Journal of Banking & Finance 40 (2014) 303–320

Contents lists available at ScienceDirect

Journal of Banking & Finance

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

The importance of the volatility risk premium for volatility forecasting ☆

Marcel Prokopczuk [a,b,*], Chardin Wese Simen [b]

[a] Zeppelin University, 88045 Friedrichshafen, Germany
[b] ICMA Centre, Henley Business School, University of Reading, Reading RG6 6BA, UK

ARTICLE INFO

Article history:
Received 23 May 2013
Accepted 1 December 2013
Available online 8 December 2013

JEL classification:
G13
G17

Keywords:
Volatility forecasting
Volatility risk premium
Implied volatility

ABSTRACT

In this paper, we study the role of the volatility risk premium for the forecasting performance of implied volatility. We introduce a non-parametric and parsimonious approach to adjust the model-free implied volatility for the volatility risk premium and implement this methodology using more than 20 years of options and futures data on three major energy markets. Using regression models and statistical loss functions, we find compelling evidence to suggest that the risk premium adjusted implied volatility significantly outperforms other models, including its unadjusted counterpart. Our main finding holds for different choices of volatility estimators and competing time-series models, underlying the robustness of our results.

1. Introduction

A plethora of academic publications compare option implied to time-series forecasts of realized volatility such as historical volatility and GARCH-type models.[1] Surprisingly, these studies pay little attention to the fact that implied volatility is obtained under the risk-neutral measure, *Q*, whereas the quantity to be forecasted, i.e. realized volatility, is observed under the physical measure, *P*. Thus, directly comparing option implied volatility to time-series models of volatility requires the assumption that the market price of volatility risk is zero. However, Carr and Wu (2009), Driessen et al. (2009), Trolle and Schwartz (2010), Mueller et al. (2011), and Prokopczuk and Wese Simen (2012) convincingly reject this assumption. They document a significant and time-varying volatility risk premium that effectively drives a wedge between the volatility forecasted under *Q* and subsequently realized under *P*. In light of this, it is natural to ask: can the forecasting performance of implied volatility be improved by adjusting for the volatility risk premium?

Our answer is "yes". In reaching this conclusion, we make two important contributions to the volatility forecasting literature. First, we build on the model-free implied volatility (MFIV) of Jiang and Tian (2005) to propose a simple and non-parametric adjustment to account for the market price of volatility risk. To the best of our knowledge, we are the first to study the role of the volatility risk premium for volatility forecasting in a model-free setting. This is in stark contrast with the approaches of Lamoureux and Lastrapes (1993) and Poteshman (2000), who rely on explicit option pricing models. Our approach also differs from that of Chernov (2007), in that we neither rely on an approximation nor on a small subset of option prices, specifically At-The-Money (ATM) options, as the author does.

Our second contribution consists of a thorough empirical assessment of the importance of our adjustment. To do this, we proceed in three stages. We begin by evaluating the relative information content of the volatility risk premium adjusted model-free implied volatility (RMFIV) vis-à-vis other models that include historical volatility (HIST) and GARCH-type models. We do this by estimating regressions of realized volatility on alternative forecasts of volatility. We then use four statistical loss functions, i.e. mean absolute errors (MAE), mean squared errors (MSE), mean absolute percentage errors (MAPE) and mean squared percentage errors (MSPE) to investigate the forecasting accuracy of each model. Lastly, we use the Diebold–Mariano and the non-parametric

☆ Marcel Prokopczuk gratefully acknowledges financial support from the British Academy. We thank two anonymous referees, Davide Avino, Christian Dorion, Olaf Korn, Ogonna Nneji and seminar participants at the FDIC Annual Derivatives Securities and Risk Management Conference (2013), the European Meeting of the Econometric Society (2013), the Meeting of the German Finance Association (2013), the Financial Management Association Meeting (2013), Zeppelin University and the University of Reading for valuable comments.

* Corresponding author at: Zeppelin University, 88045 Friedrichshafen, Germany. Tel.: +49 754160092231.

E-mail addresses: marcel.prokopczuk@zu.de (M. Prokopczuk), c.wesesimen@icmacentre.ac.uk (C. Wese Simen).

[1] See Poon and Granger (2003) for an excellent survey.

0378-4266/$ - see front matter © 2013 Elsevier B.V. All rights reserved.
http://dx.doi.org/10.1016/j.jbankfin.2013.12.002

Table 1
Univariate and encompassing forecasts for crude oil's 30-day realized volatility.

	α	β_{HIST}	β_{GJR}	β_{ATM}	β_{MFIV}	β_{RMFIV}	Adj R^2	Wald	DW	Nobs
HIST	0.11	0.69					0.47	19.18	2.33	220
	(3.85)	**(7.52)**						**[0.00]**		
GJR	0.17		0.48				0.32	64.25	1.81	220
	(3.28)		**(3.19)**					**[0.00]**		
ATM IV	0.00			0.92			0.44	4.68	1.58	220
	(0.08)			**(7.84)**				**[0.01]**		
MFIV	−0.06				1.07		0.57	11.27	1.75	220
	(−1.41)				**(8.03)**			**[0.00]**		
RMFIV	−0.01					1.08	0.62	3.83	2.05	220
	(−0.28)					**(9.67)**		**[0.02]**		
HIST + ATM IV	0.01	0.48		0.45			0.53	44.78	2.30	220
	(0.46)	**(5.89)**		**(4.53)**				**[0.00]**		
HIST + MFIV	−0.04	0.20			0.83		0.58	6.69	2.04	220
	(−1.15)	**(2.12)**			**(6.37)**			**[0.01]**		
HIST + RMFIV	−0.01	−0.06				1.15	0.62	0.39	1.99	220
	(−0.35)	(−0.39)				**(4.87)**		[0.53]		
GJR + ATM IV	0.01		0.19	0.71			0.46	12.39	1.94	220
	(0.35)		(1.64)	**(5.19)**				**[0.00]**		
GJR + MFIV	−0.06		−0.01		1.07		0.57	0.01	1.74	220
	(−1.25)		(−0.07)		**(5.17)**			[0.93]		
GJR + RMFIV	−0.02		−0.19			1.34	0.64	10.71	1.80	220
	(−0.59)		(−1.85)			**(5.47)**		**[0.00]**		
ATM + MFIV	−0.06			−0.04	1.10		0.57	0.08	1.74	220
	(−1.47)			(−0.16)	**(3.77)**			[0.77]		
ATM + RMFIV	−0.01			0.02		1.06	0.62	0.05	2.06	220
	(−0.40)			(0.11)		**(4.50)**		[0.83]		

This table presents results from regressions of realized volatility on competing forecasts for the crude oil futures market. The dependent variable is realized volatility, estimated as follows:

$$RV_{t,T} = \sqrt{\frac{252}{T}\sum_{t=1}^{T}\left(\log\frac{F_{t,T}}{F_{t-1,T}}\right)^2}$$

where $RV_{t,T}$ refers to realized volatility between t and T. $F_{t,T}$ denotes the price at time t of the futures contract maturing at T. α and β_{HIST} denote the intercept and slope coefficients of historical volatility. Likewise, $\beta_{GJR}, \beta_{ATM}, \beta_{MFIV}$, and β_{RMFIV} refer to the slope coefficients of GJR, ATM IV, MFIV, and RMFIV, respectively. We report Newey–West t-statistics in brackets, computed with two lags. Adj R^2 reports the adjusted R^2 of the corresponding regression. Column "Wald" reports the Wald test statistic and associated p-value in square brackets. In univariate regressions, we restrict the intercept and slope estimates to be equal to zero and one, respectively. In multivariate regressions, we restrict the slope estimate of the model to the left to be equal to zero. DW and Nobs report the Durbin–Watson test statistic and the number of observations, respectively. Figures in bold indicate statistical significance at 5%.

signed rank tests to assess the statistical significance of differences between models.

In conducting our empirical analysis, we are careful to select three important markets, namely crude oil, heating oil and natural gas, that are purged of the host of data issues discussed in the volatility forecasting literature. These include asynchronous trading times, irregular expiration cycles, potentially imprecise dividend yield estimates, and limited range of strike prices to name but a few. We find compelling evidence to suggest that accounting for volatility risk premium significantly improves the volatility forecasting performance of MFIV. Typically, RMFIV yields the smallest average forecasting errors of all models. This is true for all loss functions and markets. More important, our formal statistical tests show that the difference between RMFIV and its competitors is not only economically large but also statistically significant. Our results are robust to alternative proxies for realized volatility and competing time-series models, further highlighting the importance of our findings.

The remainder of this paper proceeds as follows. Section 2 provides a brief overview of extant studies on volatility forecasting. Section 3 describes our dataset and empirical methodology. Section 4 discusses our main findings. Section 5 contains robustness checks. Finally, Section 6 concludes.

2. Literature

Arguably, one of the most controversial studies on the information content of option markets for volatility forecasting in equity markets is that by Canina and Figlewski (1993). The authors study the information content of option implied volatility and find that historical volatility forecasts are superior to option implied forecasts. Their findings cast doubt on the informational efficiency of options markets. In a subsequent study, Fleming (1998) reaches a different conclusion, reporting that forecasts based on option implied volatility outperform those based on historical volatility. Using a more refined econometric methodology based on non-overlapping data, Christensen and Prabhala (1998) corroborate this finding by showing that implied volatility outperforms historical forecasts.

Studying the US Dollar/Deutsche Mark and US Dollar/Yen markets, Guo (1996) documents the superior forecasting power of implied variance extracted from the Hull and White (1987) pricing formula. Similarly, Jorion (1995) examines the information content and predictive power of option implied volatility in the Deutsche Mark, Yen, and Swiss Franc markets. He reports that time-series models underperform option implied forecasts even when given the advantage of calibration over the whole sample. Relatedly, Martens and Zein (2004) compare the information content of implied volatility to time-series models that exploit high-frequency data in several markets, including the Yen/US Dollar market. They report that, forecasts based on intra day data sometimes outperform option implied forecasts. Finally, Charoenwong et al. (2009) assess the predictive power of implied volatility extracted from options traded on different venues namely the Philadelphia Stock Exchange, the Chicago Mercantile Exchange, and the over-the-counter (OTC) market. Their study concludes that, irrespective