Inattentive Search for Currency Fundamentals*

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

March 2021

Abstract

We investigate whether rational inattention can account for the time-varying link between exchange rates and economic fundamentals. We proxy attention for different fundamentals by the search volume of related search 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 investment returns are systematically delivered by models that select the most salient fundamental. We also confront our findings with specific rational inattention hypotheses. The predicted currency returns are more persistent and less volatile than their actual counterparts, consistent with the underreaction to economic news hypothesis. Investment returns are also substantially higher during periods of elevated uncertainty. Our results suggest that superior performance during recessions is a result of increased attention to the economic fundamentals. These findings provide strong support in favor of the rational inattention theory of exchange rates.

Keywords: *exchange rate forecasting, rational inattention*

JEL Classification: F31, F41, E44

*We thank participants of the Australasian Finance and Banking Conference 2017; the Society for Nonlinear 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

There is now ample empirical evidence that the relationship between exchange rates and underlying fundamentals is unstable over time.[2] Since exchange rates are largely determined by expectations (Engel and West, 2005), the information that market participants pay attention to is a key driver of currencies' dynamics. In this paper, we show that the time-varying nature of the link between exchange rates and fundamentals can be explained by the rational inattention of market participants.

The knowledge of the future value of a currency is essential to many economic decisions, and therefore, motivates the acquisition of information about its underlying fundamentals.[3] Exchange rates are potentially driven by a large set of macroeconomic and financial determinants. Yet, tracking them all to use for inference is costly. In fact, the more attention is paid to relevant economic fundamentals, the more precise is the information about the future exchange rate, but also the higher is the associated cost. A currency market participant might therefore rationally choose to focus her attention on a small set of economic fundamentals or even on one, most salient fundamental. This intuition summarizes, in essence, the rational inattention theory, as introduced by Sims (2003). Because exchange rates have a self-referential structure, the fundamentals that market participants pay attention to become the exchange rate determinants and, as a result, can predict the future exchange rate.

Based on the rational inattention theory, we build an exchange rate forecasting model. Our aim is to capture the macroeconomic fundamentals that attract the attention of currency market participants by using Google search data in a novel way. Our key assumption is that the so-called Google Trends Index (GTI) captures the revealed attention of market participants to different economic fundamentals. The interpretation of the index in this context is rather intuitive: the higher the GTI, the higher the relative search intensity for the fundamental in question, and therefore the higher the attention paid to that fundamental. To illustrate that the GTI in fact captures attention, we study the behaviour of the index in some specific cases where we can isolate its response to shifts in attention from its response to changes in fundamentals. First, we consider the GTI dynamics around the Federal Open Market Committee (FOMC) and the ECB Governing Council meetings on the use of monetary policy instruments. Second, we study the GTI shifts around publications of the Consumer Price Index (CPI) by the U.S. Bureau of Labor Statistics (BLS). In all the event studies, we find that the index increases significantly in response to the event, regardless of whether it is associated with an actual change in the fundamental itself or not.

The GTI captures swings in attention to different economic fundamentals and we expect that these fundamentals affect economic decisions of currency market participants. We translate each of these economic fundamentals into an exchange rate model and we estimate it on a rolling basis. We then use the GTI to select the model that involves the fundamental that economic agents pay the most attention to, and use it for predictions. We start the presentation of our results by documenting that the GTI, used as an attention proxy, can successfully identify the fundamentals that carry out-of-sample predictive power, and deliver statistically and economically significant returns. In a sample ranging from 2004 to 2016, we find that the GTI-based forecasting procedure

[2] See Rossi (2013) for an excellent review.

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

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 as well as 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. In essence, the results demonstrate that the fundamentals that attract the highest attention are powerful in forecasting exchange rates.

We continue our analysis by focusing on the implications of the rational inattention theory for the relationship between exchange rates and economic fundamentals. We first investigate whether the predicted currency returns display dynamics consistent with the hypothesis of under-reaction to economic news. More specifically, rational inattention would imply that the role of shocks to fundamentals is attenuated in the currency forecasts relative to their actual counterparts. Accordingly, we show that the volatility of the forecasts is, on average, only 13% of the volatility of the actual exchange rate returns.

Since the future value of the currency and underlying fundamentals are unobserved when the exchange rate forecast is made, market participants tend to attach high importance to observed information; specifically, the past exchange rate. Accordingly, we expect predicted returns to be more persistent than their actual counterparts. We find that, in as many as 12 out of 15 actual return regressions, the $AR(1)$ coefficient is not significantly different from 0. In contrast, only 2 out of 15 predicted returns regressions display an $AR(1)$ coefficient that is not significantly different from zero. The coefficient is on average equal to 0.29, implying a substantial degree of persistence in predicted returns.

The second implication of our rational inattention framework is that fundamentals that attract high attention, receive magnified weights compared to a situation in which equal attention is paid to all relevant variables in the exchange rate determination process. If market participants' attention is limited, the best forecasts should be delivered by strategies based on only a small set of fundamentals identified by the GTI. We investigate this hypothesis by comparing our baseline results to settings in which more than one factor is considered. First, we study the simplest model with equal and fixed weights on each fundamental. Second, we consider a more flexible specification that allows for time-varying and magnified weights, derived from fundamentals' relative Google search intensity. These alternatives appear to significantly underperform: the best predictions are always delivered by the model with only one fundamental that attracts the highest attention and hence receives an extreme weight of 100%.

Finally, the rational inattention framework implies that the higher is the uncertainty about the exchange rate, the more attention market participants pay to fundamentals and the more elastic is the currency's response to economic shocks. Because increases in uncertainty are usually associated with recessions, we investigate to what extent the dynamics of our attention-based measure and resulting models change during recessions. If market participants are more attentive during high uncertainty periods, the performance of the GTI model should increase during recessions. We formally test this hypothesis by judging performance of the attention-based model during the periods defined as recessions by the OECD and periods in which the implied volatility index VIX was above its 90th percentile. We find that the performance during high uncertainty periods is substantially higher than during the entire sample. Our results suggest that the superior performance of the GTI during recessions is a result of increased attention to the fundamentals, measured as the proportion of periods in which the selected fundamental changes.

In essence, these findings provide strong support in favor of the rational inattention theory of the exchange rates.

Literature on the time varying link between exchange rates and economic fundamentals is abundant. Meese and Rogoff (1983) introduced the exchange disconnect rate puzzle by showing a poor out-of-sample performance of structural exchange rate models. Since then, there has been mixed evidence on the relation between exchange rates and economic fundamentals. 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. Several theoretical studies provide possible explanations, alternative to rational inattention, for the time varying link between currencies and fundamentals. These papers typically 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. Fratzscher et al. (2015) find empirical evidence that supports the scapegoat theory, with resulting models explaining a large fraction of the variation and directional changes in currencies in sample. However, in an out-of-sample exercise the models derived from scapegoat theory typically perform worse than the random walk.

Since our principal goal is to test the rational inattention hypothesis, this paper contributes to the growing body of literature showing that many important phenomena in economics are driven by the fact that agents cannot absorb all available and relevant information, and therefore choose which exact pieces of information to focus on. Research from the psychology literature has provided sufficient evidence to conclude that it is hard to process multiple information sources.[4] The theory of rational inattention, pioneered by the seminal work of Sims (2003), provides a model of how cognitively limited agents might simplify and summarize available information. Mackowiak et al. (2020) provide an excellent review of the existing literature on the topic with the focus on rational inattention in economics and finance. The assumption of inattentive agent has most often been implemented in the asset pricing literature. Hirshleifer and Teoh (2003), Peng and Xiong (2006), and Andrei and Hassler (2015) assume, in a theoretical framework, that investors are unable to pay attention to all the available information. Resulting limited attention can generate important features observed in asset returns that are otherwise difficult to explain with standard rational expectations models.

The subsequent empirical literature on limited attention finds support in favour of the predictions of the theoretical studies. Limited attention has been a common feature of financial markets' participants and can account for a set of unexplained asset pricing phenomena. Da, Engelberg and Gao (2011) show that an increase in the Google 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 investors' demographics (gender, age) and financial

[4]Studies on the topic include Stroop (1935), Rensink, O'Regan and Clark (1997), Pashler (1998), Pashler and Johnston (1998) and Simons and Chabris (1999).

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.

Rational inattention can be also directly tested by using data on expectations of economic agents. Coibion and Gorodnichenko (2012, 2015) document, for instance, that in the survey data expectations deviate systematically from full-information rational expectations and behave in line with the rational inattention hypothesis. We extend the empirical literature on the rational inattention by studying its implications for the nature of the relationship between exchange rates and economic fundamentals.

The remainder of this paper is organized as follows. Section 2 explains why rational inattention is a suitable framework to study the relationship between exchange rates and economic fundamentals. In this section, we also develop a set of rational inattention hypotheses that will be tested later on. Section 3 describes the construction of our measure of attention based on the Google Trends Index. In Section 4, we present the fundamental models and describe how we map them into relevant search queries. In this section, 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 test the forecasting models and rational inattention hypotheses. The results are presented in Section 5, and Section 6 concludes the paper.

2 Exchange Rate Forecasting with Rational Inattention

Rational inattention, proposed in the seminal work by Sims (2003), is motivated by the fact that people cannot pay attention to all the available information simply because it is costly. Instead, they limit their focus to what they consider important information and they optimally use it to form economic decisions. While models of information acquisition are not new, the assumption that agents can choose to get information of any form from an unrestricted menu of signals is novel to the rational inattention (RI) theory.

The RI framework is suitable to study the relationship between exchange rates and fundamentals for several reasons. First, exchange rates are largely determined by expectations of market participants (see e.g. Engel and West, 2005). The information that economic agents pay attention to is therefore a key determinant of currency dynamics. Second, although the internet makes information on a large set of currency determinants easily available, survey evidence shows that agents pay attention only to a limited number of them. In a study based on US currency traders' survey expectations, Cheung and Chinn (2001) document that the macroeconomic variables that are indicated as most important to determine the future exchange rate are few. A natural explanation for a small set of fundamentals is the cost associated with the collection and processing of information on them, a key assumption of rational inattention.

To understand the intuition behind the rational inattention framework, consider a simple model presented in Mackowiak et al. (2020), where agents decide about the precision of gathered information. Consider a manager who sets a price to maximize profits. The price depends on the market conditions which are unknown. The amount of attention the manager will pay to the current market conditions depends on the profits at stake, the degree of market uncertainty, and the cost of uncovering information about economic conditions. In particular, the more attention paid, the more precise is the information about the optimal price the manager receives, but also the higher is the associated cost. In this model, the price becomes endogenous because the choice of the level of attention determines the response of prices to market conditions.

It is easy to see how similar intuition applies to the context of exchange rate forecasting. The knowledge of the future value of a currency is essential to many economic decisions. Consider, for instance, a producer who is importer of foreign components and who needs to decide on the price of his product, conditional on the future price of the currency, or an inflation targeting central bank, who needs to decide on the short term interest rate, taking into account the future expected depreciation of its home currency. The optimal decisions of the importer and the central bank will depend on the future value of the exchange rate, and therefore, provide incentives to acquire information about future behaviour of its drivers.

Exchange rates are potentially driven by a large set of financial and macroeconomic determinants. However, tracking them all to use for inference is costly. Although internet makes all the information on exchange rate fundamentals easily accessible, even simple inquiry on their developments is time-consuming and making inference requires skilled, expensive labour. It might therefore be optimal for currency market participants to only pay attention to the most salient fundamental at a relatively low cost. Because the current exchange rate value depends on its future value expected by market participants, it is endogenous and, importantly, the fundamental paid attention to becomes the exchange rate determinant. Using the Google Trends Index as a measure of attention, we investigate if exchange rates are indeed primarily driven by fundamentals that attract the most attention. We design a set of empirical exercises to test a number of qualitative implications of rational inattention for exchange rate forecasting. These three hypotheses will be at the core of our empirical analysis:

Under-reaction In the environment where the future exchange rate is unknown, the market participants tend to put more weight on the prior knowledge (previous exchange rate return), increasing persistence of the forecasts. As a result, the role of shocks to economic fundamentals is underestimated in the exchange rate forecast and its volatility is expected to be lower that the one in the actual exchange rate returns.

Magnified relative weights Because the exchange rate has a self-referential structure, the fundamental that is paid high attention to is expected to receive a magnified weight relative to a model where equal attention is paid to all the potential exchange rate drivers.

Uncertainty increases responses The higher the uncertainty about the exchange rate, the more attention market participants pay to fundamentals and the more elastic is the currency's response to economic shocks. Higher uncertainty can trigger more attention for two main reasons. First, higher uncertainty weakens the informativeness of fundamentals, which increases market participants' incentives to search for information. Second, increases in uncertainty are usually associated with recessions during which the stakes of economic decisions are higher and lead market participants to pay extra attention.

3 Google Trends as a measure of investor attention

Search intensity on Google, or other internet sources in general, has been used before to measure investors' attention; see e.g. Da, Engelberg, and Gao (2011) and Mondria, Wu, and Zhang (2010). In this paper, we exploit the potential of search data in a novel way. Instead of measuring the aggregate investors' attention to the asset (currency), we use the GTI to capture the macroeconomic fundamentals that attract the attention of currency market participants by using data from Google Trends.[5] If market participants change the weight they attach to each fundamental by shifting their attention over time, this influences their trading behavior and thereby feeds back into the exchange rate.

Importantly, 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 participants 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 Google search data capture the attention of a non-negligible share of currency markets' participants whose decisions arguably impact prices.

In what follows, we describe the GTI and how we make use of it to construct a measure of attention to various exchange rate fundamentals. We then provide evidence that the GTI indeed measures shifts in attention. Specifically, we show that this measure captures attention that goes beyond the exchange rate responses to the shifts in fundamentals.

3.1 Construction of the Google Trends Index

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.[6] Internet users typically use a search engine to search information and Google strongly dominates the web-search market with a market share of 91%. More critically, search volume is a measure of *revealed* attention. If someone searches for a certain variable, this person is undoubtedly paying attention to it.[7]

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

[5] 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 as well (Choi and Varian, 2011). Goel, Hofman, Lahaie, Pennock and Watts (2005) show that search behavior also has predictive content further in the future.

[6] The Google Trends tool is freely accessible via www.google.com/trends.

[7] This is a key advantage compared to more indirect proxies of attention, like trading volume or extreme returns (as employed by Barber and Odean (2008), for instance).

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 option uses the "compare" function, where Google applies an additional re-scaling of the search volume data for the queries relative to each other.

Figure 1 in the appendix demonstrates why the first option 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 short-run 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 provide additional details on how we use the Google Trends tool to extract data series on search intensity, in particular on how we use punctuation to filter search volume.

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. To evaluate the possible noise introduced by this sampling error, we compute the correlation between the GTI data series downloaded twice for five different search 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). On average, the impact of this sampling error is small and, if present, would bias against finding significant results. Pairwise correlations between downloaded time series in this preliminary exercise vary between 0.985 and 0.996. This confirms the minor impact of Google drawing only a random subset of the historical search data.

3.2 Does the Google Trends Index measure attention to fundamentals?

To illustrate to what extent the GTI captures the attention of market participants, we study the behaviour of the index in some specific examples where we isolate the response to shifts in the fundamentals itself from attention shifts. To test our rational inattention framework, it is critical to show that the GTI measures attention and not simply a change in the underlying fundamentals, because the latter would be in line with a standard rational expectations theory: when an underlying fundamental changes, the exchange rate responds and the GTI could be just capturing movements in the fundamentals.

First, we investigate the GTI fluctuations around meetings of the Federal Open Market Committee (FOMC) of the Federal Reserve Board. Over the period 2004-2016, the FOMC held 104 meetings to decide about the target federal funds rate and other monetary policy instruments. In 26 of them, the FOMC decided to change the target rate, while in the remaining 77 cases no change has been implemented. To illustrate that the GTI does not solely react to a change in the fundamental itself (the target federal funds rate in this case), but captures attention of market participants, we compute the average increase in the weekly GTI for the search query "federal funds rate" during the weeks when the FOMC came together, relative to the prior week. On average, the GTI for federal funds rate increased significantly by 71.8% during weeks with an FOMC meeting on the agenda. Moreover, there is no significant difference between meetings where the FOMC decided to change the federal funds rate and meetings without such a decision. In fact, the GTI increased on average by 76.2% in the 26 cases with a change in the target rate, whereas it increased by 71.3% in the 77 cases with no change.

We also study the GTI behavior around similar meetings of the ECB Governing Council between 2011 and 2016, when the Council held 64 meetings to decide about monetary policy instruments. Table 1 in the appendix shows the average increase in the GTI for the search query "monetary policy" for the five largest economies of the Eurozone (Germany, France, Italy, Spain and the Netherlands), relative to the prior week.[8] On average, the GTI for "monetary policy" increases by 26.4% in Governing Council weeks. Surprisingly, the attention of market participants increased more in weeks without an actual change in monetary policy (on average by 28.0%) than in weeks with a change in the policy rate or other monetary policy instruments, such as purchasing programs or refinancing operations (on average by 20.4%), although both figures do not significantly differ from the average of the entire sample of meetings.

As a third example, we investigate the GTI behavior around publications of the Consumer Price Index (CPI) by the U.S. Bureau of Labor Statistics (BLS). Over the period 2004-2016, 155 publications of the CPI have been issued by the BLS. The weekly GTI for the query "inflation" in the United States is on average 5.1% higher in weeks with a publication compared to weeks without CPI publication. This is independent of the actual movement in inflation. In fact, the correlation between the actual (absolute) change in the CPI and the change in the GTI for query "inflation" is not significantly different from zero. Additionally, the average increase in the GTI is the same for CPI releases with (absolute) inflation above average (measured over the period 2004-2016) as for CPI releases with inflation below average.

These three examples provide evidence that GTI effectively captures attention and goes beyond swings in the fundamentals themselves.

4 Model selection with Google trends

The GTI captures swings in attention to different economic fundamentals and we expect that these fundamentals feed back into the currency movements. We map different economic fundamentals to exchange rate determination models, and we estimate those models in-sample. We then use the GTI to select the model that the economic agents pay the most attention to and use it for forecasting out-of-sample. In what follows, we explain in detail the fundamentals and how the

[8] We present these figures for these five countries separately, because Google Trends can only show search data for either the whole world or for individual countries. We are therefore not able to aggregate search data over the Eurozone.

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.

4.1 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).[9] All models relate a change in the natural logarithm of the exchange rate (Δs_t) to a certain fundamental. 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-month and 1-year interest rates.[10]

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 (simple average), 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 represents 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

[9] 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 models that can be easily translated into Google search queries.

[10] We estimate the model with the 3-month interest rate differential as well as with the 1-year interest rate differential. The forecast of this model is average of the two. The results are qualitatively insensitive to these choices.

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:

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

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 models. 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 exchange rates and therefore we look at changes of its prices separately in the following specification:[11]

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

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

4.2 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 attention, we select several search queries and we report them in Table 2. To select multiple search terms, we make use of Google Correlate. 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.

For the search terms shown in Table 2, we always restrict the geographical area of the Google Trends tool to the countries involved in the particular exchange rate. For instance, when we search

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

for fundamentals for the USD-JPY, we restrict the regional area of the GTI to the United States and Japan, respectively. 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.[12] 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 economic fundamentals for the United States. 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 noise that adheres to each individual search query.

4.3 Forecasting procedure

Because it is relatively easy to fit the exchange rate model in-sample, in particular with a flexible specification as ours, we bring the attention-based model of the exchange rate to a more challenging, out-of-sample test. [13] Based on the GTI, we build a forecasting procedure to predict monthly exchange rate movements between January 2004 and December 2016.[14] We perform the analysis for the fifteen currency pair combinations between the AUD, CAD, CHF, GBP, JPY and USD.[15] 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 our study. 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 country involved. We choose the monthly frequency to match the frequency of the economic 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 highest relative attention
Third, at each point in time, we select the fundamental model with the highest average GTI for its search queries, averaged over both countries involved. The model with the highest GTI for its search queries is the model containing the fundamental with the highest relative attention in the countries involved.[16]

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

[13] See Rossi (2013) for description of in-sample versus out-of-sample fit of exchange rate models.

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

[15] By including all the currency pair combinations, our results are not dependent on the choice of the base currency.

[16] More precisely, the selected model might not be the model with the highest average GTI for *both* countries. Consider the following example for illustration: the monetary model has an average GTI for its search queries of 70

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 month in the sample period. 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 model selection set out before. The second estimation period runs from 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 out-of-sample period of January 2004 - December 2016 contains 156 monthly forecasts for each of the 15 currency pairs.

4.4 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.

4.5 Statistical and economic performance measures

Once we have the predicted exchange rate series, we implement a set of forecast evaluation measures to judge the ability of the GTI to trace the fundamentals that economic agents 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:

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

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:

for the USA and 80 for the UK and the PPP-model has an average GTI of 50 for the USA and 90 for the UK. In that case, the averages for the USD/GBP exchange rate are 75 and 70, respectively. Hence, in this case we prefer the monetary model above the PPP-model, while the PPP-model would be preferred for the UK on an individual basis.

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

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.[17]

In addition to assessing model performance for each currency pair 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:

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

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 (or largest depreciation). Practically, these are investable strategies through the use of currency futures.

5 Results

Despite the abundant empirical evidence that the relationship between exchange rates and underlying fundamentals is time-varying, it is not trivial to detect the timing of the fundamentals that drive a particular currency. We start the presentation of our results by documenting that the GTI, used as an attention proxy, can identify the fundamentals that carry predictive power

[17]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.

and deliver statistically and economically significant returns. We then confront our findings with specific implications of the rational inattention theory.

5.1 Statistical and economic performance of attention-based currency forecasts

Table 4 presents the statistical evaluation 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. The p-values of the CW-statistic indicate that the MSPE reduction is significant for 9 out of 15 currencies at the 10% significance level. 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 6 out of 15 currencies, indicating outperformance in predicting the sign of currency returns.

Table 5 shows the economic evaluation of the GTI-based forecasts, which appear to generate successful investment strategies. The table shows the average annualized return, the annualized standard deviation, the annual Sharpe ratio, and the p-value of testing the significance of the average return. The 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, both at the 10% and the 5% significance level. Moreover, the three currency portfolios (equally-weighted, volatility-weighted, and long-short) all have a significant, positive mean return. Given the implemented strategies, the magnitude of the portfolio returns is considerable. The equally-weighted portfolio for instance delivers an annualized investment return of 4.9% with an annualized volatility of 3.7% and a Sharpe ratio of 1.3. 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%). For comparison, the MSCI World yielded a Sharpe ratio of 0.35 over the same sample period. In essence, the results demonstrate that the fundamentals that attract the highest attention are powerful in forecasting exchange rates and delivering substantial and significant returns.

5.2 Underreaction of exchange rate forecasts

As an implication of rational inattention, we test for underreaction of our predictions. When forming an exchange rate forecast, the future value of the currency and underlying fundamentals are unobserved and, therefore, the market participants tend to attach high importance to the observed information, specifically, the past exchange rate. We expect that, as a result, the predicted returns are more persistent than their actual counterparts. Table 6 strongly supports this hypothesis. The first four columns of the table demonstrate that the predicted currency returns are much more persistent than the actual ones. We estimate a simple $AR(1)$-model for both monthly actual exchange rate returns, as well for predicted returns. In as many as 12 out of 15 actual return regressions, the $AR(1)$-coefficient is not significantly different from 0. In contrast, 13 out of 15 predicted returns regressions display an $AR(1)$-coefficient significantly higher than 0. On average, the $AR(1)$ coefficient is 0.29, implying a substantial degree of persistence in the forecasted returns.

Because under-reaction of market participants implies that the role of shocks to economic fundamentals is attenuated in exchange rate forecasts, we also compare the volatilities of actual and predicted returns. The last column of Table 6 displays the ratio of predicted returns' volatility relative to the actual ones. The table documents that the volatility of the forecasts is only a fraction of the volatility of the actual exchange rate returns, with an average ratio of 13.4%.

5.3 Magnified relative weights

Another implication of the rational inattention framework is that fundamentals that are paid high attention to, receive magnified weights compared to a situation in which equal attention is paid to all variables in the exchange rate determination process. Our benchmark model represents the extreme case with the weight on the GTI selected fundamental being 100%. Any alternative specification including additional variables, will reduce this weight.

We investigate magnified weights hypothesis by comparing our baseline results to a setting in which more than one factor is considered.[18] First, we study the simplest naive model with equal and fixed weights on each of the fundamentals: $w_i^{ew} = \frac{1}{N}$, with N being the number of fundamentals. Since our pool of fundamental models includes eight different models, N=8. Second, we consider a more flexible specification that allows for magnified weights, derived from the fundamentals' Google search intensities. In this case, we do not select the single model with the highest GTI, but instead each model receives a non zero weight based on the relative search volume:

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

where i denotes the fundamental model and $w_{i,t}$ is the weight on the fundamental i at time t and is defined as follows:

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

where $GTI_{i,t}^{av}$ is the average GTI computed over the different search queries for fundamental i at time t.

Table 7 documents that a model with fixed and constant weights statistically outperforms the random walk in terms of the point forecast only in case of the AUD-JPY, in contrast to 9 currency pairs when predictions were based on a GTI selected fundamental that receives the full weight. Yet, this simple specification outperforms the random walk for as many as five currencies when predicting the direction of change, at a 10% significance level. Economic profits derived from the constant weights predictions, described in Table 8 are, however, substantially lower than when the weights on fundamentals are amplified. Specifically, the return on investment is significantly different from zero only for 4 individual currencies. Finally, the equally-weighted and volatility-weighted investment portfolios reported in the lower panel of Table 8 deliver returns half the size of those when only the fundamental with the highest attention is selected. Surprisingly, long-short portfolio return (last column of Table 8) is higher under equal than under magnified weights, but the difference between the two is not statistically significant.

We now consider a more flexible specification that allows for magnified weights, with weights derived from the relative search volume in Equation (16). Table 11 presents the average GTI-based weights attached to each of the fundamental models. The table documents a large difference between the mean maximum and the mean minimum weights averaged over the currencies involved. The average maximum weight (19.6%) is about five times higher than the average minimum weight (4.3%), so the relative GTI undoubtedly generates substantial differences in weights attached to

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

fundamentals. Yet, results in Table 9 demonstrate that allowing the weights to be heterogeneous and to vary over time proportionally to the relative GTIs does not improve the performance of the models' combinations forecast. The best predictions are still delivered by the model with only one fundamental that attracts the highest attention. Specifically, the MSPE of the fundamentals' combination forecast relative to the random walk is reduced by a mere 0.89% on average, and the decrease is only statistically significant for 3 currencies. Similarly, the PT-statistic reports that the extension of the model halved the number of currencies with the correct prediction of direction of change, relative to the baseline exercise where only one fundamental is selected. Although the economic evaluation of extended models presented in Table 10 documents positive and significant annualized returns for 3 currencies and all 3 investment portfolios, the absolute size of returns is substantially lower compared to the baseline specification. These results are consistent with Kouwenberg et al. (2017), who also find that a well-chosen single fundamental model outperforms a weighted average of multiple models in out-of-sample exchange rate forecasting.

5.4 Uncertainty increases responses

According to rational inattention theory, the higher the uncertainty about the exchange rate, the more attention market participants pay to fundamentals and the more elastic is the currency's response to economic shocks. Because increases in uncertainty are usually associated with recessions, we investigate to what extent the dynamics of attention-based measure and resulting models change during recessions. As an alternative measure of uncertainty, we take the implied volatility index VIX.

If market participants are more attentive during high uncertainty periods, the performance of the GTI model should increase during the recessions. Figure 2 plots 3-year rolling returns and Sharpe ratios for all three investment portfolios. Remarkably, the best results are delivered precisely during the Great Recession. We formally test this observation by judging performance of the attention-based model during the periods defined as recessions by the OECD and periods when the VIX is above the 90th percentile. Because this exercise results in a substantial reduction of the number of available observations per currency, we focus our attention on the portfolios' evaluation. Table 12 includes the results for recessions separately (rows entitled 'Recessions'), and also displays the findings for periods of financial stress (rows entitled 'VIX 90th percentile'). For comparison, we also present results from our baseline full-sample exercise. Table 12 documents large differences between portfolio returns and Sharpe ratios obtained during low and high uncertainty periods. Returns delivered during high uncertainty periods are substantially higher than those generated during the entire sample, in particular when we use the VIX definition of uncertainty. During these periods of financial distress, the GTI forecasting procedure delivers returns twice as high as during the entire sample period. Despite higher return volatility associated with economic uncertainty, the Sharpe ratios are still significantly higher in turmoil times.

In Table 13 and Table 14, we show that the superior performance of the GTI forecasting procedure during periods of elevated uncertainty is associated with increased attention to the fundamentals, measured as the percentage share of cases when the selected fundamental changes. Table 13 shows the percentage of cases where the chosen fundamental changes compared to the preceding month and the average duration (months) that the chosen model remains in place. On average, the chosen fundamental changes in 46% of the cases and the average duration is 2.26

months. Table 14 shows the model switching results for both the VIX 90th and 95th percentile as a cutoff point between high and low uncertainty regimes. Irrespective of the uncertainty definition, Table 14 documents significantly higher attention when uncertainty is elevated. For example, in case of the CAD-GBP, we report a model shift in 71% of the cases when the VIX is above its 95th percentile. When the VIX is below this percentile, the selected fundamental changes only in 41% of the cases. On average, the fundamental paid the most attention to changes 16% more often when VIX is above its 95th percentile and 12% more often when it is above its 90th percentile.

5.5 Alternative hypotheses

Scapegoat theory

Bacchetta and van Wincoop (2004, 2013) propose a scapegoat theory to explain the instability of the relationship between exchange rates and fundamentals. The scapegoat effects may arise when exchange rates move strongly in response to unobservables, but market participants blame factors that they can actually observe, and specifically those fundamentals that are out of their longer term equilibrium values and move consistently with the observed exchange rates. Fratzscher et al. (2015) provide a strong empirical evidence in favor of the scapegoat hypothesis. To investigate the presence of scapegoats, they estimate a model that conditions only on macroeconomic variables and test it against a model that includes (surveyed) scapegoats, in addition to the same macroeconomic variables. They show that the specifications with scapegoats always outperform the ones with actual macroeconomic variables. This is the case in- and out-of-sample. They also show that even the models that condition on macroeconomic variables, scapegoats and order flows do poorly out-of-sample. In contrast, we show that the actual observed fundamentals, correctly identified by our attention measure, beat the random walk out-of-sample. Because the out-of-sample exercise is more challenging, these results provide strong support in favor of the rational inattention theory rather than the scapegoat hypothesis.

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, although its performance varies over time (e.g. Hattori and Sin, 2009 and Burnside, Eichenbaum, Kleshchelski and Rebelo, 2011). Since the attention-based investment strategies perform significantly better during the crisis, it could be that they capture time-varying carry trade profits.[19] To test this hypothesis, we compare our results with carry trades taking 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.

[19] The empirical evidence rather suggests that the carry trade performed worse during the Great Recession. Still, we want to rule out the possibility that we capture it instead of attention shifts.

Table 15 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 attention-based fundamentals delivers very different results from carry strategies.[20]

5.6 Robustness tests

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 16 and 17 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 our 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 17 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.

Sample period

There is ample empirical evidence that performance of exchange rate models varies over time. To make sure that our results are not specific to to the sample period selection, Figure 2 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.

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 16 and 17 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.

Individual models

As a final robustness check, and for the sake of completeness, we describe the forecasting power of the individual fundamental models. These are the one-period ahead forecasts for each of the eight fundamental models, and therefore insensitive to the attention measure. We create portfolio over currencies; i.e., we take a long (short) position whenever the fundamental model predicts

[20] 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 trades during the financial crisis.

an appreciation (depreciation) of the currency. The portfolios are again equally weighted, value weighted, and long-short. The results in Table 18 reveal that the forecasting performance of individual models is typically quite poor, consistent with the existing literature. Only for the oil price we find a significantly positive return of 1.1% annually, for all three portfolios. This is substantially lower than the return of the GTI-based portfolio.

6 Conclusion

In this paper, we test whether the rational inattention theory of exchange rates provides a possible explanation for the time-varying link between exchange rates and fundamentals. Rational inattention postulates that economic agents cannot absorb all available information, but instead make a cost-benefit analysis in choosing which exact pieces of information to focus on. We proxy the attention of economic agents for economic fundamentals by the Google Trends Index. In a sample of macro-economic 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. The best forecasts and the highest returns are delivered by strategies that select the fundamental that is paid the highest attention to.

Our findings are in line with the implications of the rational inattention framework. The predicted currency returns are more persistent and less volatile than their actual counterparts, consistent with the under-reaction to economic news hypothesis. The forecasting procedure based on attention also results in substantially higher investment returns during periods of elevated uncertainty. Our results suggest that this superior performance is a result of increased attention to the fundamentals.

References

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

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

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

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

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

[6] Barber, B.M., and T. Odean (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.

[7] Ben-Rephael, A., 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, 3009–3047.

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

[9] Chen, Y., Rogoff, K., and B. Rossi (2010), 'Can Exchange Rates Forecast Commodity Prices?', *The Quarterly Journal of Economics*, 125, 1145–119.

[10] Cheung, Y., Chinn, M. D., and A.G. Pascual (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] Coibion, O., and Y. Gorodnichenko (2012), 'What can survey forecasts tell us about information rigidities?', *Journal of Political Economy*, 120(1), 116-159.

[13] Coibion, O., and Y. Gorodnichenko (2015), 'Information Rigidity and the Expectations Formation Process: A Simple Framework and New Facts', *American Economic Review*, 105(8), 2644-2678.

[14] Clark, T., and K.D. West (2007), 'Approximately normal tests for equal predictive accuracy in nested models', *Journal of Econometrics*, 138, 291-311.

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

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

[17] Engel, C., and K. D., West, (2005) ‘Exchange Rate and Fundamentals,’ *Journal of Political Economy* 113, 485-517.

[18] Fratzscher, M., Rime, D., Sarno, L., and G. Zinna (2015), ‘The Scapegoat Theory of Exchange Rates: The First Tests’, *Journal of Monetary Economics*, 70, 1-21.

[19] 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, 17486–17490.

[20] Hattori, M., and H. S. Shin (2009), ‘Yen Carry Trade and the Subprime Crisis’, *IMF Staff Papers*, 56, 384-409.

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

[22] Kahneman, D. (1973), “Attention and Effort”, Prentice-Hall, Englewood Cliffs, NJ.

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

[24] Mackowiak, B., Matejka, F., and M. Wiederholt (2020), ’Rational Inattention: A Review’, CEPR working paper DP15408.

[25] 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.

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

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

[28] Pashler, H.E. (1998), ’The Psychology of Inattention’. MIT Press, Cambridge.

[29] Pashler, H. and J.C. Johnston (1998), ’Attentional Limitations in Dual-Task Performance’, in H. Pashler, ed., Attention (1998), 155-189.

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

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

[32] Rensink, R.A., J.K. O’Regan, 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.

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

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

[35] Schinasi G.J., and P.A.V.B. Swamy (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.

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

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

[38] Sims, C. A. (2003), 'Implications of rational inattention', *Journal of Monetary Economics*, 50(3).

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

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

[41] 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 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. Extending our analysis to other currencies is problematic, in particular in case of smaller countries, due to limited Google coverage. 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 different languages.

Tables and Figures

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

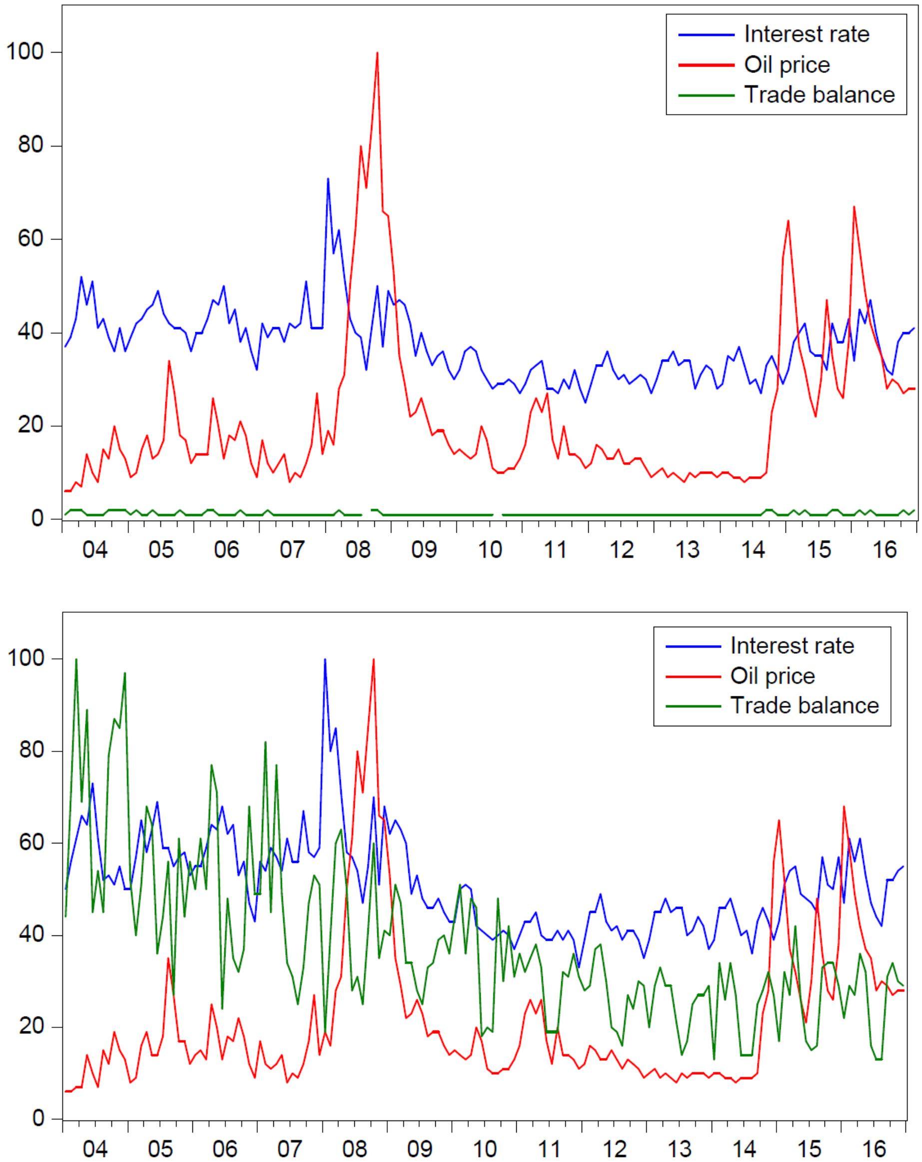

Notes: This figure illustrates the two different options to download Google Trends data for the search queries *interest rate*, *oil price* and *trade balance*. In the top panel, we use the Google Trends "compare" function. In the lower panel, data has been downloaded separately for each query.

Table 1: **Google Trends and ECB monetary policy**

	Germany	France	Italy	Spain	Netherlands	Average
Average change in GC weeks	24.2%	16.1%	24.0%	25.9%	42.0%	26.4%
Average change in ‘decision weeks’	16.9%	23.3%	19.8%	12.3%	29.8%	20.4%
Average change in ‘non-decision weeks’	26.1%	14.3%	25.0%	29.3%	45.1%	28.0%

Notes: This table shows the change in Google Trends Index for the search query *monetary policy* for the five largest European economies and for average (last column), where all five countries are equally weighted. The first row displays the average change in the weeks with an ECB Governing Council meeting. The middle row shows the same numbers, but only for the weeks in which the ECB made a decision to change policy rates or implement or change other instruments. The last row shows the average GTI change for weeks with meetings of the ECB Governing Council without a change in monetary policy. The numbers are calculated over the years 2011-2016, with 64 meetings in total.

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 (see also Appendix A).

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 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 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: **Underreaction of exchange rate forecasts**

Currency	Forecasts AR(1)	p-value	Actual AR(1)	p-value	σ_f/σ_a
CAD-GBP	0.254	0.001	-0.037	0.642	11.2%
CAD-JPY	0.358	0.000	0.042	0.600	14.0%
CAD-CHF	0.001	0.992	-0.211	0.009	7.8%
CAD-AUD	0.161	0.045	-0.051	0.525	10.7%
GBP-JPY	0.195	0.015	0.110	0.173	10.9%
GBP-CHF	0.349	0.000	-0.106	0.188	13.1%
USD-CAD	0.331	0.000	-0.077	0.338	16.5%
USD-GBP	0.367	0.000	0.058	0.470	19.0%
USD-JPY	0.772	0.000	0.056	0.492	19.2%
USD-CHF	0.333	0.000	-0.166	0.039	14.2%
USD-AUD	0.302	0.002	0.022	0.785	18.3%
CHF-JPY	0.242	0.002	-0.127	0.113	10.3%
AUD-CHF	0.141	0.077	0.058	0.472	11.4%
AUD-GBP	0.265	0.001	0.000	0.999	12.4%
AUD-JPY	0.332	0.000	0.086	0.287	12.2%

Notes: This table presents the results of AR(1) regressions of exchange rate returns. All the specifications include an intercept. The left panel, entitled "Forecasts" displays results for GTI-based predicted returns. The right panel, entitled "Actual" corresponds to the regressions with actual exchange rate returns. The last column displays the volatility of the forecasts as a percentage of the volatility of the actual exchange rate returns (σ_f/σ_a).

Table 7: **Statistical evaluation of forecast combinations with fixed and equal weights**

Currency	**MSPE ratio**	**CW p-value**	**PT p-value**
CAD-GBP	0.995	0.347	0.564
CAD-JPY	0.993	0.276	0.039
CAD-CHF	1.008	0.779	0.500
CAD-AUD	0.996	0.370	0.315
GBP-JPY	0.993	0.310	0.212
GBP-CHF	0.985	0.307	0.685
USD-CAD	0.973	0.051	0.005
USD-GBP	0.971	0.004	0.019
USD-JPY	1.010	0.803	0.100
USD-CHF	1.000	0.501	0.261
USD-AUD	0.958	0.020	0.075
CHF-JPY	1.011	0.831	0.436
AUD-CHF	0.982	0.107	0.374
AUD-GBP	0.985	0.200	0.212
AUD-JPY	0.973	0.026	0.075
Average	0.991		

Notes: This table presents the statistical evaluation of the out-of-sample forecasting performance of forecast combinations with fixed and constant weights. 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 hypothesis of equal forecast accuracy of our forecasts and the random walk.

Table 8: **Economic evaluation of forecast combinations with fixed and equal weights**

Currency	**Average**	**Stdev**	**Sharpe**	**p-value**
CAD-GBP	1.06%	9.14%	0.116	0.339
CAD-JPY	3.65%	13.31%	0.274	0.162
CAD-CHF	1.50%	11.67%	0.128	0.322
CAD-AUD	1.98%	9.62%	0.206	0.229
GBP-JPY	-2.28%	13.60%	-0.168	0.272
GBP-CHF	3.27%	10.32%	0.317	0.128
USD-CAD	4.32%	9.88%	0.437	0.059
USD-GBP	5.83%	9.32%	0.625	0.013
USD-JPY	0.25%	9.91%	0.025	0.463
USD-CHF	-1.03%	10.89%	-0.095	0.366
USD-AUD	5.96%	13.60%	0.438	0.058
CHF-JPY	0.75%	12.14%	0.062	0.412
AUD-CHF	1.64%	10.71%	0.153	0.290
AUD-GBP	-0.96%	10.78%	-0.098	0.374
AUD-JPY	10.27%	15.21%	0.675	0.008
EW-P	2.41%	4.70%	0.513	0.033
VW-P	2.09%	4.47%	0.468	0.047
LS-P	13.36%	19.26%	0.694	0.007

Notes: This table presents the economic evaluation of the out-of-sample forecasting performance of forecast combinations with fixed, equal weights. 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 9: **Statistical evaluation of forecast combinations with weights based on the relative search attention**

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 forecast weights 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 10: **Economic evaluation of forecast combinations with weights based on the relative search attention**

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 forecast weights 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 11: **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 12: **High Uncertainty and the GTI**

		Return	**St.dev.**	**Sharpe**	**p-value**
EW-P	Full sample	4.91%	3.71%	1.323	0.000
	VIX 90th 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
	VIX 90th 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
	VIX 90th 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 on different samples. The rows "EW-P", "VW-P" and "LS-P" show the statistics for the equally weighted, volatility-weighted and long-short portfolios, respectively. For each portfolio type, the first rows show the results for full sample, the second, "VIX 90th percentile", for periods of financial distress defined as periods with VIX above 90th percentile, and the third rows, "Recessions", for recessions according to the OECD definition. 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 13: **Switching behavior**

Currency	Switches	Duration
CAD-GBP	42%	2.36
CAD-JPY	59%	1.73
CAD-CHF	59%	1.73
CAD-AUD	51%	1.95
GBP-JPY	45%	2.20
GBP-CHF	33%	2.98
USD-CAD	46%	2.17
USD-GBP	39%	2.56
USD-JPY	23%	4.11
USD-CHF	55%	1.79
USD-AUD	54%	1.84
CHF-JPY	41%	2.40
AUD-CHF	62%	1.61
AUD-GBP	51%	1.96
AUD-JPY	40%	2.48
Average	46%	2.26

Notes: This table shows the summary statistics concerning the switching behavior when only the fundamental with the highest attention is taken into account. It shows the percentage of cases when the selected fundamental changes and the average duration that the selected model remains in place.

Table 14: **Model switching during uncertainty**

Currency	VIX percentile >95th	<95th	difference	VIX percentile >90th	<90th	difference
CAD-GBP	71%	41%	31%***	57%	45%	20%***
CAD-JPY	71%	57%	15%**	60%	57%	3%
CAD-CHF	71%	57%	15%**	73%	56%	18%***
CAD-AUD	86%	49%	36%***	67%	49%	17%**
GBP-JPY	57%	45%	13%*	60%	44%	16%**
GBP-CHF	71%	31%	40%***	53%	31%	23%***
USD-CAD	29%	47%	-18%	40%	46%	-6%
USD-GBP	43%	39%	4%	47%	38%	9%
USD-JPY	29%	23%	6%	33%	22%	11%*
USD-CHF	57%	55%	2%	53%	56%	-2%
USD-AUD	86%	53%	33%***	80%	51%	29%***
CHF-JPY	57%	55%	17%**	60%	39%	21%***
AUD-CHF	86%	61%	25%***	80%	60%	20%***
AUD-GBP	57%	51%	6%	53%	51%	3%
AUD-JPY	57%	39%	18%**	40%	40%	0%
Average	62%	46%	16%**	57%	45%	12%*

Notes: This table shows the percentages of cases in which the chosen fundamental model changes for different VIX percentiles. For instance: when the VIX is above its 95th percentile, the chosen fundamental model changes in 71% of the cases for the CAD-GBP, whereas it only changes in 41% of the cases when the VIX is below the 95th percentile. *, ** and *** denote statistical significance at the 10%, 5% and 1% level, respectively.

Table 15: **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.

Table 16: **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 17: **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
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
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

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). 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 18: **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.

Figure 2: Time-varying performance of investment strategies

Panel A: Equally weighted portfolio

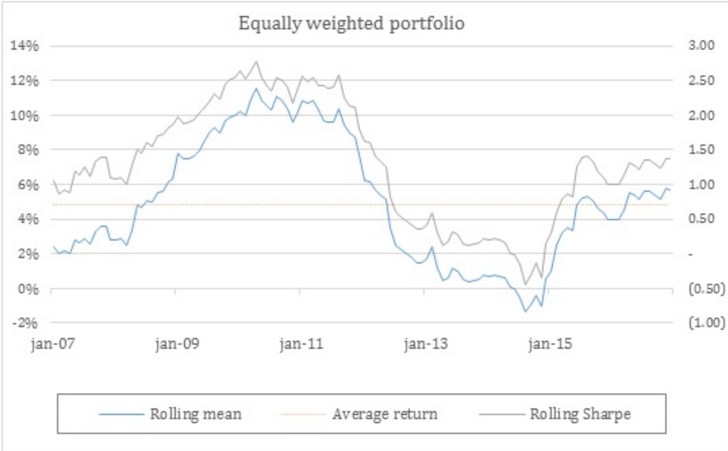

Panel B: Volatility weighted portfolio

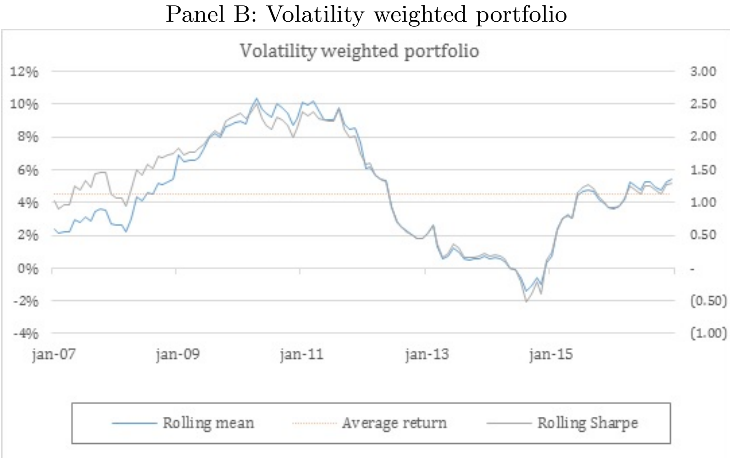

Panel C: Long - short portfolio

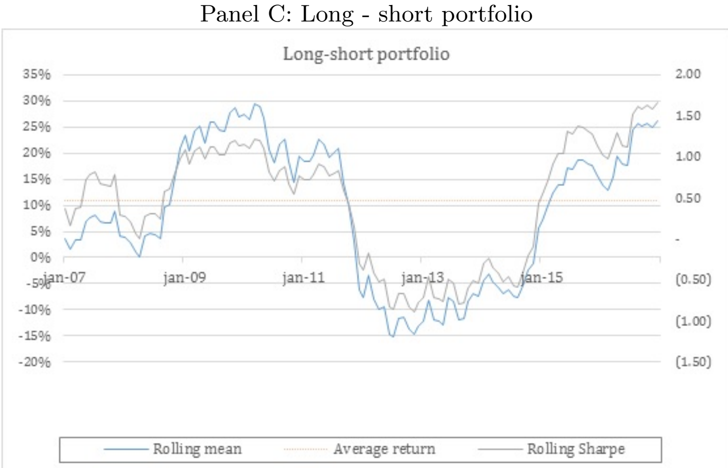

Notes: This figure presents the three-year rolling average performance of the equally weighted (Panel A), volatility weighted (Panel B), and long-short (Panel C) portfolios. The construction of portfolios is described in Section 3.5.