

**This is an electronic reprint of the original article.
This reprint *may differ* from the original in pagination and typographic detail.**

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

Title: Understanding Crowdturfing : The Different Ethical Logics Behind the Clandestine Industry of Deception

Year: 2017

Version:

Please cite the original version:

Rinta-Kahila, T., & Soliman, W. (2017). Understanding Crowdturfing : The Different Ethical Logics Behind the Clandestine Industry of Deception. In ECIS 2017 : Proceedings of the 25th European Conference on Information Systems, Guimarães, Portugal, June 5-10, 2017 (pp. 1934-1949). European Conference on Information Systems. http://aisel.aisnet.org/ecis2017_rp/124

All material supplied via JYX is protected by copyright and other intellectual property rights, and duplication or sale of all or part of any of the repository collections is not permitted, except that material may be duplicated by you for your research use or educational purposes in electronic or print form. You must obtain permission for any other use. Electronic or print copies may not be offered, whether for sale or otherwise to anyone who is not an authorised user.

UNDERSTANDING CROWDTURFING: THE DIFFERENT ETHICAL LOGICS BEHIND THE CLANDESTINE INDUSTRY OF DECEPTION

Research paper

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

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

Abstract

Crowdturfing, the dark side and usually unnoticed face of crowdsourcing, represents a form of cyber-deception in which workers are paid to express a false digital impression. While such behavior may not be punishable under the jurisdiction of formal law, its consequences are destructive to the cohesion and trustworthiness of online information. The conceptual work at hand examines the current literature on the topic, and lays the foundation for a theoretical framework that explains crowdturfing behavior. We discuss crowdturfing through three ethical normative approaches: traditional philosophical ethics, business ethics, and codified rules. We apply these lenses to an illustrative example of an online platform orchestrating the trade of paid book reviews on Amazon. The study contributes to theory by explaining the ethical logic behind crowdturfing from the perspectives of the key actors involved in the business. We argue that while crowdturfing cannot stand a critical examination through the deontological, stakeholder, or social contract perspectives, leaning on the teleological logic, the stockholder theory, or certain levels of codified rules can enable the actors involved in the business to operate with clean conscience. An increased understanding of the behavior can help both victim platforms and the Internet community at large to combat this hidden industry.

Keywords: crowdturfing, (un)ethical IS use, social contract, deterrence, cybersecurity

1 Introduction

“Our moral imperative is clear. We must insure that information technology, and the information it handles, are used to enhance the dignity of mankind. To achieve these goals we must formulate a new social contract, one that insures everyone the right to fulfill his or her own human potential” – Mason (1986), p. 11.

“In many cases the uses of computers are so unusual that no policies for their proper use exist or even have been considered. Policy vacuums exist. We need to formulate policies, even if only informally, to make sure our actions remain within ethical bounds” – Moor (2001), p. 89.

It is well established that the ultimate aim of information security is the protection of three fundamental informational qualities: integrity, availability and confidentiality (Loch, Carr, and Warkentin 1992); something that Mason (1986) outlined as four exceptionally important issues in the information age, and labeled them PAPA (short for privacy, accuracy, property, and accessibility). The manuscript at hand critically examines crowdturfing as a novel Internet phenomenon that jeopardizes information integrity or accuracy on the Web. From a formal and legal perspective, information integrity, concerns safeguarding against improper information modification or destruction, and most importantly ensuring information authenticity (U.S. Code title 44). Specifically, crowdturfing poses a unique threat to authenticity/accuracy on the Web, not only because it undermines the reliability of the circulated information, but also because it is not a punishable crime, and at best a frowned-upon behavior. After all, giving a positive comment may be justified as a ‘white lie’, at least from the commenter’s perspective. As such, there is little reason to expect that crime-focused theories like deterrence theory, which rely predominantly on applying severe and certain punishments, will be able to stem this practice anytime soon. In fact, in the past few years, news reports have surfaced showing that more and more it is becoming commonplace to crowdsource (i.e., hire people en masse) to perform a scripted behavior, i.e. to crowdturf. Whether you represent a reality show celebrity, a service provider, or even a government (Elder 2012); you have the option to gain phony support for a price. Indeed, on websites like Liftoff Social¹, Boostlikes², and Buy Likes and Followers³, Facebook likes, Instagram followers, and even YouTube video views all have a price tag. For instance, on Boostlikes \$39 gets you a one-time order of 1,000 Instagram followers. Moreover, in a recent secret investigation of the matter, *The Sunday Times* was able to hire professional fake review writers for as little as £3 for each five-star review (Henry 2015).

These services are available for anyone able and willing to pay for them, no matter what the intentions of the buyer are or the consequences that may follow. This makes crowdturfing a potentially harmful phenomenon in various respects. First of all, unauthentic online behaviors pose a major threat to the trustworthiness of information circulated on the Web. Especially online reviews have a major influence on shaping consumers’ opinions and subsequent purchase behaviors. A recent report by the British Competition and Marketing Authority estimates that annually £23 billion of UK consumer spending is influenced by online reviews (CMA 2015). Since trustworthy reviews are often essential to the credibility of a retail website, some online retailers like Amazon have been driven into a seemingly unending battle of rooting out paid reviews websites that keep constantly emerging to replace the closed ones (Rubin 2016). Secondly, since flooding an IT platform with unreliable information derogates the platform’s credibility, crowdturfing can ultimately hamper the assimilation of potentially useful ISs. For instance,

¹ www.liftoffsocial.com

² www.boostlikes.com

³ www.buylikesandfollowers.net

recent research shows that the proliferation of unethical behavior is a major factor of irreversible IS discontinuance behavior by those who feel cheated (Boukef and Charki 2014). Thirdly, such services provide actors such as political parties, government officials, or others with powerful tools to manufacture consent that is in line with their agenda (see e.g. Howard 2003). If actors with malicious intentions are able to harness the full potential of these services, the damages may not be limited to deceiving consumers into buying crummy products: crowdturfing can pose a threat to societal virtues such as freedom, integrity, or stability. For instance, it has been revealed that the Chinese government has recruited an army of turfers – also known as the ‘fifty-cent-army (Han 2015) – to push their official agenda in the virtual space (Fareed 2008). Since the human knowledge is being recorded into, and retrieved from the Internet in increasing amounts, a practice like crowdturfing can in its most extreme form facilitate the formation of an Orwellian society where knowledge is controlled by those who have acquired the tools for shaping it in their own needs. Alarming, the crowdturfing industry was estimated at millions of dollars in 2011, with expectations of exponential growth as now we only looking at the early stages of the phenomenon (Simonite 2011).

Regarding the unethical IS use in general, Chatterjee et al. (2015) have recognized the lack of research taking ethical perspective on why individuals engage in such behavior. Moreover, they highlight the need for further investigation on the role of technology in facilitating unethical behavior. Since crowdturfing is a rather novel form of unethical IS use, not much research has yet been conducted on the topic. While some efforts have been made to provide tools to detect crowdturfing activities, we do not know much about the behavioral and ethical aspects of the phenomenon. Thus, questions such as what motivates crowdturfing workers, intermediary agents, and customers?; how do they justify the behavior to themselves so that they can operate with clear conscience?; and how could we counter this problem?, remain unanswered.

To address this apparent need for investigation of the ethical and behavioral aspects of crowdturfing, we set out to find answers to the following research question: *How do the different actors involved in crowdturfing justify their behavior ethically?* We address this question by investigating the phenomenon through theoretical lenses provided by the current literature on cybersecurity and ethics in IS use. Specifically, we argue that deterrence theory, the most widely cited framework in cybersecurity research, may not be effective in dissuading crowdturfing. Thus, we turn to examine pre-kinetic events (Willison and Warkentin 2013) that precede intentions to system abuse. We discuss Smith’s (2002) meta-framework of ethics for information systems and perform a conceptual analysis of ethical logics behind crowdturfing by applying the framework to a real crowdturfing business: Paid Book Reviews, a website that promises its paying customers the arrangement of positive reviews for items sold on Amazon.

This paper is organized as follows. In Section 2, we discuss the current literature on crowdsourcing and crowdturfing. In Section 3 we establish the theoretical background of this study by discussing cybersecurity literature and ethical issues in IS use. In Section 4 we describe and analyze the crowdturfing case selected for this study. Finally, we discuss the implications of our findings in Section 5, and Section 6 is dedicated to concluding remarks.

2 Literature review

Crowdturfing represents one of the many forms crowdsourcing has taken in the past decade (Afuah and Tucci 2012; Estellés-Arolas and González-Ladrón-de-Guevara 2012; Howe 2006). Typically, they are seen as an electronic market (Alt & Klein 2011), where *seekers* (usually firms seeking solutions to specific problems) and *solvers* (potentially anyone in the crowd) are in the trade of solutions-for-rewards over the *crowdsourcing* platform (often operated by a third-party brokerage). However, the difference between crowdsourcing and crowdturfing boils down to one single characteristic: the nature of the traded tasks. On the one hand, at a general level crowdsourcing has been praised by both academics and practitioners as a desirable phenomenon (e.g. Leimeister et al. 2009; Majchrzak and Malhotra 2013; Schlagwein and Bjørn-andersen 2014), mainly due to the many economic and societal values gained from optimizing the crowd’s “under-utilized resources”, e.g., part-taking in the journalism industry

(Soliman 2015). On the other hand, crowdturfing is a specific case of crowdsourcing in which members of the crowd are paid to leave a fake digital trail, e.g., pretending to be a fan of some account on Instagram (Price 2014). While crowdsourcing in terms of value-creation has enjoyed a growing interest as a current IS research topic, less attention has been given to the dark side of crowdsourcing that often results in degrading the value of information in the Internet as its fundamental aim tends to be to further the interests of certain parties at the expense of the rest of the Internet population.

Crowdturfing, a portmanteau coined by Wang et al. (2012), combines crowdsourcing with astroturfing. The latter refers to a practice where a certain party (individual or organization) attempts to spread information and manufacture consent on a topic of its interest in a way that mimics activities of genuine grass root movements. This happens often in the form of spreading specific rumors, political messages, or false advertising. Crowdturfing is a specific form of astroturfing, where the client harnesses the power of the crowds by mobilizing people over the Internet to engage in campaigns that further the client's goal. Such campaigns are typically structured as collections of tasks that workers complete in exchange for a fee (Wang et al. 2012). In contrast to astroturfing, the fast and effective utilization of real human workers has been considered to drastically increase the impact of such campaigns (ibid.). While many crowdsourcing platforms specifically prohibit and endeavor to stem such crowdturfing campaigns (as in e.g., Amazon's Mechanical Turk), others turn a blind eye to such campaigns by allowing them to exist in vast numbers (e.g., ShortTask), and some even invite and encourage them as crowdturfing is often the core business of such sites (e.g., Buy Likes & Followers).

A crowdturfing campaign typically involves three key actors: *workers* who perform a certain task for a fee, *customers* who benefit from and pay for the task, and the *agent*, often in a form of an intermediary platform, who connects workers with customers and orchestrates the whole transaction (Wang et al. 2012). The predictable outcome of such campaigns on a target platform is the distortion of public opinion among its users. For several years, Amazon has been the target of such campaigns where a self-published author could buy a 'five stars' review for as low as five dollars (Shaffer 2013). Other businesses operating in this domain facilitate buying Facebook 'likes', following social media accounts, retweeting certain messages, and even performing precise searches on Google and clicking on a specified link to elevate its ranking.

Moreover, crowdturfing campaigns affect other stakeholders in the Internet too. Such campaigns pose a threat to the target platform where workers are assigned to execute their tasks (e.g. Amazon, Facebook, Twitter, Google), as increase in spurious information may deteriorate the credibility and usefulness of that platform. In addition, when looking beyond these four immediately affected parties, crowdturfing impacts the whole Internet community (which may ultimately translate to affect communities in the physical realm), as crowdturfing campaigns shape the nature and quality of the information available in the Internet. Parties involved in and affected by crowdturfing are illustrated in Figure 1 that demonstrates how customers' campaigns translate into information searched by Internet users.

By using real human workers, the customers counter the shortcomings of using bots for creating automatically generated accounts and content. Although means for inhibiting the malicious tasks performed automatically by bots have been developed (Abokhodair, Yoo, and McDonald 2015), they are considered as less effective against human turfing, who can appear relatively sincere and genuine. Thus far, crowdturfing has been almost exclusively addressed in the domain of computer science, largely by developing algorithms and machine learning tools that can detect and identify crowdturfers and their malicious activities. Conventional detection methods usually base their analyses on account characteristics, and have managed to produce detection rates as high as 97.35 % for crowdturfing tasks (Lee, Webb, and Ge 2015). However, a key challenge in detecting crowdturfing is that since turfing are real human workers, their account characteristics can be very similar to normal, non-malicious accounts (Song, Lee, and Kim 2015). Thus, also contemporary methods have been introduced, such as CrowdTarget that instead of turfer accounts and activities detects the very target objects that are under manipulation attempts (Song, Lee, and Kim 2015).

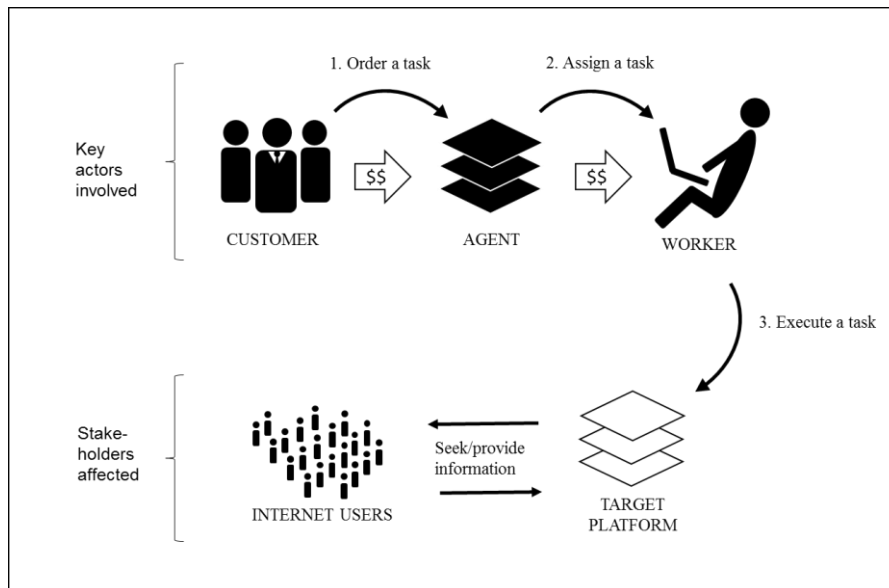


Figure 1: Information flow in crowdturfing.

While the work on crowdturfing detection has contributed invaluable in understanding of the problem and gives practical help in combating it, the approach leaves one major area unexplained: user behavior. To comprehensively explain and effectively counter such phenomenon, one should understand the ethical grounds behind the behavior. Surprisingly, to date no research exists that aims to explain the motivations and justifications behind crowdturfing. Moreover, while prior literature has succeeded in detecting crowdturfing, we are still lacking the tools for fighting and preventing it. To address this apparent research gap, we first discuss how these issues could be addressed with the existing theories on information security, and then move on to seek explanations to crowdturfing by viewing it through different normative ethics perspectives.

3 Theoretical background

Thus far, we have adopted a normative position that frames crowdturfing as the dark side of crowdsourcing, i.e., an unethical cyberspace behavior worth investigating how to stem or deter it. In the broad spectrum of behavioral information security literature, the general deterrence theory from criminology (Gibbs 1975) has been one of the most cited theories (Lebek et al. 2014) postulating that delinquent behaviors are dissuaded or deterred (hence the theory's name) through the application of severe and certain punishments (D'Arcy and Hovav 2007; Straub 1990; Straub and Nance 1990). For instance, Straub and Welke (1998) use the deterrence theory as the foundation of their overall framework for studying IS security risks. In their security action cycle, the process of reducing systems' security risks culminates in four sequential activities: 1) deterrence, 2) prevention, 3) detection, and 4) remedies. In short, this theory posits that countering risks starts from deterrence that stresses communicating the certainty and severity of sanctioning the abuser. However, if this fails to dissuade the abuser, the next step is to try to prevent the threat with active countermeasures. If also prevention fails, the system owner should be capable of detecting the threat. Finally, actions should be taken to remedy the damages and punish the abuser. Effective action in each stage strengthens the deterrence as responses to breaches on each stage inform potential abusers of the consequences of abuse (i.e., deterrence feedback loop).

However, the efficacy of sanctions (rooted in the deterrence theoretical framework) is undermined in contexts that lack strong central governance, coercion, and "mass obedience" (Kennedy 1983, p. 10), and when actors employ neutralization tactics that allow them to minimize the perceived harm of their behavior (Siponen and Vance 2010; Siponen, Vance, and Willison 2012). In terms of preventive

measures, Levchenko et al. (2011) suggest that following the flow of money could be effective in countering spamming: they find that the funding for 95 % of spam-advertised pharmaceutical, replica, and software products goes through only a handful of banks. Thus, they propose that convincing banks to refuse providing certain services to known spammers could work as an effective intervention policy that would prevent spamming. However, while following the flow of money could help in stemming crowdturfing, the fact that several agent platforms accept encrypted bitcoin payments makes this increasingly complicated. Moreover, the extant research on crowdturfing specifically addresses the third stage of the cycle: detection. Effective detection of crowdturfing has enabled target platforms to apply remedying actions such as banning certain sellers and reviewers from using their services and pursuing legal action against the agent platforms (e.g. Amazon, see Rubin 2016). But this often happens with a lag, and the malicious actors can always create new accounts for continuing the abuse.

Acknowledging these limitations, there have been recent calls to extend the research scope to go beyond deterrence and explore the “pre-kinetic” events, or phenomena that temporally occur prior to deterrence (Willison and Warkentin 2013). Such events further the formation of behavioral intentions to abuse a system, emerge from the interaction between the IS user and the surrounding context (e.g. organization), and appear in forms like feeling of injustice, disgruntlement, neutralization, and self-expression. For instance, Willison and Warkentin (2013) identify instrumental and expressive motives as prominent drivers for system abuse. Another pre-kinetic event that we advance here is the importance of understanding the various normative (ethical) arguments that those involved in crowdturfing may apply to give their behavior ethical legitimacy. The following discussion is heavily influenced by Smith’s (2002) treatise on ethics in information systems.

3.1 Ethical issues in IS use

Concerns about people’s ethical behavior in the information age are anything but new. Mason attributes this concern to a fundamental information-specific characteristic: “information is the means through which the mind expands and increases its capacity to achieve its goals, often as the result of an input from another mind” (Mason, 1986, p. 5). Based on this premise Mason calls upon a revision to the social contract governing people in the information age, most importantly with special attention to four ethical issues reflecting four vulnerabilities: (1) privacy, (2) accuracy, (3) property, and (4) accessibility. Privacy is a question of exposure, and is concerned with what information a person should (must) reveal and/or keep secret. Accuracy is a question of authenticity, and is concerned with how to make sure that decision makers (whoever they may be) are not misinformed. Property is a question of ownership, and is concerned with who should own the intellectual property of information and for what price. Accessibility is a question of admission, and is concerned with making sure of granting access rights of information only to the intended individuals. While crowdturfing may not threaten privacy, property, or accessibility, it undeniably poses a major threat to information accuracy.

3.1.1 Traditional Philosophical Ethics

Building on Mason’s work, Chatterjee et al. (2015) define unethical use of IT as “*The willful violation—by any individual, group, or organization—of privacy, and/or property, and/or accuracy, and/or access—with respect to information/information goods residing within or part of an information system, owned/controlled by any other individual, group, or organization*” (p. 55). Inherent in this definition is the assumption that a violation is unethical if it causes harm and/or violates the codified rules, representing two competing ethics schools of thought: the **teleological**⁴ and the **deontological**⁵ schools. Smith (2002) notes that the core tension between these two paradigms boils down to the conception of what is “right” or “wrong”. The teleological perspective (aka, consequential theories), on one hand, assumes

⁴ *Telos* is the Greek word for “end” or “purpose”.

⁵ *Deon* is the Greek word for “duty”.

that “*the right answer to an ethical quandary depends on the consequences that follow from the action; one should choose the alternative with consequences that produce the most “good” for some group*” (p. 11). On the other hand, the deontological perspective (aka, categorical), holds that “*the ethical “rightness” or “wrongness” of an action can never be determined solely by the consequences that flow from it. Instead, what is “right” is determined by some set of general principles. Often, the focus is on the intent behind an action. Even if things turned out badly, a categorical thinker could claim that an action was ethically right if the actor’s underlying intent had been right*” (ibid, p. 14). Based on this view, Kant and others have proposed general guidelines such as “you should never lie” (Smith 2002, p. 13).

3.1.2 Business Ethics

In addition to traditional philosophical theories, Smith (2002) discusses business ethics that provide their own guidelines to what is right and what is wrong. Firstly, **the stockholder theory** considers that decision-makers in companies should take such actions that maximize profits and generate the greatest value for stockholders in the long-term. On the contrary, **the stakeholder theory** posits that “*interests of the stakeholders other than stockholders should be considered along with those of the stockholders even if it reduces firm profitability*” (Smith 2002, p. 14). Finally, Smith (2002, p. 15) discusses **the social contract theory** of business ethics, which “*demands that managers act so that consumer and worker interests be satisfied in a manner that maximizes advantages and minimizes disadvantages*”. Moreover, social contract theory puts exceptional emphasis on the concept of justice. For instance, in a business ethics context, the justice term of a hypothetical social contract mandates that firms must avoid “fraud”, “deception”, as well as any practice that “*systematically worsens the situation of a given group in society*” (Smith 2002, p. 15). Similarly, for an individual, adherence to social contract may actualize e.g. as obedience of the law and behavior that does not cause harm to others: one would refrain from stealing from others or abusing others even though the malicious behavior would result in certain benefits for the enacting individual. Social contract theory has been found useful in investigating some controversial behaviors of the information age, e.g. Internet advertising (Gordon and De Lima-Turner 1997).

3.1.3 Codified rules

Finally, Smith (2002) explains that **codified rules** can be used to circumvent the need for normative argumentation inherent in applying the previously discussed schools of ethics. In essence, codified rules are “*a set of guidelines that can be followed*”, and which “*can be constructed at several levels, most notably for a specific corporation (often called a “Corporate Code of Conduct,” the breach of which is sometimes cause for dismissal); for a profession or industry (often denoted a “Code of Ethics”); or for a society as a whole, as witnessed in its legal structure*” (Smith 2002, p. 17). Interestingly, the IS context enables relatively easy implementation of highly specific and detailed levels of codified rules, such as terms and conditions of particular software use or different sets of rules for users that are granted different levels of modification rights on a website (e.g. domain owner vs. coder vs. moderator vs. member user vs. non-member user, etc.).

3.1.4 The collision of cyberspace and the physical world

In sum, the theories discussed above can be helpful in gaining understanding of unethical computer use, especially if we consider the distinction between virtual and physical realities. Although the virtual world is increasingly merging with the physical realm (e.g. the recent emergence of VR and certain games such as Pokémon Go), from the behavioral perspective these two spaces are still largely considered as distinct from each other. For instance, a person who would consider stealing a movie or a music record from a retail store as highly unethical might be perfectly fine with illegally downloading the same product (Siponen 2006). Overall, it appears that individuals’ behavior has less inhibitions in cyberspace than in the physical world. One explanation to this could be in the actual or perceived anonymity in Internet that reduces accountability (Davenport 2002). This implies that a certain degree of anonymity

and the lack of physical presence enabled by the cyberspace could actually help to bring out the worst in people.

Leetaru (2016) argues that the Internet might be evolving away from its original role as the anti-censorship-valuing cradle of free speech and towards a corporate-controlled medium that watches over the interests of commerce and governments that it used to rebel against. Thus, the currently ongoing conflicts and quandaries related to issues such as virtual abuse, immaterial property rights, justification of netizen surveillance, and others could relate to growing pains in the collision of virtual and physical realities before finding a new balance. It is evident that the societal role of cyberspace is still evolving: while people are making their own interpretations about what is acceptable in cyberspace, institutions are struggling to keep up with the increasing need for new legislations and ethical guidelines. As such, it may be possible that there occurs a lag between peoples' adherence to ethical guidelines like social contracts or codified rules in the cyberspace for behaviors that are otherwise unacceptable in the physical space.

4 Conceptual analysis

In order to advance our understanding of how people might justify their engagement in crowdturfing, we next discuss different ethical grounds of this behavior through a conceptual analysis of an exemplar real life case. Specifically, our analysis relies on conceptual argumentation regarding the fit between the different ethical logics presented above and the potential behaviors that ensue. We believe that this approach is suitable for illustrating purposes especially when the discussed cases have received considerable media coverage, thus allowing us to draw meaningful insights (see, e.g., Smith 2002; Wall, Lowry, and Barlow 2016). That said, our demonstrative example focuses on the continuous battle between Amazon, the American electronic commerce and cloud computing company, and the amounting number of crowdturfing platforms.

4.1 Case description

Amazon is known as the world's largest online retailer. One of Amazon's biggest assets is its user-friendly feedback system, which allows users to rate and submit reviews of the bought products to its website. In fact, Amazon was named the largest single source of Internet reviews in 2010. Thus, the credibility of the review system is paramount for Amazon's business model, as Internet shoppers value trustworthy information about products before buying them (Rubin 2016). For the past few years, mainstream media has been reporting on the difficulty Amazon is experiencing with crowdturfing sites describing it as a "never-ending game of Whac-a-Mole" (González 2016)! In the course of this battle, Amazon has singled out and sued various platforms like *amazonverifiedreviews.com*, *amazonreviews-star.com*, *buyamazonreviews.info*, *reviewconnections.com*, and *paidbookreviews.org* (González 2016; Perez 2016), in addition to 1,114 individuals advertising their paid review services on the platform *fiverr.com* (Orphanides 2015). While several sites have been shut down as a result of these legal actions (Gani 2015; Rubin 2015), many remain still in operation and new ones keep appearing. Among those still in operation at the time of writing these lines is *paidbookreviews.org*, which we use as our illustrative example. In its 37 seconds long commercial video⁶, Paid Book Reviews (PBR, hereafter) markets itself as follows:

*"Paid book reviews dot org is a team of writers who understand the effect of **positive customer reviews** on your book sales. We offer two types of book reviews: the first are unverified book reviews; we read the sample pages of your book on Amazon dot com and then post **positive comments** under the customer reviews. The second is verified book reviews; we purchase your book and read it on Kindle and then post **positive comments** about your book on Amazon. Please visit [www dotted paid book reviews dot org](http://www.dottedpaidbookreviews.org) now!"* [PBR video commercial, emphasis added].

⁶ <https://www.youtube.com/watch?v=qShlCHcMJww>

As such, PBR promises its subscribers positive reviews on their books listed on Amazon no matter what. For instance, the website promotes that authors who choose the ‘verified book reviews’ option are promised 50 positive reviews for a total of \$1250. While the company claims that their workers actually buy and read the assigned books, Amazon is not as convinced as it has requested the site to stop using Amazon's trademark and stop offering Amazon reviews for sale. Amazon has specifically stated in its Community Guidelines⁷, which explicitly forbid generating content in exchange for compensation or on behalf of anyone else. Moreover, the guidelines disallow offering compensation for such content generation.

Note that this exemplary case of crowdturfing comprises three actors that are immediately involved in the writing of paid reviews: the **customers** who are in need of (and pay for) the positive reviews; **workers**, who are hired to write the positive book reviews in exchange for a fee; and the **agent** (i.e., PBR) who orchestrates the platform and markets itself to customers and invites workers to write the reviews. The behavior of these actors affect other stakeholders who are not directly involved with writing the reviews; most prominently a) the target platform Amazon whose review system's credibility is affected by the paid reviews, and b) the user community at large whose members search and use the information available on the target platform.

4.2 Case analysis

In this section, we discuss this briefly illustrated crowdturfing example from different ethical perspectives for each actor involved traced to one of the three ethical arguments discussed in Section 3: traditional philosophical ethics, business ethics, and codified rules. The principle embedded in a certain ethical perspective can be thought to entail a behavioral guideline for any given action. Thus, since these ethical perspectives differ significantly from each other, each of these lenses is expected to provide a unique guideline on the decision to crowdturf. Table 1 summarizes the most likely behavioral guidelines.

4.2.1 Customers

Customers who buy fabricated reviews for their (or their clients') books via PBR are expecting the reviews to lead into positive economic results which can stem, for instance, from increased sales attributable to positive reviews of their own product or, in some cases negative reviews of competitors' products. Perhaps most importantly, however, the customers get more visibility for their books. In fact, recent IS research shows that fake reviews could indeed increase online visibility (Lappas, Sabnis, and Valkanas 2016). In the information overflow of today, it can be notoriously hard to stand out from the masses, so buying reviews can appear as the only conceivable way for the author to get noticed. Thus, from the teleological/consequentialist perspective, buying fabricated reviews could be considered acceptable as it entails several positive outcomes: the customers benefit from higher visibility and sales, the workers benefit from the paid salary, and PBR benefits from increased business. In addition, the customer might consider its book having a positive impact on the reader, and therefore selling more of it would be a desirable goal. Deontological perspective, however, reminds that selling one's products by paying others to lie about them is ethically wrong and as such not acceptable.

Assuming a stockholders' perspective, it could be deduced that buying reviews may be considered acceptable, since both the customers and PBR benefit from the immediate transaction. However, such actions do not stand a critical review from the stakeholder perspective, due to the same reasoning as discussed earlier. Specifically, while the customers are incentivized to maximize their own profits, they

⁷ Amazon (2016) Community Guidelines state the following:

“In order to preserve the integrity of Community content, content and activities consisting of advertising, promotion, or solicitation (whether direct or indirect) is not allowed, including:

- *Creating, modifying, or posting content in exchange for compensation of any kind (including free or discounted products) or on behalf of anyone else.*
- *Offering compensation or requesting compensation (including free or discounted products) in exchange for creating, modifying, or posting content.”*

do little to minimize disadvantages to Amazon, the Internet community, and their competitors whose sales might suffer unjustly due to potentially untruthful reviews. Lastly, while their operation might be legally sound, customers fail to heed Amazon's Community Guidelines that forbid offering compensation for creating, modifying, or posting content. These guidelines can be seen as codified rules designed to moderate the behavior and interaction of Amazon community members on the website.

4.2.2 Workers

The workers who write fabricated reviews on demand get financial compensation of their efforts, while not facing a risk of penalty. Thus, such a situation is a fertile ground for opportunistic behavior (Chatterjee et al. 2015, p. 58), which the workers can justify by using teleological reasoning. Although the workers might not be sincere in their reviews, they might convince themselves that the reviews would result in higher overall gain as the consequences are positive for several key actors. The worker gets paid, the customer gets the service, and PBR gets business. Although the dishonesty of the reviews has negative influence on the credibility of Amazon and the overall quality of information in the Internet, the workers might reason that it is unlikely that a few paid reviews would make any difference in the grand scheme of things. However, if one was to use deontological reasoning, writing fabricated reviews would be considered unacceptable since lying is categorically wrong and thus not a justifiable means, no matter how desirable the end outcome might be.

While stockholder perspective might not be the most relevant logic from a worker's perspective, writing fabricated reviews can be thought to benefit the main stockholders involved in crowdturfing, namely customers and PBR. However, the situation looks different when taking other stakeholders into account: the credibility of Amazon suffers, and the quality of information available to other Internet users deteriorates. Similarly, the lens of social contract theory reveals that while workers maximize their own advantages by crowdturfing, they do not simultaneously minimize the disadvantages to Amazon and the Internet community.

Lastly, Amazon's Community Guidelines state that "creating, modifying, or posting content in exchange for compensation of any kind (including free or discounted products) or on behalf of anyone else" is not allowed. Thus, these codified rules explicitly forbid reviewers to write reviews in exchange for compensation. However, writing an inauthentic comment on a product does not seem to be of significance to invoke law enforcement.

4.2.3 Agent

From the agent's perspective (in this case, PBR), connecting workers with customers can be thought to result in positive economic consequences, as their platform provides workers an opportunity to earn money and customers an access to affordable labor. Thus, providing such services could be considered acceptable from teleological perspective, since there are positive consequences for several stakeholders. However, from deontological perspective attracting and encouraging workers to lie and customers to pay for lies is categorically wrong and therefore not acceptable.

When taking the lens of a stockholder, PBR's activity is not problematic, since both PBR's and their clients' stockholders benefit. Then again, PBR does not take into account the needs of other stakeholders, i.e. Amazon and the Internet community, as its service stimulates behavior that damages both. Moreover, from the social contract perspective, PBR maximizes only their own and their customers' advantages. Not only it does nothing to minimize the disadvantages to Amazon and the Internet community from their actions but it in fact increases these disadvantages: the more review writers they succeed to connect with customers who are selling products in Amazon, the more unreliable Amazon's rating system becomes, and less reliable information will be available for the Internet users. Finally, since PBR does not in essence operate directly in Amazon's online platform, it does not directly violate Amazon's codified rules. However, it does invite and encourage other actors to violate them, promoting fraudulent activities.

Ethical logic	Actor		
	Customer	Worker	Agent
Teleological	DO IT: Buying favorable reviews gives the author more visibility, which would be hard to get otherwise. Also, this provides workers an opportunity to earn money and more business to PBR.	DO IT: Writing fabricated reviews generates earnings for the worker, service for the customers, and business for PBR.	DO IT: Connecting workers with customers has a positive economic effect as it gives workers an opportunity to earn money and customers an access to affordable labor.
Deontological	DO NOT DO IT: Selling one's products by lying and fraud is categorically wrong and should not be exercised, even if it would generate favorable results in certain respects.	DO NOT DO IT: Being untruthful is categorically wrong and therefore writing fabricated reviews should not be exercised, no matter what are the consequences.	DO NOT DO IT: Encouraging others to untruthfulness is categorically wrong and therefore should not be exercised, even if it would generate economically favorable results.
Stockholder	DO IT: Interests of the stockholders are considered: both PBR and customers benefit	DO IT: Interests of the stockholders are considered: both PBR and customers benefit	DO IT: Interests of the stockholders are considered: both PBR and customers benefit
Stakeholder	DO NOT DO IT: Interests of Amazon and the Internet community are not being considered.	DO NOT DO IT: Interests of Amazon and the Internet community are not being considered.	DO NOT DO IT: Interests of Amazon and the Internet community are not being considered.
Social contract	DO NOT DO IT: While bringing business for PBR by buying reviews maximizes the advantages of customers, PBR, and workers, it does not minimize the disadvantages to Amazon and the Internet community.	DO NOT DO IT: While writing fabricated reviews may maximize worker's own advantages, it does not minimize the disadvantages to Amazon and the Internet community.	DO NOT DO IT: While connecting workers with customers may maximize the advantages of workers, customers, and PBR itself, it does not minimize the disadvantages to Amazon and the Internet community.
Codified rules	DO IT: Paying people to write reviews is not against the law. DO NOT DO IT: It is against Amazon's Community Guidelines related to promotions and commercial solicitations: buying content.	DO IT: Writing fabricated reviews is not against the law. DO NOT DO IT: It is against Amazon's Community Guidelines related to promotions and commercial solicitations: creating content.	DO IT: Matching workers with customers is not against the law as these actors are not breaking the law. DO NOT DO IT: While matching workers to customers does not directly involve Amazon, PBR facilitates other actors to break Amazon's Guidelines.

Table 1: Ethical guidelines to crowdturfing

5 Discussion

In line with Smith (2002, p. 20), we recognize that a typical IS user can base their ethical argumentation on several logics within the perspectives discussed earlier, but usually one perspective is the main driver of users' argumentation. Consistently, our analysis of a crowdturfing case demonstrates that each key actor involved in crowdturfing can justify the moral of their behavior by leaning on one of the theories or levels of rules within a given domain. Since different ethical theories or rules tend to provide contradicting ethical guidelines, an actor must usually dismiss alternative views when conforming to the chosen view. Chatterjee et al. (2015) note that opportunism is something very inherent to all humans, and it is likely that actors involved in crowdturfing are mainly motivated to maximize their self-interest in the absence of existing sanctions. Overall, each actor can use a range of neutralization techniques (Siponen, Vance, and Willison 2012) grounded in certain ethical perspectives.

In the domain of traditional philosophical ethics, it appears that actors involved in crowdturfing lean on the teleological view while dismissing deontological view. Indeed, contributing to the production of phony information is hard to justify from the deontological viewpoint, which tends to entail categorical imperatives such as “you should not lie”. If this was the ethical guideline of every Internet user, crowdturfing business could not exist. However, the actors involved can justify their behavior by clinging to a subjective interpretation of the teleological perspective, albeit with certain limitations. Specifically, by focusing strictly on the short-term financial consequences for actors directly involved in crowdturfing campaign, one could argue that from the teleological point of view crowdturfing is moral. Moreover, a customer may often be motivated by other than financial aspirations, such as gaining visibility and recognition as an author.

When considering the business ethics approach, it is likely that actors involved in crowdturfing would rely on the stockholder theory, while dismissing the stakeholder and social contract theories. A strict adherence to Friedman’s doctrine that obligates one to focus on maximizing the stockholders’ advantage is a straightforward argument in favor of crowdturfing. The proponents of this view do not need to concern themselves with considering the interests of other stakeholders or minimizing the damages to other parties caused by their activity. In fact, if they would consider these issues, crowdturfing would be impossible to justify ethically. Furthermore, in terms of social contract theory, actors involved in crowdturfing might perceive that actions in cyberspace are distinct from actions in the real world, and thus they do not feel obliged to act according to same social contract or social norms when online. For them, what is not acceptable in the physical world, may be acceptable in the virtual world.

Although the stockholder theory obligates the actor to stay “within the rules of the game”, the actors can choose to draw the rules from the level of codified rules that suits them best, e.g., from effective legislation instead of target platform’s terms and conditions. Thus, the level of codified rules considered can determine the ethical grounds of crowdturfing. Selling, buying, or intermediating of manufactured support such as fake reviews is not necessarily illegal, so none of the actors have to concern themselves with legal quandaries. Thus, it seems that they deem the behavior acceptable at least as long as it does not violate rules on the level of the general law. Even though codified rules usually exist on the level of a platform (such as Amazon), these guidelines may not have a significant impact on the actors. This relates to a common problem in software services as such guidelines tend to be easily dismissed as trivial information that does not have any practical relevance. For instance, it was found that only 7 % of people read the terms and conditions when buying a product or service online (Smithers 2011).

Consistent with Willison and Warkentin (2013), our analysis highlights the importance of understanding the ethical roots and motivations behind behavior that threatens information validity and can thus become a security risk in the information age. Certain existing theories (e.g. transaction cost economics) alone may not be sufficient to fully explain why crowdturfing occurs since they lack this ethical perspective. Thus, our work serves as a foundation on which a deeper understanding of this behavior can be built. Moreover, we provide a more multifaceted perspective on the phenomenon: taking a different look at the three key actors immediately involved in crowdturfing, we find that although earlier we illustrated them as distinct entities from the rest of the Internet population (see Figure 1), eventually they are also Internet users who enjoy the benefits of freely available information. Thus, while contributing to the existence of crowdturfing campaigns generates financial rewards for the parties involved, it decreases the value of information available for all Internet users, including themselves.

This contradiction could be explained by the actors’ negligence of the universalizability thesis, that is, asking the question “what if everyone would do it?” If every Internet user would either buy or produce fake content or act as intermediary between the two, there would soon be no trustworthy information available in the shared cyber space. Moreover, there would be no logical demand for additional ‘followers’, ‘likes’, or ‘reviews’ if the default assumption would be that they are all untrustworthy. However, especially the workers may not believe that a small-scale production of spurious content matters on the grand scale as they might assume that most of the Internet users are still sincere, and their input will

balance the consensus close enough to reality. If this is case, the argument goes, then, why not make a few extra bucks on the side?

6 Conclusions and suggestions for future research

In this study, we have explored the ethical logic in crowdturfing through three normative approaches: traditional philosophical ethics, business ethics, and codified rules. While the lack of primary empirical data is a clear limitation of this study, to our best knowledge it is the first one to examine crowdturfing from a behavioral perspective, as a distinct form of unethical IS use. We have shed light on some of the potential ethical justifications for crowdturfing and provided a theoretical explanation for such online behavior. Specifically, we identify the ethical lenses that may allow the actors to operate with clean conscience. We have taken a holistic view on the phenomenon as we have considered the key parties involved in or affected by crowdturfing. Specifically, we have included the target platform and the Internet community as stakeholders affected by this practice. Our work may help target platforms to design effective educational or informational counter-measures against crowdturfing, e.g. applying a reminder pop-up that would inform about the impacts of posting false information and appeal to the potentially ignored ethical logics, when a platform user is posting a review.

It is important to note that this study is conceptual in nature, and as such, it is only a first step toward understanding the behavior. A natural next step would be to build on the insights we have advanced by conducting empirical work (e.g., by conducting both qualitative and quantitative studies). However, we acknowledge that gaining access to empirical data directly from actors involved in this line of business can be challenging for future researchers. While these businesses may operate within the rule of law, their moral grounds might be shaky, which could discourage the actors from collaborating with researchers. Therefore, future researchers may utilize certain strategies that could encourage crowdturfing workers to collaborate in empirical research, e.g. using scenario-based surveys that present hypothetical situations in the form of narrative vignettes. This could help in relaxing turfurs' discomfort which may be associated with the topic. Also, paying the workers for participating in data collection would probably further incentivize them to collaborate – after all, financial motives appear to be of great importance to turfurs. Additionally, future researchers could consider the perspective of the affected stakeholders by studying e.g., how victim platforms react and respond to crowdturfing campaigns (see e.g. Lappas, Sabnis, and Valkanas 2016).

Furthermore, we encourage future exploration of preventive measures. Following the example of the spamming case (Simonite 2011), tracking the flow of money to find the financial bottleneck of crowdturfing campaigns could be one approach to tackle the problem. Since financial gains are likely to hold a key position in motivating crowdturfing behavior, crowdturfing platforms' decreased ability to provide compensation for their workers could significantly reduce their attractiveness. It is obvious that research on detecting crowdturfing is still emerging, and more effective methods are needed (e.g., Lee, Webb, and Ge 2014; Siering, Koch, and Deokar 2016). Such research could potentially be able to provide target platforms with concrete tools to remedy the harmful effects by corrective actions such as flagging and eliminating the false content.

Finally, we acknowledge that deterrence theory is the dominant approach informing information security standards (Theoharidou, Kokolakis, and Karyda 2005), and information security behaviors in general (Lebek et al. 2014). Nevertheless, we should caution that its effectiveness is rather limited in the absence of severe sanctions (D'Arcy and Herath 2011; Levin, Dato-on, and Manolis 2007), and most importantly when individuals apply neutralization tactics (Siponen and Vance 2010). Thus, we propose that a deeper comprehension of pre-kinetic events preceding crowdturfing behavior is a good starting point in the search for means to dissuade or deter it.

References

- Abokhodair, Norah, Daisy Yoo, and David W. McDonald (2015), "Dissecting a Social Botnet," *Proceedings of the 18th ACM Conference on Computer Supported Cooperative Work & Social Computing - CSCW '15*, 839–51.
- Afuah, Allan and Christopher L. Tucci (2012), "Crowdsourcing as a solution to distant search," *Academy of Management Review*, 37 (3), 355–75.
- Amazon (2016), "Community Guidelines," [available at <https://www.amazon.com/gp/help/customer/display.html?nodeId=201929730>].
- Boukef, Nabila and Mohamed Charki (2014), "When the dark side of post-adoptive use leads to IT discontinuance: An exploration of the role of intervention," in *AMCIS*, 1–9.
- Chatterjee, Sutirtha, Suprateek Sarker, and Joseph S. Valacich (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.
- CMA (2015), "Online reviews and endorsements: Report on the CMA's call for information."
- D'Arcy, John and Tejaswini Herath (2011), "A Review and Analysis of Deterrence Theory in the IS Security Literature: Making Sense of the Disparate Findings," *European Journal of Information Systems*, 20 (6), 643–58.
- and Anat Hovav (2007), "Deterring internal information systems misuse," *Communications of the ACM*, 50 (10), 113–17.
- Davenport, David (2002), "Anonymity on the Internet: why the price may be too high," *Communications of the ACM*, 45 (4), 33.
- Elder, Miriam (2012), "Hacked emails allege Russian youth group Nashi paying bloggers," *The Guardian*.
- Estellés-Arolas, Enrique and Fernando González-Ladrón-de-Guevara (2012), "Towards an integrated crowdsourcing definition," *Journal of Information Science*, 38 (2), 189–200.
- Fareed, Malik (2008), "China joins a turf war," *The Guardian*.
- Gani, Aisha (2015), "Amazon sues 1,000 'fake reviewers,'" *The Guardian*.
- Gibbs, J. P. (1975), *Crime, punishment, and deterrence*, Elsevier.
- González, Ángel (2016), "Amazon sues alleged providers of fake reviews," *The Seattle Times*, April 25.
- Gordon, Mary Ellen and Kathryn De Lima-Turner (1997), "Consumer attitudes towards A social contract perspective," *International Marketing Review*, 14 (5), 362–75.
- Han, Rongbin (2015), "Manufacturing Consent in Cyberspace: China's 'Fifty-Cent Army,'" *Journal of Current Chinese Affairs*, 44 (2), 105–34.
- Henry, Robin (2015), "How to fake a bestseller," *The Sunday Times*.
- Howard, Philip N. (2003), "Digitizing the Social Contract: Producing American Political Culture in the Age of New Media," *The Communication Review*, 6 (3), 213–45.
- Howe, Jeff (2006), "The Rise of Crowdsourcing," *Wired*.
- Kennedy, Kevin C (1983), "Critical Appraisal of Criminal Deterrence Theory, A," *Dick. L. Rev.*, 88, 1.
- Lappas, Theodoros, Gaurav Sabnis, and Georgios Valkanas (2016), "The Impact of Fake Reviews on Online Visibility: A Vulnerability Assessment of the Hotel Industry," *Information Systems Research*, Articles i (November), 1–22.
- Lebek, Benedikt, Jörg Uffen, Markus Neumann, Bernd Hohler, and Michael H. Breitner (2014), "Information security awareness and behavior: a theory-based literature review," *Management Research Review*, 37 (12), 1049–92.
- Lee, Kyumin, Steve Webb, and Hancheng Ge (2014), "The dark side of micro-task marketplaces:

- Characterizing fiverr and automatically detecting crowdturfing,” *Association for the Advancement of Artificial Intelligence*, 275–84.
- , ———, and ——— (2015), “Characterizing and automatically detecting crowdturfing in Fiverr and Twitter,” *Social Network Analysis and Mining*, 5 (1), 1–16.
- Leetaru, Kalev (2016), “Is The Internet Evolving Away From Freedom of Speech?,” *Forbes*.
- Leimeister, Jan Marco, Michael Huber, Ulrich Bretschneider, and Helmut Krcmar (2009), “Leveraging Crowdsourcing: Activation-Supporting Components for IT-Based Ideas Competition,” *Journal of Management Information Systems*, 26 (1), 197–224.
- Levchenko, Kirill, Andreas Pitsillidis, Neha Chachra, Brandon Enright, Márk Félegyházi, Chris Grier, Tristan Halvorson, Chris Kanich, Christian Kreibich, He Liu, Damon McCoy, Nicholas Weaver, Vern Paxson, Geoffrey M. Voelker, and Stefan Savage (2011), “Click trajectories: End-to-end analysis of the spam value chain,” *Proceedings - IEEE Symposium on Security and Privacy*, 431–46.
- Levin, Aron M., Mary Conway Dato-on, and Chris Manolis (2007), “Deterring illegal downloading: The effects of threat appeals, past behavior, subjective norms, and attributions of harm,” *Journal of Consumer Behaviour*, 6 (2–3), 111–22.
- Loch, Karen D, Houston H Carr, and Merrill Warkentin (1992), “Threats to Information Systems: Today’s Reality, Yesterday’s Understanding,” *MIS Quarterly*, 16 (2), 173–86.
- Majchrzak, A. and A. Malhotra (2013), “Towards an information systems perspective and research agenda on crowdsourcing for innovation,” *Journal of Strategic Information Systems*, 22 (4), 257–68.
- Mason, Richard O. (1986), “Four ethical issues of the information age,” *MIS Quarterly*, 10 (1), 5–12.
- Moor, James H (2001), “The future of computer ethics: You ain’t seen nothin’ yet!,” *Ethics and Information Technology*, 89–91.
- Orphanides, K.G. (2015), “Can Amazon’s legal action stem the tide of fake reviews?,” *Wired*, October 19.
- Perez, Sarah (2016), “Amazon cracks down on fake reviews with another lawsuit,” *TechCrunch*, April 26.
- Price, Shayla R. (2014), “Big Business: Buying Fake Instagram Followers,” *The Huffington Post*.
- Rubin, Ben Fox (2015), “Amazon sues alleged reviews-for-pay sites,” *CNET*.
- (2016), “Amazon continues crackdown on alleged fake reviews,” *CNET*.
- Schlagwein, Daniel and Niels Bjørn-Andersen (2014), “Organizational Learning with Crowdsourcing: The Revelatory Case of LEGO,” *Journal of the Association for Information Systems*, 15 (11), 754–78.
- Shaffer, Andrew (2013), “Five Stars for Five Dollars: Buying Reviews, Reviewed,” *The Huffington Post*.
- Siering, Michael, Jascha-Alexander Koch, and Amit V. Deokar (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.
- Simonite, Tom (2011), “Hidden Industry Dupes Social Media Users,” *MIT Technology Review*.
- Siponen, Mikko (2006), “A justification for software rights,” *ACM SIGCAS Computers and Society*, 36 (3), 11–20.
- and Juhani Iivari (2006), “Six Design Theories for IS Security Policies and Guidelines,” *Journal of the Association for Information Systems*, 7 (7), 445–72.
- and Anthony Vance (2010), “Neutralization: New Insights into the Problem of Employee Information Systems Security Policy Violations,” *MIS Quarterly*, 34 (3), 487–502.

- , ———, and Robert Willison (2012), "New insights into the problem of software piracy: The effects of neutralization, shame, and moral beliefs," *Information and Management*, 49 (7–8), 334–41.
- Smith, H. Jeff (2002), "Ethics and information systems," *ACM SIGMIS Database*, 33 (3), 8–22.
- Smithers, Rebecca (2011), "Terms and conditions: not reading the small print can mean big problems," *The Guardian*.
- Soliman, Wael (2015), "People-Driven, ICT-Enabled Innovation: Crowdsourcing," Helsinki: Aalto University Publication Series; Doctoral Dissertations 173/2015, 2015.
- Song, Jonghyuk, Sangho Lee, and Jong Kim (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.
- Straub, Detmar W. (1990), "Effective IS security: An empirical study," *Information Systems Research*.
- Straub, Detmar W Jr. and William D Nance (1990), "Discovering and Disciplining Computer Abuse in Organizations: A Field Study," *Misq*, 14 (March), 45–60.
- Theoharidou, Marianthi, Spyros Kokolakis, and Maria Karyda (2005), "The insider threat to Information Systems and the effectiveness of ISO 17799," *Computers & Security*, 24 (6), 472–84.
- U.S. Code title 44 (n.d.), "44 U.S. Code § 3542."
- Wall, Jeffrey D., Paul Benjamin Lowry, and Jordan B. Barlow (2016), "Organizational Violations of Externally Governed Privacy and Security Rules: Explaining and Predicting Selective Violations under Conditions of Strain and Excess," *Journal of the Association for Information Systems*, 17 (1), 39–76.
- Wang, Gang, Christo Wilson, Xiaohan Zhao, Yibo Zhu, Manish Mohanlal, Haitao Zheng, and Ben Y. Zhao (2012), "Serf and Turf: Crowdturfing for Fun and Profit," *Proceedings of the 21st international conference on World Wide Web. ACM.*, 679–88.
- Willison, Robert and Merrill Warkentin (2013), "Beyond deterrence: An expanded view of employee computer abuse," *MIS Quarterly*, 37 (1), 1–20.