Inattentive Search for Currency Fundamentals*

Agnieszka P. Markiewicz†, Ralph C. Verhoeks‡, Willem F.C. Verschoor§
and Remco C.J. Zwinkels¶

February 2020

Abstract

We study weather agents' limited attention can account for the time varying link between exchange rates and economic fundamentals. We proxy attention for different sources of economic information by the search intensity of related queries on Google. We demonstrate that the higher the attention to a certain economic fundamental, the better its ability to forecast exchange rate movements. The best forecasts and the highest returns are systematically delivered by models that select the most salient fundamental and that assume that it changes frequently. These results suggest that attention to individual exchange rate fundamentals is limited and short-lived.

Keywords: *exchange rate forecasting, limited attention, investment strategies*

JEL Classification: F31, F41, E44

*We thank participants of the Australasian Finance and Banking Conference 2017; the Society for Non-linear Dynamics and Econometrics conference 2018; the European Economic Association 2018; and the Research in Behavioral Finance Conference 2018 for their helpful comments and suggestions. A previous version of this paper circulated under the title "The Winner Takes It All: Predicting Exchange Rates with Google Trends".

†Erasmus School of Economics and Tinbergen Institute; markiewicz@ese.eur.nl.

‡De Nederlandsche Bank and Vrije Universiteit Amsterdam; r.c.verhoeks@dnb.nl.

§Vrije Universiteit Amsterdam and Tinbergen Institute; w.f.c.verschoor@vu.nl.

¶Vrije Universiteit Amsterdam and Tinbergen Institute; r.zwinkels@vu.nl.

1 Introduction

Is the time varying relationship between exchange rates and fundamentals due to the limited attention of market participants? We investigate this question by measuring the attention to different economic fundamentals through the Google Trend Index. We demonstrate that the higher the attention for a certain economic fundamental, the better its ability to forecast exchange rate movements.

By measuring the relative attention to different fundamentals, we detect the factors that market participants use to form expectations about the future exchange rate. Due to the self-referential structure of exchange rates, the fundamental that market participants pay attention to feeds back into the exchange rate itself and hence determines its future dynamics. We exploit this feature and identify fundamentals that market participants focus on to construct predictions. We proxy investor attention for an economic variable by its relative search intensity on the Google search engine. The interpretation of the Google Trends Index (GTI) in this context is rather intuitive: the higher the GTI, the higher the relative search intensity for the fundamental in question, and the higher the attention paid to that fundamental.

By using the Google search volume as indicator of attention, we are not necessarily implying that all market participants are using Google to search for information. Instead, we assume that search volume represents a proxy of a more general measure of attention, which can affect investment decisions of various market participants. BIS (2019) decomposes the turnover in the foreign exchange markets by participant and documents that the inter-dealer market accounts for 40% of total foreign exchange market turnover while 55% of global trading volume is carried by other foreign exchange players including small non-reporting banks, retail investors and official sector financial institutions. The most heavily traded instrument is foreign exchange swap which accounts for almost half of global trading and is primarily used to manage liquidity and to hedge currency risk. Since a large share of foreign exchange turnover is driven by non-speculative activities, it is likely that the Google Trend Index captures the attention of a non-negligible share of currency markets' participants whose decisions impact prices. In fact, we show that the fundamental selected by the GTI is an excellent predictor of the future exchange rate and therefore it is likely that it captures the market-wide attention to fundamentals.

In a sample ranging from 2004 to 2016, we find that the GTI-based forecasting procedure significantly outperforms the random walk, both statistically and economically. It reduces the mean squared prediction error significantly for 9 out of 15 bilateral exchange rates. Moreover, it generates positive and significant excess returns for 8 out of 15 individual currencies and for all portfolio investment strategies. The size of economic profits is considerable: an equally-weighted currency portfolio, for instance, earns an annualized return of 4.9% with a Sharpe ratio of 1.3.

We find that the attention paid to the exchange rate fundamentals is limited. The best forecasts and the highest returns are systematically delivered by strategies that select the fundamental that is paid the highest attention to. Alternative models, including the single fundamental forecasts, equally-weighted or GTI-weighted fundamentals all underperform relative to the GTI selected predictor. We also show that the attention to the exchange rate fundamentals is short-lived and the chosen fundamental remains a predictor for on average 2.26 months.

We confirm that the GTI represents a general measure of attention as we show that it is more responsive to movements of economic fundamentals that (i) are subject to important policy decisions, (ii) for which new data is being released, or (iii) which exhibit large absolute movements.

Our results are robust to a number of alternative empirical specifications. Changing the exact set of search queries does not alter the results substantially. The results also hold throughout the entire sample period. In fact, we find that the performance of the GTI-based forecasting procedure improves in periods of financial turmoil, which is consistent with Sicherman, Loewenstein, Seppi, and Utkus (2016) who show that attention is more limited in periods of high market stress. Finally, we show that our model improves on a number of alternative benchmarks, such as the carry trade strategy as well as individual fundamental models.

Literature on time varying link between exchange rates and economic fundamentals is abundant. Schinasi and Swamy (1989) show that exchange rate models with time varying parameters outperform a random walk in an out-of-sample forecasting test. Cheung, Chinn and Pascual (2005), Sarno and Valente (2009), and Rossi (2013) find that fundamental exchange rate models have predictive power, but the performance depends on the particular currency and forecast horizon considered. Recent theoretical studies provide explanations for the time varying link between currencies and fundamentals. These papers build on the asset pricing representation of the exchange rate where its dynamics are mainly determined by fluctuating investors' expectations. These expectations may change for several reasons. For example, in the scapegoat theory of Bacchetta and Van Wincoop (2004, 2006, 2013), investors wrongly attribute currency movements to an economic variable that happens to change at the same time as the true movement-generating factor. In the model of Markiewicz (2012), investors are continuously testing the set of fundamentals and keep only the ones that significantly influence the currency dynamics. These models imply that investors pay attention to a certain fundamental (or a subset of fundamentals) at the time. The focus of attention is determined by a backward-looking mechanism and the selected fundamental is used to create future exchange rate expectations.

Recent empirical literature confirms predictions of these theoretical studies. Limited attention is a common feature of agents in financial markets and can account for a set of unexplained asset pricing phenomena. Kouwenberg, Markiewicz, Verhoeks and Zwinkels (2017) show that a strategy that puts more weight on the model that forecasted relatively well in the recent past outperforms a random walk forecast. Gholampour and van Wincoop (2019) use Twitter data to measure the opinion about the future movement of the euro - U.S. dollar exchange rate, and find evidence in support of the dispersed information model of Bachetta and Van Wincoop (2006). Da, Engelberg and Gao (2011) show that an increase in the search volume index predicts higher stock prices in the next two weeks and a price reversal over the following year. Sicherman, Loewenstein, Seppi, and Utkus (2016) find that investors pay less attention to news when the VIX is high, and that the level of attention is strongly related to investor's demographics (gender, age) and financial position (wealth, holdings). Yuan (2015) shows that attention-grabbing events predict the trading behavior of investors and hence market returns. Ben-Rephael, Da, and Israelsen (2017) study institutional investors' attention and they find that price drifts following both earnings' announcements and analysts' recommendations are driven by announcements which are ignored by institutional investors. Search intensity has been used before by Da, Engelberg, and Gao (2011), among others, to measure investors' attention. Mondria, Wu, and Zhang (2010) use, for instance, AOL dataset based on search queries to proxy for attention allocation and to partially explain the home bias of U.S. investors. Instead of measuring the aggregate investors' attention to the asset (currency), we use the GTI to capture the relative attention paid to different fundamentals at different points in time and map them into an exchange rate specification.

The remainder of this paper is organized as follows. Section 2 outlines our measure of attention based on the Google Trends data and describes the construction of our data on attention. In Section 3, we present the fundamental models and describe how we map them into the search queries. In Section 3, we also describe the data on fundamentals, explain the forecasting procedure based on the GTI, and present a set of statistical and economic measures that are used to assess the performance of the GTI-based predictions. The results are presented in Section 4 and Section 5 concludes the paper.

2 Google Trends

We build the measure of market participants' attention to fundamentals based on the Google Trends Index (GTI) data. In what follows, we describe in detail the GTI and demonstrate why it is an appropriate proxy for attention.

2.1 Google Trends as a measure of investor attention

Our aim is to capture the macroeconomic fundamentals that attract the attention of currency market participants by using data from Google Trends.[2] If market participants change the

[2] Da, Engelberg, and Gao (2011) use Google search volume to proxy for attention of individual stocks and find that it is a more timely measure than other measures of attention. Search engine data has recently been applied in several forms of so-called nowcasting of macro-economic variables (Choi and Varian, 2011). Goel, Hofman, Lahaie, Pennock and Watts (2005) show that search behavior also has predictive content further in the future.

importance they attach to each fundamental by shifting their attention over time, this influences their trading behavior and thereby feeds back into the exchange rate. With Google search volume, we aim to capture the fundamentals that investors are shifting their (scarce) attention resources to.

Google Trends is a public web facility that shows how frequently a particular term is searched on Google, relative to the total search volume across various regions and in various languages. Internet users typically use a search engine to search information, where Google holds, by far, the largest share of the web-search market (91%). More critically, search is a measure of *revealed* attention: if someone searches for a certain variable, this person is undoubtedly paying attention to it. This is a key advantage compared to indirect proxies of attention, like trading volume or extreme returns (as employed by Barber and Odean (2008), for instance). Google Trends may be thus an effective tool in identifying the fundamentals that economic agents focus on and use to predict currencies.

Importantly, we are not claiming that (all) market participants are using Google to search for information on economic fundamentals. Indeed, it is unlikely that sophisticated market participants use Google as as their primary source of information. For instance, Ben-Rephael, Da and Israelsen (2017) show that institutional investor attention for certain stocks can be measured better by the news searching and news reading activity on Bloomberg terminals. Instead, we argue that Google search volume represents a proxy of a more general measure of attention. If a particular economic fundamental attracts more search volume, it implies that is has attracted the attention of Google users and is likely to have attracted the attention of market participants as well. In addition, we do not need to capture the attention of *all* market participants, but of a non-negligible portion large enough to affect prices.

2.2 Construction and interpretation of the Google Trends Index

The data provided by the Google Trends tool consists of time series indices of search queries that Google users enter into the search engine. The index is a measure defined relative to all other search queries in a given location and period, and it ranges between 0 and 100.

For instance, Google Trends enables us to retrieve the index for the query *inflation*, that has been searched for by Canadian users between 2004 and 2016. The resulting index will only use search activity within Canada and Google Trends will re-scale the index over the 2004-2016 period.

There are two ways of computing the index to measure the attention to our set of fundamentals. The first one relies on individual searches where the Google Trends Index is computed for each search query separately. The second, preferred one, uses the "compare" function, where Google applies an additional rescaling of the search volume for the queries relative to each other.

Figure 1 in the appendix demonstrates why the relative index is appropriate in our context. It shows the difference between using the two different ways of retrieving the Google Trends data for three example search queries: *interest rate*, *oil price* and *trade balance.* The top panel of the figure plots the results of the exercise with the "compare" function, and shows that each query exhibits a different long-run mean. For example, *interest rate* has the highest search activity most of the time and the *trade balance* is always below the others. Because we are interested in developments in attention over time, we seek to subtract the long-run average from the queries, what the individual search query procedure effectively does (bottom panel of Figure 1). The reason is that the differences in the long-run average search intensity are unlikely to be related to the shifts in attention. For instance, *interest rate* is naturally a query with high average search activity in comparison with others because of its interest to broader public than *trade balance.* Its high long-run average search intensity could be for instance driven by households searching for interest rates on mortgages and other loans, activities that continuously take place and do not reflect changes in attention. In the appendix, we explain additional details on how we use Google Trends data.

To increase the response speed, Google calculates the index from a random subset of the actual historical search data. As a result, the GTI time-series on the same search term are often slightly different when they are downloaded at different points in time. On average, the impact of this sampling error is small and would work against our results. To evaluate the possible noise introduced by this sampling error, we compute the correlation between the

GTI data series download twice for five different queries (monetary policy, money supply, interest rate, inflation and GDP). This is done for two different samples and frequencies: monthly data (January 2004 - December 2015) and weekly data (January 2015 - December 2015). Correlations between downloaded time-series vary between 0.985 and 0.996.

To illustrate to what extent the GTI captures the attention of market participants, we show the behavior of the GTI around meetings of the Governing Council of the European Central Bank (ECB). Specifically, we compute the average increase in the weakly GTI, relative to the prior week, for the search query 'monetary policy' during the weeks when the council took monetary policy related decisions. The GTI increase is computed for six euro area countries and for all 20 meetings that took place during 2014 and 2015 and is reported in Table 1. The statistics indicate that the GTI captures the attention that market participants pay to monetary policy around critical Frankfurt meetings. In the vast majority of cases, the GTI increases considerably (averaged over the year) during weeks in which the council meets.

Additionally, in Figure 2, we plot the GTI dynamics for the quantitative easing (QE) query, between January 2014 until September 2016 and for Germany, France, Italy and Spain. During this period, the Governing Council decided on a set of fundamental changes in its monetary policy, involving several QE programs. The first event documented in the figure, occurred on September 4 2014, when the ECB decided to implement QE via purchases of asset backed securities and covered bonds. On January 22 2015, the ECB announced a new purchase program of sovereign bonds. While the first event was mainly captured by the spike in the German GTI, the second event generated a spike of the index in all four countries, as documented in Figure 2. The third reported event was the actual start of the purchase program, where the GTI again spikes in all four countries. These examples clearly indicate that the GTI succeeds to capture swings in attention.

3 Model selection with Google trends

The GTI captures swings in attention to different economic fundamentals and we expect that these fundamentals affect economic decisions of the currency market participants.[3] We translate each of these economic fundamentals into an exchange rate model and we estimate it. We then use the GTI to select the model that the economic agents pay the most attention to and use it for predictions. In what follows, we explain in detail the fundamentals and how the Google search intensity is used to make a selection among them. Next, we describe how we create the predictions based on the fundamental model selection.

3.1 Pool of fundamental models

We compile a pool of macroeconomic and financial fundamental models that are commonly used in the literature to forecast exchange rates. This set includes macroeconomic models that contain variables which are expected to predict exchange rate movements, as well as financial models taking into account commodity prices. This pool of models matches partly that of Kouwenberg, Markiewicz, Verhoeks and Zwinkels (2017).[4] Note that all variables are in natural logarithms, such that Δs_t is the period t currency return. Each of our eight models includes a constant term.

The first fundamental model is the uncovered interest rate parity (UIRP), which states that exchange rate movements are explained by differences in the nominal interest rate:

$$\Delta s_{t+1} = \alpha_1 + \beta_1[i_t - i_t^*] + \varepsilon_{t+1}, \tag{1}$$

in which i is the interest rate and an asterisk $(*)$ denotes a foreign variable. The sign of β_1 determines whether this model represents uncovered interest rate parity ($\beta_1 < 0$) or a carry-trade model ($\beta_1 > 0$). We take a combination of the 3 months and 1-year interest rates.

[3]By economic decisions we mean here investment and hedging decisions, central banks interventions, goods and services trading decisions, etc.

[4]We left out some common exchange rate models because of their limited applicability in combination with the Google Trends Index. One of them is the Taylor rule-based model which is difficult to translate into the GTI query because of its complexity. We focus on the simple models that can be easily translated into the Google search queries.

The second model is based on the purchasing power parity (PPP). Its absolute version implies that the expected exchange rate is a function of the deviation of the spot rate from its PPP-based fundamental value:

$$\Delta s_{t+1} = \alpha_2 + \beta_2[(p_t - p_t^*) - s_t] + \varepsilon_{t+1}, \tag{2}$$

in which p_t and p_t^* are the home and foreign price levels, respectively. Next to the absolute PPP, we also consider its relative variant, which relates exchange rate movements to the inflation differentials rather than absolute price differences:

$$\Delta s_{t+1} = \alpha_3 + \beta_3[(\pi_t - \pi_t^*) - s_t] + \varepsilon_{t+1}. \tag{3}$$

We combine the forecasts from the absolute and relative PPP models into one forecast because they are based on the same underlying fundamental. As a result, we use the same search queries for both models.

The third model is the canonical monetary model which presents exchange rates in terms of relative money demand and supply in the two countries involved:

$$\Delta s_{t+1} = \alpha_4 + \beta_4[(m_t^* - m_t) - k(y_t^* - y_t) - s_t] + \varepsilon_{t+1}, \tag{4}$$

in which m is the log money supply and y is log output (GDP). Following Molodtsova and Papell (2010), we set k equal to 0.5.

The fourth model, the international risk sharing model, relates the exchange rate to differences in consumption growth between the two countries involved:

$$\Delta s_{t+1} = \alpha_5 + \beta_5[CG_t - CG_t^*] + \varepsilon_{t+1}, \tag{5}$$

in which CG is the (annual) consumption growth. The intuition behind this model is that a relatively high consumption growth leads to a higher income, a higher money demand and, consequently, to a stronger currency.

International trade models suggest that the exchange rates respond to the trade imbalances and our fifth model is as follows:

$$\Delta s_{t+1} = \alpha_6 + \beta_{6,1}(TB_t/GDP_t) + \beta_{6,2}(TB_t^*/GDP_t^*) + \varepsilon_{t+1}, \tag{6}$$

in which TB is the trade balance. We scale the trade balance by GDP to control for the size of the economy.

The capital flows model incorporates the net foreign asset position, which is defined as the difference between purchases and sales of securities. This model states that exchange rates move in response to imbalances in the market for capital and is defined as follows:

$$\Delta s_{t+1} = \alpha_7 + \beta_7 NFA_t + \varepsilon_{t+1}, \tag{7}$$

in which NFA_t stands for the position in net foreign assets.

In addition to the fundamentals discussed above, which are all derived from macroeconomic reduced-form models, we also incorporate two financial factors. The first one relates exchange rate movements to changes in commodity prices:

$$\Delta s_{t+1} = \alpha_8 + \beta_8 \Delta COM_t + \varepsilon_{t+1}, \tag{8}$$

in which ΔCOM_t represents the commodity price index change.

The second financial factor is oil. Oil seems to be the most important commodity related to changes in the exchange rates and therefore we look at changes of its prices separately in the following specification:[5]

$$\Delta s_{t+1} = \alpha_9 + \beta_9 \Delta OIL_t + \varepsilon_{t+1}, \tag{9}$$

in which ΔOIL_t represents the movement in the oil price index.

3.2 Mapping models into search queries

In order to implement the GTI we need to map the models into relevant search queries. For each of the models, to measure investors' attention, we select several search queries and we report them in Table 2. To select multiple search terms, we make use of Google Correlate.

[5] See for instance Chen, Rogoff and Rossi (2010).

This tool suggests search queries that are highly correlated in terms of search intensity with the one entered by the user. For example, in the case of monetary policy, Google Correlate informs us that monetary supply and monetary demand display an average correlation of 0.903 between 2004 and 2015. They are thus natural candidates for search queries for the monetary model.

The search terms shown in Table 2 are always restricted to the countries of the currency involved. For instance, when we search for fundamentals for the USD-JPY, we restrict the regional area of the GTI to the United States and Japan. It is important to bear in mind that Table 2 presents the Google search queries in English. English search queries have the highest coverage, also for Japan where English is not the main language.[6] However, when the English search query does not have enough data to calculate the GTI (which is automatically indicated by Google), we use the query both in English as well as in the home language. If this procedure still does not deliver enough Google coverage, we omit the particular search term from the analysis.

Table 3 reports the correlation coefficients between the GTI series of the different search queries for each of the U.S. economic fundamentals. The table documents that, in most cases, correlations are positive and strong, confirming that individual search queries for the same fundamental are related. Still, applying multiple search queries may have important benefits: it 'diversifies' away the idiosyncratic risk that adheres to each individual search query.

3.3 Forecasting procedure

Based on the GTI, we build the forecasting procedure to predict monthly exchange rate movements between January 2004 and December 2016.[7] We perform the analysis for fifteen currency pair combinations between the CAD, GBP, JPY, CHF, AUD, and USD.[8] We have not extended the analysis to other currencies due to limitations of Google Trends. More precisely, Google coverage decreases considerably outside the countries we incorporated into

[6]The reason could be that even in Japan most investors read international news in English and hence search for queries in English.

[7]The sample period could not be extended because Google Trends data only dates back to January 2004.

[8]By taking all the currency pair combinations, are results are not dependent on the choice of base currency.

our analysis. Also, including the euro is challenging as it involves multiple countries and languages.

For clarity, we describe our forecasting procedure in the following steps:

Step 1 - Collecting GTI data on search queries

First, we collect GTI data for all search queries that are listed in Table 2, for each point in time and for each country. We choose the monthly frequency to match the frequency of the fundamentals' series. This means that we collect 210 GTI series, as we selected 35 search queries (see Table 2) for six countries.

Step 2 - Average GTI for each fundamental

Second, we calculate the average GTI over all the search queries per fundamental model for every point in the sample period and for each country separately.

Step 3 - Selection of fundamental model with relatively highest attention

Third, at each point in time, we select the fundamental model with the highest relative average GTI for its search queries, averaged over both countries involved. The model with the highest relative average GTI for its search queries is the model for which Google users have the highest attention in the country involved. Our hypothesis is that this model will have the best predictive performance because it incorporates economic information that receives the most attention.

Step 4 - Forecast selection

Finally, we take the one-step-ahead out-of-sample forecast of the model selected by the GTI. This procedure is recursive and is repeated for each data point.

3.4 Data on fundamentals

We use monthly data for 1995M1-2016M12 for Australia, Canada, Japan, Switzerland, the United Kingdom and the United States from different data sources. We use seasonally adjusted data for the GDP from the National Accounts database of the OECD. Private consumption data is taken from the Key Short Term Economic Indicators of the OECD. Data for the seasonally adjusted trade balance, net foreign assets, monetary base (M3), and CPI are retrieved from the Main Economic Indicators dataset of the OECD, where we define net foreign assets as the sum of foreign direct investments and portfolio investments. If applicable, these data are in national currency. We use two commodity data series. The composite commodity price index comes from Goldman Sachs. The oil price is measured by the price per barrel of West Texas Intermediate (WTI). Interest rates are LIBOR rates for both the three months and one-year maturity.

We estimate all eight models described above in-sample and make an out-of-sample forecast for the first observation thereafter. The first in-sample estimation runs from 1995M1 to 2003M12 and returns an out-of-sample forecast for 2004M1. This is the first month for which Google Trends data is available and we can perform the procedure set out in the previous subsection. The second estimation period runs 1995M1 to 2004M1 and delivers an out-of-sample forecast for 2004M2, and so forth. We create these rolling forecasts with expanding window for each model separately, after which we apply our selection mechanism based on the Google Trends Index. The sample period of January 2004 - December 2016 contains 156 monthly forecasts for each of the 15 currency pairs.

3.5 Statistical and economic performance measures

Once we have the predicted exchange rate series, we implement a set of measures to judge the ability of the GTI to trace the fundamentals that investors focus on. First, we measure forecast accuracy of model i by the mean squared prediction error (MSPE) relative to the MSPE of the random walk model. In comparing the forecast errors, we use the Clark and West (2007) adjustment and statistic, given by:

$$CW_i = \frac{1}{l}\sum_{t=1}^{l}\frac{(E^i_{t-1}(\Delta s_t) - \Delta s_t)^2 - E^i_{t-1}(\Delta s_t)^2}{(\Delta s_t)^2}, \tag{10}$$

in which l is the total number of forecasts. We employ the CW statistic to test the null hypothesis that the GTI-based forecast has the same predictive ability as a random walk without drift benchmark model.

In addition, we calculate the non-parametric sign test of Pesaran and Timmermann (1992), which tests the ability of our model to forecast the direction of change correctly, again relative to the random walk forecast. The Pesaran and Timmermann (1992) statistic is given by:

$$PT_i = \left(\frac{p^*(1-p^*)}{l} \right)^{-1/2} (\widehat{p} - p^*), \tag{11}$$

in which p^* is the benchmark proportion of correct sign predictions and $\widehat{p}$ the observed proportion of correct sign predictions. We set $p^* = 0.5$ to represent the random walk without drift model. The PT statistic is asymptotically distributed as $N(0,1)$.

In addition to statistical forecast evaluation measures, we also examine the economic value added of our forecasting procedure. For this purpose, we calculate the returns of an investment strategy that buys (sells) one unit of the foreign currency when the model predicts an appreciation (depreciation) of the foreign currency. In case of subsequent buy-signals, the long position is rolled over. The foreign-exchange return of the strategy i is given by:

$$r_t^i = \frac{E_{t-1}^i(\Delta s_t)}{|E_{t-1}^i(\Delta s_t)|} \Delta s_t. \tag{12}$$

The investment strategy's Sharpe ratio is calculated as $\frac{1}{l}\sum r_t^i / \frac{1}{l}\sum (r_t^i - \overline{r}^i)^2$. By definition, this measure also evaluates the performance of the strategy relative to the random walk without drift benchmark, because the benchmark random walk strategy always yields the risk-free return.[9]

In addition to assessing model performance for each currency individually, we form equally-weighted and volatility-weighted portfolios of all currencies. We measure and forecast volatility using the exponentially weighted moving average (EWMA) method, given by:

[9]For simplicity, we ignore transaction costs when calculating strategy returns. Given the liquidity of foreign exchange markets and the relatively low frequency of our analyses, however, we do not expect the transaction costs to affect the results much.

$$\sigma_{t+1}^2 = \lambda \sigma_t^2 + (1-\lambda)\Delta s_t^2, \tag{13}$$

in which we follow the RiskMetrics approach and set $\lambda = 0.94$. The return of the volatility weighted portfolio is given by:

$$r_t^{vw} = \left(\sum_p \frac{1}{\sigma_{p,t+1}^2} \right)^{-1} \sum_p \frac{1}{\sigma_{p,t+1}^2} r_t^p, \tag{14}$$

in which p denotes the currency.

Note that the equally-weighted and volatility-weighted portfolios of currencies can contain both long and short positions, depending on the signs of the return forecasts for the currencies. As an additional performance measure, we also form pure long-short portfolios based on the forecasts. Specifically, we construct a portfolio that goes long in the currency with the highest forecasted appreciation and short in the currency with the lowest forecasted appreciation.

4 Results

We first assess to what extent the GTI selected fundamental carries predictive power and delivers statistically and economically significant returns. Second, to understand how inattentive economic agents are, we include additional fundamentals into the forecasts and evaluate their predictive performance relative to the most salient, GTI selected fundamental. Finally, we investigate the properties of the GTI-based fundamentals to assess the persistence of attention.

4.1 Limited attention and currency forecasts

Table 4 presents results of the currency forecasting exercise where the fundamental is selected by the GTI at each point in time.The MSPE ratios are below unity for 12 out of 15 currencies, and significantly so for 9 out of 15 currencies. The average improvement is 3.3% based on the MSPE, ranging from 8.9% for the USD-AUD to -1.1% for the CHF-JPY. The p-value of the PT-statistic reports significance for six out of fifteen currencies.

Table 5 presents the economic evaluation of our forecasting exercise. The table shows the annualized return, the annualized standard deviation of returns, the corresponding Sharpe ratio, and the p-value of the t-test to test the significance of the average return. The investment strategy for individual currencies results in a positive investment return for all currencies except one (CHF-JPY). For eight currencies, the positive return is significant. Moreover, the three currency portfolios (equally-weighted, volatility weighted, and long-short) all have a significantly positive mean return. Given the implemented strategies, the magnitude of the portfolio returns is considerable. The equally-weighted portfolio for instance has an annualized investment return of 4.9% with an annualized volatility of 3.7% and a Sharpe ratio of 1.297. The long-short portfolio earns a very high return of 10.8%, although this is achieved at the cost of a (much) higher volatility (17.7%).

4.2 How inattentive are economic agents?

If market participants attention is limited, the best forecasts should be delivered by strategies based on only a small set of (or one) fundamentals identified by the GTI. We investigate this hypothesis by comparing benchmark, single fundamental based results to a setting in which more than one factor is considered. Specifically, we implement forecasting and investment strategies based on a set of fundamentals weighted by their Google search intensities and with equal and constant weights.[10] The weights are computed as follows:

$$E_t \Delta s_{t+1} = \sum_{i=1}^{N} w_{i,t} \left[\widehat{\alpha}_{i,t} + \widehat{\beta}_{i,t} f_{i,t} \right], \tag{15}$$

in which i is the subscript for the different fundamental models and $w_{i,t}$ is defined as follows:

$$w_{i,t} = \frac{GTI_{i,t}^{av}}{\sum_{i=1}^{N} GTI_{i,t}^{av}}, \tag{16}$$

in which $GTI_{i,t}^{av}$ is the average GTI computed over the different search queries for fundamental i for the two countries involved in a given currency. The weights in an equally weighted exercise are set to $w_i^{ew} = \frac{1}{N}$, with N being the number of fundamentals. Table 6 presents

[10] Other studies implementing model averaging to improve on forecasting performance include, for instance, Timmermann (2006), Della Corte, Sarno and Tsiakas (2009) and Wright (2008)

the average GTI-based weights attached to each of the fundamental models. It also shows the difference between the mean maximum and the mean minimum weights averaged over the currencies involved. The maximum weight is about five times higher than the minimum weight so undoubtedly the relative GTI generates substantial differences in weights attached to fundamentals. The least relevant fundamental receives the average minimum weight of 4.3%, which is not negligible.

Table 7 presents the results of the statistical evaluation of the forecast combinations. The MSPE relative to the RW reduces by 0.89% on average, and the reduction is only significant for three exchange rates. The PT-statistic reports significant out-performance for three exchange rates. The economic evaluation presented in Table 8 is somewhat more positive. The annualized returns are positive and significant for three exchange rates and all three portfolios. However, the absolute size of returns is substantially lower compared to our forecasting procedure which selects one model with the highest attention measure. The bottom row in Table 9 shows the qualitatively similar results for strategies based on the equally weighted forecasts. They deliver positive and significant Sharpe ratio, but of much smaller magnitude than the forecast based on the GTI with model selection. These findings are consistent with the results of Kouwenberg, Markiewicz, Verhoeks and Zwinkels (2017), who show that selecting a single model works best in forecasting exchange rates. These findings also suggest that the attention of economic agents is very limited.

4.3 How persistent is the attention?

Strategies with constant weights implying the absence of attention shifts perform poorly relative to the GTI based strategies. Table 9 presents the Sharpe ratios of three portfolios for static trading strategies that are based on each fundamental model separately and the last row shows results for equally weighted static strategy. The outcomes show that none of the static strategies results in a higher economic profit based on the Sharpe ratio suggesting that the attention of market participants is not very persistent. Table 10 reports how often the selected fundamental changes and implied duration in months. The selected model changes in 46% of cases (months) on average. Accordingly, the chosen fundamental remains for on average 2.26 months before it is replaced by another fundamental model. This implies that the fundamental gaining the highest average attention is rather quickly substituted by another. The extent of time variation differs slightly across currencies. For instance, for the AUD-CHF, the time-variation of the chosen fundamental model is the highest, with switches in 62% of the months and the selected model being selected for a duration of 1.61 months on average. For the USD-JPY, the model only changes in 23% of the time leading to an average duration of 4.11 months.

Table 10 also shows how often each model is chosen for each of the currencies. The consumption growth, trade balance, and commodity fundamentals are selected most often. The net-foreign assets and interest rate parity models are hardly ever chosen. There is, however, ample variation across currencies. For example, the trade balance model is chosen in 79% of the months for the USD-JPY. This might be explained by the export-oriented Japanese economy and in particular by the U.S. trade deficit of 67 billion with Japan, which attracted considerable attention in recent years.

4.4 What attracts attention?

We documented in Section 4.3 that fundamentals selected by the attention measure change rather frequently. What triggers the attention to shift between economic fundamentals? We study three types of attention triggers. First, motivated by findings in Section 2.2, we analyse the impact of policy decisions on the GTI movements. Second, we study how the releases of economic figures by statistical agencies influence the GTI dynamics. Finally, motivated by the scapegoat theory, we investigate the impact of large (absolute) shifts in fundamentals on the attention swings.

As already documented in Section 2.2, important ECB monetary policy decisions resulted in an increased attention for search terms related to monetary policy. We study several monetary policy decisions taken by the Fed. During the period 2009-2014, the Fed took several far-reaching decisions in terms of accomodative monetary policy. During this period, the Federal Open Market Committee (FOMC) had 48 meetings to decide upon the direction of monetary policy. At the time of these meetings, the GTI of the search query 'quantitative easing' increased on average by 33.7% compared to the week prior to these meetings. Figure 3 shows that the GTI of 'quantitative easing' rose considerably in response to important FOMC decisions. Hence, policy decisions of relevant institutions (such as the ECB or Federal Reserve) appear a potential trigger of investor attention.

As a second potential source of attention variation, we study whether new data publications by statistical offices attract market participants attention by investigating the GTI around a set of data releases by the U.S. Bureau of Labor Statistics (BLS) during the period 2013-2015. We focus on the BLS releases of figures on inflation (Consumer Price Index (CPI), the Producer Price Index (PPI)), and labor markets (Employment Situation). Table 11 shows that, on publication dates, the GTI of the related search term is considerably higher for all three statistical releases and for all years during the period between 2013 and 2015. For instance, using daily Google Trends data, the GTI for inflation was on average 14.1% higher on CPI release days compared to other days in 2013.

Finally, as also suggested by the scapegoat theory, large (absolute) movements in fundamentals can be a source of attention attraction as well. For instance, Figure 4 shows that high values of the GTI for the search query oil price often coincided with large movements in the oil price itself. Although the monthly absolute oil price change (right axis) is more volatile, is it clear that the largest swings in oil prices come together with an intensified attention as proxied by the GTI.

Based on this evidence, we can conclude that attention is more likely drawn towards economic fundamentals (i) that are subject to important policy decisions, (ii) of which data is being released by statistical agencies or (iii) show large absolute movements. Hence, swings in attention are at least partly caused by investors who divide their attention over different fundamental information sources over time.

4.5 Robustness tests

4.5.1 Importance of different queries

To measure the attention to a certain fundamental variable, we incorporate a number of search queries to create the Google Trends Index for each fundamental model, as listed in

Table 2. An important robustness check is to assess whether our results are sensitive to the selected queries. Tables 12 and 13 show the statistical and economic evaluations of forecasts after deleting the last query for each model in our pool from the list in Table 2, provided that at least three search queries per model remain available.

The statistical evaluation of the resulting forecasts shows that the elimination of these queries does not qualitatively impact results. On the level of individual currencies, some changes do occur. For instance, the statistical outperformance based on the PT-statistic disappears for the CAD-CHF and the AUD-GBP. The economic evaluation of portfolios' returns in Table 13 shows a similar picture. After the elimination of the queries, the estimated mean returns slightly decrease, but they remain economically and statistically significant. We also remove other search queries and find (unreported) qualitatively unaffected results. The limited impact of elimination of queries is in line with expectations since the GTIs of the queries for each fundamental are strongly correlated.

4.5.2 Sample period

To assess the robustness of our findings to the sample period selection, Figure 5 presents 3-year rolling returns and Sharpe ratios for all the three investment portfolios that we evaluate. The graphs show that the rolling returns and Sharpe ratios oscillate around the total sample mean (represented by the dashed line), but never become insignificant or negative for a longer period of time. It is worth noting that all the three portfolios perform particularly well during the Great Recession.

We also formally test for differences in performance during the first and second half of the sample period: 2004M1-2010M6 and 2010M7-2016M12, respectively. Table 12 and 13 also include the results for these two periods separately. Even if the first half of the sample covers the financial crisis period, the results do not differ much across both samples.

4.5.3 Recessions and financial distress

Motivated by visible strong performance of the GTI-based strategies during the last crisis in Figure 5, we investigate further the profitability of our attention-based prediction tool during financial distress. Table 13 includes the results for recessions separately (rows entitled

’Recessions’), where recessions are periods during which one of the countries involved was in a recession according to the OECD definition. It also displays the findings for periods of financial stress (rows entitled VIX 10-percentile), defined as the periods during which the value of the VIX was in its upper 10th percentile. The results show that our attention-based strategies perform, in fact, better during stress and recessions. The difference in returns is statistically significant when compared to calm and expansion periods. These results are consistent with Sicherman, Loewenstein, Seppi, and Utkus (2016), who show that investors pay less attention to news when the VIX is high and hence limited-attention based measures score particularly well.

4.5.4 Attention based strategies versus carry trade

Carry trading is a popular currency investment strategy that aims to exploit deviations from the uncovered interest rate parity by investing in high interest rate currencies and borrowing in low interest rate currencies. This strategy generated significant investment profits in the past, particularly in the pre-crisis period (Burnside, Eichebaum, Kleshchelski and Rebelo, 2011). To compare our results with carry trades, we take the perspectives of an American, Australian, British, Canadian, Japanese and Swiss investor, respectively, and compute the returns of a long-short currency portfolio for all of them. We form a portfolio by sorting the six currencies each month from high to low interest rate and we exclude the home currency of the investor. Next, we go long in the two currencies with the highest interest rate and go short in the two currencies with the lowest. We also calculate the returns of an investment in one leg only (long and short), again with two currencies per leg. The portfolio weights of the currencies are always equal and we re-balance each month.

Table 14 shows that the results of the carry trade strategies are characterized by lower (and sometimes even negative) Sharpe ratios. Also, in the case of positive carry trade returns, they are statistically insignificant. The one exception is the long leg for the U.K. investor, which yields a significant return. These results show that our forecasting procedure based on the relative search intensity for different fundamentals outperforms simple carry strategies. The results for the carry trade strategies reported in the literature are usually higher than reported here. This is most likely explained by the sample period, which is relatively short and includes the break-down of carry during the Financial Crisis.

5 Conclusion

The weight that economic agents attach to different fundamentals as drivers of currency movements fluctuates considerably over time. In this paper, we demonstrate how economic agents' limited attention can account for this empirical feature of the exchange rates. We proxy the attention of economic agents by the Google Trends Index. In a sample of macroeconomic data from 1995 to 2016, we find that the forecasts based on fundamental selection by Google Trends significantly outperform the random walk, both statistically and economically. The size of economic profits is considerable: an equally-weighted currency portfolio, for instance, earns an annualized return of 4.9% with a Sharpe ratio of 1.3. By comparison, between 2004 and 2016 a carry trade strategy based on the same currencies yielded a Sharpe ratio of 0.31.

The best forecasts and the highest returns are delivered by strategies that select the fundamental that is paid the highest attention to. Other fundamentals with high GTI should be discarded in the construction of the forecasts and corresponding investment strategies. Dynamic strategies significantly outperform static, single-based and equally-weighted fundamental strategies. Under dynamic GTI-based strategy, the chosen fundamental remains for on average 2.26 months, before it is replaced by another one. These findings suggest that the attention of economic agents is limited and short-lived.

References

[1] Andrei, D. and Hassler, M. (2014), 'Investor attention and stock market volatility', *Review of Financial Studies*, 28, 33-72.

[2] Bacchetta, Philippe, and Eric van Wincoop (2004), 'A Scapegoat Model of Exchange Rate Determination,' *American Economic Review, Papers and Proceedings* 94, 114-18.

[3] Bacchetta, Philippe and Eric van Wincoop (2006), 'Can Information Heterogeneity Explain the Exchange Rate Determination Puzzle?' *American Economic Review* 96, 552-76.

[4] Bacchetta, Philippe and Eric van Wincoop (2013), 'On the Unstable Relationship between Exchange Rates and Macroeconomic Fundamentals,' *Journal of International Economics*, 91, 18-26.

[5] Barber, B.M., and Odean, T. (2008), 'All that glitters: the effect of attention and news on the buying behavior of individual and institutional investors', *Review of Financial Studies* 21, 785-818.

[6] Ben-Rephael, Da Z. and R. D. Israelsen (2017), 'It Depends on Where You Search: Institutional Investor Attention and Underreaction to News' *The Review of Financial Studies*, 30, 30093047.

[7] Bank for International Settlements (2019), 'Triennial Central Bank Survey: Foreign exchange turnover in April 2019', Monetary and Economic Department, Basel.

[8] Burnside, C., M. Eichebaum, I. Kleshchelski, and S. Rebelo (2011), 'Do peso problems explain the returns to the carry trade?', *Review of Financial Studies*, 24, 853-891.

[9] Chen, Y., K. Rogoff, and B. Rossi (2010), 'Can Exchange Rates Forecast Commodity Prices?', *The Quarterly Journal of Economics*, 125, 1145119.

[10] Cheung, Y., Chinn, M. D., and Pascual, A. G., (2005), 'Empirical Exchange Rate Models of the Nineties: Are Any Fit to Survive?,' *Journal of International Money and Finance*, 24, 1150-1175.

[11] Choi, H., and H. Varian (2009), 'Predicting the present with Google Trends', *Economic Record* 88, 2-9.

[12] Clark, Todd, and Kenneth D. West (2007), 'Approximately normal tests for equal predictive accuracy in nested models', *Journal of Econometrics*, 138, 291-311.

[13] Della Corte, P.; L. Sarno; and I. Tsiakas. (2009), 'An Economic Evaluation of Empirical Exchange Rate Models.' *Review of Financial Studies*, 22, 3491-3530.

[14] Da, Z.; J. Engelberg; and P. Gao. (2011) 'In Search of Attention', *Journal of Finance*, 66, 1461-99.

[15] Gholampour, V and E. van Wincoop (2019), Exchange rate disconnect and private information: What can we learn from Euro-Dollar tweets?, *Journal of International Economics*, 119, 111-139.

[16] Goel, S., J.M. Hofman, S. Lahaie, D.M. Pennock, and D.J. Watts (2010), 'Predicting consumer behavior with web search', *Proceedings of the National Academy of Sciences*, 7, 1748617490.

[17] Hirshleifer, D.; S.S. Lim; and S.H. Teoh. (2009), 'Driven to distraction: extraneous events and underreaction to earnings news', *Journal of Finance*, 64, 2289-2325.

[18] Hirshleifer, D., and S.H. Teoh. (2003), 'Limited attention, information disclosure, and financial reporting', *Journal of Accounting and Economics*, 36, 337-386.

[19] Kahneman, D. (1973), "Attention and Effort", Prentice-Hall, Englewood Cliffs, NJ.

[20] Kouwenberg, R., A. Markiewicz, R. Verhoeks, and R.C.J. Zwinkels (2017), 'Model Uncertainty and Exchange Rate Forecasting', *Journal of Financial and Quantitative Analysis*, 52, 341-363.

[21] Markiewicz, A. (2012), 'Model uncertainty and exchange rate volatility', *International Economic Review*, 53, 815–844.

[22] Mark, N.C., (1995), 'Exchange Rates and Fundamentals: Evidence on Long-Horizon Predictability', *American Economic Review*, 85, 201-218.

[23] Meese, R.A. and K. Rogoff (1983), 'Empirical exchange rate models of the seventies: do they fit out of sample?,' *Journal of International Economics*, 14, 345-373.

[24] Molodsova, T. and Papell, D.H. (2010), 'Out-of-sample exchange rate predictability with Taylor rule fundamentals', *Journal of International Economics*, 77, 167-180.

[25] Mondria, J., T. Wu, and Yi Zhang (2010), 'The determinants of international investment and attention allocation: Using internet search query data', *Journal of International Economics*, 82(1), 85 - 95

[26] Pashlers, H.E. (1998), 'The Psychology of Inattention'. MIT Press, Cambridge.

[27] Pashler, H. and J.C. Johnston (1998), 'Attentional limitations in dual-task performance', in H. Pashler, ed., Attention (1998), 155-189.

[28] Peng, L. and W. Xiong (2006), 'Investor attention, overconfidence and category learning', *Journal of Financial Economics*, 80, 563-602.

[29] Pesaran, M.H. and A. Timmermann (1992), 'A Simple Nonparametric Test of Predictive Performance', *Journal of Business and Economic Statistics*, 10, 461-465.

[30] Rensink, R.A., J.K. ORegan, and J.J. Clark (1997), 'To see or not to see: the need for attention to perceive changes in scenes', *Psychological Science*, 8, 368-373.

[31] Rossi, B., (2013), 'Exchange Rate Predictabibility', *Journal of Economic Literature*, 51(4), 1063-1119.

[32] Sarno, L., and G., Valente, (2009), 'Exchange rates and fundamentals: footloose or evolving relationship?', *Journal of European Economic Association*, 7, 786-830.

[33] Schinasi G.,J., and Swamy P.A.V.B., (1989), 'The out-of-sample forecasting performance of exchange rate models when coefficients are allowed to change', *Journal of International Money and Finance*, 8, 375-390.

[34] Sicherman, N., Loewenstein, G., Seppi, D. J., & Utkus, S. P. (2016), 'Financial Attention', *Review of Financial Studies*, 29, 863-897.

[35] Simons, D. J. and C. Chabris (1999), 'Gorillas in our midst: sustained inattentional blindness for dynamic events', *Perception*, 28, 1059-1074.

[36] Stroop, J. (1935), 'Studies of interference in serial verbal reactions', *Journal of Experimental Psychology*, 28, 643-662.

[37] Wright, J. (2008), 'Bayesian Model Averaging and Exchange Rate Forecasting', *Journal of Econometrics*, 146, 329-341.

[38] Yuan, Y. (2015), 'Market-wide attention, trading and stock returns' , *Journal of Financial Economics*, 116, 548-564.

Appendix - How Google Trends data is used

In this appendix we describe in more detail how we use the tool of Google Trends to extract the data series of the Google Trends Index, our proxy for investor attention for the different exchange rate fundamentals. First, it is important to specify the search query correctly. Google Trends gives different possibilities to use punctuation in searches to filter the results. Table A1 below shows how the use of punctuation influences the filtering of results by means of an example query from the monetary model.

Table A1: **Punctuation and GTI search queries: monetary model example**

Variable	**Definition**
money supply	Results include searches containing both *money* and *supply* in any order. For instance, results also include queries like *increasing money supply*, *supply of money*, and *supply of cash and deposit money*.
“money supply”	Results include the exact phrase inside the double quotation marks, possibly with words before or after. For instance, the query *money supply of cash* will also be included in the results.
money + supply	Results include searches containing the words *money* or *supply*. For instance, also a not related search query like *supply of labour* or *money spender*.
money - supply	Results include search containing the query *money*, but exclude searches including the word *supply*.

For all different searches above, it should be noted that no misspellings, spelling variations, synonyms, or plural or singular versions of the search query are included. Based on the consequences of the different uses of punctuation, we decided to always apply the double quotation marks in our queries (the second line in the table above). This option most appropriately filters the outcomes and hinders unintended queries to influence the results. On the other hand, using the double quotation marks is relatively strict: if we enter *money supply*, the results will not include queries like *supply of money*, because the order of words is different. However, to be as objective as possible, we apply strict filters to the Google Trends tool, such that the Index only identifies phenomena we are interested in.

The punctuation options in the table above can also be combined. For instance, if we would enter “money supply” + “money demand” the results will include all searches containing money supply or money demand. We have used the option to combine the punctuation options in the case of Japan. In many cases, the English version of our search queries did not have enough coverage for Google to calculate the index (which is automatically indicated by Google). In that case we entered both the English and the Japanese term for the query involved (using Google Translate). For instance, the English query "balance of trade" did not result in sufficient data coverage for Japan. Therefore, we add the Japanese translation to our search query. The use of translations of English search queries is restricted to exchange rates where the Japanese Yen was involved. The other currencies did not require a different language. If our analysis would be extended to other currencies, in particular in case of smaller countries, the issue of coverage by Google becomes a more significant problem. In the case of our search queries, it was in most cases possible to calculate country specific GTIs by restricting the data to a specific country. However, unreported analysis shows that this would be more challenging for other countries, particularly when Google is not intensely used in the country involved. In the case of the Euro currency, the use of Google Trends data is also a bit more problematic as the Eurozone encloses several different languages.

Tables and Figures

Figure 1: **Retrieving the Google Trends data**

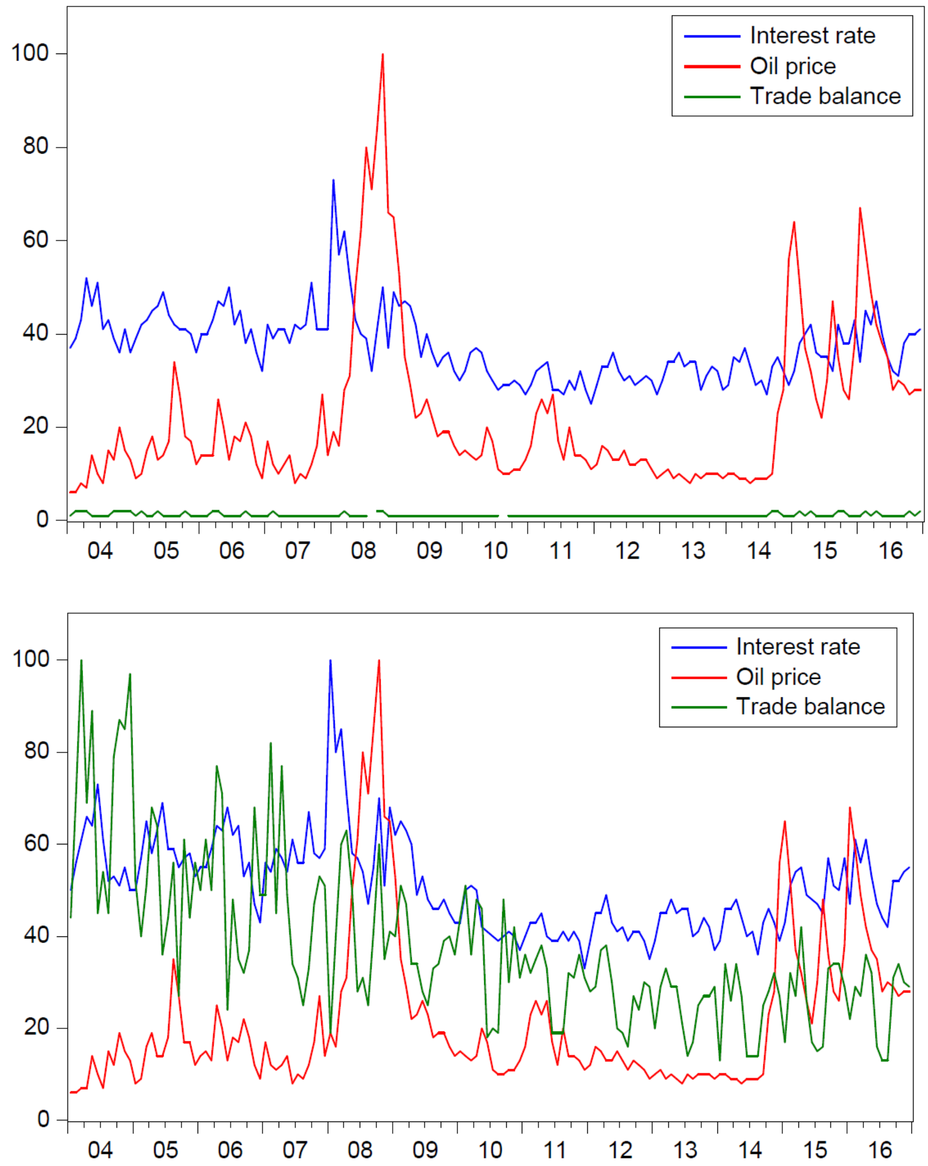

Notes: This figure shows the two different ways of downloading Google Trends data. In the top panel, we use the Google Trends "compare" function. In the lower panel, data has been downloaded separately for each query.

Table 1: **Google Trends Index illustration**

Monetary policy	Germany	France	Netherlands	Spain	Italy	Austria
2014	23.30%	40.70%	36.60%	6.50%	26.80%	19.00%
2015	-0.80%	13.20%	24.50%	-14.60%	15.40%	32.70%

Notes: This table presents the average change in the Google Trends Index for the search query 'monetary policy' per country per year in the week of ECB Governing Council meetings compared to the week prior to it.

Figure 2: **GTI behavior around ECB quantitative easing decisions**

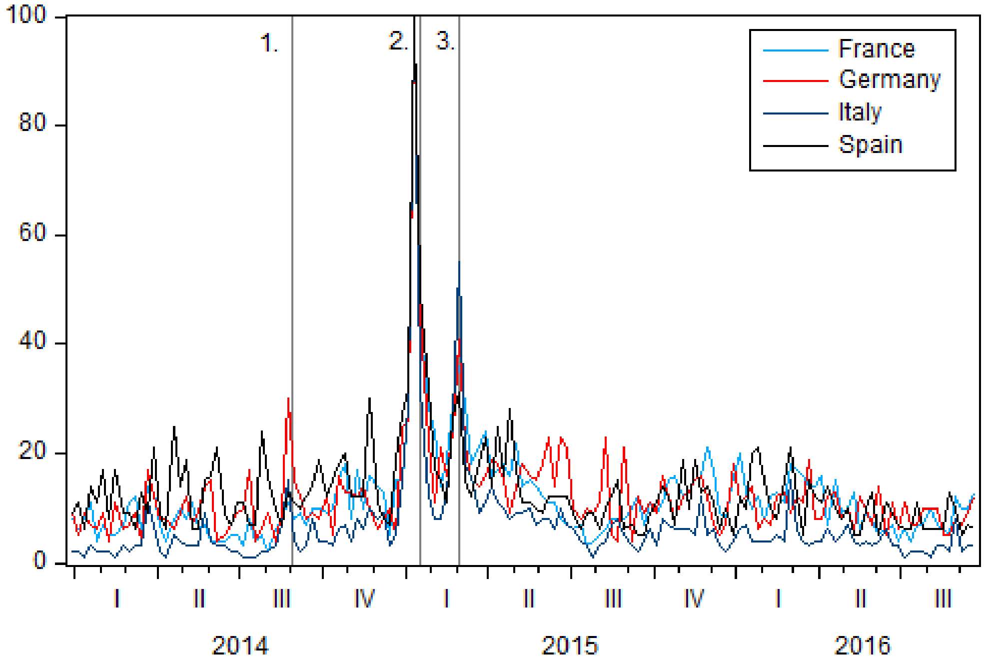

Notes: This figure shows the Google Trends Index for a number of countries on the term 'quantitative easing'. The first event in the figure is the ECB's decision on September 4 2014 to implement QE via purchases of asset backed securities and covered bonds. The second event is the ECB's announcement on January 22 2015 to start a purchase program of sovereign bonds. The third event is the actual start of the purchase program.

Figure 3: **Federal Reserve's quantitative easing and Google Trends**

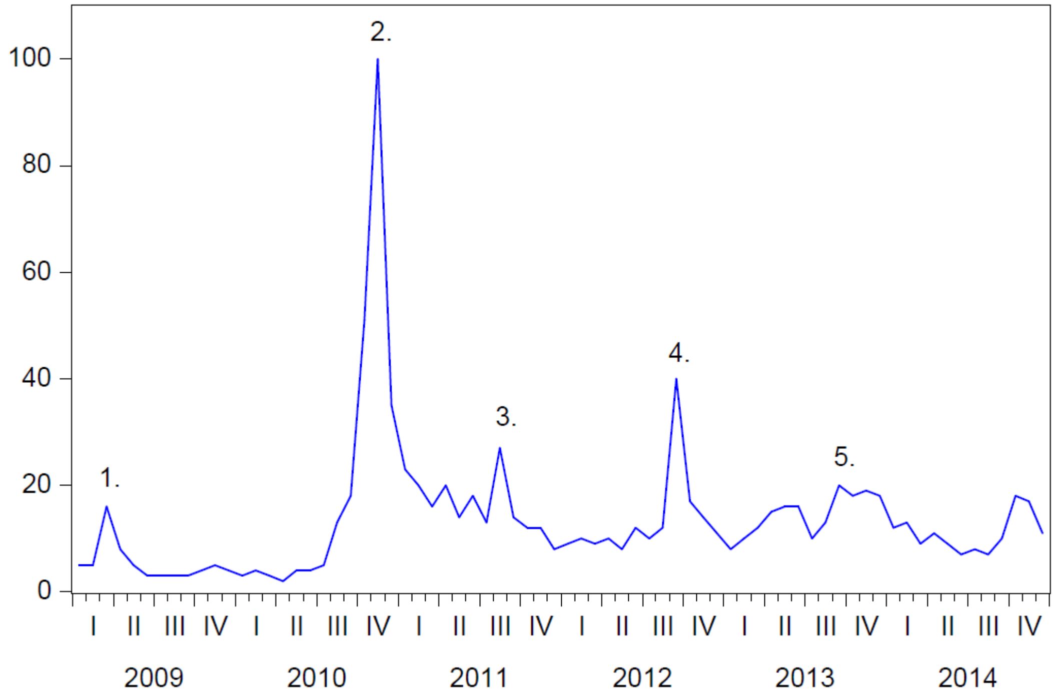

Notes: This figure shows the Google Trends Index on the search query 'quantitative easing'. The events in the graph are the following. 1: In March 2009, the Federal Reserve announces to purchase \$300 billion sovereign and to expand existing purchase programs. 2: In November 2010, the Fed announces to expand the sovereign purchase program. 3: In September 2011, Operation Twist is announced to lower long-term interest rates, 4: In September 2012, a new asset purchase program is announced, involving mortgage backed securities. 5: In September 2013, the Fed decides to continue with asset purchasing programs, whereas financial markets expected tapering.

Table 2: **Search queries for Google Trends data**

Model	**Google search queries**
UIRP	Interest rate, libor rate, interest rate parity, carry trade, central bank rate
PPP	Inflation, rate of inflation, CPI, purchasing power parity, consumer price index, price index
Monetary model	Money supply, money demand, M1, monetary base
Consumption growth	Consumption growth, GDP growth, economic growth, consumption
Trade balance	Trade balance, balance of trade, export, import, international trade
Net foreign assets	Foreign assets, foreign reserves, foreign exchange reserves, net foreign assets, capital flows, net international investment position
Commodities	Commodities, commodity price, commodity index
Oil price	Oil price, West Texas Intermediate, WTI oil, crude oil

Notes: This table lists the search queries that are employed for each fundamental exchange rate model. Google Trends data is collected for each search query and for each geographical area separately. The search queries are listed in English, but are translated if the exchange rate under consideration requires this.

Table 3: **Correlations between Google Trends series for the United States**

UIRP					
	(1)	(2)	(3)	(4)	(5)
Interest rate (1)	1				
Libor rate (2)	0.60	1			
Interest rate parity (3)	0.38	0.25	1		
Federal reserve rate (4)	0.81	0.56	0.30	1	
Carry trade (5)	0.23	0.38	0.12	0.27	1
MM					
	(12)	(13)	(14)	(15)	(16)
Money supply (12)	1				
Money demand (13)	0.79	1			
M1 money (14)	0.71	0.48	1		
Monetary base (15)	0.70	0.44	0.69	1	
Quantitative easing (16)	0.15	0.05	-0.01	0.20	1
TB					
	(21)	(22)	(23)	(24)	(25)
Trade balance (21)	1				
Balance of trade (22)	0.34	1			
International trade (23)	0.81	0.19	1		
Export (24)	0.65	0.06	0.89	1	
Import (25)	0.68	0.00	0.93	0.95	1
Commodities					
	(33)	(34)	(35)		
Commodities (32)	1				
Commodity price (33)	0.69	1			
Commodity index (34)	0.66	0.58	1		

PPP						
	(6)	(7)	(8)	(9)	(10)	(11)
Inflation (6)	1					
Rate of inflation (7)	0.63	1				
Price index (9)	0.60	0.76	1			
CPI (10)	0.61	0.72	0.93	1		
Consumer price index (11)	0.54	0.76	0.98	0.90	1	
Purchasing power parity (12)	0.67	0.63	0.73	0.68	0.72	1
CG						
	(17)	(18)	(19)	(20)		
Consumption (17)	1					
Consumption growth (18)	0.14	1				
GDP growth (19)	0.29	0.50	1			
Economic growth (20)	0.61	0.65	0.69	1		
NFA						
	(26)	(27)	(28)	(29)	(30)	(31)
Foreign assets (26)	1					
Foreign reserves (27)	0.46	1				
Foreign exchange reserves (28)	0.48	0.47	1			
Net foreign assets (29)	0.35	0.31	0.24	1		
NIIP (30)	0.50	0.43	0.44	0.37	1	
Capital flows (31)	0.71	0.48	0.55	0.17	0.43	1
Oil						
	(35)	(36)	(37)	(38)		
Oil price (35)	1					
Crude oil (36)	0.97	1				
West Texas Intermediate (37)	0.34	0.35	1			
Wti oil (38)	0.30	0.31	0.99	1		

Notes: This table presents the correlations of the Google Trends data series for the United States, for the search queries of each fundamental exchange rate model.

Table 4: **Statistical evaluation of forecasting performance with model selection based on the average Google Trends Index**

Currency	**MSPE ratio**	**CW p-value**	**PT p-value**
CAD-GBP	0.953	0.025	0.626
CAD-JPY	0.955	0.009	0.168
CAD-CHF	0.967	0.020	0.039
CAD-AUD	1.001	0.511	0.168
GBP-JPY	0.967	0.042	0.002
GBP-CHF	0.985	0.211	0.375
USD-CAD	0.927	0.012	0.075
USD-GBP	0.934	0.013	0.075
USD-JPY	0.987	0.330	0.500
USD-CHF	1.005	0.567	0.212
USD-AUD	0.911	0.009	0.075
CHF-JPY	1.011	0.772	0.500
AUD-CHF	0.984	0.275	0.500
AUD-GBP	0.975	0.056	0.039
AUD-JPY	0.931	0.009	0.027
Average	0.967		

Notes: This table presents the statistical evaluation of the out-of-sample forecasting performance with model selection based on Google Trends. The column labeled "MSPE ratio" represents the Clark-West adjusted ratio of mean squared forecast errors. The columns labeled "CW p-value" and "PT p-value" show p-values for the Clark-West and Pesaran-Timmermann statistics, respectively. The p-values test the null hypotheses of equal forecast accuracy of our forecasts and the random walk.

Table 5: **Economic evaluation of forecasting performance with model selection based on the average Google Trends Index**

Currency	**Average**	**Stdev**	**Sharpe**	**p-value**
CAD-GBP	2.42%	9.11%	0.266	0.169
CAD-JPY	8.70%	13.10%	0.663	0.009
CAD-CHF	8.55%	11.42%	0.749	0.004
CAD-AUD	2.35%	9.61%	0.244	0.190
GBP-JPY	9.38%	13.31%	0.705	0.006
GBP-CHF	2.65%	10.34%	0.256	0.179
USD-CAD	5.63%	9.82%	0.573	0.020
USD-GBP	6.86%	9.26%	0.740	0.004
USD-JPY	0.59%	9.91%	0.059	0.415
USD-CHF	0.36%	10.90%	0.033	0.453
USD-AUD	9.11%	13.45%	0.678	0.008
CHF-JPY	-1.49%	12,13%	-0.123	0.329
AUD-CHF	1.12%	10.71%	0.104	0.354
AUD-GBP	7.02%	10.59%	0.663	0.009
AUD-JPY	10.41%	15.21%	0.685	0.007
EW-P	4.91%	3.71%	1.323	0.000
VW-P	4.59%	3.61%	1.271	0.000
LS-P	11.21%	17.60%	0.637	0.011

Notes: This table presents the economic evaluation of the out-of-sample forecasting performance with model selection based on Google Trends. It shows the returns of an investment strategy that goes long (short) in the currency which is forecasted to appreciate (depreciate). The rows "EW-P", "VW-P" and "LS-P" show the statistics for the equally weighted, value weighted and long-short portfolios, respectively. The table shows the average annualized return, the annualized standard deviation, the Sharpe ratio and the p-value of testing the significance of the average return.

Table 6: **Average weight attached to individual fundamental model based on relative attention**

Model	**Avg. Weight**
UIRP	9%
PPP	16%
MM	11%
CG	17%
TB	17%
NFA	5%
COMM	16%
OIL	9%
Average minimum weight	4.30%
Average maximum weight	19.60%

Notes: This table presents the average weights attached each of the fundamental models, averaged over the forecasting period (2004M1-2016M12). Also, it highlights the difference between the average maximum and average minimum weights.

Table 7: **Statistical evaluation of forecast combinations with weights based on the relative search attention as measured by the average GTI**

Currency	MSPE ratio	CW p-value	PT p-value
CAD-GBP	0.994	0.327	0.626
CAD-JPY	0.995	0.338	0.212
CAD-CHF	1.007	0.791	0.626
CAD-AUD	0.996	0.335	0.261
GBP-JPY	0.996	0.379	0.131
GBP-CHF	0.986	0.311	0.739
USD-CAD	0.977	0.089	0.003
USD-GBP	0.970	0.007	0.002
USD-JPY	1.008	0.737	0.100
USD-CHF	1.002	0.545	0.374
USD-AUD	0.962	0.024	0.055
CHF-JPY	1.011	0.839	0.374
AUD-CHF	0.988	0.182	0.788
AUD-GBP	0.987	0.202	0.315
AUD-JPY	0.987	0.130	0.100
Average	0.991		

Notes: This table presents the statistical evaluation of the out-of-sample forecasting performance with model selection based on Google Trends. The column labeled "MSPE ratio" represents the Clark-West adjusted ratio of mean squared forecast errors. The columns labeled "CW p-value" and "PT p-value" show p-values for the Clark-West and Pesaran-Timmermann statistics, respectively. The p-values test the null hypotheses of equal forecast accuracy of our forecasts and the random walk.

Table 8: **Economic evaluation of forecast combinations with weights based on the relative search attention as measured by the average GTI**

Currency	**Average**	**Stdev**	**Sharpe**	**p-value**
CAD-GBP	2.17%	9.12%	0.238	0.196
CAD-JPY	1.66%	13.34%	0.124	0.327
CAD-CHF	0.80%	11.68%	0.069	0.402
CAD-AUD	-0.21%	9.63%	-0.021	0.469
GBP-JPY	0.42%	13.58%	0.031	0.455
GBP-CHF	2.66%	10.34%	0.258	0.177
USD-CAD	4.17%	9.88%	0.422	0.065
USD-GBP	8.31%	9.16%	0.906	0.001
USD-JPY	0.63%	9.91%	0.064	0.409
USD-CHF	-1.78%	10.88%	-0.164	0.278
USD-AUD	5.84%	13.60%	0.43	0.062
CHF-JPY	-0.24%	12.14%	-0.02	0.471
AUD-CHF	1.28%	10.71%	0.12	0.333
AUD-GBP	-0.98%	10.78%	-0.091	0.372
AUD-JPY	7.78%	15.34%	0.507	0.035
EW-P	2.17%	4.47%	0.485	0.041
VW-P	2.01%	4.25%	0.472	0.046
LS-P	9.29%	17.30%	0.537	0.059

Notes: This table presents the economic evaluation of the out-of-sample forecasting performance with model selection based on Google Trends. It shows the returns of an investment strategy that goes long (short) in the currency that is forecasted to appreciate (depreciate). The rows “EW-P”, “VW-P” and “LS-P” show the statistics for the equally weighted, value weighted and long-short portfolios, respectively. The table shows the average annualized return, the annualized standard deviation, the Sharpe ratio and the p-value of testing the significance of the average return.

Table 9: **Performance of individual models**

	EW-P	p-value	VW-P	p-value	LS-P	p-value
UIRP	-0.297	0.143	-0.303	0.138	-0.150	0.294
PPP	0.287	0.152	0.211	0.224	-0.038	0.445
MM	-0.208	0.227	-0.116	0.338	0.163	0.279
CG	-0.043	0.438	-0.024	0.465	-0.113	0.342
TB	-0.263	0.173	-0.237	0.197	0.184	0.253
NFA	0.283	0.155	0.218	0.216	0.513	0.033
COMM	0.154	0.289	0.158	0.284	0.108	0.349
OIL	1.180	0.001	1.127	0.000	1.104	0.000
Eq.Weight	0.513	0.033	0.468	0.047	0.694	0.007

Notes: This table shows the economic evaluation of the forecasts from the eight individual fundamental models as well as the equally weighted forecast in terms of the Sharpe ratio for an equally-weighted (EW), volatility-weighted (VW) and a long-short (LS) portfolio. It also shows the p-value of a t-test on the significance of the average return of each portfolio.

Table 10: **Switching behavior**

Currency	Switches	Dur	UIRP	PPP	MM	CG	TB	NFA	COMM	OIL
CAD-GBP	42%	2.36	0%	9%	1%	45%	6%	1%	29%	10%
CAD-JPY	59%	1.73	1%	15%	3%	35%	22%	3%	14%	8%
CAD-CHF	59%	1.73	1%	8%	4%	25%	13%	0%	41%	7%
CAD-AUD	51%	1.95	1%	51%	3%	26%	6%	1%	3%	9%
GBP-JPY	45%	2.20	0%	10%	1%	21%	10%	0%	49%	9%
GBP-CHF	33%	2.98	1%	1%	1%	13%	10%	1%	65%	8%
USD-CAD	46%	2.17	0%	8%	12%	44%	19%	0%	11%	7%
USD-GBP	39%	2.56	0%	7%	5%	50%	10%	0%	21%	6%
USD-JPY	23%	4.11	0%	8%	1%	4%	79%	0%	4%	4%
USD-CHF	55%	1.79	1%	5%	8%	17%	40%	0%	24%	4%
USD-AUD	54%	1.84	1%	33%	6%	24%	30%	0%	1%	6%
CHF-JPY	41%	2.40	1%	4%	0%	5%	28%	1%	54%	7%
AUD-CHF	62%	1.61	1%	15%	1%	24%	29%	0%	21%	8%
AUD-GBP	51%	1.96	0%	21%	1%	34%	8%	0%	26%	10%
AUD-JPY	40%	2.48	1%	37%	3%	15%	32%	0%	4%	8%
Average	46%	2.26	1%	15%	3%	25%	23%	0%	24%	7%

Notes: This table shows the summary statistics concerning the switching behavior when only the model with the highest attention is taken into account. "Dur" represents duration.

Table 11: **Google Trends and the publication of inflation and labor statistics**

	2013		2014		2015	
	CPI	PPI	CPI	PPI	CPI	PPI
% Difference in GTI *inflation*	14.1%	9.0%	21.6%	8.8%	13.6%	7.3%
p-value t-test	0.000	0.004	0.000	0.009	0.001	0.048
		ES		ES		ES
% Difference in GTI *employment*		7.2%		6.6%		5.0%
p-value t-test		0.028		0.002		0.026

Notes: This table presents the percentage difference in the daily US Google Trends Index for the search queries *inflation* and *employment*, respectively, on publication dates compared to non-publication dates. For *inflation*, publications involve the monthly releases of statistics on the Consumer Price Index (CPI) and the Producer Price Index (PPI) by the US Bureau of Labor Statistics. For *employment*, this considers the publication of the monthly Employment Situation (ES) by the same agency. The table also shows the p-value of a t-test to test the statistical significance of the percentage difference.

Figure 4: **Oil price movements and GTI of oil price**

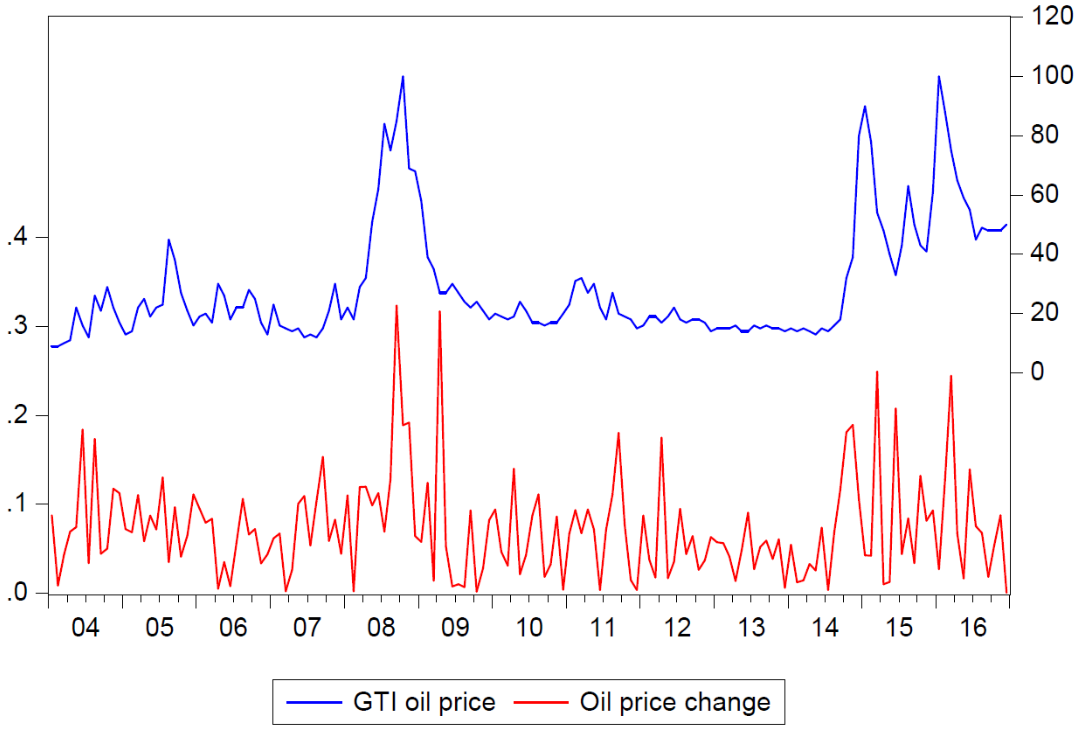

Notes: This figure shows monthly oil price movements (left axis, in percentages), together with the Google Trends Index for the search query *oil price* (right axis).

Table 12: **Robustness checks: statistical evaluation**

	Significant CW-stat	**Significant PT-stat**
Full sample	CAD-GBP, CAD-JPY, CAD-CHF, GBP-JPY, USD-CAD, USD-GBP, USD-AUD, AUD-GBP, AUD-JPY	CAD-CHF, GBP-JPY, USD-CAD USD-GBP, USD-AUD, AUD-GBP, AUD-JPY
After deleting queries	CAD-GBP, CAD-JPY, CAD-CHF, GBP-JPY, USD-CAD, USD-GBP, USD-AUD, AUD-JPY	CAD-JPY, GBP-JPY, USD-CAD, USD-AUD, AUD-JPY
1st half	CAD-GBP, CAD-JPY, CAD-CHF, USD-GBP, USD-AUD, AUD-JPY	CAD-CHF, GBP-JPY, USD-CAD, USD-AUD, AUD-GBP, AUD-JPY
2nd half	CAD-JPY, CAD-CHF, GBP-JPY, USD-CAD, USD-GBP, USD-JPY, AUD-CHF, AUD-GBP, AUD-JPY	GBP-JPY

Notes: This table presents two different robustness checks to the statistical evaluation of the out-of-sample forecasting performance. The table presents the currencies for which the out-of-sample forecasting performance is significantly better than for the random walk. The table shows the impact of deleting the last query for each fundamental model and shows the results of splitting the sample size.

Table 13: **Robustness checks: economic evaluation**

		Return	**St.dev.**	**Sharpe**	**p-value**
EW-P	Full sample	4.91%	3.71%	1.323	0.000
	After deletion of queries	3.84%	3.66%	1.048	0.000
	1st half	6.14%	3.59%	1.712	0.000
	2nd half	3.58%	3.86%	0.925	0.010
	VIX 10-percentile	11.46%	4.81%	2.383	0.000
	Recessions	7.81%	4.79%	1.630	0.000
VW-P	Full sample	4.59%	3.61%	1.271	0.000
	After deletion of queries	3.54%	3.59%	0.985	0.000
	1st half	5.64%	3.51%	1.609	0.000
	2nd half	3.41%	3.75%	0.910	0.012
	VIX 10-percentile	9.64%	4.71%	2.047	0.000
	Recessions	7.31%	4.75%	1.538	0.000
LS-P	Full sample	11.21%	17.60%	0.637	0.011
	After deletion of queries	8.74%	17.87%	0.489	0.040
	1st half	12.78%	19.11%	0.669	0.046
	2nd half	8.83%	16.08%	0.549	0.083
	VIX 10-percentile	34.10%	30.92%	1.103	0.000
	Recessions	26.36%	20.50%	1.286	0.000

Notes: This table presents the economic evaluation of the out-of-sample forecasting performance, where we apply different robustness checks. The table shows the impact of deleting the last query for each fundamental model as presented in table 2. Also, it presents the economic evaluation for both halves of the sample (2004M1-2010M6 and 2010M7-2016M12) and for periods of financial distress and recessions. The rows "EW-P", "VW-P" and "LS-P" show the statistics for the equally weighted, value weighted and long-short portfolios, respectively. The table shows the average annualized return, the annualized standard deviation, the Sharpe ratio and the p-value of testing the significance of the average return.

Table 14: **Returns to Carry Trade Strategy**

Perspective	Investment	Average	Stdev	Sharpe	p-value
American	Long-short	0.021	0.162	-0.132	0.319
	Long	-0.012	0.106	-0.116	0.338
	Short	0.009	0.086	0.103	0.357
Australian	Long-short	-0.067	0.193	-0.349	0.106
	Long	-0.034	0.093	-0.369	0.093
	Short	0.033	0.12	0.274	0.163
British	Long-short	0.054	0.171	0.316	0.129
	Long	0.039	0.089	0.433	0.061
	Short	-0.016	0.109	-0.143	0.304
Canadian	Long-short	-0.024	0.161	-0.148	0.298
	Long	-0.008	0.075	-0.102	0.358
	Short	0.016	0.109	0.147	0.299
Japanese	Long-short	0.041	0.22	0.187	0.252
	Long	0.012	0.137	0.085	0.381
	Short	-0.03	0.097	-0.305	0.137
Swiss	Long-short	-0.04	0.175	-0.228	0.207
	Long	0.016	0.098	-0.16	0.283
	Short	0.024	0.113	0.213	0.223

Notes: This table shows the returns of carry trade strategies that go long the two currencies with the highest interest rate and go short in the two currencies with the lowest. The annualized average return, the annualized return standard deviation, the Sharpe ratio and the p-value of a t-test on the significance of returns are shown.

Figure 5: Time-varying performance of investment strategies

Panel A: Equally weighted portfolio

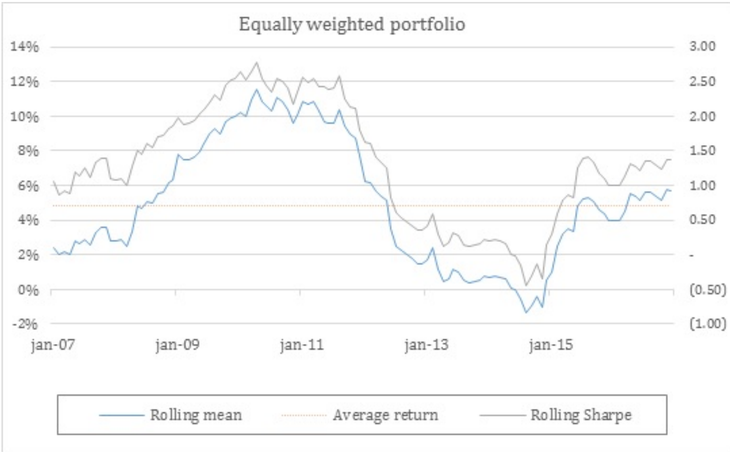

Panel B: Volatility weighted portfolio

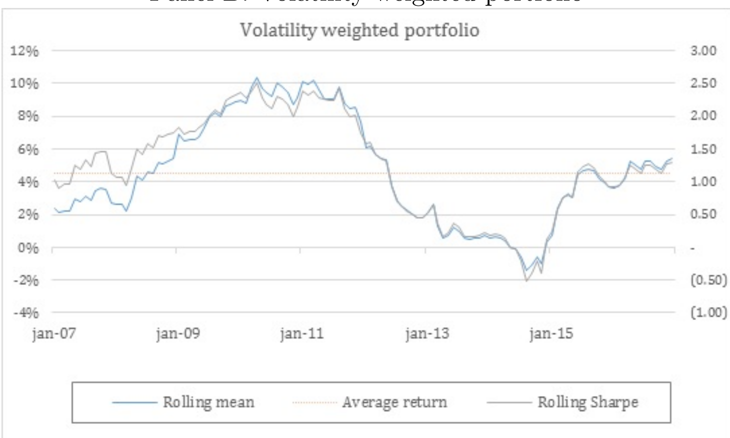

Panel C: Long - short portfolio

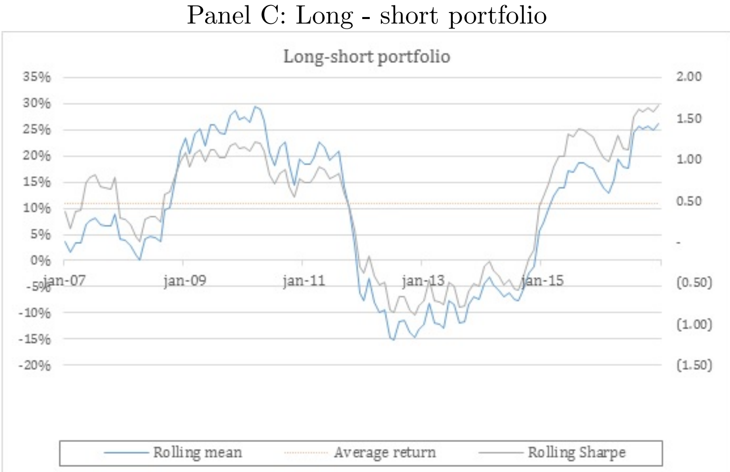

Notes: This figure presents the three-year rolling average performance of the equally weighted (Panel A), volatility weighted (Panel B), and long-short (Panel C) portfolios. The construction of portfolios is described in Section 3.5.