Journal of Economic Dynamics & Control 37 (2013) 1872–1888

Contents lists available at SciVerse ScienceDirect

Journal of Economic Dynamics & Control

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

Asymmetry in the jump-size distribution of the S&P 500: Evidence from equity and option markets

Andreas Kaeck *

Swiss Institute of Banking and Finance, University of St. Gallen, Rosenbergstrasse 52, CH-9000 St. Gallen, Switzerland

ARTICLE INFO

Article history:
Received 5 March 2012
Received in revised form
19 November 2012
Accepted 5 April 2013
Available online 23 April 2013

JEL classification:
G13
C13
C63

Keywords:
Jump-size distribution
European options
S&P 500
Model calibration

ABSTRACT

This paper studies alternative distributions for the size of price jumps in the S&P 500 index. We introduce a range of new jump-diffusion models and extend popular double-jump specifications that have become ubiquitous in the finance literature. The dynamic properties of these models are tested on both a long time series of S&P 500 returns and a large sample of European vanilla option prices. We discuss the in- and out-of-sample option pricing performance and provide detailed evidence of jump risk premia. Models with double-gamma jump size distributions are found to outperform benchmark models with normally distributed jump sizes.

The statistical properties of stock prices and equity indices have attracted researchers at least since the influential work of Louis Bachelier.[1] Accurate models for stock price dynamics are an essential ingredient for many applications in finance such as risk management and option pricing. It is now widely accepted that equity volatility is stochastic (an incomplete list of influential work includes Heston, 1993; Hull and White, 1987; Melino and Turnbull, 1990; Stein and Stein, 1991; Schöbel and Zhu, 1999; Bates, 2000; Chernov et al., 2003 or Christoffersen et al., 2009b). To adequately model the extreme return outliers in equity markets, more recent theoretical and empirical research also argues in favor of jumps in the price process. The stock market crash of 1987, for instance, caused a −23% return in the S&P 500 index on a single day. Using the insight that it is unlikely that such extreme return is generated by a simple diffusion model, Bates (1996) proposes to combine both stochastic volatility and jumps and Duffie et al. (2000) augment this model by additional discontinuities in the stochastic variance process.[2] Jump-diffusion models are ubiquitous in the finance literature and the specifications of Bates (1996) and Duffie et al. (2000) are now established benchmarks for continuous-time equity index dynamics.

This paper concentrates on Poisson-driven jump-diffusion models and investigates the empirical performance of alternative jump-size distributions for the S&P 500 index. Our main motivation for studying this distribution in more detail is that benchmark models with normally distributed return jumps still exhibit a significant degree of misspecification. For

* Tel.: +41 71 224 7029; fax: +41 71 224 7088.
E-mail address: andreas.kaeck@unisg.ch

[1] An English translation of his original work (Bachelier, 1900) can be found in Cootner (2001).

[2] The inclusion of variance jumps is often motivated by the observation that the time series of at-the-money implied volatilities exhibits occasional spikes that are not in line with a simple diffusion model. It is argued that the volatility process requires an additional jump component in order to match these empirical observations. Alternatively, one could build models from alternative variance dynamics such as in Jones (2003) or Christoffersen et al. (2010b) but closed-form option prices are not available for most of these specifications.

0165-1889/$ - see front matter
http://dx.doi.org/10.1016/j.jedc.2013.04.008

typical parameter estimates presented in the literature (for example Eraker et al., 2003 estimate an average return jump size of −2.59% with a standard deviation of 4.07%), a return outlier of the size of the stock market crash of 1987 is even inconceivable after allowing for normally distributed jumps in the return process. To generate a similar crash scenario, it would require both an unrealistically high volatility level prior to the jump event and in addition a highly unlikely draw from the jump-size distribution. Indeed, our subsequent analysis reveals that models with normally distributed jump sizes fail to portray extreme tail events as well as other facets of the data. Despite these empirical deficiencies,[3] the models of Bates (1996) and Duffie et al. (2000) feature a wide range of useful characteristics and the goal of this paper is to provide extensive empirical tests whether alternative jump-size specifications can improve the modeling of S&P 500 index returns. Recent non-parametric analysis in Bollerslev and Todorov (2011) suggests that non-standard jump models might be useful to model S&P 500 index dynamics. We show that extensions of the benchmark models are able to explain tail events more easily and they also outperform in in- and out-of-sample option pricing exercises.

Our choice of alternative specifications is built on the ideas of Kou (2002) who employs a double-exponential distribution to separate positive and negative jumps in the price process. Kou assumes that there is no volatility clustering and hence his model produces i.i.d. returns. We generalize this approach in two directions: first, we consider more general jump-size distributions by using a double-gamma, a double-exponential and a Laplace distribution; and second, we add a stochastic variance process which is augmented with exponentially distributed upward jumps. Although models with alternative jump-size distributions are not new in the literature, to date, they have been very rarely considered in a stochastic volatility framework.[4] To the best of our knowledge, we are also the first to extensively test these specifications in combination with an additional variance jump feature. We believe that it is extremely important to gage jump-size distributions based on a very general stochastic volatility model, as otherwise one cannot be sure whether empirical results are driven by the omission of the variance process, omitted jumps in variance or the superiority of a particular jump distribution.

The models introduced in this paper can produce a wide range of shapes for the price-jump distribution. First, we allow for the possibility to create positive and negative jumps separately, a feature that can enhance the modeling of the tails of the S&P 500 return distribution. Since the normal distribution has to cover positive and negative jumps, it struggles to adapt to the peculiarities of both types of outliers. For example, the parameter estimates reported above do not only underestimate the likelihood of an extreme negative jump, they also very rarely produce significant positive jumps. This effect is even compounded in option pricing applications as the addition of a jump-size risk premium would leave virtually no probability mass for positive tail events. A second interesting feature of the jump-size distributions introduced in this paper is that some variants of the double-gamma model produce bimodal density functions. This might be important because the normal distribution assigns – in an attempt to cover both positive and negative price jumps – a high probability to jump sizes near zero, values that could be easily generated by its stochastic volatility component. While frequent jumps around zero do not make a model unrealistic, jump events in a jump-diffusion model are usually introduced to produce extreme outliers. Frequent small jumps might pose a problem for estimation and filtering algorithms because such jumps are very unlikely to be detected (see Johannes et al., 2009).

Although applications of Poisson-jump models have recently attracted a considerable amount of theoretical and empirical research, there are several attractive alternatives to continuous-time jump-diffusion models. Among these are jump extensions of popular discrete-time GARCH models (see Duan et al., 2006; Christoffersen et al., 2009a or Christoffersen et al., 2010a) or continuous-time models built from general Lévy processes. The latter class involves the Lévy jump extensions of the Heston model considered in Li et al. (2008), infinite-activity Lévy processes (see for example Carr et al., 2002), or the time-changed Lévy processes introduced in Carr and Wu (2004). Loosely speaking, unlike jump-diffusion models, Lévy processes allow the possibility of an infinite number of jumps during every finite time interval. Empirical research using Lévy processes is still scarce compared to the amount of research devoted to Poisson-jump models. Recent empirical work in this area includes Huang and Wu (2004) or Bates (2012). In this study, we deliberately choose to extend the benchmark models of Bates (1996) and Duffie et al. (2000) as these have become the most widely used specifications in the finance literature. Applications include option pricing, asset allocation, equilibrium modeling, density prediction and many more. Therefore, it is crucial to understand if simple extensions of these well-studied specifications can further improve the modeling of the S&P 500 index.

There are two popular data sources to study the dynamic behavior of equity indices. Classical time-series estimation is based on a long sample of S&P 500 returns, an approach used in Andersen et al. (2002) or Eraker et al. (2003). This procedure results in estimates of the processes under the statistical probability measure. Alternatively, one can study equity dynamics using the information in a rich set of option quotes to infer structural parameters under some risk-neutral pricing measure (see for example Bakshi et al., 1997). In this paper we aim to study alternative specifications using both data sources and we adopt an estimation procedure similar to Broadie et al. (2007). We first study the time-series implications of our jump-diffusion models using Markov Chain Monte Carlo (MCMC) methods and gage their ability to explain the observed daily S&P 500 time series over a long time period from 1950 to 2010.[5] In a second step, we fix the parameters that are theoretically identical under both measures and study option pricing implications. This way, we obtain risk-neutral

[3] For further evidence see for example Bates (2006) or Li et al. (2008).

[4] A version of the double exponential jump model (without a variance jump component) was considered in Chacko and Viceira (2003). The authors provide parameter estimates for different observation frequencies, but do not test this model against alternative specifications.

[5] Such a long time series is necessary to obtain as many jump events as possible in order to estimate the jump-size distributions accurately.