



2019

Reference-Dependent Preferences Among NFL Fans: Evidence from Google Trends

Sunjae Lee

Washington University in St. Louis, sunjae.lee@wustl.edu

Follow this and additional works at: <https://digitalcommons.iwu.edu/uer>



Part of the [Behavioral Economics Commons](#)

Recommended Citation

Lee, Sunjae (2019) "Reference-Dependent Preferences Among NFL Fans: Evidence from Google Trends," *Undergraduate Economic Review*: Vol. 16 : Iss. 1 , Article 10.

Available at: <https://digitalcommons.iwu.edu/uer/vol16/iss1/10>

This Article is protected by copyright and/or related rights. It has been brought to you by Digital Commons @ IWU with permission from the rights-holder(s). You are free to use this material in any way that is permitted by the copyright and related rights legislation that applies to your use. For other uses you need to obtain permission from the rights-holder(s) directly, unless additional rights are indicated by a Creative Commons license in the record and/ or on the work itself. This material has been accepted for inclusion by faculty at Illinois Wesleyan University. For more information, please contact digitalcommons@iwu.edu.

©Copyright is owned by the author of this document.

Reference-Dependent Preferences Among NFL Fans: Evidence from Google Trends

Abstract

I look for evidence of reference-dependent preferences in the National Football League (NFL). Under reference-dependent preferences, sports fans should react more strongly to surprising wins and losses than expected wins or losses. I use Google Trends to look at the impact of NFL game outcomes on the use of positive or negative words on Google search. While search activity did respond to NFL games, I did not find that this response was sensitive to how surprising the outcome was, and so did not find evidence of reference-dependent preferences.

Keywords

Reference-Dependent Preferences, Google Trends, Behavioral Economics

Cover Page Footnote

Acknowledgements: I thank Prof. Nick Huntington-Klein for guidance and editing assistance.

Introduction

Sports fans tend to form expectations about how games will proceed and who is likely to win. Sports fans also speculate about the result of the game based on information that is available to them from sports pundits or online betting sites. However, life is unpredictable, and so are the sports game results. There could be a case in which the actual result of a game is utterly different from what people have expected. If people expected a team to win the game by ten points, but in reality the team loses by ten points, people will be shocked and enraged.

In this paper, I examine the effects of sports game results that completely go against people's "reference points on the game outcome" (Ge 2018) on people's behaviors. To be more exact, this paper looks for evidence of reference-dependent preferences, a behavioral economics concept that argues that people have "reference points" that they use to assess outcomes. People do not have preferences purely over outcomes, as the rational decision model would suggest, but evaluate outcomes relative to reference points based on what they expect.

In order to see how people's behaviors change in response to surprising sports game outcomes, I connect a data set of anticipated and actual sports outcomes, specifically from the National Football League (NFL), to a data set from Google Trends in which sports fans have an opportunity to express their frustration and loss. Anticipated game outcomes come from an online betting site, FootballLocks.com, and are linked to actual game scores from the NFL. Differences between anticipated and actual scores constitute surprises in the outcome. Negative surprises should be especially frustrating. One place people might express their frustration is on Google Search, which is not just used to search for information, but also to express hidden frustrations and opinions that may not be reported on a survey (Stephens-Davidowitz 2014).

Before examining the data, I hypothesize that an unexpected loss will result in an increase of searches of negative terms, while an unexpected win will result in an increase of searches of positive terms on Google, above and beyond what would occur from an expected loss or win. This would be evidence of reference-dependent preferences among sports fans.

This is not the first paper to look for evidence of reference-dependent preferences in sports. In the literature section of this paper, I look at four different papers on reference-dependent preferences. All of these papers find evidence of reference-dependent preferences, either among athletes or fans.

However, after completing the examination of the data collected through the Google Trends, I have found, counter to much of the rest of the published literature, no evidence of reference-dependent preferences. The results imply that football game results do have a subtle effect on Google search terms, this effect is not consistent with what reference-dependent preferences would predict.

This study is significant in a sense that it contributes to the understanding of behavioral economics, and implies that reference-dependent preferences may not be as well-supported as the prior literature would suggest. The result of the study suggests a pivotal point for further studies and applications of the reference-dependent preferences by questioning the real effects of football game results on football fans' behavior.

Literature

The primary purpose of this paper is to look for reference-dependent preferences using Google Trends, a weekly index of the volume of searches for a particular term on Google. Google Trends is an increasingly popular method for collecting data for research. In Choi & Varian (Choi and Varian 2012), the researchers used Google Trends to “forecast near-term values of economic indicators” such as “automobile sales, unemployment claims, travel destination planning, and consumer confidence” (Choi & Varian 1). They believe Google Trends can be useful in predicting near future and present economic situations. They compared the performance in prediction of an economic model with Google Trends and without Google Trends. They confirmed that “simple seasonal AR models that include relevant Google Trends variables tend to outperform models that exclude these predictors by 5% to 20%” (Choi & Varian 8).

Perhaps the most well-known application of Google Trends data in economic research is “The Cost of Racial Animus on a Black Candidate: Evidence using Google Search Data” (Stephens-Davidowitz, 2014). The author uses Google Trends to “understand the extent of contemporary prejudice” and “increase our understanding of the determinants of voting” (Stephens-Davidowitz 1). He argues that Google Trends is a “new proxy for an area’s racial animus from a non-survey source: the percent of Google search queries that include racially charged language” because of “individuals’ tendency to withhold socially unacceptable attitudes, such as negative feelings towards blacks, from surveys” (Stephens-Davidowitz 1). He further claims that “Google search query data can do more than correlate with existing measure; on socially sensitive topics; they can give better data and open new research on old questions” (Stephens-

Davidowitz 2). Unlike surveys, Google search renders more accurate data by allowing people to be much more honest about their desires and controversial opinions. Stephens-Davidowitz shows that users don't simply use Google to search for information, but also to express emotions or frustrations, which makes the data valuable as a high-frequency measure of things like frustration or anger, implying that it can be used to measure these emotions in response to unexpected sports losses under reference-dependent preferences.

Reference-dependent preferences refers to the behavioral economics concept that people have "reference points" and evaluate outcomes relative to those reference points. This means that the way someone feels about an outcome is relative to what they would have expected it to be. Finding evidence of reference-dependent outcomes can be difficult because the reference points cannot be observed. The literature on reference-dependent preferences often uses sports as a setting where winning probabilities can be calculated ahead of time, and so unexpected losses and wins can be easily identified. I highlight several studies of reference-dependent preferences in sports below.

In Bartling, Brandes, & Schunk (2015), the researchers show that "professional soccer players exhibit reference-dependent behavior during matches" (Bartling et al. 1). They used data from two soccer leagues to show evidence that players had reference-dependent preferences. When the flow of the match did not coincide with players' expectations (reference points), especially when their team was losing unexpectedly, the probability that a player would receive a red card in a given minute increased by more than 20 percent. The same pattern did not appear when the team was losing but was expected to lose, so this can be identified as reference-dependent behavior. Reference-dependent behavior was not diminished by player experience or high-stakes games.

Subsequently Pope & Schweitzer (2011) use golfer performance on the PGA tour to test for loss aversion, a feature of reference-dependent preferences. Like Bartling et al. (2015), this paper also concludes that "loss aversion, a fundamental bias, continues to persist in a highly competitive market" (Pope & Schweitzer 155), and is not eliminated by competition, large stakes, or experience.

In addition to the reference-dependent preferences shown in players' behaviors, there is also research on the effect of surprising sports game losses on the audience. Card & Dahl (2011) find an effect of unexpected wins and losses by professional football teams on family violence. They analyze police reports of

violent incidents on Sundays during football season, and find that “upset losses lead to a 10% increase in the rate of at-home violence by men against their wives and girlfriends” which contrasts with the fact that “losses when the game was expected to be close have small and insignificant effects” (Card et al. 103). Similar to my research, Card and Dahl also gathered information about reference points through the NFL betting market. Their finding not only confirms evidence of loss aversion and reference-dependent preferences, but also that reference points among football fans are formed rationally and match betting odds.

Lastly, Ge (2018) conducts research on reference-dependent preferences by analyzing the relationship between sports outcomes and passengers’ tipping behavior. Ge argues that social norms and consumer sentiment are two main factors that determine consumers’ tipping behavior. Ge uses data on New York City taxi fares, tipping, and trip information to show that passengers tend to tip more when a sports team unexpectedly wins, or win by greater score difference than the expectation, but they do not pay less tip when there is an upset loss. Ge explains the absence of loss aversion with the effect of social norms on people. As a result, Ge’s study “[demonstrates] that while consumers’ reactions can still be reference-dependent, loss averse behavior may possibly be muted in light of social norms” (Ge 5).

Evidence from a number of studies finds evidence of reference-dependent preferences in sports, both among players and fans. Standard approaches include comparing sports game outcomes to expectations drawn from outside data like betting markets, and then linking that data to outcomes collected elsewhere, like from police reports. My study will use Google Trends data, which has been shown to provide measures of animus and anger in previous work, as an outcome measure.

Data

The project as a whole performs Google Trends searches on a dictionary of words, which are coded as positive or negative, and links them to data on football games. I took a dictionary of 8,223 words available on the University of Pittsburgh’s Sentiment Lexicon website.¹ The dictionary is from work by Wiebe, Wilson, and Hoffmann (2005) and Riloff and Wiebe (2003) and contains information on whether each word is “positive” or “negative” and the strength of that polarity. For example, “abhorrent” is strongly negative, and “civility” is

¹ <http://mpqa.cs.pitt.edu/>

strongly positive. I kept only words that were either strongly negative or strongly positive, removed duplicates, and only used the “non-stemmed” versions of the words (for example, “angr” might be a stem for both “angry” and “angriest”). This results in a dictionary of 3,408 words, 1,108 of which are positive and 2,301 of which are negative. I then used the *gtrendsR* package in R (Massicotte and Eddelbuettel 2018) to perform Google Trends searches on each of the words in the dictionary.

I performed the search separately by state, gathering weekly Google Trends results from January 2015 to December 2018. Google Trends reports an index score that shows the popularity of that word in that state and how it changes over time. The score has no absolute meaning, but can be compared to itself and so provide information on whether a search has gotten more or less popular over time. Google Trends has previously been used as a measure of sentiment (Stephens-Davidowitz, 2014). Importantly, Stephens-Davidowitz (2014) finds that people use Google searches to express frustration, and so searches might be a way to pick up frustration from sports losses. Choi and Varian (2012) emphasize that Google Trends can provide relatively accurate predictions of near future and present economic situations compared to existing surveys because Google Trends eliminate the effects of the self-serving bias survey participants tend to exhibit.

I gather data on the point spread for games from 2015-2018 from FootballLocks.com, a football betting site. The spread reports the expected number of points by which a team will win or lose. I then link the point spread data to information on the actual score of each game, and the day it was played, from the NFL website, gathered by the *nflscrapR* package (Horowitz, Yurko, and Ventura 2019).

I identify the state that each team plays in by hand, and then merge the data on football spreads, scores, and dates with the Google Trends data. I assign Google Trends scores collected on a week that starts 6 or 7 days before game day as “before game” data, and Google Trends scores collected on a week that starts 0 or 1 days after game day as “after game” data. To avoid overlaps where the same week of Trends data is “before” one game but “after” another, I ignore the impact of a game on searches in a state if that same state also played a game the week previous or the week following. I also drop games in which both teams are from the same state. This results in 108 games examined, one of which could not be linked to betting spread data.

Table 1: Summary Statistics

Statistic	N	Mean	St. Dev.	Min	Pctl(25)	Pctl(75)
Surprise	107	-8.332	8.571	-33.500	-13.000	-2.000
Expected Win	107	-1.070	5.671	-16.000	-5.500	3.000
Actual Win	108	-9.370	7.974	-37	-13	-3

Table 1 shows the expected and actual score for each game in the data, from the point of view of the team that lost.² The table shows the spread (Expected Win), the actual spread (Actual Win, always negative since this table shows the losing side), and the amount of Surprise, the difference between Actual Win and Expected Win. Each observation is linked to Google Trends data for 3,408 words, both the week before and the week after the game. After dropping results for word/state combinations with too few searches to produce results, the final data set contains 1,360,956 observations, with observations uniquely identified by the combination of game, state, word, and week-before-game/week-after-game.

Methods

The primary results perform separate analyses for winning and losing teams, since the reference-dependent preferences framework suggests the results should be different for each. In each case, I regress the Google Trends score on:

- Whether the word is positive or negative
- Whether the score is collected before the game or after
- The amount of “surprise” from the results of the game
- And the interactions of all three variables

This gives us a regression equation:

$$\begin{aligned}
 Trends_{itgs} = & \beta_0 + \beta_1 Positive_{it} + \beta_2 After_t + \beta_3 Surprise_{gs} \\
 & + \beta_4 Positive_i After_t + \beta_5 Positive_i Surprise_{gs} \\
 & + \beta_6 After_t Surprise_{gs} + \beta_7 Positive_i After_t Surprise_{gs} + \varepsilon_{itgs}
 \end{aligned}$$

for word i at time t for game g with a team from state s , with only one team per game in the regression because winning and losing teams are estimated

² All tables were prepared using the stargazer package (Hlavac 2018).

separately. I also present results using state fixed effects, replacing β_0 with β_s , to account for the possibility that some states tend to see larger surprises more often and to potentially improve the precision of estimates. Regressions use robust standard errors.

I am interested mostly in the effects of surprise on how the popularity of words changes from before the game to after, and in particular on how those effects differ by polarity of the words.

The coefficients of interest are: for losing teams, I focus mostly on β_6 , which shows how surprise affects the before-to-after change in word popularity specifically among negative words (for which $Positive_i = 0$). If β_6 is negative, that means that negative words see larger increases in popularity after a particularly negatively surprising loss, supporting the reference-dependent preferences theory. I also look at β_7 , which shows the difference in how surprise affects the before-to-after change in word popularity between positive and negative words. If β_7 is zero, then both positive and negative words respond the same way to surprising games, contrary to the reference-dependent preferences theory.

For winning teams, I am interested in β_7 . If β_7 is positive, then particularly positively-surprising games lead to bigger increases in the popularity of positive words, consistent with reference-dependent preferences.

I run the analysis in a second way. The first regression uses the Google Trends score directly, but this may give us problems because the scale of the Google Trends score can't really be interpreted and may not be comparable across words and states.

So, I create an indicator variable $Increase_{igs}$ equal to 1 if the Google Trends score for word i increased in state s from before game g to after, and equal to 0 if it decreased. Words that stayed at the exact same Google Trends score are dropped.

I then run the analysis

$$Increase_{igs} = \beta_0 + \beta_1 Positive_i + \beta_2 Surprise_{gs} + \beta_3 Positive_i Surprise_{gs} + \varepsilon_{igs}$$

using a linear probability model to easily allow for state fixed effects.

For losing teams, I am interested in β_2 . If β_2 is negative, then big negative surprises make negative words more popular. I am also interested in β_3 . If β_3 is

zero, then positive and negative words change in the same way in response to surprise. Regressions use robust standard errors.

For winning teams, I am interested in β_3 . A positive β_3 shows that larger positive surprises improve the popularity of positive words more than negative.

Results

Results Section A: Main Results

Table 2: Main Regression Results

	<i>Dependent variable:</i>			
	Raw Google Index			
	Winning Teams	Losing Teams	Winning Teams (State FE)	Losing Teams (State FE)
	(1)	(2)	(3)	(4)
After the game	0.089 [*] (0.054)	0.070 (0.060)	0.089 [*] (0.053)	0.070 (0.060)
Positive	3.137 ^{***} (0.071)	3.284 ^{***} (0.080)	3.129 ^{***} (0.071)	3.269 ^{***} (0.079)
Surprise	0.011 ^{***} (0.003)	0.0003 (0.004)	-0.007 ^{**} (0.003)	0.009 ^{**} (0.004)
After the game*Positive	0.150 (0.101)	0.049 (0.113)	0.150 (0.100)	0.049 (0.112)
After the game*Surprise	0.003 (0.004)	-0.003 (0.005)	0.003 (0.004)	-0.003 (0.005)
Positive*Surprise	0.015 ^{***}	-0.002	0.014 ^{**}	-0.002

	(0.006)	(0.007)	(0.006)	(0.007)
After the game*	-0.009	-0.003	-0.009	-0.003
Positive*Surprise	(0.008)	(0.009)	(0.008)	(0.009)
Constant	10.642***	10.920***	9.107***	9.121***
	(0.038)	(0.043)	(0.074)	(0.090)
Observations	706,246	654,710	706,246	654,710
R ²	0.012	0.011	0.023	0.022
Adjusted R ²	0.011	0.011	0.023	0.022
<i>Note:</i>			*p<0.1; **p<0.05; ***p<0.01	

Table 2 shows the results of the regressions described in the Methods section, run separately for winning (Columns 1 and 3) and losing (Columns 2 and 4) teams, both without (Columns 1 and 2) and with state fixed effects (Columns 3 and 4). State fixed effects are included to account for potential state-level differences in Google search activity.

The coefficient on Positive is positive, indicating that on average positively-coded words have higher indices. However, since the trends score is not necessarily meant to be comparable across words, this is not a result of interest.

The coefficients on Surprise and Positive*Surprise are both often significant, which is interesting because it suggests that more-surprising games are related to more popular searches, especially for positive words, but both before and after the actual game. This may have something to do with excitement, but that is a speculative interpretation.

The coefficient on Surprise*After is insignificant for losing teams. This is β_6 from our regression equation. The lack of significance here indicates that more surprising losses are not related to increasing popularity of negative words. This result fails to support reference-dependent preferences.

There is also no significance on Surprise*After*Positive, β_7 , for either winning or losing teams. The lack of significance here indicates that there is no

difference between Positive and Negative words in how surprising outcomes affect popularity. This again fails to support reference-dependent preferences

State fixed effects do not change much, which is not too surprising as the Google Trends scores are within-state.

Table 3: Before-After Increase Regression Results

	<i>Dependent variable:</i>			
	Increase			
	Winning Teams	Losing Teams	Winning Teams (State FE)	Losing Teams (State FE)
	(1)	(2)	(3)	(4)
Positive	0.002 (0.002)	0.004* (0.002)	0.002 (0.002)	0.004* (0.002)
Surprise	-0.0001 (0.0001)	0.0001 (0.0001)	0.0002** (0.0001)	0.001*** (0.0001)
Positive*Surprise	-0.0001 (0.0002)	-0.0002 (0.0002)	-0.0001 (0.0002)	-0.0002 (0.0002)
Constant	0.505*** (0.001)	0.501*** (0.001)	0.501*** (0.003)	0.492*** (0.004)
Observations	353,123	327,355	353,123	327,355
R ²	0.00001	0.00004	0.001	0.002
Adjusted R ²	0.00000	0.00003	0.001	0.002

Note:

*p<0.1; **p<0.05; ***p<0.01

Table 3 shows regression results as described in the Methods section, where the dependent variable is a binary indicator equal to 1 if the word increased

in popularity from before the game to after the game. The dependent variable is equal to 0 if the word decreased in popularity. Words with no change are dropped.

When I include state fixed effects, precision increases and Surprise is now significant. However, both values are very tiny and positive. This result implies that, for winning teams, more-positive (better) surprises make negative words more popular, counter to what is expected, since reference-dependent preferences implies Surprise should improve popularity of positive words for winning teams, not negative. Moreover, for losing teams, more-positive (i.e., less-negative, better) surprises increase the popularity of negative words, when I would expect those smaller surprises to have less of an effect.

Also, unlike what I have expected, there was no effect of Positive*Surprise, which indicates that any impact of Surprise affects positive and negative equally.

Results Section B: Robustness Checks

There are two concerns I have about our results: one is a possibility that there is no response not because I am failing to replicate reference-dependent preferences, but rather because the results of sports games have no effect on the Google Searches at all. The other is that there is no effect because I am using a linear term for Surprise. To check the first concern, I bring in the Table 4, which pools together both winning and losing teams and repeats the analysis from Tables 2 and 3.

Table 4: Pooled Analysis

	<i>Dependent variable:</i>			
	Google Trends Index		Increase	
	Raw Score	Raw Score (State FE)	Increase	Increase (State FE)
	(1)	(2)	(3)	(4)
After the game	0.090**	0.090**		
	(0.042)	(0.042)		

Positive	3.295*** (0.056)	3.283*** (0.056)	0.007*** (0.002)	0.007*** (0.002)
Won	-0.204*** (0.041)	-0.171*** (0.042)	0.006*** (0.002)	0.001 (0.002)
After the game*Positive	0.072 (0.080)	0.072 (0.079)		
After the game*Won	0.027 (0.059)	0.027 (0.058)		
Positive*Won	-0.038 (0.078)	-0.038 (0.077)	-0.006** (0.003)	-0.006** (0.003)
After the game*Positive*Won	0.003 (0.110)	0.003 (0.110)		
Constant	10.938*** (0.030)	9.154*** (0.058)	0.501*** (0.001)	0.495*** (0.004)
Observations	1,366,966	1,366,966	462,692	462,692
R ²	0.011	0.023	0.00004	0.001
Adjusted R ²	0.011	0.023	0.00003	0.001
<i>Note:</i>		*p<0.1; **p<0.05; ***p<0.01		

The first two columns use raw Google Trends scores as the dependent variable. Here, search activity for both kinds of words increases after a game relative to before. Winning the game is related to search behavior, but none of the interactions are significant. These results imply that sports results have influences on the search activity, even though the previous section showed that the

relationship does not appear to be consistent with reference-dependent preferences.

The second two columns use the Increase from before variable as the binary dependent variable. They show that positive words are about .7% more likely than negative words to increase in popularity from before the game to after. Moreover, one major discovery is that winning affects positive and negative words differently, with positive words .6% less likely to increase after the game than negative words.

These results imply that Google searches do respond in some way to football results, but not in the way that reference-dependent preferences would expect. These effects are tiny but nonzero. Connecting this result to another literature on the reference-dependent preferences, Ge (2018) conducts a research on reference-dependent preferences by analyzing the relationship between sports outcomes and passengers' tipping behavior. Ge argues that there are two main factors that determine consumers' tipping behavior: "social norms and consumer sentiment" (Ge 3). Data used in the study came from a dataset of the New York City Taxi and Limousine Commission which contains "fare, tipping and trip information for taxi rides in New York City" (Ge 3). Ge found out that passengers tend to tip more when a sports team unexpectedly wins, or win by greater score difference than the expectation, but they do not pay less tip when there is an upset loss. Ge explains the absence of loss aversion with the effect of social norms on people. As a result, Ge's study "[demonstrates] that while consumers' reactions can still be reference-dependent, loss averse behavior may possibly be muted in light of social norms" (Ge, p. 5).

To check the second concern about the potential nonlinearity of the effect of Surprise, I check the Increase for each word across each value of Surprise, non-parametrically, in four graphs. I have a separate graph for positive and negative words, and for winning and losing teams.

Figures: Nonlinear Effects of Surprise

Figure 1: Positive Words for Losing Teams




Figure 2: Negative Words for Losing Teams




Figure 3: Positive Words for Winning Teams





Figure 4: Negative Words for Winning Teams



In all four graphs, the LOESS curve is never significantly different from the overall mean of .5, indicating that Surprise has no real relationship with Increase for either positive or negative words, linear or otherwise. As a result, I can conclude that our null result is not simply because of linearity.

Conclusion

The primary goal of this research paper is to find out whether the surprising results of football games affect the terms people search on Google. I was mainly interested in the effects of surprise on how the popularity of words changes from before the game to after, and how that change is related to whether those words are positive or negative. In order to do so, I utilized Google Trends data. Before the analysis, I predicted that an unexpected loss would result in increased searches for negative words, and an unexpected win would result in increased searches for positive words.

Contrary to my expectation, I found no evidence in favor of reference-dependent preferences. The results implied that Google searches do respond in some ways to football results, but not in the way that reference-dependent preferences would expect. These effects were tiny but nonzero.

This finding that the effects of football game results are insignificant in terms of influencing football fans' behavior provides a crucial point to consider in further application of the reference-dependent preferences.

A number of studies, which include studies I have discussed in the literature section, suggest that both players and fans behave according to reference-dependent preferences in sports games. However, my paper has discovered a potential weakness to these findings by demonstrating a failure to replicate, questioning the credibility of studies on the reference-dependent preferences. Consequently, governments and institutions should consider carefully before taking reference-dependent preferences into account when making decisions or establishing policies.

Bibliography

- Bartling, Björn, Leif Brandes, and Daniel Schunk. 2015. "Expectations as Reference Points: Field Evidence from Professional Soccer." *Management Science* 61 (11): 2646–61. <https://doi.org/10.1287/mnsc.2014.2048>.
- Card, David, and Gordon B. Dahl. 2011. "Family Violence and Football: The Effect of Unexpected Emotional Cues on Violent Behavior*." *The Quarterly Journal of Economics* 126 (1): 103–43. <https://doi.org/10.1093/qje/qjr001>.

- Choi, Hyunyoung, and Hal Varian. 2012. "Predicting the Present with Google Trends." *Economic Record* 88 (SUPPL.1): 2–9.
<https://doi.org/10.1111/j.1475-4932.2012.00809.x>.
- Ge, Qi. 2018. "Sports Sentiment and Tipping Behavior." *Journal of Economic Behavior & Organization* 145 (January): 95–113.
<https://doi.org/10.1016/j.jebo.2017.10.016>.
- Hlavac, Marek. 2018. *Stargazer: Well-Formatted Regression and Summary Statistics Tables*. (version 5.2.1). <https://CRAN.R-project.org/package=stargazer>.
- Horowitz, Maksim, Ron Yurko, and Samuel Ventura. 2019. *NflscrapR: Compiling the NFL Play-by-Play API for Easy Use in R* (version 1.8.1). <https://github.com/maksimhorowitz/nflscrapR>.
- Massicotte, Philippe, and Dirk Eddelbuettel. 2018. *GtrendsR: Perform and Display Google Trends Queries* (version 1.4.2). <https://CRAN.R-project.org/package=gtrendsR>.
- Pope, Devin G, and Maurice E Schweitzer. 2011. "Is Tiger Woods Loss Averse? Persistent Bias in the Face of Experience, Competition, and High Stakes." *American Economic Review* 101 (1): 129–57.
<https://doi.org/10.1257/aer.101.1.129>.
- Riloff, Ellen, and Janyce Wiebe. 2003. "Learning Extraction Patterns for Subjective Expressions." In *Proceedings of the 2003 Conference on Empirical Methods in Natural Language Processing* -, 10:105–12. Not Known: Association for Computational Linguistics.
<https://doi.org/10.3115/1119355.1119369>.
- Stephens-Davidowitz, Seth. 2014. "The Cost of Racial Animus on a Black Candidate: Evidence Using Google Search Data." *Journal of Public Economics* 118: 26–40. <https://doi.org/10.1016/j.jpubeco.2014.04.010>.
- Wiebe, Janyce, Theresa Wilson, and Claire Cardie. 2005. "Annotating Expressions of Opinions and Emotions in Language." *Language Resources and Evaluation* 39 (2): 165–210.
<https://doi.org/10.1007/s10579-005-7880-9>.