The Role of Theory-Motivated Fundamentals in Long-Run Exchange Rate Forecasting

Jared Berry

6/19/2017

Advisor(s): Dr. Jaime Marquez (Johns Hopkins SAIS)

The Research Paper is submitted
in partial fulfillment of the requirements for the
Master of Arts degree in
International Economics and Finance
The Johns Hopkins University
School for Advanced International Studies (SAIS)
Washington, D.C.

1. Introduction

Exchange rate forecasting is one of the most contentious and widely disputed topics in the modern economics literature. Following largely from the seminal work of Meese and Rogoff (1983), analysis has found, time and time again, that the random walk or random walk with drift model regularly outperforms "structural models" which incorporate actual realized observations into the development of forecasts. By construction, however, the forecasts based on the random walk implicitly assume that the only information that can be used for a forecast at any future time horizon is the realization of the variable of interest today. Accordingly, shocks accumulate at random to construct a series that serves as the counterpart random walk series for forecast evaluation.

The predominant argument that the random walk outperforms models with structural components, while convincing at shorter time-horizons, becomes increasingly suspect as that time-horizon increases beyond the initial few months. For instance, when attempting to construct forecasts at horizons of greater than one year, or with displacements of many years, it seems ill-advised to rely only on the current realization of the exchange rate—for instance, in the valuation of a multinational corporate project involving foreign cashflows. The inherent impracticality of this approach arises largely because it ignores the possible role broader theoretical or structural elements play in exchange rate determination when looking at a longer time frame.

In this spirit, there are a number of well-established structural factors that, in theory, should play a crucial role in exchange rate determination, such as money supply, economic growth, interest rates, and inflation. Moreover, it is sensible to assume that these effects would be especially pronounced over longer-time horizons (particularly those related to prices and growth), as it may take more time for them to be reflected in exchange rates which fluctuate on a moment-by-moment basis. A theory based model might incorporate these determinants, but, as Meese and Rogoff deduce, theory based structural models perform poorly against the random walk in the short-run in terms of estimate error. As such, I posit that the reason these theory-based structural models perform so poorly relative to the random walk, particularly in the context of Meese and Rogoff, is largely a function of the time horizon.

I seek to address these tensions in this analysis by employing a well-established hybrid model for exchange rate forecasting from Mark (1995) that incorporates both structural elements and time-series elements (i.e. realized exchange rate data). The motivation for the specification of the Mark model is such that exchange rates are anchored to some fundamental value and, by extension, deviations from a fundamental value ought to be self-correcting. Accordingly, a sufficiently long-run time horizon is necessary to observe this self-correction to the core, fundamental value. By employing this unique specification, I capture not only the importance of previous period's exchange rate (per the methodology that underlies the random walk), but the impact of the structural components indicated above that are likely to have an impact on these rates in the long-run. Further, this approach allows for a parsing out of the relative contributions of both the random walk and monetary factors over time, and is naturally befitting of a graphical analysis that maps these contributions over a given time-horizon.

Given the importance of the time-horizon in this context, I employ two different frequencies of data in this analysis from two different sources. In the spirit of the Mark (1995) analysis, and more directly in accordance with the work done in Meese Rogoff (1983) and other work that relies on the use of similarly specified hybrid models, I apply the model to monthly exchange rate and price data with direct, k-step ahead forecasts of exchange rates from 1 to 60 months (covering a span of 1 to 5 years) for four currencies: the British Pound Sterling, Canadian Dollar, Japanese Yen, and Swiss Franc. While this serves to bolster the external validity of this analysis and frame it within the existing literature with more current data, this analysis is novel in its use of truly long-term data at truly long-run time-horizons.

To my knowledge, little has been done in the existing literature to apply this framework to horizons beyond 16 quarters (the scope of the Mark (1995) analysis). Intuition would suggest that the value of the fundamental is most pronounced at longer-horizons, and, as such, seek to contribute to the existing literature by applying the model to data from the Bank of England's "Three Centuries Macroeconomic Dataset" (2016). The dataset, which contains (as the title suggests) almost three centuries of observations of bilateral exchange rates, price data, etc., provides a unique opportunity to not only assess the degree to which structural components can motivate existing exchange rate forecasting frameworks in long-run time-horizons, but to assess the degree to which the structural component is invariant to institutional considerations.

The Bank of England dataset spans approximately two centuries of observations of bilateral dollar-pound exchange rates and price data for use in this analysis, which offers a rich opportunity for an analysis of historical periods (chiefly the Gold Standard, Bretton Woods, and current-float periods) and the extent these periods have bearing on the fundamental. So, the scope of the analysis moves beyond exchange rate forecasting in the sense it has been applied to post-Bretton Woods quarterly or monthly data and exploits available data to provide a richer analysis of theory versus the role of institutions and a broader analysis of historical periods in this context.

As stated, in this analysis, I seek to determine the economic role of theory versus data (i.e. the random walk) at medium- and long-run exchange rate forecasting horizons and the degree to which incorporating a structural, theory-motivated component can improve upon the random walk. I assess the robustness of these results by conducting the analysis with two different frequencies of exchange rate and price data (annual and monthly) and with additional currencies in accordance with the existing work of Mark (1995). Second, regarding the exchange rate data at the annual frequency from the Bank of England database, I seek to determine the degree to which the fundamental rate vis-à-vis theory is sensitive to institutional considerations. Third, after determining the additional forecasting value this model can provide to the random walk through a number of evaluation metrics, I seek to determine to what extent the forecasts developed in the hybrid model can improve on or be improved by forward exchange rates, thereby connecting the findings of the results in this analysis to actionable behavior in the forward markets.

Section two will provide a brief review of the literature surrounding exchange rate forecasting, particularly related to the work of the hybrid model used in this analysis and other work related more tangibly to medium- and long-term exchange rate forecasting that employs the use of structural variables. Section three provides a detailed description of the data utilized in the long-term forecasting components of the analysis, related to the Bank of England data that spans two-centuries of annual observations. Section four provides a discussion of the time-series properties of the series I utilize in this analysis. Specifically, I am concerned with the stationarity of the processes and the cointegrating relationship between both the exchange rate and structural component which serves as a theoretical counterpart for the role of an underlying fundamental value for exchange rate movements. Section five recounts the methodology of the Mark (1995) model and how it is adapted for use in these data and this analysis, section six will include results and subsequent discussion, and section 7 will conclude, addressing limitations of the current analysis and offer suggested next-steps for future research on the subject.

2. Literature Review

The work of Meese and Rogoff (1983) and Mark (1995) serve as the predominant theoretical foundations upon which this analysis expounds. Meese Rogoff (1983) serves as the predominant work on exchange rate forecasting that suggests a random walk outperforms several structural models of forecast error in the 1-12-month time-horizon. Using mean error, mean absolute error and root mean square error in 1, 3,6 and 12 step-ahead forecasts, the random walk regularly outperforms structural models in terms of forecasting error in *all* time horizons. The authors declare that they, “find that a random walk model would have predicted major-country exchange rates during the recent floating-rate period as well as any of our candidate models… the structural models fail to improve on the random walk in spite of the fact that we base their forecasts on actual realized values of future explanatory variables” (3). Further, the results suggest that one possible explanation for the poor out-of-sample fit are underlying shocks that cannot be adequately forecasted using data that does not take these into account.

Per the discussion above, these results seem suspect beyond the 12-month time horizon—to believe that the random walk will continue to regularly outperform the theory-based models is to say that the only information of value of long-horizon forecasting is the value of the exchange rate today and some fundamental understanding of the shocks that accumulate in the construction of the series. Further, on potential weakness of using the theory-based models in isolation could be just that—in isolation, the models are not accounting for the more potentially volatile shocks to exchange rates. A model that accounts for both elements, therefore, could be a useful forecasting tool.

This serves as the intuitive basis for the Mark (1995) analysis that serves as the empirical basis for the analysis at hand. This work approaches the question of exchange rate forecasting as an exercise in adjusting to deviations from a “fundamental value,” (i.e. in this case, the relative ration between money supply and GDP between domestic and foreign counterparts, but, more succinctly, purchasing power parity, a widely touted “structural” model variable) in the long-run. The author postulates that, if there is some fundamental value, deviations in the exchange rate from the fundamental ought to be self-correcting in the long-run, and, more importantly, finds that “long-horizon changes in the log nominal exchange rates contain an economically significant predictable component” (Mark, 1995, 201). This analysis, and the analysis of Mark, seek to exploit this economically significant predictable component with the use of structural components that could feasibly serve as “fundamentals” in the long-run.

As such, the model postulated by Mark serves as the empirical basis for this analysis—it establishes a “hybrid model” that incorporates both a long-run fundamental component, substantiated with the use of theory, and a time-series component that captures the shorter-term fluctuations in the exchange rate, i.e. a lagged exchange rate component which directly incorporates those shocks indicated in Meese Rogoff (1983). By including both, the model is not solely an autoregressive time-series model, nor is it a pure theory model. By including both, further, I hope to improve upon the use of the two types of models in isolation, recognizing the power of the random walk/univariate time-series regressions in the short- and medium-term. The use of the theory component, therefore, ought to be motivated by self-correcting behavior of exchange rates from a fundamental like purchasing power parity. I posit that the establishment of some cointegrating relationship between the two serves as adequate evidence to justify the use of a given structural variable as a corresponding fundamental for the exchange rate series in question. In accordance with the evaluation metrics of Meese Rogoff (1983), the author utilizes RMSE, but extends the analysis to incorporate the trajectory of R^2 values and constructs Diebold-Mariano tests to assess the predictive power between models.

The work of Lothian and Taylor (1991, 1996) serve as a further theoretical precedent for the use of the long-run data samples to motivate a long-run mean-reverting, fundamental relationship in exchange rates. Using bilateral data spanning two-centuries in much the same way they are utilized in this analysis, the authors find that there is "strong evidence of mean-reverting real exchange rate behavior" (Lothian and Taylor, 1996, 488), which, while not directly related to the use of fundamentals in the hybrid model context, lends credibility to the "self-correcting" (488) behavior observed in the Mark (1995) analysis regarding long-run exchange rate forecasting. Furthermore, the authors find that "simple, stationary, autoregressive models estimated on pre-float data, easily outperform nonstationary real exchange rate models" (488). The results from both the 1996 and 1991 analyses cover similar time-horizons for two different currencies, notably the pound sterling and Japanese Yen, both of which are utilized in this analysis.

Broadly, the work serves as a precedent for analyses seeking to improve upon a weakness in existing literature related to small-sample issues—long-term horizons provide a better measure of non-stationarity, the trade-off being the inclusion of wide-ranging institutional considerations. More tangibly, (in the context of this analysis), an extension of the 1996 work entitled "Two Hundred Years of Sterling Exchange Rates and the Current Float" (Lothian and Taylor, n.d.), finds "that real exchange rates revert to equilibrium values over the long run *and correspondingly that nominal exchange rates and relative price levels converge*" (1). This analysis seeks to directly expound upon this finding by Lothian and Taylor, using a similar dataset and the application of the Mark (1995) hybrid model which incorporates both elements. Further, the authors acknowledge the potential issues surrounding out-of-sample forecasting when incorporating dramatic shocks implicit in changing political and economic institutions, further motivating an approach that splits the sample into institutional segments and then conducts the analysis on those in addition to the broader 200-year sample.

The original Mark (1995) work defers to the use of the relative (domestic to foreign) ratio of money supply/GDP for the "structural" component of the hybrid model. While this is a direct corollary for the use of purchasing power parity (i.e., the relative price ratio), in this analysis, there is a wider literature substantiating the use of a number of structural variables in the context of this analysis, including purchasing power parity as I have defined it. Engel et al. (2012) extend the logic advanced by Mark (1995) by constructing "factors from a cross section of exchange rates and [using] the idiosyncratic deviations from the factors to forecast" (2), instead seeking to exploit the deviations from the fundamental rationale on a panel of bilateral USD exchange rates with several fundamentals, including Taylor rule factors, monetary models and purchasing power parity.

The authors use a similarly specified hybrid model and more broadly define the "fundamental" as a central tendency that can be composed of a number of structural values. In doing so, the authors again find that this approach improves forecasts in the long-horizon in a 1999-2007 time-series sample. Similarly, Molodtsova and Papell (2009) conduct forecast evaluations of short-term forecasts of exchange rates (bilateral USD) using structural variables from Taylor rule equations, and find that they improve on interest rate, PPP and monetary models. This work also serves as an extension of the Mark (1995) approach, using a similar specification, but again, limits the analysis to the post-Bretton Woods float time horizon. The work of Simpson and Grossmann (2011 and 2010), Marsh and MacDonald (1997) and Taylor (2002), all lend further credibility to not only the use of fundamentals in long-run exchange rate forecasting, but do so in the particular context of purchasing power parity, and find that, with particular specifications that rely on the use of purchasing power parity, such models can outperform the random walk even at time-horizons of only 1-month displacement.

One other possible explanation for the predominance of the random walk in existing exchange rate forecasting literature is the use of conventional evaluation metrics. Engel, Mark, and West (2007), offers alternatives to the existing criteria of outperforming the random walk as the benchmark for determining the forecasting power of an alternative model in the context of this particular framework, and a vast literature exists to support the use of profitability/utility measures (Leitch and Tanner (1991); Abhyankar, Sarno and Valente (2005); Boothe and Glassman (1987); West, Edison and Cho (1993); and Diebold (2006) to name a few). Perhaps more directly related to the analysis at hand, Christofferson and Diebold (1998) recognize that mean squared error estimates do not account for the improvement provided by cointegrating relationships as I will establish.

In this spirit, Burns and Moosa (2012) propose alternative measures that focus more directly on the direction of error, rather than the magnitude. Rather than focusing specifically on these alternative measures, I will defer to the existing metrics utilized in Meese and Rogoff (1983) and Mark (1995). Utilizing existing metrics is useful for framing the analysis in the context of the existing literature, while the reporting of beta coefficients for the structural component, along with R^2 values more directly account for the rule of theory in the context of the hybrid model beyond the evaluation metrics. As stated above, to more directly relate the results of this analysis in realistic markets, I apply the forecasts to forward markets conducting forecast encompassing regressions to see to what extent patterns in the forward market are useful for these forecasts and vice versa. This is in line with existing literature regarding forward rates and their ability to predict future spot rates from Simpson and Grossman (2014, 2015).

3. Data Description

To motivate the use of purchasing power parity as a structural variable for use in the hybrid model specification for exchange rate forecasting, plots of the data are helpful for illustrative purposes. This is particularly important in the context of the annual data from the Bank of England dataset, which contains annual bilateral dollar-sterling exchange rate and price data for both the UK and the US for over 150 years of observations, with a full sample of annual data points from 1861 through 2015. Recognizing that the data covers fundamentally different, economic institutions/climates, it is appropriate to attend briefly to the character of these periods and the impact they might have on the long-run exchange rate forecasting model to be specified below. For completeness, I first attend to plots of monthly bilateral exchange rates for the dollar against the pound sterling, Japanese yen, Swiss franc, and Canadian dollar.

Figure 1: Monthly Exchange Rate and PPP – USD/Canadian Dollar

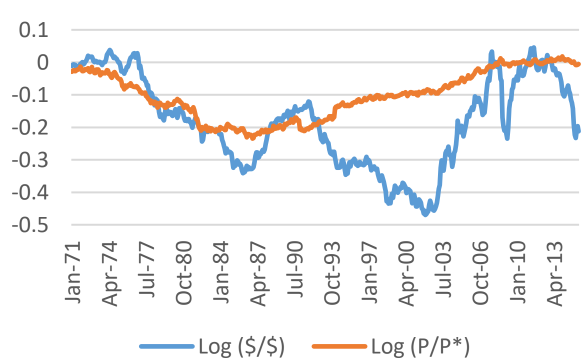

Figure 2: Monthly Exchange Rate and PPP – USD/Swiss Franc

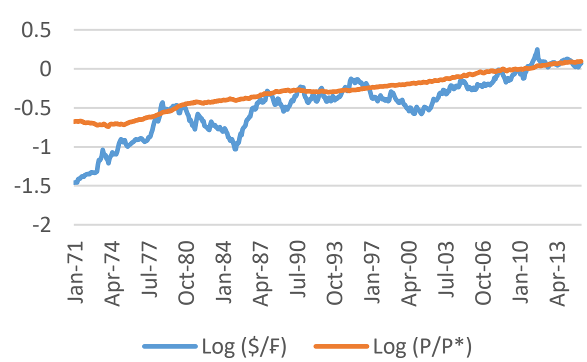

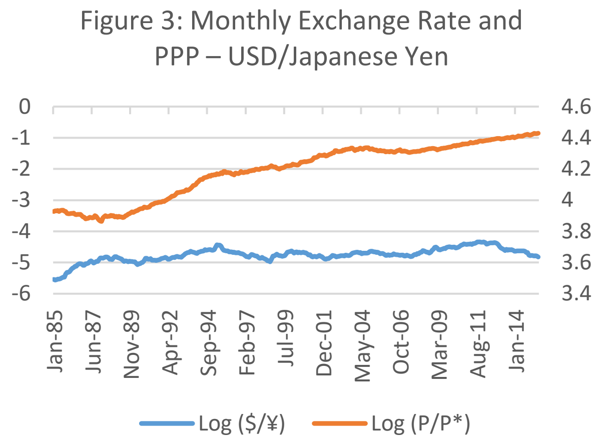

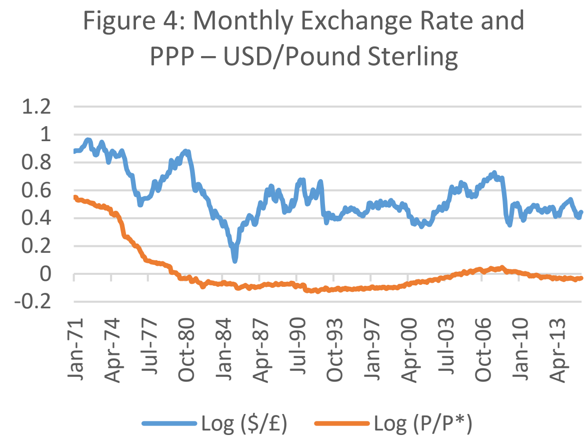

At first pass, the log purchasing power parity series tracks, more or less, the average trajectory of the bilateral exchange rates for both the Swiss franc and the Canadian dollar without any transformation. While the effect is not immediately clear in the case of the Japanese yen or pound sterling, I posit that the degree to which the two relative purchasing power parity series tracks the bilateral exchange rates would be remarkably similar to the previous two series with the simple addition of a constant. This is likely more a function of the characteristics of these particular exchange rates and their relation to the United States dollar over the course of the sample, while the Canadian dollar and Swiss franc have largely fluctuated above 1:1 value with the dollar over the course of the past three decades the same cannot be said of the pound sterling or yen. The sterling has been strong relative to the dollar throughout the sample, and the exchange rate series may be unconditionally higher than that of the purchasing power parity series as a result. In regression analysis, this should be dealt with accordingly with a mean effect vis-à-vis a constant. Shifting the purchasing power parity series up would more directly plot the trajectory of the exchange rate series. The same bias holds for the yen, particularly due to the units of the yen against the dollar. Since the relationship is never 1:1, a downward adjustment of the series would accurately track the average movement of the exchange rate series.

I posit that the trajectory of both series for all currencies, along with existing literature as discussed above, provides a firm justification for the use of purchasing power parity as a fundamental series which serves to plot an average effect of the exchange rate movements less the shocks that exchange rates are more subject to. Major differences in series seem to be the result of unconditional bias that is easily ameliorated with a constant, and does not represent a fundamental threat to the use of the series in this context. Cointegration, as discussed below in Sections 4 and 8, will further motivate the use of this structural variable and better contextualize the use of these plots.

Attending, now, to the annual data from the Bank of England dataset, I observe similar behavior in both plots throughout the sample for both series. The plot, for the log purchasing power parity and bilateral exchange rate series for both samples is as follows:

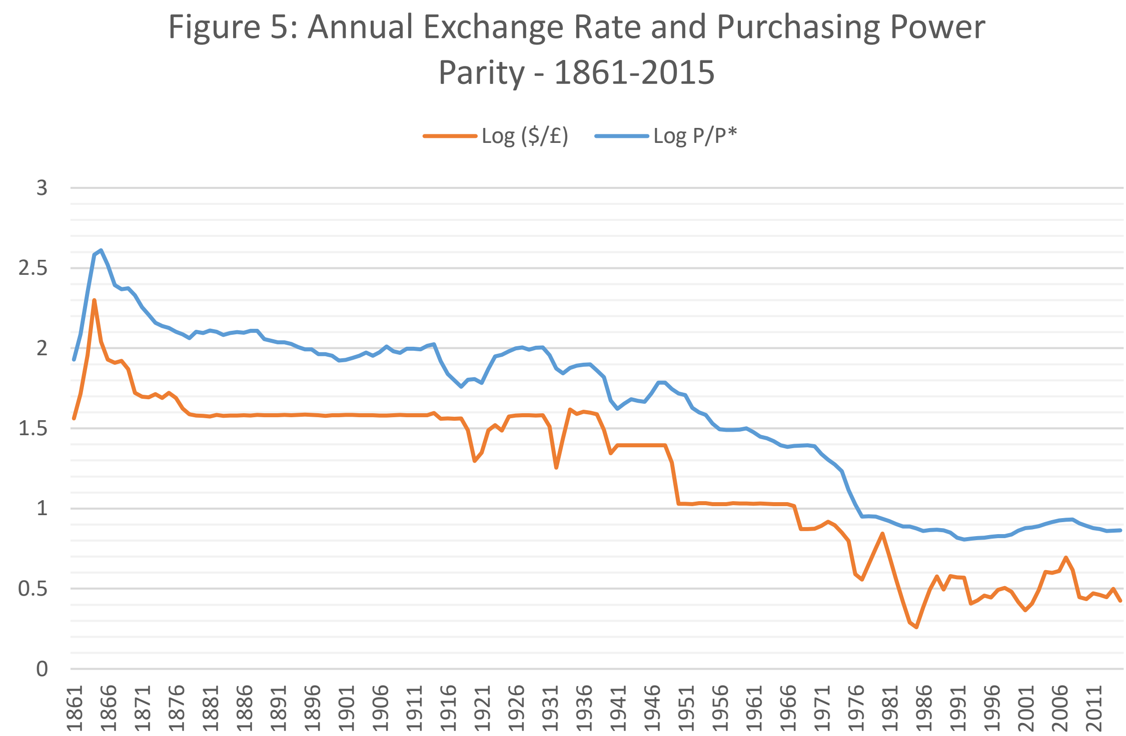

Despite the obvious impact of different institutional frameworks, the log purchasing power parity series does a surprisingly fair job of matching the relative trajectory of the log bilateral exchange rate series across all years in the 150-year sample. Again, we observe the same mean effect where one series is unconditionally above the other, but the addition of a constant would ameliorate the bias handily.

Regardless of degree to which log purchasing power parity matches the trajectory of log bilateral exchange rates throughout the sample, one of the key objectives of this analysis is to parse out the impact of institutional considerations on both exchange rate determination and the behavior of the fundamental. As such, I will subdivide the broader sample into three separate institutional frameworks: The Gold Standard, the Bretton Woods system, and the post-Bretton Woods Float. Each separate time-series represents broadly different economic climates and exchange rate regimes, which cannot be ignored when looking at a sample this broad in scope. To better empirically motivate the choice of these subperiods, I conduct Quandt-Andrews Structural Breakpoint tests[1] on the broader sample to determine whether there are empirically valid break dates in the broader sample that coincide with economic intuition.

[1] Results of Quandt-Andrews Structural Breakpoints tests are discussed in Section 7.

Figure 6: Annual Exchange Rate and Purchasing Power Parity - 1861-1939

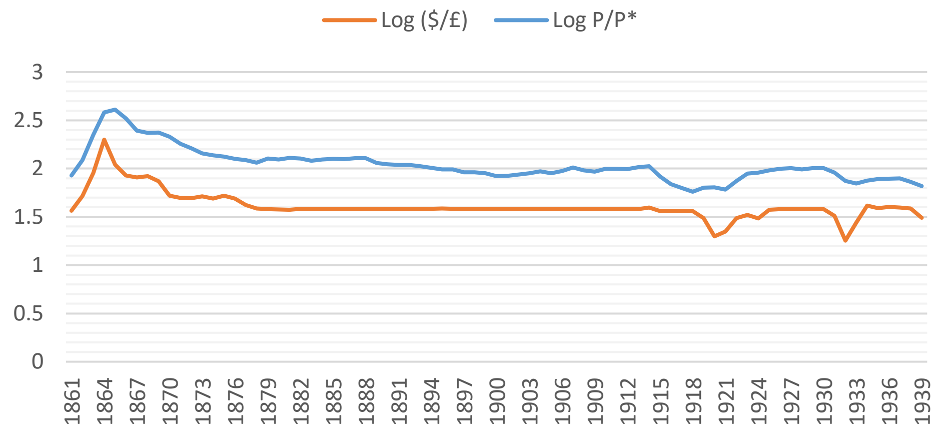

I consider the 1861-1939 period the “Gold Standard.” At a time where the United States is still a fledgling nation and the pound-sterling is king, we should expect vastly different behavior in both series than we would in, say, the post-Bretton Woods float. Notably, the inherent hegemony of the pound sterling in economic terms, and the degree to which currencies are backed by gold during this time frame, would imply that prices should be more closely linked with exchange rates and, furthermore, exchange rate movements should be less volatile. This is clearly borne out in the data- from 1870 until 1918, there is very little, if any, movement in the exchange rate and little movement in the purchasing power parity variable. Any instances of spikes in the exchange rate series (a jump at the outset of the sample and two small drops, seemingly, around World War I and the Great Depression) are closely matched with similar spikes in the price series. Given the degree to which movements in each series track one another, I would expect to find the strongest evidence of cointegration during this subperiod, relative to the other two, further motivating the use of the fundamental in this context.

Figure 7: Annual Exchange Rate and Purchasing Power Parity - 1940-1973

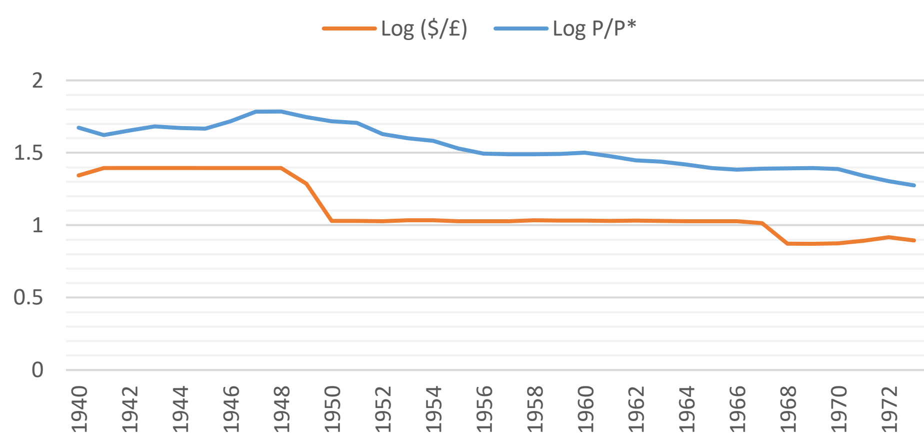

The Bretton Woods period (1944-1973 subperiod in this analysis), similarly features very little variable in either the exchange rate or prices—in fact, in stark contrast to the monthly series from above from recent years (these series largely correspond to post-Bretton Woods data), relative prices seem to exhibit somewhat more variability than the generally-volatile exchange rate counterparts. This should come as no surprise given the nature of exchange rate regimes under the Bretton Woods era. By construction exchange rates are essentially fixed in response to depression-era fluctuations, largely to stave off incipient capital flows in a fragile period for both representative economies. Thus, by design, there is little fluctuation in exchange rates much like that of the Gold Standard. However, prices do not exhibit the same degree of constancy relative to exchange rates they did during the Gold Standard. Since prices are not fixed as strictly, and neither series is backed commonly by Gold, the way in which the series behave ought to be fundamentally different than observed in the prior subperiod. Most dramatically, however, one would expect the post-Bretton Woods float to be altogether structurally different than either of these subperiods.

Figure 8: Annual Exchange Rate and Purchasing Power Parity - 1974-2015

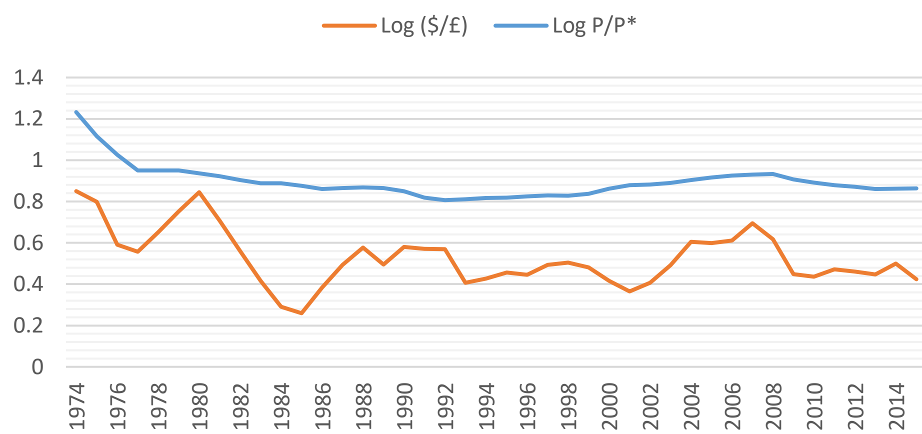

Following the breakdown of the Bretton-Woods system, we observe freely floating exchange rates through present day. While the pound sterling has remained strong against the dollar, the dollar serves as the international reserve currency and, therefore, is largely considered the predominant currency used throughout the world. These characteristics suggest the present day is, unsurprisingly, fundamentally different from the prior two periods in terms of exchange rate regimes and the relative importance of the two currencies in the global economy. Moreover, the current float incorporates several other unprecedented institutional characteristics not present in the prior two periods, including the rise of the European Exchange Rate Mechanism (ERM), the rise of the euro in its current manifestation, in the latter decades of the century, the advent of high-frequency trading, and the 2008 financial crisis.

One cannot ignore the possible impact these considerations might have on the fundamental exchange rate, and leads us to question whether the “fundamental” of the prior periods is different from the “fundamental” here. While there are inherent benefits to using such a long-run sample when seeking to better explain exchange rates with fundamentals, it is crucial to attend to these tensions that arise from using a dataset that spans inherently dissimilar structures. As such, I see this subdivision of the data as a necessary component of any analysis regarding the Bank of England dataset in this context.

4. Time-Series Properties

To fully understand the methodology established henceforth, it is crucial to understand the time-series properties of the series in question, beyond the simple plots provided in the previous section. Understanding whether bilateral exchange rates are stationary and, if non-stationary, whether they are cointegrated, is valuable for estimation purposes and interpretation of results. I begin by attending to tests of stationarity.

It is well-established in existing literature that exchange rates exhibit remarkably non-stationary behavior and are often described as random walk series, naturally giving rise to the predominance of the use of the random walk for forecasting purposes. Random walk series are remarkably persistent, and feature moments that are ill-defined. Random walks are, more generally, a special case of the broader collection of unit root series which exhibit the same characteristics. To better understand the underlying times-series properties of the data series in question, I conduct Augmented Dickey-Fuller (ADF) tests for all bilateral exchange rate time series. Results are as follows:

Figure 9: Augmented Dickey-Fuller Unit Root Tests & Regressions for Exchange Rate Series

Dependent Variable: Δ(Log Exchange Rate)

	Exchange Rate Series				
	Canadian Dollar	Swiss Franc	Japanese Yen	British Pound (Monthly)	British Pound (Annual)
ADF Test Statistic	-1.805	-2.275	-2.030	-2.791	-0.445
p-value	0.378	0.181	0.274	0.060	0.897
Log Exchange Rate (-1)	-0.008*	-0.007**	-0.005**	-0.016***	-0.006
	(0.004)	(0.003)	(0.003)	(0.006)	(0.013)
Δ(Log Exchange Rate (-1))	0.277***	0.27***	0.317***	0.355***	0.29***
	(0.042)	(0.042)	(0.041)	(0.04)	(0.079)
Δ(Log Exchange Rate (-2))					-0.246***
					(0.080)
Constant	-0.002*	-0.001	-0.025*	0.008**	-0.0003
	(0.001)	(0.002)	(0.013)	(0.003)	(0.017)
Observations	532	532	532	532	155
R^2	0.081	0.083	0.109	0.135	0.112

Note: Augmented Dickey-Fuller Tests are performed by regressing combinations of differenced and lagged exchange rate variables on the first difference of the respective exchange rate series. Both regression output and test statistic evaluation are included in the table above. Upon determining the appropriate t-statistic from the first differenced log exchange rate is tested against Dickey-Fuller distribution test critical values. As such, p-values from the ADF test statistic portion correspond to the appropriate critical values for the test. I test against the null hypothesis that each respective exchange rate series has a unit root. Standard errors in parentheses.

*** $p<0.01$, ** $p<0.05$, * $p<0.1$

As the results above imply, I fail to reject the null of a unit root for *all* series with the notable exception of the monthly log dollar/pound exchange rate for the United States and United Kingdom. This result is robust to change in frequency, as the trimmed annual sample (including only those observations used in estimation which correspond to the available price data) from the Bank of England dataset also implies there is no unit root in the bilateral exchange rate data annually. There are possible consequences for this regarding cointegration results and use of the hybrid model—I suspect that, given the apparent stationary behavior of this exchange rate data, a simple univariate autoregressive model may outperform the hybrid model. This will be tested below.

While results from the ADF test clearly indicate the underlying exchange rate series are unit root (or I(1)), this does not imply they are random walk series—random walks are, as indicated briefly above, a special case of unit root which feature more significant challenges for estimation and interpretation. Using Variance Ratio Tests for all log bilateral exchange rate series, I reject the null hypothesis of a random walk in each series at the 95% level for *all* series *except* the log dollar/pound exchange rate for the United States and United Kingdom, this time at the annual frequency.

Similarly, Augmented Dickey Fuller tests are performed on all log purchasing power parity series to determine whether they exhibit the same non-stationary characteristics:

Figure 10: Augmented Dickey Fuller Unit Root Tests & Corresponding Regressions for Purchasing Power Parity Series

Dependent Variable: Δ(Log Purchasing Power Parity)

	Purchasing Power Parity Foreign Price Series				
	Canadian Dollar	Swiss Franc	Japanese Yen	British Pound (Monthly)	British Pound (Annual)
ADF Test Statistic	-0.568	-1.067	-2.035	-3.941	-1.056
p-value	0.875	0.730	0.272	0.002	0.732
Log Purchasing Power Parity (-1)	-0.0017	-0.001	-0.002**	-0.007***	-0.006
	(0.003)	(0.001)	(0.001)	(0.002)	(0.006)
Δ(Log Purchasing Power Parity (-1))		0.061	0.124**	0.151***	0.609***
		(0.041)	(0.052)	(0.044)	(0.061)
Δ(Log Purchasing Power Parity (-2))		0.007	0.022	0.076*	
		(0.041)	(0.052)	(0.044)	
Δ(Log Purchasing Power Parity (-3))		-0.094**	-0.194***	-0.012	
		(0.041)	(0.052)	(0.039)	
Constant	0.008	0.004	0.011**	0.033***	0.006
	(0.014)	(0.003)	(0.011)	(0.009)	(0.010)
Observations	533	521	362	519	153
R^2	0.001	0.254	0.062	0.373	0.399

Note: See notes for the ADF test statistic procedure from the previous figure. Further, the results from the monthly British Pound data and Swiss Franc data are truncated as the regression estimated for determining the respective test statistic extends to 14 and 12 lagged differences, respectively. Standard errors in parentheses.

*** p<0.01, ** p<0.05, * p<0.1

In line with the results from the bilateral exchange rate data, I fail to reject the null of a unit root for all series except for the monthly United States and United Kingdom dollar/pound sterling exchange rate. Again, this is consistent with the ADF test conducted on the annual frequency data of the same exchange rate. Further, I reject the null hypothesis that the series are random walk at the 95% level for all log purchasing power parity series at the monthly frequency and at the 90% level for log purchasing power parity at the annual frequency.

Recognizing the degree to which all series in this analysis (again, apart from post-Bretton Woods USD/GBP exchange rates) are I(1), there are serious consequences for forecasting and interpretation of results. First, as discussed in previous sections, if a series is truly random walk, then forecasting beyond the information available today is inappropriate. Since a random walk series is highly persistent, the series evolves by the accumulation of random shocks in each subsequent period. As such, there is no mean reversion (the series stays constant at the current value) and variance is unbounded as time increases to infinity.

From a practical standpoint, however, as discussed, it seems illogical to assume that the only information that is of value for prediction in long-term time-horizons is the realization of the series today, so this serves more as a justification of the use of the random walk, and it's predominance as an exchange rate forecasting model, in the short term. Note, further, that while unit root series are also highly persistent, the possibility of a cointegrating relationship between two I(1) series could be helpful in motivating the inclusion of a fundamental.[2]

5. Methodology

Mark's fundamental exchange rate model serves as a natural departure from the work of Meese and Rogoff that incorporates both theory and the random walk in exchange rate forecasting. This "hybrid" model assumes a "fundamental" exchange rate, F, and further assumes that, if this is truly a fundamental exchange rate, deviations from the fundamental ought to be self-correcting in the long-run. In the context of this analysis, F, the "theory" component of the model, is assumed to be purchasing power parity (PPP) such that $F = \frac{P}{P^*}$, where P is the domestic price level and P* is the foreign price level, both measured in Consumer Price Index (CPI) terms. Taking logs, in accordance with the original approach of Mark, we have $f_t = (p_t - p_t^*)$. I defer to the log-form of these variables for the remainder of the analysis.

Based on these assumptions underlying this fundamental exchange rate, I seek to determine to what extent the gap between an observed exchange rate s_t and the fundamental f_t can be used to model *actual* changes in the exchange rate with k periods of displacement from time t. The regression is postulated as such:

$$s_{t+k} - s_t = \alpha_k + \beta_k(f_t - s_t) + \upsilon_{t+k}, \upsilon_{t+k} \sim N(0, \sigma^2)$$

Which can be further rewritten as:

$$s_{t+k} = \alpha_k + \beta_k f_t + (1 - \beta_k)s_t + \upsilon_{t+k}$$

This is the regression equation of interest in this analysis. As postulated, it implies that, when forecasting the exchange rate s_{t+k}, k periods from now, there is both a fundamental ("theory") component and random walk component which uses the past exchange rate s_t as the independent variable. Further, the regression is structured such that these components are weighted in their contribution per the coefficient β_k. This coefficient is the chief focus of this analysis.

If β_k is equal to zero, all weight is given to the past exchange rate and, as such, the random walk (or past data) is the appropriate model for forecasting this relationship. If, however, β_k is greater than zero, the "theory" component (i.e. purchasing power parity) meaningfully contributes to the predictive power of the model in modeling changes in exchange rates at across displacements of time. HAC standard

[2] See Section 8 for detailed analysis and discussion of cointegration in the context of vector error-correction models.

errors are reported in all regressions to account for the implicit serial dependence in forecasting regressions, and the impact of non-stationarity in series and residuals.

Accordingly, using this specification for both the monthly and annual series, forecast estimates are established by running k-step ahead forecast regressions—increasing lags serve as estimates of exchange rates made with increasing displacements of time, thereby making for longer-run estimations. To capture an appropriately long time-horizon, displacements of 1 through 12 will be used for the annual data, such that the longest time-horizon forecast is 12-years ahead. While this is truly a long-term horizon, it is not unrealistic—it would not be unreasonable to attempt to forecast exchange rates at such a horizon when conducting the valuation of a multinational venture with cash flows denominated in foreign exchange rates. Displacements of 1 through 60 will be used for the monthly data, such that the longest time-horizon forecast is 5-years ahead (the Mark (1995) analysis conducts analysis with the largest displacement at 64-months ahead). All data is at the monthly frequency from January, 1971 for the Canadian dollar, pound sterling, and Swiss franc. The sample begins in January, 1985 for the Japanese yen due to availability of price data.

As discussed briefly above, the predominant outcome coefficient from the regression as specified is β_k. A positive, statistically significant β_k coefficient in and of itself implies that the structural component log purchasing power parity contains information that can improve upon the predictive power of the forecast with past exchange rate data alone. While I view the interpretation of this coefficient as a sufficient measure of improvement on the random walk by including it as a special case, additional evaluation metrics are employed in accordance with existing literature. Per the methodology of Mark (1995), I report and plot R^2 values for each k-step ahead regression, which serve as indicators of the explanatory power of each k-step ahead regression with respect to the actual data, as well as constructed RMSE values to capture the relative magnitude of the forecasting error for each regression at each forecast horizon.

To further establish the predictive power of this hybrid specification, I conduct additional forecast evaluation with respect to a univariate autoregressive time-series model of the log exchange rate. This approach is utilized in Meese and Rogoff (1983) but must not be confused with a pure random walk—in the case of the random walk, the only information used for the forecast of any k-step ahead realization of the exchange rate is the information available today. The univariate unit root autoregressive model serves to model the predictive power of the lagged log exchange rate *without* the contribution of theory. This merely serves to provide an additional benchmark for the measurement of forecast quality between the hybrid model and other, more parsimonious specifications that implicitly assume theory is unhelpful in improving forecast accuracy. Regressions will be of the specification:

$$s_{t+k} = \alpha_k + \delta_k s_t + v_{t+k}$$

Which is essentially the same regression specified for the hybrid model where the original coefficient β_k is simply assumed to be equal to 0 at all forecast horizons and is therefore excluded from the forecasts. Further, in accordance with the existing literature, Diebold-Mariano tests are performed to assess the relative accuracy of the hybrid forecast model with respect to the parsimonious univariate time-series model. Difference series of squared errors for each k-step ahead forecast are constructed and, accordingly, run on a constant—as constructed, a negative, statistically significant coefficient serves as evidence of statistically valid improvement in accuracy for the hybrid model relative to the univariate time-series model.

Expanding the analysis beyond spot bilateral exchange rate data and the relative price ratio, I will conduct forecast encompassing regressions incorporating historic exchange rates from the forward market. Since the approach taken in this analysis involves conducting direct forecasts of k-displacements (i.e., a forecast for 3-months ahead is estimated directly, rather than rolling an autoregressive model forward as in the case of multistep ahead forecasting), each k-month ahead forecast will have a corresponding rate from the forward market that originated at the start date, and would allow a holder to transact at the forward rate k-periods ahead. Given information available at time t, when conducting a k-step ahead forecast, I hope to determine to what extent forward rates, which determine a corresponding k-period ahead price for a given currency, can contribute to those k-step ahead forecasts from the hybrid model specified above.

To empirically assess this contribution, I construct series of errors based on the difference between the log monthly average exchange rate at period k and the log monthly average forward rate for period k with information at time t for a given currency. I then conduct forecast encompassing regressions to determine to what extent forecast errors implicit from the difference in the forward rate and actual rate have predictive power relative to the forecast errors that result from the k-step ahead hybrid model forecasts. The broad specification of the regression run, in this framework is:

$$e_{i,t+k} = \alpha + \beta_i \eta_{i,t+k} + \xi_{i,t+k}$$

Where $e_{i,t+k}$ is the hybrid model forecast error series from the k-period ahead forecast for a given currency i , and $\eta_{i,t+k}$ is the error series constructed from the differences between log spot exchange rates at period k and the forward exchange rates established with information at time t to be transacted k-periods ahead.

6. Monthly Frequency Exchange Rate Forecasting Results

Attending first to the results at the monthly frequency most closely in line with the approach of Mark (1995), I will proceed systematically through the regression output and coefficient estimates for each currency. Broadly, across all four currencies, I find that the role of the fundamental vis-à-vis log purchasing power parity becomes increasingly economically and statistically significant as the forecast horizon increases. β_k-values that are both positive and statistically significant indicate statistically and economically valid contributions to the predictive power above and beyond what a simple univariate autoregressive model or random walk model that contain only past realizations of the exchange rate for prediction provide. Further, increasing contribution of the fundamental demonstrated in the β_k coefficients signify that the content of the forecasts of future exchange rates is being driven increasingly by the information from the fundamental alone. This will be further substantiated throughout the remainder of the section.

Figure 11: Hybrid Model Forecasting Regression Output (1- to 60-month Ahead Forecast Horizons) – Canadian Dollar/USD Exchange Rate Forecasts

Dependent Variable: Log of Dollar-Canadian Dollar Exchange Rate

	k-months ahead											
	1	3	4	6	9	12	16	24	32	36	48	60
Log(P/P*)	0.01	0.039	0.052	0.084	0.141	0.207**	0.285**	0.439***	0.596***	0.666***	0.824***	0.945***
	(0.014)	(0.038)	(0.049)	(0.067)	(0.087)	(0.103)	(0.119)	(0.148)	(0.164)	(0.168)	(0.174)	(0.171)
Log($/$)	0.991***	0.964***	0.95***	0.918***	0.873***	0.825***	0.767***	0.648***	0.522***	0.461***	0.27***	0.094
	(0.006)	(0.017)	(0.022)	(0.032)	(0.043)	(0.052)	(0.06)	(0.075)	(0.085)	(0.09)	(0.096)	(0.094)
Constant	-0.048	-0.182	-0.245	-0.398	-0.66*	-0.967**	-1.331**	-2.048***	-2.775***	-3.104***	-3.851***	-4.433***
	(0.061)	(0.172)	(0.22)	(0.304)	(0.394)	(0.464)	(0.54)	(0.673)	(0.745)	(0.765)	(0.788)	(0.775)
Observations	533	531	530	528	525	522	518	510	502	498	486	474
R^2	0.990	0.958	0.940	0.902	0.854	0.810	0.757	0.657	0.556	0.512	0.382	0.300
RMSE	0.014	0.028	0.034	0.043	0.053	0.06	0.068	0.079	0.09	0.094	0.103	0.108

HAC robust standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

Figure 12: Univariate Autoregressive Forecasting Model & Diebold-Mariano Tests (1- to 60-Month Ahead Forecast Horizons) – Canadian Dollar/USD

Dependent Variable: Log of Dollar-Canadian Dollar Exchange Rate

	k-months ahead											
	1	3	4	6	9	12	16	24	32	36	48	60
DM Constant	0.000	0.000	0.000	0.000	-0.0001	-0.0002	-0.0003	-0.0008	-0.0014*	-0.0018*	-0.003**	-0.004***
	(0.000)	(0.000)	(0.000)	(0.000)	(0.0001)	(0.0002)	(0.0003)	(0.0005)	(0.0008)	(0.0009)	(0.0012)	(0.0013)
Log($/$)	0.993***	0.974***	0.964***	0.941***	0.911***	0.88***	0.842***	0.76***	0.667***	0.621***	0.454***	0.288***
	(0.005)	(0.015)	(0.019)	(0.028)	(0.037)	(0.044)	(0.05)	(0.061)	(0.073)	(0.079)	(0.092)	(0.101)
Constant	-0.002	-0.006	-0.008*	-0.013*	-0.019**	-0.025**	-0.033***	-0.05***	-0.069***	-0.078***	-0.113***	-0.149***
	(0.001)	(0.004)	(0.005)	(0.007)	(0.009)	(0.01)	(0.011)	(0.012)	(0.015)	(0.016)	(0.018)	(0.02)
Observations	533	531	530	528	525	522	518	510	502	498	486	474
R^2	0.99	0.957	0.94	0.9	0.85	0.8	0.74	0.616	0.478	0.414	0.226	0.091
RMSE	0.014	0.029	0.034	0.044	0.053	0.061	0.07	0.084	0.097	0.103	0.115	0.123

HAC robust standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

Plotting the evolution of both the theory coefficients of interest (β_k) and the foil coefficients (that which determines the contribution of the lagged exchange rate, or the role of the past in determination of future rates) from the regression output above, the following pattern is observed:

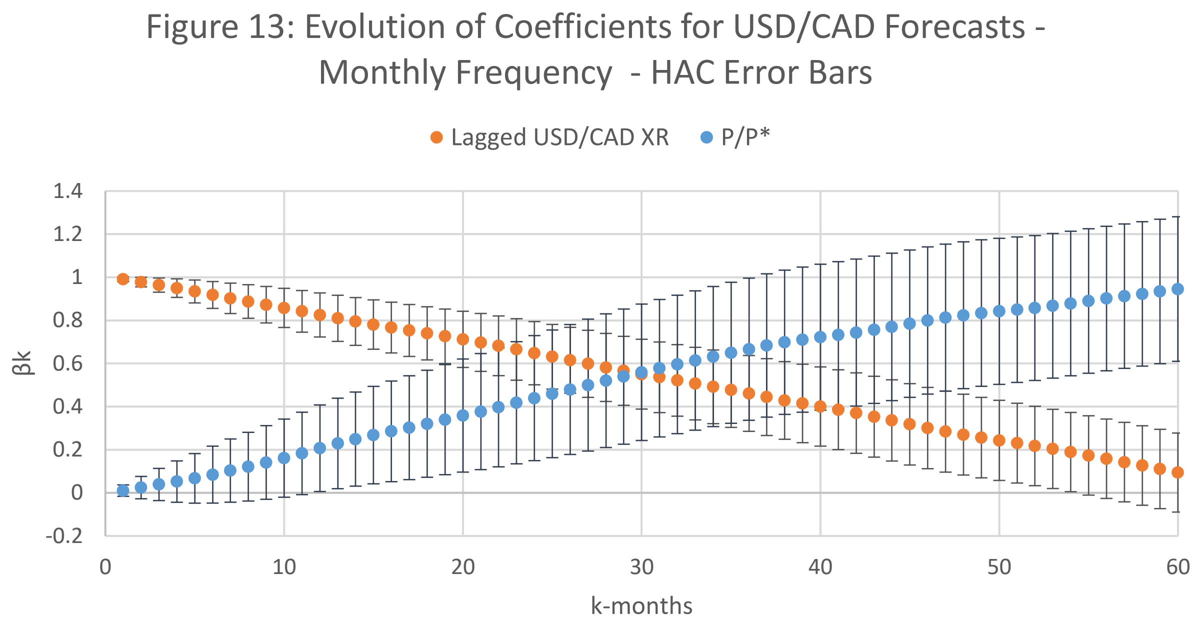

The relative contribution of both past realizations of exchange rates (the contribution that would be the sole focus of a random walk forecast) is an unequivocally downward, almost linear trend, with the exact opposite observed in the contribution of theory per the relative price ratio. Notably, the coefficients behave almost as mirror images of one another—they almost exactly fit the construction of the specification wherein the coefficient on the lagged spot rate is one minus the β_k. This is indicative of an almost unequivocal improvement upon the random walk by simply including the fundamental value exhibited by purchasing power parity. Beyond the economic significant of these plots, note that both coefficients are statistically significant in forecasts of 10 to 56 months.

Purchasing power parity does not contribute meaningfully until the 10-month ahead forecast, but becomes (and remains) significant at the 99% level from the 20-month ahead forecast and beyond. The role of the past is no longer statistically significant at the horizons beyond 56-months. This points further, to the importance of theory's relative predictive power in the *long-run*. The contribution is minimal at the outset, which is in line with the results of existing literature that points to the dominance of the random walk.

In terms of forecast performance, I defer to plots of R^2 (as a measure of explanatory power) and Root Mean Squared Error (as a measure of model fit) to graphically depict the evolution of model fit as the forecast horizon increases from 1- to 60-months ahead.

Figure 14: Model Fit for USD/CAD Forecasts - R^2 and Root Mean Squared Error

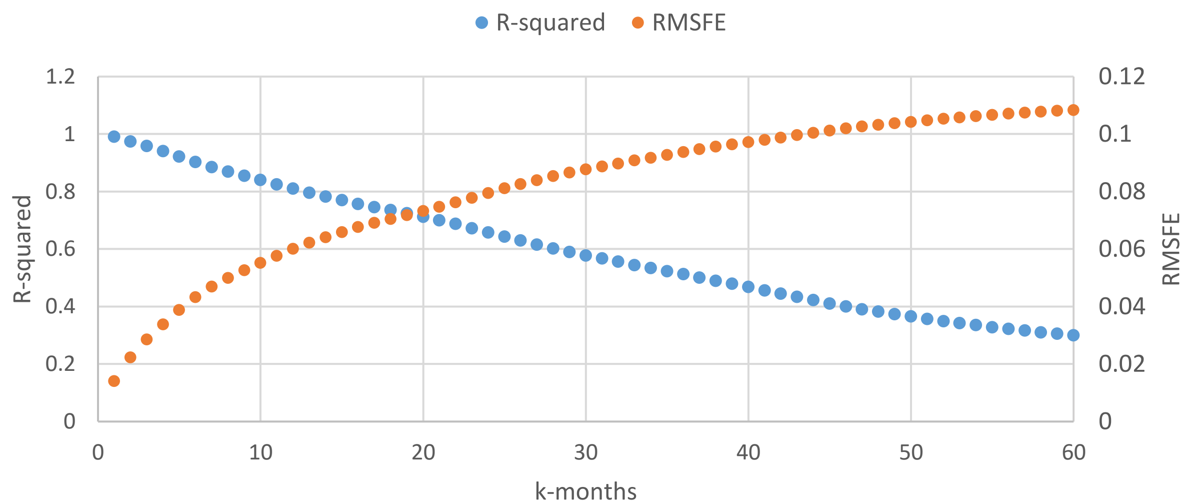

Despite the promising results exhibited in the coefficient evolution plot, RMSE increases as the forecast horizon increases, while the R^2 of the regression simultaneously decreases. The decline in R^2 can be described with a gently sloping, almost linear, curve beginning remarkably close to 1 and ending at approximately 0.30 by the 60-month ahead forecast horizon. RMSE, conversely, increases rapidly at the outset, almost logarithmically, and then growth slows. Beginning at a value of approximately 0.014, RMSE appears to plateau at approximately 0.108.

I do not see this necessarily as a repudiation of the significance of the hybrid forecasts, by likely more a product of the dramatic displacement in the data itself—as I will address more significantly in the annual results, forecasts at this horizon may be attempting to explain realizations under widely different institutional or economic climates that cause the accuracy of the forecasts to decline over time. Further, while this specification and data is not an exact replication of that used in Mark (1995), I maintain R^2 levels higher than those realized in the Canadian dollar regressions from that analysis. While these results serve to provide sufficient evidence of the contribution of theory at the monthly frequency with USD/CAD exchange rates, I assess the performance of the hybrid model relative to the univariate autoregressive exchange rate model specified in the methodology section. I find that, at all forecast horizons, the hybrid theory model outperforms the parsimonious univariate time-series model in terms of RMSE and R^2. Diebold-Mariano tests performed at each forecast horizon further imply a statistically valid difference in terms of model accuracy from the 30-month ahead forecast horizon and beyond favoring the hybrid model.[3]

[3] Refer to Figure 12 above for results of Diebold Mariano tests for USD/CAD exchange rate forecasts, represented with estimates in the DM constant row—negative, statistically significant coefficients signify better forecasts in the hybrid model, in terms of forecast accuracy measured via mean squared error, than the univariate regression.

Figure 15: Hybrid Model Forecasting Regression Output (1- to 60-month Ahead Forecast Horizons) – Swiss Franc/USD Exchange Rate Forecasts

Dependent Variable: Log of Dollar-Franc Exchange Rate

	k-months ahead											
	1	3	4	6	9	12	16	24	32	36	48	60
Log(P/P*)	0.016	0.076*	0.107**	0.165**	0.244**	0.34***	0.464***	0.723***	0.905***	0.963***	1.169***	1.443***
	(0.015)	(0.041)	(0.053)	(0.074)	(0.095)	(0.108)	(0.123)	(0.137)	(0.147)	(0.147)	(0.154)	(0.174)
Log($/₣)	0.983***	0.93***	0.902***	0.851***	0.78***	0.699***	0.594***	0.377***	0.222**	0.168**	-0.012	-0.226**
	(0.009)	(0.026)	(0.034)	(0.047)	(0.061)	(0.07)	(0.079)	(0.082)	(0.088)	(0.085)	(0.083)	(0.112)
Constant	-0.076	-0.354*	-0.497**	-0.764**	-1.132***	-1.573***	-2.147***	-3.344***	-4.181***	-4.447***	-5.393***	-6.657***
	(0.069)	(0.19)	(0.244)	(0.34)	(0.435)	(0.498)	(0.566)	(0.627)	(0.674)	(0.673)	(0.702)	(0.799)
Observations	533	531	530	528	525	522	518	510	502	498	486	474
R^2	0.994	0.976	0.966	0.949	0.924	0.897	0.865	0.81	0.769	0.756	0.729	0.739
RMSE	0.028	0.057	0.067	0.082	0.097	0.111	0.125	0.141	0.15	0.151	0.151	0.143

HAC robust standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

Figure 16: Univariate Autoregressive Forecasting Model & Diebold-Mariano Tests (1- to 60-Month Ahead Forecast Horizons) – Swiss Franc/USD

Dependent Variable: Log of Dollar-Franc Exchange Rate

	k-months ahead											
	1	3	4	6	9	12	16	24	32	36	48	60
DM Constant	0.000	-0.0001	-0.0001	-0.0003	-0.0006	-0.0011	-0.0021*	-0.005**	-0.008***	-0.009***	-0.014***	-0.022***
	(0.000)	(0.0001)	(0.0001)	(0.0002)	(0.0005)	(0.0007)	(0.0012)	(0.0021)	(0.0031)	(0.0033)	(0.0043)	(0.0063)
Log($/₣)	0.992***	0.973***	0.963***	0.945***	0.918***	0.89***	0.855***	0.779***	0.721***	0.697***	0.627***	0.559***
	(0.004)	(0.01)	(0.013)	(0.018)	(0.023)	(0.027)	(0.031)	(0.037)	(0.042)	(0.042)	(0.045)	(0.053)
Constant	-0.001	-0.004	-0.006	-0.009	-0.014	-0.019	-0.025	-0.04*	-0.05**	-0.054**	-0.063*	-0.076*
	(0.002)	(0.006)	(0.008)	(0.01)	(0.013)	(0.015)	(0.018)	(0.022)	(0.026)	(0.027)	(0.034)	(0.042)
Observations	533	531	530	528	525	522	518	510	502	498	486	474
R^2	0.994	0.975	0.966	0.946	0.919	0.888	0.847	0.759	0.683	0.655	0.56	0.456
RMSE	0.028	0.057	0.068	0.083	0.1	0.116	0.133	0.158	0.175	0.18	0.192	0.206

HAC robust standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

While I am unable to replicate the exact behavior of the USD/CAD exchange rate forecasts (an unsurprising result considering innate, fundamental differences between exchange rate series implicit in the relative relationship between the United States and the counterpart currency), broadly similar trends are observed for coefficients and model fit across currencies at the monthly level. The evolution of the coefficients for the USD/CHF forecasts are as follows:

Figure 17: Evolution of Coefficients for USD/CHF Forecasts - Monthly Frequency - HAC Error Bars

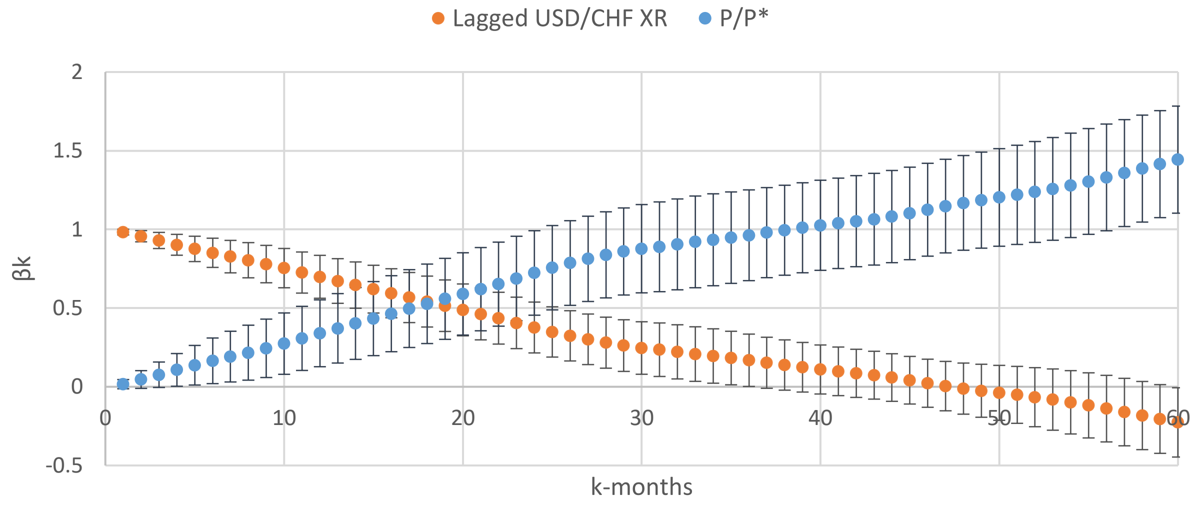

Again, the contribution of the fundamental purchasing power parity component is near zero in the initial forecast and increases linearly as the forecast horizon increases to a maximum value of 1.443 at the 60-month ahead exchange rate forecast. While I expect, based on the specification detailed in the methodology section above, the coefficient β_k to remain below one, it is not mechanically bound to behave in this fashion. The coefficient is statistically significant at the 99% level, again, from the 10-month ahead forecast and beyond, and stays remarkably significant when it is above 1.0. This implies an outsize contribution of the fundamental purchasing power parity relationship at longer-term horizons, perhaps to offset the negative impact of the exchange rates observed in the same time-frame. The estimates of β_k for the Swiss Franc in Mark (1995) also tend above 1.0 at the 16-quarter ahead forecast, so this result is not particularly surprising relative to existing literature.

Correspondingly, the contribution of the lagged exchange rate is near one at the initial one-month ahead forecast and decreases linearly as displacements increase, becoming negative at forecast horizons including and beyond four years ahead. Since these coefficients are not statistically significant, they are contributing little to no economic or statistical information to the prediction of the forecasts beyond the 38-month ahead horizon, so the negative coefficients are not troubling in this context. Regardless, the results of the Swiss franc forecasts via the hybrid model suggest that theory improves upon the random walk in terms of the β_k coefficient, and this improvement is magnified as the time-horizon of the forecasts increases.

Results for forecast evaluation and model fit are more comforting in the context of the monthly USD/CHF forecasts relative to the other currencies examined at the monthly frequency.

Figure 18: Model Fit for USD/CHF Forecasts - R^2 and Root Mean Squared Error

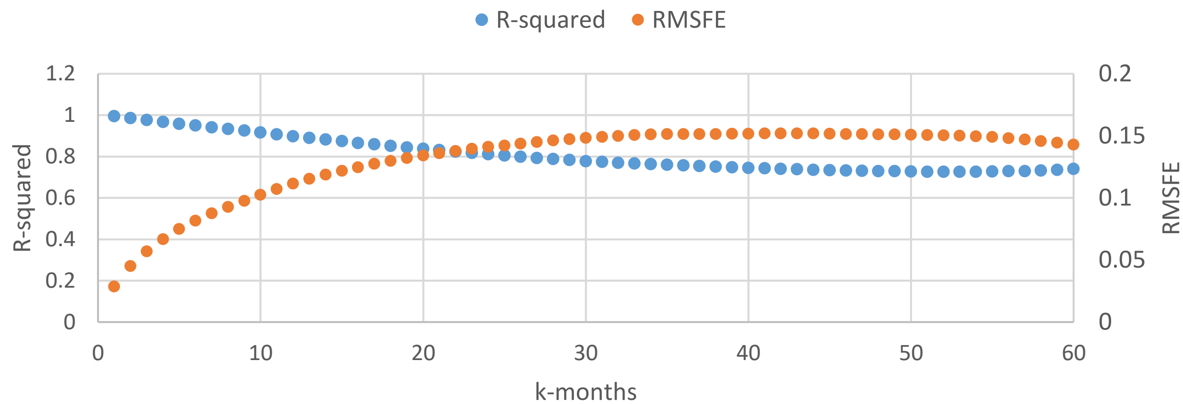

As observed in the other currencies, R^2 indicators of explanatory power decrease modestly as the forecast horizon increases, but the decline is *much* less rapid than observed in all other currencies resulting in the highest relative R^2-value of 0.739 at the 60-month ahead forecast horizon. Further, this is the only currency in this analysis that engenders an *increase* in the R^2 value at the tail-end of the forecast horizon, dropping to its lowest around the 52- and 53-month ahead forecasts and then recovering some of the explanatory power lost. This implies that the hybrid theory model can more reliably explain variation in the USD/CHF exchange rate at longer-horizons than the other currencies in question.

In terms of RMSE, the consistent pattern of marked initial increases, a near-logarithmic trend of evolution across forecast horizons, and a plateau in those increases is observed. Incidentally, while the USD/CHF forecasts with the hybrid model feature the most impressive metrics in terms of explanatory power, the model is subject to the largest incidences of RMSE across the four currencies examined at the monthly frequency, reaching 0.15184 at the 42-month ahead forecast horizon. Despite this, the USD/CHF forecast series is the only series of those sampled to exhibit a decline in the RMSE figures at the longest forecast horizons, complementing the trajectory of the R^2 values described above. Relative to the univariate autoregressive exchange rate model, the USD/CHF exchange rate forecasts generated from the hybrid theory model exhibit higher explanatory power and lower RMSE values at all forecast horizons, with statistically valid improvements in model accuracy vis-à-vis Diebold Mariano tests beginning at the forecasts beyond 1-year ahead.[4]

[4] Refer to Figure 15 above for results of Diebold Mariano tests for USD/CHF exchange rate forecasts.

Figure 19: Hybrid Model Forecasting Regression Output (1- to 60-month Ahead Forecast Horizons) – Japanese Yen/USD Exchange Rate Forecasts

Dependent Variable: Log of Dollar-Yen Exchange Rate

	k-months ahead											
	1	3	4	6	9	12	16	24	32	36	48	60
Log(P/P*)	0.002	0.027	0.043	0.071	0.102	0.161*	0.262***	0.439***	0.6***	0.652***	0.616***	0.526***
	(0.012)	(0.034)	(0.043)	(0.058)	(0.073)	(0.085)	(0.1)	(0.11)	(0.094)	(0.091)	(0.1)	(0.118)
Log($/¥)	0.973***	0.889***	0.844***	0.76***	0.654***	0.531***	0.355***	0.08	-0.138**	-0.204***	-0.137	-0.042
	(0.013)	(0.036)	(0.046)	(0.06)	(0.071)	(0.076)	(0.077)	(0.078)	(0.065)	(0.064)	(0.09)	(0.114)
Constant	-0.132	-0.637**	-0.917**	-1.436***	-2.064***	-2.90***	-4.166***	-6.225***	-7.946***	-8.473***	-7.99***	-7.135***
	(0.107)	(0.299)	(0.381)	(0.504)	(0.616)	(0.691)	(0.76)	(0.798)	(0.642)	(0.619)	(0.745)	(0.932)
Observations	365	363	362	360	357	354	350	342	334	330	318	306
R^2	0.984	0.931	0.901	0.848	0.785	0.698	0.583	0.464	0.458	0.475	0.458	0.397
RMSE	0.027	0.054	0.064	0.076	0.085	0.096	0.109	0.118	0.116	0.113	0.115	0.119

HAC robust standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

Figure 20: Univariate Autoregressive Forecasting Model & Diebold-Mariano Tests (1- to 60-Month Ahead Forecast Horizons) – Japanese Yen/USD

Dependent Variable: Log of Dollar-Yen Exchange Rate

	k-months ahead											
	1	3	4	6	9	12	16	24	32	36	48	60
DM	0	0	-0.0001	-0.0002	-0.0005	-0.001	-0.0023*	-0.005***	-0.009***	-0.011***	-0.009***	-0.0068**
Constant	(0)	(0.0001)	(0.0001)	(0.0002)	(0.0005)	(0.0008)	(0.0012)	(0.0019)	(0.0025)	(0.0029)	(0.0031)	(0.0029)
Log($/¥)	0.98***	0.924***	0.899***	0.855***	0.805***	0.75***	0.677***	0.57***	0.498***	0.475***	0.453***	0.464***
	(0.009)	(0.026)	(0.034)	(0.046)	(0.062)	(0.072)	(0.079)	(0.086)	(0.089)	(0.09)	(0.082)	(0.077)
Constant	-0.11**	-0.355***	-0.476***	-0.681***	-0.913***	-1.169***	-1.512***	-2.016***	-2.349***	-2.456***	-2.548***	-2.478***
	(0.044)	(0.124)	(0.16)	(0.221)	(0.293)	(0.343)	(0.378)	(0.411)	(0.429)	(0.434)	(0.396)	(0.372)
Obs.	366	366	366	366	366	366	366	366	366	366	366	366
R^2	0.985	0.936	0.911	0.868	0.82	0.754	0.661	0.531	0.455	0.423	0.39	0.419
RMSE	0.027	0.055	0.065	0.079	0.092	0.108	0.127	0.149	0.161	0.165	0.17	0.166

HAC robust standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

The USD/JPY and USD/GBP exchange rate forecasts do not exhibit the same uniform, linear behavior that the USD/CAD and USD/CHF forecasts do above. Moreover, the behavior of both currencies across the forecast horizon, though decidedly nonlinear, is remarkably similar between both seemingly dissimilar currencies. Attending first to the forecasts of the Japanese Yen (I attend to those of the pound sterling last to more seamlessly transition into the discussion regarding long-run forecasting using the annual data from the Bank of England dataset), I observe the following:

Figure 21: Evolution of Coefficients for USD/JPY Forecasts - Monthly Frequency - HAC Error Bars

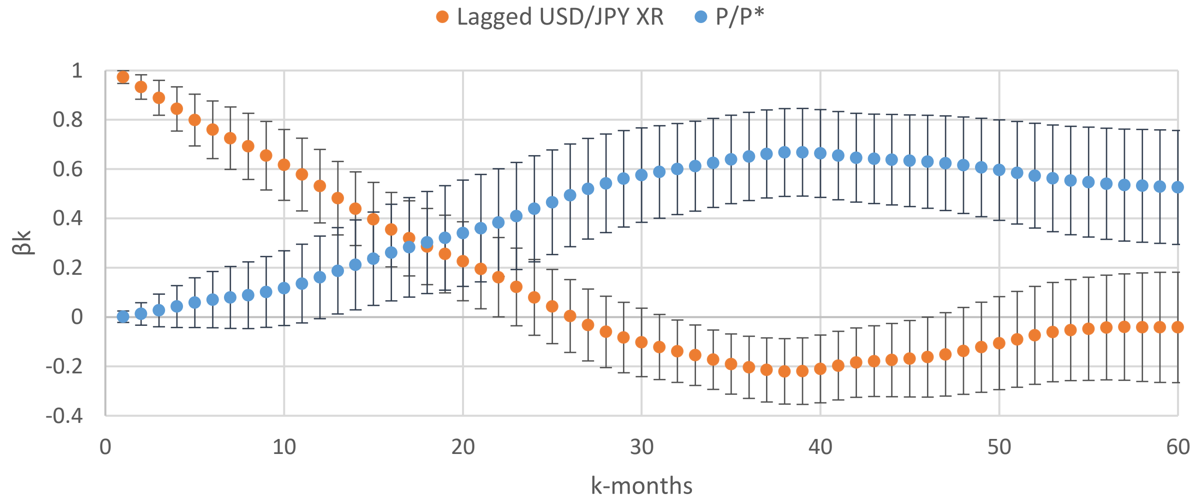

While the contribution of the fundamental and the lagged exchange rate do continue to mirror one another as observed in the previous bilateral exchange rate forecasts, the trajectory is not perfectly linear as before. There is a steady linear decline in the contribution of past exchange rate through the 38-month ahead forecast, and the contribution drops below zero shortly past the 2-year ahead forecast horizon. Unlike in the case of the Swiss franc forecasts, the negative coefficients in this context are statistically significant, implying that the past data, detracts from the contribution of the fundamental, in economic terms, at these horizons. This does effect does not persist through the remainder of the sample, however, and becomes insignificant past the 4-year ahead forecasts. Further, the decline is reversed somewhat past the 3-year ahead forecast horizon, though the coefficients remain negative indefinitely.

Correspondingly, the contribution of the theory component increases linearly, approximately through the 3-year ahead forecast, peaking at a value of 0.68, well below the maximum engendered in other currency forecasts. However, unlike the previous forecasts, this effect is not homogenous throughout the forecast horizons, and the contribution decreases somewhat beyond this horizon (though the β_k coefficients are statistically significant at the 99% level at all horizons beyond 16-months ahead). This implies there may be some optimization effect underlying the forecasts, where the contribution of theory is maximized around the 3-year forecast horizon, and the predictive power of these forecasting elements decreases summarily beyond this horizon as the contribution of the past is no longer significant and the contribution of theory decreases.

In this spirit, plots of R^2 and RMSE suggest similar patterns for the explanatory power and fit of forecasts relative to those exhibited in the USD/CHF forecasts. Note, however, while the trends are similar, the level values are altogether different.

Figure 22: Model Fit for USD/JPY Forecasts - R^2 and Root Mean Squared Error

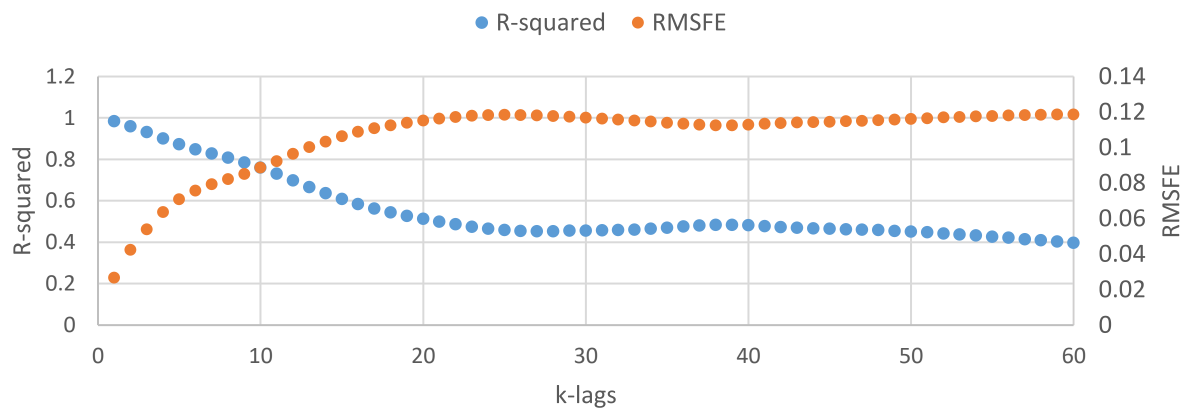

As before, the trajectory of R^2 values relative to the forecast horizon is negative—the decrease is largely linear through the 2-year ahead forecasts and then remains relatively constant at approximately 0.45 then exhibits further modest decreases past the 4-year ahead time horizon. While the plateau behavior is comforting and suggests, generally that the explanatory power of the model does not decrease indefinitely once it reaches a certain displacement, the 0.45 R^2 value is not particularly impressive, implying just less than half of the variability in the exchange rate movements can be explained by the hybrid model at these horizons. RMSE, similarly increases rapidly through the initial 2-years of forecasts, and then plateaus around 0.12 through the remainder of the sample. The relative stability of RMSE, in this context, is particularly comforting and seems to suggest that hybrid model may be better suited to accurate forecasting of the Japanese yen, consistently, for longer forecast-horizons. In this spirit, as with the prior currencies, the hybrid theory model outperforms the univariate autoregressive model *more so*, and more quickly, than the other hybrid theory models with respect to their given currency. Similarly, too, Diebold Mariano tests suggest the improvement in accuracy is statistically significant as soon as 1-year ahead relative to the univariate autoregressive model.[5]

[5] Refer to Figure 20 above for results of Diebold Mariano tests for USD/JPY exchange rate forecasts.

Figure 23: Hybrid Model Forecasting Regression Output (1- to 60-month Ahead Forecast Horizons) – British Pound/USD Exchange Rate Forecasts

Dependent Variable: Log of Dollar-Sterling Exchange Rate

	k-months ahead											
	1	3	4	6	9	12	16	24	32	36	48	60
Log(P/P*)	0.014	0.065*	0.093**	0.148**	0.218***	0.279***	0.35***	0.503***	0.613***	0.659***	0.76***	0.676***
	(0.013)	(0.037)	(0.047)	(0.062)	(0.071)	(0.073)	(0.07)	(0.08)	(0.095)	(0.098)	(0.1)	(0.093)
Log($/£)	0.98***	0.895***	0.851***	0.762***	0.645***	0.535***	0.391***	0.09	-0.166**	-0.272***	-0.534***	-0.512***
	(0.015)	(0.044)	(0.057)	(0.076)	(0.088)	(0.09)	(0.087)	(0.077)	(0.078)	(0.077)	(0.092)	(0.1)
Constant	-0.052	-0.247	-0.35*	-0.558**	-0.819***	-1.042***	-1.294***	-1.839***	-2.213***	-2.366***	-2.695***	-2.326***
	(0.053)	(0.15)	(0.191)	(0.253)	(0.289)	(0.297)	(0.291)	(0.344)	(0.411)	(0.423)	(0.421)	(0.39)
Observations	533	531	530	528	525	522	518	510	502	498	486	474
R^2	0.98	0.912	0.876	0.808	0.726	0.655	0.569	0.433	0.374	0.375	0.455	0.401
RMSE	0.024	0.049	0.058	0.071	0.084	0.093	0.102	0.113	0.114	0.111	0.098	0.098

HAC robust standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

Figure 24: Univariate Autoregressive Forecasting Model & Diebold-Mariano Tests (1- to 60-Month Ahead Forecast Horizons) – British Pound/USD

Dependent Variable: Log of Dollar-Sterling Exchange Rate

	k-months ahead											
	1	3	4	6	9	12	16	24	32	36	48	60
DM Constant	0.000	-0.0001	-0.0001	-0.0003	-0.0006	-0.001*	-0.0016**	-0.003***	-0.005***	-0.006***	-0.008***	-0.006***
	(0.000)	(0.0001)	(0.0001)	(0.0002)	(0.0004)	(0.0006)	(0.0007)	(0.0012)	(0.0017)	(0.0019)	(0.0024)	(0.0018)
Log($/£)	0.987***	0.945***	0.922***	0.876***	0.812***	0.749***	0.659***	0.476***	0.306***	0.236***	0.054	0.014
	(0.008)	(0.025)	(0.032)	(0.044)	(0.056)	(0.063)	(0.069)	(0.08)	(0.085)	(0.088)	(0.093)	(0.095)
Constant	0.007	0.028**	0.04**	0.064***	0.096***	0.129***	0.176***	0.271***	0.359***	0.396***	0.489***	0.505***
	(0.005)	(0.014)	(0.018)	(0.024)	(0.031)	(0.034)	(0.037)	(0.042)	(0.044)	(0.046)	(0.047)	(0.048)
Observations	533	531	530	528	525	522	518	510	502	498	486	474
R^2	0.98	0.91	0.872	0.797	0.703	0.615	0.503	0.284	0.129	0.08	0.005	0.000
RMSE	0.024	0.049	0.059	0.073	0.088	0.099	0.11	0.127	0.134	0.135	0.133	0.127

HAC robust standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

Transitioning now to a discussion of the pound sterling forecasts, I begin with the monthly analysis and then transition into the key discussion of annual forecasts using the Bank of England dataset. The evolution of β_k coefficients is like that of the Japanese yen forecasts.

Figure 25: Evolution of Coefficients for USD/GBP Forecasts - Monthly Frequency - HAC Error Bars

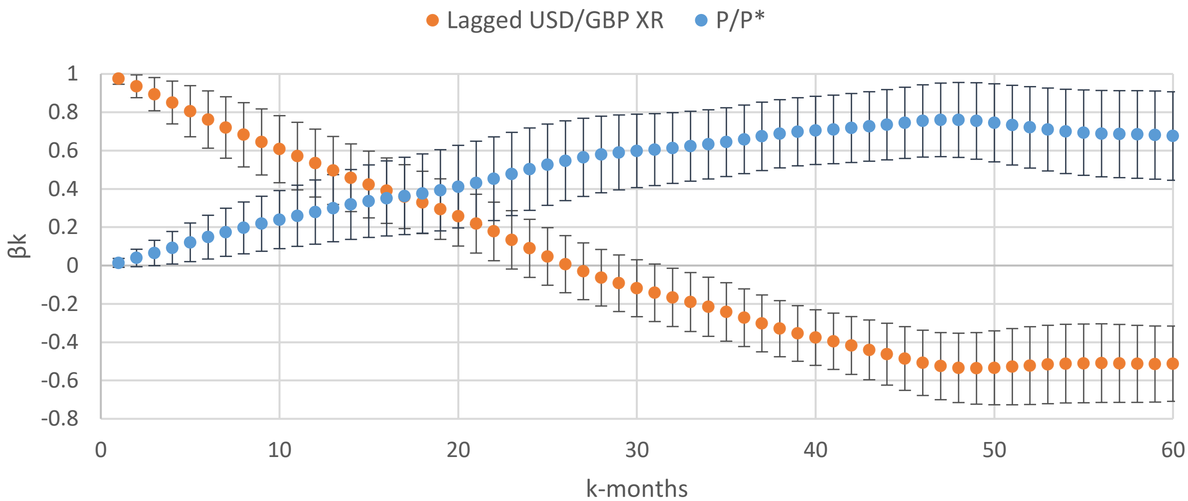

Both coefficients mirror one another in relative contribution, as observed previously, with a steady linear decline in the contribution of past exchange rate realizations that becomes negative shortly beyond the 2-year ahead forecast horizon. Unlike other exchange rate series, the hybrid model generates statistically significant negative coefficients for the lagged spot rate for the duration of the forecast horizon, which broadly reach larger magnitudes than any other series either. Similarly, there is an upward linear trend exhibited by the theory coefficient β_k representing increasing contribution through the 4-year ahead forecast of the USD/GBP exchange rate, peaking at a value of approximately 0.76. Following this, the forecasts through the fifth-year exhibit modest decreases in the contribution of the theory—the long-term trajectory beyond this point can be examined, somewhat, in the annual forecasts in the next section.

With these results, the inclusion of a theory component that serves as a fundamental for exchange rate deviations to self-correct about can unambiguously improve upon the random walk for medium-term exchange rate forecasts, with the effect becoming most prominent at longer time-horizons. This effect is robust to the choice of currency about which the exchange rate is anchored—while each coefficient series exhibits unique characteristics reflective of the underlying economic relationship between currencies, the broader trends are consistent across all four hybrid models.

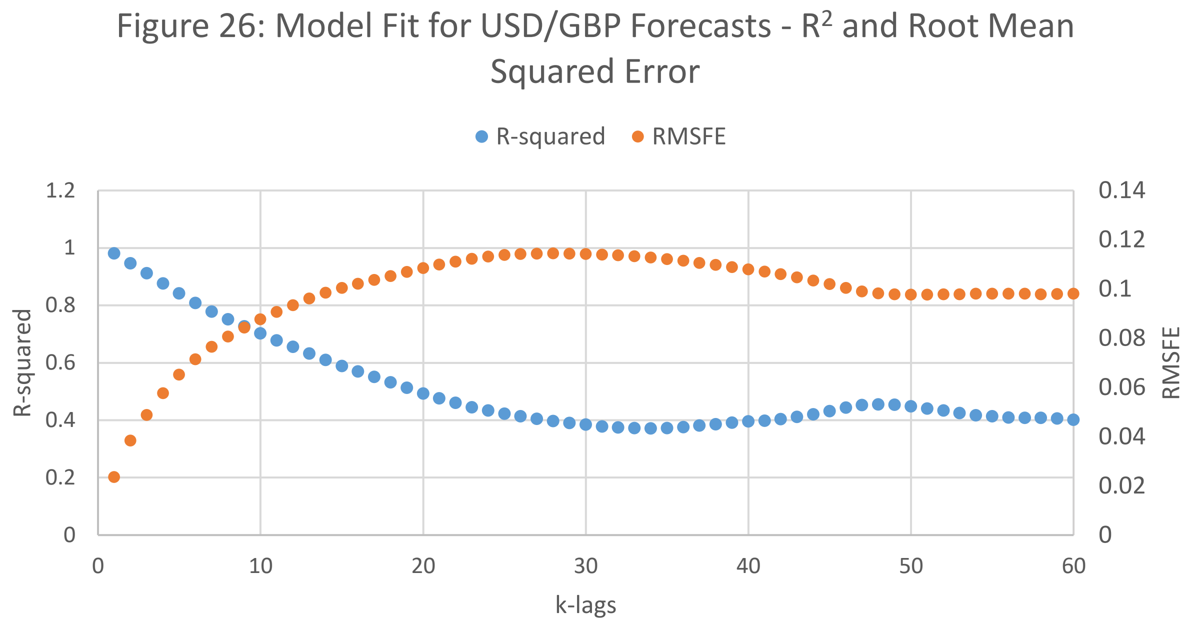

In the case of forecast accuracy/model fit for the monthly USD/GBP forecasts, R^2 again begins near to 1 and decreases as the forecast horizon increases, dropping to a low of 3.72 around the 3-year ahead forecast and then recovering slightly to 0.4 by the end of the forecasting sample. As exhibited in the other exchange rate forecasts, the explanatory power of the hybrid theory model decreases as the forecast horizon increases, despite the continued statistical and economic contribution of the theory component at these horizons. RMSE exhibits near the exact same trajectory of the other three currencies, rising rapidly through the first 2-years of k-step ahead forecasts, peaking at 0.114 and then declining somewhat to 0.098 at the end of the forecast horizon. Despite the somewhat unintuitive results borne out in the estimation of the theory coefficients, the monthly pound sterling forecasts exhibit, on average, the lowest RMSE of the forecasts in question and, unlike those of the Swiss franc and Canadian dollar, seems to stabilize in terms of forecasting accuracy and fit, with less indication of an interminable deterioration of the model as the forecast horizon increases. As with the other forecasted exchange rates, the hybrid model containing the theory component reliably outperforms the univariate autoregressive model in terms of forecast accuracy and explanatory power as soon as 2-months ahead. The difference in forecast accuracy, furthermore, is statistically significant at 1-year ahead forecasts per the results of the Diebold-Mariano tests.[6]

Broadly speaking, each series of exchange rate forecast exhibits the same general trajectory for both forecasting accuracy and coefficient evolution. The role of purchasing power parity, which serves as the fundamental in this context, through the β_k coefficient, becomes more economically and statistically significant as the forecast horizon increases. Conversely, the contribution of past exchange rate realizations (which serve as a proxy for the random walk from previous literature) declines in the presence of the fundamental at these horizons. While not all series exhibit an unequivocally positive linear trend in the β_k coefficients, this behavior is widely exhibited in all currencies through at least the first 3-year ahead forecast horizon, which constitutes a lag of 36 months—this behavior is summarily mirrored in the random walk component coefficients in all cases, which serves as a justification for our choice of specification which implies that the two will contribute to the forecast opposite of one another

[6] Refer to Figure 24 above for results of Diebold Mariano tests for USD/GBP exchange rate forecasts.

by construction. The consistency of these results across currencies suggests the specification and, by extension, the fundamental's improvement upon the random walk, is robust to exchange rate selection.

In all forecast series, I observe similar trends for model fit and forecasting accuracy measured by R^2 and RMSE. On one hand, R^2-values are highest at the lowest displacements in forecast horizon and unequivocally decrease as that horizon increases. The pace at which this deterioration in explanatory power occurs, as well as the resulting value after the 5-year ahead forecast horizon is currency specific. RMSE, on the other hand, increases rapidly as the initial forecast horizon increases and then, seems to stabilize as the horizon becomes increasingly long-term. Projections of RMSE could be neatly fit with a logarithmic curve, with rapid deterioration of forecast accuracy observed immediately, and a seemingly terminal value of error exhibited at some horizon in the future with respect to a given currency.

These results suggest that, despite the contribution of the purchasing power parity fundamental in the forecast specification increasing with forecast horizon, the relative performance of the model deteriorates across all currencies. I suspect, however, this is less a repudiation of the model itself and more a reflection of the challenges of forecasting exchange rates at such large displacements—despite the forecast horizon extending to only 5-years ahead, such a forecast encompasses 60-months of lagged data to develop the forecast. It is not terribly surprising that the quality of the model suffers somewhat. I find it more comforting that, in most cases, the effect does not seem permanent. Generally, explanatory power and accuracy stabilize throughout the forecast horizons—more research would be appropriate to determine how long these measures would remain stable for.

Based on the results of the forecasts, I consider the evolution of the theory coefficients across currencies to be a sufficient justification for an improvement upon the random walk that is robust to underlying exchange rate currency choice. To further bolster these results, and demonstrate the improvement that including purchasing power parity represents in this context, I conduct the same analysis with a univariate autoregressive model containing only lagged exchange rates as the independent variable to serve as a baseline for comparison of model fit and forecasting horizon. In all cases, the hybrid model outperforms the univariate autoregressive model in terms of R^2 and RMSE almost immediately (at either the 1-month ahead or 2-month ahead forecast horizon). The contribution of the past exchange rate data, in this specification, deteriorates both economically and statistically as the forecast horizon increases. Further, Diebold-Mariano tests of comparative forecast accuracy imply that, in all exchange rate forecasts, the difference in forecast accuracy is statistically significant, and in favor of the hybrid model, at around the 1-year ahead forecast horizon.

7. Annual Frequency (Bank of England) Exchange Rate Forecasting Results

Moving beyond what has largely been attended to in the literature, it is appropriate to test the efficacy of the hybrid theory model on truly long-run data with longer-term forecasts. While this analysis is restricted to dollar-sterling exchange rates due to availability of data, it provides possible insight into the evolution of forecasts beyond horizons typically examined in the literature today.

Figure 27: Hybrid Model Forecasting Regression Output (1- to 12-Year Ahead Forecast Horizons) – Full Sample, 1861-2015

Dependent Variable: Log of Dollar-Sterling Exchange Rate

	k-years ahead											
	1	2	3	4	5	6	7	8	9	10	11	12
Log(P/P*)	0.154***	0.407***	0.579***	0.679***	0.71***	0.673***	0.567***	0.485***	0.513***	0.574***	0.593***	0.585***
	(0.051)	(0.108)	(0.153)	(0.186)	(0.198)	(0.185)	(0.167)	(0.155)	(0.158)	(0.163)	(0.173)	(0.194)
Log($/£)	0.836***	0.562***	0.373**	0.26	0.221	0.257	0.367**	0.452***	0.418**	0.351**	0.328*	0.333
	(0.053)	(0.112)	(0.156)	(0.187)	(0.199)	(0.192)	(0.179)	(0.168)	(0.167)	(0.168)	(0.18)	(0.204)
Constant	-0.059**	-0.144***	-0.204***	-0.241***	-0.255***	-0.247***	-0.219**	-0.198**	-0.212**	-0.238**	-0.249**	-0.249**
	(0.028)	(0.054)	(0.074)	(0.089)	(0.095)	(0.09)	(0.086)	(0.087)	(0.095)	(0.102)	(0.109)	(0.116)
Observations	154	153	152	151	150	149	148	147	146	145	144	143
R^2	0.974	0.942	0.924	0.925	0.919	0.912	0.91	0.91	0.904	0.894	0.883	0.869
RMSE	0.078	0.117	0.133	0.13	0.134	0.139	0.14	0.139	0.143	0.15	0.158	0.166

HAC robust standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

Figure 28: Univariate Autoregressive Forecasting Model & Diebold-Mariano Tests (1- to 12-Year Ahead Forecast Horizons) – Full Sample, 1861-2015

Dependent Variable: Log of Dollar-Sterling Exchange Rate

	k-years ahead											
	1	2	3	4	5	6	7	8	9	10	11	12
DM Constant	-0.0002	-0.0016*	-0.0033*	-0.0047*	-0.0053*	-0.0049*	-0.0037	-0.0031	-0.0035	-0.0043	-0.0047	-0.0048
	(0.0002)	(0.0009)	(0.0019)	(0.0027)	(0.0031)	(0.0028)	(0.0023)	(0.0022)	(0.0024)	(0.0028)	(0.0031)	(0.0035)
Log($/£)	0.993***	0.979***	0.97***	0.967***	0.966***	0.970***	0.977***	0.981***	0.979***	0.977***	0.975***	0.973***
	(0.014)	(0.025)	(0.033)	(0.041)	(0.047)	(0.053)	(0.057)	(0.058)	(0.058)	(0.06)	(0.062)	(0.066)
Constant	0.001	0.012	0.015	0.012	0.006	-0.006	-0.022	-0.034	-0.037	-0.039	-0.044	-0.047
	(0.018)	(0.033)	(0.044)	(0.053)	(0.061)	(0.068)	(0.073)	(0.074)	(0.074)	(0.076)	(0.08)	(0.085)
Observations	155	155	155	155	155	155	155	155	155	155	155	155
R^2	0.973	0.935	0.906	0.888	0.874	0.866	0.865	0.865	0.855	0.841	0.829	0.815
RMSE	0.079	0.123	0.148	0.162	0.158	0.171	0.177	0.177	0.184	0.192	0.199	0.207

HAC robust standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

In applying the same methodology to the annual bilateral dollar-sterling exchange rate and price data from the Bank of England dataset, first with the full sample from 1861-2015, the following pattern of exchange rate coefficients borne out in the data is as follows:

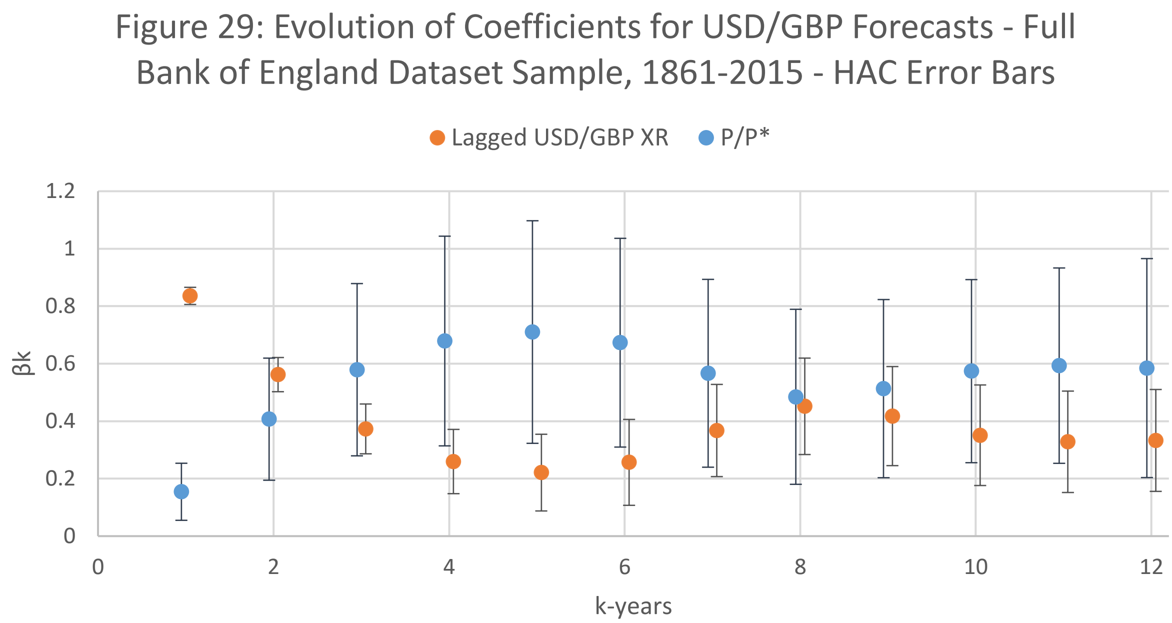

It is immediately apparent that the trajectory of coefficients observed at the monthly frequencies in the 5-year forecasting horizon from above does not hold consistently as that horizon increases. It is important to keep in mind, this analysis pertains to a vastly different dataset, with different frequencies of data and steeper increases in the forecast horizon. Institutional effects notwithstanding (more on this below), there is a steady, almost parabolic, pattern of these theory coefficients as lags increase, which peak at 0.71013 in this sample at the 5-year ahead forecast horizon. This contribution decreases modestly between the fifth and eighth periods and then rises again, seemingly steadying at approximately 0.60. Overall, the trend is positive, which suggests that the theory component of the hybrid exchange rate model contributes to predictive power in forecasting future exchange rates relative to the random walk even with this stark change in data and forecasting horizons.

This is particularly notable in the long-run as the distance between the forecasted exchange rate and the current exchange rate increases. The first-period forecasted β_k (0.15425) is significant at the 95% significance level, and all preceding forecasted β_k's from lag 2 through 12 are significant at the 99% level. As such, the role that purchasing power parity, as the theory component of this hybrid exchange rate forecasting model, is far from trivial economically, and is robustly significant, strengthening the argument that the inclusion of a fundamental does in fact improve upon the parsimonious random walk markedly.

Note, further, from the figure above that the coefficients for the random walk component mirror those of the theory component even in the parabolic pattern as the number of lags increases. As the relative role of theory increases, the role of the past decreases and vice versa. The random walk component of the regression equation and its coefficients are not as statistically robust as the theory component as detailed above. While all coefficients across the 12 lags are statistically significant, at least at the 90% level, there is less evidence of the remarkable statistical significance that prevailed in lags 2 through 11 of the theory coefficients. Thus, while the role of the past exchange rate series component becomes less relevant in its contribution to the predictive power of the forecast as lags increase, the contribution is both statistically and economically significant.

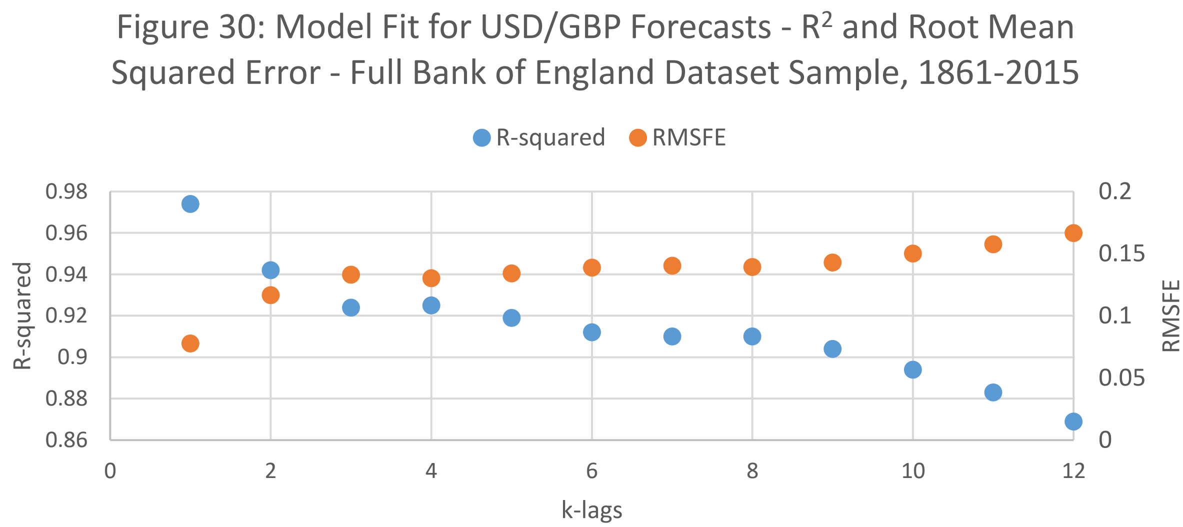

In accordance with assessments of model fit from the four currencies of analysis with monthly data, I defer to plots of R^2 and RMSE over the course of the 12 lags, as above. Consistent with the results above, the plot suggests that the model loses explanatory power over the time-horizon, and the RMSEs of the forecasts increase. The explanatory power of the model, as indicated by R^2 decreases as lags are added—from a remarkable fit of 0.97 down to 0.87 over the course of the 12 lags. Compared to the deterioration of explanatory power exhibited in the monthly analysis, however, this is the most robust trajectory of R^2 yet—note that, in conducting a forecast that is 12-years ahead given price and exchange rate data today, the forecasting model only loses 0.1 in terms of R^2, i.e. explanatory power. This is striking, especially considering the displacement in the data for these long-horizon forecasts. RMSE expectedly increases over the course of the 12 lags, beginning at 0.078 and ending at 0.166. These levels are well in-line with those of the monthly forecasts and, relatively speaking, the model has not deteriorated any more rapidly than the 5-year ahead monthly forecasts. I suspect however, given these results, the trend of increasing RMSE and decreasing R^2 continues beyond the 12-year ahead forecast. Overall, despite the possible impact of institutional considerations discussed above, the hybrid model performs remarkably well with these truly long-horizon forecasts given the long-run data sample.

One possible explanation for the deterioration of explanatory power and increasing RMSE with large displacements in time could be the role of institutions in determining exchange rates independently of the theory that underlies this model as discussed in the Data Description section. It is reasonable to assume that attempting to forecast exchange rates under one exchange rate regime may be error-prone when using exchange rate data from another. This would be especially apparent when forecasting on horizons of 10-12 years ahead when markedly different economic institutions/exchange rate regimes, in some instances, only span the length of 30 years (i.e. the Bretton Woods era).

While I provide economic justifications of the choice of years for subdividing the broader dataset in the Data Description section above, it is prudent to evaluate to what extent those choices are empirically motivated by the data itself. Quandt-Andrews unknown breakpoint tests behave, generally, as simple Chow structural break tests but utilize a more penalizing critical value table making it more difficult for the null of no structural break to be rejected. Using Quandt-Andrews unknown breakpoint tests on the entire sample (trimmed by 15% on either extreme), I reject the null hypothesis of no structural break at the 95% level with respect to an underlying structural break at 1949. Results from Quandt-Andrews unknown breakpoint test are as follows:

Figure 31: Quandt-Andrews Unknown Breakpoint Tests for annual USD/GBP Exchange Rate Series

Null Hypothesis: No breakpoints within 15% trimmed data

Equation Sample:	1861-2015			1939-2015	
Test sample:	1885-1992			1951-2004	
	t-value	p-value		t-value	p-value
Maximum Wald F-statistic (1949)	11.090	0.016	Maximum Wald F-statistic (1974)	6.503	0.129
Exp Wald F-statistic	2.634	0.024	Exp Wald F-statistic	1.480	0.101
Ave Wald F-statistic	2.990	0.045	Ave Wald F-statistic	1.662	0.160

I defer to the Maximum Wald F-statistic when drawing conclusions about the empirical evidence of structural breaks. While the 1949 structural breakpoint from the full sample is five years ahead of my choice of 1944 for the Bretton Woods period, I see it as justification, broadly, for the fundamentally different institutional and economic considerations that this period exhibits. Since the choice of 1944 is motivated specifically with the conception and implementation of the system, I will continue to use this date as the start of the period, but do argue the results of the break test substantiate the use of this separate subperiod generally.

Seeking further empirical justification for more recent structural breaks, I conduct the same test methodology on the 1939-2015 sample (Bretton Woods to present day). In this sample, the most significant break presented is for 1974, in line with the breakdown of the Bretton Woods system, but the break is not statistically significant at the 90% level with a t-value of only 1.480. Still, the choice of this year serves as further motivation for the choice of these subdivisions. I attend, now, to regression output from the subperiods and corresponding plots as above.

Figure 32: Hybrid Model Forecasting Regression Output (1- to 12-Year Ahead Forecast Horizons) – Gold Standard, 1861-1939

Dependent Variable: Log of Dollar-Sterling Exchange Rate

	k-years ahead											
	1	2	3	4	5	6	7	8	9	10	11	12
Log(P/P*)	0.12	0.295	0.321	0.292	0.242	0.171	0.058	-0.008	0.003	0.046	-0.028	-0.091
	(0.085)	(0.187)	(0.215)	(0.188)	(0.172)	(0.132)	(0.102)	(0.088)	(0.075)	(0.089)	(0.081)	(0.098)
Log($/£)	0.718***	0.333	0.189	0.16	0.129	0.123	0.205*	0.271**	0.246***	0.182**	0.278***	0.321**
	(0.136)	(0.232)	(0.238)	(0.227)	(0.221)	(0.161)	(0.106)	(0.11)	(0.084)	(0.081)	(0.096)	(0.147)
Constant	0.209*	0.473***	0.647***	0.742***	0.887***	1.039***	1.132***	1.155***	1.168***	1.181***	1.177***	1.233***
	(0.106)	(0.112)	(0.141)	(0.14)	(0.122)	(0.093)	(0.099)	(0.105)	(0.093)	(0.096)	(0.109)	(0.123)
Observations	78	77	76	75	74	73	72	71	70	69	68	67
R^2	0.706	0.427	0.32	0.37	0.308	0.217	0.183	0.203	0.219	0.201	0.239	0.214
RMSE	0.079	0.111	0.117	0.093	0.088	0.088	0.084	0.075	0.068	0.067	0.065	0.065

HAC robust standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

Figure 33: Univariate Autoregressive Forecasting Model & Diebold-Mariano Tests (1- to 12-Year Ahead Forecast Horizons) – Gold Standard, 1861-1939

Dependent Variable: Log of Dollar-Sterling Exchange Rate

	k-years ahead											
	1	2	3	4	5	6	7	8	9	10	11	12
DM Constant	-0.0001	-0.0006	-0.0007	-0.0006	-0.0004	-0.0002	0.000	0.000	0.000	0.000	0.000	-0.0001
	(0.0002)	(0.0008)	(0.001)	(0.0008)	(0.0006)	(0.0003)	(0.0001)	(0.000)	(0.000)	(0.000)	(0.000)	(0.0001)
Log($/£)	0.841***	0.63***	0.505***	0.433***	0.345***	0.261***	0.221***	0.202***	0.178***	0.157**	0.171	0.147
	(0.073)	(0.076)	(0.069)	(0.074)	(0.07)	(0.051)	(0.05)	(0.066)	(0.066)	(0.075)	(0.104)	(0.122)
Constant	0.256**	0.596***	0.797***	0.913***	1.055***	1.192***	1.255***	1.285***	1.323***	1.358***	1.335***	1.374***
	(0.113)	(0.118)	(0.114)	(0.126)	(0.122)	(0.098)	(0.103)	(0.131)	(0.131)	(0.143)	(0.185)	(0.212)
Observations	79	79	79	79	79	79	79	79	79	79	79	79
R^2	0.697	0.385	0.244	0.176	0.107	0.055	0.035	0.024	0.016	0.009	0.013	0.006
RMSE	0.08	0.113	0.12	0.096	0.091	0.089	0.084	0.075	0.068	0.067	0.065	0.065

HAC robust standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

Figure 34: Hybrid Model Forecasting Regression Output (1- to 12-Year Ahead Forecast Horizons) – Bretton Woods, 1944-1973

Dependent Variable: Log of Dollar-Sterling Exchange Rate

	k-years ahead											
	1	2	3	4	5	6	7	8	9	10	11	12
Log(P/P*)	-0.042	-0.085	-0.137	-0.079	0.051	0.17	0.271	0.445**	0.676***	0.775***	0.87***	1.061***
	(0.108)	(0.225)	(0.329)	(0.334)	(0.283)	(0.255)	(0.235)	(0.207)	(0.241)	(0.243)	(0.24)	(0.286)
Log($/£)	0.938***	0.863***	0.817***	0.736***	0.626**	0.522*	0.433*	0.319*	0.192	0.152	0.085	-0.067
	(0.095)	(0.182)	(0.233)	(0.26)	(0.271)	(0.258)	(0.23)	(0.178)	(0.154)	(0.139)	(0.123)	(0.143)
Constant	0.116	0.25	0.369	0.358	0.266	0.183	0.114	-0.043	-0.28	-0.408	-0.494	-0.636*
	(0.136)	(0.28)	(0.411)	(0.449)	(0.388)	(0.327)	(0.286)	(0.272)	(0.306)	(0.31)	(0.313)	(0.357)
Observations	30	30	30	30	30	30	30	30	30	30	30	30
R^2	0.896	0.763	0.66	0.565	0.528	0.521	0.541	0.591	0.683	0.754	0.728	0.69
RMSE	0.053	0.08	0.096	0.108	0.113	0.113	0.111	0.105	0.092	0.081	0.085	0.091

HAC robust standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

Figure 35: Univariate Autoregressive Forecasting Model & Diebold-Mariano Tests (1- to 12-Year Ahead Forecast Horizons) – Bretton Woods, 1939-1973

Dependent Variable: Log of Dollar-Sterling Exchange Rate

	k-years ahead											
	1	2	3	4	5	6	7	8	9	10	11	12
DM Constant	0.000	0.000	-0.0001	0.000	0.000	-0.0002	-0.0005	-0.0012	-0.0026*	-0.0032*	-0.004**	-0.007
	(0.0001)	(0.0003)	(0.0006)	(0.0003)	(0.0002)	(0.0006)	(0.0008)	(0.0011)	(0.0015)	(0.0016)	(0.0017)	(0.0042)
Log($/£)	0.91***	0.812***	0.740***	0.693***	0.655***	0.619***	0.589***	0.573***	0.574***	0.58***	0.560***	0.492***
	(0.084)	(0.155)	(0.2)	(0.234)	(0.233)	(0.209)	(0.185)	(0.169)	(0.162)	(0.152)	(0.153)	(0.153)
Constant	0.082	0.175	0.242	0.284	0.313	0.340	0.364*	0.372*	0.360*	0.341*	0.359*	0.442**
	(0.083)	(0.156)	(0.203)	(0.242)	(0.244)	(0.224)	(0.201)	(0.189)	(0.186)	(0.18)	(0.181)	(0.176)
Observations	30	30	30	30	30	30	30	30	30	30	30	30
R^2	0.892	0.753	0.643	0.548	0.51	0.496	0.506	0.529	0.571	0.622	0.567	0.411
RMSE	0.053	0.08	0.096	0.108	0.113	0.114	0.113	0.11	0.105	0.099	0.106	0.124

HAC robust standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

Figure 36: Hybrid Model Forecasting Regression Output (1- to 12-Year Ahead Forecast Horizons) – Current Float, 1974-2015

Dependent Variable: Log of Dollar-Sterling Exchange Rate

	k-years ahead											
	1	2	3	4	5	6	7	8	9	10	11	12
Log(P/P*)	0.141	0.504***	0.843***	1.003***	0.994***	0.825***	0.544***	0.296**	0.232*	0.29	0.319	0.241
	(0.1)	(0.119)	(0.141)	(0.173)	(0.168)	(0.149)	(0.134)	(0.113)	(0.137)	(0.193)	(0.208)	(0.171)
Log($/£)	0.684***	0.177*	-0.278***	-0.539***	-0.59***	-0.433***	-0.127	0.115	0.129	0.01	-0.07	-0.02
	(0.091)	(0.095)	(0.102)	(0.169)	(0.186)	(0.152)	(0.127)	(0.104)	(0.099)	(0.17)	(0.185)	(0.107)
Constant	0.032	-0.031	-0.099	-0.108	-0.077	-0.012	0.074	0.165*	0.216*	0.228*	0.247*	0.291**
	(0.068)	(0.097)	(0.104)	(0.092)	(0.092)	(0.104)	(0.103)	(0.097)	(0.111)	(0.128)	(0.138)	(0.14)
Observations	42	42	42	42	42	42	42	42	42	42	42	42
R^2	0.661	0.366	0.364	0.504	0.57	0.498	0.391	0.353	0.3	0.237	0.192	0.154
RMSE	0.077	0.106	0.106	0.094	0.087	0.094	0.104	0.107	0.111	0.116	0.12	0.122

HAC robust standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

Figure 37: Univariate Autoregressive Forecasting Model & Diebold-Mariano Tests (1- to 12-Year Ahead Forecast Horizons) – Current Float, 1974-2015

Dependent Variable: Log of Dollar-Sterling Exchange Rate

	k-years ahead											
	1	2	3	4	5	6	7	8	9	10	11	12
DM Constant	-0.0001	-0.0014*	-0.0047**	-0.0082**	-0.0093**	-0.0071**	-0.003**	-0.0009	-0.0005	-0.0009	-0.001	-0.0006
	(0.0001)	(0.0007)	(0.002)	(0.0037)	(0.004)	(0.0029)	(0.0015)	(0.0007)	(0.0007)	(0.0011)	(0.0013)	(0.0008)
Log($/£)	0.752***	0.461***	0.256*	0.164	0.164	0.235	0.329**	0.367***	0.331***	0.267**	0.216	0.199
	(0.062)	(0.1)	(0.152)	(0.188)	(0.206)	(0.19)	(0.151)	(0.114)	(0.105)	(0.121)	(0.134)	(0.135)
Constant	0.122***	0.275***	0.384***	0.434***	0.433***	0.389***	0.329***	0.303***	0.322***	0.36***	0.389***	0.398***
	(0.034)	(0.054)	(0.077)	(0.094)	(0.105)	(0.104)	(0.093)	(0.075)	(0.062)	(0.065)	(0.073)	(0.074)
Observations	42	42	42	42	42	42	42	42	42	42	42	42
R^2	0.648	0.268	0.075	0.019	0.022	0.076	0.197	0.284	0.251	0.169	0.112	0.099
RMSE	0.078	0.112	0.126	0.13	0.13	0.126	0.118	0.111	0.114	0.116	0.124	0.125

HAC robust standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

For the sake of brevity and ease of comparison, I plot the coefficient evolution of purchasing power parity and lagged exchange rates for each subperiod in one summary graphic—connecting lines between data points have been added for clarity. Deferring to the analysis presented in the Data Description, 1861-1939 is considered the Gold Standard ("GS"), 1944-1973 is considered the Bretton Woods era ("BW"), and 1974-2015 is the current float ("Float").

Figure 38: Evolution of Coefficients for USD/GBP Forecasts

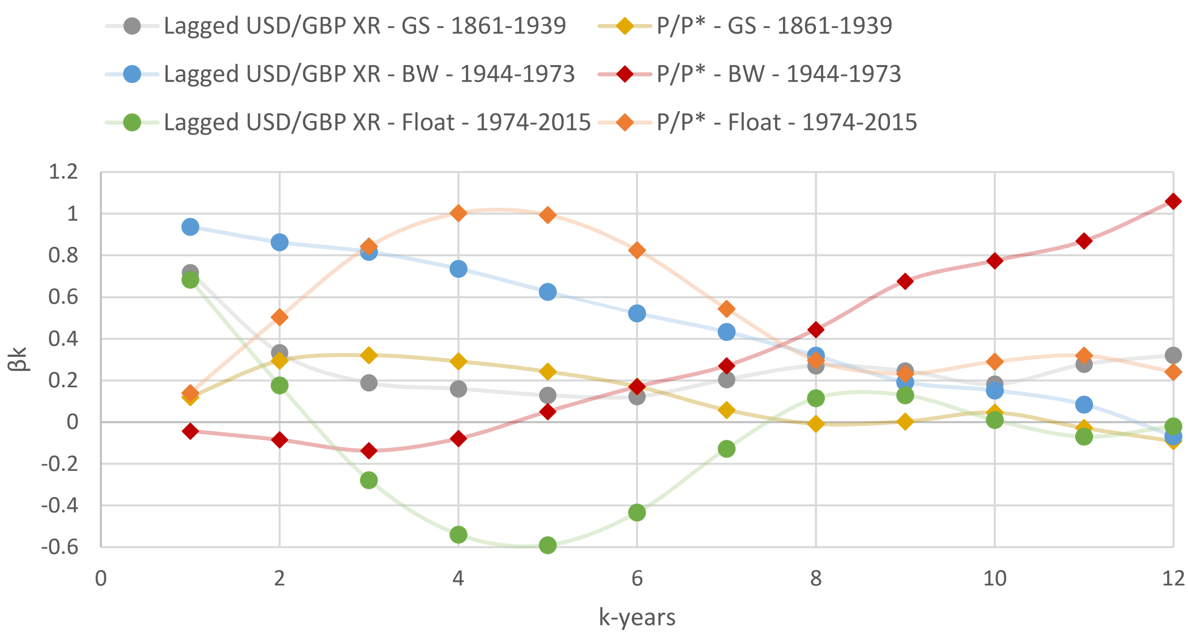

Figure 39: Model Fit for USD/GBP Forecasts - R^2 and RMSE

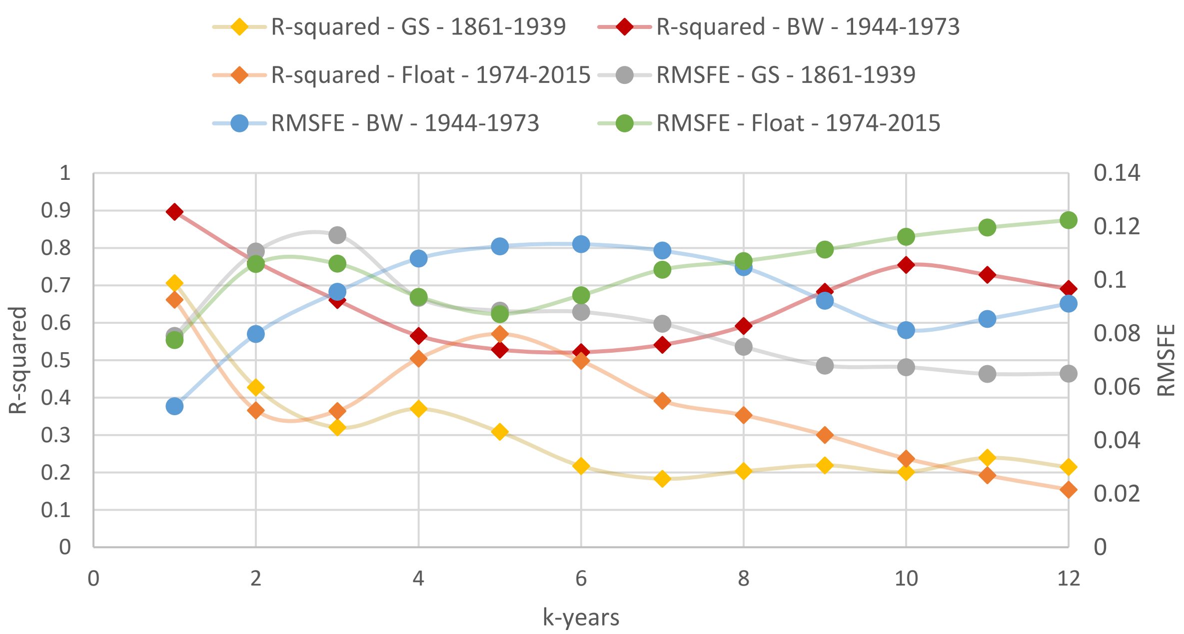

It is immediately apparent that the consistency of results from the monthly analysis and full sample annual analysis falls apart when the sample is broken up into subperiods.[7] I find dramatically different trajectories for coefficient estimates across forecast horizons and for metrics of explanatory power and forecast accuracy—this implies that the results are highly sensitive to the underlying institutional considerations and, furthermore, the "fundamental" that underlies the broader sample is likely fundamentally different than that which is observed in each distinct subperiod.

With respect to the Gold Standard, the β_k coefficient is *never* statistically significantly different from zero at any forecast horizon, and therefore does not contribute any meaningful predictive power to the hybrid model in this context. This consideration notwithstanding, the magnitude of the coefficient rises only through the 3-year ahead forecast horizon, peaking at only approximately 0.3, and then declining unequivocally as the horizon increases from there. Incidentally, the coefficient for the random walk component does not precisely mirror these movements, in stark contrast to other samples. Moreover, the coefficient only becomes significant at the 90% level at the 7-year ahead forecast horizon. Rather than exhibiting a decline in the contribution of the past, the Gold Standard sample features increasing contribution at longer-horizons in a statistical sense. This may imply that either the purchasing power parity fundamental is not applicable as a theory variable in this instance, or the specification is inappropriate for the sample. In terms of forecast accuracy and model fit, R^2 declines throughout the sample to 0.214 at the 12-year ahead forecast horizon, and RMSE, surprisingly, declines somewhat as the horizon increases. I posit this is likely due to the eventual significance of the random walk component at longer horizons as described above. Further, the hybrid theory model does not perform statistically significantly better than the univariate autoregressive model per the results of Diebold Mariano tests.[8]

Results from the Bretton Woods era are more in line with those observed in the full and monthly samples as above. The contribution of theory via the β_k coefficient increases in statistical and economic significance as the forecast horizon increases, becoming significant at the 99% level at the 9-year ahead horizon and reaching 1.06 in magnitude at the 12-year ahead forecast horizon. The contribution of the random walk more appropriately mirrors that of purchasing power parity, beginning close to 1 and decreasing in economic and statistical significance with increasing displacements in the forecast horizon. These results are highly complementary. As regression output included in the Appendix will show, only the 7-year ahead forecast incorporates statistically significant contributions from both components—prior to this horizon, only the past exchange rate contributes meaningfully, with the theory component driving the forecasts from 8-years ahead and beyond.

Evaluation metrics are far more like those of the monthly and full sample annual analyses as well. RMSE increases rapidly at the onset of the forecasts, peaking at 0.113 in the 6-year ahead forecast and then decreasing somewhat. Correspondingly, the R^2 of the forecasts steadily decreases through the 6-year ahead forecast and then increases. Note that as the role of the fundamental becomes statistically significant, the forecasts improve (in the 8- to 10-year ahead forecast horizons), suggesting the Bretton Woods era may be particularly suited to this model given the relative constancy in exchange rates and prices. Since prices are likely to fluctuate more than exchange rates given the nature of the Bretton Woods exchange rate regimes, it stands to reason that the exchange rate itself might be the best predictor in the medium-term and that the purchasing power parity component best captures what variations do present in the truly long-horizon forecasts.

---

[7] Refer to Appendix for individual coefficient trajectory graphs with HAC error bars.

[8] Refer to Figures 33, 35, and 37 above for results of Diebold Mariano tests for annual USD/GBP exchange rate forecasts.

Forecasts conducted during the current float are also more consistent with those from the monthly frequencies—I would expect to see a certain degree of similarity between the two forecasts for USD/GBP exchange rates through the 5-year horizon. In this spirit, the results are quite similar through these forecast horizons, with the contributions of the random walk and theory mirroring each other more precisely than in other subdivided samples. The β_k coefficient is significant at the 90% level for the 1-year ahead forecast horizon and rapidly approaches 1 by the four-year ahead horizon, declining in magnitude and statistical significance thereafter, and offering no meaningful contribution at the 10-year ahead forecast horizon and beyond. The contribution of the random walk, much like in the monthly analysis, is statistically significant and negative from the 3-year ahead forecast to the 6-year ahead forecast and is insignificant thereafter. R^2 values dip through the 2- and 3-year ahead forecast horizons and then recovers through the 5-year ahead, while RMSE drops somewhat at that horizon and then increases indefinitely thereafter.

These results potentially point to a general breakdown in the model past the 5- and 6-year time horizons. This does not seem surprising given the highly volatile nature of exchange rates in the modern float—while it is promising that the hybrid model can improve upon the random walk in the medium-term, it is unclear that either model perform adequately in this environment beyond those horizons. Results from the Diebold Mariano tests, similarly, support a statistically valid improvement in accuracy relative to the univariate autoregressive model in the 2-year to 7-year ahead forecasts, but neither model is more or less accurate at longer forecast horizons.

8. Cointegration and Vector Error Collection Models

As mentioned briefly in the time-series properties section of this analysis, the existence of cointegration between the fundamental and exchange rate series could provide additional predictive power to existing forecasts and provide a sounder theoretical foundation for the results of the hybrid specification I utilize here. Two series are said to be cointegrated if both underlying series are I(1), but a linear combination of the two series is I(0), or stationary. As such, while little can be said of the long-run behavior of a unit root series given its inherent non-stationarity, cointegrating series can provide information about where one of the underlying non-stationary series is headed in the long-run relative to the other. In this context, the degree to which log purchasing power parity series and log bilateral exchange rate series are cointegrated should serve as motivation for the use of the relative price ratio as the underlying fundamental to which exchange rate deviations self-correct. By construction, cointegrating series should, to some extent, be self-correcting with respect to one another.

Figure 41: Single-Equation Engle Granger Cointegration Test Results
Null Hypothesis: Series are not cointegrated

Dependent Variable	Tau-Statistic	p-value	Z-Statistic	p-value
Log USD/CAD Exchange Rate	-1.695	0.681	-5.902	0.656
Log USD/CHF Exchange Rate	-3.528	0.031	-19.577	0.059
Log USD/JPY Exchange Rate	-3.478	0.036	-17.36	0.091
Log USD/GBP Exchange Rate	-3.611	0.0248	-26.258	0.0142

Note: The independent variable for all Single-Equation Engle Granger tests for cointegration is the corresponding log purchasing power parity series for each exchange rate series.

Using Single-Equation Engle Granger tests for cointegration, I reject the null of no cointegration at the 90% significance level for all series except for that of the USD/CAD exchange rate and purchasing power parity relationship, suggesting some fundamental differences in the two series that undermine the long-term predictability of one series relative to the other. Broadly, though, the degree to which the exchange rate and purchasing power series are cointegrated is promising—we would expect that the inclusion of log purchasing power parity in the exchange rate forecast may provide additional empirically valid predictive power about the long-run behavior of the exchange rate.

To further motivate this intuition, and to more directly exploit the possible improvement in predictive power vis-à-vis the cointegrating relationships developed herein, Vector Error-Correction Models (or VECMs) *directly* incorporate the co3integrating relationship for the forecasting of changes in cointegrating series. Broadly, I begin by estimating the cointegrating relationships between the exchange rate and relative price ratio series for the pound sterling, Japanese yen, and Swiss franc (as the Canadian series are not cointegrated) by estimating:

$$s_t = \alpha + \beta_t f_t + \upsilon_t, \widehat{\upsilon_t} \sim iid$$

Where s_t is the spot exchange rate of a given currency at time t and $f_{i,t}$ is the corresponding fundamental, in this case, the relative price ratio between the given currency and the United States. By the cointegrating relationship, as indicated, the estimated residuals $\widehat{\upsilon_t}$ will be I(0) and iid, The use of this stationary residual term that results from regressions of the cointegrating relationship can provide additional predictive power for changes in the spot rate.

The specification of the VECM in the context of the hybrid theory model is as follows:

$$\Delta s_t = \alpha + \beta \Delta s_{t-1} + \delta \Delta f_{t-1} + \gamma \widehat{\upsilon_{t-1}} + \varepsilon_t$$

The primary differences here, between both the cointegrating regression and hybrid specification are the include of lagged first differences of the exchange rate and relative price ratio and the inclusion of the fitted residuals from the cointegration regression. So, rather than forecasting spot exchange rates in level terms, the VECM forecasts *changes* in the exchange rate. Of interest is the coefficient γ from the fitted residuals—a negative, significant coefficient implies that there is a tendency for deviations from the cointegrating relationship to be "pulled" back to zero in the long-run. This intuition is complementary to the specification of the hybrid model in the body of the analysis *without* cointegration elements included.

Figure 40: Modified Vector Error-Correction Model Regression Output

Dependent Variable: Change in Bilateral Spot Exchange Rate (First Difference)

	Exchange Rate Series		
	Swiss Franc	British Pound	Japanese Yen
Lagged Exchange Rate	0.283***	0.365***	0.321***
(First Difference)	(0.035)	(0.046)	(0.043)
Lagged Relative Price	-0.153	-0.074	-0.112
Ratio (First Difference)	(0.287)	(0.124)	(0.372)
Lagged Fitted	-0.021***	-0.032***	-0.032***
Cointegration Residual	(0.008)	(0.012)	(0.01)
Constant	0.002*	-0.001	0.002
	(0.001)	(0.001)	(0.002)
Observations	532	532	364
R^2	0.089	0.145	0.128

HAC robust standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

From the regression output, there is overwhelming evidence of this error-correcting behavior in the modified hybrid model with first-differenced dependent and independent variables. In all regressions, the coefficient estimate for the cointegration residual (γ) is negative and statistically significant at the 99% level. Again, the negative coefficient lends support to the intuition that large deviations from the cointegrating relationship (i.e. when the spot exchange rate is far from the linear relationship with the relative price ratio) will be largely self-correcting in the long-run.

9. Forward Market Analysis Results

Given the methodology employed to conduct forecasts of exchange rates with the hybrid model involves direct step-ahead forecasting, there are corresponding rates in the forward market from which additional predictive power may be ascertained. To further motivate this approach, it is appropriate to first illustrate the respective error series to get a sense of their relationship at first-pass. Recall the hybrid model errors are those generated from the difference between the hybrid forecasting model and the actual exchange rate series, while the forward market errors are generated from the difference between the forward rate transacted at the step-ahead horizon and the corresponding realized spot rate:

Figure 41: Forward Market & Hybrid Model Error Series (1-Month Ahead)

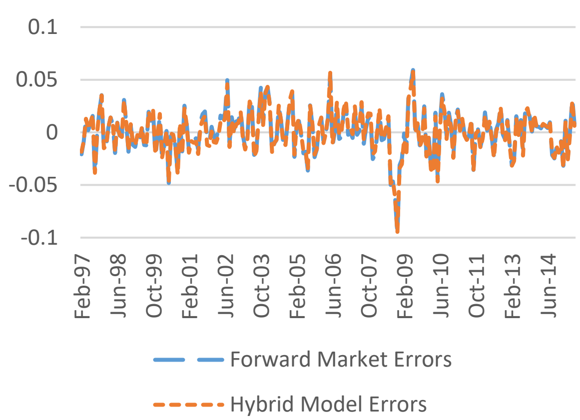

Figure 42: Forward Market & Hybrid Model Error Series (3-Month Ahead)

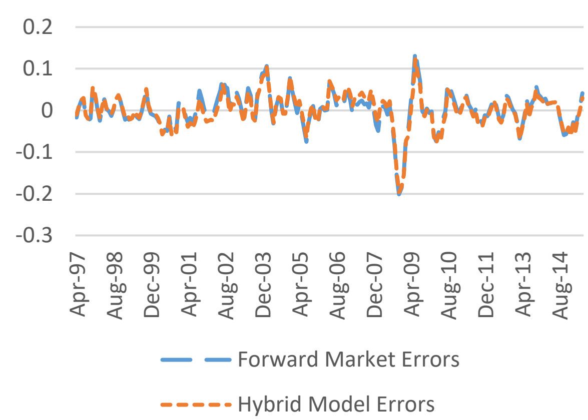

Figure 43: Forward Market & Hybrid Model Error Series (6-Month Ahead)

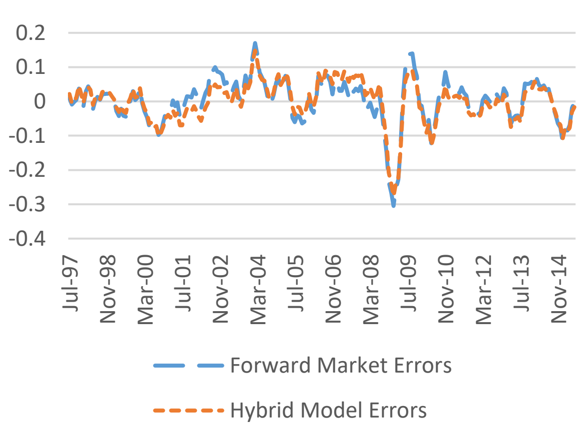

Figure 44: Forward Market & Hybrid Model Error Series (1-Year Ahead)

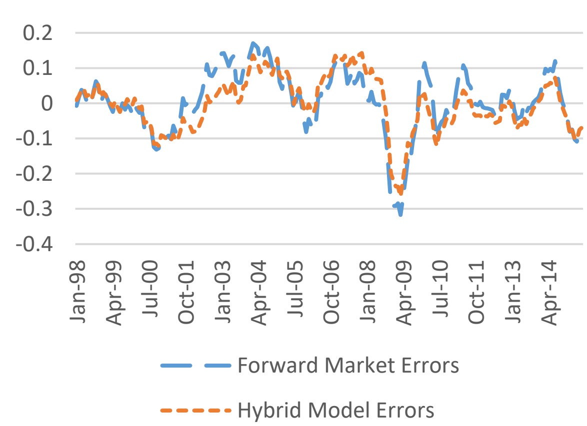

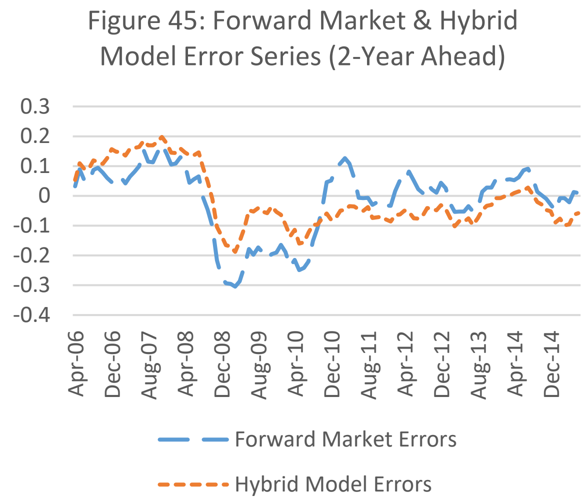

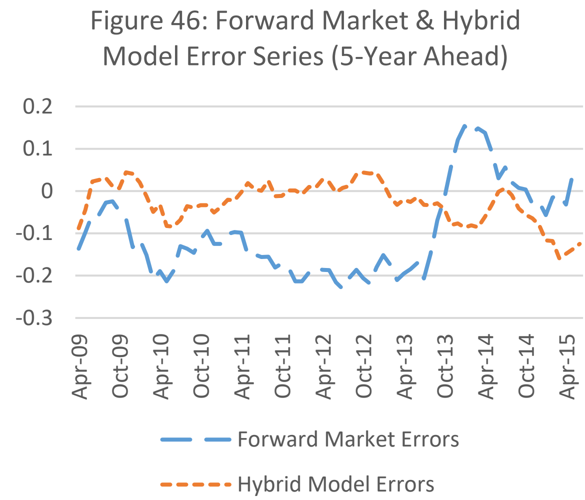

Again, to empirically assess the degree to which information in the forward market can be used to predict the errors of the hybrid forecast model above, I conduct forecast encompassing regressions of the following specification:

$$e_{i,t+k} = \alpha + \beta_i \eta_{i,t+k} + \xi_{i,t+k}$$

Where $e_{i,t+k}$ is the hybrid model forecast error series from the k-period ahead forecast for a given currency i , and $\eta_{i,t+k}$ is the error series constructed from the differences between log spot exchange rates at period k and the forward exchange rates established with information at time t to be transacted k-periods ahead. Regression results are as follows:

Figure 47: Forward Market Forecast Encompassing Regressions - USD/GBP Exchange Rates

Dependent Variable: Hybrid Model Forecast Errors, k-periods ahead

	Forecast Horizon					
	1-Month	3-Month	6-Month	1-Year	2-Year	5-Year
Forward Market Exchange Rate Errors, k-periods ahead	0.981***	0.946***	0.876***	0.767***	0.598***	-0.222***
	(0.006)	(0.02)	(0.032)	(0.044)	(0.101)	(0.071)
Constant	0.000	-0.001	-0.003	-0.007	-0.002	-0.049***
	(0.000)	(0.001)	(0.003)	(0.005)	(0.016)	(0.013)
Observations	221	219	216	210	111	75
R^2	0.988	0.950	0.877	0.776	0.500	0.227

HAC robust standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

Conducting the analysis of the forward exchange rates per the forecast encompassing methodology specified above, I find that, in all horizons (1-month, 3-month, 6-month, 1-year, 3-year, and 5-year forward rates), the errors from the forward market (with respect to the realized exchange rates to which they correspond) have a statistically significant relationship with the resulting errors of the forecast model. With all six forecast encompassing regressions, the forward errors are statistically significant at the 99% level in prediction of the errors from the hybrid model. The magnitude of the coefficient is always below 1, and decreases in magnitude as the forecast horizon increases, which is consistent with the results from the hybrid model alone. Thus, at all horizons, forward exchange rates established with the same information available for the construction of the hybrid forecast seem to contribute to the predictive power of the hybrid model insofar as they can predict the errors of the hybrid model that follow. This provides evidence of a weakness in the hybrid forecasting model which can be improved by inclusion of the forward market in some capacity—since the errors of the hybrid model can be predicted as such, it can be improved. Determining how to do so is outside the scope of the research, and is left for future research.

This effect is most pronounced with the 1-month ahead forward rates—the coefficient estimate for the forward errors is 0.981 with an R^2 of 0.988, suggesting there is significant explanatory power embodied in the forward errors with respect to the errors from the hybrid model. At the 5-year ahead forecast horizon, however, the coefficient is statistically significant and *negative*—I do not have an explanation, either theoretical or economic for this phenomenon and defer an explanation to further research on the subject. I would posit that at this horizon, given the deterioration of the hybrid model and the large displacement implicit in the forecasts, forward or otherwise, could be to blame for unintuitive results such as this. It is also possible that the small sample size (only 75 months of observations are included in the actual regression) is also biasing results. While forwards for longer time horizons exist, I have not included them in this analysis. It would be insightful to assess the broader trend of these forecast encompassing regressions at longer horizons.

Broadly, however, despite the unintuitive result from the 5-year ahead forecast, the results from the forward market analysis suggest that that errors that result from forward exchange rates relative to the realized exchange rate they are forward against have significant, statistically robust predictive power with respect to the errors from the hybrid theory model of interest in this analysis. Further research on this subject would be helpful toward understanding how to directly incorporate this supposed predictive power into more complex exchange rate forecasting models. Further, generalizing the results beyond the USD/GBP exchange rate would be valuable for assessing the robustness of this contribution to changes in currency (i.e., with the additional exchange rate forecasts conducted in this analysis).

10. Conclusion

While the random walk model as popularized by Meese and Rogoff is likely the most accurate model for forecasting exchange rates in the short run, this analysis finds strong evidence to suggest that the inclusion of an underlying fundamental like purchasing power parity can improve on this simplistic approach, particularly in the long-run. Using an adapted specification of the hybrid model first postulated in Mark (1995), this analysis finds that theory, in the form of purchasing power parity, has both a statistically and economically significant role contributing predictive power to forecasts of a number of bilateral exchange rates, at both the monthly and annual frequency. Furthermore, the contribution becomes more economically significant as the time horizon expands. In this context, the random walk serves as a foil to the fundamental—as the contribution of purchasing power parity increases over displacements in time, the contribution of the past realizations of the exchange rate diminishes, though it remains statistically significant and nontrivial in magnitude.

These results are largely robust to the choice of country-specific bilateral exchange rate, as demonstrated through forecasts conducted on USD/CAD, USD/CHF, USD/JPY, and USD/GBP bilateral exchange rates at the monthly-frequency up to 5-years ahead. This analysis serves as a contribution to the work of Mark (1995) and others in this vein by expanding upon the existing framework with additional observations that include the financial crisis. Of note in this analysis, beyond the analysis at the monthly-frequency, I conduct long-run forecasts of dollar-sterling exchange rates with the Bank of England dataset. Results from these forecasts lend further credibility to the use of the purchasing power parity fundamental with a much broader dataset and at longer horizons—however, I do find that the quality of the forecasts deteriorate beyond the 5- to 6-year forecast horizon used at the monthly frequency.

Further, there are challenges associated with conducting forecasts on a sample of data that spans so many broad global events, institutions and exchange rate regimes. To assess the robustness of the forecasts and fundamental to changes in underlying institutions, I subdivide the dataset and conduct forecasts with the hybrid specification. Results from these forecasts are dramatically different from those conducted on the full sample and dramatically different from one another—despite the degree to which purchasing power parity can be considered a fundamental to which the deviations in exchange rate self-correct, this research suggests that the fundamental is *highly* sensitive to changes in underlying institutional assumptions.

While the coefficients for the theory and past components provide sufficient evidence to reject the predominance of the random walk for exchange rate forecasting at medium- and long-run forecast horizons, I defer to conventional measures of forecast evaluation and explanatory power to more broadly assess the quality of the forecasts at each horizon. Regarding model fit and accuracy, plots of R^2 and RMSE suggest that there is an inherent trade-off between the accuracy of the model and the quality of the model as the forecast horizon increases. While the country-specific evolution of these metrics varies, the general trend borne out in the data is decreasing R^2 with the forecast horizon and increasing RMSE. Broadly this implies that the quality of the forecast deteriorates as the forecast horizon increases, despite the continued contribution of purchasing power parity at these horizons.

In some instances, this pattern seems to stabilize somewhat throughout the forecasts, and does not suggest the trend will continue indefinitely (in the case of the Swiss franc, Japanese yen and pound sterling at the monthly frequency) and other instances that seem to suggest the deterioration continues well beyond the forecast horizon analyzed (annual pound sterling and the Canadian dollar). As indicated, I see this as a reflection of the inherent challenge in forecasting a highly volatile series such as exchange rates at exceptionally long-run forecast horizons—despite this deterioration in overall forecast quality, I am confident that the inclusion of the purchasing power parity theory component improves the forecasts relative to the alternative of, say, the random walk or univariate autoregressive model.

To empirically evaluate this claim, I test the forecast accuracy of the hybrid theory model against that of a univariate autoregressive model that excludes the possible role of the fundamental in conducting the forecasts. In terms of explanatory power (measured with R^2) and model fit (measured with RMSE and AIC), the hybrid theory model improves upon the more parsimonious autoregressive model as soon as the 1-month ahead forecast horizon, but *not* relative to the forward rate. Since R^2 would, by construction, be higher with the inclusion of an additional regressor, a more robust test of forecast accuracy is in order. In this spirit, I conduct Diebold Mariano tests for all forecasts to determine whether the difference in squared forecast error between the two models is statistically significant. Except for the Gold Standard subperiod, I find statistically significant evidence of an improvement in accuracy with the use of the hybrid theory model. Generally, in the monthly frequency forecasts, this occurs approximately 1-year ahead in terms of forecast horizon.

Beyond direct evaluation per this methodology, I extend the analysis to the forward exchange rate market to determine whether there is economically valuable predictive content implicit in the errors from the forward rates. To test this assumption, I create residual series for both the hybrid forecasts and from the difference between the forward rate and realized spot rate, and then conduct forecast encompassing regressions to assess the predictive power of the forward errors relative to the forecast errors. I find that, for the pound sterling exchange rates, there is statistically and economically significant predictive power contained in forward rates toward the end of explaining the errors of the hybrid model. More research is needed to directly incorporate this contribution into more complex, theory-based exchange rate forecasting models.

In the spirit of additional research, this analysis is somewhat limited in scope—there are boundless opportunities to apply this hybrid model to a battery of other exchange rates and price data, as well as to apply other indicators in the theory component (such as the Taylor Rule variables from other research described above. Beyond an analysis of other countries in the current framework, the historical analysis of the dollar-sterling exchange rates is insightful for understanding the scope of (and limitations of) the role of a fundamental across institutions. Naturally, an analysis of this nature would be limited by the availability of data. More tangibly, perhaps, a more expansive analysis on the forecast encompassing regressions from the forward market would be valuable—since it is not the predominant focus of this research, the analysis is somewhat terse.

One key limitation I recognize in this analysis deserving of further attention is the comparison of the hybrid model to some baseline. I conclude that the inclusion of the lagged exchange rate in the specification provides sufficient evidence of the improvement upon the random walk when incorporating the theory element directly in the specification. However, a *direct* comparison against a constructed random walk series (in terms of forecast evaluation metrics) would be useful for more robustly assessing the predictive power the hybrid model provides in excess of a simple random walk.

One additional limitation worth addressing is the underlying assumptions of the choice of fundamental. In this use of the "fundamental exchange rate" is the implicit assumption that purchasing power parity is an appropriate theoretical underpinning for exchange rate forecasting in the hybrid model context. There are important historical limitations that deserve our attention if we are to continue in this vein. Purchasing power parity as we understand it was first introduced in 1918 by Gustav Cassel. As such, there are approximately 60 years of observations for price levels that pre-date the purchasing power parity model. The question that arises, then, is whether a model developed in the 20th century can appropriately describe data from the 19th century.

I would contend that purchasing power parity is substantiated with appropriate precedent to warrant the use in this context, and is powerful enough that it can be used to describe data prior to its inception. However, the issue of assuming a single "fundamental exchange rate" as described above is again called into question. If purchasing power parity is not appropriate in this context, an alternative must be provided which could be borne out from the existing literature (Taylor Rule coefficients, for instance). If there is an alternative structural model employed prior to the start of our sample, we might compare its relevancy to the use of purchasing power parity. I am not sure if such a model exists and, as such, this issue is deserving of further study.

11. Appendix – Selected Graphics

Figure A1: Evolution of Coefficients for USD/GBP Forecasts - Full Bank of England Dataset Sample, 1861-2015 - HAC Error Bars

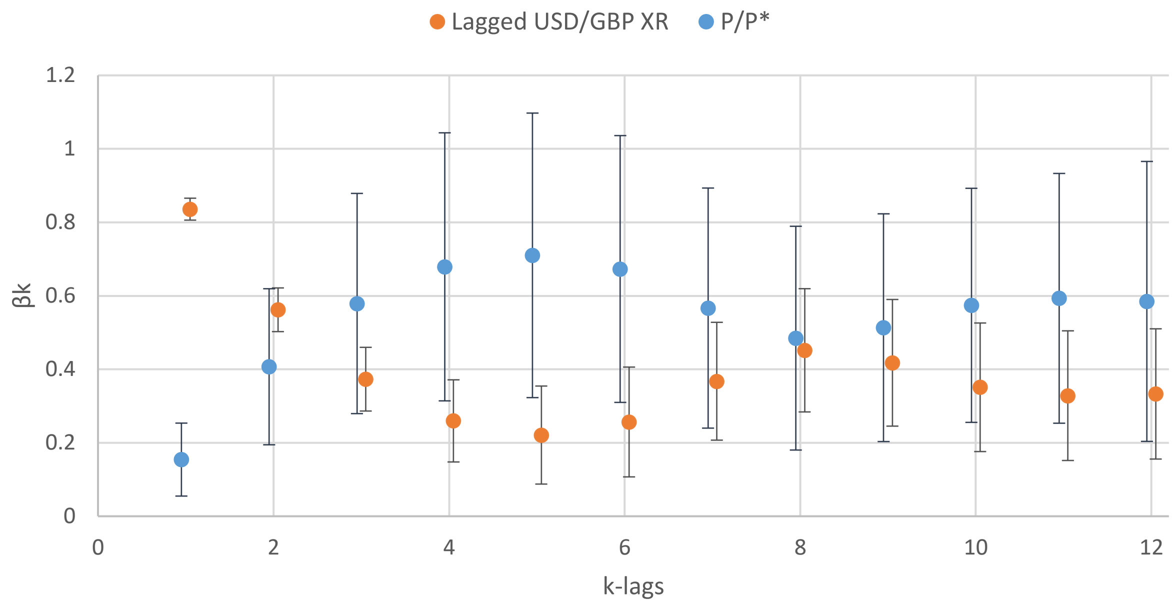

Figure A2: Evolution of Coefficients for USD/GBP Forecasts - Gold Standard, 1861-1939 - HAC Error Bars

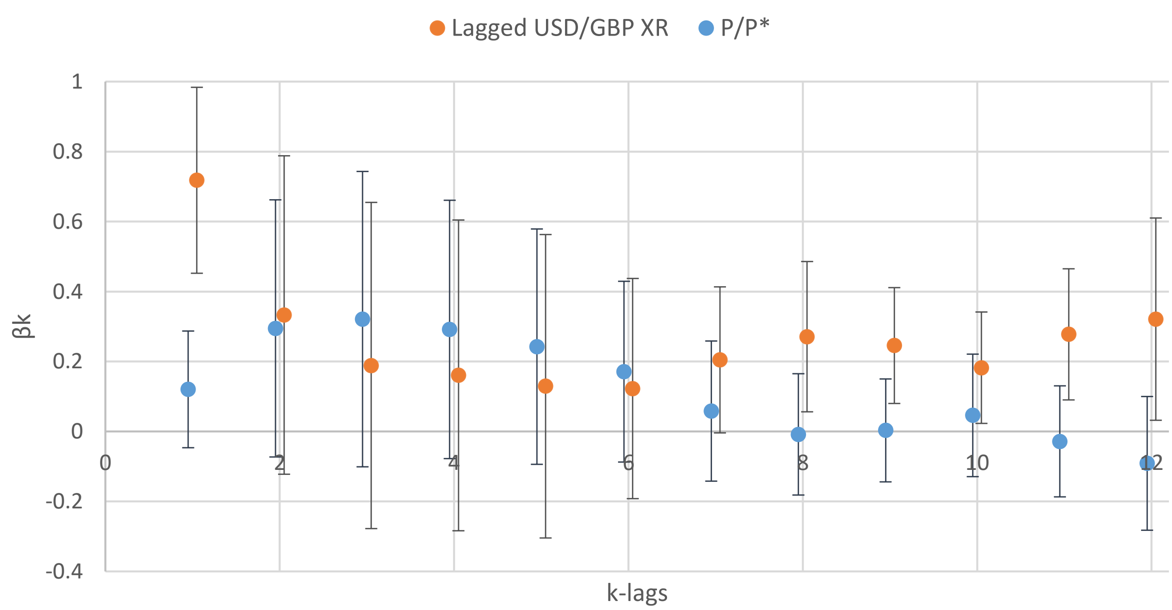

Figure A3: Evolution of Coefficients for USD/GBP Forecasts - Bretton Woods Era - 1944-1973 - HAC Error Bars

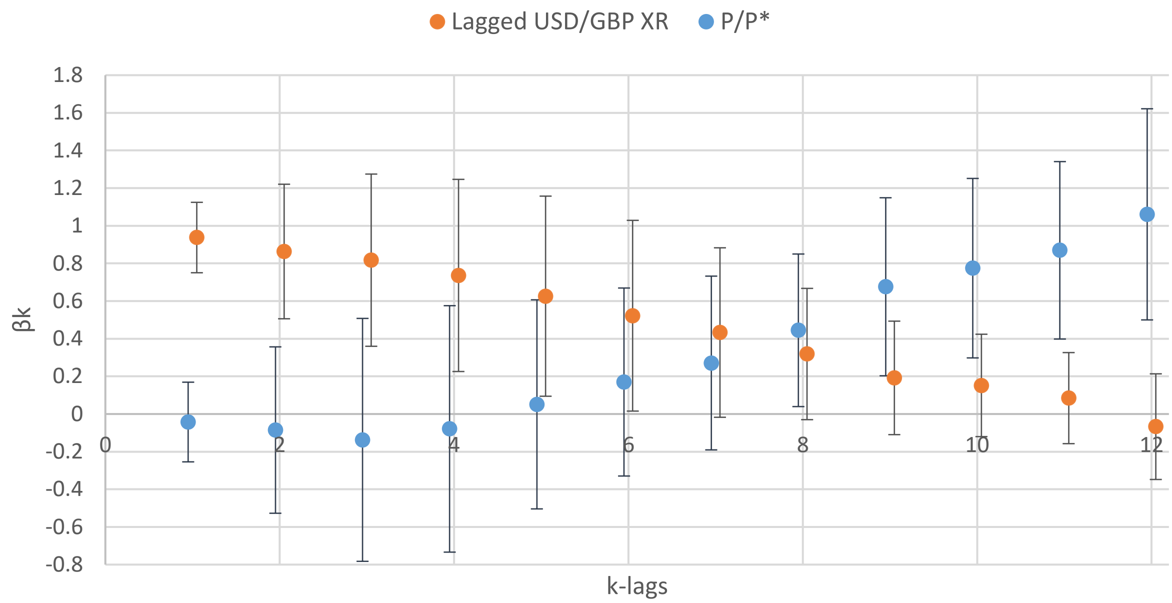

Figure A4: Evolution of Coefficients for USD/GBP Forecasts - Current Float - 1974-2015 - HAC Error Bars

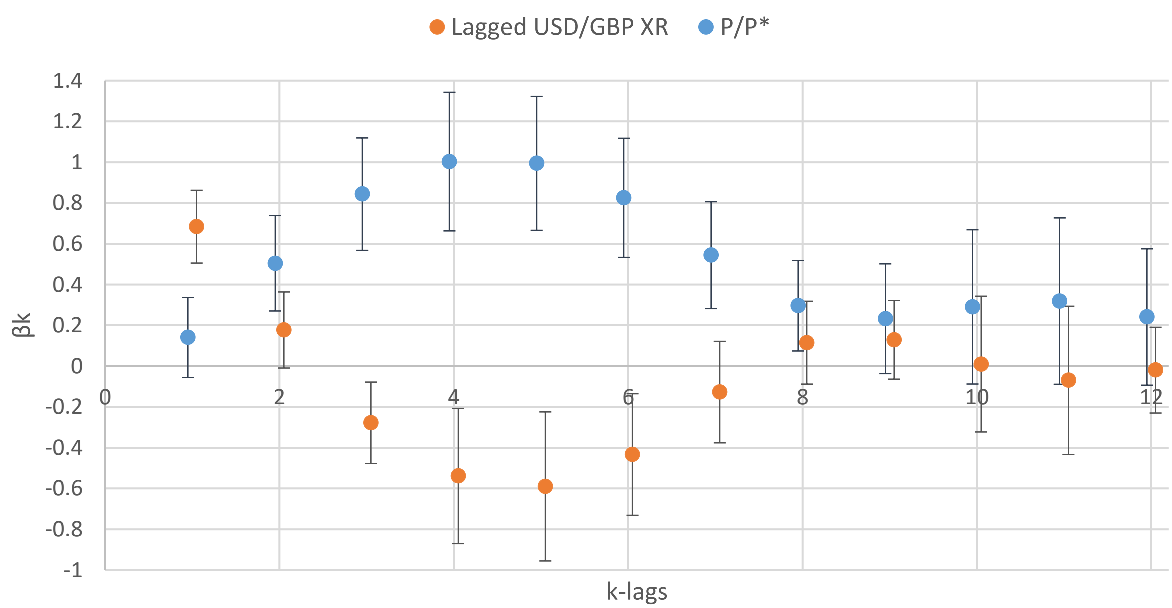

13. References

Abhyankar, Abhay, Lucio Sarno, and Giorgio Valente. 2005. "Exchange rates and fundamentals: evidence on the economic value of predictability." *Journal of International Economics* (Elsevier Science B.V.) 66: 325-348.

Boothe, Paul, and Debra Glassman. 1987. "Off the Mark: Lessons for Exchange Rate Modelling." *Oxford Economic Papers* (Oxford University Press) 39 (3): 443-457.

Burns, Kelly, and I. Moosa. 2012. "Can exchange rate models outperform the random walk? Magnitude direction and profitability as criteria." *Economia Internazionale* 65 (3): 473-490.

Choi, Doo-Yull, and Nelson, C. Mark. 1997. "Real exchange-rate prediction over long horizons." *Journal of International Economics* (Elsevier Science B.V.) 43: 29-60.

Christoffersen, P.F., and F.X. Diebold. 1998. "Cointegration and Long-Horizon Forecasting." *Journal of Business and Economic Statistics* 16 (4): 450-458.

Diebold, Francis X. 2006. *Elements of Forecasting.* 4th. South-Western College Publishers.

Dimsdale, N., and R. Thomas. 2016. "Three Centuries of Data - Version 2.3." *Bank of England.* Accessed May 24, 2017. http://www.bankofengland.co.uk/research/Pages/onebank/threecenturies.aspx.

Engel, Charles, Nelson C Mark, and Kenneth D West. 2012. "Factor Model Forecasts of Exchange Rates." *NBER Working Paper Series* (NBER) 1-20.

Engel, Charles, Nelson C Mark, and Kenneth D. West. 2007. "Exchange Rate Models Are Not As Bad As You Think." In *NBER Macro Economics Annual 2007, Volume 22*, edited by Daron Acemoglu, Kenneth Rogoff and Michael Woodford, 381-441. Chicago: University of Chicago Press.

Grossmann, Axel, and Marc W. Simpson. 2011. "Can a relative purchasing power parity-based model outperform a random walk in forecasting short-term exchange rates?" *International Journal of Finance & Economics* 16 (4): 375-392.

Grossmann, Axel, and Marc W. Simpson. 2010. "Forecasting the Yen/U.S. Dollar exchange rate: Empirical evidence from a capital enhanced relative PPP-based model." *Journal of Asian Economics* (Elsevier Science B.V.) 21 (5): 476-484.

Leitch, Gordon, and J. Ernest Tanner. 1991. "Economic Forecast Evaluation: Profits Versus the Conventional Error Measures." *The American Economic Review* (The American Economic Association) 81 (3): 580-590.

Lothian, James R. 1991. "A History of Yen Exchange Rates." *Japanese Financial Market Research* (Elsevier Science B.V.) 1-22.

Lothian, James R., and Mark P. Taylor. 1996. "Real Exchange Rate Behavior: The Recent Float from the Perspective of the Past Two Centuries." *Journal of Political Economy* (The University of Chicago Press) 104 (3): 488-509.

—. n.d. "Two Hundred Years of Sterling Exchange Rates and the Current Float." London: Irish Management Institute Centre for Economic Policy Research.

MacDonald, Ronald, and Ian W. Marsh. 1997. "On Fundamentals and Exchange Rates: A Casselian Perspective." *Review of Economics and Statistics* 79 (4): 655-664.

Mark, Nelson C. 1995. "Exchange Rates and Fundamentals: Evidence on Long-Horizon Predictability." *American Economic Review* 85 (1): 201-218.

Meese, Richard A, and Kenneth Rogoff. 1983. "Empirical Exchange Rate Models of the Seventies: Do they fit out of sample?" *Journal of International Economics* (North-Holland Publishing Company) 14: 3-24.

Meese, Richard, and Kenneth Rogoff. 1988. "Was it Real? The Exchange Rate-Interest Differential Relation OVer the Modern Floating-Rate Period." *The Journal of Finance* (Wiley for the American Finance Association) 43 (4): 933-948.

Papell, David H., and Tanya Molodtsova. 2009. "Out-of-sample exchange rate predictability with Taylor rule fundamentals." *Journal of International Economics* (Elsevier Science B.V.) 77: 167-180.

Taylor, Alan M. 2002. "A Century of Purchasing-Power Parity." *Review of Economics and Statistics* 84 (1): 139-150.