Journal of Econometrics 190 (2016) 176–196

Contents lists available at ScienceDirect

Journal of Econometrics

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

A tale of two option markets: Pricing kernels and volatility risk[☆]

Zhaogang Song[a,*], Dacheng Xiu[b]

[a] Carey Business School, John Hopkins University, 100 International Drive, Baltimore, MD 21202, USA

[b] University of Chicago, Booth School of Business, 5807 South Woodlawn Avenue, Chicago, IL 60637, USA

ARTICLE INFO

Article history:
Received 11 February 2014
Received in revised form
15 April 2015
Accepted 2 June 2015
Available online 25 September 2015

JEL classification:
G12
G13

Keywords:
Multi-factor volatility model
Pricing kernel
Variance swaps
VIX options

ABSTRACT

Using both S&P 500 option and recently introduced VIX option prices, we study pricing kernels and their dependence on multiple volatility factors. We first propose nonparametric estimates of marginal pricing kernels, conditional on the VIX and the slope of the variance swap term structure. Our estimates highlight the state-dependence nature of the pricing kernels. In particular, conditioning on volatility factors, the pricing kernel of market returns exhibit a downward sloping shape up to the extreme end of the right tail. Moreover, the volatility pricing kernel features a striking U-shape, implying that investors have high marginal utility in both high and low volatility states. This finding on the volatility pricing kernel presents a new empirical challenge to both existing equilibrium and reduced-form asset pricing models of volatility risk. Finally, using a full-fledged parametric model, we recover the joint pricing kernel, which is not otherwise identifiable.

1. Introduction

In addition to market risk, volatility risk has been well documented as an essential component of time-varying investment opportunities. A priced volatility factor leads to a pricing kernel (or stochastic discount factor) that depends on both the market return and volatility. Nevertheless, existing estimates of pricing kernels either ignore volatility factors in a nonparametric analysis or impose strong parametric restrictions because volatility is neither directly tradable nor observable.

The lack of tradable and observable volatility has changed substantially since the introduction of the Volatility Index (VIX) in 1993 by the Chicago Board of Options Exchange (CBOE), and the introduction of VIX derivatives such as futures and options in 2004 and 2006, respectively.[1] The VIX, derived from S&P 500 options as the square root of the expected average variance over the next 30 calendar days, provides investors with a direct measure of volatility; and VIX derivatives offer investors convenient instruments for trading on the volatility of the S&P 500 index.[2] As a result, the VIX is constantly exposed in the media spotlight, and VIX options have achieved abundant liquidity and become the third-most-active contracts at the CBOE as of October 2011.

Taking advantage of newly available VIX options, combined with S&P 500 index options, we propose a nonparametric framework to study the marginal pricing kernel as well as its dependence on volatility, incorporating multiple volatility factors. We also provide a complementary parametric analysis for the joint pricing kernel which is not identifiable nonparametrically. Together, we

☆ We are grateful to the Editor Oliver Linton and two anonymous referees for many insightful comments. We also benefited from discussions with Yacine Aït-Sahalia, Andrea Buraschi, Bjorn Eraker, Peter Carr, Peter Christoffersen, Fousseni Chabi-Yo, George Constantinides, Jianqing Fan, René Garcia, Kris Jacobs, Jakub Jurek, Ilze Kalnina, Ralph Koijen, Nicholas Polson, Eric Renault, Jeffrey Russell, Neil Shephard, George Tauchen (discussant), Viktor Todorov, Grigory Vilkov (discussant), and Hao Zhou, as well as seminar and conference participants at the University of Chicago, Northwestern, Princeton, Toulouse School of Economics, Liverpool School of Management, the 2012 CICF, the 5th Annual SoFiE Conference, the Measuring Risk conference 2012, the 2012 Financial Engineering and Risk Management International Symposium, and the 2012 International Symposium on Risk Management and Derivatives. Xiu acknowledges research support by the Fama-Miller Center for Research in Finance at Chicago Booth. The views expressed herein do not reflect those of the Federal Reserve System.

* Corresponding author.
E-mail addresses: zsong8@jhu.edu (Z. Song), dacheng.xiu@chicagobooth.edu (D. Xiu).

http://dx.doi.org/10.1016/j.jeconom.2015.06.024
0304-4076/

1 The VIX, from its inception, was calculated from S&P 500 index options by inverting the Black–Scholes formula. In 2003, the CBOE amended this approach and adopted a model-free method to calculate the VIX using a portfolio of S&P 500 option price quotes.

2 In order to trade volatility, investors previously had to take positions in a delta-hedged portfolio, or option portfolios such as straddles or strangles.

document several important empirical facts on the asset pricing implications of volatility risk.

First, we estimate marginal pricing kernels of market returns nonparametrically using daily data of S&P 500 options from January 1996 through December 2012. Given the documented importance of multiple volatility factors in capturing the dynamics of option prices (Christoffersen et al. (2008), Egloff et al. (2010), Mencia and Sentana (2012), and Bates (2012)), we accommodate two volatility factors in our nonparametric estimation framework. In particular, we use both the VIX and the slope of the variance swap term structure (defined as the ratio of 12-month and 3-month variance swap rates minus 1) as conditioning variables. Our estimates show that the pricing kernel of market returns strongly depends on volatility factors. Specifically, the pricing kernel is higher conditional on a low slope of the variance term structure, which signals stressful times. This provides corroborating evidence for the recently proposed asset pricing models in which the pricing kernel of market returns depends on the priced volatility risk (Bansal and Yaron (2004), Drechsler and Yaron (2011), Bollerslev et al. (2009), Zhou and Zhu (2012), and Branger and Völkert (2012)). Moreover, conditioning on volatility factors delivers a pricing kernel that exhibits a downward sloping shape in the range between negative returns and reasonably positive returns, but an upward sloping shape at the extreme right tail with a very wide confidence band. Recently, Linn et al. (2014) cast some doubt on the U-shape of the pricing kernel documented by Aït-Sahalia and Lo (2000) and Jackwerth (2000), claiming that "the 'pricing kernel puzzle' is a byproduct of econometric technique rather than a behavioral or economic phenomenon". Their main critique lies in the common practice of using unconditional physical density estimator in the construction of pricing kernels. Our finding suggests that conditioning on volatility factors helps achieve the desired downward slopping pattern within a reasonable range of returns. However, we cannot draw any definitive conclusion on the extreme right tail of the return distribution due to an insufficient amount of data.

Second, we provide nonparametric estimates of the volatility pricing kernel using daily data on VIX options from July 2007 through December 2012. Our estimates show that the volatility pricing kernel is also strongly state dependent: it is significantly higher conditional on a high lagged VIX mainly in the left tail, and conditional on a low slope factor mainly in the right tail. That is, consistent with the recently proposed asset pricing models mentioned above, the stochastic discount factor regarding volatility risk is higher in stressful times. More importantly, we find that the volatility pricing kernel exhibits a pronounced U-shape. Such a U-shape – in particular the left tail – implies that the stochastic discount factor increases as volatility decreases towards a very low level. This presents a new empirical challenge to most asset pricing models of volatility risk which lead to a monotonically increasing volatility pricing kernel. That being said, our empirical result is in agreement with the recent theoretical model proposed by Bakshi et al. (forthcoming), where they suggest that the heterogeneity in investors' beliefs can generate such a U-shape.

Third, we estimate the joint pricing kernel of the S&P 500 index return and the VIX using a full-fledged parametric model.[3] In particular, we employ a two-factor stochastic volatility model that is more flexible than most of the existing option pricing models in the literature. The estimates show that the joint pricing kernel achieves high values in the region of negative market return and relatively high VIX levels. These observations agree with the economic intuition that high marginal utility is associated with bad economic states. However, we caution that the parametric joint pricing kernel is subject to potential model misspecification error, whereas our nonparametric analysis of the marginal pricing kernels is more robust.

Finally, we compare the nonparametric and parametric marginal pricing kernels of the market return and the VIX to shed light on the limits of parametric models. We find that the parametric pricing kernel estimate of the market return is very close to the nonparametric estimate, despite that the parametric model causes slight overestimation in the region of extremely high returns. However, the parametric pricing kernel of volatility shows a monotonically increasing shape, in sharp contrast to our nonparametric VIX pricing kernel estimates that exhibit a pronounced U-shape. We also provide further evidence of potential model misspecification, e.g., the binding positivity constraints, which may prevent the parametric model from fully matching the data. Overall, such comparison reveals that the existing parametric models may capture how the market risk is priced, but largely fail to capture the price of volatility risk.

Estimating pricing kernels from option prices is discussed in Aït-Sahalia and Lo (1998), Aït-Sahalia and Duarte (2003), Jackwerth (2000), and Rosenberg and Engle (2002), which ignore the time-varying volatility risk. These studies document a puzzling U-shape of the pricing kernel. Subsequently, many studies have proposed alternative explanations for the U-shaped pricing kernel, including models with missing state variables in Chabi-Yo et al. (2008), Chabi-Yo (2012), and Christoffersen et al. (2013), and models with heterogeneous agents in Bakshi and Madan (2008) and Ziegler (2007). Our econometric analysis and empirical study contribute to this literature by documenting the dependence of the pricing kernel on multiple volatility factors. Importantly, we find that conditioning on volatility factors delivers a pricing kernel of a downward sloping shape up to the right tail where no definitive conclusion can be made due to limited amount of data.

Our paper is also related to the large body of literature on models with priced volatility risk, including both reduced-form option pricing models, e.g., Bakshi et al. (1997), Bates (2000), Pan (2002), Eraker (2004), and Broadie et al. (2007), and equilibrium models, such as Bansal et al. (2014), Bollerslev et al. (2012), and Campbell et al. (2012).[4] Unlike these studies, our framework does not depend on any parametric restrictions on volatility dynamics that may lead to misspecified pricing kernels.

A closely related study is Bakshi et al. (forthcoming), who suggest a U-shape volatility pricing kernel while exploring the link between the monotonicity of the pricing kernel and returns on VIX option portfolios. They further provide a static stylized model with heterogeneity in beliefs to account for the U-shape, in which the volatility market is dominated by investors with zero market risk. In contrast, we provide direct estimates of the volatility pricing kernel by nonparametric analysis, and further find that the U-shape kernel is state dependent. This finding implies that the price of volatility risk is dynamic, i.e., it depends on the time-varying economic states, which presents new empirical regularities that need to be incorporated into models of the volatility risk.

Methodologically, our paper is also related to Boes et al. (2007) and Li and Zhao (2009), who estimate pricing kernels of stock

[3] A follow-up paper to our study, Jackwerth and Vilkov (2013), implemented a semiparametric exercise using the parametric Frank copula along with two nonparametric marginal distributions. Their choice of copula, however, is ad hoc and inconsistent with the implied joint distribution from any dynamic asset pricing models.

[4] Several recent studies have constructed model-free measures of risk-neutral volatility from S&P 500 options, e.g., Bakshi and Kapadia (2003), Bollerslev et al. (2009), Carr and Wu (2009), and Todorov (2010), and compared them with measures of realized volatility. They focus on the sign, time variation, and the return predictability of variance risk premia, which relates only to the conditional mean of variance distributions under different measures.