Psi_{T_k}^y\sqrt{\widetilde{G}_{tT_k}}\right)z \\
&< \ln\left[\frac{I_{0T_k}}{I_{0T_{k-1}}}\frac{\exp\left[A_{T_k}\right]}{K}\right] - \phi_{T_k}^y \cdot \widetilde{x}_t - \tfrac{1}{2}\widetilde{x}_t \cdot \Psi_{T_k}^y\widetilde{x}_t\Bigg\}.
\end{aligned} \tag{6.9}$$

$\widetilde{x}_t$ and $\sqrt{\widetilde{G}_{tT_K}}$ are given in (6.5a) and (6.5b). Writing $\widetilde{\widetilde{x}}_{T_k} = \widetilde{\widetilde{x}}_t + \sqrt{\widetilde{\widetilde{G}}_{tT_k}}z$, we get in a similar way

$$\mathbb{P}_t^n\left[\widetilde{\widetilde{Y}}_{T_k}\left[\phi_{T_k}^y, \Psi_{T_k}^y\right] > K\right] = \mathcal{N}\left[\widetilde{\widetilde{\mathcal{D}}}\right],$$

where the quadratic domain $\widetilde{\widetilde{\mathcal{D}}}$ is

$$\begin{aligned}
\widetilde{\widetilde{\mathcal{D}}} = \Bigg\{z \in \mathbb{R}^{dim} : &\left(\sqrt{\widetilde{\widetilde{G}}_{tT_k}}^{\intercal}\left(\phi_{T_k}^y + \Psi_{T_k}^y\widetilde{\widetilde{x}}_t\right)\right) \cdot z + \tfrac{1}{2}z \cdot \left(\sqrt{\widetilde{\widetilde{G}}_{tT_k}}^{\intercal}\Psi_{T_k}^y\sqrt{\widetilde{\widetilde{G}}_{tT_k}}\right)z \\
&< \ln\left[\frac{I_{0T_k}}{I_{0T_{k-1}}}\frac{\exp\left[A_{T_k}\right]}{K}\right] - \phi_{T_k}^y \cdot \widetilde{\widetilde{x}}_t - \tfrac{1}{2}\widetilde{\widetilde{x}}_t \cdot \Psi_{T_k}^y\widetilde{\widetilde{x}}_t\Bigg\}.
\end{aligned} \tag{6.10}$$

$\widetilde{\widetilde{x}}_t$ and $\sqrt{\widetilde{\widetilde{G}}_{tT_k}}$ are given in (6.7a) and (6.7b). From (6.6), we now have the following proposition.

Proposition 6.4.1. *The value of a one-period inflation caplet with a payment delay $T - T_k$ at time t can then be written as*

$$\frac{1}{D_t^n}\mathbb{E}_t^n\left[D_T^n\left(\frac{I_{T_k}}{I_{T_{k-1}}} - K\right)^+\right] = \frac{1}{D_t^n}\mathbb{E}_t^n\left[D_T^n Y_{T_k}\right]\mathcal{N}\left[\widetilde{\widetilde{\mathcal{D}}}\right] - KP_{tT}^n\mathcal{N}\left[\widetilde{\mathcal{D}}\right]. \quad (6.11)$$

The formula is of a Black-Scholes-type (1973). $\mathcal{N}[\cdot]$ *is defined in* (6.8) *and the quadratic domains* $\widetilde{D}$ *and* $\widetilde{\widetilde{D}}$ *are given in* (6.9) *and* (6.10)*. The time t value of the year-on-year inflation ratio Y_{T_k} is given by*

$$\frac{1}{D_t^n}\mathbb{E}_t^n\left[D_T^n Y_{T_k}\right] = P_{tT}^n\frac{I_{0T_k}}{I_{0T_{k-1}}}\exp\left[A_{T_k}\right]\widetilde{E}_{tT_k}\left[\phi_{T_k}^y, \Psi_{T_k}^y\right]\widetilde{X}_{tT_k}\left[\phi_{T_k}^y, \Psi_{T_k}^y\right]. \quad (6.12)$$

The functions $\widetilde{E}_{tT_k}$ and $\widetilde{X}_{tT_k}$ are obtained from the functions E_{tT_k} and X_{tT_k} defined in §5.2 through the replacement of (x_t, G_{tT_k}) with $\left(\widetilde{x}_t, \widetilde{G}_{tT_k}\right)$ in (6.5a) *and* (6.5c)*:*

$$\begin{aligned}
\widetilde{E}_{tT_k}\left[\phi_{T_k}^y, \Psi_{T_k}^y\right] &= \det\left[\widetilde{M}_{tT_k}\left[\Psi_{T_k}^y\right]\right]^{-\frac{1}{2}}\exp\left[\tfrac{1}{2}\phi_{T_k}^y \cdot \widetilde{G}_{tT_k}\widetilde{M}_{tT_k}\left[\Psi_{T_k}^y\right]\phi_{T_k}^y\right], \\
\widetilde{X}_{tT_k}\left[\phi_{T_k}^y, \Psi_{T_k}^y\right] &= \exp\left[-\widetilde{M}_{tT_k}\left[\widetilde{\Psi}_{T_k}^y\right]\widetilde{\phi}_{T_k}^y \cdot \widetilde{x}_t - \tfrac{1}{2}\widetilde{x}_t \cdot \widetilde{M}_{tT_k}\left[\Psi_{T_k}^y\right]\Psi_{T_k}^y\widetilde{x}_t\right], \\
\widetilde{M}_{tT_k}\left[\Psi_{T_k}^y\right] &= \left(1 + \Psi_{T_k}^y\widetilde{G}_{tT_k}\right)^{-1}.
\end{aligned}$$

For the one-period inflation floorlet with a payment delay $T - T_k$, we use the put-call parity

$$(K - Y_{T_k})^+ = (Y_{T_k} - K)^+ - (Y_{T_k} - K)$$

and get

$$\frac{1}{D_t^n}\mathbb{E}_t^n\left[D_T^n\left(K - \frac{I_{T_k}}{I_{T_{k-1}}}\right)^+\right] = -\frac{1}{D_t^n}\mathbb{E}_t^n\left[D_T^n Y_{T_k}\right]\mathcal{N}\left[\widetilde{\widetilde{\mathcal{D}}}^c\right] + KP_{tT}^n\mathcal{N}\left[\widetilde{\mathcal{D}}^c\right].$$

We note that $1 - \mathcal{N}[D] = \mathcal{N}[D^c]$, where D^c denotes the compliment of the domain D.

Gaussian integration

When we derived the caplet formula (6.11), we obtained the domains $\widetilde{\mathcal{D}}$ and $\widetilde{\widetilde{\mathcal{D}}}$ in (6.9) and (6.10). They appeared because we had to determine for which values of the underlying driving process the year-on-year inflation ratio was larger than the strike of the caplet. These domains can be described by a multi-factor quadratic inequality and are of the form

$$\mathcal{D} = \left\{ z \in \mathbb{R}^{dim} : \phi \cdot z + z \cdot \Psi z + F < 0 \right\}$$

for some vector ϕ and matrix Ψ. In the formula for the price of the caplet in (6.11), we need to integrate a multivariate standard Gaussian distribution over the domain $\mathcal{D}$. In Appendix E, we detail a method to simplify this integral in two dimensions. We change the variables of the integral using polar coordinates and get

$$\begin{aligned}
\mathcal{N}\left[\mathcal{D}\right] &= \frac{1}{2\pi} \int_{-\infty}^{+\infty} \int_{-\infty}^{+\infty} \mathrm{e}^{-\frac{1}{2}\left(x^2+y^2\right)} \mathbb{1}_{\{Ax^2+Bxy+Cy^2+Dx+Ey+F<0\}} \, dxdy \\
&= \frac{1}{2\pi} \sum_{i=1}^{N-1} \int_{\theta_i}^{\theta_{i+1}} \left[-\,\mathrm{e}^{-\frac{1}{2}r^2} \mathbb{1}_{\{P[\theta,r]<0\}} \right]_{r=0}^{+\infty} d\theta. \qquad (6.13)
\end{aligned}$$

Here $\theta_N - \theta_1 = 2\pi$. We have defined the following quadratic function in r:

$$P\left[\theta, r\right] := \left(A\cos^2\theta + B\sin\theta\cos\theta + C\sin^2\theta\right) r^2 + \left(D\cos\theta + E\sin\theta\right) r + F.$$

We choose the angles θ_i, $i = 1, \ldots, N$ so that the coefficients of $P\left[\theta, r\right]$ do not change sign for $\theta \in [\theta_i, \theta_{i+1})$. Then it is possible to calculate the roots of $P\left[\theta, r\right]$ and to determine for which $r \in [0, +\infty)$ the quadratic inequality $P\left[\theta, r\right] < 0$. The integrand is then given in closed form. We have therefore reduced the two-dimensional integral to a finite sum of one-dimensional semi-analytical integrals. Each of these integrals can be approximated accurately using some numerical integration methods. From an example in Appendix E, we see that as few as 336 evaluations of the integrand in (6.13) are required to evaluate the integral with an error of less than 10^{-10}.

7 Three-factor reparameterisation

The driving processes in the QGY model can be of an arbitrary dimension, giving us a lot of flexibility when it comes to choosing the appropriate distribution and correlation structure of the inflation index and the year-on-year inflation rate. In this chapter, we find a low-dimensional version of the QGY model, which still captures the main distributional properties of the year-on-year inflation rate in the QGY model.

In §7.1, we describe a three-factor parameter reduction of the QGY model. In §7.2, we analyse the terminal distribution of the year-on-year inflation ratio. We are able both to reduce the number of parameters and to assume the sign of the remaining parameters, without losing the main distributional properties of the three-factor model.

In §7.3, we specify a stochastic-volatility representation of the parameters of the year-on-year inflation ratio. In §7.4, we define an alternative parameterisation that has three spherical parameters for the volatility level, the curvature of the smile and the skew of the smile. We show the behaviour of these parameters in §7.5, where we also discover the maximum curvature that is obtainable for the reduced model, as well as the maximum and minimum skew.

In §7.6, we show how to control the autocorrelation of the year-on-year inflation ratio process and, in turn, the year-on-year convexity correction of the year-on-year inflation forward rate. In §7.7, we show how to control the correlation between inflation rates and nominal interest rates.

7.1 Three-factor reduction

The modelling processes in the QGY model are defined in Chapter 5 by

$$D_t^n\left[\phi_t^n, \Psi_t^n\right] := P_{0t}^n \frac{\exp\left[-\phi_t^n \cdot x_t - \frac{1}{2} x_t \cdot \Psi_t^n x_t\right]}{\mathbb{E}_0^n\left[\exp\left[-\phi_t^n \cdot x_t - \frac{1}{2} x_t \cdot \Psi_t^n x_t\right]\right]},$$

$$Y_{T_k}\left[\phi_{T_k}^y, \Psi_{T_k}^y\right] := \frac{I_{0T_k}}{I_{0T_{k-1}}} \exp\left[A_{T_k} - \phi_{T_k}^y \cdot x_{T_k} - \frac{1}{2} x_{T_k} \cdot \Psi_{T_k}^y x_{T_k}\right].$$

Within the three-factor reduction presented here, we let one factor drive the inverse numeraire D_t^n (effectively the nominal bond) and two factors drive the year-on-year inflation ratio Y_{T_k}. We further assume that the nominal part of the model is a log-linear Gaussian process, setting $\Psi_t^n = 0$, and define a three-dimensional parameter separation for $t \geq 0$ and $k = 1, 2, \ldots$ as follows:

$$\phi_t^n := \begin{bmatrix} \phi_t^{n_1} \\ 0 \\ 0 \end{bmatrix}, \quad \phi_{T_k}^y := \begin{bmatrix} 0 \\ \phi_{T_k}^{y_1} \\ \phi_{T_k}^{y_2} \end{bmatrix} = \begin{bmatrix} 0 \\ \phi_{T_k}^{y_1} \\ 0 \end{bmatrix},$$

$$\Psi_{T_k}^y := \begin{bmatrix} 0 & 0 & 0 \\ 0 & \Psi_{T_k}^{y_1} & \Psi_{T_k}^{y_1 y_2} \\ 0 & \Psi_{T_k}^{y_1 y_2} & \Psi_{T_k}^{y_2} \end{bmatrix} = \begin{bmatrix} 0 & 0 & 0 \\ 0 & \Psi_{T_k}^{y_1} & \Psi_{T_k}^{y_1 y_2} \\ 0 & \Psi_{T_k}^{y_1 y_2} & 0 \end{bmatrix}.$$

We recall that x_t is a Gaussian process that is driven by a multi-factor Brownian motion W_t with correlation matrix ρ_t. Here we define these processes as follows:

$$x_t := \begin{bmatrix} x_t^n \\ x_t^{y_1} \\ x_t^{y_2} \end{bmatrix} = \int_0^t \operatorname{diag}\left[\sigma_s\right] dW_s = \begin{bmatrix} \int_0^t \sigma_s^n dW_s^n \\ \int_0^t \sigma_s^{y_1} dW_s^{y_1} \\ \int_0^t \sigma_s^{y_2} dW_s^{y_2} \end{bmatrix}, \qquad \sigma_t := \begin{bmatrix} \sigma_t^n \\ \sigma_t^{y_1} \\ \sigma_t^{y_2} \end{bmatrix} = \begin{bmatrix} 1 \\ \sigma_t^y \\ \sigma_t^y \end{bmatrix},$$

$$W_t := \begin{bmatrix} W_t^n \\ W_t^{y_1} \\ W_t^{y_2} \end{bmatrix}, \qquad \rho_t := \begin{bmatrix} 1 & \rho_t^{ny_1} & \rho_t^{ny_2} \\ \rho_t^{ny_1} & 1 & \rho_t^{y_1 y_2} \\ \rho_t^{ny_2} & \rho_t^{y_1 y_2} & 1 \end{bmatrix} = \begin{bmatrix} 1 & \rho_t^{ny_1} & 0 \\ \rho_t^{ny_1} & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix}.$$

The variance matrix of x_t is given by G_t, which for three factors is defined by

$$G_t := \begin{bmatrix} G_t^n & G_t^{ny_1} & G^{ny_2} \\ G_t^{ny_1} & G_t^{y_1} & G^{y_1 y_2} \\ G^{ny_2} & G^{y_1 y_2} & G_t^{y_2} \end{bmatrix} = \begin{bmatrix} \int_0^t (\sigma_s^n)^2 ds & \int_0^t \rho_s^{ny_1} \sigma_s^n \sigma_s^{y_1} ds & \int_0^t \rho_s^{y_1 y_2} \sigma_s^{y_1} \sigma_s^{y_2} ds \\ \int_0^t \rho_s^{ny_1} \sigma_s^n \sigma_s^{y_1} ds & \int_0^t (\sigma^{y_1})^2 ds & \int_0^t \rho_s^{ny_2} \sigma_s^n \sigma_s^{y_2} ds \\ \int_0^t \rho_s^{y_1 y_2} \sigma_s^{y_1} \sigma_s^{y_2} ds & \int_0^t \rho_s^{ny_2} \sigma_s^n \sigma_s^{y_2} ds & \int_0^t (\sigma_s^{y_2})^2 ds \end{bmatrix}$$

$$= \begin{bmatrix} t & \int_0^t \rho_s^{ny_1} \sigma_s^y ds & 0 \\ \int_0^t \rho_s^{ny_1} \sigma_s^y ds & \int_0^t (\sigma_s^y)^2 ds & 0 \\ 0 & 0 & \int_0^t (\sigma_s^y)^2 ds \end{bmatrix}.$$

The modelling processes can now be written as

$$D_t^n = P_{0t}^n \exp\left[-\phi_t^{n_1} x_t^n - \frac{1}{2} \left(\phi_t^{n_1}\right)^2 G_t^n\right], \tag{7.1}$$

$$Y_{T_k} = \frac{I_{0T_k}}{I_{0T_{k-1}}} \exp\left[A_{T_k} - \left(\phi_{T_k}^{y_1} + \frac{1}{2} \Psi_{T_k}^{y_1} x_{T_k}^{y_1} + \Psi_{T_k}^{y_1 y_2} x_{T_k}^{y_2}\right) x_{T_k}^{y_1}\right]. \tag{7.2}$$

As we are mainly interested in pricing inflation-linked options, we have assumed that the inverse numeraire D_t^n is a one-factor log-normal process, which does not give any curvature or skew in the nominal interest-rate volatility smile. This assumption does not limit the control we have on the inflation volatility smile. Furthermore, we can still control the correlation between inflation rates and nominal interest rates.

The year-on-year inflation ratio Y_{T_k} is now only driven by two Gaussian processes $(x_t^{y_1}, x_t^{y_2})$. We obtained a semi-analytical pricing formula for the one-period inflation caplet in §6.4. We can use this formula to price year-on-year options, assuming that we have an annual tenor structure $T \in \{T_k\}_{k=0,1,2,...}$. This means that $T_k - T_{k-1}$ equals one year for $k = 1, 2, \ldots$. For the pricing formula of the year-on-year inflation caplet in Proposition 6.4.1, any matrix multiplications and integrals are now two-dimensional. This makes calculations of the prices of year-on-year options fast.

We implemented the three-factor QGY model in the C++ quantitative pricing library at Lloyds Bank. The pricing library is linked to the Excel spreadsheet application in which we were easily able to change the model parameters and get the prices of inflation-linked swaps and year-on-year inflation options. The implementation in C++ allowed for fast pricing, which, together with the convenience of working with data in Excel, made it easy to analyse how the parameters affected the prices and, in turn, to analyse the distributional properties of the inflation index and the year-on-year inflation rate. The main properties are the year-on-year inflation volatility smile, the year-on-year inflation autocorrelation and the correlation between inflation and nominal interest rates. In the remainder of the chapter, we look further at the model parameters and see how we can control each property with a set of intuitive parameters.

7.2 Terminal distribution

In the definition of the three-factor reduction, we assumed that some of the model parameters were equal to zero. These parameters were $\phi_{T_k}^{y_2}$, $\Psi_{T_k}^{y_2}$, $\rho_t^{ny_2}$, and $\rho_t^{y_1 y_2}$. We did this because we wanted to be able to control the shape of the year-on-year smile in an intuitive way without having to adjust too many parameters. Essentially, $\phi_{T_k}^{y_1}$ contributes to the volatility level, $\Psi_{T_k}^{y_1 y_2}$ to the curvature, and $\Psi_{T_k}^{y_1}$ to the skew. To elaborate further on these assumptions, we start by looking at the quadratic driving process in the exponent of Y_{T_k}:

$$\begin{aligned} r^y_{T_k} &:= -\phi^y_{T_k}\cdot x_{T_k} - \tfrac{1}{2}x_{T_k}\cdot\Psi^y_{T_k}x_{T_k} \\ &= -\phi^{y_1}_{T_k}x^{y_1}_{T_k} - \phi^{y_2}_{T_k}x^{y_2}_{T_k} + \tfrac{1}{2}\Psi^{y_1}_{T_k}\left(x^{y_1}_{T_k}\right)^2 + \tfrac{1}{2}\Psi^{y_2}_{T_k}\left(x^{y_2}_{T_k}\right)^2 + \Psi^{y_1y_2}_{T_k}x^{y_2}_{T_k}x^{y_1}_{T_k}x^{y_2}_{T_k} \\ &= -\left(\phi^{y_1}_{T_k} + \tfrac{1}{2}\Psi^{y_1}_{T_k}x^{y_1}_{T_k} + \Psi^{y_1y_2}_{T_k}x^{y_2}_{T_k}\right)x^{y_1}_{T_k}. \end{aligned}$$

The linear term

$$\phi^y_{T_k}\cdot x_{T_k} = \phi^{y_1}_{T_k}x^{y_1}_{T_k} + \phi^{y_2}_{T_k}x^{y_2}_{T_k}$$

is a Gaussian process and does not contribute to the curvature or the skew of the year-on-year inflation volatility smile but only to its level. We therefore assume that $\phi^{y_2}_{T_k} = 0$, leaving non-zero $\phi^{y_1}_{T_k}$ to control the volatility level. The quadratic term

$$x_{T_k}\cdot\Psi^y_{T_k}x_{T_k} = \Psi^{y_1}_{T_k}\left(x^{y_1}_{T_k}\right)^2 + \Psi^{y_2}_{T_k}\left(x^{y_2}_{T_k}\right)^2 + 2\Psi^{y_1y_2}_{T_k}x^{y_1}_{T_k}x^{y_2}_{T_k}$$

contributes to both the curvature and the skew of the year-on-year smile, as well as to the volatility level. The term $\Psi^{y_1y_2}_{T_k}x^{y_1}_{T_k}x^{y_2}_{T_k}$ adds curvature but no skew. The term $\Psi^{y_1}_{T_k}\left(x^{y_1}_{T_k}\right)^2 + \Psi^{y_2}_{T_k}\left(x^{y_2}_{T_k}\right)^2$ follows a weighted chi-squared distribution and gives good control over the skew. To reduce the number of parameters, we set $\Psi^{y_2}_{T_k} = 0$ and use $\Psi^{y_1}_{T_k}$ to adjust the skew. We have assumed that the correlation between $x^{y_1}_{T_k}$ and $x^{y_2}_{T_k}$ is zero, i.e. $\rho^{y_1y_2}_t = 0$, because the parameters $\Psi^{y_1y_2}_{T_k}$ and $\Psi^{y_1}_{T_k}$ give more control than $\rho^{y_1y_2}_t$ over the terminal distribution of Y_{T_k}.[1]

We would like be able to control the correlation between inflation rates and nominal interest rates, as it affects the payment delay convexity correction incorporated in many inflation-linked derivatives. The covariance between $\ln D^n_t$ and $\ln Y_{T_k}$ is given by[2]

$$\begin{aligned} \mathrm{Cov}^n_0\left[\ln D^n_t, \ln Y_{T_k}\right] &= \mathbb{E}^n_0\left[\left(\phi^n_t\cdot x_t\right)\left(\phi^y_{T_k}\cdot x_{T_k} + \tfrac{1}{2}x_{T_k}\cdot\Psi^y_{T_k}x_{T_k}\right)\right] \\ &= \phi^n_t\cdot G_{t\wedge T_k}\phi^y_{T_k} = \phi^{n_1}_t\phi^{y_1}_{T_k}G^{ny_1}_{t\wedge T_k}. \end{aligned} \tag{7.3}$$

[1]If $x^{y_1}_{T_k}$ and $x^{y_2}_{T_k}$ are correlated with correlation ρ, we can replace $x^{y_2}_{T_k}$ with a new Gaussian process $\widetilde{x}^{y_2}_{T_k}$ with mean zero and variance $G^{y_2}_{T_k}$, independent of $x^{y_1}_{T_k}$, so that

$$x^{y_2}_{T_k} = \rho x^{y_1}_{T_k} + \sqrt{1-\rho^2}\,\widetilde{x}^{y_2}_{T_k}.$$

The quadratic term would still be of the same form as before (through redefinition of $\Psi^{y_1}_{T_k}$ and $\Psi^{y_1y_2}_{T_k}$):

$$\tfrac{1}{2}\Psi^{y_1}_{T_k}x^{y_1}_{T_k} + \Psi^{y_1y_2}_{T_k}x^{y_2}_{T_k} = \tfrac{1}{2}\left(\Psi^{y_1}_{T_k} + 2\rho\Psi^{y_1y_2}_{T_k}\right)x^{y_1}_{T_k} + \left(\sqrt{1-\rho^2}\,\Psi^{y_1y_2}_{T_k}\right)\widetilde{x}^{y_2}_{T_k}.$$

[2]The covariance does not involve the parameters of the matrix $\Psi^y_{T_k}$, because the third moments are zero, i.e. $\mathbb{E}_0\left[x^n_t x^{y_i}_{T_k}x^{y_j}_{T_k}\right] = 0$, $i,j\in\{1,2\}$.

In the final equality, we make the assumption that $\phi_t^{y_2} = 0$ so the covariance does not involve G^{ny_2}. We may therefore assume that $\rho_t^{ny_2} = 0$ and use $\rho_t^{ny_1} \in [-1, 1]$ to control the nominal-inflation correlation.

We can also make some assumptions concerning the signs of the parameters. Looking at the linear term $\phi_{T_k}^{n_1} x_{T_k}^n$ of the inverse numeraire in (7.1) and noting that x_t^n has the same distribution as $-x_t^n$, we may assume that $\phi_t^n \leq 0$. In the same way, looking at the linear term $\phi_{T_k}^{y_1} x_{T_k}^{y_1}$ of the year-on-year inflation ratio in (7.2) and noting that $x_t^{y_1}$ has the same distribution as $-x_t^{y_1}$, we may assume that $\phi_t^{y_1} \leq 0$. Then $\rho_t^{ny_1}$ gives an intuitive control over the correlation between inflation rates and nominal interest rates in (7.3).

Now $x_{T_k}^{y_1}$ and $x_{T_k}^{y_2}$ are independent and $x_t^{y_2}$ has the same distribution as $-x_t^{y_2}$. Then, from the quadratic term $\Psi_{T_k}^{y_1 y_2} x_{T_k}^{y_1} x_{T_k}^{y_2}$ in (7.2), which controls the curvature of the year-on-year smile, we may assume that $\Psi_t^{y_1 y_2} \geq 0$. The term $\Psi_{T_k}^{y_1} \left(x_{T_k}^{y_2} \right)^2$ gives negative skew if $\Psi_{T_k}^{y_1} > 0$ and positive skew if $\Psi_{T_k}^{y_1} < 0$.

Normalisation

In §2.3, we saw that the market shows a very stable year-on-year smile for different maturities. This is because the absolute volatility (see 2.3) of the year-on-year inflation ratio Y_{T_k} is proportional to the length of the period $(T_{k-1}, T_k]$, which is always one year in length. Each of the variances, $G_t^{y_1}$ and $G_t^{y_2}$, of the two inflation drivers, $x_t^{y_1}$ and $x_t^{y_2}$, increases with t. To emphasise the stability of the volatility level of the year-on-year smile across maturities, we scale these drivers so that they have unit variance. Assuming that $G_{T_k}^{y_1}$ and $G_{T_k}^{y_1}$ are strictly positive, we define the following normalised Gaussian processes for $k = 1, 2, \ldots$:

$$z_{T_k}^{y_1} := x_{T_k}^{y_1} \Big/ \sqrt{G_{T_k}^{y_1}}\,, \qquad z_{T_k}^{y_2} := x_{T_k}^{y_2} \Big/ \sqrt{G_{T_k}^{y_2}}\,.$$

They have zero mean and unit variance and are uncorrelated. We also normalise the original model parameters:

$$\widetilde{\phi}_{T_k}^{y_1} := -\phi_{T_k}^{y_1} \sqrt{G_{T_k}^{y_1}}, \qquad \widetilde{\Psi}_{T_k}^{y_1} := -\frac{1}{2} \Psi_{T_k}^{y_1} G_{T_k}^{y_1}, \qquad \widetilde{\Psi}_{T_k}^{y_1 y_2} := -\Psi_{T_k}^{y_1 y_2} \sqrt{G_{T_k}^{y_1} G_{T_k}^{y_2}}.$$

The driving process of Y_{T_k} is then of the form

$$r_{T_k}^y := -\phi_{T_k}^y \cdot x_{T_k} - \tfrac{1}{2} x_{T_k} \cdot \Psi_{T_k}^y x_{T_k} = \left(\widetilde{\phi}_k^{y_1} + \widetilde{\Psi}_{T_k}^{y_1} z_{T_k}^{y_1} + \widetilde{\Psi}_{T_k}^{y_1 y_2} z_{T_k}^{y_2} \right) z_{T_k}^{y_1}. \tag{7.4}$$

Its variance is given by

$$
\begin{aligned}
Var_0^n\left[r_{T_k}^y\right] &= \mathbb{E}_0^n\left[\left(\widetilde{\phi}_k^{y_1} z_{T_k}^{y_1} + \widetilde{\Psi}_{T_k}^{y_1}\left(z_{T_k}^{y_1}\right)^2 + \widetilde{\Psi}_{T_k}^{y_1 y_2} z_{T_k}^{y_1} z_{T_k}^{y_2}\right)^2\right] \\
&\quad - \mathbb{E}_0^n\left[\widetilde{\phi}_k^{y_1} z_{T_k}^{y_1} + \widetilde{\Psi}_{T_k}^{y_1}\left(z_{T_k}^{y_1}\right)^2 + \widetilde{\Psi}_{T_k}^{y_1 y_2} z_{T_k}^{y_1} z_{T_k}^{y_2}\right]^2 \\
&= \left(\widetilde{\phi}_{T_k}^{y_1}\right)^2 + 3\left(\widetilde{\Psi}_{T_k}^{y_1}\right)^2 + \left(\widetilde{\Psi}_{T_k}^{y_1 y_2}\right)^2 - \left(\widetilde{\Psi}_{T_k}^{y_1}\right)^2 \\
&= \left(\widetilde{\phi}_{T_k}^{y_1}\right)^2 + 2\left(\widetilde{\Psi}_{T_k}^{y_1}\right)^2 + \left(\widetilde{\Psi}_{T_k}^{y_1 y_2}\right)^2 .
\end{aligned}
\tag{7.5}
$$

Next, we will give two reparameterisations of $\left(\widetilde{\phi}_{T_k}^{y_1}, \widetilde{\Psi}_{T_k}^{y_1}, \widetilde{\Psi}_{T_k}^{y_1 y_2}\right)$ that give an intuitive way of controlling the level, curvature and skew of the year-on-year inflation volatility smile.

7.3 Stochastic-volatility parameterisation

From (7.4), we can interpret $z_{T_k}^{y_1}$ as being the main driving factor of the year-on-year inflation ratio Y_{T_k} and then $\widetilde{\phi}_k^{y_1}$ represents the deterministic volatility. The term $\widetilde{\Psi}_{T_k}^{y_1} z_{T_k}^{y_1} + \widetilde{\Psi}_{T_k}^{y_1 y_2} z_{T_k}^{y_2}$ can then be seen as a (linear normal) stochastic-volatility factor with its volatility serving as the volatility-of-volatility $\widetilde{\nu}_{T_k}^y$, and its correlation to $z_{T_k}^{y_1}$ serving as the volatility-correlation $\widetilde{\rho}_{T_k}^y$, that is[3]

$$
\begin{aligned}
\widetilde{\nu}_{T_k}^y &:= -\operatorname{sgn}\left[\widetilde{\Psi}_{T_k}^{y_1 y_2}\right]\sqrt{\operatorname{Var}\left[\widetilde{\Psi}_{T_k}^{y_1} z_{T_k}^{y_1} + \widetilde{\Psi}_{T_k}^{y_1 y_2} z_{T_k}^{y_2}\right]} = -\operatorname{sgn}\left[\widetilde{\Psi}_{T_k}^{y_1 y_2}\right]\sqrt{\left(\widetilde{\Psi}_{T_k}^{y_1}\right)^2 + \left(\widetilde{\Psi}_{T_k}^{y_1 y_2}\right)^2}, \\
\widetilde{\rho}_{T_k}^y &:= \operatorname{Corr}\left[z_{T_k}^{y_1}, \widetilde{\Psi}_{T_k}^{y_1} z_{T_k}^{y_1} + \widetilde{\Psi}_{T_k}^{y_1 y_2} z_{T_k}^{y_2}\right] = \frac{\widetilde{\Psi}_{T_k}^{y_1}}{\left|\widetilde{\nu}_{tT_k}\right|}.
\end{aligned}
$$

In terms of $\widetilde{\nu}_{T_k}^y$ and $\widetilde{\rho}_{T_k}^y \in [-1,1]$ we then have

$$
\widetilde{\Psi}_{T_k}^{y_1 y_2} = -\widetilde{\nu}_{T_k}^y\sqrt{1-\left(\widetilde{\rho}_{T_k}^y\right)^2} \qquad \text{and} \qquad \widetilde{\Psi}_{T_k}^{y_1} = \left|\widetilde{\nu}_{T_k}^y\right|\widetilde{\rho}_{T_k}^y .
$$

The driving process of Y_{T_k} defined in (7.4) can now be written as

$$
\begin{aligned}
r_{T_k}^y &= \left(\widetilde{\phi}_{T_k}^{y_1} + \widetilde{\nu}_{T_k}^y z_{T_k}^y\right) z_{T_k}^{y_1}, \\
\text{with} \quad z_{T_k}^y &= \widetilde{\rho}_{T_k}^y \operatorname{sgn}\left[\widetilde{\nu}_{T_k}^y\right] z_{T_k}^{y_1} - \sqrt{1-\left(\widetilde{\rho}_{T_k}^y\right)^2}\, z_{T_k}^{y_2}.
\end{aligned}
\tag{7.6}
$$

Here $z_{T_k}^y$ is a standard normal variable with $\operatorname{Corr}\left[z_{T_k}^y, z_{T_k}^{y_1}\right] = \widetilde{\rho}_{T_k}^y \operatorname{sgn}\left[\widetilde{\nu}_{T_k}^y\right]$.

[3] The term $\operatorname{sgn}\left[\widetilde{\Psi}_{T_k}^{y_1 y_2}\right]$ in $\widetilde{\nu}_{T_k}^y$ gives a one-to-one correspondence between $\left(\widetilde{\Psi}_{T_k}^{y_1}, \widetilde{\Psi}_{T_k}^{y_1 y_2}\right)$ and $\left(\widetilde{\nu}_{T_k}^y, \widetilde{\rho}_{T_k}^y\right)$.

This parameterisation resembles that of Piterbarg (2009). Much as with other stochastic volatility models, such as the SABR model from Hagan et al. (2002), increasing the volatility-of-volatility $\widetilde{\nu}^{y}_{T_k}$ increases the curvature of the year-on-year smile, while the volatility-correlation $\widetilde{\rho}^{y}_{T_k}$ determines how skewed the smile is. From (7.5), the variance of the year-on-year driver $r^{y}_{T_k}$ in (7.6) is given by

$$\mathrm{Var}^{n}_{0}\left[r^{y}_{T_k}\right] = \left(\widetilde{\phi}^{y_1}_{T_k}\right)^2 + \left(\widetilde{\nu}^{y}_{T_k}\right)^2 \left(1 + \left(\widetilde{\rho}^{y}_{T_k}\right)^2\right),$$

which is an increasing function of the parameters $\widetilde{\phi}^{y_1}_{T_k}$, $\widetilde{\nu}^{y}_{T_k}$ and $\widetilde{\rho}^{y}_{T_k}$. In turn, the volatility level of the year-on-year smile is an increasing function of these parameters. This level is very low in the market, which poses restrictions on the possible combinations of values of $\widetilde{\phi}^{y_1}_{T_k}$, $\widetilde{\nu}^{y}_{T_k}$ and $\widetilde{\rho}^{y}_{T_k}$. This makes it more difficult for the QGY model to match the year-on-year market smile, which has a high curvature and a significant skew.

Trovato, Ribeiro and Gretarsson (2012) use this parameterisation to calibrate the QGY model to the year-on-year smile. The model gives a good qualitative fit to the market prices, while the stochastic-volatility parameters $\widetilde{\nu}^{y}_{T_k}$ and $\widetilde{\rho}^{y}_{T_k}$ give an intuitive control of the curvature and the skew of the smile.

7.4 Spherical parameterisation

The previous parameterisation did not separate the volatility level of the year-on-year smile from its curvature and skew. A better method is to use a three-dimensional spherical transformation. We define new coordinates $\Sigma^{y}_{T_k} \geq 0$, $\nu^{y}_{T_k} \in [0, 2\pi)$, and $\rho^{y}_{T_k} \in [-\pi, \pi)$. For some $K^{y}_{T_k} > 0$, we let

$$\begin{aligned}
\widetilde{\phi}^{y_1}_{T_k} &= \frac{\Sigma^{y}_{T_k}}{K^{y}_{T_k}} \cos\left[\nu^{y}_{T_k}\right], \\
\widetilde{\Psi}^{y_1}_{T_k} &= \frac{\Sigma^{y}_{T_k}}{K^{y}_{T_k}} \sin\left[\nu^{y}_{T_k}\right] \sin\left[\rho^{y}_{T_k}\right], \\
\widetilde{\Psi}^{y_1 y_2}_{T_k} &= -\frac{\Sigma^{y}_{T_k}}{K^{y}_{T_k}} \sin\left[\nu^{y}_{T_k}\right] \cos\left[\rho^{y}_{T_k}\right].
\end{aligned}$$

The driving process of Y_{T_k} in (7.4) then becomes

$$r^{y}_{T_k} = \frac{\Sigma^{y}_{T_k}}{K^{y}_{T_k}} \left\{\cos\left[\nu^{y}_{T_k}\right] + \sin\left[\nu^{y}_{T_k}\right] \left(\sin\left[\rho^{y}_{T_k}\right] z^{y_1}_{T_k} - \cos\left[\rho^{y}_{T_k}\right] z^{y_2}_{T_k}\right)\right\} z^{y_1}_{T_k}.$$

Here $\nu^y_{T_k}$ controls how much weight is put on the quadratic terms $\left(z^{y_1}_{T_k}\right)^2$ and $z^{y_1}_{T_k} z^{y_2}_{T_k}$, and $\rho^y_{T_k}$ controls the weight between the two of them. We interpret

$$\upsilon^y_{T_k} := \sin\left[\nu^y_{T_k}\right]\left(\sin\left[\rho^y_{T_k}\right] z^{y_1}_{T_k} - \cos\left[\rho^y_{T_k}\right] z^{y_2}_{T_k}\right)$$

as the stochastic-volatility driving process. Its volatility (the volatility-of-volatility), which contributes to the curvature of the year-on-year smile, is given by

$$\sqrt{\mathrm{Var}^n_0\left[\upsilon^y_{T_k}\right]} = \left|\sin\left[\nu^y_{T_k}\right]\right|, \tag{7.7}$$

and can be controlled by adjusting $\nu^y_{T_k}$. The correlation between $\upsilon^y_{T_k}$ and the main driving factor $z^{y_1}_{T_k}$ (the volatility-correlation), which contributes to the skew of the year-on-year smile, is given by

$$\mathrm{Corr}^n_0\left[\upsilon^y_{T_k}, z^{y_1}_{T_k}\right] = \mathrm{sgn}\left[\sin\left[\nu^y_{T_k}\right]\right]\sin\left[\rho^y_{T_k}\right] \tag{7.8}$$

and can be controlled by adjusting $\rho^y_{T_k}$. Finally, from (7.5), the variance of the the driving process of Y_{T_k} is given by:

$$\begin{aligned}\mathrm{Var}^n_0\left[r^y_{T_k}\right] &= \left(\frac{\Sigma^y_{T_k}}{K^y_{T_k}}\right)^2\left\{\cos^2\left[\nu^y_{T_k}\right] + 2\sin^2\left[\nu^y_{T_k}\right]\sin^2\left[\rho^y_{T_k}\right] + \sin^2\left[\nu^y_{T_k}\right]\cos^2\left[\rho^y_{T_k}\right]\right\}\\ &= \left(\frac{\Sigma^y_{T_k}}{K^y_{T_k}}\right)^2\left\{\sin^2\left[\nu^y_{T_k}\right]\sin^2\left[\rho^y_{T_k}\right] + 1\right\}.\end{aligned}$$

By setting $\sqrt{\mathrm{Var}^n_0\left[r^y_{T_k}\right]} = \Sigma^y_{T_k}$ and solving for $K^y_{T_k}$, we get

$$K^y_{T_k} = \sqrt{\sin^2\left[\nu^y_{T_k}\right]\sin^2\left[\rho^y_{T_k}\right] + 1}.$$

The parameter $\Sigma^y_{T_k}$ therefore represents the log-volatility of the year-on-year inflation ratio and controls the level of the year-on-year inflation volatility smile.

7.5 Parameter behaviour and smile bounds

In §7.2, we noted that we could assume that $\phi^{y_1}_t \le 0$, or equivalently that $\widetilde{\phi}^{y_1}_{T_k} \ge 0$. In the spherical parameterisation, we therefore have $\cos\nu^y_{T_k} \ge 0$. Furthermore, to keep the same sign of the volatility-correlation in (7.8) and of the term $\sin\rho^y_{T_k}$, we assume $\sin\nu^y_{T_k} \ge 0$. This restricts the domain of $\nu^y_{T_k}$ to

$$\nu^y_{T_k} \in [0, \pi/2] \quad\Longrightarrow\quad \sqrt{\mathrm{Var}^n_0\left[\upsilon^y_{T_k}\right]} = \sin\left[\nu^y_{T_k}\right] \in [0,1].$$

We could also assume that $\Psi_t^{y_1 y_2} \geq 0$, and hence that $\widetilde{\Psi}_{T_k}^{y_1 y_2} \leq 0$. From the definition of the spherical parameters, we therefore have $\cos \rho_{T_k}^y \geq 0$ and restrict the domain of $\rho_{T_k}^y$ to

$$\rho_{T_k}^y \in [-\pi/2, \pi/2] \quad \Longrightarrow \quad \mathrm{Corr}_0^n \left[v_{T_k}^y, z_{T_k}^{y_1} \right] = \sin \left[\rho_{T_k}^y \right] \in [-1, 1] \, .$$

As $\Psi_{T_k}^{y_1}$ can take any value (consequently $\widetilde{\Psi}_{T_k}^{y_1}$), we do not restrict the domain of $\rho_{T_k}^y$ any further.

Increasing $\nu_{T_k}^y$ increases the curvature of the year-on-year smile, while $\rho_{T_k}^y < 0$ gives negative skew and $\rho_{T_k}^y > 0$ gives positive skew. This can be seen in Figure 7.1, which shows the year-on-year inflation volatility smile for different values of the spherical parameters $\left(\Sigma_{T_k}^y, \nu_{T_k}^y, \rho_{T_k}^y \right)$.[4] The spherical parameterisation shows, in an intuitive way, the maximum curvature and the maximum positive/negative skew that are obtainable in the three-factor reduction of the QGY model, for a given volatility level of the year-on-year inflation smile.

The maximum curvature is obtained when $\sin \nu_{T_k}^y = 1$. This reduces the driving process of Y_{T_k} to involve only quadratic terms:

$$r_{T_k}^y = \frac{\Sigma_{T_k}^y}{K_{T_k}^y} \left\{ \sin \left[\rho_{T_k}^y \right] z_{T_k}^{y_1} - \cos \left[\rho_{T_k}^y \right] z_{T_k}^{y_2} \right\} z_{T_k}^{y_1} .$$

Then there is no correlation between inflation rates and nominal interest rates, because $\phi_{T_k}^{y_1} = 0$ in (7.3). This is not a feasible property, so in practice we cap $\sin \nu_{T_k}^y$ at a value less then one. The maximum negative skew occurs when $\sin \rho_{T_k}^y = -1$. The driving process of Y_{T_k} is then only be driven by $x_{T_k}^{y_1}$:

$$r_{T_k}^y = \frac{\Sigma_{T_k}^y}{K_{T_k}^y} \left\{ \cos \left[\nu_{T_k}^y \right] - \sin \left[\nu_{T_k}^y \right] z_{T_k}^{y_1} \right\} z_{T_k}^{y_1} .$$

Then the year-on-year rate $Y_{T_k} - 1$ has an upper bound as can be seen in Figure 7.1. In the same way, the maximum positive skew occurs when $\sin \rho_{T_k}^y = 1$, and then the year-on-year rate has a lower bound.

[4] We used the calculation method in §6.4 to get prices of year-on-year inflation caplets. The relation between these prices and the year-on-year inflation volatilities is given in §2.3. Furthermore, we have used results from §6.2 to get the value for the drift parameters A_{T_k} of the year-on-year inflation ratio Y_{T_k}.

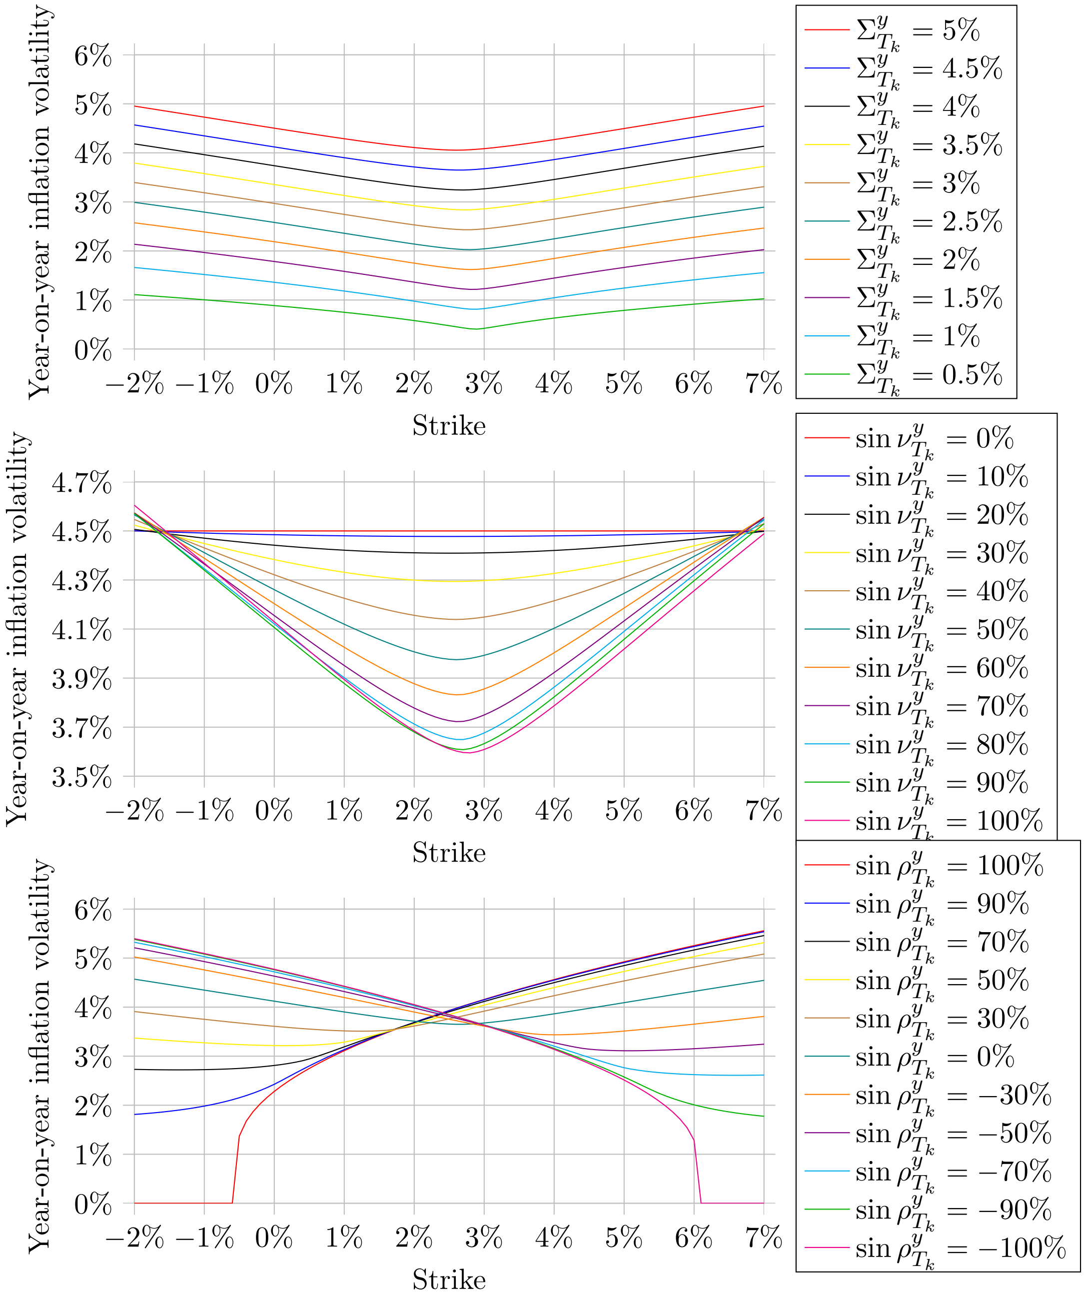

Figure 7.1: Year-on-year inflation volatility smiles for different values of the volatility level $\Sigma^y_{T_k}$ (top; $\sin\nu^y_{T_k} = 80\%$, $\sin\rho^y_{T_k} = 0\%$), the volatility-of-volatility $\sin\nu^y_{T_k}$ (middle; $\Sigma^y_{T_k} = 4.5\%$, $\sin\rho^y_{T_k} = 0\%$) and the volatility-correlation $\sin\rho^y_{T_k}$ (bottom; $\Sigma^y_{T_k} = 4.5\%$, $\sin\nu^y_{T_k} = 80\%$). Thus, the dark blue coloured curve at the top, the light blue coloured curve in the middle and the teal coloured curve at the bottom are the same. The maturity T_k is one year. The spherical parameters $\Sigma^y_{T_k} \geq 0$, $\nu^y_{T_k} \in [0, \pi/2]$ and $\rho^y_{T_k} \in [-\pi/2, \pi/2]$ are defined in §7.4 and in §7.5 we discuss their behaviour and boundary cases.

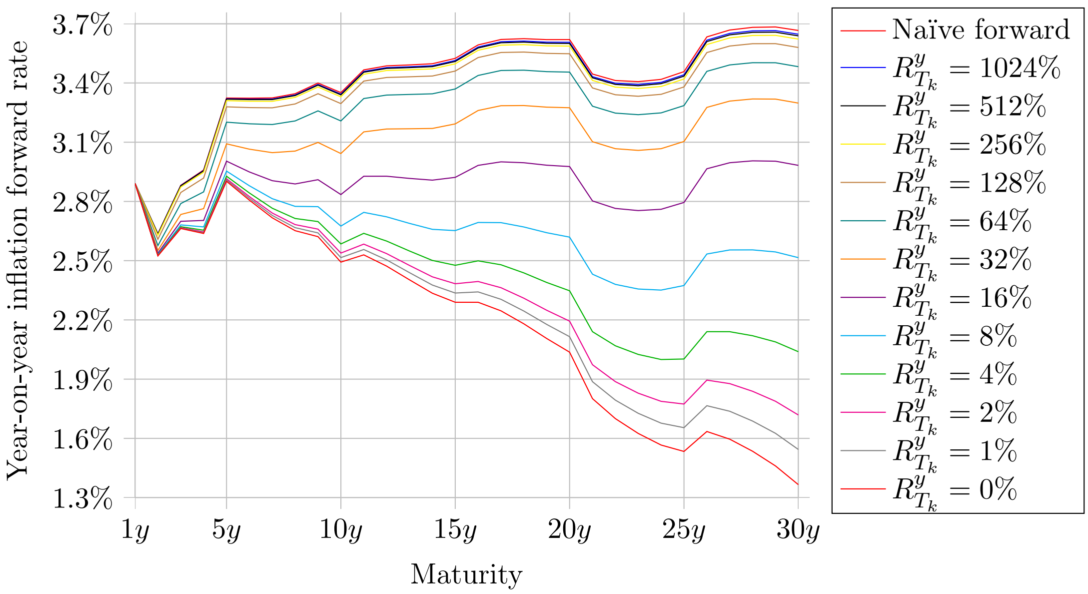

Figure 7.2: The graph shows the year-on-year inflation forward rate curves for different values of the growth rate $R^y_{T_k}$ defined in §7.6. We have assumed fixed values for the other parameters ($\Sigma^y_{T_k} = 4.5\%$, $\nu^y_{T_k} = 80\%$, $\rho^y_{T_k} = -50\%$ and $\rho^{ny_1}_t = -10\%$, for all $k = 1, 2, \ldots$ and $t \geq 0$).

7.6 Volatility growth rate and inflation autocorrelation

In the QGY model, we want to be able to control the year-on-year convexity correction inferred from year-on-year inflation forward rates (see §2.2). In the reduced three-factor version of the QGY model, we let the instantaneous volatility parameter for inflation, σ^y_t, be piecewise exponential:

$$\sigma^y_t = \exp\left[R^y_{T_k} t\right], \text{ for } t \in (T_{k-1}, T_k]\,. \tag{7.9}$$

Here $R^y_{T_k}$, $k = 1, 2, \ldots$, is a deterministic growth rate. The exponential form of σ^y_t appeared when we performed a global calibration of all the model parameters to market prices of year-on-year options. By changing the growth rate $R^y_{T_k}$, we can adjust the year-on-year inflation forward rates defined in §2.2, as can be seen in Figure 7.2. Each curve in the figure shows the year-on-year inflation forward rate for different values of the growth rate $R^y_{T_k}$. We have also plotted the naïve year-on-year inflation forward rate, i.e. $I_{0t}/\, I_{0,t-1} - 1$.

The autocorrelation of the year-on-year inflation ratio is given by $\mathrm{Corr}^n_0\left[\ln Y_t, \ln Y_T\right]$. In Figure 7.3, we plot this year-on-year autocorrelation implied by the QGY model for two different values of the instantaneous volatility growth rate, $R^y_{T_k} = 20\%, 60\%$, with $T = 1, 10, 20, 30$ years and $t = 1, 2, \ldots, 30$ years. A larger growth rate $R^y_{T_k}$ increases the decay of the year-on-year autocorrelation and for less autocorrelation we get less year-on-year convexity correction, which is consistent with the analysis in §2.2.[5]

7.7 Nominal-inflation correlation

We want to be able to adjust the payment delay convexity correction inferred from zero-coupon inflation swaps with payment delay. From the analysis in §2.1, we can control this convexity by changing the correlation between the nominal interest rate and the inflation index. In the reduced three-factor version of the QGY model, the parameter that affects this correlation is the instantaneous correlation parameter $\rho^{ny_1}_t$, which controls the correlation between the nominal driving factor x^n_t and one of the inflation driving factors $x^{y_1}_t$. To see this, we look at the following correlation between the logarithm of the inverse numeraire and the inflation index:

$$\mathrm{Corr}^n_0 \left[\ln D^n_T, \ln I_t\right],$$

for $t \leq T$, $t, T = 1, 2, \ldots, 30$ years. In Figure 7.4, we plot this correlation implied by the QGY model for two different values of the parameter $\rho^{ny_1}_t = -10\%, 20\%$. We see that a higher value of $\rho^{ny_1}_t$ coincides with an increase in the nominal-inflation correlation, as we noted in (7.3).

[5] To understand the reason for this better, we let $\Psi^{y_1}_{T_k} = \Psi^{y_1 y_2}_{T_k} = 0$. Assuming a constant $R^y_{T_k} = R > 0$, for all $k = 1, 2, \ldots$, is then the same as assuming a flat mean-reversion rate in the Hull-White equation in (4.1). We get

$$G_t = \int_0^t \exp\left[Rs\right] ds = \tfrac{1}{R} \left(\exp\left[Rt\right] - 1\right).$$

Assuming that $\phi^{y_1}_t, \phi^{y_1}_T < 0$ and $t \leq T$, we then get for large values of Rt and RT

$$\mathrm{Corr}\left[\ln Y_t, \ln Y_T\right] = \frac{\phi^{y_1}_t \phi^{y_1}_T G^{y_1}_t}{\sqrt{\left(\phi^{y_1}_t\right)^2 G^{y_1}_t \left(\phi^{y_1}_T\right)^2 G^{y_1}_T}} = \sqrt{\frac{G^{y_1}_t}{G^{y_1}_T}} = \sqrt{\frac{\exp\left[Rt\right] - 1}{\exp\left[RT\right] - 1}} \approx \exp\left[-\tfrac{1}{2} R \left(T - t\right)\right],$$

so the autocorrelation decay is exponential in $T - t$ and is more pronounced for higher values of R. This is the same observation as made by Pelsser (2000).

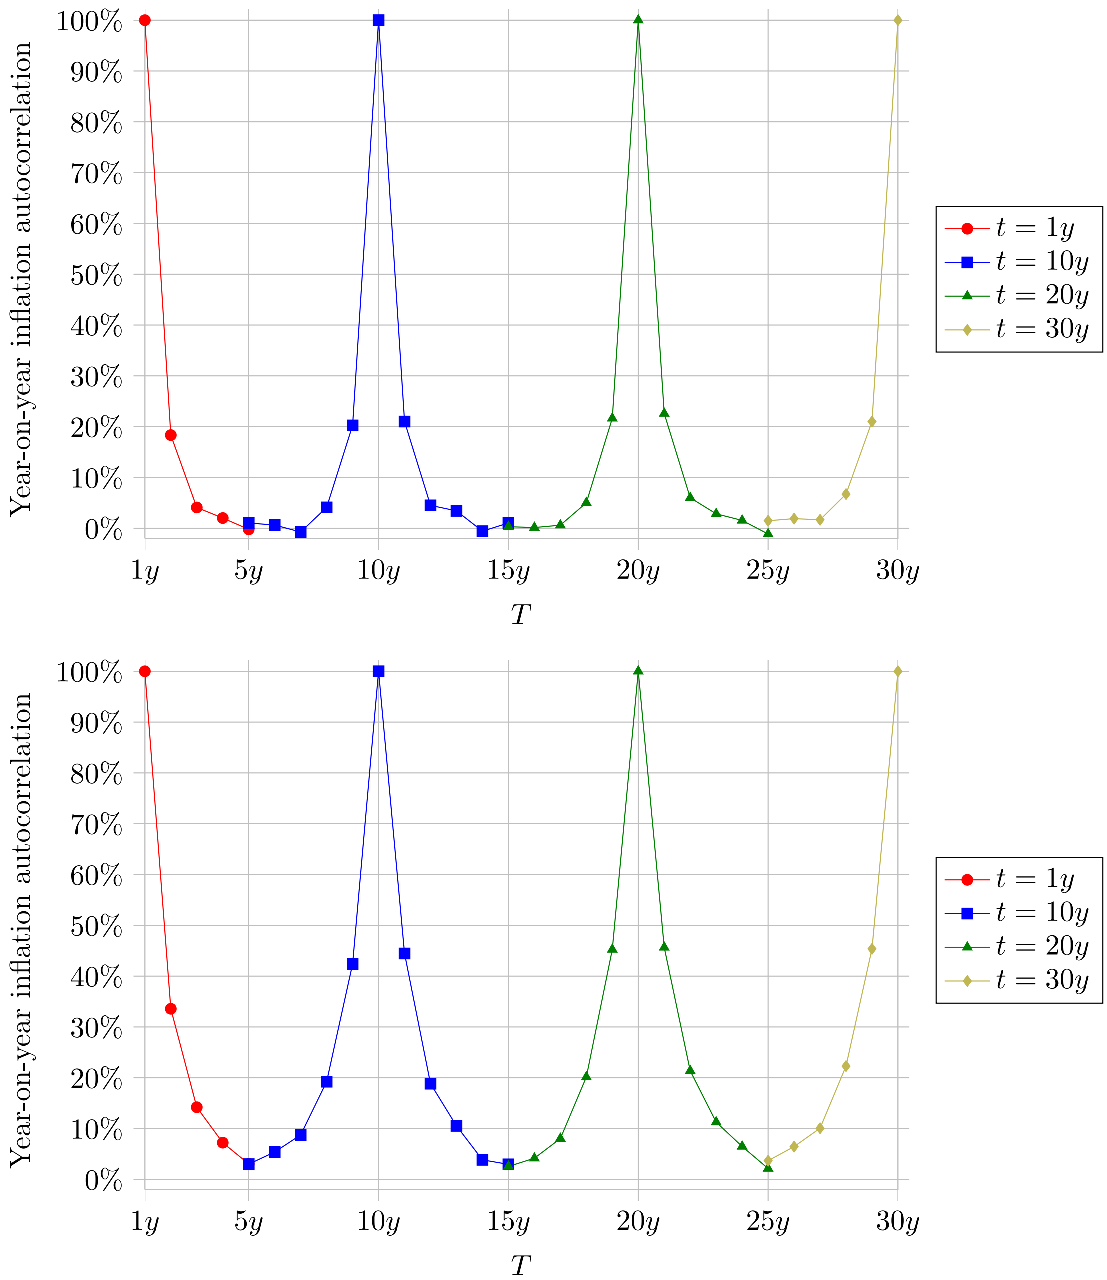

Figure 7.3: The graphs show the year-on-year inflation autocorrelation $\mathrm{Corr}_0^n\left[\ln\left[\frac{I_t}{I_{t-1}}\right], \ln\left[\frac{I_T}{I_{T-1}}\right]\right]$ for $t = 1, 10, 20, 30$ years. The autocorrelation is implied by the QGY model with a Monte Carlo simulation. We use different values for the instantaneous volatility growth rate for each graph, $R^y_{T_k} = 100\%$ (top) and $R^y_{T_k} = 50\%$ (bottom), defined in §7.6. We have assumed fixed values for the other parameters ($\Sigma^y_{T_k} = 4.5\%$, $\nu^y_{T_k} = 80\%$, $\rho^y_{T_k} = -50\%$ and $\rho^{ny_1}_t = -10\%$ for all $k = 1, 2, \ldots$ and $t \geq 0$).

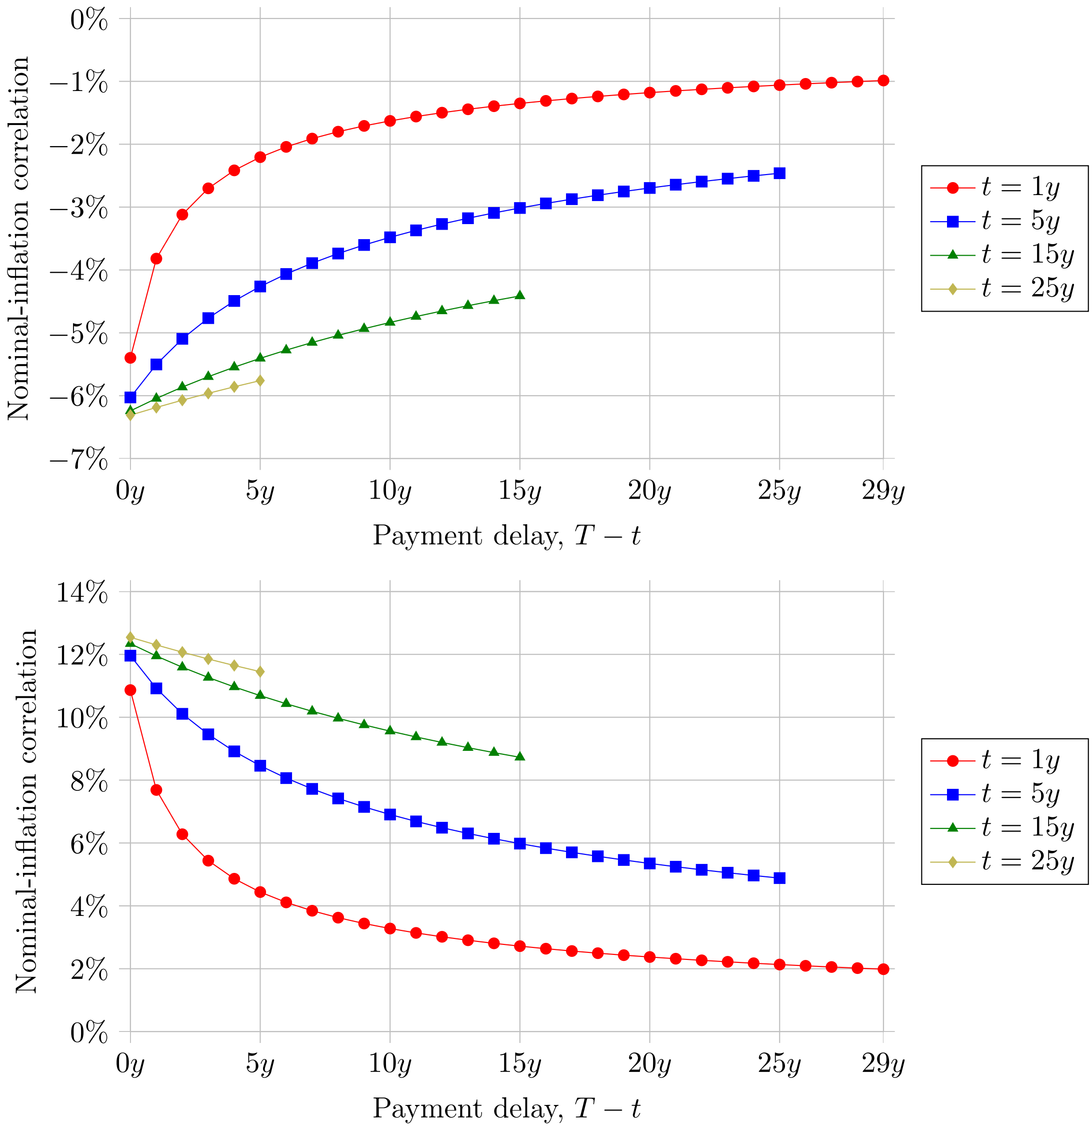

Figure 7.4: As discussed in §7.7, the graphs show the nominal-inflation correlation $\mathrm{Corr}_0^n\left[\ln D_T^n, \ln I_t\right]$ for $t = 1, 5, 10, 15, 20, 25$ years and $T = t, t+1, \ldots, 30$ years (the payment delay is defined as $T - t$). The correlation is implied by the QGY model with a Monte Carlo simulation. We use different values of the instantaneous nominal-inflation correlation for each graph, $\rho_t^{ny_1} = -10\%$ (top) and $\rho_t^{ny_1} = 20\%$ (bottom). We have assumed fixed values for the other parameters ($\Sigma_{T_k}^y = 4.5\%$, $\nu_{T_k}^y = 80\%$, $\rho_{T_k}^y = -50\%$ and $R_{T_k}^y = 100\%$ for all $k = 1, 2, \ldots$).

8 Calibration and pricing

In this chapter, we illustrate the calibration of the reduced three-factor QGY model given in Chapter 7. In §8.1, we describe the calibration method. We calibrate the model to year-on-year inflation volatilities as implied by market prices of year-on-year options on Tuesday 31 January 2012, for both the UK's Retail Prices Index (RPI) and the eurozone's Harmonised Index of Consumer Prices Excluding Tobacco (HICPxT). The results of the calibration are given in §8.2. The model fits well to the smile, with stable parameters. Using the model with the calibrated parameters, we apply Monte Carlo simulation methods to price zero-coupon inflation options on HICPxT in §8.3, and to price limited price indices (LPIs) on RPI in §8.4. Finally, in §8.5, we calibrate the model to market prices on other dates and we see that the parameters are stable for all calibration dates.

8.1 Setup

For the log-linear Gaussian inverse numeraire D_t^n, we calibrate the parameter $\phi_t^{n_1}$ to market prices of at-the-money nominal interest rate caplets, with annual maturities from 1-30 years. We are mainly interested in looking at the inflation volatility smile, so we will not discuss the properties or the results of this nominal calibration further.

Given calibrated values of $\phi_t^{n_1}$, we now turn to calibration of the inflation market. We use market prices published by Markit's Totem Service. For both RPI and HICPxT, we calibrate the three-factor QGY model from §7.1 to these market prices. The maturities are annual and we calibrate one maturity at a time, starting with the one-year maturity and going up to the thirty-year maturity. At each maturity T_k, $k = 1y, \dots, 30y$, the calibration instruments are the year-on-year inflation volatilities, defined in §2.3. The market volatilities are given in Figure 2.4. We have ten strikes at $i = -2\%, -1\%, \dots, 7\%$ for RPI and eight strikes at $i = -2\%, -1\%, \dots, 5\%$ for HICPxT. We let $V_{T_k}^{\text{MKT}}[i]$ be the market year-on-year inflation volatility at strike i, and we let $V_{T_k}^{\mathrm{QGY}}[i]$ be the corresponding year-on-year inflation volatility as given by the QGY model.[1] We calibrate the parameters of the QGY model so that we minimise the root-mean-square error

$$\sqrt{\frac{1}{N}\sum_i \left(V_{T_k}^{\mathrm{QGY}}[i] - V_{T_k}^{\mathrm{MKT}}[i]\right)^2}.$$

Using the spherical parameterisation in §7.4, we calibrate the volatility level $\Sigma_{T_k}^y$ of the year-on-year smile, the volatility-of-volatility $\sin\nu_{T_k}^y$ that controls the curvature of the smile, and the volatility-correlation $\sin\nu_{T_k}^y$ that controls the skew of the smile. We restrict the range of the parameters to $\nu_{T_k}^y \in [0, \pi/2]$ and $\rho_{T_k}^y \in [-\pi/2, \pi/2]$ as explained in §7.5. We also calibrate the volatility growth rate $R_{T_k}^y$, which is defined in §7.6 and controls the autocorrelation of the inflation index and in turn affects the year-on-year inflation forward rate.

To get some correlation between inflation rates and nominal inflation rates, we fix $\rho_t^{ny_1} = -10\%$ for RPI and $\rho_t^{ny_1} = 20\%$ for HICPxT for all values of $t \geq 0$. This correlation is consistent with the payment delay convexity corrections observed from market quotations of some inflation-linked options that have payment delays. For RPI this correction is positive, indicating a negative correlation, while for HICPxT the correction is negative, indicating a positive correlation.

The maximum value that $\sin\nu_{T_k}^y$ can take is 100%, which gives the maximum curvature of the year-on-year smile that the spherical parameterisation of the QGY model can obtain. Nevertheless, as we noted in §7.5, to get some correlation between inflation rates and nominal interest rates, we set the upper bound of $\sin\nu_{T_k}^y$ to 80%. This value is

[1] To calculate the volatility from prices of year-on-year caps, C_{0T}^k, and floors, F_{0T}^k, at given strikes k and maturities T, we first need to extract the year-on-year inflation forward rate y_{0T} using the put-call parity. If we let c_{0T}^k and f_{0T}^k be the forward prices of a year-on-year caplet and floorlet respectively, the put-call parity is $c_{0T}^k - f_{0T}^k = y_{0T} - k$. Discounting and summing from $t = 1, 2, \ldots, T$, we get

$$C_{0T}^k - F_{0T}^k = \sum_{t=1}^T P_{0T}^n c_{0T}^k - \sum_{t=1}^T P_{0T}^n f_{0T}^k = \sum_{t=1}^T P_{0T}^n y_{0T} - k\sum_{t=1}^T P_{0T}^n.$$

For two different strikes k and k', we then retrieve the year-on-year inflation forward rate and the nominal discount bond:

$$\begin{aligned} y_{0T} &= \left(k'\left(C_{0T}^k - F_{0T}^k - C_{0,T-1}^k + F_{0,T-1}^k\right) - k\left(C_{0T}^{k'} - F_{0T}^{k'} - C_{0,T-1}^{k'} + F_{0,T-1}^{k'}\right)\right) \Big/ \left(P_{0T}^n\left(k' - k\right)\right) \\ P_{0T}^n &= \left(\left(C_{0T}^k - F_{0T}^k - C_{0,T-1}^k + F_{0,T-1}^k\right) - \left(C_{0T}^{k'} - F_{0T}^{k'} - C_{0,T-1}^{k'} + F_{0,T-1}^{k'}\right)\right) \Big/ \left(k' - k\right). \end{aligned}$$

For RPI and HICPxT, we have floors with strikes $i = -2\%, \ldots, 2\%, 3\%$, and caps with strikes $i = 2\%, 3\%, \ldots, 7\%$ ($i = 6\%, 7\%$ only for RPI), so we set $k = 2\%$ and $k' = 3\%$.

high enough to get a curvature that is close to the maximum curvature that the reduced three-factor QGY model can obtain, i.e. when $\sin \nu^y_{T_k} = 100\%$.

From $R^y_{T_k}$ we retrieve $G^{y_1}_t$ and $G^{y_2}_t$ using equations in §7.6 and §7.1. From §7.4 and §7.1 we retrieve $\left(\phi^{y_1}_{T_k}, \Psi^{y_1}_{T_k}, \Psi^{y_1 y_2}_{T_k}\right)$ from $\left(\Sigma^y_{T_k}, \nu^y_{T_k}, \rho^y_{T_k}\right)$. Then we use the pricing formulas in §6.4 for the year-on-year inflation forward rates and options. Furthermore, we have used results in §6.2 to get the value for the drift parameters A_{T_k} of the year-on-year inflation ratio Y_{T_k}. All these calculations are programmed in the C++ pricing library at Lloyds Bank. The spherical parameters are changed in the Excel spreadsheet application, which provides a solver to minimise the calibration error by changing the parameters while also matching the year-on-year inflation forward rate.

Forward inflation index

The year-on-year convexity correction is the difference between the adjusted year-on-year inflation forward rate and the unadjusted year-on-year inflation forward rate (see §2.2).[2] In Figure 8.1, we plot this correction for RPI and HICPxT and we notice that it is quite unstable with varying maturity. To fit the model to this correction, we would have to change the autocorrelation of the inflation index and the correlation between inflation rates and nominal interest rates by a large amount between maturities. We fix the nominal-inflation correlation at a constant value so in our calibration we can only change the year-on-year convexity correction by changing the volatility growth rate $R^y_{T_k}$, which changes the inflation autocorrelation. We therefore replace the actual forward inflation index I_{0T} with an approximate forward inflation index $\widetilde{I}_{0T}$ ($\widetilde{I}_{00} = I_{0T}$), which is defined through the following function:

$$\frac{\widetilde{I}_{0T}}{\widetilde{I}_{0,T-1}} - 1 = y_{0T} + a \ln [T] \, .$$

We use a logarithmic function to approximate the difference between the year-on-year inflation forward rate, y_{0T}, and the year-on-year rate of the forward inflation index,

[2] To extract the forward inflation index I_{0T} from prices of zero-coupon inflation caps, C_{0T}, and floors, F_{0T}, we use the put-call parity $C^k_T - F^k_T = P^n_{0T} y_T - P^n_{0T} k$. For two different strikes k and k', we then get:

$$\begin{aligned} I_{0T} &= \left(k' \left(C^k_{0T} - F^k_{0T}\right) - k \left(C^{k'}_{0T} - F^{k'}_{0T}\right)\right) \Big/ \left(P^n_{0T} \left(k' - k\right)\right) \\ P^n_{0T} &= \left(\left(C^k_{0T} - F^k_{0T}\right) - \left(C^{k'}_{0T} - F^{k'}_{0T}\right)\right) \Big/ \left(k' - k\right) . \end{aligned}$$

For RPI and HICPxT we have floors with strikes $i = -2\%, \ldots, 2\%, 3\%$, and caps with strikes $i = 2\%, 3\%, \ldots, 7\%$ ($i = 6\%, 7\%$ only for RPI), so we set $k = 2\%$ and $k' = 3\%$.

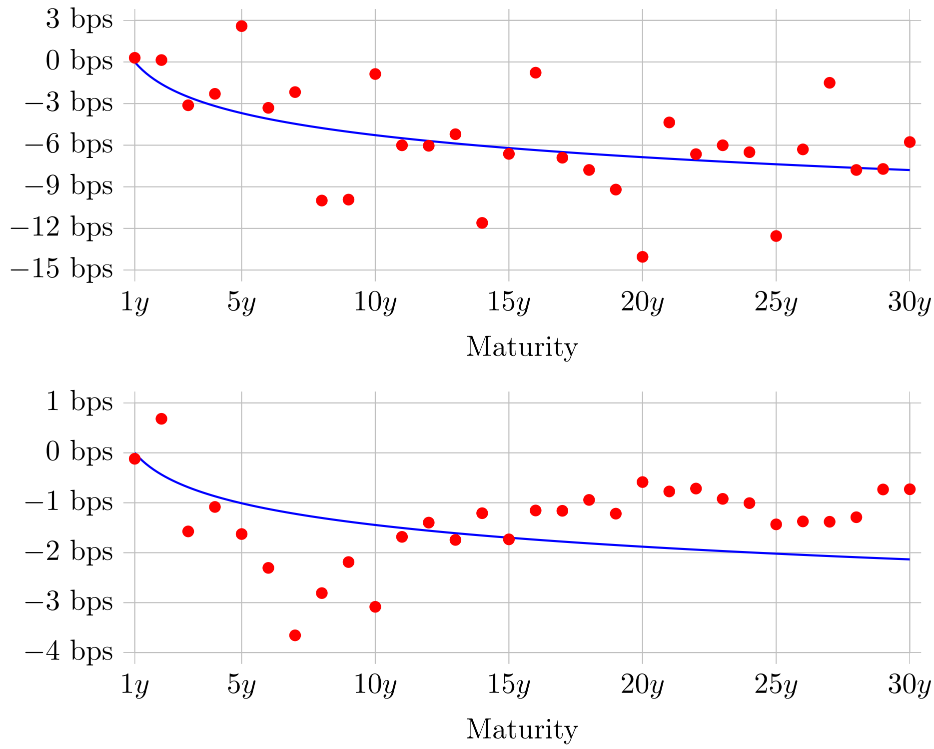

Figure 8.1: The graphs show the year-on-year convexity correction for RPI (top) and HICPxT (bottom). The red markers show the correction $y_{0T} - (I_{0T}/\, I_{0,T-1} - 1)$ inferred from market prices of zero-coupon and year-on-year inflation options that were provided by Markit on 31 January 2012. As described in §8.1, the blue curves are logarithmic curves, $-a \ln [T]$, which best fit the forward inflation index curve I_{0T}. For RPI, we have $a = 2.2913$ bps and for HICPxT, we have $a = 0.6273$ bps.

$I_{0T}/\, I_{0,T-1} - 1$. The number a is chosen so that we minimise the root-mean-square error

$$\sqrt{\frac{1}{30}\sum_{T=1}^{30}\left(\widetilde{I}_{0T} - I_{0T}\right)^2}.$$

For RPI we have $a = 2.2913$ bps, and for HICPxT we have $a = 0.6273$ bps, giving a root mean square error of 7.0903 bps and 7.2924 bps respectively. The logarithmic function describes better the shape of the year-on-year convexity correction for RPI than for HICPxT. The correction for HICPxT suggests that the nominal-inflation correlation is increasing with maturity.

8.2 Results

We now look at the result of the calibration to year-on-year inflation options and forward rates on RPI and HICPxT. The values of the calibration parameters are given in Table 8.1. The calibration is very fast as we only need to calibrate one maturity at a time. Good properties on the part of the spherical parameterisation also speed up the calibration, with all parameters being stable with varying maturity. The calibration gives us stable values for the volatility level $\Sigma^y_{T_k}$ for all maturities and lower values for HICPxT than RPI, both of which is consistent with the properties of the market. The volatility-correlation $\sin\rho^y_{T_k}$ is stable for HICPxT with values around -20%. For RPI, $\sin\rho^y_{T_k}$ gradually becomes more negative with higher maturities, which captures the properties of the market, as the skew of the RPI smile becomes more negative with higher maturities.

For maturities of less than five years, the QGY model fits the option prices well. For higher maturities, the curvature of the year-on-year smile becomes more pronounced, resulting in the volatility-of-volatility $\sin\nu^y_{T_k}$ hitting the upper bound of 80%. From experiment, we note that, even if $\sin\nu^y_{T_k}$ were allowed to go up to its maximum value of 100% (see §7.5), this would not result in a big increase in the curvature or a big decrease in the calibration error. In Figures 8.2 and 8.4, we show the RPI and HICPxT year-on-year inflation volatility surfaces given by the calibrated QGY model, as well as the difference between these volatilities and the market year-on-year inflation volatilities as defined in §2.3 (see Figure 2.4 for the corresponding market surfaces). We note that the calibration gives a better fit to the HICPxT market, where the curvature of the smile is less than for RPI.

Figures 8.3 and 8.5 show the prices of the most common options that are traded in the market. For RPI, we show the prices of year-on-year 5% cap, 3% straddle and 0% floor. For HICPxT, we show the prices of year-on-year 4% cap, 2% straddle and 0% floor. In these figures, we show the prices as given by the calibrated QGY model, as well as showing the market prices and their approximate bid and offer spreads. We can see that the largest pricing error is at the 0% strike, where the calibrated QGY model underprices the market prices, and the error increases with maturity. Consequently, to minimise the mean error, the calibrated QGY model overprices the other strikes.

	RPI					HICPxT				
T_k	$\Sigma^y_{T_k}$	$\sin\nu^y_{T_k}$	$\sin\rho^y_{T_k}$	$R^y_{T_k}$	RMS	$\Sigma^y_{T_k}$	$\sin\nu^y_{T_k}$	$\sin\rho^y_{T_k}$	$R^y_{T_k}$	RMS
$1y$	1.80	64.9	−11.7	0	0.029	1.49	59.5	−5.76	0	0.029
$2y$	2.72	80	−15.4	59.2	0.072	2.48	80	−11.95	101.6	0.067
$3y$	3.39	80	−21.3	76.7	0.130	3.22	80	−13.29	109.2	0.121
$4y$	3.84	80	−25.5	89.5	0.185	3.66	80	−12.77	118.5	0.139
$5y$	4.16	80	−32.9	97.8	0.242	3.97	80	−13.39	120.4	0.183
$6y$	4.43	80	−36.2	104.4	0.240	4.19	80	−14.97	121.5	0.201
$7y$	4.57	80	−39.5	109.2	0.250	4.27	80	−17.04	120.6	0.207
$8y$	4.48	80	−41.1	112.2	0.271	4.21	80	−18.61	119.1	0.210
$9y$	4.64	80	−47.0	114.7	0.291	4.13	80	−18.63	116.7	0.215
$10y$	4.66	80	−49.7	116.7	0.310	4.08	80	−20.64	114.3	0.241
$11y$	4.63	80	−53.1	117.9	0.358	4.08	80	−23.42	111.6	0.248
$12y$	4.61	80	−55.2	119.1	0.380	4.01	80	−20.90	108.9	0.259
$13y$	4.61	80	−57.7	119.7	0.399	3.95	80	−19.03	106.3	0.246
$14y$	4.60	80	−59.6	120.5	0.408	3.88	80	−18.02	103.6	0.237
$15y$	4.58	80	−61.9	121.0	0.421	3.86	80	−15.70	101.1	0.242
$16y$	4.58	80	−64.7	121.0	0.449	3.87	80	−14.62	99.5	0.240
$17y$	4.54	80	−65.4	121.4	0.454	3.89	80	−15.00	97.1	0.240
$18y$	4.51	80	−65.7	121.5	0.459	3.92	80	−15.76	95.4	0.239
$19y$	4.43	80	−65.0	121.7	0.460	3.96	80	−16.02	93.5	0.239
$20y$	4.38	80	−64.9	121.8	0.458	4.01	80	−17.12	91.9	0.245
$21y$	4.46	80	−62.6	121.6	0.483	4.05	80	−17.06	90.8	0.243
$22y$	4.47	80	−62.2	121.8	0.494	4.05	80	−15.20	89.1	0.236
$23y$	4.47	80	−62.6	121.9	0.503	4.05	80	−13.81	88.0	0.231
$24y$	4.46	80	−63.2	122.1	0.520	4.05	80	−12.91	86.5	0.232
$25y$	4.45	80	−63.6	122.3	0.535	4.06	80	−11.90	85.3	0.233
$26y$	4.60	80	−68.7	122.6	0.582	4.10	80	−12.20	84.1	0.240
$27y$	4.56	80	−69.0	123.0	0.589	4.14	80	−13.09	82.9	0.241
$28y$	4.56	80	−69.6	123.4	0.598	4.14	80	−15.28	81.7	0.255
$29y$	4.57	80	−70.5	123.8	0.607	4.15	80	−16.69	80.5	0.273
$30y$	4.60	80	−74.0	124.2	0.604	4.16	80	−18.31	79.5	0.289

Table 8.1: Expressed in percentages (%), the table shows the values of the calibration parameters of the QGY model and the root-mean-square calibration error of the year-on-year inflation volatilities. The parameters have been calibrated to market year-on-year inflation volatilities in RPI and HICPxT on 31 January 2012, by means of the calibration method described in §8.1. The calibration results are discussed in §8.2. We fix $\rho^{ny_1}_t = -10\%$ for RPI and $\rho^{ny_1}_t = 20\%$ for HICPxT for all $t \geq 0$. See figures 8.2-8.5 for pricing results.

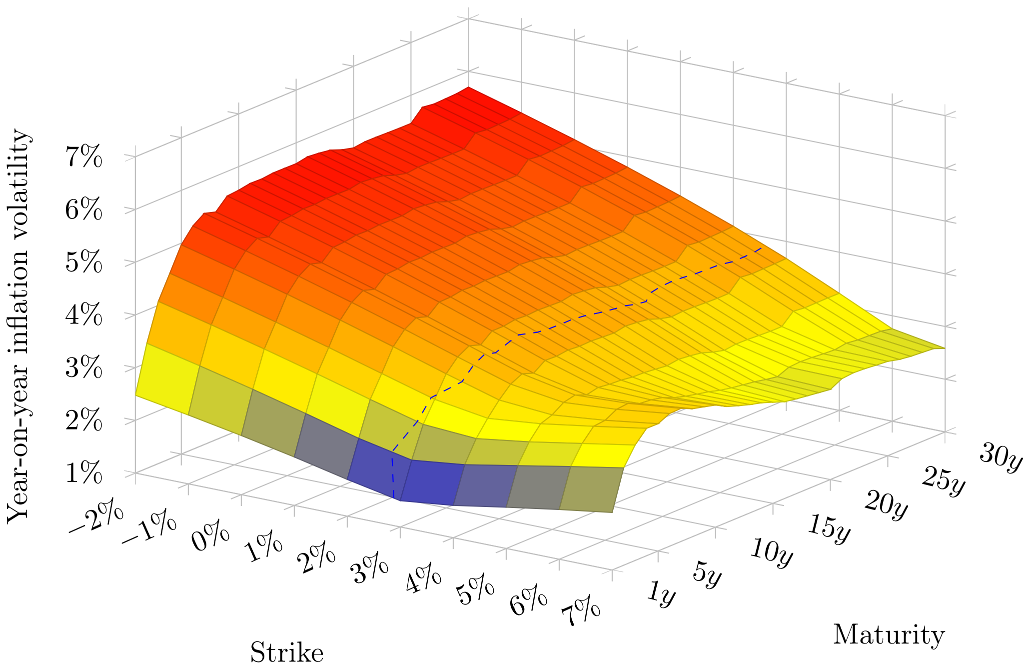

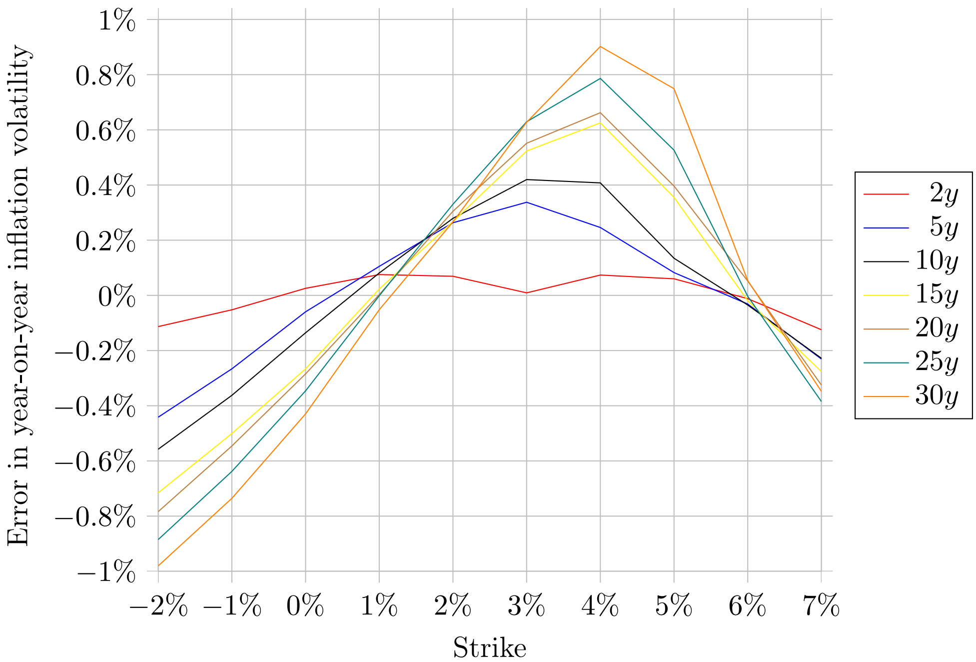

Figure 8.2: Results from the calibration of the QGY model to RPI year-on-year inflation volatilities (defined in §2.3) given on 31 January 2012, as described in §8.1. The top graph shows the year-on-year inflation volatility surface given by the QGY model, where the blue and broken curve represents the volatility at strikes that are equal to the year-on-year inflation forward rate. Each curve in the bottom graph shows, at a given maturity, the difference in the year-on-year inflation volatilities between the QGY model and the market. The calibration results are discussed in §8.2.

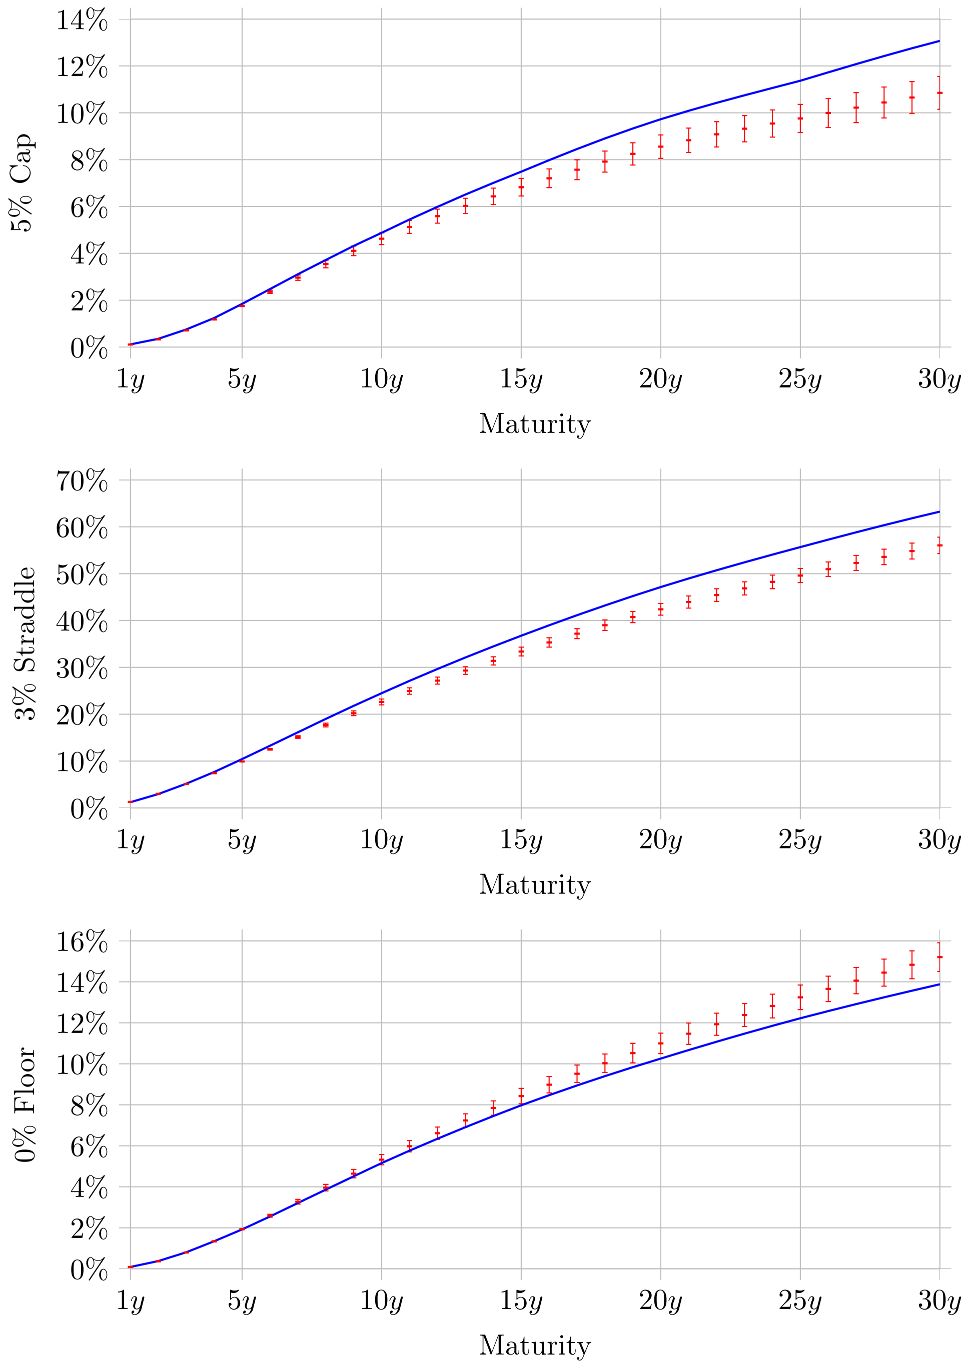

Figure 8.3: Prices of RPI year-on-year 5% cap, 3% straddle and 0% floor. The blue curves represent the result of the QGY model calibration described in §8.1. The red markers represent average market prices (the middle bar) given by Markit on 31 January 2012, with the lower and upper bars indicating the bid-offer prices. The calibration results are discussed in §8.2.

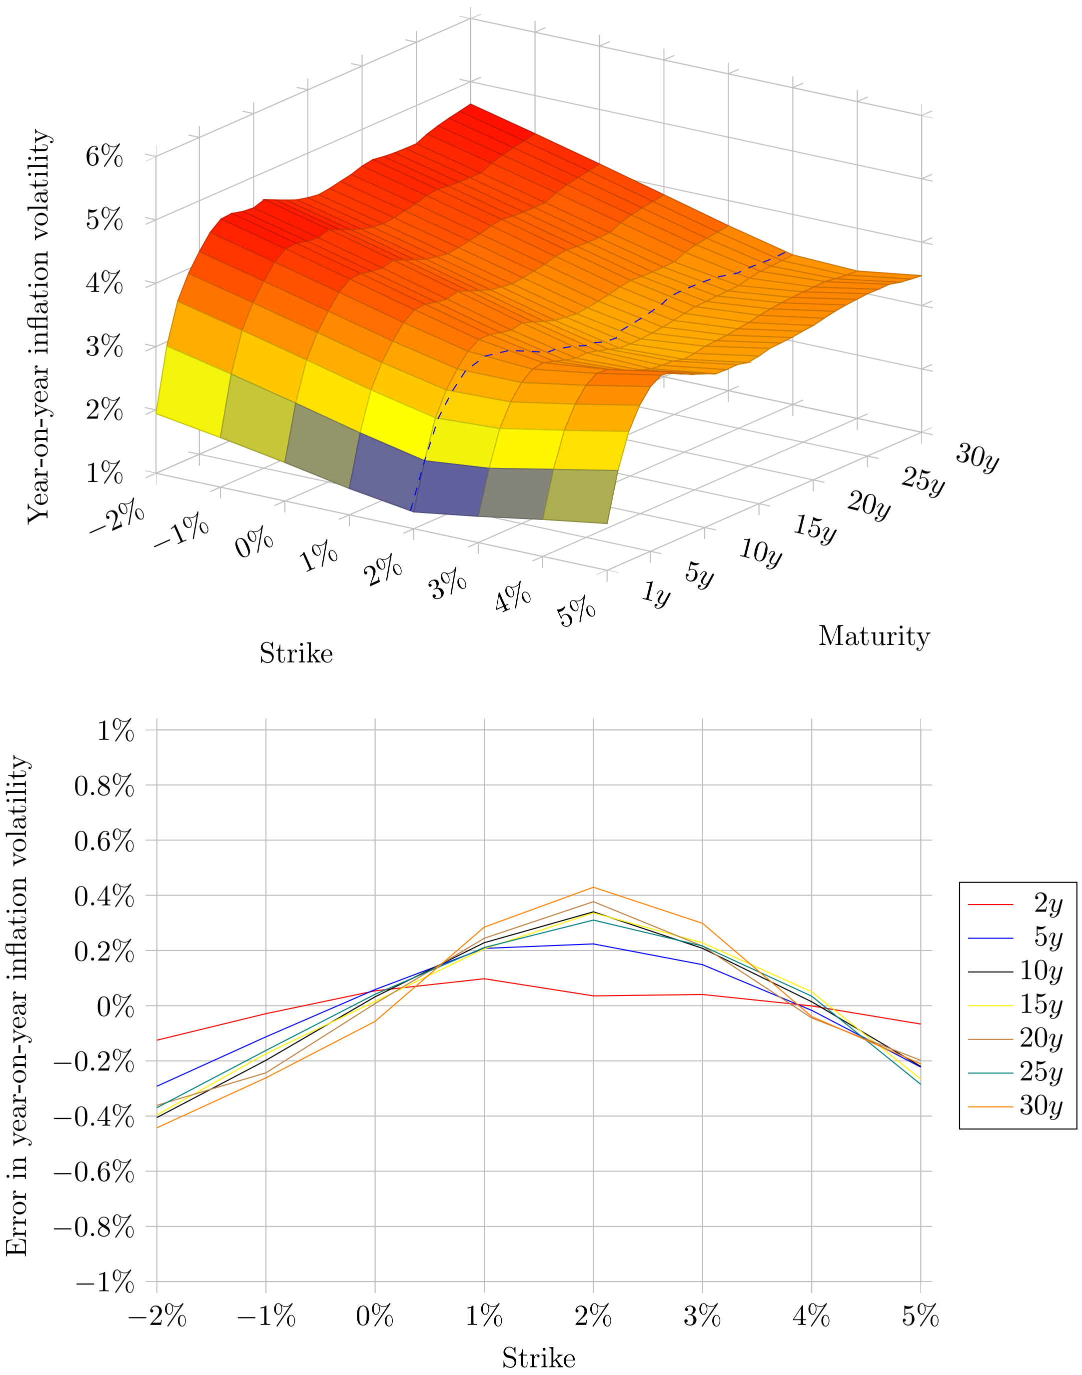

Figure 8.4: Results from the calibration of the QGY model to HICPxT year-on-year inflation volatilities (defined in §2.3) given on 31 January 2012, as described in §8.1. The top graph shows the year-on-year inflation volatility surface given by the QGY model, where the blue and broken curve represents the volatility at strikes that are equal to the year-on-year inflation forward rate. Each curve in the bottom graph shows, at a given maturity, the difference in the year-on-year inflation volatilities between the QGY model and the market. The calibration results are discussed in §8.2.

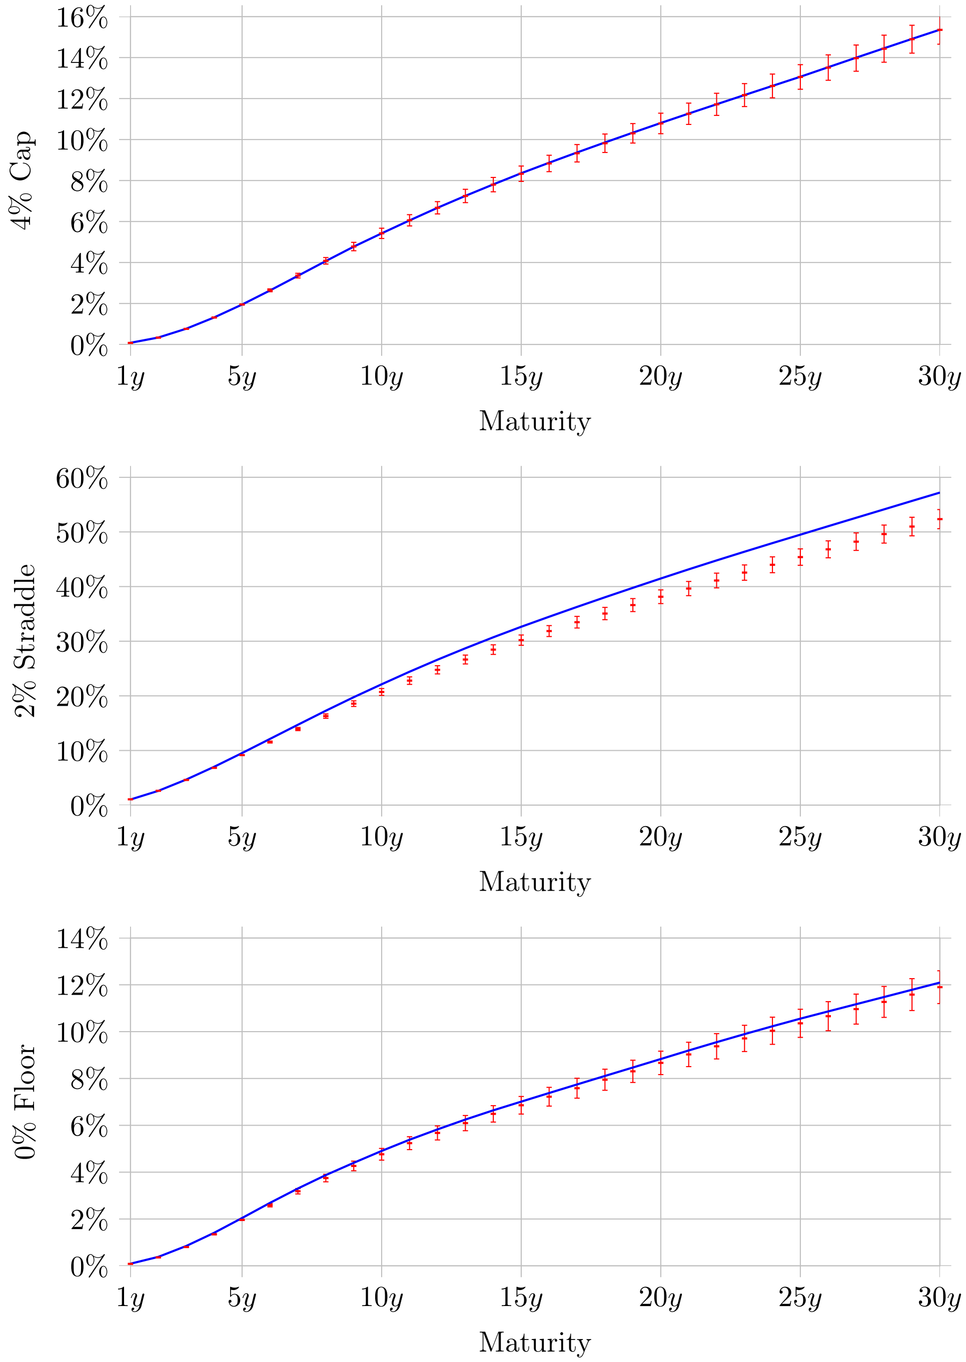

Figure 8.5: Prices of HICPxT year-on-year 4% cap, 2% straddle and 0% floor. The blue curves represent the result of the QGY model calibration described in §8.1. The red markers are average market prices (the middle bar) given by Markit on 31 January 2012, with the lower and upper bars indicating the bid-offer prices. The calibration results are discussed in §8.2.

8.3 Zero-coupon pricing

We have calibrated the QGY model to year-on-year options. We now use this calibration to price zero-coupon inflation floors and caps, which in the QGY model are path-dependent on the inflation index at discrete times T_i, $i = 1, 2, \ldots, k$:

$$\left(I_{T_k} - (1+K)^{T_k}\right)^{\pm} = \left(\prod_{i=1}^{k} Y_{T_i} - (1+K)^{T_k}\right)^{\pm}.$$

The year-on-year inflation ratio Y_{T_i} is given in (5.3). We use Monte Carlo methods to simulate the Gaussian process x_{T_i}, which is used to calculate realisations of Y_{T_i} and $D^n_{T_k}$.

In Figure 8.6, we show the zero-coupon inflation volatility surface on HICPxT as priced in the calibrated QGY model, as well as the differences in volatilities between the QGY model and the market. We define the zero-coupon inflation volatility in §2.3 and show corresponding market surface in Figure 2.3. We see that the pricing errors are larger for higher strikes, suggesting a difference between the skew of the year-on-year market and of the zero-coupon inflation market. We also see that the zero-coupon volatilities in the QGY model decrease with maturity, rather than being flat like the market volatilities. The zero-coupon inflation volatilities are controlled by the year-on-year inflation volatilities, as well as by the autocorrelation of the year-on-year inflation rate as analysed in §7.6. If we kept the year-on-year inflation volatilities fixed, in order to increase the zero-coupon option prices we would have to increase the autocorrelation,[3] which in the QGY model would lower the year-on-year inflation forward rates below the corresponding rates given by the market (see analysis in §7.6). This is a behaviour similar to that mentioned by Tan (2012). He notes that fitting both the year-on-year and zero-coupon volatility levels with a particular one-factor inflation model gives unstable mean-reversion speed—equating to unstable inflation autocorrelation. He also notes that this would result in the year-on-year inflation forward rates' being too low.

[3]The inflation index and the year-on-year inflation ratios are related through the expression $\ln I_T/I_0 = \sum_{t=1}^{T} Y_t$, with $Y_t = I_t/I_{t-1}$. The level of the zero-coupon inflation volatilities at a given maturity T is related to the square root of the following variance:

$$\mathrm{Var}_0^T[\ln I_T] = \sum_{s=1}^{T}\sum_{t=1}^{T} \mathrm{Covar}_0^T[\ln Y_s, Y_t] = \sum_{s=1}^{T}\sum_{t=1}^{T} \mathrm{Corr}_0^T[\ln Y_s, \ln Y_t]\sqrt{\mathrm{Var}_0^T[\ln Y_s]\,\mathrm{Var}_0^T[\ln Y_t]}.$$

The square root of $\mathrm{Var}_0^T[\ln Y_t]$ is related to the year-on-year inflation volatility level at time t (subtly, it incorporates some 'payment' delay $T-t$). The terms $\mathrm{Corr}_0^T[\ln Y_s, \ln Y_t]$ represent the autocorrelation of the inflation index. This autocorrelation links the volatility level of zero-coupon and year-on-year inflation volatilities.

To summarise, we have seen that calibrating the QGY model to the year-on-year inflation volatility surface and the year-on-year inflation forward rates makes the zero-coupon inflation volatilities too low and the zero-coupon inflation skew too negative. We could calibrate the QGY model to the year-on-year inflation volatility surface and the at-the-money zero-coupon inflation volatilities. This would make the year-on-year inflation forward too low, as well as being time-consuming because it requires simulation methods to price the zero-coupon inflation options. Going one step further and jointly calibrating the year-on-year and zero-coupon inflation volatility surfaces to try to even out the difference in their skew would give similar results.

8.4 LPI pricing

It is interesting to see how the calibrated QGY model prices LPIs (limited price indices) described in §2.2. We denote with $P^{f,c}_{T_k}$ the value of the LPI at time T_k, floored at f and capped at c. Assuming that we have an annual tenor structure $\{T_k\}_{k\in\mathbb{N}}$, $P^{f,c}_{T_k}$ is path-dependent on the inflation index at discrete times T_i, $i = 1, 2, \ldots, k$. At calibration time $T_0 = 0$ we set $P^{f,c}_{T_0} := I_{T_0}$ and iteratively define

$$P^{f,c}_{T_k} = P^{f,c}_{T_{k-1}}\left(1 + \max\left[f, \min\left[c, \frac{I_{T_k}}{I_{T_{k-1}}} - 1\right]\right]\right).$$

Denoting the value of $P^{f,c}_{T_k}$ at time T_0 as $P^{f,c}_{T_0T_k}$, the market quotes the LPI swap rate $s^{f,c}_{T_0T_k}$. It is the excess annual rate of return of the LPI over the forward inflation index:

$$s^{f,c}_{T_0T_k} := \left(\frac{P^{f,c}_{T_0T_k}}{P^{f,c}_{T_0}}\right)^{(T_k-T_0)^{-1}} - \left(\frac{I_{T_0T_k}}{I_{T_0}}\right)^{(T_k-T_0)^{-1}}.$$

In Figure 8.7, we show the LPI swap rates in the RPI market as given by Markit on 31 January 2012 and as priced in the calibrated QGY model. The three LPIs are floored at 0% and have 3% cap, 5% cap and no cap respectively. The calibrated QGY model underprices the year-on-year floor with 0% strike (see Figure 8.2), causing the LPI swap rates to be underpriced as they are floored at 0%. In addition the caps with 3% and 5% strikes are overpriced, which further reduces the price of the LPI swap rates capped at 3% and 5%. We also price the LPIs in the log-linear Gaussian year-on-year model, which is identical to the QGY model except we set $\Psi^y_{T_k} = 0$ and we therefore get no smile at any maturity. In that log-linear model, we have calibrated $\Sigma^y_{T_k}$ and $R^y_{T_k}$ to year-on-year inflation straddles with 3% strikes and to year-on-year inflation forward rates, while

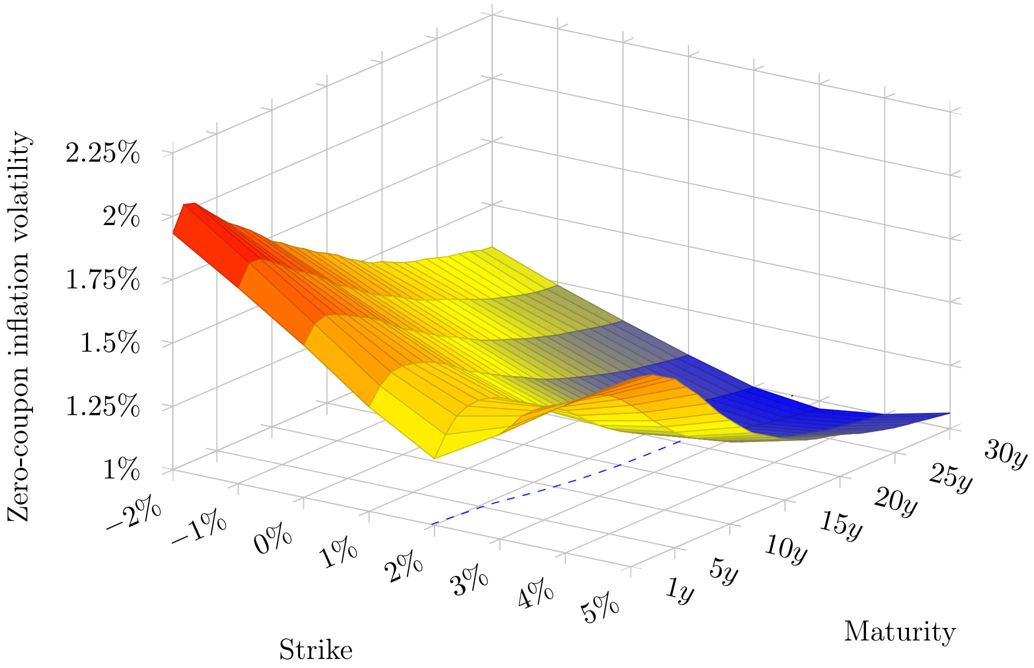

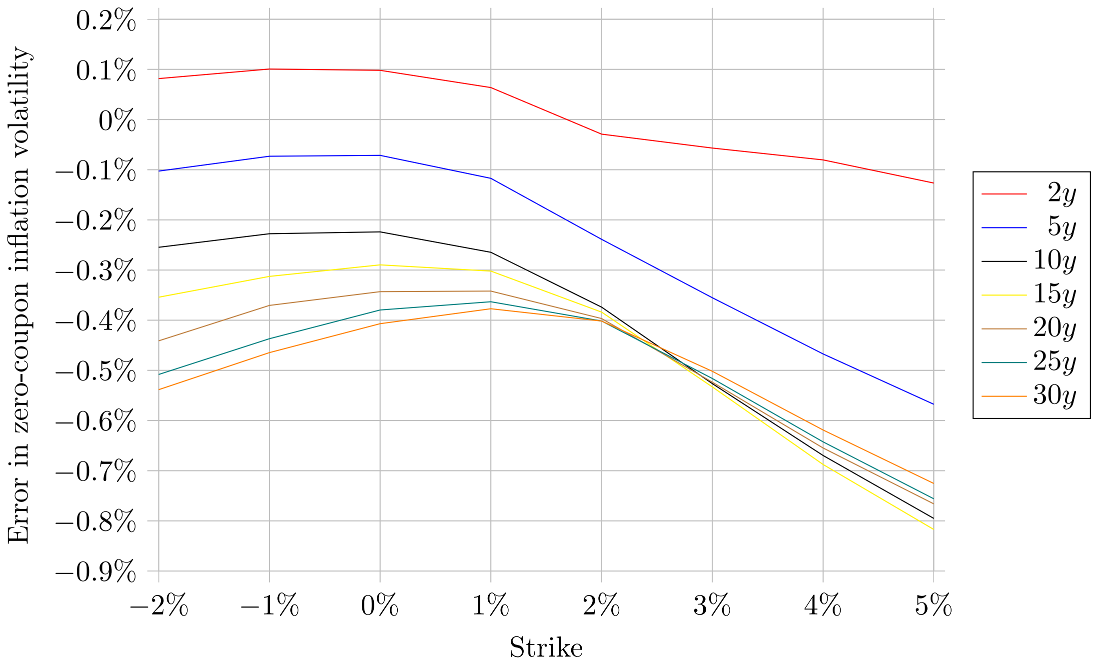

Figure 8.6: The top graph shows the zero-coupon inflation volatility surface (defined in §2.3) given by the QGY model, where the blue and broken curve represents the strikes that are equal to the annual rate of the forward inflation index. Each curve in the bottom graph shows, at a given maturity, the difference in the zero-coupon inflation volatilities between the QGY model and the market quotations given by Markit on 31 January 2012. The pricing results are discussed in §8.3. We have used the QGY model that has been calibrated to the HICPxT year-on-year inflation market on 31 January 2012, as explained in §8.1 and §8.2.

setting $\nu^y_{T_k} = \rho^y_{T_k} = 0$. In Figure 8.7 we see that the log-linear Gaussian year-on-year model underperforms the QGY model for all the LPI swap rates, because the QGY model incorporates the smile, which affects the LPI prices.

8.5 Calibration on different dates

Markit publishes market prices at the end of each month. To examine the stability of the calibrated QGY spherical parameters, we use this data for five months and perform a year-on-year calibration of the QGY model to RPI. The calibration is described in §8.1. The pricing dates are 31 January, 30 March, 31 May, 31 July and 28 September 2012.

In Figures 8.8 and 8.9, we show the calibration results at these calibration dates. The parameters are the volatility level $\Sigma^y_{T_k}$, the volatility-correlation $\sin\rho^y_{T_k}$, the growth rate $R^y_{T_k}$ and the root-mean-square calibration error (RMS). We see that the parameters are stable for different calibration dates. We do not plot the volatility-of-volatility parameter $\sin\nu^y_{T_k}$. For the one-year maturity, the value of $\sin\nu^y_{T_k}$ is around 60%-70%. For all other maturities, $\sin\nu^y_{T_k}$ always takes the maximum value of 80% to produce the maximum curvature of the year-on-year inflation smile.

We have not looked at the sensitivities of the option prices to movements in the values of the parameters—i.e. the Greeks. From the stability of $\Sigma^y_{T_k}$, $\sin\rho^y_{T_k}$, and $R^y_{T_k}$ between different calibration dates, we expect the Greeks for these parameters to be stable in time. The parameters could therefore be very suitable for hedging and risk-managing inflation-linked options.

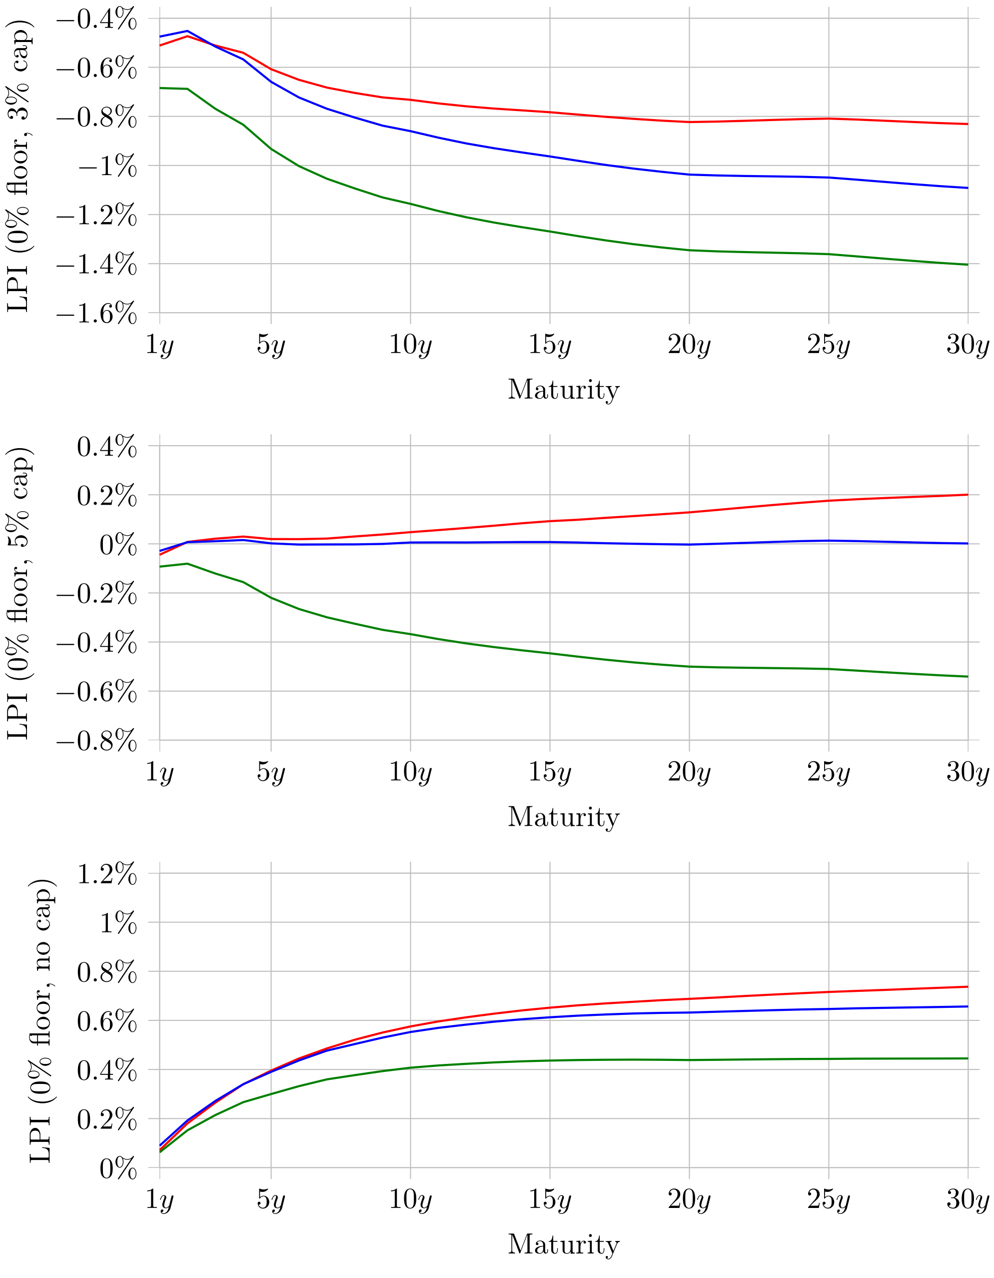

Figure 8.7: The graphs show the prices of three RPI-LPI swap rates, floored at 0% and with 3% cap (top), 5% cap (middle) and no cap (bottom). The red curves are market prices given by Markit on 31 January 2012. The blue curves are the results of the calibration described in §8.1, with parameter values given in Table 8.1. The green curves are the results of the calibration described in §8.1 but with the stochastic-volatility parameters $\nu^y_{T_k}$ and $\rho^y_{T_k}$ equal to zero (see the log-linear Gaussian year-on-year model described in §4.3). The pricing results are discussed in §8.4.

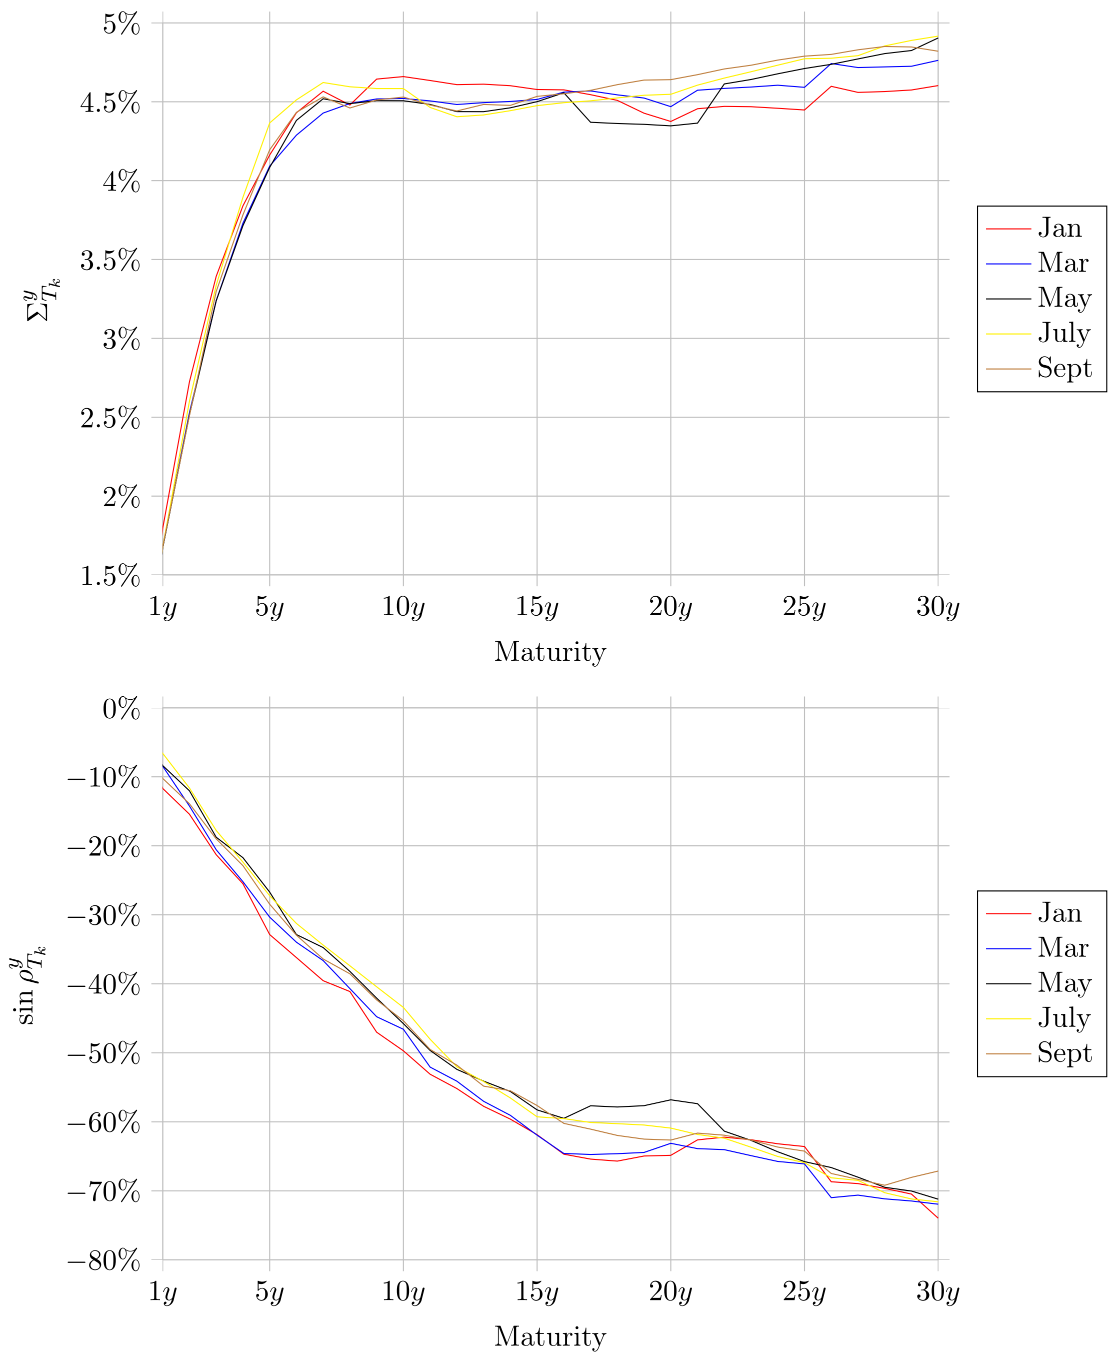

Figure 8.8: The graphs show the values of the volatility level $\Sigma^y_{T_k}$ (top) and the volatility-correlation $\sin\rho^y_{T_k}$ (bottom) for the QGY model that has been calibrated to the prices of year-on-year inflation options and forward rates on RPI at the end of five different months in 2012. The results are discussed in §8.5.

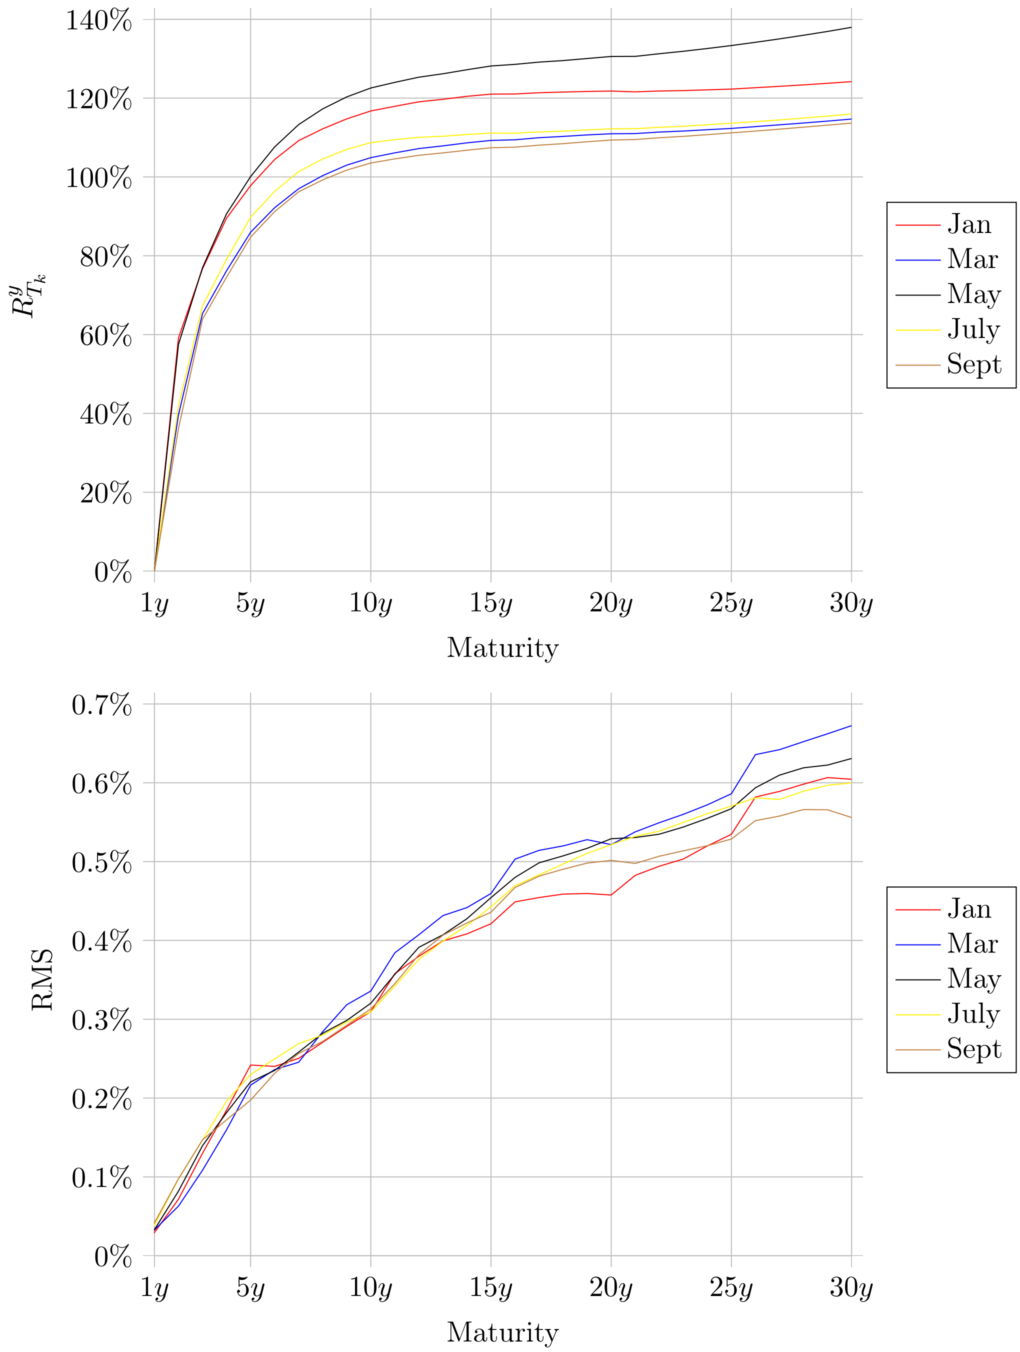

Figure 8.9: The graphs show the values of the growth rate $R^y_{T_k}$ (top) and the root-mean-square calibration errors (RMS; bottom) for the QGY model that has been calibrated to the prices of year-on-year inflation options and forward rates on RPI at the end of five different months in 2012. The results are discussed in §8.5.

9 Parameter interpolation

The value of most inflation indices, e.g. UK's Retail Prices Index (RPI) and the eurozone's Harmonised Index of Consumer Prices Excluding Tobacco (HICPxT), is calculated once a month. Inflation-linked derivatives can therefore depend on the value of an inflation index in any given month. In the QGY model, the process for the inflation index I_T is only defined discretely on the model tenor structure, that is for $T \in \{T_k\}_{k=0,1,2,\ldots}$. However, rather than assuming a monthly tenor structure when we reparameterised the model in Chapter 7, we assumed that we had an annual tenor structure. We did this so that we could use the semi-analytical pricing formula for year-on-year inflation options in §6.4.

In §9.1, we define two interpolation schemes. For each scheme we specify the process of the inflation index at any point in time in a way that is consistent when compared to the original model specifications. This enables us to calculate the drift of the inflation index and price inflation-linked derivatives that depend on the value of the inflation index on dates that are not part of the tenor structure. In the wedged interpolation, we use the original process for the inflation index, except for dates that are not part of the tenor structure. This interpolation method is well suited to pricing zero-coupon derivatives. In the shifted interpolation we shift the whole tenor structure so that two consecutive dates are one-year apart. This interpolation method is better suited to pricing of year-on-year derivatives.

In §9.2, we define parameter interpolation and extrapolation methods that are consistent with the behavior of the inflation market and of both the wedged and the shifted interpolation schemes. In §9.3, we look at some numerical examples using these interpolation schemes and we see that these methods are stable for the year-on-year inflation forward rate and the year-on-year and zero-coupon inflation volatilities.

9.1 Interpolation methods

In Chapter 5, we showed that the inflation index I_{T_k} in the QGY model is given by

$$I_{T_k} = I_{0T_k} \prod_{i=1}^{k} Y_{T_i} = I_{0T_k} \mathcal{A}_{T_k} \prod_{i=1}^{k} X_{T_i} \left[\phi_{T_i}^y, \Psi_{T_i}^y \right]. \tag{9.1}$$

Here we have defined a general drift of the inflation index for times T_k on the tenor structure, $k = 0, 1, 2, \ldots$:

$$\mathcal{A}_{T_k} := \sum_{i=1}^{k} A_{T_i}. \tag{9.2}$$

Now we would like to define $\mathcal{A}_T$ for a time T that is not on the tenor structure. One way would be to do a direct interpolation (e.g. flat/linear/cubic) from the known datapoints $\mathcal{A}_{T_k}$. Another, more consistent way is to define a process for I_T that in turn could be used for pricing different inflation-linked payoffs. The initial condition $\mathbb{E}_0^n \left[D_T^n I_T \right] = P_{0T} I_{0T}$ can then be used to determine the value of $\mathcal{A}_T$. We will specify two interpolation methods—the wedged interpolation and the shifted interpolation.

Wedged interpolation

Given a time T, we choose k so that $T_{k-1} < T < T_k$. In the wedged interpolation, we add the time T to the tenor structure, so that it is situated (wedged) between T_{k-1} and T_k. Now we need to specify the evolution of the process I_T. The process for $I_{T_{k-1}}$ is in the same form as in (9.1), so what remains is to define a process for $I_T \big/ I_{T_{k-1}}$. As for the year-on-year inflation ratio in (5.3), for consistency, we choose a log-quadratic Gaussian process that is driven by the same Gaussian process as before at time T, namely x_T:

$$\frac{I_T}{I_{T_{k-1}}} := \frac{I_{0T}}{I_{0T_{k-1}}} \exp\left[\mathcal{A}_T - \mathcal{A}_{T_{k-1}} \right] X_T \left[\widehat{\phi}_T^y, \widehat{\Psi}_T^y \right]. \tag{9.3}$$

Here we need to choose values of two new parameters $\left(\widehat{\phi}_T^y, \widehat{\Psi}_T^y \right)$ that satisfy the following limit conditions:

$$\begin{aligned} \left(\widehat{\phi}_T^y, \widehat{\Psi}_T^y \right) &\longrightarrow (0, 0) && \text{as} \quad T \searrow T_{k-1}, \\ \left(\widehat{\phi}_T^y, \widehat{\Psi}_T^y \right) &\longrightarrow \left(\phi_{T_k}^y, \Psi_{T_k}^y \right) && \text{as} \quad T \nearrow T_k. \end{aligned} \tag{9.4}$$

To determine the value of the drift $\mathcal{A}_T$, we temporarily add the time T to the tenor structure. We can then use §6.2 with the process in (9.3) to get

$$\exp\left[\mathcal{A}_T \right] = \frac{E_{0T}\left[\phi_T^n, \Psi_T^n \right]}{E_{T_k T}\left[\Xi_T^T \right] \prod_{i=1}^{k} E_{T_{i-1} T_i} \left[\Xi_{T_i}^T \right]},$$

where the parameter transformation is defined by

$$\begin{aligned}
\Xi_T^T &:= \left(\widehat{\phi}_T^y + \phi_T^n, \widehat{\Psi}_T^y + \Psi_T^n\right), \\
\Xi_{T_{k-1}}^T &:= \Theta_{T_{k-1}T}\left[\Xi_T^T\right] + \left(\widehat{\phi}_{T_{k-1}}^y, \widehat{\Psi}_{T_{k-1}}^y\right), \\
\Xi_{T_i}^T &:= \Theta_{T_iT_{i+1}}\left[\Xi_{T_{i+1}}^T\right] + \left(\phi_{T_i}^y, \Psi_{T_i}^y\right), \quad \text{for } i = 1, 2, \ldots, k-2.
\end{aligned}$$

We also want to use this method to price derivatives that depend on the general inflation ratio $I_T/\,I_t$, with $t < T$. We find $h, k \geq 0$ so that $h < k$, with $t \in [T_h, T_{h+1})$ and $T \in (T_{k-1}, T_k]$, and from (9.3) we write

$$\begin{aligned}
\frac{I_T}{I_t} &= \left(\frac{I_t}{I_{T_h}}\right)^{-1} \frac{I_{T_{k-1}}}{I_{T_h}} \frac{I_T}{I_{T_{k-1}}} \\
&= \frac{I_{0T}}{I_{0t}} \exp\left[\mathcal{A}_T - \mathcal{A}_t\right] X_t\left[-\widehat{\phi}_t^y, -\widehat{\Psi}_t^y\right] \left(\prod_{i=h+1}^{k-1} X_{T_i}\left[\phi_{T_i}^y, \Psi_{T_i}^y\right]\right) X_T\left[\widehat{\phi}_T^y, \widehat{\Psi}_T^y\right]. \qquad (9.5)
\end{aligned}$$

Let $0 \leq s \leq t < T \leq T'$. Using similar methods as in §6.3, the value at time s of a swaplet on the inflation ratio $I_T/\,I_t$ with a payment delay $T' - T$ is given by

$$\begin{aligned}
\frac{1}{D_s^n}\mathbb{E}_s^n\left[D_T^n \frac{I_T}{I_t}\right] = {}& \frac{P_{0T'}^n}{D_s^n} \frac{I_{0T}}{I_{0t}} \exp\left[\mathcal{A}_T - \mathcal{A}_t\right] X_{st}\left[\Xi_t^{TT'}\right] E_{st}\left[\Xi_t^{TT'}\right] E_{tT_{h+1}}\left[\Xi_{h+1}^{TT'}\right] \\
& \times \left(\prod_{i=h+2}^{k-2} E_{T_{i-1}T_i}\left[\Xi_{T_i}^{TT'}\right]\right) E_{T_{k-1}T}\left[\Xi_T^{TT'}\right] \frac{E_{TT'}\left[\phi_{T'}^n, \Psi_{T'}^n\right]}{E_{0T'}\left[\phi_{T'}^n, \Psi_{T'}^n\right]}.
\end{aligned}$$

The parameter transformation $\Xi^{TT'}$ is defined iteratively in the following way:

$$\begin{aligned}
\Xi_T^{TT'} &:= \Theta_{TT'}\left[\phi_{T'}^n, \Psi_{T'}^n\right] + \left(\widehat{\phi}_T^y, \widehat{\Psi}_T^y\right), \\
\Xi_{T_{k-1}}^{TT'} &:= \Theta_{T_{k-1}T}\left[\Xi_T^{T'T}\right] + \left(\phi_{T_{k-1}}^y, \Psi_{T_{k-1}}^y\right), \\
\Xi_{T_i}^{TT'} &:= \Theta_{T_iT_{i+1}}\left[\Xi_{T_{i+1}}^T\right] + \left(\phi_{T_i}^y, \Psi_{T_i}^y\right), \quad \text{for } i = h+1, h+2, \ldots, k-2, \\
\Xi_t^{TT'} &:= \Theta_{tT_{h+1}}\left[\Xi_{T_{h+1}}^{TT'}\right] - \left(\widehat{\phi}_t^y, \widehat{\Psi}_t^y\right).
\end{aligned}$$

As we only need to interpolate one set of parameters $\left(\widehat{\phi}_T^y, \widehat{\Psi}_T^y\right)$, a good property of this method is that the process for the inflation index I_T builds on the original year-on-year ratios Y_{T_i} for $i = 1, 2, \ldots, k-1$ and we only need to interpolate from T_{k-1} to T to get I_T. The wedged-interpolation method should therefore be well suited to pricing inflation-linked swaps by means of closed-form expressions, while inflation-linked options would have to be priced by simulation.

In general, $I_T/\,I_t$ is not Markovian because it depends on the driving process x_s at more than one time s. If we assume that T is not on the annual tenor structure, then the same holds for $t = T - 1$. We find $k \geq 0$ so that $T_{k-1} < t < T_k < T < T_{k+1}$. Then I_T / I_t depends on x_t, x_{T_k}, and x_T as in (9.5). This means that we cannot use the semi-analytical formula of a year-on-year option from §6.4. This is where the shifted interpolation can be put to good use.

Shifted interpolation

Given a time T, we choose k so that $T_{k-1} < T < T_k$. We then define the following time-weight:

$$\lambda_T := \frac{T - T_{k-1}}{T_k - T_{k-1}}.$$

In the shifted interpolation, we shift the original tenor structure relative to λ_T. We let $T_0^{\lambda_T} := T_0$ and define

$$T_i^{\lambda_T} := T_{i-1} + \lambda_T (T_i - T_{i-1}) \quad \text{for} \quad i = 1, 2, \ldots, k.$$

Now we need to define the year-on-year ratios $Y_{T_i^{\lambda_T}} = I_{T_i^{\lambda_T}} / I_{T_{i-1}^{\lambda_T}}$ on the new tenor structure $\{T_i^{\lambda_T}\}_{i=0,1,2,\ldots,k}$. As for the year-on-year inflation ratio in (5.3), for consistency we choose a log-quadratic Gaussian process that is driven by the same Gaussian process as before at time $T_i^{\lambda_T}$, namely $x_{T_i^{\lambda_T}}$:

$$Y_{T_i^{\lambda_T}} := \frac{I_{0T_i^{\lambda_T}}}{I_{0T_{i-1}^{\lambda_T}}} \exp\left[A_{T_i^{\lambda_T}}\right] X_{T_i^{\lambda_T}} \left[\phi^y_{T_i^{\lambda_T}}, \Psi^y_{T_i^{\lambda_T}}\right]. \tag{9.6}$$

For this method, we need to choose values for k new parameter sets $\left(\phi^y_{T_i^{\lambda_T}}, \Psi^y_{T_i^{\lambda_T}}\right)$, $i = 1, 2, \ldots, k$, which satisfy the following boundary limit conditions:

$$\begin{aligned}
\left(\phi^y_{T_i^{\lambda_T}}, \Psi^y_{T_i^{\lambda_T}}\right) &\longrightarrow \left(\phi^y_{T_{i-1}}, \Psi^y_{T_{i-1}}\right) \quad &\text{as} \quad T_i^{\lambda_T} \searrow T_{i-1}, \\
\left(\phi^y_{T_i^{\lambda_T}}, \Psi^y_{T_i^{\lambda_T}}\right) &\longrightarrow \left(\phi^y_{T_i}, \Psi^y_{T_i}\right) \quad &\text{as} \quad T_i^{\lambda_T} \nearrow T_i.
\end{aligned} \tag{9.7}$$

Now the inflation index I_T is retrieved from the relation

$$I_T = I_{0T_0} \prod_{i=1}^{k} Y_{T_i^{\lambda_T}}. \tag{9.8}$$

Therefore, determining the drift $A_{T_i^{\lambda_T}}$ is done as in §6.2, so we can use Proposition 6.2.1 directly on the new tenor structure $\{T_i^{\lambda_T}\}_{i=0,1,2,\ldots,k}$.

Here we need to interpolate k new parameter sets and use the values of the driving process x_t on the new tenor structure, $t \in \{T_i^{\lambda_T}\}_{i=0,1,2,\ldots,k}$, rather than on the old tenor structure, $t \in \{T_i\}_{i=0,1,2,\ldots,k}$. The shifted-interpolation method should be well suited to price trades that depend on the year-on-year inflation ratio $Y_T = I_T / I_{T-1}$, as the method preserves the Markovianity property of the ratio. This means that Y_T only depends on the driving process at T, namely x_T. Therefore, year-on-year options can be calculated using the pricing formula in §6.4, while limited price indices (LPIs) still need to be priced using simulation.

9.2 Parameter interpolation and extrapolation

In §7.4, we defined a three-dimensional spherical reparameterisation of the QGY model. We assumed $G^{y_1 y_2}_{T_k} = 0$ ($\rho^{y_1 y_2}_t = 0$) and $\phi^{y_2}_{T_k} = \Psi^{y_2}_{T_k} = 0$. The three polar coordinates are the volatility level $\Sigma^y_{T_k} \geq 0$, the volatility-of-volatility $\sin \nu^y_{T_k} \in [0,1]$ (with $\nu^y_{T_k} \in [0, \pi/2]$) and the volatility-correlation $\sin \rho^y_{T_k} \in [-1,1]$ (with $\rho^y_{T_k} \in [-\pi/2, \pi/2]$). These parameters give an intuitive control over the year-on-year inflation volatility smile. We will use them to interpolate and extrapolate the new parameters of the wedged- and the shifted-interpolation methods.

Covariance matrix

We start by looking at the driving factor

$$x_t = \int_0^t \operatorname{diag}(\sigma_s)\, dW^n_s.$$

We know that σ_s is strictly positive. As we only know the values σ_{T_i}, $i = 0, 1, 2, \ldots$, we assume that each of the entry of $\ln[\sigma_s]$ is linear on each of the periods $(T_{i-1}, T_i]$,

$$\begin{aligned} & \ln[\sigma_s] := \ln\left[\sigma_{T_{i-1}}\right] + \lambda_s \left(\ln[\sigma_{T_i}] - \ln\left[\sigma_{T_{i-1}}\right]\right) \\ \Longleftrightarrow \qquad & \sigma_s = \sigma_{T_{i-1}} \left(\frac{\sigma_{T_i}}{\sigma_{T_{i-1}}}\right)^{\lambda_s}, \end{aligned}$$

where the time-weight λ_s is defined by

$$\lambda_s := \frac{s - T_{i-1}}{T_i - T_{i-1}}.$$

For some $t \in (T_{i-1}, T_i]$, each element $\{p, q\}$ of the conditional covariance matrix of $x_t \,|x_{T_{i-1}}$ is given by

$$\begin{aligned} G_t^{pq} - G_{T_{i-1}}^{pq} &= \int_{T_{i-1}}^{t} \rho_s^{pq} \sigma_s^p \sigma_s^q ds = (T_i - T_{i-1})\, \rho_{T_{i-1}}^{pq} \sigma_{T_{i-1}}^p \sigma_{T_{i-1}}^q \int_0^{\lambda_t} \left(\frac{\sigma_{T_i}^p}{\sigma_{T_{i-1}}^q} \frac{\sigma_{T_i}^p}{\sigma_{T_{i-1}}^q} \right)^{\lambda_s} d\lambda_s \\ &= (T_i - T_{i-1})\, \rho_{T_{i-1}}^{pq} \sigma_{T_{i-1}}^p \sigma_{T_{i-1}}^q \ln^{-1} \left[\frac{\sigma_{T_i}^p}{\sigma_{T_{i-1}}^p} \frac{\sigma_{T_i}^q}{\sigma_{T_{i-1}}^q} \right] \left[\left(\frac{\sigma_{T_i}^p}{\sigma_{T_{i-1}}^p} \frac{\sigma_{T_i}^q}{\sigma_{T_{i-1}}^q} \right)^{\lambda_t} - 1 \right]. \end{aligned} \tag{9.9}$$

Here we changed variables from s to λ_s, as well as assuming piecewise flat values for ρ_s^{pq} on $[T_{i-1}, T_i)$ (right-continuous). For the boundary case, we get

$$\frac{\sigma_{T_i}^p}{\sigma_{T_{i-1}}^p} \frac{\sigma_{T_i}^q}{\sigma_{T_{i-1}}^q} = 1 \quad \Longrightarrow \quad G_t^{pq} - G_{T_{i-1}}^{pq} = (t - T_{i-1})\, \rho_{T_{i-1}}^{pq} \sigma_{T_{i-1}}^p \sigma_{T_{i-1}}^q.$$

Now, given $G_{T_i}^{pq}$ and $\Sigma_{T_i}^{pq} := \frac{\sigma_{T_i}^p}{\sigma_{T_{i-1}}^p} \frac{\sigma_{T_i}^q}{\sigma_{T_{i-1}}^q}$ for $i = 1, 2, \ldots$, we find for some $t \in (T_{i-1}, T_i]$ that

$$G_t^{pq} = \begin{cases} G_{T_{i-1}}^{pq} + \frac{\left(\Sigma_{T_i}^{pq}\right)^{\lambda_t} - 1}{\left(\Sigma_{T_i}^{pq}\right) - 1} \left(G_{T_i}^{pq} - G_{T_{i-1}}^{pq} \right) & \text{if } \Sigma_{T_i}^{pq} \neq 1 \\ G_{T_{i-1}}^{pq} + \lambda_t \left(G_{T_i}^{pq} - G_{T_{i-1}}^{pq} \right) & \text{if } \Sigma_{T_i}^{pq} = 1 \end{cases}$$

Here $G_{T_i}^{pq} - G_{T_{i-1}}^{pq}$ is obtained from (9.9) with $t = T_i$.

Wedged spherical parameters

For the wedged interpolations, we will determine the values of three new polar parameters $\left(\widehat{\Sigma}_T^y, \widehat{\nu}_T^y, \widehat{\rho}_T^y\right)$ for the inflation ratio $I_T / I_{T_{k-1}}$ in (9.3), where $T_{k-1} < T < T_k$. As in §7.1 and §7.4, the values of the interpolated parameters $\left(\widehat{\phi}_T^{y_1}, \widehat{\Psi}_T^{y_1}, \widehat{\Psi}_T^{y_1 y_2}\right)$ are then retrieved from the relation

$$\begin{aligned} -\widehat{\phi}_T^{y_1} \sqrt{G_T^{y_1}} &= \frac{\widehat{\Sigma}_T^y}{\widehat{K}_T^y} \cos\left[\widehat{\nu}_T^y\right], \\ -\frac{1}{2} \widehat{\Psi}_T^{y_1} G_T^{y_1} &= \frac{\widehat{\Sigma}_T^y}{\widehat{K}_T^y} \sin\left[\widehat{\nu}_T^y\right] \sin\left[\widehat{\rho}_T^y\right], \\ \widehat{\Psi}_T^{y_1 y_2} \sqrt{G_T^{y_1} G_T^{y_2}} &= \frac{\widehat{\Sigma}_T^y}{\widehat{K}_T^y} \sin\left[\widehat{\nu}_T^y\right] \cos\left[\widehat{\rho}_T^y\right], \end{aligned}$$

where $G_T^{y_1}$ and $G_T^{y_2}$ are retrieved from (9.9), and

$$\widehat{K}_T^y = \sqrt{\sin^2\left[\widehat{\nu}_T^y\right] \sin^2\left[\widehat{\rho}_T^y\right] + 1},$$

so that

$$\sqrt{\mathrm{Var}_0^n\left[\ln\left[I_T/\,I_{T_{k-1}}\right]\right]} = \widehat{\Sigma}_T^y.$$

We assume that $\widehat{\Sigma}_T^y$ is a linear scaling of $\Sigma_{T_k}^y$, that is

$$\widehat{\Sigma}_T^y := \lambda_T \Sigma_{T_k}^y \qquad \text{with} \quad \lambda_T = \frac{T - T_{k-1}}{T_k - T_{k-1}}.$$

This interpolation satisfies the conditions in (9.4) because the volatility $\widehat{\Sigma}_T^y$ vanishes at $T \searrow T_{k-1}$ and $\widehat{\Sigma}_T^y$ increases to the original volatility $\widehat{\Sigma}_{T_k}^y$ as $T \nearrow T_k$. We do not scale the parameters $\widehat{\nu}_T^y$ and $\widehat{\rho}_T^y$, as they describe the (relative) volatility-of-volatility and the volatility-correlation, which are independent of the length of the interval $T - T_{k-1}$. We therefore assume that $\widehat{\nu}_T^y$ and $\widehat{\rho}_T^y$ are piecewise flat:

$$\begin{aligned}
\widehat{\nu}_T^y &:= \nu_{T_k}^y, \\
\widehat{\rho}_T^y &:= \rho_{T_k}^y.
\end{aligned}$$

Shifted spherical parameters

In the wedged interpolation, we have three spherical parameters $(\Sigma_t^y, \nu_t^y, \rho_t^y)$ for each time on the new tenor structure, $t \in \{T_i^{\lambda_T}\}_{i=0,1,2,\ldots,k}$. As in §7.1 and §7.4, for the interpolated parameters $(\phi_t^{y_1}, \Psi_t^{y_1}, \Psi_t^{y_1 y_2})$ we have the relation

$$\begin{aligned}
-\phi_t^{y_1}\sqrt{G_t^{y_1}} &= \frac{\Sigma_t^y}{K_t^y}\cos\left[\nu_t^y\right], \\
-\frac{1}{2}\Psi_t^{y_1} G_t^{y_1} &= \frac{\Sigma_t^y}{K_t^y}\sin\left[\nu_t^y\right]\sin\left[\rho_t^y\right], \\
\Psi_t^{y_1 y_2}\sqrt{G_t^{y_1} G_t^{y_2}} &= \frac{\Sigma_t^y}{K_t^y}\sin\left[\nu_t^y\right]\cos\left[\rho_t^y\right],
\end{aligned}$$

where $G_t^{y_1}$ and $G_t^{y_2}$ are retrieved from (9.9), and

$$K_t^y = \sqrt{\sin^2\left[\nu_t^y\right]\sin^2\left[\rho_t^y\right] + 1},$$

so that

$$\sqrt{\mathrm{Var}_0^n\left[\ln\left[I_{T_i^{\lambda_T}}/I_{T_{i-1}^{\lambda_T}}\right]\right]} = \Sigma_{T_i^{\lambda_T}}^y.$$

The spherical parameters of the new tenor structure $\{T_i^{\lambda_T}\}_{i=0,1,2,\ldots,k}$ need to be calculated from the spherical parameters of the original tenor structure $\{T_i\}_{i=0,1,2,\ldots,k}$, where we have

$$T_0^{\lambda_T} = T_0 < T_1^{\lambda_T} < T_1 < \cdots < T_{i-1}^{\lambda_T} < T_{i-1} < T_i^{\lambda_T} < T_i < \cdots < T_{k-1}^{\lambda_T} < T_{k-1} < T_k^{\lambda_T} < T_k.$$

We assume that the new parameters are piecewise linear:

$$\begin{aligned}
\Sigma^y_{T_i^{\lambda_T}} &:= (1-\lambda_T)\,\Sigma^y_{T_{i-1}} + \lambda_T \Sigma^y_{T_i},\\
\nu^y_{T_i^{\lambda_T}} &:= (1-\lambda_T)\,\nu^y_{T_{i-1}} + \lambda_T \nu^y_{T_i},\\
\rho^y_{T_i^{\lambda_T}} &:= (1-\lambda_T)\,\rho^y_{T_{i-1}} + \lambda_T \rho^y_{T_i}.
\end{aligned}$$

This interpolation satisfies the conditions in (9.7). The year-on-year inflation ratio $Y_{T_i^{\lambda_T}} = I_{T_i^{\lambda_T}}/I_{T_{i-1}^{\lambda_T}}$ spans the period $(T_{i-1}^{\lambda_T}, T_i^{\lambda_T}]$, which we can split into two periods, $(T_{i-1}^{\lambda_T}, T_{i-1}]$ and $(T_{i-1}, T_i^{\lambda_T}]$. Intuitively we put a weight of $1-\lambda_T = (T_{i-1} - T_{i-1}^{\lambda_T})/(T_{i-1} - T_{i-2})$ on the parameters of $Y_{T_{i-1}}$, which spans the period $(T_{i-2}, T_{i-1}]$. In turn, we put a weight of $\lambda_T = (T_i^{\lambda_T} - T_i)/(T_i - T_{i-1})$ on the parameters of Y_{T_i}, which spans the period $(T_{i-1}, T_i]$.

Extrapolated tenor structure

In some cases, the original tenor structure is only defined up to a certain time T_k—that is, we are only given $\{T_i\}_{i=0,1,2,\ldots,k}$. We would still want to price inflation-linked trades that depend on the value of the inflation index at some later time $T > T_k$. We start by defining the last known tenor difference

$$\delta_{T_k} := T_k - T_{k-1}.$$

Then we extend the tenor structure to include new dates by recursively adding δ_{T_k}—that is,

$$T_i := T_{i-1} + \delta_{T_k} \quad \text{for} \quad i = k+1, k+2, \ldots.$$

We now need to extrapolate new values for the parameters of the model. We start by assuming that $\ln[\sigma_s]$ is linear for all $s > T_k$, defining

$$\begin{aligned}
\ln[\sigma_s] &:= \ln\left[\sigma_{T_{k-1}}\right] + \frac{s - T_{k-1}}{\delta_{T_k}}\left(\ln[\sigma_{T_k}] - \ln\left[\sigma_{T_{k-1}}\right]\right)\\
\iff \qquad \sigma_s &= \sigma_{T_{k-1}} \left(\frac{\sigma_{T_k}}{\sigma_{T_{k-1}}}\right)^{\frac{s-T_{k-1}}{\delta_{T_k}}}.
\end{aligned}$$

We can then use (9.9) to calculate the variance matrix G_t for all $t > T_k$.

Now it remains to extrapolate the other model curves. We assume flat spherical parameters for all T_i, $i = k+1, k+2, \ldots$—that is,

$$\Sigma^y_{T_i} := \Sigma^y_{T_k},$$
$$\nu^y_{T_i} := \nu^y_{T_k},$$
$$\rho^y_{T_i} := \rho^y_{T_k}.$$

This extrapolation gives similar year-on-year inflation volatility smiles for all the extrapolated tenors. As in §7.1 and §7.4, for the extrapolated parameters $\left(\phi^{y_1}_{T_i}, \Psi^{y_1}_{T_i}, \Psi^{y_1 y_2}_{T_i}\right)$ we have the relation

$$-\phi^{y_1}_{T_i}\sqrt{G^{y_1}_{T_i}} = \frac{\Sigma^y_{T_i}}{K^y_{T_i}} \cos\left[\nu^y_{T_i}\right],$$
$$-\frac{1}{2}\Psi^{y_1}_{T_i} G^{y_1}_{T_i} = \frac{\Sigma^y_{T_i}}{K^y_{T_i}} \sin\left[\nu^y_{T_i}\right] \sin\left[\rho^y_{T_i}\right],$$
$$\Psi^{y_1 y_2}_{T_i}\sqrt{G^{y_1}_{T_i} G^{y_2}_{T_i}} = \frac{\Sigma^y_{T_i}}{K^y_{T_i}} \sin\left[\nu^y_{T_i}\right] \cos\left[\rho^y_{T_i}\right].$$

Here $G^{y_1}_{T_i}$ and $G^{y_2}_{T_i}$ are retrieved from (9.9), and

$$K^y_{T_i} = \sqrt{\sin^2\left[\nu^y_{T_i}\right] \sin^2\left[\rho^y_{T_i}\right] + 1},$$

so that

$$\sqrt{\mathrm{Var}^n_0\left[\ln\left[I_{T_i}/I_{T_{i-1}}\right]\right]} = \Sigma^y_{T_i}.$$

Now, for a trade that depends on I_T for some $T > T_k$, we can then use all the pricing formulas as before, as well as using the wedged interpolation or the shifted interpolation if necessary.

9.3 Numerical examples

Here we show how the wedged and shifted interpolations perform. We calculate monthly values from one-year to fifty-year maturities. At maturities of whole integer years ($1y$, $2y$, etc.) we do not require interpolation but at the maturities in between we do require one. We use the values of the calibrated QGY parameters on 31 January 2012 (see §8.1 and §8.2). We therefore need to extrapolate the tenor structure for maturities over thirty years.

We value the year-on-year convexity corrections and the year-on-year inflation volatilities (see §2.3) for strikes equal to the year-on-year inflation forward rates (at-the-money) using the shifted interpolation. We value the zero-coupon inflation volatilities (see §2.3) for strikes equal to the forward inflation index (at-the-money) using the wedged interpolation. The results are shown in Figure 9.1.

The year-on-year convexity correction is reasonably stable. We notice that the interpolated values showed small oscillations between integer maturities. The extrapolation for maturities over thirty years shows a slightly steeper decay in the year-on-year convexity correction. One could change the rate of decay by changing the value of the extrapolated growth rate $R^y_{T_K}$.

The interpolation of the year-on-year inflation volatilities is very stable for all maturities, including the extrapolated maturities of more than thirty years. The interpolation of the zero-coupon inflation volatilities is quite stable. The oscillations between integer maturities are more pronounced from one-year to five-year maturities but even out for higher maturities, including the extrapolated maturities.

We see that both the wedged and the shifted interpolation methods perform very well. For both the year-on-year and zero-coupon inflation volatilities, we note that we get the same stability for other strikes. All oscillations are within reasonable values, which suggests that the linear interpolation of the spherical parameters is well-suited for both the wedged and the shifted interpolation methods.

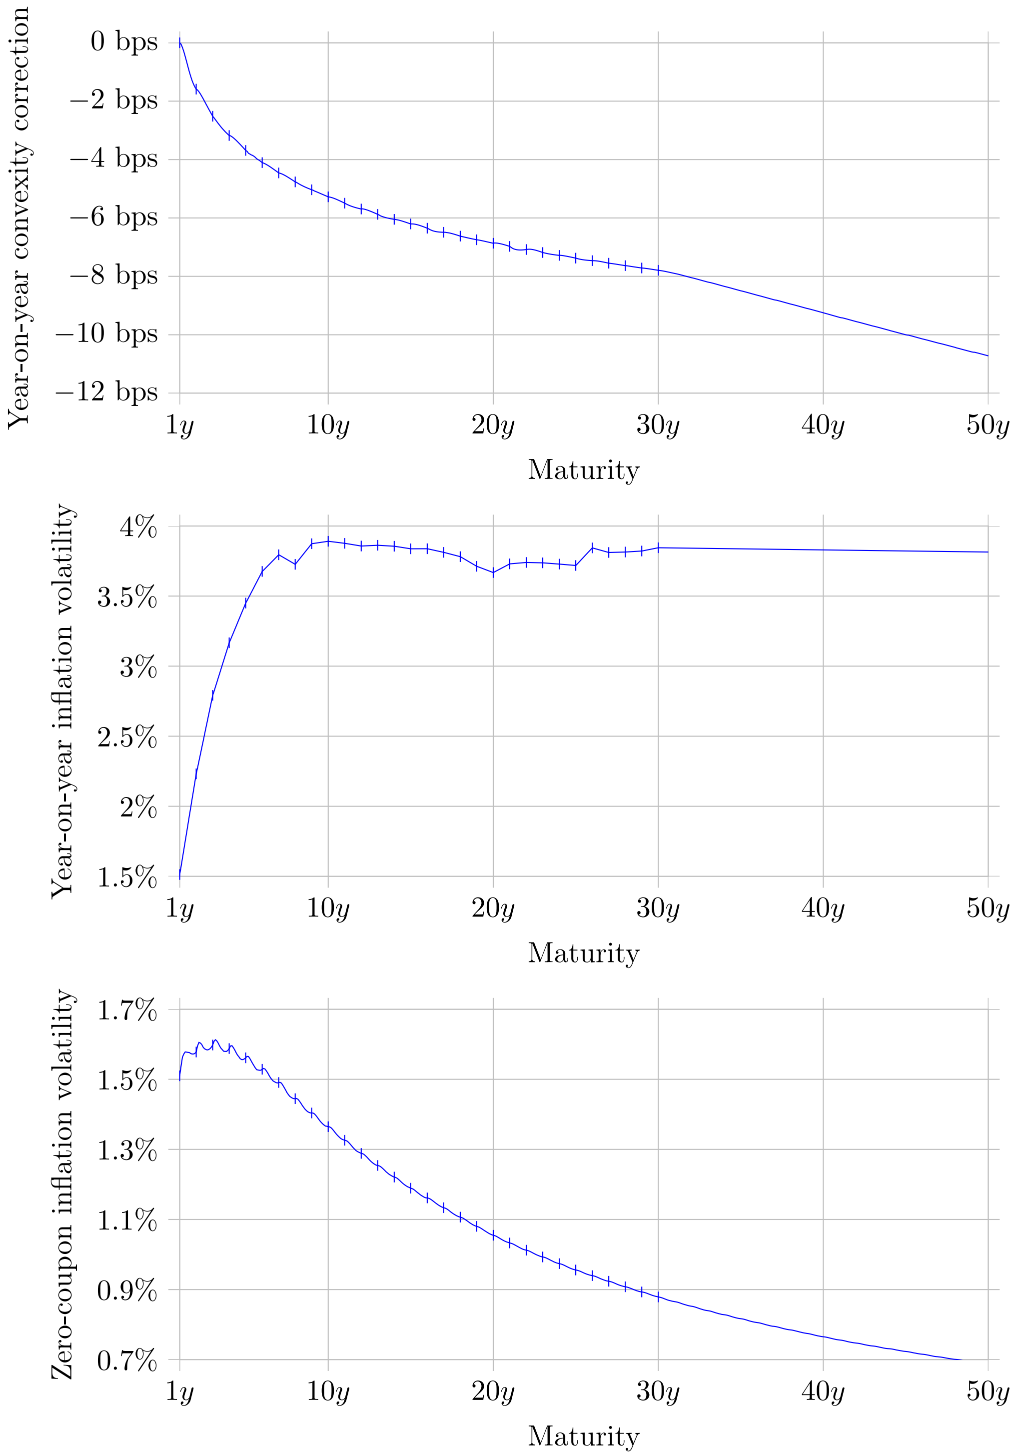

Figure 9.1: Using the shifted interpolation, the top graph shows the year-on-year convexity correction and the middle graph shows year-on-year inflation volatilities for strikes equal to the year-on-year inflation forward rates. Using the wedged interpolation, the bottom graph shows the zero-coupon inflation volatilities for strikes equal to the forward inflation index. We have used the values of the calibrated QGY parameters on 31 January 2012 (see §8.1 and §8.2).

Conclusion

In this thesis, we presented the Quadratic Gaussian Year-on-Year inflation model—the QGY model. It incorporates curvature and skew in the volatility smile implied by market prices of inflation-linked options, while having good analytical properties, intuitive parameterisation, low dimensionality and simple simulation. It overcomes some shortcomings of other inflation models, such as problems in estimating parameters that depend on the real rate in foreign-exchange analogy inflation models and computationally intensive inflation market models.

Specifications and formulas

In the QGY model, the year-on-year ratio of the inflation index is a log-quadratic function of a multi-factor Gaussian process. This ratio gives us well-defined processes for both the year-on-year inflation rate and the inflation index. We derived pricing formulas in the QGY model by using analytical expressions for the expectations of log-quadratic Gaussian processes. We derived closed-form expressions for the drift of the inflation index, as well as for swaps on the inflation index and the year-on-year rate. We derived a semi-analytical Black-Scholes-type pricing formula for year-on-year inflation options. The formula contains an integral of a standard Gaussian density over a domain bounded by a conic section. In two dimensions, we used polar coordinates to reduce this integral to a finite sum of one-dimensional integrals that could be calculated quickly and accurately.

Inflation markets and reparameterisation

By looking at the payoffs and prices of inflation-linked derivatives, we analysed the distributional properties of the inflation index and the year-on-year rate. We considered two types of volatility implied by prices of inflation-linked options with a given strike and maturity. The first is the zero-coupon inflation volatility that is implied by prices of zero-coupon inflation options. The second is the year-on-year inflation volatility that is implied by year-on-year inflation options. We looked at market prices of these options on the UK's Retail Prices Index (RPI) and the eurozone's Harmonised Index of Consumer Prices Excluding Tobacco (HICPxT). The zero-coupon inflation volatilities indicate that the volatility of the inflation index increases in approximately linearly with maturity. The year-on-year inflation volatilities indicate that the volatility of the year-on-year inflation rate increases with maturity for maturities up to five years and then stabilises for higher maturities. Because of the regular behaviour of these volatilities, we saw that the positive curvature and the negative skew in the zero-coupon and year-on-year inflation smiles are of the same magnitude for all maturities.

We also analysed the payment delay and year-on-year convexity corrections, which are incorporated into zero-coupon and year-on-year inflation swaps. We derived an explicit formula, which shows that the payment delay convexity correction decreases in the correlation between inflation rates and nominal interest rates. We approximated the year-on-year convexity correction and showed that it decreases in the autocorrelation of the inflation index and increases in the correlation between inflation rates and nominal interest rates. The QGY model can account for these convexity corrections and the inflation volatility smiles.

In order to analyse the shape of the inflation volatility smile obtained from the QGY model, we reduced the number of parameters so that we could examine the distribution of the year-on-year inflation ratio in more detail. We looked at a three-factor version of the QGY model and separated the parameters so that the nominal bond was driven by a one-factor log-linear Gaussian process and the year-on-year inflation ratio was driven by a two-factor log-quadratic Gaussian process. In the three-factor QGY model, we defined a spherical reparameterisation that had three parameters, which conveniently give separate control for the level, curvature and skew of the year-on-year inflation volatility smile. For this parameterisation, we computed the upper bound of the curvature and found the maximum and minimum skew.

Within the three-factor QGY model, we identified a parameterisation to control the inflation index autocorrelation and, in turn, the year-on-year inflation convexity correction of the year-on-year inflation forward rates. Furthermore, we were able to adjust the correlation between nominal interest rates and inflation rates, which controls the payment delay convexity correction.

Calibration and simulation

Using the spherical parameterisation, we calibrated the QGY model to year-on-year inflation volatilities implied by market prices of year-on-year options on RPI and HICPxT. We also fitted the year-on-year inflation forward rates exactly by adjusting the autocorrelation of the inflation index. The calibration was fast because of the intuitive parameters and the semi-analytical pricing formula for year-on-year inflation options. The calibrated parameters were stable across maturities and for different calibration dates. The calibration obtained the model's maximum curvature of the year-on-year smile to get close to the market's curvature. The pricing error was fairly low for HICPxT volatilities, while it was higher for RPI volatilities because of a higher curvature.

Simulating the QGY model using Monte Carlo methods was simple because the inflation index is in a functional form and driven by a Gaussian process. In the QGY model, the inflation index depends on annual observations of the process for the year-on-year inflation ratio. Accordingly, to calculate the price of zero-coupon inflation options, we simulated the value of the driving process at every year. Using the calibrated QGY model, we compared zero-coupon inflation volatilities to market volatilities in HICPxT. We observed that the zero-coupon inflation volatility skew in the QGY model was more negative than in the market. This was because the QGY model was calibrated to the year-on-year inflation volatility skew, which is more negative than the zero-coupon inflation volatility skew. The QGY model also gave a lower volatility level because the autocorrelation implied by the year-on-year inflation forward rates was too low. Using Monte Carlo simulation methods, we also priced limited price indices on RPI. We noticed that the QGY model gave smaller pricing errors than a model that did not incorporate curvature and skew into the year-on-year inflation smile.

Interpolation and implementation

In the QGY model, the value of the inflation index is only defined discretely for dates on the annual tenor structure. We provided interpolation methods that specified the process for the inflation index and for the year-on-year inflation rate for dates that were not on the tenor structure. We then used analytical and semi-analytical pricing formulas to price inflation-linked swaps and year-on-year options that depended on the value of the inflation index for dates that were not on the tenor structure. Furthermore, the interpolation methods made it easy to simulate the processes using Monte Carlo methods and to price zero-coupon options and more exotic inflation-linked derivatives. From numerical examples, we saw that the interpolation methods gave stable values of the year-on-year convexity corrections and of the year-on-year and zero-coupon inflation volatilities.

We implemented the three-factor QGY model in the quantitative library of the quantitative research team at Lloyds Banking Group. Doing so involved programming in C++ using a highly object-oriented approach. The C++ code allowed for very fast pricing of inflation-linked derivatives in the QGY model. We were able to perform extensive tests on the model, e.g. its pricing formulas and results, parameterisations, simulation and interpolation methods.

Further research

It is clear that the QGY model has many good properties and can capture the different properties of the inflation market well. Further research can be carried out on the QGY model. It remains to look at the sensitivities of prices to the parameters of the model—i.e. the Greeks. The general modelling framework can be used to price hybrid products by linking the model to other financial quantities, e.g. to nominal interest rates and inflation indices in other currencies. Because the conditional distribution of the year-on-year inflation rate is given in closed form, we could price callable year-on-year inflation bonds effectively by using accurate integration methods on a discretised grid. To price these bonds we could use calculation methods applied to Markov-functional models by Pelsser (2000).

The characteristic function of the inflation index is given in closed form in the QGY model, so, as in Carr et al. (1999), it would be possible to price zero-coupon inflation options using a fast Fourier transform method.[1] This would speed up the pricing of zero-

[1]The characteristic function of the general inflation ratio in the QGY model is given by

$$\mathbb{E}_t^n\left[\mathrm{e}^{iu\ln\left[I_{T_k}/I_{T_h}\right]}\right]=\mathbb{E}_t^n\left[\left(\frac{I_{T_k}}{I_{T_h}}\right)^{iu}\right]=\left(\frac{I_{0T_k}}{I_{0T_h}}\right)^{iu}\exp\left[iu\sum_{i=h+1}^{k}A_{T_i}\right]\mathbb{E}_t^n\left[\prod_{i=h+1}^{k}X_{T_i}\left[iu\phi_{T_i}^y,iu\Psi_{T_i}^y\right]\right].$$

It can be calculated using a formula similar to the swap in §6.3. Here $0\leq t<T_{h+1}<T_k\leq T$.

coupon inflation options and make it easier to carry out investigation into the discrepancy between the zero-coupon inflation volatility level and the autocorrelation of the inflation index. Joint calibration of year-on-year and zero-coupon inflation volatilities could be carried out to infer the autocorrelation. Nevertheless, to calibrate all three quantities—year-on-year inflation volatilities, zero-coupon inflation volatilities and autocorrelation of the inflation index—might require having more parameters and even adding more driving factors to the QGY model.

When we derived the pricing formulas in the QGY model, we mainly used calculation methods from McCloud (2008). An alternative approach has been studied by Leippold and Wu (2002), who use Fourier transform to evaluate expectations of functions of quadratic Gaussian processes. As well as comparing our method to that method, it would be interesting to look at the infinitesimal generator of a quadratic Gaussian process (see Section 7.3 in Øksendal, 2003, for a definition). One could then apply the Feynman-Kac formula to get PDEs (partial differential equations) for the different payoffs. On could subsequently solve these PDEs with finite-difference methods to find the drift of the inflation index and to calibrate the model to market prices of swaps and options. Here the non-Markovian nature of the inflation index could be dealt with in PDE form through the adoption of a method from Vecer (2001), who looks at discrete path-dependent options.

To increase the curvature of the inflation smile in the QGY model, we could look at extending the model to include a cubic Gaussian term. Lamorgese et al. (2007, 2008) have looked at models of this type in the field of fluid mechanics. The QGY model is driven by a multi-factor Brownian motion that does not have any jumps. It is possible to add jumps to the QGY model to increase the curvature of the inflation model but this would require additional research. Sigurðsson (2008) models inflation using an extended class of affine jump diffusion models with jumps at fixed times. Time-changed Lévy processes have been applied to inflation by Kenyon (2008) and Andersen (2009). Charvet and Ticot (2011) use Andersen's method and assume that the terminal distribution of the year-on-year inflation ratio follows an inverse Gaussian distribution.

We see that inflation modelling is an active field of research and with the QGY model we have captured the complex properties of inflation markets. Further research can be carried out, where one can look at other properties of the model, make improvements on the model and use it in other applications.

Part III

Appendices

A From Hull-White to log-linear

This appendix follows the paper by Hughston (1994) that details how to go from the extended Vašíček (1977) short-rate equation by Hull and White (1990, 1993) to the log-linear Gaussian equation (4.2) in §4.1.

Reparameterisation

As we show in §3.1, we can go from the universal measure $\mathbb{P}$ to the risk-neutral measure in any currency. We let r_t be the short-rate process in some fixed currency so that under the corresponding risk-neutral measure $\mathbb{Q}^*$ we have

$$dr_t = \beta_t \left(\alpha_t - r_t\right) dt + \nu_t dW_t^*.$$

Here W^* is a $\mathbb{Q}^*$-Brownian motion and α_t, β_t and ν_t are deterministic functions. The solution to the SDE for r_t is of the form

$$r_t = \mu_t + \chi_t \int_0^t \sigma_s dW_s^*,$$

where μ_t, χ_t and σ_t are deterministic functions. Using Itō calculus, we get

$$dr_t = \dot{\mu}_t dt + \dot{\chi}_t \left(\int_0^t \sigma_s dW_s^*\right) dt + \chi_t \sigma_t dW_t^* = \frac{\dot{\chi}_t}{\chi_t}\left[\mu_t - \frac{\dot{\mu}_t \chi_t}{\dot{\chi}_t} - r_t\right] dt + \chi_t \sigma_t dW_t^*,$$

where $\dot{\mu}_t = d\mu_t/dt$ and $\dot{\chi}_t = d\chi_t/dt$. Comparing coefficients, we then have the relation

$$\begin{aligned} \alpha_t &= (\mu_t \dot{\chi}_t - \dot{\mu}_t \chi_t)/\,\dot{\chi}_t \\ \beta_t &= -\,\dot{\chi}_t/\,\chi_t \\ \nu_t &= \chi_t \sigma_t \end{aligned} \qquad \Longleftrightarrow \qquad \begin{aligned} \mu_t &= r_0 \chi_t + \chi_t \int_0^t \frac{\alpha_s \beta_s}{\chi_s} ds \\ \chi_t &= \exp\left[-\int_0^t \beta_s ds\right] \\ \sigma_t &= \nu_t/\,\chi_t \end{aligned} .$$

We have assumed, without loss of generality, that $\chi_0 = 1$. The expression for μ_t is obtained by rewriting the expression for α_t as

$$\frac{d}{dt}\left(\mu_t \frac{1}{\chi_t}\right) = -\frac{\alpha_t \dot{\chi}_t}{\chi_t^2}.$$

Discount bond

In the risk-neutral measure the inverse numeraire process D_t^* is defined by

$$\begin{aligned} D_t^* = \exp\left[-\int_0^t r_s ds\right] &= \exp\left[-\int_0^t \mu_s ds - \int_0^t \chi_s \left(\int_0^s \sigma_u dW_u^*\right) ds\right] \\ &= \exp\left[-\int_0^t \mu_s ds - \int_0^t \sigma_s \left(\phi_t - \phi_s\right) dW_s^*\right], \end{aligned}$$

where we have made use of the stochastic Fubini theorem. The term volatility ϕ_t is defined by

$$\phi_t = \int_0^t \chi_s ds.$$

We now have

$$\begin{aligned} D_t^* P_{tT} &= \mathbb{E}_t^*\left[D_T^*\right] \\ &= \exp\left[-\int_0^T \mu_s ds - \int_0^t \sigma_s \left(\phi_T - \phi_s\right) dW_s^*\right] \mathbb{E}_t^*\left[\exp\left[\int_t^T \sigma_s \left(\phi_T - \phi_s\right) dW_s^*\right]\right] \\ &= \exp\left[-\int_0^T \mu_s ds - \int_0^t \sigma_s \left(\phi_T - \phi_s\right) dW_s^* + \frac{1}{2}\int_t^T \sigma_s^2 \left(\phi_T - \phi_s\right)^2 ds\right]. \end{aligned}$$

Setting $t = 0$, we get $\int_0^T \mu_s ds = -\ln P_{0T} + \frac{1}{2}\int_0^T \sigma_s^2 \left(\phi_T - \phi_s\right)^2 ds$, which gives

$$\begin{aligned} D_t^* P_{tT} &= P_{0T} \exp\left[-\int_0^t \sigma_s \left(\phi_T - \phi_s\right) dW_s^* - \frac{1}{2}\int_0^t \sigma_s^2 \left(\phi_T - \phi_s\right)^2 ds\right], \\ D_t^* &= P_{0t} \exp\left[-\int_0^t \sigma_s \left(\phi_t - \phi_s\right) dW_s^* - \frac{1}{2}\int_0^t \sigma_s^2 \left(\phi_t - \phi_s\right)^2 ds\right], \end{aligned}$$

resulting in

$$\begin{aligned} P_{tT} &= \frac{P_{0T}}{P_{0t}} \exp\left[-\left(\phi_T - \phi_t\right)\int_0^t \sigma_s dW_s^* - \frac{1}{2}\int_0^t \sigma_s^2 \left(\phi_T - \phi_s\right)^2 ds + \frac{1}{2}\int_0^t \sigma_s^2 \left(\phi_t - \phi_s\right)^2 ds\right] \\ &= \frac{P_{0T}}{P_{0t}} \exp\left[-\left(\phi_T - \phi_t\right)\int_0^t \sigma_s dW_s^* - \frac{1}{2}\left(\phi_T^2 - \phi_t^2\right)\int_0^t \sigma_s^2 ds + \left(\phi_T - \phi_t\right)\int_0^t \sigma_s^2 \phi_s ds\right]. \end{aligned}$$

Forward measure

Let $\mathbb{Q}^U$ be the U-forward measure, which has an inverse numeraire process $D_t^U = P_{0U}/P_{tU}$ and conditional expectations denoted by $\mathbb{E}_t^U[\cdot]$. The value of a price process V_t at time zero is given by

$$V_0 = \mathbb{E}_0^*\left[D_t^* V_t\right] = P_{0U}\mathbb{E}_0^*\left[\frac{P_{tU}/P_{0U}}{1/D_t^*}\frac{V_t}{P_{tU}}\right] = P_{0U}\mathbb{E}_0^*\left[\frac{d\mathbb{Q}^U}{d\mathbb{Q}^*}\frac{V_t}{P_{tU}}\right] = \mathbb{E}_0^U\left[\frac{P_{0U}}{P_{tU}}V_t\right],$$

for $t \leq U$. The change-of-measure process from $\mathbb{Q}^*$ to $\mathbb{Q}^U$ is given by

$$\begin{aligned}\left.\frac{d\mathbb{Q}^U}{d\mathbb{Q}^*}\right|_{\mathcal{F}_t} &= \frac{P_{tU}/\,P_{0U}}{1/D_t^*} = \frac{1}{P_{0U}} D_t^* P_{tU} \\ &= \exp\left[-\int_0^t \sigma_s \left(\phi_U - \phi_s\right) dW_s^* - \frac{1}{2}\int_0^t \sigma_s^2 \left(\phi_U - \phi_s\right)^2 ds\right].\end{aligned}$$

From the Girsanov theorem, we now know that

$$W_t^U = W_t^* + \int_0^t \sigma_s \left(\phi_U - \phi_s\right) ds$$

is a $\mathbb{Q}^U$-Brownian motion. Let $0 \leq t \leq T \leq U$. We then get

$$\begin{aligned}D_t^U = \frac{P_{0U}}{P_{tU}} &= P_{0t} \exp\Bigg[\left(\phi_U - \phi_t\right)\int_0^t \sigma_s \left(dW_t^U + \sigma_s\left(\phi_U - \phi_s\right) ds\right) \\ &\qquad -\frac{1}{2}\left(\phi_U^2 - \phi_t^2\right)\int_0^t \sigma_s^2 ds + \left(\phi_U - \phi_t\right)\int_0^t \sigma_s^2 \phi_s ds\Bigg] \\ &= P_{0t}\exp\left[\left(\phi_U - \phi_t\right)\int_0^t \sigma_s dW_s^U - \frac{1}{2}\left(\phi_U^2 - \phi_t^2\right)\int_0^t \sigma_s^2 ds + \left(\phi_U - \phi_t\right)\int_0^t \sigma_s^2 \phi_U ds\right] \\ &= P_{0t}\exp\left[\left(\phi_U - \phi_t\right)\int_0^t \sigma_s dW_s^U - \frac{1}{2}\left(\phi_U - \phi_t\right)^2 \int_0^t \sigma_s^2 ds\right] \\ &= P_{0t}\exp\left[\left(\phi_U - \phi_t\right) x_t^U - \frac{1}{2}\left(\phi_U - \phi_t\right)^2 G_t\right].\end{aligned}$$

Here a time-changed Brownian motion is given by

$$x_t^U = \int_0^t \sigma_s dW_s^U,$$

which has mean zero under $\mathbb{Q}^U$ and variance

$$G_t = \int_0^t \sigma_s^2 ds.$$

The inverse numeraire D_t^U can be expressed as a function of $\phi_U - \phi_t$. Now, ϕ_U is a constant, so in §4.1 we have replaced $\phi_U - \phi_t$ with $-\phi_t$.

B Quadratic Gaussian cumulants

In this appendix, we calculate the moments of a quadratic Gaussian distribution—in particular, its skewness and excess kurtosis. Let x be a Gaussian distributed vector with a mean vector μ and a (symmetric) covariance matrix G under some measure $\mathbb{Q}$. Let $\mathbb{E}[\cdot]$ denote the expectations in the measure $\mathbb{Q}$. We can write $x = \mu + \sqrt{G}z$, with z a standard normally distributed vector. A quadratic-Gaussian process is given by

$$\begin{aligned}
& Q\left[x; \widetilde{\phi}, \widetilde{\Psi}\right] \\
:= & -\widetilde{\phi} \cdot x - \frac{1}{2} x \cdot \widetilde{\Psi} x = -\widetilde{\phi} \cdot \left(\mu + \sqrt{G}z\right) - \frac{1}{2}\left(\mu + \sqrt{G}z\right) \cdot \widetilde{\Psi}\left(\mu + \sqrt{G}z\right) \\
= & -\widetilde{\phi} \cdot \mu - \sqrt{G}\widetilde{\phi} \cdot z - \frac{1}{2}\mu \cdot \widetilde{\Psi}\mu - \frac{1}{2} z \cdot \sqrt{G}\widetilde{\Psi}\sqrt{G}z - \sqrt{G}\widetilde{\Psi}\mu \cdot z \\
= & -\left(\widetilde{\phi} + \frac{1}{2}\widetilde{\Psi}\mu\right) \cdot \mu - \sqrt{G}\left(\widetilde{\phi} + \widetilde{\Psi}\mu\right) \cdot z - \frac{1}{2} z \cdot \sqrt{G}\widetilde{\Psi}\sqrt{G}z. \qquad \text{(B.1)}
\end{aligned}$$

We define $\phi := \sqrt{G}\left(\widetilde{\phi} + \widetilde{\Psi}\mu\right)$ and $\Psi := \sqrt{G}\widetilde{\Psi}\sqrt{G}$. The central moments of $Q\left[x; \widetilde{\phi}, \widetilde{\Psi}\right]$, are the same as those of the quadratic-Gaussian variable given by

$$Q[z; \phi, \Psi] = -\phi \cdot z - \frac{1}{2} z \cdot \Psi z.$$

Using (C.5), we calculate the cumulant-generating function:

$$\begin{aligned}
g[t] := & \ln \mathbb{E}\left[\exp\left[tQ[z; \phi, \Psi]\right]\right] = \ln \mathbb{E}\left[\exp\left[-t\phi \cdot z - \frac{1}{2} z \cdot t\Psi z\right]\right] \\
= & \ln\left[\det[1 + t\Psi]^{-\frac{1}{2}} \exp\left[t\phi \cdot (1 + t\Psi)^{-1} t\phi\right]\right] \\
= & \frac{1}{2}\ln\left[\det[1 + t\Psi]^{-1}\right] + \frac{1}{2} t\phi \cdot (1 + t\Psi)^{-1} t\phi \\
= & \frac{1}{2}\ln\left[\det[1 + t\Psi]^{-1}\right] + \frac{1}{2}\phi \cdot (1 + t\Psi)^{-1} t^2\phi.
\end{aligned}$$

The cumulants κ_n of $Q[z; \phi, \Psi]$ are given by the derivative of $g[t]$ at $t = 0$:

$$\kappa_n = \mathbb{E}\left[\left(Q[z; \phi, \Psi] - \mathbb{E}[Q[z; \phi, \Psi]]\right)^n\right] = g^{(n)}[0].$$

The first cumulant κ_1 is the expected value, the second cumulant κ_2 is the variance, $\kappa_3 \Big/ \kappa_2^{3/2}$ is the skewness and $\kappa_4 \big/ \kappa_2^2$ is the excess kurtosis of $Q[z;\phi,\Psi]$. We proceed to look at the following matrix and its derivatives:

$$
\begin{aligned}
M[t] :=&\ (1+t\Psi)^{-1} \\
M'[t] =&\ -(1+t\Psi)^{-1}\left(\frac{d}{dt}(1+t\Psi)\right)(1+t\Psi)^{-1} = -M[t]\,\Psi M[t] \\
M''[t] =&\ -M'[t]\,\Psi M[t] - M[t]\,\Psi M'[t] = 2M[t]\,\Psi M[t]\,\Psi M[t] \\
M^{(3)}[t] =&\ 2\left(M'[t]\,\Psi M[t]\,\Psi M[t] + M[t]\,\Psi M'[t]\,\Psi M[t] + M[t]\,\Psi M[t]\,\Psi M'[t]\right) \\
=&\ -6M[t]\,\Psi M[t]\,\Psi M[t]\,\Psi M[t] \\
M^{(4)}[t] =&\ 24M[t]\,\Psi M[t]\,\Psi M[t]\,\Psi M[t]\,\Psi M[t]\,.
\end{aligned}
$$

The moments of the matrix $M[t]$ are then given by

$$
\begin{aligned}
M^{(n)}[t] =&\ (-1)^n\, n!M[t]\prod_{i=1}^{n}(\Psi M[t]) \\
M^{(n)}[0] =&\ (-1)^n\, n!\prod_{i=1}^{n}\Psi.
\end{aligned}
$$

For the first term in the cumulant-generating function $g[t]$, we use Jacobi's formula to differentiate the determinant of a matrix:

$$
\begin{aligned}
a[t] :=&\ \ln\left[\det[1+t\Psi]^{-1}\right] = \ln[\det[M[t]]] \\
a'[t] =&\ \frac{1}{\det[M[t]]}\frac{d}{dt}\det[M[t]] = \frac{1}{\det[M[t]]}\det[M[t]]\operatorname{tr}\left[M^{-1}[t]\,M'[t]\right] \\
=&\ -\operatorname{tr}\left[M^{-1}[t]\,M[t]\,\Psi M[t]\right] = -\operatorname{tr}[\Psi M[t]] = -\operatorname{tr}[M[t]\,\Psi]\,.
\end{aligned}
$$

The n-th derivative of the term $a[t]$ for $n \geq 0$ is then given by

$$
a^{(n)}[t] = -\operatorname{tr}\left[M^{(n-1)}[t]\,\Psi\right].
$$

From the second term in the cumulant-generating function $g[t]$, we calculate the following:

$$
\begin{aligned}
b[t] :=&\ (1+t\Psi)^{-1}t^2 = M[t]\,t^2 \\
b'[t] =&\ M'[t]\,t^2 + 2M[t]\,t \\
b''[t] =&\ M''[t]\,t^2 + 2M'[t]\,t + 2M'[t]\,t + 2M[t] = M''[t]\,t^2 + 4M'[t]\,t + 2M[t] \\
b^{(3)}[t] =&\ M^{(3)}[t]\,t^2 + 2M''[t]\,t + 4M''[t]\,t + 4M'[t] + 2M'[t] \\
=&\ M^{(3)}[t]\,t^2 + 6M''[t]\,t + 6M'[t] \\
b^{(4)}[t] =&\ M^{(4)}[t]\,t^2 + 2M^{(3)}[t]\,t + 6M^{(3)}[t]\,t + 6M''[t] + 6M''[t] \\
=&\ M^{(4)}[t]\,t^2 + 8M^{(3)}[t]\,t + 12M''[t]\,.
\end{aligned}
$$

The n-th derivative of the term $b[t]$ for $n \geq 0$ is then given by

$$b^{(n)}[t] = M^{(n)}[t]\, t^2 + 2nM^{(n-1)}[t]\, t + n(n-1)\, M^{(n-2)}[t]\,.$$

By combining the expressions for $a[t]$ and $b[t]$, the cumulant-generating function is

$$g[t] = \frac{1}{2}\ln[\det[1 + t\Psi]] + \frac{1}{2}\phi\cdot(1 + t\Psi)^{-1}\, t^2\phi = \frac{1}{2}a[t] + \frac{1}{2}\phi\cdot b[t]\,\phi,$$

with the n-th derivative given by

$$\begin{aligned} g^{(n)}[t] &= \frac{1}{2}a^{(n)}[t] + \frac{1}{2}\phi\cdot b^{(n)}[t]\,\phi \\ &= -\frac{1}{2}\operatorname{tr}\left[M^{(n-1)}[t]\,\Psi\right] + \frac{1}{2}\phi\left(M^{(n)}[t]\, t^2 + 2nM^{(n-1)}[t]\, t + n(n-1)\, M^{(n-2)}[t]\right)\phi. \end{aligned}$$

The cumulants are

$$\begin{aligned} \kappa_n = g^{(n)}[0] = -\frac{1}{2}(-1)^{n-1}(n-1)!\operatorname{tr}\left[\prod_{i=1}^{n}\Psi\right] \\ + \frac{1}{2}n(n-1)(-1)^{n-2}(n-2)!\phi\left(\prod_{i=1}^{n-2}\Psi\right)\phi, \end{aligned}$$

so the expected value κ_1, the variance κ_2, the skewness $\kappa_3\big/\kappa_2^{3/2}$ and the excess kurtosis $\kappa_4\big/\kappa_2^2$ of $Q[z;\phi,\Psi]$ are deduced from

$$\begin{aligned} \kappa_1 &= -\frac{1}{2}\operatorname{tr}[\Psi]\,, \\ \kappa_2 &= \frac{1}{2}\operatorname{tr}\left[\Psi^2\right] + \phi\cdot\phi, \\ \kappa_3 &= -\operatorname{tr}\left[\Psi^3\right] - 3\phi\cdot\Psi\phi, \\ \kappa_4 &= 3\operatorname{tr}\left[\Psi^4\right] + 12\phi\cdot\Psi^2\phi. \end{aligned}$$

From (B.1), and noting that $\phi = \sqrt{G}\left(\widetilde{\phi} + \widetilde{\Psi}\mu\right)$ and $\Psi = \sqrt{G}\widetilde{\Psi}\sqrt{G}$, the first four cumulants of $Q\left[x;\widetilde{\phi},\widetilde{\Psi}\right]$ are

$$\begin{aligned} \kappa_1 &= -\left(\widetilde{\phi} + \frac{1}{2}\widetilde{\Psi}\mu\right)\cdot\mu - \frac{1}{2}\operatorname{tr}\left[\widetilde{\Psi}G\right], \\ \kappa_2 &= \frac{1}{2}\operatorname{tr}\left[\left(\widetilde{\Psi}G\right)^2\right] + \left(\widetilde{\phi} + \widetilde{\Psi}\mu\right)\cdot G\left(\widetilde{\phi} + \widetilde{\Psi}\mu\right), \\ \kappa_3 &= -\operatorname{tr}\left[\left(\widetilde{\Psi}G\right)^3\right] - 3\left(\widetilde{\phi} + \widetilde{\Psi}\mu\right)\cdot G\widetilde{\Psi}G\left(\widetilde{\phi} + \widetilde{\Psi}\mu\right), \\ \kappa_4 &= 3\operatorname{tr}\left[\left(\widetilde{\Psi}G\right)^4\right] + 12\left(\widetilde{\phi} + \widetilde{\Psi}\mu\right)\cdot G\left(\widetilde{\Psi}G\right)^2\left(\widetilde{\phi} + \widetilde{\Psi}\mu\right). \end{aligned}$$

C Quadratic Gaussian expectations

In this appendix, we will see how to calculate the expectations of log-quadratic Gaussian processes, following calculations by McCloud (2008). These calculations are used in Chapter 6 to derive pricing formulas in the QGY model. Let x be a Gaussian distributed vector with a mean vector μ and a (symmetric) covariance matrix G under some measure $\mathbb{Q}$. Let $\mathbb{E}[\cdot]$ denote the expectations in the measure $\mathbb{Q}$. We can write $x = \mu + \sqrt{G}z$, with z a standard normally distributed vector. Here $\sqrt{G}\sqrt{G}^\intercal = G$ where $\sqrt{G}^\intercal$ represents the transpose of the matrix $\sqrt{G}$. For some function $f : \mathbb{R}^{dim} \to \mathbb{R}$ for $dim > 0$, we then have

$$\begin{aligned}
&\mathbb{E}\left[\exp\left[-\phi \cdot x - \tfrac{1}{2}x \cdot \Psi x\right] f[x]\right] \qquad\qquad \text{(C.1)}\\
&= (2\pi)^{-\frac{dim}{2}} \int \exp\Big[-\tfrac{1}{2}z \cdot z - \phi \cdot \left(\mu + \sqrt{G}z\right) \\
&\qquad -\tfrac{1}{2}\left(\mu + \sqrt{G}z\right) \cdot \Psi \left(\mu + \sqrt{G}z\right)\Big] f\left[\mu + \sqrt{G}z\right] dz,
\end{aligned}$$

where ϕ is a dim-dimensional vector and Ψ is a dim-dimensional symmetric matrix. Now we look at the terms in the exponents (multiplied through by -2):

$$\begin{aligned}
&z \cdot z + 2\phi \cdot \left(\mu + \sqrt{G}z\right) + \left(\mu + \sqrt{G}z\right) \cdot \Psi \left(\mu + \sqrt{G}z\right) \\
&\overset{\text{expand}}{=} z \cdot z + 2\phi \cdot \mu + 2\sqrt{G}^\intercal \phi \cdot z + \mu \cdot \Psi\mu + 2\left(\sqrt{G}^\intercal \Psi \mu\right) \cdot z + z \cdot \left(\sqrt{G}^\intercal \Psi \sqrt{G}\right) z \\
&\overset{\text{group}}{=} z \cdot \left(1 + \sqrt{G}^\intercal \Psi \sqrt{G}\right) z + 2\left(\sqrt{G}^\intercal (\phi + \Psi\mu)\right) \cdot z + (2\phi \cdot \mu + \mu \cdot \Psi\mu) \\
&\overset{z \to \widetilde{z}}{=} \widetilde{z} \cdot \widetilde{z} - (\phi + \Psi\mu) \cdot \sqrt{G}\left(1 + \sqrt{G}^\intercal \Psi \sqrt{G}\right)^{-1} \sqrt{G}^\intercal (\phi + \Psi\mu) + (2\phi \cdot \mu + \mu \cdot \Psi\mu). \quad \text{(C.2)}
\end{aligned}$$

Here the relationship between z and $\widetilde{z}$ is obtained through completing the square and is given by

$$\begin{aligned}
\widetilde{z} &:= \left(1 + \sqrt{G}^\intercal \Psi \sqrt{G}\right)^{\frac{1}{2}} z + \left(1 + \sqrt{G}^\intercal \Psi \sqrt{G}\right)^{-\frac{1}{2}*} \sqrt{G}^\intercal (\phi + \Psi\mu), \\
z &= \left(1 + \sqrt{G}^\intercal \Psi \sqrt{G}\right)^{-\frac{1}{2}} \left(\widetilde{z} - \left(1 + \sqrt{G}^\intercal \Psi \sqrt{G}\right)^{-\frac{1}{2}*} \sqrt{G}^\intercal (\phi + \Psi\mu)\right).
\end{aligned}$$

Noting that

$$\frac{d\widetilde{z}}{dz} = \det\left[1 + \sqrt{G}^{\intercal}\Psi\sqrt{G}\right]^{\frac{1}{2}}, \tag{C.3}$$

we now change from variable z to $\widetilde{z}$ in (C.1). By using (C.2), requiring that $1+\sqrt{G}^{\intercal}\Psi\sqrt{G}$ is positive definite, we get

$$\begin{aligned}
&\mathbb{E}\left[\exp\left[-\phi\cdot x - \tfrac{1}{2}x\cdot\Psi x\right] f\left[x\right]\right] \\
&= \det\left[1 + \sqrt{G}^{\intercal}\Psi\sqrt{G}\right]^{-\frac{1}{2}} \\
&\quad\times \exp\left[\tfrac{1}{2}\sqrt{G}^{\intercal}\left(\phi + \Psi\mu\right)\cdot\left(1 + \sqrt{G}^{\intercal}\Psi\sqrt{G}\right)^{-1}\sqrt{G}^{\intercal}\left(\phi + \Psi\mu\right) - \phi\cdot\mu - \tfrac{1}{2}\mu\cdot\Psi\mu\right] \\
&\quad\times \mathbb{E}\left[f\left[\mu + \sqrt{G}\left(1 + \sqrt{G}^{\intercal}\Psi\sqrt{G}\right)^{-\frac{1}{2}}\sqrt{G}^{-1}\left(x - \mu\right)\right.\right. \\
&\qquad\qquad \left.\left. - \sqrt{G}\left(1 + \sqrt{G}^{\intercal}\Psi\sqrt{G}\right)^{-1}\sqrt{G}^{\intercal}\left(\phi + \Psi\mu\right)\right]\right].
\end{aligned} \tag{C.4}$$

Now we assume that G is invertible. For some square matrices A and B of the same dimensions, it holds that $\det\left(1 + AB\right) = \det\left(1 + BA\right)$. Also, $\left(AB\right)^{\intercal} = B^{\intercal}A^{\intercal}$ where $A^{\intercal}$ is the transpose of A, and if A and B are invertible then $\left(AB\right)^{-1} = B^{-1}A^{-1}$. We then have

$$\begin{aligned}
&\mathbb{E}\left[\exp\left[-\phi\cdot x - \tfrac{1}{2}x\cdot\Psi x\right]\right] \\
&= \det\left[1 + \Psi G\right]^{-\frac{1}{2}}\exp\left[\tfrac{1}{2}\left(\phi + \Psi\mu\right)\cdot G\left(1 + \Psi G\right)^{-1}\left(\phi + \Psi\mu\right) - \phi\cdot\mu - \tfrac{1}{2}\mu\cdot\Psi\mu\right] \\
&= \det\left[1 + \Psi G\right]^{-\frac{1}{2}}\exp\left[\tfrac{1}{2}\left(\phi + \Psi\mu\right)\cdot\left(G^{-1} + \Psi\right)^{-1}\left(\phi + \Psi\mu\right) - \phi\cdot\mu - \tfrac{1}{2}\mu\cdot\Psi\mu\right] \\
&= \det\left[1 + \Psi G\right]^{-\frac{1}{2}}\exp\left[-\phi\cdot\mu + \phi\cdot\left(G^{-1} + \Psi\right)^{-1}\Psi\mu\right. \\
&\qquad\qquad \left. - \tfrac{1}{2}\mu\cdot\Psi\mu + \tfrac{1}{2}\Psi\mu\cdot\left(G^{-1} + \Psi\right)^{-1}\Psi\mu + \tfrac{1}{2}\phi\cdot\left(G^{-1} + \Psi\right)^{-1}\phi\right].
\end{aligned}$$

Looking at the different terms in the exponent, we now get

$$\begin{aligned}
-\phi\cdot\mu + \phi\cdot\left(G^{-1} + \Psi\right)^{-1}\Psi\mu &= -\phi\cdot\left(G^{-1} + \Psi\right)^{-1}\left[\left(G^{-1} + \Psi\right) - \Psi\right]\mu \\
&= -\phi\cdot\left(1 + G\Psi\right)^{-1}\mu = -\left(1 + \Psi G\right)^{-1}\phi\cdot\mu, \\
-\mu\cdot\Psi\mu + \Psi\mu\cdot\left(G^{-1} + \Psi\right)^{-1}\Psi\mu &= -\mu\cdot\left[\left(G^{-1} + \Psi\right) - \Psi\right]\left(G^{-1} + \Psi\right)^{-1}\Psi\mu \\
&= -\mu\cdot\left(1 + \Psi G\right)^{-1}\Psi\mu,
\end{aligned}$$

so we get

$$\begin{aligned}
&\mathbb{E}\left[\exp\left[-\phi\cdot x - \tfrac{1}{2}x\cdot\Psi x\right]\right] \\
&= \det\left[1 + \Psi G\right]^{-\frac{1}{2}}\exp\left[\tfrac{1}{2}\phi\cdot G\left(1 + \Psi G\right)^{-1}\phi - \left(1 + \Psi G\right)^{-1}\phi\cdot\mu - \tfrac{1}{2}\mu\cdot\left(1 + \Psi G\right)^{-1}\Psi\mu\right].
\end{aligned} \tag{C.5}$$

We can also simplify the parameter of the function f,

$$
\begin{aligned}
&\mu+\sqrt{G}z \\
&\overset{z\to\widetilde{z}}{=}\mu+\sqrt{G}\left(1+\sqrt{G}^{\intercal}\Psi\sqrt{G}\right)^{-\frac{1}{2}}\widetilde{z}-\sqrt{G}\left(1+\sqrt{G}^{\intercal}\Psi\sqrt{G}\right)^{-1}\sqrt{G}^{\intercal}\left(\phi+\Psi\mu\right) \\
&\overset{\text{rearr.}}{=}-\sqrt{G}\left(1+\sqrt{G}^{\intercal}\Psi\sqrt{G}\right)^{-1}\sqrt{G}^{\intercal}\phi+\left(1-\sqrt{G}\left(1+\sqrt{G}^{\intercal}\Psi\sqrt{G}\right)^{-1}\sqrt{G}^{\intercal}\Psi\right)\mu \\
&\qquad+\sqrt{G}\left(1+\sqrt{G}^{\intercal}\Psi\sqrt{G}\right)^{-\frac{1}{2}}\widetilde{z} \\
&\overset{\text{simpl.}}{=}-\sqrt{G}\left(1+\sqrt{G}^{\intercal}\Psi\sqrt{G}\right)^{-1}\sqrt{G}^{\intercal}\phi+\sqrt{G}\left(1+\sqrt{G}^{\intercal}\Psi\sqrt{G}\right)^{-1}\sqrt{G}^{-1}\mu \\
&\qquad+\sqrt{G}\left(1+\sqrt{G}^{\intercal}\Psi\sqrt{G}\right)^{-\frac{1}{2}}\widetilde{z} \\
&\overset{\text{rearr.}}{=}\sqrt{G}\left(1+\sqrt{G}^{\intercal}\Psi\sqrt{G}\right)^{-1}\left(\sqrt{G}^{-1}\mu-\sqrt{G}^{\intercal}\phi\right)+\sqrt{G}\left(1+\sqrt{G}^{\intercal}\Psi\sqrt{G}\right)^{-\frac{1}{2}}\widetilde{z} \\
&\overset{\text{simpl.}}{=}G\left(1+\Psi G\right)^{-1}\left(G^{-1}\mu-\phi\right)+\sqrt{G}\left(1+\sqrt{G}^{\intercal}\Psi\sqrt{G}\right)^{-\frac{1}{2}}\widetilde{z}=\widetilde{\mu}+\sqrt{\widetilde{G}}\widetilde{z},
\end{aligned}
\tag{C.6}
$$

with

$$
\begin{aligned}
\widetilde{\mu}&:=\widetilde{G}\left(G^{-1}\mu-\phi\right), \\
\sqrt{\widetilde{G}}&:=\sqrt{G}\left(1+\sqrt{G}^{\intercal}\Psi\sqrt{G}\right)^{-\frac{1}{2}}, \\
\widetilde{G}&:=\sqrt{\widetilde{G}}\sqrt{\widetilde{G}}^{\intercal}=G\left(1+\Psi G\right)^{-1},
\end{aligned}
$$

where we note that $\widetilde{G}$ is symmetric and invertible.[1] From (C.4), (C.5) and (C.6) we can therefore write

$$
\begin{aligned}
&\mathbb{E}\left[\exp\left[-\phi\cdot x-\tfrac{1}{2}x\cdot\Psi x\right]f\left[x\right]\right] \\
&=\mathbb{E}\left[\exp\left[-\phi\cdot x-\tfrac{1}{2}x\cdot\Psi x\right]\right](2\pi)^{-\frac{dim}{2}}\int\exp\left[-\tfrac{1}{2}\widetilde{z}\cdot\widetilde{z}\right]f\left[\widetilde{\mu}+\sqrt{\widetilde{G}}\widetilde{z}\right]d\widetilde{z} \\
&=\mathbb{E}\left[\exp\left[-\phi\cdot x-\tfrac{1}{2}x\cdot\Psi x\right]\right]\mathbb{E}\left[f\left[\widetilde{\mu}+\sqrt{\widetilde{G}}\sqrt{G}^{-1}\left(x-\mu\right)\right]\right] \\
&=\mathbb{E}\left[\exp\left[-\phi\cdot x-\tfrac{1}{2}x\cdot\Psi x\right]\right]\widetilde{\mathbb{E}}\left[f\left[\widetilde{x}\right]\right].
\end{aligned}
\tag{C.7}
$$

The notation $\widetilde{\mathbb{E}}\left[\cdot\right]$ denotes the expectations under a measure where $\widetilde{x}$ is a Gaussian process with a mean vector $\widetilde{\mu}$ and a variance matrix $\widetilde{G}$.

[1] As G is invertible, the inverse of $\widetilde{G}$ is $\widetilde{G}^{-1}=\left(1+\Psi' G\right)G^{-1}=\left(G^{-1}+\Psi'\right)$.
We may therefore write $\widetilde{G}=\left(G^{-1}+\Psi'\right)^{-1}$. For a matrix A, we have $\left(A^{-1}\right)^{\intercal}=\left(A^{\intercal}\right)^{-1}$—that is, we can exchange the transpose operation with the inverse operation. As G and Ψ' are symmetric, the transpose of $\widetilde{G}$ is given by

$$
\widetilde{G}^{\intercal}=\left(\left(G^{-1}\right)^{\intercal}+\Psi'^{\intercal}\right)^{-1}=\left(\left(G^{\intercal}\right)^{-1}+\Psi'^{\intercal}\right)^{-1}=\left(G^{-1}+\Psi'\right)^{-1}=\widetilde{G},
$$

so $\widetilde{G}$ is symmetric.

D Normalised expectations

Here we simplify the expectations of a normalised log-quadratic normal distribution function by following calculations done by McCloud (2008). We use the result to get the nominal discount bond in the QGY model in §6.1. Let x_t be a Gaussian distributed vector with a mean vector 0 and a (symmetric) covariance matrix G_t. For some $t, T \geq 0$ with $t < T$, we define $G_{tT} = G_T - G_t$. We note that $x_T | \, x_t$ is a Gaussian distributed vector with a mean vector x_t and a covariance matrix G_{tT}. We define the following normalised quadratic Gaussian function D_t for any $t \geq 0$,

$$D_t\left[\phi, \Psi\right] := \frac{\exp\left[-\phi \cdot x_t - \frac{1}{2} x_t \cdot \Psi x_t\right]}{\mathbb{E}_0\left[\exp\left[-\phi \cdot x_t - \frac{1}{2} x_t \cdot \Psi x_t\right]\right]}.$$

Here, ϕ is a deterministic vector and Ψ is a deterministic symmetric matrix. For some $t, T \geq 0$ with $t < T$, we assume that $1 + \Psi G_T$ and $1 + \Psi G_t$ are positive definite and we deduce from (C.5) that

$$\begin{aligned}
&\mathbb{E}_t\left[D_T\left[\phi, \Psi\right]\right] = \frac{\mathbb{E}_t\left[\exp\left[-\phi \cdot x_T - \frac{1}{2} x_T \cdot \Psi x_T\right]\right]}{\mathbb{E}_0\left[\exp\left[-\phi \cdot x_T - \frac{1}{2} x_T \cdot \Psi x_T\right]\right]} \\
&= \frac{\det\left[1 + \Psi G_{tT}\right]^{-\frac{1}{2}}}{\det\left[1 + \Psi G_T\right]^{-\frac{1}{2}}} \frac{\exp\left[\frac{1}{2}\phi \cdot G_{tT}\left(1 + \Psi G_{tT}\right)^{-1}\phi\right]}{\exp\left[\frac{1}{2}\phi \cdot G_T\left(1 + \Psi G_T\right)^{-1}\phi\right]} \\
&\quad \times \exp\left[-\left(1 + \Psi G_{tT}\right)^{-1}\phi \cdot x_t - \frac{1}{2} x_t \cdot \left(1 + \Psi G_{tT}\right)^{-1}\Psi x_t\right] \\
&= \frac{\exp\left[-\left(1 + \Psi G_{tT}\right)^{-1}\phi \cdot x_t - \frac{1}{2} x_t \cdot \left(1 + \Psi G_{tT}\right)^{-1}\Psi x_t\right]}{\det\left[\left(1 + \Psi G_T\right)\left(1 + \Psi G_{tT}\right)^{-1}\right]^{-\frac{1}{2}} \exp\left[\frac{1}{2}\phi \cdot \left(\left(G_T^{-1} + \Psi\right)^{-1} - \left(G_{tT}^{-1} + \Psi\right)^{-1}\right)\phi\right]}.
\end{aligned} \tag{D.1}$$

Looking at the determinant we now calculate

$$\begin{aligned}
&\det\left[\left(1 + \Psi G_T\right)\left(1 + \Psi G_{tT}\right)^{-1}\right] = \det\left[\left(1 + \Psi G_{tT} + \Psi G_t\right)\left(1 + \Psi G_{tT}\right)^{-1}\right] \\
&= \det\left[1 + \Psi G_t\left(1 + \Psi G_{tT}\right)^{-1}\right] = \det\left[1 + \left(1 + \Psi G_{tT}\right)^{-1}\Psi G_t\right],
\end{aligned}$$

where, in the final step, we have used Sylvester's determinant theorem. For other terms, we have

$$\begin{aligned}
&\phi \cdot \left(\left(G_T^{-1} + \Psi\right)^{-1} - \left(G_{tT}^{-1} + \Psi\right)^{-1}\right) \phi \\
&= \phi \cdot \left(G_{tT}^{-1} + \Psi\right)^{-1} \left[\left(G_{tT}^{-1} + \Psi\right) - \left(G_T^{-1} + \Psi\right)\right] \left(G_T^{-1} + \Psi\right)^{-1} \phi \\
&= \phi \cdot \left(G_{tT}^{-1} + \Psi\right)^{-1} G_{tT}^{-1} \left(G_T - G_{tT}\right) G_T^{-1} \left(G_T^{-1} + \Psi\right)^{-1} \phi \\
&= \phi \cdot \left(1 + G_{tT}\Psi\right)^{-1} G_t \left(1 + \Psi G_T\right)^{-1} \phi \\
&= \left(1 + \Psi G_{tT}\right)^{-1} \phi \cdot G_t \left(1 + \Psi G_{tT} + \Psi G_t\right)^{-1} \phi \\
&= \left(1 + \Psi G_{tT}\right)^{-1} \phi \cdot G_t \left[1 + \left(1 + \Psi G_{tT}\right)^{-1} \Psi G_t\right]^{-1} \left(1 + \Psi G_{tT}\right)^{-1} \phi.
\end{aligned}$$

Expression (D.1) now becomes

$$\begin{aligned}
&\mathbb{E}_t\left[D_T\left[\phi, \Psi\right]\right] \\
&= \frac{\exp\left[-\left(1 + \Psi G_{tT}\right)^{-1} \phi \cdot x_t - \frac{1}{2} x_t \cdot \left(1 + \Psi G_{tT}\right)^{-1} \Psi x_t\right]}{\det\left[1 + \left(1 + \Psi G_{tT}\right)^{-1} \Psi G_t\right]^{-\frac{1}{2}} \mathrm{e}^{\frac{1}{2}\left(1 + \Psi G_{tT}\right)^{-1} \phi \cdot G_t \left[1 + \left(1 + \Psi G_{tT}\right)^{-1} \Psi G_t\right]^{-1} \left(1 + \Psi G_{tT}\right)^{-1} \phi}} \\
&= \frac{\exp\left[-\left(1 + \Psi G_{tT}\right)^{-1} \phi \cdot x_t - \frac{1}{2} x_t \cdot \left(1 + \Psi G_{tT}\right)^{-1} \Psi x_t\right]}{\mathbb{E}_0\left[\exp\left[-\left(1 + \Psi G_{tT}\right)^{-1} \phi \cdot x_t - \frac{1}{2} x_t \cdot \left(1 + \Psi G_{tT}\right)^{-1} \Psi x_t\right]\right]} \\
&= D_t\left[\left(1 + \Psi G_{tT}\right)^{-1} \phi, \left(1 + \Psi G_{tT}\right)^{-1} \Psi\right]. \qquad \text{(D.2)}
\end{aligned}$$

The expectations of a normalised quadratic Gaussian function is therefore given by the same type of function but with a transformation of the parameters (ϕ, Ψ) with the matrix $\left(1 + \Psi G_{tT}\right)^{-1}$, that is,

$$(\phi, \Psi) \longmapsto \left(\left(1 + \Psi G_{tT}\right)^{-1} \phi, \left(1 + \Psi G_{tT}\right)^{-1} \Psi\right).$$

E Gaussian integration over conics

In the formula for the price of a caplet in (6.11) we need to integrate a multi-factor standard Gaussian density over a domain D with its boundary defined by a quadratic form,

$$\mathcal{N}\left[\mathcal{D}\right] := (2\pi)^{-\frac{dim}{2}} \int_{\mathcal{D}} \exp\left[-\tfrac{1}{2} z \cdot z\right] dz,$$
$$\mathcal{D} := \left\{ z \in \mathbb{R}^{dim} : \phi \cdot z + z \cdot \Psi z + F < 0 \right\}.$$

Here ϕ is a vector, Ψ is a matrix and F is a scalar. We look at the two-dimensional case ($dim = 2$) where we can transform the variables into polar coordinates, which helps identify the boundary of the domain of integration. The domain can be visualised with a conic, where the domain is either inside or outside a conic section. Examples of three conics are shown in Figure E.1.

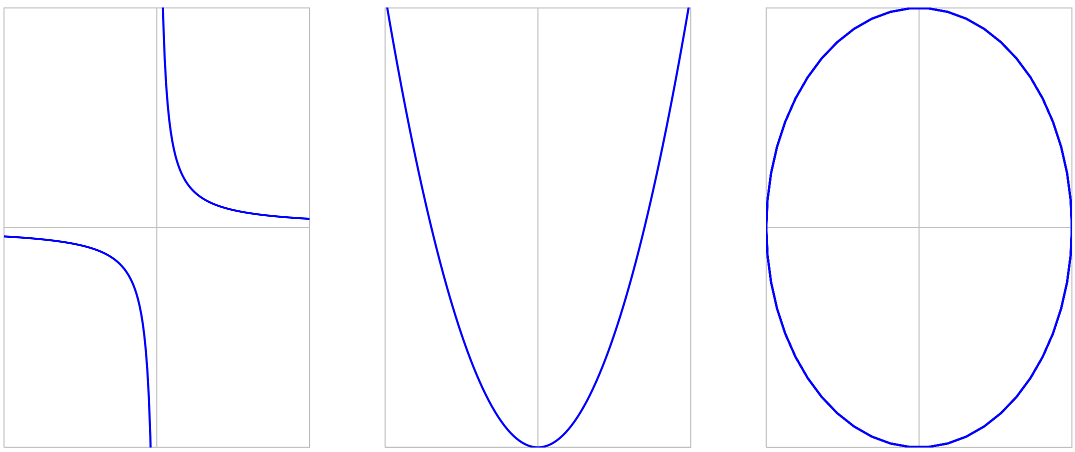

Figure E.1: The boundary of the quadratic domain in two dimensions, $\mathcal{D} = \left\{ z = (x, y) \in \mathbb{R}^{dim} : \phi \cdot z + z \cdot \Psi z + F < 0 \right\}$, can be visualised with a hyperbola (left), a parabola (middle) or an ellipse (right).

Transformation to polar coordinates

Let $A, B, C, D, E, F \in \mathbb{R}$ be some given constants and let $x, y \in \mathbb{R}$ be some variables. We then define a quadratic function in two variables with

$$P[x,y] := Ax^2 + Bxy + Cy^2 + Dx + Ey + F.$$

A quadratic domain in two dimensions is then given by

$$\mathcal{D} = \left\{(x,y) \in \mathbb{R}^2 : P[x,y] < 0\right\}.$$

Changing to polar coordinates, we let $\theta \in [0, 2\pi)$ and $r \in [0, +\infty)$ represent the angle and the length respectively of the vector (x, y). We then have the relations

$$\left.\begin{aligned} x &= r\cos\theta \\ y &= r\sin\theta \end{aligned}\right\} \implies \begin{cases} r^2 = x^2 + y^2 \\ dxdy = rdrd\theta. \end{cases}$$

Inserting this into $P[x,y]$ then gives a function in polar coordinates

$$P[\theta, r] = Q[\theta] r^2 + L[\theta] r + F.$$

Given a θ, this is a quadratic function in r with coefficients

$$\begin{aligned} Q[\theta] &:= A\cos^2\theta + B\sin\theta\cos\theta + C\sin^2\theta, \\ L[\theta] &:= D\cos\theta + E\sin\theta \end{aligned}$$

and discriminant

$$\begin{aligned} M[\theta] &:= L^2[\theta] - 4FQ[\theta] && \text{(E.1)} \\ &= \left(D^2 - 4AF\right)\cos^2\theta + (2DE - 4BF)\sin\theta\cos\theta + \left(E^2 - 4CF\right)\sin^2\theta. && \text{(E.2)} \end{aligned}$$

In Figure E.2, we plot $P[x,y]$, $P[\theta,r]$, $Q[\theta]$, $L[\theta]$ and $M[\theta]$. We set the parameters to $A = 0.001$, $B = 0.01$, $C = -0.00007$, $D = -0.01$, $E = 0.000005$ and $F = -0.03$, which are representative values for the calibrated reduced three-factor QGY model described in Chapters 7 and 8. For these values of parameters the solution to the equation $P[x,y] = 0$ is an hyperbola.

Roots of trigonometric equations

To know for which (θ, r) we have $P[\theta, r] < 0$, we start by identifying the boundary—that is, $P[\theta, r] = 0$. For a given θ, depending on the sign of $M[\theta]$, $Q[\theta]$, $L[\theta]$ and F, the

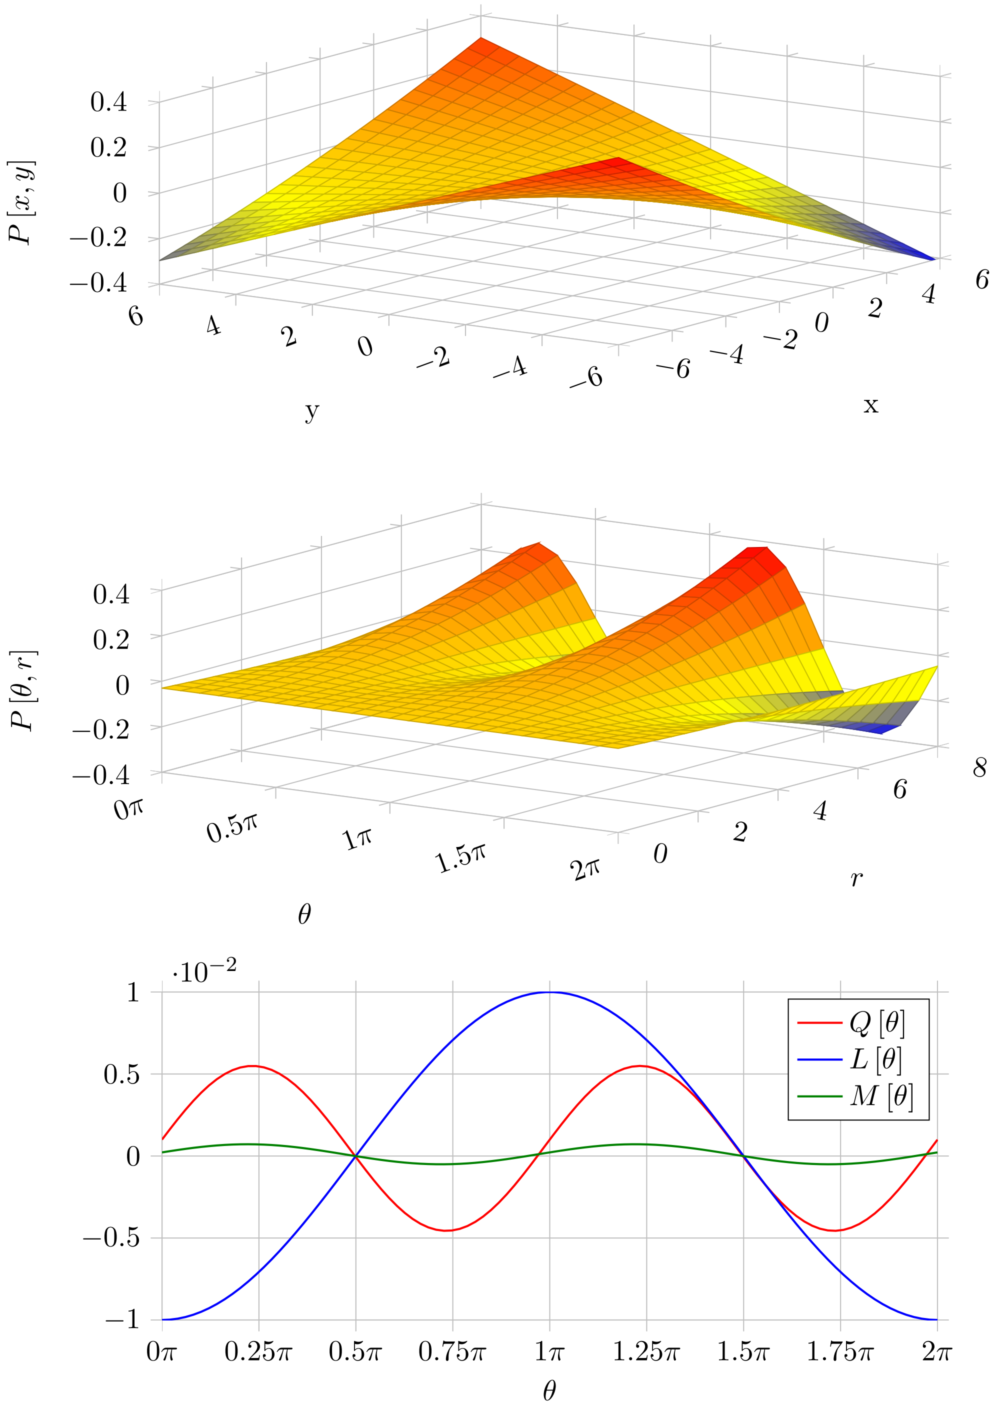

Figure E.2: Two-dimensional quadratic function $P[x,y]$ (top), corresponding quadratic function $P[\theta,r]$ in polar coordinates (middle) with coefficients $Q[\theta]$ and $L[\theta]$ and discriminant $M[\theta]$ (bottom). The values of the parameters are $A = 0.001$, $B = 0.01$, $C = -0.00007$, $D = -0.01$, $E = 0.000005$ and $F = -0.03$.

equation $P[\theta, r] = 0$ has one, two or no solutions. We therefore start determining for which θ we have $M[\theta] = 0$, $Q[\theta] = 0$ or $L[\theta] = 0$.

In Tables E.1 and E.2, we find the roots of $Q[\theta]$ and $L[\theta]$ for different values of A, B, C, D, E. From (E.2) we see that $M[\theta]$ is of the same form as $Q[\theta]$ but with different parameters, so we can use Table E.1 by replacing A, B, C with $D' = D^2 - 4AF$, $B' = 2DE - 4BF$, $C' = E^2 - 4CF$ respectively.

A	B	C	Equation	Roots
$= 0$	$= 0$	$= 0$	$0 = 0$	$\forall\, \theta$
$\neq 0$	$= 0$	$= 0$	$\cos^2 \theta = 0$	$\theta = \pi / 2 + n\pi, \quad n \in \mathbb{Z}$
$\cdot$	$\neq 0$	$= 0$	$\cos\theta\,(A\cos\theta + B\sin\theta) = 0$	$\theta = \pi / 2,\ \arctan[-\,A / B] + n\pi$
$\cdot$	$\cdot$	$\neq 0$	$A + B\tan\theta + C\tan^2\theta = 0$	$\theta = \arctan\left[\frac{-B \pm \sqrt{B^2 - 4AC}}{2C}\right] + n\pi$

Table E.1: Roots $\theta \in \mathbb{R}$ of the equation $Q[\theta] = A\cos^2\theta + B\sin\theta\cos\theta + C\sin^2\theta$ for different assumptions of the values of the parameters $A, B, C \in \mathbb{R}$, $B^2 - 4AC \geq 0$

D	E	Equation	Roots
$= 0$	$= 0$	$0 = 0$	$\forall\, \theta$
$\neq 0$	$= 0$	$\cos\theta = 0$	$\theta = \pi / 2 + n\pi, \quad n \in \mathbb{Z}$
$= 0$	$\neq 0$	$\sin\theta = 0$	$\theta = n\pi, \quad n \in \mathbb{Z}$
$\neq$	$\neq 0$	$\tan\theta = -\,D / E$	$\theta = \arctan[-\,D / E] + n\pi, \quad n \in \mathbb{Z}$

Table E.2: Roots $\theta \in \mathbb{R}$ of the equation $L[\theta] = D\cos\theta + E\sin\theta$ for different assumptions of the values of the parameters $D, E \in \mathbb{R}$

In Figure E.2 the parameters are $A = 0.001$, $B = 0.01$, $C = -0.00007$, $D = -0.01$, $E = 0.000005$ and $F = -0.03$. The roots of $Q[\theta]$ are 0.49777342465146π, 0.96829651297154π, 1.49777342465146π and 1.96829651297154π. The roots of $L[\theta]$ are 0.49984084507017π and 1.49984084507017π. The roots of $M[\theta]$ are 0.49777454085691π, 0.94235170892355π, 1.49777454085691π and 1.94235170892355π.

To know whether $M[\theta]$, $Q[\theta]$, and $L[\theta]$ are negative or positive between two consecutive roots, we check their first derivatives at the roots. If the first derivative is zero (i.e. the root is a double root) then we need to check the second derivative. A positive first (or second) derivative means the function is positive from that root to the next one, while a negative first (or second) derivative means the function is negative from that root to the next one. The derivatives are given in closed form in the following way (where

$\cos' \theta = -\sin\theta$ and $\sin' \theta = \cos\theta$):

$$
\begin{aligned}
Q'[\theta] &= B\cos^2\theta + 2(C-A)\sin\theta\cos\theta - B\sin^2\theta, \\
Q''[\theta] &= 2(C-A)\cos^2\theta - 4B\sin\theta\cos\theta - 2(C-A)\sin^2\theta, \\
L'[\theta] &= E\cos\theta - D\sin\theta, \\
L''[\theta] &= -D\cos\theta - E\sin\theta,
\end{aligned}
$$

and similarily for $M[\theta]$, which is of the same form as $Q[\theta]$.

The union of all the roots of $M[\theta]$, $Q[\theta]$ and $L[\theta]$ then gives us a set $\{\theta_i\}_{i=1,\cdots,N}$ where we know the signs of the three functions between the roots. As the roots have a period of 2π, we may assume that θ_1 is the smallest root greater than or equal to 0 and that $\theta_N = \theta_1 + 2\pi$. The roots that fall on the interval (θ_1, θ_N) are then chosen as 'candidate roots' owing to the periodicity and then placed in ascending order and labelled so that $\theta_1 < \theta_2 < \cdots < \theta_{N-1} < \theta_N$.

Gaussian integral and roots of quadratic equations

We now look at a two-dimensional integral $\mathcal{N}[\mathcal{D}]$ in $(x,y) \in \mathbb{R}^2$ of a standard normal distribution over a quadratic domain and change to polar coordinates with $\theta \in [\theta_1, \theta_N)$ and $r \in (0, +\infty)$,

$$
\begin{aligned}
\mathcal{N}[\mathcal{D}] &= \frac{1}{2\pi}\int_{-\infty}^{+\infty}\int_{-\infty}^{+\infty} \exp\left[-\tfrac{1}{2}\left(x^2+y^2\right)\right] \mathbb{1}_{\{P[x,y]<0\}}\, dxdy \\
&= \frac{1}{2\pi}\int_{\theta_1}^{\theta_N}\int_{0}^{+\infty} r\exp\left[-\tfrac{1}{2}r^2\right] \mathbb{1}_{\{P[\theta,r]<0\}}\, drd\theta \\
&= \frac{1}{2\pi}\int_{\theta_1}^{\theta_N} \left[-\exp\left[-\tfrac{1}{2}r^2\right] \mathbb{1}_{\{P[\theta,r]<0\}}\right]_{r=0}^{+\infty} d\theta.
\end{aligned} \tag{E.3}
$$

For a given $i = 1, \cdots, N-1$ we look at $\theta \in [\theta_i, \theta_{i+1})$ and, as the signs of $M[\theta]$, $Q[\theta]$ and $L[\theta]$ are known and constant on that interval, we can determine for which $r \in (0, +\infty)$ the inequality $P[\theta, r] < 0$ holds. Looking at the discriminant $M[\theta]$ of $P[\theta, r]$, if $M[\theta] < 0$ then $P[\theta, r]$ has no roots, so we have two possible cases. Either $P[\theta, r] < 0$ for all $r \in (0, +\infty)$, giving an integrand of

$$
\left[-\exp\left[-\tfrac{1}{2}r^2\right]\right]_0^{+\infty} = 1,
$$

or $P[\theta, r] \geq 0$ for all $r \in [0, +\infty)$, giving a zero integrand.

If $M[\theta] \geq 0$, we look at the signs of F, $Q[\theta]$ and $L[\theta]$, thereby obtaining the integration bounds of the integrand

$$-\exp\left[-\tfrac{1}{2}r^2\right]$$

for $r \in (0, +\infty)$. For a given θ and where we assume that $M[\theta] \geq 0$, in the first three columns of Table E.3 we assume different values of the parameters $F, Q[\theta], L[\theta] \in \mathbb{R}$. The last three columns in the table show the signs of the root(s) of $P[\theta, r] = Q[\theta] r^2 + L[\theta] r + F$, the domain of $r \in (0, +\infty)$ for which $P[\theta, r] < 0$ and the value of the integrand in (E.3).

F	$Q[\theta]$	$L[\theta]$	Roots of $P[\theta, r]$	$P[\theta, r] < 0$ for $\mathbf{r} \in$	$\left[-\mathrm{e}^{-\frac{1}{2}r^2}\mathbb{1}_{\{P[\theta,r]<0\}}\right]_{r=0}^{+\infty}$
>0	<0	$\cdot$	$r_+ < 0 < r_-$	$(r_-, +\infty)$	$\mathrm{e}^{-\frac{1}{2}r_-[\theta]^2}$
>0	>0	>0	$r_- \leq r_+ < 0$	$\emptyset$	0
>0	>0	≤ 0	$0 \leq r_- \leq r_+$	(r_-, r_+)	$\mathrm{e}^{-\frac{1}{2}r_-[\theta]^2} - \mathrm{e}^{-\frac{1}{2}r_+[\theta]^2}$
≤ 0	>0	$\cdot$	$r_- \leq 0 \leq r_+$	$(0, r_+)$	$1 - \mathrm{e}^{-\frac{1}{2}r_+[\theta]^2}$
≤ 0	<0	≤ 0	$r_+ \leq r_- \leq 0$	$(0, +\infty)$	1
≤ 0	<0	>0	$0 \leq r_+ \leq r_-$	$(0, r_+) \cup (r_-, +\infty)$	$1 + \mathrm{e}^{-\frac{1}{2}r_-[\theta]^2} - \mathrm{e}^{-\frac{1}{2}r_+[\theta]^2}$
>0	$=0$	$=0$	none	$\emptyset$	0
>0	$=0$	>0	$r_0 < 0$	$\emptyset$	0
>0	$=0$	<0	$0 < r_0$	$(r_0, +\infty)$	$\mathrm{e}^{-\frac{1}{2}r_0[\theta]^2}$
≤ 0	$=0$	$=0$	none	$(0, +\infty)$	1
≤ 0	$=0$	>0	$0 \leq r_0$	$(0, r_0)$	$1 - \mathrm{e}^{-\frac{1}{2}r_0[\theta]^2}$
≤ 0	$=0$	<0	$r_0 \leq 0$	$(0, +\infty)$	1

Table E.3: The signs of the possible roots of $P[\theta, r]$ and the value of the the integrand in (E.3) for different values of F, $Q[\theta]$ and $L[\theta]$.

Here the roots of $P[\theta, r]$ depend on θ and are given by

$$r_\pm[\theta] := \frac{-L[\theta] \pm \sqrt{L[\theta]^2 - 4FQ[\theta]}}{2Q[\theta]} \quad , \quad r_0[\theta] := -\frac{F}{L[\theta]}.$$

The method used to determine the sign of the roots is based on standard inequality calculations along with the following observation that

$$\sqrt{L[\theta]^2 - 4FQ[\theta]} \quad \begin{cases} \leq \sqrt{L[\theta]^2} = |L[\theta]| & \text{for} \quad FQ[\theta] \geq 0, \\ \geq \sqrt{L[\theta]^2} = |L[\theta]| & \text{for} \quad FQ[\theta] \leq 0. \end{cases}$$

For $i = 1, 2, \ldots, N-1$, we let $g_i[\theta]$, $\theta \in [\theta_i, \theta_{i+1})$, be the integrand in the last column of Table E.3. Our two-dimensional standard Gaussian integral I in (E.3) then becomes a sum of one-dimensional integrals,

$$\mathcal{N}\left[\mathcal{D}\right] = \frac{1}{2\pi} \sum_{i=1}^{N-1} \int_{\theta_i}^{\theta_{i+1}} g_i\left[\theta\right] d\theta.$$

For given values of $A, B, C, D, E, F \in \mathbb{R}$, each of the integrals can then be calculated by some one-dimensional numerical integration method.

Table E.4 shows an example of a Gaussian integration over a conic. The values of the parameters are $A = 0.001$, $B = 0.01$, $C = -0.00007$, $D = -0.01$, $E = 0.000005$ and $F = -0.03$ (an hyperbola). We find the roots of $Q\left[\theta\right]$, $L\left[\theta\right]$ and $M\left[\theta\right]$ using Tables E.1 and E.2, and the integrand using Table E.3. Summing up the fourth column gives the probability of 96.6623731003525%. To calculate the integral $\frac{1}{2\pi}\int_{\theta_i}^{\theta_{i+1}} g_i\left[\theta\right] d\theta$ we used an integration method provided by NAG (Numerical Algorithms Group). The last column shows the number of times $g_i\left[\theta\right]$ had to be valued to approximate the value of the integral. The error for each integral is less than 10^{-10} with only 336 valuations of the integrand.

Because the roots of $Q\left[\theta\right]$, $L\left[\theta\right]$ and $M\left[\theta\right]$ have periodicity π, we note that we could use polar coordinates such that $\theta_N - \theta_0 = \pi$ and $r \in \mathbb{R}$—that is to allow r to take both positive and negative values. Then we could pair the integrands $g_i\left[\theta\right]$, $i = 1, 2, \ldots, N/2$ together with the integrand $g_{N/2+i}\left[\theta\right]$. Table E.3 would then be simpler, we would half the number of integrals $\frac{1}{2\pi}\int_{\theta_i}^{\theta_{i+1}} g_i\left[\theta\right] d\theta$ and, in turn, make the calculation of $\mathcal{N}\left[\mathcal{D}\right]$ faster.

i	θ_i	$g_i\left[\theta\right]$	$\frac{1}{2\pi}\int_{\theta_i}^{\theta_{i+1}} g_i\left[\theta\right] d\theta$	NAG
1	0.497773424651462	1	$5.58102721790474E-07$	0
2	0.497774540856905	1	0.0010331521066328	0
3	0.499840845070171	1	0.221255431926692	0
4	0.942351708923554	$1+\mathrm{e}^{-\frac{1}{2}r_-[\theta]^2}-\mathrm{e}^{-\frac{1}{2}r_+[\theta]^2}$	0.0129326423061761	21
5	0.968296512971535	$1-\mathrm{e}^{-\frac{1}{2}r_+[\theta]^2}$	0.232158529261464	147
6	1.49777342465146	$1+\mathrm{e}^{-\frac{1}{2}r_-[\theta]^2}-\mathrm{e}^{-\frac{1}{2}r_+[\theta]^2}$	$5.58102721790474E-07$	21
7	1.49777454085691	1	0.00103315210663273	0
8	1.49984084507017	1	0.221255431926692	0
9	1.94235170892355	1	0.0129724020239904	0
10	1.96829651297153	$1-\mathrm{e}^{-\frac{1}{2}r_+[\theta]^2}$	0.263981873139802	147

Table E.4: An example of a Gaussian integration over a conic. The values of the parameters are $A = 0.001$, $B = 0.01$, $C = -0.00007$, $D = -0.01$, $E = 0.000005$ and $F = -0.03$.

Bibliography

D.-H. Ahn, R.F. Dittmar, and A.R. Gallant. Quadratic term structure models: Theory and evidence. *Review of Financial Studies*, 15(1):243–88, 2002.

A.S.T. Andersen. Inflation modelling using time changed Levy processes. Working paper, 2009. URL http://ssrn.com/abstract=1373969.

S. Assefa. Calibration and pricing in a multi-factor quadratic Gaussian model. *No. 197, Research Paper Series from the Quantitative Finance Research Centre, University of Technology, Sydney*, 2007.

D. Beaglehole and M. Tenney. General solutions of some interest rate-contingent claim pricing equations. *Journal of Fixed Income*, 1–2:69–83, 1991.

N. Belgrade, E. Benhamou, and E. Koehler. A market model for inflation. IXIS Corporate and Investment Bank working paper, January 2004. URL http://ssrn.com/abstract=576081.

F. Black and M. Scholes. The pricing of options and corporate liabilities. *Journal of Political Economy*, 81(3):637–54, 1973.

A. Brace, D. Gatarek, and M. Musiela. The market model of interest rate dynamics. *Mathematical Finance*, 7(2):127–47, April 1997.

P. Carr, D.B. Madan, and R.H. Smith. Option valuation using the fast Fourier transform. *Journal of Computational Finance*, 2:61–73, 1999.

X. Charvet and Y. Ticot. Pricing with a smile: An approach using normal inverse Gaussian distributions with a SABR-like parameterisation. Working paper, 2011. URL http://ssrn.com/abstract=1968453.

M. Deacon, A. Derry, and D. Mirfendereski. *Inflation-indexed securities: Bonds, swaps and other derivatives.* John Wiley & Sons Ltd, second edition, 2004.

D. Duffie. *Dynamic Asset Pricing Theory.* Princeton University Press, 3rd edition, 2001.

Eurostat. Compendium of HICP reference documents. *European Communities*, 2001.

I. Fisher. *The Purchasing Power of Money.* Macmillan, 1922. URL http://files.libertyfund.org/files/1165/0133_Bk_Sm.pdf.

I. Fisher. *The Money Illusion.* Adelphi, 1928.

I. Fisher. *The Theory of Interest.* Porcupine Press, 1930. URL http://files.libertyfund.org/files/1416/0219_Bk_Sm.pdf.

M. Friedman and R. Friedman. *Free to Choose — A Personal Statement.* Harcourt Brace Jovanovich, 1980.

H. Gretarsson. Inflation derivatives modelling: A Markovian multi-factor approach with linear-Gaussian applications. Master's thesis, Mathematical Finance, Imperial College London, September 2009.

P. Hagan, D. Kumar, A. Leśniewski, and D. Woodward. Managing smile risk. *Wilmott Magazine*, pages 84–108, September 2002.

D. Heath, R. Jarrow, and A. Morton. Bond pricing and the term structure of interest rates: A new methodology for contingent claims valuation. *Econometrica*, 60(1):77–105, 1992.

S. Heston. A closed-form solution for options with stochastic volatility with applications to bond and currency options. *Review of Financial Studies*, 6(2):327–43, 1993.

L. P. Hughston. Calibration of the one-factor interest rate model. Working paper, Derivative Product Risk Management, Merrill Lynch International, May 1994.

L. P. Hughston. Inflation derivatives. Working paper, Merrill Lynch and King's College London, 1998 (with added note 2004). URL http://www.mth.kcl.ac.uk/finmath/articles/Inflation_Derivatives.pdf.

J. Hull and A. White. Pricing interest-rate-derivative securities. *The Review of Financial Studies*, 3(4):573–592, 1990.

J. Hull and A. White. One-factor interest-rate models and the valuation of interest-rate derivative securities. *Journal of Financial and Quantitative Analysis*, 28(2):235–54, 1993.

P. Jäckel and J. Bonneton. Inflation products and inflation models. *Wiley Encyclopedia of Quantitative Finance*, 2010.

R. Jarrow and Y. Yıldırım. Pricing treasury inflation protected securities and related derivative securities using an HJM model. *Journal of Financial and Quantitative Analysis*, 38(2):337–358, 2003.

I. Karatzas and S.E. Shreve. *Brownian Motion and Stochastic Calculus.* Springer, second edition, 1991.

C. Kenyon. Inflation is normal. *Risk*, pages 76–82, July 2008.

J. Kerkhof. Inflation derivatives explained: Markets, products, and pricing. *Lehman Brothers, Fixed Income Quantitative Research*, July 2005.

M. King. Twenty years of inflation targeting. Speech given at the London School of Economics, October 2012. URL http://www.bankofengland.co.uk/publications/Pages/speeches/2012/606.aspx.

A.G. Lamorgese, S.B. Pope, P.K. Yeung, and B.L. Sawford. A conditionally cubic-Gaussian stochastic Lagrangian model for acceleration in isotropic turbulence. *The Journal of Fluid Mechanics*, 582:423–48, 2007.

A.G. Lamorgese, S.B. Pope, and P.K. Yeung. Analysis of the conditionally cubic-Gaussian stochastic Lagrangian model. *Physica Scripta*, T132, 2008.

M. Leippold and L. Wu. Asset pricing under the quadratic class. *Journal of Financial and Quantitative Analysis*, 37:271–95, 2002.

Markit. The Totem Service, 2012. URL http://www.markit.com/en/products/valuations/markit-totem/totem.page. Data accessed on 31 January 2012.

P. McCloud. Exponential-quadratic semimartingale models of interest rate and asset volatility. Working paper, 2008. URL http://ssrn.com/abstract=1269577.

P. McCloud. Putting the smile back on the face of derivatives. *Risk*, pages 117–123, January 2010.

F. Mercurio. Pricing inflation-indexed derivatives. *Quantitative Finance*, 5(3):289–302, June 2005.

F. Mercurio and N. Moreni. Inflation with a smile. *Risk*, pages 70–75, March 2006.

F. Mercurio and N. Moreni. Inflation modelling with SABR dynamics. *Risk*, pages 106–111, June 2009.

J. O'Donoghue, L. Goulding, and G. Allen. Consumer price inflation since 1750. *Office for National Statistics*, 2004.

Office for National Statistics. The new inflation target: The statistical perspective, December 2003a.

Office for National Statistics. Consumer price indices: Technical manual, April 2012b.

B. K. Øksendal. *Stochastic Differential Equations: An Introduction with Applications*. Springer, sixth edition, 2003.

A. Pelsser. *Efficient Methods for Valuing Interest Rate Derivatives*. Springer, 2000.

V. Piterbarg. Rates squared. *Risk*, pages 100–105, January 2009.

L.C.G. Rogers. The potential approach to the term structure of interest rates and foreign exchange rates. *Mathematical Finance*, 7:157–76, 1997.

P.A. Samuelson and W.D. Nordhaus. *Economics*. McGraw-Hill, 19th edition, 2010.

B.S. Sigurðsson. *Pricing models for inflation derivatives*. PhD thesis, Department of Management Science and Engineering, Stanford University, 2008.

C.C. Tan. Examining the latest advances in pricing and hedging inflation derivatives. Presentation given at the Global Derivatives Conference in Barcelona, April 2012.

M. Trovato, D. Ribeiro, and H. Gretarsson. Inflation derivatives modelling. Working paper, Inflation Derivatives Research, Lloyds Bank Corporate Markets, October 2009.

M. Trovato, D. Ribeiro, and H. Gretarsson. Quadratic Gaussian inflation. *Risk*, pages 86–91, September 2012.

O. Vašíček. An equilibrium characterization of the term structure. *Journal of Financial Economics*, 5:177–88, 1977.

J. Vecer. A new PDE approach for pricing arithmetic average Asian options. *The Journal of Computational Finance*, 4(4):105–13, 2001.