ARTICLE IN PRESS

Journal of Econometrics ▮ (▮▮▮▮) ▮▮▮–▮▮▮

Contents lists available at ScienceDirect

Journal of Econometrics

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

Time-varying jump tails[☆]

Tim Bollerslev[a,b,c,*], Viktor Todorov[d]

[a] Department of Economics, Duke University, Durham, NC 27708, United States

[b] NBER, United States

[c] CREATES, Denmark

[d] Department of Finance, Kellogg School of Management, Northwestern University, Evanston, IL 60208, United States

ARTICLE INFO

Article history:
Available online xxxx

JEL classification:
C13
C14
G10
G12

Keywords:
Market risk
Options
Risk-neutral distributions
Jumps
Time-varying jump tails
Extreme value theory

ABSTRACT

We develop new methods for the estimation of time-varying risk-neutral jump tails in asset returns. In contrast to existing procedures based on tightly parameterized models, our approach imposes much fewer structural assumptions, relying on extreme-value theory approximations together with short-maturity options. The new estimation approach explicitly allows the parameters characterizing the shape of the right and the left tails to differ, and importantly for the tail shape parameters to change over time. On implementing the procedures with a panel of S&P 500 options, our estimates clearly suggest the existence of highly statistically significant temporal variation in both of the tails. We further relate this temporal variation in the shape and the magnitude of the jump tails to the underlying return variation through the formulation of simple time series models for the tail parameters.

1. Introduction

Financial asset returns are not conditionally normally distributed, but instead exhibit more slowly decaying, and often asymmetric, tails. This is true even over short horizons, as most easily seen from the presence of very pronounced volatility smiles for short maturity options.[1] These fatter than normal tails are directly attributable to occasionally large absolute price changes, or "jumps". The recent financial crises has further underscored the empirical relevance of tail events, and in turn econometric techniques for more accurately estimating and modeling such risks. We add to this literature through the development of new more flexible estimation procedures that explicitly allow for the possibility of time-varying tails for the large jump moves. In comparison to the existing literature, our approach imposes much fewer structural assumption, relying on extreme-value theory approximations together with short-maturity S&P 500 options. By focusing on the risk-neutral distributions implied from options data, our estimates speak directly to the jump tail risk that is priced by the market. Consistent with the existing literature, we find that the magnitude of the left jump tail associated with dramatic market declines far exceeds that of the right jump tail corresponding to large market appreciations.[2] Our new estimation procedures also clearly point to the existence of non-trivial predictable temporal dependencies in the tail index parameters characterizing the decay in both tails.

A number of previous studies have argued that the values of the parameters for the power laws governing the tails of return distributions may be subject to structural changes; see, e.g., the studies by Quintos et al. (2001) and Galbraith and Zernov (2004) based on the traditional Hill-estimator and daily aggregate equity

☆ We would like to thank Alok Bhargava (the Editor) and two anonymous referees for their many helpful comments, and Lai Xu for excellent research assistance. The research was supported by a grant from the NSF SES-0957330 to the NBER, and CREATES funded by the Danish National Research Foundation (Bollerslev).

* Corresponding author at: Department of Economics, Duke University, Durham, NC 27708, United States.

E-mail addresses: boller@duke.edu (T. Bollerslev), v-todorov@northwestern.edu (V. Todorov).

1 The failure of the traditional Black–Scholes model and the presence of volatility smiles after the market crash of 1987 is well documented in the asset pricing literature. The impact of this failure for econometric analysis in a corporate finance setting related to executive compensation has recently been studied by Bhargava (2013).

2 The observation that the left tail inferred from aggregate equity index options dominates the right tail dates back at least to Rubinstein (1994), who attributed this to evidence of "crash-o-phobia".

http://dx.doi.org/10.1016/j.jeconom.2014.05.007
0304-4076/

index returns.[3] Relying on a large cross-section of stock returns, the more recent study by Kelly (2012) reinforces the idea of time-varying tail risks, and further argues that the temporal variation in the tail parameters may help understand aggregate market returns as well as cross-sectional differences in average returns. Similarly, the study by Bollerslev and Todorov (2011a) demonstrates how high-frequency intraday data may be used in more accurately estimating dynamically evolving tails, and how these estimates may be used in more effective risk measurement and management decisions.

All of these studies are based on directly observed return data, and in turn pertain to the objective, or statistical, return distributions. By contrast, the new estimation procedures developed here pertain to the risk-neutral distribution, explicitly reflecting the way in which the market perceives and prices tail risks. The method builds on the insight that out-of-the-money short-maturity options effectively isolate the pricing of jump risk. Formally, in the limit for decreasing times-to-maturity and fixed moneyness, the diffusive risk will not affect the price of an out-of-the-money option. Regular variation in the jump tail measure, or compensator, therefore implies a one-to-one mapping between the shape of the jump tail measure and the slope of the option price surface in the strike dimension. Consequently, the tail index parameter may be uniquely identified, and in turn estimated, from a cross-section of deep out-of-the-money short-maturity options at a given point in time without making any assumptions about the temporal variation in the overall jump intensity process.[4]

The basic idea of inferring the risk-neutral jump tails from options is related to an earlier literature that seek to better explain option prices through jump risk; see, e.g., Bates (1996), Bates (2000), Andersen et al. (2002), Pan (2002), Eraker (2004) and Broadie et al. (2007), along with the more recent work by Christoffersen et al. (2012). All of these studies are based on specific, typically affine, parametric stochastic volatility jump–diffusion models. Moreover, following Merton (1976), they postulate that conditionally on a jump occurring the size of the jump is normally distributed. Our approach is distinctly different in relying on a flexible nonparametric procedure that is able to accommodate complex dynamic tail dependencies and larger jump tails outside this classical Merton-framework.[5]

Our new estimation procedure is also related to the earlier work by Aït-Sahalia and Lo (1998), who non-parametrically estimate the entire risk-neutral state price density from options data. Their approach, however, explicitly assumes that the pricing kernel is time-invariant. On the other hand, Rosenberg and Engle (2002) do allow the pricing kernel to change over time, but rely on tightly parameterized GARCH type models for describing the dynamic dependencies. Alternatively, Metaxoglou and Smith (2012) resort to the use of conditional quantile regression techniques for estimating time-varying pricing kernels. The recent study by Song and Xiu (2013) also explicitly relates the temporal variation in the risk-neutral distribution and the pricing kernel to the VIX index and the volatility of the aggregate market. Meanwhile, none of these estimation procedures are directly geared to the tails of the distribution. By contrast, our approach explicitly focuses on the tails and the tail decay parameters, in particular, ignoring other parts of the distribution.

The current paper is related to our earlier work, Bollerslev and Todorov (2011b), in which short-maturity option data is used to estimate semiparametrically risk-neutral jump tails. From an econometric point of view, however, there are two fundamental differences. First, unlike Bollerslev and Todorov (2011b) we explicitly allow the shape of the jump tails to vary over time. Second, the estimation in the present paper is based on a fixed time span and the entire cross-section of short maturity deep out-of-the-money options, whereas the estimation in Bollerslev and Todorov (2011b) rely on long time span asymptotics and only a limited number of strikes.

At a more general level, our empirical results are also related to the recent literature by Barro (2006) and others emphasizing the importance of incorporating rare disasters in macro-finance models. The idea that rare disasters, or tail events, may help explain the equity premium and other empirical puzzles in asset pricing dates back at least to Rietz (1988). Further building on these ideas, Gabaix (2012) and Wachter (2013) have recently shown that allowing for time-varying tail risks in otherwise standard equilibrium based asset pricing models may help explain the apparent excess volatility of aggregate equity index returns. Similarly, Bollerslev and Todorov (2011b) and Aït-Sahalia et al. (2013) suggest that much of the variance risk premium is directly attributable to disaster, or jump tail risk.

The plan for the rest of the paper is as follows. We begin in the next section with a discussion of the basic setup and assumptions, including our very general time-varying jump tail formulation. Section 3 discusses how options may be used for effectively separating jumps and continuous price variation, and outlines our new estimation procedures for the jump tail parameters. Section 4 summarizes the S&P 500 options data that we use in the estimation. Our main empirical findings related to the tail risk parameters and the temporal variation therein are discussed in Section 5. Section 6 concludes. The proof of the key asymptotic approximation underlying our new estimation procedures is deferred to a technical Appendix.

2. Jump tails

The continuous-time no-arbitrage framework that underlies our new estimation procedure is very general. It includes all parametric models previously analyzed and estimated in the literature as special cases. We begin with a discussion of the basic setup and notation.

2.1. Setup and assumptions

The underlying asset price X_t is defined on the filtered probability space $(\Omega, \mathcal{F}, \mathbb{P})$, where $(\mathcal{F}_t)_{t\geq 0}$ denotes the filtration. We assume the following general dynamic specification for X_t,

$$\frac{dX_t}{X_{t-}} = \alpha_t dt + \sigma_t dW_t + \int_{\mathbb{R}} (e^x - 1)\widetilde{\mu}(dt, dx), \tag{2.1}$$

where W_t is a Brownian motion, μ is a counting measure for the jumps in X with compensator $dt \otimes \nu_t(dx)$, so that $\widetilde{\mu}(dt, dx) = \mu(dt, dx) - dt\,\nu_t(dx)$ denotes the corresponding martingale measure under $\mathbb{P}$.[6,7] The drift and volatility processes, α_t and σ_t,

[3] Fat tailed marginal daily return distributions may arise through stochastic volatility and leverage effects and/or or "jumps" possibly with time-varying intensity. As such, these earlier empirical studies are merely suggestive about the presence of temporal variation in the jump tail index.

[4] A related estimation strategy has also been proposed in independent work by Hamidieh (2011).

[5] There is also a literature on Lévy-based option pricing outside the Merton-framework, in which the underlying price is modeled as an exponential-Lévy process (see, e.g., Cont and Tankov, 2004, and the references therein), or as a time-changed Lévy process (see, e.g., Carr et al., 2003). However, these studies generally impose tight parametric structures on the volatility process and the distributions of the jumps.

[6] Recall $\mu([0,t], A) = \sum_{s\leq t} 1_{\{\log(\Delta X_s)\in A\}}$ for any measurable $A \in \mathbb{R} \setminus \{0\}$ and $\Delta X_s = X_s - X_{s-}$. Specific examples of jump compensators are given later in Eqs. (2.7), (2.12) and (5.1).

[7] We have implicitly assumed that X_t does not have fixed times of discontinuities. This assumption is satisfied by virtually all asset pricing models hitherto used in the literature. Note also that μ is the counting measure for the jumps in $\log(X_t)$.

Please cite this article in press as: Bollerslev, T., Todorov, V., Time-varying jump tails. Journal of Econometrics (2014), http://dx.doi.org/10.1016/j.jeconom.2014.05.007