Cuemacro

Seeking the cues in macro markets

CLS trusted market solutions

Cuemacro

Saeed Amen
Quantitative Strategy
+44 20 3290 9624
saeed@cuemacro.com
@cuemacro
http://www.cuemacro.com

10 May 2019

Going with the FX flow

Using CLS flow data to understand and trade FX spot

In this paper, we discuss using CLS intraday hourly flow data to understand FX markets. Using two different trading strategies (daily and hourly) against a generic trend model, we find that CLS's flow data relating to funds and non-bank financial firms (NBFIs) tends to have a positive contribution to spot returns when viewed on an aggregate contemporaneous basis. By contrast, flows from buy-side and corporates tend to have a negative contribution. Later, we develop systematic trading rules for FX which use intraday hourly flow data from funds as an input between March 2013 and March 2018 as our in-sample period. Our daily flow based FX basket has annualised returns of 2.3% and risk adjusted returns of 0.63, and outperforms a generic trend basket which has returns of 2% and risk adjusted returns of 0.5 over the same period. We find that the drawdowns of an equally weighted portfolio of trend and daily flow based strategies were smaller, than any of the individual portfolios in isolation. Out-of-sample performance for the latter part of 2018 (April 2018 to October 2018) was slightly positive for the daily flow strategy, outperforming trend which was loss-making. Our hourly based flow trading FX basket in comparison, has risk adjusted returns of 0.81 and annualised returns of 1.6% in-sample. The risk adjusted returns of a combined daily and intraday hourly FX flow trading basket is 0.92.

Introduction

It has historically been challenging to understand FX volumes and flows in a comprehensive and timely manner, given much of the flow in FX tends to be fragmented across many different venues.

Figure 1: Daily flow, trend, daily flow + trend & daily + hourly flow baskets trading FX

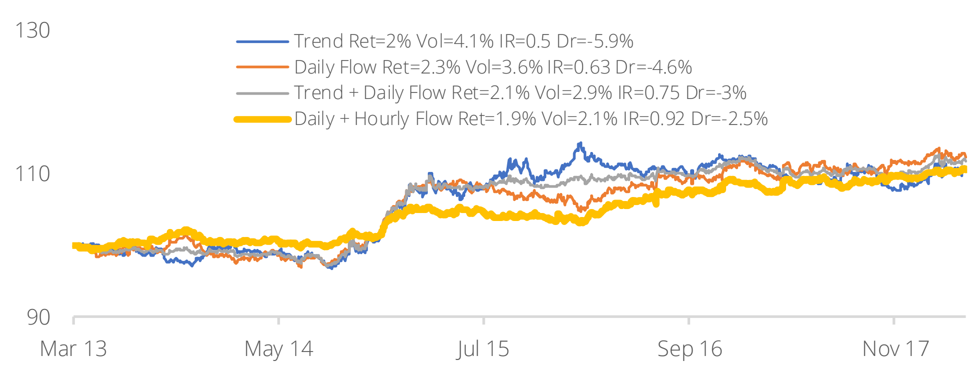

Source: Cuemacro, CLS, Bloomberg

The Bank for International Settlements (BIS) conducts a survey on FX volumes across the market for over 20 currencies, but this is produced once every three years. Certain platforms/exchanges also supply FX volume data, on a much more timely basis. However, it only covers their specific platforms or exchanges.

How can we get a much more comprehensive picture of spot FX flows, which covers a greater percentage of the FX spot market? CLS settles a large proportion of the FX spot transactions done globally. (Hasbrouck & Levich, 2018) estimate that CLS covers 37.2% of the global turnover of FX spot. We should note that if we specifically focus on those currencies which are settled by CLS this figure rises to over 50%.

(Hasbrouck & Levich, 2018) note that by contrast, EBS and Reuters combined covers 15.0%. They note that coverage of the FX spot market in the CLS data is greater, given it includes trades from EBS and Reuters as well as from other sources.

CLS now makes available flow data aggregated from its trade data collected from the market, which can give us a much clearer picture of the aggregate flow going through FX markets. We shall look at this data to understand patterns in FX flows and volumes. Additionally, we investigate the relationship between these flows and spot returns, and how CLS's flow data can be used to trade FX spot systematically.

We present the returns for a daily FX spot strategy based on flow alongside that of a generic trend model, as well a combined daily and hourly FX flow trading strategy, in Figure 1. We shall discuss all of these strategies in more detail later in the paper.

Describing CLS data

We first describe the flow datasets available from CLS which are aggregated and anonymized from spot transactions which are settled through the CLSSettlement system. First, there is the CLS- hourly order flow (HOF) dataset. This contains volume data by currency pair, aggregated hourly and delivered the next business day. The volume data is broken down by buy and sell volume (in base currency), which have matched counterparties. It is again classified for various types of price takers and price makers. There are also fields which denote the number of buy and sell trades. We can easily derive total volumes and the average trade volume from this data. Earlier CLS datasets consisted only of aggregated volume data. (Gargano, Riddiough, & Sarno, 2018) use CLS volume data to forecast FX returns and to create systematic FX trading strategies. They note that stronger reversals in spot are evident when FX volume decreases. (Ranaldo & Magistris, 2018) discuss using CLS volume to understand FX market illiquidity and also to understand the relationship between volume and volatility (vol).

However, the ability to discern the respective sizes of the buy and sell flows gives us considerably more information about the market, compared to looking at volume in isolation. In particular, we can now easily calculate the net flow.

Certain price takers in the dataset are designated as “corporate”, “fund” and “non-bank financial”. Another classification for price takers is “buy-side” which aggregates all of the above price takers and also non-market making banks. The price makers are denoted as “bank” or “sell-side”. Price makers are identified in the dataset by observing how connected they are using network analysis of the various counterparties. Transactions between market makers are omitted from the dataset. CLS-HOF is released with a lag of one day.

Secondly, there is the CLS-intraday hourly order flow (IDHOF) dataset, which is slightly different to the CLS-HOF dataset. It starts from September 2017. The main difference with CLS-HOF is that CLS-IDHOF is released typically 10 – 30 minutes past the hour, as opposed to next business day. Therefore, we shall be using CLS-IDHOF throughout this paper. It also has additional fields which show buy and sell volume, which has not been matched up to that point (i.e. CLS hasn’t received data about a transaction from both counterparties, only from one). If after a period trades remain unmatched they are removed, but this is a very small percentage of trades overall. We shall adhere to using the matched flow data in CLS-IDHOF, mainly given the dataset history is substantially longer (late 2012), up until March 2018, as our in-sample period across strategies. Later on in the paper, we shall use the period between April 2018 – October 2018 as our out-of-sample period.

Analysing CLS data on a daily basis

We first transform the original CLS-IDHOF intraday hourly dataset into individual time series for each currency pair, price taker, buy and sell volume. The original dataset is in London time, we then convert this into NYC time, because the daily spot data we shall be using is taken at NYC close, i.e. 17:00 ET time. Typically, most daily closing spot data is snapped at 17:00 ET. We then sum the hourly buy/sell volume data from 16:00 ET to 16:00 ET. This takes into account the small release lag[1] associated with the CLS intraday time series. If we collected data up to 17:00 ET, it would imply that we would need to trade it at the earliest at 17:30 ET.

With our daily buy/sell volume time series available, we can calculate net flow, as buy – sell volume for each currency pair and price taker. Before attempting to apply any sort of trading rule, we want to do some initial analysis to understand some of the dataset’s characteristics.

[1] CLS intraday data is delivered to clients within 30 minutes following the conclusion of the hour. Given that CLS data is now available intraday, we need to ensure that no “future” data is used, which would introduce look ahead bias.

Volume sizes

We first wish to understand the differences in volume sizes between the various price takers. We present our results for EUR/USD for average daily volume (in base currency) in Figure 2 shown below. Alongside this, in Figure 3, we plot the average absolute net flow. Our first observation is that the buy-side volume is by far the largest of any of the price takers. This shouldn't be a surprise, given that it is an aggregate of several different categories, as we mentioned earlier in the document.

However, on a net absolute basis, whilst the buy-side volume does appear to be the biggest category of price takers, as we might expect, it is far less than the overall volume. Hence, much of the buy-side trade consists of a very large two-way flow, i.e. not all buy-side participants are necessarily going the same way. By contrast, the absolute net flow of fund trades is relatively high as a proportion of their total volume. Hence, as a cluster they could exhibit more group-like behaviour.

We note that the corporate and non-bank financial flows both on a volume basis and an absolute net flow basis are much lower. Hence, we conclude that CLS's fund flow data is best to understand a clear group-like, directional behaviour in the FX market.

Figure 2: EUR/USD daily volume

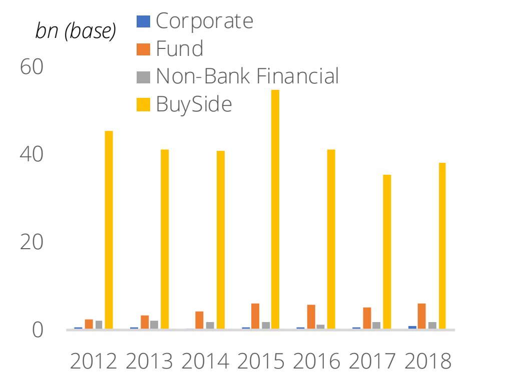

Source: Cuemacro, CLS

Figure 3: EUR/USD daily abs net flow

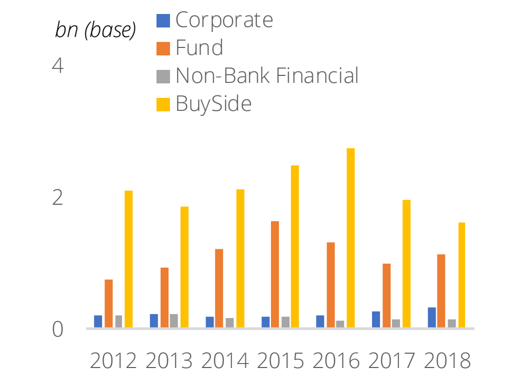

Source: Cuemacro, CLS

We also repeat the exercise for USD/JPY, plotting the average daily volume in Figure 4 and the absolute net flow in Figure 5. Our observations are very similar to those in EUR/USD. On a volume basis, the buy-side is by far the largest, followed by fund transactions. However, as with EUR/USD, once we look at the absolute net flow basis, we again observe that the buy-side is actually comparable in size to the flow from funds.

Again, this seems to suggest that the buy-side flow is more two-way, whereas as a group fund flow tends to be more directional when viewed collectively.

Figure 4: USD/JPY daily volume

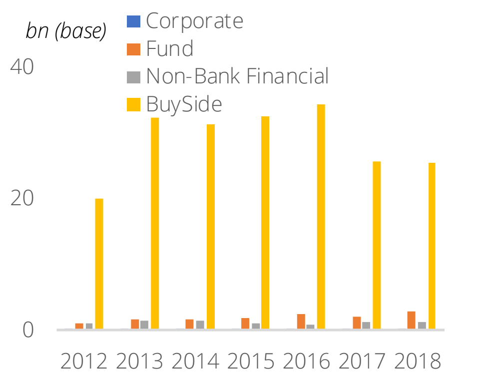

Source: Cuemacro, CLS

Figure 5: USD/JPY daily abs net flow

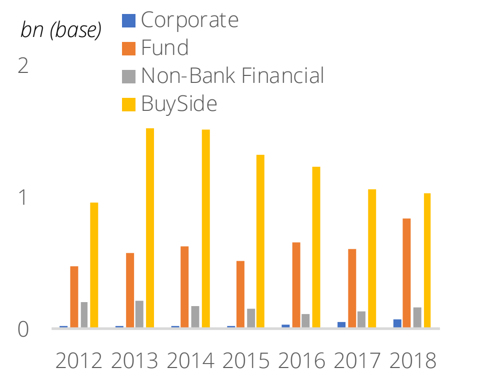

Source: Cuemacro, CLS

Price takers and FX spot

The next step is to understand the relationship between FX spot returns and volumes and the various price takers. We denote daily spot returns as our *y* variable and flow from buy-side, corporates, non-bank financials and funds as our *x* variables. We then conduct a multiple linear regression between our daily flows and spot returns for each of our currency pairs. In Figure 6, we report the T statistics from these regressions.

This can enable us to understand how flows as a whole are related to spot, rather than doing regressions or correlations in isolation on a specific type of flow, which are less likely to be able to explain spot moves.

Figure 6: Multiple regressions between spot returns and net flow

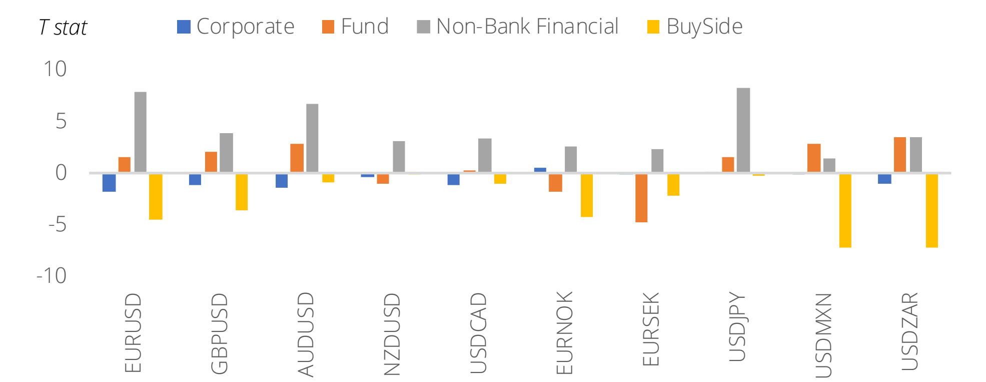

Source: Cuemacro, CLS, Bloomberg

If we examine the T statistics derived from the coefficients of these various multiple regressions, we can note several things. Generally speaking, the coefficients tend to be positive for fund and non-bank financials. Conversely, for buy-side and corporates the coefficients are negative. Admittedly, the T statistics are not always statistically significant. However, we can infer from these results the general observation that transactions by funds and non-bank financials tend to have a positive contribution to spot returns. By contrast, transactions by corporates and buy-side tend to have a negative contribution to spot returns, when aggregated on a daily basis.

There are some caveats to this sort of analysis, namely that the constant term from most of these regressions, is still fairly large (although not statistically significant), suggesting that there is still some price action which cannot be explained purely by flow data alone at this daily time scale. In addition, some flow data can be sparse, making such results more difficult to interpret.

Typically, corporates are likely to be making FX transactions to hedge and for foreign payments. In other words, FX is being traded as result of their primary business. For a corporate this could be for facilitating working in a foreign country.

By contrast, funds might be trading primarily from a speculative perspective, seeking to maximise profits from their trading, and not as a necessity to support business elsewhere. Hence, their profits are primarily tied to getting the future spot direction right.

For the buy-side category, it is an aggregate of many different flows (and also given it includes non-market making banks), so it might be trickier to explain very specific observations.

Comparing a daily flow based trading rule to trend following

We have established there is often a contemporaneous relationship between flow data and FX spot, based on hourly data, which has been aggregated into daily time series, when we view flow data in aggregate. The next step is to see if we can create profitable systematic trading rules for FX spot using flow data. In this section, we utilise the observations we made earlier to create systematic trading rules.

We noted that flow data related to funds tends to be more directional than that of buy-side data, which tends to have a net flow closer to zero. Intuitively, it should have a better relationship with the future direction spot. We have also focused on flow data from funds because on the whole it tends to be less sparse than for example corporate flow data across the various currency pairs. Hence, we should be able to create a generic trading rule to apply across a number of trading rules.

If we want to focus on sparser data, we would likely have to tailor a trading rule for each currency pair (whilst at the same time being careful to not overfit the model), or at the very least apply a trading rule which traded pairs less frequently (i.e. avoiding periods when there are low amounts of flow reported).

Our daily flow based trading rule, involves:

- buying a currency pair when fund flow is heavily positive, and then holding it until it turns more neutral
- selling a currency pair when fund flow is heavily negative, and then holding it until it turns more neutral

We create a standardised score for fund flow to trigger our trading rule, which we illustrate in Figure 7 for EUR/USD alongside the trading levels. We need to create a standardised metric to take into account the fact that the magnitude of flows can vary significantly between currency pairs and also over time. In this instance, spot has been reindexed from 100. We note that by and large, at least from this stylized example, there can often be some correlation between the flow score and spot.

Figure 7: EUR/USD index versus EUR/USD fund flow score

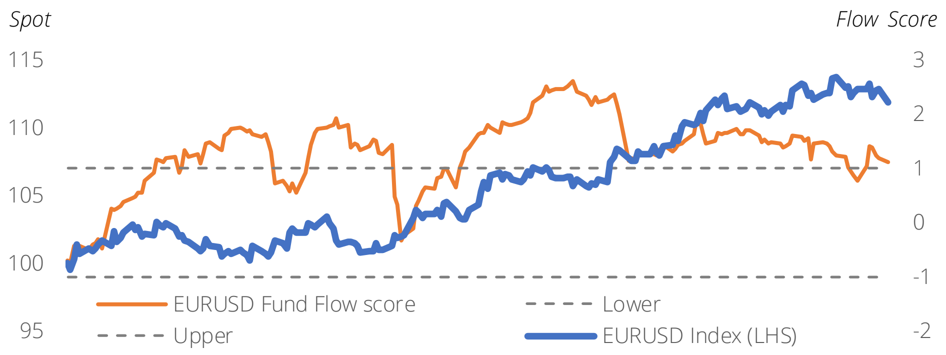

Source: Cuemacro, CLS, Bloomberg

Of course, this isn't the only approach we could have used. Other potential approaches can also include trying to create rules which attempt to fade extreme positioning (i.e. when there has been a significant amount of flow in a single direction over a long period of time). The rationale is that when market participants are heavily positioned, it increases the chances of a washout and a squeeze. In this situation, market participants attempting to square their positioning can cause large reversals in price. However, this is tricky to time, because positioning can remain extreme for a very long period. Hence, we need to be wary of attempting to fade extreme positioning prematurely. We note that our strategy is currency pair specific, rather than trading in a cross-sectional approach, which in FX is typically done with strategies as carry trades. We could also have tried to create a trading rule based on temporary divergences between spot moves and flows.

Hence, in our case, the flow based signal for each currency pair is calculated independently of any of the flow signals from any other currency pair. (Gargano, Riddiough, & Sarno, 2018) apply a cross-sectional approach to trading, based on CLS's volume data, essentially ranking comparative flows across the trading universe.

As a comparison to our daily flow based strategy, we also report the returns of a generic trend rule, which examines trends at horizons of 3 months, 6 months and 12 months, which is similar to trend horizons followed from (Hurst, Ooi, & Pedersen, 2013). This generic trend rule represents the typical type of strategies that CTAs might employ. We also apply overall a 10% volatility target for each currency pair for all trading rules, which is changed once a month. For example, if a currency pair has a 5% realised volatility over the past month, then we apply leverage of 2 times, so it can (hopefully) meet a 10% volume target. We apply a maximum leverage constraint of 5 times that of the original notional.

Applying a volatility target for each currency pair is helpful given there is a disparity in the volatility of the various currency pairs in our universe. We have focused on the primary currency crosses[2] available in the CLS data sets for the main developed markets and for some of the emerging markets.

We report risk adjusted returns for both trading rules in Figure 8. Our return history is from March 2013 till March 2018. We include both carry and transaction costs (2bp bid/ask) in our analysis.

Figure 8: Risk-adjusted returns for trend and daily flow based strategies

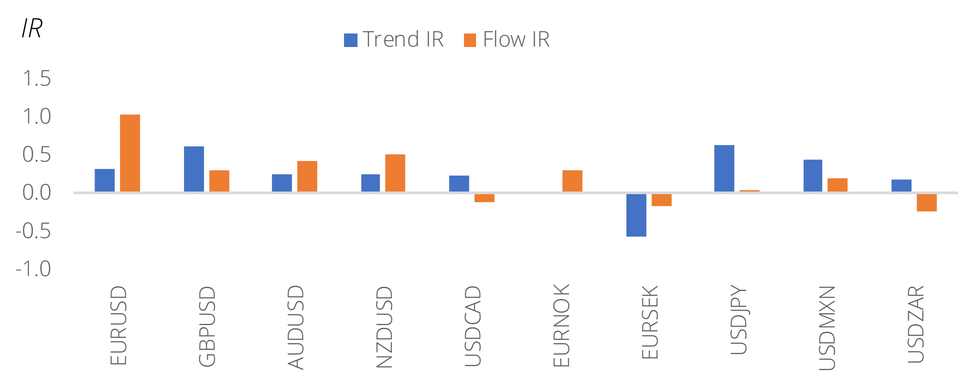

Source: Cuemacro, CLS, Bloomberg

Our daily flow based trading rule is profitable in the majority of crosses, although it is loss-making historically in USD/CAD, EUR/SEK and USD/ZAR. The trend following strategy seems to be profitable in the majority of cases, other than EUR/SEK. In Figure 9, we plot the turnover by currency pair for both trend and daily flow. On average the turnover is 32 for trend and 10 for the daily strategy per year (i.e. we execute trades equivalent to the total of 10 times the notional per year in the daily flow strategy).

[2] Typically, this will be the USD crosses, except in certain situations such as with the Nordic crosses, where they will primarily be traded against EUR. We have not included USD/CHF, given it was a heavily managed currency for a large part of our sample.

Figure 9: Turnover statistics for trend and daily flow based strategies

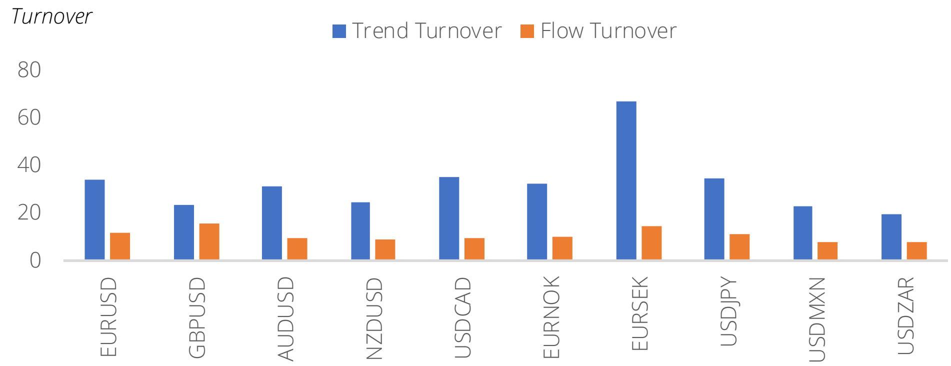

Source: Cuemacro, CLS, Bloomberg

Trading a daily FX flow basket

In Figure 10, we present the cumulative returns for a basket of all the currency pairs in our trading universe from earlier, for our trend following and daily flow based rules, as well as an equally weighted portfolio of each.

Our first observations, from a cursory glance is that there does appear to be some positive correlation between trend and CLS's flow data, as we might expect, given that our trading rule "goes with the flow". However, the size of the correlation is less than you might expect at around 11% when computed across daily returns. In order to have a negative correlation, we would need to try to construct mean-reversion style strategies which by construction fade flow or price moves.

Despite having (a small) element of correlation, there is some disparity in performance. Whereas the trend following strategy has risk-adjusted returns of 0.5 and annualised returns of 2%, the flow-based strategy has risk-adjusted returns of 0.63 and annualised returns of 2.3%.

Furthermore, when combining the two portfolios, we note that the overall peak-to-trough drawdown of 3% is less than either the trend or flow portfolio independently. This suggests that investors who are already exposed to trend following strategies in FX could benefit from adding an additional CLS's flow data based strategy to their portfolio, even if there is some element of positive correlation between the two strategies.

Figure 10: Daily flow and trend returns

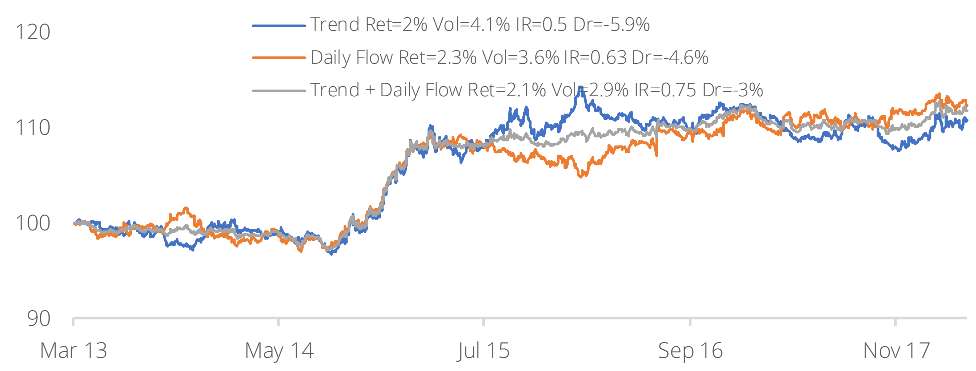

Source: Cuemacro, CLS, Bloomberg

Another important factor is to understand how stable the returns are throughout our sample. In Figure 11, we show the annualised year-on-year returns for each year during our in-sample period (returns for 2018, are annualised for the period up to April 2018).

We note that the trend model performed very strongly in 2014 and 2016. In 2018, both models performed strongly. We note that the daily flow based strategy is profitable in every year of our sample, other than 2013.

Figure 11: Year-on-year returns

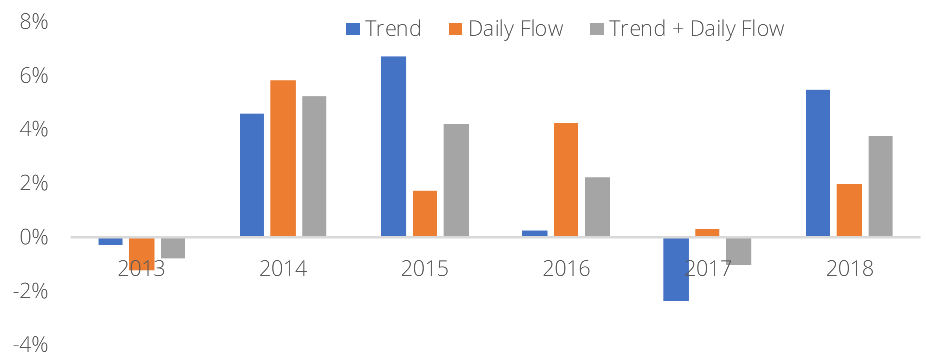

Source: Cuemacro, CLS, Bloomberg

Behaviour of FX spot volume on an intraday basis

So far, we have primarily examined FX flow data from a daily perspective and used that as a basis for creating daily systematic trading rules. However, the CLS dataset consists of hourly data. Hence, we can also use it to understand intraday behaviour, specifically around patterns in volume trading. This can also be useful from an execution point of view, to understand which times of the day liquidity is likely to be most concentrated, and how large our trades are compared with overall volume.

In Figure 12, we plot the total hourly volume for EUR/USD during the period from September 2012-March 2018 taking an average by year. We find that whilst overall volume does differ somewhat each year, the pattern throughout the day is relatively persistent, namely that volume is highest during the London afternoon (when New York is also active) and also during London morning (when Asia also has some activity).

Figure 12: Intraday hourly volume in EUR/USD, aggregated by hour of day and year

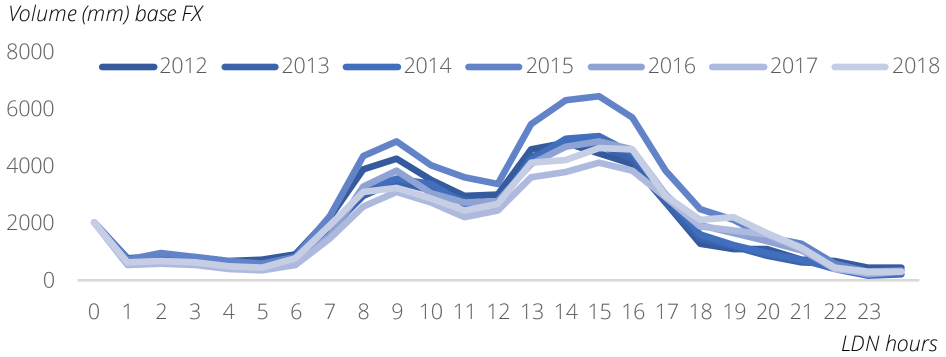

Source: Cuemacro, CLS

In Figure 13, we plot the intraday 5 minute rolling volatility for EUR/USD, USD/JPY and GBP/USD during the same in-sample period. Alongside that we plot the average EUR/USD hourly volume from CLS. We note that there does appear to be some correlation between intraday volatility and volume. As with volume, intraday volatility tends to be higher during the overlap between London afternoon and New York morning (and to some extent when London comes in during the morning with the Asia overlap). We note very specific spikes in intraday volatility, in particular at 13:30 LDN and 15:00 LDN, which often corresponds to US economic data events. There is another distinct spike at a 16:00 LDN, during the WMR fixing. The spike at 09:30 LDN seems unique to GBP/USD, which is likely related to UK economic data releases, many of which are released at this time of day throughout the month.

Figure 13: Intraday volatility patterns in EUR/USD, USD/JPY and GBP/USD with EUR/USD hourly volume

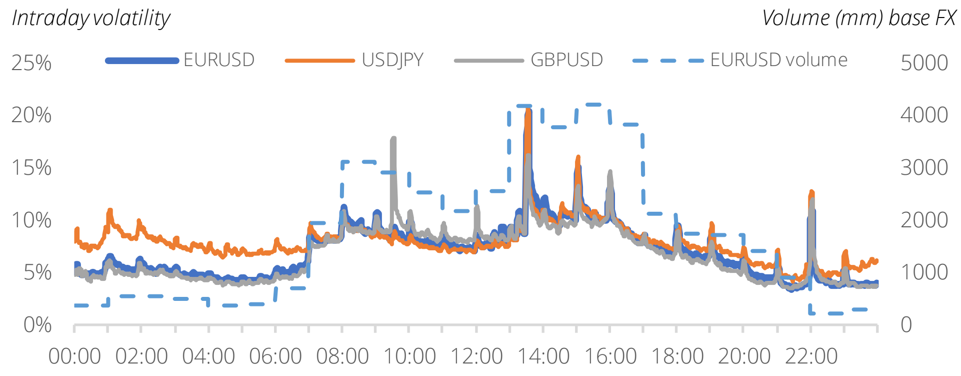

Source: Cuemacro, CLS, Bloomberg

We earlier noted that on a volume basis, the buy-side category dominates, which isn't surprising given that it's aggregated from several categories. Hence, if we do an analysis by time of day, any observations we might make are likely to be heavily skewed by buy-side transactions. In Figure 14, we instead look at each group of price takers separately for EUR/USD, plotting the average hourly volume traded by each group, for each hour of the day. Given how big buy-side volume is versus the other price takers, we have plotted that on its own axis (primary y-axis). All other price takers are plotted on the secondary y-axis.

The buy-side volume clearly matches the aggregated volume from Figures 12 and 13, as we might expect. The non-bank financial flow seems more stable during the day, but still exhibits increases the peaks in the overlapping sessions (i.e. earlier London morning and London later afternoon). The volume in corporate flow is much more concentrated around specific hours. The volume related to funds is spread mostly during the hours when New York is trading, contrasting somewhat from the buy-side which is mostly spread during London hours.

Figure 14: Average hourly volume by price takers in EUR/USD

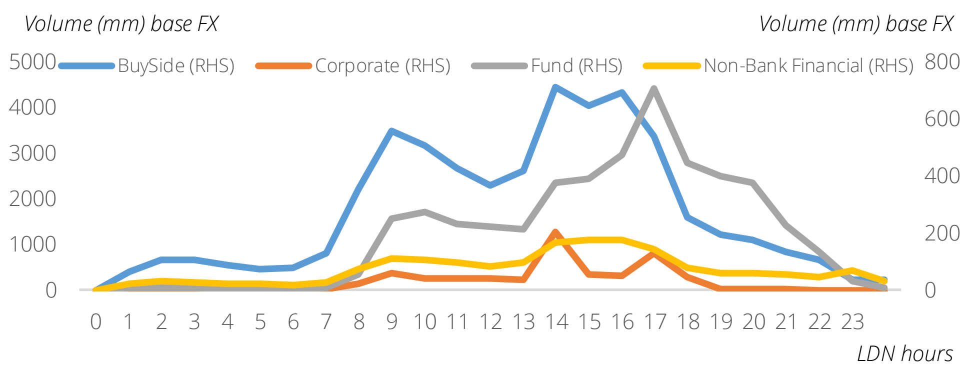

Source: Cuemacro, CLS, Bloomberg

However, we need to be aware that volume is not necessarily the whole story, when it comes to understanding price impact and the relationship with spot. We noted earlier that when we look at net flow on an absolute basis, the difference between buy-side and fund trades is actually not that much. This suggests that a lot of buy-side volume is much more two-way, as a whole. By contrast the volume of fund trades is not hugely different from the absolute net flow, which suggests that , as a group, funds tend to be doing similar trades.

Hourly based FX flow trading strategy

So far, we have examined flow based trading strategies that use CLS's data aggregated and anonymized on a daily basis and trading at New York close. However, given that we also have access to intraday hourly CLS Flow data, it is also possible to create higher frequency trading strategies, provided we put an appropriate signal delay[3] on our signal to account for the slight delay in the publication of the data.

Our approach is relatively similar to what we did for daily fund flow data, although speeding up the trading frequency. Namely, we:

- buy spot when net fund flow data is heavily positive
- sell spot when net fund flow data is heavily negative
- flat exposure otherwise

Given it is a higher frequency strategy we shall restrict ourselves to the most liquid currency pairs only, EUR/USD, USD/JPY, GBP/USD and AUD/USD, so we can assume tighter spreads (0.5bp bid/ask) in our backtest.

For simplicity, we shall assume that each currency pair is weighted by notional as opposed to volume weighting as we did previously and we have omitted carry. In Figure 15, we present the risk adjusted returns for each currency pair using this hourly based flow strategy over a sample from October 2012-March 2018. We find that in each currency pair the strategy is profitable. We also show the turnover per currency pair, which is on average 86 times the notional size. As we might expect the hourly flow trading rule has a higher frequency than the daily flow trading rule.

We also present an equally weighted portfolio of each of the strategies in Figure 16. The hourly flow based portfolio has risk adjusted returns of 0.81 and annualised returns of 1.6% and annualised volatility of 2% during this in-sample period. The volatility is lower than our daily flow basket, given that the basket tends to be underinvested more of the time.

[3] We shall be using a 1 hour delay to account for this in our backtest

Given the relatively low level of volatility for the strategy, in practice, we have quite a bit of leeway to leverage the strategy and increase returns, without having very high levels of volatility.

Figure 15: Hourly returns & turnover

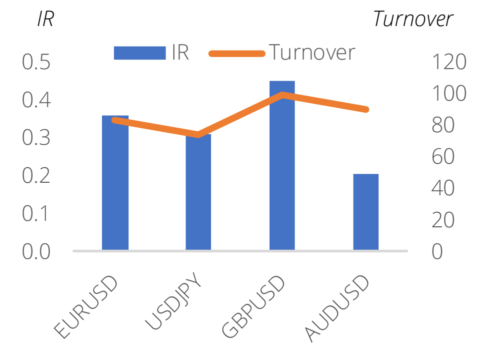

Source: Cuemacro, CLS, Bloomberg

Figure 16: Portfolio of hourly flow

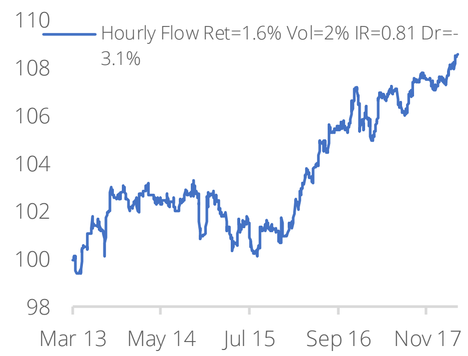

Source: Cuemacro, CLS, Bloomberg

The correlation between our hourly and daily flow baskets in the full sample (using daily returns) is around 5%, which is perhaps less than we might expect. One explanation is that whilst both trading rules "go with the flow", the different time scales can explain a difference in correlation. Indeed, this is not just the case with flow data. There is unlikely to be much correlation between high frequency "momentum" strategies that look at price data, and very long term trend following strategies which also have price data at an input. The frequency at which data is sampled can affect the final trading signals significantly.

Given the relative lack of correlation, it seems reasonable to create a combined basket of our daily and intraday hourly flow based FX baskets. In Figure 17, we present the returns for this combined basket, comparing against daily and hourly flow strategies alone between March 2013 – March 2018. We note that purely looking at returns the daily strategy appears to do best, with annualised returns of 2.3%. However, if we examine risk- adjusted returns, we note that the combined strategy has both the highest risk- adjusted returns of 0.92, and also the smallest drawdowns of 2.5%.

Figure 17: Combining hourly and daily flow strategies

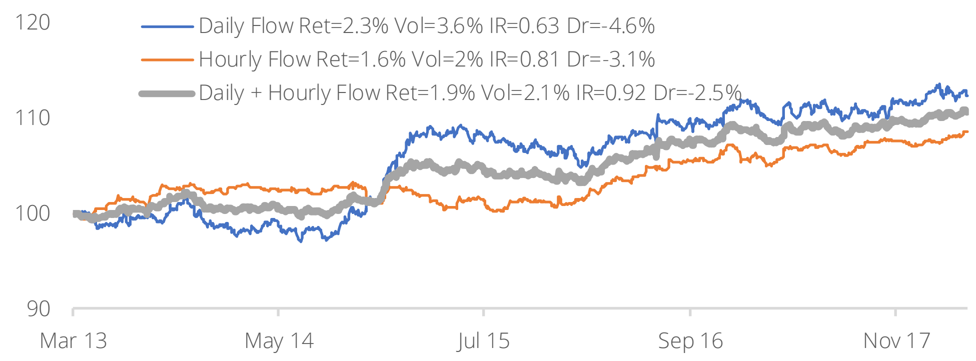

Source: Cuemacro, CLS, Bloomberg

Out-of-sample performance for daily flow, hourly flow and trend strategies

In this section, we examine out-of-sample performance, for the period from April 2018-Oct 2018 for our flow based trading strategies (both daily and hourly) and also for our generic trend basket. We present the returns in Figure 18. We observe that the trend is loss-making over the period and is also the worst performer. This contrasts with the daily flow strategy which is slightly positive over the same period. The out-of-sample return for the hourly flow strategy is slightly negative. Obviously, we need to take into account that our out-of-sample period is relatively short.

Figure 18: Out-of-sample performance Apr 2018-Oct 2018

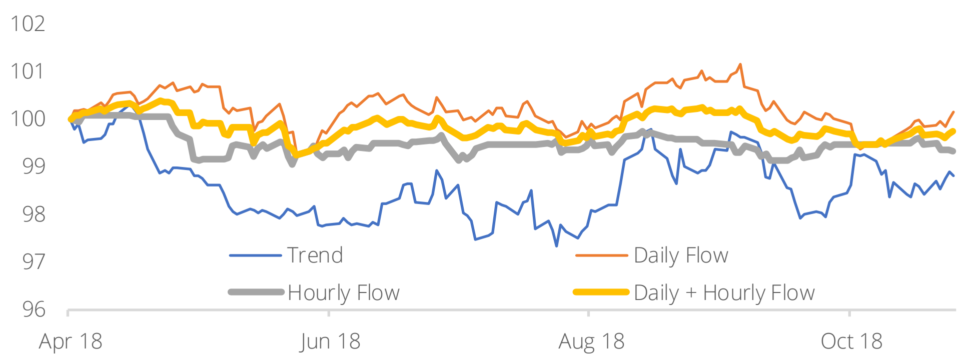

Source: Cuemacro, CLS, Bloomberg

Conclusion

We discussed the relationship between CLS FX flow data and FX spot. We noted that when aggregated on a daily basis in a multiple linear regression, typically fund flow and non-bank financial flow tended to have a positive contribution to FX spot returns. Conversely, buy-side and corporate flows tended to have a negative contribution to spot returns. These results seem fairly intuitive, although there is the obvious caveat that these were contemporaneous regressions and that in isolation flows did not explain entirely spot moves.

Buy-side and corporate flows are likely to be motivated by other transactions (ie. the currency flow is a secondary concern of this group). This contrasts with flows from fund accounts which are likely to be more speculative in nature.

We later used these observations to create trading rules using CLS FX flow data from fund accounts to trade FX spot, which had been aggregated into daily data. The idea was generally to follow these daily flows, when they were large and then to cut them when flows became neutral. We used a standardised metric for measuring the flow for each currency pair. For example, if funds were heavily buying EUR/USD as indicated by the flow data, the trading rule would also suggest buying. Conversely, heavy selling by funds would trigger a sell in the model.

Using these flow based rules, we created a daily trading FX basket which had risk adjusted returns of 0.63 and annualised returns of 2.3% between 2013 and 2018. It outperformed a generic trend following strategy over the same period (which by comparison has risk adjusted returns of 0.5). On an out-of-sample basis during the later part of 2018, the daily flow model also had slightly positive returns, outperforming trend which was negative over the same period.

Later, we also used the hourly data directly, to understand patterns in hourly volume data. We noted that times of high volume tend occur at times of high volatility (and conversely, low volume around times of low volatility). We also saw that patterns in the volumes of the various price takers did differ somewhat. The data on buy-side trades suggested that their volume was concentrated mostly during the overlap between London afternoon and New York morning trading hours (and to some extent during the overlap between London morning and Asia). Fund data also exhibited a lot of volume during London afternoon timing, however, there was also significant amount of volume during later New York hours.

We also constructed an hourly based flow trading FX basket which had risk adjusted returns of 0.81 and annualised returns of 1.6% from March 2013 – March 2018. Out-of-sample returns were slightly negative for this intraday strategy, but still better than those of the generic trend following The use of CLS's intraday hourly flow data assisted in improving the trading strategies and optimizing returns, and helped to improve the risk adjusted returns compared to running trend alone.

Our hourly and daily FX flow baskets have very little correlation. Hence, we decided to create a basket that combined both daily and hourly FX flow strategies. The resulting basket had risk adjusted returns of 0.92 in the in-sample period between March 2013 – March 2018, higher than any basket in isolation.

Our approach has been to apply the same trading rules for each currency pair in our trading universe. However, one area which is likely to be fruitful for the future is looking at flow data on a more currency specific basis. This would involve trying to understand how the various idiosyncrasies of flows in each currency pair can be traded. We of course need to be careful that we do not simply "fit" each trading rule for each currency pair, without having a fundamental understanding of why flows are different.

It seems intuitive that flows can vary between currency pairs. The flows for example in currency pairs such as NZD/USD are likely to be very different from those in EUR/USD, given the big differences in the underlying economies and capital markets of the various countries.

As well as drilling down into the data on a more currency pair specific basis, another approach which is likely to show some promise is augmenting CLS's flow data with other datasets. We already showed an aspect of this by combining our daily flow based rule with trend strategies which use CLS's price data as an input. However, we could also try using datasets derived from the FX option market such as risk reversals, which can give an indication of sentiment or news based sentiment.

It might also be interesting to have the ability to have finer granularity of the dataset, for example to be able to differentiate between different types of fund (eg. pension fund, hedge fund etc.)

Bibliography

Gargano, A., Riddiough, S. J., & Sarno, L. (2018, November 7). *Volume and Excess Returns in Foreign Exchange*. Retrieved from SSRN: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3019870

Hasbrouck, J., & Levich, R. M. (2018, March 18). *FX Market Metrics: New Findings Based on CLS Bank Settlement Data*. Retrieved from SSRN: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2912976

Hurst, B. K., Ooi, Y. H., & Pedersen, L. H. (2013, February 1). *Demystifying Managed Futures.* Retrieved from AQR: https://www.aqr.com/Insights/Research/Journal-Article/Demystifying-Managed-Futures

Ranaldo, A., & Magistris, P. S. (2018, November 22). *Trading Volume, Illiquidity and Commonalities in FX Markets.* Retrieved from SSRN: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3290504

Disclaimers

This paper offers research for idea-generation purposes and is not intended for trading purposes. It should not be used as direct advice on trading or investment and is not appropriate for the purposes of making a decision to carry out a transaction or trade. Nor does it provide any form of advice (investment, tax, legal) amounting to investment advice, or make any recommendations regarding particular financial instruments, investments or products.

Non-independent investment research disclaimer

This research paper has not been prepared in accordance with legal requirements intended to promote the independence of investment research. It is also not subject to any prohibition on dealing ahead of the dissemination of investment research. Cuemacro Limited, its affiliates or staff, may perform services for, solicit business from, hold long or short positions in, or otherwise be interested in the investments (including derivatives), of any issuer mentioned herein.

This paper is for the exclusive use of the recipient and may not be quoted, published, forwarded, copied or shared, in whole or in part, without prior written permission.

Any information or opinions in this material are not intended for distribution to, or use by, any person in any jurisdiction or country where such distribution or use would be unlawful. It is intended purely for the consumption of professional investors.

This paper is provided "as is," without any warranties of any kind. The user assumes the entire risk of any use made of this information. CLS Services Ltd and its affiliates ("CLS") and Cuemacro Limited do not warrant or represent that the document is accurate, complete, reliable, fit for any particular purpose or merchantable and do not accept liability for any act (or decision not to act) resulting from use of this document and related data. To the maximum extent permissible all express or implied warranties and other assurances by Cuemacro Limited. and CLS are hereby excluded, and Cuemacro Ltd. and CLS shall have no liability for the use, misuse, or distribution of this information. Past investment performance is no indication of future investment performance.

Cuemacro Limited, Level39, One Canada Square, Canary Wharf, London E14 5AB, UK
Tel +44 20 3290 9624, e-mail saeed@cuemacro.com
web http://www.cuemacro.com

Cuemacro Limited. is registered as a company in UK with company no. 10195390