Journal of Empirical Finance 49 (2018) 157–177

Contents lists available at ScienceDirect

Journal of Empirical Finance

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

Forecasting the term structure of option implied volatility: The power of an adaptive method

Ying Chen [a], Qian Han [b,c], Linlin Niu [b,c,d,*]

[a] *Department of Statistics & Applied Probability, National University of Singapore, Singapore*
[b] *Wang Yanan Institute for Studies in Economics (WISE), Xiamen University, China*
[c] *MOE Key Laboratory of Econometrics, Xiamen University, China*
[d] *Gregory and Paula Chow Center for Economic Research, Xiamen University, China*

A R T I C L E I N F O

JEL classification:
C32
C53

Keywords:
Term structure of implied volatility
Local parametric models
Forecasting

A B S T R A C T

We model the term structure of implied volatility (TSIV) with an adaptive approach to improve predictability, which treats dynamic time series models of globally time-varying but locally constant parameters and uses a data-driven procedure to find the local optimal interval. We choose two specifications of the adaptive models: a simple local AR (LAR) model for a univariate implied volatility series and an adaptive dynamic Nelson–Siegel (ADNS) model of three factors, each based on an LAR, to model the cross-section of the TSIV simultaneously with parsimony. Both LAR and ADNS models uniformly outperform more than a dozen alternative models with significance across maturities for 1–20 day forecast horizons. Measured by RMSE and MAE, the forecast errors of the random walk model can be reduced by between 20% and 60% for the 5 to 20 days ahead forecast. In terms of prediction accuracy of future directional changes, the adaptive models achieve an accuracy range of 60%–90%, which strictly dominates the range of 30%–59% of the alternative models.

1. Introduction

Option prices contain rich information about future underlying asset returns as well as volatilities. Canina and Figlewski (1993) and Jiang and Tian (2005) find that the implied volatility of the at-the-money option is a more efficient forecast of realized volatility than volatility measures based on historical data. Breeden and Litzenberger (1978) and voluminous following studies show that, for a given maturity, the entire risk-neutral distribution of the underlying asset returns is foreseen in option prices across the strikes as market investors with heterogeneous expectations about the future price movement trade at different strike prices.

In comparison, the information content of implied volatilities along the maturity dimension, i.e., the term structure of implied volatilities (TSIV), is less explored. Stein (1989) finds that option implied volatilities of long terms overreact to shocks to the short term volatility, which is inconsistent with the expectation hypothesis. Since then, Diz and Finucane (1993), Heynen et al. (1994), Campa and Chang (1995), Byoun et al. (2003), and Mixon (2007) debate whether the expectation hypothesis holds for the implied volatility term structure and do not reach consensus. Some studies conclude that implied volatilities with different durations do behave differently (see, for example, Xu and Taylor, 1994; Christoffersen et al., 2008; Guo et al., 2014). However, it remains unclear how useful the information from the entire term structure is for out-of-sample prediction and how performance varies for different forecast horizons.

* Correspondence to: Rm A306, Economics Building, Xiamen University, Xiamen, 361005, Fujian, China.
E-mail address: llniu@xmu.edu.cn (L. Niu).

https://doi.org/10.1016/j.jempfin.2018.09.006
Received 27 September 2016; Received in revised form 14 August 2018; Accepted 24 September 2018
Available online 3 October 2018

We examine this issue using an adaptive approach. The adaptive approach is first proposed in Mercurio and Spokoiny (2004) to forecast volatility where the basic model is simply a varying mean process, and is shown to outperform the more sophisticated GARCH model in the out-of-sample prediction. The approach takes a non-parametric view of the persistence of time series and considers nonstationarity to be generated by time-varying structures in the underlying models. Time series possessing these properties can be observationally equivalent to processes generated by long memory models or models with sophisticated time-varying mechanisms such as clustering properties, abrupt or smooth structural breaks, or regime-switches, etc., see Diebold and Inoue (2001) and Granger and Hyung (2004) for examples. Chen et al. (2010) extend the approach further with a local AR (LAR) model with time-varying parameters and demonstrate its superiority in out-of-sample forecasts for realized volatility against several popular long memory models and regime switching models. Moreover, Giacomini et al. (2009) estimate portfolio risk with the time-varying copula model. Spokoiny et al. (2013) extend the local approach to a quantile regression to investigate the tail dependence of the Hong Kong stock market and to analyze the distributions of the risk factors of temperature dynamics. Chen and Niu (2014) apply an LAR model with exogenous variables (LARX) to successfully forecast the yield curve. Haerdle et al. (2015) develop a local adaptive multiplicative error model for high-frequency forecasts. In these works, choosing a sophisticated yet simple model as the local model is the key for successful application, especially for forecast performance.

Recognizing that implied volatility, as with many other financial time series, is often highly persistent with nonstationary features, we treat the series as globally time-changing, but locally stationary, processes. We apply a LAR model as a basic specification to forecast the implied volatility of a particular term and strike price, and an adaptive dynamic term structure model built on the LAR as an extended specification for the TSIV of a strike price, one of the Adaptive Dynamic Nelson–Siegel (ADNS) models without exogenous macrovariables used in Chen and Niu (2014) for yield curve modeling. More specifically, we propose two specifications of the adaptive model. The first models the single implied volatility as a univariate time series; the other models the entire term structure given a strike price to explore information across maturity. Mimicking a real-time forecasting scenario, at each point in time, we model the time series parametrically with past historic information. We assume that the parameters are approximately constant (homogeneous), but only up to a point beyond which homogeneity will be rejected by a statistical test. Once the test procedure selects the longest possible homogeneous interval, the prediction is made, assuming that the homogeneity will remain within the forecast horizon. This adaptive approach, thus, strikes a balance between information efficiency and stationarity concerns. It not only uses the longest sample possible under homogeneity to increase information efficiency, but also limits the sample to a properly chosen interval to reduce parameter instability. As comparisons, we use thirteen models assessed in Guo et al. (2018).

The results are encouraging. The adaptive models dramatically outperform the alternative models. The ADNS model is more parsimonious with higher efficiency for the whole TSIV than the LAR model for a single IV. While Guo et al. (2018) find their models may slightly improve predictability against the random walk benchmark model up to the 5-day horizon, we find that our two adaptive models strongly improve predictability against their models up to the 20-day horizon; the longer the horizon, the better the performance. The key driving force is the adaptive process used. The LAR for a single IV and ADNS for the TSIV perform almost equally well, and both, remarkably, outperform the random walk model in all combinations of maturity and horizon. The ADNS parsimoniously combines cross-section information in three factor LAR processes without modeling each maturity separately.

In a striking comparison to the marginal improvement of other models relative to the random walk, the ADNS significantly reduces forecasting errors for almost all maturities. For example, when it comes to forecast the 20-day ahead implied volatility for 730-day maturity calls, the root mean squared errors (RMSE) and mean absolute errors (MAE) for the ADNS are only 59% and 46% of those of the random walk model, while the best forecast among the other models does not beat the random walk model. Moreover, no single model, except the ADNS, consistently beats the benchmark random walk model across all forecast horizons. Results of out-of-sample R^2 also demonstrate that only the adaptive models can outperform random walk on 5- to 20-days forecast horizons. We further verify the statistical significance of the adaptive models against all models 5- to 20-days forecast horizons by computing the Diebold–Mariano (DM) (Diebold and Mariano, 1995) and the Clark–West (CW) (Clark and West, 2006, 2007) test statistics. In terms of the directional forecast of future changes, while alternative models accurately predict between 30% and 59% of the time, the adaptive models have prediction accuracy ranged between 60% and 90%.

The clear dominance of our adaptive models over other models suggests the power of adaptive forecasting with suitable local models; prediction accuracy increases further for longer maturities and longer horizons.

Our work is closely related to the literature on implied volatility surface (IVS) modeling. Mapping option price quotes in terms of implied volatility against the maturity and strike price dimensions generates the IVS. For practitioners, forecasting the IVS has become more and more important for risk management and in developing trading strategies, and is crucial in the market of volatility derivatives. In the academic literature, the joint dynamics of the IVS are often factorized along both dimensions of strike price and maturity, as shown by Goncalves and Guidolin (2006) and Neumann and Skiadopoulos (2013), among many others. The strike price dimension, widely known as the option smile, has received much more attention than the maturity dimension. The sources and implications of the implied volatility smile on the underlying asset returns are often emphasized. For example, Bollen and Whaley (2004) proposes net hedging demand as an alternative explanation of the option smile beyond stochastic volatility and jumps in asset prices. In this paper, we focus on the information content embodied in the TSIV in terms of out-of-sample forecasting. Understanding the role of the TSIV in implied volatility forecasting is necessary and helpful for IVS modeling.

Our study is also related to the emerging literature on the VIX and VIX derivatives. Luo and Zhang (2012) find that the VIX with different times to maturity subsumes all information contained in historical volatilities. However, they do not consider using the information of the entire term structure to predict future implied volatilities. Zhu and Zhang (2007) and Lin (2013) both show that modeling the variance term structure is important for VIX derivatives pricing. In our robustness check with the VIX futures term structure data, we demonstrate the power of using the adaptive methods to predict VIX futures directly.