Quantifying Trader Response Dynamics in the Foreign-Exchange Market

Myrl G. Marmarelis*

Department of Computer Science, University of Southern California

1 Abstract

In searching for a harbinger of market movements, the author presents a technique to monitor the temporal patterns that lead to anticipated price changes. Capturing the dominant frequencies in the resulting convolution kernel functions creates a metric to reveal insights about a market as it evolves over time. Serendipitously, a few key discoveries point towards apparent predictive utility when applied to foreign exchanges: namely, multi-scale periodicity and precursor indications of actual market movements. This technique, which was inspired from the field of biomedical signal processing, opens an avenue for novel analysis of the trader/trend-response dynamics in a market.

2 Keywords

trader sentiment; non-stationary analysis; spectral dynamics; predictive modeling

3 Introduction

"Ego cogito, ergo sum," famously wrote Descartes in his *Meditations*, asserting that if there is anything of which he can be certain, it is the existence of his own mind. Hence the foundation of solipsism—how can we know that a mind-independent external world even exists? Our senses have betrayed us countless times before. Whereas the truth on this particular matter of existence is unknowable, certain other things are knowable (and indeed known) to be mere human constructs, like the global economy. The financial world is an interconnected structure composed of just one building block, the human. Society's collective whims that shape the direction of the economy are essentially a grand-scale manifestation

*445 Palos Verdes Drive West, Palos Verdes Estates, CA 90274 USA. +1(310)999-9210, mmarmare@usc.edu

of human psychology; thus, the guaranteed existence of human bias necessitates that market efficiency is an untenable idealism to which reality strives to converge.

Starting in the early 1990s there was an explosion in the field of behavioral finance, wherein empiricists worked with theoreticians to catalogue the different ways to take advantage of market inefficiencies for the sake of short-term profit. In the big picture, it is generally well-known that overall investor optimism tends to lead to lower market returns [1]. This makes intuitive sense. In 1999, Hong et al. [2] (with an updated theory in [3]) proposed a theoretical framework by which informational asymmetry would allow different groups of trading agents to profit from momentum trading. Since then, Verardo [4] and others have offered substantial empirical support for this link between price drift and belief heterogeneity, the latter of which is partly due to suboptimal access to information. Authors of the likes of Miravete [5] have also laid out rigorous methods for interpreting the effects of information asymmetry. Verardo [4] estimates trader beliefs through the dispersion of analyst forecasts of earnings. More recently, the ease of obtaining comprehensive personal records from social media websites has allowed researchers like Liu [6] to more cleanly infer actual trader sentimet by computing an individualized emotional index for someone's feelings as expressed online. Liu found a correlation between this index and each trader's daily performance. Twitter is also a widely used social-media platform wherein some individuals exert high levels of influence—there is now significant evidence towards the role of these influencers in shaping speculative performance [7]. The work in this paper complements the aforementioned studies in behavioral economics by opening an avenue for developing a new kind of active-trader-sentiment index: one that is derived solely through features intrinsic to the financial time series of the market. Since there is currently no accepted model that relates short-term sentiment to market dynamics, it is not yet possible to justify the following conclusions with formal rigor.

The fractal-market hypothesis suggests that the stability of a market hinges on the condition that the strategies of involved traders operate on diverse time horizons, in order to ensure sufficient liquidity [8]. Starting with the work of Mandelbrot, many studies point towards the existence of time-scale invariant patterns in financial time series. In other words, they seem to exhibit some degree of self-similarity. Would it follow from the assumption of self-similarity, then, that observing high-frequency oscillations in the market could provide insight about future trends? After all, practitioners of science in a wide variety of fields have found great success in revealing the underlying "collective periodic oscillations of interacting elements" within complex systems by using resonances to uncover structure [9]. To infer the behavioral pattern of each agent would be to predict the future of the economy.

The circadian (i.e. daily) periodicity of volatility in markets such as foreign currency exchanges is extremely well-documented in financial literature [10, 11, 12, 13], a phenomenon that by itself can inform viable trading strategies. Foreign exchange volatility, however, now exhibits some interesting spectral dynamics as well [14] that most likely have to do with the emotional predispositions and modus operandi of the traders.

Many researchers in physiology, when faced with the challenge of modeling the vascular system [15] or certain neural systems [16], tend to place their trust in convolutional models with "kernels" represented by Laguerre expansions. These models are philosophically appropriate because the basis functions are inspired by the input-output relationship between natural processes as they respond to phenomena and return to equilibrium. A great advantage over more black-box methods such as large artificial neural networks is that the resulting coefficients for the set of Laguerre functions are easily interpretable and, in fact, have been utilized to diagnose diseases; for instance, the effect of Alzheimer's on feedback loops in the cerebral hemodynamics has been quantified to the extent that it can be used to facilitate diagnosis [15].

This paper outlines a model in the scope of capturing the short-term "trend-response" characteristics of the active traders through Laguerre expansions of their regression kernels. Through the derivation of scalar metrics that correspond to certain key features of interest, the small-scale effects that are encapsulated by each kernel are employed to map out a bigger picture of the market trends. The Discussion section lists novel discoveries encountered along the way.

4 Methods

Global economic markets can be viewed as an accumulation of time-delayed reactions to past events, which inform the financial community's sentiment about future performance. Therefore there must be certain dynamics that govern the time frame over which the market reacts to changes. As trader sentiment evolves along with the economy, these market dynamics are non-stationary as well. If there is a metric that could be used to instantaneously track the prevailing sentiment that is most relevant to a certain market, it could potentially be useful as a forecasting tool.

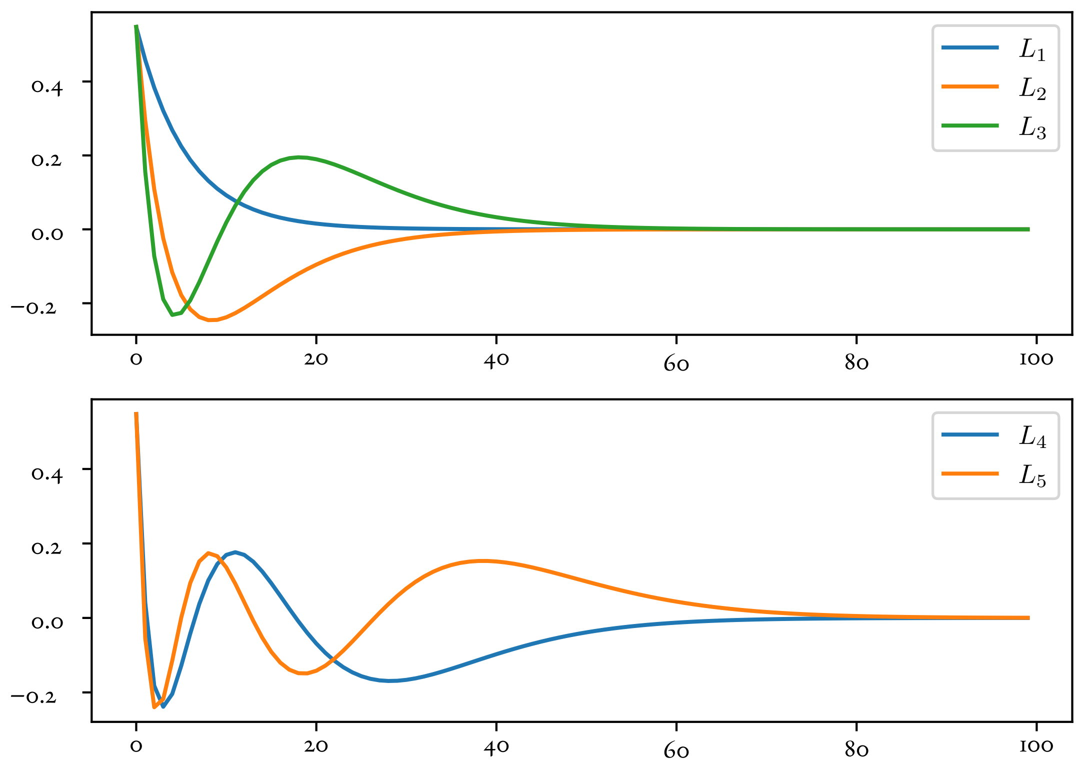

Figure 1: The first five basis functions. The system memory in this illustration is 100 units.

4.1 System Identification With Kernels

The relationship between multiple signals can be viewed as a functional operator (or a "system" that can be described by a predictive input-output model) where some of the signals are treated as inputs and the others as outputs. In the interconnected global economy, most market indices serve as both inputs and outputs to each other; however, the future is always a function of the past. This perspective leads to autoregressive modeling, where the past is taken as an input in order to predict the future. To model the time-delayed accumulation of effects, one could choose to convolve a linear-response function over the recent past of a market time series in order to generate a prediction of the near future. In the linear context, this can be achieved with a one-dimensional function, but a nonlinear context (i.e. system operator) requires multi-dimensional functions. Many studies have examined this problem, and efficient methods currently exist to estimate these functions from data [17]. Generally, the linear context is more amenable in practice.

One of the most widely used classes of simple predictive models is the representation of the future as a weighted moving average of the past. Traditionally denoted as AR(p), the model estimates the signal x_t of a process by taking into account the p previous discrete-time samples:

$$x_{t+1} = \sum_{\tau=0}^{p-1} c_\tau x_{t-\tau} + e_{t+1}, \tag{1}$$

where the coefficients $c_0, \ldots, c_{p-1}$ are the estimated parameters for the model and the residual term e_t is assumed to be white noise. The coefficients could be plotted in sequence to produce a view of the linear-response function (the "kernel"), which provides insight about the temporal dynamics of the process. Dividing p by the sampling rate of x_t effectively yields the assumed process memory; hence, it is clear that depending on the application, circumstances could lead to an undesirably large amount of free parameters. Such an AR(p) model with a high p would be unwieldy because overfitting adds noise to the resulting coefficients, rendering the kernel difficult to interpret.

To decouple the number of free parameters from the internal properties of the underlying process, one could estimate the kernel as a weighted sum of a small set of k basis functions. The accuracy of this approximation then lies on the appropriateness of the basis functions, in terms of their resemblance to the most typical scenarios in which events generate time-delayed ripple effects. For this, there has been a great amount of success resulting from the use of Laguerre polynomials multiplied by a diminishing exponential, since they capture the oscillatory mechanics that are often observed in complex systems, and they also place greater value on the more recent data. The applicability of Laguerre functions implies that only a small amount is needed to reach adequate modeling power—this is beneficial because the drastic reduction of free parameters makes the technique resilient to noise. This study will use the five basis functions shown in Figure 1. Empirically, it has been found that this number of Laguerre functions enables the model to capture up to three resonances [17].

It is clear that the resulting kernel should highlight any diminishing oscillatory effects that result from changes in the market. A basic autoregressive model that directly extends Equation 1 would have the kernel predict the immediate future as a function of the past in a discrete time series:

$$x_{t+1} = \sum_{\tau=0}^{M} x_{t-\tau} \cdot \left(\sum_{i=1}^{k} c_i L_i(\tau) \right) + e_{t+1}. \tag{2}$$

Here the basis-function coefficients are $\boldsymbol{c} \in \mathbb{R}^k$, the memory of the kernel is M, the evenly-sampled time series of interest is $\boldsymbol{x} \in \mathbb{R}^N$, the diminishing Laguerre functions are $L_i \colon \{0, \ldots, M\} \to \mathbb{R}$ for $i \in (1, \ldots, k)$, and the residual term is $\boldsymbol{e} \in \mathbb{R}^N$. The Laguerre functions have one free parameter that determines their time scale, conventionally termed α, the value of which is empirically maximized while still satisfying the following criterion for a selected system memory M: $\forall i \left[L_i(M) \approx 0 \bigwedge \dot{L}_i(M) \approx 0 \right]$. In other words, the kernel must diminish to zero towards the end of the chosen epoch. The model described so far is inadequate because its purely autoregressive kernel $\left(\sum_{i=0}^{k} c_i L_i(\tau) \right)$ would approach the identity impulse function (i.e. the Kronecker delta) that gives 1 at $\tau = 0$ and zeros everywhere else. Due to causality, the present correlates more to the immediate past than to the distant past.

Perhaps the kernels could project into the relatively distant future, with hopes that they will be able to capture more valuable information on the temporal dynamics of the market. One way to achieve this would be to attempt a prediction of the slope of a simple linear regression from the present to some point in the future. The regression slope can be viewed as a separate variable, so the model is no longer strictly autoregressive:

$$\beta_{t\to t'} = \sum_{\tau=0}^{M} x_{t-\tau} \cdot \left(\sum_{i=1}^{n} c_i L_i(\tau) \right) + e_t, \tag{3}$$

given that linear regression produces coefficients α and β that best fit the equation $x_\tau = \alpha_{t\to t'} + \beta_{t\to t'}(\tau - t)$ in some time interval $\tau \in [t, t']$, with $t' - t = \Delta t$ termed the outlook of the model and held constant. To ensure kernel effectiveness, the signal being convolved upon needs to have near-zero mean. For that reason, x_t is the difference between the market time series and a large-window non-centered and backwards-looking moving average from each past data point. The least-squares-fit kernel, which minimizes $\left(\sum_{t=1}^{N} e_{M+t}^2\right)$, can be found by efficiently solving a system of linear equations guaranteed to have a single unique solution.

Note that these kernels inform us on the dynamics of the *system*, which are different from those of the isolated signals. The two signals under analysis are future slope and past performance, which both come from the same time-series data but do not necessarily have the same characteristics. The quantified relationship between the two may also have distinct characteristics from those resulting due to pure autoregression on either signal.

4.2 Historical Data

This study utilized minute-by-minute weekday quotes of bid prices for different currency pairs, beginning in December 2011 and continuing until February 2017 (totaling 63 months). The data was obtained from GAIN Capital's public archives in weekly partitions and the pairs under scrutiny are EUR/USD, GBP/USD, and CHF/USD. Certain pairs had to be derived from existing ones; for instance, CHF/USD was calculated through $\frac{\text{EUR/USD}}{\text{EUR/CHF}}$. Each time series was then slightly smoothed (to avoid aliasing issues with noise—a moving average with a 10-minute window was applied) and subsampled at 10-minute intervals, to decrease computational costs.

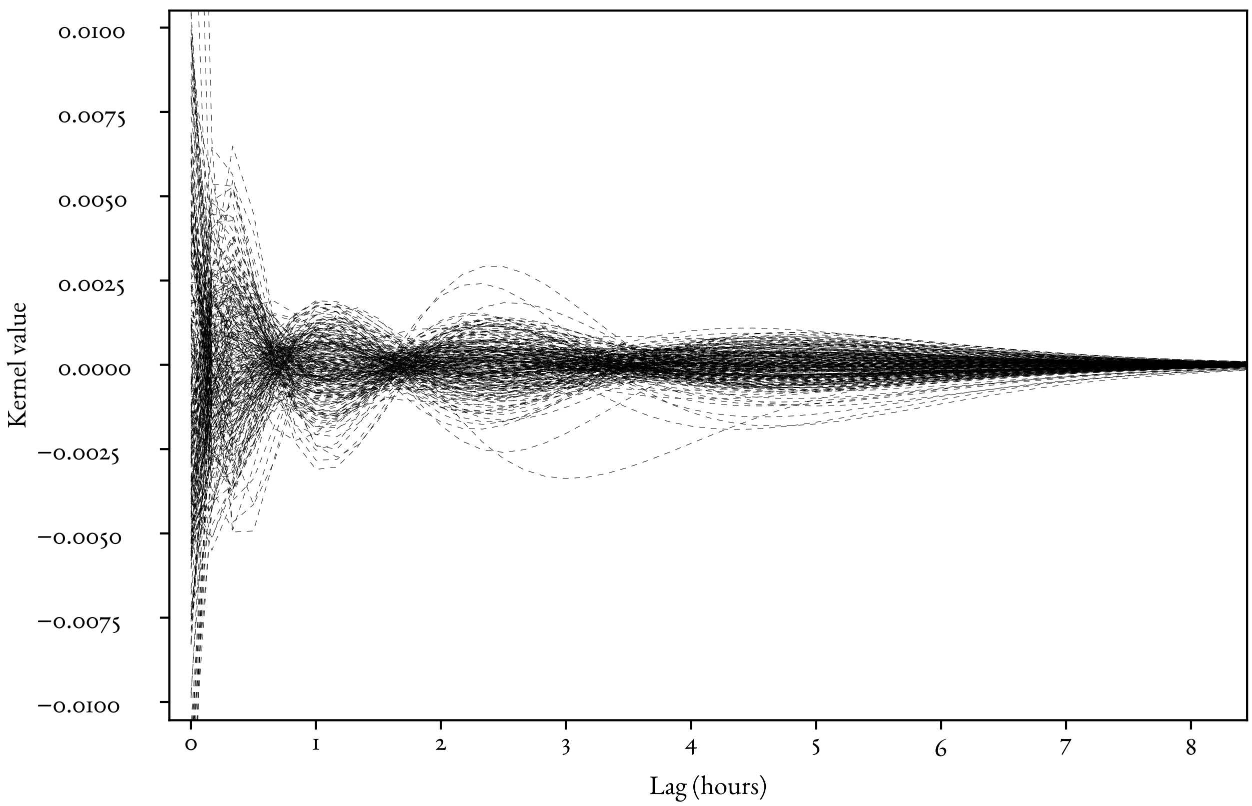

Figure 2: All of the 12-hour trend-response kernels generated from EUR/USD (zoomed in to the first eight hours).

5 Results

For each week in the data set, kernels were generated using the above method with five basis functions, a system memory of 12 hours, a future outlook of 12 hours, and a moving average of 24 hours. Figure 2 shows the resulting kernels from each week of EUR/USD: it is more common to see a kernel have a negative first value because a quick rise is likely to be met with a corresponding drop in the index. Also note that there seem to be some significantly common waveforms, with switching phases. A highly bullish time segment in the market should produce a mostly positive kernel, but the correspondence between large trends and trend-response kernels is not so cut and dry due to the distinction between system characteristics and the characteristics of individual signals.

5.1 Spectral Analysis

To plot the entire kernels as they evolve over time would not facilitate effective interpretation. Instead, there is an intermediate step of finding some scalar metric to characterize a whole kernel in a useful way. Such a metric could be the first value of each kernel—this has potential in estimating market sentiment, but does not take into account the overall kernel shape. To combat this, one could take the sum of the first few values of each kernel, but that too loses information. After multiple attempts, it was decided in this study to extract the dominant frequency of each week's kernel. This is done by identifying peaks in amplitude from the Fourier transform and using the location (i.e. corresponding frequency) of the greatest one.

In addition to the surplus of information, the erratic non-stationarity of the raw week-by-week dominant frequencies makes it difficult to interpret any meaningful features. To focus on long-term trends without over-smoothing, the 8-week moving average was plotted in blue alongside the time series under analysis in orange, to produce Figure 3.

	12-hour kernels			3-hour kernels		
NMSE	EUR/USD	GBP/USD	CHF/USD	EUR/USD	GBP/USD	CHF/USD
Mean	1.018	1.031	1.026	0.737	0.783	0.800
Median	0.991	0.998	1.005	0.386	0.395	0.392
Std. dev.	0.199	0.189	0.207	1.755	2.058	1.946
IQR	0.206	0.210	0.225	0.575	0.576	0.581
Skew	0.557	0.441	0.708	18.733	17.482	13.948

Table 1: Notable NMSE statistics—mean, median, standard deviation, interquartile range, and skewness.

5.2 Prediction Error

The normalized mean-square error (NMSE), here defined as the sum of square errors divided by the variance of the data (i.e. of the slopes), is tallied for each currency pair in Table 1, for which the columns labeled "12-hour kernels" correspond to the kernels we have discussed so far, and those labeled "3-hour kernels" refer to the much shorter and higher-frequency kernels that will be analyzed below. With respect to the 12-hour kernels, their average normalized errors are slightly greater than unit; hence the present model is a

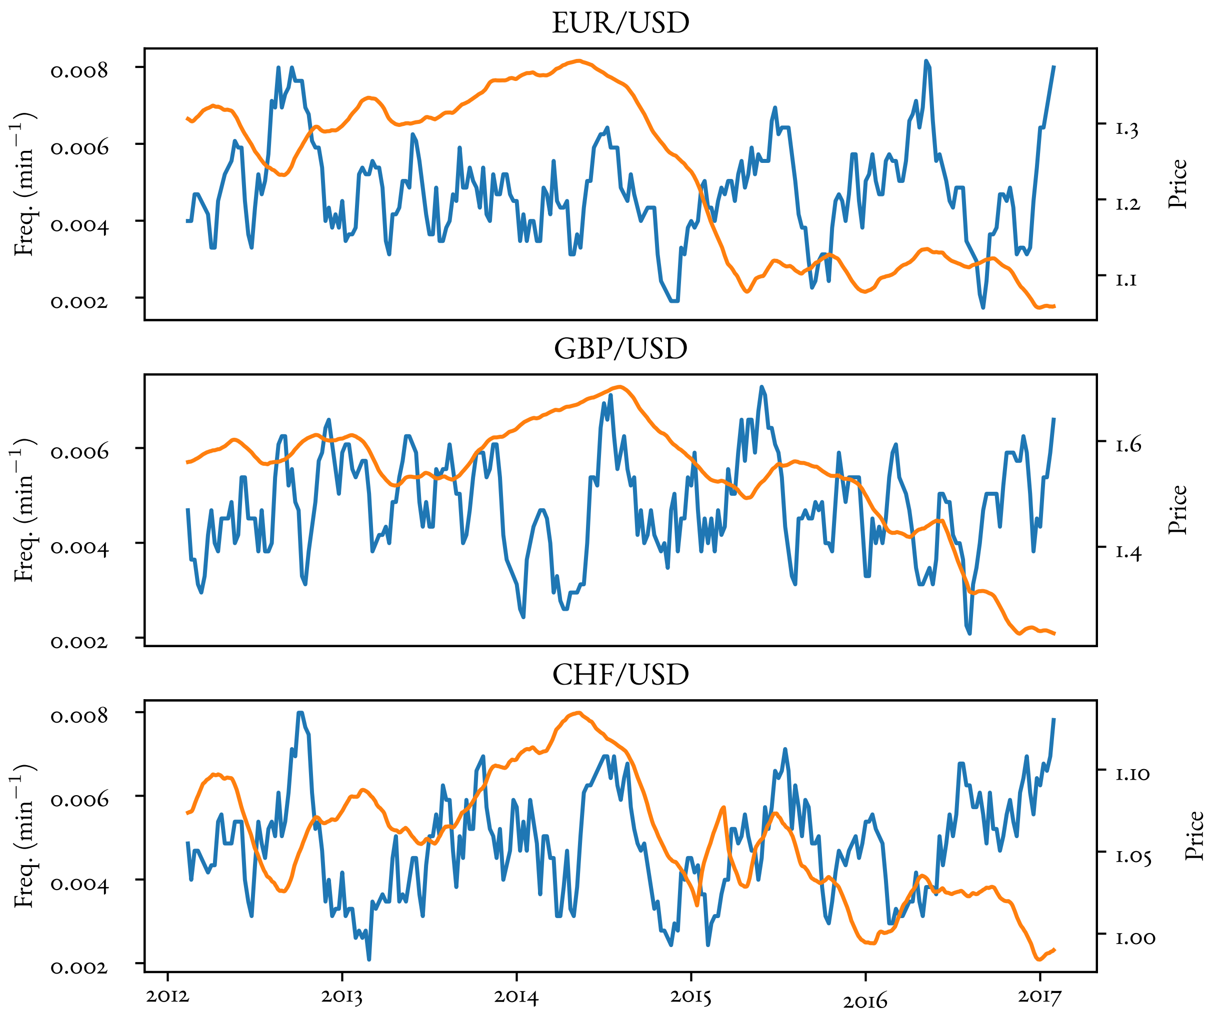

Figure 3: Smoothed dominant kernel frequencies (blue) along with trading values of the currency pairs (orange).

worse predictor than a horizontal line. This supports the hypothesis that the market is effectively Brownian motion and extremely difficult to predict at the intraday scale; regardless, this high error does not hamper the emergence of interesting large-scale patterns. Regardless, the purpose of this study is not to directly predict through the outputs of the kernels but instead to analyze the longer-term non-stationary dynamics of their shapes as a possible trend predictor.

Remarkably, decreasing both the memory and the outlook horizon to 3 hours each yields kernels with significantly lower average NMSEs, but with extremely large outliers (as is witnessed by the enormous skewness of the NMSE in Table 1). This result can be explained by the occasional drastic reaction to an exogenous event, in a market that is otherwise governed by gradual trends.

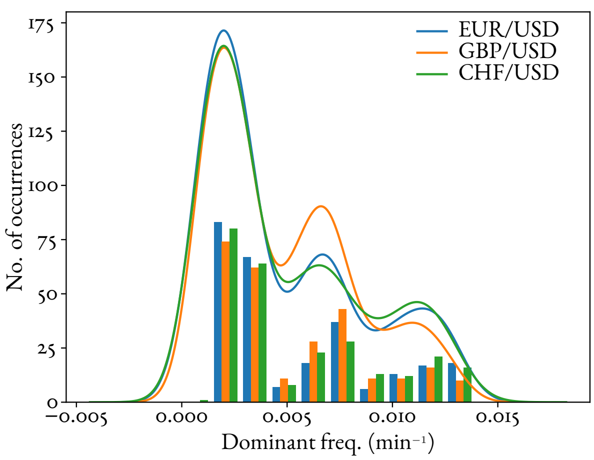

Figure 4: Histogram along with kernel-density estimates (KDEs) of the distributions of dominant frequencies of trend-response kernels for each week of all three currency pairs.

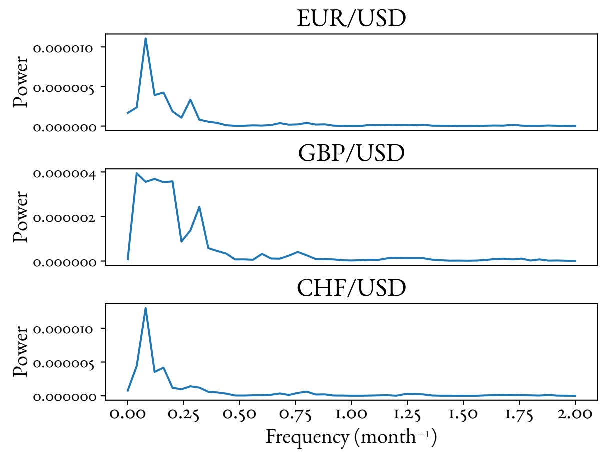

Figure 5: Spectral density estimation of the smoothed 12-hour trend-response dominant frequency signal, obtained through the use of Welch's method.

6 Discussion

The method developed in this study is hypothesized to locally quantify the emotional climate surrounding a certain market. Though a metric on short-term performance, its slow evolutionary trends over long periods of time appear especially insightful due to observed regularities. If it were feasible to determine the instantaneous trend-response kernels in real time it would allow someone to predict the performance of the underlying security and make enormous profits, and the fact that there are slow-moving trends is encouraging. Even gaining some insight into the general shape of the kernel at any given point in time has the potential to provide valuable information about the sentiment of the global community with respect to the market. Highlighted below are a few key observations, mostly with regards to the dominant trend-response frequencies as displayed in Figure 3.

6.1 Periodicity

The dominant frequency signal from sequential EUR/USD trend-response kernels is especially periodic. From 2015 to 2017, there are notable year-long cycles that seem to dip in the fall and peak in the spring or early summer. Such seasonal effects are not so apparent in the raw currency time series. As will be discussed in Section 6.3, it seems like these periodic cycles are tied to changes in the market itself. Figure 5 suggests that this almost-yearly periodicity is not unique to EUR/USD, but also manifests in the GBP/USD and the CHF/USD time series as indicated by their spectral peaks at around 0.1 month^{-1}.

6.2 Distribution of Dominant Frequencies

By taking all the dominant frequencies of the trend-response kernels through the 63-month period, histograms were compiled and plotted in Figure 4, which seems to indicate the presence of a trimodal distribution of frequencies. This inkling was tested by computing kernel-density estimates (KDE) according to Silverman's rule [18], and checking if the modes are still visible as maxima; as shown in Figure 4, the KDEs are clearly trimodal. One may conjecture that this property common to EUR/USD, GBP/USD, and CHF/USD must be a trait of the United States dollar. Further investigation makes it difficult to believe that this is the case because the EUR/CHF, for instance, also exhibits the observed multimodality.

Figure 4 can be thought of as a pseudo-spectrum, as it shows the number of occurrences (i.e. relative "intensities") of weekly trend-response dominant frequencies. The three peaks in the figure correspond to around 0.002 per min, 0.0075 per min, and 0.012 per min in that order, corresponding to 2.88 per day (8.33-hour periods), 10.80 per day (2.22-hour periods), and 17.28 per day (1.39-hour periods). By far the largest peak is the longer-term dominant frequency, making it the most valuable to predict because it is both long-term and occurring frequently. Notice, also, how this pseudo-spectrum resembles a harmonic spectrum with fundamental frequency 0.003 per min and multiples 0.006 per min and 0.012 per min. In fact, the amplitudes of the peaks are approximately whole-number multiples as well, with the first being double the second, which in turn is double the third. This result confirms the empirical wisdom that five Laguerre functions are enough to capture up to three resonances.

The finding above is significant because it hints at the presence of three distinct trading behaviors, each categorizing potentially different market dynamics. Could these specific peak frequencies be inherent to the Laguerre functions used? Potentially, but this is not a cause for concern because there are still clear and insightful patterns that emerge from the temporal alternation between the three trading behaviors.

6.3 Prediction Power

Looking back at Figure 3, it appears that a dip in the trend-response dominant frequency precedes a significant fall in the actual currency price. In EUR/USD this is visible right before the start of the years 2015, 2016, and 2017, where dips in the blue line (that is, when it approaches 0.002 per min) precede drops in the currency by about two to three months. These observed regularities are not unique to EUR/USD: in the middle of 2016, the GBP/USD currency pair also follows that pattern, as well as in the beginning of 2014. CHF/USD is a little less clear in some time periods, but in the beginning of 2015 there are two dips that adhere to our hypothesis and one more in early 2013. Both the blue and the orange lines are smoothed by the same amount, and with windows that are smaller than the hypothesized predictive horizon (2 months); therefore, it is not just retrospective.

The aforementioned findings appear to suggest that a decrease in trader responsiveness, as marked by a lowered dominant response frequency, tends to precede a drop in the market. There is no comparable observation for the upside effect of an increase in trader responsiveness: this directional asymmetry with respect to the market's trend-response dynamics is indicative of human bias, namely risk aversion.

6.4 Trying Higher Frequencies

Heretofore, we have been looking at kernels with a 12-hour memory and a 12-hour outlook, generated weekly. The multi-scale nature of this analysis begs for the inclusion of different time horizons, and of shorter ones in particular. Hence kernels were generated with 3-hour memories and 3-hour outlooks, employing 6-hour moving averages. The smaller scale of these kernels allowed the generation of a new one every 3 hours, meaning that each day had about 8 kernels in sequence. It is clear from Figure 6 that the dominant frequencies of the shorter kernels exhibit a similar trimodality as in Figure 4 with the longer

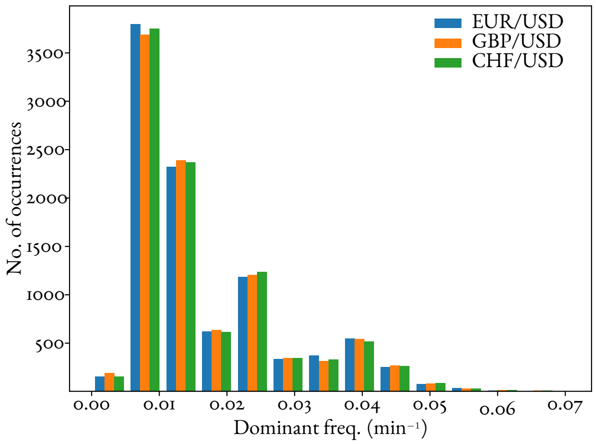

Figure 6: Histogram of the distributions of dominant frequencies of 3-hour trend-response kernels for all three currency pairs.

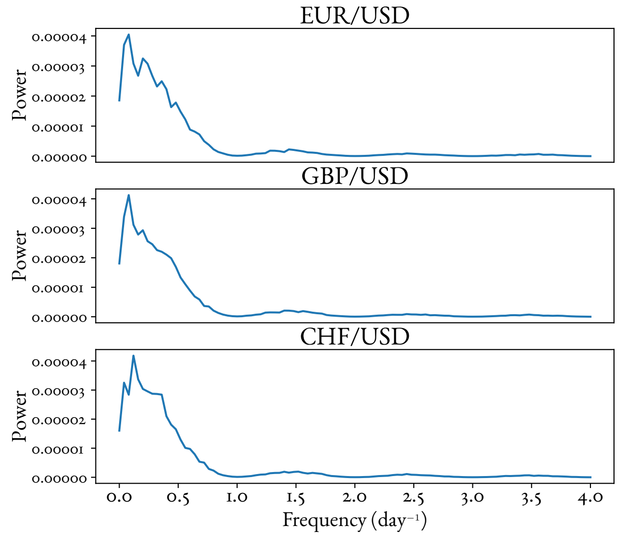

Figure 7: Spectral density estimation of the smoothed 3-hour trend-response dominant frequency signal, obtained through the use of Welch's method.

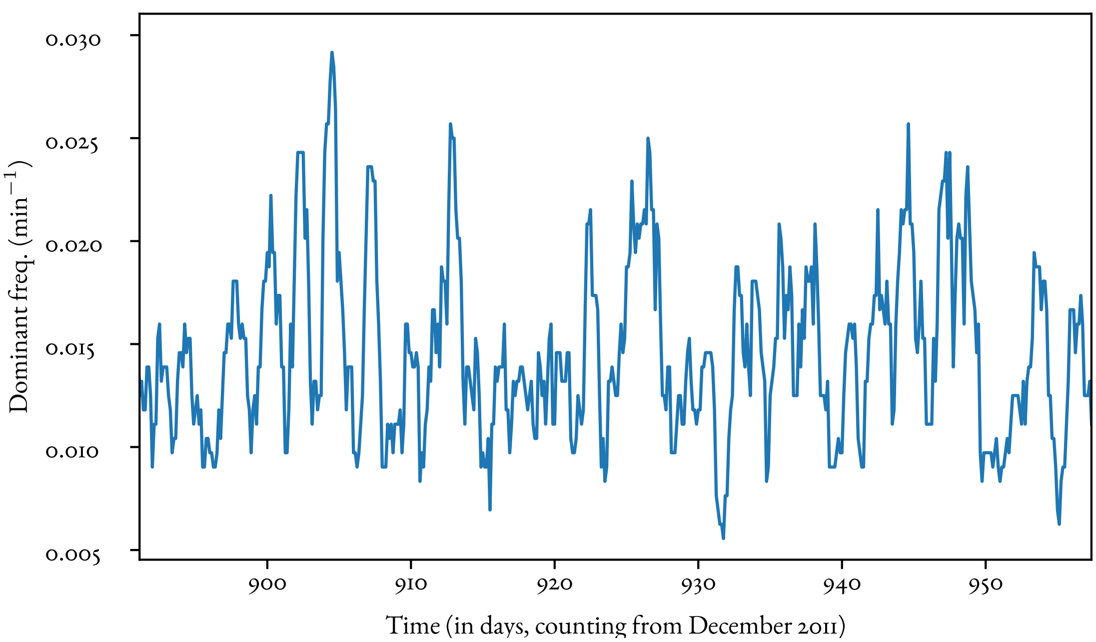

Figure 8: A snippet from the smoothed EUR/USD 3-hour trend-response dominant frequency signal, illustrating its multi-day periodicity.

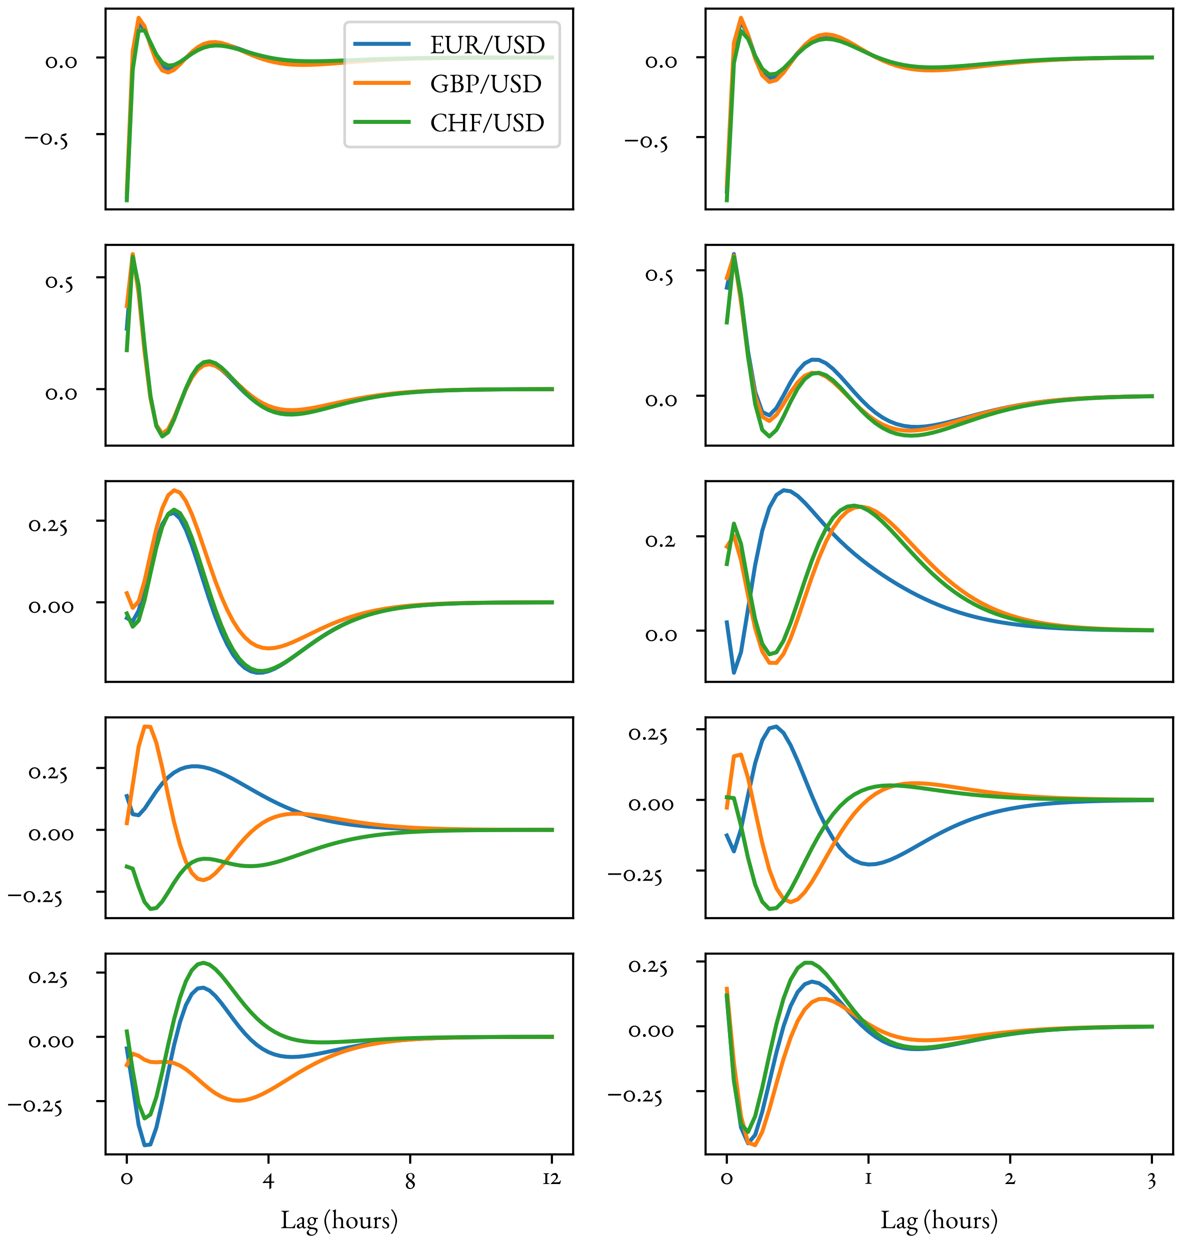

Figure 9: Reconstructed kernel basis functions through principal-component analysis (PCA). Left column is 12-hour kernels and right column is 3-hour kernels, each ranked from top to bottom in order of significance.

kernels. The new dominant-frequency signal is also periodic on the scale of a few days. With only a day-long moving average (spanning 8 kernels), one obtains a signal with characteristics as seen in Figure 8. The pronounced cycles have period length of about 2 to 3 days. Next, Figure 7 estimates the power of the aforementioned signal in Figure 8 over the frequency domain. It is likewise clear that a large part of the signal lies in the frequency range of slightly less than 0.5 cycles per day: this periodicity range is distinct from the multi-month horizon seen in Figure 3, and also from the daily cycles in volatility that are widely recognized within the financial community. It would be worth exploring how these newly observed cycles translate into actionable information about market dynamics, e.g. what it means when the day happens to be in a downward half-cycle versus in an upward one. For now, it is hypothesized that it informs us about the traders' reaction speed to sudden jumps or falls in the market.

6.5 Principal-Component Analysis

Principal-component analysis (PCA) involves the identification of "components" (i.e. vectors) that explain the highest amount of variation in a set of observations. The technique outputs a set of orthogonal bases ordered by the variance of the data when projected onto each one. This new representation of the basis functions allows us to scrutinize upon how the Laguerre functions are employed by our model to approximate kernels. PCA proceeds as follows: first, we compute the $(M \times M)$ covariance matrix of all our kernels, where M is the system memory, the length of each kernel vector. The M eigenvectors of this matrix are the new orthogonal bases, with each eigenvalue corresponding to that eigenvector's relative weight. Remarkably yet not surprisingly, there were exactly five "principal components" (i.e. eigenvectors) with any significant contribution to the variance (as indicated by the eigenvalue). It is not surprising because every kernel is a weighted sum of five orthogonal Laguerre functions. Figure 9 shows these five ordered principal components derived from the 12-hour kernels and the 3-hour kernels, with the top row of charts explaining the most variance and the bottom the least. An insightful result is the fact that the principal components more-or-less are the same for EUR/USD, GBP/USD, and CHF/USD, as shown by the blue, orange, and green curves respectively. Note also that the second, third, and fourth principal components in Figure 9 have very similar shapes.

6.6 Influence of Other Time Series

The model can easily be extended to include a contribution of exogenous variables, paired with their own kernels, to the predicted performance of a certain market. For instance, the trend-response kernels for EUR/USD can be supplemented with kernels that convolve on the recent past of GBP/USD and CHF/USD, with the combination of all three giving a prediction for the future of EUR/USD. While that particular example is not advisable because the three currency pairs are significantly collinear, resulting in the same effects being distributed across three kernels and thus increasing the proportion of noise, a new model could take into account the effects of indices on distinct sectors like oil prices, bond rates, and the S&P 500 on a given currency pair like EUR/USD. The new kernels should theoretically isolate the external market effects from the dynamics inherent to the currency, and give cleaner results.

7 Conclusion

In the present set of time series spanning the years 2012 to 2017, the dominant frequencies of the sequence of generated 12-hour kernels produce a possible forecasting signal. Whenever it dips to about $0.002 \text{ min}^{-1} \approx 2.88 \text{ day}^{-1}$, there is a remarkable likelihood that the underlying trading price will drop within the next few (2 to 3) months. Furthermore, the noteworthy feature of the 3-hour kernels appears to be the consistent multi-day periodicity in their dominant frequency. Even though this system is highly non-stationary, calculation of the kernels' principal components indicates that certain dynamics are stable across the three currencies; this inference is further supported by the pronounced trimodality in the dominant-frequency distributions of both the 12-hour and the 3-hour kernels.

The analysis described above is agnostic to any underlying market mechanics that could explain its findings—this condition is analogous to the current state of the art in machine learning, wherein model interpretability is not yet flushed out enough to match the empirical results. Nonetheless, due to a number of remarkable observations, there exists a significant basis for believing in the utility of dynamic trend-response analysis of market data.

8 Further Study

To rigorously justify the effectiveness of the outlined method, it would be helpful to examine more diverse applications of it, as well as to feed it longer records of data. As highlighted above, there are many clear paths through which this model has potential for improvement and fruitful exploration.

An interesting experiment would be to separately generate kernels for a vast array of exogenous variables that correspond to different market sectors, and rank each of these variables based on the consistency of their kernels through time (a simple approach would be to normalize each kernel and then compute the root mean-square error of the entire set). The exogenous variable with the most consistent kernels would offer the greatest predictive utility.

References

[1] M. Schmeling, "Investor sentiment and stock returns: Some international evidence," *Journal of Empirical Finance*, vol. 16, pp. 394–408, 2009.

[2] H. Hong and J. C. Stein, "A unified theory of underreaction, momentum trading, and overreaction in asset markets," *The Journal of Finance*, vol. 54, December 1999.

[3] H. Hong and J. C. Stein, "Disagreement and the stock market," *The Journal of Economic Perspectives*, vol. 21, no. 2, pp. 109–128, 2007.

[4] M. Verardo, "Heterogeneous beliefs and momentum profits," *The Journal of Financial and Quantitative Analysis*, vol. 44, pp. 795–822, August 2009.

[5] E. J. Miravete, "Convolution and composition of totally positive random variables in economics," *Journal of Mathematical Economics*, vol. 47, pp. 479–490, 2011.

[6] B. Liu, R. Govindan, and B. Uzzi, "Do emotions expressed online correlate with actual changes in decision-making?: The case of stock day traders," *PLoS ONE*, vol. 11, January 2016.

[7] F. Corea, "Emotional speculative behavior in the option market," *International Journal of Advanced Computer Research*, vol. 6, no. 22, 2016.

[8] N. Nava, T. D. Matteo, and T. Aste, "Time-dependent scaling patterns in high frequency financial data," *The European Physical Journal Special Topics*, vol. 225, 2016.

[9] A.-H. Sato and J. Hołyst, "Characteristic periodicities of collective behavior at the foreign exchange market," *The European Physical Journal B*, vol. 62, 2008.

[10] C. Brooks and M. J. Hinich, "Detecting intraday periodicities with application to high frequency exchange rates," *Journal of the Royal Statistical Society. Series C (Applied Statistics)*, vol. 55, no. 2, pp. 241–259, 2006.

[11] D. Erdemlioglu, S. Laurent, and C. J. Neely, "Which continuous-time model is most appropriate for exchange rates?," *Journal of Banking & Finance*, vol. 61, 2015.

[12] K. P. Evans and A. E. Speight, "Intraday periodicity, calendar and announcement effects in euro exchange rate volatility," *Research in International Business and Finance*, vol. 24, pp. 82–101, 2010.

[13] N. Nonejad, "Forecasting aggregate stock market volatility using financial and macroeconomic predictors: Which models forecast best, when and why," *Journal of Empirical Finance*, vol. 42, pp. 131–154, 2017.

[14] A. Grossmann, I. Love, and A. G. Orlov, "The dynamics of exchange rate volatility: A panel var approach," *Journal of International Financial Markets, Institutions & Money*, vol. 33, pp. 1–27, 2014.

[15] V. Z. Marmarelis, D. C. Shin, M. Orme, and R. Zhang, "Time-varying modeling of cerebral hemodynamics," *IEEE Transactions on Biomedical Engineering*, vol. 61, March 2014.

[16] T. P. Zanos, S. H. Courellis, T. W. Berger, R. E. Hampson, S. A. Deadwyler, and V. Z. Marmarelis, "Nonlinear modeling of causal interrelationships in neuronal ensembles," *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, vol. 16, August 2008.

[17] V. Z. Marmarelis, *Nonlinear dynamic modeling of physiological systems*. IEEE Press series on biomedical engineering, Wiley-Interscience, 2004.

[18] B. W. Silverman, "Using kernel density estimates to investigate multimodality," *Journal of the Royal Statistical Society. Series B (Methodological)*, vol. 43, no. 1, pp. 97–99, 1981.