Journal of Econometrics 157 (2010) 220–235

Contents lists available at ScienceDirect

Journal of Econometrics

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

Jumps and betas: A new framework for disentangling and estimating systematic risks☆

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

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

ARTICLE INFO

Article history:
Received 26 October 2009
Received in revised form
5 November 2009
Accepted 23 November 2009
Available online 21 December 2009

JEL classification:
C13
C14
G10
G12

Keywords:
Factor models
Systematic risk
Common jumps
High-frequency data
Realized variation

ABSTRACT

We provide a new theoretical framework for disentangling and estimating the sensitivity towards systematic diffusive and jump risks in the context of factor models. Our estimates of the sensitivities towards systematic risks, or betas, are based on the notion of increasingly finer sampled returns over fixed time intervals. We show consistency and derive the asymptotic distributions of our estimators. In an empirical application of the new procedures involving high-frequency data for forty individual stocks, we find that the estimated monthly diffusive and jump betas with respect to an aggregate market portfolio differ substantially for some of the stocks in the sample.

1. Introduction

Linear discrete-time factor models permeate academic asset pricing finance and also form the basis for a wide range of practical portfolio and risk management decisions. Importantly, within this modeling framework equilibrium considerations imply that only non-diversifiable risk, as measured by the factor loading(s) or the sensitivity to the systematic risk factor(s), should be priced, or carry a risk premium. Conversely, so-called neutral strategies that immunize the impact of the systematic risk factor(s) should earn the risk free rate.

Specifically, consider the one-factor representation,

$$r_i = \alpha_i + \beta_i r_0 + \epsilon_i, \quad i = 1, \ldots, N, \tag{1}$$

where r_i and r_0 denote the returns on the ith asset and the systematic risk factor, respectively, and the idiosyncratic risk, ϵ_i, is assumed to be uncorrelated with r_0. Then, provided sufficiently weak cross-asset dependencies in the idiosyncratic risks (see Ross (1976) and Chamberlain and Rothschild (1983)), the absence of arbitrage implies that $\mathbb{E}(r_i) = r_f + \lambda_0\beta_i$, where r_f and λ_0 denote the risk free rate and the premium for bearing systematic factor risk, respectively, so that the differences in expected returns across assets are solely determined by the cross-sectional variation in the betas. This generic one-factor setup obviously encompasses the popular market model and CAPM implications in which the betas are proportional to the covariation of the assets with respect to the aggregate market portfolio. However, the use of other benchmark portfolios in place of r_0, or more general dynamic multi-factor representations (see, e.g., the discussion in Sentana and Fiorentini (2001) and Fiorentini et al. (2004)), attach the same key import to the corresponding betas.

The beta(s) of an asset is(are), of course, not directly observable. The traditional way of circumventing this problem and estimating betas rely on rolling linear regression, typically based on five years of monthly data, see, e.g., the classical studies by Fama and MacBeth (1973) and Fama and French (1992). Meanwhile, the recent advent of readily-available high-frequency financial

☆ Bollerslev's work was supported by a grant from the NSF to the NBER and CREATES funded by the Danish National Research Foundation. We would like to thank our referees for many helpful comments, which have greatly improved the paper. We would also like to thank George Tauchen for numerous insightful discussions on closely related ideas, and Tzuo Hann Law for help with extracting and cleaning the high-frequency data used in the empirical analysis.

* Corresponding author. Tel.: +1 919 660 1846.
E-mail addresses: v-todorov@northwestern.edu (V. Todorov), boller@duke.edu (T. Bollerslev).

0304-4076/$ – see front matter © 2009 Elsevier B.V. All rights reserved.
doi:10.1016/j.jeconom.2009.11.010

prices have spurred a renewed interest into alternative ways for more accurately estimating betas. In particular, Andersen et al. (2005), Andersen et al. (2006), Bollerslev and Zang (2003), and Barndorff-Nielsen and Shephard (2004a) among others, have all explored new procedures for measuring and forecasting period-by-period betas based on so-called realized variation measures constructed from the summation of squares and cross-products of higher frequency within period returns. These studies generally confirm that the use of high-frequency data results in statistically far superior beta estimates relative to the traditional regression based procedures.

Meanwhile, another strand of the burgeoning recent empirical literature concerned with the analysis of high-frequency intraday financial data has argued that it is important to allow for the possibility of price discontinuities, or jumps, in satisfactorily describing financial asset prices; see, e.g., Andersen et al. (2007), Barndorff-Nielsen and Shephard (2004b, 2006), Huang and Tauchen (2005), Mancini (2001, 2008), Lee and Mykland (2008) and Ait-Sahalia and Jacod (2009b). Related to this, there is mounting empirical evidence from derivatives markets that options traders price the expected variation in equity returns associated with sharp price discontinuities, or jumps, differently from the expected variation associated with more smooth, or continuous, price moves; see, e.g., Bates (2000), Eraker (2004), Pan (2002) and Todorov (forthcoming). In other words, it appears as if the market rewards erratic price moves differently from more orderly or smooth price variation, and implicitly treating the risk premia for two different types of price variation to be the same, as it is commonly done in most existing pricing models, is too simplistic.

Combining these recent ideas and empirical observations naturally suggests decomposing the return on the benchmark portfolio(s) within the linear factor model framework into the returns associated with continuous and discontinuous price moves (r_0^c and r_0^d, respectively). In particular, for the one-factor model in Eq. (1),

$$r_i = \alpha_i + \beta_i^c r_0^c + \beta_i^d r_0^d + \epsilon_i, \quad i = 1, \ldots, N, \tag{2}$$

where by definition $r_0 = r_0^c + r_0^d$, and the two separate betas represent the systematic risks attributable to each of the two return components.[1] Of course, for $\beta_i^c = \beta_i^d$ the model trivially reduces to the standard one-factor model in Eq. (1). However, there is no a priori theoretical reason to restrict, let alone expect, the two betas to be the same. Indeed, the classical paper by Merton (1976) hypothesized that in the context of the market model, jump risks for individual stocks are likely to be non-systematic, so that effectively $\beta_i^d = 0$. On the other hand, the evidence for larger cross-asset correlations for extreme returns documented in Ang and Chen (2002) among others, indirectly suggests non-zero jump sensitivities, or $\beta_i^d > 0$. Despite the potential importance of such a decomposition, both from a theoretical asset pricing as well as a practical portfolio management perspective, direct empirical assessment has hitherto been hampered by the lack of formal statistical procedures for actually estimating different types of beta. The present paper fills this void by developing a general theoretical framework for disentangling and separately estimating the sensitivity towards systematic continuous and systematic jump risks. For simplicity and ease of notation, we will focus on the one-factor representation in Eq. (2), but the same ideas and estimation procedures extend to more general multi-factor representations.

The asymptotic theory underlying our results rely on the notion of increasingly finer sampled returns over a fixed time-interval. Our estimation and inference procedures thus extend the results in Barndorff-Nielsen and Shephard (2004a) on realized covariation measures for continuous sample path diffusions. The derivation of our results directly builds on and extends the work of Jacod (2008) on power variation for general semimartingales (containing jumps) as well as the recent work of Ait-Sahalia and Jacod (2009b) and Jacod and Todorov (2009) on testing for jumps in discretely sampled univariate and multivariate processes. Related ideas have also recently been explored by Mancini (2008) and Gobbi and Mancini (2008). Additionally, we also utilize the procedures of Barndorff-Nielsen and Shephard (2004b) and Barndorff-Nielsen et al. (2005) for measuring the continuous sample path variation in the construction of feasible estimates for the asymptotic variances of the betas.

To illustrate the practical usefulness of the new procedures, we estimate separate continuous and jump betas with respect to an aggregate market portfolio for a sample of forty individual stocks, focussing on the monthly horizon. Consistent with the aforementioned studies on high-frequency based beta estimates, which implicitly restrict the two kinds of beta to be the same, we find overwhelming empirical evidence that both kinds of beta vary non-trivially over time. Our findings of systematically positive jump betas for all of the stocks directly contradict the notion that jump risk is diversifiable. Our results also show that for some of the stocks the two types of beta can be quite different, with the estimated jump betas typically being *larger* and *less* persistent than their continuous counterparts.

The calendar time span of high-frequency data available for the empirical analysis is too short to allow for the construction of meaningful statistical tests for whether the separate betas truly reflect differences in priced systematic risks. However, the differences in the magnitudes of the estimates for some of the companies are such that the new betas developed here could make a material difference in terms of pricing and similarly allow for more informed portfolio and risk management decisions.

The rest of the paper proceeds as follows. Section 2 details our theoretical setup and assumptions, along with the intuition for how to calculate continuous and jump betas in the unrealistic situation when continuous price records are available. Our new procedures for actually estimating separate betas based on discretely sampled high-frequency observations and the corresponding asymptotic distributions allowing for formal statistical inference are presented in Section 3. Our empirical application entailing estimates of the betas for the extended market model for the forty individual stocks is discussed in Section 3. Section 4 concludes. All of the proofs are relegated to a technical Appendix.

2. The continuous record case and assumptions

Discrete-time models and procedures along the lines of the simple one-factor model in Eq. (1) are commonly used in finance for describing returns over annual, quarterly, monthly or even daily horizons. Our goal here is to make inference for the separate betas in the extended one-factor model in (2) under minimal assumptions about the processes that govern the returns within the discrete time intervals.

To this end, assume that within some fixed time-interval $[0, T]$ the log-price p_i is generated by the following general process (defined on some filtered probability space $(\Omega, \mathcal{F}, (\mathcal{F})_{t\geq 0}, \mathbb{P})$),

$$\begin{aligned} \mathrm{d}p_{it} &= \alpha_{it}\mathrm{d}t + \beta_i^c \sigma_{0t}\mathrm{d}W_{0t} + \sigma_{it}\mathrm{d}W_{it} \\ &\quad + \beta_i^d \int_{E^0} \kappa(\delta_0(t,x))\tilde{\mu}_0(\mathrm{d}t, \mathrm{d}\mathbf{x}) \\ &\quad + \beta_i^d \int_{E^0} \kappa'(\delta_0(t,x))\mu_0(\mathrm{d}t, \mathrm{d}\mathbf{x}) + \int_{E_i} \kappa(\delta_i(t,x))\tilde{\mu}_i(\mathrm{d}t, \mathrm{d}\mathbf{x}) \\ &\quad + \int_{E_i} \kappa'(\delta_i(t,x))\mu_i(\mathrm{d}t, \mathrm{d}\mathbf{x}), \quad i = 1, \ldots, N, \end{aligned} \tag{3}$$

[1] A very different economically motivated decomposition of the beta within the context of the one factor market model into so-called cash-flow and discount rate betas has recently been proposed by Campbell and Vuolteenaho (2004).