Journal of Financial Economics 118 (2015) 113–134

Contents lists available at ScienceDirect

Journal of Financial Economics

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

Tail risk premia and return predictability[☆]

Tim Bollerslev[a,b,c,*,1], Viktor Todorov[d,2], Lai Xu[e,3]

[a] Department of Economics, Duke University, Durham, NC 27708, United States
[b] NBER, United States
[c] CREATES, Denmark
[d] Department of Finance, Kellogg School of Management, Northwestern University, Evanston, IL 60208, United States
[e] Department of Finance, Whitman School of Management, Syracuse University, Syracuse, NY 13244-2450, United States

ARTICLE INFO

Article history:
Received 9 May 2014
Received in revised form
29 September 2014
Accepted 18 February 2015
Available online 26 July 2015

JEL classification:
C13
C14
G10
G12

Keywords:
Variance risk premium
Time-varying jump tails
Market sentiment and fears
Return predictability

ABSTRACT

The variance risk premium, defined as the difference between the actual and risk-neutral expectations of the forward aggregate market variation, helps predict future market returns. Relying on a new essentially model-free estimation procedure, we show that much of this predictability may be attributed to time variation in the part of the variance risk premium associated with the special compensation demanded by investors for bearing jump tail risk, consistent with the idea that market *fears* play an important role in understanding the return predictability.

When the VIX is high, it's time to buy, when the VIX is low, it's time to go.

Wall Street adage

1. Introduction

The *VIX* is popularly referred to by market participants as the "investor fear gauge." Yet, on average, only a small fraction of the *VIX* is arguably attributable to market *fears*.

☆ The research was supported by a grant from the NSF to the NBER, and CREATES funded by the Danish National Research Foundation (Bollerslev). We are grateful to an anonymous referee for her/his very useful comments. We would also like to thank Caio Almeida, Reinhard Ellwanger, Asaf Manela and seminar participants at NYU Stern, the 2013 SETA Meetings in Seoul, South Korea, the 2013 Workshop on Financial Econometrics in Natal, Brazil, the 2014 SCOR/IDEI conference on Extreme Events and Uncertainty in Insurance and Finance in Paris, France, the 2015 Midwest Finance Association meetings, Chicago and University of California Berkeley for their helpful comments and suggestions.

* Corresponding author at: Department of Economics, Duke University, Durham, NC 27708, USA.
E-mail addresses: boller@duke.edu (T. Bollerslev), v-todorov@northwestern.edu (V. Todorov), lxu100@syr.edu (L. Xu).
[1] Tel.: +1 919 660 1846.
[2] Tel.: +1 847 467 0694.
[3] Tel.: +1 315 443 3602.

http://dx.doi.org/10.1016/j.jfineco.2015.02.010

We show that rather than simply buying (selling) when the *VIX* is high (low), the genuine *fear* component of the index provides a much better guide for making "good" investment decisions.

Volatility clustering in asset returns is ubiquitous. This widely documented temporal variation in volatility (Schwert, 2011; Andersen, Bollerslev, Christoffersen, and Diebold, 2013) represents an additional source of risk over and above the variation in the actual asset prices themselves.[4] For the market as a whole, this risk is also rewarded by investors, as directly manifest in the form of a wedge between the actual and risk-neutralized expectations of the forward variation of the return on the aggregate market portfolio (Bakshi and Kapadia, 2003). Not only is the variance risk premium on average significantly different from zero, like the variance itself it also fluctuates non-trivially over time (Carr and Wu, 2009; Todorov, 2010). Mounting empirical evidence further suggests that unlike the variance, the variance risk premium is useful for predicting future aggregate market returns over and above the predictability afforded by more traditional predictor variables such as the dividend–price and other valuation ratios, with the predictability especially strong over relatively short quarterly to annual horizons (Bollerslev, Tauchen, and Zhou, 2009).[5]

The main goals of the present paper are twofold. First, explicitly recognizing the prevalence of different types of market risks, we seek to nonparametrically decompose their sum total as embodied in the variance risk premium into separate diffusive and jump risk components with their own distinct economic interpretations. Second, relying on this new decomposition of the variance risk premium, we seek to clarify where the inherent market return predictability is coming from and how it plays out over different return horizons and for different portfolios with different risk exposures.

Extending the long-run risk model of Bansal and Yaron (2004) to allow for time-varying volatility-of-volatility, Bollerslev, Tauchen, and Zhou (2009) and Drechsler and Yaron (2011) have previously associated the temporal variation in the variance risk premium with notions of time-varying economic uncertainty. On the other hand, extending the habit formation type preferences of Campbell and Cochrane (1999), Bekaert and Engstrom (2010) and Bekaert, Hoerova, and Lo Duca (2013) have argued that the variance risk premium may be interpreted as a proxy for aggregate risk-aversion. Meanwhile, as emphasized by Bollerslev and Todorov (2011b), the variance risk premium formally reflects the compensation for two very different types of risks: continuous and discontinuous price moves. The possibility of jumps, in particular, adds an additional unique source of market variance risk stemming from the locally non-predictable nature of jumps. This risk is still present even if the investment opportunity set does not change over time (i.e., even in a static economy with independent and identically distributed returns), and it remains a force over diminishing investment horizons (i.e., even for short time-intervals where the investment opportunity set is approximately constant). As discussed more formally below, these distinctly different roles played by the two types of risks allow us to uniquely identify the part of the variance risk premium attributable to market *fears* and the special compensation for jump tail risk.

Our estimation of the separate components of the variance risk premium builds on and extends the new econometric procedures recently developed by Bollerslev and Todorov (2014). The basic idea involves identifying the shape of the risk-neutral jump tails from the rate at which the prices of short maturity options decay for successively deeper out-of-the-money contracts. Having identified the shape of the tails, their levels are easily determined by the actual prices of the options. In contrast to virtually all parametric jump-diffusion models hitherto estimated in the literature, which restrict the shape of the tail decay to be constant over time, we show that the shapes of the nonparametrically estimated jump tails vary significantly over time, and that this variation contributes non-trivially to the temporal variation of the variance risk premium. The statistical theory underlying our new estimation procedure is formally based on an increasing cross-section of options. Importantly, this allows for a genuine predictive analysis avoiding the look-ahead bias which invariably plagues other more traditional parametric-based estimation procedures relying on long-span asymptotics for the tail estimation.

The two separately estimated components of the variance risk premium each exhibit their own unique dynamic features. Although both increase during times of financial crisis and distress (e.g., the 1997 Asian crisis, the 1998 Russian default, the 2007–08 global financial crisis, and the 2010 European sovereign debt crisis), the component due to jump risk typically remains elevated for longer periods of time.[6] By contrast, the part of the variance risk premium attributable to "normal" risks rises significantly during other time periods that hardly register in the jump risk component (e.g., the end of the dot-com era in 2002–03). Counter to the implications from popular equilibrium-based asset pricing models, nonparametric regression analysis also suggests that neither of the two components of the variance risk premium can be fully explained as nonlinear functions of the aggregate market volatility.[7]

[4] Following the classical Intertemporal Capital Asset Pricing Model (ICAPM) of Merton (1973), variance risk has traditionally been associated with changes in the investment opportunity set, which in turn induce a hedging component in the asset demands.

[5] Recent studies corroborating and extending the predictability results in Bollerslev, Tauchen, and Zhou (2009) include Drechsler and Yaron (2011), Han and Zhou (2011), Du and Kapadia (2012), Eraker and Wang (2015), Almeida, Vicente, and Guillen (2013), Bekaert and Hoerova (2014), Bali and Zhou (2015), Camponovo, Scaillet, and Trojani (2013), Kelly and Jiang (2014), Li and Zinna (2014), Vilkov and Xiao (2013), and Bollerslev, Marrone, Xu, and Zhou (2014), among others. The empirical results in Andreou and Ghysels (2013) and Bondarenko (2014) also suggest that the variance risk premium cannot be explained by other traditional risk factors.

[6] The overall level of the market volatility also tends to mean revert more quickly than the jump risk premia following all of these events.

[7] The habit persistence model of Campbell and Cochrane (1999), for example, and its extension in Du (2010), imply such a nonlinear relationship.