

Building Longitudinal *Google Trends* to Measure Dynamic Local-Level Issue Attention*

Taeyong Park[†]

Haewoon Kwak[‡]

Jisun An[§]

Abstract

Google search indices can be useful for measuring dynamic local-level attention to public issues for which survey data are rarely available. However, there is a practical difficulty with generating longitudinal Google Trends. Google Trends provides normalized counts from zero to 100 instead of absolute counts, thereby placing its cross-sectional indices across different times on different scales. Thus, merely pooling cross-sectional data fails to create desirable longitudinal data. We develop a method for rescaling Google Trends indices to build longitudinal data at the United States media market level. We illustrate this method with applications to the issues of employment and the coronavirus. This new tool opens the door to using Google searches merged with various kinds of time-series cross-sectional data, which has not been possible.

*This paper is supported by the Qatar Foundation Seed Research Project Fund.

[†]Arts and Sciences, Carnegie Mellon University in Qatar.

[‡]School of Information Systems, Singapore Management University.

[§]School of Information Systems, Singapore Management University.

Introduction

Google search data have become a popular source for measuring public attention to political issues. Google searches reveal people’s strong interest in specific issues.¹ The traces people leave during internet searching are naturally occurring, not affected by researcher intervention (Japac et al. 2015). Additionally, while people are honest when searching for information on the internet, they could lie or misreport in surveys to give socially desirable answers (Stephens-Davidowitz 2017; Tourangeau, Groves and Redline 2010).

The massive volume of Google searches can be especially useful to measure issue attention at local levels, such as cities and media markets. Nationally representative “most important problem” (MIP) surveys, a conventional method for measuring issue attention, mostly lack proper sample sizes for local communities. This problem often leads researchers to make inferences about communities based on a few people. In contrast, Google searches’ vast amounts of information make it possible to investigate how issue attention depends on local contexts.

Furthermore, instant access to Google searches permits tracking issue attention in real time. It is critical to capture real-time trends because issue attention tends to be dynamic and ephemeral (Dennison 2019; Jones 1994; Ripberger 2011). For example, the coronavirus pandemic continues to evolve and affects different regions across the United States differently. Moreover, public attention to the virus might soar and subside rather quickly. Unfortunately, though, longitudinal panel surveys that enable researchers to compare local-level dynamic issue attention are nonexistent to the best of our knowledge. For this reason, we pay attention to the potential that longitudinal Google search data, if available, fill the gap.²

Despite this potential benefit, little research has explored longitudinal Google search data. Although Bento et al. (2020), Mellon (2014) and Ripberger (2011) study time-varying Google Trends, none of the studies uses longitudinal data. The lack of research utilizing longitudinal Google searches is primarily due to the lack of longitudinal data per se. One of the major obstacles to generating longitudinal Google Trends is that Google Trends provides normalized

indices from zero to 100, instead of absolute counts. While this normalization scheme allows for comparing relative search interest across geographic regions for a specific time period, it places cross-sectional indices at different times on different scales. Consequently, pooled cross-sectional indices are different from desirable longitudinal data.

To resolve this problem, we develop a novel method for rescaling Google search indices and provide guidance on constructing longitudinal Google Trends. It is imperative to create longitudinal Google Trends although the feasibility of the longitudinal data as a measure of time-varying local issue attention should be tested in future research. Indeed, testing for the feasibility is impossible unless the longitudinal data exist.

In the following section, we discuss the importance of measuring local-level dynamic issue attention. Then, we introduce the rescaling method with illustrative examples of employment and the coronavirus. Finally, we conclude with some caveats.

Measuring Local Public Attention

We focus on longitudinal Google Trends at the media-market level in the U.S.³ Google Trends provides numerous countries' subregion data, including U.S. state searches. However, data on the 210 U.S. media markets, or Designated Market Areas (DMAs), are the only metro-level data accessible by Google Trends. Although city-level searches are also available, Google Trends does not cover all cities in any single country. Thus, currently, the U.S. DMA data permit exhaustive comparisons of Google searches at the lowest subnational level.

Local-level comparisons are important in the study of political attitudes towards public issues. How people perceive public issues may vary by the extent to which people acquire information from their local experiences. Previous research suggests local contexts matter for people's perceptions of the national economy (Ansolabehere, Meredith and Snowberg 2014; Reeves and Gimpel 2012) and racial groups (Wong 2007). Cramer (2016) argues that local communities function as a lens through which people interpret politics.

Scholars have also examined the role of local media in influencing political attitudes and behavior. Gilliam and Iyengar (2000) show that the racial element of the crime script in local television news increases support for punitive approaches to crime and heightens negative attitudes about African-Americans among white viewers. Filla and Johnson (2010) and Althaus and Trautman (2008) suggest that people’s exposure to local news affects voter turnout. Additionally, Darr, Hitt and Dunaway (2018) argue that residents of communities where local newspapers close tend to be politically polarized as they rely more on national news or partisan heuristics.

Given that local communities and local media influence how people perceive and respond to public issues, DMAs are arguably better subnational indicators of political heterogeneity than states. States are vast spatial containers. Statewide conditions sometimes do not reflect the circumstances local communities experience. DMAs are smaller and reflect inter-state areas that receive the same local television and radio station offerings. For instance, St. Louis City, Missouri, probably has more in common with East St. Louis, Illinois, than Kansas City, Missouri, and St. Louis City and East St. Louis belong to the same DMA.

Therefore, measuring public attention to political issues at media markets helps future studies about what a local community is and how it engages with political information. The method introduced below allows researchers to compare DMA-level time-varying issue attention. Despite our focus on DMAs, it is worth noting that the method is also applicable to U.S. state-level and cross-national Google Trends. Accordingly, the proposed method opens the door to using Google Trends merged with various kinds of time-series cross-sectional data.

Building Longitudinal Google Trends

It is straightforward to compile time-series data for a particular region or cross-sectional data for a specific period using normalized Google Trends indices. However, cross-sectional indices

obtained in different time windows have different scales, and merely pooling them does not generate desirable longitudinal data. Our solution is to rescale cross-sectional indices across different times to convert them into a common scale.

We consider two different types of Google searches separately because they require different approaches. One is a single-search index for one search term. Two is a comparison-search index for comparing multiple search terms.

Longitudinal Single-Search Index

To describe how to rescale Google single-search indices, we take an example from the issue of employment. Following Mellon (2014), suppose we use the search term “jobs” to measure weekly DMA-level prominence of employment issues. For simple discussion, we randomly selected 10 DMAs’ cross-sectional search results across 10 weeks as Table 1 displays. Each week’s search indices are placed on a normalized 0–100 scale.⁴

Google normalizes its search results following two steps. First, each DMA’s search interest for “jobs” is given as a percentage of the query’s search volumes out of the DMA’s all searches for the time selected. After that, the jobs search-interest percentages for individual DMAs are scaled on a 0-100 range such that the highest percentage becomes 100.⁵ In this way, users can compare relative interest in jobs among all DMAs. For instance, Rochester’s index 50 and Spokane’s index 60 in Week 9 indicate that Rochester’s relative interest in jobs is approximately 5/6 of Spokane’s.

Such inter-regional comparisons in the same week are possible. However, comparing normalized indices over time makes no sense. For example, Austin’s search indices 38 and 55 for the first and second weeks are proportional to the highest percentage for each of the two weeks, respectively. It is uncertain which index, 38 or 55, corresponds to higher search interest. Consequently, simply combining the 10 columns in Table 1 does not create longitudinal data. Before combining them, we must convert the cross-sectional indices across different weeks to the same scale.

Table 1: Weekly Search Index for “jobs” for September 18 - November 26, 2016: 10 Designated Market Areas

DMA	Week									
	1	2	3	4	5	6	7	8	9	10
Austin TX	38	55	55	55	42	52	50	46	55	19
Baltimore MD	43	66	62	62	47	66	65	56	67	23
Baton Rouge LA	39	55	59	66	40	65	51	54	64	22
Honolulu HI	38	57	55	57	38	58	54	51	57	19
Lexington KY	44	60	61	62	44	72	71	53	67	21
Rochester NY	33	49	44	48	30	46	44	44	50	17
Spokane WA	40	50	49	50	39	56	57	51	60	19
Syracuse NY	32	49	50	53	34	49	47	49	49	18
Toledo OH	37	54	57	55	40	56	47	55	64	19
Wilmington NC	37	67	51	57	46	62	48	51	63	24

[Notes: Week 1 refers to September 18 - 24, Week 2 means September 25 - October 1, and so forth.]

We present a three-step guideline with an example of the weekly jobs indices between December 31, 2006, and December 31, 2016. The first step is to select a reference that supplies a common scale. Time-series search data for each DMA are accessible by Google Trends, and they are on the same scale across different times. We use the time-series data for Honolulu, Hawaii, as our reference without loss of generality.⁶

Second, we may need to divide the reference time-series into subset series and concatenate them after rescaling, depending on the time-series length. We call this the divide-and-concatenate process. For instance, our example covers 522 weeks. This period is long enough to require the divide-and-concatenate process because we found that Google Trends returned a monthly time-series, instead of a weekly series, for any period exceeding 269 weeks.⁷ Hence, we split the entire period into two: December 31, 2006–February 26, 2012, and February 26, 2012–December 31, 2016. Then, neither of the two periods exceeds 269 weeks. Contingent upon the reference time-series length, researchers may need to decompose it into more than two periods or may not need to divide it at all.

When splitting the reference time-series, it is crucial to have the preceding time window’s end date identical to the following window’s start date to use the date as a node connecting the two windows. Since each window produces 0-100 normalized counts, the node belonging

Table 2: Concatenating Reference Time-Series: Honolulu “jobs” Search

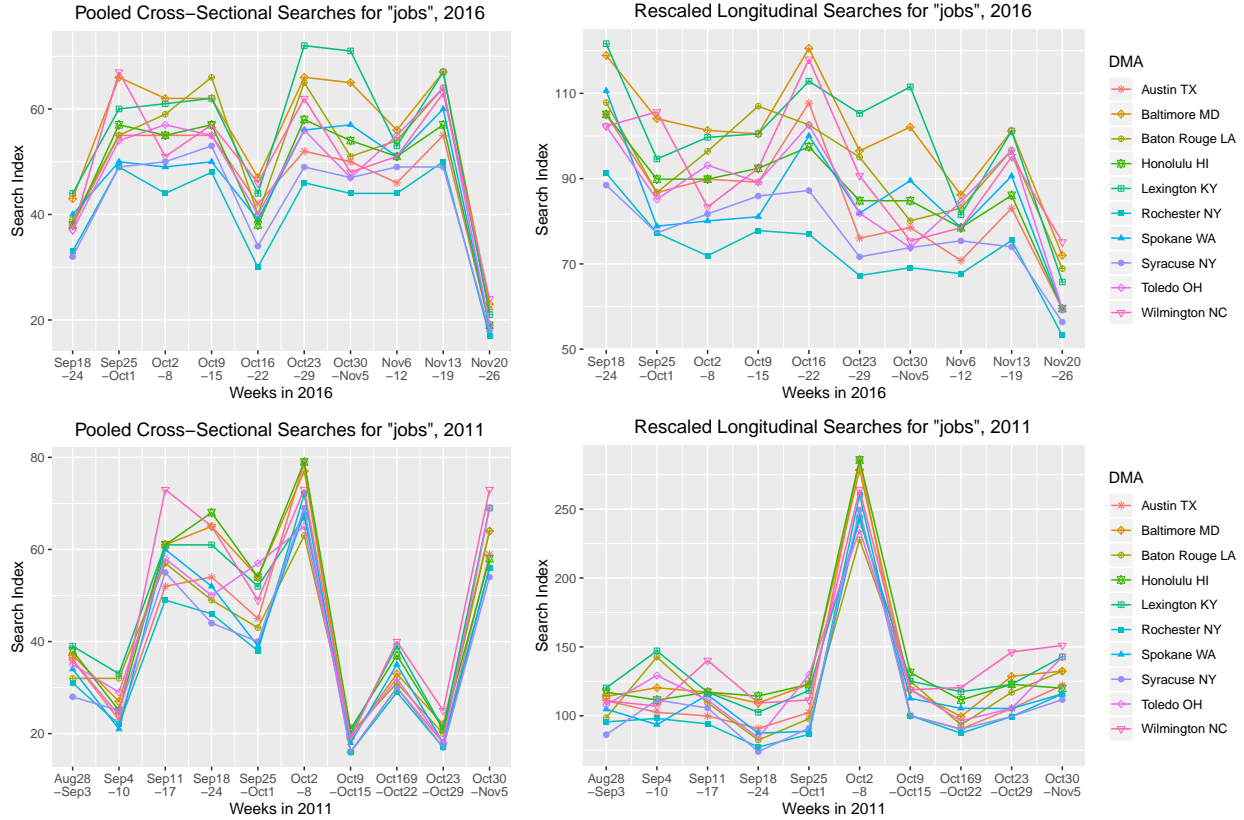
Search Window (Weekly)		Before Rescaling	After Rescaling
	⋮	⋮	⋮
Window 1	2/12/12	34	97.14286
	2/19/12	35	100
	2/26/12	35	100
Window 2	2/26/12	79	100
	3/4/12	80	101.2658
	3/11/12	74	93.67089
	⋮	⋮	⋮

to both windows should be used to transform normalized counts to the same scale. Our example’s node is 2/26/2012 as Table 2 shows. We begin by converting the node’s value into 100 in both windows to obtain the scaling ratios: $100/35=2.857143$ and $100/79=1.265823$.⁸ Then, we multiply all the first-window values by 2.857143 and all the second-window values by 1.265823.

Once the rescaled time-series is ready, as shown in the rightmost column of Table 2, the third step uses the rescaled reference time-series to rescale all the cross-sectional indices. For example, consider the first week from Table 1. Baltimore has a value of 43. Honolulu’s index is 38. For the same week, the reference Honolulu time-series gives a value of 105.0632911. We begin by calculating the ratio between the Honolulu time-series and the Honolulu cross-sectional index. This ratio is computed as $105.0632911/38 = 2.764823$. Then, we multiply all the cross-sectional indices in the corresponding week by this ratio. Accordingly, Honolulu’s index becomes 105.0632911, and Baltimore’s index turns into $43 \times 2.764823 = 118.8874$. For the second week from Table 1, the scaling ratio is $89.87341772/57 = 1.576727$, where 89.87341772 comes from the reference Honolulu time-series. Subsequently, Honolulu’s cross-sectional index is adjusted from 57 to 89.87341772. Baltimore’s index is rescaled to $66 \times 1.576727 = 104.064$. In this way, we convert all the cross-sectional indices into the same scale to make them comparable with one another.

Following the three-step guideline, we built longitudinal jobs search indices for the 210

Figure 1: Comparing Pooled Cross-Sectional Google Searches and Longitudinal Google Searches for “jobs”: Examples for 10 DMAs and 10 Weeks, 2016 and 2011



[Notes: We obtained the pooled cross-sectional Google searches directly from Google Trends. The longitudinal Google searches were generated by the method proposed in this paper. We selected the displayed examples from the jobs search indices for the 210 DMAs across the 522 weeks, from December 31, 2006, to December 31, 2016.]

DMAs across the 522 weeks. For simple discussion, the upper panel in Figure 1 corresponds to the 10 DMAs and 10 weeks in Table 1. The lower panel shows another subset for the 10 weeks from 2011. For each panel, the right plot presents rescaled longitudinal searches whereas the left plot presents pooled cross-sectional searches readily accessible by Google Trends.

We discuss two important implications by comparing the pooled cross-sectional data and the longitudinal data. First, the longitudinal data track the actual trend over time that the pooled cross-sectional data fail to capture. For instance, according to the upper-panel left graph, the first week’s search indices are lower than the second week’s indices. Contrarily,

the longitudinal data demonstrate that the first week has a higher search interest overall. The pattern revealed by the longitudinal data is the true trend because our rescaling scheme interweaves the time-series information with the cross-sectional data.

Similarly, the lower panel’s longitudinal indices accurately track the sudden rise and fall in the week of October 2, 2011. This dramatic surge seems due to the death of Apple co-founder Steve Jobs on October 5. By contrast, the pooled cross-sectional data fail to identify the real trend before and after the week. When using search indices as a measure of issue attention, it is critical to check their content validity. For example, as a measure of the prominence of employment issues, the job search index must not be confounded by searches for “Steve Jobs” (Mellon 2014). In this regard, tracking accurate trends can help researchers detect the presence of confounders.

The second implication to be discussed is that longitudinal data mitigate the problem that the least populated DMAs’ search indices misinform the relative search interest. For example, Glendive, Montana’s cross-sectional job search index is 100 in 240 weeks of the 522 weeks in our analysis. Several other DMAs, including Alpena, Michigan, and Zanesville, Ohio, have similar patterns. These results come from Google’s use of the search-interest ratio in calculating normalized indices. This ratio is the query’s search volume divided by the total number of searches on all topics for the selected region and time. Those least populated DMAs have relatively small amounts of total searches, giving a small value for the search-interest ratio’s denominator. Thus, their search-interest ratio easily surpasses larger DMAs’ when a certain issue becomes just slightly salient. Specifically, consider the deep valleys in the lower-left plot in Figure 1. In the weeks of October 9 and 23 in 2011, Glendive’s search-interest rate is far higher than all other DMAs’ rates and forces the other DMAs’ indices to be less than 30. In the week of October 30, the 10 DMAs’ indices rebound since no extremely high search-interest rate appears. These bumpy ups and downs misrepresenting the true trend do not occur in the longitudinal data because our rescaling scheme’s time-series component takes into account all weeks’ search indices.

Longitudinal Comparison-Search Index

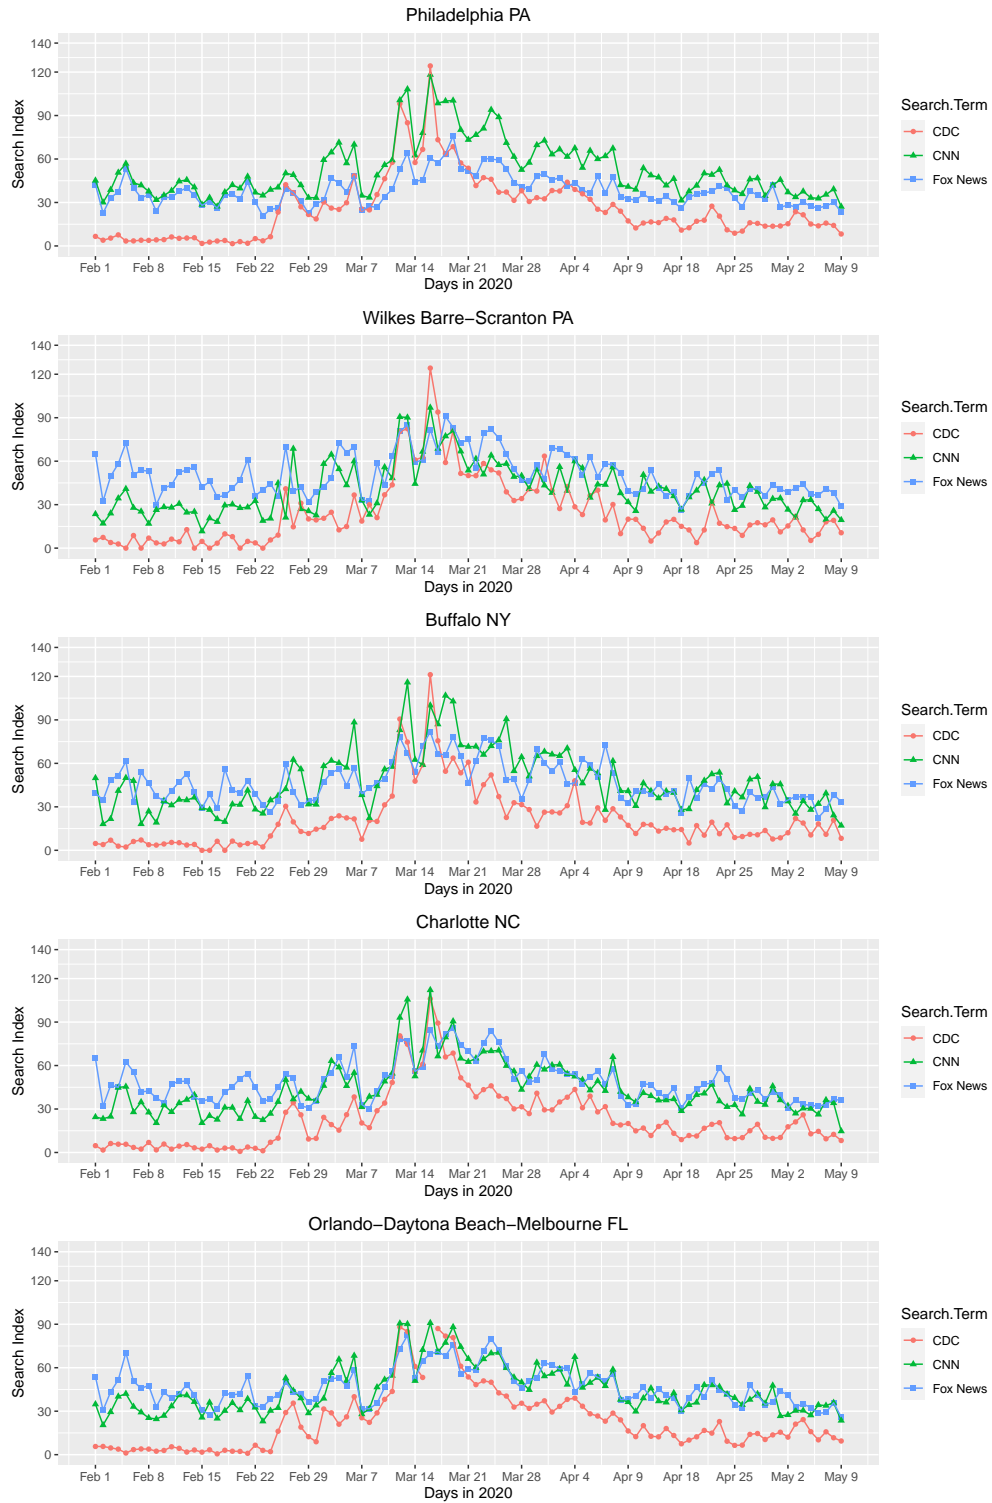
We now turn to a comparison-search index. Google Trends allows users to compare up to five distinct search terms simultaneously. As an illustrative example, we use daily searches for the Centers for Disease Control and Prevention (CDC), CNN, and Fox News over the period between February 1 - May 9, 2020. We use the search terms “cdc,” “cnn,” and “fox news.” Our preliminary analysis discusses whether people seek different information sources as their attention to the coronavirus evolves.

A comparison-search index is given by a ratio of each query’s searches to all queries’ total searches. For example, Philadelphia’s comparison-search indices for CDC, CNN, and Fox News for February 1 are 7, 48, and 45. These numbers indicate that the search volume for CDC is approximately $7/48$ and $7/45$ of those for CNN and Fox News, respectively. Those numbers are on the same scale. However, different days have different scales, depending on each day’s total searches for all queries. The three-step guidance below describes how to rescale comparison-search indices across times.

First, select a reference time-series. It is essential to choose a time-series from the search for a single query, not from the comparison search. Each day’s comparison-search indices for CDC, CNN, and Fox News are proportional to one another. Therefore, if we know a single query’s time-series for the whole period, whether CDC, CNN, or Fox News, we can rescale the three queries’ searches.⁹ We use the CNN time-series for Atlanta because it does not have a value of zero over the entire period. The second step performs the divide-and-concatenate process if necessary. Our example covering less than 269 days does not require this process. Finally, compute scaling ratios based on the reference time-series to rescale all the cross-sectional indices.

Figure 2 shows five examples of the rescaled longitudinal comparison-search indices for CDC, CNN, and Fox News. From the five examples and many other DMAs, we find that the search for CDC (red) relative to CNN (green) and Fox News (blue) starts increasing around February 24 when Italy reported a surge in their coronavirus cases. The CDC index hits

Figure 2: Longitudinal Daily Comparison-Searches for CDC, CNN, and Fox News during the COVID-19 Pandemic: Examples for Five DMAs, February 1 - May 9, 2020



its peak in mid-March and then gradually subsides. This pattern is similar to the overall longitudinal trend for the “coronavirus” search that we present in the online appendix. Thus, we deem that the comparison-search indices for CDC serve as a proxy for public attention to the pandemic.

We notice a subtle information-seeking pattern that would not be identified without the longitudinal data. When public attention to the coronavirus was relatively low until February 23, the search volumes for Fox News were higher than or similar to those for CNN in the five DMAs. After that time, people’s concern appears to push up the search for CNN more than Fox News as their attention to the virus outbreak increases. One hypothetical explanation is that people tend to seek out politically less biased information when searching for updates on the public health crisis, in that Fox News has been alleged of having Republican Party bias.¹⁰ This analysis is more or less anecdotal though we found similar patterns from multiple other DMAs, where the search for Fox News was higher than CNN before the virus outbreak. Future research should conduct statistical analyses to delve into further the pattern described here, while the present study focuses on illustrating the potential application of longitudinal Google Trends.

Additionally, we explore one aspect of the fault line between the rural and urban regions. According to Figure 2, the residents of the rural rust-belt Wilkes Barre-Scranton media market in Pennsylvania tend to search for Fox News more than CNN. In contrast, the residents of the Philadelphia DMA, about 120 miles away from Wilkes Barre-Scranton, tend to search for CNN more than Fox News. This rural-urban gap characterized by the different preferred news channels echoes previous research that refers to America’s rural-urban political divide as the main dimension of political conflict (Cramer 2016; Rodden 2019). The gap is also in line with the 2016 presidential election results. Trump performed well in Wilkes Barre-Scranton, receiving 69.7% of the votes in the Lycoming county and 69.4% in the Schuylkill county. By contrast, his vote share was 15.3% in the Philadelphia county within the Philadelphia DMA.¹¹ This example shows an avenue for the potential use of longitudinal Google Trends

in measuring local-level political heterogeneity.

Conclusion

In this article, we develop a new method for constructing DMA-level longitudinal Google Trends. We also discuss several avenues for the application of the longitudinal data. Measuring time-varying local-level issue attention can open the door to new research questions.

We conclude with some caveats. First, Google Trends returns only integers without clarifying how it rounds decimals. Consequently, a value of zero for the search index might range from zero to values like 0.49, which could cause an error when rescaling normalized indices. Such an error may be large or small, depending on search terms and the type of analysis. We recommend researchers beware of this potential problem and analyze their data with caution.

Second, DMA-level Google searches lack enough data for those topics not widely recognized among the public. For example, national-level Google Trends shows that the American public's interest in Universal Basic Income increased during the coronavirus outbreak. Yet, DMA-level Google Trends omits many regions that do not generate sufficient searches.

With these caveats in mind, we hope that the guidance we provide helps researchers construct and analyze longitudinal Google Trends. We also hope that the present study promotes further discussions about the feasibility of internet search data as a measure of dynamic issue attention.

Notes

1. We focus on Google, among other search engines. Google fields about 92 percent of the world's online queries as of 2020 (<https://gs.statcounter.com/search-engine-market-share>).
2. By longitudinal data, we mean repeated cross-sectional data. Google Trends provides Google search data

across various regions. Longitudinal Google Trends means repeated cross-regional data of Google searches.

3. A media market is a geographic area where people receive the same television station offerings. We found that Google Trends media market data corresponded to Designated Market Areas as defined by Nielsen: <https://www.nielsen.com/us/en/intl-campaigns/dma-maps/>.
4. We use the open-source library Pytrends to collect Google Trends data. Our Python code is available in the online appendix.
5. For further details, refer to https://support.google.com/trends/answer/4365533?hl=en&visit_id=637320875923427946-2062678754&rd=1 and <https://medium.com/google-news-lab/what-is-google-trends-data-and-what-does-it-mean-b48f07342ee8>.
6. We recommend researchers select a time-series that does not have a value of zero throughout the whole period. We discuss the rationale behind this recommendation in Endnote 8.
7. When Google Trends returns a time-series with more than 269 data points, Google appears to change its time structure to limit the number of data points. For example, when Google Trends is asked to return a time-series for 269 days, it gives daily data. If the period becomes 270 days, it gives weekly data. Similarly, when the period spans 270 weeks, it returns monthly data instead of weekly data.
8. If the node has a value of zero, the scaling ratio is not computable. One solution is to find the first nonzero from the following weeks and make the second time window start from there. Another solution is to select a different reference DMA that does not have a value of zero.
9. In theory, the individual CDC, CNN, and Fox News time-series must produce the same rescaled data. In practice, however, differences could exist because of rounding errors.
10. Consistent with this explanation, monthly U.S. traffic data show that people's visit to partisan websites was stagnant or falling between February and March. Fox News saw only modestly increasing traffic, compared to other large outlets' big jumps. Retrieved from <https://www.nytimes.com/interactive/2020/04/07/technology/coronavirus-internet-use.html>.

11. The vote shares are retrieved from the New York Times website: <https://www.nytimes.com/elections/2016/results/president>.

References

- Althaus, Scott L and Todd C Trautman. 2008. “The impact of television market size on voter turnout in American elections.” *American Politics Research* 36(6):824–856.
- Ansolabehere, Stephen, Marc Meredith and Erik Snowberg. 2014. “Mecro-Economic Voting: Local Information and Micro-Perceptions of the Macro-Economy.” *Economics and Politics* 26(3):380–410.
- Bento, Ana I, Thuy Nguyen, Coady Wing, Felipe Lozano-Rojas, Yong-Yeol Ahn and Kosali Simon. 2020. “Evidence from internet search data shows information-seeking responses to news of local COVID-19 cases.” *Proceedings of the National Academy of Sciences* 117(21):11220–11222.
- Cramer, Katherine J. 2016. *The politics of resentment: Rural consciousness in Wisconsin and the rise of Scott Walker*. University of Chicago Press.
- Darr, Joshua P, Matthew P Hitt and Johanna L Dunaway. 2018. “Newspaper closures polarize voting behavior.” *Journal of Communication* 68(6):1007–1028.
- Dennison, James. 2019. “A review of public issue salience: Concepts, determinants and effects on voting.” *Political Studies Review* 17(4):436–446.
- Filla, Jackie and Martin Johnson. 2010. “Local news outlets and political participation.” *Urban Affairs Review* 45(5):679–692.
- Gilliam, Franklin D and Shanto Iyengar. 2000. “Prime suspects: The influence of local television news on the viewing public.” *American Journal of Political Science* 44(3):560–573.

- Japiec, Lilli, Frauke Kreuter, Marcus Berg, Paul Biemer, Paul Decker, Cliff Lampe, Julia Lane, Cathy O’Neil and Abe Usher. 2015. “Big data in survey research: AAPOR task force report.” *Public Opinion Quarterly* 79(4):839–880.
- Jones, Bryan D. 1994. *Reconceiving decision-making in democratic politics: Attention, choice, and public policy*. University of Chicago Press.
- Mellon, Jonathan. 2014. “Internet search data and issue salience: The properties of Google Trends as a measure of issue salience.” *Journal of Elections, Public Opinion & Parties* 24(1):45–72.
- Reeves, Andrew and James G Gimpel. 2012. “Ecologies of Unease: Geographic Context and National Economic Evaluations.” *Political Behavior* 34(3):507–534.
- Ripberger, Joseph T. 2011. “Capturing curiosity: Using internet search trends to measure public attentiveness.” *Policy Studies Journal* 39(2):239–259.
- Rodden, Jonathan A. 2019. *Why cities lose: The deep roots of the urban-rural political divide*. Hachette UK.
- Stephens-Davidowitz, Seth. 2017. *Everybody lies: Big data, new data, and what the internet can tell us about who we really are*. HarperCollins New York, NY.
- Tourangeau, Roger, Robert M Groves and Cleo D Redline. 2010. “Sensitive topics and reluctant respondents: Demonstrating a link between nonresponse bias and measurement error.” *Public Opinion Quarterly* 74(3):413–432.
- Wong, Cara J. 2007. ““Little” and “big” pictures in our heads: Race, local context, and innumeracy about racial groups in the United States.” *Public Opinion Quarterly* 71(3):392–412.

Supplementary Document for

“Building Longitudinal *Google Trends* to Measure Dynamic Local-Level Issue Attention”

1. Python and Jupyter Notebook code for the “jobs” index

Available at the following website with no identifying author information.

- Python code for cross-sectional indices, 31 December 2006 - 31 December 2016: <https://cutt.ly/7hY7H6Q>
- Python code for the reference time-series, 31 December 2006 - 31 December 2016: <https://cutt.ly/8h0Iu2M>
- Jupyter Notebook code for rescaling the cross-sectional indices (Step 2 and Step 3): <https://cutt.ly/JhY7NxI>

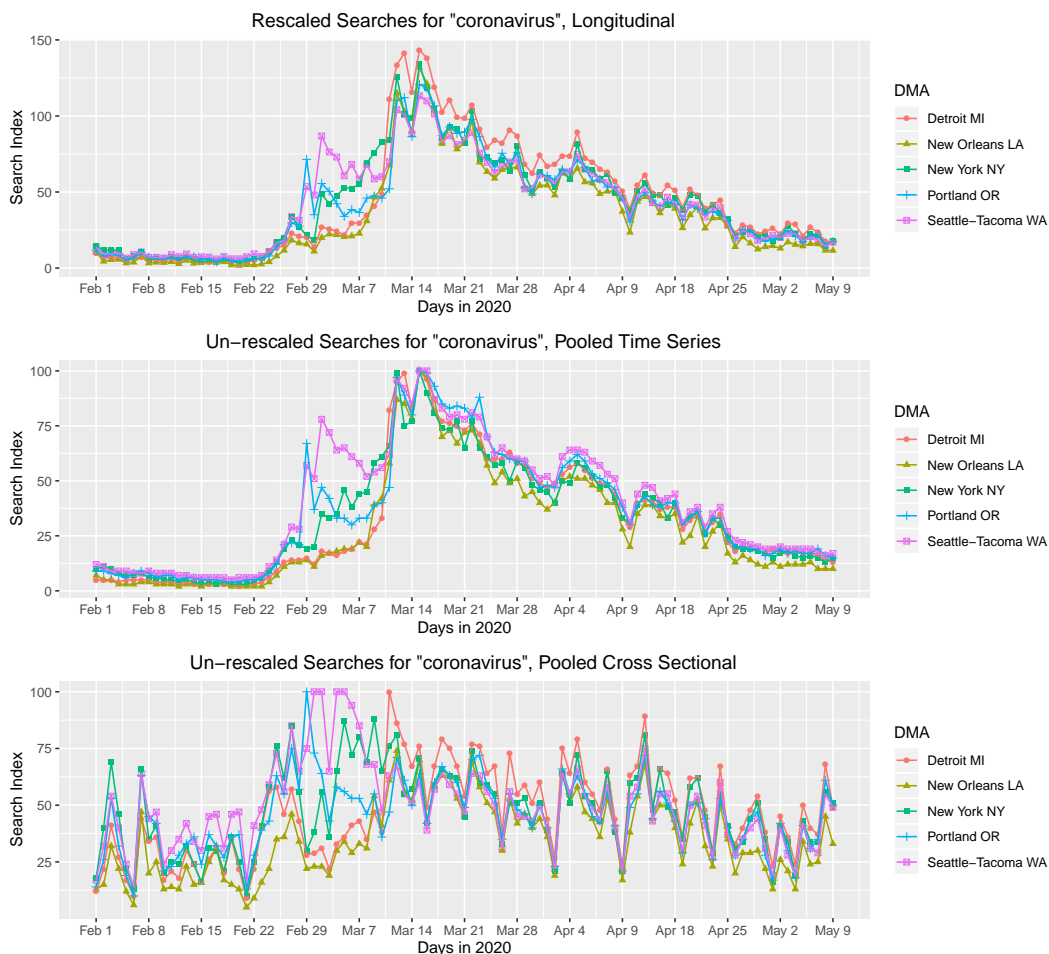
2. Python and Jupyter Notebook code for the “CDC-CNN-Fox News” index

Available at the following website with no identifying author information.

- Python code for cross-sectional indices, 1 February 2020 - 9 May 2020: <https://cutt.ly/mhUhrXI>
- Python code for the reference time-series, 1 February 2020 - 9 May 2020: <https://cutt.ly/0hUhbZu>
- Jupyter Notebook code for rescaling the cross-sectional indices (Step 2 and Step 3): <https://cutt.ly/VhUhUQj>

3. Dynamic public attention to the coronavirus pandemic

Figure 3: Comparing Rescaled Longitudinal Searches, Pooled Time-Series Searches, and Pooled Cross-Sectional Searches for “coronavirus”: Examples of Daily Google Searches from Five DMAs



We use daily searches for “coronavirus” over the period between February 1 - May 9, 2020 to track how public attention to the pandemic varies over time in different DMAs. We use the Atlanta time-series data as our reference to rescale the daily searches. The time period in our analysis spans 99 days. Google Trends returns a daily time-series for the entire period, and the divide-and-concatenate process is unnecessary.

The top plot in Figure 3 presents five examples of longitudinal daily searches. It appears that the American public had been relatively less concerned about the coronavirus until

late February when Italy started seeing significant increases in confirmed cases. Since then, the coronavirus had been a more salient issue in Seattle-Tacoma and Portland until early March than any other DMAs. These results make sense because Seattle was the first U.S. epicenter of COVID-19. As New York City and New Orleans emerged as next epicenters of the pandemic, the issue became more salient in those regions than Seattle-Tacoma and Portland in mid-March. At the same time, Detroit started seeing higher search volumes than New York. Though we do not present here, we also found that the trends of the Lansing and Grand Rapids-Kalamazoo-Battle Creek DMAs were similar to Detroit's. Perhaps, the coronavirus outbreak was particularly salient to Michigan residents. Consistent with this inference, a survey using geotagged twitter data shows that Michigan residents complained about depression and anxiety during the virus crisis more than people in any other state.¹²

These dynamic rises and falls of public attention to COVID-19 in different DMAs are not accurately captured by pooling time-series searches across DMAs or pooling cross-sectional searches across time. As the middle plot in Figure 3 displays, the pooled time-series searches provide a trend that resembles the longitudinal data. Yet, in each time point, the searches fail to deliver the correct ranks among different DMAs. For example, the pooled time-series data do not detect Detroit's surge at the peak and afterward. In the bottom plot, while the pooled cross-sectional searches present the correct ranks among different DMAs in each day, they fail to capture the accurate trend over time.