EUROPEAN CENTRAL BANK

EUROSYSTEM

WORKING PAPER SERIES

NO 1505 / JANUARY 2013

INTEREST RATE VOLATILITY

A CONSOL RATE-BASED MEASURE

Vincent Brousseau and Alain Durré

In 2013 all ECB publications feature a motif taken from the €5 banknote

Acknowledgements
The authors would like to thank colleagues in the Directorate General Research and the Directorate General Economics of the ECB for fruitful discussion and suggestions. A special thank to the anonymous referee of the ECB Working Paper Series who allows to improve the paper with his helpful remarks. The opinions expressed in this paper are those of the authors and do not reflect the views of the European Central Bank or the Eurosystem.

Vincent Brousseau
European Central Bank; e-mail: vincent.brousseau@ecb.europa.eu

Alain Durré
European Central Bank, IÉSEG-School of Management and LEM-CNRS (U.M.R. 8179); e-mail: alain.durre@ecb.europa.eu

Address	Kaiserstrasse 29, 60311 Frankfurt am Main, Germany
Postal address	Postfach 16 03 19, 60066 Frankfurt am Main, Germany
Telephone	+49 69 1344 0
Internet	http://www.ecb.europa.eu
Fax	+49 69 1344 6000

ISSN	1725-2806 (online)
EU Catalogue No	QB-AR-13-002-EN-N (online)

Abstract

In this paper we propose a new methodology to estimate the volatility of interest rates in the euro area money market. In particular, our approach aims at avoiding the limitations of currently available measures, i.e. the dependency on arbitrary choices in terms of maturity and frequencies and/or of factors other than pure interest rates, e.g. credit risk or liquidity risk. The measure is constructed as the implied instantaneous volatility of a consol bond that would be priced on the EONIA swap curve over the sample period from 4 January 1999 to 20 November 2012.

We show that this measure tracks well the historical volatility, in the sense that dividing the consol excess returns by this volatility removes nearly entirely excess of kurtosis and volatility clustering, bringing them close to an ordinary Gaussian white noise.

Keywords: consol rate, historical volatility, overnight money market, interbank offered interest rates
JEL classification: E43, E58, C22, C32

Non-technical summary

Excess returns of any asset traded in a liquid financial market are usually characterised by excess ‘peakedness’ (the so-called excess kurtosis) and by the fact that large fluctuations are followed by similar ones – commonly defined as the autocorrelation of squares or of absolute values of excess returns – (the so-called volatility clustering). These two statistical properties distinguish excess returns from random walk process increments (or, in the continuous time language, to be Gaussian white noises). This substantially complicates their representation and thus their estimation. In this paper, we test the assumption according to which the presence of excess kurtosis and volatility clustering could solely result from the variability of the instantaneous volatility. Should this be the case, dividing the excess returns by the corresponding instantaneous volatility would allow removing excess kurtosis and volatility clustering, hence allowing the normalized excess returns to be Gaussian white noise by construction. This test is applied to the case of the euro area money market interest rates in this paper over the sample period from 4 January 1999 to 20 November 2012.

For this purpose, the first step is to estimate the instantaneous volatility, which may prove to be a somewhat complex task. Indeed, an accurate measure of historical volatilities – computed as the standard deviation of excess returns across a time window demands a large time window while estimation of instantaneous volatility requests a small time window. Although an implied instantaneous volatility could be inferred from a sufficiently rich data set of quoted volatilities, such an approach involves additional technical difficulties when applied to market interest rates in opposition with other markets like equities or foreign exchange. To overcome these technical challenges, we propose to reconstruct the implied instantaneous volatility from available market data on the basis of the definition of a perpetual bond (namely a consol bond) that would have been priced on the overnight index swap (OIS) curve (namely the EONIA swap curve in the case of the euro area). The calculation of this EONIA swap curve consol volatility is then used to compare the statistical behaviour of the raw excess returns of such a perpetual (consol) bond with the normalised ones, i.e. obtained by dividing the raw excess returns by this volatility. Tests of our assumption are first made on the basis of simulations based on arbitrage free models. The method is then applied to the actual euro yield curve data. The results confirm our initial theoretical intuition, namely excess kurtosis and volatility clustering are (almost) entirely removed from EONIA excess returns.

The interest of this analysis is thus twofold. First, by removing the excess kurtosis and volatility clustering, it allows a much easier representation of the EONIA cumulated excess returns in the form of a Brownian motion. Second, it allows an accurate approximation of the instantaneous volatility over a long time horizon. Our findings are of relevant importance for those who have to monitor the dynamics of the market interest rates, both central bankers and practitioners.

1 Introduction

The information content of interest rates is important not only for practitioners, but also for monetary authorities. The estimation of the volatility of interest rates is equally important as it allows to gauge uncertainty surrounding market's expectations, notably as regards the future path of the monetary policy rate.

Despite the importance of the concept, no measure of interest rates volatility currently available in the literature is entirely satisfactory for several reasons. Most of them do not reflect solely the interest rate risk, but are usually contaminated by other factors that a rigorous analysis may not ignore, but should keep separated from the interest rate risk per se. Such factors can be liquidity funding risk, credit risk, convenience yield, collateral cost and so on. Finally – and perhaps more importantly – they rely on arbitrary choices in terms of time to maturity of some underlying rate instrument, possibly also of coupon frequency of that underlying instrument, and, in the case of market-implied volatilities measures, time to maturity of the option itself - those measures do not depict instantaneous volatilities. By construction, the common measures of historical (also said empirical or realized) volatilities, cannot capture the instantaneous underlying volatility variable. Renewed interest of volatility measures has occured with the works of Andersen and Bollerslev (1997) and Andersen and Bollerslev (1998) on realized volatility. However, despite the advantages of this measure in comparison with parametric measures, the (arbitrary) choice of the intraday time intervals remains a shortcoming. Furthermore, when applied to the short-term interest rate itself, the measure loses its interpretation in terms of a standard deviation of forthcoming excess returns[1].

One remains left with market implied volatilities. By contrast, market-implied volatilities can potentially deliver information about instantaneous volatility, even if options with intraday expiry date are not commonly traded: the possibility remains to extrapolate the instant value from the spectrum of the effectively traded values.

Implied volatilities remain nevertheless subject to the other drawbacks that we discuss in details in the paper. In particular, as they are derived from the (centralised market-based) options on future contracts on the interbank offered rate (BOR) fixings or bond interest rates (e.g. the German Bund which is largely used) and those derived from the (OTC-based) options on interest rate swaps (swaptions), they never represent a comprehensive measure of the pure volatility of interest rates, as they are usually derived from instruments that are not only affected by the movements of the EONIA curve, but also by funding liquidity risk, credit risk, convenience yield etc. Beneath, they still rely on the arbitrary choices in terms of maturity and/or of coupon frequency that we would judge preferable to eliminate.

The purpose of this article is thus to propose a measure of implied intan-

[1] This is because no instrument exists that would necessarily be equal to the EONIA of a trading day, on that trading day, and would also be led to be equal to the EONIA of the next day the following day. It follows that the difference between to subsequent EONIA fixings cannot be rigorously interpreted as a return.

taneous volatility that encompasses as much information as possible, while remaining as independent as possible of arbitrary choices such those related to e.g. maturity or frequency. In practice, estimating intantaneous volatility of interest rates may prove to be somewhat challenging. An accurate measure of historical volatilities – computed as the standard deviation of excess returns across a time window demands a large time window while estimation of instantaneous volatility requests a small time window. Although an implied instantaneous volatility could be inferred from a sufficiently rich data set of quoted volatilities, such an approach involves additional technical difficulties when applied to market interest rates in opposition with other markets like equities or foreign exchange. To overcome these technical challenges, we propose to reconstruct the implied instantaneous volatility from available market data on the basis of the definition of a perpetual bond (namely a consol bond) that would have been priced on the overnight index swap (OIS) curve (namely the EONIA swap curve in the case of the euro area). The calculation of this EONIA curve consol volatility is then used to compare the statistical behaviour of the raw excess returns of such a perpetual (consol) bond with the normalised ones, i.e. obtained by dividing the raw excess returns by this volatility. Our underlying assumption is that the presence of excess kurtosis and volatility clustering that characterised excess returns of any asset traded in a liquid financial market could solely result from the variability of the instantaneous volatility. Should this be the case, dividing the excess returns by the corresponding instantaneous volatility would allow removing excess kurtosis and volatility clustering, yielding the normalized excess returns to display Gaussian white noise process. Our findings demonstrate that our measure of instantaneous volatility captures pretty well the magnitude of the interest rates fluctuations in the near future while allowing to remove excess of kurtosis and volatility clustering from the excess returns of interest rates. The latter finding is powerful at least for two reasons. First, this finding suggests a significant efficiency of the interest rates option markets since our measure is an implied volatility. Second, it indicates that the actual dynamics of excess returns can be factorised in a very simple manner. It follows that their statistical properties can be explained, in a parsimonious manner, by the sole variability of the instantaneous volatility, which takes responsibility for both the excess kurtosis and the volatility clustering.

The remainder of the paper is organised as follows. Section 2 briefly recalls the main measures of volatility used in the literature and the standard measures of volatility of the euro area interest rates. Section 3 presents the construction of the new volatility measure against the theoretical background. Section 4 presents the data while Section 5 tests the statistical performance of the new measure against, first, simulated data and, second, the true empirical data for the case of the euro. Finally, Section 6 concludes.

2 Why and how to measure volatility of interest rates ?

Measuring and analysing volatility of interest rates is an important element of any financial market analysis. In the case of central banks, analysing the volatility of interest rates is of paramount importance, since monetary policy is usually implemented by steering short-term interest rates and by shaping the market expectations of the future values of those short rates. There are several reasons for the central bank's interest in this type of analysis.

First, it may offer insights into the effects played by the microstructure of markets and the efficiency with which they operate. For instance, comparing the volatility of interest rates at specific maturities with the average volatility across the whole maturity spectrum allows the central bank to detect atypical movements in some segments of financial markets, in particular the money market, which, in turn, could be related to imperfections in the market's structure or disturbances of the market functionning. Second, it may allow to assess ex-post the effectiveness and efficiency of the central bank's operational framework through which monetary policy decisions are implemented. It can thus help the central bank to draw conclusions about whether its liquidity management is implemented effectively (i.e. neutral from a market viewpoint with respect to the policy stance decided by the central bank decision-making bodies) and well understood by market participants. In some cases, it may lead to revision of the design of its operational framework as discussed in Durré and Nardelli (2008). Third, analysis based on rates volatility may help central banks to check the understanding of their own actions and communication by the market. Money market interest rates are henceforth of particular importance to central banks in signalling the monetary policy stance[2]. For example, increasing volatility (due to surprises) may blur the transmission of policy decisions along the yield curve over time, and this increased volatility may eventually translate into higher risk premium. Last but not least, interest rates volatility at longer maturities provides information about the underlying macroeconomic uncertainty and market participants' perceptions of the economic outlook.

Despite the importance of having an accurate and precise measure of volatility, the empirical literature is based on various types of measures which generally present the features of containing only a partial set of information and of being contaminated by the reflections of factors distinct from interest rates. Broadly speaking, one can classify the various existing measures of volatility within two main classes: model-based measures, and historical volatility measures. The

[2] For example, investigating how interest rate volatility evolves along the money market yield curve on days in which the central bank announces its decision regarding the appropriate level of policy interest rates offers a timely feedback on how the monetary policy decision is perceived by market participants. If money market interest rate volatility is relatively low both prior to and after the announcement of the monetary policy decision, this could suggest that the decision was expected, i.e. no surprise. It can thus be used as an indication of market's understanding of the central bank's operational framework. See for details the discussion in ECB (2006).

former are function of a specific model or econometric specification (including conditional volatility models and stochastic volatility models). The latter, quite simpler, essentially boils down to standard deviation of realised excess returns.

Somewhat outside the scope of the literature, there exists a third category of measures, the market based implied volatilities.

As pointed out in Andersen, Bollerslev and Diebold (2002), when measuring the volatility of a financial asset's return, the fist step is to determine the excess return. In financial literature, continuous-time stochastic volatility models assume an arbitrage-free process for prices (see Andersen, Bollerslev and Meddahi (2005)), and provide price processes related to a particular class of semi-martingales, which allow a unique decomposition of returns into a local martingale and a predictable finite variation process (see e.g. Andersen, Bollerslev, Diebold and Labys (2001) and Hansen and Lunde (2004)).

Suppose that p_t denotes the continuous-time logarithmic price process of some financial asset over a time interval T in which t denotes a compact time interval. The class of continuous-time stochastic volatility models for continuous price process of this asset can be generally expressed in terms of a stochastic differential equation (Ito's process) as:

$$dp_t = \mu_t \; d_t + \sigma_t \; dW_t \tag{1}$$

where μ_t and σ_t are time-varying random functions and W_t a standard Brownian motion. The drift term μ_t is (locally) predictable and of finite variation while σ_t satisfies technical conditions such as the requirement that $\int_0^t \sigma_u^2 du < \infty$ for any $t > 0$. As a result, the process $\int_0^t \sigma_\mu dW_\mu$ is a local martingale and p_t is a semi-martingale (see also Back (1991) and Protter (2004)). Note that the σ_t in equation (1)has the character of an instantaneous volatility, i.e. with maturity dt. The integrated squared volatility for such processes is defined as:

$$IV_t \equiv \int_0^t \sigma^2(t)dt \tag{2}$$

i.e. the cumulative sum of the square of instantaneous volatility σ. Being the cumulative instantaneous variance process of the asset, it is not directly observable (Andersen et al. (2002)). Analoguous specifications exist in discrete time. The continuous-time nature of equation (1) nonetheless presents the theoretical advantage of ensuring that the model's properties will not depend in a hidden manner of some particular discretization step. Naturally, as explained in Andersen et al. (2005), the estimation of returns and volatility necessarily involves a discretization of time, and available empirical data cannot be discretized at any arbitrary small step of discretization, even if Hansen and Lunde (2004) recalls that equation (2) could indeed be accurately approximated by the sum of high-frequency intra-daily squared returns over small contiguous intervals.[3] Hence, even the historical volatilities type of measures has necessarily a limited accuracy. Furthermore, one cannot a priori rule out that they depend on the choice of the sampling frequency, which makes them quite difficult to interpret. In effect, that dependency appears very strong when measured for what remains the proptotype of the highly liquid market, namely the exchange rate between euro and dollar (Brousseau (2006)). Finally, let us also mention that microstructure effects are likely to result into measurement bias.[4]

Turning now to the model-based measures, one way of approximating the true integrated volatility would be by estimating some abitrage-free factor model. This approach presents some practical drawbacks:

- First, it does not easily allow to exploit the information content of high-frequency data.
- Second, the tractability of this model often occurs at the cost of the presence of some unrealistic properties. For instance, some models assume that the yield curve is only explained by the short-term interest rate whereas it does not allow parallel shift in the yield curve (while, in the data, 90% of the variance of the yield curve is explained by parallel shift). This is particularly obvious in the case of the generic 1-factor affine model, namely the Duffie and Kan (1996) model with one factor (hereafter abbreviated as DK1), which can be seen as a generalisation of Vasicek (1977) and Cox, Ingersoll and Ross (1979). In the particular case of Vasicek (1977), although it does explain 100% of the variance of the curve by the parallel shift, this is at the cost of not being able to exclude negative interest rates of potentially high absolute values. The accuracy of volatility measures based on models with unrealistic assumptions is questionable.
- Third, the estimation procedure itself may be numerically unstable.

In this context, one is led to consider market-implied volatility measures. Although these measures involve some specific parametric option-pricing model, this model plays only a conventional role, which can be described as an encryption key allowing to translate an option price into a volatility or conversely. The

[3] Using the theory of quadratic variation (Protter (2004)), if the (log) price process p_t is a semi-martingale and $r_t = p_t - p_{t-1}$ denotes the return of this asset during a time period t, assuming that the trading day t is partitioned into m intervals of equal length $x = t/m$, then equation (2) corresponds to the following equation: $RV_t^m = \sum_{t=1}^{m} r_{t,m}^2$ where $r_{t,m}$ refers to intraday returns, $r_{t,m} = ({}_t p_{m+1} - {}_t\, p_m)$ is the return (or log-difference of prices) for the interval of equal length m and RV_t^m is defined as 'realised volatility'. It is thus assumed that m to be an integer and that RV_t^m converges uniformly in probability to IV_t, as $m \longrightarrow 0$. Simply stated, when assuming that the sampling frequency of $r_{t,m}^2$ approaches zero, then the realised volatility *consistently estimates the true (latent) integrated volatility* as mentioned in Andersen et al. (2005).

[4] In fact, based on observed prices from transaction data and/or quotes, the realised volatility can be affected by the microstructure of market, involving autocorrelation in intraday returns. The lower the sampling frequency, the stronger the effect of the microstructure. In that case, RV_t^m may significantly deviate from the (latent) true value of IV_t. It is thus crucial to define the data in a equally time-spaced observations with an adequate frequency, which remains however based on a arbitrary choice.

model's realism or lack thereof does not per se constitute much of a concern.[5] In addition, implied volatilities present the advantage to be forward looking in contrast with backward-looking historical volatilities. If one admits (as we tend to do) that at any given time the market as a whole is in possession of the best available information about the likely evolutions of the underlying instrument, then there is a case for using implied volatilities. Nonetheless, these measures are not immune from technical limitations:

- First, their availability depends on the existence of traded options, and may thus not exist for all instruments of interest (e.g. implied volatility for the euro area money market is only available for the 3-month EURIBOR futures contracts, and not for its 3-month EONIA swap counterpart).
- Second, the specifications of the underlying options limit by nature the volatility measure (e.g. its availability subject to the remaining duration of actual option contracts, or in other words, they are never instantaneous).[6]
- Third, the interpretation of the related volatility measure requires the knowledge of the specific option pricing model, which, by the ruling market convention, must be used for converting volatilities in option prices and conversely. This complicates the translation of volatility measures from one instrument to another, due to different market conventions referring to different market segments.[7]

In light of the pros and cons of each different existing volatility measure, the purpose of this paper is to construct a "pure" measure of volatility for the euro area money market. In contrast with existing measures, such a "pure" measure would meet the following characteristics: It would be (i) an implied volatility, (ii) an instantaneous volatility; and its underlying instrument (iii) would not reflect any other risk than that related to the sole interest rate risk; and (iv) would be independent of any arbitrarily chosen date of maturity, or time to maturity, or (in the case of a swap or a bond) coupon frequency. All those requirements leave us with virtually only one choice, the *implied instantaneous volatility of a consol bond* that would be priced on the EONIA swap curve. No consol bond priced on the EONIA curve exists in the market, and a fortiori no option on it, from which an implied volatility could be derived. Nevertheless,

[5] See also in particular Chapter 9 in Campbell, Lo and MacKinlay (1997) for a detailed discussion on volatility measures within derivative pricing models.

[6] For instance, market-exchanged options on futures contracts (like BOR traded on the LIFFE) have a fixed maturity date (which means, as time elapses, a decreasing time to maturity over the life of the option). By contrast, OTC-based options (like the swaptions) are usually quoted for some fixed times to maturity (which in turn means, as time elapses, increasing maturity dates).

[7] In particular, the implied volatility for the swaptions is based on the Black and Scholes formula for the interest rates whereas the implied volatility of bond instruments like the German Bund is based on the Black and Scholes formula for the price (a similar market convention is also used for other markets like the stock price indexes or the foreign exchange markets). Furthermore, the implied volatility derived from the EURIBOR futures contracts is based on the Back and Scholes for the quantity (100-interest rate).

the values that both the consol rate and the consol volatility would have if they were liquidly traded can effectively be reconstructed from available market data. This reconstruction is technically difficult, but conceptually simple.

3 A volatility measure based on the consol bond rate: theoretical background

The main objective of the construction of the consol volatility is to have a measure removing the main limitations discussed in the previous sections: (i) the dependency of arbitrary choices of maturity and / or frequency of the measure; (ii) the dependency of the underlying instrument of factors other than interest rates, e.g. credit risk, liquidity risk. But before proceeding to that reconstruction, we should examine what is the exact meaning on instantaneous consol volatility in a rigorous mathematical framework, and how it relates, within this framework, to other usual notions, such as zero-coupon prices or rates, or forward rates. We will focus on the continuous time, continuous prices framework.

Let us first recall some key mathematical notions related to the concepts of yield curve and of consol rate, before discussing the basic properties of a consol volatility. From those basic properties, it follows that the consol volatility reduces the consol excess returns to a Gaussian white noise, within the mathematical frameworkof a continuous-time arbitrage-free model. This hints that, in the real world, a correct measure of volatility should also be able to reduce the excess returns to a Gaussian white noise. This can be tested by examining whether dividing the excess returns by the volatility actually reduces the leptokurticity and the volatility clustering. This test will be illustrated with simulated data, for which the true volatility can be known ex ante, before being presented for the actual euro data. In this case, a success of the test indictes that the true volatility can be, and has effectively been, recovered.

3.1 Notational convention

The purpose of this paragraph is to recall and define the key notions of this paper within the continuous-time framework, as well as their basic mathematical properties. For convenience, the following convention is adopted: Latin letters denote dates while Greek letters refer to delay/duration between two dates.[8]

3.1.1 Zero-coupon and forward rates

Yield curves are formally defined as functions of a continuous time parameter, which associate an interest rate to a maturity within a (theoretically unbounded) maturity set. Yield curves are also usually expressed in mainly two ways: (i) as *zero-coupon* interest rate curves; or (ii) as *instantaneous forward* interest rate curves. These two ways are equivalent and convey the same quantity of

[8] For example, a bond observed at time t and having maturity date T shall have a time to maturity τ satisfying to the relation: $\tau = T - t$.

information. With $z(\tau)$ the zero-coupon interest rate at maturity τ, $f(\tau)$ the forward interest rate at maturity τ (whereby both interest rates are continuously compounded) and $P(\tau)$ the spot price of the zero coupon of maturity τ, i.e. the present value of one currency unit to be paid over τ, one can write that:

$$P(\tau) = e^{-\tau \; z(\tau)} \tag{3}$$

and

$$\frac{P'(\tau)}{P(\tau)} = -f(\tau), \tag{4}$$

which implies the following mathematical relationship between zero-coupon interest rate and the forward interest rate:

$$f(\tau) = z(\tau) + \tau \frac{dz(\tau)}{d\tau} \tag{5}$$

or, conversely, as the change of variables is duly invertible:

$$z(\tau) = \frac{1}{\tau} \int_0^{\tau} f(\theta) d\theta \tag{6}$$

From eq. (5) and (6), it follows two basic but important properties follow:

(a) $z(.)$ is constant if and only if $f(.)$ is constant. In this particular case, it means that the constant value taken by $f(.)$ is the same as the constant value taken by $z(.)$, which would imply a *flat* yield curve while the constant value taken by both $z(.)$ and $f(.)$ is called the *level* of the flat yield curve;

(b) irrespective of the shape of the yield curve, $z(0)$ and $f(0)$ are always equal and the corresponding unique value defines the short-term interest rate r, i.e. $r := z(0) = f(0)$.

Note also that one defines as parallel shift in the remainder of the analysis a transformation of the yield curve such as a constant is added to the zero-coupon interest rates, or, equivalently, of the forward rates following from eq. (5).

3.1.2 Consol bond and consol rate

A consol bond is defined as a perpetual (infinite horizon) bond paying continuously a constant rate of money, which is called the coupon flow. By definition, the consol price C is defined as the price of the consol bond divided by the coupon flow[9], which can be expressed, in terms of the yield curve, as follows:

$$C = \int_0^{\infty} P(\theta) \, d\theta \tag{7}$$

By substituting $P(\theta)$ by its value using eq. (3), the consol price becomes:

$$C = \int_0^{\infty} e^{-\theta \; z(\theta)} d\theta \tag{8}$$

[9] Such a normalisation is needed given the perpetual nature of this bond.

When using eq. (6), eq. (7) can also be alternatively expressed as follows:

$$C = \int_0^\infty \frac{dP(\theta)}{f(\theta)} d\theta \tag{9}$$

Similarly, the consol rate, y, defined as the yield of the consol bond, is the inverse value of the consol price, i.e.:[10]

$$y = \frac{1}{\displaystyle\int_0^\infty e^{-\theta\ z(\theta)} d\theta} \tag{10}$$

Furthermore, the consol duration, D, that is the duration (in the sense of Fisher and Weil (1971)) of the consol bond reflecting the sensitivity of the logarithm of the consol price to a parallel shift of the consol yield curve, can be expressed as:

$$D = \frac{\displaystyle\int_0^\infty \theta\ e^{-\theta\ z(\theta)} d\theta}{\displaystyle\int_0^\infty e^{-\theta\ z(\theta)} d\theta} \tag{11}$$

where the sensitivity of the consol rate to a parallel shift of the yield curve, χ, is given by the following product:

$$\chi = y\ D \tag{12}$$

This dimensionless number χ is equal to 1 in case of a flat curve. Empirical evidence shows that yield curves have usually χ smaller than, but close to, 1. We say that that a financial price is in constant terms (e.g. in constant euros, dollars, and so on), by opposition to current terms, when it is expressed in currency values (say e.g. EUR or USD) of a fixed reference date in the past rather than in currency values of the current date. The conversion of the value of an euro of one date into euros of another date is made with the compounding of the short-term interest rate between these two dates.

Finally, the consol wealth process, denoted A hereafter, is defined as the wealth, expressed in constant terms, of an ideal investor facing no transaction costs or short-selling restrictions who holds a portfolio of consol bonds in which she reinvests automatically the whole coupon flow at the prevailing market price. Therefore, the corresponding (infinitesimal) consol excess return, denoted (dA_t/A_t), is written:

$$\frac{dA_t}{A_t} = \frac{dC_t}{C_t} + \left(\frac{1}{C_t} - r_t\right) dt = y_t d\frac{1}{y_t} + (y_t - r_t)\ dt \tag{13}$$

whereby the first two terms refer to the nominal gain (or loss) due respectively to the change in the market price of the consol bond (dC_t/C_t) and to the coupon flow (dt/C_t) while the third term (rdt) refers to the carry cost of the position, i.e. holding the portfolio of consol bonds.

[10] By construction, the consol rate corresponding to a flat consol yield curve is equal to the level of that yield curve.

3.1.3 Volatility of a consol bond

The quadratic variation of a process X_t is denoted $[X_t]$, so that Ito's formula is:

$$df_t(X_t) = f'(X_t)\,dX_t + \frac{1}{2}f''(X_t)\,d[X_t] \tag{14}$$

The volatility of the consol bond, σ_t, is defined as:

$$\sigma_t^2 dt = \frac{d[C]_t}{C_t^2} = \frac{d[y]_t}{y_t^2} \tag{15}$$

The differential dL_t (with L_t denoting the logarithm of the consol wealth process, hereafter referenced as to consol performance) can be obtained by applying the Ito's iteration to eq. (13), i.e.:

$$dL_t = y_t\ d\frac{1}{y_t} + \left(y_t - r_t - \frac{\sigma_t^2}{2}\right)\ dt \tag{16}$$

which leads to the following identity using the Ito's formula:

$$y_t\ d\frac{1}{y_t} = -\frac{dy_t}{y_t} + \sigma_t^2 dt \tag{17}$$

Substituting eq. (17) into eq. (16) leads to express the consol excess return as:

$$\frac{dA_t}{A_t} = -\frac{dy_t}{y_t} + \left(y_t - r_t + \sigma_t^2\right)\ dt \tag{18}$$

which, by applying Ito's iteration to eq. (18), yields:

$$dL_t = -\frac{dy_t}{y_t} + \left(y_t - r_t + \frac{\sigma_t^2}{2}\right)\ dt \tag{19}$$

Finally, the normalized excess return is defined as:

$$dN_t = \frac{dA_t}{\sqrt{\frac{d[A_t]}{dt}}} \tag{20}$$

which, by combining eqs. (18), (19) and (20), yields:

$$dN_t = \frac{dL_t}{\sigma_t} + \frac{\sigma_t}{2}dt \tag{21}$$

We will then apply the mathematical specification of the consol rate, and the calculation of the corresponding volatility as discussed from eqs. (14) to (21), to standardise the volatility measure for interest rates in the money market.

3.1.4 Risk-neutral probability

Within the framework of Heath, Jarrow and Morton (1992), the dynamics of y_t, L_t and A_t are fully specified by the stochastic process of σ_t under risk-neutral probability. A_t must be a martingale, so eqs. (13) and (15) imply:

$$\frac{dA_t}{A_t} = \sigma_t dW_t \tag{22}$$

with W_t denoting a Wiener process, from which it immediately follows that:

$$dL_t = \sigma_t \; dW_t - \frac{\sigma_t^2}{2} dt \tag{23}$$

By combining eqs. (18) and (22) under the risk-neutral probability, the change of the yield of the consol, bond, dy_t, becomes:

$$dy_t = -\sigma_t \;\; y_t \; dW_t + y_t \; \left(y_t - r_t + \sigma_t^2\right) \; dt \tag{24}$$

By combining eqs. (15), (18), (20) and (24), it follows:

$$dN_t = dW_t \tag{25}$$

where N_t itself is also a Wiener process under risk-neutral probability. This implies that:

$$d[N]_t = dt \tag{26}$$

Eq. (26) is valid not only under risk-neutral probability but also under any equivalent probability to the risk neutral probability, as it pertains to the quadratic variation only. Finally, eqs. (21) and (25) yield that:

$$\frac{dL_t}{\sigma_t} + \frac{\sigma_t}{2} dt = dW_t \tag{27}$$

Note the left-hand side of equation (27) is defined as the normalized excess return.

It follows from equation (27) that the the normalized excess returns - i.e. the excess returns divided by the true volatility - should be akin a Gaussian white noise under risk-neutral probability. As we know, empirical excess returns of any financial asset usually differ from a Gaussian white noise by two properties: (i) their empirical distribution more kurtosis than the Gaussian distribution (the property of *leptokurticity*); and (ii) their absolute values (or their square values) have a positive serial correlation (the property of *volatility clustering*). To the extent that the Ito-process modelisation is a realistic representation of the consol rate dynamics, one should be able to remove those two properties, by dividing the excess returns by a correct measure of the underlying volatility, and that operation of normalization would recover the underlying Gaussian white noise process.

Attemps to approximate the true (unobserved implicit) volatility, namely σ_t in equation (27), are commonly done through non-parametric or parametric volatility measures of assets other than the consol rate. This way of doing is subject to two limitations. First, the time-to-maturity, and henceforth the midified duration, of the financial asset does not remain constant across the window used for the calculation of the corresponding volatility measure. Second, the window length involved in the calculation of the volatility measure being not instantaneous, one recovers at best an estimate of the average, in some sense, of the instantaneous volatility inside the window. Both things make the estimate of σ_t necessarily unprecise.

Our proposed measure of volatility circumvents those difficulties. On the one hand, unlike historical volatility, implied volatility can be extrapolated into instantaneous volatility. On the other hand, the consol bond remains identical as time slides given its perpetuity nature.

4 Data

The dataset contains only TARGET working days; it contains all the TARGET working days from 4 January 1999 to 20 November 2012, which represents 3560 TARGET working days. The financial instruments taken into account are handled in the OTC market. They consist into: short-term unsecured deposit of maturity 1-day (overnight, tom-next and spot-next), EONIA swaps from 1-week to 30-year, 6-month EURIBOR swaps for the corresponding maturities, at-the-money implied volatilities of options on the EURIBOR swaps.

We used the options on 6-month EURIBOR swaps with option maturity 1-month, 3-month, 6-month, 1-year, 2-year, 3-year, 4-year and 5-year, and with underlying EURIBOR swap maturity 1-year, 2-year, 3-year, 4-year, 5-year, 7-year, 10-year, 15-year, 20-year, 25-year and 30-year. Other maturities of options or of underlying swaps are represented in the quotes contributed by brokers, but their history may start at relatively recent dates, which makes preferable not to use them. Besides, the EONIA fixing is included in the dataset. Each instrument or fixing is identified in the Reuters database by a unique RIC (Reuters Instrument Code).

As the financial instruments taken into account are handled on the OTC market, we made use of quoted data, generally given as a bid-ask spread from which we retained only the mid. We gave a preference to quotes issued by the broker ICAP, and when not available, our primary fallback was the generic quote of Reuters, which contains the latest quote issued by a bank or broker at the time of its snapshot or of its contribution. In case of missing data, the data set is completed by a reconstruction of data as described in detail in Annex I.

5 Testing the robustness of the benchmark rule

To assess the correctness of a measure of consol volatility ('benchmark rule' hereafter), we will test whether the excess returns of the consol bond, when normalized by our volatility measure, resemble a Gaussian white noise process, i.e. that both leptokurticity and volatility clustering are essentially reduced. Removing the volatility clustering is not sufficient to assess the correctness of the measure. It is easy to see that for an indicator X that oscillates rapidly enough, the absolute value, and the square, of the ratio dL/X have small autocorrelations: the rapid oscillations may remove the volatility clustering but only at the expense of an increase of the leptokurticity. This makes necessary to require the reduction of *both* elements (volatility clustering and leptokurticity) in order to assess the quality of the volatility indicator.

Two formal tests are presented in this section. First, we conduct simulation of affine factor models (which allow the knowledge of the true volatility ex ante) and apply the test on the result of those simulations: this aims at checking that our implementation of the test actually behaves as it is supposed to. Second, we apply the test to the actual history of the EONIA curve, and to our reconstruction of the consol volatility, and assess by this way the quality of this reconstruction of the consol volatility.

5.1 Test based on simulation

The consol performance L is also well defined in the case of simulated data. In simulating the evolution of the yield curve under an arbitrage-free affine factor model, one can check whether (i) dL_t exhibits (or not) leptokurticity and volatility clustering; and (ii) $\frac{dL_t}{\sigma_t} + \frac{\sigma_t}{2}\, dt$ does exhibit less (or no) leptokurticity and less less (or no) volatility clustering. This type of exercise is of particular interest as it would provide a benchmark for the results of our test and hence a natural comparison point for the results based on the empirical data set.

5.1.1 General setting of the simulation exercise

Denote with t_1 and t_2 two consecutive TARGET working days. Excess returns are constructed by setting the cost of carry equal to the EONIA observed at the close of business of t_1. Normalized excess returns are constructed using the volatility computed at the close of business of t_1, (and not t_2).

To assess the existence of leptokurticity, we examine whether the excess of kurtosis differs from zero (contrary to a normal distribution where its value is zero). To test the existence of volatility clustering, we use the correlation of the absolute values of two consecutive returns, and the correlation of the square of two consecutive returns.

Both simulations are run on 3560 TARGET working days.

5.1.2 Simulation models

For the sake of simplicity, we will use the special case of arbitrage-free models known as affine models with constant parameters, with continuous time setting, and continuous trajectories. We will perform our simulations in the case of the generic 1-factor affine model and in the case of the generic 2-factor affine model.

In an arbitrage-free model, the zero-coupon bond price takes necessarily the form of the expectation:

$$P(t) = E_u[e^{-\int_u^t r_s ds}] \tag{28}$$

whereby the expectation refers to the so-called *risk-neutral probability*. Furthermore, the probability under which the short-term rate r_s actually diffuses should be equivalent, in the probabilistic sense[11], to the risk-neutral one. We will refer to that second probability as to the *data-generating probability*.

We have performed several attempts with different choices of parameters. The results are always that the normalised excess returns behave close to a gaussian white noise, and that the raw excess returns behave less close to a gaussian white noise. Yet the contrast between the normalized excess return's and the raw excess return's behaviours may be more or less pronounced. Typically, we find a excess of kurtosis ranging between zero and six for the raw case, and close to zero for the normalized case. The simulations that we present here will roughly correspond to a median case.

Simulation with 1-factor model - The generic 1-factor affine model is coincident with the Duffie and Kan 1-factor model (hereafter DK1). In the DK1 model, the risk-neutral probability is the solution[12]of the stochastic differential equation (SDE):

$$dr_s = (a - br_s)\ ds + \sqrt{c + r_s\ \nu^2}\ dW_s \tag{29}$$

with a and b positive constants, c a real constant, and ν a positive or zero constant, c and ν not being both equal to zero at the same time. By construction, this model has thus four parameters (a, b, c and ν) and one unique factor which can be identified to the short-term interest rate, r_s. It evolves between $-c\ /\ \nu^2$ and $+\infty$.

The data-generating probability should be equivalent (in probabilistic terms) to the risk-neutral probability. Since we are here in a continuous-time setting, the equivalence condition implies that the Brownian part of the data-generating probability is also provided by the expression $\sqrt{c + r_s\ \nu^2}\ dW_s$ with the same c and ν whereas it has no implication whatsoever as regards the form of the drift terms of the data-generating probability. Nevertheless, again for the sake of simplicity, one focuses on the specific case for which the stochastic differential equation defining the data-generating probability takes a form similar to equation (29), only allowing for different values for the drift parameters. For instance, denoting new values for the drift parameters by a^* and b^*, we have another data-generating probability given by:

$$dr_s = (a^* - b^* r_s)\ ds + \sqrt{c + r_s\ \nu^2}\ dW_s$$

[11] By definition, two probabilities are said to be equivalent if and only if they are defined on the same σ-algebra and, furthermore, they attribute the weight zero or the weight one to the same measurable sets.

[12] The solution of a SDE consists of a measure on the set of future trajectories - endowed with some suitable σ-algebra - and is therefore the relevant concept when defining expectations, which are essentially integrations w.r.t. that measure.

with a^* and b^* being strictly positive.

Note that the DK1 model is the generic case of the 1-factor affine model with constant coefficients. It is also the simplest possible arbitrage-free model of the yield curve, which includes two specific cases: (a) the original Vasicek (1977) model when ν is set to zero; (b) the Cox et al. (1979) (the CIR model hereafter) when c is set to zero.[13]

The functional form resulting from the DK1 model happens to be in itself rather realistic (see Brousseau (2002)). Yet, the DK1 model appears in practice relatively far from the real motion of actual yield curves, for at least four reasons.

First, the DK1 model assumes that the curve is entirely specified by the knowledge of its short-term rate (or also, of its zero-coupon rate or forward rate of any given maturity).

Second, the DK1 model implies that the zero-coupon rates or forward rates tend to a constant value for maturities tending to infinity. This constant value is called the long-term rate and is a function of the four parameters).

Third, it is well-known that the motion of empirical yield curve is constituted in major part by parallel shifts. This behaviour cannot be replicated in the frame of a DK1 model.

Fourth, the consol χ defined in equation (12) as the sensibility of the consol rate to a parallel shift of the curve, is always close to one for empirical yield curves. It is considerably smaller for DK1 curves.

Each of those features appears at odds with empirical observations. Our previous reasoning leads us to expect that the normalized consol excess return, as defined in equation (27), should display less excess of kurtosis and less volatility clustering than the raw consol excess return. Further, the normalized excess return should also have a standard deviation close to unity.

To perform the simulation, we need to choose values for the parameters and for the initial value of the factor. There is no compelling reason to choose one set of parameters rather than another. We have adopted values producing yields which are realistic for the euro, but this, strictly speaking, is not a constraint for the particular purpose that we how have of testing the test. We perform the simulation on the basis of the following values for the parameters and initial value of the factor:

parameters	annualized values
a^*	0.028
b^*	0.5
a	0.022
b	0.35
c	0.0002
ν	0.25
r_s *initial*	0.03

Based on this simulation model, the following Table reports the results for

[13] When ν is not equal to zero, the DK1 model can be rewritten as a parallel shift of the CIR model.

the (raw) consol excess returns and the normalized consol excess returns:

	Raw	Normalized
Std. Deviation	0.009	0.984
Excess of kurtosis	3.959	0.138
ACF lag 1 of abs. values	36%	1%
ACF lag 1 of squares	25%	3%

Simulation with 2-factor model - The generic 2-factor affine model is coincident with the model presented in Gourieroux and Sufana (2006) (hereafter GS2), but we will rephrase it with another choice of parameters and factors, in order to ensure formal consistency with the previous discussion.[14]

In the GS2 model, excluding again the case where the short-term rate is bounded away from zero, the risk-neutral probability can be seen as the solution of SDE described in equation (29):

$$\begin{pmatrix} dr_s \\ dp_s \end{pmatrix} = \left(\begin{pmatrix} a_1 \\ a_2 \end{pmatrix} - \begin{pmatrix} b_{11} & b_{12} \\ 0 & b_{22} \end{pmatrix} \begin{pmatrix} r_s \\ p_s \end{pmatrix} \right) ds + \sqrt{\begin{pmatrix} c + r_s\,\nu^2 & \frac{p_s \nu^2}{2} \\ \frac{p_s \nu^2}{2} & \frac{\nu^2}{4} \end{pmatrix}} d\mathbf{W}_s \qquad (30)$$

with a_1, a_2, b_{11}, b_{12}, b_{22}, c, ν, p_s and r_s satisfying certain constraints (as described in Annex II). The square root sign over the matrix has to be interpreted as the matrix square root operator, as opposed to an operator acting component by component. The model has seven parameters, a_1, a_2, b_{11}, b_{12}, b_{22}, c and ν, and two factors of which the first one is identified to the short-term rate r_s. As was the case for the 1-factor affine model, r_s evolves between $-c\,/\,\nu^2$ and $+\infty$. The second factor, denoted with p, has no particular economic interpretation. It evolves between $-\infty$ and $+\infty$, and its physical dimension is the same one as for a volatility, or equivalently, as the square root of a rate.

From the equation (30), it appears that if b_{12} is set to zero, the GS2 model is reduced to the DK1 model. The factor r_s in this case is not influenced by the dynamics of the other factor p and follows simply the solution of equation (29).

Again, while the mathematic structure of the model only obliges us to have the same Brownian part for the risk-neutral and data-generating probabilities, we focus nevertheless on data-generating probabilities sharing the same algebraic form with the risk-neutral one. The data-generating probability is then

[14] The connection between those parameters and factors and the ones appearing in the original paper is elaborated in Annex II.

given by the following equation:

$$\begin{pmatrix} dr_s \\ dp_s \end{pmatrix} = \left(\begin{pmatrix} a_1^* \\ a_2^* \end{pmatrix} - \begin{pmatrix} b_{11}^* & b_{12}^* \\ 0 & b_{22}^* \end{pmatrix} \begin{pmatrix} r_s \\ p_s \end{pmatrix} \right) ds + \sqrt{\begin{pmatrix} c + r_s\,\nu^2 & \frac{p_s \nu^2}{2} \\ \frac{p_s \nu^2}{2} & \frac{\nu^2}{4} \end{pmatrix}} d\mathbf{W}_s \tag{31}$$

where the parameters with an asterisk follow similar constraints than the parameters without asterisk.

We perform the simulation on the basis of the following values for the parameters and initial value of the factor:

a_1^*	0.028
a_2^*	0.03
b_{11}^*	0.5
b_{12}^*	−0.06
b_{22}^*	0.37
a_1	0.022
a_2	0.027
b_{11}	0.35
b_{12}	0.045
b_{22}	0.5
c	0.0002
ν	0.25
$r_{init.}$	0.03
$p_{init.}$	−0.15

The results for the consol excess returns on the left side and for the normalized excess return on the right side of equation (27) become:

Variables	Raw data	Normalized data
Standard deviation	0.009	0.991
Excess of kurtosis	4.136	0.122
ACF lag 1 of abs. v.	36%	1%
ACF lag 1 of squares	23%	3%

We obtain again the expected results, regarding the fact that leptokurticity and volatility clustering are present in the excess returns and removed from the normalized excess returns. Leptokurticity and volatility clustering reach values comparable to those of the 1-factor model. Yet, as we will see, they still cannot be compared with what is observed on empirical excess returns.

5.2 Test based on empirical data

Similarly to the general specifications recalled in the previous sections, the consol rate and the corresponding volatility for the EONIA is calculated over the period from 4 January 1999 to 20 November 2012, i.e. 3560 TARGET working days. Table 2 reports the results based on empirical data. A graphical representation of the volatility measure based on the consol rate specification is presented in Figure 1 at the end of the paper.

	Raw data	Normalized data
Standard Deviation	0.013	0.967
Excess of kurtosis	15.866	1.047
ACF lag 1 of abs. variation	36%	7%
ACF lag 1 of squares	42%	6%

As shown by Table 2, the excess of kurtosis and the volatility clustering exhibited by the normalised consol excess returns is substantially lower than those exhibited by (raw) excess returns. It is also interesting to underline that the excess of kurtosis and the volatility of the normalized consol excess returns based on empirical data appears even lower that the value they take for the raw excess returns (i.e. before nomalisation) in the case of the simulations.

Trends - In qualitative terms, the most striking difference between the case of empirical data and of simulated data is the following: in the case of empirical data, the consol volatility and the consol rate exhibit trends. The log volatility follows an increasing trend; the rate follows a decreasing trend.

This decreasing trend of the consol rate is reflected in an ascending trend of the normalized performance.

But a trend of performance could be engineered also in the simulations with a suitable choice of the difference between the risk-neutral probability and the data-generating probability. While this has not been the case for the parameters of the simulations reported above, we could generate such trends with other trial choices of parameters.

By contrast, a trend in the consol rate or volatility cannot. This is because, with our specifications (29) and (30), and the constraints that the parameters must satisfy (see Annex II), the factors follow an ergodic process, with a stationary distribution. Henceforth, while we could engineer ascending trends in performance, they did not appear as resulting from decreasing trends in the consol rate, as the consol rate did not follow any trend at all. This clearly demonstrates that the two things are in fact different even if, in the case of the empirical data, the trend of the performance appears as a clear reflection of the trend of the consol rate.

The strategy consisting into borrowing at short maturities and lending at long maturities has the reputation of being effective and of being a major source of revenues for the banking system. Our analysis of the data samples proves that the strategy has indeed be effective, as normalized performance increased at a regular pace. But the comparison between the simulated case and the empirical case also indicate that the source of that effectiveness is not necessarily the presence of a term premium.

Special events - The normalized excess returns are much closer, in terms of statistical signature, of the Gaussian white noise than of the raw excess returns.

Furthermore, they are constructed so as to present a standard deviation of one. A standardized and centred normal variable will only exceptionally be above 3.5 in absolute value. What is it with the consol normalized excess returns?

We find eight dates at which this has effectively happened. The ancient ones are now difficult to associate with recognizable market events, but in other cases the trigger is pretty clear. Table 3 reports these dates with their (most likely) corresponding economic event.

Date	Possible event
14 May 1999	
2 January 2001	
12 September 2001	9/11 attacks (CET)
1 November 2001	
2 January 2008	First new year after the turmoil
29 September 2008	Freezing of the EA money market
27 April 2010	Greek debt crisis
11 July 2011	Greek debt crisis (2^{nd} wave)

'Rich' vs. 'Poor' volatility - While over the whole sample, the standard deviation of the normalized excess returns is close to unity, two relatively large sub-periods are identified over which it deviates by some ten or fifteen percent from that value. The first sub-period ranges from 3 January 2005 to 31 May 2006 (where kurtosis and standard deviation amounts respectively to 0.04 and 1.10%). A second period spans from 2 January 2009 to 20 November 2012 (with kurtosis and standard deviation at respectively 2.21 and 0.84%). A period where the implied volatility is higher than the historical one is colloquially referred as a period of "rich" volatility. Conversely, the period is said to be of "cheap" volatility. Our data sample, covering the history of the euro, contains one period of cheap volatility covering the year 2005 and the first half of the year 2006, a time at which the risk perception by the market actors was generally subdued (also the FX volatility creped at historically low levels). In early 2009, when it was realized that the disruptions having followed the Lehman collapse were here to last, a period of "rich" volatility started, which is still prevailing today.

A parsimonious representation of the yield curve dynamics - Despite the existence of those periods of relatively rich or poor volatility, the kurtosis and the clustering are substantially removed when the excess returns are normalized by our reconstructed volatility measure. A convenient way to visualize to what extent this is cocnlusion holds true is to apply to the normalized excess returns, the inverse normal distribution function, yielding a number comprised between zero and one. We refer to that number as to the implied uniform variable; a name whose justification follows the following reasoning: if the normalization has effectively brought the statistical signature of the returns closer to the one of a Gaussian white noise process, then those numbers comprised between zero and one should appear uniformly distributed. The implied uniform variable thus appears uniformly distributed on any large enough subset of the time span of the sample.[15] It follows that the variability of the volatility suffices to explain the greatest part of both the kurtosis and the volatility clustering of the excess returns, i.e. two features that usually lead to deviations from a Gaussian white noise. The description of the dynamics of the yield curve becomes then strongly parsimonious.

6 Concluding remarks

Despite the importance to have an accurate measure of volatility to monitor financial markets, the standard measures currently available contain several shortcomings. This also applies to the measures of volatility used in general by most central banks to assess market's reactions (either to refinancing operations or to communication) as to monitor money market interest rates. By nature, standard measures of volatility only provide limited information related to one specific money market segments, which impedes an easy comparison across them.

In this context, this paper proposes a new measure of volatility derived from the specification of the consol rate for the EONIA swap curve in order to have an accurate estimation of volatility free from any model-based specifications and relaxed from maturity and frequency constraints. In addition, we demonstrate that this volatility measure is very close to the true (unobserved implicit) instantaneous volatility as it allows the excess returns of the consol rate to display a Gaussian white noise process (under risk-neutral probability or any similar probability) once normalised by this measure. This finding is quite powerful for several reasons.

First, our measure allows an homogeneisation of volatility measure (with a forward looking feature), hence providing information for the entire market without being restricted to one particular maturity. In the empirical part of the paper, we estimate this measure of volatility for the entire unsecured money market in the euro area which is obtained from the EONIA swap curve.

Second, by allowing to remove the leptokurticity and volatility clustering (which are natural features of excess returns of any financial assets), the excess returns normalised by our volatility measure is only sensitive to a limited number of exceptional events. In the illustrative example used in the paper, only eight special events are reported above the threshold of 3.5 standard deviation (out of 3560 observations, i.e. from 4 January 1999 to 20 November 2012), each of them related to exterme events.

Last but not least, our measure solves, by encompassing the volatility fluctuations of the entire market, the complexity of introducing a volatility measure in econometric models. More specifically, the restrictive nature of standard volatility measures (due to the strong link to a certain maturity and/or frequency) usually limit the use of volatility measure in times series regressions. It thus

[15] We thank the referee for his comment which has allowed us to add this demonstration in our analysis.

offers new research avenues as regards volatility transmission and/or assessment of market stress with a more manageable measure of volatility.

References

Andersen, T. and Bollerslev, T.: 1997, Intraday periodicity and volatility persistence in financial markets, *Journal of Empirical Finance* **4**, 115–158.

Andersen, T. and Bollerslev, T.: 1998, Answering the skeptics: yes, standard volatility models do provide accurate forecasts, *International Economic Review* **39**, 885–905.

Andersen, T., Bollerslev, T. and Diebold, F.: 2002, Parametric and nonparametric measurement. NBER Technical Working Paper No. 279.

Andersen, T., Bollerslev, T., Diebold, F. and Labys, P.: 2001, The distribution of realized exchange rate volatility, *Journal of the American Statistical Association* **96**.

Andersen, T., Bollerslev, T. and Meddahi, N.: 2005, Correcting the errors: volatility forecast evaluation using high-frequency data and realized volatilities, *Econometrica* **73**.

Back, K.: 1991, Asset prices for general processes, *Journal of Mathematical Economics* **20**, 317–395.

Brousseau, V.: 2002, The functional form of yield curves. ECB Working Paper Series No. 148.

Brousseau, V.: 2006, *The spectrum of Euro-Dollar*, Long Memory in Economics, Springer-Verlag edn, in Teyssière, Gilles. and Alan P. Kirman (eds.), chapter 3, pp. 69–107.

Campbell, J., Lo, A. and MacKinlay, A.: 1997, *The Econometrics of Financial Markets*, Princeton University Press, Princeton.

Cox, J., Ingersoll, J. and Ross, S.: 1979, Duration and the measurement of basis risk, *Journal of Business* **52**(1), 51–61.

Duffie, D. and Kan, R.: 1996, A yield factor model of interest rates, *Mathematical Finance* **6**(4), 379–406.

Durré, A. and Nardelli, S.: 2008, Volatility in the euro area money market: effects from the monetary policy operational framework, *International Journal of Finance and Economics* **13**(4), 307–322.

Fisher, L. and Weil, R. L.: 1971, Coping with the risk of interest-rate, *Journal of Business* **44**(October), 408–31.

Gourieroux, C. and Sufana, R.: 2006, A classification of two-factor affine diffusion term structure models, *Journal of Financial Econometrics* **4**(1), 31–52.

Hansen, P. and Lunde, A.: 2004, An unbiased measure of realized variance. Mimeo, Brown University.

Heath, D., Jarrow, R. and Morton, A.: 1992, Bond pricing and the term structure of interest rates: A new methodology for contingent claims valuation, *Econometrica* **60**, 77–105.

Protter, P.: 2004, *Stochastic integration and differential equations*, Springer-Verlag, New York.

Vasicek, O.: 1977, An equilibrium characterization of the term structure, *Journal of Financial Economics* **5**, 177–88.

ANNEX

Annex I - Determination of the consol volatility

1. - Data

1.1. - Working days and instruments

The dataset contains only TARGET working days; it contains all the TARGET working days from the 4 January 1999 to the 20 November 2012, so 3560 TARGET working days. The financial instruments taken into account are handled in the OTC market. They consist into: short-term unsecured deposit of maturity 1-day (overnight, tom-next and spot-next), EONIA swaps from 1-week to 30-year, 6-month EURIBOR swaps for the corresponding maturities, at-the-money implied volatilities of options on the EURIBOR swaps.

We used the options on 6-month EURIBOR swaps with option maturity 1-month, 3-month, 6-month, 1-year, 2-year, 3-year, 4-year and 5-year, and with underlying EURIBOR swap maturity 1-year, 2-year, 3-year, 4-year, 5-year, 7-year, 10-year, 15-year, 20-year, 25-year and 30-year. Other maturities of options or of underlying swaps are represented in the quotes contributed by brokers, but their history may start at relatively recent dates, which makes preferable not to use them.

Besides, the EONIA fixing is included in the dataset. Each instrument or fixing is identified in the Reuters database by a unique RIC (Reuters Instrument Code).

1.2. - Raw data

As the financial instruments taken into account are handled on the OTC market, we made use of quoted data, generally given as a bid-ask spread from which we retained only the mid. We gave a preference to quotes issued by the broker ICAP, and when not available, our primary fall-back was the generic quote of Reuters, which contains the latest quote issued by a bank or broker at the time of its snapshot or of its contribution.

1.3. - Completion

The data have been completed by reconstructed figures in three cases.

- When the history of long-term EONIA swaps was missing in the Reuters database, we reconstructed it on the following basis: before the 9 August 2007, it was assumed to have a constant and small spread with the corresponding EURIBOR swap, and between the 9 August 2007 and the first occurrence of the long-term EONIA swap in the Reuters database, it was reconstructed on the basis of both the corresponding Bloomberg data and of the Reuters data of the most similar instruments. The results of those reconstructions have been permanently integrated in the dataset used by the application.
- When the history of the options was missing, we reconstructed it according to the following procedure. The missing data, as it turned out, always pertained to short-term option on long term EURIBOR swaps, while longer

options on the same underlying swap were nevertheless available, as well as options of the same maturity on swaps of shorter maturity. We reconstructed then recursively, following the array of volatilities in the order of decreasing option maturities, then increasing underlying swaps maturity, assuming always that the reconstructed volatility and the previous one had the same ratio than the two corresponding volatilities of the next option maturity. The results of those reconstructions have not been permanently integrated in the dataset used by the application; instead, the application proceeds to their reconstruction at every session.

- When the history of an instrument was missing due to a London closing day, a case which occurred only in the recent years, the missing rates were reconstructed from the corresponding yield curve of the previous TARGET day, and the missing volatilities were reconstructed by copying their values of the previous TARGET working day. The results of those reconstructions have not been permanently integrated in the dataset used by the application; instead, the application proceeds to their reconstruction at every session.

1.4. - Descriptive statistics of the sample

The resulting data sample is described by the following descriptive statistics:

Number of TARGET working days	3560
Start	4 January 1999
End	20 November 2012
Number of RICs	187
Number of recomputed rates	72
Size	505868

2. - Algorithms

2.1. - Yield curves

2.1.1. - Composition

The EONIA curve is made of short-term unsecured deposits of maturity 1-day (overnight, tom-next and spot-next) and of EONIA swaps from 1-week to 30-year. It therefore spans a maturity interval ranging from 1-day to 30-year.

The EURIBOR swap curve is made of short-term unsecured deposits of maturity 1-day (overnight, tom-next and spot-next) and longer (between 1-week and 6-month) and of swaps versus 6-month EURIBOR from 1-year to 30-year. It therefore spans a maturity interval ranging from 1-day to 30-year.

The two curves evolve in a somewhat independent manner since the beginning of the crisis, 9 August 2007. Before that date, they were keeping a spread of small sized and which could be regarded, in first approximation, as being constant.

2.1.2. - Bootstrapping

We turn now to the construction of the yield curve from a collection of interest rate instruments. This construction, which consists into the successive calculation of zero-coupon prices or "discount factors", is termed "bootstrapping". For a detailed description of the bootstrapping, we refer to Brousseau (2002).[16] In a nutshell, the algorithm should be such that:

- it re-prices all the instruments contained in the above described yield curves to their exact original observed price,
- it can be entirely described through a finite (albeit not a priori specified) number of rates and dates.

The instruments to be integrated in the curve are sorted by ascending maturity. One constructs the curve by recurrence up to each maturity. Each of those maturities is termed a "knot point". Rates at intermediate points on the curve can be estimated by assuming a shape for the curve either in zero-coupon price or rate space. The choice of that interpolation rule constitutes the signature or the identification of the bootstrapping method. This choice is not conditioned by any theoretical reason, but it is conditioned by several practical reasons. The chosen rule should be such that:

1. The curve is smooth.
2. The curve does not have strong oscillations.

To those requirements, we add the supplementary one that:

3. The integral of the zero-coupon price between two knots can be computed in closed analytical form. The same holds for the integral of the zero-coupon price multiplied by the time to maturity.

The two first conditions are somewhat antagonist: an interpolation rule that favours one requirement will generally disfavour the other one.

The third requirement finds its justification in the necessity of performing consol-related calculations, which will be described in paragraph 2.3.2. below.

We have implemented and tested four different rules, satisfying to those three criteria, among which one is the method known as "unsmoothed Fama-Bliss" bootstrapping, whose interpolation rule results into stepwise constant forward rates. All of the methods performed rather close in term of the realistic aspect of the constructed curve. On the basis of some minute differences, we have chosen as default bootstrapping method, and used in this study, one of the three remaining methods.

2.1.3. - Extrapolation

It order to handle swaptions of maturity 5-year on 30-year swaps, one needs a yield curve covering a range of 35 years. Yet, quotes for the EONIA swaps stop at the 30-year tenor, at least for the price source that we have chosen to privilege. We have therefore extrapolated the EONIA curve.

While 35-year EURIBOR swaps quotes are actually available, we have nevertheless chosen to also extrapolate the EURIBOR swap curve to 35-year, to ensure the similarity of their treatment.

The extrapolation has always been made according to the following rule: denote with $z(\tau)$ the zero-coupon rate of maturity τ, continuously compounded and with day count actual/365 (i.e. whereby one year always means 365 calendar

[16] See paragraph 3.2.2. pp. 19-20.

days). The extrapolation of the curve has been achieved by adding the rate z(35) defined as 2*z(30)-z(25).

2.2. - Implied volatilities

2.2.1. - Converted volatilities

The primary input for the calculation of the consol volatility is a set of implied volatilities quoted in the market. The implied volatilities that we have used as raw data are those of EURIBOR swaptions, which are standardized options on 6 month EURIBOR swaps. They cannot be directly used, and this, for two reasons.

Firstly, the swaptions volatilities are implied by a Black and Scholes model in which the logarithm of the swap rate is assumed to be a Brownian motion. By contrast, while the consol volatility, resulting from the equations presented in the text, has to be implied by a Black and Scholes model of the standard variety, i.e. in which the logarithm of the zero coupon prices are assumed to be Brownian motions. A conversion will thus be necessary, changing the raw swaptions volatilities into other ones implied by the second model, and therefore directly comparable to, e.g., Bund options implied volatilities, or FX implied volatilities.

Secondly, the swaptions volatilities pertain to EURIBOR-linked instruments, whereby the consol volatility pertains to the EONIA curve. Also for that reason, a conversion will be needed.

The change of model cannot be done by an exact calculation (or that exact calculation would be too complicated). However, as we already mentioned, we know that the motion of the empirical yield curves are primarily composed of parallel shifts. We will then make an approximation and assume that those movements consist purely of parallel shifts. In this case, we only need to do the calculation at the first order, i.e. to multiply the raw swaptions volatility by the sensitivity of the log zero coupon price w.r.t. the log swap rate.

The two conversions are simultaneously achieved as follows:

The underlying of the option – which is a forward EURIBOR swap as the tenor of the option itself is not zero – is priced from the EURIBOR curve. Then, one computes the flat spread of that forward swap over the EONIA curve: this is the quantity of parallel shift to apply to the EONIA curve to let it price the forward swap at its present market price. That quantity is called the "z-spread" or the "yield curve spread" of the forward swap over the EONIA curve. We denote it with h. Remark that, like the zero coupon and forward rates of the yield curve, this z-spread is a continuously compounded rate with day-count actual/365, while an EURIBOR swap rate is neither continuously compounded nor has day-count actual/365. The needed conversion was operated at the stage of the construction by bootstrapping of the EURIBOR swap curve.

The sensitivity of the log zero coupon price (of the EONIA curve) w.r.t. the log swap rate (non compounded, with day-count actual/360) is then given by:

$$\frac{(\tau_2 - \tau_1)\, S\,(\tau_1, \tau_2, x)}{\frac{\partial S(\tau_1,\tau_2,x)}{\partial x}\,|_{x=h}} \qquad (1)$$

where:

- τ_1 refers to the time-to-maturity of the option,
- τ_2 is the time-to-maturity of the swap (e.g. 35-year for a 30-year swap forward in 5-year),
- S(τ_1,τ_2) is the swap rate (non compounded, with day-count actual/360), h is the swap's z-spread to EONIA curve (continuously compounded, with day-count actual/365), and
- S(τ_1,τ_2, x) is the swap rate (non compounded, with day-count actual/360) of a forward swap having z-spread x w.r.t. the EONIA curve (so that S(τ_1,τ_2, 0)= S(τ_1,τ_2)).

The product of the quoted volatility by this sensitivity yields a first order approximation of the "converted volatility". To be on the safew side, we actually used a higher order approximation (up to the 5-th power of the quoted volatility: a polynomial expression with three non-zero terms corresponding to the thirst, the third and the fivth power of the qulted volatility). However, that higher precision does not bring any visible difference.

We denote the "converted volatility" with $\sigma(\tau_1,\tau_2)$. For an instantaneous volatility $\sigma(0,\tau)$, we will use the short-hand $\sigma(\tau)$.

2.2.2. - Instantaneous volatilities

We are then left with a collection of implied volatilities for options tenors ranging between 1-month and 5-year.

We need to obtain instantaneous volatilities, i.e. the limit of the at the money volatility tends to zero. Because we reconstruct the history of consol performance and of consol normalized performance with a daily frequency, we need to understand instantaneity as meaning, in concrete terms, the interval between two subsequent TARGET working days. This implies that we reconstruct at the money volatilities of tenor one TARGET working day.

This is achieved by creating the cubic spline of the converted volatilities of tenors ranging between 1-month and 5-year and by extrapolating the resulting splined function to the tenor 1-day. The quantity to be splined is not directly the converted volatility, but the squared converted volatility multiplied by the time to maturity: this is the proper way to interpolate a term structure of volatilities.

2.3. - Consol-related calculations

2.3.1. - Formulas

The consol rate, consol duration, consol ksi and consol volatility rely on the computation of three integrals:

$$I_1 = \int_0^\infty P\,(\tau)\;\; d\tau \qquad (2)$$

$$I_2 = \int_0^{\infty} \tau \ P(\tau) \ d\tau \tag{3}$$

$$I_3 = \int_0^{\infty} \sigma(\tau) \ P(\tau) \ d\tau \tag{4}$$

where P(τ) is the zero coupon price for time-to-maturity τ, and $\sigma(\tau)$ is the zero coupon price instantaneous volatility for time-to-maturity τ. We keep the notation σ, without argument, to designate the consol volatility.

The computation of the consol rate, denoted with y, consol duration, denoted with D, and consol ksi, denoted with χ, following the derivation presented in the main text, boils down to the formulas:

$$y = \frac{1}{I_1} \tag{5}$$

$$D = \frac{I_2}{I_1} \tag{6}$$

$$\chi = \frac{I_2}{I_1^2} \tag{7}$$

The computation of the consol volatility involves the integral I3, but in all rigor the $\sigma(\tau)$ under the integral sign should be interpreted as vectors living in a space of some unknown, and possibly high, dimension, so that the quantity $\sigma(\tau_1)\ \sigma(\tau_2)$ is interpreted as the covariance of the increments of the log zero coupon prices of time-to-maturity τ_1 and τ_2 over the time interval dt.

Consistently with what we did for the conversion of volatilities, we will assume that the yield curves movements consist purely of parallel shifts. With the help of that approximation, that we use for the second time, the $\sigma(\tau)$ under the integral sign in (4) can be interpreted as scalars instead of vectors. The consol volatility is given by the formula:

$$\sigma = \frac{I_3}{I_1} \tag{8}$$

The $\sigma(\tau)$, even when we made them scalar numbers, have to be determined on the basis of the instantaneous volatilities resulting from the procedure described in 6. Yet those ones exist for eleven tenors of the underlying swaps, ranging from 1-year to 30-year. We take recourse again to a spline procedure, whereby we add to the list of available tenors the artificial tenor zero-day, associated with a value of the volatility of zero. For now, the quantity to be splined is directly the instantaneous volatility, and not the squared one multiplied by the time to maturity: we are not interpolating any more a term structure of volatilities, since all the involved ones are already instantaneous volatilities.

To compute the consol volatility of the theoretical yield curves of the affine models, used in the simulations, we do not need to make that approximation, because we do not need to remove the vector character of $\boldsymbol{\sigma}(\tau)$. We have indeed that:

$$\sigma(\tau) = V \cdot \boldsymbol{\nabla} \log\left(P(\tau)\right) \tag{9}$$

where V designates the volatility of the vector of factors: V is either $\sqrt{c + r\ \nu^2}$ for the 1-factor case or $\sqrt{\begin{pmatrix} c + r\ \nu^2 & \frac{p\ \nu^2}{2} \\ \frac{p\ \nu^2}{2} & \frac{\nu^2}{4} \end{pmatrix}}$ for the 2-factor case (see Annex II). The operator $\boldsymbol{\nabla}$ represents the gradient w.r.t. the vector of factors.

For the 1-factor model, the $\sigma(\tau)$ is anyway a scalar, and for the 2-factor model, the $\boldsymbol{\sigma}(\tau)$ is a vector of dimension 2, but both components can be explicitly computed.

2.3.2. - Numerical implementation

For what regards the consol rate, duration and ksi, in the case of the curves bootstrapped from empirical data, we have analytically computed the integrals between subsequent knot points of the bootstrapping. This was possible because the interpolation rules underlying the bootstrapping algorithms were chosen so as to allow for it, as was explained above. Then, by summing those partial integrals, we obtained the integral up to the latest available maturity: given the extrapolation procedure described in 6, that latest available maturity was 35-year. To complement the integral between 35-year and infinity, we have assumed that the zero-coupon rate remained constant. This way, all the complement integrals also could be obtained by closed form analytical formulas.

For what regards the consol rate, duration and ksi, in the case of the affine model curves, and for what regards the consol volatility, in both cases of empirical curve and model curve, we have discretized the integrals with a step of 30 calendar days and we have computed 1000 steps, reaching a maturity of circa 80-year. To complement the integral between that latest maturity and infinity, we have assumed that the zero-coupon rate remained constant, and that the instantaneous volatility remained strictly proportional to the time-to-maturity. This way, all the complement integrals also could be obtained by closed form analytical formulas.

2.3.3. - Consol performance and numerical performance

Let t_1 and t_2 to consecutive TARGET working days. The excess return between t_1 and t_2 is defined as the logarithm of the consol price in t_2 divided by the consol price in t_1 and by the actualisation coefficient (1+EONIA(t_1)/(t_2-t_1)/36000), where EONIA(t_1) is the EONIA published at the close of business of t_1 - so: non compounded, with a day count actual/360, expressed in percentage points - and (t_2-t_1) is the difference in calendar days between t_2 and t_1.

The consol excess return thus writes:

$$\log\left(\frac{\frac{100}{CONSOL(t_2)} + \frac{t_2 - t_1}{365}}{1 + \frac{(t_2-t_1)\ EONIA(t_1)}{36000}} \frac{CONSOL\,(t_1)}{100}\right) \tag{10}$$

where CONSOL(t), is the consol rate observed at the close of business of t, continuously compounded, with a day count actual/365, and expressed in percentage points. The normalized excess return is the quotient of the excess return between t_1 and t_2 and of the consol volatility observed at the close of business of t_1.

The consol performance and normalized performance designate respectively the cumulated sum of the excess returns and of the normalized excess returns.

Annex II - The parameters, factors and constraints of the generic 2-factor Affine model

The model, hereafter denoted GS2 model, is introduced in Gourieroux and Sufana (2006), where its main properties are explored, and the relationship with the 2-factor models belonging to the Dai and Singleton classification are elucidated. The GS2 model is proven in this paper to be the sole non-degenerated case of the 2-factor affine models; as non-degenerated models can be expected to produce more leptokurtic returns than degenerated ones, this property justifies that we use it in our simulation.

Gourieroux and Sufana (2006) writes the model using 8 parameters α_{01}, α_{02}, α_{11}, α_{21}, α_{22}, β_0, β_1 and β_2, and 2-factors x and y.

With those notations, the SDE defining the process writes (Proposition 2 page 37):

$$\begin{pmatrix} dx \\ dy \end{pmatrix} = \begin{pmatrix} \alpha_{11}x + \beta_1 \\ \alpha_{21}\ x + \alpha_{22}\ y + \beta_2 \end{pmatrix} dt + \sqrt{\begin{pmatrix} 1 & 2x \\ 2x & 4y \end{pmatrix}} d\mathbf{W} \tag{1}$$

where the short-term interest rate r is defined as an affine combination of the factors x and y, as (page 32 line 12):

$$r = \alpha_{01}\ x + \alpha_{02}\ y + \beta_0 \tag{2}$$

and where the parameters and factors satisfy constraints, among which the condition that α_{02} is positive (see page 42 line -10).

We change the parameters and the factors in order to have equations that visibly generalise the equations used for the 1-factor case. The transformation rules between the original parameters and the new parameters are:

$$\begin{aligned} \theta &= -\tfrac{\alpha_{01}}{\alpha_{02}} \\ a_1 &= \tfrac{\alpha_{01}^2(\alpha_{22}-\ \alpha_{11}) - \alpha_{01}\ \alpha_{02}(\alpha_{21}-\ 2\ \beta_1) + 2\ \alpha_{02}\ (\alpha_{02}\ \beta_2\ -\ \alpha_{22}\ \beta_0)}{2\ \alpha_{02}} \\ a_2 &= \tfrac{2\alpha_{02}\ \beta_1\ -\ \alpha_{01}\ \alpha_{11}}{2\sqrt{\alpha_{02}}} \\ b_{11} &= -\ \alpha_{22} \\ b_{12} &= \tfrac{\alpha_{01}\ \alpha_{22} - \alpha_{01}\ \alpha_{11} - \alpha_{02}\ \alpha_{21}}{\sqrt{\alpha_{02}}} \\ b_{22} &= -\ \alpha_{11} \\ c &= \alpha_{01}^2 - 4\ \alpha_{02}\beta_0 \\ \nu &= 2\ \sqrt{\alpha_{02}} \end{aligned} \tag{3}$$

The inverse transformation rules are:

$$\begin{array}{c} \alpha_{01} = -\frac{\theta\ \nu^2}{4} \\ \alpha_{02} = \frac{\nu^2}{4} \\ \alpha_{11} = -b_{22} \\ \alpha_{21} = (b_{11} - b_{22})\ \theta - \frac{2\ b_{12}}{\nu} \\ \alpha_{22} = -b_{11} \\ \beta_0 = -\frac{c}{\nu^2} + \frac{\theta^2 \nu^2}{16} \\ \beta_1 = \frac{b_{22}\ \theta}{2} + \frac{2\ a_2}{\nu} \\ \beta_2 = \frac{4\ b_{11}\ c}{\nu^4} + \frac{4\ a_1}{\nu^2} + \frac{(2\ a_2 + b_{12})}{\nu}\ \theta + \frac{(2\ b_{22} - b_{11})}{4}\ \theta^2 \end{array} \tag{4}$$

The transformation rules between the original factors and the new factors are:

$$\begin{array}{c} r = \alpha_{01} x + \alpha_{02} y + \beta_0 \\ p = \frac{\alpha_{01} + 2\alpha_{02} x}{2\sqrt{\alpha_{02}}} \end{array} \tag{5}$$

in which one naturally recognises in the second equation of (5) the formula (2). The inverse transformation rules are:

$$\begin{array}{c} x = \frac{2\ p}{\nu} + \frac{\theta}{2} \\ y = \frac{4\ \left(c + r_s\ \nu^2\right)}{\nu^4} + 2\ \frac{p}{\nu}\ \theta + \frac{1}{4}\ \theta^2 \end{array} \tag{6}$$

The SDE becomes:

$$\begin{array}{rcl} \begin{pmatrix} dr_s \\ dp_s \end{pmatrix} & = & \left(\begin{pmatrix} a_1 \\ a_2 \end{pmatrix} - \begin{pmatrix} b_{11} & b_{12} \\ 0 & b_{22} \end{pmatrix} \begin{pmatrix} r_s \\ p_s \end{pmatrix} \right) ds + \\ & & \sqrt{\begin{pmatrix} c + r_s\ \nu^2 & \frac{p_s \nu^2}{2} \\ \frac{p_s \nu^2}{2} & \frac{\nu^2}{4} \end{pmatrix}} d\mathbf{W}_s \end{array} \tag{7}$$

It is apparent that the parameter θ plays no role, and that consequently the model has truly seven parameters only. We have written the model using the seven parameters a_1, a_2, b_{11}, b_{12}, b_{22}, c and ν, and the two factors r and p. One makes appear the 1-factor model as a particular case by setting simply the parameter b_{12} to zero. Conversely, any system of parameters and factors such that one can obtain the 1-factor model as a particular case by setting simply one parameter to zero is necessarily isomorphic to (3)-(6).

The lowest possible value, that we will denote as $r_{\min}$, attainable by the short-term rate, admits the following expressions in terms of the original parameters and in terms of the transformed parameters:

$$r_{\min} \doteqdot \beta_0 - \frac{\alpha_{01}^2}{4\ \alpha_{02}} = -\frac{c}{\nu^2} \tag{8}$$

We turn now to the constraints that the parameters and variable should satisfy in the model. They are slightly more complicated than in the 1-factor affine model. Gourieroux and Sufana (2006) specify, page 37, that either:

$$(\alpha_{21} - 2\ \beta_1)^2 < 4\,(\alpha_{22} - 2\ \alpha_{11})\,(\beta_2 - 1)\ \ and\ \ (\alpha_{22} > 2\ \alpha_{11}) \tag{9}$$

or:

$$(\alpha_{21} - 2\ \beta_1)\ \ and\ \ (\beta_2 > 1)\ \ and\ \ (\alpha_{22} = 2\ \alpha_{11}) \tag{10}$$

Under the transformations (3) and (5) condition (9) is rewritten as:

$$\begin{gathered} 2\ b_{22} - b_{11} > 0 \\ a_1 \geq a_{1\,\min}^{(0)} \end{gathered} \tag{11.1}$$

where:

$$a_{1\,\min}^{(0)} = b_{11}\ r_{\min}\ + \frac{\nu^2}{4} + \frac{(2\ a_2\ +\ b_{12})^2}{4\ (2\ b_{22} - b_{11})} \tag{11.2}$$

whereby condition (10) is rewritten as:

$$\begin{gathered} 2\ b_{22} - b_{11} = 0 \\ b_{12} = -\ a_2\ \tfrac{b_{11}}{b_{22}} \\ a \geq a_{\min}^{(1)} \end{gathered} \tag{11.3}$$

where:

$$a_{\min}^{(1)} = b_{11}\ r_{\min}\ + \frac{\nu^2}{4} \tag{11.4}$$

We will focus on the cases that are interesting for our simulation: No degeneracy, existence of a lower bound and inexistence of an upper bound for the short-term rate, existence of a stationary distribution for the factors. We use for the simulation a parameter set falling into the first case (11.1) (11.3), as it is generic, while the second case (11.2) (11.4) is degenerated.

Relations (11.2) and (11.4) require that ν is not zero, but this is granted by the already mentioned positivity of α_{02}.

The parameter ν intervene always raised at the power 2 or 4, so under assumption $\alpha_{01}^2 > 4\alpha_{02}\ \beta_0$ and without further loss of generality, and in order to preserve the interpretation of ν as the volatility parameters of a Vasicek model and a CIR model, we add the conditions:

$$\nu > 0 \tag{12}$$

It is easy to see that the spectrum of the Jacobian of the drift field is $\{-b_{11}, -b_{22}\}$. So, taking into account that by (11.1) we have $b_{11}\ <\ 2\ b_{22}$, it is sufficient to add:

$$b_{11} > 0 \tag{13}$$

to obtain the case with existence of a stationary probability. Finally, the factors must satisfy a condition, written in the notations of Gourieroux and Sufana (2006) as $y > x^2$ and therefore, in the current notations, as conditions:

$$c + \left(r - p^2\right)\ \nu^2 > 0 \tag{14}$$

To conclude, the conditions fulfilled by the parameters and factors in order to implement our simulation shall be (11.1) with (11.2), (12), (13) and (14).

Figure 1 – Historical and implied consol volatility measures (logarithmic scale)