UNIVERSITAT DE BARCELONA

# 3-month Euribor expectations and uncertainty using option-implied probability densities

Josep Maria Puigvert

Universitat de Barcelona

# 3-MONTH EURIBOR EXPECTATIONS AND UNCERTAINTY USING OPTION-IMPLIED PROBABILITY DENSITIES

Doctoral thesis presented by Josep Maria Puigvert
in the Doctoral Program of Mathematics and Informatics
Directed by Dr. Josep Fortiana Gregori

The doctoral candidate | The director

Barcelona, October 2015

*3-month Euribor expectations and uncertainty using option-implied probability densities*

Author: Josep Maria Puigvert Gutiérrez
Advisor: Dr Josep Fortiana Gregori

The views expressed in this doctoral thesis are the ones of the author and do not necessarily represent those of the European Central Bank.
Text printed in Frankfurt am Main.
First edition, October 2015

*to Ramon and Maria Lluïsa*
*to Eulàlia, Pau, Laia and Teo*

# Preface

This work aims to contribute to the study and interpretation of the 3-month Euribor risk-neutral option-implied probability density functions (PDFs). The focus is mainly on the study of daily options on 3-month Euribor futures, but the estimation is also extended to tick by tick data.

This type of models have already been widely used in some other markets, like the equity and index markets, and several authors have already contributed to the existing literature. However, this work focuses on the interest rate derivatives market and, in particular, on the Euribor one, which so far was rather unexplored. In particular, the study analyses Euribor options from the first trading date until the present time. Moreover, the study also focuses on how this type of functions reacted during periods of economic crisis, financial turbulences and around ECB Governing Council decisions.

In addition, this works shows how this type of functions can also be used to forecast the Euribor futures rates by using real-world PDFs which are derived from risk-neutral PDFs. The ratio between the risk and the real-world PDFs, i.e. the state price densities, allows us also to analyse investor behaviours during different financial time periods.

Hence, this works contributes not only to the analysis of the implied PDFs and their forecasts but also shows how this type of models can be used for monetary policy and financial stability purposes.

Acknowledgements

The drafting of this thesis has been a rather long process and I am thankful to many people who have contributed to this final outcome.

I initially started this thesis as a PhD. Student at the Institut Municipal d'Investigació Médica (IMIM) in Barcelona. From this initial period, I am deeply in debt to Jordi Alonso, who gave me a PhD. grant and encouraged me to start this thesis in his Research department. I also would like to express my gratitude to all the wonderful colleagues I had during my stay at IMIM: Gemma, Berta, Eli, Miquel, Carmen, Josué, Montse, Antònia, Olatz, Àngels, Sali, Sandra, Puri, Vladimir. I remember with true nostalgia the good times we spent together.

I am also really thankful to Holger Neuhaus who, in fact, introduced me to the risk-neutral option-implied probability density functions, most probably without him all this work would not have been possible. I also would like to express my gratitude to my three co-authors: Rupert, Olivier and Vesela. We have had many interesting discussions and we have enjoyed the time we spent while drafting our articles and preparing for international conferences. I have learnt a lot from the three of you and working with you has just been a true pleasure.

I am very thankful to Juan Ángel, who introduced me to two of my three co-authors and with whom I had many discussions on options data quality and how to further extend the implied-density functions to inflation options data.

To all my European Central Bank colleagues with whom I discussed this topic over the years: Jean-Marc, Patrick, Björn, Nicolàs, Enric, Pucho, Cristian, Enrique, Gemma and Andrej. Moreover, I cannot forget Ron, who in addition of being an extraordinary colleague during the last year also helped me solving some doubts on LaTeX.

I cannot forget to mention my thesis director, Josep Fortiana. Thank you for your patience, your guidance and the understanding of my family situation which has delayed the publication of this thesis. You have been truly a model director and an outstanding person. I very much appreciated the time we spent together in our first conference and the time you devoted to this thesis.

To all my friends from Frankfurt and Barcelona who encouraged me to finish this thesis and for all the unforgettable moments we spent together. To my family in Barcelona, París, Olot, Burgos, Madrid, Singapur and London for being such an important part of my life.

Finally, I am truly thankful to my parents, this work is the result of their support and education over the years. In particular, to my father who passed away last year who I am sure would have been very proud of this. To Eulàlia, who after more than twenty years being together has given me all the support I needed to finalise this thesis. Finally to my three kids: Pau, Laia and Teo. I just hope that I manage to educate them as well as my parents did with me and that the four of you keep on being the sunshine of my life.

Thank you! Gràcies! Gracias! Danke! Merci! Xiè xiè!

Pep Puigvert
April 2015

# Contents

# List of Figures

# List of Tables

# Abstract

The evolution of market interest rates is a key component of the transmission of monetary policy. Central Banks, market participants and monetary policy practitioners make use of the information contained in financial prices to better understand market interest rates developments. Such a comprehensive and quantitative assessment might also be derived from option-implied probability density functions (PDFs), and in particular when applied to Euribor options, which constitute a natural complement to the existing financial market indicators.

A number of methods for constructing these option-implied PDFs have already been developed in the literature. In general, although these methods might differ in the extremes of the tails of the distribution, there is no major difference in the central section of the estimated option-implied PDFs. And, arguably it is the central section of the option-implied PDFs which is more likely to be useful for monetary policy purposes, in contrast to financial stability analysis, where there may be greater focus on the tails of the distribution. In particular, such option-implied PDFs have not been studied in detail during periods of financial crisis, where arguably they may be the most useful.

In general, the methods that have been used to construct and estimate implied densities are ”risk-neutral”. Hence, they are indifferent regarding the investor behaviour and do not include a risk premium component. Some authors have already extended these methods to create ”real-world” option-implied PDFs which incorporate the investor behaviour and take into account the risk premium component. However,

there is very little research analysing and comparing the differences between these two densities in the Euribor market and, in particular, around episodes of crisis or monetary policy decisions.

By using a non-parametric technique, based on the Bliss and Panigirzoglou methodology, this thesis presents an analysis of PDFs for Euribor outturns in three months' time, using ”risk-neutral” and ”real-world” option-implied PDFs. This type of analysis allows us to reveal typical market reactions which could be potentially used by central banks as a complement to the already existing tools that allow them to take monetary policy decisions.

This thesis consists of the following 3 articles, which were published in international peer-reviewed journals:

- **A quantitative mirror on the Euribor market using implied probability density functions.** Puigvert-Gutiérrez J., de Vincent-Humphreys R. Eurasian Economic Review 2(1), 1-31, Spring 2012.
- **Interest rate expectations and uncertainty during ECB Governing Council days: Evidence from intraday implied densities of 3-month Euribor.** Vergote O., Puigvert-Gutiérrez J. Journal of Banking and Finance 36 (2012) 2804-2823.
- **Interest rate forecasts, state price densities and risk premium from Euribor options.** Ivanova V., Puigvert-Gutiérrez J. Journal of Banking and Finance 48 (2014) 210-223.

The first two articles above have been also published in the ECB Working Paper Series and were additionally peer-reviewed by two anonymous referees.

The first article deals with the estimation and subsequent analysis of the option implied PDFs for the three-month Euribor futures, from 13 January 1999, when options on Euribor futures first started trading. The article analyses for the first time and using daily data how three-month Euribor option-implied PDFs evolved during periods of prolonged stability as well as during periods of turbulence. Additionally, this article presents a detailed analysis of the daily trading volumes for the Euribor call and put options in terms of moneyness during different years of trading. Based on the volumes analysis, the article also explains how data need to be filtered out by considering only those options which are out-of-the-money.

The first article also sets out the estimation technique that is used to compute the option-implied PDFs, which is based on the Bliss Panigirtzoglou method. The article also explains how to construct constant maturity option-implied PDFs which are needed to better analyse and compare option-implied PDFs across different periods. Particular attention is given to how the constant maturity PDFs, and their associated summary statistics, reacted to the unfolding financial crisis between 2007 and 2009. In doing so, this article shows how the higher moments of the option-implied PDFs can provide timely and quantitative indicators not only of the amount of uncertainty around forward Euribor, the mean of the option-implied PDF, but the directional bias within that. The latter finding shows how option-implied PDFs can be effectively used as an uncertainty measure for monetary policy and financial stability analysis purposes.

The second article analyses changes in short-term interest rate expectations and uncertainty during ECB Governing Council days. For this purpose, the article extends the estimation of option-implied PDFs from daily frequency, which was introduced in the previous paper, up to tick frequency. This is the first time that option-implied PDFs are extended to this frequency since, so far, only the daily frequency had been explored for a wide set of instruments. In particular, the non-parametric estimator of these densities is applied to estimate intraday expectations of three-month Euribor of three month constant maturities. In addition, the paper assesses the impact of the ECB communication during 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 third article studies the development of interest rate risk premium and option implied state price densities in the Euribor futures option market. We investigate the period from the introduction of the Euro in 1999 until December 2012. The estimation of the risk-neutral PDFs from the Euribor options prices is derived again by using the same methodology as applied in the previous two articles. However, in this article we focus on one-month constant maturity density functions rather than on the three or six-month ones. Once we have a set of risk-neutral PDFs we test their ability to forecast Euribor futures rate. Using parametric and non-parametric statistical calibration, we also transform the risk-neutral option implied densities for the Euribor futures rate into real-world densities. The purpose of this transformation is twofold. First, to try to improve the forecast of the Euribor futures rate by using real-world densities; and, additionally, to compare the ratio of these two functions, i.e. the state price densities, by analysing the general investor preference on different state prices.

The option-implied PDFs have shown to be a robust indicator that helps measuring uncertainty during periods of crisis. In general, terms the conclusions and findings are that:

- Option-implied PDFs constitute a quantitative assessment to measure the monetary policy transmission and a natural complement to the wide range of financial indicators considered by central banks and monetary policy practitioners. Furthermore, they can provide an easily-accessible tool for visualizing how market reacts to specific events, and may thus contribute to both monetary and financial stability analysis.
- In particular, the relevance of the European Central Bank press release and conference as a communication tool is confirmed by the analysis of the moments of the option-implied PDFs at a higher frequency. The latter holds for both the introductory statement of the European Central Bank press conference but also the following question and answer session.
- The option-implied PDFs, which are by definition risk-neutral, cannot be used to forecast possible outcomes of the 3-month Euribor futures rates. However, the transformation of the risk-neutral PDFs into real-world PDFs allows us to forecast 3-month Euribor futures rates. Moreover, the analysis of the ratio between risk-neutral and real-world PDFs, i.e. the state price densities, suggest that investors price higher states with high and low rates compared to the expected spot rate. We found that, in general, state prices have a more pronounced right tail, implying that investors are more risk averse to increasing interest rates.

# Resum

L'evolució dels tipus d'interès de mercat és un dels components principals del mecanisme de transmissió de la política monetària. Els bancs centrals, els participants del mercat i els professionals de la política monetària recorren a la informació continguda en els preus financers per entendre millor l'evolució dels tipus d'interès de mercat. També és possible obtenir una avaluació completa i quantitativa d'aquestes característiques a través de les funcions de densitat de probabilitat (PDFs, per les seves sigles en anglès) implícita en opcions, en particular quan s'apliquen a opcions sobre l'Euribor, la qual cosa constitueix un complement natural dels indicadors del mercat financer existents.

La literatura recull diversos mètodes per a construir aquestes PDFs implícita basades en opcions. En general, si bé els mètodes poden presentar diferències als extrems de les cues de la distribució, no s'observen diferències significatives a la secció central de les PDFs implícites basades en opcions calculades. I, precisament, es pot afirmar que la secció central de les PDFs implícita basades en opcions és la que pot ser més útil a efectes de la política monetària, al contrari del que passa amb l'anàlisi de l'estabilitat financera, que s'acostuma a fixar més en les cues de la distribució. Concretament, aquestes PDFs implícita basades en opcions no s'han estudiat a fons durant períodes de crisi financera, que és precisament quan podrien resultar més útils.

En general, els mètodes que s'han emprat per construir i calcular densitats implícites són «neutrals al risc». Per tant, són indiferents al comportament dels inversors i no inclouen el component de la prima

de risc. Alguns autors ja han ampliat aquests mètodes, la qual cosa ha donat lloc a PDFs implícita basades en opcions «de condicions reals», que incorporen el comportament dels inversors i tenen en compte el component de la prima de risc. No obstant això, hi ha molts pocs estudis que analitzin i comparin les diferències entre aquestes dues densitats en el mercat de l'Euribor i, en particular, en relació amb episodis de crisi o decisions de política monetària.

En recórrer a una tècnica no paramètrica, basada en la metodologia de Bliss i Panigirzoglou, aquesta tesi presenta una anàlisi de PDFs per als resultats de l'Euribor a tres mesos, a partir de PDFs implícita basades en opcions «neutrals al risc» i «de condicions reals». Una anàlisi d'aquestes característiques permet posar de manifest reaccions típiques dels mercats, que els bancs centrals podrien emprar com a complement de les eines de les quals ja disposen per prendre decisions de política monetària.
En recórrer a una tècnica no paramètrica, basada en la metodologia de Bliss i Panigirzoglou, aquesta tesi presenta una anàlisi de PDFs per als resultats de l'Euribor a tres mesos, a partir de PDFs implícita basades en opcions «neutrals al risc» i «de condicions reals». Una anàlisi d'aquestes característiques permet posar de manifest reaccions típiques dels mercats, que els bancs centrals podrien emprar com a complement de les eines de les quals ja disposen per prendre decisions de política monetària.

Aquesta tesi consta dels tres articles següents, publicats en revistes internacionals arbitrades:

- **A quantitative mirror on the Euribor market using implied probability density functions.** Puigvert-Gutiérrez J., de Vincent-Humphreys R. Eurasian Economic Review 2(1), 1-31.

- **Interest rate expectations and uncertainty during ECB Governing Council days: Evidence from intraday implied densities of 3-month Euribor.** Vergote O., Puigvert-Gutiérrez J. Journal of Banking and Finance 36 (2012) 2804-2823.
- **Interest rate forecasts, state price densities and risk premium from Euribor options.** Ivanova V., Puigvert-Gutiérrez J. Journal of Banking and Finance 48 (2014) 210-223.

Els dos primers s'han publicat també a la ECB Working Paper Series i van ser revisats, a més, per dos avaluadors anònims.

El primer article aborda el càlcul i la subsegüent anàlisi de les PDFs implícita en opcions per als futurs de l'Euribor a tres mesos, a partir del 13 de gener de 1999, data en la qual es va iniciar la cotització de les opcions sobre futurs de l'Euribor. S'hi analitza per primera vegada, i a partir de dades diàries, l'evolució de les PDFs implícita basades en opcions de l'Euribor a tres mesos al llarg de períodes tant d'estabilitat prolongada com de turbulències. A més, presenta una anàlisi detallada dels volums d'operacions diàries de les opcions de compra i venda de l'Euribor, en termes de monetització, en diferents anys de negociació. A partir de l'anàlisi de volums, l'article explica també de quina manera cal filtrar les dades, tenint en compte només les opcions fora de diner («out of the money»).

Així mateix, el primer article descriu la tècnica d'estimació emprada per calcular les PDFs implícita en opcions, que es basa en el mètode Bliss-Panigirtzoglou, i explica com construir PDFs implícita basades en opcions de venciment constant, necessàries per millorar l'anàlisi i la comparació de PDFs implícita basades en opcions que pertanyen a períodes diferents. Es para particular atenció a la manera com van reaccionar les PDFs a venciment constant, així com les estadístiques resumides associades, davant el desenvolupament de la crisi financera entre 2007 i 2009. Així, l'article mostra de quina manera els moments de les PDFs implícita en opcions poden ser indicadors oportuns i quantitatius no només del nivell d'incertesa de l'Euribor a llarg termini i la mitjana de les PDFs implícita basades en opcions, sinó també del biaix direccional entre ells. Aquesta darrera observació revela que les PDFs implícita en opcions poden emprar-se de manera eficaç com a mesura de la incertesa a l'efecte de l'anàlisi de la política monetària i l'estabilitat financera.

El segon article examina els canvis en les expectatives de tipus d'interès a curt termini i la incertesa els dies de reunió del Consell de Govern del BCE. Amb aquesta finalitat, l'article amplia el càlcul de PDFs implícita en opcions amb freqüència diària, introduït en l'article anterior, fins a la freqüència tick. Per primera vegada, les PDFs implícita en opcions s'amplien a aquesta freqüència, ja que fins ara només s'havia aplicat la freqüència diària a un ampli conjunt d'instruments. Concretament, s'aplica l'estimador no paramètric d'aquestes densitats al càlcul de les expectatives intradia de l'Euribor a tres mesos amb venciments constants de tres mesos. A més, l'article avalua els efectes de les comunicacions del BCE els dies de reunió del Consell de Govern. En primer lloc, l'article aborda una sèrie de consideracions pràctiques i estadístiques que es plantegen en aplicar l'extracció de la densitat implícita a la freqüència elevada. En segon lloc, a partir d'estudis de casos i l'anàlisi de models intradia, l'article mesura també el contingut informatiu de les densitats i el nivell d'incertesa obtinguts. També, duu a terme una anàlisi de regressió per identificar els factors desencadenants de les reaccions observades en el mercat, expressades en forma de variacions de la densitat.

El tercer article estudia l'evolució de la prima de risc de tipus d'interès i les densitats dels preus d'estat implícites en opcions en el mercat d'opcions sobre futurs de l'Euribor. El període analitzat comprèn des de la introducció de l'euro el 1999 fins a desembre de 2012. L'estimació de les densitats neutrals al risc a partir dels preus de les opcions sobre l'Euribor es calcula un cop més a partir de la metodologia emprada en els dos articles anteriors. Tanmateix, en aquesta ocasió ens centrem en les funcions de densitat amb un venciment constant d'un mes, i no de tres o sis mesos. Un cop creat el conjunt de PDFs neutrals al risc, en provem la capacitat per predir els preus dels futurs de l'Euribor. Mitjançant el calibratge estadístic paramètric i no paramètric, transformem a més les densitats implícites en opcions neutrals al risc dels preus dels futurs de l'Euribor en densitats de condicions reals. Aquesta transformació respon a dos objectius: El primer és intentar millorar la previsió dels preus dels futurs de l'Euribor a partir de densitats de condicions reals i, a més, comparar els coeficients d'aquestes dues funcions; és a dir, les densitats dels preus d'estat, mitjançant l'anàlisi de les preferències dels inversors en general davant diversos preus d'estat.

Les PDFs implícita en opcions han demostrat ser un indicador sòlid, capaç de contribuir a mesurar la incertesa en períodes de crisi. En termes generals, les conclusions i observacions van ser les següents:

- Les PDFs implícita en opcions representen una avaluació quantitativa per al mesurament de la transmissió de la política monetària, així com un complement natural de l'ampli ventall d'indicadors financers emprats pels bancs centrals i els professionals de la política monetària. A més, poden emprar-se com a eina de fàcil accés per visualitzar la reacció del mercat a un fet concret i, per tant, es poden fer servir per a l'anàlisi de l'estabilitat financera i monetària.
- En particular, la rellevància de la nota i la conferència de premsa del Banc Central Europeu com a eina de comunicació queda corroborada per l'anàlisi dels moments de les PDFs implícita en

opcions a una freqüència superior. Aquesta afirmació és aplicable tant al comunicat preliminar de la conferència de premsa del Banc Central Europeu com a la sessió de preguntes i respostes posterior.

- Les PDFs implícita en opcions, neutrals al risc per definició, no es poden fer servir per predir els resultats possibles dels preus dels futurs de l'Euribor a tres mesos. Això no obstant, la transformació de les PDFs neutrals al risc en PDFs de condicions reals ens permet preveure els preus dels futurs de l'Euribor a tres mesos. A més, l'anàlisi de la relació entre les PDFs neutrals al risc i les de condicions reals, és a dir, les densitats de preu d'estat, indica que els inversors assignen preus més elevats a estats amb tipus alts i baixos en comparació amb el tipus al comptat esperat. Observem que, en general, els preus d'estat presenten una cua a la dreta més pronunciada, la qual cosa implica que els inversors tenen una aversió al risc més gran davant un increment dels tipus d'interès.

# Resumen

La evolución de los tipos de interés de mercado es uno de los principales componentes del mecanismo de transmisión de la política monetaria. Los bancos centrales, los participantes en el mercado y los profesionales de la política monetaria recurren a la información contenida en los precios financieros para comprender mejor la evolución de los tipos de interés de mercado. También es posible obtener una evaluación completa y cuantitativa de estas características a través de las PDFs implícita en opciones (PDFs, por sus siglas en inglés), en particular cuando se aplican a opciones sobre el Euribor, lo que constituye un complemento natural de los indicadores del mercado financiero existentes.

La literatura recoge varios métodos para la construcción de estas PDFs implícita en opciones. En general, si bien los métodos pueden presentar diferencias en los extremos de las colas de la distribución, no se observan diferencias significativas en la sección central de las PDFs implícita en opciones calculadas. Y, precisamente, puede afirmarse que la sección central de las PDFs implícita en opciones es la que puede resultar más útil a efectos de la política monetaria, al contrario de lo que ocurre en el análisis de la estabilidad financiera, que suele prestar mayor atención a las colas de la distribución. Concretamente, estas PDFs implícita basadas en opciones no han sido estudiadas en profundidad durante períodos de crisis financiera, que es precisamente cuando podrían resultar de mayor utilidad.

En general, los métodos a los que se ha recurrido para construir y calcular densidades implícitas son ≪neutrales al riesgo≫. Por tanto, son indiferentes al comportamiento de los inversores y no incluyen un componente de prima de riesgo. Algunos autores han ampliado ya dichos métodos, dando lugar a PDFs implícita en opciones ≪de condiciones reales≫, que incorporan el comportamiento de los inversores y tienen en cuenta el componente de la prima de riesgo. No obstante, existen muy pocos estudios donde se analicen y comparen las diferencias entre estas dos densidades en el mercado del Euribor y, en particular, en relación con episodios de crisis o decisiones de política monetaria.

Al recurrir a una técnica no paramétrica, basada en la metodología de Bliss y Panigirzoglou, esta tesis presenta un análisis de PDFs para los resultados del Euribor a tres meses, a partir de PDFs implícita basadas en opciones ≪neutrales al riesgo≫ y ≪ de condiciones reales≫. Un análisis tal nos permite poner de manifiesto reacciones típicas de los mercados, que los bancos centrales podrían utilizar como complemento de las herramientas de las que ya disponen para la toma de decisiones de política monetaria.

Esta tesis consta de los tres artículos siguientes, publicados en revistas internacionales arbitradas:

- **A quantitative mirror on the Euribor market using implied probability density functions.** Puigvert-Gutiérrez J., de Vincent-Humphreys R. Eurasian Economic Review 2(1), 1-31, Spring 2012.
- **Interest rate expectations and uncertainty during ECB Governing Council days: Evidence from intraday implied densities of 3-month Euribor.** Vergote O., Puigvert-Gutiérrez J. Journal of Banking and Finance 36 (2012) 2804-2823.

- **Interest rate forecasts, state price densities and risk premium from Euribor options.** Ivanova V., Puigvert-Gutiérrez J. Journal of Banking and Finance 48 (2014) 210-223.

Los dos primeros han sido publicados también en la ECB Working Paper Series y fueron revisados además por dos evaluadores anónimos.

El primer artículo aborda el cálculo y el subsiguiente análisis de las PDFs implícita en opciones para los futuros del Euribor a tres meses, a partir del 13 de enero de 1999, fecha en la que se inició la cotización de las opciones sobre futuros del Euribor. En él se analizan por primera vez y a partir de datos diarios la evolución de las PDFs implícita en opciones del Euribor a tres meses a lo largo de períodos tanto de estabilidad prolongada como de turbulencias. Además, presenta un análisis detallado de los volúmenes de operaciones diarias de las opciones de compra y venta del Euribor, en términos de monetización, en distintos años de negociación. A partir del análisis de volúmenes, el artículo explica también de qué modo es necesario filtrar los datos, teniendo en cuenta solo las opciones fuera de dinero (≪out of the money≫).

Asimismo, el primer artículo describe la técnica de estimación utilizada para calcular las PDFs implícita en opciones, que se basa en el método Bliss-Panigirtzoglou, y explica cómo construir PDFs implícita de vencimiento constante basadas en opciones, necesarias para mejorar el análisis y la comparación de PDFs implícita basadas en opciones pertenecientes a períodos distintos. Se presta particular atención a cómo reaccionaron las PDFs a vencimiento constante, así como las estadísticas resumidas asociadas, ante el desarrollo de la crisis financiera entre 2007 y 2009. Así, el artículo muestra de qué modo los momentos de las PDFs implícita basadas en opciones pueden ser indicadores oportunos y cuantitativos no solo del nivel de incertidumbre del Euribor a largo plazo y la media de la función de densidad de probabilidad implícita basada en opciones, sino también del sesgo direccional en ellos. Esta última observación revela que las PDFs implícita en opciones pueden utilizarse de manera eficaz como medida de la incertidumbre a efectos del análisis de la política monetaria y la estabilidad financiera.

El segundo artículo examina los cambios en las expectativas de tipos de interés a corto plazo y la incertidumbre en los días de reunión del Consejo de Gobierno del BCE. Para ello, el artículo amplía el cálculo de PDFs implícita en opciones con frecuencia diaria, introducido en el artículo anterior, hasta la frecuencia tick. Por primera vez, las PDFs implícita en opciones se amplían a esta frecuencia, ya que hasta la fecha solo se había aplicado la frecuencia diaria a un amplio conjunto de instrumentos. Concretamente, se aplica el estimador no paramétrico de estas densidades al cálculo de las expectativas intradía del Euribor a tres meses con vencimientos constantes de tres meses. Además, el artículo evalúa los efectos de las comunicaciones del BCE en los días de reunión del Consejo de Gobierno. En primer lugar, el artículo aborda una serie de consideraciones prácticas y estadísticas que se plantean al aplicar la extracción de la densidad implícita a la frecuencia elevada. En segundo lugar, a partir de estudios de casos y el análisis de modelos intradía, el artículo mide también el contenido informativo de las densidades y el nivel de incertidumbre obtenidos. Asimismo, lleva a cabo un análisis de regresión para identificar los factores desencadenantes de las reacciones observadas en el mercado, expresadas en forma de variaciones de la densidad.

El tercer artículo estudia la evolución de la prima de riesgo de tipo de interés y las densidades de los precios de estado implícitas en opciones en el mercado de opciones sobre futuros del Euribor. El período analizado abarca desde la introducción del euro en 1999 hasta diciembre de 2012. La estimación de las densidades neutrales al riesgo a partir de los precios de las opciones sobre el Euribor se calcula una vez más a partir de la metodología utilizada en los dos artículos anteriores. Sin embargo, en esta ocasión nos centramos en las funciones de densidad con un vencimiento constante de un mes, y no de tres o seis meses. Una vez creado el conjunto de PDFs neutrales al riesgo, probamos su capacidad para predecir los precios de los futuros del Euribor. Mediante la calibración estadística paramétrica y no paramétrica, transformamos además las densidades implícitas en opciones neutrales al riesgo de los precios de los futuros del Euribor en densidades de condiciones reales. Esta transformación responde a dos objetivos: El primero de ellos es intentar mejorar la previsión de los precios de los futuros del Euribor a partir de densidades de condiciones reales y, además, comparar los coeficientes de estas dos funciones; es decir, las densidades de los precios de estado, mediante el análisis de las preferencias de los inversores en general ante distintos precios de estado.

Las PDFs implícita en opciones han demostrado ser un indicador sólido, capaz de contribuir a medir la incertidumbre en períodos de crisis. En términos generales, las conclusiones y observaciones fueron las siguientes:

- Las PDFs implícita en opciones constituyen una evaluación cuantitativa para la medición de la transmisión de la política monetaria, así como un complemento natural del amplio abanico de indicadores financieros utilizados por los bancos centrales y los profesionales de la política monetaria. Asimismo, pueden utilizarse como herramienta de fácil acceso para visualizar la reacción del mercado a un hecho concreto y, por tanto, pueden utilizarse para el análisis de la estabilidad financiera y monetaria.

- En particular, la relevancia de la nota y la conferencia de prensa del Banco Central Europeo como herramienta de comunicación se ve corroborada por el análisis de los momentos de las PDFs implícita en opciones a una frecuencia superior. Esta afirmación es aplicable tanto al comunicado preliminar de la conferencia de prensa del Banco Central Europeo como a la sesión de preguntas y respuestas posterior.

- Las funciones de densidad de probabilidad implícita en opciones, neutrales al riesgo por definición, no pueden emplearse para predecir los resultados posibles de los precios de los futuros del Euribor a tres meses. No obstante, la transformación de las funciones de densidad de probabilidad neutrales al riesgo en funciones de densidad de probabilidad de condiciones reales nos permite prever los precios de los futuros del Euribor a tres meses. Además, el análisis de la relación entre las funciones de densidad de probabilidad neutrales al riesgo y las de condiciones reales, es decir, las densidades de precio de estado, indica que los inversores asignan precios más elevados a estados con tipos altos y bajos en comparación con el tipo al contado esperado. Observamos que, en general, los precios de estado presentan una cola a la derecha más pronunciada, lo cual implica que los inversores presentan una mayor aversión al riesgo ante un incremento de los tipos de interés.

# Part I

# Introduction

*Finance is wholly different from the rest the economy.*

Alan Greenspan

CHAPTER

1

# Preliminary theory

This preliminary Chapter aims at introducing the underlying option pricing and financial theory used in this thesis to construct the option-implied probability density functions (PDFs). In this respect, this Chapter includes a set of basic financial definitions which are the starting point for the construction of the risk-neutral option-implied PDFs.

This Chapter also presents the Breeden and Litzenberg (1978) theorem, which is used in a vast amount of methods to derive the option-implied PDFs. The Breeden and Litzenberger (1978) theorem makes use of two previous financial results: the Arrow-Debreu (1954) Securities and the Cox-Ross (1976) theorem, which are also presented in this Chapter.

## 1.1 Basic definitions

**Definition 1.** *A future is a financial contract that gives the obligation to buy or sell an underlying asset $S$ at a pre-agreed price at a certain time $T$ in the future.*

**Definition 2.** *A European call option with strike price $K$ and maturity $T$ is a financial instrument that gives the holder the right, but not the obligation, to buy a* *particular asset S at price K at time T. The payoff of a European call option at time T is then given by*

$$P_T(S) = max(S_T - K, 0) \tag{1.1}$$

*where S represents the price of the underlying asset.*

*Similarly, a European put option with strike price K and maturity T is a financial instrument that gives the holder the right, but not the obligation, to sell a particular asset S at price K time T. The payoff of a European put option at time T is given by*

$$P_T(S) = max(K - S_T, 0) \tag{1.2}$$

**Definition 3.** *The strike price of an option is the price at which the underlying asset S of an option can be exercised, i.e. either bought or sold.*

**Definition 4.** *A risk-free rate interest rate is the theoretical rate of return of a particular investment in which there is no risk of financial loss. An example of risk-free interest rate assets or proxies are US government bonds, AAA government or company rated bonds, since there is in principle no perceived risk of default associated with this type of investments.*

**Theorem 1.** ***Call-put parity relationship.*** *Let's suppose that $C_t$ is the value of a European call option of and underlying asset S with strike price K and maturity T. Let's suppose that $P_t$ is the value of a European put option on the same asset S with the same strike price and expiration. Finally, let's suppose that S has a final price at expiration of $S_T$, and let B(t, T) represent the value of a risk-free zero-coupon bond at time t with final value 1 at expiration time T. If these assumptions hold and there is no arbitrage, then*

$$C_t + K * B(t, T) = P_t + S_t \tag{1.3}$$

*The equality, which is known as the call-put parity relationship shows that the value of European call can be derived from the value of European put both with the certain price exercise and exercise date.*

**Definition 5.** *The* moneyness *of a given option is the difference between the strike price K of the option minus the current market price of the underlying asset S. A call option is:*

- in-the-money *if the strike price K is below the market price of the underlying asset S,*
- at-the-money *if the strike price K is equal the market price of the underlying asset S, and*
- out-of-the-money *if the strike price K is above the market price of the underlying asset S.*

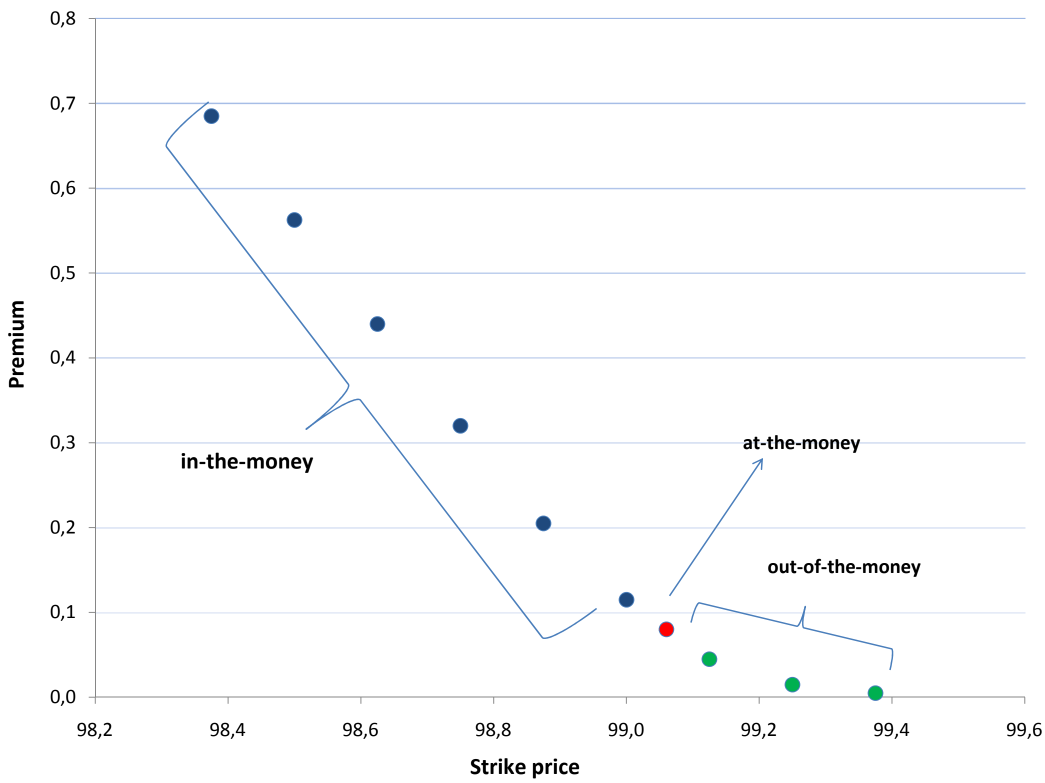


**Figure 1.1:** Moneyness of a call option.

**Definition 6.** *A* risk-neutral probability density measure *is a probability measure such that the current price of a given security is exactly equal to the present value* *of the discounted expected value of its future payoffs. A* risk-neutral probability density measure *exists* if and only if *the market is* arbitrage-free.

## 1.2 Arrow-Debreu Securities

**Definition 7.** *An Arrow-Debreu (1954) Security, also called an elementary claim, is a derivative security that pays 1 unit at a future time T if the value of the underlying asset, or portfolio assets, takes a particular state* $S_T = K$ *at this time, and 0 otherwise.*

The Arrow-Debreu (1954) security is the simplest example of a *risk-neutral probability density measure*. However, Arrow-Debreu (1954) securities are not traded on any market exchange, so their price is not directly observable. In this particular context they are used to derive the Breeden and Litzenberger (1978) results and to better understand the economic equilibrium in an uncertain environment.

In this respect, in the proof of 1.4, an Arrow-Debreu (1954) security is replicated by combining call options that have the same time to maturity but different strike prices into a strategy called *butterfly spread*. The price of this *butterfly spread* will reflect the probabilities that investors attribute to those particular states in the future.

## 1.3 Cox-Ross Theorem

Cox and Ross (1976) showed that there is a relationship between the price of an option and the expected value of its futures values discounted with risk free rate. This result is also used to proof the Breeden and Litzenberg (1978) theorem.

**Theorem 2.** *The price of a call option,* $C_t$*, at time* $t$ *on a given asset with price* $F_t$ *is the expectation under the option-implied PDF,* $f(F_T)$*, of its future option values*

$$C_t(F_t, K, \tau) = e^{-r\times\tau} \times E_t^{\mathbb{Q}}[(F_T - K)^+] = e^{-r\times\tau} \times \int_K^{\infty} f(F_T)(F_T - K) dF_T, \quad (1.4)$$

*Similarly a put option, $P_t$, can be expressed as*

$$P_t(F_t, K, \tau) = e^{-r\times\tau} \times E_t^{\mathbb{Q}}[(K - F_T)^+] = e^{-r\times\tau} \times \int_0^K f(F_T)(K - F_T)dF_T, \quad (1.5)$$

*where $K$ is the strike price, $T$ is the option expiry date, r is the risk-free rate, and $\tau$ is the remaining time to maturity.*

## 1.4 Breeden and Litzenberger Theorem

A direct relationship between call option prices and option-implied PDFs was firstly introduced by Breeden and Litzenberger (1978). In this regard, the authors proved that under certain conditions the second derivative of a call option price with respect to the strike price is directly proportional to the option-implied PDF.

**Theorem 3.** *Let be $C_t(K)$ a call function, monotonic decreasing and convex, and twice differentiable with respect to K. Let's also assume that markets are perfect, i.e. there are no restrictions on short-sales, no transactions costs or taxes, and investors borrow considering the risk-free interest rate. Then a relationship between the call price function and the risk-neutral option implied PDF can by expressed by equation (1.6)*

$$\frac{\partial^2 C_t(F_t, K, \tau)}{\partial^2 K} = e^{-r\times\tau} \times f(F_T) \quad (1.6)$$

*where $K$ is the strike price, $T$ is the option expiry date, $F_T$ is the underlying future at time $T$, $r$ is the risk-free rate, $\tau$ is the remaining time to maturity and $f$ is the option-implied PDF.*

*Proof.* The Breeden and Litzenberger (1978) equation can be proved directly by using the expression of a European call option or Arrow-Debreu (1954) securities.

Let's prove first the Breeden and Litzenberger (1978) equation by considering the expression of a European call option which is given by the Cox and Ross (1976) equation. A relationship between the call option and the option-implied PDF can be established by:

$$C_t(F_t, K, \tau) = e^{-r\times\tau} \times E_t^{\mathbb{Q}}[(F_T - K)^+] = e^{-r\times\tau} \times \int_K^{\infty} f(F_T)(F_T - K) dF_T, \quad (1.7)$$

By differentiating now the expression 1.7 in respect to $K$ the expression in equation 1.8 is obtained.

$$\frac{\partial C_t(F_t, K, \tau)}{\partial K} = -e^{-r\times\tau} \times \int_K^{\infty} f(F_T) dF_T, \quad (1.8)$$

Finally, the expression 1.6 in the statement of the theorem is obtained by differentiating now equation 1.8 with respect to $K$.

The Breeden and Litzenberger (1978) equation can also be proven by using Arrow-Debreu (1954) securities.

In order to do so, an Arrow-Debreu (1954) security needs to be constructed first by using a *butterfly spread*, which consists of a combination of two European call options $C_t(F_t, K, \tau)$ with strike price $K = F_T$ and two additional European call options, one with strike price $K = F_T + \Delta F_T$ and the other one with strike price $K = F_T - \Delta F_T$, where $\Delta F_T$ is the step between two adjacent calls.

The butterfly spread is given in this case by:

$$\frac{[C_t(F_t, F_T + \Delta F_T, \tau) - C_t(F_t, F_T, \tau)] - [C_t(F_t, F_T, \tau) - C_t(F_t, F_T - \Delta F_T, \tau)]}{\Delta F_T} \quad (1.9)$$

Letting $\Delta F_T$, the step between two adjacent calls, tend to 0, the butterfly spread tends to a Dirac delta measure centered at $K = F_T$.

$$\left.\frac{[C_t(F_t, F_T + \Delta F_T, \tau) - C_t(F_t, F_T, \tau)] - [C_t(F_t, F_T, \tau) - C_t(F_t, F_T - \Delta F_T, \tau)]}{\Delta F_T}\right|_{K=F_T} = 1 \quad (1.10)$$

The equation 1.10 represents an Arrow-Debreu (1954) security, paying 1 unit if $F_T = K$ and 0 otherwise. This result would allow us in the next step to proof the Breeden-Litzenberger (1978) equation.

The result above can be better understood with an example as suggested by Breeden and Litzenberger (1978) or Bahra (1997). A portfolio of a butterfly centered in 1 and with a unit step between adjacent options would pay unit 1 only if the state $F_T = 1$ and can be constructed by

$$[C_t(F_t, 2, \tau) - C_t(F_t, 1, \tau)] - [C_t(F_t, 1, \tau) - C_t(F_t, 0, \tau)]. \tag{1.11}$$

In particular, the different payoffs of this portfolio can be obtained for different values of the exercise price.

| $K$ | $C_t(F_t, 0, \tau)$ | $C_t(F_t, 1, \tau)$ | $C_t(F_t, 2, \tau)$ |
|---|---|---|---|
| 1 | 1 | 0 | 0 |
| 2 | 2 | 1 | 0 |
| 3 | 3 | 2 | 1 |
| .... | | | |
| n | n | n-1 | n-2 |

**Table 1.1:** Butterfly spread centered in 1.

Table 1.1 can be further generalised for the case described in 1.11. If we denote $P(F_t, F_T, \tau)$, the current price of the Arrow-Debreu (1954) security centered in $F_T$, we can write the price of this butterfly spread divided by the step size between two consecutive calls as follows:

$$\frac{P(F_t, F_T, \tau)}{\Delta F_T} = \frac{[C_t(F_t, F_T + \Delta F_T, \tau) - C_t(F_t, F_T, \tau)] - [C_t(F_t, F_T, \tau) - C_t(F_t, F_T - \Delta F_T, \tau)]}{(\Delta F_T)^2} \tag{1.12}$$

| $K$ | $C_t(F_t, F_T - \Delta F_T, \tau)$ | $C_t(F_t, F_T, \tau)$ | $C_t(F_t, F_T + \Delta F_T, \tau)$ | Payoff of the Arrow-Debreu security |
|---|---|---|---|---|
| $F_T - \Delta F_T$ | 0 | 0 | 0 | 0 |
| $F_T$ | $\Delta F_T$ | 0 | 0 | $\Delta F_T$ |
| $F_T + \Delta F_T$ | $2\Delta F_T$ | $\Delta F_T$ | 0 | 0 |
| $F_T + 2\Delta F_T$ | $3\Delta F_T$ | $2\Delta F_T$ | $\Delta F_T$ | 0 |
| ... | .... | .... | .... | .... |
| $F_T + n\Delta F_T$ | $(n+1)\Delta F_T$ | $n\Delta F_T$ | $(n-1)\Delta F_T$ | 0 |

**Table 1.2:** Butterfly spread centered in $F_T$.

If $\Delta F_T$ tends to 0, the right side of expression 1.12 tends to the second derivative of the call price with respect to the exercise price but evaluated at the state $K = F_T$.

If $\Delta F_T$ tends to 0, we can re-write 1.12 as

$$\lim_{\Delta F_T \to 0} \frac{P(F_t, F_T, \tau)}{\Delta F_T} = \frac{\partial^2 C_t(F_t, K, \tau)}{\partial^2 K} \tag{1.13}$$

The price of a butterfly spread at $F_T = K$ can be expressed as the discounted expected future payoff of an Arrow-Debreu (1954) security, i.e. 1 multiplied by the risk-neutral implied-option probability of the state $F_T = K$. By applying the price in equation 1.13 across the continuum of possible values of $F_T = K$, this gives the result that the exercise price is equal to the risk-neutral implied-option probability of $F_T$ conditioned on the underlying price at time t. $\square$

The conditions defined in the the Breeden and Litzenberger (1978) theorem imply that there are no arbitrage opportunities. Hence, the observed call prices need to be convex and monotonically decreasing. This implies, in practice, that in order to fulfill these conditions the original call and put price data need to be filtered as shown in Chapter 3.

## 1.5 Black and Scholes

In this Section, the Black and Scholes (1973) model is presented. This model is important since it gives an estimate to price European call options. Additionally, in this thesis, from this model, the sensitivity of a financial instrument with respect to several parameters such as the spot price, volatility or interest rate is derived. These parameters, which are known as the *Greeks*, are used to derive the Bliss and Panigirtzoglou (2004) methodology.

The Black and Scholes (1973) model makes an assumption on the evolution of the price of the underlying asset $F_t$ of the option. In this respect, it assumes that this price evolves according to the geometric Brownian motion and can be described by the stochastic differential Black and Scholes (1973) equation:

$$dF = \mu F dt + \sigma F dW \tag{1.14}$$

where $\mu F$ is the instantaneous expected drift rate ($\mu$ is a constant of the model), $\sigma^2 F^2$ is the instantaneous variance rate ($\sigma$ is a constant of the model), $W$ is a Wiener process, and $dW$ is an increment of the Wiener process.

The equation 1.14 can be also re-written in the following way:

$$F_t = F_0 + \int_0^t \mu F_u du + \int_0^t \sigma F_u dW_u \tag{1.15}$$

The solution of the Black-Scholes (1973) equation presented by the authors is motivated by the construction of a portfolio containing the option and the underlying asset $F_t$. In this case the authors assume that the returns of this portfolio are equal to the risk-free rate interest rate.

The authors showed that the solution of equation 1.14 can be given by:

$$ln(F_t) = ln(F_0) + (\mu - \frac{\sigma^2}{2})t + \sigma W_t \tag{1.16}$$

which can be also expressed as:

$$F_t = F_0 \exp\{(\mu - \frac{\sigma^2}{2})t + \sigma W_t\} \tag{1.17}$$

In fact, given the case that a Black-Scholes (1973) markets is complete, a European call can be also priced by using equation 1.17. In particular, the price of a call option on a future contract is given by:

$$C_t(F_t, K, \tau) = e^{-r\times\tau}[F_0 N(d_1) - KN(d_1)], \tag{1.18}$$

where $F_0$ is the spot price of the futures contract, $K$ is the strike price, $T$ is the option expiry date, r is the risk-free rate, $\tau$ is the remaining time to maturity, and:

$$d_1 = \frac{ln(\frac{F_0}{K}) + \frac{\sigma^2}{2}\tau}{\sigma\sqrt{\tau}} \tag{1.19}$$

and

$$d_2 = d_1 - \sigma\sqrt{\tau} \tag{1.20}$$

Similarly, the corresponding price of a European put option can be obtained from the *put-call parity*:

$$P_t(F_t, K, \tau) = P_t(F_t, K, \tau) - F_t + e^{-r\times\tau}K. \tag{1.21}$$

## 1.6 The Greeks

The Greeks represent the sensitivity of the price of the European options defined in 1.18 and 1.21 to a change in underlying parameters such as the spot price, volatility or interest rate. The Greeks are used in the Bliss and Panigirzoglou (2004) method to transform the original *premium* and *strike price* into the *delta* and *sigma* space defined by the Greeks. Despite that further Greeks are calculated, we will only present the two which are used in the Bliss and Panigirzoglou (2004) method.

Given a European call option $C_t(F_t, K, \tau)$, the *delta* of an option measures the rate of change of the theoretical option value with respect to changes in the underlying asset's price $F_t$ and can be derived by calculating the first derivative of $C_t(F_t, K, \tau)$ with respect to the underlying asset's price $F_t$:

$$\delta_t(F_t, K, \tau) = \frac{\partial C_t(F_t, K, \tau)}{\partial F_t} = N(d_1). \tag{1.22}$$

Given a European call option $C_t(F_t, K, \tau)$, the *vega* of an option measures the sensitivity to the volatility of the model and can be derived by calculating the first derivative of $C_t(F_t, K, \tau)$ with respect to the $\sigma$ of the underlying asset:

$$\nu_t(F_t, K, \tau) = \frac{\partial C_t(F_t, K, \tau)}{\partial \sigma} = \frac{N(d_1)}{F_t \sigma \sqrt{\tau}}. \tag{1.23}$$

*Derivatives are financial weapons of mass destruction.*

Warren Buffet

# CHAPTER 2 Context

Option-implied probability density functions (PDFs) are part of the set of econometric tools used by national central banks, economists and financial researchers to evaluate market expectations and uncertainty on the basis of the prices for option contracts on different market instruments.

Option-implied PDFs are also often referred as risk-neutral densities or risk-neutral PDFs. They are PDFs based on options for a particular asset, i.e. 3-month Euribor, 10-year government bond, assuming that investors are risk-neutral.

Option-implied PDFs summarise the total set of likely outcomes and probabilities in the near future for a particular asset attached by the market around specific economic and financial events, such as Governing Council decisions, financial crisis, etc...

Option-implied PDFs go a step further compared to the analysis of forward rates, because they can capture not only market expectations, but uncertainty in the near future. This uncertainty can be estimated by analysing for instance the percentiles or the statistical moments of the distribution over time.

## 2.1 Literature review and methodology classification of option-implied probability density functions

The option-implied PDFs first appeared on the last quarter of the twentieth century, but they started to become very popular late in the nineties, when several financial and economic authors tried to explore how to analyse market expectations and uncertainty from option prices.

The common denominator of the current existing methodologies is mainly derived from the results provided by Cox and Ross (1976) and Breeden and Litzenberger (1976). In this respect, Cox and Ross (1976), showed that assuming that investors are risk-neutral, the price of a call option, $C_t$, at time $t$ on a given asset with price $F_t$ is the expectation under the option-implied PDF, $f(F_T)$, of its future option values

$$C_t(F_t, K, \tau) = e^{-r\times\tau} \times E_t^{\mathbb{Q}}[(F_T - K)^+] = e^{-r\times\tau} \times \int_K^{\infty} f(F_T)(F_T - K)dF_T, \quad (2.1)$$

Similarly a put option, $P_t$, can be expressed as

$$P_t(F_t, K, \tau) = e^{-r\times\tau} \times E_t^{\mathbb{Q}}[(K - F_T)^+] = e^{-r\times\tau} \times \int_0^{K} f(F_T)(K - F_T)dF_T, \quad (2.2)$$

where $K$ is the strike price, $T$ is the option expiry date, r is the risk-free rate, and $\tau$ is the remaining time to maturity. However, in practice, the Cox and Ross (1976) result does not allow to directly calculate the option-implied PDF.

Based on the Cox and Ross (1976) result, Breeden and Litzenberger (1978) went a step further and showed that the relationship of the PDF can be directly obtained by differentiating twice the call option in respect to the strike price. By differentiating twice the call option in equation (2.1) Breeden and Litzenberg (1978) formally obtained

$$\frac{\partial^2 C_t(F_t, K, \tau)}{\partial^2 K} = e^{-r\times\tau} \times f(F_T) \quad (2.3)$$

The results provided by Cox and Ross (1976) and Breeden and Litzenberger (1978) and the underlying literature were further detailed in Chapter 1.

The Breeden and Litzenberger (1978) result is in fact the basis of the existing techniques for constructing the risk-neutral option-implied PDFs. Bahra (1997), Bliss and Panigirtzoglou (2002) and more recently Jondeau (2006) *et al.* have reviewed and classified the different option-implied PDFs methodologies.[1] According to Bliss and Panigirtzoglou (2002), the methodologies can be largely classified into five groups: stochastic process methods, implied binomial trees, finite-difference methods, option-implied PDF approximating function methods, and implied volatility smoothing methods. Bahra (1997) classifies the option-implied PDFs methodologies in four groups, similar to the ones defined by Bliss and Panigirtzoglou (2002). Jondeau *et al.* (2006) classify the methodologies in two categories: structural and non-structural. The authors define the structural methodologies as those that take into account a specific stock price dynamics structure or, in some cases, the volatility of the process, as opposite to the non-structural methodologies, which do not take into account the dynamics of the price.

In particular, the first group as described by Bliss and Panigirtzoglou (2002), i.e. *stochastic process methods*, assumes that the price of the underlying asset follows a stochastic process and make use of market option prices to estimate the parameters of the stochastic process. The parameters of the stochastic process are used to obtain an option-implied PDF. This approach has been used for instance by Bates (2000) and Malz (1997). Bates (2000) uses this methodology to derive a model for pricing American options on jump-diffusion processes with systematic jump risk and uses the jump-diffusion parameters implicit in option prices to indicate that a stock market crash was expected in October 1987. Similarly, Malz (1997) fits a jump-diffusion model of exchange rate behavior. By fitting this model to option price date the author retrieves the parameters of the jump-diffusion process, which are used to estimate the option-probability distribution. According to

[1] Jackwerth (1999) presents a partial and selective review of the existing literature previous to 1999, particularly on parametric methods.

Banbula (2008), the stochastic process methods are probably among the least popular ones, due to its relatively small flexibility. Making the assumption concerning the stochastic process of the underlying instrument implies strong restrictions on the type of option-implied PDFs. This is due to the fact that a stochastic process can only follow one single distribution. However, the advantage of this approach over other methods is that once a stochastic process is identified, this method can be used to replicate options and hedge its exposure.

The *implied binomial tree* was first developed by Jackwerth and Rubistein (1996), Rubinstein (1999) and most recently also applied by Cícha (2009). This method assumes that the price of the underlying asset follows the process of a binomial tree. In this respect, this methodology assigns probability p if the underlying asset moves up and probability 1-p otherwise. According to Rubinstein (1999), the approach provides a way to generalize to arbitrary ending risk-neutral PDFs. Interpreted in terms of continuous-time diffusion processes, the model assumes that the drift and local volatility are at most functions of the underlying asset price and time by endogenously fitting current option prices.

The *finite-difference methods* try to calculate the second derivative of the call function of the Breeden and Litzenberg (1978) method described in (2.3) by using finite difference methods. These methods compute option prices numerically, by approximating the partial derivatives of the call function with finite differences. For instance, Neuhaus (1995) uses the finite difference method to the first derivative, rather than the second derivative, of the option-implied PDF. Although these methods may be the easiest to implement from the computationally point of view, they are less popular, since they require evenly spaced strike prices and their output is a discrete option-implied PDF rather than a continuous one. These methods also assume that a large number of option prices is traded for many different strikes, which in practice is not generally the case.

The *approximating function methods* assume that the option-implied PDF has a particular parametric form. The parametric form of the option-implied PDF depends on a set of parameters that are calculated by minimising the fitted error of both the observed and the theorical puts and calls. The theoretical call and put prices are derived by using Cox and Ross (1976) formula which is described in equation (2.1). The most popular approach within this methodology is the so-called ”double lognormal” introduced by Melick and Thomas (1997) in where the two authors make use of a mixture of two lognormal distributions. The double lognormal method has been also used by Söderlind and Svensson (1997), among others.

In practice, the Breeden and Litzenberger (1978) result cannot always be used to derive directly the option-implied PDF. This is particularly the case when the call function, which is a discrete function, is not twice differentiable or when a set of three strikes are close to a straight line. The *implied volatility smoothing methods*, rather than applying directly the Breeden and Litzenberger (1978) result, transform the initial strike price-option premium space into the strike price-implied volatility space by using the Black and Scholes (1973) formula. This allows approximating a continuous smoothing function in the implied volatilities strike price space. This continuous implied volatility function is converted back into a continuous call price function, which is now twice differentiable. Hence, the Breeden and Litzenberger (1978) result can now be applied to the set of continuous call price functions to obtain the underlying option-implied PDF. One of the advantages of the implied volatility smoothing methods is that they do not assume a parametric distribution of the underlying density function.

Shimko (1993) first introduced this approach by transforming the option premium into implied volatilities, using the Black and Scholes (1973) formula. Malz (1997) transformed not only the option premium but also the strike prices into delta and fitted the implied volatility smile in the delta-implied volatility space. Campa, Chang and Reider (1998) used the transformation of option premium into implied volatilities described by Shimko (1993). However, the authors made use of cubic splines, which ensure that the first derivative is continuous and differentiable through the range of strike prices. Finally, Bliss and Panigirtzoglou (2002) made use of the delta space transformation suggested by Malz (1997) and combined it with the natural spline estimation used by Campa, Chang and Reider (1998). In this thesis, the Bliss and Panigirtzoglou (2002) methodology is used.

Another classification of option-implied PDFs methodologies is provided by Jondeau *et al.* (2006). The authors classify the different models into structural and non-structural, depending on whether they take into account a specific pattern for the price or volatility of the process. The structural category would include for instance jump diffusion models or models based on stochastic volatility. The non-structural category would include parametric models, like the double log-normal method, semi-parametric (or non-parametric) models, which try to approximate to the true option-implied PDF, like the one proposed by Madan and Milne (1994) and using a Hermite polynomial approximation, and non-parametric models, in which no assumption about the distribution of the option-implied PDF is made, like for instance the methods suggested by Shimko (1993), Malz (1997) and Bliss and Panigirtzoglou (2002).

## 2.2 Methodology comparison of option-implied probability density functions

Over time there has also been an extensive discussion on the different existing methodologies and the differences between them. Since the last quarter of the twentieth century the number of methodologies has been growing and there has been no general consensus on which technique should be used in which situation. Moreover, opinions differ as to how option-implied PDFs should be used and interpreted. In this respect, several authors have tried to compare the strengths and weaknesses of the existing methodologies. For instance, Campa, Chang and Reider (1998) compared the following three methods: implied binomial trees, a smoothed implied volatility smile and a mixture of two lognormals; Coutant, Jondeau and Rockinger (1998) compared a single lognormal method, a mixture of two lognormals, Hermite polynomials and maximum entropy methods; similarly, McManus (1999) compared jump diffusion models, a mixture of lognormals, Hermite polynomials and maximum entropy methods on a set of one week data; Jondeau and

Rockinger (2000) also compared jump diffusion models, mixture of lognormals, Hermite polynomials and a Heston's approach assuming a stochastic model; Syrdal (2002) compared monthly estimations of a single lognormal method with a double lognormal method with three variants of the smoothed implied volatility smile method; Dutta and Babel (2005) derived a closed form option pricing formula for pricing European options and compared it with option prices based on the lognormal, Burr-3, Weibull, and GB2 distributions; most recently, Lai (2014) compared the performance of three non-parametric methods, i.e. kernel regression, spline interpolation and neural network models by using simulated data instead of real data.

To compare the methodologies, the authors based their results on goodness-of-fit or on the direct comparison of the PDFs percentiles and statistical moments. In general, the studies concluded that there is not much difference between the different methodologies. In particular, the differences are almost insignificant in the first two moments of the distribution. However, some authors also indicated a preference for a particular methodology. In this respect, McManus (1999) concluded that the double lognormal distribution is preferable to jump diffusion models, Hermite polynomials or maximum entropy methods; Dutta and Babel (2005) that the g-and-h distribution were preferable compared to the other distributions analysed in their study and Lai (2014) pointed out that the kernel regression yields the best performance, followed by the spline interpolation.

Additionally to these studies, the Bank for International Settlements (BIS) decided to organise a one-day workshop on the estimation of option-implied PDFs methods that were used in the central banking community. At the BIS workshop (1999) 14 different methodologies to estimate the implied PDFs were compared by using a common dataset. This common data set referred to 61 trading days of options on Eurodollar futures traded on the Chicago Mercantile Exchange from September 1, 1998 through November 30, 1998 for the December 1998 contract. In order to compare the implied PDFs of the different methodologies, the participants were asked to provide the mean, the standard deviation and the following 11 percentiles 0.5%, 0.1%, 5%, 10%, 25%, 50%, 75%, 90%, 95%, 99% and 99.5%

of their methodology.

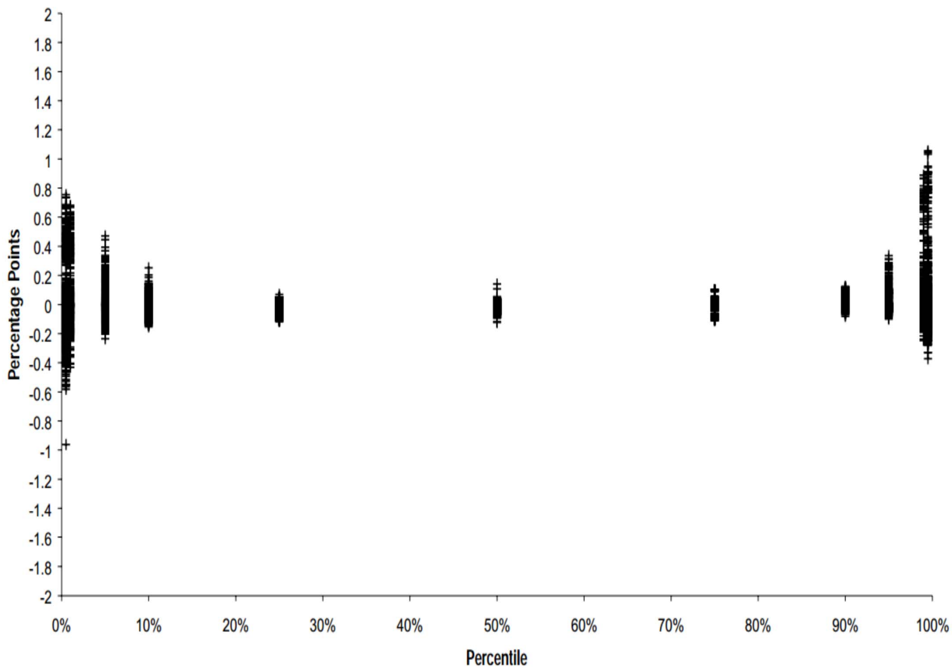


Source: BIS

**Figure 2.1:** Scatter plot of scaled percentiles of different methodologies as presented at the BIS.

The results shown in Figure 2.1 indicate that between the 10% and 90% percentiles there is not much difference between the techniques. In fact, all of the large outliers from the 11 medians of the 11 percentiles (for the 14 methods) that were observed, are due to the inclusion of the mixture of lognormal methodology in the median calculation (which is excluded from Figure 2.1). By excluding the mixture of lognormal methodology, the largest deviation from any of the median of the 11 percentiles amounts to 96 basis points which occur in the 0.5% percentile. The range of deviations from the 50% percentile amounts to only 25 basis points. That is, as concluded in the BIS Workshop (1999), practitioners can have some confidence that the results they report are not overly sensitive to the particular method they use to estimate the option-implied PDF. Outside of these percentiles, the sensitivity to the technique increases dramatically.

Figure 2.1 shows that, in general, although these methods might differ in the extreme of the tails of the distribution, there is generally no major difference in the central section of the estimated option-implied PDFs. And, arguably, it is the central section of the option-implied PDF which is more likely to be useful for monetary policy purposes, in contrast to financial stability analysis, where there may be greater focus on the tails of the distribution. In this respect, and as concluded in the BIS Workshop, an analysis for a value-at-risk calculation using option-implied PDF estimation to provide a measure of the future short-term interest rate below, having less than a 1% chance of falling, will be quite sensitive to the choice of option-implied PDF estimation technique.

Of particular interest are the results presented by Bliss and Panigirzoglou (2002) comparing the double lognormal method with the smoothed implied volatility smile methods. The authors give in this case particular attention to the stability of the option-implied probability functions and the robustness of the estimates produced. They examined the extent to which small perturbations in actual options prices generated large changes in the estimated option-implied PDFs. The parametric and non-parametric methods were then evaluated by comparing the sample distributions of a number of summary statistics. The authors concluded that the smooth implied volatility smile is more robust to small perturbations in actual options prices than the double lognormal method.

## 2.3 Bliss and Panigirtzoglou method to estimate the option-implied probability density functions

Up to date there is no major consensus on which is the best methodology to derive option-implied PDFs. For this reason, the option-implied PDFs which are presented in this study are based on the Bliss and Panigirtzoglou (2002) method. The reason behind is that although the computational implementation of this method might not be straightforward, this methodology does not rely on a functional form and hence on a set of initial parameters. In this respect, this methodology can run on a set of data without possible convergence problems when trying to optimise the function parameters. This phenomena, as pointed out by Bliss and Panigirtzoglou (2002), can generate implausibly large changes in the shape of the option-implied PDF between consecutive days. This is true particularly for measures of the skewness and kurtosis of the distribution, which are of particular interest when comparing different periods of data. The methodology derived by Bliss and Panigirtzoglou (2002) is described in detail in Chapter 4.

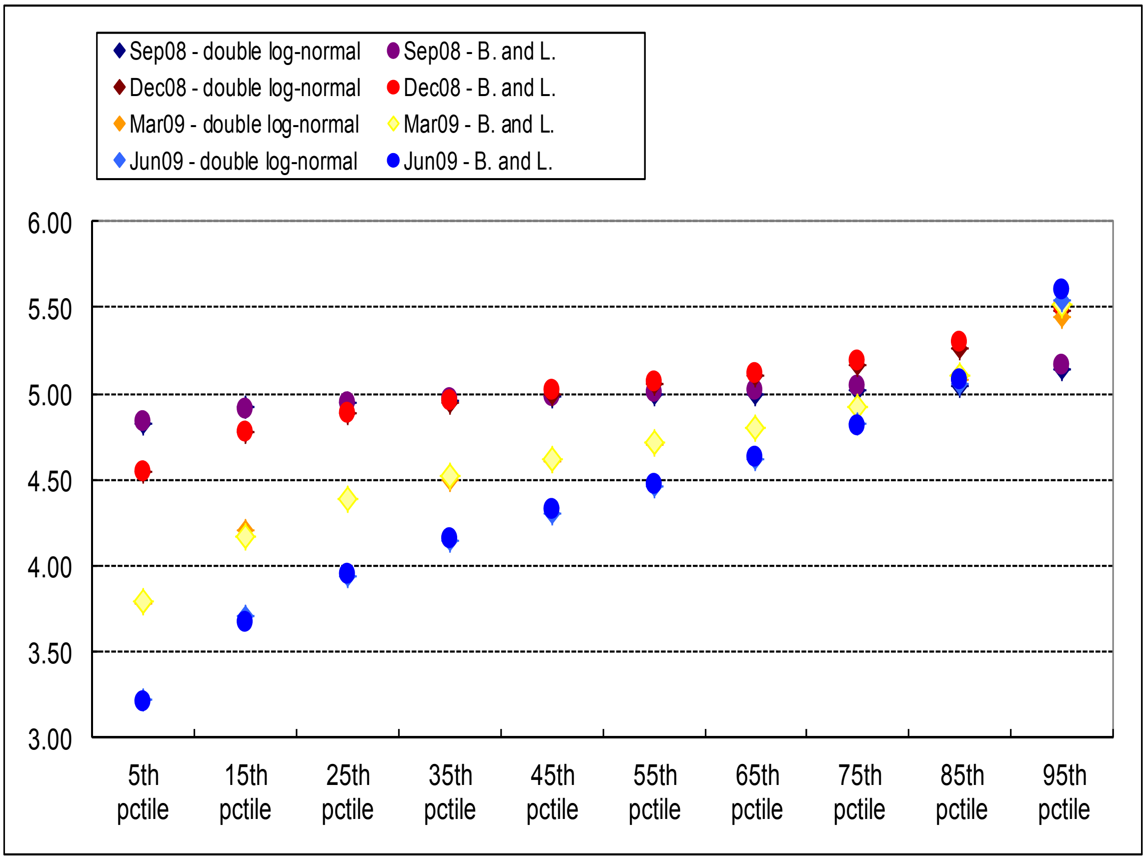


**Figure 2.2:** Percentiles comparison calculated with the Breeden and Litzenberg methodology and the double lognormal method, 18 August 2008.

Previous to applying this methodology, and similar to what other authors have done, the Breeden and Litzenberger (1978) method was compared to the Neuhaus (1995) and the double log-normal methods covering the period between January to August 2008 for four different contracts. The results were slightly better than the ones presented in the BIS workshop (1999). This can be simply explained by the fact that the comparison of three methodologies was made by using a different time period. Overall, the differences between the three methodologies in the observed percentiles were around -12 to 11 basis points and -5 to 2 basis points when excluding the extreme of the tails of the distribution. In sum, the comparison showed that there is generally no major difference in the central section of the estimated option-implied PDFs and that the three methodologies could be used effectively for monetary policy purposes. Figure 2.2 above presents the percentiles comparison of the Breeden and Litzenberger (1978) result with the double lognormal method for a given day.

By applying the double lognormal method one can see that a parametric approach has also another important caveat to be taken into account. Figure 2.3 shows two cases where the double lognormal method generated option-implied PDFs with a sharp spike. In this case, the initial parameters that were used did not allow the optimisation function to converge to the global minimum. Instead, the optimisation function converged to a set of parameters of two lognormal distributions, but one of them with a very small standard deviation. Additionally, this does not support the use of the lognormal method with a large set of data by using an iterative process. This is particularly the case for this study, where option-implied PDFs are derived from a set of 10 years of daily options.

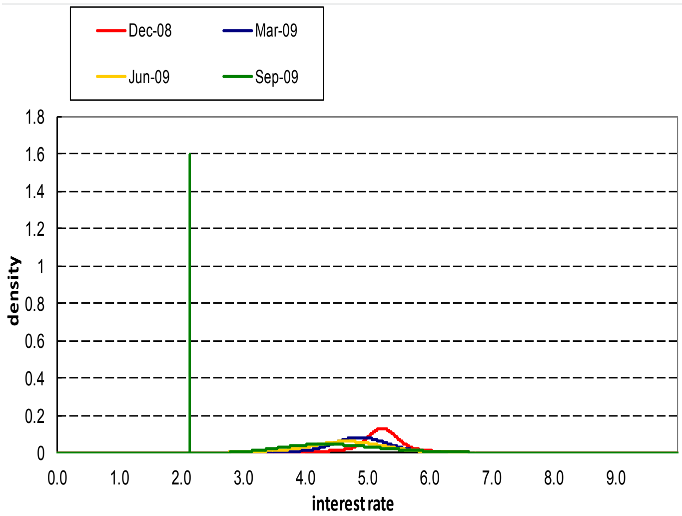


**(a)** 25 September 2008

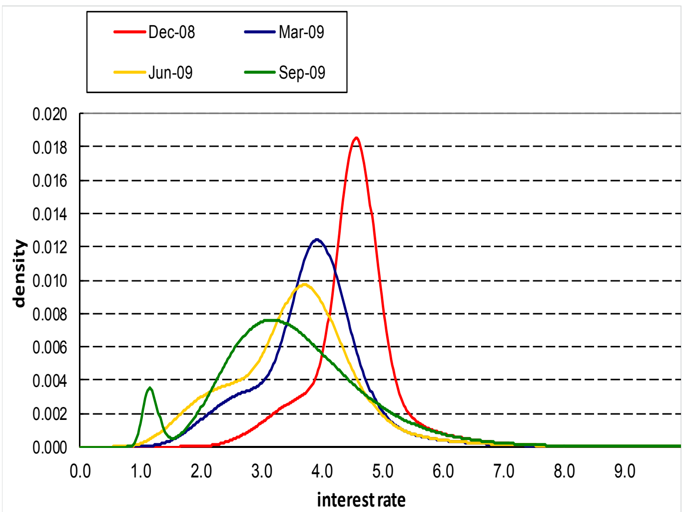


**(b)** 6 October 2008

**Figure 2.3:** Caveat of the usage of the double lognormal method on two given days.

*In God we trust. All others must bring data.*

W. Edwards Deming

CHAPTER

# 3

# The Data

## 3.1 3-Months Euribor futures and options

The data used to derive the option-implied probability density functions (PDFs) refer to daily settlement prices of futures on the 3-month Euribor and on prices of options written on the 3-month Euribor futures. The dataset presented in this Chapter covers the complete history of Euribor futures options until 2014, thus consisting of sixteen years of daily observations, i.e. from January 1999 until December 2014. However, each journal paper on which this thesis is based covers the period of data available at the time that the paper was published. In particular, the second paper extends the underlying methodology to intra-day settlement prices in a set of specific days rather than daily ones.

The Euribor, or Euro Interbank Offered Rate, was established in 1999 after the introduction of the Euro as a daily reference rate within the Economic and Monetary Union (EMU). The Euribor is based on the average interest rate at which banks offer to lend unsecured funds to other banks in the interbank market. Euribor futures and options on Euribor futures are financial derivatives, traded on NYSE Liffe, whose terminal value depends on the outturn of Euribor.

The daily settlement prices of futures and options on the 3-month Euribor futures are published by Euronext.liffe, formed in January 2002 from the takeover of the London International Financial Futures and Option Exchange (LIFFE).[1] According to LIFFE, these contracts were developed in response to the economic and monetary union within Europe, and the emergence of Euribor as the key cash market benchmark within Europe's money markets. Since its launch, LIFFE's Euribor contracts have come to dominate the euro denominated short-term interest rate (STIR) derivatives market, capturing over 99% of the market share; they are now the most liquid and heavily traded euro-denominated STIR contracts in the world.

Delivery months for the 3-month Euribor futures contracts are March, June, September and December; the last trading day is two business days prior to the third Wednesday of the delivery month, and the delivery day is the first business day after the last trading day. The settlement price will be 100 minus the EBF Euribor Offered Rate rounded to three decimal places. On a given day, twelve option contracts with fixed expiry date are traded. The first five options expire on the following five closest months from the respective day, while the remaining contracts expire in the next quarters.[2] Due to the quarterly delivery structure of the future contracts, the options are settled with the assignment of a futures contract at the exercise price and with the respective quarterly delivery. For example, the futures delivery month associated with options expiring in January, February and March is March, while for those expiring in April, May and June is June.

Each file used as input for our model on a given day contains the following variables: the reference date, e.g. the day in which options were traded, the maturity date when the option contract can be exercised, the strike price of the option, the

[1] In the past, one could easily download the data directly from the LIFFE internet website via the following link http://www.liffe.com/reports/eod?item=Histories. However, since November 2014, after a series of takeovers from LIFFE, the data can be downloaded from https://www.theice.com/products/38527989/Options-on-Three-Month-Euribor-Futures.

[2] For instance, on the 3rd of October 2014 the initial five option contracts were expiring in October, November, December of 2014 and January and February of 2015. The remaining seven contracts were expiring in March, June, September and December 2016 and March, June, September 2017.

option type indicator (e.g. 1 indicates that the option is a call option and 2 indicates a put option), the volume (e.g. the total number of call or puts being traded for this particular strike), the premium of the put and call options and the spot price of the underlying future. Table 3.1 presents an example of the first contract expiring on 19 September 2009 for the 7 July 2009 data.

| Reference Date | Maturity Date | Strike Price | Option Type Indicator | Volume | Premium | Spot Price |
|---|---|---|---|---|---|---|
| 20090707 | 20090914 | 98.375 | 1 | 0 | 0.685 | 99.06 |
| 20090707 | 20090914 | 98.375 | 2 | 11000 | 0.005 | 99.06 |
| 20090707 | 20090914 | 98.500 | 1 | 0 | 0.5625 | 99.06 |
| 20090707 | 20090914 | 98.500 | 2 | 0 | 0.0025 | 99.06 |
| 20090707 | 20090914 | 98.625 | 1 | 0 | 0.44 | 99.06 |
| 20090707 | 20090914 | 98.625 | 2 | 0 | 0.005 | 99.06 |
| 20090707 | 20090914 | 98.750 | 1 | 1000 | 0.32 | 99.06 |
| 20090707 | 20090914 | 98.750 | 2 | 0 | 0.01 | 99.06 |
| 20090707 | 20090914 | 98.875 | 1 | 6750 | 0.205 | 99.06 |
| 20090707 | 20090914 | 98.875 | 2 | 200 | 0.025 | 99.06 |
| 20090707 | 20090914 | 99.000 | 1 | 30775 | 0.115 | 99.06 |
| 20090707 | 20090914 | 99.000 | 2 | 0 | 0.05 | 99.06 |
| 20090707 | 20090914 | 99.125 | 1 | 58250 | 0.045 | 99.06 |
| 20090707 | 20090914 | 99.125 | 2 | 0 | 0.115 | 99.06 |
| 20090707 | 20090914 | 99.250 | 1 | 36850 | 0.015 | 99.06 |
| 20090707 | 20090914 | 99.250 | 2 | 0 | 0.21 | 99.06 |
| 20090707 | 20090914 | 99.375 | 1 | 3500 | 0.005 | 99.06 |
| 20090707 | 20090914 | 99.375 | 2 | 0 | 0.32 | 99.06 |
| 20090707 | 20091214 | 97.750 | 1 | 0 | 1.1975 | 98.945 |
| 20090707 | 20091214 | 97.750 | 2 | 0 | 0.0025 | 98.945 |
| 20090707 | 20091214 | 97.875 | 1 | 0 | 1.075 | 98.945 |
| 20090707 | 20091214 | 97.875 | 2 | 0 | 0.005 | 98.945 |
| 20090707 | 20091214 | 98.000 | 1 | 1100 | 0.945 | 98.945 |
| 20090707 | 20091214 | 98.000 | 2 | 0 | 0.01 | 98.945 |

**Table 3.1:** Input data on 7 July 2009 - extract.

Figure 3.1 presents a chart with the eight fixed expiring contracts that were traded on 7 July 2009. The eight contracts traded on this date expire in September and December of 2009, in March, June, September and December of 2010, and in March and June of 2011. For each contract, the chart presents the initial calls and puts (in red and blue) on interest rates. Additionally, the chart also shows the volumes that were traded, in number of transactions, for some of the strikes. In particular, there is heavy trading for options expiring in less than 3 months and less than 6 months, little trading for those contracts expiring in more than one year and no trading for contracts expiring in around two years. In the absence of trades, the settlement prices used for the options were based on quotes directly given by LIFFE.

Section 3.2 presents a detailed analysis of the volumes being traded and in particular focuses on in out-of-the-money options, those used as input in our methodology.

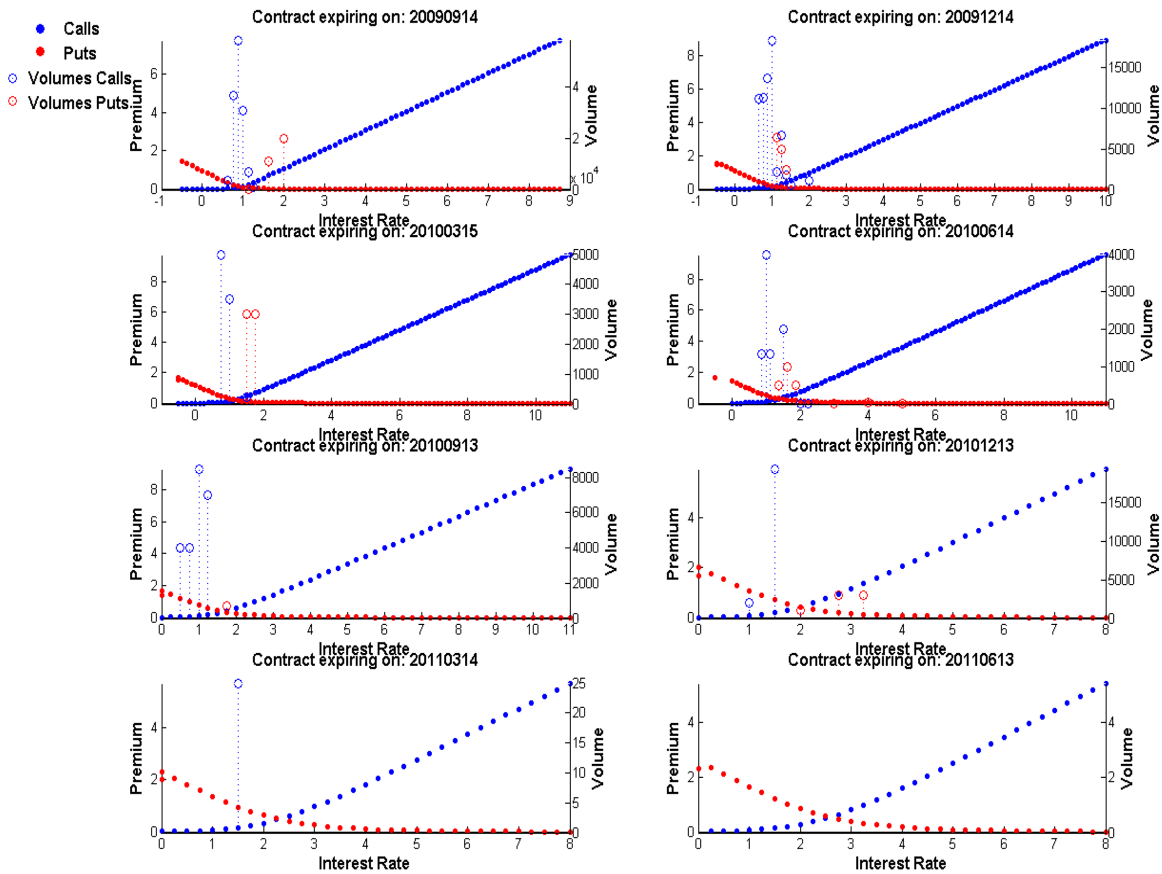


**Figure 3.1:** Input data for 8 contracts on 7 July 2009.

## 3.2 Volume analysis and out-of-the-money options

This Section presents the trading volume for all Euribor options from the first day of trading, 13 January 1999 until December 2014. The results of this analysis are shown in Table 3.2 and are classified by contract maturity.

As initially presented in Figure 3.1, the results show that trading is in fact more concentrated in those options contracts maturing in nine months or less. The trading for these contracts accounts for more than the 85% of the total trading. For this reason, the option-implied probabilities which are derived and studied make always reference to contracts expiring in less than one year; in particular, the option-implied PDFs are derived by using contracts expiring in less than 6-months, accounting for almost 65% of the total trading.

The number of traded contracts has increased steadily since this instrument was first introduced, with most options being traded in the most recent years. In particular, as presented in Table 3.2, the maximum trading occurs between 2009 and 2010. This can be partially explained in the aftermath of the collapse of Lehman Brothers and the decline of the ECB official interest rates.

In absolute terms, 18% of the options are traded in-the-money, 1% of the options are traded at-the-money and 81% of the options are traded out-of-the-money. This is in line with the results presented by Bliss and Panigirtzoglou (2004), where out-of-the-money calls (puts) tend to be more liquid than puts (calls) of the same strike. Furthermore, some of the in-the-money options are traded not independently, but as part of a bundled trading strategy, e.g. straddles or strangles, which combine out-of-the-money options with in the-money options. For this reason, the Bliss and Panigirtzoglou (2004) methodology, described in Chapter 4, is applied to those option prices which are either at- or out-of-the-money, but not in-the-money.

| | | | | | | |
|---|---|---|---|---|---|---|
| | **Contract expiring in less than 3-months** | | | | | |
| | **Call** | | | **Put** | | |
| | In-the-money | Out-of-the-money | At-the-money | In-the-money | Out-of-the-money | At-the-money |
| Volume traded | 35,061,134 | 122,072,867 | 1,682,755 | 12,771,469 | 59,878,167 | 824,466 |
| Mean | 8,627 | 30,038 | 414 | 3,143 | 14,734 | 203 |
| Std. Dev. | 22,407 | 64,081 | 3,202 | 9,657 | 31,183 | 1,877 |
| Maximum | 561,194 | 1,675,662 | 104,500 | 339,499 | 544,707 | 48,851 |
| Maximum Date | 20081013 | 20090109 | 20080411 | 20080925 | 20100507 | 20110113 |
| Volume per option type | 158,816,756 | | | 73,474,102 | | |
| Total volume | **232,290,858 (27.82%)** | | | | | |
| | **Contract expiring between 3 and 6 months** | | | | | |
| | **Call** | | | **Put** | | |
| | In-the-money | Out-of-the-money | At-the-money | In-the-money | Out-of-the-money | At-the-money |
| Volume traded | 36,168,922 | 159,708,826 | 1,154,422 | 17,589,698 | 88,082,940 | 785,691 |
| Mean | 8,900 | 39,298 | 284 | 4,328 | 21,674 | 193 |
| Std. Dev. | 19,109 | 57,117 | 3,373 | 13,484 | 38,261 | 2,306 |
| Maximum | 226,775 | 592,700 | 116,436 | 421,595 | 595,077 | 65,576 |
| Maximum Date | 20080207 | 20091102 | 20070810 | 20070215 | 20090623 | 20110303 |
| Volume per option type | 197,032,170 | | | 106,458,329 | | |
| Total volume | **303,490,499 (36.34%)** | | | | | |
| | **Contract expiring between 6 and 9 months** | | | | | |
| | **Call** | | | **Put** | | |
| | In-the-money | Out-of-the-money | At-the-money | In-the-money | Out-of-the-money | At-the-money |
| Volume traded | 18,068,333 | 90,719,909 | 508,909 | 11,470,843 | 58,882,240 | 510,042 |
| Mean | 4,446 | 22,323 | 125 | 2,823 | 14,489 | 126 |
| Std. Dev. | 13,073 | 38,572 | 1,192 | 9,354 | 34,039 | 1,653 |
| Maximum | 286,500 | 449,700 | 38,901 | 222,122 | 880,283 | 58,805 |
| Maximum Date | 20100304 | 20100111 | 20070502 | 20070719 | 20091211 | 20110308 |
| Volume per option type | 109,297,151 | | | 70,863,125 | | |
| Total volume | **180,160,276 (21.57%)** | | | | | |
| | **Contract expiring between 9 months and 12 months** | | | | | |
| | **Call** | | | **Put** | | |
| | In-the-money | Out-of-the-money | At-the-money | In-the-money | Out-of-the-money | At-the-money |
| Volume traded | 6,063,629 | 41,021,534 | 182,604 | 4,382,082 | 22,304,807 | 176,738 |
| Mean | 1,492 | 10,094 | 45 | 1,078 | 5,488 | 43 |
| Std. Dev. | 6,624 | 27,549 | 576 | 5,073 | 17,087 | 705 |
| Maximum | 240,900 | 765,140 | 19,520 | 183,000 | 453,200 | 36,200 |
| Maximum Date | 20100303 | 20031210 | 20030317 | 20060912 | 20100514 | 20060914 |
| Volume per option type | 47,267,767 | | | 26,863,627 | | |
| Total volume | **74,131,394 (8.88%)** | | | | | |
| | **Contract expiring between 12 months and 15 months** | | | | | |
| | **Call** | | | **Put** | | |
| | In-the-money | Out-of-the-money | At-the-money | In-the-money | Out-of-the-money | At-the-money |
| Volume traded | 1,692,030 | 10,955,276 | 63,585 | 1,356,319 | 6,693,531 | 52,069 |
| Mean | 416 | 2,696 | 16 | 334 | 1,647 | 13 |
| Std. Dev. | 1,602 | 8,438 | 212 | 1,920 | 6,600 | 213 |
| Maximum | 30,015 | 177,325 | 5,500 | 68,750 | 237,900 | 10,000 |
| Maximum Date | 20120103 | 20131115 | 20051121 | 20131115 | 20091210 | 20050907 |
| Volume per option type | 12,710,891 | | | 8,101,919 | | |
| Total volume | **20,812,810 (2.49%)** | | | | | |
| | **Contract expiring between 15 months and 18 months** | | | | | |
| | **Call** | | | **Put** | | |
| | In-the-money | Out-of-the-money | At-the-money | In-the-money | Out-of-the-money | At-the-money |
| Volume traded | 882,891 | 4,804,957 | 25,570 | 733,761 | 3,713,815 | 29,858 |
| Mean | 217 | 1,182 | 6 | 181 | 914 | 7 |
| Std. Dev. | 1,180 | 5,450 | 103 | 901 | 3,789 | 126 |
| Maximum | 50,010 | 170,500 | 3,250 | 24,010 | 84,300 | 5,000 |
| Maximum Date | 20030527 | 20040805 | 20070831 | 20091222 | 20110715 | 20100811 |
| Volume per option type | 5,713,418 | | | 4,477,434 | | |
| Total volume | **10,190,852 (1.22%)** | | | | | |
| | **Contract expiring between 18 months and 21 months** | | | | | |
| | **Call** | | | **Put** | | |
| | In-the-money | Out-of-the-money | At-the-money | In-the-money | Out-of-the-money | At-the-money |
| Volume traded | 653,584 | 3,120,354 | 46,181 | 763,689 | 3,008,411 | 25,425 |
| Mean | 161 | 768 | 11 | 188 | 740 | 6 |
| Std. Dev. | 775 | 3,540 | 245 | 1218 | 3,738 | 117 |
| Maximum | 27,000 | 135,750 | 10,000 | 52,000 | 135,120 | 4,500 |
| Maximum Date | 20120718 | 20110909 | 20120531 | 20130301 | 20130301 | 20100317 |
| Volume per option type | 3,820,119 | | | 3,797,525 | | |
| Total volume | **7,617,644 (0.91%)** | | | | | |
| | **Contract expiring between 21 months and 24 months** | | | | | |
| | **Call** | | | **Put** | | |
| | In-the-money | Out-of-the-money | At-the-money | In-the-money | Out-of-the-money | At-the-money |
| Volume traded | 595,195 | 2,217,296 | 17,004 | 719,958 | 2,788,972 | 18,208 |
| Mean | 146 | 546 | 4 | 177 | 686 | 4 |
| Std. Dev. | 723 | 3,072 | 83 | 852 | 3,287 | 87 |
| Maximum | 23,500 | 124,000 | 2,500 | 27,000 | 81,011 | 2,500 |
| Maximum Date | 20100129 | 20090305 | 20070529 | 20060210 | 20130814 | 20070529 |
| Volume per option type | 2,829,495 | | | 3,527,138 | | |
| Total volume | **6,356,633 (0.76%)** | | | | | |

**Table 3.2:** Total volume traded between 1999 and 2014.

Additionally, we also analyse whether the volume of options which are out-of-the money is concentrated closer to the future or either far away, e.g. how deep out-of-the-money are the options being analysed. In order to do so, the call and put options are classified in three groups: out-of-the-money options, referring to those options which are between 25 basis points and 100 basis from the underlying future, deep out-of-the-money, refering to those options which are between 100 basis points and 175 points from the underlying future, and very deep out-of-the-money, referring to those options whose difference with the underlying future is bigger than 175 basis points.

Table 3.3 presents the percentage of trading for puts and calls, in respect with the total volume, in each of the three categories described above. In absolute terms, the volume of options traded out-of-the-money accounts for 97% of the total volume of options traded out-of-the-money. Additionally, 2% of the volume of the options is traded deep out-of-the money with respect with the total volume of options traded out-of-the-money, and only 1% (or even less) of the options which are out-of-the-money are traded very deep out-of-the-money. For this reason, a minimum premium threshold of 25 basis points is set up to avoid collecting those options which are very deep out-of-the-money or any option with premium either zero or negative.

| | CALLS | | | PUTS | | |
|---|---|---|---|---|---|---|
| | Out-of-the-money | Deep out-of-the-money | Very deep out-of-the-money | Out-of-the-money | Deep out-of-the-money | Very deep out-of-the-money |
| 1999 | 71.94% | 1.66% | 0.00% | 26.04% | 0.33% | 0.03% |
| 2000 | 57.93% | 2.15% | 0.08% | 38.91% | 0.56% | 0.37% |
| 2001 | 62.96% | 1.28% | 0.15% | 34.45% | 0.97% | 0.19% |
| 2002 | 71.52% | 2.95% | 0.08% | 24.04% | 1.33% | 0.08% |
| 2003 | 73.96% | 1.78% | 0.20% | 23.66% | 0.23% | 0.17% |
| 2004 | 73.84% | 1.14% | 0.00% | 24.52% | 0.37% | 0.12% |
| 2005 | 63.04% | 0.21% | 0.00% | 36.46% | 0.19% | 0.10% |
| 2006 | 53.20% | 0.32% | 0.02% | 46.29% | 0.18% | 0.01% |
| 2007 | 54.11% | 0.46% | 0.01% | 45.35% | 0.07% | 0.01% |
| 2008 | 59.43% | 5.06% | 0.79% | 32.33% | 2.03% | 0.36% |
| 2009 | 70.00% | 1.69% | 0.13% | 24.85% | 2.11% | 1.21% |
| 2010 | 54.55 | 0.53% | 0.00% | 42.58% | 1.61% | 0.73% |
| 2011 | 65.47% | 0.77% | 0.00% | 32.17% | 1.21% | 0.38% |
| 2012 | 55.63% | 0.00% | 0.00% | 40.40% | 2.90% | 1.07% |
| 2013 | 58.78% | 0.00% | 0.00% | 37.20% | 3.26% | 0.76% |
| 2014 | 58.13% | 0.00% | 0.00% | 40.80% | 0.97% | 0.10% |

**Table 3.3:** Out-of-the-money data.

## 3.3 Convexity and monotonicity

In addition to deleting those options which are not out-of-the money, the convexity and monotonicity of the call and put options is checked. This owes to the option-pricing theory, whereby a call price function should be both monotonic and convex in order to yield non-negative probability estimates. Hence, those options which do not fulfill the convexity and monotonicity criteria are deleted.

In order to check their convexity, calls and puts are analysed separately. Both of them are sorted by strike price in an ascending order. The convexity restriction can only be checked in case three or more options are available per contract and is applied to each set of 3 given strike prices, i.e. calls and puts, $K_1$, $K_2$, $K_3$. If $y = mK + n$ is the straight line that passes through points $K_1$ and $K_3$ then $m = \frac{f(K_3)-f(K_1)}{K_3-K_1}$ and $n = f(K_3) - \frac{f(K_3)-f(K_1)}{K_3-K_1}K_3$. In this case, a function $f$ is convex, if and only if, $f(K_2) \leq mK_2 + n$.

The last inequation of the paragraph above is equivalent to fulfill:

$$f(K_2) \leq \frac{f(K_3) - f(K_1)}{K_3 - K_1}K_2 + f(K_3) - \frac{f(K_3) - f(K_1)}{K_3 - K_1}K_3 \tag{3.1}$$

To avoid numerical problems, 3.1 is expressed as follows:

$$(K_3 - K_1)f(K_2) \leq f(K_3)(K_2 - K_1) + f(K_1)(K_3 - K_2) \tag{3.2}$$

The monotonicity inequation 3.2 is applied from the second sorted call (put) to the last but one sorted call (put), and each call (put) is analysed with the two surrounding calls (puts). For each group of three calls (puts) with strikes $K_1$, $K_2$, $K_3$ the premium of the call (put) that is in the middle, i.e. $f(K_2)$, needs to be less than or equal to the value that would be obtained interpolating linearly the two surrounding calls (puts). After checking all the triples, those that do not hold the restriction are treated one by one. Not necessarily, the point which is in the middle is the one not fulfilling the convexity criteria. For each non-convex triple, each point is analysed, and then the point of the triple that may cause the largest number of convexity violations is removed. Therefore, every point of each non-convex triple is analysed. Then the point of the triple with the largest number of convexity violations is removed. This is done since not necessarily the point which is in the middle is the one not fulfilling the convexity criteria. This action is repeated until all triples hold the convexity checks.

In addition, the monotonicity of the call and put functions is also checked. Puts and calls are again separated into two groups and sorted by the strike in an ascending way. The monotonicity restriction cannot be checked when less than 2 options are available per contract.[1] Therefore, if there are more than two calls (puts), the following inequation is checked to verify the monotonicity $\frac{f(K_i)-f(K_i+1)}{K_i-K_i+1} \leq 0$.

Moreover, additionally to the monotonicity and convexity check and to avoid the instability of the smile function which is calculated in Chapter 4, options with a delta outside the range of 0.01 to 0.99 are removed. Furthermore, call options and put options with the same delta are deleted. Finally, if a contract on a given day has less than two calls and two puts the option-implied PDF for this particular day cannot be derived.

[1] A function f is called monotonically increasing if for all $x$ and $y$ such that $x \leq y$ one has $f(x) \leq f(y)$, so f preserves the order. Likewise, a function is called monotonically decreasing if, whenever $x \geq y$ one has $f(x) \geq f(y)$, so it reverses the order.

*All models are wrong, but some are useful.*

George E. P. Box

CHAPTER

# 4

# Underlying methodology

## 4.1 Bliss-Panigirzoglou method

The Bliss and Panigirtzoglou (2002) method is part of the *implied volatility methods* described in Section 2.1. This method makes use of the Breeden and Litzenberger (1978) result described in 1.6. However, in this case, the method does not take directly the second derivative of the call price function with respect to the strike price. In fact, the method previously transforms the initial data on option premia and strike prices into implied volatility and delta values. This is done to avoid taking directly the second derivative of a call price function and interpolating through a discrete set of calls which can lead to unstable or inaccurate results.

### 4.1.1 Methodology description

The procedure for computing the option-implied probability density function (PDF) as defined by Bliss and Panigirtzoglou (2002) can be described in several steps which are set out in the following subsections.

#### 4.1.1.1 **Delta, Sigma and Vega calculation**

After having filtered the data[1], from the observed call option prices of a given contract, the *first step* consists in transforming the observed call option prices (see Chapter 3) into implied volatilities. Implied volatilities are computed by numerically solving for the value of $\sigma$, which solves the Black (1976) futures options pricing model for each option contract at time *t*:

$$C_t(F_t, K_i, \tau) = (F_t\Theta(\frac{\log(\frac{F_t}{K_i}) + \frac{\sigma^2}{2}\tau}{\sigma\sqrt{\tau}}) - \Theta(\frac{\log(\frac{F_t}{K_i}) - \frac{\sigma^2}{2}\tau}{\sigma\sqrt{\tau}})) \tag{4.1}$$

where $C$ is the call price function, $F_t$ is the price of the underlying future at time $t$, $K_i$ is the i-th strike price, $T$ is the option expiry date, r is the risk-free rate, $\tau$ is the remaining time to maturity and $\Theta$ is the standard normal cumulative distribution function.

Following this calculation, we use the implied volatilities, i.e. the values of $\sigma$ obtained in 4.1, to calculate the values of the delta and vega parameters, i.e. $\delta$ and $\vartheta$ respectively, by using the following two formulae:

$$\delta_t(F_t, K_i, \tau) = \frac{\partial C_t(F_t, K_i, \tau)}{\partial K} = (\Theta(\frac{\log(\frac{F_T}{K_i}) + \frac{\sigma^2}{2}\tau}{\sigma\sqrt{\tau}})) \tag{4.2}$$

$$\nu_t(F_t, K_i, \tau) = \frac{\partial C_t(F_t, K_i, \tau)}{\partial F_t} = \frac{K_i\sqrt{T}}{\sqrt{2\pi}} e^{\frac{-(\frac{\log(\frac{F_T}{K_i}) + \frac{\sigma^2}{2}\tau}{\sigma\sqrt{\tau}})^2}{2}} \tag{4.3}$$

The values for the implied volatilities, deltas and vega parameters are also calculated from the observed put options using put-call parity.

#### 4.1.1.2 **Delta-Sigma space**

The purpose of the results derived in the previous section is to construct a delta-sigma grid of 1,000 points.[2] To do so, the methodology first calculates the parameters of the cubic smoothing natural spline of the initial pair of (*delta*, *sigma*).

[1]Before applying the methodology, the data are filtered as described in Chapter 3.

[2]The number of points in the grid is arbitrary but at least 1,000 points are needed to avoid instability.

Following the calculation of these parameters, a grid of 1,000 deltas is constructed. Finally, by applying the cubic smoothing natural spline in the grid of 1,000 deltas the methodology also derives a grid of 1,000 sigmas.

To construct the parameters of the cubic smoothing spline the *(delta,sigma)* points obtained in 4.1 and 4.2 are interpolated by fitting a cubic smoothing spline function which minimizes:

$$\lambda \sum_{i=1}^{n} \omega_i (\sigma_i - g(\delta_i))^2 + (1-\lambda) \int g^2(t)^2 dt \tag{4.4}$$

where $\lambda$ is the smoothing roughness parameter, equal to 0.99,[1] $delta$ is the Black-Scholes *delta* and represents the x-axis of the spline, $\sigma$ is the Black-Scholes *sigma* and represents the y-axis of the spline and the weights $\omega_i$ are calculated using $\omega_i = \frac{\nu_i^2}{mean(\nu_i^2)}$ where $\nu_i$ is Black-Scholes *vega*.

The value of *vega* is almost negligible for options which are deep out-of-the-money and deep in-the-money and sequentially increases as we get near-the-money. In particular, it reaches a maximum for at-the-money options. Hence, the $\omega_i$ used in 4.4 place most weight on near-the-money options, and therefore lesser weight on away-from-the-money options. Hence, the advantage of interpolating in the implied volatility and delta space rather than in the initial premia and strike price space is also that the options which are deep out-of-the-money (i.e. class with a delta close to zero and puts with a delta close to one) are grouped together. This is consistent with using these option-implied PDFs to support monetary policy analysis, where interest is likely to lie in the centre of the distribution, i.e. close to the underlying interest rate, rather than the distribution's tails.[2]

[1]The optimal smoothing roughness parameter is the one that minimizes the observed deltas with the fitted deltas by the smoothing spline.

[2]Although the Black-Scholes (1973) formulae is being used, no assumptions of the Black-Scholes (1973) option pricing paradigm are being assumed - in particular the implicit underlying asset price dynamics - hold true. They merely provide convenient transformation allowing the option data to be interpolated in a way that produces more stable results. That transformation is later undone.

Although *delta* can take values between 0 and $exp(r\tau)$, the traded contracts may not span that complete range. Therefore, the cubic smoothing spline is extrapolated outside the range of traded price points with a second order polynomial, i.e. a quadratic equation. As a result of the extrapolation, the piecewise cubic curve obtained using interpolation is extended with a quadratic curve at each end-point so that the full *delta* range, defined in the interval [0, 1], is covered. Figure 4.1 presents the extrapolated part of the the cubic smoothing spline in yellow.

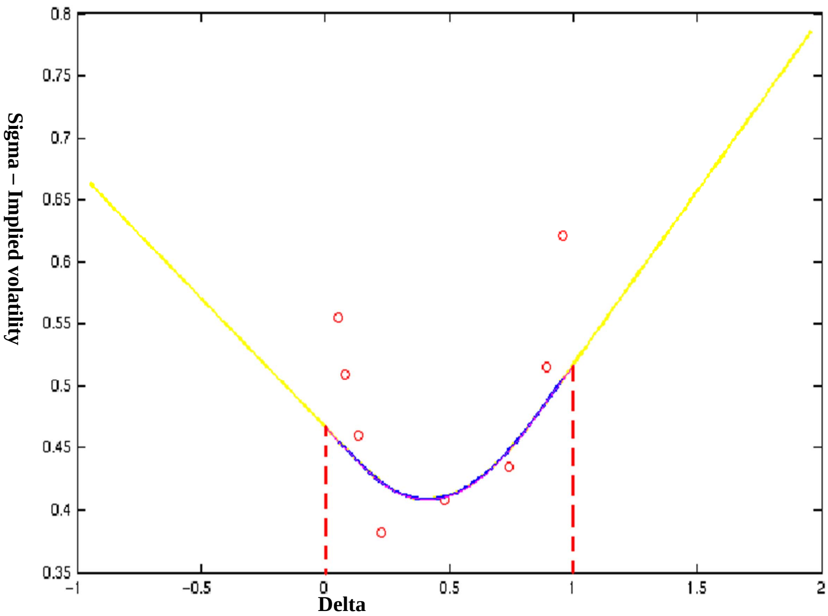


**Figure 4.1:** Input points, fitted spline, extrapolated line and evaluated 1,000 deltas.

In the next step, we calculate a grid of 1,000 delta points by dividing the full delta range in 1,000 equally-spaced points. Following this calculation, the 1,000 equally-spaced *deltas* are evaluated using the parameters of the cubic smoothing spline to obtain a grid of 1,000 sigma points. Figure 4.2 presents the fitted spline evaluated at the grid of 1,000 deltas.

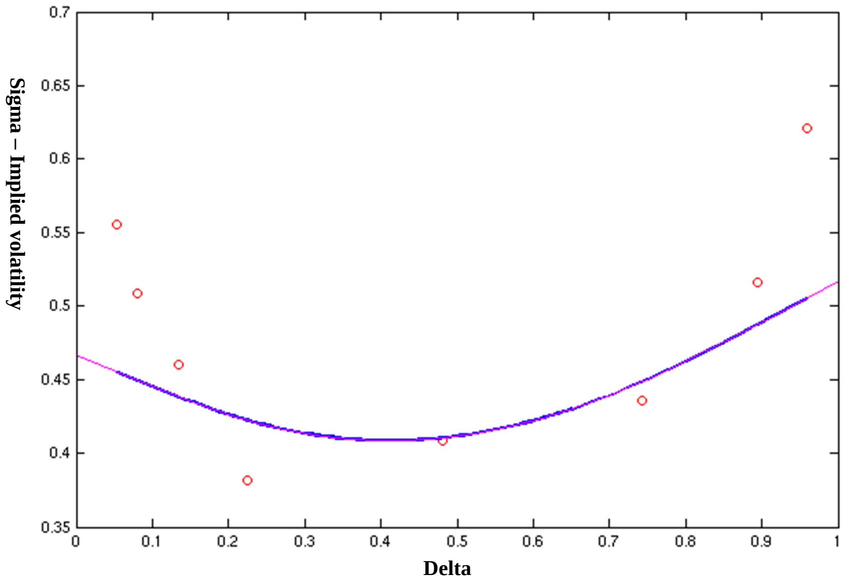


**Figure 4.2:** Input points, fitted spline and 1,000 evaluated deltas.

#### 4.1.1.3 **Back to premium strike space**

In the next step, the grid of 1,000 (*delta*, *sigma*) points are transformed back into a grid of 1,000 (*delta*, *sigma*) points.

The 1,000 *delta* values are then transformed back into strike prices using the inverse of equation 4.2:

$$F_T exp((\frac{\sigma_{ATM}^2}{2}\tau) - \sigma\sqrt{\tau}\Theta^{-1}(\delta)) \tag{4.5}$$

where $\Theta^{-1}$ is the inverse of the cumulative density function of a standardised *Normal* distribution and $\sigma_{ATM}^2$ is the implied volatility at-the-money, e.g. the volatility with delta equal to 0.5. The implied volatility values of the spline can easily be translated back into call prices using the Black-Scholes (1976) option pricing equation 4.1.

#### 4.1.1.4 **Option-implied probability density functions**

The Breeden and Litzenberger (1978) result described in 1.6 is finally applied to the 1,000 premium (strike, prices). In order to calculate the second derivative of the call function, cubic polynomials are fitted through triplets of consecutive *(strike price, call price)* pairs; from the coefficients of the fitted polynomials the second derivative is evaluated, which allows deriving the option-implied PDF.

## 4.2 Fitting option-implied probability density functions for constant maturity contracts

A total of eight option contracts on the three-month Euribor futures are traded daily on NYSE LIFFE. Each of them expires on the same day as the underlying future contract cycle of March, June, September or December. As each option contract gets closer to the expiry date, the uncertainty about possible future Euribor outcomes declines. Therefore, the amount of uncertainty embodied in the option-implied PDF also tends to decline as we approach the expiry date. Moreover, very little trading, if any, typically takes place on the days immediately prior to the expiry date. This regular time-to-maturity feature makes it very difficult to compare option-implied probability statistics on the same fixed expiry contract over time.

A solution to this time pattern is to estimate constant maturity option-implied PDFs interpolating over the eight fixed expiry option-implied PDFs calculated in 4.1.1. Based on this interpolation, the three-month, six-month, nine-month, one-year and one-year and six-months constant maturity contracts are calculated. For any given day, each of these option-implied probability densities always represents the same constant period ahead.

The method does not interpolate directly over the option-implied PDFs but over the implied volatility curves with the same delta, but with different maturities that were calculated in 4.1.1.3. The advantage of this method is that the same delta, but for contracts expiring on different dates, is always defined by the non-parametric technique. In addition, by construction the delta always ranges between 0 and 1.

### 4.2.1 Constant maturity option-implied probability density functions

In detail, to construct the constant maturity option-implied PDFs a vector containing the nine delta values from 0.1 to 0.9, with a step-width of 0.1, is first created for every fixed expiry contract. For each delta in this vector, the value of the corresponding sigma is then calculated by evaluating the previously estimated volatility smiles. This is done by using the grid of 1,000 two-component points defined in 4.1.1.3, where the first coordinate is the delta and the second is the sigma. From this grid, the nine sigmas are calculated using linear interpolation. Figure 4.3 presents the initial nine delta and sigma values used to calculate the constant maturity option-implied PDFs.

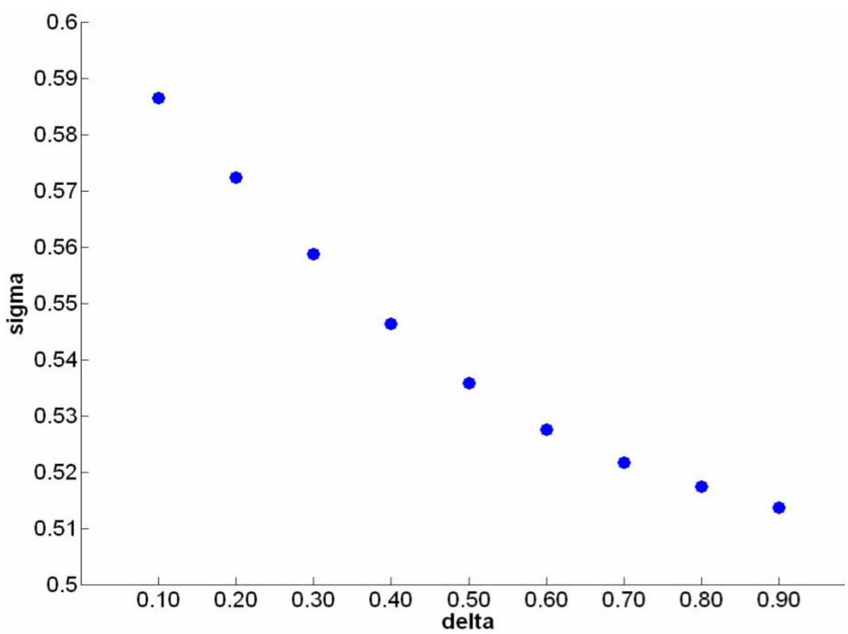


**Figure 4.3:** Nine points delta-sigma space for the December 2009 contract; 27 October 2009.

This procedure is repeated for each of the fixed expiry contracts. Hence, for each of the nine deltas, the value of sigmas for different times to maturity, e.g. the fixed expiry ones, are calculated for all the fixed expiry contracts. Figure 4.4 below presents the initial nine delta and sigma values used to calculate the constant maturity option-implied PDFs for all fixed expiry contracts.

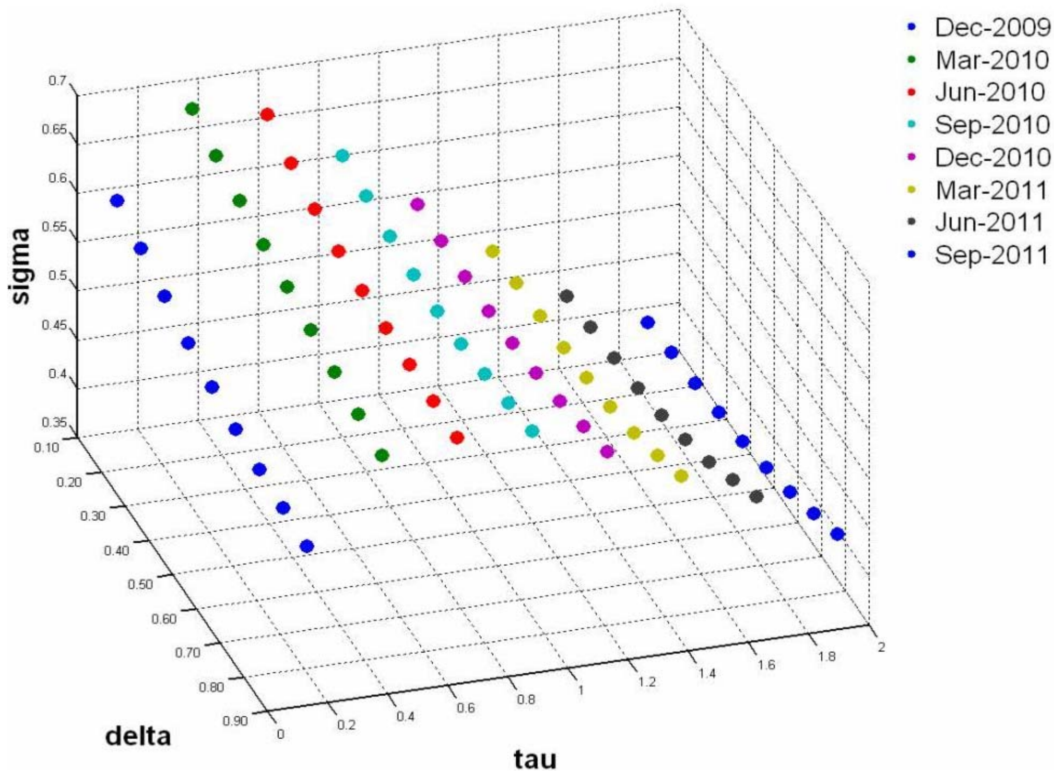


**Figure 4.4:** Nine points delta-sigma space for the all fixed expiry contracts; 27 October 2009.

For each of the nine deltas, nine smoothing splines are constructed by interpolating the sigmas of all fixed expiry contracts with the same delta. For instance, a smoothing spline is calculated by interpolating all the sigmas of all the fixed expiry contracts with delta equal to 0.1. Figure 4.5 below shows the smoothing spline for deltas equal to 0.1 and 0.2.

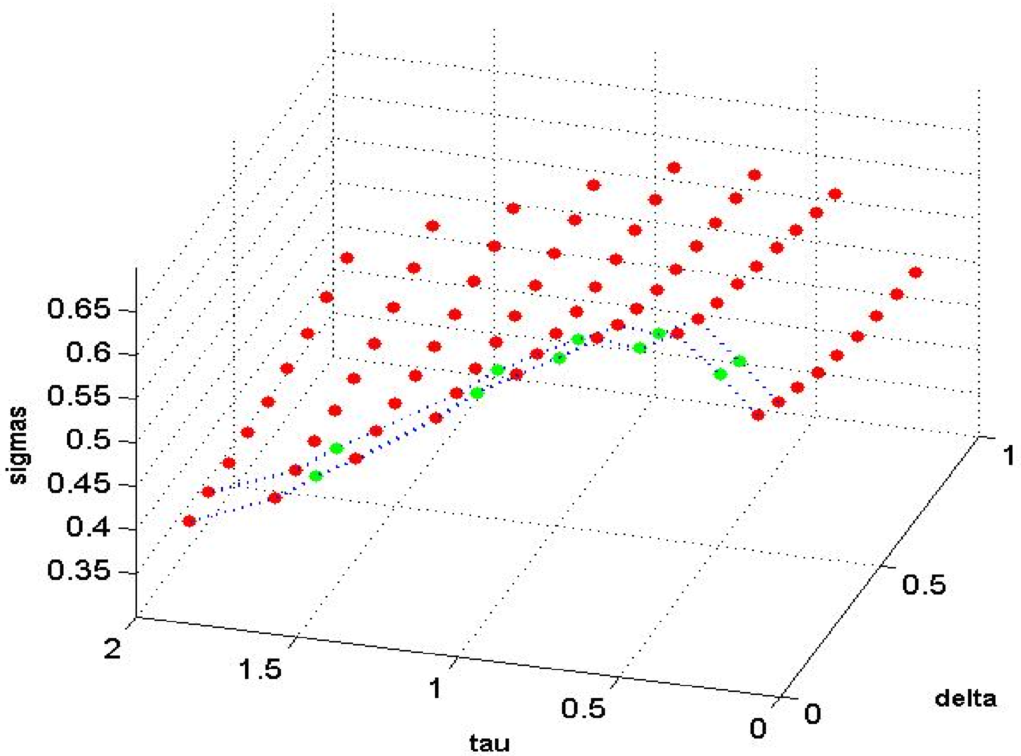


**Figure 4.5:** Splines at delta equal 0.1 and delta equal 0.2.

Each of the 9 splines is evaluated at the constant maturity values, e.g. 3-months, 6-months, 9-months, in order to get the corresponding sigmas at these points. The green dots in Figure 4.5 are, in fact, the sigmas of deltas equal to 0.1 and 0.2 for the constant horizon contracts. All the required data are derived for each constant horizon, e.g. nine deltas, nine sigmas, tau, risk-free interest rate, and the underlying value, which is obtained by interpolating the two closest underlying contracts with a smoothing spline.

Later on, the deltas are converted into strikes, and the premium of every artificially-created option is calculated using the Black-Scholes (1973) model. Then, the non-parametric model is used again to calculate the option-implied PDFs, as defined in 4.1.1. Summarizing, by repeating the same process described in Section 4.1.1, the exactly at-the-money implied volatility is calculated, a 1,000-point delta grid is generated, a 1,000-point sigma grid is calculated using splines, then the deltas are transformed back into the strike space to calculate the premium, and finally the constant maturity option-implied PDF is calculated.

## 4.3 Risk-neutral and real world option-implied probability density functions

### 4.3.1 Risk neutrality assumption

Option-implied PDFs estimated under a Black-Scholes (1973) option-pricing derivation, such as those defined in the previous section, are by construction risk-neutral. The option-implied PDFs represent the set of probabilities under which the expectation of the terminal asset price must be discounted by the risk-free rate, in order to equate with the market price. Such option-implied PDFs correspond to the probabilities that risk-neutral investor would have, but the agents that price the options might in fact be risk averse. If that were the case, risk premia would lead to differences in both the location and shape of the risk-neutral and actual distributions. The extent of such differences is likely to vary with both asset class and maturity.

### 4.3.2 Real world option-implied probability density functions

Different techniques can be used to transform the option-implied PDFs into estimates of the actual distribution. Bliss and Panigirtzoglou (2004) and Alonso *et al.* (2009) exploit the fact that the risk-neutral and actual distributions are related to each other via the marginal rate of substitution of the representative investor to define the functional form of the transformation. They estimate then the parameters of that transformation function for different assumed forms of the utility function by maximising the forecasting ability of the transformed PDFs. In contrast, Liu (2007) *et al.*, following Fackler and King (1990), define their transformation in terms of the beta function. The additional flexibility of the beta function might better align the transformed PDFs with the pattern of past outturns, but perhaps at the cost of economic insight. The real-world PDFs techniques are described and used in detail in the third article of this thesis.

*It is a hypothesis that the sun will rise tomorrow: and this means that we do not know whether it will rise.*

Ludwig Wittgenstein

CHAPTER

5

# Hypotheses and aims

## 5.1 Hypotheses

The work that has been conducted in this thesis can be summarised in the following three main hypotheses:

- A number of methods for constructing option-implied probability density functions (PDFs) have already been developed in the literature. Although these methods might differ in the extreme values of the tails of the distribution, there is generally, when using and comparing these methods, no major difference in the central section of the estimated PDF which makes these methods suitable for monetary policy purposes.

- Option-implied PDFs have never been used with intraday data and they have never been used during periods of extended crisis or to compare long periods of prolonged stability against long periods of economic crisis or financial turbulences. Additionally, the market's quantitative assessment of the risks around inter-banks rates around central bank monetary policy decisions may be studied using these types of functions.

- Option-implied probability densities may also be used for forecasting purposes to provide estimates of the interest rate risk premium in the futures market and to compare those estimates against historical time series of asset prices. Additionally, the estimation of real-world option implied PDFs may allow us to study and compare the investors' risk behaviour. In this respect, the state price densities, as the ratio of the risk neutral and real world densities, may allow us to analyse whether investors are more risk neutral or rather risk-averse to a possible increase or decrease of interest rates during times of financial stability or during periods of financial crisis.

## 5.2 Aims

The general aims of the present work can be summarised as follows:

- To evaluate option-implied PDFs for monetary policy purposes during periods of extended crisis and periods of financial stability.
- To evaluate option-implied PDFs around monetary policy decisions and in particular when using intraday data.
- To assess the adequacy of option-implied PDFs to calculate forecasts of interest rate risk premium and analyse whether state price densities can be also used to analyse investors' risk behaviour.

In article 1:

- To first estimate the complete history of daily data for the Euribor market, from 1999 until the most recent available data, by using daily option-implied PDFs on the 3-month Euribor futures.
- In this respect, to compare the statistical moments of different economic periods and to give particular attention to how these types of statistics reacted during the financial crisis in the period between 2007-2009.

In article 2:

- To better understand, based on case studies, the reaction of market participants around ECB Governing Council decisions by analysing option-implied PDFs at intraday frequency.

In article 3:

- To estimate real-world implied PDFs at daily frequency in order to take into account the investors' risk behaviour.
- To evaluate whether risk-neutral and real-world option-implied PDFs can be used to forecast estimates of the risk premium component in the future.
- To make use of state price densities to evaluate the risk behaviour of the investors in different states of the economy.

# Part II

# Contributions

*If we knew what it was we were doing, it would not be called research, would it?*

Albert Einstein

CHAPTER

# 6

# Summary of articles

The three articles presented in this chapter, which constitute the basis of this thesis, follow a line of research to analyse the development of option-implied probability density functions (PDFs) of 3-months Euribor futures. In this regard, the first article presents the results of using daily risk-neutral option implied PDFs on more than 10 years of data, the second one analyses risk-neutral option implied PDFs by using tick-by-tick data around ECB Governing Council Decisions, and the last one shows the results of using daily real-world option implied PDFs to calculate forecasts of the 3-months Euribor futures. The three articles use the methodology that has been described in Chapter 4.

The following sections included in this chapter present a detailed summary of the 3 articles that constitute this thesis. Each section of this chapter is devoted to a different article and sets out as well the motivation and the main research findings. However, a complete version of these 3 articles is included in the Annex of this thesis (see Chapter 9).

## 6.1 Article 1: Daily risk-neutral option-implied probability density functions

The first article, entitled *A quantitative mirror on the Euribor market using implied probability density functions* (Annex Section 9.1) presents an analysis of daily PDFs based on 3-months option Euribor futures, from 13 January 1999, when options on Euribor futures first started trading, until April 2010. By including different periods of prolonged stability as well as period of turbulence, the article provides a comprehensive picture of the market's quantitative assessment of the risks around interbank rates in the future.

### 6.1.1 Motivation

In the past several authors have used option-implied PDFs to analyse market uncertainty. For instance, Blanco *et al.* (2009) to try to analyse the value of information contained in prices of options on the IBEX 35 index at the Spanish Stock Exchange Market; Melick and Thomas (1997) to estimate the distribution for crude oil in particular during the Persian Gulf crisis. Mandler (2002) also focuses on the 3-months Euribor, however only during the period 1999-2000, when the economy was growing and there were no major market fluctuations; Anagnou-Basioudis *et al.* (2005) use option-implied PDFs for both currency and index future contracts, and Bliss and Panigirtzoglou (2002) estimate option-implied PDFs using short sterling futures options and the FTSE 100 index options.

The previous literature mainly focused in the equity and commodity derivatives and very little in short-term interest rates, which was the purpose of our study. Moreover, the option-implied PDFs had not been used neither in an extended period of crisis, when they would have appeared to be most useful, nor for comparing periods of economic stability with those with financial market crisis. In this respect, this first paper can be seen as a reference to monetary policy makers, economic practitioners and researchers on the analysis of the uncertainty of the 3-months Euribor rates, allowing to identify the reaction of option-implied PDFs over a long time period which includes both episodes of prolonged stability and episodes of turbulence. The evolution of market interest rates is a key component of the transmission of monetary policy. In this regard, the paper provides a comprehensive, quantitative assessment of the option-implied PDFs, which may be a natural complement to the wide range of financial market indicators already considered by monetary policy makers.

Although the methodology relies on a previous result presented by Bliss and Panigirtzoglou (2002) and Cooper (1999), the paper further specifies and elaborates the technicalities of this methodology. In particular, it analyses the trading volume for this ten-year data period by option type, maturity date, year and moneyness type. The study of the trading volume based on these categories may allow to understand better the interest rate option market and to select the most liquid strikes, avoiding the possible bias that could be introduced to the option-implied PDF if the whole set of option strikes were selected.

### 6.1.2 Summary of results

The first paper illustrates how information from option-implied PDFs can be used to inform and add value to economic analysis. In this respect, it analyses how option-implied PDFs covers the whole spam between 1999 and 2010. Figure 6.1 shows the percentile evolution of the 3-month constant maturity PDFs over time. The figure shows that in periods of financial turmoil, in particular after 2008, there is not only a bigger standard deviation, but also the skewness becomes more positive in comparison to other periods, in particular between 2005 and 2007, where the skewness is close to zero.

Moreover, the first paper also describes in detail how Euribor PDFs reacted to the unfolding financial crisis between 2007 and 2009. For the latter period, two important events that occurred in 2008 are described. In doing so, it demonstrates how option-implied PDFs can provide timely and quantitative indicators not only of the amount of uncertainty around forward Euribor rates, but also of the directional bias of this uncertainty.

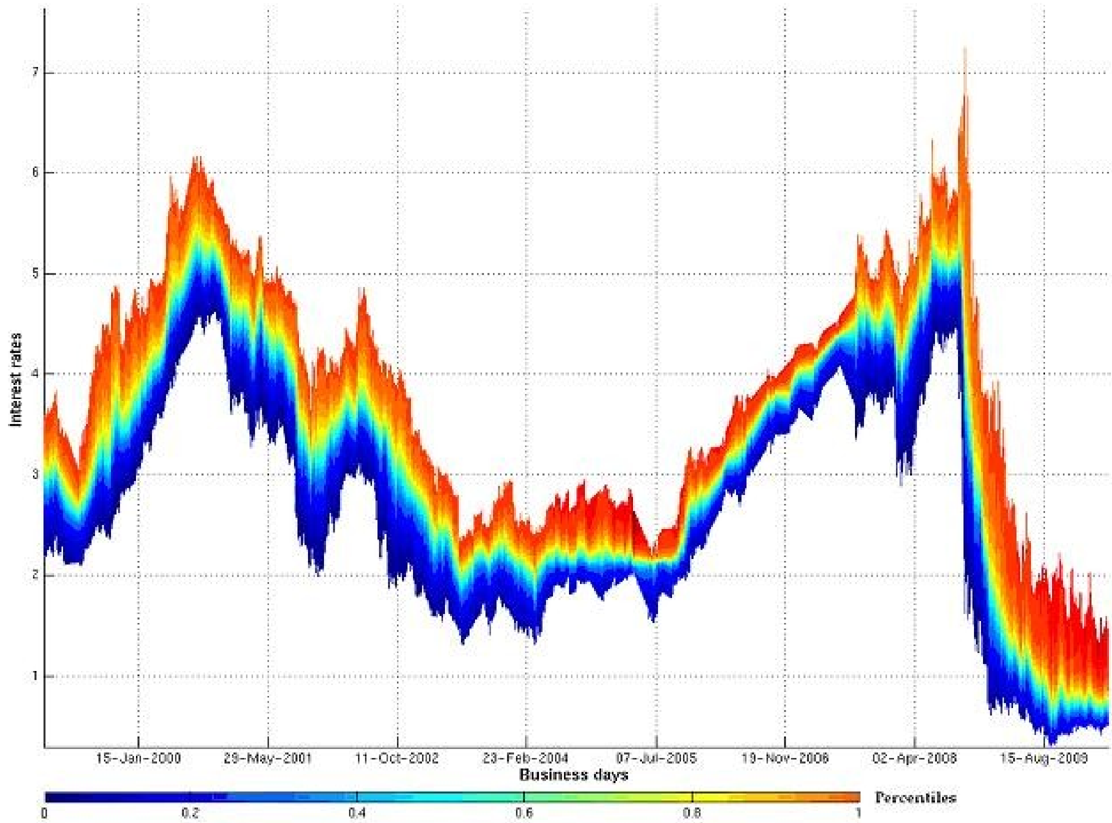


**Figure 6.1:** Projection in the density plane of the three dimension option-implied probability density function.

The autumn of 2008 was especially tumultuous, but two events stand out: the failure of Lehman Brothers on 15 September and the coordinated policy interest rate reductions of the Bank of Canada, the Bank of England, the European Central Bank (ECB), the Federal Reserve, the Sveriges Riksbank and the Swiss National Bank, strongly supported by the Bank of Japan. The failure of Lehman Brothers led to material changes in the three-month-ahead Euribor distribution (see Figure 6.2). While there had been little movement in the PDF in the preceding week, Euribor option prices assigned a significantly greater weight to interest rate outturns much less than the prevailing forward rate. And that left-tail continued to grow.

Stress in the cash markets increased markedly too and the spread between forward Euribor and EONIA increased. But while it could also be argued that the large negative skew reflects in part the view that the Euribor-EONIA spread could be much narrower than the forward spread, the sheer magnitude of the left tail suggests that it is also likely to reflect beliefs about future policy rates.

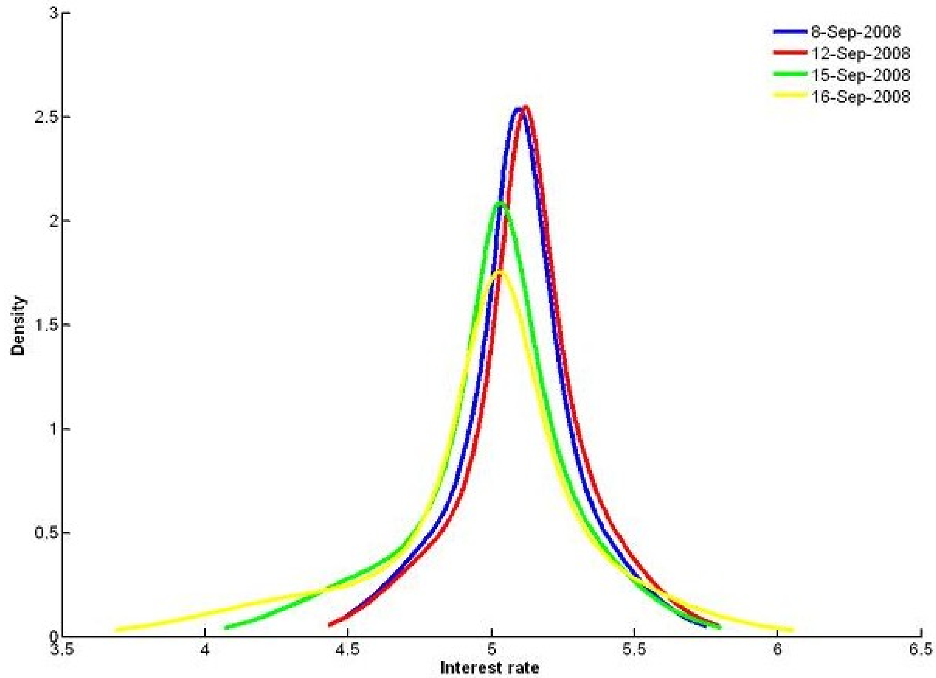


**Figure 6.2:** Three-month constant maturity Euribor option-implied probability density functions before and after the failure of Lehman Brothers on 15 September 2008.

On 8 October, as part of internationally-coordinated monetary policy action, the ECB announced that, from the operation settled on 15 October, the weekly main refinancing operations would be carried out through a fixed rate tender procedure with full allotment at the interest rate on the main refinancing operation, i.e. at that moment 3.75%. That rate is 50 basis points below the minimum bid rate agreed at the previous Governing Council meeting on 2 October. An examination of the Euribor option-implied PDF in the days leading up to that announcement and shortly afterwards reveals two interesting observations (see Figure 6.3). First, it appears that the impact of both the 2 October Press Conference or the 8 October announcement on the option-implied Euribor distribution was small compared to that of the accumulation of news during the intervening days (in particular, over the weekend). The fact that even by 7 October the option-implied PDF had shifted so much to the left, and become more negatively skewed, suggests that market participants were already placing more weight on Euribor outturns in three months' time being much less than the current forward rate, even though the precise timing and details of the 8 October announcement took the market by surprise. The second interest observation is that the although the bulk implied three-month Euribor distribution continued to move towards lower interest rates, there was no movement in the tail of the distribution. One possible explanation is that, despite the unprecedented events of the preceding month, market participants still did not attach any weight to the possibility that Euribor would be 2% or less in three months' time.

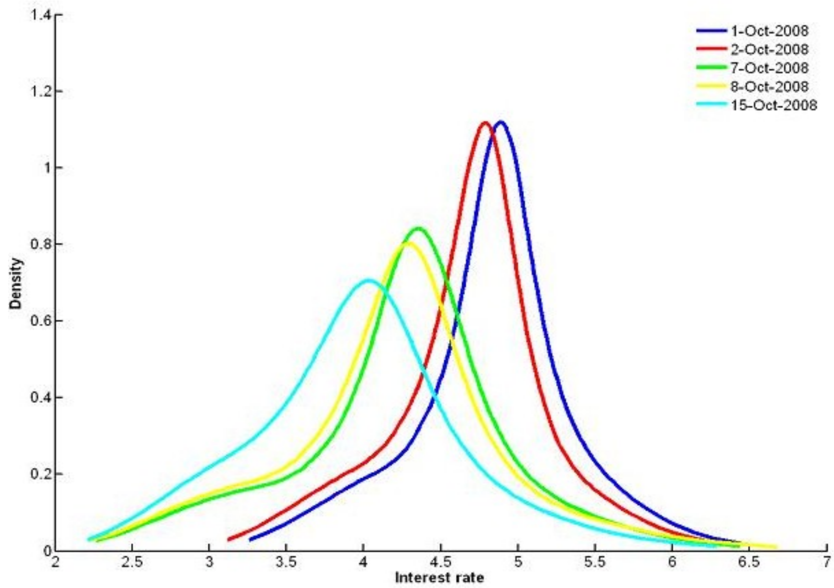


**Figure 6.3:** Three-month constant maturity Euribor option-implied probability density functions before and after the change in monetary policy on 8 October 2008.

Regarding the analyses of volumes, the paper shows that in absolute terms, 81% of the options are traded out-of-the-money, whereas only 18% are traded in-the-money. Furthermore, some of the in-the-money options are traded not independently, but as part of a bundled trading strategy, e.g. straddles or strangles, which combine options out-of-the-money with options in-the-money. With this confirmation, the analysis in the first article was made by using those option prices which were either at- or out-of-the money, but not in-the-money. Moreover, trading was much more concentrated in the options contracts maturing in nine months or less. These accounted for more than 85% of the total trading. For those option contracts with longer maturities, whose trading was very seldom, perhaps even being untraded for some days, the settlement prices were directly assigned by LIFFE. Finally, and as may be expected, Figure 6.4 the number of traded contracts has increased steadily since this instrument was first introduced, with most trading in the most recent years.

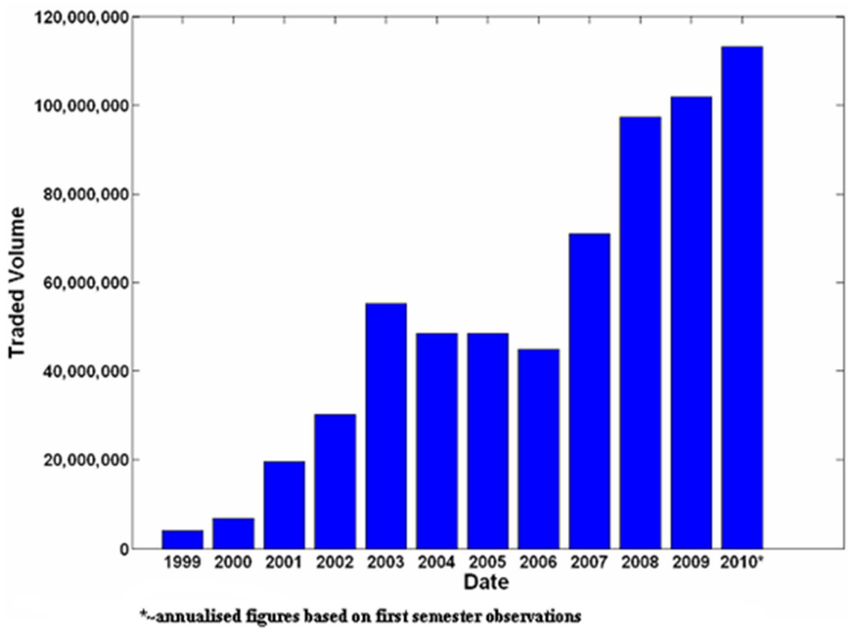


**Figure 6.4:** Total traded volume for all contracts per year.

### 6.1.3 Caveat on the risk-neutrality assumption

Option-implied PDFs estimated under a Black-Scholes (1973) option-pricing derivation, which are the ones used in the first article, are risk-neural by construction. The option-implied PDFs represent the set of probabilities under which the expectation of the terminal asset price must be discounted by the risk-free rate, in order to equate with the market price. This is because as previously mentioned such PDFs correspond to the probabilities risk-neutral investor would have, but the agents that price the options might in fact be risk-averse. If that were the case, risk premia would lead to differences in both the location and shape of the risk-neutral and actual distributions. The extent of such differences is likely to vary with both asset class and maturity.

Different techniques can be used to transform the risk-neutral PDFs by options into estimates of the actual distribution. Estimating the actual PDFs from the risk-neutral PDFs would merit further research. This is, in fact, the scope of the third article. The first article instead focuses directly on the option-implied PDFs themselves.

### 6.1.4 General assessment

The first article shows that option-implied PDFs applied to the Euribor provide a timely and quantitative indication of the market's assessment of the risks around the forward Euribor: not just how much uncertainty there is, but precisely how that is distributed over different possible outturns. These can be used to analyse trends such as the extent to which the balance of risks is skewed to the upside or the downside, or to analyse how specific events affected the entire spectrum of views.

This indicator may appeal to those interested in monetary policy or financial stability. Moreover, such a comprehensive dataset, covering the complete history of the euro, is particularly valuable because it gives the context against which the current situation, or recent developments, may be compared, and provides a benchmark to help judge whether the current situation is 'normal' or 'extreme'.

## 6.2 Article 2: Intraday risk-neutral option-implied probability density functions

The second article entitled *Interest rate expectations and uncertainty during ECB Governing Council days: Evidence from intraday implied densities of 3-month Euribor* (Annex Section 9.2) further extends the frequency of the analysis of the option-implied PDFs. In particular, this second article focuses on the study of short-term interest rates expectations and uncertainty up to tick frequency, and aims to identify the drivers of the market reactions during ECB Governing Council days.

### 6.2.1 Motivation

Central bank communication receives widespread attention by financial market participants. Reaction 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, as suggested by Amato *et al.* (2002) or ECB (2007).

When analysing central bank communication, the first challenge is to define appropriate measures and determine the relevant indicators to look at. A second challenge is to determine the factors driving the markets expectations and uncertainty. In this regard, several authors have tried to study these two challenges which are fundamental to central bank communication. For instance, Brand *et al.* (2010) analyse high-frequency changes in the euro area money market yield curve on dates when the ECB regularly sets and communicates decisions on policy interest rates to construct different indicators on monetary policy news relating to policy decisions and to central bank communication. Indicators based on the ECB yield curve show that ECB communications during at press conference may result in significant changes in market expectations of the path of monetary policy. Similarly to Brand *et al.* (2010), Castren (2004) focuses on changes in the currency options market's assessment of likely future exchange rate developments around the times of official interventions in the JPY/USD exchange rate. By using daily option-implied risk-neutral density functions, similarly to the previous article, it is concluded that episodes of interventions on the JPY/USD exchange rate coincide with systematic changes in all moments of the estimated option-implied risk-neutral density functions on the JPY/USD currency pair, and in several of the moments of the estimated option-implied risk-neutral density functions on the JPY/EUR and USD/EUR currency pairs. Ehrmann and Fratzcher (2009) analyse the question of how to best communicate monetary policy decisions, which remains an important topic among central banks. Focusing on the experience of the European Central Bank, the authors study how explanations of monetary policy decisions at press conferences are perceived by financial markets. The empirical findings show that ECB press conferences provide substantial additional information to financial markets, beyond that contained in the monetary policy decisions, and that the information content is closely linked to the characteristics of the decisions. Andersen and Wagener (2002) analyse the impact of the 11 September attacks on the expectations of future Euribor interest rates around ECB Governing Council Decisions.

Following the research of these authors, 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. Clearly the benefit of the approach is that - as few case studies show - one can zoom in on certain events and judge the immediate market reactions, thereby minimizing the bias from any other information hitting the market.

### 6.2.2 Summary of results

The analysis is based on expectations three months ahead about a money market interest rate, namely the 3-month Euribor. The option-implied PDFs are computed based on the same non-parametric estimator which was used in the first article, but applied to tick data on three-month Euribor futures and options. The results show that the estimator is robust to market microstructure noise and produces stable risk-neutral option-implied PDFs. At the same time, when information hits the market, the option-implied PDFs 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.

Overall, the relevance of the press release and conference as communication tools 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 practised by the ECB.

Moreover, 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 option-implied PDFs show a significant shock in activity following the press release and significantly increased activity during the press conference. 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.

Finally, by using a regression analysis the paper 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 for driving expectations.

### 6.2.3 General assessment

The second article shows that option-implied PDFs can be also extended to tick frequency data. Moreover, the results obtained by using tick data options can be used to detect changes in the statistical moments of the option-implied PDFs better than when using daily option-implied PDFs.

In particular, this indicator appears to be very useful when applied around ECB Governing Council decisions or during ECB communication on short-term interest rate expectations.

## 6.3 Article 3: Real world option-implied probability density functions and interest rate forecasts

The third article entitled *Interest rate forecasts, state price densities and risk premium from Euribor options* (Annex Section 9.3), largely focuses on daily real world option-implied PDFs.

The real world option-implied PDFs incorporate the risk premium component and are adjusted for the individual risk profile of the different investors regardless of whether they are risk-neutral or risk-averse. They are constructed from the risk-neutral density functions described in the first article by using two types of transformations: a beta statistical transformation as described by Fackler and King (1990) and a non-parametric statistical transformation following the approach introduced by Shackleton *et al.* (2010).

The article further studies and compares the interest rate forecasts obtained by using risk-neutral and real-world density functions and the development of state price densities, which are by definition constructed as a ratio of the two density functions.

### 6.3.1 Motivation

The first and second articles discussed in the previous sections try to estimate the option-implied PDFs by assuming that investors are risk-neutral. This assumption has the major disadvantage that it does not incorporate the risk premium component and does not correspond to investors' actual sentiment. For this reason, we try to explore and analyse in a third article the evolution of the daily real world option-implied PDFs over different crisis periods.

In addition, we also go one step further in this article and compare the forecasting power with respect to the 3-months Euribor futures of the risk-neutral and real-world PDFs. The forecasting ability of option prices has already been studied by some authors, who have contributed to the existing literature. However, the authors have studied those forecasts in the equity and index markets, but not in the interest rate market, which is the scope of our study. In this respect, Jiang and Tian (2005) and Martens and Zein (2004) show that option forecasts of index volatility can be more useful than historical forecasts. Shackleton *et al.* (2010) and Liu *et al.* (2007) show also similar results after using implied probability density forecasts.

Moreover, the calculation of the real world option-implied PDFs may also be used to compute state price densities, as a ratio between the real world and risk-neutral PDFs. These may be used to analyse the different states of the markets on different periods of time.

### 6.3.2 Summary of results

By investigating the period from the introduction of the Euro in 1999 until December 2012 we conclude that real world option-implied PDFs can be used to forecast the 3-months Euribor futures rate. However, we have not enough statistical evidence to come to the same conclusion regarding the forecasting power for the risk-neutral option-implied PDFs.

Similar to the results described by Liu *et al.* (2007) we also find that the state price densities in the market show a U-shape curve. This shape suggests that investors give in particular high price to states with high rates but also, although to a lesser extent, to states with low rates compared to the expected spot rate. The periods where there is not a high fluctuation in the markets and rates are high, a further increase in the rates is perceived as a bad state of inflation. This may indicate that investors are in fact more risk-averse to an increase of interest rates when rates are already relatively high.

### 6.3.3 General assessment

The last paper of this thesis shows that there is a benefit in further transforming the risk-neutral option-implied PDFs into real-world ones. In this respect, the real-world option-implied PDFs can be used not only to forecast 3-months Euribor futures rates but also to analyse the different state price densities for different periods of data.

However, one needs to take into account that real-world option-implied PDFs need a larger amount of data to be constructed in comparison to the risk-neutral option-implied PDFs. In this regard, in order to avoid overlapping, information monthly maturity is used, and not 3-months or 6-months like in the previous papers. This very much limits the forecasts and the study of the uncertainty to the near future. A detailed description on how real-world option-implied density functions have been constructed can be found in the paper.

# Part III

# Discussion

*If all the economist were laid end to end, they would never reach a conclusion.*

George Bernard Shaw

CHAPTER

7

# Conclusions

- There is no major difference between existing option-implied probability density (PDF) models in the central section of the distribution.

- Different option-implied PDFs models might differ in the tails of the distribution, i.e. outside the range between the 0.1 and 0.9 percentiles. In this regard, option-implied PDFs models are better fit for monetary policy rather than for financial stability analysis.

- Option-implied PDFs models, which rely on a computational form, might have convergence problems depending on the initial set of parameters used.

- The option volume analysis shows that options that expire out-of-the-money account for 81% of the total volume compared to options which expire either at-the-money or in-the-money. In addition, trading is much more concentrated in options with contracts maturing in nine months or less, i.e. short-term contracts, than in contracts maturing in more than nine months and up to two-years, i.e. long-term contracts.

- Option-implied PDFs provide an easily-accessible tool for visualizing how market participants react to the relevance of the European Central Bank press release and, in particular, during periods of extended crisis. Additionally, they can be used to compare long periods of prolonged stability against long periods of economic crisis or financial turbulence.

- Intraday option-implied PDFs allow to better capture the uncertainty attached by market participants to a given day before and after European Central Bank Governing Council decisions.

- The option-implied PDFs, which are by definition risk-neutral, cannot be used to forecast possible outcomes of the 3-month Euribor futures rates. However, the transformation of the risk-neutral PDFs into real-world PDFs can be used to forecast 3-month Euribor futures rates.

- The analysis of the ratio between the risk-neutral and the real-world PDFs, i.e. the state price densities, suggest that investors price higher states with high and low rates compared to the expected spot rate.

*The best thing about the future is that it comes one day at a time.*

Abraham Lincoln

CHAPTER

# 8

# Future research

The option-implied probability density functions (PDFs) analysed in this thesis allow us to study, in particular, the reactions of market participants to the 3-month Euribor interest rate around ECB Governing Council decisions. This is largely due to the important role that the money market plays in the monetary policy transmission mechanism, since changes in monetary policy instruments affect, at first instance, the money market, and the 3-month Euribor is together with the EONIA one of the two important reference rates for the unsecured market.

The Treaty on the Functioning of the European Union (see for instance The Monetary Policy of the ECB), clearly establishes the maintenance of price stability as the primary objective of the Eurosystem. In this sense, inflation rates should be maintained below, but close to, 2% over the medium term. Since the future reactions of inflation expectations are also important for monetary policy implementation, one could also study inflation option-implied PDFs as a complement to the 3-month Euribor ones.

This type of market analysis would allow monetary policy authorities to assess directly uncertainty about future inflation. As pointed out by Bahra (1997), the most important limitation from the point of view of monetary policy authorities is that, in the past, there were no markets where options on inflation could be traded. However, this is no longer the case and there are currently markets which options on inflation expectations are being traded.

Some authors have recently used inflation options to construct option-implied PDFs for inflation and studied how these functions responded to certain announcements. In this regard, Kitsul and Wright (2013) have recently used inflation options to construct implied PDFs in the US market. By using this type of options, the authors analyse uncertainty around inflation rates in the future. Smith (2012) has also used these inflation options for the UK market. However, there is no research published using this inflation options in the euro area.

## 8.1 Estimating inflation option-implied probability density functions

Inflation options need to be constructed using inflation cap options or inflation floor options. An inflation cap option is a financial asset that hedges against inflation being higher than a given percentage rate of inflation. Inflation cap options are therefore used by investors to guarantee a maximum level of inflation. Similarly, inflation floor options are used to hedge against downside risks to inflation.

The structure of inflation options is more complex than that of the the 3-month Euribor options. In this regard, an inflation cap (floor) is bundled to a series of consecutive options, which are called caplets (floorlets), and with a maturity of one year. Each caplet is related to the same rate of inflation (the option strike). For example, a 4-year inflation cap on 1 September 2014 would be bundled to 4 series of caplets of one year maturity expiring respectively on 1 September 2015, 1 September 2016, 1 September 2017 and 1 September 2018. In the euro area, the option strike price ranges between -2% and 5%. The important caveat about this type of options is that liquidity remains still very limited and concentrated only on a few option strikes on each trading day.

An inflation caplet works in a similar way to that of a 3-month Euribor call option or an interest rate cap: the buyer pays the seller a premium upfront (the option price) and, in exchange, the seller pays the buyer the difference between actual inflation in a given period (e.g. one year in the case of a year-on-year option) and a pre-specified rate of inflation (the strike rate) multiplied by the notional amount if the actual inflation rate is higher than the strike rate. Hence, inflation options offer protection against inflation being higher than the strike rate. Similarly, a floorlet works in the same way if inflation is lower than the strike rate.

Formally speaking, an inflation cap can be described by the sum of the inflation caplets:

$$Cap(t, t_N, K) = \sum_{t_k=t_1}^{t_N} Caplet(t, t_k, K) \tag{8.1}$$

where t is the current point in time, $t_N$ is the maturity of the cap, $t_k$ is the settlement time of the individual caplets and K is the strike price.

Kruse (2011), by inhereting the assumptions of the Black-Scholes (1973) model, such as, for instance, the assumption of constant volatility over time and all strikes, describes how inflation caplets can be evaluated.

**Proposition 1.** ***Kruse***
*Assuming that the dynamics of the inflation index are given by a Geometric Brownian motion, the price of a caplet on the inflation rate $i(T_i - 1, T_i)$ over a future interval from time $T_i - 1$ up to the maturity of the caplet at time $T_i$ with $0 < T_i - 1 < T_i$ and payoff at time $T_i$*

$$C_{FIR}(T_i; i(T_i; T_i - 1)) = (i(T_i - 1, T_i) - k)^+ = (\frac{I(T_i) - I(T_i - 1)}{I(T_i - 1)} - k)^+ \tag{8.2}$$

*is at time* $t = 0$ *given by*

$$C_{FIR}(0, I(0)) = e^{-r_R(T_i - T_i - 1)} e^{-r_N T_i - 1} N(d) - (1+k) e^{-r_N T_i} N(d - \sigma_I \sqrt{T_i - T_{i-1}}) \tag{8.3}$$

*with*

$$d = \frac{-ln(1+k) + (r_N - r_R + \frac{1}{2}\sigma_I^2)(T_i - T_{i-1})}{\sigma_I \sqrt{T_i - T_{i-1}}} \tag{8.4}$$

where $r_N$ is the constant continuous nominal interest rate, $r_R$ the constant continuous real interest rate, $\sigma_I$ is the volatility of the inflation index under the geometric Brownian motion assumption and $I(t)$ the inflation index at time $t$ and the inflation rate $i$ is given by

$$i(t, T) = \frac{I(T) - I(t)}{I(t)}. \tag{8.5}$$

As noted by Kruse (2014), the model is rather simplistic and easy to use for practitioners, central bankers and researchers. Intuitively, the price of an option on the inflation rate over a future time interval is dependent on today's market position on the future behaviour of inflation. Using the formula above and the observed implied inflation option volatilities, the option-implied PDFs can be derived by calculating the implied volatility smile as described in Chapter 4.

# Part IV

# Annexes

*It is well enough that people of the nation do not understand our banking and monetary system, for if they did, I believe there would be a revolution before tomorrow morning.*

Henry Ford

CHAPTER

9

# Published Articles

## 9.1 Article 1: A quantitative mirror on the Euribor Market Using Implied Probability Density Functions

*Eurasian Economic Review, 2(1), 2012, 1-31*

# A QUANTITATIVE MIRROR ON THE EURIBOR MARKET USING IMPLIED PROBABILITY DENSITY FUNCTIONS†

**Josep Maria Puigvert-Gutiérrez* and Rupert de Vincent-Humphreys****

*Abstract:* This paper presents a set of probability density functions for EURIBOR outturns in three months' time, estimated from the prices of options on EURIBOR futures. It is the first official and freely available dataset to span the complete history of EURIBOR futures options, thus comprising over ten years of daily data, from 13 January 1999 onwards. Time series of the statistical moments of these option-implied probability density functions are documented until April 2010. Particular attention is given to how these probability density functions, and their associated summary statistics, reacted to the unfolding financial crisis between 2007 and 2009. The latter shows how option-implied probability density functions can be used as an uncertainty measure for monetary policy and financial stability analysis purposes.

*Keywords:* Financial Market, Probability Density Functions, Options, Financial Crisis
*JEL Classification:* C13, C14, G12, G13

## 1. Introduction

Forward interest rates reflect the market's aggregate risk-neutral expectation of spot interest rates in the future.[1] From the prices of options on interest rate futures, it is possible to construct the entire probability distribution for the interest rate in the future. And because that probability distribution describes the (risk-neutral) likelihood the market ascribes to all possible outcomes, it provides a quantitative measure of the market's assessment of the risks around the forward rate, in terms of both magnitude and directional bias. Therefore, such option-implied probability density functions (PDFs) constitute a natural complement to the many other

† We would like to thank Björn Fischer, Juan Ángel García, Jean-Marc Israël, Holger Neuhaus, Jochen Schmitz, Olivier Vergote, and two anonymous referees for methodological comments, suggestions and discussions. We are grateful to participants at the Financial Markets Statistics Task Force at the ECB for their comments. We are also grateful to NYSE Liffe for providing us with all the data since the initial trading date of the EURIBOR and to Enric Gironés, Cristian González, Nicolas González and Josep Maria Vendrell for their helpful work during the project implementation. The views expressed in this paper are those of the authors and do not necessarily represent those of the European Central Bank or of the Bank of England.

* Corresponding author: European Central Bank, Germany. Email: josep_maria.puigvert@ecb.europa.eu

** Bank of England, London, U.K. Email: rupert.devincent-humphreys@bankofengland.co.uk

[1] At short horizons, where term premia are likely to be negligible, forward rates could therefore represent a good approximation of the market's actual expectation of interest rates in the future.

financial market indicators already considered by central banks and monetary policy practitioners. For instance, the yield curve - an important component of monetary policy transmission - is influenced by how short-term rates are expected to evolve over time. Furthermore, option-implied PDFs can provide an easily-accessible tool for visualizing how the market reacts to specific events, and may thus contribute to both monetary policy and financial stability analysis.

A number of methods for constructing these option-implied PDFs have already been developed. Bliss and Panigirtzoglou (2002) have classified them into five groups for comparison: stochastic process methods, implied binomial trees, option-implied PDF approximating function methods, finite-difference methods, and implied volatility smoothing methods. To date, there has been a large discussion on the different possible methods and the differences among them. For instance, Campa *et al.* (1997) compared implied binomial trees, smoothed implied volatility smile and a mixture of lognormal methods. Coutant *et al.* (1999) compared single lognormal, mixtures of lognormals, Hermite polynomials and maximum entropy methods. In general, although these methods might differ in the very tails of the distribution, there is no major difference in the central section of the estimated PDFs. Arguably, it is the central section of the PDFs which is more likely to be useful for monetary policy purposes, in contrast to financial stability analysis, where there may be greater focus on the tails of the distribution.

This paper uses a non-parametric technique, based on the Bliss and Panigirtzoglou (2002) and the Cooper (1999) results, to estimate the option-implied PDFs. This method was preferred because, according to the previous authors it is much more stable than other techniques and avoids the possible existence of spikes in the distribution. In fact, Cooper states that by using the non-parametric technique, the small errors in the prices cause only small local errors in the estimated probability density function; while for the mixture of lognormals, the errors can be sufficient for the minimisation to reach very different parameter estimates, with large changes in the shape of the estimated probability density function. The results produced by the non-parametric technique are in general not materially different to those of other existing techniques.[2] The dataset produced by this methodology may be considered the first official EURIBOR dataset large enough for practitioners and researchers to extract useful information in support of their macroeconomic analysis.[3] In particular, such option-implied PDFs have not been studied in detail during periods of financial crisis, where arguably they may be the most useful. This paper presents an analysis of probability density functions for EURIBOR outturns in three months' time, from 13 January 1999, when

[2] A detailed comparison of the different techniques is presented by Chang and Melick in the *Proceedings of the workshop on implied PDFs* held at the BIS on 14 June 1999.
[3] Option-implied PDFs for GBP Libor since 1988 published by the Bank of England are available at: http://www.bankofengland.co.uk/statistics/impliedpdfs/index.htm.

options on EURIBOR futures first started trading, until April 2010. With more than ten years of daily data, this dataset provides a comprehensive picture of the market's quantitative assessment of the risks around interbank rates in the future. Importantly, this dataset includes periods of prolonged stability as well as periods of turbulence. The evolution of market interest rates is a key component of the transmission of monetary policy, and so such a comprehensive, quantitative assessment may be a natural complement to the wide range of financial market indicators already considered by monetary policy makers. In addition, in Section 2 we study in detail the evolution of the options on three-month EURIBOR futures trading volume. The trading volume for this ten-year data period is analysed by option type, maturity date, year and moneyness type. The study of the trading volume based on these categories allows us to select the most liquid strikes and to avoid the possible bias that could be introduced to the option-implied PDF if the whole set of option strikes was selected.

The remainder of the paper is organised as follows: Section 2 describes the type of data used and how the data are filtered. Section 3 sets out the estimation technique used to compute the option-implied probability density functions. Section 4 shows how information from option-implied probability density functions can be used to inform and add value to economic analysis. Section 5 describes in detail how the EURIBOR-implied PDFs reacted to the unfolding financial crisis between 2007 and 2009. In doing so, it demonstrates how implied PDFs can provide timely and quantitative indicators of not only the amount of uncertainty around forward EURIBOR, but also the directional bias within that. Section 6 concludes.

## 2. Data

Two types of data are required by the methodology presented in this paper to estimate probability density functions. The first is daily settlement prices for futures on the three-month EURIBOR, and the second is daily settlement prices for options on those three-month EURIBOR futures. The EURIBOR, or Euro Interbank Offered Rate, was established in 1999, after the introduction of the Euro as a daily reference rate within the Economic and Monetary Union (EMU) zone. The EURIBOR is based on the average interest rate at which banks offer to lend unsecured funds to other banks in the interbank market. EURIBOR futures, and options on EURIBOR futures are financial derivatives, traded on NYSE Liffe, whose terminal value depends on the outturn of EURIBOR. Hull (2000) explains the properties of financial futures and options in more detail. The key property of a financial option is that an option holder will only receive a payout if a certain condition is met, for example, if EURIBOR turns out to be higher or lower than a specified threshold value. The price of such a financial option will therefore embody a measure of the probability of that condition being fulfilled. This property allows probability density functions to be estimated from option prices. In the specific example presented in this paper, we estimate probability density functions for EURIBOR outturns in three months' time, from options on EURIBOR futures.

These daily settlement prices on the three-month EURIBOR futures are published by NYSE Liffe. According to NYSE Liffe, these contracts were developed in response to the economic and monetary union within Europe, and the emergence of EURIBOR as the key cash market benchmark within Europe's money markets. Since its launch, NYSE Liffe's EURIBOR contracts have come to dominate the euro-denominated short-term interest rate (STIR) derivatives market, capturing over 99% of the market share; they are now the most liquid and heavily traded euro-denominated STIR contracts in the world. Delivery months for the three-month EURIBOR futures are March, June, September and December; the last trading day is two business days prior to the third Wednesday of the delivery month, and the delivery day is the first business day after the last trading day. The Exchange Delivery Settlement Price (EDSP) is based on the European Bankers Federations' EURIBOR Offered Rate (EBF EURIBOR) for three-month euro deposits at 11 a.m. CET on the last trading day. The settlement price will be 100.000 minus the EBF EURIBOR Offered Rate rounded to three decimal places. The minimum size price movement is 0.05, which equates to EUR 12.50.[4]

Bliss and Panigirtzoglou (2002) state that out-of-the-money calls (puts) tend to be more liquid than puts (calls) of the same strike. We began by analysing the trading volume for all EURIBOR options since the first day of trading, 13 January 1999. In absolute terms, 81% of the options are traded out-of-the-money, whereas only 18% are traded in-the-money. Furthermore, some of the in-the-money options are traded not independently, but as part of a bundled trading strategy, e.g. straddles or strangles, which combine options out-of-the-money with options in-the-money. With this confirmation, we also applied our methodology to those option prices which were either at- or out-of-the money, but not in-the-money.

Table 1 shows the total trading volume (in number of transactions) per contract and per option. Trading was much more concentrated in out-of-the-money options whose contracts mature in nine months or less. These accounted for more than 85% of the total trading. For those option contracts for longer maturities, in which trading was very seldom, perhaps even going untraded some days, the settlement prices were directly assigned by NYSE Liffe. Finally, and as may be expected, the number of traded contracts has increased steadily since this instrument was first introduced, with most trading seen in the more recent years as described in Figure 1.

[4] The data are available at: http://www.liffe.com/reports/eod?item=Histories.

# Table 1. Trading volume descriptive statistics (in number of transactions)

| | CALL | | | PUT | | |
|---|---|---|---|---|---|---|
| **Panel A - Contract expiring in 3 months or less** | | | | | | |
| | At the money | In the money | Out of the money | At the money | In the money | Out of the money |
| Volume Traded | 1,165,039 | 23,522,170 | 81,332,548 | 529,526 | 8,302,374 | 34,738,846 |
| Percentage Traded | 1.1% | 22.2% | 76.7% | 1.2% | 19.1% | 79.7% |
| Maximum | 104,500 | 561,194 | 1,675,662 | 41,540 | 339,499 | 437,522 |
| Maximum Date | 11-Apr-08 | 13-Oct-08 | 9-Jan-10 | 10-Apr-08 | 25-Sep-08 | 14-Oct-08 |
| Mean | 412 | 8,318 | 28,760 | 187 | 2,936 | 12,284 |
| Std. Dev | 3,367 | 23,258 | 65,531 | 1,745 | 9,704 | 25,452 |
| Volume per option type | | 106,019,757 | | | 43,570,746 | |
| Total volume | 149,590,503 (26.7%) | | | | | |
| **Panel B - Contract expiring between 3 and 6 months** | | | | | | |
| | CALL | | | PUT | | |
| | At the money | In the money | Out of the money | At the money | In the money | Out of the money |
| Volume Traded | 720,182 | 24,323,437 | 108,736,808 | 508,490 | 12,966,668 | 57,225,845 |
| Percentage Traded | 0.5% | 18.2% | 81.3% | 0.7% | 18.3% | 80.9% |
| Maximum | 116,436 | 226,775 | 592,700 | 58,900 | 421,595 | 595,077 |
| Maximum Date | 10-Aug-07 | 7-Feb-08 | 2-Nov-09 | 22-Apr-08 | 15-Feb-07 | 23-Jun-09 |
| Mean | 255 | 8,601 | 38,450 | 180 | 4,585 | 20,235 |
| Std. Dev | 3,092 | 17,925 | 54,940 | 2,043 | 14,088 | 36,799 |
| Volume per option type | | 133,780,427 | | | 70,701,003 | |
| Total volume | 204,481,430 (36.5%) | | | | | |
| **Panel C - Contract expiring between 6 and 9 months** | | | | | | |
| | CALL | | | PUT | | |
| | At the money | In the money | Out of the money | At the money | In the money | Out of the money |
| Volume Traded | 322,074 | 13,060,952 | 66,188,484 | 330,656 | 8,646,703 | 40,888,158 |
| Percentage Traded | 0.4% | 16.4% | 83.2% | 0.7% | 17.3% | 82.0% |
| Maximum | 38,901 | 286,500 | 449,700 | 51,396 | 222,122 | 880,283 |
| Maximum Date | 2-May-07 | 4-Mar-10 | 11-Jan-10 | 2-May-07 | 19-Jul-07 | 11-Dec-09 |
| Mean | 114 | 4,618 | 23,405 | 117 | 3,058 | 14,458 |
| Std. Dev | 1,205 | 13,912 | 39,838 | 1,427 | 9,918 | 35,774 |
| Volume per option type | | 79,571,510 | | | 49,865,517 | |
| Total volume | 129,437,027 (23.1%) | | | | | |
| **Panel D - Contract expiring between 9 months and 1 year** | | | | | | |
| | CALL | | | PUT | | |
| | At the money | In the money | Out of the money | At the money | In the money | Out of the money |
| Volume Traded | 100,805 | 3,992,276 | 30,327,167 | 140,485 | 3,041,959 | 14,115,998 |
| Percentage Traded | 0.3% | 11.6% | 88.1% | 0.8% | 17.6% | 81.6% |
| Maximum | 19,520 | 240,900 | 765,140 | 36,200 | 183,000 | 326,210 |
| Maximum Date | 17-Mar-03 | 3-Mar-10 | 10-Dec-03 | 14-Sep-06 | 12-Sep-06 | 27-Feb-07 |
| Mean | 36 | 1,412 | 10,724 | 50 | 1,076 | 4,992 |
| Std. Dev | 482 | 6,846 | 30,284 | 825 | 5,553 | 14,310 |
| Volume per option type | | 34,420,248 | | | 17,298,442 | |
| Total volume | 51,718,690 (9.2%) | | | | | |

## Table 1. Continued

| **Panel E - Contract expiring between 1 year and 1 year and 3 months** | | | | | | |
|---|---|---|---|---|---|---|
| | **CALL** | | | **PUT** | | |
| | At the money | In the money | Out of the money | At the money | In the money | Out of the money |
| Volume Traded | 40,370 | 976,863 | 7,681,420 | 38,400 | 663,486 | 3,813,320 |
| Percentage Traded | 0.5% | 11.2% | 88.3% | 0.9% | 14.7% | 84.5% |
| Maximum | 5,500 | 24,800 | 120,900 | 10,000 | 64,000 | 237,900 |
| Maximum Date | 21-Nov-05 | 3-Dec-02 | 12-Sep-03 | 7-Sep-05 | 22-Dec-09 | 10-Dec-09 |
| Mean | 14 | 345 | 2,716 | 14 | 235 | 1,348 |
| Std. Dev | 210 | 1,348 | 8,127 | 244 | 1,469 | 6,375 |
| Volume per option type | | 8,698,653 | | | 4,515,206 | |
| Total volume | 13,213,859 (2.40%) | | | | | |
| **Panel F -Contract expiring between 1 year and 3 months and 1 year and 6 months** | | | | | | |
| | **CALL** | | | **PUT** | | |
| | At the money | In the money | Out of the money | At the money | In the money | Out of the money |
| Volume Traded | 15,750 | 441,905 | 3,472,557 | 18,550 | 409,348 | 1,589,324 |
| Percentage Traded | 0.4% | 11.2% | 88.4% | 0.9% | 20.3% | 78.8% |
| Maximum | 3,250 | 50,010 | 170,500 | 3,250 | 24,010 | 80,000 |
| Maximum Date | 31-Aug-07 | 27-May-03 | 5-Aug-04 | 31-Aug-07 | 22-Dec-09 | 14-Dec-09 |
| Mean | 6 | 156 | 1,228 | 7 | 145 | 562 |
| Std. Dev | 93 | 1,160 | 6,274 | 107 | 833 | 2,746 |
| Volume per option type | | 3,930,212 | | | 2,017,222 | |
| Total volume | 5,947,434 (1.1%) | | | | | |
| **Panel G - Contract expiring between 1 year and 6 months and 1 year and 9 months** | | | | | | |
| | **CALL** | | | **PUT** | | |
| | At the money | In the money | Out of the money | At the money | In the money | Out of the money |
| Volume Traded | 10,895 | 271,009 | 1,707,171 | 13,475 | 336,323 | 1,044,415 |
| Percentage Traded | 0.5% | 13.6% | 85.8% | 1.0% | 24.1% | 74.9% |
| Maximum | 4,500 | 7,500 | 62,750 | 4,500 | 12,000 | 87,350 |
| Maximum Date | 17-Mar-10 | 8-Jan-08 | 9-Jul-09 | 17-Mar-10 | 7-Jan-10 | 9-Dec-09 |
| Mean | 4 | 96 | 604 | 5 | 119 | 369 |
| Std. Dev | 106 | 423 | 2,662 | 121 | 568 | 2,370 |
| Volume per option type | | 1,989,075 | | | 1,394,213 | |
| Total volume | 3,383,288 (0.6%) | | | | | |
| **Panel H - Contract expiring between 1 year and 9 months and 2 years** | | | | | | |
| | **CALL** | | | **PUT** | | |
| | At the money | In the money | Out of the money | At the money | In the money | Out of the money |
| Volume Traded | 15,129 | 293,580 | 1,599,628 | 13,029 | 415,902 | 871,069 |
| Percentage Traded | 0.8% | 15.4% | 83.8% | 1.0% | 32.0% | 67.0% |
| Maximum | 2,500 | 23,500 | 124,000 | 2,500 | 27,000 | 34,000 |
| Maximum Date | 29-May-07 | 29-Jan-10 | 5-Mar-09 | 29-May-07 | 10-Feb-06 | 20-Aug-09 |
| Mean | 5 | 104 | 566 | 5 | 147 | 308 |
| Std. Dev | 95 | 605 | 3,535 | 88 | 778 | 1,470 |
| Volume per option type | | 1,908,337 | | | 1,300,000 | |
| Total volume | 3,208,337 (0.6%) | | | | | |

**Source:** NYSE Liffe

In addition, three other types of quality assurance checks are performed on the price data. First, a basic plausibility check: any option prices that are either zero or negative are immediately rejected. The second check is founded in option-pricing theory. In order to yield non-negative probability estimates, a call price function should be both monotonic and convex. In practice, this may not be the case if the difference between the “true” price of options with adjacent strikes is less than the minimum tick size, or if there are sufficiently large variations in the bid-ask spread. So any option prices that do not meet these monotonicity and convexity requirements are also excluded. Finally, if, after the application of the preceding two filters, there are less than three out-of-the-money option prices for a particular expiry date, then no PDF will be estimated for that expiry date.

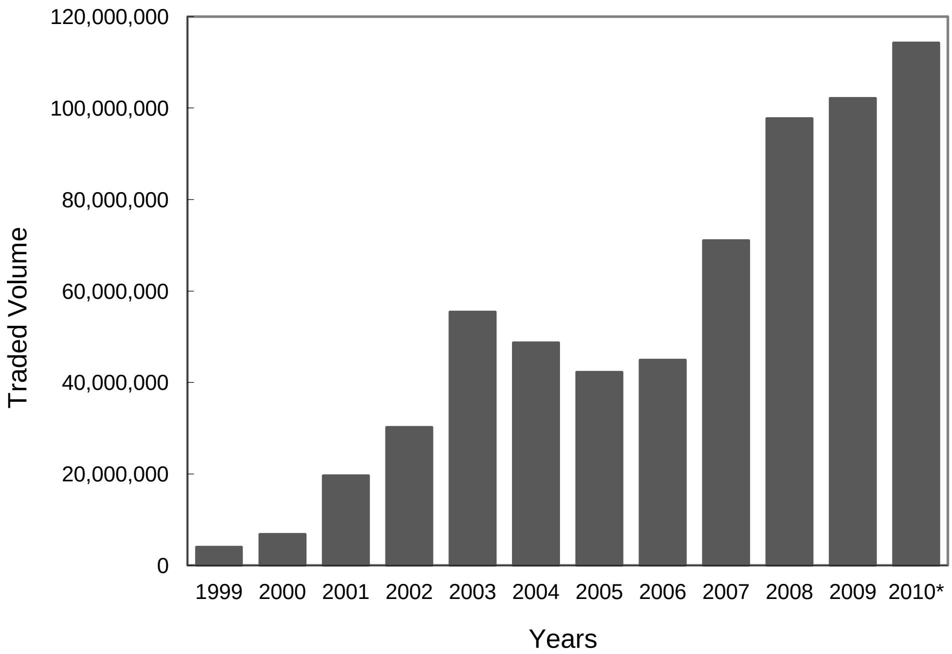


**Figure 1. Total traded volume for all contracts per year**
**Source:** LIFFE.
* annualised figures based on first semester observations

## 3. Methodology[5]

### *3.1. Fixed-Expiry Probability Density Functions*

The non-parametric technique used in this paper to derive the PDF is based on both Bliss and Panigirtzoglou (2002) and Cooper (1999). These two articles make use of 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.

The Breeden and Litzenberger result follows from the Cox and Ross (1976) pricing model, and is set out below:

$$C(F,K,\tau) = e^{-r\tau}\int_K^\infty f(F_T)(F_T - K)dF_T \tag{1}$$

$$\frac{\partial C(F,K,\tau)}{\partial K} = -e^{-r\tau}\int_K^\infty f(F_T)dF_T \tag{2}$$

$$\frac{\partial^2 C(F,K,\tau)}{\partial^2 K} = e^{-r\tau} f(F_T) \tag{3}$$

where C is the call function, K is the option's strike price, r is the risk-free rate, $F_T$ is the value of the underlying future at time T and *f* ($F_T$) is the implied probability density function which describes the possible outturns for the underlying futures at time T. The option's time to maturity, τ is equal to T-t. So in practice, the task of estimating a PDF using the Breeden and Litzenberger result amounts to estimating a twice-differential call price function, i.e. the blue dots in Figure 2 which represent the call price function. However, equation (3) cannot be applied directly to obtain *f* ($F_T$), because we only observe option prices for a discrete set of strike prices or interest rates, rather than a twice-differentiable continuum.

[5] White (1973), Neuhaus (1995), Bahra (1996, 1997), Soderlind and Svensson (1997), Melick and Thomas (1997, 1998), Clews *et al.* (2000), De Boor (2001), Andersen and Wagener (2002), and ECB (2011) are used for methodological and conceptual purposes while modelling the implied probability density functions.

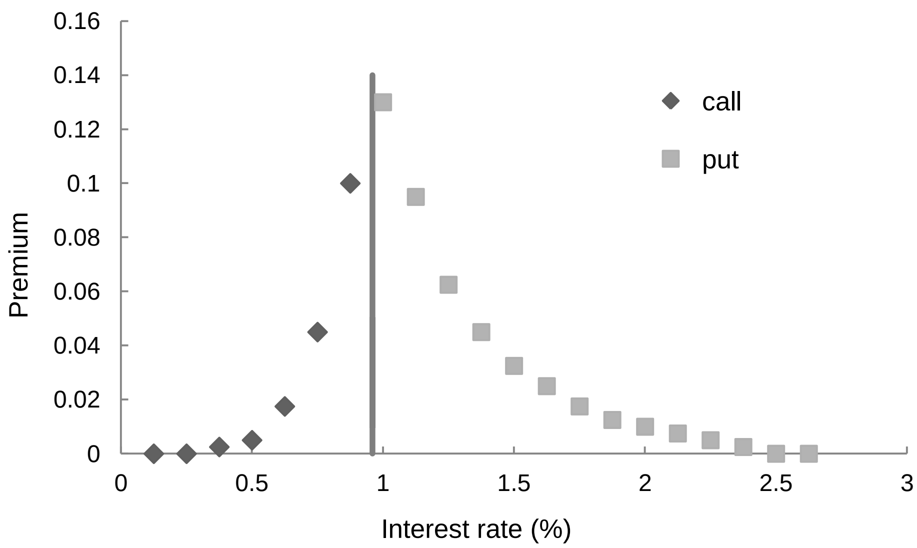


**Figure 2. Out-of-the-money calls and puts (March 2010 contract); 27 October 2009**

**Source:** NYSE Liffe

Taking the second derivative of a call price function estimated directly, interpolating through the discrete set of data on option premia and strike price as represented in Figure 2, can sometimes lead to unstable or inaccurate PDFs. This can also be seen in practice in the same figure, where some of the consecutive triplets of calls appear in a straight line, avoiding the possibility of calculating a twice-differential function. Instead, Bliss and Panigirtzoglou (2002), following the results derived from Malz (1997) and Shimko (1993), have suggested that smoother results might be obtained if the data on option premia and strike price are transformed into implied volatility and delta values prior to interpolation.[6]

The procedure for computing the PDF, $f(F_T)$, as defined by Bliss and Panigirtzoglou, can be described in several steps which are detailed below.[7] The first step consists in transforming the option strike prices into implied volatilities. Implied volatilities can be computed by

[6] The implied volatility or sigma of an option is defined as the volatility of the price of the underlying asset that is implied by the market price of the option based on the Black-Scholes model, and is calculated by inverting the Black-Scholes formula in the sense that, given an observed option price, a value for the volatility can be found that produces an option price which corresponds to the market price. Furthermore, the delta of an option measures the rate of change in the option price relative to changes in the underlying asset price. For example, with call options, a delta of 0.4 means that for every increase of unit in the underlying asset, the call option will increase by 0.4 unit. For call options, the delta is always defined in the [0,1] interval, whereas for put options, it is defined in the [-1,0] interval.

[7] The methodology described in this section is programmed in Matlab R2007a by using both built-in functions and proprietary code.

numerically solving for the value of σ which solves the Black (1976) futures options pricing model, for each option contract at time t :[8]

$$C(F,K,\tau) = e^{-r\tau}\left(F_t\Theta\left(\frac{\ln(\frac{F_t}{K}) + \frac{\sigma^2}{2}\tau}{\sigma\sqrt{\tau}}\right) - K\Theta\left(\frac{\ln(\frac{F_t}{K}) - \frac{\sigma^2}{2}\tau}{\sigma\sqrt{\tau}}\right)\right) \quad (4)$$

In the second step, the implied volatilities, i.e. the values of σ obtained in (4), are used to calculate the delta values by using the following formula:

$$\delta(F,K,\tau) = \frac{\partial C(F,K,\tau)}{\partial F} = e^{-r\tau}\Theta\left(\frac{\ln(\frac{F_t}{K}) + \frac{\sigma^2}{2}\tau}{\sigma\sqrt{\tau}}\right) \quad (5)$$

where Θ is the standard normal cumulative distribution function. The next step consists in interpolating the implied volatilities and the deltas calculated in (4) and (5). The advantage of interpolating in the implied volatility and delta space rather than in the initial premia and strike price space is also that the options which are deep out-of-the-money – i.e. calls with a delta close to zero and puts with a delta close to one – are grouped together as presented in Figure 3. This allows, as presented in the same figure, the interpolation to better fit calls and puts close to near-the-money, i.e. calls or puts with a delta close to 0.5. This makes the final PDF more suited for monetary policy purposes.

The interpolation of the implied volatility and delta is done, following Campa *et al.* (1997), using a cubic smoothing spline, which minimizes the following function:

$$\lambda\sum_{i=1}^{n}\omega_i(\sigma_i - g(\delta_i))^2 + (1-\lambda)\int_0^1 g''(t)^2\,dt \quad (6)$$

where λ is the smoothing roughness parameter, equal to 0.99[9], delta is the Black-Scholes *delta* and represents the x-axis of the spline, σ is the Black-Scholes sigma and represents the y-axis of the spline and the weights $\omega_i$ are calculated using

$\omega_i = \dfrac{v_i^2}{mean(v_i^2)}$ where $v_i$ is the Black-Scholes vega. The value of vega is almost negligible for options which are deep out-of-the-money and

[8] See, for instance, Hull (2000) for an overview of option pricing and related quantities.
[9] The optimal smoothing roughness parameter is the one that minimizes the observed deltas with the fitted deltas by the smoothing spline.

deep in-the-money and sequentially increases as we get near-the-money. In particular, it reaches a maximum for at-the-money options. Hence, the $\omega_i$ used in (6) place most weight on near-the-money options, and therefore lesser weight on away-from-the-money options. This is consistent with using these PDFs to support monetary policy analysis, where interest is likely to lie in the centre of the distribution, i.e. close to the underlying interest rate, rather than the distribution's tails. Figure 3 shows the interpolated volatility smile as a function of delta. Although delta can take values between 0 and exp(rτ), the traded contracts may not span that complete range. Therefore, the smoothing spline is extrapolated outside the range of traded price points with a second order polynomial, i.e. a quadratic equation. As a result of the extrapolation, the piecewise cubic curve obtained using interpolation is extended with a quadratic curve at each endpoint so that the full delta range, defined in the interval [0,1], is covered. As previously mentioned, the initial calculation of the second derivative of a call price can produce unstable PDF due to possible curvature constraints. Delta, vega and sigma are only used in this case, to avoid initially calculating the second derivative of a call price. However, note that although we are using the Black-Scholes formulae to calculate delta, vega, and sigma we do not necessarily need the assumptions of the Black-Scholes option model to hold true. The transformation from the initial premium and strike prices space into the delta and sigma space is only done to allow the option data to be interpolated such that the final PDF is more stable.

In the next step, the interpolated volatility smile is transformed back from volatility versus delta values to premium versus strike price values. This is done by evaluating the interpolated volatility smile, calculated in the previous step, at 1000 equally-spaced delta values between zero and one. The 1000 delta values are then transformed back into strike prices using the inverse of equation (5):

$$K = F_t \exp\left(\frac{\sigma^2 \tau}{2} - \sigma\sqrt{\tau}\,\Theta^{-1}(e^{r\tau}\delta)\right), \tag{7}$$

where $\Theta^{-1}$ is the inverse of the cumulative density function of a standardised Normal distribution. The implied volatility values of the spline are translated back into call prices using the Black-Scholes option pricing equation (4) as presented in Figure 4. In fact, this figure shows the fitted call price function and the fitted put price function after the interpolation. In the last step, we obtain the PDF by calculating the second derivative of the call function, applying the Breeden and Litzenberger result. To do so, we fit cubic polynomials through triplets of consecutive (strike price, call price) pairs; from the coefficients of the fitted polynomials we evaluate the second derivative, which gives us the final PDF.

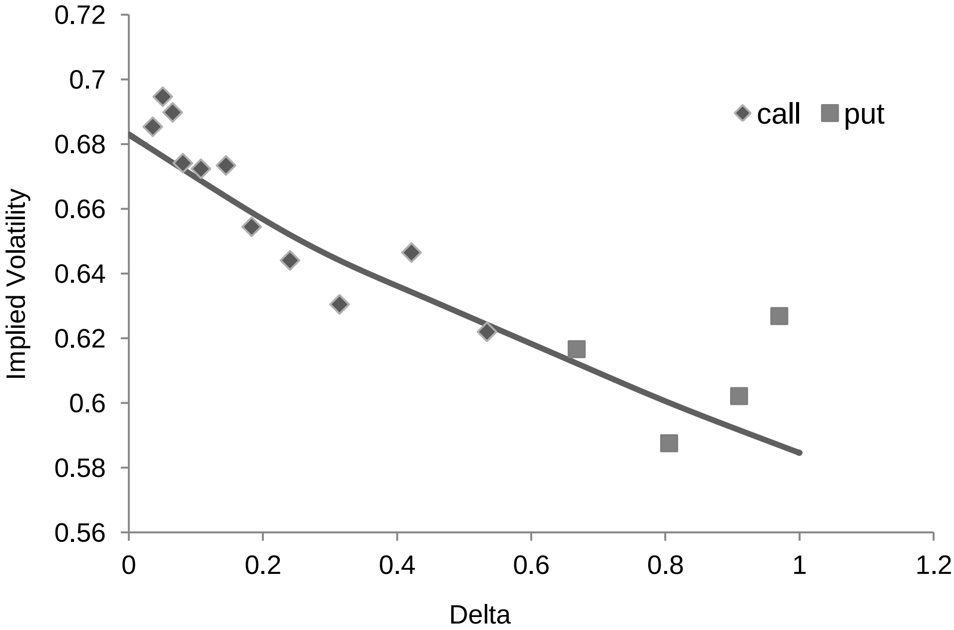


**Figure 3. Delta-implied volatility smile for the out-of-the-money calls and puts (March 2010 contract); 27 October 2009**

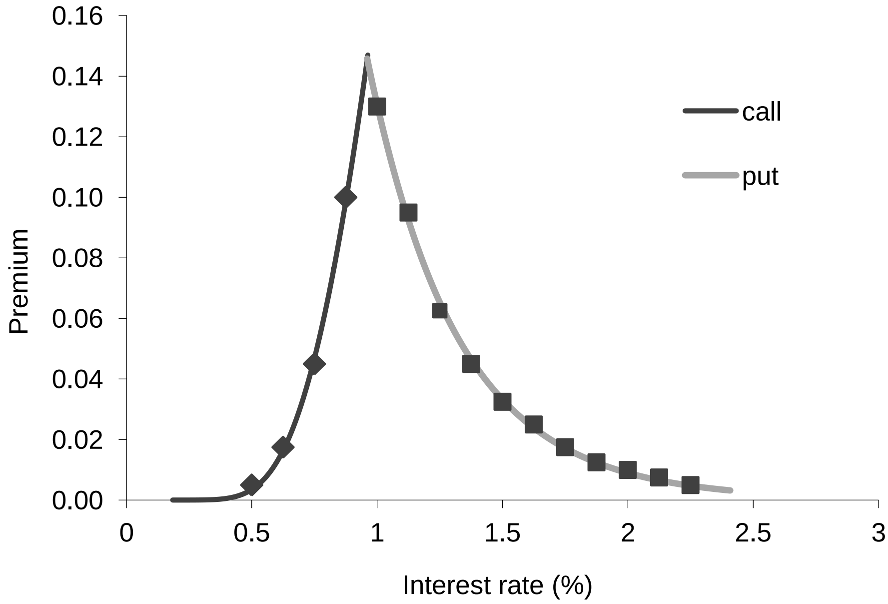


**Figure 4. Fitted call and put option function for the out-of-the-money calls and puts (March 2010 contract); 27 October 2009**

### 3.2. Constant Maturity Probability Density Functions

A total of eight quarterly option contracts on the three-month EURIBOR futures are traded daily on NYSE Liffe. Each of these eight contracts expires on the same day as the underlying future contract cycle of March, June, September or December. As each option contract gets closer to the expiry date, the uncertainty about possible future EURIBOR outcomes declines. Therefore, the amount of uncertainty embodied in the PDF also tends to decline as we approach the expiry date. Moreover, very little trading, if any, typically takes place on the days immediately prior to the expiry date. This regular time-to-maturity feature makes it very difficult to compare PDF statistics on the same fixed expiry contract over time. A possible solution to this time pattern is to estimate constant maturity PDFs interpolating over the eight fixed expiry PDFs. Based on this interpolation, we calculate three-month, six-month, nine-month, one-year and eighteen-month constant maturity contracts. For any given day, each of these PDFs always represents the same constant period ahead. An example of a 6-month constant maturity PDF is shown in Figure 5, where it is also compared with the two closest fixed expiry contracts.

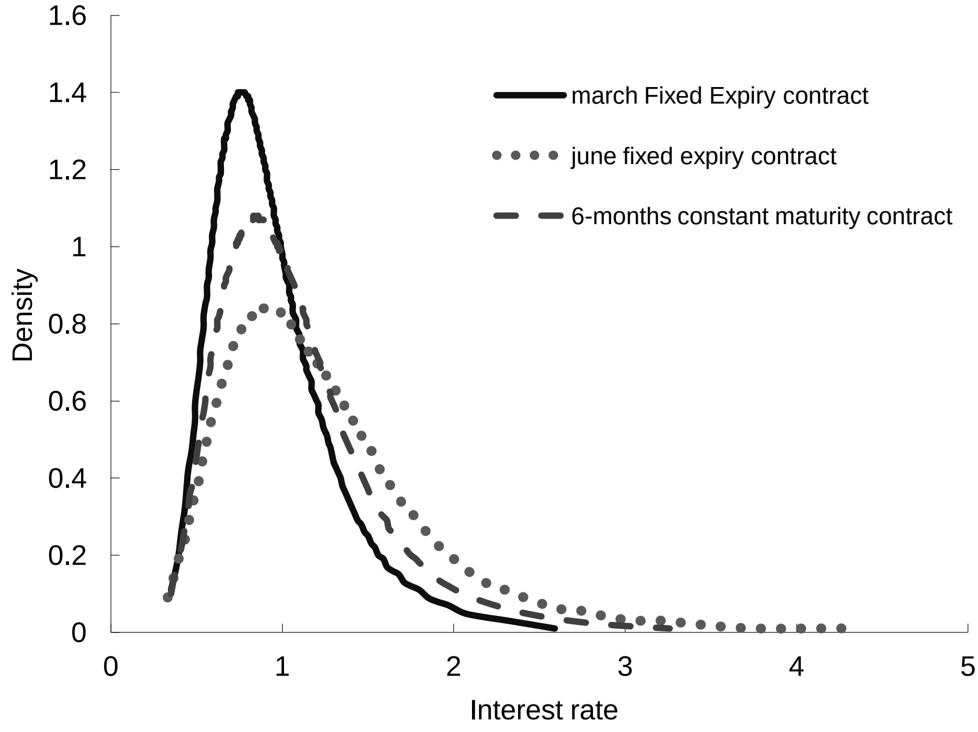


**Figure 5. Interpolation of the six-month constant maturity PDF on 27 October 2009**

The method does not interpolate directly over the PDFs but over the implied volatility curves with the same delta but with different maturities. The advantage of this method is that the same delta, but for contracts expiring on different dates, is always defined by the non-parametric technique. In addition, the delta always ranges between 0 and 1. In detail, to construct the constant maturity PDF a vector containing the nine delta values from 0.1 to 0.9, with a step-width of 0.1, is first created for every fixed expiry contract. For each delta in this vector, the value of the corresponding sigma is then calculated by evaluating the previously estimated volatility smiles. This is done by using the grid of 1000 two-component points defined in the previous section, where the first coordinate is the delta and the second is the sigma. From this grid, the nine sigmas are calculated using linear interpolation.

For each of the nine deltas, the value of the sigmas for different times to maturity (the fixed expiry ones) is calculated and a smoothing spline is constructed by interpolating the sigmas of all the fixed expiry contracts. Each of the nine splines is evaluated in the constant maturity values in order to obtain the corresponding sigmas at these points. Figure 6 presents the final nine pairs of deltas and sigmas points for the initial constant maturity contracts and also for the constant maturity ones.

We now have all the required data for each constant horizon: nine deltas, nine sigmas, tau, risk-free interest rate, and the underlying value, which is obtained by interpolating the two closest underlying contracts with a smoothing spline. The deltas are later converted into strikes and the premium of every artificially-created option is calculated using the Black- Scholes model. The non-parametric model is then used again to calculate the PDF, as defined in the previous section. Summarizing, we calculate the exactly ATM implied volatility, generate a 1000-point delta grid, calculate a 1000-point sigma grid using splines, then transform the deltas back into the strike space to calculate the premium, and finally calculate the constant maturity PDF by applying the Breedon and Litzenberger theorem.

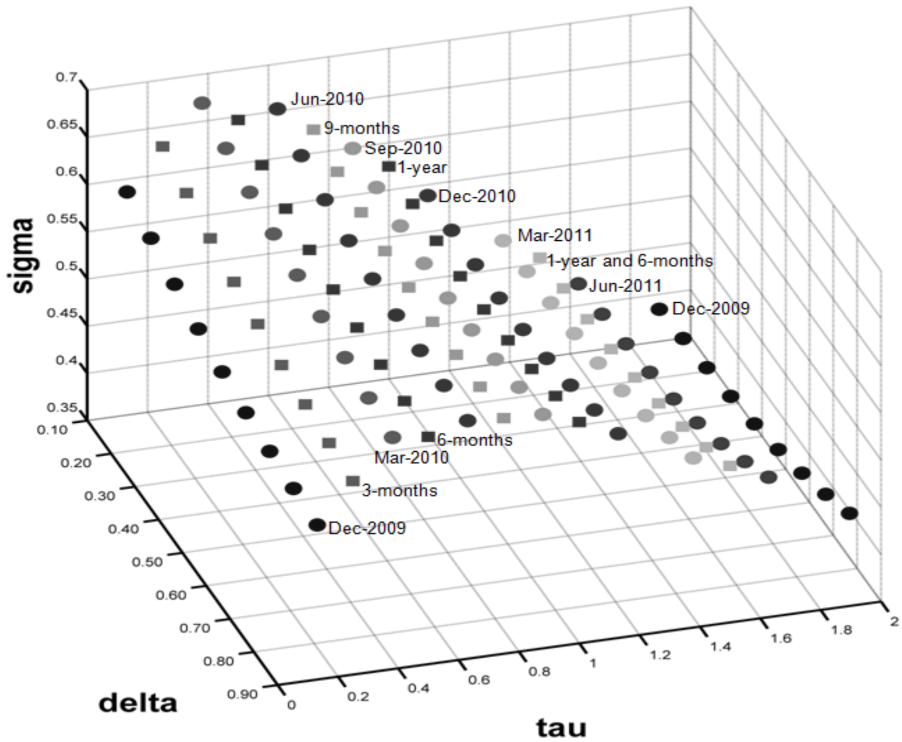


**Figure 6. Nine delta-sigma space for all the fixed expiry and constant maturity contracts; 27 October 2009**

## 4. How Can We Use Option-Implied PDF Derived Indicators?

This section provides a number of examples to demonstrate how option-implied PDFs may be able to enhance our analysis. Indicators derived from implied PDFs may be better quality than those derived from (single) option prices. Furthermore, option-implied PDFs may offer the possibility of new indicators, e.g. the most likely outturn implied by option prices (i.e. the mode of the implied distribution). Finally, PDFs are a powerful communication tool: they provide a concise, visual summary of risk and uncertainty - both magnitude and directional bias - embodied in option prices. Being able to visualize the distribution can be particularly useful when the associated risk parameters are changing rapidly. However, when interpreting option-implied PDFs two caveats must be borne in mind. The first is a general caveat, applicable to the interpretation of all option-implied PDFs; the second is specific to option-implied PDFs for EURIBOR, the subject of this paper. These two caveats are first discussed below.

### *4.1. Two Important Caveats to Interpreting EURIBOR PDFs*
#### *4.1.1. Option-Implied PDFs are Risk Neutral*

Option-implied PDFs estimated under a Black-Scholes option-pricing derivation (such as this one) are by construction risk-neutral. The option-implied PDF represents the set of probabilities under which the expectation of the final asset price must be discounted by the risk-free rate in order to equate with the market price. Such PDFs correspond to the probabilities that an agent would have if he were risk-neutral, but the agents that price the options might in fact be risk averse. If that were the case, risk premia would lead to differences between both the location and shape of the risk-neutral and actual distributions.[10] The extent of such differences is likely to vary with both asset class and maturity.

Different techniques can be used to transform the risk-neutral PDFs implied by options into estimates of the actual distribution. Bliss and Panigirtzoglou (2004) and Alonso *et al.* (2009) exploit the fact that the risk-neutral and actual distributions are related to each other via the marginal rate of substitution of the representative investor to define the functional form of the transformation. They then estimate the parameters of that transformation function for different assumed forms of the utility function by maximising the forecasting ability of the transformed PDFs. In contrast, Liu *et al.* (2004) following Fackler and King (1990), define their transformation in terms of the beta function. The additional flexibility of the beta function might better align the transformed PDFs with the pattern of past outturns, but perhaps at the cost of economic insight. Estimating the actual PDFs from the risk-neutral PDFs is outside the scope of this paper; instead it focuses directly on the option-implied PDFs themselves. The issue of how best to extract the actual distribution of possible asset price outcomes in the future is one that would merit further research.

[10] The concept of risk neutrality and how it might matter for the estimated PDF can be understood by drawing a parallel with insurance. The question is: how much does someone pay for, say, house insurance? Suppose the probability of your house burning down is $p$, and if that happens you will receive a payout $V$ from your insurance policy. A fair price is simply the payout multiplied by the probability of receiving the payout, $pV$. That is the price a risk-neutral agent would pay. However, if the agent were risk-averse, e.g. because he would value a payout of $V$ more in a state of the world where his house had burned down then he would be willing to pay more than $pV$ for the insurance policy. In other words, he would pay: $pV$ + *risk premium*. Similarly, a risk-loving agent would not even pay $pV$ (his risk premium would be negative). If market option prices embodied such a risk premium, then that would be captured by this methodology as a higher risk-neutral probability. This phenomenon is more clearly visible in the implied-PDFs for equity indices. Since the 1987 stock market crash, market participants have become more averse to further crashes and use options to insure against them. This increases the risk-neutral probability associated with large falls in the index, relative to market participants' actual assessment of the probability of such falls. Implied PDFs for equity indices, since 1987, are typically far more left-skewed than a histogram of actual outturns (Bates, 2000).

### *4.1.2. EURIBOR PDFs Pertain to the Inter-Bank Rate, not the Policy Rate*

Rates on overnight index swaps (OIS) are considered to provide the best market-based indication of market participants' expected path of average official policy rates. This is because although OIS may still include a premium to compensate for term risk and liquidity risk, the element that compensates for credit risk is minimal: it pertains to only overnight, rather than three-month, credit risk.

Before the financial turbulence, the spread between EURIBOR and EONIA, the euro OIS rate, had been small and stable (Figure 7): over the first half of 2007 it averaged 53 basis points, with a standard deviation of 0.7 basis points. At that time, therefore, the path of forward EURIBOR could also be considered a reasonable (if slightly upward biased) proxy of the market's expectations of average future policy rates. More importantly, the stability of the EURIBOR-EONIA spread meant that the risks around future EURIBOR outturns, as captured by EURIBOR PDFs, were driven by the perceived risks around the outlook for expected policy rates, rather than the outlook for that spread. But that spread became large and volatile with the onset of the financial turbulence, reaching over 200 basis points at its peak. This means that EURIBOR PDFs no longer characterize the risks purely around expected policy rates. Instead, they can be thought of as conflating the risks around central expectations for both the official policy rate and the inter-bank credit spread.[11] This does not diminish the value of EURIBOR PDFs because EURIBOR is still a fundamental element of the transmission mechanism.

### *4.2. PDFs May Offer New Information*

Theoretically, the mean of the PDF - the risk-neutral expectation of the outturn - is equal to the futures rate, by definition[12, 13].So differences in the mean of the PDF can simply be observed from movements in the (interpolated) futures rates. However, without information about the skewness of the distribution, one cannot determine whether differences in the futures rate are simply because the whole distribution has undergone a shape-preserving translation, or whether the weight on one of the tails has increased.

[11] Bank of England (2009) provides an indicative illustration as to how the PDF for the expected average policy rate and the PDF for the spread could be separated, but only if a simplifying assumption is made about the functional forms of both distributions.
[12] Although this methodology does not impose that condition.
[13] Ignoring the small difference between a forward rate and a futures rate that arises because of the margin requirement for exchange-traded futures.

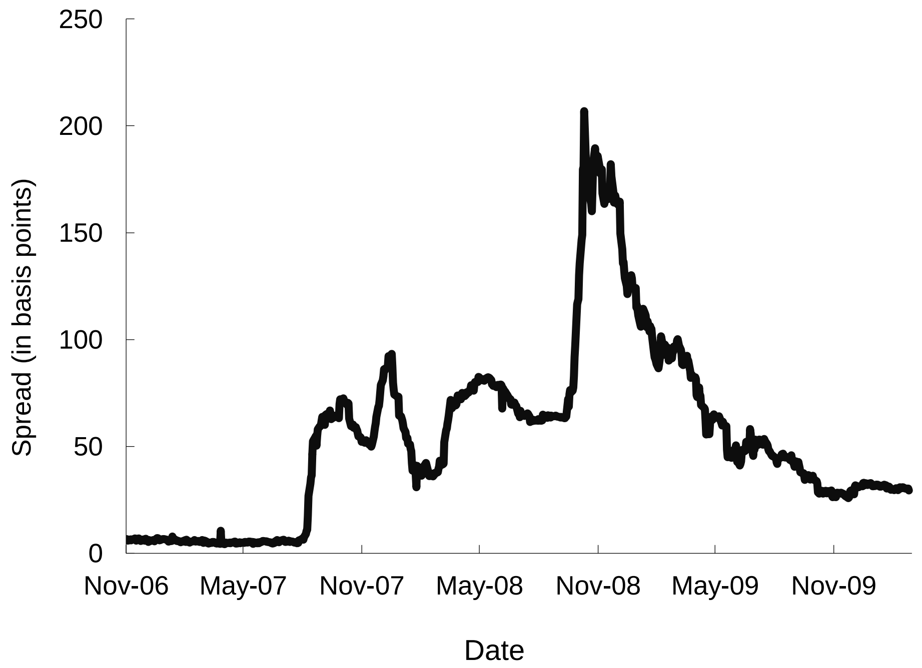


**Figure 7. Three-month EURIBOR and three-month forward EONIA sport rate spread**

**Source:** EURIBOR-EBF

It may therefore be useful also to consider differences in the mode of the distribution before interpreting differences in the mean (the futures rate). Note that this applies to interpreting differences across maturity as well as changes in one (constant) maturity over time. The PDFs estimated for 30 October 2009 in Figure 8 are a good example. Because of the strong positive skewness, the mean of the one-year PDF is notably higher than that of the three-month PDF. However, the modes are not so different. Figure 9 compares the mean path for EURIBOR, i.e. the futures curve, with the modal path implied by options. This shows that the most likely outcomes implied by options prices were for much weaker rises in EURIBOR over the coming year than suggested by the futures curve. That may be of interest to policy makers, although the caveat about risk-neutrality should be borne in mind.

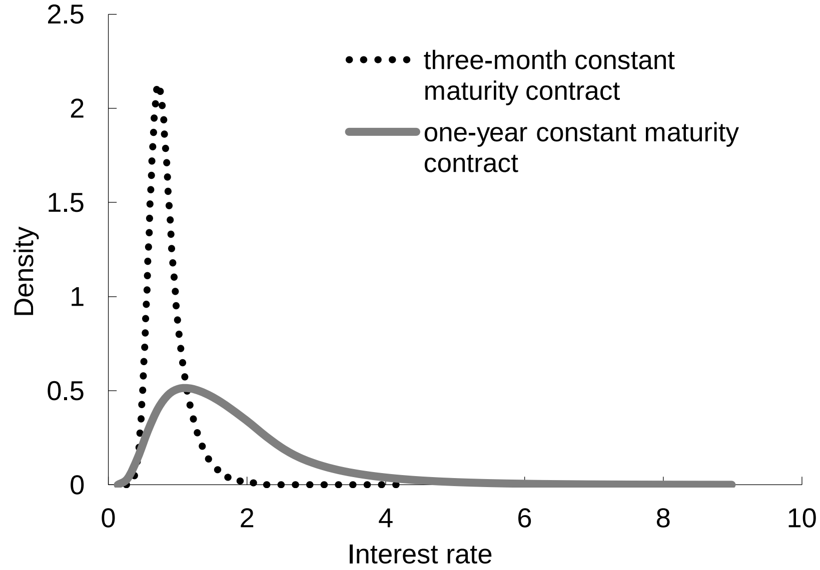


**Figure 8. Three-month and one-year constant maturity PDFs for 30 Oct. 2009**

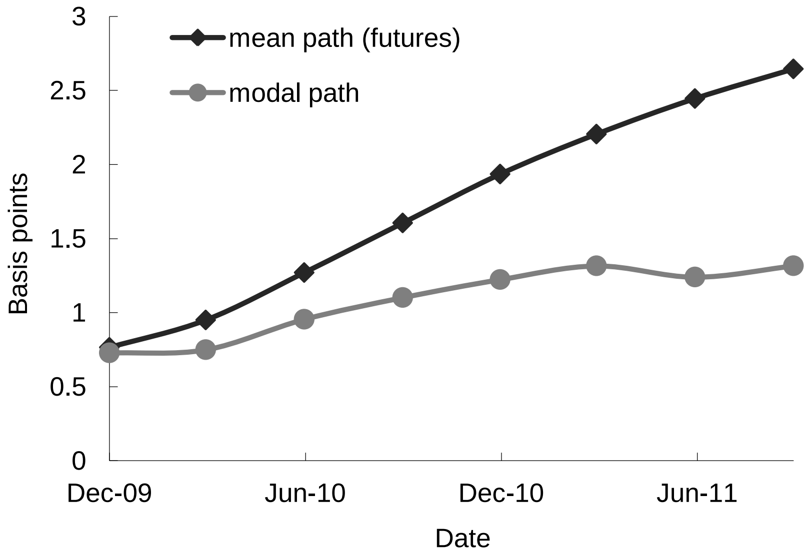


**Figure 9. Mean and mode interest rate paths**

*4.3. PDFs are a Powerful Tool for Conveying Information on Risk and Uncertainty*

The autumn of 2008 was especially tumultuous, but two events stand out: the failure of Lehman Brothers on 15 September and the internationally-coordinated monetary policy actions on 8 October. Option-implied PDFs are a powerful tool for succinctly capturing how such events affect market participants' views on the likely evolution of EURIBOR. They may also be used to assess the extent to which option prices anticipated such events.

The failure of Lehman Brothers led to material changes in the three-month-ahead EURIBOR distribution as shown in Figure 10. While there had been little movement in the PDF in the preceding week, EURIBOR option prices assigned a significantly greater weight to interest rate outturns much lower than the prevailing forward rate. And that left-tail continued to grow. Stress in the cash markets increased markedly too, and the spread between forward EURIBOR and EONIA increased. But while it could also be argued that the large negative skew reflects in part the view that the EURIBOR-EONIA spread could be much narrower than the forward spread, the sheer magnitude of the left tail suggests that it is also likely to reflect beliefs about future policy rates. However, it is difficult to be sure that such developments were not influenced by changes in risk aversion. For instance, if investors' intrinsic assessment of the actual probabilities of such outturns had not changed, but rather they decided that they would now require protection (in the form of options) against outturns at that particular probability, then that would also increase the estimated risk-neutral probabilities.

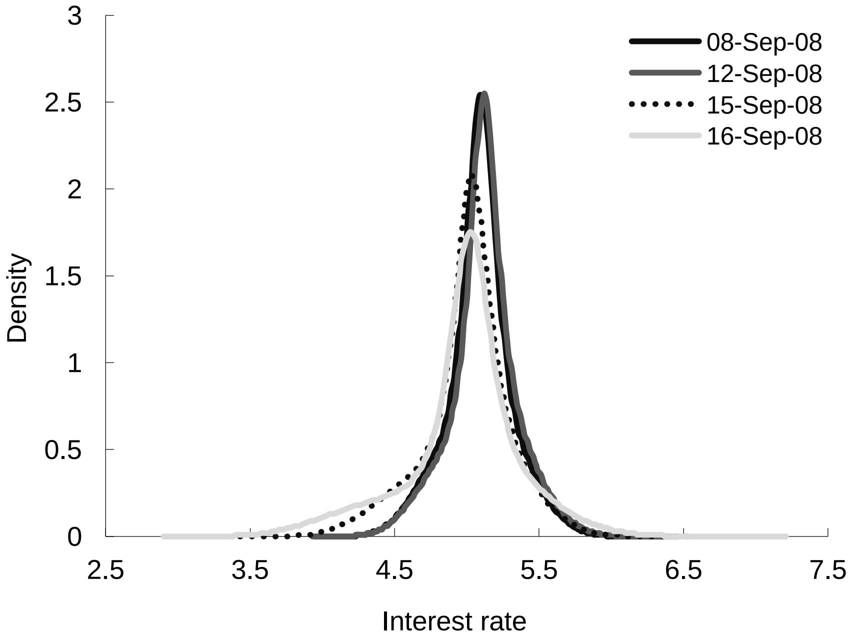


**Figure 10. Three-month constant maturity EURIBOR PDFs before and after the failure of Lehman Brothers on 15 September 2008**

On 8 October, as part of internationally-coordinated monetary policy action, the ECB announced that, from the operation settled on 15 October, the weekly main refinancing operations would be carried out through a fixed rate tender procedure with full allotment at the interest rate on the main refinancing operation, i.e. 3.75%.[14] That rate was 50 basis points below the minimum bid rate affirmed at the previous Governing Council meeting on 2 October. An examination of the EURIBOR PDFs in the days leading up to that announcement and shortly afterwards reveals two interesting observations as presented in Figure 11. First, it appears as if the impact of both the press conference on 2 October and the announcement on 8 October on the option-implied EURIBOR distribution was small compared to that of the accumulation of news during the intervening days (in particular, over the weekend). The fact that even by 7 October, the PDF had shifted so much to the left, and become more negatively skewed, suggests that market participants were already placing more weight on EURIBOR outturns in three months' time being much lower than the current forward rate, even though the precise timing and details of the 8 October announcement took the market by surprise. The second interesting observation is that although the bulk of the implied three-month EURIBOR distribution continued to move towards lower interest rates, there was no movement in the tail of the distribution. One possible explanation is that, despite the unprecedented events of the preceding month, market participants still did not attach any weight to the possibility that EURIBOR would be 2% or less in three months' time.

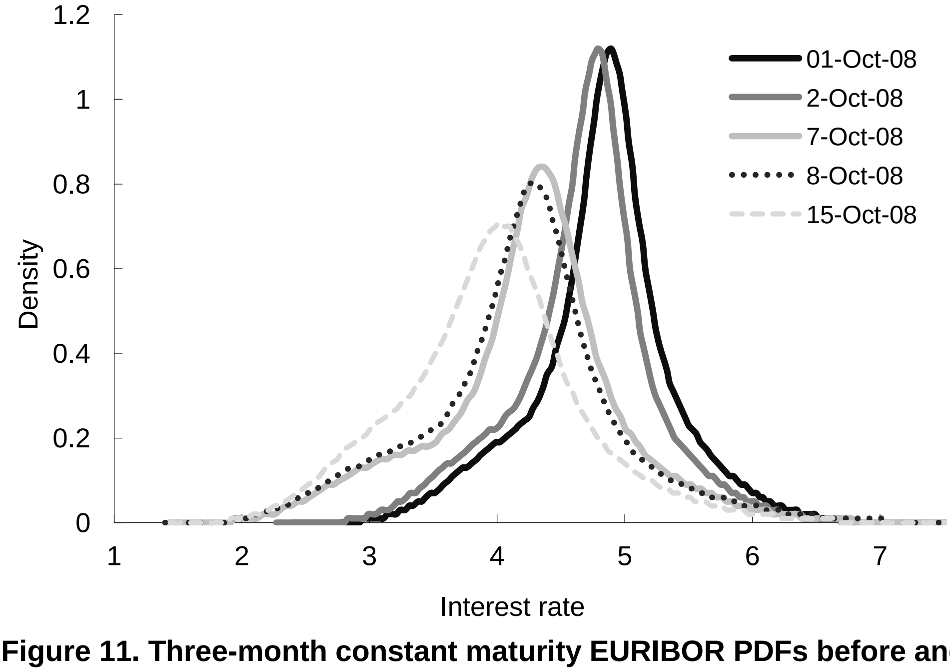


**Figure 11. Three-month constant maturity EURIBOR PDFs before and after the change in monetary policy on 8 October 2008**

---

[14] It was also announced on 8 October that, as of 9 October, the ECB would reduce the corridor of standing facilities from 200 basis points to 100 basis points around the interest rate on the main refinancing operation. Further details on both announcements can be found at:
http://www.ecb.europa.eu/press/pr/date/2008/html/pr081008.en.html

## 5. The Evolution of Option Implied PDF Statistics during the Financial Crisis

This final section documents in detail how EURIBOR PDFs reacted to the unfolding financial crisis between 2007 and 2009. In doing so, it demonstrates how the higher moments of the option-implied PDFs can provide timely and quantitative indicators of not only the amount of uncertainty around forward EURIBOR, the mean of the PDF, but the directional bias within that. The data are introduced in Figures 12 to 14 to provide a general overview. Figure 12 first shows the mean of option-implied distributions, which is simply equal to the forward rate around which the risks are measured. Figure 13 presents a measure of the amount of uncertainty and Figure 14 a measure of its directional bias. Note that the options price data in early 2007 did not always meet the quality criteria outlined in Section 2 to estimate PDFs. The following two episodes are then examined more closely:

1. The onset of financial market turbulence
2. February to August 2008: The tension between declining demand and rising prices

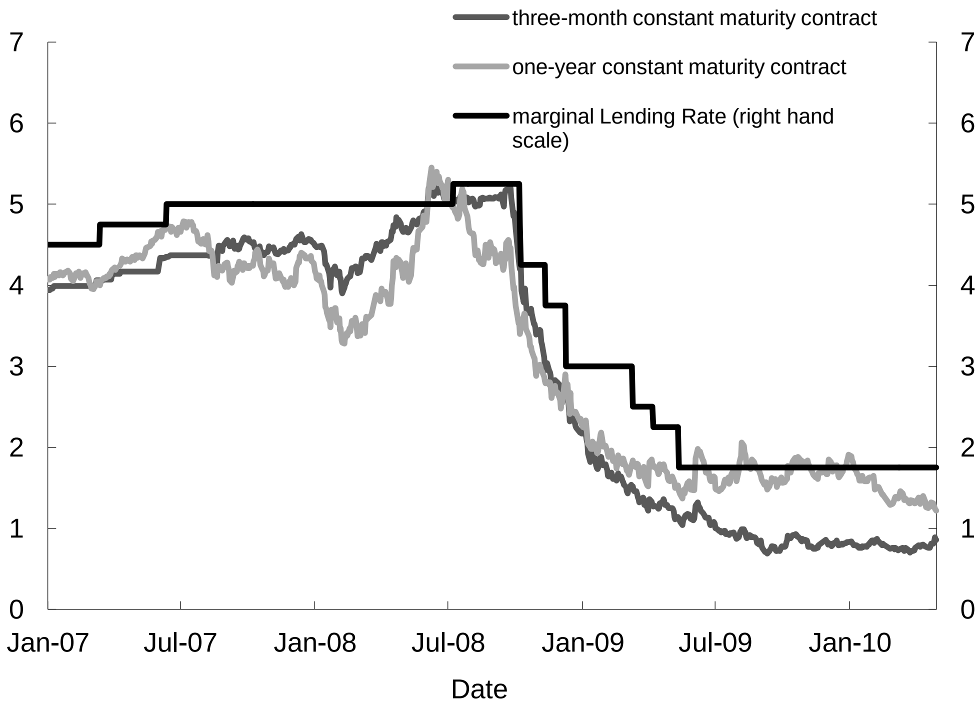


**Figure 12. Mean of the three-month and one-year EURIBOR constant maturity PDFs from 2007**

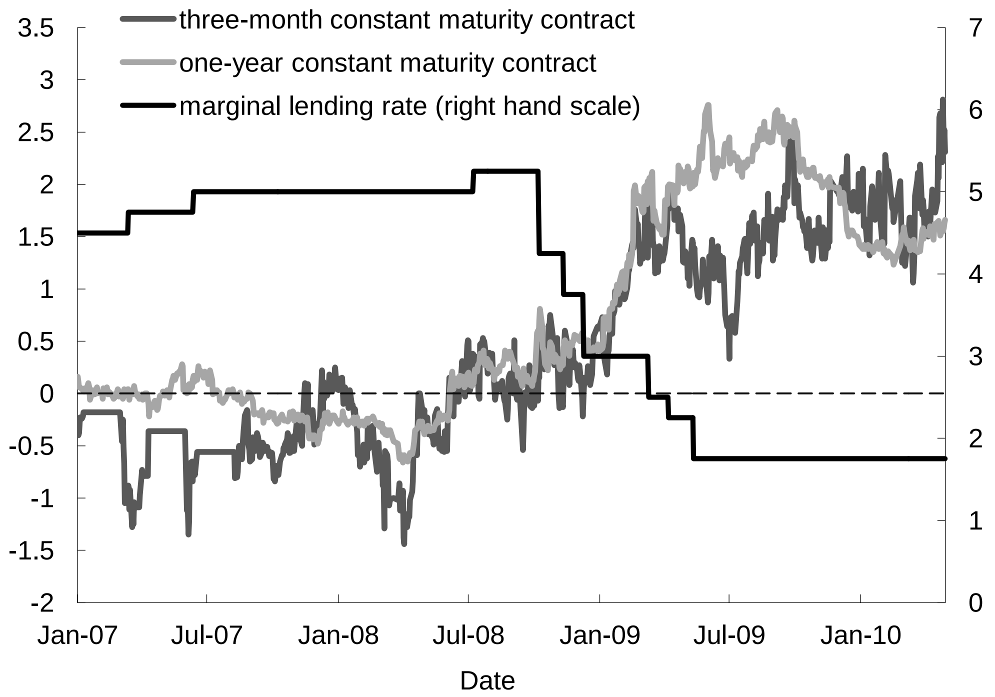


**Figure 13. Implied Volatility of the three-month and one-year EURIBOR constant maturity PDFs**

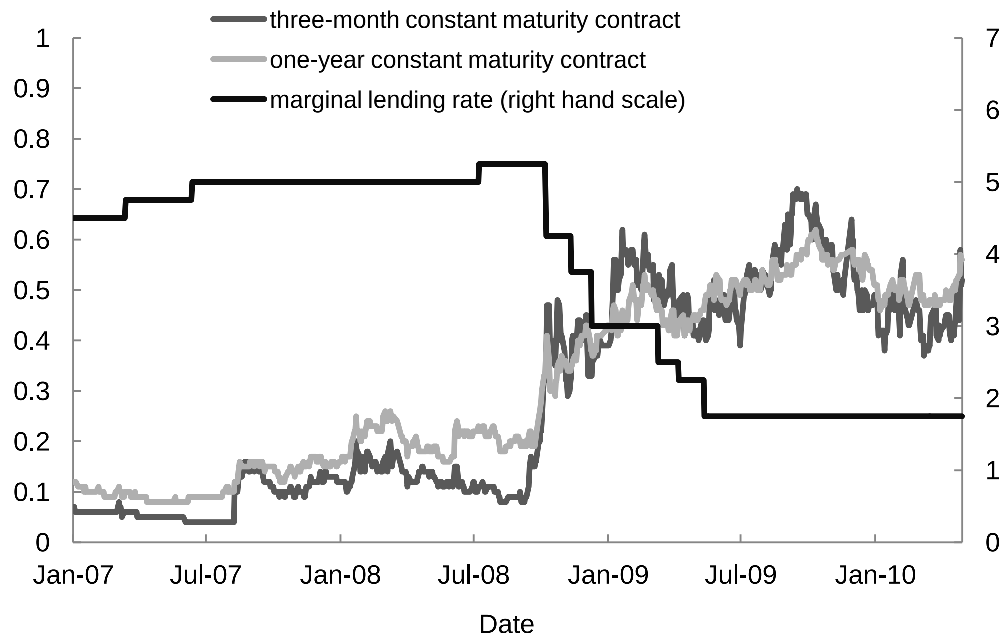


**Figure 14. Skewness of the Three-month and One-year EURIBOR constant maturity PDFs**

### *5.1. The Onset of Financial Market Turbulence*

One striking feature of the onset of financial market turbulence was the dislocation that occurred in short-term money markets on 9 August 2007, when EURIBOR and the EONIA rate diverged (Figure 15). Both the EURIBOR and EONIA curves flattened, but the EURIBOR curve lying unusually far above the EONIA curve (Figure 16). The fact that the spread between the red and yellow lines narrows, suggests that the market expected the situation of abnormally high three-month EURIBOR-EONIA spreads to ease only slowly, over the coming year.

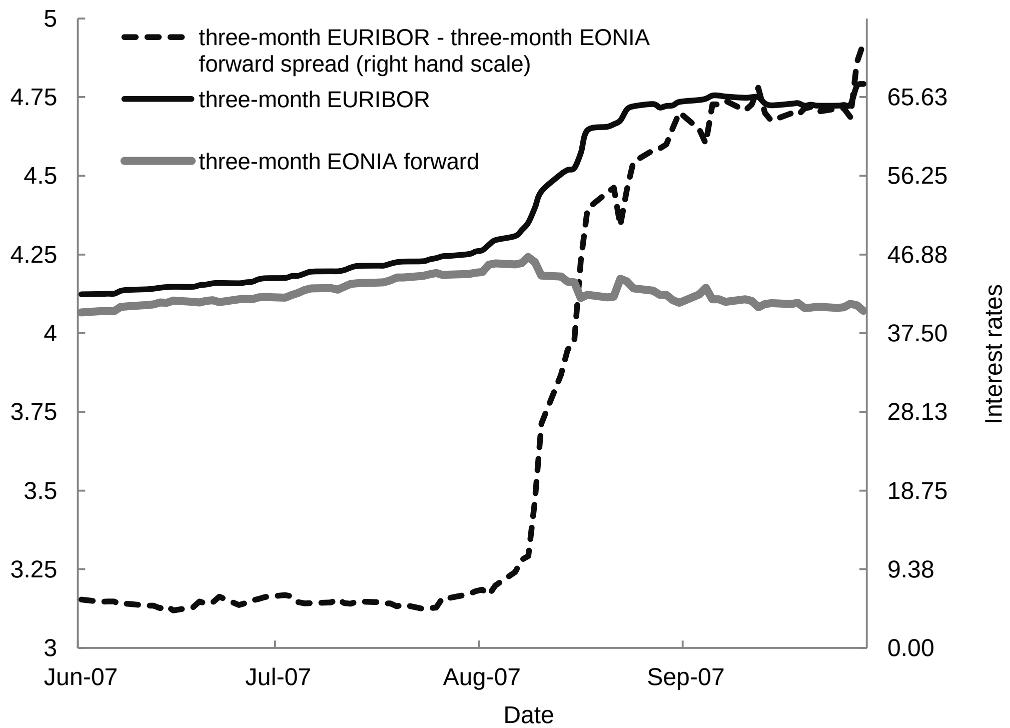


**Figure 15. Three-month interest rates, three-month EONIA forward, and the spread between them**
**Source:** Bloomberg and EURIBOR-EBF

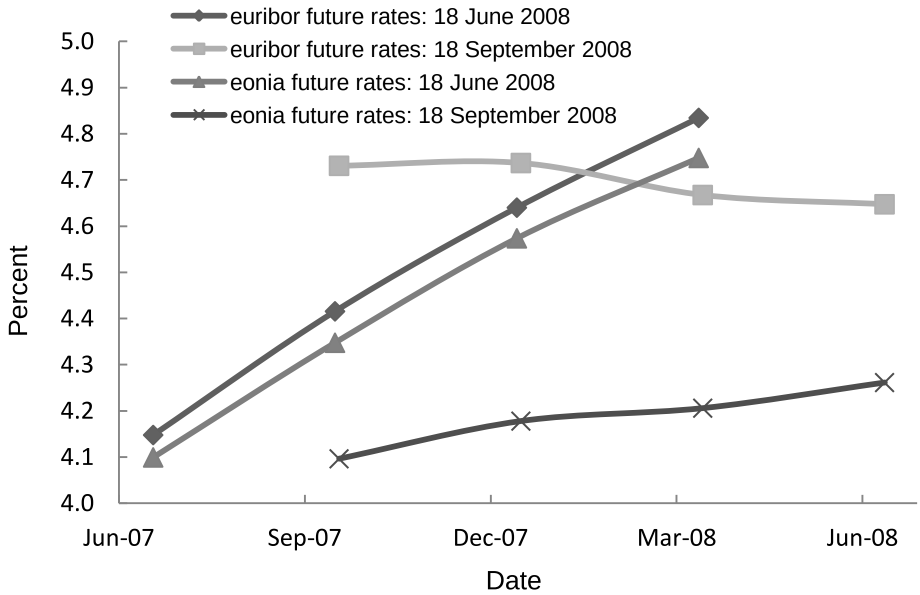


**Figure 16. EURIBOR and EONIA three-month forward curves, before and after the money market dislocation**

**Source:** Bloomberg and EURIBOR-EBF

Option-implied EURIBOR PDFs offer insight on the market's assessment of the risks around the EURIBOR curve. Figure 17 shows the estimated three-month PDFs before and after the onset of market turbulence. The moments of these PDFs, and additional information, are presented in Table 2. These data show how the width of both distributions increased considerably, reflecting in part the abrupt and unprecedented divergence from EONIA swap rates and ensuing uncertainty about the speed and magnitude of any subsequent convergence. This increase in width was predominantly a near-term phenomenon. However, as already noted in Section 4.1.1, these PDFs are risk-neutral, so an increase in width could be because of an increase in risk aversion as well as an increase in the actual amount of risk. In this context, both factors may well have played a role. Movements in the skewness of the distribution indicate how market participants perceive the balance of risks to be changing.[15] At short horizons there was little change in the balance of risks. However, at longer horizons the balance of risks moved to the downside. This suggests that market participants placed more weight on outturns that were much lower than the prevailing

[15] Here, 'balance of risks' has the precise economic meaning as set out in Lynch and Panigirtzoglou (2004). In summary, it is the difference between expected conditional losses, depending on whether the outturn is greater or less than the central estimate, for an agent with rates forecast error with a quadratic loss function.

forward rate at that time. And that could be consistent with an even more rapid return to more normal spread levels than the interest rate curves alone suggested.

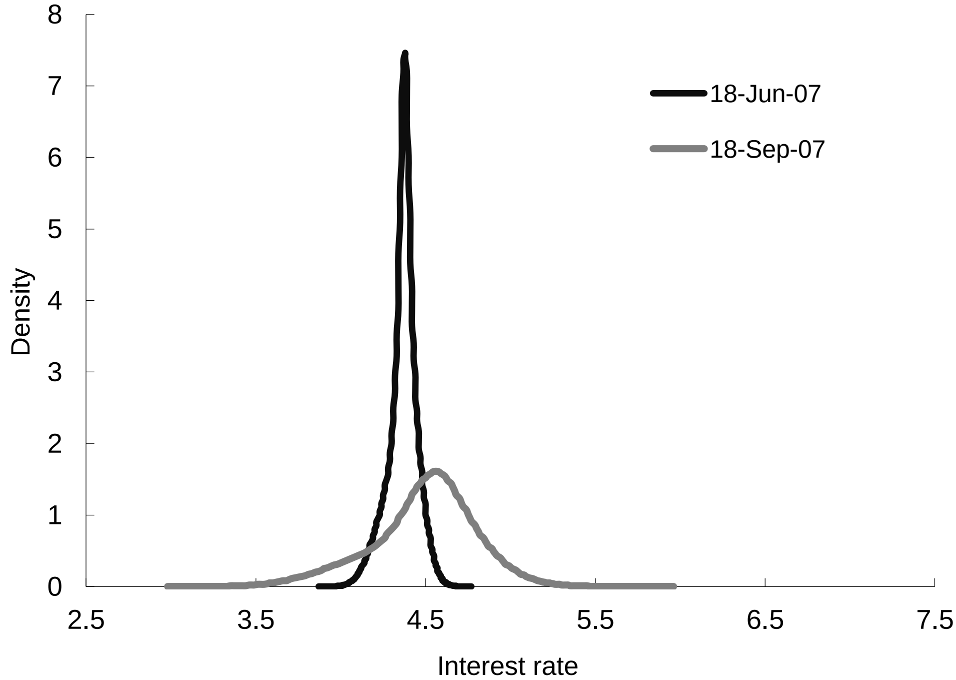


**Figure 17. Three-month constant maturity EURIBOR PDFs, before and after the onset of market turbulence**

**Table 2. Moments of the three-month constant maturity in Figure 17 and related information**

| **Moments** | **3-months** | **1-year** | **change** |
|---|---|---|---|
| Mean | 4.36 | 4.5 | 0.14 |
| Standard deviation (percentage points) | 0.08 | 0.3 | 0.22 |
| Skewness | -0.55 | -0.47 | -0.08 |
| Memo | | | |
| Implied Volatility | 0.03 | 0.13 | 0.1 |
| Implied Volatility (basis points) | 0.15 | 0.58 | 0.43 |
| Forward rates from the EURIBOR spot rate | 4.42 | 4.74 | 0.32 |
| Forward rates from the EONIA spot rate | 4.35 | 4.18 | -0.17 |
| EURIBOR - EONIA (basis points) | 7 | 56 | 49 |

So the onset of financial market turbulence led market participants to reappraise their view on longer-term rates, and their assessment of the uncertainty around shorter-term rates. These developments were captured by movements in option-implied EURIBOR distributions. In particular, the standard deviation and skewness of these distributions inform us about the quantity and balance of risk, subject to the risk-neutral caveat. One advantage of these indicators is that they are a quantitative measure and can therefore be used to put the latest developments into an historical context. The movements in these option-implied indicators directly following the outbreak of the financial market turbulence did not appear to be exceptionally notable compared to their own history. That, and the fact that implied uncertainty did not change much at longer horizons, may suggest that the market did not, at first, believe that the overall impact of the turbulence would be severe.

### *5.2. February to August 2008: A Tension between Declining Demand and Rising Prices*

By February 2008 it was becoming clear that despite the possible demand implications of the financial crisis, risks to inflation over the medium term were still to the upside, and forward rates began to increase once more. The rise in three-month forward EONIA over February to May broadly unwound the policy cuts that had been implicitly priced in during January, whereas, as presented in Figure 18, EURIBOR rose significantly above its year-end level, thus widening the gap between these two interest rates. However, as we show in Figure 19, the balance of risks around forward EURIBOR moved significantly to the downside between January and April. So although market participants were revising up their central expectation for EURIBOR outturns, they were initially still attaching increasing weight to EURIBOR outcomes below the forward rate. This may reflect views on either the ECB policy rate in the future, or the evolution of the EURIBOR-EONIA spread.

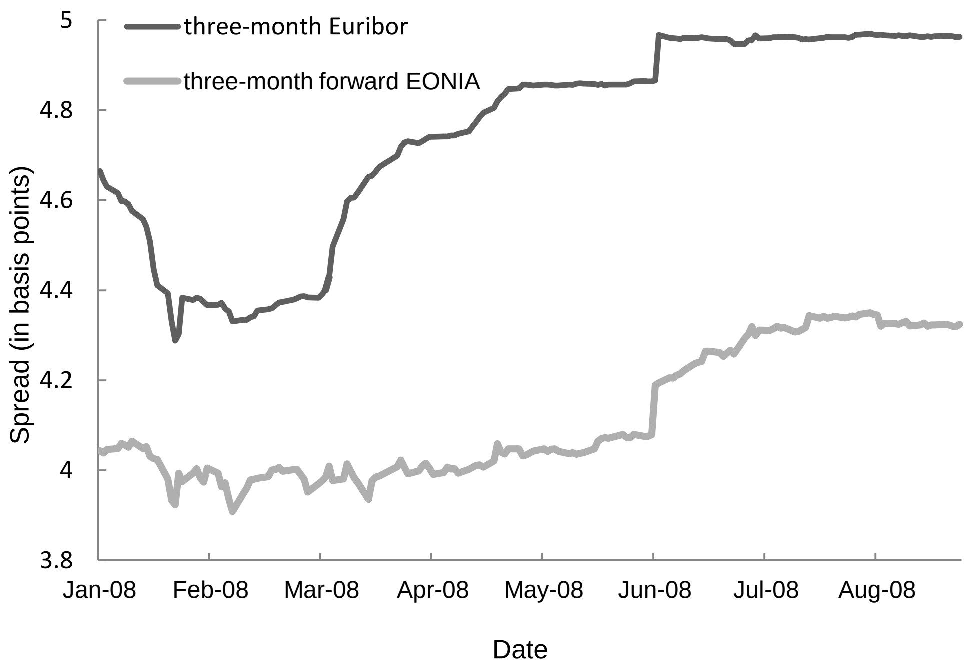


**Figure 18. Three-month EURIBOR and three-month forward EONIA**
**Source:** Bloomberg and EURIBOR-EBF

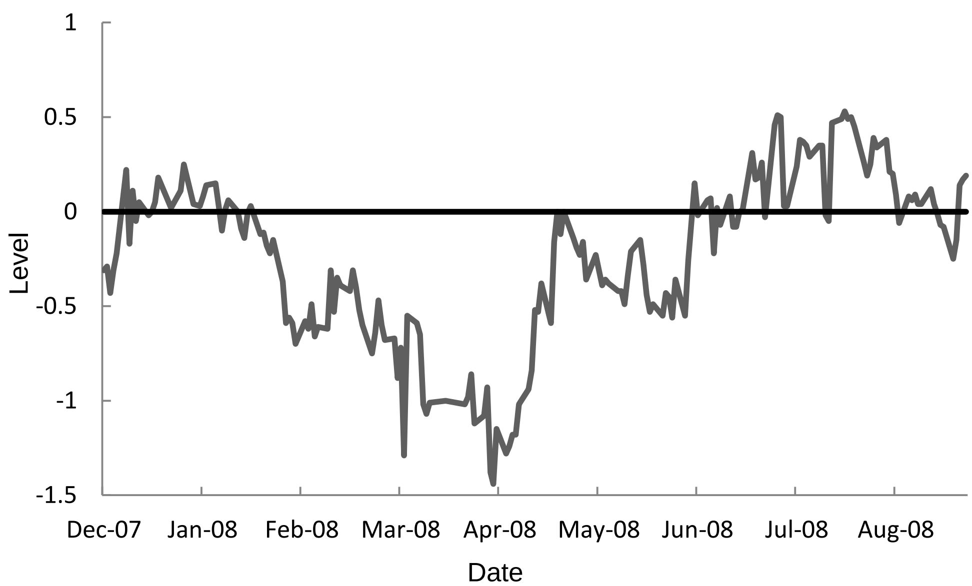


**Figure 19. Skewness of the three-month EURIBOR constant maturity PDF**
**Source:** Bloomberg and EURIBOR-EBF

## 6. Concluding Remarks

This paper has shown how the methodology for extracting probability distributions from the prices of financial options, as first developed by Bliss and Panigirtzoglou (2002) and Cooper (1999), can be applied to EURIBOR. Using this methodology, we have estimated probability distributions for EURIBOR outturns three months in the future; the resulting dataset which is to be made publicly available via the ECB's Statistical Data Warehouse comprises over ten years of daily data. These PDFs provide a timely and quantitative indication of the market's assessment of the risks around forward EURIBOR: not just how much uncertainty there is, but precisely how that is distributed over different possible outturns. These can be used to analyse trends such as the extent to which the balance of risks is skewed to the upside or the downside, or to analyse how specific events affected the entire spectrum of views. Therefore, this indicator may appeal to those interested in monetary policy or financial stability. Moreover, such a comprehensive dataset, spanning the complete history of the euro, is particularly valuable because it gives the context against which the current situation or recent developments may be compared and provides a benchmark to help judge whether the current situation is 'normal' or 'extreme'. For most of the euro's history, the balance of risks around EURIBOR was driven primarily by the perceived balance of risks around the key policy rate. However, following the financial market turbulence of August 2007 that originated in the US sub-prime mortgage market and the exceptional consequences for banking systems worldwide, expectations of interbank rates in the future diverged from expectations of the policy rate. Therefore, EURIBOR PDFs must be interpreted more carefully during this period because they combine a view on possible future values of the policy rate with possible future values of the EURIBOR-EONIA spread. Nevertheless, they still provide a good quantitative indication of the balance of risks around this key part of the monetary policy transmission mechanism.

## References

Alonso, F., Blanco, R., and Rubio, G., 2009. Option-implied preferences adjustments, density forecasts and the equity risk premium. *Spanish Economic Review*, 11(2), pp.141-164.

Andersen, A.B. and Wagener, T., 2002. Extracting risk neutral probability densities by fitting implied volatility smiles: some methodological points and an application to the 3m EURIBOR futures option prices. *ECB Working Paper Series*, 198.

Bank of England, 2009. An indicative decomposition of the option- implied probability distribution for Libor. *Quarterly Bulletin*, 49(3), pp.158-159.

Bates, D.S., 2000. Post-'87 crash fears in the S&P 500 futures option market. *Journal of Econometrics*, 94(1-2), pp.181-238.

Bahra, B., 1996. Probability distributions of future asset prices implied by option prices. *Bank of England Quarterly Bulletin*, 36(3), pp.299-311.

Bahra, B., 1997. Implied risk-neutral probability density functions from option prices: theory and application. *Bank of England Working Paper*, 66.

BIS, 1999. Estimating and interpreting probability density functions, *Proceedings of the workshop on implied PDFs held at the BIS on 14 June 1999*. Bank for International Settlements, Basel.

Black F., 1976. The pricing of commodity contracts. *Journal of Financial Economics,* 3(1-2), pp.167-179.

Bliss, R., and Panigirtzoglou, N., 2002. Testing the stability of implied probability density functions. *Journal of Banking and Finance,* 26(2-3), pp.381-422.

Bliss, R. and Panigirtzoglou, N., 2004. Option-implied risk aversion estimates. *Journal of Finance*, 59(1), pp.407-446.

Breeden, D.T. and Litzenberger, R.H., 1978. Prices of state-contingent claims implicit in option prices. *Journal of Business*, 51(4), pp.621-51.

Campa, J.M., Chang, P.H.K., and Reider, R.L., 1997. Implied exchange rate distributions: evidence from OTC option markets. *NBER Working Paper*, 6179.

Clews, R., Panigirtzoglou, N., and Proudman, J., 2000. Recent developments in extracting information from option markets. *Bank of England Quarterly Bulletin,* 40(1), pp.50-60.

Cooper, N., 1999. Testing techniques for estimating implied RNDs from the prices of European-Style options. *Proceedings of the workshop on implied PDFs held at the BIS on 14 June 1999*. Bank of International Settlements, Basel.

Coutant, S., Jondeau, E., and Rockinger, M., 1999. Reading interest rate and bond futures options smiles: How PIBOR and notational operators appreciated the 1997 French snap election. *Banque de France Working Paper NER*, 54.

Cox, J.C. and Ross, S.A., 1976. The valuation of options for alternative stochastic processes. *Journal of Financial Economics,* 3(1-2), pp.145-166.

De Boor, C., 2001. *A practical guide to splines*. Revised Edition. New York: Springer-Verlag.

ECB, 2011. The information content of option prices during the financial crisis. *ECB Monthly Bulletin*, February, pp.87-99.

Fackler, P.L. and King, R.P., 1990. Calibration of option-based probability assessments in agricultural commodity markets. *American Journal of Agricultural Economics,* 72(1), pp.73-83.

Hull J., 2000. *Options, futures, and other derivatives.* 4th edition. New Jersey: Prentice Hall.

Lynch, D.P.G. and Panigirtzoglou, N., 2004. Using option prices to measure financial market views about balances of risk to future asset prices. *Bank of England Quarterly Bulletin,* 44(4), pp.442-454.

Liu, X., Shackleton, M.B., Taylor, S.J., and Xu, X., 2007. Closed-form transformations from risk-neutral to real-world distributions. *Journal of Banking and Finance*, 31(5), pp.1501-1520.

Malz, A.M., 1997. Estimating the probability distribution of the future exchange rate from option prices. *Journal of Derivatives*, 5(2), pp.18-36.

Melick, W.R. and Thomas, C.P., 1997. Recovering an asset's implied PDF from option prices: an application to crude oil during the Gulf crisis. *Journal of Financial and Quantitative Analysis*, 32(1), pp.91-115.

Melick, W.R. and Thomas, C.P., 1998. Confidence intervals and constant maturity series for probability measures extracted from option prices. *Proceedings of the Bank of Canada Conference on 'Information Contained in Prices of Financial Assets'.*

Neuhaus, H., 1995. The information content of derivatives for monetary policy-implied volatilities and probabilities. *Economic Research Group of Deutsche Bundesbank Discussion Paper,* 3/95.

Shimko, D.C., 1993. Bounds of probability. *Risk,* 6(4), pp.33-37.

Soderlind, P. and Svensson, L.E.O., 1997. New techniques to extract market expectations from financial instruments. *Journal of Monetary Economics,* 40(2), pp.383-429.

White, A., 1973. *Pricing options with futures-style margining: A genetic adaptive neural network approach (Financial sector of the American economy)*. New York: Garland Publishing.

## 9.2 Article 2: Interest rate expectations and uncertainty during ECB Governing Council days

Journal of Banking & Finance 36 (2012) 2804–2823

Contents lists available at SciVerse ScienceDirect

# Journal of Banking & Finance

journal homepage: www.elsevier.com/locate/jbf

# Interest rate expectations and uncertainty during ECB Governing Council days: Evidence from intraday implied densities of 3-month EURIBOR

Olivier Vergote *, Josep Maria Puigvert Gutiérrez

*European Central Bank, Frankfurt am Main D-60311, Germany*

ARTICLE INFO

*Article history:*
Received 23 June 2011
Accepted 29 June 2012
Available online 8 July 2012

*JEL classification:*
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

ABSTRACT

This paper analyses changes in short-term interest rate expectations and uncertainty during ECB Governing Council days. For this purpose, it 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 3-month EURIBOR 3 months ahead. Estimates of the noise impact on the statistical moments of the densities 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.



## 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

* Corresponding author. Tel.: +49 69 1344 0; fax: +49 69 1344 6000.
*E-mail addresses:* olivier.vergote@ecb.europa.eu (O. Vergote), josep_maria.puigvert@ecb.europa.eu (J.M. Puigvert Gutiérrez).

0378-4266/$ - see front matter 
http://dx.doi.org/10.1016/j.jbankfin.2012.06.014

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 Puigvert-Gutiérrez and de Vincent-Humphreys (2012) 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. Following this approach, the extraction of a density from option prices can be seen to consist of four steps. Fig. 1 presents an example of these estimation steps.

The method starts with the selection of option price observations, where 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; see also Bliss and Panigirtzoglou (2002) and Puigvert-Gutiérrez and de Vincent-Humphreys (2012). As a confirmation, Table 1 presents the trading volume for all Euribor options since the first day of trading, 13 January 1999 up to 2009. From this table it can be derived that in absolute terms 81% of the options are traded out of the money whereas only 18% are traded in the money. Furthermore, some of the in-the-money options are not traded independently, but as part of a bundled trading strategy, e.g. straddles or strangles, which combine out-of-the-money options with in-the-money options.

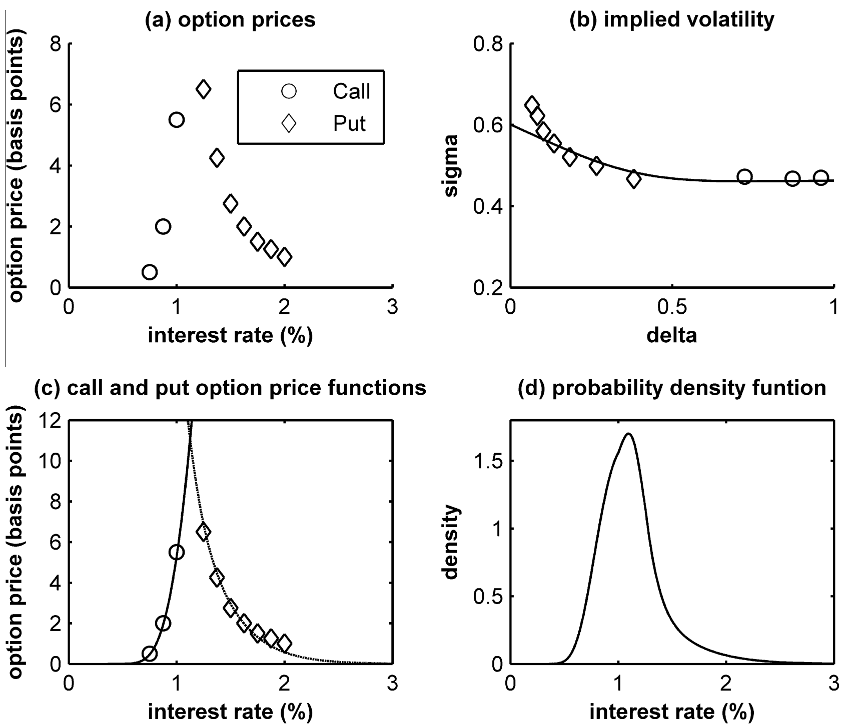


**Fig. 1.** Estimation steps.

**Table 1**
Total traded volume by type of option and intrinsic value.[a]

| Year | Total volume | Totals by PUT/CALL | | | | | | | | |
|---|---|---|---|---|---|---|---|---|---|---|
| | | Total | | | Call | | | Put | | |
| | | At-the-money (% of the total volume) | In-the-money (% of the total volume) | Out-of-the-money (% of the total volume) | At-the-money | In-the-money | Out-of-the-money | At-the-money | In-the-money | Out-of-the-money |
| 1999 | 4018 | 22 (0.54%) | 587 (14.6%) | 3410 (84.86%) | 10 | 374 | 2497 | 11 | 213 | 913 |
| 2000 | 6805 | 81 (1.18%) | 1320 (19.34%) | 5404 (79.42%) | 41 | 577 | 3203 | 39 | 743 | 2201 |
| 2001 | 19,617 | 184 (0.94%) | 3678 (18.75%) | 15,756 (80.32%) | 142 | 2808 | 10,198 | 41 | 870 | 5558 |
| 2002 | 30,210 | 149 (0.49%) | 5241 (17.35%) | 24,820 (82.16%) | 95 | 4083 | 18,467 | 54 | 1158 | 6353 |
| 2003 | 55,417 | 352 (0.64%) | 9972 (17.99%) | 45,093 (81.37%) | 263 | 7183 | 34,242 | 89 | 2789 | 10,851 |
| 2004 | 48,662 | 338 (0.69%) | 8476 (17.42%) | 39,848 (81.89%) | 194 | 6005 | 29,706 | 144 | 2471 | 10,142 |
| 2005 | 42,271 | 449 (1.06%) | 7584 (17.94%) | 34,238 (81%) | 251 | 4341 | 21,993 | 198 | 3243 | 12,245 |
| 2006 | 44,879 | 273 (0.61%) | 10,572 (23.56%) | 34,034 (75.83%) | 83 | 4283 | 18,080 | 190 | 6289 | 15,954 |
| 2007 | 71,030 | 801 (1.28%) | 11,623 (16.36%) | 58,606 (82.51%) | 390 | 4882 | 31,907 | 412 | 6742 | 26,700 |
| 2008 | 97,434 | 719 (0.74%) | 19,285 (19.77%) | 77,430 (79.49%) | 461 | 13,349 | 50,400 | 258 | 5936 | 27,030 |
| 2009 | 102,065 | 424 (0.41%) | 16,478 (16.15%) | 85,163 (83.44%) | 324 | 13,908 | 61,237 | 100 | 2570 | 23,926 |

*Unit*: Contracts traded per year in thousands, rounded to the nearest thousand.
[a] As regards the transaction amount, it needs to be noted that since their launch LIFFE Euribor contracts have come to dominate the euro denominated short-term interest rate (STIR) derivatives market, capturing over 99% of the market share; they are now the most liquid and heavily traded euro-denominated STIR contracts.

Focusing on the out-of-the-money (and at-the-money) options also implies that 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 Fig. 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 derivative 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 volatilities. The implied volatilities are calculated by numerically solving the Black's (1976) version of the options pricing model for the value of $\sigma$

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

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 $\Phi$ 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

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

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

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

where $\sigma_i, \hat{\sigma}_i$, and $\omega_i$ are respectively the observed sigma, fitted sigma and weight of observation $i$ $(i = 1, \ldots, n, \sum_{i=1}^{n}\omega_i = 1)$, $\delta$ represents the observed deltas, $\Theta$ is the matrix of polynomial parameters belonging to the spline, $g(\cdot)$ is the cubic spline function, and $\lambda$ is the smoothing parameter fixed at 0.99. The weighting is based on Black–Scholes vega $(\nu), \omega_i = \upsilon_i^2/\sum_{i=1}^{n}\upsilon_i^2, i = 1, \ldots, n$. 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 (1000) of delta values, transforming the delta values back into strike prices using the inverse of Eq. (2):

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

where $\Phi^{-1}$ 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 of 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 3 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 Puigvert-Gutiérrez and de Vincent-Humphreys (2012) 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.

## 3. Intraday implied densities

This section explains the application of the non-parametric density estimator to tick data on futures of 3-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. Section 3.1 explains the way prices are measured for the empirical study in Sections 5–7. 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 withdrawals 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 Puigvert-Gutiérrez and de Vincent-Humphreys (2012), 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, Fig. 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 mid-quote exists. This data filter helped to obtain more density estimates within days of low market activity.

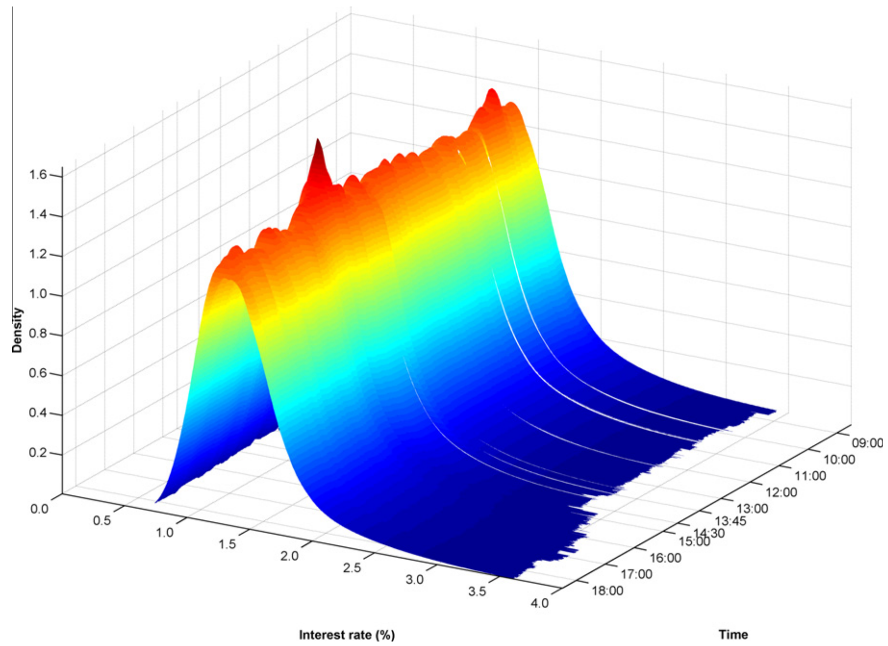


**Fig. 2.** Implied densities at tick frequency during 5 March 2009.

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.

Fig. 3 presents several central moments at tick frequency. The impact of the noise is clear here with eratic 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 $i$, $i = 1, \ldots, n$, 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, $r_i = m_i - m_{i-1} = u_i + e_i$, $i = 2, \ldots, n$. 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, $\omega^2$, is given by

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

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 Fig. 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.

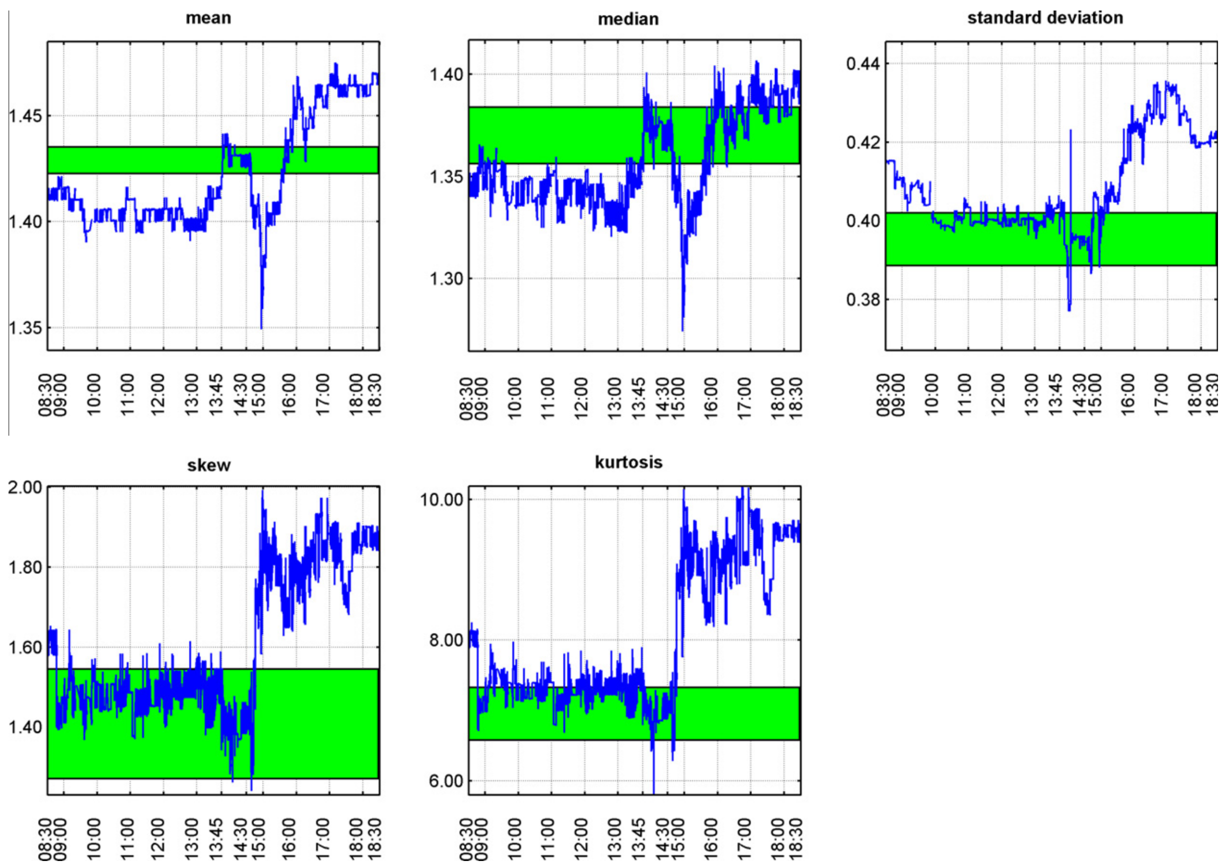


Fig. 3. Selected density moments at tick frequency during 5 March 2009 and noise band estimate centred around the moment level reached at 14:30 C.E.T.

In this example, the mean and median are found to decrease sharply during the press conference and to end at a significantly higher level by the end of the day. Also the standard deviation, skew and kurtosis move to higher levels than what could be considered noise induced.

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.

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. The stability of the estimates confirms the robustness results of Bliss and Panigirtzoglou (2002) now at an intraday level. 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 interpreting 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 3-month EURIBOR rate 3 months ahead. This implies that what counts are interest rate expectations between now and 6 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 6 months from now for a 3-month contract, the estimation may also pick up interest rate expectations between 6 and 9 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 introductory statement were seen as predicting the mentioning of vigilance at the next meeting or hence a hike in 2 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. Fig. 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.

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

Fig. 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 3-month EURIBOR 3 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.,

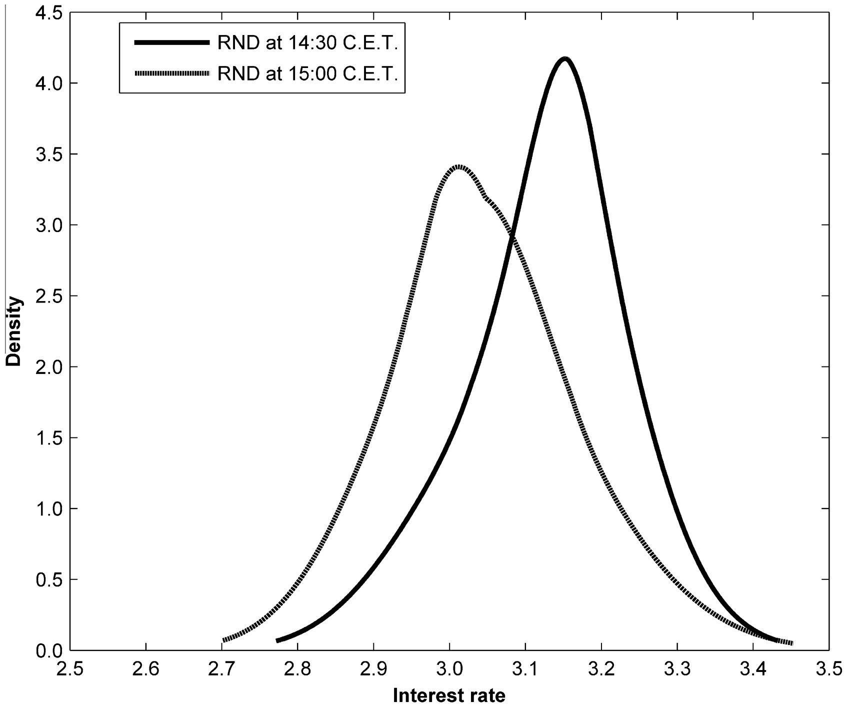


**Fig. 4.** Two intraday implied densities on 6 April 2006.

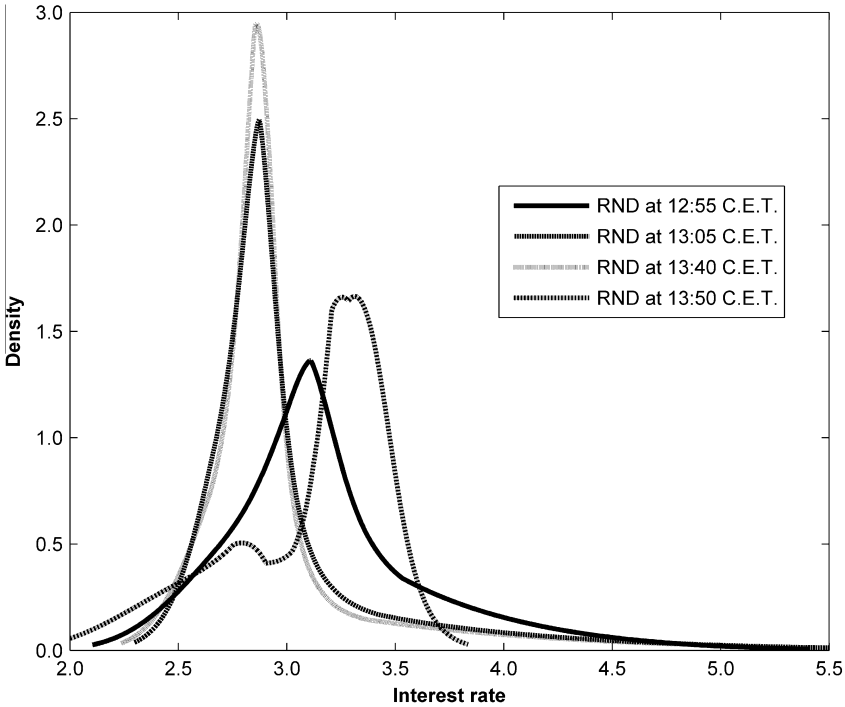


**Fig. 5.** Implied density mid-day developments on 6 November 2008.

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 min and strongly increased rate expectations of 3-month EURIBOR 3 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.

### 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

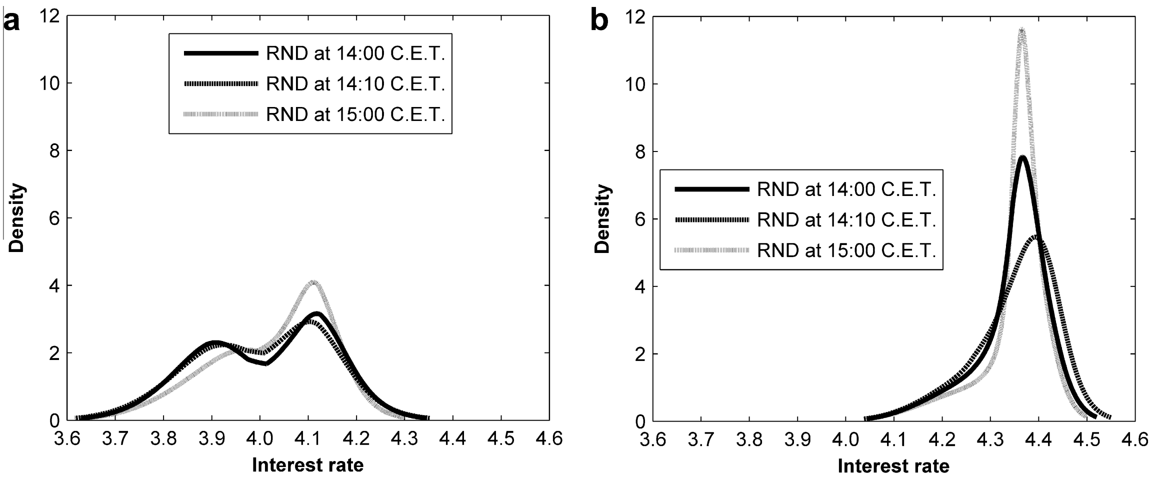


**Fig. 6.** Implied densities during 8 March 2007 (lhs) and 6 June 2007 (rhs).

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 2 days, Fig. 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.

## 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 3-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 A 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.

Section 6.1 analyses the intraday pattern of the density moments during policy rate decision days. Section 6.2 assesses the range of the intraday moment changes. Section 6.3 studies the persistence of the Governing Council impact on daily density moments.

### 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-min change in each of the moments was computed per minute and is presented in Appendix D. For this purpose, densities were computed at 1-min equally-spaced time intervals. The intraday moment patterns are closely connected to the tick arrival patterns presented in Appendix C with shocks at 13:00 and 13:45 C.E.T., and elevated levels during the press conference in all moments. 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 D 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.

The patterns derived by judging the figures in Appendix D are also confirmed by statistical tests. Table 2 presents averages for

**Table 2**
*T*-test results for difference in averages.

| | Average value during | | | | |
|---|---|---|---|---|---|
| | Morning | Press release[c] | Introductory statement | Q&A | Press conference |
| *Abs. change*[a] | | | | | |
| Mean | 0.0015 | 0.0042***/** | 0.004*** | 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*** |
| Skew | 0.0365 | 0.0785***/ | 0.1004*** | 0.0774*** | 0.0826*** |
| Kurtosis | 0.0536 | 0.1578***/ | 0.1725*** | 0.1353*** | 0.1434*** |
| *Level*[b] | | | | | |
| 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), 10 min following the press release and the actual duration of the initial statement and Q&A session.
[a] Test if mean absolute changes are higher than in the morning.
[b] Test if mean level is smaller than in the morning.
[c] The second test for the press release tests its mean to the press conference mean.
** Significance at the 95% level.
*** Significance at the 99% level.

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 2 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 Puigvert Gutiérrez and de Vincent-Humphreys (2012) provide strong support by showing that the Governing Council days imply significantly bigger reaction than other days. For this purpose, Appendix B 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. Fig. 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-min 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 10–20 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 Puigvert Gutiérrez and de Vincent-Humphreys (2012) 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

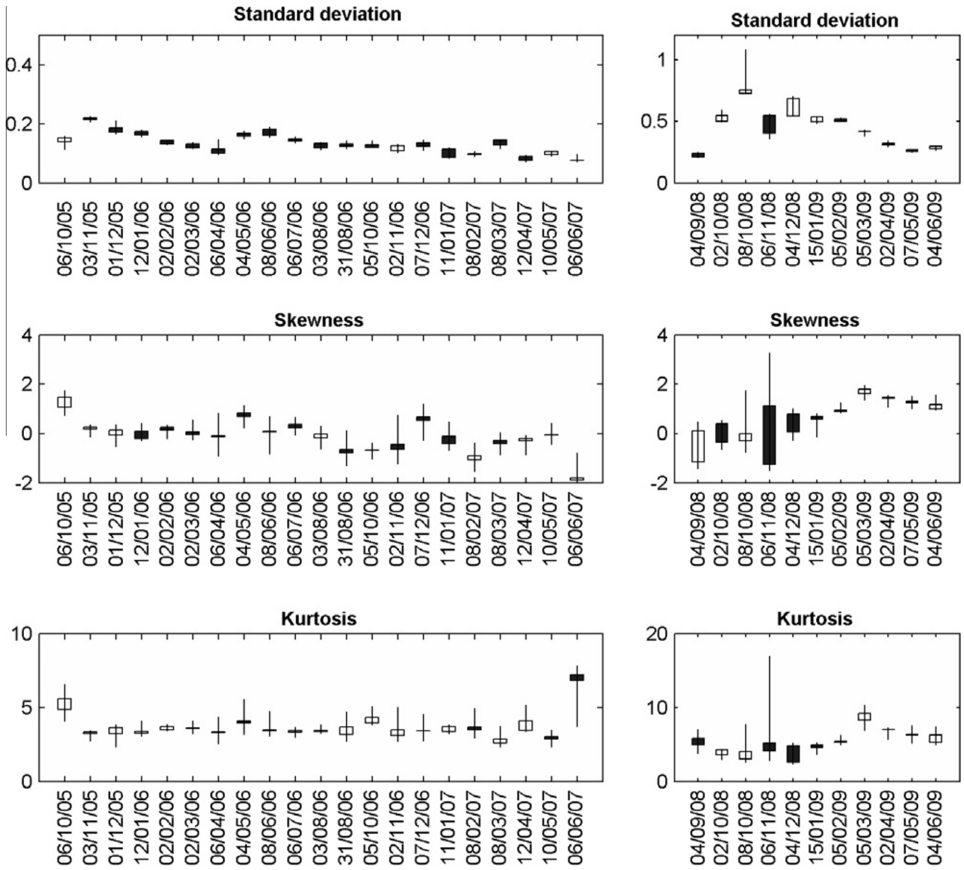


**Fig. 7.** Candle plots of moment developments within a day: sub-period 2005-2007 (lhs) and sub-period 2008–2009 (rhs).

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 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.

### *6.3. Persistence of the Governing Council impact*

While the previous sections showed how density reactions to a Governing Council meeting can be sizeable and persist within the day, this sub-section takes a closer look at the persistence of the reactions beyond the Governing Council day. If we assume that financial markets are efficient in aggregating relevant information, then all the reactions to news are permanent. However, if we do not want to impose this assumption, then there is the possibility that reactions are temporary and expectations revert to previous levels without news arriving.

The following exercise provides important indications about the persistence of the impact of central bank communication on Governing Council days on market expectations. As a long sample period is required for this type of analysis and no intraday dataset of such length is available, the exercise relies on daily densities estimated at close of business for each day of the sample period 1999–2011. In order to measure the persistence of density adjustments with precision, one would need to control for news arriving in the period between Governing Council meetings. However, collecting and controlling explicitly for news goes beyond the scope of this paper, importantly because a detailed news dataset is lacking that would allow such a study.

Nevertheless, the persistence of the density moments is telling as such. First, for each of the five statistical moments the change over the Governing Council day is taken, i.e. close of the Governing Council day minus close of the previous day, and next only the upper quartile of the observations is kept. This selection relies on the assumption that days with larger reactions correspond to days with Governing Council surprises. Second, a measure is computed that equals 100% if the statistical moment remains at the level reached directly after the Governing Council meeting or strengthened in the direction of the surprise, 0% in case the moment reverts completely to the level prior to the Governing Council meeting or beyond that, and a proportional value between 0% and 100% if

**Table 3**
Descriptive statistics of changes in statistical moments between Governing Council meetings.

| | Changes of density | | | | |
|---|---|---|---|---|---|
| | Mean | Median | St. dev. | Skew | Kurtosis |
| Sample average | −0.03 | −0.05 | −0.01 | 0.06 | 0.43 |
| Sample st. dev. | 0.20 | 0.24 | 0.07 | 0.52 | 1.93 |

**Table 4**
Persistence of Governing Council impact per statistical moment and time horizon.

| No. days | Statistical moments | | | | |
|---|---|---|---|---|---|
| | Mean | Median | St. dev. | Skew | Kurtosis |
| 1 | 90 | 89 | 90 | 49 | 47 |
| 5 | 83 | 85 | 65 | 50 | 51 |
| 10 | 74 | 74 | 56 | 57 | 52 |
| 27 | 71 | 69 | 62 | 48 | 53 |

the moment reached a level between those before and after the Governing Council meeting. As this exercise does not control explicitly for new information arriving after the Governing Council, it makes the assumption that such information cancels out on average over the sample. Table 3 presents descriptive statistics of the change in statistical moments between Governing Council meetings. As the sample average is close to zero, in particular in relation to the standard deviation of the changes, the assumption appears reasonable.

The results show that the Governing Council impact fades gradually for the mean, median and standard deviation over time, but fast for skewness and kurtosis. Table 4 shows the results taking time horizons of 1, 5 and 10 working days following the Governing Council and also until the day before the next Governing Council (which is a time horizon of 27 days on average). The results show that just before the next Governing Council, the impact on the mean, median and standard deviation of the previous meeting has not faded out and on average still contains 62-71% of the last Governing Council's impact, which appears strong given the caveats of the exercise. In other words, the impact on the average expectation, typical expectation, and the uncertainty about the future level of the interest rate is persistent according to the results. On the contrary, the impact on the skewness and kurtosis is already much smaller on the day following the Governing Council and remains around that level afterwards. The smaller degree of persistence of skewness and kurtosis is not surprising, because these moments may be seen as capturing how expectations of market participants are positioned *relative* to each other. A shock leading to 'disagreement' within the set of expectations held by market participants may be interesting to monitor as the shock unfolds, but the skew and kurtosis may also soon revert to their mean level, as expectations reposition around the (new) mean of the density.

## 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-min change of the corresponding moment. These activity measures are computed over three time intervals: the 10 min 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 two market reaction variables act as dependent variables in regressions of the type used in Ehrmann and Fratzscher (2009):

$$Y_{\mathrm{PR},m,t} = \alpha_{1,m,t} + \sum_{i=1}^{z_1} \beta_{1,i,m} x_{i,t} + \varepsilon_{1,m,t}$$

$$Y_{\mathrm{IS},m,t} = \alpha_{2,m,t} + \sum_{i=1}^{z_2} \beta_{2,i,m} x_{i,t} + \gamma_{1,m}\varepsilon_{1,m,t} + \varepsilon_{2,m,t},$$

$$Y_{\mathrm{QA},m,t} = \alpha_{3,m,t} + \sum_{i=1}^{z_3} \beta_{3,i,m} x_{i,t} + \gamma_{2,m}\varepsilon_{1,m,t} + \delta_{1,m}\varepsilon_{2,m,t} + \varepsilon_{3,m,t},$$

$$Z_{\mathrm{PR},m,t} = \alpha_{4,m,t} + \sum_{i=1}^{z_4} \beta_{4,i,m} x_{i,t} + \varepsilon_{1,m,t},$$

$$Z_{\mathrm{IS},m,t} = \alpha_{5,m,t} + \sum_{i=1}^{z_5} \beta_{5,i,m} x_{i,t} + \gamma_{3,m}\varepsilon_{1,m,t} + \varepsilon_{2,m,t}$$

$$Z_{\mathrm{QA},m,t} = \alpha_{6,m,t} + \sum_{i=1}^{z_6} \beta_{6,i,m} x_{i,t} + \gamma_{4,m}\varepsilon_{1,m,t} + \delta_{2,m}\varepsilon_{2,m,t} + \varepsilon_{3,m,t}$$

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 surrounding the decision:* measure for the direction of the uncertainty based on policy rate expectation 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 1000 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 1000 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.

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 3 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 Appendix E 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 Eqs. (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 Eqs. (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 (Eqs. (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 (Eqs. (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 (Eqs. (6)–(10)), and when allowing for direction only proofs significant for the total change in the median and skew (Eqs. (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 3 months ahead.

Higher uncertainty surrounding the 3-month EURIBOR 3 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 (Eqs. (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 3 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 (Eqs. (16)–(20)). It also no longer drives the total change in any of the moments during the introductory statement (Eqs. (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 (Eqs. (19) and (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 (Eqs. (12) and (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 (Eqs. (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 (Eqs. (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 (Eqs. (21) and (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 (Eqs. (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 (Eqs. (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, is applied to tick data on 3-month EURIBOR futures and options to estimate option-implied densities representing expectations of 3-month EURIBOR 3 months ahead. The paper discusses this estimator in an intraday setting and introduces an efficient method of pre-filtering the data to impose no-arbitrage conditions as required by option pricing theory. The density estimates indicate 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 succeeds in demonstrating the feasibility of intraday estimation.

An economic assessment of the announcement effects of ECB communication on short-term interest rate expectations is carried out. The sample consists 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 increase 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, the results indicate that there can be large movements in moments within a Governing Council day, in particular during the financial crisis period. Evidence based on daily densities also suggest that the impact of Governing Council surprises on the mean, median and standard deviation persists into the period following the meeting and is still present to an important degree at the start of the next Governing Council meeting.

Finally, 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, is 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 are also found to be relevant, but the evidence for this is not as strong. A code word, as perceived by the market in predicting rate hikes, is 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 results indicate that the overall content of the introductory statement and Q&A session are relevant drivers of expectations. While the sample size is 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

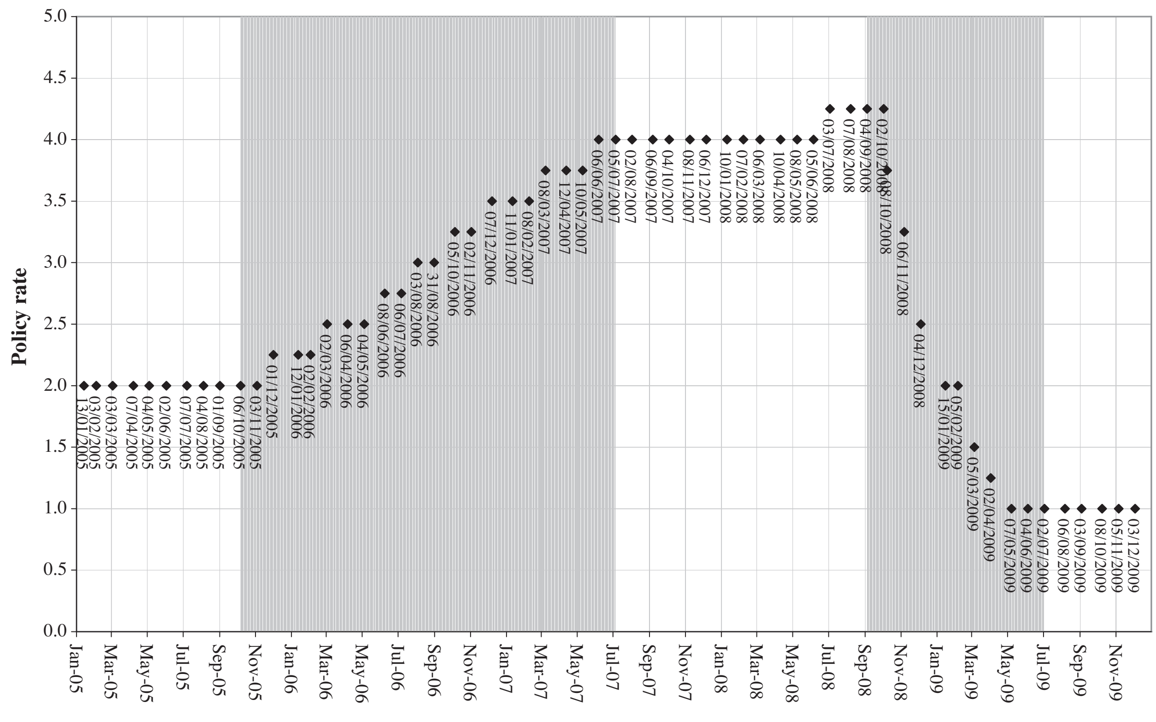


**Fig. A1.** Sample of Governing Council days.

**Table B1**
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*** | 0.3 | −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 3 months' time. The values for the mean, median and standard deviation are in basis points.

*** Significance at the 99% level for a *t*-test of difference in averages. 'Governing Council days' are tested against 'Other Thursdays'.

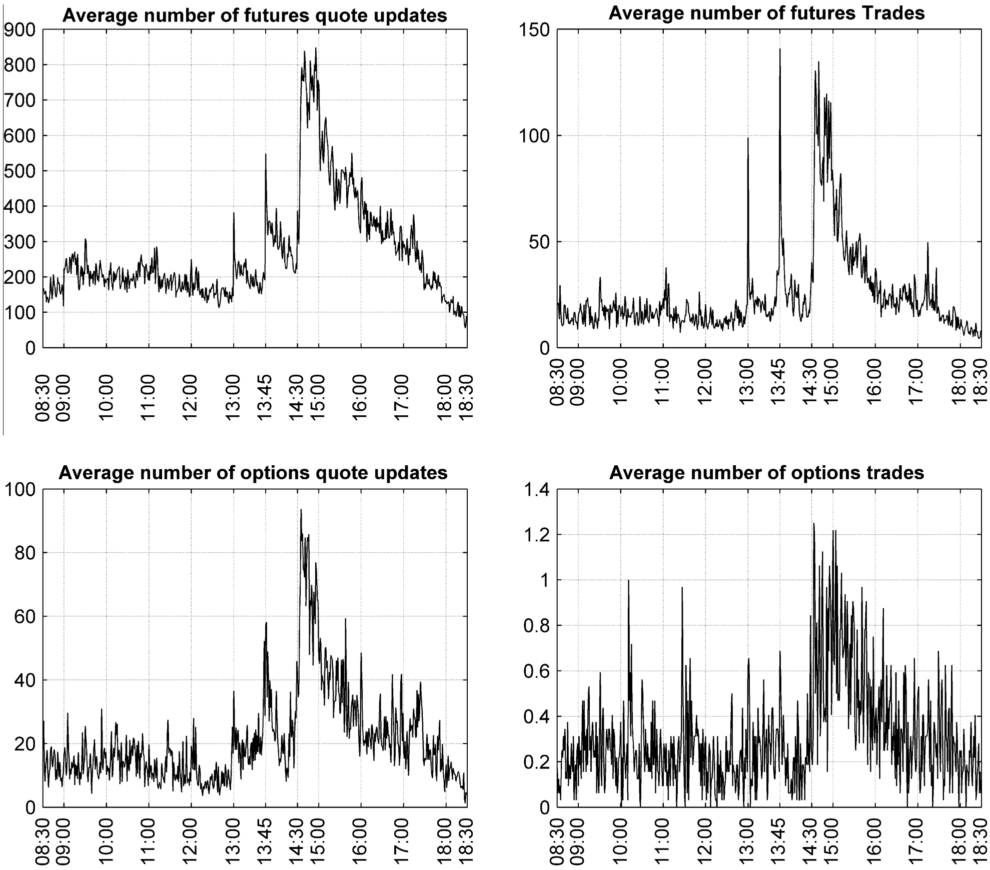


Fig. C1. Intraday pattern of tick activity (per minute).

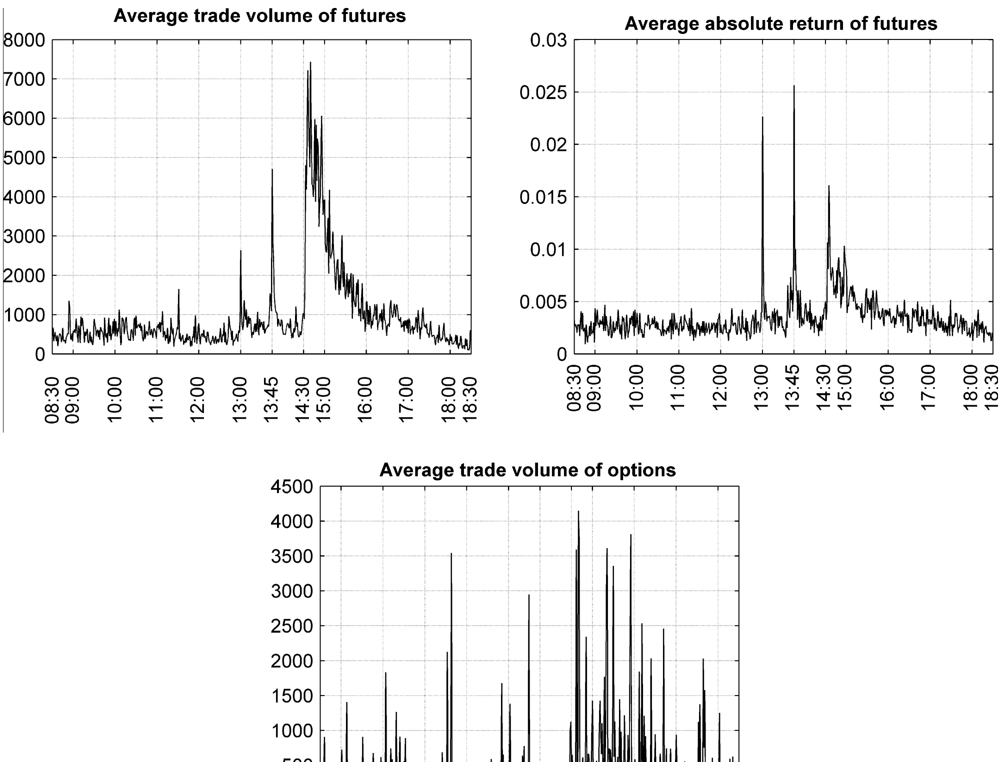


Fig. C2. Intraday pattern of tick activity (per minute).

may provide further statistical evidence of the impact of the content.

Overall, the relevance of the press release and conference as a communication tool is 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 holding a press conference following policy rate announcements. 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.

## Acknowledgments

We are grateful to Thomson Reuters for providing the tick data. We thank participants at the EC-squared conference, 2011 at the European University Institute (Florence), at the Macro and Financial Econometrics Conference 2011 at Heidelberg University, ICABE 2011 at the University of Piraeus, and EEFS 2011 at Queen Mary University of London for fruitful discussions. We are grateful to an anonymous referee for valuable comments. All remaining errors are ours.

## Appendix A. Sample of Governing Council days

(Sample in shaded area)
Fig. A1.

## Appendix B. Impact of Governing Council days on the density moments

Table B1.

## Appendix C. Intraday pattern of tick activity (per minute)

Figs. C1 and C2.

## Appendix D. Intraday pattern of implied density moments activity and level (per minute)

Figs. D1–D3.

## Appendix E. Regression results

Table E1.

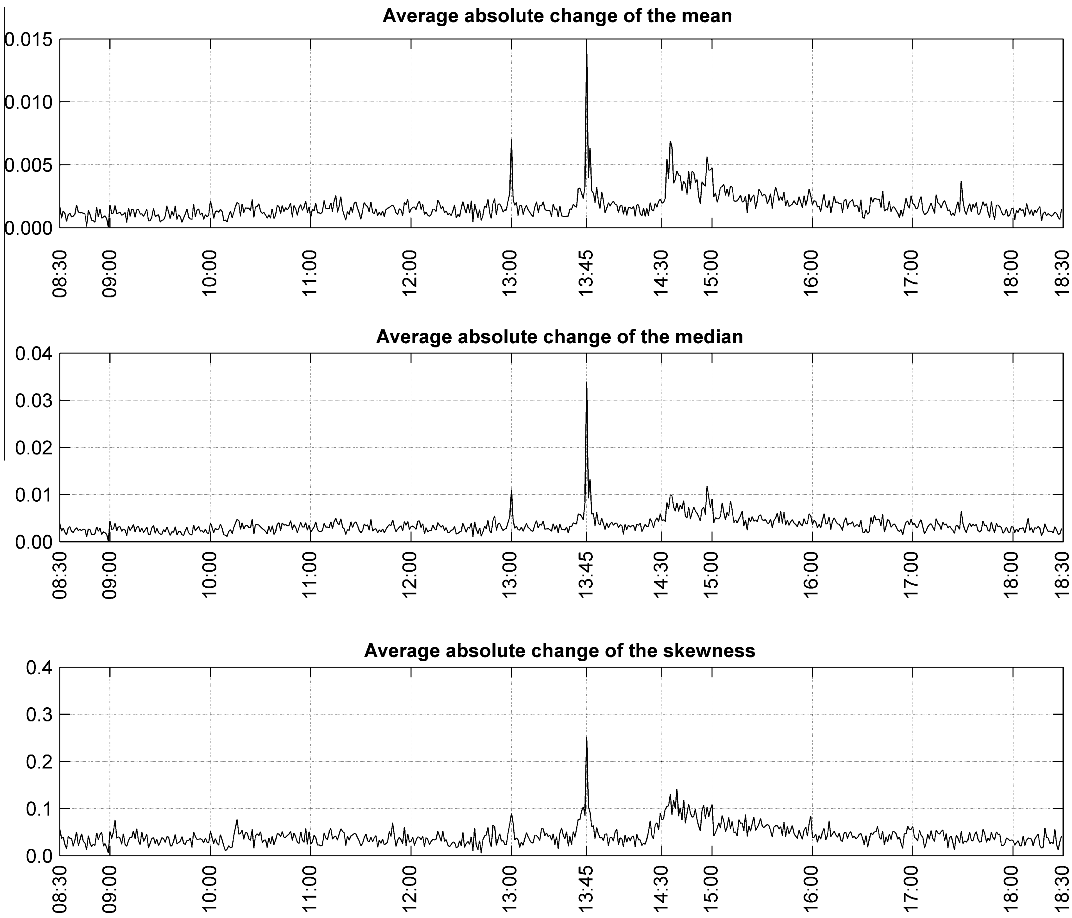


**Fig. D1.** Intraday pattern of implied density moments activity and level.

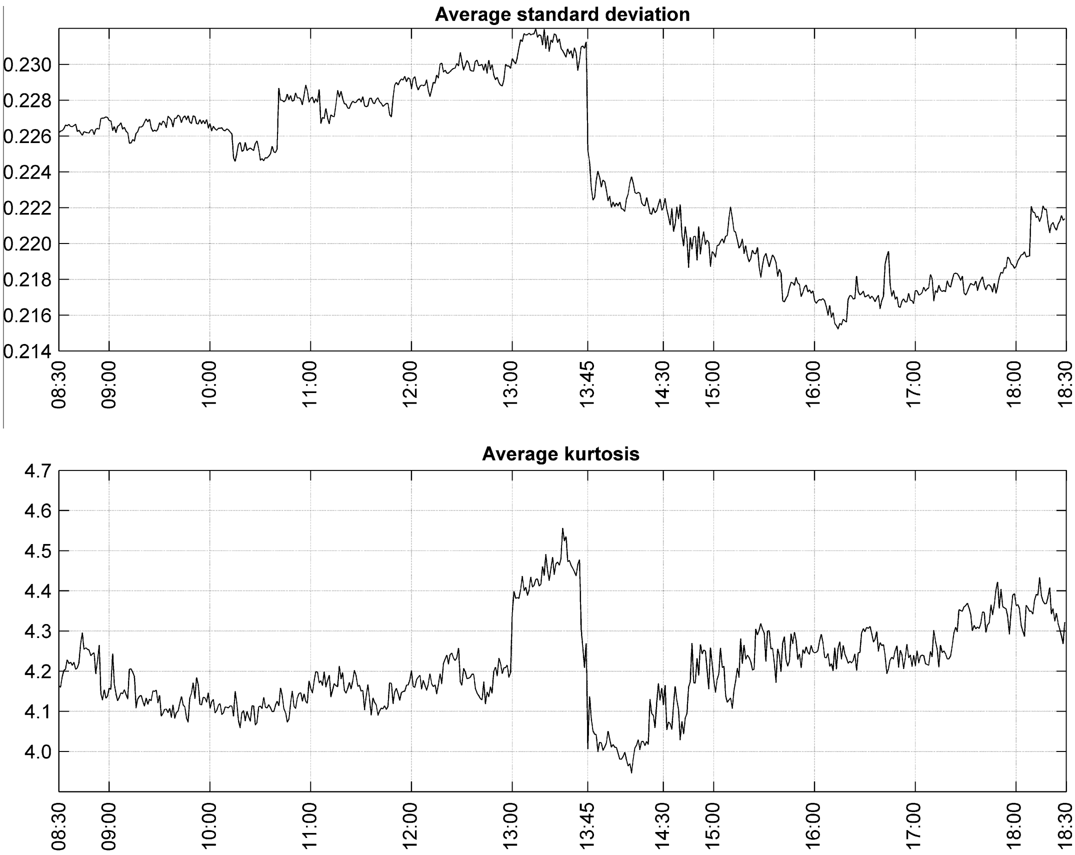


**Fig. D2.** Intraday pattern of implied density moments activity and level.

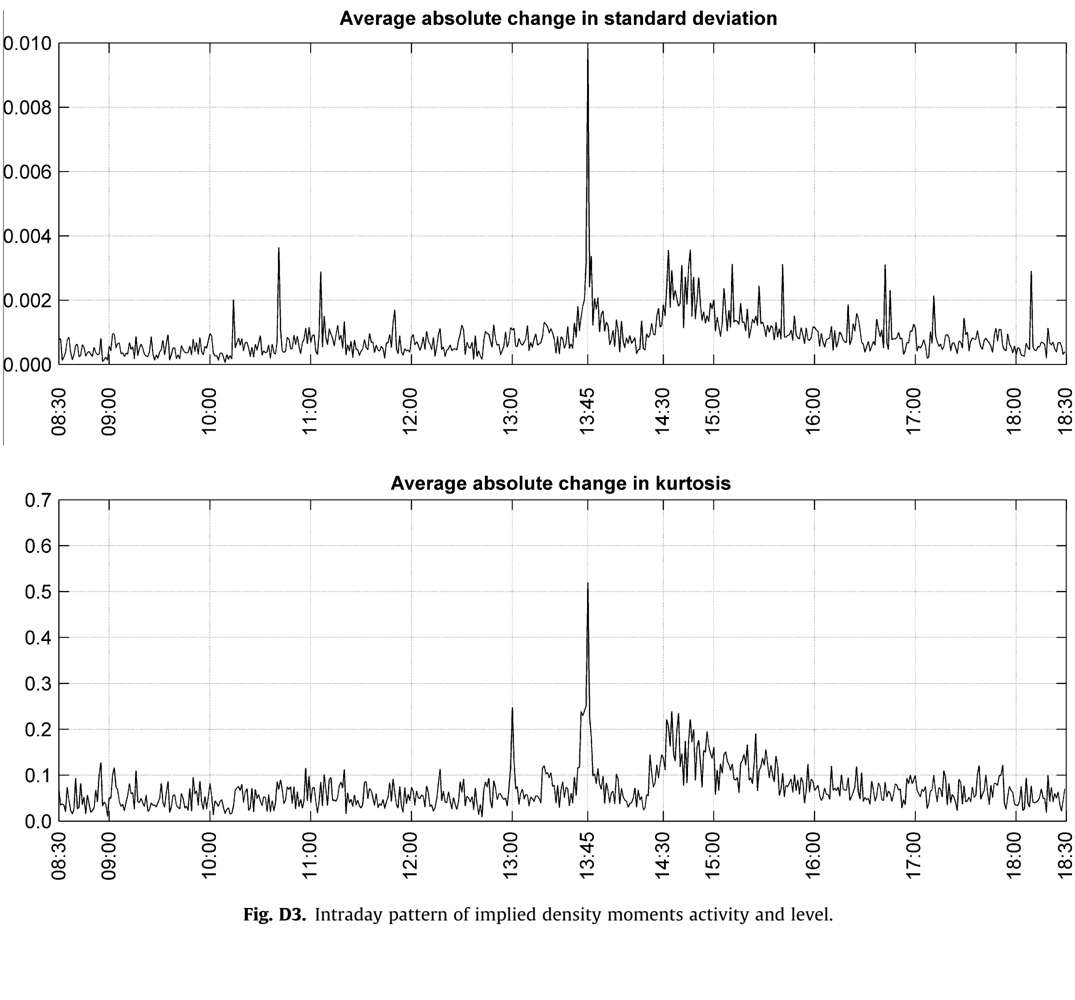


**Fig. D3.** Intraday pattern of implied density moments activity and level.

**Table E1**
Regression results.

| Press release | Dependent variable: $Y$ | | | | | | Dependent variable: $Z$ | | | | |
|---|---|---|---|---|---|---|---|---|---|---|---|
| | (1) | (2) | (3) | (4) | (5) | | (6) | (7) | (8) | (9) | (10) |
| | Mean | Median | St. dev.[a] | Skew | Kurtosis[a] | | Mean | Median | St. dev. | Skew | Kurtosis |
| Constant | 0.306 | 0.507 | −0.549 | −0.023 | 0.285 | Constant | 0.142 | 0.672 | 0.321** | 0.097*** | 0.045 |
| | (0.249) | (0.984) | (0.474) | (0.075) | (0.213) | | (0.177) | (0.410) | (0.128) | (0.023) | (0.045) |
| Decision | −0.035*** | −0.066* | 0.048 | 0.005 | −0.015 | \|Decision\| | 0.009*** | 0.017** | 0.009*** | 0.001 | 0.003 |
| | −0.014 | (0.038) | (0.058) | (0.004) | (0.010) | | (0.002) | (0.007) | (0.002) | (0.001) | (0.002) |
| Decision surprise | 0.347*** | 1.143*** | 0.441** | −0.056** | 0.041 | \|Decision surprise\| | 0.107*** | 0.263*** | 0.079*** | 0.010*** | 0.023*** |
| | (0.098) | (0.275) | (0.187) | (0.022) | (0.051) | | (0.013) | (0.042) | (0.011) | (0.003) | (0.006) |
| Decision uncertainty | −0.034 | −0.129** | −0.135 | 0.009* | −0.025 | \|Decision uncertainty\| | 0.010 | 0.020 | 0.001 | 0.001 | −0.001 |
| | −0.033 | (0.061) | (0.084) | (0.005) | (0.027) | | (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) |
| $R^2$ | 0.59 | 0.59 | 0.42 | 0.30 | 0.13 | $R^2$ | 0.79 | 0.69 | 0.79 | 0.37 | 0.48 |
| Initial statement | Dependent variable: $Y$ | | | | | | Dependent variable: $Z$ | | | | |
| | (11) | (12) | (13) | (14) | (15) | | (16) | (17) | (18) | (19) | (20) |
| | Mean | Median | St. dev.[b] | Skew | Kurtosis[b] | | Mean | Median | St. dev. | Skew | Kurtosis |
| Constant | 0.253 | 0.382 | −0.072 | −0.038 | 0.024 | Constant | 0.330* | 0.798** | 0.121 | 0.039 | −0.184* |
| | (0.347) | (0.462) | (0.300) | (0.061) | (0.093) | | (0.191) | (0.360) | (0.179) | (0.068) | (0.107) |
| Decision | 0.019 | −0.001 | −0.028*** | 0.004** | 0.003 | \|Decision\| | 0.003 | 0.002 | 0.002 | 0.001 | 0.002 |
| | (0.015) | (0.016) | (0.007) | (0.002) | (0.004) | | (0.002) | (0.003) | (0.002) | (0.000) | (0.002) |
| Decision surprise | −0.104 | −0.125* | −0.111 | 0.010 | −0.014 | \|Decision surprise\| | 0.046** | 0.068** | 0.029 | 0.003 | 0.010 |
| | (0.094) | (0.069) | (0.078) | (0.009) | (0.020) | | (0.021) | (0.033) | (0.022) | (0.002) | (0.009) |
| decision uncertainty | 0.013 | 0.010 | 0.023*** | 0.002 | 0.004 | \|Decision uncertainty\| | −0.002 | −0.006 | −0.002 | −0.002** | −0.003** |
| | (0.020) | (0.020) | (0.008) | (0.002) | (0.006) | | (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 | 0.004 | −0.014 | 0.000 | 0.007** | −0.003 | \|Initial claims\| | −0.002 | −0.003 | −0.003 | −0.001 | −0.001 |
| | (0.017) | (0.022) | (0.009) | (0.004) | (0.009) | | (0.001) | (0.004) | (0.003) | (0.001) | (0.003) |
| Continuing claims | −0.003 | −0.003 | 0.003 | −0.000 | −0.002* | \|Continuing claims\| | 0.000 | 0.002 | −0.001 | −0.000 | −0.001 |
| | (0.006) | (0.006) | (0.003) | (0.0006) | (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 | 2.377*** | 3.252*** | −0.195 | −0.293*** | 0.056 | \|Code word surprise\| | 0.101 | 0.091 | −0.066 | 0.053* | −0.002 |
| | (0.652) | (0.763) | (0.307) | (0.111) | (0.139) | | (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 | 0.115 | 0.084* | −0.086** | 0.036 | −0.008 | Resid. PR | −0.078 | −0.047 | 0.014 | 0.288 | 0.502 |
| | (0.162) | (0.047) | (0.039) | (0.067) | (0.030) | | (0.086) | (0.046) | (0.210) | (0.259) | (0.420) |
| $R^2$ | 0.32 | 0.35 | 0.32 | 0.24 | 0.11 | R2 | 0.53 | 0.43 | 0.39 | 0.49 | 0.43 |
| Q&A | Dependent variable: $Y$ | | | | | | Dependent variable: $Z$ | | | | |
| | (21) | (22) | (23) | (24) | (25) | | (26) | (27) | (28) | (29) | (30) |
| | Mean | Median | St. dev.[c] | Skew | Kurtosis[c] | | Mean | Median | St. dev. | Skew | Kurtosis |
| Constant | −0.126 | −0.081 | −0.476* | −0.086** | 0.152 | Constant | 0.107 | 0.346 | 0.198** | 0.054 | −0.092 |
| | (0.366) | (0.404) | (0.286) | (0.041) | (0.097) | | (0.093) | (0.261) | (0.084) | (0.033) | (0.124) |
| Decision | −0.062*** | −0.042*** | 0.038 | −0.003 | 0.013 | \|Decision\| | 0.002 | 0.003 | 0.003*** | 0.000 | 0.001 |
| | (0.012) | (0.012) | (0.024) | (0.002) | (0.012) | | (0.001) | (0.002) | (0.001) | (0.000) | (0.001) |
| Decision surprise | −0.385*** | −0.353*** | 0.263 | −0.010 | −0.022 | \|Decision surprise\| | 0.014*** | 0.050*** | 0.008*** | 0.001 | 0.002 |
| | (0.062) | (0.054) | (0.260) | (0.008) | (0.032) | | (0.005) | (0.014) | (0.002) | (0.001) | (0.003) |
| Decision uncertainty | −1.780 | −0.460 | −6.021 | −0.118 | −1.609 | \|Decision uncertainty\| | 0.281 | 0.321 | −0.063 | −0.086** | −0.295*** |
| | (1.351) | (1.387) | (4.040) | (0.172) | (1.335) | | (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 | 1.891 | 1.710 | −0.238 | 0.123∗ | 0.101 | \|Code word surprise\| | 0.074 | 0.064 | 0.119 | 0.026 | 0.075 |
| | (1.423) | (1.375) | (0.323) | (0.075) | (0.161) | | (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 | 0.345** | 0.018 | 0.159 | −0.024 | −0.062 | Resid. PR | −0.023 | −0.012 | 0.127*** | 0.165*** | 0.419*** |
| | (0.141) | (0.032) | (0.154) | (0.068) | (0.041) | | (0.089) | (0.051) | (0.047) | (0.061) | (0.106) |
| Resid. IS | 0.031 | −0.041 | −1.241* | −0.004 | −0.446 | Resid. IS | 0.246** | 0.252 | 0.384*** | 0.453*** | 0.643*** |
| | (0.207) | (0.187) | (0.665) | (0.111) | (0.394) | | (0.113) | (0.159) | (0.072) | (0.099) | (0.153) |
| $R^2$ | 0.60 | 0.49 | 0.37 | 0.11 | 0.22 | $R^2$ | 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.
* Significance at 90% level.
** Significance at 95% level.
*** Significance at 99% level.
[a] Regressed on absolute value of decision, decision surprise and decision uncertainty.
[b] Regressed on absolute value of decision, decision surprise, decision uncertainty, initial claims, continuing claims and code word surprise.
[c] Regressed on absolute value of decision, decision surprise, decision uncertainty and code word surprise.

## References


Andersen, A.B., Wagener, T., 2002. Extracting Risk Neutral Probability Densities by Fitting Implied Volatility Smiles: Some Methodological Points and Application to the 3M EURIBOR Futures Option Prices. ECB Working Paper 198.

Andersson, M., 2007. Using Intraday Data to Gauge Financial Market Responses to Fed and ECB Monetary Policy Decisions. ECB Working Paper 726.

Bahra, B., 1997. Implied Risk-neutral Probability Density Functions from Option Prices: Theory and Application. Bank of England Working Paper 66.

Black, F., 1976. The pricing of commodity contracts. Journal of Financial Economics 3, 167–179.

Black, F., Scholes, M.J., 1973. The pricing of options and corporate liabilities. Journal of Political Economy 81 (3), 637–654.

Bliss, R., Panigirtzoglou, N., 2002. Testing the stability of implied probability density functions. Journal of Banking and Finance 26, 381–422.

Brand, C., Bruncic D., Turunen J., 2006. The Impact of ECB Monetary Policy Decisions and Communication on the Yield Curve. ECB Working Paper 657.

Breeden, D., Litzenberger, R., 1978. Prices of state-contingent claims implicit in option prices. Journal of Business 51, 621–652.

Campa, J.M., Chang, P.H.K., Reider, R.L., 1997. ERM bandwidths for EMU and after: evidence from foreign exchange options. Economic Policy 24, 55–89.

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.

Cooper, N., 1999. Testing techniques for estimating implied RNDs from the prices of European-style options. In: Proceedings of the Workshop on Estimating and Interpreting Probability Density Functions held at the BIS on 14 June 1999.

Cox, J.C., Ross, S.A., 1976. The valuation of options for alternative stochastic processes. Journal of Financial Economics 3, 145–166.

ECB, 2007. Communicating Monetary Policy to Financial Markets. ECB Monthly Bulletin Article April.

ECB, 2011. The Information Content of Option Prices during the Financial Crisis. ECB Monthly Bulletin Article February.

Ehrmann, M., Fratzscher, M., 2009. Explaining monetary policy in press conferences. International Journal of Central Banking 5 (2), 41–84.

Lynch, D., Panigirtzoglou, N., Kapetanios, G., 2004. Using option prices to measure financial market views about balances of risk to future asset prices. Bank of England Quarterly Bulletin Winter, 442–454.

Malz, A.M., 1998. Option prices and the probability distribution of exchange rates. In: Zhaohui, C. (Ed.), Currency Options and Exchange Rate Economics. World Scientific, Singapur, pp. 108–137.

Mandler, M., 2002. Comparing Risk-neutral Probability Density Functions Implied by Option Prices – Market Uncertainty and ECB-Council Meetings. Working Paper, Universität Giessen.

NYSE Liffe, 2009. NYSE Liffe Trading Procedures. <www.euronext.com>.

Puigvert Gutiérrez, J.M., de Vincent-Humphreys, R., 2012. A quantitative mirror on the EURIBOR market using implied probability density functions. Eurasian Economic Review 2 (1), 1–31.

Shimko, D., 1993. Bounds of probability. Risk 6 (4), 33–37.

Zhang, L., Mykland, P.A., Aït-Sahalia, Y., 2005. A tale of two time scales: determining integrated volatility with noisy high frequency data. Journal of the American Statistical Association 472, 1394–1411.

## 9.3 Article 3: Interest rate forecasts, state price densities and risk premium from Euribor options

Journal of Banking & Finance 48 (2014) 210–223

Contents lists available at ScienceDirect

# Journal of Banking & Finance

journal homepage: www.elsevier.com/locate/jbf

# Interest rate forecasts, state price densities and risk premium from Euribor options ☆

Vesela Ivanova, Josep Maria Puigvert Gutiérrez *

*European Central Bank, Kaiserstrasse 29, D-60311 Frankfurt am Main, Germany*

ARTICLE INFO

*Article history:*
Received 25 June 2013
Accepted 18 March 2014
Available online 30 March 2014

*JEL classification:*
G12
G13
G17

*Keywords:*
Euribor futures
Forecast density
Interest rate premium
State price densities

ABSTRACT

In this paper we study option-implied interest rate forecasts and the development of risk premium and state prices in the Euribor futures option market. Using parametric and non-parametric statistical calibration, we transform the risk-neutral option implied densities for the Euribor futures rate into real-world densities. We investigate the period from the introduction of the Euro in 1999 until December 2012. The estimated densities are used to provide a measure for the interest rate risk premium and state prices implicit in the futures market. We find that the real-world option-implied distributions can be used to forecast the futures rate, while the forecasting ability of the risk-neutral distributions is rejected. The state price densities in the market show a U-shaped curve suggesting that investors price higher states with high and low rates compared to the expected spot rate. However, we show that, in general, state prices have a more pronounced right tail, implying that investors are more risk averse to increasing interest rates. We also document a negative market price of interest rate risk which generates positive premium for the futures contract.



## 1. Introduction

Derivative contracts are the main instruments used by investors to efficiently hedge risk and speculate on perceived market uncertainty. Therefore, they convey information on the likelihood that the market assigns to different future outcomes. As a result, option implied probability density functions are used by market participants and central bankers to estimate aggregated market expectations and uncertainty. In this study, we aim to improve forecasts and extract market fundamentals in the Euro Area interest rate derivatives market.

The introduction of a common currency in the Euro Area in 1999 created a new market for interest rate derivatives. Two of the most heavily traded products are 3-month Euribor futures and options on 3-month Euribor futures. Euribor is the rate at which Euro interbank deposits are offered from one bank to another within the European Monetary Union. As a result, Euribor futures and options are derivatives on the 3-month Euro denominated short-term interest rate. Fig. 1 shows that the trading volume of Euribor options increased more than tenfold in the current years of financial turmoils, which not only indicates their usefulness for hedging purposes but also implies valuable information for investors perceived risk in the market. This is why we use the history of Euribor option prices until the end of 2012 to produce reliable interest rate forecasts and to study the implied state price densities and risk premium in the market.

A vast amount of literature has already compared the forecasting ability of option prices, which are forward-looking by construction, against historical time series of asset prices. Among others Jiang and Tian (2005) and Martens and Zein (2004) show that option forecasts of index volatility perform better than historical forecasts, while Shackleton et al. (2010) and Liu et al. (2007) document the same for density forecasts. However, most of these studies are focused on the equity and index markets. We contribute to the literature by analyzing the information content of an interest rate derivatives market, which is heavily traded and unexplored.

We start our study by estimating the risk-neutral densities from the Euribor options prices and testing their ability to forecast the Euribor futures rate. A major disadvantage of the risk-neutral distribution, however, is that it does not incorporate risk premium,

☆ We would like to thank the two anonymous referees, Christian Schlag and the conference participants in "Financial Globalisation and Sustainable Finance: Implications for Policy and Practice", Cape Town, 2013 for the valuable comments. For correspondence, please contact any of the authors.

* Corresponding author. Tel.: +49 6913448660.

*E-mail addresses:* vesela.ivanova@ecb.europa.eu (V. Ivanova), josep_maria.puigvert@ecb.europa.eu (J.M. Puigvert Gutiérrez).

http://dx.doi.org/10.1016/j.jbankfin.2014.03.028

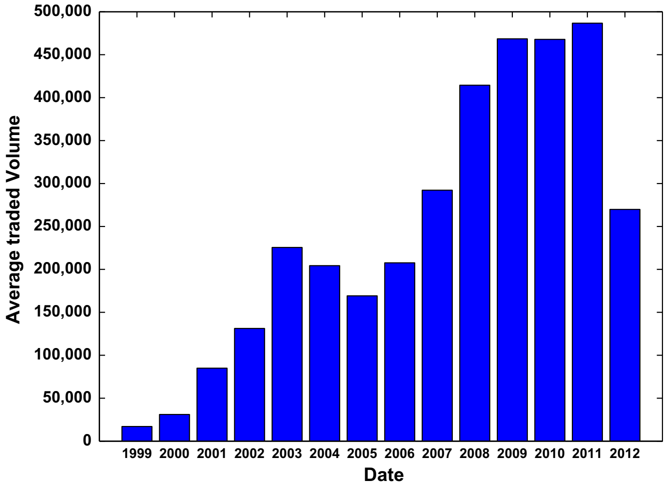


**Fig. 1.** Total trading of EURIBOR options. The figure shows the average number of transactions per year for the 3-month EURIBOR futures options.


**Fig. 2.** EURIBOR futures rate. The figure shows the development of the realized 3-month EURIBOR futures rate (in percentage) at the option maturity date over the sample period. The futures rate is defined as 100 minus the underlying futures price.

since it is a distorted probability measure whereby an option can be priced with the expectation of its discounted future values. Hence, it would be identical to the true probability measure if investors were indifferent to risk. Therefore, risk-neutral forecasts do not correspond to investors' actual forecasts.

We correct for this bias by estimating the real-world densities using statistical calibration as initially proposed by Fackler and King (1990). Similarly to Bliss and Panigirtzoglou (2004) and Shackleton et al. (2010) we analyze density forecasts for the option underlying futures prices and not directly the spot prices. The risk-neutral forecasts are constructed four weeks before option maturity and risk-adjusted with parametric and non-parametric calibration methods to obtain the real-world distributions. We then explore the market price of interest rate risk and the state prices in the Euribor futures market from its introduction in 1999 through the turbulent times of the recent financial crises.

We show that risk-neutral densities exhibit location and dispersion bias and their forecasting ability is rejected at a 5%-significance level. The forecasting power of the adjusted real-world densities, however, is strongly improved and cannot be rejected at significance level higher than 50%. We also document a U-shaped form of the state price densities (SPD) in the Euribor interest rate market, which shows that investors assign high prices to states of high and low rates compared to the expected spot rate. In this regard, the U-shaped curve has a more pronounced right tail, implying that investors are more risk-averse to increasing interest rates. We further investigate the development of the state prices during the current years of financial turmoil. We find that given the recently higher likelihood of recession, the SPD curve became more symmetric for a short period of time and investors risk perception to extreme interest rate changed. We also document an economically significant interest rate risk premium in the Euribor futures market, which is generated through a negative market price of interest rate risk.

Therefore, our study contributes to the literature devoted to understanding the market for short-term interest rates (STIR) and option-implied density forecasts. While most of the finance literature focuses its attention on the interest rate risk premium and pricing kernel in the bonds market,[1] those papers that analyze the premium in the interest rate futures market do not use the information content of option futures prices.[2] A similar study on option-implied densities from STIR contracts is the one by Li and Zhao (2009). The authors use interest rate caps and the locally polynomial estimator of Ait-Sahalia and Duarte (2003) for non-parametric estimation of the probability distribution of the LIBOR under the forward martingale measure. They also document a pronounced U-shape of the state price densities as a function of the LIBOR rates. However, the authors did not test the forecasting power of the estimated densities nor analyzed the premium in the market.

On the other hand, Shackleton et al. (2010) compared historical density forecasts for the S&P 500 index with option-implied risk-neutral and real-world densities and showed that option-based forecasts are superior for horizons of two, four or more weeks. Liu et al. (2007) used utility transformation and parametric calibration to risk-adjust the risk-neutral densities for the FTSE-100 index four weeks before maturity. They also showed that the transformed real-world densities have more explanatory power than the historical and risk-neutral densities. Similar studies have also been applied to different markets and option instruments like crude oil by Hog and Tsiaras (2011), Japanese Nikkei 225 index by Takkabutr (2012), FTSE 100 and 3-month LIBOR by De Vincent-Humphreys and Noss (2012), short-term interest rates by Cesari and Sevini (2004), Spanish IBEX 35 index by Alonso et al. (2006), Brazilian Real/US dollar by Ornelas et al. (2012) and S&P 500 index and British pound-US dollar by Anagnou-Basioudis et al. (2005).[3] None of these papers, however, studies or analyzes in detail pricing fundamentals in the respective markets. In this light, market risk aversion is studied by Bliss and Panigirtzoglou (2004) by adjusting the risk-neutral densities on S&P 500 and FTSE 100 indices with power and exponential utility functions. The authors also confirmed the superior forecasting ability of risk-adjusted option-implied forecasts at a four-weeks horizon.

The paper is structured as follows: Section 2 describes the data used for estimating the forecast densities. Section 3 explains the estimation method for the risk-neutral densities, Section 4 provides details on how the forecasting power is evaluated, while Section 5 presents the statistical calibration that is used to transform the risk-neutral to real-world densities. Section 6 elaborates on the main findings and the empirical results. Section 7 concludes.

[1] See Stanton (1997), Cochrane and Piazzesi (2005), and Fama and French (1993).
[2] See Hirscleifer (1988), Hess and Kamara (2003), and Grinblatt and Jegadeesh (1996).
[3] For other fundamental studies on option-implied distributions see Jackwerth and Rubinstein (1996) and Soderlind and Svensson (1997).

## 2. Data

The data refers to daily settlement prices of 3-month Euribor futures and prices on options written on the 3-month Euribor futures. The dataset covers the complete history of Euribor futures options until 2012, thus consisting of fourteen years of daily observations from January 1999 until December 2012.

The Euribor was created in 1999 following the introduction of the euro and is a daily reference benchmark based on the average rate at which prime banks in the Eurozone lend unsecured funding in the euro interbank market for a given period. Euribor is used as a reference rate for euro-denominated forward rate agreements, short term interest rate futures, option contracts and interest rate swaps, in very much the same way as the LIBOR is commonly used for Sterling and US dollar-denominated instruments. Thus they provide the basis for some of the world's most liquid and active interest rate markets.

The Euribor is currently constructed by a representative panel of 31 banks based on market criteria. The contributors to Euribor are the banks with the highest volume of business in the euro area money markets. These banks have been selected to ensure that the diversity of the euro money market is adequately reflected, thereby making Euribor an efficient and representative benchmark.[4] On 11 January 2013, the European Securities and Markets Authority (ESMA) and the European Banking Authority (EBA) published a joint set of recommendations made to Euribor-EBF aiming to improve the governance and transparency of the rate-setting process and to avoid fraudulent actions similar to the ones that occurred with the LIBOR.

Euribor options and their underlying asset, Euribor futures, are derivatives on the 3-month euro denominated short-term interest rate, which allow us to have a better picture of market expectations for the 3-month Euribor rate. The contracts are traded on the London International Financial Futures and Options Exchange (LIFFE). Futures daily settlement prices are published by Euronext.liffe,[5] formed in January 2002 from the takeover of LIFFE. According to LIFFE, these contracts were developed as a response to the Economic and Monetary Union in Europe, and the emergence of Euribor as the key cash market benchmark in Europe's money markets. Since its launch, LIFFE's Euribor contracts have come to dominate the euro-denominated STIR derivatives market, capturing over 99% of its volume. They are now the most liquid and heavily traded euro-denominated STIR contracts in the world. The number of traded contracts has increased steadily since this instrument was first introduced, with the highest trading in the most recent years, as presented in Fig. 1.

A 3-month Euribor future allows an investor to lock the interest rate on one million euro for a future period of 3 months. The contracts are delivered quarterly in March, June, September and December; their last trading day is two business days prior to the third Wednesday of the delivery month, while the delivery day is the first business day after the last trading day. A call option written on a 3-month Euribor future gives the holder the right, but not the obligation, to buy the underlying futures contract. On any given day twelve option contracts with fixed expiry date are traded. The first five options expire on the following five closest months from the respective day, while the remaining contracts expire in the next quarters.[6] Due to the quarterly delivery structure of the future contracts, the options are settled with the assignment of a futures contract at the exercise price and with the respective quarterly delivery. For example, the futures delivery month associated with options expiring in January, February and March is March, while for those expiring in April, May and June is June.

An important characteristic of the Euribor futures options is that the option price is not paid at the time of purchase, since the option positions are marked-to-market in a futures-style margining. As a result, the buyer and seller of the option pay or receive valuation margins according to the daily changes in the option price.[7] This characteristic has several important implications that simplify the structure of the Euribor futures options, and as such also the model complexity. First, for American style options, which is also the case here, the premature exercising is never optimal. As shown by Lieu (1990) and Chen and Scott (1993), due to the margining, there is no opportunity cost for holding the option, which implies that American options can be priced as standard European style options. This result stems from the fact that, as shown by Duffie (1989), the future-style margining turns an option on a futures contract into a futures contract on an option. Second, Chen and Scott (1993) also proved that there is no discounting in the pricing of futures-style margined options, even in a general equilibrium setting with stochastic interest rates.

For the empirical analysis, we choose to work with out-of-the-money (OTM) options, as they are known to be more liquid. In this regard, De Vincent-Humphreys and Puigvert Gutiérrez (2012) analysed the trading volume for all Euribor options from 13 January 1999 until April 2010 and showed that 81% of the options are traded OTM, whereas only 18% are traded in-the-money (ITM). Furthermore, some of the ITM options are not traded independently, but as part of a bundled trading strategy, e.g. straddles or strangles. As a result, we estimate risk-neutral and real-world probability density functions (PDF) using option contracts which were either at- or out-of-the money. To construct the forecast densities, we concentrate on options expiring in between 3 and 4 weeks, as later elaborated in the paper. OTM options are also the most liquid for this time to expiration as shown in Table 1.

In addition, we apply three other types of quality data checks to the option price data. First, as a basic plausibility check, we delete any option prices that are either zero or negative. The second check is motivated by the option-pricing theory, whereby a call price function should be both monotonic and convex in order to yield non-negative probability estimates. In practice, this may not be the case if the difference between the "true" price of options with adjacent strikes is less than the minimum tick size, or if there are sufficiently large variations in the bid-ask spread. Hence, option prices that do not meet the monotonicity and convexity requirements are also excluded. Finally, if after the application of the preceding two filters there are less than three OTM option prices for a particular expiration date, no PDF will be estimated for that expiration date.

## 3. Risk-neutral densities

The approach applied to the study of forecasts on the futures rate and the associated market fundamentals involves first the

[4] Every panel bank is required to provide daily quotes of the rate that each panel bank believes one prime bank is quoting to another prime bank for interbank term deposits within the euro area, for a maturity ranging from one week to one year. The published rate is a rounded, truncated mean of the quoted rates: the highest and lowest 15% of quotes are eliminated, the remainder are averaged and the result is rounded to 3 decimal places. Euribor rates are spot rates, i.e. for a start two working days after measurement day. Like US money-market rates, they are Actual/360, i.e. calculated with an exact day count over a 360-day year.

[5] Data for the last ten days can be downloaded directly from the Internet via the following link http://www.liffe.com/reports/eod?item=Histories.

[6] For instance, on the 2nd of January 2012 five option contracts were expiring in January, February, March, April and May of 2012. The remaining contracts were expiring in June, September, December 2012 and March, June, September and December 2013.

[7] Note that under the standard option system, the option buyer must pay the option price when the transaction is initiated. With a futures-style margining of the options, the buyer and the seller are required to put up a margin and the accounts are marked to market daily, during the life of the option, as the settlement price of the option changes.

**Table 1**
Summary statistics of option trading volumes. The table presents the number of transactions made for option contracts expiring in 3–4 weeks. The table is based on a breakdown of call and put options, and then further of at-the-money, in-the-money and out-of-the-money contracts. We report the total traded volume as well as the mean and standard deviation of the number of transactions for the respective year.

| | | CALL | | | PUT | | |
|---|---|---|---|---|---|---|---|
| | | At-the-money | In-the-money | Out-of-the-money | At-the-money | In-the-money | Out-of-the-money |
| 1999 | Volume Traded | 17,313 | 10,198 | 26,555 | 6453 | 50 | 7969 |
| | Mean | 279 | 164 | 428 | 104 | 1 | 129 |
| | Std. Dev. | 1141 | 889 | 1480 | 370 | 6 | 511 |
| 2000 | Volume Traded | 28,166 | 15,443 | 113,235 | 36,369 | 14,195 | 57,296 |
| | Mean | 486 | 266 | 1952 | 627 | 245 | 988 |
| | Std. Dev. | 1130 | 824 | 3488 | 1765 | 584 | 2241 |
| 2001 | Volume Traded | 27,268 | 78,375 | 353,259 | 27,644 | 19,092 | 158,927 |
| | Mean | 440 | 1264 | 5698 | 446 | 308 | 2563 |
| | Std. Dev. | 1339 | 3020 | 7732 | 1397 | 693 | 4781 |
| 2002 | Volume Traded | 88,989 | 260,722 | 676,381 | 29,439 | 46,652 | 224,133 |
| | Mean | 1369 | 4011 | 10,406 | 453 | 718 | 3448 |
| | Std. Dev. | 4403 | 10,492 | 16,639 | 1174 | 1729 | 6530 |
| 2003 | Volume Traded | 49,059 | 210,243 | 457,229 | 10,653 | 17,978 | 208,474 |
| | Mean | 791 | 3391 | 7375 | 172 | 290 | 3362 |
| | Std. Dev. | 2076 | 11,766 | 14,942 | 596 | 1010 | 7929 |
| 2004 | Volume Traded | 53,298 | 158,131 | 470,926 | 6260 | 10,200 | 123,011 |
| | Mean | 833 | 2471 | 7358 | 98 | 159 | 1922 |
| | Std. Dev. | 3368 | 7414 | 17,078 | 323 | 499 | 4958 |
| 2005 | Volume Traded | 192,289 | 3846 | 87,963 | 60,614 | 18,213 | 158,669 |
| | Mean | 3101 | 62 | 1419 | 978 | 294 | 2559 |
| | Std. Dev. | 9448 | 213 | 3819 | 3678 | 1265 | 9871 |
| 2006 | Volume Traded | 32,393 | 7782 | 164,800 | 71,751 | 23,713 | 166,573 |
| | Mean | 522 | 126 | 2658 | 1157 | 382 | 2687 |
| | Std. Dev. | 1278 | 474 | 5305 | 4576 | 1072 | 6273 |
| 2007 | Volume Traded | 111,685 | 31,128 | 612,253 | 238,150 | 117,688 | 564,855 |
| | Mean | 1745 | 486 | 9566 | 3721 | 1839 | 8826 |
| | Std. Dev. | 5749 | 2295 | 23,992 | 9765 | 5886 | 24,404 |
| 2008 | Volume Traded | 154,436 | 309,507 | 1,537,408 | 102,462 | 75,500 | 735,234 |
| | Mean | 2376 | 4762 | 23,652 | 1576 | 1162 | 11,311 |
| | Std. Dev. | 6103 | 18,927 | 50,919 | 3649 | 2978 | 19,681 |
| 2009 | Volume Traded | 271,405 | 563,109 | 1,890,354 | 144,258 | 71,090 | 888,011 |
| | Mean | 4378 | 9082 | 30,490 | 2327 | 1147 | 14,323 |
| | Std. Dev. | 9141 | 16,709 | 57,550 | 8012 | 4466 | 28,326 |
| 2010 | Volume Traded | 181,927 | 153,106 | 440,951 | 153,250 | 129,315 | 780,053 |
| | Mean | 2888 | 2430 | 6999 | 2433 | 2053 | 12,382 |
| | Std. Dev. | 7774 | 5125 | 9157 | 5178 | 7221 | 24,378 |
| 2011 | Volume Traded | 311,808 | 286,565 | 2,464,451 | 111,946 | 132,245 | 1,829,871 |
| | Mean | 4872 | 4478 | 38,507 | 1749 | 2066 | 28,592 |
| | Std. Dev. | 11,914 | 13,792 | 62,690 | 3886 | 4693 | 49,167 |
| 2012 | Volume Traded | 267,351 | 252,665 | 505,181 | 76,335 | 23,839 | 452,902 |
| | Mean | 4051 | 3828 | 7654 | 1157 | 361 | 6862 |
| | Std. Dev. | 7987 | 5787 | 10,198 | 2981 | 1203 | 13,032 |

extraction of the risk-neutral distribution (RND) from option prices. The price of a European call option, $C_t$, at time $t$ written on a Euribor futures contract with price $P_t$ is the expectation under the RND, $f(P_T)$, of its future option values

$$C_t(P_t, K, \tau) = E_t^{\mathbb{Q}}[(P_T - K)^+] = \int_K^{\infty} f(P_T)(P_T - K)dP_T, \quad (1)$$

where $K$ is the strike price, $T$ is the option expiry date, and $\tau$ is the remaining time to maturity. As mentioned before, due to the futures-style margining of the Euribor futures options, the standard discounting function of the option future values drops out from (1). The result has been proven by Chen and Scott (1993) to hold also in a general equilibrium setting with stochastic interest rates. For a brief presentation, recall that Cox et al. (1981) showed that with stochastic interest rates, the price of a general futures contract is equal to

$$P(t,T) = E_t^{\mathbb{Q}}\left[\exp\left(-\int_t^T R_u du\right)\exp\left(\int_t^T R_u du\right)S(T)\right] = E_t^{\mathbb{Q}}[S(T)],$$

where $S(T)$ is the spot price and $R$ denotes the instantaneous interest rate. Therefore, the futures price $P$ at time $t$ is the expectation of the terminal spot value under the risk-adjusted measure, while the risk-adjustment involves substraction of the risk premium from the mean of the underlying process. As the Euribor futures option is equivalent to a futures contract on the option, due to the futures-style margining, the value of the option is simply the expected value, under the RND, of its terminal payoff as shown in (1).[8]

As shown by Breeden and Litzenberger (1978), the extraction of the implicit risk-neutral probabilities, $f(P_T)$, for the value of the option underlying asset at maturity can be inferred from the second partial derivative of the option price function with respect to the strike price.

$$\frac{\partial^2 C_t(P_t, K, \tau)}{\partial^2 K} = f(P_T). \quad (2)$$

[8] See Lieu (1990) and Chen and Scott (1993) for further technical proofs on the implication of the futures-style margining of option contracts.

So in practice, the task of estimating the RND amounts to estimating a twice-differential option price function. However, Eq. (2) cannot be applied directly to obtain $f(P_T)$, because we only observe option prices for a discrete set of strike prices, and not a twice-differentiable continuum. Moreover, estimating directly the second derivative of a call price function can also lead to unstable or inaccurate density. Instead, Bliss and Panigirtzoglou (2002), following the results derived from Shimko (1993) and Malz (1997), have suggested that smoother results can be obtained if, prior to the interpolation, *option prices* and *strike prices* are transformed into implied *volatility* and *delta* values.

To derive the RND, we use the non-parametric approach developed by Cooper (1999) and Bliss and Panigirtzoglou (2002). On a given day $t$ of our sample, we collect the option prices for all traded strike prices $K_i,\ i = 1, \dots, m$. The observed strike prices are then transformed into option implied volatilities by numerically solving the Black (1976) futures options pricing model for the value of $\sigma_i$.

$$C_t(P_t, K_i, \tau) = P_t \Phi\left(\frac{\ln(\frac{P_t}{K_i}) + \frac{\sigma_i^2}{2}\tau}{\sigma_i\sqrt{\tau}}\right) - K_i \Phi\left(\frac{\ln\left(\frac{P_t}{K_i}\right) - \frac{\sigma_i^2}{2}\tau}{\sigma_i\sqrt{\tau}}\right). \tag{3}$$

As a second step, the obtained *implied volatilities* derived from (3) are used to calculate the option *delta* values using the formula below:

$$\delta_i(P_t, K_i, \tau) = \frac{\partial C(P_t, K_i, \tau)}{\partial P_t} = \Phi\left(\frac{\ln(\frac{P_t}{K_i}) + \frac{\sigma_i^2}{2}\tau}{\sigma_i\sqrt{\tau}}\right), \tag{4}$$

where $\Phi$ is the standard normal cumulative distribution function (CDF).

As a result, the observed option and strike prices are converted to raw data points in the implied volatility-delta space. The raw data is then interpolated, as in Campa et al. (1997), using a cubic smoothing spline, which minimizes the following function:

$$\lambda \sum_{i=1}^{n} \omega_i (\sigma_i - g(\delta_i))^2 + (1 - \lambda) \int g^2(t)^2 dt, \tag{5}$$

where $\lambda$ is a smoothing roughness parameter, equal to 0.99,[9] and the weights $\omega_i$ are calculated using $\omega_i = \frac{v_i^2}{mean(v_i^2)}$ where $v_i$ is the option sensitivity to volatility, known as *vega*. The value of *vega* is almost negligible for options which are deep out-of-the-money and deep in-the-money and sequentially increases as we get near-the-money. In particular, it reaches a maximum for at-the-money options. Hence, the $\omega_i$ used in (5) places most weight on near-the-money options. This is consistent with using these PDFs to support monetary policy analysis, where interest is likely to lie in the center of the distribution, i.e. close to the underlying interest rate, rather than in the distribution's tails. However, although *delta* can take values between 0 and 1, the traded contracts may not span that complete range. Therefore, the smoothing spline is further extrapolated outside the range of traded price points with a second order polynomial. As a result, the piecewise cubic curve obtained using the interpolation is extended with a quadratic curve at each endpoint so, that the full *delta* range is covered.

In the last step, the interpolated volatility smile is transformed back from the *volatility-delta* space to *strike price* and *option price* values. This is done by evaluating the interpolated volatility smile at 3000 equally-spaced delta points, $\delta_j,\ j = 1, \dots, 3000$, between zero and one. The delta values are then transformed back into strike prices $K_j$ by spacing the moneyness, with respect to the current future price, evenly in the delta space with

$$\frac{K_j}{P_t} = \exp\left(\left(\frac{\sigma_j^2}{2}\tau\right) - \sigma_j\sqrt{\tau}\Phi^{-1}(\delta_j)\right), \tag{6}$$

where $\Phi^{-1}$ is the inverse of the standard normal cumulative density function. The implied volatility values of the spline are translated back into call prices using the Black (1976) option pricing equation in (3). Finally, to obtain the risk-neutral PDF, we calculate the second derivative of the option price function by fitting cubic polynomials through triplets of consecutive *strike price-call price* pairs and evaluating the second derivative from the coefficients of fitted polynomials.

Given that the object of interest in this paper is the *futures rate*, we estimate directly the RND for the futures rate $r_t^F$ by using the fact that the futures price $P_t$ is quoted as $P_t = 100 - r_t^F$.[10] We can then re-write (1) into

$$C_t(r_t^F, r^K, \tau) = \int_{r^K}^{\infty} f(r_T^F)(r_T^F - r^K) dr_T^F, \tag{7}$$

so that $f(r_T^F)$ is the implied probability density function, which describes the likelihood of all perceived possible out-turns of the underlying futures rate at time $T$. We extract the option-implied RND by following the approach described above.[11]

## 4. Forecasting tests

The probability density estimated at time $t$, $f_t(r_T^F)$, denotes the density forecast for the option underlying futures rate at option maturity $T$. We choose to estimate the forecast densities four weeks before the option maturity. Therefore, with a four-weeks forecast horizon $h$, the option-implied densities are estimated at time $t = T - h$, over all maturity dates in our sample, $T_i,\ i = 1, \dots, N$. If the densities could not be estimated on the target observation date, $t$, due to a lack of sufficient trading, we use the density at the nearest trading date, if this is not more than three days before or after the target date.

The main reason for choosing to work with a four-weeks horizon is that our sample consists of options with a monthly maturity, and hence densities estimated with a longer forecast horizon than a month will contain overlapping information.[12] Therefore, a four-week horizon avoids serial dependence arising from overlapping data, which is crucial for the tests we perform later. On the other hand, option trading shortly before the maturity date is significantly lower; thus, using a shorter horizon would reduce our sample size and may be subject to an illiquidity bias.

To test whether the estimated densities correctly capture the distribution of the ex-post futures rate, we use the probability integral transform (PIT) introduced by Diebold et al. (1998) as a major constructive element. The PIT, $y_t^i$, represents the cumulative probability, $F_t(\cdot)$, of the estimated forecast density $f_t(\cdot)$ given the ex-post realization of the futures rate at option maturity $r_{T_i}^F$

$$y_t^i = \int_{-\infty}^{r_{T_i}^F} \hat{f}_{t_i}(u) du = F_{t_i}(r_{T_i}^F). \tag{8}$$

[9] The optimal smoothing roughness parameter is the one that minimizes the observed deltas with the fitted deltas by the smoothing spline.

[10] In this respect, it should be noted that $P_t$ is not bounded at 100, as negative market rates can also be traded in the market. The issue of negative interest rates is postponed for Section 6 of the paper.

[11] Note that the transformation of a call on the futures contract into a call on the explicit futures rate with $r^K = 100 - K$ leads to a put option according to market definition, i.e. a futures call option is equivalent to a put on the futures rate. This is accounted for in the extraction of the density.

[12] For example, consider a forecast horizon of two months, $h = 2m$, and two options maturing in March and April, $T = Mar, Apr$. In this case, the path between the forecast dates $t$ and the realization dates $T$ will overlap and hence the two forecast densities will contain common information about the innovations in the futures rate development.

Note that to compute the forecast CDF in (8), the lower limit of the integral is minus infinity, as we also allow for negative interest rates, following the option trading in 2012, also in the range of negative Euribor rates.

To highlight the rationale behind the use of the PIT here, let $Y = F(r^F)$ be a random variable with realizations the PITs described in (8) and with a cumulative density function denoted as $Q(\cdot)$. As shown by Rosenblatt (1952) and Diebold et al. (1998), only when the forecast CDF, $F(\cdot)$, is correctly estimated and hence identical with the CDF of the true data-generating process, $\overline{F}(\cdot)$, the random variable $Y$ will be independently and uniformly distributed:

$$Q(y) = Prob(Y \leqslant y) = Prob(F(r_T^F) \leqslant y) = Prob(r_T^F \leqslant F^{-1}(y)) \\ = \overline{F}(F^{-1}(y)) = y. \tag{9}$$

Therefore, the forecasting ability is tested with the null hypothesis that the sequence of estimated density forecasts coincides with the true density, i.e. $\bar{f}_t(\cdot) \equiv f_t(\cdot)$. Under this hypothesis, the series of the PIT must be independently and uniformly distributed between 0 and 1. We obtain a total set of 127 PIT observations for the complete sample of available Euribor options.[13] To test the null hypothesis, we use the parametric joint test developed by Berkowitz (2001) and the direct non-parametric Kolmogorov–Smirnov test to check the uniformity of the PIT series.

The Berkowitz test is based on a simple transformation of the PIT to normality

$$z_t = \Phi^{-1}(y_t),$$

where $\Phi^{-1}(\cdot)$ is the inverse of the standard normal CDF. Under the null hypothesis, the series $z_t$ must be *iid* standard normal distributed. The hypothesis is tested against the first-order autoregressive model

$$z_t - \mu = \rho(z_{t-1} - \mu) + \epsilon_t$$

with mean and variance different from $(0, 1)$. The model parameters are estimated using the maximum likelihood approach. As a result, the joint hypothesis for normality and independence translates to $\mu = 0, \rho = 0$ and $Var(\epsilon) = 1$ with the statistic

$$LR_3 = -2\left[L(0,1,0) - L(\hat{\mu}, \hat{\sigma}^2, \hat{\rho})\right],$$

where $L(\mu, \sigma^2, \rho)$ denotes the likelihood ratio; under the null hypothesis, the test has a $\chi^2(3)$ distribution. A single check for the independence assumption can also be performed with the likelihood ratio statistic

$$LR_1 = -2\left[L(\hat{\mu}, \hat{\sigma}^2, 0) - L(\hat{\mu}, \hat{\sigma}^2, \hat{\rho})\right],$$

which is $\chi^2(1)$ distributed. The test results are to be read as follows: if $LR_3$ rejects normality, failure to reject independence provides evidence that the densities do not possess forecasting power; if both $LR_1$ and $LR_3$ reject the hypothesis, then no constructive conclusion can be drawn; however, failure to reject both tests evidences that the estimated densities have forecasting power.

Bliss and Panigirtzoglou (2004) show that the Berkowitz test is more powerful than the Kolmogorov–Smirnov test in small samples, because it is a joint test of uniformity and independence. This is why we focus our discussion on the results of the Berkowitz test and we report the Kolmogorov–Smirnov results for completeness and robustness.

## 5. Real-world densities

A major disadvantage of option implied distributions is that they are risk-neutral and hence do not incorporate risk premium. However, if investors are not risk-neutral, in the presence of market uncertainty they will require premium for bearing risk. As a result, the risk-neutral distribution assigns greater weight to bad states than the true physical density.[14] The discrepancy between the two densities has warranted further research on the forecast accuracy of the option-implied distribution and its transformation into the true density.

We use statistical calibration, as introduced by Bunn (1984) and Diebold et al. (1999), to convert the risk-neutral densities to their real-world counterparts. The objective of statistical calibration is to transform forecast densities to probability assessment methods that generate reliable forecast statements. Hence, this method is a convenient way to risk-adjust the option-implied risk-neutral density to a real-world density, as described by Fackler and King (1990).

The transformation is based on the cumulative density function $Q(\cdot)$ of the variable $Y$, whose realizations are the PITs in (8). Let $F(\cdot)$ and $G(\cdot)$ denote respectively the cumulative distributions of the risk-neutral (RN) and the real-world (RW) functions. As shown below, the CDF of the real-world density can be expressed as a function of the RND CDF.[15]

$$G(r_T^F) = Prob(r^F \leqslant r_T^F) = Prob(F(r^F) \leqslant F(r_T^F)) \\ = Prob(Y \leqslant F(r_T^F)) = Q(F(r_T^F)). \tag{10}$$

Therefore, the function $Q(\cdot)$ is the calibration function that provides the mean for transforming the RND into real-world probability density

$$g_t(r_T^F) = \frac{\partial Q(F_t(r_T^F))}{\partial r_T^F} = \frac{\partial Q}{\partial F_t}\frac{\partial F_t}{\partial r_T^F} = \frac{\partial Q}{\partial F_t} f_t(r_T^F). \tag{11}$$

However, the function $Q(\cdot)$ is unknown and we apply a parametric and a non-parametric methods to estimate it, as described in the following two subsections.

### 5.1. Beta statistical transformation

Fackler and King (1990) proposed to approximate the calibration function with the standard Beta distribution defined in the interval $[0, 1]$. The CDF of the Beta distribution is defined as

$$Q^B(F_t(r_T)|p,q) = \frac{\int_0^{F_t(r_T)} v^{p-1}(1-v)^{q-1}dv}{B(p,q)}, \tag{12}$$

where $B(p, q)$ is the Beta function with parameters $p$ and $q$. The main advantage of the Beta distribution is that it can have different shapes and is flexible enough to apply different corrections to the risk-neutral density. It can be applied when the RND is well calibrated and as such has no bias, which occurs when the estimated parameters $p$ and $q$ are both equal to 1.

Given Eq. (11) and the proposed calibration function in (12), the resulting real-world PDF is a function of the PITs and is defined as

$$g_t(r_T^F) = \frac{F_t(r_T^F)^{p-1}\left(1 - F_t(r_T^F)\right)^{q-1}}{B(p,q)} f_t(r_T^F). \tag{13}$$

The transformation of the RND to RWD boils down to the estimation of the parameters $p$ and $q$ by maximizing the likelihood of the observed futures rates at option maturity

$$(p,q) \equiv \arg\max \log L(p,q;F_t(r_T^F)) = \sum_{i=1}^{N} \log\left(g_t\left(r_{T_i}^F|p,q\right)\right). \tag{14}$$

Fig. 3 presents the shape of the transformation factor in (13) for different parameter combinations. As can be seen, the transformation factor can correct for both over- and underestimation of prob-

[13] Some observations in the period 2005–2007 are lost mainly due to lack of enough available OTM options to estimate the risk-neutral distribution.

[14] See Cochrane (2001) for a further elaboration on this relation.

[15] See Fackler and King (1990) for more details.

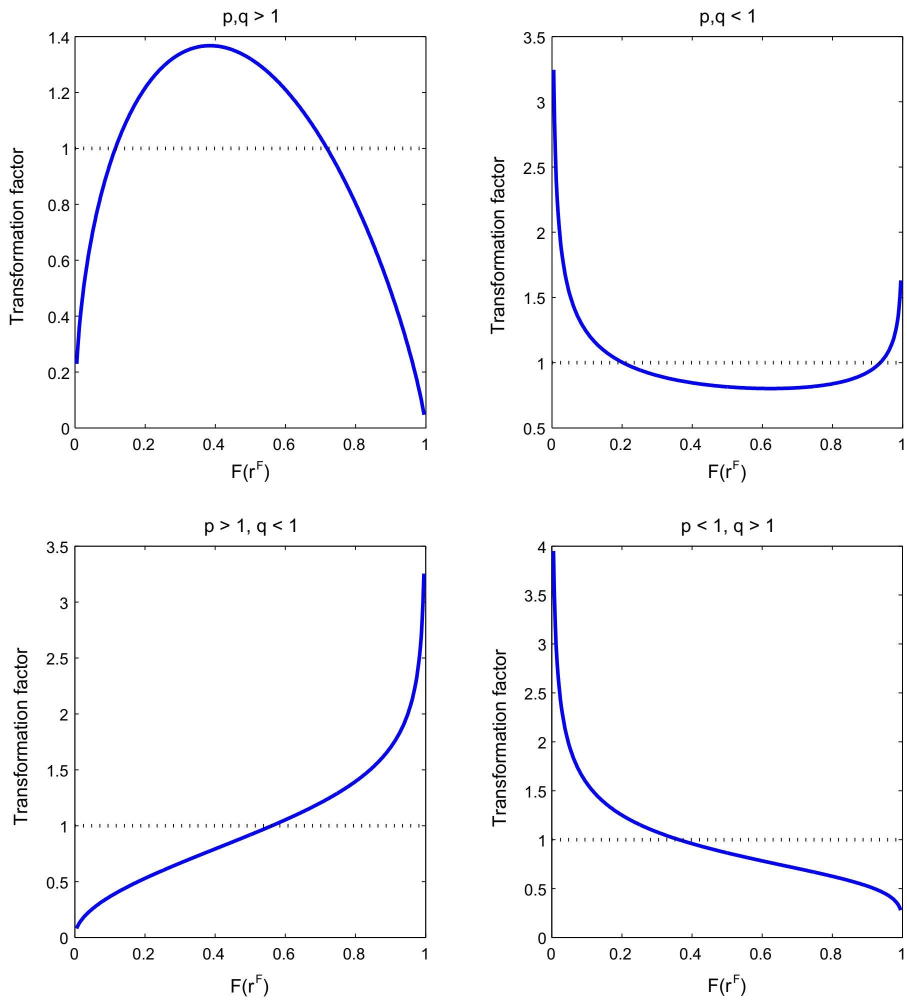


**Fig. 3.** Shapes of the parametric transformation for different parameters. The figure shows four different shapes of the transformation factor when using the Beta distribution. The shapes are generated over the general range [0,1] of a cumulative density function and for different parameter combinations.

ability in the tails (dispersion bias) and in the center of the density (location bias). The Beta distribution also nests the special case when the RND has no location basis, i.e. when the two parameters $p$ and $q$ are equal.

### 5.2. Non-parametric statistical transformation

The most straightforward method to derive the calibration function is to estimate the empirical CDF of the set of PIT, which motivates our attempt here. As stated before, given the observed interest rates at option maturity, we obtain a set of $N$ probability integral transforms $y_t^i$, $i = 1, \ldots, N$. Following the approach of Shackleton et al. (2010), we transform the observations $y_t^i$ using the inverse of the standard normal cumulative distribution, i.e. $z_t^i = \Phi^{-1}(y_t^i)$. Then, assuming that the observations are identically and independently distributed, we fit a normal kernel to the set $\{z_t^1, \ldots, z_t^2\}$. The kernel density estimator is defined to be

$$\hat{h}(z) = \frac{1}{NW}\sum_{i=1}^{N}\phi\left(\frac{z - z_i}{W}\right),$$

where $\phi(\cdot)$ denotes the standard normal density and $W$ denotes a smoothing parameter. Then the estimator for the kernel CDF is

$$\widehat{H}(z) = \frac{1}{N}\sum_{i=1}^{N}\Phi\left(\frac{z - z_i}{W}\right).$$

We select the smoothing parameter $W$, as described in Silverman (1986) and Shackleton et al. (2010), to be $W = 0.9\frac{\sigma_z}{N^{0.2}}$, where $\sigma_z$ is the standard deviation of the observation set $z_t^i$. Once the kernel estimation is performed the calibration function is defined by

$$Q(y) = \widehat{H}(\Phi^{-1}(y)), \tag{15}$$

and the real-world density results in

$$\begin{aligned} g_t(r_T^F) &= \frac{\partial G(r_T^F)}{\partial r_T^F} = \frac{\partial \widehat{H}(z)}{\partial r_T^F} = \frac{\partial z}{\partial r_T^F}\frac{\partial \widehat{H}(z)}{\partial z} = \frac{\partial y}{\partial r_T^F}\frac{\partial z}{\partial y}\frac{\partial \widehat{H}(z)}{\partial z} \\ &= \frac{\hat{h}(z)}{\phi(z)} f_t(r_T^F). \end{aligned} \tag{16}$$

The kernel density estimator is less stable for small samples, as it is also the case here. However, the method allows us to conduct a robustness check for the main findings of our study.

## 6. Empirical results

The first step in our empirical analysis is to use option prices four weeks before option maturity to estimate the risk-neutral probability density forecasting the futures rate at the option expiry date. We then evaluate its ability to produce reliable forecasts and check it for potential biases.

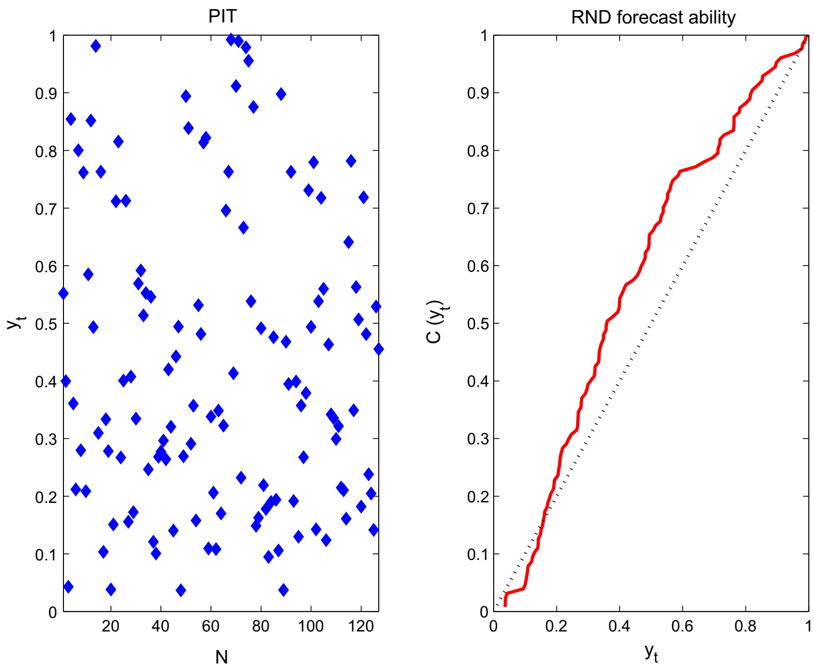


**Fig. 4.** Risk-neutral distribution forecast ability. The plots represent the forecasting power of the risk-neutral distributions. The left panel is a scatter plot of the empirical PITs calculated from the risk-neutral distribution (RND PIT). The right panel plots the empirical cumulative distribution function of the RND PIT (in red) against the cumulative distribution function of uniformly distributed random variables (dashed line). (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

As a second step we re-calibrate the risk-neutral densities to real-world densities aiming to correct the RND for the risk premium induced by the uncertain development of the futures rate. For the transformation, we estimate the calibration function in (11) following the parametric and non-parametric approach described in Section 5. The sample we use includes two crisis periods that may introduce potential regime changes and inconsistencies in the estimation. To control for this, we start the estimation using the set of PIT from the period 1999–2006, since these years show a period without major market turbulences. We then apply a rolling estimation by adding the next month PIT observation in the estimation process to obtain the calibration function for the next month. The forecasting ability of the transformed densities is then tested as described in Section 4. In this regard, note that while the calibration function is estimated in-sample for a given set of PITs, it is then used to transform the next out-of-sample forecast density.[16] As a result, in the context of traditional predictability tests, the analysis on the re-calibrated densities are not purely in-sample also in view of the rolling estimation. The purpose of the transformation is to obtain a well-calibrated probability assessment method based on the option-implied densities already extracted.

Once the densities are estimated, we proceed further by analyzing the state prices and the risk premium in the market.

### 6.1. Forecasting performance

Fig. 4 illustrates the forecast ability of the RND. The left panel is a scatter plot of the full series of probability integral transforms from Eq. (8) obtained with the risk-neutral densities, and the right panel plots the empirical CDF of the estimated PIT. As shown in Section 4, if the risk-neutral densities coincide with the true densities, the CDF of the PIT will lie on the 45%-degree line. However, it becomes obvious that the risk-neutral densities exhibit a location and dispersion bias. We quantify the dispersion bias with the inter-quartile range (IQR), defined as the difference between the 75%- and 25%-quartiles of the PIT CDF, while we evaluate the location bias by calculating the proportion of PIT values lying in the lower 50% interval of $[0,1]$.[17] In the ideal case, both tests will result in a value of 0.5, indicating equal probability mass in both halfs of the forecast density and no overestimation of volatility leading to high weights in the tails. Here, the IQR and the location tests are 0.52 and 0.65 respectively, which implies only a slight over-assessment of dispersion but significantly more weight to the left of the density.

The quantitative assessment of the forecast ability of the estimated densities is presented in Table 2. We report the test results based on the estimation using the full history of PIT, but the conclusion holds for all the sub-periods used in the rolling-window estimation of the calibration function. The left column shows the test statistics and corresponding $p$-values, when the RND is tested. The null hypothesis is that the forecast densities coincide with the true densities generating the futures rates. For the Berkowitz test this translates to the hypothesis that the series of transformed PITs, $z_t$, is independently and normally distributed. Here, the joint LR3 test indicates that the null hypothesis can be rejected at 5%-significance level. The Kolmogorov–Smirnov test also confirms this result. Therefore, we cannot find evidence that the risk-neutral densities provide reliable forecasts of the EURIBOR futures rate.

The next step is to use the parametric and non-parametric statistical calibration as presented in Section 5 to risk adjust the risk-neutral to real-world densities. Fig. 5 presents the development over time of the estimated parameters for the parametric statistical calibration. Two main results can be derived from this plot. First, both parameters $p$ and $q$ have values above 1. This implies that the resulting transformation factor has the form presented in the

[16] Only the initial calibration function from the sample covering the period 1999–2006 is used to transform not only the next density forecast but also all backward densities, so that we obtain a full set of PITs under the two measures.

[17] See Fackler and King (1990) for a more detailed discussion on the quantification of the location and dispersion bias.

**Table 2**
Forecasting power test statistics. The table presents the test results for the forecasting power of the risk-neutral (RND), the parametric real-world (Prm RWD) and non-parametric real-world (NonPrm RWD) distributions. The first test is the Berkowitz joint test (LR3) for normality and independence of the transformed PIT $z_t$, while the second one is Berkowitz's single test (LR1) only for independence of the observations. The null hypothesis for the Kolmogorov–Smirnov test is that the PIT transform $y_t$ has a uniform distribution. *P*-values are reported in parentheses. The last row shows the negative log-likelihoods for the realized ex-post futures rates from the three distributions.

| | RND | Prm RWD | NonPrm RWD |
|---|---|---|---|
| Berkowitz LR3 test | 12.691<br>(0.005) | 1.832<br>(0.608) | 2.241<br>(0.524) |
| Berkowitz LR1 test | 1.722<br>(0.189) | 1.764<br>(0.184) | 0.887<br>(0.346) |
| Kolmogorov–Smirnov test | 0.236<br>(0.001) | 0.087<br>(0.708) | 0.063<br>(0.957) |
| Log Likelihood | −90.571 | −95.451 | −102.048 |

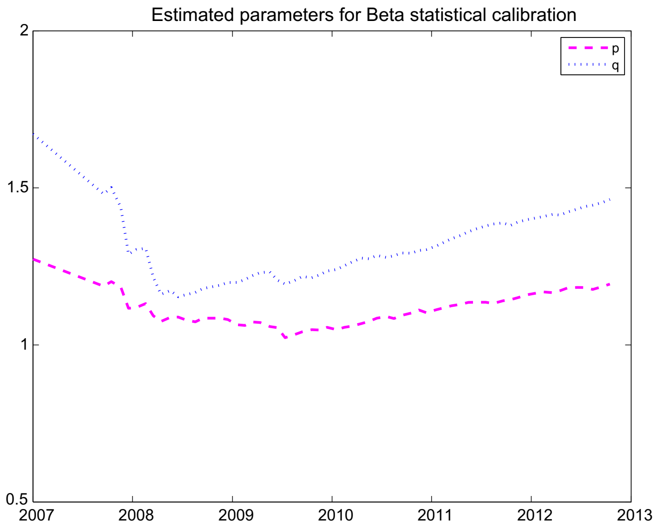


**Fig. 5.** Stability of estimated parameters for parametric transformation. The figure shows the development of the estimated parameters *p* and *q* for the parametric calibration. Initial parameters are based on the data from January 1999 until December 2006. The following parameters are obtained based on a rolling window estimation by adding a next-month PIT observation to the sample. As a result, the last parameters are based on the full history of PIT in the sample.

top-left panel of Fig. 3. As discussed earlier, the risk-neutral probabilities assign higher weight to states that the investor perceives as bad states of the world and the purpose of the transformation factor is to correct for this overestimation. From the shape of the transformation factor, one can read that states with very low and very high interest rates, i.e. $F(r_T)$ close to 0 and 1, are perceived as bad states, and the transformation reduces the weight by assigning a factor lower than one. Second, with the start of the US mortgage market turmoil in November 2007, we observe a slight drop of the parameter values. However, this did not result in a drastic change in the state price densities, as the general form of the transformation factor remained the same, but the market-assigned probability in the tails changed. We further elaborate on these results in Section 6.2.

After obtaining the real-world densities, we test their forecasting power. The forecasting performance should not be read in the context of in-sample versus out-of-sample predictability, as the approach we follow aims to use the available information set to re-calibrate only the next out-of-sample density so as to produce a reliable probability statement. As said before, this process is

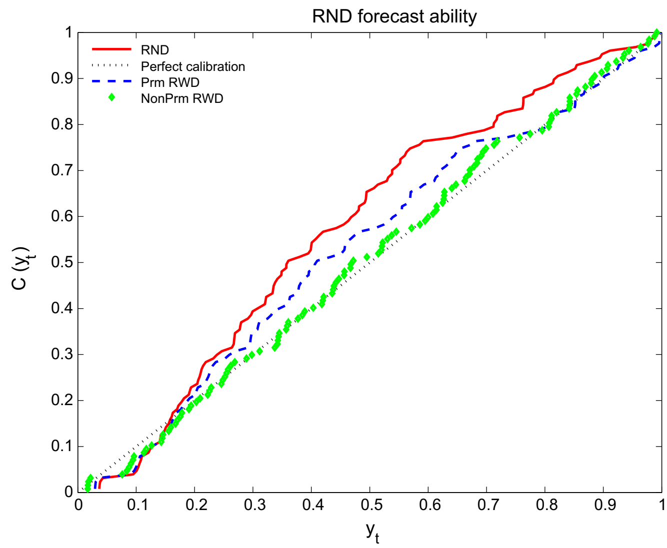


**Fig. 6.** Real-world densities forecast ability. The figure shows the cumulative distribution functions of the PIT transforms obtained from the risk-neutral distribution (red line), parametric (in blue) and non-parametric (in green) real-world distributions. The dashed line presents the CDF of a uniformly distributed variable. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

performed based on information of RND PITs. Results are reported in the second and third column of Table 2. The last row of the table also reports the negative log likelihood of observing the ex-post futures rates given the three forecast densities we evaluate. We consider the joint Berkowitz test for our main hypothesis that the real-world forecast densities coincide with the true densities generating the futures rate. The likelihood ratio statistic *LR*3 is $\chi^2(3)$-distributed; thus, the critical level for rejecting the null hypothesis at 5%-significance level is a value of 7.81. We find that the we cannot reject the null hypothesis at a significance level of above 50% for both transformations. Therefore, we conclude that only after a risk adjustment, through statistical re-calibration the option-implied densities can correctly capture the distribution of the Euribor futures rate.

Fig. 6 plots the CDF of the PIT transforms obtained from the real-world densities based on the full data set. One can see that both series of PIT from RW-densities manage to get closer to the 45-degree line than the RND, which also supports the superior forecasting ability of the RW-distributions. Both RWD also reduce the dispersion and location biases of the RND, with an IQR of 0.49 and 0.52 and location tests of 0.57 and 0.51, for parametric and non-parametric real-world densities respectively.

As a result, we can produce good forecasts for the price of the underlying instrument using option-implied information, but only after risk adjustment.[18] Statistical calibration is one possible adjustment method that allows for significant improvement of the density forecasting performance, since it does not require major economic assumptions such as a choice of the investor's utility function. In addition, the fact that the transformation using the Beta distribution as a calibration function provides as good forecasts as the pure non-parametric empirical calibration confirms its flexible use as argued by Fackler and King (1990).

[18] See also Bliss and Panigirtzoglou (2004), Liu et al. (2007), and Shackleton et al. (2010).

### 6.2. State price densities

Once the transformation from risk-neutral to real-world densities is applied, we can use both distributions to construct the state price densities. The difference between the two distributions arises from the fact that risk-averse investors value wealth in different states of the world differently. An extra euro in states where the investor is already wealthy is valued less than in states when wealth is low. The risk-neutral probabilities distort the true probability measure by assigning greater weight to bad states.

To understand the relationship between the two distributions, consider the price of a call option written on a Euribor future contract. Again, since the economic variable of interest in our paper is the futures rate implied by the price of the futures contract, we express the fundamentals accordingly:

$$C_t(r_t^F, r^K, T) = \int_{r^K}^{\infty} m_T(s)\left(r_T^F(s) - r^K\right)g\left(r_T^F(s)\right)ds$$
$$= \int_{r^K}^{\infty} \left(r_T^F(s) - r^K\right)f\left(r_T^F(s)\right)ds. \quad (17)$$

As a result, in each state of the world $s$, the risk-neutral density $f(s)$ is a product of the real-world distribution $g(s)$ and a stochastic discount function implying investors' risk-aversion $m(s)$.[19] Therefore, by constructing the ratio of the two distributions as in (18) below, known also as state price density (SPD) or pricing kernel, we can read the price at $t$ that investors assign to one euro in different states of the world at $T$.

$$m_T(s) = \frac{f(r_T^F(s))}{g(r_T^F(s))}. \quad (18)$$

We can also see from Eqs. (11) and (18) that the SPD is in fact the inverse of the transformation factor, $\frac{\partial Q}{\partial F}$, used in the calibration. The estimates are here obtained on a monthly basis due to the 4-weeks horizon used in the construction of the forecast densities. Several observations can be drawn by analyzing the shape and development of the state price densities over time.

Fig. 7 presents the transformation from risk-neutral density to real-world density on two random days in the sample, i.e. 22-December-2008 and 17-September-2012. First, both transformation methods correct the risk-neutral density by slightly shifting its center to the left, while reducing the probability mass in the tails of the distributions. This could also be read from the corresponding SPDs, plotted in the right panel of Fig. 7,[20] since the SPDs are higher in the tails, implying a higher probability mass in the tails of the RND. It should be noted that also the interquartile range of the RND has identified a strong dispersion bias, which is significantly corrected after the transformation to RWD.

Second, the bottom subplot indicates that the forecast density may also assign a positive probability for zero and negative futures rates. This is observed in 2012, when not only option prices were recorded for strikes higher and equal to 100, implying zero and negative futures rate respectively, but also these options were traded with high volumes and outstanding open interest. This gives an important insight into market perceptions, as interest rates have been traditionally thought to be limited at zero. Nevertheless, the sharp drop of interest rates with the US sub-prime crises, followed by the Euro sovereign debt crisis, have left market participants with a significant exposure to low interest rates. This led to heightened market expectations for further potential interest rate cuts, even in the negative territory, and triggered the need for downside Euribor protection, which resulted in option trading also for negative interest rates.

Third, the figures show that the state price densities are always positive and the parametric SPD exhibits a pronounced U-shaped curve. Therefore, investors assign higher state prices to payoffs in states with very low and high interest rates compared to the expected rate. This result is also confirmed by the non-parametric transformation. Nevertheless, state prices for high interest rates are higher than those for low rates, implying that market participants are actually more risk-averse to high interest rates. This conclusion holds even stronger for the period after 2008, for which the right tail of the SPD is much more pronounced, as shown on the bottom-right plot.

For comparison, empirical literature on pricing kernel estimation has been mainly focusing on equity and index markets. Several authors have documented a pricing kernel puzzle in that, opposite to major economic theory assumptions, the estimated pricing kernels take negative values for a wide range of states and are not monotonically downward-sloping across wealth.[21] Jackwerth (2000) explains these observations with a dramatic change in investors risk aversion after the 1987 crash. Rosenberg and Engle (2002) also confirmed, with a study on the S&P 500 index, that the level of risk aversion varies over time and showed that the pricing kernel is steeply upward-sloping for large negative returns and downward-sloping for large positive returns.

In this respect, our result is consistent with the economic theory in that the pricing kernel is positive across all states. In a related study, Li and Zhao (2009) applied a non-parametric estimation of the state prices in the LIBOR caps market and also found a pronounced U-shape of the SPDs, which they explained by means of term structure factors. However, they did not investigate the development of the pricing kernel through time. We go one step further and report the development of the parametric state prices from 2006 onwards in Table 3. The values correspond to the estimated SPD in the last month of the year. The first and the last column present the SPD for the minimum and maximum futures rate, while the columns in the middle refer to the $5^{th}\%$, $50^{th}\%$ and $95^{th}\%$-quantiles of the cumulative RND.

The table clearly indicates that the state prices for high interest rates are the largest. This is very pronounced in 2006, with the state price for the maximum rate being 50 times higher than for the minimum rate, and two times higher for the 95%-quantile rate versus the 5% one. This shows how in tranquil times, when rates are high, a further increase of the rate is perceived as a very bad state of inflation, which induces state prices higher for high rates than for low rates. It also becomes obvious that the futures rates outside the 90% inter-quartile range are the extreme states of the world that carry huge state prices. Ross (2011) justifies this with potential inconsistency of the two distributions in the tails.

We observe two main changes with the US mortgage market concerns in 2007, the resulting sub-prime crisis and the following economic turbulence in Europe. First, the state prices decreased significantly from 2007 onwards for extreme interest rates but not in the center of the distribution. This highlights the perceived market expectations for extreme central banks interventions and the changes in investors risk aversion to extreme interest rate adjustments. Therefore, in times of financial turmoils investors assign lower prices to extreme realizations than in tranquil times. Hence, in a crisis period actual probabilities in the market for tail realizations are closer to the risk-neutral counterparts. Second, the U-shaped SPD became more symmetric. The state prices for

[19] As its name reveals, the stochastic discount factor, besides incorporating the change of measure, also involves discounting. However, as a result of the futures-style margining of the Euribor futures option, no discounting applies here.

[20] Given that the state prices are too large for extreme tail realizations, we plot only the SPDs corresponding to the RN-probability density mass between the 5th and 95th percentiles, in order to produce a readable graphical output.

[21] See for instance the study of Jackwerth (2000) on the S&P 500 index market from 1986 until 1995.

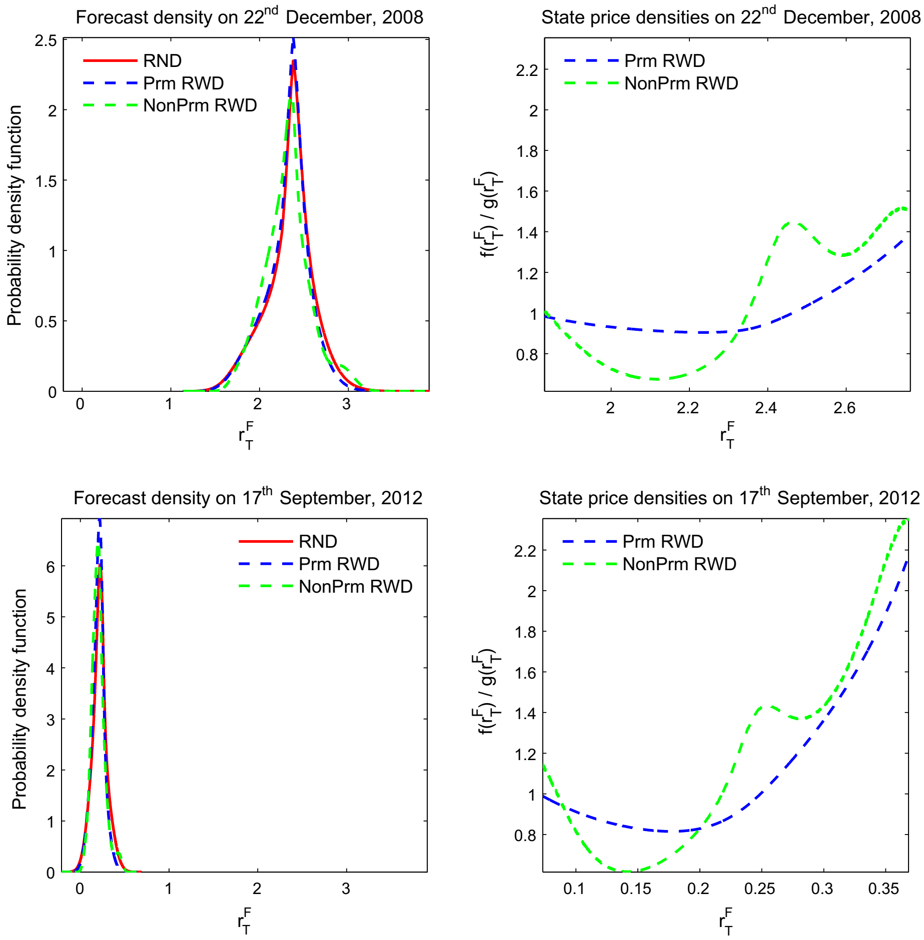


**Fig. 7.** Transformation from risk-neutral to real-world densities. The figure shows the transformation from risk-neutral to real-world densities on two particular days in the sample. The densities forecast the futures rate in 4 weeks. The left panel plots the probability density functions, i.e. the estimated risk-neutral density from option prices on the respective day (RND) and the real-world densities (RWD) obtained through the parametric (Prm) and non-parametric (NonPrm) transformation of the RND. The right panel plots the state-price densities, i.e. the ratio of the RND to RWD densities, corresponding to the RN-probability density mass between the $5^{th}$ and $95^{th}$ percentiles. Futures rates are reported in percentage.

**Table 3**
Development of parametric state price densities over time. The table presents the state price densities (SPD) for the respective futures rates derived from the parametric transformation at the last month of the reported year. Futures rates in percentages are shown in parenthesis. The first column refers to the SPD at the minimum futures rate in the observation month, while the last column refers to the maximum futures rate. The second, third and fourth columns report the SPD and the futures rates that correspond to the $5^{th}\%$, $50^{th}\%$ and $95^{th}\%$- quantiles of the cumulative risk-neutral distribution.

| | Min ($r^F$) | Q (5) | Q (50) | Q (95) | Max ($r^F$) |
|---|---|---|---|---|---|
| 2006 | 6.268 | 1.015 | 0.816 | 2.213 | 326.440 |
| | (3.223) | (3.514) | (3.659) | (3.750) | (4.038) |
| 2007 | 1.997 | 0.981 | 0.903 | 1.342 | 9.732 |
| | (3.578) | (4.213) | (4.603) | (4.766) | (5.328) |
| 2008 | 1.608 | 0.983 | 0.928 | 1.219 | 4.765 |
| | (1.144) | (1.830) | (2.368) | (2.655) | (3.908) |
| 2009 | 1.273 | 0.910 | 0.931 | 1.315 | 6.630 |
| | (0.326) | (0.621) | (0.760) | (0.922) | (1.474) |
| 2010 | 1.756 | 0.943 | 0.906 | 1.390 | 11.203 |
| | (0.648) | (0.936) | (1.160) | (1.337) | (2.301) |
| 2011 | 2.421 | 0.967 | 0.880 | 1.515 | 22.279 |
| | (0.636) | (1.057) | (1.393) | (1.660) | (3.623) |
| 2012 | 3.057 | 0.989 | 0.864 | 1.609 | 35.974 |
| | (−0.212) | (0.073) | (0.218) | (0.327) | (0.683) |

low rates were not so small compared to the state prices for high rates. Low realizations of the rates are a monetary policy attempt to stimulate economic growth in times of recession. A value of €1 in times of recession is higher than in normal times, which together with the higher likelihood for recession in this period caused more significant state prices for low rates.

Last but not least, to check the significance of our state price density estimates we also construct the confidence bounds of the parametric SPD. The top panel of Fig. 8 presents the state prices for the last month in the sample with their 90%-confidence bounds. To construct the bounds we use the asymptotic variance–covariance matrix of the estimated Beta distribution parameters from the maximum-likelihood estimation. The resulting confidence levels indicate the significance of the higher state prices for extreme futures rates. The bottom panel of the figure plots the futures rate range and the corresponding state prices extracted from the full history of Euribor futures options.[22] The range of option-implied interest rate expectations has shifted downward, as the Euribor rate decreased with the outburst of the US sub-prime crisis.[23] The plot confirms that U-shape state price densities have been observed dur-

[22] Note that the color map indicates the magnitude of the state prices. The blue nuances should be read as low state prices, while the red ones as very high.
[23] See Fig. 2 for the development of the 3-month Euribor futures rate over time.

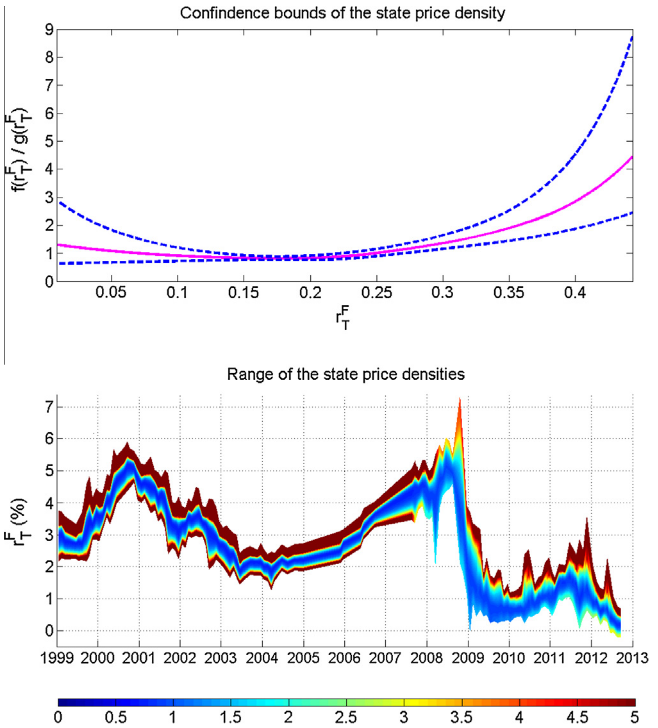


**Fig. 8.** Confidence bounds and range of the parametric state price densities. The top panel of the figure shows the 90% confidence bounds of the parametric state price density (SPD) obtained in the last month of the sample. The confidence bounds are constructed using the variance–covariance matrix of the estimated calibration parameters. The plot presents the range of futures rates (in percentages) that fall between the 1st and 99th percentiles of the risk-neutral distribution. The lower panel presents the development of the parametric SPD over time and futures rates. It plots the probability ratio, $\frac{f(r_T^F)}{g(r_T^F)}$, for the range of futures rates (on the left axis, in percentages) over which the forecast densities are defined. The color map indicates the values of the SPDs, with values between 0 and 5 ranging from blue to red, while all SPDs higher that 5 are depicted in dark red. For example, an SPD value in very light red indicates that the RND assigns a probability four times higher than the parametric RWD for this range of $r_T^F$. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

ing the full history both in tranquil and turbulent times, as the state prices have been continuously higher for the extreme interest rate realizations. It also indicates that starting from 2008, the right tail of the SPDs is significantly more pronounced. In addition, one can also read from the plot that, on several occasions, market participants were also pricing interest rate levels close to zero, but only in 2012 negative interest rates were traded.

### 6.3. Interest rate risk premium

Risk premium in the futures market is generated through the uncertainty about the spot interest rate at the futures maturity date. For illustrative purposes on the following discussion of the empirically observed interest risk premium in the market, let us assume any general dynamics for the interest rate $r$

$$dr_t = \mu(r_t, t)dt + \sigma(r_t, t)dW_t. \tag{19}$$

Consider now the futures rate as a function of the interest rate, $r_t^F = J(r, t)$.[24] By applying Ito's lemma, we can analyze directly the dynamics of the futures rate

[24] Note that the presentation here assumes a general model set-up and abstracts from the specific dynamics assumed in Black (1976) and Cox et al. (1981). Also, since the futures price is defined as $P_t = 100 - r_t^F$, the problem here is equivalent to considering the futures price as a function of the interest rate but with an inverse relationship, i.e. $P_t = 100 - J(r, t)$.

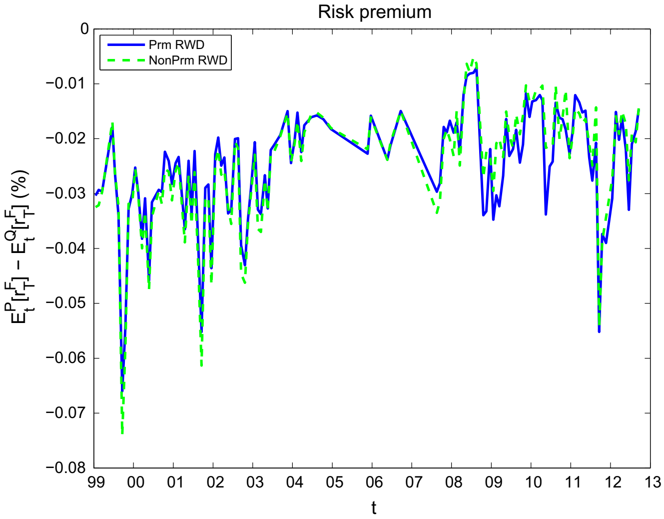


**Fig. 9.** Interest rate risk premium. The figure shows the development of the interest rate risk premium over time obtained with both calibration methods. Risk premium is defined as the difference between the mean of the real-world density and the risk-neutral density. The $y$-axis is in percentages.

$$dr_t^F = (J_r\mu(r_t, t) + J_{rr}\sigma^2(r_t, t) + J_t)dt + J_r\sigma(r_t, t)dW_t. \tag{20}$$

Therefore, under the true probability measure, the expected value of the futures rate increments is

$$E^{\mathbb{P}}[dr^F] = J_r\mu(r_t, t) + J_{rr}\sigma^2(r_t, t) + J_t. \tag{21}$$

As shown by Cox et al. (1981, 1985), in the presence of interest rate uncertainty the futures rate will solve

$$E^{\mathbb{Q}}[dr^F] = J_r(\mu(r_t, t) - \lambda_t) + J_{rr}\sigma^2(r_t, t) + J_t = 0, \tag{22}$$

under the terminal condition that the futures rate is equal to the observed spot Euribor rate at the futures maturity date. Here the expectation is taken under the risk-neutral probability measure and $\lambda$ is the factor risk premium, also called market price of risk, associated with interest rate uncertainty. Hence, when we compare the mean of the futures rate distributions, we measure the interest rate risk premium as of the forecast date $t$ for the uncertain innovations in the interest rate path until the option maturity date $T$.

$$E_t^{\mathbb{P}}[r_T^F] - E_t^{\mathbb{Q}}[r_T^F] = r_t^F + \int_t^T \lambda_\tau J_r d\tau - r_t^F = \int_t^T J_r\lambda_\tau d\tau. \tag{23}$$

Market price of interest rate risk is the excess return investors require to bear an additional unit of risk. Therefore, it is the market-embedded driver of the difference between the risk-neutral and real-world distribution of the futures rate. Given that the futures rate is the expected value of the spot rate, its partial derivative is positive, thus we can infer the sign of the market price of risk, $\lambda$, from the sign of the risk premium in (23).[25]

As shown by Cox et al. (1985), in a general equilibrium setting with interest rate uncertainty, the market price of risk, $\lambda$, is the covariance of changes in the spot rate with percentage changes in an investor's wealth. As a result, a negative covariance implies that high interest rates are associated with low wealth and thus with a market-perceived bad state of the economy. Fig. 9 plots the difference of the means of the real-world and risk-neutral distributions from both transformation methods. It shows that the market price of risk is indeed negative in Euribor futures market.

[25] The inverse relationship between the futures rate and the price of the futures contract does not change the reasoning here for the market price of interest rate risk: $E_t^{\mathbb{P}}[P_T] - E_t^{\mathbb{Q}}[P_T] = \int_t^T -J_r\lambda_\tau d\tau$ and $J_r < 0$.

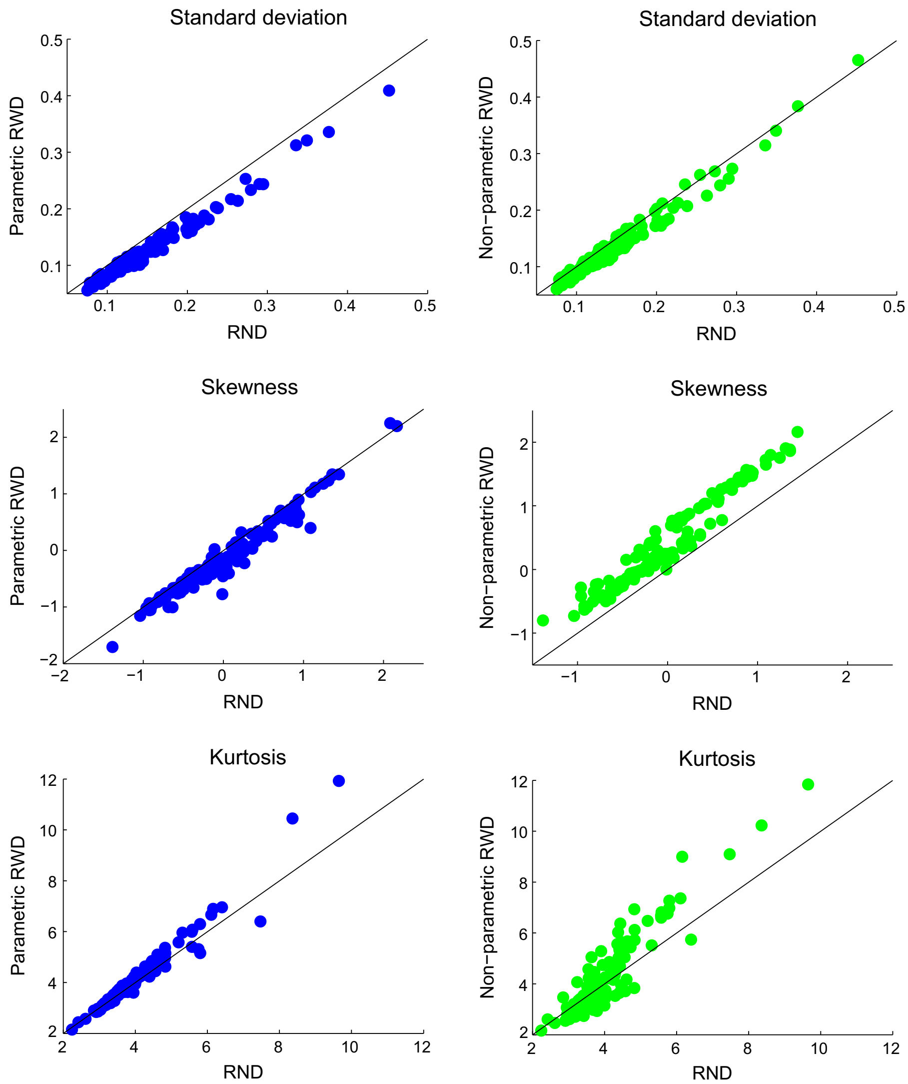


**Fig. 10.** Comparison of moments. The figure shows the comparison of moments derived from the set of risk-neutral PDFs and the transformed real-world densities. The left panels compare the moments using the parametric transformation, while the right panel is based on the non-parametric transformation. The moments are derived for the set of densities obtained four weeks before option maturity.

This finding holds under both transformation methods. As discussed earlier, regardless of the U-shaped SPDs, high futures rates are associated with state prices higher than those for low rates, which confirms that they are considered as a worse state of the economy. Therefore, a negative market price of risk reads that high rates were linked to bad states with low wealth.

Our result is consistent with the negative market price of interest risk derived in the bond market. Stanton (1997) estimates a negative market price of interest rate risk, using daily Treasury Bills yields, decreasing in the level of the rate. In this paper, we do not decompose the market price of risk as a function of the interest rate and thus we cannot infer this. However, an interesting observation here is that the mean difference reaches its lowest absolute value when the futures rate reach its maximum peak around July 2008. As shown in Section 6.2, this corresponds to a period when fears of recession were driving the U-shaped SPD curve more symmetric. Following the interest rate declines after the end of 2008, the mean difference increased also on average.

It is also important to highlight that while we document here the premium on the futures rate directly, it is the futures contract which is the actually traded asset in the market, and not directly the futures rate. As a result, given the inverse relation between the futures rate and futures price, the futures contract carries a positive premium, while the market price of interest rate risk remains negative.

Last but not least, after finding a significant difference in the mean of the real-world and risk-neutral distributions, we proceed with the analysis of other moments of the forecast distributions. Fig. 10 presents standard deviation, skewness and kurtosis of the RND against their RWD counterparts with both transformation methods. It shows that the rest of the moments also differ substantially after the risk adjustment, which is not surprising in view of the significant differences in the forecast ability of the risk-neutral and real-world probability density functions. Note that the parametric transformation has led to a small but consistent reduction of the standard deviation of the RWD forecast densities as compared to the RND, as also empirically observed by Liu et al. (2007). The sign of the skewness was preserved with the transformation, while its magnitude was also slightly reduced. Nevertheless, the skewness changed its sign during the market developments observed in the Euribor history. The skewness in the market was on average negative before 2008, but then turned positive indicating, higher expectations for low interest rates due to the heavier probability mass in this region. This change in the distributional moments between the RND and the parametric RWD can be also seen in the bottom subplot of Fig. 8, as the state price densities represent the probability ratio of under both measures. It shows clearly that the RND has assigned higher probability mass to both tails until 2008, when the right tail of both distributions became longer, but remained persistently heavier under the RND. The kurtosis also increased after the middle of 2007, indicating growing concerns for the economic outlook in the euro area. For example, market expectations around the first bailout of Greece in May 2010 and of Ireland in November 2010 produced the two significant outliers, observed in the middle and bottom left subplots, with very high positive skewness and kurtosis.

## 7. Conclusion

In this paper we study the information content of Euribor futures options regarding interest rate forecasts, states prices and risk premium in the market. We estimate the risk-neutral distributions from option prices using the spline method suggested by Bliss and Panigirtzoglou (2002). To adjust for potential interest rate risk premium we further transform the risk-neutral densities into their real-world counterparts. We use parametric and non-parametric statistical calibration to adjust the densities as proposed by Fackler and King (1990). The estimated distributions are then used to generate forecasts for the futures rate four weeks before option maturity. We find that the risk-neutral densities cannot produce reliable forecasts, while the adjusted real-world densities adequately capture the distribution of the ex-post futures rates.

Using the two sets of option-implied distributions we proceed to analyze market fundamentals. We find a positive state price density with a U-shape form with respect to the Euribor futures rate. This implies that investors fear both extremes of high and low rates compared to the expected spot rate, and assign higher state prices to them. However, we show that state prices have a more pronounced right tail, implying that investors are more risk-averse to increasing interest rates. In addition, with the series of financial turmoils after 2007, we provide evidence for a change in the investors' risk perceptions and a more symmetric SPD curve for a short period of time. Last but not least, we confirm the negative market price of interest rate risk, as also found in the bonds market, and show that it generates an economically significant risk premium, which is positive for the futures contract.

## Disclaimer

This paper should not be reported as representing the views of the European Central Bank (ECB). The views expressed are those of the authors and do not necessarily reflect those of the ECB.

## References


Ait-Sahalia, Y., Duarte, J., 2003. Nonparametric option pricing under shape restrictions. Journal of Econometrics 116, 9–47.

Alonso, F., Blanco, R., Rubio, G., 2006. Option-Implied Preference Adjustments, Density Forecasts, and the Equity Risk Premium. Bank of Spain Working Paper Series 0630.

Anagnou-Basioudis, I., Bedendo, M., Hodges, S., Tompkins, R., 2005. Forecasting accuracy of implied and GARCH-based probability density functions. Review of Futures Markets 11, 41–66.

Berkowitz, J., 2001. Testing density forecasts with applications to risk management. Journal of Business and Economic Statistics 19, 465–474.

Black, F., 1976. The pricing of commodity contracts. Journal of Financial Economics 3, 167–179.

Bliss, R., Panigirtzoglou, N., 2002. Testing the stability of implied probability density functions. Journal of Banking and Finance 26 (2–3), 381–422.

Bliss, R., Panigirtzoglou, N., 2004. Option-implied risk aversion estimates. Journal of Finance 59, 407–446.

Breeden, D., Litzenberger, R., 1978. Prices of state-contingent claims implicit in option prices. Journal of Business 51 (4), 621–651.

Bunn, D., 1984. Applied Decision Analysis. McGraw-Hill.

Campa, J., Chang, K., Reider, R., 1997. Implied Exchange Rate Distributions: Evidence from OTC Option Markets. NBER Working Papers 6179.

Cesari, R., Sevini, L., 2004. Using Options to Forecast Libor. University of Bologna, Working Paper. <http://ssrn.com/abstract=587343>.

Chen, R., Scott, L., 1993. Pricing interest rate futures options with futures-style margining. Journal of Futures Markets 13, 15–22.

Cochrane, J., 2001. Asset Pricing. Princeton University Press.

Cochrane, J., Piazzesi, M., 2005. Bond risk premia. The American Economic Review 95, 138–160.

Cooper, N., 1999. Testing Techniques for Estimating Implied RNDs from the Prices of European and American Options. Mimeo, Bank of England.

Cox, J., Ingersoll, J., Ross, S., 1981. The relation between forward price and futures prices. Journal of Financial Economics 9, 321–346.

Cox, J., Ingersoll, J., Ross, S., 1985. A theory of the term structure of interest rates. Econometrica 53, 385–408.

De Vincent-Humphreys, R., Noss, J., 2012. Estimating Probability Distributions of Future Asset Prices: Empirical Transformations from Option-Implied Risk-Neutral to Real-World Density Functions. Bank of England Working Paper Series 455.

De Vincent-Humphreys, R., Puigvert Gutiérrez, J., 2012. A quantitative mirror on the Euribor market using implied probability density functions. Eurasian Economic Review 2 (1), 1–31.

Diebold, F., Todd, A., Tay, A., 1998. Evaluating density forecasts with applications to financial risk management. International Economic Review 39, 863–883.

Diebold, F., Hahn, J., Tay, A., 1999. Multivariate density forecast evaluation and calibration in financial risk management: high-frequency returns on foreign exchange. Review of Economics and Statistics 81, 661–673.

Duffie, D., 1989. Futures Markets. Prentice-Hall, Englewood Cliffs, NJ.

Fackler, P., King, R., 1990. Calibration of option-based probability assessments in agricultural commodity markets. American Journal of Agricultural Economics 72, 73–83.

Fama, E., French, K., 1993. Common risk factors in the returns on stocks and bonds. Journal of Financial Economics 33, 3–56.

Grinblatt, M., Jegadeesh, N., 1996. Relative pricing of eurodollar futures and forward contracts. Journal of Finance 51, 1499–1522.

Hess, A., Kamara, A., 2003. Conditional Time-Varying Interest Rate Risk Premium: Evidence from the Treasury Bill Futures Market. Working paper.

Hirscleifer, D., 1988. Residual risk, trading costs, and commodity futures risk premia. Review of Financial Studies 1, 173–193.

Hog, E., Tsiaras, L., 2011. Density forecasts of crude-oil using option-implied and ARCH-type models. Journal of Futures Markets 31, 727–754.

Jackwerth, J., 2000. Recovering risk aversion from option prices and realized returns. Review of Financial Studies 13, 433–451.

Jackwerth, J., Rubinstein, M., 1996. Recovering probability distributions from option prices. Journal of Finance 51, 1611–1631.

Jiang, G., Tian, Y., 2005. The model-free implied volatility and its information content. Review of Financial Studies 18, 1305–1342.

Li, H., Zhao, F., 2009. Nonparametric estimation of state-price densities implicit in interest rate cap prices. Review of Financial Studies 22, 4335–4376.

Lieu, D., 1990. Option pricing with futures-style margining. The Journal of Futures Markets 10, 327–338.

Liu, X., Shackelton, M., Taylor, S., Xu, X., 2007. Closed-form transformations from risk-neutral to real-world distributions. Journal of Banking and Finance 31, 501–520.

Malz, A., 1997. Estimating the probability distribution of the future exchange rate from option prices. Journal of Derivatives 5, 18–36.

Martens, M., Zein, J., 2004. Predicting financial volatility: high-frequency time-series forecasts vis-a-vis implied volatility. Journal of Futures Markets 24, 1005–1028.

Ornelas, J., Barbachan, J., de Farias, A., 2012. Estimating Relative Risk Aversion, Risk-Neutral and Real-World Densities Using Brazilian Real Currency Options. Banco Central do Brasil Working Paper Series.

Rosenberg, J., Engle, R., 2002. Empirical pricing kernels. Journal of Financial Economics 64, 341–372.

Rosenblatt, M., 1952. Remarks on a multivariate transformation. Annals of Mathematical Statistics 23, 470–472.

Ross, S., 2011. The Recovery Theorem. Working Paper.

Soderlind, P., Svensson, L., 1997. New techniques to extract market expectations from financial instruments. Journal of Monetary Economics 40, 383–429.

Shackleton, M., Taylor, S., Peng, Y., 2010. A multi-horizon comparison of density forecasts for the S&P 500 using index returns and option prices. Journal of Banking and Finance 34, 2678–2693.

Shimko, D., 1993. Bounds of probability. RISK 6, 33–37.

Silverman, B., 1986. Density Estimation for Statistics and Data Analysis. Chapman Hall.

Stanton, R., 1997. A nonparametric model of term structure dynamics and the market price of interest rate risk. Journal of Finance LII, 1973–2002.

Takkabutr, N., 2012. Density Estimation and Option-Implied Risk Aversion of the nikkei 225: Evidence before and after the Subprime Crisis. GCOE Discussion Paper Series 242.

# Bibliography

[1] Francisco Alonso, Roberto Blanco, and Gonzalo Rubio. Option-implied preferences adjustments, density forecasts, and the equity risk premium. *Spanish Economic Review*, 11(2):141–164, 2009.

[2] Jeffery D Amato, Stephen Morris, and Hyun Song Shin. Communication and monetary policy. *Oxford Review of Economic Policy*, 18(4):495–503, 2002.

[3] Yacine Aït-Sahalia and Jefferson Duarte. Nonparametric option pricing under shape restrictions. *Journal of Econometrics*, 116(1):9–47, 2003.

[4] Iliana Anagnou-Basioudis, Mascia Bedendo, Stewart D Hodges, and Robert Tompkins. Forecasting accuracy of implied and GARCH-based probability density functions. *Review of Futures Markets*, 11:41–66, 2005.

[5] Allan Bødskov Andersen and Tom Wagener. Extracting risk neutral probability densities by fitting implied volatility smiles: some methodological points and an application to the 3M Euribor futures option prices. *ECB Working Paper Series*, 198, 2002.

[6] Magnus Andersson. Using intraday data to gauge financial market responses to Fed and ECB monetary policy decisions. *ECB Working Paper Series*, 726, 2007.

[7] Kenneth J Arrow and Gerard Debreu. Existence of an equilibrium for a competitive economy. *Econometrica: Journal of the Econometric Society*, pages 265–290, 1954.

[8] Piotr Banbula. Options and Market Expectations: Implied Probability Density Functions on the Polish Foreign Exchange Market. *Bank i Kredyt*, (5):21–35, 2008.

[9] David S Bates. Post-'87 crash fears in the S&P 500 futures option market. *Journal of Econometrics*, 94(1-2):181–238, 2000.

[10] Bhupinder Bahra. Probability distributions of future asset prices implied by option prices. *Bank of England Quarterly Bulletin*, 36(3):299–311, 1996.

[11] Bhupinder Bahra. Implied risk-neutral probability density functions from option prices: theory and application. *Bank of England Working Paper*, 66, 1997.

[12] Jeremy Berkowitz. Testing density forecasts, with applications to risk management. *Journal of Business & Economic Statistics*, 19(4):465–474, 2001.

[13] BIS. Estimating and interpreting probability density functions, Proceedings of the workshop on implied PDFs held at the BIS on 14 june 1999. *Bank for International Settlements, Basel*, 1999.

[14] Fischer Black. The pricing of commodity contracts. *Journal of financial economics*, 3(1):167–179, 1976.

[15] Fischer Black and Myron Scholes. The pricing of options and corporate liabilities. *The journal of political economy*, 81(3):637–654, 1973.

[16] Robert R Bliss and Nikolaos Panigirtzoglou. Testing the stability of implied probability density functions. *Journal of Banking & Finance*, 26(2):381–422, 2002.

[17] Robert R Bliss and Nikolaos Panigirtzoglou. Option-implied risk aversion estimates. *The journal of finance*, 59(1):407–446, 2004.

[18] Claus Brand, Daniel Buncic, and Jarkko Turunen. The impact of ECB monetary policy decisions and communication on the yield curve. *Journal of the European Economic Association*, 8(6):1266–1298, 2010.

[19] Douglas T Breeden and Robert H Litzenberger. Prices of state-contingent claims implicit in option prices. *Journal of business*, 51:621–651, 1978.

[20] Derek W. Bunn. *Applied decision analysis*, volume 8. McGraw-Hill New York, 1984.

[21] José M Campa, PH Chang, and Robert L Reider. Implied exchange rate distributions: evidence from OTC option markets. *Journal of International Money and Finance*, 17(1):117–160, 1998.

[22] Olli Castren. Do options-implied RND functions on G3 currencies move around the times of interventions on the JPY/USD exchange rate? *ECB Working Paper Series*, 410, 2004.

[23] Riccardo Cesari and Lorenzo Sevini. Using Options to Forecast Libor. *University of Bologna, Working Paper*, 2004.

[24] PH Kevin Chang and William R Melick. Background note. In *Proceedings of the BIS workshop on implied PDFs, Bank of International Settlements, Basel*, Bank of International Settlements, Basel, 1999. Bank of International Settlements.

[25] Ren-Raw Chen and Louis Scott. Pricing interest rate futures options with futures-style margining. *Journal of Futures Markets*, 13(1):15–22, 1993.

[26] Martin Cícha. Risk-neutral Option Pricing and Information Embedded in CEZ Warrant Prices: Implied Volatility Smile, Implied Density Function. *Research Centre for Competitiveness of Czech Economy Working Paper*, (5), 2009.

[27] Roger Clews, Nikolaos Panigirtzoglou, and James Proudman. Recent developments in extracting information from options markets. *Bank of England Quarterly Bulletin*, 40(1):50–59, 2000.

[28] John Howland Cochrane. *Asset pricing*, volume 41. Princeton University Press Princeton, 2005.

[29] John H Cochrane and Monika Piazzesi. Bond risk premia. *The American Economic Review*, 95:138–160, 2005.

[30] Neil Cooper. Testing techniques for estimating implied rnds from the prices of european-style options. In *Proceedings of the workshop on implied PDFs held at the BIS on 14 June 1999*, Bank of International Settlements, Basel, 1999. Bank of International Settlements.

[31] Sophie Coutant, Eric Jondeau, and Michael Rockinger. Reading interest rate and bond futures options' smiles: How pibor and notional operators appreciated the 1997 french snap election. (54), 1998.

[32] Kabir K Dutta and David F Babbel. Extracting Probabilistic Information from the Prices of Interest Rate Options: Tests of Distributional Assumptions. *The Journal of Business*, 78(3):841–870, 2005.

[33] Eric Jondeau, Ser-Huang Poon, and Michael Rockinger. *Financial Modeling Under Non-Gaussian Distributions*. Springer Finance, 2006.

[34] John C Cox, Jonathan E Ingersoll Jr, and Stephen A Ross. The relation between forward prices and futures prices. *Journal of Financial Economics*, 9(4):321–346, 1981.

[35] John C Cox, Jonathan E Ingersoll Jr, and Stephen A Ross. A theory of the term structure of interest rates. *Econometrica: Journal of the Econometric Society*, 53:385–407, 1985.

[36] John C Cox and Stephen A Ross. The valuation of options for alternative stochastic processes. *Journal of financial economics*, 3(1):145–166, 1976.

[37] Carl De Boor. *A practical guide to splines*, volume 27. Applied Mathematical Sciences. Springer, 1978.

[38] Francis X Diebold, Todd A Gunther, and Anthony S Tay. Evaluating density forecasts with applications to financial risk management. *International Economic Review*, 39:863–883, 1998.

[39] Francis X Diebold, Jinyong Hahn, and Anthony S Tay. Multivariate density forecast evaluation and calibration in financial risk management: high-frequency returns on foreign exchange. *Review of Economics and Statistics*, 81(4):661–673, 1999.

[40] Darrell Duffie. *Futures markets*. Prentice Hall Englewood Cliffs,NJ, 1989.

[41] ECB. Communicating Monetary Policy to Financial Markets. *ECB. Monthly Bulletin Article*, pages 61–71, April 2007.

[42] ECB. The information content of option prices during the financial crisis. *ECB. Monthly Bulletin Article*, pages 87–99, February 2011.

[43] Michael Ehrmann and Marcel Fratzscher. Explaining monetary policy in press conferences. *International Journal of Central Banking*, 5(2):41–83, 2009.

[44] Paul L Fackler and Robert P King. Calibration of option-based probability assessments in agricultural commodity markets. *American Journal of Agricultural Economics*, 72(1):73–83, 1990.

[45] Eugene F Fama and Kenneth R French. Common risk factors in the returns on stocks and bonds. *Journal of financial economics*, 33(1):3–56, 1993.

[46] Mark Grinblatt and Narasimhan Jegadeesh. Relative pricing of Eurodollar futures and forward contracts. *The Journal of Finance*, 51(4):1499–1522, 1996.

[47] Alan C Hess and Avraham Kamara. Conditional time-varying interest rate risk premium: evidence from the treasury bill futures market. *Journal of Money, Credit and Banking*, pages 679–698, 2005.

[48] David Hirshleifer. Residual risk, trading costs, and commodity futures risk premia. *Review of Financial Studies*, 1(2):173–193, 1988.

[49] Esben Høg and Leonidas Tsiaras. Density forecasts of crude-oil prices using option-implied and ARCH-type models. *Journal of Futures Markets*, 31(8): 727–754, 2011.

[50] Jens Carsten Jackwerth. Option-implied risk-neutral distributions and implied binomial trees: A literature review. *The Journal of Derivatives*, 7(2):66–82, 1999.

[51] Jens Carsten Jackwerth. Recovering risk aversion from option prices and realized returns. *Review of Financial Studies*, 13(2):433–451, 2000.

[52] Jens Carsten Jackwerth and Mark Rubinstein. Recovering probability distributions from option prices. *The Journal of Finance*, 51(5):1611–1631, 1996.

[53] George J Jiang and Yisong S Tian. The model-free implied volatility and its information content. *Review of Financial Studies*, 18(4):1305–1342, 2005.

[54] Eric Jondeau and Michael Rockinger. Reading the smile: the message conveyed by methods which infer risk neutral densities. *Journal of International Money and Finance*, 19(6):885–915, 2000.

[55] Yuriy Kitsul and Jonathan H Wright. The economics of options-implied inflation probability density functions. *Journal of Financial Economics*, 110(3): 696–711, 2013.

[56] Susanne Kruse. On the pricing of inflation-indexed caps. *European Actuarial Journal*, 1(2):379–393, 2011.

[57] John Hull. *Options, futures and other derivatives*. Pearson education. Prentice Hall, 4th edition, 2009.

[58] Wan-Ni Lai. Comparison of methods to estimate option implied risk-neutral densities. *Quantitative Finance*, 14(10):1839–1855, 2014.

[59] Haitao Li and Feng Zhao. Nonparametric estimation of state-price densities implicit in interest rate cap prices. *Review of Financial Studies*, 22:4335–4376, 2009.

[60] Derming Lieu. Option pricing with futures-style margining. *Journal of Futures Markets*, 10(4):327–338, 1990.

[61] Nyse Liffe. *NYSE Liffe Trading Procedures*. www.euronext.com, 2009.

[62] Xiaoquan Liu, Mark B Shackleton, Stephen J Taylor, and Xinzhong Xu. Closed-form transformations from risk-neutral to real-world distributions. *Journal of Banking & Finance*, 31(5):1501–1520, 2007.

[63] Damien Lynch. Using option prices to measure financial market views about balances of risk to future asset prices. *Bank of England Quarterly Bulletin*, 44(4):442–454, 2004.

[64] Allan M Malz. Estimating the probability distribution of the future exchange rate from option prices. *The Journal of Derivatives*, 5(2):18–36, 1997.

[65] Martin Mandler. Comparing risk-neutral probability density functions implied by option prices: market uncertainty and ECB council meetings. *Working Paper. Universität Giessen*, 2002.

[66] Martin Martens and Jason Zein. Predicting financial volatility: High-frequency time-series forecasts vis-à-vis implied volatility. *Journal of Futures Markets*, 24(11):1005–1028, 2004.

[67] Desmond John McManus. The information content of interest rate futures options. *Working Paper. Bank of Canada*, 1999.

[68] William R Melick and Charles P Thomas. Recovering an asset's implied PDF from option prices: an application to crude oil during the Gulf crisis. *Journal of Financial and Quantitative Analysis*, 32(01):91–115, 1997.

[69] William R Melick and Charles P Thomas. Confidence intervals and constant maturity series for probability measures extracted from option prices. In *Information in Financial Asset Prices: Proceedings of a Conference at the Bank of Canada*, pages 293–320. Citeseer, 1998.

[70] Holger Neuhaus. The information content of derivatives for monetary policy: Implied volatilities and probabilities. *Economic Research Group Discussion Paper. Deutsche Bundesbank*, 3(95), 1995.

[71] Jose Renato Haas Ornelas, José Fajardo Barbachan, and Aquiles Rocha de Farias. Estimating relative risk aversion, risk-neutral and real-world densities using brazilian real currency options. *Banco Central do Brasil Working Paper Series*, 2012.

[72] Joshua V Rosenberg and Robert F Engle. Empirical pricing kernels. *Journal of Financial Economics*, 64(3):341–372, 2002.

[73] Murray Rosenblatt. Remarks on a multivariate transformation. *The annals of mathematical statistics*, 23:470–472, 1952.

[74] Steve Ross. The recovery theorem. *The Journal of Finance*, 2013.

[75] Mark B Shackleton, Stephen J Taylor, and Peng Yu. A multi-horizon comparison of density forecasts for the S&P 500 using index returns and option prices. *Journal of Banking & Finance*, 34(11):2678–2693, 2010.

[76] David Shimko. Bounds of probability. *Risk*, 6(4):33–37, 1993.

[77] Tom Smith. Option-implied probability distributions for future inflation. *Bank of England Quarterly Bulletin*, page Q3, 2012.

[78] B. Silverman. *Density Estimation for Statistics and Data Analysis*. Chapman Hall, 1986.

[79] Paul Söderlind and Lars Svensson. New techniques to extract market expectations from financial instruments. *Journal of Monetary Economics*, 40(2): 383–429, 1997.

[80] Stig Arild Syrdal. A study of implied risk-neutral density functions in the Norwegian option market. *Working Paper*, 2002.

[81] Richard Stanton. A nonparametric model of term structure dynamics and the market price of interest rate risk. *The Journal of Finance*, 52(5):1973–2002, 1997.

[82] N Takkabutr. Density estimation and option-implied risk aversion of the Nikkei 225: Evidence before and after the subprime crisis. Technical report, GCOE Discussion paper series 242, 2012.

[83] Alan White. Pricing options with futures-style margining: A genetic adaptive neural network approach. (Financial sector of the American economy). *Garland Publishing Inc*, 2000, October 1973.

[84] Lan Zhang, Per A Mykland, and Yacine Aït-Sahalia. A tale of two time scales. *Journal of the American Statistical Association*, 100:1394–1411, 2005.

# Declaration

I herewith declare that I have produced this work without the prohibited assistance of third parties and without making use of aids other than those specified; notions taken over directly or indirectly from other sources have been identified as such. This work has not previously been presented in identical or similar form to any examination board.

The dissertation work was conducted from 2008 to 2015 under the supervision of Josep Fortiana at the University of Barcelona.

Frankfurt am Main, Saturday $18^{th}$ April, 2015

Pep Puigvert.

Josep Maria Puigvert Gutiérrez

*This page is intentionally left blank*