Currencies in the US Session: A Replication of Zhang (2018) and the Persistent Weakness of the Yen

Arthur Pegliasco
Undergraduate Student in Economics, Université Côte d'Azur
arthur.pegliasco@gmail.com

November 2026

Abstract

Zhang (2018) reports that major currencies depreciate against the dollar during their own domestic trading hours and appreciate during the US session, with the Japanese yen as a notable exception. I replicate the study on seven major currencies with one-minute data. Over the original sample (2010–2015) the session means match the published table almost to the hundredth of a basis point. Out of sample (2016 to August 2026), the two central hypotheses no longer hold: none of fourteen tests survives a Holm correction, and the US-session appreciation of the euro and the pound fades. The exception, however, becomes the rule. The yen falls against the dollar between 11:00 and 17:00 New York time by 1.7 basis points a day (Newey–West $t = -4.3$), and it is negative in sixteen of the seventeen years from 2010 to 2026, the remaining one being flat. The drift is a shift of the whole distribution rather than the result of a few large days, it is not explained by the dollar, by US equities or by futures carry, and it vanishes on days of risk-off stress. It is concentrated on days that follow a rise of the yen. Net of costs it is small, and the only evidence that can confirm it is new data, which are reserved under a pre-registered plan.

Keywords: foreign exchange, intraday seasonality, trading sessions, Japanese yen, replication, pre-registration
JEL classification: F31, G14, G15

1. Introduction

The foreign exchange market trades around the clock, but its activity moves with the sun: Asian centres open first, then London, then New York. A line of research documents that currency returns are not spread evenly across these sessions. Ranaldo (2009) and Breedon and Ranaldo (2013) find that a currency tends to depreciate during its own domestic trading hours and to appreciate during foreign hours, a pattern they relate to order flow: domestic investors are net buyers of foreign currency when their own market is open. Zhang (2018) extends this result to 2010–2015 and to seven major currencies, and adds a third claim linking session returns to realised volatility.

Patterns of this kind are fragile. They are easy to find in a given sample, attractive to trade, and often gone once they are published. This paper does two things. First, it replicates Zhang (2018) exactly on its own period and then tests its hypotheses on the ten years that followed, under a pre-registered protocol. Second, it studies in detail the one pattern that survives: the yen, which Zhang's own table shows falling against the dollar during the US session, against the direction of the hypothesis.

The results are clear on both counts. The replication works on the original sample but fails out of sample, for every currency. The yen effect does the opposite: it holds out of sample with a t-statistic above four, it is present in almost every year, and it cannot be traced to a handful of days. Section 2 states the hypotheses, Section 3 describes the data, Sections 4 and 5 report the replication and the out-of-sample test, Section 6 studies the yen, Section 7 discusses explanations, Section 8 sets out the pre-registered tests still to come, and Sections 9 and 10 give limitations and conclusions.

2. Hypotheses

Following Zhang (2018), each day is divided into four sessions, all defined in local time so that they follow daylight saving changes: the *domestic* session, from 08:00 to 16:00 in the currency's home market; the *London–New York overlap*, from 08:00 to 11:00 New York time; the *US session*, from 11:00 to 17:00 New York time; and the *non-US* session, which covers the remaining hours. Prices are quoted in dollars per unit of foreign currency, so a negative return means that the currency falls against the dollar. The three hypotheses tested are those of the paper:

1. H1: a currency depreciates during its domestic session;
2. H2: a currency appreciates during the US session;
3. H3: session returns are negatively related to realised volatility during the domestic session.

The yen is a counter-example to H2 in Zhang's own table (−2.12 bp a day in the US session over 2010–2015, significant at 1 %). It was not one of the pre-registered hypotheses, but it is reported throughout and becomes the subject of Section 6.

3. Data

The seven currencies are the euro, pound, yen, Swiss franc, Australian dollar, Canadian dollar and New Zealand dollar, all against the US dollar. Table 1 summarises the sources. For the replication

period I use Dukascopy spot mid-quotes. For the out-of-sample period I use CME futures (6E, 6B, 6J, 6A) where they exist, back-adjusted at each of their 42 rolls, and Dukascopy spot for the other three currencies. Nothing after 28 August 2026 is used: later data are reserved for the pre-registered tests of Section 8.

Table 1. Samples and sources

Sample	Period	Source
Replication (R)	3 May 2010 – 31 December 2015	Dukascopy spot, mid-quote, one-minute bars
Out of sample (O)	1 January 2016 – 28 August 2026	CME futures 6E, 6B, 6J, 6A (one-minute bars, back-adjusted); Dukascopy spot for CHF, CAD and NZD
Yen, earlier spot check	January 2012 – December 2015	Dukascopy spot USD/JPY, bid and ask
Reserved (not used)	After 28 August 2026; USD/JPY spot 2003–2009	Kept for the pre-registered tests of Section 8

Notes: All times are New York time unless stated. The futures roll at midnight UTC, outside the US session. Trading cost for futures is one tick plus $5 per round trip, about 1.8 to 2 ticks.

Inference uses Newey–West standard errors (Newey and West, 1987) for the out-of-sample tests, Holm's (1979) correction within each family, and a GARCH(1,1) specification for H3, as in the original paper. The protocol, the hypotheses and the code were committed before any computation.

4. Replication on the original sample

Table 2 compares my session means for 2010–2015 with Zhang's Table 1. For the euro, pound, yen, Canadian and New Zealand dollars, the numbers agree almost to the hundredth of a basis point. Of the nine cells that the paper marks as significant at 5 % or better for these seven currencies, seven are reproduced with the same sign and $p < 0.05$.

Table 2. Mean session returns, 2010–2015: this paper against Zhang (2018), basis points per day

	Domestic		London–NY		US session		Non-US	
Currency	Here	Zhang	Here	Zhang	Here	Zhang	Here	Zhang
EUR	−2.17*	−2.10*	−0.65	−0.63	+2.20***	+2.21***	−0.57	−0.67
GBP	−1.46	−1.48	−0.50	−0.44	+1.65***	+1.66***	−0.27	−0.39
JPY	+0.88	+0.95	+0.76	+0.90	−2.03***	−2.12***	−0.65*	−1.36*
CHF	+1.51	+3.74**	+0.70	+2.70**	+1.51	+1.34	−2.19***	−4.12***
AUD	−1.89*	−1.65	−1.16	−1.61*	+1.25	+0.71	+0.21	+1.85*
CAD	−1.62	−1.69	−2.40***	−2.50***	+0.83	+0.80	−0.55	−0.41
NZD	−2.56**	−2.60**	−1.45	−1.85*	+2.02**	−0.47	+2.44**	+3.10***

Notes: Stars use a simple t-test, as in the original paper: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. A negative value means the currency falls against the dollar.

Both failures concern the Swiss franc and come from a single day, 15 January 2015, when the Swiss National Bank abandoned its floor against the euro. On that day the franc moves by about +1,410 bp in the domestic session, −395 bp in the overlap, +801 bp in the US session and −436 bp in the remaining hours, enough to shift each session mean by about 1 bp over 1,430 days. How that shock is split across hours depends on the data vendor, so this is a data issue, not a disagreement of substance. The only unexplained gap is the New Zealand dollar in the US session (+2.02** here against −0.47 in the paper), probably also a matter of source.

Within the pre-registered family of 21 tests, two pass Holm's correction: H2 for the euro (adjusted p = 0.029) and for the pound (0.031). H3 is not confirmed anywhere, even on this sample. The volatility coefficients have the same size as in the paper (for instance −0.029 against −0.031 for the euro, −0.081 against −0.092 for the Canadian dollar), but none is significant under GARCH (best t = −1.75), whereas the paper reports them at 1 %. Since the relation is contemporaneous, it would not allow any forecast in any case.

5. Out of sample, 2016–2026

Table 3 repeats the exercise on the following ten years. The domestic-session weakness has disappeared: no currency falls significantly during its own hours. The US-session strength has mostly gone too. The euro and the pound still gained about +0.9 bp a day in the US session over 2016–2020, but the euro drops to +0.06 bp over 2021–2026, and the pound to +0.8 bp (t = 1.4). None of the fourteen tests of H1 and H2 passes Holm's correction; the best raw p-values are 0.046 (pound) and 0.048 (Canadian dollar), which become 0.97 after adjustment.

Table 3. Mean session returns out of sample, January 2016 – August 2026, basis points per day

Currency	Domestic	London–NY	US session	Non-US	Newey–West t (dom. / LDN–NY / US / non-US)
EUR	−0.78	+0.70	+0.44	+0.22	−1.1 / 1.3 / 1.1 / 0.6
GBP	+0.06	+0.59	+0.85	−1.49	0.1 / 1.0 / 1.7 / −2.6
JPY	−0.07	+0.60	**−1.72**	−0.42	−0.1 / 1.0 / **−4.3** / −1.7
CHF	+0.34	+1.48	−0.06	+0.33	0.5 / 2.8 / −0.1 / 1.0
AUD	−0.07	+0.74	−0.16	−0.35	−0.1 / 1.2 / −0.3 / −0.9
CAD	+0.76	+0.43	+0.66	−1.10	1.2 / 0.9 / 1.7 / −2.4
NZD	−0.06	+0.89	−0.06	−0.89	−0.1 / 1.4 / −0.1 / −1.3

Notes: EUR, GBP, JPY and AUD from CME futures, the others from spot. Removing weekend returns changes no conclusion.

Three cells are significant before correction, with the same sign in both halves of the sample, but do not pass Holm's correction over the 28 cells: the pound in the non-US session (−1.5 bp; halves −1.9 and −1.1), the Canadian dollar in the non-US session (−1.1 bp; −0.7 and −1.4) and the Swiss franc in the London–New York overlap (+1.5 bp; +1.5 and +1.4). I keep them on a watch list. None of them, and no other out-of-sample move apart from the yen, exceeds the cost of a round trip, which ranges from about 0.4 bp (one tick on the euro future) to 1.6–4.2 bp for spot sessions that touch the 17:00 rollover.

The yen is the exception. In the US session it falls by −1.72 bp a day (Newey–West t = −4.28, two-sided p = 2×10^{-5}, Holm-adjusted over 28 cells about 5×10^{-4}). The sign holds in both halves of the sample (−1.14 and −2.24 bp), and across 2010–2026 the effect is negative in sixteen years out of seventeen; the exception, 2015, is flat.

6. The yen in the US session

6.1 A trading rule and its cost

To measure the effect in tradable terms, I use the simplest rule that captures it: sell the yen future (6J) at 11:00 and buy it back at 16:55, just before the CME's daily break, every trading day. One contract is worth ¥12.5 million, about $85,000 to $110,000 over the sample. Table 4 shows the result under three cost assumptions, and on spot data before the futures sample.

Table 4. Selling the yen from 11:00 to 16:55 New York time

Sample and cost	Return per day	Winning days	Annual Sharpe	P&L, one contract
6J futures 2016–2026, no cost	+2.03 bp	54.6 %	1.37	+$52,600
6J futures 2016–2026, 1 tick + $5	+0.90 bp	52.7 %	0.61	+$22,300
6J futures 2016–2026, 2 ticks + $5	+0.28 bp	51.4 %	0.19	+$5,500
Spot USD/JPY 2012–2015, mid-quote	+2.34 bp		1.52	
Spot USD/JPY 2012–2015, quoted bid–ask	+1.55 bp		1.01	
Spot USD/JPY 2012–2015, bid–ask + 0.5 bp	+1.05 bp		0.68	

Notes: 2,635 trading days for the futures. A positive return means the yen fell. The median quoted spot spread is 0.67 bp.

The gross effect is about two basis points a day, and costs take more than half of it. One extra tick per round trip divides the net result by three: the whole question is one of execution. At standard cost, seven years in eleven are positive, the worst drawdown is about $7,100 per contract (2020–2022) and the longest one lasts about two and a half years. Fridays are weak (−0.9 bp net) and the other days range from +0.8 to +1.8 bp, a difference that is not significant.

6.2 What the effect is not

Several mechanical explanations can be ruled out. The effect is not an artefact of futures carry or rolls: spot shows the same move (−2.0 bp over 2010–2015), and the spot swap is paid at 17:00, outside the window. It is not a general rise of the dollar: over the same hours the euro gains +0.41 bp against the dollar, and a regression of the yen's return on the euro's and on the S&P 500 future's leaves a constant of −2.05 bp (t = −5.7). It is not, on average, the US equity market either. The yen does fall more on days when equities rise (−3.6 bp, against 0.0 on days when they fall), as expected from its role as a funding currency, but equities explain only about −0.23 bp of the mean (correlation −0.15). Finally, it is not just the trend of the yen. Over 2016–2020 the yen rose by +0.86 bp a day over 24 hours while falling by −1.17 bp in the US window, which means it rose by about +2 bp during the other hours: an intraday rotation. Over 2021–2026 the yen fell by −3.58 bp a day, of which −2.65 bp, about three quarters, in the US window.

6.3 The shape of the drift

A small average can come from a steady drift or from a few large days. Table 5 shows that it is the former. The median equals the mean, and symmetric trimmed means stay at about +2.1 bp whatever the share trimmed. The quantile-shift profile, $(q(p) + q(1-p))/2$, is flat between +1.7 and +2.2 bp for p from 0.05 to 0.45: calm and volatile days are shifted by the same amount. The ten best and ten worst days cancel out (+1,342 against −1,343 bp), so the total of +5,463 bp comes from the rest.

Table 5. Distribution of the gross US-session return from selling the yen, 6J 2016–2026

Statistic	Value
Mean (Newey–West t)	+2.03 bp (4.9)
Median	+2.19 bp
Trimmed means, 1 / 5 / 10 / 25 % each side	+2.07 / +2.05 / +2.10 / +2.11 bp
Standard deviation; skewness; kurtosis	23.6 bp; −0.09; 14.5
Positive days	54.6 %
Mean without the 1 % / 5 % most extreme days	+2.01 (t 5.5) / +1.71 (t 5.4)
Years with a positive mean	11 of 11 (from +0.47 to +4.42 bp)
Quarters with a positive mean	36 of 43

Notes: 2,694 trading days. Spot data for 2012–2015 give the same shape: mean +1.55, median +1.70 bp.

Quarterly means are even less dispersed than chance would imply ($\chi^2 = 23.7$ with 42 degrees of freedom against a constant mean, $p = 0.99$; variance ratios of 0.69 at 21 days and 0.66 at 63 days), because of a slight negative autocorrelation (−0.07 at the first lag): a large fall of the yen is often partly reversed the next day. The effect is stronger over 2022–2026 (+3.1 bp on average against +1.2 over 2016–2021), but present throughout.

6.4 Conditions

I examined 34 cuts of the sample, fixed in advance. None of the conditions known before 11:00 concentrates the effect: weekday, month-end, volatility measured beforehand and the yen's move earlier in the morning all give differences that are not significant. The only clear contrast is with stress on the day itself. When the S&P 500 future falls during the window, the yen does not fall (+0.01 bp, against +3.45 bp when equities rise), and in the most volatile quintile of the window the drift turns negative (−1.20 bp). Only this last contrast survives Holm's correction across the seven F-tests (adjusted $p = 0.03$). It fits the yen's role as a safe haven, but it is only known after the session. Days of FOMC decisions go the other way (−4.4 bp, $t = -1.1$, 90 days) and days of FOMC minutes are strong (+8.0 bp, $t = 2.4$, 86 days); with 34 cuts, neither counts.

6.5 A reversal filter

The slight negative autocorrelation suggests that the drift might be stronger after the yen has risen. I tested eight filters that use only information available before 11:00, fixed and committed before computation (Table 6): four based on volatility and four based on the yen's recent rise. A "day" for the reversal filters runs from 16:55 to 16:55.

Table 6. Filters for selling the yen in the US session, 6J February 2017 – August 2026, net of costs

Filter (trade only if…)	Days	Net, traded	Net, skipped	t (diff.)	Daily Sharpe	Years +
No filter (reference)	100 %	+0.90 bp			0.64	
V1 20-day volatility above median	42 %	+1.31	+0.60	0.75	0.52	8/10
V2 20-day volatility in top tercile	30 %	+1.85	+0.50	1.19	0.58	7/10
V3 hourly ATR above median	43 %	+1.58	+0.39	1.25	0.62	6/10
V4 hourly ATR in top tercile	29 %	+2.20	+0.36	1.59	0.67	7/10
R1 yen rose in the previous US window	43 %	+1.98	+0.09	2.07	0.91	7/10
R2 yen rose over the three previous US windows	40 %	+1.67	+0.38	1.42	0.78	8/10
R3 yen rose over the previous day	46 %	+1.97	+0.00	2.16	0.90	8/10
R4 yen rose over the three previous days	46 %	+2.09	−0.11	2.43	1.00	9/10

Notes: 2,425 days, the common sample after 250-day rolling windows. Net returns in basis points per day traded. The daily Sharpe ratio counts zero on days without a position, so it compares directly with the unfiltered rule. ATR is the average true range.

Volatility filters do not help. Volatile days earn more in basis points, but this is the same drift at a larger scale, with proportionally more noise: the daily Sharpe ratio does not rise, and the difference with skipped days is never significant. Reversal filters, in contrast, all point the same way. After the yen has risen, selling it in the US session earns about +2 bp net (+3.1 bp gross, one and a half times the unfiltered +2.0 bp); on the other days it earns nothing after costs. Trading fewer than half of the days gives the same total profit, and the daily Sharpe ratio rises from 0.64 to about 1.0. The best variant, R4, has a Newey–West t of 3.5, balanced halves (+1.80 and +2.38 bp), nine positive years out of ten, and a Deflated Sharpe Ratio (Bailey and López de Prado, 2014) of 0.98 with 42 trials. This echoes Piccotti (2014), who describes a reversal of 24-hour dollar–yen returns during New York trading. The choice of R4 among the four reversal filters was not written down in advance, which flatters it somewhat.

Two further attempts to concentrate the drift failed. Within the day, nothing happens before 11:00, and although the last hour carries about half of the gross drift with little noise, a fixed cost per trade makes shorter windows unprofitable: no window among 48 beats the full 11:00–16:55 window. Across days, the largest rises of the yen are reversed most (the top decile of the three-day rise earns +5.5 bp net, t = 3.0), but the relation is irregular from one decile to the next and does not pass the threshold of interest fixed beforehand, nor the Deflated Sharpe Ratio with 140 trials (0.40).

7. Interpretation

The yen's US-session weakness runs against the domestic-hours pattern of Breedon and Ranaldo (2013) and Zhang (2018), who leave it unexplained. The evidence above narrows the possibilities without settling them. Three explanations remain, none of which can be tested with the data used here. The first is portfolio or dealer inventory flows: net purchases of dollars against yen concentrated in US hours, which would require flow data to verify. The second is the carry trade:

yen-funded positions opened or rebuilt during US hours, consistent with the stronger fall on days when US equities rise and with its disappearance on risk-off days. The third is US interest rates: the dollar–yen rate follows Treasury yields, and if yields rise on average during US hours the yen falls mechanically; testing this would require Treasury futures. The reversal result adds one element: when the yen has just risen, typically in risk-off episodes or on repatriation flows, the US session undoes part of the move.

8. Tests still to come

Every observation of the yen futures up to 28 August 2026 has now been examined, and every result in Section 6 is therefore exploratory. Two pre-registered tests were written and committed before any later data were downloaded. The first tests the unfiltered rule of Table 4: H1, the gross return is positive (the phenomenon), and H2, the net return is positive (the strategy), on data after 28 August 2026 only. It will be read on 29 February 2028 and 31 August 2029, with success only if $z \geq 3.0$ (a Haybittle–Peto boundary; Haybittle, 1971; Peto et al., 1976), and finally on 29 August 2031 with Holm's correction. Its power is low for the net hypothesis: at the historical effect size, about 18 % at the first read and 37 % after five years; the gross hypothesis is the real test. The second pre-registration tests the reversal filter R4 in two steps: first on dollar–yen spot data from 2003 to 2009, which have never been downloaded and are reserved for a single read, then on new futures data at the same dates as the first test. No price, return or signal on the yen after 28 August 2026 is computed before those dates.

9. Limitations

The yen effect is small relative to costs, and its net value depends on execution more than on anything else. All the evidence about its shape, its conditions and the reversal filter comes from data that have been looked at repeatedly; only the reserved samples can confirm it. The out-of-sample period mixes futures and spot data across currencies, which complicates comparisons with the original paper. The watch-list cells of Section 5 were found in a scan of 28 cells and do not pass correction. Finally, the economic mechanism behind the yen's weakness is not identified.

10. Conclusion

The session pattern documented by Zhang (2018) replicates closely on its own sample and disappears afterwards: currencies no longer weaken during their own hours, and the US-session strength of the euro and the pound has faded. The single exception in the original table has become the most robust result. The yen falls against the dollar during US trading hours by about two basis points a day, a steady shift of the whole distribution, present in nearly every year, not explained by the dollar, by equities or by futures carry, and absent on days of market stress. It is strongest after the yen has risen. Whether it survives is now a question for new data, which have been set aside under a fixed plan.

References

Bailey, D. H. and López de Prado, M. (2014). The deflated Sharpe ratio: Correcting for selection bias, backtest overfitting and non-normality. *Journal of Portfolio Management*, 40(5), 94–107.

Breedon, F. and Ranaldo, A. (2013). Intraday patterns in FX returns and order flow. *Journal of Money, Credit and Banking*, 45(5), 953–965.

Haybittle, J. L. (1971). Repeated assessment of results in clinical trials of cancer treatment. *British Journal of Radiology*, 44(526), 793–797.

Holm, S. (1979). A simple sequentially rejective multiple test procedure. *Scandinavian Journal of Statistics*, 6(2), 65–70.

Newey, W. K. and West, K. D. (1987). A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix. *Econometrica*, 55(3), 703–708.

Peto, R., Pike, M. C., Armitage, P., Breslow, N. E., Cox, D. R., Howard, S. V., Mantel, N., McPherson, K., Peto, J. and Smith, P. G. (1976). Design and analysis of randomized clinical trials requiring prolonged observation of each patient. I. Introduction and design. *British Journal of Cancer*, 34(6), 585–612.

Piccotti, L. R. (2014). Pricing errors and the geography of trade in the foreign exchange market. SSRN Working Paper 2446584.

Ranaldo, A. (2009). Segmentation and time-of-day patterns in foreign exchange markets. *Journal of Banking and Finance*, 33(12), 2199–2206.

Zhang, H. (2018). Intraday patterns in foreign exchange returns and realized volatility. *Finance Research Letters*, 27. SSRN Working Paper 3138756.

This paper is a research note. It is not investment advice. All figures are historical estimates under the assumptions stated above.