

Measuring inflation expectations and the effect of
monetary policy
Euro area inflation compensation and ECB policy
announcements

Economics
Master's thesis
David Cowen
2016

Author	David Cowen	
Title of thesis	Measuring inflation expectations and the effect of monetary policy	
Degree	Master of Science (Economics and Business Administration)	
Degree programme	Economics	
Thesis advisor(s)	Pekka Ilmakunnas	
Year of approval	Number of pages	Language
2016	71	English

Abstract

This thesis looks at ways of measuring the inflation expectations of economic agents, by constructing a market-based measure from the inflation compensation of government bonds. Inflation compensation is defined as the difference in yields between nominal and index-linked government bonds leaving investors indifferent between holding the two. Inflation compensation is a true measure of expectations only if investors are risk neutral, hence the presence of liquidity and inflation risk premia is evaluated and these premia are extracted. After deriving inflation expectations in this way, the effect of monetary policy measures and announcements on expectations is analysed.

The focus is on the Euro area experience for the time period 2007 – 2015. The data used in the analysis consists of daily prices of French government bonds, both nominal and linked to the Euro area harmonized consumer index (HICP). The monetary policy measures evaluated are those conducted by the European Central Bank (ECB) during the sample period, including changes to key interest rates and non-standard monetary policy measures.

The results show that there are clear and time-varying liquidity and inflation risk premia in the inflation compensation. The effect of liquidity is not substantial and is stable throughout the sample reflecting good liquidity in index-linked French government bonds. The extraction of the inflation risk premium is challenging due to the lack of observations and unrealistically low variance of the survey expectations used to identify the true inflation expectations. Hence, the final expectations and inflation risk premium are obtained only at a quarterly frequency.

The inflation compensation measure exhibits strong volatility during the sample period, at all of 5-, 10- and 15-year horizons. This would suggest that Euro area inflation expectations are weakly anchored. Expectations fell dramatically during the height of the financial crisis in 2008, but rebounded quickly afterwards. The trends for the different time horizons began dispersing thereon, reflecting growing uncertainty and deflationary fears in the short term.

The daily series of both liquidity-adjusted and non-adjusted compensation is used in an event study context to determine whether inflation expectations react to the ECB's monetary policy measures and announcements. There appears to be a positive correlation between past values of expectations and the present ECB rate, and a negative correlation between the current rate and future expectations, pointing to monetary easing being able to increase expectations when they are low. Rate change announcements look to be well anticipated, with little change in expectations on days of the announcements. Non-standard measures are less anticipated and they succeed in raising expectations, albeit briefly. Finally, the effects of the ECB's expanded asset purchase programme (EAPP) are evaluated. It appears the launch of the programme raised expectations significantly for a while, before they fell back to previous levels during the summer of 2015. The expectations were rising again at the beginning of autumn 2015.

Keywords Inflation expectations, inflation compensation, monetary policy, yield curve, variance ratio test, state-space-model

Contents

1.	Introduction.....	5
1.1	Research question and methodology.....	8
1.2	Thesis structure	9
2.	Review of the literature	11
2.1	Survey measures	11
2.2	Inflation compensation.....	14
2.3	Options-implied probability density functions	16
3.	Methodology	20
3.1	Yield curve fitting.....	21
3.1.1	Yield-to-maturity	21
3.1.2	The Nelson-Siegel-Svensson yield curve.....	22
3.2	Inflation compensation and inflation expectations.....	25
3.2.1	Test of random walk hypothesis of inflation compensation.....	27
3.2.2	Extracting liquidity and inflation risk premia	28
3.3	Event study.....	33
4.	Data	35
4.1	French government inflation-linked bonds	35
4.2	ECB policy announcements	37
4.2.1	The Harmonized Index of Consumer Prices	39
5.	Estimation results	41
5.1	Yield curve estimation	41
5.2	Inflation expectations	44
5.2.1	Variance ratio test	45
5.2.2	Liquidity premium	47
5.2.3	Inflation risk premium.....	49
5.3	Event study.....	53
5.3.1	Expanded asset purchase programme (EAPP).....	59

6. Conclusions.....	63
References.....	67
Appendix.....	71

List of Tables

Table 1. Characteristics of outstanding OAT€i issues.....	36
Table 2. Variance ratio test of the random walk hypothesis of the inflation compensation.	46
Table 3. Regression of liquidity premium in inflation compensation	47
Table 4. Regressions of liquidity-adjusted and non-adjusted 5-year inflation compensation on the ECB's main refinancing operations rate.	53
Table 5. Dates of non-standard monetary policy measures conducted by the ECB 2009-2015.....	56
Table 6. Regressions of changes to inflation compensation on time dummies of ECB monetary policy announcements.	58
Table 7. Regressions of the liquidity-adjusted 5-year inflation compensation on a 120-day pre-announcement and post announcement period.	60

List of Figures

Figure 1. Issued OAT€i securities.	36
Figure 2. The ECB's key policy rates.	39
Figure 3. French government index-linked yield curve 13.11.2014.....	42
Figure 4. Inflation compensation on 13.11.2014.....	43
Figure 5. Euro area inflation compensation 14.3.2007-18.8.2015	44
Figure 6. Estimated liquidity premium of French government 5-year index-linked bonds.....	48
Figure 7. Estimated 5-year inflation expectations from liquidity-adjusted and non-adjusted inflation compensation.	51
Figure 8. Decomposition of 5-year inflation compensation into inflation expectations and inflation risk premium.	52
Figure 9. Impulse responses of shocks to inflation expectations and the ECB's key rate.	55
Figure 10. Dates of announcements of changes to the ECB rate (upper panel) and of ECB non-standard monetary policy measures (lower panel).	59
Figure 11. Trends of liquidity-adjusted 5-year inflation compensation before and after announcement of the EAPP on 22 January 2015.....	61

1. Introduction

An integral part of almost any model in economics is that economic agents make decisions based on what they think is going to happen in the future. In other words they must make a prediction of the future values of key economic variables, such as income (wages in the case of individuals, profits in the case of firms), output, taxes, asset prices (such as stock prices, interest rates and exchange rates) and, being most relevant to this study, prices. Estimating the future price level is the same process as estimating the rate at which prices change, hence the variable of interest is inflation (or if prices are falling, deflation).

Given that economic agents base their behaviour on expectations, any policymaker wishing to affect that behaviour, or even merely to analyse it, should have a way to measure the prevailing expectations among agents to be able to decide on the appropriate policy measures. When focusing on prices, the relevant policymaker is the monetary authority charged with conducting monetary policy, usually the central bank of a country or currency union.

Price stability is almost universally a key consideration among central banks' main policy goals. Price stability is usually stated as a specific target rate of inflation that the central bank has pledged to maintain. The most common inflation target is around 2 % annually. The European Central Bank (ECB) has defined its target as less than but close to 2 % annual growth over the medium term. The Bank of England has a target of exactly 2 %. The Bank of Canada has a band of 1 % to 3 % for inflation and the Reserve Bank of Australia has a target of 2 – 3 %. Conspicuously lacking an explicit target rate of inflation is the U.S. Federal Reserve, which has merely stated long-term price stability as the second part of its dual mandate.

A key way for a central bank to determine whether the policies and communication it is currently conducting are maintaining inflation close to its target, is looking at the inflation expectations of agents at different time horizons. More precisely looking at the variability of long-run expectations gives some indication about

how credible the central bank's inflation target is and how well anchored inflation expectations are to it.

The anchoring of expectations means how much long-run expectations react to new information, including price shocks. If expectations are perfectly anchored, long-run expectations do not change at all over time, even if the economy experiences current rates of inflation that are higher than the target rate. This means that a sudden increase in for example the price of oil will not have a persistent effect on inflation. On the other hand, weakly anchored expectations do react to new information and a higher current rate of inflation will likely cause agents to increase their long-run expectations. This would in turn result in higher wage demands, which will lead to higher prices and the cycle repeats until the long-run expectation level is met. The anchoring of inflation expectations is therefore an essential consideration for a central bank attempting to maintain long-run price stability. (Bernanke 2007)

The anchoring of inflation expectations raise the importance of central bank credibility, in other words, do economic agents believe that the measures the central bank is taking are sufficient to maintain price stability? This is usually connected to a discussion of a policy rule for monetary policy. The most widely used policy rule is the Taylor rule and related Taylor principle that states that the nominal interest rate set by the central bank should respond to changes in inflation by more than one to one. This should then be sufficient to keep long-run inflation expectations anchored. (Gonzalez-Paramo 2007)

The problem for central banks in recent years, and currently especially for the European Central Bank, is that firstly, the nominal interest rate has reached the zero lower bound and cannot adjust anymore. Secondly, the problem in maintaining price stability has traditionally been concerned with too high inflation, whereas presently, in the Euro area, the problem is that of non-existent inflation and borderline deflation.

This means that in order to affect inflation expectations central banks have had to resort to non-standard monetary policy measures, i.e. measures that do not involve changes to the nominal interest rate. The most prominent of these has been the purchase of public sector bonds using printed money, or what is known as quantitative easing. In addition there is increased scrutiny and importance on the

communication of central banks, including informing the public of the schedule of current and future monetary policy measures, which is called forward guidance.

Besides the practical considerations discussed so far, expectations matter to central banks also from a theoretical perspective. The main macroeconomic models used by central banks today belong to the class of new Keynesian models. These models are derived from the utility and profit maximizing behaviour of households and firms respectively, hence they are said to have microfoundations, and they are forward-looking by incorporating agents' expectations of future inflation and output. The Keynesian aspect of the models comes from nominal wage or price rigidities and the consequential assumption that counter-cyclical policy measures can ease short-term fluctuations in aggregate demand. (Walsh 2010, 329)

The basic form of the new Keynesian model is expressed in two equations; the expectations augmented IS curve (equation 1.1) and the new Keynesian Phillips curve (equation 1.2):

$$x_t = E_t x_{t+1} - \left(\frac{1}{\sigma}\right) (i_t - E_t \pi_{t+1}) + u_t \quad (1.1)$$

$$\pi_t = \beta E_t \pi_{t+1} - \kappa x_t + e_t \quad (1.2)$$

where x_t is the output gap, i_t is the nominal interest rate, π_t is the inflation rate, σ is a parameter of how the output gap responds to changes in the real interest rate, β is a parameter of how current inflation depends on expected inflation, κ is a term describing the degree of price rigidity and u_t and e_t are disturbances. The IS curve describes the demand side of the economy and is derived from the Euler condition placed on the representative households' decision problem and the Phillips curve describes the supply side and is derived from firms' optimal pricing decisions. (Walsh 2010, 330-340)

This model has become the standard framework for monetary policy analysis since the mid 1990s (Walsh 2010, 329). As both the demand and supply side equations contain the term for expected future inflation, the measurement and understanding of these expectations is surely a key concern for a central bank when analysing its monetary policy measures in the context of the new Keynesian model.

1.1 Research question and methodology

The purpose of this thesis is to provide a practical empirical measure of inflation expectations in the Euro area and to appraise the interaction between the European Central Bank's monetary policy and the expectations. The main motivation for the relevance of the question has been provided in the discussion so far. Particularly having a daily (and in some cases even intraday) market based measure of expectations allows a robust analysis of the immediate effects of monetary policy measures.

The specific relevance of inflation expectations in the Euro area is due to the historically low levels of inflation and inflation expectations of the past few years and the introduction of the non-standard monetary policy measures by the ECB as mentioned before. Chief among these was the introduction of the ECB's expanded asset purchase programme (EAPP) in January 2015. Indeed the two issues are intricately linked given that the ECB's chairman Mario Draghi mentioned historically low inflation expectations as one of the main reasons behind the ECB's decision to implement the EAPP. It is therefore adequate to look at market based inflation expectations leading up to the beginning of 2015 and gauge what impact if any the asset purchase programme has had so far.

The emphasis in the study is on the technical considerations in observing the markets' forecast for inflation. The discussion on expectations formation mechanisms is therefore excluded. We will simply follow convention and assume that agents have rational expectations. We touch briefly upon the criticism levelled at rational expectations, especially in the context of inflation expectations and its relevance to the potency of monetary policy in section 2.1.

We will also not attempt to evaluate the accuracy or predictive power of the market based expectations we produce. We are interested only, much as a central bank would, on the current level of expectations and its implications for monetary policy. The forecasting performance of different methods, especially surveys, is again briefly discussed in section 2.1.

The method for estimating the expectations is to look at the price differences of nominal and inflation-linked government bonds, a concept known as inflation compensation, or break-even inflation. Inflation compensation is defined as the extra yield investors require to hold nominal government bonds compared with otherwise identical inflation-linked bonds. The exact methodology follows Gur-

kaynak, Sack & Wright (2010), who place emphasis on the likelihood of liquidity and inflation risk premia inherent in inflation compensation. The true inflation expectations can therefore only be produced by extracting the aforementioned premia from the measure of compensation.

After computing this measure of inflation expectations, we assess the effectiveness of the ECB's monetary policy. We look at both changes to interest rates and non-standard measures as well as the ECB's communication on both measures. Combined with the expectations for different horizons, this analysis will give some indication whether inflation expectations have remained well anchored through the turbulence in the Euro area in recent years and if not, has the ECB been able to bring medium to long-run expectations back closer to its inflation target.

The specific contribution of this thesis to the literature is to look at the Euro area experience of the interaction between monetary policy and market based measures of inflation expectations. There are studies focusing on the interaction of market based expectations and macroeconomic news announcements (for example Gürkaynak, Sack & Swanson 2005 and Kitsul & Wright 2012 for U.S. data) and the surprise component in ECB announcements (Winkelmann, Bibinger & Lintzert 2014), but not specifically between inflation expectations and monetary policy. There is also a dearth of prominent studies on Euro area inflation expectations during the sovereign debt crisis period. In particular, there has not yet been a meaningful analysis of the effects of the EAPP given the short period of time since its announcement.

1.2 Thesis structure

The thesis is structured as follows. Chapter 2 contains a review of the existing literature on empirical measurement of inflation expectations. This starts off by looking at results gained from surveys conducted to both households and professional forecasters, including observations about the nature of inflation expectations. Next we look at results gained using similar methodology to the one used in this study, i.e. studies using inflation compensation as a measure for inflation expectations. Finally we look at the most recent development in market based measures of inflation expectations, namely options-implied probability density functions for inflation, which provide an estimate of the uncertainty linked to future inflation in addition to a point estimate.

Chapter 3 introduces the main methodology used in the empirical analysis. It begins with the formation of the nominal and real yield curves for government bonds from which the inflation compensation can be calculated. Next we present the variance ratio test to determine the existence of the liquidity and inflation risk premia and then how to extract them from the compensation using a simple linear regression for the liquidity premium and a state-space model for the inflation risk premium. Finally we go through the methods of looking at the interaction between inflation expectations and monetary policy by running an event study on the ECB monetary policy announcements. This includes several linear regression models regressing inflation expectations on the ECB's key rate and time dummies for announcements as well as a VAR model looking at the interaction between the key rate and expectations. The section ends with an appraisal of the effects of the EAPP.

Chapter 4 introduces the data to be used in the analysis, mainly the characteristics of the French government bonds, in particular the bonds linked to the Euro area consumer price index, used for the yield curve estimation, and the ECB's monetary policy variables. Chapter 5 contains the results from the empirical analysis following the methodology laid out in Chapter 3 and Chapter 6 concludes.

2. Review of the literature

The measurement of prevailing inflation expectations among economic agents has been fixed closely with the development of financial markets and instruments. Especially in recent times, since the mid-2000s, there has been an increase in the amount of assets and derivatives, whose payoff is directly linked to inflation. This has followed from an increased willingness among investors to have direct hedging against future inflation outcomes. This development is of advantage to economists who now have better and more flexible methods to gauge a market based measure of inflation expectations from the price movements of these assets.

This chapter presents the existing literature of estimating a direct measure of expectations. It is structured in chronological order, beginning with survey expectations, which have been for most of the time the common as well as the only available method of measuring expectations¹. We then look at results gained from similar methods that will be used in this thesis, using inflation compensation of government bonds. This methodology became available when governments began issuing debt tied to the relevant consumer price index of the Euro area around the turn of the millennium. Finally we review the most recent development in expectations measurement, that of inflation options and the probability density functions for future inflation that can be extracted from them using options pricing theory.

2.1 Survey measures

Before the development of financial instruments tied to inflation, the only way for policymakers to evaluate actual inflation expectations of agents was through direct

¹ Besides surveys, there is a rich literature of measuring expectations indirectly by using a theoretical model, such as the Phillips curve, containing expectations and some assumption on the formation mechanism of these expectations. The model and expectations hypothesis are then empirically tested jointly. (Kuismanen & Spolander 1994, 7) However the focus of this study is on direct measurement of expectations, hence we omit these kinds of indirect measures from the review.

surveys. Surveys of expectations are still gathered however, to produce an independent measure to complement market based measures. The participants of surveys can be both professional forecasters, such as economists and financial markets professionals, as well as the general public.

In general there are three types of inflation surveys. Firstly, broader categorical surveys asking respondents whether they believe prices will go up, down or remain the same during a specific time period, often one year into the future. Secondly, quantitative surveys asking for a range for future inflation, such as between 1 % and 3 %. Thirdly, quantitative point estimates of future inflation giving a discrete forecast, such as 2 %. (Kuismanen & Spolander 1994, 8). Currently, in addition to a point estimate, respondents can be asked for density forecasts around the point estimate, giving more information about the uncertainty of the forecasts. This is usually only included in surveys directed at professional forecasters, as is the case for the Survey of Professional Forecasters (SPF) conducted by the ECB that will be used later in the analysis. (European Central Bank 2015b)

The evidence of the usefulness of survey measures of inflation expectations as a source of information for a central bank's monetary policy decisions is mixed. Berk (1999) uses survey data of Dutch inflation expectations and finds that the derived expectations are cointegrated with future realized inflation and the forecast errors are stationary, which support the notion that survey expectations have informational value for monetary policy. However this conclusion is tempered with the observation that there appears to be no causal relationship between expected inflation and future inflation, hence care is advised in the use of survey data for policy purposes.

In general, survey measures have performed very well in predicting future inflation, at least in the short-run. Ang, Bekaert and Wei (2007) find that survey forecasts for inflation outperform three other forecast methods, including ARIMA time-series models of historical inflation, regression forecasts based on the Phillips curve using real economic activity measures and VAR models using the term structure of interest rates. They use survey measures of both professional forecasters and the general public and find that, though the professional forecasters perform better, even the surveys of consumers outperform the other aforementioned forecasts. Fama and Gibbons (1984) compare forecasts from the term structure of interest rates and surveys, and find the surveys providing inferior estimates, alt-

though their study was conducted at the time when inflation had only just ceased to be high, making comparisons between the two studies difficult. In any case it seems worthwhile using survey expectations as a benchmark in this study, although we will employ expectations of medium to long-term inflation and not the short-term.

The other main reason to look at survey expectations beside for forecasting is to determine if inflation expectations are rational. Mankiw, Reis and Wolfers (2003) look at whether the survey estimate of inflation can predict the forecast errors of the same surveys, which under rationality should not be possible. They report that rationality holds for two out of the four surveys they use in the study. They also find autocorrelation in the forecast errors for all surveys, implying that errors are persistent and that there is information in past errors that agents are not correcting for, again a violation of rationality. Thomas (1999) performs a similar analysis and finds that agents do not incorporate all available information on macroeconomic variables into their forecasts, hence also rejecting rationality.

The rationality or lack thereof of expectations does matter for the purposes of a central bank's monetary policy. This is due to the main class of new classical models, which argue that, if agents are fully rational, monetary policy cannot be effective unless it is unexpected, in other words unless the policymaker is able to surprise the public. However if the expectations of individuals are not fully rational, this would no longer be the case, and even rules-based, predictable monetary policy can be effective. (Thomas 1999, 125)

Furthermore, one direct consequence of rational expectations is the efficient market hypothesis, which leads directly to the random walk hypothesis. These are mainly applied to stock market fluctuations, but can be extended to expectations of any stochastic process. The reasoning is that rational investors will use all available information to make forecasts on future stock prices. They will then buy the stocks they expect will have higher returns and sell stocks they expect will have lower returns. From this it follows that the expectations of the investors become built into the stock prices and the risk-adjusted expected return of all stocks are equal. Hence the only factors that can change stock prices are surprise random shocks that are unknown at the time of the forecasts. (Sargent 2008). The random walk property of inflation expectations, assuming rational expectations, is used to determine the existence of liquidity and inflation risk premia in section 3.2.1.

There are a few well documented practical issues with surveys, which in particular can be improved upon by market based measures of expectations. Firstly, as with all survey based data, the measure of expectations is susceptible to sampling errors and possible departures from the assumed underlying probability distribution. Furthermore there is evidence of sensitivity regarding the phrasing and framing of the questionnaire. Finally and most relevantly compared to market based measures, survey expectations are not based on observed behaviour, nor is there evidence that the respondents act corresponding to the views they present. This would be confirmed only through market transactions, which surveys are unable to capture, but market-based measures by definition can. (Berk 1999, 1467-1468)

2.2 Inflation compensation

Most of the problems associated with survey measures of inflation expectations can be corrected by using a measure derived from the actual behaviour of financial markets participants. Given that investors are staking their own funds on bets on future inflation it seems evident that these bets represent their best estimate of future inflation.

The simplest and most common assets linked to inflation are index-linked government bonds, whose coupon and principal are linked to a consumer price index². The idea of inflation linked bonds is nothing new, indeed the first known bonds issued in real terms were issued by the State of Massachusetts in 1780. Until roughly the last 30 years though, only a very small proportion of government debt was in fact linked to inflation. The main reason for issuing real bonds before that time was mainly out of necessity for governments to raise funds during periods of high inflation. (Garcia & Van Rixtel 2007, 11)

In the 1980s governments began issuing inflation-linked bonds for policy reasons. By issuing the debt, the governments were showing that they were committed to ensuring price stability and also they were able to reduce their borrowing costs by eliminating the inflation risk premium that is priced into nominal bonds. Furthermore governments explain the issuance of these bonds as completing markets by giving investors a credible way to hedge against inflation outcomes. The final catalyst for the expansion of inflation-linked debt markets was the arrival of

² A more extended description of the technical aspects of inflation-linked bonds is provided in chapters 3 and 4.

the United States on the scene with the issuance of the first treasury inflation protected security (TIPS) in 1997 and, with the formation of the Euro area, bonds linked to Euro area inflation. (Garcia & Van Rixtel 2007, 9)

The development of the index-linked bonds market, with the increase in the number of issues and growing trading activity and liquidity, enables central banks to gauge the markets expectation of future inflation. When comparing the price or yield of an index-linked bond to a nominal bond with identical characteristics, the only difference should be the expected rate of inflation at maturity. In other words, this measure of inflation compensation is the rate of expected inflation that if realized, means the return of a nominal and index-linked bond is equal.

As mentioned before, inflation compensation can be artificially high due to inflation risk premia in the nominal bonds and artificially low due to liquidity premia in index-linked bonds. Much of the literature on using inflation compensation to estimate inflation expectations concentrates on the extraction or elimination of these premia. The methodology used in this thesis follows Gurkaynak et al. (2010) and is presented in section 3.2.2. They look at U.S. Treasury bonds for the time period 1999-2009 and find clear liquidity and inflation risk premia in the TIPS. Their final measure of five and ten year ahead inflation expectations is fairly steady between 2,3 % and 2,9 % throughout the sample only dropping sharply at the onset of the financial markets turmoil in 2008.

Many of the earlier studies using mainly data on U.K. index-linked bonds, the only regularly traded index-linked bonds at the time, simply assume that there is no inflation risk premium or that it is constant over time. Evans (1998) also uses U.K. data and derives a yield curve formulation from an asset pricing model to compare the term structure of nominal and index-linked bonds. He finds a significant and time-varying inflation risk premium in the nominal bonds. Ejsing, Garcia and Werner (2007) compare Euro area inflation compensation to inflation swaps and find that the measure of inflation expectations derived from the government bond prices contains seasonal variation. This is due to the reference price index having seasonal variation and it affects the shape of the real yield curve at short maturities. Given that we will look at medium to long-term expectations and eliminate short-term maturity bonds from the yield curve due to their high price volatility, we will omit correcting for seasonal effects from our analysis. It is however an important consideration for studies looking at shorter-term expectations.

Jochman, Koop and Potter (2008) consider the process of inflation expectations by comparing short-run and long-run expectations derived from TIPS. To see how anchored expectations are, they look at the inflation pass through coefficient, which determines how changes in expectations at one horizon transfer to another horizon. They find that expectations are not perfectly anchored, but rather contained inside a reasonably narrow band.

As mentioned earlier, the majority of the studies using the measure of inflation compensation to estimate inflation expectations use data from the time before the Euro area sovereign debt crisis, with most samples barely containing the U.S. sub-prime crisis. Our sample on the other hand begins just before the crisis in 2007 and comparing the level of inflation compensation we derive with much of the existing literature, it seems clear inflation expectations have decreased significantly, especially during the past few years.

2.3 Options-implied probability density functions³

Since the mid-2000s, there has been significant development in the number and variety of financial instruments and derivatives where the payoff is linked to future inflation. Inflation swaps for instance are over-the-counter contracts providing inflation protection and are currently the most popular and liquid inflation derivative available. In an inflation swap, one party wishing to hedge against inflation, called an inflation receiver, agrees to pay a fixed rate to another party called an inflation payer, who in return makes a floating rate payment which is linked to the realized inflation at the time of maturity. The observed market rate is the fixed rate payment and, under risk-neutrality, it is the direct market expectation of inflation at the time of maturity. If investors are risk-averse it will be the expectation plus a risk premium. (Hurd & Relleen 2006, 25-26)

Haubrich, Pennachi and Ritchken (2012) note that real yields derived from the difference between yields of nominal government bonds and the market rate of inflation swaps depend less on the variability of liquidity than index-linked government bonds and are hence a superior method to estimate inflation compensation than the methods discussed in section 2.2. They use data from all three measures of inflation expectations discussed so far, surveys, index-linked bonds

³ The methods described in this section would have been the preferred methodology for this thesis, but unfortunately data on the over-the-counter prices of the options were not publically available.

and swaps, to perform a comprehensive decomposition of the derived inflation compensation into real interest rate, expected inflation and long-term inflation components as well as a time varying stochastic risk premium. They find that the short-term real interest rate and expected inflation are highly volatile, but long-term expectations fairly stable and decreasing through their sample of 1982-2010.

An even more recent development in inflation derivatives are inflation options, commonly referred to as inflation caps and floors. These are contracts where the buyer of a cap will be compensated if inflation is above a certain strike rate k at time n . More precisely they will receive a fraction of a nominal underlying principal given by:

$$\max((1 + \pi(n))^n - (1 + k)^n, 0) \quad (2.1)$$

An inflation floor works similarly, except the payment will be made if inflation is below a certain threshold, hence the terms in (2.1) will be reversed. Naturally if inflation turns out to be below the cap or above the floor, the option will expire out-of-the money. (Kitsul & Wright 2012, 3)

In practice, an inflation cap or floor is often a bundle of caplets and floorlets that have the above payment structure with identical strikes and notional and which expire in consecutive years. The option then pays out only in years when realized inflation is above or below the strike rate. These are called year-on-year inflation options, while the caplets and floorlets trade as zero coupon options. As with inflation swaps, these instruments are also traded only on a bilateral basis over-the-counter. (Smith 2012, 225-226)

The option-like structure of these contracts makes it possible to use option pricing theory to produce the entire probability density function (pdf) of investors' beliefs about future inflation and not just a point estimate. This provides policymakers with much richer information about the dispersion and uncertainty linked to prevailing inflation expectations than inflation compensation or inflation swaps. (Kitsul & Wright 2012, 1-2)

The investors will enter into the option contract only if they believe that the probability of the payoff matches the price of the option. In other words a buyer will only agree to pay the price of the option if they believe inflation will be sufficiently higher than the strike, and likewise the seller is willing to enter the contract

only if they believe the price will be higher than the payoff. The process of generating the pdfs consists then of fitting a probability distribution to the set of observed prices and strikes. (Smith 2012, 225)

The standard technical procedure for fitting pdfs to observed options prices is presented by Breeden and Litzenberger (1978). They find that the pdf is related to the second partial derivative of the call-price function with respect to the strike. Using the pricing model developed by Cox and Ross (1976) the price of a call option (or inflation cap) C is defined as:

$$C(S, X, T) = e^{-rt} \int_X^{\infty} (S_T - X) g(S_T) dS_T \quad (2.2)$$

where S_T is the value of the underlying asset at time T , X is the strike price, r is the risk-free rate and t is the maturity of the option. The pdf $g(S_T)$ is then defined as the second partial derivative of (2.2) with respect to X , yielding:

$$g(S_T) = e^{-rt} \frac{\partial^2 C(X, S, T)}{\partial^2 X^2} \quad (2.3)$$

The actual process of fitting the pdfs is far from straightforward owing to the complex payoff structure of the options and limited amount of observed prices, meaning that a continuous call-price function is not observed. To counter these issues some form of interpolation of the observed prices must be performed in order to produce fitted prices for the missing maturities and then the individual caplets and floorlets prices can be extracted. (Smith 2012, 233)

Furthermore it should be noted that the pdfs obtained in this way will again only reflect investors true expectations about inflation if they are risk-neutral, hence these are called risk neutral densities. A physical density function, assuming risk aversion, will take into account the fact that investors will value positive payoffs more in adverse states of the world. In other words investors are willing to pay more for an inflation option that provides positive cash flow in the event of high inflation or deflation. This increase in the price reflects a risk premium and not a view that these adverse conditions are more probable. Since the risk neutral densities do not account for that, the tails of the distribution will be artificially fat. (Smith 2012, 227-228)

Smith (2012) produces options-implied pdfs for U.K. inflation for the period of 2008-2012 using year-on-year options based on U.K. CPI inflation. He finds that the mean of the distributions, which corresponds to the fixed rate of an inflation swap, is relatively flat throughout the sample, especially after the first few years from the day of the observation. The standard deviation of the distributions however, measuring the uncertainty of investors about future inflation, increased substantially after 2008, with a peak approximately five-years into the future.

Kitsul and Wright (2012) look at U.S. data using a simpler methodology to extract discrete pdfs from zero coupon caps and floors. They then run an event study to see which events affect the probability of deflation or high inflation. They also produce physical density functions using a variety of time series models designed to forecast inflation and compare the ratio between the risk neutral densities and the physical densities to construct an empirical pricing kernel that accounts for the risk premium.

They report that surprise macroeconomic news announcements on the whole do not have a significant effect on the pdfs. Interestingly in the context of this thesis, they find that meetings of the federal open market committee (FOMC) of the Federal Reserve, which announce changes to the Fed's monetary policy, have no effect on the risk of deflation. The estimated pricing kernels confirm the theory that investors have higher marginal utility of payoffs in the tails of the distributions, i.e. during deflation or high inflation.

3. Methodology

As discussed in section 2.2, a standard method of estimating the inflation expectations of a single country has long been comparing the difference between nominal and real yields of government bonds issued by that country, called inflation compensation or breakeven inflation. This is done by comparing the yield on a nominal bond to the yield of a similar inflation index-linked bond. A nominal government bond pays the holder a coupon annually or biannually and the principal at maturity. An index-linked government bond pays a coupon similarly, but the principal payment is multiplied by the reference consumer price index of the country in question. The inflation compensation is calculated from nominal and index-linked bonds of similar maturity and coupon, so that the only difference between the bonds is the inflation compensation.

With the formation of the Euro area, a new consumer price index was formed for it (the harmonized index for consumer prices excluding tobacco, hereon referred to as HICPxT) and subsequently a whole range of new financial markets products based on it. Eurozone governments have begun issuing index-linked sovereign debt based on the Euro area HICPxT and so now it is possible to estimate inflation compensation for the Euro area by comparing the yields on these bonds to nominal bonds issued by the same government.

Governments only issue bonds for a limited (in the case of inflation linked bonds very limited) and discrete number of maturities. Thus it is necessary to fit a yield curve for both the nominal and index-linked bonds to calculate the nominal and real yield (or discount factor) for any horizon and then derive the inflation compensation from these.

The structure of the chapter is as follows: first we describe the method for fitting the yield curve to the observed yields and maturities of the different government

bonds. We then concentrate on how to calculate the inflation compensation from the yields and how to extract inflation expectations from the compensation when taking into account the inflation risk premium and liquidity premium. Finally we look at the event study regression of the derived measure of inflation expectations on the key ECB policy rate and time dummies relating to ECB policy announcements as well as a VAR model and a trend comparison of time periods before and after the announcement of the ECB's expanded asset purchase programme.

3.1 Yield curve fitting

The first step when beginning to fit a yield curve to the market data is to convert the observed bond prices, coupons and maturities into yields. For a zero-coupon bond the continuously compounded yield is:

$$y_t(n) = -\frac{\ln P_t(n)}{n} \quad (3.1)$$

where $P_t(n)$ represents the value today to an investor of a €1 payment n years from now. (Gurkaynak et al. 2010, 72-73)

For a coupon bearing bond the calculation of the yield is more complicated as it has to take into account the coupon rate as well as the price. The main concept we look at in this section regarding the yield of a coupon bearing bond is the yield-to-maturity measure. After that we look at the formulation for the yield curve itself, in this case utilizing the Nelson-Siegel-Svensson yield curve function.

3.1.1 Yield-to-maturity

A common method used among finance practitioners is the yield-to-maturity measure, which connects the price of a bond at any given time to the face value of the bond. The yield-to-maturity can be calculated from the formula for the present value of a bond:

$$PV = \sum_{t=1}^n \frac{C}{(1+r)^t} + \frac{FV}{(1+r)^n} \quad (3.2)$$

where the present value PV corresponds to the observed bond price, C is the coupon rate, the face value FV corresponds to the principal payment at maturity and r is the discount factor. Solving (3.2) for r gives the yield-to-maturity. (Tuckman & Serrat 2011, 100)

For more than one coupon payment the yield-to-maturity cannot however be calculated algebraically, so numerical methods must be used instead. We calculate the yield-to-maturity using the YIELD-function in Excel, which utilizes Newton's method. The basic formula for Newton's method is:

$$r_b = r_a - \frac{f(r_a)}{f'(r_a)} \quad (3.3)$$

where r_a and r_b correspond to the first and second approximation for the true yield respectively. $f(r_a)$ is defined as $PV_a - P$, where PV_a is the value of (3.2) when using the discount factor r_a and P is the true observed bond price. $f'(r_a)$ is the first derivative of $f(r_a)$. The value for r_a can be obtained from an approximation for the yield-to-maturity. A commonly used formula for this approximation is:

$$r_a = \frac{C - \frac{PV - FV}{n}}{\frac{PV + FV}{2}} \quad (3.4)$$

To calculate the true yield, (3.4) is plugged into (3.3) and the algorithm is then run for as many iterations as necessary to achieve a close enough result. The YIELD-function in Excel automatically runs 100 iterations. (Deeley 2008, 2-6)

3.1.2 The Nelson-Siegel-Svensson yield curve

The methods for constructing a yield curve from observed market data can be roughly divided into parametric and non-parametric methods. Non-parametric methods, sometimes referred to as spline-based methods, make no assumption about the functional form of the curve and the number of parameters is not fixed, but growing with the number of data observations. Parametric methods on the

other hand have a fixed number of parameters and a clearly defined functional form. (Gurkaynak et al. 2006, 11-13)

In the case of yield curve fitting the choice between a parametric and non-parametric method comes down to the intended use of the curve. For excess return forecasting, a non-parametric method can prove to be superior due to the lack of information loss that is present in a parametric model. The resulting curve usually fits the data well, but the smoothness of the curve suffers. (Gurkaynak et al. 2006, 12)

For a broader macroeconomic analysis of the yield curve, the loss of information is less relevant, but there are clear advantages of using a defined functional form. Firstly it eliminates the effects of small changes to the yields of individual bonds that are not due to macroeconomic shocks, making the analysis easier. Secondly a correctly chosen functional form will have appealing properties, such as beginning and ending at estimated parameters and having a typical yield curve shape, namely being monotonic and S-shaped or humped. (Gurkaynak et al. 2010, 74) (Nelson & Siegel 1987, 473)

A commonly used functional form for the yield curve, used by most central banks, is the Nelson-Siegel-Svensson yield curve (hereon referred to as NSS), as proposed originally by Nelson and Siegel (1987) and later complemented by Svensson (1994). To derive a function that gives the desired shapes described above, Nelson and Siegel (1987, 474-475) begin by assuming that the final functional form must be a solution to differential equation. A further assumption is that spot interest rates follow from a differential equation and forward rates, being forecasts of spot rates, must be solutions to these difference equations.

They then propose that the instantaneous forward rate at maturity m , $r(m)$, is the solution to a second-order differential equation, with two identical real roots, of the form:

$$r_t(m) = \beta_0 + \beta_1 \cdot \exp\left(-\frac{m}{\tau}\right) + \beta_2 \left[\frac{m}{\tau} \cdot \exp\left(-\frac{m}{\tau}\right)\right] \quad (3.5)$$

To calculate yield as a function of maturity, $y(m)$, (3.5) is integrated from zero to m and divided by m to obtain the following function⁴:

⁴ See Appendix A.

$$y_t(m)^{NS} = \beta_0 + (\beta_1 + \beta_2) \cdot \frac{1 - \exp\left(-\frac{m}{\tau}\right)}{\frac{m}{\tau}} - \beta_2 \cdot \exp\left(-\frac{m}{\tau}\right) \quad (3.6)$$

The curve consists of three parts: long, short and medium -term components corresponding to the parameters β_0 , β_1 and β_2 respectively. This is apparent from the forms of the different sections in (3.5). The long-term component is simply a constant given by the value of β_0 . This is often interpreted as the long-run prevailing interest rate. The medium-term component has the form $f(x) = xe^{-x}$, which begins at zero and decays monotonically back to zero. The short-term component has the form $f(x) = e^{-x}$, which has a non-zero starting value and has a faster decay monotonically to zero. The different monotonic, S-shaped or humped-shaped curves can be created through different values of β_0 , β_1 and β_2 . The parameter τ is a time constant, which determines the rate at which the short and medium -term components decay to zero.

To improve the fit and flexibility of the model especially at long maturities, Svensson (1994, 6) introduces a fourth component to the yield function giving it a second hump. The functional form for the instantaneous forward rate now becomes:

$$r_t(m) = \beta_0 + \beta_1 \cdot \exp\left(-\frac{m}{\tau_1}\right) + \beta_2 \left[\frac{m}{\tau_1} \cdot \exp\left(-\frac{m}{\tau_1}\right) \right] + \beta_3 \left[\frac{m}{\tau_2} \cdot \exp\left(-\frac{m}{\tau_2}\right) \right] \quad (3.7)$$

As before to obtain the yield as a function of the maturity, (3.7) is integrated from zero to m and divided by m to get:

$$\begin{aligned}
y_t(m)^{NSS} = & \beta_0 + \beta_1 \frac{1 - \exp\left(-\frac{m}{\tau_1}\right)}{\frac{m}{\tau_1}} \\
& + \beta_2 \left[\frac{1 - \exp\left(-\frac{m}{\tau_1}\right)}{\frac{m}{\tau_1}} - \exp\left(-\frac{m}{\tau_1}\right) \right] \\
& + \beta_3 \left[\frac{1 - \exp\left(-\frac{m}{\tau_2}\right)}{\frac{m}{\tau_2}} - \exp\left(-\frac{m}{\tau_2}\right) \right]
\end{aligned} \tag{3.8}$$

There are now six parameters in the model to be estimated: β_0 , β_1 , β_2 , β_3 , τ_1 and τ_2 . The interpretation of the parameters is the same as before, only now β_3 controls the shape of the second hump and τ_2 its rate of decay to zero. For computational reasons we note that τ_1 and τ_2 should be non-zero.

This is the final model that will be used in fitting the yield curve. Following Gurkaynak et al. (2010, 74-75), the parameters in (3.8) are chosen so as to minimize the sum of squared deviations between the fitted and actual yields⁵:

$$\min \sum (y_t(m)^{NSS} - y_t(m)^{obs})^2 \tag{3.9}$$

As this cannot be done analytically we must again use numerical methods to find the minimum.

3.2 Inflation compensation and inflation expectations

The NSS-curve gives a continuous set of yields for any maturity, making the comparison of real and nominal yields now possible. The NSS-curve is fitted both to the index-linked and nominal government bonds, from which $y_t(m)^{real}$ and $y_t(m)^{nom}$ can be calculated. Following from the Fisher equation, which defines the ex-ante real interest rate as the nominal interest rate minus expected inflation, inflation compensation is simply defined as:

⁵ The authors use the deviation between actual and fitted *prices*, weighted by the inverse duration of each individual bond. The use of unweighted *yields* is in essence the same procedure. (Gurkaynak et al. 2006, 15)

$$\pi_t^E(m) = y_t(m)^{nom} - y_t(m)^{real} \quad (3.10)$$

In other words, inflation compensation gives the rate of inflation that if realized would give investors an identical return on an index-linked and nominal bond. (Gurkaynak et al. 2010, 76)

However inflation compensation is a direct estimate of inflation expectations only if investors are assumed to be risk neutral and do not require extra compensation for risks associated with real and nominal bond yields (Gurkaynak et al. 2010, 84). The consensus in the literature is that this is an unrealistic assumption and that there will be some premia attached to the yields. As bond yields reflect the real interest rate prevalent in the economy, a clear uncertainty relating to these yields is interest rate risk, hence they are likely to incorporate some form of interest rate risk premium. The assumption though is that this risk will affect nominal and real yields in equal measure so it will not be present in the inflation compensation. (Pflueger & Viceira 2011, 1)

The two risk premia that are commonly associated with the inflation compensation are inflation risk and liquidity risk premia. Inflation risk arises from the fact that the inflation rate at the time of maturity is unknown and so it creates uncertainty to the yield of the nominal bond and hence risk-averse investors will require a higher yield as compensation for this uncertainty. The liquidity premium follows from the poorer liquidity of the index-linked bonds compared with nominal government bonds due to the fact that the trading volume of the former is markedly lower than that of the latter (Pflueger & Viceira 2011, 1-2)

If we were to assume that the inflation and liquidity premia are constant over time, their presence in the inflation compensation would not necessarily be a problem, as the changes to the compensation would then reflect directly changes to inflation expectations, which is the primary interest in this thesis (Jochmann et al. 2008, 3-4). However this too is not a reasonable assumption. The uncertainty about future inflation is strongly affected by changes in the macroeconomic environment and it is not very credible to think that inflation risk in the Euro Area was the same say in 2004-2006 and 2009-2011.

3.2.1 Test of random walk hypothesis of inflation compensation

The presence of the inflation and liquidity risk premia, or whether they are constant over time, can be determined by looking at the volatility of the inflation compensation and more specifically comparing short-term and long-term volatility to determine whether inflation compensation is a random walk, as should be the case under rational expectations. We again follow Gurkaynak et al. (2010, 84-85) and look at the five-year and five-year minus one day forward rates of inflation compensation, denoted $\pi_t^f(5)$ and $\pi_t^f(5-)$ respectively. If inflation compensation follows a random walk, then $E_t(\pi_{t+1}^f(5-) - \pi_t^f(5)) = 0$ and $x_t = \pi_{t+1}^f(5-) - \pi_t^f(5)$ is a martingale difference sequence.

This can be tested using the variance ratio test developed by Lo and MacKinlay (1988) and following the presentation of Chen (2008, 98-99). The variance of differences of x_t taken at multiple intervals should be linear across the data series. In other words the variance of $x_t - x_{t-q}$ should be q times the variance of $x_t - x_{t-1}$. This is tested by comparing $1/q$ times the variance of $x_t - x_{t-q}$ to the variance of $x_t - x_{t-1}$. The Variance ratio is hence defined as:

$$VR(q) = \frac{\sigma^2(q)}{\sigma^2(1)} \quad (3.11)$$

where $\sigma^2(q)$ is $1/q$ times the variance of $x_t - x_{t-q}$ and $\sigma^2(1)$ is the variance of $x_t - x_{t-1}$. The null hypothesis is that $VR(q) = 1$.

The equations for $\sigma^2(q)$ and $\sigma^2(1)$ follow from the definition of variance:

$$\sigma^2(1) = \frac{1}{nq-1} \sum_{t=1}^{nq} (x_t - x_{t-1} - \hat{\mu})^2 \quad (3.12)$$

$$\sigma^2(q) = \frac{1}{m} \sum_{t=1}^{nq} (x_t - x_{t-q} - q\hat{\mu})^2 \quad (3.13)$$

where $\hat{\mu}$ is the average of x_t across the whole sample and $m = q(nq - q + 1)(1 - \frac{q}{nq})$. The null hypothesis is tested by the standard normal test statistic under the assumption of both homo- and heteroscedasticity. The test statistic is of the form:

$$Z(q) = \frac{(VR(q) - 1)}{\sqrt{\theta(q)}} \sim N(0,1) \quad (3.14)$$

The definition of $\theta(q)$ depends on the assumption of homo- or heteroscedasticity with:

$$\theta(q)^{ho} = \frac{2(2q - 1)(q - 1)}{3q(nq)} \quad (3.15)$$

assuming homoscedasticity and:

$$\theta(q)^{he} = \sum_{j=1}^{q-1} \left[\frac{2(q-j)}{q} \right]^2 \hat{\delta}(j) \quad (3.16)$$

assuming heteroscedasticity, with:

$$\hat{\delta}(j) = \frac{\sum_{t=j+1}^{nq} (x_t - x_{t-1} - \hat{\mu})^2 (x_{t-j} - x_{t-j-1} - \hat{\mu})^2}{\left[\sum_{t=1}^{nq} (x_t - x_{t-1} - \hat{\mu})^2 \right]^2} \quad (3.17)$$

3.2.2 Extracting liquidity and inflation risk premia

The next step in order to derive the true inflation expectations is to decompose the inflation compensation into the inflation and liquidity risk premia and inflation expectations components. The liquidity premium arises from the fact that inflation-linked bonds are not as liquid as nominal bonds, hence investors require a larger return as compensation. From this it follows that inflation-linked bonds have a higher yield relative to nominal bonds and inflation compensation is less than it would be without the liquidity premium. Conversely, inflation risk is present only in the nominal bonds again causing investors to require a higher yield

and widening inflation compensation. Combining these two effects yields that inflation compensation is decomposed into these elements as follows:

$$\pi_t^{Comp} = \pi_t^{Exp} + \pi_t^{RP} - \pi_t^{LP} \quad (3.18)$$

Following Gurkaynak et al. (2010, 86-88), the liquidity risk premium can be extracted using the following OLS regression:

$$\pi_t^{Comp} = \alpha + \beta X_t + \varepsilon_t \quad (3.19)$$

where X is a vector containing proxies for the liquidity of index-linked bonds. The fitted values of the regression represent the changes to inflation compensation due to changes in the liquidity of the inflation-linked and nominal government bonds.

The proxies for liquidity can include the difference of the respective daily average bid-ask spreads of the inflation-linked and nominal government bonds or the relative trading volume of the inflation-linked debt, as a percentage of all traded government debt of a particular country. These both measure the liquidity of the government bonds in the secondary markets. The bid-ask spread is the most commonly used measure of the liquidity of a financial asset as it incorporates both the direct and indirect transaction costs of trading activity and it will be the measure used in the analysis. (Sarr & Lybek 2002, 9)

The bid-ask spread S can be defined as the absolute difference of the bid price P_B and the ask price P_A at the end of a given trading day, or as a percentage of their average:

$$S = \frac{P_A - P_B}{\left(\frac{P_A + P_B}{2}\right)} \quad (3.20)$$

The advantage of using the percentage of the average price is that it standardizes the spread across high and low prices and across different markets and securities allowing for robust comparisons. (Sarr & Lybek 2002, 10). The final measure of the relative liquidity of the index-linked and nominal bonds used in the regression in equation 3.19 is the difference of their respective average bid-ask spreads:

$$X_t = \frac{1}{N} \sum_{i=1}^N S_{i,t}^{real} - \frac{1}{N} \sum_{j=1}^N S_{j,t}^{nom} \quad (3.21)$$

The remaining inflation risk premium can now be extracted, with the help of survey inflation expectations, using a state space model. This is a common way of estimating stochastic models with measurement errors, which is the case here, as we will be combining the daily compensation series with discrete quarterly survey estimations (Jalles 2009, 18).

The liquidity-adjusted measure for inflation compensation should now only contain the true inflation expectations and the inflation risk premium:

$$\pi_t^{Adj} = \pi_t^{Exp} + \pi_t^{RP} \quad (3.22)$$

The survey expectations of inflation are considered to be an approximation of the true expectations on the days that the survey is published (Gurkaynak et al. 2010, 88-89). The survey to be used in the analysis is the Survey of Professional Forecasters (SPF), which produces a quarterly forecast for inflation published on the ECB's website.

These survey expectations can be expressed in the following way:

$$\pi_t^{SPF} = \pi_t^{Exp} + u_t^{SPF} \quad (3.23)$$

where u_t^{SPF} is an independently and identically distributed measurement error. Equations (3.22) and (3.23) are the measurement equations and they are combined to form a state space form representation, where the vector $(\pi_t^E, \pi_t^{RP})'$ is the vector of the unobserved states. Again following from the rational expectations hypothesis we assume that long-run inflation expectations follow a random walk. We further assume that the inflation risk premium follows an AR(1) process. The system of equations can then be written in the form:

$$\begin{pmatrix} \pi_t^{Exp} \\ \pi_t^{RP} \end{pmatrix} = \begin{pmatrix} 1 & 0 \\ 0 & \varphi \end{pmatrix} \begin{pmatrix} \pi_{t-1}^{Exp} \\ \pi_{t-1}^{RP} \end{pmatrix} + \begin{pmatrix} v_{1t} \\ v_{2t} \end{pmatrix} \quad (3.24)$$

where v_{1t} and v_{2t} are independently and identically distributed, mutually uncorrelated random variables, with mean zero and variances σ_1^2 and σ_2^2 respectively. σ_1^2 is set to the data and the rest of the parameters are estimated using the Kalman filter. As a result the true inflation expectations π_t^{Exp} can be produced. (Gurkaynak et al. 2010, 89-90)

The Kalman filter is a method introduced by Kalman (1960) as a recursive solution to produce a linear filtering of noisy discrete data. It is in the form of an algorithm consisting of a prediction for the new state of the model and a correction component that minimizes the difference of that prediction and the actual new state when it is confirmed. The filter assumes that the state vector and error terms follow a normal distribution which allows for the estimation of the unknown parameters of the system to be estimated by maximum likelihood. (Jalles 2009, 18)

Following the presentation of Harvey and Shephard (1993, 267-270) an observed multivariate time series y_t of N elements is related to a $p \times 1$ state vector α_t by the measurement equation:

$$y_t = Z_t \alpha_t + \varepsilon_t \quad (3.25)$$

where Z_t is a non-stochastic $N \times p$ matrix, ε_t is a $\sim N(0, H_t)$ distributed $N \times 1$ error vector and H_t is a $N \times p$ variance-covariance matrix. We note that equation (3.25) corresponds to equations (3.22) and (3.23), with $y_t = (\pi_t^{Adj}, \pi_t^{SPF})'$ and $\alpha_t = (\pi_t^{Exp}, \pi_t^{RP})'$. As N and p both equal 2, Z_t and H_t are 2×2 matrices of the form $\begin{pmatrix} z_{11} & z_{12} \\ 1 & 0 \end{pmatrix}$ and $\begin{pmatrix} 0 & 0 \\ 0 & h_{22} \end{pmatrix}$ respectively.

In the state space model α_t is allowed to change through time according to the transition equation:

$$\alpha_t = T_t \alpha_{t-1} + \omega_t \quad (3.26)$$

where T_t is a non-stochastic $p \times p$ matrix, ω_t is a $\sim N(0, Q_t)$ distributed $N \times 1$ error vector and Q_t is a $p \times p$ variance-covariance matrix. Again we note that (3.26) corresponds to (3.24). T_t and Q_t are 2×2 matrices of the form $\begin{pmatrix} 1 & 0 \\ 0 & t_{22} \end{pmatrix}$ and

$\begin{pmatrix} q_{11} & 0 \\ 0 & q_{22} \end{pmatrix}$ respectively. All in all, depending on the specification of the variances of the unobserved states in Q_t , the parameters to be estimated in the model are $z_{11}, z_{12}, h_{22}, t_{22}, q_{11}, q_{22}$.

To begin with, some assumptions about the initial state of the model must be made. We assume that α_0 is distributed $\sim N(a_0, P_0)$. Also α_0 is assumed to be uncorrelated with ε_t and ω_t . The Kalman filter provides the mean of the conditional density of $\alpha_t | Y_t$, where Y_t is the set of information up until time t $\{Y_0, \dots, Y_t\}$ and the knowledge of the initial state α_0 .

Combining equations (3.25) and (3.26) and taking into account the above assumptions yields the distributions for α_t and y_t and the three sets of equations necessary to perform the filter:

$$\begin{pmatrix} \alpha_t \\ y_t \end{pmatrix} | Y_t \sim N \left(\begin{pmatrix} a_{t|t-1} \\ Z_t a_{t|t-1} \end{pmatrix}, \begin{pmatrix} P_{t|t-1} & P_{t|t-1} Z'_t \\ Z_t P_{t|t-1} & F_t \end{pmatrix} \right), \quad (3.27)$$

where

$$a_{t|t-1} = T_t a_{t-1}, \quad P_{t|t-1} = T_t P_{t-1} T'_t + Q_t \quad (3.28)$$

are the prediction equations and:

$$v_t = y_t - Z_t a_{t|t-1}, \quad F_t = Z_t P_{t|t-1} Z'_t + H_t \quad (3.29)$$

are the one-step ahead forecast equations, where v_t is the one-step ahead forecast error and F_t is its variance, hence $v_t \sim N(0, F_t)$. From the conditions placed on multivariate normal distributions we get:

$$\alpha_t | Y_t \sim N(a_t, P_t), \quad (3.30)$$

where

$$a_t = a_{t|t-1} + P_{t|t-1} Z'_t F_t^{-1} v_t, \quad P_t = P_{t|t-1} - P_{t|t-1} Z'_t F_t^{-1} Z_t P_{t|t-1} \quad (3.31)$$

are the updating equations.

The filter is run so that some initial values are assumed for the unobserved variables and the rest of the parameters are estimated by maximum likelihood. One can choose how closely the filter should follow either the observations or the predictions, which is measured by the Kalman gain. With a high gain the filter places more emphasis on the measurements and conversely a low gain more emphasis to the model predictions. (Bishop & Welch 1997, 3)

3.3 Event study

The approximation for the inflation expectations extracted from the inflation compensation of government bond yields can now be used to gauge the effectiveness of monetary policy in affecting the expectations of future inflation. This is done in an event study context, where the evolution of inflation expectations is looked at before and after an event, in this case relating to new information about ECB policy.

The first approximation of this effect is done simply by regressing the inflation expectations extracted in section 3.2.2 on the key ECB policy rate:

$$\pi_t^{Exp} = \alpha + \beta r_t^{ECB} + \varepsilon_t \quad (3.32)$$

This will give some indication of the direction of correlation between expected inflation and monetary policy.

To gauge the longer-term effects of changes to the ECB's rate we run a basic vector autoregression (VAR) model of the form:

$$\begin{pmatrix} \pi_t^{Exp} \\ r_t^{ECB} \end{pmatrix} = \sum_{j=1}^p \Phi_j \begin{pmatrix} \pi_{t-j}^E \\ r_{t-j}^{ECB} \end{pmatrix} + \begin{pmatrix} \varepsilon_{1t} \\ \varepsilon_{2t} \end{pmatrix} \quad (3.33)$$

The optimal lag length p will be determined by looking at the standard information criteria. The main objective is to look at how shocks manifesting through ε_{1t} and ε_{2t} affect the ECB rate and inflation expectations respectively over an extended period of time.

However the changes of the key ECB rate are not the only channel of monetary policy, as for example the decision to keep rates unchanged can be a policy measure in the same way as changing the rate is. This effect is not captured in the regressions in (3.32) and (3.33). Hence it is necessary to look at a model containing time dummies relating to the release of news about monetary policy conditions, whether that is the changing the key rate, keeping it unchanged or merely informing about the future strategy of the ECB. Announcements such as Mario Draghi's famous "anything it takes to save the Euro" can have a much larger effect on the expectations of financial market participants than a rate change, which is usually expected and hence does not necessarily have much of an impact when confirmed.

We run two time dummy regressions, one where the dummies relate to dates of meetings of the ECB governing council when a change to the ECB's rate were decided, and the second for dates of announcements of unconventional monetary policy measures by the ECB. The regressions are of the form:

$$\Delta\pi_t^{Exp} = \alpha + \beta D_t + \varepsilon_t \quad (3.34)$$

where D_t is the dummy variable with the value 1 for days of the governing council meetings or unconventional monetary policy announcements and 0 for all other days. The dependent variable in these regressions is the first difference of expectations, rather than the level, in order to see if there is a consistent change in expectations on the day of the announcements.

The final section of the event study looks at what effect of the ECB's expanded asset purchase programme has had so far on inflation expectations. This is done by looking at the trend of expectations for a period before the announcement and after. We set the periods to approximately 6 months, or 120 trading days, which should be enough to establish the pre-event trend. The regressions are of the form:

$$\pi_t^{Exp} = \alpha + \beta\tau_t + \varepsilon_t \quad (3.35)$$

where τ is the trend for the pre- and post-event periods.

4. Data

The main data used in the thesis is the prices and maturities of inflation-linked bonds of Eurozone governments. For liquidity purposes the focus is on large Eurozone countries, whose debt is traded most frequently. The inflation-linked bond market is a comparatively recent development, with most of the first debt issues having taken place in the late 1990s. The first bonds indexed to the Euro area inflation index, the HICPxT, were issued soon afterwards in the early 2000s.

The main idea of an index-linked bond is that the principal is adjusted daily to the reference index so that at maturity investors retain their purchasing power. The coupon payments are calculated from this index referenced principal, hence investors' income from holding the bond is also protected against inflation. The principal payed at maturity cannot fall below the initial principal, hence index-linked bonds further protect against deflation. (Tuckman & Serrat 2011, 7)

The data on the French government bonds, including the price and bid and ask rates used for the liquidity premium extraction, are from Thomson Reuters Datastream. All prices are mid rates of the closing quotes of each trading day. The data on realized HICPxT inflation is from Eurostat and the Survey of Professional Forecasters is from the ECB.

4.1 French government inflation-linked bonds

The main data used to build the model is on French government debt, as the French government has as of October 2015 eight outstanding bonds indexed to the HICPxT with two bonds already matured. Other possibilities could be the debt of other core Euro area countries such as Germany and Italy. The German government has issued only five bonds linked to the HICPxT however and the lower credit rating of Italian bonds would mean that the implied credit risk would have to be extracted during the analysis. Naturally the more outstanding bonds there are for a

given day, the better the fit and shape of the yield curve for that day. As will be seen later on, nine bonds produce a remarkably accurate yield curve.

The French government issues two kinds of inflation-linked debt, based on French and Euro area inflation. We use the latter of these, which is called OAT€i (standing for Obligations Assimilables du Trésor € inflation). The first OAT€i was issued in October 2001 and apart from the years 2003, 2006, 2008 and 2009 there has been a new issue every year since. The characteristics of the individual currently outstanding OAT€i issues are described in Table 1.

Table 1. Characteristics of outstanding OAT€i issues.

Maturity	2018	2020	2022	2024	2027	2030	2032	2040
Real Coupon	0,25 %	2,25 %	1,10 %	0,25%	1,85%	0,70%	3,15 %	1,80%
Base index year	2011	2003	2009	2012	2010	2013	2002	2006

Note: Both the maturities and base index dates are the 25th July of each year. (AFT, 2014)




Figure 1. Issued OAT€i securities.

To further demonstrate the range of maturities available for the analysis, as in Gurkaynak et al. (2010), all the issues of OAT€i are shown graphically in Figure 1 including bonds that matured in 2012 and 2015. Each line corresponds to an individual issue with the dates on the horizontal axis and years to maturity in years on the vertical axis.

From this graphical representation it is quickly observed that most of the bond issues have taken place in more recent years. Again as the robustness of the yield curve is the better the more individual securities there are available to form it, the analysis will inevitably have to be focused on the last couple of years. That is not necessarily a bad thing, as said before, the issue of very low medium term inflation expectations in the Euro area has been most relevant during the last few years.

The OAT€is have a nominal principal of 100€ and pay a coupon that is a fixed percentage of the index-linked principal. The principal is linked to a daily reference that is calculated by linear interpolation for day j and month m using the following formula:

$$reference_j = IPC_{m-3} + \frac{nj - 1}{NJ_m} \cdot (IPC_{m-2} - IPC_{m-3}) \quad (4.1)$$

Where IPC is the value of the HICPxT for months $m-2$ and $m-3$ respectively, nj is the number of day j in month m and NJ is the total number of days in month m .

The indexation lag is due to the fact that the reference applicable for the principal is the HICPxT of the third previous month to the month of the payment. For example if the payment of the principal were to fall on June 1 the reference index would be the HICPxT for March. For any other day the reference is the one given by (4.1). The indexation lag is due to the fact that the value of the HICPxT for a given month is subject to possible revisions by Eurostat at the end of the subsequent month. (Agence France Tresor, 2015). Given that we eliminate bonds with maturities of less than two years, the indexation lag should not affect the results of the analysis.

4.2 ECB policy announcements

The European Central Bank is the central bank that conducts monetary policy for all 19 member states of the Euro area. The tasks of the European System of Central

Banks (ESCB) and the Eurosystem are laid down in the Treaty on the Functioning of the European Union. Article 127(1) of this treaty defines that the primary objective of the Eurosystem and hence the ECB "shall be to maintain price stability". (European Central Bank, 2011)

The main decision-making body of the ECB is the Governing Council which consists of the six members of the ECB's executive board and the governors of the 19 national central banks of the Euro area. In order to specify its objective more precisely, the Governing Council announced the following quantitative definition in 1998: "Price stability shall be defined as a year-on-year increase in the Harmonised Index of Consumer Prices (HICP) for the euro area of below 2%. Price stability is to be maintained over the medium term". Following a thorough evaluation of its monetary policy strategy in 2003, the Governing Council further clarified that, within the definition, it aims to maintain inflation rates "below, but close to, 2% over the medium term". (European Central Bank, 2011)

The Governing council meets twice every month, with decisions regarding monetary policy taken during the first of these meetings and the second one covering other tasks and responsibilities of the ECB. The monetary policy decisions made every six weeks are explained in detail at a press conference taking place at the same intervals. It is at these meetings that the Governing Council concludes its key interest rate decisions. (European Central Bank, 2011)

There are three main policy interest rates that the ECB controls. They are the interest rate for the main refinancing operations (MRO), which provide the bulk of liquidity to the banking system. Secondly the rate on the deposit facility, which banks may use to make overnight deposits with the ECB. Third is the rate on the marginal lending facility, which offers overnight credit to banks from the Euro area. The evolution of the three rates from the inception of the monetary union in 1999 is presented in Figure 2. (European Central Bank, 2011)

As the MRO constitute most of the refinancing of the financial sector and are essential to the implementation of the ECB's open market operations, it is the rate that has the largest effect on the Euro area economy and is the one watched most closely by financial market participants. Hence it will be the one used in the first regression of the event study section. Though as can be seen from the graph, all three rates are usually changed at the same time and by the same magnitude, so any one of the three would apply just as well.

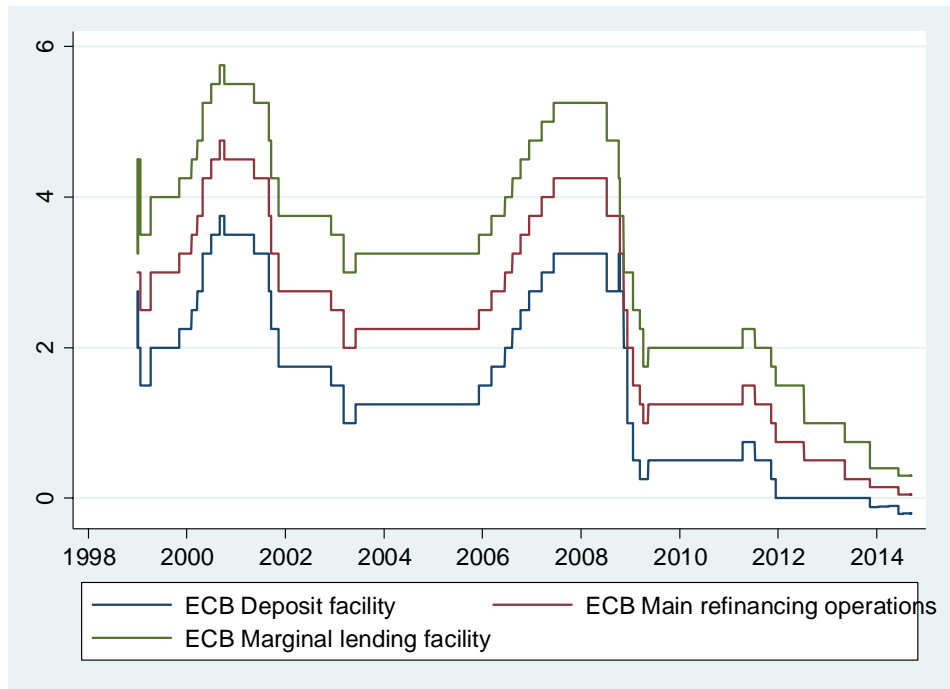


Figure 2. The ECB's key policy rates.

4.2.1 The Harmonized Index of Consumer Prices

As both the inflation protected government bonds and the Governing Council of the ECB refer to inflation in terms of the Harmonised Index of Consumer Prices for the euro area, it is worth saying something about the composition and compilation of this index. First of all it is worth noting that the ECB's price stability target is tied to the full index (HICP), whereas the reference index for the index-linked French government bonds excludes tobacco (HICPxT). The latter index is the one used for all Euro area inflation-linked financial products and has become the one most looked at by market participants. As tobacco has a weight of only 2,4 % in the HICP index, the difference in inflation rates between the indices is negligible.

The HICP index is compiled by Eurostat in collaboration with the individual national statistics institutes of the Eurozone countries. The ECB is also directly involved in the compilation of the index, given its use in monetary policy. Data on HICP inflation is available from 1995 onwards, with not entirely comparable back data available from 1990. (European Central Bank, 2011)

The different Eurozone countries have different weights in the index corresponding to population and the individual items have different weights depending on the

country. For the Euro area as a whole, goods make up 56,5 % and services 43,5 % respectively of the index as of November 2015 . These are broken down into smaller categories to identify the different economic factors that impact on consumer price developments. Energy prices are strongly affected by the oil price for instance and foods are divided into processed and unprocessed foods given that weather patterns have more of an impact on the latter. (European Central Bank, 2011)

5. Estimation results

The empirical estimation is conducted by following the steps laid down in chapter 3. Firstly the continuous set of yields are obtained by fitting a yield curve for the individual bond prices of both the nominal and index-linked bonds for each trading day. The measure of inflation compensation is calculated as the difference between the yields of these bonds for each maturity. The inflation compensation is then tested to see if it represents true inflation expectations, in other words if it follows a random walk. In the event that compensation does not represent expectations, the remaining liquidity and inflation risk –premia are extracted. Finally the hopefully now accurate measure of inflation expectations is regressed on the ECB key interest rate and on time dummies for ECB policy announcements. The interaction between rate and expectations is further looked at with a VAR model and the chapter concludes with an appraisal of the effects of the EAPP.

5.1 Yield curve estimation

The first step of the empirical analysis is fitting a yield curve to the observed yields of the French government bonds. This is done by utilizing the Nelson-Siegel-Svensson yield curve expressed in equation (3.8). The actual minimization of (3.9) is done by the Solver function in Excel. The only practical issue in calibrating the curve is choosing the initial parameter values for β_0 , β_1 , β_2 , β_3 , τ_1 and τ_2 ⁶. For most of the sample we choose initial values of $\beta_0 = 0,02$, $\beta_{1,2,3} = -0,05$ and $\tau_{1,2} = 2$ as they provide a smooth traditional yield curve shape. For some days, however, the NSS yield curve gives quite erratic values for short maturities, especially for the

⁶ The Solver function uses a generalized reduced gradient algorithm for the optimization problem, which comes with the drawback that the algorithm stops when it finds the closest minimum to the initial values, which might not be the global optimum of the function (Frontline Solvers, 2015). This results in the squared error between the observed and estimated yields being slightly larger if the algorithm finds only a locally optimal solution. Given the already high accuracy of the method, this small loss of accuracy is worth accepting since the algorithm is highly efficient, which given the large dataset is desirable.

time period of late 2008 during the height of the financial crisis, when the prices of the bonds were volatile and there were fewer outstanding bonds to use for the yield curve fitting. For these days the curve can be made to behave better by choosing slightly different individual initial parameter values. The average cumulative squared error between the NSS model yields and the observed yields (equation 3.9) is approximately $1,3 \times 10^{-6}$ for the index-linked bonds and $1,1 \times 10^{-6}$ for the nominal bonds, for the whole sample, so the accuracy of the model is very high.

Gurkaynak et al. (2010, 76) note that due to the indexation lag in inflation protected government bonds, the price movements of bonds with less than 18 months to maturity can be large and subsequently they drop these bonds for their estimation. Also, as discussed in section 2.2, the elimination of short-term maturities from the curve reduces the possibility of seasonal fluctuations distorting the results. Therefore, we chose not to include any bonds with less than two years to maturity from the construction of the yield curves.

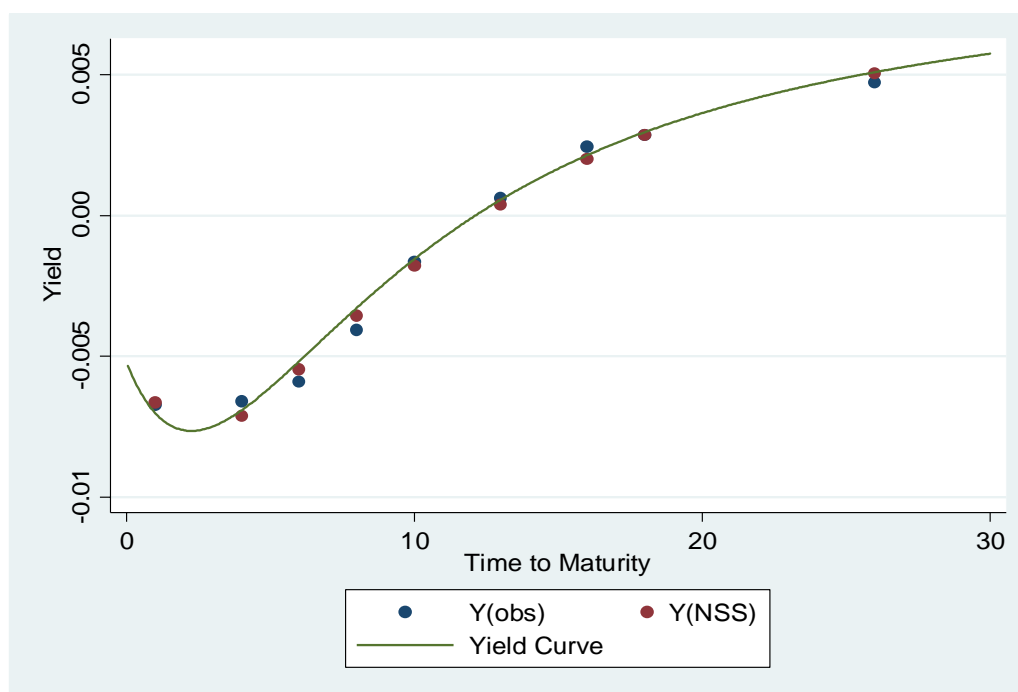


Figure 3. French government index-linked yield curve 13.11.2014

Figure 3 presents the French government index-linked yield curve for 13 November 2014. As can be seen, the NSS yield curve does a good job fitting the data, both at short and long maturities, providing a traditional yield curve shape. The overall

low level of yields and in particular the negative yields for the shorter maturities arise from high prices of the bonds reflecting a strong demand for them following extensive bond purchases by the ECB. The negative yields at short maturities also correspond to the general low level of interest rates. Naturally the level of yield for these bonds are lower than those on nominal bonds, as seen from Figure 4, due to inflation being eliminated from the index-linked bonds by construction.

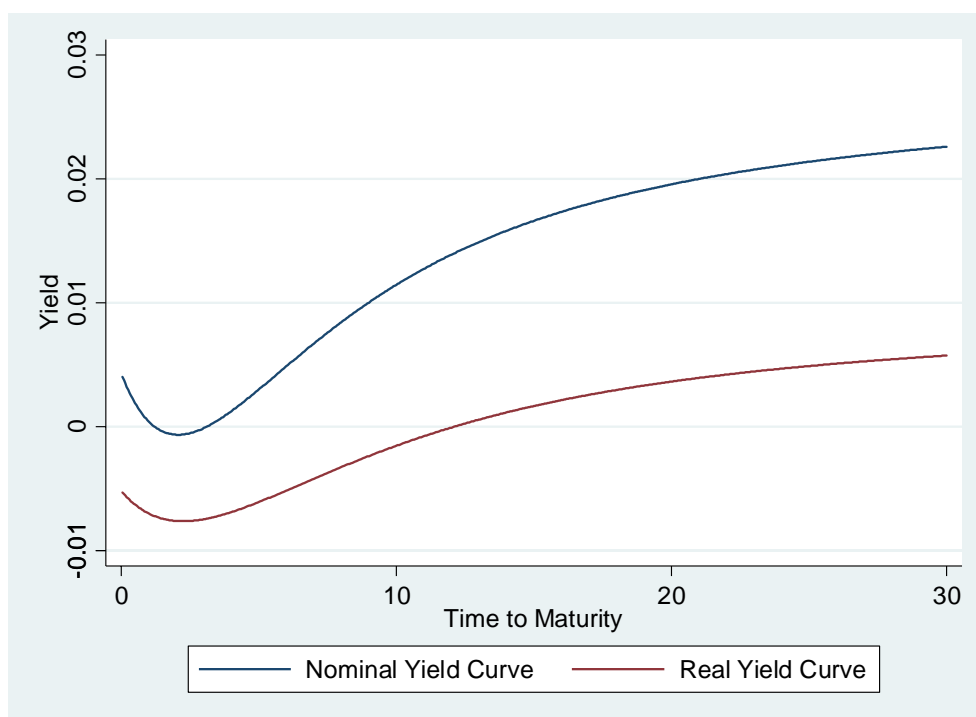


Figure 4. Inflation compensation on 13.11.2014

In Figure 4 are the respective yield curves for the nominal and index-linked French government bonds on 13.11.2014. The inflation compensation is defined as the difference between the two yields, in the figure it is the vertical distance between the two curves. As can be seen the two curves have roughly the same shape, with the amount of compensation required by investors smaller at short maturities and growing towards longer maturities reflecting the greater uncertainty in the longer run level of inflation.

From these two curves the measure of inflation compensation for specific maturities, for example 5, 10 or 15-years, can now be calculated for the whole sample period. Initially at least this period is chosen to begin from the issue of the 1,80 % OAT€i maturing in 2040, which provides at least five bonds for both the nominal

and real yield curve estimations on any given day. Hence the time period for the analysis runs from 14.3.2007 – 18.8.2015. This is convenient for a robust analysis as it contains a brief period of time before the financial crisis, the crisis period itself and the post crisis period. As this study focuses on the Euro area of particular interest are the sovereign debt crises of the early 2010s, which naturally are also contained in the sample period.

5.2 Inflation expectations

The measure of inflation compensation obtained from the yield curves is presented in Figure 5. The compensation is calculated for maturities of 5, 10 and 15-years giving an estimate for the expected level of inflation at those points of time in the future. We will use the five-year ahead compensation for the analysis, as it is consistent with the longest forecast given by the Survey of Professional forecasters.

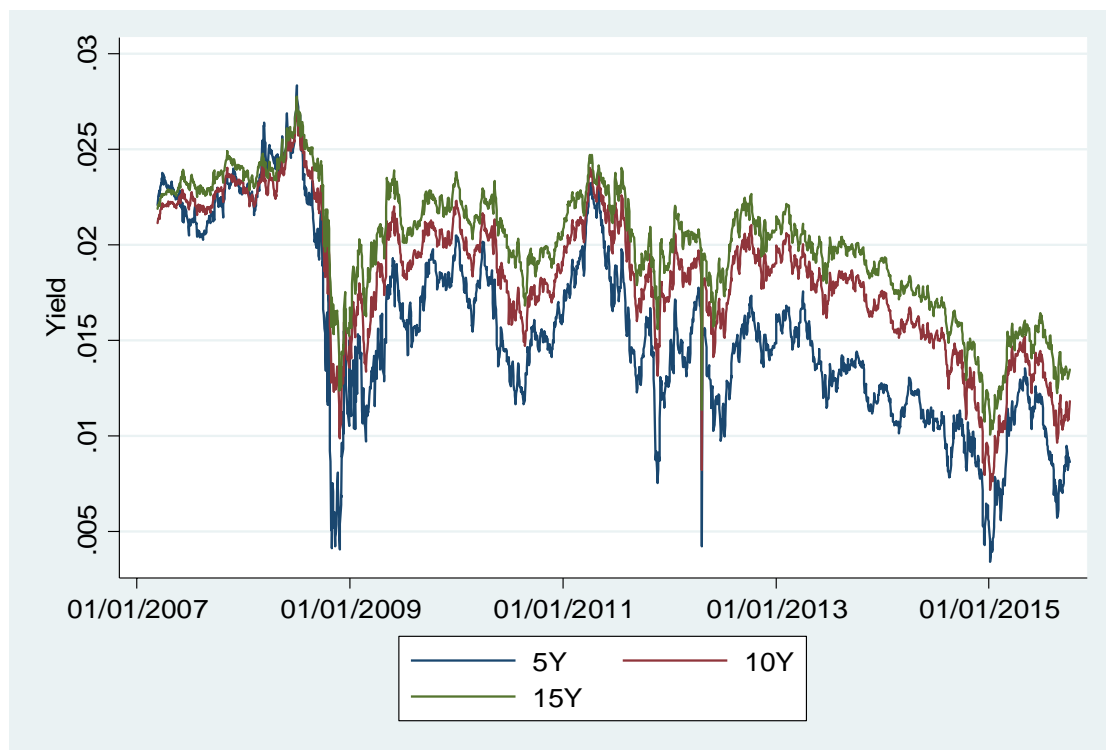


Figure 5. Euro area inflation compensation 14.3.2007-18.8.2015

Several observations can already be made from this measure of inflation expectations, mainly concerning the relative magnitudes of the different maturities assum-

ing that the possible liquidity and inflation risk premia have an equal effect on all of them. Of particular interest are the developments for the period 2007-2009 as the data for the liquidity adjustment used later on is restricted to August 2009 and so this figure will give some indication of what happened to expectations before that date.

Firstly the compensation for the three maturities was remarkably converged during the period before the global financial crisis began in earnest in the autumn of 2008. This is a period that might be characterized as “normal times”. Secondly, the crisis caused a clear collapse in the compensation for all time horizons, but afterwards the outlook for shorter-term inflation compensation began to disperse compared to the longer-term. This would seem to point to the growing deflationary fears at least for the shorter-term and overall uncertainty about the future path of inflation. Also looking at this figure, given the strong variation of even the 15-year ahead compensation, it would seem that Euro area inflation expectations are poorly anchored. This is however too early to judge at this point, as much of the variation can be due to the liquidity and inflation risk premia.

Thirdly, probably the most significant event during the sample period from a monetary policy perspective was the announcement of the ECB’s quantitative easing programme in early 2015. The effect of the announcement can be seen as a clear bump in the inflation compensation, but interestingly it seems to have only had a temporary effect, with the compensation falling with uncertainty over Greek debt arrangements during the summer of 2015. We will analyse the effects of the asset purchase program in more detail in the event study in section 5.3.

A final note is that there is a clear outlier in the data on the date of 23.4.2012 for the prices of the nominal bonds causing an extreme deviation in the inflation compensation. This date is also an outlier in the data on liquidity used later on. The observation corresponds to the French presidential elections that took place on the previous day. As the elections have no direct consequence on Euro area inflation expectations or the monetary policy of the ECB we will omit this observation from the analysis as it would otherwise cause some distortions to the results.

5.2.1 Variance ratio test

This measure of inflation compensation equals inflation expectations only if investors are risk neutral. If this is the case, then under the rational expectations hy-

pothesis inflation compensation should follow a random walk as discussed in section 2.2. To test this we run the Lo-Mackinlay variance ratio test described in section 3.2.1, by looking at whether the variance of inflation compensation over a longer period is proportionally larger than the daily changes to it.

Following Gurkaynak et al. (2010, 84-85) we calculate the five-year and five-year-less-one-day measures of inflation compensation from the yield curves assuming 260 business days a year. To look if the expectation of inflation in five-years today is the same as it will be tomorrow, as should be the case by the law of iterated expectations, we define $x_t = \pi_{t+1}^f(5-) - \pi_t^f(5)$, where 5- denotes five-years-less-one-day. The task is now to compare whether the variance of x_t is proportional to the cumulative variance of x_t over a period of one, three and six months. In other words if the variance of $\sum_{j=1}^k x_{t+j}$ is k times the variance of x_t , with k equal to 22, 66 and 132 assuming 22 business days in a month.

Table 2. Variance ratio test of the random walk hypothesis of the inflation compensation.

	Standard deviation	Variance ratio	Test statistic	P-value
1 day	0,04 %			
1 month	0,21 %	0,050	-3,99	0,000
3 month	0,33 %	0,019	-2,59	0,001
6 month	0,75 %	0,010	-2,03	0,042

Notes: The null hypothesis is $VR = 1$. The variance ratio test statistic is assumed to have a standard normal asymptotic distribution and it is computed allowing for time-varying conditional heteroscedasticity.

The results of the variance ratio test are presented in Table 2. The variance ratio is calculated using the expression in equation (3.11). The test statistics used in the table are heteroscedasticity robust, as expressed in equations (3.16) and (3.17). As can be seen the variance ratios of the respective months are significantly different from 1 at the 1 % level for k equal to 22 and 66, and at the 5 % level for k equal to 132. From this we can draw the fairly robust conclusion that inflation compensation does not follow a random walk. More precisely it seems that the daily changes of inflation are too large compared to longer-term variation, given that

the test statistics are negative. We also run the variance ratio test using the homoscedasticity consistent test statistics (equation 3.15). In that case all the variance ratios are significantly different from 1 at the 1 % level, which would support the rejection of the null hypothesis, although the assumption of heteroscedasticity seems clear from a graphic analysis of x_t .

5.2.2 Liquidity premium

Now that we have confirmed the likely existence of the liquidity and inflation risk premia in the inflation compensation, we can go about trying to extract them to produce the true measure of expectations. Firstly the liquidity effect on inflation compensation is gauged as presented in section 3.2.2.

The results of the regression in equation (3.19) are presented in Table 3. As can be seen the changes in the difference of the liquidity of the inflation-linked and nominal bonds, as measured by their respective average bid-ask spreads, has a negative effect on the inflation compensation that is statistically significant at the 5 % level. In other words when the spread difference grows, the compensation narrows, which is consistent with the idea of a liquidity premium. However it is worth noting that the changes in liquidity explain only a very small amount of the overall variation of inflation compensation.

Table 3. Regression of liquidity premium in inflation compensation

Regressor	5-year inflation compensation
Bid-ask spread	-0,108** (0,048)
Intercept	0,0144
R ²	0,0024

Notes: The heteroscedasticity robust standard error is in parentheses. The number of observations is 1 612. ** Significant at the 5 % level.

In fact if we correct for possible autocorrelation in the regression errors by using Newey-West error terms for lags greater than one, the coefficient for the bid-ask spread becomes insignificant. Hence, it is possible that the liquidity premium is

not an issue for the French government inflation-linked bonds. Looking at the respective bid-ask spreads, the spread for the inflation-linked bonds is consistently higher compared to the nominal spread, so we will proceed with adjusting the compensation to liquidity effects. In any case, the changes to the inflation compensation caused by adjusting for liquidity are fairly cosmetic.

The fitted values of the regression represent the effect of liquidity on five-year inflation compensation, but do not give the level of the liquidity premium. Following Gurkaynak et al. (2010, 88) we normalize the fitted value to zero at its lowest point in time, which is on 20 September 2011, giving an approximation for the level of the liquidity premium, relative to that point in time.




Figure 6. Estimated liquidity premium of French government 5-year index-linked bonds

The estimated liquidity premium is expressed in terms of the yield compensation in Figure 6. The sign of the liquidity premium is changed from the negative of the regression to positive in order to show the extra yield demanded by investors to hold the inflation-linked bonds. This presentation confirms the low impact of the liquidity; for most of sample period, liquidity only adds approximately 5-10 basis points to the yield of the inflation-linked French government bonds.

Overall, the liquidity premium has been fairly steady during the sample period, with some spikes. Of interest is again the effect of the ECB's asset purchase programme at the end of the sample. Reflecting the initial effect on the compensation, the initiation of the programme of buying sovereign bonds looks to have added to the relative liquidity of the inflation-linked bonds, before the effect is reversed during the summer of 2015.

5.2.3 Inflation risk premium

From equation 3.18 it follows that the measure for the adjusted inflation compensation that contains only the true expectation and the inflation risk premium is the sum of the inflation compensation and the liquidity premium: $\pi_t^{ADJ} = \pi_t^{COMP} + \pi_t^{LP}$. This is the measure that is then used in equation (3.22). Using this adjusted inflation compensation and the point estimate for future inflation provided by the Survey of Professional Forecasters, we can now estimate the state-space model described in section 3.2.2.

The SPF forecasts are made only on a quarterly basis, and as the data on liquidity is only from 2009 onwards, these combine to make the estimation of the model decidedly tricky. We can either use the adjusted compensation and impose restrictions on the covariance structure of the model, or assume that the liquidity premium is insignificant and use the original measure of inflation compensation from March 2007 onwards without any restrictions on the covariance matrix.

Beginning with the liquidity-adjusted compensation, the model to be estimated is described in equations (3.22) – (3.24). The coefficient of the inflation expectations on lagged expectations in equation (3.24) is constrained to 1 as per the assumption that it follows a random walk. The coefficient on expectations in equation (3.23) is likewise constrained to 1 following from survey expectations being a noisy measure of true expectations. To be able to estimate the model using only the 24 survey observations the variance-covariance matrix Q in equation (3.26) is constrained to an identity matrix, leaving only four parameters to be estimated, $z_{11}, z_{22}, h_{22}, t_{22}$ as described in equations (3.25) and (3.26).

The estimation of the model using the unadjusted compensation follows an otherwise identical procedure, except that no structure is imposed on Q and the variance of the error term of the inflation expectations in equation (3.24) is set to the variance of actual HICPxT inflation. This latter procedure is done by fitting an un-

observed components model to the HICPxT data in the spirit of Stock and Watson (2007, 16), who add stochastic volatility to the model and find it to provide good forecasts for inflation. For simplicity we omit the stochastic volatility and estimate the variance of inflation with the basic random walk model:

$$\pi_t^{HICP} = \pi_{t-1}^{HICP} + \eta_t \quad (5.1)$$

From this we get that the variance of η_t is approximately 0,05. The element q_{11} in matrix Q in equation (3.26) is then constrained to this value when estimating the state-space model.

Though in theory the state-space model is able to work with missing observations, it would seem that the amount of missing data in the quarterly survey expectations, if they are treated as daily observations, is too much for an algorithm to estimate the unobserved state variables. This means that we are only able to run the state-space model for the quarterly data and hence the true expectations and risk premium are estimated only at a quarterly frequency, which severely limits the usefulness of the analysis. This problem could possibly be overcome by using a longer sample period or with the addition of another set of observed survey expectations. Unfortunately, neither of these measures was possible for this study.

The addition of another measure of survey expectations would be welcome also to compensate for the fact that the Survey of Professional Forecasters, conducted by the ECB, gives a very static measure of long-term expectations, which closely follows the ECB's own target of inflation of less than but close to 2 %. The level of the true expectations is identified by the survey expectations, given that they are assumed to represent a noisy measure of the true expectations. This means that the true expectations estimated from the model are very flat around that 2 % mark for the whole sample period. It then comes down to which measure is given more credibility; the market derived inflation compensation or the survey expectations. Given the communication of the ECB referring to historically low inflation expectations and the fairly drastic measures it has taken in order to raise expectations, it seems likely that the compensation, which as seen from Figure 5 dipped below the 0,5 % mark in late 2014, is closer to reality. True we are dealing with the longer-term 5-year ahead expected level of inflation, but even that measure should exhibit

a clear deviation from the 2 % level to warrant the extensive asset purchase programmes.

Figure 7 depicts the final estimated inflation expectations derived from both the liquidity-adjusted and non-adjusted measures of the 5-year inflation compensation. As can be seen the two measures follow each other closely, with the liquidity-adjusted expectations having a slightly higher volatility. This is another indication of the limitations placed on the state-space model by the lack of observations, as adjusting for the liquidity premium should in fact decrease the volatility of expectations. The fact that the two measures follow each other closely allows us to use them interchangeably, which enables the use of the longer time series of the non-adjusted compensation in the later analysis.

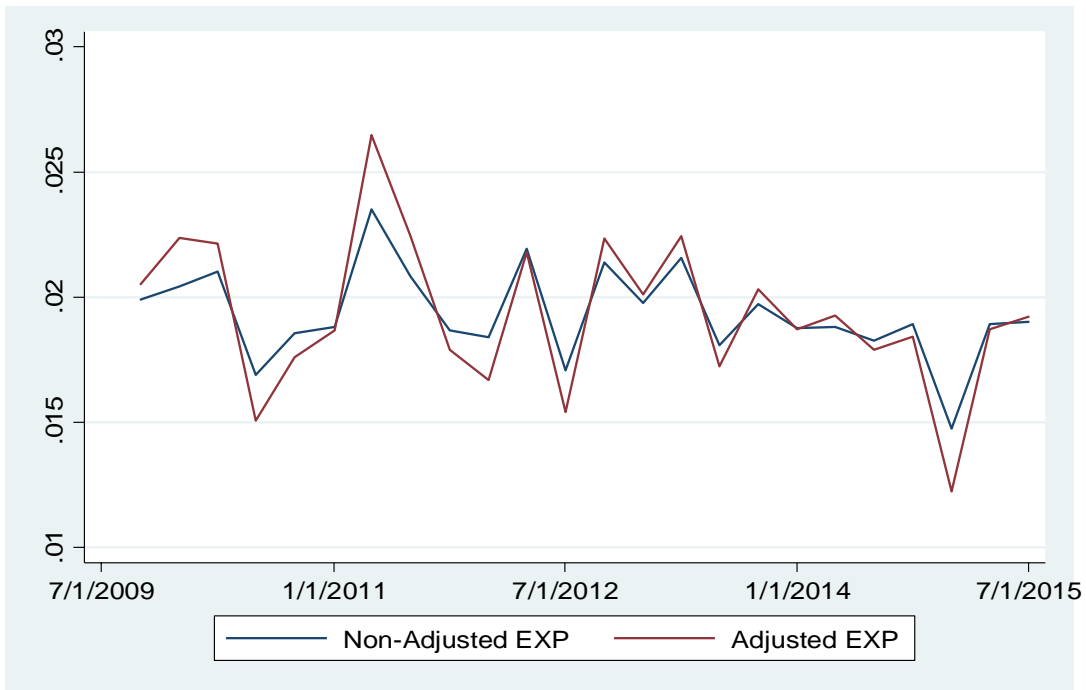


Figure 7. Estimated 5-year inflation expectations from liquidity-adjusted and non-adjusted inflation compensation.

When extracting the risk premium it turns out that fixing the variance of inflation expectations to the variance of actual HICPxT inflation is not the best specification owing again to the limitations of the data. When the variance of expectations is fixed to 0,05 the estimated variance of the risk premium is 0,42, which is unrealistically high. Furthermore, the sign of the coefficient on the risk premium is nega-

tive, which is against the assumption that a rise in the risk premium increases inflation compensation (equation 3.22).

We therefore run the state-space model again using the non-adjusted compensation and impose no restrictions on the variances in the model. This means that q_{11} and q_{22} in equation (3.26) are both estimated in the model. The resulting decomposition of the compensation into the expectations and risk premium is presented in Figure 8. As can be seen, most of the volatility of the compensation is captured in the risk premium, which is due to the very low volatility of the SPF survey expectations to which the estimated expectations are linked.

This also mainly explains why the risk premium is negative for the latter part of the sample. It is also possible however, that if the true expectations are as anchored as the survey forecasts suggest, the negative risk premium would be consistent with the growing deflationary fears of the past few years. Again though, both measures of the true inflation expectations and inflation risk premium, extracted in this way, are so crude that we should avoid drawing too many meaningful conclusions from them.

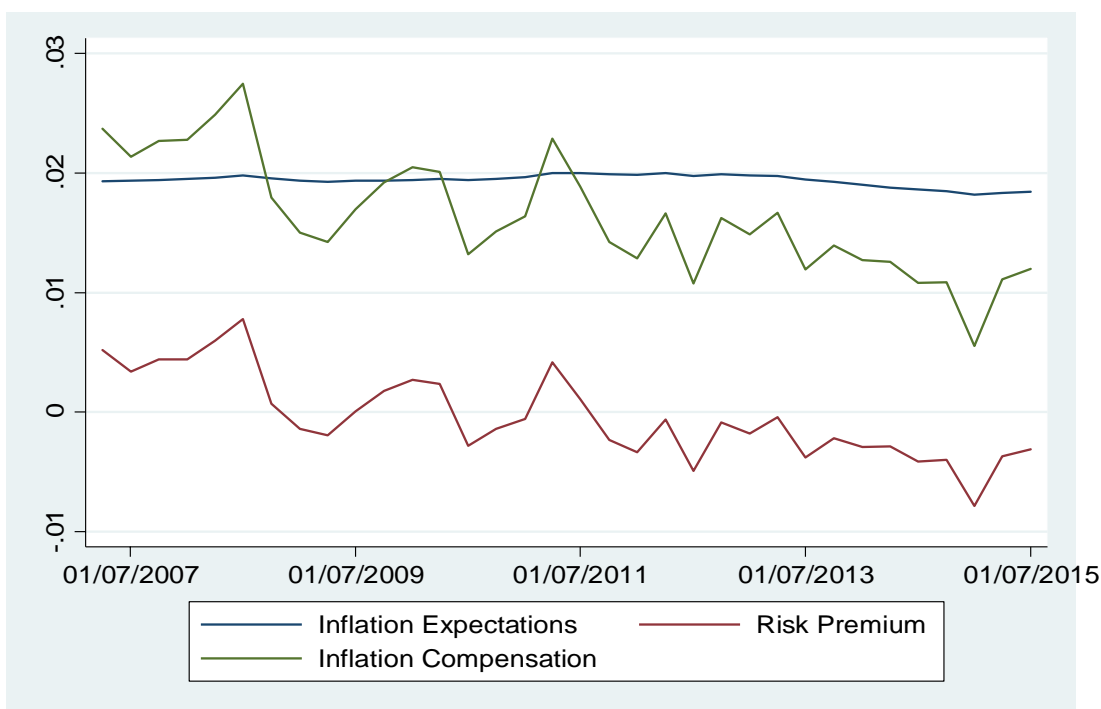


Figure 8. Decomposition of 5-year inflation compensation into inflation expectations and inflation risk premium.

5.3 Event study

We now turn our attention to the interaction of Euro area inflation expectations and the monetary policy conducted by the ECB. To begin with, we look at the effect of changes to the ECB's key rates on expectations by simply regressing the inflation expectations on the ECB's main refinancing operations rate, as described in equation 3.32. We consider also adding the first lag of the rate as a regressor, taking into account the fact that the rate change might take effect after trading has been completed and hence be reflected in the bond prices of the next day.

As we were able to derive the true expectations only for a quarterly basis, we cannot use them in the event study analysis. We hence turn back to the daily time series of the 5-year inflation compensation and run the regression for both the liquidity-adjusted compensation and non-adjusted compensation. The adjusted series runs only from August 2009 and contains more of the European sovereign debt crisis time period, whereas the non-adjusted series contains the rapid lowering of the ECB's rate from around 4 % to 1 % during the second half of 2008.

Table 4. Regressions of liquidity-adjusted and non-adjusted 5-year inflation compensation on the ECB's main refinancing operations rate.

ECB Main Refinancing Operations Rate	Non-adjusted Compensation	Adjusted Compensation
t	1,148** (0,31)	1,154** (0,36)
t-1	-0,886** (0,32)	-0,634* (0,36)
Intercept	0,012	0,012
R ²	0,60	0,56

Notes: Newey-West standard errors with maximum lag of 10 are in parentheses. The number of observations are 2 231 and 1 611 for the non-adjusted and adjusted series respectively. ** Significant at the 5 % level. * Significant at the 10 % level.

The results of the regressions are presented in Table 4. As can be seen, all the coefficients are significant and the explanatory power of the model is high. The fact that the coefficient of the first lag of the ECB's rate is negative would support the

conclusion that the effect of a rate change is observed by the markets on the next trading day, given that convention suggests that a loosening in monetary policy leads to higher expectations of inflation. The larger, more significant and positive effect of the same period ECB rate on inflation compensation might point to the fact that, if we assume the effect is felt on the next trading day, a rate change is most likely a response to either too low or too high inflation or inflation expectations. If this is the case we would expect to see past values of inflation expectations have a positive correlation with the ECB's rate and conversely (if monetary policy is effective) the future values of expectations having a negative correlation, which is consistent with the results⁷.

To look at a more dynamic relationship between expectations and the ECB rate, we run the VAR model described in equation (3.33). The model is run for both the liquidity-adjusted and non-adjusted series, with the results being very similar. Hence, for simplicity we present only the results of the non-adjusted compensation. The lag length is specified as $p = 2$, which is the optimum given by each of the final prediction error (FPE), Akaike's information criterion (AIC), Schwartz's Bayesian information criterion (SBIC) and the Hannan-Quinn information criterion (HQIC).

The ordering of the variables follows naturally from the setup of looking at the effects of changes in monetary policy on expectations, hence the assumed direction of causality runs from the ECB rate to expectations. We note that empirically there is no clear direction of causality given by either the Granger causality test or correlograms. Again this is most likely due to rate changes being a response to expectations and expectations in turn responding to rate changes.

The results of the VAR model are presented in the impulse response functions in Figure 9. The graphs of interest are the effect of a shock in the expectations on the ECB rate (top-right hand corner) and the effect of a shock in the ECB rate on expectations (bottom-left hand corner). The impulse responses seem to support the hypothesis that a negative shock in inflation expectations leads to a lowering in the ECB's main refinancing operations rate and conversely, a negative shock in the rate (i.e. monetary easing) causes initially an increase in expectations that levels

⁷ The fact that past values of expectations correlate positively with present values of the rate is further confirmed by regressing compensation on the first forward of the ECB rate. The coefficients are positive and significant at the 1 % level for both the adjusted and non-adjusted series.

out through time. It has to be noted however that the latter effect is very small, just about significant at the 5 % level and not persistent.

Moreover, the results should be tempered by the fact that, as with all macroeconomic variables, there are so many forces affecting each other at the same time that drawing any kind of causal link between phenomena is practically impossible. This is especially the case with regressions as simple as the ones presented here. The robustness of the models could be improved for example by adding surprise components of announcements of changes to other macroeconomic variables, but this is beyond the scope of this study.

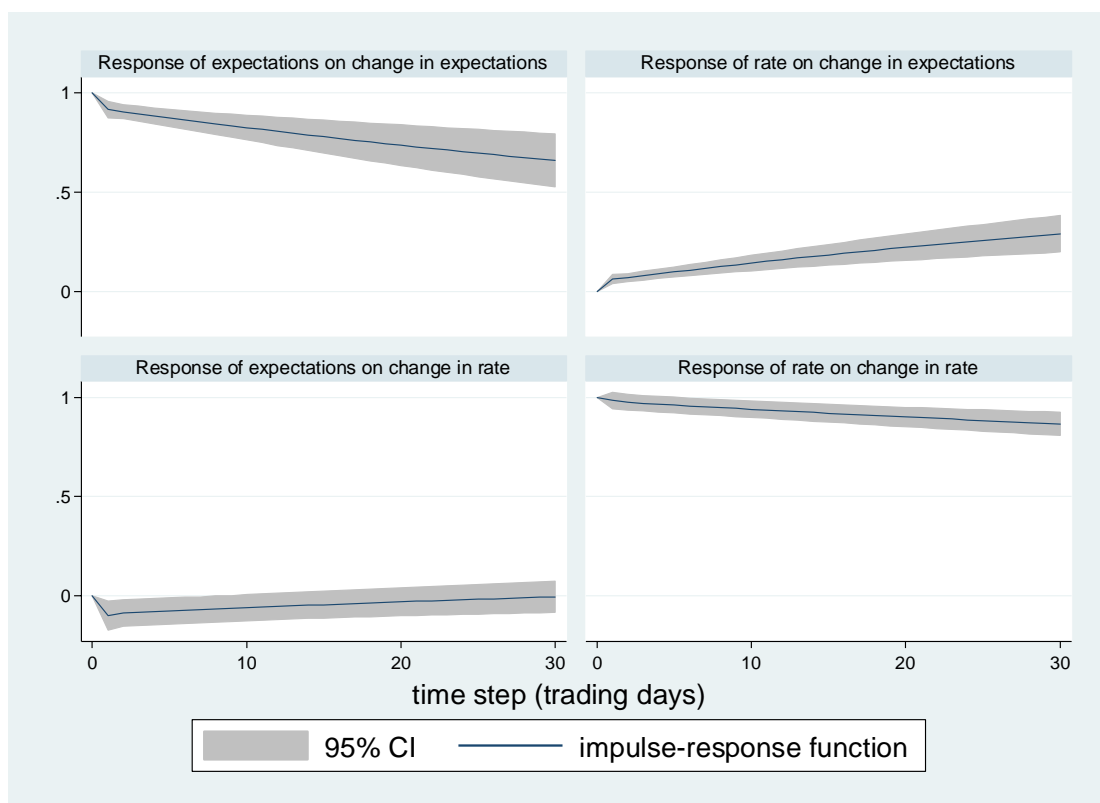


Figure 9. Impulse responses of shocks to inflation expectations and the ECB's key rate.

Another issue is the identification of the timing of the rate change, or more precisely, the timing of the release of information on the rate change. As mentioned in section 4.2, the monetary policy decisions of the ECB are made at the meetings of the Governing Council, and published during a press conference on the same day. The announcement of the rate change precedes the implementation by about a

week, for example the last change to the rate, a cut of 10 basis points, was announced on 4 September 2014, but not put into effect until 10 September 2014. According to conventional notions, efficient markets will absorb the information as soon as it is made public, hence the effect on expectations should occur straight after the announcement.

Furthermore, a decision to not change rates can be as impactful as a decision to change them depending on what the expectations of the ECB's actions are. The ECB can also inform of other monetary policy measures besides rate changes during these meetings, which can likewise have an impact on inflation expectations. Neither of these effects is captured in the regressions using the physical rate changes.

In order to take these effects into account, we regress the change of expectations against a dummy variable containing the dates of the monetary policy announcements as presented in equation (3.34). The dummy variable is 1 for the days of the announcements, again we will look at both the day of the announcement and the next trading day after it, and 0 for all other days. Since the governing council meets approximately once a month to decide on the key interest rate decisions, including every meeting date in the regression will likely not produce any meaningful results. Instead, we first include only the dates of the meetings when the rate was changed to complement the previous regression and VAR. The regression is run for both the liquidity-adjusted and non-adjusted series.

We then look at specific dates when the ECB announced significant non-standard monetary policy measures to counter the challenging economic environment of the sample period. These measures are detailed in Table 5. As before, we use a dummy for both the date of the announcement and the next trading day

Table 5. Dates of non-standard monetary policy measures conducted by the ECB 2009-2015.

Date	Measure announced
4.6.2009	Covered bond purchase programme (CBPP)
10.5.2010	Measures to address severe tensions in financial markets, including securities markets programme and long-term refinancing operations (LTRO)

6.10.2011	New covered bond purchase programme (CBPP2)
8.12.2011	Measures to support bank lending and money market activity
26.7.2012	Mario Draghi: "Within its mandate, the ECB will do whatever it takes to preserve the Euro"
6.9.2012	Outright monetary transactions (OMT)
4.7.2013	Forward guidance: ECB expects key rates to remain at current or lower levels for an extended period of time
5.6.2014	Measures to enhance the functioning of the monetary policy transmission mechanism, including targeted longer-term refinancing operations (TLTRO) and negative deposit rates
4.9.2014	Asset-backed purchase programme (ABSPP) and new covered bond purchase programme (CBPP3)
22.1.2015	Extended asset purchase programme (EAPP)

The results of the regressions are presented in Table 6. The regressions A) and B) are run separately. On the whole it seems that announcements of the changes to the ECB rate have no effect on expectations. Half of the rate changes occurred before the beginning of the adjusted series, hence the significant coefficient on the next trading day after the announcement for that series is most likely a coincidence due to a lack of observations.

The ineffectiveness of the announcements is most likely explained by the fact that the rate changes are well anticipated even before the announcement. This could be verified by regressing the compensation on a surprise component, in other words comparing the consensus ex ante view of the rate change to the ex post outcome. Winkelman et al. (2014) for example look at simultaneous movements in intraday interest rates to identify surprises in the ECB's monetary policy announcements. They find that on the whole ECB policy decisions are well anticipated with few significant surprises in the announcements, which would seem to be in line with the results of regression A).

The coefficients in regression B) on the other hand are more significant for the day of the announcement, so it looks like the non-standard monetary policy measures have more of a direct effect on expectations. It seems intuitive that non-standard measures would be harder to anticipate than rate changes given their unorthodox nature. The positive sign on the change to both the adjusted and non-

adjusted series is in line with the purpose of the non-standard measures to raise expectations when they are low. The fact that the day after the announcement has no effect on expectations indicates that the most of the immediate effect of the measures is instantly priced into the government bonds.

Table 6. Regressions of changes to inflation compensation on time dummies of ECB monetary policy announcements.

Regressor	Δ Non-adjusted Compensation	Δ Adjusted Compensation
A) Date of rate change		
t	-0,00001 (0,0002)	0,00004 (0,00008)
t+1	-0,00010 (0,0002)	-0,00314** (0,0015)
R ²	0,0002	0,0029
B) Date of non-standard measure		
t	0,00040** (0,0002)	0,00056*** (0,0002)
t+1	0,00010 (0,0002)	-0,0001 (0,0001)
R ²	0,0022	0,0095

Notes: Regressions of the non-adjusted 5-year inflation compensation on dummies for the dates of ECB key rate change announcements (A) and ECB non-standard monetary policy measure announcements (B). Heteroscedasticity robust standard errors are in parentheses. Number of observations is 2 233 and 1 610 for the non-adjusted and adjusted series respectively. *** Significant at the 1 % level. ** Significant at the 5 % level.

One possibility not accounted for in the various models run so far is that the monetary policy measures, both standard and non-standard, might not have an instantaneous effect, but might work their way into expectations more gradually. This could be the case if the French government bond markets are not fully efficient meaning that information takes time to be included in the prices. To answer this question we look at the effects of the announcements on the non-adjusted compensation graphically, as presented in Figure 10.

Again, it should be remembered that with macroeconomic variables looking for natural experiments and “treatment effects” is next to impossible given the multi-

tude of variables affecting each other simultaneously. The graphs should therefore not be taken at face value. Even so, there are no clear universal effects after either rate change, or policy announcements. There are a few instances where a trough in the compensation is coupled with a monetary policy measure, but others when the post-announcement trend is steeply downwards or flat. It should be noted that the last rate change of 2009 and the first change of 2011 are actually rate hikes, both of which are roughly followed by a drop in compensation, as would be expected.

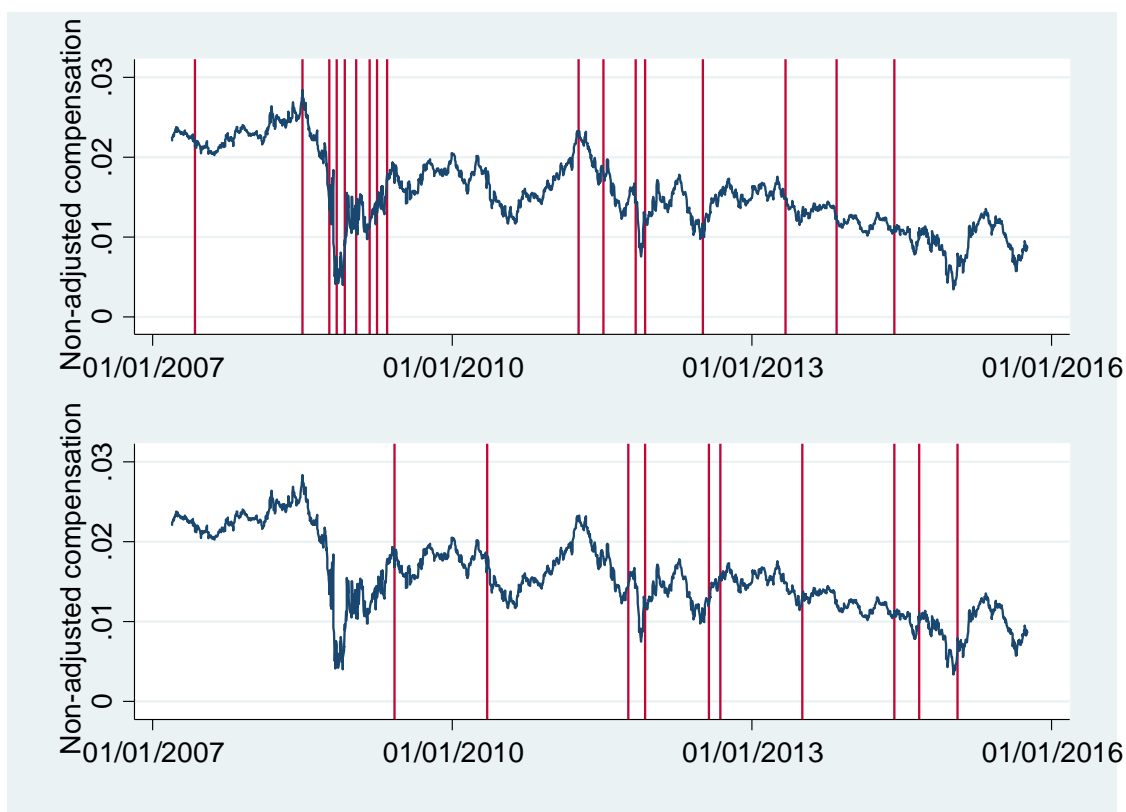


Figure 10. Dates of announcements of changes to the ECB rate (upper panel) and of ECB non-standard monetary policy measures (lower panel).

5.3.1 Expanded asset purchase programme (EAPP)

To conclude the section on empirical analysis we look at the most significant single monetary policy measure conducted by the ECB during the period of interest. On 22 January 2015, the ECB announced the initiation of an expanded asset purchase programme involving purchases of bonds issued by the Euro area central governments, agencies and European institutions in secondary markets. The programme is intended to run at least until September 2016 with monthly purchases

totalling 60 billion Euros. The programme has been likened to similar asset purchase programmes run by the Federal Reserve, Bank of Japan and Bank of England, generally dubbed quantitative easing (QE). (European Parliament 2015)

The main objective of the programme is to help the ECB fulfil its price stability mandate, more specifically bringing inflation back to the ECB's target of less than but close to 2 %. In the press conference of the announcement, the ECB's chairman, Mario Draghi, when discussing the motivation behind the programme, remarked that: "this assessment is underpinned by a further fall in market-based measures of inflation expectations over all horizons, and the fact that most indicators of actual or expected inflation stand at, or close to, their historical lows" (European Central Bank 2015a). Our market based measure of expectations backs up this claim as the lowest value of the non-adjusted compensation in our sample is on 6 January 2015. Hence, given that targeting inflation expectations is a specific goal of the EAPP, it will be interesting to look at its effect so far.

We do this simply by comparing the pre-announcement and post-announcement trends, both by regressing inflation expectations on time, and graphically. A more precise method to correctly identify the specific effect of the announcement would be to use high frequency intraday data for the day of the announcement, but again this data was not available to us.

Table 7. Regressions of the liquidity-adjusted 5-year inflation compensation on a 120-day pre-announcement and post announcement period.

Regressor	Liquidity-adjusted compensation
Trend in 7.8.2014 - 21.1.2015	-0,00003*** (8,70e-06)
Intercept	0,667
R ²	0,55
Trend in 22.1.2015 - 9.7.2015	0,00003*** (7,77e-06)
Intercept	-0,616
R ²	0,53

Newey-West standard errors with maximum lag of 10 are in parentheses. Number of observations is 121. *** Significant at the 1 % level.

The results of the regression are shown in Table 7. The pre- and post-announcement periods are set at 120 trading days running from 7 August 2014 to 21 January 2015 and 22 January 2015 to 9 July 2015 respectively. The liquidity-adjusted inflation compensation is chosen as the dependent variable as we are not interested in the period prior to August 2009. As can be seen, the pre-announcement trend is almost perfectly reversed after the announcement, with the coefficient of the same magnitude, but with the sign reversed to positive.

The same effect can be seen from the graphical representation in Figure 11. Here the post announcement period is extended to the end of the sample to look at the situation at the time of writing. The announcement of the EAPP coincides with a trend reversal in inflation expectations, though they had begun to rise already before the announcement, most likely in anticipation of a significant measure by the ECB. There is also a clear jump in expectations on the exact day of the announcement.

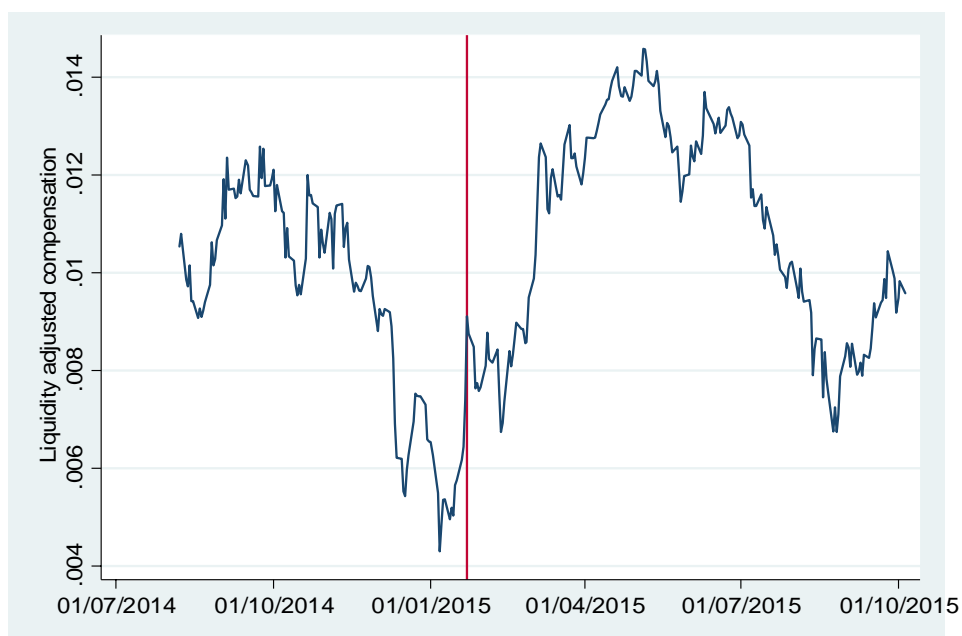


Figure 11. Trends of liquidity-adjusted 5-year inflation compensation before and after announcement of the EAPP on 22 January 2015.

However, as mentioned earlier, expectations fall back down close the pre-announcement level after a summer containing uncertainties about Greek debt repayment and the growth of the Chinese economy. Expectations begin to pick up

once more at the start of autumn and it will be interesting to see if they reach a permanently higher level in the near future.

6. Conclusions

This thesis has looked at the estimation of market-based measures of Euro area inflation expectations and the interaction of those expectations with the monetary policy conducted by the ECB for the period of 2007-2015. The method used in the analysis was the measure of inflation compensation, calculated as the difference between nominal and real yield curves extracted from French government bond prices. We find that there appears to be time-varying liquidity and inflation risk premia in the inflation compensation, meaning that it is not a pure measure of inflation expectations. We go about extracting the premia using proxies for liquidity and a state-space model, though due to the limits of available data, we are only able to produce the true inflation expectations at a quarterly frequency.

The evolution of inflation expectations as measured by inflation compensation is volatile during the sample period. There was a sharp fall in expectations during the height of the global financial crisis in 2008 followed by a rebound back to the ECB's target of close to 2 %. Expectations began to fall again during 2013 and reached the low point of the sample in January 2015. The quantitative easing programme initiated also in January 2015 raised expectations for a while, until a volatile summer brought them back down. Expectations have been rising again during the autumn of 2015.

Judging by the volatility of the long-run compensation series of 10- and 15-year ahead inflation, inflation expectations are weakly anchored. We note also that liquidity seems to account for only a small amount of the variation in the compensation of French government bonds, hence the liquidity-adjusted compensation follows closely the path of the non-adjusted series with only the level slightly higher.

We use the daily compensation series to analyse the effectiveness of the ECB's monetary policy with a series of regressions, which constitute an event study. We find a positive correlation between past expectations and the current ECB key rate and a negative correlation between the past rate and current expectations, con-

sistent with the notion that expansive monetary policy is conducted to raise expectations when they are low, and that expectations subsequently pick up as a result. We find no relationship between the announcements of rate changes and expectations, pointing to the conclusion that ECB rate changes are well anticipated. Announcements of non-standard monetary policy measures are less anticipated and they seem to succeed in raising expectations, at least initially. Finally the quantitative easing programme initiated by the ECB in January 2015 lifted expectations for a while, but its long-run effects are yet unknown.

These findings would seem quite encouraging from the perspective of a central bank's monetary policy. The fact that the rate change announcements are well anticipated indicates that the ECB's communication is consistent and predictable, meaning that it has credibility among financial markets participants. The success of the non-standard measures in affecting a positive change in expectations is a good complement to the predictability of rates, especially since the latter have ceased to be a policy measure now that we are at the zero lower bound. This ensures that the ECB's monetary policy can still be effective even in this kind of environment. Whether the measures can help affect a permanent rise in expectations remains to be seen. This is especially true of the effectiveness of the expanded asset purchase programme (EAPP).

The limitations of the thesis concern mainly the availability and quality of the data used in the analysis. Due to the relatively low number of government bonds linked to Euro area inflation, the beginning of the sample period is constricted to a point in time when there were enough outstanding bonds from which to construct the yield curves. This means that most of the sample contains times of severe financial markets turbulence, which makes the comparison to normal times impossible, and might cause fluctuations in the inflation compensation that are not related to changes in expectations.

Particularly problematic is the survey measure of inflation expectations, the Survey of Professional Forecasters, used to identify the true expectations, both in frequency and variation. It seems unlikely that true Euro area inflation expectations are as well anchored as the survey seems to suggest, given the communication and measures of the ECB. These problems could have possibly been overcome with the use of another survey series, which unfortunately was not available. The inclusion of high frequency intraday data on the movements of bond prices would help bet-

ter identify the impact of monetary policy announcements. Adding surprise components of macroeconomic news announcements would also make the results of the regressions more robust.

A longer sample of both bonds and surveys will produce more robust results, hence future studies using the same methodology can be more informative. In addition, the further development of financial assets and derivatives linked to inflation will give researchers a richer source of data to directly measure market-based expectations, including an accurate gauge of the uncertainty of expectations, as well as a point estimate. The effects of the EAPP can also be better appraised in a few years' time. Inflation expectations are a key consideration for a central bank, and so ways to more accurately measure these expectations should remain an important consideration for future research.

References

- Agence France Tresor (2015): *Characteristics of OAT€is*.
<http://www.aft.gouv.fr/rubriques/oateurois-and-btaneurois_178.html>, retrieved 14.1.2015.
- Ang, A, & Bekaert, G. & Wei, M. (2007): *Do macro variables, asset markets, or surveys forecast inflation better?* Journal of monetary Economics. Vol 54 No 4(2007) 1163-1212.
- Berk, J. M. (1999): *Measuring inflation expectations: a survey data approach*. Applied Economics. Vol 31, No 11(1999) 1467-1480.
- Bernanke, B. (2007): *Inflation expectations and inflation forecasting*. Speech at the Monetary Economics Workshop of the National Bureau of Economic Research Summer Institute, Cambridge, Massachusetts, July 10, 2007.
- Bishop, G. & Welch, G. (1997): *An Introduction to the Kalman Filter*. Department of Computer Science, University of North Carolina.
- Breedon, D. T. & Litzenberger, R. H. (1978): *Prices of state-contingent claims implicit in option prices*. Journal of Business. Vol 51, No 4 (1978), 621-651.
- Chen, J.H. (2008): *Variance ratio tests of random walk hypothesis of the euro exchange rate*. International Business & Economics Research Journal. Vol 7, No 12(2008) 97-106.
- Cox, J. C. & Ross, S. A. (1976): *The valuation of options for alternative stochastic processes*. Journal of Financial Economics. Vol 3, No 1 (1976), 145-166.
- Deeley, C. (2008): *Superseding Newton with a superior bond yield algorithm*. 21st Australasian Finance and Banking Conference 2008 Paper.
- Ejsing, J. & Garcia, J. A. & Werner, T. (2007): *The term structure of euro area break-even inflation rates: the impact of seasonality*. Working paper 830. European Central Bank.
- European Central Bank (2011): *The Monetary Policy of the ECB 2011*. <<https://www.ecb.europa.eu/pub/pdf/other/monetarypolicy2011ee.pdf>>, retrieved 31.1.2015.

- European Central Bank (2015a): *Introductory statement to the press conference (with Q&A): Mario Draghi, President of the ECB*. Frankfurt am Main, 22 January 2015
<<https://www.ecb.europa.eu/press/pressconf/2015/html/is150121.en.html>> , retrieved 3.11.2015.
- European Central Bank (2015b): *ECB survey of professional forecasters*.
<<http://www.ecb.europa.eu/stats/prices/indic/forecast/html/indin.en.html>> , retrieved 11.11.2015.
- European Parliament (2015): *The ECB's Expanded Asset Purchase Programme: Will quantitative easing revive the euro area economy?*. Briefing February 2015. <<http://www.europarl.europa.eu/EPRS/EPRS-Briefing-548976-The-ECBs-EAPP-FINAL.pdf>> , retrieved 3.11.2015
- Evans, M. D. (1998): *Real rates, expected inflation, and inflation risk premia*. Journal of Finance. Vol 53, No 1, 187-218.
- Fama, E. F. & Gibbons, M. R. (1984): *A comparison of inflation forecasts*. Journal of Monetary Economics. Vol 13, No 3, 327-348.
- Frontline Solvers (2015): *Standard Excel Solver – GRG nonlinear solver stopping conditions*. <<http://www.solver.com/standard-excel-solver-grg-nonlinear-solver-stopping-conditions>> , retrieved 2.12.2015.
- Garcia, J. A., & Van Rixtel, A. A. (2007): *Inflation-linked bonds from a central bank perspective*. ECB Occasional Paper, No 62.
- Gonzalez-Parado, J. M. (2007): *Expectations and credibility in central banking: a practitioner's view*. Speech at conference on inflation targeting, central bank independence and transparency, Cambridge, United Kingdom, 15 June 2007.
- Gürkaynak, R. S. & Sack, B. & Swanson, E. (2005): *The sensitivity of long-term interest rates to economic news: Evidence and implications for macroeconomic models*. American Economic Review. Vol 95, No 1, 425-436.
- Gürkaynak, R. S. & Sack, B. & Wright, J. H. (2006): *The U.S. treasury yield curve: 1961 to the present*. Finance and Economics Discussion Series. Staff Working Paper. Federal Reserve Board, Washington D.C.
- Gürkaynak, R. S. & Sack, B. & Wright, J. H. (2010): *The TIPS yield curve and inflation compensation*. American Economic Journal: Macroeconomics. Vol 2, No 1 (2010) 70-92.

- Harvey, Andrew C. & Shephard, Neil (1993): *Structural Time Series Models*. Handbook of Statistics. Vol 11(1993) 261-302.
- Haubrich, J. & Pennacchi, G. & Ritchken, P. (2012): *Inflation expectations, real rates, and risk premia: evidence from inflation swaps*. Review of Financial Studies. Vol 25, No 5, 1588-1629.
- Hurd, M & Relleen, J (2006): *New information from inflation swaps and index-linked bonds*. Bank of England Quarterly Bulletin. Vol 46, No 1 (2006), 24-34.
- Jalles, J. T. (2009): *Structural Time Series Models and the Kalman Filter: a concise review*. FEUNL Working Paper Series.
- Jochmann, M. & Koop, G. & Potter, S. M. (2008): *Modelling the dynamics of inflation compensation*. Working Paper 15-09. The Rimini Centre for Economic Analysis.
- Kitsul, Y. & Wright, J. H. (2012): *The economics of options-implied inflation probability density functions*. Working Paper 18195. National Bureau of Economic Research.
- Kuismanen, M. & Spolander, M. (1994): *Measuring inflation expectations in Finland: a survey data approach*. Bank of Finland discussion papers. 21/94.
- Lo, A. W. & Mackinlay, C. A. (1988): *Stock market prices do not follow random walks: evidence from a simple specification test*. Review of Financial Studies. Vol 1, No 1(1988) 41-66.
- Mankiw, N. G. & Reis, R. & Wolfers, J. (2004). Disagreement about inflation expectations. NBER Macroeconomics Annual 2003. Volume 18, 209-270. The MIT Press.
- Nelson, C. R. & Siegel, A. F. (1987): *Parsimonious modeling of yield curves*. The Journal of Business. Vol 60, No 4(1987) 473-489.
- Pflueger, C. E. & Viceira, L. M. (2011): *An empirical decomposition of risk and liquidity in nominal and inflation-indexed government bonds*. Working Paper 11-094. Harvard Business School.
- Sargent, T. J. (2008): *Rational Expectations*. The Concise Encyclopedia of Economics. Library of Economics and Liberty. <
<http://www.econlib.org/library/Enc/RationalExpectations.html>>, retrieved 5.2.2015

- Sarr, A. & Lybek, T. (2002): *Measuring liquidity in financial markets*. Working Paper 02/232. International Monetary Fund.
- Smith, T (2012): *Option-implied probability distributions for future inflation*. Bank of England Quarterly Bulletin. Vol 52, No 3 (2012), 224-233.
- Stock, J. H. & Watson, M. W. (2007): *Why has U.S. inflation become harder to forecast?* Journal of Money, Credit & Banking. Vol 39, No 1 (2007) 3-33.
- Svensson, L. E. O. (1994): *Estimating and interpreting forward interest rates: Sweden 1992-1994*. Seminar Paper No 579. Institute for International Economic Studies, University of Stockholm.
- Thomas, L. B. (1999): *Survey measures of expected US inflation*. The Journal of Economic Perspectives. Vol 13, No 4, 125-144.
- Tuckman, B. & Serrat, A. (2011): *Fixed income securities: tools for today's markets*. 3rd edition. John Wiley & Sons. Hoboken New Jersey.
- Walsh, C. E. (2010): *Monetary theory and policy*. 3rd edition. The MIT Press Cambridge Massachusetts.
- Winkelmann, L. & Bibinger, M. & Linzert, T. (2014): *ECB monetary policy surprises: identification through cojumps in interest rates*. European Central Bank. Working paper no 1674. May 2014.

Appendix

A. Deriving the Nelson-Siegel function (equation 3.6)

Equation 3.5 is integrated from 0 to m and divided by m :

$$\begin{aligned} & \frac{1}{m} \int_0^m \beta_0 + \beta_1 \cdot e^{\left(-\frac{m}{\tau}\right)} + \beta_2 \left(\frac{m}{\tau} \cdot e^{\left(-\frac{m}{\tau}\right)}\right) dm \\ &= \frac{1}{m} \left[\beta_0 \cdot m + \beta_1 \left(\tau - \tau e^{\left(-\frac{m}{\tau}\right)} \right) + \beta_2 \left(\tau - e^{\left(-\frac{m}{\tau}\right)} (m + \tau) \right) \right] \\ &= \frac{1}{m} \left[\beta_0 \cdot m + \beta_1 \left(\tau \left(1 - e^{\left(-\frac{m}{\tau}\right)} \right) \right) + \beta_2 \left(\tau \left(1 - e^{\left(-\frac{m}{\tau}\right)} \right) - m \cdot e^{\left(-\frac{m}{\tau}\right)} \right) \right] \\ &= \beta_0 + \beta_1 \left(\frac{1 - e^{\left(-\frac{m}{\tau}\right)}}{\frac{m}{\tau}} \right) + \beta_2 \left(\frac{1 - e^{\left(-\frac{m}{\tau}\right)}}{\frac{m}{\tau}} - e^{\left(-\frac{m}{\tau}\right)} \right) \\ &= \beta_0 + (\beta_1 + \beta_2) \cdot \frac{1 - e^{\left(-\frac{m}{\tau}\right)}}{\frac{m}{\tau}} - \beta_2 \cdot e^{\left(-\frac{m}{\tau}\right)} \blacksquare \end{aligned}$$

The derivation of the Svensson extension (equation 3.8) follows essentially the same procedure.