International Journal of Economics and Financial Issues
Vol. 3, No. 3, 2013, pp.723-733
ISSN: 2146-4138
www.econjournals.com

Mean Aversion in and Persistence of Shocks to the US Dollar: Evidence from Nine Foreign Currencies

Samih Antoine Azar
Faculty of Business Administration & Economics,
Haigazian University, Mexique Street, Kantari, Beirut, Lebanon.
E-mail: samih.azar@haigazian.edu.lb
Telephone/fax: +9611349230

ABSTRACT: This paper analyses the statistical behavior of the US dollar, against nine different currencies, over the float period, with a monthly data set. The martingale hypothesis is rejected for all currencies. However, all currencies have a unit root. There is overwhelming evidence for significant positive serial correlation and ARCH effects in the logged series and in the changes of the logs. In addition, the nine logged series are characterized by structural breaks in both the intercept and the slope. Surprisingly, when the changes of the logs are considered, these do not show up any structural breaks, although the sample period has witnessed more than one political, social, and economic regime. All the nine logged currencies are well described by a low-order ARIMA model, with a parcimonious GARCH specification of the conditional variance. These ARIMA models imply mean aversion, rather than mean reversion, and high persistence of shocks. Mean aversion and persistence of shocks are formally tested and found to be significant for all nine currencies. This justifies the title of the paper.

Keywords: US dollar; mean aversion; persistence of shocks; market efficiency; martingale; structural breaks; ARIMA; GARCH
JEL Classifications: C12; C13; C22; F31; G14

1. Introduction

The purpose of this paper is to study the statistical behavior of the US dollar against nine different currencies on a monthly basis for the float period (1971/2013). The topic is of interest to a wide audience, including policy makers, academicians, practitioners, investors, and the lay man. The original feature of this paper is in the use of monthly data over a long span. This is imperative especially because most of the research on the subject has used either weekly or daily data (Baillie and Bollerslev, 1989; Hsieh, 1989; Liu and He, 1991; Van De Gucht et al., 1996; Chang, 2004; Ibrahim et al., 2011; Abdalla, 2012; Verbeek, 2012). Very few consider monthly data. One exception is Baillie and Bollerslev (1989). Moreover the literature on spot foreign exchange rates is scant, and little has been published recently on the issue. This paper tries to remedy for this paucity.

The Efficient Market Hypothesis states that all asset prices, like foreign exchange rates, should be unpredictable or uncorrelated (Fama, 1965, 1970, 1991). Therefore the starting point is to test whether (log) foreign exchange rates follow a random walk or a martingale. The evidence points to the contrary. The rejection can be due to serial correlation and/or heteroscedasticity, and/or the presence of high persistence of shocks. Log returns are calculated, by taking the first difference of the logs. These log returns measure the continuously compounded rates of change. Application of the Ljung-Box Q-statistics on the demeaned log returns provides support to the existence of both serial correlation and ARCH effects, or conditional heteroscedasticity (Engel, 1982). Serial correlation is significant for all nine foreign currencies. However ARCH effects are absent for two currencies, the Swiss franc and the Danish krone. The variance ratios, from the variance ratio tests, are all significantly higher than 1, indicating that the statistical problem is that of pervasive positive, and not negative, serial correlation. This hints towards mean aversion and persistence of shocks.

Augmented Dickey-Fuller (ADF) tests are carried out on the logged currencies (Dickey and Fuller, 1979). The choice of the ADF test is dictated by the need to find all the autoregressive terms that have to be included and this is done by minimizing the Schwarz information criterion (Schawrz, 1978). The null of a unit root fails to be rejected for all nine currencies. A characteristic of the results is that the log returns of all these nine currencies are positively auto-correlated. One autoregressive lag is sufficient for all currencies, except for the Swedish krona which requires two autoregressive lags. Hence there is a unit root in each currency but log returns are nonetheless non-random. This would seem to reject market efficiency, but may be due to risk aversion and risk premiums, to official intervention by central banks, or to undershooting (Dornbusch, 1976), among other reasons.

Since the data span such a long period of time structural breaks may exist. In fact the Perron (1997) test fails to reject the null hypothesis of a unit root that has a structural break in both the intercept and the trend. The estimated breaks are recovered from these tests. In order to find out whether the recovered breaks apply to the log returns, in addition to the logged variables, Chow break point tests are carried out on the log returns. The results show that there are no breaks in these log returns, except for the Canadian dollar. Further tests on the Canadian dollar find that the recovered breakpoint is also insignificant when the conditional variance is modeled as a GARCH(1,1). See Bollerslev (1986) on the GARCH model. Moreover, if the Quandt-Andrews tests for an unknown break, Andrews (1993) and Andrews and Ploberger (1994), are applied to the nine currency log returns, the data show no breaks at all. This is true for all the six different test statistics. This is surprising because the estimation period is long and encompasses many different economic, social, political, and policy regimes.

Having established positive serial correlation and conditional heteroscedasticity in the log returns these are regressed on their lagged values. One lag is enough for all currencies, except for the Swedish krona which requires two lags for obtaining white noise residuals. The conditional mean regression is supplemented with a conditional variance equation. In five currencies the model of the conditional variance is integrated of order 1, i.e. it is IGARCH. From the results the standardized residuals are retrieved, by dividing the regression residuals by the conditional standard deviation, and these are tested for serial correlation, further conditional heteroscedasticity, BDS independence (Brock et al., 1996), and normality. In general the standardized residuals are identically and independently distributed, are BDS independent, but normality is rejected for most currencies. From these regressions mean aversion and persistence of shocks are tested, following the footsteps of Campbell and Mankiw (1987a, 1987b). There is strong evidence for mean aversion and for persistence of shocks. This explains the reason behind the title of this paper.

The paper is organized as follows. The first part of the following section, section 2, provides for the source of the data. Next, in that same section, the tests that are chosen are each described and explained, and their implementations are derived, stated and interpreted. Section 2 is subdivided into 6 subsections. Section 3 is a conclusion.

2. The Empirical Results

2.1 Source of the data

The data is taken from the web site of the Federal Reserve Bank of Saint Louis, and span the monthly period from January 1971 to January 2013, i.e. 505 observations per variable. Nine currencies are selected and these are the Australian dollar (AUD), the Canadian dollar (CAD), the Swiss franc (CHF), the Danish krone (DKK), the British pound (GBP), the Norwegian krone (NOK), the New Zealand dollar (NZD), the Swedish krona (SEK), and the Japanese yen (JPY). The notations of these nine foreign currencies follow conventional international standards. Three currencies are quoted in terms of the number of units of the US dollar per one unit of the foreign currency, and these are the AUD, the GBP, and the NZD. The other six series are quoted in terms of the number of units of the foreign currency per one unit of the US dollar. Since all the series are logged, and log returns are calculated, the problem of Jensen's inequality, or of Siegel's paradox, does not arise (Azar, 2008).

2.2 Variance ratio tests

The variance ratio test is for the null hypothesis of a martingale. A martingale is defined as a process whereby the expected future value of the variable, based on current information, is equal to its current value. Another term for such a process is a random walk. If Z is the foreign exchange rate, α is a constant, ε is a well-behaved residual, and E is the expectation operator for information at time t, then a (sub)martingale adheres to the following relation in period $t+1$:

$$\log(Z_{t+1}) = \alpha + \log(Z_t) + \varepsilon_{t+1} \quad \Rightarrow \quad \mathrm{E}(\log(Z_{t+1})) = \alpha + \log(Z_t) \tag{1}$$

The reference to this test is Lo and MacKinlay (1988) which allows for heteroscedastic robust standard errors. The test is a joint test on the maximum absolute value of a z-statistic under the normal distribution. The actual p-values are reported in Table 1, column 2. Obviously the null hypothesis of a martingale is easily rejected at marginal significance levels less than 1%. The joint Fisher combined test on the stacked logs of the nine foreign currencies also rejects a martingale. If the distribution is bootstrapped from the normal distribution (Kim, 2006) the resulting p-values are immaterially different (Table 1, column 3). The variance ratios are then calculated as follows:

$$\frac{variance(\log(Z_{t+k+1})-\log(Z_t))}{variance(\log(Z_{t+1})-\log(Z_t))}\frac{1}{k+1} \qquad (2)$$

Table 1. Variance ratio test statistics on $\log(Z)$. The null hypothesis is a martingale.

Variable Z	Joint test for maximum $\lvert z-statistic \rvert$ (Lo and MacKinlay, 1988)	Joint test for maximum $\lvert z-statistic \rvert$ using wild bootstrap from normal distribution (Kim, 2006)	Variance ratio for period:				
			2	6	12	18	48
AUD	0.0000	0.0002	1.3385	1.6247	1.7089	1.6953	1.7203
CAD	0.0003	0.0014	1.2550	1.6442	1.7584	1.7191	2.2677
CHF	0.0000	0.0000	1.2870	1.5259	1.6000	1.6464	1.5385
DKK	0.0000	0.0000	1.3181	1.6481	1.8151	1.9516	2.2335
GBP	0.0000	0.0000	1.3588	1.7376	1.7572	1.7816	1.8308
NOK	0.0000	0.0000	1.3651	1.6845	1.6547	1.5267	1.7015
NZD	0.0000	0.0000	1.3570	1.8100	2.1351	2.2050	2.3588
SEK	0.0000	0.0002	1.3945	1.8495	2.0015	2.0392	2.1742
JPY	0.0000	0.0000	1.3172	1.5996	1.6511	1.7983	1.5101
All 9 stacked series (Fisher combined)	0.0000						

Notes: AUD stands for the Australian dollar, CAD for the Canadian dollar, CHF for the Swiss franc, DKK for the Danish krone, GBP for the British pound, NOK for the Norwegian krone, NZD for the New Zealand dollar, SEK for the Swedish krona, and JPY for the Japanese yen. The sample goes from January 1971 to January 2013, i.e. 505 monthly observations per variable. The specified lags for the two joint tests are from 2 to 18 with a step equal to 1. Bootstrapping is carried out with 5,000 replications. All standard errors are heteroscedastic robust. Actual p-values for the maximum absolute z-statistics are reported in columns 2 and 3.

This variance ratio is exactly equal to +1 for a martingale. In Table 1 the values of these variance ratios are reported for k that takes the values 2, 6, 12, 18, and 48 months. All variance ratios are higher than 1, giving support to the hypothesis that there is significant positive serial correlation. This is true because a given variance ratio is a weighted-average of the autocorrelation coefficients. Variance ratios, which are significantly above 1, do not only denote deviations from a random walk, but are also interpreted to denote high persistence of shocks to the underlying variables. Since the standard errors are heteroscedastically robust, nothing can be inferred yet about the existence of heteroscedasticity.

2.3 Ljung-Box Q-statistics

Since serial correlation is expected to prevail, and to see whether there are ARCH heteroscedastic effects, the demeaned log returns are calculated, and the Ljung-Box Q-statistics are applied on these demeaned log returns and their squares (Ljung and Box, 1978). Since the data is monthly four values for the lag length are used: 6, 12, 18, and 24. The results are depicted in Table 2. Whatever the lag length the null hypotheses of no serial correlation are easily rejected for all 9 currencies at marginal significance levels less than 1%. Therefore the evidence in support of serial correlation in log returns is quite strong. It must be mentioned that the Ljung-Box Q-statistic is a weighted-average of the squares of the autocorrelation coefficients, which implies that the autocorrelation coefficients are relatively large in absolute values. Since it is already known that the serial correlations are positive from the previous subsection, then the conclusion is strong that log returns have positive and relatively large autocorrelation coefficients.

Next the Q-statistics on the squares of the demeaned log returns are computed. The same four different lag lengths are also selected: 6, 12, 18, and 24. The results in Table 2 reject homoscedasticity for seven currencies out of nine. The Swiss franc and the Danish krone are not tainted by ARCH effects. While serial correlation is evidence of linear dependence, ARCH effects are evidence of non-linear dependence. These non-linear effects do not contradict the Efficient Market Hypothesis, but the serial correlation does, even though the latter can be rationalized by risk aversion, risk premiums, failure of arbitrage or limits to arbitrage, official intervention, undershooting, persistence in differential inflation rates, or transactions costs.

Table 2. Serial correlation and ARCH effects on the demeaned $\Delta(\log(Z))$.

Variable Z	Q-statistics on the demeaned variables				Q-statistics on the squares of the demeaned variables			
	K=6	K=12	K=18	K=24	K=6	K=12	K=18	K=24
AUD	0.000	0.000	0.000	0.000	0.000	0.000	0.002	0.012
CAD	0.000	0.000	0.000	0.000	0.018	0.000	0.000	0.000
CHF	0.000	0.000	0.000	0.000	0.441	0.846	0.777	0.785
DKK	0.000	0.000	0.000	0.000	0.057	0.132	0.302	0.364
GBP	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
NOK	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
NZD	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
SEK	0.000	0.000	0.000	0.000	0.000	0.000	0.006	0.032
JPY	0.000	0.000	0.000	0.000	0.000	0.001	0.004	0.017

Notes: See Table 1 for the definition of the variables. All Q-statistics are Ljung-Box Q-statistics (Ljung and Box, 1978). Actual p-values are reported.

2.4 Unit root tests

Although there are more recent tests for unit roots than the Augmented Dickey-Fuller test, or ADF test, this test is picked up because it automatically reports the optimal number of lags in the log returns that must be included by minimizing the Schwarz information criterion (Schwarz, 1978). This criterion identifies one lag for all log returns, but two lags for the Swedish krona. Table 3 presents the results.

Table 3. Augmented Dickey-Fuller unit root tests, with a constant and a trend, on $\log(Z)$ (Dickey and Fuller, 1979). The dependent variable is $\Delta(\log(Z))$.

Variable Z	Coefficient on				
	intercept	$\log(Z_{t-1})$	$\Delta(\log(Z_{t-1}))$	$\Delta(\log(Z_{t-2}))$	trend
AUD	-0.001310	-0.006510	0.335978	-	7.30E-07
	(0.5605)	(0.8971)	(0.0000)		(0.9376)
CAD	0.002236	-0.004324	0.246884	-	-5.71E-06
	(0.0922)	(0.9460)	(0.0000)		(0.1859)
CHF	0.019269	-0.022264	0.287026	-	-4.10E-05
	(0.0141)	(0.1018)	(0.0000)		(0.0194)
DKK	0.027466	-0.014215	0.320179	-	-4.50E-06
	(0.0270)	(0.4445)	(0.0000)		(0.5585)
GBP	0.012153	-0.019365	0.361702	-	-7.47E-06
	(0.0203)	(0.1885)	(0.0000)		(0.3468)
NOK	0.030006	-0.016649	0.368729	-	3.23E-06
	(0.0138)	(0.3214)	(0.0000)		(0.6478)
NZD	-0.003070	-0.005748	0.352056	-	2.82E-06
	(0.1762)	(0.8926)	(0.0000)		(0.7667)
SEK	0.017667	-0.010550	0.439925	-	9.04E-06
	(0.0574)	(0.6930)	(0.0000)		(0.3650)
JPY	0.100096	-0.018088	0.327357	-0.114545	-4.34E-05
	(0.0123)	(0.2808)	(0.0000)	(0.0106)	(0.0334)

Notes: See Table 1 for the definition of the variables. In parenthesis are actual p-values. The actual p-values are all for the ordinary t-distribution, except for the coefficients on $\log(Z_{t-1})$ for which the p-values are based on the Dickey-Fuller t-distribution.

The minimum actual Dickey-Fuller p-value on the lagged logs of each currency is 0.1018, and this is for the Swiss franc. Hence the presence of a unit root is strongly supported. The coefficients on the lagged log returns are all positive and hover between 0.213 and 0.440. Non-reported ADF tests on the second difference of the logs, or the first difference of the log returns, reject strongly the hypotheses of a unit root. Non-reported Phillips and Perron (1988) tests, which are robust in the presence of heteroscedasticity, that is a feature of the data as evidenced from the previous subsection, also endorse the fact that the logged series are integrated of order 1, or I(1). Therefore all these tests ascertain that the logged series are integrated of order 1, or I(1), but that log returns are nonetheless non-random. This is already discernible from the variance ratio tests on the logged series and the Ljung-Box Q-statistics on the log returns. This kind of positive serial correlation in the log returns alludes towards mean aversion, persistence of shocks and undershooting.

2.5 Structural breaks

Since the span of the data is long the data may suffer from structural breaks. The Perron (1997) test is applied on the nine logged series. The null hypothesis of this test is a unit root with a structural break in both the intercept and the trend. The results for all nine currencies are the same. In all cases the null hypotheses fail to be rejected. Details are in Table 4. However the breakpoints from the tests are rather disparate. In two cases the same break arises, and this is twice in January 1981 and twice in September 1980. There is another case close to the September 1980 break point and this is July 1980. These breakpoints coincide approximately with the Reagan US presidency, and the change in operating procedures by the US Fed. The remaining three breakpoints are August 1996, January 2000, and December 2002.

Having a structural break in the logged series does not imply that the same break occurs in the log returns. In order to resolve this issue autoregressions are first estimated, without a GARCH model, and the Chow break point test is applied on the autoregression residuals for the same break as the one obtained by the Perron test on the logged series, and this is repeated for each currency. For details on the Chow test see Verbeek (2012). The null hypotheses that there are no breaks in the log returns fail to be rejected for eight currencies out of nine (Table 5). Only the Canadian dollar fails the test. A categorical variable is then generated that takes the value 1 before the CAD break point (December 2002), and zero otherwise. This variable is included in the autoregression of the CAD as a shift variable in the intercept, and as a shift variable in the slope, i.e. it appears interactively with the constant and with the lagged log returns. In addition, instead of running a simple regression, an additional GARCH model is specified. See Table 6. The coefficient on the slope interactive variable has an insignificant t-statistic (-0.9422). The joint null hypothesis that there are no shifts in the intercept and in the slope fails to be rejected with an actual p-value of 0.0594. Therefore the breaks in the logged series do not permeate to the series of log returns.

Table 4. Unit root Perron test on $\log(Z)$ (Perron, 1997). The null hypothesis is a unit root with a structural break in both the intercept and the trend.

Variable Z	Perron t-statistic
AUD	-3.916550 (2000M01)
CAD	-3.779244 (2002M12)
CHF	-3.936062 (1996M08)
DKK	-3.845181 (1980M07)
GBP	-4.389166 (1981M01)
NOK	-3.784578 (1980M09)
NZD	-3.322763 (1981M01)
SEK	-3.548546 (1980M09)
JPY	-4.762038 (1985M09)

Notes: See Table 1 for the definition of the variables. The critical values for the Perron t-statistic are -6.32 (1%), -5.59 (5%), and -5.29 (10%). All nulls fail to be rejected at the 10% marginal significance level. Break points are in parenthesis.

Table 5. Stability tests on the autoregressions of $\Delta(\log(Z))$. All regressions include the first-order lag, except the Swedish autoregression which includes both the first-order and the second-order lags.

Variable Z	Quandt-Andrews unknown breakpoint test						Chow breakpoint test		Reset test
	Max LR F-statistic	Max Wald F-statistic	Exp LR F-statistic	Exp Wald F-statistic	Ave LR F-statistic	Ave Wald F-statistic	breakpoint	Wald statistic	
AUD	0.3908	0.3908	0.5139	0.4299	0.4506	0.4506	2000M01	0.2687	0.1465
CAD	0.1131	0.1131	0.1992	0.1229	0.1634	0.1634	2002M12	0.0115	0.1867
CHF	0.4269	0.4269	0.4371	0.4125	0.3188	0.3188	1996M08	0.3351	0.6228
DKK	0.9601	0.9601	1.0000	1.0000	1.0000	1.0000	1980M07	0.5989	0.0808
GBP	0.2978	0.2978	0.5268	0.4698	0.4485	0.4485	1981M01	0.5568	0.0205
NOK	0.9283	0.9238	0.9647	0.9432	0.9359	0.9359	1980M09	0.2769	0.0044
NZD	0.2070	0.2070	0.4468	0.3456	0.3942	0.3942	1981M01	0.7623	0.3704
SEK	0.9049	0.9049	0.8864	0.8635	0.8138	0.8138	1980M09	0.6660	0.1162
JPY	0.7976	0.7976	0.8239	0.7986	0.7648	0.7648	1985M09	0.9801	0.4441

Notes: See Table 1 for the definition of the variables. Actual p-values are reported for all tests. The p-values for the Quandt-Andrews tests are calculated using Hansen's (1997) method. The Quandt-Andrews tests compare 352 breakpoints each, for the null hypothesis that there are no breakpoints within 15% trimmed data. Max stands for maximum, LR for likelihood ratio, Exp for exponential, and Ave for average. The Chow breakpoints are taken from Table 4. The Reset test follows Ramsey (1969).

To be more certain about the existence of breaks Quandt-Andrews unknown break point tests are implemented (Andrews, 1993; Andrews and Ploberger, 1994). Six different test statistics are computed (Table 5). In total, and for each currency autoregression, 352 break points are compared. The actual p-values are reported in Table 5. Since the null hypothesis is the absence of breaks, and since all p-values are larger than 10%, the conclusion is strong that there are no breaks, which is a surprising fact. It seems that a parcimonious ARIMA model is sufficient to describe the short and long run statistical behavior of all nine foreign exchange rates.

Table 6. Regression on $\Delta(\log(Z))$, where the variable Z is the Canadian dollar (CAD).

Conditional mean equation:	
Constant	-0.002664 (-1.732914)
D	0.003329 (2.084640)
$\Delta(\log(Z_{t-1}))$	0.287010 (3.126641)
$D*\Delta(\log(Z_{t-1}))$	-0.100535 (-0.942243)
Conditional variance equation:	
Constant	0.00000251 (1.974426)
ARCH(-1)	0.094315 (2.357699)
GARCH(-1)	0.898055 (32.26257)
Adjusted R-Square	0.072515
F-test that the coefficients on D and on $D*\Delta(\log(Z_{t-1}))$ are jointly zero	0.0594

Notes: In parenthesis are t-statistics. The variable D is a categorical variable that takes the value 1 between January 1971 and December 2002, and zero thereafter. The actual p-value of the joint test is reported. Robust standard errors are calculated following the Bollerslev-Wooldridge adjustment (Bollerslev and Wooldridge, 1992).

From the previous subsections it is known that there are ARCH effects, and that an equation for the conditional variance needs to be specified in the autoregressions. This means that the estimated autoregressions, without a GARCH model for the conditional variance, are badly specified. This should lead to the rejection of the Ramsey Reset stability test (Ramsey, 1969). The null hypothesis of this test is stability or proper specification. Surprisingly the Reset test fails for only two currencies, the British pound and the Norwegian krone (Table 5). This is unexpected. The test should fail for all currencies except for the Swiss franc and the Danish krone, which do not show ARCH effects (see

Table 2). Incidentally, Baillie and Bollerslev (1989) did not find ARCH effects in their monthly sample, which implies that their data set is structurally different.

2.6 Final models and final econometric diagnostics

The final models are autoregressions that include an equation for the conditional variance. The conditional mean equation is compactly specified as follows:

$$\varphi(L)\Delta(\log(Z_t)) = \varepsilon_t \quad \Rightarrow \quad \Delta(\log(Z_t)) = \frac{1}{\varphi(L)}\varepsilon_t = \phi(L)\varepsilon_t \qquad (3)$$

where $\varphi(L)$ and $\phi(L)$ are both polynomials in the lag operator L. The results are in Table 7. One lag suffices to eight autoregressions, i.e. $\varphi(L) = (1 - \lambda L)$, where λ is the autoregressive coefficient. The Swedish kronor autoregression requires two lags, i.e. $\varphi(L) = (1 - \lambda_1 L - \lambda_2 L^2)$, where λ_1 and λ_2 are the first-order and the second-order autoregressive coefficients respectively. The autoregressive coefficients are estimated to be between 0.279 and 0.333, which is a tight range. This implies that the statistical behavior of the nine currencies is quite similar. Moreover these coefficients are all statistically significant. The minimum t-statistic is 4.657. In five cases out of nine the conditional variance is integrated of order 1, i.e. the model is IGARCH (integrated GARCH). Most coefficients in the variance equations are statistically significant. If the autoregression residual at time t is denoted by ε_t, and if the conditional standard deviation at time t is denoted by σ_t, then the standardized residual at time t is ε_t / σ_t.

Table 7. Regressions on $\Delta(\log(Z))$.

Variable Z	Conditional mean equation: Coefficient on			Conditional variance equation: Coefficient on		
	intercept	$\Delta(\log(Z_{t-1}))$	$\Delta(\log(Z_{t-2}))$	intercept	ARCH(-1)	GARCH(-1)
AUD	0.000526	0.325135	-	0.000153	0.593219	0.300510
	(0.517029)	(7.035812)		(2.960904)	(3.246016)	(2.427378)
CAD	0.000367	0.223209	-	2.41E-06	0.097891	0.895666
	(0.841046)	(4.656988)		(1.830820)	(2.783882)	(37.88361)
CHF	-0.001537	0.279151	-	-	0.023550	0.976450
	(1.161787)	(5.277624)			(1.313407)	(54.45852)
DKK	-0.006056	0.308616	-	-	0.079698	0.920302
	(1.956384)	(6.135774)			(2.300742)	(26.56757)
GBP	-0.000340	0.332968	-	2.93E-05	0.100016	0.844588
	(0.394660)	(6.935623)		(2.533685)	(2.906350)	(20.87928)
NOK	0.004102	0.312710	-	-	0.051204	0.948796
	(1.551621)	(5.184649)			(2.188617)	(40.55485)
NZD	-0.000232	0.354849	-	1.95E-05	0.047483	0.924076
	(0.278176)	(7.836354)		(1.929150)	(2.632452)	(36.97580)
SEK	-0.004900	0.431193	-0.145232	-	0.079492	0.920508
	(2.073660)	(8.228501)	(2.301445)		(2.963200)	(34.31348)
JPY	-0.003659	0.292939	-	-	0.034177	0.965823
	(2.304325)	(5.111069)			(2.194659)	(62.01987)

Notes: See Table 1 for the definition of the variables. Absolute t-statistics are in parenthesis. Robust standard errors are calculated following the Bollerslev-Wooldridge adjustment (Bollerslev and Wooldridge, 1992).

Table 8 provides tests on the standardized residuals for all nine currencies. If the final models are well-specified it is expected that the standardized residuals are well-behaved. The Ljung-Box Q-statistics, for lag lengths of 12 and 24 months, on the standardized residuals, and their squares, provide support to the hypotheses that these adjusted residuals are independently and identically distributed for all nine foreign currencies. The BDS tests, (Brock et al., 1996), for dimensions 2, 4, 8, and 12, find also linear and non-linear independence in general. Non-reported runs tests fail to reject randomness and therefore the hypothesis that the nine series of adjusted residuals are independently and identically distributed. For this test the lowest asymptotic two-tailed p-value is 0.211. Five different normality tests are carried out. See the Eviews7 (2009) user's guides for details. The most stringent normality test is the Jarque-Bera test, (Jarque and Bera, 1980, 1987), which rejects normality for all nine standardized residuals. However there is evidence of normality, from the other four normality tests, for the Canadian dollar and the British pound. The central limit theorem can be invoked to allow valid hypothesis testing.

Table 8. Tests on the standardized residuals from the regressions in Table 5.

Variable Z	Q-statistic with k: 12	Q-statistic with k: 24	Q^2-statistic with k: 12	Q^2-statistic with k: 24	Jarque-Bera test	BDS test with dimension: 2	4	8	12	Lilliefors (D)	Cramer-von Mises	Watson (U2)	Anderson-Darling (A2)
AUD	0.653	0.752	0.996	0.983	0.000	0.035	0.447	0.569	0.234	0.000	0.000	0.000	0.000
CAD	0.196	0.185	0.615	0.688	0.000	0.392	0.309	0.385	0.269	> 0.1	0.599	0.680	0.349
CHF	0.804	0.768	0.142	0.280	0.000	0.906	0.090	0.038	0.040	0.006	0.001	0.000	0.000
DKK	0.399	0.917	0.764	0.984	0.000	0.169	0.688	0.750	0.975	0.031	0.001	0.000	0.000
GBP	0.322	0.121	0.758	0.449	0.008	0.306	0.133	0.416	0.473	> 0.1	0.148	0.166	0.070
NOK	0.727	0.734	0.961	0.998	0.000	0.374	0.008	0.036	0.124	0.005	0.003	0.001	0.001
NZD	0.210	0.232	0.986	0.841	0.000	0.978	0.291	0.152	0.227	0.000	0.000	0.000	0.000
SEK	0.392	0.436	0.954	0.912	0.000	0.623	0.246	0.379	0.572	0.000	0.000	0.000	0.000
JPY	0.234	0.249	0.802	0.001	0.000	0.027	0.084	0.071	0.032	0.000	0.000	0.000	0.000

Notes: See Table 1 for the definition of the variables. Actual p-values are reported for all tests. The Q-statistics and the Q^2-statistics are the Ljung-Box Q-statistics on the standardized residuals and on the squares of the standardized residuals respectively. All other tests are on the standardized residuals. The p-values of the BDS tests, Brock et al. (1996), are from the bootstrap distribution with 5,000 repetitions.

From the autoregressive coefficients one can find out the long run behavior of each currency. Without loss of generality, if λ is equal to the coefficient on the autoregressive variable, as in equation (3), then $1/(1-\lambda)$ measures the long run behavior of the currency. Since λ is estimated to be between zero and one, $1/(1-\lambda)$ is definitely higher than one, providing evidence for *mean aversion*. In Table 9 this aversion coefficient is estimated for each currency, and a t-test is undertaken on whether this coefficient is significantly higher than +1. In all nine cases this coefficient, which is a measure of the long run impact of shocks, is significantly higher than +1. In the same table the variance ratios, that denote persistence of shocks, are calculated by applying the following formulae (Campbell and Mankiw, 1987a, 1987b):

$$Variance\ \ ratio = \left(1 - R^2\right)\left(\frac{1}{\varphi(1)}\right)^2 \tag{4}$$

In equation (4) R^2 stands for the autoregression R-Square, and $\varphi(L)$ is the polynomial of the lag operator L that is estimated from the autoregressions (see equation (3)). All nine values of the variance ratio are much higher than 1 (Table 9, column 4) and are generally close to the variance ratios, for period 48, computed by equation (2) above, and reported in Table 1 and repeated in column 5 of Table 9. Mean aversion and persistence of shocks are thus features of all nine currencies. While there are many explanations for such a finding, the latter is nevertheless troublesome, especially since it does not seem to depend on the foreign currency as it is common to all nine of them.

Table 9. Mean aversion based on $\varphi(L)\Delta(\log(Z_t))$ where φ is a polynomial in L , i.e. $\varphi(L) = 1 - \lambda_1 L - \lambda_2 L^2$ in the case of the Swedish kronor, and $\varphi(L) = 1 - \lambda L$ for the other eight currencies, L being the lag operator.

$$\varphi(L)\Delta(\log(Z_t)) = \varepsilon_t \;\Rightarrow\; \Delta(\log(Z_t)) = \frac{\varepsilon_t}{\varphi(1)} \textbf{ in the long run}$$

Variable Z	$1/\varphi(1)$	Autoregression R-Square (R^2)	Estimate of the variance ratio = $(1-R^2)*(1/\varphi(1))^2$	Estimate from Table 1 for a period of 48 months
AUD	1.481778 (0.101465) [4.748222]	0.110444	1.95316	1.7203
CAD	1.287348 (0.079433) [3.617505]	0.061042	1.55610	2.2677
CHF	1.387252 (0.101791) [3.804372]	0.079002	1.77243	1.5385
DKK	1.446375 (0.105223) [4.242175]	0.049250	1.98900	2.2335
GBP	1.499178 (0.107900) [4.626284]	0.124560	1.96758	1.8308
NOK	1.454990 (0.127686) [3.563357]	0.093023	1.92007	1.7015
NZD	1.550026 (0.108795) [5.055629]	0.123738	2.10529	2.3588
SEK	1.400483 (0.125958) [3.179487]	0.121952	1.72216	2.1742
JPY	1.414305 (0.114644) [3.613837]	0.096439	1.80736	1.5101

Notes: See Table 1 for the definition of the variables. The estimators of $1/\varphi(1)$ are obtained by the delta method using analytic derivatives. In parenthesis are standard errors. In brackets are t-statistics for the null that $1/\varphi(1)$ is equal to 1.

Mean aversion and persistence of shocks can also be measured by the accumulated impulse response effects of the autoregressions (Campbell and Mankiw, 1987a). The effects up to 12 months in the future, based on the nine separate ARMA models for $\Delta(\log(Z_t))$, are reproduced in Table 10, column 2. The evidence for mean aversion is even stronger than that in Table 9 because all the t-statistics for the nulls that the effects are +1 are higher than those in Table 9. Nevertheless the accumulated impulse response effects are very close in magnitude to the mean aversion coefficients in column 2 of Table 9 with a slight upward bias. It seems that a horizon of 12 months is enough to represent the long run. Taking a horizon of 24 months instead of 12 does not materially change the accumulated impulse response effects and their standard errors. The difference between the estimates in column 2 of Table 9 and those in column 2 of Table 10 is that, in the former, a GARCH model is appended to the estimation while, in the latter, the model is just a simple autoregression. Econometrically, the estimates from the former model ought to be more precise because they take directly into consideration more information about the statistical processes that the data follow.

If one adopts a simultaneous vector autoregression (VAR) of the first-order for the nine currencies together, the results are substantially different (Table 10, column 3). Three currencies have now aversion coefficients $(1/\varphi(1))$ that are insignificantly different from +1, and these are the CAD, the DKK, and the NZD. Moreover, when testing for an aversion coefficient that is equal to +1 all the nine t-statistics become smaller which implies a loss in precision. However, little confidence can be put in these results because such a VAR suffers from the following four specification biases: the non-reported sensitivity of the estimates to the lag order, the non-reported fact that most cross effects are insignificant, the lack of a supplementary conditional variance equation, and the fact that constraining the VAR to be first-order for all currencies is unacceptable because it is already established that the SEK autoregression is of a second-order. From this it is concluded that, despite the above, the evidence for mean aversion and persistence of shocks remains very strong.

Table 10. Mean aversion based on the accumulated effect of a nonfactorized one unit innovation in $\Delta(\log(Z_t))$ on $\Delta(\log(Z_t))$, assuming a first-order autoregression for all currencies, except for the Swedish krona which follows a second-order autoregression (column 2), and assuming a first-order vector autoregression (column 3).

Variable Z	Accumulated effect up to 12 months from individual autoregressions	Accumulated effect up to 12 months from first-order vector autoregression
AUD	1.499919 (0.09476) [5.275633]	1.586255 (0.12830) [4.569408]
CAD	1.333530 (0.07692) [4.336063]	1.162037 (0.08649) [1.873477]
CHF	1.392800 (0.08315) [4.723993]	1.617589 (0.17101) [3.611420]
DKK	1.455709 (0.08993) [5.067375]	0.921306 (0.23729) [-0.331636]
GBP	1.547176 (0.10003) [5.470119]	1.419937 (0.13596) [3.088680]
NOK	1.562268 (0.10174) [5.526519]	1.642617 (0.20747) [3.097397]
NZD	1.542815 (0.09954) [5.453235]	1.130291 (0.12994) [1.002701]
SEK	1.457830 (0.10441) [4.384925]	1.579645 (0.19037) [3.044834]
JPY	1.470294 (0.09200) [5.111891]	1.390791 (0.10898) [3.585896]

Notes: See Table 1 for the definition of the variables. In parenthesis are standard errors. In brackets are t-statistics for the nulls that the accumulated effects are equal to 1.

3. Conclusion

This paper analyses the statistical behavior of the US dollar, against nine different currencies, over the float period, using a long monthly data set. The martingale, or random walk, hypothesis is rejected for all nine logged currencies. Although all these logged currencies have a unit root the changes in the logs, or log returns, are not white noise and suffer from significant positive serial correlation and ARCH effects. The nine logged series are characterized by structural breaks in both the intercept and the slope. But, surprisingly, the retreived breaks from the logged series do not show up in the series of log returns. In addition, and although the sample period is long and has witnessed different economic, social and political regimes, there is no evidence for any breaks in these log returns. The nine logged currencies are well described by a low-order ARIMA model, with a parcimonious GARCH specification of the conditional variance. These ARIMA models imply mean aversion and high persistence of shocks. Mean aversion and persistence of shocks are formally tested and found to be significant for all nine currencies, whether one adjusts for GARCH effects or not. This provides justification for the title of the paper.

References

Abdalla, S.Z.S. (2012), *Modeling exchange rate volatility using GARCH models: empirical evidence from Arab countries*, International Journal of Economics and Finance, 4, 216-229.

Andrews, D.W.K. (1993), *Tests for parameter instability and structural change with unknown change point,* Econometrica, 61, 821-856.

Andrews, D.W.K., Ploberger, W. (1994), *Optimal tests when a nuisance parameter is present only under the alternative,* Econometrica, 62, 1383-1414.

Azar, S. (2008), *Jensen's inequality in finance*, International Advances in Economic Research, 14, 433-440.

Baillie, R.T., Bollerslev, T. (1989), *The message in daily exchange rates: A conditional-variance tale*, Journal of Business & Economic Statistics, 7, 297-305.

Bollerslev, T. (1986), *Generalized autoregressive conditional heteroskedasticity*, Journal of Econometrics, 31, 307-327.

Bollerslev, T., Wooldridge, J.M. (1992), *Quasi-maximum likelihood estimation and inference in dynamic models with time varying covariances*, Econometric Reviews, 11, 143-172.

Brock, W., Dechert, D., Sheinkman, J., LeBaron, B. (1996), *A test for independence based on the correlation dimension*, Econometric Reviews, 15, 197-235.

Campbell, J.Y., Mankiw, N.G. (1987a), *Are output fluctuations transitory?* Quarterly Journal of Economics, 102, 857-880.

Campbell, J.Y., Mankiw, N.G. (1987b), *Permanent and transitory components in macroeconomic fluctuations*, American Economic Review, 77, 111-117.

Chang, Y. (2004), *A re-examination of variance-ratio test of random walks in foreign exchange rates*, Applied Financial Economics, 14, 671-679.

Dickey, D.A., Fuller, W.A. (1979), *Distribution of the estimators for autoregressive time series with a unit root*, Journal of the American Statistical Association, 74, 427-431.

Dornbusch, R. (1976), *Expectations and exchange rate dynamics*, Journal of Political Economy, 84, 1161-1176.

Engle, R. (1982), *Autoregressive conditional heteroskedasticity with estimates of the variance of U.K. inflation,* Econometrica, 50, 987-1008.

EViews 7 User's Guides (2009), *Quantitative Micro Software*. Irvine, CA, USA.

Fama, E. (1965), *The behavior of stock market prices*, Journal of Business, 38, 34-105.

Fama, E. (1970), *Efficient capital markets: A review of theory and empirical work*, Journal of Finance, 25, 383-417.

Fama, E. (1991), *Efficient markets II*, Journal of Finance, 46, 1575-1618.

Hansen, B.E. (1997), *Approximate asymptotic p-values for structural-change tests,* Journal of Business & Economic Statistics, 15, 60-67.

Hsieh, D.A. (1989), *Modeling heteroscedasticity in daily foreign-exchange rates*, Journal of Business & Economic Statistics, 7, 307-317.

Ibrahim, J., Long, Y., Ghani, H.A., Salleh, S.I.M. (2011), *Weak-form efficiency of foreign exchange market in the Organization for Economic Cooperation and Development countries: unit root test*, International Journal of Business and Management, 6, 55-65.

Jarque, C.M., Bera, A.K. (1980), *Efficient tests for normality, homoskedasticity and serial dependence of regression residuals*, Economics Letters, 6, 255-259.

Jarque, C.M., Bera, A.K. (1987), *A test for normality of observations and regression residuals,* International Statistical Review, 55, 163-172.

Kim, J.H. (2006), *Wild bootstrapping variance ratio tests*, Economics Letters, 92, 38-43.

Liu, C.Y., He, J. (1991), *A variance-ratio test of random walks in foreign exchange rates*, Journal of Finance, 46, 773-785.

Ljung, G., Box, G. (1978), *On a measure of lack of fit in time series models*, Biometrika, 66, 265-270.

Lo, A.W., MacKinlay, A.C. (1988), *Stock market prices do not follow random walks: evidence from a simple specification test*, The Review of Financial Studies, 1, 41-66.

Perron, P. (1997), *Further evidence on breaking trend functions in macroeconomic variables*, Journal of Econometrics, 80, 355-385

Phillips, P.C.B., Perron, P. (1988), *Testing for a unit root in time series regression*, Biometrika, 75, 335-346.

Ramsey, J.B. (1969), *Tests for specification errors in classical linear least squares regression analysis,* Journal of the Royal Statistical Society*,* Series B*,* 31, 350–371.

Schwarz, G.E. (1978), *Estimating the dimension of a model*, Annals of Statistics, 6, 461–464.

Van De Gucht, L., Dekimpe, M.G., Kwok, C.C.Y. (1996), *Persistence in foreign exchange rates*, Journal of International Money and Finance, 15, 191-220.

Verbeek, M. (2012), *A Guide to Modern Econometrics*. Chichester: Wiley, 4th edition.