Journal of Banking & Finance 37 (2013) 3388–3400

Contents lists available at SciVerse ScienceDirect

Journal of Banking & Finance

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

The cross-sectional relation between conditional heteroskedasticity, the implied volatility smile, and the variance risk premium

Louis H. Ederington [a,*], Wei Guan [b]

[a] Finance Division, Michael F. Price College of Business, University of Oklahoma, 205A Adams Hall, Norman, OK 73019, USA
[b] College of Business, University of South Florida St. Petersburg, 140 Seventh Avenue South, St. Petersburg, FL 33701, USA

ARTICLE INFO

Article history:
Received 11 October 2012
Accepted 14 April 2013
Available online 17 May 2013

JEL classification:
G13
G10
G12

Keywords:
Implied volatility
Volatility smile
Variance risk premium
GARCH
Conditional heteroskedasticity

ABSTRACT

This paper estimates how the shape of the implied volatility smile and the size of the variance risk premium relate to parameters of GARCH-type time-series models measuring how conditional volatility responds to return shocks. Markets in which return shocks lead to large increases in conditional volatility tend to have larger variance risk premia than markets in which the impact on conditional volatility is slight. Markets in which negative (positive) return shocks lead to larger increases in future volatility than positive (negative) return shocks tend to have downward (upward) sloping implied volatility smiles. Also, differences in how volatility responds to return shocks as measured by GARCH-type models explain much, but not all, of the variations in excess kurtosis and multi-period skewness across different markets.

1. Introduction

Along with jump risk, a leading explanation for both the implied volatility smile and the variance risk premium is stochastic volatility, i.e., the possibility that volatility will change in the future. While changes in volatility could be associated with many factors or occur randomly, one pattern which has been well documented by GARCH model estimations is that in many markets the conditional return variance responds to return shocks. For instance, in most markets, the conditional variance tends to rise following large absolute return shocks and fall following periods of very small price movements. To the extent that fluctuations in the variance of returns are associated with return shocks, as GARCH model estimations indicate, investors should anticipate greater future variance fluctuations in markets in which the variance responds sharply to return shocks than in markets in which the impact of the same size shock on the conditional variance is slight. This is the idea explored in the present paper. In particular, we investigate how variations in the implied volatility smile and in the variance risk premium across different markets relate to differences in how conditional volatility responds to surprise return shocks as measured by GARCH-type models.

Consider, for instance, the most studied implied volatility smile – that on options on US equity indices, such as the S&P 500. It is well-known that Black–Scholes (BS) implied volatilities on US stock market index options decline with the strike price in a smirk pattern. While there are other possible explanations, such as jump risk, the hedging pressure hypothesis of Bollen and Whaley (2004) and Ederington and Guan (2002), the Bakshi et al. (2003) skewness hypothesis, or the transaction cost hypothesis of Peña et al. (1999), one popular explanation of this smirk pattern is that it arises because volatility is stochastic and tends to be negatively correlated with recent stock market returns.[1] According to this explanation, because volatility tends to increase when the market drops, large multi-period market declines are more likely than large multi-period market increases imparting a negative skewness to multi-period returns and increasing the likelihood that far out-of-the-money (OTM) puts will finish in the money thereby raising their price and implied

* Corresponding author. Tel.: +1 405 325 5591.
E-mail addresses: lederington@ou.edu (L.H. Ederington), wguan@mail.usf.edu (W. Guan).

[1] For prominent examples, see Heston (1993), Duan (1995), Bakshi et al. (1997), Bates (2000), Heston and Nandi (2000), Andersen et al. (2002), Christoffersen et al. (2006) and Christoffersen et al. (2010). See also the smile chapter in Hull (2006). However, Bates (2000) and Andersen et al. (2002) conclude that volatility changes alone are insufficient to explain the stock index smile. For a similar correlation argument for jump risk, see Câmara et al. (2011).

0378-4266/$ - see front matter © 2013 Elsevier B.V. All rights reserved.
http://dx.doi.org/10.1016/j.jbankfin.2013.04.017

volatility. Consistent with this explanation, estimations of asymmetric GARCH models consistently find that conditional US stock market volatility rises following stock market declines,[2] and other studies find that implied volatilities calculated from options on stock indices are inversely correlated with stock market movements.[3]

Consider the implications of this stock market smile argument for other markets. While equity market volatility tends to increase more following a large market decline than following an equally large market advance, in several commodity markets this pattern is reversed. In other words, in these markets, volatility tends to increase more following large positive return shocks than following equally large negative return shocks.[4] Applied to these markets, the equity market smile theory would imply that in these markets multi-period returns should be positively skewed and that an upward sloping smile should be observed.

In other markets, GARCH-type models find that conditional volatility tends to increase roughly equally following positive and negative price shocks, i.e., a U-shaped news impact curve. However, within this set, the convexity of the news impact curve may differ. Suppose the same size absolute return shock tends to be followed by a much larger change in future volatility in market A than in market B. If this is the case, then rational investors should anticipate more variance changes, i.e., a higher variance of the variance, in market A. Thus if anticipated volatility changes are responsible for the smile, one would expect a more convex smile in market A and a flatter smile in B – also more kurtosis in A.

In summary, we expect downward (upward) sloping implied volatility smiles and negatively (positively) skewed multi-period returns in markets where time-series GARCH-type models indicate that negative return shocks impact future conditional volatility more (less) than positive shocks. Likewise, we expect more convex smiles and greater return kurtosis in markets in which conditional volatility is strongly impacted more by an absolute return shock than in markets where GARCH-type models find that the impact on conditional volatility is slight.

Are in fact these smile, skewness, and kurtosis patterns observed? This paper explores how cross-sectional variations in the shape of the implied volatility smile across 43 different equity, foreign exchange, and commodity markets relate to how conditional volatility in these different markets is impacted by market return shocks as estimated by asymmetric GARCH models, specifically Nelson's (1991) EGARCH model and the GJR (or TGARCH) model of Glosten et al. (1993). In other words, this paper asks how cross-sectional variations in both the slope and convexity of the implied volatility smile relate to differences in how conditional volatility responds to positive and negative return shocks in various markets. We are unaware of any previous similar cross-sectional exploration.[5] We also explore how cross-sectional differences in return skewness and kurtosis relate to parameters of the EGARCH and GJR models.

Like the smile, it is well established that implied volatilities tend to exceed ex post or realized volatilities – a difference which has been termed the variance risk premium. The usual explanation of this variance risk premium is that investors seek to hedge against the large losses associated with tail events by purchasing options thus driving implied volatilities above expected volatilities. As with the smile, the two main theories for why extreme events might be more likely than the current return distribution would predict are: (1) the risk that the return variance might increase in the future, and (2) the possibility of jumps. Bakshi and Kapadia (2003), Bakshi and Madan (2006), Carr and Wu (2009), and Bollerslev et al. (2009) attribute the variance risk premium to changing variances; Todorov (2010) to jumps; Bakshi et al. (2012) to both. If the variance risk premium is indeed due to the risk that the variance will change and if many changes in the variance are due to return shocks as captured by GARCH type models, then (like the smile) the variance risk premium should be related to the GARCH models' parameters. In particular, we should observe higher risk premia in markets in which return shocks cause large changes in the variance than in markets in which the impact of equally large return shocks on the conditional variance is small. We test whether this is the case.

The paper proceeds as follows. First, we estimate EGARCH models and news impact curves, for 43 equity, foreign exchange, and commodity markets using data from January 1996 (or shortly thereafter) through June 2006. We then explore the implications of the EGARCH parameter estimates for return skewness and kurtosis. As hypothesized, we find that kurtosis tends to be higher in markets where return shocks lead to large changes in conditional volatility and lower in markets where return shocks have less impact on conditional volatility. In 39 of our 43 markets, standardized returns, r_t/σ_t, have less kurtosis than r_t where r_t is the daily log return and σ_t is the conditional standard deviation for day t predicted by the EGARCH model. Nonetheless, in most markets, r_t/σ_t is still characterized by excess kurtosis. In short, our results suggest that a substantial part of the observed excess kurtosis is due to the impact of return shocks on conditional volatility as captured by the EGARCH model but that part is either due to variance changes not associated with return shocks or due to jumps. As hypothesized, we further find that multi-day returns tend to be positively (negatively) skewed in markets in which volatility increases more following positive (negative) return shocks than following equal negative (positive) return shocks.

Next we relate the implied volatility smile patterns in the various markets to parameters of the EGARCH model. As hypothesized, we find that in markets in which negative shocks have a stronger impact on future volatility than positive shocks, the smile tends to be downward sloping. When positive shocks have a stronger impact on future volatility than negative shocks, the smile tends to be upward sloping. In contrast to our findings for the smile's slope, we find no evidence that the convexity of the smile varies with the magnitude of the conditional volatility response to return shocks as measured by the EGARCH model. However, smile convexity does vary with the kurtosis not caused by the impact of return shocks on volatility. These results suggest that the convexity of the smile might be due more to jumps or to variance changes not associated with return shocks than to variance changes due to return shocks. Measures of how volatility responds to return shocks using the GJR model yield the basically the same results as the EGARCH model.

Finally, we relate the variance risk premium to parameters of the EGARCH and GJR models. As hypothesized, we find that the variance risk premium tends to be larger in markets in which a given size market return shock leads to a large change in the conditional variance than in markets in which the same size return shock leads to only a slight change in the conditional variance. After controlling for the GARCH and GJR model parameters, there is little relation between the variance risk premium and the remaining return kurtosis suggesting that the variance risk premium is due more to the risk that the variance will change as measured by the GARCH models, than to jumps. We also find evidence

[2] For a review of pre-2000 estimations of such models see Bekaert and Wu (2000). More recent examples include: Bekaert and Wu (2000), Wu and Xiao (2002), Li et al. (2005), Caporin and McAleer (2006), Koulakiotis et al. (2006) and Ederington and Guan (2010).

[3] Fleming et al. (1995), Whaley (2000), Low (2004), Bates (2000), Poteshman (2001), Pan (2002), Dennis et al. (2006) and Ederington and Guan (2010).

[4] Beck (2001) and Tansuchat et al. (2009). A possible reason for this pattern is that a large increase in the price of a commodity may signal that inventories are low so that supply is less elastic so that future shifts in demand or supply tend to have larger price impacts than prior to the price increase.

[5] However, Tompkins (2001) estimates smiles for a number of different markets and how they relate to extreme shocks.