Journal of Banking & Finance 37 (2013) 448–462

Contents lists available at SciVerse ScienceDirect

Journal of Banking & Finance

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

Incomplete information, idiosyncratic volatility and stock returns

Tony Berrada [a,*], Julien Hugonnier [b]

[a] University of Geneva and Swiss Finance Institute, Unimail, Boulevard du Pont d'Arve 40, 1211 Geneva 4, Switzerland
[b] EPFL and Swiss Finance Institute, Quartier UNIL Dorigny, Batiment Extranef, 1015 Lausanne, Switzerland

A R T I C L E I N F O

Article history:
Received 11 January 2012
Accepted 11 September 2012
Available online 18 September 2012

JEL classification:
G12
D83
D92

Keywords:
Idiosyncratic volatility
Incomplete information
Cross-section of stock returns

A B S T R A C T

When investors have incomplete information, expected returns, as measured by an econometrician, deviate from those predicted by standard asset pricing models by including a term that is the product of the stock's idiosyncratic volatility and the investors' aggregated forecast errors. If investors are biased this term generates a relation between idiosyncratic volatility and expected stocks returns. Relying on forecast revisions from IBES, we construct a new variable that proxies for this term and show that it explains a significant part of the empirical relation between idiosyncratic volatility and stock returns.

1. Introduction

According to textbook asset pricing theory, investors are only rewarded for bearing aggregate risk and, consequently, idiosyncratic volatility should not be priced in the cross section of stock returns. However, numerous recent empirical studies have documented a relation between stock returns and idiosyncratic volatility. In particular, Ang et al. (2006), Jiang et al. (2009), Brockman and Yan (2008), and Guo and Savickas (2010) provide evidence of a negative relation for the US stock market, and Ang et al. (2009) confirm that a similar relation also holds in other markets. There is however no consensus as to the direction of this effect. Indeed, Malkiel and Xu (2001), Spiegel and Wang (2005), and Fu (2005) find positive relations between idiosyncratic volatility and expected returns, while Longstaff (1989) finds a weakly negative relation.[1]

We propose a simple model of firm valuation under incomplete information that sheds some light on the ambiguous link between idiosyncratic volatility and stock returns. Specifically, we assume that investors observe aggregate shocks as well as the cash flows of all firms but have incomplete information about idiosyncratic shocks and therefore have to estimate the growth rates of cash flows. Rather than modeling the learning mechanism at the individual level, we assume that the investors' perceptions can be summarized by a single subjective probability measure that is equivalent to the objective or "true" probability measure. Because investors behave rationally the CAPM holds under their subjective probability measure in the sense that expected returns under this measure only reflect exposure to aggregate risk. However, it fails under the objective measure as expected returns under that measure also depend on the investors' forecast errors. Indeed, the idiosyncratic shocks perceived by investors are a combination of the true idiosyncratic shocks and forecast errors that cannot be disentangled given the available information. Since idiosyncratic volatility is defined as the loading of the firm's stock returns on the perceived idiosyncratic shocks, this implies that expected returns under the objective probability measure contain an additional term that is given by the product of the firm's idiosyncratic volatility and the investors' aggregated forecast error. This additional term, which we refer to as the idiosyncratic volatility effect, is the basis for our explanation of the relation between idiosyncratic volatility and stock returns. As explained by Timmermann (1993), Timmermann (1996), and Leuwellen and Shanken (2002) among others, unconditional tests do not capture the expected returns as perceived by investors. Rather, they measure a combination of these perceived expected returns that are solely due to aggregate

* Corresponding author. Tel.: +41 22 379 81 26.
E-mail addresses: tony.berrada@unige.ch (T. Berrada), julien.hugonnier@epfl.ch (J. Hugonnier).

[1] Peterson et al. (2011) find a negative relation between realized idiosyncratic volatility and stock returns and a positive relation between expected idiosyncratic volatility and stock returns. Chabi-Yo (2011) proposes an explanation for the relation based on coskewness and Chen and Petkova (2012) finds that idiosyncratic volatility is a proxy for risk exposure.

0378-4266/$ - see front matter
http://dx.doi.org/10.1016/j.jbankfin.2012.09.004

risk exposure, and forecast errors that are due to incomplete information. Since the weight of the forecast errors in this combination is given by the firm's idiosyncratic volatility it follows that idiosyncratic volatility can have an impact on expected returns as measured by regressions. It is important to observe that the deviation from the CAPM which is implied by our model under the objective measure is not due to a missing factor. The idiosyncratic volatility effect that we identify is generated by the investors' aggregated forecast errors and, hence, does not represent a reward for exposure to a systematic risk factor. The presence of such a component in expected returns is entirely due to incomplete information and cannot be generated by introducing additional state variables into an otherwise standard model.

If investors are unbiased in aggregated terms, that is if they consider the correct underlying model and use Bayes'rule to update their beliefs, or equivalently if there exists a representative agent with unbiased beliefs, then their aggregated forecast errors are zero on average. In this case the idiosyncratic volatility effect predicted by our model is by construction equal to zero on average and, therefore, does not affect unconditional estimates of expected returns. While it may be natural to assume that investors are Bayesian at the individual level, this assumption does not necessarily imply that the aggregation of their beliefs is itself Bayesian (see e.g. Detemple and Murthy, 1997; Berrada, 2006; Jouini and Napp, 2007) and one should therefore expect that their perceptions are biased in aggregate terms. If that is indeed the case then the idiosyncratic volatility effect is different from zero on average and, hence, affects unconditional estimates of expected returns. The existence and direction of this bias, and whether or not it can help us understand the empirical relation between idiosyncratic volatility and stock returns are the main questions we address in the empirical part of the paper.

We focus our empirical investigation on two important implications of the model. First, firms with high idiosyncratic volatility should underperform when news are bad, and overperform when news are good. In the context of our model, where the growth rates of cash flows are unobserved, bad news correspond to situations where realized earning growth is smaller than expected and induce negative perceived shocks on returns through the mechanism highlighted above. Since stocks with high idiosyncratic volatility are more exposed to such shocks, they should underperform following bad news. An identical reasoning suggests that high idiosyncratic volatility stocks should overperform following good news. This implication of the model relates to the vast literature on the post-earning announcement drift, see e.g. Ball and Brown (1968), Watts (1978), Foster et al. (1984), and Bernard and Thomas (1990). We contribute to this literature by proposing a model that not only explains the response of returns to news but also predicts a stronger effect on the return of high idiosyncratic volatility stocks.

Second, our model predicts that if there appears to be a relation between idiosyncratic volatility and stock returns in the data, then this relation should not remain significant when controlling for the idiosyncratic volatility effect. This implication of our model provides a potential explanation for the empirical results documenting a cross sectional relation between idiosyncratic volatility and stock returns. Note that while our model predicts the existence of such a relation it is silent about its direction and, therefore, can be consistent with both a negative relation (e.g. Ang et al., 2006, 2009; Jiang et al., 2009; Brockman and Yan, 2008) and a positive relation (e.g. Malkiel and Xu, 2001; Spiegel and Wang, 2005; Fu, 2005).

To test the predictions of our model, the first step is to construct a proxy for the idiosyncratic volatility effect. Since this effect is defined as the product of a stock's idiosyncratic volatility and the investors' aggregated forecast errors, we need proxies for both quantities. Following standard practice we measure a stock's idiosyncratic volatility in a given month by the standard deviation of the residuals from the 3 factor model of Fama and French (1993) run at a daily frequency. To approximate the investors' aggregated forecast errors about the growth rates of cash flows we use the average of analyst forecast revisions for end-of-year earning growth obtained from the I/B/E/S database and standardize this measure to obtain comparable quantities across firms. Our proxy for the idiosyncratic volatility effect is computed for each firm at a monthly frequency using all analyst forecasts from January 1982 to December 2007.

When we split the sample in good and bad news groups, the first implication of the model is remarkably well verified. Indeed, we find that portfolios of high and low idiosyncratic volatility stocks behave very differently following good and bad news and that the risk-adjusted effect goes in the direction predicted by the model. In particular, portfolios of high idiosyncratic volatility stocks have a significant and largely positive alpha after good news and a significant and largely negative alpha after bad news. Furthermore, the magnitude of the average idiosyncratic volatility effect for ten idiosyncratic volatility-sorted portfolios is very close to the magnitude of the alphas. In particular, the difference between alpha and our proxy for the idiosyncratic volatility effect is not statistically significant for eighteen out of the 20 portfolios we construct.

In the split sample, the evidence is mixed relative to the second implication of the model. Controlling for the idiosyncratic volatility effect reduces the magnitude and statistical significance of the alphas of the decile portfolios but fails to completely explain the cross-sectional relation between idiosyncratic volatility and stock returns. As this may be due to the presence of outliers we repeat the same regressions after applying a monthly filter that eliminates 1% of most extreme idiosyncratic volatility effects as well as those firms which are followed by less than five analysts in the given month. The results for the filtered sample are more in line with the predictions of the model. In particular, controlling for the idiosyncratic volatility effect now makes the alpha insignificant on 5 of the decile portfolios in the bad news group. Unfortunately, controlling for the idiosyncratic volatility effect still has a marginal impact on the alphas in the good news group and we therefore cannot conclude that the second implication of the model holds in the split sample. These results are confirmed by a detailed analysis of the returns on a portfolio that is long in high idiosyncratic volatility stocks and short in low idiosyncratic volatility stocks. In particular, we show that controlling for the idiosyncratic volatility effect completely eliminates the alpha of the long/short portfolio in the bad news group but has little effect on that of the good news group.

To investigate the validity of the model's second implication we perform a number of tests on the whole sample. Specifically, we follow Ang et al. (2006, 2009) in constructing ten portfolios sorted on the previous month's idiosyncratic volatility and compare the alphas of these portfolios to the idiosyncratic volatility effects implied by the model. As in Ang et al. (2006) we find a negative relation between idiosyncratic volatility and stock returns in the whole sample. In particular, a portfolio that is long in high idiosyncratic volatility stocks and short in low idiosyncratic volatility stocks produces a negative and significant alpha of −66 basis points per month. When comparing the alphas of the decile portfolios to the predictions of our model we find that the idiosyncratic volatility effect decreases as idiosyncratic volatility increases, and that its magnitude explains about half of the negative abnormal risk-adjusted return on the long/short portfolio. To confirm this finding we conduct a regression analysis which shows that controlling for the idiosyncratic volatility effect eliminates the alpha of the long/short portfolio. Remarkably, the estimated coefficient on the idiosyncratic volatility effect is not significantly different from the value predicted by the model.