International Review of Financial Analysis 23 (2012) 30–34

Contents lists available at ScienceDirect

International Review of Financial Analysis

On the nonstationarity of the exchange rate process

Takaaki Ohnishi [a,b,*], Hideki Takayasu [c], Takatoshi Ito [b], Yuko Hashimoto [d], Tsutomu Watanabe [a,e], Misako Takayasu [f]

[a] The Canon Institute for Global Studies, 11F, Shin-Marunouchi Bldg., 1-5-1, Marunouchi, Chiyoda-Ku, Tokyo 100-6511, Japan
[b] Graduate School of Economics, University of Tokyo, 7-3-1 Hongo, Bunkyo-Ku, Tokyo, 113-0033, Japan
[c] Sony Computer Science Laboratories, 3-14-13 Higashigotanda, Shinagawa-ku, Tokyo, 141-0022, Japan
[d] Statistics Department, International Monetary Fund, 700 19th Street, N.W., Washington, D.C. 20431, United States
[e] Institute of Economic Research, Hitotsubashi University, 2-1 Naka, Kunitachi-city, Tokyo 186-8603, Japan
[f] Department of Computational Intelligence and Systems Science, Interdisciplinary, Graduate School of Science and Engineering, Tokyo Institute of Technology, 4259-G3-52, Nagatsuta-cho, Midori-ku, Yokohama 226-8502, Japan

ARTICLE INFO

Article history:
Received 15 November 2010
Received in revised form 10 May 2011
Accepted 30 June 2011
Available online 15 July 2011

Keywords:
Econophysics
Foreign exchange market
Strict stationarity
Nonstationarity
Two-sample Kolmogorov–Smirnov test
Pearson's chi-square test
Poisson process

ABSTRACT

We empirically investigate the nonstationarity property of the USD–JPY exchange rate by using a high frequency data set spanning 8 years. We perform a statistical test of strict stationarity based on the two-sample Kolmogorov–Smirnov test for the absolute price changes, and Pearson's chi square test for the number of successive price changes in the same direction, and find statistically significant evidence of nonstationarity. Further, we study the recurrence intervals between the days in which nonstationarity occurs and find that the distribution of recurrence intervals is well approximated by an exponential distribution. In addition, we find that the mean conditional recurrence interval hTjT0i is independent of the previous recurrence interval T0. These findings indicate that the recurrence intervals are characterized by a Poisson process. We interpret this observation as a reflection of the Poisson property regarding the arrival of news.

1. Introduction

In econophysics, financial time series data have been extensively investigated using a wide variety of methods. These studies tend to assume, explicitly or implicitly, that a time series is stationary, since stationarity is a requirement for most of the mathematical theories underlying time series analysis. However, despite its nearly universal assumption, few previous studies seek to test stationarity in a reliable manner (Tóthla, et al., 2010).

For low-frequency financial data (i.e., monthly or daily data), a number of procedures to test stationarity have been advocated and applied to various time series processes in econometrics. Most of them focus on the first two moments of a process; in other words, they test covariance stationarity. These tests work well for normally distributed random variables. However, for high-frequency financial data such as tick-by-tick data, it is well known that price change distributions are fat-tailed and substantially deviate from a normal distribution. These fat-tailed distributions cannot be dealt with by the above stationarity tests.

In this paper, we advocate a test for strict stationarity that considers the entire distribution of a process rather than the first two moments of the process, and apply this test to the USD–JPY exchange rate.

We describe the data used in this paper in Section 2. In Section 3, we explain our procedure to test stationarity, which is based on the two-sample Kolmogorov–Smirnov test and Pearson's chi-square test. In Section 4, we present the empirical results. In Section 5, we discuss some implications of our results.

2. Data description

The tick-by-tick data we study is the USD–JPY exchange rate provided by ICAP EBS with a recording frequency of every 1 s, for the period of January 1998 through December 2005. The foreign exchange market is the world's largest and liquid financial market. Most spot interbank transactions are executed through global electronic broking systems such as ICAP EBS and Reuters. In the USD–JPY exchange rate, the ICAP EBS has a strong market share.

We exclude observations for special days such as Mondays, weekends, holidays, and official intervention days (i.e., the government

* Corresponding author at: The Canon Institute for Global Studies, 11F, Shin-Marunouchi Bldg., 1-5-1, Marunouchi, Chiyoda-Ku, Tokyo 100-6511, Japan.
E-mail address: ohnishi.takaaki@canon-igs.org (T. Ohnishi).

1057-5219/$ – see front matter © 2011 Elsevier Inc. All rights reserved.
doi:10.1016/j.irfa.2011.06.010

and/or the central bank intervenes in the foreign exchange market in order to stabilize the rate), which are obviously different from regular business days. We analyze the time series of 1-tick price changes of the mid-quote price, which is defined as the average of the best bid and the best ask. The best bid and the best ask, representing the lowest sell offer and highest buy offer respectively, are recorded at the end of one-second time slice.

In this paper, we focus on the following two time series: The first one is the time series for the absolute price changes, which we refer to as G; second, the time series for the number of successive price changes in the same direction, which we refer to as D. Note that in producing these time series, we drop observations with no price changes. For example, a particular sequence of 14 1-tick price changes

$$\{0.01, 0.02, 0.01, -0.02, 0, -0.03, -0.01, 0.02, 0, 0.02, -0.04, 0.01, -0.02, -0.03\}$$

is represented by

$$\{0.01, 0.02, 0.01, 0.02, 0, 0.03, 0.01, 0.02, 0, 0.02, 0.04, 0.01, 0.02, 0.03,\}$$

in G sequence and

$$\{3,\ 3,\ 2,\ 1,\ 1,\ 2\}$$

in D sequence.

3. Stationarity test

One can test stationarity in Gaussian time series processes by measuring any number of simple statistics such as the mean or standard deviation and employing a standard statistical test. However, such an approach is not particularly effective for high-frequency financial time series, because from the seminal work by Mantegna and Stanley(1995), we know that the distributions of price changes have fat tails often approximated by a power law (Ohnishi et al., 2008). Therefore, the procedure for Gaussian processes cannot be applied to high-frequency financial data.

Our analysis is based on a precise definition of stationarity: The joint distribution of any two segments of data of the same length should be identical. Formally, a stochastic process X_t is called strictly stationary if for any set of times $t_1, t_2, \ldots, t_n$ and for any k, the joint probability distributions of $\{X_{t_1}, X_{t_2}, \ldots, X_{t_n}\}$ and of $\{X_{t_1+k}, X_{t_2+k}, \ldots, X_{t_n+k}\}$ coincide. That is, it requires that the joint distribution depends only on time lags. It follows that the mean remains constant, and that the autocorrelation function depends on only time lags, and not on the time index.

Given this definition of stationarity, the test of stationarity may look straightforward: namely, all we have to do is to pick up any two segments of data of the same length, and then to see whether the distributions of G and D are identical across the two segments. However, the test of stationarity is not so simple since the exchange rate exhibits a strong seasonality. It is well known by practitioners and researchers that trading occurs differently even within a day, depending on, for instance, whether it is conducted in the morning or afternoon session. Specifically, we know from the previous studies that the absolute price changes (Ohnishi et al., 2008) and activities (Ito & Hashimoto, 2006) display an intraday seasonality. Figs. 1 and 2 present the cumulative distributions of G and D respectively, showing clearly that these distributions differ depending on the hour of the day.

One may want to apply some traditional methods of seasonal adjustment, such as X-12-ARIMA, in order to eliminate this intraday seasonality. However, such methods may not be appropriate in this context, since the time series property of G and D may be altered substantially by applying these methods. To avoid such risk, we eliminate the intraday seasonality in a different way. First, we assume that the time series can be regarded as approximately stationary at least during the one-hour period. Second, on the basis of this assumption, we divide the entire time series into the subsets with one-hour periods, each of which is identified by hour $h = 0, 1, \ldots, 23$ and day t. Third, we compare the distributions of G and D for the subset (h, t) (i.e., the set of observations belonging to hour h of day t) with those for the subset (h, t') (i.e., the set of observations belonging to hour h and day t'), as illustrated schematically in Fig. 3. Note that we compare the observations belonging to the *same* hour h, although they come from different days. In this way, we conduct the stationarity test separately for each h $(h = 0, 1, \ldots, 23)$.

To test for stationarity, we compare the distribution of observations belonging to the subset (h, t) and that belonging to the subset (h, t') to examine whether the two distributions are identical. We perform tests by using the two-sample Kolmogorov–Smirnov test for continuous distributions of G and Pearson's chi-square test for discrete distributions of D. The stationarity is determined at a conventional significance level of 5%. The two-sample Kolmogorov–Smirnov test compares two cumulative distribution functions of G; then, the maximum difference between these two cumulative distribution functions yields the P-value. Pearson's chi-square test is performed by considering a histogram of D having 4 bins, that is, $D = 1, D = 2, D = 3$, and $D \geq 4$. These two tests have the advantage of being nonparametric, and without making assumptions about the distribution function of the data, we get the probability that the two sets of data are drawn from the same distribution.

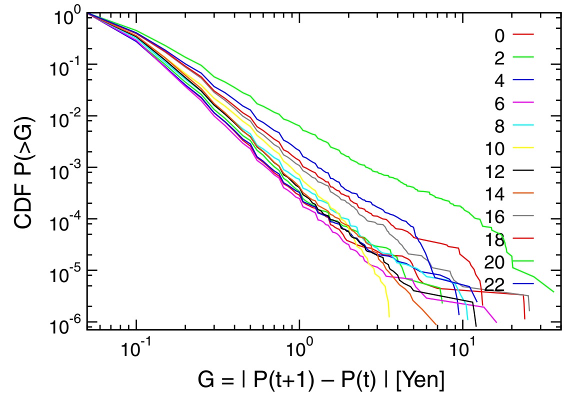

Fig. 1. Cumulative probability distributions of absolute price changes G. The colors represent the different hours of the day.

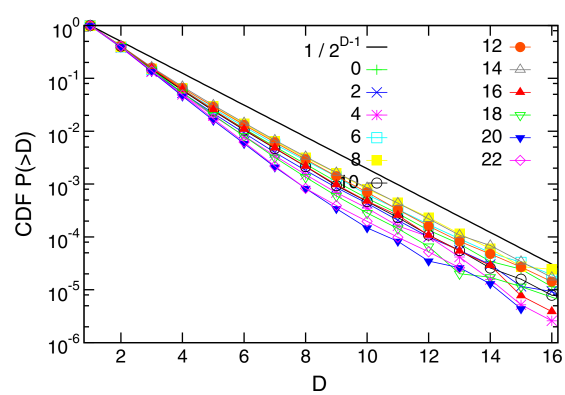

Fig. 2. Cumulative probability distributions of the number of successive price changes in the same direction D. The colors represent the different hours of the day.