

This is a self-archived version of an original article. This version may differ from the original in pagination and typographic details.

Author(s): Soliman, Wael; Rinta-Kahila, Tapani

Title: Unethical But Not Illegal : Uncovering the Persuasive Messages Leveraged by Providers of the 'Real' Online Social Impressions

Year: 2018

Version: Accepted version (Final draft)

Copyright: © The Authors, 2018.

Rights: In Copyright

Rights url: <http://rightsstatements.org/page/InC/1.0/?language=en>

Please cite the original version:

Soliman, W., & Rinta-Kahila, T. (2018). Unethical But Not Illegal : Uncovering the Persuasive Messages Leveraged by Providers of the 'Real' Online Social Impressions. In ECIS 2018 : Proceedings of the 26th European Conference on Information Systems, Portsmouth, UK, 23-28 June (Article 172). European Conference on Information Systems.
https://aisel.aisnet.org/ecis2018_rp/172

UNETHICAL BUT NOT ILLEGAL: UNCOVERING THE PERSUASIVE MESSAGES LEVERAGED BY PROVIDERS OF THE “REAL” ONLINE SOCIAL IMPRESSIONS

Research paper

Soliman, Wael, Faculty of Information Technology, University of Jyväskylä, Finland,
wael.soliman@jyu.fi

Rinta-Kahila, Tapani, Aalto University, Helsinki, Finland, tapani.rinta-kahila@aalto.fi

Abstract

Due to the drastically elevated prominence of social networking sites (SNS), online social impressions such as views, comments, followers, subscribers, likes, and dislikes have become a valuable currency that translates to popularity, credibility, and even financial gains. Aside from machine-generated impressions, a growing industry known as crowdurfing utilizes human workers to provide “real” social impressions as-a-service. Although crowdurfing platforms are often seen as a clear example of deceptive conduct, they justify their business by leveraging well-crafted persuasive strategies and ethical appeals. Given the increasingly significant role of online impressions on shaping people’s views and opinions, the servitization of these impressions calls for a clearer understanding. To address this call, we set out to investigate 1) What persuasive strategies do crowdurfing agents leverage to promote their service offerings?; and 2) To what extent these offerings can be ethically justified? Our analysis reveals utilization of three key persuasive strategies – namely, educational messages, bragging messages, and reassuring messages. Moreover, we find that they use various ethical appeals which largely depend on the conception of what ‘real’ means. The theoretical and practical significance of these findings are discussed.

Keywords: crowdurfing, ethics theory, online social impressions, buying real views, qualitative content analysis

1 Introduction

“It’s always a cat and mouse game between YouTube and the unethical black hat social media marketing gurus ... It’s unethical but not illegal. At the end of the day I am not scamming people. I am providing people with a service that is unethical but helpful in the same sense. I have helped thousands of YouTubers get noticed on YouTube”.

– Manager of a crowdurfing firm¹

Inspired by the Wild West movies, information systems (IS) and security (InfoSec) scholars make a symbolic distinction between white hat and black hat activities (Mahmood et al., 2010; Warkentin et al., 2012). White hat activities represent the good, legal and/or ethically grounded deeds, while black hat activities represent the bad, illegal and/or ethically questionable ones. Such ethically questionable use

¹ The quote above is an excerpt from a correspondence chain with one of the firms included in our study. The correspondence was aimed to convey to the firm the ethical critique against the commercialization of social impressions, and to give the firm a chance to respond to such critique. Out of 10 firms we contacted (also the subject of this article), only one firm responded.

of technology has been a subject of concern for decades, exemplified by Mason (1986) who has identified four key aspects of information threatened by such use: privacy, accuracy, property, and accessibility. Today, such concerns could not hold higher relevance as our lives are increasingly immersed in online spaces, and social media platforms like YouTube, Twitter, Facebook, and LinkedIn are compellingly overthrowing traditional media as the main source of information and entertainment. The fundamental promise of these online services is appealing: virtually anyone or anything can become a success story, with no production companies, record labels, publishing houses, or other major sponsors needed! Another key difference to traditional media relates to the social element of these platforms. They typically provide their users with a powerful measure of popularity, quality, and relevance of the content: feedback. We call such feedback *social impressions* as they come in the form of likes, views, followers, retweets, reviews, or comments, etc. If a tremendous majority has given a positive feedback, then it is generally seen as a sign of good quality, and conversely, if a significant part of the feedback is negative, then the content is likely to be suspect at least. This measure of quality is a powerful one as it produces concrete instrumental value: for instance, online *reviews* and Facebook *likes* have been shown to have a significant effects on sales (Duan et al., 2008; Li & Wu, 2013; Zhu & Zhang, 2010). Also online popularity in itself may be a source of revenue, considering recent estimations that every 1000 views on a YouTube video could yield its creator royalties between 25 cents and \$4 (McIntyre, 2017).

However, the soaring popularity of online media is making competing and getting noticed online harder than ever. As content producers engage in fierce battles for audiences in these social media spaces, some lose their faith in the original promise of these services and resort to dubious marketing methods, such as those offering fake social impressions (Rinta-Kahila & Soliman, 2017) or reputation manipulation services (Farooqi et al., 2017). Indeed, the increasing importance of online popularity has given rise to online astroturfing and its crowdsourced form, crowdturfing, where impressions of popular support become an object of trade. Thus, selling fabricated social impressions such as likes, views, comments, or even online friends has become a commonplace business in the modern day. In essence, a crowdturfing agent arranges the buying and selling of online popularity by offering their services to customers and inviting crowd workers to execute these paid requests. The increasing eminence of these services is highlighted by the fact that there are now websites dedicated to ranking such services based on various KPIs, such as pricing, portfolio size, delivery time, customer support, to mention a few (e.g., buyviews-review.com).

Crowdturfing poses an undeniable threat to the accuracy of online information: not only it may mislead the public to follow content producers with limited merit, but it can also be used for spreading misinformation for commercial or political purposes. Such misinformation can have dangerous real-world consequences, e.g., fake reviews have been used to encourage unaware Internet users to use hazardous dietary supplements (Fayazi et al., 2015). These potential harms become even more unsettling when considering the fact that the service sold by crowdturfing agents is inherently deceitful: the service offering is a lie or an exaggeration meant to mislead its target audience. One could easily argue that this critique applies to any reputation management or advertising service, but crowdturfing differs inherently from such traditional services in one crucial way: crowdturfing activities do not appear as the voice of the service provider but rather as the voice of the people. While the advertisers of any product or service are logically expected to drive their own self-interests above anything else, the collective opinion of the internet public – whether manifested as reviews, likes, views, or other impressions – has been perceived as the independent and undominated voice of the people (Mitra & Watts, 2002).

Recent studies have revealed that the spread of fabricated social impressions may influence consumer behavior (Lappas et al., 2016), making understanding the logic in which the business operates critically important. While the current literature is mostly focused on the detection of fraudulent users and their activities, the human element in crowdturfing remains mostly uncharted. Moreover, the perspective of crowdturfing agents who orchestrate the trade of fabricated activities has received scarce attention (Farooqi et al., 2017). Specifically, we lack understanding of the communication strategies the agents use in marketing their services. Considering the controversial nature of such services, and that these

controversies are probably apparent to the potential customers, one could expect that the agents attempt to alleviate such concerns in their communication.

Thus, to gain a better understanding of how such ethically questionable platforms present themselves to their potential customers, we build on recent work on the topic (Farooqi et al., 2017; Lappas et al., 2016; Rinta-Kahila & Soliman, 2017). Specifically, we set out to tackle the following research questions: 1) What persuasive strategies do crowdturfing agents leverage to promote their service offerings?; and 2) To what extent these offerings can be ethically justified? We investigate these questions in the context social impressions traders who primarily target YouTube as a leading social media destination. We answer these questions based on the analysis of ten leading platforms that orchestrate the sale and delivery of social impressions. In our investigation, we focus on the persuasive strategies these agents leverage to promote and legitimate their service offerings. We find that these firms provide a range of fabricated social impressions for multiple target platforms, and leverage communication strategies that typically consist of three types of messages: educational messages, bragging messages and reassuring messages. We distinguish specific communication styles and techniques behind each message type. We find that the use of reassurance techniques depends, above all, on the agents' construction of what 'real' means. Moreover, we find that the agents use specific ethical appeals, drawing on consequentialist logic and the stockholder theory, as well as agents' self-serving interpretation of target platforms' codified rules.

This paper is organized as follows. In Section 2, we discuss the current literature on crowdturfing. In Section 3 we establish the theoretical background of this study by discussing ethical perspectives in IS use. In Sections 4 and 5 we describe our methods and findings respectively. Finally, we discuss the implications of our findings, and present some concluding remarks, in Sections 6 and 7 respectively.

2 Literature review

Crowdturfing represents the sinister side of crowdsourcing (Wang et al., 2012). Generally speaking, crowdsourcing has been acclaimed as a positive phenomenon in both academic and business contexts (e.g. Leimeister et al. 2009; Majchrzak & Malhotra 2013; Schlagwein & Bjørn-andersen 2014) as it generates economical value by harnessing the under-utilized resources of the crowd. However, consequences are less desirable when the logic of crowdsourcing is applied to astroturfing. Astroturfing is an old and simple idea: secretly paying individuals to publicly demonstrate their support. Beder (1998) describes astroturfing operations as "artificially created grassroots coalitions", or more formally as "grassroots program that involves the instant manufacturing of public support for a point of view in which either uninformed activists are recruited or means of deception are used to recruit them" (p. 21). Following the same logic, in crowdturfing, the members of the crowd are paid to leave a fake social impression, e.g., writing a positive review of an unread book (Rinta-Kahila & Soliman, 2017), or pretending to be a fan of a particular Instagram account (Price, 2014).

Like in crowdsourcing, crowdturfing campaigns consist of collections of specific tasks, and they typically involve three main actors: *customers*, who order and pay for the tasks; *workers*, who execute tasks in exchange for financial compensation; and *agents*, who market and orchestrate the trade of such tasks (Wang et al., 2012). Agents appear usually in the form of a two-sided platform (Eisenmann et al., 2006), i.e., websites that provide services and employment opportunities for customers and workers, respectively. Using crowds is by no means the only available method of manipulating online appearances: for instance, certain reputation management services provide advice on how to remove or undermine inconvenient customer feedback (e.g., www.udemy.com). Online manipulation has been also practiced by leveraging automated bots (Abokhodair et al., 2015; Yao et al., 2017). However, thus far utilizing crowds has offered a significant advantage over 'automated' turfing: since the activities are executed by real human workers, they appear more genuine and are thus more difficult to detect and stem (Wang et al., 2012). Current literature on crowdturfing is still scarce but the academic interest toward the topic is growing. Most of the prior work resides in the domain of computer science and focuses on characterizing the phenomenon and detecting crowdturfers.

Typical tasks involve social media manipulation, such as buying views or likes, has been found as the most common type of crowdturfing (56% of campaigns), followed by paid sign-up tasks to services (26%), search engine optimization (SEO, 7%), and paid votes (4%) (Lee et al., 2013). Worryingly similar methods have also been leveraged for spreading political propaganda (Han, 2015; Pham, 2013), executing coordinated cyberattacks (Lee et al., 2013), and conning money for non-existent projects (Siering et al., 2016). Social media manipulation campaigns range from buying views on YouTube and retweets on Twitter to purchasing positive reviews for books on Amazon (Shaffer, 2013) or hotels on TripAdvisor (Lappas et al., 2016). The customers who order such campaigns may be, for instance, artists who want to get their work noticed, businesses looking to increase their revenues or defame their competitors, or political actors who attempt to shape public opinions.

Customers are typically based in English-speaking Western countries, predominantly from the US, while crowdturfing workers come mostly from South Asian countries, although many of them are US-based too (Lee et al., 2013). For workers, executing crowdturfing tasks is a source of income that can be remarkably profitable one, for instance, some Bangladeshi crowdturfers have been able to generate earnings that exceed the national average income (Lee et al., 2013). This implies that crowdturfing can become a full-time job for some, and consistently it has been found that such professional workers are actively involved in multiple campaigns simultaneously, while casual workers take upon tasks only occasionally as a side job (Lee et al., 2013). In addition, when turfing content or messages like tweets and posts, some workers act as middlemen who spread the content directly from its origin, i.e., the customer, and this content will then be spread onwards by other workers (along with unaware members of the Internet public who agree with the message).

Agents' roles vary from passive to active depending on the platform. Due to the controversial nature of such tasks, distinguished crowdsourcing platforms tend to explicitly prohibit the use of their platform for crowdturfing. For instance, Amazon's Mechanical Turk, a popular crowdsourcing platform, has made conscious efforts to stem crowdturfing campaigns from their service. However, others overlook such activities and let them operate uninterrupted. For example, almost 90% of all the tasks sold on the two largest Chinese crowdsourcing sites were found to be linked to crowdturfing campaigns (Wang et al., 2012). On the other hand, many agents, commonly referred to as online black-hat marketplaces, have specifically positioned themselves as crowdturfing service providers, openly advertising purchasable social impressions as their main offering. Some of these platforms are strictly focused on only one type of task on one target platform (e.g., PaidBookReviews.org sells only book reviews for Amazon listings), others provide a wide range of offerings ranging from subscriptions to YouTube channels to endorsements for LinkedIn profiles and SEO in Google (e.g., YTview.com). Often agents sell the services under their own name utilizing workers behind the scenes but in some cases an agent platform acts as a marketplace where customers can buy tasks directly from freelancer workers (e.g., SEOClerks.com).

In addition to characterizing the phenomenon, previous work on the topic attempts to provide concrete tools for curbing crowdturfing through developing algorithms and machine learning tools that can detect and identify crowdturfers and their malicious activities. Latest such work involves crowdturfing detection, traffic classification, and monitoring systems through security algorithm optimization, encryption schemes, abnormal behavior discovering, and ciphertext updates (Li et al., 2017). Crowdturfing detection methods may base their analyses on account characteristics, yielding detection rates as high as 97.35% (Lee et al., 2015). However, since turfing are real human workers, their account characteristics can be very similar to those of sincere users, and thus some focus on detecting the target objects that are under manipulation attempts (Song et al., 2015). In the domain of news reporting, similar AI-based methods have been touted as a highly potential countermeasure against fabricated news content, i.e., fake news (Bloomberg, 2017). While such initiatives are promising, they might also mean "the beginning of an automated arms race" (Snow, 2017), as fake news producers too could adopt similar machine learning algorithms. Likewise, in context of online feedback manipulation, the development of AI has started to rapidly bridge the gap between automated and human-produced fabrication (Price, 2017; Yao et al., 2017), raising concerns on whether AI-based flagging and protection mechanisms can keep up with corresponding novel methods of generating misinformation. Although using AI to generate complex forms of fake online feedback such as reviews has not taken off yet, it might be a growing concern

in the near future (Price, 2017). Still, as of today, leveraging humans-based crowdturfing remains a significant means of manipulating public opinions.

In sum, most current work on crowdturfing has characterized the network of actors who participate into the business as well as developed detection methods of crowdturfing workers and tasks. However, the ethical and behavioral aspects of the phenomenon have received less attention, and thus approaches outside the computer science domain have been called for (Li et al., 2017). In addition, prior work has examined the types of services offered by black-hat marketplaces (Farooqi et al., 2017; Wang et al., 2012) but thus far, little attention has been given to *how* they offer them. Thus, in line with recent investigation on the ethical rationales behind crowdturfing (Rinta-Kahila & Soliman, 2017), we turn to examine how crowdturfing agents, or what Beder (1998) calls ‘public relations firms’, persuade potential customers to buy their services and how they discuss the ethical implications of the business.

3 Theoretical background: What is ethical IS use?

Revolutionary advances in ICTs have posed serious questions regarding the moral compass guiding technology use behavior in the information age. These questions include “what is ethical”, what is legal”, “on what grounds an action is right or wrong”, and so forth. Such quandaries have undoubtedly occupied information systems and security scholars for decades (see, e.g., Mason, 1986; Moor, 2001; Rinta-Kahila & Soliman, 2017; Siponen & Vartiainen, 2002). Already in 1986, Richard Mason identified four key ethical issues in the information age: 1) privacy, 2) accuracy, 3) property, and 4) accessibility of information. These aspects have provided a foundation for studying information technology ethics ever since (Peslak, 2006), and they still hold significant relevance (Chatterjee et al., 2015). Each of the four issues concerns a potential point of information vulnerability that may be threatened by unethical online behavior, thus reflecting prominent IS security issues. For instance, consumers’ willingness to trade their private information for increased convenience has put the spotlight on large corporations like Google and Amazon, as critical voices have questioned the ethics of these companies’ behavior (Zuboff, 2015), arguing that they are posing a threat to the privacy of information. Likewise, buying fabricated online social impressions can be considered as a threat to the accuracy of information, not least because decreased authenticity of online information may misinform decision-makers, whether they are consumers, merchants, or politicians.

Indeed, crowdturfing campaigns threaten the credibility and usefulness of their target platforms by derogating their trustworthiness, possibly resulting in negative financial impacts. In addition, users of the platforms may become misled by the deceptive information and end up making harmful decisions, ranging from buying a bad product to adopting unhealthy life habits (e.g., avoiding recommended vaccinations) or otherwise endangering their health (Fayazi et al., 2015). These effects may be amplified by the unaware ‘following crowds’, who accept the fabricated information and continue spreading it. Hence, on a larger scale, crowdturfing may erode the accuracy of information available to the Web at large, which may then translate into undesired real-world effects (e.g., mass confusion), since online interactions are increasingly interacting with events occurring in the physical world. Thus, it is safe to argue that participating into crowdturfing entails some burning ethical conundrums.

To resolve these dilemmas, we turn to ethics theory for guidance. It has been noted that “ethics theory offers an understanding about competing views in ethics and the difficulties encountered in decision-making in human life” (Siponen & Vartiainen, 2002, p. 427). However, literature on ethical theory is rich, diversified, and rests on different moral philosophical schools, including ethical egoism, utilitarianism, deontology, feminist ethics and virtue ethics just to mention a few (Jones et al., 2007). However, despite the many differences and divergences among these ethical foundations, “they converge on essential point—their emphasis on concern for others over self-interest” (ibid, p. 140). Smith (2002) provides a useful normative ethical framework that classifies the various theories into three levels that he termed: *traditional philosophical ethics*, *business ethics*, and *codified rules*.

The traditional philosophical schools of ethics, namely, consequentialist (or teleological) and categorical (or deontological) theories, provide simple, generalizable guidelines about right and wrong. The consequentialist school argues that the rightness of an action is determined by its resulting consequences:

from alternative courses of action, the one that produces most good for a given group, is ethically right. Thus, a content producer could justify buying YouTube views for her video as acceptable because that will give the content a chance to be seen, leading to positive outcomes for their referent group. The categorical perspective is often portrayed to be in sharp contradiction with the former view as it asserts that the ethics of a behavior are determined by the intent behind it, and especially whether the action condones with some generally agreed categorical principles, such as “deception is wrong”. Thus, using the categorical logic it would be difficult to ethically accept buying views since it is based on deceiving the Internet public.

Business ethics offer a more contextualized set of principles in the form of ethical guidelines for managers and other decision-makers in companies, which rest on three prominent theories: the stockholder theory, the stakeholder theory, and the social contract theory. First, based on Milton Friedman’s doctrine, the stockholder theory asserts that decision-makers should only take such actions that maximize their stockholders’ interests. This would give any online business relatively free hands to operate as their only concern would be producing profits to the owners (albeit within the frames set by the law). By contrast, the stakeholder theory highlights that the interests of every stakeholder, i.e., those affected by the company’s business, should be taken into account in decision-making. Thus, companies engaged in fraudulent marketing, like crowdturfing, would have difficulties in justifying their behavior drawing on this logic. For instance, those content producers who play fair and do not use crowdturfing services could be considered as negatively affected stakeholders since they become disadvantaged due to others’ unethical behavior. In a similar vein, the social contract theory emphasizes the fulfilment of justice through collective avoidance of fraud and deception, positing that businesses should aim to maximize the advantages of their actions and minimizing their disadvantages (Smith 2002, p. 15).

While each perspective discussed above arms one with a behavioral guideline, it comes with a normative stance that can always be counter-argued by leaning to the opposite perspective. Thus, societies have developed codified rules that materialize as laws, regulations, codes of conduct, legal contracts, terms and conditions, and other written univocal norms. In the context of digital services these usually manifest as terms and conditions or disclaimers that the user must approve before using the service. This is also the case with platforms like YouTube, where posting a video requires registering to the service, which further requires accepting its terms and conditions. For instance, in its Terms of Service, YouTube specifically forbids deceptive or automatic means of generating views, including “*purchasing views from third-party websites*”.

4 Methodology

In this study, we utilize the analysis of textual content method (Hsieh & Shannon, 2005). Analysis of text is at the heart of qualitative research. In fact, Gephardt argues that “qualitative research starts from and returns to words, talk, and texts as meaningful representations of concepts” (Gephardt, 2004, p. 455). It is noted that “research using qualitative content analysis focuses on the characteristics of language as communication with attention to the content or contextual meaning of the text ... Text data might be in verbal, print, or electronic form and might have been obtained from narrative responses, open-ended survey questions, interviews, focus groups, observations, or print media such as articles, books, or manuals” (Hsieh & Shannon, 2005, p. 1278). Adopting the content analysis approach for analyzing websites has been used extensively in various fields. Examples include analyzing how websites coordinate the selling of sex online (Castle & Lee, 2008), advocate harmful eating disorder practices (Norris et al., 2006), mobilize youth involvement (Gerodimos, 2008), as well as how websites advertise the selling of e-cigarettes (Grana & Ling, 2014).

Qualitative content analysis is defined as “a research method for the subjective interpretation of the content of text data through the systematic classification process of coding and identifying themes or patterns” (Hsieh & Shannon, 2005, p. 1278). Coding is a crucial component of such analysis. Codes in qualitative content are more likely to transpire based on the analyst’s careful reading of the text under study, as opposed to using of algorithms that generate codes from the data automatically that might be more common in quantitative or summative approaches. As such, the process is characterized by a

search for meaning in the text, not merely word counts (Morgan, 1993). Hsieh and Shannon (2005) remind us that the qualitative approach is similar in spirit to grounded theory (Glaser & Strauss, 1967) and “is usually appropriate when existing theory or research literature on a phenomenon is limited” (Hsieh & Shannon, 2005, p. 1279).

4.1 Data collection

Readily available (online) data represents a rich, yet under-utilized approach to data collection for IS scholars (Ghazawneh & Henfridsson, 2013). Moreover, it is acknowledged that when the studied phenomenon poses a challenge regarding the production of primary data, researchers are encouraged to utilize unconventional approaches to data collection (Mahmood et al., 2010; Warkentin et al., 2012). We collected readily available data from ten views-buying websites that were identified by buyviews-review.com as the top-ten agents in the business. First, we extracted data about the companies’ background, service offering, pricing, marketing communications, arguments for legitimacy and credibility, and acknowledgement of ethical issues. We tabulated this data into an Excel file for initial analysis. Second, we extracted data into a text file from two specific sections of each website: their “About us” section and “Frequently Asked Questions (FAQ)” section. The reason for selecting these specific sections was that they typically provide a comprehensive overview on the business, along with arguments for legitimacy and credibility, and company’s take on potentially most relevant customer concerns.

4.2 Data analysis

We utilized thematic analysis to establish the main categories or themes. Similar to other qualitative approaches to data analysis, thematic analysis relies on an iterative process of careful reading of the data, generating codes that best describe the text, searching for themes or patterns in the data, and developing a coherent report that is most faithful to the data (Braun & Clarke, 2006). The qualitative data analysis tool Atlas.ti assisted us in organizing, tabulating, and visualizing the code maps, while the actual coding was done entirely by the researchers. Both descriptive and interpretive coding have been utilized. Whereas descriptive coding was largely ‘data-driven’; ‘theory-driven’ coding assisted us in the development of interpretive coding (Fereday & Muir-Cochrane, 2006). An example of the former includes our classification of the “**Bragging Messages**” – used by the platforms as part of their communication strategies – into “*We-Have Claims*”; “*We-Are Claims*” and “*We-Are-Superior-To-Competition Claims*”. An example of the latter (i.e., theory-driven coding) includes our development of the “**Ethical Appeals**” category – largely influenced by our understanding of the different schools of thoughts on ethics theories (Jones et al., 2007; Rinta-Kahila & Soliman, 2017; Smith, 2002) – into “*Appeal to Codified Rules*”; “*Appeal to Stockholder View*”; “*Appeal to Rhetoric*”; and “*Confessionals*”. In addition to utilizing data-driven and theory-driven coding, our analysis was influenced by Van Maanen’s (1979) distinction between ‘first-order’ and ‘second-order’ concepts. Specifically, first-order concepts in our analysis are seen to be more reflective of the empirical data, and represent as close as possible the actual text generated by the studied firms themselves. As such, they are more descriptive in nature. For example, “we have claims” reflect messages the firms use to convey objects they possess. On the other hand, second-order concepts are more analytic in nature, as they reflect the analysts’ (us researchers) interpretation of certain concepts. This means that it is our own interpretation, based on our analysis, that firms use bragging messages as a persuasive strategy, which is something the firms themselves might deny or are unaware of.

5 Findings

In this section, we present a brief overview of the key findings our analysis, after which we address our research questions.

5.1 Overview

Table 1 gives a holistic view on our findings. While two firms do not disclose information about their physical location, we find that six firms are located in the US, one in the UK, and one in Thailand. In

terms of services offered, most firms provide a rich portfolio of social support services, including views, shares, followers, subscribers, comments, likes and even dislikes. A minority of the firms target a single platform, whereas it is common that the service portfolio targets a wide spectrum of the most popular social networking platforms like Facebook, LinkedIn, Twitter, and YouTube, just to mention a few. Surprisingly, these services come at rather affordable prices. For instance, the price range for 1000 YouTube dislikes is between \$3.60 on QQTube and \$79 on Devumi. We note that comments are rarest and most expensive item on their menu (ranging between \$15 on YTview and \$189,95 on Marketing-Heaven for 100 comments), probably due the labor needed for its orchestration. Overall, the studied firms leverage a combination of the three persuasive strategies – educating, bragging and reassuring – manifested by their constituent components. As we will elaborate in the following sections, each firm has its own mix of persuasive messages, with varying emphasis on different components. For instance, a firm may dedicate its focus chiefly to educational messages, while relaying reassuring and bragging messages to the background (e.g., AudienceGain). By contrast, other firms may put their emphasis on reassuring, giving far less emphasis to educational messages (e.g., QQTube).

Firm	Location	Target Platforms	Social Impressions 'on Sale'							Key Communications Strategies							Ethical appeals				What Is Real?		
			VI	LI	DI	SH	FA	SU	CO	(1) Bragging Messages			(2) Educational Messages			(3) Reassuring Messages							
										We are superior	We are claims	We have claims	Gen. Info.	Why buy	How it works	Quality	Security	Persistence	Authenticity	Rhetoric		Stockholder	Codified Rules
500views.com	USA	YouTube; Instagram	✓	✓					✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓		
audiencegain.com	USA	YouTube; Facebook; Twitter & Others	✓					✓		✓		✓	✓				✓	✓	✓				✓
buildmyviews.org	n/a	YouTube only	✓	✓					✓	✓	✓	✓	✓	✓			✓	✓	✓	✓	✓	✓	
buyrealmarketing.com	USA	YouTube; Twitter; Facebook & Others	✓	✓		✓	✓	✓	✓	✓	✓		✓	✓	✓		✓	✓	✓	✓	✓	✓	✓
buyviews.co	USA	YouTube; Dailymotion	✓	✓	✓	✓	✓	✓		✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
devumi.com	USA	YouTube; Twitter; SoundCloud & Others	✓	✓	✓	✓	✓	✓	✓	✓	✓		✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
qqtube.com	USA	YouTube only	✓	✓				✓	✓		✓		✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
socialshop.co	Thailand	YouTube; Instagram; Facebook & Others	✓	✓	✓	✓	✓	✓	✓	✓		✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
themarketingheaven.com	UK	YouTube; Twitter; Facebook & Others	✓					✓		✓		✓		✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
ytview.com	n/a	YouTube; Twitter; Facebook & Others	✓	✓	✓	✓	✓	✓	✓			✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓

VI: Viewing DI: Disliking FA: Favoriting SH: Sharing
 LI: Liking SU: Subscribing CO: Commenting

Table 1. Summary of findings

5.2 Key persuasive strategies

Our analysis shows that all firms use a combination of three key communication strategies or persuasive messages. We labeled them educational messages, bragging messages, and reassuring messages (For an overview of the coding map, see Figure A1 in the Appendix). Interestingly, these three persuasive strategies may be generally observed in most legitimate business communications. However, it is in scrutinizing, and attention to details, that one is able to uncover areas of inconsistency and reasons to question their legitimacy, as we will point out next.

5.2.1 Educational messages

Educational messages mainly reflect a communication strategy in which the firm transmits three types of information. We label them ‘general information about social media’, ‘importance of buying noticeability’, and ‘how our service works’.

The ‘**general information about social media**’ messages are utilized by six out of the ten firms. As the name implies, this strategy entails providing potential clients with introductory information about the business and social media campaigns in general. Specifically, such messages aim to educate about the key target platforms in this domain (e.g., YouTube and Facebook), how these platforms operate, and

most importantly they attempt to normalize crowdturfing business. For instance, SocialShop notes the following *“Becoming the next viral sensation doesn’t have to be so hard. Did you pour hours into your videos only to struggle to get the traction you hoped for? ... Buying Youtube views isn’t shady, it’s merely a way of kickstarting the popularity of awesome videos”*.

Closely related is leveraging the **‘importance of buying noticeability’** strategy, which is utilized by seven out of the ten firms. This strategy primarily reflects messages intended to explain to potential customers why buying the firm’s services (views, likes, comments, etc.) is so important. The central messages here rest on the argument that buying the services has a tremendous positive impact on the buyer’s *“fame”*, *“popularity”*, *“credibility”*, and most importantly, on *“social proof”*. Explaining why buying views is important, BuyViews notes the following: *“Buying YouTube Views is a quick, easy and affordable method of kickstarting your video and gaining a powerful force known as Social Proof. With more views, your video will attract more views naturally, from ranking better in search results to getting shared more often. People are drawn to popular things and that’s exactly what your video will be after buying views”*.

The third educational communication strategy is what we labeled **‘how our service works’**. Interestingly, all firms (in our sample) have leveraged this strategy. It typically involves information regarding the service buying process, and is probably the closest that it gets to how the service operations run behind the curtain. Information disseminated via these messages usually describe how to place an order, delivery time, conditions for refund or replacement, and most importantly where do views, likes, comments come from. For instance, Devumi.com writes: *“Our Likes and Comments are added manually by our dedicated and knowledgeable YouTube team. This ensures that Likes and Comments will be positive, as well as that comments will be related to your video and encouraging of discussion”*.

5.2.2 Bragging messages

Bragging messages mainly reflect a communication strategy in which the firms make self-promotional claims. We divided those into ‘superiority over competition’ claims; ‘we are’ claims; and ‘we have’ claims.

The key distinguishing factor between the **‘superiority’** claims on the one hand, and the ‘we are’ and ‘we have’ claims on the other, is that the former makes a claim which invokes a comparative process between the business in question and its peers, the outcome of which is always in the firm’s favor. The two latter claims do not. In other words, ‘superiority’ claims typically leverage statements like *“we are the only provider ... of guaranteed non-botted real views.”* (MarketingHeaven).

By contrast, **‘we are’** statements, like *“you have checked that we are a credible company, the next step would be to purchase YouTube comments”* (BuildMyViews); and **‘we have’** statements like *“We have optimized methods over the past 3 years to keep costs at the lowest possible. Our production-costs drop as our production-volume and connections increase”* (BuyRealMarketing), do not invite such comparison with peers.

5.2.3 Reassuring messages

The third persuasive strategy we identified relies on the use of reassuring messages. We could distinguish between four types of messages reassuring potential clients regarding four distinct service attributes: quality, security, persistence, and authenticity.

Quality-focused messages are utilized by seven out of the ten firms, and they are intended to reassure potential customers of the high quality of *“the service”* in general, its *“design”*, the *“user experience”*, as well as to emphasize *“customer satisfaction”*. The following message by SocialShop sums up quality-oriented communication: *“In business, price timing and quality is key and we at SocialShop understands that which is why we deliver this on a daily basis to every client that comes to us. This allows our clients to buy YouTube views cheap knowing that you are receiving top quality”*.

Security-focused messages are utilized by seven out of the ten firms, and they are intended to reassure potential customers that the service delivery is *“secured”*, *“safe”* and *“confidential”*. In fact, reassuring the *“anonymity”* or *“secrecy”* of the service seems like an indispensable requirement for the success of

such these transactions. The following messages reflect the latter notion: *“No one will know you’ve hired a professional company to stimulate your social media success”* (MarketingHeaven); and *“We never share or sell your confidential information to third parties. Your orders and purchases through us are not shared with anyone else”* (YTview).

Persistence-focused messages are utilized by seven out of the ten firms. We labelled it as such to highlight the emphasis on continuity. Specifically, these messages are intended to reassure potential customers that buying the service will not cause harm to their target accounts. Leveraging these persuasive messages emphasize that potential customers should rest assured that their target accounts – be it on Facebook, YouTube, Twitter, etc. – will not be *“suspended”*, *“banned”*, or *“removed”* as a consequence to buying the service. For instance, QQTube reassures its customers that *“because our views are from real people, your account is safe and we have never caused a customers’ account to be suspended”*. Similarly, 500views reassures its customers that *“your video or channel will not get banned. We have never experienced any of our clients get banned”*. Moreover, these messages also aim to convince that the bought impressions will persist on the target account: views will not be pruned off by YouTube and subscribers will not unsubscribe.

Authenticity-focused messages are probably the most important persuasive communications in this type of business. In fact, they have been utilized by all the firms in our study. These messages are intended to reassure potential customers that the services sold (e.g., views, likes, comments, etc.) are *“real”*, *“not fake”*, *“come from real people”*, and *“not from bots”*. For instance, Devumi emphasizing the authenticity of their subscriber service notes that *“they are all real, however not all of them are active. We deliver both active and inactive YouTube Subscribers ... The purpose of our subscriber service is to increase your perceived popularity to convince other YouTube users to subscribe”*. Corroborating this line of reasoning, BuyViews declares that *“the YouTube Views we deliver are always from REAL PEOPLE. We have never and will never deliver fake, bot or computer-generated views – it goes against our code of honor and puts your video at risk”*.

What is real? Intimately related to authenticity reassurance messages is how the studied firms use the terms *“real”* and *“fake”*. There is a unanimous consensus among those firms that they only provide real services, and that that using bots is something they cannot tolerate, as emphasized earlier in the previous section, that using bots is against their *“code of honor”*. Astonishingly, the same firm also declares that *“YouTube Subscribers are usually inactive users, so they are unlikely to watch your videos. However, they do look real and will drive up your popularity and social credibility”* (BuyViews). The narrative the studied firms adopt points to a perplexing definition of real. As long as there is a human being responsible for generating the likes, dislikes, comments, etc., then they are real. In fact, some firms allow customers to customize their own real comments to ensure that the comments meet the expectations. For instance, BuyViews advertises *“what type of monsters would we be if you couldn’t customize your comments? Of course you can. When ordering comments, just enter any custom comments in the Text Box provided”*. Eventually, what counts as real is what *“appears to be real”* or what *“appears to be genuine”*. A lie is true as long as it is told by a human being!

This quandary cannot be resolved without understanding the ethical grounds on which this type of business rests, which brings us to the second research question of this study. In the next sections, we identify and critically assess the ethical appeals utilized by the studied firms.

5.3 Ethical appeals

To answer the second question – to what extent the trade in social impressions may be ethically justified – we relied predominantly on ethics theory (Jones et al., 2007; Rinta-Kahila & Soliman, 2017; Smith, 2002) in our interpretation of the firms’ narratives. Specifically, the distinction Smith (2002) makes between philosophical ethics, business ethics, and codified rules acted as a central sensitizing device in our analysis. We find that the agents appear to acknowledge that ethical concerns might discourage their customers from buying their services, as each actively employs specific ethical appeals in their advertising. Our analysis points to four distinct narratives or appeals that we labeled: appeal to stockholder view, appeal to codified rules, appeal to rhetoric and confessionals.

First, appeal to stockholder view represent narratives emphasizing that the firm's ultimate goal is to serve its (paying) customers. This type of appeal or message portrays 'achieving the customer goals' and 'customer satisfaction' as the agent's ultimate goal. For instance, 500Views announces that the company "*was developed to help increase your video's popularity on YouTube. We want to get your video noticed by Corporations and would like to see you become successful through what you do best on your video.*" Similarly, QQTube writes: "*We want you to succeed in your YouTube marketing career and we will do everything possible to meet your needs.*" Eight out of the ten firms adopt this narrative (see Table 1). In general, this finding demonstrates that the studied firms justify their service offerings mainly by aligning their goals with those of their paying customers, with no apparent regard to the implications of their service on the wider audience. As such, we found no evidence to imply that the firms utilize either stakeholder or social contract as a basis for their justification.

Second, six out of the ten firms appeal to the codified rules of the target platform, i.e., YouTube's Terms of Service. This is interesting since YouTube explicitly forbids actions like fabrication of views. Nevertheless, the agents make statements like: "*our services are 100% compliant with YouTube's Terms and Policies*" (BuyViews), "*our methods are compliant with Terms of the networks we work on*" (Buy-RealMarketing), and "*our services are completely compliant with YouTube's Terms and Conditions*" (Devumi). It appears that YTVview attempts to distract concerned clients by referring them to Google's Terms of Service (which do not contain YouTube's terms). Moreover, Devumi circumvents issues with breaking YouTube's rules by referring to their extensive network of partners: "*We acquire real views from a large network of partnered websites that display your video to their visitors, so our YouTube Views services does not violate any of YouTube's Terms and Conditions*".

The third type of appeal is rather interesting since it justifies the service offering merely rhetorically, or by simply stating that 'appropriateness' of the service as a 'matter of fact'. For instance, the sites claim that "*buying YouTube views isn't shady*" (SocialShop); that "*there is nothing wrong with getting a promotional campaign set up for your video*" (MarketingHeaven); and YTVView's statement that "*We never sell any fake YouTube views. Our views (and other services) are 100% organic, real human*" (see the previous discussion on 'what is real'). We interpret these statements as unfounded ethical justification, and they typically aim to provide potential customers with general assurance that they need not worry about the ethical implications of buying social impressions. Seven out of the ten firms adopt this approach (see Table 1).

Finally, we found that four firms adopt an approach that is best described as confessionals. Narratives belonging to this approach usually reflect a revelation of sorts; an admission of something that stands in contradiction with one or more of the three persuasive narratives discussed earlier. These messages seem to emphasize the transparency and sincerity of the agent. Consider Buy Real Marketing for instance: "*...the Followers, Views and Play that we offer are in large part from inactive user accounts. Yes, we admit it, they are mostly inactive. This means that they will not engage, like, comment or share on your social profiles. They are for credibility and vanity purposes*". Others are less dramatic with a realist tone, and can be seen as disclaimers of service, e.g., "*Results vary widely based on video content and quality, so we can't guarantee every campaign will be a raging success*" (Buyviews).

6 Discussion

In most modern economies, corporations have been seeking professional assistance to help them popularize their brand. These public relation arrangements typically result in customized campaigns targeting certain customer segments, and using tailored persuasive messages accentuating the benefits of the products or services they offer. One of the core ingredients of such practice is transparency. Obscuring the fact that a certain message has a (paying) sponsor shifts the discussion from traditional marketing and public relations to what the literature describes as fabricated public relations or astroturfing (Beder, 1998; Goldschein, 2011; Kraemer et al., 2013). Just like astroturfing in its traditional form, crowdturfing campaigns are orchestrated as genuine online activities performed by real people and this makes such campaigns especially challenging to expose. Even when being able to detect and abolish such activities, it is difficult to dissuade the involved parties from re-engaging into same activities, especially if those

involved use ethically justified arguments. Therefore, we believe that understanding the common strategies used for marketing and selling these services is undoubtedly useful in the battle against the trade of fake social impressions. Specifically, being aware of the ethical grounds on which crowdturfing business rests allows us to scrutinize its soundness and develop counter arguments.

Uncovering the communications discussed above reveals that these firms do acknowledge the ethical dilemmas in their business and are actively trying to mitigate or circumvent them in order to attract customers. When doing this, they are practicing implicitly collective construction and circulation of a specific online myth, namely, that it would be next to impossible to rise to fame in social media without the use of crowdturfing services. We observed that this myth was consistently repeated on the websites of several firms included in our study. The myth, along with specific ethical appeals, is then leveraged as a justification (or a scapegoat) for the existence and necessity of their services. Most notably, the agents (a) draw on a consequentialist logic by focusing exclusively on self-interest, (b) rhetorically assure their customers that there is ethically justified, and (c) appeal to a self-serving interpretation of codified rules. At the same time, these communication approaches and the flexible interpretation of the target platform's codified rules allow the agents to dismiss the viewpoints of other stakeholders and the negative consequences they and the rest of the Internet community have to bear because of their service offerings. Thus, the firms clearly embrace a self-interest leaning ethical perspective as opposed to taking others into consideration (Jones et al., 2007).

By uncovering persuasion strategies used in the trade of fabricated social impressions, several theoretical implications are worth highlighting. First, from an ethical theory perspective, our findings reveal that crowdturfing agents generally adopt a consequentialist, stockholder rationale, with a clear focus on serving the paying customers, with little or no regard to other stakeholders (Smith, 2002; Jones et al., 2007). This is reflected by the agents' extensive emphasis on aligning their goals with those of their paying customers, and the adoption of various appeals to support this objective. The stockholder view is especially apparent in the persuasive messages of the agents, as they educate their customers about the "reality" of the business of getting online visibility, brag about their services in an exaggerated manner, and actively reassure the customers about the high standards of their service. Consistent with this classical economics view, crowdturfing agents may have grounds to argue that business ventures exist "to maximize the present value of profits over the long term" (Armstrong & Green, 2013, p. 1922), and that they are doing exactly that by nurturing and serving their paying customers. This perspective, however, has long been criticized for its egoistic, self-centered orientation. For instance, research on corporate social responsibility (CSR) and irresponsibility (CSI) considers ignoring the wider implications of corporate decision-making to be unethical. Considering the wide recognition of the consequentialism "shortcomings", Moor (1999) has proposed advancing the consequentialist thought by incorporating an impartiality test, or what Bernard Gert (1998) calls the 'blindfold of justice'. This hypothetical blindfold "removes all knowledge of who will benefit or will be harmed by one's choices" (Moor, 1999, p. 67), thus subjecting decision makers to the consequences of their own choices. If the crowdturfing agents indeed subject their business decision to the impartiality test, they might start questioning the justness (i.e., ethical grounds) on which their service offerings stand.

Second, our findings pinpoint the emergence of a novel phenomenon surrounding the utilization of fake public relations by 'the layperson'. Traditionally, literature on fabricated public relations (e.g., astroturfing) has focused on studying how institutions such as governments (Han, 2015), organizations (Cho et al., 2011; Kraemer et al., 2013), or even political parties (Beder, 1998), have been utilizing fake grassroots activities to serve their (hidden) agendas (see also, Goldschein, 2011). Such new direction rings a warning signal for our modern connected societies.

Third, and closely related, existing research on fabricated public relations points to two central deceptive strategies utilized by orchestrating firms (i.e., the agents): confusion and fronting (Cho et al., 2011). On the one hand, confusion refers to agents' strategies aiming at distorting public opinion and manipulating their perceptions surrounding a particular subject (e.g., that global warming is a hoax). On the other hand, fronting refers to attempts made by the agents to organize what appears to be a genuine frontline group (hence the name) but is in fact dedicated to defending "hidden [paying] corporate interests" (ibid,

p. 582). Our findings show that the sale of confusion is a central strategy in modern crowdturfing services, as demonstrated by their ability to provide massive amounts of people-generated likes, dislikes, comments (positive or negative), as well as views and followership. Fronting, by contrast, seems quite unlikely. Fronting requires courage and willingness from the members of the fronting group to engage in tough debates, while often adopting an unpopular side (e.g., the artificial sweetener has no side effect, Beder, 1998). We could not trace the marketing of such elaborate strategies in the data. In fact, in their confessionals, the crowdturfing agents inform their (potential) customers that their crowd-workers are mostly inactive users who will not engage with the customer social profile beyond the requested task.

Overall, our findings related to the myth being circulated in the websites and the agents' conception of what 'real' is point out to a shared construction of online reality. It may be that the ethical perspectives crowdturfers lean on serve as the foundation of this reality, where those assumptions that support crowdturfers' businesses are accepted as facts without criticism. Recent media scandals related to questionable online behaviors, ranging from copyright infringements to online harassment and irresponsible tweeting practiced by high-level civil servants indicate that the distinction between physical world and online space remains elusive and as a subject of debate. One critical explanation to this relates to perceived anonymity in Internet that could be thought to reduce accountability (Davenport, 2002), further implying that actions that are considered real in the physical space are not necessarily so in cyberspace. It is evident that societal norms in cyberspace are still evolving: institutions are trying to keep up with the increasing need for new legislations while people's values are constantly being shaped by latest disruptive IT innovations.

This study comes with certain limitations that we should acknowledge. Our sample is limited to the top-10 agents in the views-buying business. Although we deem this sample as representative and sufficient for our analytical purposes, it is possible that a larger sample would yield additional insights. Further, while we uncovered a set of key persuasive messages applied by crowdturfing agents, we cannot make direct conclusions about their relative effectiveness and their roles in agents' strategy. Future research may build on our findings and investigate how the identified persuasion strategies work on potential customers. Moreover, the current study is limited to investigating crowdturfing agent platforms. Future research could widen the scope to the perspectives of workers and buyers. It would be particularly interesting to learn about customers' motivations to use such services and whether workers have other incentives than financial rewards behind their behavior. Considering the ethical aspects to IS use could help to explain why these actors participate in this business. One potential direction would be to study the maliciousness or trustworthiness of crowdturfing workers in line with recent crowdsourcing research (Gadiraju et al., 2015) – it would be intriguing to find out how the ethical quandaries associated with crowdturfing business interact with workers' work ethics.

7 Conclusion

The business of selling 'real' (i.e., human-generated) online social impressions for a fee (aka, crowdturfing) is booming. This study was set out to explore two central questions: 1) What persuasive strategies do crowdturfing agents leverage to promote their service offerings?; and 2) To what extent these offerings can be ethically justified? The quest to answering these questions has led us to identify ten leading platforms operating in this domain and analyze their persuasive strategies. Our findings revealed that crowdturfing platforms provide a range of fabricated social impressions for multiple target platforms, and leverage communication strategies that typically consist of three types of messages: educational messages, bragging messages, and reassuring messages. We distinguished specific communication styles and techniques behind each message type. Moreover, we found that the agents resort to specific ethical appeals, drawing largely on a combination of consequentialist logic and a mechanistic interpretation of target platforms' codified rules. Most importantly, we found that defining 'what real is' is not as straightforward as one might have thought. Although this article is by no means intended to spread paranoia, it certainly opens our eyes that in cyberspace, things are not always as 'real' as they may appear!

References

- Abokhodair, N., Yoo, D., & McDonald, D. W. (2015). Dissecting a Social Botnet: Growth, Content and Influence in Twitter. *Proceedings of the 18th ACM Conference on Computer Supported Cooperative Work & Social Computing - CSCW '15*, 839–851. <https://doi.org/10.1145/2675133.2675208>
- Armstrong, J. S., & Green, K. C. (2013). Effects of corporate social responsibility and irresponsibility policies. *Journal of Business Research*, 66(10), 1922–1927.
- Beder, S. (1998). Public relations' role in manufacturing artificial grass roots coalitions. *Public Relations Quarterly*, 43, 20–23.
- Bloomberg, J. (2017, January 8). Fake News? Big Data And Artificial Intelligence To The Rescue. *Forbes*. Retrieved from <https://www.forbes.com/sites/jasonbloomberg/2017/01/08/fake-news-big-data-and-artificial-intelligence-to-the-rescue/#69205f14a303>
- Braun, V., & Clarke, V. (2006). Using thematic analysis in psychology. *Qualitative Research in Psychology*, 3(2), 77–101.
- Castle, T., & Lee, J. (2008). Ordering sex in cyberspace: a content analysis of escort websites. *International Journal of Cultural Studies*, 11(1), 107–121. <https://doi.org/10.1177/1367877907086395>
- Chatterjee, S., Sarker, S., & Valacich, J. S. (2015). The Behavioral Roots of Information Systems Security: Exploring Key Factors Related to Unethical IT Use. *Journal of Management Information Systems*, 31(4), 49–87. <https://doi.org/10.1080/07421222.2014.1001257>
- Cho, C. H., Martens, M. L., Kim, H., & Rodrigue, M. (2011). Astroturfing Global Warming: It Isn't Always Greener on the Other Side of the Fence. *Journal of Business Ethics*, 104(4), 571–587. <https://doi.org/10.1007/s10551-011-0950-6>
- Davenport, D. (2002). Anonymity on the Internet: why the price may be too high. *Communications of the ACM*, 45(4), 33. <https://doi.org/10.1145/505248.505267>
- Duan, W., Gu, B., & Whinston, A. B. (2008). Do online reviews matter? - An empirical investigation of panel data. *Decision Support Systems*, 45(4), 1007–1016. <https://doi.org/10.1016/j.dss.2008.04.001>
- Eisenmann, T., Parker, G., & Alstyne, M. W. Van. (2006). Strategies for Two- Sided Markets. *Harvard Business Review*, 84(10), 12. <https://doi.org/10.1007/s00199-006-0114-6>
- Farooqi, S., Ikram, M., De Cristofaro, E., Friedman, A., Jourjon, G., Kaafar, M. A., ... Zaffar, F. (2017). Characterizing Key Stakeholders in an Online Black-Hat Marketplace. *arXiv:1505.01637v2*. Retrieved from <http://arxiv.org/abs/1505.01637>
- Fayazi, A., Lee, K., Caverlee, J., & Squicciarini, A. (2015). Uncovering Crowdsourced Manipulation of Online Reviews. *Proceedings of the 38th International ACM SIGIR Conference on Research and Development in Information Retrieval - SIGIR '15*, 233–242. <https://doi.org/10.1145/2766462.2767742>
- Fereday, J., & Muir-Cochrane, E. (2006). Demonstrating Rigor Using Thematic Analysis: A Hybrid Approach of Inductive and Deductive Coding and Theme Development. *International Journal of Qualitative Methods*, 5(1), 80–92. <https://doi.org/10.1177/160940690600500107>
- Gadiraju, U., Kawase, R., Dietze, S., & Demartini, G. (2015). Understanding Malicious Behavior in Crowdsourcing Platforms : The Case of Online Surveys. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems* (pp. 1631–1640).
- Gephardt, R. (2004). What is qualitative research and why is it important? *Academy of Management Journal*, 7(4), 454–462.
- Gerodimos, R. (2008). Mobilising young citizens in the UK: A content analysis of youth and issue websites. *Information Communication and Society*, 11(7), 964–988. <https://doi.org/10.1080/13691180802109014>

- Gert, B. (1998). *Morality: Its Nature and Justification*. Oxford: Oxford University Press.
- Ghazawneh, A., & Henfridsson, O. (2013). Balancing platform control and external contribution in third-party development: The boundary resources model. *Information Systems Journal*, 23(2), 173–192.
- Glaser, B. G., & Strauss, A. L. (1967). *The discovery of grounded theory: strategies for qualitative research*. New Brunswick, London: AldineTransaction.
- Goldschein, E. (2011). 10 Fake Grassroots Movements Started by Corporations to Sway Your Opinion. *Business Insider*. Retrieved from <http://www.businessinsider.com/astroturfing-grassroots-movements-2011-9?r=US&IR=T&IR=T>
- Grana, R. A., & Ling, P. M. (2014). “Smoking revolution”: A content analysis of electronic cigarette retail websites. *American Journal of Preventive Medicine*, 46(4), 395–403. <https://doi.org/10.1016/j.amepre.2013.12.010>
- Han, R. (2015). Manufacturing Consent in Cyberspace: China’s “Fifty-Cent Army.” *Journal of Current Chinese Affairs*, 44(2), 105–134.
- Hsieh, H.-F., & Shannon, S. E. (2005). Three Approaches to Qualitative Content Analysis. *Qualitative Health Research*, 15(9), 1277–1288. <https://doi.org/10.1177/1049732305276687>
- Jones, T. M., Felps, W., & Bigley, G. A. (2007). Ethical theory and stakeholder-related decisions: The role of stakeholder culture. *Academy of Management Review*, 32(1), 137–155.
- Kraemer, R., Whiteman, G., & Banerjee, B. (2013). Conflict And Astroturfing In Niyamgiri: The Importance Of National Advocacy Networks In Anti-Corporate Social Movements. *Organization Studies*, 34(5–6), 823–852. <https://doi.org/10.1177/0170840613479240>
- Lappas, T., Sabnis, G., & Valkanas, G. (2016). The Impact of Fake Reviews on Online Visibility: A Vulnerability Assessment of the Hotel Industry. *Information Systems Research, Articles i*(November), 1–22. <https://doi.org/10.1287/isre.2016.0674>
- Lee, K., Tamilarasan, P., & Caverlee, J. (2013). Crowdturfers, campaigns, and social media: tracking and revealing crowdsourced manipulation of social media. *Icwsn 2013*, 331–340. Retrieved from <http://www.aaai.org/ocs/index.php/ICWSM/ICWSM13/paper/viewPDFInterstitial/5988/6372>
- Lee, K., Webb, S., & Ge, H. (2015). Characterizing and automatically detecting crowdturfing in Fiverr and Twitter. *Social Network Analysis and Mining*, 5(1), 1–16. <https://doi.org/10.1007/s13278-014-0241-1>
- Leimeister, J. M., Huber, M., Bretschneider, U., & Krcmar, H. (2009). Leveraging Crowdsourcing: Activation-Supporting Components for IT-Based Ideas Competition. *Journal of Management Information Systems*, 26(1), 197–224. <https://doi.org/10.2753/MIS0742-1222260108>
- Li, G., Niu, W., Batten, L., & Liu, J. (2017). New advances in securing cyberspace and curbing crowdturfing. *Concurrency and Computation: Practice and Experience*, 29(20), 1–3. <https://doi.org/10.1002/cpe.4162>
- Li, X., & Wu, L. (2013). Measuring effects of observational learning and social-network Word-of-Mouth (WOM) on the sales of daily-deal vouchers. *Proceedings of the Annual Hawaii International Conference on System Sciences*, 2908–2917. <https://doi.org/10.1109/HICSS.2013.397>
- Mahmood, M. A., Siponen, M., Straub, D., Rao, H. R., & Raghu, T. S. (2010). Moving toward black hat research in information systems security: An editorial introduction to the special issue. *MIS Quarterly*, 34(3), 431–433.
- Majchrzak, A., & Malhotra, A. (2013). Towards an information systems perspective and research agenda on crowdsourcing for innovation. *Journal of Strategic Information Systems*, 22(4), 257–268. <https://doi.org/10.1016/j.jsis.2013.07.004>
- Mason, R. O. (1986). Four ethical issues of the information age. *MIS Quarterly*, 10(1), 5–12. <https://doi.org/10.2307/248873>

- McIntyre, H. (2017). How Much Money Does 3 Billion YouTube Views Actually Generate For A Musician? *Forbes*. Retrieved from <https://www.forbes.com/sites/hughmcintyre/2017/09/18/how-much-money-does-3-billion-youtube-views-bring-in/#283e29f24aec>
- Mitra, A., & Watts, E. (2002). Theorizing cyberspace: The idea of voice applied to the internet discourse. *New Media and Society*, 4(4), 479–498. <https://doi.org/10.1177/146144402321466778>
- Moor, J. H. (1999). Just consequentialism and computing. *Ethics and Information Technology*, 1, 65–69.
- Moor, J. H. (2001). The future of computer ethics: You ain't seen nothin' yet! *Ethics and Information Technology*, 89–91. <https://doi.org/10.1023/a:1011881522593>
- Morgan, D. L. (1993). Qualitative Content Analysis: A Guide to Paths not Taken. *Qualitative Health Research*, 3(1), 112–121.
- Norris, M. L., Boydel, K. M., Pinhas, L., & Katzman, D. K. (2006). Ana and the Internet: A Review of Pro-Anorexia Websites. *International Journal of Eating Disorders*, 39(6), 443–447. <https://doi.org/10.1002/eat>
- Peslak, A. R. (2006). Papa Revisited: a Current Empirical Study of the Mason Framework. *Journal of Computer Information Systems*, 46(3), 117–124. <https://doi.org/10.1080/08874417.2006.11645905>
- Pham, N. (2013). Vietnam admits deploying bloggers to support government. *BBC News*. Retrieved from <http://www.bbc.com/news/world-asia-20982985>
- Price, R. (2017). Researchers taught AI to write totally believable fake reviews, and the implications are terrifying. *Business Insider*. Retrieved from <http://www.businessinsider.com/researchers-teach-ai-neural-network-write-fake-reviews-fake-news-2017-8?r=UK&IR=T&IR=T>
- Price, S. R. (2014). Big Business: Buying Fake Instagram Followers. *The Huffington Post*. Retrieved from http://www.huffingtonpost.com/shayla-r-price/big-business-buying-fake_b_6322362.html
- Rinta-Kahila, T., & Soliman, W. (2017). Understanding Crowdturfing: the Different Ethical Logics Behind the Clandestine Industry of De- Ception. *Twenty-Fifth European Conference on Information Systems (ECIS), 2017*, 1–16.
- Schlagwein, D., & Bjørn-Andersen, N. (2014). Organizational Learning with Crowdsourcing: The Revelatory Case of LEGO. *Journal of the Association for Information Systems*, 15(11), 754–778.
- Shaffer, A. (2013). Five Stars for Five Dollars: Buying Reviews, Reviewed. *The Huffington Post*. Retrieved from http://www.huffingtonpost.com/andrew-shaffer/five-stars-for-five-dolla_b_3997107.html
- Siering, M., Koch, J.-A., & Deokar, A. V. (2016). Detecting Fraudulent Behavior on Crowdfunding Platforms: The Role of Linguistic and Content-Based Cues in Static and Dynamic Contexts. *Journal of Management Information Systems*, 33(2), 1–35. <https://doi.org/10.1080/07421222.2016.1205930>
- Siponen, M., & Vartiainen, T. (2002). Teaching end-user ethics: Issues and a solution based on universalizability. *Communications of the Association for Information Systems*, 8, 422–444.
- Smith, H. J. (2002). Ethics and information systems. *ACM SIGMIS Database*, 33(3), 8–22. <https://doi.org/10.1145/569905.569908>
- Snow, J. (2017, December 13). Can AI Win the War Against Fake News? *MIT Technology Review*.
- Song, J., Lee, S., & Kim, J. (2015). CrowdTarget: Target-based Detection of Crowdturfing in Online Social Networks. *Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security*, (i), 793–804. <https://doi.org/10.1145/2810103.2813661>
- Van Maanen, J. (1979). The Fact of Fiction in Organizational Ethnography. *Administrative Science Quarterly*, 24(4), 539–550.
- Wang, G., Wilson, C., Zhao, X., Zhu, Y., Mohanlal, M., Zheng, H., & Zhao, B. Y. (2012). Serf and Turf: Crowdturfing for Fun and Profit. *Proceedings of the 21st International Conference on World*

Wide Web. ACM., 679–688. <https://doi.org/10.1145/2187836.2187928>

Warkentin, M., Straub, D., & Malimage, K. (2012). Featured talk: Measuring secure behavior. Annual Symposium on Information Assurance & Secure Knowledge Management. Albany, NY.

Yao, Y., Viswanath, B., Cryan, J., Zheng, H., & Zhao, B. Y. (2017). Automated Crowdturfing Attacks and Defenses in Online Review Systems. *arXiv Preprint arXiv:1708.08151*. <https://doi.org/10.1145/3133956.3133990>

Zhu, F., & Zhang, X. (2010). Impact of Online Consumer Reviews on Sales: The Moderating Role of Product and Consumer Characteristics. *Journal of Marketing*, 74(March), 133–148. <https://doi.org/10.1509/jmkg.74.2.133>

Zuboff, S. (2015). Big other: Surveillance capitalism and the prospects of an information civilization. *Journal of Information Technology*, 30(1), 75–89. <https://doi.org/10.1057/jit.2015.5>

Appendix: Coding map

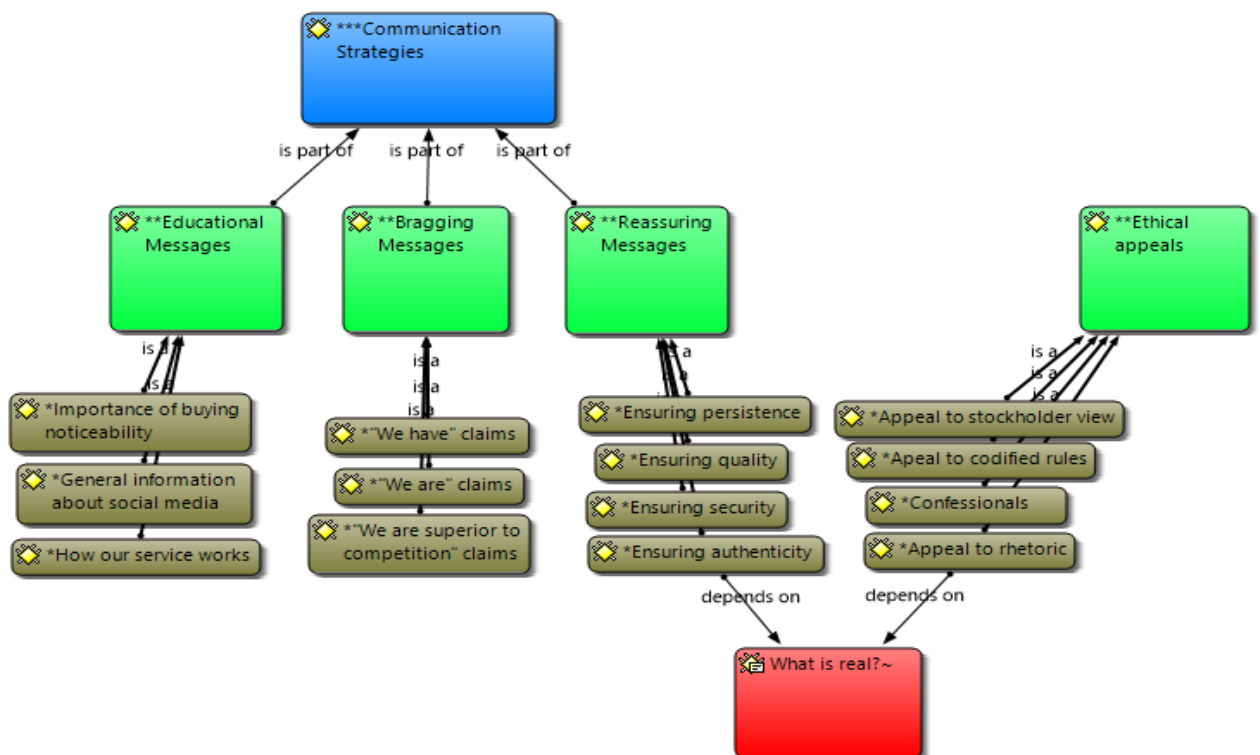


Figure A1. Coding map for overall analysis