EUROPEAN CENTRAL BANK

EUROSYSTEM

WORKING PAPER SERIES

NO 1391 / OCTOBER 2011

INTEREST RATE EXPECTATIONS AND UNCERTAINTY DURING ECB GOVERNING COUNCIL DAYS

EVIDENCE FROM INTRADAY IMPLIED DENSITIES OF 3-MONTH EURIBOR

by Olivier Vergote and
Josep Maria Puigvert Gutiérrez

EUROPEAN CENTRAL BANK

EUROSYSTEM

WORKING PAPER SERIES

NO 1391 / OCTOBER 2011

INTEREST RATE EXPECTATIONS AND UNCERTAINTY DURING ECB GOVERNING COUNCIL DAYS

EVIDENCE FROM INTRADAY IMPLIED DENSITIES OF 3-MONTH EURIBOR[1]

by Olivier Vergote[2]

and Josep Maria Puigvert Gutiérrez[2]

In 2011 all ECB publications feature a motif taken from the €100 banknote.

1 The authors would like to thank Thomson Reuters for providing the tick data. An earlier version of the paper benefited from comments by B. Fischer, M. Ehrmann and the ECB Graduate Programme Committee.

2 European Central Bank, Kaiserstrasse 29, D-6031, Frankfurt am Main, Germany; e-mails: olivier.vergote@ecb.europa.eu and josep_maria.puigvert@ecb.europa.eu.

Address
Kaiserstrasse 29
60311 Frankfurt am Main, Germany

Postal address
Postfach 16 03 19
60066 Frankfurt am Main, Germany

Telephone
+49 69 1344 0

Internet
http://www.ecb.europa.eu

Fax
+49 69 1344 6000

Information on all of the papers published in the ECB Working Paper Series can be found on the ECB's website, http://www.ecb.europa.eu/pub/scientific/wps/date/html/index.en.html

ISSN 1725-2806 (online)

CONTENTS

Abstract

This paper analyses changes in short-term interest rate expectations and uncertainty during ECB Governing Council days. For this purpose, it first extends the estimation of risk-neutral probability density functions up to tick frequency. In particular, the non-parametric estimator of these densities, which is based on fitting implied volatility curves, is applied to estimate intraday expectations of three-month EURIBOR three months ahead. The estimator proves to be robust to market microstructure noise and able to capture meaningful changes in expectations. Estimates of the noise impact on the statistical moments of the densities further enhance the interpretation. In addition, the paper assesses the impact of the ECB communication during Governing Council days. The results show that the whole density may react to the communication and that such repositioning of market participants' expectations will contain information beyond that of changes in the consensus view already observed in forward rates. The results also point out the relevance of the press conference in providing extra information and triggering an adjustment process for interest rate expectations.

JEL: C14, E43, E52, E58, E61
Keywords: risk-neutral probability density functions, option-implied densities, interest rate expectations, central bank communication, intraday analysis, announcement effects, tick data.

Non-technical summary

Central bank communication receives widespread attention by financial market participants. Reactions to central bank messages can take several forms: surprises, changes in uncertainty or the absence thereof when announcements were already anticipated. The extent of such market reactions and its drivers are of interest to both market participants and policy makers. A first challenge in analysing the reactions is to define appropriate measures and determine the relevant indicators to look at. A second challenge is to determine the factors that drive the market expectations and uncertainty. This paper contributes to both areas.

In order to study expectations held by market participants, the paper relies on the estimation of risk-neutral probability density functions. Such densities summarise the total set of likely outcomes and probabilities attached by the market. The densities can be extracted from option prices and are therefore also referred to as option-implied densities. The densities capture not only the consensus expectation as present in forward rates, but also the uncertainty around it, thereby allowing a much broader view based on analysis of for example the statistical moments of the densities. So far, such densities were estimated at daily frequency, while most announcement effect analysis now takes place at intraday frequencies. Therefore the paper also brings the estimation of implied densities to intraday frequency.

In particular, this paper studies changes in the expectations and uncertainty up to tick frequency and aims to identify drivers of the market reactions during ECB Governing Council days. First, the paper tackles a number of practical and statistical considerations that appear when bringing implied density extraction to high frequency. Second, based on case studies and analysis of intraday patterns the paper also measures the information content of the obtained densities and uncertainty measures. In addition, it carries out a regression analysis to identify drivers of the observed market reactions as expressed in the density changes.

The benefit of the approach clearly is that - as is demonstrated with a few case studies - one can zoom in on certain events and judge the immediate market reactions, thereby minimising the bias by any other information hitting the market. Furthermore, the intraday densities reflect the dispersion and symmetry of the expectations, thereby giving the policy maker an idea about the relative expectations and uncertainty in the market, and market participants about the risks in the market.

The analysis is based on expectations three months ahead about a money market interest rate, namely the 3-month EURIBOR. The densities are computed based on a non-parametric estimator applied to tick data on three-month EURIBOR futures and options. The paper introduces an efficient method to pre-filter data to impose no-arbitrage conditions as required by option pricing theory. The results show that the estimator is robust to market microstructure noise by producing stable risk-neutral densities. At the same time, when information hits the market the densities adapt quickly and meaningfully, indicating that the estimator is flexible enough to capture changes in expectations. Estimates of the noise impact point to a relatively small influence and allow it to be taken into account when interpreting developments.

An economic assessment of the announcement effects of ECB communication on short-term interest rate expectations is carried out based on a sample of 32 days on which the ECB Governing Council took a policy rate decision. The intraday patterns of the statistical moments of the implied densities show a significant shock in activity following the press release and significantly increased activity during the press conference showing the relevance of both their content. All considered moments (mean, median, standard deviation, skew and kurtosis) show such patterns. Furthermore, apart from reaching very distinct levels between days, it is shown that the moments can also strongly move within a Governing Council day, in particular during the financial crisis period.

A regression analysis identifies a number of drivers of the expectation changes following the press release and during the press conference. A surprise in the policy rate decision, as perceived by the

market, was found to significantly affect the entire density, hence not only the consensus view but also the relative positioning of expectations. Uncertainty surrounding the decision and the EURIBOR itself was also found to be relevant, but here evidence was less strong. A code word, as perceived by the market in predicting rate hikes, was found to have guided expectations. This confirms the value attached by markets to perceived patterns in the wording by the central bank and rate decisions. In addition, indications were found that the overall content of the introductory statement and Q&A session was relevant in driving expectations.

Overall, the relevance of the press release and conference as communication tool is confirmed. This holds for both the introductory statement and the question and answer session of the press conference, which given the (continued) high activity during these sessions appears to provide additional information to markets. The information is not simply adding noise that could offer an alternative explanation for the increased activity. Instead expectations are guided in specific directions. This provides support to the use of a press conference following policy rate announcements as is practised by the ECB. In principle, the tools and analysis could be extended to the quarterly press briefings recently introduced by the Federal Reserve.

1 Introduction

Policy rate announcements by central banks are renowned for their widespread financial market and media attention given the relevance of rate setting for asset prices and economic developments. The expectations and uncertainty that prevail among market participants about these announcements, and the extent to which surprises occur, are informative to both market participants and the policy maker. Central banks typically hold the responsibility of contributing to the efficient allocation of funds in the economy and hence have an incentive to avoid market surprises. At the same time, monetary authorities typically will not want to pre-commit to following through on any policy signal they may have given, and thus surprises remain possible to some extent.

Unsurprisingly, given the usually swift reactions observed in asset prices, there has been a move to ever higher frequencies in analysing market reactions to news. This allows the asset price reaction to be observed more directly and 'contamination' of the signal by reactions to other news arriving around the same time can be kept to a minimum. An extensive literature has established the significance of various macroeconomic announcements and assessed it for a number of financial markets. In the context of this paper, market reactions to ECB policy announcements at intraday frequency have been studied by e.g. Andersson (2007), who analyses the impact on asset price volatility, while Brand et al. (2006) study the reaction of the money market yield curve, and Ehrmann and Fratzscher (2009) the reaction in EURIBOR futures prices.

However, intraday research has focused merely on changes in the consensus expectation, as expressed in forward rates, while changes in the uncertainty surrounding this average expectation have been broadly ignored. Still, uncertainty measures such as implied volatility have been analysed intensively at daily frequency. Likewise, the literature on implied densities, which looks at the entire density of expectations and the developments of the statistical moments of such densities, has provided useful insights. More specifically, these densities capture the likelihood attached that market participants attach to specific outcomes; see e.g. Bahra (1997) for an overview.

This paper studies changes in the expectations and uncertainty up to tick frequency and aims to identify drivers of the market reactions during ECB Governing Council days. The paper thus contributes in two distinct ways to the literature. First, the estimation of implied densities is brought to the intraday frequency. There are a number of practical and statistical considerations that need to be tackled for this purpose. In particular, market microstructure effects, which are known to challenge high-frequency inference, need to be taken into account. Second, the paper assesses the information content of the obtained densities and uncertainty measures based on case studies and analysis of intraday patterns. In addition, it carries out a regression analysis to identify drivers of the observed market reactions as expressed in the density changes. While the sample size is limited, the regression results do allow an assessment of the impact of ECB communication, without claiming to be exhaustive. Importantly, this final part of the paper aims to promote further research on this topic and the collection of the necessary detailed data. The tools presented here could also easily be extended to other financial instruments and used to evaluate the communication of other central banks, e.g. the quarterly press briefings which the Federal Reserve recently introduced.

The paper is structured as follows. Section 2 presents the estimator used to extract densities from option prices. Section 3 discusses the way in which the estimation is brought to the intraday setting. Section 4 discusses the statistical moments of the densities and the impact of market microstructure noise. In Section 5, the added value of these densities as a monitoring tool is demonstrated with a few case studies. In Section 6, intraday patterns of the density moments are analysed to gauge the impact of the press release and conference. Next, Section 7 carries out a regression analysis to identify a number of drivers of the changes in the density moments. Section 8 concludes.

2 Implied density estimation

The estimation of option-implied densities - capturing market participants' expectations - is based on futures and options prices of a specific underlying instrument, e.g. EURIBOR. Since the payoff of these securities depends on the future outcome of the underlying instrument, the current price of these securities contains information about market participants' expectations about that future outcome. These expectations can be seen as a set of likely outcomes with different probabilities attached to them, hence defining a probability density function. Consequently, the whole idea behind the estimation is to extract this density from the observed prices.

The estimation method applied in this paper belongs to the non-parametric class of estimators. The literature has step-by-step suggested further improvements to the non-parametric estimation of implied densities. The implementation by de Vincent-Humphreys and Puigvert (2010) builds on recent suggestions in the literature and importantly on the estimator presented by Bliss and Panigirtzoglou (2002). Their estimator is also the one applied in this paper, apart from being brought to the intraday frequency as discussed in the next section. This section briefly presents the non-parametric estimator, while the above two articles contain further details on the implementation.

In short, the estimation of the implied probability density function is based on the Breeden and Litzenberger (1978) result, which states that the implicit interest rate probabilities can be inferred from the second partial derivative of the call price function with respect to the strike price. However, instead of directly interpolating the call price function and calculating the second partial derivative, smoother results can be obtained if the data on option premia and strike prices are transformed into implied volatility and delta values prior to interpolation.

Following this approach, the extraction of a density from option prices can be seen to consist of four steps. Figure 1 presents an example of these estimation steps. It starts with the selection of option price observations. Here only the out-of-the-money (and at-the-money) options are selected. The reason is that the market for these options is more liquid than for in-the-money options, which may not be traded and hence lacking a price or less actively traded and therefore more prone to measurement error. This way a single option price is selected per strike price (i.e. interest rate) taken from either call or put options as presented in panel (a) of Figure 1. The second step consists of estimating the implied volatility curve. Abstracting from this for a moment, it is natural that the estimation of a continuous density function requires the interpolation between discrete observations at some stage of the estimation. In short, this is done here. However, instead of fitting a price function for the option price observations in panel (a), the literature has shown that more stable results are obtained by fitting instead the implied volatility curve in 'delta-sigma' space as presented in panel (b), where delta is the derivate of the Black-Scholes (1973) price with respect to the price of the underlying asset. This approach is motivated by the work of Shimko (1993) and Malz (1998), and since the reliance on the Black-Scholes pricing formula is only used as a tool it does not make the density estimation parametric. The option strike prices are transformed into deltas and the option prices into implied volatilies. The implied volatilities are calculated by numerically solving the Black's (1976) version of the options pricing model for the value of σ

$$C(F_t, K, t) = e^{-r\tau}\left[F_t\Phi\left(\frac{\ln\left(\frac{F_t}{K}\right)+\frac{\sigma^2\tau}{2}}{\sigma\sqrt{\tau}}\right) - K\Phi\left(\frac{\ln\left(\frac{F_t}{K}\right)-\frac{\sigma^2\tau}{2}}{\sigma\sqrt{\tau}}\right)\right], \tag{1}$$

where C is the call price, K is the strike price, r is the risk-free rate, F_t is the value of the underlying future at time t, $\tau = T - t$ is the time to maturity T, and Φ is the standard normal distribution function. And similarly for put options. To transform strike prices into deltas, the implied volatilities are used to calculate the delta values

$$\delta = \frac{\partial C}{\partial F} = e^{-r\tau}\Phi\left(\frac{\ln\left(\frac{F_t}{K}\right)+\frac{\sigma^2\tau}{2}}{\sigma\sqrt{\tau}}\right). \tag{2}$$

Following Campa et al. (1997), a cubic smoothing spline is fitted in delta-sigma space resulting in a volatility curve, also referred to as the 'volatility smile'. The cubic smoothing spline minimises the objective function

$$\min \lambda \sum_{i=1}^{n} \omega_i(\sigma_i - \hat{\sigma}_i(\Theta))^2 + (1-\lambda)\int_0^1 g''(\delta,\Theta)^2 d\delta, \tag{3}$$

where σ_i, $\hat{\sigma}_i$, and ω_i are respectively the observed sigma, fitted sigma and weight of observation i ($i = 1, ..., n$, $\sum_{i=1}^{n}\omega_i = 1$), δ represents the observed deltas, Θ is the matrix of polynomial parameters belonging to the spline, $g(\cdot)$ is the cubic spline function, and λ is the smoothing parameter fixed at 0.99. The weighting is based on Black-Scholes vega (ν), $\omega_i = \nu_i^2 / \sum_{i=1}^{n}\nu_i^2$, $i = 1, ..., n$. The value of vega approaches zero for deep out-of-the-money options and reaches a maximum for at-the-money options. More specifically, the weight attached to the observations in this estimation decreases towards the end points of the curve. This way, the impact of measurement error that the underlying price observations typically contain is minimised. This explains why the fitted implied volatility curve may (intentionally) deviate somewhat from the observations as in panel (b). The third step consists of moving the fitted curve back to 'interest rate - option price' space as shown in panel (c). This is done by evaluating the interpolated volatility smile at a large number (1,000) of delta values, transforming the delta values back into strike prices using the inverse of equation (2):

$$K = F_t \exp\left(\frac{\sigma^2\tau}{2} - \sigma\sqrt{\tau}\Phi^{-1}(e^{r\tau}\delta)\right), \tag{4}$$

where Φ^{-1} is the inverse of the standard Normal distribution function, and computing call option prices at those strike prices using (1). A put option premium function is obtained similarly. In a fourth step, the second derivative of the premium function of panel (c) is taken, which provides the implied density as shown in panel (d). This last step relies on the analytical result of Breeden and Litzenberger (1978) which motivated the first steps in the estimation.

The implied density estimates are in fact estimates of the so-called risk-neutral probability density function as for example appears in Cox and Ross (1976) option valuation formula. Therefore, it needs to be kept in mind that the expectations as presented by the risk-neutral densities differ to some extent from the density of 'real' expectations. The reason is that the density extraction relies on a simplifying assumption, i.e. all assets have the same expected return, namely the risk-free rate. The literature usually considers the risk-neutral densities to be close enough proxies to make inference as if it concerned real expectations. Therefore, the paper also abstracts from this difference and for simplicity refers to implied densities in the remainder of the paper.

Another fact that is important for the interpretation of the implied densities is that they are computed for a constant maturity; in this case expectations three months ahead. Since the underlying futures contracts have fixed expiry dates, the computation of constant maturity densities involves interpolation between the fixed expiry dates. The method does this interpolation between implied volatility curves, instead of implied densities. The advantage is mainly practical as the delta range [0,1] is the same for different contracts trading at different strike price ranges. See Vincent-Humphreys and Puigvert (2010) for further details and ECB (2011) for a less technical discussion about daily density estimates.

Compared to other implied density estimators presented in the literature, the non-parametric method appears to have a few advantages. Cooper (1999) and Bliss and Panigirtzoglou (2002) found that the estimator based on fitting the volatility smile is more stable and robust to pricing errors than the parametric

approach based on a mixture of lognormals, which has also received broad attention in the literature. As microstructure noise is expected to introduce more measurement error, the stability of the method is important, and motivates the choice here for the non-parametric estimator based on the volatility smile. Furthermore, since the aim is not to estimate specific parameters of the underlying asset price process, the estimation benefits from the flexibility provided under the non-parametric setting.

Figure 1: Estimation steps

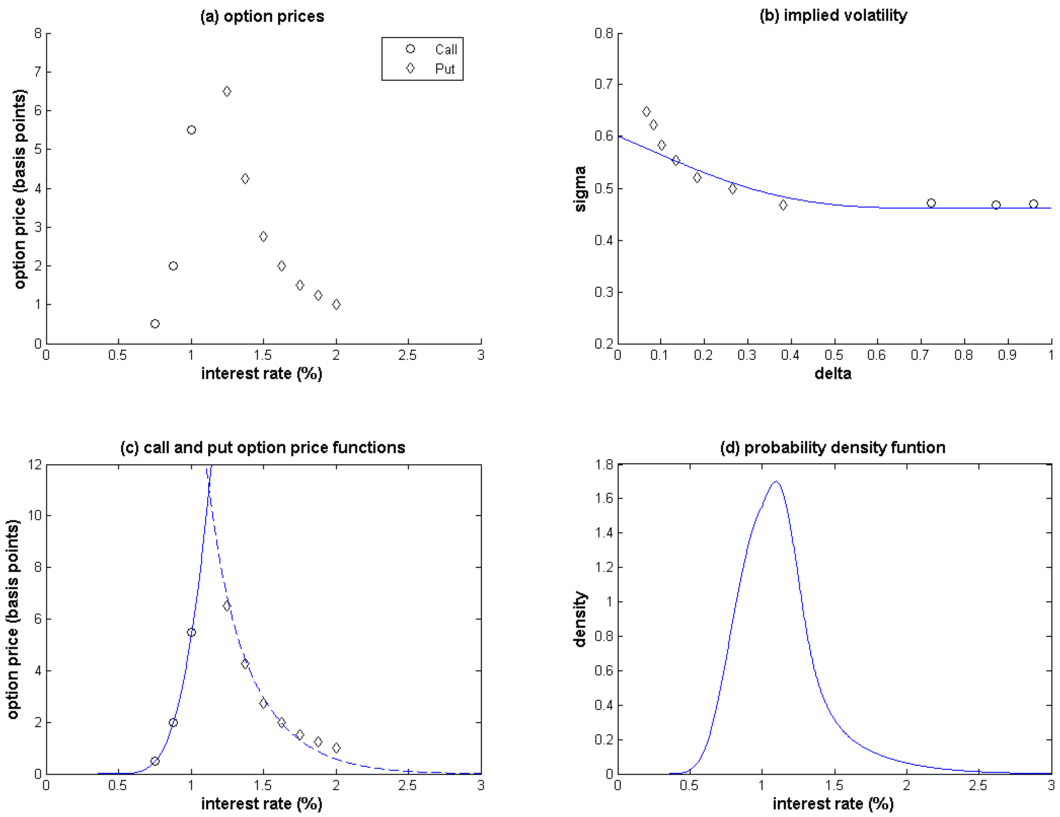

Compared to other implied density estimators presented in the literature, the non-parametric method appears to have a few advantages. Cooper (1999) and Bliss and Panigirtzoglou (2002) found that the estimator based on fitting the volatility smile is more stable and robust to pricing errors than the parametric approach based on a mixture of lognormals, which has also received broad attention in the literature. As microstructure noise is expected to introduce more measurement error, the stability of the method is important, and motivates the choice here for the non-parametric estimator based on the volatility smile. Furthermore, since the aim is not to estimate specific parameters of the underlying asset price process, the estimation benefits from the flexibility provided under the non-parametric setting.

3 Intraday implied densities

This section explains the application of the non-parametric density estimator to tick data on futures of three-month EURIBOR and options on these futures. These instruments are traded at LIFFE.

The estimation of implied densities is subject to a number of assumptions that may not entirely hold in practice. First, the underlying analytical results assume perfectly competitive markets. For example, the Cox-Ross option valuation assumes short-selling is allowed and there are no transaction costs or taxes. In reality, certain rigidities are in place. Still, the EURIBOR derivatives market studied here is very liquid.

Even during the recent financial crisis when money markets were under pressure, liquidity in EURIBOR derivatives remained vivid. Second, the analytical results of Breeden and Litzenberger (1978) underlying the estimation method were derived based on no-arbitrage conditions. From empirical studies, it is clear that these do not always hold. In particular, the observed call and put option premium functions are not always monotonic and convex as would be required under no-arbitrage conditions. Therefore, previous studies have often pre-filtered the data before estimating (daily) implied densities.

It is natural to expect that moving to intraday frequencies brings new challenges. First, the price of each asset needs to be determined for any given moment in time during the day, which differs substantially from the daily setting where one can conveniently rely on the daily settlement prices provided by the exchange. Second, as is known from financial research at high-frequencies, market microstructure effects are likely to add noise to the estimates. In particular, the rules of the exchange determine how trades and quotes can take place and affect among others the observed price process. One could expect to find more violations of monotonicity and convexity conditions when studying tick data. Sub-section 3.1 explains the way prices are measured for the empirical study in Sections 5 to 7. Sub-section 3.2 presents an efficient filter to impose no-arbitrage conditions on the data.

3.1 Prevailing prices

Prices can be derived from both transactions and quotes. EURIBOR futures trade very frequently within the day. Therefore, the transaction price is used as their price indicator. In contrast, the corresponding options do not trade so actively. However, they are actively quoted within the day. Therefore, the transaction prices, but also the mid-quotes are used in the case of options. Since quotes are binding and spreads are tight, this average of the best bid and ask price offers a good price indicator. The use of quotes is actually common in case of applications to exchange rates, see e.g. Castrén (2004). Moreover, the LIFFE rules state that the settlement prices for EURIBOR futures and options, which are commonly used for daily inference, can also be based on quotes in the absence of trades (NYSE LIFFE, 2009).

The next sections present implied densities estimated up to tick frequency between 8:30 and 18:30 C.E.T. A new implied density is computed each time the price of the future or a related option changes. For this purpose, the price of all the securities needs to be known at each tick time. This is done by computing the price that prevails at each tick for each security. The prevailing price is determined by looking back in time for the last price update found in the tick data for that security. The fact that many options with different strike prices are considered in the estimation of an implied density, and that the time of the last price update can differ substantially between each of these options raises the issue of non-synchronous trading/quoting. If certain instruments had recent price updates while other instruments had not, then this would bias the estimation if the latter quotes could be considered outdated. Fortunately, the LIFFE tick data allow one to control for this to an important extent, because they also contain indications when quotes seize to exist for a security (i.e. its order book is empty because of order withdrawels or executions in a trade). In such case, the security can be taken out of the estimation, and one can be confident that the remaining best quotes are still active, even if they were entered some while ago. The fact that quotes are binding also contributes to them being representative.

In this context, it needs to be remarked that the intraday data do not conceptually differ from the daily data that typically rely on settlement prices. The computation of settlement prices also comes down to determining the prevailing quote, but towards the end of the trading day. As different methods are considered to compute settlement prices for EURIBOR futures and the computation is at the exchange's discretion, the settlement price may be even seen as more opaque. On the other hand, the settlement prices for options undergo some pre-filtering since a consistency check is carried out on the implied volatilities of adjacent contracts (NYSE LIFFE, 2009). Daily settlement price data may therefore better satisfy no-arbitrage conditions.

3.2 No-arbitrage conditions

The non-parametric implied density estimator is based on the Cox-Ross option pricing equation, which relies strongly on no-arbitrage conditions. The presence of any arbitrage opportunities as reflected in the option premia, which can occur through pricing errors or genuine market conditions, distorts the implied volatilities, volatility smile and implied density. In particular, the no-arbitrage conditions require the option premia as a function of the strike price to be monotonous and convex. A common example is the violation due to the price grid, which turns the premium function into a step function in the area where the premium function becomes relatively flat, and thus not convex. In addition, the so-called 'bid-ask bounce' of prices combined with asynchronous trading and quoting could lead to violations. In contrast, a favourable feature is that the LIFFE trading system implements price limits when entering orders thereby already avoiding pure price errors to a certain extent.

The raw data are usually filtered before an implied density estimator is applied. First, as pointed out by de Vincent-Humphreys and Puigvert (2010), the estimation is best applied to out-of-the-money and at-the-money options since these are the most liquid options. Second, deep out-of-the-money options with the smallest possible premium, i.e. the tick size, for more than one consecutive strike price are deleted. The case for deletion is that discreteness strongly blurs their price signal and that in delta-sigma space these options have about the same delta but different values for sigma, which is inconsistent with the rest of the observations. Third, monotonicity and convexity are then tested for the call and put premium functions separately.

The solution adopted in the literature is to exclude observations from the observed option premia that - as a function of the strike price - do not satisfy the monotonicity and convexity conditions. However, the method to select the observations to be excluded is usually not presented and is likely to have been fairly arbitrary. This observation is strengthened by the fact that any attempt to run a filter through the premium function that tests these conditions sequentially on observations will not be able to guarantee that the optimal selection has been made and not even that all violations have been cleared. Iterating such a filter may achieve the latter, but will not be able to guarantee the optimal selection in terms of a minimum amount of observations deleted.

This paper suggests the following optimal method. Instead of sequential operations on adjacent observations, it is feasible to consider all the observations at once. Let n be the total number of observations and consider first all possible combinations of $n-1$ observations out of n. For each such combination a test of monotonicity and convexity can be applied. If one or more combinations pass the tests, then the implied density estimator can be applied to one of these combinations. If all combinations fail the tests, then all possible combinations of $n-2$ observations out of n are considered. This sequence of selecting and checking continues by reducing the number of observations considered until a combination is found that passes the tests. This way, the identified combination of observations is also known to minimise the number of observations excluded from the total set of n. Furthermore, the monotonicity and convexity tests can easily be set up by checking their mathematical definition sequentially on sets of (two and three) adjacent observations of the combination to be checked. Although the set of all possible combinations grows fast in n, in practice the number of observations is normally not that big to cause numerical problems.

3.3 A chain of densities

The density estimator allows estimates for the wide majority of ticks considered leading to a relatively stable chain of densities. The estimator is the one presented in Section 1 adapted to the intraday setting as discussed in the previous sub-sections. As an example, Figure 2 presents the chain of densities estimated at tick frequency throughout 5 March 2009. Densities could be computed throughout the day, which proved to be a relatively calm day given the modest changes in the densities. If the estimator happened to break down when applied to other days, it usually meant that market activity was so low that too few option price observations were available to allow estimation; minimum three observations are needed to allow an implied volatility curve. The market was also found to halt occasionally during some Governing Council days, e.g. within the minute before the press release or start of the press conference, making estimation infeasible for an instant. In cases with too few active options, it was found useful to attach a price to far out-of-the-money options such as to guide the estimation. In particular, if there is no quote on the bid side while there is a quote on the ask side at a low price (e.g. up to twice the tick size), it could be assumed that the bid is zero such that a midquote exists. This data filter helped to obtain more density estimates within days of low market activity.

Figure 2: Implied densities at tick frequency during 5 March 2009

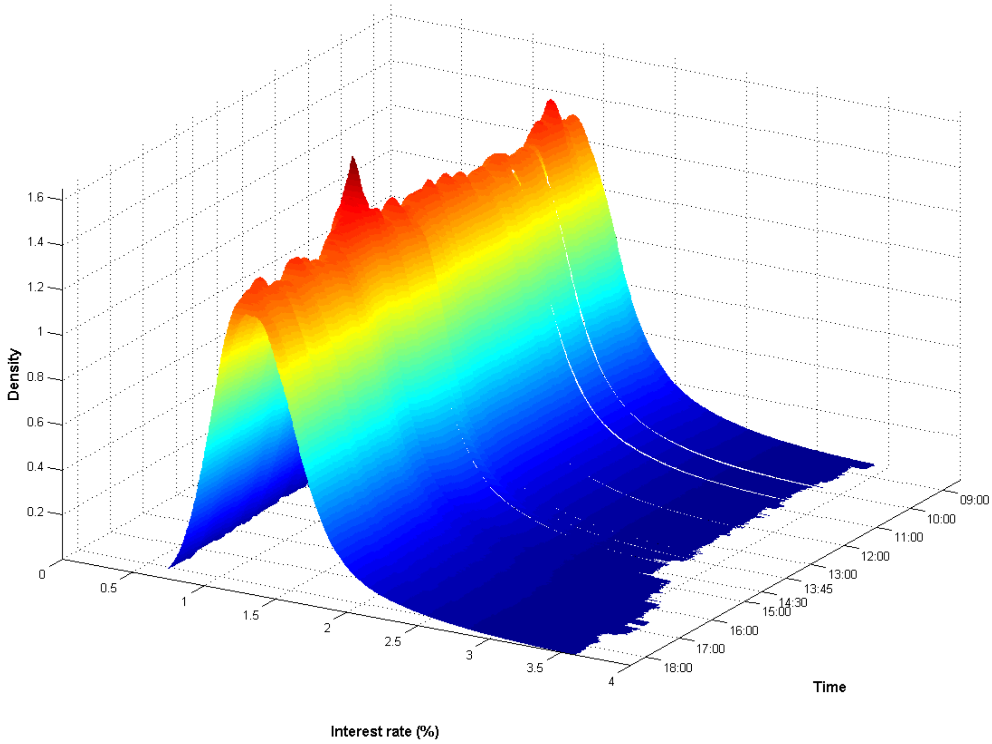

The stability of the estimated densities suggests that the estimator is robust to market microstructure noise to a large extent, although robustness checks would be needed to quantify and confirm this formally. Given the lack of an agreed upon benchmark model for implied densities and noise definitions in the literature, however, such a robustness study would easily become extensive. Therefore, this is considered to go beyond the scope of this paper, while the literature would benefit from such studies. Furthermore, any changes are hard to judge based on 3-D graphs of the densities and monitoring the density moments usually makes it easier to interpret developments.

4 Density moments and dealing with noise

The density moments quantify different properties of an implied density and make it easier and more intuitive to interpret changes over time than having to judge series of density shapes visually. The width of a density reflects uncertainty and this dispersion of expectations can be measured by the standard deviation. Furthermore, the (a)symmetry of a density reflects the probability attached to outcomes above versus that below the average expectation. For market participants this represents relative risks and asymmetry informs them about the 'balance of risks' (Lynch et al., 2004). Different symmetry measures are in use, but here the statistical skew (i.e. the normalised third central moment) is used. Next, the probability present in the density tails reflects the likelihood attached to extreme moves and provides another uncertainty measure. Kurtosis (i.e. the normalised fourth central moment) is used to capture this.

Figure 3 presents several central moments at tick frequency. The impact of the noise is clear here with erratic behaviour of the observed moments. For interpretation purposes, it remains difficult to judge what can be discarded as noise and what represents the signal.

The signal can be distinguished from the noise based on estimates of the size of this noise. Let m_i be a density moment observed at tick i, $i = 1, ..., n$, with n the number of observations in a fixed time interval. The changes in the moment (r) can be seen as composed of signal (u) and noise (e) components, $r_i = m_i - m_{i-1} = u_i + e_i$, $i = 2, ..., n$. According to asset pricing theory, signal changes will be very small at high frequency. The observed changes in the moments are substantial, however, implying that observed high-frequency changes are dominated by the noise component. Under general noise distribution, we know from Zhang et al. (2005) that a consistent estimator of the variance of the noise, ω^2, is given by

$$\widehat{\omega}_n^2 = \frac{1}{2n} \sum_{i=1}^{n} r_i^2 \xrightarrow{p} \omega^2,$$

as $n \to \infty$. Computing an average over all days in the sample (see Section 6) and taking the square root, the noise impact on the mean, median and standard deviation are estimated to be respectively 0.31, 0.69 and 0.33 basis points. This implies that the impact of the noise is small in absolute value. Furthermore, the median is more affected than the mean. The noise impact on the skewness and kurtosis, estimated to be 0.0685 and 0.1865, are also small.

As an example of using the noise estimate in practice, the horizontal shaded band presented for each moment in Figure 3 presents twice the noise size centred around the level the moment reached at the start of the press conference (14:30 C.E.T.). As long as the moment stays within the band, its changes can be considered as noise, but when it leaves the band is very likely to be signal related. In this example, the policy rate was increased by 25 basis points on that day and the events can most likely be interpreted as follows. As the mean did not react at 13:45 C.E.T., the decision can be seen as broadly anticipated. However, the skewness dropped at that point in time while the median increased, which suggests that the bulk of the probability mass of the density moved towards higher EURIBOR levels. Thus, while the policy rate decision was anticipated, its actual occurrence and confirmation did impact the expectations among market participants about future interest rates as now more participants saw higher EURIBOR outcomes as more likely. During the press conference that followed, all moments moved significantly. In particular, the mean jumped up at the start of the conference, which may be a reaction to statements such as "monetary policy continues to be on the accommodative side" and "the Governing Council will monitor very closely all developments" that could have been interpreted by market participants as further rate hikes are to come. At the same time, uncertainty about future EURIBOR as captured by the standard deviation decreased, while uncertainty increased in terms of kurtosis, i.e. more probability attached to bigger changes. The skew and median continued their adjustment initiated earlier, which suggests that the expectations driving those changes were strengthened during the press conference. The noise band suggests that the changes in all moments that took place after 14:30 C.E.T. can be attributed to a signal, but also that the skew and standard deviation closed the day at levels that are not distinguishable from their 14:30 C.E.T. level when noise is taken into account.

Naturally, the simple noise assumption made above does not need to hold and one may want to expand on this if deemed necessary, but may come at the expense of simplicity. In addition, instead of comparing developments to one particular moment in time, one could attempt to place a band around the entire moment series as a confidence interval. However, this would require the choice of a smoothing parameter which remains arbitrary.

Figure 3: Selected density moments at tick frequency during 8 March 2007 and noise band estimate centred around the moment level reached at 14:30 C.E.T.

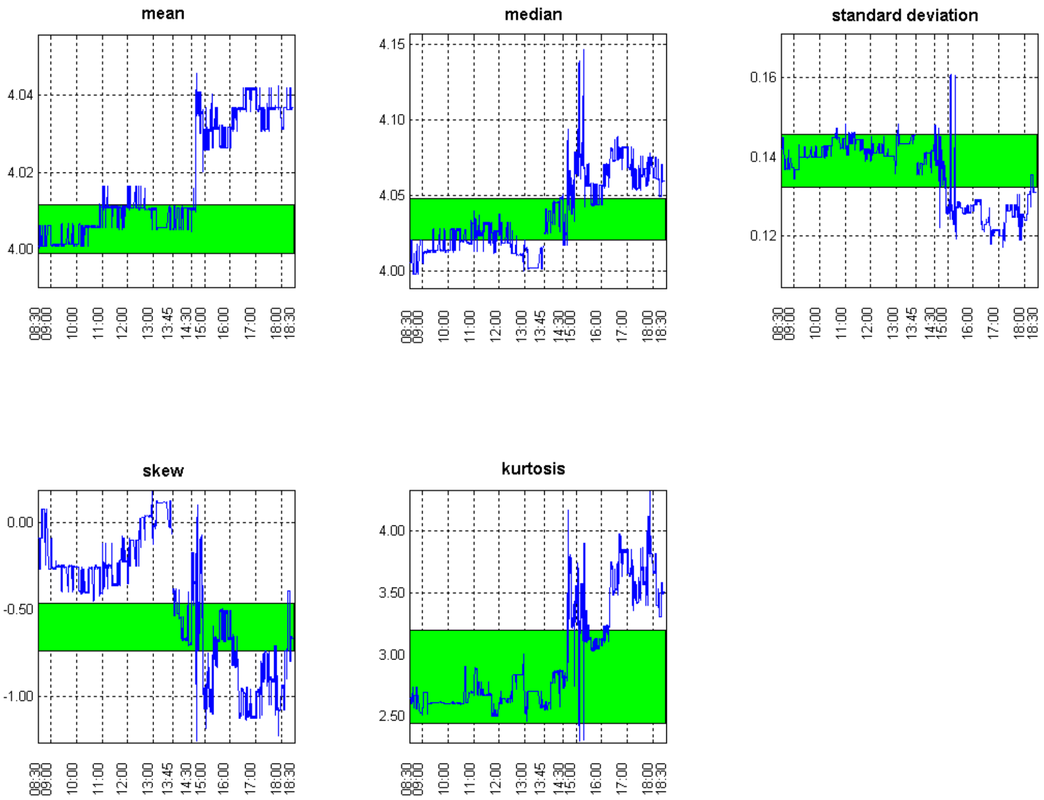

The noise size estimates support the view that the estimator is robust in the sense that moments can be seen to follow an underlying path over time around which noise causes small and very-short-lived changes. However, a robustness study as in Bliss and Panigirtzoglou (2002) at an intraday level would be beneficial. Cases where the estimator produced outliers proved to be rare and rather due to exceptional market situations.

Finally, it needs to be remarked that the noise is relevant at all frequencies. Although the noise is best estimated at the highest frequency available, its impact remains relevant for densities computed at lower frequencies. Consider sampling the densities or moments at lower densities from the tick frequency series. These series will look less eratic, but the selected observations are still the same and hence the noise impact remains equally present. Especially for stationary series such as interest rates this is an important issue since when the interest rate reverts towards a level previously reached, the noise may make it impossible to consider the levels as being different. For stock prices this would be less of an issue since they usually follow a drift making the signal the dominant part at low frequencies. Overall, this implies that even when comparing implied densities separated in time (e.g. daily) one would need to take the noise into account when interpreting relatively small changes as representing a signal versus being noise induced.

5 Implied densities as a monitoring tool

This section presents a few case studies that show how the implied densities can be used to monitor expectation and uncertainty developments over time and assess the developments around specific events. Andersen and Wagener (2002) pointed this out already by analysing the change in expectations about the next policy rate decision around the 9/11 attacks based on implied densities. Here, this monitoring is extended to the intraday frequency, and the ease in interpretating the results also provides a view on the quality of the estimates and shows that the estimator has the necessary flexibility to capture meaningful developments. The first case discusses the situation where the market attached value to a perceived code word concerning future policy rate decisions. The second case discusses the occurrence of a strong change in expectations when monetary policy information was released just before the ECB press release; in particular the effect of a strong rate cut by the Bank of England on 6 November 2008. The third case looks at more subtle changes in expectations.

Before interpreting the market reactions, it is important to realise that the size of the reactions of the implied densities may not necessarily look large in absolute value owing to the maturity of the expectations they capture. What is captured here are expectations about the three-month EURIBOR rate three months ahead. This implies that what counts are interest rate expectations between now and six months. Furthermore, since what is estimated are constant-maturity implied densities derived by interpolation of the implied densities around the first two futures expiry contracts, and the second futures contract may settle up to six months from now for a 3-month contract, the estimation may also pick up interest rate expectations between six and nine months from now. Consequently, what is observed in the implied densities are not only expectations about the next rate decision, but those for several consecutive months. Even if a policy rate decision was not fully priced in before the release, the interest rate and uncertainty reaction may not be strong in absolute terms because - following the expectations hypothesis - it is partly averaged out with the rate expected for the coming months. At the same time, this helps to explain the high activity typically observed during the press conference as it can be expected to contain information about the path of monetary policy in the short to medium term.

In addition, the financial crisis created special money market conditions. In particular, the interest rate on (unsecured) EURIBOR loans contained an elevated spread (above secured EONIA swap contracts) driven by perceived credit and liquidity risk. Therefore, expectations concerning this spread also played a role apart from policy rate expectations. At the same time, the spread and the implied density moments became also indicators of money market tensions.

5.1 Code word surprise

Market participants attached special value to the mentioning of the expression 'vigilance' during the introductory statement of the press conference, which was perceived as a code word for a rate hike at the next meeting during the rate hike cycle of 2005-2007. Other expressions in the intodsuctory statement were seen as predicting the mentioning of vigilance at the next meeting or hence a hike in two months time. The case presented here captures the events during 6 April 2006, a day when the perceived code word did not occur as was expected and thus a rate hike at the next meeting became less likely than previously thought. Figure 4 presents the implied densities at 14:30 and 15:00 C.E.T., i.e. just when the introductory statement is about to start and after half an hour of press conference.

Figure 4: Two intraday implied densities on 6 April 2006

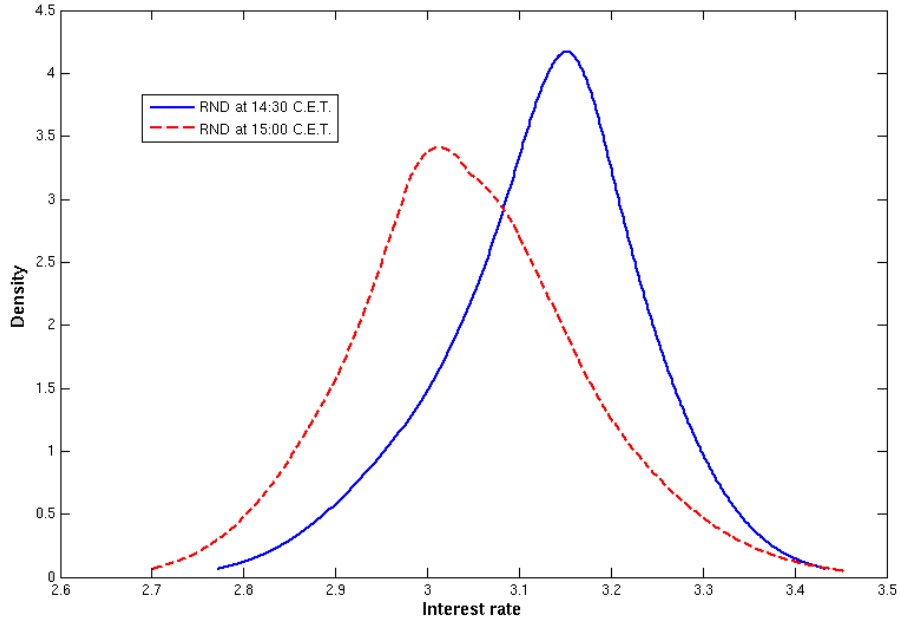

The results show a move of the implied density to the left as probability mass moved towards smaller interest rates. In economic terms, the change in the implied density is strong, especially given the small time interval. It is visually clear that the mean of this implied density - capturing the consensus expectation - decreased. However, the implied density also clearly contains more information. The change was not a mere shift of the entire density. Instead the support widened to the left indicating an increase in uncertainty and the skewness increased implying that the bulk of the expectations moved to lower rates, but leaving a longer tail behind at higher rates. Overall, a case study cannot control for other factors that may have played a role in the reaction, but judging from the comments during the question and answer (Q&A) session, the perceived code word surprise was surely an important element.

5.2 Expectation formation before the press release

Figure 5 presents the developments on 6 November 2008 when the ECB cut rates for the second time by 50 basis points. The first implied density shows the expectations at 12:55 C.E.T. whose mean clearly represented lower expected interest rates for three-month EURIBOR three months ahead, given that the policy rate was still at 3.75% at that point in time. At 13:00 C.E.T., the Bank of England announced a big rate cut of 150 basis points. Five minutes later, the implied density had moved tremendously to the left indicating that the Bank of England decision surprise made part of the market participants believe that the ECB would also come with a rate cut bigger than previously expected. At 13:40 C.E.T., the implied density still represented those expectations. At 13:45 C.E.T., the ECB announced (only) a 50 basis points rate cut. Five minutes later, the implied density looked completely different and rate expectations had moved up again. Clearly, the rate cut was smaller than what many had expected since the past 45 minutes and strongly increased rate expectations of three-month EURIBOR three months ahead. In fact, the implied density was bi-modal with the bulk of expectations around 3.25% and a smaller part around 2.75%, thereby still expressing uncertainty about the coming rate decisions. However, uncertainty in the form of a long right tail had disappeared.

Figure 5: Implied density mid-day developments on 6 November 2008

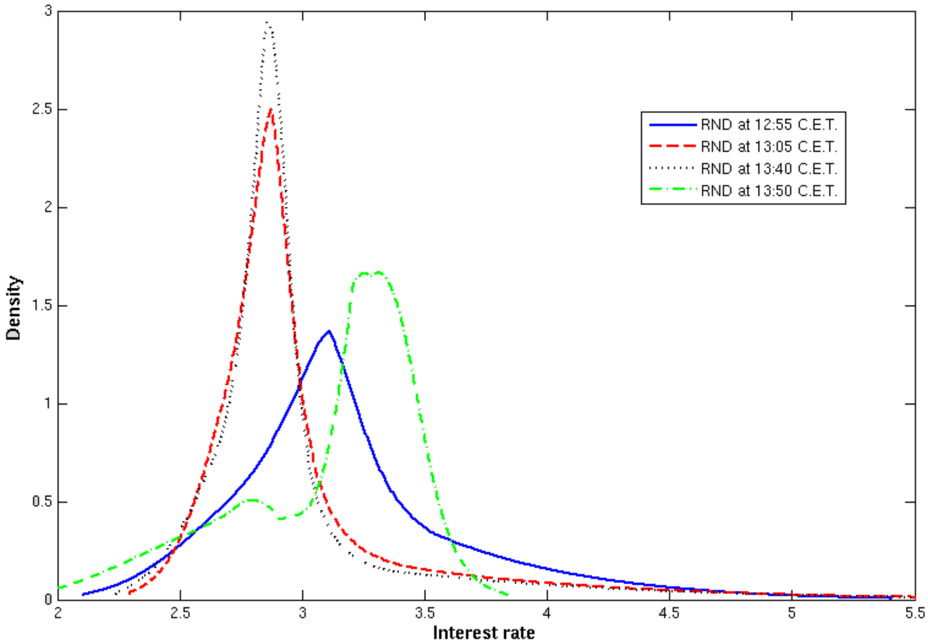

5.3 Changes in asymmetry and uncertainty

The cases discussed above represent exceptional expectation movements within a day where shifts of the density and its mean play an important role. The expectation and uncertainty changes are usually more subtle, however, and also captured by other density moments.

The developments during the press conferences of 8 March 2007 and 6 June 2007 serve as good examples of how the whole density can add value to the interpretation of expectation developments. For those two days, Figure 6 presents the implied density at 14:00 C.E.T. when the press conference is about to start, at 14:10 C.E.T. when the introductory statement is about to end, and at 15:00 C.E.T. close to the end of the question and answer session. On 8 March 2007, a policy rate increase to 3.75% was announced at 13:45 C.E.T. We notice that the density had two modes at 14:00 C.E.T. representing different views on future policy rate decisions. Interesting here is that during the press conference more probability mass of the subsequent implied densities moved towards the mode situated at the higher rate. A policy maker could check this type of developments in the modes and asymmetry of the density against its own believes and the intentions of its own communication. On 6 June 2007, a policy rate increase to 4.00% was announced and changes in interest rate uncertainty could be observed during the press conference that followed. After the introductory statement, uncertainty had increased somewhat as can be observed from the density width, but by the end of the Q&A session had decreased materially compared to the start of the press conference. For a policy maker it may be interesting to take note of such changes in uncertainty and may allow it to be checked against the content and intentions of its communication. As discussed in Section 4, however, it would be prudent to keep track of the impact of noise when judging relatively small changes in the densities. Overall, these examples show how the analysis of the whole density can be informative about changes in market expectations.

Figure 6: Implied densities during 8 March 2007 (left) and 6 June 2007 (right)

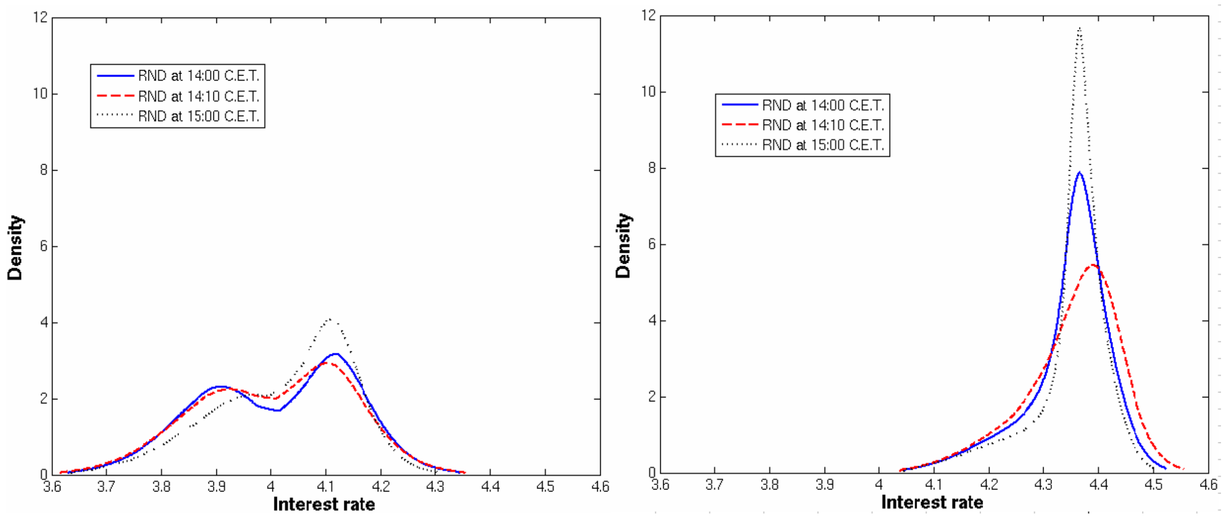

6 Relevance of press release and conference

In order to assess the impact of the press release and conference, this section extends the analysis to a set of Governing Council days. A unique tick-dataset on three-month EURIBOR futures and options was obtained from Thomson Reuters. The sample consists of 32 days on which ECB policy rate decisions were made and covers two sub-periods. The sub-periods are October 2005 - June 2007 and September 2008 - June 2009, which cover the two latest rate cycles of a gradually increasing policy rate and strong policy rate cuts, respectively. Appendix 1 presents the sample as part of a chart of the policy rate decisions during the period 2005-2009 where those occurring within the sample period are shaded. All except one decision occurred on regular Governing Council meetings, i.e. on 8 October 2008 a policy rate decision was taken in between meetings. The sample is limited owing to the restriction on the amount of days that could be obtained for research purposes. Therefore, the sample focuses on the last two policy rate cycles where most action in terms of market reaction and expectation adjustments can be expected.

Sub-section 6.1 analyses the intraday pattern of the density moments during policy rate decision days. Sub-section 6.2 assesses the range of the intraday moment changes.

6.1 Patterns during Governing Council days

The intensity of moment changes throughout a Governing Council day depends on the time of the day. As measure of density activity, the average absolute 1-minute change in each of the moments was computed per minute and is presented in Appendix 4. For this purpose, densities were computed at 1-minute equally-spaced time intervals. The patterns of intraday moment activity show a sharp spike at 13:45 C.E.T., when the policy rate decision is released, for all moments. The reaction is short-lived, which indicates that on average price adjustments occur swiftly when news hits the market. At 14:30 C.E.T., at the start of the press conference, activity jumps up again in all moments, but to a smaller extent on average while it stays elevated for a considerable period. The latter suggests that news may arrive at different points during the press conference. Activity only moderates to the morning level after half to one and a half hour. Furthermore, the intraday moment patterns are closely connected to the tick arrival patterns presented in Appendix 3 with a shock at 13:45 C.E.T., and elevated levels during the press conference in all moments. It.needs to be remarked that the smaller spike at 13:00 C.E.T. in these charts relates to certain Bank of England decisions that coincided with an ECB Governing Council day, and that U.S. jobless claim announcements coincide with the ECB press release and thus may contribute to the spike at 13:45 C.E.T.

The results show that not only the mean changes during a Governing Council meeting day, which is the part studied in the literature so far based on futures prices, but that all parts of the density and the expectations they represent change on average during Governing Council meeting days. Therefore, tracking the entire density is likely to improve our understanding of the market and expectations developments around announcements.

However, on average, Governing Council meetings are not found to significantly reduce the level of uncertainty. Appendix 4 also presents the average level of the standard deviation and kurtosis within the day. The increase in the average kurtosis between 13:00 and 13:45 C.E.T. is caused by the speculation on 6 November 2008, discussed as a case study above, and whose outlier status affects the average. More important, however, is that the average standard deviation and kurtosis hardly change within the day. The average standard deviation stayed close to 22 basis points, while the average kurtosis was close to 4 and hence somewhat more leptokurtic than the Normal density.

Table 1: T-test results for difference in averages

	Average value during:				
	Morning	Press release[3]	Introductory statement	Q&A	Press conference
Abs. change[1]:					
mean	0.0015	0.0042*** ’ **	0.0040***	0.0032***	0.0034***
median	0.0031	0.0089*** ’ ***	0.0071***	0.0061***	0.0063***
standard dev.	0.0007	0.0027*** ’ ***	0.0023***	0.0017***	0.0018***
skewness	0.0365	0.0785*** ’	0.1004***	0.0774***	0.0826***
kurtosis	0.0536	0.1578*** ’	0.1725***	0.1353***	0.1434***
Level[2]:					
standard dev.	0.228	0.224	0.227	0.224	0.225
kurtosis	4.153	4.041 ’ **	4.167	4.207	4.199

Notes: the averages are computed for a morning hour (11:00-12:00 C.E.T.), ten minutes following the press release and the actual duration of the introductory statement and Q&A session.
1. Test if mean absolute change is higher than in the morning,
2. Test if mean level is smaller than in the morning,
3. The second test for the press release tests the difference in mean with the press conference,
*, **, and *** denote significance at the 90%, 95% and 99% level respectively.

The patterns derived by judging the figures in Appendix 4 are also confirmed by statistical tests. Table 1 presents averages for several indicators measured over a number of intraday time intervals and t-test results for differences in those averages. First, the average of the absolute change in a moment during the press release, introductory statement, Q&A and total press conference are tested against the average absolute change during the morning. All test results show that the activity levels reached during the press release and (the parts of the) press conference were significantly higher than during the morning for all density moments. An extra test comparing the press release reaction to that of the total press conference shows that the mean, median and standard deviation are significantly higher due to the press release, but that this is not the case for the skewness and kurtosis. This result is consistent with the hypothesis that decision surprises will particularly involve changes in the mean and median as the consensus view adapts, while the skewness and kurtosis are relatively more affected by shifts in expectations related to the outlook discussed during the press conference. Furthermore, Table 1 presents the average standard deviation and kurtosis levels in the different time intervals. Tests show that their level was not significantly lower than during the morning, and thus that uncertainty did not decrease on average. Only the kurtosis following the press release is found to be significantly lower than the kurtosis over the total press conference, but in economic terms the difference in average kurtosis is small.

The intraday patterns are specific to the Governing Council days and not common to other days. The sample consisting of only Governing Council does not allow this to be tested directly. However, the difference in tick activity levels presented in Chart 1 of ECB (2007) point strongly in this direction. Moreover, tests based on the daily densities estimated by de Vincent-Humphreys and Puigvert (2010) provide strong support by showing that the Governing Council days imply significantly bigger reaction than other days. For this purpose, Appendix 2 presents a table of t-test results comparing reactions on Governing Council days to those on other Thursdays. In fact, these results are stronger than the results obtained by Mandler (2002) when studying the impact in a similar way for the period 1999-2000.

6.2 Range of intraday changes

Although the previous results indicate that the standard deviation and kurtosis - on average - move little throughout a Governing Council day, this hides that those moments do move substantially within individual days – and at times dramatically. This sub-section briefly discusses the main developments apparent from these statistics. Figure 7 presents candle plots for the changes of the standard deviation, skewness and kurtosis within each individual day of the sample. The two sub-periods are presented in separate charts. Each chart shows the range between the maximum and minimum reached each day (as a line) and the difference between the opening and closing observation (as a box, which is filled in case of a daily decrease). Again 1-minute spaced densities were used here, which reduces the impact of outliers among the tick densities.

Starting with the first period, the standard deviation was remarkably small and stable with an average level of 10 to 20 basis points. Even within days, this measure of dispersion moved little. This finding corresponds to the relatively narrow corridor of daily densities presented in Chart 11 of de Vincent-Humphreys and Puigvert (2010) for the same period. The kurtosis was also stable across the period, but here substantial movements within the day can be observed. The intraday changes in skewness are also more pronounced and significant deviations from zero, i.e. symmetry, are reached. The first and last day in this period happen to correspond to the days with the biggest and smallest skewness. On 6 October 2005, the positive skew after a long period of constant policy rate suggests that part of the market participants anticipated a rate hike in the near future, which did not (yet) correspond to the consensus view. The following month, this positive skew was no longer there suggesting these expectations had become a more central view. On 6 June 2007, after a period of gradual rate hikes, the pronounced negative skew suggests that part of the market participants anticipated at least a halt of the rate hikes, which was not represented by the consensus view yet. At the same time, the increased kurtosis for both of these days confirms the built-up of diverging views.

Turning to the second period, much higher levels of dispersion and kurtosis (note the difference in scale) and dramatic movements of all higher moments within certain days are observed. Uncertainty in terms of expectation dispersion was highest on 8 October 2008 the day the ECB announced its first rate cut. This picture is consistent with reports on the general financial market uncertainty at that time. With kurtosis still close to its average level and skewness moving around zero, the market was clearly divided. Uncertainty about the heightened risk premium contained by EURIBOR (above EONIA) around that period had very likely contributed to the interest rate uncertainty.

As time continued, dispersion decreased, but skewness and kurtosis reached high levels towards the end of the second period. As the policy rate approaches its trough the density has a tendency to become positively skewed. The strong positive skew captured the presence of a long tail on the right-hand side of the density, representing the adverse expectations for the money market kept by some market participants. Apart from this, the prominence of this tail implied that skewness and kurtosis captured market tensions more than the standard deviation during that period. Remaining uncertainty about the risk premium and about the bottom level for the policy rate were the likely drivers of uncertainty as also expressed during the Q&A session. Apparent are also the strong intraday movements on 6 November 2008, as discussed in Sub-section 5.2.

Overall, the intraday volatility of these moments is evidence of the ongoing price discovery in the market and the observed announcement effects may provide valuable information to the central bank. In particular, the changes in the skewness at high-frequency appear relevant since they capture the direction expectations are taking. Its interpretation is subtle though since it is important to make the difference between, on the one hand, a long tail building up on one end of the density as a result of certain market participants developing discordant expectations and, on the other hand, the bulk of the density shifting in the opposite direction and possibly "leaving a tail behind". It is clear that this will be easier to interpret if the skewness developments can be followed at high-frequency around a specific announcement. Similarly, high-frequency monitoring of the standard deviation and kurtosis would help to interpret increases in uncertainty and diverging expectations.

Figure 7: Candle plots of moment developments within a day: sub-period 2005-2007 (lhs) and sub-period 2008-2009 (rhs)

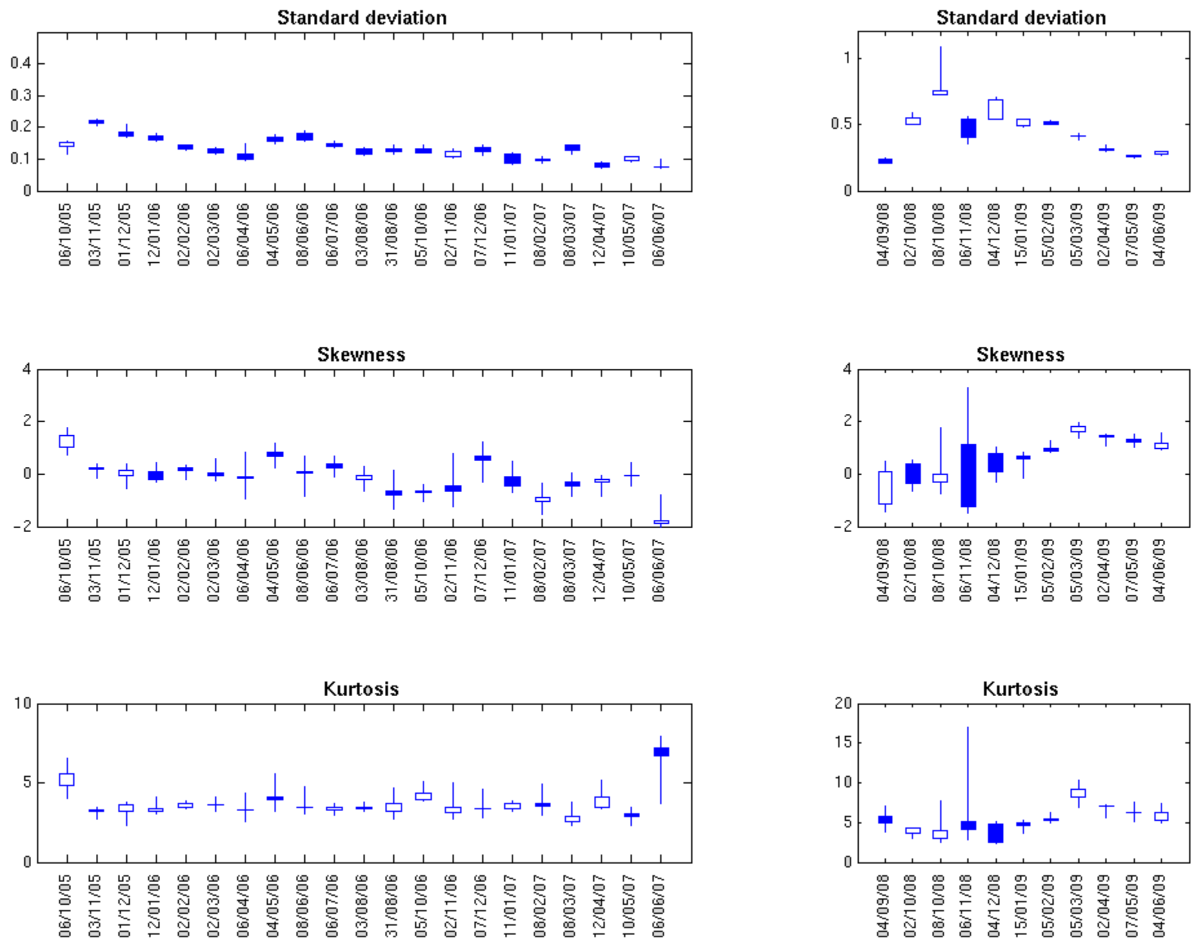

7 Determinants of market reactions

This section aims to identify drivers of short-term interest rate expectations as captured by the different moments of the implied densities during Governing Council days. For this purpose, a regression analysis is carried out to look for stronger statistical support for the drivers that were put forward as part of the case studies and intraday analysis in the previous sections.

Activity and total changes of expectations are computed based on the mean, median, standard deviation, skewness, and kurtosis of the density. Activity is measured as the average absolute 1-minute change of the corresponding moment. These activity measures are computed over three time intervals: the ten minutes following the press release, the duration of the introductory statement and the duration of the Q&A session. Furthermore, the total change in each moment is computed over each time interval, thus also allowing for direction of the change.

These market reaction variables act as dependent variables in regressions of the type used in Ehrmann and Fratzscher (2009):

$$
\begin{aligned}
Y_{PR,m,t} &= \alpha_{1,m,t} + \sum\nolimits_{i=1}^{z_1} \beta_{1,i,m} x_{i,t} + \varepsilon_{1,m,t} \quad , \\
Y_{IS,m,t} &= \alpha_{2,m,t} + \sum\nolimits_{i=1}^{z_2} \beta_{2,i,m} x_{i,t} + \gamma_{1,m}\varepsilon_{1,m,t} + \varepsilon_{2,m,t} \quad , \\
Y_{QA,m,t} &= \alpha_{3,m,t} + \sum\nolimits_{i=1}^{z_3} \beta_{3,i,m} x_{i,t} + \gamma_{2,m}\varepsilon_{1,m,t} + \delta_{1,m}\varepsilon_{2,m,t} + \varepsilon_{3,m,t} \quad ,
\end{aligned}
$$

$$
\begin{aligned}
Z_{PR,m,t} &= \alpha_{4,m,t} + \sum\nolimits_{i=1}^{z_4} \beta_{4,i,m} x_{i,t} + \varepsilon_{1,m,t} \quad , \\
Z_{IS,m,t} &= \alpha_{5,m,t} + \sum\nolimits_{i=1}^{z_5} \beta_{5,i,m} x_{i,t} + \gamma_{3,m}\varepsilon_{1,m,t} + \varepsilon_{2,m,t} \quad , \\
Z_{QA,m,t} &= \alpha_{6,m,t} + \sum\nolimits_{i=1}^{z_6} \beta_{6,i,m} x_{i,t} + \gamma_{4,m}\varepsilon_{1,m,t} + \delta_{2,m}\varepsilon_{2,m,t} + \varepsilon_{3,m,t} \quad ,
\end{aligned}
$$

where Y is the total change and Z is the activity measure described above, computed for moment m, $m = mean$, *median*, *standard deviation*, *skewness or kurtosis*, at the time of the *press release* (PR), introductory statement (IS) *or Q&A session* (QA), on Governing Council day t, $t = 1, \ldots, 32$. Two dependent variables, three intraday intervals and five moments imply 30 individual regressions in total.

Each equation explains the market reaction of one of the moments in one of the time intervals based on a number of explanatory variables. The set of explanatory variables (x_i, $i = 1, \ldots z$) varies per equation. If the total change (Y) is the dependent variable, then they are selected among the following variables:

- Decision: change of the policy rate (in basis points)
- Decision surprise: difference between the decision and what was priced in just before the decision (as derived from forward EONIA, in basis points)
- Uncertainty about the policy rate decision: measure for the direction of the uncertainty based on policy rate expectations among economists participating in a Bloomberg survey; computed as $maximum + minimum - 2 * mean$ of the survey (in basis points)
- Surprise in the release of U.S. initial jobless claims: difference between actual and forecast. Released on Thursdays at 14:30 C.E.T. (in 1,000 claims)
- Surprise in the release of U.S. continuing jobless claims: difference between actual and forecast. Released on Thursdays at 14:30 C.E.T. (in 1,000 claims)
- Code word surprise: a signed dummy indicating if mentioning or not of 'vigilance' during the introductory statement was reported as a surprise in the Q&A.

Table 2: Regression results

Press release	(1) mean	(2) median	(3) st. dev.[1]	(4) skew	(5) kurtosis[1]
	Dependent variable: *Y*				
constant	0.306	0.507	-0.549	-0.023	0.285
	(0.249)	*(0.984)*	*(0.474)*	*(0.075)*	*(0.213)*
decision	-0.035 ***	-0.066 *	0.048	0.005	-0.015
	-0.014	*(0.038)*	*(0.058)*	*(0.004)*	*(0.010)*
decision surprise	0.347 ***	1.143 ***	0.441 **	-0.056 **	0.041
	(0.098)	*(0.275)*	*(0.187)*	*(0.022)*	*(0.051)*
decision uncertainty	-0.034	-0.129 **	-0.135	0.009 *	-0.025
	-0.033	*(0.061)*	*(0.084)*	*(0.005)*	*(0.027)*
R^2	0.59	0.59	0.42	0.30	0.13

	(6) mean	(7) median	(8) st. dev.	(9) skew	(10) kurtosis
	Dependent variable: *Z*				
constant	0.142	0.672	0.321 **	0.097 ***	0.045
	(0.177)	*(0.410)*	*(0.128)*	*(0.023)*	*(0.045)*
|decision|	0.009 ***	0.017 **	0.009 ***	0.001	0.003
	(0.002)	*(0.007)*	*(0.002)*	*(0.001)*	*(0.002)*
|decision surprise|	0.107 ***	0.263 ***	0.079 ***	0.010 ***	0.023 ***
	(0.013)	*(0.042)*	*(0.011)*	*(0.003)*	*(0.006)*
|decision uncertainty|	0.010	0.020	0.001	0.001	-0.001
	(0.006)	*(0.018)*	*(0.003)*	*(0.001)*	*(0.002)*
kurtosis	-0.057	-0.221 *	-0.092 ***	-0.014 **	0.005
	(0.050)	*(0.113)*	*(0.034)*	*0.007)*	*(0.014)*
R^2	0.79	0.69	0.79	0.37	0.48

Initial statement	(11) mean	(12) median	(13) st. dev.[2]	(14) skew	(15) kurtosis[2]
	Dependent variable: *Y*				
constant	0.253	0.382	-0.072	-0.038	0.024
	(0.347)	*(0.462)*	*(0.300)*	*(0.061)*	*(0.093)*
decision	0.019	-0.001	-0.028 ***	0.004 **	0.003
	(0.015)	*(0.016)*	*(0.007)*	*(0.002)*	*(0.004)*
decision surprise	-0.104	-0.125 *	-0.111	0.010	-0.014
	(0.094)	*(0.069)*	*(0.078)*	*(0.009)*	*(0.020)*
decision uncertainty	0.013	0.010	0.023 ***	0.002	0.004
	(0.020)	*(0.020)*	*(0.008)*	*(0.002)*	*(0.006)*
initial claims	0.004	-0.014	0.000	0.007 **	-0.003
	(0.017)	*(0.022)*	*(0.009)*	*(0.004)*	*(0.009)*
continuing claims	-0.003	-0.003	0.003	-0.000	-0.002 *
	(0.006)	*(0.006)*	*(0.003)*	*(0.0006)*	*(0.001)*
code word surprise	2.377 ***	3.252 ***	-0.195	-0.293 ***	0.056
	(0.652)	*(0.763)*	*(0.307)*	*(0.111)*	*(0.139)*
resid PR	0.115	0.084 *	-0.086 **	0.036	-0.008
	(0.162)	*(0.047)*	*(0.039)*	*(0.067)*	*(0.030)*
R^2	0.32	0.35	0.32	0.24	0.11

	(16) mean	(17) median	(18) st. dev.	(19) skew	(20) kurtosis
	Dependent variable: *Z*				
constant	0.330 *	0.798 **	0.121	0.039	-0.184 *
	(0.191)	*(0.360)*	*(0.179)*	*(0.068)*	*(0.107)*
|decision|	0.003	0.002	0.002	0.001	0.002
	(0.002)	*(0.003)*	*(0.002)*	*(0.000)*	*(0.002)*
|decision surprise|	0.046 **	0.068 **	0.029	0.003	0.010
	(0.021)	*(0.033)*	*(0.022)*	*(0.002)*	*(0.009)*
|decision uncertainty|	-0.002	-0.006	-0.002	-0.002 *	-0.003 **
	(0.002)	*(0.004)*	*(0.003)*	*(0.001)*	*(0.001)*
kurtosis	-0.055	-0.110	-0.053	-0.011	0.018
	(0.045)	*(0.071)*	*(0.045)*	*(0.013)*	*(0.028)*
|initial claims|	-0.002	-0.003	-0.003	-0.001	-0.001
	(0.001)	*(0.004)*	*(0.003)*	*(0.001)*	*(0.003)*
|continuing claims|	0.000	0.002	-0.001	-0.000	-0.001
	(0.001)	*(0.002)*	*(0.001)*	*(0.000)*	*(0.001)*
code word	0.047	0.121	0.039	0.057 *	0.099 *
	(0.084)	*(0.151)*	*(0.070)*	*(0.030)*	*(0.060)*
|code word surprise|	0.101	0.091	-0.066	0.053 *	-0.002
	(0.072)	*(0.149)*	*(0.068)*	*(0.028)*	*(0.064)*
time IS	0.012	0.012	0.025 **	0.008 **	0.021 **
	(0.014)	*(0.020)*	*(0.011)*	*(0.004)*	*(0.011)*
resid PR	-0.078	-0.047	0.014	0.288	0.502
	(0.086)	*(0.046)*	*(0.210)*	*(0.259)*	*(0.420)*
R^2	0.53	0.43	0.39	0.49	0.43

Q&A

	Dependent variable: Y				
	(21) mean	(22) median	(23) st. dev.[3]	(24) skew	(25) kurtosis[3]
constant	-0.126	-0.081	-0.476 *	-0.086 **	0.152
	(0.366)	*(0.404)*	*(0.286)*	*(0.041)*	*(0.097)*
decision	-0.062 ***	-0.042 ***	0.038	-0.003	0.013
	(0.012)	*(0.012)*	*(0.024)*	*(0.002)*	*(0.012)*
decision surprise	-0.385 ***	-0.353 ***	0.263	-0.010	-0.022
	(0.062)	*(0.054)*	*(0.260)*	*(0.008)*	*(0.032)*
decision uncertainty	-1.780	-0.460	-6.021	-0.118	-1.609
	(1.351)	*(1.387)*	*(4.040)*	*(0.172)*	*(1.335)*
code word surprise	1.891	1.710	-0.238	0.123 *	0.101
	(1.423)	*(1.375)*	*(0.323)*	*(0.075)*	*(0.161)*
resid PR	0.345 **	0.018	0.159	-0.024	-0.062
	(0.141)	*(0.032)*	*(0.154)*	*(0.068)*	*(0.041)*
resid IS	0.031	-0.041	-1.241 *	-0.004	-0.446
	(0.207)	*(0.187)*	*(0.665)*	*(0.111)*	*(0.394)*
R^2	0.60	0.49	0.37	0.11	0.22

	Dependent variable: Z				
	(26) mean	(27) median	(28) st. dev.	(29) skew	(30) kurtosis
constant	0.107	0.346	0.198 **	0.054 *	-0.092
	(0.093)	*(0.261)*	*(0.084)*	*(0.033)*	*(0.124)*
|decision|	0.002	0.003	0.003 ***	0.000	0.001
	(0.001)	*(0.002)*	*(0.001)*	*(0.000)*	*(0.001)*
|decision surprise|	0.014 ***	0.050 ***	0.008 ***	0.001	0.002
	(0.005)	*(0.014)*	*(0.002)*	*(0.001)*	*(0.003)*
|decision uncertainty|	0.281	0.321	-0.063	-0.086 **	-0.295 ***
	(0.190)	*(0.239)*	*(0.054)*	*(0.036)*	*(0.089)*
kurtosis	0.016	-0.032	-0.017 **	-0.005	0.033 ***
	(0.014)	*(0.025)*	*(0.008)*	*(0.004)*	*(0.010)*
|code word surprise|	0.074	0.064	0.119	0.026	0.075
	(0.112)	*(0.163)*	*(0.103)*	*(0.016)*	*(0.052)*
time QA	0.001	0.004	-0.000	0.001	0.002
	(0.002)	*(0.005)*	*(0.002)*	*(0.001)*	*(0.002)*
resid PR	-0.023	-0.012	0.127 ***	0.165 ***	0.419 ***
	(0.089)	*(0.051)*	*(0.047)*	*(0.061)*	*(0.106)*
resid IS	0.246 **	0.252	0.384 ***	0.453 ***	0.643 ***
	(0.113)	*(0.159)*	*(0.072)*	*(0.099)*	*(0.153)*
R^2	0.60	0.59	0.59	0.66	0.75

Notes:
coefficient estimates and their standard errors in brackets
the standard errors are Newey-West heteroskedasticity and autocorrelation consistent standard errors
*, **, and *** denote significance at the 90%, 95% and 99% level respectively
1: regressed on absolute value of decision, decision surprise and decision uncertainty
2: regressed on absolute value of decision, decision surprise, decision uncertainty, initial claims, continuing claims and code word surprise
3: regressed on absolute value of decision, decision surprise, decision uncertainty and code word surprise

If activity is the dependent variable, then the absolute values of the above variables are used as explanatory variables and also the following variables are considered:

- Uncertainty surrounding the 3-month EURIBOR three months ahead: average kurtosis of the implied density in the morning of the Governing Council day (measured between 11:00 and 12:00 C.E.T.).
- Perceived code word: dummy variable indicating if the perceived code word 'vigilance' was mentioned during the introductory statement.
- Duration of the introductory statement (in minutes)
- Duration of the question and answer session (in minutes)

In addition to the explanatory variables, the residuals of the press release equation feed into the initial statement equation, and both residuals of the press release and initial statement equations feed into the Q&A equation. Introducing these residuals as explanatory variables in the next time window allows testing whether unobserved factors that drove market reactions have persistent effects on the next time windows.

The regression analysis identifies a number of these drivers of expectation developments as statistically and economically significant. The regression results are reported in Table 2 with robust standard errors. The impact of each explanatory variable – ceteris paribus – is discussed below one-by-one.

7.1 Press release

Turning first to the activity following a press release in equations 6-10, activity in all the moments is significantly higher if the decision contains a surprise component. This is completely in line with the hypothesis that a rebalancing of positions takes place following a surprise. It shows that the whole density and the expectations it represents can change in case of a surprise. The fact that not only the consensus adapts (i.e. a simple shift of the density) following a surprise supports the hypothesis that a surprise triggers changes in the relative views market participants hold as information about the future path of the policy rate is interpreted differently. Next, looking at the impact on total changes in each of the moments in equations 1-5, a surprise in the decision is also statistically significant here, except for the kurtosis. The coefficients also have the expected sign with a positive surprise increasing the mean and median, decreasing skewness with the hump of the density or bulk of the expectations moving right, and increasing the dispersion of expectations and hence uncertainty about the future 3-month EURIBOR outcome. Furthermore, the size of the impact is also in line with expectations, where e.g. one third of the policy rate surprise is passed on to the mean of the implied density.

Rather surprising is that activity in the mean, median and standard deviation increases in case of a rate change compared to when the rate remains unchanged (eq. 6-10). A priori, one would not expect such an effect to be statistically significant and instead only the surprise component to matter. A similar relation for the mean was also found in Ehrmann and Fratzscher (2009). The result suggests that the density shifts in case of a rate change, but that its shape and relative expectations as expressed in terms of skewness and kurtosis are not significantly affected. Turning to the impact on the total change (eq. 1-5), however, the change in the policy rate is found to have a significant effect only on the mean, and its sign is negative and hence different from what expected. By taking a closer look at the data, it becomes clear that this sign is likely to be driven by the events in the second sub-period where big rate cuts were associated with positive changes in the mean and median. The latter is consistent with the increases observed in the futures rates, which were seen as evidence that some rate cuts were smaller than what markets had expected. If in addition the decision surprise variable did not capture the surprise effect fully, then the decision variable would capture this effect partly with a negative coefficient. Therefore, the statistical significance of the decision variable may be somewhat spurious and driven by surprises in the second period. In any case the economic significance is small.

Uncertainty about a policy rate decision, as captured by the survey indicator, did not significantly trigger activity in any of the moments (eq 6-10), and when allowing for direction only proofs significant for the total change in the median and skew (eq. 1-5). The insignificance for the mean can be expected since the market consensus does not capture the heterogeneity in market expectations. For the other moments, one may expect stronger changes following a rate announcement when there was prior uncertainty about the decision. The coefficients carry the expected signs, but any prior disagreement on the policy rate among the economists participating in the survey did not predict statistically significant movements in the expectations surrounding the 3-month EURIBOR three months ahead.

Higher uncertainty surrounding the three-month EURIBOR three months ahead, as expressed by implied density kurtosis in the morning of a Governing Council day, is significantly linked with lower activity in the median, standard deviation and skewness (eq. 6-10). This finding may be explained by high kurtosis (which was especially present in the second period) being more persistent and hence correlated with smaller changes of the density. In this respect, the effect may be somewhat mechanic where stronger expectation shifts are required to achieve the same changes in the median, standard deviation and skewness when kurtosis is big than when kurtosis is small. Overall, the weak impact of uncertainty is consistent with persistence in the expectations around 3-month EURIBOR in three months' time once the decision surprise has been controlled for in the regression.

7.2 Introductory statement

The explanatory variables of the reactions during the introductory statement can be divided in three groups. As a first set of explanatory variables, the variables and residual from the press release equation are used to check if the decision has further effects during the press conference and if this is related to the size of the decision, the surprise in the decision, prior uncertainty or persistent effects of unobserved variables driving the press release reaction. As a second set of explanatory variables, the two U.S. jobless claim variables enter the equation that allow checking if the activity during the introductory statement is related to the release of this data which coincides with the start of the introductory statement. As a third set, two dummy variables are added to check how relevant the use of code words during the first period was. In particular, markets are believed to have attached significance to the mentioning of 'vigilance' during the statement as an indicator of a hike at the next meeting.

The decision variable comes in strongly statistically significant as a driver of the standard deviation (eq. 13), where Governing Council days with (bigger) policy rate changes decrease dispersion of the expectations during the introductory statement. This would support the hypothesis that the explanation given during a statement that follows a rate change reduces uncertainty about future rates more than when rates remained unchanged. Given that the sample focuses on periods of policy rate changes, the result is not driven by long periods of constant policy rates associated with little change in policy rate uncertainty.

A surprise component in the decision is found to explain the increase in activity of the mean and median, but not of the other moments (eq. 16-20). It also no longer drives the total change in any of the moments during the introductory statement (eq. 11-15). This result suggests that any explanation as part of the introductory statement following a surprise still affects the consensus expectation, but not the shape of the density around it in a significant way.

The uncertainty variable and the residual from the press-release regression show mixed results about the role of the introductory statement. Higher uncertainty about a policy rate decision is significantly associated with lower activity in the skewness and kurtosis during the statement (eq. 19-20). Moreover, it is found to significantly increase the dispersion during the statement. Thus, initial uncertainty about the policy rate decision is associated with persistent expectations or even increasing dispersion of expectations about the future EURIBOR during the statement (eq. 13). This result would rather speak against the hypothesis that the statement brings clarifications and decreases uncertainty, and instead potentially even raises new issues. In contrast, the residual from the press release regression is statistically significant for the median and standard deviation (eq. 12-13), where a positive dispersion residual of the press release is followed by a significant decrease in the dispersion during the introductory statement. This suggests that an 'excess' increase in uncertainty about the future EURIBOR due to the press release is offset during the introductory statement, suggesting the statement's explanations do matter.

The fact that the start of the press conference coincides with the release of U.S. jobless claims motivates two U.S. jobless claim surprise variables to be added to the initial-statement equations. Turning to the regression results, surprises contained in the release of initial and continuing U.S. jobless claims appear to have had less power in explaining reactions during the statement than what results in the literature suggested so far based on earlier samples. The releases may hence have been somewhat less relevant in the present sample. Still, an upward surprise in the initial jobless claims is found to significantly shift the hump of the density to the left (eq. 14). In other words, negative U.S. news implies lower expected euro area rates for market participants. Also a bigger surprise in the continuing jobless claims is found to significantly decrease kurtosis which would mean that it focuses expectations slightly better (eq. 15).

The presence of perceived code words clearly guided expectations about the next decision. The dummy indicating the mentioning of a perceived code word is not significant, but this is also what would be expected once surprises in the perceived code word are controlled for. By contrast, code word surprises significantly move the mean, median and skew (eq 11-15). The dummy indicating the code word surprise only identifies three cases based on comments in the Q&A. In particular, on 6 April 2006 and 11 January 2007, 'vigilance' was expected, but did not arrive, while on 6 July 2006 it arrived earlier than expected. For the other Governing Council meetings of the first period, the occurrence of 'vigilance' was apparently correctly anticipated by most market participants. Since code word surprises help explain the overall change, the existence of perceived code words clearly guided expectations.

Overall, judging from the coefficient of determination (R^2) of the regressions, the explanatory variables leave big parts of the variation of the moments during the introductory statement unexplained. As most variables (apart from the code word dummies) rather capture conditions surrounding the introductory statement instead of what is being said during the statement, this is not surprising. More detailed high-frequency analysis comparing topical phrases to their immediate market reaction is likely to confirm the importance of the content. Interesting in this respect, however, is the significance of the length of the introductory statement for the increased activity of the dispersion, skew and kurtosis. As the dependent variable is an average over the statement, one would a priori not expect its size to be dependent on time. Still, this variable is found to be significant. Since longer statements are likely to indicate that additional economic information is provided, this variable would load on the content of this extra information and its significance point out the relevance of the content.

7.3 Question and answer session

Finally, a number of explanatory variables for the market reaction during the Q&A are tested. As for the reactions during the introductory statement, much will depend on the content of the session, but as this is difficult to quantify we concentrate here on the impact of surrounding conditions and proxy variables. As explanatory variables, the indicators describing the decision, the code word surprise and the residuals from the two previous equations are considered.

A decision surprise is found to explain increased activity in the mean, median and standard deviation, which at first suggests that the questions posed during the session still concern the interpretation of the surprise (eq. 26-30). However, turning to the total change of the moments during the Q&A, the decision and its surprise come in as statistically significant drivers of the mean and median, but with an unexpected negative sign (eq. 21-22). This result suggests that many direction reversals took place during the Q&A when there was a decision surprise. This is not unlikely since the Q&A typically concentrates on extracting information about next decisions and the outlook and the reaction may differ from the current decision surprise. In this context, Brand et al. (2006) confirmed the view that ECB communication during the press conference may result in significant changes in market expectations of the path of monetary policy. Looking closer at the underlying data, decision surprises belonged particularly to the second period.

In contrast, a code word surprise during the introductory statement is not really found to have a significant effect during the Q&A, which suggests that these events were clearly understood and the market took them into account quickly.

The residuals of the previous regressions are often statistically significant indicating persistent effects in the activity of most moments during the Q&A (eq. 26-30). Given that the impact of the content of the introductory statement is captured by its residual this is not a surprise. However, looking at the equations for the total changes, hardly any significant impact is found (eq. 21-25). The insignificance of the residuals in explaining the change of moments, while significantly explaining increased activity in those moments, again suggests that reversals in the direction were common during the Q&A.

8 Conclusion

Measures of the expectations held by financial market participants about the outcome of a certain asset price have been refined over time in the literature. The estimation of risk-neutral probability density functions has proved a powerful tool in this field since it summarises the total set of likely outcomes and probabilities attached by the market. Another advantage - stressed even more by this paper - is that they can be extracted at almost any moment in time since the estimation is based on financial market data. So far, only the daily frequency had been explored for a wide set of instruments.

This paper extracts such densities based on option prices up to tick frequency for the first time in the literature. They have the clear benefit that – as was demonstrated in a few case studies – one can zoom in on certain events or announcements and judge the immediate market reactions, thereby minimising the bias caused by any other information hitting the market. Furthermore, the intraday densities are shown to offer additional information to the interpretation of intraday futures and forward rates, which in fact capture only the average or consensus view of the market. More specifically, the densities reflect the dispersion and symmetry of the expectations, thereby giving the policy maker an idea of the relative expectations and uncertainty in the market, and market participants an idea of the risks in the market.

A non-parametric estimator based on fitting implied volatility curves, as was developed in the literature, was applied to tick data on three-month EURIBOR futures and options to estimate option-implied densities representing expectations of three-month EURIBOR three months ahead. The paper discussed this estimator in an intraday setting and introduced an efficient method of pre-filtering the data to impose no-arbitrage conditions as required by option pricing theory. The density estimates indicated that the estimator is robust to market microstructure noise by producing stable risk-neutral densities. At the same time, when information hits the market the densities adapt quickly and meaningfully, indicating that the estimator is flexible enough to capture changes in expectations. An estimator of the noise size shows a relatively small impact and allows it to be taken into account when interpreting developments. As a result, the application succeeded in demonstrating the feasibility of intraday estimation.

An economic assessment of the announcement effects of ECB communication on short-term interest rate expectations was carried out. The sample consisted of 32 days on which the ECB Governing Council took a policy rate decision. The intraday patterns of the statistical moments of the implied densities show a significant shock in activity following the press release and significantly increased activity during the press conference, showing the relevance of both their content. All considered moments (mean, median, standard deviation, skew and kurtosis) show such patterns. Furthermore, apart from reaching very distinct levels between different days, it was shown that there can be large movements in moments within a Governing Council day, in particular during the financial crisis period.

Finally, a regression analysis identified a number of drivers of the expectation changes following the press release and during the press conference. A surprise in the policy rate decision, as perceived by the market, was found to significantly affect the entire density, hence not only the consensus view but also the relative positioning of expectations. Uncertainty about the policy rate decision and about the future EURIBOR outcome were also found to be relevant, but the evidence for this was not as strong. A code word, as perceived by the market in predicting rate hikes, was found to have guided expectations. This confirms the value attached by markets to perceived patterns in the wording used by the central bank and rate decisions. In addition, the paper found indications that the overall content of the introductory statement and Q&A session was relevant in driving expectations. While the sample size was limited and the study is hence not exhaustive, the results are telling. Future research that explicitly tests the economic statements during the press conference against the immediate market reaction may provide further statistical evidence of the impact of the content.

Overall, the relevance of the press release and conference as a communication tool has been confirmed by analysis of market expectation developments at high frequency. This also holds for both the introductory statement and the question and answer session of the press conference, which, given the (continued) high activity during these sessions, appears to provide additional information to markets. The sensible interpretation that can be given to the regression results when identifying drivers of the reactions also indicates that the information is not simply adding noise that could offer an alternative explanation for the increased activity. Instead, expectations are guided in specific directions. This provides support for the use of a press conference following policy rate announcements, as is currently practised solely by the ECB. For example, the recent introduction of quarterly press briefings by the Federal Reserve falls into the same category and their impact on the relevant asset prices could be studied with the tools presented here.

Appendix

A.1. Sample of Governing Council days

(sample in shaded area)

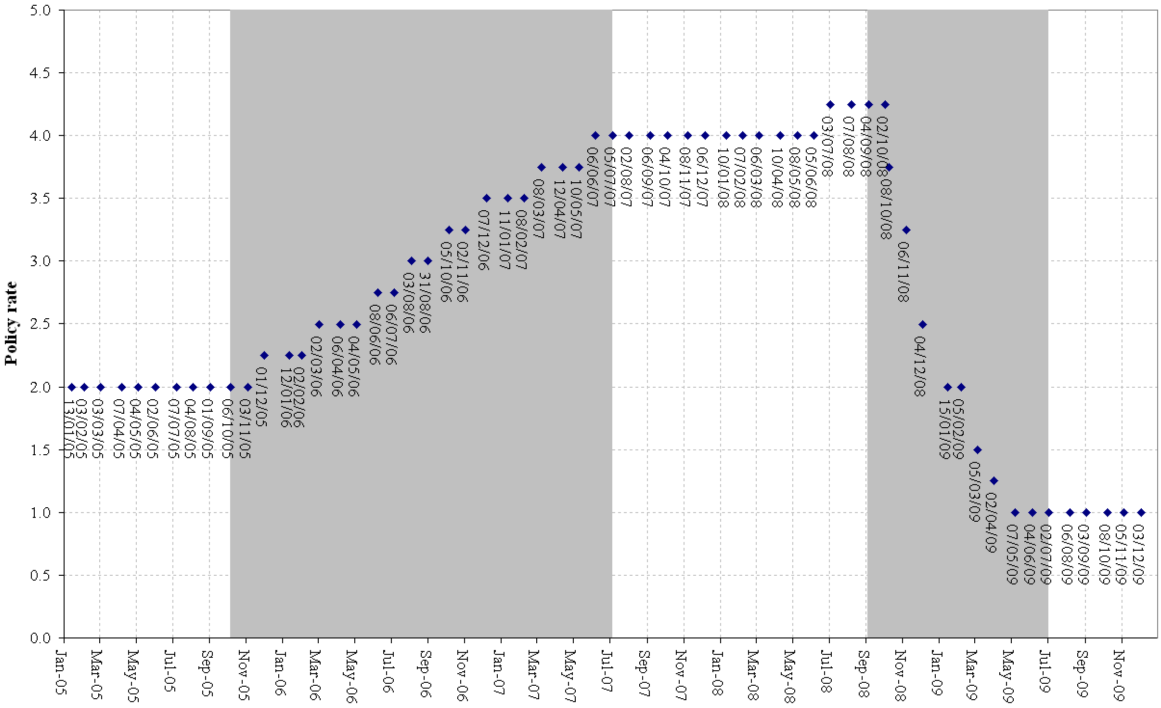

A.2. Impact of Governing Council days on the density moments

Table 3: Average (absolute) change in moments on Governing Council days versus other days

	average absolute change					average change	
	mean	**median**	**st. dev.**	**skew**	**kurtosis**	**st. dev.**	**kurtosis**
GC days	3.8***	3.7***	1.7***	0.16***	0.30	-1.1***	-0.06
Other Thursdays	2.3	2.3	0.9	0.12	0.29	0.1	-0.02

Notes: the averages are computed based on daily implied density estimates during the period 1999-2010 for 3-month EURIBOR in three months' time. The values for the mean, median and standard deviation are in basis points. *** denotes significance at the 99% level for a t-test of difference in averages. 'Governing Council days' are tested against 'Other Thursdays'.

A.3. Intraday pattern of tick activity (per minute)

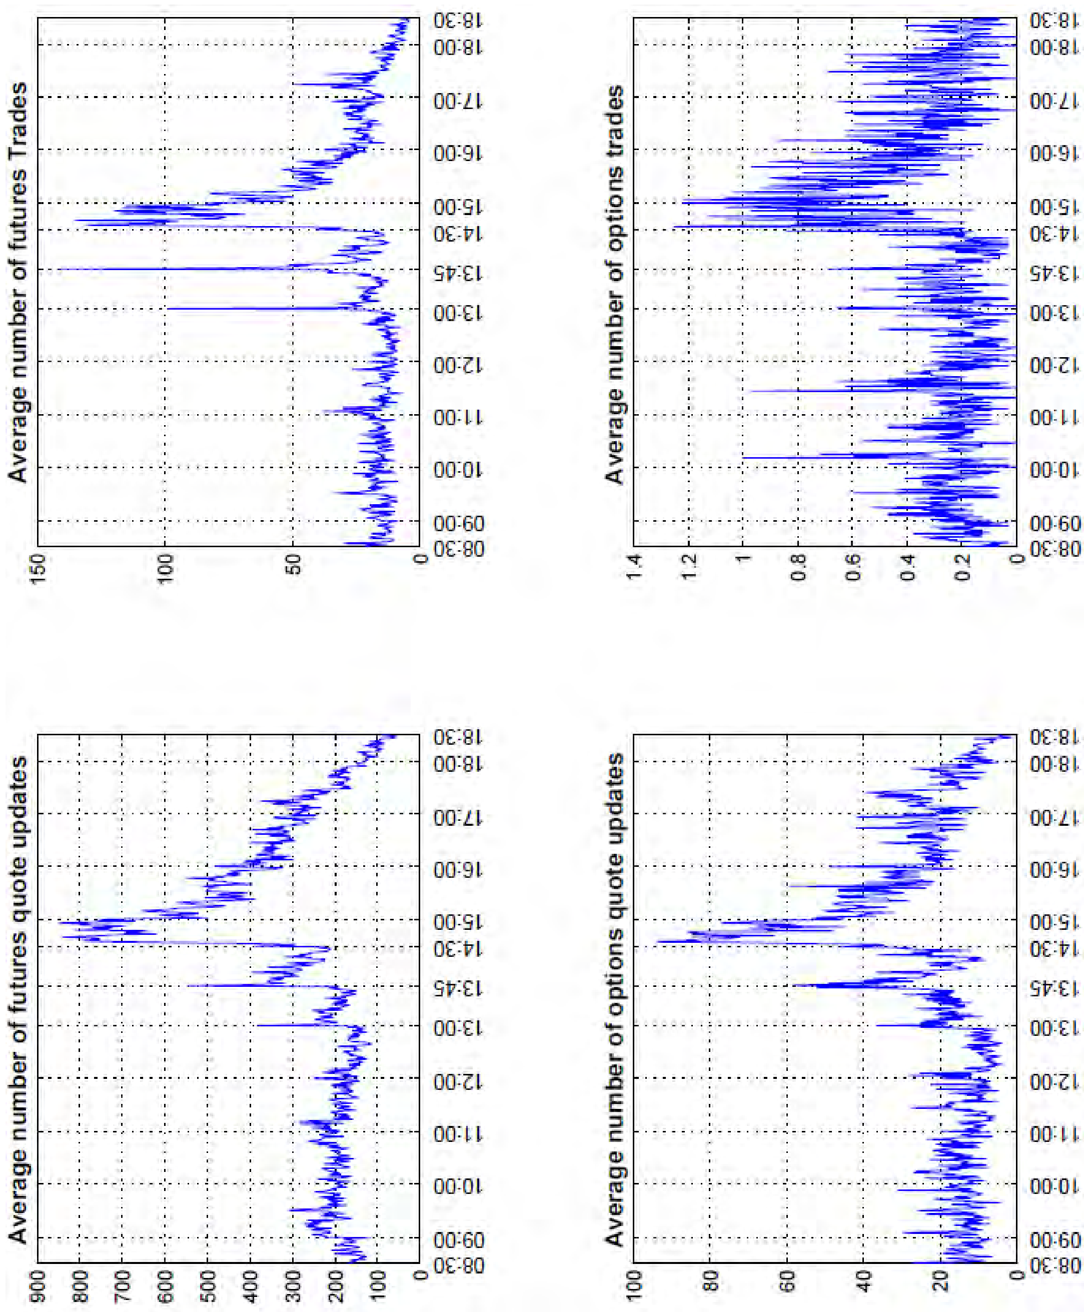

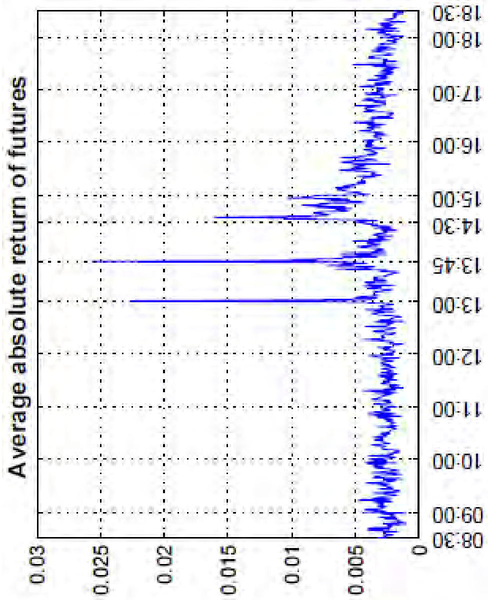
Average absolute return of futures
0.03
0.025
0.02
0.015
0.01
0.005
0
08:30
09:00
10:00
11:00
12:00
13:00
13:45
14:30
15:00
16:00
17:00
18:00
18:30

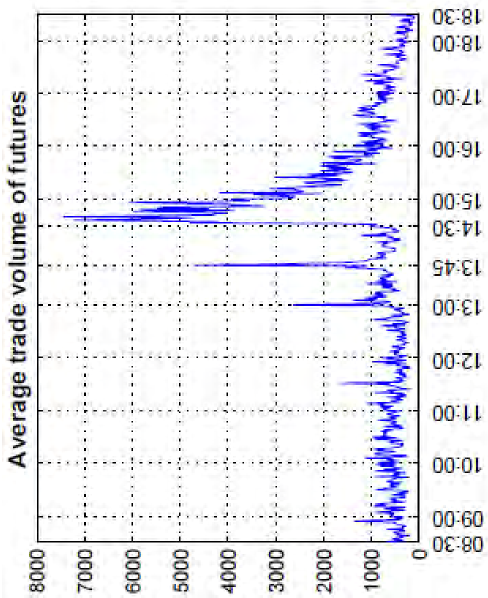
Average trade volume of futures
8000
7000
6000
5000
4000
3000
2000
1000
0
08:30
09:00
10:00
11:00
12:00
13:00
13:45
14:30
15:00
16:00
17:00
18:00
18:30

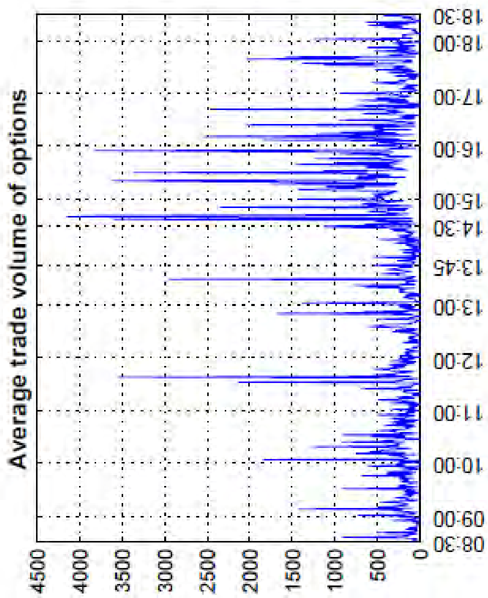
Average trade volume of options
4500
4000
3500
3000
2500
2000
1500
1000
500
0
08:30
09:00
10:00
11:00
12:00
13:00
13:45
14:30
15:00
16:00
17:00
18:00
18:30

A.4. Intraday pattern of implied density moments activity and level (per minute)

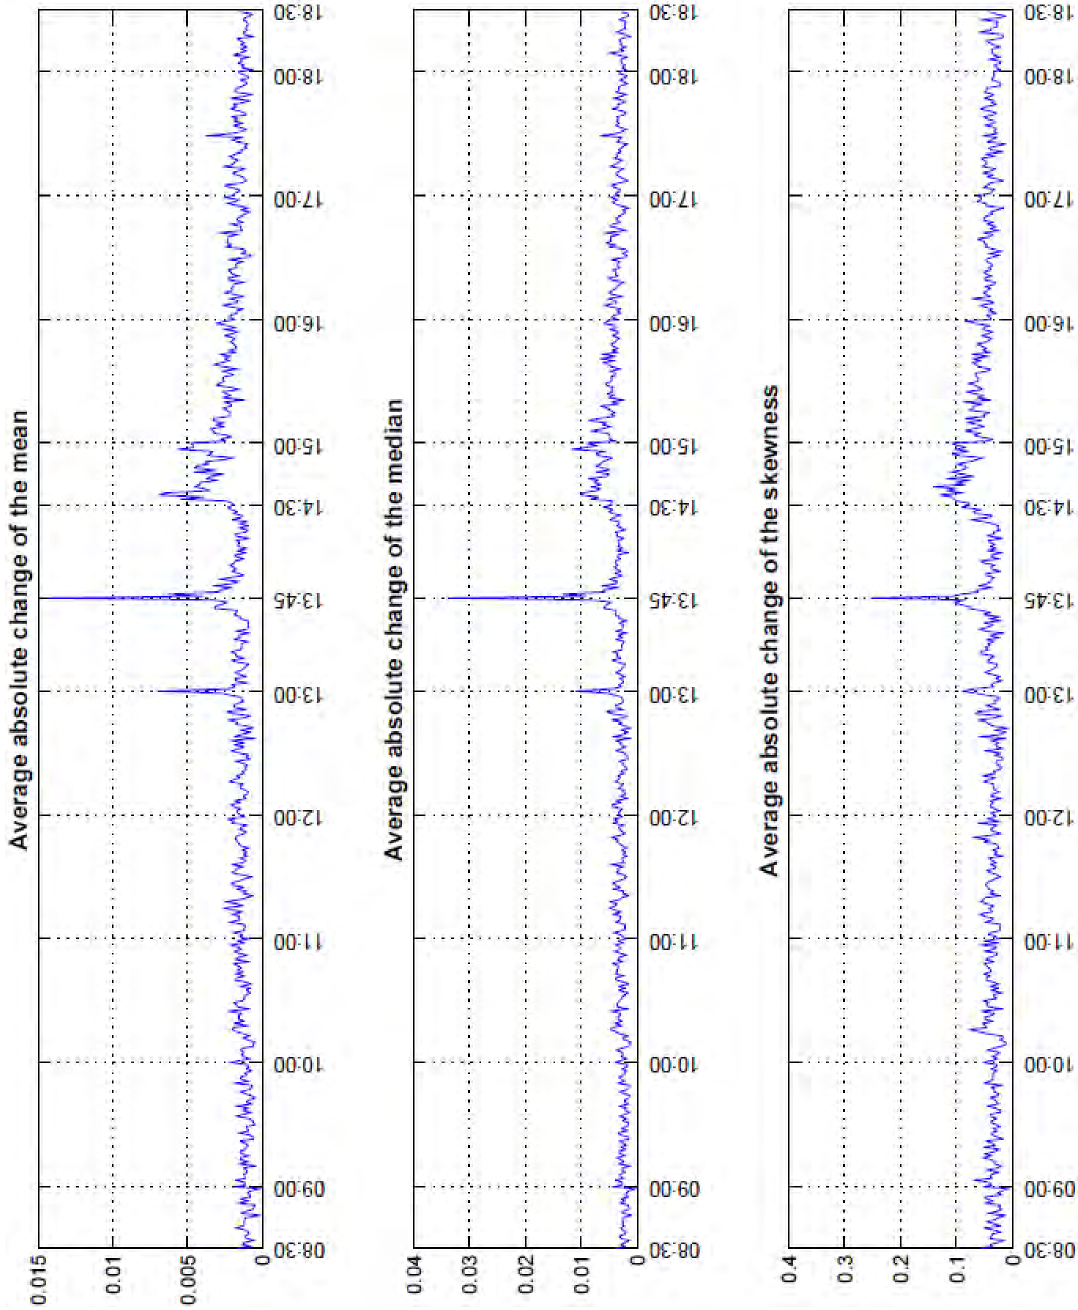

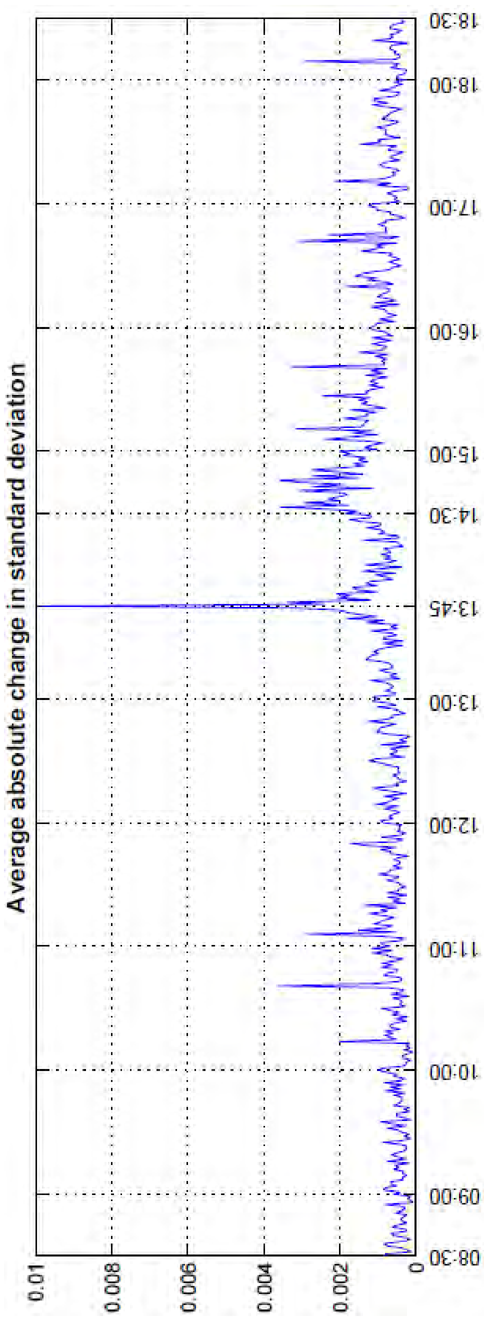
Average absolute change in standard deviation
0.01
0.008
0.006
0.004
0.002
0
08:30
09:00
10:00
11:00
12:00
13:00
13:45
14:30
15:00
16:00
17:00
18:00
18:30

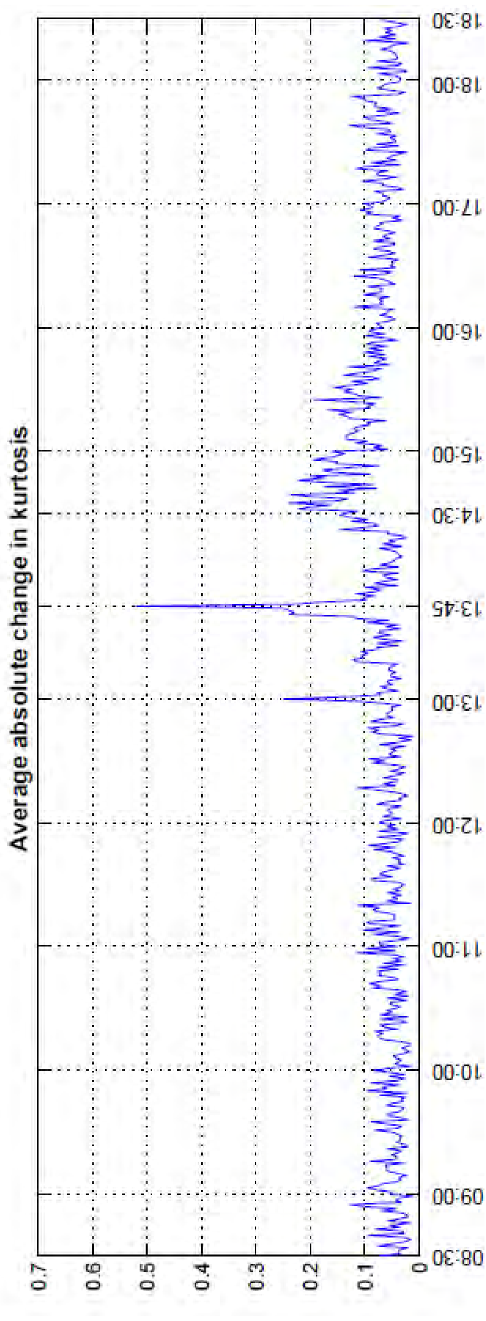
Average absolute change in kurtosis
0.7
0.6
0.5
0.4
0.3
0.2
0.1
0
08:30
09:00
10:00
11:00
12:00
13:00
13:45
14:30
15:00
16:00
17:00
18:00
18:30

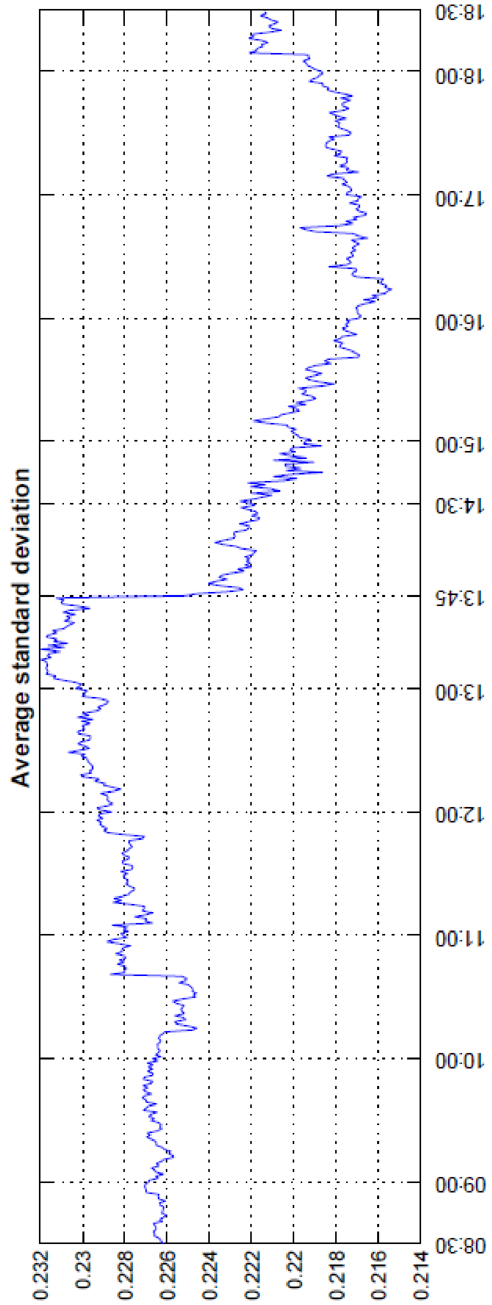
Average standard deviation
0.232
0.23
0.228
0.226
0.224
0.222
0.22
0.218
0.216
0.214
08:30
09:00
10:00
11:00
12:00
13:00
13:45
14:30
15:00
16:00
17:00
18:00
18:30

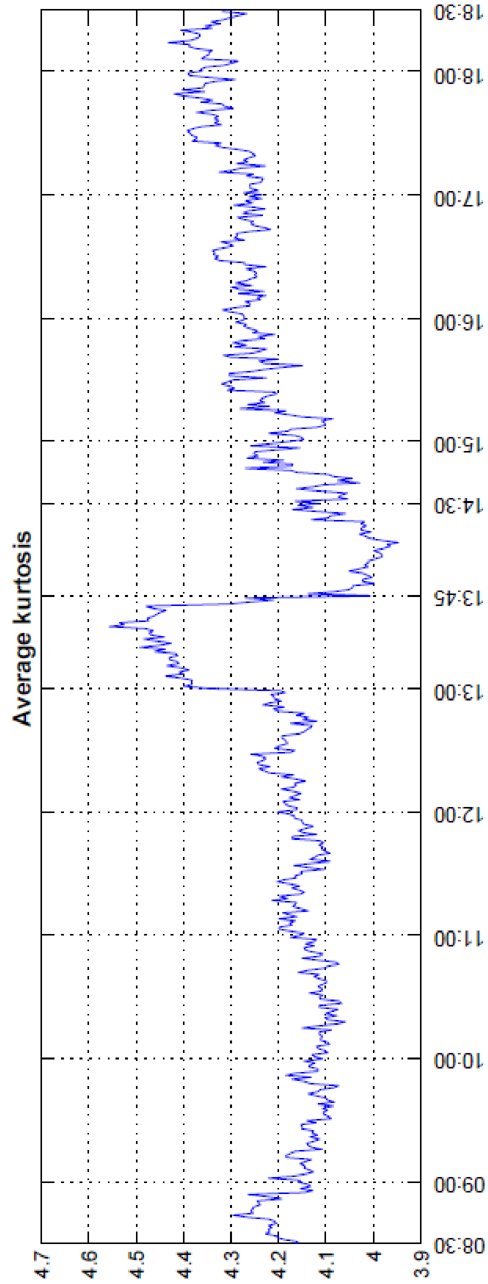
Average kurtosis
4.7
4.6
4.5
4.4
4.3
4.2
4.1
4
3.9
08:30
09:00
10:00
11:00
12:00
13:00
13:45
14:30
15:00
16:00
17:00
18:00
18:30

References

[1] ANDERSEN, A. B., AND T. WAGENER (2002): "Extracting Risk Neutral Probability Densities by Fitting Implied Volatility Smiles: Some Methodological Points and and Application to the 3M EURIBOR Futures Option Prices," *ECB Working Paper*, 198.

[2] ANDERSSON, M. (2007): "Using Intraday Data to Gauge Financial Market Responses to Fed and ECB Monetary Policy Decisions," *ECB Working Paper*, 726.

[3] BAHRA, B. (1997): "Implied Risk-Neutral Probability Density Functions from Option Prices: Theory and Application," *Bank of England Working Paper*, 66.

[4] BLACK, F., AND M.J. SCHOLES (1973): "The Pricing of Options and Corporate Liabilities," *Journal of Political Economy*, 81(3), 637-654.

[5] BLACK, F. (1976): "The Pricing of Commodity Contracts," *Journal of Financial Economics*, 3, 167-179.

[6] BLISS, R., AND N. PANIGIRTZOGLOU (2002): "Testing the Stability of Implied Probability Density Functions," *Journal of Banking and Finance*, 26, 381-422.

[7] BRAND, C., D. BUNCIC, AND J. TURUNEN (2006): "The Impact of ECB Monetary Policy Decisions and Communication on the Yield Curve," *ECB Working Paper*, 657.

[8] BREEDEN, D., AND R. LITZENBERGER (1978): "Prices of State-Contingent Claims Implicit in Option Prices," *Journal of Business*, 51, 621-652.

[9] CAMPA, J.M., P.H.K. CHANG, AND R.L. REIDER (1997): "ERM Bandwidths for EMU and After: Evidence from Foreign Exchange Options," *Economic Policy*, 24, 55-89.

[10] CASTRÉN, O. (2004): "Do Options-Implied RND Functions on G3 Currencies Move around the Times of Interventions on the JPY/USD Exchange Rate?," *ECB Working Paper*, 410.

[11] COOPER, N. (1999): "Testing Techniques for Estimating Implied RNDs From The Prices of European-Style Options," Proceedings of the Workshop on Estimating and Interpreting Probability Density Functions held at the BIS on 14 June 1999.

[12] COX, J.C., AND S.A. ROSS (1976): "The Valuation of Options for Alternative Stochastic Processes," *Journal of Financial Economics*, 3, 145-166.

[13] DE VINCENT-HUMPHREYS, R., AND J.M. PUIGVERT GUTIÉRREZ (2010): "A Quantitative Mirror on the EURIBOR Market using Implied Probability Density Functions," ECB Working Paper, 1281.

[14] ECB (2007): "Communicating Monetary Policy to Financial Markets," *ECB Monthly Bulletin Article*, April 2007.

[15] ECB (2011): "The Information Content of Option Prices during the Financial Crisis", *ECB Monthly Bulletin Article*, February 2011.

[16] EHRMANN, M., AND M. FRATZSCHER (2009): "Explaining Monetary Policy in Press Conferences," *International Journal of Central Banking*, 5(2), 41-84.

[17] LYNCH, D., N. PANIGIRTZOGLOU AND G. KAPETANIOS (2004): "Using Option Prices to Measure Financial Market Views about Balances of Risk to Future Asset Prices," Bank of England Quarterly Bulletin, Winter 2004, 442-454.

[18] Malz, A.M. (1998): "Option Prices and the Probability Distribution of Exchange Rates," in: *Currency Options and Exchange Rate Economics*, ed. C. Zhaohui, World Scientific, Singapur, 108-137.

[19] Mandler, M. (2002): "Comparing Risk-Neutral Probability Density Functions Implied by Option Prices – Market Uncertainty and ECB-Council Meetings," Universität Giessen.

[20] NYSE LIFFE (2009): "NYSE Liffe Trading Procedures," www.euronext.com.

[21] Shimko, D. (1993): "Bounds of Probability," *Risk*, 6(4), 33-37.

[22] Zhang, L., P.A. Mykland and Y. Aït-Sahalia (2005): "A Tale of Two Time Scales: Determining Integrated Volatility with Noisy High Frequency Data, *Journal of the American Statistical Association*, 472, 1394-1411.