



Electronic Cigarettes and Twitter: Sentiments, Categorization, and Hedonometrics

Eric M. Clark^{1,2,3,4,5}, Chris Jones^{5,8,9}, Diann Gaalema^{6,7,8}, Ryan Redner^{6,8}, Thomas J. White^{6,8}, Allison Kurti^{6,8}, Andrew Schneider⁸, Marion Couch⁵, Peter Dodds^{1,2,3,4}, and Chris Danforth^{1,2,3,4}

Computational Story Lab¹, Department of Mathematics & Statistics², Vermont Complex Systems Center³, Vermont Advanced Computing Core⁴, Department of Surgery⁵, Department of Psychiatry⁶, Department of Psychology⁷, Vermont Center on Behavior and Health⁸, & Global Health Economics Unit of the Vermont Center for Clinical and Translational Science⁹


Abstract

Electronic cigarettes, or e-cigs for short, have become a popular alternative to traditional tobacco products. The vaporization technology present in e-cigarettes allows consumers to simulate tobacco smoking without igniting the carcinogens found in tobacco. The health risks, marketing regulations, and the potential of these devices as a form of nicotine replacement therapy are hotly debated both politically and clinically. Twitter, a mainstream social media outlet, provides a means to survey the popularity and sentiment of consumer opinions regarding e-cigarettes. Approximately 700,000 tweets containing mentions of e-cigarettes were collected from a 10% sample of Twitter spanning from January 2012 to July 2014.

All tweets mentioning e-cigarettes were categorized as Commercial, Infermorial, or Organic. Tweets in the commercial category ($\approx 70\%$) contained atleast 3 marketing key words (e.g. 'free trial', 'buy', 'coupon', 'starter kit', ...) , a key word along with a URL, or are from SPAM accounts. The Infermorial category ($\approx 15\%$) contains all tweets with URLs that omit these key words. The Organic category ($\approx 15\%$) contains the remaining tweets.

The emotionally charged words that contribute to the positivity of various subsets of tweets from each category are quantitatively measured, a hedonometrics. Outliers in both the positivity and frequency time-series distributions correspond to political debates regarding the regulation of e-cigarettes. Time-series analysis techniques are implemented to determine the effect promotional tweets have on organic sentiments. Due to the high youth presence on twitter as well as the clinical uncertainty regarding the risks associated with e- cigarettes, understanding the effect of promotionally marketing vaporization products across social media is relevant to public health agendas.

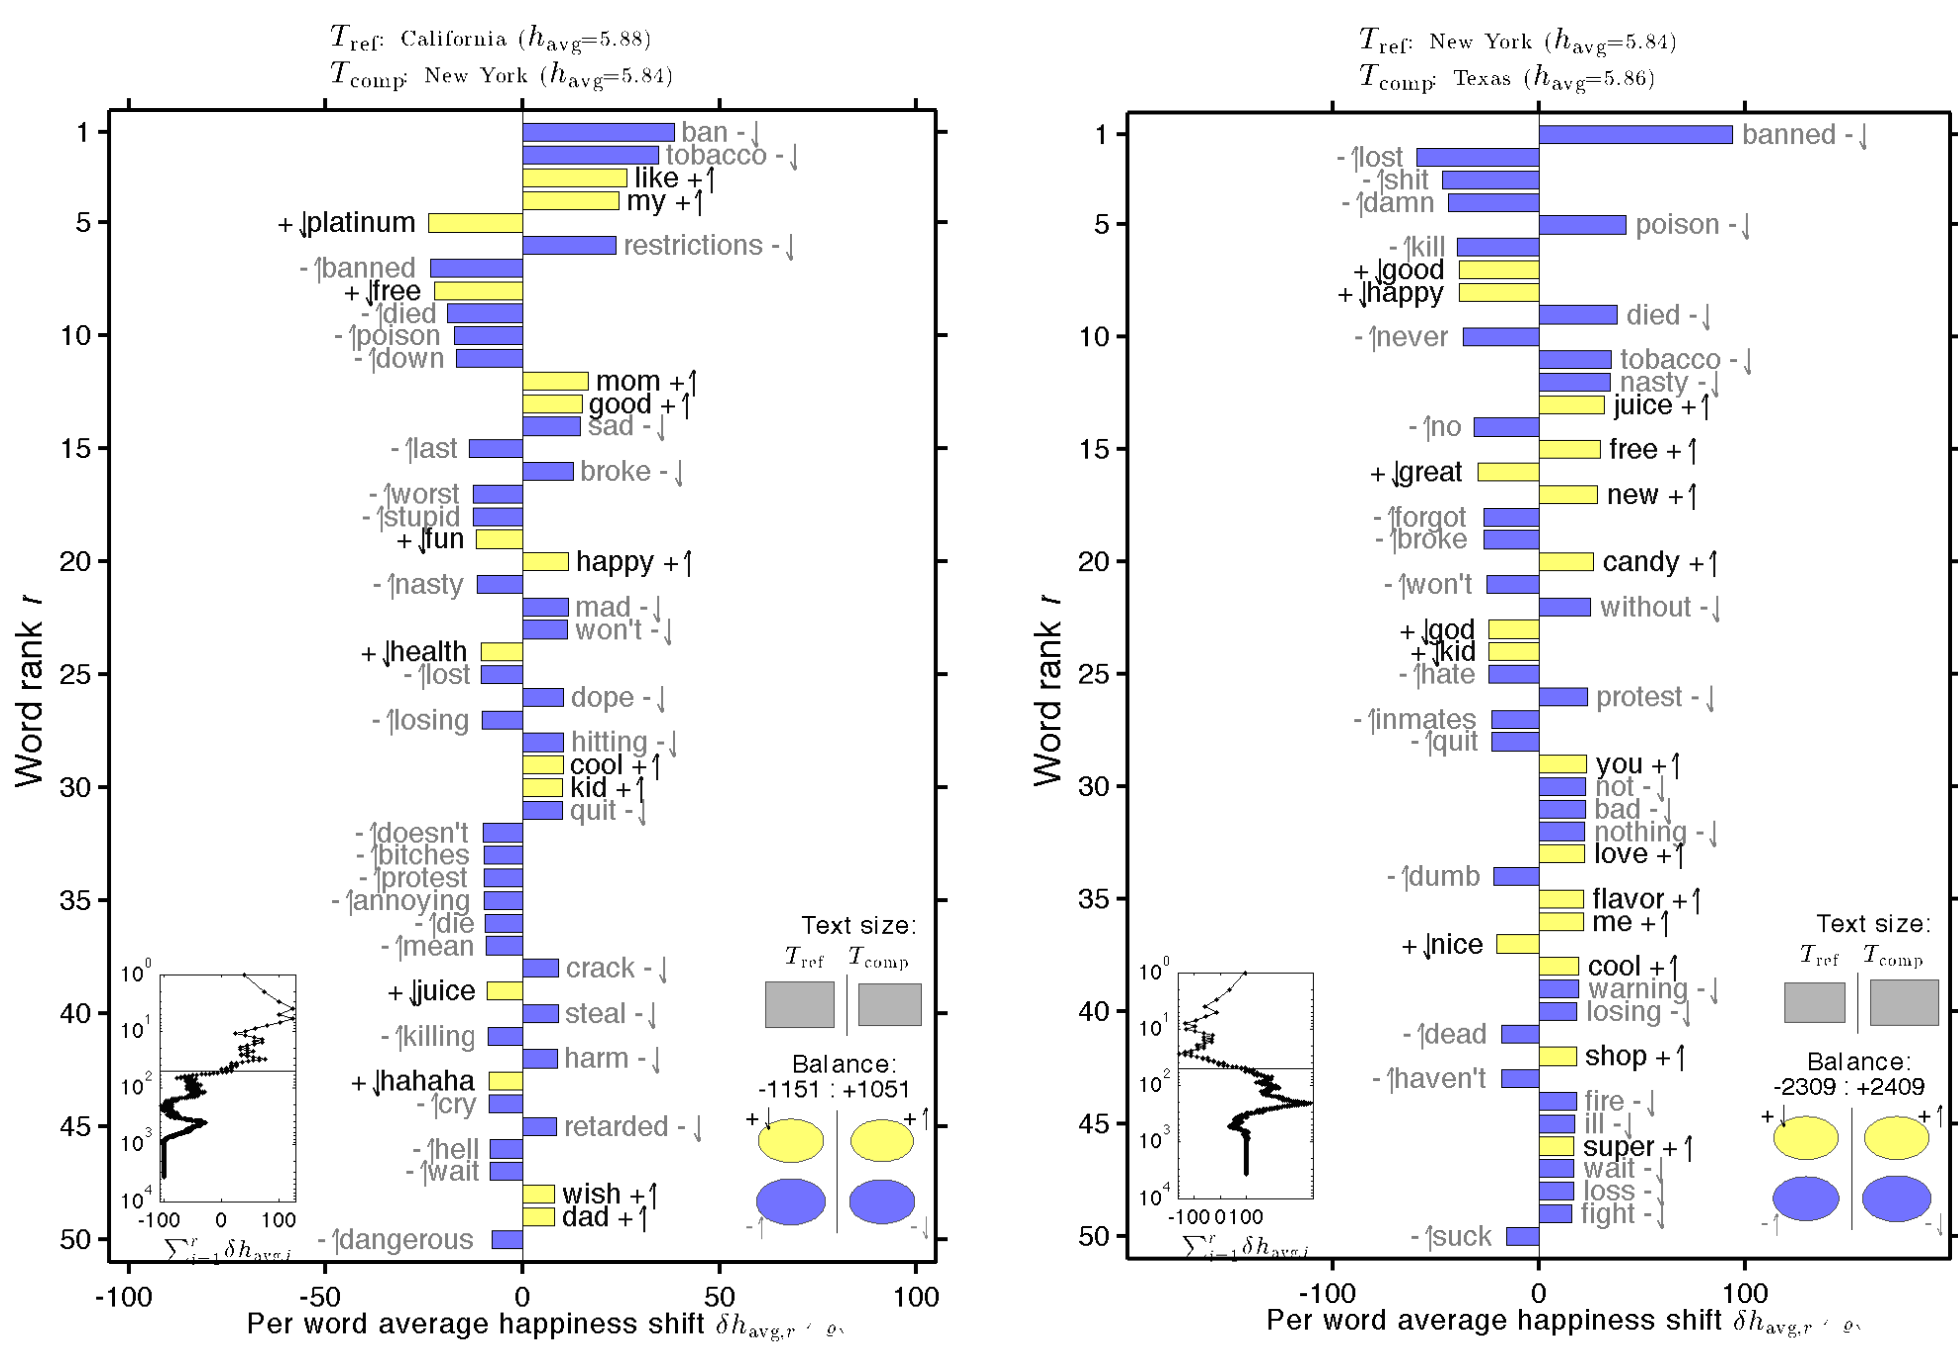
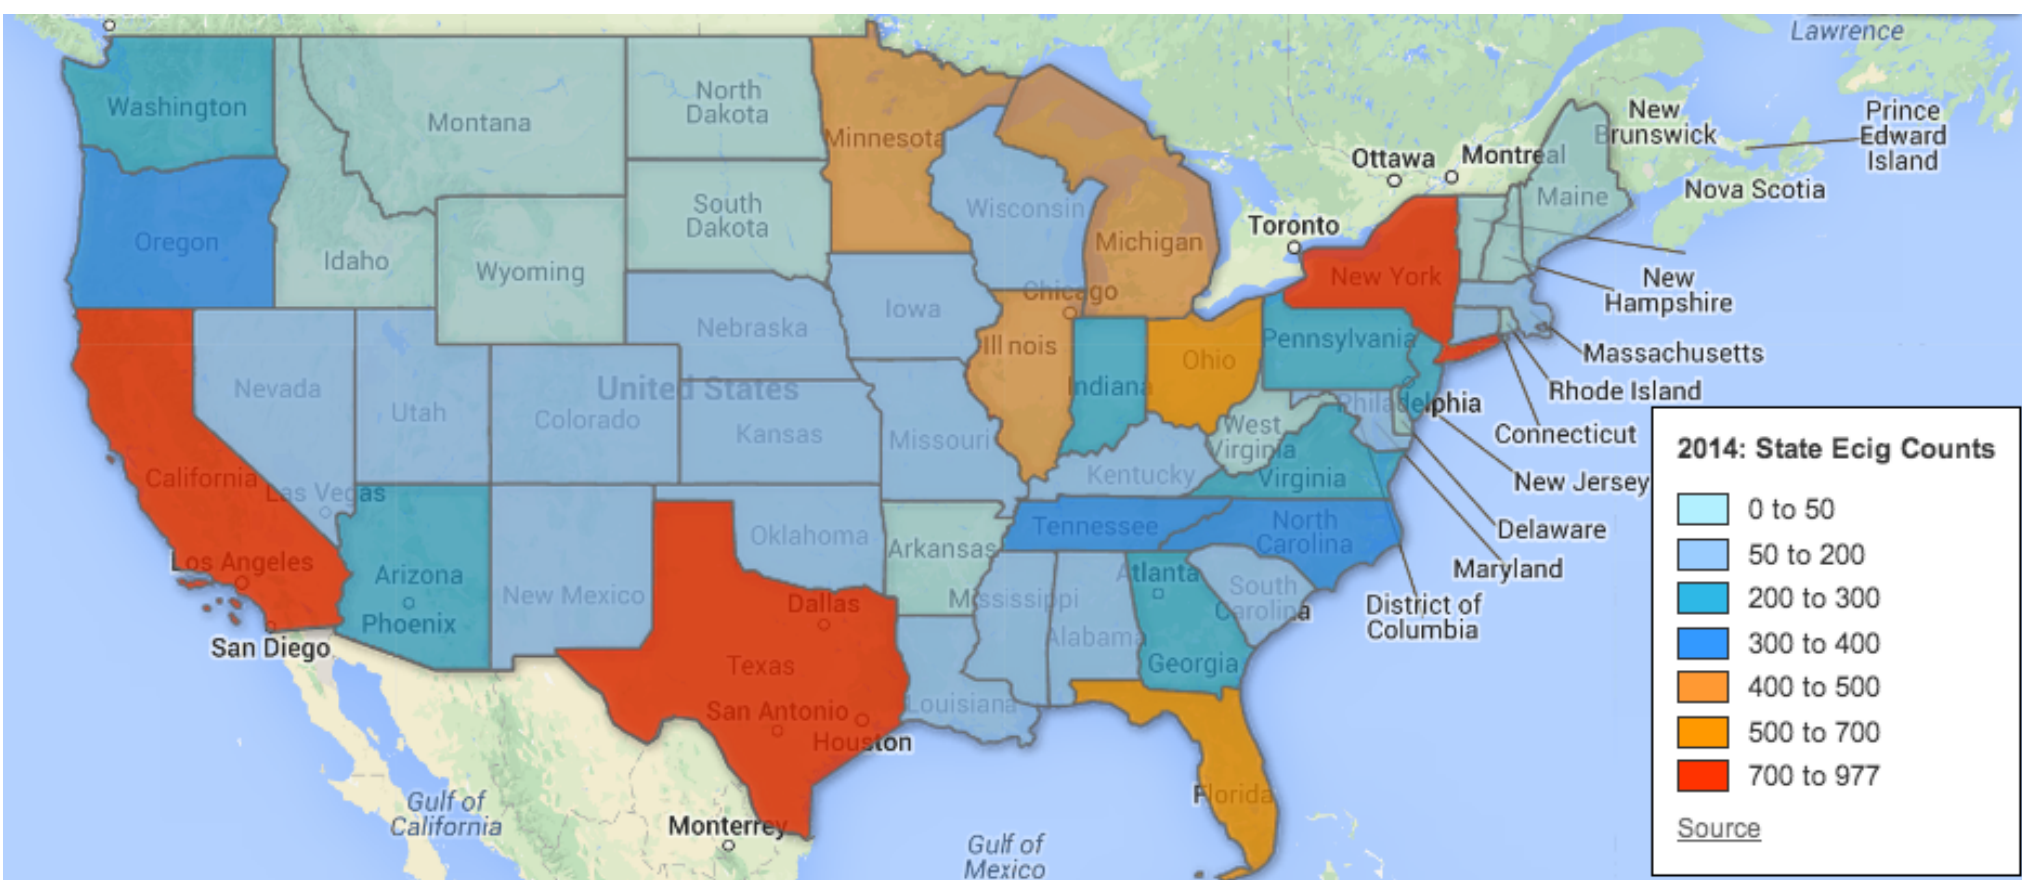
Hedonometrics: Measuring the Happiness of a Text



LabMT is a happiness distribution of the most frequently occurring 10,000 English words that were compiled through frequency distributions from literature,(Google Books), websites (Google Web Crawl), and Twitter. Surveys were created mimicking the self affective mannequin method, a sample of which is given above. Fifty participants were recruited using the online survey tool, Amazon Mechanical Turk, to identify the face that best matched the emotional response elicited by each word, which were then converted to a 9 point scale. On the numeric scale, 1 corresponded to the face with the largest frown and 9 to the face with the largest smile. The average happiness score, h_{avg} , for each word was then calculated via the arithmetic mean of 50 user reported ratings per word. Using the average happiness scores of each word, the average positivity of a subset of tweets can be quantified and used to compare different tweet distributions. To increase the emotional signal, neutral words ($4 \leq h_{avg} \leq 6$) are removed from the analysis. The standard approach to perform a hedonometric analysis on twitter is to create a happiness time-series. Outliers on the time series correspond to time-periods containing an overabundance of emotionally charged words. These outliers can then be investigated with word-shift graphs to help illuminate what is driving the emotional shift.

US Geo-tagged E-cig Tweets: March-August 2014

Approximately 1% of all tweets report the geo-location to within ten meters of accuracy of the user. This Geo-tagged data-set allows for regional comparisons of Electronic Cigarette mentions across the United States. All tweets that mention E-cigarette keywords between March and August 2014 were collected and binned by their U.S. state. Below we present a heat map of the counts of these tweets per each state. A substantial number of tweets is required to perform a meaningful hedonometric analysis of a region. This heat map shows us that California, Texas, and New York are most prevalently mentioning E-cigarettes, which in part is due to their larger populations relative to the rest of the US.



Using word shift graphs (see right pane for a more thorough explanation), the different types of words influencing the average happiness of several regions are displayed below. The leftmost shift compares the tweets from New York relative to California. In New York there are less occurrences of the negative words 'ban', 'restrictions', a higher occurrence of the negative words 'poison', 'die', 'worst', 'stupid'. On the right, tweets from Texas are compared to tweets from New York. In Texas there are less occurrences of 'banned', 'poison', 'protest', and 'nasty', and more occurrences of 'juice', 'flavor', 'candy', and 'quit'.

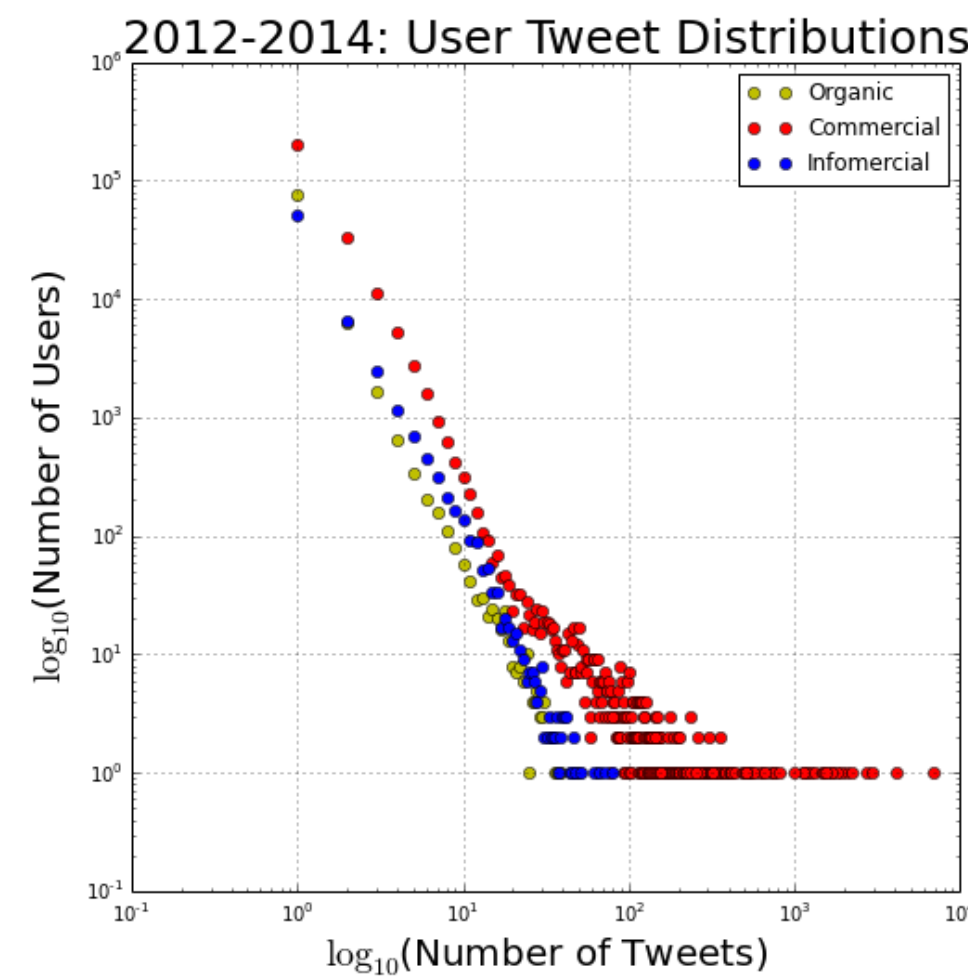

Acknowledgments

The authors wish to acknowledge the Vermont Advanced Computing Core, which is supported by NASA (NNX-08AO96G) at the University of Vermont which provided High Performance Computing resources that contributed to the research results reported within this poster. EMC was supported by the UVM Complex Systems Center, PSD was supported by NSF Career Award # 0846668. CMD and PSD were also supported by a grant from the MITRE Corporation. CJ, DG, RR, TJW, AK, and AS are supported in part by the National Institute of Health (NIH) Research wards R01DA014028 & R01HD075669, and by the Center of Biomedical Research Excellence Award P20GM103644 from the National Institute of General Medical Sciences.

Ecig Categorical Tweet Happiness Distributions

Using the happiness scores from LabMT, the average emotional rating of a corpus is calculated by tallying the appearance of words found in the intersection of the wordlist and a given corpus, in this case subsets of tweets. A weighted arithmetic mean of each word's frequency, f_{word} , and corresponding happiness score, h_{word} for each of the N words in a text yields the average happiness score for the corpus, \tilde{h}_{text} :

$$\tilde{h}_{text} = \frac{\sum_{w=1}^N f_w \cdot h_w}{\sum_{w=1}^N f_w}$$

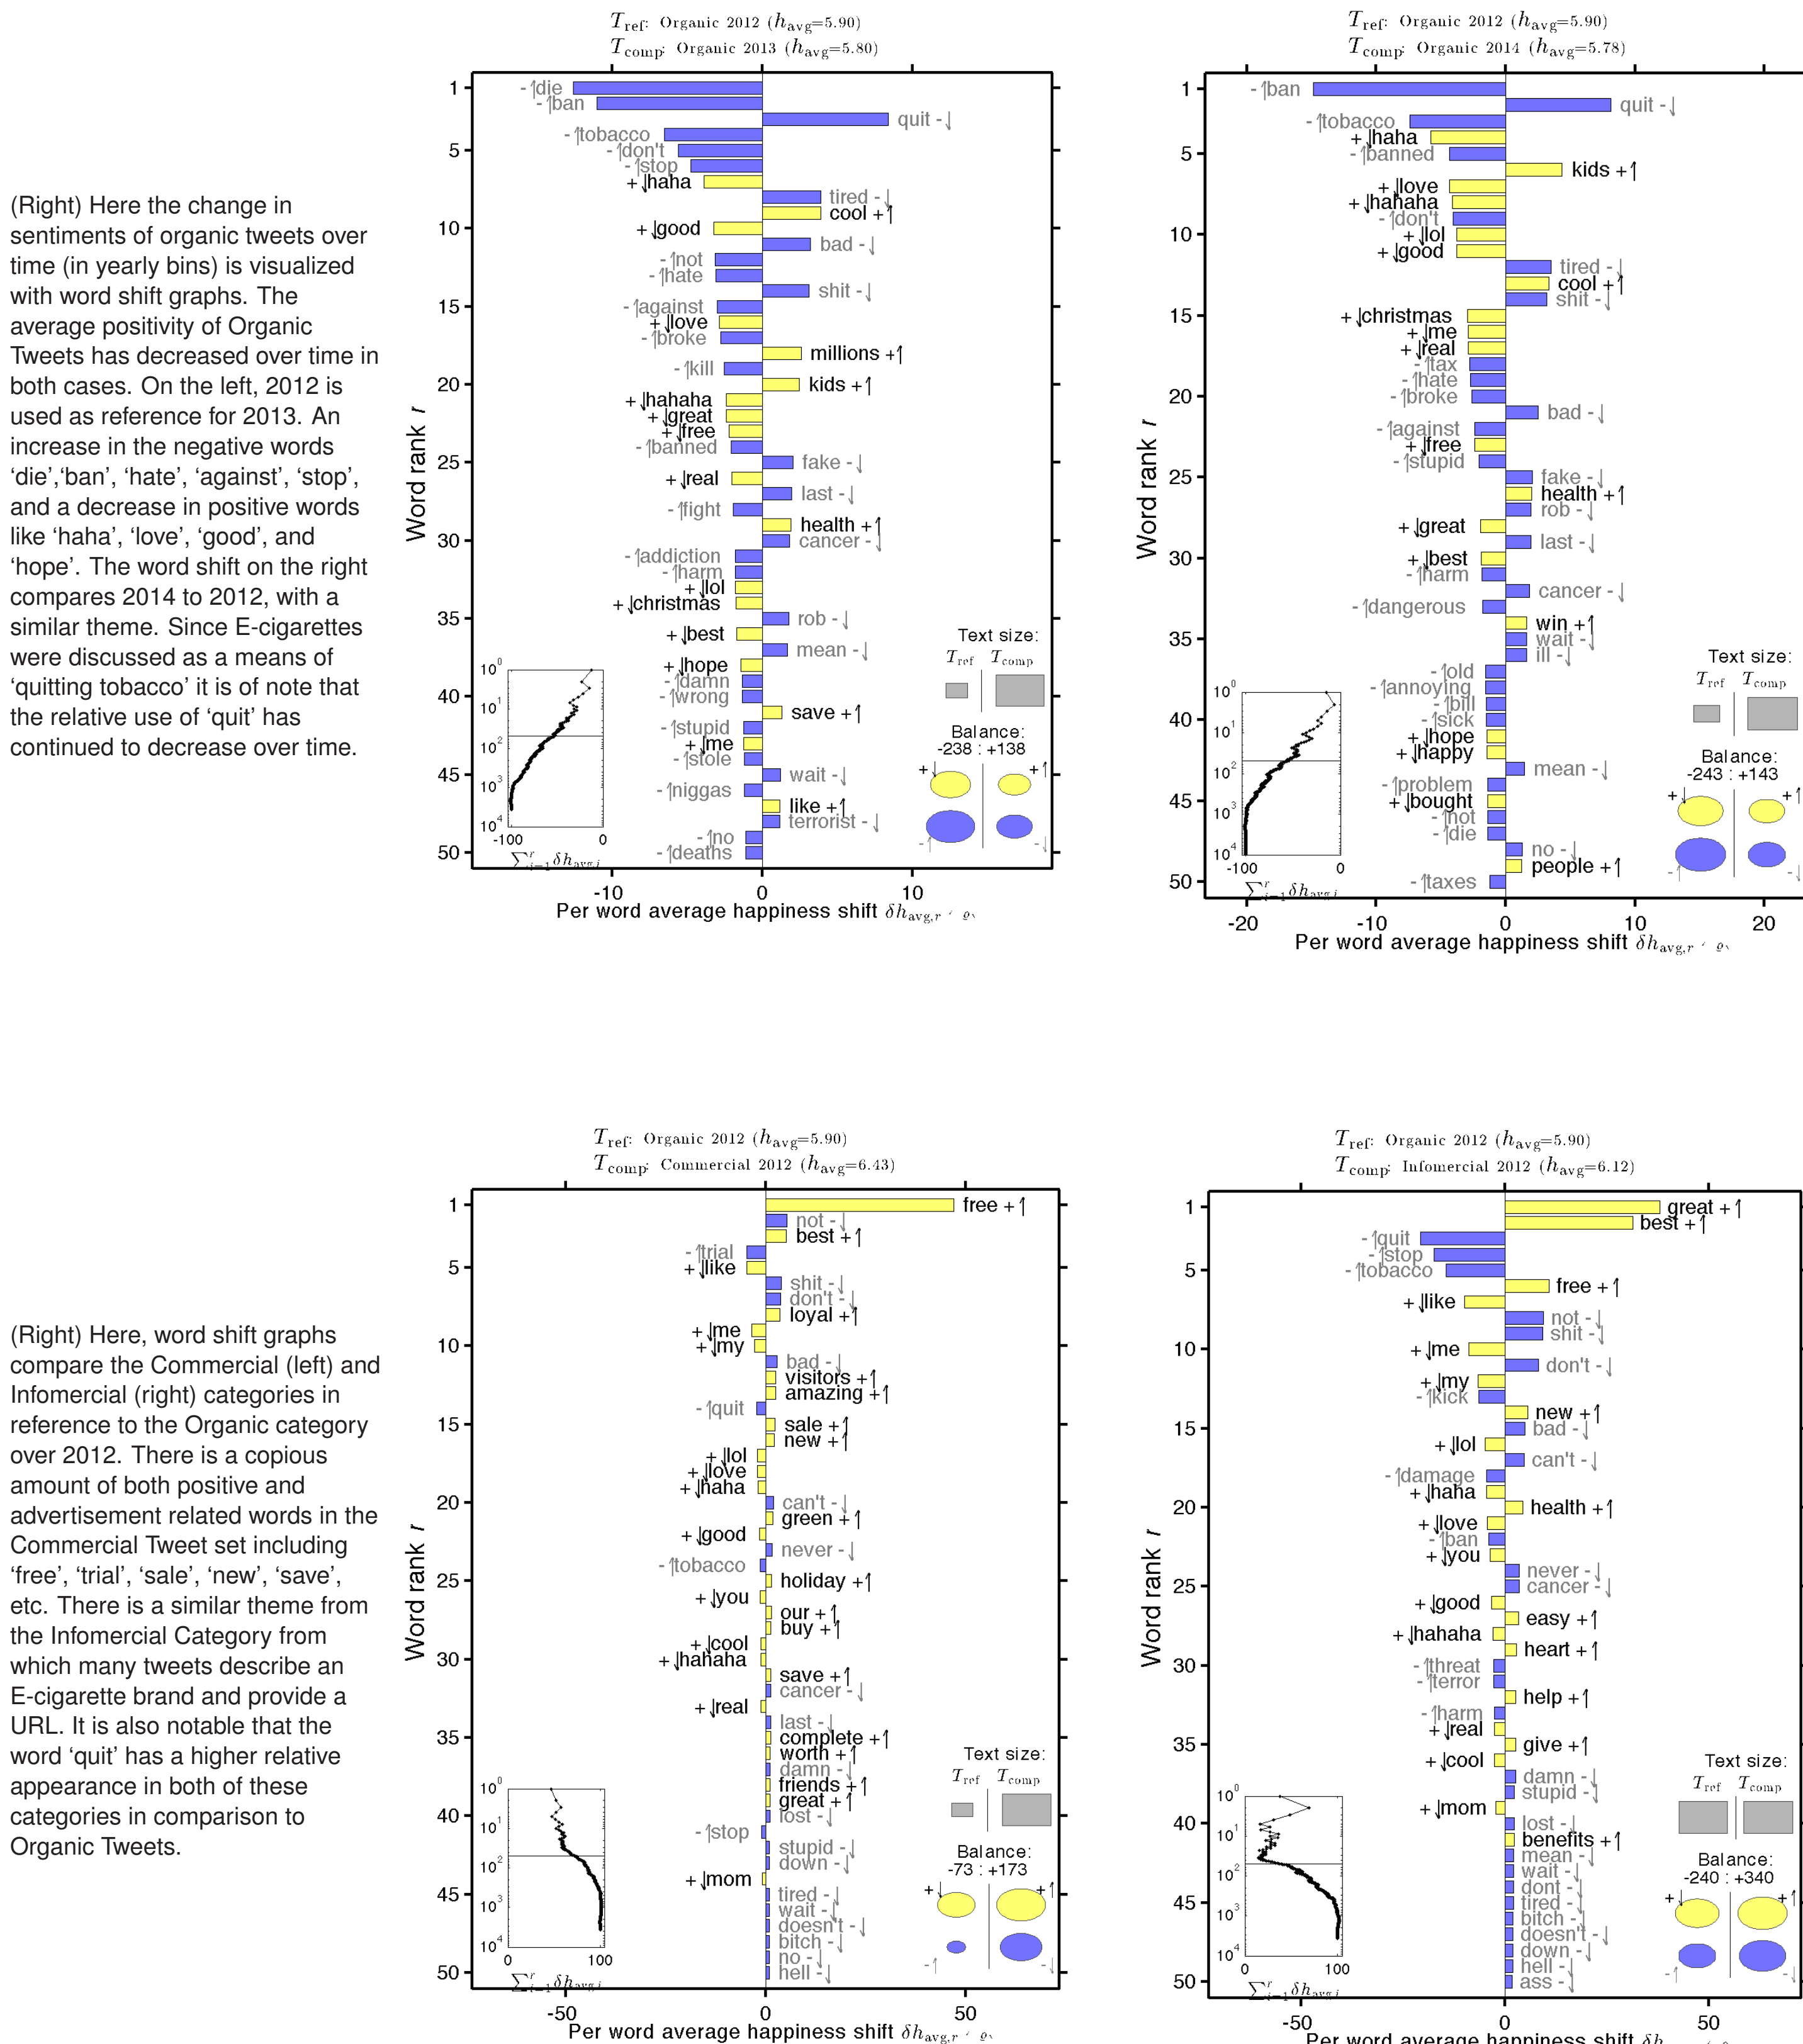


All E-cigarette mentions spanning January 2012 to July 2014 from the Twitter firehose, a 10% sample of all tweets, were collected and plotted as a function of time (upper left). The tweets were categorized into three classes: Organic, Commercialized, and Infermorial. Tweets with an abundance of marketing keywords were classified as Commercial. Tweets without these commercial keywords but containing a URL were classified as Infermorial. The remaining tweets make up the Organic Category. Categorizing and analyzing these categories separately is important to isolate true user sentiments pertaining to E-cigarettes. Marketing tweets use many overly positive words to advertise the product. There are also orders of magnitude more commercialized tweets than Organic and Infermorial Tweets. Since the use of Social Media as a marketing outlet for E-cigarettes is currently a hot political issue, it's important to isolate each of these categories and analyze each separately.

The number of E-cigarette tweets from each user in this study is displayed on logarithmic axes to the right. The Commercial distribution is quite different from the Organic and Infermorial in terms of its size and max number of individual user tweets. These marketing (SPAMers) tweet high volumes of E-cigarette related advertising, some of which are directed to Organic twitter patrons.

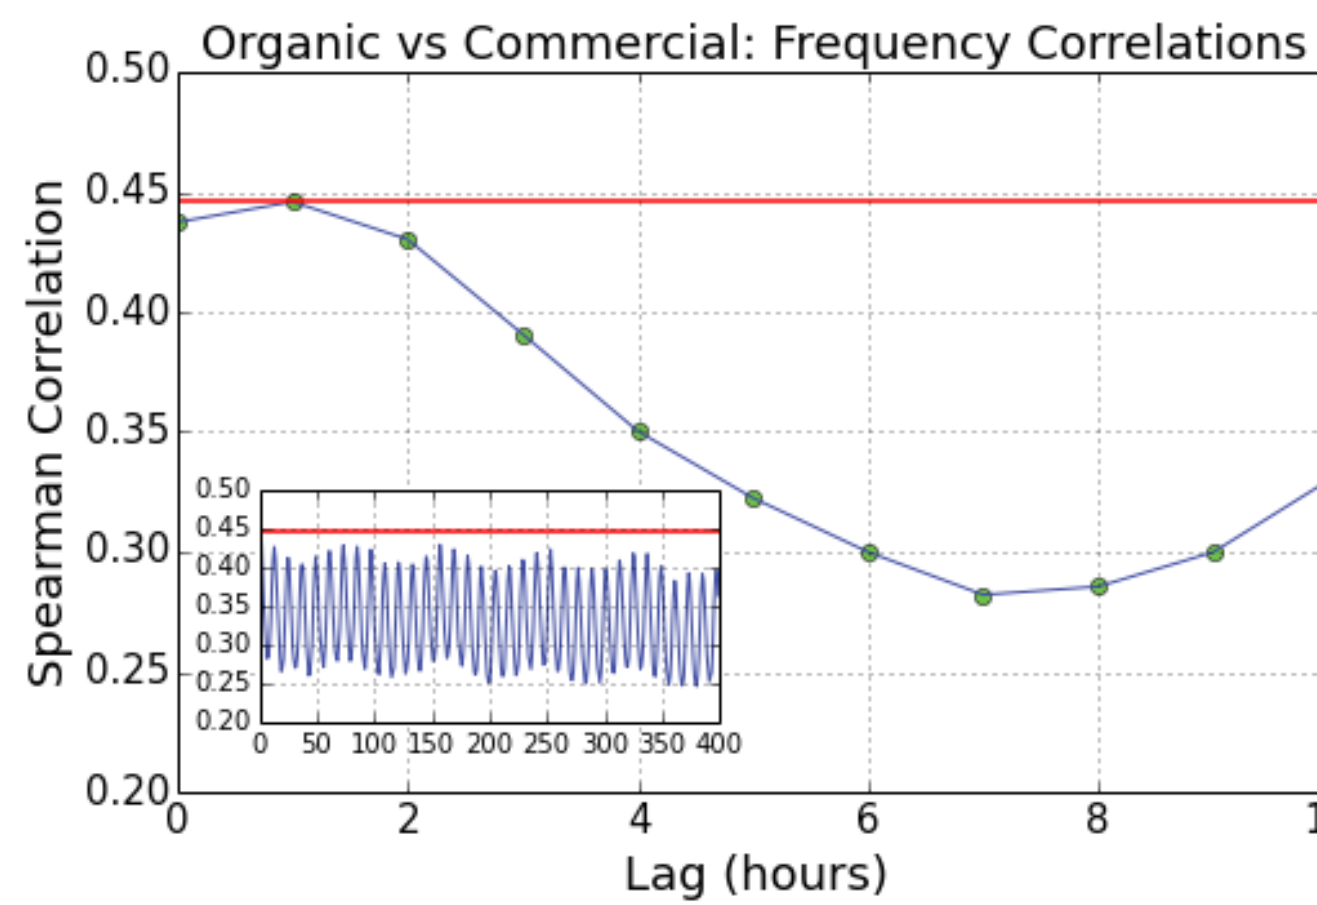
Categorical Word-Shifts: Sentiments over Time

Word-shift graphs illustrate two separate word frequency distributions. A reference period (T_{ref}), creates a basis of the emotional words being used to compare with another period, (T_{comp}). The top 50 words responsible for a happiness shift between the two periods are displayed, along with their contribution to shifting the average happiness of the tweet-set. The arrows (\uparrow , \downarrow) next to a word indicate an increase or decrease, respectively, of the word's frequency during the comparison period with respect to the reference period. The addition and subtraction signs indicate if the word contributes positively or negatively, respectively, to the average happiness score. Here we can identify the words contributing to the change in happiness between each category and over time.



(Right) Here, word shift graphs compare the Commercial (left) and Infermorial (right) categories in reference to the Organic category over 2012. There is a copious amount of both positive and advertisement related words in the Commercial Tweet set including 'free', 'trial', 'sale', 'new', 'save', etc. There is a similar theme from the Infermorial Category from which many tweets describe an E-cigarette brand and provide a URL. It is also notable that the word 'quit' has a higher relative appearance in both of these categories in comparison to Organic Tweets.

Categorical Time Series Correlations: Sentiments and Counts



Spearman Monthly Frequency Correlations

	Commercial	Infermorial
Organic	0.587	0.879
Infermorial	0.468	($p < 0.05$)


Spearman Monthly Happiness Correlations

	Commercial	Infermorial
Organic	0.434	0.762
Infermorial	0.406	($p < 0.05$)

Here, the relationship between each of these categorical tweets is explored. Some possible evidence that the Commercial or Infermorial categories are having an effect on Organic sentiments and frequencies is quantified. A nontrivial number of Commercial and Infermorial tweets are directed at Organic users. On the left, all Commercial and Organic tweets spanning January 2012 to July 2014 are binned into their hourly distributions and correlated as a function of an hourly lag. Each of these correlations are significant for the first 10 hours ($p < 0.01$). The correlation is maximized with a lag of one hour. The subgraph presents the correlations for a lag of up to 400 hours. The cyclic nature is due to the daily cycle of twitter activity, and although the correlation cyclically returns to above 0.40 it is maximized within the first hour. On the right, tweets from each distribution are binned by month and correlated against each other. Both the frequency and happiness distributions exhibit a strong positive Spearman correlation.

2013 Daily Resolution: Political Responses to Ecig Regulation

Each categorical time-series exhibits a severe negative trend occurring in January of 2013. Observing tweets at the daily resolution for Organic users, a spike in the frequency distribution occurred in December of 2013. There is an inverse relationship with the average happiness daily scores during this time period. This was during the time that the EU was debating a possible e-cigarette ban. Many tweets in this time frame were tagged with #EUcigban. The sentiments of Organic users as well as those from Commercialized accounts are visualized from this time period with word shift graphs.



(Right) On the leftmost word shift Organic Tweets from December 2013 (during the debate) are compared against tweets from January 2013 for reference. There is a plethora of negative words including 'ban', 'stop', 'against', 'disaster', 'deaths', and 'corruption' among others. The word shift on the right depicts the same time period, but are taken from the Commercial tweet category. Here there is an increase of the negative words 'die', 'stop', 'no', 'not', 'ban', and a decrease in positive words (related to marketing) 'free', 'happy', 'win', and 'thanks'.

