Journal of Financial Economics 000 (2016) 1–26

Contents lists available at ScienceDirect

Journal of Financial Economics

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

The price of variance risk☆

Ian Dew-Becker[a,e], Stefano Giglio[b,f,*], Anh Le[c], Marius Rodriguez[d]

[a] Kellogg School of Management, Northwestern University, USA
[b] Booth School of Business, University of Chicago, USA
[c] Smeal College of Business, Pennsylvania State University, USA
[d] Board of Governors of the Federal Reserve System, USA
[e] National Bureau of Economic Research
[f] National Bureau of Economic Research and the Centre for Economic Policy Research

ARTICLE INFO

Article history:
Received 7 August 2015
Revised 9 March 2016
Accepted 11 April 2016
Available online xxx

JEL classifications:
G12
G13

Keywords:
Variance risk premium
Recursive preferences
Disasters

ABSTRACT

Between 1996 and 2014, it was costless on average to hedge news about future variance at horizons ranging from 1 quarter to 14 years. Only unexpected, transitory realized variance was significantly priced. These results present a challenge to many structural models of the variance risk premium, such as the intertemporal CAPM and recent models with Epstein–Zin preferences and long-run risks. The results are also difficult to reconcile with macro models in which volatility affects investment decisions. At the same time, the data allows us to distinguish between different disaster models; a model in which the stock market has a time-varying exposure to disasters and investors have power utility fits the major features of the variance term structure.

1. Introduction

The recent explosion of research on the effects of volatility in macroeconomics and finance shows that economists care about uncertainty shocks. It appears that investors, on the other hand, do not. In the period since 1996, it has been costless on average to hedge news about future volatility in aggregate stock returns; in other words investors have not been required to pay for insurance against volatility news. Many economic theories—both in macroeconomics and in finance—have the opposite prediction. The recent consumption-based asset pricing literature is heavily influenced by Epstein and Zin (1991) preferences, which in standard calibrations, with a preference for early resolution of uncertainty, imply that investors have a strong desire to hedge news about future uncertainty, and hence should be willing to pay large premia for insurance against volatility shocks. Furthermore, in recent macroeconomic models, shocks to uncertainty about the future can induce large fluctuations in the economy.[1] But if increases

☆ The views in this paper are solely the those of the authors and do not reflect the views of the Board of Governors of the Federal Reserve System. We appreciate helpful comments from and conversations with Dante Amengual, Jean-Francois Bacmann, Brandon Bates, Nick Bloom, John Campbell, John Cochrane, Steve Davis, Itamar Drechsler, Greg Duffee, Darrell Duffie, Cam Harvey, Jing-zhi Huang, Dana Kiku, Loriano Mancini, Ian Martin, Tim McQuade, Ivan Shaliastovich, Fabio Trojani, Harald Uhlig, Joel Vanden, Andrea Vedolin, Jessica Wachter, Liuren Wu, and seminar participants at the Booth Econometrics Workshop, University of New South Wales, University of Sydney, University of Technology, Sydney, University of Adelaide, Northwestern Kellogg, the 2014 Macro Finance Society Meeting, the Office for Financial Research, the Midwest Finance Association, ASU Sonoran Winter Finance Conference, Oxford, LSE, HEC, CSEF, Bocconi, MIT Sloan, Ohio State, Southern Denmark University, Copenhagen Business School, Penn State Smeal, SFS Cavalcade, SITE, ITAM, WFA, FIRS, SoFie, CEPR Gerzensee and AFA. We are particularly grateful to Gerald Lim and Markit Group for providing variance swap data.

* Corresponding author.
E-mail address: Stefano.giglio@chicagobooth.edu (S. Giglio).

http://dx.doi.org/10.1016/j.jfineco.2016.04.003
0304-405X/

[1] See, e.g., Bloom (2009), Bloom, Floetotto, Jaimovich, Saporta-Eksten, and Terry (2014), Christiano, Motto, and Rostagno (2014), Fernandez-Villaverde, Guerron, Rubio-Ramirez, and Uribe (2011), and Gourio (2012) Gourio (2013).

in economic uncertainty can drive the economy into a recession, we would expect that investors would want to hedge those shocks.[2] The fact that shocks to expected volatility have not earned a risk premium thus presents a challenge to a wide range of recent research.

As a concrete example, consider the legislative battles over the borrowing limit of the US in the summers of 2010 and 2011. Those periods were associated with increases in both financial measures of uncertainty, e.g., the Chicago Board Options Exchange's Volatility Index (VIX), and also the measure of policy uncertainty from Baker, Bloom, and Davis (2014). Between July and October, 2011, the 1-month variance swap rate—a measure of investor expectations for Standard & Poor's (S&P) 500 volatility over the next month—rose every month, from 16.26 to 42.32% (annualized, computed at the beginning of the month). But those shocks also had small effects on *realized* volatility in financial markets; for example, realized volatility actually decreased between August and September of 2011. The debt ceiling debate caused uncertainty about the future to be high during the whole period, but did not correspond to high contemporaneous volatility during the same period. It is precisely this imperfect correlation between realized volatility and expectations of future volatility that allows us to disentangle the pricing of their shocks. In this paper, we directly measure how much people pay to hedge shocks to expectations of future volatility. We find that news shocks have been unpriced: any investor could have bought insurance against volatility shocks for free, and therefore any investor could have freely hedged the increases in uncertainty during the debt ceiling debate.

We measure the price of variance risk using novel data on a wide range of volatility-linked assets both in the US and around the world, focusing primarily on variance swaps with maturities between one month and ten years. The data cover the period 1996–2014. Variance swaps are assets that pay to their owner the sum of daily squared stock market returns from their inception to maturity. They thus give direct exposure to future stock market volatility and are the most natural and direct hedge for the risks associated with increases in aggregate economic uncertainty. Importantly, though, we show that our results hold in a range of other markets, including index options, which are both more liquid and traded on exchanges.

The analysis of the pricing of variance swaps yields two simple but important results. First, news about future volatility is unpriced in our sample—exposure to volatility news did not earn a risk premium. Second, exposure to *realized variance* is strongly priced in our data, with an annualized Sharpe ratio of −1.3—four times larger than the Sharpe ratio on equities. We find that it is the downside component of realized volatility that investors are specifically trying to hedge, consistent with the results of Bollerslev and Todorov (2011) and Segal, Shaliastovich, and Yaron (2015). We conclude that over our sample, investors paid a large amount of money for protection from extreme negative shocks to the economy (which mechanically generate spikes in realized volatility), but they did not pay to hedge news that uncertainty or the probability of a disaster has changed.

The results present a challenge to a wide range of models. From a finance perspective, Merton's (1973) intertemporal capital asset pricing model says that assets that have high returns in periods with good news about future investment opportunities are viewed as hedges and thus earn low average returns. Since expected future volatility is a natural state variable for the investment opportunity set, the covariance of an asset's returns with shocks to future volatility should affect its expected return, but it does not.[3]

Consumption-based models with Epstein and Zin preferences have similar predictions. Under Epstein–Zin preferences, marginal utility depends on lifetime utility, so that assets that covary positively with innovations to lifetime utility earn high average returns.[4] If high expected volatility is bad for lifetime utility (either because volatility affects the path of consumption or because volatility reduces utility simply due to risk aversion), then volatility news should be priced.[5]

As a specific parameterized example with Epstein–Zin preferences, we study variance swap prices in Drechsler and Yaron's (2011) calibrated long-run risk model. While that model represents a major innovation in being able to both generate a large variance risk premium (the average gap between the 1-month variance swap rate and realized variance) and match results about the predictability of market returns, we find that its implications for the term structure of variance swap prices and returns are distinctly at odds with the data: it predicts that shocks to future expected volatility should be strongly priced, counter to what we observe empirically.

We obtain similar results in a version of Wachter's (2013) model of time-varying disaster risk with Epstein–Zin preferences. The combination of fluctuations in the probability of disaster and Epstein–Zin preferences results in a counterfactually high price for insurance against shocks to expected future volatility relative to current volatility. Du's (2011) model of disaster risk and habit formation also fails to match the data.[6]

[2] See Berger, Dew-Becker, and Giglio (2016) for an analysis of the effects of volatility shocks on the real economy, finding that news about future volatility is not contractionary.

[3] Recently, Campbell, Giglio, Polk, and Turley (2013) and Bansal, Kiku, Shaliastovich, and Yaron (2013) estimate an ICAPM model with stochastic volatility and find that shocks to expected volatility (and especially long-run volatility) are priced in the cross-section of returns of equities and other asset classes. Although the focus on their paper is not the variance swap market, Campbell, Giglio, Polk, and Turley (2013) test their specification of the ICAPM model also on straddle returns and synthetic volatility claims, and find that the model manages to explain only part of the returns on these securities. This suggests that the model is missing some high-frequency features of the volatility market.

[4] This is true in the most common calibrations with a preference for early resolution of uncertainty. When investors prefer a late resolution of uncertainty the risk prices are reversed.

[5] Also see Branger and Volkert (2010) and Zhou and Zhu (2012) for discussions. Barras and Malkhozov (2014) study the determinants of changes in the variance risk premium over time.

[6] Similar problems with matching term structures of Sharpe ratios in structural models have been studied in the context of claims to aggregate market dividends by van Binsbergen, Brandt, and Koijen (2012). Our results thus support and complement theirs in a novel context. See also van Binsbergen and Koijen (2015) for a recent review of the broad range of evidence on downward sloping term structures. Our paper also

Please cite this article as: I. Dew-Becker et al., The price of variance risk, Journal of Financial Economics (2016), http://dx.doi.org/10.1016/j.jfineco.2016.04.003