Journal of Empirical Finance 34 (2015) 45–59

Contents lists available at ScienceDirect

Journal of Empirical Finance

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

The economic value of volatility timing with realized jumps☆

Ingmar Nolte [a], Qi Xu [b,*]

[a] Accounting and Finance Department, Lancaster University Management School, Lancaster University, Lancaster LA1 4YX, United Kingdom
[b] Finance Group, Warwick Business School, University of Warwick, Coventry CV4 7AL, United Kingdom

ARTICLE INFO

Article history:
Received 17 March 2014
Received in revised form 30 January 2015
Accepted 16 March 2015
Available online 1 September 2015

JEL classification:
C58
C53
G11

Keywords:
High frequency data
Jumps
Nonparametric tests
Asset allocation
Volatility forecasting
Realized volatility

ABSTRACT

This paper comprehensively investigates the role of realized jumps detected from high frequency data in predicting future volatility from both statistical and economic perspectives. Using seven major jump tests, we show that separating jumps from diffusion improves volatility forecasting both in-sample and out-of-sample. Moreover, we show that these statistical improvements can be translated into economic value. We find that a risk-averse investor can significantly improve her portfolio performance by incorporating realized jumps into a volatility timing based portfolio strategy. Our results hold true across the majority of jump tests, and are robust to controlling for microstructure effects and transaction costs.

1. Introduction

The importance of jumps in asset pricing, option pricing, and risk management is widely recognized (Ait-Sahalia, 2004). Although, resorting on jumps as a modeling device is not new, realized jumps were generally overlooked until recently. In this paper, we comprehensively investigate the role of realized jumps detected from high frequency data for the prediction of future volatility. Different from previous studies with similar focus, we not only conduct an extensive statistical evaluation of volatility forecasting using all major jump tests, but also provide new economic insights in the form of whether a risk-averse investor can significantly benefit from considering realized jumps in volatility timing based portfolio allocation strategies.

The literature on this topic can be broadly categorized into two streams: The parametric literature starting with Merton (1976) includes jump-diffusion and stochastic-volatility with jumps (SVJ) models in continuous time (Chernov et al., 2003; Eraker, 2004; Eraker et al., 2003) and GARCH-J models in discrete time (Christoffersen et al., 2012; Duan et al., 2006; Maheu and McCurdy, 2004). These parametric models are widely used in portfolio choice, option pricing, and risk management applications and the jumps introduced in models are ex ante in nature. As Backus et al. (2011) admit: “A jump component, in this context, is simply a mathematical device that produces nonnormal distributions”.

☆ We thank Christian Wolff (editor), an anonymous associate editor, two anonymous referees, Marcelo Fernandes, Sandra Nolte, Alessandro Palandri, and conference and seminar participants at EEA-ESEM 2013 (Gothenburg), IFABS 2013 (Nottingham), and Warwick Business School for their helpful comments.

* Corresponding author.
E-mail addresses: I.Nolte@lancaster.ac.uk (I. Nolte), Qi.Xu.11@mail.wbs.ac.uk (Q. Xu).

http://dx.doi.org/10.1016/j.jempfin.2015.03.019
0927-5398/

The second stream of the literature considers nonparametric approaches. Recently, many nonparametric jump tests (Ait-Sahalia and Jacod, 2009; Andersen et al., 2007b; Barndorff-Nielsen and Shephard, 2006) use high frequency data to estimate ex post realized jumps. This stream of the literature primarily focuses on issues such as why asset prices jump (e.g. macroeconomic new announcements) or how often asset prices jump (e.g. less than one per day). However, only a very few studies consider economic applications of realized jumps. We therefore aim to fill this gap between two related but different streams of literature by considering economic applications of realized jumps.

We focus on two research questions: Firstly, we are interested in whether realized jumps can forecast future volatility. We apply seven main stream nonparametric jump tests to identify realized jumps, decompose realized variance into jump and diffusion components, and then adapt them into a forecasting framework. Our findings suggest that realized jumps do contain predictive information for future volatility for the majority of jump tests both in-sample and out-of-sample. We find that jump models in general generate higher adjusted R squares and lower Mean Squared Prediction Errors (MSPE) compared to the benchmark model, which does not separate jumps from diffusion. Results hold true across the majority of jump specifications, and different forecasting horizons. Existing studies investigate similar issues. However, they mainly rely on one particular jump test and their results are mixed. For example, Andersen et al. (2007a) find a negative (but insignificant) relationship between jumps and one period ahead volatility. Corsi et al. (2010) on the contrary show statistically significant evidence to support a positive relationship if a modified jump test is applied. By using all major jump tests, our results contribute to the debate whether in general realized jumps help to forecast volatility.

Incorporating realized jumps into volatility forecasting require accessing intraday high frequency data and applying sophisticated nonparametric jump tests. Therefore, a natural question arises whether it is worth to estimate and use realized jumps. Even though separating jumps from diffusion improves volatility forecasting, it is interesting to know whether the improvement is large, and more importantly whether the improvement is economically valuable. Therefore, our second research question explicitly asks whether the potential statistical forecasting improvement obtained by separating jumps from diffusion can be translated into tangible economic benefits for a risk-averse investor. We construct a mean-variance portfolio strategy based on the predicted volatility obtained from the previous step. Our findings suggest that the statistical improvements are also economically significant. Under different risk aversion levels and jump specifications, jump strategies can in general generate positive and statistically significant performance fees relative to the benchmark strategy. A few existing papers also consider the role of jumps in asset allocations. For example, Liu et al. (2003) provide an analytical solution to the optimal portfolio choice problem when event risk or jumps are considered. They find that jumps play an important role in determining the optimal portfolio choice. Two recent studies by Chen et al. (2010) and Maheu et al. (2012) are also close to us in considering jumps in asset allocation. However, we differ from those studies on a few aspects. Firstly, our nonparametric framework enables us to exploit the information embedded in jump variations in a model free fashion while previous papers rely on a parametric specification. Secondly, we use high frequency data to separate the jumps and the diffusion component precisely, while they mainly rely on daily data to obtain relatively noisy proxies for jumps (i.e. large extraordinary movement or middle size jumps etc). The high frequency data we use also allows us to access intraday information, which is overlooked by previous studies.

We then conduct comprehensive robustness checks, and find that further controlling for market microstructure effects and transaction costs does not change our main results. We also investigate the predictive ability of realized jumps on alternative realized moments. We find that realized jumps can predict realized volatility and its signed components, but can hardly predict realized higher moments. We further show that a mean-variance portfolio strategy based on predicting positive and negative conditional volatilities separately can outperform the benchmark strategy based on predicting total volatility, and incorporating realized jumps can additionally improve economic benefits.

A few other studies are also related to ours. Firstly, our paper can be viewed as a natural extension of the stream of literature considering the economic value of volatility timing. Previous studies (Bandi and Russell, 2006; Bandi et al., 2008; Fleming et al., 2003; Liu, 2009) already document that volatility timing performance can be improved by using high frequency data, optimal sampling, and optimal rebalancing frequencies. We extend the above studies by considering realized jumps. Secondly, our paper is also related to other uses of realized jumps or applications of jump tests. For example, Dumitru and Urga (2012) and Theodosiou and Zikes (2011) conduct comprehensive simulation studies to compare size and power of jump tests. Tauchen and Zhou (2011) and Jiang and Yao (2013) use detected realized jumps to predict bond risk premia and the cross-section of stock returns respectively. We distinguish ours from previous studies by focusing on the role of realized jumps in volatility timing.

The rest of the paper is structured as follows: Section 2 discusses the theoretical setup and the jump tests. Section 3 describes the data and methodology. Section 4 discusses empirical findings from both statistical and economic perspectives. Section 5 conducts comprehensive robustness checks. Section 6 concludes.

2. Jumps in asset prices

2.1. Theoretical setup

Let p_t denote the logarithmic price which follows a jump diffusion process given by

$$dp_t = \mu_t dt + \sigma_t dW_t + dJ_t \quad (1)$$