

Characterizing Online Vandalism: A Rational Choice Perspective

Kaylea Champion
University of Washington
Seattle, WA, USA
kaylea@uw.edu

ABSTRACT

What factors influence the decision to vandalize? Although the harm is clear, the benefit to the vandal is less clear. In many cases, the thing being damaged may itself be something the vandal uses or enjoys. Vandalism holds communicative value: perhaps to the vandal themselves, to some audience at whom the vandalism is aimed, and to the general public. Viewing vandals as rational community participants despite their antinormative behavior offers the possibility of engaging with or countering their choices in novel ways. Rational choice theory (RCT) as applied in value expectancy theory (VET) offers a strategy for characterizing behaviors in a framework of rational choices, and begins with the supposition that subject to some weighting of personal preferences and constraints, individuals maximize their own utility by committing acts of vandalism. This study applies the framework of RCT and VET to gain insight into vandals' preferences and constraints. Using a mixed-methods analysis of Wikipedia, I combine social computing and criminological perspectives on vandalism to propose an ontology of vandalism for online content communities. I use this ontology to categorize 141 instances of vandalism and find that the character of vandalistic acts varies by vandals' relative identifiability, policy history with Wikipedia, and the effort required to vandalize.

CCS CONCEPTS

- **Security and privacy** → *Social aspects of security and privacy;*
- **Human-centered computing** → **Wikis**; *Empirical studies in collaborative and social computing; Computer supported cooperative work; Collaborative content creation; Social media.*

KEYWORDS

rational choice, RCT, vandalism, wikipedia, anonymity, graffiti, online community

ACM Reference Format:

Kaylea Champion. 2020. Characterizing Online Vandalism: A Rational Choice Perspective. In *International Conference on Social Media and Society (SMSociety '20)*, July 22–24, 2020, Toronto, ON, Canada. ACM, New York, NY, USA, 11 pages. <https://doi.org/10.1145/3400806.3400813>

1 INTRODUCTION

Vandalism offers asymmetric power to the vandal. The object being damaged may originate in the effort of a powerful entity such as a

large institution or a government or a large group of well-meaning people. A vandal's momentary action with a hammer or pen can be profoundly disruptive. In physical environments, communities may expend substantial resources to restore vandalized property: public works must be rebuilt, damaged artwork may be irretrievably lost, and time and money spent removing defacement could be deployed in other ways.

Online communities also struggle with vandalism. Automation in digitally-mediated environments can aid vandals who use tools to vandalize with incredible speed and volume [10]. Although automated tools support rapid restoration of the digital work to its original state—often within minutes [2], automated vandal fighting is not without costs. It requires development, invocation, and maintenance of tools and it leads to impersonal messages and improper rejection of contributions which has been implicated in a decline of well-intentioned newcomers [15], a trend which requires substantial effort to remediate [16, 25]. The digital nature of online vandalism does not mean it is victimless—besides offending and burdening hard-working creators, maintainers, and moderators [11], the general public suffers as well. On highly visited websites, some information-seekers are likely to read the damaged page regardless of how quickly vandalism is removed.

Given the social and individual cost and the difficult to discern source of benefit to the vandal, we may ask why it occurs at all. Competing theories in criminology seek to explain the motivations for and causes of crime, ascribing criminal behavior to such factors as lack of impulse control, lack of morals, or to societal failure. Alternatively, rational choice theory proposes that behaviors are the product of rational choices. In order to apply rational choice theory to vandalism, this project seeks to understand vandal decision-making in terms of preferences and constraints. Building on past examinations of Wikipedia contributors, we examine vandalism from four groups: users of a privacy tool, those contributing without an account, those contributing with an account for the first time, and those contributing with an account but having some prior edit history [4, 34].

This paper makes three contributions: I describe the extension of rational choice theory and value expectancy theory to Wikipedia, I offer an extension to the existing ontologies for characterizing online vandalism, and I demonstrate the application of both rational choice theory and this vandalism ontology in a preliminary series of findings. The paper is structured as follows. I offer background on the topic of vandalism and theories of rational choice in Å§2. I describe the comparison groups in Å§3 and my approach to analysis in Å§4. I share the results of classifying vandalism and examining the relative prevalence of each type in Å§5 with limitations as described in Å§6. In Å§7 I explore the implications of these findings, as well as how these results can inform design and policy in online communities. I conclude in Å§8.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

SMSociety '20, July 22–24, 2020, Toronto, ON, Canada

© 2020 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-7688-4/20/07.

<https://doi.org/10.1145/3400806.3400813>

2 BACKGROUND

Wikipedia calls itself “the free encyclopedia that anyone can edit” and seeks to present encyclopedic information from a neutral point of view, with freely usable, editable, and distributable content, developed by a respectful community with no firm rules.¹ Like a public park, Wikipedia is a public good—defined as non-rivalrous in that one person’s use does not consume it and non-excludable in that it is open to everyone. Technological advances have allowed for the collaborative development of new public goods through a process called commons-based peer production in which self-directed volunteers contribute according to their interests, whether small and casual or vast and serious [3]. However, in opening editing power to almost anyone, Wikipedia opens itself up to the possibility that not all participants engage in good faith to develop a high-quality, neutral encyclopedia. Instead, some contributors seek to do harm.

As one of the top 10 most visited sites on the internet, damage to Wikipedia holds the potential to harm information seekers in multiple ways. Vandalism that deletes or disrupts a page will block access to information. Insertion of deliberate misinformation can generate rumors or distort understanding of a topic. Vandalism can expose the general public to threats or harassment.

Other vandalism in Wikipedia is targeted to damage the production community itself, such as threatening and harassing contributors and administrators, wasting volunteer time by spamming requests for help when no help is needed, or by disrupting the behaviors of automated tools.

Some Wikipedia vandals write their names, add jokes, or make rude comments. These kinds of contributions are also characterized as vandalism in Wikipedia,² although the content of the text is in some cases indistinguishable from what in the physical world would be called graffiti. Although graffiti has long been analyzed as a source of meaning in contexts ranging from ancient Pompeii to modern university library bathrooms [7], the addition of unwelcome or disruptive contributions to Wikipedia is contrary to its use, making vandalism an apt description regardless of text. The Wikipedia community has a tradition of joking about persistent vandalism.³

Just as volunteers built the public good that is Wikipedia, they work to defend its utility from damage. However, countering vandalism in Wikipedia is a substantial undertaking. Vandal-fighters and their tools are recognized as an important element of the community [11, 15]. Vandalism-related research has tended to focus on the detection and removal of vandalism [33], with relatively little attention paid to understanding vandals themselves. Notable exceptions include Shachaf and Hara [29] who considered the cases of four people who Wikipedia community administrators (“sysops”) had identified as “trolls” who had engaged in acts of vandalism, including replacing main page photos with pornography and writing threats. In their interviews with sysops, Shachaf and Hara found that the sysops interpreted trolls as having engaged in harmful

edits intentionally, repetitively, and in violation of policy, and that they tended to target the community itself for damage. The sysops further interpreted the trolls as being motivated by “boredom, attention seeking, revenge....fun and entertainment....[and desire to do] damage to the community” [29, p. 357].

Sierra and Castanedo [30] asked focus groups of students and Wikipedia editors to speculate as to why people vandalize. Students thought vandals might be making a joke, acting out of boredom, or engaging in ideological protest. Wikipedia editors added to these potential motives the possibility of experiencing a sense of pride at defeating Wikipedia’s defenses against these vandalism. These interpretations of vandal behavior suggest that increasing the time required to vandalize may reduce more casual and sociable forms of vandalism, but that some kinds of vandals may find obstacles part of the appeal. Cruz et al. [6] was able to interview individuals identified as trolling in an online community, as well as witnesses to trolling, and found that effective trolling requires knowledge of a community’s rules and norms. This suggests that to the extent that some vandalism qualifies as trolling, it may emanate from community members who know the norms well enough to intentionally transgress them. Such individuals may have more to lose if their behaviors lead to their accounts being blocked and may seek out ways to avoid identification.

There are many types of vandalism. Some have only a mildly negative impact, and others might even carry a positive impact alongside their harm. Acts of vandalism may be amusing, insightful, or comprise a political protest. Some scholars have found evidence that unwanted behavior on Wikipedia serves as a signal of attention from the general public, which in turn may inspire further efforts or improvements from the community. [14] That said, just as janitorial staff scrub away even the most supportive and friendly bathroom scribble, even the most innocuous forms of vandalism are removed from Wikipedia to protect the integrity of the resource.

2.1 Theoretical Framing

One may consider vandalism on Wikipedia as an example of criminality and antinormative behavior in a broad sense and seek insight from social theories used to understand deviance. One important perspective is rational choice theory (RCT), which has found applications in economics, biology, psychology, and sociology. Although multiple articulations exist, the version of RCT I use states that people select actions which they believe will maximize their utility with consideration of their preferences and any constraints associated with that action [27]. I use a ‘wide’ version of RCT which Goldthorpe [13] recommends as most suitable for understanding social action and relations. The wide version incorporates an understanding that people make choices with incomplete information, subject to bounded rationality, with concern for benefits beyond material wealth (such as social approval and a positive emotional state), and nuanced models of personal preferences, traits, and fears [24, 27].

To assist in the application of RCT to analysis of specific acts, I use Value Expectancy Theory (VET), which is derived from RCT. Riker and Ordeshook [28] articulate value expectancy theory as expected utility E , of some behavioral alternative a (identified

¹https://en.wikipedia.org/wiki/Wikipedia:Five_pillars

²<https://en.wikipedia.org/wiki/Wikipedia:Vandalism>

³For example, this humorous essay describes categories of vandals as ‘typing students’, ‘the curious’, ‘critics’, ‘men with big penises’, ‘cheerleaders’, and ‘friends of gays’, the last being a group of people who the essay calls “overly proud friends” of gay and lesbian people, i.e. those whose vandalism consists of texts like “[name] is gay.” The essay is available at: https://meta.wikimedia.org/wiki/Friends_of_gays_should_not_be_allowed_to_edit_articles

specifically as alternative i among some list of alternatives) calculated as the sum of the product of two considerations: the probability of some outcome O (identified specifically as outcome j among some list of outcomes 1 through N) and the utility valuation U of that same outcome, O_j . Expected utility is thus $E(a_i) = \sum_{j=1}^N p_{ij}(O_j)U(O_j)$. With this definition in hand, we can express RCT in terms of a simple relationship among expected utilities. If the expected utility of an action i exceeds the utility of some other action j , then rational choice predicts the selection of action i : $E(a_i) > E(a_j) \Rightarrow a_i$.

Utility maximization of this kind need not be assessed numerically for the proposed relationship to hold. A would-be vandal might select among alternatives iteratively, examining each possible pair in turn and choosing the preferred option. After this round-robin process, if one's thinking process is consistent, a preference ranking of activities would emerge [12]. In order to discern whether online vandals make rational choices consistent with VET, I examine vandalism and consider the factors relevant to VET in building up hypotheses: what constraints, what preferences, and what utility applies?

2.2 Constraints

Given the significance of Wikipedia, the freedom with which contributors improve or damage its content, and the substantial challenge of protecting it, we might ask what factors act as constraints against vandals. Although editing is free of charge, anyone may click the edit button, make some changes, and save them. Although this barrier sounds relatively low, relatively few people actually contribute to Wikipedia compared to the number who read it. For example, English Wikipedia, the largest of the 302 different language editions, with over 6 million articles, received over 11 billion page views in April 2020 from almost 1 billion unique devices, but the total number of unique editors in that month was only 424,155, of whom only 68,481 made five or more contributions⁴. Hence it may be that even a very low level of effort is nonetheless meaningful. The Wikipedia sysops interviewed about their impression of trolls' motives by Shachaf and Hara [29] included "ease of execution" as an important factor in troll behavior, suggesting that time constraints may be relevant. Creating an account costs more in time and energy than making the same contribution without an account because a person must fill out a brief form and choose a unique name and acceptable password.

Because VET predicts that users who have expended more effort are less likely to choose deviant behavior, and I hypothesize that **(H1) users who have created accounts will vandalize less frequently**.

A second potential constraint against vandalism is fear of detection. Given that internet activities are routinely monitored by employers, school officials, parents, internet service providers, and governments, vandals may be discouraged by a concern that they will be discovered. One component of being detected is being identified. Friedman and Resnick [9] describes the decision between using an anonymous pseudonym and a persistent identifier as "a

strategic variable" in the interaction between participants in a co-operative online platform because it allows control over the spread of a user's reputation (p. 174).

Wikipedia offers two types of identifiability for contributors. First, users can contribute without making an account in which case their IP address is publicly associated with their actions. Second, users can create and then log in to an account to which all subsequent contributions will be attributed.⁵ Despite its arcane appearance as a sequence of numbers and/or letters, a user's IP address is identifiable information, and is regarded as such under privacy regulations such as the GDPR. Although individuals may not be aware of the fact, IP address can be used to identify a user's geographic location. With assistance from a service provider, the IP address can identify the individual home or even computer. Wikipedia's interface encourages contributors to create an account, telling them that doing so will keep their IP address private from other contributors and the general public.

In addition to choosing between these two levels of identifiability, users can take other steps to hide their identity, including naming themselves according to some pseudonym unlinked to their identity, avoiding sharing personal information online, using multiple, shared, or public IP addresses, or using an anonymity service to mask their IP address. Detecting whether users mitigating against detection is difficult. We can address this limitation with a unique dataset from Tran et al. [34], who identified contributions made using the Tor anonymity service.

The Tor anonymity network allows users to access the Internet while maintaining IP privacy by relying on volunteers to transport network traffic using a method called "onion routing." Typical network routing creates a traceable path of intermediate hops over which each network packet travels such that each hop knows the identity of every other node in the network. Onion routing only gives each layer—each step in the network path—the identity of the next "hop." Hence, none of the intermediate steps know the identity of the point of origin, and none of them know the destination; they only know the next step [18].

Although Tor is sometimes characterized as a tool for criminals and is blocked or limited by some online services [23], privacy-protecting tools are a crucial line of defense for people living under oppressive regimes (e.g. Iran, [26] or Turkey [1]), or whose identities or interests subject them to harassment and attacks, or who simply seek relief from surveillance and the monetization of their online behavior [8, 20, 35]. Although this prior work often tends to explore broadly normative behavior that is nonetheless potentially costly to the speaker, privacy techniques can certainly be used to protect high-risk anti-normative behaviors as well. Because VET suggests that identifiability acts as a constraint on deviant behavior, I hypothesize that **(H2) populations which differ in their identifiability produce different types of vandalism, such that the least identifiable individuals are more likely to produce vandalism that has high-risk repercussions**.

2.3 Preferences

The second component of a VET model is the role of preferences. In the case of vandalism, preferences can be observed in the types of

⁴<https://stats.wikimedia.org>

⁵<https://en.wikipedia.org/wiki/Special:CreateAccount>




Figure 1: Regardless of whether they are logged in or not, a Tor user will see a message like the one above if they attempt to edit Wikipedia. In this case, the author is seeking to edit her own personal user page while logged in. The message states that editing one's own User Talk page and e-mailing administrators is still allowed.

vandalism a vandal chooses. For example, if a vandal were simply desiring to do damage, any kind of vandalism might do. However, if there is some communicative intent or a specific target to be harmed, we might expect to see themes in the impact or victim/s of vandalism.

This study compares vandals to understand whether, pursuant to **H2**, they differ in their willingness to undertake higher risk vandalism based on their expressed privacy preferences. In order to do so, I examine privacy-seekers using Tor. Although Wikipedia forbids contributing via Tor, this policy was not always consistently or correctly applied, with the bulk of edits evading the block occurring between 2008 and 2013. During this time, about 11,000 edits were either lucky enough to fall through the cracks or able to evade the block [34]. Some of the individuals using Tor to edit Wikipedia may be privacy-seekers who simply got lucky and have no idea that Tor is blocked. Others from the sample are clearly aware of the contentious relationship between Tor and Wikipedia, and taunt Wikipedia administrators about their inability to block them [4].

As described, the groups under study differ by how they are treated by community policies. Newcomers are targeted for social interventions to welcome, train, and retain them. Wikipedia invites IP-based editors to create accounts as well as welcoming them. However, Tor-based editors generally experience rejection, despite the fact that their contributions (when they can make them) are of similar quality to newcomers and IP-based editors [34]; Tor-based editors can only contribute through a combination of luck and determination. Thus it seems reasonable to examine an additional relationship between specific group traits and vandalism type. VET predicts that preferences will drive target selections, and this suggests a final hypothesis: **(H3) Members of excluded groups are more likely to strike against the community targeting them.**

2.4 Conceptual model

Figure 2 uses VET to elaborate a conceptual model of vandalism. Read from left to right, the figure suggests that the decision to vandalize Wikipedia (as opposed to read it, contribute to it, or do nothing) begins with personal preferences and motives as well as consideration and experience of constraints. The model anticipates an individual weighing the combinations of preferences and constraints with their relative likelihoods and choosing the action that they find most suits their preferences.

3 DATA

The data for this study was assembled from a random sample of article revisions made by members of four different groups of editors that differ in their identifiability and in their relative time investment in editing Wikipedia: (1) users of the Tor privacy service whose identity is fully masked that I call "Tor-based editors" ($n = 500$); (2) users who contribute without making an account but whose IP addresses are disclosed in place of a username making it highly likely that their location can be identified [21] (although they may not know this) that I call "IP-based editors" ($n = 200$); (3) users who have made an account and who are making their first contribution that I call "New Contributors" ($n = 200$); and (4) users with accounts for whom a given edit is not their first that I call "Experienced Contributors" ($n = 200$). The available population of Tor-based edits is relatively small (approximately 7,700), and the Tor-based sample is a random sample of this population. For the other three groups, I drew a random sample to time-match the Tor population on a monthly basis (e.g., if our Tor population made 30 edits in January 2009, I would randomly draw 30 edits from registered editors in January 2009, and so on.) I then drew a random sample from these time-matched samples.




Figure 2: Conceptual model of vandalism. In the ‘maximize utility’ stage, Value Expectancy Theory argues that individuals weigh out their possible actions against their possible outcomes, their value, and their likelihood, then choose the action they find the most optimal.

These user categories differ both in their privileges in Wikipedia [34] and in their privacy concerns [8].

All IP editors have fewer privileges on the site and are automatically scrutinized at higher levels. Contributors with accounts receive progressively less scrutiny and increase in their site privileges as time passes. The Wikipedia community also elevates some members for additional administrative powers. These individuals have additional abilities to enforce norms and sanction violators, including banning them from contributing. Established contributors are also typically scrutinized less by algorithms [17]. The four groups are summarized in Table 1.

4 METHODS

This study followed a mixed-methods approach, starting with a two-stage qualitative content analysis followed by statistical analysis. I made use of digital trace data made public by the Tor Project and by Wikipedia, benefiting from a sample prepared in the course of Tran et al.’s [34] comparison of Tor-based contributors to the same contributor groups I examine in this paper. Because the research involved only public data and did not involve interaction or intervention, the research was determined to not be human subjects research by the IRB at the University of Washington. Despite this, I have omitted names of most editors and articles and, in some cases, paraphrased quotes in order to make reidentification more difficult.

Identifying vandalism from the sample described above made use of my experience as both an editor of Wikipedia and a researcher who has observed the community for several years.

I defined vandalism using a three-part definition, where all three criteria must be met. First, an edit must be something which does not comply with community norms or is not encyclopedic, and hence should be removed. In the language of Wikipedia, these are “damaging” edits. Second, the edit must not seem to be intended to contribute to the article in a direct and constructive way. An edit can be damaging but made in good faith if it violates some rule about formatting but otherwise appears well-intentioned. Third, an edit should not be part of an “edit war.”

An edit war involves multiple parties changing the content of a particular article back and forth. Edit wars may be due to competing visions of what is correct (e.g., moving an article back and forth between English spelling and British spelling), or due to competing beliefs about what is true (e.g., a dispute about how best to summarize the plot of a novel or the reasons for a historical event), or due to competing beliefs about the nature of knowledge (e.g., describing herbal medicine as pseudoscientific versus a reasonable and safe

alternative). According to Wikipedia policy, disputes over content should be resolved through discussion and, if needed, arbitration—not by continually doing and undoing the work of others. Edit wars, while damaging to the article and contrary to community norms about engaging in good faith, are ultimately content-level disputes. By contrast, vandalism disrupts content and seeks to amuse, offend, disrupt, make comments to, or deliberately misinform the reader. In summary, vandalism in Wikipedia is damaging, bad-faith, and communicative or meta-communicative. The proposed analysis required that I identify and characterize acts of vandalism, a process which I will now describe. The categorization schema is used is drawn from research on graffiti and vandalism in physical locations as well as specifically within Wikipedia. After categorizing each act of vandalism by its type, I used statistical tests to assess if observed differences were statistically significant.

4.1 Constructing an Ontology of Vandalism

In order to classify vandalism into types, I extend two ontologies—one from Chin et al. [5] which describes vandalism in the context of Wikipedia, and one from White [37] which elaborates types of graffiti in the physical world. My proposed synthesis is summarized with working definitions in Table 2.

One extension in this proposed ontology is the addition of “attack graffiti.” While attacks on groups is combined with political graffiti by White, my contention is that while racial slurs obviously have a politics, so might many statements of opinion or fact. However, racial slurs are qualitatively different from, say, calling on the ruling party to resign, and collapsing the two together does not improve clarity. Similarly, White combines attacks on individuals with social graffiti and toilet humor. I distinguish between “social graffiti” including goofy nonsense and lighthearted toilet humor and “attack graffiti” including intimidation, insults, and threats. These distinctions are not always immediately clear and require some interpretation, just as some nuance exists between teasing and bullying.

My proposed ontology does not distinguish between what White calls “political graffiti” and “protest graffiti.” White defines protest graffiti as a subset of political graffiti targeting the content of “mainstream commercial visual objects.” In the context of an online encyclopedia, versus a physical environment full of advertisements and signs, a distinction singling out commercial targets is less salient. Additionally, White adopts the emic vocabulary of street graffiti to describe tagger graffiti versus gang graffiti. White describes tagging as a message of “I’m here” while gang graffiti asserts power and

Table 1: Constraints and preferences: vandalism has differing costs and risks for different kinds of uses.

Editor Population	Effort Required	Privacy Risk	Privacy Awareness
Tor-based editors	low to high: may require several minutes to initiate or fail due to blocks	little to none	likely to be privacy-aware
IP-based editors	none: no sign-up required	location made public	may be unaware
New Account	low to moderate: registration process	location admin-visible	may be unaware
Experienced Account	moderate to high: account access may be lost or reputation damaged	location admin-visible	may be unaware

control. My proposed ontology follows this underlying distinction in the categorization scheme, but renames “gang graffiti” to “pro-group graffiti” to avoid confusion. I also omit White’s category of “Graffiti Art”.

Applying the ontology proposed in this study is necessarily interpretive, derived from the coder’s personal perspective, background, and knowledge, including social norms about gender and what qualifies as a joke. The vandalistic acts which fell within the “attack graffiti” category in this analysis are varied in their severity, including gross insults, repeated name substitutions, racial slurs, and rape threats against a specific Wikipedia administrator. Indeed, all categories contain variation in their severity.

Wikipedia includes a range of features for reviewing the history of articles as well as the contributions of individual users. Figure 3 shows the interface that anyone can use to walk through the revision history of a given article and replay the timestamped sequence of the actions that built the page. As seen in Figure 4, anyone can search for and view the list of contributions made by a given user. These features help readers, contributors, and researchers understand users’ history in Wikipedia, trace interactions among contributors, and record the evolution of each page through time.

5 FINDINGS

5.1 Examples of Vandalism From Wikipedia

Despite the small sample of 141 examples of vandalism found by the manual content analysis on 1100 revisions described in Åg4, I identified examples of all but one of the categories in my proposed ontology of vandalism.

Blanking vandalism removes sections of text and entire articles. Sometimes the text is replaced, other times it is simply removed. For example, on June 19, 2010, an IP-based editor deleted a section of an article about a historical figure. The section was written in an encyclopedic tone and had several references, and the IP editor did not leave any explanation. In another example, a newly-created account using the name of a rap group deleted all of the text on the page of another rap group which appeared to be from the same US city. The lengthy text of the victimized group’s page was replaced with a brief description and a link to a myspace.com page with the same name as the vandal.

Large-scale editing involves adding a substantial amount of the text to the page. For example, a Tor-based editor inserted multiple paragraphs of repeated text stating over and over that “[Celebrity stalker] died.” in the text of several television and celebrity articles on December 24, 2008. In other cases, the text added by a Tor-based editor is varied and seemed to be making an argument, but

not one relevant to the page. One example involved a Tor user critiquing an administrative decision in Wikipedia, claiming to be an individual who had been identified and banned for misusing dozens of accounts (a practice called “sockpuppeting”), and calling other editors “evil” on multiple occasions from February to April 2008.

Misinformation vandalism refers to attempts to deliberately introduce false information or to assert information based on rumors. One example of misinformation vandalism occurred on February 26, 2013 when a Tor-based editor changed two letters in a word in a long quotation, changing the meaning of the quote. I looked at multiple primary sources and all of them contradicted this change. This misinformation persisted until November 5, 2013. Another incident of misinformation vandalism occurred on February 1, 2013, when an IP-based editor altered the biography of a broadcaster stating that he had been shot to death in his home. News reports contemporary to the edit described his death as occurring in a hospital after an illness. Wikipedia includes an option for editors to summarize their change as a note for other editors to see. In this case the vandal described their change as fixing a “minor spelling error.” However, their edit did not change any spellings.

Image attack vandalism involves substituting an irrelevant or harassing image for the legitimate image on a page. This type of vandalism is difficult to assess retrospectively because generally a removed image is no longer available online. In some cases, the nature of the attack is clearer: for example, on September 23, 2009, a new editor replaced a female politician’s photo with an image with the word “horse” in the name, which was reversed 33 minutes later.

Link spam is the insertion of links to external sites in violation of the Wikipedia external links policy. These links may be commercial in nature or connected to malware sites. For example, on February 11, 2012, a Tor-based editor inserted a link to a spam website into an article about a card game. The links were removed 7 days later as part of a group of edits removing similar links.

Irregular formatting refers to vandalism which specifically targets the wiki-specific syntax of a page. This includes special code-like symbols such as square brackets [], curly braces {}, and so on which control how the page appears. Altering a single symbol can cause the entire page to display incorrectly or to be unreadable. I observed this type of vandalism in two cases: for example, on December 17, 2012, an IP-based editor removed “}}” from a biography of a prominent female scientist, causing the top portion of the page to appear jumbled and difficult to read.

Political graffiti is vandalism that invokes or protests public policy, law, power structures, and social norms. I observed only

Table 2: Proposed Ontology of Vandalism

Category	Description	Source
Blanking	Large content deletion	Chin et al. [5]
Large-scale editing	“Massive” insert or change	Chin et al. [5]
Misinformation	Replace content with false information	Chin et al. [5]
Image attack	Insert or replace an image with an irrelevant one.	Chin et al. [5]
Link spam	Adding irrelevant links	Chin et al. [5]
Irregular Formatting	Insert or remove tags incorrectly	Chin et al. [5]
Political Graffiti	public policy, law, power structures, social norms	combines two categories from White [37]: political graffiti and protest graffiti
Attack graffiti	attack an individual or group	new
Pro-group graffiti	asserting group power	White [37]
Tagging	assertion of personal presence, bragging	White [37]
Community-related Graffiti	opposition to community, norms, or policies	new
Social Graffiti	Gossip, jokes, anat omical comments	called “Toilet and Other Public” in White [37]




Figure 3: This screenshot is taken from Wikipedia’s “diff” view. It highlights the differences between two versions of the article on Social Media. The left side shows the “before” view; “after” is on the right. The highlighted blue area with a + sign indicates what was added by the version. In this case, an IP-based editor adds a type of vandalism this paper would characterize as social (“whoop whoop big man tings”). Arrows above the text allow navigation through the revision history of the article.

one example of this type of vandalism in which an IP-based editor on December 13, 2011 added to an article which described a particular substance as having been historically believed to diminish homosexual desire stating that “these...insane theories still persist to the current day” and that “People need to learn to mind their own [expletive] business.” This mostly empty category suggests further refinement to the ontology is needed.

Attack graffiti insults, seeks to intimidate, or seeks to humiliate individuals or groups. One example of attack graffiti occurred on September 20, 2008 when a newly-created accounts edited a Wikipedia administrator’s personal messages page (called their

“User Talk” page within Wikipedia) to include a description of the administrator being raped by an individual who had recently been banned. I found that threats of this kind were made repeatedly from both IP-based user identities and newly created accounts. The administrator used a female-presenting name and the vandal in question used a male-presenting name. Another example of attack graffiti occurred on December 31, 2007, and altered a page about medieval history with text which altered the word Wikipedia to contain a reference to pedophilia and then to state that they should “...should probably start blocking Tor.” This edit was also coded as community-related graffiti.




Figure 4: This screenshot is taken from a Wikipedia user’s contribution page, in this case the contributions from the co-founder of Wikipedia, Jimmy (Jimbo) Wales. This information is available for all contributors.

Pro-group graffiti, the redefinition of “gang” graffiti for the context of Wikipedia is graffiti which asserts some group identity or ownership. The sample as coded did not contain any examples of this type of graffiti, suggesting that the ontology may need to be refined in this dimension as well.

Tagging in the context of Wikipedia is defined as actions which assert personal presence or brag in some way about an individual, presumably oneself. Of course, it cannot be determined if tagging is perpetrated by the person whose name is used or if it is done by a rival seeking to draw negative attention to an apparent braggart. One example of tagging graffiti in Wikipedia occurred on an article about a beer company, in which on February 5, 2008 a newly registered account added the phrase “[name] is a pimp” in capital letters. Another example of tagging graffiti occurred on December 17, 2007, when an IP-based editor altered a page about an online first-person-shooter video game to include five sentences describing a particular player by name as being a dominant player in the game, ending with the phrase that he “pwns noobs out da club.”

Community graffiti is vandalism that references the specific community where it appears. In my case, this type of vandalism referred directly to Wikipedia or its policies, or made use of sophisticated understanding of Wikipedia features to target Wikipedia administrators with vandalism. One example occurred on April 30, 2008 when a Tor-based editor repeatedly updated the user page associated with their own Tor IP in such a way as to trigger a malfunction in an automated script (“bot”) used by the Wikipedia community to monitor and block Tor-based users. The malfunction persisted until November 10, 2008 when a different bot cleaned up the page.

Social graffiti refers to joking comments, anatomical references, and exchanges of greetings. One example of social graffiti occurred on December 9, 2007, when an IP-based user altered a page about a

Table 3: The rate of vandalism in the sample.

Editor Population	Vandalism Prevalence
Registered Editors	1%
First-time Editors	15%
IP-based Editors	24%
Tor-based Editors	13%

technology to include the words “stupid people are so stupid” and some random characters. Another example occurred on December 3, 2011, when a newly-created account updated an article about a particular species of plant to claim that it grows in a fictional videogame universe.

5.2 Vandalism Rate

The relative prevalence of vandalism in each group is summarized in Table 3. I observed vandalism in 13% of the Tor-based edits, 24% of the IP-based edits, 1% of registered user edits, and 15% of the new editor edits. In order to understand whether the variation in rates by group is statistically significant, I used a χ^2 test of a frequency table containing the counts of vandalism and non-vandalism for each type. The χ^2 reported $p < 0.001$, indicating that the difference among user types is statistically significant.

As a robustness check, I verified that the results of the χ^2 test of the statistical significance of the differences between vandalism rates still held without considering registered editors (pursuant to **H1**); this result was confirmed with a p-value estimated at $p < .002$. This provides evidence in support of **H1** that populations which have expended more effort will vandalize less frequently.

The highest-effort group (registered editors) engaged in almost no vandalism. Tor-based editors arguably must work harder than

IP-based editors, and their rate of vandalism was lower. This observed variation in behavior by user types is somewhat surprising given prior work which found that levels of non-damaging editing behavior are roughly equivalent among the first-time, IP-based, and Tor-based editor populations [34].

The full sample of revisions examined through content analysis ($n = 1100$) was used for assessing **H1**. However, because the rate of observed vandalism was so low in registered editors (2 instances in the sample of 200 edits), further analysis was impractical. The registered editor group is omitted from consideration in **H2** and **H3**, and the only the identified acts of vandalism in the other three groups ($n = 141$) are used for these latter two hypothesis tests.

5.3 Vandalism Types

To test **H2**, that types of vandalism will vary such that groups with lower identifiability will be more likely to engage in vandalism, I determined the relative prevalence of vandalism types.

Table 4 reports both the average incidence rate of types of vandalism and whether or not the differences in the reported means are statistically significant as assessed at the $\alpha = .05$ level via a one-way ANOVA. I find partial support for **H2**. Groups that vary in their identifiability also vary in their vandalism type, with low identifiability associated with higher risk types of vandalism. Tor-based users are substantially more likely than other groups to engage in large-scale vandalism and least likely to engage in the lowest risk type of vandalism, that which communicates friendly and social intent. Although other potential high-risk types of vandalism, such as blanking, misinformation, link spamming, or attacking people, was higher among Tor-based editors, this difference was not statistically significant.

This analysis also supports **H3**. Excluded groups are more likely to vandalize in ways that strike against the community: Tor-based editors target the community itself at a much higher rate as seen in Table 4.

6 LIMITATIONS & FUTURE WORK

This study is limited in important ways. First, I rely on a qualitative coding process conducted by a single investigator. Of course, assessment of vandalism in Wikipedia is often a series of judgment calls from vandal-fighters working at high speed on their own. I am familiar with both the empirical context and several of the results are consistent with independent findings in other work.

Future work should not only use a larger number of observations and a multi-investigator coding process but also should consider additional potential constraints, such as the effort required to enact a specific act of vandalism (e.g., by counting keystrokes).

Additionally, while I have assessed each act of vandalism independently and from a random sample, some vandalistic acts are committed by serial vandals. These serial vandals may learn from the community response and adjust their approach in order to make their actions more effective. For example, Matsueda et al.'s [22] study of at-risk and delinquent juveniles found evidence for the influence of learning. Future research on social deviance in Wikipedia looking for evidence of learning might yield interesting results. For example, a vandal seeking to spread misinformation

about a celebrity death might observe that their efforts were immediately reversed and fine-tune their language to avoid words which trigger automated vandalism detection systems, or make use of conventions which build the confidence of other editors, such as citing phony sources.

This work is also limited because the study does not include direct discussion with vandals. Past studies of crime have benefited tremendously from interviewing, for example at-risk adolescents as in Matsueda et al. [22], or people observed as trolling in an online community as in Cruz et al. [6]. Although such conversations might be enlightening, the generally anonymous Wikipedia vandal population might be tremendously difficult to contact since no contact or communication mechanism is required in order to contribute to Wikipedia. Finally, this study makes only a preliminary use of RCT and VET and a future investigation could both ground the phenomena more deeply in these perspectives as well as develop expanded empirical models.

7 DISCUSSION

Taken together, the results for my three hypotheses add texture to the more general explanation often offered for antinormative behavior online in terms of online disinhibition: the notion that being more distant and less identifiable may reduce one's reluctance to give in to one's worst impulses [19, 31, 32]. Instead, we see that the least identifiable group (Tor-based editors) does not have the highest rate of vandalism. The supportive evidence found for **H1**, **H2**, and **H3** reinforce the utility of applying Rational Choice and Value Expectancy theories and suggest that identifiability and time may serve as operative constraints on vandal behavior.

One way to use these results to counter vandalism is to consider how potential interventions might change the kinds of vandalism a community receives. For example, adding a CAPTCHA [36] to slow down contributions increases costs for would-be contributors. However, given that the two groups for whom contributing is most costly (new editors and Tor-based users) seem to have a different pattern of vandalism types than other groups, increased time requirements alone may not change the rate or quality of the vandalism. Friedman and Resnick [9] proposed an alternative approach to the question of deterring antinormative behavior without requiring identifiability. They suggested anonymous certificates, using cryptographic techniques both to protect identity and to guarantee a single individual will only have one anonymous certificate. Another intervention might be to increase the sophistication of tools that detect and filter contributions before they become visible: for example, disallowing some groups from making large-scale edits or blanking pages might be desirable in some communities. Tools seeking to detect vandalism may find that the patterns that can be identified within each category allow for more accurate targeting algorithms. Although the differences in the prevalence of several other categories in this sample were not statistically significant, a larger-scale analysis or additional analytic variables to represent preferences and constraints may reveal additional trends.

For privacy advocates seeking policy changes, anti-community vandalism poses a difficult challenge. Although frustration from banned users is not surprising, targeting the community is disruptive and the experiences of harassment victims should not be

Table 4: Relative prevalence of observed vandalism types. Percentages do not sum to 100% because multiple codes were applied to individual instances of vandalism (n=141). Bold text indicates that the difference in means among the groups is statistically significant at the .05 level as determined via a one-way ANOVA. Registered editors were eliminated as a comparison group due to their low overall incidence rate of vandalism in the sample. No instances of “pro group” graffiti were identified in the sample, hence the category does not appear here.

Group	Blank	Large	Misinfo.	Image	Link	Format	Political	Attack	Wikipedia	Social	Tag
IP-based	13%	4%	9%	0%	0%	2%	2%	26%	0%	55%	9%
Tor-based	19%	14%	13%	2%	11%	2%	0%	36%	17%	19%	9%
First-time	13%	0%	7%	3%	7%	0%	0%	17%	3%	57%	17%

discounted. Given that many of Wikipedia’s defense mechanisms against damage make use of IP-based identification to block abusers as precisely as possible, accommodating those who seek IP privacy poses a substantial challenge. However, viewing vandals as rational and having diverse motives opens potential avenues for dialogue.

8 CONCLUSION

This study characterized vandalism in Wikipedia using Rational Choice Theory and the specific considerations of Value Expectancy Theory. Based on consideration of preferences and constraints, this study assessed the types and prevalence of vandalism and found that both vary among groups of vandals in ways that are predicted by hypotheses drawn from RCT and VET.

These findings contribute to the emerging understanding of on-line norm violations that are low-risk to violators but harmful to victims, communities, and public goods. Identifiability and effort, together with exclusionary policies that affect some vandals, offer some explanation for both the rates of vandalism and types of vandalism perpetrated. Interventions that target these factors independently may have unintended consequences and could deter newcomers and valuable casual contributions [15, 16].

This study contributes to the study of online vandalism by applying theories of Rational Choice and Value Expectancy in order to develop a framework for understanding and modeling vandal behavior. Doing so required an extension and synthesis of two vandalism ontologies and bringing together perspectives rooted in online and offline vandalism. This work suggests that treating vandalism as the product of an antinormative but rational decision-making process supports further consideration of different strategies for countering and preventing it.

8.1 Acknowledgements

I am tremendously grateful for the guidance of Karl-Dieter Opp; all errors are mine, but his encouragement and mentorship were invaluable. Three anonymous reviewers provided helpful suggestions and the resulting manuscript is much improved by their feedback. This work would not have been possible without the support of Chau Tran, who kindly shared his dataset of Tor users with me and assisted in sample preparation. This work was born out of early discussions with Nora McDonald, Stephanie Bankes, and Joseph Zhang as we puzzled out what qualified as vandalistic behavior in Wikipedia. This work was supported by the National Science Foundation (awards CNS-1703736 and CNS-1703049).

REFERENCES

- [1] Mustafa Akg  j and Melih K  sr  do     . 2015. Internet censorship in Turkey. *Internet Policy Review* 4, 2 (June 2015), 1–22. <https://doi.org/10.14763/2015.2.366>
- [2] Abdulwhab Alkharashi and Jeomon Jose. 2018. Vandalism on Collaborative Web Communities: An Exploration of Editorial Behaviour in Wikipedia. In *Proceedings of the 5th Spanish Conference on Information Retrieval - CERI '18*. ACM Press, Zaragoza, Spain, 1–4. <https://doi.org/10.1145/3230599.3230608>
- [3] Yochai Benkler. 2006. *The wealth of networks: How social production transforms markets and freedom*. Yale University Press, New Haven, CT.
- [4] Kaylea Champion, Nora McDonald, Stephanie Banks, Joseph Zhang, Rachel Greenstadt, Andrea Forte, and Benjamin Mako Hill. 2019. A forensic qualitative analysis of contributions to Wikipedia from anonymity seeking users. *Proceedings of the ACM: Human-Computer Interaction* 3, CSCW (Nov. 2019), 531:1–53:26. <https://doi.org/10.1145/3359155>
- [5] Si-Chi Chin, W. Nick Street, Padmini Srinivasan, and David Eichmann. 2010. Detecting Wikipedia vandalism with active learning and statistical language models. In *Proceedings of the 4th workshop on Information credibility - WICOW '10*. ACM Press, Raleigh, North Carolina, USA, 3. <https://doi.org/10.1145/1772938.1772942>
- [6] Angela Gracia B. Cruz, Yuri Seo, and Mathew Rex. 2018. Trolling in online communities: A practice-based theoretical perspective. *The Information Society* 34, 1 (Jan. 2018), 15–26. <https://doi.org/10.1080/01972243.2017.1391909>
- [7] Quinn Dombrowski. 2011. Walls that Talk: Thematic Variation in University Library Graffiti. *Journal of the Chicago Colloquium on Digital Humanities and Computer Science* 1, 3 (2011), 1–13.
- [8] Andrea Forte, Nazanin Andalibi, and Rachel Greenstadt. 2017. Privacy, anonymity, and perceived risk in open collaboration: a study of Tor users and Wikipedians. In *Proceedings of the 2017 ACM Conference on Computer Supported Cooperative Work and Social Computing (CSCW '17)*. ACM, New York, NY, 1800–1811. <https://doi.org/10.1145/2998181.2998273>
- [9] Eric J. Friedman and Paul Resnick. 2001. The Social Cost of Cheap Pseudonyms. *Journal of Economics & Management Strategy* 10, 2 (June 2001), 173–199. <https://doi.org/10.1111/j.1430-9134.2001.00173.x>
- [10] R. Stuart Geiger and Aaron Halfaker. 2013. When the levee breaks: Without bots, what happens to Wikipedia’s quality control processes?. In *Proceedings of the 9th International Symposium on Open Collaboration (OpenSym '13)*. ACM, New York, NY, 6:1–6:6. <https://doi.org/10.1145/2491055.2491061>
- [11] R. Stuart Geiger and David Ribes. 2010. The Work of Sustaining Order in Wikipedia: The Banning of a Vandal. In *Proceedings of the 2010 ACM Conference on Computer Supported Cooperative Work (CSCW '10)*. ACM, New York, NY, 117–126. <https://doi.org/10.1145/1718918.1718941>
- [12] Itzhak Gilboa. 2010. *Rational choice*. MIT Press, Cambridge, Mass.
- [13] John H. Goldthorpe. 1998. Rational Action Theory for Sociology. *The British Journal of Sociology* 49, 2 (June 1998), 167–192.
- [14] Andreea D. Gorbat  ci. 2014. *The paradox of novice contributions to collective production: evidence from Wikipedia*. SSRN Scholarly Paper ID 1949327. Social Science Research Network, Rochester, NY. <https://papers.ssrn.com/abstract=1949327>
- [15] Aaron Halfaker, R. Stuart Geiger, Jonathan T. Morgan, and John Riedl. 2013. The rise and decline of an open collaboration system: how Wikipedia’s reaction to popularity is causing its decline. *American Behavioral Scientist* 57, 5 (May 2013), 664–688. <https://doi.org/10.1177/0002764212469365>
- [16] Aaron Halfaker, Aniket Kittur, and John Riedl. 2011. Don’t bite the newbies: How reverts affect the quantity and quality of Wikipedia work. In *Proceedings of the 7th International Symposium on Wikis and Open Collaboration (WikiSym '11)*. ACM, New York, NY, 163–172. <https://doi.org/10.1145/2038558.2038585>
- [17] Aaron Halfaker, Jonathan Morgan, Amir Sarabadani, and Adam Wight. 2016. *ORES: Facilitating re-mediation of Wikipedia’s socio-technical problems*. Working Paper. Wikimedia Research. https://meta.wikimedia.org/wiki/Research:ORES:_Facilitating_re-mediation_of_Wikipedia%27s_socio-technical_problems
- [18] Hsiao Ying Huang and Masooda Bashir. 2016. The Onion Router: Understanding a Privacy Enhancing Technology Community. In *Proceedings of the 7th AISAC&T*

- Annual Meeting: Creating Knowledge, Enhancing Lives Through Information & Technology (ASIST '16)*. American Society for Information Science, Silver Springs, MD, USA, 34:1–34:5. <http://dl.acm.org/citation.cfm?id=3017447.3017481>
- [19] Adam Joinson. 1998. Causes and implications of disinhibited behavior on the Internet. In *Psychology and the Internet: Intrapersonal, interpersonal, and transpersonal implications*, Jayne Gackenbach (Ed.). Academic Press, San Diego, CA, US, 43–60.
 - [20] Ruogu Kang, Stephanie Brown, and Sara Kiesler. 2013. Why Do People Seek Anonymity on the Internet?: Informing Policy and Design. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '13)*. ACM, New York, NY, USA, 2657–2666. <https://doi.org/10.1145/2470654.2481368>
 - [21] Michael D. Lieberman and Jimmy Lin. 2009. You Are Where You Edit: Locating Wikipedia Contributors through Edit Histories.. In *Proceedings of the Third International ICWSM Conference*. 106–113. http://www.pensivepuffin.com/dwmcphd/syllabi/infx598_wi12/papers/wikipedia/lieberman-lin.YouAreWhereYouEdit.ICWSM09.pdf
 - [22] Ross L. Matsueda, Derek A. Kreager, and Davis Huizinga. 2006. Detering Delinquents: A Rational Choice Model of Theft and Violence. *American Sociological Review* 71, 1 (2006), 95–122.
 - [23] Nora McDonald, Benjamin Mako Hill, Rachel Greenstadt, and Andrea Forte. 2019. Privacy, anonymity, and perceived risk in open collaboration: a study of service providers. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems (CHI '19)*. ACM, New York, NY, 671:1–671:12. <https://doi.org/10.1145/3290605.3300901>
 - [24] Guido Mehlkop and Peter Graeff. 2010. Modelling a Rational Choice theory of Criminal Action: Subjective Expected Utilities, Norms, and Interaction. *Rationality and Society* 22, 2 (2010), 189–222.
 - [25] Jonathan T. Morgan and Aaron Halfaker. 2018. Evaluating the impact of the Wikipedia teahouse on newcomer socialization and retention. In *Proceedings of the 14th International Symposium on Open Collaboration (OpenSym '18)*. ACM, New York, NY, 20:1–20:7. <https://doi.org/10.1145/3233391.3233544>
 - [26] Nima Nazeri and Collin Anderson. 2013. *Citation filtered: Iran's censorship of Wikipedia*. Technical Report. University of Pennsylvania.
 - [27] Karl-Dieter Opp. 1999. Contending Conceptions of the Theory of Rational Action. *Journal of Theoretical Politics* 11 (1999), 171–202.
 - [28] William H. Riker and Peter C. Ordeshook. 1973. *An Introduction to Positive Political Theory*. Prentice Hall, Englewood Cliffs, N.J.
 - [29] Pnina Shachaf and Noriko Hara. 2010. Beyond vandalism: Wikipedia trolls. *Journal of Information Science* 36, 3 (June 2010), 357–370. <https://doi.org/10.1177/0165551510365390>
 - [30] Ángel Obregón Sierra and Jorge Ocea Castanedo. 2018. University students in the educational field and Wikipedia vandalism. In *Proceedings of the 14th International Symposium on Open Collaboration - OpenSym '18*. ACM Press, Paris, France, 1–7. <https://doi.org/10.1145/3233391.3233540>
 - [31] John Suler. 2004. The online disinhibition effect. *Cyberpsychology & behavior* 7, 3 (2004), 321–326.
 - [32] John R. Suler and Wende L. Phillips. 1998. The Bad Boys of Cyberspace: Deviant Behavior in a Multimedia Chat Community. *CyberPsychology & Behavior* 1, 3 (Jan. 1998), 275–294.
 - [33] Jesús Tramullas, Piedad Garrido-Picazo, and Ana I. Sánchez-Casabán. 2016. Research on Wikipedia Vandalism: a brief literature review. In *Proceedings of the 4th Spanish Conference on Information Retrieval - CERI '16*. ACM Press, Granada, Spain, 1–4. <https://doi.org/10.1145/2934732.2934748>
 - [34] Chau Tran, Kaylea Champion, Andrea Forte, Benjamin Mako Hill, and Rachel Greenstadt. 2020. Tor Users Contributing to Wikipedia: Just Like Everybody Else?. In *Conference Proceedings, IEEE Security and Privacy (Oakland 2020)*, Vol. 1. 974–990. <https://doi.org/10.1109/SP40000.2020.00053>
 - [35] Eric C. Turner, Subhasish Dasgupta, and George Washington. 2003. Privacy on the web: An examination of user's concerns, technology, and implications for business organizations and individuals. *Information Systems Management* 20, 1 (2003), 8–18. <https://doi.org/10.1201/1078/43203.20.1.20031201/40079.2>
 - [36] Luis von Ahn, Manuel Blum, Nicholas J. Hopper, and John Langford. 2003. CAPTCHA: Using hard AI problems for security. In *Advances in Cryptology – EUROCRYPT 2003 (Lecture Notes in Computer Science)*, Eli Biham (Ed.). Springer, Berlin, Germany, 294–311. https://doi.org/10.1007/3-540-39200-9_18
 - [37] Rob White. 2001. Graffiti, Crime Prevention & Cultural Space. *Current Issues in Criminal Justice* 12, 3 (2001), 253–268.