Time-varying Volatility and the Power Law Distribution of Stock Returns

Missaka Warusawitharana*

Board of Governors of the Federal Reserve System

February 19, 2016

JEL Classifications: C58, D30, G12

Keywords: Tail distributions, high frequency returns, power laws, time-varying volatility

*I thank Yacine Aït-Sahalia Dobislav Dobrev, Joshua Gallin, Tugkan Tuzun, Toni Whited, Jonathan Wright and Yesol Yeh for helpful comments on the paper. Contact: Division of Research and Statistics, Board of Governors of the Federal Reserve System, Mail Stop 97, 20th and C Street NW, Washington, D.C. 20551. m1mnw00@frb.gov, (202)452-3461.

Time-varying Volatility and the Power Law Distribution of Stock Returns

Abstract

While many studies find that the tail distribution of high frequency stock returns follow a power law, there are only a few explanations for this finding. This study presents evidence that time-varying volatility can account for the power law property of high frequency stock returns. The power law coefficients obtained by estimating a conditional normal model with nonparametric volatility show a striking correspondence to the power law coefficients estimated from returns data for stocks in the Dow Jones index. A cross-sectional regression of the data coefficients on the model-implied coefficients yields a slope close to one, supportive of the hypothesis that the two sets of power law coefficients are identical. Further, for most of the stocks in the sample taken individually, the model-implied coefficient falls within the 95 percent confidence interval for the coefficient estimated from returns data.

1 Introduction

A growing literature has documented that the tail distributions of a broad range of data in the natural and social sciences follow a power law, which implies that the tail probability of a series declines proportionally as it increases in value.[1] This finding has drawn attention both due to the simplicity and scalability of the power law distribution as well as for the ubiquity of this relationship across many fields and data sets. As noted by Stumpf and Porter (2012), however, our understanding of the causes and mechanisms that underpin many of these relationships lag the empirical evidence on them. This study examines whether time-varying volatility can help explain one such finding—the power law property of high-frequency stock returns documented by Plerou, Gopikrishnan, Amaral, Meyer, and Stanley (1999), among others.

Gabaix, Gopikrishnan, Plerou, and Stanley (2003) and Gabaix, Gopikrishnan, Plerou, and Stanley (2006) provide the leading explanation for the power law property of stock returns. These studies argue that the tail distribution of stock returns arises from the price impact of trades initiated by different market participants, who themselves have a tail distribution of assets. Using data on the tail distributions of trades and assets, they demonstrate how a model of price impact can simultaneously explain the tail properties of trades and returns, as these processes in the model inherit the power law distribution of investor assets.

Building on the insight of Clark (1973) and others that mixture distributions have heavy tails, this study argues that time-varying volatility can generate the power law property. The key insight of the study is that a conditional normal model with time-varying volatility may exhibit a power law distribution in the tails. This result is established formally using a stylized conditional normal model with an inverse-gamma distribution for return variance. As such a model is unlikely to accurately capture the distribution of volatility in the data, the empirical analysis focuses on a conditional normal model with a nonparametric volatility distribution.

[1] Gopikrishnan, Plerou, Gabaix, and Stanley (2000) find that returns on stock indices exhibit a power law with a power coefficient around 3; Plerou, Gopikrishnan, Amaral, Meyer, and Stanley (1999) find that the returns on individual stocks follow a power law with coefficients ranging from 2.5 to 4; Axtell (2001) shows that the firm size distribution follows a power law with a coefficient of 1; Rozenfeld, Rybski, Gabaix, and Makse (2011) present evidence that the distribution of populated areas follows a power law with a coefficient of 1; Toda (2012) demonstrates that income follows a double power law; and Walsh and Toda (2014) examine the implications of the double power law distribution of consumption growth on asset pricing models.

Using more than 925 million distinct observations of stock prices at 1 second intervals, this study examines whether such a model can explain the cross-sectional variation in power law coefficients for the 41 stocks that were included in the Dow Jones Industrial Index at some point from 2003 to 2014. The analysis reveals strong evidence in favor of the conditional normal model with nonparametric volatility. Specifically, focusing on the cross-section of stocks, one finds that the model-implied power law coefficients co-move one-for-one with those obtained from return data. More strikingly, examining the stocks in the sample individually one finds evidence that, for the bulk of these stocks, one would not reject the hypothesis that the model-implied power law coefficient differs from that obtained from return regressions.

The empirical analysis begins by estimating power law coefficients using 15 minute returns from 2003 to 2014 for the 41 stocks in the sample. The results indicate that a power law fits the tail distribution of stocks returns. The estimated power law coefficients are centered around 3, consistent with the findings of Plerou, Gopikrishnan, Amaral, Meyer, and Stanley (1999). The cross-sectional distribution of the estimated coefficients, which range from 2.09 to 3.46, provides the key source of variation for the subsequent analysis in the study.

Estimating the conditional normal models with time-varying volatility requires measuring volatility at the high frequencies used to measure returns. This poses a challenge, as the realized volatility method developed by Andersen, Bollerslev, Diebold, and Labys (2003) and Barndorff-Nielsen and Shephard (2002) runs into difficulty with market microstructure noise when applied at such high frequencies.[2] As such, the two-scales realized volatility method developed by Zhang, Mykland, and Aït-Sahalia (2005) is used to obtain return volatility within 15 minute intervals, the primary time interval used in the study. Applying this method to returns at 30, 45, or 60 second intervals offset at intervals of 1 second each, one obtains robust measures of return volatility at 15 minute intervals.

The first test examines a conditional normal model with an inverse-gamma distribution for 15-minute return volatility; this model has the parametric property that the tail distribution is given by a power law distribution with a power coefficient equal to twice the shape coefficient of the inverse-gamma distribution for volatility. Comparing the power law coefficient obtained from 15-minute returns with that obtained by estimating the model re-

[2]Barndorff-Nielsen and Shephard (2004) extends this method to estimate realized covariances and regression coefficients.

veals a positive relationship between these two sets of coefficients. However, one rejects the hypothesis that the two sets of coefficients are equal.

The subsequent tests examine a conditional normal model with a nonparametric volatility distribution. For each stock, the power law coefficient implied by this model is obtained from a simulated data set constructed using the empirical distribution for return volatility at 15 minute intervals. A comparison of the model-implied power law coefficients and the corresponding estimates from return data reveals a close relationship. The correlation coefficient between these two sets of power law coefficients exceeds 0.95. A regression through the origin of the data coefficients on the model-implied coefficients yields slope coefficients between 1.00 and 1.01. These coefficients are statistically indistinguishable from one, providing support for the hypothesis that the two sets of power law coefficients are identical. Testing the model one stock at a time, one finds that the model-implied coefficient falls within the 95 percent confidence interval of the power law coefficient estimated from returns data for the bulk of the stocks in the sample. The ability of the model to fit the data, both in the cross-section and at the individual stock level, provides striking evidence that the conditional normal model with time-varying volatility can help explain the power law property of stock returns.

One limitation of the above model is that it abstracts from the clustering of volatility observed in the data. As the model-implied power law coefficient is determined by the unconditional distribution of volatility, in principle, this may not be a concern if there are sufficient observations to accurately capture the unconditional distribution of volatility, regardless of the underlying conditional distribution. That said, in order to address this concern, the nonparametric model is estimated using a block sampling scheme that preserves the observed time-series dependency of volatility. One obtains the same key findings from this setting: the slope coefficients of the regression through the origin of the power law coefficients from returns on the model-implied coefficients equal one and, for most of the stocks in the sample taken individually, the model-implied coefficient falls within the 95 percent confidence interval for the coefficient from the return data.

The above findings may reflect the fact that periods of high volatility are associated with tail returns. Investigations of related models reveal that these models can explain some dispersion in observed power law coefficients; however, they do not exhibit the close fit of the conditional normal model with nonparametric volatility. Further, Monte Carlo analysis of hypothetical data that capture the correlation between returns and volatility in the data show that the association between returns and volatility explains some, but not all of the patterns observed in the data. The latter result indicates that the key findings of the study are unlikely to occur spuriously.

Taken together, these findings suggests that time-varying volatility, a key property of financial markets emphasized in influential studies such as Engle (1982), Bollerslev (1986) and Schwert (1989), can explain the power law property of stock returns. The results indicate that economic forces such as time-varying economic uncertainty (see Bloom (2009)) or changes in the arrival rate of information (see Mandelbrot and Taylor (1967) and Kyle and Obizhaeva (2013)) may account for the remarkable regularity of the tail distribution of stock returns.

This study is organized as follows. Section 2 describes the statistical models used in the study; Section 3 details the data used to test these models; Section 4 presents the results of the empirical analysis and Section 6 concludes.

2 Models

This section provides a brief overview of power law distributions and presents the main parametric and nonparametric models analyzed in the study.

2.1 Power law distribution

A random variable X follows a power law distribution in the tail if, for large x, it has the following probabiliby density function

$$P(X = x) \propto x^{-\alpha-1},$$

where α is referred to as the power coefficient. A power law distribution has the property that its tail distribution is given by the following:

$$P(X > x) \propto x^{-\alpha}. \tag{1}$$

This relationship is typically tested using a log-log plot of the tail distribution $P(X > x)$ on x. A linear regression of $\log(P(X > x))$ on $\log(x)$ provides an estimate of the power law coefficient, α. When the primary focus is on whether the tail distribution of a given series follows a power law, this regression is run only on values of x that exceed a given threshold.

2.2 Parametric model

Consider the following highly stylized parametric distribution for high frequency log returns:

$$\begin{aligned} r_t &\sim N(0, \sigma_t^2) \\ \sigma_t^2 &\sim \text{Inv-Gamma}(\alpha, \beta). \end{aligned} \tag{2}$$

That is, log equity returns follow a conditional normal distribution, with variance (σ_t^2) drawn from an inverse-gamma distribution with shape paramater α and scale parameter β. The conditional normal distribution has a long history in understanding asset returns dating back to Engle (1982).[3] The inverse-gamma distribution for variance provides a parsimonious specification that captures time-variation in return volatility.[4]

The joint density function of the above model can be written as:

$$p(r_t, \sigma_t^2) = \frac{1}{\sqrt{2\pi}\sigma_t} \exp\left(\frac{-r_t^2}{2\sigma_t^2}\right) \times \frac{\beta^\alpha}{\Gamma(\alpha)} (\sigma_t^2)^{-\alpha-1} \exp\left(\frac{-\beta}{\sigma_t^2}\right), \tag{3}$$

where $\Gamma(\cdot)$ denotes the gamma function. Following the derivation in Poirier (1995), one can show that the marginal distribution of returns follows a Student's t-distribution:

$$p(r_t) = \frac{\Gamma((2\alpha+1)/2)}{\Gamma(\alpha)\sqrt{(2\pi\beta)}} \left(1 + \frac{r_t^2}{2\beta}\right)^{-(2\alpha+1)/2}. \tag{4}$$

The above t-distribution has a degree-of-freedom equal to twice the shape parameter of the inverse-gamma distribution, 2α, a mean of 0, and, for $\alpha > 1$, a variance equal to $\frac{\beta}{\alpha-1}$.

[3] Andersen, Bollerslev, Diebold, and Ebens (2001) provide evidence that daily stock returns are well described by a conditional normal distribution.

[4] The above specification is also used in the Bayesian literature as a conjugate informative prior.

Using the fact that the tail probability of a t-distribution is given by a power law with power coefficient equal to the degree-of-freedom of the t-distribution, one can write the unconditional tail distribution of returns as:

$$P(r_t > r) \propto r^{-2\alpha}. \tag{5}$$

The above derivation shows that the parsimonious conditional normal distribution shown in Equation (2) implies a power law distribution for returns.[5] Strikingly, the power law coefficient is inherited directly from the unconditional distribution of the variance of returns, providing a simple testable prediction.

2.3 Nonparametric model

An inverse-gamma distribution is unlikely to accurately capture the distribution of stock return variance over time. As such, this study also examines the following model:

$$\begin{aligned} r_t &\sim N(0, \sigma_t^2) \\ \sigma_t^2 &\sim f(\cdot), \end{aligned} \tag{6}$$

where $f(\cdot)$ denotes a nonparametric distribution for variance. While this model lacks the clear parametric implications of the previous model, the nonparametric specification has the flexibility to match the observed empirical distribution of return volatility. This model is subsequently tested by comparing the power law coefficient obtained from simulating this model using the empirical distribution for return volatility at high frequencies with the power law coefficient obtained using data on returns.

One limitation of the power law literature in general, and the models (1), (2) and (6) in particular, is that they are not invariant to time-scales, as these distributions are not invariant under addition. In practice, this appears to have little impact, as Plerou, Gopikrishnan, Amaral, Meyer, and Stanley (1999) find that power law coefficients remain mostly constant for returns measured over time intervals ranging from 5 minutes to 1 day.

[5] Weitzman (2007) shows that a t-distribution for consumption growth arising from parameter uncertainty may help explain the equity premium.

2.4 Sources of time-variation in volatility

The above models emphasize the role of time-varying volatility. Economically, this likely reflects time-varying uncertainty about stock returns. There are many sources of uncertainty in an economy. One potential source is uncertainty about actions that may be taken by policy makers, sometimes referred to as policy uncertainty (see Baker, Bloom, and Davis (2013)). In addition, the magnitude of the shocks faced by the economy may be larger in periods of stress such as the recent financial crisis than others. Changes in the quantity of information produced in the economy can also lead to changes in uncertainty over time (see Gorton and Ordonez (2014)). Finally, as emphasized by Kyle and Obizhaeva (2013), the rate of arrival of new information may vary over time, leading to time-variation in uncertainty over a fixed time interval.[6] All of these sources of uncertainty can lead to time-variation in the volatility of stock returns.

3 Data

The data used in the study are obtained from the Consolidated Trades segment of the Trades and Quotes database accessed through Wharton Research Data Services. The sample comprises all stocks that were included in the Dow Jones Industrial Average at some point from 01/01/2003 to 12/31/2014. Including additions and deletions to the index, this list comprises 41 stocks. These stocks were selected as they are among the most liquid stocks traded in the US equity markets. As such, market microstructure concerns are likely to have a smaller impact on measured returns and volatilities for these stocks.

For each of the stocks in the sample, stock prices are measured at 1 second intervals daily from 01/01/2003 to 12/31/2014. For stocks that had multiple trades within the same second, the stock price is obtained from the median of the reported transaction prices. The sample excludes dates with an early close to trading activity.[7] In order to avoid the effect of information released at the beginning or the end of the trading day, stock prices are measured from 9:45am to 3:45pm. Even taking into account that not all stocks are traded every second, the data set consists of more than 925 million distinct stock price observations.

[6] See Andersen, Bollerslev, Diebold, and Vega (2003) for evidence that news announcements affect prices.

[7] I thank Yesol Huh for providing data on early closing dates.

Following the convention in the realized volatility literature, throughout the study, returns are measured as logarithmic returns, rather than arithmetic returns.[8] Given the short time period over which returns are measured, the logarithmic returns are almost identical to arithmetic returns, and this choice has little effect on the results (see also Kelly (2014)).

The empirical analysis in the study primarily uses stock returns and volatilities measured over 15 minute intervals, the time interval used in Gabaix, Gopikrishnan, Plerou, and Stanley (2006).[9] A 15 minute interval provides a sufficient number of price observations to measure volatility within that interval—see the discussion below—while also providing a large number of return observations that enable a precise estimate of the power law coefficient. For robustness, the analysis also examines results at 10 and 30 minute intervals.

3.1 Measurement of volatility

The analysis requires measurement of stock return volatility within 15 minute intervals. This poses a challenge, as the traditional approach in the realized volatility literature measures volatility at daily frequencies using either 1- or 5-minute returns (see Andersen, Bollerslev, Diebold, and Labys (2003)). There are too few observations of 1- or 5-minute returns to derive a robust measure of volatility at 15 minute intervals. One cannot address this issue by measuring returns at much higher frequencies, as that increases the market microstructure component of observed returns, thereby adding noise to the realized volatility measures.[10]

This study tackles this challenge by measuring volatility at 15 minute frequencies using the two-scales realized volatility (hereafter, TSRV) method developed by Zhang, Mykland, and Aït-Sahalia (2005) and Aït-Sahalia, Mykland, and Zhang (2011). They note that measuring realized volatility using returns over 5 minute intervals involves a considerable loss of information and develop a method to measure realized volatility utilizing stock price data at much higher frequencies. This method involves first computing realized volatility using returns at a frequency typically used in the literature, such a 1 minute frequency. Based on the observation that there are many such overlapping 1-minute return intervals within

[8]Given a stock price process, Y_t, logarithmic returns are given by $\log\left(\frac{Y_{t+1}}{Y_t}\right)$, while arithmetic returns are given by $\frac{Y_{t+1}}{Y_t} - 1$.

[9]The stock price at the end of each 15 minute interval is obtained as the last price at which the stock traded prior to the end of the 15 minute interval.

[10]Hansen and Lunde (2006) discusses the effect of market microstructure noise on realized volatility measures as the interval over which returns are measured shrinks.

a period of interest, the above studies show that the average of the realized volatility measure over these overlapping intervals, adjusted for a bias correction term involving realized volatility calculated using returns at the higher frequency, provides an estimate of realized volatility over the desired interval.[11] For instance, using a smaller time scale of 1 second, one could obtain realized volatility based on the average of 60 different 1-minute returns, each offset by 1 second. This method enables one to measure realized volatility using stock price information at the much higher frequency by which the returns are offset rather than the lower frequency at which the squared returns are measured. The Appendix provides further details on the measurement of volatility using this method.

The results reported in this study are obtained using the TSRV method applied at frequencies of 30, 45 and 60 seconds offset at 1 second intervals. This enables one to use up to 900 observations of stock prices at 1 second intervals to measure realized volatility within each 15 minute interval. Observations for which the bias correction term causes one of the TSRV measures to be negative are discarded from the sample. Finally, the realized volatility measures in the sample are Winsorized at the 0.01th percentile to remove the effect of extreme outliers.[12] This affects only a very small fraction (0.02%) of the observations, and is done primarily to limit the effect of outliers on the parameter estimates of the inverse-gamma model.

3.2 Summary statistics

Table 1 reports summary statistics for the 15-minute return and volatility data used in the subsequent analysis. Panel A reports univariate statistics; Panel B reports the correlation coefficients among the TSRV volatility measures obtained using different return frequencies; and Panel C reports the correlation coefficients among the corresponding realized volatility measures. Taken together, the data set used in the analysis includes more than 2 million observations of stock returns and volatilities measured over 15-minute intervals (note that many more stock price observations were used to construct these volatility measures). On a per stock basis, the data set contains about 70,000 15-minute return and volatility data for

[11]The TSRV measure used in the study also incorporates the small sample refinement proposed in Aït-Sahalia, Mykland, and Zhang (2011).

[12]That is, observations with realized volatility above the 99.99th percentile are set to the value at this percentile while observations below the 0.01th percentile are set to that value.

most of the 41 stocks in the sample. The large sample size helps ensure that small sample concerns do not affect the estimated power law coefficients and that the data set captures the unconditional distributions of return volatility.

The univariate statistics indicate that the mean 15-minute return among the stocks in the sample is slightly positive, reflecting the fact that overall stock returns were positive over the sample period. The corresponding median return equals zero. The statistics also indicate that the TSRV measures provide a robust estimate of realized volatility at 15 minute intervals over the three return frequencies used in the study. Moving from measuring squared returns at a 60 second interval to a 30 second interval leads to only a slight increase in the mean and dispersion of realized volatilities, suggesting that the TSRV approach manages to successfully account for market microstructure noise in the data while incorporating a sizeable quantity of stock price information. The robustness of this approach is further supported by the correlation matrix shown in Panel B, which indicates that the three TSRV measures have a very high correlation with one another.

Panel C reports the corresponding correlation matrix for realized volatility measures at 15-minute intervals obtained using the approach in Andersen, Bollerslev, Diebold, and Labys (2003) applied to squared returns at 30, 45 and 60 second intervals. The results indicate that while the realized volatility measures are closely correlated, the correlations among them are not as strong as those for the TSRV volatility measures. This suggests that the TSRV approach provides a more robust measure of volatility within 15 minute intervals than the realized volatility method.

4 Results

The empirical analysis tests the models presented in Section 2 by comparing the cross-section of power law coefficients obtained from the tail distribution of returns with the corresponding model-implied coefficients. That is, the tests examine whether a conditional normal model with time-varying volatility can explain the observed variation in the power law coefficients across the stocks in the sample. In comparison, prior studies such as Gabaix, Gopikrishnan, Plerou, and Stanley (2006) have focused on explaining a single power law coefficient obtained by aggregating information from different stocks.

4.1 Power law coefficients from stock returns

The power law coefficients for the tail distribution of stock returns are obtained from regressions of the log probability that the absolute stock return exceeds a given value on the log of the absolute return. These regressions are carried out for absolute stock returns greater than the 95$^{\text{th}}$ percentile. The standard error of each regression coefficient is obtained using the method in Gabaix (2009). As the sample consists of about 70,000 returns for most stocks, each regression uses about 3,500 observations on tail probabilities, mitigating small sample concerns.

Figure 1 presents log-log plots of the tail distribution of the absolute value of 15-minute returns for each of the stocks in the sample. It also shows the fitted line obtained from a regression of the log tail probability on the absolute stock return. As the figure shows, a linear regression fits the tail distribution for the absolute returns of the stocks used in the sample well, indicating that the tail distribution of returns follows a power law, consistent with the literature. For some of the stocks in the sample, the fitted line exceeds the implied probability that the abolute return will be greater than a given value for the small number of observations near the very tail of the distribution.[13]

Figure 2 gathers the above results and plots the distribution of the estimated power law coefficients for all the stocks in the sample. As the figure indicates, the distribution of the estimated power law coefficients is centered around 3, with a mean value of 2.92 and a median of 3.00. This result matches the finding in Gabaix, Gopikrishnan, Plerou, and Stanley (2003) and others that high frequency stock returns exhibit a power law coefficient of around 3. The estimated power law coefficients exhibit some dispersion, with a cross-sectional standard deviation of 0.32. The subsequent empirical analysis examines whether conditional normal models can capture this cross-sectional variation. In comparison, prior studies have not focused on examining the cross-section of power law coefficients observed in the data.

[13]This effect is visually enhanced by the log-log plot, which gives equal prominence to values that have tail probabilities between 10^{-4} and 10^{-5} as those that have a tail probability between 10^{-2} and 10^{-3}, even though there are far fewer observations of the former.

4.2 Estimation of parametric model

As noted in Section 2.2, a conditional normal model with an inverse-gamma volatility distribution implies a power law distribution for stock returns. This section reports the results from testing this model by comparing the power law coefficients obtained from the data with the model-implied coefficients, given by 2 times the shape coefficient of the inverse-gamma distribution for stock return variance measured at 15 minute intervals. The shape and scale parameters of the inverse-gamma distributions are obtained using maximum likelihood estimation.

Figure 3 plots the power law coefficient obtained from the regression analysis in Section 4.1 on two times the shape parameters of the inverse-gamma distributions for return variance. Panels A, B and C, respectively, plot this correspondence for the variance distributions obtained using the TSRV method with volatility measured within each 15-minute window using returns at 30, 45 and 60 second intervals offset at 1 second. The figures also report the slope of a regression through the origin of the power law coefficients from the data on the model-implied coefficients.

The null hypothesis in this (and subsequent related) analysis is whether the power law coefficients implied by the model are equal to their data counterparts. This section examines the null hypothesis by testing whether the slope coefficient of a regression through the origin of the power law coefficients from the data on the model-implied coefficients equals 1. Notably, the constant term in the regression equals zero by construction to reflect the fact that a simple linear model with a non-zero constant term would be a misspecified model for examining whether the two sets of coefficients are identical. Imposing the restriction that the constants term in the regression equals zero implies that the fitted line obtained from the regression shown in Figure 3 (and in subsequent figures) does not line up with a more conventional best-fit line obtained from a linear regression with a non-zero constant.

As the figure indicates, there exists a clear positive relationship between the power law coefficients obtained from stock returns and the power law coefficient implied by the conditional normal model given in Equation (2). The slope of the linear regressions through the origin for the three volatility measures are positive and significant, albeit significantly greater than 1, implying that one rejects the null that the two sets of coefficients are identical. While stocks with high/(low) model-implied power law coefficients also have a high/(low) coefficient in the data, there is a greater dispersion in the model-implied coefficients than in the data. Further, the model-implied coefficients are lower, on average, than the coefficients from the returns data.

The estimated model uses a highly stylized distribution for volatility. Comparing the empirical distribution for variance to that obtained by fitting the inverse-gamma distribution, one finds that the inverse-gamma distribution has greater mass at the center of the distribution and less mass at the tails than in the data. The positive relationship between the power law coefficients implied by the model and those from the return regressions suggests examining whether a conditional normal model with a more general volatility distribution can help explain the power law property of stock returns.

4.3 Estimation of nonparametric model

This section examines whether the conditional normal model with a nonparametric distribution for volatility shown in Equation (6) can generate the power law coefficients observed in the data. Using a nonparametric specification allows one to match the observed volatility distribution in the data. However, this poses the limitation that one no longer obtains a parametric derivation of the power law coefficient implied by the model. As such, the model-implied power law coefficients are obtained using simulation methods.

Specifically, for each stock, one obtains the model-implied power law coefficient from a simulated data set of returns, $r_i, i = 1, ..., S$ constructed as follows:

1. Obtain a draw for variance, σ_i^2, by sampling with replacement from the empirical distribution for variance, $f(\cdot)$.

2. Obtain a draw for returns, r_i, by sampling from a normal distribution with mean zero and variance σ_i^2.

This simulated sample has a size, S, of 100 times the number of 15-minute return observations in the data, enabling a precise estimate of the power law coefficient implied by the model; the standard errors of the power law coefficients obtained from the simulated data are all less than 0.01. As such, the subsequent analysis abstracts from the estimation error for the model-implied power law coefficients.[14]

[14]A Monte Carlo analysis using 400 different simulation samples confirms that there is little simulation error in the model-implied coefficients and that this simulation error has no impact on the subsequent analysis.

4.3.1 Cross-sectional evidence

Figure 4A plots the power law coefficients obtained from regression analyses of 15-minute returns data on the power law coefficients obtained using simulated data from the conditional normal model with nonparametric volatility distribution. The volatility distribution underlying the model is given by the empirical distribution of 15-minute return volatility obtained using the TSRV method with 30-second returns offset at 1-second intervals. Figures 4B and 4C, respectively, plot the correspoding results for the simulated data sets from the conditional normal models with volatility obtained using the TSRV method applied to 45- and 60-second returns offset at 1-second intervals. The figures also present the slope coefficients, associated 95 percent confidence intervals and fitted lines obtained from a regression through the origin of the power law coefficients from the returns data on the model-implied power law coefficients. As noted earlier, imposing the restriction that the constant term of the regression equals zero enables one to test the hypothesis that the two sets of power law coefficients are identical.

The figures indicate a striking concordance between the power law coefficients obtained from 15-minute returns and those obtained from the conditional normal models. For each of the three 15-minute return volatility measures, stocks with high/(low) model-implied power law coefficients also have a high/(low) high power law coefficient in the return data. Indeed, the correlation coefficient between the two sets of power law coefficients ranges from 0.96 to 0.98 across the three volatility measures. In addition, the dispersion of the model-implied power law coefficients is also quite similar to that of the data coefficients.

Turning to the regression analysis, one finds that the slope of a regression line through the origin of the power law coefficient from 15-minute returns on the corresponding power law coefficient from the conditional normal models ranges from 1.004 to 1.007. For each of the three volatility measures, the 95 percent confidence interval of the estimated slope coefficient encompasses one. The corresponding p-values for the tests that the slope coefficients reported in Figures 4A, 4B and 4C equal one are 0.41, 0.27, and 0.06, respectively. This indicates that, for all these volatility measures, one would not reject the hypothesis that the two sets of power law coefficients are identical. This close fit indicates that the model can explain the cross-section of power law coefficients observed in the data.

Note that the above finding is a stronger result than that typically obtained in empirical studies, which focus on examining whether a given coefficient is positive or not. Instead, the analysis tests the hypothesis that the two sets of coefficients are identical by examining whether the slope coefficient equals a specific value, one, and finds that one fails to reject this hypothesis.

What explains the improved performance of the nonparametric model compared to the results obtained using the parametric model? Examining the fit of the inverse-gamma distribution for volatility, one finds that the model implies less frequent realizations for inverse volatility than observed in the data around the mode of the distributions and more frequent realizations for very low values. As such, the inverse-gamma distribution predicts a higher frequency of very high volatility realizations than observed in the data. As high volatility realizations are associated with high absolute returns, the return distributions implied by the parametric model has fatter tails than the corresponding distributions implied by the nonparametric model. The fatter tails in the parametric model leads to lower estimates for the power law coefficients α, as can be seen by comparing the results reported in Figure 3 with those reported in Figures 4A, 4B and 4C.

The above confidence intervals and p-values are obtained using OLS standard errors applied to a t-distribution with 40 degrees of freedom. The analysis does not use heteroskedasticity robust standard errors due to the fact that the finite sample properties of the resulting t-statistics are not well known. As such, it is necessary to examine whether the above regressions suffer from heteroskedasticity in the error term. Applying the White (1980) and Breusch and Pagan (1979) heteroskedasticity tests to the regressions for the three volatility measures, one finds little significant evidence of heteroskedasticity.[15] This finding is confirmed by a visual examination of the residuals. Consistent with these findings, the heteroskedasticity robust standard errors for the slope coefficients are within 2 basis points of the OLS standard errors.

4.3.2 Evidence at the individual stock level

Another way of evaluating the model is to investigate whether, for a given stock, the power law coefficient obtained from the simulated data falls within the 95 percent confidence interval of the power law coefficient estimated using 15-minute returns data. This examines whether the model can explain the data at the individual stock level, thus providing a sterner test

[15]The p-values for the White test equals 0.04, 0.21 and 0.45, respectively, while the p-values for the Breusch-Pagan test range from 0.34 to 0.52.

of the model. Table 2 presents the results of this analysis. For each of the stocks in the sample, it reports the power law coefficient estimated from the returns data, the associated 95 percent confidence interval, and the model-implied power law coefficients obtained using simulated data for each of the three measures of 15-minute return volatility used in the study. * and **, respectively, denote stocks for which one can reject the hypothesis the model-implied power law coefficient equals that obtained from the return regression at the 95 and 99 percent confidence intervals.

The results show that the model-implied power law coefficient falls within the 95 percent confidence interval for the regression estimate from returns data for between 36 to 38 of the 41 stocks in the sample. This indicates that, for most of the stocks in the sample taken individually, one would not reject the hypothesis that the power law coefficient implied by the conditional normal model is different from the corresponding coefficient obtained from the regression analysis using 15-minute returns. The ability of the model to fit the power law coefficients obtained from returns at the individual stock level provides strong evidence that the conditional normal model with time-varying volatility can explain the power law property of stock returns.

4.3.3 Volatility clustering

The conditional normal models used in the study abstract from the time-series dependency of volatility, highlighted in Engle (1982), Bollerslev (1986) and Bollerslev and Wright (2001), among others. I.e., the estimation of the non-parametric model employs random sampling of the empirical volatility distribution. This approach is used as the models' implications are based on the unconditional volatility distribution and, in principle, one would be able to accurately obtain the model-implied power law coefficient if one had sufficient observations to capture the unconditional distribution for volatility, regardless of the underlying conditional distribution. The empirical distribution includes about 70,000 15-minute volatility observations over 10 years for most of the stocks in the sample, suggesting that one has sufficient observations to capture the unconditional distribution. Nonetheless, this section uses an alternate sampling scheme to examine whether the results differ when one estimates the model while accounting for the time-series dependency of the volatility distribution.

The nonparametric conditional normal model with volatility clustering is estimated using a block sampling method. That is, the model-implied volatility distribution is obtained by sampling from blocks of 120 consecutive observations for 15-minute volatility from the corresponding empirical distribution.[16] This sampling approach, akin to the block bootstrap method, helps preserve the time series correlation of volatility in the data. The model-implied power law coefficient is obtained from a simulated sample of returns obtained from a conditional normal distribution with these volatility draws.

Figure 5 presents scatter plots of the power law coefficients obtained from the return regression against the coefficients implied by the nonparametric model with volatility clustering. Panels A, B, and C, respectively, present results from measuring volatility at 15 minute intervals using the TSRV method applied to 30, 45, and 60 second intervals offset at one second. Across the three volatility measures and the 41 stocks in the sample, the volatility distribution in the simulated data exhibit an auto-correlation of 0.581, close to the mean auto-correlation coefficient in the volatility data of 0.588. This indicates that the sampling method captures the time-series dependency of volatility in the data.

As the figure indicates, one continues to obtain the key result using this specification. That is, the slope coefficient of the regression through the origin of the power law coefficient from the return regression on the model-implied coefficient is statistically and economically indistinguishable from one for each of the three volatility measures. The 95 percent confidence intervals for the slope coefficient encompass one, and the p-values for testing whether the slope coefficient equals one ranges from 0.11 to 0.45. This indicates that, for all the three volatility measures, one would not reject the hypothesis that the model-implied power law coefficients differ from those obtained from the return regressions.

Furthermore, comparing the results at the individual stock level one finds that for the three volatility measures, the model-implied power law coefficient falls within the 95 percent confidence interval for the coefficient from the return regression for between 36 to 38 stocks in the sample, similar to that obtained from the nonparametric model without volatility clustering.

One limitation of the analysis is that it does not account for the leverage effect of returns on future volatility (see Christie (1982)). This partly reflects the challenge of nonparametric sampling from the unconditional distribution of volatility while incorporating a leverage effect. It also reflects the fact that, empirically, lagged volatility has much stronger predictive

[16] The overall simulation sample size is approximately equal to 100 times the number of 15-minute return observations, similar to the simulation sample size in the estimation without volatility clustering.

power for future volatility than lagged returns, which suggests that incorporating a leverage effect is unlikely to substantially change the above results.

4.3.4 Results from different time intervals

The previous results were obtained using a 15 minute time interval to measure stock returns and volatility. This time interval was chosen partly as it has been used elsewhere in the literature, and partly as it enabled sufficient data on prices at a one-second interval to measure return volatility using the TSRV method. This section examines whether the findings in the study are robust to changes in the time interval over which returns and volatility are measured.

Figure 6 compares the power law coefficients from the data and the corresponding coefficients from the nonparametric conditional normal model using returns and volatility measured at 10 minute intervals.[17] Panels A, B, and C, respectively, report results using the empirical distribution for return volatility within 10 minute intervals obtained using the TSRV method applied to 30, 45, and 60 second returns offset at 1 second intervals. The panels also report the slope coefficient and associated 95 percent confidence interval from a regression through the origin of the power law coefficients in the data on the coefficients from the nonparametric model.

The results indicate that while the model continues to closely match the dispersion of power law coefficients observed in the data, the fit is not quite as good as that for the 15-minute returns. The slope coefficients for the regressions through the origin of the data coefficients on the model-implied coefficients derived from volatility measured using the TSRV method are statistically different from one, with p-values for the test that the coefficient equals one less than 0.01. However, in economic magnitudes, the slope coefficients—which range from 1.019 to 1.024—are quite close to one. Comparing the model-implied power law coefficients with the coefficients obtained from the return data at the individual stock level reveals that, for each of the three volatility measures, the model implied coefficient falls within the confidence interval of the return coefficient for between 30 to 33 of the 41 stocks examined in the study.

Figure 7 presents the corresponding results using stock returns and volatility measured

[17] These results are obtained by sampling volatility observations one at a time from the unconditional distribution.

at 30 minute intervals. As above, Panels A, B, and C, respectively, report results using the empirical distribution for volatility obtained using the TSRV method applied to 30, 45, and 60 second returns offset at 1 second intervals. The cross-sectional results are quite similar to those obtained using 10 minute data. The slope coefficients of the regressions through the origin are statistically different from one for each of the three volatility measures, with p-values for the test that the slope coefficients equal one less than 0.05. While statistically different from one, the slope coefficients—which range from 0.982 to 0.985—are quite close to one in economic terms.

Evaluating the model at the individual stock level, one finds that the fit is similar to that obtained using 15-minute returns (and better than that for the 10-minute returns). Specifically, the model-implied power law coefficient falls within the 95 percent confidence interval for the regression estimate for 37 or 38 of the 41 stocks in the sample. Taken together, these findings indicate that the conditional normal model with a nonparametric distribution for volatility can help explain the power law coefficients observed in the data over a range of time intervals.

4.3.5 Power law distribution for volatility

Gopikrishnan, Plerou, Amaral, Meyer, and Stanley (1999) find that the distribution of stock return volatility itself follows a power law. As such, one possibility is that the power law coefficient for returns inherits this power law coefficient. This possibility is examined by estimating the power law coefficients for the distribution of volatility within 15 minute intervals, and comparing them to the coefficients obtained from stock returns. The comparison provides evidence against this hypothesis, as the slope of a line through the origin from a regression of the power law coefficient on returns on the corresponding power law coefficients on volatility equals 0.88, significant less than one. Further, for the majority of the stocks in the sample, the power law coefficient for the volatility distribution is significantly different from the corresponding coefficient from the return regression. These findings indicate that the ability of the conditional normal model to match the observed distribution of power law coefficients arises from the particular features of that distribution, and is not inherited from the power law distribution for volatility.

5 Monte Carlo analysis of alternate models

One potential concern is that the above results may arise spuriously.[18] A large literature documents a link between stock returns and volatility (for instance, see French, Schwert, and Stambaugh (1987)). Further, one may expect that periods of high volatility would be naturally associated with returns that are large in absolute value, leading to a relationship between the volatility of a stock and its power law coefficient. This section uses a Monte Carlo approach to examine whether two alternate experiments would generate the close fit observed in the data between the power law coefficients from return regressions and those implied by the conditional normal model with a nonparametric volatility distribution. The first experiment takes as given the observed return distribution while the second experiment takes as given the observed volatility distribution.

5.1 Monte Carlo analysis using stock returns distribution

One approach to examining whether the results reported in Section 4.3 may be spurious is to use a Monte Carlo approach to examine whether one would necessarily obtain similar results using a stock returns distribution that matches the data and a volatility distribution that matches some, but not all of the features observed in the data. Specifically, for each stock, fit a linear regression of volatility on the absolute stock return to obtain the constant, regression coefficient and the standard deviation of the regression residual. Randomly sample with replacement from the observed distribution of absolute stock returns to create a hypothetical sample of returns of equal size as the data, and derive an associated volatility sample based on the estimates of the above regression. Repeating this exercise across the 41 stocks in the sample yields a hypothetical data set that matches the return distribution in the data and the mean, standard deviation and the above regression coefficient for volatility observed in the data. The Monte Carlo analysis examines the results that would be obtained by estimating the power law coefficients implied by the conditional normal model with nonparametric volatility on 400 such hypothetical data sets.

[18]Note that the observed results cannot arise purely by construction due to the fact that the returns are measured over a 15-minute interval, while volatilities are measured using prices within that interval. Specifically, a given return over a 15 minute window could be associated with a number of different stock price paths over that window, which would correspond to different volatilities, thereby implying that the observed relationship cannot be purely due to mechanical reasons.

Table 3 reports the results from this exercise. The first two rows report the statistics across the 400 Monte Carlo simulations of the mean and standard deviation of the model-implied power law coefficients. As the table indicates, the model-implied power law coefficients are, on average, larger than the true value of 2.92, with the mean model-implied power law coefficient equal to 3.79. The cross-sectional standard deviation is about 0.39, also above the corresponding data value of 0.32. There is little variation across the Monte Carlo simulations of the above statistics, reflecting the fact that the size of the data set results in little sampling error.

Following the approach of Section 4.3.1, estimating a regression through the origin of the data coefficients on the model-implied coefficients from the Monte Carlo simulations yields a slope coefficient of about 0.77, statistically and economically different from one. This suggests that the conditional normal model with a hypothetical volatility distribution obtained as above would not be able to generate the power law coefficients observed in the data. This arises due to the fact that the above linear relationship used to create the hypothetical volatility data does not generate as many realizations of very high volatility draws as seen in the data; the lower frequency of these high volatility draws lead to few realizations of high stock returns in the conditional normal model and thus a smaller power law coefficient. More generally, the above results indicate that the ability of the conditional normal model to fit the returns data shown in Section 4.3.1 does not arise spuriously due to the linear relationship between volatility and absolute stock returns observed in the data.

5.2 Monte Carlo analysis using volatility distribution

The second approach examines whether an alternate model for returns, that takes as given the observed volatility distribution, would generate similar power law coefficients. Specifically, this section examines a model specification where stock returns are drawn from a conditional Student's t-distribution, with mean zero and volatility drawn from the empirical distribution. The degree-of-freedom for the t-distribution is set equal to 2.34, obtained by fitting the empirical distribution of all stock returns to a t-distribution.[19] The power law coefficients implied by this model are obtained using a simulation approach similar to that

[19] An alternate approach would be to set the degree-of-freedom equal to the corresponding value from returns data for that stock. However, this runs into the difficulty that the implied degree-of-freedom for some stocks is less than 2.

employed in Section 4.3. The Monte Carlo analysis examines the power law coefficients implied by this model using 400 hypothetical data sets obtained by sampling with replacement from the empirical distribution for return volatility.

Table 4 reports the results of this exercise. The first two rows report the statistics across the 400 Monte Carlo simulations of the mean and standard deviation of the model-implied power law coefficients. The mean of the power law coefficient obtained from this approach equals 2.04, lower than the corresponding data value of 2.92. This indicates that the conditional Student's t model implies more fat tails than observed in the data. Consistent with this finding, the slope coefficient from a regression through the origin of power law coefficients from returns data on the model-implied coefficients yields estimates of about 1.44, significantly greater than the estimate of one obtained in Section 4.3.1. This arises from the fact that the conditional Student's t-distribution with volatility given by the empirical distribution contains two sources of tail returns: one associated with the fat tails in the t-distribution, and another associated with high volatility realization. The combination of these two sources of fat tails lead to smaller power law coefficients (greater fat tails) than observed in the data.

Overall, the findings from these two Monte Carlo experiments demonstrate that the association of high volatility realizations with large absolute stock returns leads to a positive relationship between the volatility of a stock and its power law coefficient in a number of different models. Indeed, a similar finding arose from the conditional normal model with the inverse-gamma distribution as well as the model with a power law distribution for volatility. Yet, none of these models could match the observed close fit between the power law coefficients from the returns data and those implied by the conditional normal model with nonparametric volatility, document above in Sections 4.3.1 and 4.3.2. These findings suggest that one needs the combination of conditional normality with the nonparametric volatility distribution to best fit the observed tail distribution of stock returns.

6 Conclusion

This study examines whether a conditional normal model with time-varying volatility can help explain the observed tail distribution of stock returns. Examining a simple parametric model with an inverse-gamma distribution for variance, one finds a positive relationship between the model-implied power law coefficients and the power law coefficients obtained from the data; however, the model-implied coefficients exhibit greater dispersion than the data coefficients. The investigation of a conditional normal model with a nonparametric distribution for volatility reveals a close relationship between the model-implied power law coefficients and the corresponding data values. A cross-sectional analysis using a regression through the origin finds evidence that the two sets of power law coefficients are identical. Furthermore, comparing the model-implied power law coefficients with the corresponding coefficients from the data at the individual stock level, one fails to reject the hypothesis that the model-implied coefficient equals the data coefficient for most of the stocks in the sample, providing evidence in favor of the model at the individual stock level. These results are generally robust to different time intervals for measuring volatility, different intervals for estimating power laws, and to a sampling approach that accounts for volatility clustering. Examining a few alternate models, one finds that while these models capture some of the dispersion in power law coefficients in the data, their performance is not as good as that of the conditional normal model with nonparametric volatility.

These findings indicate that a conditional normal model with time-varying volatility can explain the power law property of stock returns. Economically, this suggests that either changes in uncertainty over time or changes in the arrival rate of information may explain the tail distribution of stock returns. Econometrically, the findings indicate that mixture distributions may help provide insights into why many financial and economic data series exhibit power law properties. Further research into the channels underpinning the observed power law distributions of other series, such as the distribution of firm and city sizes, may prove fruitful.

References

Aït-Sahalia, Yacine, Per A. Mykland, and Lan Zhang, 2011, Ultra high frequency volatility estimation with dependent microstructure noise, *Journal of Econometrics* 160, 160–175.

Andersen, Torben G., Tim Bollerslev, Francis X. Diebold, and Heiko Ebens, 2001, The distribution of realized stock return volatility, *Journal of Financial Economics* 61, 43–76.

Andersen, Torben G., Tim Bollerslev, Francis X. Diebold, and Paul Labys, 2003, Modeling and forecasting realized volatility, *Econometrica* 71, 579–625.

Andersen, Torben G., Tim Bollerslev, Francis X. Diebold, and Clara Vega, 2003, Micro effects of macro announcements: Real-time price discovery in foreign exchange, *American Economic Review* 93, 38–62.

Axtell, Robert L., 2001, Zipf distribution of U.S. firm sizes, *Science* 293, 1818–1820.

Baker, Scott R., Nicholas Bloom, and Steven J. Davis, 2013, Measuring economic policy uncertainty, Working paper, Stanford University.

Barndorff-Nielsen, Ole E., and Neil Shephard, 2002, Econometric analysis of realized volatility and its use in estimating stochastic volatility models, *Journal of the Royal Statistical Society B* 64, 253–280.

Barndorff-Nielsen, Ole E., and Neil Shephard, 2004, Econometric analysis of realized covariation: high frequency based covariance, regression, and correlation in financial economics, *Econometrica* 72, 885–925.

Bloom, Nick, 2009, The impact of uncertainty shocks, *Econometrica* 77, 623–685.

Bollerslev, Tim, 1986, Generalized autoregressive conditional heteroskedasticity, *Journal of Econometrics* 31, 307–327.

Bollerslev, Tim, and Jonathan H. Wright, 2001, High-frequency data, frequency domain inference and volatility forecasting, *Review of Economics and Statistics* 83, 596–602.

Breusch, Trevor, and Adrian R. Pagan, 1979, A simple test for heteroskedasticity and random coefficient variation, *Econometrica* 47, 1287–1294.

Christie, Andrew A., 1982, The stochastic behavior of common stock variances: Value, leverage and interest rate effects, *Journal of Financial Economics* 10, 407–432.

Clark, Peter K., 1973, A subordinated stochastic process model with finite variance for speculative prices, *Econometrica* 41, 135–155.

Engle, Robert F., 1982, Autoregressive conditional heteroskedasticity with estimates of the variance of United Kingdom inflation, *Econometrica* 50, 987–1007.

French, Kenneth R., G. William Schwert, and Robert F. Stambaugh, 1987, Expected stock returns and volatility, *Journal of Financial Economics* 19, 3–29.

Gabaix, Xavier, 2009, Power laws in economics and finance, *Annual Reviews of Economics* 1, 255–293.

Gabaix, Xavier, Parameswaran Gopikrishnan, Vasiliki Plerou, and H. Eugene Stanley, 2003, A theory of power-law distributions in financial market fluctuations, *Nature* 423, 267–300.

Gabaix, Xavier, Parameswaran Gopikrishnan, Vasiliki Plerou, and H. Eugene Stanley, 2006, Institutional investors and stock market volatility, *Quarterly Journal of Economics* 121, 461–504.

Gopikrishnan, Parameswaran, Vasiliki Plerou, Luis A. Nunes Amaral, Martin Meyer, and H. Eugene Stanley, 1999, Scaling of the distribution of fluctuations of financial market indices, *Physics Review E* 60, 5303–5316.

Gopikrishnan, Parameswaran, Vasiliki Plerou, Xavier Gabaix, and H. Eugene Stanley, 2000, Statistical properties of share volume traded in financial markets, *Physics Review E* 62, 4493–4496.

Gorton, Gary, and Guillermo Ordonez, 2014, Collateral Crises, *American Economic Review* 104, 343–378.

Hansen, Peter R., and Asger Lunde, 2006, Realized variance and market microstructure noise, *Journal of Business Economics and Statistics* 24, 127–161.

Kelly, Bryan, 2014, The dynamic power law model, Working paper, University of Chicago.

Kyle, Albert S., and Anna A. Obizhaeva, 2013, Market microstructure invariance: Theory and empirical tests, Working paper, University of Maryland.

Mandelbrot, Benoit, and Howard M. Taylor, 1967, On the distribution of stock price differences, *Operations Research* 15, 1057–1062.

Plerou, Vasiliki, Parameswaran Gopikrishnan, Luis A. Nunes Amaral, Martin Meyer, and H. Eugene Stanley, 1999, Scaling of the distribution of price fluctuations of individual companies, *Physics Review E* 60, 6519–6529.

Poirier, Dale J., 1995, *Intermediate Statistics and Econometrics.* (MIT Press Cambridge, MA).

Rozenfeld, Hernan D., Diego Rybski, Xavier Gabaix, and Hernan A Makse, 2011, The area and population of cities: New insights from a different perspective on cities, *American Economic Review* 101, 2205–2225.

Schwert, William G., 1989, Why does stock market volatility change over time?, *Journal of Finance* 44, 1115–1153.

Stumpf, Michael P. H., and Mason A. Porter, 2012, Critical truths about power laws, *Science* 335, 665–666.

Toda, Alexis A., 2012, The double power law in income distribution: Explanations and evidence, *Journal of Economic Behavior and Organization* 84, 364–381.

Walsh, Kieran, and Alexis A. Toda, 2014, The double power law in consumption and implications for testing Euler equations, Forthcoming, *Journal of Political Economy.*

Weitzman, Martin L., 2007, Subjective expectations and asset-return puzzles, *American Economic Review* 97, 1102–1130.

White, Halbert, 1980, A heteroskedasticity-consistent covariance matrix estimator and a direct test for heteroskedasticity, *Econometrica* 24, 817–838.

Zhang, Lan, Per A. Mykland, and Yacine Aït-Sahalia, 2005, A tale of two time scales: Determining integrated volatility with noisy high-frequency data, *Journal of the American Statistical Association* 100, 1394–1411.

Appendix

Measurement of volatility using the TSRV method

This section provides a brief summary of the TSRV method developed in Zhang, Mykland, and Aït-Sahalia (2005) and Aït-Sahalia, Mykland, and Zhang (2011) that is used in the study to measure realized volatility at 15-minute intervals. For further details, see the above studies.

Let $\tilde{r}_t$ denote the true return process of a security and $\sigma_{t,15}$ denote the volatility of the true return process over a 15-minute window. Let $r^1_{t,15}$ denote the observed returns of the security over 1-second intervals within the 15-minute window. In the absence of market microstructure noise, one could estimate the volatility of the true return process based on the realized volatility of observed returns.

$$RV^{\text{all}}_{t,15} = \frac{1}{n}\Sigma\left(r^1_{t,15}\right)^2, \tag{A.1}$$

where the number of 1-second returns within a 15-minute window, n, equals 900.

The above estimate is biased in the presence of market microstructure noise, leading one to instead estimate realized volatility by measuring returns over a sparse interval, such as a 60-second or a 300-second window. Letting $r^{60}_{t,15}$ denote returns over 60 second windows, the corresponding realized volatility estimate is given by

$$RV^{\text{sparse}}_{t,15} = \frac{1}{\bar{n}}\Sigma\left(r^{60}_{t,15}\right)^2, \tag{A.2}$$

where $\bar{n}$ denotes the number of such returns. While this estimate is much less influenced by market microstructure noise, it involves a loss of information.

The method of Aït-Sahalia, Mykland, and Zhang (2011) incorporates additional information into the estimate of realized volatility by measuring returns at 60-second windows offset at k seconds. Denote the correspoding 60-second returns offset at k seconds by $r^{60,k}_{t,15}$. The realized volatility measure obtained from these returns is given by:

$$RV^{\text{sparse},k}_{t,15} = \frac{1}{\bar{n}}\Sigma\left(r^{60,k}_{t,15}\right)^2. \tag{A.3}$$

One can construct K such measures for $k \in 1, 2, ..., K$.

Aït-Sahalia, Mykland, and Zhang (2011) show that one could obtain a consistent estimate of the volatility of the true return process by taking the average of the above realized volatility values and adjusting for a bias-correction term involving the realized volatility measure obtained using 1-second returns. Their TSRV estimator is given by

$$TSRV_{t,15} = \frac{1}{K}\Sigma_{k=1}^{K} RV_{t,15}^{\text{sparse},k} - \frac{\bar{n}}{n} RV_{t,15}^{\text{all}}. \tag{A.4}$$

One limitation of the above estimator is that it convergences only at $n^{1/6}$. As such, this study measures realized volatility using the following small-sample refinement proposed by Aït-Sahalia, Mykland, and Zhang (2011):

$$TSRV_{t,15}^{\text{adj}} = \left(1 - \frac{\bar{n}}{n}\right) TSRV_{t,15}. \tag{A.5}$$

The analysis in the study employs the above method to measure volatility at 10, 15 and 30 minute intervals. The time period at which returns are offset, k, increments at one second. For each time interval, three estimates of TSRV are obtained using squared returns measured over 30, 45 and 60 second intervals (the corresponding K values equal 30, 45 and 60, respectively).

Table 1: Data moments

The table reports summary statistics for intraday stock returns and return volatility measured at 15 minute intervals. Panel A reports univariate statistics; Panel B reports the correlation coefficients among return volatility measured using the TSRV method of Zhang, Mykland, and Aït-Sahalia (2005); and Panel C reports the correlation coefficients among return volatility measured using the realized volatility approach. The sample consists of all 41 stocks that were included in the Dow Jones Industrial Average at some point from 01/01/2003 to 12/31/2014. The volatility measures for 15-minute returns reported in Panels A and B are obtained using the TSRV method applied to stock returns at 30, 45, and 60 second intervals, respectively, offset at 1 second intervals. The volatility measures for 15-minute returns reported in Panel C are obtained using the realized volatility method applied to stock returns at 30, 45, and 60 second intervals.

Panel A: Univariate statistics

	Mean	Median	Std. dev.
Stock return	5.63e-06	0.00	0.0030
Return volatility			
30 seconds	0.0025	0.0019	0.0021
45 seconds	0.0024	0.0019	0.0021
60 seconds	0.0023	0.0018	0.0020
Number of observations:		2796632	

Panel B: Correlation of volatility measures obtained using TSRV method

	Time interval for TSRV measure		
	30 seconds	45 seconds	60 seconds
30 seconds	1.000	0.994	0.984
45 seconds	0.994	1.000	0.997
60 seconds	0.984	0.997	1.000

Panel C: Correlation of realized volatility measures

	Time interval for returns		
	30 seconds	45 seconds	60 seconds
30 seconds	1.000	0.934	0.947
45 seconds	0.934	1.000	0.932
60 seconds	0.947	0.932	1.000

Table 2: Power law coefficients from data and nonparametric model

The table compares the power law coefficients obtained from the returns data with those obtained from the conditional normal model with a nonparametric volatility distribution for each of the 41 stocks that were included in the Dow Jones Industrial Average at some point from 01/01/2003 to 12/31/2014. The second column reports the power law coefficient for stock returns obtained from the 15-minute return data, the third column reports the corresponding 95 percent confidence interval, and the fourth to sixth columns report the coefficients obtained from the conditional normal model with volatility within 15 minutes measured using the TSRV method of Zhang, Mykland, and Aït-Sahalia (2005). The latter results are obtained by applying the TSRV method to stock returns at 30, 45, and 60 second intervals offset at 1 second intervals, respectively. * and **, respectively, denote observations for which the model implied power law coefficient falls outside the 95 and 99 percent confidence intervals for the power law coefficient from the return data.

Stock	Estimated power law coefficient	Confidence interval	Power law coefficient from model		
			30 seconds	45 seconds	60 seconds
AA	2.93	[2.79, 3.06]	2.95	2.93	2.90
AIG	2.29	[2.14, 2.45]	2.37	2.33	2.35
AXP	2.60	[2.48, 2.72]	2.62	2.59	2.60
BA	3.17	[3.03, 3.32]	3.06	3.09	3.09
BAC	2.28	[2.17, 2.38]	2.35	2.32	2.31
C	2.21	[2.11, 2.31]	2.36**	2.31	2.27
CAT	3.01	[2.87, 3.14]	2.92	2.94	2.95
CSCO	3.26	[3.11, 3.41]	3.28	3.28	3.26
CVX	3.00	[2.87, 3.14]	2.92	2.92	2.92
DD	3.10	[2.96, 3.24]	3.01	3.03	3.02
DIS	3.08	[2.93, 3.22]	3.02	3.03	3.03
EK	3.05	[2.88, 3.23]	3.26*	3.21	3.18
GE	2.50	[2.38, 2.61]	2.56	2.53	2.53
GM	2.09	[1.96, 2.22]	2.34**	2.30**	2.25*
GS	2.50	[2.39, 2.62]	2.51	2.50	2.50
HD	3.07	[2.92, 3.21]	3.00	3.00	2.98
HON	3.10	[2.96, 3.25]	3.01	3.02	3.02
HPQ	3.26	[3.11, 3.41]	3.22	3.22	3.22
IBM	2.93	[2.79, 3.06]	2.82	2.84	2.83
INTC	3.46	[3.30, 3.62]	3.38	3.38	3.38
IP	2.92	[2.79, 3.05]	2.83	2.84	2.83

Contd.

Table 2(contd.): Power law coefficients from data and nonparametric model

Stock	Estimated power law coefficient	Confidence interval	Power law coefficient from model		
			30 seconds	45 seconds	60 seconds
JNJ	3.00	[2.86, 3.14]	3.04	3.03	3.02
JPM	2.59	[2.47, 2.71]	2.59	2.58	2.56
KFT	3.10	[2.94, 3.26]	3.08	3.11	3.12
KO	3.17	[3.02, 3.32]	3.11	3.13	3.12
MCD	3.18	[3.03, 3.33]	3.04	3.06	3.08
MMM	3.06	[2.92, 3.21]	3.00	3.01	3.01
MO	2.76	[2.64, 2.89]	2.95^{**}	2.95^{**}	2.93^{**}
MRK	2.93	[2.79, 3.06]	2.94	2.94	2.94
MSFT	3.16	[3.01, 3.31]	3.26	3.23	3.19
NKE	3.14	[3.00, 3.29]	3.05	3.07	3.10
PFE	3.20	[3.05, 3.35]	3.29	3.27	3.25
PG	3.08	[2.94, 3.22]	3.02	3.03	3.04
T	2.89	[2.75, 3.02]	2.92	2.91	2.90
TRV	2.63	[2.48, 2.78]	2.57	2.57	2.57
UNH	2.78	[2.65, 2.91]	2.76	2.79	2.81
UTX	3.15	[3.00, 3.29]	2.98^{*}	2.98^{*}	2.98^{*}
V	2.83	[2.65, 3.00]	2.86	2.87	2.87
VZ	2.97	[2.84, 3.11]	2.96	2.96	2.94
WMT	3.21	[3.06, 3.35]	3.09	3.09	3.09
XOM	3.00	[2.86, 3.14]	2.96	2.96	2.95

Table 3: Monte Carlo analysis of model with power law returns

The table reports results from a Monte Carlo analysis of data obtained from a model that takes as given the observed distribution of 15-minute returns. The hypothetical volatility samples in each Monte Carlo simulation are constructed using a linear regression that matches the mean and standard deviation of volatility observed in the data, and the sensitivity of volatility to absolute returns. These volatility samples are used to compute power law coefficients implied by the conditional normal model with nonparametric volatility distribution. For each Monte Carlo simulation, the following cross-sectional statistics are obtained: mean, standard deviation, and the slope coefficient from a regression through the origin of the power law coefficients from the returns data on the model-implied coefficients. The reported results comprise the mean, the 5^{th} percentile, the median and the 95^{th} percentile of the above statistics from 400 such Monte Carlo simulations.

Simulation statistics	Mean	Percentile		
		5	50	95
Power law coefficient				
Mean	3.79	3.78	3.79	3.80
Standard deviation	0.39	0.38	0.39	0.40
Slope coefficient	0.77	0.77	0.77	0.77

Table 4: Monte Carlo analysis of conditional t-distribution model

The table reports results from a Monte Carlo analysis of a model where 15-minute returns are given by a condition Student's t-distribution with zero mean and volatility given by the empirical distribution observed in the data. Each Monte Carlo simulation computes the power law coefficients implied by such a model and derives the following statistics from these coefficients: mean, standard deviation, and the slope coefficient from a regression through the origin of the power law coefficients from the returns data on the model-implied coefficients. The reported results comprise the mean, the 5^{th} percentile, the median and the 95th percentile of the above statistics from 400 such Monte Carlo simulations.

Simulation statistics	Mean	Percentile		
		5	50	95
Power law coefficient				
Mean	2.04	2.04	2.04	2.04
Standard deviation	0.07	0.07	0.07	0.07
Slope coefficient	1.44	1.44	1.44	1.44

Figure 1: Tail distribution of stock returns

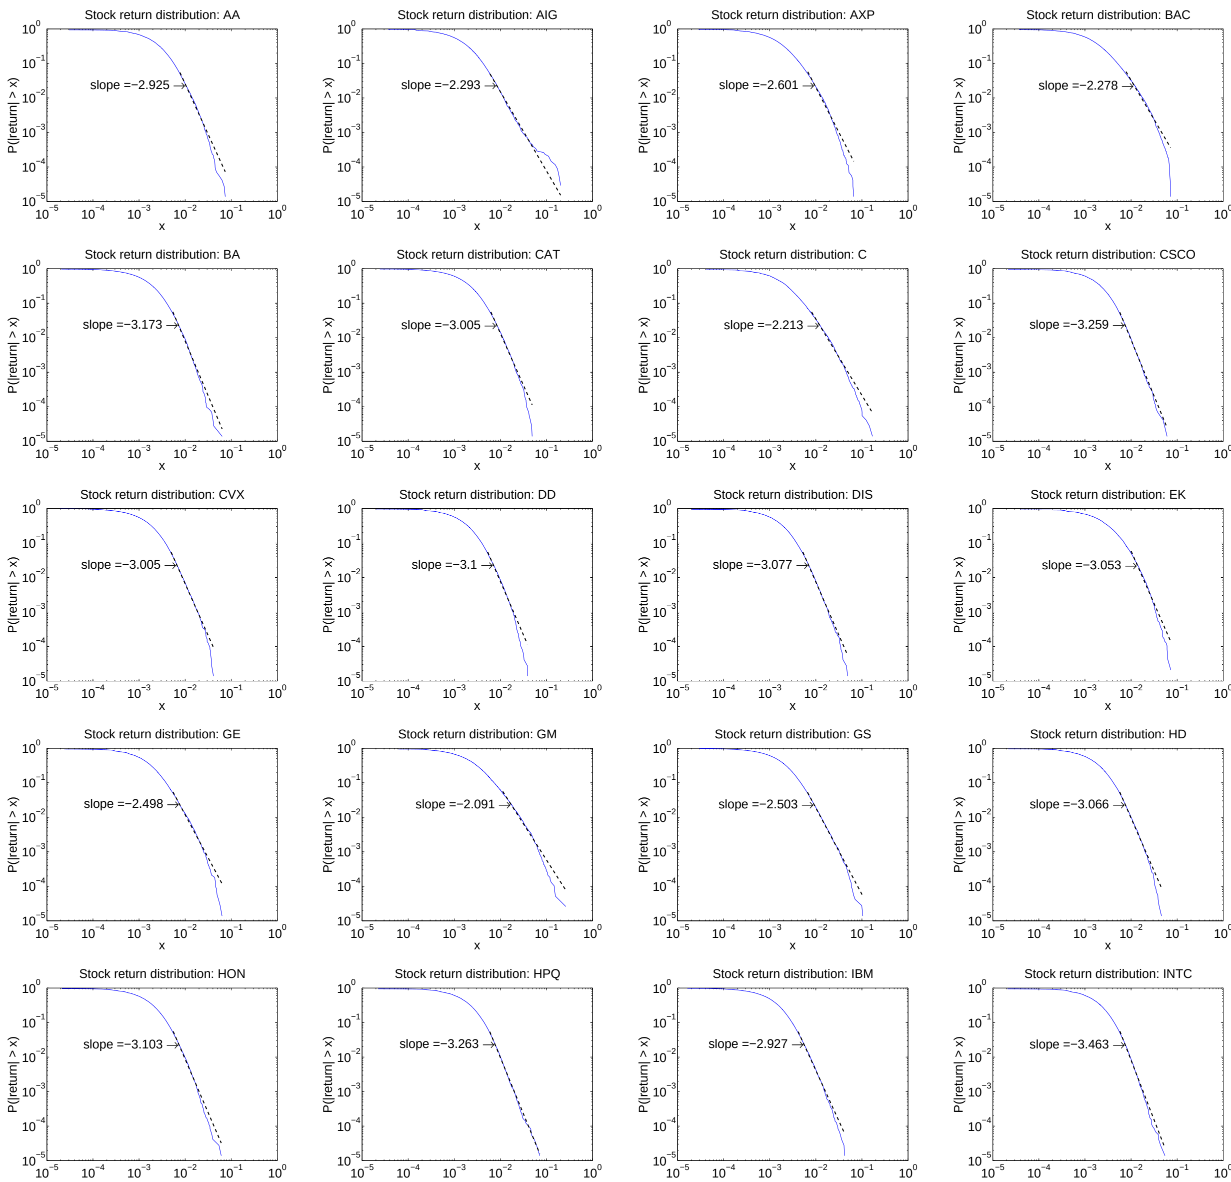

The figure plots the tail distribution of the absolute value of 15-minute stock returns from 01/01/2003 to 12/31/2014 for each of the stocks in the sample. The x-axis plots the absolute return and the y-axis plots the probability that the absolute return exceeds the given value. Both axes are on log scales. The figures also report the slope and fitted line obtained from a regression of the log tail probability on the log absolute return estimated on returns exceeding the 95th percentile of the empirical distribution for the absolute value of stock returns.

Figure 1(contd.): Tail distribution of stock returns

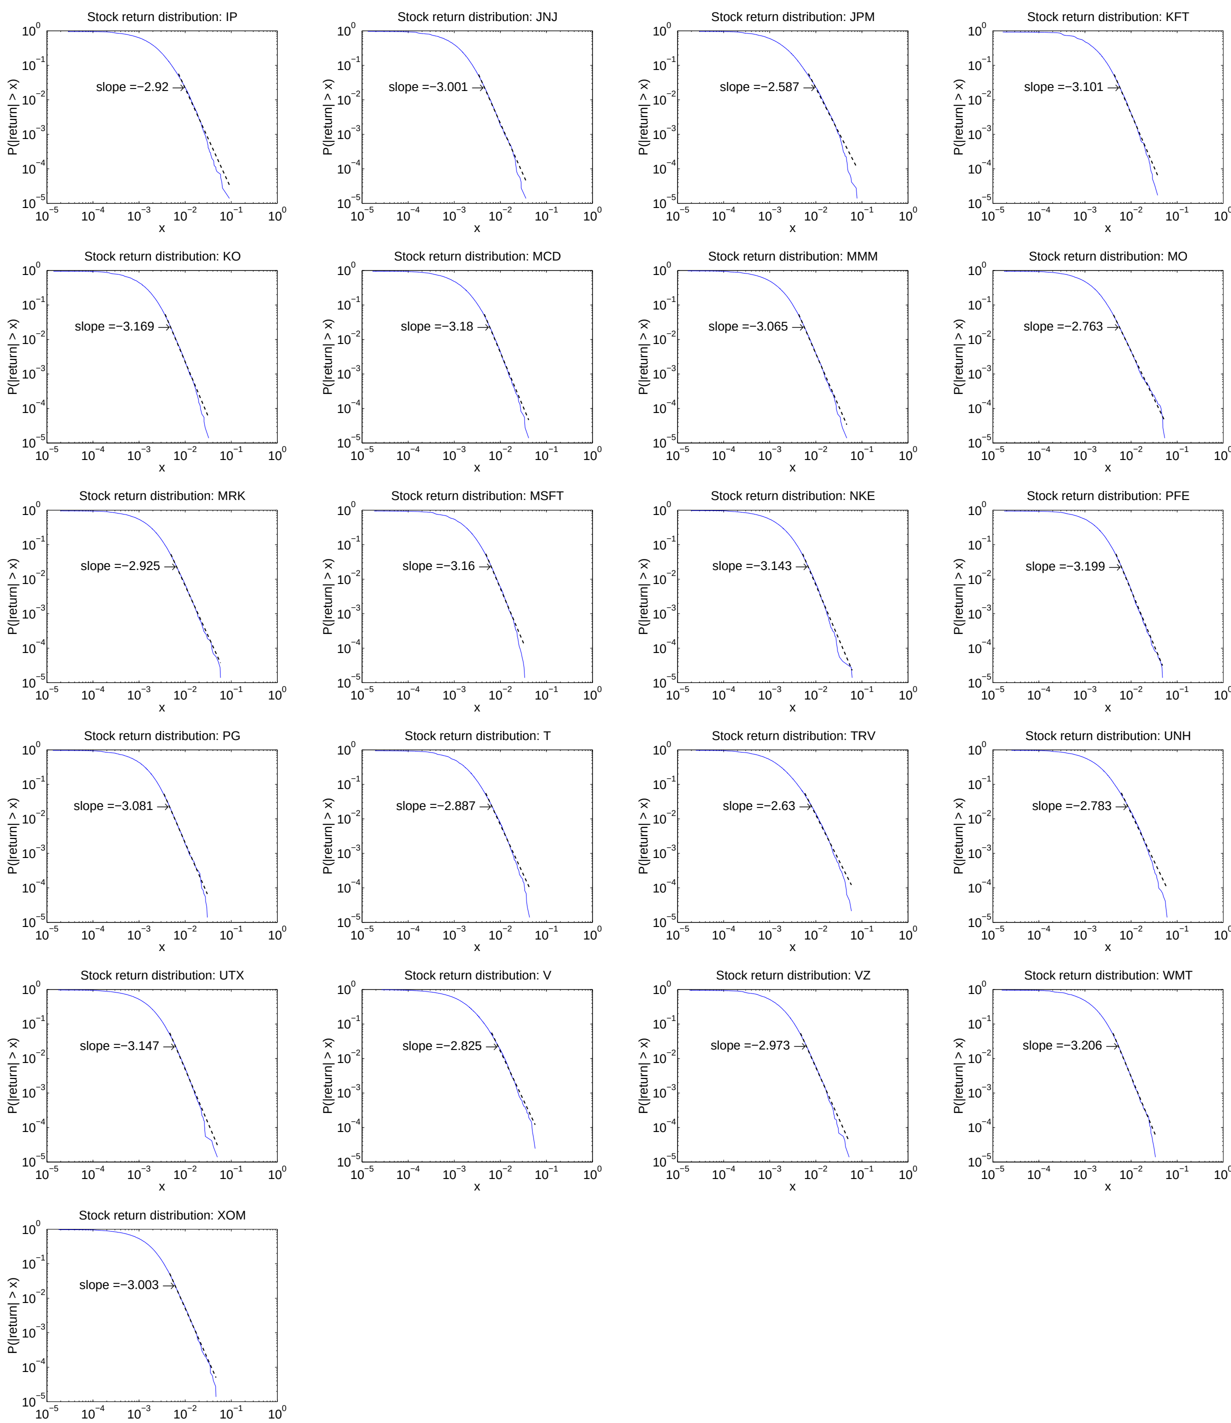

Figure 2: Distribution of power law coefficients

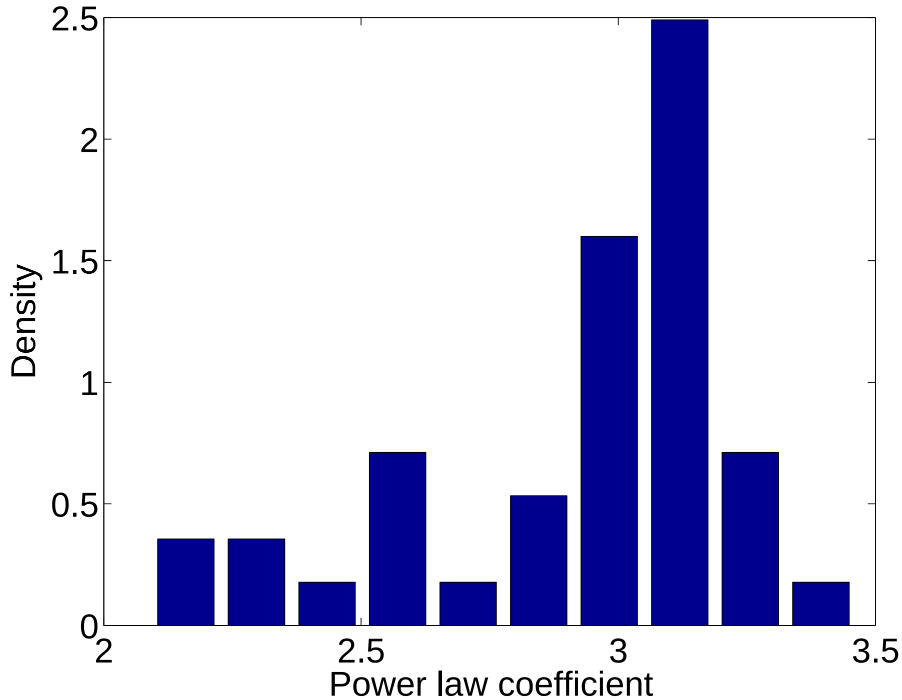

The figure reports the cross-sectional distribution of the power law coefficients for stock returns at 15 minute intervals for the 41 stocks that were, at some point, included in the Dow Jones Industrial Average from 01/01/2003 to 12/31/2014.

Figure 3: Comparison of power law coefficients: Data and parametric model

Panel A: Volatility measured using 30-second returns

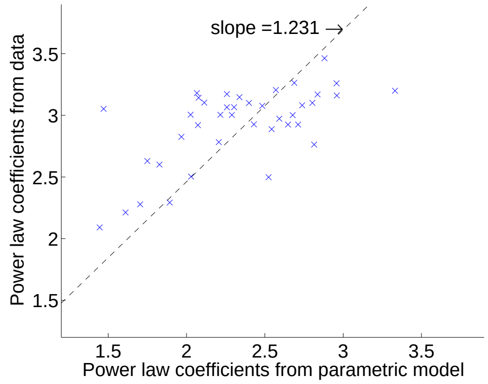

Panel B: Volatility measured using 45-second returns

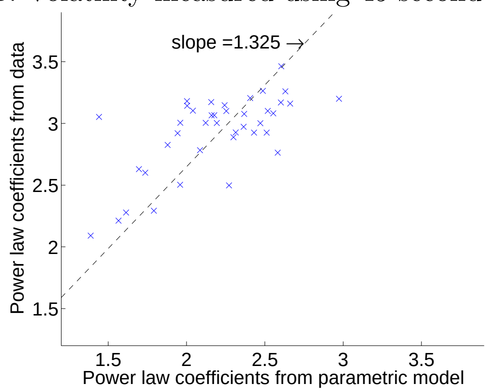

Panel C: Volatility measured using 60-second returns

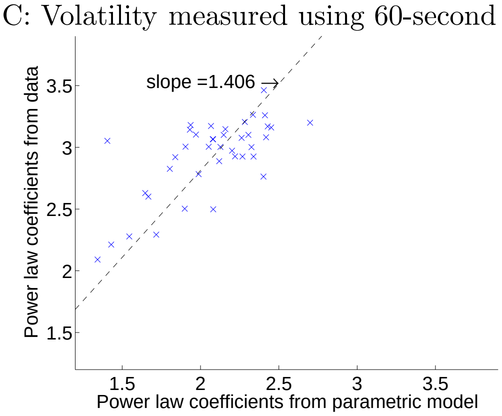

The figures present scatter plots of the power law coefficient for stock returns over 15 minute intervals plotted against twice the shape coefficient of the inverse-gamma distribution for stock return volatility within 15 minute intervals. The sample comprises the 41 stocks that were, at some point, included in the Dow Jones Industrial Average from 01/01/2003 to 12/31/2014. Panels A, B, and C, respectively, measure volatility using the TSRV method of Zhang, Mykland, and Aït-Sahalia (2005) applied to returns over 30, 45, and 60 seconds offset at 1 second. The dashed lines represent OLS estimates of a regression through the origin of the power law coefficients from the data on twice the shape coefficients of the inverse-gamma distribution.

Figure 4A: Comparison of power law coefficients: Data and nonparametric model

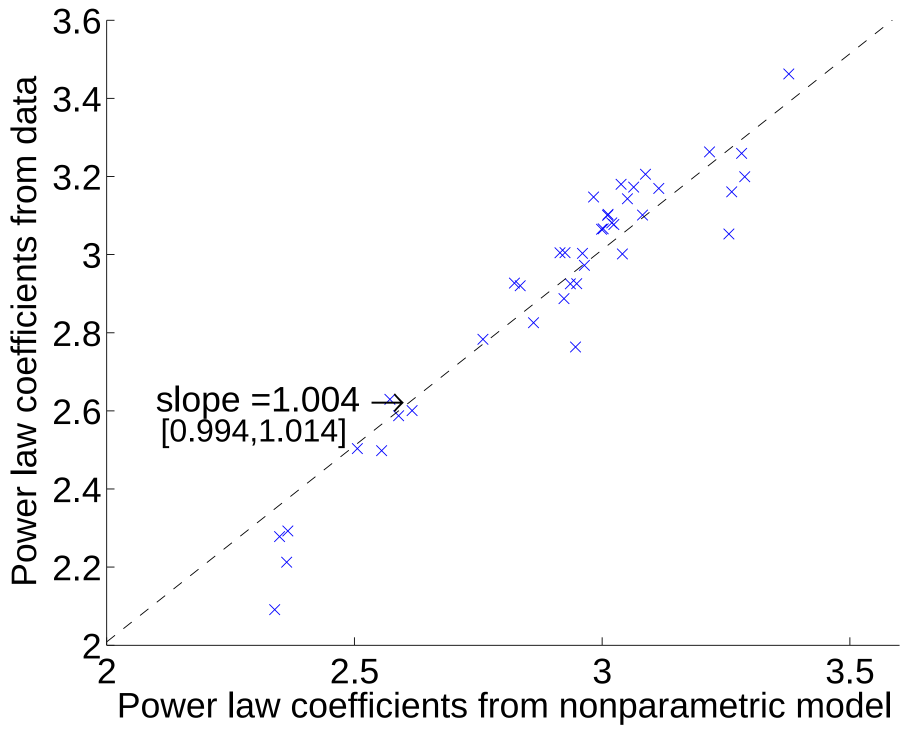

The figure presents a scatter plot of the power law coefficient for stock returns over 15 minute intervals plotted against the power law coefficient obtained from simulating the conditional normal model with a nonparametric volatility distribution for stock return volatility within 15 minute intervals. The sample comprises the 41 stocks that were, at some point, included in the Dow Jones Industrial Average from 01/01/2003 to 12/31/2014. Volatility is measured using the TSRV method of Zhang, Mykland, and Aït-Sahalia (2005) applied to returns over 30 second intervals offset at 1 second. The dashed line represents the OLS fit of a regression through the origin of the power law coefficient from the data on the power law coefficient from the simulated nonparametric model. The square brackets report the 95 percent confidence interval for the slope coefficient of the regression.

Figure 4B: Comparison of power law coefficients: Data and nonparametric model

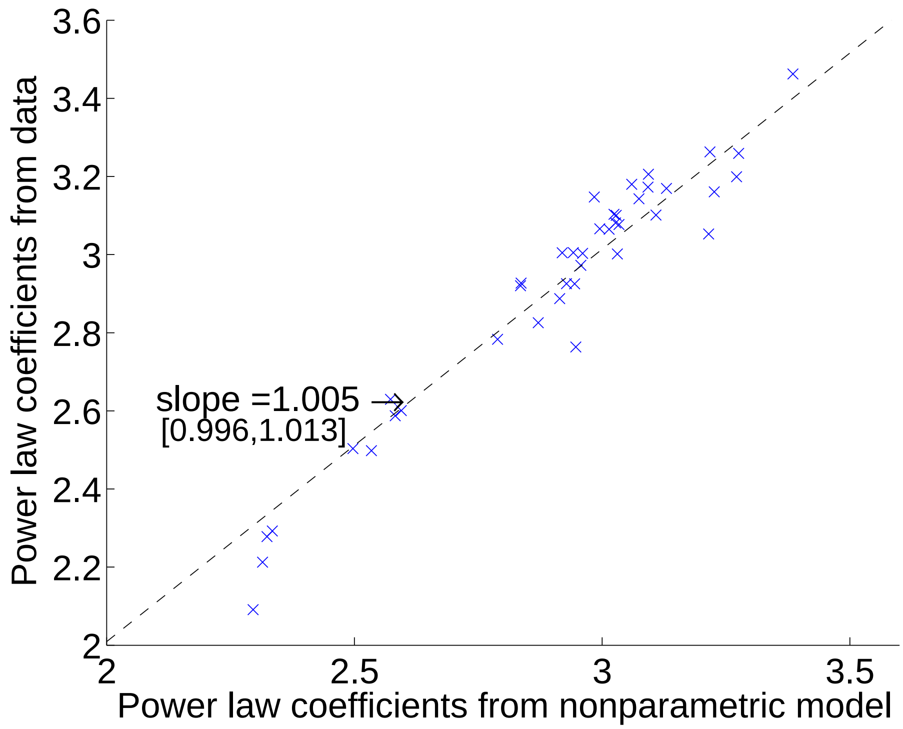

The figure presents a scatter plot of the power law coefficient for stock returns over 15 minute intervals plotted against the power law coefficient obtained from simulating the conditional normal model with a nonparametric volatility distribution for stock return volatility within 15 minute intervals. The sample comprises the 41 stocks that were, at some point, included in the Dow Jones Industrial Average from 01/01/2003 to 12/31/2014. Volatility is measured using the TSRV method of Zhang, Mykland, and Aït-Sahalia (2005) applied to returns over 45 second intervals offset at 1 second. The dashed line represents the OLS fit of a regression through the origin of the power law coefficient from the data on the power law coefficient from the simulated nonparametric model. The square brackets report the 95 percent confidence interval for the slope coefficient of the regression.

Figure 4C: Comparison of power law coefficients: Data and nonparametric model

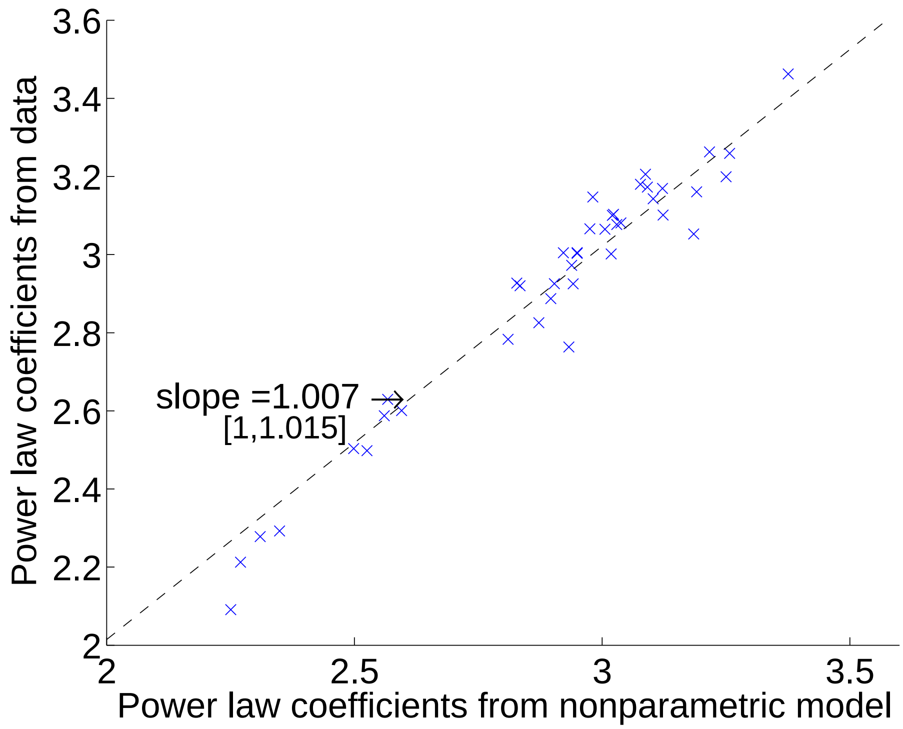

The figure presents a scatter plot of the power law coefficient for stock returns over 15 minute intervals plotted against the power law coefficient obtained from simulating the conditional normal model with a nonparametric volatility distribution for stock return volatility within 15 minute intervals. The sample comprises the 41 stocks that were, at some point, included in the Dow Jones Industrial Average from 01/01/2003 to 12/31/2014. Volatility is measured using the TSRV method of Zhang, Mykland, and Aït-Sahalia (2005) applied to returns over 60 second intervals offset at 1 second. The dashed line represents the OLS fit of a regression through the origin of the power law coefficient from the data on the power law coefficient from the simulated nonparametric model. The square brackets report the 95 percent confidence interval for the slope coefficient of the regression.

Figure 5: Comparison of data and nonparametric model with volatility clustering

Panel A: Volatility measured using 30-second returns

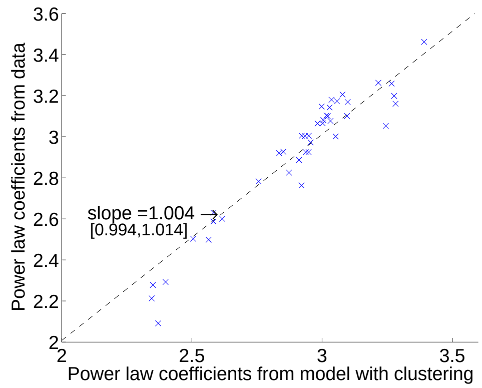

Panel B: Volatility measured using 45-second returns

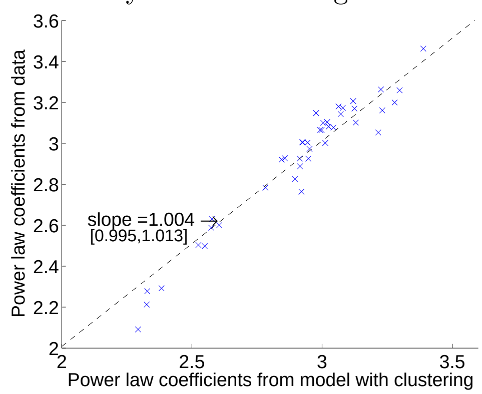

Panel C: Volatility measured using 60-second returns

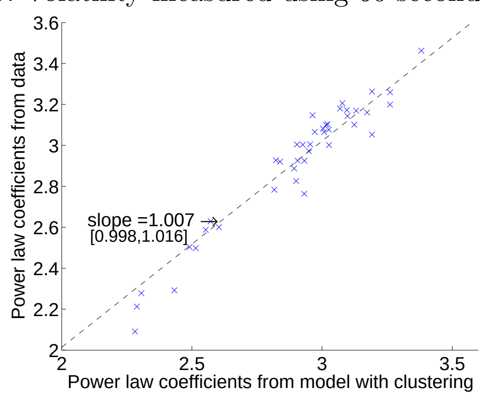

The figures present scatter plots of the power law coefficient for 15-minute stock returns plotted against the power law coefficient obtained from simulating a nonparametric conditional normal model with clustering by sampling volatility from blocks of 120 consecutive 15-minute volatility observations. The sample comprises the 41 stocks that were, at some point, included in the Dow Jones Industrial Average from 01/01/2003 to 12/31/2014. Panels A, B, and C, respectively, measure volatility using the TSRV method of Zhang, Mykland, and Aït-Sahalia (2005) applied to returns over 30, 45, and 60 seconds offset at 1 second. The dashed line represents the OLS fit of a regression through the origin.

Figure 6: Comparison of data and nonparametric model with 10 minute returns

Panel A: Volatility measured using 30-second returns

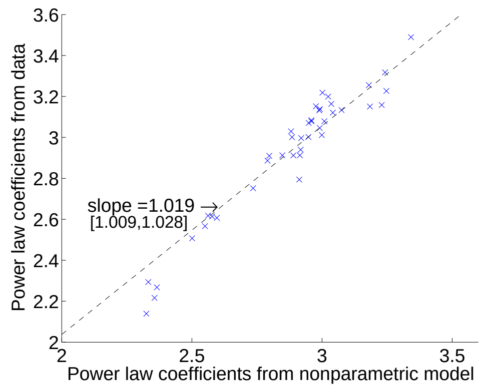

Panel B: Volatility measured using 45-second returns

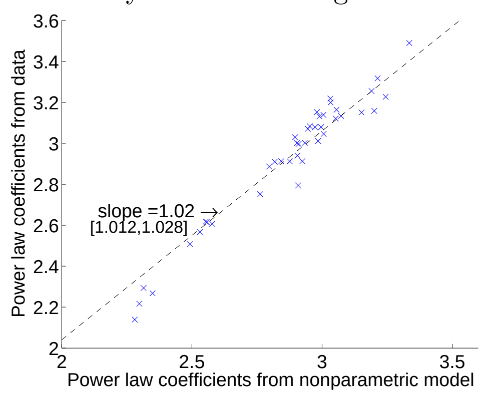

Panel C: Volatility measured using 60-second returns

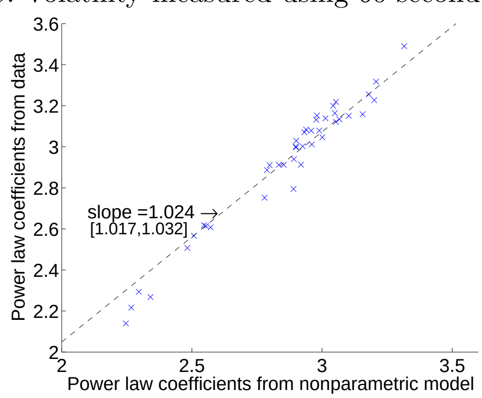

The figures present scatter plots of the power law coefficient for 10-minute stock returns plotted against the power law coefficient obtained from simulating the conditional normal model with a nonparametric volatility distribution. The sample comprises the 41 stocks that were, at some point, included in the Dow Jones Industrial Average from 01/01/2003 to 12/31/2014. Panels A, B, and C, respectively, measure volatility using the TSRV method of Zhang, Mykland, and Aït-Sahalia (2005) applied to returns over 30, 45, and 60 seconds offset at 1 second. The dashed line represents the OLS fit of a regression through the origin.

Figure 7: Comparison of data and nonparametric model with 30 minute returns

Panel A: Volatility measured using 30-second returns

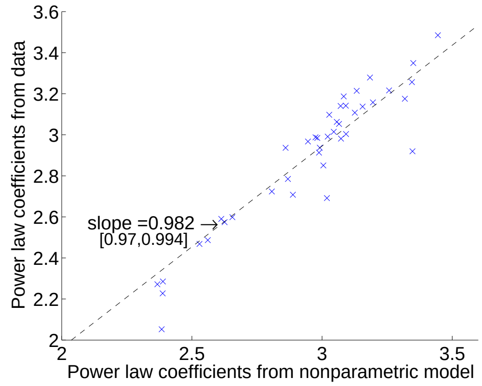

Panel B: Volatility measured using 45-second returns

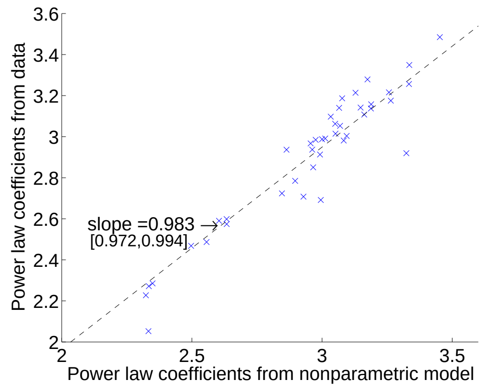

Panel C: Volatility measured using 60-second returns

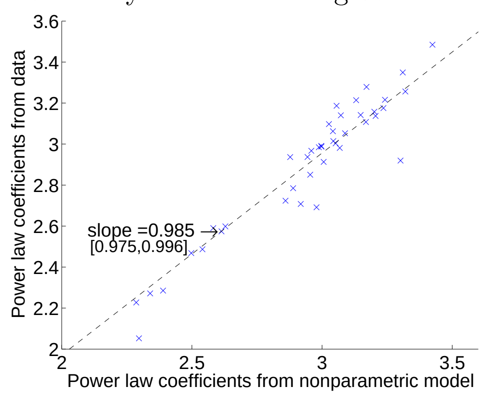

The figures present scatter plots of the power law coefficient for 30-minute stock returns plotted against the power law coefficient obtained from simulating the conditional normal model with a nonparametric volatility distribution. The sample comprises the 41 stocks that were, at some point, included in the Dow Jones Industrial Average from 01/01/2003 to 12/31/2014. Panels A, B, and C, respectively, measure volatility using the TSRV method of Zhang, Mykland, and Aït-Sahalia (2005) applied to returns over 30, 45, and 60 seconds offset at 1 second. The dashed line represents the OLS fit of a regression through the origin.