Journal of Financial Economics 105 (2012) 229–259

Contents lists available at SciVerse ScienceDirect

Journal of Financial Economics

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

U.S. stock market crash risk, 1926–2010[☆]

David S. Bates[*]

University of Iowa and the National Bureau of Economic Research, Henry B. Tippie College of Business, Iowa City, IA 52242-1000, USA

ARTICLE INFO

Article history:
Received 27 January 2010
Received in revised form
31 March 2011
Accepted 25 April 2011
Available online 9 March 2012

JEL classification:
C22
C46
G01
G13

Keywords:
Lévy processes
Time-changed Lévy processes
Stock market crashes
Option pricing

ABSTRACT

This paper examines how well alternate time-changed Lévy processes capture stochastic volatility and the substantial outliers observed in U.S. stock market returns over the past 85 years. The autocorrelation of daily stock market returns varies substantially over time, necessitating an additional state variable when analyzing historical data. I estimate various one- and two-factor stochastic volatility/Lévy models with time-varying autocorrelation via extensions of the Bates (2006) methodology that provide filtered daily estimates of volatility and autocorrelation. The paper explores option pricing implications, including for the Volatility Index (VIX) during the recent financial crisis.

1. Introduction

What is the risk of stock market crashes? Answering this question is complicated by two features of stock market returns: the fact that conditional volatility evolves over time, and the fat-tailed nature of daily stock market returns. Each issue affects the other. Which returns are identified as outliers depends upon that day's assessment of conditional volatility. Conversely, estimates of current volatility from past returns can be disproportionately affected by outliers such as the 1987 crash. In standard generalized autoregressive conditional heteroskedasticity (GARCH) specifications, for instance, a 10% daily change in the stock market has one hundred times the impact on conditional variance revisions of a more typical 1% move.

This paper explores whether recently proposed continuous-time specifications of time-changed Lévy processes are a useful way to capture the twin properties of stochastic volatility and fat tails. The use of Lévy processes to capture outliers dates back at least to the Mandelbrot (1963) use of the stable Paretian distribution, and many specifications have been proposed, including the Merton (1976) jump-diffusion, the Madan and Seneta (1990) variance gamma, the Eberlein, Keller, and Prause (1998) hyperbolic Lévy, and the Carr, German, Madan, and Yor (2002) CGMY process. As all of these distributions assume identically and independently distributed (i.i.d.)

[☆] I am grateful for comments on earlier versions of the paper from the referee, from seminar participants at the universities of Iowa, Turin, London (Birkbeck), Colorado, Calgary, and Massachusetts-Amherst, and from conference participants at McGill University's 2008 Risk Management Conference, the 18th Annual Derivatives Securities and Risk Management Conference, Princeton University's 2008 Conference on Implied Volatility Models, the 2009 Western Finance Association meetings, the 2009 Conference on Liquidity, Credit Risk, and Extreme Events at the University of Chicago, the 2010 Derivatives, Volatility, and Correlation Conference at Warwick University, and the 2010 Financial Econometrics Conference at the University of Toulouse. I also thank Ken French for providing the stock turnover data from French (2008).

[*] Tel.: +1 319 353 2288; fax: +1 319 335 3690.
E-mail address: david-bates@uiowa.edu
URL: http://www.biz.uiowa.edu/faculty/dbates

0304-405X/$ - see front matter © 2012 Elsevier B.V. All rights reserved.
doi:10.1016/j.jfineco.2012.03.004

returns, however, they are unable to capture stochastic volatility.

More recently, Carr, German, Madan, and Yor (2003) and Carr and Wu (2004) have proposed combining Lévy processes with a subordinated time process. The idea of randomizing time dates back at least to Clark (1973). Its appeal in conjunction with Lévy processes reflects the increasing focus in finance – especially in option pricing – on representing probability distributions by their associated characteristic functions. Lévy processes have log characteristic functions that are linear in time. If the time randomization depends on underlying variables that have an analytic conditional characteristic function, then the resulting conditional characteristic function of time-changed Lévy processes is also analytic. Conditional probability densities, distributions, and option prices can then be numerically computed by Fourier inversion of simple functional transforms of this characteristic function.

Thus far, empirical research on the relevance of time-changed Lévy processes for stock market returns has largely been limited to the special cases of time-changed versions of Brownian motion and the Merton (1976) jump-diffusion. Furthermore, there has been virtually no estimation of newly proposed time-changed Lévy processes solely from time series data.[1] Papers such as Carr, German, Madan, and Yor (2003) and Carr and Wu (2004) rely on option pricing evidence to provide empirical support for their approach, instead of providing direct time series evidence. The reliance on options data is understandable. Because the state variables driving the time randomization are not directly observable, time-changed Lévy processes are hidden Markov models, creating a challenging problem in time series econometrics. Using option prices potentially identifies realizations of those latent state variables, converting the estimation problem into the substantially more tractable problem of estimating state space models with observable state variables.

While options-influenced parameter and state variable estimates should be informative under the hypothesis of correct model specification, the objective of this paper is to provide estimates of crash risk based solely upon time series analysis. Such estimates are of interest in their own right, and are useful for testing the central empirical hypothesis in option pricing: whether option prices are, in fact, compatible with the underlying time series properties of the underlying asset, after appropriate risk adjustments. Testing the compatibility hypothesis is more difficult under joint options/time series estimation approaches that are premised upon compatibility. Furthermore, option-based and joint estimation approaches are constrained by the availability of options data only since the 1980s, whereas time series estimation can exploit a longer history of extreme stock market movements.[2] For instance, it has been asserted that deep out-of-the-money index put options appear overpriced, based on their surprisingly large negative returns since the 1987 crash. But all such tests require reliable estimates of downside risk; and it can be difficult to establish whether puts are overpriced based only on post-1987 data.[3]

Risk-adjusted time series estimates of conditional distributions can also provide useful real-time valuations of option prices, for comparison with observed option prices. At the end of the paper I compare the options-based Volatility Index (VIX) measure of volatility with time series estimates, during a 2007–2010 period spanning the recent financial crisis.

This paper uses the Bates (2006) approximate maximum likelihood (AML) methodology for estimation of various time-changed Lévy processes over 1926–2006, and for out-of-sample fits over 2007–2010. AML is a filtration methodology that recursively updates conditional characteristic functions of latent variables over time given observed data. Filtered estimates of the latent variables are directly provided as a by-product, given the close link between moments and characteristic functions. The methodology's focus on characteristic functions makes it especially useful for estimating Lévy processes, which typically lack closed-form probability density functions. The paper primarily focuses on the time-changed CGMY process, which nests other Lévy processes as special cases. The approach is also compared with the stochastic volatility processes with and without normally distributed jumps previously estimated in Bates (2006).

A concern with any extended data set is the possibility that the data generating process might not be stable over time. Indeed, this paper identifies substantial instability in the autocorrelation of daily stock market returns. Autocorrelation estimates appear to be nonstationary, and peaked at the extraordinarily high level of 35% in 1971 before trending downward to the near-zero values observed since the 1980s. The instability is addressed directly, by treating autocorrelation as another latent state variable to be estimated from observed stock market returns. The paper also uses subsample estimation to test for (and find) apparent instabilities or specification issues in the one-factor volatility process used. Given these issues, I estimate a two-factor concatenated model of volatility evolution, which can be interpreted as a model of parameter drift in the unconditional mean of the one-factor variance process. Finally, I examine the sensitivity of volatility filtration and option prices to the use of different data sets and volatility models.

Overall, the time-changed CGMY process is found to be a slightly more parsimonious alternative to the Bates (2006) approach of using finite-activity stochastic-intensity jumps

---

[1] Li, Wells, and Yu (2008) use Markov chain Monte Carlo (MCMC) methods to estimate some models in which Lévy shocks are added to various stochastic volatility models. However, the additional Lévy shocks are independently and identically distributed, not time-changed.

[2] The -20% and $+11½$ movements on October 19, 1987 and October 13, 2008, respectively, were the only daily stock market movement over 1945–2010 in the Center for Research in Security Prices (CRSP) value-weighted index to exceed 10% in magnitude, whereas there were seven such movements over 1929–1932.

[3] See Broadie, Chernov, and Johannes (2009) for a Monte Carlo study of unhedged one-month returns for puts on S&P 500 futures over August 1987 to June 2005. They find that their large excess return estimates often lack statistical significance, especially when volatility is stochastic.