Journal of Financial Economics 106 (2012) 447–472

Contents lists available at SciVerse ScienceDirect

Journal of Financial Economics

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

Dynamic jump intensities and risk premiums: Evidence from S&P500 returns and options ☆

Peter Christoffersen [a,b], Kris Jacobs [c,d,*], Chayawat Ornthanalai [a]

[a] Rotman School of Management, University of Toronto, Canada
[b] Copenhagen Business School and CREATES, Denmark
[c] C.T. Bauer College of Business, 334 Melcher Hall, University of Houston, Houston, TX 77204-6021, United States
[d] Tilburg University, The Netherlands

ARTICLE INFO

Article history:
Received 13 June 2010
Received in revised form
23 September 2011
Accepted 27 November 2011
Available online 7 June 2012

JEL classification:
G13

Keywords:
Compound Poisson jumps
Analytical filtering
Fat tails
Risk premiums

ABSTRACT

We build a new class of discrete-time models that are relatively easy to estimate using returns and/or options. The distribution of returns is driven by two factors: dynamic volatility and dynamic jump intensity. Each factor has its own risk premium. The models significantly outperform standard models without jumps when estimated on S&P500 returns. We find very strong support for time-varying jump intensities. Compared to the risk premium on dynamic volatility, the risk premium on the dynamic jump intensity has a much larger impact on option prices. We confirm these findings using joint estimation on returns and large option samples.

1. Introduction

This paper provides a new modeling framework that allows for general specifications of dynamic volatility models with fat-tailed innovations that are potentially time-varying themselves. These models are much easier to implement than existing models. As suggested in Fleming and Kirby (2003), we directly use GARCH processes as filters for the unobservable state variables. We allow for dynamic fat tails in the standard GARCH model by adding a dynamic Compound Poisson jump component to the conventional normal innovation. Our framework allows for dynamic jumps in returns as well as in volatility. The resulting specifications of complex fat-tailed models can be estimated on return data using a standard maximum likelihood estimation procedure and an analytical filtering technique to identify the two return innovations. Estimation using extensive option data sets is also feasible. Our approach allows for separate risk premiums for the fat-tailed jump and normal innovations. We provide the risk-neutral processes for use in option valuation, and conduct an extensive empirical investigation of return and option fit.

Because implementation of our models is relatively straightforward, we are able to estimate models with very complex tail characteristics. We consider four different

☆ Christoffersen and Jacobs are grateful for financial support from FQRSC, IFM2 and SSHRC. Ornthanalai was supported by CIREQ and J.W. McConnell Fellowships. Part of this research was done while Ornthanalai was at the Georgia Institute of Technology. For helpful comments we would like to thank Pedro Santa-Clara (the referee), Giovanni Barone-Adesi, David Bates, Tom McCurdy, Anders Trolle, participants at the AFA, the EFA, the NFA, the EFMA, the FDIC Derivatives Conference, CESG, and seminar participants at the Federal Reserve Board, the Atlanta Fed, CU Boulder, HKUST, the Helsinki School of Economics, York, INSEAD, and Rotman. Any remaining inadequacies are ours alone.

* Corresponding author at: C.T. Bauer College of Business, 334 Melcher Hall, University of Houston, Houston, TX 77204-6021, United States. Tel.: +1 713 743 2826; fax: +1 713 743 4622.

E-mail address: kjacobs@bauer.uh.edu (K. Jacobs).

0304-405X/$ - see front matter
http://dx.doi.org/10.1016/j.jfineco.2012.05.017

but nested models. The simplest model, labeled DVCJ (dynamic volatility with constant jumps), has a constant jump intensity. The CVDJ (constant volatility with dynamic jumps) model has time-varying jump intensity, but the normal innovation to the return process is assumed to be homoskedastic. The DVSDJ (dynamic volatility with separate dynamic jumps) model is the most general model we investigate: both the jump intensity and the normal innovations are time-varying, and the dynamics are parameterized separately. The DVDJ (dynamic volatility with dynamic jumps) model is a special case of DVSDJ: both the jumps and the normal innovation are time-varying, but parameterized identically. We first estimate these models using 47 years of daily returns. After estimating the models, we risk-neutralize the parameters and compare their option valuation performance using 13 years of index option data. Subsequently, we estimate the models using returns and option data jointly.

Our empirical results emphasize the importance of time-varying jump intensities to generate dynamic fat tails, fit returns, and value options. The index return data reveal strong support for the DVSDJ and DVDJ models. When using the risk-neutralized estimates from returns to value options, the DVSDJ model yields a 45% improvement in the option-implied volatility root mean squared errors over the standard GARCH model. The DVDJ model also performs well. It yields a 34% improvement in the option-implied volatility root mean squared errors over the GARCH model. When estimating the models using return data and option data jointly, the DVSDJ model again performs best, and the DVDJ model also performs well.

Our findings therefore indicate the benefit of specifications with time-varying jump intensities and dynamic fat tails, with the more richly parameterized model performing relatively better, in-sample as well as out-of-sample. Our results also indicate that fat-tailed models provide superior option pricing performance, particularly during medium and high volatility periods. The contribution of the jump component to total equity volatility varies with the specification of the jump innovation. The more flexible the jump innovation, the more important it is, and the larger is the contribution of the jump component to total equity volatility.

We also investigate the importance of the risk premium on the jump innovation. We conclude that to produce significant improvements in option valuation, models must allow for a risk premium on the jump innovation. We investigate if risk premiums can generate plausible shapes and levels of the implied volatility term structure, and we find that the implied volatility term structure is highly responsive to the jump risk premium. On the other hand, unrealistically large risk premiums for the normal innovation are required to generate levels and slopes of implied volatility found in the data.

Our approach is closely related to the discrete-time jump approach in Maheu and McCurdy (2004), who find strong evidence of time-varying jumps in various individual equity and index returns. We find similar evidence using S&P500 equity index returns, and importantly, we provide theory and empirics on option valuation using our fat-tailed models. Hansen (1994) develops a class of dynamic fat-tailed density models that has subsequently been generalized by Jondeau and Rockinger (2003). However, neither study considered option valuation. Duan, Ritchken, and Sun (2006) provide a risk-neutralization of a discrete-time model with jumps, but they do not allow for time-varying jump intensities and higher-moment dynamics. Option valuation using GARCH models with normal innovations was initially developed in Duan (1995), and explored further in Ritchken and Trevor (1999) and Heston and Nandi (2000).

Our results are also closely related to the existing literature on continuous-time stochastic volatility jump-diffusions (SVJ). The DVCJ model has features similar to the stochastic volatility with correlated jumps in returns and volatility (SVCJ) model underlying most existing empirical estimates. Our DVDJ and DVSDJ models and estimation results are related to the most complex dynamics investigated in the continuous-time literature, as in Santa-Clara and Yan (2010) and Eraker (2004), who specify time-varying intensities, and estimate their models using returns and options jointly.[1]

The remainder of the paper proceeds as follows. Section 2 presents our modeling approach and discusses the four nested specifications. Section 3 provides empirical results from estimating the models on daily returns. Section 4 develops the theoretical framework for risk-neutralization and option valuation. Section 5 provides the empirical results on option valuation using parameters estimated on returns, and Section 6 estimates the models using a joint likelihood composed of returns and options data. Section 7 concludes.

2. Daily returns with jump dynamics

In this section we present the proposed asset return process, including the structure of the jump innovation and the dynamics for volatility and jump intensity.

2.1. The return process

The process contains two components. The first we model using a normal innovation and the second is a jump component. Here, we discuss some aspects of the general structure of the jump and normal components. The variance and jump intensity dynamics are discussed in more detail in Section 2.5.

The return process is given by

$$R_{t+1} \equiv \log\left(\frac{S_{t+1}}{S_t}\right) = r + \left(\lambda_z - \frac{1}{2}\right) h_{z,t+1} + (\lambda_y - \xi) h_{y,t+1} + z_{t+1} + y_{t+1}, \tag{1}$$

[1] For other empirical estimates of jump processes in the SVJ literature using returns and/or options, see, for example, Bates (2000, 2006), Andersen, Benzoni, and Lund (2002), Pan (2002), Huang and Wu (2004), Eraker, Johannes, and Polson (2003), Broadie, Chernov, and Johannes (2007), Li, Wells, and Yu (2007, 2011), and Chernov, Gallant, Ghysels, and Tauchen (2003). Bollerslev and Todorov (2011) demonstrate the need for a time-varying risk premium with a model-free investigation.