Journal of Econometrics 172 (2013) 307–324

Contents lists available at SciVerse ScienceDirect

Journal of Econometrics

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

Jump tails, extreme dependencies, and the distribution of stock returns[☆]

Tim Bollerslev[a,b,c,1], Viktor Todorov[d,*], Sophia Zhengzi Li[a]

[a] Department of Economics, Duke University, Durham, NC 27708, United States
[b] NBER, United States
[c] CREATES, United States
[d] Department of Finance, Kellogg School of Management, Northwestern University, Evanston, IL 60208, United States

ARTICLE INFO

Article history:
Available online 23 August 2012

JEL classification:
C13
C14
G10
G12

Keywords:
Extreme events
Jumps
High-frequency data
Jump tails
Non-parametric estimation
Stochastic volatility
Systematic risks
Tail dependence

ABSTRACT

We provide a new framework for estimating the systematic and idiosyncratic jump tail risks in financial asset prices. Our estimates are based on in-fill asymptotics for directly identifying the jumps, together with Extreme Value Theory (EVT) approximations and methods-of-moments for assessing the tail decay parameters and tail dependencies. On implementing the procedures with a panel of intraday prices for a large cross-section of individual stocks and the S&P 500 market portfolio, we find that the distributions of the systematic and idiosyncratic jumps are both generally heavy-tailed and close to symmetric, and show how the jump tail dependencies deduced from the high-frequency data together with the day-to-day variation in the diffusive volatility account for the "extreme" joint dependencies observed at the daily level.

1. Introduction

Tail events and non-normal distributions are ubiquitous in finance. The earliest comprehensive empirical evidence for fat-tailed marginal return distributions dates back more than half a century to the influential work of Mandelbrot (1963) and Fama (1965). It is now well recognized that the fat-tailed unconditional return distributions first documented in these, and numerous subsequent studies may result from time-varying volatility and/or jumps in the underlying stochastic process governing the asset price dynamics. Intuitively, periods of high-volatility can result in seemingly "extreme" price changes, even though the returns are drawn from a normal distribution with light tails, but one with an unusually large variance; see e.g., Bollerslev (1987), Mikosch and Starica (2000), and the empirical analyses in Kearns and Pagan (1997) and Wagner and Marsh (2005) pertaining to the estimation of tail parameters in the presence of GARCH effects. On the other hand, the aggregation of multiple jump events over a fixed time interval will similarly result in fat-tailed asset return distributions, even for a pure Lévy-type jump processes with no dynamic dependencies; see, e.g., Carr et al. (2002). As such, while fundamentally different, these two separate mechanisms will both manifest themselves in the form of apparent "tail" events and leptokurtic marginal return distributions.[2]

These same general issues carry over to a multivariate context and questions related to "extreme" dependencies across assets. In particular, it is well documented that the correlations between equity returns, both domestically and internationally, tend to be higher during sharp market declines than during "normal" periods[3]; see e.g., Longin and Solnik (2001) and Ang and Chen (2002).

☆ We would like to thank the Editor (Marc Paolella) and two anonymous referees for many helpful comments and suggestions. The research was supported by a grant from the NSF to the NBER, and CREATES funded by the Danish National Research Foundation (Bollerslev).

* Corresponding author.
E-mail addresses: boller@duke.edu (T. Bollerslev), v-todorov@northwestern.edu (V. Todorov), zhengzi.li@duke.edu (S.Z. Li).
1 Tel.: +1 919 660 1846; fax: +1 919 684 8974.

2 Importantly, these different mechanisms also have very different pricing implications and risk premia dynamics, as recently explored by Bollerslev and Todorov (2011b).

3 The use of simple linear correlations as a measure of dependence for "extreme" observations has been called into question by Embrechts et al. (2002), among others.

0304-4076/$ – see front matter
doi:10.1016/j.jeconom.2012.08.014

Similarly, Starica (1999) documents much stronger dependencies for large currency moves compared to "normal-sized" changes, while Jondeau (2010) based on an explicit parametric model reports much stronger tail dependence on the downside for several different equity portfolios.

In parallel to the marginal effects discussed above, it is generally unclear whether these increased dependencies in the tails are coming from commonalities in time-varying volatilities across assets and/or common jumps. Poon et al. (2004), for instance, report that "devolatilizing" the daily returns for a set of international stock markets significantly reduces the joint tail dependence, while Bae et al. (2003) find that time-varying volatility and GARCH effects cannot fully explain the counts of coincident "extreme" daily price moves observed across international equity markets. More closely related to the present paper, recent studies by Bollerslev et al. (2008), Jacod and Todorov (2009), and Gobbi and Mancini (2009), based on high-frequency data and nonparametric methods, have all argued for the presence of common jump arrivals across different assets, thus possibly inducing stronger dependencies in the "extreme".

In light of these observations, one of the goals of the present paper is to separate jumps from volatility to more directly assess the "extreme" dependencies inherent in the jump tails. Motivated by the basic idea from asset pricing finance that only non-diversifiable systematic jump risks should be compensated, we further dissect the jumps into their systematic and idiosyncratic components. This decomposition in turn allows us to compare and contrast the behavior of the two different jump tails and how they impact the return distributions.[4]

Our estimation methodology is based on the idea that even though jumps and time-varying volatility may have similar implications for the distribution of the returns over coarser sampling frequencies, the two features manifest themselves very differently in high-frequency returns. Intuitively, treating the volatility as locally constant over short time horizons, it is possible to perfectly separate jumps from the price moves associated with the slower temporally varying volatility through the use of increasingly finer sampled observations. Empirically, this allows us to focus directly on the high-frequency "filtered" jumps. Relying on the insight from Bollerslev and Todorov (2011a) that regardless of *any* temporal variation in the jump intensity, the jump compensator for the "large" jumps behaves like a probability measure, we non-parametrically estimate the decay parameters for the univariate jump tails using a variant of the Peaks-Over-Threshold (POT) method.[5]

Going one step further, we characterize the extreme joint behavior of the "filtered" jump tails through non-parametric estimates of Pickands (1981) dependence function as well as the residual tail dependence coefficient of Ledford and Tawn (1996, 1997). The Pickands dependence function succinctly characterizes the dependence of the limiting bivariate extreme value distribution. When the latter has independent marginals, the residual tail dependence coefficient further discriminates among the dependencies that disappear in the limit.[6] We implement several different estimators for the Pickands dependence function and the residual tail dependence coefficient. Together with the estimated decay parameters for each of the underlying univariate extreme distributions, these summary measures effectively describe the key features of the bivariate joint tail behavior.[7]

Our actual empirical analysis is based on high-frequency observations for fifty large capitalization stocks and the S&P 500 aggregate market portfolio spanning the period from 1997 through 2010. We find that the number of "filtered" idiosyncratic jumps exceeds the number of systematic jumps for all of the stocks in the sample, and typically by quite a large margin. Nonetheless, the hypothesis of fully diversifiable individual jump risk is clearly not supported by the data, thus pointing to more complicated dependence structures in the tails than hitherto entertained in most of the existing asset pricing literature.[8]

Even though the assumption of "light" Gaussian jump tails cannot necessarily be rejected for many of the individual estimates, the combined evidence for all of the stocks clearly supports the hypothesis of heavy jump tails. Our estimates for the individual jump tail decay parameters also suggest that the tails associated with the systematic jumps are slightly fatter than those for the idiosyncratic jumps, albeit not uniformly so. Somewhat surprisingly, we also find that the right tail decay parameters for both types of jumps in quite a few cases exceed those for the left tail.

Our estimates of various dependence measures reveal a strong degree of tail dependence between the market-wide jumps and the systematic jumps in the individual stocks. This therefore calls into question the assumption of normally distributed jumps previously used in the asset and derivatives pricing literature.

Further, comparing our high-frequency based estimation results with those obtained from daily returns, we find that the latter indicate much weaker tail dependencies. Intuitively, while the estimates based on the daily returns represent the tail dependence attributable to both systematic jumps and common volatility factors, both of which may naturally be expected to be associated with positive dependence, the idiosyncratic jumps when aggregated over time will tend to weaken the dependence. In contrast, by focusing directly on the high-frequency "filtered" systematic and idiosyncratic jumps, we are able to much more accurately assess the true extreme jump tail dependencies, and assess how the different effects impart the dependencies in the lower-frequency daily returns.

The rest of the paper is organized as follows. Section 2 introduces the formal setup and assumptions. Section 3 outlines the statistical methodology and econometric procedures, beginning in Section 3.1 with the way in which we disentangle jumps from continuous prices moves, followed by a discussion of our univariate tail estimation procedures in Section 3.2, and the framework that we rely on for assessing the joint jump tail dependencies in Section 3.3. Section 4 presents the results from an extensive Monte Carlo simulation study designed to assess the properties of the different estimators in an empirically realistic setting. Section 5 summarizes our main empirical results, starting in Section 5.1 with a brief description of the data, followed by our findings pertaining to the individual jump tails in Section 5.2, and the bivariate jump tail dependencies in Section 5.3. Section 6 concludes.

4 In a related context, Barigozzi et al. (2010) have recently explored a factor structure for disentangling the total realized variation for a large panel of stocks into a single systematic component and remaining idiosyncratic components, while Todorov and Bollerslev (2010) propose a framework for the estimation of separate continuous and jump CAPM betas.

5 The POT method for characterizing extremes dates back to Fisher and Tippett (1928). It has been formalized more recently by Balkema and deHaan (1974) and Pickands (1975); for general textbook discussions see also Embrechts et al. (2001) and Jondeau et al. (2007).

6 In technical terms, the Pickands dependence function captures asymptotic tail dependence, while the residual tail dependence coefficient captures pre-asymptotic tail dependence.

7 For a general textbook discussion of the relevant concepts, see, e.g., Coles (2001) and Beirlant et al. (2004). Existing applications of these ideas have primarily been restricted to climatology and insurance. Steinkohl et al. (2010), for instance, have recently employed this approach to characterize the asymptotic dependence for high-frequency wind speeds across separate geographical locations.

8 The mere existence of market-wide jumps, of course, refutes the hypothesis of fully diversifiable jump risk as in Merton (1976). The estimates reported in, e.g., Eraker et al. (2003), also suggest large risk premia for systematic jump risk.