A Descriptive Study of the Prevalence and Typology of Alcohol-Related Posts in an Online Social Network for Smoking Cessation

AMY M. COHN, PH.D.,^{a,b,*} KANG ZHAO, PH.D.,^c SARAH CHA, M.S.P.H.,^d XI WANG, M.S.,^e MICHAEL S. AMATO, PH.D.,^d JENNIFER L. PEARSON, PH.D.,^{f,g} GEORGE D. PAPANDONATOS, PH.D.,^h & AMANDA L. GRAHAM, PH.D.^{b,d}

^aBattelle Memorial Institute, Arlington, Virginia

^bDepartment of Oncology, Georgetown University Medical Center, Washington, DC

^cDepartment of Management Sciences, Tippie College of Business, The University of Iowa, Iowa City, Iowa

^dSchroeder Institute for Tobacco Research and Policy Studies, Truth Initiative, Washington, DC

^eSchool of Information, Central University of Finance and Economics, Beijing, China

^fSchool of Community Health Sciences, University of Nevada, Reno, Nevada

^gDepartment of Health, Behavior and Society, Johns Hopkins Bloomberg School of Public Health, Baltimore, Maryland

^hCenter for Statistical Sciences, Brown University School of Public Health, Brown University, Providence, Rhode Island

ABSTRACT. Objective: Alcohol use and problem drinking are associated with smoking relapse and poor smoking-cessation success. User-generated content in online social networks for smoking cessation provides an opportunity to understand the challenges and treatment needs of smokers. This study used machine-learning text classification to identify the prevalence, sentiment, and social network correlates of alcohol-related content in the social network of a large online smoking-cessation program, BecomeAnEX.org. **Method:** Data were analyzed from 814,258 posts (January 2012 to May 2015). Posts containing alcohol keywords were coded via supervised machine-learning text classification for information about the user's personal experience with drinking, whether the user self-identified as a problem drinker or indicated problem drinking, and negative sentiment about drinking in

the context of a quit attempt (i.e., alcohol should be avoided during a quit attempt). **Results:** Less than 1% of posts were related to alcohol, contributed by 13% of users. Roughly a third of alcohol posts described a personal experience with drinking; very few (3%) indicated "problem drinking." The majority (70%) of alcohol posts did not express negative sentiment about drinking alcohol during a quit attempt. Users who did express negative sentiment about drinking were more centrally located within the network compared with those who did not. **Conclusions:** Discussion of alcohol was rare, and most posts did not signal the need to quit or abstain from drinking during a quit attempt. Featuring expert information or highlighting discussions that are consistent with treatment guidelines may be important steps to ensure smokers are educated about drinking risks. (*J. Stud. Alcohol Drugs*, 78, 665–673, 2017)

WEB-BASED SMOKING-CESSATION programs play an increasingly central role in tobacco control. They have broad reach to smokers (Pew Internet and American Life Project, 2006), can be accessed 24/7 and in an on-demand format, and are cost effective (An et al., 2010; Graham et al., 2013) and scalable (An et al., 2010; Bock et al., 2008; Cobb et al., 2005). Large-scale trials report quit rates of 18%–20% at 12 months (Graham et al., 2011; Muñoz et al., 2009), and several meta-analyses and systematic reviews provide evidence of effectiveness (Civljak et al., 2013; Graham et al., 2016; Hutton et al., 2011; Shahab & McEwen, 2009).

A unique aspect of some web-based cessation programs is their ability to facilitate the exchange of information and support throughout the quitting process via online social networks. Participation in these forms of online social interaction has been shown to have both physical and psycho-

logical benefits (Eysenbach et al., 2004; Idriss et al., 2009; Qiu et al., 2011; Wang et al., 2014; Zhao et al., 2014). User-generated content in online social networks also provides a unique opportunity to understand the specific challenges and potential treatment needs of smokers.

Alcohol use is one topic that may be especially relevant to explore in an online network for smoking cessation (Cunningham et al., 2006, 2008). Alcohol use is robustly associated with heavier smoking, greater nicotine dependence, lower motivation to quit (Cargill et al., 2001), and poor smoking-cessation outcomes (Hughes & Kalman, 2006; Leeman et al., 2008). It has also been identified as a primary trigger for smoking relapse (Kahler et al., 2010; McKee et al., 2006). Research shows that nearly half of all smoking-related slips occur during a drinking episode, and even after a quit attempt, individuals are more likely to slip or relapse back to smoking on days in which any alcohol is consumed, compared with nondrinking days, and on heavy drinking days (Leeman et al., 2008).

Because alcohol use is associated with relapse, tobacco cessation treatment guidelines recommend that smokers consider limiting or abstaining from alcohol while quitting (Fiore et al., 2008). Consuming alcohol is normative and routine for many smokers (Falk et al., 2006; Piasecki et al.,

Received: August 29, 2016. Revision: December 22, 2016.

This work is supported by National Cancer Institute Grant R01CA192345 to Amanda L. Graham/Kang Zhao.

*Correspondence may be sent to Amy M. Cohn at Battelle Memorial Institute, 2111 Wilson Blvd., Suite 1000, Arlington, VA 22201, or via email at: cohn@battelle.org.

2011), including those accessing online smoking-cessation programs (Cunningham et al., 2006), making it likely that discussion of alcohol use during a quit attempt would be common in an online community for smoking cessation.

Despite the proliferation of studies investigating adherence to and outcomes of web-based smoking-cessation programs, few have examined the content and structure of online social networks for cessation (Bondy & Bercovitz, 2013; Brandt et al., 2013; Burri et al., 2006; Cobb et al., 2010, 2013; Myneni et al., 2013; Selby et al., 2010; van Mierlo et al., 2012; Zhang et al., 2013), and none have focused on alcohol-related content within the network. Nicotine replacement therapy and coping with cravings are popular topics in cessation networks (Burri et al., 2006), and requests for support in these areas are common (Zhang et al., 2013). Several studies have identified characteristics of key network members who act as leaders in online cessation communities, providing advice and support to others (Cobb et al., 2010; Selby et al., 2010; van Mierlo et al., 2012). These individuals generate the most posts and participate in a wide range of threads; however, they comprise only a small proportion of active members.

Only one study of which we are aware has examined smoking and drinking conjointly in an online intervention. Cunningham et al. (2008) found that among registered users of an online smoking-cessation program, one third of current daily smokers were problem drinkers, 44% were social drinkers, and only 22% were nondrinkers. In sum, very little is known about user-generated content related to alcohol in online communities for smoking cessation.

Most studies of sentiment and network structure in online cessation programs have relied on manual coding of only a fraction of the available social network content (Brandt et al., 2013; Burri et al., 2006; Zhang et al., 2013). Manual coding of the hundreds of thousands of exchanges that occur within online social networks is impractical. However, advanced text analytics techniques, such as supervised machine learning, can allow for scalable coding of a large amount of unstructured data (Krippendorff, 2012) with relatively little human supervision. Text classification is one analytic technique that examines the content of documents (in this case, users' posts) and classifies them into different categories.

Text classification requires input from human domain experts at the outset to "train" the computer algorithm. For example, to decide whether a post has provided information about drinking behavior, human experts must first annotate, or code, a small number of posts as either providing this information or not. The classification algorithm then extracts various features from these posts (Guo et al., 2009; Zhao et al., 2014), such as the appearance of certain terms in the post (*wine*, *beer*). From human annotations, the computer algorithm then "learns" which features are more predictive of the category of posts. These powerful social computing methods have been used to automate the analysis of social

ties and sentiment in other areas of health (Chee et al., 2009; Guo et al., 2009; Zhao et al., 2014, 2016) but have yet to be applied broadly to online cessation programs.

This study sought to address five key research questions designed to understand the content and sentiment of alcohol-related posts in an online community for smoking cessation, as follows: (a) What is the prevalence of alcohol-related content? (b) How common are alcohol-related posts that suggest personal experience with drinking (vs. virtual celebratory "toasts" for smoking abstinence milestones or general comments about alcohol)? (c) How common are alcohol-related posts that indicate problem drinking behavior or a user who self-identifies as currently or previously in alcohol recovery? (d) Is the normative sentiment about drinking during a quit attempt focused on limiting alcohol consumption or on abstaining from alcohol completely (i.e., negative sentiment), in alignment with tobacco dependence treatment guidelines? (e) Is sentiment about drinking during a quit attempt associated with social network position? That is, are those who express negative sentiment about drinking during a quit attempt more centrally connected—and thus influential within the network—or are they more likely to be users who are only peripherally connected? Our analyses blend powerful machine-learning computational methods with traditional statistical approaches to examine these questions.

Method

Data source

We analyzed longitudinal data from BecomeAnEX.org (EX), a web-based smoking-cessation program developed and managed by Truth Initiative (formerly American Legacy Foundation). Launched in 2008, EX was developed in collaboration with the Mayo Clinic Nicotine Dependence Center in accordance with the Clinical Practice Guidelines for Treating Tobacco Dependence (Fiore et al., 2008). A national mass media campaign (McCausland et al., 2011) and ongoing online advertising have resulted in more than 700,000 registered users since its inception. Core elements of the site include setting a quit date, tracking cigarettes and identifying smoking triggers, building a support system, and providing information about pharmacotherapy.

EX also includes a large online community of current and former smokers who connect via multiple communication channels: private messages, public posts on member profile pages ("message boards"), group discussions, and blog posts. Blog posts and group discussions elicit many-to-many communications, whereas message board posts and private messages elicit one-to-one communications. Communication via blogs (and comments), message boards, and threaded group discussions are all public communications that can be accessed by all EX users; private messages occur only between two users.

The data set used in the current study spanned January 1, 2012, to May 31, 2015, and included records of 814,258 online activities by 9,377 users, including both posting and reading events. This period was selected because of the community's migration from a different platform before this period, which resulted in a slightly different user experience. Our analyses focus on this time frame given the stability of the social network feature set. The content of private messages was not included in the data set to protect privacy. The study protocol for these analyses was reviewed and approved by Chesapeake Institutional Review Board (protocol #00010302).

Data reduction—text classification

Manual annotation. We recruited three longstanding members of the EX community to serve as domain experts to first manually annotate (i.e., code) a subset of posts, which would later be used for machine-learning coding. Two experts in alcohol- and smoking-cessation research worked closely with the domain experts to develop and refine an annotation/coding guide based on the key research questions we sought to address.

For the first step in creating an annotation guide, we generated a list of alcohol-related keywords (e.g., *drink, alcohol, beer*; Appendix 1) and searched for these keywords in all public posts using an automated search process. We retrieved 19,547 posts that contained at least one of the alcohol-related keywords. Domain experts manually coded samples of the posts in batches of 200 posts; weekly group meetings were held following each batch to review annotations, clarify disagreements, and agree on revisions to the guide. This process continued iteratively until a clear set of guidelines was finalized with consensus among domain experts and the researchers, and κ of $> .70$ (Fleiss κ) was reached for each of the four domains. The final annotation guide included the following content domains (CD), each coded as yes/no:

CD1. "Did this post mention anything related to alcohol?"

CD2. "Did the author describe personal experience with drinking in the post?"

CD3. "Did the author describe him/herself as a problem drinker/alcoholic/in recovery or was problem drinking mentioned in the post?"

CD4. "Does the post express negative sentiment about alcohol as it relates to quitting?"

Following the practice period, the three domain experts received a random selection of 1,850 posts with alcohol-related keywords and completed annotations for the four content domains listed above. The rating given by the majority of domain experts (2 out of 3) was used as the final code when a rating disagreement arose. The final inter-rater reliability measured by Cohen's κ for each of the four content domains was adequate: CD1 = .74, CD2 = .85, CD3 = .88, CD4 = .71.

Machine learning. Four machine-learning-based binary classifiers, one for each of the content domains, were then trained from the set of 1,850 manually annotated posts. Of these 1,850 posts, domain experts identified 672 as being related to alcohol use. A classifier uses machine-learning algorithms to automatically detect the label of a post (yes/no for each of the content domains). For each classifier, we first extracted different types of characteristics (i.e., features), including meta-features and text features.

Meta-features were characteristics (but not content) of the post including post length (in word count) and post type (blog post/comment, message board post, group discussion post/reply). Text features included unigrams (the frequency of each word in a post), bigrams (the frequency of each possible two-word sequence in a post), and Term-Frequency-Inverse Document Frequency (TF-IDF) score. For example, "smoking cessation" and "cessation program" would be two possible bigrams in a post containing the phrase *smoking-cessation program*. TF-IDF score is a numeric value that is the product of term frequency, which measures how frequently a term appears in a single post, and inverse document frequency, which measures how rare a phrase is across all posts (Manning et al., 2008).

Next, we evaluated the performance of various classification algorithms on different combinations of feature sets (unigram, bigram, or TF-IDF scores) using standard 10-fold cross validation (Picard & Cook, 1984). This form of validation uses 90% of the posts to train a classifier and the remaining 10% to test the classifier's performance. The validation and training sets were rotated in 10 different trials. The performance of the different classification algorithms on each of the four content domain questions was evaluated using accuracy, F1 score, and area under the curve (AUC).

Accuracy is the percentage of posts whose labels were predicted correctly by the classifier. F1 score is also an accuracy measure, defined as

$$F1 = \frac{2 * \text{true positive}}{2 * \text{true positive} + \text{false negative} + \text{false positive}}.$$

F1 ranged from 0 to 1, with higher values indicating better performance. AUC is the area under the receiver operating characteristic curve, which plots the true-positive rate against the false-positive rate. AUC also has a range of 0 to 1, with 1 representing the perfect classifier based on the classifications provided by the human domain experts.

If no algorithm dominated others on all three metrics, the best-performing algorithm with the highest AUC was used, as AUC is more robust against skewed prior distributions (Gao et al., 2007). Finally, the best performing classifier algorithms were applied to 17,697 posts (19,547 minus 1,850) that contained alcohol-related keyword(s), but were not annotated, and a label of Yes or No was given to each post for each of the four content domains.

TABLE 1. The best performing classifier for each content domain

Content domain	Features	Algorithm	Accuracy	F1	AUC
CD1: "Did this post mention anything about alcohol?"	Meta features, unigrams	J48 decision tree	.86	.86	.89
CD2: "Did the author describe personal experience with drinking in the post?"	Meta features, unigrams	AdaBoost with naive Bayesian weak learners	.81	.81	.81
CD3: "Did the author describe himself/herself as a problem drinker/alcoholic/in recovery, or was problem drinking mentioned in the post?"	Frequency of 17 words/phrases related to problem drinking	Naive Bayesian with undersampling	.96	.95	.87
CD4: "Does the post express negative sentiment about alcohol as it related to quitting?"	Meta features, unigrams weighted by TF-IDF scores and distances to alcohol-related keywords	Random forest with undersampling	.83	.81	.81

Notes: F1 = a measure of the accuracy of posts whose labels correctly predicted the classifier (range: 0–1); AUC = area under the receiver operating characteristic curve; J48 = an open source implementation of C4.5 decision tree algorithm, which uses normalized information gain to decide how to split a node of the decision tree.

Data reduction—Social network variables

A user's social network position in the EX community was calculated by examining each user's reading and posting behaviors and constructing a social network metric based on the flow of information either "toward" or "away from" that user. In this approach, each node represents a user; a directed tie from user A to user B indicates that a post or message published by user A was accessed by user B. One way to measure how central or important an individual is within such a social network is to use in-degree and out-degree centralities (Krippendorff, 2012). In degree reflects the total number of EX members a user has been exposed to by reading their posts; out degree reflects the total number of EX members who have read the content posted by a user (Zhao et al., 2016). Higher values for in/out degree reflect greater centrality in the social network.

Data analysis

We first examined the accuracy of algorithms and results of machine-learning coding of the four content domains. Next, we calculated the frequency and prevalence of any posts containing alcohol-related content (CD1) and alcohol-related content that mentioned personal experience with drinking (CD2) or problem drinking (CD3), as well as negative sentiment about drinking in the context of a quit attempt (CD4). Univariate analysis of variance tests were then conducted to examine whether expressing negative sentiment about drinking (yes/no) was associated with higher or lower levels of network connectivity, as measured by in-degree and out-degree centralities (dependent variables). To meet model assumptions, dependent variables were evaluated for nonnormality, and both square root and logarithmic transformations were examined to correct for

skewness of data. The log transformation best minimized skewness for both in degree and out degree and was used in the analyses of variance below.

Results

Machine-learning classification

Classifier performance is reported in Table 1. Standard classification schemes worked well for CD1 and CD2 classifiers, as evidenced by their satisfactory performance as measured by accuracy, F1 score, and AUC. However, for CD3 and CD4, posts with positive labels ("yes") were rare (3.30% for CD3 and 13.46% for CD4 among annotated posts). For these two classification tasks, standard classification schemes were characterized by low recall for the positive ("yes") labels as well as low AUC, even though the overall accuracy rates were higher than 85%.

To improve CD3 and CD4 classifiers for the positive labels, we also adopted undersampling techniques for CD3 and CD4 to mitigate the effect of low prior probabilities. Undersampling the negative label can generate a training data set that is more balanced between both positive ("yes") and negative ("no") labels so that a classifier can better learn differences between the two.

In addition, for CD3, we found that using a list of high-frequency keywords and phrases (Appendix 2) improved performance beyond using all words and phrases (i.e., unigrams). For CD4, because the focus was on sentiment toward alcohol use—instead of the sentiment of the whole post—we leveraged a popular method for aspect-based sentiment analysis (Feldman, 2013). We weighted unigrams by their distance to alcohol-related keywords in a post, assuming that words appearing closer to these keywords in texts contribute more to the author's sentiment toward alcohol

use. With these modifications to standard text classification schemes, we were able to improve the performance of CD3 and CD4 classifiers. Overall, all classifiers achieved satisfactory performance, with more than 80% accuracy for all three performance metrics.

Research Question 1

After applying the four classifiers to all 17,697 unannotated posts with at least one alcohol-related keyword (e.g., *drink*), our CD1 classifier identified 6,527 posts that were correctly related to alcohol use (e.g., *drink alcohol* instead of *drink water/coffee*). Adding the 672 posts that were annotated as “positive” for alcohol content by domain experts, we found a total of 7,199 alcohol-related posts (6,527 plus 672). The final number of posts that contained alcohol-related content ($n = 7,199$) differed from the number of posts that contained all possible mentions of alcohol keywords ($n = 19,547$) because some mentions were not related to alcohol. For example, although the word *drink* was an alcohol keyword, posts that had the word *drink* were dropped from the final data set if they referred to drinking coffee, water, soda, or any other nonalcoholic beverage.

Alcohol-related posts represented just under 1% of all posts made during the study period ($N = 814,258$ posts). The alcohol-related posts were made by 13% of users ($N = 9,377$ users) who made any post in the community during the study period.

Research Questions 2 and 3

Among the alcohol-related posts, 33.02% ($n = 2,377$) described a personal experience with drinking (e.g., “I drank a lot of beer Sunday night, stayed up late & smoked double the cigarettes I usually smoke”; “I’m also making a goal to not drink until I’m completely confident that I will not smoke!!!”). Among the alcohol-related posts, only 3.65% ($n = 263$) were identified as being related to problem drinking (e.g., “I am going to go to an AA meeting today”; “I’m drinking way too much and blacking out”).

Research Question 4

Among all alcohol-related posts, 33.07% ($n = 2,381$) expressed negative sentiment about drinking alcohol when quitting smoking; that is, the user focused on limiting or abstaining from drinking alcohol completely during a quit attempt (e.g., “No more drinking for a good while, as that’s a huge trigger”; “I absolutely promise you that drinking alcohol at this early stage of your quit especially when you are struggling is self-sabotage”). Alcohol-related posts indicating negative sentiment were roughly equally distributed across three of the subnetworks, with 33% occurring in blog posts, 31% in message boards, and 31% in group discussions.

Research Question 5

For the final research questions, we were interested in examining whether users who posted negative sentiment about alcohol use during a quit attempt were more likely to be highly connected and active users on the website compared with users who did not post negative sentiment about alcohol. If highly connected users regard drinking alcohol as an obstacle to quit success and communicate this to others on the website, this would suggest a possible mechanism for spreading information about alcohol’s negative effects on cessation success in future interventions.

The EX social network during the study period included 2.58 million ties and 71,251 nodes with at least one tie. Figure 1 shows differences in network characteristics as a function of negative sentiment about drinking. Those who expressed negative sentiment about drinking (the circles), compared with those who did not express negative sentiment (the triangles), showed greater centrality both in degree, $F(1, 1230) = 107.55$ ($M [SD] = 701.09 [845.06]$ vs. $M [SD] = 281.72 [394.19]$, respectively), and out degree, $F(1, 1242) = 111.66$ ($M [SD] = 852.61 [1,174]$ vs. $M [SD] = 308.20 [451.92]$, respectively).

Discussion

Our analyses showed that alcohol-related content was present in 1% of all posts and made by 13% of all contributing users. A third of the alcohol-related posts described personal experience with drinking, whereas only 4% indicated some level of past or current problem drinking behavior by the user. The majority of posts did not express negative sentiment about drinking alcohol during a quit attempt. However, users who did express negative sentiment about drinking during a quit attempt were more highly connected to others in the network, with roughly double the number of network ties compared with those who did not express negative sentiment.

It was surprising that alcohol-related content was this rare. Research shows about a quarter to a third of smokers enrolling in smoking-cessation trials (Fridberg et al., 2014; Leeman et al., 2008; O’Malley et al., 2009) or calling a state quit line report moderate to hazardous drinking behavior (Toll et al., 2012). It is possible that members of an online smoking-cessation community are reluctant to discuss the topic of alcohol, especially if they perceive that alcohol is not commonly discussed in the community or that users, particularly more influential ones, hold negative sentiment about drinking and may disparage or judge users for relapsing because of a drink. Alternatively, smokers who are concurrent problem drinkers may be seeking other, more intensive forms of cessation treatment and may not be represented in an online cessation program.

Our findings are somewhat consistent with other work of alcohol-related posts on Twitter. For example, one study by

FIGURE 1. Complementary cumulative distributions (a) in-degree centrality and (b) out-degree centrality for users who expressed negative sentiment toward alcohol use (circles) versus those who did not (triangles).

West et al. (2006) found that less than 1% of Twitter posts culled over a 36-day period indicated problem drinking behavior, whereas another study found that 7.8% of Twitter posts culled over a similar time frame indicated alcohol-related content using a select set of five alcohol-related keywords (*alcohol, beer, liquor, vodka, and hangover*) (Cavazos-Rehg et al., 2015). It is important to note that these studies differed from ours in that they focused on a select group of alcohol-related keywords and coded posts over a much shorter time frame.

The low prevalence of negative sentiment about drinking alcohol during a quit attempt was also surprising, espe-

cially because this is at odds with the tobacco dependence treatment guidelines that encourage smokers to avoid or eliminate smoking triggers, including alcohol (Fiore et al., 2008). However, our result is consistent with analysis of alcohol-related posts on Twitter, where anti-alcohol tweets were found to be much less common than pro-alcohol tweets (Cavazos-Rehg et al., 2015). Given that users who did express negative sentiment about drinking were more likely to be centrally connected in the network, we may also presume that they are more likely to be abstinent from smoking (Cobb et al., 2010). If this is indeed the case, the higher prevalence

of negative sentiment among well-connected users may reflect that those users are confident exsmokers who see any alcohol use as a threat to the success of one's quitting and are able to see the clear benefits of abstaining from alcohol completely during the quit-smoking process.

Because any level of drinking—especially problem drinking—increases the risk of smoking relapse (Dawson, 2000; Leeman et al., 2008), it is important to augment user-generated content with didactic “expert-generated” information that adequately addresses the links between smoking and drinking. The role of the community administrator could be leveraged to seed discussions about the risks of drinking alcohol and encourage users to share their experiences with managing alcohol during their quit attempt. Further, it may be useful to prominently feature content about alcohol to make more apparent the challenges about drinking, particularly so that less well-connected users can make an informed decision about whether or when to drink.

This study had several limitations. First, machine text classifiers achieved satisfactory but not perfect accuracy. This could be because of the type of qualitative data being extracted, which was highly nuanced even for domain experts. It is possible that we were unable to code all possible posts referencing alcohol use, particularly if the wording used in the post was highly nuanced or colloquial (i.e., “tossing a few back at happy hour”).

Second, annotations for CD4 focused only on negative sentiment about drinking during a cessation attempt and did not code for positive or neutral sentiment. We were specifically interested in the extent to which user-generated content aligned with tobacco treatment guidelines (i.e., expressed negative sentiment that alcohol should be limited or avoided) but may have missed important insights about the extent to which community members expressed overtly positive sentiment about alcohol. Future work should examine the extent to which both negative and positive sentiment about alcohol use during a quit attempt occur in the network.

Third, we did not link the creation of an alcohol-related post to the user's social network position at the time of the post, nor examine whether or to what degree network position may have changed over time as a function of posting about alcohol-related content.

Finally, data on user characteristics were not available, so we were unable to characterize the individuals included in these analyses. The goal of our study was to understand the extent to which alcohol was present in user-generated content rather than identifying the individual characteristics of those who posted about alcohol use. This is an important topic for future research.

This study adds to the literature in several ways. This is the first and largest study that has mined the prevalence and sentiment of alcohol-related content in an online smoking-cessation program. At an aggregated level, prior studies have examined associations of mental health status or psychologi-

cal distress with adherence to online smoking-cessation programs and characteristics associated with users (Lukowski et al., 2015; Talati et al., 2016; Vickerman et al., 2015).

Second, this study used machine learning for text classification as an innovative and novel approach to study alcohol-related conversations at a scale that is not possible with manual coding (Krippendorff, 2012). Such scalability on a large-scale data set reduces sampling bias and enables analyses at both the community and individual level.

Third, this work is consistent with Substance Abuse and Mental Health Services Administration initiatives to integrate substance use and tobacco dependence treatments (Santhosh et al., 2014) as well as the priorities of the National Institutes of Health Collaborative Research on Addictions to understand factors associated with addiction comorbidities through multidisciplinary work (National Institutes of Health, 2016). Our findings provide an important foundation for understanding the interplay between smoking and alcohol use among members of an online smoking-cessation program and for identifying potential touchpoints for intervention.

Conclusions

This study breaks new ground as the first exploration of alcohol-related content in an online social network for smoking cessation. We have blended social network analyses with rigorous methods in machine learning to focus on alcohol use as an important comorbidity of tobacco dependence. Our analyses provide a methodological framework that could be applied to other important topics related to smoking cessation (e.g., use of pharmacotherapy or other nicotine products). Our overarching goal is to build a rich understanding of the individual-, interpersonal-, and network-level influences on tobacco use behavior that may be amenable to intervention. Our findings underscore the notion that individuals who are centrally connected in a social network can play an important role in the spread of information. With the evolution of Web 2.0 technologies and the increasingly customizable array of communication channels, our findings lay important groundwork for exploring ways to better address the comorbid use of alcohol and tobacco.

Conflict of Interest Statement

Amy M. Cohn and Jennifer L. Pearson were employees of Truth Initiative, and Sara Cha, Michael S. Amato, and Amanda L. Graham are current employees of Truth Initiative, which runs the BecomeAnEX.org smoking-cessation website.

References

- An, L. C., Betzner, A., Schillo, B., Luxenberg, M. G., Christenson, M., Wendling, A., . . . Kavanaugh, A. (2010). The comparative effectiveness of clinic, work-site, phone, and Web-based tobacco treatment programs. *Nicotine & Tobacco Research*, 12, 989–996. doi:10.1093/ntr/ntq133

- Bock, B. C., Graham, A. L., Whiteley, J. A., & Stoddard, J. L. (2008). A review of Web-Assisted Tobacco Interventions (WATIs). *Journal of Medical Internet Research*, 10(5), e39. doi:10.2196/jmir.989
- Bondy, S. J., & Bercovitz, K. L. (2013). Hike up yer skirt, and quit. What motivates and supports smoking cessation in builders and renovators. *International Journal of Environmental Research and Public Health*, 10, 623–637. doi:10.3390/ijerph10020623
- Brandt, C. L., Dalum, P., Skov-Ettrup, L., & Tolstrup, J. S. (2013). "After all — it doesn't kill you to quit smoking": An explorative analysis of the blog in a smoking cessation intervention. *Scandinavian Journal of Public Health*, 41, 655–661. doi:10.1177/1403494813489602
- Burri, M., Baujard, V., & Etter, J.-F. (2006). A qualitative analysis of an internet discussion forum for recent ex-smokers. *Nicotine & Tobacco Research*, 8, Supplement 1, S13–S19. doi:10.1080/14622200601042513
- Cargill, B. R., Emmons, K. M., Kahler, C. W., & Brown, R. A. (2001). Relationship among alcohol use, depression, smoking behavior, and motivation to quit smoking with hospitalized smokers. *Psychology of Addictive Behaviors*, 15, 272–275. doi:10.1037/0893-164X.15.3.272
- Cavazos-Rehg, P. A., Krauss, M. J., Sowles, S. J., & Bierut, L. J. (2015). "Hey everyone, I'm drunk." An evaluation of drinking-related Twitter chatter. *Journal of Studies on Alcohol and Drugs*, 76, 635–643. doi:10.15288/jsad.2015.76.635
- Chee, B. W., Berlin, R., & Schatz, B. R. (2009). Measuring population health using personal health messages. *AMIA Annual Symposium Proceedings Archive*, 2009, 92–96.
- Civiljak, M., Stead, L. F., Hartmann-Boyce, J., Sheikh, A., & Car, J. (2013, July 10). Internet-based interventions for smoking cessation. *Cochrane Database of Systematic Reviews*, Issue 7, Article No. CD007078. doi:10.1002/14651858.CD007078.pub4
- Cobb, N. K., Graham, A. L., & Abrams, D. B. (2010). Social network structure of a large online community for smoking cessation. *American Journal of Public Health*, 100, 1282–1289. doi:10.2105/AJPH.2009.165449
- Cobb, N. K., Graham, A. L., Bock, B. C., Papandonatos, G., & Abrams, D. B. (2005). Initial evaluation of a real-world Internet smoking cessation system. *Nicotine & Tobacco Research*, 7, 207–216. doi:10.1080/14622200500055319
- Cobb, N. K., Mays, D., & Graham, A. L. (2013). Sentiment analysis to determine the impact of online messages on smokers' choices to use varenicline. *Journal of the National Cancer Institute Monographs*, 2013, 224–230. doi:10.1093/jncimonographs/igt020
- Cunningham, J. A., Selby, P., & van Mierlo, T. (2006). Integrated online services for smokers and drinkers? Use of the check your drinking assessment screener by participants of the Stop Smoking Center. *Nicotine & Tobacco Research*, 8, Supplement 1, S21–S25. doi:10.1080/14622200601039865
- Cunningham, J. A., van Mierlo, T., & Fournier, R. (2008). An online support group for problem drinkers: AlcoholHelpCenter.net. *Patient Education and Counseling*, 70, 193–198. doi:10.1016/j.pec.2007.10.003
- Dawson, D. A. (2000). Drinking as a risk factor for sustained smoking. *Drug and Alcohol Dependence*, 59, 235–249. doi:10.1016/S0376-8716(99)00130-1
- Eysenbach, G., Powell, J., Englesakis, M., Rizo, C., & Stern, A. (2004). Health related virtual communities and electronic support groups: Systematic review of the effects of online peer to peer interactions. *BMJ*, 328, 1166. doi:10.1136/bmj.328.7449.1166
- Falk, D. E., Yi, H. Y., & Hiller-Sturmhöfel, S. (2006). An epidemiologic analysis of co-occurring alcohol and tobacco use and disorders: Findings from the National Epidemiologic Survey on Alcohol and Related Conditions. *Alcohol Research & Health*, 29, 162–171.
- Feldman, R. (2013). Techniques and applications for sentiment analysis. *Communications of the ACM*, 56, 82–89. doi:10.1145/2436256.2436274
- Fiore, M. C., Jaén, C. R., Baker, T. B., Bailey, W. C., Benowitz, N. N., Curry, S. J., . . . Wewers, M. E. (2008). *Treating tobacco use and dependence: 2008 update*. Rockville, MD: Agency for Healthcare Research and Quality. Retrieved from <https://bphc.hrsa.gov/buckets/treatingtobacco.pdf> p 5
- Fridberg, D. J., Cao, D., Grant, J. E., & King, A. C. (2014). Naltrexone improves quit rates, attenuates smoking urge, and reduces alcohol use in heavy drinking smokers attempting to quit smoking. *Alcoholism: Clinical and Experimental Research*, 38, 2622–2629. doi:10.1111/acer.12513
- Gao, J., Fan, W., Han, J., & Philip, S. Y. (2007). A general framework for mining concept-drifting data streams with skewed distributions. *Proceedings of the 2007 SIAM International Conference on Data Mining*, 3–14. doi:10.1137/1.9781611972771.1
- Graham, A. L., Carpenter, K. M., Cha, S., Cole, S., Jacobs, M. A., Raskob, M., & Cole-Lewis, H. (2016). Systematic review and meta-analysis of Internet interventions for smoking cessation among adults. *Substance Abuse and Rehabilitation*, 7, 55–69. doi:10.2147/SAR.S101660
- Graham, A. L., Chang, Y., Fang, Y., Cobb, N. K., Tinkelman, D. S., Niaura, R. S., . . . Mandelblatt, J. S. (2013). Cost-effectiveness of internet and telephone treatment for smoking cessation: An economic evaluation of The iQUIT Study. *Tobacco Control*, 22, e11. Retrieved from <http://tobaccocontrol.bmj.com/content/22/6/e11.info>
- Graham, A. L., Cobb, N. K., Papandonatos, G. D., Moreno, J. L., Kang, H., Tinkelman, D. G., . . . Abrams, D. B. (2011). A randomized trial of Internet and telephone treatment for smoking cessation. *Archives of Internal Medicine*, 171, 46–53. doi:10.1001/archinternmed.2010.451
- Guo, L., Tan, E., Chen, S., Zhang, X., & Zhao, Y. E. (2009). Analyzing patterns of user content generation in online social networks. *Proceedings of the 15th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 369–378. Retrieved from <https://dl.acm.org/citation.cfm?id=1557064>
- Hughes, J. R., & Kalman, D. (2006). Do smokers with alcohol problems have more difficulty quitting? *Drug and Alcohol Dependence*, 82, 91–102. doi:10.1016/j.drugalcdep.2005.08.018
- Hutton, H. E., Wilson, L. M., Apelberg, B. J., Tang, E. A., Odelola, O., Bass, E. B., & Chander, G. (2011). A systematic review of randomized controlled trials: Web-based interventions for smoking cessation among adolescents, college students, and adults. *Nicotine & Tobacco Research*, 13, 227–238. doi:10.1093/ntr/ntq252
- Idriss, S. Z., Kvedar, J. C., & Watson, A. J. (2009). The role of online support communities: Benefits of expanded social networks to patients with psoriasis. *Archives of Dermatology*, 145, 46–51. doi:10.1001/archdermatol.2008.529
- Kahler, C. W., Spillane, N. S., & Metrik, J. (2010). Alcohol use and initial smoking lapses among heavy drinkers in smoking cessation treatment. *Nicotine & Tobacco Research*, 12, 781–785. doi:10.1093/ntr/ntq083
- Krippendorff, K. (2012). *Content analysis: An introduction to its methodology*. Atlanta, GA: Sage.
- Leeman, R. F., McKee, S. A., Toll, B. A., Krishnan-Sarin, S., Cooney, J. L., Makuch, R. W., & O'Malley, S. S. (2008). Risk factors for treatment failure in smokers: Relationship to alcohol use and to lifetime history of an alcohol use disorder. *Nicotine & Tobacco Research*, 10, 1793–1809. doi:10.1080/14622200802443742
- Lukowski, A. V., Morris, C. D., Young, S. E., & Tinkelman, D. (2015). Quit-line outcomes for smokers in 6 states: Rates of successful quitting vary by mental health status. *Nicotine & Tobacco Research*, 17, 924–930. doi:10.1093/ntr/ntu252
- Manning, C., Raghavan, P., & Schütze, H. (2008). *Introduction to information retrieval* (1st ed.). New York, NY: Cambridge University Press.
- McCausland, K. L., Curry, L. E., Mushro, A., Carothers, S., Xiao, H., & Vallone, D. M. (2011). Promoting a web-based smoking cessation intervention: Implications for practice. *Cases in Public Health Communication & Marketing*, 5, 3–26. Retrieved from <http://www.casesjournal.org/volume5summer>
- McKee, S. A., Krishnan-Sarin, S., Shi, J., Mase, T., & O'Malley, S. S. (2006). Modeling the effect of alcohol on smoking lapse behavior. *Psychopharmacology*, 189, 201–210. doi:10.1007/s00213-006-0551-8

- Muñoz, R. F., Barrera, A. Z., Delucchi, K., Penilla, C., Torres, L. D., & Pérez-Stable, E. J. (2009). International Spanish/English Internet smoking cessation trial yields 20% abstinence rates at 1 year. *Nicotine & Tobacco Research, 11*, 1025–1034. doi:10.1093/ntr/ntp090
- Myneni, S., Cobb, N. K., & Cohen, T. (2013). Finding meaning in social media: Content-based social network analysis of QuitNet to identify new opportunities for health promotion. *Studies in Health Technology and Informatics, 192*, 807–811. doi:10.3233/978-1-61499-289-9-807
- National Institutes of Health. (2016). *CRAN: Collaborative Research on Addiction at the National Institutes of Health, Strategic Plan 2016–2021*. Retrieved from <https://www.addictionresearch.nih.gov/cran-strategic-plan-2016-2021>
- O'Malley, S. S., Krishnan-Sarin, S., McKee, S. A., Leeman, R. F., Cooney, N. L., Meandzija, B., . . . Makuch, R. W. (2009). Dose-dependent reduction of hazardous alcohol use in a placebo-controlled trial of naltrexone for smoking cessation. *International Journal of Neuropsychopharmacology, 12*, 589–597. doi:10.1017/S146114570800936X
- Pew Internet and American Life Project. (2006). *Online health search 2006: Most internet users start at a search engine online when looking for health information online*. Retrieved from http://www.pewinternet.org/files/old-media/Files/Reports/2006/PIP_Online_Health_2006.pdf
- Piasecki, T. M., Jahng, S., Wood, P. K., Robertson, B. M., Epler, A. J., Cronk, N. J., . . . Sher, K. J. (2011). The subjective effects of alcohol–tobacco co-use: An ecological momentary assessment investigation. *Journal of Abnormal Psychology, 120*, 557–571. doi:10.1037/a0023033
- Picard, R. R., & Cook, R. D. (1984). Cross-validation of regression models. *Journal of the American Statistical Association, 79*, 575–583. doi:10.1080/01621459.1984.10478083
- Qiu, B., Zhao, K., Mitra, P., Wu, D., Caragea, C., Yen, J., . . . Portier, K. (2011, October). *Get online support, feel better--sentiment analysis and dynamics in an online cancer survivor community*. Paper presented at the Privacy, Security, Risk and Trust (PASSAT) and 2011 IEEE Third International Conference on Social Computing (SocialCom), Boston, MA. doi:10.1109/PASSAT/SocialCom.2011.127
- Santhosh, L., Meriwether, M., Saucedo, C., Reyes, R., Cheng, C., Clark, B., . . . Schroeder, S. A. (2014). From the sidelines to the frontline: How the Substance Abuse and Mental Health Services Administration embraced smoking cessation. *American Journal of Public Health, 104*, 796–802. doi:10.2105/AJPH.2013.301852
- Selby, P., van Mierlo, T., Voci, S. C., Parent, D., & Cunningham, J. A. (2010). Online social and professional support for smokers trying to quit: An exploration of first time posts from 2562 members. *Journal of Medical Internet Research, 12*(3), e34. doi:10.2196/jmir.1340
- Shahab, L., & McEwen, A. (2009). Online support for smoking cessation: A systematic review of the literature. *Addiction, 104*, 1792–1804. doi:10.1111/j.1360-0443.2009.02710.x
- Talati, A., Keyes, K. M., & Hasin, D. S. (2016). Changing relationships between smoking and psychiatric disorders across twentieth century birth cohorts: Clinical and research implications. *Molecular Psychiatry, 21*, 464–471. doi:10.1038/mp.2015.224
- Toll, B. A., Cummings, K. M., O'Malley, S. S., Carlin-Menter, S., McKee, S. A., Hyland, A., . . . Celestino, P. (2012). Tobacco quitlines need to assess and intervene with callers' hazardous drinking. *Alcoholism: Clinical and Experimental Research, 36*, 1653–1658. doi:10.1111/j.1530-0277.2012.01767.x
- van Mierlo, T., Voci, S., Lee, S., Fournier, R., & Selby, P. (2012). Superusers in social networks for smoking cessation: Analysis of demographic characteristics and posting behavior from the Canadian Cancer Society's smokers' helpline online and StopSmokingCenter.net. *Journal of Medical Internet Research, 14*(3), e66. doi:10.2196/jmir.1854
- Vickerman, K. A., Schauer, G. L., Malarcher, A. M., Zhang, L., Mowery, P., & Nash, C. M. (2015). Quitline use and outcomes among callers with and without mental health conditions: A 7-month follow-up evaluation in three states. *BioMed Research International, 2015*, Article ID 817298. doi:10.1155/2015/817298
- Wang, X., Zhao, K., & Street, N. (2014). Social support and user engagement in online health communities. In D. Zeng, H. Chen, Y. Zhang, C. Xing, & D. B. Neill (Eds.), *International Conference on Smart Health 2014, Lecture Notes in Computer Science, Vol. 8549* (pp. 97–110). doi:10.1007/978-3-319-08416-9_10
- West, J. H., Hall, P. C., Hanson, C. L., Giraud-Carrier, C., Neeley, E. S., Barnes, M. D. (2012). Temporal variability of problem drinking on Twitter. *Open Journal of Preventive Medicine, 2*, 43–48.
- Zhang, M., Yang, C. C., & Gong, X. (2013). *Social support and exchange patterns in an online smoking cessation intervention program*. Paper presented at the 2013 IEEE International Conference on Healthcare Informatics (ICHI). Retrieved from <http://ieeexplore.ieee.org/stamp/stamp.jsp?arnumber=6680481>
- Zhao, K., Wang, X., Cha, S., Cohn, A. M., Papandonatos, G. D., Amato, M. S., . . . Graham, A. L. (2016). A multirelational social network analysis of an online community for smoking cessation. *Journal of Medical Internet Research, 18*(8), e233. doi:10.2196/jmir.5985
- Zhao, K., Yen, J., Greer, G., Qiu, B., Mitra, P., & Portier, K. (2014). Finding influential users of online health communities: A new metric based on sentiment influence. *Journal of the American Medical Informatics Association, 21*(e2), e212–e218. doi:10.1136/amiajnl-2013-002282

Appendices

APPENDIX 1. List of keywords used to retrieve alcohol-related posts (case insensitive):

AA, alcohol, alcoholic, alcoholics anonymous, alcopop, annihilated, bar, beer, bender, beverage, binge, black russian, blitzed, bottle, booze, brew, brewski, bud, budweizer, buzz, buzzed, cabernet, chardonnay, cider, cocktail, cold one, coors, drink, drinkin, drinkin', drinking, drunk, faded, g and t, gin, glass, hair of the dog, hammered, hangover, heineken, henny, hooch, hosed, hungover, inebriated, inhibition(s), intoxicated, keg, liquor, loaded, looped, malt, margarita, merlot, nightcap, pina colada, pinot grigio, pint, plastered, pub, quit killer, rum, sauce/d, smashed, soused, spirit, stewed, tanked/up, tequila, tipsy, under the influence, vodka, wasted, wine, wrecked

APPENDIX 2. List of keywords and phrases used for CD3 classifier (case insensitive):

12 step (12-step), AA (A.A.), alcohol being (a) problem, alcoholic, battle alcohol, black(ed/ing) out(s) (blackout), DUI (D.U.I), get (got/getting) wasted, problem(s) with alcohol