

A perspective to Uncovered Interest Rate Parity (UIP): FX Forward Contracts as a tool to hedge long-run currency risk

A THESIS IN THE JOHN MOLSON SCHOOL OF BUSINESS MASTER OF SCIENCE PROGRAM

Presented in Partial Fulfilment of the Requirements for the

Master of Science Finance (Co-op) at Concordia University

Montreal, Quebec, Canada

March 21, 2026

Supervisor: Professor Dr. Kun Ho Kim

©Aakash Kumar, 2026

School of Graduate Studies

This is to certify that the thesis prepared

By: Aakash Kumar

Entitled: A perspective to Uncovered Interest Rate Parity (UIP): FX Forward Contracts as a tool to hedge long-run currency risk

and submitted in partial fulfillment of the requirements for the degree of

Master of Science (Finance)

complies with the regulations of the University and meets the accepted standards with respect to originality and quality.

Signed by the final examining committee:

_____ Thesis Supervisor

Dr. Kun Ho Kim

_____ Thesis Committee Member

Dr. Parianen Veeren

_____ Thesis Committee Member

Dr. Sergey Isaenko

Approved by _____ Graduate Program Director

Dr. Frederick Davis

Dr. Anne-Marie Croteau, Associate Dean of John Molson School of Business

A perspective to Uncovered Interest Rate Parity (UIP): FX Forward Contracts as a tool to hedge long-run currency risk

Aakash Kumar

Abstract

This work investigates whether foreign exchange (FX) forward contracts offer an effective hedge for long-run currency risk, defined as the cumulative horizon over which short-run deviations dissipate, constructed by systematically rolling 3-month and 6-month forward contracts, considering the constant empirical failures of Uncovered Interest Rate Parity (UIP) at short horizons. Drawing on a chronological literature review of the UIP and International Fisher Effect (IFE), it traces how early short-term rejections, post-crisis regime shifts, and recent methodological advances, particularly Durbin regressions, have reshaped my understanding of the link between interest differentials and expected exchange rate movements. Using 3-month / 6-month forward data and econometric techniques that explicitly address overlapping observations and autocorrelation, the study reassesses whether UIP holds over slightly longer horizons when currency exposure is hedged through systematically rolled FX forward contracts. The study evaluates the extent to which long-run UIP validity restores the theoretical equivalence between hedged and unhedged positions, and identifies conditions under which risk premia, structural breaks, and low-interest-rate environments continue to generate meaningful deviations. Taken together, the findings provide an updated perspective on the long-run viability of FX forwards as a hedging instrument, clarifying when they deliver unbiased protection against currency risk and when they may embed systematic costs or lost opportunities for international investors.

Table of Contents

Section A.	Introduction	1
Section B.	Chronological Literature Review	2
	1. Early tests and rejections in 1980s - 1990s	3
	2. Evolution to Long-Horizon Studies and Partial Validation (2000s).....	4
	3. The Impact of Financial Crises, Regime Shifts, and Risk (Post-2008).....	5
	4. Recent advances in the 2020s using Durbin regressions and data	7
Section C.	Theoretical Foundation, and Methodology	10
	1. The UIP, IFE, and Forward Rate Framework	10
	2. Forward Premium Puzzle	11
	3. Durbin Regression.....	12
	a. Long-Run Coefficient (γ)	13
Section D.	Data Description	14
	Data Sources and Sample	14
Section E.	Analysis of Data and Preliminary Findings	14
	1. Summary Statistics – 3M returns/premiums.....	14
	a. Descriptive Statistics	14
	b. Distributions	15
	c. Time-series assessment.....	16
	2. Summary Statistics – 6M returns/premiums.....	17
	a. Descriptive Statistics	17
	b. Distributions	18
	c. Time-series assessment.....	19
	3. Volatility comparison	19
	4. Correlation analysis of spot returns and forward premiums	20
Section F.	Lag(s) Selection Methodology and Results.....	20
Section G.	Testing and Inferences	21
Section H.	Robustness.....	24

Section I. Practical Implications and Contributions	25
Section J. Conclusion and Limitations	25
Section K. Citations and Tables	27
Appendix	39

List of Figures

Figure 1 - Distribution of Spot Returns (3-month).....	29
Figure 2 - Time series: 3M Spot return vs Forward premium	30
Figure 3 - Distribution of Spot Returns (6-month).....	31
Figure 4 - Time series: 6M Spot return vs Forward premium	32
Figure 5 - Volatility comparison (3m vs 6m).....	33
Figure 6 - Volatility comparison (3m vs 6m).....	33
Figure 7 - BIC Values Across Lags.....	34

List of Tables

Table 1 - Data Characteristics	35
Table 2 - Observations post variable construction.....	35
Table 3 - 3M spot returns and 3M forward premiums (summary statistics).....	35
Table 4 - 6M spot returns and 6M forward premiums (summary statistics).....	36
Table 5 - BIC and AIC (Optimal Lag(s))	36
Table 6 - Test of UIP (Conventional Fama).....	37
Table 7 - Test of UIP (Durbin Regression - Long-run Coefficient).....	37
Table 8 - Subsample Test of UIP Pre/Post 2008 Financial Crises (Durbin Regression - Long-run Coefficient)	38

Section A. Introduction

One of the most enduring puzzles in international finance is whether exchange rates adjust in accordance with interest rate differentials, as predicted by Uncovered Interest Rate Parity (UIP). UIP posits that a currency with a higher nominal interest rate should depreciate against a lower-interest currency by an amount that offsets the interest advantage, leaving no systematic excess return from cross-border investments. Despite its theoretical elegance, decades of empirical research have overwhelmingly rejected UIP at short horizons, a result known as the forward premium puzzle. High-interest-rate currencies have historically tended to appreciate rather than depreciate in the short run, generating profitable carry-trade opportunities that appear to contradict the parity condition.

However, a growing body of evidence suggests that UIP may hold over longer horizons. Seminal work by Chinn and Meredith (2004) demonstrated that multi-year interest differentials predict exchange rate movements in a manner consistent with UIP, and more recent studies employing Durbin regression techniques (Baillie, Kapetanios & Kim, 2026) have confirmed that the apparent UIP failure at short horizons may be an artefact of econometric issues, specifically, the serial correlation induced by overlapping observations, rather than a true economic violation.

This thesis investigates the validity of UIP over long horizons and its implications for the use of foreign exchange (FX) forward contracts as hedging instruments. Specifically, I ask: do FX forward contracts provide an effective, unbiased hedge for long-run currency risk? If UIP holds in the long run, then the forward rate is an unbiased predictor of the future spot rate, and a strategy of systematically rolling forward contracts should, on average, yield the same outcome as remaining unhedged. If UIP fails, forward hedges may embed systematic costs or missed opportunities.

To address this question, I employ daily data for six major USD currency pairs (AUD, CAD, GBP, NZD, CHF, and JPY) from January 1990 to November 2025, using both 3-month and 6-month forward contracts sourced from Bloomberg. I apply two complementary econometric frameworks: the conventional Fama (1984) regression and the recently developed Durbin regression approach. The Fama regression establishes the well-known short-run failure of UIP, while the Durbin regression accounts for the serial correlation inherent in overlapping-horizon data to recover the long-run relationship between forward premiums and exchange rate changes.

The remainder of this thesis is organized as follows. Section B provides a chronological literature review tracing the evolution of UIP testing from the early rejections of the 1980s through recent advances. Section C lays out the theoretical foundation and econometric methodology, including the Durbin regression framework and the derivation of the long-run UIP coefficient. Section D describes the data. Section E presents descriptive statistics and preliminary analysis. Section F details the lag selection procedure. Section G reports the main empirical results and inferences. Section H discusses robustness, while Section I draws out practical implications. Section J concludes and identifies limitations and avenues for future research. Section K covers citations and respective tables/figures.

Section B. Chronological Literature Review

Uncovered Interest Rate Parity (UIP) is formally defined as the condition under which the expected change in the spot exchange rate between two currencies equals the nominal interest rate differential between those two countries, assuming risk-neutral and rational investors. In practical terms, Uncovered Interest Parity (UIP) confirms that the expected change in exchange rates between two currencies matches their difference of interest rates, assuming rational expectations and no risk premium. An investor should not earn excess return from borrowing in a low-interest-rate country and investing in a high-interest-rate economy, since the currency of high-interest country is expected to depreciate. The International Fisher Effect (IFE) is a related concept, expecting that currencies with higher nominal interest rates will depreciate in the long run by the interest differential, roughly. Under conditions of no arbitrage, UIP implies that the forward exchange rate, that is, the rate at which two parties agree today to exchange currencies at a future date (the core nature of forward contract), should be an unbiased predictor of the future spot rate. If UIP holds, a foreign exchange (FX) forward contract provides a hedge that on average yields the same outcome as an unhedged position, as the forward price already reflects the expected depreciation. If UIP fails, however, it suggests the presence of a risk premium or systematic forecast error, meaning forward rates may systematically overshoot or undershoot subsequent spot rates. This has critical implications for hedging, defined in this context as the practice by which international investors, multinational corporations, and portfolio managers use FX forward contracts to lock in a future exchange rate, thereby eliminating (or reducing) the uncertainty of converting foreign-currency cash flows back into their domestic currency. While hedgers seek to reduce risk, speculators in the FX forward market take the opposite side of these contracts, bearing currency risk in pursuit of profit. Speculators who exploit UIP deviations through carry trades borrowing in low-interest-rate currencies and investing in high-interest-

rate currencies, provide the liquidity that makes hedging possible, but their activity also contributes to the short-run deviations from UIP documented in the literature. A failure of UIP could mean that consistently hedging currency exposure with forwards might incur a cost or missed opportunity relative to staying unhedged (as seen in the “forward premium puzzle”).

I organize this review into four periods: (1) early tests and rejections in the 1980s–1990s, (2) evidence at long horizons in the early 2000s, (3) impacts of the 2008 global financial crisis and structural changes in the 2010s, and (4) recent advances in the 2020s using new method (for e.g. Durbin regressions) and data, with a focus on what they imply for using FX forward contracts to hedge long-run currency risk.

1. Early tests and rejections in 1980s - 1990s

Academic curiosity in UIP surged in the late 1970s and 1980s, leading to a series of empirical tests at short horizons (typically 1 to 3 months). These early studies almost homogeneously rejected UIP. A representative approach is the (Fama, 1984) regression, which regresses the spot exchange rate on the interest differential. Fama found that the slope coefficient was significantly negative. Practically, this meant high-interest-rate currencies managed to appreciate rather than depreciate, yielding gainful carry trades. (Hansen & Hodrick, 1980) offered one of the initial formal tests of UIP using multi-period forward predictions, and they too recognized systematic forecast errors; they presented the practice of using OLS with heteroskedasticity-and-autocorrelation-consistent (HAC) corrections to deal with overlapping data in long-horizon regressions. Despite these adjustments, the forward rate continued to look like a biased predictor of future spot rates in the data. The immediate implication was that forward contracts did not perform as theory expected in the short run – a forward premium often signaled an eventual appreciation of the high-rate currency, rather than a depreciation, undermining the forward’s value as a hedging instrument for near-term exposures.

By the mid-1980s, many studies, for example e.g., (Hakkio’, 1981), (R. T. Baillie et al., 1983), and (Levy & Nobay, 1986) among others applied more sophisticated methods (treating exchange rate changes and interest differentials jointly) and rejected UIP strongly. These rejections were mainly attributed to one of two factors: (a) time-varying risk premium in the FX market, and/or (b) deviations from rational expectations. In other words, investors wanted a risk premium for holding high-interest currencies (so the forward rate was influenced predictor of future spot rates), or investors mis-anticipated future exchange rates. By 1990, the consensus view was that UIP failed at short horizons, which by extension meant a

short-run failure of the IFE, currencies did not depreciate enough to offset interest differentials, leading to carry trade profits.

Overall, this period firmly established the baseline result i.e., short-term UIP does not hold, and short-term forward rates are biased predictors of future exchange rates, casting doubt on FX market efficiency. The challenge going forward was explaining why and discovering if UIP might fare better under different conditions (different horizons, currencies, or sample periods). Two wide-ranging explanatory paths appeared: risk-based explanations (carry trade risk premia) and the possibility that UIP might hold in the “long run” even if it fails at short horizons. The latter path set the point for research in the late 1990s and 2000s focusing on longer-horizon tests and the long-run IFE.

2. Evolution to Long-Horizon Studies and Partial Validation (2000s)

After two decades of short-horizon rejections, researchers in the late 1990s and early 2000s revisited UIP by looking at longer forecast horizons. The motivation was that short-run exchange rate movements are dominated by noise, speculative flows, and policy shocks, which might decouple them from interest differentials. Over longer horizons, however, cumulative interest differentials might eventually be offset by exchange rate movements (as implied by IFE), especially if price levels and fundamentals adjust. Early evidence in this direction was provided by [\(Alexius, 2001\)](#), who examined long-term bond yields for 14 countries and found UIP held better at long horizons than at the 1-month horizon. The findings essentially suggested that when looking at multi-year periods (and using long-term bond yields as proxies for expected future short rates), the forward premium puzzle diminished or disappeared

A landmark contribution came from [\(Chinn & Meredith, 2004\)](#) who explicitly tested UIP using long-maturity bond yields (for the G7 countries) instead of the usual short-term rates. In contrast to the nearly universal short-horizon rejections, these long-horizon regressions yielded support for UIP. Chinn and Meredith reported that all the slope coefficients on long-term interest differentials had the correct positive sign, and most were much closer to 1 (the UIP-expected value) than to 0. In many cases, one could not reject the hypothesis that the slope = 1 when using 5-year or 10-year interest rate differentials, suggesting that a high-interest-rate country’s currency tended to depreciate in line with UIP predictions over multi-year periods. This finding was a stark departure from the [\(Fama, 1984\)](#) result for 1-month horizons, and it provided partial rehabilitation for UIP. It suggested that while short-run markets deviate from parity (perhaps due to risk premia or short-term inefficiencies), over longer horizons exchange rates realign with fundamentals.

Several contemporaneous studies reinforced or refined this conclusion. For example, (Bekaert et al., 2007) took a comprehensive look at UIP across horizons and in conjunction with the expectations hypothesis of the term structure. Using data on both short- and long-term interest rates, they found that different currency pairs showed different behaviors (some adhering closer to UIP than others), and while UIP was not perfect at any horizon, the degree of failure at long horizons was modest. In some cases, the forward premium puzzle almost vanished when using multi-year horizons.

(Snaith et al., 2013), conducted the multi-horizon study, spanning 1-month up to 10-year forwards for major USD currency pairs (AUD, CAD, CHF, EUR, GBP) over 1980–2006. They took care to address econometric issues in long-horizon regressions by using a novel bootstrapped HAC estimator for inference. Their findings confirmed a clear “horizon effect”: as the forward contract maturity is extended, the Fama regression slope coefficient rises toward unity. In fact, by the 3-year horizon, the forward premium puzzle essentially disappears, the slope coefficients were not statistically different from 1 (often around 0.8 to 1.2) and the usual negative bias was gone. The results imply that while a 1-month or 1-quarter investment in a high-yield currency tended to earn excess returns, a 3+ year investment did not, as the currency depreciation over that period roughly equaled the interest gain.

The perception by the mid-2000s was that UIP fares better at long horizons. In practical terms, this means that if an investor engaged in a multi-year carry trade (e.g. borrowing in USD and investing in GBP for 5 years) or hedges a long-term foreign investment with forward contracts, the expected exchange rate movement might roughly balance out the interest advantage, consistent with the IFE, whereas one-month or quarterly trades were still yielding systematic profits in the data. This offered a more optimistic outlook for the effectiveness of forward contracts such as hedging tools over longer horizons. The challenge was then to understand what changed over long horizons (were risk premia decaying? were expectations more accurate?), and to confirm these patterns with expanding datasets and improved econometric techniques.

3. The Impact of Financial Crises, Regime Shifts, and Risk (Post-2008)

The global financial crisis of 2008–2009 and its aftermath provided a natural experiment to observe UIP under extreme conditions. The crisis led to unprecedented monetary easing, flight-to-safety capital flows, and occasional unwinding of carry trades.

(Bussiere et al., 2018) document a structural break in the UIP regression around the time of the Global Financial Crisis. They find that while the joint hypothesis of UIP and rational expectations is still often rejected, it is rejected “with much less frequency” in the post-2008 era. In fact, in the period after 2008,

the Fama regression slope coefficients turned positive and large, closer to +1 or even above, in contrast to the significantly negative coefficients typically observed pre-2008. In other words, during 2009–2015, a high-interest-rate currency tended (if anything) to depreciate, as UIP predicts, rather than appreciate. (Bussiere et al., 2018) refer to this as the “New Fama Puzzle” and indicated a regime shift in FX market post 2008 crisis. This suggests that part of the post-crisis shift was due to changes in how expectation errors or risk premia correlate with interest rates. They conclude that the reversal of the puzzle post-2008 is linked to a change in the correlation structure: after the crisis, high-interest currencies no longer provided positive unexpected returns, consistent with perhaps a change in investor behavior or risk pricing.

One contributing factor to this shift is the zero-lower-bound (ZLB) environment that prevailed in the US, EU, and Japan for many years after 2008. A recent study by (Engel et al., 2022) takes a deep look at UIP for the U.S. dollar in this context. They examine short, medium, and long horizons for USD exchange rates and find that the predictive power of interest rate differentials is not stable over time and disappears altogether when interest rates are near the zero lower bound. Essentially, when all countries have near-zero rates, interest differences carry little information and UIP can’t be meaningfully tested. However, (Engel et al., 2022) make a fascinating observation: year-on-year inflation differentials did predict excess currency returns at a time when interest rates were uninformative. Specifically, when U.S. inflation was higher than foreign inflation, subsequent returns on USD deposits tended to be higher (i.e. the USD did not depreciate as much as one might expect). In other words, if the U.S. has higher inflation, investors might anticipate the Fed will eventually raise rates, but initially they underreact and do not bid up the dollar enough, leading the dollar to earn high returns until it adjusts. This finding links to the IFE concept: in a world of inflation targeting, inflation differentials become a key driver of expected exchange rate changes when nominal rates are constrained.

Another critical development in the post-2008 era was a deeper investigation into risk-based explanations for UIP deviations. The carry trade was extremely popular pre-2008, and its sudden unwinding during the crisis gave credence to the idea that excess returns to carry are compensation for risk. (Brunnermeier et al., 2008) documented that carry trades are subject to “crash risk”, high-interest currencies can plunge sharply in value during crises or flights to safety, erasing years of carry profits in a short time. In their framework, speculators require a premium to bear the risk of rare but extreme depreciations of high-yield currencies, which results in the forward rate being a biased predictor of the future spot rate. This line of

reasoning connects UIP failure to the broader asset pricing literature. Indeed, several studies in the 2010s demonstrated that currency excess returns can be linked to risk factors: for example, (Bansal & Shaliastovich, 2012) proposed a long-run risk model that could explain predictability puzzles in currency and bond markets.

Of relevance is the work by (Lustig et al., 2019), who examined the term structure of currency carry trade risk premia. They found that the excess return from carry trade is a declining function of the investment. In practical terms, if an investor consistently rolls over 1-month investments in high-interest-rate bonds, they earn a substantial risk premium. But if the investor instead engages in a longer-term strategy using, say, 5-year or 10-year bonds (thus committing to the trade for a longer horizon), the average excess return is much smaller. This result dovetails with the earlier findings that long-horizon UIP holds better: it suggests that the risk premium that cause UIP to fail are concentrated at short horizons. At longer horizons, either hedging costs or offsetting local bond premiums erase the carry advantage. Their no-arbitrage condition shows that if long-term bonds carry higher term premiums in the high-interest country, it can offset the currency premium, resulting in UIP holding for long-term rates.

During normal times with modest interest differentials, carry trade strategies thrived and UIP failures were evident; but during low-rate regimes, these patterns changed. Risk considerations (crash risk, liquidity constraints) help explain why arbitrage does not force UIP to hold at short horizons, yet the accumulation of macroeconomic fundamentals over time helps restore UIP in the long run. For someone looking to hedge currency exposure over, say, a 5-year investment horizon, the evidence implies that forward contracts may perform as intended, delivering roughly the exchange rate one would expect based on interest/inflation fundamentals, with any risk premium effects washing out over time. This naturally leads to the question: can I more rigorously test long-run UIP with modern data and methods?

4. Recent advances in the 2020s using Durbin regressions and data

The 2020s have seen significant progress in empirical methodology for testing UIP/IFE, as well as the use of new data sources. One breakthrough is the development of the Durbin regression approach for predictive regressions, which has been applied to the UIP question by Baillie, Diebold, Kapetanios, and Kim.

(Baillie, Diebold, Kapetanios & Kim 2023) and (Baillie, Kim, Diebold, Kapetanios, & Mora 2024) suggest a new single-equation test for UIP based on a Durbin regression that explicitly models the autocorrelation in regression. In essence, this approach augments the standard UIP regression with lags of

the dependent variable to capture the dynamics. The result is a consistent and asymptotically efficient estimator for the slope coefficient under UIP, which does not rely on the assumption of strict exogeneity and avoids the potential inconsistency of OLS in this context. Applying this methodology, (Baillie, Diebold, Kapetanios & Kim 2023) and (Baillie, Kim, Diebold, Kapetanios, & Mora 2024) found differences in the estimated UIP relationship. In their re-examination of the classic 1-month horizon Fama regression (using modern data and multiple currencies), the slope coefficients were all positive and remarkably stable across currencies, and larger than those found in previous studies. Whereas traditional OLS estimates often found negative slopes, the dynamic regression approach yielded slope estimates much closer to zero or even slightly positive for short horizons and moving toward the theoretical value of +1 for longer overlapping horizons.

(Baillie, Kapetanios & Kim 2025) further demonstrate the versatility of Durbin regressions by extending the analysis to cases where the regression errors may exhibit long-memory or nonlinear pattern. Their findings suggest that the Durbin approach remains robust in those settings, whereas the conventional OLS-HAC inference cannot perform well. In simulations and empirical comparisons, the Durbin estimator consistently recovers the true UIP slope even when errors are highly persistent, while HAC-based OLS can severely underestimate the slope. These advances imply that many earlier UIP tests might have been hampered by statistical artifacts, which the new methods can correct.

Armed with these improved techniques, researchers have revisited the question of whether UIP holds in the long run with fresh eyes in their paper (Baillie, Kapetanios & Kim 2026). Their study used weekly data for several major currencies and considered forward contracts that were constantly rolled over (creating synthetic long-horizon returns) to simulate long-horizon investment. Crucially, by using 30-day forward markets at high frequency, they avoid the complications of using actual long-term bonds (which may introduce credit risk or illiquidity) while still capturing long-run information by aggregating the weekly data over many months. The findings are striking: Under conventional OLS-HAC, they replicate the usual UIP rejection at short horizons. But using the Durbin regression with robust inference, they fail to reject UIP over long horizons. In their words, “UIP is valid over long horizons” according to the Durbin test, and the results provide “important confirmatory evidence for the long-run international Fisher condition”. This directly supports the notion that the IFE is alive and well in the long run. The paper’s highlights emphasize that while traditional methods would reject UIP, the improved approach finds no rejection, aligning with theory. In practical terms, (Baillie, Kapetanios & Kim 2026) maintain the idea that if you hold a currency position over many months or years, any interest rate advantage will be offset by

expected depreciation, leaving no arbitrage profit. It also vindicates the use of FX forwards as hedging tools for long-term exposure: if UIP holds, the forward rate is an unbiased forecast of the future spot, meaning a hedger who locks in the forward essentially receives the fair expected value and avoids any systematic cost or benefit.

In summary, the frontier of research suggests a more optimistic outlook for UIP in the long run. With advanced techniques like Durbin regressions providing robust inference, I have a view that much of the apparent failure of UIP was due to short-run noise, bias, and risk premium effects. When those are accounted for, UIP does appear to hold over long horizons, meaning the exchange rate eventually “corrects” for any interest differential.

Section C. Theoretical Foundation, and Methodology

1. The UIP, IFE, and Forward Rate Framework

Uncovered Interest Parity (UIP) suggests that the variance in nominal interest rates between two countries must equal the expected difference in the exchange rate between those countries' currencies. UIP implies:

- Currencies that have higher interest rates are expected to depreciate in the future, offsetting the advantage of interest
- Conversely, currencies with lower interest rates should appreciate, compensating for the lower yield.

The International Fisher Effect (IFE) is a closely linked concept, asserting that differences in nominal interest rates reflect potential inflation differentials, which in turn predict currency depreciation/appreciation. In fact, the IFE leads me to the same testable prediction as UIP i.e., in the long run, interest rate differentials (or forward premia) should unbiasedly explain exchange rate movements. Under these contexts, the forward exchange rate plays a central role. Covered Interest Parity (CIP) is the no-arbitrage condition stating that the forward premium on a currency equals the interest rate differential between the two countries, because any discrepancy would be exploited through riskless arbitrage involving simultaneous borrowing, investing, and forward contracting. Unlike UIP, CIP is an arbitrage-enforced relationship that holds consistently in practice (barring credit risk or capital controls). If Covered Interest Parity holds, the forward premium is approximately equal to the interest rate differential between the two countries. Thus, UIP/IFE can be tested by examining whether the forward exchange rate correctly predicts future spot rates. A forward rate, consistent with UIP, would mean that using forward contracts to hedge currency risk yields no systematic cost or profit. In other words, the real cost of hedging would be near zero in the long run. This underlines the practical importance of my tests: confirming UIP not only corroborates a core theory of international finance but also implies that forward contracts effectively neutralize exchange rate risk for long-term hedging. The relationship is mentioned as:

$$E[s_{t+k}] - s_t = i_t - i_t^* \approx f_t - s_t \quad (1)$$

Where:

s_t : natural log of the spot exchange rate at time t

f_t : natural log of the forward exchange rate at time t for delivery at time $t + k$

i_t : domestic (US) interest rate

i_t^* : foreign interest rate

$E[s_{t+k}]$: expectation conditional on information at time t

To explain in other words, if US interest rates are higher than foreign rates, the US dollar should be expected to depreciate. Or else, investors could earn risk-free profits by borrowing in the low-interest currency and investing in the high-interest regime (carry trade).

2. Forward Premium Puzzle

The Forward Premium Puzzle (Fama, 1984) refers to the observed finding that:

- The forward premium $f_t - s_t$ often predicts exchange rate movements in the wrong direction.
- When regressing spot turns on forward premiums, the slope coefficient β is either negative or significantly different from 1.

The standard Fama Regression is:

$$s_{t+k} - s_t = \alpha + \beta(f_t - s_t) + u_{t+k} \quad (2)$$

UIP hypothesis: $H_0: \alpha = 0, \beta = 1$

Empirical findings constantly show $\beta < 1$ and often $\beta < 0$, which have been called one of the robust findings in international finance.

Why might UIP fail in short run?

- *Time-Varying Risk Premia*: Investors may expect a risk premium for holding foreign currency, causing forward rates to deviate from subsequent spot changes. If this risk premium changes over time, it can lead to the observed bias in short-run UIP tests.
- *Peso Problems*: Rare but unusual events (e.g. a large devaluation that is expected but hasn't occurred in the sample) can alter expectations and UIP tests. These low-probability events (called "peso problems") mean the sample may not be a true reflection of long-run relationship.
- *Market Inefficiencies*: Behavioral biases or limits to arbitrage might allow deviations from UIP in the short run. For example, traders might constantly underreact / overreact to information, or capital controls and transaction costs might make it difficult to arbitrage.
- *Econometric Issues*: Overlapping data and other statistical problems can induce biases. For example, using high-frequency data to test multi-month horizons creates serial correlation in errors, or measurement errors can lead to spurious rejections of UIP.

3. Durbin Regression

With this backdrop, recent research by (Baillie, Kapetanios & Kim 2026) has revisited UIP testing with improved econometric methods. They maintain that the apparent failure of UIP in many studies may be due to econometric issues rather than a true economic violation. In particular, the use of overlapping forward dates (e.g. using daily data to examine a 3-month horizon) brings strong autocorrelation in the regression residuals, invalidating standard OLS inference. Traditional tests (such as Fama regression of future exchange rate change on the forward premium) might therefore wrongly reject UIP because of influenced standard errors and inefficient estimates, rather than because UIP truly fails in the population.

This study takes inspiration from the work of (Baillie, Kapetanios & Kim 2026) who uses Durbin regression approach, a dynamic time-series regression that comprises lagged dependent and independent variables, to capture the persistence in the data and to distinguish short-run versus long-run effects. Using this approach, they find that UIP does hold in the long run even if it appears to fail in the short run, thus solving the puzzle by attributing it to short-run noise or risk premia that average out over time.

The purpose of this study is to build on that insight and rigorously test UIP on an extensive dataset and with longer horizons. The model can be written as:

Durbin regressions:

$$s_{t+k} - s_t = \alpha_0 + \sum_{j=1}^p \alpha_j (s_{t+k-j} - s_{t-j}) + \sum_{j=0}^p \beta_j (f_{t-j} - s_{t-j}) + \varepsilon_{t+k} \quad (3)$$

where I include p lags of the dependent variable (past spot changes) and p lags of the independent variable (past forward premiums) with error term ε_{t+k} . This autoregressive distributed lag (ARDL) model is proposed to address the econometric issues in the standard UIP test. Including lagged terms helps absorb the serial correlation arising from overlapping horizons (e.g. a 3-month spot return computed daily means today's 3M return shares 63 days in common with tomorrow's 3M return, inducing strong autocorrelation). Including lagged forward premiums allows for delayed effects of the forward premium on the exchange rate, acknowledging that the relationship may play out over several periods rather than immediately.

The Durbin regression approach follows the recommendation of the authors in response to their key insight: *"The use of overlapping observations in UIP regressions creates strong serial correlation in the error term, which biases standard errors and leads to spurious rejections of UIP."* By explicitly incorporating the dynamics (lags), I mitigate this problem and obtain more dependable inference. In

principle, the dynamic model separates short-run noise from long-run signal. Short-run deviations can cause β to deviate from 1, but the cumulative long-run effect might average out to unity if UIP holds in the long run.

a. Long-Run Coefficient (γ)

An important outcome of the dynamic regression is the computation of the long-run UIP coefficient (γ). This is the long-run multiplier that tells me the total effect of a permanent change in the forward premium on the cumulative change in the spot rate. In the above model, γ is given by:

$$\gamma = \frac{\sum_{j=0}^p \beta_j}{1 - \sum_{j=1}^p \alpha_j} = \frac{\beta(1)}{\alpha(1)} \quad (4)$$

The “long run” in this context is defined empirically as the cumulative horizon over which all lagged effects in the Durbin regression have fully played out. Given the use of 3-month and 6-month forward contracts with daily overlapping observations spanning 35 years, the long-run coefficient captures the total exchange rate adjustment that unfolds as short-run noise dissipates. By instinct, if the forward premium were to permanently increase by 1, γ is the predicted total change in the spot rate over the long run (after all lagged effects have played out). Under UIP, I expect $\gamma = 1$. This long-run coefficient γ is my focus for testing UIP in the long run. I will test below hypothesis:

$$\text{UIP hypothesis: } H_0: \gamma = 1$$

In simpler terms, I check whether γ is statistically different from 1. A failure to reject the null hypothesis would indicate that, despite any short-run anomalies, the long-run relationship between the forward premium and exchange rate changes aligns with UIP.

In the dynamic model, I also obtain a *contemporaneous* slope β (the coefficient on the current forward premium, comparable to the β in the Fama regression) and possibly lagged coefficients $\beta_1, \beta_2, \beta_3, \dots, \beta_{j=p}$ captures the immediate response of the exchange rate to the forward premium within the horizon, while the sum of β yields the total response over time. I anticipate β may differ substantially from 1 because it captures only short-run effects; it would not be surprising if $\beta \neq 1$ due to short-term risk premia or slow reaction of exchange rates. The hypothesis of UIP in the long run requires $\gamma = 1$. This distinction is important: a rejection of $\beta \neq 1$ in the dynamic model would support the forward premium puzzle in the short run, but a failure to reject $\gamma = 1$ would indicate UIP holds when short-run noise is accounted for.

Section D. Data Description

Data Sources and Sample

My analysis covers six major currencies (all priced against the U.S. dollar as the base currency). I focus on two maturities for forward contracts: the 3-month forward and the 6-month forward for each currency. These horizons are potentially long enough to see UIP, with adequate number of observations. Daily frequency is used to increase the effective sample size for constructing overlapping multi-month returns and yield more precise parameter estimates in the Durbin regression framework (Table 1).

All exchange rate series are converted to natural logarithms before analysis to linearize returns and yield symmetrical percentage changes. From these, I construct the key variables:

- **Spot Return $\Delta s_{(t,k)}$** : The realized log return from holding the foreign currency for k periods. I compute this as $\Delta s_{(t,k)} = s_{t+k} - s_t$. For the 3-month horizon, $k = 63$ trading days, and for 6-month, $k = 126$ trading days.
- **Forward Premium $f_{t+k} - s_t$** : The forward premium is the variance between the log forward rate and the current log spot rate. This term is essentially the market-implied expected depreciation of the foreign currency with risk-neutral pricing: in theory $f_{t+k} - s_t \approx i_t - i_t^*$ (interest rate differential). A positive forward premium means the forward price of foreign currency is higher than the current spot (often because the foreign interest rate is higher than U.S. interest rate), indicating the foreign currency is expected to depreciate (lose value) relative to USD. A negative forward premium means the foreign currency is at a forward discount, implying the foreign currency is expected to appreciate relative to USD.

Due to the forward-looking characteristics of spot returns, observations at the end of the sample stand as outlined in Table 2.

Section E. Analysis of Data and Preliminary Findings

1. Summary Statistics – 3M returns/premiums

a. Descriptive Statistics

Per

Table 3, the 3-month spot return series for all six USD pairs exhibit mean returns that are effectively zero. Indeed, each currency's average spot return is on the order of 0.1% or less, with values ranging from about +0.15% (GBP) to -0.45% (CHF). Such near-zero mean returns are consistent with market efficiency, indicating no clear systematic drift in exchange rates over time.

High interest “carry” currencies, AUD, NZD, GBP, and CHF, have positive average forward premiums (e.g. about +0.38% per quarter for AUD, +0.51% for NZD, +0.59% for CHF), reflecting that foreign interest rates were generally higher than U.S. rates. Notably, these differentials in interests are not mirrored by equivalent opposite-signed mean spot returns, for example, NZD offered a sizable positive forward premium, yet its currency saw essentially zero average appreciation or depreciation, a potential sign of UIP deviations (the forward premium puzzle).

In terms of risk, spot returns are far more volatile than forward premiums. The standard deviation of quarterly spot returns ranges from about 3.6% (CAD, the least volatile) to ~5.5% (AUD, NZD, JPY). By contrast, the forward premium series are relatively steadier, with standard deviations around 0.45–0.55% for most currencies, and only CHF's forward premium volatility is somewhat higher (~0.88%). These figures suggest that the exchange rate contributes the dominant share of uncertainty in the currency investment return. Overall, the stats indicate all six currencies had minimal average excess returns with marked differences in risk. This aligns with the aspect that higher-yield currencies tend to be more volatile, theoretically compensating investors with risk for their interest advantage.

b. Distributions

All spot return distributions exhibit clear non-normal features, with significant skewness and kurtosis across currencies. The returns display fat tails much thicker than the normal distribution, which is evidenced by excess kurtosis values. AUD stands out with the highest kurtosis (~5.73)

Table 3, indicating a heavy-tailed return distribution. This implies that AUD/USD experienced more frequent extreme movements (large outliers) compared to other currencies, pointing to substantial tail risk. GBP also shows pronounced excess kurtosis (~3.98), while CAD and JPY have more moderate tail heaviness (kurtosis ~3.1 and 0.94, respectively). Notably, CHF's spot returns have the lowest kurtosis

(~ 0.68), suggesting the fewest extreme outliers in the sample. These patterns are consistent with [Figure 1](#), which reveal sharper peaks and fatter tails than a normal curve for most currencies.

Skewness varies across the currencies, indicating asymmetric return regimes. Most of the spot return distributions are positively skewed, meaning they have a propensity for occasional large positive returns. AUD and GBP are skewed to the right (skew $\approx +0.89$ and $+0.92$), signifying episodes where these currencies saw abrupt appreciations against the USD. By contrast, JPY exhibits a slight negative skew (-0.30), the only currency in the set with a negative-skewed return distribution. CHF's distribution is nearly symmetric (skew ~ 0.04 , essentially zero), indicating balanced risks. In summary, the distributional shape metrics highlight that high-yield, risk-sensitive currencies (AUD, NZD, GBP) have more pronounced tail risks and skewness, whereas low-yield currencies (JPY, CHF) have more symmetric or left-skewed distributions, showing different risk features.

c. Time-series assessment

Over time, the 3-month spot returns show noticeable volatility clustering. Periods of relative calm in exchange rate movements are punctuated by episodes of heightened volatility, often coinciding with global financial stress. For instance, all currencies experienced spikes in volatility during the 2008–2009 financial crisis and again during the 2020 COVID-19 shock. In these crisis periods, spot returns swung dramatically, then rebounded, whereas in quiet periods, exchange rates fluctuated in a more contained band around their long-run equilibrium. This clustering implies that exchange rate risk is time-varying.

In stark contrast to the oscillating path of spot returns, the 3-month forward premium series appear remarkably stable over the same sample horizon. Plotting the two together ([Figure 2](#)), the forward premium traces a gentle curve, almost flat line (usually staying within a few basis points of zero). This comparison highlights a key aspect of the forward premium puzzle: there is no explicit correspondence between the level of the forward premium and corresponding exchange rate changes. Periods when the interest differential is wide, do not reliably precede any trend in the spot rate. For example, even when the USD-JPY forward premium was persistently negative (reflecting Japan's near-zero rates vs higher US rates), there was no continued drift in USD/JPY; instead, the yen's value would often respond to unrelated risk-off events. A lack of predictive co-movement is apparent: high forward premiums did not foreshadow systematic depreciations of the foreign currency, nor did negative premiums precede appreciations, on a timeline of days or months. In other words, the short-term currency swings are driven by shocks and

market sentiment rather than slowly evolving interest differentials. These time-series dynamics suggest that any UIP-implied relation is obscured by dominant noise and volatility in exchange rates.

2. Summary Statistics – 6M returns/premiums

a. Descriptive Statistics

Table 4, the summary statistics for 6-month spot returns and forward premiums reveal several key features. Across six major USD currency pairs, the average half-year spot returns are all very close to zero, especially in comparison to their volatility, indicating that on a 6-month basis, exchange rates have no strong drift, consistent with the near-random-walk trend. For example, the mean log return for GBP/USD is about +0.32% per 6-month period, and for AUD/USD and CAD/USD around +0.24%, whereas JPY/USD and NZD/USD average essentially 0%, and CHF/USD has a slightly negative mean around –0.85%. In other words, typical 6-month exchange rate moves are on the order of several percentage points, shadowing the negligible average appreciation or depreciation over the long run.

The forward premium varies more markedly across currencies. High-interest-rate currencies like AUD and NZD exhibit positive forward premiums on average, whereas low-interest-rate currencies such as JPY show negative forward premiums (JPY's is about –1.25%). CHF also shows a positive average forward premium (~+1.21%), reflecting periods where U.S. interest rates were lower than Swiss rates. CAD's forward premium is close to zero on average (slightly negative at –0.08%), consistent with Canada's interest rates being often in line with U.S. rates. Standard deviations of forward premiums are smaller than those of spot returns, indicating that interest differentials evolve much more smoothly than exchange rates. Largely, the forward premium series are highly persistent and slowly varying, unlike spot returns.

Comparing spot returns and forward premiums side by side highlights the classic UIP deviation. If UIP held, a currency with a positive 6-month forward premium should, on average, depreciate by about the same magnitude over 6 months, yielding a zero net gain. My data do not show this: for instance, AUD and NZD offered significant interest advantages, yet their average 6-month returns are slightly positive or near zero instead of the sizeable negative values UIP would predict. Conversely, JPY, consistently a low-interest currency with a –1.25% half-year forward premium, did *not* appreciate by +1.25% on average; its mean return is essentially zero. This indicates that the high-interest currencies have earned positive excess returns, whereas low-interest currencies have faced losses on average, violating UIP. These summary statistics already hint at the forward premium puzzle: currencies at a forward discount (high interest) have not depreciated as expected, and those at a forward premium (low interest) have often depreciated when

they should have appreciated under UIP theory. In the sections below, I go deeper into the distributional and time-series properties of these returns to further understand this deviation and its risk characteristics.

b. Distributions

Examining the distribution of 6-month spot returns (Figure 3) confirms that returns are not perfectly normally distributed, and there are obvious differences across currencies. The histograms of most currencies are roughly bell-shaped and centered near zero, but their tails and asymmetry vary. Several currencies exhibit positive skewness. For instance, GBP/USD has a pronounced right skew (skew $\approx +1.07$), having few instances of exceptionally large positive returns. Indeed, GBP's maximum half-year log return ($\sim +36\%$) far exceeds its largest drop ($\sim -20\%$), indicating occasional sharp appreciations of GBP against the USD. AUD and NZD also show relative right skew (skew $\approx +0.58$ and $+0.44$, respectively), meaning these currencies had some instances of outsized gains relative to their losses. This tells that while crashes do occur for these high-yield currencies, there have also been periods of dramatic rebounds or appreciations. By contrast, JPY/USD has essentially zero skew (≈ 0.00), and CHF/USD's skew is mild ($+0.19$), indicating their return distributions are more symmetric with alike probabilities of large ups and downs over 6-month periods.

The tail heaviness of the return distributions varies widely. AUD/USD and GBP/USD returns are leptokurtic with kurtosis ~ 3.6 and 3.3 , respectively, slightly above the normal scale. This implies a sharper peak and fatter tails than normal: an elevated likelihood of extreme 6-months moves in these currency pairs. Notably, CAD/USD and NZD/USD have kurtosis around 2.3 – 2.23 , demonstrating *flatter* distributions with thinner tails than a normal. Even more striking are CHF/USD and JPY/USD, with kurtosis as low as ~ 0.45 and 0.29 . Such low kurtosis informs me that their return distributions are quite flat-topped and light-tailed, in contrast to what I classically see in respective daily FX returns.

On the other hand, the forward premiums are much less dispersed and closer to symmetric than spot returns, since interest rates do not usually swing as wildly as exchange rates. For most currencies, the forward premium would be narrow and approximately bell-shaped around the mean differential. In short, spot return distributions are much wider and exhibit varying degrees of skewness and tail thickness, whereas forward premium distributions are narrower / normal-like. These distributional properties underscore that exchange rate movements introduce considerable randomness and risk beyond what interest differentials alone would predict.

c. Time-series assessment

Figure 4 shows the trend of the 6-month spot returns and forward premiums. An observation similar to that of 3-month return trend is found in scale and volatility between the two series: the spot return is extremely volatile, while the forward premium is relatively smooth and trend-like. The forward premium for each currency tends to drift slowly as interest rate variances evolve through different monetary policy regimes. JPY's forward premium is persistently negative throughout most of the sample, though its magnitude varies from about -2% in the late 1990s to around -0.5% in recent years. The spot return series shows that exchange rate changes are largely unpredictable and prone to abrupt swings. The spot return jumps up and down around zero with no long-term trend, but with obvious volatility clusters.

In sum, the 6-month trend show that while UIP fails on average (interest differentials are not offset by proportional exchange rate moves), the distributional and temporal patterns of spot returns point to compensation for risk: investors earn the carry during stable periods but face occasional sharp corrections. This sets the platform for further analysis of how these properties compare to the 3-month horizon and how volatility and covariance across currencies contribute to the broader UIP puzzle.

3. Volatility comparison

A clear pattern emerges when comparing the volatility of spot exchange rate changes to that of forward premiums across all currencies and horizons (Figure 5). Spot return volatility is in magnitude higher than that of forward premium for every currency examined. For instance, at the 3-month horizon, exchange rate changes are roughly 5–12 times more volatile than the corresponding interest differentials, and this disparity persists at 6 months too. This means that exchange rates turn far more wildly than interest rate gaps, which themselves move gradually. All currencies show this behavior, i.e., the interest rate spread is relatively smooth and persistent, whereas spot rates are noisy and turbulent. In fact, the sensitivity of exchange rates to news is dramatically larger than that of interest differentials, as illustrated below.

This contrast has important implications for UIP testing. The volatility gap implies that interest differentials influence little to explain exchange rate variance. Exchange rate changes are dominated by noise (or risk premia expectations) rather than by the slow-moving interest rate advantage between countries. Moreover, specific currency patterns highlight the distinction: for example, the CHF's forward premium volatility is notably higher than others, yet still small in absolute terms. This persistent volatility gap is a key feature in understanding why UIP often fails in short run: exchange rates surround risk and

expectation shocks that overshadow the information in interest differentials, making currency returns much harder to predict despite the relatively mild behavior of forward premiums.

4. Correlation analysis of spot returns and forward premiums

The 3M and 6M correlation heatmap ([Error! Reference source not found.](#)) reveal that spot exchange rate returns and forward premiums are essentially not correlated for all examined currencies. In fact, each currency's spot return vs. forward premium correlation is very close to zero, ranging from about -0.15 to $+0.17$ (3M) and -0.26 to 0.22 (6M).

The consistency of these correlations across both 3-month and 6-month horizons indicate that the forward premium (specially 3M) contains little predictive information about short-term exchange rate movements in my sample.

From an economic perspective, these findings underscore the forward premium puzzle. Under UIP, a positive forward premium should predict a corresponding currency depreciation so that the forward rate is an unbiased predictor of the future spot rate. However, my correlation analysis shows no meaningful linear association, if anything, higher forward premiums are not reliably followed by depreciations. This aligns with extensive empirical literature documenting UIP failures: for short horizons, the slope coefficients from regressing currency returns on forward premiums are typically far below 1 and often negative.

It is important to note, however, that a zero correlation in the short run does not definitively invalidate UIP in the long run. It simply indicates that short-horizon exchange rate changes are dominated by noise or risk premia. In subsequent econometric analysis, I therefore use the Durbin regression approach to capture any long-run relationship hidden behind the short-run noise.

Section F. Lag(s) Selection Methodology and Results

Selecting the optimal number of lags in the Durbin regression framework is crucial to properly specify the model for UIP testing. The Durbin regression comprises lags of the dependent variable (spot return) to account for serial correlation arising from overlapping observations. Using intersecting multi-month return data can induce autocorrelation in the regression residuals, so including too few lags may leave this serial correlation unaddressed, biasing coefficients and standard errors. On the other hand, including too many lags wastes degrees of freedom and can complicate the model. Furthermore, the long-run UIP coefficient in the Durbin regression is a nonlinear function of all the lagged terms' coefficients, so

increasing the lag structure could make the long-run estimate less reliable. For these reasons, a systematic lag selection procedure is required to balance model fit.

I use information criteria to determine the optimal number of lags for each currency and horizon. I computed both the Bayesian Information Criterion (BIC) and the Akaike Information Criterion (AIC) for Durbin regressions with varying lags. BIC and AIC trade off goodness-of-fit against model complexity, but BIC imposes a stricter penalty for additional parameters than AIC. As a result, BIC tends to favor more economical models and is consistent. I therefore interpret BIC as the primary guide, while still noting AIC results for comparison.

Lag selection results were strongly consistent. For all six currencies and both horizons, BIC unequivocally selected a single lag ($p = 1$) as the optimal lag length (Figure 7). The AIC criterion largely agreed in most cases, also favoring one lag, but in a few instances, AIC showed a second lag might marginally improve fit (Table 5). Nonetheless, even in those cases the AIC improvement for two lags was very slight, and BIC's more rigorous penalty judged that the additional complexity was not justified. This consistency across currencies and horizons simplifies my modeling, I can adopt a single-lag Durbin regression for all currencies, confident that higher-order lags do not materially contribute to explaining exchange rate changes.

An extra lag or two does not significantly add explanatory power because the additional serial correlation from overlapping intervals is minimal once the first-order autocorrelation is considered. By avoiding overfitting, I try to enhance the robustness of my results and ensure that the test of UIP is driven by genuine economic signal rather than products of model complexity.

Section G. Testing and Inferences

The conventional Fama regression results (

Table 6) confirm that UIP fails decisively at short horizons for several major currencies. Across six USD currency pairs and 3M/6M horizons, the estimated slope coefficients show substantial deviations from the UIP-predicted value of 1.

In fact, half of the currencies exhibit a pronounced “slope bias”: the AUD and JPY have negative β estimates (around -0.8) that are highly significant. Even where the slope is positive, it is far below unity, for instance, the CAD 3M β is only ~ 0.15 , suggesting that only a fraction of the interest gap is offset by exchange rate changes. No short-run regression comes close to $\beta = 1$, and in most cases, I can reject the UIP hypothesis at the 5% level. Consistent with earlier correlation analysis, these Fama regressions have near-zero explanatory power (R^2 on the order of 0.1%–0.5%), underscoring that interest differentials explain virtually none of the day-to-day or quarter-to-quarter FX fluctuations. Short-term exchange rate movements are evidently dominated by noise and expectations of risk premia.

Not all currencies demonstrate the same degree of UIP deviation. AUD and JPY show the most extreme bias (β significantly below 0, opposite to UIP) in my sample. This associates with their risk profiles: these currencies had some of the highest return volatility and more pronounced tail risks (e.g. AUD’s return distribution has exceptionally heavy tails, and JPY exhibits notable left-skew). Investors ask for excess returns to hold volatile currencies, resulting in forward rates that methodically “overshoot” the actual depreciation realized. By contrast, currencies with lower interest differentials (GBP, CHF, NZD) showed slopes closer to UIP expectations (e.g. β around 0.5–0.7 for GBP) and were statistically indistinguishable from 1 in my tests. In these cases, I “fail to reject” UIP in the conventional regression, suggesting that any deviation is too noisy to confirm. It is said that the CHF and GBP did not exhibit a strong forward premium puzzle. In short, the short-term evidence of UIP failure is currency-specific: pairs characterized by continued interest-rate gaps tend to show the strongest bias. This pattern supports the idea that time-varying risk premia are driving the deviations.

When I turn to the Durbin regression results (Table 7), a very different picture emerges in the long run. The Durbin specification augments the Fama setup with a lagged dependent variable to capture persistence, and indeed I find exchange rate changes to be highly constant over multi-month. This dynamic approach yields an estimated long-run slope (γ), representing the total response of the exchange rate after accounting for cumulative lags, that is much closer to the UIP ideal. In fact, none of the six currencies shows a long run γ significantly different from 1. Despite point estimates of γ that vary across a wide range (some even

negative or above 1), their confidence intervals are large due to the persistent noise, and all include the value 1. For example, the 3M USD/JPY regression that had a short-run β of about -0.8 yields a long-run coefficient $\gamma \approx -0.9$, but with such uncertainty that I cannot reject $H_0: \gamma = 1$ ($p \approx 0.13$). Even currencies that showed little, short-run bias (GBP, CHF) have γ estimates near 1 (often slightly above unity) with no significant departure. In short, once I account for autocorrelation and allow the interest differential effect to accumulate over time, the evidence fails to reject UIP. *This finding aligns with prior long-horizon studies that reported slopes trending toward unity at longer horizons. It appears that short-run deviations are temporary: exchange rates may wander or even move perversely relative to rate differentials in the short term, but over the long haul they “catch up” to reflect cumulative interest differences.*

The contrast between the Fama and Durbin regressions highlights short-term vs. long-term dynamics in UIP. The short-term regressions indicate frequent deviations from parity i.e., the forward rates are poor predictors and carry trade opportunities abound over 3–6-month periods. These deviations coincide with the fact that FX returns are extremely volatile at high frequency. I observe exchange rates driven by market sentiment, risk-on/risk-off flows, and other shocks, which can overwhelm and even invert the expected interest rate effect in the short run. However, the long-run results indicate a reversion toward UIP: as these short-lived shocks average out, currencies that offered higher interest eventually do depreciate.

In other words, forward premium *does* get “compensated back” in the long term, even though in the interim investors might earn excess returns. This reconciliation suggests that UIP holds as a long-run equilibrium condition (consistent with the international Fisher effect). In practical terms, hedging currency exposure with forwards may not neutralize short-term currency risk, but over extended horizons the interest differential tends to be offset by exchange rate movement, restoring parity. The general pattern across my six currencies, therefore, is one of short-term divergence but long-term convergence: an investor can gain from deviations for a while (the carry trade profits), but over time those gains are eroded as exchange rates, realigning with fundamentals. This underscores a key insight for both academia and practitioners, short-run inefficiencies or risk premiums can mask an underlying long-run parity in the currency market.

Sub-sample analysis – Financial Crises 2008

To test whether these findings are robust to the structural breaks, I split the sample to study and re-estimate the Durbin regression on each subsample (Table 8). The results remain consistent: for ten of the twelve currency-horizon pairs, the long-run coefficient γ remains statistically indistinguishable from unity in both

the pre-crisis (January 1990 – September 2008) and post-crisis (September 2008 – November 2025) subsamples. The sole exception is JPY, where long-run UIP is rejected in the pre-crisis subsample at both the 3-month ($\gamma = -2.85$, $p = 0.03$) and 6-month ($\gamma = -3.25$, $p < 0.01$) horizons, consistent with the yen's role as a funding currency in carry-trade strategies during that era. Notably, this rejection vanishes entirely in the post-crisis period ($\gamma \approx 0.3$, $p \approx 0.72$ at both horizons), suggesting that the post-2008 regime shift, with compressed interest rate differentials and changing investor risk behavior, brought the yen into conformity with long-run UIP. This subsample evidence confirms that the full-sample finding of long-run UIP validity is not an object of averaging across structurally different regimes; rather, the result holds broadly within each period, reinforcing the conclusion that short-run deviations from parity are temporary and that exchange rates ultimately adjust in line with interest differentials.

Section H. Robustness

My empirical tests cover a thorough design: I estimate UIP regressions at both 3-month and 6-month horizons for six major USD exchange rates, using two econometric approaches. I run the standard (Fama) regression and a Durbin for each currency-horizon combination. This confirms my results are not driven by one specific specification or timeframe.

The two methods naturally generate different numeric estimates, but they tell a consistent story. The short-horizon Fama regressions yield highly biased slope coefficients: for example, several currencies show negative or near-zero slopes that are statistically far from the UIP-predicted value of 1. By contrast, the Durbin regressions produce long-run slope estimates that are much closer to unity and statistically indistinguishable from 1.

This layered approach (multiple horizons, currencies, and regression forms) improves the credibility of my findings. Showing that the long-run Durbin slopes reliably incline toward unity reinforces that the short-run Fama anomalies are not spurious artifacts. Overall, the robustness of this evidence paves the way to practical implications of UIP's horizon-dependent behavior, that I discuss in the next section.

Additionally, the subsample analysis reported in [Table 8](#) confirms that the long-run UIP result is not driven by pooling structurally different pre- and post-crisis regimes.

Section I. Practical Implications and Contributions

The empirical results have practical view for currency investors and risk managers. In the **short run**, my Fama-style regressions reproduce the well-known “forward premium puzzle”, i.e., high-interest-rate currencies tend to appreciate rather than depreciate, yielding positive carry-trade returns. This implies that carry-trade profits are prevalent and sizable over these horizons but come with a considerable risk premium. Indeed, theory and evidence suggest these excess returns compensate for risk: carry trades are known to be volatile and to suffer losses during market stress. In practical terms, a firm hedging a 3–6-month FX exposure should recognize that forward hedges may be systematically mis-priced by this risk premium. On the other hand, Durbin-regression results indicate that short-term forward contracts do *not* fully neutralize expected currency moves. In other words, hedging via forwards may incur persistent costs or gains in the short run, and managing *short-run* currency risk thus requires active adjustment for this bias. By contrast, when exposures are held over longer horizons my findings imply that the forward premium bias disperses. Exchange rates tend to revert over time so that high-interest currencies eventually depreciate as UIP predicts. This proposes that long-dated hedges and strategic placement can rely more confidently on parity: in the aggregate, the interest differential will be offset by exchange-rate moves in the long run.

From an academic standpoint, this thesis contributes a novel multi-method, multi-horizon framework for UIP testing. I apply both the Fama regression and the recent Durbin regression approach at two short-to-medium horizons (3-month and 6-month) across several major currency pairs. This layered design is, to my knowledge, relatively unique in the literature. It builds on long-horizon UIP studies ([Baillie, Kapetanios & Kim 2026](#)) and the latest time-series techniques but extends them by jointly examining multiple horizons and currencies in a single coherent analysis. In doing so, I clarify that these two approaches are complementary. My evidence confirms that the “Fama puzzle” of near-zero short-run coefficients coexists with a long-run slope near unity, as others have begun to find using Durbin methods.

Section J. Conclusion and Limitations

My analysis of six major USD currency pairs using both conventional Fama regressions and Durbin regressions yields a clear pattern. In the short run, the Fama regressions produce significantly negative or near-zero slopes on the forward premium, indicating a failure of UIP and the presence of the classic forward-premium puzzle. By contrast, once I employ the Durbin regression, the long-run slope estimates are much closer to unity and statistically indistinguishable from 1. In other words, the short-run UIP

“failure” gives way to mean reversion over longer horizons, consistent with the finding that UIP holds in the long run. These results suggest that high-interest-rate currencies appreciate over 3–6 months (yielding carry-trade profits), but over time exchange rates realign with interest differentials as UIP predicts.

However, several limitations qualify these findings, including

- Only six major USD currency pairs were analyzed, so the results may not generalize to less-traded currencies.
- I considered only 3-month and 6-month forward horizons; further long maturities (e.g. 1-year or beyond) were not tested and might exhibit different dynamics, if available.

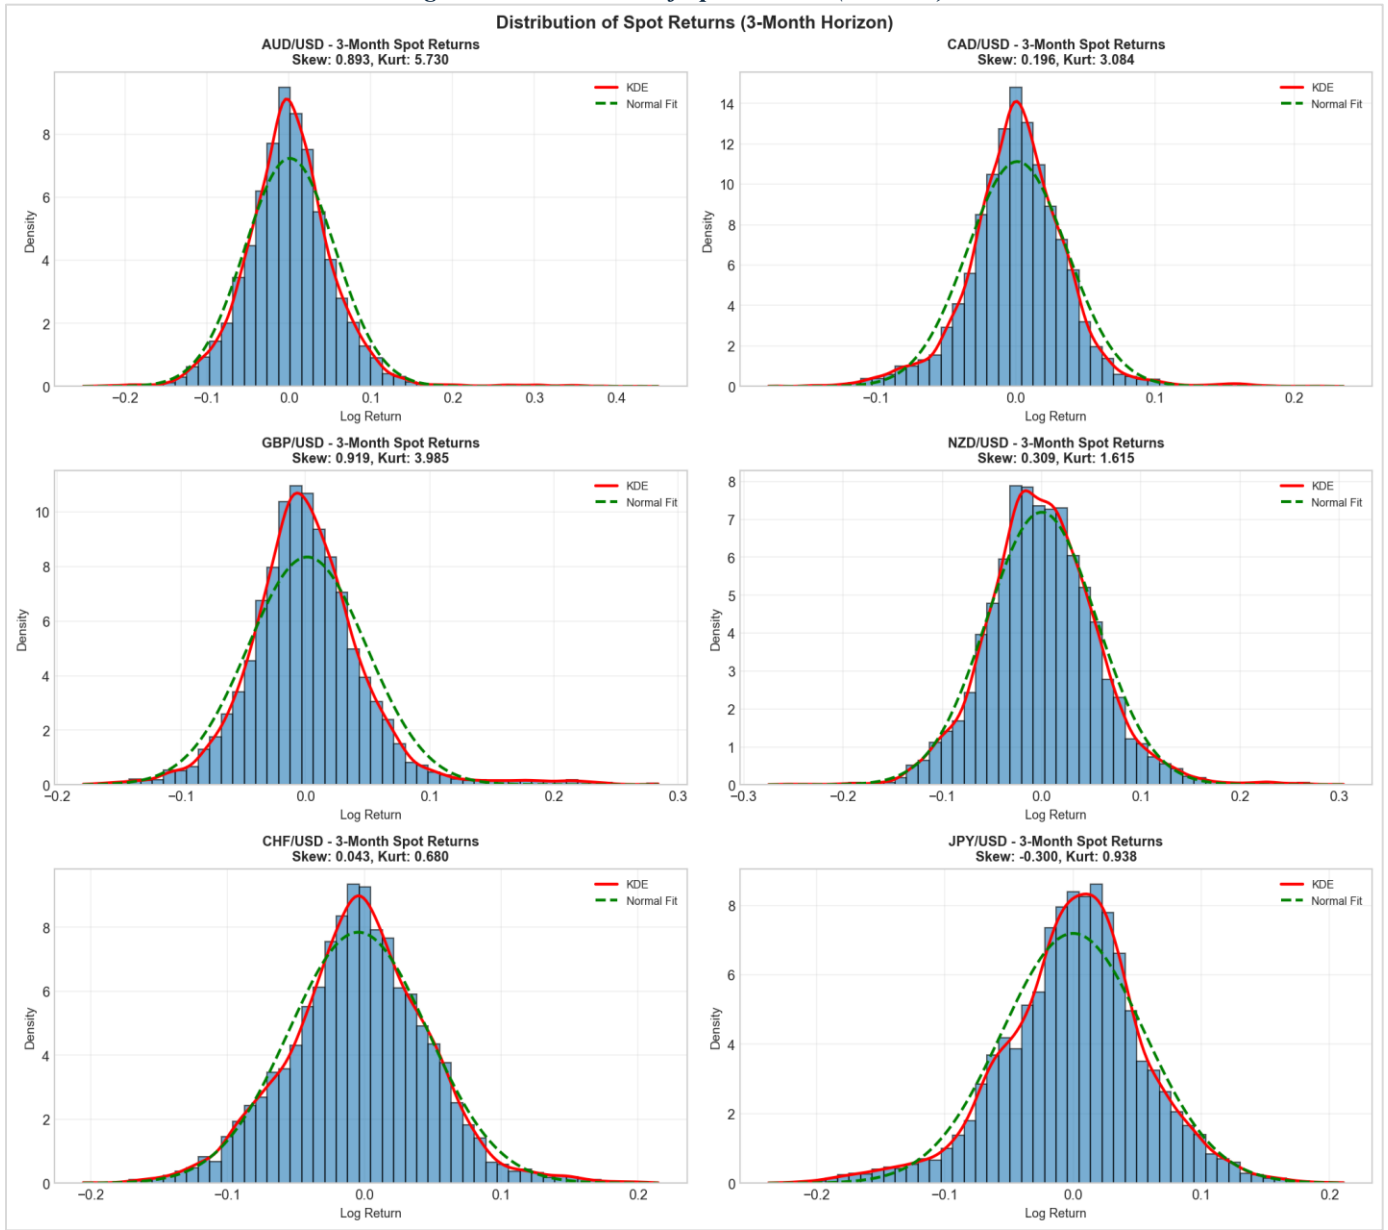
Future research could extend this work by including additional currencies and regime-dependent UIP models to capture any horizon-dependent or regime-sensitive effects.

Section K. Citations and Tables

1. Alexius, A. (2001). Uncovered Interest Parity Revisited. *Review of International Economics*, 9(3), 505–517.
2. Baillie, R., Kapetanios, G., & Kim, K. (2026). Yes! Uncovered Interest Parity Does Hold in the Long Run. *Journal of International Money and Finance*, 160, 103455.
3. Baillie, R. T., Diebold, F. X., Kapetanios, G., & Kim, K. H. (2023). A new test for market efficiency and uncovered interest parity. *Journal of International Money and Finance*, 130, 102765.
4. Baillie, R. T., Kapetanios, G., & Kim, K. H. (2025). Amazingly versatile Durbin regressions with persistent and nonlinear errors: HAC comparisons. *Economics Letters*, 257, 112696.
5. Baillie, R. T., Kim, K. H., Diebold, F. X., Kapetanios, G., & Mora, A. (2024). On Robust Inference in Time Series Regression. *National Bureau of Economic Research, Working Paper 32554*. <http://arxiv.org/abs/2203.04080>
6. Baillie, R. T., Lippens, R. E., & McMahon, P. C. (1983). Testing Rational Expectations and Efficiency in the Foreign Exchange Market. *Econometrica*, 51(3), 553.
7. Bansal, R., & Shaliastovich, I. (2012). A Long-Run Risks Explanation of Predictability Puzzles in Bond and Currency Markets. *National Bureau of Economic Research, Working Paper 18357*.
8. Bekaert, G., Wei, M., & Xing, Y. (2007). Uncovered interest rate parity and the term structure. *Journal of International Money and Finance*, 26, 1038e1069.
9. Brunnermeier, M. K., Nagel, S., & Pedersen, L. H. (2008). Carry Trades and Currency Crashes. *National Bureau of Economic Research, Working Paper 14473*.
10. Bussiere, M., Chinn, M. D., Ferrara, L., & Heipertz, J. (2018). The New Fama Puzzle. *National Bureau of Economic Research, Working Paper 24342*.
11. Chinn, M. D., & Meredith, G. (2004). Monetary Policy and Long-Horizon Uncovered Interest Parity. *IMF Staff Papers*, 51(3), 409–430.
12. Engel, C., Kazakova, K., Wang, M., & Xiang, N. (2022). A reconsideration of the failure of uncovered interest parity for the U.S. dollar. *Journal of International Economics*, 136, 103602.
13. Fama, E. F. (1984). Forward and spot exchange rates. *Journal of Monetary Economics*, 14(3), 319–338.


14. Hakkio', C. S. (1981). Expectations and the Forward Exchange Rate. *International Economic Review*, 22(3), 663–678.
15. Hansen, L. P., & Hodrick, R. J. (1980). Forward Exchange Rates as Optimal Predictors of Future Spot Rates: An Econometric Analysis. *Journal of Political Economy*, 88(5), 829–853.
16. Levy, E., & Nobay, A. R. (1986). The Speculative Efficiency Hypothesis: A Bivariate Analysis. *The Economic Journal*, 96, 109.
17. Lustig, H., Stathopoulos, A., & Verdelhan, A. (2019). The Term Structure of Currency Carry Trade Risk Premia. *American Economic Review*, 109(12), 4142–4177.
18. Snaith, S., Coakley, J., & Kellard, N. (2013). Does the forward premium puzzle disappear over the horizon? *Journal of Banking & Finance*, 37(9), 3681–3693.

Figure 1 - Distribution of Spot Returns (3-month)



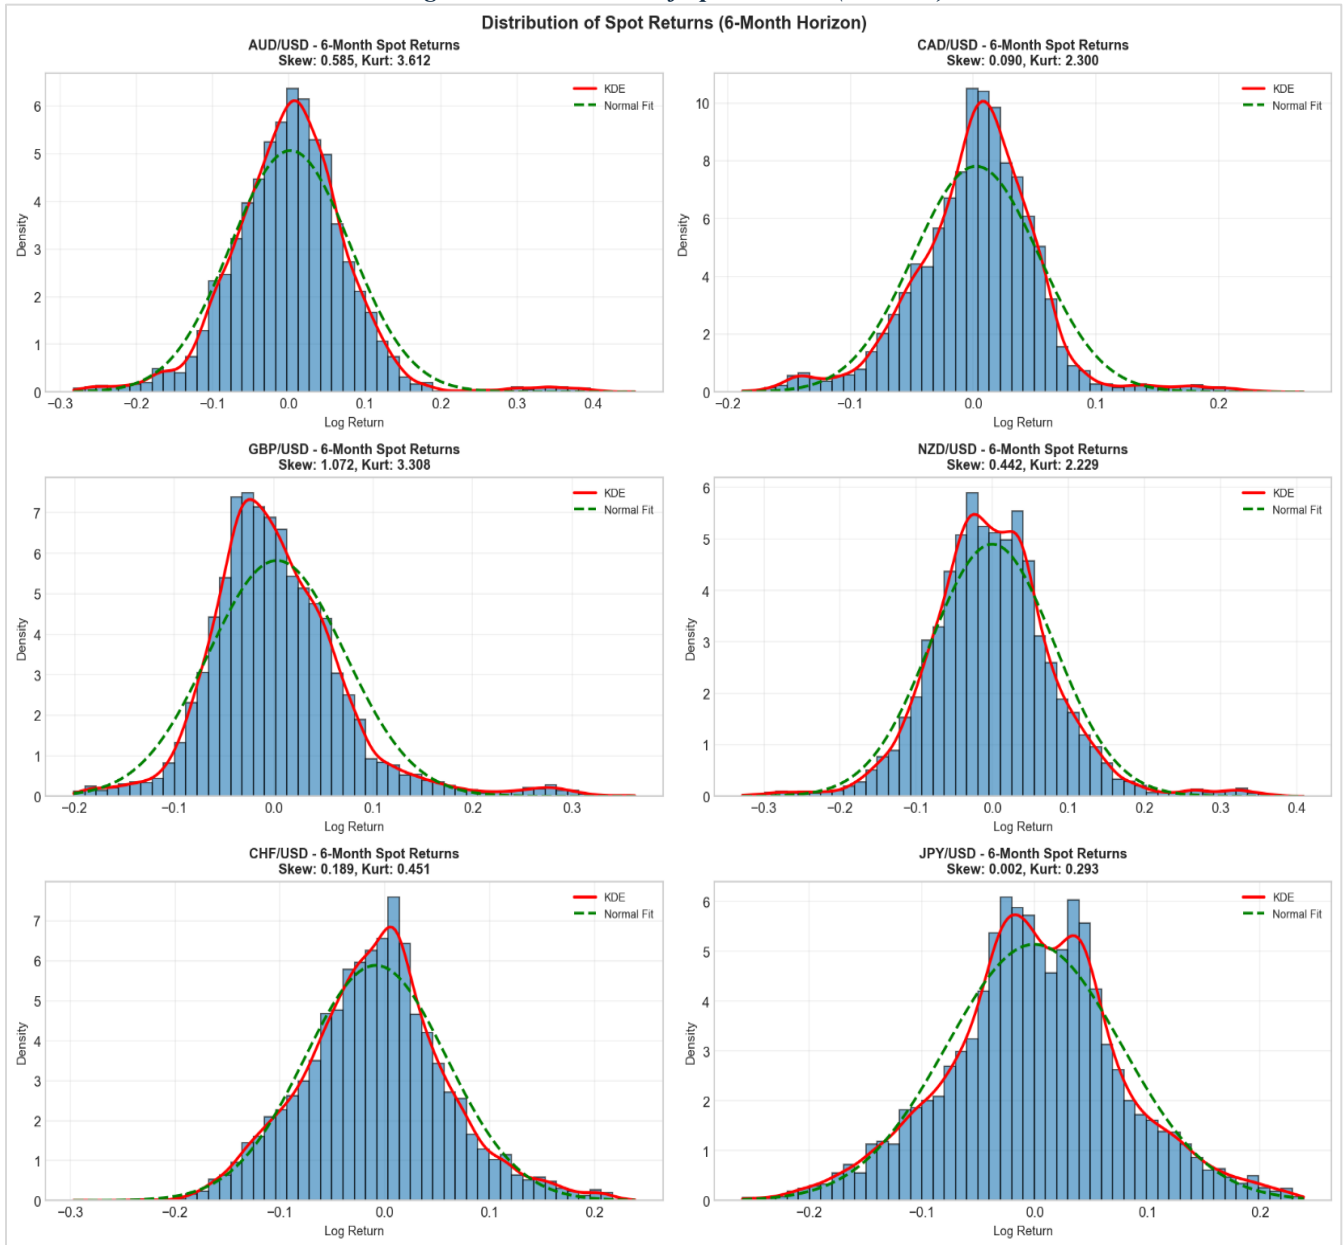
Distribution of spot returns (3-month) for all six currencies

Figure 2 - Time series: 3M Spot return vs Forward premium



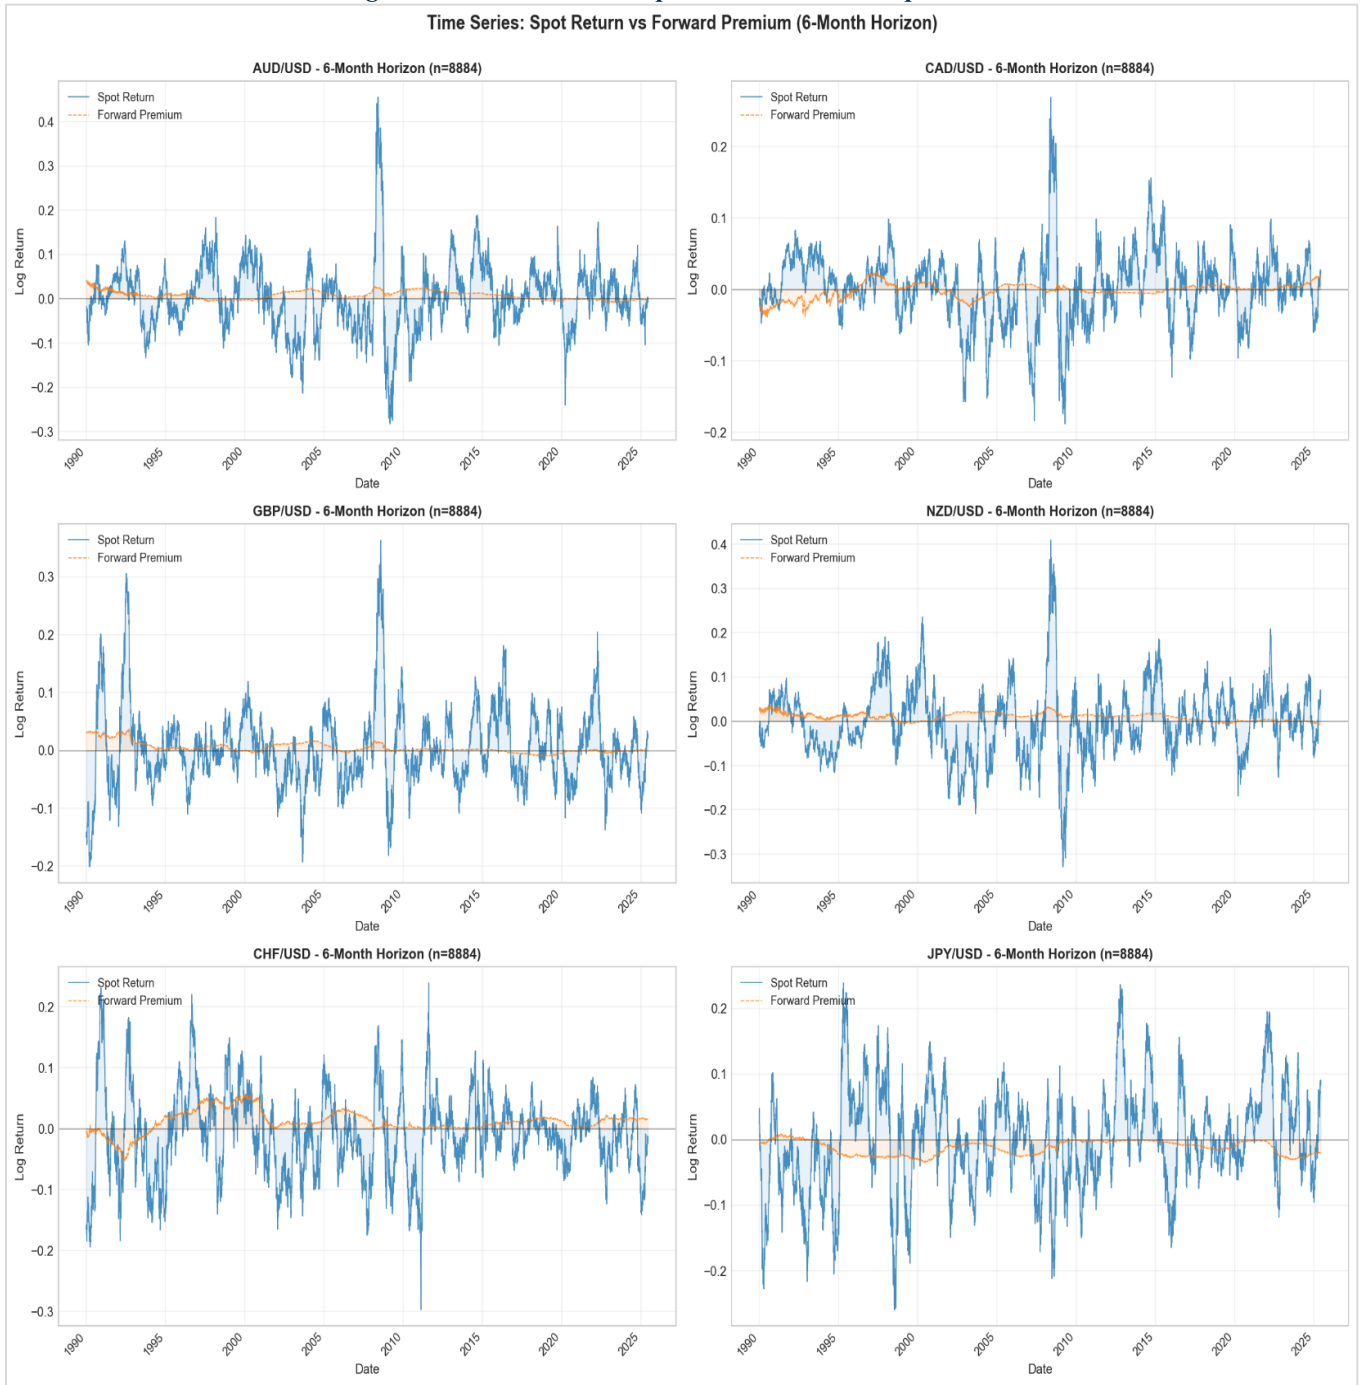
Timeseries of spot and forward returns (3-month) for all six currencies

Figure 3 - Distribution of Spot Returns (6-month)



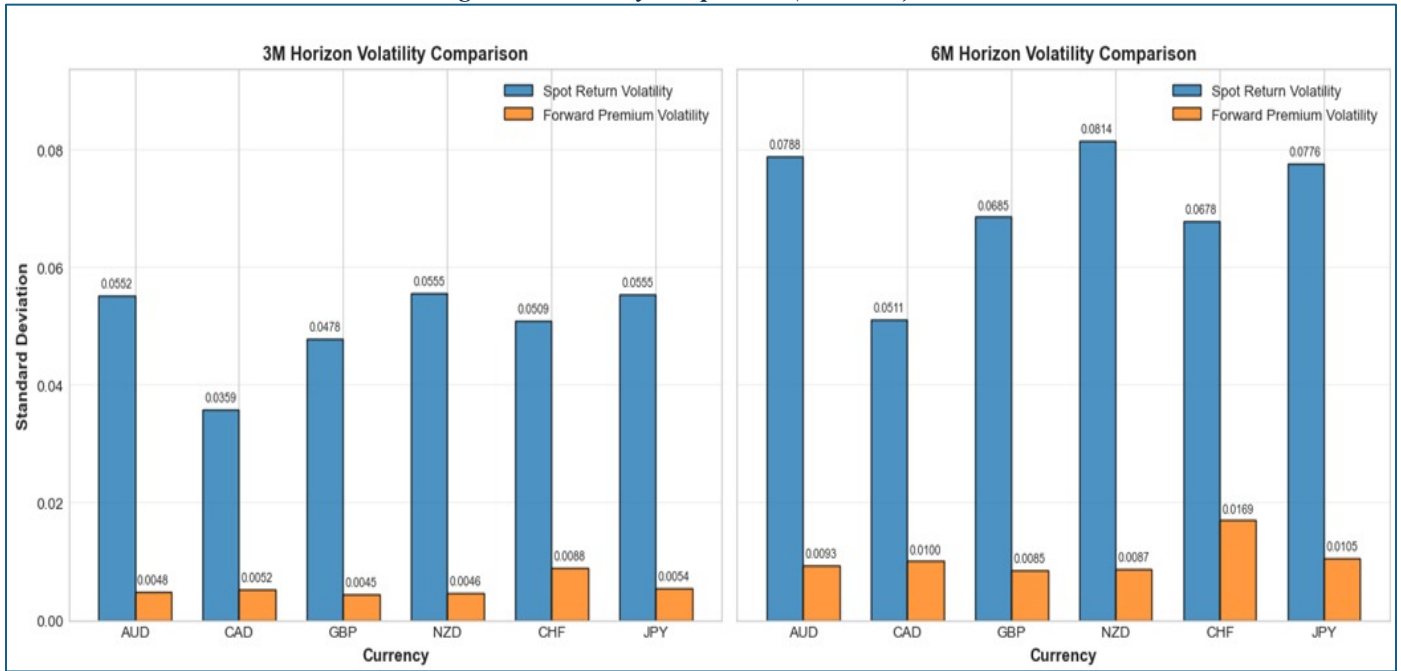
Distribution of spot returns (6-month) for all six currencies

Figure 4 - Time series: 6M Spot return vs Forward premium



Timeseries of spot and forward returns (6-month) for all six currencies

Figure 5 - Volatility comparison (3m vs 6m)



The comparison of volatilities (3m vs 6m) for all six currencies

Figure 6 - Volatility comparison (3m vs 6m)

Correlation heatmap for all six currencies, both horizons

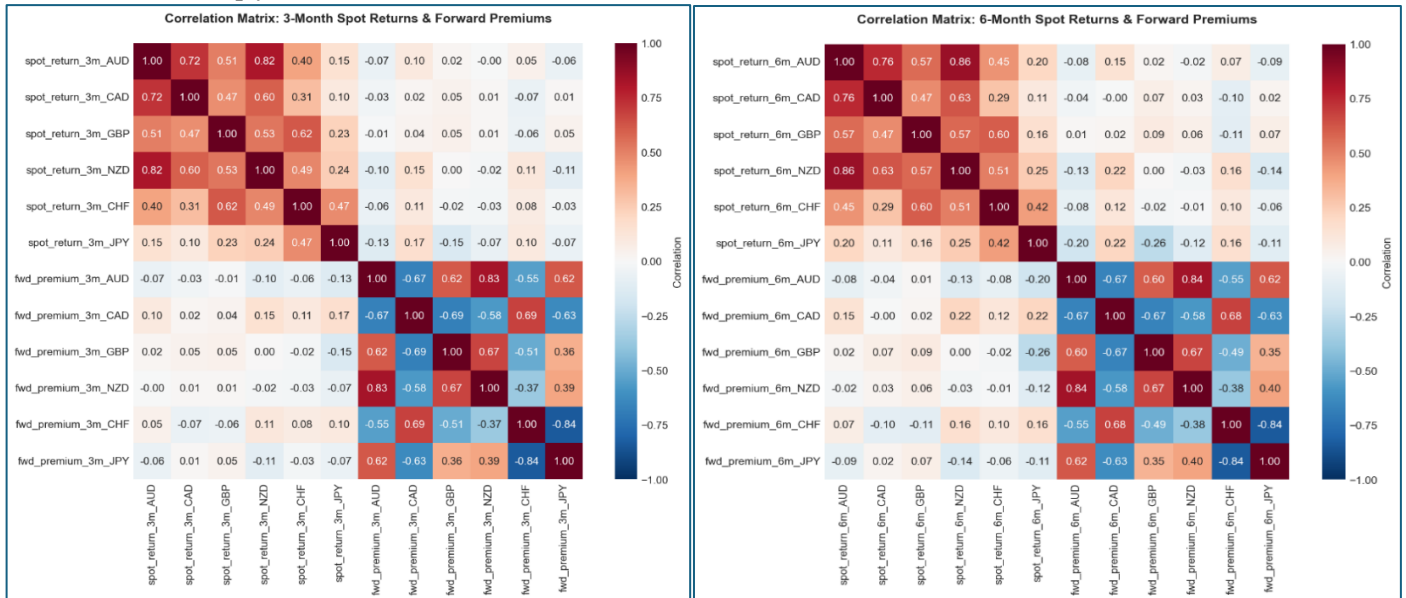
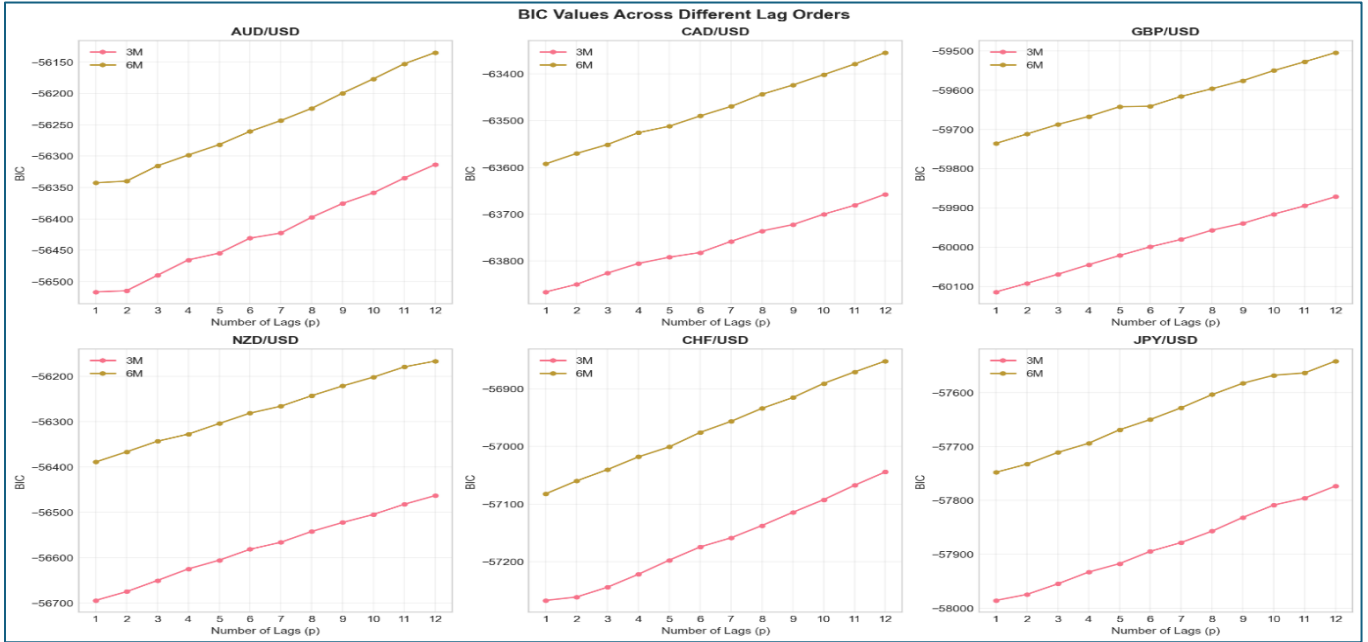


Figure 7 - BIC Values Across Lags



The comparison of volatilities (3m vs 6m) for all six currencies

Table 1 - Data Characteristics

Characteristics	Description
Source	Bloomberg Terminal
Sample period	02-Jan-1990 to 28-Nov-2025
Frequency	Daily (trading days)
Total observations	9,010 per currency ¹
Currencies	AUD, CAD, GBP, NZD, CHF, JPY (all vs. USD)
Instruments	Spot rates, 3-month forwards, 6-month forwards ²
Quote type	Mid-price quotes

The comparison of volatilities (3m vs 6m) for all six currencies

Table 2 - Observations post variable construction

Horizon	Total Observations	Valid Observations	Missing (end of sample)
3-month	9,010	8,947	63
6-month	9,010	8,884	126

Table 3 - 3M spot returns and 3M forward premiums (summary statistics)

	Obs.	Mean	Median	Std. Dev.	Skewness	Kurtosis	Min	Max	IQR
<i>AUD_Spot</i>	8947	0.001116	-0.00055	0.055155	0.892705	5.730439	-0.25157	0.451329	0.061035
<i>AUD_Fwd</i>	9010	0.003811	0.003479	0.004792	0.489432	-0.34985	-0.00407	0.024782	0.007541
<i>CAD_Spot</i>	8947	0.001193	0.001268	0.035877	0.19634	3.084478	-0.17744	0.235573	0.040261
<i>CAD_Fwd</i>	9010	-0.00041	9.50E-05	0.005205	-0.6532	1.213894	-0.02208	0.012634	0.005227
<i>GBP_Spot</i>	8947	0.001507	-0.00145	0.047797	0.918646	3.984823	-0.179	0.284156	0.051216
<i>GBP_Fwd</i>	9010	0.001976	0.000643	0.004464	1.531318	2.413698	-0.00513	0.018432	0.004341
<i>NZD_Spot</i>	8947	0.000194	-0.00155	0.055502	0.309358	1.614894	-0.27517	0.304563	0.068447
<i>NZD_Fwd</i>	9010	0.005106	0.00601	0.004557	0.085823	-0.55599	-0.01421	0.021153	0.007531
<i>CHF_Spot</i>	8947	-0.00447	-0.00413	0.05087	0.043175	0.679878	-0.20564	0.214801	0.063009
<i>CHF_Fwd</i>	9010	0.005937	0.005017	0.008829	-0.15713	1.739766	-0.03021	0.02877	0.007864
<i>JPY_Spot</i>	8947	8.58E-05	0.002809	0.055467	-0.30038	0.93764	-0.23755	0.211035	0.06524
<i>JPY_Fwd</i>	9010	-0.00614	-0.00478	0.005361	-0.23404	-1.26425	-0.01743	0.005438	0.010731

¹ Original sample size was 9,370 but it's reduced to 9,010 due to missing quotes.

² Forward rates are calculated using the pips points. All currencies are scaled by 10,000 except of JPY, which is scaled by 100.

Table 4 - 6M spot returns and 6M forward premiums (summary statistics)

	Obs.	Mean	Median	Std. Dev.	Skewness	Kurtosis	Min	Max	IQR
<i>AUD_Spot</i>	8884	0.002364	0.002883	0.078773	0.584545	3.61205	-0.28248	0.454937	0.09018
<i>AUD_Fwd</i>	9010	0.007418	0.006754	0.009283	0.434968	-0.6168	-0.00822	0.042529	0.014695
<i>CAD_Spot</i>	8884	0.002356	0.005155	0.051075	0.090031	2.299593	-0.18834	0.269386	0.057253
<i>CAD_Fwd</i>	9010	-0.00084	7.71E-05	0.010019	-0.60276	0.941361	-0.03856	0.024606	0.010395
<i>GBP_Spot</i>	8884	0.003236	-0.00502	0.068531	1.071991	3.307975	-0.20101	0.363126	0.076319
<i>GBP_Fwd</i>	9010	0.0038	0.001223	0.008528	1.482642	2.347286	-0.00964	0.036903	0.008426
<i>NZD_Spot</i>	8884	3.72E-05	-0.0033	0.081446	0.442227	2.228933	-0.32856	0.409259	0.094955
<i>NZD_Fwd</i>	9010	0.009985	0.011719	0.008698	0.041017	-0.72127	-0.00773	0.0362	0.014423
<i>CHF_Spot</i>	8884	-0.00854	-0.00756	0.067763	0.189343	0.450818	-0.29685	0.238779	0.082111
<i>CHF_Fwd</i>	9010	0.012149	0.010241	0.016942	-0.04361	1.624113	-0.05404	0.056025	0.015356
<i>JPY_Spot</i>	8884	-0.00023	-0.001	0.077608	0.00216	0.293458	-0.25942	0.238846	0.092426
<i>JPY_Fwd</i>	9010	-0.0125	-0.00993	0.010453	-0.25673	-1.28295	-0.03541	0.008726	0.020705

Table 5 - BIC and AIC (Optimal Lag(s))

Currency	Horizon	Optimal_Lag_BIC	BIC_Value	Optimal_Lag_AIC	AIC_Value
AUD	3M	1	-56516.6	2	-56557.4
AUD	6M	1	-56342.5	2	-56382.4
CAD	3M	1	-63865.9	1	-63894.3
CAD	6M	1	-63592.2	1	-63620.6
GBP	3M	1	-60114.1	1	-60142.5
GBP	6M	1	-59735.7	1	-59764.1
NZD	3M	1	-56694	1	-56722.4
NZD	6M	1	-56388.7	1	-56417.1
CHF	3M	1	-57266.9	2	-57303.6
CHF	6M	1	-57082.2	1	-57110.5
JPY	3M	1	-57985.7	2	-58017.2
JPY	6M	1	-57747.7	1	-57776.1

Table 6 - Test of UIP (Conventional Fama)

Fama Regression | $s_{t+k} - s_t = \alpha + \beta(f_t - s_t) + u_{t+k}$

Currency	Horizon	α	(s.e.)	β	p-value	R ²	Adj. R ²	n	UIP (H ₀ : $\beta = 1$)
AUD	3M	0.0042	-0.8982	-0.8071**	0.0442**	0.0049	0.0048	8946	Reject UIP
AUD	6M	0.0077	-0.7665	-0.7104**	0.0257**	0.007	0.0069	8883	Reject UIP
CAD	3M	0.0013	-0.3626	0.1482**	0.0188**	0.0005	0.0003	8946	Reject UIP
CAD	6M	0.0023	-0.3321	-0.0178***	0.0022***	0	-0.0001	8883	Reject UIP
GBP	3M	0.0005	-1.3199	0.5094	0.7102	0.0023	0.0022	8946	Fail to Reject
GBP	6M	0.0005	-1.0175	0.7098	0.7755	0.0079	0.0078	8883	Fail to Reject
NZD	3M	0.0013	-0.9612	-0.208	0.2089	0.0003	0.0002	8946	Fail to Reject
NZD	6M	0.0027	-0.9123	-0.2585	0.1677	0.0007	0.0006	8883	Fail to Reject
CHF	3M	-0.0071	-0.4874	0.4513	0.2603	0.0062	0.0061	8946	Fail to Reject
CHF	6M	-0.0135	-0.3216	0.4080*	0.0657*	0.0106	0.0104	8883	Fail to Reject
JPY	3M	-0.004	-0.7182	-0.6758**	0.0197**	0.0043	0.0042	8946	Reject UIP
JPY	6M	-0.0105	-0.5649	-0.8258***	0.0012***	0.0125	0.0124	8883	Reject UIP

Note: *** p<0.01, ** p<0.05, * p<0.10

UIP (H₀: $\beta=1$): Reject UIP if β significantly different from 1; Fail to Reject if $\beta \approx 1$

Table 7 - Test of UIP (Durbin Regression - Long-run Coefficient)

Durbin Regression | $s_{t+k} - s_t = \alpha_0 + \sum_{j=1}^p \alpha_j (s_{t+k-j} - s_{t-j}) + \sum_{j=0}^p \beta_j (f_{t-j} - s_{t-j}) + \varepsilon_{t+k}$ | $\gamma = \frac{\sum_{j=0}^p \beta_j}{1 - \sum_{j=1}^p \alpha_j} = \frac{\beta(1)}{\alpha(1)}$

Currency	Horizon	p	γ	(s.e.)	β_j	α_j	p-value	R ²	n	Long-run UIP (H ₀ : $\gamma = 1$)
AUD	3M	1	-0.3149	-1.3032	-0.0055	0.9826	0.313	0.9654	8946	Fail to Reject
AUD	6M	1	0.715	-1.4576	0.0058	0.9919	0.845	0.9835	8883	Fail to Reject
CAD	3M	1	-0.0568	-0.7688	-0.001	0.9819	0.1693	0.9641	8946	Fail to Reject
CAD	6M	1	-0.5545*	-0.8359	-0.0048	0.9913	0.0630*	0.9826	8883	Fail to Reject
GBP	3M	1	1.1299	-1.2823	0.0175	0.9845	0.9193	0.9692	8946	Fail to Reject
GBP	6M	1	2.3176	-1.3622	0.018	0.9922	0.3334	0.9851	8883	Fail to Reject
NZD	3M	1	0.3072	-1.4171	0.0052	0.9831	0.6249	0.9665	8946	Fail to Reject
NZD	6M	1	0.1652	-1.6374	0.0013	0.9923	0.6102	0.9846	8883	Fail to Reject
CHF	3M	1	0.1425	-0.636	0.0026	0.9815	0.1776	0.9626	8946	Fail to Reject
CHF	6M	1	0.0386	-0.597	0.0004	0.9898	0.1073	0.9794	8883	Fail to Reject
JPY	3M	1	-0.9077	-1.2607	-0.0134	0.9852	0.1303	0.971	8946	Fail to Reject
JPY	6M	1	-1.1595*	-1.2756	-0.0086	0.9926	0.0905*	0.9855	8883	Fail to Reject

Note: *** p<0.01, ** p<0.05, * p<0.10

$\Sigma\beta$: Sum of all forward premium coefficients; $\Sigma\alpha$: Sum of lagged spot return coefficients

Formula: $\gamma = \Sigma\beta / (1 - \Sigma\alpha)$

Table 8 - Subsample Test of UIP Pre/Post 2008 Financial Crises (Durbin Regression - Long-run Coefficient)

Currency	Horizon	GFC 2008	γ	(s.e.)	β_j	α_j	p-value	R ²	n	Long-run UIP (H ₀ : $\gamma = 1$)
AUD	3M	Pre-Crisis	1.331968	2.435855	0.017281	0.987026	0.891603	0.970603	4532	Fail to Reject
AUD	3M	Post-Crisis	-1.31116	1.566369	-0.02824	0.978461	0.140153	0.959818	4413	Fail to Reject
AUD	6M	Pre-Crisis	3.649095	3.133817	0.020908	0.99427	0.397974	0.986097	4532	Fail to Reject
AUD	6M	Post-Crisis	-1.00558	1.700958	-0.00968	0.990377	0.238428	0.981134	4350	Fail to Reject
CAD	3M	Pre-Crisis	0.201659	1.115075	0.002469	0.987758	0.474057	0.970862	4532	Fail to Reject
CAD	3M	Post-Crisis	-1.82643	1.824905	-0.04219	0.9769	0.1215	0.957877	4413	Fail to Reject
CAD	6M	Pre-Crisis	-0.46752	1.452055	-0.00236	0.994956	0.312238	0.98685	4532	Fail to Reject
CAD	6M	Post-Crisis	-2.19306	1.781121	-0.02698	0.987698	0.073087	0.9782	4350	Fail to Reject
GBP	3M	Pre-Crisis	1.666674	2.090376	0.021145	0.987313	0.749798	0.972206	4532	Fail to Reject
GBP	3M	Post-Crisis	1.052416	3.876018	0.019894	0.981097	0.989211	0.965057	4413	Fail to Reject
GBP	6M	Pre-Crisis	3.6068	2.463496	0.021664	0.993994	0.290034	0.98685	4532	Fail to Reject
GBP	6M	Post-Crisis	0.021702	3.530869	0.000222	0.989769	0.78174	0.982254	4350	Fail to Reject
NZD	3M	Pre-Crisis	2.427351	2.510611	0.032322	0.986684	0.569705	0.970568	4532	Fail to Reject
NZD	3M	Post-Crisis	-1.75579	2.163351	-0.03593	0.979538	0.202783	0.962061	4413	Fail to Reject
NZD	6M	Pre-Crisis	4.102006	4.060563	0.019177	0.995325	0.444946	0.98809	4532	Fail to Reject
NZD	6M	Post-Crisis	-2.39279	1.823435	-0.02929	0.987759	0.06286	0.979802	4350	Fail to Reject
CHF	3M	Pre-Crisis	0.237528	0.832691	0.003757	0.984185	0.359887	0.968093	4532	Fail to Reject
CHF	3M	Post-Crisis	-2.38954	1.905725	-0.05741	0.975974	0.075373	0.951789	4413	Fail to Reject
CHF	6M	Pre-Crisis	0.069805	0.775675	0.000619	0.99113	0.23051	0.982403	4532	Fail to Reject
CHF	6M	Post-Crisis	-1.88673	1.778393	-0.02432	0.987111	0.104614	0.973693	4350	Fail to Reject
JPY	3M	Pre-Crisis	-2.85166	1.812712	-0.04277	0.985003	0.033657	0.970993	4532	Reject UIP
JPY	3M	Post-Crisis	0.352254	1.816656	0.005599	0.984105	0.721439	0.970112	4413	Fail to Reject
JPY	6M	Pre-Crisis	-3.2503	1.600281	-0.02824	0.991313	0.007936	0.985203	4532	Reject UIP
JPY	6M	Post-Crisis	0.299642	1.966636	0.002253	0.992482	0.721768	0.984887	4350	Fail to Reject

Note: *** p<0.01, ** p<0.05, * p<0.10

$\Sigma\beta$: Sum of all forward premium coefficients; $\Sigma\alpha$: Sum of lagged spot return coefficients

Formula: $\gamma = \Sigma\beta / (1 - \Sigma\alpha)$

Appendix

Uncovered Interest Rate Parity (UIP): The condition under which the expected change in the spot exchange rate between two currencies equals the nominal interest rate differential between those two countries, assuming risk-neutral and rational investors.

International Fisher Effect (IFE): The proposition that currencies with higher nominal interest rates will depreciate in the long run by approximately the interest rate differential, linking nominal interest rate differences to expected exchange rate movements through inflation expectations.

Covered Interest Parity (CIP): The no-arbitrage condition stating that the forward premium on a currency equals the interest rate differential between two countries, enforced by riskless arbitrage involving simultaneous borrowing, investing, and forward contracting.

Forward Premium Puzzle: The empirically observed phenomenon, first documented by Fama (1984), in which the forward premium predicts exchange rate movements in the wrong direction, the slope coefficient β in the Fama regression is typically negative or significantly below 1, contrary to UIP's prediction of $\beta = 1$.

FX Forward Contract: A binding agreement between two parties to exchange a specified amount of one currency for another at a predetermined exchange rate (the forward rate) on a future settlement date.

Forward Exchange Rate: The rate agreed upon today at which two parties will exchange currencies at a specified future date, reflecting both the current spot rate and the interest rate differential between the two currencies.

Spot Exchange Rate: The current market price at which one currency can be exchanged for another for immediate delivery.

Hedging (in this context): The practice by which international investors, multinational corporations, and portfolio managers (including banks) use FX forward contracts to lock in a future exchange rate, thereby reducing or eliminating the uncertainty of converting foreign-currency cash flows back into their domestic currency.

Carry Trade: A strategy in which an investor borrows in a low-interest-rate currency and invests the proceeds in a high-interest-rate currency, profiting from the interest rate differential if the high-yield currency does not depreciate as much as UIP would predict.

Fama Regression: The standard empirical test of UIP that regresses the realized change in the spot exchange rate on the forward premium.

Durbin Regression: A dynamic time-series regression that augments the standard UIP test with lagged dependent and independent variables to capture persistence and serial correlation arising from overlapping observations, yielding consistent and efficient estimates of both short-run and long-run UIP coefficients.

Risk Premium (Time-Varying): The additional expected return that investors demand for bearing the uncertainty of holding a foreign currency, which varies over time and is a leading explanation for short-run UIP failures.

Peso Problem: A situation in which rare but anticipated events (e.g., a large currency devaluation) that have not yet occurred within the sample period distort expectations and cause empirical tests to reject UIP, even though UIP may hold in the true long-run distribution.

Overlapping Observations: A feature of long-horizon regressions using high-frequency data, where successive return observations share common time periods (e.g., daily 3-month returns overlap by 62 of 63 days), inducing strong serial correlation in regression residuals.

BIC (Bayesian Information Criterion) / AIC (Akaike Information Criterion): Model selection criteria that balance goodness-of-fit against model complexity; BIC imposes a stricter penalty for additional parameters and is used in this study as the primary guide for selecting the optimal lag length in Durbin regressions.