Documentos de Trabajo

Departamento de Análisis Económico y Finanzas

Department of Economics and Finance

Working Papers

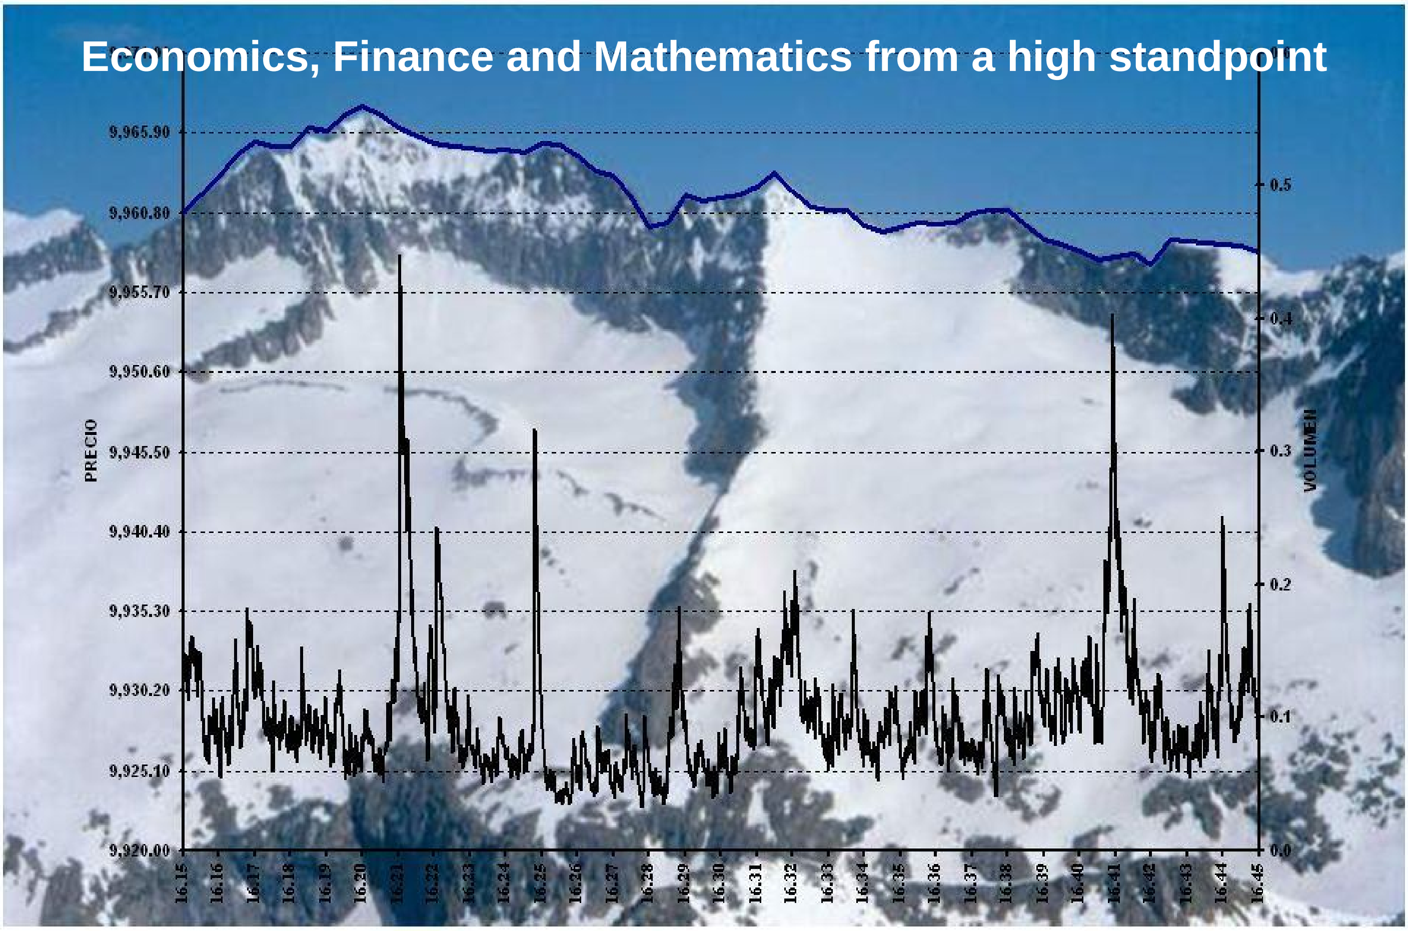

Constructing interest rate volatility indices over short- and long-term horizons

Raquel López
Universidad de Castilla-La Mancha

Eliseo Navarro
Universidad de Castilla-La Mancha

DT-DAEF 2011/3

Autores:
Raquel López (raquel.lopez@uclm.es)
Eliseo Navarro (eliseo.navarro@uclm.es)

Edita:
Departamento de Análisis Económico y Finanzas de la UCLM
Teléfono: 34 967 599 200. Extensión 2143
http://www.uclm.es/dep/daef/

Depósito Legal: AB-351-2009
ISSN: 1989-4856

Constructing interest rate volatility indices over short- and long-term horizons*

This version: May 23, 2011

Raquel López[a], Eliseo Navarro[b]

Abstract

We suggest the construction of a set of interest rate volatility indices (IRVIXs) that measure the future volatility of three-month tenor forward rates over horizons ranging from one to ten years ahead. This is a very important difference with respect to other indices such as VIX or VDAX in the equity market or the MOVE Indices for interest rates, which measure volatility over very short horizons (from one to six months ahead). We observe that the current financial crisis has had a severe impact on both short- and long-term IRVIXs. The potential uses of these indices are broad and include the introduction of derivative contracts, the estimation of the volatility term structure or use as leading indicators of the business cycle.

JEL codes: G11, G13, G17

Keywords: interest rates, volatility indices, caps and floors

* The authors would like to thank Anouk Claes, Marc de Ceuster, Antonio Díaz, Matías Gámez, Francisco Jareño, and the participants at the 2009 Spanish-Italian Congress of Financial and Actuarial Mathematics, the 2010 Conference of Mathematical and Statistical Methods for Actuarial Science and Finance, the 2010 Financial Management Association Doctoral Student Consortium, the European Financial Management Association 2010 Annual Meetings and the Business Economics Seminars in the field of finance at the University of Antwerp for useful discussions and comments. Raquel and Eliseo acknowledge the financial support provided by Ministerio de Educación y Ciencia grant ECO2008-05551/ECON. Raquel also acknowledges the financial support provided by Junta de Comunidades de Castilla-La Mancha grant PCI08-0089.

[a] Raquel López is in the Department of Economic Analysis and Finance, University of Castilla-La Mancha (E-mail: raquel.lopez@uclm.es).
[b] Eliseo Navarro is in the Department of Economic Analysis and Finance, University of Castilla-La Mancha (E-mail: eliseo.navarro@uclm.es).

1. Introduction

In 1993, the Chicago Board Options Exchange (CBOE) introduced the first implied volatility index for a stock market using data from options on the S&P100 Index: the S&P 100 Volatility Index (VIX). VIX very quickly became the benchmark risk measure for stock market volatility. It represents the implied volatility of an at-the-money (ATM) "theoretical" S&P 100 option with constant time to maturity (30 calendar days).[1] In 2003, the CBOE introduced the new VIX, which was constructed through an updated methodology and computed from options on the S&P 500 rather than the S&P 100.

Following the example of the CBOE, other options markets introduced their own volatility indices in Europe. In 1994, the Deutsche Börse created a volatility index for the German stock market, VDAX, constructed from options on the DAX Index.[2] In 1997, the MONEP (Marché des Options négociables de Paris) introduced the VX1 and VX6 indices to measure the uncertainty concerning the French stock market from options on the CAC-40 Index.[3] VDAX has a forecast horizon of 45 calendar days, whereas VX1 and VX6 look ahead 31 and 185 calendar days, respectively. An attempt to create implied volatility indices in the context of emerging markets can also be found in Skiadopoulos (2004) for the Greek derivatives market.

Nevertheless, far less attention has been paid by researchers and practitioners to equivalent indices for interest rate volatility. In particular, the two indices of reference for volatility in the fixed-income market are the Merrill Option Volatility Expectations (MOVE) Index constructed by Merrill Lynch and the Lehman Brothers Swaption Volatility

[1] Details regarding the construction of VIX are available in Fleming, Ostdiek, and Whaley (1995).
[2] See Lyons (2005) for an in-depth description of the construction process of VDAX.
[3] See Moraux, Navatte, and Villa (1999) for further details about these indices.

Index (LBPX). MOVE is constructed as a weighted average of the normalized implied volatility of one-month options on the two-year, five-year, ten-year, and 30-year Treasury securities.[4] Notice that although the underlying instruments are long-term assets, the maturity of the options and so the forecast horizon for volatility is only up to one month. Regarding the LBPX, it represents the weighted average of the normalized implied volatilities of a basket of liquid swaptions.[5] Thus, MOVE measures implied volatility over a forecast horizon similar to VIX, whereas LBPX does not have a specific forecast horizon for volatility.

An important advantage of interest rate derivatives markets over equity derivatives markets is that the formers contain option-like instruments with very long terms to maturity which allows us to develop a method to measure the uncertainty about the future evolution of a forward rate with horizons ranging from one to ten years ahead. Thus, IRVIXs are aimed at providing a pure measure of the future volatility of forward rates with specific maturities. For this purpose we work with data from the U.S. cap (floor) market. Although this information could be directly obtained from caplet (floorlet) quotations, these contracts are quite illiquid; thus, obtaining a complete enough range of caplets (floorlets) with different maturities can be complicated. Moreover, their prices, if available, can be distorted by this lack of liquidity.

It is important to point out that one of the requirements of any volatility index is that it must be constructed from a sufficiently liquid instrument. Consequently, IRVIXs are implemented using data from one of the most liquid interest rate derivatives: caps and

[4] Less known indices constructed by Merrill Lynch include the MOVE indices computed from options with time to maturity of three and six months, and the swaps MOVE (SMOVE) indices based upon similar swaps with maturities of one, three and six months.

[5] In addition to LBPX, Lehman Brothers provides a different index named LBOX. The former index computes volatility based on prices while the latter computes volatility based on yields.

floors. The construction of IRVIXs from these instruments can give a much more accurate indication of the actual uncertainty regarding the future behavior of interest rates for a wide range of maturities, without the intrusion of the noise caused by liquidity.

Of course, we might be tempted then to directly use cap (floor) implied volatilities to measure uncertainty about the future behavior of interest rates. However, they present an important disadvantage over IRVIXs. As we show in the next section, cap (floor) volatilities, often referred to as flat volatilities, are in fact a mixture of the future volatilities over different term horizons of a set of forward rates with diverse terms to maturity. Thus, it is difficult to assess what is really behind these flat volatilities.

This way, the use of data from the cap (floor) market poses the problem of having to deal with a contract where the underlying rate is not a single forward rate but a set of forward rates with consecutive maturities. Recall that caps (floors) consist of portfolios of concatenated caplets (floorlets), i.e., options on forward rate agreements with consecutive maturities.

Therefore, the construction of each IRVIX, which is aimed at providing a measure of the uncertainty of the future evolution of a particular forward rate, will require the extraction of information about the prices of the caplets (floorlets) that compose caps (floors), as caplets (floorlets) are the contracts that do have an underlying specific forward rate.

In any case, apart from its high liquidity, the use of caps (floors) presents another advantage. The resulting indices can provide information about the future performance of interest rates for very long horizons. In fact, we propose the construction of a set of alternative indices with horizons ranging from one to ten years ahead. This constitutes one of the main differences with respect to MOVE, which captures the uncertainty about long-term yields over a very short horizon (one month). Thus, IRVIXs can be particularly useful for obtaining information about the long-term level of uncertainty of economic agents with respect to one of the main economic and financial variables: interest rates. Short-term IRVIXs changes would indicate temporary changes regarding the uncertainty of the future behavior of interest rates. However, if changes in short-term IRVIXs are accompanied by changes in long-term IRVIXs, this would be a signal of long and permanent periods of turbulence in interest rate markets. In fact, short-term IRVIXs reacted very quickly to the beginning of the current financial crisis, whereas long-term IRVIXs responded more slowly to the financial turmoil.

The set of indices is obtained through the implementation of a methodology similar to that used in the equity derivatives market to construct well-known volatility indices, such as VIX or VDAX.

Although their most direct application consists of providing a measure of the uncertainty regarding the future evolution of forward interest rates over different horizons, there is also a wide range of other potential applications.

One of the most promising applications is the possibility of introducing futures and options on the indices, as occurred after the launch of VIX. In February 2006, options on VIX began trading on the CBOE, following the previous introduction of VIX futures on the CBOE Futures Exchange (CFE) in 2004. According to Areal (2008), in practice, these derivatives can be used in turn to create hedging strategies against changes in volatility or to speculate on changes in the market volatility. As shown later on, the dramatic changes in interest rate volatility experienced since the beginning of the current financial crisis can make IRVIX a powerful instrument to protect against these interest rate volatility changes.

In fact, the statistical analysis of IRVIXs shows that they share many similarities with VIX, thus indicating their potential use as the underlying instruments of derivative contracts. In this sense, it must be pointed out that short-term IRVIXs present an outstanding negative relationship with their underlying forward rates. This finding is particularly important and should be taken into account when using some derivative contracts, such as caps or floors, for hedging against interest rate movements, since these interest rate changes will take place together with important changes in interest rate volatility.

Other applications of IRVIXs are their potential use for the estimation of the volatility term structure, which is one of the main topics in the valuation of interest rate derivative securities. IRVIXs can also be used for analyzing the role of implied volatility in forecasting future realized volatility[6] or for testing the additional information content of interest rate volatility with respect to the future state of the economy, improving the broadly documented forecasting ability of the term structure of interest rates with respect to the business cycle.

The structure of the paper is as follows. The next section is focused on caps (floors) valuation within the Libor Market Model (LMM) framework. This model is used by practitioners to quote these contracts. Moreover, caps and floors are quoted in terms of implied volatilities, and contract prices are then obtained through the application of the well-known Black pricing formula. Thus, we provide a brief description of the assumptions of LMM. In Section Three, we discuss how to implement IRVIX, using the methodology

[6] See Poon and Granger (2003) for a wide review on this topic for different markets and time periods. See, among others, Fleming, Ostdiek, and Whaley (1995), Moraux, Navatte, and Villa (1999), Bluhm and Yu (2001), Corrado and Miller (2005) and Becker, Clements, and White (2007) for an analysis of the forecasting ability of implied volatility indices.

applied in equity derivatives markets as a reference. Section Four describes the database and the methodology applied for the construction of the set of indices. In Section Five, the behavior and statistical properties of the volatility indices are analyzed, and finally, Section Six provides a summary of the study.

2. Caps and floors valuation. The LMM and the Black formula

A forward rate agreement (FRA) is the underlying of one of the simplest interest rate options: caplets (floorlets). A FRA can be defined as an agreement between two parties at time t to exchange at time $T_i+\tau$ an amount of money proportional to the difference between a strike, K, agreed upon at time t, and the floating interest rate outstanding at T_i, $L(T_i,T_i+\tau)$. Thus, the cash flow generated at $T_i+\tau$ by this contract is given by

$$NP \cdot \left[L(T_i,T_i+\tau)-K\right] \cdot \tau \, , \tag{1}$$

where NP is the notional principal of the contract, and τ is the tenor interval.

A caplet[7] is an option on this FRA that will be exercised only if $L(T_i,T_i+\tau)$ is greater than the strike K. Thus, the cash flow of this option at $T_i+\tau$ will be

$$NP \cdot \left[L(T_i,T_i+\tau)-K\right]^{+} \cdot \tau \, , \tag{2}$$

where

$$\left[L(T_i,T_i+\tau)-K\right]^{+} = Max\left\{L(T_i,T_i+\tau)-K,0\right\}. \tag{3}$$

The LMM and market quotations assume that the forward interest rate $f(t,T_i,T_i+\tau)$[8] follows a lognormal stochastic process.[9] Taking into account that the limiting value of the

[7] Floorlets are defined in an analogous way, but, in this case, cash flows are generated when the floating rate is below the strike K.

forward rate when t approaches T_i is equal to the floating interest rate $L(T_i,T_i+\tau)$, and assuming there are no arbitrage opportunities, it is easy to derive the well-known Black's formula used to value this sort of contracts.[10] According to this formula, the value of a caplet per unit of notional principal is given by

$$Caplet(t,T_i,\tau,K,\sigma_{i,Black}^{K}) = \left[f(t,T_i,T_i+\tau)\cdot N(h_1) - K\cdot N(h_2)\right]\cdot P(t,T_i+\tau)\cdot\tau\,, \tag{4}$$

where

T_i is the exercise date of the caplet (and the maturity date of the underlying forward rate),

τ is the tenor of the underlying forward rate (and $T_i+\tau$ is the maturity date of the caplet),

K is the strike rate,

$P(t,T_i+\tau)$ is the price at t of a unit-zero coupon bond with maturity at $T_i+\tau$,

$N(\cdot)$ is the cumulative normal distribution,

$$h_1 = \frac{\ln\left[f(t,T_i,T_i+\tau)/K\right] + \frac{1}{2}\cdot(\sigma_{i,Black}^{K})^2\cdot(T_i-t)}{\sigma_{i,Black}^{K}\cdot\sqrt{(T_i-t)}} \tag{5}$$

$$h_2 = \frac{\ln\left[f(t,T_i,T_i+\tau)/K\right] - \frac{1}{2}\cdot(\sigma_{i,Black}^{K})^2\cdot(T_i-t)}{\sigma_{i,Black}^{K}\cdot\sqrt{(T_i-t)}}, \tag{6}$$

and

---

[8] We will refer to T_i as the maturity date of the forward rate and τ as the tenor of the forward rate.
[9] See Brigo and Mercurio (2006) for an extensive review of LMM.
[10] See, for instance, Díaz, Meneu, and Navarro (2009) for a brief overview of the derivation of the Black formula to value these securities.

$\sigma_{i,Black}^{K}$ is the so-called Black implied volatility of a caplet with exercise date T_i and strike K.

Black implied volatility can be understood, within the LMM, as an average of the instantaneous volatility of the log of the forward rate $f(t,T_i,T_i+\tau)$ over the period $[t,T_i]$. In particular,

$$(\sigma_{i,Black}^{K})^2 = \frac{\int_t^{T_i} \sigma^2(u,T_i)du}{(T_i - t)}, \tag{7}$$

where $\sigma(t,T_i)$ is the instantaneous volatility at t of the lognormal process followed by the forward rate $f(t,T_i,T_i+\tau)$.

According to the LMM and Equation (7), the implied volatility of caplets should be the same for all caplets with the same term to maturity, independent of the strike K. However, in practice, the market implied volatility of caplets (with everything else equal) varies with the strike rate K.

In any case, the instruments we are going to use to construct the interest rate volatility indices are not caplets (floorlets), but much more liquid and popular interest rate derivatives in the over-the-counter (OTC) markets: caps (floors). According to the Bank for International Settlements[11], the outstanding notional amount of OTC interest rate options (caps, floors, collars and corridors) in December 2009 was $48.8 trillion.

A cap can be understood as a portfolio of caplets with the same strike and tenor but with concatenated maturities. For instance, a two-year cap consists of a chain of seven caplets with exercise dates in three, six, nine, 12, 15, 18 and 21 months, respectively, all of

[11] http://www.bis.org/

them with a tenor of three months (τ = 3 months). The maturity date of each caplet coincides with the exercise date of the following one.

The payoffs generated by a cap can be described as follows. On the exercise date of the first caplet that composes the cap, the floating rate of the contract is observed and compared to the strike. If the floating rate is greater than the strike, then on the second reset date the seller of the cap pays the holder the difference between the floating rate and the strike multiplied by the notional principal and the tenor. If the floating rate is less than the strike, there is no payoff from the cap. Thus, through the life of a cap, payments are due at the end of each tenor interval, although the amount is known at the reset date (at the beginning of the tenor interval) when the floating interest rate is observed.[12]

Then, the price at t of an n-year cap with strike K can be obtained as the sum of the values of the caplets that make it up. That is,

$$Cap(t, T_{n\cdot k}, K) = \sum_{i=1}^{n\cdot k-1} caplet(t, T_i, \tau, K, \sigma_{i,Black}^{K}), \tag{8}$$

where k can take values of 2 or 4 depending on the length of the tenor interval, which is six or three months, respectively,[13] and $T_1, T_2, \ldots, T_{n\cdot k-1}$ are the reset dates of the cap that coincide with the exercise dates of the caplets that compose the cap and $T_{n\cdot k} = T_{n\cdot k-1}+\tau$, i.e., the date that the last cash flow will be due if $L(T_{n\cdot k-1}, T_{n\cdot k-1}+\tau) > K$.

However, quotations in the cap market are computed assuming that the volatility of all the caplets that compose a particular cap is the same. In fact, an n-year cap with strike K is quoted by the market through the so-called flat volatility, which is the constant value

[12] Caps are usually defined so that the initial floating rate, even if it is greater than the cap rate, does not lead to a payoff on the first reset date (Hull 2009).
[13] In the case of the US market, caps have a three-month tenor.

$\sigma^K_{n,flat}$ that equals the sum of the values of all of the caplets that compose the cap to its market price, i.e., the value $\sigma^K_{n,flat}$ such that

$$Cap(t, T_{n\cdot k}, K) = \sum_{i=1}^{n\cdot k-1} caplet(t, T_i, \tau, K, \sigma^K_{n,Black}). \tag{9}$$

Therefore, flat volatilities cannot be considered to be a pure measure of the future evolution of volatility of a forward rate; rather, they are a mixture of the average future volatilities of a set of forward rates with consecutive terms to maturity.[14]

Theoretically, flat volatilities should be the same for all caps with the same term to maturity, but in practice we find that flat volatilities depend heavily on the strike rate *K*, giving rise to volatility surfaces.

3. Implementing IRVIX

As stated above, the objective of IRVIX is to provide an index that reflects the uncertainty about the future behavior of a forward rate with a given term to maturity.

To obtaining this information from the cap (floor) market, two problems have to be overcome. The first one is, as mentioned before, that flat volatilities are not a pure measure of the expected future volatility of a specific forward rate but rather a mixture of the future volatility of a set of forward rates. For instance, following the example of the former section, the flat volatility of a two-year cap is a mixture of the average future volatility of three-month tenor forward rates with maturities in three, six, nine, 12, 15, 18 and 21 months.

[14] The difference between $\sigma^K_{i,Black}$ and $\sigma^K_{n,flat}$ is similar to the difference between zero-coupon rates and the yields to maturity of coupon-bearing bonds.

The second problem is that the flat volatility of an *n*-year cap varies considerably with the strike rate K, i.e., we have to deal with the problem of the smiles and smirks, a problem with specific features in this market due to the different meaning of what should be understood, in this case, by an ATM option. Recall that caps do not have a single underlying but a set of them (all forward rates involved in the valuation of the set of caplets that makes up the cap). Thus, we have to determine the strikes of the caps to be used for the construction of IRVIX.

These two problems will be faced in a two-step process. The first step consists of recovering the implied prices (and volatilities) of the individual caplets that compose the caps using a stripping procedure.[15] The second step will be to apply interpolation techniques, similar to those used in the stock market to construct VIX or VDAX from the implied volatilities of options with different strikes.[16]

In particular, the methodology applied in the stock market for the construction of these indices consists of interpolating the implied volatilities of a set of nearest to-the-money call and put options at the two nearest maturities to the constant time to expiration established for the construction of the indices (30 calendar days in the case of VIX and 45 calendar days in the case of VDAX). However, in our case, caplets recovered from caps have a constant life period (from t to T_i, exercise date of the option), and thus, the only criterion we need to consider in the selection of the caps involved in the construction of IRVIX is the strike.

[15] See Hernández (2005).
[16] Since 2003 a new methodology has been developed to construct volatility indices based on variance swap replication techniques (see Carr and Madan, 1998). However, this new method assumes that short-term interest rates are non-stochastic. If this assumption can be considered reasonable for valuing stock options or options on long-term bonds, it is unacceptable for IRVIXs, particularly for those with a long-term horizon.

The stripping process consists of obtaining the price at time t of a caplet with a strike K and reset date T_i, $caplet(t,T_i,\tau,K,\sigma_{i,Black}^{K})$, by subtracting the prices of two consecutive caps with the same strike K. That is,

$$Caplet(t,T_i,\tau,K,\sigma_{i,Black}^{K}) = Cap(t,T_{i+1},K) - Cap(t,T_i,K)\,, \tag{10}$$

where $Cap(\cdot)$ are defined as in Equation (9). It must be emphasized that these two caps must have the same strike K and consecutive maturities.

Once the price of the caplet is obtained, it is easy to derive the corresponding implied volatility from the Black pricing formula. Note that we obtain different implied Black volatilities for the same caplet, depending upon the strike rate K.

To determine the strikes of the caps to be used for the construction of IRVIX, we should point out that the most liquid caps are the ATM ones. In this market, an n-year cap is said to be ATM if the strike of this instrument equals the fixed rate of a swap that has the same payment days as the cap.[17] However, if we have to apply the technique described above to obtain caplet prices, we cannot use ATM caps because two consecutive caps would have different strikes to the extent that swaps with different maturities usually have different fixed rates.

Therefore, we have to deal with the problem of determining the strike of an ATM caplet. We propose to use those available caps with strikes closest to the outstanding forward rate $f(t,T_i,T_i+\tau)$ defined as

$$f(t,T_i,T_i+\tau) = \left(\frac{P(t,T_i)}{P(t,T_i+\tau)} - 1 \right) \cdot \frac{1}{\tau}\,, \tag{11}$$

[17] See, for instance, Hull (2009).

where $P(t,T_i)$ and $P(t,T_i+\tau)$ are the prices at t of unit zero-coupon bonds with maturities at T_i and $T_i+\tau$, respectively.

Particularly, we will use caps with strikes immediately over and below $f(t,T_i,T_i+\tau)$, and we will refer to them as K^A and K^B, respectively, with $K^B < f(t,T_i,T_i+\tau) < K^A$.

By using Equation (10), we can obtain the prices of caplets with strikes K^A and K^B. These two caplets are the first out-of-the-money caplet (the one with strike K^A) and the first in-the-money caplet (the one with strike K^B). Then, we proceed to compute their implied volatilities using the Black formula. We denote these two implied volatilities by $\sigma_{i,Black}^{K^A}$ and $\sigma_{i,Black}^{K^B}$.

Finally, we obtain the implied volatility of a theoretical caplet with the strike equal to the current forward rate $f(t,T_i,T_i+\tau)$ by linear interpolation, using the following formula:

$$IRVIX(t,T_i) = \sigma_{i,Black}^{K^B} \cdot \left(\frac{K^A - f(t,T_i,T_i+\tau)}{K^A - K^B} \right) + \sigma_{i,Black}^{K^A} \cdot \left(\frac{f(t,T_i,T_i+\tau) - K^B}{K^A - K^B} \right), \tag{12}$$

where IRVIX(t,T_i) represents the annualized implied volatility of an ATM caplet with a constant horizon (T_i-t). Thus, IRVIX(t,T_i) can be understood as a measure of the uncertainty about the future evolution of the forward rate $f(t,T_i,T_i+\tau)$ over the period $[t,T_i]$.

4. Data and methodology

In conducting this study, we used two sets of data from the U.S. fixed-income market provided by *Reuters*. The first one includes market-implied flat volatilities of caps (floors) for different strikes and terms to maturity.[18] The second one consists of zero-coupon curves (discount factors bootstrapped from the most liquid rate instruments available, a combination of deposits, liquid futures and interest rate swaps). Daily data were collected for the period from August 26, 2004 to January 30, 2009.

Flat volatilities correspond to caps (floors) with maturities in one, two, three, four, five, six, seven, eight, nine, ten, 12, 15 and 20 years and with the following range of strike rates: 0.01, 0.015, 0.02, 0.025, 0.03, 0.035, 0.04, 0.05, 0.06 and 0.07. These strikes cover the range of values of the forward rates during the sample period such that there will be always a strike above and below the outstanding forward rates.

One of the problems that appear when we try to obtain caplet prices using the stripping technique is that we need the prices of consecutive caps with the same strike. In the case of the American market, the tenor interval is three months; thus, we need prices of caps with maturities every three months. However, markets only provide caps with annual terms to maturity, i.e., with an integer number of years to maturity.

Therefore, interpolation and extrapolation techniques must be applied to obtain flat volatilities of caps with maturities non available. To accurately capture the hump usually observed in the shape of the term structure of flat volatilities, we eventually proceed to interpolate flat volatilities by using cubic splines. We proceed to use linear interpolation or extrapolation only when the number of available flat volatilities corresponding to caps with the same strike is less than six. It is important to highlight that these interpolation/extrapolation techniques must produce uniquely determined values of

[18] Information provided by *Reuters* consists of flat volatility quotes of caps/floors; at a particular strike and for a concrete term to maturity, traders may contract the same instrument as a cap or a floor depending on their expectations.

unobservable flat volatilities with any term to maturity up to ten years and three months (the maturity date of the caplet with an exercise date in ten years time) in such a way that anyone should be able to exactly reproduce the same outcomes given a set of volatility quotations. See the Appendix for a detailed description of the interpolation/extrapolation procedure.

In any case, this method provides the flat cap volatilities needed to strip caplet prices according to the process described in the former section to construct the set of IRVIXs; for a given strike K we get flat volatilities of caps with maturities every three months, and then the stripping technique is applied to obtain the prices of caplets with any exercise date from one to ten years and strike K.

Once we have obtained caplet prices, we can proceed to the second step. For a given exercise date T_i, we select the two caplets with strikes closest to the outstanding forward interest rate $f(t,T_i,T_i+\tau)$ and then recover their implied volatility using the Black formula. Eventually, we apply the linear interpolation, as shown in Equation (12), to obtain the theoretical ATM implied volatility of a caplet with a particular exercise date T_i.

According to Fleming, Ostdiek, and Whaley (1995), this linear interpolation of implied volatilities from OTM and ITM options to create an ATM implied volatility implicitly assumes that the "volatility smile" is well approximated by a line. Thus, this approximation is considered reasonable when the interpolation is made for a small range of strikes. In this case, the two strikes closest (above and below) to the forward rate $f(t,T_i,T_i+\tau)$ can differ only by 100 or 50 basis points. Figures 1 through 3 show the surfaces of flat volatilities corresponding to caps (floors) with maturities in one, five and ten years, respectively, as a function of time and strikes K.[19] According to these pictures, smiles have had very different shapes and magnitudes.

[19] Flat volatility quotes for the strike 0.07 are not provided for any term to maturity from August 26, 2004 to November 11, 2004. Thus, surface graphs are constructed using the same flat volatility quotes for the strikes 0.06 and 0.07 up to November 12, 2004.

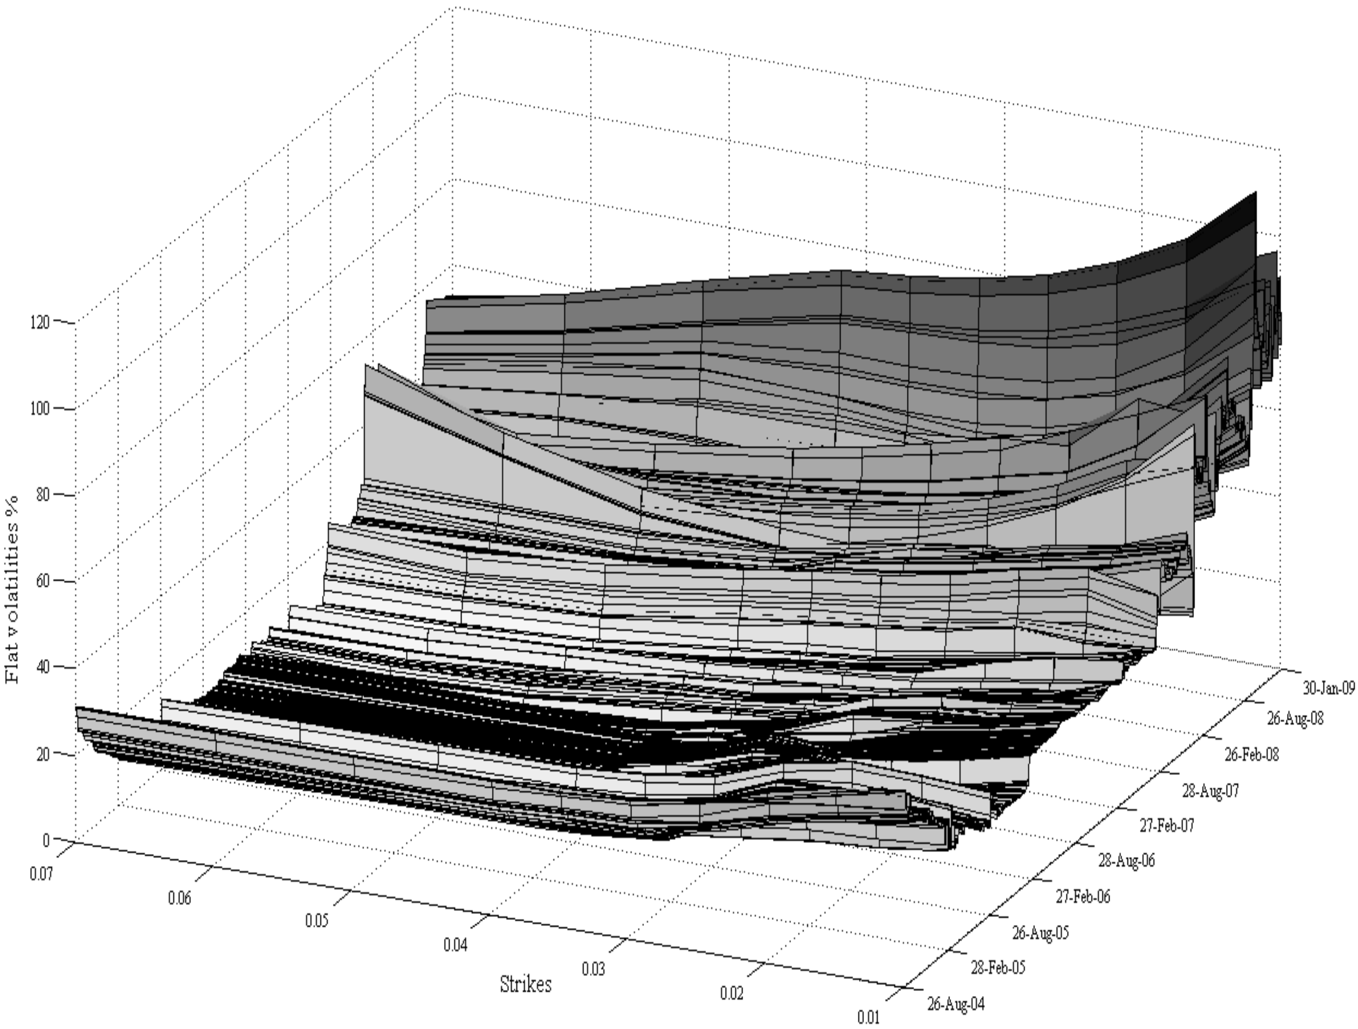

FIGURE 1. One-year flat cap (floor) volatilities as a function of strikes across the sample period.

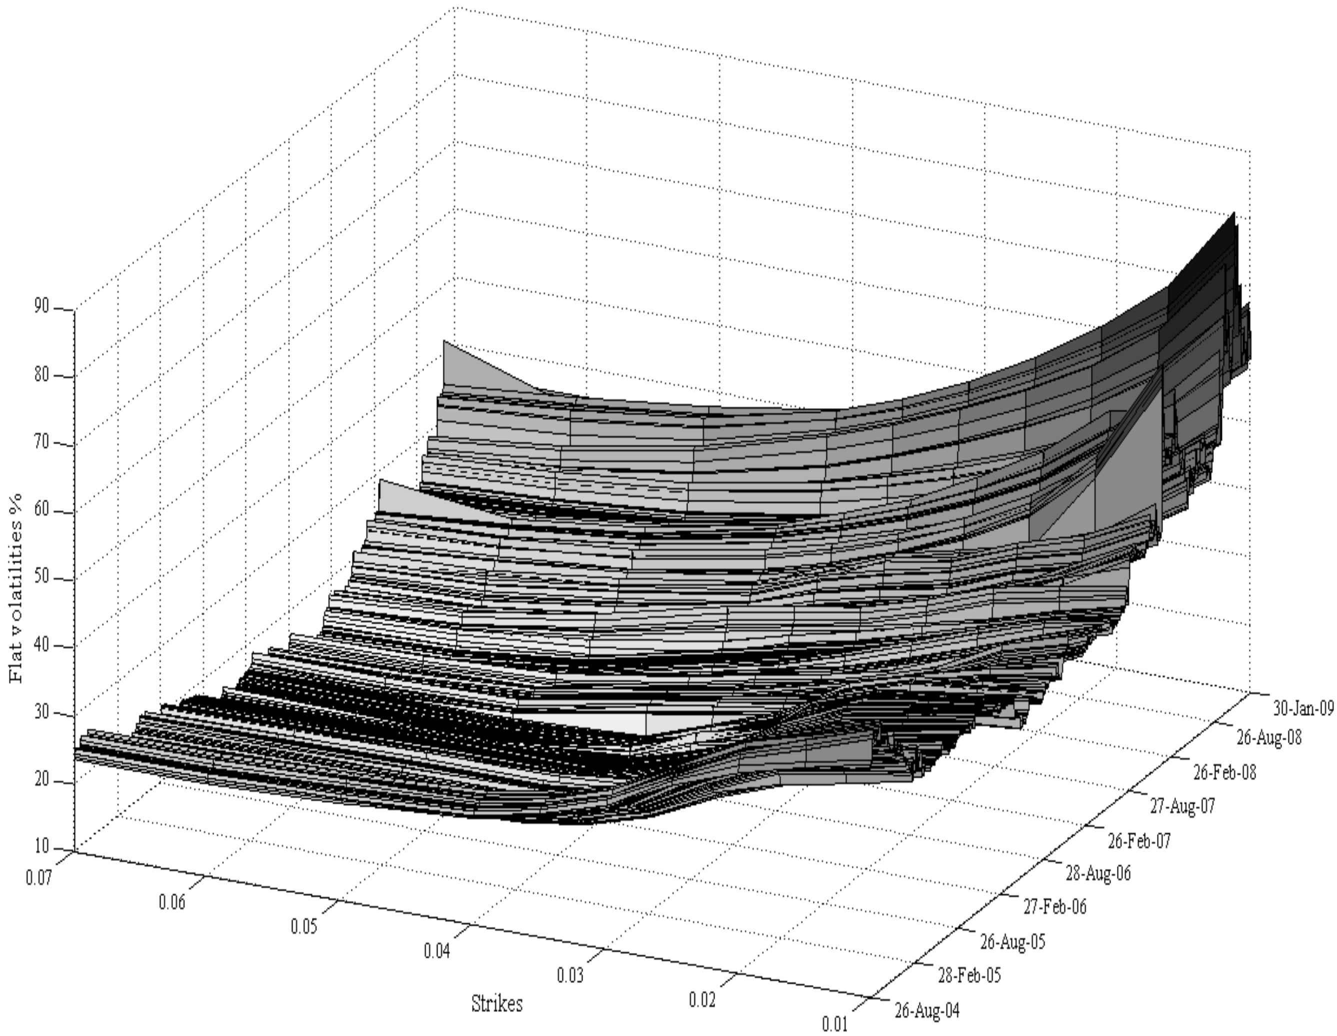

FIGURE 2. Five-year flat cap (floor) volatilities as a function of strikes across the sample period.

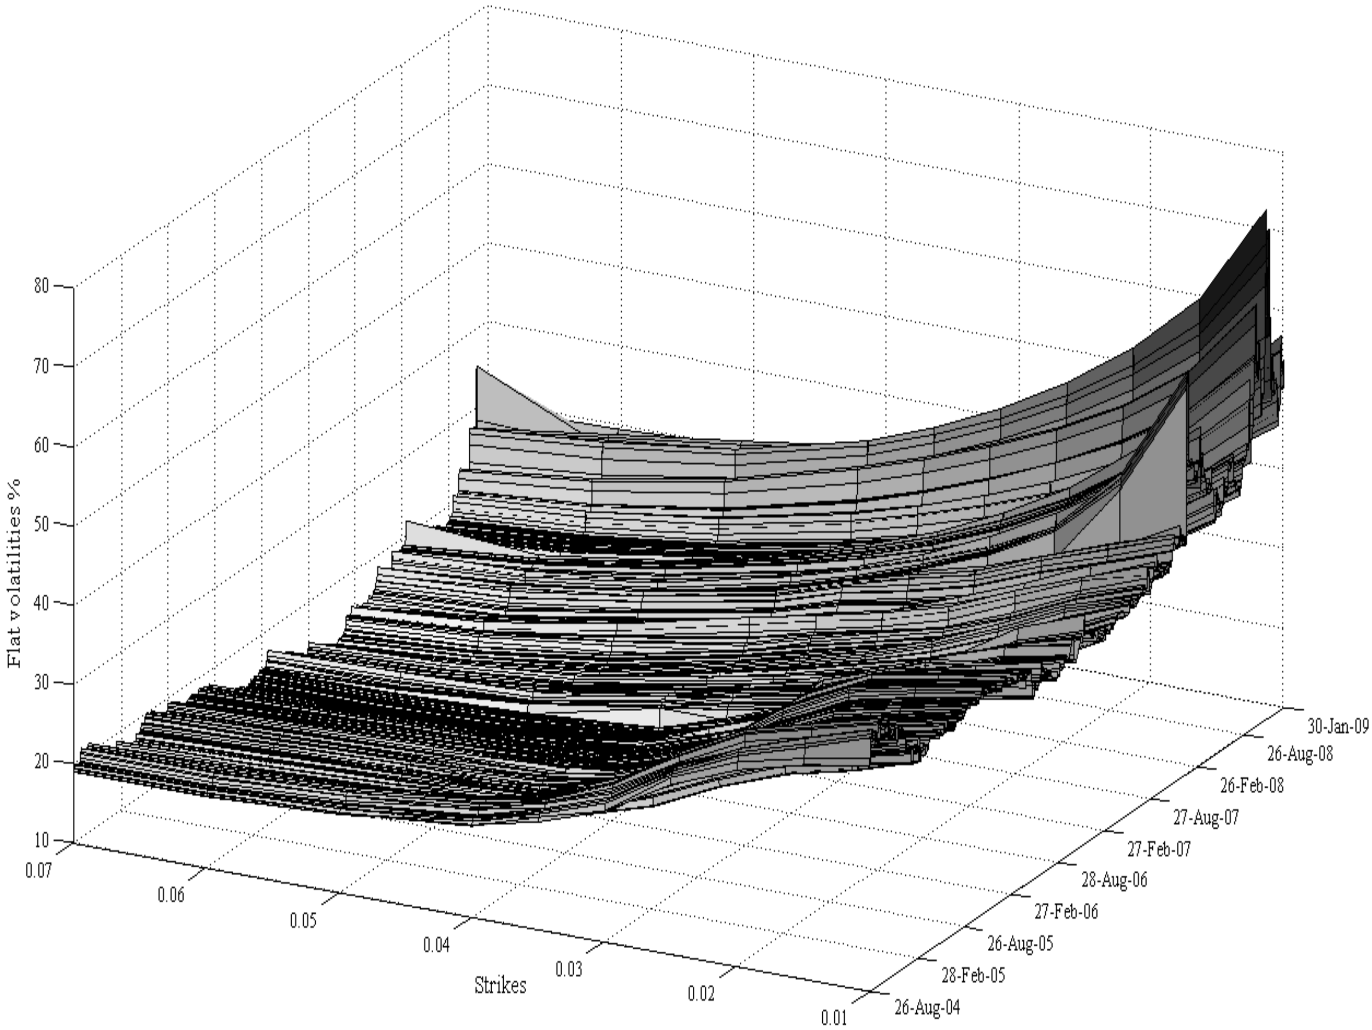

FIGURE 3. Ten-year flat cap (floor) volatilities as a function of strikes across the sample period.

Using the methodology proposed in this paper, we construct a daily set of interest rate volatility indices for forward rates with a three-month tenor and terms to maturity of one, two, three, four, five, seven and ten years. According to Duarte, Longstaff, and Yu (2007), these are the most-liquid cap maturities.

In particular, and according to Equation (7), IRVIX provides the average future volatility of a forward interest rate up to its maturity. For instance, the implied volatility index IRVIX(t,1Y) measures the market's assessment at any time *t* of the uncertainty regarding the evolution of the forward rate *f*(t,t+1Y,t+1Y+3M) during the next year, and IRVIX(t,10Y) would indicate the average volatility of the forward rate *f*(t,t+10Y,t+10Y+3M) during the next ten years.

This way, IRVIXs provide a set of forward-looking measures of expected volatility different to already existing volatility indices like those constructed by Merrill Lynch. The Merrill Lynch´s indices measure the implied volatility of long-term yields of Treasury securities over a relatively short horizon (the most popular is the one constructed from one-month options although indices with three and six-month horizons are also available). As far as IRVIXs are concerned, they represent market uncertainty about the expected three-month interest rate in the distant future (over horizons ranging from one to ten years ahead).

Another advantage of IRVIXs is that they provide a pure measure of the volatility of forward rates compared with flat volatilities that, according to LMM, are a mixture of the volatilities of forward rates with different tenors and terms to maturity. For instance, $\sigma_{1,flat}^{K}$ would be some sort of average of the future volatilities of the forward rates *f*(t,t+3M,t+6M), *f*(t,t+6M,t+9M) and *f*(t,t+9M,t+1Y) up to their respective maturities.

5. Empirical analysis

In this section we analyze the behavior and statistical properties of a set of interest rate volatility indices covering different forecast horizons (one year, two years, three years, four years, five years, seven years and ten years). The period covered by this study extends from August 26, 2004 to January 30, 2009. The fact that the sample period comprises the origin of the current financial crisis is especially relevant for visualizing the impact of the financial turmoil on the indices and thus on uncertainty with respect to the future behavior of forward rates. This outcome suggests the potential use of IRVIXs as leading indicators of the business cycle. We also analyze the relationship between IRVIXs and other volatility indices; in particular, with respect to MOVE, LBPX, and VIX. The empirical analysis additionally provides evidence of an overwhelming negative relationship between the indices and their underlying forward rates, similar to how VIX or VDAX are related to their underlying stock market indices. This result should be taken into account in the modeling of interest rate dynamics and in the design of hedging strategies.

Figure 4 plots the daily levels of IRVIXs with the two closest and two furthest forecast horizons: one year, two years, seven years and ten years.

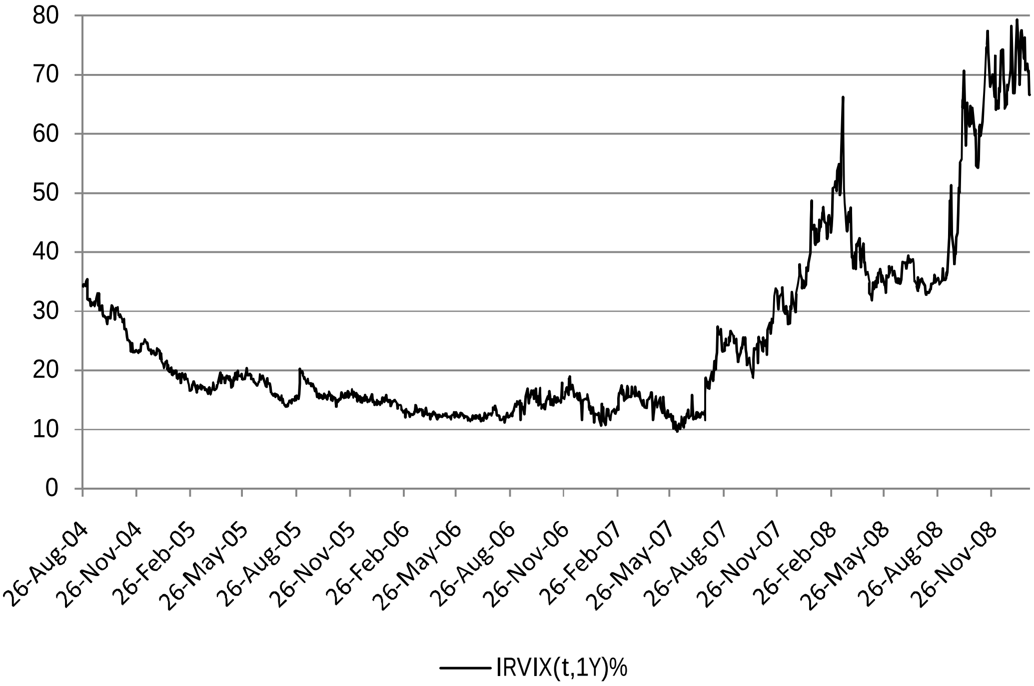
80
70
60
50
40
30
20
10
0
26-Aug-04
26-Nov-04
26-Feb-05
26-May-05
26-Aug-05
26-Nov-05
26-Feb-06
26-May-06
26-Aug-06
26-Nov-06
26-Feb-07
26-May-07
26-Aug-07
26-Nov-07
26-Feb-08
26-May-08
26-Aug-08
26-Nov-08
IRVIX(t,1Y)%

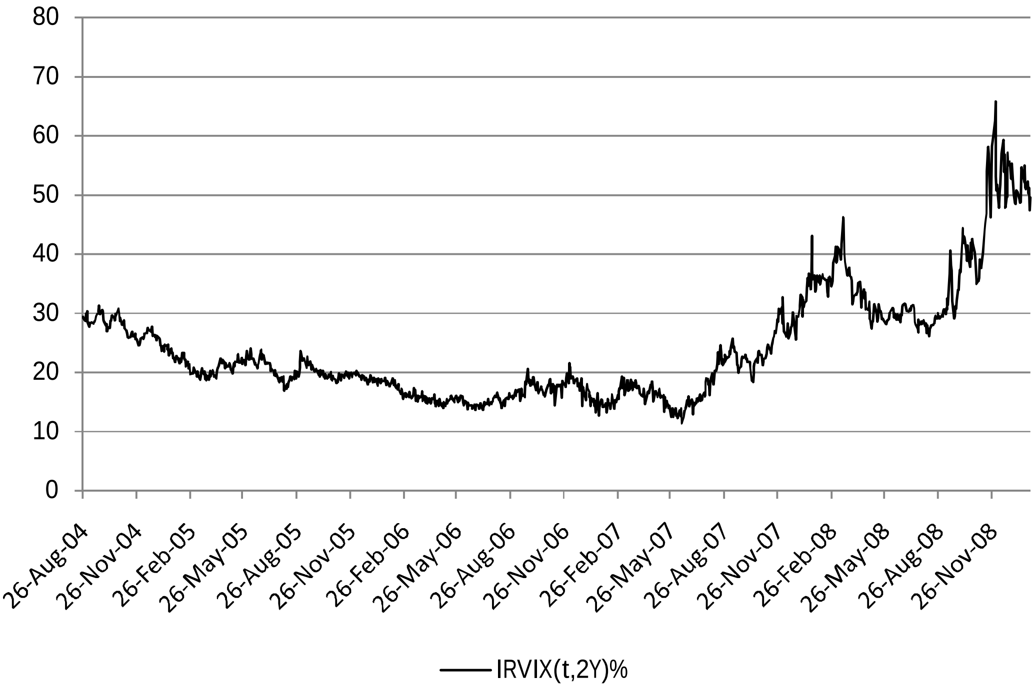
80
70
60
50
40
30
20
10
0
26-Aug-04
26-Nov-04
26-Feb-05
26-May-05
26-Aug-05
26-Nov-05
26-Feb-06
26-May-06
26-Aug-06
26-Nov-06
26-Feb-07
26-May-07
26-Aug-07
26-Nov-07
26-Feb-08
26-May-08
26-Aug-08
26-Nov-08
IRVIX(t,2Y)%

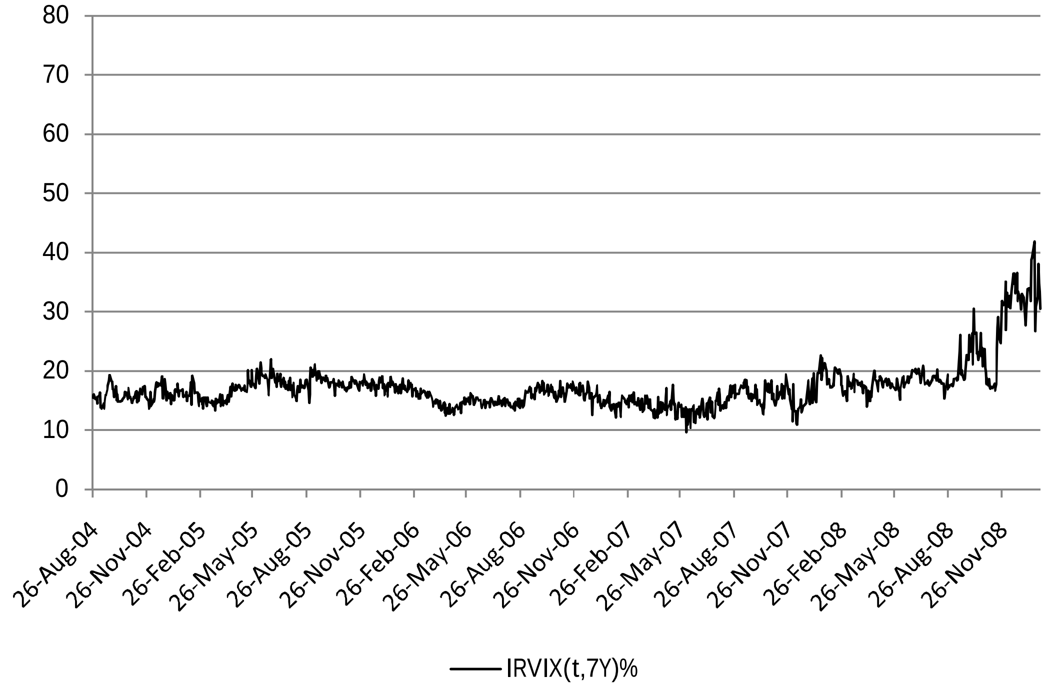

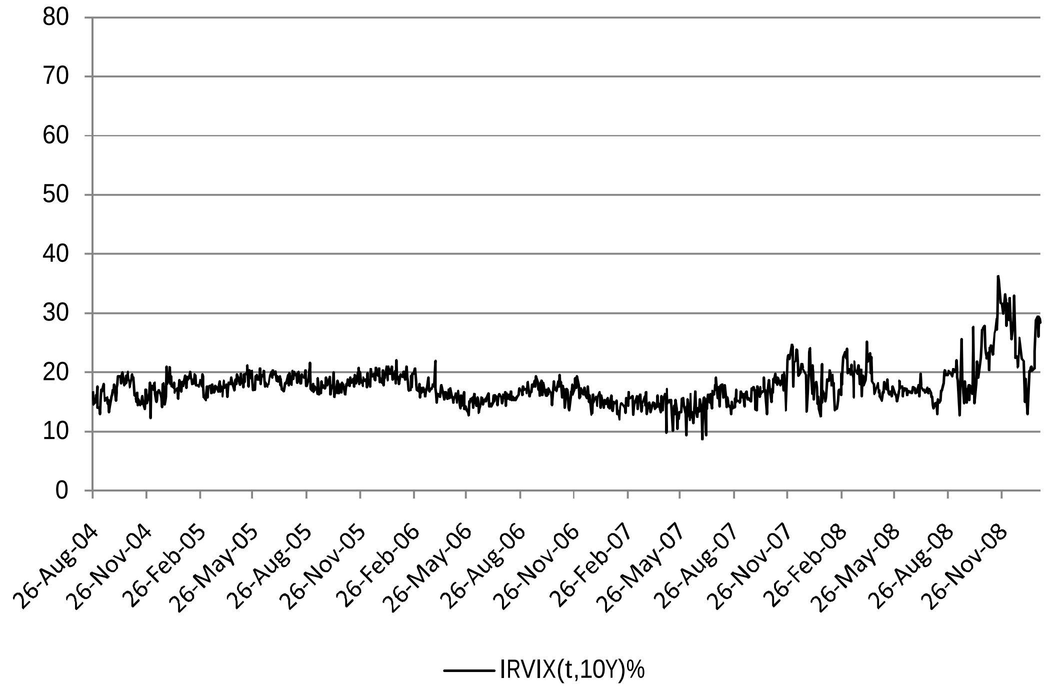

FIGURE 4. Daily levels of IRVIX(t,1Y), IRVIX(t,2Y), IRVIX(t,7Y) and IRVIX(t,10Y) during the period from August 26, 2004 to January 30, 2009.

IRVIX(t,1Y) and IRVIX(t,2Y) show a decreasing trend from the beginning of the sample period up to approximately mid-2006. During the following months, from September 2006 onwards, the series started to show frequent and larger spikes. Finally, after the beginning of the current financial crisis in the summer of 2007, the levels of the volatility indices remarkably increased and larger (up and down) spikes were observed. During the period associated with the crisis, we observe a dramatic increase in IRVIX(t,1Y) from 10% to nearly 80% in January 2009, while IRVIX(t,2Y) experienced a smoother upwards trend, reaching the maximum level of 66% in December 2008. It is interesting to observe that financial markets, particularly cap (floor) markets, appeared to be seized by uncertainty from September 2006.[20] This is a noteworthy finding that suggests the possible use of IRVIXs as potential leading indicators of the business cycle.

On the whole, graphs show that the average level of the indices decreases as the forecast horizon increases. Long-term expectations, as measured by IRVIX(t,7Y) and IRVIX(t,10Y), remain quite stable around a mean value along the entire sample. Graphs only exhibit an increase in the indices levels at the end of the period,[21] above all from mid-2008, when the levels of the volatility indices more than doubled with respect to the values prior to the Great Recession.

Although these volatility increases observed in long-term volatility indices are small compared with those experienced by short-term volatility indices, they may indicate a deeper and more permanent increase in the market uncertainty about interest rates. Recall that the IRVIX provides the average level of future volatility until the maturity of the

[20] Although the crisis burst out during mid-2007, some analysts started to detect and warn a turning point in home prices and in the demand for loans to buy new houses during the summer of 2006. See, for instance, *The Economist*´s articles “What´s that hissing sound?” and “Gimme shelter” (August 26, 2006).

[21] Intermediate maturities of the indices also tend to reflect this feature; the levels and evolution of the indices become more moderate as the forecast horizons move further away from today.

underlying forward rate (see Equation (7)). Therefore, a transitory increment in this volatility should have a limited impact on long-term IRVIXs. Only a general and permanent increment in the future volatility of forward rates would have a significant impact on these long-term, forward-looking measures of volatility.

Table 1 shows the summary statistics of the set of interest rate volatility indices. We suggest analyzing descriptive statistics for the full sample as well as before and after the beginning of the subprime crisis. Thus, the first subsample extends from August 26, 2004 to July 31, 2007, whereas the second subsample comprises the period from August 01, 2007 to January 30, 2009.[22] See panels B and C in Table 1 for summary statistics corresponding to each of these two periods.

The more steady evolution of the indices around a mean value when making longer-term predictions also explains the decreasing standard deviation (volatility of interest rate volatility) of the indices as the forecast horizon increases. Application of the Jarque-Bera test makes us reject the null hypothesis of a normal distribution (series show positive skewness and excess kurtosis).[23] Regarding to first-order autocorrelation coefficients, they show a progressive decline.

We also report the results from the Augmented Dickey Fuller (ADF) test performed on the logarithm of the volatility indices. Values of the test for its most general specification (i.e., with intercept and linear trend) indicate that the unit root hypothesis only

[22] August 2007 is usually referred to as the onset date of the subprime crisis, and a change in the values of the indices is also especially perceptible around this date.
[23] In line with studies on stock return volatility, such as Christensen and Prabhala (1998) and Andersen, Bollerslev, Diebold, and Ebens (2001), we find that the logarithms of the volatility indices (i.e. log implied volatilities) are closer to a normal distribution. Nevertheless, the null hypothesis of normality continues being rejected.

can be rejected for IRVIX(t,10Y).[24] Nevertheless, the time series plot might suggest a possible structural break in the series instead of a unit root. Further analysis on this issue is included next.

From the evidence reported by Perron (1989) on the data generating process of a set of macroeconomic series, we know that failure to account for possible breaking points may significantly reduce the power of traditional unit root tests that are biased towards the non-rejection of the null hypothesis. Thus, in addition to the ADF test, we apply the test developed by Zivot and Andrews (1992) that allows for a one-time structural break in the data. Unlike the test proposed by Perron (1989), the breaking point is endogenously determined as the outcome of the test instead of being considered as an exogenous occurrence.

Results of the Zivot-Andrews test applied to the logarithm of the set of indices are shown in Table 2. Following the analysis implemented in the work by Chaudhuri and Wu (2003), we report the estimation results of the specification that gives the most significant test statistic on α^i (i.e., the specification that reports the strongest evidence against the unit root hypothesis). Inference concerning the unit root hypothesis is based on the asymptotic critical values provided by Zivot and Andrews (1992).

We observe that the unit root hypothesis continues being accepted for the volatility indices with forecast horizons ranging from one to five years; however, IRVIX(t,7Y) might be characterized as a trend-stationary process at the 5% significance level after allowing for a possible structural break.

---

[24] The same results hold for the trendless version of the test, as well as by applying the modified Dickey-Fuller test (known as the DF-GLS test) proposed by Elliot, Rothenberg, and Stock (1996).

An additional advantage of this test is that it provides valuable information for analyzing whether a structural break in a certain variable is associated with the occurrence of a particular event. We find that the coefficients of the break dummy variables are statistically significant at the 5% level or better (based on critical values from the standard normal distribution) in all the cases. Moreover, as expected, structural breaks in short-term IRVIXs identified by the Zivot-Andrews test occur during the summer of 2007.

Concerning the two subsamples, summary statistics reveal the higher average value of all the volatility indices for the second subsample in comparison to the first subsample. In addition, as expected, much of the volatility of the indices computed for the entire sample is due to the greater variability observed for the second part of the sample.

In any case, we proceed to apply a set of tests to check if there are significant differences in the behavior of the indices between both subsamples in Table 3. First, we apply the Anova-F test for the equality of means, finding that the means of IRVIXs during the two subperiods are significantly different for all the forecast horizons. However, as we have found evidence of non-normality in the distribution of the indices, we cannot rely on this result.[25] Thus, we have proceeded to apply non-parametric tests to find evidence of differences in the distribution of the indices before and after the beginning of the crisis. In particular, we apply Wilcoxon/Mann-Withney test for the equality of medians and Brown-Forsythe test for the equality of variances. As expected, outcomes show evidence of significant differences in both the median and the variance between the first and second subsamples at the 1% significance level for all the indices.

[25] Jarque-Bera test applied to the log series in both subsamples also makes us reject the null hypothesis of normality.

In Table 4 we also show the statistical properties of the first log-difference transformation applied on the interest rate volatility indices with terms to maturity from one to five years to induce stationarity.

The excess kurtosis observed in the original series still remains, but skewness has been reduced. The first-order autocorrelation coefficients reveal a statistically significant negative autocorrelation that suggests the presence of mean reversion in the first log-differences of the indices. The same results (excess kurtosis and negative first-order autocorrelation) are also reported for different implied volatility indices introduced in stock markets in Dotsis, Psychoyios, and Skiadopoulos (2007). According to these authors, the evidence of non-normality may be attributed to the presence of jumps in implied volatility. Finally, the ADF test allows rejection of the null hypothesis of a unit root in the series.

For the series in first log-differences, statistically significant differences in the medians between the first and second subsamples are not proved for any forecast horizon (see Table 5). The null hypothesis of equality of variances is only rejected for IRVIX(t,4Y) and IRVIX(t,5Y).

The key difference between the proposed IRVIXs and the MOVE Index constructed by Merrill Lynch has been stressed along the article. IRVIXs measure uncertainty about the expected three-month interest rate over both short- and long-term (up to ten years) horizons, whereas MOVE measures one-month uncertainty about long-maturity yields. Next we show descriptive statistics for MOVE and analyze the contemporaneous relationship between IRVIXs and some other volatility indices.

Table 6 summarizes the statistical properties of MOVE. The mean, median and standard deviation computed for MOVE are not comparable to the values estimated for the set of IRVIXs since MOVE is constructed by applying a normalization process. Similar to IRVIXs, the series shows positive skewness and excess kurtosis. Finally, according to the ADF test performed on the logarithm of MOVE, the hypothesis of a unit root cannot be rejected. Nevertheless, after allowing for a structural break in the series, results from the Zivot-Andrews test suggest the rejection of the unit root hypothesis at the 5% significance level.[26]

In Figure 5 we show the contemporaneous relationship between weekly changes in the logarithm of IRVIX with the shortest term to maturity, IRVIX(t,1Y), and some other volatility indices. In particular, with respect to IRVIX(t,10Y), MOVE, LBPX, and VIX. On the one hand, a relationship between IRVIXs with the closest and furthest forecast horizons is barely perceptible. On the other hand, scatter plots suggest that IRVIX(t,1Y) is positively related to both other interest rate volatility indices (MOVE and LBPX) and VIX.

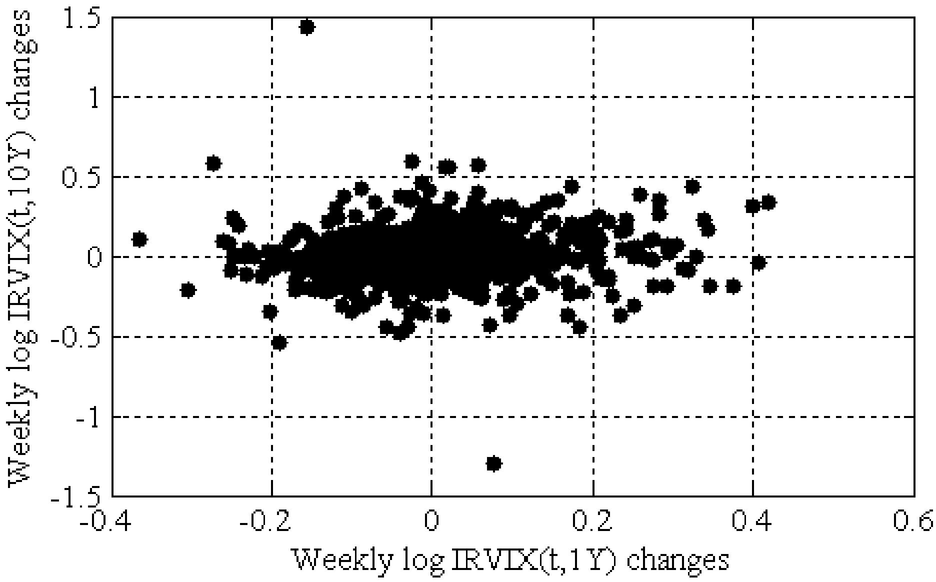

[26] The t-statistic for $\alpha = 0$ in Model C is -5.37. According to the test, the timing of the structural break is set at T_B = 06/13/2007.

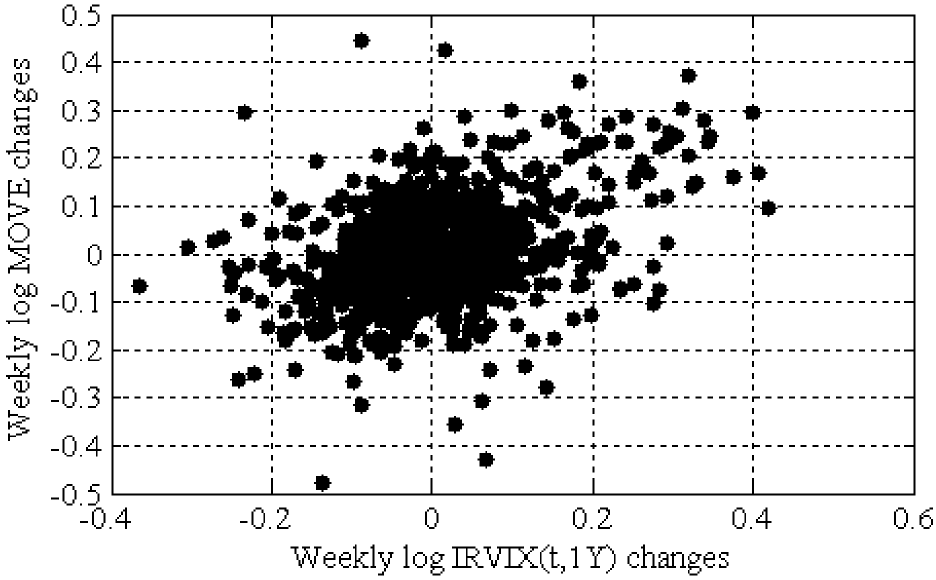
0.5
0.4
0.3
0.2
0.1
0
-0.1
-0.2
-0.3
-0.4
-0.5
Weekly log MOVE changes
-0.4
-0.2
0
0.2
0.4
0.6
Weekly log IRVIX(t,1Y) changes

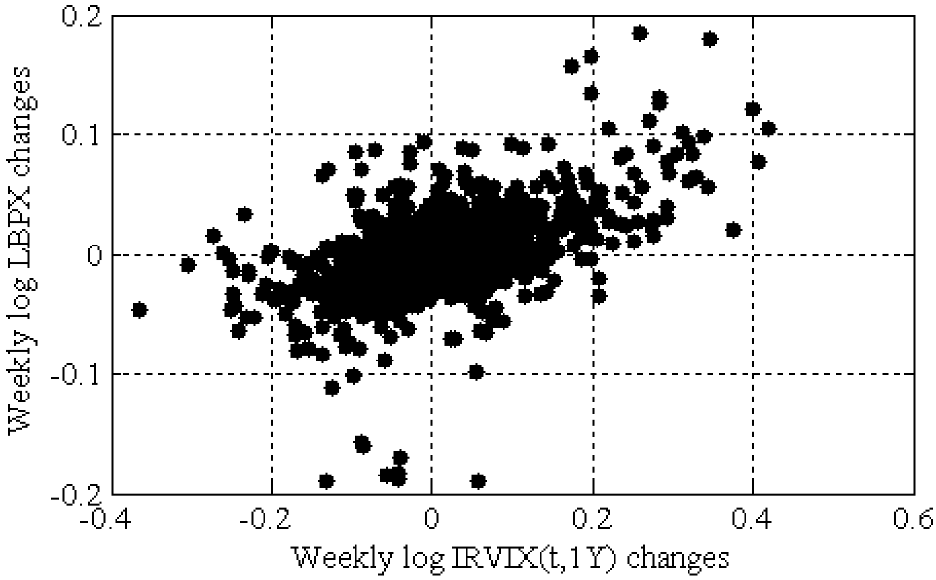
0.2
0.1
0
-0.1
-0.2
Weekly log LBPX changes
-0.4
-0.2
0
0.2
0.4
0.6
Weekly log IRVIX(t,1Y) changes

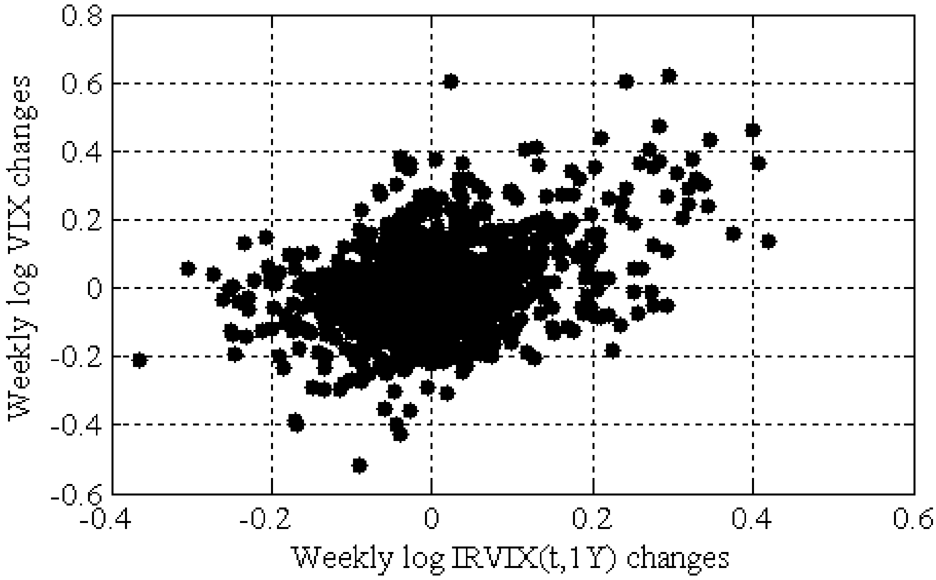

FIGURE 5. Scatter plots of weekly log-changes in IRVIX(t,1Y) against weekly log-changes in IRVIX(t,10Y), MOVE, LBPX, and VIX

In order to quantify the relationships between the different indices, we carry out a correlation analysis. We are also interested in analyzing whether the indices tend to move more closely during periods of financial instability. Thus, linear correlation coefficients are reported for the entire sample and also before and after the beginning of the subprime crisis (see Table 7).

Outcomes show that all the indices are positively correlated. Not surprisingly, we find that IRVIX(t,1Y) is clearly more strongly correlated to IRVIX(t,2Y) than to IRVIX(t,10Y). Correlation coefficients are 0.71 and 0.10, respectively. Let´s remember that IRVIX(t,1Y) captures interest rate volatility one year ahead, whereas IRVIX(t,2Y) and IRVIX(t,10Y) do the same over the next two and ten years, respectively. This fact also explains why IRVIX(t,1Y) is the most correlated index with MOVE, which captures market expectations of volatility over the next month. Moreover, the correlation coefficient between IRVIX(t,1Y) and MOVE is very similar to the one computed with respect to VIX, which measures stock market uncertainty over the same forecast horizon. We also observe that the strongest correlation between IRVIXs and the rest of indices is reported for short-term IRVIXs and LBPX (linear correlation is 0.51). Finally, Panels B and C reveal that short-term IRVIXs, MOVE, LBPX and VIX tend to move more closely in a context of higher financial uncertainty.

Next, we focus on the relationship between the interest rate volatility indices and their underlying forward rates. In the equity market literature, a negative contemporaneous relationship between implied volatility indices and the returns of the underlying stock market indices has been usually reported. Studies by Fleming, Ostdiek, and Whaley (1995), Whaley (2000), Skiadopoulos (2004), Giot (2005), and González and Novales (2009) document this feature for different countries and time periods.

Figure 6 shows weekly log-changes in interest rate volatility indices against weekly log-changes in forward interest rates at the two closest forecast horizons (one and two years) over the entire sample.

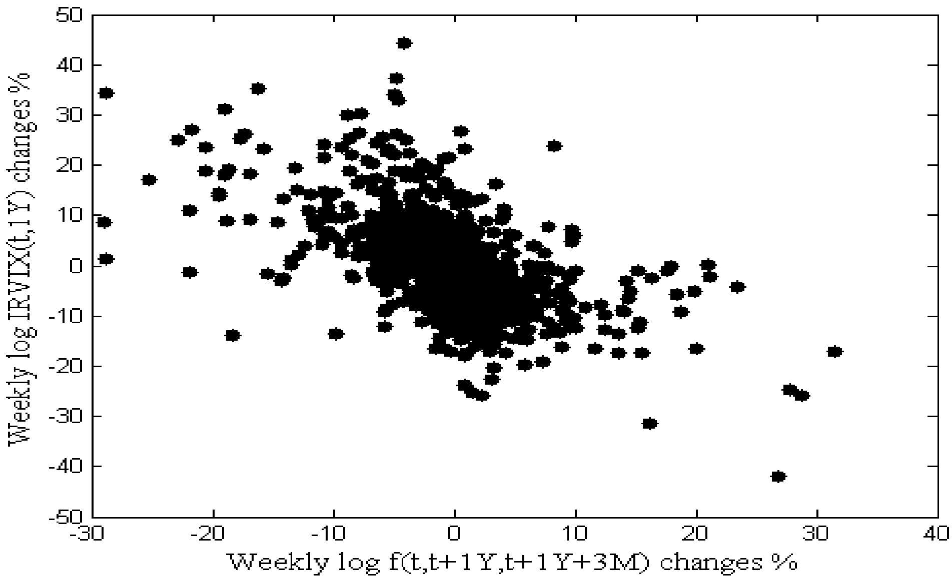

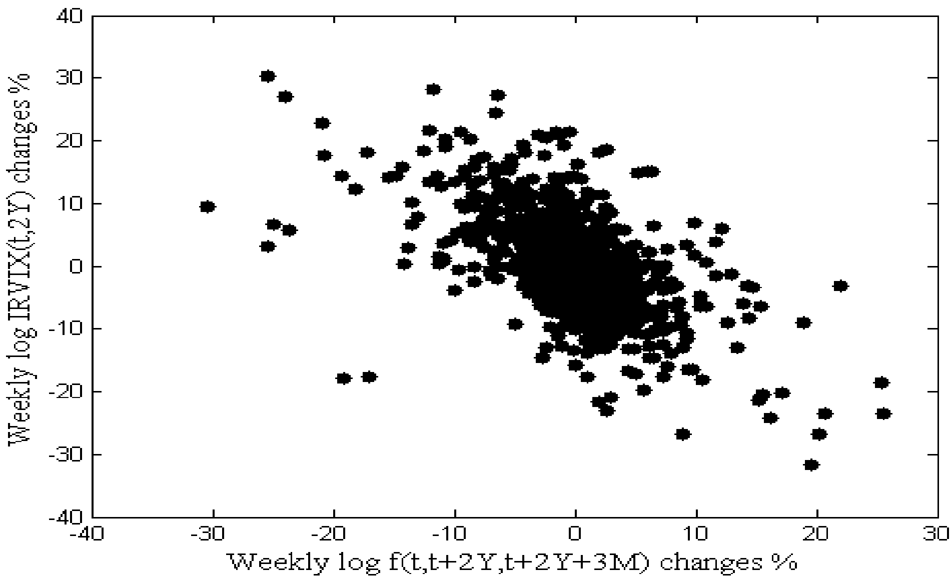

Figure 6. Scatter plots of weekly log-changes in interest rate volatility indices against weekly log-changes in forward interest rates at the one-year and two-year forecast horizons.

Graphs suggest an evident negative contemporaneous relationship between changes in short-term IRVIXs and changes in forward rates. The linear fitting releases correlation coefficients of -0.57 at the one-year forecast horizon and -0.58 at the two-year forecast horizon (see Table 8). Two important conclusions can be drawn from this analysis. First, we prove that the negative relationship between changes in implied volatility indices and the respective underlying also holds in the fixed-income market. Moreover, we show that this relationship holds when making predictions of volatility for longer forecast horizons than those used in stock markets. Second, we find that volatility indices and forward rates seem to be more negatively correlated in periods of financial turmoil even though both variables are also strongly correlated in periods of financial stability.

This tight relationship between changes in volatility and forward rates must be taken into account when designing hedging strategies. According to our results, an increase (decrease) in interest rates has a double impact on the prices of caps, floors, caplets or floorlets. On the one hand, as interest rates rise (fall), caps become deeper in the money (out of the money). On the other hand, the consequent decrease (increase) in volatility as a result of interest rates rising (falling) would pull down (up) the price of the option. Therefore, delta-hedging would not be adequate, as the cap or caplet prices would react less (more) than expected against interest rate changes.

The large shifts in volatility observed in the behavior of IRVIXs (volatility of volatility) in periods such as Winter 2007 (especially at the one-year forecast horizon) and Fall of 2008 reveal the growing need for derivatives to hedge volatility risk. Currently, VIX options and VIX futures are among the most actively traded contracts at the CBOE and the CBOE Futures Exchange (CFE), and we also foresee good expectations about the use of derivatives on an interest rate volatility index. Similar to VIX futures and options in the equity market, the introduction of futures and options on IRVIXs would offer a very useful tool to hedge against changes in interest rate volatility.

6. Summary and conclusions

In this paper, we propose a methodology for the construction of a set of indices that reflect market expectations about the volatility of three-month tenor forward rates up to ten years ahead.

IRVIXs have been extracted from one of the most liquid interest rate derivatives markets: the cap (floor) market. Apart from their high liquidity, these instruments allow measuring interest rate volatility over both short- and long-term horizons. This constitutes one of the main differences with respect to other volatility indices in the fixed-income market like MOVE, which measures the volatility of long-term yields of Treasury securities over a one-month horizon, or LBPX, which does not have a specific forecast horizon.

We construct a daily set of interest rate volatility indices with different forecast horizons (one year, two years, three years, four years, five years, seven years and ten years) over the period from August 26, 2004 to January 30, 2009.

We find that interest rate volatility indices have reacted very sharply to the uncertainty caused by the current financial crisis, suggesting their potential use as business cycle indicators. The impact of the current financial turmoil seemed to initially affect short-term market predictions of volatility, but as the crisis deepened, it also eventually had a very relevant impact on long-term IRVIXs, which suggests an increase in the investors´ long-term uncertainty regarding the future evolution of interest rates.

We also analyze the contemporaneous relationship between IRVIXs and other volatility indices for both the fixed-income and the equity markets: MOVE, LBPX, and VIX. We find that short-term IRVIXs are most highly correlated with these three indices. In addition, we observe that these indices tend to move more closely in a context of higher financial uncertainty.

Another remarkable finding is the strong negative correlation between changes in IRVIXs and in the underlying forward rates, particularly for short-term horizons (maximum level of linear correlation is found at the two-year forecast horizon with a correlation coefficient of -0.58). These outcomes are in line with the earlier empirical evidence in stock markets, but furthermore, we prove that the negative relationship between volatility indices and underlying financial variables also holds for longer forecast horizons.

Moreover, we find that the strong negative correlation holds for the two subperiods into which we divide the sample period, i.e., before and after the beginning of the present financial crisis. This fact should be taken into account when proceeding to hedge against interest rate risk using option-like contracts and particularly when designing delta-hedging strategies. An increase (decrease) in interest rates will cause changes in the value of caps and floors, but this change can be offset by the impact of consequent movements in volatility.

Concerning the statistical properties of the set of indices, the analysis of the first log-differences of the original indices shows excess kurtosis (leptokurtosis) and significant negative first-order autocorrelation. The same findings were obtained for most of the implied volatility indices in stock markets, where the non-normality is sometimes attributed to the presence of jumps, and the negative first-order autocorrelation supports the modeling of implied volatility indices as mean-reverting processes.

The distributional similarities with other volatility indices such as VIX or VDAX, and the dramatic changes experienced by interest rate volatility during the current financial crisis suggest that the introduction of futures and options on IRVIXs would offer a very useful tool to hedge against this important source of risk, similar to how futures and options on stock volatility indices hedge against changes in stock market volatility.

Appendix

The objective of the interpolation/extrapolation techniques used in the paper is twofold. On the one hand, we need a good fitting of the term structure of flat volatilities as a function of the term to maturity, particularly to capture the hump usually observed in the mid-term maturities. This requires that simple linear interpolation techniques be avoided, as they can lead to underestimation of flat volatilities around the peak of the hump. On the other hand, we need to use a method that determines these intermediate flat volatilities without ambiguity, i.e., a method that leads to uniquely determined values for intermediate flat volatilities.

Let us recall that the maximum number of flat volatility quotes available for a particular strike corresponds to caps (floors) with maturities of one to ten years plus 12, 15 and 20 years (13 maturities). When the number of available flat volatility quotes is equal to or greater than six, we use cubic splines. Otherwise, we propose simple linear interpolation/extrapolation.

In the first case, we distinguish two possibilities. If the number of observations is greater than nine, we use cubic splines with two intermediate knots. Otherwise, we use a single knot. Knots are positioned in such a way that the observations are uniformly distributed between knots. In particular, the position of the knots is set as follows.

Let N denote the number of available flat volatilities for a particular strike and t_1 and t_2 denote the positions of the knots. Then, we have the following:

If N= 13, t_1 = 4.5 and t_2= 8.5.

If N=12, t_1 is settled in the midpoint between the fourth and fifth observations and t_2 in the midpoint between the eighth and ninth observations.

If N=11, t_1 is positioned in the midpoint between the third and fourth observations and t_2 in the midpoint between the seventh and eighth observations.

If N=10, t_1 is settled at the midpoint between the third and fourth observations and t_2 in the midpoint between the sixth and seventh observations.

If N=9 or N=8, the unique knot is settled at the midpoint between the fourth and fifth observations.

If N=7 or N=6, the knot is positioned at the midpoint between the third and fourth observations.

When N is less than 6, we apply linear interpolation. For those caps with maturities out of the range of available maturities, we proceed to extrapolate. Let us denote the flat volatilities of the caps (floors) with the shortest and greatest terms to maturity by $\sigma^K_{Min,\,flat}$ and $\sigma^K_{Max,\,flat}$, respectively. Then, for caps (floors) maturing before the first available cap (floor), we assume that flat volatilities are equal to $\sigma^K_{Min,\,flat}$. For caps with maturity greater than the last available cap (floor), we assume that flat volatilities are equal to $\sigma^K_{Max,\,flat}$.

To give a hint of the completeness of the sample, Table 9 shows the proportion of flat volatilities available, corresponding to a single day and a given strike during the whole sample. As it is considered desirable in 89% of cases, the sample is complete (thirteen observations), and only in roughly 2% of cases do we have to apply linear interpolation or extrapolation techniques.

References

Andersen, T.G., Bollerslev, T., Diebold, F.X., Ebens, H., 2001. The distribution of realized stock return volatility. Journal of Financial Economics 61, 43-76.

Areal, N.M., 2008. FTSE-100 implied volatility index. Working paper (available at SSRN: http://ssrn.com/abstract=1102135).

Becker, R., Clements, A.E., White, S.I., 2007. Does implied volatility provide any information beyond that captured in model-based volatility forecasts? Journal of Banking and Finance 31, 2535-2549.

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

Bluhm, H.H.W., Yu, J., 2001. Forecasting volatility: Evidence from the German stock market. Working paper, University of Auckland.

Brigo, D., Mercurio, F., 2006. Interest Rate Models – Theory and Practice: With Smile, Inflation and Credit. Springer-Verlag, Berlin.

Carr, P., Madan, D., 1998. Towards a theory of volatility trading. In Jarrow, R.A. (Ed.), Volatility: New Estimation Techniques for Pricing Derivatives. Risk Books, London.

Chaudhuri, K., Wu, Y., 2003. Random walk versus breaking trend in stock prices: Evidence from emerging markets. Journal of Banking and Finance 27, 575-592.

Christensen, B.J, Prabhala, N.R., 1998. The relation between implied and realized volatility. Journal of Financial Economics 50, 125-150.

Corrado, C.J., Miller, T.W. Jr., 2005. The forecast quality of CBOE implied volatility indexes. The Journal of Future Markets 25, 339-373.

Díaz, A., Meneu, V., Navarro, E., 2009. International evidence on alternative models of the term structure of volatilities. The Journal of Future markets 29, 653-683.

Dotsis, G., Psychoyios, D., Skiadopoulos, G., 2007. An empirical comparison of continuous-time models of implied volatility indices. Journal of Banking and Finance 31, 3584-3603.

Duarte, J., Longstaff, F.A., Yu, F., 2007. Risk and Return in Fixed-income Arbitrage: Nickels in front of a streamroller? Review of Financial Studies 20, 769-811.

Elliott, G., Rothenberg, T.J., Stock, J.H., 1996. Efficient tests for an autoregressive unit root. Econometrica 64, 813-836.

Fleming, J., Ostdiek, B., Whaley, R.E., 1995. Predicting stock market volatility: A new measure. The Journal of Futures Markets 15, 265-302.

Giot, P., 2005. Relationships Between Implied Volatility Indexes and Stock Index Returns. Are implied volatility indexes leading indicators? The Journal of Portfolio Management 31, 92-100.

González, M.T., Novales, A., 2009. Are volatility indices in international stock markets forward looking? Matemática Aplicada 103, 339-352.

Hernández, L.G., 2005. Pricing of Game Options in a market with stochastic interest rates. Thesis (available at http://etd.gatech.edu/theses/available/etd-03152005-044640/).

Hull, J.C., 2009. Options, futures and other derivatives. Prentice Hall, Upper Saddle River.

Lyons, v.v.L., 2005. Volatility and its Measurements: The Design of a Volatility Index and the Execution of its Historical Time Series at the DEUTSCHE BÖRSE AG. Thesis (available at http://www.eurexchange.com/download/documents/ publications/Volatility_and_its_Measurements.pdf).

Moraux, F., Navatte, P., Villa, C., 1999. The Predictive Power of the French Market Volatility Index: A Multi Horizons Study. European Finance Review 2, 303-320.

Perron, P., 1989. The Great Crash, the Oil Price Shock, and the Unit Root Hypothesis. Econometrica 57, 1361-1401.

Poon, S-H., Granger, C.W.J., 2003. Forecasting Volatility in Financial Markets: A Review. Journal of Economic Literature XLI, 478-539.

Schwert, G. W., 1989. Tests for Unit Roots: A Monte Carlo Investigation. Journal of Business and Economic Statistics 7, 147-159.

Skiadopoulos, G., 2004. The Greek implied volatility index: Construction and properties. Applied Financial Economics 14, 1187-1196.

The Economist 380, N. 8492 (August 26, 2006).

Whaley, R. E., 2000. The Investor Fear Gauge. The Journal of Portfolio Management 26, 12-17.

Zivot, E., Andrews, D.W.K., 1992. Further evidence on the Great Crash, the Oil-Price Shock, and the Unit-Root Hypothesis. Journal of Business and Economic Statistics 10, 251-270.

TABLE 1.- Summary statistics of IRVIXs across the entire sample (Panel A) and for two subsamples: from August 26, 2004 to July 31, 2007 (Panel B) and from August 01, 2007 to January 30, 2009 (Panel C)

	IRVIX (t,1Y)	IRVIX (t,2Y)	IRVIX (t,3Y)	IRVIX (t,4Y)	IRVIX (t,5Y)	IRVIX (t,7Y)	IRVIX (t,10Y)
Panel A: August 26, 2004 to January 30, 2009							
Observations	1125	1125	1125	1125	1125	1125	1125
Mean	0.24	0.23	0.24	0.22	0.19	0.17	0.17
Median	0.17	0.20	0.22	0.21	0.19	0.16	0.17
Std. Deviation	0.15	0.09	0.06	0.05	0.04	0.03	0.03
Skewness	1.63	1.54	1.04	1.33	2.48	2.82	1.74
Kurtosis	5.16	5.54	4.22	5.95	11.39	13.43	8.93
Jarque-Bera	722.10	752.87	275.01	745.70	4468.85	6596.56	2223.26
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
ρ_1	0.99^*	0.98^*	0.97^*	0.96^*	0.96^*	0.92^*	0.83^*
ADF	-1.90	-1.61	-1.63	-1.36	-1.67	-2.20	-3.67^*
Panel B: August 26, 2004 to July 31, 2007							
Observations	756	756	756	756	756	756	756
Mean	0.16	0.19	0.21	0.20	0.17	0.15	0.16
Median	0.15	0.18	0.21	0.20	0.18	0.15	0.17
Std. Deviation	0.04	0.04	0.04	0.03	0.02	0.01	0.02
Skewness	1.66	0.91	0.62	0.40	-0.11	0.03	-0.38
Kurtosis	5.59	3.23	2.89	2.50	2.09	2.88	3.27
Jarque-Bera	560.84	106.51	50.12	28.82	27.41	0.57	20.83
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.75)	(0.00)
Panel C: August 01, 2007 to January 30, 2009							
Observations	369	369	369	369	369	369	369
Mean	0.41	0.33	0.29	0.26	0.22	0.19	0.19
Median	0.36	0.30	0.28	0.25	0.21	0.17	0.17
Std. Deviation	0.15	0.09	0.06	0.05	0.05	0.05	0.04
Skewness	0.77	0.95	0.68	1.17	1.81	1.76	1.34
Kurtosis	2.63	3.50	3.20	4.48	5.43	5.63	4.82
Jarque-Bera	39.33	59.41	29.69	118.41	293.78	299.50	162.84
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)

Notes:

[a] *p*-values of the Jarque-Bera test are inside parenthesis.

[b] ρ_1 denotes the first-order autocorrelation coefficient.

[c] The ADF test is performed on the logarithm of the indices. The optimal lag length is determined according to the Akaike information criterion, with the maximum lag set to 21 based on Schwert (1989).

[d] One asterisk denotes statistical significance at the 5% significance level. Two asterisks denote statistical significance at the 1% significance level.

TABLE 2.-Zivot-Andrews structural break test performed on the logarithm of IRVIXs

	Model	T_B	k	θ	γ	α
log IRVIX 1Y	C	07/26/2007	4	0.0427 (4.666)	0.0001 (2.597)	-0.0431 (-4.696)
log IRVIX 2Y	C	07/05/2007	4	0.0263 (3.699)	0.0001 (2.804)	-0.0448 (-4.204)
log IRVIX 3Y	C	07/20/2007	11	0.0328 (3.691)	0.0001 (2.371)	-0.0640 (-3.928)
log IRVIX 4Y	C	07/20/2007	8	0.0265 (3.511)	0.0001 (2.933)	-0.0701 (-4.141)
log IRVIX 5Y	B	05/29/2007	4		0.0001 (4.069)	-0.0664 (-4.330)
log IRVIX 7Y	B	05/21/2008	9		0.0004 (4.463)	-0.098* (-4.813)
log IRVIX 10Y	B	06/06/2007	7		0.0002 (4.126)	-0.1363** (-5.471)

Notes:

[a] IRVIX 1Y indicates the index maturing in one year.

[b] Optimal lag length, *k*, is selected according to the Akaike information criterion, with the maximum lag set to 21 based on Schwert (1989).

[c] Numbers inside parenthesis are *t*-ratios.

[d] The 10%, 5% and 1% asymptotic critical values for the *t*-statistic for $\alpha^i = 0$ (i = A, B or C), obtained from Zivot and Andrews (1992), are respectively as follows. Model A: -4.58, -4.80, and -5.34; Model B: -4.11, -4.42, and -4.93; Model C: -4.82, -5.08, and -5.57.

[e] One asterisk and two asterisks denote statistical significance at the 5% and 1% levels, respectively, based on the asymptotic critical values.

The test is described as follows. Let T_B be a potential breaking point in the series y_t. The test may be performed based on three alternative specifications of the process followed by the series:

Model (A): $$\Delta y_t = \mu^A + \theta^A DU_t + \beta^A t + \alpha^A y_{t-1} + \sum_{j=1}^{k} \phi_j^A \Delta y_{t-j} + \varepsilon_t,$$

Model (B): $$\Delta y_t = \mu^B + \beta^B t + \gamma^B DT_t + \alpha^B y_{t-1} + \sum_{j=1}^{k} \phi_j^B \Delta y_{t-j} + \varepsilon_t,$$

Model (C): $$\Delta y_t = \mu^C + \theta^C DU_t + \beta^C t + \gamma^C DT_t + \alpha^C y_{t-1} + \sum_{j=1}^{k} \phi_j^C \Delta y_{t-j} + \varepsilon_t,$$

where t is the time trend variable; the two dummy variables are defined as follows: $DU_t = 1$ if $t > T_B$, zero otherwise, and $DT_t = t - T_B$ if $t > T_B$, zero otherwise; and ε_t are white noise processes. This way, Model (A) allows for a one-time change in the intercept; Model (B) permits a break in the slope of the trend function; and Model (C) admits both changes. Like in the ADF test, the *k* extra lags of the dependent variable are added to correct for serial correlation in the error term.

The null hypothesis in all the three models is $\alpha^i = 0$ (i = A, B or C), i.e., the series follows a unit root process; while the alternative hypothesis implies that the series is a trend-stationary process with a one-time break in the trend function occurring at an unknown point in time. The goal of the procedure is to search for the break point that gives the most weight to the trend-stationary alternative. In a sample with T observations, for each specification (A, B or C), to determine the break and to compute the test statistic for a unit root, an ordinary least squares regression is run with a potential break point at T_B (for T_B ranging from 2 to $T-2$).

TABLE 3.- Tests of equality of means, medians and variances between the first and second subsamples of IRVIXs

	IRVIX (t,1Y)	IRVIX (t,2Y)	IRVIX (t,3Y)	IRVIX (t,4Y)	IRVIX (t,5Y)	IRVIX (t,7Y)	IRVIX (t,10Y)
Anova F- test	1617.44**	1159.20**	598.00**	517.09**	398.80**	261.47**	154.28**
Wilcoxon/Mann-Withney test	25.92**	24.07**	19.04**	18.54**	17.86**	13.92**	8.38**
Brown-Forsythe test	328.77**	181.40**	67.57**	47.10**	75.17**	114.01**	122.07**

Notes:

[a] One asterisk denotes statistical significance at the 5% significance level. Two asterisks denote statistical significance at the 1% significance level.

TABLE 4.- Summary statistics of first-log differences of IRVIXs across the entire sample (Panel A) and for two subsamples: from August 26, 2004 to July 31, 2007 (Panel B) and from August 01, 2007 to January 30, 2009 (Panel C)

	IRVIX (t,1Y)	IRVIX (t,2Y)	IRVIX (t,3Y)	IRVIX (t,4Y)	IRVIX (t,5Y)
Panel A: August 26, 2004 to January 30, 2009					
Observations	1124	1124	1124	1124	1124
Mean	0.0005	0.0004	0.0001	0.0002	0.0005
Median	0.0002	0.0013	0.0001	0.0001	-0.0002
Std. Deviation	0.05	0.05	0.06	0.05	0.05
Skewness	0.33	-0.12	0.15	-0.02	-0.16
Kurtosis	6.31	5.69	5.82	6.00	6.00
Jarque-Bera	535.67 (0.00)	343.31 (0.00)	376.89 (0.00)	423.12 (0.00)	428.62 (0.00)
ρ_1	-0.28*	-0.37*	-0.36*	-0.37*	-0.44*
ADF	-20.54**	-20.43**	-12.98**	-15.68**	-22.21**
Panel B: August 26, 2004 to July 31, 2007					
Observations	755	755	755	755	755
Mean	-0.0009	-0.0006	-0.0005	-0.0004	-0.0001
Median	-0.0011	0.0003	0.0002	-0.0001	-0.0005
Std. Deviation	0.05	0.05	0.06	0.05	0.05
Skewness	0.45	-0.17	0.12	-0.007	-0.19
Kurtosis	7.23	5.97	6.71	7.22	6.61
Jarque-Bera	591.17 (0.00)	283.03 (0.00)	437.15 (0.00)	561.53 (0.00)	415.06 (0.00)
Panel C: August 01, 2007 to January 30, 2009					
Observations	368	368	368	368	368
Mean	0.0035	0.0027	0.0015	0.0015	0.0021
Median	0.0034	0.0042	-0.0009	0.0016	0.0000
Std. Deviation	0.05	0.05	0.06	0.06	0.06
Skewness	0.06	-0.04	0.20	-0.05	-0.15
Kurtosis	4.18	5.17	3.69	4.16	5.05
Jarque-Bera	21.71 (0.00)	72.44 (0.00)	10.11 (0.00)	20.98 (0.00)	66.50 (0.00)

Notes:

[a] *p*-values of the Jarque-Bera test are inside parenthesis.

[b] ρ_1 denotes the first-order autocorrelation coefficient.

[c] The optimal lag length in the ADF test is determined according to the Akaike information criterion, with the maximum lag set to 21 based on Schwert (1989).

[d] One asterisk denotes statistical significance at the 5% significance level. Two asterisks denote statistical significance at the 1% significance level.

TABLE 5.- Tests of equality of means, medians and variances between the first and second subsamples of IRVIXs with terms to maturity from one to five years in first log-differences

	IRVIX (t,1Y)	IRVIX (t,2Y)	IRVIX (t,3Y)	IRVIX (t,4Y)	IRVIX (t,5Y)
Anova F- test	1.53	1.00	0.26	0.26	0.38
Wilcoxon/Mann-Withney test	1.44	1.10	0.05	0.74	0.54
Brown-Forsythe test	0.06	0.59	0.68	5.81*	4.70*

Notes:

[a] One asterisk denotes statistical significance at the 5% significance level. Two asterisks denote statistical significance at the 1% significance level.

TABLE 6.- Summary statistics of the MOVE Index constructed by Merrill Lynch across the entire sample

	MOVE
Observations	1109
Mean	0.009
Median	0.008
Std. Deviation	0.004
Skewness	1.46
Kurtosis	4.67
Jarque-Bera	527.88 (0.00)
ρ_1	0.99*
ADF	-2.50

Notes:

[a] *p*-value of the Jarque-Bera test is inside parenthesis.

[b] ρ_1 denotes the first-order autocorrelation coefficient.

[c] The ADF test is performed on the logarithm of MOVE. The optimal lag length is determined according to the Akaike information criterion, with the maximum lag set to 21 based on Schwert (1989).

[d] One asterisk denotes statistical significance at the 5% significance level. Two asterisks denote statistical significance at the 1% significance level.

TABLE 7. Correlation coefficients between weekly log-changes in IRVIX(t,1Y), IRVIX(t,2Y), IRVIX(t,10Y), MOVE, LBPX, and VIX across the entire sample (Panel A) and for two subsamples: from August 26, 2004 to July 31, 2007 (Panel B) and from August 01, 2007 to January 30, 2009 (Panel C)

Panel A: August 26, 2004 to January 30, 2009						
	IRVIX(t,1Y)	IRVIX(t,2Y)	IRVIX(t,10Y)	MOVE	LBPX	VIX
IRVIX(t,1Y)	1	0.71	0.10	0.37	0.51	0.40
IRVIX(t,2Y)	0.71	1	0.13	0.36	0.51	0.34
IRVIX(t,10Y)	0.10	0.13	1	0.09	0.19	0.16
MOVE	0.37	0.36	0.09	1	0.55	0.38
LBPX	0.51	0.51	0.19	0.55	1	0.42
VIX	0.40	0.34	0.16	0.38	0.42	1
Panel B: August 26, 2004 to July 31, 2007						
	IRVIX(t,1Y)	IRVIX(t,2Y)	IRVIX(t,10Y)	MOVE	LBPX	VIX
IRVIX(t,1Y)	1	0.61	0.12	0.29	0.46	0.27
IRVIX(t,2Y)	0.61	1	0.06	0.27	0.49	0.23
IRVIX(t,10Y)	0.12	0.06	1	0.12	0.16	0.14
MOVE	0.29	0.27	0.12	1	0.58	0.39
LBPX	0.46	0.49	0.16	0.58	1	0.36
VIX	0.27	0.23	0.14	0.39	0.36	1
Panel C: August 01, 2007 to January 30, 2009						
	IRVIX(t,1Y)	IRVIX(t,2Y)	IRVIX(t,10Y)	MOVE	LBPX	VIX
IRVIX(t,1Y)	1	0.79	0.08	0.43	0.57	0.54
IRVIX(t,2Y)	0.79	1	0.18	0.44	0.53	0.45
IRVIX(t,10Y)	0.08	0.18	1	0.06	0.20	0.18
MOVE	0.43	0.44	0.06	1	0.57	0.37
LBPX	0.57	0.53	0.20	0.57	1	0.50
VIX	0.54	0.45	0.18	0.37	0.50	1

TABLE 8. Correlation coefficients between weekly log-changes in IRVIXs and weekly log-changes in underlying forward interest rates from one-year to ten-year forecast horizons

	1Y	2Y	3Y	4Y	5Y	7Y	10Y
Entire sample	-0.57	-0.58	-0.49	-0.51	-0.42	-0.33	-0.20
First subsample	-0.55	-0.53	-0.34	-0.35	-0.35	-0.25	-0.15
Second subsample	-0.63	-0.63	-0.62	-0.63	-0.49	-0.40	-0.23

Notes:

[a] 1Y stands for the time to maturity of both IRVIXs and forward rates.

TABLE 9. Frequency of available flat volatility quotes along the sample

Number of available flat volatilities	Frequency
N = 13	0.8970
N = 12	0.0332
N = 11	0.0110
N = 10	0.0100
N = 9	0.0088
N = 8	0.0059
N = 7	0.086
N = 6	0.056
N = 5	0.0049
N = 4	0.0036
N = 3	0.0041
N = 2	0.0034
N = 1	0.0032