Journal of Banking & Finance 33 (2009) 1156–1165

Contents lists available at ScienceDirect

Journal of Banking & Finance

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

The dynamics of the volatility skew: A Kalman filter approach

Mascia Bedendo [a,*], Stewart D. Hodges [b]

[a] Universita' Bocconi and Carefin, Via Roentgen 1, 20136 Milan, Italy
[b] Cass Business School, City University, 106 Bunhill Row, EC1Y 8TZ London, UK

ARTICLE INFO

Article history:
Received 22 January 2008
Accepted 21 December 2008
Available online 8 January 2009

JEL classification:
C53
G13

Keywords:
Implied volatility
Kalman filter
Density forecasting
Value-at-Risk

ABSTRACT

Much attention has been devoted to understanding and modeling the dynamics of implied volatility curves and surfaces. This is crucial for both trading, pricing and risk management of option positions. We suggest a simple, yet flexible, model, based on a discrete and linear Kalman filter updating of the volatility skew. From a risk management perspective, we assess whether this model is capable of producing good density forecasts of daily returns on a number of option portfolios. We also compare our model to the sticky-delta and the vega–gamma alternatives. We find that it clearly outperforms both alternatives, given its ability to easily account for movements of different nature in the volatility curve.

1. Introduction

It is well known that the volatilities implied from observed option prices are not constant across strikes and time to maturity, as the Black–Scholes model would predict. Instead, they exhibit a smile/skew pattern across strikes for a given time to maturity, which extends to an entire volatility surface when different expiries are examined. These implied volatility curves and surfaces also change through time, raising the need for an accurate modeling of their dynamics (Becker et al., 2007; Deuskar et al., 2008; Konstantinidi et al., 2008).

From the standpoint of a vanilla option market maker, producing reliable forecasts for the evolution of the volatility curve is essential for two reasons. First, it provides an up-to-date indication of where the market stands, which is essential for trading. Second, it allows an efficient and dynamic risk management of the option portfolios. In practice we do not observe simultaneous updates of the entire volatility skew, on the market, as new quotes for individual trades become available at different points in time. New information is incorporated into the new trades and it becomes crucial for market practitioners to be able to learn about the volatility skew from the quotes of those new trades available, which are typically very few. A natural tool that allows to quickly update the estimates of the volatility curve following any new information, albeit partial, released on the market, is the Kalman filter technique.

In this paper, we suggest a simple model for the dynamic evolution of the volatility skew for a given contract based on a linear Kalman filter approach. In our specification, the observable variables are implied volatilities and the unobservable state variables are the stochastic coefficients driving the skew. The skew is modeled as a cubic polynomial in the moneyness measure, whose dynamics follows a gaussian Ornstein–Uhlenbeck process correlated with the stock price process. By using a state-space model, we can exploit the information on both: (1) the temporal evolution of implied volatilities; (2) the cross-section of implied volatilities with respect to moneyness.[1] The specification we suggest is intuitively simple and easy to implement. Most importantly, it accounts for stochastic non-parallel changes in the volatility skew, which makes it appealing not only for trading, but also for risk management purposes.

Even though the Kalman filter seems a natural tool for financial problems of this kind, our study constitutes, to our knowledge, the first application of this powerful and robust econometric technique to the updating of the volatility skew. While several works have proposed using the Kalman filter to estimate and model the dynamics of the term structure of interest rates, very few applications to risk management problems can be found in the literature.

* Corresponding author. Tel.: +39 02 5836 5973; fax: +39 02 5836 5920.
E-mail addresses: mascia.bedendo@unibocconi.it (M. Bedendo), stewart.hodges.1@city.ac.uk (S.D. Hodges).

[1] In order to keep the analysis simple, we do not impose constraints to ensure an arbitrage-free specification.

0378-4266/$ - see front matter
doi:10.1016/j.jbankfin.2008.12.014

Brealey et al. (1975) were the first to employ this technique for the estimation of equity betas. Recently, Sommacampagna (2002) has proposed a Kalman filter method for VaR estimation purposes.

Traditionally, the dynamic properties of implied volatility have been studied mainly by focussing on either the term structure of the at-the-money (ATM) implied volatility, or the volatility skew for a given maturity. Investigations of the dynamics followed by the entire volatility surface have begun to appear recently. Empirical studies on the volatility dynamics normally consist in identifying the number and shapes of the shocks in the implied volatility via principal component analysis (PCA) (see, for example, Skiadopoulos et al., 1999; Alexander, 2001; Cont and Da Fonseca, 2002). Fengler (2006) describes both these and other methods.

More recent contributions involve the specification of a deterministic or stochastic model for the implied volatility smile or surface, which fully describes its evolution through time. The implied volatility models introduced by Derman (1999) assume that either the per-delta or the per-strike implied volatility surface has a deterministic evolution. Rosenberg (2000) proposed a stochastic process for the ATM implied volatility, while keeping the shape of the curve fixed. A stochastic evolution of the entire volatility surface characterizes the stochastic implied volatility models by Schönbucher (1999) and Ledoit and Santa-Clara (1998).

General equilibrium models which explain the stochastic volatility and leverage behind the volatility surface with the investors' uncertainty about the economy have also been proposed, amongst others, by David and Veronesi (2002), Guidolin and Timmermann (2003), and Hibbert et al. (2008). All these models generate time-varying implied volatility skews in line with the empirical evidence, but they are difficult to implement and not designed for risk measurement purposes.

A paper related to ours is by Gonçalves and Guidolin (2006), who propose a two-stage approach to model the dynamic properties of the S&P 500 index options volatility surface by exploiting both cross-sectional and time series information. However, their paper differs from ours in various respects. First, we adopt a one-step Kalman filter approach instead of a two-stage model. From a theoretical standpoint, a two-step procedure is suboptimal relative to a simultaneous estimation, when the model can be expressed in a state-space form (Harvey, 1989). The Kalman filter approach we follow provides optimal estimates (via Maximum Likelihood) and forecasts of the latent variables, and can be extended to allow for exogenous variables and conditional heteroskedasticity in the latent factors. Second, while Gonçalves and Guidolin (2006) assume that today's prices of the underlying asset are tomorrow's best forecasts, we specify a model for the underlying asset returns. Also, we link the dynamics of the underlying returns to that of the volatility skew, after finding evidence of a significant correlation between the two. Finally, our analysis is focused mainly on the volatility skew, while Gonçalves and Guidolin (2006) investigate the dynamics of the entire volatility surface.

In what follows we investigate a simple application of the Kalman filter technique to the risk management of portfolios of futures options on the S&P 500.[2] By implementing a standard Monte Carlo technique, we produce density forecasts of the daily changes in the marked-to-market value of four option portfolios. The distributions of the returns predicted by our model are then compared with the actual daily profits and losses (P&L) on the option portfolios, across 1 year. A comparison is also drawn with the portfolio returns estimated according to two methods which are widely used in practice to model changes in implied volatility. The first benchmark is the sticky-delta model by Derman (1999). The second benchmark applies a Taylor expansion of Black–Scholes option prices to account for first and second order changes in the underlying (delta and gamma), and first order changes in the volatility level (vega). Various evaluation techniques are employed to assess the goodness of the P&L density forecasts. A special emphasis is placed on the accuracy of the tails, which are relevant for VaR computations.

Our main findings can be summarized as follows. First, we find that the density forecasts generated from the Kalman filter model display a good fit to the actual portfolio returns, superior to that observed for the sticky-delta and vega–gamma alternatives. This result holds for all the kinds of option portfolios we consider, and indicates that more than one factor is needed to explain the dynamics of the volatility curve itself. The Kalman filter approach suggested in this paper seems to represent a simple, yet effective, way of taking this into account. Secondly, daily forecasts can be quite noisy. Producing accurate density forecasts of daily returns on option portfolios seems to be a hard task, even for fairly basic portfolios of plain-vanilla options.

The paper is structured as follows. Section 2 presents the data set that we use. In Section 3 we describe the Kalman filter model for the dynamics of the volatility skew. The two alternative methods investigated in this study are presented in Section 4. Section 5 illustrates the criteria for building the option portfolios. In Sections 6 and 7 we discuss how to derive the density forecasts from the Monte Carlo exercise and assess their goodness-of-fit performance. Section 8 presents an extension of our approach to model the dynamics of multiple volatility skews for different maturities. Section 9 concludes.

2. The data set

Our data set consists of daily data on quarterly futures options (March, June, September and December expiries) on the S&P 500 index traded at the Chicago Mercantile Exchange over the period 1999–2003. The first four years are employed to estimate the parameters of the models, whereas the assessment of their forecasting properties is carried out over the last year.

We only use closing option prices on the two quarterly contracts closest to expiry, except for the days within 2 weeks to expiration, when we roll on to the next contract. This ensures that we always use a very liquid contract with a wide range of strikes. The usual no-arbitrage restrictions for futures options are applied to filter the option data. Also, we employ the Barone-Adesi and Whaley (1987) approximation for American options to derive pseudo-European prices. Throughout the entire analysis, we exclude in-the-money (ITM) options, which are less liquid and more sensitive to non-synchronicity pricing errors than out-of-the-money (OTM) options.

We also filter out some options with extreme strikes, which may have very low liquidity. In particular, we eliminate the options traded at the minimum price (tick), as well as those options for which a change in the premium equal to the tick size yields a change in the corresponding implied volatility larger than 15% of the volatility itself. After applying all the relevant filters, we end up with an average number of strikes/options of 78 a day on the relevant futures contract, with a minimum of 44 and a maximum of 99. These options on average span a range of deltas between 0.13 for OTM calls and −0.05 for OTM puts.

[2] From the risk management perspective, very little attention has been dedicated to the effects of the dynamics of the volatility skew/surface on the vega risk of a portfolio of options, and on the possible interaction with other risk factors. Malz (2001) incorporates smile effects on VaR by letting implied volatilities of options with a different delta vary in correlated fashion. Panigirtzoglou and Skiadopoulos (2004) assess the performance of a smile-consistent option pricing approach within a VaR framework.