

Medeiros, Marcelo C.; Pires, Henrique

Working Paper

The proper use of Google Trends in forecasting models

Texto para discussão, No. 683

Provided in Cooperation with:

Departamento de Economia, Pontifícia Universidade Católica do Rio de Janeiro

Suggested Citation: Medeiros, Marcelo C.; Pires, Henrique (2021) : The proper use of Google Trends in forecasting models, Texto para discussão, No. 683, Pontifícia Universidade Católica do Rio de Janeiro (PUC-Rio), Departamento de Economia, Rio de Janeiro

This Version is available at:

<https://hdl.handle.net/10419/249731>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

TEXTO PARA DISCUSSÃO

No. 683

The Proper Use of Google Trends in
Forecasting Models

Marcelo C. Medeiros
Henrique Pires



The Proper Use of Google Trends in Forecasting Models

Marcelo C. Medeiros*

Pontifical Catholic University of Rio de Janeiro

Henrique F. Pires

Pontifical Catholic University of Rio de Janeiro

First version: Fev 2021

This version: Mar 2021

Abstract

It is widely known that **Google Trends** have become one of the most popular free tools used by forecasters both in academics and in the private and public sectors. There are many papers, from several different fields, concluding that **Google Trends** improve forecasts' accuracy. However, what seems to be widely unknown, is that each sample of Google search data is different from the other, even if you set the same search term, data and location. This means that it is possible to find arbitrary conclusions merely by chance. This paper aims to show why and when it can become a problem and how to overcome this obstacle.

Keywords: Google trends, forecasting, nowcasting, big data.

*Updated versions of this paper will be available on the corresponding author's webpage.

1 Introduction

Google Trends is an incredible Google tool, available at no cost, that shows the most popular terms searched in the recent past. It provides access to a sample of actual search requests made to Google, allowing one to look into the interest in a particular topic from around the globe or down to city-level geography. Several studies from different used this resource mainly to improve forecast accuracy. See, for example, Askitas and Zimmermann [2009], Artola and Galan [2012], Baker and Fradkin [2017], Naccarato et al. [2018], and Ferrara and Simoni [2019].

The tool provides the frequency in which a particular term is searched for in several languages from various regions of the world. However, Google normalizes the search data to make comparisons between terms easier. This means that search results are normalized to the time and location of a query by the following process. First, each data point is divided by the total searches of the location and time range chosen by the user. Then, the resulting numbers are scaled from 0 to 100 based on the topic's proportion to all searches on all topics.

By doing so, **Google Trends** data represent the relative popularity. This feature avoid the situation where places with the most search volume would always be ranked the highest. Naturally, this normalization also implies that, if different regions show the same search interest for a term, it does not always mean that they actually have the same total search volumes.

All these Google Trends features described above are widely known by its users. However, what most do not seem to know (or to ignore) is the fact that almost every time one looks for a **Google Trend's** term and downloads it, she gets a different series of numbers. The reason for this is that Google handle billions of searches per day, which means that providing access to the entire data set would be too large to process quickly. To circumvent this obstacle, only a small sample of Google searches are actually used in Google Trends. By sampling data, one can look at a dataset - generally - representative of all Google searches and that can be processed within minutes of an event happening in the real world.

However, there are some cases when this feature can actually become a real problem if not treated. We are going to show in this paper in which cases it is important to look at **Google Trend** data with some care. Moreover, we are going to show with both simulated and real data that it is possible to keep using **Google Trends** as covariates (and indeed important ones) without having the problem of your dataset not representing the actual population of the term's search for some specific period of time and location.

The article is divided as follows. Section 2 explains the problem, highlighting how this sampling feature of **Google Trends** may be specially harmful for exercises involving nowcasts and the use of data vintages. Section 3 shows with a simulation exercise how can one strongly improves model's accuracy when using **Google Trends** as covariates. Section 4, the last one before the conclusion, will illustrate with real data how our proposed strategy to deal with **Google Trends** can improve forecast accuracy.

2 How representative of the actual data is each sample?

As previously explained, Google makes available only a small sample of its search database. What most researchers and practitioners seem not to know (or to ignore) is the fact that this small sample is not always the same. In fact, it is constantly changing. This means that someone who downloads **Google Trends** data today will not download the same data tomorrow, even if she filters the same topics, languages and location.

We show with two examples of search topics (“Refined Petroleum” and “GDP Growth”) in two different regions (US and Brazil) that this difference in the data will be higher the less often the term is searched. Even though we do not have access to the actual number of searches for each term in each region, it is reasonable to assume that “Refined Petroleum” is less searched than “GDP Growth” and that less people use Google in Brazil than in the US (smaller population and percentage of people with access to the internet). This translates as the “Refined Petroleum” **Google Trends**’ search data in Brazil being the more volatile and the “GDP Growth” data in the US the less.

It is important to note that this feature is not something that Google tries to hide. One can find information about the sampling of **Google Trends** as easy as in its FAQ. Nevertheless, most studies using **Google Trends** do not seem to take it into consideration. For example, Pelat et al. [2009], Vicente et al. [2015], Ferrara and Simoni [2019], and Doerr and Gambacorta [2020] do not even mention this fact that each **Google Trends** sampling may result in different outcomes. Others like Heikkinen [2019], do recognize the importance of what this feature can imply, but do not seem to do anything about it in their estimations.

Objectively speaking, how bad can these different samples be to one’s research or forecasting model? To answer this question, we gathered many different samples from Google searches in the US and in Brazil between January, 2009 to January, 2019 (this way we avoid the period from both the 2008 crisis and the COVID-19 pandemic). In Figure 1 we report the values from three different samples for the same term (“GDP Growth”) in Brazil and same timestamp. The figure Figure illustrates well the fact that each **Google Trends** sample is different from the other, even when we set fixed the same topic, timestamp and region. As one can observe above, the black curve (random sample 1) does not even reach its maximum value in the same date as the other two. In case it is not very clear in the plot the lack of strong correlation between each series.

The correlation among series are reported in Panel (a) in Table 1. The table shows that the correlation between two different **Google Trends** sample can be as low as 0.496, even when we consider a relatively popular topic in Economics. Moreover, as the following plot and correlation matrix show, this finding is not an exclusivity of Brazil. See, Figure 2 and Panel (b) in Table 1, where we report the results for the case of US.

But how can a researcher use such a volatile dataset in her forecast? We propose a very simple solution in the next subsection.




Figure 1: Three different samples of same topic and date in Brazil.

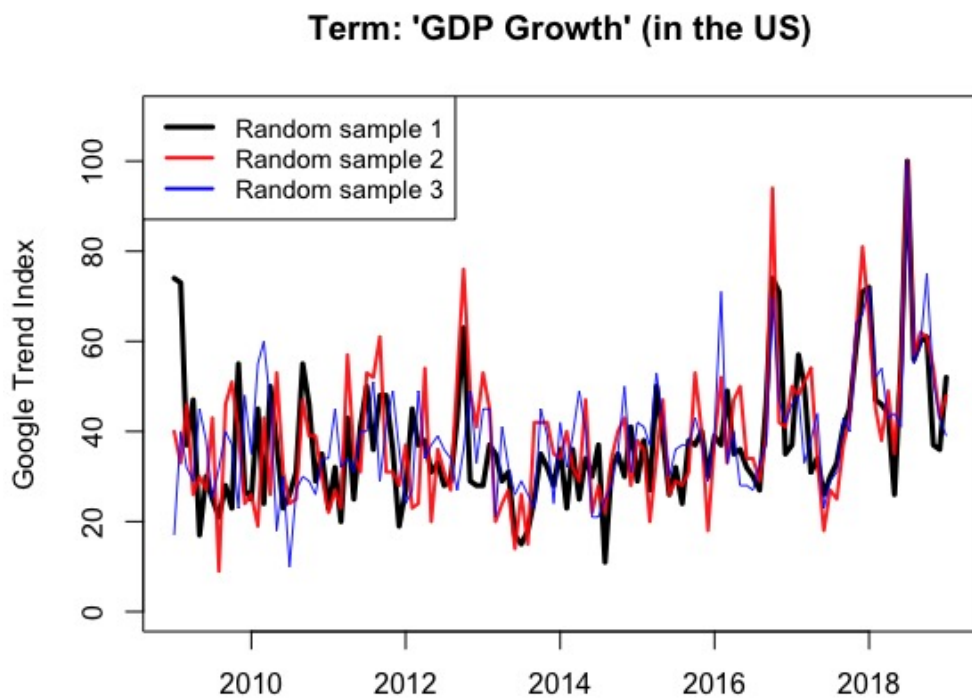


Figure 2: Three different samples of same topic and date in the US.

Table 1: Correlation between three different samples (S).

The table reports the correlation matrix among three different samples for “GDP Growth” searches. Panel (a) reports the case of Brazil and Panel (b) shows the US case.

Panel (a): Brazil			
	GDP Growth - S1	GDP Growth - S2	GDP Growth - S3
GDP Growth - S1	1	0.496	0.545
GDP Growth - S2	0.496	1	0.564
GDP Growth - S3	0.545	0.564	1

Panel (b): US			
	GDP Growth - S1	GDP Growth - S2	GDP Growth - S3
GDP Growth - S1	1	0.655	0.516
GDP Growth - S2	0.655	1	0.575
GDP Growth - S3	0.516	0.575	1

2.1 A simple way to circumvent the problem

The plots previously displayed¹ illustrate well the motivation of this article, by showing that each **Google Trend** sample may be very different from each other. But does it imply that **Google Trends** should not be used as research tools? Not at all.

First of all, it is important to mention that very popular terms (e.g. COVID-19 in 2020) do not vary so much among different samples. However, in many situations our terms of interest are not these very popular ones. So how to overcome the possible problem shown in the plots above? The answer is very simple. By gathering many different samples and averaging across every term, one can get a more reliable time series of that term.

To illustrate how taking averages of many samples improve the series consistency, Figures 3 and 4 below plot the curve for the same term and same timestamp, but now comparing averages from different samples. Each curve represent the average of “GDP Growth” searches in seven different random samples. It is important to note that if some sample is in one average, it necessarily is not in the other one. As one can observe in the figures, taking averages strongly raise the correlation between the (averaged) samples. Besides that, as expected, the correlation between each series is a little higher in the US (0.95) than in Brazil (0.92).

¹In the Appendix we show another setup in which the sampling problem with **Google Trends** may arise and should be treated with care.




Figure 3: Averaged series using different samples in Brazil.




Figure 4: Averaged series using different samples in the US.

3 Simulation: Model selection improvement

The current section will exemplify through a simulation exercise how can one improve model selection performance simply by using as covariates averages of multiple **Google Trends** samples instead of a single sample. The idea here is to observe the capacity of LASSO (Least Absolute Shrinkage and Selection Operator, Tibshirani [1996]) to select the correct variables (i.e., variables that indeed make part of the Data Generation Process or DGP) in two different linear regression setups.

For this exercise, we gathered 28 different **Google Trends** samples (14 for Brazil and 14 for the US) collected in different days but using the same timestamps and terms (in its respective language). We consider 20 different search terms related to Economics. When the location is Brazil, we search for terms in Portuguese, while when the location is the US, we search for the same terms in English. The searches comprehend the period from January 2009 to December 2018 (120 months). The searches are conducted at the monthly frequency. In the first setup², we construct different dependent variables for each replication using only a few variables extracted from a single random sample (from the 14 downloaded) of our **Google Trends** data. Then, we run 13 different LASSO regressions, each of which having as covariates only one sample of our **Google trends** data (except the one that was used to construct the dependent variables in that iteration) to try to identify which variables (among the 20 terms in the database) are indeed in that DGP.

In the second setup we do a similar replication exercise, but in this case in each iteration we construct the dependent variables using a few variables from the average of seven (randomly selected) **Google Trends** samples to then run a single LASSO regression (using the remaining seven samples among the 14) in order to try to select the correct variables in each DGP. Finally, we compare the average (across replications) of correctly selected variables between setups 1 and 2 and conclude that using the average strongly increases model selection performance.

More specifically, in the first setup, we simulate 1,000 replications with 120 observations each. In each iteration, a random number (s) ranging from 1 to 14 is drawn. This will indicate which one of the 14 **Google Trends** sample we will use to generate the data. Then, we randomly pick five of the 20 search terms to form the set of relevant variables of replication i that will actually be used to construct using sample s ($\mathbf{X}_{0,i,t,s}$). Finally, for each location (Brazil or US) we generate three different dependent variables given as

$$\begin{aligned} Y_{i,k,t,s}^{BR} &= \beta_{i,k}^{BR} \mathbf{X}_{0,i,t,s}^{BR} + \epsilon_{i,k,t,s}^{BR} \\ Y_{i,k,t,s}^{US} &= \beta_{i,k}^{US} \mathbf{X}_{0,i,t,s}^{US} + \epsilon_{i,k,t,s}^{US}, \end{aligned} \tag{1}$$

where $i = 1, 2, \dots, 1000$, $k = 1, 2$ or 3 indexes the DGP and $t = 1, 2, \dots, 120$. The linear

²In Appendix B we write the methodology in an algorithm format.

coefficients are determined as follows. $\beta_{i,1}$ is a vector of integers sampled from a discrete uniform distribution in the interval $[-10, 10]$; $\beta_{i,2}$ has elements set to 1 or 2 with equal probabilities; and $\beta_{i,3}$ is sampled from a continuous uniform distribution in the interval $[0, 1]$. Finally, $\epsilon_{i,k,t,s}^{BR}$ and $\epsilon_{i,k,t,s}^{US}$ are independent and normally distributed with zero mean and are set to have the same variance as $\beta_{i,k}^{BR} \mathbf{X}_{0,i,t,s}^{BR}$ and $\beta_{i,k}^{US} \mathbf{X}_{0,i,t,s}^{US}$.

Then, for each replication in setup 1, after all six dependent variables are constructed, we run 13 different LASSO regressions (each one using one of the 13 remaining samples that are not used in the GDP of that iteration) to then save the topics each one of them selected as predictors. It is important to remember that each of these 13 models predicts each $Y_{i,k,t,s}$ using the full database (which has 20 variables, each one being a time series related to a different **Google Trends** term), whereas each $Y_{i,k,t,s}$ is constructed using only five random terms among the 20 possible.

As previously explained, the second setup brings a similar idea. As in setup 1, in every replication the dependent variables will be constructed using only five random terms (among the 20 topics in each database). However, instead of randomly picking only one among the 14 **Google Trends** data, in each iteration we will randomly draw seven databases, take their average for each topic and then use the five previously selected topics to construct the dependent variables. Then, after all six dependent variables are constructed, the algorithm will now run only one LASSO regression (instead of 13) which will use the average of the remaining 7 samples (the ones that are used to construct each Y) as covariates. Finally, the LASSO regression selects which variables (from the 20 topics available - each one being the average of that topic across the seven remaining samples) are in that iteration's DGP. The rest of the process is analogous to the first setup.

Our conclusions - which are summarized in Table 2 - are that one can highly increase model selection performance (up to 27% in our simulations) with LASSO when using the average of multiple terms instead of a single sample from **Google Trend's** website.

Table 2: Simulation Results.

The table reports the mean percentage of correctly selected variables in each model and each simulation setup.

Setup	US^1	US^2	US^3	BR^1	BR^2	BR^3
1	51.1	62.7	58.3	56.2	68.1	61.1
2	64.72	73.84	66.82	67.62	76.08	68.76

4 Empirical Application

This section will advocate both for the use of **Google Trends** as a powerful prediction tool as well for the use of the average of many samples of the same search pattern as your independent

variable, instead of only one sample.

During the pandemic, many authors tried to predict and understand the pattern of COVID-19 cases and death counts with all sort of models, including both epidemiological and statistical ones; see, for example, Medeiros et al. [2020], Bertozzia et al. [2020], Anastassopoulou et al. [2020], Carneiro et al. [2020], and Matjaž et al. [2020]. One of the most important indicators of the actual severity of the pandemic, is the number of cases organized by the date of first symptoms felt by each person as well as the daily new deaths sorted by the day of the event (and not the date of registry). These numbers give a clear indication how the disease is evolving in real-time.

However, in most countries and specially in Brazil, the actual numbers are available with major delays (more or less 2 weeks). For some locations in Brazil, the delay can exceed a month. Part of the delay is due to the coronavirus cycle but most of it is related to the bureaucracy in the register system.³ This delay means that by using **Google Trends** (which is available almost in real-time) one could have a good insight of the direction in which the numbers would be moving towards two weeks in advance. We call this *nowcasting*.

However, would it be possible to correctly nowcast the trend of SARS (syndrome respiratoire aigu sévère, strongly associated with COVID-19 in 2020) cases in Brazil using merely a handful of **Google Trends** topics as predictors?

To answer this question, we consider a simple LASSO regression to predict the trend of SARS new cases and new deaths in 2020 in Brazil. For each dependent variable (new cases and new deaths) we estimate eight LASSO regressions (each one using a different **Google Trends**’ sample of the same topics, timestamps and location, but collected in different days) from February to May 2020 using only **Google Trends** data with terms related to the pandemic, such as “COVID symptoms”, “COVID ICU” and “coronavirus hospital” as predictors⁴. Figure 5 displays the results of the nowcasting model for SRAG’s new cases. The black curve in the figure represents the actual number of new SARS cases aggregated by the reported day of the first symptoms felt by each person. The colored curves represent the point forecasts of the eight different samples of the same **Google Trends** terms starting on June 1, 2020.

³For further reference with respect to the Brazilian case, please visit the following link.

⁴Naturally each of these terms translated to Portuguese.




Figure 5: Nowcast of SARS daily cases in Brazil using only Google Trends.

As it is clear in Figure 5, it would have been possible to nowcast in real-time the trend of SARS new cases using either one of the samples. In other words, by using only **Google Trends** data, one could predict that the number of new cases would continue on a high plateau until August, then it would start falling until Mid-November to then start rising again until the last day of the year. By looking at Figure 6 below, we could conclude that it would also have been possible to predict the trend of SARS new deaths by using only **Google Trends** data.




Figure 6: Nowcast of SARS daily deaths in Brazil using only Google Trends.

But now putting on the “black hat”, there is another thing to consider. Even though the eight predictions (using different samples) are similar, they are not the same. This is specially meaningful if we take into consideration that all of the topics used in the models described above could be considered very popular terms in the year of 2020.

To understand how using the average of the eight samples across each term would improve the forecasts results for both dependent variables, we compared the Root Mean Squared Error (RMSE) of each model (New cases and New deaths) using averaged series (Proposed Model) with the other eight models that used **only one Google Trends** sample. The results are displayed in Table 3.

Table 3: Nowcasting Results.

The table reports the Root Mean Squared Error (RMSE) of the nowcasting LASSO regression using the average of the **Google Trend** samples as regressors as well as the RMSE for the sample-specific models with the worst, best and average performance.

	Proposed Model	Worst	Best	Average
New cases’ RMSE	888.6	1340.97	900.6	1058.13
New deaths’ RMSE	149.2	175.3	144.4	159.3

As one can infer above, for the eight new cases’ models, not only none of them had a smaller RMSE than the average, but we found out that the RMSE was up to 51% higher in the models using only one sample. For the new death’s models, the proposed method had a RMSE more than 6% smaller than the average RMSE of the other models and more than 17% smaller than the worse one.

These numbers above illustrate the fact that merely by chance an uninformed researcher could find herself in one of the following situations when using only a single sample of **Google Trends** in her studies. First, it could be possible that a given term is actually a good predictor of a variable of interest, but she may unfortunately find herself working with a specific random sample of **Google Trends** data that does not represent very well the true population. In this case, she may end up concluding that there is not a significant relationship between the searches on the term and the dependent variable. The other situation is the one that she may find an apparent significant relation when there is not. By averaging multiple samples, one can diminish this risk.

5 Conclusions

This article has reinforced the usefulness of internet search data as a powerful forecast tool. However, it has also highlighted some limitations of these data, such as the lack of information

on the actual volume of searches, but mainly the fact that Google Trends’ index is based on a subsample that is changing all the time. We explained that this feature might become a real problem for the forecaster when dealing with less frequent searched topics and/or when filtering the query to many years in the past.

We then illustrated with two examples using both simulated and real data that simply by taking averages of many different samples with the same specifications (i.e. same topic, timestamp and location) it is possible to improve both model selection and its forecast accuracy.

References

- M. Alexander, K. Polimis, and E. Zagheni. Combining social media and survey data to nowcast migrant stocks in the united states. *Population Research and Policy Review*, 2020.
- C. Anastassopoulou, L. Russo, A. Tsakris, and C. Siettos. Data-based analysis, modelling and forecasting of the covid-19 outbreak. *PLoS ONE*, 15(3), 2020.
- C. Artola and E. Galan. Tracking the future on the web: Construction of leading indicators using internet searches. *Documentos Ocasionales (Bank of Spain)*, 1203:255—271, 2012.
- N. Askitas and K. Zimmermann. Google econometrics and unemployment nowcasting. *DIW Berlin*, Discussion Paper 899, 2009.
- S. Baker and A. Fradkin. The impact of unemployment insurance on job search: Evidence from google search data. *The Review of Economics and Statistics*, 99 (5):756–768, 2017.
- A. Bertozzia, E. Franco, G. Mohlerd, M. Shorte, and D. Sledgef. The challenges of modeling and forecasting the spread of covid-19. *Proceedings of the National Academy of Sciences*, 117 (29): 16732–16738, 2020.
- C. B. Carneiro, I. H. Ferreira, M. C. Medeiros, H. F. Pires, and E. Zilberman. Lockdown effects in US states: an artificial counterfactual approach. *arXiv e-prints*, art. arXiv:2009.13484, Sept. 2020.
- M. Carvalho and A. Rua. Real-time nowcasting the us output gap: Singular spectrum analysis at work. *International Journal of Forecasting*, 33 (1):185–198, 2017.
- S. Doerr and L. Gambacorta. Identifying regions at risk with google trends: the impact of covid-19 on us labour markets. *BIS Bulletin*, 8:1–6, 2020.
- L. Ferrara and A. Simoni. When are google data useful to nowcast gdp? an approach via pre-selection and shrinkage. Working paper, Center for Research in Economics and Statistics, 2019.

- J. Heikkinen. Nowcasting gdp growth using google trends. Master's thesis, Jyväskylä University School of Business and Economics, 2019.
- P. Matjaž, G. Nina, S. Mitja, and S. Andraž. Forecasting covid-19. *Frontiers in Physics* ., 8, 2020.
- M. Medeiros, A. Street, D. Valladão, G. Vasconcelos, and E. Zilberman. Short-Term Covid-19 Forecast for Latecomers. *arXiv e-prints*, art. arXiv:2004.07977, Apr. 2020.
- A. Naccaratoa, S. Falorsib, S. Lorigab, and A. Pierinia. Combining official and google trends data to forecast the italian youth unemployment rate. *Technological Forecasting & Social Change*, 130:114–122, 2018.
- C. Pelat, C. Turbelin, A. Bar-Hen, A. Flahault, and A. Valleron. More diseases tracked by using google trends. *Emerging Infectious Diseases*, 15(8):1327–1328, 2009.
- A. Richardson, T. Mulder, and T. Vehbi. Nowcasting new zealand gdp using machine learning algorithms. Working Paper 47, Centre for Applied Macroeconomic Analysis, 2018.
- R. Tibshirani. Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society.*, 58 (1):267–288, 1996.
- M. Vicente, A. López-Menéndez, and R. Pérez. Forecasting unemployment with internet search data: Does it help to improve predictions when job destruction is skyrocketing? *Technological Forecasting & Social Change*, 92:132–139, 2015.
- N. Woloszko. Tracking activity in real time with google trends. Working Paper 1634, OECD Economics Department, 2020.

A Forecast models in Real Time

In 2015, **Google Trends** data was made available in real time. It is no doubt that it is helping people around the world explore the global reaction to major events. But not only that. The way **Google Trends** is constructed allows one to verify how was the search patterns in the past for a specific term and location. This feature makes possible the construction of vintages of the same search patterns, reproducing for any given day in that range, information that was available to a real-time forecaster.

That being said, it becomes clear that due to the fact that the use of real-time forecast and nowcast has been increasingly rising (some examples are Carvalho and Rua [2017], Richardson et al. [2018], Woloszko [2020], and Alexander et al. [2020]), the possibility of constructing vintages of **Google Trends** is a very powerful feature of this tool. However, the same possible issue shown in previous sections may arise in a setup of real-time forecast and/or nowcast.

To illustrate the issue, we display below a series of figures. Each one contains three search pattern for the same term one month ahead. For example, the first figure below brings the time series of “Refined Petroleum” searches in the US from: 1) Jan, 2004 to Jan, 2014; 2) Feb, 2004 - Feb, 2014; 3) Mar, 2004 - Mar, 2014. This would be what a real-time forecaster could gather when constructing the vintages.




Figure 7: Vintages of “Refined Petroleum” topic in the US.

It is very clear from Figure 7 that each series differ a lot from each other. As a matter of fact, the highest correlation between these series is 0.27 (between black and blue lines) and the lowest is 0.01! It is important to remember that the values of the volume index displayed by **Google Trends** is normalized and therefore could indeed differ from one vintage to another if the period

of search was different. However, it is clear from the plot above that no new maximum search volume index is achieved neither in February, 2014 nor in March, 2014, i.e., the maximum value (100) of the three plots should be the same, when it is not.

Another important thing to mention is that we are showing a very extreme example for two main reasons: 1) we went back very further in time, when fewer people used Google, which means their database should be less consistent; 2) we searched for a very-low volume term, contributing to the low consistency of the data. However, if one intends to perform real-time analysis with **Google Trends**, it is very likely that she will go back that further in time. Besides that, depending on what she is trying to forecast (e.g., volume or value of exports), she may try to use those low-volume terms in **Google Trends** as covariates in her forecast.

In Figure 8 we display the same plot as above, but now with timestamps going from Jan, 2009 - Jan, 2019 to Mar, 2009 - Mar, 2019. We can see that the correlations between each sample becomes a little bit higher: the highest one is one 0.43 and the lowest 0.29.




Figure 8: More recent vintages of 'Refined Petroleum' topic in the US.

Figure 9 displays search patterns of "US inflation", a far more popular term. We can see that even for the 2004-2014 period, the correlations are already higher (maximum of 0.64 and minimum of 0.44) than the ones from the "Refined Petroleum" topic.




Figure 9: Vintages of 'US Inflation' topic in the US.

It is clear that in recent years correlations can be as high as 0.85. However, this does not mean that taking averages of many samples of these series could not improve their consistency with respect to the actual search pattern of the term in each period of time.




Figure 10: More recent vintages of "US inflation" topic in the US.

B Algorithms used in the simulation

B.1 Setup 1

Algorithm 1

- 1: **for** $iteration = 1, 2, \dots, 1000$ **do**
 - 2: sample $s =$ a random integer from 1 to 14
 - 3: sample $\vec{c} = 5$ random variables from the 20 available
 - 4: for $k = 1, 2$ and 3 , construct each $Y_{k,t,s}^{US}$ and each $Y_{k,t,s}^{BR}$ using only variables $\in \vec{c}$
 - 5: **for** $m \in \{1, 2, \dots, 14\} - s$ **do**
 - 6: Run a LASSO using sample m to predict each $Y_{k,t,s}^{BR}$ and $Y_{k,t,s}^{US}$
 - 7: Save which variables were selected by the model to forecast each dependent variable
 - 8: **end for**
 - 9: **end for**
 - 10: Compute average percentage of variables corrected selected by each LASSO.
-

B.2 Setup 2

Algorithm 2

- 1: **for** $iteration = 1, 2, \dots, 1000$ **do**
 - 2: sample $\vec{s} = 7$ random integers from 1 to 14 without replacing
 - 3: for every topic in each sample (20), compute its average across each sample $\in \vec{s}$
 - 4: sample $\vec{c} = 5$ random variables from the 20 available
 - 5: for $k = 1, 2$ and 3 , construct each $Y_{k,t,s}^{US}$ and each $Y_{k,t,s}^{BR}$ using variables $\in \vec{c}$
 - 6: Do the same as step 3 but for the other 7 samples $\ni \vec{s}$
 - 7: Run a single LASSO using the averaged sample from step 6 to predict each $Y_{t,s}^{BR,US}$
 - 8: Save which variables were selected by the model to forecast each dependent variable
 - 9: **end for**
 - 10: Compute average percentage of variables corrected selected by each LASSO.
-