

Choosing a Data Frequency to Forecast the Quarterly Yen-Dollar
Exchange Rate

by

Benjamin Cann
B.Sc, University of Victoria, 2013

A Thesis Submitted in Partial Fulfillment of the
Requirements for the Degree of

MASTER OF ARTS

in the Department of Economics

©Benjamin Cann, 2016
University of Victoria

All rights reserved. This thesis may not be reproduced in whole or in part,
by photocopy or other means, without the permission of the author.

Choosing a Data Frequency to Forecast the Quarterly Yen-Dollar Exchange Rate

by

Benjamin Cann
B.Sc, University of Victoria, 2013

Supervisory Committee

Dr. David Giles, Supervisor
Department of Economics

Dr. Judith Clarke, Departmental Member
Department of Economics

ABSTRACT

Potentially valuable information about the underlying data generating process of a dependent variable is often lost when an independent variable is transformed to fit into the same sampling frequency as a dependent variable. With the mixed data sampling (MIDAS) technique and increasingly available data at high frequencies, the issue of choosing an optimal sampling frequency becomes apparent. We use financial data and the MIDAS technique to estimate thousands of regressions and forecasts in the quarterly, monthly, weekly, and daily sampling frequencies. Model fit and forecast performance measurements are calculated from each estimation and used to generate summary statistics for each sampling frequency so that comparisons can be made between frequencies. Our regression models contain an autoregressive component and five additional independent variables and are estimated with varying lag length specifications that incrementally increase up to five years of lags. Each regression is used to forecast a rolling, one and two-step ahead, static forecast of the quarterly Yen and U.S Dollar spot exchange rate. Our results suggest that it may be favourable to include high frequency variables for closer modeling of the underlying data generating process but not necessarily for increased forecasting performance.

Keywords: mixed data sampling, forecasting, model selection criteria, time-series, yen dollar exchange rate

Table of Contents

Supervisory Committee	ii
Abstract	iii
Table of Contents	iv
1. Introduction	1
2. Midas Modelling	1
3. Data	4
4. Methodology	9
5. Results	13
6. Discussion	18
7. Further Research	20
8. Conclusions	21
9. References	23
Appendices	26
Appendix A: Model Selection Criterion and Forecasting Error	27
Appendix B: Data Plots.....	28
Appendix C: Results by Frequency.....	35

List of Tables

Table 1: Regression Data Sources	5
Table 2: ADF Unit Root test's DF and P-Value Across Frequencies.....	13
Table 3: PP Unit Root test's DF and P-Value Across Frequencies	13
Table 4: KPSS Unit Root test's KPSS Level and P-Value Across Frequencies	13
Table 5: High- Frequency Independent Variable Observations per 1 Dependent Variable Observation ..	14
Table 6: Comparing Results Across Frequencies: Horizon 1	15
Table 7: Comparing Results Across Frequencies: Horizon 2	16
Table 8: Summary Statistic Results by Frequency: Horizon 1	33
Table 9: Summary Statistic Results by Frequency: Horizon 2	34

1. INTRODUCTION

Potentially valuable information about the underlying data generating process of a dependent variable is often lost when an independent variable is transformed to fit into the same sampling frequency as a dependent variable. Mixed data sampling (MIDAS) is a technique developed by Ghysels, Sinko, and Valkanov (2004) that allows independent variables to appear at a higher frequency than the dependent variable parsimoniously. Since data is increasingly being recorded and becoming available at higher frequencies, the next issue to explore is which sampling frequency is optimal for modeling, forecasting, and whether sampling frequencies behave differently as the distance between the dependent variable sampling frequency and independent variable sampling frequencies increase. Our experiment evaluates the general performance of sampling frequencies as a whole in estimating and forecasting with mixed data sampling.

We build a representation of the performance of an entire sampling frequency for a limited data range. We do this by generating summary statistics for each frequency from model selection criterion and forecasting performance measurements calculated from thousands of regression and forecast estimations differing in varying lag lengths. The dependent variable is kept in the quarterly sampling frequency while we change the sampling frequency of the independent variables between four sampling frequencies: quarterly, monthly, weekly, and daily. Each of our regressions contain an autoregressive component and five additional independent variables. The independent variables are modelled in each of the four sampling frequencies with a comprehensive set of different lag length model specifications combinations starting low and incrementally increasing to five years of lags. Our variables are chosen from financial data sets due to their consistency and availability at higher frequencies. We forecast the Yen- U.S dollar quarterly spot exchange rate using rolling, one and two-step ahead forecasts.

The thesis proceeds as follows. Section 2 outlines mixed frequency data sampling regressions. Section 3 describes and graphs our data. The methodology framework we've created to compare the different sampling frequencies and their relative performance is introduced in Section 4. Section 5 presents the results and is followed by a discussion in section 6. Future research considerations and potential expansions are outlined in Section 7 followed by a set of summarized conclusions in Section 8.

2. MIDAS MODELING

The mixed frequency data sampling method allows independent variables occurring at a higher frequency than the dependent variable to be included in estimation at their original higher frequency

without needing to transform them to the same frequency as the dependent variable. The motivation behind MIDAS is to include a variable's higher frequency observations to take advantage of the increased information regarding its underlying data-generating process to potentially increase estimation and forecasting performance (e.g., Galvao, 2013, p. 397). Simply adding lagged independent variables per a distributed lag model quickly exhausts a model's degrees of freedom, especially if the distance of sampling frequencies is great. With parsimonious models in mind, Ghysels, Santa-Clara, and Valkanov (2004) introduced a framework that enables the forecaster to estimate a relatively small number of parameters representing a larger set of parameters, defined as hyperparameters, by using exogenously determined distributed lag polynomials as weighting functions (see, e.g., Armesto, Engemann and Owyang, 2010, p. 521).

MIDAS has been used to estimate the state of the economy using factor models with unbalanced datasets (see, e.g., Marcellino and Schumacher, 2010), integrate with Bayesian techniques and forecast GDP using data that is available after the last in-sample dependent variable observation (nowcasting) (see, e.g., Duan, 2015), forecast Chinese inflation with Google user generated data (see, e.g., Li, Shang, Wang, and Ma, 2015), utilize financial data's high frequency observations to forecast economic indicators (see, e.g., Andreou, Ghysels, and Kourtellis, 2013), forecast exchange rate volatility (see, e.g., Chortareas, Jiang, and Nankervis, 2011), and incorporate the lagged dependent regressor into MIDAS to create the MIDAS-AR model (see, e.g., Clements and Galvão, 2008).

Computationally, MIDAS regression statistical packages are currently available in MATLAB, the open source R programming statistical software, and in EViews starting from version 9.5 (EViews, 2016). The original MATLAB package written by Eric Ghysels and collaborators was updated in August 2016 by Hang Qian. The R MIDAS package, named *midasr*, was written by Ghysels, Kvedaras, and Zemlys in 2016. This thesis uses the *midasr* package written in R. Ghysels, Kvedaras, and Zemlys (2016) also wrote an extensive user guide demonstrating much of *midasr*'s functionality in MIDAS model restriction specifications, model selection criterion, forecasting without new data, and nowcasting incorporating available high frequency data before the forecasted lower frequency variable.

The fundamental MIDAS model for a single explanatory variable and h step ahead forecasting horizon can be described by

$$y_t = \beta_0 + \beta_1 B(L^{1/m}; \theta) x_{t-h}^{(m)} + \varepsilon_t, \quad (1)$$

where $B(L^{1/m}; \theta)$ is a lag distribution and equal to $\sum_{k=1}^K b(k; \theta)L^{(k-1)/m}$ where $L^{s/m}x_{t-1}^m = x_{t-1-s/m}^m$ (Clements and Galvao, 2008, p. 547). Note that s is added for simplicity and is equal to $k - 1$, x is the independent variable, m is the frequency, h is the forecasting horizon, lower case k is the high frequency lag order, and upper case K is the maximum amount of high frequency lags included in the estimation. For example, if our dependent variable is quarterly and our explanatory variables are monthly, $m = 3$ as there are three monthly observations for every quarter. $b(k; \theta)$ can be parametrized by functional restriction such as

$$b(k; \theta) = \frac{\exp(\theta_1 k + \theta_2 k^2)}{\sum_{k=1}^K \exp(\theta_1 k + \theta_2 k^2)}, \quad (2)$$

the exponential Almon lag function (Clements and Galvao, 2008, p. 547). This is not the only or necessarily the best structure to impose on the high frequency observations. The midasr package provides many different weighting options (see e.g., Ghysels, Kvedaras, and Zemyls, 2016, p.20). Note that, by construction, the θ_1 and θ_2 parameters in equation 2 cannot be negative and are normalized to sum to 1. The θ parameters impose the structural form that the high frequency independent variables take and are the estimated output of the midasr package we use (Ghysels, Kvedaras, and Zemyls, 2016, p. 20). Of course, by normalization, equation 2 is nonlinear and renders the θ parameters to be non-informative regarding the marginal effect each parameter has on the dependent variable. It is important to note that despite the potential for considerable efficiency gains from a proper specification (Ghysels, Kvedaras, and Zemyls, 2016, p. 3), imposing a priori assumptions on the underlying form of the high frequency variable observations risk constraint misspecification leading to asymptotic bias and incorrect distribution assumptions potentially rendering further econometric tests invalid. Fortunately, Ghysels, Kvedaras, and Zemyls (2016) show that when even when degrees of freedom are low, an incorrect parameterization constraint may lead to efficiency gains compared to an unconstrained MIDAS model (p. 4).

A notable strength of MIDAS forecasting is the ability to include new high frequency observations that occur between the last available low frequency observation and the first horizons forecasted value. Suppose that the value of x is available for the first two months of the quarter being forecasted. This can be represented in the MIDAS framework simply by setting $h = 1/3$. This then indicates that 2/3 of the current quarter's monthly information is known. Algebraically, this can be represented by

$$y_t = \beta_0 + \beta_1 B(L^{1/3}; \theta) x_{t-1/3}^{(3)} + \varepsilon_t, \quad (3)$$

where h is now less than one (Clements and Galvao, 2008, p. 457). This paper represents this notion of including intra-quarter monthly observations as nowcasting. Nowcasting is a novel feature of MIDAS techniques due to the nature of higher frequency data being available before lower frequency publication dates; however, nowcasting is not explored in this paper because of the ready availability of our dependent variable in higher frequencies. Nowcasting is most appropriate for models whose dependent variable fits two non-mutually exclusive criteria: it is available only at lower frequencies or has a publication lag. Because our dependent variable is an exchange rate it is readily available at a relatively high frequency. Being so readily available provides very limited opportunity to include new independent data observations between dependent variable observations to improve forecasts. For this thesis, if two months of available higher-frequency variable observations were included after our last in-sample quarterly exchange rate dependent variable observation, it may make more sense to, instead, change the model from a one quarter dependent variable forecast to a thirty-day forecast of a daily exchange rate dependent variable.

Autoregressive (AR) components are often added to forecasting models to improve forecasts. Clements and Galvao (2008) explain how a seasonal response is inadvertently generated of y to $x_{t-h}^{(m)}$ when a lagged dependent variable is simply added to the MIDAS model (p. 547). Their solution is to add AR dynamics of the dependent variable as a common factor so that the response of y to $x_{t-h}^{(m)}$ does not become seasonal (Clements and Galvao, 2008, p. 547). This results in equation (4)

$$y_t = \lambda y_{t-d} + \beta_1 + \beta_2 B(L^{1/m}; \theta) (1 - \lambda L^d) x_{t-h}^{(m)} + \varepsilon_t, \quad (4)$$

which is estimated in the multiple steps described by Clements and Galvao (2008, p. 547). The parameter d is an integer when h is an integer, but when nowcasting information is available on, say, a monthly basis and $h = 1/3$, d remains equal to 1 (Clements and Galvao, 2008, p.547).

3. DATA

Each data set was obtained through the databases summarized in Table 1 via the Quandl finance and economic data API (Quandl, 2016). All datasets are available in daily, weekly, monthly, and quarterly frequencies. The eXtensible Time Series (xts) package in r changes periods of the data by choosing the endpoint observation for all periods except for monthly and quarterly periods where the starting value is chosen (Ryan & Ulrich, 2015, p.44). The xts package in r also provides different methods to change a datasets frequency (see e.g., Ryan & Ulrich, 2015, p. 44). Data is in nominal rates, not de-seasonalized, nor are they de-trended. Nominal rates are used because nominal rates are being forecasted. Nominal

rates are being forecasted to represent what a currency trader may be facing when deciding to trade or not.

Table 1

Regression Data Sources

<u>Data Set</u>	<u>Source</u>
JPY/USD	US Federal Reserve ¹
Nikkei	Nikkei ²
S&P500	Yahoo Finance ³
GBP/USD	US Federal Reserve ⁴
JPY/GBP	Calculated from GBP/USD and JPY/USD
Federal Reserve Overnight Rate	St. Louis Federal Reserve ⁵

Note: Data source citations are as follows,

1: (Board of Governors of the Federal Reserve System, 2016)

2: (Nikkei, 2016)

3: (Yahoo Finance, 2016)

4: (Board of Governors of the Federal Reserve System, 2016)

5: (St. Louis Federal Reserve, 2016)

To explain the long-term U.S dollar (USD) and Japanese yen (JPY) exchange rate, we use the Nikkei stock index, S&P 500 stock index, USD and British pound (GBP) exchange rate, JPY and GBP exchange rate, and the U.S's federal funds overnight interest rate. The Nikkei is used as an approximate proxy for the health of the Japanese economy while the S&P 500 is used for the same purpose for the U.S economy. The U.S dollar – Great British Pound is included because it's considered a major currency pair and the Yen-Pound is included because it's a minor currency pair both providing influence to the Yen and U.S dollar rates separately. It would be best to include the interest rate differential between Japan and the U.S, but because the Bank of Japan's rate has historically been constantly near zero, the federal funds rate itself is considered to be sufficient. The DOW Jones Industrial average is not used to represent the U.S market due to its price-weighted index calculation method that over emphasizes stocks with a higher price per share and consequently does not take into account the volume of stocks available for a company (Nagarajan, 2009).

The following pages display plots of the variables in the daily frequency. Differenced plots are present because, as seen in the unit-root test results presented in Section 5, the data is believed to be integrated. It's also worth noting that the S&P 500 differenced plot has clustering typical of GARCH models; but, because the S&P 500 data is not a dependent variable, no GARCH filtering is applied.

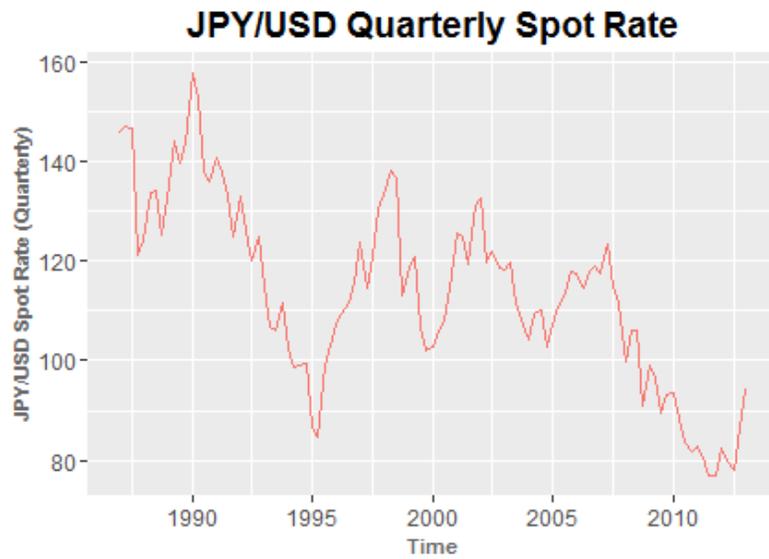


Figure 1: JPY/USD Quarterly Spot Rate

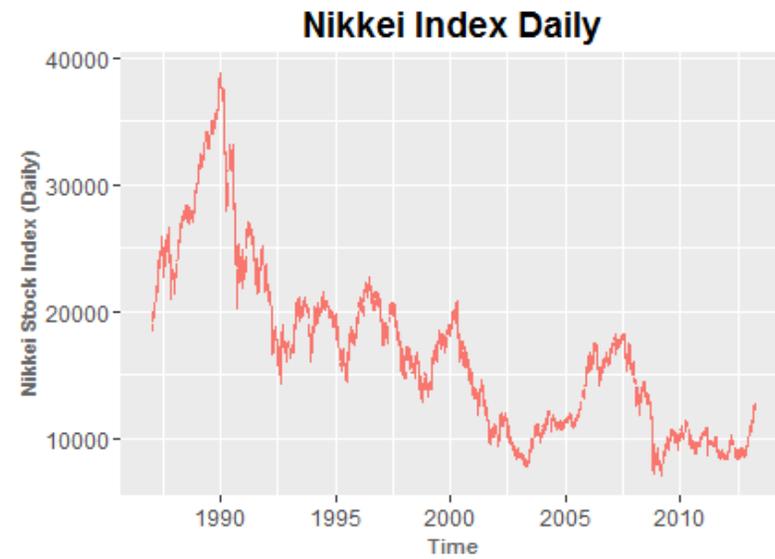


Figure 2: Nikkei Index Daily

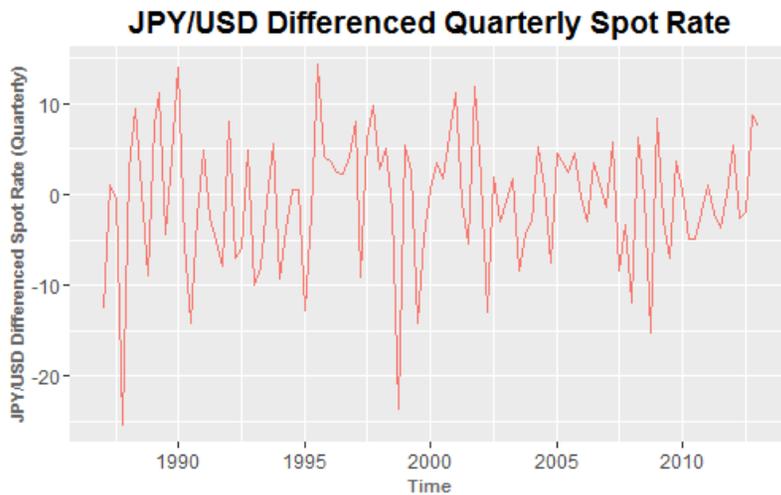


Figure 3: JPY/USD Differenced Quarterly Spot Rate

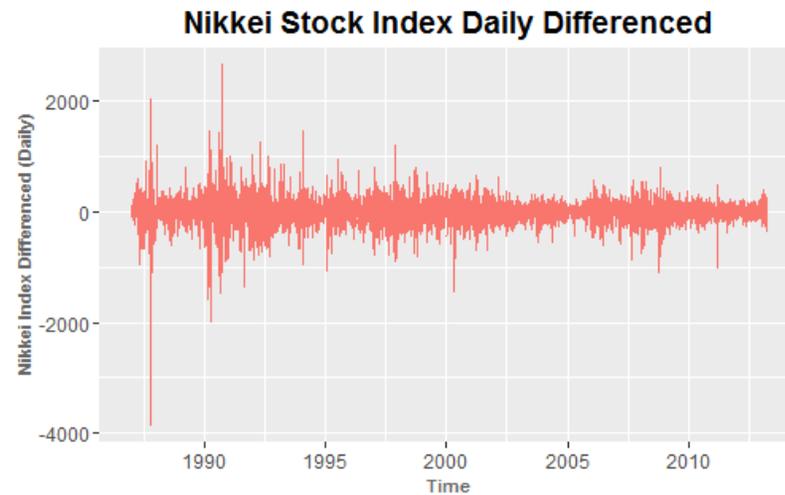


Figure 4: Nikkei Index Daily Differenced

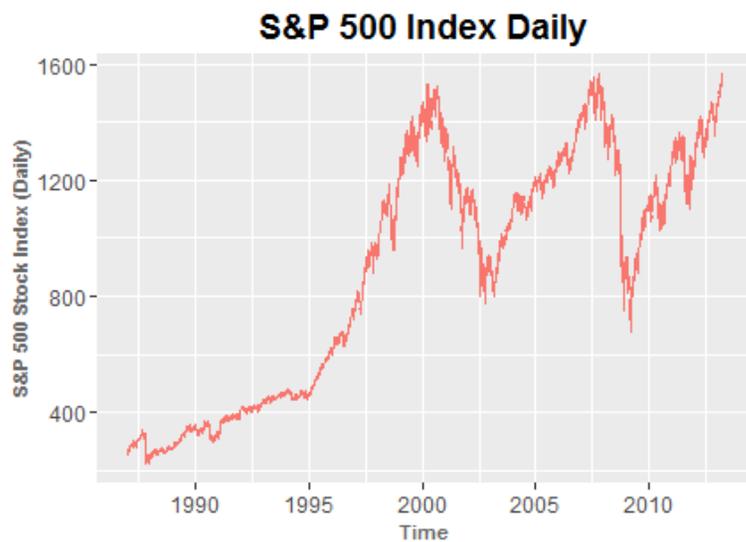


Figure 5: S&P 500 Index Daily

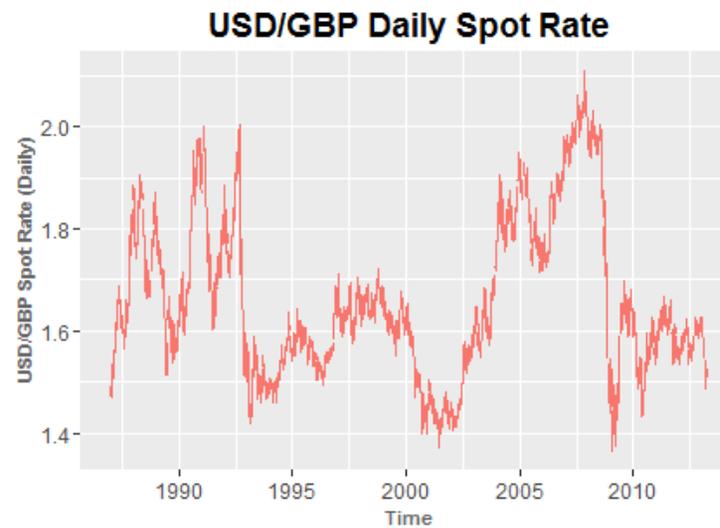


Figure 6: USD/GBP Daily Spot Rate

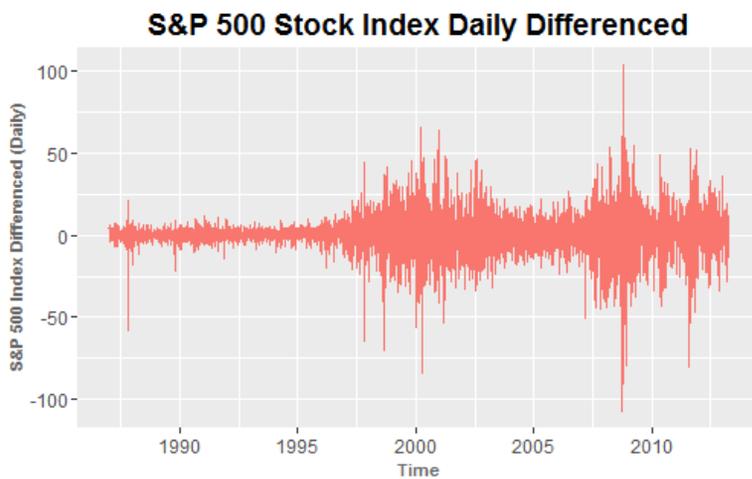


Figure 7: S&P 500 Index Daily Differenced

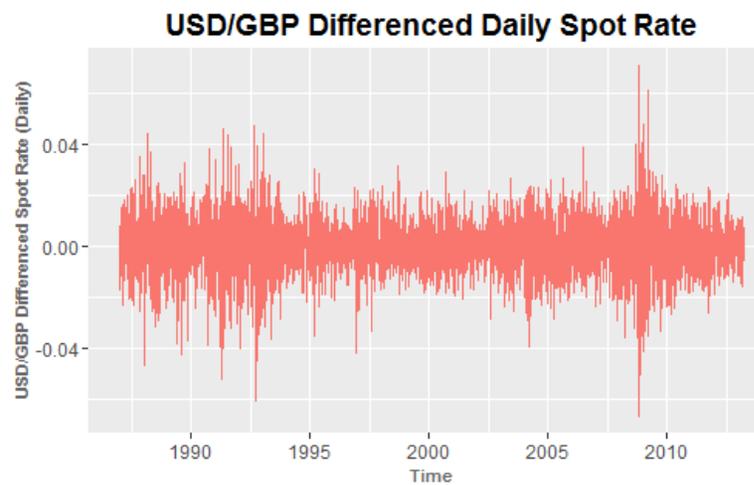


Figure 8: USD/GBP Differenced Daily Spot Rate

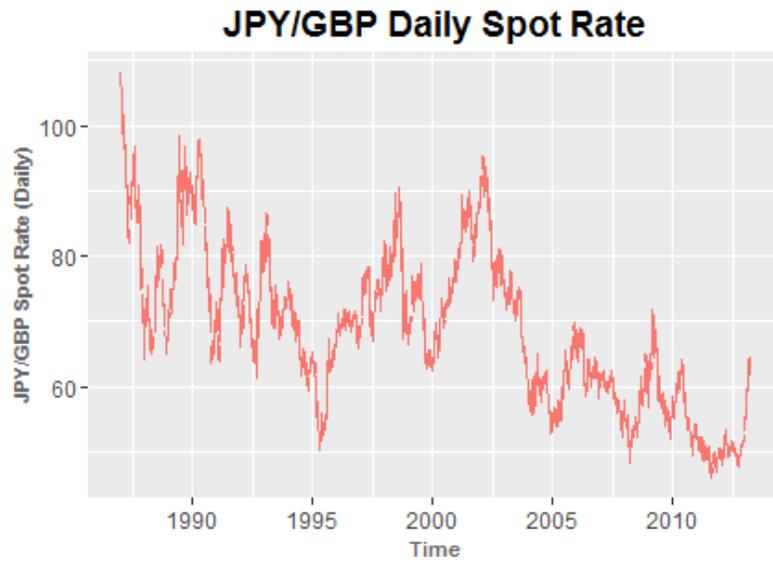


Figure 9: JPY/GBP Daily Spot Rate

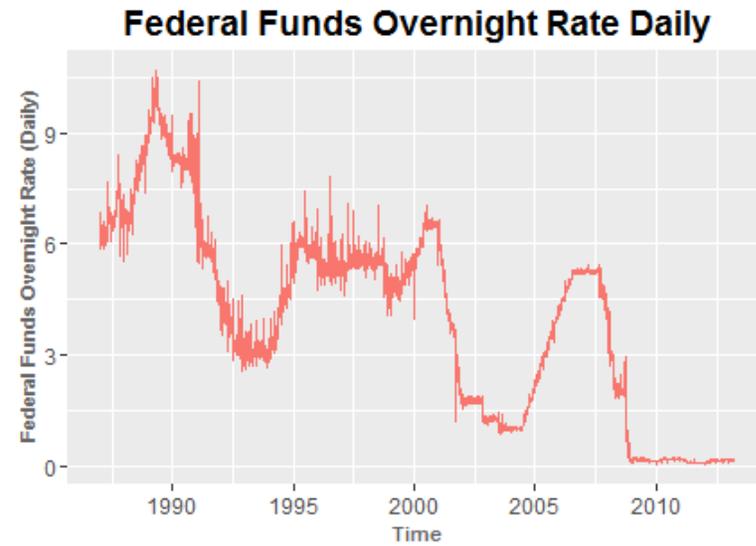


Figure 10: Federal Funds Overnight Rate Daily

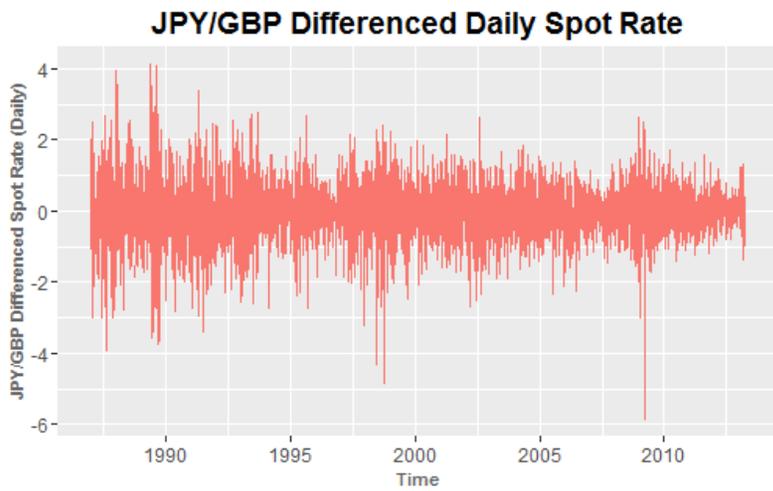


Figure 11: JPY/GBP Differenced Daily Spot Rate

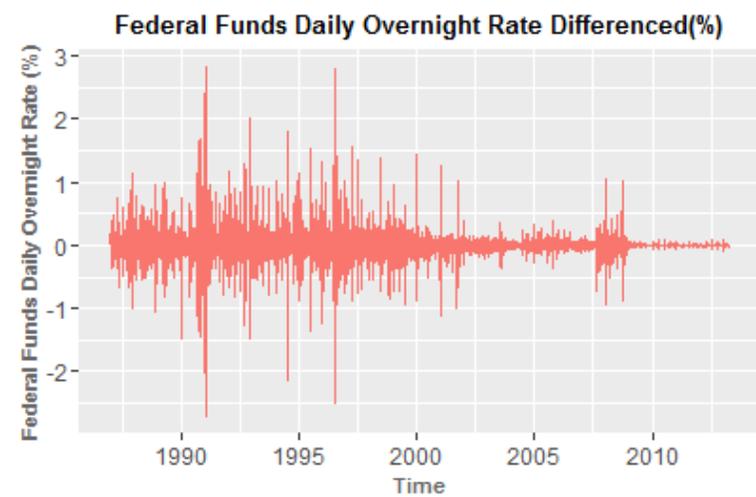


Figure 12: Federal Funds Daily Overnight Rate Differenced

4. METHODOLOGY

The Yen – U.S Dollar was chosen instead of another major currency pair due to the recent history of Japan's economy of the last 20 years. The asset pricing bubble collapse in the late 1980's and early 1990's (Okina, Shirakawa, and Shiratsuka, 2001), Tohoku megathrust 9.0 earthquake and tsunami in 2011 (USGS, 2011), contested Yen depreciation in 2016 (Mogi & Komiya, 2016), and the Bank of Japan's adoption of a negative interest rate monetary policy in 2016 are the reasons why we decided to use the Yen-U.S Dollar currency pair instead of another major pair. The reason the Yen-U.S Dollar exchange rate is chosen to be forecasted in its quarterly sampling frequency versus a typical higher frequency is an attempt to model longer term trends without the extreme variation typical of high-frequency exchange rates. Having knowledge of the longer term trends of the quarterly exchange rate would be potentially valuable information for importers of Japanese goods. For example, there is a market for Japanese automobile imports due to the Japan's strict emission regulations (Cabinet, 2003) that make a vehicle unfit for Japan but seen by many countries to be a relatively new vehicle in good condition. When importing, a bill of lading is issued as a receipt to the importer from the shipping carrier as a successfully received good. An order bill of lading contract is a common type of bill of lading that doesn't require the importer to purchase the goods prior to shipment (Buckley, 2004). Often these contracts allow payment to be made within 6 months of the bill of lading receipt. An importer then has reason to investigate how the exchange rate will act near the time of delivery to the maximum amount of time afterwards in which they must pay in order to purchase at an exchange rate most favourable to them. This then opens the possibility of the importer entering the futures market to hedge against spot rate volatility and unfavourable appreciations or devaluations of either currency pair. In these type of international trade situations, forecasting a quarterly exchange rate instead of a daily rate is of interest and a main reason why we chose to forecast the quarterly Yen-U.S Dollar spot exchange rate.

While forecasting the quarterly JPY/USD foreign exchange rate, we calculate model fit and forecast performance measures reasonably representative of different sampling frequencies to evaluate the behaviour of those different sampling frequencies in how they forecast the quarterly exchange rate. The Akaike information criterion (AIC) is used to evaluate model fit (Akaike, 1974); and, the mean absolute percentage forecasting error (MAPE), root mean squared error (RMSE), and mean absolute error (MAE) are calculated to evaluate forecast performance. To obtain reasonably representative measures for each sampling frequency, we perform a computationally intensive number of regressions within each sampling frequency. Computationally intensive here is defined as possibly needing to consider parallel

programming techniques or different statistical software when estimating MIDAS regressions with high frequency observations. This is due to the sheer amount of iterations indicated by equation 4 and the memory requirements of these iterations on higher sampling frequencies.

The sum of regressions and forecasts estimated for all frequencies totals over ten thousand regressions and forecasts. In addition to these regressions and forecasts, model fit and forecast performance measurements are calculated for each iteration indicated in equation 4. Summary statistics are calculated for measurements and sampling frequencies to attempt to construct a reasonable representation of the behaviour of each sampling frequency in forecasting the quarterly exchange rate.

The series of regressions conducted in each sampling frequency differ by their lag lengths. In each MIDAS regression model there are six explanatory variables. Of these six, one is the lagged dependent variable with a lag order defined as p . The other five variables are high frequency variables. The high frequency variables include high frequency lags, lower-case k , up to a maximum lag order previously defined in equation 1 as an upper-case K . The maximum lag order can be different for each high frequency variable. Forecasting results differ depending on the autoregressive p order and the high frequency k lag length of each independent variable. It would be ideal to test every combination of high frequency lag lengths for all five high frequency variables, but is too computationally intensive for this thesis. We set the high frequency variable lag orders to be the same as each other and increase them together as a set instead of increasing each high frequency variable's lag order individually. We set the lag orders to be the same for all high frequency variables and set the lagged dependent variable lag order independently. The optimal number of lags to include for each high frequency variable depends on the underlying structure of the particular model and the specific forecast being conducted. This may be worth looking into in future experiments.

For a single MIDAS regression to be performed in the `midasr` package, the high frequency (HF) data sets must be complete (Ghysels, Kvedaras and Zemlys, 2015, p. 5). That is, theoretically, there must be a consistent number of HF variables for every low frequency (LF) dependent variable observation, i.e., seven days for every week, four weeks for every month, or three months for every quarter. Realistically, there is often a different amount of naturally occurring HF observations per LF date range, i.e., the number of business days per quarter will likely change depending on the amount of holidays occurring. Because of this, our daily sampling frequency variables need to be made into balanced data sets in order to be able to perform MIDAS regressions.

Because of the midasr package requirements for balanced data sets, we create a template to consistently obtain a balanced data set in a way that least affects the results. From the beginning of the data range starting in 1971, this template initially sets NA values for Saturdays and Sundays until the variable's data set is complete. This typically amounts to only two to three NA values inserted at the beginning of a 40-year sample containing over 12,000 daily observations with forecasts that only include 5 years of lagged values. Alternatively, beginning observations could simply be excluded until the data set becomes complete, but this would result in removing 60-63 daily observations and have a much larger impact than our template. The midasr package rightly requires that MIDAS regression have no NA values. Our inserted NA values are replaced with the previous Friday's value. In small data sets this method is discouraged because it places undue emphasis on the Friday's value; but because it is extremely unlikely the memory of our daily time series extends 40 years, the result of inserting NA's values and replacing them with the previous Friday's values at the beginning of a 40-year data set is decidedly insignificant. In smaller samples, further efforts could be made to randomly allocate these NA values within the whole range of sampling data if deemed necessary.

We chose to estimate rolling forecasts instead of recursive forecasts. The reason for this is because of the balanced data set requirements of MIDAS regressions and the midasr package. It takes a considerable amount of time to balance each variable's data set and recursive forecasting would require us to rebalance for every subsequent forecasting horizon. The balanced data set requirement is also a major reason why this thesis only forecasts one date range. For these reasons we use rolling forecasts. It is important to note that the horizon we are forecasting is an extremum, a change in direction of the JPY/USD exchange rate. It's important for a model to be able to forecast changes in direction. Having a model successfully forecast that a long-standing trend will continue to trend in the same direction isn't a strong test of a forecasting model. Evaluating model performance at extreme points in the data is a better indication of performance. The ability for the sampling frequencies to forecast a change in direction and the magnitude of the change is being evaluated. The JPY/USD reaches a low at 2012 Q3 and 2012 Q4 will be the one horizon forecast. The JPY/USD value at 2012 Q4 is not only a change in direction, but also a significant change in terms of magnitude.

Calculating regressions for all combinations of autoregressive and higher frequency lags for five years of observations is extremely computationally intensive. As the sampling frequency increases, the time required for each set of regressions also increases significantly. For example, except for the quarterly sampling frequency, with the immediate computer resources available at the time of this thesis each

frequency could take anywhere up to 18 hours to finish calculating a full set of regression and forecast estimations represented by the table in Figure 13 with the midasr package. The testing associated with calibrating the program to execute took equally as long. Altogether, it's a very long process to obtain model fit and forecast performance calculations that are reasonably representative of each sampling frequency for twenty years of financial data. In the future, it may be pertinent to pursue parallel computing programming methods if using R's midasr package as it will drastically decrease the time required for testing and executing each sampling frequency's set of regressions.

Mathematically, a regression is estimated for every iteration of the following form:

$$E^j = \sum_{p=1}^P \sum_{k=1}^K. \quad (5)$$

Here E is the number of regression estimations, j is the sampling frequency, P is the maximum dependent lag order, and upper-case K is the maximum high frequency lag order. For example, one year of data at the quarterly sampling frequency is a maximum of four lags, i.e., $K = 4$, and at the weekly sampling frequency it is a maximum of fifty-two lags, i.e., $K = 52$. A general example of the process is that starting at $p = 1$, regressions are estimated for high-frequency variables at each lag order of $k = 1, 2, 3 \dots K$. Once K is reached and, consequently, K regressions have been estimated, p is incremented to $p = 2$ and regressions are again estimated for every value of k up to the maximum K . This process, and Equation 4, is visually summarized in Figure 13. Figure 13 displays a table visualizing the lag specifications and the number of measurements calculated for frequency.

For each regression and forecast in equation (5), the Akaike information criterion is calculated (Akaike, 1974). To evaluate forecasting performance, three measures of forecasting errors are calculated: mean absolute percentage error (MAPE), root mean squared error (RMSE), and the mean absolute error (MAE). The formulas used are in Appendix A. Each regression is used to forecast a one-step, rolling, one and two horizon, static forecast. Our original intention of estimating dynamic second horizon forecasts was unsuccessful due to limitations of the midasr R package. This is an area for further research.

		AR(p) → $P = y * q = (5) * (4) = 20$																				
		1	2	3	4	5	6	7	8	9	10	11	...	20								
k	1																					
	2																					
	3																					
	4																					
	5																					
	6																					
	...																					
	$20 * m$																					$20 * m * P$

$K = y * q * m = 20 * m$

Figure 13: Measurement tables visualizing the lag specification from which each measurement is calculated from

Summary statistics are calculated from each table in Figure 13 for each of the AIC, MAPE, RMSE, and MAE measures. The summary statistics for each measure consists of the standard deviation, mean, median, kurtosis, minimum value, and maximum value. Since the minimum value is desirable for all AIC, MAPE, RMSE, and MAE, the minimum value of these measures also represents the best model for that particular frequency. The best model specification in terms of p AR lags and k high frequency lags is also included for each measure. The summary statistics are compared across sampling frequencies to see if any patterns exist in each frequency's distribution.

5. RESULTS

Before analyzing the data, stationarity of the data needs to be achieved. Unit root tests are performed for each variable in each of the four sampling frequencies to determine if data needs to be differenced. The Kwiatkowski-Phillips-Schmidt-Shin (KPSS) (1992), Augmented Dickey-Fuller (ADF) (1984), and the Phillips-Perron (PP) (1988) tests regarding stationarity are used. The KPSS test has a null hypothesis of trend stationarity against the alternative of a unit root (Kwiatkowski et al., 1992). The ADF (1984) and PP (1988) tests have a null hypothesis of a unit root against the alternative of stationarity. The ADF and PP unit root test regressions contain a constant and a linear trend.

Table 2*ADF Unit Root and P-Values*

Sampling						
Frequency	<u>JPY/USD</u>	<u>Nikkei</u>	<u>S&P500</u>	<u>USD/GBP</u>	<u>JPY/GBP</u>	<u>F. Funds</u>
Quarterly	-2.1 (0.536)	-1.9 (0.62)	-2.69 (0.29)	-2.97 (0.17)	-3 (0.159)	-3.86* (0.018)
Monthly	-	-1.95 (0.599)	-2.22 (0.49)	-2.79 (0.24)	-3.06 (0.13)	-2.7 (0.28)
Weekly	-	-1.73 (0.69)	-1.68 (0.715)	-2.85 (0.219)	-2.71 (0.28)	-3.59* (0.03)
Daily	-	-1.73 (0.69)	-1.86 (0.636)	-2.57 (0.337)	-2.52 (0.36)	-3.63* (0.03)

Note: The ADF test has a constant and linear trend

Table 3*PP Unit Root test and P-Values*

Sampling						
Frequency	<u>JPY/USD</u>	<u>Nikkei</u>	<u>S&P500</u>	<u>USD/GBP</u>	<u>JPY/GBP</u>	<u>F. Funds</u>
Quarterly	-10.85 (0.495)	-7.05 (0.71)	-9.42 (0.58)	-14.18 (0.3)	-17.31 (0.126)	-29.92* (0.001)
Monthly	-	-5.91 (0.78)	-7.52 (0.69)	-14.05 (0.326)	-15.21 (0.26)	-28.28* (0.012)
Weekly	-	-5.43 (0.807)	-7.625 (0.684)	-12.03 (0.439)	-12.12 (0.434)	-22.66* (0.041)
Daily	-	-5.05 (0.828)	-7.94 (0.667)	-10.59 (0.52)	-10.56 (0.521)	-25.12* (0.02)

Note: The PP test has a constant and linear trend

Table 4*KPSS Unit Root test and P-Values*

Sampling						
Frequency	<u>JPY/USD</u>	<u>Nikkei</u>	<u>S&P500</u>	<u>USD/GBP</u>	<u>JPY/GBP</u>	<u>F. Funds</u>
Quarterly	3.8* (0.001)	0.71* (0.01)	0.40* (0.01)	0.48* (0.01)	0.22* (0.01)	0.32* (0.01)
Monthly	-	1.41* (0.01)	0.78* (0.01)	0.92* (0.01)	0.40* (0.01)	0.55* (0.01)
Weekly	-	3.04* (0.01)	1.64* (0.01)	1.95* (0.01)	0.83* (0.01)	1.00* (0.01)
Daily	-	6.64* (0.01)	3.62* (0.01)	4.25* (0.01)	1.76* (0.01)	2.27* (0.01)

Note: The null hypothesis is for trend stationarity

Except for the federal funds overnight rate, the unit-root test results in Table 2, Table 3, and Table 4 suggest that all of the variables are non-stationary and that the data should be differenced. The federal funds rate has mixed results. We difference the federal funds rate to avoid the negative effects of not differencing a variable that needs to be versus the lesser negative effects of differencing a variable that doesn't need to be differenced. A complete set of plots of the differenced variables can be seen in Appendix B.

Recall that in Equation 4, m represents the number of high frequency observations that occur per one observation of the dependent variable. Table 5 summarizes the m that we use. Our quarterly and monthly m are straightforward and don't require any explanation. We choose 12 weekly observations and 65 daily observations per quarter for the weekly and daily sampling frequencies. 65 daily observations are chosen because in our data range there was, on average, 65 business trading days per quarter for the last 20 years of data.

Table 5
High-Frequency Independent Variable Observations per 1 Dependent Variable Observation

<u>Frequency</u>	<u>Observations per Quarter (m)</u>
Quarterly	1
Monthly	3
Weekly	12
Daily	65

The computational requirements for estimating every regression representing by Equation 4 and Figure 13 proved to be too great for our daily sampling frequency using normal looping programming methods. Numerous methods were attempted to obtain all combinations of 5 years-worth of daily sampling frequency MIDAS regressions but were unsuccessful. For this reason, we estimate daily sampling frequency regressions for every HF lag order in intervals of five instead of one. For example, instead of estimating regressions with (p, k) combination of $(1,1), (1,2), (1,3) \dots (P, K)$ we estimate regressions for every fifth k in a pattern following $(1,5), (1,10), (1,15) \dots (P, K)$ until maximum K is reached. As displayed in Table 5 and Table 6, this resulted in 5200 regressions and forecasts and 1300 high-frequency lags for each quarter in the daily sampling frequency.

Frequency Comparisons Horizon 1

Table 6: Comparing Results Across Sampling Frequency: Horizon 1 Forecasts

Number of Regressions & Forecasts Per Frequency								
Quarterly	400	Weekly	5200					
Monthly	1200	Daily	5200					
Akaike Information Criterion (AIC)								
	SD	Mean	Median	Kurtosis ²	Minimum	Maximum	Best Model	
							AR(p)	0: k ¹
Quarterly	15.21	369.84	372.11	0.024	327.52	421.88	2	20*
Monthly	15.44	370.26	371.88	0.098	319.14	421.02	2	60*
Weekly	12.56	358.50	359.62	0.668	302.16	400.38	2	260*
Daily	16.02	359.36	358.75	-0.109	311.57	418.80d4	20*	1005
MEAN ABSOLUTE PERCENTAGE ERROR (MAPE)								
	SD	Mean	Median	Kurtosis ²	Minimum	Maximum	Best Model	
							AR(p)	0: k ¹
Quarterly	2.603	8.78	9.37	1.36	0.024	13.01	5	17
Monthly	2.07	9.37	9.8	0.709	0.79	14.2	5	52
Weekly	2.5	8.27	8.21	-0.11	0.102	17.85	8	221
Daily	2.75	15.86	15.84	1.27	1.26	29.22	7	1160
ROOT MEAN SQUARED ERROR (RMSE)								
	SD	Mean	Median	Kurtosis ²	Minimum	Maximum	Best Model	
							AR(p)	0: k ¹
Quarterly	2.26	7.61	8.12	1.36	0.02	11.27	5	17
Monthly	1.8	8.12	8.49	0.709	0.69	12.31	5	52
Weekly	2.16	7.16	7.11	-0.11	0.09	15.47	8	221
Daily	2.38	13.74	13.72	1.273	1.09	25.32	7	1160
MEAN ABSOLUTE ERROR (MAE)								
	SD	Mean	Median	Kurtosis ²	Minimum	Maximum	Best Model	
							AR(p)	0: k ¹
Quarterly	2.26	7.61	8.12	1.36	0.02	11.27	5	17
Monthly	1.8	8.12	8.49	0.709	0.69	12.31	5	52
Weekly	2.16	7.16	7.11	-0.11	0.09	15.47	8	221
Daily	2.38	13.74	13.72	1.273	1.09	25.32	7	1160

Note: The maximum amount of p and K lags are equivalent to 20 quarters worth

1: 0:k is defined as zero to k high frequency lags in units of respective independent variable's sampling frequency

2: Excess kurtosis

*: Maximum lags used

Frequency Comparisons Horizon 2

Table 7: Comparing Results Across Sampling Frequencies: 2 Horizon Forecasts

Number of Regressions & Forecasts Per Frequency								
Quarterly	400	Weekly	5200					
Monthly	1200	Daily	5200					
Akaike Information Criterion (AIC)								
	<u>SD</u>	<u>Mean</u>	<u>Median</u>	<u>Kurtosis²</u>	<u>Minimum</u>	<u>Maximum</u>	<u>Best Model</u>	
							<u>AR(p)</u>	<u>0: k¹</u>
Quarterly	15.17	365.05	367.05	0.144	316.47	396.72	1	20*
Monthly	14.71	365.51	366.5	0.024	317.46	409.15	1	57
Weekly	16.51	361.07	360.84	-0.36	311.81	417.84	1	260*
Daily	16.14	362.4	361.46	-0.513	316.42	412.71	1	1255
MEAN ABSOLUTE PERCENTAGE ERROR (MAPE)								
	<u>SD</u>	<u>Mean</u>	<u>Median</u>	<u>Kurtosis²</u>	<u>Minimum</u>	<u>Maximum</u>	<u>Best Model</u>	
							<u>AR(p)</u>	<u>0: k¹</u>
Quarterly	4.29	16.92	17.19	2.002	5.187	34.6	13	17
Monthly	3.89	17.91	18.2	3.24	0.88	37.44	10	54
Weekly	3.54	17.46	17.75	2.88	0.65	44.03	12	242
Daily	3.38	18.51	18.46	2.164	2.5	36.796	8	1185
ROOT MEAN SQUARED ERROR (RMSE)								
	<u>SD</u>	<u>Mean</u>	<u>Median</u>	<u>Kurtosis²</u>	<u>Minimum</u>	<u>Maximum</u>	<u>Best Model</u>	
							<u>AR(p)</u>	<u>0: k¹</u>
Quarterly	4.14	16.06	16.31	2.2	4.68	33.82	13	17
Monthly	3.72	17.02	17.2	3.53	0.86	36.78	10	54
Weekly	3.37	16.7	16.93	3.06	0.71	42.67	12	242
Daily	3.32	17.64	17.55	2.12	2.84	35.83	19	460
MEAN ABSOLUTE ERROR (MAE)								
	<u>SD</u>	<u>Mean</u>	<u>Median</u>	<u>Kurtosis²</u>	<u>Minimum</u>	<u>Maximum</u>	<u>Best Model</u>	
							<u>AR(p)</u>	<u>0: k¹</u>
Quarterly	3.92	15.45	15.68	2.03	4.68	31.72	13	17
Monthly	3.55	16.36	16.61	3.29	0.78	34.33	10	54
Weekly	3.23	15.96	16.22	2.91	0.57	40.32	12	242
Daily	3.11	16.88	16.84	2.14	2.17	33.71	8	1185

Note: The maximum amount of p and K lags are equivalent to 20 quarters worth

1: 0:k is defined as zero to k high frequency lags in units of respective independent variable's sampling frequency

2: Excess kurtosis

*: Maximum lags used

Table 6 and Table 7 display the results from the first and second-horizon summary statistics. Recall that the summary statistics are calculated from the model fit and performance measurement tables represented by Figure 13. Table 6 and Table 7 are arranged by model fit and performance measurements. For results that are arranged by sampling frequency, see Appendix C.

6. DISCUSSION

If including high-frequency data doesn't add to explaining the variation in the underlying data generating process for this time-series experiment, then we'd expect to see no significant difference between the quarterly AIC results in comparison to the other higher-frequency results. There would be no relative difference because the included observations would not lend to reducing the sum of squared residuals and the AIC (see Appendix A for AIC formula used). However, what we observe in Table 8 is a difference in the average AIC values for the weekly frequency. The monthly and daily frequencies produce AIC's that are generally the same as the quarterly AIC's thus suggesting they don't necessarily fit better models than the quarterly frequency, on average. This suggests that, for this experiment, including high-frequency variables at the weekly sampling frequency may result in better fitting models and a lower sum of squared residuals than not using MIDAS.

If including high-frequency data helps our forecasts, then we'd expect the mean forecasting performance errors to be smaller than the quarterly frequency results. Mean forecasting errors larger than the quarterly errors suggest a sampling frequency where higher-frequency observations actually hinder forecasting the dependent variable. Results in Table 8 show that forecasts using variables in the daily frequency consistently produce poorer forecasts than the forecasts in the other frequencies. The degree to which they are poorer is quite large as well. Often, the forecast errors from variables in the daily frequency are almost twice in magnitude relative to the other frequencies. These results may occur because variables in the daily frequency include a large amount of observations for every one quarterly dependent variable observation. There may be an excessive amount of irrelevant observations in the daily frequency that skew the quarterly forecast.

Another reason why the significant amount of observations from variables in the daily frequency may produce poor forecasts might be because of the way that MIDAS estimates hyper-parameters and how we've estimated our forecast models. Too many irrelevant high-frequency daily observations may skew the MIDAS hyper-parameters; and, because our models don't remove statistically insignificant variables before using the regression model to forecast, this could become an issue. If an over-abundance of daily

observations causes the hyper-parameters to be incorrect due to incorrect weighting estimations and these statistically insignificant hyper-parameter variables are included in the forecast, the resulting forecast may be largely incorrect. In this way, statistically insignificant noise from daily observations may be incorporated into our daily frequency models resulting in largely incorrect first-horizon quarterly forecasts.

The weekly sampling frequency was able to produce the lowest in-sample forecasting errors for horizon one. This is particularly interesting because the weekly frequency also produced the lowest in-sample AIC values. It is well known that the best fitting econometric models don't necessarily produce the best forecasts, but the data from Table 8 seems to suggest the opposite for this thesis in terms of the weekly first horizon forecasting results. The weekly horizon was able to produce both the lowest in-sample AIC values and the lowest in-sample performance errors. This strongly suggests that if a forecaster were to forecast the quarterly Yen-Dollar spot exchange rate with the variables we've used, they should use the weekly sampling frequency. Moreover, it seems to support the statement that modeling in different frequencies does have the potential to produce results that are characteristic to one sampling frequency that the other frequencies do not have.

There is evidence to suggest that the inclusion of high frequency variables help model the underlying data generating process better than if the independent variables were all at the same low frequency as the dependent variable. The frequencies that produced the lowest AIC mean and median summary statistics are at a higher frequency than the dependent variable. Weekly and daily sampling frequency AIC summary statistic results from horizon 1 are much lower than quarterly and monthly values shown in Table 6. As stated, this may be because there is information about the underlying data generating process available in the higher frequencies that is inadvertently being removed when transforming independent variable observations to a lower frequency. The same relationship is also observed in the horizon 2 results, but to a lesser degree. It's worth noting that better modeled data doesn't necessarily mean better forecasts. This is evident in Table 6 where the daily frequency has a relatively low in-sample AIC's but significantly poorer forecasts than the other frequencies.

The explaining power of high frequency variables may decrease as the forecasting horizon increases. The lag lengths of the models that produced the best in-sample MAPE for horizon 2 forecasts have much higher autoregressive orders and much lower high frequency lag length orders than their horizon 1 counterpart as shown in Table 7. Furthermore, results in Table 7 show that the quarterly sampling

frequency produces the best in-sample forecasts. This may be because the lower frequency variables capture the longer term trends better while the observations in the higher frequencies skew the results.

7. FURTHER RESEARCH

There are certain elements to this experiment that were outside the scope of this thesis, due to complexity or time constraints, which may be worth looking into for further research. For instance, instead of only including lagged dependent variables in their original quarterly frequency, it may be interesting to instead include lagged dependent variables at various higher frequencies within the MIDAS framework. As stated earlier, this would not work for the nowcasting environment, because if you were to include a higher frequency version of the dependent variable in an intra-period forecast, you might reconsider the original dependent variable frequency and subsequent forecasting horizon. A higher frequency lagged dependent variable estimated via MIDAS regressions may add to forecasting performance.

In this experiment we held the sampling frequency of the dependent variable constant while changing the frequencies of the independent variable. However, if the dependent variable frequency was also changed, different insights may arise. The behaviour of AIC's and forecast errors of a quarterly dependent variable to other independent variable sampling frequencies may be different than that of a monthly dependent variable to other independent sampling frequencies. The same may be true for weekly as well. The distance between dependent and independent variable sampling frequency may shed insight as to which sampling frequency is optimal for each dependent variable sampling frequency.

Due to time constraints from the intensive computational nature of this experiment and the quantity of regressions and forecasts that are estimated, our experiment only tested a single date range. A model that performs well in one date range will not necessarily perform well in another date range. A model that consistently performs well across various date ranges is stronger than one that does not. Further research might include robustness checks surrounding differing date ranges. For each date range, an additional extension of this thesis would do well to also include predictive accuracy testing for our forecasts. To further expand upon improve forecasting performance, it isn't necessary to only choose one sampling frequency, combining forecasts can also improve forecasting performance.

This experiment had the same lag order for all of the explanatory variables in each regression. The high frequency lag orders did not differ between explanatory variables. This may not be optimal. Ideally, hyper-parameters will correctly capture the influence of an explanatory variable with too many high

frequency lags without the negative effects of over-specification. Though extremely computationally intensive, it may be worth investigating all possible combinations of high frequency lag lengths.

It's important to keep in mind that our explanatory variables were chosen largely due to their high frequency availability and are not intended to be structural models. From the results, we see that they appear to be able to explain the variance in quarterly JPY/USD spot rates to some degree. Forecasting performance might likely be increased if a highly optimized selection of explanatory variables known to explain the underlying structure of the JPY/USD exchange rate was chosen. This choice of variables might consider incorporating variables like U.S and Japan's quantitative easing market interventions and relative trade balance. Of course, if structural modelling was the primary goal, repeating this experiment with different sets of restrictions would be ideal.

8. CONCLUSIONS

Transforming high frequency data sets into low frequency data sets potentially removes valuable information about the underlying data generating process of a series. Given an overabundance of data resources and mixed frequency data sampling (MIDAS) techniques, researchers have to make a decision whether this potentially valuable information should be included in their models and which higher frequency is optimal. We use financial data and the mixed data sampling technique to estimate thousands of regressions and forecasts in four different sampling frequencies with models of low and high frequency lag length combinations up to five year's worth. Our regressions contain a quarterly autoregressive component and five independent variables that, as a set, are estimated in quarterly, monthly, weekly, and daily sampling frequencies. These regressions are used to forecast the quarterly Yen- U.S Dollar spot exchange with rolling, one and two-step ahead, static forecasts. From each regression, we calculate the Akaike (1974) information criterion for a model fit measure; and, from each forecast, we calculate the mean absolute percentage error, root mean squared error, and mean absolute error for forecast performance error measures. We arrange these measures by frequency and calculate summary statistics to compare relative performance of each sampling frequency.

We find that the weekly sampling frequency produces the lowest in-sample AIC measures and also the best in-sample forecasts for horizon 1 forecasts. However, the explanation of the best fitting models producing the best forecasts for our sample does not correspond with the extremely poor in-sample forecast results from the daily frequency despite their relatively low in-sample AIC measures. These relatively low in-sample AIC measures for our highest sampling frequencies also give evidence to suggest

that including high frequency variables may model the underlying data generating process better than if all variables were transformed to the dependent variable's sampling frequency. However, our results show that including high frequency variables don't necessarily produce better forecasts.

9. REFERENCES

- Akaike, H. (1974). A new look at the statistical model identification. *IEEE Transactions on Automatic Control*, 19(6), 716–723. doi:10.1109/TAC.1974.1100705
- Andreou, E., Ghysels, E., & Kourtellos, A. (2010). Regression models with mixed sampling frequencies. *Journal of Econometrics*, 158, 246-261.
- Andreou, E., Ghysels, E., & Kourtellos, A. (2013). Should macroeconomic forecasters use daily financial data and how? *Journal of Business & Economic Statistics*, 31(2), 240-251. doi:10.1080/07350015.2013.767199
- Armesto, M., Engemann, K., & Owyang, M. (2010). Forecasting with mixed frequencies. *Federal Reserve Bank of St. Louis Review*, 92(6), 521-536.
- Atkins, A. B., & Dyl, E. A. (1997). Market Structure and Reported Trading Volume: NASDAQ versus the NYSE. *The Journal of Financial Research*(3), 291-304.
- Bank of Japan. (2016, January 29). *Quantitative and Qualitative Monetary Easing with a Negative Interest Rate*. Retrieved from Bank of Japan: <https://www.boj.or.jp/en/mopo/outline/qqe.htm/>
- Board of Governors of the Federal Reserve. (2016, 05 01). *Economic Research & Data: USD/JPY Exchange Rate*. Retrieved from Federal Reserve: <http://www.federalreserve.gov/datadownload/Choose.aspx?rel=H10>
- Buckley, A. (2004). *Multinational Finance*. Harlow, UK: Pearson Education Limited. ISBN: 978-0-27-368209-7.
- Cabinet, P. M. (2003, March 28). *Opening Statement by Prime Minister Junichiro Koizumi at the Press Conference on the Passage of the FY2003 Budget*. Retrieved from Prime Minister of Japan and His Cabinet: http://japan.kantei.go.jp/koizumispeech/2003/03/28yosan_e.html
- Chortareas, G., Jiang, Y., & Nankervis, J. C. (2011). Forecasting exchange rate volatility using high-frequency data: Is the euro different? *International Journal of Forecasting*, 27(2011), 1089-1107.
- Clements, M., & Galvão, A. (2008). Macroeconomic forecasting with mixed-frequency data: Forecasting output growth in the United States. *Journal of Business & Economic Statistics*, 26(4), 546-554. doi:10.1198/073500108000000001
- 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.
- Duan, J. (2015, September 23). *University of Victoria Electronic Theses and Dissertations*. Retrieved May 15, 2016, from University of Victoria Library: <http://hdl.handle.net/1828/6711>
- Eviews. (2016). Retrieved from What's New in EViews 9.5: <http://www.eviews.com/>
- Galvão, A. (2013). Changes in predictive ability with mixed frequency data. *International Journal of Forecasting*, 29, 395-410.
- Ghysels, E., Kvedaras, V., & Zemlys, V. (2016). Mixed frequency data sampling regression models: The R package midasr. *Statistical Software*, 74(4), 1-35. doi:10.18637/jss.v072.i04

- Ghysels, E., Santa-Clara, P., & Valkanov, R. (2004). The MIDAS touch: Mixed data sampling regression models. *Working Paper*.
- Kwiatkowski, D., Phillips, P. C., Schmidt, P., & Shin, Y. (1992). Testing the null hypothesis of stationarity against the alternative of a unit root. *Journal of Econometrics*, 54(1-3), 159-178. doi:10.1016/0304-4076(92)90104-Y
- Li, X., Shang, W., Wang, S., & Ma, J. (2015). A MIDAS modelling framework for Chinese inflation index forecast incorporating Google search data. *Electronic Commerce Research and Applications*, 14, 112-125.
- Marcellino, M. a. (2010). Factor MIDAS for Nowcasting and Forecasting with Ragged-Edge Data: A Model Comparison for German GDP. *Oxford Bulletin of Economics and Statistics*, 72, 518-550. doi: 10.1111/j.1468-0084.2010.00591.x.
- Mogi, C., & Komiya, H. (2016, March 2). *Japan's Three Biggest Banks Declare Yen's Depreciation Is Over*. Retrieved from Bloomberg News: <http://www.bloomberg.com/news/articles/2016-03-01/japan-s-three-biggest-banks-declare-yen-s-depreciation-is-over>
- Nagarajan, R. (2009, February 24). *DJIA: A Highly Flawed Benchmark*. Retrieved August 10, 2016, from The Rational Walk: <http://www.rationalwalk.com/?p=146>
- Nikkei. (2016, 05 01). *Historical Data (Nikkei)*. Retrieved from Nikkei Indexes: <http://indexes.nikkei.co.jp/en/nkave/statistics/>
- Okina, K., Shirakawa, M., & Shiratsuka, S. (2001, February). The asset price bubble and monetary policy: Japan's experience in the late 1980s and the lessons. *Monetary and Economic Studies at the Bank of Japan*.(Special Edition), 395-450.
- Phillips, P. C., & Perron, P. (1988). Testing for a unit root in time series regression. *Biometrika*, 75(2), 335-346. doi:10.1093/biomet/75.2.335
- Qian, H., & Ghysels, E. (2016, August 6). *MIDAS Matlab Toolbox*. Retrieved August 10, 2016, from Mathworks File Exchange: <https://jp.mathworks.com/matlabcentral/fileexchange/45150-midas-matlab-toolbox>
- Quandl. (2016, 05 01). Retrieved from Quandl Finance & Economic Data: <https://www.quandl.com/>
- Ryan, J. A., & Ulrich, J. M. (2015, February 20). *xts: eXtensible Time Series*. Retrieved from CRAN R-Project: <https://cran.r-project.org/package=xts>
- Said, S. E., & Dickey, D. A. (1984). Testing for unit roots in autoregressive-moving average models of unknown order. *Biometrika*, 71(3), 599-607. doi:10.1093/biomet/71.3.599
- St. Louis Federal Reserve. (2016, 05 01). *Federal Funds Rate Data*. Retrieved from St. Louis Federal Reserve Research: <https://research.stlouisfed.org/fred2/data/DFF.csv>
- USGS. (2011, April 5). *Magnitude 9.03 – Near The East Coast Of Honshu, Japan*. Retrieved from United States Geological Survey (USGS): <http://www.webcitation.org/5xgj6FuHC>

Yahoo Finance. (2016, 05 01). *S&P 500 Index*. Retrieved from
[http://ichart.finance.yahoo.com/table.csv?s=%5EGSPC&a=1&b=1&c=1902&d=07&e=4&f=2039
&g=d&ignore=.csv](http://ichart.finance.yahoo.com/table.csv?s=%5EGSPC&a=1&b=1&c=1902&d=07&e=4&f=2039&g=d&ignore=.csv)

APPENDICES

APPENDIX A: MODEL SELECTION CRITERIA AND FORECASTING ERRORS

Akaike Information Criterion (AIC)

The AIC's are calculated using equation (A.1).

$$\text{AIC} = n \log(\hat{\sigma}^2) + 2k \quad (\text{A.1})$$

Where $\hat{\sigma}^2 = \frac{SSR}{n-k}$, SSR is the sum of squared residuals, n is the number of observations, and k is the number of model parameters.

Mean Absolute Percentage Error (MAPE)

The MAPE's are calculated using equation (A.2).

$$\text{MAPE} = 100 * \frac{1}{n} \sum_{t=1}^n \left| \frac{A_t - F_t}{A_t} \right| \quad (\text{A.2})$$

Where n is the number of forecasts, A is the actual value, and F is the forecasted value. The multiplication of 100 makes it a percentage error.

Root Mean Squared Error (RMSE)

The RMSE's are calculated using equation (A.3).

$$\text{RMSE} = \sqrt{\frac{1}{n} \sum_{t=1}^n (A_t - F_t)^2} \quad (\text{A.3})$$

Where n is the number of forecasts, A is the actual value, and F is the forecasted value.

Mean Absolute Error (MAE)

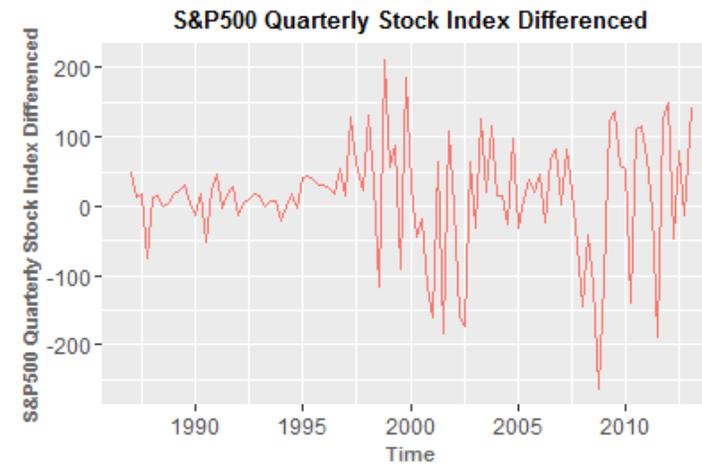
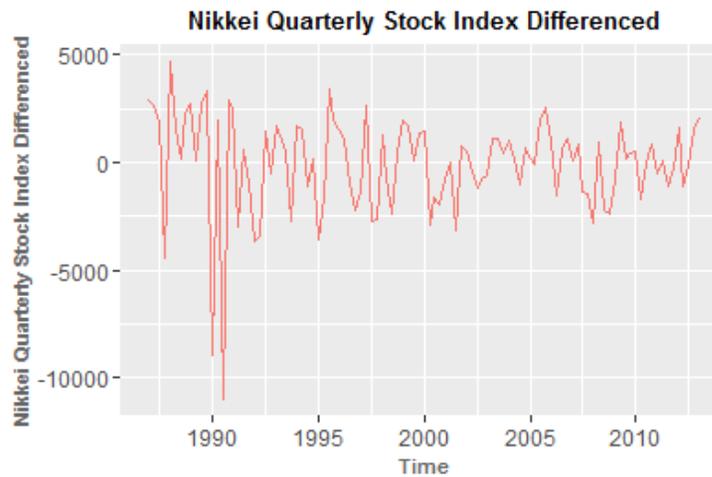
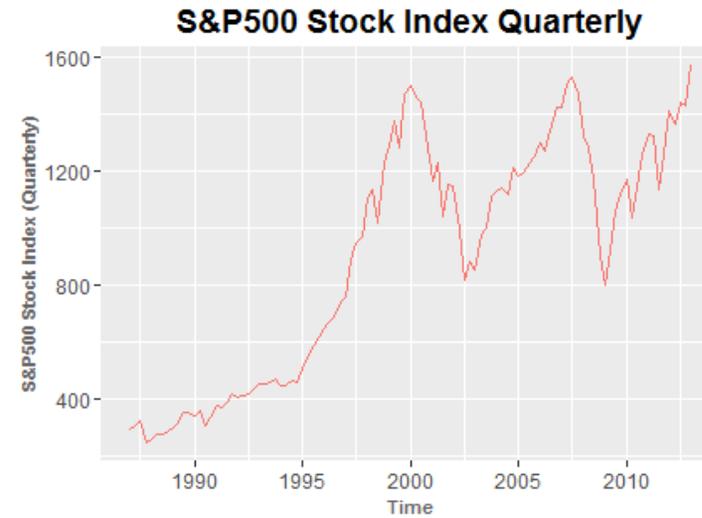
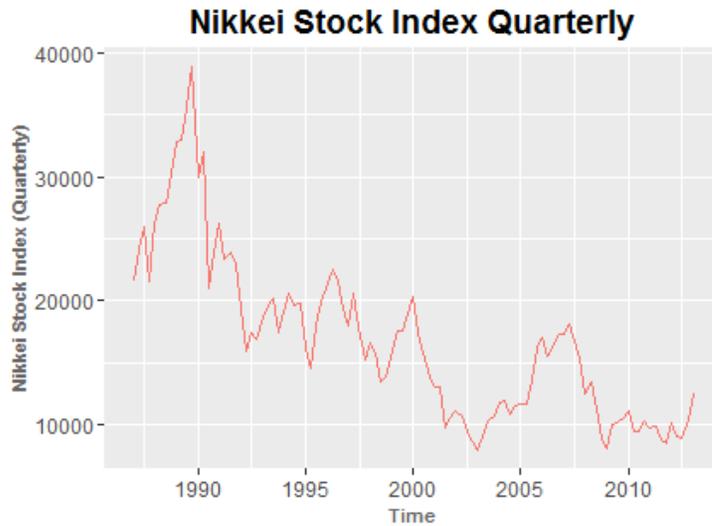
The MAE's are calculated using equation (A.4).

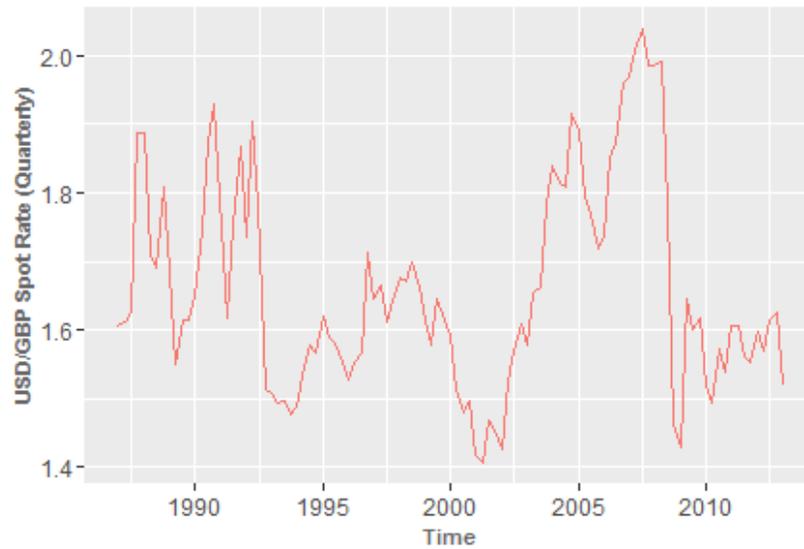
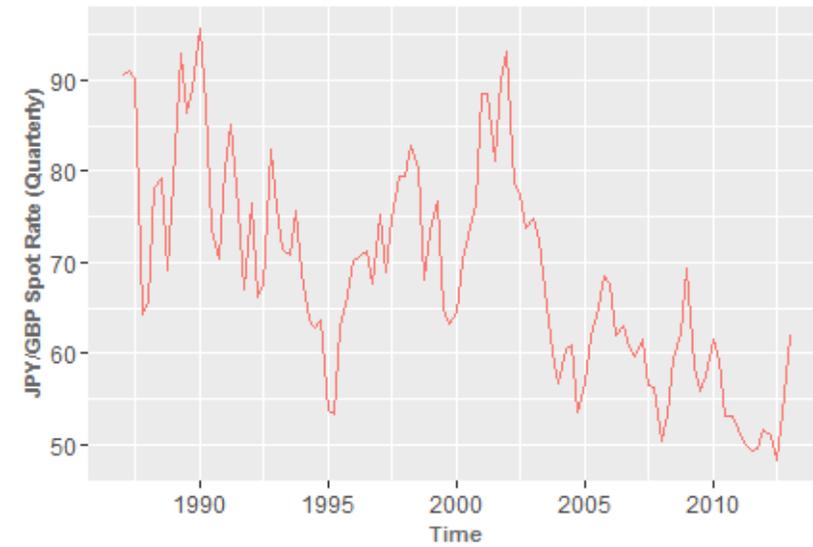
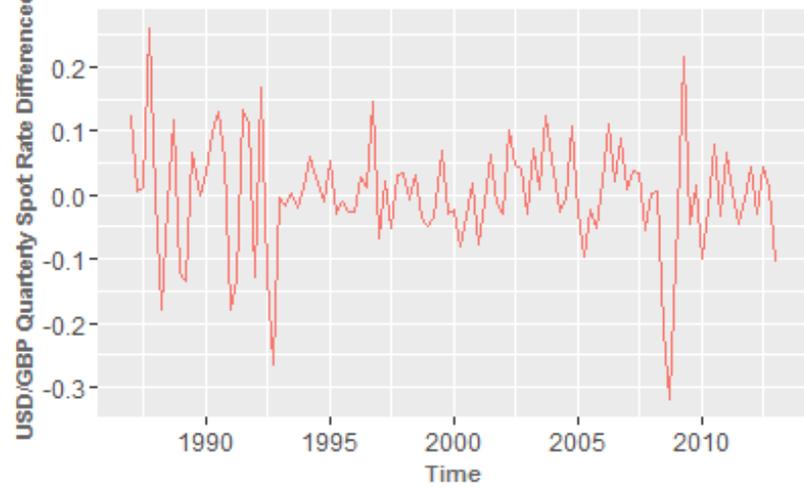
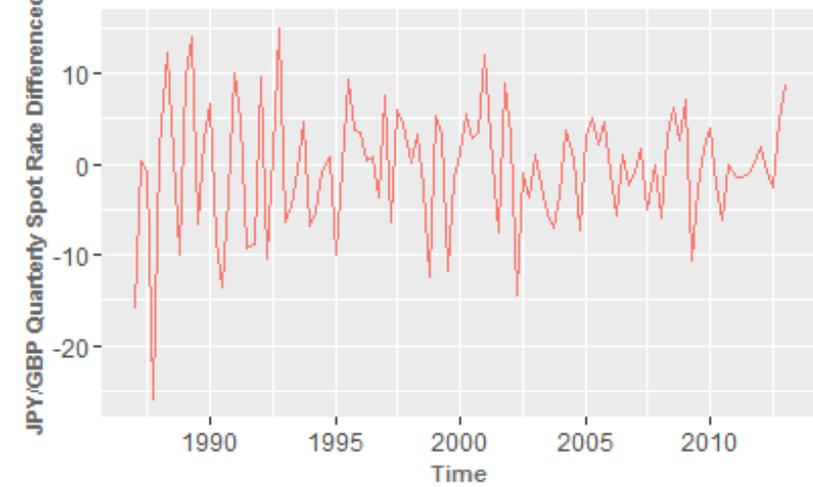
$$\text{MAE} = \frac{1}{n} \sum_{t=1}^n |A_t - F_t| \quad (\text{A.4})$$

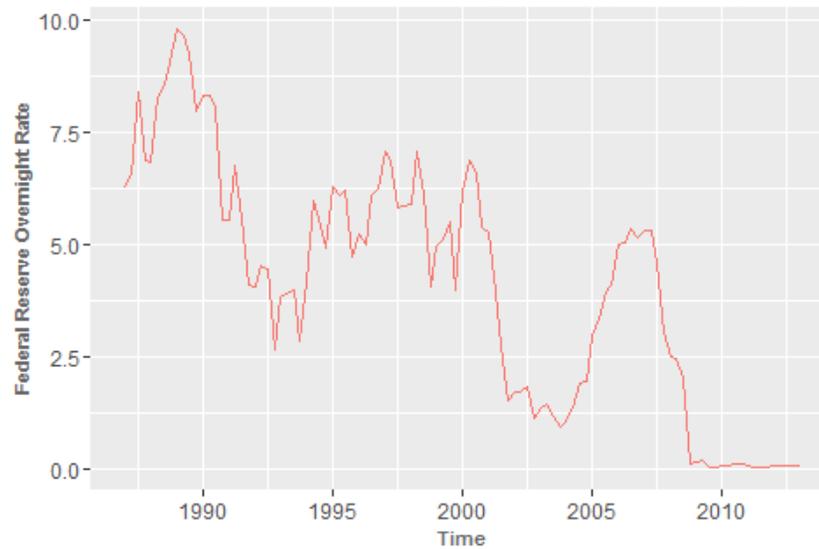
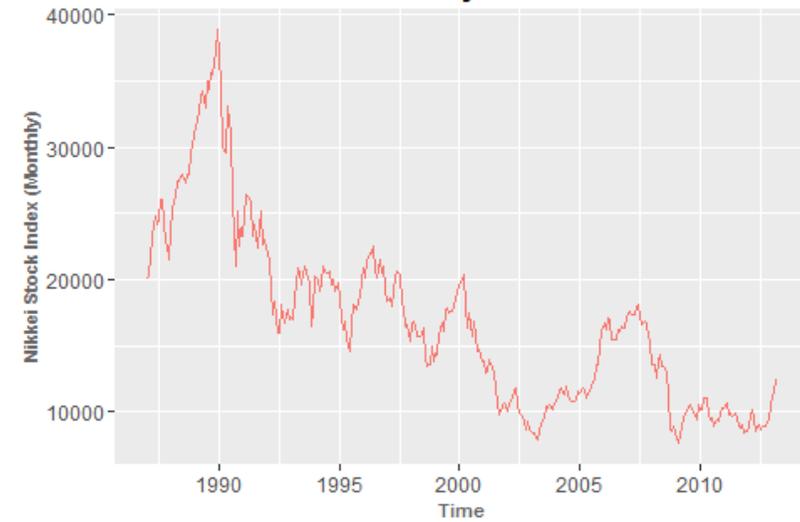
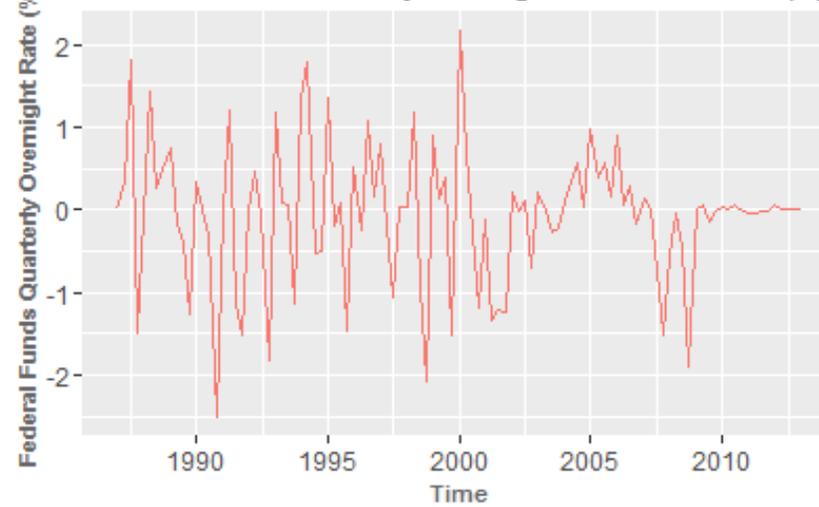
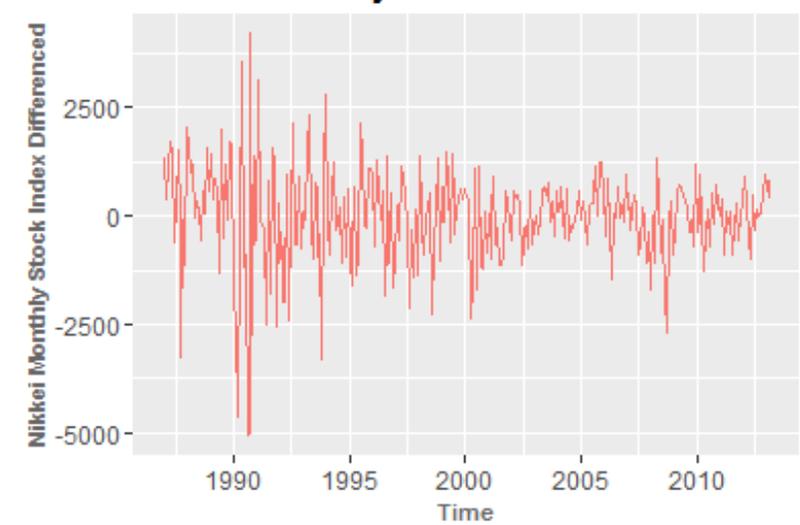
APPENDIX B: DATA PLOTS

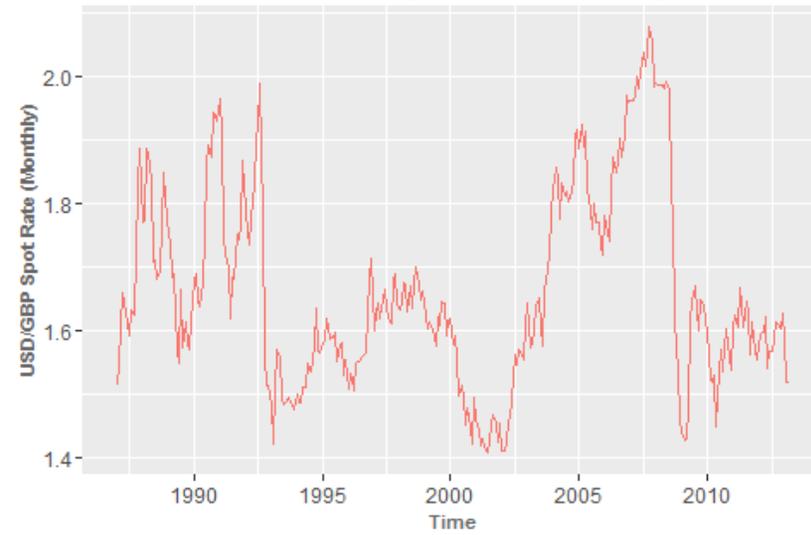
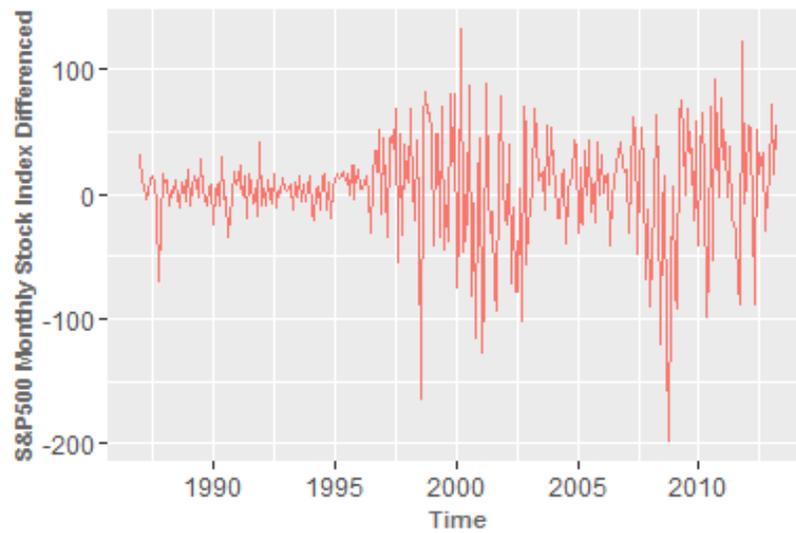
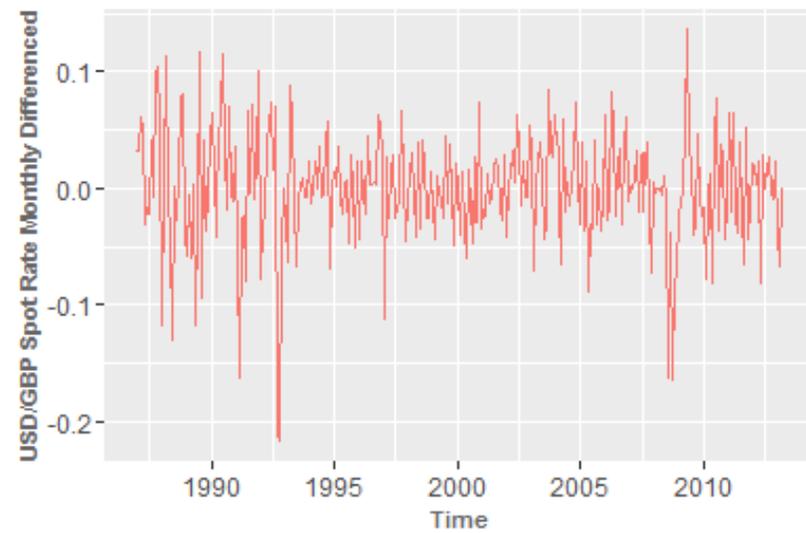
The following are plots for the independent variables sampling frequencies, differenced and undifferenced

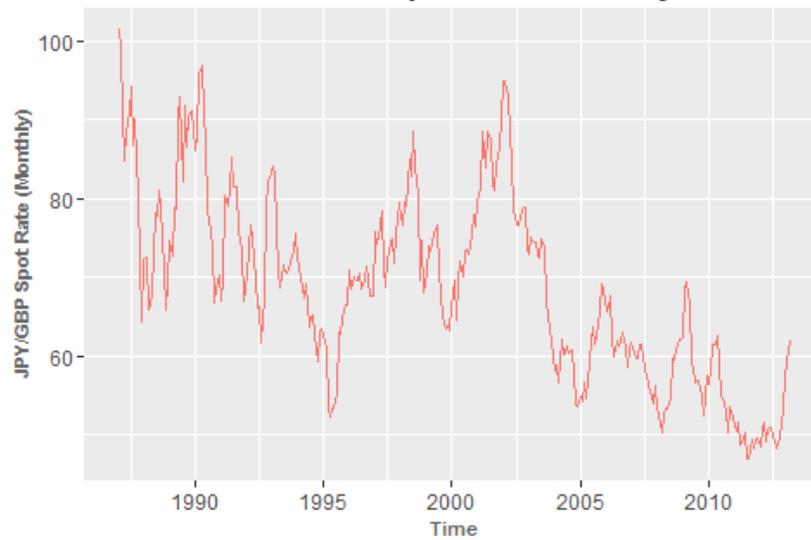
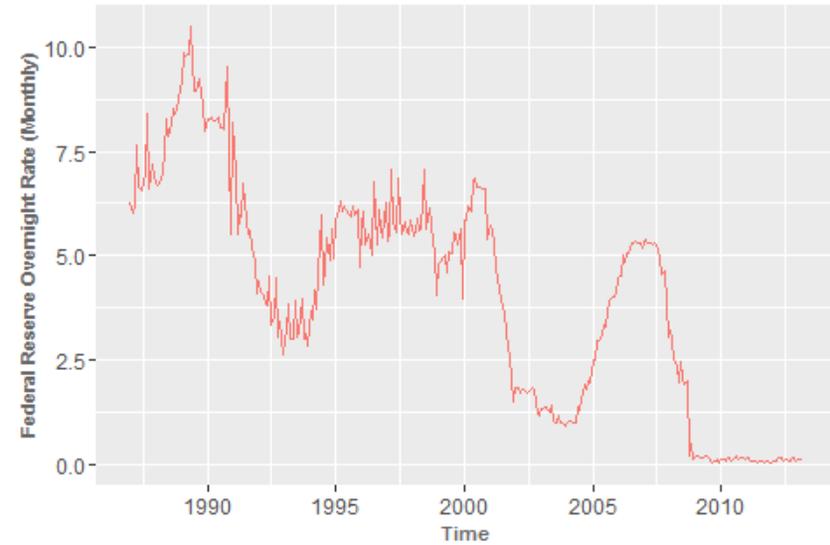
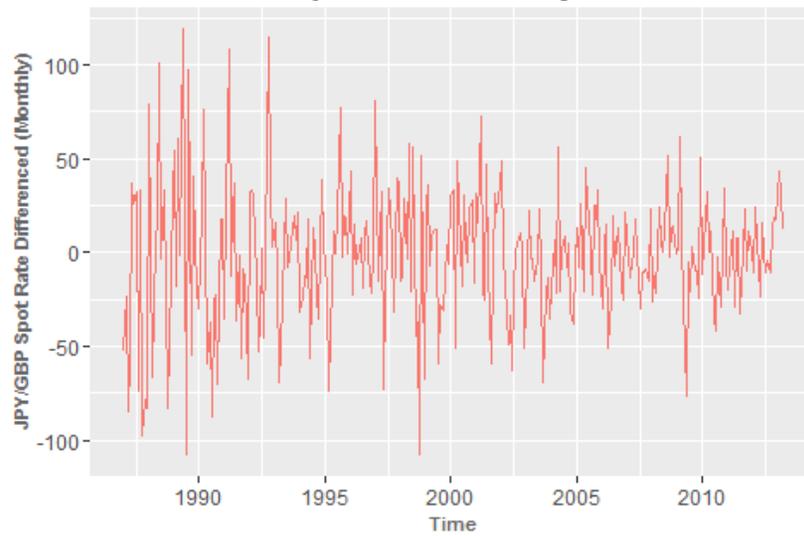
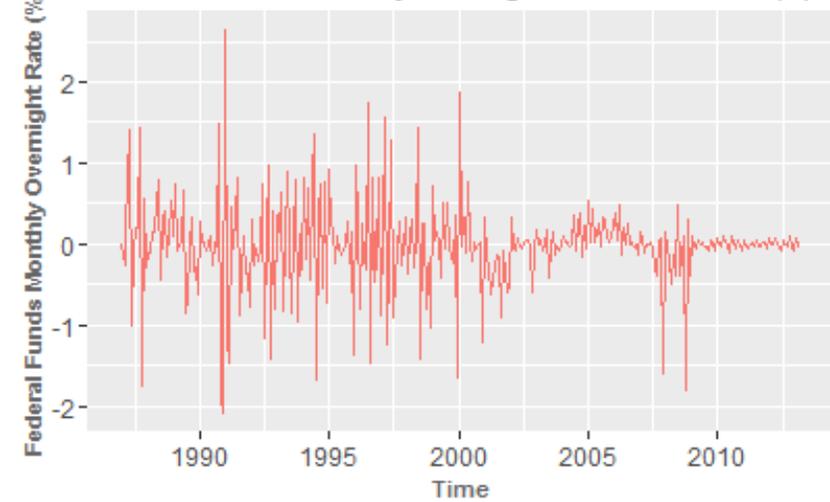
QUARTERLY FREQUENCY

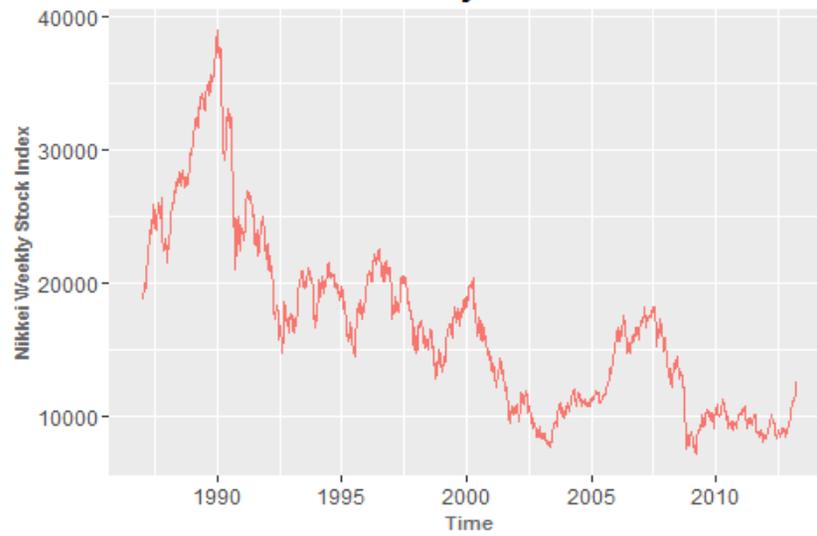
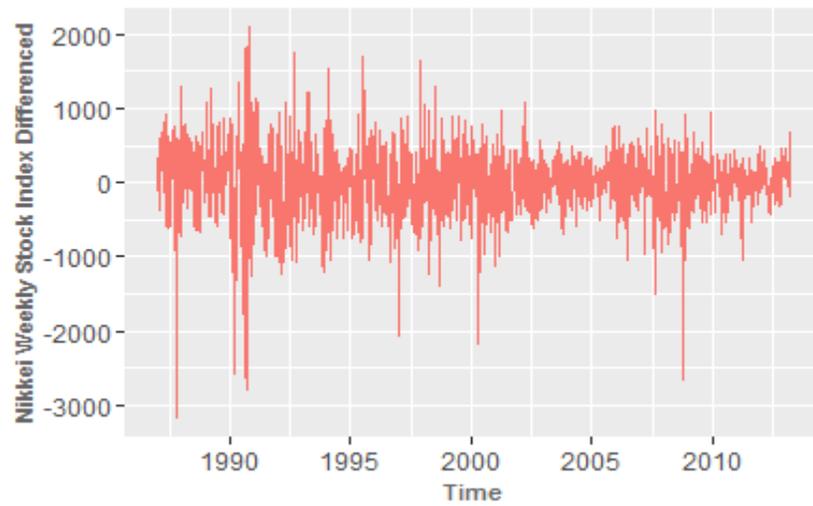
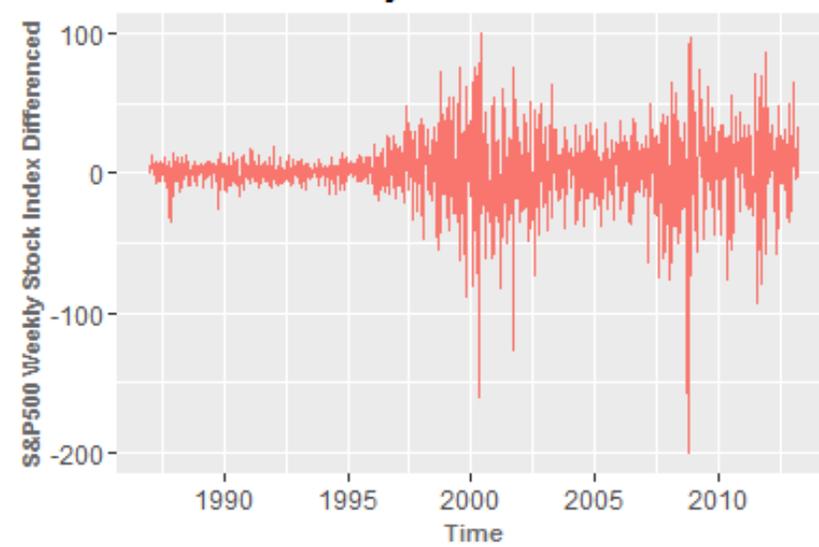


USD/GBP Quarterly Spot Rate**JPY/GBP Quarterly Spot Rate****USD/GBP Quarterly Spot Rate Differenced****JPY/GBP Quarterly Spot Rate Differenced**

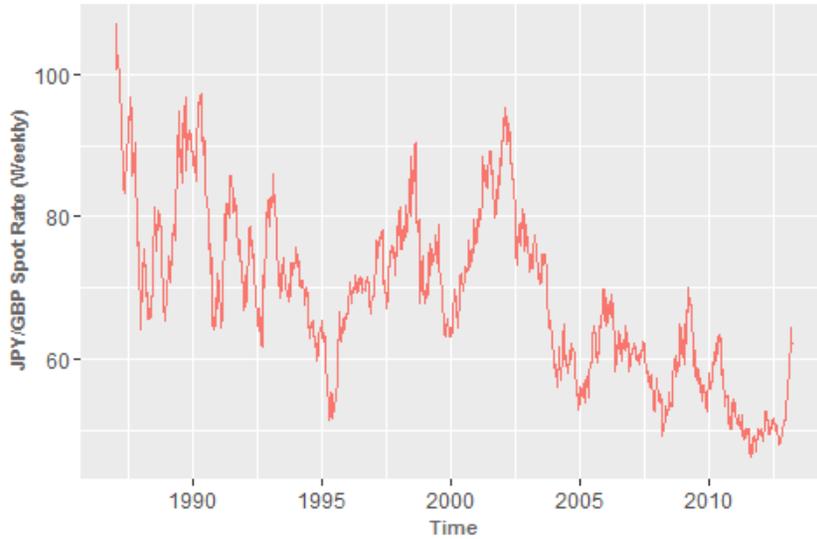
Federal Reserve Overnight Rate (Quarterly)**MONTHLY FREQUENCY****Nikkei Monthly Stock Index****Federal Funds Quarterly Overnight Rate Differenced (%)****Nikkei Monthly Stock Index Differenced**

S&P 500 Monthly Stock Index**USD/GBP Spot Rate Monthly****S&P 500 Monthly Stock Index Differenced****USD/GBP Spot Rate Monthly Differenced**

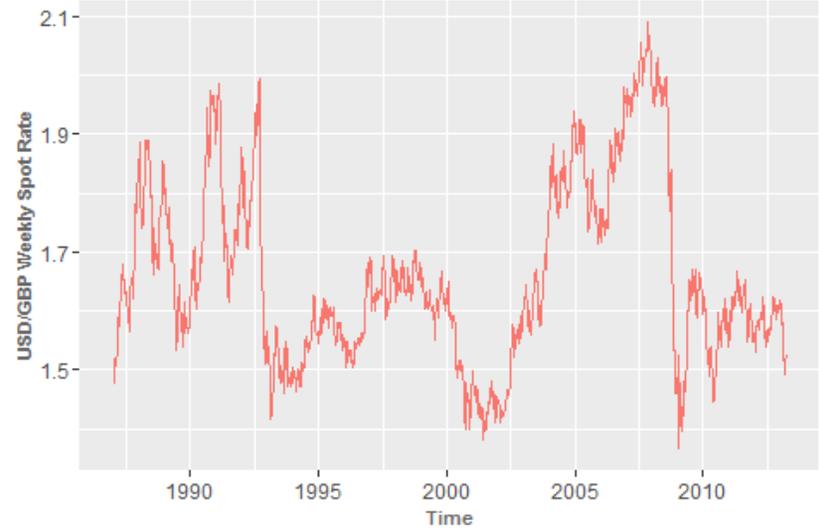
JPY/GBP Spot Rate Monthly**Federal Reserve Overnight Rate (Monthly)****JPY/GBP Spot Rate Monthly Differenced****Federal Funds Monthly Overnight Rate Differenced(%)**

Nikkei Weekly Stock Index**S&P500 Weekly Stock Index****Nikkei Weekly Stock Index Differenced****S&P500 Weekly Stock Index Differenced**

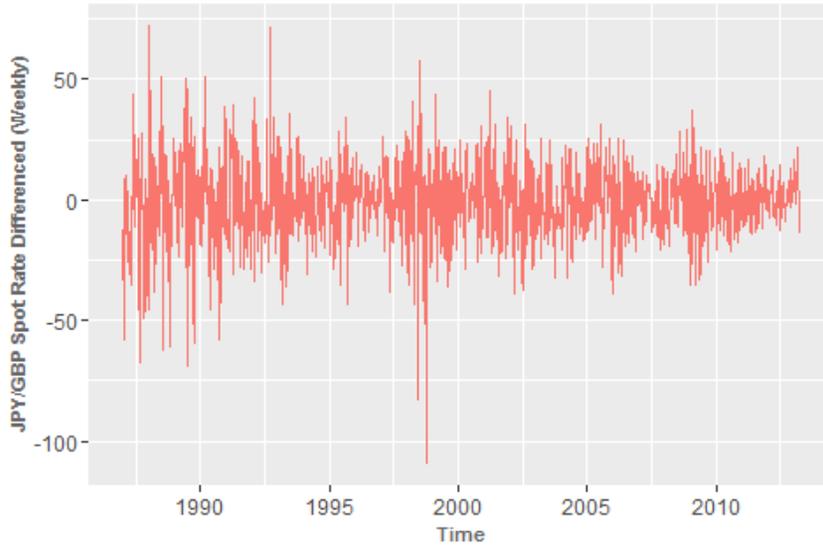
JPY/GBP Spot Rate Weekly



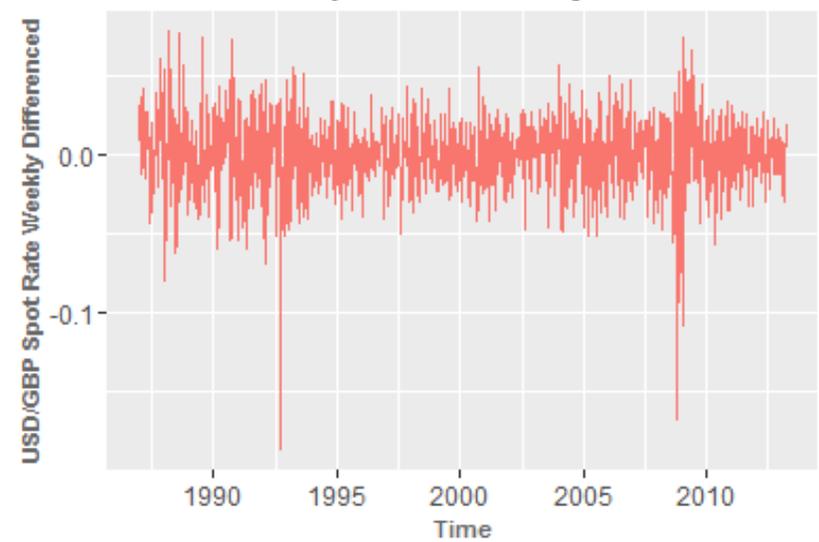
USD/GBP Weekly Spot Rate

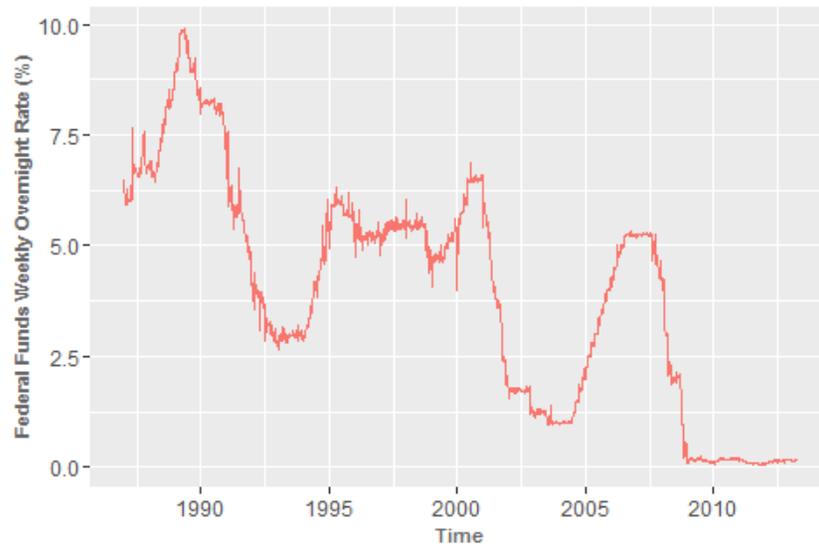
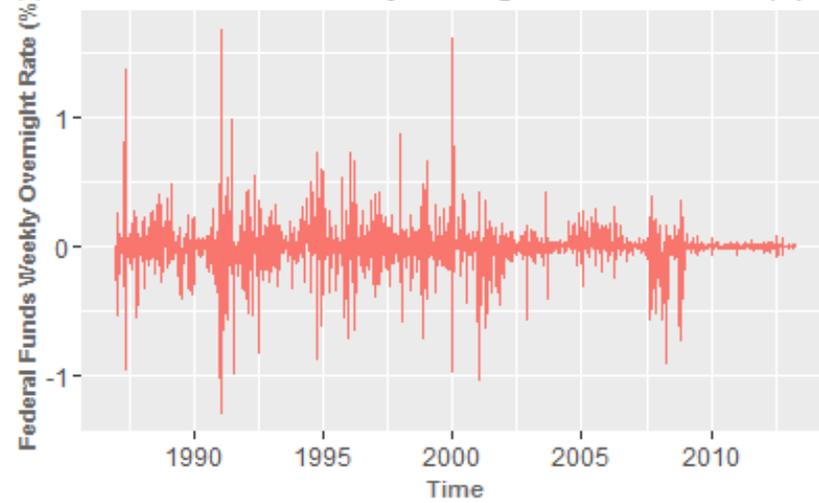


JPY/GBP Spot Rate Weekly Differenced



USD/GBP Spot Rate Weekly Differenced



Federal Funds Weekly Overnight Rate (%)**Federal Funds Weekly Overnight Rate Differenced(%)**

APPENDIX C: PERFORMANCE MEASURE SUMMARY STATISTICS BY FREQUENCY

Horizon 1 Forecasts

Table 8: Summary Statistic Results by Frequency: Horizon 1

QUARTERLY								
<u>#Regressions & Forecasts</u>	<u>Maximum Possible P</u>				<u>Maximum Possible HF Lags (K)</u>			
400	20				20			
	<u>Best Model</u>							
	<u>SD</u>	<u>Mean</u>	<u>Median</u>	<u>Kurtosis²</u>	<u>Minimum</u>	<u>Maximum</u>	<u>AR(p)</u>	<u>0: k¹</u>
AICS	15.21	369.84	372.11	0.0235	327.52	421.88	2	20*
MAPES	2.603	8.78	9.37	1.36	0.024	13.01	5	17
RMSES	2.26	7.61	8.12	1.36	0.02	11.27	5	17
MAES	2.26	7.61	8.12	1.36	0.02	11.27	5	17
MONTHLY								
<u>#Regressions & Forecasts</u>	<u>Maximum Possible P</u>				<u>Maximum Possible HF Lags (K)</u>			
1200	20				60			
	<u>Best Model</u>							
	<u>SD</u>	<u>Mean</u>	<u>Median</u>	<u>Kurtosis²</u>	<u>Minimum</u>	<u>Maximum</u>	<u>AR(p)</u>	<u>0: k¹</u>
AICS	15.44	370.26	371.88	0.098	319.14	421.02	2	60*
MAPES	2.07	9.37	9.8	0.709	0.79	14.2	5	52
RMSES	1.8	8.12	8.49	0.709	0.69	12.31	5	52
MAES	1.8	8.12	8.49	0.709	0.69	12.31	5	52
WEEKLY								
<u>#Regressions & Forecasts</u>	<u>Maximum Possible P</u>				<u>Maximum Possible HF Lags (K)</u>			
5200	20				260			
	<u>Best Model</u>							
	<u>SD</u>	<u>Mean</u>	<u>Median</u>	<u>Kurtosis²</u>	<u>Minimum</u>	<u>Maximum</u>	<u>AR(p)</u>	<u>0: k¹</u>
AICS	12.562	358.498	359.62	0.668	302.16	400.38	2	260*
MAPES	2.5	8.27	8.21	-0.11	0.102	17.85	8	221
RMSES	2.16	7.16	7.11	-0.11	0.09	15.466	8	221
MAES	2.16	7.16	7.11	-0.11	0.09	15.47	8	221
DAILY								
<u>#Regressions & Forecasts</u>	<u>Maximum Possible P</u>				<u>Maximum Possible HF Lags (K)</u>			
5200	20				1300			
	<u>Best Model</u>							
	<u>SD</u>	<u>Mean</u>	<u>Median</u>	<u>Kurtosis²</u>	<u>Minimum</u>	<u>Maximum</u>	<u>AR(p)</u>	<u>0: k¹</u>
AICS	16.018	359.36	358.75	-0.109	311.57	418.804	20*	1005
MAPES	2.75	15.86	15.84	1.27	1.26	29.22	7	1160
RMSES	2.38	13.74	13.72	1.273	1.09	25.32	7	1160
MAES	2.38	13.74	13.72	1.273	1.09	25.32	7	1160

1: 0:k is defined as zero to k high frequency lags in units of respective independent variable's sampling frequency

2: Given values are excess kurtosis

*: Maximum lags used

Horizon 2 Forecasts

Table 9: *Summary Statistic Results by Frequency: Horizon 2*

QUARTERLY														
<i>#Regressions & Forecasts</i>					<i>Maximum Possible P</i>				<i>Maximum Possible HF Lags (K)</i>					
	400				20				20					
													Best Model	
	SD	Mean	Median	Kurtosis²	Minimum	Maximum	AR(p)	0: k¹						
AICS	15.17	365.05	367.05	0.144	316.47	396.72	1	20*						
MAPES	4.29	16.92	17.19	2.002	5.187	34.6	13	17						
RMSES	4.14	16.06	16.31	2.2	4.68	33.82	13	17						
MAES	3.92	15.45	15.68	2.03	4.68	31.72	13	17						
MONTHLY														
<i>#Regressions & Forecasts</i>					<i>Maximum Possible P</i>				<i>Maximum Possible HF Lags (K)</i>					
	1200				20				60					
													Best Model	
	SD	Mean	Median	Kurtosis²	Minimum	Maximum	AR(p)	0: k¹						
AICS	14.71	365.51	366.5	0.024	317.46	409.15	1	57						
MAPES	3.89	17.91	18.2	3.24	0.88	37.44	10	54						
RMSES	3.72	17.02	17.2	3.53	0.86	36.78	10	54						
MAES	3.55	16.36	16.61	3.29	0.78	34.33	10	54						
WEEKLY														
<i>#Regressions & Forecasts</i>					<i>Maximum Possible P</i>				<i>Maximum Possible HF Lags (K)</i>					
	5200				20				260					
													Best Model	
	SD	Mean	Median	Kurtosis²	Minimum	Maximum	AR(p)	0: k¹						
AICS	16.51	361.07	360.84	-0.36	311.81	417.84	1	260*						
MAPES	3.54	17.46	17.75	2.88	0.65	44.03	12	242						
RMSES	3.37	16.7	16.93	3.06	0.71	42.67	12	242						
MAES	3.23	15.96	16.22	2.91	0.57	40.32	12	242						
DAILY														
<i>#Regressions & Forecasts</i>					<i>Maximum Possible P</i>				<i>Maximum Possible HF Lags (K)</i>					
	5200				20				1300					
													Best Model	
	SD	Mean	Median	Kurtosis²	Minimum	Maximum	AR(p)	0: k¹						
AICS	16.14	362.4	361.46	-0.513	316.42	412.71	1	1255						
MAPES	3.38	18.51	18.46	2.164	2.5	36.796	8	1185						
RMSES	3.32	17.64	17.55	2.118	2.84	35.83	19	460						
MAES	3.11	16.88	16.84	2.142	2.17	33.71	8	1185						

1: 0:k is defined as zero to k high frequency lags in units of respective independent variable's sampling frequency

2: Given values are excess kurtosis

*: Maximum lags used