May 2018

Improving Relevancy in FX Media Indicators with Supervised Machine Learning

State Street MediaStats' Indicators ® provide a timely measure of media sentiment drawn from hundreds of thousands of curated, unstructured data sources.

Over the past seven years, we have found that this source of data can generate timely insights into the performance of a range of investments.

This real-time pulse of broad media sentiment enables investors to anticipate and evaluate the impact of online chatter, including the impact of unusual media coverage and the degree of disagreement in opinion.

Executive Summary

In this White Paper, we introduce a modeling approach that improves the relevancy of media articles related to foreign exchange currencies and demonstrate that media indicators can be used to predict future currency returns. More specifically, using supervised machine learning techniques, we design a filter model to determine the probability that a news article is classified as relevant by a human examiner. We then apply the filter on our entire set of currency articles and measure its impact on performance.

Introduction

Natural language processing (NLP) is the discipline that studies how to enable machines to digest and interpret the language that people use—the natural language. Today, much of the information investors rely on to identify trends is conveyed through human speech and writing—e.g. media articles, online blogs and other types of social media. Extensive research has been conducted by practitioners and academics alike to extract insights on equities using NLP techniques. Yet, certain challenges remain in identifying systematically news that is relevant for foreign currencies (FX). For instance, a currency is both a financial asset and a unit of measurement—and therefore, it may be questionable whether articles discussing large development projects costing millions or billions of USD are relevant to the outlook for the Greenback. Other examples include crude NLP algorithms classifying news articles on a Chinese ballerina named Yuan as relevant for the Chinese currency. State Street MediaStats' state-of-the-art NLP solution is engineered to account for such noise and allows investors to parse through millions of news articles per year, extracting those trends that are most compelling with efficiency and objectivity. In this paper, we describe the process used to enhance our NLP algorithm using supervised machine learning and analyze the impact on predicting currency returns.

Data Description

We leveraged State Street MediaStats proprietary technology to collect an initial set of over 300,000 curated articles pertaining to 33 foreign currencies. The dataset covers the period between March 2013 and December 2017. The panel of sources from which articles are collected is large, allowing for a diversity of opinions; some 10,000 individual sources from 73 countries contribute to our panel. The 15 countries whose sources contribute the highest percentage of articles are presented in Figure 1 overleaf.

Figure 1: Percentage of Articles per Country Source

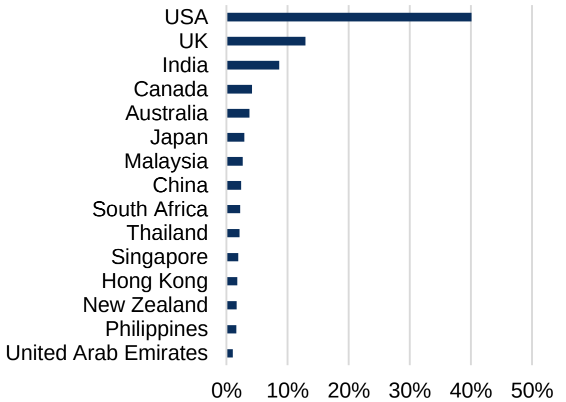

Source: State Street Global Exchange, MKT MediaStats (Mar 2013-Dec 2017).

The following currencies constitute our coverage universe: i) developed market currencies: AUD, CAD, CHF, DKK, EUR, GBP, ILS, JPY, NOK, NZD, SEK, and SGD; and ii) emerging market currencies: ARS, BRL, CLP, CNY, COP, CZK, EGP, HUF, IDR, INR, KRW, MXN, MYR, PEN, PHP, PLN, RUB, THB, TRY, TWD, and ZAR.

We then extracted a random subset of 2,753 articles from our initial set and asked humans to determine whether each article covered at least one of the following subjects: i) currency trading, ii) financial markets, or iii) macro economy. Humans determined that 31% of the articles in the random set are not related to either of these topics.

Methodology

In this study, we construct media sentiment indicators for FX using filtered and filtered-out articles and examine the impact of filtering on the predictive ability of media sentiment indicators on future currency returns. To accomplish this, we employ a three-step analysis: i) conduct supervised learning, ii) enhance our model, and iii) validate our results. In Step 1, we select a random subset of articles from our FX article database and ask humans to determine relevancy—a binary outcome of "Yes" or "No"—with respect to at least one of the following topics: currency trading, financial markets, or macro economy. In Step 2, we estimate the probability of receiving a "Yes" answer using various filters. Lastly, in Step 3, we apply the model to the entire sample of FX articles and compare correlations between article sentiment scores and contemporaneous returns with and without our model.

Supervised Machine Learning

We select a random subsample of articles from the sample of FX articles and ask graduate students in the disciplines of Finance and Economics to read them and submit a "Yes" or "No" answer to the question whether the article discusses at least one of the following subjects: currency trading, financial markets, or macro economy. Among the 2,753 random articles, 841 (31%) are assigned a "No" response—in other words, are classified as irrelevant by the human examiners.

The Model

In this step, we model the probability of receiving a "Yes" answer using various filters. Suppose if p is the probability of an article being relevant, then the estimated model is:

$$\log\left(\frac{p}{1-p}\right) = \alpha + \beta' X$$

The list of predictive variables X includes:

- The number of times a currency is mentioned in the text
- The number of times a currency is mentioned with words like "billion" or "million"
- A dummy variable indicating if currency name or code appears in title
- A dummy variable indicating if currency pair appears in the title
- A dummy variable indicating if title does not include words such as 'sport', and 'weather'
- A dummy variable indicating if article URL does not include words such as 'sport' and 'weather'

- The number of times a pair of currency is mentioned in the article
- Article source is classified as a Top International Source
- Article source is classified as a Top US Source
- Article source is classified as a Domestic Source
- Article source is classified as a Trading Focused Source
- Article source is classified as a Member of a Major Type

Using the estimated coefficients, we calculate the probability of each article being relevant. For each probability threshold level, each article is classified as relevant (non-relevant) depending on whether the calculated probability is above (below) our chosen threshold probability level.

Figure 2 shows the classification outcome of our model using 52% as the cutoff probability for our random sample. Among the 2,753 articles, 2,199 (80%) articles are considered relevant by our model at this probability level. We also present the interaction of classifications between the human readers and our model. If we rely on human judgement as the correct benchmark, then 1,726 (78%) articles are correctly classified as relevant, and 473 (22%) are incorrectly categorized as relevant. The false positive is therefore at 22%. Among the 554 articles categorized as non-relevant by our model, 66% are also classified as non-relevant by humans. The false negative is therefore at 34%.

Figure 2: Classification of the Random Sample

			Model	
Student		**All**	**No**	**Yes**
Total	#	2,753	554	2,199
			20%	80%
No	#	841	368	473
	% of All		13%	17%
	% of No / Yes	31%	66%	22%
Yes	#	1,912	186	1,726
	% of All		7%	63%
	% of No / Yes	69%	34%	78%

Source: State Street Global Exchange, MKT MediaStats (Mar 2013-Dec 2017).

Figure 3 below illustrates the clear tradeoff between the false positive (Type I error) and the false negative (Type II error) as we vary the probability cutoff level. As the probability threshold increases, fewer articles are classified as relevant, and the false positive decreases—accompanied by culling potentially relevant articles and thus increasing the false negative.

Figure 3: False Positive and False Negative as a Function of Probability Threshold

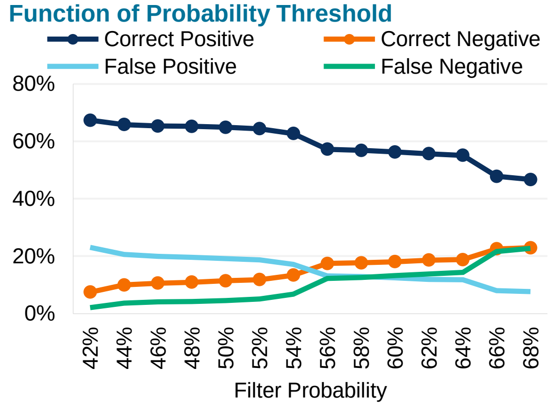

Source: State Street Global Exchange, MKT MediaStats (Mar 2013-Dec 2017).

Out-of-Sample Model Validation: Correlation with Contemporaneous Returns

In this section, we examine whether filtering impacts the correlations between article sentiment and contemporaneous currency returns. If our model does indeed successfully improve the relevancy of media articles, then the correlations between the sentiment of remaining articles and currency returns should be higher as the remaining filtered sample is potentially more informative. We apply our model on each of the articles in our entire sample of media articles related to both developed and emerging market currencies. Each sample is divided into two subsamples: a filtered—or relevant—subsample containing articles with relevancy probability above the selected threshold, and a filtered-out—or non-relevant—subsample including those articles with relevancy probability below the threshold. Figures 4 and 5 report the contemporaneous correlations between media article sentiment and currency returns for the filtered and filtered-out samples at different levels of probability thresholds. We show results for developed market (DM) currencies and emerging market (EM) currencies separately.

Figure 4: Correlations between Article Sentiment and Contemporaneous DM FX Returns

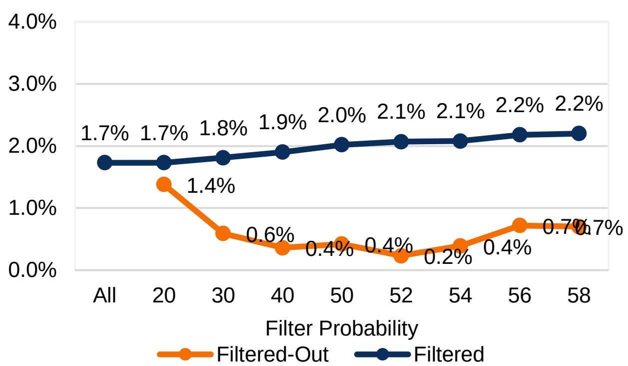

Source: State Street Global Exchange, MKT MediaStats (Mar 2013-Dec 2017).

Figure 5: Correlations between Article Sentiment and Contemporaneous EM FX Returns

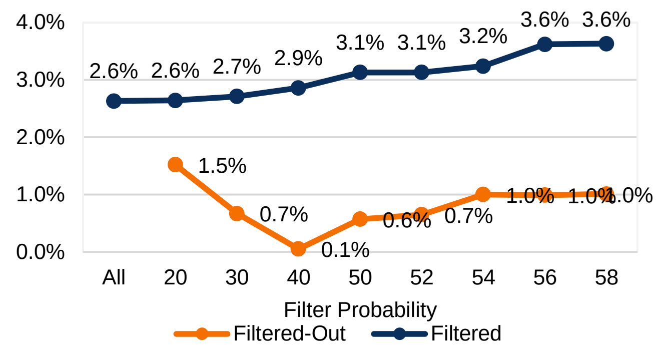

Source: State Street Global Exchange, MKT MediaStats (Mar 2013-Dec 2017).

Figure 4 shows that for DM currencies, the correlation between media sentiment of the unfiltered sample and the 1-day contemporaneous currency return is 1.7%, which is statistically significant with a p-value less than 0.01%, as shown in Figure 6. Similarly, Figure 5 shows the correlation between media sentiment of the unfiltered sample and the 1-day contemporaneous currency return is 2.6% for EM currencies, statistically significant with a p-value less than 0.01%, as shown in Figure 6 overleaf. These results provide validation for our media sentiment indicators.

Figure 6: Correlations between Article Sentiment and Contemporaneous FX Returns

		Filter Probability								
		All	20	30	40	50	52	54	56	58
Developed Market FX										
Filtered Out	% Total		5.60%	7.61%	10.40%	16.56%	17.41%	20.64%	30.09%	30.63%
	Corr		0.0138	0.0059	0.0036	0.0042	0.0023	0.0039	0.0072	0.0070
	p-value		0.21	0.53	0.65	0.51	0.71	0.49	0.13	0.13
Filtered	% Total	100%	94.40%	92.39%	89.60%	83.44%	82.59%	79.36%	69.91%	69.37%
	Corr	0.0173	0.0173	0.0181	0.0190	0.0202	0.0207	0.0208	0.0218	0.0220
	p-value	0	0	0	0	0	0	0	0	0
Emerging Market FX										
Filtered Out	% Total		3.84%	5.91%	10.76%	21.15%	21.94%	28.57%	40.03%	40.56%
	Corr		0.0152	0.0067	0.0005	0.0057	0.0065	0.01	0.0099	0.0101
	p-value		0.25	0.53	0.95	0.31	0.25	0.04	0.02	0.01
Filtered	% Total	100%	96.16%	94.09%	89.24%	78.85%	78.06%	71.43%	59.97%	59.44%
	Corr	0.0263	0.0264	0.0271	0.0286	0.0313	0.0313	0.0324	0.0362	0.0363
	p-value	0	0	0	0	0	0	0	0	0

Source: State Street Global Exchange, MKT MediaStats (Mar 2013-Dec 2017).

Figures 4 and 5 also demonstrate that correlations between the filtered article sentiment and currency returns increase as we increase the threshold probability. When the threshold probability level increases, fewer articles remain in the filtered sample. For example, for DM currencies, when we set the threshold probability at 54%, 20.64% of the articles are removed, as shown in Figure 6. The correlation between the filtered-out—or non-relevant—articles and currency return is 0.39%, with a p-value of 49% indicating that currency price movement does not respond significantly to the sentiment expressed in these articles. The correlation for the filtered sample, in contrast, increases to 2.08%, consistent with the prior that increasing the relevancy of articles improves correlation with currency returns. The results for EM currencies are quite similar. Indeed, as we remove more articles using higher cutoff probability levels, the correlations between the sentiment of remaining relevant articles and currency returns increase from 2.63% in the unfiltered sample to 3.63% for the threshold probability of 58%.

Predicting Future Returns

Following the validation of our model efficacy to improve article relevancy, we proceed to investigate the impact of filtering on return predictability of the media sentiment indicators. To that end, we construct portfolios based on the sentiment indicators generated from three samples of media articles: i) the non-filtered, ii) the filtered, and iii) the filtered-out samples.

The sentiment indicators are generated daily and incorporate media articles from a rolling window of the past 28 days. At the end of each trading day, we sort currencies into two groups using the latest sentiment indicators available. A zero-cost long/short portfolio is then constructed buying the high-sentiment currencies and shorting the low-sentiment ones. The portfolio is held for 10 trading days. We utilize a ladder approach: every day, we hold 10 long/short portfolios, and only one-tenth of the aggregate portfolio is rebalanced reflecting the most recent sentiment available. The daily return of the aggregate long/short portfolio under our ladder approach represents the equal-weighted average of the returns of the 10 portfolios. In Figure 7 and Figure 8, we report the annualized return spread of the zero-cost aggregate portfolio for various levels of filter probabilities, in DM and EM separately.

Figure 7: Return Spread (Ann %) between High- and Low-Sentiment DM Currencies

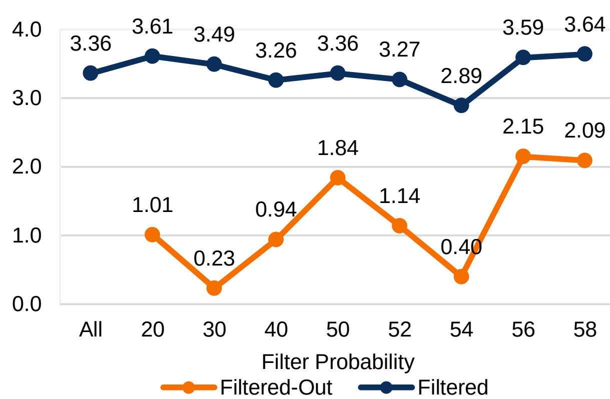

Source: State Street Global Exchange, MKT MediaStats (May 2011-Dec 2017).

Figure 8: Return Spread (Ann %) between High- and Low-Sentiment EM Currencies

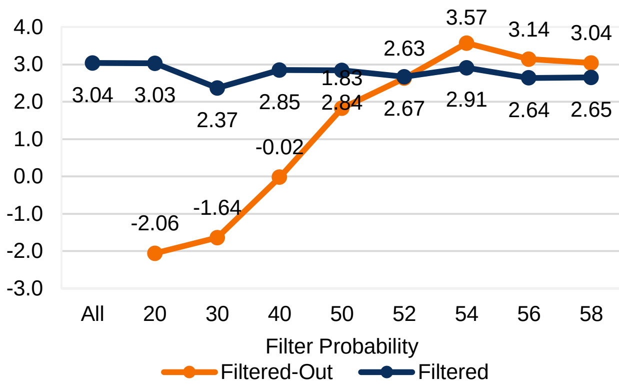

Source: State Street Global Exchange, MKT MediaStats (May 2011-Dec 2017).

The results in Figure 7 and Figure 8 indicate that media sentiment predicts positively future currency returns. For example, using sentiment indicators generated using the original non-filtered articles, an equal-weighted portfolio of high-sentiment currencies outperforms an equal-weighted portfolio of low-sentiment currencies by 3.36% annualized for DM currencies. The corresponding return spread for EM currencies stands at 3.04%.

Using more filtered media articles increases the predictability of media sentiment for DM FX. The return spread of the zero-cost long/short portfolio expands to 3.61% and 3.49% annually when we increase the filter probability to 20% and 30%, respectively. EM FX, in contrast, exhibit a different behavior as increasing the filter probability does not seem to impact predictability.

Figure 7 and Figure 8 also present the return spreads of an identical trading strategy using the filtered-out—or non-relevant—articles, as defined by our filter model. The associated return spreads are considerably lower—especially in DM FX—consistent with our thesis that these articles contain less pertinent information for currencies. In EM FX, however, we find evidence that the non-relevant articles do in fact predict positively future currency returns and generate better results than the relevant set of media articles at filter probabilities higher than 52%. These results fuel the debate spanning the continuum between practitioners (who predicate trading decisions on deep intuition) and statisticians (who find true potential in being able to detect data patterns that have not been discovered yet). Indeed, our results indicate that quantitative algorithms may be trading currencies on hidden, spurious data patterns in emerging markets.

Conclusions

In this White Paper, we set out to evaluate whether using supervised machine learning to filter out noise in media articles enhances FX trading strategies. We demonstrate that correlations between article sentiment scores and contemporaneous currency returns increase using machine learning techniques. However, filtering does not affect return predictability of media-based sentiment indicators consistently across currencies. While we found a statistically significant improvement in developed market currencies, there is limited value in emerging market currencies over time.

For more information, please contact:

North America:

Jenni Lee

jenni.lee@statestreet.com

Jamie Stewart

jstewart2@statestreet.com

Globally:

mktmediastats@statestreet.com

Disclaimers and Important Risk Information

Important Legal Information – Please Read

State Street Global Exchange℠ is a trademark of State Street Corporation (incorporated in Massachusetts) and is registered or have registrations pending in multiple jurisdictions.

Use of report

This document and information herein (together, the "Content") is for informational, illustrative and/or marketing purposes only and it does not constitute investment research or investment, legal, or tax advice. The Content provided is not, nor should be construed as, any offer or solicitation to buy or sell any product, service, or securities or any financial instrument, and it does not constitute any binding contractual arrangement or commitment for State Street Corporation and its subsidiaries and affiliates ("State Street") of any kind. The Content provided does not purport to be comprehensive nor intended to replace the exercise of a client's own careful independent review regarding any corresponding investment or other financial decision.

The material contained in this document is for general information purposes only. It relates only to hypothetical models of past performance of the proposed State Street MediaStats strategies, and not to any actual portfolios or funds in existence. Prospective application of the methodology used to construct actual portfolios may not result in performance commensurate with the backtest returns shown. Backtest results are based on criteria applied retroactively with the benefit of hindsight and knowledge of factors that may have positively affected performance, and cannot account for all financial risk that may affect the actual performance of the strategy. No allowance has been made for management fees which would reduce investment performance. Actual results may differ. This material is not intended as an offer or a solicitation for the purchase and/or sale of any security or financial instrument, nor is it advice or a recommendation to enter into any transaction. Returns represent simulated performance based on rules used in the creation of the hypothetical model, are not a guarantee of future performance and are not indicative of any specific investment.

Distribution

The Content provided is not intended for retail clients, nor is intended to be relied upon by any person or entity, and is not intended for distribution to or use by any person or entity in any jurisdiction where such distribution or use would be contrary to applicable law or regulation. No permission is granted to reprint, sell, copy, distribute, or modify the Content in any form or by any means without the prior written consent of State Street.

Other Important Disclosures

This document and information herein (together, the "Content") is subject to change without notice based on market and other conditions and in any event may not reflect the views of State Street Corporation and its subsidiaries and affiliates ("State Street"). The Content provided is for informational, illustrative and/or marketing purposes only; it does not constitute investment research or investment advice. The Content provided was prepared and obtained from sources believed to be reliable at the time of preparation, however it is provided "as-is" and State Street makes no guarantee, representation, or warranty of any kind including, without limitation, as to its accuracy, suitability, timeliness, merchantability, fitness for a particular purpose or otherwise. State Street disclaims all liability, whether arising in contract, tort or otherwise, for any claims, losses, liabilities, damages (including direct, indirect, special or consequential), expenses or costs arising from or connected with the Content. The Content provided does not constitute any binding contractual arrangement or commitment for State Street of any kind. The Content provided is not intended for retail clients, nor is intended to be relied upon by any person or entity, and is not intended for distribution to or use by any person or entity in any jurisdiction where such distribution or use would be contrary to applicable law or regulation. The Content provided may contain certain statements that could be deemed forward-looking statements; any such statements or forecasted information are not guarantees or reliable indicators for future performance and actual results or developments may differ materially from those depicted or projected. Past performance is no guarantee of future results. No permission is granted to reprint, sell, copy, distribute, or modify the Content in any form or by any means without the prior written consent of State Street.